Intelligent Human Systems Integration (IHSI 2024), Vol. 119, 2024, 127–136
https://doi.org/10.54941/ahfe1004476

Causal Discovery for Observational Image Datasets: A Vision Paper

Atul Rawal[1], Adrienne Raglin[2], Ziying Tang[1], and Qianlong Wang[1]

[1]Department of Computer and Information Science, Towson University, Towson, MD 21009, USA
[2]US ARMY DEVCOM Army Research Laboratory, Adelphi MD USA

ABSTRACT

Artificial intelligence (AI) and machine learning (ML) systems have seen tremendous growth within the last few decades. Even with unprecedented new levels of autonomy for artificial reasoning systems, there are still challenges that remain. Challenges related to causal reasoning act as a roadblock for AI/ML systems to achieve human-like intelligence. For these systems to achieve human-like intelligence they must be able to gather causal information from given information. While causality for machine learning has made progress within the past years, there is still a lack of ability for AI/ML systems to generate causal relations from image datasets. To this end, this paper proposes a novel new perspective on discovering causal relations with image data by utilizing existing tools and methodologies.

Keywords: Artificial intelligence (AI), Machine learning (ML), Causal learning, Causal discovery image datasets

INTRODUCTION

Artificial Intelligence (AI) and Machine Learning (ML) have made tremendous advancements within the past decade leading to unprecedented levels of integration within human lives. AI/ML systems have been utilized for a plethora of applications ranging from recommender systems to biomedical systems for tumour/cancer detection. This level of integration has led to an increased interest for autonomous systems capable displaying of human-like intelligence. Research and development for AI/ML systems capable of human-like intelligence is hindered by the challenges associated with causal learning. The vital relationship between the cause and effect which comes naturally to human beings, has yet to be perfected by AI/ML systems. Even with the deployment of new state-of-the-art autonomous systems such as Tesla's fleet of electric vehicles and Boston Dynamics fleet of Spot and Atlas robots, these AI/ML systems still lack capability to identify the fundamental cause and effect relationships that are embedded into every action (Rawal et al., 2022). Towards this new frontier of human-like intelligence for AI/ML systems there has been an increased interest in causality for AI/ML systems leading to the coining of new terms such as causal learning (CL).

Causality is the cause-and-effect relation between a treatment and an outcome. It has been a fundamental scientific concept for numerous discoveries.

Here the cause describes the "*why*" whereas the effect describes the "*what*". It is often used interchangeably with correlation in an incorrect context. However, it has been well defined that correlation does not imply causation. While a critical concept for AI/ML and computer science, misrepresenting causation as correlation can have unwanted effects. Correlation is the relationship between two variables that follow a specific trend, while causality is the cause-and-effect relation between the two variables. Here the cause is responsible for the effect, and the effect relies on the cause (Morgan and Winship, 2015, Pearl, 2018). Causal learning refers to the study of these cause-and-effect relations between different variables in any given datasets for AI/ML systems. Therefore, causal learning should be viewed as a critical core component of any artificial reasoning systems, not just an ad-hoc feature.

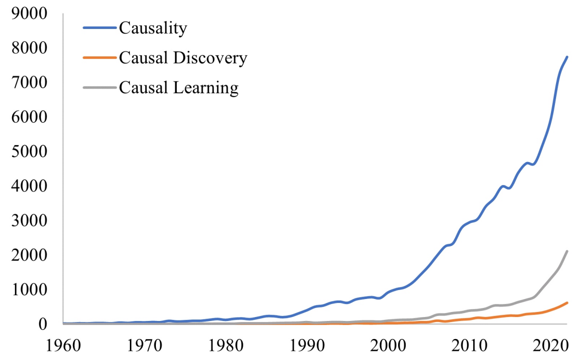

Figure 1: Yearly publications for causality, causal discovery, and causal learning. (Data derived from SCOPUS.)

The field of causal learning or causality for AI/ML systems is still fairly new and in its infant stage. While causality as a field of study has been around for quite some time and is well represented in literature, causal learning has only recently seen an increase in the research outcome. Even then, when compared to the other research subjects the number of publications each year highlights the relatively young age of the research field (Figure 1). Due to the relatively young age of the field, there are still challenges such as the lack of ground truth and methods for evaluating causality for multi-modal datasets that limit its progress. Challenges associated with the lack of ground truth for observational datasets have been widely acknowledged in literature (Rawal et al., 2021, Cheng et al., 2022, Rawal et al., 2022). To achieve causal learning for AI/ML systems novel new methods and perspectives need to be explored, as experimental data is not always feasible, and researchers must rely on observational datasets. While there have been studies on the use of observational datasets, most investigations have focused on causality from tabular datasets (Rohrer, 2018, Nichols, 2007, Maathuis et al., 2010). There is a gap of knowledge in literature for both methods and perspectives on generating causal relations in observational image datasets. In this vision paper, we present a novel perspective for causal learning to generate causal relations from observational image data via causal discovery. The paper is organized as follows: Section 2 provides the overview of causal learning and causal discovery, Section 3 provides the current state of the art for causal discovery for image datasets and provides our perspective, while Section 4 highlights the challenges and perspectives. Section 5 includes concluding remarks.

OVERVIEW OF CAUSAL LEARNING

As mentioned in the previous section causality is the relation between a cause and its effect (Guo et al., 2020). For AI/ML this goes beyond statistical correlation and association in the data. While correlation highlights specific trends between variables in the data, causation defines the cause-and-effect relations between those variables. Causality for AI/ML is the investigation of the change in the output(prediction/classification) of an AI/ML system caused by the change in a variable when another variable is modified/manipulated. Here the variable being modified is called the *treatment*, and the variable whose change is being investigated is called the *outcome*. Variables in the data that can affect both the treatment and the outcome are called *confounders*, and other background/noise variables within the data are referred to as the *covariates*.

Causal relations between the variables in data can be classified into three categories as listed by Judea Pearl's causal hierarchy: *association, intervention*, and *counterfactuals* (Pearl, 2009a, Pearl, 2019, Pearl, 2018). The first level of the hierarchy, association refers to the simple statistical correlation between variables in the data and is the building block for AI/ML systems where correlations are derived from the data to make informed predictions. The second level of the hierarchy, intervention is where each action's effect is investigated. Here specific treatment's modification/manipulation is investigated via the causal structure between variables. The final level of the hierarchy is called the counterfactuals, and it encompasses both association and intervention. Counterfactuals are utilized to generate causal relations underlying both the association and intervention levels to make predictions based on unknown outcomes. For AI/ML causal learning is utilized to answer two basic questions within a plethora of applications:

- To what extent does a change in one variable (treatment) have an impact on the target (outcome)?
- To cause a specific change in the target (outcome) which variable/s (treatment) needs to be modified upon?

These two questions are basis of causal learning which can be classified into two categories: *causal discovery* and *causal inference* (Gelman, 2011, Peters et al., 2017). Causal discovery is utilized in AI/ML to highlight the causal relations between the various variables within the data, whereas causal inference can be applied to investigate the causal effects of different treatments on an outcome. Investigations of causal relations via either of the two causal inference or causal discovery can be done through the two available formal frameworks called *structural causal models* (SCMs) and *potential outcome framework*. Structural causal models consist of *causal graphs* and *structural equations* and provide a holistic theory for causality (Pearl, 2009a, Guo et al., 2020, Pearl, 2009b, Yao et al., 2021). For this article, we focus on generating causal graphs from observational image datasets. They describe the causal effects between the different variables via a directed graph where different nodes represent different variables like outcome, treatment, confounders and covariates (Guo et al., 2020).

Causal Discovery for Machine Learning

Causal discovery can be utilized in numerous applications to generate causal relations within the data. This can be done via several available methods within literature. Here the focus is on the variables that have a cause-and-effect relation, where modifying variable A can have effects on another variable B. This is done by generating causal relations between the variables for the data identified by statistical relations (Spirtes et al., 2000, Schölkopf et al., 2012). Three general algorithms are available to identify causal relations via causal discovery: Constraint-based, Score-based and Functional Causal Model-based algorithms (Malinsky and Danks, 2018). Functional causal models are based on the structural equations to identity causal relations, whereas score and constraint-based models are based on statistical relations to identify causal relations and generate causal graphs (Yao et al., 2021).

Constraint-based algorithms identify the causal relations within data that satisfy the conditional independence based on the faithfulness assumption to generate causal graphs. Score-based algorithms were developed to overcome the faithfulness assumption of the constraint-based models and replace these with the goodness of fit tests. Based on the three general algorithms various models have been proposed in literature such as the *Peter-Clarke Algorithm* (Spirtes et al., 2000), *IC algorithm* and its variants (Spirtes et al., 2000, Pearl, 2009b, Fukumizu et al., 2007, Kalisch and Bühlman, 2007, Le et al., 2016, Ramsey, 2014, Sejdinovic et al., 2013, Zhang et al., 2012) for constraint-based models. Examples of score-based models in literature include the *Bayesian Information Criterion* (BIC) score (Schwarz, 1978), *Factorized Normalized Maximum Likelihood* (NML) universal model (Roos et al., 2008), *Bayesian Dirichlet* score (Heckerman et al., 1995), *Greedy Equivalent Search* (GES) (Chickering, 2002), *Fast GES* (Ramsey et al., 2017), and the adaptation of the Greedy SP algorithm by Wang et al (Wang et al., 2017). For functional causal models, examples include the Linear Non-Gaussian Acyclic Model (LiNGAM), and the *Independent Component Analysis LiNGAM* (ICA-LiNGAM) (Shimizu et al., 2006), *DirectLiNGAM* (Shimizu et al., 2011), and *auto-regressive LiNGAM* (Hyvärinen et al., 2010). Other additional methods have also been proposed for investigating causal relations such as the re-weighting, stratification, tree-based, representation learning and meta-learning methods (Yao et al., 2021). Stratification methods utilize classification/blocking to modify confounders (Imbens and Rubin, 2015). Matching based methods reduce the estimation bias to estimate the confounders, whereas tree-based models utilize decision trees. Examples of these methods include Classification and Regression Tree (CART) (Athey and Imbens, 2016), Bayesian Additive Regression Tree (BART) (Chipman et al., 2006, Chipman et al., 2010), and Random Forest methods (Wager and Athey, 2018).

Causal Discovery for Image Data

Since the field of causal learning is still in its infancy, all the methods and techniques mentioned above for causal discovery have been mainly investigated for tabular data. Numerous studies have reported on the use of causal discovery for various different applications. However, the utilization of causal discovery in image datasets has not been investigated as much. Even though there are some studies that investigate the use of causal discovery for images, more datasets, methods, techniques and concepts for causal discovery of images are needed. In this section we provide a brief vision for utilizing causal discovery for images using a novel perspective of existing methods and techniques. We also provide a brief overview of existing studies that have investigated causal discovery for images.

Castro et al., highlighted the importance of causality and causal learning for utilization in medical imaging applications (Castro et al., 2020). They presented the use of causal relationships between images and their annotations using causal graphs. For given medical images the study presented the use of *causal* and *anti-causal* tasks, where causal tasks are when an effect is predicted from a cause whereas, the cause is predicted from the effect in the anti-causal task. An example of a causal graph for cancer classification using medical images is presented. The authors also highlighted the advantages of causal learning for tackling challenges for medical imaging including data scarcity and data mismatch.

Li et al., presented a novel causal image synthesis method of generating causal models between MRI images and clinical/demographic variables for patients with Alzheimer's disease (Li et al., 2023). The study utilized structural causal models in conjunction with a styled generator to highlight causal relations and synthesize the images. The authors proposed the use of low dimensional latent feature representations of high dimensional 3D images to build causal relations between the image and the tabular data. Using these techniques, the authors are able to generate counterfactual 3D Brain MRI images and causal relationships between the tabular variables and the MRI images.

Chalupka et al., presented a framework to generalize causal learning in settings where the causal variables are reconstructed from the micro-variables for visual causes in images (Chalupka et al., 2014). The authors defined the visual cause as "*a function/feature of raw image pixels that has a causal effect on a target behaviour of a perceiving system of interest*". Here causal reasoning is employed to obtain improvement on the performance of correlation-based classifiers. The causal manipulator network presented in the papers can generate causal features from the data and perform causal modifications based on these features. From a causal discovery perspective, the authors present the utility of well-defined causal macro-variables to be used in causal graphs.

Anciukevicius et al., presented a scene understanding framework for object-centric 3D scenes capable of generalizing to out-of-distribution images (Anciukevicius et al., 2022). A causal generative model is designed to replicate the physical process of camera's image production for a multi-object scene.

A compositional generative causal model is utilized by the framework over multi-view images and the scenes depicted in the images.

Lopez-Paz et al., presented the Neural Causation Coefficient (NCC) to learn causation from labelled large image datasets (Lopez-Paz et al., 2017). The authors present the footprints to highlight the existence of "causal disposition" of object categories in images. The authors investigated causal discovery of observational image datasets by using a classifier to generate causal relations between pairs of random variables with samples from joint distribution. They employed the causal direction classifier to differentiate between object features and context features in static images. Using experiments on the MNIST dataset the study demonstrated the existence of observational signals which highlight the causal relations between objects. The study highlights the presence of causal information between objects and contexts within images using high order statistical properties of the datasets.

While these studies all present novel and practical techniques to utilize causal discovery for images, there is still a lack of a proper framework for generating causal graphs directly from image datasets. Therefore, we propose a novel framework of existing methodologies to generate causal relations via causal graphs from observational image datasets. Mainly we propose to use existing tools/techniques into a workflow to generate causal graphs from observational image datasets. This consists of the following four components:

- Feature extraction from images.
- Conversion of extracted features to tabular data.
- Causal discovery on the tabular data.
- Comparison of the causal features to ML model.

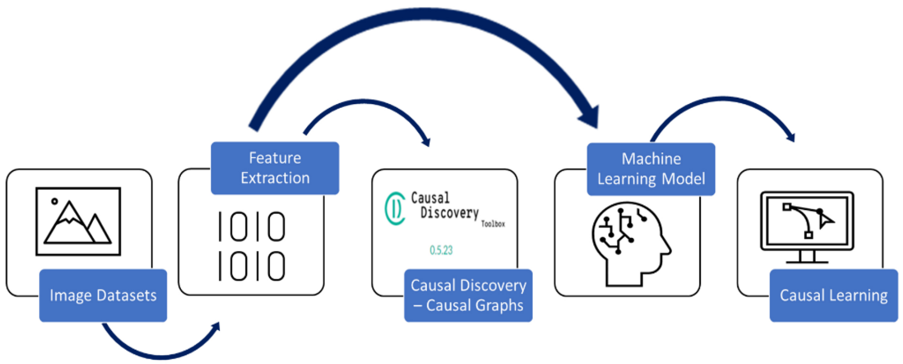

Figure 2: Proposed workflow for investigating causal discovery with image datasets.

Using the framework displayed in Figure 2 we propose to investigate the causal discovery in image datasets from existing tools and methodologies. The first step of the proposed framework is feature extraction from images. Feature representations from raw images are to be extracted and saved in a tabular format to be utilized for causal discovery using existing methods such as auto-encoders or deep neural networks. Once the extracted features are saved as tabular data, causal discovery can be performed in the tabular data via the generation of the causal graphs. The tabular data should consist of multiple features for each label. For this step a multitude of existing tools can be utilized such as the Causal Discovery Toolbox, Tetrad toolbox, and gCastle. The features with a direct causal link to the label can then be explored/investigated further using ML models. The comparison of the causal features to the features from correlation-based feature relevance can provide insights into potential biases within the data or the model. While there are some studies in literature that have focused on extracting causal relations from image datasets, the proposed study here is the first to extract causal features by generating causal graphs directly from the observational image datasets.

CHALLENGES AND PERSPECTIVES

Even though methods for causal discovery have made great progress, there are still challenges that hinder research for the use of causality for image data. Some of the challenges associated with causality in-general are also relevant for image datasets such as the lack of experimental datasets, the lack of ground-truth for observational datasets, causal discovery for time-series data, and the issues related to bias within the data and the model algorithms.

The lack of experimental image data has been a major challenge for causal reasoning research for image modality. For example, in biomedical applications of causality, such as tumor detection using MRI/CAT scan images, the lack of datasets large enough or fit enough to be trained and tested by ML models has been a major issue. Even though causal learning from observational tabular data has made great strides in the past few years, there is still a scarcity in literature highlighting the successful utilization of observational image data. It is not always feasible or practical for researchers to get access to experimental data. In lieu of experimental datasets synthetic data, emulating real world conditions that imitate experimental data can be utilized.

The absence of ground truth for observational image datasets is another major challenge for causal learning. This relates back to the lack of experimental data accompanied by the ground truth. However, for observational image datasets, when used for any AI/ML applications with causal learning, the lack of ground truth questions the validity of the data and hence the models themselves. An alternative for image datasets is the use of annotated labels as ground truth when experimental data with ground truth is not available. *Model selection* plays a crucial role for causal learning via observational data. Due to the lack of explicit AI/ML models for causal learning, the choice and utilization thereof will depend on the specific applications. Current studies in literature, such as the one by Moraffah et al., provide an excellent review of available causal interpretable models with criterions and metric for model evaluation (Moraffah et al., 2021, Kusner et al., 2017, Arjovsky et al., 2019).

Bias in data – has been a major challenge for most (if not all) applications of AI/ML systems. For causal learning, the ability to detect and mitigate biases from raw images will also be crucial for the practical deployment of these causal AI/ML systems. New laws and regulations around the world, such as European Union's GDPR, have increased the need for fair and unbiased predictive models. While experimental studies can mitigate bias to a certain degree via randomization, data imperfection in image datasets such as class imbalance, scarce/weak annotations, noise, and human errors can lead to unwanted bias within the data. From a big-data perspective, the presence of sample bias in observational data is also of concern and needs to be accounted and mitigated (Guyon et al., 2010, Moraffah et al., 2020, Stips et al., 2016).

CONCLUSION

AI/ML systems need to go beyond simple correlation to causation to achieve human-like intelligence. Therefore, causal learning is bound to play a crucial role in taking artificial reasoning systems into the next generation of advancements. For any AI/ML system dealing with image data, the ability to generate causal relations from observational images will be of great importance. This paper presented an overview of causality and causal learning for AI/ML systems. We proposed a new framework/workflow to investigate causality in observational image datasets. By utilizing existing tools and methodologies causal graphs can be generated via feature extraction for large image datasets. We also highlighted some of the challenges that must be addressed for the field of causal learning to advance into the next stage.

REFERENCES

Anciukevicius, T., Fox-Roberts, P., Rosten, E. & Henderson, P. 2022. Unsupervised Causal Generative Understanding of Images. *Advances in Neural Information Processing Systems,* 35, 37037–37054.

Arjovsky, M., Bottou, L., Gulrajani, I. & Lopez-Paz, D. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893.*

Athey, S. & Imbens, G. 2016. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences,* 113, 7353–7360.

Castro, D. C., Walker, I. & Glocker, B. 2020. Causality matters in medical imaging. *Nature Communications,* 11, 3673.

Chalupka, K., Perona, P. & Eberhardt, F. 2014. Visual causal feature learning. *arXiv preprint arXiv:1412.2309.*

Cheng, L., Guo, R., Moraffah, R., Sheth, P., Candan, K. S. & Liu, H. 2022. Evaluation methods and measures for causal learning algorithms. *IEEE Transactions on Artificial Intelligence,* 3, 924–943.

Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of machine learning research,* 3, 507–554.

Chipman, H., George, E. & Mcculloch, R. 2006. Bayesian ensemble learning. *Advances in neural information processing systems,* 19.

Chipman, H. A., George, E. I. & Mcculloch, R. E. 2010. BART: Bayesian additive regression trees.

Fukumizu, K., Gretton, A., Sun, X. & Schölkopf, B. 2007. Kernel measures of conditional dependence. *Advances in neural information processing systems,* 20.

Gelman, A. 2011. Causality and statistical learning. University of Chicago Press Chicago, IL.

Guo, R., Cheng, L., Li, J., Hahn, P. R. & Liu, H. 2020. A survey of learning causality with data: Problems and methods. *ACM Computing Surveys (CSUR),* 53, 1–37.

Guyon, I., Janzing, D. & Schölkopf, B. Causality: Objectives and assessment. Causality: Objectives and Assessment, 2010. PMLR, 1–42.

Heckerman, D., Geiger, D. & Chickering, D. M. 1995. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine learning,* 20, 197–243.

Hyvärinen, A., Zhang, K., Shimizu, S. & Hoyer, P. O. 2010. Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research,* 11.

Imbens, G. W. & Rubin, D. B. 2015. *Causal inference in statistics, social, and biomedical sciences*, Cambridge University Press.

Kalisch, M. & Bühlman, P. 2007. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research,* 8.

Kusner, M. J., Loftus, J., Russell, C. & Silva, R. 2017. Counterfactual fairness. *Advances in neural information processing systems,* 30.

Le, T. D., Hoang, T., Li, J., Liu, L., Liu, H. & Hu, S. 2016. A fast PC algorithm for high dimensional causal discovery with multi-core PCs. *IEEE/ACM transactions on computational biology and bioinformatics,* 16, 1483–1495.

Li, Y., Shi, J. & Zhou, S. K. 2023. Causal Image Synthesis of Brain MR in 3D. *arXiv preprint arXiv:2303.14349*.

Lopez-Paz, D., Nishihara, R., Chintala, S., Scholkopf, B. & bottou, L. Discovering causal signals in images. Proceedings of the IEEE conference on computer vision and pattern recognition, 2017. 6979–6987.

Maathuis, M. H., Colombo, D., Kalisch, M. & Bühlmann, P. 2010. Predicting causal effects in large-scale systems from observational data. *Nature methods,* 7, 247–248.

Malinsky, D. & Danks, D. 2018. Causal discovery algorithms: A practical guide. *Philosophy Compass,* 13, e12470.

Moraffah, R., Karami, M., Guo, R., Raglin, A. & Liu, H. 2020. Causal interpretability for machine learning-problems, methods and evaluation. *ACM SIGKDD Explorations Newsletter,* 22, 18–33.

Moraffah, R., Sheth, P., Karami, M., Bhattacharya, A., Wang, Q., Tahir, A., Raglin, A. & Liu, H. 2021. Causal inference for time series analysis: Problems, methods and evaluation. *Knowledge and Information Systems,* 63, 3041–3085.

Morgan, S. L. & Winship, C. 2015. *Counterfactuals and causal inference*, Cambridge University Press.

Nichols, A. 2007. Causal inference with observational data. *The Stata Journal,* 7, 507–541.

Pearl, J. 2009a. Causal inference in statistics: An overview.

Pearl, J. 2009b. *Causality*, Cambridge university press.

Pearl, J. 2018. Theoretical impediments to machine learning with seven sparks from the causal revolution. *arXiv preprint arXiv:1801.04016*.

Pearl, J. 2019. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM,* 62, 54–60.

Peters, J., Janzing, D. & Schölkopf, B. 2017. *Elements of causal inference: foundations and learning algorithms*, The MIT Press.

Ramsey, J., Glymour, M., Sanchez-Romero, R. & Glymour, C. 2017. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International journal of data science and analytics,* 3, 121–129.

Ramsey, J. D. 2014. A scalable conditional independence test for nonlinear, non-gaussian data. *arXiv preprint arXiv:1401.5031*.

Rawal, A., Mccoy, J., Rawat, D. B., Sadler, B. M. & Amant, R. S. 2021. Recent advances in trustworthy explainable artificial intelligence: Status, challenges, and perspectives. *IEEE Transactions on Artificial Intelligence,* 3, 852–866.

Rawal, A., Raglin, A., Rawat, D. B., Sadler, B. M. & Laboratory, D. A. R. 2022. Causality and Machine Learning Review.

Rohrer, J. M. 2018. Thinking clearly about correlations and causation: Graphical causal models for observational data. *Advances in methods and practices in psychological science,* 1, 27–42.

Roos, T., Silander, T., Kontkanen, P. & Myllymaki, P. Bayesian network structure learning using factorized NML universal models. 2008 Information Theory and Applications Workshop, 2008. IEEE, 272–276.

Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K. & Mooij, J. 2012. On causal and anticausal learning. *arXiv preprint arXiv:1206.6471*.

Schwarz, G. 1978. Estimating the dimension of a model. *The annals of statistics*, 461–464.

Sejdinovic, D., Sriperumbudur, B., Gretton, A. & Fukumizu, K. 2013. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The annals of statistics*, 2263–2291.

Shimizu, S., Hoyer, P. O., Hyvärinen, A., Kerminen, A. & Jordan, M. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research,* 7.

Shimizu, S., Inazumi, T., Sogawa, Y., Hyvarinen, A., Kawahara, Y., Washio, T., Hoyer, P. O., Bollen, K. & Hoyer, P. 2011. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *Journal of Machine Learning Research-JMLR,* 12, 1225-1248.

Spirtes, P., Glymour, C. N. & Scheines, R. 2000. *Causation, prediction, and search*, MIT press.

Stips, A., Macias, D., Coughlan, C., Garcia-Gorriz, E. & Liang, X. S. 2016. On the causal structure between CO2 and global temperature. *Scientific reports,* 6, 21691.

Wager, S. & Athey, S. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association,* 113, 1228–1242.

Wang, Y., Solus, L., Yang, K. & Uhler, C. 2017. Permutation-based causal inference algorithms with interventions. *Advances in Neural Information Processing Systems,* 30.

Yao, L., Chu, Z., Li, S., Li, Y., Gao, J. & Zhang, A. 2021. A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data (TKDD),* 15, 1–46.

Zhang, K., Peters, J., Janzing, D. & Schölkopf, B. 2012. Kernel-based conditional independence test and application in causal discovery. *arXiv preprint arXiv:1202.3775*.