Causal decision-making in static and dynamic settings

Virginia Aglietti[1]

Collaborators: Neil Dhir[2], Theo Damoulas[2,3], Javier González[4]

1.DeepMind 2.The Alan Turing Institute 3. The University of Warwick 4.Microsoft Research Cambridge

Systems decompose in sets of interconnected nodes

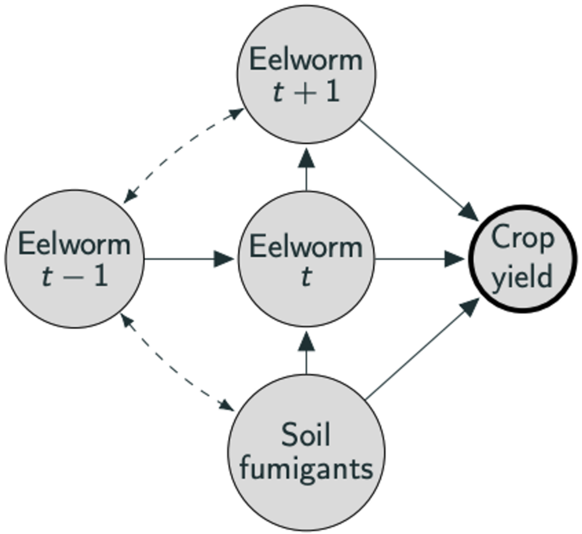

Figure 1: Causal Graph for Crop Yield.

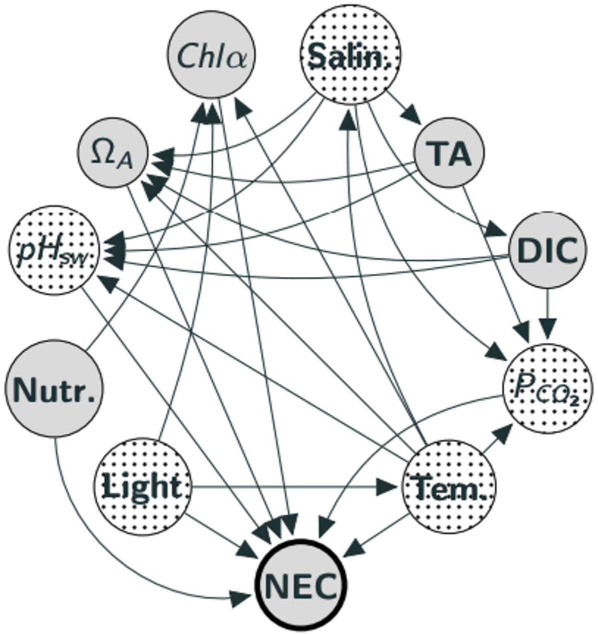

Figure 2: Causal Graph for Net Ecosystem Calcification (NEC).

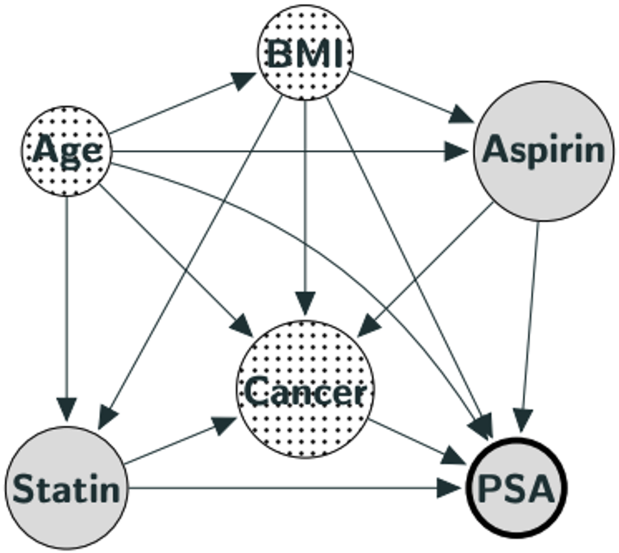

Figure 3: Causal Graph for Prostate Specific Antigen (PSA) level.

Setting and goal

- A causal graph (Directed Acyclic Graph - DAG).
- Observational data from all (non hidden) nodes.
- Ability of running experiments (in reality or in simulation).
- Cost of experiments depends on the number and type of nodes in which we intervene.

Goal: Efficiently find the optimal intervention to perform.

Exploiting all available source of information.

Intervention optimizing a target node in a graph.

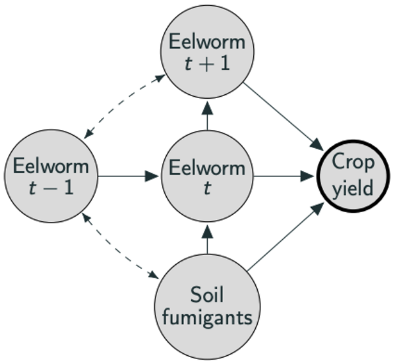

E.g. In order to optimize the crop yield should we intervene on soil fumigants or on eel-worm population? If the optimal intervention is soil fumigants (intervention set), what level should we set them to (intervention level)?

Causal model and the do-calculus (1/3)

Causal Model : DAG $\mathcal{G}$ + four-tuple $\langle \mathrm{U}, \mathrm{V}, F, P(\mathrm{U}) \rangle$

- U: independent *exogenous* background variables.
- $P(\mathrm{U})$ distribution of U.
- V: *endogenous* variables (non-manipulative, manipulative, target).
- $F = \{f_1, ..., f_{|\mathrm{V}|}\}$: functions $v_i = f_i(pa_i, u_i)$, pa_i are the parents of V_i.

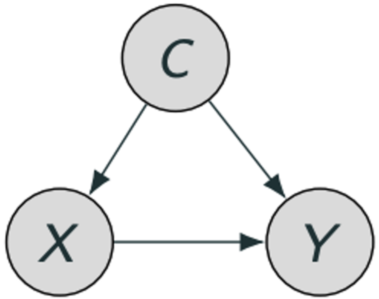

$$
\begin{aligned}
C &= f_c(U_c),\ U_c \sim \mathcal{N}(0, \sigma_c^2) \\
X &= f_x(C, U_x),\ U_x \sim \mathcal{N}(0, \sigma_x^2) \\
Y &= f_y(X, C, U_y),\ Y_c \sim \mathcal{N}(0, \sigma_y^2)
\end{aligned}
$$

Causal model and the do-calculus (2/3)

Intervention : Setting a manipulative variable X to a value x, $do(X = x)$

Observed universe

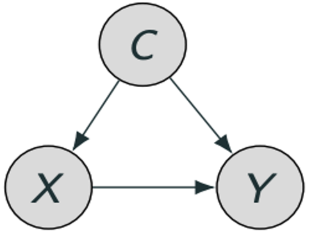

$$\begin{aligned} C &= f_c(U_c) \\ X &= f_x(C, U_x) \\ Y &= f_y(X, C, U_y) \end{aligned} \quad \Leftrightarrow \quad P(X, C, Y)$$

Post-intervention universe

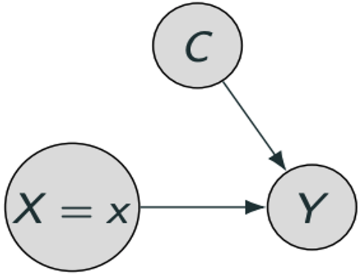

$$\begin{aligned} C &= f_c(U_c) \\ X &= x \\ Y &= f_y(x, C, U_y) \end{aligned} \quad \Leftrightarrow \quad P^{do(X=x)}(C, Y)$$

$$\boxed{P(Y|do(X = x)) := P^{do(X=x)}(Y|X = x)}$$

Causal model and the do-calculus (3/3)

Key question : How to do inference in the post-intervention universe.

- Intervene $\Rightarrow$ Interventional data $\Rightarrow$ $P(Y | \mathrm{do}(X = x))$
- Observe $\Rightarrow$ Observational data $\Rightarrow$ **do-calculus** $\Rightarrow$ $\hat{P}(Y | \mathrm{do}(X = x))$

Do-calculus: algebra to emulate the post-intervention universe in terms of conditionals P(Y |X = x) in the observed universe.

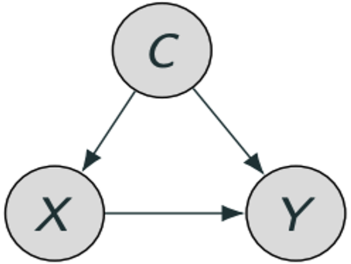

Back-door adjustment

$$p(Y|do(X = x)) = \int P(Y|c, X = x)P(c)dc$$

Take home messages

- Many real systems decompose in interconnected nodes.
- Optimization of an experimental output requires “intervening” in the manipulative nodes.
- Do-calculus allows “simulating” experiments with observational data.

Causal Global Optimization

$$\mathbf{X}_s^\star, \mathbf{x}_s^\star = \underset{\mathbf{X}_s \in \mathcal{P}(\mathbf{X}), \mathbf{x}_s \in D(\mathbf{X}_s)}{\arg\min} \mathbb{E}_{P(\mathbf{Y} | \mathrm{do}(\mathbf{X}_s = \mathbf{x}_s))}[\mathbf{Y}]$$

- Optimal intervention set
- Optimal intervention level

Exploring all possible interventions in a causal graph.

With the final goal of optimizing a target variable.

- $\mathbf{X}_s$ and $\mathbf{x}_s$ are one possible intervention set and value.
- $\mathbf{X}_s^\star$ and $\mathbf{x}_s^\star$ are the optimal intervention set and value.

Causal Global Optimization

Global Optimization

$$\mathbf{x}^{\star} = \arg\min_{\mathbf{x} \in D(\mathbf{X})} \mathbb{E}[Y \mid \mathrm{do}(\mathbf{X} = \mathbf{x})]$$

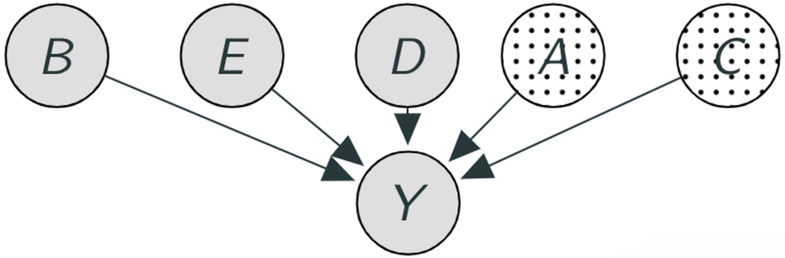

Causal Global Optimization

$$\mathbf{X}_s^{\star}, \mathbf{x}_s^{\star} = \arg\min_{\substack{\mathbf{X}_s \in \mathcal{P}(\mathbf{X}) \\ \mathbf{x}_s \in D(\mathbf{X}_s)}} \mathbb{E}[Y \mid \mathrm{do}(\mathbf{X}_s = \mathbf{x}_s)]$$

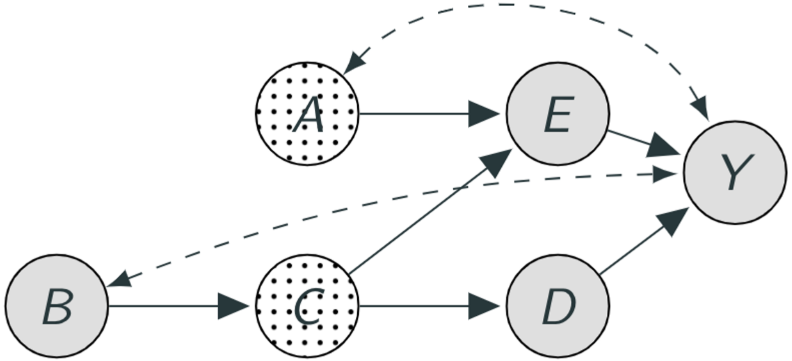

Causal Global Optimization

Global Optimization

$$\mathbf{x}^\star = \arg\min_{\mathbf{x} \in D(\mathbf{X})} \mathbb{E}[Y | \mathrm{do}(\mathbf{X} = \mathbf{x})]$$

- Target function is explicitly unknown and multimodal
- Evaluations are perturbed by noise
- Evaluations are expensive

Bayesian Optimization

Causal Global Optimization

$$\mathbf{X}_s^\star, \mathbf{x}_s^\star = \arg\min_{\substack{\mathbf{X}_s \in \mathcal{P}(\mathbf{X}) \\ \mathbf{x}_s \in D(\mathbf{X}_s)}} \mathbb{E}[Y | \mathrm{do}(\mathbf{X}_s = \mathbf{x}_s)]$$

- Target function is explicitly unknown and multimodal
- Evaluations are perturbed by noise
- Evaluations are expensive

+ Causal Graph

Causal Bayesian Optimization

Causal Global Optimization

1. Limit the sets to explore by identifying interventions worth exploring;
2. Construct a surrogate model incorporating observational and interventional data;
3. Extend the expected improvement acquisition function to explore different intervention sets;
4. Allow the agent to observe or intervene.

Causal Global Optimization

$$\mathbf{X}_s^\star, \mathbf{x}_s^\star = \underset{\substack{\mathbf{X}_s \in \mathcal{P}(\mathbf{X}) \\ \mathbf{x}_s \in D(\mathbf{X}_s)}}{\arg\min}\ \mathbb{E}[Y | \mathrm{do}\,(\mathbf{X}_s = \mathbf{x}_s)]$$

- Target function is explicitly unknown and multimodal
- Evaluations are perturbed by noise
- Evaluations are expensive

+ Causal Graph

Causal Bayesian Optimization

Causal Bayesian Optimization

$$\mathbf{X}_s^\star, \mathbf{x}_s^\star = \underset{\boxed{\mathbf{X}_s\in\mathcal{P}(\mathbf{X}),\mathbf{x}_s\in D(\mathbf{X}_s)}}{\arg\min} \mathbb{E}_{P(\mathbf{Y}|\mathrm{do}(\mathbf{X}_s=\mathbf{x}_s))}[\mathbf{Y}]$$

1) Identify sets worth intervening on

Definition 3.1. Minimal Intervention set (MIS). Given $\langle \mathcal{G}, \mathbf{Y}, \mathbf{X}, \mathbf{C} \rangle$, a set of variables $\mathbf{X}_s \in \mathcal{P}(\mathbf{X})$ is said to be a MIS if there is no $\mathbf{X}'_s \subset \mathbf{X}_s$ such that $\mathbb{E}[Y|\mathrm{do}(\mathbf{X}_s = \mathbf{x}_s)] = \mathbb{E}[Y|\mathrm{do}(\mathbf{X}'_s = \mathbf{x}'_s)]$.

$$\mathbb{E}[Y|\mathrm{do}(X=x),\mathrm{do}(Z=z)] = \mathbb{E}[Y|\mathrm{do}(Z=z)]$$

Definition 3.2. Possibly-Optimal Minimal Intervention set (POMIS). Given $\langle \mathcal{G}, \mathbf{Y}, \mathbf{X}, \mathbf{C} \rangle$, let $\mathbf{X}_s \in \mathbb{M}^{\mathbf{C}}_{\mathcal{G},Y}$. $\mathbf{X}_s$ is a POMIS if there exists a SEM conforming to $\mathcal{G}$ such that $\mathbb{E}[Y|\mathrm{do}(\mathbf{X}_s = \mathbf{x}^*)] > \forall_{\mathbf{W}\in\mathbb{M}^{\mathbf{C}}_{\mathcal{G},Y}\setminus\mathbf{X}_s} \mathbb{E}[Y|\mathrm{do}(\mathbf{W} = \mathbf{w}^*)]$ where $\mathbf{x}^*$ and $\mathbf{w}^*$ denote the optimal intervention values.

$$\begin{aligned}
\mathbb{E}[Y|\mathrm{do}(X=x^\star)] &= \int_{\mathcal{Z}} \mathbb{E}[Y|\mathrm{do}(Z=z)]p(z|\mathrm{do}(X=x^\star))dz \\
&\leq \int_{\mathcal{Z}} \mathbb{E}[Y|\mathrm{do}(Z=z^\star)]p(z|\mathrm{do}(X=x^\star))dz \\
&= \mathbb{E}[Y|\mathrm{do}(Z=z^\star)]
\end{aligned}$$

Lee, Sanghack, and Elias Bareinboim. "Structural causal bandits: where to intervene?." *Advances in Neural Information Processing Systems 31* 31 (2018).

Causal Bayesian Optimization

$$\mathbf{X}_s^\star, \mathbf{x}_s^\star = \underset{\mathbf{X}_s \in \mathcal{P}(\mathbf{X}), \mathbf{x}_s \in D(\mathbf{X}_s)}{\arg\min} \boxed{\mathbb{E}_{P(\mathbf{Y}|\mathrm{do}(\mathbf{X}_s=\mathbf{x}_s))}[\mathbf{Y}]}$$

2) Construct surrogate models

$$f(x_s) \sim \mathcal{GP}(m(x_s), k(x_s, x_s'))$$

$$m(x_s) = \hat{\mathbb{E}}[Y|\mathrm{do}\,(X_s = x_s)]$$

$$k(x_s, x_s') = k_{RBF}(x_s, x_s') + \sigma(x_s)\sigma(x_s')$$

- $k_{RBF}(x_s, x_s') := \exp(-\frac{||x_s - x_s'||^2}{2l^2})$
- $\sigma(x_s) = \sqrt{\hat{\mathbb{V}}(Y|\mathrm{do}\,(X_s = x_s))}$ with $\hat{\mathbb{V}}$ is the variance of the causal effects estimated from observational data.

Toy Example

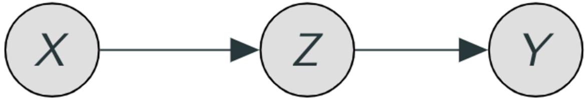

$$X = \epsilon_X$$

$$Z = \exp(-X) + \epsilon_Z$$

$$Y = \cos(Z) - \exp(-\frac{Z}{20}) + \epsilon_Y$$

$$\boxed{\begin{aligned} \mathbb{M}_{\mathcal{G},Y} &= \{\emptyset, \{X\}, \{Z\}\} \\ \mathbb{P}_{\mathcal{G},Y} &= \{\{Z\}\} \\ \mathbb{B}_{\mathcal{G},Y} &= \{\{X, Z\}\} \end{aligned}}$$

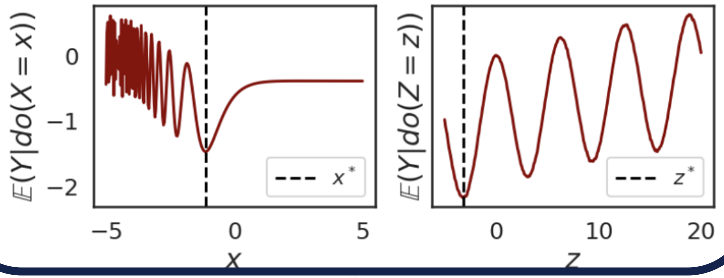

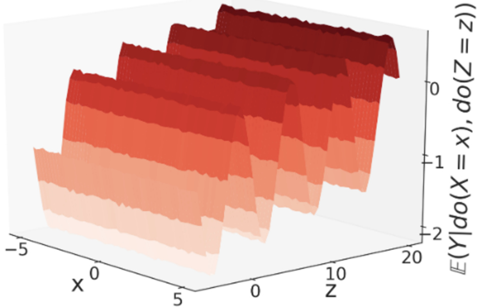

Sets worth intervening on based on the causal graph structure.

Toy Example

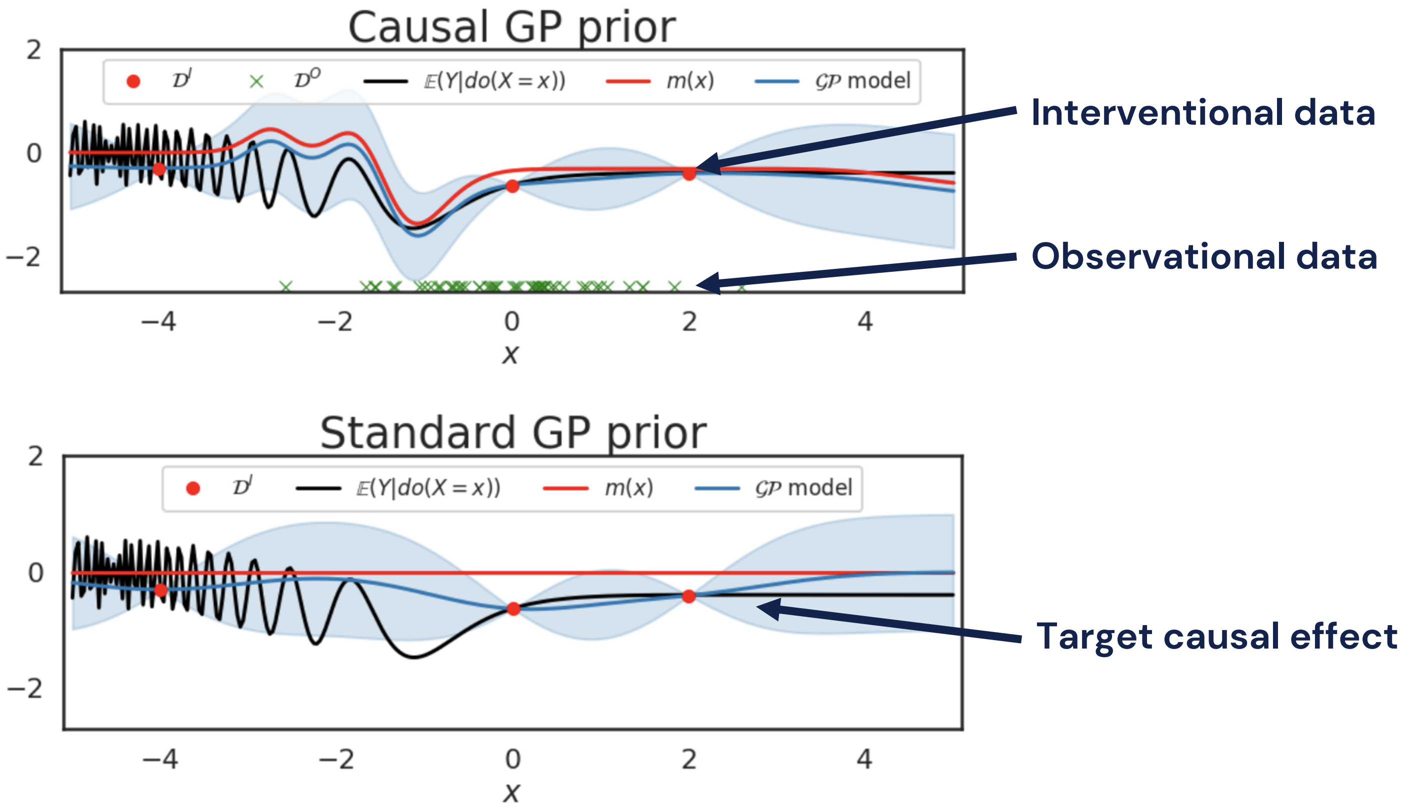

Causal GP prior
$\mathcal{D}^I$
$\mathcal{D}^O$
$\mathbb{E}(Y|do(X=x))$
$m(x)$
$\mathcal{GP}$ model
2
0
−2
−4
−2
0
2
4
X
Interventional data
Observational data
Standard GP prior
$\mathcal{D}^I$
$\mathbb{E}(Y|do(X=x))$
$m(x)$
$\mathcal{GP}$ model
2
0
−2
−4
−2
0
2
4
X
Target causal effect

Causal Bayesian Optimization

Select Actions

$$EI^s(\mathbf{x}) = \mathbb{E}_{p(y_s)}[\max(y_s - y^\star, 0)]/Co(\mathbf{X}_s, \mathbf{x}_s)$$

- $y_s = \mathbb{E}[Y|\mathrm{do}(\mathbf{X}_s = \mathbf{x}_s)]$
- $y^\star = \max_{\mathbf{X}_s \in \mathbf{es}, \mathbf{x} \in D(\mathbf{X}_s)} \mathbb{E}[Y|\mathrm{do}(\mathbf{X}_s = \mathbf{x}_s)]$

Optimize EI for every set and select the set giving the highest expected improvement.

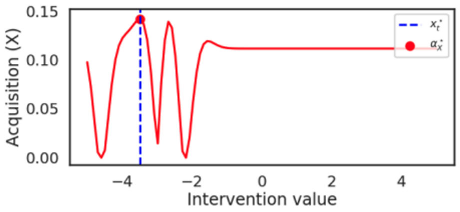

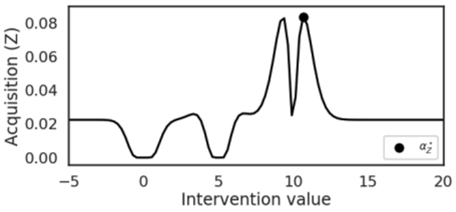

Address the intervention-observation trade-off

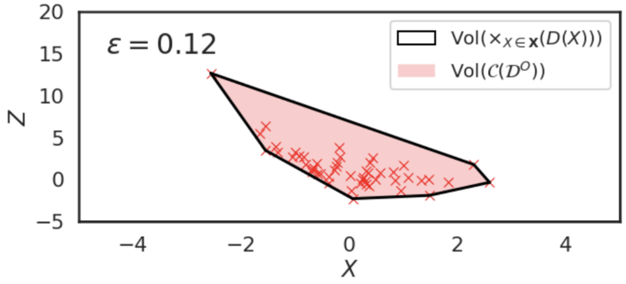

$$\epsilon = \frac{\mathrm{Vol}(\mathcal{C}(\mathcal{D}^O))}{\mathrm{Vol}(\times_{X\in\mathbf{X}}(D(X)))} \times \frac{N}{N_{\max}}$$

Causal Bayesian Optimization

Algorithm: Causal Bayesian Optimization

Data: $\mathcal{D}^{O}$, $\mathcal{D}^{I}$, $\mathcal{G}$, **es**, number of steps T

Result: $\mathbf{X}_s^\star, \mathbf{x}_s^\star, \hat{\mathbb{E}}[\mathbf{Y}^\star | \mathrm{do}(\mathbf{X}_s^\star = \mathbf{x}_s^\star)]$

Initialise: Set $\mathcal{D}_0^{I} = \mathcal{D}^{I}$ and $\mathcal{D}_0^{O} = \mathcal{D}^{O}$

BO Loop

for t=1, ..., T **do**

Compute ϵ and sample $u \sim \mathcal{U}(0, 1)$

if $\epsilon > u$ **then**

(Observe)

1. Observe new observations $(\mathbf{x}_t, c_t, \mathbf{y}_t)$.
2. Augment $\mathcal{D}^{O} = \mathcal{D}^{O} \cup \{(\mathbf{x}_t, c_t, \mathbf{y}_t,)\}$.
3. Update prior of the causal GP.

Update prior distributions

end

else

(Intervene)

1. Compute $EI^s(\mathbf{x})$ for each element $s \in$ **es**.
2. Obtain the optimal interventional set-value pair $(s^\star, \alpha^\star)$.
3. Intervene on the system.
4. Update posterior of the interventional GP.

Update posterior distribution

end

end

Simulation Results

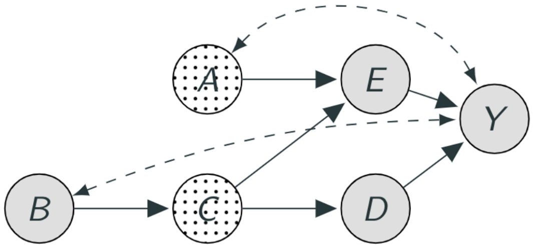

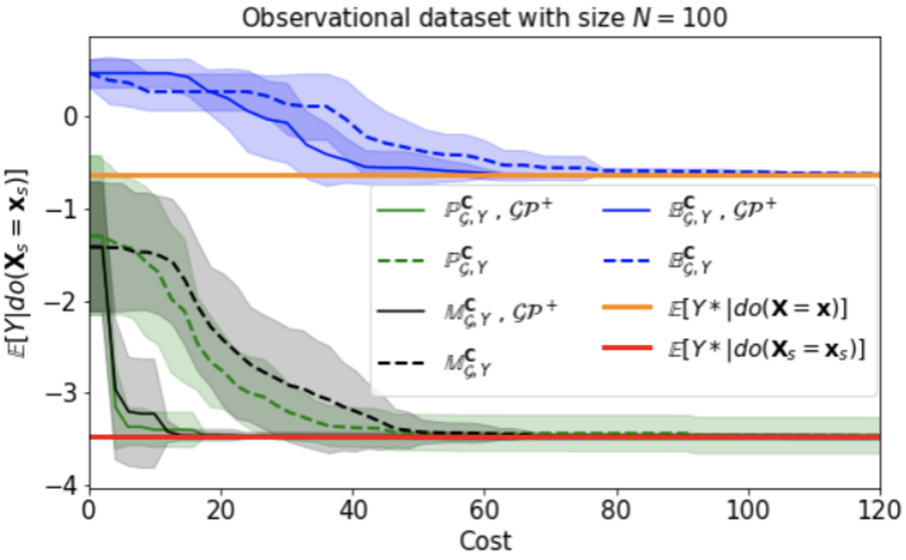

- BO is slower and identifies a suboptimal intervention
- CBO achieves the best result when using the Causal GP model

CBO for healthcare

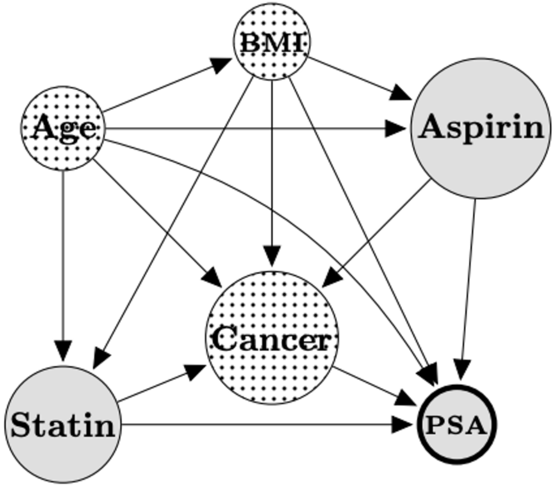

Decide whether to intervene on Statin, Aspirin or both and select the best intervention level in order to minimize PSA.

We found the optimal intervention to be **({aspirin, statin}, (0.0, 1.0))**.

Take home messages

- Standard BO ignores causal assumptions.
- Causal Global Optimization requires a new approach which we call CBO.
- CBO avoids exploring all possible sets.
- CBO merges observational and interventional data via the Causal GP prior.
- CBO solves both the exploration-exploitation and the observation-intervention trade off.

CBO limitations

- The number of GPs we require is determined by the number of sets to explore which is potentially huge.
- We don't transfer interventional information across GPs e.g. we don't account for the fact that intervening on X might give us some information about an intervention on X and Z

→ **The DAG-GP framework**

- We do not account for time and dynamic changes in the causal effects.

→ **Dynamic CBO**

CBO limitations

Virginia Aglietti, Theodoros Damoulas, Mauricio Álvarez, and JavierGonzález. Multi-task Causal Learning with Gaussian Processes. In *Neural Information Processing Systems* (NeurIPS), volume 33, pages 6293–6304.PMLR, 2020a.

- We do not account for time and dynamic changes in the causal effects.

Dynamic CBO

Virginia Aglietti, Neil Dhir, Javier González, and Theodoros Damoulas.Dynamic Causal Bayesian Optimization. In *Neural Information Processing Systems* (NeurIPS) 2021.

Dynamic Causal Bayesian Optimization

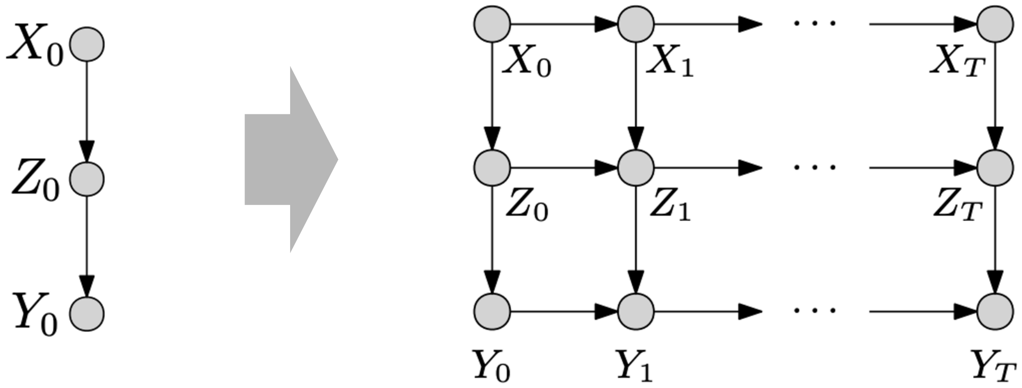

$$\mathbb{E}[Y_0 | \mathrm{do}(X_0 = x)] \neq$$
$$\neq \mathbb{E}[Y_1 | \mathrm{do}(X_1 = x)]$$
$$\neq \mathbb{E}[Y_2 | \mathrm{do}(X_2 = x)]$$

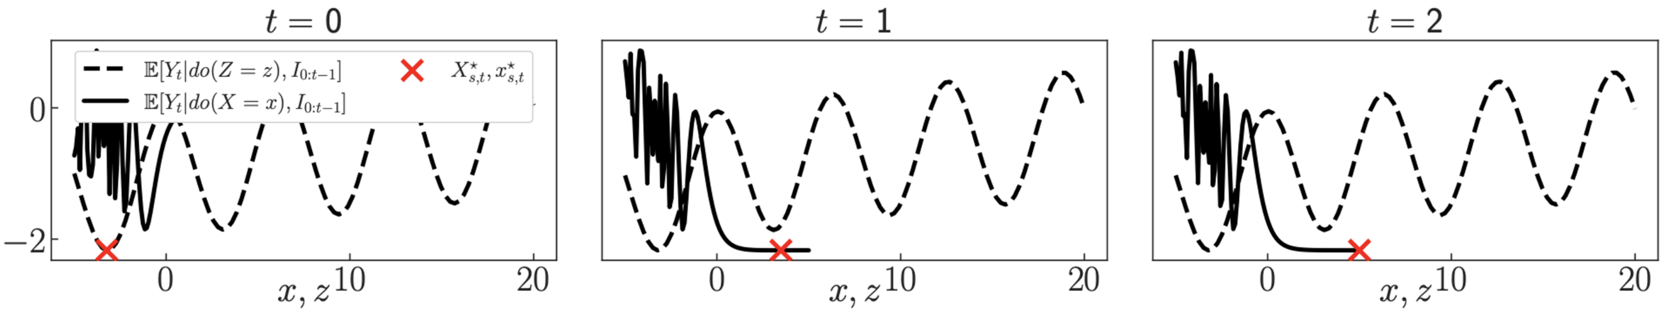

Dynamic Causal Bayesian Optimization

$$\mathbf{X}^\star_{s,t}, \mathbf{x}^\star_{s,t} = \underset{\substack{\mathbf{X}_{s,t} \in \mathcal{P}(\mathbf{X}_t) \\ \mathbf{x}_{s,t} \in D(\mathbf{X}_{s,t})}}{\arg\min} \mathbb{E}[Y_t \mid \mathrm{do}\left(\mathbf{X}_{s,t} = \mathbf{x}_{s,t}\right), \mathbb{1}_{t>0} \cdot I_{0:t-1}]$$

For every time step:

- Optimal intervention set
- Optimal Intervention level

Exploring all possible interventions in a causal graph.

Accounting for previous interventions.

- Refine a **recursion linking causal effects** across time steps thus allowing to share interventional information.
- **Construct a surrogate model** incorporating observational and interventional data, both at the current time step and at previous time steps.

Dynamic Causal Bayesian Optimization

DAG	M_t	

$\mathcal{G}_{t=0}$

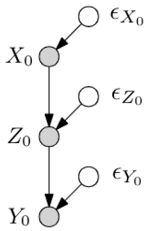

$$X_0 = f_{X_0}(\epsilon_{X_0})$$
$$Z_0 = f_{Z_0}(X_0, \epsilon_{Z_0})$$
$$Y_0 = f_{Y_0}(Z_0, \epsilon_{Y_0})$$

$$\mathbf{F}_0 = \{f_{X_0}, f_{Z_0}, f_{Y_0}\} \qquad \mathbf{C}_0 = \varnothing$$
$$\mathbf{U}_0 = \{\epsilon_{X_0}, \epsilon_{Z_0}, \epsilon_{Y_0}\} \qquad \mathbf{Y}_0 = Y_0$$
$$\mathbf{X}_0 = \{X_0, Z_0\}$$

$$\boxed{I_0 = (Z_0, z_0) = \underset{\substack{\mathbf{X}_{s,0} \in \mathcal{P}(\mathbf{X}_0),\\ \mathbf{x}_s \in D(\mathbf{X}_{s,0})}}{\arg\min} \; \mathbb{E}\left[Y_0 \mid \mathrm{do}(\mathbf{X}_{s,0} = \mathbf{x}_{s,0})\right]}$$

$\mathcal{G}_{t=1}$

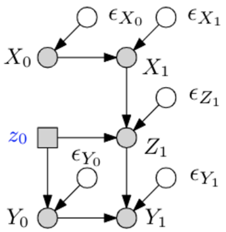

$$X_0 = f_{X_0}(\epsilon_{X_0})$$
$$Z_0 = f^I_{Z_0}(\cdot) = z_0$$
$$Y_0 = f_{Y_0}(z_0, \epsilon_{Y_0})$$
$$X_1 = f_{X_1}(X_0, \epsilon_{X_0})$$
$$Z_1 = f_{Z_1}(z_0, X_1, \epsilon_{Z_1})$$
$$Y_1 = f_{Y_1}(Y_0, Z_1, \epsilon_{Y_1})$$

$$\mathbf{F}_{0:1} = \{f_{X_0}, f^I_{Z_0}, f_{Y_0, X_1}, f_{Z_1}, f_{Y_1}\} \qquad \mathbf{C}_{0:1} = \{X_0, Y_0\}$$
$$\mathbf{U}_{0:1} = \{\epsilon_{X_0}, \epsilon_{Y_0}, \epsilon_{X_1}, \epsilon_{Z_1}, \epsilon_{Y_1}\} \qquad \mathbf{Y}_{0:1} = Y_1$$
$$\mathbf{X}_{0:1} = \{X_1, Z_1\}$$

$$\boxed{I_1 = (X_1, x_1) = \underset{\substack{\mathbf{X}_{s,1} \in \mathcal{P}(\mathbf{X}_1),\\ \mathbf{x}_s \in D(\mathbf{X}_{s,1})}}{\arg\min} \; \mathbb{E}\left[Y_1 \mid \mathrm{do}(\mathbf{X}_{s,1} = \mathbf{x}_{s,1}), I_0\right]}$$

$\mathcal{G}_{t=2}$

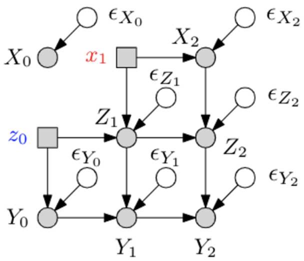

$$X_0 = f_{X_0}(\epsilon_{X_0})$$
$$Z_0 = f^I_{Z_0}(\cdot) = z_0$$
$$Y_0 = f_{Y_0}(z_0, \epsilon_{Y_0})$$
$$X_1 = f^I_{X_1}(\cdot) = x_1$$
$$Z_1 = f_{Z_1}(z_0, x_1, \epsilon_{Z_1})$$
$$Y_1 = f_{Y_1}(Y_0, Z_1, \epsilon_{Y_1})$$
$$X_2 = f_{X_2}(x_1, \epsilon_{X_2})$$
$$Z_2 = f_{Z_2}(Z_1, X_2, \epsilon_{Z_2})$$
$$Y_2 = f_{Y_2}(Y_1, Z_2, \epsilon_{Y_2})$$

$$\mathbf{F}_{0:2} = \{f_{X_0}, f^I_{Z_0}, f_{Y_0}, f^I_{X_1}, f_{Z_1}, f_{Y_1}, f_{X_2}, f_{Z_2}, f_{Y_2}\}$$
$$\mathbf{U}_{0:2} = \{\epsilon_{X_0}, \epsilon_{Y_0}, \epsilon_{Z_1}, \epsilon_{Y_1}, \epsilon_{X_2}, \epsilon_{Z_2}, \epsilon_{Y_2}\}$$
$$\mathbf{X}_{0:2} = \{X_2, Z_2\}$$
$$\mathbf{C}_{0:2} = \{X_0, Y_0, Z_1, Y_1\}$$
$$\mathbf{Y}_{0:2} = Y_2$$

$$\boxed{I_2 = (Z_2, z_2) = \underset{\substack{\mathbf{X}_{s,2} \in \mathcal{P}(\mathbf{X}_2),\\ \mathbf{x}_s \in D(\mathbf{X}_{s,2})}}{\arg\min} \; \mathbb{E}\left[Y_2 \mid \mathrm{do}(\mathbf{X}_{s,2} = \mathbf{x}_{s,2}), I_1, I_0\right]}$$

Dynamic Causal Global Optimization

- *Step (1):* Study the correlation among objective functions for two consecutive time steps and use it to derive a recursion formula that, based on the topology of the graph, **expresses the causal effects at time t as a function of previously implemented interventions.**
- *Step (2):* Develop a **new surrogate model** for the objective functions that can be used within a CBO framework to find the optimal sequence of interventions.

Assumptions

1. Invariance of causal structure: $\mathcal{G}(t) = \mathcal{G}(0), \forall t > 0$.
2. Additivity of $f_{Y_t}(\cdot)$ that is $Y_t = f_{Y_t}(\mathrm{Pa}(Y_t)) + \epsilon$ with $f_{Y_t}(\mathrm{Pa}(Y_t)) = f_Y^{Y}(Y_t^{\mathrm{PT}}) + f_Y^{\mathrm{NY}}(Y_t^{\mathrm{PNT}})$ where f_Y^{Y} and f_Y^{NY} are two generic unknown functions and $\epsilon \sim \mathcal{N}(0, \sigma^2)$.
3. Absence of unobserved confounders in $\mathcal{G}_{0:T}$.

Dynamic Causal Global Optimization

Assumptions

1. Invariance of causal structure: $\mathcal{G}(t) = \mathcal{G}(0), \forall t > 0$.
2. Additivity of $f_{Y_t}(\cdot)$ that is $Y_t = f_{Y_t}(\text{Pa}(Y_t)) + \epsilon$ with $f_{Y_t}(\text{Pa}(Y_t)) = f_Y^Y(Y_t^{\text{PT}}) + f_Y^{\text{NY}}(Y_t^{\text{PNT}})$ where f_Y^Y and f_Y^{NY} are two generic unknown functions and $\epsilon \sim \mathcal{N}(0, \sigma^2)$.
3. Absence of unobserved confounders in $\mathcal{G}_{0:T}$.

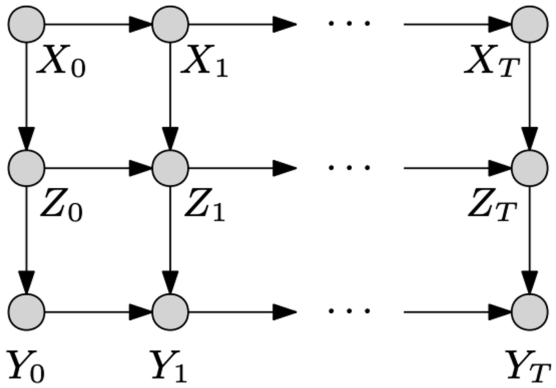

Example 1.

1. Same variables Y, X, Z at all time steps and edges oriented in the same way.
2. Functional relationship for Y at time step $t = 1$ is $Y_1 = f_Y^Y(Y_0) + f_Y^{NY}(Z_1)$.
3. No dashed edges.

Characterization of the time structure in a DAG with time dependent variables

Theorem: The Time Operator

Consider a DAG $\mathcal{G}_{0:T}$ and the related SEM satisfying the assumptions. $\forall \mathbf{X}_{s,t} \in \mathcal{P}(\mathbf{X}_t)$, the intervention function $f_{s,t}(\mathbf{x}) = \mathbb{E}[Y_t \mid \mathrm{do}(\mathbf{X}_{s,t} = \mathbf{x}), \mathbb{1}_{t>0} \cdot I_{0:t-1}]$ with $f_{s,t}(\mathbf{x}) : D(\mathbf{X}_{s,t}) \to \mathbb{R}$ can be written as:

$$f_{s,t}(\mathbf{x}) = \boxed{f_Y^Y(\mathbf{f}^\star)} + \boxed{\mathbb{E}_{p(\mathbf{w} \mid \mathrm{do}(\mathbf{X}_{s,t}=\mathbf{x}), \mathbf{i})}\left[f_Y^{\mathrm{NY}}(\mathbf{x}^{\mathrm{PY}}, \mathbf{i}^{\mathrm{PY}}, \mathbf{w})\right]} \qquad (2)$$

where $\mathbf{f}^\star = \{\mathbb{E}\left[Y_i \mid \mathrm{do}\left(\mathbf{X}^\star_{s,i} = \mathbf{x}^\star_{s,i}\right), I_{0:i-1}\right]\}_{Y_i \in Y_t^{\mathrm{PT}}}$ that is the set of previously observed optimal targets that are parents of Y_t. f_Y^Y denotes the function mapping Y_t^{PT} to Y_t and f_Y^{NY} represents the function mapping Y_t^{PNT} to Y_t.

Example 1.

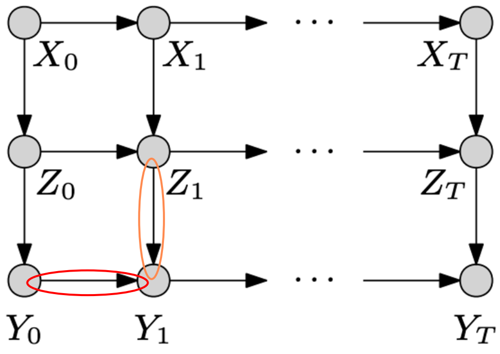

$$I^V_{0:t-1} = \{Z_0\}$$

$$\mathbb{E}[Y \mid \mathrm{do}(Z_1 = z), I_0]$$

$$= \boxed{f_Y^Y(y_0^\star)} + \boxed{f_Y^{\mathrm{NY}}(z)}$$

$$\mathbf{W} = \varnothing$$

Are the parents of Y that are not intervened nor previous targets.

Characterization of the time structure in a DAG with time dependent variables

Theorem: The Time Operator

Consider a DAG $\mathcal{G}_{0:T}$ and the related SEM satisfying the assumptions. $\forall \mathbf{X}_{s,t} \in \mathcal{P}(\mathbf{X}_t)$, the intervention function $f_{s,t}(\mathbf{x}) = \mathbb{E}[Y_t \mid \mathrm{do}(\mathbf{X}_{s,t} = \mathbf{x}), \mathbb{1}_{t>0} \cdot I_{0:t-1}]$ with $f_{s,t}(\mathbf{x}) : D(\mathbf{X}_{s,t}) \to \mathbb{R}$ can be written as:

$$f_{s,t}(\mathbf{x}) = \boxed{f_Y^Y(\mathbf{f}^\star)} + \boxed{\mathbb{E}_{p(\mathbf{w} \mid \mathrm{do}(\mathbf{X}_{s,t}=\mathbf{x}),\mathbf{i})}\left[f_Y^{\mathrm{NY}}(\mathbf{x}^{\mathrm{PY}}, \mathbf{i}^{\mathrm{PY}}, \mathbf{w})\right]} \qquad (2)$$

where $\mathbf{f}^\star = \{\mathbb{E}\left[Y_i \mid \mathrm{do}\left(\mathbf{X}^\star_{s,i} = \mathbf{x}^\star_{s,i}\right), I_{0:i-1}\right]\}_{Y_i \in Y_t^{\mathrm{PT}}}$ that is the set of previously observed optimal targets that are parents of Y_t. f_Y^Y denotes the function mapping Y_t^{PT} to Y_t and f_Y^{NY} represents the function mapping Y_t^{PNT} to Y_t.

Example 1.

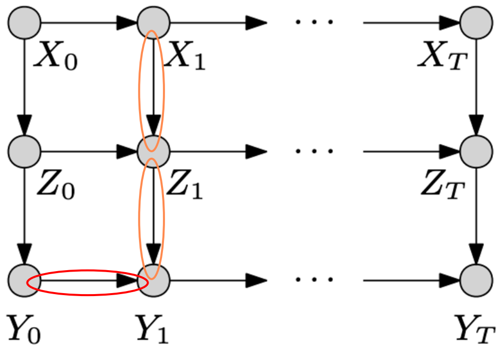

$$I^V_{0:t-1} = \{Z_0\}$$

$$\mathbb{E}[Y_1 \mid \mathrm{do}(X_1 = x), I_0]$$
$$= \boxed{f_Y^Y(y_0^\star)} + \boxed{\mathbb{E}_{p(z_1 \mid \mathrm{do}(X_1=x), I_0)}\left[f_Y^{\mathrm{NY}}(z_1)\right]}$$

$$\mathbf{W} = \{Z_1\}$$

Are the parents of Y that are not intervened nor previous targets.

The Dynamic Causal GP model

$$f_{s,t}(\mathbf{x}) = \mathbb{E}[Y_t \mid \mathrm{do}(\mathbf{X}_{s,t} = \mathbf{x})\,, 1_{t>0} \cdot I_{0:t-1}]$$

$$f_{s,t}(\mathbf{x}) \sim \mathcal{GP}(m_{s,t}(\mathbf{x}), k_{s,t}(\mathbf{x}, \mathbf{x}'))$$

$$m_{s,t}(\mathbf{x}) = \mathbb{E}\left[\boxed{f_Y^Y(\mathbf{f}^\star)} + \boxed{\widehat{\mathbb{E}}[f_Y^{NY}(\mathbf{x}^{PY}, \mathbf{i}^{PY}, \mathbf{w})]}\right]$$

$$k_{s,t}(\mathbf{x}, \mathbf{x}') = k_{\mathrm{rbf}}(\mathbf{x}, \mathbf{x}') + \sigma_{s,t}(\mathbf{x})\sigma_{s,t}(\mathbf{x}')$$

$$\text{with } \sigma_{s,t}(\mathbf{x}) = \sqrt{\mathbb{V}[f_Y^Y(\mathbf{f}^\star) + \boxed{\hat{\mathbb{E}}[f_Y^{NY}(\mathbf{x}^{PY}, \mathbf{i}^{PY}, \mathbf{w})]}}.$$

The surrogate model integrates observational data and interventional data at previous time steps in the prior.

Interventional data at the current time step in the posterior.

Dynamic Causal Bayesian Optimization

Algorithm: Causal Bayesian Optimization

Data: $\mathcal{D}^O$, $\mathcal{D}^I$, $\mathcal{G}$, **es**, number of steps T
Result: $\mathbf{X}_s^\star, \mathbf{x}_s^\star, \hat{\mathbb{E}}[\mathbf{Y}^\star|\mathrm{do}\,(\mathbf{X}_s^\star = \mathbf{x}_s^\star)]$
Initialise: Set $\mathcal{D}_0^I = \mathcal{D}^I$ and $\mathcal{D}_0^O = \mathcal{D}^O$

BO Loop

for t=1, ..., T **do**
  Compute ϵ and sample $u \sim \mathcal{U}(0,1)$
  if $\epsilon > u$ **then**
    (Observe)
    1. Observe new observations $(\mathbf{x}_t, c_t, \mathbf{y}_t)$.
    2. Augment $\mathcal{D}^O = \mathcal{D}^O \cup \{(\mathbf{x}_t, c_t, \mathbf{y}_t,)\}$.
    3. Update prior of the causal GP.
  end
  else
    (Intervene)
    1. Compute $EI^s(\mathbf{x})$ for each element $s \in$ **es**.
    2. Obtain the optimal interventional set-value pair $(s^\star, \alpha^\star)$.
    3. Intervene on the system.
    4. Update posterior of the interventional GP.
  end
end

Update prior distributions

Update posterior distribution

Algorithm 1: DCBO

Data: $\mathcal{D}^O$, $\{\mathcal{D}^I_{s,t=0}\}_{s\in\{0,\ldots,|\mathbb{M}_0|\}}$, $\mathcal{G}_{0:T}$, H.
Result: Optimal intervention path $\{\mathbf{X}^\star_{s,t}, \mathbf{x}^\star_{s,t}, y^\star_t\}_{t=1}^T$
Initialise: $\mathbb{M}$, $\mathcal{D}^I_0$ and initial optimal $\mathcal{D}^I_\star = \varnothing$.

Time steps

for $t = 0, \ldots, T$ **do**
  1. Initialise dynamic causal GP models for all $\mathbf{X}_{s,t} \in \mathbb{M}_t$ using $\mathcal{D}^I_{\star,t-1}$ if $t > 0$.
  2. Initialise interventional dataset $\{\mathcal{D}^I_{s,t}\}_{s\in\{0,\ldots,|\mathbb{M}_t|\}}$
  for $h = 1, \ldots, H$ **do**
    1. Compute $\mathrm{EI}_{s,t}(\mathbf{x})$ for each $\mathbf{X}_{s,t} \in \mathbb{M}_t$.
    2. Obtain $(s^\star, \alpha^\star)$
    3. Intervene and augment $\mathcal{D}^I_{s=s^\star,t}$
    4. Update posterior for $f_{s=s^\star,t}$
  end
  3. Return the optimal intervention $(\mathbf{X}^\star_{s,t}, \mathbf{x}^\star_{s,t})$
  4. Append optimal interventional data $\mathcal{D}^I_{\star,t} = \mathcal{D}^I_{\star,t-1} \cup ((\mathbf{X}^\star_{s,t}, \mathbf{x}^\star_{s,t}), y^\star_t)$
end

BO Loop

Causal decision-making in *dynamic* settings

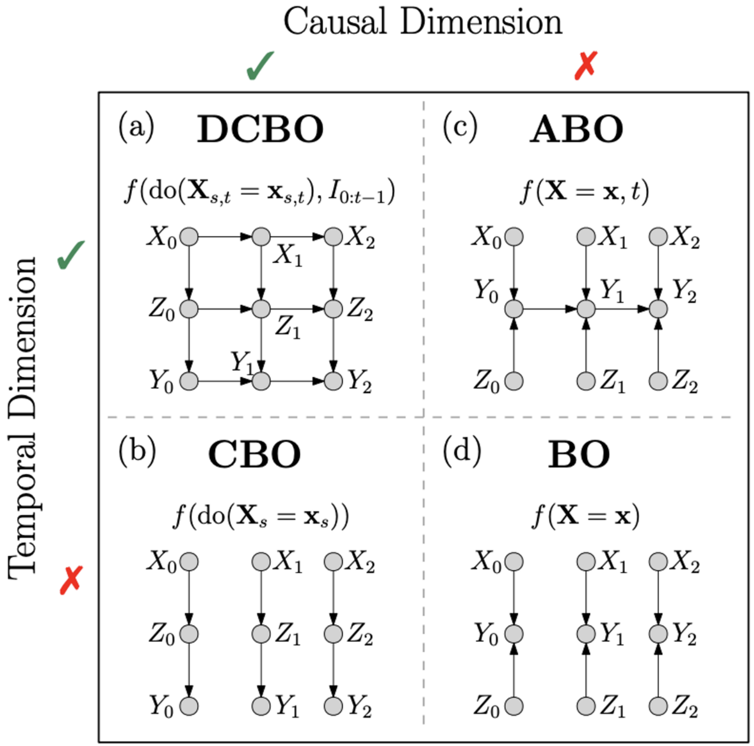

DAG representation of a *dynamic causal global optimisation problem* and the DAG considered when using CBO, ABO or BO to address the same problem.

The Dynamic Causal GP model

$X = \epsilon_X$

$Z = \exp(-X) + \epsilon_Z$

$Y = \cos(Z) - \exp(-\frac{Z}{20}) + \epsilon_Y$

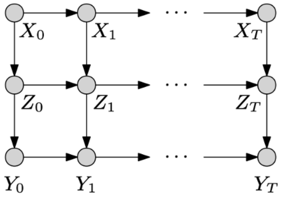

$X_t = x_{t-1} + \epsilon_X$

$Z_t = \exp(-X_t) + z_{t-1} + \epsilon_Z$

$Y_t = \cos(Z_t) - \exp\left(\frac{Z_t}{20}\right) + y_{t-1} + \epsilon_Y$

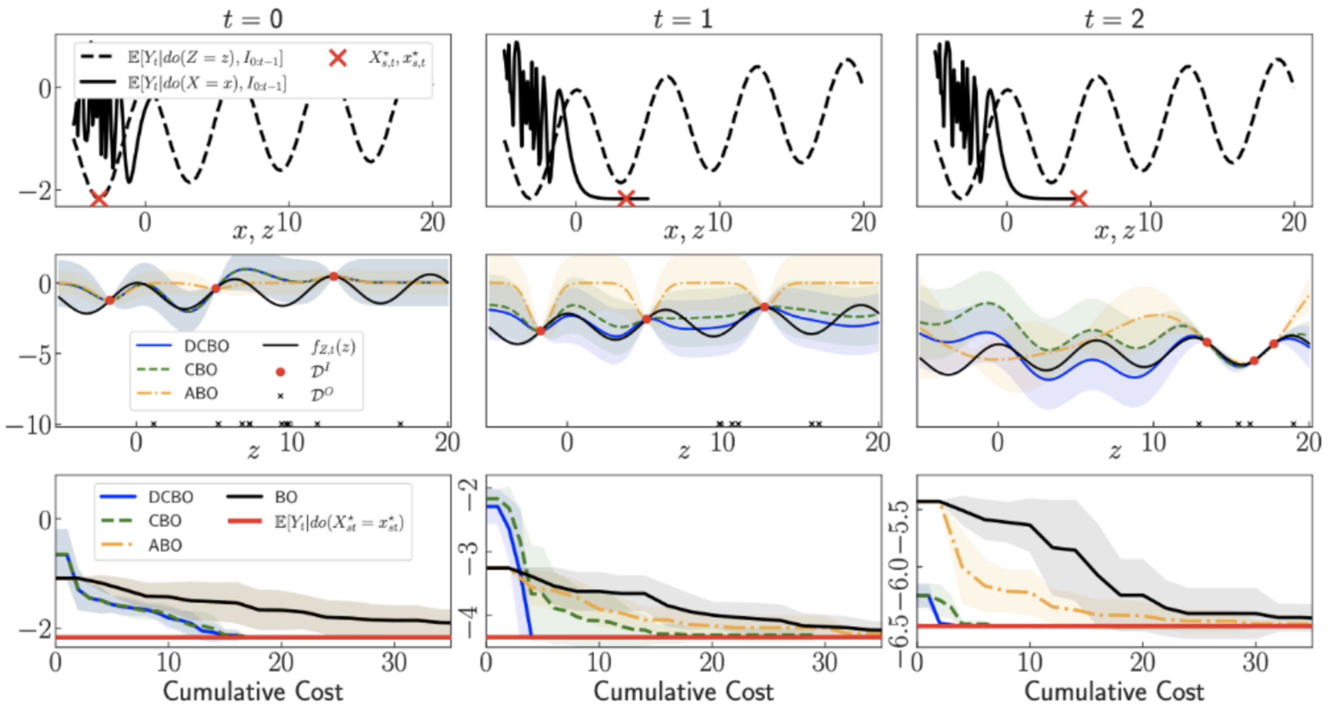

Experimental results

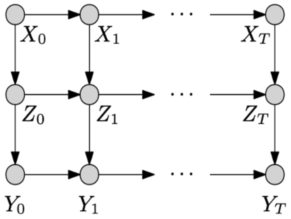

+

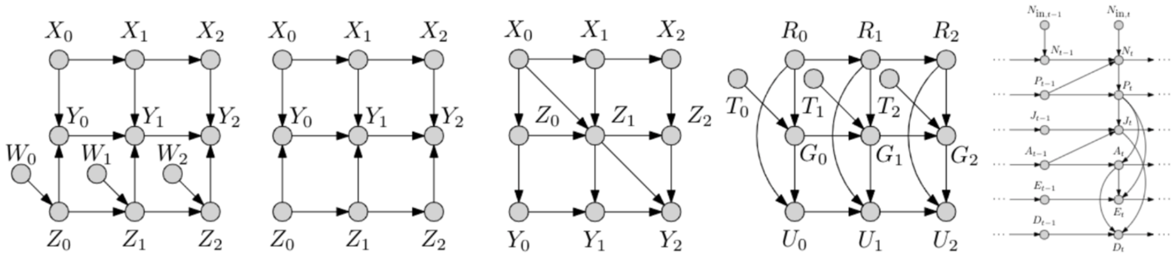

(a) MULTIV. (b) IND. (c) NONSTAT. (d) ECON. (e) ODE.

Figure 3: DAGs used in the experimental sections for the real (§4.2) and synthetic data (§4.1).

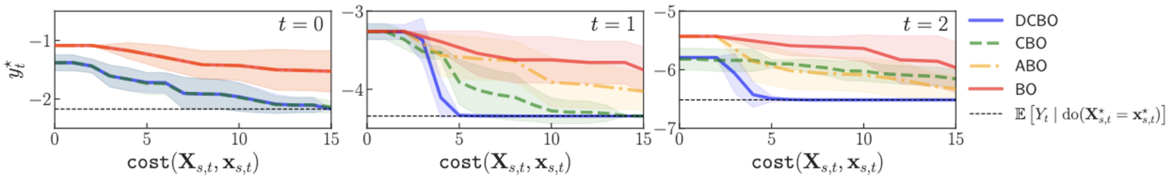

Figure 4: Convergence of dcbo and competing methods across replicates. The dashed black line (- - -) gives the optimal outcome $y_t^*, \forall t$. Shaded areas are $\pm$ one standard deviation.

Experimental results

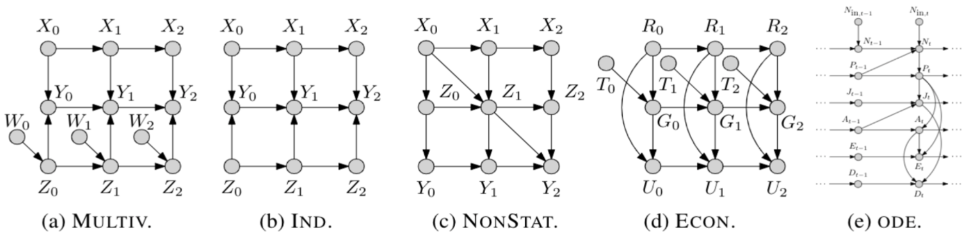

Figure 3: DAGs used in the experimental sections for the real (§4.2) and synthetic data (§4.1).

GAP metric

$$G_t = \left[\frac{y(\mathbf{x}^\star_{s,t}) - y(\mathbf{x}_{\text{init}})}{y^\star - y(\mathbf{x}_{\text{init}})} + \frac{H - H(\mathbf{x}^\star_{s,t})}{H}\right] \Big/ \left(1 + \frac{H-1}{H}\right)$$

Table 1: Average G_t across 10 replicates and time steps. See Fig. 1 for a summary of the baselines. Higher values are better. The best result for each experiment in bold. Standard errors in brackets.

	Synthetic data						Real data	
	STAT.	MISS.	NOISY	MULTIV.	IND.	NONSTAT.	ECON.	ODE
DCBO	**0.88** (0.00)	**0.84** (0.01)	**0.75** (0.00)	**0.49** (0.01)	0.48 (0.04)	**0.69** (0.00)	**0.64** (0.01)	**0.67** (0.00)
CBO	0.70 (0.01)	0.70 (0.02)	0.51 (0.02)	0.48 (0.09)	0.47 (0.07)	0.61 (0.00)	0.61 (0.01)	0.65 (0.00)
ABO	0.56 (0.01)	0.49 (0.02)	0.49 (0.04)	0.39 (0.21)	**0.54** (0.01)	0.38 (0.02)	0.57 (0.02)	0.48 (0.01)
BO	0.54 (0.02)	0.48 (0.03)	0.38 (0.05)	0.35 (0.08)	0.50 (0.01)	0.38 (0.03)	0.50 (0.01)	0.44 (0.03)

Experimental results

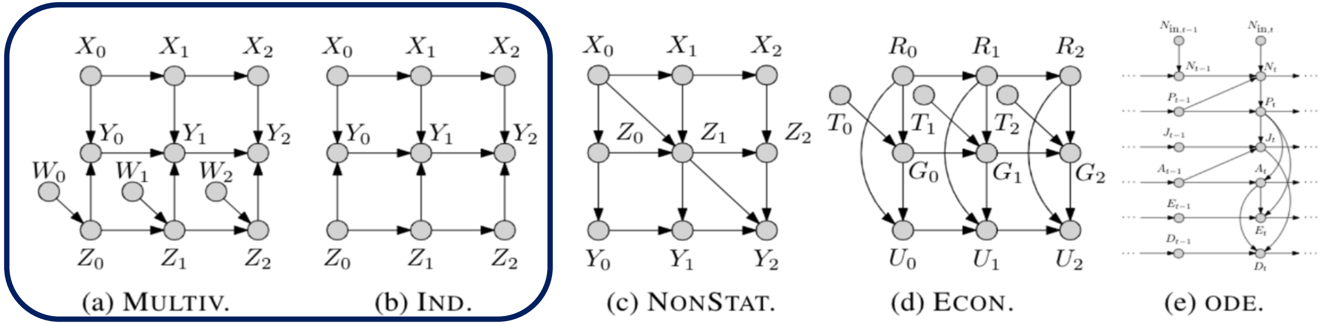

Figure 3: DAGs used in the experimental sections for the real (§4.2) and synthetic data (§4.1).

- **MULTIV.** : When the optimal intervention set is multivariate, both DCBO and CBO convergence speed worsen.

- **IND.** : Having to explore multiple intervention sets significantly penalises DCBO and CBO when there is no causal relationship among manipulative variables which are also the only parents of the target.

Experimental results

- **MULTIV.** : When the optimal intervention set is multivariate, both DCBO and CBO convergence speed worsen.

- **IND.** : Having to explore multiple intervention sets significantly penalises DCBO and CBO when there is no causal relationship among manipulative variables which are also the only parents of the target.

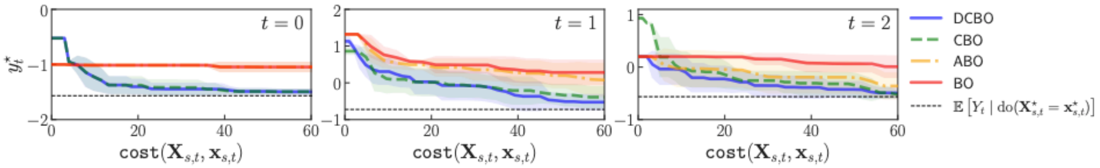

Figure 8: Experiment MULTIV.. Convergence of DCBO and competing methods across replicates. The red line gives the optimal $y_t^*, \forall t$. Shaded areas are $\pm$ standard deviation.

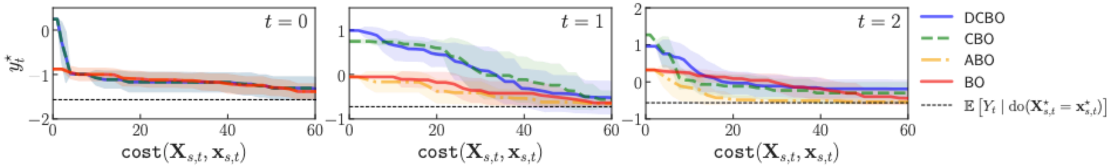

Figure 9: Experiment IND. Convergence of DCBO and competing methods across replicates. The red line gives the optimal $y_t^*, \forall t$. Shaded areas are $\pm$ standard deviation.

Experimental results

We repeat all experiments in the paper allowing the algorithms to perform a **lower number of trials at every time steps**. For t > 0, when moving to step t the convergence of the algorithm at step t − 1 is not guaranteed. This affect the optimum value that the algorithm can reach at subsequent steps.

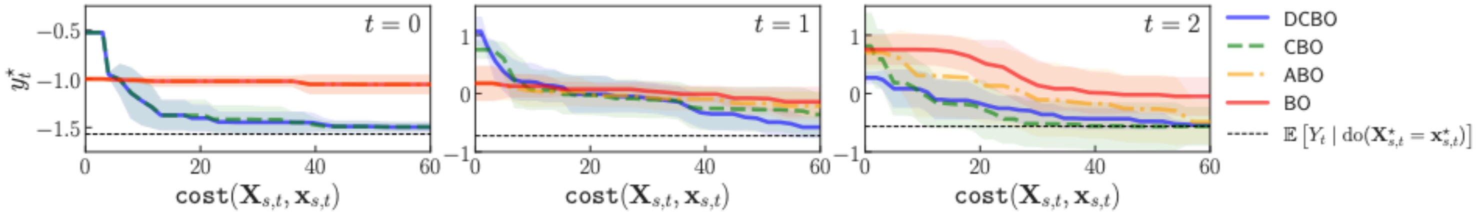

Figure 15: Experiment MULTIV. with maximum number of trials $H = 30$. Convergence of DCBO and competing methods across replicates. The black line gives the optimal y_t^*, $\forall t$. Shaded areas are $\pm$ one standard deviation.

Take home messages

- Identifying an optimal intervention at every time step requires solving a Dynamic Causal Global Optimization.
- DCBO solves the Dynamic Causal Global Optimization problem.
- DCBO proposes a surrogate model integrating all available data across time steps thus identifying interventions faster than CBO in dynamic settings.

Future research directions …

- Multi-objective causal BO to jointly maximize different interventional functions or deal with multi-dimensional outputs.
- A non-myopic causal BO to be used in dynamical systems where interventions performed at one time step affect the rewards an agent can obtain at future time steps.
- Causal BO to deal with discrete outputs and more generally non-Gaussian likelihoods.
- Connection between Causal RL, Causal Bandits and Causal BO.
- CBO for ITE and individual decision making.
- Offline/Off policy CBO.

and many more ….

Thank You!

Causal Bayesian Optimization

Virginia Aglietti
University of Warwick
The Alan Turing Institute
V.Aglietti@warwick.ac.uk

Xiaoyu Lu
Amazon
Cambridge, UK
luxiaoyu@amazon.com

Andrei Paleyes
Amazon
Cambridge, UK
paleyes@amazon.com

Javier González
Amazon
Cambridge, UK
gojav@amazon.com

Abstract

This paper studies the problem of globally optimizing a variable of interest that is part of a causal model in which a sequence of interventions can be performed. This problem

manipulating variables in order to optimize an outcome of interest. For instance, in strategic planning, companies need to decide how to allocate scarce resources across different projects or business units in order to achieve performance goals. In biology, it is common to change the phenotype of organisms by acting on individ-

Dynamic Causal Bayesian Optimization

Virginia Aglietti*
University of Warwick
The Alan Turing Institute
V.Aglietti@warwick.ac.uk

Neil Dhir*
The Alan Turing Institute
ndhir@turing.ac.uk

Javier González
Microsoft Research Cambridge
Gonzalez.Javier@microsoft.com

Theodoros Damoulas
University of Warwick
The Alan Turing Institute
T.Damoulas@warwick.ac.uk

Abstract

This paper studies the problem of performing a sequence of optimal interventions in a causal dynamical system where both the target variable of interest and the inputs evolve over time. This problem arises in a variety of domains e.g. system biology and operational research. Dynamic Causal Bayesian Optimization (DCBO) brings together ideas from sequential decision making, causal inference and Gaussian process (GP) emulation. DCBO is useful in scenarios where all causal effects