# High-Dimensional Bayesian Optimization via Additive Models with Overlapping Groups 

Paul Rolland ${ }^{1}$, Jonathan Scarlett ${ }^{2}$, Ilija Bogunovic ${ }^{1}$, Volkan Cevher ${ }^{1}$<br>${ }^{1}$ Laboratory for Information and Inference Systems (LIONS), EPFL<br>${ }^{2}$ Department of Computer Science \& Department of Mathematics, National University of Singapore<br>\{paul.rolland, ilija.bogunovic, volkan.cevher\}@epfl.ch, scarlett@comp.nus.edu.sg


#### Abstract

Bayesian optimization (BO) is a popular technique for sequential black-box function optimization, with applications including parameter tuning, robotics, environmental monitoring, and more. One of the most important challenges in BO is the development of algorithms that scale to high dimensions, which remains a key open problem despite recent progress. In this paper, we consider the approach of Kandasamy et al. (2015), in which the high-dimensional function decomposes as a sum of lower-dimensional functions on subsets of the underlying variables. In particular, we significantly generalize this approach by lifting the assumption that the subsets are disjoint, and consider additive models with arbitrary overlap among the subsets. By representing the dependencies via a graph, we deduce an efficient message passing algorithm for optimizing the acquisition function. In addition, we provide an algorithm for learning the graph from samples based on Gibbs sampling. We empirically demonstrate the effectiveness of our methods on both synthetic and real-world data.


## 1 Introduction

Bayesian optimization (BO) is a powerful method for sequentially optimizing an unknown function $f$ that is costly to evaluate, for which noisy point evaluations are available. Since its introduction, BO has successfully been applied to a variety of applications, including algorithm parameter tuning (e.g., deep neural networks) [1], 2, [3] and robotics [4, [5. However,

Proceedings of the $21^{\text {st }}$ International Conference on Artificial Intelligence and Statistics (AISTATS) 2018, Lanzarote, Spain. PMLR: Volume 84. Copyright 2018 by the author(s).
most successful applications of BO have involved lowdimensional input spaces. Efficiently scaling to high dimensions remains a key open challenge, and is crucial in applications such as computer vision [6], biology [7, and larger-scale parameter tuning.

High-dimensional BO comes with two key inter-related challenges [8]: Identifying "low-dimensional structure" in the high-dimensional function, and choosing an acquisition function that can efficiently be optimized. There is an inherent tension between these goals, with richer forms of structure often leading to acquisition functions that are harder to optimize.

### 1.1 Related Work

A recent overview of BO can be found in [9. Most BO algorithms can be posed as choosing the next point to maximize an acquisition function, which in turn depends on the current posterior of the function. Popular choices include upper confidence bound (GP-UCB) [10, probability of improvement (PI) [11, expected improvement (EI) [12], [13], and (predictive) entropy search (ES) 14 (15). In this paper, we are particularly interested in high-dimensional extensions of GP-UCB.
The earliest works on high-dimensional BO considered functions that only vary along a low-dimensional subspace [16, [17, [18. While such approaches can be effective, this assumption on the function is rather strong, and achieving the full potential highdimensional BO requires moving to richer classes. A promising alternative approach was recently proposed by Kandasamy et al. 19, who modeled the function by a sum of independent low-dimensional functions, each defined on a fixed subset of the underlying variables.

Crucially, the work of 19 assumed that these subsets are disjoint. This constraint considerably simplifies the optimization of the acquisition function, but restricts the space of functions that can be modeled. A generalization of this approach was proposed in [20] that allows for possible rotations within each term of
the decomposition.
Within the framework of [19, a crucial challenge is learning the underlying additive decomposition from samples (typically while simultaneously performing optimization). In [19], this was done based on randomly sampling the decompositions and choosing one to maximize the likelihood. More efficient approaches were proposed by Wang et al. [21] based on Gibbs sampling, and by Gardner et al. [22] based on a simple Markov Chain Monte Carlo (MCMC) method.

Additive models, not necessarily making use of Gaussian processes, have also found extensive use in other contexts such as function learning, e.g., see [23, 24] and the references therein.

### 1.2 Contributions

The main contributions of this paper are as follows:

- We generalize the additive model of [19] by allowing the additive model to consist of functions defined on general subsets of the underlying variables that need not be disjoint.
- By representing the interactions between the groups (i.e., the subsets on which the lowdimensional functions are defined) using a graph, we deduce an efficient high-dimensional variant of GP-UCB based on message passing.
- We present a Gibbs sampling algorithm for learning the structure of the graph from data, having a similar flavor to that of 21] while addressing new challenges for the case of overlapping groups.
- We demonstrate the effectiveness of our approach on both synthetic and real data sets, including improved versatility compared to the work of [19].
- In the supplementary material, we generalize certain aspects of the mathematical analysis from [19], and discuss their possible implications towards providing regret bounds.

We very recently learned of a closely related independent parallel work [25] adopting a similar model, but using an approximate decentralized approach to optimize the acquisition function, and learning the graph via an alternative approach building on [22].

## 2 Generalized Additive GP Model

We consider the optimization of a $D$-dimensional function $f(x)$, where $x=\left(x_{1}, \ldots, x_{D}\right)$ is the input vector. We focus on discrete domains, where each variable $x_{i}$ takes values in some finite set $\mathcal{X}_{i}$ (though this set could represent the quantization of a real interval such as $[0,1])$. Hence, the high-dimensional domain is
$\mathcal{X}=\mathcal{X}_{1} \times \ldots \times \mathcal{X}_{D}$. Each time we query the function $f$ at some point $x \in \chi$, we obtain a noisy observation $y=f(x)+\epsilon$ where $\epsilon \sim \mathcal{N}\left(0, \eta^{2}\right)$. We aim at maximizing this function over the domain $\mathcal{X}$, i.e. finding $x_{\mathrm{opt}}=\arg \max _{x \in \mathcal{X}} f(x)$.

Additive structure: Following the work of Kandasamy et al. [19], we assume that the target function can be decomposed into a sum of low-dimensional components as follows:

$$
\begin{equation*}
f(x)=f^{(1)}\left(x^{(1)}\right)+f^{(2)}\left(x^{(2)}\right)+\ldots+f^{(M)}\left(x^{(M)}\right), \tag{1}
\end{equation*}
$$

where each $x^{(i)} \in \mathcal{X}^{(i)} \subseteq \mathcal{X}$ is a low-dimensional component, with $\mathcal{X}^{(i)}$ being the product of a small number of $\mathcal{X}_{k}$. The variables involved in this product are referred to as the $i$-th group, and this set of variables is denoted by $\mathcal{G}^{(i)} \subseteq\{1, \ldots, D\}$.
In [19], it was assumed that the different variable sets $\mathcal{G}^{(i)}$ do not overlap, i.e. $\mathcal{G}^{(i)} \cap \mathcal{G}^{(j)}=\emptyset$ for all $(i, j)$. In our setting, we allow for arbitrary overlaps between these variable sets, thereby permitting a significantly richer model class that is suited to interacting groups.

Prior and posterior: We assume that each term $f^{(i)}$ is an independent sample from a Gaussian process $\mathrm{GP}\left(\mu^{(i)}, \kappa^{(i)}\right)$. As a result, the overall target function is also a sample from a $G P: f \sim \operatorname{GP}(\mu, \kappa)$ where

$$
\begin{align*}
\mu(x) & =\sum_{i=1}^{M} \mu^{(i)}\left(x^{(i)}\right)  \tag{2}\\
\kappa\left(x, x^{\prime}\right) & =\sum_{i=1}^{M} \kappa^{(i)}\left(x^{(i)}, x^{\prime(i)}\right) \tag{3}
\end{align*}
$$

Let $\mathcal{D}_{t}=\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{t}$ be the data observed from the target function $f$ where $\mathbf{y}=\left(y_{1}, \ldots, y_{n}\right)$ are the noisy observations corresponding to $\mathbf{x}=\left(x_{1}, \ldots, x_{n}\right)$, i.e. $\quad y_{i} \sim \mathcal{N}\left(f\left(x_{i}\right), \eta^{2}\right), i=1, \ldots, n$. Conditioned on these observations $\mathcal{D}_{t}$, we can infer the posterior mean and variance for each term $f^{(i)}$ at an arbitrary point $x_{*}$. We show in the supplementary material that $\left(f_{*}^{(j)} \mid \mathbf{y}\right) \sim \mathcal{N}\left(\mu_{t-1}^{(i)},\left(\sigma_{t-1}^{(i)}\right)^{2}\right)$, where

$$
\begin{align*}
\mu_{t-1}^{(j)}= & \kappa^{(j)}\left(x_{*}^{(j)}, \mathbf{x}^{(j)}\right) \Delta^{-1} \mathbf{y} \\
\left(\sigma_{t-1}^{(j)}\right)^{2}= & \kappa^{(j)}\left(x_{*}^{(j)}, x_{*}^{(j)}\right)  \tag{4}\\
& -\kappa^{(j)}\left(x_{*}^{(j)}, \mathbf{x}^{(j)}\right) \Delta^{-1} \kappa^{(j)}\left(\mathbf{x}^{(j)}, x_{*}^{(j)}\right)
\end{align*}
$$

with $\Delta=\kappa(\mathbf{x}, \mathbf{x})+\eta^{2} I_{n} \in \mathbb{R}^{n \times n}$. Here $\kappa^{(j)}\left(\mathbf{x}^{(j)}, x_{*}^{(j)}\right)$ is a column vector of size $n$ whose $i$-th entry is $\kappa^{(j)}\left(x_{i}^{(j)}, x_{*}^{(j)}\right)$, and $\kappa(\mathbf{x}, \mathbf{x})$ is a matrix of size $n \times n$ whose $\left(i, i^{\prime}\right)$-th entry is $\kappa\left(x_{i}, x_{i^{\prime}}\right)$.
Dependency graph: An additive decomposition can be represented by a dependency graph. The dependency graph is built by joining variables $i$ and $j$


Figure 1: Example dependency graph. There are 4 maximal cliques: $(1,2,3),(1,3,4),(4,5)$, and ( 6 ), each of which are associated to a term of the decomposition.
with an edge whenever they appear together within some set $x^{(k)}$. For example, the graph associated with the decomposition of $f(x)=f^{(1)}\left(x_{1}, x_{2}, x_{3}\right)+$ $f^{(2)}\left(x_{1}, x_{3}, x_{4}\right)+f^{(3)}\left(x_{4}, x_{5}\right)+f^{(4)}\left(x_{6}\right)$ is shown in Figure1. In the special case of [19], the dependency graph consists of disjoint fully connected components.

While different decompositions may lead to the same graph (e.g., consider the case of three functions on $\left(x_{1}, x_{2}\right),\left(x_{2}, x_{3}\right)$, and $\left(x_{1}, x_{3}\right)$ vs. a single function on $\left(x_{1}, x_{2}, x_{3}\right)$ ), one can always adopt the more general decomposition for a given graph to avoid any loss of generality (e.g., take the single function on $\left.\left(x_{1}, x_{2}, x_{3}\right)\right)$. To this end, one can let each lowdimensional function correspond to a single maximal clique (i.e., fully-connected component) of the graph.

## 3 Generalized Additive GP-UCB

### 3.1 Description of algorithm

Acquisition function: We assign to each term $f^{(i)}$ a low-dimensional UCB acquisition function [26]:

$$
\begin{equation*}
\phi_{t}^{(i)}\left(x^{(i)}\right)=\mu_{t-1}\left(x^{(i)}\right)+\beta_{t}^{\frac{1}{2}} \sigma_{t-1}\left(x^{(i)}\right) \tag{5}
\end{equation*}
$$

where $\beta_{t}$ is an exploration parameter. The global acquisition function is the sum of low-dimensional ones:

$$
\begin{equation*}
\phi_{t}(x)=\sum_{i=1}^{M} \phi_{t}^{(i)}\left(x^{(i)}\right) \tag{6}
\end{equation*}
$$

The bulk of this section is devoted to methods for efficiently maximizing this function. While this is straightforward for disjoint groups [19], it is non-trivial in our generalized model.

Full algorithm: Our algorithm, which we call generalized additive $G P-U C B$ (G-Add-GP-UCB), is given in Algorithm 1 As stated, the algorithm is suited to the case that the kernel (and hence the dependency graph) is known; this assumption is dropped in Section 4.

```
Algorithm 1 G-Add-GP-UCB
    Pick \(\left(x_{t}\right)_{i=1}^{N_{\text {init }}}\) at random; evaluate \(f\) at these points
    to get \(\left(y_{t}\right)_{i=1}^{N_{\text {init }}}\) and add these pairs to \(D_{0}\).
    for \(t=1, \ldots, N_{\text {iter }}\) do
        Perform Bayesian posterior updates conditioned
        on \(D_{t-1}\) to obtain \(\mu_{t-1}^{(j)}, \sigma_{t-1}^{(j)}\) for \(j=1, \ldots, M\)
        Optimize the acquisition function (cf., Section
        (3.2) to obtain \(x_{t} \leftarrow \arg \max _{x \in \chi} \tilde{\phi}_{t}(x)\)
        Evaluate \(x_{t}\) and observe \(y_{t} \leftarrow f\left(x_{t}\right)+\epsilon\)
        Augment the data set : \(D_{t}=D_{t-1} \cup\left(x_{t}, y_{t}\right)\)
    end for
```


### 3.2 Maximizing the acquisition function

The key to maximizing the acquisition function (when the dependency graph is known) is to connect the optimization with the problem of maximizing probability in Markov random fields [27]. In particular, in the same way as the latter setting, we can make use of message passing for efficient maximization. We first explain how this is done when the dependency graph is triangulated (i.e., when the graph has no chordless cycles of length greater than 3 ), and then extend the algorithm to general graphs.
Triangulated dependency graphs: We start by constructing a junction tree [27] of the dependency graph. The nodes of the junction tree are then the different maximal cliques of the dependency graph. By the construction of the dependency graph in the previous section, there is a low-dimensional acquisition function $\phi^{(C)}$ of the form (5) associated to each of these maximal cliques $C$.
Starting from the leaves of the junction tree, we sequentially maximize, for each clique $C$, the associated low-dimensional acquisition functions $\phi^{(C)}$, and pass "messages" $m_{C_{p} \leftarrow C}$ to the parent node $C_{p}$. Each message is a function of the variables of the parent node, with the variables of the previous nodes already optimized. This process propagates up to the root, with a given node adding the messages from its children (cf., Algorithm 2). The running intersection property of the junction tree ensures that there is no conflict between the maximizations performed, and that this algorithm returns the maximum of the global acquisition function [27.
Arbitrary dependency graphs: When the dependency graph is not triangulated, there is no junction tree that can be constructed from the original dependency graph. We thus first need to triangulate it, and construct a junction tree for the processed dependency graph. However, since the dependency graph has been modified, we no longer directly have a unique lowdimensional acquisition function for each clique of the

```
Algorithm 2 Message Passing algorithm on a trian-
gulated dependency graph
    Root the junction tree at a random node \(C_{R}\).
    \(d \leftarrow\) depth of rooted tree G.
    while \(d \geq 1\) do
        for each node \(C\) of the junction tree at distance
        \(d\) from the root do
            \(C_{p} \leftarrow \operatorname{parent}(\mathrm{C}), I \leftarrow C \cap C_{p}, J \leftarrow C \backslash C_{p}\)
            Compute messages to be passed to the parent
            node:
            \(m_{C_{p} \leftarrow C}\left(x^{(I)}\right)=\max _{x^{(J)}} \phi^{C}\left(x^{(C)}\right)+\)
            \(\sum_{C_{c} \in \operatorname{Children}(C)} m_{C \leftarrow C_{c}}\left(x^{\left(C \cap C_{c}\right)}\right)\)
            for each \(x^{(I)} \in \chi^{(C)}\)
        end for
        \(d \leftarrow d-1\).
    end while
    Return \(\quad \max _{x^{\left(C_{R}\right)}} \phi^{\left(C_{R}\right)}\left(x^{\left(C_{R}\right)}\right)+\)
    \(\sum_{C_{c} \in \operatorname{Children}\left(C_{R}\right)} m_{C_{R} \leftarrow C_{c}}\left(x^{\left(C_{R} \cap C_{c}\right)}\right)\).
```

junction tree. In order to ensure that each function $\phi^{(i)}$ is maximized once and only once during the optimization process, we assign to each clique $C$ the following "acquisition function":

$$
\begin{equation*}
\phi^{(C)}\left(x^{(C)}\right)=\sum_{\substack{\text { clique } c \\ c \notin C_{c} \forall C_{c} \in \text { of } G, c \subset C, \operatorname{Children}^{(C)}}} \phi^{(c)}\left(x^{(c)}\right), \tag{7}
\end{equation*}
$$

where $G$ corresponds to the original dependency graph. By doing so, we can then apply Algorithm 2 on the triangulated graph.

Complexity: The complexity of running the message-passing algorithm on a junction tree $J$ is exponential in the size of the maximum clique of the triangulated graph associated with junction tree $J$. This quantity depends on the chosen triangulation, so it would be desirable to compute a triangulation that yields a small maximal clique [28, 29].

## 4 Learning the Dependency Graph

One of the main important practical challenges of BO is choosing a suitable kernel. In the high-dimensional setting, this challenge is even more difficult, as we need to learn not only the kernel parameters (e.g., length scales), but also the structure associated with the highdimensional function (i.e., the dependency graph). In this section, we present a Gibbs sampling procedure for this purpose, building on the approach of 21 for the case of non-overlapping groups.

Preliminaries: As discussed previously, any decomposition can be represented by an undirected graph, in which each low-dimensional kernel $\kappa^{(j)}$ is associated to a maximal clique. For convenience, here we
represent this decomposition by an adjacency matrix $Z \in\{0,1\}^{D \times D}$, where $Z_{i j}=1$ if variable $i$ is connected to $j$ in the graph, and 0 otherwise. This matrix is symmetric and has zeros on the diagonal, so the number of free parameters is $\frac{D(D-1)}{2}$.
The dependency graph defines the kernel decomposition, and thus influences the number of lowdimensional kernels and their dimensions. Therefore, in general, the kernel parameters can be defined only once the decomposition is known. In order to learn these parameters simultaneously with the additive structure, we assume that these kernels only depend on a constant number $n_{\text {param }}$ of parameters independent of the decomposition (see Section 5 for a concrete example). We then group them into the set $L=$ $\left\{L_{i}\right\}_{i=1, \ldots, n_{\text {param }}}$. Overall, the parameters that must be optimized are $\theta=\left\{\left\{Z_{i j}\right\}_{1 \leq i<j \leq D},\left\{L_{j}\right\}_{1 \leq j \leq n_{\text {param }}}\right\}$, which results in $\frac{D(D-1)}{2}+n_{\text {param }}$ parameters.

Maximum likelihood: In order to infer the structure of the dependency graph and the kernels' parameters, we seek to maximize data likelihood:

$$
\begin{aligned}
\log p\left(D_{n} \mid Z, L\right) & =-\frac{1}{2} y^{T}\left(K_{n}^{G(Z), L}+\sigma^{2} I\right)^{-1} y \\
& -\frac{1}{2} \log \left|K_{n}^{G(Z), L}+\sigma^{2} I\right|-\frac{n}{2} \log 2 \pi
\end{aligned}
$$

where $G(Z)$ is the graph corresponding to the adjacency matrix $Z, y$ is the vector containing the current observations, and $K_{n}^{G, L} \in \mathbb{R}^{n \times n}$ is the kernel matrix of the observed data points, supposing a dependency graph $G$ and kernel parameters $\left\{L_{i}\right\}_{1 \leq i \leq n_{\text {param }}}$, i.e., $\left(K_{n}^{G, L}\right)_{i j}=\kappa^{G, L}\left(x_{i}, x_{j}\right)$. Here $\kappa^{G, L}$ represents the high-dimensional kernel determined by $G$ and $L$.
Gibbs sampling: We adopt a Bayesian approach, in which we place a prior distribution on the parameters $\left\{\theta_{i}\right\}_{1 \leq i \leq N}$, and seek to sample from the posterior distribution $p\left(\theta_{1}, \ldots, \theta_{N} \mid D\right)$. Since we cannot directly sample from this high-dimensional probability distribution, we use the Gibbs sampling method (cf., Algorithm (3). This provides a means for approximately sampling from the joint distribution, as long as we can sample from the 1-dimensional conditional distributions $p\left(\theta_{i} \mid \theta_{-i}, D\right)$ with $\theta_{-i}=\left\{\theta_{j} \mid j \neq i\right\}$.
The algorithm starts from a set of parameters $\theta^{(0)}$, and iteratively samples new sets of parameters $\theta^{(j)}$ by modifying one coordinate at a time based on the posterior (Algorithm 3). Once the sampling is performed, we choose the set of parameters that achieves the highest data likelihood. It can be shown [30] that under soft assumptions on the probability distribution $p(\theta \mid D)$, the Gibbs sampler tends to sample the same way as if we were to sample directly from $p(\theta \mid D)$.

Prior distributions: We model the variables

```
Algorithm 3 Structure learning via Gibbs Sampling
    \(\Theta \leftarrow\left\{\theta^{(0)}\right\}, \theta \leftarrow \theta^{(0)}, j \leftarrow 0\)
    for \(j=1,2, \ldots\) do
        for \(i=1, \ldots, N\) do
            \(\theta^{(j+1)} \leftarrow \theta^{(j)}\).
            Sample \(\theta_{i}^{\text {new }}\) from \(p\left(\theta_{i} \mid \theta_{-i}^{(j)}, D\right)\)
            \(\theta_{i}^{(j+1)} \leftarrow \theta_{i}^{\text {new }}\)
            Augment the data set \(: ~ \Theta \leftarrow \Theta \cup\left\{\theta^{(j+1)}\right\}\)
        end for
    end for
```

$\left\{Z_{i j}\right\}_{1 \leq i<j \leq D}$ as Bernoulli random variables with parameter $p: Z_{i j} \sim \operatorname{Bernoulli}(p)$ where $p$ gives the probability of an edge joining variables $i$ and $j$. This parameter $p$ can be used to control the sparsity of the graph, or set to $\frac{1}{2}$ if no prior information is available.
Using this model, the posterior distribution for $Z_{i j}$ is

$$
\begin{aligned}
& p\left(Z_{i j}=1 \mid Z_{-(i j)}, L, D_{n} ; p\right) \\
& \propto p\left(D_{n} \mid Z_{-(i j)}, Z_{i j}=1, L\right) \cdot p\left(Z_{i j}=1 \mid Z_{-(i j)}, L ; p\right) \\
& =p\left(D_{n} \mid Z_{-(i j)}, Z_{i j}=1, L\right) \cdot p \\
& \propto p \cdot e^{\phi\left(Z_{-(i j)} \cup Z_{i j}=1, L\right)}
\end{aligned}
$$

with $\phi(Z, L)=-\frac{1}{2} y^{T}\left(K_{n}^{G(Z), L}+\sigma^{2} I\right)^{-1} y-$ $\frac{1}{2} \log \left|K_{n}^{G(Z), L}+\sigma^{2} I\right|$. Concerning the kernels' parameters, we simply model them as uniform variables over pre-defined sets of possible values: $L_{i} \sim \operatorname{Uniform}\left(\mathcal{L}_{i}\right)$ so that for each $l \in \mathcal{L}_{i}$, we have

$$
\begin{equation*}
p\left(L_{i}=l \mid Z, L_{-i}, D_{n} ; p\right) \propto e^{\phi\left(Z, L_{-i} \cup\left\{L_{i}=l\right\}\right)} \tag{8}
\end{equation*}
$$

Using these posterior distributions, we can apply Gibbs sampling to the parameter set $\theta=$ $\left\{\left\{Z_{i j}\right\}_{1 \leq i<j \leq D},\left\{L_{i}\right\}_{1 \leq i \leq n_{\text {param }}}\right\}$, and select the set which produces the highest data likelihood. For the binary variables $Z_{i j}$, we simply compute $p_{0}=(1-$ $p) \cdot e^{\phi\left(Z_{-(i j)} \cup\left\{Z_{i j}=0\right\}, L\right)}$ and $p_{1}=p \cdot e^{\phi\left(Z_{-(i j)} \cup\left\{Z_{i j}=0\right\}, L\right)}$ and then sample from the binary distribution $\operatorname{Bernoulli}\left(\frac{p_{1}}{p_{0}+p_{1}}\right)$. For the kernels' parameters, we proceed the same way by computing $e^{\phi\left(Z, L_{-i} \cup\left\{L_{i}=l\right\}\right)}$ for values $l$ in the set $\mathcal{L}_{i}$ and then sample from the normalized probability distribution.

Stopping criterion: We stop the process after some number $N_{\text {Gibbs }}$ of data likelihoods have been computed. The more data we have, the better the algorithm will perform to find the true dependency graph, or a closely related one. Therefore, this learning process is repeated throughout the Bayesian optimization algorithm every $N_{\text {cyc }}$ iterations, for some choice of $N_{\text {cyc }}$. Each time we learn new parameters (i.e., graph structure and kernel parameters), we start the sampling from the previously learned parameters.


Figure 2: True dependency graphs for synthetic experiments. For both graphs, each edge forms a maximal clique, and thus corresponds to a term in the additive decomposition of the target function $f$. These graphs are not composed of disjoint cliques, and thus cannot be modeled by the constrained model without overlap.

## 5 Experiments

In this section, we experimentally compare our generalized algorithm against the one of [19]. We focus on squared exponential kernels for modeling the lowdimensional components of the target function:

$$
\begin{aligned}
\kappa^{(i), L^{(i)}} & \left(x^{(i)}, x^{\prime(i)}\right) \\
& =\sigma^{(i)} \exp \left(-\frac{1}{2}\left(x^{(i)}-x^{\prime(i)}\right)^{T} L^{(i)}\left(x^{(i)}-x^{\prime(i)}\right)\right)
\end{aligned}
$$

The scales $\sigma^{(i)}$ are set to $\sigma^{(i)}=\frac{d_{i}}{\sum_{j} d_{j}}$, so that $\kappa(x, x)=1$. The matrices $L^{(i)}$ are all diagonal and generated from one $D$-dimensional lengthscale vector $l:\left(L^{(i)}\right)^{-1}=\operatorname{diag}\left(l_{\chi^{(i)}}\right)^{2}$ where $l_{\chi^{(i)}}$ is the vector containing the lengthscales of variables within $\chi^{(i)}$. This vector is then learned from the data via Gibbs sampling together with the dependency graph.

### 5.1 Experiments on synthetic data

We first test our algorithm on synthetic data by sampling functions from Gaussian processes, according to the dependency graphs shown in Figure 2. We use a squared exponential kernel with lengthscale matrix $L^{(i)}=0.2 * I_{2}$. When applying Gibbs sampling use $p=\frac{1}{2}$ (i.e., no prior knowledge) for the sampling of the dependency graphs.
We apply our algorithm when considering 3 different cases, described as follows.

Overlap: In this case, we learn both the lengthscales and the dependency graph, and allow for overlaps between the groups in the function decomposition. Thus, any dependency graph is allowed, and junction trees are used when optimizing the acquisition function.

No Overlap: This case is similar to the previous one, except that we do not allow for overlaps between
the groups in the function decomposition. This algorithm is thus the same as the one designed by Kandasamy et al. [19]. We place no "hard constraints" on the group sizes (i.e., restrictions on the maximum number of groups $M$ and maximum group size $\mathrm{d}_{\text {max }}$ therein), but we apply the Gibbs sampling procedure of [21, which only samples graphs that are sufficiently likely according to the samples.

Oracle: in the last situation, we assume that we know the true dependency graph and lengthscales, and hence we do not need to learn them throughout the algorithm. We expect this case to perform the best.

We also compare the results with a random algorithm that simply evaluates points at random. We perform 10 runs, starting from different initial situations where 10 points are chosen at random (for each run, these randomly chosen points are the same for each model), and then run the simulation for a certain number of iterations. All the parameters are the same for each model, and are summarized in Table 1 .

$$
\begin{array}{c|c|c|c}
\beta_{t}^{(i)} & \mathrm{N}_{\mathrm{cyc}} & \mathrm{~N}_{\text {Gibbs }} & \max _{\text {eval }} \\
\hline 0.5 \log 2 t & 30 & 200 & 1000
\end{array}
$$

Table 1: Parameters for the synthetic BO examples.
Optimization performance: For each run and each iteration $t$, we compute the simple regret $S_{t}=$ $\min _{i \leq t} r_{i}$ as well as the average cumulative regret $\frac{R_{t}}{t}=\frac{1}{t} \sum_{i=1}^{t} r_{i}$. We average these quantities over all trials (Figures 3 and (4).

We observe that the knowledge of the variable dependencies indeed greatly improves the performance of the algorithm. Around the beginning of the algorithm, the learning process does not perform well as it uses too few data. Therefore, there is no significant difference between the models "Overlap" and "No Overlap", while the "Oracle" is much more efficient as it uses the true dependency graph. As we get more data, the learning process improves and the "Overlap" model becomes more efficient compared to "No Overlap".

Graph learning accuracy: It is also interesting to assess the learning of the dependency graph throughout the algorithm. To do so, we define two quantities in order to evaluate how close a graph $G$ is from a reference graph $G_{\text {true }}$. We first define the Correct Connections quantity:

$$
\mathrm{CC}\left(G, G_{\text {true }}\right)=\frac{\# \text { edges in both } \mathrm{G} \text { and } G_{\text {true }}}{\# \text { edges in } G_{\text {true }}}
$$

which describes how well the variables connections in the true graph are represented by the learned graph.

This quantity is between 0 and 1 , and is equal to 1 if and only if all edges of the true graph are part of the edges of the learned graph. The second quantity that we define is the Correct Separation quantity:

$$
\mathrm{CS}\left(G, G_{\text {true }}\right)=\frac{\# \text { non-edges in both } \mathrm{G} \text { and } G_{\text {true }}}{\# \text { non-edges in } G_{\text {true }}}
$$

which describes how well the variables separations in the true graph are represented by the learned graph. This quantity is also between 0 and 1 , and is equal to 1 if and only if all edges of the learned graph are part of the edges of the true graph. It follows that $G_{1}=G_{2} \Leftrightarrow C C\left(G_{1}, G_{2}\right)=C S\left(G_{1}, G_{2}\right)=1$.

In Figures 5] and 6, we plot these quantities as we obtain more data, and observe that the learning process steadily improves over time. However, in the "No Overlap" case, the graph constraint induces a saturation in the learning, while in the "Overlap" case, we observe converges to the true graph. We note that the high number of correct separations at early iterations is not an indicator of good performance, as the number of correct connections is very low.

We can observe that the time at which the "Overlap" model starts to be more efficient than the "No Overlap" one corresponds to the time where the dependency graph starts to be closer to the true graph.



Figure 3: Star dependency graph: Simple and average cumulative regret averaged over 10 runs.


Figure 4: Grid dependency graph: Simple and average cumulative regret averaged over 10 runs.

### 5.2 Experiments on real data

Face recognition: We consider the optimization of parameters for the Viola and Jones (VJ) Cascade Classifier 31. This algorithm aims at recognizing faces in images. It is based on a cascade of increasingly complex classifiers, which all detect particular features on the image to classify. At each stage, the classifier searches on the image for some features. If it does not find it, the image is rejected, directly classified and no more processing is performed. Otherwise, it continues to the next stage, and so on. It is thus necessary that the first stages have a low rate of false negatives.

Each classifier has a threshold parameter that controls this rejection. The problem is then to optimize these thresholds in order to obtain the highest possible classification accuracy. For each set of parameters, we evaluate the classification accuracy on 1500 images. Among these images, 1000 of them contain exactly one face, and the 500 others do not contain any face. The algorithm correctly classifies an image containing a face if it detects exactly one face. The classification accuracy is defined as the proportion of correctly classified images. We use the OpenCV implementation of the Cascade VJ algorithm 32, which uses a 22 stages cascade algorithm. OpenCV also provides an optimized set of thresholds, which gives a classification


Figure 5: Star dependency graph: Correct Connections and Separations as a function of the number of data used for learning the graph.
accuracy of $92.6 \%$ on our data set. When applying our Bayesian optimization algorithm to this problem, we choose as the domain a neighborhood of the optimized parameter set given by OpenCV, and set the target function as the classification accuracy.

Methods: As in the synthetic data example, we want to compare our "Overlap" model to the "No Overlap" one. In their paper [19, Kandasamy et al. applied their model to this same data set, and observed that the optimal extra parameters for this optimization problem were $M=4$ and $\mathrm{d}_{\max }=6$. We thus use these parameters for this model.

As discussed above, consecutive stages will have more similar rejection thresholds than stages that are far from each other. Therefore, we may suppose that two consecutive stages have more similar thresholds than stages that are far apart. This assumption is enforced by the fact that the set of parameters provided by OpenCV is increasing with the number of stages, justifying this possible correlation between close stages. Therefore, for our "Overlap" model, we set a dependency graph where each variable $i$ is connected to variables $i-1$ and $i+1 \forall i=2, \ldots, 21$, and do not apply the graph structure learning throughout the BO.


Figure 6: Grid dependency graph: Correct Connections and Separations as a function of the number of data used for learning the graph.

The Oracle method cannot be applied here since we do not know the true structure of the target function.

| $\mathrm{N}_{\text {iter }}$ | $\beta_{t}^{(i)}$ | $\mathrm{N}_{\mathrm{cyc}}$ | $\mathrm{N}_{\text {Gibbs }}$ | $\max _{\text {eval }}$ |
| :---: | :---: | :---: | :---: | :---: |
| 200 | $0.5 \log 2 t$ | 50 | 300 | 2000 |

Table 2: BO parameters for face recognition data.

Optimization performance: We apply the two algorithms using the same set of parameters, described in Table 2. We also compare them to a random algorithm which queries random points on the same domain. The results are given in Figure 7 where we average over 15 runs.

We observe that the "Overlap" model performs slightly better than the "No Overlap" algorithm both in terms of convergence speed and optimal value. Both algorithms reach better classification accuracy than with the parameters provided by OpenCV. We can also observe that the "Overlap" model is much better in terms of the average cumulative regret.

Astrophysical data: In the supplementary material, we provide a second example for real-world data based on maximizing likelihood in an astrophysical model.


Figure 7: Face detection using VJ algorithm: classification accuracy and average cumulative regret.

## 6 Conclusion

We introduced a novel model and algorithm for highdimensional Bayesian optimization, in which the target function to optimize can be represented by a sum of possibly overlapping low-dimensional components. By defining an acquisition function with the same additive structure, we can efficiently maximize it using message passing. In addition, we proposed a Gibbs sampling method for learning the structure from samples. By expanding the space of functions that can be modeled, we observed through experiments that the efficiency of Bayesian optimization is improved.

We have not discussed the theoretical aspects of Bayesian optimization (e.g., see [33]). In the supplementary material, we provide some mathematical analysis in this direction, including the calculation of the fundamental information gain, and discuss difficulties in attaining regret bounds for G-Add-GP-UCB.
An interesting direction for future research is to develop variants of message passing that are targeted to continuous settings. Our techniques can readily be applied by discretizing each continuous variable, but it would potentially be more effective (albeit non-trivial) to combine the message passing idea with a continuous global optimization procedure such as DiRect 34.

## 7 Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement $\mathrm{n}^{\circ} 725594$ - time-data). This work was also supported by the Swiss National Science Foundation (SNSF) under grant number 407540_167319.

## References

[1] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In Advances in Neural Information Processing Systems, pages 29512959. 2012.
[2] James Bergstra, R. Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In Conference on Neural Information Processing Systems (NIPS), volume 24, Granada, Spain, December 2011.
[3] Nimalan Mahendran, Ziyu Wang, Firas Hamze, and Nando de Freitas. Adaptive MCMC with Bayesian optimization. Journal of Machine Learning Research - Proceedings Track for Artificial Intelligence and Statistics (AISTATS), 22:751-760, 2012.
[4] Daniel Lizotte, Tao Wang, Michael Bowling, and Dale Schuurmans. Automatic gait optimization with Gaussian process regression. In International Joint Conference on Artificial Intelligence (IJCAI), pages 944-949, 2007.
[5] 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, Atlanta, GA, USA, June 2007.
[6] D. Yamins Bergstra, J. and D. D. Cox. Making a Science of Model Search: hyper-parameter optimization in hundreds of dimensions for vision architectures. In International Conference on Machine Learning (ICML), Atlanta, Gerorgia, 2013), 2013.
[7] David C. James Neil D. Lawrence Javier Gonzlez, Joseph Longworth. Bayesian optimization for synthetic gene design. arXiv:1505.01627, 2015.
[8] Nando. de Freitas. Talk on current challenges and open problems in Bayesian optimization. 2015.
[9] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. Proc. IEEE, 104(1):148-175, 2016.
[10] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. J. Mach. Learn. Res., 3:397-422, March 2003.
[11] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. Journal of Global Optimization, 13(4):455-492, 1998.
[12] Jonas Mockus. Application of Bayesian approach to numerical methods of global and stochastic optimization. Journal of Global Optimization, 4(4):347-365, 1994.
[13] Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. CoRR, abs/1012.2599, 2010.
[14] Philipp Hennig and Christian J Schuler. Entropy search for information-efficient global optimization. J. Mach. Learn. Research, 13(1):1809-1837, 2012.
[15] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In Adv. Neur. Inf. Proc. Sys. (NIPS), pages 918-926, 2014.
[16] Andreas Krause Bo Chen, Rui Castro. Joint optimization and variable selection of highdimensional gaussian processes. arXiv:1206.6396, 2012.
[17] Ziyu Wang, Masrour Zoghi, Frank Hutter, David Matheson, N Freitas, et al. Bayesian optimization in high dimensions via random embeddings. International Joint Conferences on Artificial Intelligence (AAAI), 2013.
[18] Josip Djolonga, Andreas Krause, and Volkan Cevher. High-dimensional Gaussian process bandits. In Conference on Neural Information Processing Systems (NIPS), pages 1025-1033. 2013.
[19] Kirthevasan Kandasamy, Jeff G. Schneider, and Barnabas Poczos. High dimensional Bayesian optimisation and bandits via additive models. In International Conference on Machine Learning, volume 37, pages 295-304, 2015.
[20] Chun-Liang Li, Kirthevasan Kandasamy, Barnabas Poczos, and Jeff Schneider. High dimensional Bayesian optimization via restricted projection pursuit models. In International Conference on Artificial Intelligence and Statistics, Cadiz, Spain, 2016.
[21] Zi Wang, Chengtao Li, Stefanie Jegelka, and Pushmeet Kohli. Batched high-dimensional Bayesian optimization via structural kernel learning. arXiv preprint arXiv:1703.01973, 2017.
[22] Jacob Gardner, Chuan Guo, Kilian Weinberger, Roman Garnett, and Roger Grosse. Discovering and exploiting additive structure for Bayesian optimization. In International Conference on Artificial Intelligence and Statistics, Fort Lauderdale, FL, USA, 20-22 Apr 2017.
[23] Pradeep Ravikumar, Han Liu, John Lafferty, and Larry Wasserman. Spam: Sparse additive models. In Conf. Neur. Inf. Proc. Sys. (NIPS). Curran Associates Inc., 2007.
[24] Hemant Tyagi, Anastasios Kyrillidis, Bernd Gärtner, and Andreas Krause. Algorithms for learning sparse additive models with interactions in high dimensions. Information and Inference, 00:1-67, August 2017.
[25] Trong Nghia Hoang, Quang Minh Hoang, Ruofei Ouyang, and Kian Hsiang Low. Decentralized high-dimensional bayesian optimization with factor graphs. http://arxiv.org/abs/1711.07033v3, 2018.
[26] Niranjan Srinivas, Andreas Krause, Matthias Seeger, and Sham M. Kakade. Gaussian process optimization in the bandit setting: No regret and experimental design. In International Conference on Machine Learning (ICML), pages 1015-1022, 2010.
[27] Martin J. Wainwright. Graphical Models and Message-Passing Algorithms: Some Introductory Lectures, pages 51-108. Springer International Publishing, Cham, 2015.
[28] Stefan Arnborg, Derek G Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in ak-tree. SIAM Journal on Algebraic Discrete Methods, 8(2):277-284, 1987.
[29] Andrés Cano and Serafín Moral. Heuristic algorithms for the triangulation of graphs. Advances in Intelligent Computing (IPMU), pages 98-107, 1995.
[30] G.O. Roberts and A.F.M. Smith. Simple conditions for the convergence of the Gibbs sampler and Metropolis-Hastings algorithms. Stochastic Processes and their Applications, 49(2):207-216, 1994.
[31] Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features, 2001.
[32] Dr. Gary Rost Bradski and Adrian Kaehler. Learning OpenCV, 1st Edition. O'Reilly Media, Inc., first edition, 2008.
[33] N. Srinivas, A. Krause, S.M. Kakade, and M. Seeger. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. IEEE Trans. Inf. Theory, 58(5):32503265, May 2012.
[34] D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. Journal of Optimization Theory and Applications, 79(1):157-181, 1993.
[35] Carl Edward Rasmussen. Gaussian processes for machine learning. MIT Press, 2006.
[36] Kirthevasan Kandasamy, Jeff G. Schneider, and Barnabs Pczos. High dimensional Bayesian optimisation and bandits via additive models. arXiv:1503.01673 [v3], 2017.
[37] MW. Seeger, SM. Kakade, and DP. Foster. Information consistency of nonparametric gaussian process methods. IEEE Transactions on Information Theory, 54(5):2376-2382, May 2008.

# Supplementary Material <br> High-Dimensional Bayesian Optimization via Additive Models with Overlapping Groups (AISTATS 2018) 

Paul Rolland, Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher

All citations below are to the reference list in the main document.

## A Derivation of posterior mean and variance

In this section, we prove equation (4) for the posterior mean and variance conditioned on the observations. The derivation is similar to that of Kandasamy et al. [19], which in turn is similar to the standard Gaussian process posterior derivation [35].

Recall that instead of directly computing the posterior mean and variance on the high-dimensional function, we are considering the terms $f^{(j)}$ in the additive decomposition of $f$ separately. We are thus interested in the distributions of $f_{*}^{(j)}=f^{(j)}\left(x_{*}^{(j)}\right), j=1, \ldots, M$ conditioned on the noisy samples $\mathbf{y}=y_{1}, \ldots, y_{n}$ at points $\mathbf{x}=x_{1}, \ldots, x_{n}$, for some query points $x_{*}^{(j)}$. We claim that the joint distribution of $f_{*}^{(j)}$ and $\mathbf{y}$ can be written as

$$
\binom{f_{*}^{(j)}}{\mathbf{y}} \sim \mathcal{N}\left(\mathbf{0},\left(\begin{array}{ll}
\kappa^{(j)}\left(x_{*}^{(j)}, x_{*}^{(j)}\right) & \kappa^{(j)}\left(x_{*}^{(j)}, \mathbf{x}^{(j)}\right)  \tag{9}\\
\kappa^{(j)}\left(\mathbf{x}^{(j)}, x_{*}^{(j)}\right) & \kappa(\mathbf{x}, \mathbf{x})+\eta^{2} I_{n}
\end{array}\right)\right)
$$

To see this, we recall that distinct functions in the additive decomposition (11) are independent given $\mathbf{x}$. Hence, for any observation $y_{p}=f\left(x_{p}\right)+\epsilon$ and $j=1, \ldots, M$, we have

$$
\begin{aligned}
\operatorname{Cov}\left(f_{*}^{(j)}, y_{p}\right) & =\operatorname{Cov}\left(f_{*}^{(j)}, \sum_{i=1}^{M} f^{(i)}\left(x_{p}^{(i)}\right)+\epsilon\right) \\
& =\operatorname{Cov}\left(f_{*}^{(j)}, f^{(j)}\left(x_{p}^{(j)}\right)\right) \\
& =\kappa^{(j)}\left(x_{*}^{(j)}, x_{p}^{(j)}\right),
\end{aligned}
$$

which establishes (9).
With (9) in place, we can use a standard conditional Gaussian formula (as used in standard GP posterior derivations [35], as well as the non-overlapping setting of [19]) to derive the posterior mean and variance. Specifically, defining the matrix $\Delta=\kappa(\mathbf{x}, \mathbf{x})+\eta^{2} I_{n} \in \mathbb{R}^{n \times n}$, we have for past query points $\mathbf{x}$ and next query point $x_{*}^{(j)}$ that

$$
\begin{align*}
\left(f_{*}^{(j)} \mid \mathbf{y}\right) \sim \mathcal{N} & \left(\kappa^{(j)}\left(x_{*}^{(j)}, \mathbf{x}^{(j)}\right) \Delta^{-1} \mathbf{y}\right. \\
& \left.\kappa^{(j)}\left(x_{*}^{(j)}, x_{*}^{(j)}\right)-\kappa^{(j)}\left(x_{*}^{(j)}, \mathbf{x}^{(j)}\right) \Delta^{-1} \kappa^{(j)}\left(\mathbf{x}^{(j)}, x_{*}^{(j)}\right)\right) \tag{10}
\end{align*}
$$

under the notation in (4). This concludes the derivation.

## B Mathematical analysis and theoretical challenges

## B. 1 Discussion on existing theory

Guarantees of GP-UCB. A notable early work providing theoretical guarantees on Bayesian optimization (without the high-dimensional aspects) is that of Srinivas et al. [33], who considered the cumulative regret $R_{T}=\sum_{t=1}^{T}\left(f\left(x_{\mathrm{opt}}\right)-f\left(x_{t}\right)\right)$ with $x_{\mathrm{opt}}=\arg \max _{x \in \mathcal{X}} f(x)$. In the case of a finite domain $\mathcal{X}$, it was shown that the GP-UCB algorithm with exploration parameter $\beta_{t}=2 \log \left(\frac{|\mathcal{X}| t^{2} \pi^{2}}{6 \delta}\right)$ achieves

$$
\begin{equation*}
R_{T} \leq \sqrt{\frac{8}{\log \left(1+\sigma^{-2}\right)} T \beta_{T} \gamma_{T}} \tag{11}
\end{equation*}
$$

with probability at least $1-\delta$. Here the kernel-dependent quantity $\gamma_{T}$ is known as the information gain, and is defined as the maximum of a mutual information quantity:

$$
\begin{equation*}
\gamma_{T}=\max _{A:|A|=T} I\left(y_{A} ; f_{A}\right) \tag{12}
\end{equation*}
$$

for $A=\left\{x_{1}, \ldots, x_{T}\right\}$ with corresponding function values $f_{A}$ and observations $y_{A}$. Analogous results were presented for the continuous setting in [33] under mild technical assumptions, and explicit bounds on $\gamma_{T}$ were provided for the squared exponential and Matérn kernels.

Extension to non-overlapping additive models. Kandasamy et al. [19] attempted to upper bound the cumulative regret of their algorithm Add-GP-UCB in the high-dimensional setting with additive models. In particular, they sought a bound whose complexity is only exponential in the maximal dimension $d$ of the lowdimensional kernels, instead of the full dimension $D$ as in Srinivas et al. 33. However, they subsequently stated in an updated version of their paper that the proof contains an error 36. In our understanding, the error is due to the fact that the approximated standard deviation $\sum_{i=1}^{M} \sigma_{t-1}^{(i)}$ is different from the true standard deviation $\sigma_{t-1}$, and that the ratio $\frac{\sum_{i=1}^{M} \sigma_{t-1}^{(i)}(x)}{\sigma_{t-1}(x)}$ cannot be upper bounded for all $x$ (see below).
In the parallel independent work of Hoang et al. [25], it was shown that a sufficient condition that leads to sublinear regret bounds in high-dimensional additive models is as follows [25, Assumption 4]: The posterior variance $\sigma_{t}^{(i)}\left(x^{(i)}\right)$ of each component $i$ given the observations can be upper bounded by a constant times $\widehat{\sigma}_{t}^{(i)}\left(x^{(i)}\right)$, defined to be the posterior variance as if the corresponding function $f^{(i)}$ had been sampled directly instead. However, it remains an open problem to determine specific models and kernels for which this assumption is true.

Outline of this appendix. We further discuss the relation between the true and approximate posterior standard deviations in Section B.2 and then provide a novel bound on the information gain for our setting in Section B.3. While the latter is only one step towards attaining a regret bound for G-Add-GP-UCB, it also quantifies the regret bound (11) when GP-UCB is applied to the high-dimensional setting. Unfortunately, GP-UCB is not computationally feasible in high dimensions, so establishing a similar regret bound G-Add-GP-UCB remains an important direction for future research.

## B. 2 Relation between true posterior variance and its approximation

Our algorithm G-Add-GP-UCB is based on an acquisition function that can be computed efficiently in high dimensions. This property comes from the fact that it that can be decomposed into the sum of low-dimensional components (see (6)). Each term in the sum consists of a mean and standard deviation corresponding to a low-dimensional function.
We observe that $\tilde{\mu}_{t-1}(x)=\sum_{j=1}^{M} \mu_{t-1}^{(j)}\left(x^{(j)}\right)=\mu_{t-1}(x)$. Therefore, this way of splitting the posterior mean into several lower dimensional components does not involve any approximation. However, $\tilde{\sigma}_{t-1}(x)=$ $\sum_{j=1}^{M} \sigma_{t-1}^{(j)}\left(x^{(j)}\right) \neq \sigma_{t-1}(x)$ in general; this can be viewed as being due to the non-linearity of the quadratic term $\kappa\left(x_{*}, \mathbf{x}\right) \Delta^{-1} \kappa\left(\mathbf{x}, x_{*}\right)$ in the posterior variance (4).

Our analysis below reveals that

$$
\begin{equation*}
\sum_{j=1}^{M} \sigma_{t-1}^{(j)}\left(x^{(j)}\right) \geq \sigma_{t-1}(x) \tag{13}
\end{equation*}
$$

Thus, this splitting of the posterior standard deviation into low-dimensional components generally over-estimates the true variance. An example where the inequality is strict is as follows: In the case of zero noise, the true posterior standard deviation at a point $x_{\text {evaluated }}$ that has already been evaluated is $\sigma_{t-1}\left(x_{\text {evaluated }}\right)=0$. However in general, $\sum_{j=1}^{M} \sigma_{t-1}^{(j)}\left(x_{\text {evaluated }}\right)>0$. Therefore, the ratio $\frac{\sum_{j=1}^{M} \sigma_{t-1}^{(j)}(x)}{\sigma_{t-1}(x)}$ can sometimes diverge, and it is thus not possible to upper bound it for all $x$.

Derivation of the upper bound (13). The true posterior variance based on observations at the points $\mathbf{x}=\left(x_{1}, \ldots, x_{t}\right)$ is given by

$$
\begin{equation*}
\sigma_{t}(x)^{2}=\kappa(x, x)-\kappa(x, \mathbf{x}) \Delta^{-1} \kappa(\mathbf{x}, x) \tag{14}
\end{equation*}
$$

where $\kappa$ is the full dimensional kernel $\kappa\left(x, x^{\prime}\right)=\sum_{i=1}^{M} \kappa^{(i)}\left(x^{(i)}, x^{\prime(i)}\right), k(x, \mathbf{x})$ and $k(\mathbf{x}, x)$ are the corresponding vectors of kernel values, and $\Delta$ is a matrix such that $\Delta_{i j}=\kappa\left(x_{i}, x_{j}\right)$ for $i, j=1, \ldots, t$. The approximated posterior
variance based on the same points is as follows:

$$
\begin{align*}
\sum_{i=1}^{M} \sigma_{t}\left(x^{(i)}\right)^{2} & =\sum_{i=1}^{M} \kappa^{(i)}\left(x^{(i)}, x^{(i)}\right)-\kappa^{(i)}\left(x^{(i)}, \mathbf{x}^{(i)}\right) \Delta^{-1} \kappa^{(i)}\left(\mathbf{x}^{(i)}, x^{(i)}\right)  \tag{15}\\
& =\kappa(x, x)-\sum_{i=1}^{M} \kappa^{(i)}\left(x^{(i)}, \mathbf{x}^{(i)}\right) \Delta^{-1} \kappa^{(i)}\left(\mathbf{x}^{(i)}, x^{(i)}\right) \tag{16}
\end{align*}
$$

under the notation following (4).
By definition, the matrix $\Delta$ is symmetric and positive definite, and hence so is the matrix $\Delta^{-1}$. We can thus define a norm induced by this matrix on the space $\mathbb{R}^{t}$; for $\vec{k} \in \mathbb{R}^{t}$, we have

$$
\begin{equation*}
\|\vec{k}\|_{\Delta^{-1}}^{2}=\vec{k}^{T} \Delta^{-1} \vec{k} \tag{17}
\end{equation*}
$$

The fact that $\Delta^{-1}$ is symmetric positive definite implies that this has all the properties of a norm. For any $x \in \mathbb{R}^{D}$, we define the $t$-dimensional vector $\vec{k}^{(i)}(x)$ as $\vec{k}^{(i)}(x)_{j}=\kappa^{(i)}\left(x^{(i)}, x_{j}^{(i)}\right)$. We also recall that $\kappa(x, x)=1$ for all $x \in \mathbb{R}^{D}$. Using this notation, we can rewrite the expressions for the true and approximated posterior variances as

$$
\begin{equation*}
\sigma_{t}(x)^{2}=1-\left\|\sum_{i=1}^{M} \vec{k}^{(i)}(x)\right\|_{\Delta^{-1}}^{2} \tag{18}
\end{equation*}
$$

and

$$
\begin{equation*}
\sum_{i=1}^{M} \sigma_{t}\left(x^{(i)}\right)^{2}=1-\sum_{i=1}^{M}\left\|\vec{k}^{(i)}(x)\right\|_{\Delta^{-1}}^{2} \tag{19}
\end{equation*}
$$

By the triangle inequality, we have

$$
\begin{aligned}
\sigma_{t}(x)^{2} & =1-\left\|\sum_{i=1}^{M} \vec{k}^{(i)}(x)\right\|_{\Delta^{-1}}^{2} \\
& \leq 1-\left(\sum_{i=1}^{M}\left\|\vec{k}^{(i)}(x)\right\|_{\Delta^{-1}}\right)^{2} \\
& \leq 1-\sum_{i=1}^{M}\left\|\vec{k}^{(i)}(x)\right\|_{\Delta^{-1}}^{2} \\
& =\sum_{i=1}^{M} \sigma_{t}\left(x^{(i)}\right)^{2}
\end{aligned}
$$

As $\sigma_{t}\left(x^{(i)}\right) \geq 0 \forall x \in \mathbb{R}^{D}$, we have that $\sum_{i=1} \sigma_{t}\left(x^{(i)}\right)^{2} \leq\left(\sum_{i=1} \sigma_{t}\left(x^{(i)}\right)\right)^{2}$, which implies

$$
\begin{equation*}
\sigma_{t}(x) \leq \sum_{i=1}^{M} \sigma_{t}\left(x^{(i)}\right) \tag{20}
\end{equation*}
$$

as desired.
Numerical evaluation. In order to observe the difference between the true posterior variance $\sigma_{t}^{2}$ and the approximated variance $\left(\sum_{j=1}^{M} \sigma_{t-1}^{(j)}\right)^{2}$, we generate a 10 dimensional synthetic function via Gaussian processes with the star dependency graph shown in Figure 2. We first evaluate this function at 200 randomly-selected points. Based on these observations, we then evaluate the true and approximated posterior variances at 1000 randomly selected points. We then compare the results obtained with these two different methods (Figure 8).

From the figure on the right, we can observe that the obtained values can be significantly different. However, there is a clear correlation between these two quantities, in the sense that points with higher true variance tend to also have a higher approximated variance. The figure on the left shows that the ratio between approximated and true variance increases as the true variance becomes smaller. This approximation error strongly depends on the decomposition, and in particular on $M$. The higher the value of $M$, the higher the approximation error.


Figure 8: Evaluation of the posterior standard deviation from the full dimensional kernel (true posterior variance), vs. separating the kernel into lower dimensional additive components (approximated posterior variance). Left: Ratio between the two computed posterior standard variations as a function of the true posterior standard variation. Right: Approximated posterior standard deviation versus true posterior standard deviation.

## B. 3 Bounding the information gain with overlapping groups

As outlined above, the information gain $\gamma_{T}$ in (12) plays a crucial role in the regret bounds for Bayesian optimization. While we do not claim any regret bounds for G-Add-GP-UCB, bounding $\gamma_{T}$ may provide an initial step towards this, and also allows us to understand the performance of GP-UCB with our structured kernels (cf., (11)). We provide such a bound for our setting, focusing on the squared exponential kernel, and ultimately showing that $\gamma_{T}=O\left(D d^{d}(\log T)^{d+1}\right)$ analogously to [19]. Here $d$ is the highest dimension of any low-dimensional function in the additive decomposition.

We follow the high-level steps of [19] with suitable modifications for our setting. It was shown in [26] that under some mild assumptions on the target function $f$, the maximal information gain can be bounded as

$$
\begin{equation*}
\gamma_{T} \leq \inf _{\tau}\left(\frac{1 / 2}{1-e^{-1}} \max _{r \in\{1, \ldots, T\}}\left(T_{*} \log \left(r n_{T} / \eta^{2}\right)+C_{9} \eta^{2}\left(1-\frac{r}{T}\right)\left(T^{r+1} B_{\kappa}\left(T_{*}\right)+1\right) \log T\right)+O\left(T^{1-\tau / D}\right)\right) \tag{21}
\end{equation*}
$$

for any $T_{*} \in\left\{1, \ldots, \min \left(T, n_{T}\right)\right\}$, where $C_{9}=4 D+2, n_{T}=C_{9} T^{\tau} \log T$, and $B_{\kappa}\left(T_{*}\right)=\sum_{s>T_{*}} \lambda_{s}$. Here $\left\{\lambda_{n}\right\}_{n \in \mathbb{N}}$ are the eigenvalues of $\kappa$ with respect to the uniform distribution.

In order to bound $\gamma_{T}$, it therefore suffices to bound $B_{\kappa}\left(T_{*}\right)$, i.e., to bound the sum of the eigenvalues of $\kappa$ at the tail. Unlike the setting of [19], the eigenfunctions corresponding to different kernels $\kappa^{(i)}$ and $\kappa^{(j)}$ are not necessarily orthogonal, since overlaps between kernel variables are possible. To circumvent this difficulty, we can make use of Weyl's inequality.
Lemma 1. (Weyl's inequality) Let $H, P \in \mathbb{R}^{n \times n}$ be two Hermitian matrices, and define $M=H+P$. Let $\mu_{i}, \nu_{i}, \rho_{i}, i=1, \ldots, n$ be the eigenvalues of $M, H$ and $P$ respectively such that $\mu_{1} \geq \ldots \geq \mu_{n}, \nu_{1} \geq \ldots \geq \nu_{n}$ and $\rho_{1} \geq \ldots \geq \rho_{n}$. Then for all $i \geq r+s-1$, we have

$$
\begin{equation*}
\mu_{i} \leq \nu_{r}+\rho_{s} \tag{22}
\end{equation*}
$$

This result immediately generalizes to a sum with an arbitrary number of matrices. In particular, we will use the following corollary.
Corollary 1. Let $K_{i} \in \mathbb{R}^{n \times n}, i=1, \ldots, M$ be Hermitian matrices, and define $K=\sum_{i=1}^{M} K_{i}$. Let $\left\{\lambda_{j}^{(i)}\right\}_{j=1, \ldots, n}$, be the eigenvalues of $K_{i}$ such that $\lambda_{1}^{(i)} \geq \ldots \geq \lambda_{n}^{(i)} \forall i=1, \ldots, M$, and let $\left\{\lambda_{i}\right\}_{i=1, \ldots, n}$ be the eigenvalues of $K$ such that $\lambda_{1} \geq \ldots \geq \lambda_{n}$. Then for all $k \in \mathbb{N}$ such that $k M+1 \leq n$, we have

$$
\begin{equation*}
\lambda_{k M+1} \leq \sum_{i=1}^{M} \lambda_{k+1}^{(i)} \tag{23}
\end{equation*}
$$

Let $\left\{\lambda_{s}\right\}_{s \in \mathbb{N}}, \lambda_{1} \geq \lambda_{2} \geq \ldots$ denote the eigenvalues of $\kappa$, and for all $j=1, \ldots, M$, let $\left\{\lambda_{s}^{(j)}\right\}_{s \in \mathbb{N}}, \lambda_{1}^{(j)} \geq \lambda_{2}^{(j)} \geq \ldots$ denote the eigenvalues of $\kappa^{(j)}$. It was shown by Seeger et al. 37] that the eigenvalues $\lambda_{s}^{(j)}$ for the square
exponential kernel $\kappa^{(j)}$ satisfy $\lambda_{s}^{(j)} \leq c^{d} B^{s^{1 / d}}, B<1$, where each $\kappa^{(j)}$ is a $d_{j}$-dimensional kernel, and $d_{j} \leq d$. Defining $T_{+}=\left\lfloor\frac{T_{*}}{M}\right\rfloor$ we have the following:

$$
\begin{align*}
B_{\kappa}\left(T_{*}\right) & =\sum_{s>T_{*}} \lambda_{s}  \tag{24}\\
& \leq \sum_{k>T_{+}} \sum_{i=1}^{M} \lambda_{(k-1) M+l}  \tag{25}\\
& \leq \sum_{k>T_{+}} \sum_{i=1}^{M} \sum_{j=1}^{M} \lambda_{k}^{(j)}  \tag{26}\\
& \leq M^{2} c^{d} \sum_{k>T_{+}} B^{k^{1 / d}} \tag{27}
\end{align*}
$$

where the second line uses the fact that the eigenvalues are increasingly ordered, the third line follows from Weyl's inequality, and the final line follows from the bound on the tail eigenvalues given in [37].

The rest of the proof follows via a similar analysis to [19. One difference is that we get an extra $M$ term in our bound for $B_{\kappa}$ compared to the setting of [19]. However, this does not affect the bound for $\gamma_{T}$, since the leading term on the right hand side of (21) is $T_{*} \log \left(r n_{T} / \eta^{2}\right)$ which does not involve $B_{\kappa}$. We thus obtain the same bound for $\gamma_{T}$ as in 19, namely,

$$
\begin{equation*}
\gamma_{T}=O\left(D d^{d}(\log T)^{d+1}\right) \tag{28}
\end{equation*}
$$

Note that this bound only has linear dependence on the dimension $D$, while being exponential in the maximal group size $d$.

## C Astrophysical data experiment

In this appendix, we consider an additional experiment on real-world data that aims at estimate a set of 9 cosmological parameters (e.g., Hubble's constant, proportion of helium in the universe, etc) in order to best match reality. These constants are involved in the theoretical model of physics, but are estimated experimentally. To do so, programs model the dynamics of the universe given these parameters, and compare the results of the simulations with the observed data.

For each set of parameters, we can compute the likelihood that the chosen parameters match the reality. The aim is thus to find the set of parameters that maximize this likelihood, or equivalently that minimize the negative log-likelihood. We use the LRG DR7 Likelihood Software released by NASA 1 in order to compute likelihoods given these cosmological parameters based on experimental data released by the Sloan Digital Sky Survey.
We note that this data was used by both Kandasamy et al. [19] and Gardner et al. [22] for testing high-dimensional BO algorithms, but it was used in somewhat different ways. We adopt the approach of [22], and we avoid adding additional "dummy dimensions" as in [19.

The software provides a set of parameters which achieves a negative log-likelihood of 23.7. We thus apply the two Bayesian algorithms "Overlap" and "No Overlap" in a range of $75 \%-125 \%$ of this default set of parameters. Unlike the previous real world experiment, we do not set a fixed dependency graph for the "Overlap" model and learn it throughout the algorithm using Gibbs sampling ( $c f$. ., Section 44). Similarly, for the "No Overlap" model, we use the Gibbs sampling approach of [21], placing no "hard constraints" (i.e., choices of $M$ and $d$ in [19]).
The remaining parameters for the Bayesian optimization algorithms are the same as for the first real-world experiment (Table 2), except that we run the simulation for 1000 iterations. The results are shown in Figure 9 We observe that both algorithms achieve a higher likelihood than with the default parameters, and that the "Overlap" algorithm achieves a higher likelihood than the "No Overlap" one across the entire time horizon. In particular, by the final iteration, the gap to "random" has increased from approximately 0.09 to 0.15 , i.e., an increase of over $60 \%$.

[^0]

Figure 9: Results on the astrophysical experiment. The lower the vertical axis value, the more likely it is that the chosen constants match the observed data.


[^0]:    ${ }^{1}$ http://lambda.gsfc.nasa.gov/toolbox/lrgdr/

