

CausalPlayground: Addressing Data-Generation Requirements in Cutting-Edge Causality Research

Anonymous CVPR submission

Paper ID *****

Abstract

001 *Research on causal effects often relies on synthetic data due*
002 *to the scarcity of real-world datasets with ground-truth ef-*
003 *fects. Since current data-generating tools do not always*
004 *meet all requirements for state-of-the-art research, ad-hoc*
005 *methods are often employed. This leads to heterogeneity*
006 *among datasets and delays research progress. We address*
007 *the shortcomings of current data-generating libraries by in-*
008 *troducing CausalPlayground, a Python library that provides*
009 *a standardized platform for generating, sampling, and shar-*
010 *ing structural causal models (SCMs). CausalPlayground*
011 *offers fine-grained control over SCMs, interventions, and*
012 *the generation of datasets of SCMs for learning and quan-*
013 *titative research. Furthermore, by integrating with Gym-*
014 *nasium, the standard framework for reinforcement learning*
015 *(RL) environments, we enable online interaction with the*
016 *SCMs. Overall, by introducing CausalPlayground we aim*
017 *to foster more efficient and comparable research in the field.*
018 *All code is available at [https://anonymous.4open.](https://anonymous.4open.science/r/CausalPlayground-D1B2/)*
019 *science/r/CausalPlayground-D1B2/. The API*
020 *documentation will be released after de-anonymization.*

021 1. Introduction

022 Ever since the formalization of causality [18] the field has
023 gained significant attention for improving inference from
024 data, scientific discovery, and others. Progress in causal-
025 ity research relies heavily on data collection about the sce-
026 nario that is being investigated. However, a major challenge
027 is the lack of real-world data with known causal ground
028 truth [6]. As an example, even after decades of research,
029 only a few strong real-world benchmark datasets are avail-
030 able for causal structure discovery [6, 32]. Consequently,
031 many state-of-the-art methods use synthetically generated
032 data, such as in causal representation learning [16], causal
033 discovery [9], among others.

034 Naturally, each research comes with its specific require-
035 ments for the data-generating process that is being inves-

tigated. In causality particularly though, some require- 036
ments can be observed that are common among many in- 037
stances. We identify them as requirements for a causal 038
data-generating library to be useful for a broad variety of 039
research questions as follows: 040

R1 Interventional data generation: Intuitively, interven- 041
tions are experiments in an environment that are crucial 042
for identifying causal effects [4]. Therefore, a general 043
causal data-generation framework must facilitate the sam- 044
pling of interventional data. 045

R2 Interaction with the causal model: It is becoming in- 046
creasingly clear, that interacting with the causal model is 047
beneficial for many tasks. For example, intervening after 048
every sample can improve sample efficiency for causal 049
discovery methods [22, 23] and causal inference [29]. 050
This requirement is further highlighted by the popularity 051
of interactive frameworks like Causal World [1] and var- 052
ious interactive causal models created ad-hoc for specific 053
research questions [15, 29, 35]. 054

R3 Fine-grained control over the causal model: As causal 055
inference and discovery often investigate settings with 056
clearly defined assumptions on the data, like the popu- 057
lar linear functions with additive with non-Gaussian noise 058
[19] assumption, a general data-generation library must 059
provide detailed control over the functional relations of 060
the causal model. 061

R4 Causal model generation for quantitative results: Re- 062
cent research in causal methods, such as [9, 17, 22, 23], 063
emphasize the need for training and evaluation not merely 064
on a single causal model, but rather on data-sets of mod- 065
els. This results in the requirement for a general causal 066
data-generating library to be able to easily create many 067
causal models. 068

As our comparison in Sec. 2 shows existing tools fall 069
short of addressing all current requirements within one 070
framework. This leads to many researchers having to rely 071
on ad-hoc data generation processes, introducing undesir- 072
able heterogeneity of datasets amongst methods, more dif- 073
ficult comparisons of results, and generally obstructing the 074
research progress of the field. 075

| | | interventional data | DAG-Generation | SCM-Generation | interactive | SCM-controlled | API-Documentation | Language |
|-------------------------|---|---------------------|----------------|----------------|-------------|----------------|-------------------|----------|
| Do-Why [25] | ✓ | X | X | X | X | ✓ | | P |
| TETRAD [21, 24] | X | ✓ | X | X | X | ✓ | | J/P/R |
| cause2e [8] | ✓ | ✓ | X | X | X | ✓ | | P |
| Lawrence et al. [14] | X | ✓ | X | X | ~ | X | | P |
| MANM-CS [11] | ✓ | ✓ | X | X | ~ | X | | P |
| BNlearn [27] | ~ | X | X | X | X | ✓ | | P |
| causalDag [5] | ~ | ✓ | X | X | X | ✓ | | P |
| gCastle [33] | X | ✓ | X | X | ~ | ✓ | | P |
| CDT [12] | X | ✓ | X | X | ~ | ✓ | | P |
| SCModels [2] | ✓ | X | X | X | ✓ | X | | P |
| R6causal [13] | ✓ | X | X | X | ✓ | ✓ | | R |
| CausalWorld [1] | ✓ | X | X | ✓ | X | ✓ | | P |
| SynTReN [31] | X | ✓ | X | X | X | X | | J |
| pcalg [3] | X | ✓ | X | X | X | ✓ | | R |
| JustCause [10] | X | X | X | X | ✓ | ✓ | | P |
| CausalPlayground | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | | P |

Table 1. Overview of software libraries that can be used for data-generation for causality research. Symbols ✓, X, and ~ indicate whether a requirement is met, not met, or partially met, respectively. J, P, R refer to Java, Python, and R, respectively. We evaluate the methods on the ability to sample interventional data, generate directed acyclic graphs (DAG), generate SCMs, interact with the SCM, the availability of a detailed API documentation, and the programming language that it is written in.

076 To offer a solution to this problem we created
077 *CausalPlayground*, a Python library for generating syn-
078 thetic data and providing generation processes for causal
079 models. Our library enables fine-grained control over the
080 models, integrates processes within the popular RL frame-
081 work *Gymnasium* [30] to accommodate the requirement for
082 interacting with the causal model, and implements the gener-
083 ation of many causal models at once to enable learn-
084 ing and quantitative research. It thereby constitutes a
085 framework for causal data generation that allows for more
086 standardized and easy-to-share causal data-generation pro-
087 cesses.

088 In the remainder of the paper we will compare current
089 causal data-generating libraries in Sec. 2, introduce funda-
090 mental notions of causality more formally in Sec. 3, pro-
091 vide an overview of *CausalPlayground* in Sec. 4, outline
092 a simple use-case in Sec. 5, and give an outlook on future
093 versions of this library in Sec. 6.

094 2. Related Work

095 Based on the requirements we identified in Sec. 1, we in-
096 vestigate the current landscape of general purpose causal
097 data-generation tools. We restrict our comparison to pack-
098 ages specifically designed for causality-related research, ac-
099 knowledging that other general-purpose tools have been
100 used, such as common RL environments like MuJoCo [28].

The comparison is summarized in Tab. 1.

Regarding interventional data-generation (R1), many
available packages, including Do-Why [25], cause2e [8],
MANM-CS [11], SCModels [2], R6causal [13], and
CausalWorld [1], allow for sampling from interventional
distributions for specifiable functional interventions. BN-
learn [27] and causalDag [5] offer less versatile methods for
intervening, and facilitating interventions by changing the
node conditional distributions or the graphical structure, re-
spectively.

Considering interaction with causal models (R2), only
CausalWorld [1] provides out-of-the-box functionality for
actively interacting with the model, highlighting an oppor-
tunity for useful new tools to meet this requirement.

With respect to generating causal models with specific
assumptions on the causal functions (R3) SCModels [2],
R6causal [13], and JustCause [10] give full control over
functional relations and noise distributions, while Lawrence
et al. [14], MANM-CS [11], gCastle [33], and CDT [12]
provide coarse-grained control that simultaneously applies
to all functions and distributions.

Lastly, to the best of our knowledge, no general-purpose
framework for causal data-generation provides methods to
generate sets of causal models with specific functional rela-
tions and noise distributions (R4).

Overall, this comparison shows a gap in libraries that si-
multaneously address some of the most important require-

128 ments for causal data-generation. This opens the opportu-
129 nity to develop a tool that is relevant to a broader set of
130 research questions.

131 3. Fundamental Concepts of Causality

132 In this section, we provide an overview of the necessary for-
133 mal concepts for our approach. Let an SCM M be a tuple
134 $M = (\mathcal{X}, \mathcal{U}, \mathcal{F}, \mathcal{P})$, where \mathcal{X} and \mathcal{U} are sets of endoge-
135 nous and exogenous random variables, respectively, \mathcal{F} is a
136 set of functions f_i that map the direct causes of $X_i \in \mathcal{X}$
137 to X_i , and \mathcal{P} is a set of pairwise independent probability
138 distributions P_i s.t. U_i follows distribution P_i with $U_i \in \mathcal{U}$.
139 Furthermore, we denote interventions as $do(X_i = g(\dots))$,
140 where g is an arbitrary function that takes as input a sub-
141 set of $\mathcal{X} \cup \mathcal{U}$. Such an intervention replaces function f_i
142 with function g . Applying a set of N arbitrary interventions
143 $a_t = \{do_0(\dots), \dots, do_N(\dots)\}$ to an SCM, means replac-
144 ing all corresponding functions simultaneously. Further-
145 more, each SCM M induces a distribution $P^M(\mathcal{X}, \mathcal{U} \mid a_t)$
146 that we can sample via ancestral sampling.

147 4. CausalPlayground

148 In this section we provide an overview of the *CausalPlay-*
149 *ground* library. By implementing the functionalities that ad-
150 dress the requirements R1 - R4, we provide an off-the-shelf
151 causal data-generation platform that can be used for state-
152 of-the-art causality research.

153 4.1. Fine-grained control over the SCM

154 For each endogenous variable X_i , the structural equation f_i
155 can be arbitrarily defined. Similarly, the distribution of the
156 exogenous variables P_j can be arbitrarily defined. Notably,
157 this entails the possibility of defining pixel-based environ-
158 ments. An example for creating the SCM with $\mathcal{F} = \{A \leftarrow$
159 $U + 5, Effect \leftarrow 2A\}$ and $U \sim Uniform(3, 8)$ is pro-
160 vided below:

```
161 scm = StructuralCausalModel()
162 scm.add_endogenous_var('A',
163     lambda noise: noise+5,
164     {'noise': 'U'})
165 scm.add_exogenous_var('U',
166     random.randint,
167     {'a': 3, 'b': 8})
168 scm.add_endogenous_var('Effect',
169     lambda x: x*2,
170     {'x': 'A'})
```

171 This level of control over the SCM is a significant im-
172 provement compared to other libraries, which often provide
173 more coarse-grained control.

174 4.2. Arbitrary interventions

Interventions can be applied to existing SCMs using arbi-
175 trary functions g . The only restriction is that the new func-
176 tion does not induce cyclic causal relations. In the example
177 code snippet:
178

```
scm.do_interventions([( "Effect",
179     (lambda a: a+1,
180     {'a': 'A'}) ]))
```

181 the intervention sets the value of the variable `Effect` to
182 the value of `A + 1`, using a lambda function that takes `a` as
183 an argument and a dictionary that maps `a` to the variable
184 name `A`. This, effectively, applies $do(Effect = A + 1)$ to
185 the SCM.
186

187 4.3. Interaction with the SCM

188 Users can intervene actively in existing SCMs. At each step,
189 a set of interventions a_t can be applied, after which a sample
190 of the endogenous and exogenous variables is drawn from
191 the induced distribution. The interventions are undone be-
192 fore the next step. The interaction is enabled by wrapping
193 the SCMs in the popular *Gymnasium* environment, facilitat-
194 ing standardized integration with (deep) RL methods [20].
195 In the example code snippet:

```
env = SCMEnvironment(scm,
196     possible_interventions=
197     [( "A", (lambda: 5, {})),
198     ("Effect", (lambda a: a+1,
199     {'a': 'A'}) ]))
```

201 An `SCMEnvironment` object is created by passing the
202 `scm` object and a list of two possible interventions with the
203 same syntax as described before. The resulting `env` object
204 can be used to interact with the SCM using the standard
205 *Gymnasium* environment interface, allowing for interactive
206 exploration and experimentation with the causal model.

207 More specifically, calling `env.step(action)` ap-
208 plies the interventions defined in the action via their index,
209 samples the intervened SCM, determines the new observa-
210 tion, termination flag, truncated flag, and reward, and, fi-
211 nally undoes the interventions. Importantly, the precise ob-
212 servation, termination, truncated, and reward functions can
213 be specified by the user.

214 4.4. SCM generation

215 *CausalPlayground*'s SCM generation procedure provides a
216 flexible approach to creating datasets of SCMs. Users can
217 specify the number of endogenous variables and exogenous
218 variables, the presence of confounders, and define the class
219 of possible causal relations. This is exemplified by the fol-
220 lowing code:

```
gen = SCMGenerator(all_functions=
221
```

```

222         {'linear': f_linear})
223 scm_unconfounded = gen.create_random(
224     possible_functions=["linear"],
225     n_endo=5, n_exo=4,
226     exo_distribution=random.gauss,
227     exo_distribution_kwargs=
228         {mu=0, sigma=1},
229     allow_exo_confounders=
230     False)[0]

```

231 that creates a random SCM with the causal relations follow-
 232 ing a predefined `f_linear` function, 5 endogenous vari-
 233 ables, 4 exogenous variables, a Gaussian exogenous distri-
 234 bution with $\mu = 0, \sigma = 1$ for each exogenous variable, and
 235 no confounders. Similarly, an SCM can also be generated
 236 from a given causal structure.

237 This comprehensive control over the SCM generation
 238 process facilitates the systematic evaluation and compari-
 239 son of causal models.

240 4.5. Further Implementation Details

241 *CausalPlayground* is implemented in Python with
 242 five main classes. The `StructuralCausalModel`
 243 class represents an SCM and provides methods for
 244 manipulation and sampling. The `SCMGenerator`
 245 creates random SCMs based on specified param-
 246 eters, while the `CausalGraphGenerator` and
 247 `CausalGraphSetGenerator` generate random causal
 248 graphs and sets of unique causal graphs, respectively. The
 249 `SCMEnvironment` class, allows interaction with an SCM.
 250 Furthermore, *CausalPlayground* relies on external classes
 251 such as `networkx.DiGraph`, `pandas.DataFrame`,
 252 `Gymnasium.Box`, and `Gymnasium.Sequence` for
 253 various functionalities. All code is published under
 254 the open-source MIT license and we highly welcome
 255 contributions.

256 5. Use-Case: Comparison of Causal Discovery 257 Algorithms

258 As one of many use-cases, we outline how to use the
 259 *CausalPlayground* library to compare different causal dis-
 260 covery algorithms. We generate synthetic data using the
 261 library’s functions and evaluate the performance of the al-
 262 gorithms on both confounded and unconfounded, automati-
 263 cally generated SCMs. We use the causal discovery algo-
 264 rithms provided by the `gCastle` [33] library. The detailed
 265 implementation can be found in the repository examples.

266 First, we generate distinct confounded
 267 and unconfounded causal graphs using the
 268 `CausalGraphSetGenerator` class from *CausalPlay-*
 269 *ground*. Each of the sets has 4 endogenous variables and
 270 4 exogenous variables. Based on these graphs we generate
 271 the SCMs for evaluation using the `SCMGenerator` class.




Figure 1. Experimental results of the causal discovery algorithm comparison use-case.

272 The functional relations are determined by specifiable
 273 functional relations. For our experiment we use linear
 274 additive functions and functions that add all causes and
 275 multiply the value of two randomly selected causes of a
 276 variable. More specifically, for every endogenous variable,
 277 the function is drawn from either of the two classes.

278 We then iterate over the 30 generated SCMs and draw
 279 100 samples per SCM. We evaluate the performance of each
 280 causal discovery algorithm on both confounded and uncon-
 281 founded datasets. More specifically, we compare PC [26],
 282 GES [7], and NOTEARS [34] and measure F1 score and the
 283 true-positive rate with `gCastle` [33]. The results can be seen
 284 in Fig. 1.

285 This use-case provides a glimpse into the usefulness of
 286 *CausalPlayground* for causality research, not the least be-
 287 cause the data-generating processes could now be shared
 288 among experiments and researchers.

289 6. Conclusion and Outlook

290 In this work, we identified the core requirements for data-
 291 generation in causal research and found that they are unsat-
 292 isfactorily met by the current landscape of data-generation
 293 tools. To offer a solution, we introduced *CausalPlayground*,
 294 a Python library that allows for sampling of, generation of,
 295 and online interaction with SCMs. By providing functionali-
 296 ty to address some of the main requirements of causality re-
 297 search, this library can enable a broad research community
 298 to easily share their SCMs and more rigorously compare its
 299 methods.

300 For future versions of this library, we are anticipating the
 301 integration of more diverse causal models such as Bayesian
 302 networks, and optimizations regarding parallelization for
 303 fast deployment on GPUs. Ultimately, we envision that our
 304 proposed library contributes to faster scientific progress in
 305 fields that rely on causality.

References

- [1] Ossama Ahmed, Frederik Träuble, Anirudh Goyal, Alexander Neitz, Manuel Wüthrich, Yoshua Bengio, Bernhard Schölkopf, and Stefan Bauer. CausalWorld: A Robotic Manipulation Benchmark for Causal Structure and Transfer Learning, 2020. 1, 2
- [2] Michael Aichmüller. Structural Causal Models, 2023. 2
- [3] Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464, 2012. 2
- [4] Elias Bareinboim, Juan D. Correa, Duligur Ibeling, and Thomas Icard. On Pearl’s Hierarchy and the Foundations of Causal Inference. In *Probabilistic and Causal Inference*, pages 507–556. ACM, New York, NY, USA, 2022. 1
- [5] Chandler Squires. causal dag: creation, manipulation, and learning of causal models, 2018. 2
- [6] Lu Cheng, Ruocheng Guo, Raha Moraffah, Paras Sheth, K Selçuk Candan, Huan Liu, Hong Kong, and Huan Liu are. Evaluation Methods and Measures for Causal Learning Algorithms. *IEEE Transactions on Artificial Intelligence*, 3(6), 2022. 1
- [7] David Maxwell Chickering. Optimal Structure Identification with Greedy Search. *J. Mach. Learn. Res.*, 3(null):507–554, 2003. 4
- [8] Daniel Grünbaum. cause2e: A Python package for end-to-end causal analysis., 2021. 2
- [9] Tristan Deleu, Mizu Nishikawa-Toomey, Jithendaraa Subramanian, Nikolay Malkin, Laurent Charlin, and Yoshua Bengio. Joint Bayesian Inference of Graphical Structure and Parameters with a Single Generative Flow Network. In *Advances in Neural Information Processing Systems*, pages 31204–31231. Curran Associates, Inc., 2023. 1
- [10] Maximilian Franz, Florian Wilhelm, and Tanmay Kulkarni. JustCause, 2020. 2
- [11] Johannes Huegle, Christopher Hagedorn, Lukas Böhme, Mats Pörschke, Jonas Umland, and Rainer Schlosser. MANM-CS: Data generation for benchmarking causal structure learning from mixed discrete-continuous and nonlinear data. In *WHY-21 at NeurIPS 2021*, 2021. 2
- [12] Diviyani Kalainathan and Olivier Goudet. Causal Discovery Toolbox: Uncover causal relationships in Python. 2019. 2
- [13] Juha Karvanen. R6causal: R6 Class for Structural Causal Models, 2023. 2
- [14] Andrew R. Lawrence, Marcus Kaiser, Rui Sampaio, and Maksim Sipos. Data Generating Process to Evaluate Causal Discovery Techniques for Time Series Data. 2021. 2
- [15] Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M Asano, Taco Cohen, and Stratis Gavves. Citris: Causal identifiability from temporal intervened sequences. In *International Conference on Machine Learning*, pages 13557–13603, 2022. 1
- [16] Phillip Lippe, Sara Magliacane, Sindy Löwe, Yuki M Asano, Taco Cohen, and Efstratios Gavves. BISCUIT: Causal Representation Learning from Binary Interactions. In *The 39th Conference on Uncertainty in Artificial Intelligence*, 2023. 1
- [17] Lars Lorch, Scott Sussex, Jonas Rothfuss, Andreas Krause, and Bernhard Schölkopf. Amortized Inference for Causal Structure Learning. In *Advances in Neural Information Processing Systems*, pages 13104–13118. Curran Associates, Inc., 2022. 1
- [18] Judea Pearl. *Causality*. Cambridge university press, 2009. 1
- [19] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017. 1
- [20] Aske Plaat. *Deep reinforcement learning*. Springer, 2022. 3
- [21] Joseph Ramsey and Bryan Andrews. Py-Tetrad and RPy-Tetrad: A New Python Interface with R Support for Tetrad Causal Search. In *Causal Analysis Workshop Series PMLR*, pages 40–51, 2023. 2
- [22] Andreas W M Sauter, Erman Acar, and Vincent Francois-Lavet. A Meta-Reinforcement Learning Algorithm for Causal Discovery. In *Proceedings of the Second Conference on Causal Learning and Reasoning*, pages 602–619. PMLR, 2023. 1
- [23] Andreas W M Sauter, Nicolò Botteghi, Erman Acar, and Aske Plaat. CORE: Towards Scalable and Efficient Causal Discovery with Reinforcement Learning. *arXiv preprint arXiv:2401.16974*, 2024. 1
- [24] Richard Scheines, Peter Spirtes, Clark Glymour, Christopher Meek, and Thomas Richardson. The TETRAD project: Constraint based aids to causal model specification. *Multivariate Behavioral Research*, 33(1):65–117, 1998. 2
- [25] Amit Sharma and Emre Kiciman. DoWhy: An End-to-End Library for Causal Inference. 2020. 2
- [26] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000. 4
- [27] Erdogan Taskesen. Learning Bayesian Networks with the bnlearn Python Package., 2020. 2
- [28] Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. *IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033, 2012. 2
- [29] Christian Toth, Lars Lorch, Christian Knoll, Andreas Krause, Franz Pernkopf, Robert Peharz, and Julius von Kügelgen. Active Bayesian Causal Inference. In *Advances in Neural Information Processing Systems*, pages 16261–16275. Curran Associates, Inc., 2022. 1
- [30] Mark Towers, Jordan K Terry, Ariel Kwiatkowski, John U Balis, Gianluca de Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Arjun KG, Markus Krimmel, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Andrew Tan Jin Shen, and Omar G Younis. Gymnasium, 2023. 2
- [31] Tim Van den Bulcke, Koenraad Van Leemput, Bart Naudts, Piet van Remortel, Hongwu Ma, Alain Verschoren, Bart De Moor, and Kathleen Marchal. SynTReN: a generator of synthetic gene expression data for design and analysis of structure learning algorithms. *BMC Bioinformatics*, 7(1):43, 2006. 2
- [32] Matthew J. Vowels, Necati Cihan Camgoz, and Richard Bowden. D’ya Like DAGs? A Survey on Structure Learning and Causal Discovery. *ACM Computing Surveys*, 55(4): 1–36, 2023. 1

- 419 [33] Keli Zhang, Shengyu Zhu, Marcus Kalander, Ignavier Ng,
420 Junjian Ye, Zhitang Chen, and Lujia Pan. gCastle: A Python
421 Toolbox for Causal Discovery. 2021. [2](#), [4](#)
- 422 [34] Xun Zheng, Bryon Aragam, Pradeep Ravikumar, and Eric P
423 Xing. DAGs with NO TEARS: Continuous Optimization for
424 Structure Learning. 2018. [4](#)
- 425 [35] Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal
426 Discovery with Reinforcement Learning. *ArXiv*,
427 abs/1906.04477, 2019. [1](#)