Towards MFACBO: Multi-Fidelity Abstraction Causal Bayesian Optimization in the Context of the Abstraction-Fidelity Connection

Jakob Zeitler[1]

[1]Department of Statistics, University of Oxford

Abstract

We investigate a novel integration of Causal Bayesian Optimization (CBO, Aglietti et al. [2020]) and Multi-Fidelity Bayesian Optimization (MFBO, Poloczek et al. [2017]) at the intersection of Causal Abstraction. MFBO enables cost-effective exploration by pooling information from fidelities with differing costs, while CBO introduces structural assumptions through incorporating causal knowledge—particularly Directed Acyclic Graphs (DAGs) encoding intervention relationships that can enhance multi-fidelity optimization. This fusion, which we term Multi-Fidelity Abstraction Causal Bayesian Optimization (MFACBO), is expected to improve decision-making efficiency in resource-constrained settings, such as healthcare or physical simulation, by guiding both the selection of intervention sets and fidelity levels. At the core, we expect Causal Abstraction to formalize the relationship between CBO and MFBO, where different fidelities are assumed to exist on differing levels of abstraction. In return, MFACBO emerges as a data-driven method to learn approximate causal abstraction mappings. We will evaluate our approach through synthetic experiments and real-world inspired scenarios using the recently introduced Causal Chambers, with particular attention to fidelity correlation modeling and acquisition strategies. This work lays the foundation for a framework that integrates causal reasoning into fidelity-aware optimization, subsumed by the theory of causal abstraction.

1 INTRODUCTION

Optimization of expensive to evaluate and complex functions is a pressing issues across the sciences and industry such as in healthcare, simulations, hyperparamter optimisation and more. Consider the graph in Figure 1 which describes a healthcare problem evaluating the influence of statin drugs on the levels of prostate specific antigen (PSA). The optimization goal is to minimize PSA by intervening with optimal levels on each statin and aspirin. The remaining variables, Age, BMI and Cancer, are measured but cannot be intervened on. Especially in the healthcare setting patient outcomes can be measured at different levels. For example, a patient might be sent to the hospital for expensive daily blood tests for PSA. This high-fidelity measurement of the PSA outcome could be supplemented by low-fidelity measurements using a smartwatch, which will only give approximate measures of PSA, but at a much lower cost. This low-fidelity measurement will be a coarsening of the high-fidelity PSA, an approximate causal abstraction.

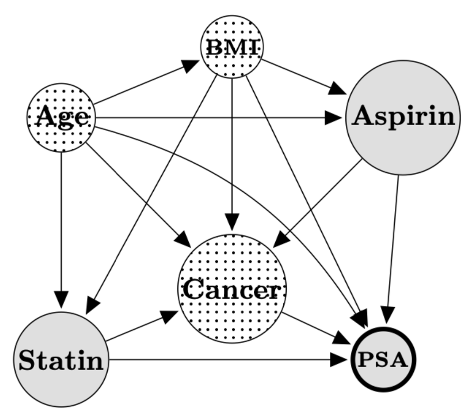

Figure 1: A common CBO example is the causal graph of PSA levels by Aglietti et al. [2020], where shaded nodes indicate variables that can be manipulated, and dotted nodes represent non-manipulable variables. The target variable, PSA, is shown as a thick shaded node.

In the remainder of this paper, we will start out with the most simple possible DAG in Figure 2 where X can be seen as the level of statin, Z as the level of aspirin and Y the

Accepted for the 1st Workshop on Causal Abstractions and Representations (CAR) at UAI 2025 (CAR 2025).

outcome PSA. Our goal will be to find the right intervention set and levels, such that the optimum is

$$\mathbf{X}_{\mathbf{s}}^*, \mathbf{x}_{\mathbf{s}}^*, = \underset{\mathbf{X}_s \in P(\mathbf{X}), \mathbf{x}_s \in D(\mathbf{X}_{\mathbf{s}})}{\text{argmin}} E_{P(Y_{high}|do(\mathbf{X}_{\mathbf{s}}=\mathbf{x}_{\mathbf{s}}))[Y_{high}]} \quad (1)$$

where $P(\mathbf{X})$ is all possible intervention sets, i.e. the power set of $\mathbf{X}$, and $D(\mathbf{X}_{\mathbf{s}} = \times_{X\in\mathbf{X}_{\mathbf{s}}}(D(\mathbf{X}))$ the domain of the intervention of $\mathbf{X}$, as defined in [Aglietti et al., 2020]. Furthermore, we also have access to one or more fidelities of the outcome $Y_{high-fidelity}$, e.g. $Y_{low-fidelity}$, often referred to as information sources. It will become evident that there are conceptual similarities between how we model differences between fidelities and how we model differences between causal abstractions. In the following sections we will

- **Introduce the key ingredients of MFACBO**: Bayesian Optimization, Causal Bayesian Optimization, Multi-fidelity Bayesian Optimization and Causal Abstraction
- **Attempt a problem statement** for MFACBO
- **Suggest initial steps** of a solution for MFACBO
- **Characterize** the *abstraction-fidelity connection* underpinning MFACBO
- **Present experimental design**, both synthetic and real-world, and their evaluation strategy
- **Conclude with a discussion** of next steps.

1.1 BAYESIAN OPTIMIZATION

Bayesian Optimization is a probabilistic, model-based optimization technique particularly well-suited for optimizing objective functions that are expensive to evaluate or black-box in nature. It operates by constructing a surrogate model—typically a Gaussian Process (GP) (Williams and Rasmussen [2006])—to approximate the true objective function, capturing both mean predictions and uncertainty. An acquisition function, such as Expected Improvement, Upper Confidence Bound, or Probability of Improvement, guides the selection of query points by balancing exploration and exploitation (Frazier [2018]). Bayesian Optimization has demonstrated strong performance in domains such as hyperparameter tuning for machine learning models and automated experiment design due to its sample efficiency and principled handling of uncertainty (Brochu et al. [2010]).

In vanilla Bayesian Optimization, no structural knowledge about the variables is used such that the graphical representation reduces to the DAG shown in Figure 2.

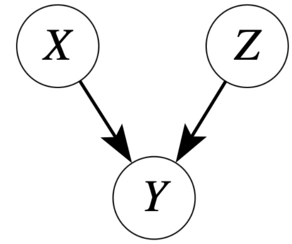

Figure 2: BO Example Graph

1.2 CAUSAL BAYESIAN OPTIMIZATION

In Causal Bayesian Optimization (CBO), we incorporate knowledge about the causal relationships of optimization variables. For example, we might have evidence that X is a parent of Z, and does not influence the outcome Y directly, but only through a mediating variable Z. This leaves us with four possible interventions sets

$\{\emptyset, \{X\}, \{Z\}, \{X, Z\}\}$

where $\emptyset$ indicates the observational state of the system, i.e. with no interventions. This is described in Figure 3, and extensively discussed in [Aglietti et al., 2020].

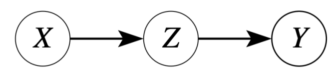

Figure 3: CBO Example graph

1.3 MULTI-FIDELITY BAYESIAN OPTIMIZATION

Multi-fidelity Bayesian Optimization (MFBO) extends traditional Bayesian Optimization to settings where evaluations of the objective function are available at multiple levels of cost and accuracy (i.e., fidelities). By leveraging cheaper, lower-fidelity approximations—such as coarser simulations, reduced training epochs, or subsampled datasets—MFBO can significantly reduce the overall optimization cost while maintaining accuracy. The core idea is for example to model the correlations between fidelities using techniques like multi-output Gaussian Processes (Kennedy and O'Hagan [2000]), enabling informed decisions about which fidelity to query at each iteration. Acquisition functions are adapted to account for both information gain and query cost. MFBO has been successfully applied in engineering design, hyperparameter tuning, and scientific experiments where high-fidelity evaluations are prohibitively expensive.

In this paper, we will follow the approach of multi-information source BO by Poloczek et al. [2017]. Here, we assume that there is access to M information sources, possibly biased and/or noisy, indexed by $l \in [M]_0$. We can query any IS l at design point x yielding i.i.d measurements with mean $f(l, x)$ and finite variance $\lambda_l(x)$, independent conditional on $f(l, x)$. These can also be referred to as auxililary tasks, while the primary task is considered to be function g. Therefore,

$$f(0, x) = g(x)$$

The bias of each IS is then $f(l, x) - g(x) = \delta_l$. Each IS is also associated with a cost function $c_l(x) : D \mapsto R$. We assume both $\lambda_l(x)$ and $c_l(x)$ to be known, e.g. as provided by domain experts.

1.4 CAUSAL ABSTRACTION

Causal abstraction concerns the study of causal systems at different levels of abstraction, from the fundamental micro-'atomistic' to the macro-'coarsened'. Beckers et al. [2019] attempt a formal formulation of approximate causal abstraction, assuming that abstract models will capture only ever capture the underlying fundamental system in an approximate way, as exemplified in Figure 4.

First, we very briefly revisit the notation of approximate causal abstraction as introduced by Beckers et al. [2019]. A signature $\mathcal{S}$ is a tuple $(\mathcal{U}, \mathcal{V}, \mathcal{R})$ where $\mathcal{U}$ is a set of exogenous variables, $\mathcal{V}$ is a set of endogenous variables and $\mathcal{R}$ is a function that associates the range to every variable $X_i \in \mathcal{U} \cup \mathcal{V}$, i.e. a non-empty set of possible values of X_i. A signature $\mathcal{S}$ is the basis for a causal model M which is a pair $(\mathcal{S}, \mathcal{F})$, where $\mathcal{F}$ defines a function for each X_i, i.e. a structural equation F_{X_i} that maps all possible values of $\mathcal{U}$ and $\mathcal{V}$ to X_i, i.e. F_{X_i} maps the range $\mathcal{R}(\mathcal{U} \cup \mathcal{V} - X_i)$ to $\mathcal{R}(X_i)$.

Central to this is the formulation of an abstraction function $\tau : \mathcal{R}_L(\mathcal{V}_\mathcal{L}) \mapsto \mathcal{R}_\mathcal{H}(\mathcal{V}_\mathcal{H})$ that maps endogenous states of the low-level model M_L to the endogenous states of M_H, where

$$\begin{aligned} M_L &= ((\mathcal{U}_\mathcal{L}, \mathcal{V}_\mathcal{L}, \mathcal{R}_\mathcal{L}), \mathcal{F}_\mathcal{L}, \mathcal{I}_\mathcal{L}) \\ M_H &= ((\mathcal{U}_\mathcal{H}, \mathcal{V}_\mathcal{H}, \mathcal{R}_\mathcal{H}), \mathcal{F}_\mathcal{H}, \mathcal{I}_\mathcal{H}) \end{aligned} \tag{2}$$

as defined in [Beckers et al., 2019]. For the PSA examples introduced earlier in Figure 1, the abstraction models might be defined as follows, where aspirin and statin are approximate by statin-drugs, and PSA is measured via skin instead of blood:

$$\begin{aligned} &M_L = \\ &= ((\mathcal{U}, (statin, aspirin, PSA_{blood}), \mathcal{R}_\mathcal{L}), (g_L, f_L), \mathcal{I}_\mathcal{L}) \\ &M_H = \\ &= ((\mathcal{U}, (statin - drugs, PSA_{skin}), \mathcal{R}_\mathcal{H}), (g_H, f_H), \mathcal{I}_\mathcal{H}) \end{aligned} \tag{3}$$

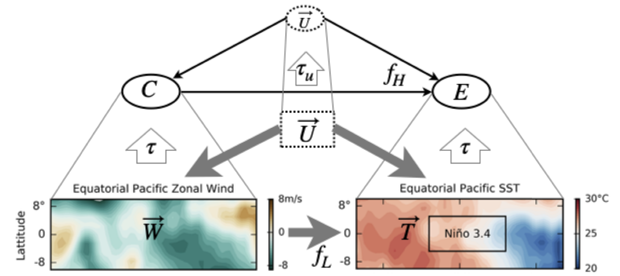

Figure 4: This climate example is apapted from Chalupka et al. [2016] by Beckers et al. [2019], where a high-level causal model of the El Niño phenomenon is derived from low-level (high-dimensional) measurements of wind W and sea surface temperature T. Here, U represents an unmeasured confounder, and τ denotes the mapping between the two models.

1.5 STRUCTURAL LEARNING, CAUSAL DISCOVERY & CAUSAL ENTROPY OPTIMISATION

In this paper, we will assume that the causal graph is known, either via expert knowledge or data-driven methods. CBO has been extended to also learn the structure of the graph as discussed in [Mamaghan et al., 2024, Branchini et al., 2022, Lorch et al., 2021, Tigas et al., 2023], but due to the increased model complexity require much more data while also being numerically more instable. Notably, [Zhang et al., 2023] have already described a first step towards leveraging multiple fidelities for learning DAGs. We leave a formulation that simultaneously learns the graphs as a future step. It's also worth mentioning work that attempts to learn abstraction maps, as we will introduce later, with out the knowledge of the causal graph, as dicussed in Felekis et al. [2024].

The work closest to our paper tackles the problem of "Causally Abstracted Multi-armed Bandits" [Zennaro et al., 2024] with their point of departure not rooted in Bayesian optimization as in [Aglietti et al., 2020], but the bandit literature which first tackles causal problems in [Lattimore et al., 2016].

2 PROBLEM STATEMENT

As a starting point for MFACBO, we consider the integration of CBO and MFBO, as described in the DAG of Figure 5. Here, we assume that a low-fidelity outcome measurement Y_{low} is a causal descendant of a high-fidelity outcome Y_{high}. In the problem setting of MFBO, we would expect a δ difference between these two information sources, often referred to as 'model discrepancy'. In terms of causal abstraction, we consider Y_{low} to be a coarsening Y_{high} that is cheaper to inquire, and, easily confused, Y_{low} is a *higher abstraction* of the base causal model Y_{high}.

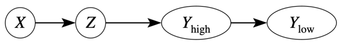

Figure 5: MFACBO setup where the low-fidelity is a direct consequence of the high-fidelity outcome.

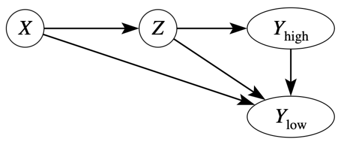

Figure 6: MFACBO DAG where the low-fidelity is both a causal consequence of the intervention variables and the high-fidelity outcome.

There are also other possible DAGs for MFACBO, such as show in Figure 7, where we consider the possibility of low and high fidelity interventions for X and Z.

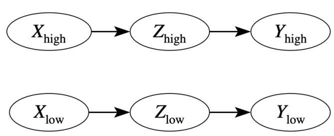

Figure 7: MFACBO graph where there are also differing fidelities of intervention variables. It is unclear which DAG representation is most adequate to model fidelity-levels of this kind, and whether there should be edges between level outcomes and/or interventions.

Related to that is also the possibility to treat each possible intervention set as its individual fidelity, e.g. such that for $\{\emptyset, \{X\}, \{Z\}, \{X, Z\}\}$ there exists fidelities $IS_{\emptyset}, IS_X, IS_Z, IS_{X,Z}$.

Finally, there is also a connection to recent work on functional networks and Bayesian optimization as discussed by Astudillo and Frazier [2021]. This case is shown in Figure 8 and a possible extension with soft interventions on M could be derived from [Massidda et al., 2023].

3 SOLUTION

As a first step, we will approach the integration of CBO and MFBO assuming the DAG structure shown in Figure 6, i.e. we assume that there is a high-fidelity causal process $IS_0 = f(0, x) = g(x) = Y_{high}(x)$ that is expensive to intervene on and/or measure, and a low-fidelity causal process $IS_1 = f(1, x) = g(x) - \delta_1(x) = Y_{low}(x)$ that is cheap, but approximate, and costs $c_0(x) = 5 \;\; c_1(x) = 1$.

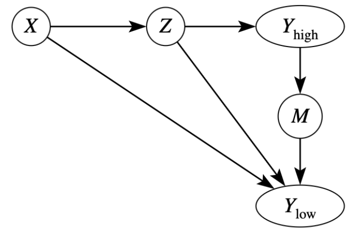

Figure 8: Functional networks Bayesian optimization considers intermediate outcomes which in our case we could consider Y_{high} to be, including a (soft)-intervention variable M governing the relationship of the high and low fidelity.

3.1 MODEL DEFINITION

For each intervention set in the minimal intervention sets, e.g. $\{\emptyset, \{X\}, \{Z\}, \{X, Z\}\}$ for Figure 5 we will continue with the MISO model design as introduced before in [Poloczek et al., 2017]: We define our surrogate model for f as a:

- Gaussian process prior on f
- with mean functions $\mu : [M] \times D \mapsto R$
- and covariance kernels $\Sigma : ([M] \times D)^2 \mapsto R$

To model the information source discrepancies, we assume that for each $l > 0$ a function $\delta_l : D \mapsto R$ was drawn from a separate GP, $\delta_l \sim GP(\mu_l, \Sigma_l)$ where the primary task $\delta_0 = 0$. This generative model δ_l models the bias $f(l, x) - g(x)$ for all IS l. Therefore,

$$f(l, x) = f(0, x) + \delta_l(x)$$

Furthermore, $g \sim GP(\mu_0, \Sigma_0)$, and also:

$$\mu(l, x) = E[f(l, x)] = E[g(x)] + E[\delta_l(x)] = \mu_0(x)$$

$$\begin{aligned} \Sigma((l, x), (m, x')) &= \\ = Cox(g(x) + \delta_l(x), g(x') + \delta_m(x')) &= \\ = \Sigma_0(x, x') + 1_{l,m} \times \Sigma_l(x, x') \end{aligned} \tag{4}$$

3.2 ACQUISITION FUNCTIONS

In MISO (Poloczek et al. [2017]) both the design $x \in D$ and information source $l \in [M]_0$ are selected at the beginning of each round. For each IS, the value of information at design point x is the expected gain in quality of the best design given all samples measured so far. This gain is normalized using the cost function $c_l(x)$, yielding which IS l should be sampled at what design point x. If costs are identical, this means we would query the true objective $g(x)$ at every turn. This can be formalized into the follow acquisition function:

$$MKG^n(l,x) = E[\frac{max_{x' \in D}\mu^{(n+1)}(0,x') = max_{x' \in D}\mu^{(n)}(0,x')}{c_l(x)}| \quad |l^{(n+1)} = l, x^{(N=1)} = x] \tag{5}$$

In CBO (Aglietti et al. [2020]) the acquisition function is defined for each intervention set. The expected improvement (EI) is:

$$EI^s(x) = E_{p(y_s)}[max(y_s - y*)]/Co(x) \tag{6}$$

As a next step in future work, we will combine these acquisition functions, such that they inform and predict the next best query for maximal gain, in regard to which fidelity IS_l to query over which intervention set.

4 THE ABSTRACTION-FIDELITY CONNECTION

Through the above exposition, it becomes evident that there is a conceptual connection between multi-fidelity optimization incorporating causal knowledge and recent work in causal abstraction [Beckers et al., 2019]. We identify the parallels in Table 1.

MFBO per [Poloczek et al., 2017] models the model discrepancy between fidelities, s.t. $\delta_l(x) = f(l,x) - f(0,x)$. Once learned from data, we can access approximations of lower and higher fidelities via shifts of $\delta_l(x)$. Similarly, in causal abstraction we define τ as a function to map between abstraction levels. To avoid confusion it is important to note that in MFBO high-fidelity information sources are considered to be the most precise and therefore most costly to query, whereas in causal abstraction high-level models are considered to be approximation, and therefore are equivalent to low-fidelities and cheaper to query.

Furthermore, MFBO is not necessarily restricted to problems with a ranking of information sources, but also allows for problems where relationships of information sources aren't necessarily ranked. [Beckers et al., 2019] also introduce distance measures to quantify the approximation loss of abstract models from their low-level underpinnings. Such a measure, d_v, might be conceptually closer to MFBO's δ_l than the abstraction function τ, which will be part of further inquiry.

Example

A brief worked out example is shown in Table 2, where we take the synthetic test function of Equation 9 and show the difference between the low and high abstraction level. Formally, this model or fidelity discrepancy equates to

$$\delta_1 = sin(2Z_H) = d(IS_0, IS_1) = d(M_L, M_H) = \tag{7}$$

$$= \sqrt{0 \times (X_L - X_H)^2 + 0 \times (Z_L - Z_H) + 1 \times (Y_L - Y_H)} \tag{8}$$

where we use the most simple definition of approximate causal abstraction distance, i.e. weighted squared distance, and only focus on the difference of outcomes Y_l.

5 EXPERIMENTS

5.1 SYNTHETIC EXPERIMENTS

We will run experiments on the toy example introduced by Aglietti et al. [2020] in their seminal CBO paper, i.e. using the DAG in Figure 6. We will extend its associated structural equations loosely based on the MFBO example in Poloczek et al. [2017] such that:

$$\begin{aligned} X &= \epsilon_X \\ Z &= exp(-X) + \epsilon_Z \\ Y_{high} &= cos(Z) - exp(-\frac{Z}{20}) + \epsilon_Z \\ Y_{low} &= Y_{high} + sin(2Z) \end{aligned} \tag{9}$$

Experiments will be run to determine the influence of different acquisition functions and surrogate models, and their influence on convergence performance such as cumulative regret, also in the context of CBO. We choose random design generation as a baseline strategy to outperform, also comparing against off-the-shelf algorithms such as single fidelity BO, CBO, MFBO. Finally, we run MFACBO, which we hypothesize should outperform both CBO and MFBO individually as it has access to both differing fidelities as well as the causal graph, suggesting information theoretic advantage.

5.2 REAL-WORLD EXPERIMENTS

Synthetic experiments have several limitations in terms of assessing real-world performance of Bayesian Optimization. Therefore, we will also run experiments on the recently introduced Causal Chambers (CC, Gamella et al. [2025]). The CCs are built upon physical mechanisms where the underlying causal DAG is well known and studied, allowing us to measure close-to ground truth data. We will focus on the

	Abstraction	MFBO
Entities	High and low levels M_H, M_L	High and low fidelities IS_0, IS_1
Distance	$d_{max}(M_1, M_2) = max_{X \leftarrow x \in \mathcal{I}, u \in \mathcal{R}(\mathcal{U})}\{d_v(M_1(u, X \leftarrow x), M_2(u, X \leftarrow x))\}$	$\delta_l(x) = f(l, x) - f(0, x)$

Table 1: Causal Abstraction and Multi-fidelity Bayesian Optimization in terms of their different entities of study and distance between those entities.

IS	X_l	Z_l	Y_l
0 (Low)	ϵ_{X_L}	$\exp(-X_L) + \epsilon_{Z_L}$	$cos(Z_L) - exp(-\frac{Z_L}{20}) + \epsilon_{Y_L}$
1 (High)	ϵ_{X_H}	$\exp(-X_H) + \epsilon_{Z_H}$	$cos(Z_H) - exp(-\frac{Z_H}{20}) + \epsilon_{Y_H} + \underline{sin(2Z_H)}$

Table 2: Comparison table showcasing the model discrepancy (underlined) in the structural equations from the low fidelity to high fidelity information source (i.e. the abstraction error from the low model to the high abstract model). Formally, this discrepancy is defined as $\delta_1 = sin(2Z_H)$.

wind tunnel experiment, shown in Figure 9, where have control over the resolution of the sensors, representing differing levels of fidelity, as seen in Figure 11. The wind tunnel can be described via a DAG shown in Figure 10, where have parent variables such as fan load (L_{in}, L_{out}) and their causal children, air pressure $(\tilde{P}_{dw}, \tilde{P}_{up})$. As we do in the synthetic experiments, we choose random as a baseline to outperform, and compare against off-the-shelf algorithms such as single fidelity BO, CBO, MFBO. Then, we run MFACBO, which we hypothesize should outperform both CBO and MFBO individually as it has access to both different levels of information sources via the sensor resolutions, as well as the causal graph. Its information theoretic advantage should help it outperform both CBO and MFBO individually.

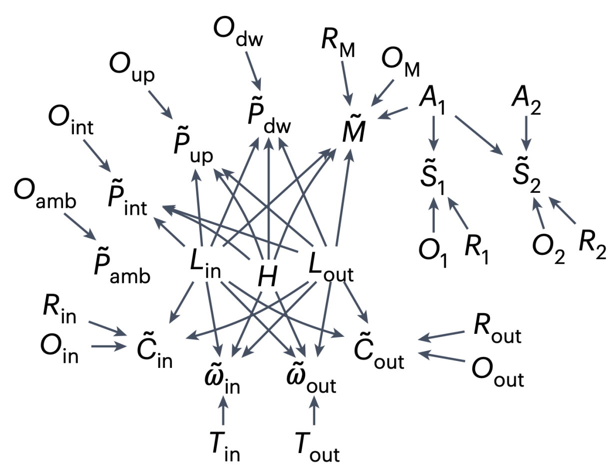

Figure 10: This DAG represents the control and measurement variables, and will be assumed to be true for the benchmarking of MFACBO.

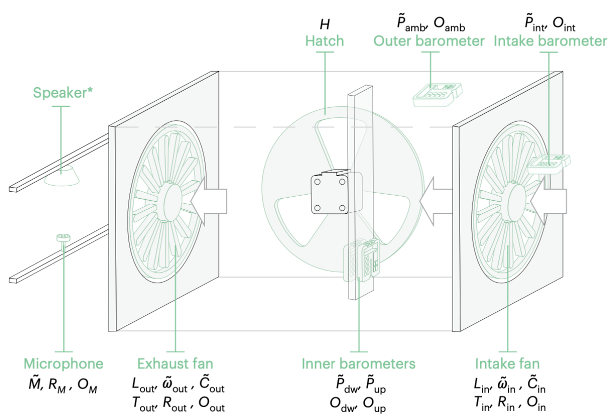

Figure 9: The wind tunnel CC is a real-world simulation of air pressure differences regulated by fan speed and pressure valves, see Gamella et al. [2025]

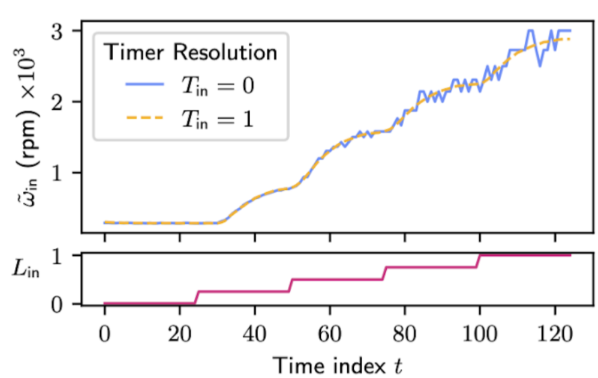

Figure 11: Causal Chambers allow for differing fidelities by varying the resolution of its sensors, as in their Figure 8 (a).

6 DISCUSSION

In this workshop paper, we have introduced initial ideas for the description of a new method called MFACBO, that integrates both causal graphical knowledge via CBO as well as fidelity pooling via MFBO. We state possible model choices and acquisition functions with key steps still open to be explored. We also state our experimental strategy, both synthetic using established examples, as well as real-world facilitated via causal chambers. Most importantly, we expect significant conceptual overlap of CBO and MFBO to be characterized via recent work in causal abstraction. Here, the exact connection of causal abstraction and multi-fidelity modeling will be found at the intersection of model discrepancy δ_l, causal abstraction mapping τ and approximation error distance d_v, see Table 1. In return, MFACBO emerges as a method for data-driven learning of abstraction mappings τ in optimisation settings.

References

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

Raul Astudillo and Peter I. Frazier. Bayesian optimization of function networks. *CoRR*, abs/2112.15311, 2021. URL https://arxiv.org/abs/2112.15311.

Sander Beckers, Frederick Eberhardt, and Joseph Y. Halpern. Approximate causal abstraction, 2019. URL https://arxiv.org/abs/1906.11583.

Nicola Branchini, Virginia Aglietti, Neil Dhir, and Theodoros Damoulas. Causal entropy optimization, 2022. URL https://arxiv.org/abs/2208.10981.

Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

Krzysztof Chalupka, Tobias Bischoff, Pietro Perona, and Frederick Eberhardt. Unsupervised discovery of el nino using causal feature learning on microlevel climate data, 2016. URL https://arxiv.org/abs/1605.09370.

Yorgos Felekis, Fabio Massimo Zennaro, Nicola Branchini, and Theodoros Damoulas. Causal optimal transport of abstractions. In Francesco Locatello and Vanessa Didelez, editors, *Proceedings of the Third Conference on Causal Learning and Reasoning*, volume 236 of *Proceedings of Machine Learning Research*, pages 462–498. PMLR, 01–03 Apr 2024. URL https://proceedings.mlr.press/v236/felekis24a.html.

Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

Juan L Gamella, Jonas Peters, and Peter Bühlmann. Causal chambers as a real-world physical testbed for ai methodology. *Nature Machine Intelligence*, pages 1–12, 2025.

Marc C Kennedy and Anthony O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.

Finnian Lattimore, Tor Lattimore, and Mark D Reid. Causal bandits: Learning good interventions via causal inference. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/b4288d9c0ec0a1841b3b3728321e7088-Paper.pdf.

Lars Lorch, Jonas Rothfuss, Bernhard Schölkopf, and Andreas Krause. Dibs: Differentiable bayesian structure learning, 2021. URL https://arxiv.org/abs/2105.11839.

Amir Mohammad Karimi Mamaghan, Panagiotis Tigas, Karl Henrik Johansson, Yarin Gal, Yashas Annadani, and Stefan Bauer. Challenges and considerations in the evaluation of bayesian causal discovery, 2024. URL https://arxiv.org/abs/2406.03209.

Riccardo Massidda, Atticus Geiger, Thomas Icard, and Davide Bacciu. Causal abstraction with soft interventions. In Mihaela van der Schaar, Cheng Zhang, and Dominik Janzing, editors, *Proceedings of the Second Conference on Causal Learning and Reasoning*, volume 213 of *Proceedings of Machine Learning Research*, pages 68–87. PMLR, 11–14 Apr 2023. URL https://proceedings.mlr.press/v213/massidda23a.html.

Matthias Poloczek, Jialei Wang, and Peter Frazier. Multi-information source optimization. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/df1f1d20ee86704251795841e6a9405a-Paper.pdf.

Panagiotis Tigas, Yashas Annadani, Desi R. Ivanova, Andrew Jesson, Yarin Gal, Adam Foster, and Stefan Bauer. Differentiable multi-target causal Bayesian experimental design. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 34263–34279. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/tigas23a.html.

Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

Fabio Massimo Zennaro, Nicholas Bishop, Joel Dyer, Yorgos Felekis, Anisoara Calinescu, Michael Wooldridge, and Theodoros Damoulas. Causally abstracted multi-armed bandits, 2024. URL https://arxiv.org/abs/2404.17493.

Zeyu Zhang, Chaozhuo Li, Xu Chen, and Xing Xie. Bayesian active causal discovery with multi-fidelity experiments. *Advances in Neural Information Processing Systems*, 36:63957–63993, 2023.