

Causal Inference in Non-linear Time-series using Deep Networks and Knockoff Counterfactuals

Wasim Ahmad
Computer Vision Group
Friedrich Schiller University Jena
Jena, Germany
wasim.ahmad@uni-jena.de

Maha Shadaydeh
Computer Vision Group
Friedrich Schiller University Jena
Jena, Germany
maha.shadaydeh@uni-jena.de

Joachim Denzler
Computer Vision Group
Friedrich Schiller University Jena
German Aerospace Center (DLR)
Institute of Data Science, Jena, Germany

Abstract—Estimating causal relations is vital in understanding the complex interactions in multivariate time series. Non-linear coupling of variables is one of the major challenges in accurate estimation of cause-effect relations. In this paper, we propose to use deep autoregressive networks (DeepAR) in tandem with counterfactual analysis to infer nonlinear causal relations in multivariate time series. We extend the concept of Granger causality using probabilistic forecasting with DeepAR. Since deep networks can neither handle missing input nor out-of-distribution intervention, we propose to use the Knockoffs framework (Barber and Candès, 2015) for generating intervention variables and consequently counterfactual probabilistic forecasting. Knockoff samples are independent of their output given the observed variables and exchangeable with their counterpart variables without changing the underlying distribution of the data. We test our method on synthetic as well as real-world time series datasets. Overall our method outperforms the widely used vector autoregressive Granger causality and PCMCI in detecting nonlinear causal dependency in multivariate time series.

Index Terms—Causal Inference, Deep networks, Knockoffs, Non-linear Time series, Counterfactuals

I. INTRODUCTION

Improving our understanding of cause-effect relationships when analyzing real-world complex systems, e.g., economy, climate- and neuro-sciences, helps in addressing questions like what factors affect our health, the economy, climate, and which actions need to be taken to mitigate the adverse effects and enhance beneficial effects? One of the main challenges in inferring cause-effect relations in such systems is the non-linear dependency among variables, which often leads to inaccurate causal discovery. To mitigate this issue, we propose to use deep recurrent networks in combination with counterfactual analysis to infer nonlinear causal relations in multivariate time series. We apply the concept of Granger causality (GC) [1] using probabilistic forecasting with deep autoregressive (DeepAR) networks [2]. GC makes use of the following two properties: (i) Temporal precedence: a cause is followed by its effect and (ii) Physical influence: changing cause alters the effects. It estimates the extent by which a variable X_i helps in the prediction of another variable X_j . The impact of the inclusion of X_i in the model is estimated by comparing the prediction of X_j with and without X_i . Due to its practical implementation, GC has been intensively used

in various fields i.e. neuroscience, climatology, finance, etc., for discovering causal graphs in time series [1], [3]–[5].

A major shortcoming of Granger causality however is in its assumption of linear dependency in the observational data. Deep networks excel in learning complex and nonlinear relations. In this paper, we use DeepAR networks [2] to model non-linear relations in multivariate time series. DeepAR characterizes conditioning on observed variables and extracts hidden features from the time series which helps in dealing with spurious associations that might occur due to confounders such as periodic cycles and trends. However, deep networks can not handle missing input or out-of-distribution data. To deal with this issue, we propose to use the theoretically well-established Knockoffs framework [6] to generate intervention variables. The Knockoffs framework was originally developed as a variable selection tool with a controllable false discovery rate in a prediction model. Knockoff samples are statistically null-variables, i.e., independent of the model output, and can be swapped with the original variables without changing the underlying distribution of the data. The in-distribution nature of the knockoffs is important when used to generate counterfactuals for causal discovery because the trained deep networks expect test data to be within the same distribution as training data. Moreover, knockoff samples control false discovery rate in causal inference as it holds as low correlation with the candidate variable as possible.

Causal hierarchy operates in three layers (i) Association (ii) Intervention and (iii) Counterfactual [7]. We use counterfactuals where questions like *what if certain thing has been done differently?* Or *is it certain action that causes the change?* are addressed. A counterfactual outcome describes a cause-effect scenario such that if event X_i had not occurred, event X_j would not have happened. Counterfactuals have been used for feature selection, time series explanation, and model interpretability [8], [9]. We generate counterfactual output by substituting one of the observed variables at a time with its knockoff samples and estimate its causal impact by measuring its effect on the forecast error assuming that the used multivariate time series are stationary. Moreover, we assume causal sufficiency, i.e., the set of observed variables includes all possible causes, and there exist no hidden confounders. We compare the causal inference performance of the

proposed method with the widely used vector autoregressive Granger causality (VAR-GC) and PCMCI [10] using synthetic as well as real river discharges datasets. We show that our method outperforms these methods in estimating nonlinear causal links. To highlight the advantage of using the Knockoffs framework, we carry out a comparison using different methods for generating counterfactuals. More specifically, we compare the causal inference results of when generating counterfactuals using out-of-distribution, a constant value of the distribution-mean, and Knockoffs samples.

The remainder of this paper is organized as follows. In Section II of the paper, we cover the related work. Methodological background are presented in Section III. In Section IV, we describe in detail the proposed causal inference method. Experimental results are presented and discussed in Section V. The work is concluded in Section VI of the paper.

II. RELATED WORK

Several methods in the literature have addressed the challenge of causal inference in nonlinear multivariate time series. The method of [11] uses structured multilayer perceptrons and recurrent neural networks combined with sparsity-inducing penalties on the weights. It monitors the weights assigned to each time series when predicting the target time series. The authors of [12] also proposed a similar solution for causal discovery using deep learning where the architecture of the neural network is partially utilized from data to infer causality among variables. They chose to interpret the weights of a neural network under some constraints. The method trains as many neural networks models as the number of variables in the dataset, where one variable is the target and the rest of the variables are inputs to predict the target. Analyzing the weights of the network reveals whether there is a relationship between a variable and the target: if any of the paths between a variable and the target has a weight value as 0, the target does not have a causal link with that particular variable. However, it is challenging to extract any clear structure of the data from the deep networks that can be used for interpretability [13]. Attention-based convolutional neural networks have also been used for causal discovery, where they utilize the internal parameters of the networks to find out time delay in the causal relationship between variables [14].

The PCMCI [10] method estimates causality in multivariate time series using either linear or non-linear conditional independent tests. PCMCI works in two steps, first, it determines the causal parents of each variable and then applies conditional independence test to estimate causality between variables. PCMCI has been implemented in [15] to estimate causal networks in biosphere-atmosphere interaction. PCMCI is also used along with dimensionality reduction to infer causality in ecological time series data [16] where conditional independence tests are applied on the reduced feature space.

Cause-effect variational autoencoders (CEVAE) [17] estimate non-linear causal links in the presence of hidden confounders by training two separate networks, with and without intervention. This method however estimates the causal link of

two variables only. The extended version of CEVAE to time series has been developed by integrating the domain-specific knowledge for inferring nonlinear causality in ecological time series [18]. Additive models are also used to estimate nonlinear interactions in time series, where the past of each series may have an additive non-linear effect that decouples across time series [19], however, it may overlook important nonlinear relation between variables and therefore may fail to properly identify Granger causal links [11]. Counterfactuals are used by [9] to estimate feature importance. They substitute features to obtain counterfactuals and estimate feature importance by finding their impact on the output. Unlike our approach, it does not estimate causal relations in time series and does not use knockoff in-distribution samples to reduce the false discovery rate. Knockoffs have been used to generate counterfactuals to explain the decision of deep image classifiers [20]. To deal with spurious causal links that may result from periodic patterns and trends in time series, the method in [21], [22] proposed to use time-frequency causality analysis. However, these methods use the spectral representation of linear VAR-GC and hence may not be able to deal with nonlinear causal relations.

III. METHODOLOGICAL BACKGROUND

In this section we will briefly introduce the necessary background on DeepAR [2], the fundamental concept of Knockoffs framework [12], as well as VAR-GC.

A. DeepAR

DeepAR [2] is a powerful method mainly used for multivariate non-linear non-stationary time series forecasting. It is based on autoregressive recurrent neural networks, which learn a global model from the history of all the related time-series in the dataset. DeepAR is based on previous work on deep learning for time series data [23] and uses a similar LSTM-based recurrent neural network architecture to the probabilistic forecasting problem. Few key attributes of DeepAR compared to classical approaches and other global forecasting methods are: It extracts hidden features in the form of seasonal patterns and trends in the data, and characterizes conditioning on observed variables and extracted features during inference. Moreover, in this probabilistic forecasting method, lesser hand-crafted feature engineering is needed to capture complex, group-dependent behavior.

The training and forecast processes of the DeepAR is shown in Fig. 1. Let $z_i, i = 1, \dots, N$ be the N -variate time series. Each time series $z_{i,t}, t = 1, \dots, r$ is a realization of length r real-valued discrete stochastic process $Z_i, i = 1, \dots, N$. At each time step t , the inputs to the network are the observed target value at the previous time step $z_{i,t-1}$ as well as the previous network output $\mathbf{h}_{i,t-1}$. The network output $\mathbf{h}_{i,t} = h(\mathbf{h}_{i,t-1}, z_{i,t-1}, x_{i,t}, \Theta)$ then used to compute the parameters $\theta_{i,t} = \theta(\mathbf{h}_{i,t-1}, \Theta)$ of the likelihood $\ell(z|\theta)$, which is used for training the model parameters. For prediction, the past of the time series $z_{i,t}$ is fed in for $t < t_0$, where t_0 represents the first value in the forecast horizon, then in the forecast horizon

for $t \geq t_0$ a sample $\hat{z}_{i,t} \sim \ell(\cdot | \theta_{i,t})$ is drawn and fed back for the next point until the end of the prediction range $t = t_0 + T$ generating one sample trace. This sampling process is repeated iteratively to generate multiple traces of the variables in the forecast horizon representing the joint predicted distribution.

B. Knockoffs

The Knockoffs framework was developed by Barber and Candès in 2015 as a tool for estimating feature importance using conditional independence testing with controllable false discovery rate [24]. Given a set of observed variables $Z = (Z_1, \dots, Z_n)$ with known distribution P_Z , and a predictive model, the knockoffs of the observed variables, defined as $\tilde{Z} = (\tilde{Z}_1, \dots, \tilde{Z}_n)$, are constructed to be in-distribution null-variables and as decorrelated as possible from the original data. They can be swapped with the original variables without changing the underlying distribution of the data. This is accomplished by ensuring that the correlation between the generated knockoffs is the same as the correlation between the original variables [24]. To be null-variables, knockoffs do not contain any information about the target variable. More specifically, knockoff variables should satisfying the pairwise exchangeability condition [24]

$$(Z, \tilde{Z})_{\text{swap}(\mathbb{A})} \stackrel{d}{=} (\tilde{Z}, Z) \quad (1)$$

for any subset $\mathbb{A} \subseteq 1, \dots, n$, here $\stackrel{d}{=}$ represents equal distributions. The $(Z, \tilde{Z})_{\text{swap}(\mathbb{A})}$ is obtained from (Z, \tilde{Z}) by swapping the entries Z_j and \tilde{Z}_j for each $j \in \mathbb{A}$.

The knockoff mechanism can be thought of as generating a probability distribution $P_{\tilde{Z}|Z}(\cdot | z)$ which is the conditional distribution of \tilde{Z} given $Z = z$ chosen such that the obtained joint distribution of (Z, \tilde{Z}) which is equal to

$$P_Z(z)P_{\tilde{Z}|Z}(\tilde{z}|z),$$

is pairwise (z_j, \tilde{z}_j) symmetric and satisfy the exchangeability condition in (1). For Gaussian variables with Gaussian distribution $P_Z = \mathcal{N}_n(\mathbf{0}_n, \Sigma)$, [25] shows that knockoffs $\mathbf{Z}_{i,*}$ can be drawn from the conditional distribution

$$P_{\tilde{Z}|Z}(\cdot | \mathbf{Z}_{i,*}) = \mathcal{N}_n((\mathbf{I}_n - S\Sigma^{-1})\mathbf{Z}_{i,*}, 2S - S\Sigma^{-1}S) \quad (2)$$

for any fixed diagonal matrix S satisfying $0 \leq S \leq 2\Sigma$. This leads to the joint distribution of $(\mathbf{Z}_{i,*}, \tilde{\mathbf{Z}}_{i,*})$ being equal to

$$\mathcal{N}_{2n}\left(\mathbf{0}_{2n}, \begin{pmatrix} \Sigma & \Sigma - S \\ \Sigma - S & \Sigma \end{pmatrix}\right),$$

which satisfies the pairwise exchangeability condition in (1).

C. VAR Granger Causality

GC is based on the idea that cause precedes its effects and can help in their prediction. Here we briefly describe how VAR is used to infer GC. Let $z_i, i = 1, \dots, N$ be the time series of N variables. Each time series $z_{i,t}, t = 1, \dots, r$ is a realization of length r real-valued discrete stationary stochastic process $Z_i, i = 1, \dots, N$. These N time-series can be represented by a p th order VAR of the form

$$\begin{bmatrix} z_{1,t} \\ \vdots \\ z_{N,t} \end{bmatrix} = \sum_{m=1}^p A_m \begin{bmatrix} z_{1,t-m} \\ \vdots \\ z_{N,t-m} \end{bmatrix} + \begin{bmatrix} \epsilon_1(t) \\ \vdots \\ \epsilon_N(t) \end{bmatrix} \quad (3)$$

The residuals $\epsilon_i, i = 1, \dots, N$ form a white noise stationary process with covariance matrix Σ . The model parameters at time lags $m = 1, \dots, p$ comprise the matrix $A_m = [a_{ij}(m)]_{N \times N}$. Let \sum_j be the covariance matrix of the residual ϵ_j associated to z_j using the model in (3), and let \sum_j^{i-} denote the covariance matrix of this residual after missing out the i th row and column in A_m . The VAR-GC of z_i on z_j conditioned on all other variables is defined by [26].

$$\gamma_{i \rightarrow j} = \ln \frac{|\sum_j^{i-}|}{|\sum_j|} \quad (4)$$

IV. CAUSAL EFFECT ESTIMATION USING DEEPAR AND KNOCKOFF COUNTERFACTUALS

A. Granger causality using probabilistic forecasting

Let $z_i, i = 1, \dots, N$ be the N -variate time series. Each time series $z_{i,t}, t = 1, \dots, r$ is a realization of length r real-valued discrete stochastic process $Z_i, i = 1, \dots, N$. Throughout our study, we make use of the following two assumptions.

- Stationarity: $Z_i, i = 1, \dots, N$ is a stationary stochastic process.
- Causal Sufficiency: The set of observed variables $Z_i, i = 1, \dots, N$ includes all of the common causes of pairs in Z .

We train and optimize DeepAR by providing multivariate time series $z_{i,t}, i = 1, \dots, N; t = 1, \dots, r$ as input. For each realization of the time series $z_{i,t}$, we calculate forecast error as mean absolute percentage error (MAPE) using the trained model.

$$\text{MAPE} = \frac{1}{r} \sum_{t=1}^r \frac{|z_{i,t} - \hat{z}_{i,t}|}{|z_{i,t}|} \quad (5)$$

Where $z_{i,t}$ is the actual value of time series i at time step t , and $\hat{z}_{i,t}$ is the predicted value which is defined by the mean of the probability distribution $P(z_{i,t_0:T} | z_{i,1:t_0-1})$ at time step t in the forecast horizon. To know the cause-effect relationship of z_i on z_j , we intervene on variable $z_i, i = 1, \dots, N$ with its knockoffs \tilde{z}_i , and notice the influence on the prediction of a target variable $z_j, j \neq i$ by comparing the obtained counterfactual output with the actual output. In analogy to the definition of VAR-GC of z_i on z_j in (4), we define our metric for estimating the non-linear Granger causality of z_i on z_j conditioned on all other variables by the causal significance score (CSS $_{i \rightarrow j}$) as follows.

$$\text{CSS}_{i \rightarrow j} = \ln \frac{\text{MAPE}_j^i}{\text{MAPE}_j} \quad (6)$$

where MAPE_j and MAPE_j^i denote the mean absolute percentage error when predicting z_j before and after intervention on z_i .

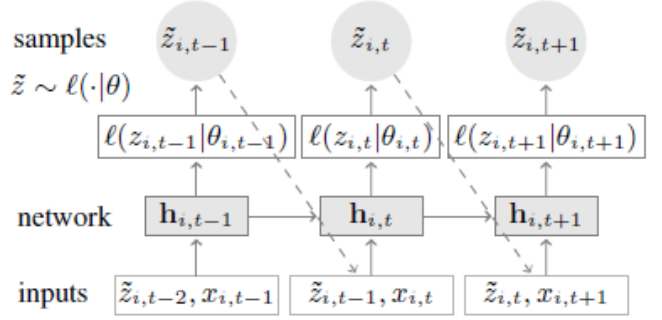



Fig. 1: Architecture of DeepAR network showing training (left) and forecast (right) phases [2].

B. Counterfactual generation

In this work, we investigate multiple ways to generate counterfactuals, i.e. distribution mean, out-of-distribution, and knockoff samples as described below. However, we mainly emphasize the use of knockoffs to generate counterfactuals for causal discovery in time series.

- Distribution mean: Here we replace each value $z_{i,t}$ of the time series z_i with the mean $\bar{z}_i = \frac{1}{r} \sum_{t=1}^r z_{i,t}$.
- Out-of-distribution: In this type of intervention, the candidate variable z_i that belongs to data distribution D_i is replaced with a variable \bar{z} from another data distribution \bar{D} such that $\bar{D} \neq D_i$. Moreover, \bar{z}_i is selected to be as uncorrelated with the original variable z_i as possible.
- Knockoffs: The Knockoffs of the observed variables $Z = (Z_1, \dots, Z_n)$ are constructed to be in-distribution null-variables as decorrelated as possible from the original data. To this end, we implement DeepKnockoffs [12], a framework for sampling approximate knockoffs using deep generative models. The knockoff samples are generated using a mixture of Gaussian models as implemented in [27]. The procedure consists of first sampling the mixture assignment variable from the posterior distribution. The knockoffs are then sampled from the conditional distribution of the knockoffs given the original variables and the sampled mixture assignment such that the exchangeability condition is satisfied. We obtain knockoffs $\tilde{Z}_j, j \subseteq 1, \dots, n$, for all variables in the multivariate time series and substitute one variable at a time Z_j with \tilde{Z}_j for each $j \subseteq 1, \dots, n$ for causal inference.

C. Causal hypothesis testing

We calculate $CSS_{i \rightarrow j}$ using (6) by incorporating $MAPE_j^i$ and $MAPE_j^j$ for each realization $z_{j,t}$ of the stochastic process $Z_j, j = 1, \dots, N$ before and after intervention on $z_{i,t}, i \neq j$. Since we have multiple realizations for each time series, as a result, we get a distribution of values of $CSS_{i \rightarrow j}$ for each pair of time series $z_{i,t}$ and $z_{j,t}$. The null hypothesis $H_0: z_i$ does not Granger cause z_j , is accepted if the mean of the distribution of $CSS_{i \rightarrow j}$ is close to zero. In case the mean of the distribution is significantly different from zero, the alternate hypothesis $H_1:$

z_i Granger causes z_j is accepted. The causal relationship for each pair of time series is estimated iteratively and ultimately the full causal graph of the multivariate time series is extracted.

V. EXPERIMENTS

We conducted experiments on synthetic as well as a real datasets. Experiments are performed on multiple realizations of the time series. For synthetic data, we set $r = 200$ as the length of each realization $z_{i,t}$ of the time series. The prediction length T is set to 14. We trained our model for 150 epochs. The number of layers and cells in each layer of the network is set to 4 and 40 respectively while dropout is 0.05–0.1. We estimate causal link among time series by hypothesis testing on the distribution of $CSS_{i \rightarrow j}$ as explained in Section IV. To quantify the performance of our method, we calculate the two metrics: false positive rate (FPR) and F-score where

$$FPR = \frac{FP}{FP + TN} \quad (7)$$

$$F\text{-score} = \frac{TP}{TP + 0.5(FP + FN)} \quad (8)$$

Here, TP is the number of causal links that exist in the causal graph and also detected by the method; TN is the number of causal links that neither exist in the causal graph nor detected by the method; FP is the number of causal links that do not exist in the causal graph but detected by the method and FN is the number of the causal links that exist in the causal graph but not detected by the method.

A. Synthetic data

To test the proposed method on synthetic data we generate time series $X_i(t), i = 1, \dots, 4$, with seasonal pattern and non-linear relationships among them as follows. This synthetic model simulate a test model of climate-ecosystem interactions as suggested in [15], where X_1 denotes global radiation R_g , X_2 represents air temperature T_{air} while X_3 and X_4 denote gross primary production GPP and ecosystem respiration R_{eco} respectively.

$$\begin{aligned} X_1(t) &= \mathcal{N}(0, 1) + |\cos(2\pi ft)| \\ X_2(t) &= c_1 X_2(t - \tau_1) + c_2 X_1(t - \tau_2) + \eta_1(t) \\ X_3(t) &= c_3 X_1(t - \tau_3) * X_2(t - \tau_4) + \eta_2(t) \\ X_4(t) &= c_4 X_3(t - \tau_5) * \beta^{\frac{X_2(t - \tau_6) - Q}{10}} + \eta_3(t) \end{aligned} \quad (9)$$

The coupling coefficients and time lags are represented as $c_1, c_2, c_3, c_4, \beta$ and $\tau_1, \tau_2, \dots, \tau_5$ respectively. The notation η , termed as ‘‘intrinsic’’ or ‘‘dynamical noise’’, which represents data from uncorrelated, normally distributed noise. We set η_1, η_2, η_3 as Gaussian with 0 mean and variances [0.30, 0.35, 0.25]. The value for the coupling coefficients c_1 to c_4 and β were set to [0.95, 0.80, 0.50, 0.75] and [0.2, 0.4, ..., 1.0], respectively. The lags were given only integer values in the range [0, 10] and Q is set to 10 which is referred to reference temperature in the test data model used in [15]. The seasonal part of $X_1(t)$ is $|\cos(2\pi ft)|$ where f represents the frequency which takes the value 150 and t is the time that varies from 0 to 3000. We keep changing β , the coefficient of non-linear relationship between X_2 and X_4 , in the test model every simulation while rest of the parameters are kept intact. We are interested to see the consistency and causal estimation capability of our method in comparison with PCMCI and VAR-GC when β is increased.

B. Real data

As a real-world application of our method, we carry out causal analysis of average daily discharges of rivers in the upper Danube basin, provided by the Bavarian Environmental Agency ¹. We use measurements of three years (2017-2019) from the Iller at Kempten K_t , the Danube at Dillingen D_t , and the Isar at Lenggries L_t as considered by [28]. While the Iller flows into the Danube upstream of Dillingen with the water from Kempten reaching Dillingen within a day approximately, the Isar reaches the Danube downstream of Dillingen. Based on the stated scenario, we expect a contemporaneous link $K_t \rightarrow D_t$ and no direct causal links between the pairs K_t, L_t and D_t, L_t . Since all variables may be confounded by rainfall or other weather conditions, this choice allows testing the ability of the methods to detect and distinguish directed and bidirectional links.

C. Results

For synthetic data model, we demonstrate results in Fig. 2 which shows the F-score values for PCMCI, VAR-GC, and DeepAR with all three types of substitution methods represented as DeepAR-Knockoffs, DeepAR-OutDist, DeepAR-Mean. The figure illustrates the performance of these methods in response to an increase in the coefficient of non-linearity β between X_2 and X_4 . VAR-GC assumes linearity and hence yields a lower F-score as the dataset contains non-linear relations among variables in the synthetic data model. PCMCI however has better results compare to VAR-GC because of using a non-linear method, Gaussian process distance correlation (GDPC), for testing conditional independence. DeepAR-Knockoffs, yields better results in terms of F-score compare to VAR-GC and PCMCI. DeepAR-OutDist suffers from a high false discovery rate because of using out-of-distribution variables as substitution and hence yield lower F-score throughout the experiment. DeepAR-Mean performs better than DeepAR-OutDist but falls behind DeepAR-Knockoffs. The plot in Fig




Fig. 2: F-score of given methods in response to increase in the coefficient of non-linear relationship β between X_2 and X_4 .




Fig. 3: False positive rate (FPR) of all methods for increasing β between X_2 and X_4 in the synthetic data model.

3 shows FPR in response to increase in the non-linearity coefficient β . It can be seen that VAR-GC yields a high FPR due to its assumption of linearity. Besides that, DeepAR-OutDist also produces a high number of false positives, because outside-distribution samples introduce bias in the model and hence results in detection of false causal links. PCMCI, DeepAR-Mean, and DeepAR-Knockoffs however have better false discovery rates.

For real data, the contemporaneous link $K_t \rightarrow D_t$ is correctly detected by DeepAR-Knockoffs, VAR-GC and PCMCI as shown in Table I. Our method does not report any causal link $D_t \rightarrow K_t$ however VAR-GC and PCMCI wrongly finds $D_t \rightarrow K_t$. PCMCI and VAR-GC detected a causal link $K_t \rightarrow L_t$ and $L_t \rightarrow K_t$ which may be because of the weather acting as a confounder. DeepAR-Knockoffs does not report a causal link between K_t and L_t in any direction as the method has the capability of handling confounders, such as seasonal patterns, in a better way. Our method wrongly finds a causal

¹<https://www.gkd.bayern.de>

link $L_t \rightarrow D_t$, however, VAR-GC and PCMCI do not detect such a causal relationship. Due to extreme events in the river discharges data caused by heavy rainfall, the assumption of stationarity is expected to be violated, which is why VAR-GC and PCMCI report these false links. In comparison DeepAR-Knockoffs correctly discovers the expected links in the river discharges data, except a single false link.

TABLE I: Expected and detected causal links in river discharges data by VAR-GC, PCMCI and DeepAR-Knockoffs

| Causal links | Expected | VAR-GC | PCMCI | DeepAR-Knockoffs |
|-----------------------|----------|--------|-------|------------------|
| $K_t \rightarrow D_t$ | Yes | Yes | Yes | Yes |
| $K_t \rightarrow L_t$ | No | Yes | Yes | No |
| $D_t \rightarrow K_t$ | No | Yes | Yes | No |
| $D_t \rightarrow L_t$ | No | Yes | No | No |
| $L_t \rightarrow K_t$ | No | Yes | Yes | No |
| $L_t \rightarrow D_t$ | No | No | No | Yes |

VI. CONCLUSION

We proposed a novel method for inferring cause-effect relations in non-linear multivariate time series. Since deep networks can not handle out-of-distribution intervention, we proposed to use probabilistic forecasting with DeepAR in combination of knockoffs-based counterfactuals for estimating nonlinear Granger causality. We applied our method on synthetic and real river discharges datasets. Our results confirm that using knockoff samples to generate counterfactuals yields better results in comparison to using the mean- and out-of-distribution as intervention methods. Results also indicate that the proposed method outperforms VAR-GC and PCMCI methods in estimating non-linear relations and dealing with confounders. It should be noted however that the better performance of the proposed method comes along with the higher computational load associated with using deep networks specially for large multivariate time series.

VII. ACKNOWLEDGMENT

This work is funded by the Carl Zeiss Foundation within the scope of the program line “Breakthroughs: Exploring Intelligent Systems” for “Digitization — explore the basics, use applications” and the DFG grant SH 1682/1-1.

REFERENCES

- [1] C. W. Granger, “Some recent development in a concept of causality,” *Journal of econometrics*, vol. 39, no. 1-2, pp. 199–211, 1988.
- [2] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, “Deepar: Probabilistic forecasting with autoregressive recurrent networks,” *International Journal of Forecasting*, vol. 36, no. 3, pp. 1181–1191, 2020.
- [3] C. W. Granger, “Testing for causality: a personal viewpoint,” *Journal of Economic Dynamics and control*, vol. 2, pp. 329–352, 1980.
- [4] R. Goebel, A. Roebroek, D.-S. Kim, and E. Formisano, “Investigating directed cortical interactions in time-resolved fmri data using vector autoregressive modeling and granger causality mapping,” *Magnetic resonance imaging*, vol. 21, no. 10, pp. 1251–1261, 2003.
- [5] A. Brovelli, M. Ding, A. Ledberg, Y. Chen, R. Nakamura, and S. L. Bressler, “Beta oscillations in a large-scale sensorimotor cortical network: directional influences revealed by granger causality,” *Proceedings of the National Academy of Sciences*, vol. 101, no. 26, pp. 9849–9854, 2004.

- [6] R. F. Barber, E. J. Candès, and R. J. Samworth, “Robust inference with knockoffs,” *The Annals of Statistics*, vol. 48, no. 3, pp. 1409–1431, 2020.
- [7] J. Pearl and D. Mackenzie, “The book of why,” Harlow, England: Penguin Books, 2019.
- [8] E. Delaney, D. Greene, and M. T. Keane, “Instance-based counterfactual explanations for time series classification,” in *International Conference on Case-Based Reasoning*. Springer, 2021, pp. 32–47.
- [9] S. Tonekaboni, S. Joshi, D. Duvenaud, and A. Goldenberg, “Explaining time series by counterfactuals,” 2019.
- [10] J. Runge, P. Nowack, M. Kretschmer, S. Flaxman, and D. Sejdinovic, “Detecting and quantifying causal associations in large nonlinear time series datasets,” *Science Advances*, vol. 5, no. 11, p. eaau4996, 2019.
- [11] A. Tank, I. Covert, N. Foti, A. Shojaie, and E. B. Fox, “Neural granger causality,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, p. 1–1, 2021. [Online]. Available: <http://dx.doi.org/10.1109/TPAMI.2021.3065601>
- [12] Y. Romano, M. Sesia, and E. Candès, “Deep knockoffs,” *Journal of the American Statistical Association*, vol. 115, no. 532, pp. 1861–1872, 2020.
- [13] Y. Luo, J. Peng, and J. Ma, “When causal inference meets deep learning,” *Nature Machine Intelligence*, vol. 2, no. 8, pp. 426–427, 2020.
- [14] M. Nauta, D. Bucur, and C. Seifert, “Causal discovery with attention-based convolutional neural networks,” *Machine Learning and Knowledge Extraction*, vol. 1, no. 1, pp. 312–340, 2019.
- [15] C. Krich, J. Runge, D. G. Miralles, M. Migliavacca, O. Perez-Priego, T. El-Madany, A. Carrara, and M. D. Mahecha, “Estimating causal networks in biosphere–atmosphere interaction with the pcmci approach,” *Biogeosciences*, vol. 17, no. 4, pp. 1033–1061, 2020.
- [16] C. Krich, M. Migliavacca, D. G. Miralles, G. Kraemer, T. S. El-Madany, M. Reichstein, J. Runge, and M. D. Mahecha, “Functional convergence of biosphere–atmosphere interactions in response to meteorological conditions,” *Biogeosciences*, vol. 18, no. 7, pp. 2379–2404, 2021.
- [17] C. Louizos, U. Shalit, J. Mooij, D. Sontag, R. Zemel, and M. Welling, “Causal effect inference with deep latent-variable models,” *arXiv preprint arXiv:1705.08821*, 2017.
- [18] V. T. Trifunov, M. Shadaydeh, J. Runge, V. Eyring, M. Reichstein, and J. Denzler, “Nonlinear causal link estimation under hidden confounding with an application to time series anomaly detection,” in *German Conference on Pattern Recognition*. Springer, 2019, pp. 261–273.
- [19] K. Vantas, E. Sidiropoulos, and C. Evangelides, “Estimating rainfall erosivity from daily precipitation using generalized additive models,” in *Environmental Sciences Proceedings*, vol. 2, no. 1. Multidisciplinary Digital Publishing Institute, 2021, p. 21.
- [20] O.-I. Popescu, M. Shadaydeh, and J. Denzler, “Counterfactual generation with knockoffs,” *arXiv preprint arXiv:2102.00951*, 2021.
- [21] M. Shadaydeh, Y. G. Garcia, M. D. Mahecha, M. Reichstein, and J. Denzler, “Causality analysis of ecological time series: a time-frequency approach,” in *8th International Workshop on Climate Informatics: CI 2018*, 2018, pp. 111–114.
- [22] M. Shadaydeh, J. Denzler, Y. G. García, and M. Mahecha, “Time-frequency causal inference uncovers anomalous events in environmental systems,” in *German Conference on Pattern Recognition*. Springer, 2019, pp. 499–512.
- [23] A. Graves, “Generating sequences with recurrent neural networks,” *arXiv preprint arXiv:1308.0850*, 2013.
- [24] R. F. Barber and E. J. Candès, “Controlling the false discovery rate via knockoffs,” *The Annals of Statistics*, vol. 43, no. 5, Oct 2015. [Online]. Available: <http://dx.doi.org/10.1214/15-AOS1337>
- [25] E. Candès, Y. Fan, L. Janson, and J. Lv, “Panning for gold: Model-x knockoffs for high-dimensional controlled variable selection,” 2017.
- [26] J. Geweke, “Measurement of linear dependence and feedback between multiple time series,” *Journal of the American statistical association*, vol. 77, no. 378, pp. 304–313, 1982.
- [27] J. R. Gimenez, A. Ghorbani, and J. Zou, “Knockoffs for the mass: new feature importance statistics with false discovery guarantees,” in *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 2125–2133.
- [28] A. Gerhardus and J. Runge, “High-recall causal discovery for autocorrelated time series with latent confounders,” *34th Conference on Neural Information Processing Systems*, 2020.