# Functional Causal Bayesian Optimization

**Limor Gultchin**[1,*] **Virginia Aglietti**[2,*] **Alexis Bellot**[2] **Silvia Chiappa**[2]

[1]University of Oxford, The Alan Turing Institute, Work done at DeepMind, London, UK
[2]DeepMind, London, UK

## Abstract

We propose *functional causal Bayesian optimization* (fCBO), a method for finding interventions that optimize a target variable in a known causal graph. fCBO extends the CBO family of methods to enable functional interventions, which set a variable to be a deterministic function of other variables in the graph. fCBO models the unknown objectives with Gaussian processes whose inputs are defined in a reproducing kernel Hilbert space, thus allowing to compute distances among vector-valued functions. In turn, this enables to sequentially select functions to explore by maximizing an expected improvement acquisition functional while keeping the typical computational tractability of standard BO settings. We introduce graphical criteria that establish when considering functional interventions allows attaining better target effects, and conditions under which selected interventions are also optimal for conditional target effects. We demonstrate the benefits of the method in a synthetic and in a real-world causal graph.

## 1 INTRODUCTION

Finding interventions in a system that optimize a target variable is key to many scientific disciplines, including medicine, biology, and social sciences. Causal graphs [Pearl, 2000, Koller and Friedman, 2009], in which an intervention on a variable is represented as modifying the casual influence from its incoming edges, offer a powerful tool for dealing with the effects of interventions, and are therefore increasingly integrated into approaches to learning optimal policies such as bandits [Lattimore et al., 2016, Lee and Bareinboim, 2018, 2019, Lu et al., 2020, Nair et al., 2021, De Kroon et al., 2022], reinforcement learning [Lu et al., 2018, Zhang and Bareinboim, 2019, Zhang, 2020, Gasse et al., 2021, Zhang and Bareinboim, 2022], and Bayesian optimization [Aglietti et al., 2020, 2021, Sussex et al., 2023].

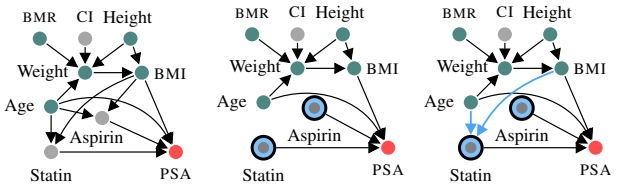


Figure 1: *Left:* Graph representing causal relationships between prostate specific antigen (PSA) and other variables. Red, grey and green nodes indicate target, intervenable, non-intervenable variables respectively. *Middle:* Modified graph describing a policy made of hard interventions on Aspirin and Statin. *Right:* Modified graph describing a policy with an intervention on Statin that retains dependence on Age and BMI.

Most works in causal Bayesian optimization (CBO) have focused on the *hard intervention* $\text{do}(X = x)$, which consists in setting variable $X$ to a constant value $x$. However, in many practical scenarios the investigator may be able to implement policies that also contain other types of interventions. Consider, for example, the graph in Fig. 1(left) representing causal relationships between prostate specific antigen (PSA) and other variables. An investigator wishing to find a policy for prescribing Aspirin and Statin dosages, as well as Calories Intake (CI), that minimizes PSA might be able to consider, in addition to policies made of only hard interventions (as the one represented in Fig. 1(middle)), also policies where e.g. Statin dosage retains a dependence on Age and BMI (as the one represented in Fig. 1(right)).

Contextual interventions are achieved in Arsenyan et al. [2023] and in Sussex et al. [2023] by searching for different hard interventions in separate sub-groups defined by some contexts and by inducing changes in the parametrization of a node's conditional distribution via action variables, respectively. However, the first approach learns an implicit

[*]Equal contribution.

*Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence* (UAI 2023), PMLR 216:756–765.

mapping between contexts and intervention values, and requires extrapolating to unseen or rarely explored areas of the context space; while the second approach can only induce some modifications of the parametrization and does not allow choice of context.

In this work, we introduce an extension of the CBO family of methods that considers a more flexible and general type of contextual intervention, consisting in making variable $X$ a *deterministic function* of other nodes in the graph. Such a *functional intervention* is implemented via new techniques for computing distances among functions of different variables. Our contributions can be summarized as follows:

- We formalize the problem of finding policies made of hard and functional interventions optimizing the expectation of a target variable as the *functional causal global optimization* (fCGO) problem.
- We introduce two graphical criteria that establish when functional interventions could be necessary to solve the fCGO problem and when policies made of only hard interventions are sufficient, respectively.
- We introduce conditions in which a policy solving the fCGO problem also optimizes conditional expectations of the target variable.
- We propose *functional causal Bayesian optimization* (fCBO), a method for solving the fCGO problem that models the expectation of the target variable under each policy scope with a Gaussian process model whose inputs are defined in a reproducing kernel Hilbert space.
- We validate fCBO in a synthetic and in a real-world setting with respect to target effects, conditional target effects, and costs of interventions.

## 2 BACKGROUND AND SETTING

We consider a system of observable random variables $\boldsymbol{V}$ with *target variable* $Y \in \boldsymbol{V}$ and *intervenable variables* $\boldsymbol{I} \subseteq \boldsymbol{V} \backslash Y$, and the problem of finding a subset of $\boldsymbol{I}$ and interventions on it that optimize the expectation of $Y$. Our goal is to introduce a method that allows two types of interventions on a variable $X \in \boldsymbol{I}$: (i) the *hard intervention* $\text{do}(X = x)$ consisting in setting $X$ to value $x$; and (ii) the *functional intervention*[1] $X = \pi_{X|\boldsymbol{C}_X}(\boldsymbol{C}_X)$ that makes $X$ a deterministic function of a set of variables $\boldsymbol{C}_X \subseteq \boldsymbol{V} \backslash \{X, Y\}$, called the *context* of $X$, where $\pi_{X|\boldsymbol{C}_X} : \mathcal{R}_{\boldsymbol{C}_X} \mapsto \mathcal{R}_X$ with e.g. $\mathcal{R}_{\boldsymbol{C}_X}$ indicating the range of $\boldsymbol{C}_X$. Both hard and functional interventions make $X$ a deterministic function of a context $\boldsymbol{C}_X$ (the hard intervention $\text{do}(X = x)$ can be viewed as a functional intervention with empty context $\boldsymbol{C}_X = \emptyset$, setting $X$ to value $x = \pi_{X|\emptyset}(\emptyset)$ where $\pi_{X|\emptyset} : \emptyset \mapsto x$ is the empty function), and are therefore referred to as *deterministic interventions* Lee and Bareinboim [2020].

We specify the system's behavior using a *structural causal model* (SCM) $\mathcal{M}$ defined by the tuple $\langle \boldsymbol{V}, \boldsymbol{U}, \mathcal{F}, p(\boldsymbol{U}) \rangle$, where $\boldsymbol{U}$ is a set of exogenous, mutually-independent, unobserved random variables with distribution $p(\boldsymbol{U})$, and $\mathcal{F} = \{f_V\}_{V \in \boldsymbol{V}}$ is a set of deterministic functions such that $V = f_V(\text{pa}(V), \boldsymbol{U}_V)$ with $\text{pa}(V) \subseteq \boldsymbol{V} \backslash V$ and $\boldsymbol{U}_V \subseteq \boldsymbol{U}$, $\forall V \in \boldsymbol{V}$. A deterministic intervention on $X$ therefore replaces $f_X$ with $\pi_{X|\boldsymbol{C}_X}$.

$\mathcal{M}$ has associated a *directed graph*, which we assume to be *acyclic*[2], with nodes $\boldsymbol{V} \cup \boldsymbol{U}$ and with an edge from $A$ to $B$ if $A \in \text{pa}(B)$ or $A \in \boldsymbol{U}_B$. A node $A$ with an edge into $B$ is called a *parent* or *direct cause* of $B$ (in this case $B$ is called a *child* of $A$). A node $A$ with a directed path ending at $B$ is called an *ancestor* of $B$ (in this case $B$ is called a *descendant* of $A$). We consider the projection of this graph into the graph that contains only nodes $\boldsymbol{V}$ and that has a directed edge from $V$ to $W$ if $V$ is a parent of $W$ and a bi-directed edge between $V$ and $W$ if $\boldsymbol{U}_V \cap \boldsymbol{U}_W \neq \emptyset$ ($\boldsymbol{U}_V \cap \boldsymbol{U}_W$ is an unobserved *confounder* between $V$ and $W$), and refer to it as *causal graph* associated with $\mathcal{M}$. Given a causal graph $\mathcal{G}$, we say that $\mathcal{M}$ is *compatible* with $\mathcal{G}$ if all edges that are in the causal graph associated with $\mathcal{M}$ are also in $\mathcal{G}$. We indicate the set of parents, ancestors, and descendants of $V$ in $\mathcal{G}$ with $\text{pa}_{\mathcal{G}}(V)$, $\text{an}_{\mathcal{G}}(V)$ and $\text{de}_{\mathcal{G}}(V)$, respectively. We indicate the nodes connected to $V$ by a bi-directed edge with $\text{sp}_{\mathcal{G}}(V)$. We refer to the joint distribution of $\boldsymbol{V}$ determined by $p(\boldsymbol{U})$, which we denote by $p(\boldsymbol{V})$, as *observational distribution*.

The space of deterministic interventions for a casual graph $\mathcal{G}$ can be formalized using the concepts of *mixed policy scope* (MPS) and *deterministic mixed policy* (DMP) introduced in Lee and Bareinboim [2020].

**Definition 2.1** (Mixed Policy Scope (MPS))**.** A mixed policy scope $\mathcal{S}$ for a causal graph $\mathcal{G}$ is a collection of pairs $\langle X, \boldsymbol{C}_X \rangle$ such that (i) $X \in \boldsymbol{I}$, $\boldsymbol{C}_X \subseteq \boldsymbol{V} \backslash \{X, Y\}$; and (ii) the graph $\mathcal{G}_{\mathcal{S}}$ obtained by removing from $\mathcal{G}$ the incoming edges into $X$ and by adding to $\mathcal{G}$ directed edges from $\boldsymbol{C}_X$ to $X$, for every $\langle X, \boldsymbol{C}_X \rangle \in \mathcal{S}$, is acyclic.

An MPS specifies the variables in $\boldsymbol{I}$ on which interventions are performed and their contexts. For example, MPS $\mathcal{S} = \{\langle \text{Aspirin}, \emptyset \rangle, \langle \text{Statin}, \{\text{Age}, \text{BMI}\} \rangle\}$ for $\mathcal{G}$ in Fig. 1(left) specifies that interventions are performed on Aspirin and Statin, and with context $\emptyset$ and $\{\text{Age}, \text{BMI}\}$ respectively, as graphically represented in Fig. 1(right).

**Definition 2.2** (Deterministic Mixed Policy (DMP))**.** A deterministic mixed policy $\pi_{\mathcal{S}}$ compatible with MPS $\mathcal{S}$ is defined

[1] Functional interventions are also called *conditional interventions* in Correa and Bareinboim [2020a,b].

[2] A directed graph is acyclic if it has no *directed paths* starting and ending at the same node. A directed path is a sequence of linked nodes whose edges are directed and point from preceding towards following nodes in the sequence.

as $\pi_{\mathcal{S}} = \{\pi_{X|\boldsymbol{C}_X}\}_{\langle X,\boldsymbol{C}_X\rangle \in \mathcal{S}\setminus\mathcal{S}_{\text{hard}}} \bigcup \{\pi_{X|\emptyset}(\emptyset)\}_{\langle X,\boldsymbol{C}_X\rangle \in \mathcal{S}_{\text{hard}}}$, where $\pi_{X|\boldsymbol{C}_X} : \mathcal{R}_{\boldsymbol{C}_X} \mapsto \mathcal{R}_X$, $\pi_{X|\emptyset}(\emptyset)$ denotes the value returned by the empty function, and $\mathcal{S}_{\text{hard}} = \{\langle X, \boldsymbol{C}_X\rangle \in \mathcal{S} \colon \boldsymbol{C}_X = \emptyset\}$.

A DMP specifies the function $\pi_{X|\boldsymbol{C}_X}$ or the value $\pi_{X|\emptyset}(\emptyset)$ that replaces $f_X \in \mathcal{F}$ in $\mathcal{M}$, $\forall \langle X, \boldsymbol{C}_X\rangle \in \mathcal{S}$. The replacements induce a variant $\mathcal{M}_{\pi_{\mathcal{S}}}$ of $\mathcal{M}$ with joint distribution over $\boldsymbol{V}$ denoted by $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$. We refer to $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$ as *interventional distribution* induced by $\pi_{\mathcal{S}}$, and to an observation from $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$ as an *interventional data sample*.

## 3 FCGO PROBLEM

Let $\mu^Y_{\pi_{\mathcal{S}}} = \mathbb{E}_{p_{\pi_{\mathcal{S}}}}[Y]$ denote the expectation of $Y$ w.r.t. the interventional distribution induced by $\pi_{\mathcal{S}}$, which we refer to as the *target effect*. Our goal is to introduce a method for solving the problem of minimizing $\mu^Y_{\pi_{\mathcal{S}}}$ over the space $\Sigma$ of MPSs for $\mathcal{G}$ and the space $\Pi_{\mathcal{S}}$ of DMPs that are compatible with MPS $\mathcal{S}$, formally defined below.

**Definition 3.1.** (fCGO problem) The *functional causal global optimization* (fCGO) problem is the problem of identifying a tuple $(\mathcal{S}^*, \pi^*_{\mathcal{S}^*})$ such that

$$\mathcal{S}^*, \pi^*_{\mathcal{S}^*} = \underset{\mathcal{S}\in\Sigma, \pi_{\mathcal{S}}\in\Pi_{\mathcal{S}}}{\arg\min}\ \mu^Y_{\pi_{\mathcal{S}}}. \tag{1}$$

Importantly, Proposition 1 in Lee and Bareinboim [2020] implies that the target effect $\mu^Y_{\pi^*_{\mathcal{S}^*}}$ given by a solution of the fCGO problem $(\mathcal{S}^*, \pi^*_{\mathcal{S}^*})$ equals the one that would be obtained by also considering *stochastic interventions* [Correa and Bareinboim, 2020a].

The fCGO problem extends the *causal global optimization* (CGO) problem defined in Aglietti et al. [2020] that only considers hard interventions. In Section 3.1 we introduce graphical criteria that establish when only considering hard interventions might lead to a bigger target effect and when this is not the case. In addition, in Section 3.2 we introduce conditions under which a policy solving the fCGO problem is also optimal for conditional target effects.

Solving the fCGO problem requires computing distances between functions defined over different contexts. In Section 4.2 we propose to model each target effect via a Gaussian process whose kernel allows computing such distances. We discuss how this approach enables us to keep the computational tractability of standard Bayesian optimization (BO) methods while allowing to flexibly specify functional interventions.

### 3.1 HARD INTERVENTIONS (SUB-)OPTIMALITY

Let $\Sigma_{\text{hard}}$ denote the set of MPSs in $\Sigma$ that contain only hard interventions, i.e. $\Sigma_{\text{hard}} = \{\mathcal{S} \in \Sigma \colon \mathcal{S} = \mathcal{S}_{\text{hard}}\}$. In this section, we introduce graphical criteria that establish when restricting the search space in the fCGO problem from $\Sigma$ to $\Sigma_{\text{hard}}$ might lead to a bigger target effect and when this is not the case, thereby informing the investigator about when functional interventions should be considered. The proofs are given in Section 1 of the supplementary material.

**Proposition 3.2** (Sub-optimality of hard interventions)**.** *Let $\mathcal{G}$ be a causal graph such that (i) $\exists C \in pa_{\mathcal{G}}(Y)$ with $C \notin \boldsymbol{I}$; or (ii) $\exists C \in sp_{\mathcal{G}}(Y)$. If $\exists X \in an_{\mathcal{G}}(Y) \cap \boldsymbol{I}$ such that $\{\langle X, C\rangle\}$ is an* MPS*, then there exists at least one* SCM *compatible with $\mathcal{G}$ for which $\min_{\mathcal{S}\in\Sigma_{hard}, \pi_{\mathcal{S}}\in\Pi_{\mathcal{S}}} \mu^Y_{\pi_{\mathcal{S}}} > \min_{\mathcal{S}\in\Sigma, \pi_{\mathcal{S}}\in\Pi_{\mathcal{S}}} \mu^Y_{\pi_{\mathcal{S}}}$.*

**Proposition 3.3** (Optimality of hard interventions)**.** *In a casual graph $\mathcal{G}$, if $pa_{\mathcal{G}}(Y) \subseteq \boldsymbol{I}$ and $sp_{\mathcal{G}}(Y) = \emptyset$ there exists a* DMP *compatible with* MPS *$\mathcal{S} = \{\langle X, \emptyset\rangle : X \in pa_{\mathcal{G}}(Y)\}$ that solves the f*CGO *problem.*

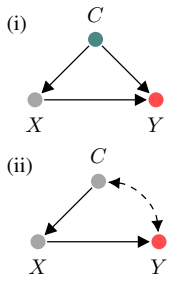


Proposition 3.2 captures two conditions for sub-optimality of hard interventions: the existence of a non-intervenable variable $C$ in $\text{pa}_{\mathcal{G}}(Y)$ that can serve as context for a functional intervention on a variable $X$, as in the causal graph (i) on the right (for which $\boldsymbol{I} = \{X\}$); and the existence of a variable $C$ with an unobserved confounder between it and $Y$ that can serve as context for a functional intervention on a variable $X$, as in the casual graph (ii) on the right (for which $\boldsymbol{I} = \{X, C\}$). In both cases, a hard intervention on $X$ would cut the paths from $X$ to $Y$ passing through $C$ (i.e. $X \leftarrow C \rightarrow Y$ and $X \leftarrow C \leftrightarrow Y$ respectively). Instead, a functional intervention on $X$ with context $C$ would keep such paths open and therefore could assign intervention values to $X$ informed by values of $C$, potentially leading to a smaller target effect. Below, we provide two SCMs and functional interventions for which this is the case.

Consider graph (i), with SCM $\mathcal{M}$ with $\boldsymbol{U} = \{U_C, U_X, U_Y\}$ such that $p(U_C) = p(U_X) = \mathcal{N}(0,1)$ and $p(U_Y) = \mathcal{N}(1,1)$, and functional assignments $C = U_C, X = CU_X, Y = CXU_Y$. $\Sigma_{\text{hard}} = \{\mathcal{S}^1 = \{\langle X, \emptyset\rangle\}\}$ with DMP $\pi_{\mathcal{S}^1} = \{x = \pi_{X|\emptyset}(\emptyset)\}$ induces the modified SCM $\mathcal{M}_{\pi_{\mathcal{S}^1}}$ where $Y = U_C x U_Y$ and $\mu^Y_{\pi_{\mathcal{S}^1}} = 0$. In contrast, MPS $\mathcal{S} = \{\langle X, C\rangle\}$ with DMP $\pi_{\mathcal{S}} = \{\pi_{X|C}(C) = -1/C\}$ induces $\mathcal{M}_{\pi_{\mathcal{S}}}$ with $Y = -U_Y$, giving $\mu^Y_{\pi_{\mathcal{S}}} = -1.0$. Therefore, $\pi_{\mathcal{S}}$ achieves a smaller target effect than $\pi_{\mathcal{S}^1}$.

Consider graph (ii), with SCM $\mathcal{M}$ with $\boldsymbol{U} = \{U_{CY}, U_X, U_Y\}$ such that $p(U_{CY}) = p(U_X) = \mathcal{N}(0,1)$ and $p(U_Y) = \mathcal{N}(1,1)$, and functional assignments $C = U_{CY}, X = CU_X, Y = U_{CY}XU_Y$. In this case, $\Sigma_{\text{hard}} = \{\mathcal{S}^1 = \{\langle X, \emptyset\rangle\}, \mathcal{S}^2 = \{\langle C, \emptyset\rangle\}, \mathcal{S}^3 = \{\langle X, \emptyset\rangle, \langle C, \emptyset\rangle\}\}$ with DMPs $\pi_{\mathcal{S}^1} = \{x = \pi_{X|\emptyset}(\emptyset)\}$, $\pi_{\mathcal{S}^2} = \{c = \pi_{C|\emptyset}(\emptyset)\}$, and $\pi_{\mathcal{S}^3} = \{x = \pi_{X|\emptyset}(\emptyset), c = \pi_{C|\emptyset}(\emptyset)\}$. In $\mathcal{M}_{\pi_{\mathcal{S}^1}}$, $Y = xU_{CY}U_Y$ thus $\mu^Y_{\pi_{\mathcal{S}^1}} = 0$. In $\mathcal{M}_{\pi_{\mathcal{S}^2}}$, $Y = cU_X U_{CY} U_Y$

thus $\mu^Y_{\pi_{\mathcal{S}^2}} = 0$. In $\mathcal{M}_{\pi_{\mathcal{S}^3}}$, $Y = xU_{CY}U_Y$ thus $\mu^Y_{\pi_{\mathcal{S}^3}} = 0$. In contrast, MPS $\mathcal{S} = \{\langle X, C\rangle\}$ with DMP $\pi_{\mathcal{S}} = \{\pi_{X|C}(C) = -1/C\}$ induces $\mathcal{M}_{\pi_{\mathcal{S}}}$ with $Y = -U_Y$ giving $\mu^Y_{\pi_{\mathcal{S}}} = -1$. Therefore $\pi_{\mathcal{S}}$ achieves a smaller target effect than any other DMP containing only hard interventions.

### 3.2 CONDITIONAL TARGET EFFECTS

In addition to potentially leading to a smaller target effect, considering functional interventions allows to deal with settings in which the investigator might wish to minimize the target effect *conditioned* on a set of variables. For instance, in the health example of Fig. 1(left), the investigator might want to find interventions minimizing the expectation of PSA in a given population as well as in a specific sub-group made of individuals aged over 65, i.e. $\mu^{\text{PSA}}_{\pi_{\mathcal{S}},\text{Age}>65} := \mathbb{E}_{p_{\pi_{\mathcal{S}}}}[\text{PSA} \mid \text{Age} > 65]$ – since a high percentage of prostate cancer cases are diagnosed within this sub-group [Rawla, 2019] – while still not negatively affecting individuals of other ages. Such settings can be formalized as wishing to minimize the conditional target effect $\mu^Y_{\pi_{\mathcal{S}},\boldsymbol{C}=\boldsymbol{c}} = \mathbb{E}_{p_{\pi_{\mathcal{S}}}}[Y \mid \boldsymbol{C} = \boldsymbol{c}]$ for $\boldsymbol{C} \subset \boldsymbol{V}\backslash Y$ and $\boldsymbol{c} \in \mathcal{R}_{\boldsymbol{C}}$.

Let $\boldsymbol{X}_{\mathcal{S}}$ denote the intervention variables included in MPS $\mathcal{S}$, i.e. $\boldsymbol{X}_{\mathcal{S}} = \{X \colon \langle X, \boldsymbol{C}_X\rangle \in \mathcal{S}\}$, and $\boldsymbol{C}^{\mathcal{S}}_X$ the context variables in MPS $\mathcal{S}$ for an intervention on $X$. Unlike when considering only hard interventions, the following proposition shows that, under some conditions, a solution of the fCGO problem also minimizes $\mu^Y_{\pi_{\mathcal{S}},\boldsymbol{C}=\boldsymbol{c}}$ in a restricted MPSs space (the proof is given in Section 1 of the supplementary material).

**Proposition 3.4** (Optimizing conditional target effects)**.** *If* $\mathcal{S}^*, \pi^*_{\mathcal{S}^*} = \arg\min_{\mathcal{S}\in\Sigma,\pi_{\mathcal{S}}\in\Pi_{\mathcal{S}}} \mu^Y_{\pi_{\mathcal{S}}}$*, then* $\mathcal{S}^*, \pi^*_{\mathcal{S}^*} = \arg\min_{\mathcal{S}\in\Sigma^{\boldsymbol{C}},\pi_{\mathcal{S}}\in\Pi_{\mathcal{S}}} \mu^Y_{\pi_{\mathcal{S}},\boldsymbol{C}=\boldsymbol{c}}\ \forall \boldsymbol{C} \subset \boldsymbol{V}\backslash Y$ *such that* $\boldsymbol{C} \cap de_{\mathcal{G}}(\boldsymbol{I}) = \emptyset$ *and* $\forall \boldsymbol{c} \in \mathcal{R}_{\boldsymbol{C}}$ *with* $\Sigma^{\boldsymbol{C}} = \{\mathcal{S} \in \Sigma : \boldsymbol{X}_{\mathcal{S}} = \boldsymbol{X}_{\mathcal{S}^*}$ *and* $\{\langle X, \boldsymbol{C}^{\mathcal{S}^*}_X \cup \boldsymbol{C}^{\mathcal{S}}_X \cup \boldsymbol{C}\rangle : X \in \boldsymbol{X}_{\mathcal{S}^*}\}$ *is an* MPS$\}$.

## 4 METHODOLOGY

We propose to solve the fCGO problem using the *functional causal Bayesian optimization* (fCBO) method summarized in Algorithm 1, which assumes known casual graph $\mathcal{G}$ and continuous variables $\boldsymbol{V}$. fCBO first reduces the search space from $\Sigma$ to a subset $\mathbb{M}_\Sigma$ using the `NRMPSReduce` procedure described in Section 4.1; and then solves the minimization problem in Eq. (1) using a Gaussian process (GP) $g_{\mathcal{S}}(\pi_{\mathcal{S}})$ to model the unknown target effect $\mu^Y_{\pi_{\mathcal{S}}}, \forall \mathcal{S} \in \mathbb{M}_\Sigma$, as described in Section 4.2, with the following sequential strategy. At each trial $t = 1, \dots, T$: (1) MPS $\mathcal{S}_t$ and DMP $\pi^t_{\mathcal{S}_t}$ are selected via the expected improvement acquisition functional (fEI) described in Section 4.3; (2-3) a set of $S$ interventional data samples is obtained and used to compute a sample mean estimate, $\hat{\mu}^Y_{\pi^t_{\mathcal{S}_t}}$, of $\mu^Y_{\pi^t_{\mathcal{S}_t}}$; (4) $(\pi^t_{\mathcal{S}_t}, \hat{\mu}^Y_{\pi^t_{\mathcal{S}_t}})$ is added to the interventional dataset $\mathcal{D}^I_{\mathcal{S}_t}$ of the MPS $\mathcal{S}_t$; (5) the posterior distribution of the GP $g_{\mathcal{S}_t}$, denoted by $\tau(g_{\mathcal{S}_t} \mid \mathcal{D}^I_{\mathcal{S}_t})$, is updated. Once the maximum number of trials is reached, a tuple $(\mathcal{S}^*, \pi^*_{\mathcal{S}^*})$ giving the smallest estimated target effect in $\mathcal{D}^I = \{\mathcal{D}^I_{\mathcal{S}}\}_{\mathcal{S}\in\Sigma}$ is returned.

**Algorithm 1** fCBO

**Inputs:** $\mathcal{G}$, $\boldsymbol{I}$, $Y$, $\mathcal{D}^I = \{\mathcal{D}^I_{\mathcal{S}}\}_{\mathcal{S}\in\Sigma}$, $T$, $S$
$\mathbb{M}_\Sigma \leftarrow$ `NRMPSReduce`$(\mathcal{G}, \boldsymbol{I}, Y)$
Initialise GPs $g_{\mathcal{S}}(\pi_{\mathcal{S}})\ \forall \mathcal{S} \in \mathbb{M}_\Sigma$ with $\mathcal{D}^I_{\mathcal{S}}$
**for** $t = 1, \dots, T$ **do**
  1. Select MPS $\mathcal{S}_t$ and DMP $\pi^t_{\mathcal{S}_t}$ via the fEI
  2. Obtain samples $\{y^{(s)}\}_{s=1}^S$ from $p_{\pi^t_{\mathcal{S}_t}}(Y)$
  3. Compute sample mean estimate $\hat{\mu}^Y_{\pi^t_{\mathcal{S}_t}}$ using $\{y^{(s)}\}_{s=1}^S$
  4. $\mathcal{D}^I_{\mathcal{S}_t} \leftarrow \mathcal{D}^I_{\mathcal{S}_t} \cup (\pi^t_{\mathcal{S}_t}, \hat{\mu}^Y_{\pi^t_{\mathcal{S}_t}})$
  5. Update $\tau(g_{\mathcal{S}_t} \mid \mathcal{D}^I_{\mathcal{S}_t})$
**end**
**Output:** $(\mathcal{S}^*, \pi^*_{\mathcal{S}^*})$ with min $\hat{\mu}^Y_{\pi^*_{\mathcal{S}^*}}$ over $\mathcal{D}^I$

Notice that Algorithm 1 only requires realizations from $p_{\pi^t_{\mathcal{S}_t}}(Y)$ (and could also operate if given directly $\hat{\mu}^Y_{\pi^t_{\mathcal{S}_t}}$ instead). This is a considerable practical advantage compared to context-specific reward approaches such as the one in Arsenyan et al. [2023] that, similarly to non-causal contextual BO methods [Krause and Ong, 2011], require values of the contexts and of the target variable resulting from the intervention at that specific context values. Similarly to recent approaches in contextual BO [Feng et al., 2020], fCBO can directly operate on aggregate rewards.

### 4.1 SEARCH SPACE REDUCTION

The cardinality of $\Sigma$ grows exponentially with the cardinality of $\boldsymbol{I}$ and the number of possible context sets $\boldsymbol{C}_X$ for each $X$. Therefore, solving the fCGO problem by exploring the entire set could be prohibitively expensive. Even if $\Sigma$ has small cardinality, reducing the search space would simplify the problem by reducing the number of target effects to be modelled. We propose to use the results in Lee and Bareinboim [2020] to reduce the search to the subset of *non-redundant* MPSs included in $\Sigma$, denoted by $\mathbb{M}_\Sigma$, which is guaranteed to contain a solution to the fCGO problem. For completeness and clarity, in this section we describe these results in the setting of DMPs.

Let $\mathcal{S}' \subseteq \mathcal{S}$ indicate that $\boldsymbol{C}'_X \subseteq \boldsymbol{C}_X, \forall \langle X, \boldsymbol{C}'_X\rangle \in \mathcal{S}'$ with $\boldsymbol{X}_{\mathcal{S}'} \subseteq \boldsymbol{X}_{\mathcal{S}}$. Furthermore, let $\pi_{\mathcal{S}'} \subseteq \pi_{\mathcal{S}}$ indicate that $\pi_{X|\boldsymbol{C}'_X}(\boldsymbol{c}'_X) = \int \pi_{X|\boldsymbol{C}_X}(\boldsymbol{c}'_X \cup \boldsymbol{c}''_X) p_{\pi_{\mathcal{S}}}(\boldsymbol{c}''_X \mid \boldsymbol{c}'_X) d\boldsymbol{c}''_X$, $\forall X \in \boldsymbol{X}_{\mathcal{S}'}, \boldsymbol{c}''_X \in \mathcal{R}_{\boldsymbol{C}_X\backslash\boldsymbol{C}'_X}$. Finally, let $\perp\!\!\!\perp_{\mathcal{G}}$ denote *d-separation* in $\mathcal{G}$, and $\mathcal{G}_{\mathcal{S}}\backslash X$ the modification of $\mathcal{G}$ obtained by removing node $X$ and its incoming and outgoing edges.

**Definition 4.1** (Non-redundant MPS)**.** An MPS $\mathcal{S}$ is said to be non-redundant if there exists an SCM compatible with $\mathcal{G}$

and $\pi_{\mathcal{S}} \in \Pi_{\mathcal{S}}$ such that $\mu^Y_{\pi'_{\mathcal{S}}} \neq \mu^Y_{\pi_{\mathcal{S}}}$ $\forall \mathcal{S}' \subset \mathcal{S}$ and $\pi'_{\mathcal{S}} \subset \pi_{\mathcal{S}}$.

The following proposition gives a graphical criterion for identifying $\mathbb{M}_\Sigma$.

**Proposition 4.2** (Characterization of non-redundant MPS)**.** *An* MPS $\mathcal{S}$ *is non-redundant if and only if (1)* $\boldsymbol{X}_{\mathcal{S}} \subseteq an_{\mathcal{G}_{\mathcal{S}}}(Y)$ *and (2)* $Y \not\perp\!\!\!\perp_{\mathcal{G}_{\mathcal{S}}\setminus X} C \,|\, \boldsymbol{C}_X \setminus C$ *for every* $X \in \boldsymbol{X}_{\mathcal{S}}$ *and* $C \in \boldsymbol{C}_X$.

## 4.2 GAUSSIAN PROCESS SURROGATE MODELS

We model the unknown target effect $\mu^Y_{\pi_{\mathcal{S}}}$ for each $\mathcal{S}$ using a GP $g_{\mathcal{S}}(\pi_{\mathcal{S}})$. Differently from existing works on Bayesian functional optimization that focus on univariate functional inputs, $\pi_{\mathcal{S}}$ can include scalar values as well as functions potentially defined on different input spaces.

For instance, for MPS $\mathcal{S} = \{\langle X, \{C_1, C_2\}\rangle, \langle Z, \{C_2\}\rangle\}$ with $\mathcal{G}_{\mathcal{S}}$ given on the top right, $\pi_{X|\{C_1,C_2\}}$ is defined over $\mathcal{R}_{C_1} \times \mathcal{R}_{C_2}$, while $\pi_{Z|C_2}$ over $\mathcal{R}_{C_2}$. Alternatively, for MPS $\mathcal{S} = \{\langle X, \{C_1\}\rangle, \langle Z, \{C_2\}\rangle\}$ with $\mathcal{G}_{\mathcal{S}}$ given on the bottom right, $\pi_{X|C_1}$ is defined over $\mathcal{R}_{C_1}$, while $\pi_{Z|C_2}$ over $\mathcal{R}_{C_2}$. We address this complexity by introducing a kernel function for $g_{\mathcal{S}}(\pi_{\mathcal{S}})$ that allows to compute distances among the mixed inputs while handling the different input dimensionality.

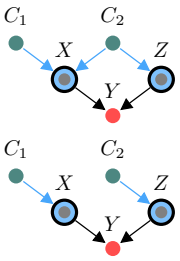


More specifically, $g_{\mathcal{S}}(\pi) \sim \mathcal{GP}(m_{\mathcal{S}}(\pi), K^\theta_{\mathcal{S}}(\pi, \pi'))$, where $\pi, \pi' \in \Pi_{\mathcal{S}}$ (we omit the subscript $\mathcal{S}$ to simplify the notation[3]), and $m_{\mathcal{S}}$ and $K^\theta_{\mathcal{S}}$ denote the prior mean and covariance functional with hyperparameters $\theta$. Notice that $\Pi_{\mathcal{S}} := \mathbb{R}^{|\mathcal{S}_{\text{hard}}|} \times \mathcal{B}(\boldsymbol{C}_{\mathcal{S}})$ where[4] $\mathbb{R}^{|\mathcal{S}_{\text{hard}}|}$ is the space of scalar values for $\boldsymbol{X}_{\mathcal{S}_{\text{hard}}}$ while $\mathcal{B}(\boldsymbol{C}_{\mathcal{S}})$ is the space of bounded vector-valued functions on $\boldsymbol{C}_{\mathcal{S}} = \bigcup_{X \in \boldsymbol{X}_{\mathcal{S}}} \boldsymbol{C}_X$. Given an interventional dataset $\mathcal{D}^I_{\mathcal{S}}$ for $\mathcal{S}$, for which we assume a Gaussian likelihood, the posterior distribution $\tau(g_{\mathcal{S}} \,|\, \mathcal{D}^I_{\mathcal{S}})$ can be computed by standard GP updates [Williams and Rasmussen, 2006]. We initialize $m_{\mathcal{S}}$ to a zero mean functional and extend the RBF kernel to consider mixed inputs as detailed below.

**Kernels for Functional GP.** We define $K^\theta_{\mathcal{S}}$ as the RBF kernel $K^\theta_{\mathcal{S}}(\pi, \pi') = \sigma_f^2 \exp(-||\pi - \pi'||^2/2\ell^2)$, where $\theta = (\sigma_f^2, \ell)$ and where $||\pi - \pi'||$ represents a distance between mixed inputs to the GP[5]. Let $\pi_{\text{hard}}$ and $\pi_{\text{func}}$ denote the vectors whose elements are the scalar values and the functions included in $\pi$, respectively. We define $||\pi - \pi'||^2$ as $||\pi - \pi'||^2 = ||\pi_{\text{hard}} - \pi'_{\text{hard}}||^2 + ||\pi_{\text{func}} - \pi'_{\text{func}}||^2_{\mathcal{H}_{\kappa_{\mathcal{S}}}}$, with $||\pi_{\text{hard}} - \pi'_{\text{hard}}||^2$ indicating the square of the Euclidean distance in $\mathbb{R}^{|\mathcal{S}_{\text{hard}}|}$, and $||\pi_{\text{func}} - \pi'_{\text{func}}||^2_{\mathcal{H}_{\kappa_{\mathcal{S}}}}$ the distance between functions in the vector-valued reproducing kernel Hilbert space (RKHS, Aronszajn [1950]) $\mathcal{B}(\boldsymbol{C}_{\mathcal{S}}) = \mathcal{H}_{\kappa_{\mathcal{S}}}$ described below.

Specifically, $\mathcal{H}_{\kappa_{\mathcal{S}}}$ is an RKHS with vector-valued reproducing kernel $\kappa^\xi_{\mathcal{S}} : \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \times \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \to \mathbb{R}^{|\mathcal{S}_{\text{func}}| \times |\mathcal{S}_{\text{func}}|}$ where $\xi$ denotes the hyper-parameters and $\mathcal{S}_{\text{func}} = \{\langle X, \boldsymbol{C}_X\rangle \in \mathcal{S} \colon \boldsymbol{C}_X \neq \emptyset\}$. We refer to $\kappa^\xi_{\mathcal{S}}$ as the *functional intervention kernel* to distinguish it from $K^\theta_{\mathcal{S}}$. We thus have $||\pi_{\text{func}} - \pi'_{\text{func}}||^2_{\mathcal{H}_{\kappa_{\mathcal{S}}}} = \langle \pi_{\text{func}} - \pi'_{\text{func}}, \pi_{\text{func}} - \pi'_{\text{func}} \rangle_{\mathcal{H}_{\kappa_{\mathcal{S}}}}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ denotes the inner product in the space $\mathcal{H}$. Evaluating this quantity requires computing $\kappa^\xi_{\mathcal{S}}$ at different input values for the variables in $\boldsymbol{C}_{\mathcal{S}}$, say $\mathbf{c}_{\mathcal{S}}$ and $\mathbf{c}'_{\mathcal{S}}$, for $\pi_{\text{func}}$ and $\pi'_{\text{func}}$ respectively.

We write the vector of functions $\pi_{\text{func}}$ included in the RKHS $\mathcal{H}_{\kappa_{\mathcal{S}}}$ as $\pi_{\text{func}}(\cdot) = \sum_{i=1}^{N_\alpha} \kappa^\xi_{\mathcal{S}}(\mathbf{c}^i_{\mathcal{S}}, \cdot)\boldsymbol{\alpha}_i$ with $\boldsymbol{\alpha}_i \in \mathbb{R}^{|\mathcal{S}_{\text{func}}|}$ and $\mathbf{c}^i_{\mathcal{S}} \in \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}}$ and let $\pi'_{\text{func}}(\cdot) = \sum_{i=1}^{N_\beta} \kappa^\xi_{\mathcal{S}}(\mathbf{c}^i_{\mathcal{S}}, \cdot)\boldsymbol{\beta}_i$ with $\boldsymbol{\beta}_i \in \mathbb{R}^{|\mathcal{S}_{\text{func}}|}$. This implies that the inner product $\langle \pi_{\text{func}} - \pi'_{\text{func}}, \pi_{\text{func}} - \pi'_{\text{func}} \rangle_{\mathcal{H}_{\kappa_{\mathcal{S}}}}$ can be written as

$$\sum_{i=1}^{N_\alpha}\sum_{j=1}^{N_\alpha} \boldsymbol{\alpha}_i^\top \kappa^\xi_{\mathcal{S}}(\mathbf{c}^i_{\mathcal{S}}, \mathbf{c}^j_{\mathcal{S}})\boldsymbol{\alpha}_j + \sum_{i=1}^{N_\beta}\sum_{j=1}^{N_\beta} \boldsymbol{\beta}_i^\top \kappa^\xi_{\mathcal{S}}(\mathbf{c}^i_{\mathcal{S}}, \mathbf{c}^j_{\mathcal{S}})\boldsymbol{\beta}_j - 2\sum_{i=1}^{N_\alpha}\sum_{j=1}^{N_\beta} \boldsymbol{\alpha}_i^\top \kappa^\xi_{\mathcal{S}}(\mathbf{c}^i_{\mathcal{S}}, \mathbf{c}^j_{\mathcal{S}})\boldsymbol{\beta}_j.$$

To construct $\kappa^\xi_{\mathcal{S}}$, we propose to augment the input space by including a task index for each function $\pi_{X|\boldsymbol{C}_X}$ in $\mathcal{S}$, i.e. we redefine $\kappa^\xi_{\mathcal{S}} \colon (\mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \times \mathcal{T}) \times (\mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \times \mathcal{T}) \to \mathbb{R}^{|\mathcal{S}_{\text{func}}| \times |\mathcal{S}_{\text{func}}|}$ where $\mathcal{T}$ is the space of integer values from 1 to $|\mathcal{S}_{\text{func}}|$. For every realization of the context variables and the task index, say $(\mathbf{c}_{\mathcal{S}}, t)^i$, we can then evaluate $\kappa^\xi_{\mathcal{S}}((\mathbf{c}_{\mathcal{S}}, t)^i, (\mathbf{c}_{\mathcal{S}}, t)^j)$. We assume the covariance between functions defined on different input spaces, i.e. for which $t^i \neq t^j$, to be 0[6]. Instead, we let the covariance structure across function values associated with different inputs for $t^i = t^j$ be determined by a task-specific kernel, which we denote by $k^{t^i}$. Denote by $\mathbf{c}^i_{\mathcal{S}}[t^i]$ the subset of values included in $\mathbf{c}^i_{\mathcal{S}}$ for the contexts of the $t^i$ task and by $\xi[t^i]$ the subset of hyper-parameters for $t^i$ included in $\xi$. We have that $\kappa^\xi_{\mathcal{S}}((\mathbf{c}_{\mathcal{S}}, t)^i, (\mathbf{c}_{\mathcal{S}}, t)^j)$ is equal to $k^{t^i}(\mathbf{c}^i_{\mathcal{S}}[t^i], \mathbf{c}^j_{\mathcal{S}}[t^j])$ with hyper-parameters $\xi[t^i]$ if $t^i = t^j$ and to 0 otherwise. The kernel $k^{t^i}$ might differ across tasks both in terms of functional form and hyper-parameter values. This allows to impose different characteristics in terms of e.g. smoothness for each function $\pi_{X|\boldsymbol{C}_X}$ included in $\pi$.

[3] In this section, a DMP $\pi_{\mathcal{S}}$ indicates a vector, rather than a set, of interventions.

[4] $|\boldsymbol{X}|$ indicates the cardinality of the set $\boldsymbol{X}$.

[5] While we discuss the RBF kernel, this procedure can be used to compute any stationary kernel involving the distance between functional inputs similarly to Vien et al. [2018].

[6] Alternative kernel constructions where this assumption is relaxed are discussed in Section 2 of the supplementary material.

### 4.3 ACQUISITION FUNCTIONAL

We sequentially select interventions by numerically[7] maximizing the expected improvement (EI) per unit of cost $\texttt{Co}_{\mathcal{S}}(\cdot)$ *across* the MPSs in $\mathbb{M}_{\Sigma}$. Given an interventional dataset $\mathcal{D}^{\mathcal{I}}_{\mathcal{S}}$, for each $\mathcal{S} \in \mathbb{M}_{\Sigma}$ the functional EI (fEI) is given by:

$$\text{fEI}_{\mathcal{S}}(\pi) = \sigma^2_{\mathcal{S}}(\pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}})[\gamma(\pi)\mathbf{\Phi}(\gamma(\pi)) + \phi(\gamma(\pi))]/\texttt{Co}_{\mathcal{S}}(\pi),$$

where $\sigma^2_{\mathcal{S}}(\pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}}) = K^{\theta}_{\mathcal{S}}(\pi, \pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}})$, $\mathbf{\Phi}(\cdot)$ and $\phi(\cdot)$ are the CDF and PDF of a standard Gaussian random variable respectively, and $\gamma(\pi) = \frac{m_{\mathcal{S}}(\pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}}) - g^*}{K^{\theta}_{\mathcal{S}}(\pi, \pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}})}$ with $g^*$ denoting the optimum observed for $g_{\mathcal{S}}$ across MPSs in $\mathbb{M}_{\Sigma}$. $m_{\mathcal{S}}(\pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}})$ and $K^{\theta}_{\mathcal{S}}(\pi, \pi \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}})$ denote the posterior parameters of $\tau(g_{\mathcal{S}} \mid \mathcal{D}^{\mathcal{I}}_{\mathcal{S}})$. At every trial $t$ of the optimization, the MPS and the DMP are chosen by numerically solving $\mathcal{S}_t, \pi^t_{\mathcal{S}_t} = \arg\max_{\mathcal{S} \in \mathbb{M}_{\Sigma}, \pi_{\mathcal{S}} \in \Pi_{\mathcal{S}}} \text{fEI}_{\mathcal{S}}(\pi_{\mathcal{S}})$.

$\texttt{Co}_{\mathcal{S}}(\pi)$ denotes the cost associated to $\pi$. We consider two types of costs: (i) $\texttt{Co}_{\mathcal{S}}(\pi) = |\mathcal{S}|$; and (ii) $\texttt{Co}_{\mathcal{S}}(\pi) = \sum_{X \in \boldsymbol{X}_{\mathcal{S}}} \int_{\mathcal{R}_{\boldsymbol{C}_X}} \pi_{X|\boldsymbol{C}_X}(\boldsymbol{c}_X) d\boldsymbol{c}_X$, (i.e. the sum of the area under $\pi_{X|\boldsymbol{C}_X}$ over all $X \in \boldsymbol{X}_{\mathcal{S}}$), which can be seen as a measure of the units of intervention given to a population whose context values $\boldsymbol{c}_X$ are uniformly distributed in $\mathcal{R}_{\boldsymbol{C}_X}$. Notice that the second cost requires knowledge of $\mathcal{R}_{\boldsymbol{C}_X}$ at initialization. We use the first cost in the CHAIN experiments of Section 6.1, and the second cost in the HEALTH experiments of Section 6.2.

## 5 RELATED WORK

There exist two other CBO-type methods in the literature that can achieve contextual interventions, namely CoCa-BO [Arsenyan et al., 2023] and MCBO [Sussex et al., 2023]. CoCa-BO performs different hard interventions in separate sub-groups defined by some contexts *after* observing context values. Interventional data samples, formed by context values, intervention values, and target effect, are used to fit a GP model over the potentially high-dimensional context-intervened variables space. Therefore, CoCa-BO can only be used in settings in which the investigator observes the values of the context variables, say $\boldsymbol{C} = \boldsymbol{c}$, selects an intervention and observes the resulting target effect across units with $\boldsymbol{C} = \boldsymbol{c}$, rather than an aggregate target effect across all possible context values in a population. This is not feasible in many applied problems (e.g. in A/B testing platforms, in which outcomes are often measured as an aggregate across a large population that spans an entire distribution of contexts), and might lead to sup-optimal policies for unseen or rarely observed context values. In addition, this method defines the GP surrogate model for each MPS $\mathcal{S}$ on $\boldsymbol{C}_{\mathcal{S}}$ thus reducing the flexibility of the learned policy by not encoding the existence of different $\boldsymbol{C}_X$ for each $X$ in $\boldsymbol{X}_{\mathcal{S}}$. MCBO considers systems described by SCMs in which $X \in \boldsymbol{I}$ is of the form $X = f_X(\text{pa}_{\mathcal{G}}(X), \boldsymbol{A}_X) + U_X$, where $\boldsymbol{A}_X$ is a set of action variables that parametrize $f_X$ whose values can be set by the investigator to induce a change in the parametrization. Therefore, a contextual intervention in MCBO *modifies* a node's original functional assignment rather than *replacing* it as in fCBO. This might lead to more limited interventions and does not allow change of contexts. In addition, this method can achieve contextual interventions only in settings in which the system's SCM contains action variables. When this is not the case, MCBO can only implement hard interventions (see the HEALTH experiment of Section 6.2). Finally, unlike fCBO, MCBO does not reduce the search space and cannot handle unobserved confounders.

Extensions of BO [Shahriari et al., 2015] to solve functional global optimization (FGO) problems have been studied by searching over the space of Bernstein polynomials [Vellanki et al., 2019], by constructing a sequence of low-dimensional search spaces [Shilton et al., 2020], or by representing the functional inputs as elements in an RKHS (BFO) [Vien et al., 2018]. This work takes an approach similar to BFO, but considers a varied search space and its causal reduction. More importantly, thanks to a simple kernel construction, it enables functional BO, which has generally focused on univariate functional inputs, to deal with settings where the inputs are multi-task functions.

## 6 EXPERIMENTS

We compare[8] fCBO with CBO, MCBO, BO, and BFO on the synthetic graph in Section 6.1 (CHAIN), and on the healthcare graph in Fig. 1(a) (HEALTH). The experiments aim at highlighting three main advantages of using fCBO to find optimal interventions. The first advantage is the ability to achieve smaller target effects compared to methods that use only hard interventions. We assess this by looking at the convergence to the optimum. The second advantage is the ability to perform well w.r.t. conditional target effects. We demonstrate this in the CHAIN experiments, by computing the *performance gain* for DMP $\pi_{\mathcal{S}}$ on sub-group $\boldsymbol{C} = \boldsymbol{c}$, which is defined as $\text{PGain}(\pi_{\mathcal{S}}, \boldsymbol{C} = \boldsymbol{c}) = \hat{\mu}^{Y}_{\boldsymbol{C}=\boldsymbol{c}} - \hat{\mu}^{Y}_{\pi_{\mathcal{S}}, \boldsymbol{C}=\boldsymbol{c}}$, where $\hat{\mu}^{Y}_{\boldsymbol{C}=\boldsymbol{c}}$ denotes an estimate of the conditional expectation of $Y$ given $\boldsymbol{C} = \boldsymbol{c}$ w.r.t. the observational distribution and $\hat{\mu}^{Y}_{\pi_{\mathcal{S}}, \boldsymbol{C}=\boldsymbol{c}}$ an estimate of the conditional target effect. The third advantage is the ability to craft flexible and more targeted DMPs that can incur similar or lower cost, while still ensuring a smaller target effect than policies made of only hard interventions. We exemplify this in the HEALTH experiments where we assume a cost function given by

[7] Alternatively, the functional gradient w.r.t. functions in a RKHS could be derived analytically (see Vien et al. [2018]).

[8] We cannot compare to CoCa-BO as: (i) in our settings the values of the contexts are not observed before intervening, and only an aggregate target effect across contexts is observed post intervention; (ii) this method does not allow considering MPSs that do not share the same contexts.

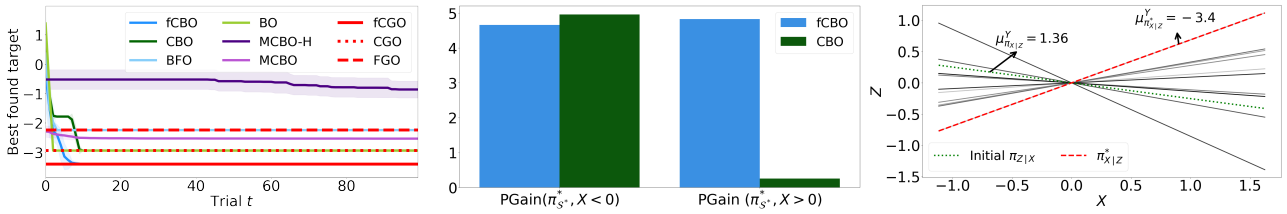


Figure 2: CHAIN experiments. *Left*: Average convergence of fCBO to the fCGO optimum (solid red line); of CBO, BO, and MCBO-H to the CGO optimum (dotted red line); of BFO to the FGO optimum (dashed red line); and of MCBO to the CGO$^*$ optimum (across 20 initializations of $\mathcal{D}^I$ for fCBO, CBO, BO, BFO, and across 20 seeds for MCBO-H and MCBO – shaded areas give $\pm$ standard deviation). *Middle*: Average performance gains PGain$(\pi^*_{\mathcal{S}^\star}, X < 0)$ and PGain$(\pi^*_{\mathcal{S}^\star}, X > 0)$ obtained by the optimal DMP $\pi^*_{\mathcal{S}^\star}$. *Right*: Initial $\pi_{Z|X}$ included in $\mathcal{D}^I_Z$, $\pi^*_{Z|X}$ found by fCBO and associated target effect values.

$\texttt{Co}_{\mathcal{S}}(\pi) = \sum_{X \in \boldsymbol{X}_{\mathcal{S}}} \int_{\mathcal{R}_{\boldsymbol{C}_X}} \pi_{X|\boldsymbol{C}_X}(\boldsymbol{c}_X) d\boldsymbol{c}_X$.

**Search Space and Optimization Problem**

| fCBO | CBO | MCBO-H | MCBO | BO | BFO |
|---|---|---|---|---|---|
| $\Sigma$ | $\Sigma_{\text{hard}}$ | $\Sigma_{\text{hard}}$ | $\mathcal{P}_{\boldsymbol{A}}$ | $\mathcal{S}_{\boldsymbol{I},\boldsymbol{C}_X=\emptyset}$ | $\mathcal{S}_{\subseteq \boldsymbol{I},\boldsymbol{C}_X \neq \emptyset}$ |
| fCGO | CGO | CGO | CGO$^*$ | GO | FGO |

The different search spaces of fCBO, CBO, MCBO, BO, and BFO are summarized in the table above. An intervention in BO and BFO is performed on all variables or on a subset of variables in $\boldsymbol{I}$ simultaneously: BO considers only hard interventions, thus its search space contains only MPS $\mathcal{S}_{\boldsymbol{I},\boldsymbol{C}_X=\emptyset} = \{\langle X, \boldsymbol{C}_X \rangle \colon X \in \boldsymbol{I}, \boldsymbol{C}_X = \emptyset\})$; while BFO considers functional interventions with a fixed $\boldsymbol{C}_X \neq \emptyset$ over trials, i.e. its search space contains only one MPS formed by tuples $\langle X, \boldsymbol{C}_X \rangle$ with $\boldsymbol{X}_{\mathcal{S}} \subseteq \boldsymbol{I}$, denoted by $\mathcal{S}_{\subseteq \boldsymbol{I},\boldsymbol{C}_X \neq \emptyset}$. CBO and MCBO with hard interventions, denoted by MCBO-H, consider the space of MPSs containing only hard interventions $\Sigma_{\text{hard}}$. Finally, MCBO performs interventions via actions variables $\boldsymbol{A} = \{\boldsymbol{A}_X\}_{X \in \boldsymbol{I}}$ thus exploring the power set $\mathcal{P}_{\boldsymbol{A}}$ (with the convention that no intervention on $X$ corresponds to removing $\boldsymbol{A}_X$ from the SCM). While fCBO aims at solving the fCGO problem, CBO and MCBO-H target the CGO problem, and BFO the FGO problem. Finally, BO solves a global optimization problem (GO), while MCBO a CGO problem in the action variable space, denoted by CGO$^*$. In all experiments, we consider settings where the fCGO, CGO, and FGO problems have unique solutions, and the GO optimum coincides with the CGO optimum.

**fCBO, CBO, BO, and BFO.** While fCBO does not impose restrictions in terms of context variables used for functional interventions beyond acyclicity of $\mathcal{G}_{\mathcal{S}}$, for ease of demonstration and for computational reasons, in the experiments we only consider keeping the original parents as contexts. In other words, we set $\boldsymbol{C}_X = \text{pa}_{\mathcal{G}}(X)$ for each functional intervention. We make the same choice for BFO. To demonstrate performance on different choices for $\Pi_{\mathcal{S}}$, we consider linear and RBF functional intervention kernels $\kappa^{\xi}_{\mathcal{S}}$ in the CHAIN and HEALTH experiments, respectively. We use the same functional intervention representation for BFO. For each $\mathcal{S} \in \mathbb{M}_{\Sigma}$ we numerically optimize the acquisition functions on a grid whose size is set to $\texttt{GridSize}^{|\mathcal{S}_{\text{hard}}|+1}$ where $\texttt{GridSize}$ is a hyper-parameter. We initialize $\mathcal{D}^I$ by randomly generating a single DMP and associated target effect for each $\mathcal{S} \in \Sigma$. We provide average results across the 20 different initializations.

**MCBO.** In the CHAIN experiments, we consider both MCBO restricted to hard interventions (MCBO-H) and MCBO with contextual interventions (by augmenting the SCM with an action variable for each variable in $\boldsymbol{I}$). In the HEALTH experiments, the SCM is given and does not contain action variables. Therefore, we follow Sussex et al. [2023] and consider only hard interventions on Aspirin, Statin, and CI. We run the algorithm[9] by setting the random seed controlling both the initial interventional data and the optimization of the acquisition function to values $1, \dots, 20$. We report results across the 20 different seeds. Cross-validation with values $0.05$, $0.5$, and $5$ on the hyper-parameter $\beta$ for the UCB acquisition function, as done in Sussex et al. [2023], does not give major differences in the performance (we report the results for $\beta = 5$).

## 6.1 CHAIN EXPERIMENTS

We first experiment on the chain graph with associated SCM given on the right (see Section 3 of the supplementary material for details). Fig. 2(left) shows how considering mixes of hard and functional interventions allows fCBO to reach the smallest target effect.

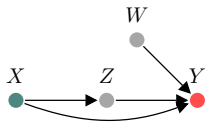


$X = U_X,\ W = U_W$
$Z = -0.5X + U_Z$
$Y = -W - 3ZX + U_Y$

Fig. 2(middle) shows how fCBO and CBO differ in terms of conditional target effects defined for $X < 0$ and $X > 0$.

[9] We used the code companion to Sussex et al. [2023] available at https://github.com/ssethz/mcbo.

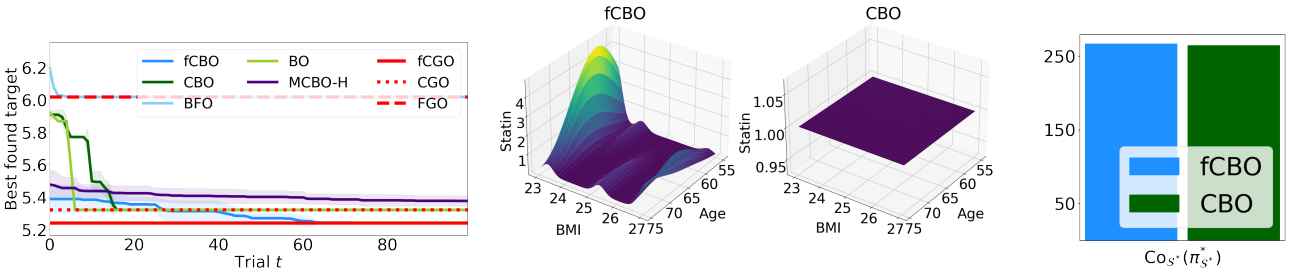


Figure 3: HEALTH experiments with `GridSize` $= 5$. *Left*: Average convergence of fCBO to the fCGO optimum (solid red line); of CBO, BO, and MCBO-H to the CGO optimum (dotted red line); and of BFO to the FGO optimum (dashed red line) (across 20 initializations of $\mathcal{D}^I$ for fCBO, CBO, BO, BFO, and across 20 seeds for MCBO-H – shaded areas give $\pm$ standard deviation). *Middle*: $\pi^*_{\text{Statin}|\text{Age,BMI}}$ found by fCBO (left) and $\pi^*_{\text{Statin}|\emptyset}$ found by CBO (right) across levels of Age and BMI. *Right*: Cost associated to the optimal MPS and associated optimal DMP found by fCBO and CBO.

Due to the existence of the interaction term $-3ZX$, minimizing $Y$ would require setting $Z$ to a negative value when $X < 0$ and to a positive value when $X > 0$. However, this cannot be achieved via hard interventions that set $Z$ to a fixed value irrespective of $X$ as in CBO. As a consequence CBO, which selects MPS $\mathcal{S}^* = \{\langle Z, \emptyset\rangle, \langle W, \emptyset\rangle\}$ and DMP $\pi^*_{\mathcal{S}^*} = \{-1, 1\}$, achieves a very low performance gain for $X > 0$, $\text{PGain}(\pi^*_{\mathcal{S}^*}, X > 0)$. Instead, fCBO selects MPS $\mathcal{S}^* = \{\langle Z, X\rangle, \langle W, \emptyset\rangle\}$ and DMP $\pi^*_{\mathcal{S}^*} = \{\pi^*_{Z|X}, 1\}$, where the linear function $\pi^*_{Z|X}$ (shown as a dashed red line in Fig. 2(right)) has a slope that gives an optimal $Z$ value for both sub-groups thus leading to an evenly distributed performance gain.

### 6.2 HEALTH EXPERIMENTS

For the HEALTH experiments, we use the SCM by Ferro et al. [2015] (see Section 4 of the supplementary material for details). Fig. 3 shows the results obtained with `GridSize` $= 5$. In these experiments, fCBO achieves the smallest target effect by selecting MPS $\mathcal{S}^* = \{\langle \text{Aspirin}, \emptyset\rangle, \langle \text{Statin}, (\text{Age}, \text{BMI})\rangle, \langle \text{CI}, \emptyset\rangle\}$ and DMP $\pi^*_{\mathcal{S}^*} = \{0.1, \pi^*_{\text{Statin}|\text{Age,BMI}}, 1\}$. BO and CBO select MPS $\mathcal{S}^* = \{\langle \text{Aspirin}, \emptyset\rangle, \langle \text{Statin}, \emptyset\rangle, \langle \text{CI}, \emptyset\rangle\}$, and DMP $\pi^*_{\mathcal{S}^*} = \{0.1, 1, 1\}$. MCBO-H does not reach convergence.

Fig. 3(middle) displays $\pi^*_{\text{Statin}|\text{Age,BMI}}$ selected by fCBO (left) and $\pi^*_{\text{Statin}|\emptyset}(\emptyset) = 1$ selected by CBO as a constant function over Age and BMI (right). These two plots show that, while methods that consider only hard interventions are forced to assign intervention values uniformly across the context space, methods that also allow functional interventions can concentrate on specific sub-groups, in this case characterized by lower values of Age and BMI. Being able to differentiate among interventions assigned to different sub-groups has important implications in terms of cost $\text{Co}_{\mathcal{S}^*}(\pi^*_{\mathcal{S}^*})$. Fig. 3(right) shows that fCBO incurs almost the same cost as CBO. This result demonstrates another key property of functional interventions: taking the context values into account allows the investigator to assign interventions to units in the population characterized by context values that lead to smaller target effects. Similar results are observed with `GridSize` $= 8$ (Fig. 4). fCBO achieves the smallest target effect (Fig. 4, left), and incurs a lower cost compared to CBO (Fig. 4, right). In this setting fCBO converges to $\mathcal{S}^* = \{\langle \text{Aspirin}, \emptyset\rangle, \langle \text{Statin}, (\text{Age}, \text{BMI})\rangle\}$ with $\pi^*_{\mathcal{S}^*} = \{0.1, \pi^*_{\text{Statin}|\text{Age,BMI}}\}$. Due to the more complex $\pi^*_{\text{Statin}|\text{AgeBMI}}$ (Fig. 4(middle, left)), which allocates the highest Statin dosages to mid-range value of Age and BMI, the investigator can avoid intervening on CI thus lowering the overall cost of the intervention while still achieving an overall smaller target effect.

## 7 CONCLUSION

We proposed the fCBO method for finding policies made of hard and functional interventions that optimize a target effect. We introduced graphical criteria that establish when functional interventions could be necessary to achieve optimal target effects and when hard interventions are sufficient. Furthermore, we showed that optimizing a target effect by considering functional interventions allows the investigator to identify policies that are also optimal w.r.t. conditional target effects. We demonstrated the benefit of the proposed approach on a synthetic and on a real-world causal graph. Future work will explore the use of gradient-based optimization methods for the acquisition functional, as well as the development of more flexible kernel construction for the GP functionals (see Section 2 of the supplementary material). These extensions would enable the identification of more flexible functional interventions while speeding up the convergence of the algorithm.

**Acknowledgements**

The authors would like to thank Michalis Titsias, Alan Malek, and Eleni Sgouritsa for valuable discussions.

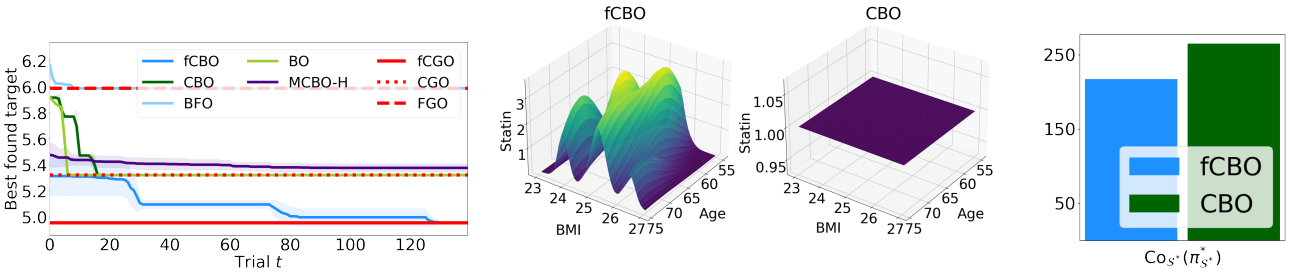


Figure 4: HEALTH experiments with `GridSize` $= 8$. *Left*: Average convergence of fCBO to the fCGO optimum (solid red line); of CBO, BO, and MCBO-H to the CGO optimum (dotted red line); and of BFO to the FGO optimum (dashed red line) (across 20 different initializations of $\mathcal{D}^I$ for fCBO, CBO, BO, BFO, and across 20 seeds for MCBO-H – shaded areas give $\pm$ standard deviation). *Middle*: $\pi^*_{\text{Statin}|\text{Age,BMI}}$ found by fCBO (left) and $\pi^*_{\text{Statin}|\emptyset}$ found by CBO (right) across levels of Age and BMI. *Right*: Cost associated to the optimal MPS and associated optimal DMP found by fCBO and CBO.

## References


Virginia Aglietti, Xiaoyu Lu, Andrei Paleyes, and Javier González. Causal Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3155–3164, 2020.

Virginia Aglietti, Neil Dhir, Javier González, and Theodoros Damoulas. Dynamic causal Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 10549–10560, 2021.

Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3): 337–404, 1950.

Vahan Arsenyan, Antoine Grosnit, and Haitham Bou-Ammar. Contextual causal Bayesian optimisation. *arXiv preprint arXiv:2301.12412*, 2023.

Juan Correa and Elias Bareinboim. A calculus for stochastic interventions: Causal effect identification and surrogate experiments. In *AAAI Conference on Artificial Intelligence*, pages 10093–10100, 2020a.

Juan Correa and Elias Bareinboim. General transportability of soft interventions: Completeness results. In *Advances in Neural Information Processing Systems*, pages 10902–10912, 2020b.

Arnoud De Kroon, Joris Mooij, and Danielle Belgrave. Causal bandits without prior knowledge using separating sets. In *Conference on Causal Learning and Reasoning*, pages 407–427, 2022.

Qing Feng, Ben Letham, Hongzi Mao, and Eytan Bakshy. High-dimensional contextual policy search with unknown context rewards using Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 22032–22044, 2020.

Ana Ferro, Francisco Pina, Milton Severo, Pedro Dias, Francisco Botelho, and Nuno Lunet. Use of statins and serum levels of prostate specific antigen. *Acta Urológica Portuguesa*, 32(2):71–77, 2015.

Maxime Gasse, Damien Grasset, Guillaume Gaudron, and Pierre-Yves Oudeyer. Causal reinforcement learning using observational and interventional data. *arXiv preprint arXiv:2106.14421*, 2021.

Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.

Andreas Krause and Cheng Ong. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, 2011.

Finnian Lattimore, Tor Lattimore, and Mark D Reid. Causal bandits: Learning good interventions via causal inference. In *Advances in Neural Information Processing Systems*, 2016.

Sanghack Lee and Elias Bareinboim. Structural causal bandits: Where to intervene? In *Advances in Neural Information Processing Systems*, 2018.

Sanghack Lee and Elias Bareinboim. Structural causal bandits with non-manipulable variables. In *AAAI Conference on Artificial Intelligence*, pages 4164–4172, 2019.

Sanghack Lee and Elias Bareinboim. Characterizing optimal mixed policies: Where to intervene and what to observe. In *Advances in Neural Information Processing Systems*, pages 8565–8576, 2020.

Chaochao Lu, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Deconfounding reinforcement learning in observational settings. *arXiv preprint arXiv:1812.10576*, 2018.

Yangyi Lu, Amirhossein Meisami, Ambuj Tewari, and William Yan. Regret analysis of bandit problems with

causal background knowledge. In *Conference on Uncertainty in Artificial Intelligence*, pages 141–150, 2020.

Vineet Nair, Vishakha Patil, and Gaurav Sinha. Budgeted and non-budgeted causal bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2017–2025, 2021.

Judea Pearl. *Causality: Models, Reasoning and Inference*. Springer, 2000.

Prashanth Rawla. Epidemiology of prostate cancer. *World Journal of Oncology*, 10(2):63, 2019.

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. *Proceedings of the IEEE*, 104(1):148–175, 2015.

Alistair Shilton, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Sequential subspace search for functional Bayesian optimization incorporating experimenter intuition. *arXiv preprint arXiv:2009.03543*, 2020.

Scott Sussex, Anastasiia Makarova, and Andreas Krause. Model-based causal Bayesian optimization. In *International Conference on Learning Representations*, 2023.

Pratibha Vellanki, Santu Rana, Sunil Gupta, David Rubin de Celis Leal, Alessandra Sutti, Murray Height, and Svetha Venkatesh. Bayesian functional optimisation with shape prior. In *AAAI Conference on Artificial Intelligence*, pages 1617–1624, 2019.

Ngo Anh Vien, Heiko Zimmermann, and Marc Toussaint. Bayesian functional optimization. In *AAAI Conference on Artificial Intelligence*, pages 4171–4178, 2018.

Christopher KI Williams and Carl Edward Rasmussen. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

Junzhe Zhang. Designing optimal dynamic treatment regimes: A causal reinforcement learning approach. In *International Conference on Machine Learning*, pages 11012–11022, 2020.

Junzhe Zhang and Elias Bareinboim. Near-optimal reinforcement learning in dynamic treatment regimes. In *Advances in Neural Information Processing Systems*, 2019.

Junzhe Zhang and Elias Bareinboim. Online reinforcement learning for mixed policy scopes. In *Advances in Neural Information Processing Systems*, 2022.