

CAUSALPROFILER: GENERATING SYNTHETIC BENCHMARKS FOR RIGOROUS AND TRANSPARENT EVALUATION OF CAUSAL MACHINE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Causal machine learning (Causal ML) aims to answer “what if” questions using machine learning algorithms, making it a promising tool for high-stakes decision-making. Yet, empirical evaluation practices in Causal ML remain limited. Existing benchmarks often rely on a handful of hand-crafted or semi-synthetic datasets, leading to brittle, non-generalizable conclusions. To bridge this gap, we introduce CausalProfiler, a synthetic benchmark generator for Causal ML methods. Based on a set of explicit design choices about the class of causal models, queries, and data considered, the CausalProfiler randomly samples sets of data, assumptions, and ground truths constituting the synthetic causal benchmarks. In this way, Causal ML methods can be rigorously and transparently evaluated under a variety of conditions. This work offers the first random generator of synthetic causal benchmarks with coverage guarantees and transparent assumptions operating on the three levels of causal reasoning: observation, intervention, and counterfactual. We demonstrate its utility by evaluating several state-of-the-art methods under diverse conditions and assumptions, both in and out of the identification regime, illustrating the types of analyses and insights the CausalProfiler enables.

1 INTRODUCTION

Causal machine learning (Causal ML) seeks to estimate the effects of interventions and counterfactuals using machine learning techniques (Kaddour et al., 2022), enabling principled decision making—for example in medicine and policy. Despite the theoretical maturity and growing relevance of Causal ML, current research practices lack rigorous evaluations of how proposed methods would perform under realistic and diverse conditions, limiting their practical utility (Curth et al., 2024; Feuerriegel et al., 2024; Poinot et al., 2025; Berrevoets et al., 2024).

In Causal ML, evaluation is particularly challenging due to the unobservability of counterfactual outcomes (Holland, 1986). Researchers can rely only on scarce real-world data sources such as randomized controlled trials, considered the gold standard, which are expensive, are ethically constrained, and often encompass a low amount of data (Greenland & Brumback, 2002; Tennant et al., 2021). As a result, existing benchmarks often rely on a few semi-synthetic datasets (e.g., Syntren (den Bulcke et al., 2006), ACIC2016 (Dorie et al., 2019)) or model-driven synthetic datasets generated from fitted causal mechanisms (Neal et al., 2020; Parikh et al., 2022; Athey et al., 2024; de Vassimon Manela et al., 2024). However, these datasets encode assumptions that are rarely made explicit and whose validity is difficult to generalize beyond the original study context (Poinot et al., 2025). In parallel, many researchers define handcrafted synthetic datasets, useful for theory but fragile for empirical evaluation: a few manually chosen models can overstate performance by aligning with method-specific assumptions (Gentzel et al., 2019). Moreover, lessons from predictive machine learning show that narrow, static benchmarks can give a false sense of reliability (Geirhos et al., 2020; Herrmann et al., 2024; Freiesleben & Grote, 2023; Longjohn et al., 2024), underscoring the need for structured diversity: systematic variation of tasks under explicit, controllable assumptions.

In this work, we take a concrete step toward addressing these fundamental concerns about the field. Specifically, we introduce a synthetic benchmark generator, the CausalProfiler, that enables robust empirical evaluations grounded in transparently defined synthetic causal datasets. Central to our

054 approach is the notion of a *Space of Interest (SoI)* (Definition 5.1), defining the domain from which
055 causal datasets are sampled. Given an *SoI*, CausalProfiler samples Structural Causal Models (SCMs),
056 data, and queries, and estimates the ground truth value of the queries to enable the evaluation of
057 Causal ML methods. The assumptions are explicit; and dataset characteristics can be systematically
058 varied through the *SoI*. Hence, CausalProfiler enables transparent, controlled, repeatable, and diverse
059 sampling of synthetic causal datasets.

060 CausalProfiler shifts the focus of empirical evaluation from performance on individual datasets to
061 trends and patterns across a well-characterized *SoI*, reframing the evaluation question from “what
062 dataset to use” to specifying a *SoI* that defines the scope of evaluation. This enables researchers to
063 evaluate performance across a well-defined set of conditions—on graph density, or causal mechanisms
064 complexity, for instance—and to understand under which conditions a method succeeds or fails,
065 helping practitioners identify methods that remain reliable when their causal assumptions are violated.
066 Compared to conventional evaluations in the current literature, an evaluation with CausalProfiler
067 yields more robust and reliable performance estimates; it uncovers failure modes, generalization
068 limits, and assumption sensitivities that remain hidden in conventional evaluations.

069 Although synthetic evaluation cannot replace real data, it provides the only reliable access to ground-
070 truth causal queries, since counterfactuals are unobservable and many assumptions are unfalsifi-
071 able (Holland, 1986; Poinsot et al., 2025). CausalProfiler brings about a much needed complement to
072 real-world studies by enabling transparent, diverse, and controlled synthetic experiments.

073 We make two primary contributions. First, we present CausalProfiler¹ (Section 5), the first open-
074 source benchmark generator that enables principled sampling of synthetic causal datasets, with
075 coverage guarantees, that promote transparency and reproducibility for Causal ML evaluation over
076 the three levels of causal reasoning. Secondly, we demonstrate through experiments (Section 6)
077 how evaluation with CausalProfiler yields richer and more robust insights than the current standard
078 practice.

080 2 RELATED WORK

082 **Evaluating Causal ML methods.** Causal ML currently lacks a rigorous, systematic paradigm for
083 empirical evaluation, whether synthetic or semi-synthetic. Semi-synthetic datasets, such as synthetic
084 outcome datasets (Dorie et al., 2019; Shimoni et al., 2018; Hill, 2011) and model-based semi-synthetic
085 datasets (Neal et al., 2020; Parikh et al., 2022; Athey et al., 2024; de Vassimon Manela et al., 2024),
086 combine real covariates and simulated outcomes under assumed structural models. On the other hand,
087 fully synthetic datasets are generated entirely from researcher-defined SCMs, allowing for greater
088 control and access to ground truth. Yet both synthetic and semi-synthetic approaches share critical
089 limitations. First, synthetic evaluations often lack realism, relying on overly simplistic mechanisms
090 such as additive noise or linear functions, and frequently omitting robustness analyses (Gentzel et al.,
091 2019; Curth et al., 2024; Poinsot et al., 2024; 2025). Such evaluations rarely reflect the complexity of
092 real-world causal processes and are insufficient to test the limits of modern causal inference methods.
093 Secondly, synthetic and semi-synthetic datasets are shaped by researcher-defined design choices,
094 including the causal graph structure, the form of the outcome function, and the noise distribution.
095 These decisions, often made implicitly, can unintentionally introduce hidden biases that favor certain
096 methods (Curth et al., 2021; Cheng et al., 2022; Feuerriegel et al., 2024). Such assumptions are rarely
097 documented or systematically varied, hindering reproducibility and fair method comparison (Poinsot
098 et al., 2024; 2025). Additionally, these benchmarks are typically small in scale and narrow in scope,
099 often covering only a limited range of causal settings. As a result, empirical evaluations raise concerns
100 about overfitting and generalization (Gentzel et al., 2019; Berrevoets et al., 2024). For instance, it
101 has been shown that even small changes to the data-generating process can lead to dramatic shifts in
102 performance rankings (Curth et al., 2021). Moreover, methods are often evaluated only under the very
103 conditions that guarantee their identifiability, offering little insight into robustness under assumption
104 violations, as is common in real-world settings (Petersen, 2024; Hutchinson et al., 2022). In short,
105 without broader and more transparent evaluation across diverse causal settings, the field risks drawing
106 conclusions that do not generalize. For Causal ML to have wide impact in practice, there is a need to
107 move beyond fixed benchmarks toward frameworks that support transparent, controlled, and diverse
experimentation across well-defined spaces of causal assumptions.

¹The code is provided in the supplementary material and will be publicly available after the review process.

Recent benchmarking efforts. Recent works have sought to address some of these problems, introducing tools to generate synthetic SCMs for causal discovery (Kalainathan et al., 2020; Gupta et al., 2023; Rudolph et al., 2023) or support query estimation from hand-specified models (Sharma & Kiciman, 2020; Textor et al., 2017; Abril-Pla et al., 2023). However, none of these frameworks support all components required for robust evaluation of causal machine learning methods. First, the causal discovery benchmarks do not compute ground truth for intervention or counterfactual queries. Further, query estimation frameworks often require manual SCM specification and do not support random sampling, diversity control, or analysis of the distribution of tasks. Even when SCMs are sampled (Rudolph et al., 2023; Xia et al., 2023), key properties (e.g., positivity) are neither reported nor constrained. Moreover, the absence of randomness in the graph structures limits generalization. In contrast, CausalProfiler integrates SCM sampling, query ground-truth computation, and coverage guarantees into a unified framework. To the best of our knowledge, this is the first benchmark generator that enables systematic exploration of how Causal ML methods behave across spaces of SCMs and queries defined by user-specified constraints.

3 BACKGROUND & NOTATION

We use capital letters for random variables (e.g., X), lowercase for realizations (e.g., x), and boldface for vectors (e.g., \mathbf{x}). For a more complete background, please refer to Appendix B and Pearl (2009).

The **Pearl Causal Hierarchy (PCH)** (Pearl & Mackenzie, 2018) classifies causal reasoning into three levels: \mathcal{L}_1 (associational), \mathcal{L}_2 (interventional), and \mathcal{L}_3 (counterfactual). Associative questions use only observed data, whereas interventional and counterfactual questions require assumptions about the data-generating process. Importantly, lower levels are insufficient to answer higher-layer questions in almost all causal models (Bareinboim et al., 2022).

The class of **Structural Causal Models (SCMs)** (Pearl, 2009) provide a representation allowing reasoning on the three levels of the PCH. An SCM is a tuple $\mathcal{M} := \{\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\}$, where \mathbf{V} is a set of endogenous variables, \mathbf{U} is a set of exogenous variables, \mathcal{F} is a set of structural equations $V_i = f_i(\mathbf{PA}(V_i), \mathbf{U}_{V_i})$, also called *causal mechanisms*, and $P(\mathbf{U})$ defines a distribution over the exogenous variables \mathbf{U} . SCMs induce a distribution $P_{\mathcal{M}}(\mathbf{V})$ over the endogenous variables \mathbf{V} , called the *entailed distribution*. We consider two types of endogenous variables: the observed variables, denoted \mathbf{V}_O , and the unobserved variables, denoted \mathbf{V}_H , where $\mathbf{V} = \mathbf{V}_O \cup \mathbf{V}_H$ and $\mathbf{V}_O \cap \mathbf{V}_H = \emptyset$. We represent causal relationships using the **causal graph** \mathcal{G} of a SCM. This is an acyclic directed mixed graph over the endogenous variables. Directed edges $X \rightarrow Y$ encode causal dependencies via causal mechanisms where $X \in \mathbf{PA}(Y)$ is called a parent of Y , while bidirected edges $X \leftrightarrow Y$ indicate latent confounding due to shared exogenous causes. With SCMs one can represent intervention and counterfactual questions. On the one hand, an **intervention** replaces one or more structural equations to model external manipulations. A common example is a *hard intervention*, $\mathbf{do}(T = t)$, which fixes a variable’s value, disconnecting it from its causes. This defines a new SCM and alters the induced distribution. On the other hand, **counterfactual** questions reason about what would have happened under a different intervention, given an observed outcome called a *factual* realization. They are evaluated by conditioning on observed variables (abduction), modifying the SCM with the intervention (action), and predicting outcomes under the new distribution (prediction)—a process known as the *three-step procedure* (Pearl, 2009).

More generally, a **causal query** refers to a probabilistic statement about the effect of hypothetical manipulations of the data-generating process. This includes *intervention queries*, such as Average Treatment Effect (ATE), and *counterfactual queries*, such as Counterfactual Total Effect (Ctf-TE). A query is *identifiable* if its value can be uniquely determined from data, given a set of assumptions (e.g., a causal sufficiency) (Pearl, 2009). In other words, **identifiability** refers to whether causal queries can be empirically estimated, and under what assumptions.

4 PROBLEM FORMULATION

Causal inference aims to answer causal queries using data drawn from an unknown SCM. Let $\mathcal{M}^* = (\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}))$ denote the unknown ground truth SCM. A causal query Q (e.g., ATE) defined over \mathcal{M}^* has ground truth value $Q^* = Q(\mathcal{M}^*)$. As \mathcal{M}^* is unknown, causal estimators rely on causal assumptions \mathbf{H} (e.g., causal sufficiency) and available data D drawn from \mathcal{M}^* to produce

an estimate \hat{Q} of the target quantity Q^* . Definition 4.1 below formalizes the elements of a causal dataset.

Definition 4.1 (Causal Dataset). A **causal dataset** is a tuple $\mathcal{D} = \{Q, Q^*, D, \mathcal{G}^*, \mathbf{H}^*\}$ constructed from a known SCM $\mathcal{M}^* = (\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}))$ where:

- Q is a causal query defined over \mathbf{V} ;
- $Q^* = Q(\mathcal{M}^*)$ is the exact value of the query Q ;
- $D = \{D_k \sim P_{\mathcal{M}^*}(\mathbf{V} \mid \text{do}(\mathbf{V}_k) = \mathbf{v}_k)\}_{k=1}^I$ is a collection of samples under I interventional (or observational) settings;
- \mathcal{G}^* is the causal graph associated with \mathcal{M}^* ;
- \mathbf{H}^* is the set of assumptions satisfied by \mathcal{M}^* .

In this work, we wish to develop a generator of causal datasets following the Definition 4.1 such that given an error metric $E(\hat{Q}, Q^*)$, Causal ML methods can be evaluated under the identification-consistent regime, where the assumed causal graph and assumptions coincide with the ground truth $(\mathcal{G}^*, \mathbf{H}^*)$, and under controlled assumption violations. This setup enables systematic comparison of Causal ML methods not only under ideal conditions, where identification is guaranteed, but also in more realistic settings that test robustness to misspecification.

Remark on causal discovery. Causal datasets, as defined above, can also be used for evaluating causal discovery algorithms. Each dataset already includes the ground-truth causal graph \mathcal{G}^* , allowing direct assessment of discovery methods. Thus, the query Q can be left empty.

5 SAMPLING CAUSAL DATASETS WITH THE CAUSALPROFILER

To generate causal datasets, CausalProfiler relies on a parametric specification of the sampling domain, called the *Space of Interest (SoI)*. Given an *SoI*, CausalProfiler samples an SCM (Section 5.2) and generates a corresponding causal dataset (Section 5.3). Appendices C to F contain the pseudocode for the sampling algorithms. Appendix I presents a visual overview of the sampling strategy.

5.1 DEFINING A SPACE OF INTEREST

The central abstraction of our framework is the *Space of Interest (SoI)* (Definition 5.1), which provides a standardized way to specify synthetic causal datasets (Definition 4.1). Table 3 in Appendix C lists all configurable *SoI* parameters.²

Definition 5.1 (Space of Interest). A **Space of Interest (SoI)** is a tuple $\mathcal{S} = \{\mathbb{M}, \mathbb{Q}, \mathbb{D}\}$, where \mathbb{M} is a class of SCMs, \mathbb{Q} a class of causal queries, and \mathbb{D} a class of data.

5.2 SAMPLING STRUCTURAL CAUSAL MODELS

Causal Graphs. CausalProfiler first samples a directed acyclic graph over a set of endogenous variables, defining the SCM’s causal structure. If specified in the *SoI*, CausalProfiler samples a subset of endogenous variables, \mathbf{V}_H , to be treated as unobserved and excluded from the observed dataset. To expose only the visible causal structure to the user, we apply Verma’s latent projection algorithm (Verma, 1993) to the full causal graph, which produces an acyclic directed mixed graph.

Mechanisms. Given the causal graph, CausalProfiler assigns each endogenous variable a mechanism based on its parents and an exogenous noise distribution set by the *SoI*. It supports two types of mechanisms. First, **discrete mechanisms**, also called regional discrete mechanisms (see Appendix E.1 for a formal definition) which support binary and categorical treatments, are defined tabularly by associating each element of a partition of the exogenous noise with distinct parents-to-child

²While the current implementation of CausalProfiler supports only \mathcal{L}_1 training data and Average Treatment Effect (ATE), Conditional Average Treatment Effect (CATE), and Counterfactual Total Effect (Ctf-TE) queries, the *SoI* abstraction can, in principle, be defined over any class of queries, datasets, and SCMs.

mappings. This enables controllable stochasticity and complexity, including highly non-linear and non-invertible behavior. The *SoI* also specifies how such mechanisms are sampled (e.g., with rejection-based sampling, see Appendix E.2). Secondly, **continuous mechanisms** are defined using parametric function families—such as neural networks or linear functions—with randomly initialized parameters.

5.3 SAMPLING CAUSAL DATASETS

Data D . Given an SCM \mathcal{M}^* sampled from the *SoI*, we generate an observational dataset D by sampling i.i.d. data points from the entailed distribution of \mathcal{M}^* over observed variables. This involves forward-sampling from the structural equations in topological order, using the noise distributions specified for each variable and marginalizing out any latent variables.

Query Q . We first sample endogenous observable variables to serve as treatment, outcome, covariates, and factuials, depending on the query class of the *SoI*. By default, realizations are drawn from a large, separately sampled observational dataset, rather than from the theoretical variable domains. This ensures that queries are well-defined and correspond to realizable variable configurations under the SCM. To support different research goals, *SoIs* can be configured to relax this behavior (e.g., to include NaN queries) to stress-test robustness. For causal discovery, query sampling can be disabled to generate datasets more efficiently given that they already include their ground-truth graph \mathcal{G}^* .

Query ground truth Q^* . Each query is estimated by drawing samples from the (manipulated) ground truth SCM: interventional queries via do-operations (action and prediction), and counterfactual queries via the three-step procedure (Pearl, 2009).

Ground truth causal graph \mathcal{G}^* . As presented in Section 5.2, \mathcal{G}^* is built as the latent projection of the ground-truth SCM’s causal graph over the observed variables.

Ground truth causal assumptions \mathbf{H}^* . To characterize the properties of the ground-truth SCM from the user’s perspective, we provide an analysis module that computes summary metrics related to common causal assumptions (e.g., measuring linearity via Pearson correlation). A full list of available metrics is provided in Appendix G.

Coverage guarantee. Proposition 5.1 (proof in Appendix J) shows that, with sufficiently expressive discrete mechanisms, CausalProfiler’s sampling strategy can theoretically generate any causal dataset within a given *SoI*, guaranteeing \mathcal{L}_3 -expressivity. In addition, Appendix H provides an analysis exploring the empirical distribution of the sampled datasets.

Proposition 5.1 (Coverage). For a Space of Interest $\mathcal{S} = \{\mathbb{M}, \mathbb{Q}, \mathbb{D}\}$, whose class of Structural Causal Models is a class of Regional Discrete SCMs¹ with the maximum number of noise regions, denoted $\mathbb{M}_{\text{RD-SCM}, r=R_{\text{max}}}$, any causal dataset $\mathcal{D} = \{Q, Q^*, D, \mathcal{G}^*, \mathbf{H}^*\}$ has a strictly positive probability to be generated.

$$\forall \mathcal{S} = \{\mathbb{M}, \mathbb{Q}, \mathbb{D}\} \text{ s.t. } \mathbb{M} \subseteq \mathbb{M}_{\text{RD-SCM}, r=R_{\text{max}}}, P(\mathcal{D}|\mathcal{S}) > 0$$

¹A formal definition can be found in Appendix E.1.

Benchmark Design. Taken together, these design choices reflect four key properties that are considered essential for rigorous synthetic evaluation in Causal ML (Poinsot et al., 2025): **transparency**, by making all assumptions explicit via the parametrization of the *SoI*, which serves as a declarative specification of the evaluation domain; **repeatability**, through randomized but seed-controlled sampling procedures, ensuring that SCMs and queries can be exactly reproduced across runs; **bias awareness**, supported by the coverage guarantee and the empirical distribution analysis module; and **control over experiments**, by exposing a wide range of configurable parameters in the *SoI* that allow users to tailor the causal dataset generation to their assumptions and research goals.

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

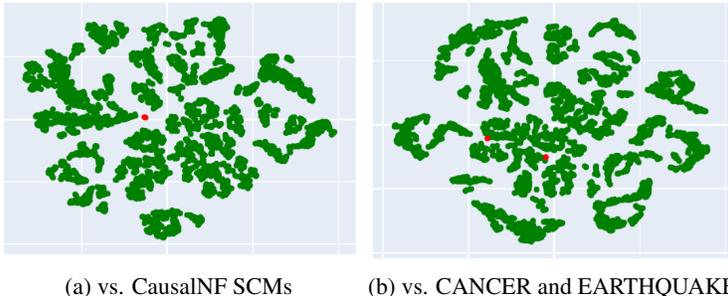


Figure 1: Two-dimensional t-SNE plots of CausalProfiler’s SCMs (green) and established benchmarks (red), characterized by metrics from the analysis module.

6 EXPERIMENTS

6.1 VERIFICATION OF BENCHMARK CORRECTNESS

To validate the soundness of our benchmark generator, we perform consistency checks across the three levels of the PCH. Using the SCM sampler and query estimator of the CausalProfiler, we test whether sampled SCMs satisfy the Markov condition, do-calculus rules, and the structural counterfactual axioms (Pearl, 2009). We use discrete SCMs to allow exhaustive enumeration of conditioning sets for statistical tests. To ensure robustness, we iterate over a *SoI* parameter grid spanning the number of variables, edge density, cardinalities, and noise regions. See Appendix K for full details and results.

L1: Markov Property Verification. We test whether d-separations in the causal graph imply conditional independencies in the entailed observational distribution of the sampled SCMs. For each SCM, we enumerate d-separated triplets (A, B, C) and test $A \perp B \mid C$ with Pearson’s χ^2 test (Pearson, 1900), filtering low-sample strata (Koehler & Larntz, 1980) and correcting for multiple tests (Benjamini & Hochberg, 1995). The Markov property holds in about 95% of the tested cases, with most violations due to finite-sample variability.

L2: Do-Calculus Verification. We test whether the three rules of do-calculus hold empirically. For each rule, we identify variable tuples satisfying its graphical preconditions. We then use the query estimator to generate two interventional datasets corresponding to the rule’s left- and right-hand sides. We compare the resulting distributions with Pearson’s χ^2 test, filtering low-sample strata (Koehler & Larntz, 1980) and correcting for multiple tests (Benjamini & Hochberg, 1995). About 5.5% of tests fail, mostly due to finite-sample noise.

L3: Structural Counterfactual Axiom Verification. We test whether the axioms of *composition*, *effectiveness*, and *reversibility* hold for sampled SCMs. Since the axioms involve deterministic functional relationships, we count only exact matches of the query estimator. All axioms hold exactly across our samples, confirming the estimator’s consistency with structural counterfactual semantics.

6.2 COMPARISON TO EXISTING BENCHMARKS

Comparison. To illustrate CausalProfiler’s contribution to SCM diversity for evaluating Causal ML methods, we compare its SCMs (sampled over a *SoI* grid spanning number of variables, edge density, cardinalities, noise regions, and dataset size) with two existing benchmarks: the synthetic SCMs from the Causal Normalizing Flows (CausalNF) work (Javaloy et al., 2023) and the CANCER and EARTHQUAKE models from bnlearn (Scutari, 2019). For interpretable visualization, we apply two-dimensional t-SNE (Maaten & Hinton, 2008) to the computable metrics of the analysis module (Appendix G), with a perplexity set to 30, see Figure 1.

Findings. CausalProfiler’s SCMs are more diverse than those of CausalNF and bnlearn. The eleven CausalNF SCMs are so homogeneous relative to CausalProfiler’s diversity that dimensionality reduction projects them onto a confined space, appearing as a single point. Moreover, some generated SCMs are similar to those of CausalNF and bnlearn, showing that CausalProfiler can reproduce existing benchmarks. Further details and results are presented in Appendices H.3 and H.4.

6.3 METHOD EVALUATION USING CAUSALPROFILER

We demonstrate the utility of our framework by evaluating several recent causal inference methods across diverse *SoIs*. Our goal is not to exhaustively benchmark each method but to showcase the kinds of structured empirical investigations CausalProfiler enables — especially on exploring robustness and violations of causal assumptions. Accordingly, we keep most *SoI* parameters fixed and vary only one (or a small subset) at a time, so observed differences can be attributed to the parameter under study rather than confounded by simultaneous changes.

For each *SoI*, we evaluate every method using five random seeds, sampling 100 SCMs per seed. Each SCM yields one training set and five queries with ground-truth values, and results are aggregated across SCMs and seeds (see Algorithm 12 in Appendix L). Experiments were run on a single Intel Core i9-14900K machine (24 cores, 32 threads, 96GB RAM), fully parallelized on CPU. Although some methods (e.g., DCM) could benefit from GPU acceleration, none was used here.

Performance is assessed by mean squared error between predicted and true query values, with mean error, standard deviation, runtime, and failure rate (due to numerical issues or exceptions) for each method and *SoI*. We compare Causal Normalizing Flows (CausalNF) (Javaloy et al., 2023), Neural Causal Models (NCM) (Xia et al., 2023), Variational Causal Graph Autoencoder (VACA) (Sánchez-Martin et al., 2022), and Diffusion-based Causal Models (DCM) (Chao et al., 2023).

Additional experiments, extended results, and *SoI* configurations are provided in Appendix L.

6.4 EXPERIMENT 1: GENERAL EVALUATION ACROSS DIVERSE SCMS

To showcase CausalProfiler’s flexibility, we evaluate ATE estimates of VACA, CausalNF, DCM, and NCM on continuous-variable SCMs across four *SoIs*: **Linear-Medium**, linear SCMs (15-20 nodes, 1000 samples); **NN-Medium**, neural SCMs with a 2-layer ReLU network (8 hidden units, 15-20 nodes, 1000 samples); **NN-Large**, larger neural SCMs (20-25 nodes, 1000 samples); and **NN-Large-LowData**, identical to NN-Large but with 50 samples. See Table 1 for results.

Table 1: Performance summary of CausalNF, DCM, NCM, and VACA on the general experiments.

Space	Method	Mean Error	Std Error	Max Error	Runtime (s)	Fail Rate (%)
Linear-Medium	CausalNF	0.4625	0.8985	9.6079	13790.4	0.00
	DCM	0.1530	1.5289	33.9766	16541.2	0.00
	NCM	0.4618	0.9001	9.6134	7384.7	0.00
	VACA	0.4209	0.6195	2.3807	2734.5	53.40
NN-medium	CausalNF	0.0160	0.0107	0.1209	10732.7	0.00
	DCM	0.0276	0.0114	0.0746	15894.4	0.00
	NCM	0.0111	0.0121	0.1484	7322.8	0.00
	VACA	0.0090	0.0077	0.0479	5759.6	5.00
NN-Large	CausalNF	0.0159	0.0105	0.1535	15114.8	0.00
	DCM	0.0267	0.0100	0.0739	19166.2	0.00
	NCM	0.0101	0.0103	0.1161	9450.6	0.00
	VACA	0.0090	0.0094	0.0535	5690.8	11.60
NN-Large-LowData	CausalNF	0.0359	0.0146	0.1712	22138.2	0.00
	DCM	0.0777	0.0445	0.3701	2412.1	0.00
	NCM	0.0097	0.0107	0.1263	404.7	0.00
	VACA	0.0103	0.0134	0.1043	5217.4	0.00

Findings (Linear-Medium vs. NN-Medium). In the **Linear-Medium** setting, DCM achieves the lowest average error (0.1530), indicating excellent performance, but its error standard deviation is notably high (1.5289), driven by a few extreme outliers (max error 33.98). This suggests DCM is effective for most queries but can produce large errors in rare cases—potentially problematic in safety-critical applications matching this *SoI*. VACA performs competitively with lower max error and faster runtime, but suffers a high failure rate (53.4%) due to NaNs. In the **NN-Medium** setting, where the causal mechanisms are small neural networks, DCM’s advantage disappears. VACA emerges as the best performer, with the lowest error mean (0.0090) and standard deviation (0.0077),

while reducing its failure rate to 5%. Interestingly, DCM becomes the weakest performer in this setting, showing that method rankings are highly sensitive to the underlying functional form of the mechanisms. This underscores the need for practitioners to evaluate methods within the *SoI* most relevant to their application. Lastly, NN SCMs surprisingly yield lower errors than linear ones. A plausible explanation is an inductive-bias match with the evaluated neural methods and the tendency of small randomly initialized NNs to produce relatively smooth, low-frequency functions that are easier to estimate from finite data (Rahaman et al., 2019).

Findings (NN-Large vs. NN-Large-LowData). In this comparison, we investigate the effect of reducing data availability. Comparing **NN-Large** (1000 samples) to **NN-Large-LowData** (50 samples), DCM is strongly affected: its error nearly triples (from 0.0267 to 0.0777) and its IQR expands noticeably. CausalNF also shows greater sensitivity to low-data regimes. In contrast, both VACA and NCM maintain stable performance, with nearly unchanged mean and standard deviation. Notably, VACA achieves a 0% failure rate, with unexpectedly strong robustness under limited data.

Insights. While not intended as a comprehensive benchmark, these experiments illustrate the types of insights enabled by our framework. Across the selected *SoIs*, DCM performs well on average but can produce large outlier errors or become less stable in low-data settings. Conversely, VACA shows promising generalization even with limited data, though it occasionally fails on certain SCMs. These findings are specific to the explored *SoIs* and should not be taken as general conclusions. Rather, they show how our framework enables structured, *SoI*-specific evaluations, helping practitioners assess which methods may be more suitable for their own modeling context.

6.5 EXPERIMENT 2: COUNTERFACTUAL ESTIMATION ON DISCRETE SCMS

This experiment evaluates counterfactual estimation on discrete-variable SCMs as a robustness check, testing CausalNF and DCM—originally designed for continuous settings—motivated by prior work showing that CausalNF can approximate discrete distributions (Javaloy et al., 2023; de Vassimon Manela et al., 2024). We consider three discrete *SoIs*: **Disc-C2-Reject**, with 10-15 nodes, binary variables, and rejection-based mechanism sampling; **Disc-C4-Unbias**, with the same graph size but 4-category variables and unbiased random mechanism sampling; and **Disc-L-C2-Unbias**, with larger graphs (20-30 nodes), binary variables, and unbiased random mechanism sampling (Table 2).

Table 2: Performance summary of CausalNF and DCM on the discrete experiments.

Space	Method	Mean Error	Std Error	Max Error	Runtime	Fail Rate
Disc-C2-Reject	CausalNF	0.0415	0.1116	0.6240	212.8 s	08.08 %
	DCM	0.0424	0.1123	0.6240	4406.2 s	04.28 %
Disc-C4-Unbias	CausalNF	0.0431	0.1270	0.7071	190.7 s	40.68 %
	DCM	0.0411	0.1199	0.7071	3839.4 s	22.60 %
Disc-L-C2-Unbias	CausalNF	NaN	NaN	NaN	0.0 s	100.00 %
	DCM	0.0183	0.0814	0.5000	8192.7 s	11.32 %

Findings. On **Disc-C2-Reject**, both CausalNF and DCM perform well and comparably, with low error means (~ 0.04) and low failure rates (8% for CausalNF, 4% for DCM). This suggests that both methods can produce reliable estimates even outside their original assumptions when the functional mechanisms are simple and binary. However, when moving to **Disc-C4-Unbias**, where variables have 4 categories and mechanisms are sampled with unbiased random sampling, the failure rates increase significantly, especially for CausalNF, which fails on over 40% of SCMs (typically with NaN errors). This highlights the sensitivity of some methods to mechanism sampling or variable cardinality, even when mean errors remain similar. To further probe robustness, we scale the graph size in **Disc-L-C2-Unbias** while reverting to binary variables. CausalNF fails on all runs, returning NaNs. DCM has an 11% failure rate, indicating greater resilience in this setting.

Insights. These results underscore the utility of our framework in systematically stress-testing methods beyond their nominal design assumptions. While CausalNF is not built for discrete data, prior work suggested it could work in practice. Our framework can help clarify *when* and *how* it fails: certain function classes and discrete configurations are more likely to cause divergence or failure.

DCM appears more robust across these tests, though not immune. Importantly, this evaluation is not meant as a definitive comparison, but as a demonstration of how failure cases can be surfaced and studied in a principled way using the CausalProfiler.

7 LIMITATIONS AND FUTURE WORK

We note that any open-source framework such as CausalProfiler is never a completely finished project, but rather continuously evolving to meet community needs, with new features added as the field advances through contributions to the repository.

Diversify Spaces of Interest. Several directions remain open for extending the supported *SoIs* in CausalProfiler, such as support for scaled and mixed-variable SCMs, sampling interventional training data, and more realistic data-generating scenarios, including selection bias or measurement noise. Another direction is to extend beyond tabular data to time-series, images, and text.

Causal Datasets Distribution. While the coverage proposition (Proposition 5.1) guarantees that any causal dataset has a positive probability of being sampled within a given *SoI* with sufficiently expressive discrete mechanisms, it does not characterize the distribution of generated datasets. As presented in Appendix H, certain classes of SCMs remain unlikely to be sampled unless explicitly specified in the *SoI* (e.g., linear SCMs). Hence, when aggregating results, users should bear in mind that causal datasets are not distributed uniformly to avoid misleading interpretations. We strongly recommend users to use the analysis module, presented in Appendix G, to identify the underrepresented attributes, as these vary from one *SoI* specification to another.

Reducing distributional bias is an important future research direction. Achieving a perfectly balanced distribution over all metrics is inherently impossible. For instance, uniform sampling over discrete mechanism functions biases toward non-bijective ones, since bijections are not dense in the function space. Future work may enable finer control over dataset distributions and underrepresented attributes, depending on the guarantees one wishes to enforce. One promising avenue is stratified sampling, which would provide weighted coverage of selected attributes. Currently, controllable *SoI* parameters (e.g., number of nodes) are sampled uniformly, but emergent attributes follow skewed distributions induced by generation. For controllable *SoI* parameters, stratification could be achieved constructively via weighted sampling over groups of *SoIs*. For emergent properties, approximate stratification may require rejection sampling or, more efficiently, new sampling algorithms that enforce global constraints during generation.

Bridging the simulation-to-real gap. While synthetic evaluation is indispensable (Poinsot et al., 2025), it is insufficient to fully assess method capabilities, as results may not transfer to real-world settings. In CausalProfiler, alignment with real domains currently relies on manually specified *SoIs*, guided by domain expertise or empirical features. A key direction for future work is to develop methods that automatically map real data to *SoIs*, enabling principled semi-synthetic evaluation pipelines where *SoIs* are shaped by empirical evidence rather than fixed assumptions. However, mapping from observational data to *SoIs* is a fundamentally underconstrained problem, and any such inference must be handled with care, given the challenges around identifiability and inductive bias.

8 CONCLUSION

This work introduces CausalProfiler, a synthetic causal dataset generator for evaluating Causal ML methods across the three levels of the Pearl Causal Hierarchy. At its core is the notion of a *Space of Interest*, which replaces the ad hoc choice of fixed evaluation datasets with a principled specification of the entire evaluation scope, i.e., classes of causal models, queries and data. This shift enables transparent, repeatable, and assumption-aware assessments under diverse causal conditions. After demonstrating that the causal datasets generated by CausalProfiler are correct and can be similar to existing benchmarks while also being considerably more diverse, we show that the performance of state-of-the-art Causal ML methods varies substantially across different *Spaces of Interest*, underscoring the importance of rigorous, distribution-level evaluation. CausalProfiler is not intended to replace real-data studies or targeted evaluations, but to complement them. By enabling systematic exploration, it helps uncover failure modes, expose robustness to violated assumptions, and highlight unexpected strengths that may motivate new research directions. In this way, CausalProfiler marks a first step toward a more complete evaluation ecosystem for Causal ML.

486 REPRODUCIBILITY STATEMENT
487

488 We have taken extensive measures to ensure the reproducibility of our results. The paper specifies
489 fully the steps required to reproduce our experiments, with pseudocode for all algorithms provided
490 in the appendices. All experimental configurations are also documented in the appendices. An
491 anonymized zip archive containing the full code and reproduction instructions is included in the
492 supplementary materials. The codebase reflects the exact setup used in the reported experiments.
493 Upon acceptance, we will publicly release the codebase on GitHub. We note that no external datasets
494 are required to reproduce the experiments. We also specify the hardware used and report runtime
495 metrics, making computational requirements transparent.

496
497 ETHICS STATEMENT
498

499 This work introduces CausalProfiler, a synthetic benchmark generator for evaluating Causal ML
500 methods. As a methodological tool rather than an application-facing system, it does not directly
501 raise societal impact concerns to the best of our knowledge. Furthermore, to prevent naive use of
502 CausalProfiler, this work transparently outlines its guarantees and limitations. We also remind readers
503 of the simulation-to-real gap inherent to any synthetic system. To mitigate the risk of inadvertent
504 misuse of CausalProfiler, it is emphasized that evaluation results should not be aggregated and
505 interpreted naively without exploring the distribution of the generated causal datasets.

506 Finally, we do not release pretrained models or real-world datasets. We provide code that generates
507 fully synthetic data, thereby avoiding issues related to privacy, fairness, and security. The paper
508 involves no human subjects, crowdsourcing, or sensitive data.

509
510 REFERENCES
511

- 512 Oriol Abril-Pla, Virgile Andreani, Colin Carroll, Larry Dong, Christopher J Fongesbeck, Maxim
513 Kochurov, Ravin Kumar, Junpeng Lao, Christian C Luhmann, Osvaldo A Martin, et al. PyMC: a
514 modern, and comprehensive probabilistic programming framework in python. *PeerJ Computer
515 Science*, 9, 2023.
- 516 Alejandro Almodóvar, Adrián Javaloy, Juan Parras, Santiago Zazo, and Isabel Valera. DeCaFlow: A
517 deconfounding causal generative model. *arXiv:2503.15114*, 2025.
- 518 Susan Athey, Guido W Imbens, Jonas Metzger, and Evan Munro. Using wasserstein generative
519 adversarial networks for the design of monte carlo simulations. *Journal of Econometrics*, 240(2):
520 105076, 2024.
- 521 Elias Bareinboim, Juan D. Correa, Duligur Ibeling, and Thomas Icard. *On Pearl’s Hierarchy and the
522 Foundations of Causal Inference*, pp. 507–556. 2022.
- 523 Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful
524 approach to multiple testing. *Journal of the Royal Statistical Society*, 57(1):289–300, 1995.
- 525 Jeroen Berrevoets, Krzysztof Kacprzyk, Zhaozhi Qian, Mihaela van der Schaar, et al. Causal deep
526 learning: Encouraging impact on real-world problems through causality. *Foundations and Trends
527 in Signal Processing*, 18(3):200–309, 2024.
- 528 Patrick Chao, Patrick Blöbaum, and Shiva Prasad Kasiviswanathan. Interventional and counterfactual
529 inference with diffusion models. *arXiv:2302.00860*, 2023.
- 530 Lu Cheng, Ruocheng Guo, Raha Moraffah, Paras Sheth, K. Selçuk Candan, and Huan Liu. Evaluation
531 methods and measures for causal learning algorithms. *IEEE Transactions on Artificial Intelligence*,
532 3(6):924–943, 2022.
- 533 Alicia Curth, David Svensson, Jim Weatherall, and Mihaela van der Schaar. Really doing great at
534 estimating CATE? a critical look at ML benchmarking practices in treatment effect estimation. In
535 *Advances in Neural Information Processing Systems*, 2021.

- 540 Alicia Curth, Richard W. Peck, Eoin McKinney, James Weatherall, and Mihaela van der Schaar.
541 Using machine learning to individualize treatment effect estimation: Challenges and opportunities.
542 *Clinical Pharmacology & Therapeutics*, 115(4):710–719, 2024.
- 543
544 Daniel de Vassimon Manela, Laura Battaglia, and Robin J. Evans. Marginal causal flows for validation
545 and inference. In *Advances in Neural Information Processing Systems*, 2024.
- 546
547 Tim Van den Bulcke, Koenraad Van Leemput, Bart Naudts, Piet van Remortel, Hongwu Ma, Alain
548 Verschoren, Bart De Moor, and Kathleen Marchal. Syntren: a generator of synthetic gene
549 expression data for design and analysis of structure learning algorithms. *BMC bioinformatics*, 7(1),
550 2006.
- 551
552 Vincent Dorie, Jennifer Hill, Uri Shalit, Marc Scott, and Dan Cervone. Automated versus do-it-
553 yourself methods for causal inference: Lessons learned from a data analysis competition. *Statistical
554 Science*, 34:43–68, 2019.
- 554
555 Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Alicia
556 Curth, Stefan Bauer, Niki Kilbertus, Isaac S. Kohane, and Mihaela van der Schaar. Causal machine
557 learning for predicting treatment outcomes. *Nature Medicine*, 30(4):958–968, 2024.
- 558
559 Timo Freiesleben and Thomas Grote. Beyond generalization: A theory of robustness in machine
560 learning. *Synthese*, 202(4):109, 2023.
- 560
561 Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias
562 Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine
563 Intelligence*, 2(11):665–673, 2020.
- 564
565 Amanda Gentzel, Dan Garant, and David Jensen. The case for evaluating causal models using
566 interventional measures and empirical data. *Advances in Neural Information Processing Systems*,
567 2019.
- 567
568 Sander Greenland and Babette Brumback. An overview of relations among causal modelling methods.
569 *International Journal of Epidemiology*, 31(5):1030–1037, 2002.
- 570
571 Shantanu Gupta, Cheng Zhang, and Agrin Hilmkil. Learned causal method prediction.
572 *arXiv:2311.03989*, 2023.
- 572
573 Moritz Herrmann, F. Julian D. Lange, Katharina Eggensperger, Giuseppe Casalicchio, Marcel Wever,
574 Matthias Feurer, David Rügamer, Eyke Hüllermeier, Anne-Laure Boulesteix, and Bernd Bischl.
575 Position: Why we must rethink empirical research in machine learning. In *International Conference
576 on Machine Learning*, 2024.
- 576
577 Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational
578 and Graphical Statistics*, 20(1):217–240, 2011.
- 579
580 Paul W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81
581 (396):945–960, 1986.
- 582
583 Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):
584 251–257, 1991.
- 584
585 Ben Hutchinson, Negar Rostamzadeh, Christina Greer, Katherine Heller, and Vinodkumar Prab-
586 hakaran. Evaluation gaps in machine learning practice. In *ACM Conference on Fairness, Account-
587 ability, and Transparency*, 2022.
- 588
589 Adrián Javaloy, Pablo Sanchez-Martin, and Isabel Valera. Causal normalizing flows: from theory to
590 practice. In *Advances in Neural Information Processing Systems*, 2023.
- 590
591 Jean Kaddour, Aengus Lynch, Qi Liu, Matt J Kusner, and Ricardo Silva. Causal machine learning: A
592 survey and open problems. *arXiv:2206.15475*, 2022.
- 593
Diviyani Kalainathan, Olivier Goudet, and Ritik Dutta. Causal discovery toolbox: Uncovering causal
relationships in python. *Journal of Machine Learning Research*, 21(37):1–5, 2020.

- 594 Kenneth J. Koehler and Kinley Larntz. An empirical investigation of goodness-of-fit statistics for
595 sparse multinomials. *Journal of the American Statistical Association*, 75(370):336–344, 1980.
596
- 597 Rachel Longjohn, Markelle Kelly, Sameer Singh, and Padhraic Smyth. Benchmark data repositories
598 for better benchmarking. In *Advances in Neural Information Processing Systems Datasets and*
599 *Benchmarks Track*, 2024.
- 600 Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine*
601 *learning research*, 9(Nov):2579–2605, 2008.
602
- 603 Brady Neal, Chin-Wei Huang, and Sunand Raghupathi. Realcause: Realistic causal inference
604 benchmarking. *arXiv:2011.15007*, 2020.
- 605 Harsh Parikh, Carlos Varjao, Louise Xu, and Eric Tchetgen Tchetgen. Validating causal inference
606 methods. In *International Conference on Machine Learning*, 2022.
607
- 608 Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition,
609 2009.
- 610 Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic books,
611 2018.
612
- 613 Karl Pearson. On the criterion that a given system of deviations from the probable in the case of
614 a correlated system of variables is such that it can be reasonably supposed to have arisen from
615 random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of*
616 *Science*, 50(302):157–175, 1900.
617
- 618 Anne Helby Petersen. Are you doing better than random guessing? a call for using negative controls
619 when evaluating causal discovery algorithms. *arXiv:2412.10039*, 2024.
- 620 Audrey Poinot, Alessandro Leite, Nicolas Chesneau, Michèle Sébag, and Marc Schoenauer. Learning
621 structural causal models through deep generative models: Methods, guarantees, and challenges. In
622 *International Joint Conference on Artificial Intelligence*, 2024.
623
- 624 Audrey Poinot, Panayiotis Panayiotou, Alessandro Leite, Nicolas CHESNEAU, Özgür Şimşek, and
625 Marc Schoenauer. Position: Causal machine learning requires rigorous synthetic experiments for
626 broader adoption. In *International Conference on Machine Learning, Position Paper Track*, 2025.
- 627 Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua
628 Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International conference*
629 *on machine learning*, 2019.
- 630 Kara E Rudolph, Nicholas T Williams, Caleb H Miles, Joseph Antonelli, and Ivan Diaz. All models
631 are wrong, but which are useful? comparing parametric and nonparametric estimation of causal
632 effects in finite samples. *Journal of Causal Inference*, 11(1), 2023.
633
- 634 Marco Scutari. Package ‘bnlearn’. bayesian network structure learning, parameter learning and
635 inference, r package version, 2019. URL <https://www.bnlearn.com/>.
636
- 637 Rajen D. Shah and Jonas Peters. The hardness of conditional independence testing and the generalised
638 covariance measure. *The Annals of Statistics*, 48(3), 2020.
- 639 Amit Sharma and Emre Kiciman. DoWhy: An end-to-end library for causal inference.
640 *arXiv:2011.04216*, 2020.
641
- 642 Yishai Shimoni, Chen Yanover, Ehud Karavani, and Yaara Goldschmidt. Benchmarking framework
643 for performance-evaluation of causal inference analysis. *arXiv:1802.05046*, 2018.
- 644 Beate Sick and Oliver Dürr. Interpretable neural causal models with tram-dags. In *Conference on*
645 *Causal Learning and Reasoning*, 2025.
646
- 647 Pablo Sánchez-Martin, Miriam Rateike, and Isabel Valera. Vaca: Designing variational graph
autoencoders for causal queries. *AAAI Conference on Artificial Intelligence*, 2022.

648 Peter WG Tennant, Eleanor J Murray, Kellyn F Arnold, Laurie Berrie, Matthew P Fox, Sarah C Gadd,
649 Wendy J Harrison, Claire Keeble, Lysie R Ranker, Johannes Textor, et al. Use of directed acyclic
650 graphs (DAGs) to identify confounders in applied health research: review and recommendations.
651 *International Journal of Epidemiology*, 50(2):620–632, 2021.

652 Johannes Textor, Benito van der Zander, Mark S Gilthorpe, Maciej Liškiewicz, and George TH Ellison.
653 Robust causal inference using directed acyclic graphs: the r package ‘dagitty’. *International
654 Journal of Epidemiology*, 45(6):1887–1894, 2017.

655
656 Jin Tian and Judea Pearl. A general identification condition for causal effects. In *AAAI Conference
657 on Artificial Intelligence*, 2002.

658 Thomas Sadanand Verma. Graphical aspects of causal models. In *UCLA Cognitive Systems Labora-
659 tory, Technical Report (R-191)*, 1993.

660
661 Kevin Xia, Yushu Pan, and Elias Bareinboim. Neural causal models for counterfactual identification
662 and estimation. In *International Conference on Learning Representations*, 2023.

663
664 Junzhe Zhang, Jin Tian, and Elias Bareinboim. Partial counterfactual identification from observational
665 and experimental data. In *International Conference on Machine Learning*, 2022.

666
667 Qingyang Zhou, Kangjie Lu, and Meng Xu. Causally consistent normalizing flow. *AAAI Conference
668 on Artificial Intelligence*, 2025.

669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

A USE OF LARGE LANGUAGE MODELS (LLMs)

We used LLMs, specifically ChatGPT, as a writing assistant. The model was used only to help with language-related aspects of the paper, including:

- Rephrasing existing content without changing its meaning
- Improving clarity and flow
- Identifying issues such as unclear points, unintended tones, or awkward phrasing.

All scientific contributions originate from the authors, who take full responsibility for the paper.

B ADDITIONAL DEFINITIONS & NOTATIONS

Definition B.1 (Semi-Markovian and Markovian SCMs). An SCM is said to be **semi-Markovian** (Pearl, 2009) if its set of structural equations is acyclic, meaning there exists an ordering of the equations such that for any two functions $f_i, f_j \in \mathcal{F}$, if $f_i < f_j$, then $V_j \notin \mathbf{PA}(V_i)$. This condition ensures that the causal dependencies among endogenous variables form a Directed Acyclic Graph.

An SCM is **Markovian** (Pearl, 2009) if the exogenous variables influencing different endogenous variables are mutually independent. Formally, for all distinct $V_i, V_j \in \mathbf{V}$, we have $\mathbf{U}_{V_i} \perp\!\!\!\perp \mathbf{U}_{V_j}$. This implies the absence of latent confounding, allowing the model to be fully described by a DAG with independent noise terms.

Definition B.2 (Causal Graph of a Semi-Markovian SCM). The causal graph of a Semi-Markovian (Bareinboim et al., 2022) SCM is an acyclic directed mixed graph with:

- Directed edge $V_i \rightarrow V_j$ if $V_i \in \mathbf{PA}(V_j)$
- Bi-directed edge $V_i \leftrightarrow V_j$ if $\mathbf{U}_{V_i} \not\perp\!\!\!\perp \mathbf{U}_{V_j}$

B.1 INTERVENTIONAL QUANTITIES (\mathcal{L}_2)

Average Treatment Effect (ATE):

$$\text{ATE}_{T \rightarrow Y} = \mathbb{E}[Y|\mathbf{do}(T = 1)] - \mathbb{E}[Y|\mathbf{do}(T = 0)]$$

Conditional Average Treatment Effect (CATE):

$$\text{CATE}_{T \rightarrow Y}(\mathbf{x}) = \mathbb{E}[Y|\mathbf{do}(T = 1), \mathbf{X} = \mathbf{x}] - \mathbb{E}[Y|\mathbf{do}(T = 0), \mathbf{X} = \mathbf{x}]$$

B.2 COUNTERFACTUAL QUANTITIES (\mathcal{L}_3)

A counterfactual query such as $P(Y_{\mathbf{do}(T=t)}|\mathbf{V}_F = \mathbf{v}_F)$ is computed by abduction (conditioning on factual data), action (intervening), and prediction (computing the outcome) (Pearl, 2009).

Counterfactual Total Effect (Ctf-TE):

$$\text{Ctf-TE}_{T \rightarrow Y}(y, t, c, \mathbf{v}_F) = P(y_{\mathbf{do}(T=t)}|\mathbf{V}_F = \mathbf{v}_F) - P(y_{\mathbf{do}(T=c)}|\mathbf{V}_F = \mathbf{v}_F)$$

C SPACE OF INTEREST

Each Space of Interest is defined by a set of parameters that control the *SCM space*, the causal queries of interest (*Query space*), and the dataset used for estimation (*Data space*). Table 3 provides an overview of all configurable parameters in a Space of Interest instance, along with their default values. Some parameters are only relevant under specific conditions—for instance, kernel parameters are used only with continuous variables (e.g., when evaluating conditional expectations), function

sampling strategies apply exclusively to discrete mechanisms, noise regions apply only for discrete SCMs, and noise mode is ignored for tabular mechanisms (noise is already embedded in the table). Note that one can use symbolic expressions involving N (the number of nodes) and V (the cardinality of a variable) to define parameters that depend on sampled values. For example, the expected number of edges can be set as $0.5 * N$, or the number of noise regions in a discrete SCM can be set to V .

Table 3: Parameters defining a Space of Interest instance and their default values. The double lines in the table conceptually separate the SCM space, Query space, and Data space.

Category	Parameter	Default Value
SCM structure	Number of endogenous variables	[5, 15]
	Variable dimensionality	[1, 1]
	Expected number of edges (required)	—
	Proportion of hidden variables	0.0
	Markovian boolean flag	True
	Semi-Markovian boolean flag	False
	Predefined causal graph	—
Mechanisms	Mechanism family (e.g., Linear, NN, Tabular)	Linear
	Mechanism arguments (used to define custom NN/tabular mechanisms)	—
	Endogenous variable cardinality (for discrete variables only)	2
	Variable type	Continuous
	Discrete function sampling (for discrete variables only)	Sample Rejection
Noise	Noise mode	Additive
	Noise distribution	Uniform
	Noise distribution arguments	[-1, 1]
Query	Number of noise regions (for discrete variables only)	N
	Number of queries per sample	1
	Query type	ATE
	Specific query (overrides random query sampling)	—
	Whether to allow queries that evaluate to NaN	False
Kernel	Whether to disable query sampling (e.g., for causal discovery)	False
	Kernel type	Gaussian
	Kernel bandwidth	0.1
Data	Custom kernel function	—
	Number of samples in the set of observed data	1000

D CAUSAL GRAPH SAMPLING

We first generate a random Directed Acyclic Graph (DAG) that specifies causal relations between variables. This structure is then extended by designating a subset of variables as hidden/unobserved, enabling the creation of both Markovian and semi-Markovian SCMs depending on the *SoI* spec. We separate these two steps in separate algorithms for clarity (Algorithm 2 uses Algorithm 1).

First, Algorithm 1 samples a DAG over a unique type of variables, not yet distinguishing between observable and unobservable variables. To do so, the list of nodes is defined as a list of integers imposed to be the topological order of the DAG (line 1). Then, for each node (line 4), its number of parents is sampled from a Binomial law of parameters $i - 1$ and p_{edge} with i the rank of the node in the topological order (line 5). The actual parents are sampled from the set of nodes having a smaller topological rank (line 6) which guarantees that the generated graph is a DAG.

Second, from the generated DAG, Algorithm 2 simply creates the two sets of observables and unobservable variables by sampling $p_h \cdot |\mathbf{V}|$ unobservable node among the total set of nodes (line 3).

Algorithm 1 Generate a Random DAG with Expected Degree**Inputs:** number of nodes N , expected degree d

```

810 1:  $V \leftarrow \{1, \dots, N\}$ 
811 2:  $E \leftarrow \{\}$ 
812 3:  $p_{edge} \leftarrow \frac{2d}{N-1}$ 
813 4: for  $i \in [1, N]$  do
814 5:    $N_{PA(i)} \sim B(i-1, p_{edge})$ 
815 6:    $PA(i) \leftarrow N_{PA(i)}$  nodes sampled without replacement from  $V$ 
816 7:    $E \leftarrow E \cup \{j \rightarrow i \mid j \in PA(i)\}$ 
817 8: end for
818 Output:  $\mathcal{G} = \{V, E\}$ 

```

Algorithm 2 Generate a DAG with Observed and Hidden Variables**Inputs:** number of nodes N , expected degree d , proportion of hidden variables p_h

```

826 1:  $\mathcal{G} = (V, E) \leftarrow DAG\_sampling(N, d)$  (see Algorithm 1)
827 2:  $N_h \sim B(N, p_h)$ 
828 3:  $V_h \leftarrow N_h$  nodes sampled without replacement from  $V$ 
829 4:  $V_o \leftarrow V \setminus V_h$ 
830 Output:  $\mathcal{G} = \{V = V_o V_h, E\}$ 

```

Because some variables in the DAG are unobserved, we expose only the observed structure to the user in the form of an acyclic directed mixed graph. To obtain this, we apply Verma’s latent projection algorithm to the causal graph of each sampled regional discrete SCM (see Algorithm 3). If a method requires the true SCM, including the hidden confounders, that can be accessed as well.

Algorithm 3 Projection Algorithm (Verma, 1993)**Input:** an acyclic directed mixed graph $\mathcal{G} = \{\mathbf{V}_O, \mathbf{V}_H, \mathbf{E}\}$, with \mathbf{V}_O the set of observed variables, \mathbf{V}_H the set of hidden variables and \mathbf{E} the mixed edges

```

841 1:  $\mathbf{E}' \leftarrow \{\}$ 
842 2: for  $A, B \in \mathbf{V}_O$  do
843 3:   if there is a directed path  $A \rightarrow \dots \rightarrow B$  in  $\mathcal{G}$  with all intermediate nodes belonging to  $\mathbf{V}_H$  then
844 4:    $\mathbf{E}' \leftarrow \mathbf{E}' \cup \{A \rightarrow B\}$ 
845 5:   end if
846 6:   if there is a collider-free path  $A \leftarrow \dots \rightarrow B$  in  $\mathcal{G}$  with all intermediate nodes belonging to  $\mathbf{V}_H$  then
847 7:    $\mathbf{E}' \leftarrow \mathbf{E}' \cup \{A \leftrightarrow B\}$ 
848 8:   end if
849 9: end for
850 10:  $\mathbf{G}' \leftarrow \{\mathbf{V}_O, \mathbf{E}'\}$ 
851 Output:  $\mathbf{G}'$  the latent projection of  $\mathbf{G}$  over  $\mathbf{V}_O$ 

```

E SAMPLING DISCRETE SCMS

E.1 REGIONAL DISCRETE SCMS

Regarding discrete SCMs, we sample discrete Markovian SCMs which we refer to as **Regional discrete SCMs** as presented in definition E.1 below.

Definition E.1. Regional discrete SCM

A **regional discrete SCM** is a markovian SCM $\mathcal{M} := \{\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\}$ where:

- $\mathbf{V} = \{V_1, \dots, V_d\}$ the set of finite discrete endogenous variables is divided into two sets \mathbf{V}_o and \mathbf{V}_h respectively representing the set of observed and hidden variables such that $\mathbf{V} = \mathbf{V}_o \cup \mathbf{V}_h$ and $\mathbf{V}_o \cap \mathbf{V}_h = \emptyset$
- $\mathbf{U} = \{U_1, \dots, U_d\}$ the set of mutually independent continuous exogenous variables is such that $\forall i \in [1, d], U_{V_i} = U_i$
- \mathcal{F} the structural equations are regional discrete mechanisms as defined in Definition E.2

The class of regional discrete SCMs is denoted $\mathbb{M}_{\text{RD-SCM}}$.

Definition E.2. Regional discrete mechanism

Given $\mathbf{I}_V = \{I_V^r\}_{r \in [1, R]}$ a partition of R parts of Ω_{U_V} and $m_V = \{m_V^r : \Omega_{\mathbf{PA}(V)} \mapsto \Omega_V\}_{r \in [1, R]}$ a set of R distinct mappings from $\Omega_{\mathbf{PA}(V)}$ to Ω_V , the **regional discrete mechanism** of an endogenous variables V is a function $f_V : \Omega_{\mathbf{PA}(V)}, \Omega_{U_V} \mapsto \Omega_V$ such that:

$$f_V(\mathbf{pa}(V), u_V) = m_r(\mathbf{PA}(V) \mapsto V) \text{ when } u_V \in I_V^r$$

I_V^r and m_r are called the r^{th} noise region and mapping of the regional discrete mechanism f_V .

Remark on Ω_{U_V} and R : In the definition of a regional discrete mechanism (Definition E.2), no constraints are imposed on Ω_{U_V} . However, if Ω_{U_V} is discrete, then $|\Omega_{U_V}| \geq R$ is required to form a partition of R elements of Ω_{U_V} . Consequently, in order to be able to constitute such a partition for any finite R , we decided to consider continuous exogenous variables in the definition of a regional discrete SCM (Definition E.1). In addition, since the m_V^r mappings are considered distinct and there are exactly $|\Omega_V|^{|\Omega_{\mathbf{PA}(V)}|}$ different mappings from V to $\mathbf{PA}(V)$, $R \leq |\Omega_V|^{|\Omega_{\mathbf{PA}(V)}|}$ is required.

Even if regional discrete SCMs are Markovian, the fact that they contains two types of endogenous variables (i.e., observed and unobserved by the user) enables the representation of complex situations where not all variables are observable. This induces the presence of potential hidden confounders from the user’s perspective. As a result, the causal sufficiency assumption is no longer always respected. In our parametric definition of a *SoI*, this phenomenon is controlled by the parameter specifying the proportion of unobserved variables among the endogenous variables. Thus, if this parameter is set to 0, the *SoI*’s class of SCMs is included in the class of causally sufficient discrete SCMs.

The complexity of discrete mechanisms can be controlled by the number of noise regions R . Indeed, as the number of noise regions increases, so does the complexity of the causal mechanism, in the sense that it becomes a mixture of a larger number of mappings. The distribution of a variable given its parents is, hence, more stochastic. As a result, the user-defined class of regional discrete SCMs can be very broad. This provides an additional degree of complexity to make our synthetic causal datasets less trivial.

The class of regional discrete SCMs got inspired by the class of Regional Canonical Models by (Xia et al., 2023) and the class of canonical SCMs by (Zhang et al., 2022). We decided to define our own class rather than using one of these two classes for two reasons. First, canonical SCMs are very expensive to sample particularly because of the presence of confounded components. Second, even if Regional Canonical Models are designed to be less expensive because their expressivity can be regulated via the number of noise regions to consider, they lose some interesting properties such as the non overlapping of the noise regions which is crucial to favor a strong dependence between the user choice of the number of noise regions and the complexity of the generated mechanisms. Moreover, Regional Canonical Models still rely on confounded components which is the major source

of complexity at the sampling stage. Hence, we defined the class of Regional discrete SCMs to not have to deal with confounded components at the sampling stage (instead we rely on a projection algorithm after sampling, see Appendix D) and to regulate mechanisms expressivity through the use of non-overlapping noise regions.

E.2 DISCRETE MECHANISM SAMPLING STRATEGIES

We use *regional discrete mechanisms* (Definition E.2), which define tabular mappings from parent variables to a target variable, conditioned on regions of the exogenous noise space. By default, each region induces a distinct mapping, enabling both stochasticity and high functional expressivity.

To generate these mechanisms, we support three sampling strategies described below. All methods define a partition of the exogenous noise domain Ω_U into R regions, and assign a parent-to-child mapping to each region. Let C be the cardinality of the variables, and $\Omega_{\text{pa}(V)}$ the space of parent configurations for variable V .

Controlling complexity. The number of possible mappings from parent configurations to output values grows as $|\Omega_V|^{|\Omega_{\text{pa}(V)}|}$. To keep simulations tractable, users can control the number of noise regions R . When R is small, sampling provides diverse but lightweight mechanisms. When R approaches the total number of mappings, full enumeration becomes feasible but computationally expensive.

We now describe the three supported sampling strategies.

EXHAUSTIVE PARTITION

This strategy enumerates all possible mappings from parent configurations to output values and assigns each one to a distinct noise region ($R = |\Omega_V|^{|\Omega_{\text{pa}(V)}|}$), ensuring complete coverage of the function space. This method guarantees maximal functional diversity across regions and can serve as a stress test for generalization under highly non-linear mechanisms. This is the only strategy where the number of noise regions is not decided by the user but rather set to the maximum. The exhaustive partition sampling strategy is the one to use if one wants the coverage guarantee (Proposition 5.1) to apply.

SAMPLE REJECTION

This strategy samples parent-to-output mappings uniformly at random, rejecting duplicates to ensure that each region corresponds to a distinct function. As mappings are sampled with replacement, rejection may require several attempts when R approaches the number of possible mappings.

We provide below, in Algorithm 4, a pseudocode version of this strategy. The algorithm proceeds as follows. For each endogenous variable V (line 2) a regional discrete mechanism is created. To do so, the domain of V is first initialized with a list of integers corresponding of the cardinality specified in the *SoI* (line 3). Then, if the number of noise regions R specified in the *SoI* is larger than the maximum number of noise regions, the maximum number of noise regions is used to generate the regional discrete mechanism (lines 4-5). The partition of the noise regions is built as consecutive intervals of random size resulting from the ordering of $R - 1$ sampled realizations of the uniform exogenous distribution (lines 6 to 9 and 13). Finally, for each noise region r (line 12), mappings m_V^r are sampled till one mapping not already used for other noise regions is sampled (lines 15 to 18). This is why this algorithm is denoted as the “sample rejection” approach. One can note that there are two sources of randomness in this algorithm: the size of the noise regions and the sampled mappings whenever the number of noise regions is not maximal.

Algorithm 4 Generating regional discrete mechanisms with sample rejection

Inputs: set of endogenous variables \mathbf{V} of cardinality C , causal graph \mathcal{G} , Ω_U domain of exogenous variables, number of noise regions R

```

1:  $\mathcal{F} \leftarrow \{\}$ 
2: for  $V \in \mathbf{V}$  do
3:    $\Omega_V \leftarrow \{1, \dots, C\}$ 
4:    $\Omega_{PA_{\mathcal{G}}(V)} \leftarrow \{1, \dots, C\}^{|\text{PA}_{\mathcal{G}}(V)|}$ 
5:    $R \leftarrow \min(R, |\Omega_V|^{|\Omega_{PA_{\mathcal{G}}(V)}|})$ 
6:    $l_{\min} \leftarrow \inf(\Omega_U)$ 
7:    $l_{\max} \leftarrow \sup(\Omega_U)$ 
8:    $\mathbf{L} = \{l_i \sim \mathcal{U}[l_{\min}, l_{\max}] \mid i \in [1, R-1]\} \cup \{l_{\min}, l_{\max}\}$ 
9:   Sort  $\mathbf{L}$  in ascending order
10:   $f_V \leftarrow \{\}$ 
11:   $m_V \leftarrow \{\}$ 
12:  for  $r \in [1, R]$  do
13:     $I_V^r \leftarrow [\mathbf{L}_r, \mathbf{L}_{r+1}[$  with  $\mathbf{L}_r$  the  $r^{\text{th}}$  element of  $\mathbf{L}$ 
14:     $m_V^r \leftarrow \{\}$ 
15:    while  $m_V^r = \{\}$  or  $m_V^r \in m_V$  do
16:       $m_V^r \leftarrow |\Omega_{PA(V)}|$  elements sampled with replacement from  $\Omega_V$ 
17:    end while
18:     $m_V \leftarrow m_V \cup m_V^r$ 
19:     $f_V \leftarrow f_V \cup \{m_V^r; I_V^r\}$ 
20:  end for
21:   $\mathcal{F} \leftarrow \mathcal{F} \cup f_V$ 
22: end for
Output:  $\mathcal{F}$ 

```

UNBIASED RANDOM ASSIGNMENT

In this strategy, each noise region is assigned a mapping sampled independently and without enforcing uniqueness. As a result, multiple regions may correspond to the same function from parent configurations to outputs.

For example, suppose a variable has one binary parent taking values in $\{0, 1\}$, and the output variable takes values in $\{0, 1, 2\}$. One randomly sampled mapping might assign output 0 to parent value 0, and output 2 to parent value 1. Since mappings are sampled independently for each region, this same function ($0 \rightarrow 0, 1 \rightarrow 2$) may appear in multiple regions by chance.

This approach reflects scenarios where mechanisms are drawn independently from a distribution over functions, without enforcing any requirements on uniqueness or coverage. As a result, the effective variability in the entire system may be lower compared to other strategies, but the sampling is a lot more computationally efficient.

F QUERY SAMPLING AND ESTIMATION

In this work, we consider the following types of queries: Average Treatment Effect (ATE), Conditional Average Treatment Effect (CATE) and Counterfactual Total Effect (Ctf-TE). Their definitions can be found in Appendix B. All the queries can be defined for sets of covariates and factuials belonging to the set of endogenous variables. In other words, we do not implement multi-interventions, but we consider conditioning and observing factuials on several variables. Finally, the values taken by these variables (e.g., treatment and control values for ATE) must belong to their definition domain. The only parameter that controls the queries class is the type of queries chosen by the user (i.e., ATE, CATE and Ctf-TE). Thus, the class of considered queries can be defined as follows:

$$\mathcal{Q}_{\text{ATE}} = \{\text{ATE}_{T \rightarrow Y}(t, c) \mid T, Y \subseteq \mathbf{V} \text{ and } t, c \in \Omega_T\}$$

$$\mathcal{Q}_{\text{CATE}} = \{\text{CATE}_{T \rightarrow Y | \mathbf{X}}(t, c, \mathbf{x}) \mid T, Y \subseteq \mathbf{V}, \mathbf{X} \subseteq \mathbf{V} \setminus \{T, Y\} \text{ and } t, c \in \Omega_T, \mathbf{x} \in \Omega_{\mathbf{X}}\}$$

$$\mathcal{Q}_{\text{Ctf-TE}} = \{\text{Ctf-TE}_{T \rightarrow Y}(y, t, c, \mathbf{v}_F) \mid T, Y, \mathbf{V}_F \subseteq \mathbf{V} \text{ and } t, c \in \Omega_T, y \in \Omega_Y, \mathbf{v}_F \in \Omega_{\mathbf{V}_F}\}$$

Formally speaking, we have not integrated the causal graph as a causal query but rather as a hypothesis or prior knowledge. Indeed, except for causal discovery tasks, the causal graph is most often assumed to be known (or at least some information derived from the graph, such as the constitution of a valid adjustment set, or a valid causal ordering). Nevertheless, one can use our random causal dataset generator to evaluate causal discovery or causal representation learning methods. To do so, one just needs to retrieve the causal graph from the causal dataset directly instead of using a query.

Finally, a user can also implement a specific query and use it to generate synthetic causal datasets. To do this, the user has to use the `Query` class in our code base.

F.1 QUERY SAMPLING

As the values taken by variables in the queries have to belong to their definition domain, we draw realizations from a large, separately sampled observational dataset. Indeed, given the randomness of the causal mechanisms, we cannot know in advance the domain over which the SCMs are defined. Even when variable cardinalities are fixed, the sampled mechanisms may be non-surjective, making certain values impossible to observe. For this reason, we approximate the domain of definition through data sampling, ensuring that queries are computed only for realizable variable configurations. Moreover, since the dataset given to the user is smaller to the one we use for query sampling and estimation, it is possible that queries use values outside of the observational dataset or that they are non-identifiable. Explicitly enabling queries to be outside the observed dataset can be useful for studying generalization—especially in settings where the support is known, such as linear SCMs. However, we let for future work the development of a user-configurable option in *SoIs*, for instance, allowing users to define a custom domain for the query variables.

The following algorithms detail the procedures for sampling ATE, CATE, and Ctf-TE queries. In these algorithms, given a dataset D , a variable X and a realization x of X , we use the notation $D|_X$ (resp. $D|_{X=x}$) to represent the dataset D restricted to the variable X (resp. restricted to the samples whose X realization equals x). In addition, $B(n, p)$ denotes the Binomial law of parameters n and p .

Algorithm 5 Generating sets of observed data

Inputs: causal graph \mathcal{G} , causal mechanisms \mathcal{F} , distribution of the exogenous variables $P(\mathbf{U})$, dataset size N

```

1:  $D \leftarrow \{\}$ 
2:  $D_o \leftarrow \{\}$ 
3:  $\{\mathbf{u}_1, \dots, \mathbf{u}_N\} \sim P(\mathbf{U})$ 
4: for  $V \in \mathbf{V}$  following a causal order given by  $\mathcal{G}$  do
5:    $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D|_{\mathbf{pa}(V)}$ 
6:    $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D|_{\mathbf{U}_V}$ 
7:    $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$ 
8:    $D \leftarrow D \cup \{v_1, \dots, v_N\}$ 
9:   if  $V \in \mathbf{V}_o$  then
10:     $D_o \leftarrow D_o \cup \{v_1, \dots, v_N\}$ 
11:   end if
12: end for
Output:  $D_o$ 

```

Algorithm 6 Generating ATE queries

Inputs: set of observable endogenous variables \mathbf{V}_o , training set D

```

1:  $T \leftarrow$  one variable randomly sampled from  $\mathbf{V}_o$ 
2:  $Y \leftarrow$  one variable randomly sampled from  $\mathbf{V}_o$ 
3:  $t \leftarrow$  one realization of  $T$  randomly sampled from  $D|_T$ 
4:  $c \leftarrow$  one realization of  $T$  randomly sampled from  $D|_T$ 
Output:  $Q_{ATE} = \{T, Y, t, c\}$ 

```

Algorithm 7 Generating CATE queries**Inputs:** set of observable endogenous variables \mathbf{V}_o , training set D

- 1: $T \leftarrow$ one variable randomly sampled from \mathbf{V}_o
- 2: $Y \leftarrow$ one variable randomly sampled from \mathbf{V}_o
- 3: $d_{\mathbf{X}} \leftarrow$ an integer randomly sampled from $[1, \dots, |\mathbf{V}_o| - 2]$
- 4: $\mathbf{X} \leftarrow d_{\mathbf{X}}$ variables randomly sampled from $\mathbf{V}_o \setminus \{T, Y\}$
- 5: $t \leftarrow$ one realization of T randomly sampled from $D_{|T}$
- 6: $c \leftarrow$ one realization of T randomly sampled from $D_{|T}$
- 7: $\mathbf{x} \leftarrow$ one realization of \mathbf{X} randomly sampled from $D_{|\mathbf{X}}$

Output: $Q_{CATE} = \{T, Y, \mathbf{X}, t, c, \mathbf{x}\}$ **Algorithm 8** Generating Ctf-TE queries**Inputs:** set of observable endogenous variables \mathbf{V}_o , training set D

- 1: $T \leftarrow$ one variable randomly sampled from \mathbf{V}_o
- 2: $Y \leftarrow$ one variable randomly sampled from \mathbf{V}_o
- 3: $d_{\mathbf{V}_F} \leftarrow$ an integer randomly samples from $[1, \dots, |\mathbf{V}_o|]$
- 4: $\mathbf{V}_F \leftarrow d_{\mathbf{V}_F}$ variables randomly sampled from \mathbf{V}_o
- 5: $t \leftarrow$ one realization of T randomly sampled from $D_{|T}$
- 6: $c \leftarrow$ one realization of T randomly sampled from $D_{|T}$
- 7: $\mathbf{v}_F \leftarrow$ one realization of \mathbf{V}_F randomly sampled from $D_{|\mathbf{V}_F}$

Output: $Q_{CTF-TE} = \{T, Y, \mathbf{V}_F, t, c, \mathbf{v}_F\}$

F.2 SCM-BASED QUERY ESTIMATION

Each query is evaluated by modifying the SCM, sampling the exogenous variables, and computing expectations over the outcomes. In practice, we simulate interventions and counterfactuals by directly manipulating structural equations and conditioning on sampled variables. Our implementation supports efficient batch estimation using the same random seeds for reproducibility.

Queries that yield NaN estimates can optionally be rejected and resampled, depending on the *SoI* settings. NaN estimates appear if the corresponding sampled query is undefined (e.g., conditioning on a zero-probability event). However, to evaluate the ability of some models to identify if the query is undefined instead of trying to answer it, NaN estimates can be interesting to keep. This is why we decided to let users choose this option through a parameter of the *SoI*.

The following algorithms detail the procedures for estimating ATE, CATE, and Ctf-TE queries.

1134 **Algorithm 9** Estimating ATE queries

1135 **Inputs:** ATE query to estimate $Q = \{T, Y, t, c\}$, causal graph \mathcal{G} , causal mechanisms \mathcal{F} , distribution of the

1136 exogenous variables $P(\mathbf{U})$, number of samples to draw for estimation N

1137 1: $\{\mathbf{u}_1, \dots, \mathbf{u}_N\} \sim P(\mathbf{U})$

1138 2: $D_t \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$

1139 3: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1140 4: **if** $V = T$ **then**

1141 5: $\{v_1, \dots, v_N\} \leftarrow \{t, \dots, t\}$

1142 6: **else**

1143 7: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_{t|\mathbf{pa}(V)}$

1144 8: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_{t|\mathbf{U}_V}$

1145 9: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1146 10: **end if**

1147 11: $D_t \leftarrow D_t \cup \{v_1, \dots, v_N\}$

1148 12: **end for**

1149 13: $D_c \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$

1150 14: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1151 15: **if** $V = T$ **then**

1152 16: $\{v_1, \dots, v_N\} \leftarrow \{c, \dots, c\}$

1153 17: **else**

1154 18: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_{c|\mathbf{pa}(V)}$

1155 19: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_{c|\mathbf{U}_V}$

1156 20: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1157 21: **end if**

1158 22: $D_c \leftarrow D_c \cup \{v_1, \dots, v_N\}$

1159 23: **end for**

1160 24: $Q^* \leftarrow \text{avg}(D_{t|Y}) - \text{avg}(D_{c|Y})$

1161 **Output:** Q^*

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188 **Algorithm 10** Estimating CATE queries

1189 **Inputs:** CATE query to estimate $Q = \{T, Y, \mathbf{X}, t, c, \mathbf{x}\}$, causal graph \mathcal{G} , causal mechanisms \mathcal{F} , distribution of

1190 the exogenous variables $P(\mathbf{U})$, number of samples to draw for estimation N

1191 1: $\{\mathbf{u}_1, \dots, \mathbf{u}_N\} \sim P(\mathbf{U})$

1192 2: $D_t \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$

1193 3: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1194 4: **if** $V = T$ **then**

1195 5: $\{v_1, \dots, v_N\} \leftarrow \{t, \dots, t\}$

1196 6: **else**

1197 7: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_{t|\mathbf{pa}(V)}$

1198 8: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_{t|\mathbf{U}_V}$

1199 9: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1200 10: **end if**

1201 11: $D_t \leftarrow D_t \cup \{v_1, \dots, v_N\}$

1202 12: **end for**

1203 13: $D_c \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$

1204 14: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1205 15: **if** $V = T$ **then**

1206 16: $\{v_1, \dots, v_N\} \leftarrow \{c, \dots, c\}$

1207 17: **else**

1208 18: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_{c|\mathbf{pa}(V)}$

1209 19: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_{c|\mathbf{U}_V}$

1210 20: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1211 21: **end if**

1212 22: $D_c \leftarrow D_c \cup \{v_1, \dots, v_N\}$

1213 23: **end for**

1214 24: $D_t \leftarrow D_{t|\mathbf{X}=\mathbf{x}}$

1215 25: $D_c \leftarrow D_{c|\mathbf{X}=\mathbf{x}}$

1216 26: $Q^* \leftarrow \text{avg}(D_{t|Y}) - \text{avg}(D_{c|Y})$

1217 **Output:** Q^*

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

1242 **Algorithm 11** Estimating Ctf-TE queries

1243 **Inputs:** Ctf-TE query to estimate $Q = \{T, Y, \mathbf{V}_F, t, c, \mathbf{v}_F\}$, causal graph \mathcal{G} , causal mechanisms \mathcal{F} , distribution

1244 of the exogenous variables $P(\mathbf{U})$, number of samples to draw for estimation N

1245 1: $\{\mathbf{u}_1, \dots, \mathbf{u}_N\} \sim P(\mathbf{U})$

1246 2: $D_{\mathbf{U}_{\mathbf{v}_F}} \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$

1247 3: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1248 4: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_{\mathbf{U}_{\mathbf{v}_F} | \mathbf{pa}(V)}$

1249 5: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_{\mathbf{U}_{\mathbf{v}_F} | \mathbf{U}_V}$

1250 6: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1251 7: $D_{\mathbf{U}_{\mathbf{v}_F}} \leftarrow D_{\mathbf{U}_{\mathbf{v}_F}} \cup \{v_1, \dots, v_N\}$

1252 8: **end for**

1253 9: $D_{\mathbf{U}_{\mathbf{v}_F}} \leftarrow D_{\mathbf{U}_{\mathbf{v}_F} | \mathbf{V}_F = \mathbf{v}_F}$

1254 10: $M \leftarrow |D_{\mathbf{U}_{\mathbf{v}_F}}|$

1255 11: $\{\mathbf{u}_1, \dots, \mathbf{u}_M\} \leftarrow D_{\mathbf{U}_{\mathbf{v}_F} | \mathbf{U}}$

1256 12: $D_t \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_M\}$

1257 13: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1258 14: **if** $V = T$ **then**

1259 15: $\{v_1, \dots, v_N\} \leftarrow \{t, \dots, t\}$

1260 16: **else**

1261 17: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_t | \mathbf{pa}(V)$

1262 18: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_t | \mathbf{U}_V$

1263 19: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1264 20: **end if**

1265 21: $D_t \leftarrow D_t \cup \{v_1, \dots, v_N\}$

1266 22: **end for**

1267 23: $D_c \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_M\}$

1268 24: **for** $V \in \mathbf{V}$ following a causal order given by \mathcal{G} **do**

1269 25: **if** $V = T$ **then**

1270 26: $\{v_1, \dots, v_N\} \leftarrow \{c, \dots, c\}$

1271 27: **else**

1272 28: $\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\} \leftarrow D_c | \mathbf{pa}(V)$

1273 29: $\{u_{V_1}, \dots, u_{V_N}\} \leftarrow D_c | \mathbf{U}_V$

1274 30: $\{v_1, \dots, v_N\} \leftarrow f_V(\{\mathbf{pa}(V)_1, \dots, \mathbf{pa}(V)_N\}, \{u_{V_1}, \dots, u_{V_N}\})$

1275 31: **end if**

1276 32: $D_c \leftarrow D_c \cup \{v_1, \dots, v_N\}$

1277 33: **end for**

1278 34: $Q^* \leftarrow \text{avg}(D_t | Y) - \text{avg}(D_c | Y)$

1279 **Output:** Q^*

1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

G ANALYSIS MODULE’S METRICS

In order to analyze the characteristics of the sampled SCMs we implemented the following metrics. Let us imagine we sampled an SCM $\mathcal{M} := \{\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\}$ with $\mathbf{V} = (\mathbf{V}_o, \mathbf{V}_h)$ and whose causal graph is denoted \mathcal{G} . The projection of \mathcal{G} over the observable variables \mathbf{V}_o is denoted $\mathcal{G}_{\mathbf{V}_o}$.

Analysis of the causal graph \mathcal{G} :

- Average in-degree: $\bar{d}_{in} = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} |\mathbf{PA}(V)|$
- Variance of in-degree: $\text{var}(d_{in}) = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} (|\mathbf{PA}(V)| - \bar{d}_{in})^2$
- Average number of ancestors: $|\overline{An(V)}| = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} |An(V)|$ where $An(V)$ denotes the set of ancestors of V
- Variance of number of ancestors: $\text{var}(|An(V)|) = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} (|An(V)| - |\overline{An(V)}|)^2$
- Average number of descendants: $|\overline{De(V)}| = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} |De(V)|$ where $De(V)$ denotes the set of descendants of V
- Variance of number of descendants: $\text{var}(|De(V)|) = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} (|De(V)| - |\overline{De(V)}|)^2$
- Average length of causal paths: $\bar{L} = \frac{1}{|\mathbf{p}_{\mathcal{G}}|} \sum_{p \in \mathbf{p}_{\mathcal{G}}} |p|$ where $\mathbf{p}_{\mathcal{G}}$ denotes the set of directed paths in \mathcal{G}
- Variance length of causal paths: $\text{var}(L) = \frac{1}{|\mathbf{p}_{\mathcal{G}}|} \sum_{p \in \mathbf{p}_{\mathcal{G}}} (|p| - \bar{L})^2$
- Maximum length of causal paths: $L_{\max} = \max_{p \in \mathbf{p}_{\mathcal{G}}} |p|$

Analysis of the projected causal graph $\mathcal{G}_{\mathbf{V}_o}$:

- Average number of siblings³: $|\overline{Si(V)}| = \frac{1}{|\mathbf{V}_o|} \sum_{V \in \mathbf{V}_o} |Si(V)|$ where $Si(V)$ denotes the set of siblings of V
- Variance of number of siblings: $\text{var}(|Si(V)|) = \frac{1}{|\mathbf{V}_o|} \sum_{V \in \mathbf{V}_o} (|Si(V)| - |\overline{Si(V)}|)^2$
- Number of maximal confounded components (c-comps)⁴: $|\mathbf{C}|$ where \mathbf{C} denotes the set of maximal c-comps in $\mathcal{G}_{\mathbf{V}_o}$
- Average size of maximal c-comps: $|\overline{\mathbf{C}}| = \frac{1}{|\mathbf{C}|} \sum_{C \in \mathbf{C}} |C|$
- Variance of the size of maximal c-comps: $\text{var}(|\mathbf{C}|) = \frac{1}{|\mathbf{C}|} \sum_{C \in \mathbf{C}} (|C| - |\overline{\mathbf{C}}|)^2$

Analysis of the observational distribution $P_{\mathcal{M}}(\mathbf{V}_o)$:

- Minimum probability of the joint distribution: $p_{\mathbf{V}_o, \min} = \min_{\mathbf{v}_o \in \Omega_{\mathbf{V}_o}} P_{\mathcal{M}}(\mathbf{V}_o = \mathbf{v}_o)$
- Proportion of events with a null probability: $p_0 = \frac{1}{|\Omega_{\mathbf{V}_o}|} \sum_{\mathbf{v}_o \in \Omega_{\mathbf{V}_o}} \mathbf{1}_{P_{\mathcal{M}}(\mathbf{V}_o = \mathbf{v}_o) = 0}$ where $\mathbf{1}_\cdot$ denotes the indicator function
- Minimum probability of the marginal distributions:

$$p_{\min} = \min_{V \in \mathbf{V}_o} \min_{v \in \Omega_V} P_{\mathcal{M}}(V = v)$$

- Average minimum probability of the marginal distributions:

$$\bar{p}_{\min} = \frac{1}{|\mathbf{V}_o|} \sum_{V \in \mathbf{V}_o} \frac{1}{|\Omega_V|} \min_{v \in \Omega_V} P_{\mathcal{M}}(V = v)$$

³Two variables are considered siblings if they are linked by a bi-directed edge.

⁴We use (Tian & Pearl, 2002) definition of (maximal) confounded components.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

- Variance of the minimum probability of the marginal distributions:

$$\text{var}(p_{\min}) = \frac{1}{|\mathbf{V}_o|} \sum_{V \in \mathbf{V}_o} (\min_{v \in \Omega_V} P_{\mathcal{M}}(V = v) - \bar{p}_{\min})^2$$

- Distance (L_1) of the joint distributions to the uniform one:

$$d(P_{\mathcal{M}}; \mathcal{U}) = \sum_{\mathbf{v}_o \in \Omega_{\mathbf{V}_o}} |P_{\mathcal{M}}(\mathbf{V}_o = \mathbf{v}_o) - \frac{1}{|\Omega_{\mathbf{V}_o}|}|$$

- Average distance (L_1) of the marginal distributions to the uniform one:

$$\overline{d}(P_{\mathcal{M}}; \mathcal{U}) = \frac{1}{|\mathbf{V}_o|} \sum_{V \in \mathbf{V}_o} \sum_{v \in \Omega_V} |P_{\mathcal{M}}(V = v) - \frac{1}{|\Omega_V|}|$$

- Variance of the distance (L_1) of the marginal distributions to the uniform one:

$$\text{var}(d(P_{\mathcal{M}}; \mathcal{U})) = \frac{1}{|\mathbf{V}_o|} \sum_{V \in \mathbf{V}_o} \left(\sum_{v \in \Omega_V} |P_{\mathcal{M}}(V = v) - \frac{1}{|\Omega_V|}| - \overline{d}(P_{\mathcal{M}}; \mathcal{U}) \right)^2$$

- Entropy of the joint distribution: $H(P_{\mathcal{M}}(\mathbf{V}))$

All the above-mentioned probabilities are computed from a set of 1M samples drawn from the SCM \mathcal{M} .

Let us note that p_{\min} enables the user to check if the strong positivity assumption holds. If $p_{\mathbf{V}_o, \min} > 0$, then strong positivity is respected. In addition, if strong positivity does not hold, $p_{\mathbf{V}_o, \min}$ and p_0 indicate the extent to which the assumption is not met – the higher the metrics, the less the hypothesis is respected. On the other hand, p_{\min} indicates whether the weak positivity assumption holds. If $p_{\min} > 0$, then weak positivity is respected. Finally, $d(P_{\mathcal{M}}; \mathcal{U})$, $\overline{d}(P_{\mathcal{M}}; \mathcal{U})$ and $\text{var}(d(P_{\mathcal{M}}; \mathcal{U}))$ enables the user to assess to which extent the observational distribution is imbalanced.

Analysis of the causal mechanisms \mathcal{F} :

- Average Pearson’s correlation between the parent-child pairs⁵:

$$\bar{\rho}_P = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} \frac{1}{|\mathbf{PA}(V) \cup U_V|} \sum_{V_j \in \mathbf{PA}(V) \cup U_V} \rho_P(V, V_j)$$

- Variance of Pearson’s correlation between the parent-child pairs:

$$\text{var}(\rho_P) = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} \frac{1}{|\mathbf{PA}(V) \cup U_V|} \sum_{V_j \in \mathbf{PA}(V) \cup U_V} (\rho_P(V, V_j) - \bar{\rho}_P)^2$$

- Average Spearman’s correlation between the parent-child pairs³

$$\bar{\rho}_S = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} \frac{1}{|\mathbf{PA}(V) \cup U_V|} \sum_{V_j \in \mathbf{PA}(V) \cup U_V} \rho_S(V, V_j)$$

- Variance of Spearman’s correlation between the parent-child pairs:

$$\text{var}(\rho_S) = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} \frac{1}{|\mathbf{PA}(V) \cup U_V|} \sum_{V_j \in \mathbf{PA}(V) \cup U_V} (\rho_S(V, V_j) - \bar{\rho}_S)^2$$

- Average conditional entropy of a variable given its parents:

$$\bar{H} = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} H(V | \mathbf{PA}(V))$$

⁵ ρ_P and ρ_S respectively denote the Pearson’s and Spearman’s correlation

- Variance of conditional entropy of a variable given its parents:

$$\text{var}(\mathbf{H}) = \frac{1}{|\mathbf{V}|} \sum_{V \in \mathbf{V}} (\mathbf{H}(V|\mathbf{PA}(V)) - \bar{\mathbf{H}})^2$$

In order to be able to use person correlations, spearman correlations, and conditional entropy as indicators of degrees of linearity, monotonicity, and stochasticity of causal mechanisms, we do not derive these quantities from samples drawn from the entailed distribution. Instead, for each variable, we create a dataset resulting from the application of its causal mechanism to the cartesian product of the values taken by its endogenous and exogenous parents⁶. In other words, we analyze the mechanisms' images of their input space. This allows us to analyze each mechanism independently of the others.

Thus, $\bar{\rho}_P$ and $\text{var}(\rho_P)$ can be interpreted as the average degree of linearity of causal mechanisms and their variance. Furthermore, $\bar{\rho}_S$ and $\text{var}(\rho_S)$ can be interpreted as the average degree of monotonicity of causal mechanisms and their variance. Finally, $\bar{\mathbf{H}}$ and $\text{var}(\mathbf{H})$ can be interpreted as the average level of stochasticity of causal mechanisms and its variance.

H ANALYSIS OF THE EMPIRICAL DISTRIBUTION OF THE GENERATED SCMS

As we do not provide the user with an expression of the distribution of the sampled regional discrete SCMs, we need to investigate if some SCMs classes are over/underrepresented. This analysis is important to identify the potential biases CausalProfiler might create in order to take them into account when evaluating Causal ML methods. Indeed, as our goal is to provide a tool for rigorous empirical evaluation of causal methods, we need to be transparent on the limitations of our generator so that researchers and practitioners can interpret the results of their methods with full knowledge of the potential biases coming from CausalProfiler.

H.1 EXPERIMENT

To visualize the distribution of the SCMs generated, we analyze the distribution of the metrics of the analysis module characterizing the SCMs. For each SCM sampled, all the implemented metrics (see Appendix G) are computed.

The studied SCMs are sampled from the SoIs defined by the cartesian product of the following parameters:

- **Number of endogenous variables:** $\{3, 4, 5\}$
- **Expected edge probability:** $\{0.2, 0.4, 0.6, 0.8\}$
- **Proportion of unobserved endogenous variables:** $\{0, 0.1, 0.2, 0.3\}$
- **Number of noise regions:** $\{2, 5, 10, 20, 50\}$
- **Cardinality of endogenous variables:** $\{2, 3, 4, 7\}$
- **Distribution of exogenous variables:** set to $\mathcal{U}[0, 1]$

For each *SoI* 10 SCMs are sampled, making a total of 9600 SCMs studied. Let us mention that we sample more SCMs than for verification (Section 6.1 for two reasons. First, it enables us to have a better approximation of the SCMs distribution. Second, the computation of all the assumptions and characteristics metrics is, in fact, less computationally expensive than computing all the independence tests that were required for verification.

⁶For continuous SCMs, we first discretize the variables' domains of definition and then build the cartesian product.

H.2 RESULTS

The first conclusion, based on Figures 2 to 6, is that the generated SCMs do indeed belong to the specified SoIs and that their characteristics are consistent with the latter.

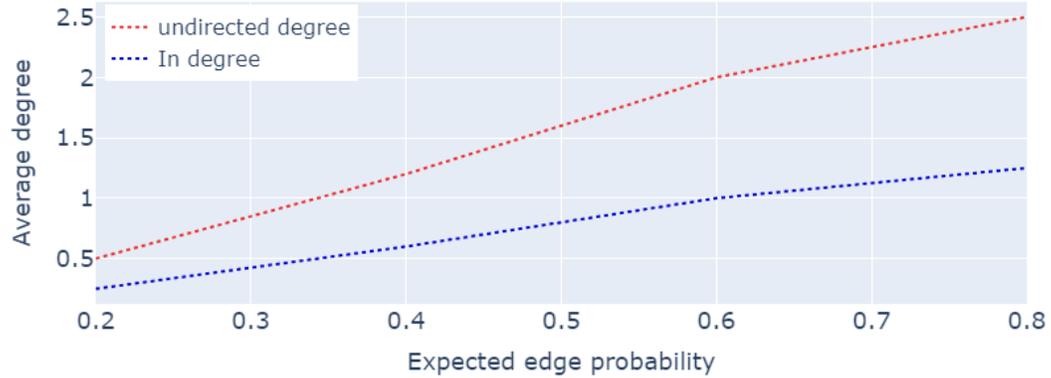


Figure 2: Average degree of the causal graphs for the generated SCMs depending on the expected edge probability. Observation: The average degree corresponds on average to the degree of the generated causal graphs.

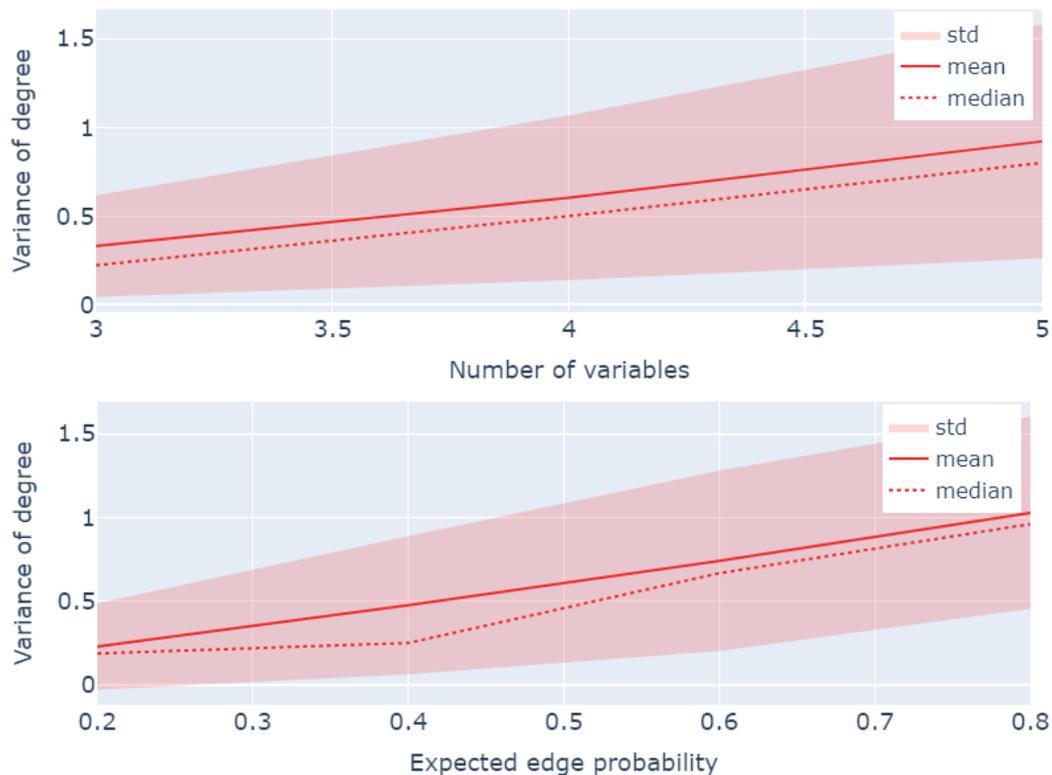


Figure 3: Variance of the causal graphs' degree of the generated SCMs depending on the number of variables and the expected edge probability. Observation: The variance of the degree increases with the size of the graph and its density.

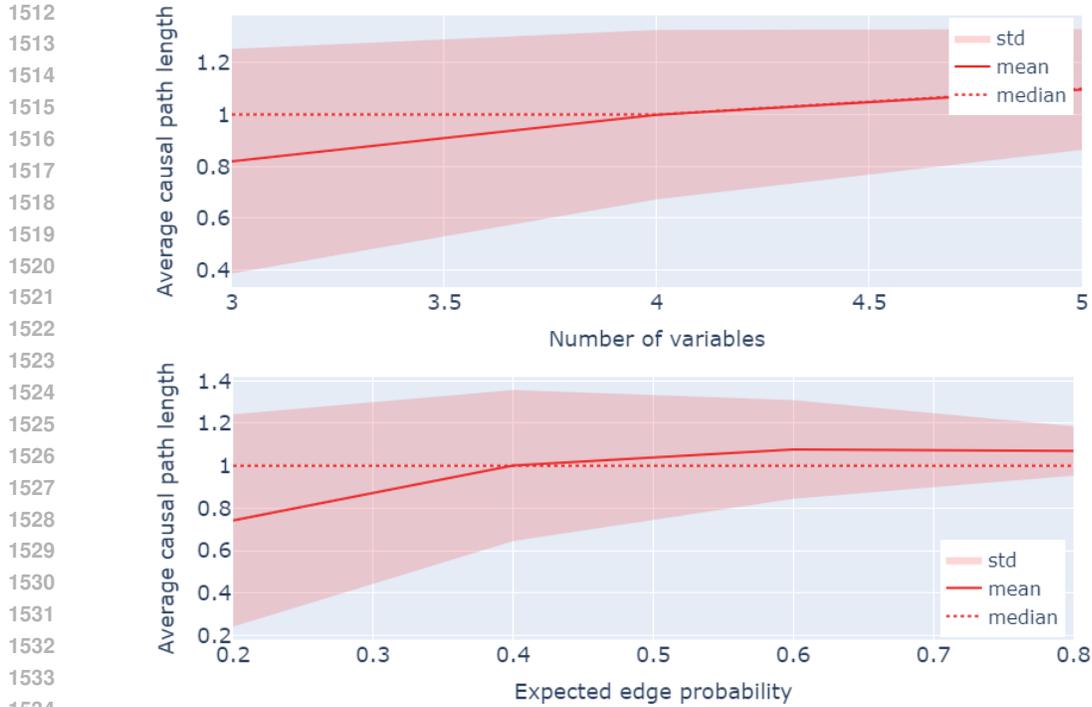


Figure 4: Average causal paths length of the causal graphs of the generated SCMs depending on the number of variables and the expected edge probability. Observation: The length of causal paths increases with the size of the causal graph and its density.

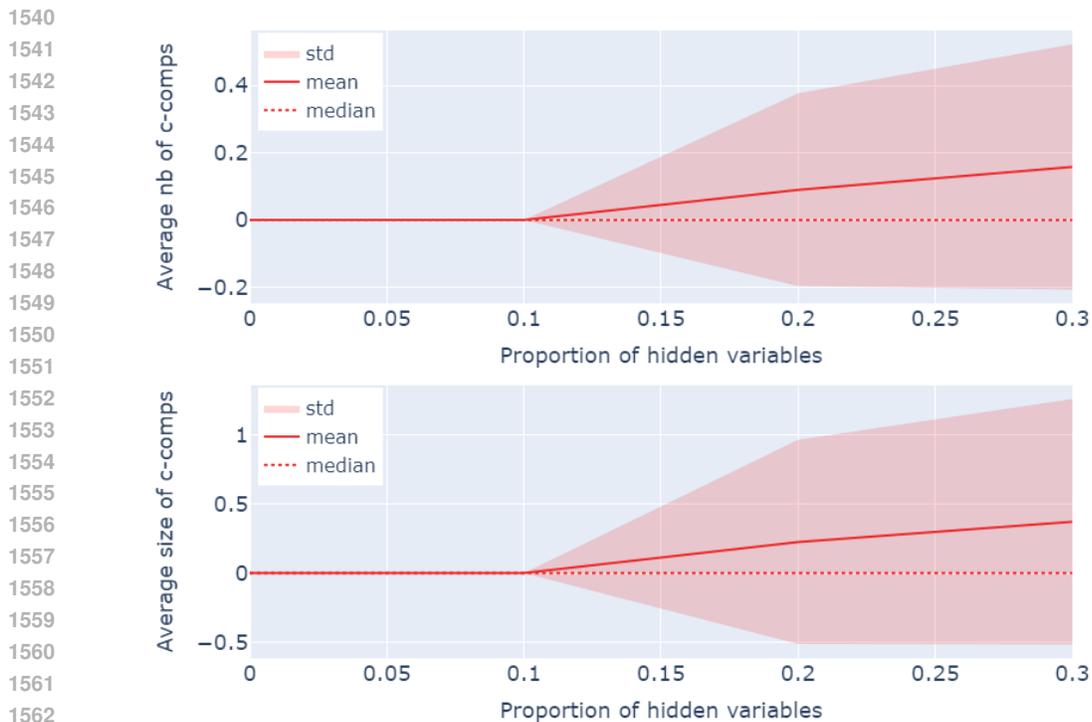


Figure 5: Average number and size of maximally confounded components in the projected causal graphs of the generated SCMs depending on the number of unobserved variables. Observation: The number and size of confounded components increase with the proportion of unobserved variables.

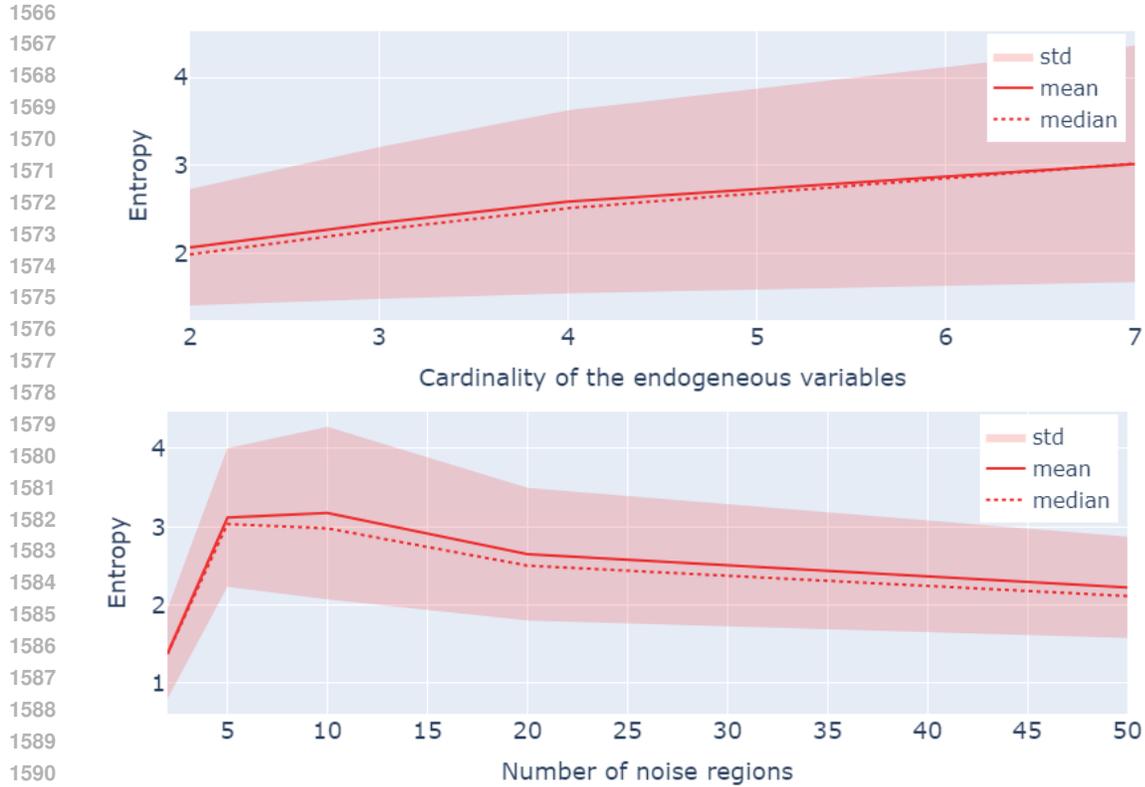


Figure 6: Average conditional entropy of a variable given its parents in the generated SCMs depending on the variables’ cardinality and the number of noise regions. Observation: The stochasticity of causal mechanisms increases with the cardinality of endogenous and exogenous variables.

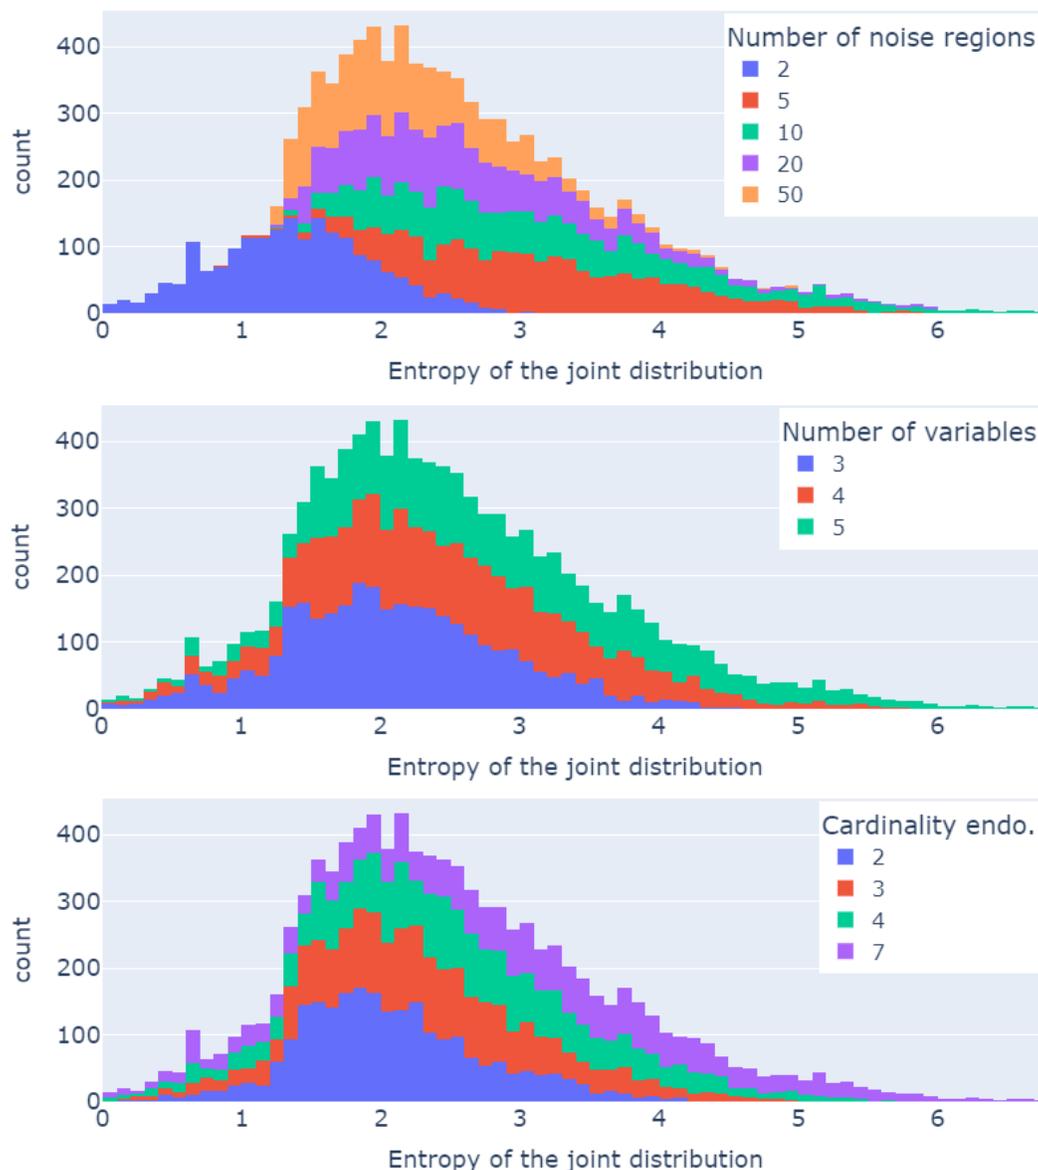
Table 4: Percentage of SCMs with confounded components depending on their proportion of unobserved endogenous variables.

Unobserved endo. variables (%)	Number of maximally confounded components		
	0	1	>1
0	100	0	0
10	100	0	0
20	90.9	9.1	0
30	83.3	16.7	0

Table 5: Percentage of SCMs with confounded components of different sizes depending on their proportion of unobserved endogenous variables. The size 1 of confounded components is not referenced, as if a confounded component is not empty, it is at least composed of two variables.

Unobserved endo. variables (%)	Avg. size of maximally confounded components				
	0	2	3	4	>4
0	100	0	0	0	0
10	100	0	0	0	0
20	90.9	5.7	2.3	1.0	0
30	83.3	10.8	4.9	1.0	0

1620 In addition, a number of findings about the distribution of the sampled SCMs can also be drawn. First,
 1621 the number and size of confounded components often equals zero (see also Tables 4 and 5). As highly
 1622 confounded SCMs are rare, we recommend that users sample SCMs with a large enough number of
 1623 variables and edge probability, if they want to consider graphs containing hidden confounders. For
 1624 instance, we recommend at least 10 variables with a 50% edge probability to have a large proportion
 1625 of graphs with at least one confounded component when setting the proportion of hidden endogenous
 1626 variables to 30%.



1666 Figure 7: Stacked histograms of the stochasticity level (measured through the entropy of the \mathcal{L}_1
 1667 joint distribution) of the sampled SCMs depending on the number of noise regions, the number of
 1668 variables, and their cardinality. Mean, standard deviation, and skewness of the distributions can be
 1669 found in Tables 6 to 8.

1670

1671

1672 Second, analyzing the stochasticity level (measured through the entropy of the \mathcal{L}_1 joint distribution,
 1673 see Appendix G) of the generated SCMs, one can see that the latter can be controlled in part by the
 parameters of the *SoI*. Indeed, increasing the number of endogenous variables and their cardinality

tends to increase the level of stochasticity, see Figure 7 and Tables 6 and 7. This behavior is expected as the discrete mechanisms are randomly sampled with an almost null probability of being deterministic (i.e., the probability of sampling a noise region with an empty support is almost null).

Table 6: Mean, standard deviation, and skewness of the distribution of stochasticity level (measured through the entropy of the \mathcal{L}_1 joint distribution) over the sampled SCMs depending on their number of endogenous variables. The distribution is displayed in Figure 7.

Number of endogenous variables	Entropy of the joint distribution		
	Mean	Std	Skewness
3	2.09	0.77	0.19
4	2.54	1.03	0.28
5	2.88	1.21	0.46

Table 7: Mean, standard deviation, and skewness of the distribution of stochasticity level (measured through the entropy of the \mathcal{L}_1 joint distribution) over the sampled SCMs depending on the cardinality of their endogenous variables. The distribution is displayed in Figure 7.

Cardinality	Entropy of the joint distribution		
	Mean	Std	Skewness
2	2.06	0.67	0.50
3	2.34	0.84	0.36
4	2.57	1.02	0.19
7	3.03	1.35	0.05

Table 8: Mean, standard deviation, and skewness of the distribution of stochasticity level (measured through the entropy of the \mathcal{L}_1 joint distribution) over the sampled SCMs depending on their number of noise regions. The distribution is displayed in Figure 7.

Number of noise regions	Entropy of the joint distribution		
	Mean	Std	Skewness
2	1.35	0.56	0.08
5	3.12	0.87	0.37
10	3.15	1.10	0.74
20	2.65	0.86	0.88
50	2.24	0.65	0.84

In addition, increasing the number of noise regions and the number of variables tends to increase the asymmetry of the distribution, see Figure 7 and Tables 6 and 8. This illustrates the fact that the number of degrees of freedom is increasing, and that it is therefore possible to generate increasingly stochastic mechanisms, although their probability of being sampled remains low. On the contrary, increasing the cardinality of the endogenous variables seems to reduce the asymmetry of the distribution, which may seem surprising. In reality, the distribution flattens out at higher stochasticity levels, making it more symmetrical. Indeed, both the mean and the standard deviation increase.

This analysis also reveals a surprising result: The number of noise regions does not seem to increase the level of stochasticity, cf. Figure 7 and Table 8. Theoretically, the more noise regions, the higher the number of mappings defining a causal mechanism. By complementing this mixture, we could expect to obtain a higher level of stochasticity. Further analysis is therefore required here to clarify the effect of the noise region parameter on stochasticity.

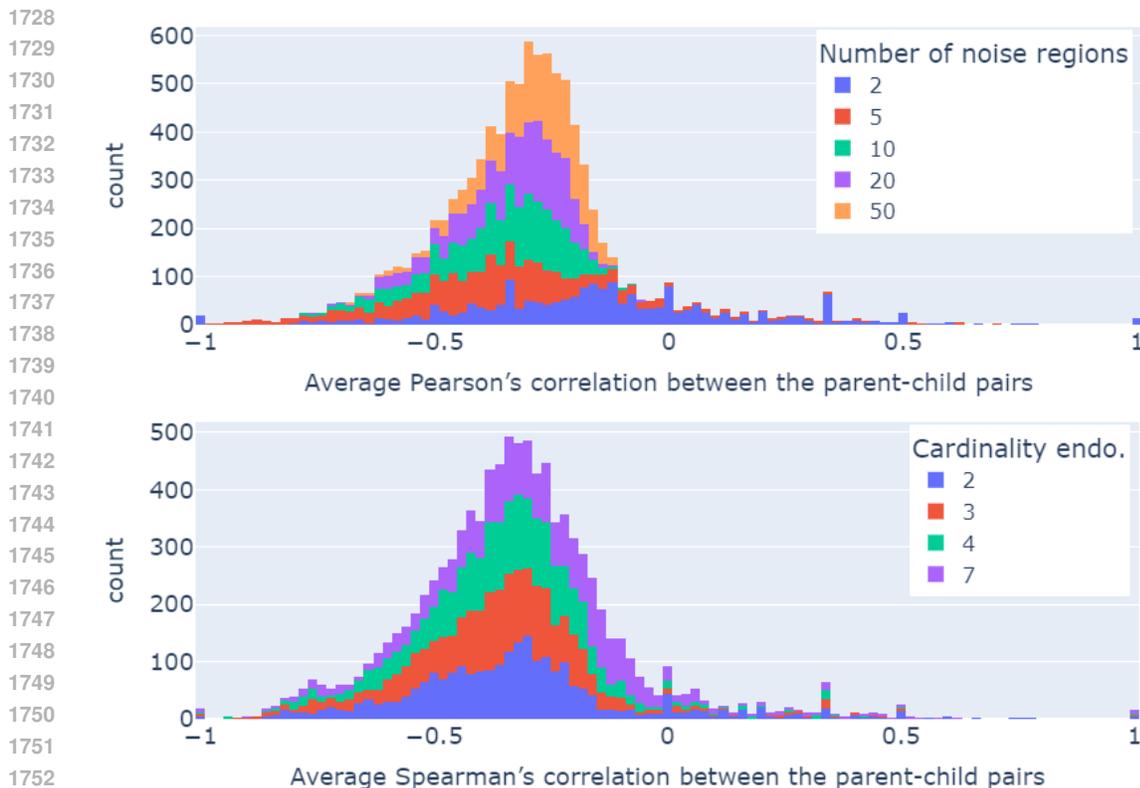


Figure 8: Stacked histograms of the average Pearson’s and Spearman’s correlation between the parent-child pairs of the generated SCMs. Mean, standard deviation, and skewness of the distributions can be found in Tables 9 and 10.

Table 9: Mean, standard deviation, and skewness of the distribution of the average Pearson’s correlation between the parent-child pairs of the generated SCMs depending on their number of noise regions. The distribution is displayed in Figure 8.

Number of noise regions	Pearson’s correlation		
	Mean	Std	Skewness
2	-0.15	0.30	0.54
5	-0.38	0.22	0.60
10	-0.36	0.13	-0.73
20	-0.34	0.11	-0.68
50	-0.30	0.10	-0.80

Table 10: Mean, standard deviation, and skewness of the distribution of the average Spearman’s correlation between the parent-child pairs of the generated SCMs depending on the cardinality of their endogenous variables. The distribution is displayed in Figure 8.

Cardinality	Pearson’s correlation		
	Mean	Std	Skewness
2	-0.32	0.25	1.10
3	-0.36	0.20	0.97
4	-0.35	0.19	1.09
7	-0.27	0.22	0.57

Third, the analysis of the levels of linearity and monotonicity (measured using Pearson and Spearman correlations) reveals that the sampled causal mechanisms are mostly neither linear nor monotonic, see Figure 8. Even if this result is to be expected, as the regional discrete mechanisms are discrete mappings without any notion of ordering, the fact that all the distributions are constituted of one peak on the negative side instead of two peaks, symmetric with respect to 0 is surprising. Hence, more investigation remains to be done to understand if our sampling algorithm tends to favor the generation of monotonically decreasing mechanisms.

One can also notice from Tables 9 and 10 that neither the cardinality of the endogenous variables nor the number of noise regions seems to affect the mean of the distributions, which is close to 0.35. In particular, the cardinality seems to have no effect on the distribution, while increasing the number of noise regions seems to increase the asymmetry of the distribution towards more linear mechanisms and decrease the standard deviation. Hence, we warn the users that choosing a high number of noise regions, hoping to be very diverse when generating mechanisms, might create the opposite effect over some metrics, as the distributions of Spearman’s and Pearson’s correlations seem to narrow down in this analysis.

Table 11: Percentage of SCMs respecting the strong positivity assumption depending on the number of endogenous variables.

Number of variables	Avg. min. proba. of the joint distribution	
	0	>0
3	88.6	11.4
4	95.6	4.4
5	97.0	3.0

Table 12: Percentage of SCMs respecting the strong positivity assumption depending on the cardinality of the endogenous variables.

Cardinality	Avg. min. proba. of the joint distribution	
	0	>0
2	95.7	4.3
3	97.5	2.5
4	95.5	4.5
7	89.8	10.2

Finally, Tables 11 and 12 illustrate that the assumption of strong positivity is rarely respected for all kinds of SCMs, whereas weak positivity is respected for all the sampled SCMs. More precisely, strong positivity hold on average in 6% of the generated datasets. This figure should be interpreted as a conservative lower bound. Indeed, our check uses finite samples, while strong positivity is defined in the infinite-sample regime: we reported a violation whenever any realization had an empirical frequency of 0 in 10,000 samples. In addition, there does not seem to be a correlation between the cardinality of the endogenous variables and the validation of the positivity assumption. It seems to mainly depend on the number of variables, which makes sense as the number of possible observations increases exponentially with the number of variables. Failure to respect the strong positivity assumption is a direct consequence of working with finite data, where infinitesimal probabilities are rounded to zero.

We therefore recommend that in order to evaluate Causal ML methods taking the strong positivity assumption, users use our analysis module to classify the sampled SCMs into two groups, depending on their compliance with the strong positivity assumption or not, and analyze them separately. This isolates the performance analysis within the theoretical validity framework of the method, and the analysis of its robustness to the violation of this assumption.

As a result, the generated SCMs belong mainly to the non-identifiable domain of Causal ML methods, as positivity is poorly respected. Users must, therefore, be careful in their interpretations when evaluating methods, as identifiable SCMs are much less represented than non-identifiable ones. We recommend starting the evaluation on small SoIs close to the identifiable domain, before progressively increasing the complexity of the causal datasets generated.

Let us highlight that this study was only carried out on regional discrete SCMs. We reserve for future work its extension to continuous SCMs.

H.3 COMPARISON TO CAUSALNF SYNTHETIC SCMs USED FOR EVALUATION

To illustrate the contribution in SCMs diversity that CausalProfiler can give to practitioners wishing to evaluate Causal ML methods, we compare the SCMs sampled in the previous Section with those used in the CausalNF work (Javaloy et al., 2023) for evaluation. We decided to first focus on the CausalNF synthetic SCMs because they have been reused by other papers (Sick & Dürr, 2025; Zhou et al., 2025) to evaluate new methods as if they were classical synthetic benchmarks for counterfactual evaluation.

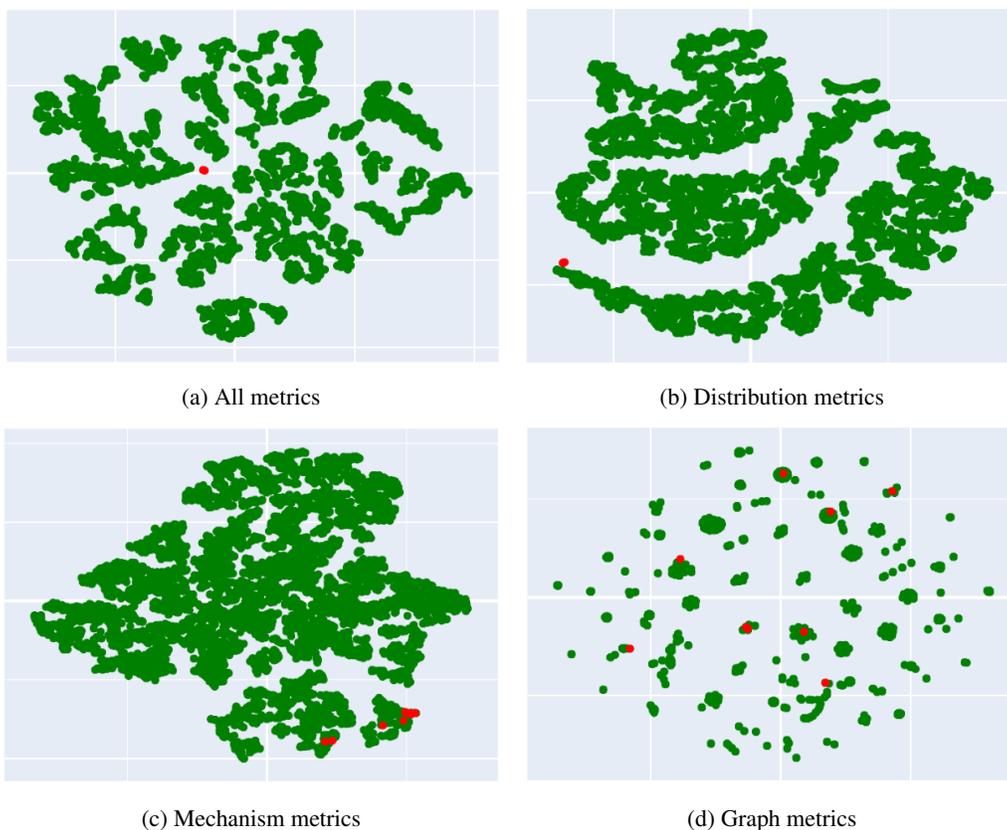


Figure 9: Two-dimensional t-SNE plots representing our sampled SCMs (green) and the synthetic SCMs used for evaluation of CausalNF (red). The latter, less numerous, have been plotted in the foreground to highlight their distribution in relation to our SCMs. The SCMs are described using characterization metrics from the analysis module. (a) t-SNE plot using all metrics (b) t-SNE plot using distribution metrics only (c) t-SNE plot using mechanism metrics only (d) t-SNE plot using graph metrics only.

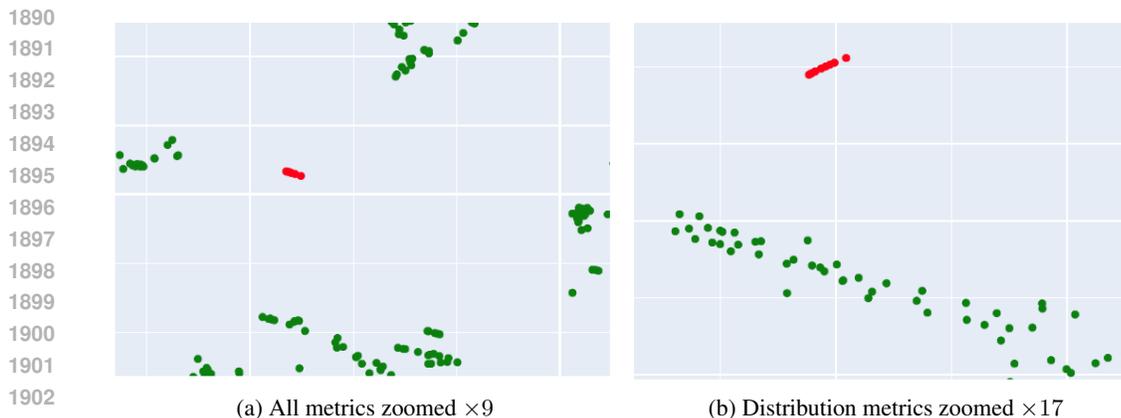


Figure 10: t-SNE plots of Figures 9a and 9b zoomed on CausalNF synthetic SCMs.

For this comparison, we reimplemented the synthetic SCMs of CausalNF using CausalProfiler, and applied all the metrics of the analysis module (cf. Appendix G). In this way, the CausalNF SCMs were processed in the same way as our SCMs. We then used these metrics to compare the two groups of SCMs. For the sake of having a fair comparison, not penalizing the fact that some assumptions were taken by the authors, we removed some metrics from the analysis: the hidden confounders and positivity metrics. Indeed, all CausalNF SCMs satisfy the causal sufficiency and strong positivity hypotheses, whereas, as presented in Appendix H.2, our SCMs do not by design. Finally, in order to obtain an easily interpretable visual result, we applied a two-dimensional t-SNE projection (Maaten & Hinton, 2008) to all these metrics and subgroups of metrics (Figure 9). Each t-SNE has been applied here with a perplexity of 30.

It can be seen that our SCMs are more diverse than those of CausalNF. Regarding graph metrics, it seems that CausalNF already has good diversity. The fact that we have greater support could mainly stem from the fact that we sampled a large number of SCMs. On the other hand, regarding distributions and mechanisms metrics, the increase in diversity is clear: The CausalNF SCMs are so similar compared to the total diversity that the dimension reduction projected them onto a confined space, cf. Figure 10.

As a result, we can conclude that CausalProfiler can enable practitioners to evaluate Causal ML methods on a more diverse set of SCMs and naturally derive more conclusions.

H.4 COMPARISON TO BNLEARN SEMI-SYTHETIC GRAPHICAL CAUSAL MODELS

This section also illustrates the contribution CausalProfiler makes to SCMs' diversity by comparing them with other causal models used in the literature: CANCER and EARTHQUAKE from bnlearn (Scutari, 2019). Unlike the synthetic and continuous SCMs from CausalNF, CANCER and EARTHQUAKE are discrete causal graph models. The following analysis, therefore, enriches the conclusions of the previous section.

CANCER and EARTHQUAKE were compared to the SCMs sampled by CausalProfiler in the same way as in the previous section: a two-dimensional t-SNE projection is applied to the metrics from the analysis module. The only difference here is that the mechanisms metrics cannot be computed on CANCER and EARTHQUAKE, as they are not proper SCMs, but graphical causal models. For the sake of having a fair comparison, we also excluded the hidden confounders metrics (as both bnlearn graphs are DAGs) but kept the positivity metrics for this analysis.

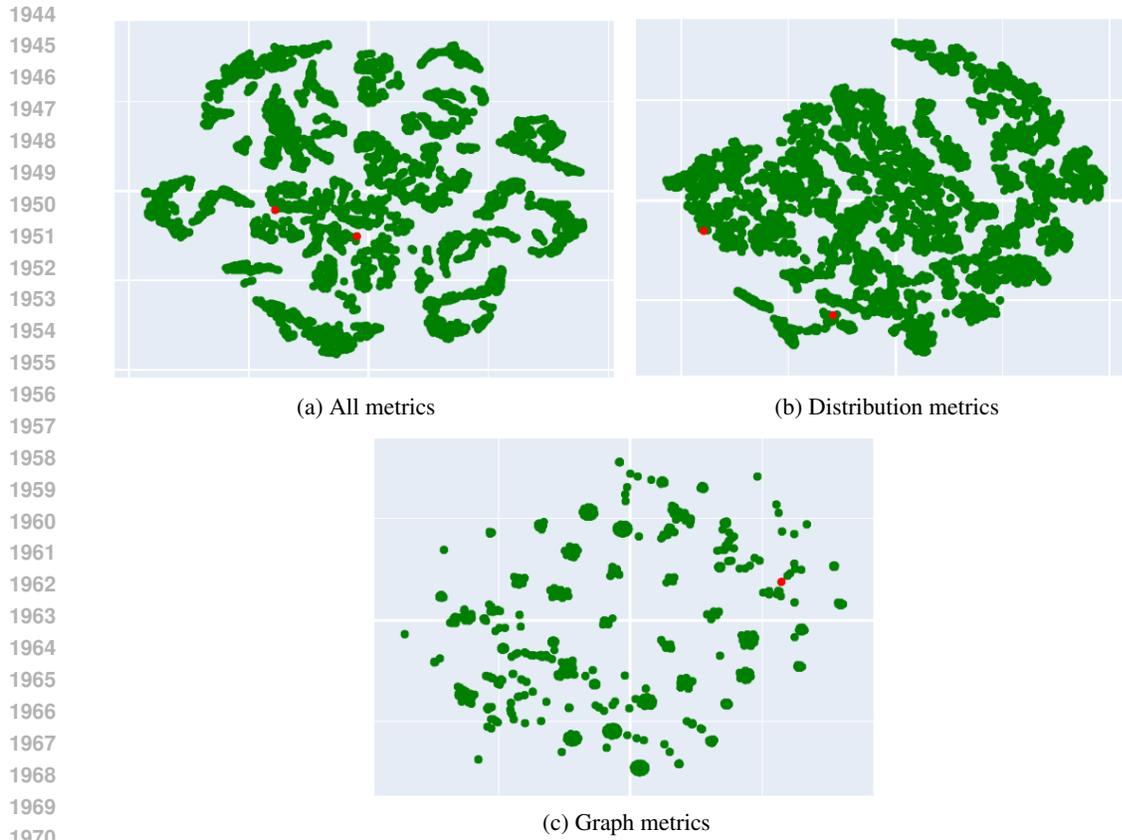


Figure 11: Two-dimensional t-SNE plots representing our sampled SCMs (green) and the semi-synthetic graphical causal models CANCER and EARTHQUAKE from bnlearn (red). The SCMs are described using characterization metrics from the analysis module. (a) t-SNE plot using all metrics (b) t-SNE plot using distribution metrics only (c) t-SNE plot using graph metrics only. Mechanism metrics cannot be used as bnlearn models do not model mechanisms but rather distributions.

The results, presented in Figure 11, show that the two bnlearn datasets are not confined to a small region of the two-dimensional space. Instead, they fall within the bottom left region of the t-SNE plot, overlapping with some of our generated SCMs. Hence, the conclusion of this analysis is similar to the previous one: CausalProfiler can generate SCMs producing similar causal datasets to existing ones while also generating more diverse sets of SCMs.

I VISUAL OVERVIEW OF CAUSALPROFILER'S SAMPLING STRATEGY

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

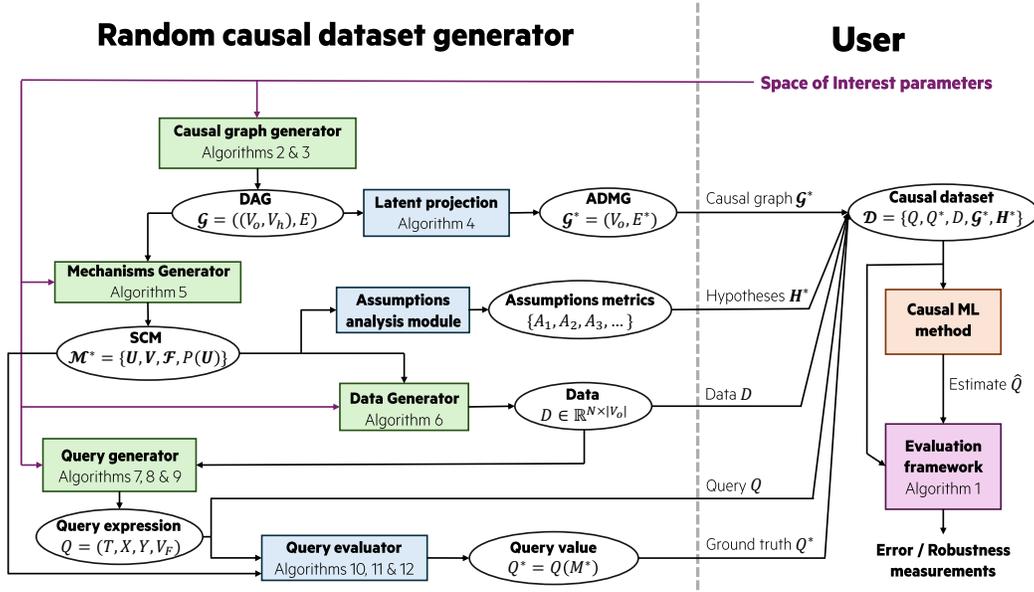


Figure 12: CausalProfiler structure. The left-hand side of the figure represents the code structure of the causal dataset generator. The right-hand side represents the user code. It illustrates how CausalProfiler can be used to evaluate a Causal ML method.

J PROOF OF PROPOSITION 5.1 (COVERAGE)

This section presents the proof of Proposition 5.1 stating that: For a Space of Interest $\mathcal{S} = \{\mathbb{M}, \mathbb{Q}, \mathbb{D}\}$, whose class of SCMs is a class of Regional Discrete SCMs with the maximum number of noise regions, any causal dataset $\mathcal{D} = \{Q, Q^*, D, \mathcal{G}^*, \mathbf{H}^*\}$ has a strictly positive probability to be generated.

Firstly, let us note that:

- Stating that any query Q can have any ground truth value Q^* given \mathcal{S} is equivalent to saying that the class of considered SCMs, i.e., the class of Regional Discrete SCMs with the maximum number of noise regions, is \mathcal{L}_3 -expressive with regards to the class of Markovian discrete SCMs (i.e., any \mathcal{L}_3 -distribution of the class of Markovian discrete SCMs can be expressed with a Regional Discrete SCM).
- As the set of hypotheses \mathbf{H}^* can contain at most \mathcal{L}_3 conditions, if the class of considered SCMs is \mathcal{L}_3 -expressive, then any set of hypotheses \mathbf{H}^* can be represented.
- If the class of considered SCMs is \mathcal{L}_3 -expressive, then it is also \mathcal{L}_1 -expressive, hence, D can be sampled from any distribution

As a result, our proof consists of showing that $P(Q, \mathcal{G}^* | \mathcal{S}) > 0$ and that the class of Regional Discrete SCMs with the maximum number of noise regions, denoted $\mathbb{M}_{\text{RD-SCM}, r=R_{\text{max}}}$, is \mathcal{L}_3 -expressive with regards to the class of Markovian discrete SCMs given an *SoI* \mathcal{S} and a causal graph \mathcal{G} .

Let us consider a *SoI* $\mathcal{S} = \{\mathbb{M}, \mathbb{Q}, \mathbb{D}\}$ with $\mathbb{M} \subseteq \mathbb{M}_{\text{RD-SCM}, r=R_{\text{max}}}$.

Proving $P(\mathcal{G}^* | \mathcal{S}) > 0$:

\mathcal{G}^* is built through Algorithm 2 as the latent projection of a DAG $\mathcal{G} = \{(\mathbf{V}_H, \mathbf{V}_O), E\}$ over \mathbf{V}_O where \mathcal{G} is sampled using Algorithm 1. As a result, following the steps of Algorithms 1 and 2:

$$\begin{aligned}
P(\mathcal{G}^*|\mathcal{S}) &= P(\{(\mathbf{V}_H, \mathbf{V}_O), E\}|\mathcal{S}) \\
&= P(E|\mathbf{V})P(\mathbf{V}_H, \mathbf{V}_O|\mathcal{S}) && \text{Edges are sampled independently of the} \\
& && \text{observability of the variables} \\
&= P(E|\mathbf{V})P(\mathbf{V}_H, \mathbf{V}_O|\mathbf{V})P(|\mathbf{V}|) && |\mathbf{V}| \text{ and } p_h \text{ are the only parameters influ-} \\
& && \text{encing the observability of the variables} \\
&= P(E|\mathbf{V})P(\mathbf{V}_H, \mathbf{V}_O|\mathbf{V})\frac{1}{N_{\max} - N_{\min}} && |\mathbf{V}| \sim \mathcal{U}[N_{\min}, N_{\max}] \\
&= P(E|\mathbf{V})\frac{|\mathbf{V}_H|!}{|\mathbf{V}|!}\frac{1}{N_{\max} - N_{\min}} && \mathbf{V}_H \subseteq \mathbf{V} \text{ sampled without replacement} \\
&= \frac{|\mathbf{V}_H|!}{|\mathbf{V}|!(N_{\max} - N_{\min})}P(E|\mathbf{V})
\end{aligned}$$

As $E = \{V_k \rightarrow V_i \mid V_k \in \mathbf{PA}(V_i), \forall V_i \in \mathbf{V}\}$ and the edges are sampled along the causal order $[1, N]$ with probability p_{edge} :

$$\begin{aligned}
P(\mathcal{G}^*|\mathcal{S}) &= \frac{|\mathbf{V}_H|!}{|\mathbf{V}|!(N_{\max} - N_{\min})} \prod_{i=1}^N P(\{V_k \rightarrow V_i \mid V_k \in \mathbf{PA}(V_i)\}) \\
&= \frac{|\mathbf{V}_H|!}{|\mathbf{V}|!(N_{\max} - N_{\min})} \prod_{i=1}^N p_{edge}^{|\mathbf{PA}(V_i)|} (1 - p_{edge})^{i-1-|\mathbf{PA}(V_i)|}
\end{aligned}$$

Let us note that $p_{edge} = 0 \implies |\mathbf{PA}(V_i)| = 0$ and $p_{edge} = 1 \implies |\mathbf{PA}(V_i)| = i - 1$. As a result, $P(\mathcal{G}^*|\mathcal{S}) > 0$.

Proving that $\mathbb{M}_{\text{RD-SCM}, r=R_{\max}}$ is \mathcal{L}_3 -expressive with regards to the class of Markovian discrete SCMs: Regional discrete SCMs are, by construction, Markovian Canonical SCMs (Zhang et al., 2022). Furthermore, if the number of noise regions is chosen to be large enough (typically set to its maximum value), any Markovian Canonical SCM can be represented using a Regional Discrete SCM⁷. Thus, applying Zhang et al. (2022) Theorem 2.4, we can assert that: for an arbitrary Markovian discrete SCM, there exists a Regional Discrete SCM such that they both have the same causal graph and the same \mathcal{L}_3 -distribution. Consequently, the class of Regional Discrete SCMs is \mathcal{L}_3 -expressive with respect to the class of Markovian discrete SCMs given the causal graph \mathcal{G} . Moreover, $P(\mathcal{G}) > 0$ for all \mathcal{G} because $\prod_{i=1}^N p_{edge}^{|\mathbf{PA}(V_i)|} (1 - p_{edge})^{i-1-|\mathbf{PA}(V_i)|} > 0$ (cf. previous paragraph). Thus, more generally, the class of Regional Discrete SCMs sampled by our CausalProfiler is \mathcal{L}_3 -expressive with respect to the class of Markovian SCMs.

Proving $P(Q|\mathcal{G}^*, \mathcal{S}) > 0$: Q is sampled given \mathbb{Q}, D and \mathcal{G}^* . Even though we currently only implement queries sampling for the classes $\mathcal{Q}_{\text{ATE}}, \mathcal{Q}_{\text{CATE}}$ and $\mathcal{Q}_{\text{Cf-TE}}$ (cf. Appendix F and Algorithms 6, 7 and 8), we can generalize our proof to any other query class (e.g., CDE, NDE). We simply assume that these classes translate the set of constraints on the variables under consideration (e.g., conditioning variables have to be distinct from treatment variables or any other graphical constraints that can be checked with \mathcal{G}^*) and express the probabilistic causal formula to be estimated. Once such a query class \mathbb{Q} is defined, our method randomly samples variables from \mathbf{V}_O in accordance with \mathbb{Q} constraints and by sampling realizations from D . We showed in the previous paragraph that $\mathbb{M}_{\text{RD-SCM}, r=R_{\max}}$ is \mathcal{L}_3 -expressive implying that it is \mathcal{L}_1 -expressive too. So, any realization can be present in D . As a result, for a given query class \mathbb{Q} , any Q can be generated. Hence, $P(Q|\mathcal{G}^*, \mathcal{S}) > 0$.

⁷The distinction between \mathbf{V}_O and \mathbf{V}_H is of no importance for \mathcal{L}_3 -expressiveness. \mathbf{V}_O and \mathbf{V}_H are only used to determine what will be visible to the user as benchmark.

2106 **Proving Proposition 5.1 by combining previous results:** We proved that $\mathbb{M}_{\text{RD-SCM}, r=R_{\max}}$
 2107 is \mathcal{L}_3 -expressive, hence any training set D , ground truth query Q^* and set of hypotheses
 2108 \mathbf{H}^* can be generated given an *SoI* \mathcal{S} , a causal graph \mathcal{G} and a causal query Q . In addition,
 2109 $P(Q, \mathcal{G}^* | \mathcal{S}) = P(Q | \mathcal{G}^*, \mathcal{S})P(\mathcal{G}^* | \mathcal{S})$ and we also prove that $P(Q | \mathcal{G}^*, \mathcal{S}) > 0$ and $P(\mathcal{G}^* | \mathcal{S}) > 0$.
 2110 Hence, $P(Q, \mathcal{G}^* | \mathcal{S}) > 0$. As a result, any causal dataset \mathcal{D} has a strictly positive probability to be
 2111 generated.

2112
 2113 **Remark on continuous SCMs.** The universal approximation theorem (Hornik, 1991) states that
 2114 NNs (with non-polynomial activation functions) are dense in the space of continuous functions,
 2115 meaning that any continuous function can be approximated by a sequence of NNs converging to this
 2116 function. However, this does not guarantee that they strictly cover the space of continuous functions.
 2117 In particular, whenever the number of layers and neurons is finite, one can always build a continuous
 2118 function too complex to be represented with this finite number of parameters. Hence, Proposition 5.1
 2119 cannot be extended to any class of continuous SCMs. However, it could potentially be adapted not to
 2120 ask for strict coverage but rather density. We leave this question for future work.

2121

2122

2123 K VERIFICATION RESULTS

2124

2125 We design and run verification experiments targeting each level of the PCH.

2126

2127 All following experiments are done on discrete SCMs to reduce approximations. Indeed, distributions
 2128 over continuous variables can only be approximated (e.g., using kernel methods) while discrete ones
 2129 can be computed exactly. In addition, the experiments rely on conditional independence testing,
 2130 which has been proven to be particularly difficult to use with continuous variables. Indeed, (Shah &
 2131 Peters, 2020) proved that no conditional independence test with a continuous conditioning variable
 2132 can have both a valid significance level and power.

2133

2134 K.1 \mathcal{L}_1 VERIFICATION

2135

2136 Consistency with \mathcal{L}_1 level of the PCH is tested through the verification that the Markov property
 2137 holds on randomly sampled regional discrete SCMs. Below is a description of the experimental
 2138 design choices made and the associated results.

2139

2140

2141 K.1.1 EXPERIMENT

2142

2143 For a given SCM $\mathcal{M} := \{\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\}$, we check that the Markov property is satisfied by
 2144 assessing whether there is a statistically significant amount of d-separations not leading to conditional
 2145 independence in the entailed distribution.

2146

2147 To do so, we first enumerate the list of sets of variables $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ in \mathbf{V} corresponding to d-separations
 2148 in \mathcal{M} 's causal graph $\mathcal{G}_{\mathcal{M}}$, ie $\mathbf{A} \perp\!\!\!\perp_{\mathcal{G}_{\mathcal{M}}} \mathbf{B} | \mathbf{C}$. Second, for each d-separated set $(\mathbf{A}, \mathbf{B}, \mathbf{C})$, we test
 2149 whether $\mathbf{A} \perp\!\!\!\perp_{P_{\mathcal{M}}} \mathbf{B} | \mathbf{C}$ by sampling 50k data points from the entailed distribution $P_{\mathcal{M}}$.

2150

2151 In practice, enumerating all the d-separations can be very costly. Moreover, as the set of variables
 2152 \mathbf{C} increases, it becomes increasingly complicated to robustly test the conditional independence
 2153 $\mathbf{A} \perp\!\!\!\perp_{P_{\mathcal{M}}} \mathbf{B} | \mathbf{C}$. Indeed, as the cardinality of \mathbf{C} increases, so does the number of combinations of
 2154 values for which to test independence between variables \mathbf{A} and \mathbf{B} . Running the statistical test becomes
 2155 costly, and the data volume required for robust independence test results increases exponentially. This
 2156 is why we limit ourselves to listing the d-separated sets (A, B, C) such that $A \in \mathbf{V}$, $B \in \mathbf{V} \setminus A$, and
 2157 $C \in \mathbf{V} \cup \mathbf{V}^2 \cup \mathbf{V}^3$ by enumerating all the possible (A, B, C) tuples, and testing whether they are
 2158 d-separated in $\mathcal{G}_{\mathcal{M}}$.

2159

2155 As the sampled SCMs are regional discrete, the conditional independence $A \perp\!\!\!\perp_{P_{\mathcal{M}}} B | C$ can be
 2156 tested with Pearson's χ^2 independence tests (Pearson, 1900). More precisely, A and B are considered
 2157 independent conditionally to C if for all values c of C , the H_0 hypothesis "A and B are independent"
 2158 is not rejected. Since Pearson's χ^2 test is based on the assumption that the number of samples is
 2159 large, we decide to skip tests where the Koehler criterion (Koehler & Larntz, 1980) is not met. Based
 on empirical analyses, this criterion indicates whether the χ^2 test is reliable depending on the number

of samples considered. In addition, as we conduct tests for each observed value c , we need to control for the expected proportion of false positives (represented by the Type I error of the test). To do so, we apply the Benjamini-Hochberg correction (Benjamini & Hochberg, 1995).

For each SoI , defined by the Cartesian product of the following parameters, we sample 5 SCMs:

- **Number of endogenous variables:** $\{4, 5, 6\}$
- **Expected edge probability:** $\{0.1, 0.4\}$
- **Proportion of unobserved endogenous variables:** set to 0 because the Markov property only hold for Markovian SCMs
- **Number of noise regions:** $\{5, 10\}$
- **Cardinality of endogenous variables:** $\{2, 3, 10\}$
- **Distribution of exogenous variables:** set to $\mathcal{U}[0, 1]$
- **Number of data points:** 50000

K.1.2 RESULTS

Table 13: Conditional independence tests based on χ^2 independence tests to assess compliance of sampled SCMs with the Markov property. Results are expressed as a percentage of the total of each test type for each conditioning set size. The number of tests is also shown in brackets.

Conditioning set size	$A \perp\!\!\!\perp_{P_{\mathcal{M}}} B C$ tests				χ^2 independence tests			
	Total	Pass	Fail	Skip	Total	Pass	Fail	Skip
$ C = 1$	100 (2 391)	91.76 (2 194)	4.94 (118)	3.3 (79)	100 (9 130)	85.4 (7 797)	1.43 (131)	13.17 (1 202)
$ C = 2$	100 (2 986)	91.16 (2 722)	5.63 (168)	3.22 (96)	100 (53 040)	45.2 (23 976)	0.33 (177)	54.46 (28 887)
$ C = 3$	100 (1 693)	91.08 (1 542)	5.67 (96)	3.25 (55)	100 (145 320)	18.49 (26 874)	0.07 (106)	81.43 (118 340)
TOTAL	100 (7 070)	91.34 (6 458)	5.40 (382)	3.25 (230)	100 (207 490)	28.26 (58 647)	0.2 (414)	71.54 (148 429)

The experimental results are summarized in Table 13, where it can be seen that 5.4% of the conditional independence tests failed. Despite the use of the Koehler criterion and Benjamini-Hochberg correction, some tests can still be rejected due to the random nature of finite data sampling, which can produce slight artificial correlations in the data. Moreover, on closer inspection, the majority of the failed tests (at least 350 out of 382)⁸ are unsuccessful because of a single failed χ^2 independence test. This reinforces our previous argument about the random nature of finite data sampling.

One can also notice that the number of skipped χ^2 independence tests increases with the size of the conditioning set. Such behavior is to be expected, since the number of realizations of the conditioning set increases exponentially with its cardinality, while the number of observations sampled to perform the independence tests remains constant. As a result, there are fewer and fewer observations available to perform each χ^2 test. In contrast, the number of fully skipped conditional independence tests remains constant. This means that the χ^2 skipped tests are relatively homogeneously distributed across all the conditional independence tests.

Someone might argue that the number of sampled observations should simply be automatically computed to verify the Koehler criterion. However, in general, such a calculation is complicated, if not impossible, to automate, as causal mechanisms are randomly sampled. As a result, all kinds of observational distributions can be induced with potentially very low probability realizations, for which the Koehler criterion could never be validated because the number of data to be sampled would be too large.

⁸Indeed, there is a total of 414 χ^2 tests that failed corresponding to 382 failed conditional independence tests. It mean that, at most 32(=414-382) conditional independence tests can have more than one failed χ^2 independence test.

To conclude, these results are sufficient to conclude that the Markov property is empirically verified by the sampled SCMs.

K.2 \mathcal{L}_2 VERIFICATION

Consistency with \mathcal{L}_2 level of the PCH is tested through the verification that the Do-calculus rules hold on randomly sampled regional discrete SCMs. Below is a description of the experimental design choices made (Appendix K.2.1) and the associated results (Appendix K.2.2).

K.2.1 EXPERIMENT

Definition K.1. Do-Calculus rules (Pearl, 2009)

Given an SCM $\mathcal{M} := \{\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\}$ whose causal graph \mathcal{G} is a DAG, and disjoint subsets $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, and \mathbf{W} of \mathbf{V} , the rules of the **Do-Calculus** are defined as follows:

1. **Insertion/deletion of observation:** if \mathbf{Y} and \mathbf{Z} are d-separated by $\mathbf{X} \cup \mathbf{W}$ in $\mathcal{G}_{\overline{\mathbf{X}}}$, then $P(\mathbf{Y}|\mathbf{do}(\mathbf{X} = \mathbf{x}), \mathbf{W}, \mathbf{Z}) = P(\mathbf{Y}|\mathbf{do}(\mathbf{X} = \mathbf{x}), \mathbf{W})$
2. **Action/observation exchange:** if \mathbf{Y} and \mathbf{Z} are d-separated by $\mathbf{X} \cup \mathbf{W}$ in $\mathcal{G}_{\overline{\mathbf{X}}, \underline{\mathbf{Z}}}$, then $P(\mathbf{Y}|\mathbf{do}(\mathbf{X} = \mathbf{x}), \mathbf{do}(\mathbf{Z} = \mathbf{z}), \mathbf{W}) = P(\mathbf{Y}|\mathbf{do}(\mathbf{X} = \mathbf{x}), \mathbf{Z}, \mathbf{W})$
3. **Insertion/deletion of action:** if \mathbf{Y} and \mathbf{Z} are d-separated by $\mathbf{X} \cup \mathbf{W}$ in $\mathcal{G}_{\overline{\mathbf{X}}, \overline{\mathbf{Z}(\mathbf{W})}}$, then $P(\mathbf{Y}|\mathbf{do}(\mathbf{X} = \mathbf{x}), \mathbf{do}(\mathbf{Z} = \mathbf{z}), \mathbf{W}) = P(\mathbf{Y}|\mathbf{do}(\mathbf{X} = \mathbf{x}), \mathbf{W})$

where $\mathcal{G}_{\overline{\mathbf{X}}}$ (resp. $\mathcal{G}_{\underline{\mathbf{X}}}$) represents the graph \mathcal{G} where the incoming edges in (resp. outgoing edges from) \mathbf{X} have been removed and $\mathbf{Z}(\mathbf{W})$ is the subset of nodes in \mathbf{Z} that are not ancestors of any node in \mathbf{W} in $\mathcal{G}_{\overline{\mathbf{X}}}$

For a given SCM, we check each rule by first enumerating the sets of d-separated variables of interest. Second, for each d-separated set, we test whether the distributions are statistically significantly similar by sampling 50k data points from the intervened SCMs and testing whether they are drawn from the same distribution.

For the same computational cost reasons as for \mathcal{L}_1 verification, we consider only univariate sets of variables X, Y, Z , and W . In addition, the studied SCMs are sampled from the same *SoIs* as defined in the \mathcal{L}_1 -verification experiment (Appendix K.1.1). Finally, to assess whether two conditional distributions are identical, we used Pearson’s χ^2 goodness of fit tests (Pearson, 1900). As done in Section K.1, we also use the Koehler criterion (Koehler & Larntz, 1980) and the Benjamini-Hochberg correction (Benjamini & Hochberg, 1995).

For each *SoI*, defined by the Cartesian product of the following parameters, we sample 2 SCMs:

- **Number of endogenous variables:** $\{4, 5, 6\}$
- **Expected edge probability:** $\{0.1, 0.4\}$
- **Proportion of unobserved endogenous variables:** set to 0 because the Markov property only hold for Markovian SCMs
- **Number of noise regions:** $\{5, 100\}$
- **Cardinality of endogenous variables:** $\{2, 5\}$
- **Distribution of exogenous variables:** set to $\mathcal{U}[0, 1]$
- **Number of data points:** 50000

Compared to the previous experiment (Appendix K.1.1), we reduce the number of sampled SCMs because comparing distributions two by two is more computationally expensive than conditional independence tests.

Table 14: Conditional independence tests based on χ^2 goodness of fit tests to assess compliance of sampled SCMs with the Do-Calculus rules. Results are expressed as a percentage of the total of each test type for each conditioning set size. The number of tests is also shown in brackets.

Do-Calculus Rule	Cond. goodness of fit				χ^2 goodness of fit			
	Total	Pass	Fail	Skip	Total	Pass	Fail	Skip
Rule 1 Insertion/deletion of observation	100 (3 378)	96.15 (3 248)	3.85 (130)	0 (0)	100 (171 092)	88.84 (152 004)	0.1 (172)	11.06 (18 916)
Rule 2 Action/observation exchange	100 (5 065)	94.04 (4 763)	5.96 (302)	0 (0)	100 (259 509)	83.84 (217 578)	0.09 (241)	16.06 (41 690)
Rule 3 Insertion/deletion of action	100 (5 169)	93.75 (4 846)	6.25 (323)	0 (0)	100 (282 184)	89.21 (251 731)	0.06 (157)	10.74 (30 296)
TOTAL	100 (13 612)	94.45 (12 857)	5.55 (755)	0 (0)	100 (712 785)	87.17 (621 313)	0.08 (570)	12.75 (90 902)

K.2.2 RESULTS

The experimental results are summarized in Table 14 where it can be seen that they are very similar to the \mathcal{L}_1 verification ones: roughly 6% of the conditional goodness of fit tests were not validated, some tests are rejected due to the random nature of finite data sampling but the majority them (at least 570 out of 755) are unsuccessful because of a single failed χ^2 goodness of fit test.

One can also notice that the percentage of skipped χ^2 goodness of fit tests is similar for rules 1 and 3 but increases by roughly 50% for rule 2. Such behavior is to be expected as rule 2 is the only rule to have conditioning sets of size 3 on both sides of the equality. However, the number of skipped tests remains low, with a maximum of 16%.

As a result, we estimate that these results are sufficient to conclude that the Do-calculus rules are respected by the sampled SCMs.

K.3 \mathcal{L}_3 VERIFICATION

Consistency with \mathcal{L}_3 level of the PCH is tested through the verification that the axiomatic characterization of structural counterfactuals holds on randomly sampled regional discrete SCMs. Below is a description of the experimental design choices made (Appendix K.3.1) and the associated results (Appendix K.3.2).

Definition K.2. Axiomatic characterization of structural counterfactuals (Pearl, 2009)
Given an SCM $\mathcal{M} := \{\mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\}$ whose causal graph \mathcal{G} is a DAG, the **axioms of structural counterfactuals** are defined as follows:

- Composition:** For any sets of endogenous variables \mathbf{X} , \mathbf{Y} , and \mathbf{W} in \mathbf{V} and any realization \mathbf{u} of \mathbf{U} , if $\mathbf{W}_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u}) = \mathbf{w}$ then $\mathbf{Y}_{do(\mathbf{X}=\mathbf{x}), do(\mathbf{W}=\mathbf{w})}(\mathbf{u}) = \mathbf{Y}_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u})$
- Effectiveness:** For any disjoint sets of endogenous variables \mathbf{X} , and \mathbf{W} in \mathbf{V} and any realization \mathbf{u} of \mathbf{U} , $\mathbf{X}_{do(\mathbf{X}=\mathbf{x}), do(\mathbf{W}=\mathbf{w})}(\mathbf{u}) = \mathbf{x}$
- Reversibility:** For any two distinct variables Y and W and any sets of other variables \mathbf{X} in \mathbf{V} and any realization \mathbf{u} of \mathbf{U} , if $Y_{do(\mathbf{X}=\mathbf{x}), do(\mathbf{W}=\mathbf{w})}(\mathbf{u}) = y$ and $W_{do(\mathbf{X}=\mathbf{x}), do(\mathbf{Y}=\mathbf{y})}(\mathbf{u}) = w$ then $Y_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u}) = y$

Note that we do not write $P(\mathbf{W}_{do(\mathbf{X}=\mathbf{x})}|\mathbf{U})$ but rather $\mathbf{W}_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u})$ as it is a deterministic expression. Indeed, if \mathbf{U} is fixed, there is no stochastically anymore, so we no longer need to reason in distributions but rather in functional forms.

K.3.1 EXPERIMENT

For a given SCM, using Definition K.1 notations, we check that:

1. The **Composition** axiom is satisfied by assessing whether $\mathbf{W}_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u}) = \mathbf{w}$ implies $\mathbf{Y}_{do(\mathbf{X}=\mathbf{x}),do(\mathbf{W}=\mathbf{w})}(\mathbf{u}) = \mathbf{Y}_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u})$ for any sets of endogenous variables \mathbf{X} , \mathbf{Y} , and \mathbf{W} in \mathbf{V} and any realization \mathbf{u} of \mathbf{U}
2. The **Effectiveness** axiom is satisfied by assessing whether $\mathbf{X}_{do(\mathbf{X}=\mathbf{x}),do(\mathbf{W}=\mathbf{w})}(\mathbf{u}) = \mathbf{x}$ for any sets of endogenous variables \mathbf{X} , and \mathbf{W} in \mathbf{V} and any realization \mathbf{u} of \mathbf{U}
3. The **Reversibility** axiom is satisfied by assessing whether $Y_{do(\mathbf{X}=\mathbf{x}),do(\mathbf{W}=\mathbf{w})}(\mathbf{u}) = y$ and $W_{do(\mathbf{X}=\mathbf{x}),do(\mathbf{Y}=\mathbf{y})}(\mathbf{u}) = w$ implies $Y_{do(\mathbf{X}=\mathbf{x})}(\mathbf{u}) = y$ for any two (distinct) variables Y and W and any sets of variables \mathbf{X} in \mathbf{V} and any realization \mathbf{u} of \mathbf{U}

For each *SoI*, defined by the Cartesian product of the following parameters, we sample 5 SCMs:

- **Number of endogenous variables:** $\{3, 5, 10\}$
- **Expected edge probability:** $\{0.1, 0.5, 0.7\}$
- **Proportion of unobserved endogenous variables:** set to 0 because the Markov property only hold for Markovian SCMs
- **Number of noise regions:** $\{3, 5, 10\}$
- **Cardinality of endogenous variables:** $\{2, 5, 7\}$
- **Distribution of exogenous variables:** set to $\mathcal{U}[0, 1]$
- **Number of data points:** 50000

For each SCM, instead of enumerating all the possible four sets of variables \mathbf{X} , \mathbf{Y} and \mathbf{W} , we sample a partition of three elements of a randomly sampled subset of \mathbf{V} of a size randomly picked in $[3, |\mathbf{V}|]$. This sampling strategy enables us to make sure the three sets are disjoint and of randomly varying size. In addition, for each four sets, we sample 50k realizations of \mathbf{U} .

Let us note that the axioms now correspond to exact realizations and not equal probabilities. As a result, we expect no failure as no approximation is made in this experiment.

K.3.2 RESULTS

As expected, all the tested equalities are verified in our experiments. We can, therefore, consider that the SCMs created by our generator allows the estimation of any structural counterfactual queries.

L EXTENDED EXPERIMENTAL RESULTS

This appendix complements Section 6 with extended setup details and results. We first provide further details for Experiments 1 and 2, including the Algorithm 12 describing our evaluation protocol. We then include an additional experiment on ATE estimation under hidden confounding, and then an evaluation of runtime scalability on larger graphs.

L.1 EXPERIMENT 1: ADDITIONAL INFORMATION

Table 15 details the *SoI* used in our experiments, Table 16 reports extended performance metrics complementing Table 1, and Figure 13 shows box plots of ATE estimation errors.

Parameters not explicitly listed for a given *SoI* are set to their default values as per the benchmark configuration. Neural Networks for our experiments have two 8-neuron layers and use ReLU activation. Unless otherwise specified, we use 1000 samples per SCM in our experiments. This value was chosen as a stable default for these *SoIs* after testing several dataset sizes. More precisely, after testing the stability of the methods we evaluate (i.e., CausalNF, DCM, NCM, VACA) over the following dataset sizes, 50, 100, 200, 1000, and 2000, we found that 1000 samples was the smallest dataset size not drastically degrading the performance of the methods. This is why we decided to take this value as default for our experiments. We only vary it explicitly when studying the effect of limited data (e.g., in NN-Large-LowData).

2376 **Algorithm 12** Evaluation process for causal machine learning methods

2377 1: **Input:** List of Spaces of Interest *SoIs*, list of seeds *seeds* number of examples per SCM

2378 *num_examples*

2379 2: **Initialize:** *method* \leftarrow CausalMLMethod()

2380 3: **for** each *SoI* in *SoIs* **do**

2381 4: **for** each *seed* in *seeds* **do**

2382 5: setGlobalSeed(*seed*)

2383 6: **for** each *examples* in *num_examples* **do**

2384 7: Generate samples, queries, and targets from the profiler

2385 8: Get estimates using the *method* on the generated samples and queries

2386 9: Calculate (and store) error by comparing estimates with targets

2387 10: **end for**

2388 11: Compute performance statistics for seed

2389 12: **end for**

2390 13: Compute performance statistics for *SoI*

2391 14: **end for**

2392 15: **Output:** Final summary with evaluation results

2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410

2411 Table 15: Specification of each *SoI* used in the general experiments. N denotes the sampled number
2412 of nodes.

2413		2414		2415	
2414		Linear-Medium	2415		NN-Medium
2415	# Nodes	15-20	2415	# Nodes	15-20
2416	Mechanism	Linear	2416	Mechanism	NN
2417	Expected Edges	$2 \times N$	2417	Expected Edges	$2 \times N$
2418	Variable Type	Continuous	2418	Variable Type	Continuous
2419	Samples	1000	2419	Samples	1000
2420	Query Type	ATE	2420	Query Type	ATE
2421	Seeds	[10, 11, 12, 13, 14]	2421	Seeds	[10, 11, 12, 13, 14]
2422		2423		2424	
2423		NN-Large	2424		NN-Large-LowData
2424	# Nodes	20-25	2424	# Nodes	20-25
2425	Mechanism	NN	2425	Mechanism	NN
2426	Expected Edges	$2 \times N$	2426	Expected Edges	$2 \times N$
2427	Variable Type	Continuous	2427	Variable Type	Continuous
2428	Samples	1000	2428	Samples	50
2429	Query Type	ATE	2429	Query Type	ATE
	Seeds	[10, 11, 12, 13, 14]		Seeds	[10, 11, 12, 13, 14]

2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483

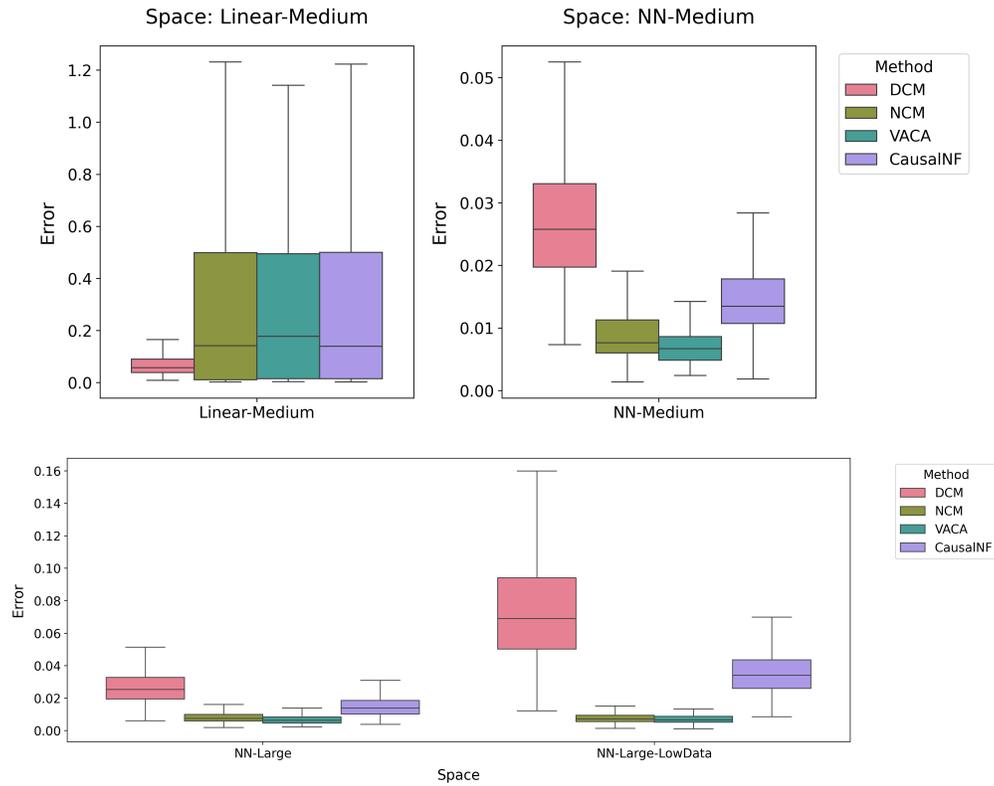


Figure 13: Box plots showing ATE estimation errors across different *Sols*

Table 16: Additional performance metrics of CausalNF, DCM, NCM, and VACA on the general experiments.

Space	Method	Min Error	Total Fail	Runtime Mean	Runtime Std
Linear-Medium	CausalNF	0.0024	0	27.58 s	18.33 s
	DCM	0.0086	0	33.08 s	9.71 s
	NCM	0.0024	0	14.77 s	1.42 s
	VACA	0.0038	1335	11.69 s	4.54 s
NN-Medium	CausalNF	0.0019	0	21.47 s	19.52 s
	DCM	0.0073	0	31.79 s	10.62 s
	NCM	0.0014	0	14.65 s	1.43 s
	VACA	0.0024	125	12.13 s	4.41 s
NN-Large	CausalNF	0.0038	0	30.23 s	25.33 s
	DCM	0.0060	0	38.33 s	14.02 s
	NCM	0.0018	0	18.90 s	1.38 s
	VACA	0.0023	290	12.88 s	4.31 s
NN-Large-LowData	CausalNF	0.0086	0	44.28 s	17.10 s
	DCM	0.0121	0	4.82 s	1.34 s
	NCM	0.0013	0	0.81 s	0.11 s
	VACA	0.0010	0	10.43 s	4.59 s

L.2 EXPERIMENT 2: ADDITIONAL INFORMATION

We provide more details about the *SoI* used in our experiments in Table 17 and present extended performance metrics in Table 18, complementing those already shown in Table 2. Parameters not explicitly listed for a given *SoI* are set to their default values as per the benchmark configuration.

Table 17: Specification of the Spaces of Interest used for evaluating discrete SCMs with Ctf-TE queries. N denotes the sampled number of nodes.

Name	Disc-C2-Reject	Name	Disc-C4-Unbias
# Nodes	10–15	# Nodes	10–15
# Categories	2	# Categories	4
Mechanism	Tabular	Mechanism	Tabular
Sampling Strategy	Rejection	Sampling Strategy	Random
Edges	N	Edges	N
Samples	500	Samples	500
Query Type	Ctf-TE	Query Type	Ctf-TE
Seeds	[1, 2, 3, 4, 5]	Seeds	[1, 2, 3, 4, 5]

Name	Disc-L-C2-Unbias
# Nodes	20–30
# Categories	2
Mechanism	Tabular
Sampling Strategy	Random
Edges	N
Samples	500
Query Type	Ctf-TE
Seeds	[1, 2, 3, 4, 5]

Table 18: Additional performance metrics of CausalNF and DCM on the discrete experiments.

Space	Method	Min Error	Total Fail	Runtime Mean	Runtime Std
Disc-C2-Reject	CausalNF	0.0000	202	0.46 s	0.04 s
	DCM	0.0000	107	8.81 s	3.55 s
Disc-C4-Unbias	CausalNF	0.0000	1017	0.42 s	0.03 s
	DCM	0.0000	565	7.68 s	3.43 s
Disc-L-C2-Unbias	CausalNF	NaN	2500	0 s	0 s
	DCM	0.0000	283	16.39 s	6.42 s

L.3 EXPERIMENT 3: ATE ESTIMATION UNDER HIDDEN CONFOUNDING

In this experiment, we demonstrate how our framework can be used to evaluate methods in the presence of latent confounders — a common challenge in real-world causal inference. A key goal here is not only to confirm theoretical limitations but to investigate how quickly and severely performance degrades when assumptions are violated. While theory can tell us whether identification holds, it is often agnostic to the *degree* of failure. See Table 20 for a summary of results, Table 21 for a few additional performance metrics, and Figure 14 for a boxplot of ATE estimation errors over the different *SoI*.

We focus on two linear SCM settings:

- **Linear-No-Hidden:** Linear SCMs with 10-15 nodes and full observability (no hidden confounders), using 1000 data points per SCM.
- **Linear-60-Hidden:** Same setup as above, but with 60% of the variables unobserved (hidden).

2538 We provide more details about the *SoI* used in our experiments in Table 19. Parameters not explicitly
 2539 listed for a given *SoI* are set to their default values as per the benchmark configuration.
 2540

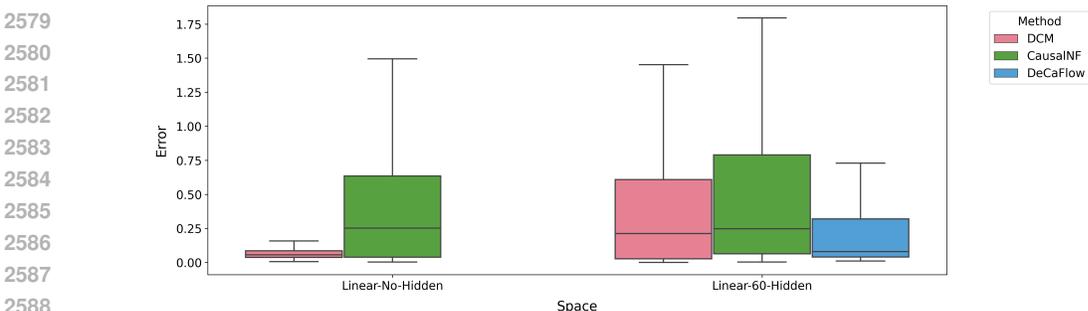
2541
 2542 Table 19: Specification of the *SoIs* used to evaluate performance under hidden confounding. N
 2543 denotes the sampled number of nodes.
 2544

Linear-No-Hidden		Linear-60-Hidden	
Name		Name	
# Nodes	10-15	# Nodes	10-15
Mechanism	Linear	Mechanism	Linear
Expected Edges	$2 \times N$	Expected Edges	$2 \times N$
Variable Type	Continuous	Variable Type	Continuous
Prop. Hidden Nodes	0%	Prop. Hidden Nodes	60%
Samples	1000	Samples	1000
Query Type	ATE	Query Type	ATE
Seeds	[42, 43, 44, 45, 46]	Seeds	[42, 43, 44, 45, 46]

2553
 2554
 2555 **Setup.** We evaluate three methods: CausalNF (Javaloy et al., 2023), DCM (Chao et al., 2023), and
 2556 DeCaFlow (Almodóvar et al., 2025). The first two methods assume causal sufficiency, and therefore
 2557 cannot, in theory, handle hidden confounding. DeCaFlow, in contrast, is explicitly designed for this
 2558 setting but requires access to the full causal graph (including hidden variables) and does not run when
 2559 all variables are observed. Thus, we include it only in the hidden confounding *SoI*.
 2560

2561
 2562 **Results (Linear-No-Hidden).** As expected, both CausalNF and DCM perform well when all
 2563 variables are observed. DCM achieves lower mean error (0.0845) and standard deviation (0.1515),
 2564 with a maximum error of 2.89. The upper whisker of DCM’s box plot lies below the median of
 2565 CausalNF, indicating consistent superior performance. These results serve as a reference point for
 2566 comparison when introducing hidden variables.
 2567

2568
 2569 **Results (Linear-60-Hidden).** With 60% of variables hidden, method performance degrades signif-
 2570 icantly. DeCaFlow performs reliably, with an error mean of 0.3405 and low variance. In contrast,
 2571 CausalNF—despite a box plot that visually appears well-behaved—has a massive error mean of
 2572 2.67×10^{12} and a maximum error exceeding 10^{15} . This is due to a small subset of SCMs producing
 2573 extremely large errors (14 with error > 1000), illustrating that, when assumptions are violated, error
 2574 can become arbitrarily large. While DCM does not show such instability on this particular sample, its
 2575 theoretical limitations under hidden confounding still hold — the expectation is that if we evaluate
 2576 over enough SCMs we will eventually also get arbitrarily large errors due to the violation of the
 2577 causal sufficiency assumption.
 2578



2579
 2580
 2581
 2582
 2583
 2584
 2585
 2586
 2587
 2588
 2589
 2590 Figure 14: Box plots of ATE estimation errors in the presence and absence of hidden confounding.
 2591 Each box shows the interquartile range and median, with whiskers extending to $1.5 \times$ IQR. CausalNF
 and DCM are shown for both *Sols*; DeCaFlow is shown only for the hidden setting.

2592 Table 20: Performance summary of CausalNF, DCM, and DeCaFlow on the hidden confounder
 2593 experiments.

2594

Space	Method	Mean Error	Std Error	Max Error	Runtime (s)
Linear-No-Hidden	CausalNF	0.5538	0.9866	14.2495	8570.0
	DCM	0.0845	0.1515	2.8954	12144.6
Linear-60-Hidden	CausalNF	2.667e+12	5.497e+13	1.225e+15	293.2
	DCM	0.5584	1.2122	17.2049	4187.6
	DeCaFlow	0.3405	0.6799	5.9435	2264.0

2601

2602

2603 Table 21: Additional performance metrics of CausalNF, DCM, and DeCaFlow on the hidden con-
 2604 founder experiments.

2605

Space	Method	Min Error	Total Fail	Runtime Mean	Runtime Std
Linear-No-Hidden	CausalNF	0.0036	0	17.14 s	10.61 s
	DCM	0.0068	0	24.29 s	7.64 s
Linear-60-Hidden	CausalNF	0.0029	0	0.59 s	0.02 s
	DCM	0.0000	0	8.38 s	3.45 s
	DeCaFlow	0.0108	0	4.53 s	1.27 s

2612

2613

2614 L.4 RUNTIME SCALABILITY ON LARGER GRAPHS

2615

2616 We additionally evaluate the scalability of CausalProfiler with respect to the number of variables.
 2617 Batch processing and vectorized operations enable efficient dataset generation even for graphs with
 2618 hundreds of variables. Table 22 reports the average generation time (over 5 runs) for producing
 2619 10,000 samples and 50 queries (each estimated using 10,000 additional datapoints), using the same
 2620 CPU hardware described in Section 6.3.

2621 Table 22: Average runtime (seconds) of CausalProfiler for generating datasets across increasing
 2622 numbers of variables. Each value is the mean over 5 runs with standard deviation in parentheses.

2623

Num Variables	Mean Time (s)	Std Dev (s)
10	0.19	0.01
50	0.89	0.03
100	1.81	0.03
500	9.61	0.11
1000	19.24	0.21

2630

2631 For completeness, Table 23 reports the runtime of each evaluated method in Experiment 1 (Section
 2632 6.4) on the *NN-Large Sol* as the number of nodes increases (with the expected number of edges fixed
 2633 to N , the number of nodes). While some methods scale better than others, dataset generation with
 2634 the CausalProfiler remains efficient.

2635 Table 23: Runtime scaling of causal inference methods (in seconds). Each entry reports mean and
 2636 standard deviation across runs.

2637

Node Range	CausalNF	DCM	NCM	VACA
30–40	(1, 0.6)	(24, 8.4)	(12, 1.2)	(11, 4.6)
50–70	(2, 0.4)	(43, 12.8)	(22, 2.6)	(12, 4.8)
70–90	(3, 0.3)	(53, 18.0)	(29, 2.4)	(12, 4.7)
90–110	(4, 0.4)	(60, 23.0)	(36, 2.5)	(9, 2.5)

2644

2645 The apparent reduction in average VACA runtime is explained by its increasing failure rate. All other
 methods exhibit a 0% failure rate.