Functional Causal Bayesian Optimization

Abstract

We propose the *functional causal Bayesian optimization method* (fCBO) for finding interventions that optimize a target variable in a known causal graph. fCBO extends CBO to perform, in addition to hard interventions, functional interventions which consist in setting a variable to be a deterministic function of a set of other variables in the graph. This is achieved by modelling the unknown objective 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 show that functional interventions can attain better target effects compared to hard interventions and ensure that the found optimal policy is also optimal for sub-groups. We demonstrate the benefits of the method in a synthetic setting 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, such as medicine, biology, and the social sciences. The framework of causal graphical models [Pearl, 2000, Koller and Friedman, 2009] offers a powerful tool for formalizing and dealing with interventions, and is therefore increasingly integrated into policy optimization methods like 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].

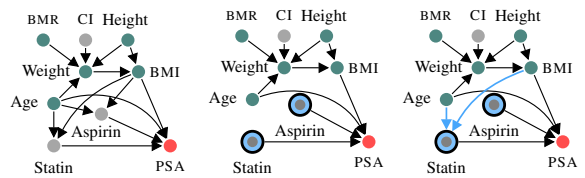

Figure 1: *Left:* Causal graph $\mathcal{G}$ representing causal relationships between prostate specific antigen (PSA) and other variables. Intervenable nodes are indicated in grey, non-intervenable nodes in dark green and target node in red. *Middle:* Modified $\mathcal{G}$ with hard interventions. *Right:* Modified $\mathcal{G}$ with hard and functional interventions.

A causal graphical model represents a system of interconnected variables with a casual graph and an associated distribution. Nodes represent random variables and directed edges causal relationships between them. Nodes with an edge pointing into node X are *direct causes* of X, i.e. they dictate the behavior of X through a causal mechanism specified by a conditional distribution. Within this formalism, an intervention on X corresponds to replacing its associated conditional distribution with a different one possibly governed by different direct causes.

Most current works on finding optimal interventions have focused on the *hard intervention* $\text{do}(X = x)$, which consists in setting X to a constant value x, thereby removing the influence of direct causes. However, in many practical scenarios an investigator may also be able to perform other types of interventions. Consider, for example, the causal graph in Fig. 1(left). An investigator wishing to find a policy for prescribing Statin and Aspirin dosages, as well as calories intake (CI), that minimize prostate specific antigen (PSA) might be able to perform, in addition to hard interventions on Statin and Aspirin that ignore any information about Age and BMI (Fig. 1(middle)), also interventions that still considers this contextual information (Fig. 1(right)).

Recent concurrent works in the causal Bayesian optimization (CBO) literature have introduced approaches that aim

Submitted to the 39th Conference on Uncertainty in Artificial Intelligence (UAI 2023). **To be used for reviewing only**.

at achieving contextual interventions, either by searching for different hard interventions in separate sub-groups of the population defined by some context variables [Arsenyan et al., 2023] or by operating on the system with "action" nodes that induce changes in the parametrization of the conditional distributions of the system's nodes [Sussex et al., 2023]. However, both approaches are constrained in the type of contextual interventions that can be achieved, as Arsenyan et al. [2023] only learns an implicit mapping between contexts and intervention values, while Sussex et al. [2023] only allows some limited modification of the parametrization of the original conditional distribution. The former is also limited in its choice of context and its ability to extrapolate to unseen on rarely explored areas of the context space, while the latter does not consider choice of context at all. In this work, we introduce an extension of the CBO family of methods [Aglietti et al., 2020] that considers a more flexible and general type of contextual intervention, consisting in replacing the original casual mechanism with a *deterministic function* of other observed variables in the graph. Such a *functional intervention* is implemented via new techniques for computing distances among functions of different variables. Together with hard interventions, functional interventions cover the full landscape of interventions required for solving the problem of optimizing a target variable in a system. Our contributions are the following:

- We frame the problem of finding optimal policies as a functional casual global optimization (fCGO) problem.
- We introduce graphical criteria that establish when functional interventions could be necessary to solve the fCGO problem as well as when hard interventions are sufficient.
- We introduce graphical criteria showing that, under mild conditions, a globally optimal policy containing functional interventions is simultaneously optimal across sub-groups, ensuring performance gains from an optimal policy are equally distributed across groups.
- We propose the functional causal Bayesian optimization (fCBO) method for solving the fCGO problem by modelling the unknown objective via a set of Gaussian process models whose inputs are defined in a reproducing kernel Hilbert space.
- We validate fCBO in a synthetic and in a real-world setting, presenting its advantage over the exclusive use of hard interventions in terms of target effects attained, effect on sub-groups, and intervention cost.

2 BACKGROUND AND SETTING

We consider a system of interconnected variables $\boldsymbol{V}$, which include a set of *intervenable* variables $\boldsymbol{I}$, a *target variable* Y, and a set on non-intervenable variables $\boldsymbol{V} \backslash \boldsymbol{I} \cup Y$.

We represent $\boldsymbol{V}$ as the nodes of a *causal graph* $\mathcal{G}$ with directed and bi-directed edges. A directed edge from node V to node W expresses that V is a *direct cause* of W, while a bidirected edge between node V and node W expresses the presence of an unobserved common cause of V and W (unobserved *confounder*). We assume that $\mathcal{G}$ is *acyclic*, namely that there are no *directed paths*[1] starting and ending at the same node. We indicate with $\text{pa}_{\mathcal{G}}(V)$ the *direct causes* of V in $\mathcal{G}$, and refer to them as the *parents* of V in $\mathcal{G}$; with $\text{an}_{\mathcal{G}}(V)$ the *ancestors* of V in $\mathcal{G}$, namely the nodes with a directed path ending at V; with $\text{de}_{\mathcal{G}}(V)$ the *descendants* of V in $\mathcal{G}$, namely the nodes with a directed path starting at V; with $\text{sp}_{\mathcal{G}}(V)$ the variables connected to V with a bi-directed edge; and with $\perp\!\!\!\perp_{\mathcal{G}}$ *d-separation* in $\mathcal{G}$.

We describe the joint distribution of $\boldsymbol{V}$, $p(\boldsymbol{V})$, using SCM $\mathcal{M}$ compatible with $\mathcal{G}$, 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 latent random variables, and $\mathcal{F} = \{f_V\}_{V \in \boldsymbol{V}}$ where f_V is a function that assigns a value to $V \in \boldsymbol{V}$ based on the values of $\text{pa}_{\mathcal{G}}(V) \subset \boldsymbol{V}$ and $\boldsymbol{U}_V \subseteq \boldsymbol{U}$, i.e. $f_V \colon \mathcal{R}_{\text{pa}_{\mathcal{G}}(V)} \times \mathcal{R}_{\boldsymbol{U}_V} \mapsto \mathcal{R}_V$ (where e.g. $\mathcal{R}_{\text{pa}_{\mathcal{G}}(V)}$ denotes the range of $\text{pa}_{\mathcal{G}}(V)$) and $V = f_V(\text{pa}_{\mathcal{G}}(V), \boldsymbol{U}_V)$. Therefore f_V dictates the behavior of V based on $\text{pa}_{\mathcal{G}}(V)$ and the latent direct causes $\boldsymbol{U}_V$. The presence of a bi-directed edge between V and W corresponds to $\boldsymbol{U}_V \cap \boldsymbol{U}_W \neq \emptyset$. The joint distribution $p(\boldsymbol{V})$ is therefore given by $p(\boldsymbol{V}) = \int p(\boldsymbol{U}) \prod_{V \in \boldsymbol{V}} \delta_{V = f_V(\text{pa}_{\mathcal{G}}(V), \boldsymbol{U}_V)} d\boldsymbol{U}$.

We consider two types of interventions on a variable $X \in \boldsymbol{I}$: i) setting X to value x, corresponding to the classic *hard intervention* $\text{do}(X = x)$, and (ii) setting X as a deterministic function of a set of variables $\boldsymbol{C}_X \subseteq \boldsymbol{V} \backslash \{X, Y\}$, that we call *functional intervention*[2]. More precisely a functional intervention replaces f_X with another function $\pi_{X|\boldsymbol{C}_X} \colon \mathcal{R}_{\boldsymbol{C}_X} \mapsto \mathcal{R}_X$ such that $X = \pi_{X|\boldsymbol{C}_X}(\boldsymbol{C}_X)$. The set $\boldsymbol{C}_X$ therefore becomes the new direct causes of X. To comply with the literature used in this work, we call $\boldsymbol{C}_X$ the *context* of X [Lee and Bareinboim, 2020]. Hard and functional interventions can be classified as *deterministic interventions*, as they replace the original causal mechanism with a constant and with a deterministic function of new causes respectively. In contrast, a *stochastic intervention* would change f_X with $\pi_{X|\boldsymbol{C}_X, \boldsymbol{U}'_X} \colon \mathcal{R}_{\boldsymbol{C}_X} \times \mathcal{R}_{\boldsymbol{U}'_X} \mapsto \mathcal{R}_X$ where $\boldsymbol{U}'_X \cap \boldsymbol{U} = \emptyset$, such that $X = \pi_{X|\boldsymbol{C}_X, \boldsymbol{U}'_X}(\boldsymbol{C}_X, \boldsymbol{U}'_X)$.

The space of deterministic interventions can be formalized using the concepts of *mixed policy scope* (MPS) and *deterministic mixed policy* (DMP) introduced in Lee and Bareinboim [2020]. In the definitions below, we view the hard intervention $\text{do}(X = x)$ as a functional intervention with empty context that replaces f_X with the empty function $\pi_{X|\emptyset} \colon \emptyset \mapsto x$, corresponding to the assignment $X = \pi_{X|\emptyset}(\emptyset)$.

[1] A directed path is a sequence of linked nodes whose edges are directed and point from preceding towards following nodes in the sequence.

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

Definition 2.1 (Mixed Policy Scope (MPS))**.** A mixed policy scope $\mathcal{S}$ for $\mathcal{G}$ 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 the context for each intervention. For example, the 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, describing the interventions carried out in Fig. 1(right). We denote the intervention variables and contexts included in $\mathcal{S}$ by $\boldsymbol{X}_{\mathcal{S}} = \{X \colon \langle X, \boldsymbol{C}_X \rangle \in \mathcal{S}\}$ and $\boldsymbol{C}_{\mathcal{S}} = \bigcup_{X \in \boldsymbol{X}_{\mathcal{S}}} \boldsymbol{C}_X$ respectively. The acyclicity constraint formalizes the requirement that contexts of interventions are not descendants of intervened variable.

Definition 2.2 (Deterministic Mixed Policy (DMP))**.** Let $\mathcal{S}_{\text{hard}} = \{\langle X, \boldsymbol{C}_X \rangle \in \mathcal{S} \colon \boldsymbol{C}_X = \emptyset\}$. A deterministic mixed policy $\pi_{\mathcal{S}}$ compatible with MPS $\mathcal{S}$ is defined as the set $\pi_{\mathcal{S}} = \{\pi_{X|\boldsymbol{C}_X}\}_{\langle X, \boldsymbol{C}_X \rangle \in \mathcal{S} \backslash \mathcal{S}_{\text{hard}}} \bigcup \{\pi_{X|\emptyset}(\emptyset)\}_{\langle X, \boldsymbol{C}_X \rangle \in \mathcal{S}_{\text{hard}}}$, where $\pi_{X|\boldsymbol{C}_X} \colon \mathcal{R}_{\boldsymbol{C}_X} \mapsto \mathcal{R}_X$ while $\pi_{X|\emptyset}(\emptyset)$ is the value returned by an empty function.

A deterministic mixed policy specifies, $\forall \langle X, \boldsymbol{C}_X \rangle \in \mathcal{S} \backslash \mathcal{S}_{\text{hard}}$, the function $\pi_{X|\boldsymbol{C}_X}$ that replaces $f_X \in \mathcal{F}$ in $\mathcal{M}$, and $\forall \langle X, \boldsymbol{C}_X \rangle \in \mathcal{S}_{\text{hard}}$, the value to which X is set. These replacements induce a variant $\mathcal{M}_{\pi_{\mathcal{S}}}$ of $\mathcal{M}$ with joint distribution denoted by $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$. We refer to $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$ as *interventional distribution* and to an observation from $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$ as an *interventional data sample*.

3 FCGO PROBLEM

Let $\mu^Y_{\pi_{\mathcal{S}}}$ denote the expectation of $Y \in \boldsymbol{V}$ with respect to $p_{\pi_{\mathcal{S}}}(\boldsymbol{V})$, 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 Σ of MPSs for $\mathcal{G}$ and the space $\Pi_{\mathcal{S}}$ of DMPs that are compatible with MPS $\mathcal{S}$, which we call the *functional causal global optimization* (fCGO) problem.

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

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

The fCGO problem can be seen as an extension of the CGO problem introduced in Aglietti et al. [2020] to consider functional interventions in addition to hard interventions. In Section 3.1, we introduce graphical criteria that establish when such an extension might lead to a smaller target effect (the proofs are given in Appendix A). In addition, in Section 3.2, we show that solving the fCGO problem is equivalent to solving an optimization problem where the target effects are conditioned on variables defining sub-groups of a population, which guarantees that the found optimal policy does not underperform on those sub-groups (the proofs are given in Appendix B). Importantly, the result in Lee and Bareinboim [2020](Proposition 1) implies the solution of the fCGO problem leads to the same target effect that would be obtained but extending the search space to also considering stochastic interventions. Therefore, solving the fCGO solves the most general problem of finding interventions that minimize a target effect over all possible intervention types.

Solving the fCGO problem is more complex than solving the CGO problem and requires computing distances between functions defined over different contexts. In Section 4.2 we propose to model each target effect via a GP whose kernel allows computing such distances. We discuss how this approach allows us to keep the computational tractability of standard BO methods while allowing to flexibly specify the type of functional interventions.

3.1 HARD INTERVENTIONS (SUB-)OPTIMALITY

Let Σ_{hard} denotes the subset of Σ whose elements are the MPSs made of only hard interventions, i.e. $\Sigma_{\text{hard}} = \{\mathcal{S} \in \Sigma \colon \mathcal{S} \backslash \mathcal{S}_{\text{hard}} = \emptyset\}$. In this section, we introduce two propositions that provide graphical conditions under which restricting the search space in the fCGO problem from Σ to Σ_{hard} leads to a smaller or the same target effect, thereby informing an investigator about whether to consider functional interventions or not.

Proposition 3.2 (Sub-optimality of hard interventions)**.** *Let $\mathcal{G}$ be an arbitrary graph in which (i) $\exists C \in pa_{\mathcal{G}}(Y)$ with $C \notin \boldsymbol{I}$, or (ii) $\exists C \in sp_{\mathcal{G}}(Y)$. If there 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 restricting the search space in f*CGO *from Σ to Σ_{hard} would lead to a suboptimal target effect.*

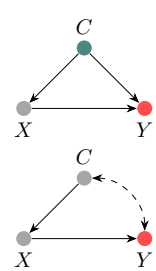

The above proposition captures two conditions for the sub-optimality of hard interventions: (i) the existence of non a intervenable variable C among the parents of Y that can also serve as context, and (ii) the existence of an unobserved confounder between a potential context C and the target node Y, forming all graphs in which a direct cause (i.e. a parent or an unobserved cause) of Y is non-intervenable and can serve as or inform a context for a functional intervention. An example of graph satisfying each of the two conditions is given on the right. It is thus possible to find a SCM such that hard interventions are suboptimal: we give two such SCMs in Appendix A. Indeed, a hard intervention on X would cut the incoming edges into

X thus disregarding the information that C provides about Y along the paths $C \to Y$ or $C \leftrightarrow Y$. Instead, a functional intervention on X with context C would keep the paths $X \leftarrow C \to Y$ and $X \leftarrow C \leftrightarrow Y$ open and could assign intervention values to X informed by values of C.

Proposition 3.3 (Optimality of hard interventions)**.** *If all directed causes of Y are in $\boldsymbol{I}$ then there exists a hard intervention on $pa_{\mathcal{G}}(Y)$ that is optimal.*

3.2 SUB-GROUP PERFORMANCE

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 context variables. For instance, in the example of Fig. 1, the investigator might want to find the interventions minimizing the level of PSA in a given population as well as in a specific sub-population of individuals aged over 65, i.e. the conditional target effect $\mu^{\text{PSA}}_{\pi_{\mathcal{S}},\text{Age}>65} := \mathbb{E}_{p_{\pi_{\mathcal{S}}}}[\text{PSA} \,|\, \text{Age} > 65]$, since a high percentage of prostate cancer cases are diagnosed within this group [Rawla, 2019]. Additionally, she may want to ensure that the new optimal policy $\pi^*_{\mathcal{S}^*}$ would not negatively affect other groups, e.g. younger patients (Age $<= 65$). This can be seen as the problem of minimizing the conditional target effect $\mu^Y_{\pi_{\mathcal{S}},\boldsymbol{C}=\boldsymbol{c}} := \mathbb{E}_{p_{\pi_{\mathcal{S}}}}[Y \,|\, \boldsymbol{C} = \boldsymbol{c}]$ where $\boldsymbol{C}$ is a subset of variables that can be used as contexts for $\boldsymbol{X}_{\mathcal{S}}$ with corresponding values[3] $\boldsymbol{c} \in \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}}$. Remarkably, under some mild conditions on $\boldsymbol{C}$, a DMP that solves the fCGO problem also minimizes the conditional target effect.

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 $\mathcal{S}^*, \pi^*_{\mathcal{S}^*} = \arg\min_{\mathcal{S}\in\Sigma,\pi_{\mathcal{S}}\in\Pi_{\mathcal{S}}} \mu^Y_{\pi_{\mathcal{S}},\boldsymbol{C}=\boldsymbol{c}}$ $\forall \boldsymbol{C} \subset \boldsymbol{V}\backslash Y$ such that $\boldsymbol{C} \cap de_{\mathcal{G}}(\boldsymbol{I}) = \emptyset$ and $\forall \boldsymbol{c} \in \mathcal{R}_{\boldsymbol{C}}$.*

Therefore, a method that solves the fCGO problem also solves the problem of optimizing conditional target effects. Importantly, this does not hold for methods that only consider hard interventions.

4 METHODOLOGY

We propose to solve the fCGO problem with the following *functional causal Bayesian optimization* (fCBO) method. First, we reduce the search space from Σ to a subset $\mathbb{M}_\Sigma$ as described in Section 4.1. Second, for each MPS $\mathcal{S} \in \mathbb{M}_\Sigma$, we initialize a Gaussian process (GP) $g_{\mathcal{S}}(\pi_{\mathcal{S}})$ that models the unknown target effect $\mu^Y_{\pi_{\mathcal{S}}}$ as described in Section 4.2. Third, we propose solving the minimization problem with the following a sequential strategy over T trials. At each trial t: (1) MPS $\mathcal{S}_t$ and DMP $\pi^t_{\mathcal{S}_t}$ are selected via the EI 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 MPS $\mathcal{S}_t$; (5) the posterior distribution of the GP $g_{\mathcal{S}_t}$, denoted by $\tau(g_{\mathcal{S}_t} \,|\, \mathcal{D}^I_{\mathcal{S}_t})$, is updated. Once the max number of trials is reached, $(\mathcal{S}^*, \pi^*_{\mathcal{S}^*})$ giving the minimum estimated target effect in $\{\mathcal{D}^I_{\mathcal{S}}\}_{\mathcal{S}\in\mathbb{M}_\Sigma}$ is returned. This procedure is summarized in Algorithm 1.

Algorithm 1 fCBO

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

Notice that fCBO only requires observing realizations from the marginal $p_{\pi^t_{\mathcal{S}_t}}(Y)$ (or been given $\hat{\mu}^Y_{\pi^t_{\mathcal{S}_t}}$). This is a considerable, practical, advantage compared to context-specific reward approaches such as Arsenyan et al. [2023] that, similarly to non-causal contextual BO methods [Krause and Ong, 2011], requires observing context pre-intervention and receiving context values at every trial. Similarly to recent approaches in contextual BO [Feng et al., 2020], our method handles this issue out-of-the-box, and can operate on aggregate rewards directly.

4.1 SEARCH SPACE REDUCTION

The cardinality of Σ grows exponentially with $\boldsymbol{I}$ and the number of possible contexts $\boldsymbol{C}_X$ for each X. We use the results in Lee and Bareinboim [2020] to reduce the search to the subset of *non-redundant* MPSs included in Σ, denoted by $\mathbb{M}_\Sigma$, which is guaranteed to contain a solution of the fCGO problem. For completeness and clarity, we describe these results in the deterministic mix policy setting.

$\mathcal{S}' \subseteq \mathcal{S}$ is used to indicate that $\boldsymbol{X}_{\mathcal{S}'} \subseteq \boldsymbol{X}_{\mathcal{S}}$ and $\boldsymbol{C}'_X \subseteq \boldsymbol{C}_X$, $\forall \langle X, \boldsymbol{C}'_X \rangle \in \mathcal{S}'$. Further, $\pi_{\mathcal{S}'} \subseteq \pi_{\mathcal{S}}$ is used to indicate that $\pi_{X|\boldsymbol{C}'_X}(\boldsymbol{c}'_X) = \int \pi_{X|\boldsymbol{C}_X}(\boldsymbol{c}'_X \cup \boldsymbol{c}''_X) p_{\pi_{\mathcal{S}}}(\boldsymbol{c}''_X \,|\, \boldsymbol{c}'_X) d\boldsymbol{c}''_X$, $\forall X \in \boldsymbol{X}_{\mathcal{S}'}, \boldsymbol{c}''_X \in \mathcal{R}_{\boldsymbol{C}_X \backslash \boldsymbol{C}'_X}$.

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 $\pi'_{\mathcal{S}} \subset \pi_{\mathcal{S}}$.

[3] We assume that the range of values observed for $\boldsymbol{C}_{\mathcal{S}}$ does not change after performing $\pi_{\mathcal{S}}$.

Proposition 4.2. *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)* $Y \not\perp\!\!\!\perp_{\mathcal{G}_{\mathcal{S}} \setminus \{X\}} C \mid \boldsymbol{C}_X \setminus \{C\}$ *for every* $X \in \mathbf{X}_{\mathcal{S}}$ *and* $C \in \mathbf{C}_X$.

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 work on Bayesian functional optimization that focuses on univariate functional inputs, $\pi_{\mathcal{S}}$ can include scalar values as well as functions potentially defined on different input spaces.

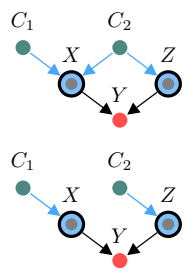

For instance, for MPS $\mathcal{S} = \{\langle X, \{C_1, C_2\}\rangle, \langle Z, \{C_2\}\rangle\}$ with $\mathcal{G}_{\mathcal{S}}$ on the top right, $\pi_{X|\{C_1,C_2\}}$ is defined over $\mathcal{R}_{C_1} \times \mathcal{R}_{C_2}$, while $\pi_{Z|C_2}$ over $\mathcal{R}_{C_2}$. Alternatively, for $\mathcal{S} = \{\langle X, \{C_1\}\rangle, \langle Z, \{C_2\}\rangle\}$ with $\mathcal{G}_{\mathcal{S}}$ 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 $m_{\mathcal{S}}$ and $K^{\theta}_{\mathcal{S}}$ denote the prior mean and covariance functional with hyperparameters θ and $\pi, \pi' \in \Pi_{\mathcal{S}}$ (we omit the subscript $\mathcal{S}$ to simplify the notation[4]). $\Pi_{\mathcal{S}} := \mathbb{R}^{|\mathcal{S}_{\text{hard}}|} \times \mathcal{B}(\boldsymbol{C}_{\mathcal{S}})$ where[5] $\mathbb{R}^{|\mathcal{S}_{\text{hard}}|}$ 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 function on $\boldsymbol{C}_{\mathcal{S}}$. Given an interventional dataset $\mathcal{D}^I_{\mathcal{S}}$ for $\mathcal{S}$ for which we assume a Gaussian likelihood, the posterior distribution $\tau(g_{\mathcal{S}} \mid \mathcal{D}^I_{\mathcal{S}})$ can be computed by standard GP updates (see 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 $\theta = (\sigma_f^2, \ell)$ are the hyperparameters and $||\pi - \pi'||$ represents a distance between mixed inputs to the GP[6]. Let π_{hard} denote the vector whose elements are the scalar values in π, and π_{func} the vector whose elements are the functions included in π. We define $||\pi - \pi'||^2$ as $||\pi - \pi'||^2 := ||\pi_{\text{hard}} - \pi'_{\text{hard}}||^2 + ||\pi_{\text{func}} - \pi'_{\text{func}}||^2_{H^{\mathcal{S}}_{\kappa}}$, with $||\pi_{\text{hard}} - \pi'_{\text{hard}}||^2$ indicating the square of the Euclidean distance in $\mathbb{R}^{|\mathcal{S}_{\text{hard}}|}$, and $||\pi_{\text{func}} - \pi'_{\text{func}}||^2_{H^{\mathcal{S}}_{\kappa}}$ the distance between functions in a vector-valued reproducing kernel Hilbert space (RKHS, Aronszajn [1950]) $\mathcal{B}(\boldsymbol{C}_{\mathcal{S}}) = H^{\mathcal{S}}_{\kappa}$ described below.

More specifically, $H^{\mathcal{S}}_{\kappa}$ is a RKHS with vector-valued reproducing kernel $\kappa^{\phi} : \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \times \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \to \mathbb{R}^{|\mathcal{S}_{\text{func}}| \times |\mathcal{S}_{\text{func}}|}$ where ϕ denotes the hyperparameters. We refer to this kernel as the *functional intervention kernel* to distinguish it from $K^{\theta}_{\mathcal{S}}$. We thus have $||\pi_{\text{func}} - \pi'_{\text{func}}||^2_{H^{\mathcal{S}}_{\kappa}} = \langle \pi_{\text{func}} - \pi'_{\text{func}}, \pi_{\text{func}} - \pi'_{\text{func}} \rangle_{H^{\mathcal{S}}_{\kappa}}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ denotes the inner product in the space $\mathcal{H}$. Evaluating this quantity requires computing κ^{ϕ} at different input values for the variables in $\boldsymbol{C}_{\mathcal{S}}$, say $\mathbf{c}_{\mathcal{S}}$ and $\mathbf{c}'_{\mathcal{S}}$, for π_{func} and π'_{func}.

We write the vector of functions π_{func} included in the RKHS $H^{\mathcal{S}}_{\kappa}$ as $\pi_{\text{func}}(\cdot) = \sum_{i=1}^{N} \kappa(\mathbf{c}^i_{\mathcal{S}}, \cdot)\boldsymbol{\alpha}_i$ where $\boldsymbol{\alpha}_i \in \mathbb{R}^{|\mathcal{S}_{\text{func}}|}$ and $\mathbf{c}^i_{\mathcal{S}} \in \mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}}$. This implies

$$\langle \pi_{\text{func}} - \pi'_{\text{func}}, \pi_{\text{func}} - \pi'_{\text{func}} \rangle_{H^{\mathcal{S}}_{\kappa}} = \sum_{i=1}^{N}\sum_{j=1}^{N} \boldsymbol{\alpha}_i^{\top} \kappa^{\phi}(\mathbf{c}^i_{\mathcal{S}}, \mathbf{c}^j_{\mathcal{S}}) \boldsymbol{\alpha}_j$$
$$+ \sum_{i=1}^{M}\sum_{j=1}^{M} \boldsymbol{\beta}_i^{\top} \kappa^{\phi}(\mathbf{c}^i_{\mathcal{S}}, \mathbf{c}^j_{\mathcal{S}}) \boldsymbol{\beta}_j - 2\sum_{i=1}^{N}\sum_{j=1}^{M} \boldsymbol{\alpha}_i^{\top} \kappa^{\phi}(\mathbf{c}^i_{\mathcal{S}}, \mathbf{c}^j_{\mathcal{S}}) \boldsymbol{\beta}_j,$$

where we have exploited $\pi'_{\text{func}}(\cdot) = \sum_{i=1}^{M} \kappa^{\phi}(\mathbf{c}^i_{\mathcal{S}}, \cdot)\boldsymbol{\beta}_i$ with $\boldsymbol{\beta}_i \in \mathbb{R}^{|\mathcal{S}_{\text{func}}|}$.

To construct $\kappa^{\phi}(\cdot, \cdot)$, we propose to augment the input space by including a task index for each function $\pi_{X|\boldsymbol{C}_X}$, i.e. $\kappa^{\phi} \colon (\mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \times \mathcal{T}) \times (\mathcal{R}_{\boldsymbol{C}_{\mathcal{S}}} \times \mathcal{T}) \to \mathbb{R}^{|\mathcal{S}_{\text{func}}| \times |\mathcal{S}_{\text{func}}|}$ where $\mathcal{S}_{\text{func}} = \{\langle X, \boldsymbol{C}_X \rangle \in \mathcal{S} \colon \boldsymbol{C}_X \neq \emptyset\}$ and $\mathcal{T}$ is the space of integer values from one to $|\mathcal{S}_{\text{func}}|$. For every realisation of the context variables and the task index, say $(\mathbf{c}_{\mathcal{S}}, t)_i$, we can then evaluate $\kappa^{\phi}((\mathbf{c}_{\mathcal{S}}, t)_i, (\mathbf{c}_{\mathcal{S}}, t)_j)$. We assume the covariance between functions defined on different input spaces, i.e. for which $t_i \neq t_j$, to be zero[7]. 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. Thus $\kappa^{\phi}((\mathbf{c}_{\mathcal{S}}, t)_i, (\mathbf{c}_{\mathcal{S}}, t)_j)$ is equal to $k^{\phi[t_i]}_{t_i}(\mathbf{c}_{\mathcal{S}}[t_i], \mathbf{c}_{\mathcal{S}}[t_j])$ if $t_i = t_j$ and to 0 otherwise, where $\mathbf{c}_{\mathcal{S}}[t_i]$ denote the subset of values included in $\mathbf{c}_{\mathcal{S}}$ for the contexts of the t_i task and $\phi[t_i]$ denotes the subset of hyper-parameters for t_i included in ϕ. The kernel $k^{\phi[t_i]}_{t_i}$ might differ across tasks both in terms of functional form and hyperparameter values. This allows to impose different characteristics in terms of e.g. smoothness for each function $\pi_{X|\boldsymbol{C}_X}$ included in π.

4.3 ACQUISITION FUNCTIONAL

We sequentially select interventions by numerically maximizing the expected improvement (EI) per unit of cost $\mathtt{Co}(\cdot)$ *across* the MPS in $\mathbb{M}_{\Sigma}$[8]. Given an interventional dataset $\mathcal{D}^I_{\mathcal{S}}$,

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

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

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

[7] Alternative kernel constructions where this assumption is relaxed are discussed in Appendix C.

[8] While here we numerically optimize fEI, its functional gradient with respect to functions in a RKHS can be derived analytically (see Vien et al. [2018]) thus potentially speeding convergence.

		Search Space over MPSs	Search Space over DMPs
Causal BO	fCBO	Any MPS $\mathcal{S}$	Functional
	CBO	Any MPS $\mathcal{S}$, s.t., $\boldsymbol{C}_X = \emptyset$	Hard
	MCBO	Only $\mathcal{P}(\{a_i\}_{i \in \boldsymbol{I}})$	Hard$^{\#}$
	CoCa-BO	Any MPS $\mathcal{S}$, s.t., $\boldsymbol{C}_X$ the same across X	Hard^{+}
BO	BO	$\mathcal{S} = \{\langle X, \boldsymbol{C}_X \rangle \colon \forall X \in \boldsymbol{I}, \boldsymbol{C}_X = \emptyset\}$	Hard
	BFO	$\mathcal{S} = \{\langle X, \boldsymbol{C}_X \rangle \colon \forall X \in \boldsymbol{I}, \boldsymbol{C}_X \neq \emptyset\}$	Functional

Table 1: Comparison of fCBO against alternatives. $\mathcal{P}$ indicates the power set. Hard$^{\#}$ stands for hard interventions parameterizing fixed functions. Hard^{+} denotes the parameterization of conditional target effect via hard interventions.

for each $\mathcal{S} \in \mathbb{M}_\Sigma$ the functional EI (fEI) is given by:

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

where $\sigma^2_{\mathcal{S}}(\pi \,|\, \mathcal{D}^I_{\mathcal{S}}) = K^\theta_{\mathcal{S}}(\pi, \pi \,|\, \mathcal{D}^I_{\mathcal{S}})$, $\mathbf{\Phi}(\cdot)$ and $\phi(\cdot)$ are the CDF and PDF of a standard Gaussian random variable respectively, and $\gamma(\pi) = \frac{m_{\mathcal{S}}(\pi \,|\, \mathcal{D}^I_{\mathcal{S}}) - g^*}{K^\theta_{\mathcal{S}}(\pi, \pi \,|\, \mathcal{D}^I_{\mathcal{S}})}$ with g^* denoting the optimum value observed for $g_{\mathcal{S}}$ across MPSs in $\mathbb{M}_\Sigma$. $m_{\mathcal{S}}(\pi \,|\, \mathcal{D}^I_{\mathcal{S}})$ and $K^\theta_{\mathcal{S}}(\pi, \pi \,|\, \mathcal{D}^I_{\mathcal{S}})$ denote the posterior parameters of $\tau(g_{\mathcal{S}} \,|\, \mathcal{D}^I_{\mathcal{S}})$. 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 π. 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 experiment of Section 6.1, and exemplify the second cost in the HEALTH experiment of Section 6.2.

5 RELATED WORK

We introduced an extension of CBO [Aglietti et al., 2020] that also considers functional interventions. Our extension covers the full landscape of interventions required for solving the problem of optimizing a target variable in a system. Rather than using functional interventions, other CBO methods that can perform contextual interventions achieve this by either searching for hard interventions in separate sub-groups of a population (CoCa-BO) [Arsenyan et al., 2023], or by using "action" nodes that induce changes in the parametrization of the conditional distributions (MCBO) [Sussex et al., 2023] associated to a given SCM.

Specifically, MCBO considers interventions on a set of parent-less "action" nodes that, by taking different values, induce a change in the parameters of the distribution of their child node. Therefore, the type of contextual intervention that can be achieved with this approach corresponds to a *modification* of a node's fixed conditional distribution rather than to a *replacement* of its functional assignment in the SCM as in our case. This might lead to interventions that are less general and flexible. To use MCBO in its contextual mode, the authors assume a pre-existing separation between action nodes and observed variables in the given SCM. Such action nodes are all described as parent-less intervenable nodes. However, in many common real-world settings all possible contexts have parents, as for example in the HEALTH experiment of Section 6.2. In these cases only hard interventions are possible. In addition, MCBO does not consider changes of contexts, reduction of search spaces, and cannot handle unobserved confounders. A longer discussion on this method can be found in Appendix D.3.

Instead, CoCa-BO considers settings where investigators select a context specific intervention *after observing the context values* and observe the resulting context specific output. This can again be seen as an indirect method to obtain functional interventions via hard interventions that, by returning a discretized version of the optimal policy, might lack flexibility and lead to sup-optimal interventions for unseen or rarely observed context values. CoCa-BO assumes all intervention values to be determined by the same contexts, which significantly simplifies the optimization problem by allowing to model the target effect via a GP over the potentially high-dimensional action-context input space.

Finally, in the BO literature, functional optimization has been studied by searching over the space of Bernstein polynomials [Vellanki et al., 2019], constructing a sequence of low-dimensional search spaces [Shilton et al., 2020] or representing the functional inputs as elements in an RKHS [Vien et al., 2018]. This work takes an approach similar to Vien et al. [2018], 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 evaluate fCBO on a synthetic chain graph visualized in Section 6.1 (CHAIN), and on a real-world setting based on the healthcare graph in Fig. 1(a) (HEALTH). The experiments aim at highlighting three main advantages of fCBO

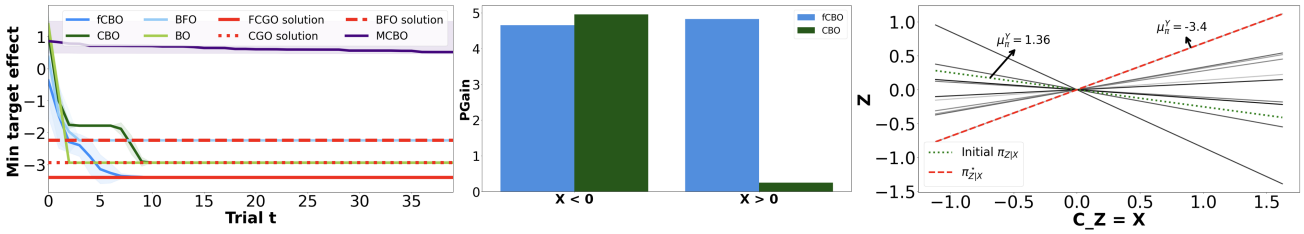

Figure 2: CHAIN experiment. *Left*: Mean of minimum target effect in $\{\mathcal{D}^I_{\mathcal{S}}\}_{\mathcal{S}\in\mathbb{M}_\Sigma}$ over 20 replicates obtained by fCBO and the baselines over 100 trials (shaded areas give $\pm$ std. dev., plot cropped at 40 trials due to early convergence). Red lines give optimal target effects corresponding to the solution of the CGO problem (achievable by CBO, BO), the fCGO problem (achievable by fCBO), and the BFO problem (achievable by fCBO). *Middle*: Mean of performance gain over 20 replicates. *Right*: The initial functional intervention included in $\mathcal{D}^I$ and the optimal one found for Z by fCBO.

w.r.t to other approaches to find optimal interventions. The first advantage is the ability to achieve lower target effects compared to methods that only allow hard interventions. We assess this by looking at the convergence to the global optimum. The second advantage is the ability to perform well on sub-groups. We demonstrate this in the CHAIN experiment, by computing the *performance gain* for implementing DMP $\pi_{\mathcal{S}}$ for a subgroup with $\boldsymbol{C} = \boldsymbol{c}$. This 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 (i.e. $\hat{\mu}^Y_{\boldsymbol{C}=\boldsymbol{c}} \approx \mathbb{E}_p[Y \mid \boldsymbol{C} = \boldsymbol{c}]$) and $\hat{\mu}^Y_{\pi_{\mathcal{S}},\boldsymbol{C}=\boldsymbol{c}}$ an estimate of the conditional target effect defined in Section 3.2. PGain can be interpreted as the improvement in the target effect attained by observing the system versus intervening according to $\pi_{\mathcal{S}}$. The third advantage is the ability to craft flexible and more targeted DMPs that can incur similar or lower cost, while still ensuring a better target effect than DMPs only made of hard interventions. We exemplify this in the HEALTH experiment with $\text{Co}_{\mathcal{S}}(\pi) = \sum_{X\in\boldsymbol{X}_{\mathcal{S}}} \int_{\mathcal{R}_{\boldsymbol{C}_X}} \pi_{X|\boldsymbol{C}_X}(\boldsymbol{c}_X) d\boldsymbol{c}_X$ (see Section 4.3). Specifically, in the main text we show how, given the same budget in terms of $\text{Co}_{\mathcal{S}}(\pi)$, fCBO achieves a lower target effect. In addition, in Appendix E, we show a setting in which fCBO simultaneously achieves a lower target effect while incurring a lower cost.

Comparison. We compare fCBO with BO [Shahriari et al., 2015], CBO, BFO [Vien et al., 2018] and MCBO. While all these algorithms use GPs to model the unknown target effects, the search spaces they explore differ. BO and BFO, which do not represent the underlying system via a casual graph, only perform intervention on all intervenable nodes $\boldsymbol{I}$ simultaneously. In addition, BO only considers hard interventions (i.e. $\mathcal{S} = \{\langle X, \boldsymbol{C}_X\rangle \colon X \in \boldsymbol{I}, \boldsymbol{C}_X = \emptyset\}$), while BFO only considers functional interventions where $\boldsymbol{C}_X$ is pre-specified and fixed over trials (i.e. $\mathcal{S} = \{\langle X, \boldsymbol{C}_X\rangle \colon X \in \boldsymbol{I}, \boldsymbol{C}_X \neq \emptyset\}$ for a pre-specified and fixed $\boldsymbol{C}_X$). MCBO operates on the system with "action nodes", which are parentless nodes that are set to different values in order to induce a change in the distribution of their child nodes. Sussex et al. do not provide guidelines on how their method could be used in general environments, and so we proceeded as follows. For the CHAIN experiment we interpret that to mean W is the only variable that could be considered as an action node, and provide a longer discussion on this choice in Section D.3.2. For the HEALTH experiment, we followed Sussex et al. [2023] and considered hard intervention on $\boldsymbol{I} = \{\text{CI}, \text{Statin}, \text{Aspirin}\}$. Note that we cannot compare to CoCa-BO as: (i) in our settings the contexts values 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. The different search spaces of all these methods are summarized in Table 1.

Experimental Setup and Details. 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 of intervened variables as context. 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 $\Pi_{\mathcal{S}}$, as functional intervention kernel κ^ϕ, we consider the linear and RBF kernel in the CHAIN and HEALTH experiments respectively. We use the same functional intervention representation for BFO. We initialize $\mathcal{D}^I_{\mathcal{S}}$ for each $\mathcal{S} \in \mathbb{M}_\Sigma$ by randomly generating a single initial DMP and associated target value. We provide average results across 20 different initializations of $\mathcal{D}^I_{\mathcal{S}}$. More experimental details are given in Appendix D.

6.1 CHAIN EXPERIMENTS

We first experiment on the augmented chain graph with associated SCM given on the right. Fig. 2(left) shows how considering mixes of hard and functional interventions allows fCBO to reach the lowest target effect

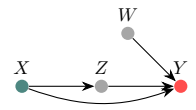

$$X = U_X, W = U_W$$
$$Z = -0.5X + U_Z$$
$$Y = -W - 3ZX + U_Y$$

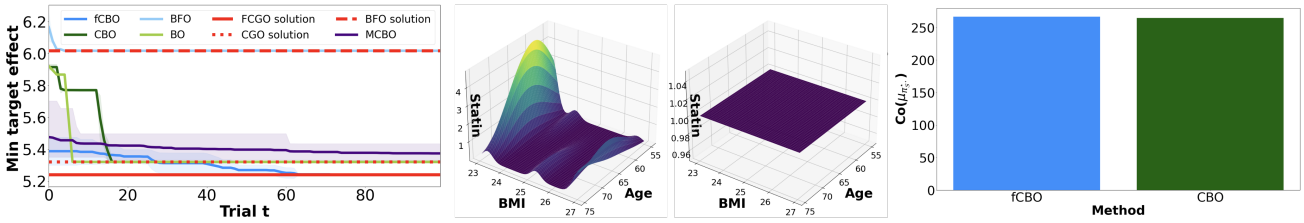

Figure 3: HEALTH experiment. *Left*: Mean of minimum target effect in $\{\mathcal{D}^I_{\mathcal{S}}\}_{\mathcal{S}\in\mathbb{M}_\Sigma}$ over 20 replicates obtained by fCBO and the baselines over 100 trials (shaded areas give $\pm$ std. dev.). *Middle*: Intervention determining allocation of Statin across levels of Age and BMI. *Right*: Comparison of overall cost (in terms of sum of area under the curve for all intervention nodes) for the best intervention found by fCBO vs. CBO.

value compared to algorithms that consider only hard interventions (BO, CBO and MCBO[9]) or only functional interventions (BFO). fCBO and CBO also differ in terms of target effects achieved across different values for the context variables. Fig. 2(middle) shows the performance gain obtained by implementing the optimal intervention picked by fCBO and CBO when the context variable X takes positive values ($\text{PGain}(\pi^*_{\mathcal{S}^*}, X > 0)$) or negative values ($\text{PGain}(\pi^*_{\mathcal{S}^*}, X < 0)$). Due to the interaction term $-3ZX$ in the structural equation for Y, minimizing Y would require setting Z to a negative value when $X < 0$ and to a positive value when $X > 0$. This cannot be achieved via hard interventions that set Z to a fixed value irrespective of X thus distributing the performance gain unevenly across subgroups. This is indeed observed for CBO which converges to $\mathcal{S}^* = \{\langle Z, \emptyset\rangle, \langle W, \emptyset\rangle\}$ and, by picking a negative values for Z ($\pi^*_{\mathcal{S}^*} = \{-1, 1\}$), achieves a very low gain for $X > 0$. On the contrary, fCBO allows tailoring a functional intervention that leads to a high gain for both $X < 0$ and $X > 0$. In particular, fCBO converges to $\mathcal{S}^* = \{\langle Z, X\rangle \langle W, \emptyset\rangle\}$, with $\pi^*_{\mathcal{S}^*} = \{\pi_{Z|X}, 1\}$ where $\pi_{Z|X}$ is a linear function highlighted in red in Fig. 2(right). The chosen function has a slope that gives an optimal Z value for both groups, corresponding to a higher target effect overall.

6.2 HEALTH EXPERIMENTS

For the HEALTH experiment, we use the SCM by Ferro et al. [2015] (see Appendix D.2). Fig. 3(left) shows how considering functional interventions allows fCBO to achieve the lowest target effect. In this case, both BO and CBO converge to $\mathcal{S}^* = \{\langle \text{CI}, \emptyset\rangle, \langle \text{Statin}, \emptyset\rangle, \langle \text{Aspirin}, \emptyset\rangle\}$ with $\pi^*_{\mathcal{S}^*} = \{1, 1, 0.1\}$ while fCBO converges to $\mathcal{S}^* = \{\langle \text{CI}, \emptyset\rangle, \langle \text{Statin}, \{\text{Age}, \text{BMI}\}\rangle, \langle \text{Aspirin}, \emptyset\rangle\}$ with $\pi^*_{\mathcal{S}^*} = \{1, \pi^*_{\text{Statin}|\text{Age},\text{BMI}}, 0.1\}$. MCBO intervenes on Statin, Aspirin and CI attaining an target effect that is close to that found by others hard intervention based methods. Fig. 3(middle) displays the optimal intervention function $\pi^*_{\text{Statin}|\text{Age},\text{BMI}}$ found for Statin by fCBO (left) and the value 1 found by CBO as a constant function over Age and BMI (right). These two plots show that, while hard intervention based methods are forced to assign interventions uniformly across the context space, methods that also allow functional interventions can concentrate on specific subspaces, in this case on lower values of Age and BMI that is on younger and lower-BMI individuals in a population. This has important cost implications (Fig. 3(right)). Indeed, not only fCBO outperforms the alternative methods in terms of target effects achieved, but it can also do so while incurring the same cost or lower compared to CBO. This result demonstrates another key property of functional interventions: taking the context values into account allows an investigator to allocate units of intervention to areas of the contexts space that are the most promising, thus achieving the lowest target effect at a desirable cost.

[9]The poor performance of MCBO may also be due to the number of trials taken and the hyperparameter configuration, which the authors did not provide means to tune. However, notice that this result is inline with the suboptimal behaviour reported by Sussex et al. [2023](Fig. 6 in the appendix) in a setting similar to the CHAIN graph we present here.

7 CONCLUSION

We presented fCBO, a method for exploring both hard and functional interventions with the goal of finding those optimizing a target effect. We introduced graphical criteria that establish when functional interventions could be necessary to achieve optimal target effects, as well as when hard interventions are sufficient. Furthermore, we introduced graphical criteria showing that, under mild conditions, functional interventions can unlock policies that are optimal with respect to sub-groups in a population under an optimal intervention for the average target effect. We demonstrated these takeaways with empirical results. Future directions of this work include the consideration of gradient-based optimization methods for the acquisition functional and the consideration of more flexible kernel construction for the GP functionals (see Appendix C). These directions would allow the identification of more flexible and general optimal interventions while speeding up the convergence of the algorithm.

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. *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. *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. PMLR, 2022.

Qing Feng, Ben Letham, Hongzi Mao, and Eytan Bakshy. High-dimensional contextual policy search with unknown context rewards using bayesian optimization. *Advances in Neural Information Processing Systems*, 33:22032–22044, 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. *Advances in Neural Information Processing Systems*, 24, 2011.

Finnian Lattimore, Tor Lattimore, and Mark D Reid. Causal bandits: Learning good interventions via causal inference. *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. PMLR, 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. PMLR, 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*, volume 2. MIT Press Cambridge, MA, 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. *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.