Guest Editorial: Special Issue on Causal Discovery and Causality-Inspired Machine Learning

CAUSALITY is a fundamental notion in science and engineering. It has attracted much interest across research communities in statistics, machine learning (ML), healthcare, and artificial intelligence (AI), and is becoming increasingly recognized as a vital research area. One of the fundamental problems in causality is how to find the causal structure or the underlying causal model. Accordingly, one focus of this Special Issue is on *causal discovery*, i.e., how can we discover causal structure over a set of variables from observational data with automated procedures? Besides learning causality, another focus is on using causality to help understand and advance ML, that is, causality-inspired ML.

There has been impressive progress in theoretical and algorithmic developments on causal discovery from various types of data, including independent and identically distributed (i.i.d.) data with or without latent confounding or selection bias, and non-i.i.d. data under distribution shifts, in nonstationary settings, or with missing data. Moreover, recent years have also seen its practical applications in several scientific fields, such as neuroscience, climate, biology, and epidemiology. However, a number of practical issues, including confounding effects, the large scale of the data, the presence of measurement error, and complex causal mechanisms, are still to be properly addressed, in order to achieve reliable causal discovery in real-world scenarios.

On the other hand, causality-inspired ML (in the context of transfer learning, reinforcement learning, deep learning (DL), etc.) leverages ideas from causality to improve generalization, adaptivity, robustness, interpretability, and sample efficiency and is attracting more and more interest in ML and AI. For instance, off-policy evaluation, which is fundamentally a causal intervention issue, has received much attention in the deep reinforcement learning community. Despite the benefit of the causal view in transfer learning and reinforcement learning, several tasks in ML, such as dealing with adversarial attacks and learning disentangled representations, are closely related to the causal view and worth careful investigation, and cross-disciplinary efforts may facilitate the anticipated progress.

Inspired by such achievements and challenges, this Special Issue aims at reporting progress in fundamental principles, practical methodologies, efficient implementations, and applications of causal discovery and inference methods. The Special Issue also welcomes contributions to causality-inspired ML. We thank the authors of the 53 submissions to the special issues, some of which were extended versions of the papers accepted to the 2020 Conference on Neural Information Processing Systems (NeurIPS 2020) Workshop on Causal Discovery & Causality-Inspired Machine Learning. Thanks to the reviewers of the Special Issue, after their careful reviews, this Special Issue finally selected 13 papers to publish.

Six of them are on the classical problem of causal discovery, aiming to address various practical issues. Constraint-based causal discovery aims to find a Markov equivalence class, which may contain multiple directed acyclic graphs (DAGs), to satisfy conditional independence constraints discovered from data [3], and is usually applied to find causal relations among random variables. In contrast, Thams and Hansen [A1] are concerned with constraint-based causal structure learning for point processes and develop a test of local independence in point process data. The test relies on approximating point process intensities by basis expansions and using higher order interaction terms of events to fit intensities. The authors applied their approach to a real-world dataset on neuron spiking in turtles and found that their test led to sparser, more informative estimated networks.

Finding a causal direction between two variables is a fundamental causal discovery problem. The developments in this line of research made in the past 20 years mainly deal with continuous variables. In [A2], Wei et al. consider the causal direction between discrete variables. It exploits a discrete additive noise model, compares the dissimilarity of conditional distributions of the estimated noise across the two possible directions, and chooses the direction with a smaller dissimilarity as the correct causal direction.

Despite several advances in recent years, learning causal structures represented by DAGs remains a very challenging task in high-dimensional settings when the graphs to be learned are not sparse. In [A3], Fang et al. exploit a low-rank assumption regarding the (weighted) adjacency matrix of a DAG causal model to address this problem. Specifically, it establishes several useful results relating interpretable graphical conditions to the low-rank assumption and shows that interestingly, the maximum rank is highly related to hubs, suggesting that scale-free networks, which are frequently encountered in practice, tend to be low-rank.

In causal analysis, given two variables, a basic problem is to distinguish between direct causal influences, confounding effects (dependence between them because of their unmeasured common causes), and dependence induced by selection

Digital Object Identifier 10.1109/TNNLS.2024.3365968

bias. This Special Issue has two papers to address the issue of latent confounders. In [A4], Gilligan-Lee et al. develop a general heuristic that takes a causal discovery algorithm that can only distinguish purely directed causal relations and modifies it to also detect latent common causes. Their experimental results demonstrate that the modified algorithms can not only detect latent confounders but also preserve the performance in causal direction determination.

In addition to the detection of latent confounders, the discovery of true causal relations in the confounding case is an essential problem in causal discovery. Bellot and Schaar [A5] focus on applications where latent variables are known to have a widespread effect on many measured ones and shows that latent confounders, in this setting, leave a statistical footprint in the measured data distribution that allows for disentangling spurious and causal effects. It, accordingly, demonstrates that a sparse linear Gaussian DAG among measured variables may be recovered approximately and proposes a modified score-based algorithm, which may be implemented with general-purpose solvers and scale to high-dimensional problems.

The development of causal discovery is continually being enriched with new algorithms for learning causal graphs; each one of them usually requires a set of hyperparameters to tune. Given that the true graph is unknown and the learning task is unsupervised, the challenge to a practitioner is how to tune these choices. To address this challenge, Biza et al. [A6] propose out-of-sample causal tuning (OCT) that aims to select an optimal combination of hyperparameters.Technically, it treats a causal model as a set of predictive models and uses out-of-sample protocols for supervised methods. The method adopts an information-theoretic approach to be able to generalize to mixed data types and a penalty for dense graphs to penalize for complexity.

Thanks to the temporal constraint that effects cannot precede causes, causal discovery from time series is a traditional, natural way to find causality, as pioneered by the so-called Granger causal analysis [1]. Gao et al. [A7] notice both the directed and hierarchical features of brain functional networks and propose a new approach called stepwise multivariate Granger causality (SMGC) to estimate them. Their simulation studies demonstrate that the diverse and complex hierarchical organization can be embedded in simple directed networks, and the proposed SMGC could capture the multiple hierarchies of the directed network. It successfully revealed interesting properties of the multilevel hierarchical brain network.

Multivariate time series prediction is a classical problem closely related to causality in time series, and it relies on how to extract and exploit temporal dependence patterns in the multivariate time series. Yuan et al. [A8] propose a joint spatiotemporal feature learning framework for multivariate time series prediction, leveraging both temporal and spatial dependencies to enhance prediction accuracy. It consists of a module with multiple sparse autoencoders to extract latent spatial features and another module with multiple high-order fuzzy cognitive maps to model these spatial features and capture their temporal dynamics, and shows better/competitive results on four real-world datasets.

Observational studies of causal effects (often known as causal inference) and counterfactual inference [2], as traditional problems in the causality field, aim to estimate causal effects at the group or individual level even without random trials. Grecov et al. [A9] introduce a new method to estimate the causal effects of an intervention over multiple treated units by combining the techniques of probabilistic forecasting with global forecasting methods using DL models. The paper presents a framework for producing accurate quantile probabilistic forecasts for the counterfactual outcome, based on training a global autoregressive recurrent neural network model with conditional quantile functions on a large set of related time series. The authors further demonstrate how this probabilistic methodology added to the global DL technique to forecast the counterfactual trend and distribution outcomes overcomes many challenges faced by the baseline approaches to the policy evaluation problem.

In causal inference via randomized trials, one splits the population into control and treatment groups and then compares the average response across the two groups. To ensure that the difference between the two groups is caused by only the treatment, it is required that the control and the treatment groups have similar statistics. Covariate balancing methods have been widely adopted to increase such similarity. Interestingly, Babaei et al. [A10] empirically demonstrate that covariate balancing with the standardized mean difference covariate balancing measure is susceptible to worst-case treatment assignments. The authors further provide an optimization-based algorithm to find the adversarial treatment assignments. The findings suggest researchers be more cautious when using covariate balancing methods to investigate causal effects.

A traditional application of causal inference is to identify treatment effects. Rao et al. [A11] investigate causal modeling of a randomized clinical trial (RCT)-established causal association: the effect of classes of antihypertensive on incident cancer risk. The authors develop a transformer-based model, targeted bidirectional EHR transformer (T-BEHRT), coupled with doubly robust estimation to estimate the average risk ratio. They tested the model on simulated data and situations of limited data and found that their model provides more accurate estimates of relative risk least sum absolute error from ground truth, compared with benchmark estimations.

Causality-inspired ML is an emerging research direction in the field of ML, benefiting from the (causal) process perspective or well-designed causal techniques. Neto [A12] is concerned with static anticausal ML tasks (i.e., prediction tasks where the outcome causally influences the inputs), and proposes a counterfactual approach to train "causality-aware" predictive models that are able to leverage causal information in the anticausal setting. In applications plagued by confounding, it can generate predictions free from the influence of observed confounders. In the presence of observed mediators, it generates predictions that only capture the direct or the indirect causal influences. It is achieved by training supervised learners on (counterfactually) simulated inputs that retain only the associations generated by the causal relations of interest. The authors validated their approach on various synthetic data and illustrated its application to a real dataset.

Current ML schema typically uses a one-pass model inference (e.g., forward propagation) to make predictions in the testing phase. It is inherently different from human students who double-check their answers during examinations especially when the confidence is low. To bridge this gap, Deng et al. [A13] propose a learning to double-check (L2D) framework, which formulates double-check as a learnable procedure with two core operations: recognizing unreliable predictions and revising predictions. To judge the correctness of a prediction, it resorts to counterfactual faithfulness in causal theory and designs a contrastive faithfulness measure. Furthermore, the authors design a simple and effective revision module to revise the original model prediction according to faithfulness. The effectiveness of L2D in prediction correctness judgment and revision was validated on three classification models and two public datasets for image classification.

It is becoming increasingly clear that causal modeling benefits many tasks such as disease treatment, decision-making, recommender systems, adaptive/robust prediction, anomaly detection, and data generation. They involve finding causality from data, identifying causal effects, and reformulating and addressing ML problems with a causal view, which are the focus of this Special Issue. This collection of papers showcases an overview of the current development in those directions. We hope that the work presented here can motivate more exciting future works from readers to push causality research and its applications to a new height.

KUN ZHANG, *Guest Editor*
ILYA SHPITSER, *Guest Editor*
SARA MAGLIACANE, *Guest Editor*
DAVIDE BACCIU, *Guest Editor*
FEI WU, *Guest Editor*
CHANGSHUI ZHANG, *Guest Editor*
PETER SPIRTES, *Guest Editor*

APPENDIX: RELATED ARTICLES

[A1] N. Thams and N. R. Hansen, "Local independence testing for point processes," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4902–4910, Apr. 2024.

[A2] Y. Wei, X. Li, L. Lin, D. Zhu, and Q. Li, "Causal discovery on discrete data via weighted normalized Wasserstein distance," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4911–4923, Apr. 2024.

[A3] Z. Fang, S. Zhu, J. Zhang, Y. Liu, Z. Chen, and Y. He, "On low-rank directed acyclic graphs and causal structure learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4924–4937, Apr. 2024.

[A4] C. M. Gilligan-Lee, C. Hart, J. Richens, and S. Johri, "Leveraging directed causal discovery to detect latent common causes in cause-effect pairs," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4938–4947, Apr. 2024.

[A5] A. Bellot and M. V. D. Schaar, "Linear deconfounded score method: Scoring DAGs with dense unobserved confounding," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4948–4962, Apr. 2024.

[A6] K. Biza, I. Tsamardinos, and S. Triantafillou, "Out-of-sample tuning for causal discovery," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4963–4973, Apr. 2024.

[A7] Q. Gao et al., "A stepwise multivariate Granger causality method for constructing hierarchical directed brain functional network," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4974–4984, Apr. 2024.

[A8] K. Yuan, K. Wu, and J. Liu, "Is single enough? A joint spatiotemporal feature learning framework for multivariate time series prediction," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4985–4998, Apr. 2024.

[A9] P. Grecov et al., "Probabilistic causal effect estimation with global neural network forecasting models," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 4999–5013, Apr. 2024.

[A10] H. Babaei, S. Alemohammad, and R. G. Baraniuk, "Covariate balancing methods for randomized controlled trials are not adversarially robust," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 5014–5026, Apr. 2024.

[A11] S. Rao et al., "Targeted-BEHRT: Deep learning for observational causal inference on longitudinal electronic health records," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 5027–5038, Apr. 2024.

[A12] E. C. Neto, "Causality-aware predictions in static anticausal machine learning tasks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 5039–5053, Apr. 2024.

[A13] X. Deng, F. Feng, X. Wang, X. He, H. Zhang, and T.-S. Chua, "Learning to double-check model prediction from a causal perspective," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 4, pp. 5054–5063, Apr. 2024.

REFERENCES

[1] C. W. J. Grange, "Testing for causality: A personal viewpoint," *J. Econ. Dyn. Control*, vol. 2, pp. 329–352, Jan. 1980.

[2] J. Pearl, *Causality: Models, Reasoning and Inference*. Cambridge, U.K.: Cambridge Univ. Press, 2000

[3] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search* (Lecture Notes in Statistics). New York, NY, USA: Springer-Verlag, 1993.