

VCDF: A Validated Consensus-Driven Framework for Time Series Causal Discovery

Gene Yu¹[0000-0003-4867-1340]*, Ce Guo¹[0000-0002-0272-9175], and Wayne Luk¹[0000-0002-6750-927X]

Department of Computing, Imperial College London, United Kingdom
gene.yu23@alumni.imperial.ac.uk, c.guo@imperial.ac.uk,
w.luk@imperial.ac.uk

Abstract. Time-series causal discovery is essential for understanding dynamic systems, yet many existing methods remain sensitive to noise, non-stationarity, and sampling variability. We propose the Validated Consensus-Driven Framework (VCDF), a simple and method-agnostic layer that improves robustness by evaluating the stability of causal relations across blocked temporal subsets. VCDF requires no modification to base algorithms and can be applied to methods such as VARLiNGAM and PCMCI. Experiments on synthetic datasets show that VCDF improves VARLiNGAM by approximately 0.08–0.12 in both window and summary F1 scores across diverse data characteristics, with gains most pronounced for moderate-to-long sequences. The framework also benefits from longer sequences, yielding up to 0.18 absolute improvement on time series of length 1000 and above. Evaluations on simulated fMRI data and IT-monitoring scenarios further demonstrate enhanced stability and structural accuracy under realistic noise conditions. VCDF provides an effective reliability layer for time-series causal discovery without altering underlying modeling assumptions.

Keywords: Causal Discovery · Reliable Statistics · Vector Autoregression · Cross-validation · High-dimensional Data Analysis

1 Introduction

Causal discovery for multivariate time series plays a central role in understanding dynamic systems across domains such as finance [8,16], neuroscience [23], and climate science [21]. By identifying directional relations and temporal dependencies, causal models support downstream analyses including forecasting, diagnosis, and system interpretation. However, real-world time series frequently exhibit non-stationarity, non-linearity, heterogeneous noise, and distributional shifts [13]. In addition, many systems involve large numbers of interacting components whose complex network structure further complicates causal inference [4]. These characteristics make reliable structural estimation difficult.

* Corresponding author

Typical causal discovery algorithms such as VARLiNGAM and PCMCI remain widely used due to their interpretability and theoretical grounding [14,22]. Yet extensive evaluations indicate that their outputs may vary when noise levels, temporal regimes, or other data characteristics change [22,2]. Sensitivity to sampling variability and structural assumptions can lead to inconsistent causal graphs across data subsets [2,7]. This instability undermines reliability and challenges practical adoption in applications that require consistent structural estimates.

A key observation motivating this work is that such variability across data partitions is itself informative. Relations that consistently appear across multiple subsets are more likely to reflect genuine causal structure, whereas highly variable edges may arise from estimation noise or structural ambiguity. Similar ideas appear in robustness studies and ensemble-based approaches, where agreement across models is used as a measure of reliability [5]. However, the implications of these principles for validating causal structures in time series remain under-explored.

We introduce the **Validated Consensus-Driven Framework (VCDF)**, a simple and method-agnostic validation layer that evaluates the stability of causal relations across blocked temporal partitions. VCDF computes directional consistency and variability metrics for each candidate edge and removes those that fail to meet stability thresholds. Because VCDF only requires the base algorithm to output directed relations with an associated score or effect estimate, it can be applied to diverse causal discovery paradigms—including linear non-Gaussian models, constraint-based methods, and non-linear approaches—without modifying their modeling assumptions.

Extensive experiments on synthetic data demonstrate that VCDF improves structural accuracy across a wide range of temporal dynamics, noise conditions, and data-generating processes. Benefits are particularly pronounced for long sequences, where consensus validation can exploit richer temporal evidence. Experiments on simulated fMRI data and IT-monitoring scenarios further show that VCDF enhances robustness in complex real-world settings.

Our main contributions are:

- **Consensus-driven refinement:** A validation framework that assesses the stability of causal relations across blocked data partitions.
- **Method-agnostic integration:** VCDF applies to a broad range of causal discovery algorithms without altering their assumptions.
- **Comprehensive empirical evaluation:** Across synthetic, fMRI, and IT-monitoring datasets, VCDF consistently improves both window and summary F1 scores.

2 Background and Related Work

2.1 Time-Series Causal Discovery

Causal discovery in multivariate time series aims to uncover lagged and contemporaneous dependencies that govern system behavior [26]. Two complementary

representations are commonly used: *window causal graphs*, which model lag-specific effects, and *summary causal graphs*, which aggregate influences across all lags [2]. Compared with static causal inference, the temporal dimension introduces additional challenges, including variable delays, evolving relationships, and indirect pathways [22,24]. These challenges are amplified in the presence of non-stationarity, non-linear relations, non-Gaussian disturbances, and heterogeneous temporal regimes [13,27]. High dimensionality further increases estimation noise and computational complexity [6].

2.2 Representative Methods

A range of methods has been developed for time-series causal discovery. VAR-LiNGAM extends linear non-Gaussian acyclic models with vector autoregression, enabling identification of linear causal structures under non-Gaussian noise [14]. PCMCI and related approaches use conditional independence tests and momentary conditional independence measures to handle high-dimensional and partially non-linear dependencies [22,20]. Neural and optimization-based methods such as TCDF and DYNOTEARS extend causal modeling capabilities to non-linear architectures and continuous optimization formulations [18,19]. Despite their differences in modeling assumptions, these methods remain sensitive to noise, temporal heterogeneity, and sampling variability [17,7].

2.3 Robustness and Stability

Achieving robustness in causal structure estimation remains a significant challenge. Small perturbations in data, distributional shifts, or changes in sampling can lead to substantial variations in inferred graphs [2]. This instability is particularly problematic in time series, where dependencies span multiple lags and temporal regimes [27]. Recent studies highlight the need for scalable and reliable methods for practical deployment [12,11,10,9]. The idea of leveraging consistency across models also appears in ensemble and robustness research, such as negative-correlation learning [5], where agreement across multiple predictors is used as a measure of reliability. These insights motivate validation strategies that assess stability across temporal partitions.

2.4 Cross-Validation for Time Series

Cross-validation is widely used to evaluate generalization in predictive modeling [25,15]. However, time-series data require preserving temporal order, as random shuffling breaks autocorrelation structures [1]. Blocked K-fold cross-validation has been shown to be appropriate for autoregressive processes under mild assumptions [3]. While existing approaches focus primarily on prediction error, causal discovery requires validating the *structural* stability of inferred relations. This motivates VCDF, which uses blocked folds not for model selection or prediction assessment, but for edge-stability validation. In each run, one contiguous

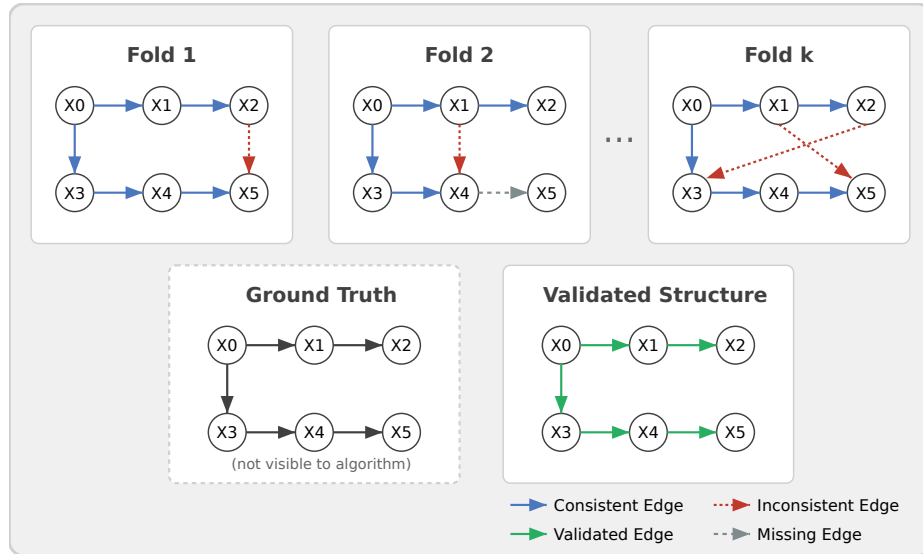


Fig. 1. Illustration of fold-to-fold variability and validated edges in VCDF. Consistent relations across partitions are retained; unstable ones are removed.

block is held out and the base method is applied to the remaining data, yielding leave-one-block-out causal estimates. VCDF retains relations that are stable across runs, measured by $C(r_{ij})$ and $V(r_{ij})$.

3 Validated Consensus-Driven Framework

3.1 Motivation

Outputs of causal discovery methods can vary substantially when data are noisy, non-stationary, or sampled differently over time. Applying the same algorithm to different partitions of the same dataset often yields inconsistent edges due to sensitivity to sampling variation. We treat this variability as informative: relations that appear consistently across subsamples are more likely to reflect genuine causal structure, whereas unstable edges may reflect estimation noise or structural ambiguity. This motivates a validation procedure that examines how each relation behaves across multiple blocked partitions of the data, as illustrated in Fig. 1.

3.2 Blocked Partitioning

The overall workflow of VCDF is shown in Fig. 2. Given a multivariate time series \mathcal{D} , VCDF divides it into k contiguous folds. For fold m , the held-out block $\mathcal{D}_m^{\text{val}}$ is excluded, and the remaining segments form the training set:

$$\mathcal{D}_m^{\text{train}} = \mathcal{D} \setminus \mathcal{D}_m^{\text{val}}.$$

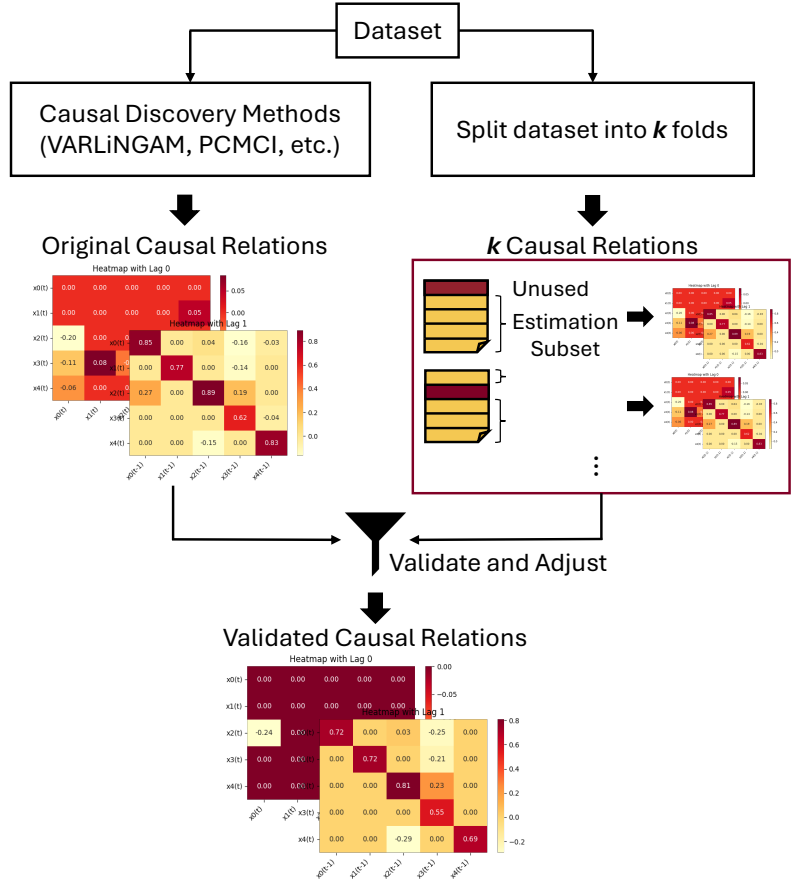


Fig. 2. Workflow of VCDF. Each fold produces a causal estimate; stability validation aggregates them into a final graph.

Temporal order is preserved, and the validation block is never used during estimation. Running a base causal discovery method (e.g., VARLiNGAM or PCMCI) independently on each $\mathcal{D}_m^{\text{train}}$ yields graphs $\{G_1, \dots, G_k\}$ as well as the full-sample estimate G_0 . Although removing a contiguous block may create a small number of cross-boundary pairs when concatenating segments, their impact is negligible when the maximum lag is small relative to sequence length.

3.3 Stability Metrics

Let r_{ij}^0 denote the effect estimate for edge $i \rightarrow j$ from the full dataset, and r_{ij}^m the corresponding estimate from fold m . VCDF evaluates each relation with two

metrics. First, directional consistency:

$$C(r_{ij}) = \frac{1}{k} \sum_{m=1}^k \mathbf{1}\{\text{sign}(r_{ij}^m) = \text{sign}(r_{ij}^0)\},$$

and second, relative variability:

$$V(r_{ij}) = \frac{\text{std}(\{r_{ij}^m\})}{|r_{ij}^0| + \epsilon}.$$

VCDF assumes that effect estimates are comparable across folds. Stable relations should exhibit high directional consistency and low normalized variability.

3.4 Filtering and Optional Adjustment

VCDF retains only relations that satisfy the stability thresholds:

$$C(r_{ij}) \geq \tau_c, \quad V(r_{ij}) \leq \tau_v.$$

In our experiments, thresholds are set as method-specific fixed defaults selected once via a synthetic sweep and then kept unchanged across datasets: for VAR-LiNGAM, $\tau_c=0.4, \tau_v=0.4$; for PCMCI, $\tau_c=0.7, \tau_v=0.4$. Optionally, effect-size refinement interpolates between the full-sample estimate and the cross-fold mean:

$$r_{ij}^{\text{adj}} = (1 - w) r_{ij}^0 + w \cdot \text{mean}(\{r_{ij}^m\}),$$

where $w \in [0, 1]$. In this work we use $w = 0$ to focus on structural validation; effect-size refinement is optional when magnitude estimation is of interest.

3.5 Algorithm

The procedure is summarized in Algorithm 1. VCDF produces a refined version of G_0 by removing unstable edges but does not introduce new ones, as it relies on signals supported by the full dataset.

VCDF introduces moderate runtime overhead that scales approximately linearly with k and depends primarily on the cost of the base causal discovery method. In typical settings, k ranges from 5 to 10, offering a practical balance between stability and computational cost.

4 Experimental Evaluation

We evaluate VCDF on both synthetic and domain-specific time-series datasets. The goal is to assess whether consensus-based validation improves structural accuracy across diverse data characteristics and temporal scales. Following Asaad et al. [2], we report both window F1 (lag-specific) and summary F1 (time-aggregated) scores as standard structural metrics, while additional measures show consistent trends.

Algorithm 1 Validated Consensus-Driven Framework (VCDF)

Require: Dataset \mathcal{D} , folds k , thresholds τ_c, τ_v

- 1: $G_0 \leftarrow \text{BaseCausalDiscovery}(\mathcal{D})$
- 2: **for** $m = 1$ to k **do**
- 3: Split \mathcal{D} into $\mathcal{D}_m^{\text{train}}$ and $\mathcal{D}_m^{\text{val}}$
- 4: $G_m \leftarrow \text{BaseCausalDiscovery}(\mathcal{D}_m^{\text{train}})$
- 5: **end for**
- 6: **for** each relation $i \rightarrow j$ in G_0 **do**
- 7: Compute $C(r_{ij})$ and $V(r_{ij})$
- 8: **if** $C(r_{ij}) < \tau_c$ or $V(r_{ij}) > \tau_v$ **then**
- 9: Remove $i \rightarrow j$ from G_0
- 10: **end if**
- 11: **end for**
- 12: **return** G_0

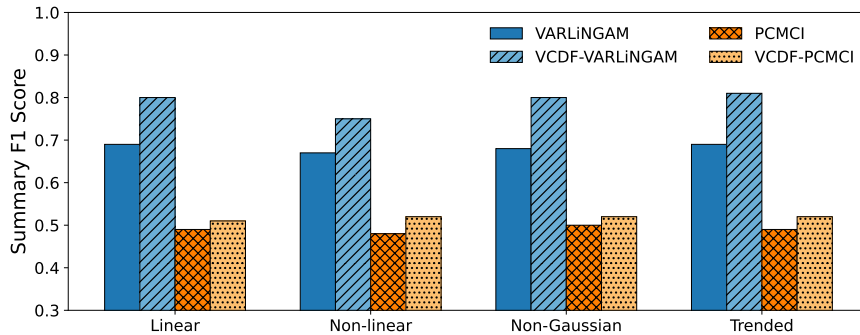


Fig. 3. Comparison of VARLiNGAM, VCDF-VARLiNGAM, PCMCI, and VCDF-PCMCI across representative synthetic settings.

4.1 Synthetic Data

We use four representative synthetic settings widely adopted in causal discovery benchmarks, capturing (1) linear vs. nonlinear relations, (2) Gaussian vs. non-Gaussian errors, and (3) trended vs. stationary temporal dynamics. Each configuration contains 15 variables with multiple dataset realizations.

We compare VCDF against four causal discovery paradigms: VARLiNGAM, PCMCI, TCDF, and DYNOTEARS. Since TCDF and DYNOTEARS rely on regularization or neural filtering, their variability across folds is smaller; thus, we focus detailed analysis on VARLiNGAM and PCMCI, where consensus validation yields clearer improvements.

Table 1 summarizes results across characteristics. VCDF consistently boosts both window and summary F1 scores. In the linear setting, VCDF-VARLiNGAM improves summary F1 from 0.69 to 0.80. Across all four characteristics, VARLiNGAM gains 0.08–0.12 summary F1 after applying VCDF. PCMCI also benefits from consensus filtering, with summary F1 improvements of 0.02–0.04 across

Table 1. Performance across four representative synthetic settings.

Setting	Method	F1 score	
		Window	Summary
Linear	PCMCI	0.37 ± 0.02	0.49 ± 0.03
	VCDF-PCMCI	0.43 ± 0.02	0.51 ± 0.02
	VARLiNGAM	0.64 ± 0.03	0.69 ± 0.03
	VCDF-VARLiNGAM	0.75 ± 0.03	0.80 ± 0.03
	TCDF	0.38 ± 0.03	0.43 ± 0.03
	DYNOTEARS	0.48 ± 0.04	0.53 ± 0.04
Non-linear	PCMCI	0.36 ± 0.01	0.48 ± 0.01
	VCDF-PCMCI	0.44 ± 0.02	0.52 ± 0.02
	VARLiNGAM	0.61 ± 0.04	0.67 ± 0.04
	VCDF-VARLiNGAM	0.69 ± 0.03	0.75 ± 0.03
	TCDF	0.35 ± 0.01	0.38 ± 0.01
	DYNOTEARS	0.54 ± 0.04	0.60 ± 0.03
Non-Gaussian	PCMCI	0.38 ± 0.02	0.50 ± 0.02
	VCDF-PCMCI	0.43 ± 0.02	0.52 ± 0.03
	VARLiNGAM	0.62 ± 0.04	0.68 ± 0.05
	VCDF-VARLiNGAM	0.74 ± 0.04	0.80 ± 0.04
	TCDF	0.40 ± 0.04	0.45 ± 0.04
	DYNOTEARS	0.47 ± 0.04	0.52 ± 0.04
Trended	PCMCI	0.38 ± 0.02	0.49 ± 0.02
	VCDF-PCMCI	0.42 ± 0.03	0.52 ± 0.02
	VARLiNGAM	0.64 ± 0.03	0.69 ± 0.03
	VCDF-VARLiNGAM	0.76 ± 0.04	0.81 ± 0.04
	TCDF	0.36 ± 0.03	0.40 ± 0.04
	DYNOTEARS	0.46 ± 0.05	0.51 ± 0.05

all characteristics. These results confirm that VCDF effectively filters unstable edges while preserving true causal relations.

Figure 3 visualizes the improvements. VCDF-VARLiNGAM shows the largest gains across settings, while VCDF-PCMCI reliably enhances its baseline.

Effect of Sequence Length. To evaluate temporal scalability, we vary sequence length from $T = 250$ to $T = 2000$. Table 2 shows that baseline VARLiNGAM saturates beyond $T \approx 500$, whereas VCDF-VARLiNGAM continues to improve, reaching window F1 ≈ 0.79 and summary F1 ≈ 0.82 – 0.83 at $T = 1000$ – 2000 . These results indicate that VCDF leverages additional temporal information more effectively, enhancing robustness without overfitting to noise. For very short sequences (e.g., $T=250$), base methods may be underpowered, and VCDF’s estimates are computed from even shorter effective samples, which can reduce the benefit of stability filtering and occasionally lead to slight degradation.

Table 2. Performance across different time-series lengths (synthetic, $n = 15$).

Length T	Method	F1 score	
		Window	Summary
250	VARLiNGAM	0.65 ± 0.05	0.71 ± 0.03
	VCDF-VARLiNGAM	0.65 ± 0.04	0.67 ± 0.03
	PCMCI	0.39 ± 0.03	0.52 ± 0.03
1000	VARLiNGAM	0.64 ± 0.02	0.71 ± 0.03
	VCDF-VARLiNGAM	0.79 ± 0.04	0.82 ± 0.04
	PCMCI	0.38 ± 0.02	0.53 ± 0.01
2000	VARLiNGAM	0.61 ± 0.03	0.66 ± 0.02
	VCDF-VARLiNGAM	0.79 ± 0.02	0.83 ± 0.02
	PCMCI	0.37 ± 0.03	0.52 ± 0.03

4.2 Domain-Specific Datasets

We further evaluate VCDF on simulated fMRI data [23] and three IT monitoring systems (antivirus, middleware, and web services). Ground-truth causal structures correspond to system topology, enabling reliable evaluation.

On the fMRI benchmark, VCDF-PCMCI achieves the best summary F1, while VCDF-VARLiNGAM improves temporal precision with the highest window F1. In IT-monitoring datasets, VCDF-VARLiNGAM improves summary F1 by approximately 10–15% across all systems, demonstrating that consensus-based refinement remains effective under realistic, noisy operational conditions.

4.3 Runtime Considerations

VCDF incurs moderate computational overhead due to repeated estimation on k partitions. Runtime grows approximately linearly with T and k . For 15-variable systems with $T = 1000$, VCDF-VARLiNGAM runs in several seconds, remaining practical for typical time-series analysis. Table 4 reports runtime comparisons.

4.4 Summary

Across synthetic and domain-specific datasets, VCDF consistently improves the stability and accuracy of time-series causal discovery. Synthetic benchmarks show that, for VARLiNGAM, VCDF yields gains of 0.08–0.12 in both window and summary F1 scores across linear, nonlinear, non-Gaussian, and trended settings. VCDF benefits particularly from longer sequences, where baseline methods tend to saturate. Results on fMRI and IT-monitoring datasets further confirm that consensus filtering improves robustness under realistic noise conditions. Overall, VCDF acts as an effective reliability layer that complements a range of causal discovery algorithms.

Table 3. Performance on domain-specific datasets (summary F1 shown where applicable).

Dataset	Method	F1 score	
		Window	Summary
fMRI	VARLiNGAM	0.48 ± 0.07	0.61 ± 0.12
	VCDF-VARLiNGAM	0.51 ± 0.14	0.57 ± 0.17
	VCDF-PCMCI	0.39 ± 0.12	0.64 ± 0.11
IT Antivirus	VARLiNGAM	—	0.35
	VCDF-VARLiNGAM	—	0.49
	VCDF-PCMCI	—	0.36
IT MoM	VARLiNGAM	—	0.44
	VCDF-VARLiNGAM	—	0.56
	VCDF-PCMCI	—	0.47
IT Web	VARLiNGAM	—	0.45
	VCDF-VARLiNGAM	—	0.55
	VCDF-PCMCI	—	0.47

Table 4. Runtime comparison (seconds) on synthetic datasets with 15 variables.

Length T	VARLiNGAM	VCDF-VARLiNGAM
250	0.50	2.83
500	0.65	3.62
1000	0.96	5.15
2000	1.75	8.66

5 Conclusion

We introduced VCDF, a consensus-driven validation framework that enhances the robustness of time-series causal discovery by evaluating stability across temporal blocks. As a model-agnostic post-processing layer, VCDF can be applied to existing algorithms without altering their assumptions, and improves both window and summary F1 scores across diverse synthetic conditions, particularly for moderate-to-long sequences. Experiments on fMRI and IT-monitoring datasets further demonstrate its practical effectiveness in noisy, real-world-like settings. VCDF is a validation layer, so recall is bounded by the base method. When causal structure genuinely changes over time, cross-partition inconsistency may reflect true change rather than noise.

These results suggest that consensus validation provides a simple yet reliable mechanism for strengthening causal discovery pipelines. Future work includes extending VCDF to scenarios with latent confounders, exploring adaptive stability thresholds, and developing more scalable implementations for high-dimensional or streaming data.

Our source code is available at <https://github.com/sonnets-project/vcdf-time-series-causal-discovery>.

Acknowledgments. We thank the anonymous reviewers and the dataset providers. We are also grateful to Dr. Paul Labonne (Bank of England) for his valuable feedback. This work is supported by the United Kingdom EPSRC (grant numbers UKRI256, EP/V028251/1, EP/N031768/1, EP/S030069/1, and EP/X036006/1), Altera, Intel, AMD, and SiliconFlow.

References

1. Arlot, S., Celisse, A.: A survey of cross-validation procedures for model selection. *Statistics Surveys* **4**, 40–79 (2010). <https://doi.org/10.1214/09-SS054>
2. Assaad, C.K., Devijver, E., Gaussier, E.: Survey and evaluation of causal discovery methods for time series. *Journal of Artificial Intelligence Research* **73**, 767–819 (2022). <https://doi.org/10.1613/jair.1.13428>
3. Bergmeir, C., Hyndman, R.J., Koo, B.: A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis* **120**, 70–83 (2018). <https://doi.org/10.1016/j.csda.2017.11.003>
4. Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., Hwang, D.U.: Complex networks: Structure and dynamics. *Physics Reports* **424**(4-5), 175–308 (2006). <https://doi.org/10.1016/j.physrep.2005.10.009>
5. Chen, H., Yao, X.: Regularized negative correlation learning for neural network ensembles. *IEEE Transactions on Neural Networks* **20**(12), 1962–1979 (2009). <https://doi.org/10.1109/TNN.2009.2034144>
6. Colombo, D., Maathuis, M.H., Kalisch, M., Richardson, T.S.: Learning high-dimensional directed acyclic graphs with latent and selection variables. *The Annals of Statistics* **40**(1), 294–321 (2012). <https://doi.org/10.1214/11-AOS940>
7. Glymour, C., Zhang, K., Spirtes, P.: Review of causal discovery methods based on graphical models. *Frontiers in Genetics* **10**, 524 (2019). <https://doi.org/10.3389/fgene.2019.00524>
8. Granger, C.W.J.: Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37**(3), 424–438 (1969). <https://doi.org/10.2307/1912791>
9. Guo, C., Cupello, D., Luk, W., Levine, J., Warren, A., Brookes, P.: FPGA-accelerated causal discovery with conditional independence test prioritization. In: *Proceedings of the 33rd International Conference on Field-Programmable Logic and Applications*. pp. 182–188 (2023). <https://doi.org/10.1109/FPL60245.2023.00033>
10. Guo, C., Luk, W.: Accelerating constraint-based causal discovery by shifting speed bottleneck. In: *Proceedings of the 2022 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. pp. 169–179. *FPGA '22*, Association for Computing Machinery (2022). <https://doi.org/10.1145/3490422.3502363>
11. Hagedorn, C.: Parallel execution of causal structure learning on graphics processing units. Ph.D. thesis, Universität Potsdam (2023)
12. Hagedorn, C., Lange, C., Huegle, J., Schlosser, R.: Gpu acceleration for information-theoretic constraint-based causal discovery. In: *KDD Workshop on Causal Discovery*. pp. 30–60 (2022)

13. Huang, B., Zhang, K., Zhang, J., Ramsey, J., Sanchez-Romero, R., Glymour, C., Schölkopf, B.: Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research* **21**, 1–53 (2020)
14. Hyvärinen, A., Zhang, K., Shimizu, S., Hoyer, P.O.: Estimation of a structural vector autoregression model using non-Gaussianity. *Journal of Machine Learning Research* **11**, 1709–1731 (2010)
15. Kohavi, R.: A study of cross-validation and bootstrap for accuracy estimation and model selection. In: *Proceedings of the 14th International Joint Conference on Artificial Intelligence*. pp. 1137–1145 (1995)
16. Kristjanpoller, W., Michell, K., Llanos, C., Minutolo, M.C.: Incorporating causal notions to forecasting time series: a case study. *Financial Innovation* **11**(1), 15 (2025). <https://doi.org/10.1186/s40854-024-00681-9>
17. Mooij, J.M., Peters, J., Janzing, D., Zscheischler, J., Schölkopf, B.: Distinguishing cause from effect using observational data: Methods and benchmarks. *Journal of Machine Learning Research* **17**, 1–102 (2016)
18. Nauta, M., Bucur, D., Seifert, C.: Causal discovery with attention-based convolutional neural networks. *Machine Learning and Knowledge Extraction* **1**(1), 312–340 (2019). <https://doi.org/10.3390/make1010019>
19. Pamfil, R., Sriwattanaworachai, N., Desai, S., Pilgerstorfer, P., Georgatzis, K., Beaumont, P., Aragam, B.: DYNOTEARS: Structure learning from time-series data. In: *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*. *Proceedings of Machine Learning Research*, vol. 108, pp. 1595–1605 (2020)
20. Runge, J.: Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets. In: *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence*. pp. 1388–1397 (2020)
21. Runge, J., Gerhardus, A., Varando, G., Eyring, V., Camps-Valls, G.: Causal inference for time series. *Nature Reviews Earth & Environment* **4**(7), 487–505 (2023). <https://doi.org/10.1038/s43017-023-00431-y>
22. Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., Sejdinovic, D.: Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances* **5**(11), eaau4996 (2019). <https://doi.org/10.1126/sciadv.aau4996>
23. Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W.: Network modelling methods for fMRI. *NeuroImage* **54**(2), 875–891 (2011). <https://doi.org/10.1016/j.neuroimage.2010.08.063>
24. Spirtes, P., Glymour, C., Scheines, R.: *Causation, prediction, and search*. MIT Press (2001)
25. Stone, M.: Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society Series B* **36**(2), 111–133 (1974). <https://doi.org/10.1111/j.2517-6161.1974.tb00994.x>
26. Yao, L., Chu, Z., Li, S., Li, Y., Gao, J., Zhang, A.: A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data* **15**(5), 1–46 (2021). <https://doi.org/10.1145/3444944>
27. Zhang, K., Huang, B., Zhang, J., Glymour, C., Schölkopf, B.: Causal discovery from nonstationary/heterogeneous data: Skeleton estimation and orientation determination. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. pp. 1347–1353 (2017)