

Leveraging Pre-Images to Discover Nonlinear Relationships in Multivariate Environments

M. Ali Vosoughi

Graduate Student Member, IEEE

Dept. of Electrical and Comp. Eng.

University of Rochester, Rochester, NY, USA

Email: mvosough@ece.rochester.edu

Axel Wismüller^{1,2,3,4}

¹Dept. of Imaging Sciences, ²Electrical and Comp. Eng., ³Biomed. Eng.,

University of Rochester Medical Center, Rochester, NY, USA

⁴Faculty of Medicine, Ludwig Maximilian University, Munich, Germany

Abstract—Causal discovery, beyond the inference of a network as a collection of connected dots, offers a crucial functionality in scientific discovery using artificial intelligence. The questions that arise in multiple domains, such as physics, physiology, the strategic decision in uncertain environments with multiple agents, climatology, among many others, have roots in causality and reasoning. It became apparent that many real-world temporal observations are nonlinearly related to each other. While the number of observations can be as high as millions of points, the number of temporal samples can be minimal due to ethical or practical reasons, leading to the curse-of-dimensionality in large-scale systems. This paper proposes a novel method using kernel principal component analysis and pre-images to obtain nonlinear dependencies of multivariate time-series data. We show that our method outperforms state-of-the-art causal discovery methods when the observations are restricted by time and are nonlinearly related. Extensive simulations on both real-world and synthetic datasets with various topologies are provided to evaluate our proposed methods.

Index Terms—causal graphical models, discovering nonlinear interactions, network inference, causality, large-scale time-series

I. INTRODUCTION

Causal relationships are meaningful and apply to many problems, as evidenced by the DARPA's recent challenge on machine common-sense (MCS) [1]. Such applications require an AI machine with a mechanism to understand and reason the system's underlying relations, which solo statistics cannot plead (see [2] for the condition that statistics fail). One may argue that such AI machines are causal because AI that finds rules is equivalent to AI that finds causality: when a causal relationship is strong enough, we call it a rule [3, Chapter 1]. An inherent advantage of learning reasons is learning with few samples. Causality gained the most traction recently to build models with sufficient robustness, generalizability, interpretability, and the capability to discover semantically meaningful representations.

The completeness of identifiability and development of *do-calculus* paved the way for causal inference using observa-

tional data [2]. However, using solo temporal measurements, the generative process underneath data is unknown, making the use of *do-calculus* unsettled. Granger causality (GC) is widely used to infer causality in the time-series data given the faithfulness assumption. GC initially stated for linear relations [4], and states that a time-series A is a cause for time-series B if A causes a better prediction for B.

Although there is a rich body of literature on discovering nonlinear relations from time-series data in the recent decade, yet it is a challenging problem for two reasons: 1) the detected nonlinear relations can be spurious, which requires a multivariate approach, and 2) the curse-of-dimensionality and ill-posedness becomes the main problem with an increasing number of nodes [5], [6]. Various algorithms have been proposed to recover nonlinear relations in networks, including Ancona's work [7], where they present conditions to detect nonlinear relationships, and in [8] Marinazzo develops a method, called kernel Granger causality (KGC) to consider spurious relations using kernel functions. Unfortunately, their approach in [7] is bivariate and cannot be extended to multivariate settings, and [8] cannot be applied to large-scale problems due to over-fitting. In parallel with the Marinazzo approach, another method called transfer entropy (TE) was proposed by Schreiber [9], which formulates the causal relationship among time-series from the information theory's perspective. In his 2009 paper, Barnett proved that the TE and GC were equivalent for Gaussian variables [10]. Therefore, the KGC under the Gaussian assumption is equal to TE. These two lines of work (TE and KGC) have received considerable attention.

Several algorithms have been proposed to extended the previously discussed approaches to multivariate time-series [5], [11]–[15]; however, the curse-of-dimensionality of the large-scale problem yet is an open problem [5]. Note that in real-world studies such as climatology, the dimensions of the problem can reach millions of variables [16], so in the presence of short time-series (just a few decades of temperature measurement), the inverse problems quickly become ill-posed due to the presence of redundant variables. In large-scale problems, TE-based methods are costly based on their dependency on estimating the probability distribution functions. Our analysis of a recent TE algorithm based on the Kraskov estimation [12] for a graph with 34 nodes takes 26 hours on an Nvidia GeForce

This research was partly funded by an Ernest J. Del Monte Institute for Neuroscience Award from the Harry T. Mangurian Jr. Foundation. This work was conducted as a Practice Quality Improvement (PQI) project related to American Board of Radiology (ABR) Maintenance of Certificate (MOC) for Prof. Dr. Axel Wismüller. We thank Dr. John Foxe, Institute for Neuroscience, University of Rochester, for his support.

1080-Ti GPU and tens of times more on a CPU, while the cost increases drastically with the number of nodes.

The closest study to our work in terms of the ongoing open problem is the PCMCI method [5]. Peter and Clark Momentary Conditional Independence (PCMCI) is a two-step algorithm, with an statistical independence test in the PC stage, and obtaining causality strengths of the significant links in the MCI stage [5]. Besides, because our method is a nonlinear Granger causality under Gaussian assumption, it is equivalent to the TE [10]–[12] for Gaussian variables. However, our proposed method differs from the above methods in two ways. First, since it is a nonlinear GC method, it is equivalent to the TE method for Gaussian variables, while it is also suitable for large-scale problems. On the other hand, although PCMCI is proposed for large-scale problems, the method we present provides the ability to apply to short time-series on a large scale and provides all steps without examining nonlinear independencies using a partial correlation as is in the PCMCI. We believe our contribution adds numerous capabilities to the existing literature, affirmed through extensive simulations on synthetic and real-world datasets.

II. PRELIMINARIES

In an abuse of notation, we use i, j and Y_i, Y_j to point to the nodes and random variables on the nodes. Scalars are small letters, vectors are in small bold letters, and matrices are in bold capital letters. Calligraphic letters $\mathcal{N}, \mathcal{E}, \mathcal{H}$, and \mathcal{Y} represent the set of nodes, set of weighted edges, feature space (Hilbert subspace), and input space, correspondingly.

a) Causality assumptions: A variable i is said to be a *cause* of a variable j if j can change in response to change in i . Consider a *directional graphical model* $G = (\mathcal{N}, \mathcal{E})$ consisting of the set of nodes $\mathcal{N} = \{i | i = 1, \dots, N\}$ and the set of directed edges \mathcal{E} , whose topology is unknown, but time-series $\{y_{it}\}_{t=1}^T$ are observed per node i , over T time intervals. We assume that *edges are causal*, meaning that every parent i is a direct cause of all its children $\mathcal{C}(i)$, and therefore i is a parent of j , denoting as $i \in \mathcal{P}(j)$. Furthermore, we assume that measurements $\{y_{it}\}_{t=1}^T$ are *causally sufficient*, meaning that there are no unobserved confounders of any of the variables in the graph¹. Our goal is to estimate causal quantities δ_{ij} between each pair of the nodes i and j in the presence of all other nodes $Z = \mathcal{N} \setminus \{i, j\}$, and derive the topology of the graph as $\mathcal{E} = \{\delta_{ij} | i, j \in \mathcal{N}\}$, such that $i \in \mathcal{P}(j)$ if $\delta_{ij} > 0$ else $i \notin \mathcal{P}(j)$. We assume that δ_{ij} is *identifiable*, therefore we can compute it from a purely statistical quantity² [2]. We use *faithfulness* assumption to obtain causal relations among the nodes \mathcal{N} , which allows us to infer *d-separations* in the graph from its dependencies in the distributions $Y_i \perp\!\!\!\perp_G Y_j | Z \Leftarrow Y_i \perp\!\!\!\perp_P Y_j | Z$. The notion

$\perp\!\!\!\perp$ implies independence in the graph G , or the distribution $P(Y_1, \dots, Y_N)$ where Y_i denotes the random variable corresponding to node i . Note that faithfulness is a weak assumption; however, without it, we cannot infer the topology of the causal graph from the observational time-series³.

b) Causal quantities: Linear vector autoregressive model (VARM) is widely adopted framework to infer Granger causality in multivariate time-series, in which a *Granger causality index* δ_{ij} can be defined as a quantification of prediction quality as follows. VARM states that each observed $\mathbf{y}_t := [y_{1t}, \dots, y_{Nt}]^\top$ is a linear combination of the time-lagged versions of the measurements $\{\{y_{i(t-\ell)}\}_{i=1}^N\}_{\ell=1}^L$. Let $\mathbf{A}^\ell \in \mathbb{R}^{N \times N}$ denote the “time-lagged” VARM parameters matrix, with $[\mathbf{A}^\ell]_{ij} = a_{ij}^\ell$, and a_{ij}^ℓ as model coefficients over a lag of ℓ time points. Given the time-lagged multivariate time-series $\{\mathbf{y}_{(t-\ell)}\}_{t=\ell}^L$, where $\mathbf{y}_{(t-\ell)} = [y_{1(t-\ell)}, \dots, y_{N(t-\ell)}]^\top$, the goal is to estimate the model parameter matrices $\{\mathbf{A}^\ell\}_{\ell=1}^L$ in:

$$\mathbf{y}_t = \sum_{\ell=1}^L \mathbf{A}^\ell \mathbf{y}_{(t-\ell)} + \mathbf{e}_t, \quad (1)$$

and minimize the residual errors \mathbf{e}_t using an optimization method, such as ordinary least squares (OLS), the lasso, the ridge, or the elastic-net regressions [18].

Let $\mathbf{E} := [\mathbf{e}_1, \dots, \mathbf{e}_T]$ be the matrix of residuals of VARM for the full system including all nodes \mathcal{N} , and let $\mathbf{E}^{i-} := [\mathbf{e}_1^{i-}, \dots, \mathbf{e}_T^{i-}]$ to be the same matrix when the node i is excluded $\mathcal{N} \setminus \{i\}$ from the VARM, where each \mathbf{e}_t^{i-} is obtained using $\mathbf{y}_t^{i-} := [y_{1t}, \dots, y_{(i-1)t}, y_{(i+1)t}, \dots, y_{Nt}]^\top$ in expression (1). Then, the derivation of error covariance matrices $\Sigma = \text{cov}(\mathbf{E}, \mathbf{E})$ and $\Sigma^{i-} = \text{cov}(\mathbf{E}^{i-}, \mathbf{E}^{i-})$ will be straightforward. Based on Σ of the full VARM and Σ^{i-} of the VARM without i , the degree of information flow from node i to node j can be quantified by $\ln(\Sigma_j^{i-}/\Sigma_j)$, where Σ_j^{i-} and Σ_j denote the diagonal entries of Σ^{i-} and Σ associated to node j , respectively. Consequently, the topology of the causal graph can be inferred using:

$$\delta_{ij} = \begin{cases} \ln(\Sigma_j^{i-}/\Sigma_j), & \ln(\Sigma_j^{i-}/\Sigma_j) > 0 \\ 0, & \ln(\Sigma_j^{i-}/\Sigma_j) \leq 0, \end{cases} \quad (2)$$

as $\mathcal{E} = \{\delta_{ij} | i, j \in \mathcal{N}, i \in \mathcal{P}(j) \text{ if } \delta_{ij} > 0 \text{ else } i \notin \mathcal{P}(j)\}$.

c) Problem Statement: Given nodal measurements in (1), the aim is to solve VARM in the reproducing kernel Hilbert space (RKHS) and then revert the solution (residuals) to the input space to infer the topology in the input space. The benefit is to encompass nonlinearities while reducing dimensionality to address the curse-of-dimensionality.

III. THE PROPOSED METHOD

The proposed method consists of three steps using three off-the-shelf modules, as shown in Fig. 1. The first step involves a

¹Note that we cannot test unconfoundedness, and cannot guarantee that it is satisfied. In other words, no unobserved confounding is unrealistic assumption, and we will always have unobserved confounders [17]

²Identifiability requires to have four assumptions: unconfoundedness, positivity, consistency, and no interference.

³The reason is related to global Markov assumption, where given that \mathbf{P} is Markov with respect to graph G one can use G to infer independencies in \mathbf{P} . However, in case of the unstructured time-series data, the graph topology is priory unknown and we cannot use Markov assumption before obtaining the graph topology.

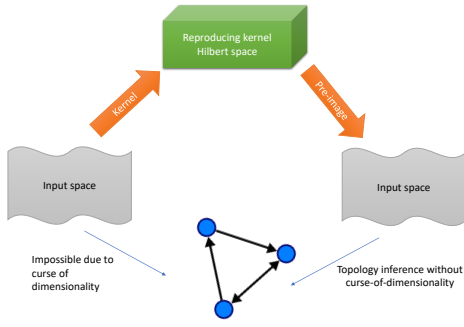


Fig. 1: The proposed method involves three steps: kernel PCA, solving linear vector autoregressive model (VARM) in feature space, and using pre-images to infer topology in input space.

transformation of the input space to the dimensionality reduced (or lifted) feature space using kernel principal component (kernel PCA), the second step involves solving the VARM of (1) in the feature space, and the third step which is the essential part of the method, is to revert the residuals of the VARM in the feature space into the input space using pre-images, and finally infer the topology of the graph \mathcal{E} in the input space.

a) *Dimensionality reduction using kernel PCA:* The aim is to transfer the equation (1) to the space of the principal components with lower dimensions, encompassing nonlinear dependencies. We define RKHS function $\phi(\cdot)$ as $\phi(y_{jt}) : \mathcal{Y} \rightarrow \mathcal{H}$ to project time-series measurements into the feature space. A feature space related to the principal components with the maximum explained variance of data would be desired; therefore, kernel PCA is of interest. Although $P < N$ as the dimension of the feature space is desired, $N \leq P < \infty$ can be also used to lift the space (rather than reducing). Therefore, $\{\phi_p(\cdot)\}_{p=1}^P$ builds the basis of feature space, and $\phi(y_{it}) := [\phi_1(y_{it}), \dots, \phi_P(y_{it})]^\top$. A projection of input data y_{jt} onto the p -th principal component would be desired, such that the major proportion of data variation is explained by a few principal components [19].

b) *Solving the VARM in the reproducing kernel Hilbert space:* Using the RKHS function transformation, the feature space representation of (1) will be:

$$\eta_{pt} = \sum_{\ell=1}^L \sum_{j=1}^N \alpha_{jp}^\ell \phi_p(y_{j(t-\ell)}) + \varepsilon_{pt} \quad (3)$$

where $\{\eta_{pt}\}_{t=1}^T$ for $p = 1, \dots, P$ is the transformed time-series into the feature space, and $\{\varepsilon_{pt}\}_{p=1}^P$ is the error at time slot t in the feature space. The expression (3) is a linear VARM and can be derived with the same approach as discussed in section II. Once the $\{\eta_{pt}\}_{p=1}^P$ in (3) are obtained, the predictions $\{\hat{\eta}_{pt}\}_{p=1}^P$ will be transformed into the input space to infer the graph topology.

c) *Estimating pre-images:* Let $\mathbf{Y} := [\mathbf{y}_1, \dots, \mathbf{y}_T]$, $\gamma(\eta_{pt}) := [\gamma_1(\eta_{pt}), \dots, \gamma_N(\eta_{pt})]^\top$, $\mathbf{\Gamma}_p := [\gamma(\eta_{p1}), \dots, \gamma(\eta_{pT})]^\top$, $\mathbf{\Gamma} := [\mathbf{\Gamma}_1, \dots, \mathbf{\Gamma}_P]$, $\mathbf{\Phi}_j := [\phi(y_{j1}), \dots, \phi(y_{jT})]^\top$,




Fig. 2: The topologies of datasets, which have a different number of nodes $\{2, 3, 5\}$, and different characteristics $\{\text{chaotic (a), linear (d), nonlinear (b, c, e)}\}$. Equations governing these networks are detailed in [25].

and $\mathbf{\Phi} := [\mathbf{\Phi}_1, \dots, \mathbf{\Phi}_P]$, and the notations $\mathbf{Y}^{j-} := [\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_{j-1}, \tilde{\mathbf{y}}_{j+1}, \dots, \tilde{\mathbf{y}}_N]$, $\tilde{\eta}_p := [\eta_{p1}, \dots, \eta_{pT}]^\top$, $\mathbf{H} := [\tilde{\eta}_1, \dots, \tilde{\eta}_P]$ for the sake of simplicity. To transform $\{\hat{\eta}_{pt}\}_{p=1}^P$ into the input space and obtain the pre-image $\gamma(\hat{\eta}_{pt}) : \mathcal{H} \rightarrow \mathcal{Y}$, one needs to estimate $\hat{y}_{jt} = \arg \min_{\hat{y}_{jt} \in \mathcal{Y}} \|\hat{y}_{jt} - \gamma(\phi(\hat{y}_{jt}))\|^2$, which is a nonlinear optimization problem, and is not interesting as a solution to find pre-images [19], [20]. However, interestingly, both the inputs $\{y_{it}\}_{t=1}^T$ and their corresponding mappings are present for $\phi(y_{jt}) : \mathcal{Y} \rightarrow \mathcal{H}$, easing the pre-image problem to a simple learning problem. Therefore, we use the learning problem $\mathcal{L} = \arg \min_{\mathbf{\Gamma}} \|\mathbf{Y} - \mathbf{\Gamma} \mathbf{\Phi}^\top\|^2$ instead, which is adopted from [20] to directly obtain coefficients $\gamma_{jp} := [\mathbf{\Gamma}]_{jp}$, and consequently the predictions $\hat{y}_{jt} = \sum_{p=1}^P \gamma_{jp} \hat{\eta}_{pt}$ will be obtained. Simultaneously, the topology of the graph can be obtained as discussed in section II. The complete algorithm is shown in Algorithm. 1. Note that novel statistical frameworks such as [21]–[24] can be used to formulate the problem in a neural network, which can be as a future work.

Algorithm 1 The proposed method leveraging pre-images

Require: $L, P, \mathbf{Y}, \phi(\cdot)$
Ensure: $\mathcal{E} \leftarrow \{\ln(\frac{\Sigma_j^{i-}}{\Sigma_j})\}_{j=1}^N \leftarrow \ln(\frac{\text{diag}\{\text{cov}(\mathbf{Y}^{i-} - \hat{\mathbf{Y}}^{i-})\}}{\text{diag}\{\text{cov}(\mathbf{Y} - \hat{\mathbf{Y}})\}})$
 $\mathbf{\Gamma} \leftarrow \mathcal{L} \leftarrow \mathbf{Y} \leftarrow \text{normalize}(\mathbf{Y})$
for $j = 1$ to N **do**
 $\mathbf{H} \xleftarrow{\mathbf{\Phi}} \mathbf{Y}$, and $\mathbf{H}^{j-} \xleftarrow{\mathbf{\Phi}} \mathbf{Y}^{j-}$
for $\ell = 1$ to L **do** $\hat{\mathbf{Y}} \xleftarrow{\mathbf{\Gamma}} \hat{\mathbf{H}}$ and $\hat{\mathbf{Y}}^{j-} \xleftarrow{\mathbf{\Gamma}} \hat{\mathbf{H}}^{j-}$
end for
end for

IV. NUMERICAL TESTS

Extensive simulations on a benchmark of the synthetic datasets with known ground-truth and a real-world dataset in a downstream task are presented in this section.




Fig. 3: Five different datasets vs. ROC-AUC of different methods are shown: a benchmark to compare PC-momentary conditional independence (PCMCI, 2019), mutual information (MI, 2018), transfer entropy (TE, 2018), kernel Granger causality (KGC, 2011) with our method. The number of time samples is $T=500$. Synthetic datasets with known ground truth and various topologies and dependencies (Fig. 2) each are replicated for 50 different random implementations. Our proposed method's performance (ROC-AUC) (light blue, first on the right side of each dataset) is significantly better than the competitor methods for all datasets. Boxplots represent [0.25,0.75] quartiles and median.

a) *Tests on synthetic datasets:* We used five various networks with various topologies, each with 50 random realizations, to test our method and some of the state-of-the-art methods, namely PCMCI (2019) [5], mutual information using Kraskov estimation (2018) [12], transfer entropy (TE) using Kraskov method (2018) [11], [12], and kernel Granger causality 2011 [13]. The topology of the datasets are shown in Fig. 2. We used one 2-nodes chaotic master-slave, two 3-nodes with fork and immortality topologies, and two 5-nodes datasets in the linear and nonlinear dependencies mode. Datasets are detailed in [25].

We used datasets to address two challenges: 1) Fig. 3 shows the case of which datasets have enough time samples for different topologies, and 2) Fig. 4 shows when the trend of different algorithms by reducing the number of time samples. ROC-AUC is used to compare the performance in the box plots. It is clearly seen in Fig. 3 and 4 that our algorithm is more robust to the size of the network, nonlinearity in dependencies, and *short time-series*. Note that PCMCI is also proposed for large-scale problems, but in our simulations, we show that our algorithm's robustness for short time-series is better than the PCMCI.

b) *Tests on real datasets:* We examine our method on climatology data of average daily discharges of rivers in the upper Danube basin by three stations located on the Iller at Kempten (IK), the Danube at Dillingen (DD), and the Isar at Lenggries (IL). The data are available through Bavarian Environmental Agency at <https://www.gkd.bayern.de>, and we use the measurements of three years (2017-2019). As shown in Fig. 5, there is a causal relationship $IK \rightarrow DD$ since IK discharges into the DD upstream after a day, while there is no causal relationship between $IL \rightarrow IK$ and $IL \rightarrow DD$ (and *vice versa*). Due to confounder variables such as rainfall or other weather conditions, the statistical analyses are vulnerable to detecting spurious connections. Our method correctly detects a causal relationship between IK and DD while detecting no relationships between IL-IK and IL-DD,

unconfounding spurious variables. In contrast, the Kraskov's TE [9], [12] wrongly detects the spurious connections as the $DD \rightarrow IK$, $DD \rightarrow IL$, $IL \rightarrow IK$, and $IL \rightarrow DD$ altogether, for the best parameters. Our analysis of rivers' underlying dynamics conforms to [26], which shows our method's capability to distinguish between confounding and causal variables.

V. CONCLUSIONS

This paper puts forth a novel method for discovering non-linear dependencies between time-series in multivariate environments. The method aims to solve curse-of-dimensionality while preserving nonlinear dependencies using combination of kernel principal component analysis, linear vector autoregressive model, and estimating pre-images to obtain topology of the graph in the input space. The method's advantage is its simplicity, while the results on the synthetic and real-world datasets show the merit of the method compared to state-of-the-art causality inference methods.

REFERENCES

- [1] M. Turek. Machine Common Sense. June 1, 2021. [Online]. Available: <https://www.darpa.mil/program/machine-common-sense>
- [2] I. Shpitser and J. Pearl, "Complete Identification Methods for the Causal Hierarchy," *Journal of Machine Learning Research*, Vol. 9, pp. 1941–1979, September 2008.
- [3] J. Pearl and D. Mackenzie, *The Book of Why: the New Science of Cause and Effect*. Basic Books, 2018.
- [4] C. Granger, "Some Recent Development in a Concept of Causality," *Journal of Econometrics*, Vol. 39, No. 1-2, pp. 199–211, September 1988.
- [5] J. Runge *et al.*, "Detecting and Quantifying Causal Associations in Large Nonlinear Time Series Datasets," *Science Advances*, Vol. 5, No. 11, pp. 1–15, November 2019.
- [6] Y. Antonacci, L. Astolfi, and L. Faes, "Testing Different Methodologies for Granger Causality Estimation: A Simulation Study," in *European Signal Processing Conference*, January 2021, pp. 940–944.
- [7] N. Ancona, D. Marinazzo, and S. Stramaglia, "Radial Basis Function Approach to Nonlinear Granger Causality of Time Series," *Physical Review E*, Vol. 70, No. 5, pp. 1–7, November 2004.
- [8] D. Marinazzo, M. Pellicoro, and S. Stramaglia, "Kernel-Granger Causality and the Analysis of Dynamical Networks," *Physical review E*, Vol. 77, No. 5, pp. 1–9, May 2008.




Fig. 4: Simulations of various algorithms on different datasets are shown. Each column shows a different dataset (see titles). From left to right: 2-Logistic, 3-Fan out, 3-Fan in, 5-Linear, and 5-Nonlinear. Five different methods (PCMCi, MI, TE, KGC, and the proposed method) are tested for various time samples $T = \{50, 100, 200, 500\}$ to compare the performance on the short time-series. Each dataset has 50 different replications. The confidence intervals of 95 percent are shown as shaded areas around the means. The time samples are increasing from left to right. Our proposed method (light blue, denoted by *Ours*) is significantly more robust than the counterparts to the short time series and for all topologies.




Fig. 5: The average daily discharges of rivers in the upper Danube basin, the Iller at Kempten (IK), the Danube at Dillingen (DD), and the Isar at Lenggries (IL) are nonlinearly related. Only IK causes DD, which is correctly detected by the proposed method.

[9] T. Schreiber, "Measuring Information Transfer," *Physical Review Letters*, Vol. 85, No. 2, pp. 1–4, July 2000.

[10] L. Barnett, A. B. Barrett, and A. Seth, "Granger Causality and Transfer Entropy are Equivalent for Gaussian Variables," *Physical Review Letters*, Vol. 103, No. 23, pp. 1–4, December 2009.

[11] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating Mutual Information," *Physical Review E*, Vol. 69, No. 6, pp. 1–16, June 2004.

[12] L. Novelli *et al.*, "Large-Scale Directed Network Inference with Multivariate Transfer Entropy and Hierarchical Statistical Testing," *MIT Press*, Vol. 3, No. 3, pp. 827–847, July 2019.

[13] D. Marinazzo *et al.*, "Nonlinear Connectivity by Granger Causality," *Neuroimage*, Vol. 58, No. 2, pp. 330–338, September 2011.

[14] J. Runge *et al.*, "Inferring Causation from Time Series in Earth System Sciences," *Nature Communications*, Vol. 10, No. 1, pp. 1–13, June 2019.

[15] M. A. Vosoughi and A. Wismüller, "Large-scale extended Granger

causality for classification of marijuana users from functional MRI," in *Medical Imaging 2021: Biomedical Applications in Molecular, Structural, and Functional Imaging*, Vol. 11600, International Society for Optics and Photonics. SPIE, February 2021, pp. 72 – 84.

[16] G. Camps-Valls, D. Sejdinovic, J. Runge, and M. Reichstein, "A Perspective on Gaussian Processes for Earth Observation," *National Science Review*, Vol. 6, No. 4, pp. 616–618, July 2019.

[17] C. Manski, *Partial Identification of Probability Distributions*. Springer Science & Business Media, 2003.

[18] J. Friedman, T. Hastie, and R. Tibshirani, "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software*, Vol. 33, No. 1, p. 1, August 2010.

[19] B. Schölkopf *et al.*, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT press, 2002.

[20] G. Bakır, J. Weston, and B. Schölkopf, "Learning to Find Pre-Images," pp. 449–456, December 2003.

[21] N. Mohammadi, M. M. Dooley, and M. Cetin, "Finite element reconstruction of stiffness images in mr elastography using statistical physical forward modeling and proximal optimization methods," *arXiv preprint arXiv:2103.14632*, 2021.

[22] N. Mohammadi *et al.*, "Ultrasound elasticity imaging using physics-based models and learning-based plug-and-play priors," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 1165–1169.

[23] N. Mohammadi, M. M. Dooley, and M. Cetin, "A statistical framework for model-based inverse problems in ultrasound elastography," *arXiv preprint arXiv:2010.10729*, 2020.

[24] N. Mohammadi, M. M. Dooley *et al.*, "Mr elasticity reconstruction using statistical physical modeling and explicit data-driven denoising regularizer," *arXiv preprint arXiv:2105.12922*, 2021.

[25] A. Wismüller *et al.*, "Large-scale nonlinear Granger causality for inferring directed dependence from short multivariate time-series data," *Scientific Reports*, Vol. 11, No. 1, pp. 1–11, 2021.

[26] L. Mhalla, V. Chavez-Demoulin, and D. J. Dupuis, "Causal Mechanism of Extreme River Discharges in the Upper Danube Basin Network," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, Vol. 69, No. 4, pp. 741–764, August 2020.