

On Counterfactual Interventions in Vector Autoregressive Models

Kurt Butler, Marija Iloska, and Petar M. Djurić

Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY 11794, USA
kurt.butler@stonybrook.edu; marija.iloska@stonybrook.edu; petar.djuric@stonybrook.edu

Abstract—Counterfactual reasoning allows us to explore hypothetical scenarios in order to explain the impacts of our decisions. However, addressing such inquiries is impossible without establishing the appropriate mathematical framework. In this work, we introduce the problem of counterfactual reasoning in the context of vector autoregressive (VAR) processes. We also formulate the inference of a causal model as a joint regression task where for inference we use both data with and without interventions. After learning the model, we exploit linearity of the VAR model to make exact predictions about the effects of counterfactual interventions. Furthermore, we quantify the total causal effects of past counterfactual interventions.

Index Terms—causal models, counterfactuals, interventions, time series, vector autoregressive models

I. INTRODUCTION

In many situations, we may be interested in “what-if” questions. As decisions in real life are rarely made with perfect knowledge of their effects, it is only natural that later one may question the optimality of their choices and ask how things would be if a different decision had been made. This process of reasoning about decisions and actions that contradict what actually occurred pertains to problems of counterfactual reasoning.

A counterfactual event is an event which does not agree with a particular outcome that was observed in a given experiment [1, p. 29]. For example, a student applying to a college might be rejected due to a low score on an entrance exam. A counterfactual question could then be “would this student have been accepted if their score were higher?” In asking this question, we are imagining a hypothetical scenario in which we modify the score of this student, and keep all other variables (i.e., the scores of all the other applicants) constant. In general, counterfactual questions may help us understand whether our actions are effective.

For the problem of detecting the presence of causal relationships and learning the graphical structure of a causal system, a large variety of approaches have been proposed, often arising from different fundamental principles [2]–[5]. Many of these methods attempt to rely on entirely *observational* data, where the experimenter cannot *intervene* upon the system under study to probe into its cause and effect relationships. While it is sometimes possible to detect causalities without intervention, observational approaches often require strong

assumptions that are difficult to verify empirically or are otherwise restrictive. For instance, identification of a linear causal model from only observations is ill-posed if the noises are Gaussian [6, pp.50–51]. Methods based on predictability, such as Granger causality [7], can deliver false positives when not all relevant variables are included in the model. Other methods may require strong assumptions; for example, convergent cross mapping [8] requires the existence of a dynamical attractor, which is difficult to test on small data sets [9]. In comparison, when it is possible to perform interventions, causal inference becomes considerably simpler, with several standard algorithms [3], [10]. Under ideal conditions, when experiments can be carried out perfectly and there are no hidden processes, all causal relationships in system of N variables can be detected with N single-variable interventions [11].

To answer counterfactual questions, more information than just the direction of causalities is required: one must also have a functional model of how the effect is determined by all of its causes [1]. For the college student example, we would need to know the entrance exam scores of the other applicants to determine the functional model for acceptance and subsequently if the given student would be accepted, as their acceptance would depend on all of the scores.

In this work, we consider the counterfactual prediction problem: Given a particular realization of a causal system, i.e., a multivariate time series and a causal model that explains it, can we predict what we would have observed if a previous, hypothetical intervention occurred?

We present several contributions relevant to studying counterfactual events with time series. Firstly, We introduce an approach to learning the causal structure of a system by leveraging multiple data sets with different interventions. We apply this approach to the setting of vector autoregressive (VAR) models. After learning the causal model, we address the problem of making counterfactual predictions. Due to linearity of the VAR model, we are able to make exact predictions without needing to estimate unobserved variables. Finally, we validate our approach using several examples.

II. PROBLEM FORMULATION

We consider the problem of modeling causal relationships from a multivariate time series $\mathbf{x}_t = [x_{1,t}, \dots, x_{D,t}]^\top$. For each node i and time instant t , we assume that the value of $x_{i,t}$ is determined by a function of the values of the system at

This work was supported by the National Science Foundation under Awards 2021002 and 2212506.

previous time steps, $\mathbf{x}_{t-1}, \dots, \mathbf{x}_{t-Q}$, up to some order Q , as well as a separate independent noise term $w_{i,t}$. As a simple example, we examine the vector autoregressive (VAR) model:

$$\mathbf{x}_t := \sum_{q=1}^Q \mathbf{B}_q \mathbf{x}_{t-q} + \mathbf{w}_t, \quad (1)$$

where the \mathbf{B}_q are matrices of linear coefficients, and $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$, for some given noise covariance \mathbf{C} . The notation \rightarrow expresses the direction of causality, where we consider (1) as a causal generative process.

According to (1), the components of \mathbf{x}_{t-1} , \mathbf{x}_{t-2} , and \mathbf{x}_{t-Q} exert causal influences on \mathbf{x}_t . It is traditional in causality to express these relationships in terms of a graph, although there are several options as to how this can be done. In Figure 1, we present one approach that treats each time instance as a separate vertex, with the purpose of understanding the temporal dependencies of the VAR model, which will later simplify the calculation and prediction of counterfactuals.

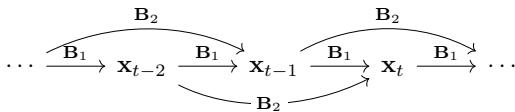


Fig. 1: A causal graph obtained from *unfolding* the causal relationships in (1) to show their dependencies in time, in the $Q = 2$ case. Each vertex represents a random vector, and each edge represents the presence of a cause-effect relationship.

To answer counterfactual queries about a system, we must first have a **structural causal model** (SCM) of the system [1]. A SCM is a hypothesis about the generative process that created a data set. The SCM represents each variable in the system as a function of its parents, a subset of variables directly influencing a given variable, along with an exogenous noise variable. In this way, the VAR model in (1) can be interpreted as a SCM after we declare that the equation explicitly describes how the value of \mathbf{x}_t is determined.

Asserting that a VAR model is a causal model is a stronger statement than asserting it is a statistical model. A statistical model needs only to generate the correct joint probability distribution over the family of variables, but a causal model is a hypothesis about the generative process as it exists in reality. In the same way that multiple generative processes can represent the same probability distribution, there might be multiple SCMs that are compatible with a given data set. Additionally, different causal models will produce different answers to counterfactual questions, so care should be taken to learn the most accurate causal model before discussing counterfactuals. In the next section, we introduce interventions and how SCMs respond to them.

A. Interventions

A major feature of causal models is that they are modularized, meaning that we may think of them as a set of

individual components which can be changed or modified. In our example, each variable $x_{i,t}$ receives a value according to the formula

$$x_{i,t} := \sum_{q=1}^Q \mathbf{b}_{q,i}^\top \mathbf{x}_{t-q} + w_{i,t}, \quad (2)$$

where $\mathbf{b}_{q,i}^\top$ is the i -th row of \mathbf{B}_q , and $w_{i,t}$ is the i -th element of \mathbf{w}_t in (1). An **intervention at node i and time t** is a modification to the causal model which replaces the function that assigns a value to $x_{i,t}$ with a new one. For example, we may intervene by replacing the function in (2) with a new one:

$$x_{i,t} := \sum_{q=1}^Q \tilde{\mathbf{b}}_{q,i}^\top \mathbf{x}_{t-q} + \tilde{\sigma} w_{i,t} + \tilde{u}_{i,t}, \quad (3)$$

where $\tilde{\mathbf{b}}_{q,i}^\top$ are modified linear coefficients, $\tilde{u}_{i,t}$ is a signal injected by the experimenter, and $\tilde{\sigma} \in [0, 1]$ is a parameter that may be used to control the level of noise. Under ideal conditions, the experimenter's intervention perfectly controls the value of $x_{i,t}$ by tuning the parameters in (3) to achieve a desired effect.

In this paper, we assume that all interventions are able to perfectly control the value of the node intervened upon, so that $x_{i,t}$ satisfies (3) whenever an intervention on $x_{i,t}$ is being performed. While intervening on a node i , it is assumed that the functions governing other nodes in the system remain unchanged, which is a form of modularity of the system. In real life, there might be some uncertainty in the values of $\tilde{\mathbf{b}}_{q,i}^\top$ and $\tilde{\sigma}$ when the intervention is performed, due to imperfect experimental conditions. An intervention can be applied for a period of time, during which we can actively observe the behavior of the rest of the system, i.e., the unperturbed nodes. We illustrate this process visually in Figure 2.

Data that we record while performing an intervention are called *interventional*, contrasting with *observational* data that we record while not altering the system under study. While performing an experiment with multiple interventions, we allocate recorded data into separate data sets according to the type of intervention being applied (cf. Figure 2). In Section III-A, we introduce a method to leverage multiple data sets to improve estimation of \mathbf{B}_q over purely observational data. The caveat is that this approach is only possible when obtaining interventional data is feasible.

B. The problem of predicting counterfactuals

Having introduced SCMs and interventions, we may now formulate **counterfactuals** in our framework. In short, a counterfactual is a hypothetical intervention. It is not an intervention that occurred in reality, but rather one whose behavior we can predict, having already observed a realization of the system.

Now that the necessary causal background has been introduced, we can state the primary problem of this work. Consider an experiment in which we observe a multivariate signal \mathbf{x}_t until a certain time T , and suppose that we have assumed a particular SCM for \mathbf{x}_t . Now assume we are interested in

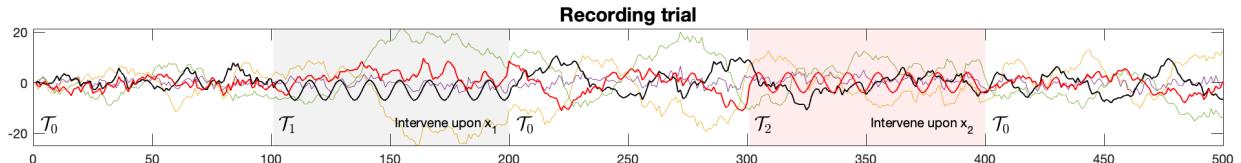


Fig. 2: Interventions applied to a multivariate time series. The underlying causal model is a VAR model, (1), with $Q = 2$ and $D = 5$. Two separate interventions were performed during the experiments in which we drove a node in the system with a sinusoidal stimulus. For $t = 101, \dots, 200$, $x_{1,t}$ (black curve) was forced to take on the value $4 \sin(t/2)$, and for $t = 301, \dots, 400$, $x_{2,t}$ (red curve) was set to $4 \sin(t/2)$. After recording, we partitioned the trial into three data sets, $\mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_2$, according to the intervention periods. Explicitly, $\mathcal{T}_1 = \{101, \dots, 200\}$, $\mathcal{T}_2 = \{301, \dots, 400\}$, and $\mathcal{T}_0 = \{1, \dots, 500\} \setminus (\mathcal{T}_1 \cup \mathcal{T}_2)$

a hypothetical intervention that would have occurred in the past, and want to predict what we would have observed in this hypothetical scenario. That is,

Given: A particular realization $\{\mathbf{x}_t | t = 1, \dots, T\}$,
and an SCM for \mathbf{x}_t ,
(4)

Predict: $\tilde{\mathbf{x}}_t$,
(5)

where $\tilde{\mathbf{x}}_t$ are the signals in the case that we performed a hypothetical intervention upon $x_{i,t}$ for some t in the range $1 < T_1 \leq t \leq T_2 < T$, and for a fixed i . Here we make an important assumption that all conditions of the experiment, such as the particular realization of the noise process \mathbf{w}_t , remains the same in both cases. This is necessary, because in considering the counterfactual problem, asserting that all other conditions were the same implies that the noise too, should have the same realization.

In general models, the values of \mathbf{w}_t need to be estimated before counterfactual analysis can be performed. However, this estimation is straightforward in any additive noise model since \mathbf{w}_t enters the formulas additively. In our approach, we exploit linearity of the system to circumvent the need to estimate \mathbf{w}_t at all because the influence of the noise will cancel out.

III. PROPOSED SOLUTION

Before making counterfactual predictions, it is essential to first learn the causal structure of the system. To learn the parameters of the causal model in a way that leverages both observational and interventional data, we propose a joint regression task that allows us to infer the parameters of each module shared across datasets. We make the tacit assumption that these interventions are ideal, and that the parameters of the model are unchanged when they are not related to the intervention. We now introduce some notation. If at time t we perform an intervention to control the value of $x_{i,t}$, then the time point t should be allocated to a set of time indices \mathcal{T}_i . If no intervention was performed, then the index t of \mathbf{x}_t is allocated to the index set \mathcal{T}_0 . We assume that only one node can be intervened upon per time step t , so $\mathcal{T}_i \cap \mathcal{T}_j = \emptyset$ whenever $i \neq j$. We define

$$\mathcal{T}_{-i} = \left(\bigcup_{j=0}^D \mathcal{T}_j \right) \setminus \mathcal{T}_i \quad (6)$$

to be the set of time indices of x_i , where x_i was not subject to intervention. It is permissible that some index sets may be empty, $\mathcal{T}_j = \emptyset$. In principle, these sets could be partitions of a single recording trial, or obtained from multiple recording trials, with no major modification to our approach.

A. Causal structure learning

For clarity in presentation, we will first describe the case of one lag ($Q = 1$), and from there we extend to the multiple lag case.

a) **SINGLE LAG:** The goal is to learn the linear coefficient matrix \mathbf{B} of the true causal model jointly across datasets by minimizing the following objective,

$$\min_{\mathbf{B}} \sum_{t \in \mathcal{T}_0} \|\mathbf{x}_t - \mathbf{B}\mathbf{x}_{t-1}\|_2^2 + \sum_{i=1}^D \sum_{t \in \mathcal{T}_i} \|\mathbf{x}_t^{[-i]} - \mathbf{B}^{[-i]}\mathbf{x}_{t-1}\|_2^2 \quad (7)$$

where the operation $\mathbf{A}^{[-i]}$ removes row i from the matrix \mathbf{A} .

Using the identity $\|\mathbf{x}\|_2^2 = x_1^2 + \dots + x_D^2$, we can simplify this optimization task considerably. Recall that whenever we intervene on a particular variable, all other linear coefficients remain unchanged. Thus, the function that assigns a value to $x_{i,t}$ is given by

$$x_{i,t} := \mathbf{b}_i^\top \mathbf{x}_{t-1} + \mathbf{w}_t$$

whenever $t \notin \mathcal{T}_i$, or equivalently, when $t \in \mathcal{T}_{-i}$. These considerations lead us to propose the following objective, which is equivalent to (7):

$$\min_{\mathbf{B}} \sum_{i=1}^D \sum_{t \in \mathcal{T}_{-i}} \|x_{i,t} - \mathbf{b}_i^\top \mathbf{x}_{t-1}\|_2^2, \quad (8)$$

where, again, \mathbf{b}_i^\top is the i -th row of \mathbf{B} .

The revised objective in (8) has multiple advantages. Firstly, it can be solved for each variable i separately as a distinct subproblem. Secondly, it makes it clear that minimum mean squared estimate of \mathbf{B} can be obtained using classical methods. If we assume some level of sparsity in \mathbf{B} , then we can use LASSO [12], or ℓ_1 regularization, to get a sparse estimate:

$$\min_{\mathbf{B}} \sum_{i=1}^D \sum_{t \in \mathcal{T}_{-i}} \|x_{i,t} - \mathbf{b}_i^\top \mathbf{x}_{t-1}\|_2^2 + \lambda \|\mathbf{B}\|_1, \quad (9)$$

where λ is a regularization parameter.

b) *MULTIPLE LAGS*: To generalize to the case in which $Q > 1$, we only need to modify the approach slightly. Consider the model in (1). If we define

$$\Theta = [\mathbf{B}_1 \ \cdots \ \mathbf{B}_Q], \quad (10)$$

and

$$\bar{\mathbf{x}}_t = [\mathbf{x}_{t-1} \ \cdots \ \mathbf{x}_{t-Q}], \quad (11)$$

then we can express (1) as

$$\mathbf{x}_t := \Theta \bar{\mathbf{x}}_t + \mathbf{w}_t. \quad (12)$$

Now we can formulate a similar optimization task as before, giving us the multiple lag version of (7):

$$\min_{\Theta} \sum_{t \in \mathcal{T}_0} \|\mathbf{x}_t - \Theta \bar{\mathbf{x}}_t\|_2^2 + \sum_{i=1}^D \sum_{t \in \mathcal{T}_i} \|\mathbf{x}_t^{[-i]} - \Theta^{[-i]} \bar{\mathbf{x}}_t\|_2^2,$$

which can again be split into tractable sub-problems. Considering sparsity, the multiple lag version of (9) is given to be

$$\min_{\Theta} \sum_{i=1}^D \sum_{t \in \mathcal{T}_i} \|\mathbf{x}_{i,t} - \theta_i^\top \bar{\mathbf{x}}_t\|_2^2 + \lambda \|\Theta\|_1, \quad (13)$$

where θ_i^\top is the i -th row of Θ . The use of sparsity constraints becomes increasingly relevant when considering multiple lags, since strongly autocorrelated signals can have multiple distinct models with similar predictive power [13].

B. Estimation of total causal effects

We now consider the problem of estimating the effects of a past counterfactual intervention. Suppose that we have observed \mathbf{x}_t up until some time T . For some $t < T$, consider an intervention that modifies $x_{i,t}$ to take on a new value $x_{i,t}^*$. If we define $\Delta x_{i,t} = x_{i,t}^* - x_{i,t}$, then we can reason about the hypothetical intervention as an additive change in $x_{i,t}$:

$$x_{i,t} \longrightarrow x_{i,t} + \Delta x_{i,t}.$$

Since the VAR model is linear, the additive change in $x_{i,t}$ induces an additive change in all variables “downstream” from it in the graph. For example, if we move one step into the future, then \mathbf{x}_{t+1} changes according to

$$\mathbf{x}_{t+1} \longrightarrow \mathbf{x}_{t+1} + \mathbf{B}_1 \mathbf{e}_i \Delta x_{i,t},$$

where \mathbf{e}_i is the i -th unit vector in \mathbb{R}^D . We observe that the linear model coefficients in \mathbf{B}_1 measure the sensitivity of \mathbf{x}_{t+1} to changes in \mathbf{x}_t , and can be interpreted as a measure of the strength of the causal relationship [14].

If we perform an intervention at time t and attempt to predict its effect at time $t+k$, we need to propagate the changes through the graph in Figure 1, that is, through each intermediate time step. When $Q > 1$, the path in the graph from \mathbf{x}_t to \mathbf{x}_{t+k} is not unique, and the total sensitivity of \mathbf{x}_{t+k} to \mathbf{x}_t is given by a sum over all paths in the graph [15], [16]. Since the VAR model is time-invariant, the total sensitivity is a matrix \mathbf{T}_k which only depends upon k . We refer to this matrix as a **total causal effect** matrix.

To compute the total causal effect matrix for the VAR model, we can exploit recursion. Suppose that $k > Q$. Any path from \mathbf{x}_t to \mathbf{x}_{t+k} must go through a node between \mathbf{x}_{t+k-Q} and \mathbf{x}_{t+k-1} . As a result, the total causal effect decomposes according to the same rule:

$$\mathbf{T}_k = \mathbf{B}_1 \mathbf{T}_{k-1} + \mathbf{B}_2 \mathbf{T}_{k-2} + \cdots + \mathbf{B}_Q \mathbf{T}_{k-Q}. \quad (14)$$

If $k \leq Q$, then a similar rule holds, but we define $\mathbf{T}_0 = \mathbf{I}$ to be the identity matrix, and $\mathbf{T}_k = 0$ for any $k < 0$, since causation can only move forward in time.

IV. EXPERIMENTS

In this section, we demonstrate our approach to modeling counterfactuals with two examples. First, we demonstrate how joint regression can be leveraged to more accurately learn the causal model. Second, we explore an example in which we study the effects of a past intervention on future, forecasted events. The presented results are of single outcomes of the experiments.

A. Learning causal models with joint regression

We simulated a time-series of length $T = 390$, and system dimension $D = 5$ with lag $Q = 2$. Each element $b_{q,i,j}$ in the matrix coefficients \mathbf{B}_q was generated as $b_{q,i,j} \sim \mathcal{U}(-0.5, 0.5)$, for $q = 1, 2$. We set about 30% sparsity for each matrix. The dynamics were perturbed with Gaussian noise $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$, where the covariance \mathbf{C} is a Toeplitz matrix with $C_{ii} = 1$, $C_{ij} = 0.5$ if $|i-j| = 1$, and 0 otherwise. We intervened on the node $x_{1,\tau}$ within the interval $\tau \in \{101, 170\}$ and the node $x_{2,\tau}$ in the interval $\tau \in \{201, 270\}$. For this example, we let the intervened nodes be driven by a sine wave, specifically $x_{i,\tau} := \tilde{x}_{i,t}$, where $\tilde{x}_{i,t} = 10 \sin(\frac{\tau}{2})$, for $i = 1, 2$ in their respective time intervals. The rest of the nodes ($i = 3, 4, 5$) were left unperturbed. We learned the \mathbf{B} coefficients in two ways: i) by using only the observational (unintervened data \mathcal{T}_0), and ii) using intervened data as well - the proposed method. To insure a fair comparison, we opted to use the same number

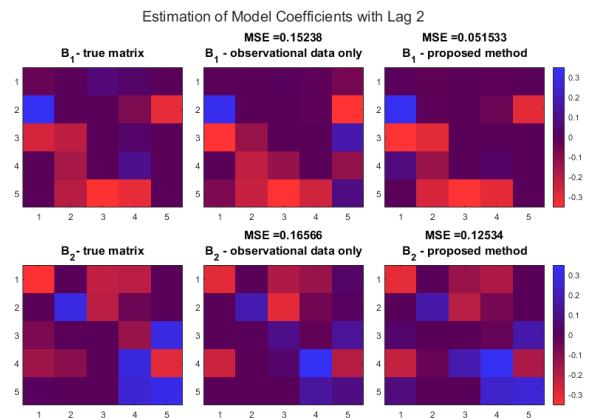


Fig. 3: Learning the causal model of a small system. The performance is quantified by the MSE of the estimated entries of \mathbf{B}_q , defined by $\text{MSE} = \frac{1}{D^2} \sum_{i,j} (b_{q,i,j} - \hat{b}_{q,i,j})^2$.

of data points $T_0 = 250$ for learning in both i) and ii), where in the case of i), 140 of those points are the interventional dataset. With the current setup, the proposed method is at a slight disadvantage in regards to amount of data used, since we lose 140 points for learning the coefficients $\mathbf{b}_{q,i}$ for $i = 1, 2$, $q = 1, 2$. However, those same 140 points of intervention allow us to better estimate the coefficients $\mathbf{b}_{q,i}$ for $q = 3, 4, 5$.

In Figure 3, we show the true matrices \mathbf{B}_q , for $q = 1, 2$, their estimate using \mathcal{T}_0 only, and their estimate with the proposed method, i.e., using interventional data as well. We can see there is a drop in MSE when using the proposed method, which is also reflected visually in the heat maps. The effect becomes even more prominent when learning larger systems.

B. Predicting the effects of a counterfactual intervention

In this experiment, we consider the problem of predicting the effects of a counterfactual intervention in the past. Consider a multivariate time series of length $T = 100$ and dimension $D = 2$. We assume that the causal model is given by a VAR model, (1), with $Q = 2$, whose parameters are known from prior investigation. Let us consider a counterfactual situation in which, at times $t = 240$ until the end of the experiment, the experimenter intervenes upon $x_{1,t}$ by driving the signal with a new signal, $u(t) \sim \mathcal{N}(1, 1)$.

In Figure 4, we predict the effects of the counterfactual intervention on $x_{1,t}$, as perceived by another variable $x_{2,t}$. We plot the signal $x_{2,t}$ that we observed originally alongside its variant in an “alternate universe,” in which the counterfactual intervention is performed. For times $t < 40$, both the observed time series and the counterfactual one coincide, as no intervention has not occurred yet. For $40 \leq t$, we see the counterfactual intervention begin to make the two time series diverge due to the effects of the intervention. Again, the model is able to well anticipate the counterfactual universe.

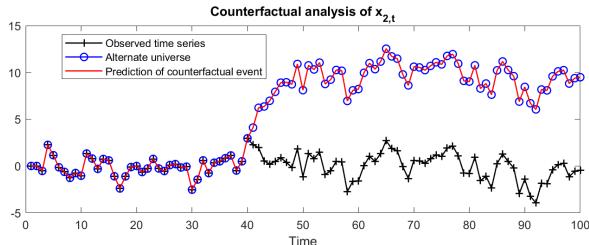


Fig. 4: Predicting the effects of a past counterfactual intervention. We consider a counterfactual (hypothetical) intervention in which, starting at time 40. We visualize three curves: the original observed signal $x_{2,t}$ (black), the signal which would have been observed if the hypothetical intervention upon $x_{1,t}$ was performed (blue), and our prediction of what would have been observed in the counterfactual situation (red). The black and blue curves were generated precisely by simulation of the system, and the red curve is produced by making predictions using the observed realization.

V. CONCLUSION

In this work, we studied the problems of counterfactual reasoning and causal inference in the context of vector autoregressive models. By considering the system as modular and assigning separate functions to each variable, we can approach regression and model learning in a nonlinear manner. Learning of the causal model is a necessary requirement to reason about counterfactual interventions and other what-if scenarios. The analysis of counterfactuals for nonlinear settings is analogous to what we have presented here, although additional care must be taken to estimate unobserved variables like the noises. Future work involves understanding the interventions for improved learning of the studied system and combining counterfactual analysis with interventions to achieve desired outcomes. Other directions include extending the counterfactual analysis to nonlinear models and models where the functional relationships are unknown and need to be estimated.

REFERENCES

- [1] J. Pearl, *Causality*. Cambridge University Press, 2009.
- [2] D. B. Rubin, “Estimating causal effects of treatments in randomized and nonrandomized studies,” *Journal of Educational Psychology*, vol. 66, no. 5, p. 688, 1974.
- [3] P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman, *Causation, Prediction, and Search*. MIT Press, 2000.
- [4] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing, “DAGs with NO TEARS: Continuous optimization for structure learning,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [5] T. Schreiber, “Measuring information transfer,” *Physical Review Letters*, vol. 85, no. 2, p. 461, 2000.
- [6] J. Peters, D. Janzing, and B. Schölkopf, *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- [7] C. W. Granger, “Investigating causal relations by econometric models and cross-spectral methods,” *Econometrica: Journal of the Econometric Society*, pp. 424–438, 1969.
- [8] G. Sugihara, R. May, H. Ye, C.-h. Hsieh, E. Deyle, M. Fogarty, and S. Munch, “Detecting causality in complex ecosystems,” *Science*, vol. 338, no. 6106, pp. 496–500, 2012.
- [9] K. Butler, G. Feng, and P. M. Djurić, “On causal discovery with convergent cross mapping,” *IEEE Transactions on Signal Processing*, vol. 71, pp. 2595–2607, 2023.
- [10] J. M. Ogario, P. Spirtes, and J. Ramsey, “A hybrid causal search algorithm for latent variable models,” in *Conference on Probabilistic Graphical Models*, pp. 368–379, PMLR, 2016.
- [11] F. Eberhardt, C. Glymour, and R. Scheines, “N-1 experiments suffice to determine the causal relations among N variables,” *Innovations in Machine Learning: Theory and Applications*, pp. 97–112, 2006.
- [12] R. Tibshirani, “Regression shrinkage and selection via the Lasso,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 58, no. 1, pp. 267–288, 1996.
- [13] C. Cui, P. Banelli, and P. M. Djurić, “Topology inference of directed graphs by gaussian processes with sparsity constraints,” *IEEE Transactions on Signal Processing*, 2024.
- [14] K. Butler, G. Feng, and P. M. Djurić, “A differential measure of the strength of causation,” *IEEE Signal Processing Letters*, vol. 29, pp. 2208–2212, 2022.
- [15] Y. Liu, C. Cui, D. Waxman, K. Butler, and P. M. Djurić, “Detecting confounders in multivariate time series using strength of causation,” in *2023 31st European Signal Processing Conference (EUSIPCO)*, pp. 1400–1404, IEEE, 2023.
- [16] S. Salehkaleybar, A. Ghassami, N. Kiyavash, and K. Zhang, “Learning linear non-Gaussian causal models in the presence of latent variables,” *Journal of Machine Learning Research*, vol. 21, no. 39, pp. 1–24, 2020.