

# DoWhy-GCM: An Extension of DoWhy for Causal Inference in Graphical Causal Models

**Patrick Blöbaum**

BLOEBP@AMAZON.COM

**Peter Götz**

PEGO@AMAZON.COM

**Kailash Budhathoki**

KAIBUD@AMAZON.COM

**Atalanti A. Mastakouri**

ATALANTI@AMAZON.COM

**Dominik Janzing**

JANZIND@AMAZON.COM

*Amazon Web Services*

**Editor:** Joaquin Vanschoren

## Abstract

We present DoWhy-GCM, an extension of the DoWhy Python library, which leverages graphical causal models. Unlike existing causality libraries, which mainly focus on effect estimation, DoWhy-GCM addresses diverse causal queries, such as identifying the root causes of outliers and distributional changes, attributing causal influences to the data generating process of each node, or diagnosis of causal structures. With DoWhy-GCM, users typically specify cause-effect relations via a causal graph, fit causal mechanisms, and pose causal queries—all with just a few lines of code. The general documentation is available at <https://www.pywhy.org/dowhy> and the DoWhy-GCM specific code at <https://github.com/py-why/dowhy/tree/main/dowhy/gcm>.

**Keywords:** Python Library, Causal Inference, Graphical Causal Models, Structural Causal Models

## 1. Introduction

Analyzing complex systems beyond mere correlations is becoming increasingly important in modern data science (Peters et al., 2017; Pearl, 2009). In order to analyze a system and to understand the interaction between components, modeling their causal interactions is crucial. One of the widespread applications of causality is in *effect estimation* problems, where the goal is to estimate the effect of an intervention on one variable on another (Imbens and Rubin, 2015; Rubin, 1974). The Python library DoWhy (Sharma and Kiciman, 2020) initially only focused on addressing such type of problems, but causal questions often span beyond effect estimation problems (Peters et al., 2017; Janzing et al., 2024; Budhathoki et al., 2022b, 2021; Janzing et al., 2020).

DoWhy-GCM is a complementary extension of DoWhy and leverages graphical causal models (GCM) (Pearl, 2009) to answer a broader range of causal questions (see Figure 1 for an overview). To get started with DoWhy-GCM, users provide two objects: a directed acyclic graph (DAG) representing the causal relationships between variables in a system, and *tabular observational data* corresponding to the variables in the DAG, which can be continuous, discrete or categorical. They then fit a GCM, and ask causal questions, such as identifying root causes of outliers (Budhathoki et al., 2022b), quantitatively attribute distributional changes to the data generating mechanisms in place at each node (Budhathoki et al., 2021), quantify strength of causal influences (Janzing et al.,

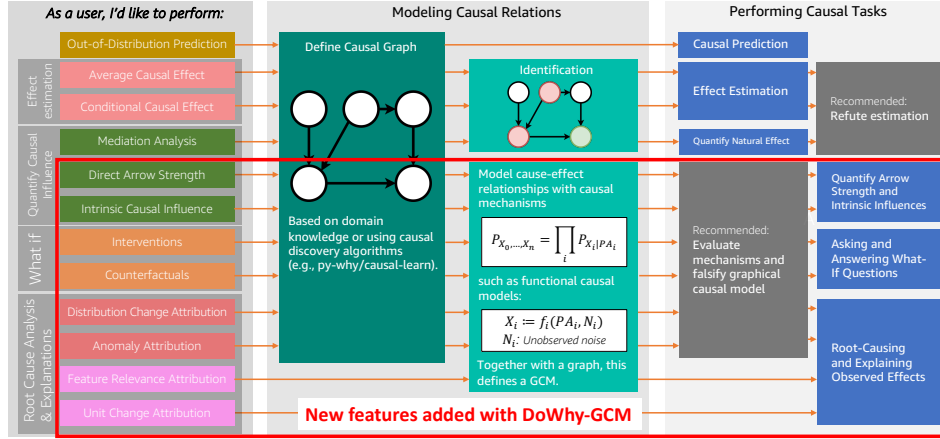


Figure 1: DoWhy-GCM complements DoWhy by allowing to address a wide range of causal questions by utilizing graphical causal models. The graphical structure is a common data type across most of the features, which ensures interoperability.

2013, 2024, 2020), or point- and population-level based estimation of counterfactuals (Pearl, 2009; Peters et al., 2017) (see Section 3 for more).

DoWhy-GCM has a modular design to allow easy-integration with third-party libraries. To this end, we define interfaces for key components of the library, such as causal graph, causal model, and causal mechanism. This allows users to contribute their own algorithms and models. For example, if a user knows the functional relationship between a variable and its parents, they can assign a custom implementation of the causal mechanism to that variable. In addition, for several causal queries, we provide wrappers on top of popular existing third-party libraries, such as `scikit-learn` (Pedregosa et al., 2011) and `SciPy` (Virtanen et al., 2020) and, thus, support wide range of parametric as well as non-parametric models. While DoWhy-GCM offers a general framework for modeling a GCM, it comes with various native implementations of certain algorithms (see Section 3).

## 2. DoWhy-GCM Three-Step Recipe

DoWhy-GCM implements a general purpose framework for modeling graphical causal models, which is based on the formal framework developed by Judea Pearl (Pearl, 2009). Cause and effect relationships are represented via directed edges and the data generation process of each node in the graph given its parents is described by a so-called *causal mechanism*. These mechanisms are assumed to be modular, i.e., we can independently change the causal mechanism of one variable without affecting the causal mechanisms of other variables in the system (Peters et al. (2017, Ch. 2)). Typically answering a causal question using GCMs involves 3 key steps, which DoWhy-GCM embraces.

**1/ Model cause-effect relationships in a graphical causal model:** The first step is to model the cause-effect relationships between variables through a causal graph using any supporting graph library that abides by our interface, e.g. `NetworkX` (Hagberg et al., 2008). In addition, we assign a causal mechanism at each node in the causal graph that allows us to model the data generation process. A causal mechanism of a node yields its conditional distribution given its parents and modeling these mechanisms explicitly at each node allows us to answer a wide range of causal questions beyond typical effect estimation tasks. As a root node does not have any parents, we obtain its distribution from either the empirical distribution or a parametric model. We refer to the causal graph and its causal mechanisms as a GCM.

Modeling the generating process of each node explicitly, however, requires additional assumptions about how the causal mechanisms look like. While the user has the flexibility to define their own custom mechanisms, DoWhy-GCM also provides an API to automatically infer the causal mechanisms from observational data. For this, DoWhy-GCM supports additive noise model (ANM) (Hoyer et al., 2009) and post non-linear models (Zhang and Hyvärinen, 2009) out of the box, which come with certain modeling assumptions but are flexible and restrictive enough to render tasks like estimation of counterfactuals solvable. To further illustrate this, let us use a simple causal graph  $X \rightarrow Y$ , where we model the marginal distribution  $P_X$  and the conditional distribution  $P_{Y|X}$ . Using, e.g., an ANM here for  $Y$ , we obtain  $X := N_X$ , and  $Y := f(X) + N_Y$ , where  $N_X$  and  $N_Y$  are assumed to be independent. In DoWhy-GCM, noises  $N_X$  and  $N_Y$  can be represented by a probability distribution, whereas the function  $f$  can be any prediction function (e.g., (non-)linear regression). If the causal mechanism permits, DoWhy-GCM can even reconstruct the noise based on a given observation, which is a requirement for estimating counterfactuals in Pearl’s framework.

**2/ Fit the causal mechanisms:** After assigning causal mechanisms to nodes in the causal graph, the next step is to learn the parameters of those mechanisms from data in case of parametric models. Note that DoWhy-GCM also allows users to provide their own "ground truth" causal mechanism. After this step, the GCM is ready for causal reasoning tasks.

**3/ Ask a causal question:** With a fitted GCM, users can then answer a wide range of causal questions, where the same fitted GCM can be reused for different queries.

All these step above can be done with a few lines of code:

---

```
# Define structural causal model and automatically assign models
causal_model = gcm.StructuralCausalModel(nx.Digraph([("X", "Y"), ("Y", "Z")]))
gcm.auto.assign_causal_mechanisms(causal_model, samples_df)

# Fit generative models
gcm.fit(causal_model, samples_df)

# Optional: Evaluate causal model
gcm.evaluate_causal_model(causal_model, samples_df)

# Causal model is now ready for different causal queries
gcm.<causal query>(causal_model, ...)
```

---

### 3. DoWhy-GCM Features: `gcm.<causal query>`

This section covers some exemplary causal queries users can ask with a fitted GCM object:

**Causal Model Evaluation:** These methods check statistically falsifiable assumptions in a causal model, e.g., reject a graph or evaluate assigned causal mechanisms given data.

- » `gcm.falsify_graph`: Falsify a given DAG using observational data (Eulig et al., 2023).
- » `gcm.evaluate_causal_model`: Evaluate causal mechanisms and modeling assumptions using observational data.

**Quantify Causal Influence:** These quantify causal influences, such as the strength of an edge or the intrinsic influence of nodes on a target.

- » `gcm.arrow_strength`: Quantifies the strength of an edge (Janzing et al., 2013).

- » `gcm.intrinsic_causal_influence`: Quantifies contribution of a source node to the uncertainty in a target node (Janzing et al., 2024).

**What-If:** These methods estimate the effect of interventions which can be more general beyond standard atomic interventions, and compute counterfactuals (Peters et al., 2017).

- » `gcm.interventional_samples`: Performs intervention on nodes in the causal graph.
- » `gcm.counterfactual_samples`: Estimates counterfactual data for observed data if we were to perform specified interventions.

**Attribution:** These methods attribute observed effects (e.g., outliers, distributional changes) to their root causes (i.e., nodes upstream in the causal graph).

- » `gcm.attribute_anomalies`: Quantifies contributions to anomalies (Budhathoki et al., 2022b).
- » `gcm.distribution_change`: Quantifies the contribution of each node to the change in the distribution of the target (Budhathoki et al., 2021).
- » `gcm.parent_relevance`: Quantifies the relevance of a causal parent with respect to the causal mechanism by explicitly incorporating the noise (Lundberg and Lee, 2017; Janzing et al., 2020).
- » `gcm.unit_change`: Quantifies the contributions of input features and prediction mechanisms to the change in the value of the target for a statistical unit (Budhathoki et al., 2022a).

## 4. Code design

DoWhy-GCM follows functional programming principles to focus on causal questions and avoid issues like code delegation, duplication, bookkeeping, statefulness, and scalability problems with ever-growing class APIs. It leverages existing, commonly used libraries like `NetworkX` (Hagberg et al., 2008), `NumPy` (Harris et al., 2020), and `Pandas` (McKinney, 2010) to avoid reinventing the wheel. The API is designed to be functional - operating on a GCM object and returning results to prevent statefulness. It provides sensible defaults for methods, offers various convenience functions to simplify usage while allowing full customization, and enables inspection by exposing the various GCM components and models explicitly.

## 5. Discussion

We introduced DoWhy-GCM, an extension of DoWhy, that complements DoWhy’s existing effect estimation features by adding features for graphical causal models operating on the same causal graphs. Graphical causal models allow us to look at a system as a whole, not just the effect of one variable on another, allowing us to address causal questions beyond effect estimation. That is, we model cause-effect relationships between all variables explicitly with causal mechanisms representing probabilistic models and fit them based on observational tabular data. The API is designed with a focus on modularity, allowing easy integration of custom models, algorithms, or other third-party libraries. The scalability of the provided algorithms heavily depend on the inference complexity of the used models, the number of variables in the causal graph, the sample size, or the structure of the causal graph. In contrast to other Python causality libraries such as `EconML` (Battocchi et al., 2019), `CausalML` (Chen et al., 2020), `cdt` (Kalainathan and Goudet, 2019), `DiCE` (Mothilal et al., 2020), `Tetrad` (Ramsey et al.), `CausalNex` (Beaumont et al., 2021) or `WhyNot` (Miller et al., 2020) that primarily focus on effect estimation, DoWhy-GCM offers features and novel algorithms that explicitly leverage graphical causal models.

## References

- K. Battocchi, E. Dillon, M. Hei, G. Lewis, P. Oka, M. Oprescu, and V. Syrgkanis. EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/microsoft/EconML>, 2019. Version 0.x.
- P. Beaumont, B. Horsburgh, P. Pilgerstorfer, A. Droth, R. Oentaryo, S. Ler, H. Nguyen, G. A. Ferreira, Z. Patel, and W. Leong. CausalNex, 10 2021. URL <https://github.com/quantumblacklabs/causalnex>.
- K. Budhathoki, D. Janzing, P. Bloebaum, and H. Ng. Why did the distribution change? In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1666–1674. PMLR, 13–15 Apr 2021. URL <https://proceedings.mlr.press/v130/budhathoki21a.html>.
- K. Budhathoki, G. Michailidis, and D. Janzing. Explaining the root causes of unit-level changes, 2022a. URL <https://arxiv.org/abs/2206.12986>.
- K. Budhathoki, L. Minorics, P. Blöbaum, and D. Janzing. Causal structure-based root cause analysis of outliers. In *International Conference on Machine Learning*, pages 2357–2369. PMLR, 2022b.
- H. Chen, T. Harinen, J. Lee, M. Yung, and Z. Zhao. CausalML: Python package for causal machine learning, 2020. URL <https://arxiv.org/abs/2002.11631>.
- E. Eulig, A. A. Mastakouri, P. Blöbaum, M. Hardt, and D. Janzing. Toward falsifying causal graphs using a permutation-based test, 2023. URL <https://arxiv.org/abs/2305.09565>.
- A. A. Hagberg, D. A. Schult, and P. J. Swart. Exploring network structure, dynamics, and function using networkx. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.
- C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. Fernández del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- P. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Proceedings of the conference Neural Information Processing Systems (NIPS) 2008*, Vancouver, Canada, 2009. MIT Press.
- G. W. Imbens and D. B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, 2015.
- D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. Quantifying causal influences. *Annals of Statistics*, 41(5):2324–2358, 2013.

- D. Janzing, L. Minorics, and P. Blöbaum. Feature relevance quantification in explainable AI: A causal problem. In S. Chiappa and R. Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2907–2916, Online, 26–28 Aug 2020. PMLR.
- D. Janzing, P. Blöbaum, A. A Mastakouri, P. M Faller, L. Minorics, and K. Budhathoki. Quantifying intrinsic causal contributions via structure preserving interventions. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li, editors, *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 2188–2196. PMLR, 02–04 May 2024. URL <https://proceedings.mlr.press/v238/janzing24a.html>.
- D. Kainathan and O. Goudet. Causal discovery toolbox: Uncover causal relationships in Python, 2019. URL <https://arxiv.org/abs/1903.02278>.
- S. Lundberg and S. Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017.
- W. McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56 – 61, 2010. doi: 10.25080/Majora-92bf1922-00a.
- J. Miller, C. Hsu, J. Troutman, J. Perdomo, T. Zrnic, L. Liu, Y. Sun, L. Schmidt, and M. Hardt. WhyNot, 2020. URL <https://doi.org/10.5281/zenodo.3875775>.
- R. K Mothilal, A. Sharma, and C. Tan. Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 607–617, 2020.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, NY, 2nd edition, 2009.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference – Foundations and Learning Algorithms*. MIT Press, 2017.
- J. D Ramsey, K. Zhang, M Glymour, R. S. Romero, B. Huang, I. Ebert-Uphoff, S. Samarasinghe, E. A Barnes, and C. Glymour. Tetrad—a toolbox for causal discovery.
- D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701, 1974.
- A. Sharma and E. Kiciman. DoWhy: An end-to-end library for causal inference, 2020. URL <https://arxiv.org/abs/2011.04216>.

- P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C J Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, Montreal, Canada, 2009.