
Constrained Causal Bayesian Optimization

Virginia Aglietti¹ Alan Malek¹ Ira Ktena¹ Silvia Chiappa¹

Abstract

We propose constrained causal Bayesian optimization (cCBO), an approach for finding interventions in a known causal graph that optimize a target variable *under some constraints*. cCBO first reduces the search space by exploiting the graph structure and, if available, an observational dataset; and then solves the restricted optimization problem by modelling target and constraint quantities using Gaussian processes and by sequentially selecting interventions via a constrained expected improvement acquisition function. We propose different surrogate models that enable to integrate observational and interventional data while capturing correlation among effects with increasing levels of sophistication. We evaluate cCBO on artificial and real-world causal graphs showing successful trade off between fast convergence and percentage of feasible interventions.

1. Introduction

The problem of understanding which interventions in a system optimize a target variable is of relevance in many scientific disciplines, including biology, medicine, and social sciences. Often, the investigator might want to solve this problem while also ensuring that interventions satisfy some constraints. As an example, consider the protein-signalling network of Fig. 1(a), which describes the causal pathways linking several phosphorylated proteins and phospholipids in human immune system cells, including the mitogen-activated protein kinases Erk studied in cancer therapy (Frémin & Meloche, 2010; Sachs et al., 2005). An investigator might wish to find which variables among Mek, PKC, PKA and Akt to perturb and to which levels in order to minimize Erk, under the constraints of not inhibiting PKA and PKC which play an important role in the functioning of healthy cells. As another example, consider the graph of Fig. 1(b) describing

¹DeepMind, London, UK. Correspondence to: Virginia Aglietti <aglietti@deepmind.com>.




Figure 1: (a) Protein-signaling network and (b) causal PSA graph. Red, grey and pink nodes indicate target, intervenable, and constrained variables respectively.

the causal relationships between prostate specific antigen (PSA) and its risk factors, some of which might be due to an existing policy at the time of study (Ferro et al., 2015). An investigator might wish to understand which variables among Aspirin, Statin, and Calories Intake (CI) to fix and to which values in order to minimize PSA, under the constraint of keeping BMI below a certain value.

In this paper, we propose *constrained causal Bayesian optimization* (cCBO), a sequential approach for efficiently solving such constrained optimization problems in the setting of known causal graph that builds on the literature on causal Bayesian optimization (Aglietti et al., 2020), constrained Bayesian optimization (Gardner et al., 2014), and multi-task learning with Gaussian processes (GPs) (Alvarez et al., 2011). Our contributions are summarized as follows:

- We provide a formalization of the constrained optimization problem using the framework of structural causal models (SCMs).
- We introduce a novel theory for reducing the search space of interventions that eliminates intervention sets of higher cardinality.
- We propose different GP surrogate models for the unknown target and constraint quantities that leverage observational data, interventional data and the SCM structure, capturing correlations with increasing level of sophistication. This enables uncertainty quantification thereby speeding up the identification of an optimal intervention while limiting the number of *infeasible* interventions, i.e. not satisfy the constraints.
- We propose a constrained expected improvement acquisition function for sequentially selecting interventions

based on both output optimization and constraint satisfaction. This enables finding the optimal solution with a smaller number of interventions.

- We evaluate cCBO on synthetic and real-world graphs with different SCM characteristics showing how it successfully trades off achieving a fast convergence and collecting a high percentage of feasible interventions.

2. Constrained Causal Global Optimization

We consider a system of observable random variables \mathbf{V} with *target variable* $Y \in \mathbf{V}$ and *intervenable variables* $\mathbf{I} \subseteq \mathbf{V} \setminus Y^1$, and the problem of finding an intervention set $\mathbf{X} \subseteq \mathbf{I}$ and intervention values \mathbf{x} that optimize the expectation of Y while ensuring that the expectations of *constrained variables* $\mathbf{C} \subseteq \mathbf{V} \setminus Y$ are above/below certain values.

We assume that the system can be described by a *structural causal model* (SCM) (Pearl, 2000; Pearl et al., 2016) $\mathcal{M} = \langle \mathbf{V}, \mathbf{U}, \mathbf{F}, p(\mathbf{U}) \rangle$, where \mathbf{U} is a set of exogenous unobserved random variables with distribution $p(\mathbf{U}) = \prod_{U \in \mathbf{U}} p(U)$, and $\mathbf{F} = \{f_V\}_{V \in \mathbf{V}}$ is a set of deterministic functions such that $V = f_V(\text{pa}(V), \mathbf{U}_V)$ with $\text{pa}(V) \subseteq \mathbf{V} \setminus V$ and $\mathbf{U}_V \subseteq \mathbf{U}$, $\forall V \in \mathbf{V}$. \mathcal{M} has associated a *causal graph* \mathcal{G} with nodes \mathbf{V} , and with a directed edge from V to W if $V \in \text{pa}(W)$, in which case V is called a *parent* of W , and a bi-directed edge between V and W if $\mathbf{U}_V \cap \mathbf{U}_W \neq \emptyset$ ($\mathbf{U}_V \cap \mathbf{U}_W$ is an unobserved *confounder* between V and W). We assume that \mathcal{G} is *acyclic*, namely that there are no *directed paths*² starting and ending at the same node. 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 refer to the joint distribution of \mathbf{V} determined by $p(\mathbf{U})$, which we denote by $p(\mathbf{V})$, as *observational distribution*, and to an observation from $p(\mathbf{V})$ as *observational data sample*.

An *intervention* on $V \in \mathbf{I}$ setting its value to v , denoted by $\text{do}(V = v)$, corresponds to modifying \mathcal{M} by replacing $f_V(\text{pa}(V), \mathbf{U}_V)$ with v . We refer to the joint distribution of \mathbf{V} in the modified SCM under such an intervention, indicated by $p_{\text{do}(V=v)}(\mathbf{V})$, as *interventional distribution*, and to an observation from it as *interventional data sample*. Notice that, in the case of no unobserved confounders, $p_{\text{do}(V=v)}(\mathbf{V}) = \prod_{W \in \mathbf{V} \cap V} p(W | \text{pa}(W)) \delta_{V=v}$ with $\delta_{V=v}$ denoting a delta function centered at v . An intervention on a set of variables is defined similarly.

Let $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^Y := \mathbb{E}_{p_{\text{do}(\mathbf{X}=\mathbf{x})}}[Y]$ denote the expectation of Y

¹To simplify the notation we write Y to denote the set $\{Y\}$, and similarly throughout the paper.

²A directed path is a sequence of linked nodes whose edges are directed and point from preceding towards following nodes in the sequence.

Algorithm 1 cCBO

Inputs: $\mathcal{G}, \mathbf{I}, \mathbf{C}, Y, \mathcal{D}^I := \{\mathcal{D}_{\mathbf{X}}^I\}_{\mathbf{X} \in \mathbf{I}}, \mathcal{D}^O, \lambda^{\mathbf{C}}, S, T$
 $n\mathcal{M}_{\mathcal{C}[\mathbf{Y}, \mathcal{G}]} \leftarrow \text{cMI SReduce}(\mathcal{G}, \mathbf{I}, \mathbf{C}, Y, \mathcal{D}^O)$
 Initialise GPS $g_{\mathbf{X}}(\mathbf{x}) \forall \mathbf{X} \in n\mathcal{M}_{\mathcal{C}[\mathbf{Y}, \mathcal{G}]}$ with $\mathcal{D}_{\mathbf{X}}^I$ and \mathcal{D}^O
for $t = 1, \dots, T$ **do**
 1. Select intervention $(\mathbf{X}_t, \mathbf{x}_t)$ with cEI
 2. Obtain samples $\{\mathbf{c}_{\mathbf{X}_t}^{(s)}, y^{(s)}\}_{s=1}^S$ from the distribution $p_{\text{do}(\mathbf{X}_t=\mathbf{x}_t)}(\mathbf{C}_{\mathbf{X}_t}, Y)$, and use them to compute the sample mean estimate $\hat{\mu}_{\text{do}(\mathbf{X}_t=\mathbf{x}_t)}$
 3. $\mathcal{D}_{\mathbf{X}_t}^I \leftarrow \mathcal{D}_{\mathbf{X}_t}^I \cup (\mathbf{x}_t, \hat{\mu}_{\text{do}(\mathbf{X}_t=\mathbf{x}_t)})$
 4. Update $\tau(g_{\mathbf{X}_t} | \mathcal{D}_{\mathbf{X}_t}^I)$

end

Output: $(\mathbf{X}^*, \mathbf{x}^*)$ with min feasible $\hat{\mu}_{\text{do}(\mathbf{X}^*=\mathbf{x}^*)}^Y$ over \mathcal{D}^I

w.r.t. the interventional distribution³, which we refer to as the *target effect*. The goal of the investigator is to find a set $\mathbf{X} \subseteq \mathbf{I}$ and values \mathbf{x} that optimize the target effect while ensuring that $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^{\mathbf{C}}$, which we refer to as *constraint effect*, is greater/smaller than a threshold $\lambda^{\mathbf{C}}$, $\forall \mathbf{C} \in \mathbf{C} \subseteq \mathbf{V} \setminus Y$. Target and constraint effects are unknown.

The constraint for a variable $X \in \mathbf{C} \cap \mathbf{X}$ can directly be enforced by restricting the range of intervention values, $D(X)$, to be in accordance with the threshold. Instead, the constraints for $\mathbf{C}_{\mathbf{X}} := \mathbf{C} \setminus (\mathbf{C} \cap \mathbf{X})$ need to be included in the optimization problem. The constrained optimization problem can thus be formalized as follows.

Definition 2.1 (cCGO Problem). The *constrained causal global optimization* (cCGO) problem is the problem of finding a tuple (\mathbf{X}, \mathbf{x}) such that⁴

$$\mathbf{X}, \mathbf{x} = \arg \min_{\substack{\mathbf{X} \subseteq \mathcal{P}_{\mathbf{I}}, \\ \mathbf{x} \in D(\mathbf{X})}} \mu_{\text{do}(\mathbf{X}=\mathbf{x})}^Y \quad \text{s.t.} \quad \mu_{\text{do}(\mathbf{X}=\mathbf{x})}^{\mathbf{C}_{\mathbf{X}}} \geq \lambda^{\mathbf{C}_{\mathbf{X}}}, \quad (1)$$

where $\mathcal{P}_{\mathbf{I}}$ indicates the power set⁵ of \mathbf{I} , $D(\mathbf{X}) := \times_{X \in \mathbf{X}} D(X)$, and $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^{\mathbf{C}_{\mathbf{X}}}$ and $\lambda^{\mathbf{C}_{\mathbf{X}}}$ denote the constraint effects and corresponding thresholds on all variables in $\mathbf{C}_{\mathbf{X}}$.

The cCGO problem extends the *causal global optimization* (CGO) problem defined in Aglietti et al. (2020) to incorporate constraints.

3. Constrained Causal Bayesian Optimization

We propose to solve the cCGO problem using the *constrained causal Bayesian optimization* (cCBO) method sum-

³This expectation is commonly indicated with $\mathbb{E}[Y | \text{do}(\mathbf{X} = \mathbf{x})]$ (Pearl, 2000).

⁴Depending on the application, the investigator might be interested in maximizing $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^Y$ and/or ensuring that some or all constraints effects are smaller than the thresholds. In such cases Eq. (1) would need to be adjusted accordingly.

⁵Excluding \emptyset .

marized in Algorithm 1, which assumes known casual graph \mathcal{G} , continuous variables \mathbf{V} , and full observation of the system after an intervention. First, the search space is reduced to a subset $nM_{\mathcal{C}[Y,G]}$ of \mathcal{P}_I via the cMI SReduce procedure described in Section 3.1. Then, the restricted cCGO problem is solved by modelling, $\forall \mathbf{X} \in nM_{\mathcal{C}[Y,G]}$, the unknown target and constraint effects $\mu_{\text{do}(\mathbf{X}=\mathbf{x})} := (\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^Y, \mu_{\text{do}(\mathbf{X}=\mathbf{x})}^{\mathbf{C}_X})$ using Gaussian processes (GPs) (Rasmussen & Williams, 2006) $g_{\mathbf{X}}(\mathbf{x}) := (g_{\mathbf{X}}^V(\mathbf{x}))_{V \in \mathcal{C}_X[Y]}$, as described in Section 3.2, and via the following sequential strategy. At each trial $t = 1, \dots, T$: (1) a set of intervenable variables \mathbf{X}_t and intervention values \mathbf{x}_t are selected via a constrained expected improvement (cEI) acquisition function that accounts for both the target and constraint effects, as described in Section 3.3; (2) a set of S interventional data samples is obtained and used to get a sample mean estimate $\hat{\mu}_{\text{do}(\mathbf{X}_t=\mathbf{x}_t)}$ of $\mu_{\text{do}(\mathbf{X}_t=\mathbf{x}_t)}$; (3) $(\mathbf{x}_t, \hat{\mu}_{\text{do}(\mathbf{X}_t=\mathbf{x}_t)})$ is added to the interventional dataset $\mathcal{D}_{\mathbf{X}_t}^I$ of \mathbf{X}_t ; (4) the posterior distribution of $g_{\mathbf{X}_t}(\mathbf{x})$, denoted by $\tau(g_{\mathbf{X}_t} | \mathcal{D}_{\mathbf{X}_t}^I)$, is updated. Once the maximum number of trials T is reached, a tuple $(\mathbf{X}^*, \mathbf{x}^*)$ giving the minimum feasible target effect value in $\mathcal{D}^I := \{\mathcal{D}_{\mathbf{X}}^I\}_{\mathbf{X} \in I}$ is returned.

3.1. Reducing the Search Space

The cardinality of the power set, $|\mathcal{P}_I|$, grows exponentially with the cardinality of I , thus solving the cCGO problem by exploring the entire set could be prohibitively expensive. Even when $|\mathcal{P}_I|$ is small, reducing the search space would simplify the problem as a smaller number of comparisons between constraint and target effects would be required. In this section, we propose a procedure, which we refer to as cMI SReduce, that leverages invariances of the target and constraint effects w.r.t. different intervention sets to eliminate intervention sets of higher cardinality. This is achieved by extending the theory of minimal intervention sets of Lee & Bareinboim (2018) to account for the presence of constraints (a summary of the procedure and all proofs are given Appendix A).

The search space is first reduced from \mathcal{P}_I to the set of *constrained minimal intervention sets* (cMISs) relative to $(\mathbf{C} \cup Y, \mathcal{G})$, denoted by $M_{\mathcal{C}[Y,G]}$.

Definition 3.1 (cMIS). A set $\mathbf{X} \subseteq I$ is said to be a cMIS relative to $(\mathbf{C} \cup Y, \mathcal{G})$ if there is no $\mathbf{X}^\theta \subset \mathbf{X}$ with $\mathbf{C}_X = \mathbf{C}_{X^\theta}$ such that $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^W = \mu_{\text{do}(\mathbf{X}^\theta=\mathbf{x}^\theta)}^W$, where \mathbf{x}^θ indicates the subset of \mathbf{x} corresponding to variables $\mathbf{X}^\theta, \forall \mathbf{x} \in D(\mathbf{X}), \forall W \in \mathbf{C}_X \cup Y$, and \forall SCM compatible with \mathcal{G} .

The following theorem justifies this reduction.

Theorem 3.2 (Sufficiency of $M_{\mathcal{C}[Y,G]}$). $M_{\mathcal{C}[Y,G]}$ contains a solution of the cCGO problem (if a solution exists), \forall SCM compatible with \mathcal{G} .




Figure 2: (a) Causal graph with $I = \{A, D, E\}$ and $\mathbf{C} = \{C, D, E\}$. (b) Causal graph with $I = \mathbf{C} = \{X, Z\}$.

Let $\text{an}(W, \mathcal{G})$ be the set of *ancestors* of W in \mathcal{G} , i.e. the nodes with a directed path to W ; $\text{an}(W, \mathcal{G}) := \cup_{W \neq Z} \text{an}(W, \mathcal{G})$; and $\mathcal{G}_{\overline{\mathbf{X}}}$ the graph with all incoming links onto all elements of \mathbf{X} removed. The following proposition gives a graphical criterion for identifying $M_{\mathcal{C}[Y,G]}$.

Proposition 3.3 (Characterization of cMIS). $\mathbf{X} \subseteq I$ is a cMIS relative to $(\mathbf{C} \cup Y, \mathcal{G}) \iff \mathbf{X} \subseteq \text{an}(\mathbf{C} \cup Y, \mathcal{G}_{\overline{\mathbf{X}}}) \cup (\mathbf{C} \cap \mathbf{X})$.

If an observational dataset \mathcal{D}^O is available, the search space is further reduced from $M_{\mathcal{C}[Y,G]}$ to $nM_{\mathcal{C}[Y,G]}$ by checking, $\forall \mathbf{X} \in M_{\mathcal{C}[Y,G]}$, if all reducible variables in \mathbf{C}_X are null-feasible.

Definition 3.4 (Reducibility). $C \in \mathbf{C}_X$ is reducible if $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^C = \mu^C := \mathbb{E}_p[C], \forall \mathbf{x} \in D(\mathbf{X})$.

Definition 3.5 (Null-feasibility). $C \in \mathbf{C}_X$ is null-feasible if $\mu^C \geq \lambda^C$.

If $\mathbf{X} \cap \text{an}(C, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$, $C \in \mathbf{C}_X$ is reducible. Indeed, if $\mathbf{X} \cap \text{an}(C, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$, by Lemma A.1 in Appendix A (with $\mathbf{W}_1 = \mathbf{X}, \mathbf{W}_2 = C, \mathbf{W}_3 = \emptyset$) we have $C \perp_{\mathcal{G}_{\overline{\mathbf{X}}}} \mathbf{X}$, which implies $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^C = \mu^C$ by rule 3 of do-calculus (Pearl, 2000).

If a reducible $C \in \mathbf{C}_X$ is not null-feasible, \mathbf{X} is removed from the search space as it is not a solution of the cCGO problem. If instead C is null-feasible, all $\mathbf{X}^\theta \supset \mathbf{X}$ in $M_{\mathcal{C}[Y,G]}$ that are of the form $\mathbf{X}^\theta = \mathbf{X} \cup \mathbf{X}_1$ where $\mathbf{X}_1 \cap \text{an}(\mathbf{C}_X \setminus \mathbf{C} \cup Y, \mathcal{G}_{\overline{\mathbf{X}^\theta}}) = \emptyset$, and for which (i) $\mathbf{C}_{X^\theta} = \mathbf{C}_X$ or (ii) all variables in $\mathbf{C}_X \setminus \mathbf{C}_{X^\theta}$ are reducible and null-feasible, are eliminated from the search space. Indeed, in such cases, if \mathbf{X}^θ were a solution of the cCGO problem, then \mathbf{X} would also be a solution as: (a) $\forall W \in (\mathbf{C}_X \setminus \mathbf{C}) \setminus (\mathbf{C} \cap \mathbf{X}_1) \cup Y$, $\mu_{\text{do}(\mathbf{X}^\theta=\mathbf{x}^\theta)}^W = \mu_{\text{do}(\mathbf{X}=\mathbf{x})}^W$ (Lemma A.1 in Appendix A (with $\mathbf{W}_1 = \mathbf{X}_1, \mathbf{W}_2 = \mathbf{C}_X \setminus \mathbf{C} \cup Y, \mathbf{W}_3 = \mathbf{X}$) gives $\mathbf{C}_X \setminus \mathbf{C} \cup Y \perp_{\mathcal{G}_{\overline{\mathbf{X}^\theta}}} \mathbf{X}_1 | \mathbf{X}$, which implies $\mu_{\text{do}(\mathbf{X}^\theta=\mathbf{x}^\theta)}^W = \mu_{\text{do}(\mathbf{X}=\mathbf{x})}^W$ by rule 3 of do-calculus); (b) $\forall W \in \mathbf{C} \cap \mathbf{X}_1$ the constraint effects are satisfied for \mathbf{X} as $\mathbf{C} \cap \mathbf{X}_1 = \mathbf{C}_X \setminus \mathbf{C}_{X^\theta}$.

Example: We describe the search space reduction for the casual graph of Fig. 2(a). In this graph $M_{\mathcal{C}[Y,G]} = \mathcal{P}_I = \{\{A\}, \{D\}, \{E\}, \{A, D\}, \{A, E\}, \{D, E\}, \{A, D, E\}\}$, as $\mathbf{X} \subseteq \text{an}(\mathbf{C} \cup Y, \mathcal{G}_{\overline{\mathbf{X}}}) \cup (\mathbf{C} \cap \mathbf{X}), \forall \mathbf{X} \in \mathcal{P}_I$ (e.g. $A \in \text{an}(\mathbf{C} \cup Y, \mathcal{G}_{\overline{A}}) \cup (\mathbf{C} \cap A)$). Consider $\mathbf{X} = \{D, E\} \in M_{\mathcal{C}[Y,G]}$ and $C = X$. $\mathbf{X} \cap \text{an}(C, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$,

and therefore C is reducible. If C is not null-feasible, \mathbf{X} is removed from the search space. If instead C is null-feasible, the set $\mathbf{X}^\theta = \{A, D, E\} \supset \mathbf{X}$ in $\mathcal{M}_{C|Y,G}$ is of the form $\mathbf{X}^\theta = \mathbf{X} \cup A$ where $A \not\subseteq \text{an}(C_{\mathbf{X}} \setminus C \cup Y, \mathcal{G}_{\overline{\mathbf{X}}^\theta})$. Therefore \mathbf{X}^θ is removed from the search space. All other sets in $\mathcal{M}_{C|Y,G}$ have constrained variables that are non reducible and thus need to be included in $n\mathcal{M}_{C|Y,G}$.

3.2. Gaussian Processes Surrogate Models

For each $\mathbf{X} \in n\mathcal{M}_{C|Y,G}$ and $V \in C_{\mathbf{X}} \cup Y$, we model $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^V$ with a GP $g_{\mathbf{X}}^V(\mathbf{x}) \sim \mathcal{GP}(m_{\mathbf{X}}^V(\mathbf{x}), S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta))$, as GPs allow constructing flexible surrogate models while enabling uncertainty quantification and closed form updates. We propose one single-task GP (STGP) and two multi-task GPs (MTGP and \mathcal{G} -MTGP) which capture the correlation across effects with increasing level of sophistication.

For $V \neq W$, the single-task GP treats $g_{\mathbf{X}}^V(\mathbf{x})$ and $g_{\mathbf{X}}^W(\mathbf{x}^\theta)$ as independent, while the multi-task GPs model their correlation via a covariance matrix $S_{\mathbf{X}}^{V,W}(\mathbf{x}_1, \mathbf{x}_2)$, with (i, j) -th element given by $\mathbb{E}[g_{\mathbf{X}}^V(\mathbf{x}_i)g_{\mathbf{X}}^W(\mathbf{x}_j)] - \mathbb{E}[g_{\mathbf{X}}^V(\mathbf{x}_i)]\mathbb{E}[g_{\mathbf{X}}^W(\mathbf{x}_j)]$, either by assuming a common latent structure among the GPs (MTGP) or, for the setting of no unobserved confounders and under the assumption $V = f_V(\text{pa}(V)) + U_V$ with $p(U_V) = \mathcal{N}(0, \sigma_V^2)$, by explicitly exploiting the SCM structure (\mathcal{G} -MTGP). We propose different prior parameters constructions for STGP and MTGP, including one that leverages the availability of an observational dataset \mathcal{D}^O (STGP⁺ and MTGP⁺). The different GP constructions allow the investigator to model the target and constraint effects in both settings where no information about the system is available and black-box models are preferred (STGP and MTGP) and settings in which one can leverage different sources of information and integrate them in a structured prior formulation that quantifies uncertainty in a principled way (STGP⁺, MTGP⁺, and \mathcal{G} -MTGP).

For all surrogate models, after an intervention (\mathbf{X}, \mathbf{x}) is selected by the acquisition function, as described in Section 3.3, a set of S interventional data samples $\{\mathbf{c}_{\mathbf{X}}^{(s)}, y^{(s)}\}_{s=1}^S$ from $p_{\text{do}(\mathbf{X}=\mathbf{x})}(C_{\mathbf{X}}, Y)$ is obtained. For each $V \in C_{\mathbf{X}} \cup Y$, this set is used to form a sample mean estimate $\hat{\mu}_{\text{do}(\mathbf{X}=\mathbf{x})}^V$ of $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^V$, which is treated as a noisy realization of $g_{\mathbf{X}}^V(\mathbf{x})$ with additive Gaussian noise. The tuple $(\mathbf{x}, \hat{\mu}_{\text{do}(\mathbf{X}=\mathbf{x})}^V)$ is then added to $\mathcal{D}_{\mathbf{X}}^I$ and the posterior distribution $\tau(g_{\mathbf{X}}^V | \mathcal{D}_{\mathbf{X}}^I)$ is computed via standard GP updates (Rasmussen & Williams, 2006). Full details are given in Appendix B.

STGP. For the single-task GP, we either assume $m_{\mathbf{X}}^V(\mathbf{x}) = 0$ and radial basis function (RBF) kernel $S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta) = \sigma_f^2 \exp(-\frac{j\|\mathbf{x} - \mathbf{x}^\theta\|^2}{2l^2})$ with (σ_f^2, l) hyper-parameters, or the prior construction proposed in CBO (Aglietti et al., 2020). In the latter case, \mathcal{D}^O is used to obtain estimates of the target

and constraint effects which are then used as prior mean functions (we refer to this variant as STGP⁺).

MTGP. Our first multi-task GP, inspired by the linear coregionalization model of Alvarez et al. (2011), assumes that the target and constraint GPs are linear combinations of shared independent GPs, i.e. $g_{\mathbf{X}}^V(\mathbf{x}) = \sum_{q=1}^Q a_{\mathbf{X},q}^V u_{\mathbf{X},q}(\mathbf{x})$ with $u_{\mathbf{X},q}(\mathbf{x}) \sim \mathcal{GP}(m_{\mathbf{X},q}(\mathbf{x}), S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta))$. In this case, the variance and covariance terms across functions and intervention values are given by $S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta) = \sum_{q=1}^Q (a_{\mathbf{X},q}^V)^2 S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$ and $S_{\mathbf{X}}^{V,W}(\mathbf{x}, \mathbf{x}^\theta) = \sum_{q=1}^Q a_{\mathbf{X},q}^V a_{\mathbf{X},q}^W S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$ respectively. The scalar parameters $a_{\mathbf{X},q}^V$ are learned with a standard type-2 ML approach together with the kernel hyper-parameters. We either consider $m_{\mathbf{X},q}(\mathbf{x}) = 0$ and RBF kernel for each $S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$, or the prior construction using \mathcal{D}^O as in STGP⁺ (we refer to this variant as MTGP⁺).

\mathcal{G} -MTGP. For the setting of no unobserved confounders and under the assumption $V = f_V(\text{pa}(V)) + U_V$ with $p(U_V) = \mathcal{N}(0, \sigma_V^2)$, we propose to model each f_V as an independent GP with an RBF kernel $f_V(\text{pa}(v)) \sim \mathcal{GP}(0, S_{\text{RBF}}^V(\text{pa}(v), \text{pa}(v)^\theta))$, where $\text{pa}(v)$ denotes a value taken by $\text{pa}(V)$, and to fit it using \mathcal{D}^O . Consider the intervention $\text{do}(\mathbf{X} = \mathbf{x})$ and let $\mathbf{U}_{\mathbf{X}}^{\text{an}(V)}$ denote the subset of \mathbf{U} corresponding to the ancestors of V in \mathcal{G} that are not d-separated from V by \mathbf{X} , and similarly for $\mathbf{f}_{\mathbf{X}}^{\text{an}(V)}$. We can write V as an explicit function of $\mathbf{U}_{\mathbf{X}}^{\text{an}(V)}$ and $\mathbf{f}_{\mathbf{X}}^{\text{an}(V)}$ by recursively replacing parents with their functional form in the modified SCM under $\text{do}(\mathbf{X} = \mathbf{x})$. Taking the expectation w.r.t. $\mathbf{U}_{\mathbf{X}}^{\text{an}(V)}$ therefore gives $g_{\mathbf{X}}^V(\mathbf{x})$ as a function of $\mathbf{f}_{\mathbf{X}}^{\text{an}(V)}$.

For example, for the causal graph in Fig. 2(b) with $X = U_X$, $Z = f_Z(X) + U_Z$ and $Y = f_Y(Z) + U_Y$, we can write $Y = f_Y(f_Z(X) + U_Z) + U_Y$, which gives $g_{\mathbf{X}}^Y(\mathbf{x}) = \mathbb{E}_{p(U_Z)}[f_Y(f_Z(\mathbf{x}) + U_Z)]$ and $g_{\mathbf{X}}^Z(\mathbf{x}) = f_Z(\mathbf{x})$. We can then obtain realizations $\{g_{\mathbf{X}}^{V,(s)}\}_{s=1}^{S^O}$ of $g_{\mathbf{X}}^V$ by using samples of $\mathbf{f}_{\mathbf{X}}^{\text{an}(V)}$ and $\mathbf{U}_{\mathbf{X}}^{\text{an}(V)}$ to form an approximation of $S_{\mathbf{X}}^{V,W}(\mathbf{x}_1, \mathbf{x}_2)$ as $\frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{V,(s)}(\mathbf{x}_1)g_{\mathbf{X}}^{W,(s)}(\mathbf{x}_2) - \left(\frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{V,(s)}(\mathbf{x}_1)\right)\left(\frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{W,(s)}(\mathbf{x}_2)\right)$, and similarly for $S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta)$ and $m_{\mathbf{X}}^V(\mathbf{x})$.

3.3. Acquisition Functions

To select interventions accounting for both the target and constraint effects we propose acquisition functions based on those used for constrained BO (Gardner et al., 2014) and noisy BO (Letham et al., 2019). We define the constrained expected improvement (cEI) per unit of intervention cost at an input point \mathbf{x} as $\text{cEI}_{\mathbf{X}}(\mathbf{x}) := \mathbb{E}_{\tau(g_{\mathbf{X}}^Y | \mathcal{D}_{\mathbf{X}}^I)} \left[\frac{\max(0, g^Y - g_{\mathbf{X}}^Y(\mathbf{x}))}{j_{\mathbf{X}}^Y} \mathbb{1}_{g_{\mathbf{X}}^Y(\mathbf{x}) < \lambda_{C_{\mathbf{X}}}} \right]$, where g^Y is the minimum feasible value attained by $g_{\mathbf{X}}^Y$ across interventional dataset \mathcal{D}^I and $\mathbb{1}_{g_{\mathbf{X}}^Y(\mathbf{x}) < \lambda_{C_{\mathbf{X}}}}$ is an indicator variable

equal to one if $g_{\mathbf{X}}^{C_{\mathbf{X}}} \geq \lambda^{C_{\mathbf{X}}}$ and to zero otherwise. The division by $|\mathbf{X}|$ is due to assuming an intervention cost for \mathbf{X} equal to its cardinality. Alternative costs can be considered as long as they are greater than zero.

STGP. When using a single-task GP as surrogate model, we can exploit the factorization $\tau(g_{\mathbf{X}} | \mathcal{D}_{\mathbf{X}}^I) = \tau(g_{\mathbf{X}}^Y | \mathcal{D}_{\mathbf{X}}^I) \prod_{C \in \mathcal{C}_{\mathbf{X}}} \tau(g_{\mathbf{X}}^C | \mathcal{D}_{\mathbf{X}}^I)$ to obtain $cEI_{\mathbf{X}}(\mathbf{x}) = \mathbb{E}_{\tau(g_{\mathbf{X}}^Y | \mathcal{D}_{\mathbf{X}}^I)} [\max(0, g^Y - g_{\mathbf{X}}^Y(\mathbf{x}))] \prod_{C \in \mathcal{C}_{\mathbf{X}}} \mathbb{P}(g_{\mathbf{X}}^C \geq \lambda_C)$ where $\mathbb{P}(g_{\mathbf{X}}^C \geq \lambda_C) = \left[1 - \Phi\left(\frac{\lambda_C - m_{\mathbf{X}}^C(\mathbf{x})}{S_{\mathbf{X}}^C(\mathbf{x}, \mathbf{x})}\right)\right]$ with $\Phi(\cdot)$ denoting the CDF of a standard Gaussian distribution. In the case of noiseless observations of $g_{\mathbf{X}}^{V_k}(\mathbf{x})$, g^Y is known thus we can compute the first term in closed form as $(g^Y - m_{\mathbf{X}}^Y(\mathbf{x})) \Phi\left(\frac{g^Y - m_{\mathbf{X}}^Y(\mathbf{x})}{S_{\mathbf{X}}^Y(\mathbf{x}, \mathbf{x})}\right) + S_{\mathbf{X}}^Y(\mathbf{x}, \mathbf{x}) \phi\left(\frac{g^Y - m_{\mathbf{X}}^Y(\mathbf{x})}{S_{\mathbf{X}}^Y(\mathbf{x}, \mathbf{x})}\right)$ with $\phi(\cdot)$ denoting the PDF of a standard Gaussian distribution. In the case of noisy observations of $g_{\mathbf{X}}^{V_k}(\mathbf{x})$, g^Y is unknown as we observe noisy values of the target effects. We can get an estimate of g^Y using samples from $\tau(g_{\mathbf{X}} | \mathcal{D}_{\mathbf{X}}^I)$ for all $\mathbf{X} \in nM_{\mathcal{C}[Y, G]}$, and use the estimate to compute $cEI_{\mathbf{X}}(\mathbf{x})$ in closed form using the terms derived above. For every $\mathbf{X} \in nM_{\mathcal{C}[Y, G]}$ the values of $cEI_{\mathbf{X}}(\mathbf{x})$ obtained with different g^Y are then averaged to integrate out the uncertainty on the optimal feasible value observed across all intervention sets.

MTGP. When using a multi-task model, $cEI_{\mathbf{X}}(\mathbf{x})$ cannot be computed in closed form as $\tau(g_{\mathbf{X}} | \mathcal{D}_{\mathbf{X}}^I)$ does not factorize. We thus compute it via Monte Carlo integration with a similar procedure as for the noisy STGP setting.

3.4. Computational Aspects

The computational cost of cCBO is dominated by the algebraic operations needed to compute the posterior parameters for the GP models of the target and constraints effects. Let N_I denote the largest among the cardinalities of the interventional datasets collected for the sets in $nM_{\mathcal{C}[Y, G]}$, i.e. $N_I = \max_{\mathbf{X} \in nM_{\mathcal{C}[Y, G]}} |\mathcal{D}_{\mathbf{X}}^I|$, and let M denote the highest among the number of target and constraint effects for the sets in $nM_{\mathcal{C}[Y, G]}$, i.e. $M = \max_{\mathbf{X} \in nM_{\mathcal{C}[Y, G]}} 1 + |C_{\mathbf{X}}|$. As the GPs corresponding to different intervention sets are updated independently, the computational cost scales as $O(N_I^3)$ when using a single-task model and as $O(M^3 N_I^3)$ when using a multi-task model. Independent GPs updates also imply linear scaling of the computational complexity with respect to the cardinality of $nM_{\mathcal{C}[Y, G]}$. Therefore, a larger \mathcal{G} might induce a higher number of sets to explore, but does not induce higher computational complexity.

In terms of convergence to the true global optimum, cCBO inherits the properties of BO algorithms. While any alternative acquisition function for constrained BO can be used within cCBO, in this work we resort to a constrained expected improvement function due to its computationally

tractability. The expected improvement acquisition function was shown to have strong theoretical guarantees (Vazquez & Bect, 2010; Bull, 2011) while performing well in practice (Snoek et al., 2012). However, performance guarantees have yet to be established for the constrained version.

3.5. Related Work

cCBO is related to the vast literature on constrained BO methods, which find feasible solutions either by using an acquisition function that accounts for the probability of feasibility (Gardner et al., 2014; Gelbart et al., 2014; Griffiths & Hernández-Lobato, 2017; Hernández-Lobato et al., 2014; Keane et al., 2008; Schonlau et al., 1998; Sóbester et al., 2014; Tran et al., 2019), or by transforming the constrained problem into an unconstrained one (Ariafar et al., 2019; Picheny et al., 2016), or by exploiting trust regions (Eriksson & Poloczek, 2021). While these works disregard the causal aspect of the optimization and model the unknown functions independently, multi-task surrogate models have been considered by multi-objective BO methods (Dai et al., 2020; Feliot et al., 2017; Hakhamaneshi et al., 2021; Mathern et al., 2021; Swersky et al., 2013) or, more recently, in the context of safe BO (Bergmann & Graichen, 2020; Berkenkamp et al., 2016; 2021; Kirschner et al., 2019; Sui et al., 2015; 2018). cCBO is also related to the works combining causality with decision-making frameworks which have mainly focused on finding optimal interventions using observational data (Atan et al., 2018; Håkansson et al., 2020; Zhang et al., 2012) or on designing interventions for causal discovery (Tigas et al., 2022). The idea of identifying optimal interventions through sequential experimentation has been explored in causal bandits (Lattimore et al., 2016), causal reinforcement learning (Zhang, 2020) and, more recently, in CBO (Aglietti et al., 2020) and model-based causal BO (MCBO, Sussex et al. (2023)). All these works tackle unconstrained settings and disregard the effects that interventions optimizing a target variable might have on the constrained variables.

4. Experiments

We evaluate cCBO with surrogate models STGP, STGP⁺, MTGP, MTGP⁺, and \mathcal{G} -MTGP on the causal graphs of Fig. 2(b) (SYNTHETIC-1), Fig. 2(a) (SYNTHETIC-2), Fig. 1(b) (HEALTH), and Fig. 1(a) (PROTEIN-SIGNALING). We assume an initial \mathcal{D}^I that includes one point per intervention set, and consider different settings with respect to the SCM characteristics, null-feasibility, $|nM_{\mathcal{C}[Y, G]}|$, and $N_O = |\mathcal{D}^O|$. See Appendix C for full experimental details⁶.

To the best of our knowledge, there are no other constrained

⁶Code for reproducing the experiments is available at <https://github.com/deepmind/ccbo>.

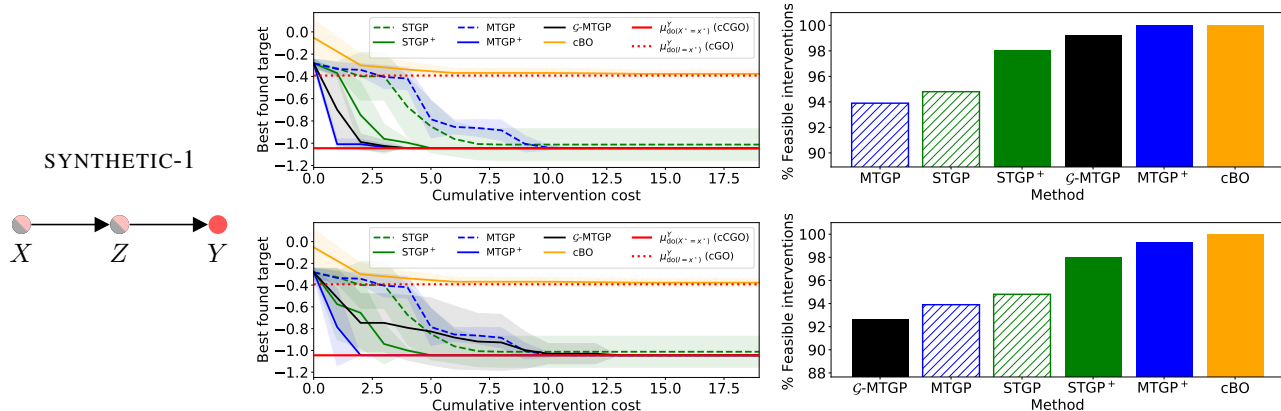


Figure 3: SYNTHETIC-1 with $N_O = 500$ (top row) and $N_O = 100$ (bottom row) and $\lambda^Z = 2$. *Left*: Causal graphs. *Center*: Convergence to the cCGO (solid red line) and CGO (dotted red line) optima. Lines give average results across different initialization of \mathcal{D}^I . Shaded areas represent \pm standard deviation. *Right*: Average percentage of feasible interventions collected over trials.

methods in the literature that exploit causal structure. Therefore, we compare to the closest method aiming at solving the non-causal *constrained global optimization* (cGO) problem, namely the constrained BO algorithm (cBO) proposed by Gardner et al. (2014). cBO intervenes on all variables in \mathbf{I} simultaneously, models the target and constraints effects via independent GPs, and selects interventions by maximizing the constrained EI acquisition function.

For each model, we show the average convergence to the optimal target effect (\pm standard deviations given by shaded areas), and the mean percentage of feasible interventions collected over trials across 20 initialization of \mathcal{D}^I and for varying values of N_O . The combination of these two metrics allows us to understand how each method balances convergence speed and feasibility. To gain more insights into this balance, we also include in the analysis CBO, an algorithm that randomly picks interventions (RANDOM) and, for SYNTHETIC-1 and SYNTHETIC-2, MCBO (see Appendix C). Notice that, as CBO and MCBO solve the CGO problem, a comparison with them is only relevant for cases in which the CGO and cCGO optima are equal.

4.1. Synthetic Causal Graphs

SYNTHETIC-1. For the SYNTHETIC-1 graph with $\lambda^X = 1$, $\lambda^Z = 2, 10$, and $N_O = 500, 100, 10$ observational data samples from the SCM in Appendix C.1, we obtain $nM_{C|Y,G} = \{\{X\}, \{Z\}\}$.

With $N_O = 500$ and $\lambda^Z = 2$, using a surrogate model that exploits \mathcal{D}^O when constructing the prior GP parameters for functions in $g_X(\mathbf{x})$, as in STGP⁺, MTGP⁺ and \mathcal{G} -MTGP, leads to faster average convergence (Fig. 3, top row). The convergence speed is further improved when capturing the covariance structure among the target and constraint effects

as done by MTGP⁺ and \mathcal{G} -MTGP. High convergence speed for STGP⁺, MTGP⁺ and \mathcal{G} -MTGP is associated with a higher percentage of feasible interventions collected over trials compared to the other methods. Similar results are obtained with $N_O = 100$ and $\lambda^Z = 2$ (Fig. 3, bottom row) and with $N_O = 10$ and $\lambda^Z = 2$ (Fig. 9 in Appendix C.1). In these cases, the lower cardinality of \mathcal{D}^O leads to a less accurate estimation of the effects which affects the prior mean functions for STGP⁺, MTGP⁺ and \mathcal{G} -MTGP, and kernel function for \mathcal{G} -MTGP. In turns, this leads to a lower number of feasible interventions and convergence speed, particularly for \mathcal{G} -MTGP. As a consequence, MTGP⁺ outperforms all other models in this setting. By intervening on all constrained variables ($\mathbf{I} = \mathbf{C}$) with intervention domains that are in accordance with the threshold values, cBO only collects feasible interventions. However, it converges to the higher cGO optimum, as causal structure is disregarded.

Interestingly, the results obtained with $N_O = 500$ and $\lambda^Z = 10$ (Fig. 10 (top row) in Appendix C.1) for which the cCGO and CGO optima are equal, making CBO, MCBO and cCBO comparable, show that both CBO and MCBO converge at a slower pace compared to MTGP⁺ and \mathcal{G} -MTGP while collecting a lower number of feasible interventions. Indeed, CBO and MCBO disregard the values taken by the constrained variables. Finally, when $\lambda^Z = 10$, \mathcal{G} -MTGP and MTGP⁺ show fast convergence performance and high percentage of feasible interventions collected over trials. This is even when $N_O = 100$ (Fig. 11 in Appendix C.1).

SYNTHETIC-2. For the SYNTHETIC-2 graph with $\lambda^C = 10, \forall C \in \mathbf{C}$, and $N_O = 500, 100, 10$ observational data samples from the SCM in Appendix C.2, we obtain null-feasibility $\forall C \in \mathbf{C}$ and thus $nM_{C|Y,G} = \mathcal{P}_I \setminus \{A, D, E\}$. As in SYNTHETIC-1, when $N_O = 500$ using a prior GP

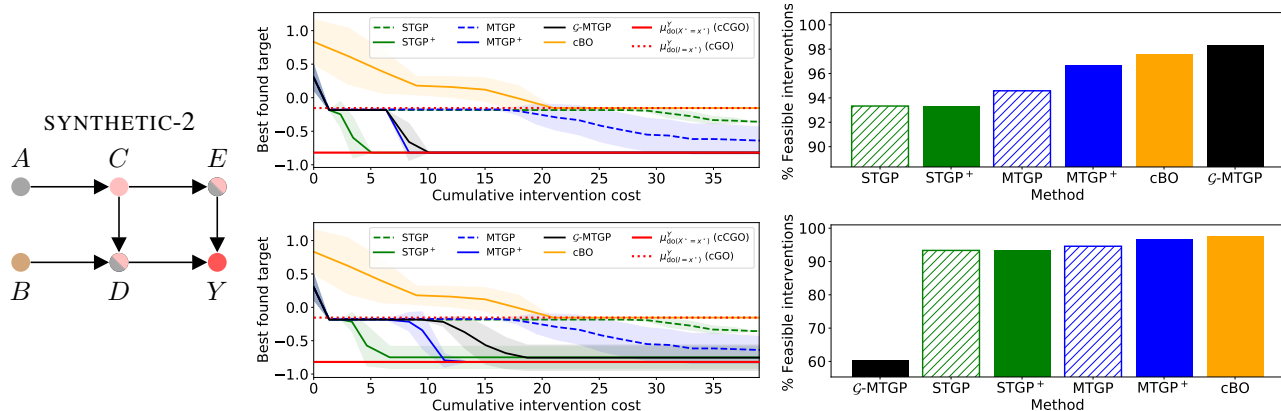


Figure 4: SYNTHETIC-2 with $N_O = 500$ (top row) and $N_O = 100$ (bottom row). *Left*: Causal graphs. *Center*: Convergence to the cCGO (solid red line) and cGO (dotted red line) optima. Lines give average results across different initialization of \mathcal{D}^I . Shaded areas represents \pm standard deviation. *Right*: Average percentage of feasible interventions collected over trials.

construction that exploits \mathcal{D}^O leads to faster convergence (Fig. 4, top row). In particular, the accurate estimation of the target and constraint effects with \mathcal{D}^O , which are used as prior mean functions, favours the prior formulation of STGP⁺ which identifies the optimal interventions after ~ 5 trials. However, by disregarding the correlation among the effects, STGP⁺ collects a higher percentage of infeasible interventions compared to MTGP⁺ and G-MTGP. In this setting, accurate estimation of the functions in the SCM translates to better uncertainty estimation around the effects given by the kernel function of G-MTGP. Therefore G-MTGP successfully trades off improvement and feasibility showing a similar convergence to MTGP⁺ but the highest percentage of feasible interventions collected.

With $N_O = 100$ (Fig. 4, bottom row) and $N_O = 10$ (Fig. 12 of Appendix C.2), the estimation of the target and constraint effects and the functions in the SCM from \mathcal{D}^O deteriorates leading MTGP⁺, which learns the correlations directly from \mathcal{D}^I , to outperform all other methods.

As in SYNTHETIC-1, intervening on all variables in \mathbf{I} simultaneously blocks the propagation of causal effects in the graph thus leading cBO to achieve a sub-optimal solution compared to cCBO and a lower number of feasible interventions compared to G-MTGP.

4.2. Real-world Causal Graphs

HEALTH. For the HEALTH graph, we use the SCM in Ferro et al. (2015) (given in Appendix C.3). We consider the constraint that BMI must be lower than 25, which is considered the maximum healthy level. When both $N_O = 100$ and $N_O = 10$, BMI is not null-feasible thus all sets in $nM_{\mathcal{C}[Y,G]} = \{\{CI\}, \{\text{Aspirin}, CI\}, \{\text{Statin}, CI\}, \{\text{Aspirin}, \text{Statin}, CI\}\}$ include CI. In this setting, \mathbf{I} is the optimal intervention set,

and therefore the cCGO and cGO optima coincide.

When $N_O = 100$, cBO is significantly slower than G-MTGP and, by disregarding causal information, collects a lower percentage of feasible interventions over trials (Fig. 5, top row). Even though the effects are simple linear functions, the stochasticity in \mathcal{D}^O determined by $p(\mathbf{U})$ is high (Fig. 15 of Appendix C.3) compared to the previous examples and obscures their estimation. This leads to a less accurate prior mean function for G-MTGP, STGP⁺ and MTGP⁺ which translates to a lower convergence speed for the latter two models. Despite the less accurate prior mean functions, G-MTGP achieves the fastest convergence and the highest percentage of feasible interventions (100%) by capturing the correlation among target and constraint effects induced by the SCM thus properly quantifying uncertainty around them. Collecting 100% of feasible interventions is particularly important in this setting as infeasible interventions might negatively affect patients' health status. While G-MTGP performs well in settings with $N_O = 100$, it does not reach convergence when $N_O = 10$ (Fig. 5, bottom row). Indeed, the prior mean and kernel functions estimation from \mathcal{D}^O deteriorates when the size of the observational dataset is very small leading to an incorrect uncertainty quantification around the effects and preventing the exploration of the interventional space. This is also observed for STGP⁺ and MTGP⁺ as their prior mean function is affected by the incorrect estimation of the SCM functions and thus of the target and constraint effects. This leads cBO to outperform all other methods in this setting.

PROTEIN-SIGNALING. For the PROTEIN-SIGNALING graph, we use the observational dataset from Sachs et al. (2005), to construct an SCM (see Appendix C.4). Given $N_O = 100$, 10 observational data samples from the SCM, all constrained variables are null-feasible, giving $nM_{\mathcal{C}[Y,G]} =$

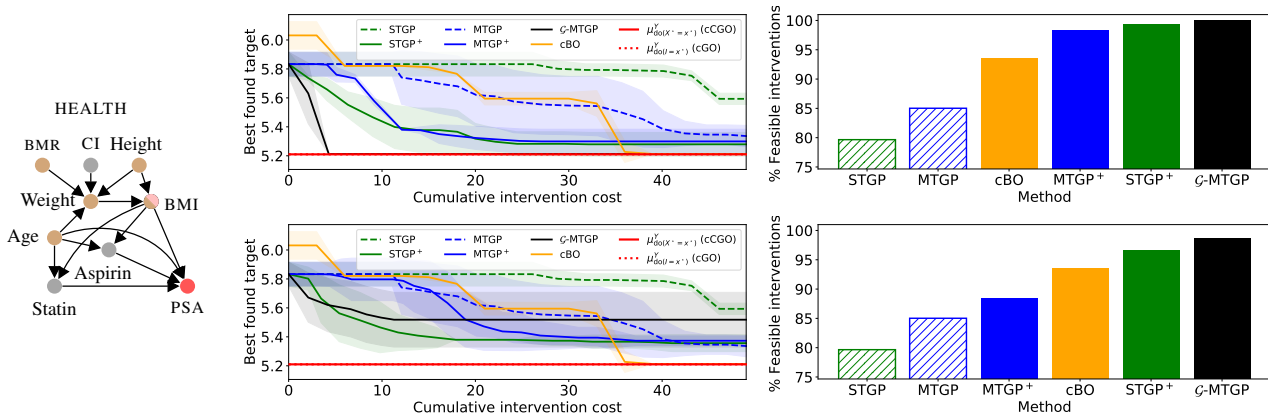


Figure 5: HEALTH with $N_O = 100$ (top row) and $N_O = 10$ (bottom row). *Left*: Causal graphs. *Center*: Convergence to the cCGO (solid red line) and cGO (dotted red line coinciding with the solid red line in this experiment) optima. Lines give average results across different initialization of \mathcal{D}^I . Shaded areas represents \pm standard deviation. *Right*: Average percentage of feasible interventions collected over trials.

$\{\{PKC\}, \{PKA\}, \{Mek\}, \{PKC, PKA\}, \{PKC, Mek\}, \{PKA, Mek\}\}$. All methods converge to the same optimum (Fig. 6). Indeed, the target effect achieved by intervening on I is equal to the one achieved by intervening on the optimal intervention set $\{PKA, Mek\}$ (this is the result of Erk being independent on PKC and Akt given Mek and PKA). In addition, in this example $C_I = \emptyset$ thus by intervening on I with interventional domains that are in accordance with the threshold values cBO achieves a 100% average feasibility.

Overall cBO performs comparably to \mathcal{G} -MTGP and $MTGP^+$ (Fig. 6, top row) when $N_O = 100$, while cBO is faster but selects more infeasible interventions (Fig. 18 in Appendix C.4). Despite achieving slower convergence than single task models, $MTGP^+$ and \mathcal{G} -MTGP select a higher percentage of feasible interventions (99.13%) with \mathcal{G} -MTGP converging slightly faster than $MTGP^+$. The feasibility aspect is again particularly important as every non feasible intervention results in the inhibition of PKA, which can impede healthy functions of the cell. Hence, a method that yields greater than 99% feasible interventions at the cost of (~ 10) additional interventions ($MTGP^+$ and \mathcal{G} -MTGP) is deemed preferable compared to methods that explore the intervention space more aggressively but yield a lower number of feasible interventions (cBO, STGP and $STGP^+$).

As in the other experiments, the convergence performance of \mathcal{G} -MTGP and $MTGP^+$ slightly deteriorates when $N_O = 10$ (Fig. 6, bottom row) due to an incorrect uncertainty quantification around the effects which prevents the exploration of the interventional space. As in the results for $N_O = 100$, notice how cBO achieves a 100% average feasibility as all constrained variables are intervened ($C \subset I$, $C_I = \emptyset$) with interventional ranges that are in accordance with the threshold values. The algorithm quickly identify the cGO

optimum, which is in this case equal to the cCGO optimum.

5. Discussion

In this paper, we introduced the ccBO approach for identifying interventions optimizing a target effect under constraints. We proposed different GP surrogate models for the target and constraint effects that leverage observational data \mathcal{D}^O and the structure of the SCM. Our results show that incorporating \mathcal{D}^O in the GP prior construction leads to faster identification of optimal interventions and higher percentage of feasible interventions selected. They also show that further performance improvement can be obtained using multi-task GP models which capture the correlation among target and constraint effects. Accounting for the SCM structure in the covariance matrix, as done by \mathcal{G} -MTGP, is especially beneficial with more complex correlation structures or with a high number of effects. We found \mathcal{G} -MTGP to successfully trade off improvement and feasibility, achieving a fast convergence while collecting the highest percentage of feasible interventions. However, as \mathcal{G} -MTGP exploits the fitted SCM functions in the computation of the prior parameters, when these functions cannot be learned accurately $MTGP^+$ is preferable as it can learn the correlation directly from \mathcal{D}^I .

Our search space reduction procedure can be thought as placing an intervention cost for each $X \in \mathcal{P}_I$ that is equal to its cardinality. This procedure could be modified using different cost structures or augmented with budget constraints. For example, one could exclude from \mathcal{P}_I the intervention sets non satisfying a budget and then proceed by excluding sets according to the proposed procedure.

ccBO requires knowledge of the true causal graph underlying the system of interest, an assumption that might not be

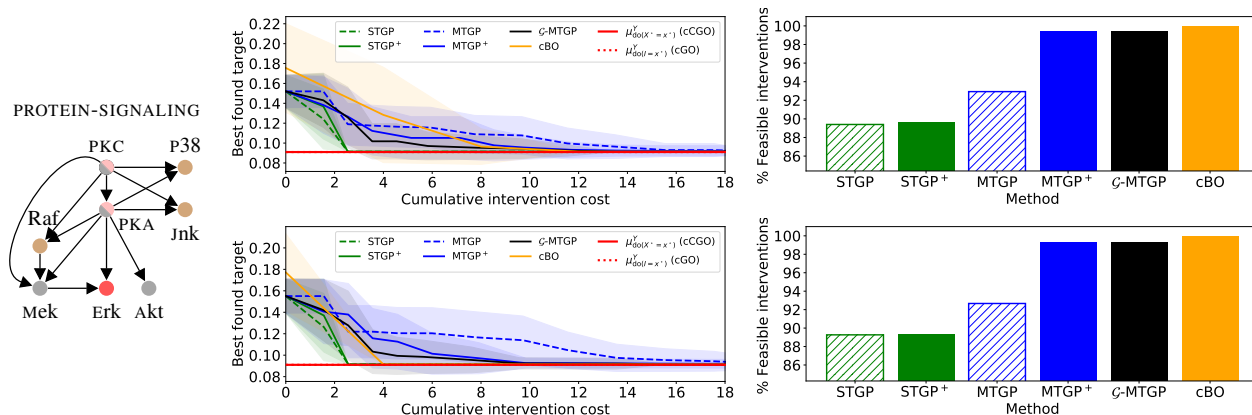


Figure 6: PROTEIN-SIGNALING with $N_O = 100$ and $N_O = 10$ (bottom row). *Left*: Causal graphs. *Center*: Convergence to the cCGO (solid red line) and cGO (dotted red line coinciding with the solid red line in this experiment) optima. Lines give average results across different initialization of \mathcal{D}^I . Shaded areas represents \pm standard deviation. *Right*: Average percentage of feasible interventions collected over trials.

satisfied in practice. Using cCBO with an incorrect graph \mathcal{G}^0 could lead to inaccurate GP prior parameters and invalid search space reduction. Indeed, there is no guarantee that $nM_{\mathcal{C}|Y,\mathcal{G}^0}$ would contain the optimal intervention set when the independence and the ancestor-type relationships encoded in \mathcal{G}^0 differ from those in \mathcal{G} . Extending the methodology to deal with settings characterized by misspecified or unknown \mathcal{G} , similarly e.g. to Branchini et al. (2023), represents an important future direction.

From a computational perspective, cCBO might suffer from poor scalability when $|C|$ or the number of collected interventional data samples are high. In the latter case, sparse GP methods e.g. inducing inputs (Titsias, 2009), could be used to significantly reduce the computational complexity of both single-task and multi-task GPs. Alternatively, one could consider sharing information across the surrogate models of different sets in $nM_{\mathcal{C}|Y,\mathcal{G}}$ in order to reduce the total number of interventions. Evaluating cCBO on real-world settings where $|C|$ is high would require simulators characterized by a high number of variables. In terms of convergence properties, while the constrained EI does not currently provide theoretical guarantees, one could extend cCBO to use a constrained acquisition function achieving asymptotic convergence, e.g. GP-UCB, (Srinivas et al., 2009; Lu & Paulson, 2022) or a non-myopic acquisition function to avoid exploration issues (Lam & Willcox, 2017).

Another future direction for safe-critical applications would be to extend the current framework to deal with safety constraints in order to guarantee that the number of feasible interventions collected never falls below a critical value. Finally, while cCBO focuses on hard interventions in which variables are set to specific values, an interesting extension would be to consider more general soft-interventions.

Acknowledgements

The authors would like to thank Alexis Bellot for his valuable feedback.

References

Aglietti, V., Lu, X., Paleyes, A., and González, J. Causal Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 3155–3164, 2020.

Alvarez, M. A., Rosasco, L., and Lawrence, N. D. Kernels for vector-valued functions: A review. *arXiv preprint arXiv:1106.6251*, 2011.

Ariaifar, S., Coll-Font, J., and Brooks, Dana Hand Dy, J. G. ADMMBO: Bayesian optimization with unknown constraints using ADMM. *Journal of Machine Learning Research*, 20(123):1–26, 2019.

Atan, O., Jordon, J., and Van der Schaar, M. Deep-treat: Learning optimal personalized treatments from observational data using neural networks. In *AAAI Conference on Artificial Intelligence*, pp. 2071–2078, 2018.

Bergmann, D. and Graichen, K. Safe Bayesian optimization under unknown constraints. In *IEEE Conference on Decision and Control*, pp. 3592–3597, 2020.

Berkenkamp, F., Schoellig, A. P., and Krause, A. Safe controller optimization for quadrotors with Gaussian processes. In *IEEE International Conference on Robotics and Automation*, pp. 491–496, 2016.

Berkenkamp, F., Krause, A., and Schoellig, A. P. Bayesian optimization with safety constraints: Safe and automatic

- parameter tuning in robotics. *Machine Learning*, pp. 1–35, 2021.
- Branchini, N., Aglietti, V., Dhir, N., and Damoulas, T. Causal entropy optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 8586–8605, 2023.
- Bull, A. D. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(10), 2011.
- Curi, S., Berkenkamp, F., and Krause, A. Efficient model-based reinforcement learning through optimistic policy search and planning. *Advances in Neural Information Processing Systems*, 33:14156–14170, 2020.
- Dai, S., Song, J., and Yue, Y. Multi-task Bayesian optimization via Gaussian process upper confidence bound. In *ICML 2020 Workshop on Real World Experiment Design and Active Learning*, 2020.
- Eriksson, D. and Poloczek, M. Scalable constrained Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 730–738, 2021.
- Feliot, P., Bect, J., and Vazquez, E. A Bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67(1-2):97–133, 2017.
- Ferro, A., Pina, F., Severo, M., Dias, P., Botelho, F., and Lunet, N. Use of statins and serum levels of prostate specific antigen. *Acta Urológica Portuguesa*, 32(2):71–77, 2015.
- Frémin, C. and Meloche, S. From basic research to clinical development of MEK1/2 inhibitors for cancer therapy. *Journal of Hematology & Oncology*, 3(1):1–11, 2010.
- Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. Bayesian optimization with inequality constraints. In *International Conference on Machine Learning*, pp. 937–945, 2014.
- Gelbart, M. A., Snoek, J., and Adams, R. P. Bayesian optimization with unknown constraints. *arXiv preprint arXiv:1403.5607*, 2014.
- Griffiths, R.-R. and Hernández-Lobato, J. M. Constrained Bayesian optimization for automatic chemical design. *arXiv preprint arXiv:1709.05501*, 2017.
- Håkansson, S., Lindblom, V., Gottesman, O., and Johansson, F. D. Learning to search efficiently for causally near-optimal treatments. In *Advances in Neural Information Processing Systems*, volume 33, pp. 1333–1344, 2020.
- Hakhamaneshi, K., Abbeel, P., Stojanovic, V., and Grover, A. JUMBO: Scalable multi-task Bayesian optimization using offline data. *arXiv preprint arXiv:2106.00942*, 2021.
- Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. *arXiv preprint arXiv:1406.2541*, 2014.
- Keane, A., Forrester, A., and Sobester, A. *Engineering design via surrogate modelling: a practical guide*. American Institute of Aeronautics and Astronautics, Inc., 2008.
- Kirschner, J., Mutny, M., Hiller, N., Ischebeck, R., and Krause, A. Adaptive and safe Bayesian optimization in high dimensions via one-dimensional subspaces. In *International Conference on Machine Learning*, pp. 3429–3438, 2019.
- Lam, R. and Willcox, K. Lookahead bayesian optimization with inequality constraints. In *Advances in Neural Information Processing Systems*, 2017.
- Lattimore, F., Lattimore, T., and Reid, M. D. Causal bandits: Learning good interventions via causal inference. In *Advances in Neural Information Processing Systems*, 2016.
- Lee, S. and Bareinboim, E. Structural causal bandits: where to intervene? In *Advances in Neural Information Processing Systems*, pp. 2568–2578, 2018.
- Letham, B., Karrer, B., Ottoni, G., and Bakshy, E. Constrained Bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 2019.
- Lu, C. and Paulson, J. A. No-regret bayesian optimization with unknown equality and inequality constraints using exact penalty functions. *IFAC-PapersOnLine*, 55(7):895–902, 2022.
- Mathern, A., Steinholtz, O. S., Sjöberg, A., Önnheim, M., Ek, K., Rempling, R., Gustavsson, E., and Jirstrand, M. Multi-objective constrained Bayesian optimization for structural design. *Structural and Multidisciplinary Optimization*, 63(2):689–701, 2021.
- Pearl, J. *Causality: Models, Reasoning and Inference*, volume 29. Springer, 2000.
- Pearl, J., Glymour, M., and Jewell, N. P. *Causal Inference in Statistics: A Primer*. John Wiley & Sons, 2016.
- Picheny, V., Gramacy, R. B., Wild, S., and Digabel, S. L. Bayesian optimization under mixed constraints with a slack-variable augmented Lagrangian. In *Advances in Neural Information Processing Systems*, pp. 1443–1451, 2016.

- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D. A., and Nolan, G. P. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721): 523–529, 2005.
- Schonlau, M., Welch, W. J., and Jones, D. R. Global versus local search in constrained optimization of computer models. *Lecture Notes-Monograph Series*, pp. 11–25, 1998.
- Snoek, J., Larochelle, H., and Adams, R. P. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, volume 25, 2012.
- Sóbester, A., Forrester, A. I., Toal, D. J., Tresidder, E., and Tucker, S. Engineering design applications of surrogate-assisted optimization techniques. *Optimization and Engineering*, 15(1):243–265, 2014.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Sui, Y., Gotovos, A., Burdick, J., and Krause, A. Safe exploration for optimization with Gaussian processes. In *International Conference on Machine Learning*, pp. 997–1005, 2015.
- Sui, Y., Zhuang, V., Burdick, J., and Yue, Y. Stagewise safe Bayesian optimization with Gaussian processes. In *International Conference on Machine Learning*, pp. 4781–4789, 2018.
- Sussex, S., Makarova, A., and Krause, A. Model-based causal Bayesian optimization. In *International Conference on Learning Representations*, 2023.
- Swersky, K., Snoek, J., and Adams, R. P. Multi-task Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2013.
- Tigas, P., Annadani, Y., Jesson, A., Schölkopf, B., Gal, Y., and Bauer, S. Interventions, where and how? Experimental design for causal models at scale. In *Advances in Neural Information Processing Systems*, 2022.
- Titsias, M. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pp. 567–574, 2009.
- Tran, A., Sun, J., Furlan, J. M., Pagalthivarthi, K. V., Visintainer, R. J., and Wang, Y. pBO-2GP-3B: A batch parallel known/unknown constrained Bayesian optimization with feasibility classification and its applications in computational fluid dynamics. *Computer Methods in Applied Mechanics and Engineering*, 347:827–852, 2019.
- Vazquez, E. and Bect, J. Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and Inference*, 140(11):3088–3095, 2010.
- Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012.
- Zhang, J. Designing optimal dynamic treatment regimes: A causal reinforcement learning approach. In *International Conference on Machine Learning*, pp. 11012–11022, 2020.

A. Reducing the Search Space

Recall the definition of cMIS relative to $(C \cup Y, \mathcal{G})$.

Definition 3.1 (cMIS). A set $X \subseteq I$ is said to be a cMIS relative to $(C \cup Y, \mathcal{G})$ if there is no $X^0 \subset X$ with $C_X = C_{X^0}$ such that $\mu_{\text{do}(X=x)}^W = \mu_{\text{do}(X^0=x^0)}^W$, where x^0 indicates the subset of x corresponding to variables X^0 , $\forall x \in D(X)$, $\forall W \in C_X \cup Y$, and \forall SCM compatible with \mathcal{G} .

Also recall that $M_{C|Y,G}$ is the subset of \mathcal{P}_I whose elements are cMISs relative to $(C \cup Y, \mathcal{G})$, that $\text{an}(W, \mathcal{G})$ denotes the set of ancestors of W in \mathcal{G} , that $\text{an}(W, \mathcal{G}) := \cup_{W \succ W'} \text{an}(W', \mathcal{G})$, and that $\mathcal{G}_{\overline{X}}$ is the graph with all incoming edges onto all elements of X removed.

Proposition 3.3 (Characterization of cMIS). $X \subseteq I$ is a cMIS relative to $(C \cup Y, \mathcal{G}) \iff X \subseteq \text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)$.

Proof.

(\implies) We first prove that if $X \subseteq \text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)$ then X is a cMIS relative to $(C \cup Y, \mathcal{G})$.

$X \subseteq \text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)$ implies that, for any $X^0 \subset X$, the set $Q = X \setminus X^0$ is also a subset of $\text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)$. Therefore (i) $Q \subset \text{an}(C \cup Y, \mathcal{G}_{\overline{X}})$, or (ii) $Q \subset C \cap X$, or (iii) Q includes both variables in $\text{an}(C \cup Y, \mathcal{G}_{\overline{X}})$ and $C \cap X$. In case (ii), $Q \subseteq C$ and, as $X = Q \cup X^0$, we obtain $C \cap X = C \cap (Q \cup X^0) = (C \cap Q) \cup (C \cap X^0) = Q \cup (C \cap X^0) \supset C \cap X^0$ implying $C_X \subset C_{X^0}$ which violates the requirement $C_X = C_{X^0}$ of Definition 3.1. In cases (i) and (iii), Q includes at least one variable, say Q , with at least one directed path, say from Q to $W \in C \cup Y$, that is not passing through X^0 (as the incoming edges into X^0 are removed in $\mathcal{G}_{\overline{X}}$). Consider a SCM with $V = \sum_{i=1}^{\text{pa}(V)} \text{pa}(V)_i + U_V$, where $\text{pa}(V)_i$ is the i -th parent, and $U_V \sim \mathcal{N}(0, 1)$. In this SCM, $\mu_{\text{do}(X^0=x^0)}^W = ax^0 + b\mu_{\text{do}(X^0=x^0)}^Q$ where a and b are two positive constants, while $\mu_{\text{do}(Q=q, X^0=x^0)}^W = ax^0 + bq$, so that taking $q = \mu_{\text{do}(X^0=x^0)}^Q + 1$ gives $\mu_{\text{do}(Q=q, X^0=x^0)}^W > \mu_{\text{do}(X^0=x^0)}^W$. Therefore, for any $X^0 \subset X$ we can construct a SCM such that $\mu_{\text{do}(X=x)}^W > \mu_{\text{do}(X^0=x^0)}^W$ for at least one $W \in C \cup Y$. As there is no $X^0 \subset X$ such that $\mu_{\text{do}(X=x)}^W = \mu_{\text{do}(X^0=x^0)}^W \forall W \in C \cup Y$ and \forall SCM compatible with \mathcal{G} , X satisfies the requirements of Definition 3.1 for being a cMIS relative to $(C \cup Y, \mathcal{G})$.

(\impliedby) We now prove that if X is a cMIS relative to $(C \cup Y, \mathcal{G})$ then $X \subseteq \text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)$.

If X were not a subset of $\text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)$, then we could define the non empty set $Q =$

$X \setminus (X \cap (\text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)))$ and the set $X^0 = X \setminus Q \subset X$ such that (1) $C_X = C_{X^0}$ and (2) all effects on variables in $C \cup Y$ for X^0 and X are equal, contradicting X being a cMIS relative to $(C \cup Y, \mathcal{G})$. Condition (1) would hold as $X \cap (\text{an}(C \cup Y, \mathcal{G}_{\overline{X}}) \cup (C \cap X)) = (X \cap \text{an}(C \cup Y, \mathcal{G}_{\overline{X}})) \cup (X \cap C)$ implies that we can express Q as $Q = X \setminus (B \cup (X \cap C))$ for $B = X \cap \text{an}(C \cup Y, \mathcal{G}_{\overline{X}})$, showing that Q does not include variables in X that are in C , and therefore that $C_X = C_{X^0}$. Condition (2) would follow from the fact that the non-overlapping sets Q and $C \cup Y$ satisfy $(C \cup Y) \perp_{\mathcal{G}_{\overline{X}}} Q \mid X^0$ as: (i) there could not be causal paths from any variable in Q to any variable in $(C \cup Y)$ in $\mathcal{G}_{\overline{X}}$ (as Q does not contain the ancestors of $(C \cup Y)$ by definition), (ii) non-causal frontdoor paths would be closed as the incoming edges into the colliders that might be included in X^0 would be removed in $\mathcal{G}_{\overline{X}}$, and (iii) all backdoor paths would be removed in $\mathcal{G}_{\overline{X}}$. Therefore, by the Rule 3 of do-calculus, we would have $\mu_{\text{do}(Q=q, X^0=x^0)}^W = \mu_{\text{do}(X^0=x^0)}^W \forall W \in C \cup Y$.

Theorem 3.2 (Sufficiency of $M_{C|Y,G}$). $M_{C|Y,G}$ contains a solution of the cCGO problem (if a solution exists), \forall SCM compatible with \mathcal{G} .

Proof. Consider a set X that satisfies Eq. (1). By Definition 3.1, if $X \notin M_{C|Y,G}$ there exists a set $X^0 \subset X$ with $C_X = C_{X^0}$ such that $\mu_{\text{do}(X=x)}^Y = \mu_{\text{do}(X^0=x^0)}^Y$ and $\mu_{\text{do}(X=x)}^C = \mu_{\text{do}(X^0=x^0)}^C$, $\forall C \in C_{X^0}$ and therefore X^0 also satisfies Eq. (1). If $X^0 \notin M_{C|Y,G}$, we can apply a similar reasoning, and proceed until we find a set for which there is no subset that gives the same effects, which therefore is in $M_{C|Y,G}$.

Lemma A.1. For any disjoint set of variables W_1, W_2, W_3 , if $W_1 \cap \text{an}(W_2, \mathcal{G}_{\overline{W_1, W_3}}) = \emptyset$ then $W_2 \perp_{\mathcal{G}_{\overline{W_1, W_3}}} W_1 \mid W_3$.

Proof. $W_1 \cap \text{an}(W_2, \mathcal{G}_{\overline{W_1, W_3}}) = \emptyset$ implies that there are no directed paths from (any element of) W_1 to (any element of) W_2 in $\mathcal{G}_{\overline{W_1, W_3}}$. In addition, all backdoor paths from W_1 to W_2 in \mathcal{G} are removed in $\mathcal{G}_{\overline{W_1, W_3}}$. Therefore, all paths from W_1 to W_2 in $\mathcal{G}_{\overline{W_1, W_3}}$ are non-directed frontdoor paths. As W_3 cannot be a collider on paths from W_1 to W_2 in $\mathcal{G}_{\overline{W_1, W_3}}$ (W_3 does not have incoming edges in $\mathcal{G}_{\overline{W_1, W_3}}$), these paths cannot be opened by conditioning on W_3 . In conclusion, all paths from W_1 to W_2 are closed in $\mathcal{G}_{\overline{W_1, W_3}}$ given W_3 , therefore $W_2 \perp_{\mathcal{G}_{\overline{W_1, W_3}}} W_1 \mid W_3$.

Let $nM_{C|Y,G}(X, C) := \{X^0 \in M_{C|Y,G} : X^0 \supset X, X^0 = X \cup X_1 \text{ with } X_1 \cap \text{an}(C_X \setminus C \cup Y, \mathcal{G}_{\overline{X^0}}) = \emptyset \text{ and } C_{X^0} = C_X \text{ or } X \cap \text{an}(C^0, \mathcal{G}_{\overline{X^0}}) = \emptyset \text{ and } \mu^{C^0} \geq \lambda^{C^0} \text{ for all } C^0 \in C_X \setminus C_{X^0}\}$; let $nM_{C|Y,G}(X) := \left\{ \bigcup_{C \in C_X, X \setminus \text{an}(C, \mathcal{G}_{\overline{X}})} [X, nM_{C|Y,G}(X, C)]_{\mu^C \geq \lambda^C} \right\}$ with $[X, nM_{C|Y,G}(X, C)]_{\mu^C \geq \lambda^C} = nM_{C|Y,G}(X, C)$ if $\mu^C \geq \lambda^C$ and X otherwise.

Table 1: Summary of the prior mean $m_{\mathbf{X}}^V(\mathbf{x})$ and kernel $S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta)$ parameters associated to the different surrogate models. $\{g_{\mathbf{X}}^{V,(s)}\}_{s=1}^{S^0}$ denote a set of S^0 realisations of $g_{\mathbf{X}}^V$ obtained by sampling from $p(U)$ and the posterior distributions of f_V given \mathcal{D}^O , $\hat{\sigma}_{\text{do}(\mathbf{X}=\mathbf{x})}^V = \sqrt{\frac{1}{S^0} \sum_{s=1}^{S^0} (g_{\mathbf{X}}^{V,(s)}(\mathbf{x}) - m_{\mathbf{X}}^V(\mathbf{x}))^2}$, σ_f^2 and l are the kernel hyper-parameters, and $S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$ is the kernel function for $u_{\mathbf{X},q}$ with associated scalar coefficient $a_{\mathbf{X},q}^{V,k}$.

	Prior parameters for $g_{\mathbf{X}}^V$	
	$m_{\mathbf{X}}^V(\mathbf{x})$	$S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta)$
cBO	0	$\sigma_f^2 \exp(-\frac{ij\mathbf{x} \cdot \mathbf{x}^\theta jj^2}{2l^2})$
CBO	0	$\sigma_f^2 \exp(-\frac{ij\mathbf{x} \cdot \mathbf{x}^\theta jj^2}{2l^2})$
STGP	0	$\sigma_f^2 \exp(-\frac{ij\mathbf{x} \cdot \mathbf{x}^\theta jj^2}{2l^2})$
STGP ⁺	$\frac{1}{S^0} \sum_{s=1}^{S^0} g_{\mathbf{X}}^{V,(s)}(\mathbf{x})$	$\sigma_f^2 \exp(-\frac{ij\mathbf{x} \cdot \mathbf{x}^\theta jj^2}{2l^2}) + \hat{\sigma}_{\text{do}(\mathbf{X}=\mathbf{x})}^V \times \hat{\sigma}_{\text{do}(\mathbf{X}=\mathbf{x}^\theta)}^V$
MTGP	0	$\sum_{q=1}^Q (a_{\mathbf{X},q}^V)^2 S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$
MTGP ⁺	$\frac{1}{S^0} \sum_{s=1}^{S^0} g_{\mathbf{X}}^{V,(s)}(\mathbf{x})$	$\sum_{q=1}^Q (a_{\mathbf{X},q}^V)^2 S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$
\mathcal{G} -MTGP	$\frac{1}{S^0} \sum_{s=1}^{S^0} g_{\mathbf{X}}^{V,(s)}(\mathbf{x})$	$\frac{1}{S^0} \sum_{s=1}^{S^0} g_{\mathbf{X}}^{V,(s)}(\mathbf{x}) g_{\mathbf{X}}^{V,(s)}(\mathbf{x}^\theta) - \left(\frac{1}{S^0} \sum_{s=1}^{S^0} g_{\mathbf{X}}^{V,(s)}(\mathbf{x})\right) \left(\frac{1}{S^0} \sum_{s=1}^{S^0} g_{\mathbf{X}}^{V,(s)}(\mathbf{x}^\theta)\right)$

Theorem A.2 (Sufficiency of $nM_{\mathcal{C}[Y,G]}$). $nM_{\mathcal{C}[Y,G]} := M_{\mathcal{C}[Y,G]} \setminus \bigcup_{\mathbf{X} \in 2M_{\mathcal{C}[Y,G]}} nM_{\mathcal{C}[Y,G]}(\mathbf{X})$ contains a solution of the cCGO problem (if a solution exists), \forall SCM compatible with \mathcal{G} .

Proof. Consider a set $\mathbf{X}^\theta \in M_{\mathcal{C}[Y,G]}$ that satisfies Eq. (1). If $\mathbf{X}^\theta \notin nM_{\mathcal{C}[Y,G]}$, then $\mathbf{X}^\theta \subseteq \bigcup_{\mathbf{X} \in 2M_{\mathcal{C}[Y,G]}} nM_{\mathcal{C}[Y,G]}(\mathbf{X})$, and therefore there exist at least one set $\mathbf{X} \in M_{\mathcal{C}[Y,G]}$ and at least one $C \in \mathbf{C}_{\mathbf{X}}$ with $\mathbf{X} \cap \text{an}(C, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$ (implying $\mu_{\text{do}(\mathbf{X}=\mathbf{x})}^C = \mu^C$). If $\mu^C < \lambda^C$ then $[\mathbf{X}, nM_{\mathcal{C}[Y,G]}(\mathbf{X}, C)]_{\mu^C} \lambda^C = \mathbf{X}$, and thus $\mathbf{X}^\theta \subseteq \mathbf{X}$ implying that $\mathbf{C}_{\mathbf{X}^\theta} \supseteq \mathbf{C}_{\mathbf{X}}$ and therefore that \mathbf{X}^θ does not satisfies Eq. (1) as $\mu^C < \lambda^C$. If instead $\mu^C \geq \lambda^C$, $\mathbf{X}^\theta \in nM_{\mathcal{C}[Y,G]}(\mathbf{X}, C)$, and thus $\mathbf{X}^\theta \supset \mathbf{X}$, $\mathbf{X}^\theta = \mathbf{X} \cup \mathbf{X}_1$ with $\mathbf{X}_1 \cap \text{an}(\mathbf{C}_{\mathbf{X}} \setminus C \cup Y, \mathcal{G}_{\overline{\mathbf{X}^\theta}}) = \emptyset$. In addition, \mathbf{X} has either the same constraint set of \mathbf{X}^θ (i.e. $\mathbf{C}_{\mathbf{X}^\theta} = \mathbf{C}_{\mathbf{X}}$) or is such that its additional constraints are satisfied ($\mathbf{X} \cap \text{an}(C^0, \mathcal{G}_{\overline{\mathbf{X}}})$ and $\mu^{C^0} \geq \lambda^{C^0}$ for all $C^0 \in \mathbf{C}_{\mathbf{X}} \setminus \mathbf{C}_{\mathbf{X}^\theta}$). Therefore \mathbf{X} also satisfies Eq. (1). If $\mathbf{X} \notin nM_{\mathcal{C}[Y,G]}$ we can apply a similar reasoning and proceed until we find a set in $nM_{\mathcal{C}[Y,G]}$.

The cMI SReduce procedure is given below.

cMISReduce. First construct $M_{\mathcal{C}[Y,G]}$ as: $M_{\mathcal{C}[Y,G]} = \emptyset$, $\forall \mathbf{X} \in \mathcal{P}_{\mathbf{I}}$, if $\mathbf{X} \subseteq \text{an}(\mathbf{C} \cup Y, \mathcal{G}_{\overline{\mathbf{X}}}) \cup (\mathbf{C} \cap \mathbf{X})$ add \mathbf{X} to $M_{\mathcal{C}[Y,G]}$. Then, set $nM_{\mathcal{C}[Y,G]}$ to $M_{\mathcal{C}[Y,G]}$. If $\mathcal{D}^O \neq \emptyset$, $\forall \mathbf{X} \in M_{\mathcal{C}[Y,G]}$, $\forall C \in \mathbf{C}_{\mathbf{X}}$ with $\mathbf{X} \cap \text{an}(C, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$, if C is not null-feasible (according to an estimate $\hat{\mu}^C$ of μ^C) remove \mathbf{X} from $nM_{\mathcal{C}[Y,G]}$. If instead C is null-feasible, remove from $nM_{\mathcal{C}[Y,G]}$ all $\mathbf{X}^\theta \supset \mathbf{X}$ with $\mathbf{X}^\theta = \mathbf{X} \cup \mathbf{X}_1$ where $\mathbf{X}_1 \cap \text{an}(\mathbf{C}_{\mathbf{X}} \setminus C \cup Y, \mathcal{G}_{\overline{\mathbf{X}^\theta}}) = \emptyset$ and such that $\mathbf{C}_{\mathbf{X}^\theta} = \mathbf{C}_{\mathbf{X}}$ or $\mathbf{X} \cap \text{an}(C^0, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$ and $\mu^{C^0} \geq \lambda^{C^0}$ for all $C^0 \in \mathbf{C}_{\mathbf{X}} \setminus \mathbf{C}_{\mathbf{X}^\theta}$.

B. Baselines and Surrogate Models

In the experimental section we compare the performance of cCBO using different surrogate models against CBO, cBO and, for SYNTHETIC-1 and SYNTHETIC-2, also against MCBO.

For the surrogate models of CBO, cBO, and STGP we assume a zero prior mean function $m_{\mathbf{X}}^V(\mathbf{x}) = 0$ and an RBF kernel $S_{\mathbf{X}}^V(\mathbf{x}, \mathbf{x}^\theta) = \sigma_f^2 \exp(-\frac{ij\mathbf{x} \cdot \mathbf{x}^\theta jj^2}{2l^2})$ for all $\mathbf{X} \subseteq \mathcal{P}_{\mathbf{I}}$ and $V \in \mathbf{C}_{\mathbf{X}} \cup Y$. RBF kernels are also considered for the kernel functions $S_{\mathbf{X},q}(\mathbf{x}, \mathbf{x}^\theta)$, $q = 1, \dots, Q$, associated to the latent GPs of MTGP and MTGP⁺. Notice that the surrogate model parameters for cBO, CBO and STGP are equal. Indeed, these methods only differ in terms of search space and acquisition functions used. cBO solves a non-causal constrained global optimization problem therefore considers only one intervention set, i.e. $\mathbf{X} = \mathbf{I}$, models the associated target $\mu_{\text{do}(\mathbf{I}=\mathbf{x})}^Y$ and constraint effects $\mu_{\text{do}(\mathbf{I}=\mathbf{x})}^{C_I}$ independently and selects interventions via the constrained expected improvement acquisition function (Gardner et al., 2014). CBO uses the same surrogate models construction but explores the intervention sets included in $nM_{\mathcal{C}[Y,G]}$ via a standard expected improvement acquisition function thus solving an unconstrained optimization problem. Finally, STGP considers the surrogate model constructions of CBO and cBO but explores the sets included in $nM_{\mathcal{C}[Y,G]}$ selecting interventions via a constrained expected improvement acquisition function.

For all surrogate models exploiting \mathcal{D}^O in the GP prior mean functions (STGP⁺, MTGP⁺ and \mathcal{G} -MTGP), we compute $m_{\mathbf{X}}^V(\mathbf{x})$ for all $\mathbf{X} \in nM_{\mathcal{C}[Y,G]}$ and $V \in \mathbf{C}_{\mathbf{X}} \cup Y$ by averaging the values of V obtained by sampling from the fitted SCM where the functions for \mathbf{X} are fixed to




Figure 7: SYNTHETIC-1 with $N_O = 500$. Scatter plots for the observational data together with interventional ranges $D(X)$ (top) and $D(Z)$ (bottom).

\mathbf{x} . Specifically, we model each function f_V with a GP $f_V(\text{pa}(v)) \sim \mathcal{GP}(0, S^V(\text{pa}(v), \text{pa}(v)^\theta))$, where $\text{pa}(v)$ denotes a value taken by $\text{pa}(V)$. We use an RBF kernel defined as $S^V(\text{pa}(v), \text{pa}(v)^\theta) = \sigma_f^2 \exp(-\frac{j\text{pa}(v) - \text{pa}(v)^\theta j^2}{2l^2})$. We assume a Gaussian likelihood for \mathcal{D}^O and compute the posterior distribution for all f_V in closed form via standard GP updates. We then sample from $p(U)$ and these posteriors to obtain a set of samples $\{g_{\mathbf{X}}^{V,(s)}\}_{s=1}^{S^O}$ for all $\mathbf{X} \in n\mathcal{M}_{C|Y,G}$. The mean functions for the target and constraint effects are then obtained as $m_{\mathbf{X}}^V(\mathbf{x}) = \frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{V,(s)}$ with $V = Y$ and $V = C, \forall C \in \mathcal{C}$ respectively. A similar procedure is used to compute the kernel function $S_{\mathbf{X}}^{V,W}(\mathbf{x}, \mathbf{x}^\theta)$ for \mathcal{G} -MTGP. In particular, given the samples $\{g_{\mathbf{X}}^{V,(s)}\}_{s=1}^{S^O}$ and $\{g_{\mathbf{X}}^{W,(s)}\}_{s=1}^{S^O}$ for each couple of variables (V, W) in $\mathcal{C} \cup Y$, we compute the correlation across the associated effects as $\frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{V,(s)}(\mathbf{x}) g_{\mathbf{X}}^{W,(s)}(\mathbf{x}^\theta) - \left(\frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{V,(s)}(\mathbf{x})\right) \left(\frac{1}{S^O} \sum_{s=1}^{S^O} g_{\mathbf{X}}^{W,(s)}(\mathbf{x}^\theta)\right)$. See Table 1 for a summary of the prior parameters used for each surrogate model.

Similarly to CBO, MCBO solves an unconstrained optimization problem. Instead of explicitly modelling the target effect, MCBO assumes each $V \in \mathcal{V}$ to be of the form $V = f_V(\text{pa}_G(V), \mathbf{A}_V) + U_V$, where \mathbf{A}_V is a set of action variables whose values can be set by the investigator. It then places a vector-valued GP prior on the functions $\{f_V\}_{V \in \mathcal{V}}$ and exploits the posterior distribution of this GP together with a reparameterization trick introduced by Curi et al. (2020) to derive confidence bounds for the target effect. These bounds are then used within an upper confidence bound acquisition function. Notice that, by explicitly modelling the functions in the SCM, MCBO cannot deal with settings in which there are unobserved confounders.

C. Experimental Details and Additional Results

For all experiments we initialize the kernel hyper-parameters (σ_f^2, l) of the surrogate models to 1 and optimize them with a standard type-2 ML approach. When sampling from the modified SCM is required in order to compute the prior mean or kernel functions we use $S^\theta = 10$ samples. The same




Figure 8: SYNTHETIC-2 with $N_O = 500$. Scatter plots for the observational data together with interventional ranges $D(A)$ (top), $D(D)$ (center) and $D(E)$ (bottom).

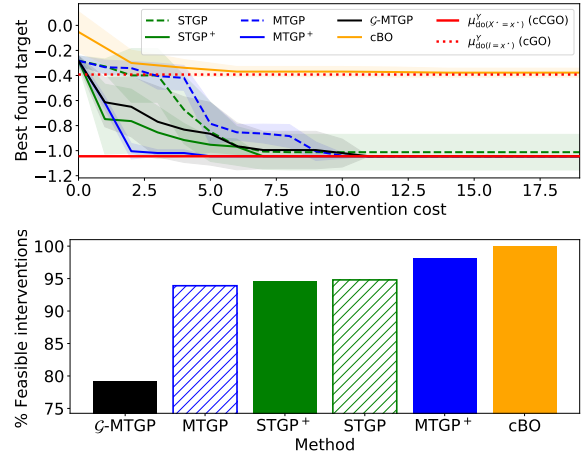


Figure 9: SYNTHETIC-1 with $N_O = 10$ and $\lambda^Z = 2$. *Top*: Convergence to cCGO (solid red line) and cGO (dotted red line) optima. Lines give average results across different initialization of \mathcal{D}^I . Shaded areas represents \pm standard deviation. *Bottom*: Average percentage of feasible interventions collected over trials.

number of samples is used when the computation of the acquisition function requires a Monte Carlo approximation.

C.1. SYNTHETIC-1

For the causal graph in Fig. 2(b) with $\mathbf{I} = \mathcal{C} = \{X, Z\}$ we consider the SCM

$$\begin{aligned} X &= U_X, \quad Z = \exp(-X) + U_Z, \\ Y &= \cos(Z) - \exp(-Z/20) + U_Y, \end{aligned}$$

with $U_X, U_Z, U_Y \sim \mathcal{N}(0, 1)$. We set the interventional ranges to $D(X) = [-3, 2]$ and $D(Z) = [-1, 1]$, the constraint thresholds to $(\lambda^X, \lambda^Z) = (1, 2)$ for Fig. 3, Fig. 9 and Fig. 10 (bottom row) and to $(\lambda^X, \lambda^Z) = (1, 10)$ for Fig. 10 (top row) and Fig. 11, and require the constraint effects to be lower than the thresholds. Notice that, even in cases when N_O is high, there exists a mismatch between the observational and interventional ranges. Fig. 7 shows how

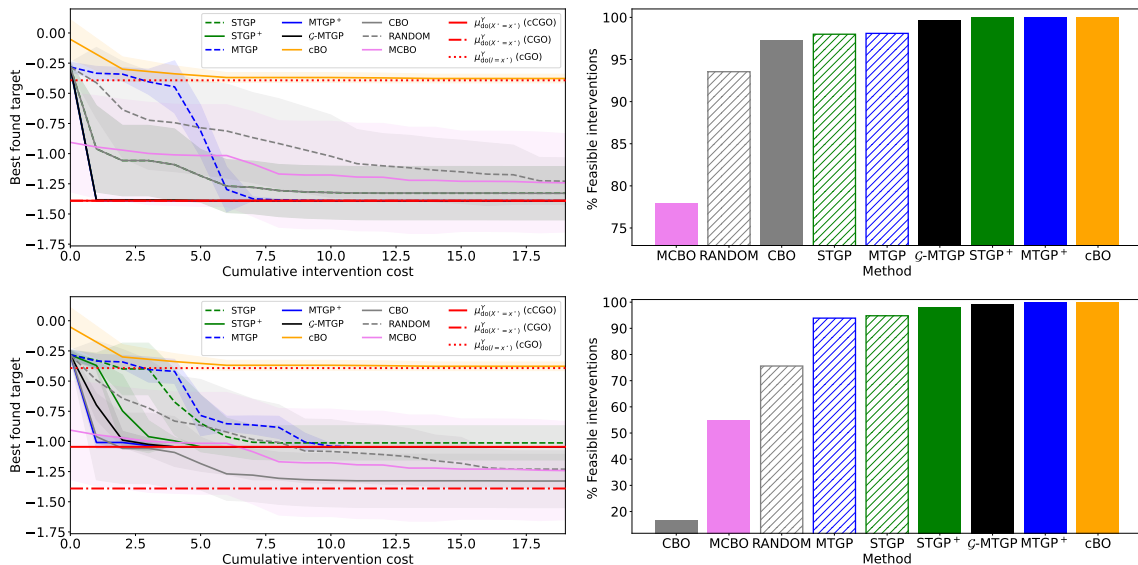


Figure 10: SYNTHETIC-1 with $N_O = 500$ and $\lambda^Z = 10$ (top row), and with $N_O = 500$ and $\lambda^Z = 2$ (bottom row). *Left*: Convergence to the cCGO (solid red line), CGO (dash-dotted red line) and cGO (dotted red line) optima. *Right*: Average percentage of feasible interventions collected over trials.

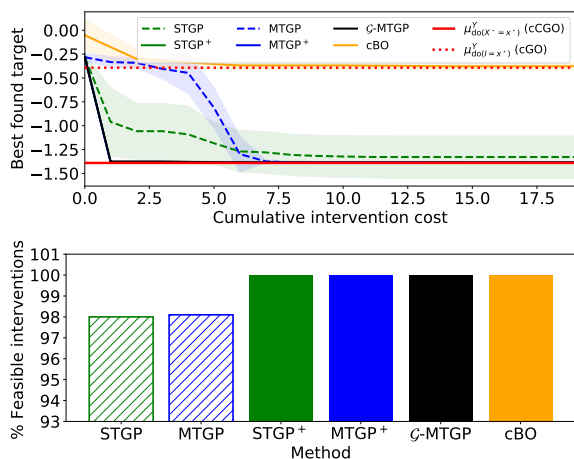


Figure 11: SYNTHETIC-1 with $N_O = 100$ and $\lambda^Z = 10$. *Top*: Convergence to the cCGO (solid red line) and cGO (dotted red line) optima. *Bottom*: Average percentage of feasible interventions collected over trials.

the interventional ranges are only partially covered by the observational data, especially for X . Therefore, estimating the target and constraint effects with \mathcal{D}^O for interventions values outside of the observational ranges, say for $X = -2$, might lead to inaccurate results which translate to unreliable prior parameters for the surrogate models exploiting \mathcal{D}^O . A combination of interventional and observational data is needed in order to quantify uncertainty around the effects and to choose the intervention to perform next while en-

suring satisfaction of the constraints, and thus to efficiently identify an optimal feasible intervention.

In this graph, $\mathcal{P}_I = \{\{X\}, \{Z\}, \{X, Z\}\}$, and $M_{\mathcal{C}|Y,G} = \mathcal{P}_I$ as $\mathbf{X} \subseteq \text{an}(\{X, Z, Y\}, \mathcal{G}_{\overline{\mathbf{X}}}) \cup (\{X, Z\} \cap \mathbf{X})$, $\forall \mathbf{X} \in \mathcal{P}_I$. The set $\mathbf{X} = \{Z\}$ has only one constrained variable, X , that is reducible (as $\mathbf{X} \cap \text{an}(X, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$) and null-feasible for $N_O = 500, 100, 10$. We can thus exclude $\mathbf{X}^0 = \{X, Z\} \supset \mathbf{X}$ from $M_{\mathcal{C}|Y,G}$. Indeed, $\{X, Z\}$ is of the form $\mathbf{X}^0 = \mathbf{X} \cup X$ with $X \notin \text{an}(\emptyset \cap Y, \mathcal{G}_{\overline{\mathbf{X}}})$. We therefore obtain $nM_{\mathcal{C}|Y,G} = \{\{X\}, \{Z\}\}$.

Fig. 10 (top row) shows the convergence results for $N_O = 500$ and $\lambda^Z = 10$, for which the cCGO and CGO problems have the same optimum. Notice how both CBO and MCBO converge to the optimum at a slower pace compared to STGP+, MTGP+ and \mathcal{G} -MTGP. Indeed, CBO and MCBO only take into account the target value observed after performing each intervention. CBO also models each function individually and, by disregarding the values of the constraints, needs to perform more interventions before identifying an optimal one. More importantly, MTGP+, STGP+ and \mathcal{G} -MTGP collect more than 99% of feasible interventions over trials. This is a critical aspect of the method as in real-world applications the investigator might not know whether an optimal solution is in a feasible region or not. Using cCBO in such settings does not slow down the identification of an optimal solution, but rather improves it, while allowing to efficiently restrict the exploration regions. For completeness, in Fig. 10 (bottom row) we also report the comparison with CBO, MCBO, and RANDOM for $N_O = 500$ and $\lambda^Z = 2$.

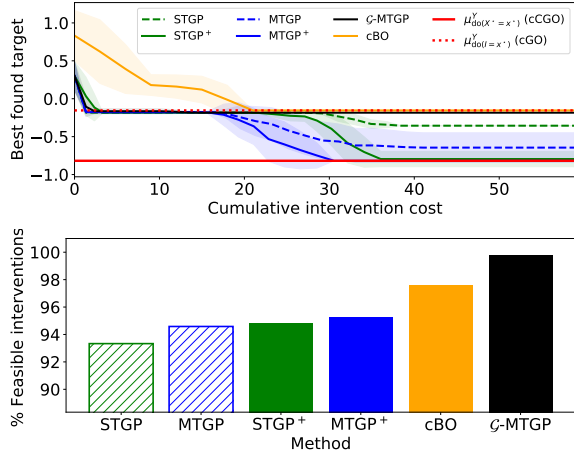


Figure 12: SYNTHETIC-2 with $N_O = 10$. *Top*: Convergence to the cCGO (solid red line) and cGO (dotted red line) optima. *Bottom*: Average percentage of feasible interventions collected over trials.

C.2. SYNTHETIC-2

For the causal graph in Fig. 2(a) with $I = \{A, D, E\}$ and $C = \{C, D, E\}$ we consider the SCM

$$\begin{aligned} A &= U_A, \quad B = U_B, \\ C &= \exp(-A)/5 + U_C, \\ D &= \cos(B) + C/10 + U_D, \\ E &= \exp(-C)/10 + U_E, \\ Y &= \cos(D) - D/5 + \sin(E) - E/4 + U_Y, \end{aligned}$$

with $U_A, U_B, U_C, U_D, U_E, U_Y \sim \mathcal{N}(0, 1)$. We set the interventional ranges to $D(A) = [-5, 5]$, $D(D) = [-1, 1]$ and $D(E) = [-1, 1]$, the constraint thresholds to $\lambda^C = 10$, $\lambda^D = 10$, $\lambda^E = 10$, and require all constraint effects to be smaller than the thresholds. For $N_O = 500, 100, 10$ all variables in C are null-feasible, giving $nM_{C|Y,G} = \mathcal{P}_I \setminus \{A, D, E\} = \{\{A\}, \{D\}, \{E\}, \{A, D\}, \{A, E\}, \{D, E\}\}$ (see the example in Section 3.1). As in SYNTHETIC-1, even in cases when N_O is high, there exists a mismatch between the observational and interventional ranges (Fig. 8) for all interventions sets in $nM_{C|Y,G}$.

Fig. 12 shows the convergence plots and associated percentage of feasible interventions when $N_O = 10$. As discussed in Section 5, when the functions in the SCM cannot be learned accurately due to limited or noisy observational data, MTGP+ should be preferred. Indeed, a lower value for N_O leads to a less accurate estimation of the target and constraint effects from \mathcal{D}^O which negatively affects the prior mean functions of STGP+, MTGP+ and G-MTGP. G-MTGP is further penalized in this case as also the kernel functions are computed exploiting the fitted SCM functions. As a

consequence, we observe slower convergence and a lower percentage of feasible interventions collected when using G-MTGP. In particular, notice how the inaccurate estimation of the constraint effects translates into a significantly lower convergence speed for G-MTGP compared to the settings in which N_O is higher. On the contrary, MTGP+ successfully trades off feasibility and improvement in these experiments reaching convergence while collecting a high percentage of feasible interventions over trials. Finally, by breaking the existing causal relationships and intervening on all variables in I , cBO blocks the propagation of causal effects in the graph and converges to a solution (cGO) that is sub-optimal with respect to the one achieved by all cCGO instances.

Finally, Fig. 13 shows how, in settings where the cCGO and CGO optima differ, CBO, MCBO, and RANDOM converge to a solution which is lower than the cCGO one. In addition, by disregarding the constraints, CBO, MCBO, and RANDOM collect a high number of infeasible interventions over trials.

C.3. HEALTH

For the causal graph in Fig. 1(b) with $I = \{\text{Statin}, \text{Aspirin}, \text{CI}\}$ and $C = \{\text{BMI}\}$, the SCM is taken from Ferro et al. (2015) and can be written as

$$\begin{aligned} \text{Age} &= U_{\text{Age}}, \quad \text{CI} = U_{\text{CI}}, \quad \text{BMR} = 1500 + 10 \times U_{\text{BMR}}, \\ \text{Height} &= 175 + 10 \times U_{\text{Height}}, \\ \text{Weight} &= \frac{\text{BMR} + 6.8 \times \text{Age} - 5 \times \text{Height}}{13.7 + \text{CI} \times 150/7716}, \\ \text{BMI} &= \text{Weight}/(\text{Height}/100)^2, \\ \text{Aspirin} &= \sigma(-8 + 0.1 \times \text{Age} + 0.03 \times \text{BMI}), \\ \text{Statin} &= \sigma(-13 + 0.1 \times \text{Age} + 0.2 \times \text{BMI}), \\ \text{PSA} &= 6.8 + 0.04 \times \text{Age} - 0.15 \times \text{BMI} - 0.60 \times \text{Statin} \\ &\quad + 0.55 \times \text{Aspirin} + \sigma(2.2 - 0.05 \times \text{Age} \\ &\quad + 0.01 \times \text{BMI} - 0.04 \times \text{Statin} \\ &\quad + 0.02 \times \text{Aspirin}) + U_{\text{PSA}}, \end{aligned}$$

with $U_{\text{Age}} \sim \mathcal{U}(55, 75)$, $U_{\text{CI}} \sim \mathcal{U}(-100, 100)$, $U_{\text{BMR}} \sim t\mathcal{N}(-1, 2)$, $U_{\text{Height}} \sim t\mathcal{N}(-0.5, 0.5)$, $U_{\text{PSA}} \sim \mathcal{N}(0, 0.4)$, where $\mathcal{U}(\cdot, \cdot)$ denotes a uniform distribution, $t\mathcal{N}(a, b)$ a standard Gaussian distribution truncated between a and b , and $\sigma(\cdot)$ the sigmoidal transformation defined as $\sigma(x) = \frac{1}{1 + \exp(-x)}$. We set the interventional ranges to $D(\text{Aspirin}) = [0, 1]$, $D(\text{Statin}) = [0, 1]$, and $D(\text{CI}) = [-400, 400]$, and require BMI to be lower than 25.

For $N_O = 100$, these interventional ranges are only partially covered by the observational data (Fig. 14) thus significantly complicating the estimation of the effects with \mathcal{D}^O , especially when the stochasticity determined by $p(U)$ is high (see scatter plots for CI, Aspirin, Statin, and PSA in Fig. 15).

We have $\mathcal{P}_I = \{\{\text{Aspirin}\}, \{\text{Statin}\}, \{\text{CI}\}, \{\text{Aspirin}, \text{Statin}\},$




Figure 13: SYNTHETIC-2 with $N_O = 500$. *Top*: Convergence to the cCGO (solid red line), CGO (dash-dotted red line) and cGO (dotted red line) optima. *Bottom*: Average percentage of feasible interventions collected over trials.




Figure 14: HEALTH with $N_O = 100$. Scatter plots for the observational data together with interventional ranges $D(CI)$ (top), $D(Statin)$ (center) and $D(Aspirin)$ (bottom).

$\{\text{Aspirin}, CI\}, \{\text{Statin}, CI\}, \{\text{Aspirin}, \text{Statin}, CI\}$ which is equal to $M_{C|Y,G}$. BMI is reducible for $\mathbf{X} = \{\text{Aspirin}\}$, $\mathbf{X} = \{\text{Statin}\}$ and $\mathbf{X} = \{\text{Aspirin}, \text{Statin}\}$ as $\mathbf{X} \cap \text{an}(\text{BMI}, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$. Given an observational dataset of size $N_O = 100, 10$, $\hat{\mu}^{\text{BMI}} \approx 26$ making BMI not null-feasible. We therefore obtain $nM_{C|Y,G} = \{\{CI\}, \{\text{Aspirin}, CI\}, \{\text{Statin}, CI\}, \{\text{Aspirin}, \text{Statin}, CI\}\}$.

For completeness we report the full comparison with CBO and RANDOM in Fig. 16. As in the previous experiments, the cCGO and CGO optima differ thus CBO and RANDOM converge, at a much lower speed, to different values while collecting a large number of infeasible interventions.

C.4. PROTEIN-SIGNALING

For the causal graph⁷ of Fig. 1(a) with $\mathbf{I} = \{\text{PKC}, \text{PKA}, \text{Mek}, \text{Akt}\}$ and $\mathbf{C} = \{\text{PKC}, \text{PKA}\}$, Sachs et al. (2005) give a dataset but no SCM. Thus we first fit a SCM by placing a GP prior with zero mean and RBF kernel on f_V , $\forall V \in \mathbf{V}$, and using the observational data. We assume a Gaussian distribution with zero mean and unit variance for PKC. We use the learned SCM to generate observational and interventional data.

⁷We exclude the Plc γ phospho-protein, the PIP2 and the PIP3 phospho-lipids, as these form a separate graph with no connections to Erk.




Figure 15: HEALTH. $N_O = 100$ observational data samples from the SCM in Section C.3.

As in the previous experiments, the interventional ranges are only partially covered by $N_O = 100$ observational samples (Fig. 17) which complicates the estimation of the effects with \mathcal{D}^O . The power set of \mathbf{I} is given by $\mathcal{P}_{\mathbf{I}} = \{\{\text{PKC}\}, \{\text{PKA}\}, \{\text{Mek}\}, \{\text{Akt}\}, \{\text{PKC}, \text{PKA}\}, \{\text{PKC}, \text{Mek}\}, \{\text{PKC}, \text{Akt}\}, \{\text{PKA}, \text{Mek}\}, \{\text{PKA}, \text{Akt}\}, \{\text{Mek}, \text{Akt}\}, \{\text{PKC}, \text{PKA}, \text{Mek}\}, \{\text{PKC}, \text{PKA}, \text{Akt}\}, \{\text{PKA}, \text{Mek}, \text{Akt}\}\}$. Note that, for $\mathbf{X} \in \{\{\text{Akt}\}, \{\text{PKC}, \text{Akt}\}, \{\text{PKA}, \text{Akt}\}, \{\text{Mek}, \text{Akt}\}, \{\text{PKC}, \text{PKA}, \text{Akt}\}, \{\text{PKA}, \text{Mek}, \text{Akt}\}\}$, we have $\mathbf{X} \not\subseteq \text{an}(\{\text{PKC}, \text{PKA}, \text{Erk}\}, \mathcal{G}_{\overline{\mathbf{X}}}) \cup (\{\text{PKC}, \text{PKA}\} \cap \mathbf{X})$ thus we can exclude these sets from $\mathcal{P}_{\mathbf{I}}$ and obtain $M_{C|Y,G} = \mathcal{P}_{\mathbf{I}} \setminus \{\{\text{Akt}\}, \{\text{PKC}, \text{Akt}\}, \{\text{PKA}, \text{Akt}\}, \{\text{Mek}, \text{Akt}\}, \{\text{PKC}, \text{PKA}, \text{Akt}\}, \{\text{PKA}, \text{Mek}, \text{Akt}\}\}$. Given \mathcal{D}^O with size $N_O = 100, 10$, all constrained variables are null-feasible.

The set $\mathbf{X} = \{\text{PKC}, \text{PKA}\}$ in $M_{C|Y,G}$ has only one constrained variable PKC that is reducible (as $\mathbf{X} \cap \text{an}(\text{PKC}, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$) and null-feasible for $N_O = 100, 10$. We can thus exclude $\mathbf{X}^0 = \{\text{PKC}, \text{PKA}, \text{Mek}\} \supset \mathbf{X}$ from $M_{C|Y,G}$ as it is of the form $\mathbf{X}^0 = \mathbf{X} \cup \text{PKC}$

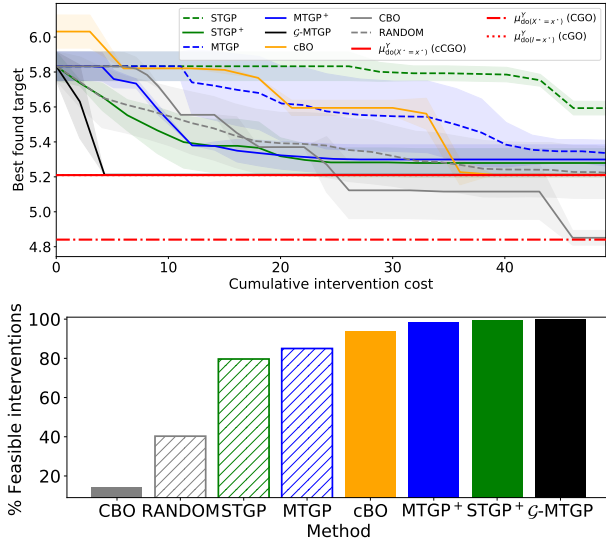


Figure 16: HEALTH with $N_O = 100$. *Top*: Convergence to the cCGO (solid red line), CGO (dash-dotted red line) and cGO (dotted red line) optima. *Bottom*: Average percentage of feasible interventions collected over trials.

with $PKC \notin \text{an}(\emptyset \cup Y, \mathcal{G}_{\overline{\mathbf{X}}}) = \emptyset$. We therefore obtain $nM_{\mathcal{C}[Y, \mathcal{G}] = M_{\mathcal{C}[Y, \mathcal{G}]\{PKC, PKA, Mek\}}$.

In addition, as the CGO optimum equals the cCGO optimum, CBO and RANDOM converge to the same value achieved by cCBO (Fig. 18, top row) but collect a significantly higher number of infeasible interventions (Fig. 18, bottom row).

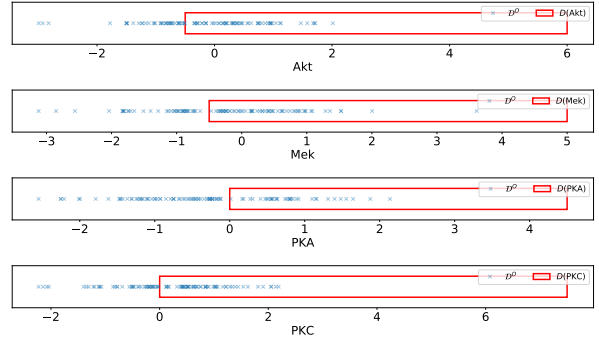


Figure 17: PROTEIN-SIGNALING with $N_O = 100$. Scatter plots for the observational data together with interventional ranges $D(\text{Akt})$ (top row), $D(\text{Mek})$ (second row), $D(\text{PKA})$ (third row) and $D(\text{PKC})$ (bottom row).

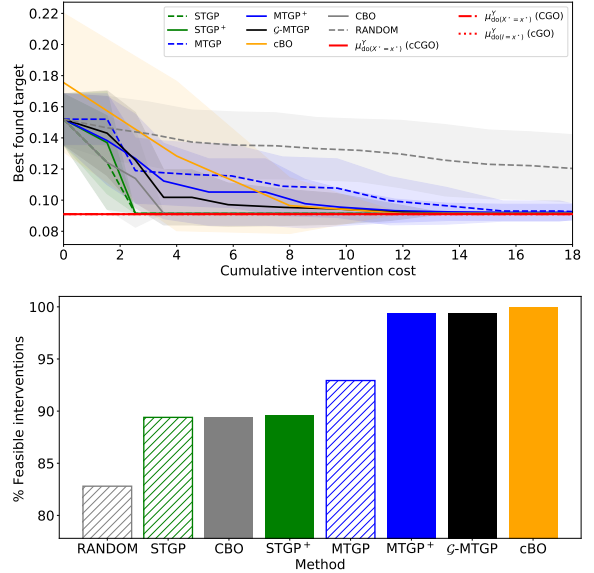


Figure 18: PROTEIN-SIGNALING with $N_O = 100$. *Top*: Convergence to the cCGO (solid red line), CGO (dash-dotted red line) and cGO (dotted red line) optima. *Bottom*: Percentage of feasible interventions collected over trials.