Extending Multi-Source Bayesian Optimization With Causality Principles

Extended Abstract

Luuk Jacobs
Radboud University
Nijmegen, The Netherlands
luuk.jacobs@ru.nl

Mohammad Ali Javidian
Appalachian State University
Boone, United States
javidianma@appstate.edu

ABSTRACT

Multi-Source Bayesian Optimization (MSBO) serves as a variant of the traditional Bayesian Optimization (BO) framework applicable to situations involving optimization of an objective black-box function over multiple information sources such as simulations, surrogate models, or real-world experiments. However, traditional MSBO assumes the input variables of the objective function to be independent and identically distributed, limiting its effectiveness in scenarios where causal information is available and interventions can be performed, such as clinical trials or policy-making. In the single-source domain, Causal Bayesian Optimization (CBO) extends standard BO with the principles of causality, enabling better modeling of variable dependencies. This leads to more accurate optimization, improved decision making, and more efficient use of low-cost information sources. In this article, we propose a principled integration of the MSBO and CBO methodologies in the multi-source domain, leveraging the strengths of both to enhance optimization efficiency and reduce computational complexity in higher-dimensional problems. We present the theoretical foundations of both Causal and Multi-Source Bayesian Optimization, and demonstrate how their synergy informs our Multi-Source Causal Bayesian Optimization (MSCBO) algorithm. We compare the performance of MSCBO against its foundational counterparts for both synthetic and real-world datasets with varying levels of noise, highlighting the robustness and applicability of MSCBO. Based on our findings, we conclude that integrating MSBO with the causality principles of CBO facilitates dimensionality reduction and lowers operational costs, ultimately improving convergence speed, performance, and scalability.

KEYWORDS

Bayesian Optimization, Causality; Interventions; Multi-source information; Gaussian Process; Structural equation model; DAG

ACM Reference Format:
Luuk Jacobs and Mohammad Ali Javidian. 2026. Extending Multi-Source Bayesian Optimization With Causality Principles: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/JWIE9181

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/JWIE9181

1 THEORETICAL FRAMEWORK

In this paper, we introduce Multi-Source Causal Bayesian Optimization (MSCBO), a framework for efficient decision making in systems that must operate under limited budgets while reasoning over heterogeneous information sources. MSCBO addresses a key limitation of existing approaches: causal Bayesian optimization [1] exploits structural knowledge to identify impactful interventions but ignores source cost and information gain, while multi-source Bayesian optimization [5] accounts for heterogeneous sources yet neglects causal dependencies, often resulting in redundant or inefficient actions. Our method unifies these perspectives by modeling each information source as a causal graph equipped with its own Gaussian Process and an adjusted cost-sensitive Knowledge Gradient acquisition function, enabling principled trade-offs between intervention impact and cost. To ensure scalability in high-dimensional settings, MSCBO restricts optimization to Posterior Optimal Minimal Intervention Sets (POMIS) [4], which leverage the causal structure to reduce the intervention space without sacrificing optimality. An adaptive ϵ-greedy policy further balances exploration through low-cost observations and exploitation via targeted interventions under a global budget constraint.

The algorithm requires the user to provide the information sources (the DAG's and associated observations), their costs, an intervention- and observation cost, a ground truth (for performance evaluation) and a general budget for interventions/observations. It operates in iterations and is structured as follows.

1. Optimize and Evaluate the Cost-sensitive information Gradient (CKG) for each of the information sources.
2. Collect the source producing the maximum CKG-value.
3. Determine whether to intervene or observe based on the ϵ-greedy policy.
    (a) **IF** observe, collect new observational data from the previously collected best source's network (k samples, where k is user defined) and update the hyperparameters of its causal model with this data.
    (b) **IF** intervene, collect the optimal intervention set from the previously collected best source, perform the intervention to compute a posterior value, update the hyperparameters of its causal model and update the global optimum if the existing optimum is exceeded by the computed posterior value.
4. Update running cost, terminate operation if the total budget is exceeded and repeat step 1 if the budget is not exceeded.

After completion of the algorithmic process, one is left with the optimal intervention set - value pair along with the information source

PSA Example

E. coli Example

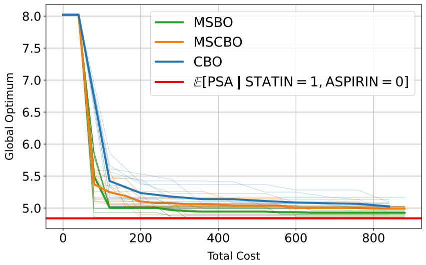

(a) The Base Case

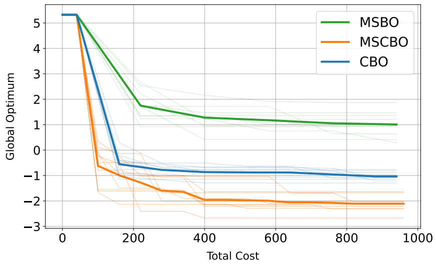

(b) The Base Case

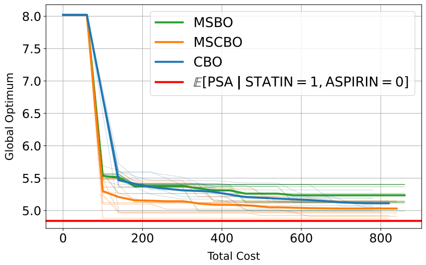

(c) Scenario 3: Added/removed nodes

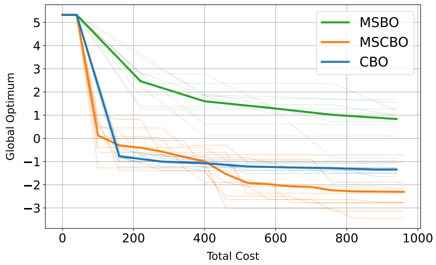

(d) Scenario 3: Added/removed nodes

that the pair originates from. Its features make MSCBO particularly well-suited for real-world applications where interventions are expensive or limited, such as clinical trials and industrial process optimization. By explicitly incorporating causal knowledge, multi-source information, and cost considerations, MSCBO provides a practical and principled framework for optimization, enabling the prioritization of interventions that are both informative and impactful while efficiently managing resources in environments where causal relationships can be leveraged for better outcomes.

2 EXPERIMENTAL RESULTS

We empirically evaluated the performance of MSCBO by using an experimental setup encompassing a range of three increasingly noisy testing scenarios applied to Bayesian networks sampled from related publications and the Bayesian Network Repository [2]. The code and data are available at https://github.com/LuukJacobs1/MSCBO.git This setup served the purpose of assessing the robustness and scalability of our implementation while also determining its comparative performance to both a non-causal multi-source and a causal single-source approach. This paper presents results on two real-world causal networks, one examining the effect of statins' (a medicine group) use on Prostate Specific Antigen (PSA) levels [3, 7] and the other describing regulatory interactions among 46 genes in E. coli [6]. For brevity, we depict only the base case and the most noisy testing scenario for each network. A complete overview of the methodology and experimental results is available at https://arxiv.org/abs/2602.14791. Results highlight MSCBO as a flexible and reliable optimization framework. In baseline and worst-case scenarios, MSCBO achieves performance comparable to its foundational counterparts MSBO and CBO, demonstrating that incorporating causal structure and source selection does not introduce performance regressions. When causal information can be effectively exploited, MSCBO consistently improves cost efficiency and convergence behavior, particularly in larger and more complex networks. These improvements are driven by the algorithm's ability to reduce the effective search space using causal reasoning while selectively intervening in the most cost-efficient information source, allowing optima to be identified at substantially lower cost.

Across noisy and perturbed environments, MSCBO and CBO exhibit increased robustness relative to a non-causal approach, suggesting that the usage of causal information provides stability when the underlying system is partially misspecified or uncertain. In high-dimensional settings, the advantages of MSCBO become more pronounced, as traditional Bayesian optimization methods struggle with rapidly growing intervention costs and inefficient exploration. By combining exploration set reduction with intelligent source selection, MSCBO scales more effectively with network size and the number of information sources.

3 CONCLUSIONS

In this paper, we proposed the concept of MSCBO to translate the advantage of a causality-based approach to the multi-source domain setting. We empirically evaluated the performance of MSCBO against both its foundational counterparts, concluding that its performance is competitive in worst-case settings and exceeds that of both alternatives when its causal strengths can be effectively leveraged. Moreover, MSCBO is expected to scale more favorably with increasing problem dimensionality and number of sources, making it a particularly well-suited optimization framework for multi-source Bayesian optimization scenarios where causal information is available.

REFERENCES

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

[2] Gal Elidan. 2001. Bayesian Network Repository. https://www.cse.huji.ac.il/ galel/Repository/.

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

[4] Sanghack Lee and Elias Bareinboim. 2018. Structural causal bandits: Where to intervene? *Advances in neural information processing systems* 31 (2018).

[5] Matthias Poloczek, Jialei Wang, and Peter Frazier. 2017. Multi-information source optimization. *Advances in neural information processing systems* 30 (2017).

[6] Juliane Schäfer and Korbinian Strimmer. 2005. A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Statistical applications in genetics and molecular biology* 4, 1 (2005).

[7] Clay Thompson. 2019. Causal graph analysis with the causalgraph procedure. In *Proceedings of SAS Global Forum*.