

Optimal model-free prediction from multivariate time seriesJakob Runge,^{1,2} Reik V. Donner,¹ and Jürgen Kurths^{1,2,3,4}¹*Potsdam Institute for Climate Impact Research, P. O. Box 60 12 03, 14412 Potsdam, Germany*²*Department of Physics, Humboldt University, Newtonstr. 15, 12489 Berlin, Germany*³*Institute for Complex Systems and Mathematical Biology, University of Aberdeen, Aberdeen AB24 3UE, United Kingdom*⁴*Department of Control Theory, Nizhny Novgorod State University, Gagarin Avenue 23, 606950 Nizhny Novgorod, Russia*

(Received 9 July 2014; revised manuscript received 18 February 2015; published 13 May 2015)

Forecasting a time series from multivariate predictors constitutes a challenging problem, especially using model-free approaches. Most techniques, such as nearest-neighbor prediction, quickly suffer from the curse of dimensionality and overfitting for more than a few predictors which has limited their application mostly to the univariate case. Therefore, selection strategies are needed that harness the available information as efficiently as possible. Since often the right combination of predictors matters, ideally all subsets of possible predictors should be tested for their predictive power, but the exponentially growing number of combinations makes such an approach computationally prohibitive. Here a prediction scheme that overcomes this strong limitation is introduced utilizing a causal preselection step which drastically reduces the number of possible predictors to the most predictive set of causal drivers making a globally optimal search scheme tractable. The information-theoretic optimality is derived and practical selection criteria are discussed. As demonstrated for multivariate nonlinear stochastic delay processes, the optimal scheme can even be less computationally expensive than commonly used suboptimal schemes like forward selection. The method suggests a general framework to apply the optimal model-free approach to select variables and subsequently fit a model to further improve a prediction or learn statistical dependencies. The performance of this framework is illustrated on a climatological index of El Niño Southern Oscillation.

DOI: [10.1103/PhysRevE.91.052909](https://doi.org/10.1103/PhysRevE.91.052909)

PACS number(s): 05.45.Tp, 02.50.-r, 89.70.Cf, 92.60.Wc

I. INTRODUCTION

Predicting the future behavior of complex systems from measured time series constitutes a major goal in many fields of science. Traditionally, this problem has been mainly addressed in terms of model-based approaches, often using linear models [1]. As an alternative, since the late 1980s also model-free predictions have been developed using nearest neighbors in state space [2–7] or neural networks [8–10]. In the nearest-neighbor technique, states similar to the present state are searched for in the past of a time series Y and a future value Y_{t+h} at a prediction step h is forecasted by simply averaging the Y values h steps ahead of the nearby past states or using local-linear models [2]. Model-free predictions have so far been mostly univariate where states are usually reconstructed from embedding a single time series using Takens's theorem [7,11,12]. However, the intertwined nature of complex systems calls for multivariate approaches taking into account more available information. Now the problem is that the curse of dimensionality makes useful nearest-neighbor predictions almost impossible for more than a few predictors [13], especially if many of these predictors carry redundant information.

From an information-theoretic perspective, the minimal set of variables that maximizes the (multivariate) mutual information [14] with a target variable is most predictive [6,15]. Minimality is required to avoid the curse of dimensionality. It is important to note that this set of variables can differ from those with *individually* large mutual information with the target variable. Indeed, sometimes the right combination of predictors matters. For example, if Y is driven multiplicatively by XZ , the mutual information of each of these predictors with Y can be very low and only the mutual

information of the combined set (X,Z) with Y is very high. In general, such *synergetic* sets can only be detected by searching through all subsets of variables. However, the number of possible combinations for taking into account more variables and larger time lags grows exponentially making such a search strategy prohibitive due to computational constraints.

Therefore, simple search strategies such as ranking the predictors by their mutual information with a target variable or the *CMI forward selection* using conditional mutual information (CMI) have been proposed recently [16]. Here we demonstrate that such approaches can fail already in simple cases where one cannot avoid to test different subsets of predictors. However, we information-theoretically prove that the search can be restricted to *causal* drivers. To obtain these drivers, we exploit a recently developed model-free algorithm to consistently reconstruct causal drivers from multivariate time series that keeps the estimation dimension as low as possible with typically low computational complexity [17]. The much smaller set of causal drivers then allows for a globally optimizing search strategy which is optimal also for synergetic cases. In this contribution, we additionally provide a practical criterion for selecting the optimal size of the subset of predictors which compares well even with computationally expensive cross-validation approaches. In numerical experiments we found that our optimal scheme yields much improved predictions and often even runs faster than forward selection.

Our method suggests a general framework not only for prediction but also for general statistical inference problems for data sets (not only time series) where the underlying mechanisms are poorly understood: First, the optimal

model-free approach can be applied for selecting not only causal but also possibly synergetic driving variables and, second, these variables can be used to fit a model to learn the functional form of the dependencies on these causal predictors. This approach combines the advantage of a model-free approach to detect relevant variables with the advantage of model-based methods to efficiently harness these variables to further improve predictions or understand mechanisms. Our framework is illustrated on a sea-surface-temperature-based index of the El Niño Southern Oscillation (ENSO) in the tropical Pacific.

This paper is organized as follows. After deriving the optimality of causal predictors in Sec. II, we discuss common approaches for information-theoretic variable selection for predictions in Sec. III. The optimal scheme is explained in Sec. IV including the causal preselection algorithm and selection criteria. Section V discusses the computational complexity of the different schemes. In Sec. VI we compare our scheme with other approaches in a model example. Extensive numerical experiments on multivariate nonlinear stochastic delay processes are conducted in Sec. VII. In Sec. VIII we analyze the combination of the model-free selection with a model-based prediction scheme which is applied in Sec. IX to predict future values of the considered ENSO index.

II. OPTIMAL PREDICTION

The discrete-time evolution of a subprocess Y of a multivariate N -dimensional stochastic process \mathbf{X} can be described by

$$Y_{t+1} = f(\mathcal{P}_{Y_{t+1}}, \eta_{t+1}^Y), \quad (1)$$

with some function $f(\cdot)$, where $\mathcal{P}_{Y_{t+1}} \subset \mathbf{X}_{t+1}^- = (\mathbf{X}_t, \mathbf{X}_{t-1}, \dots)$ are the deterministically driving variables including the past of Y at possibly different time lags and η_{t+1}^Y represents stochastic noise driving Y . The uncertainty about the outcome of Y_{t+1} can be quantified by the Shannon entropy [14] which decomposes into

$$H(Y_{t+1}) = I(Y_{t+1}; \mathbf{X}_{t+1}^-) + H(Y_{t+1} | \mathbf{X}_{t+1}^-), \quad (2)$$

where the latter term is the *source entropy* [18,19]. This conditional entropy quantifies the minimum level of uncertainty that cannot be predicted even if the whole past (and present) \mathbf{X}_{t+1}^- is known. If the dependency of f on the noise term η_{t+1}^Y is additive, the source entropy equals the entropy of the noise: $H(Y_{t+1} | \mathbf{X}_{t+1}^-) = H(\eta_{t+1}^Y)$. The infinite-dimensional multivariate mutual information (MMI) $I(Y_{t+1}; \mathbf{X}_{t+1}^-)$, on the other hand, quantifies the predictable part by measuring by how much the uncertainty about the outcome of Y_{t+1} can be maximally reduced if \mathbf{X}_{t+1}^- was perfectly measured.

In practice, a prediction using the entire set \mathbf{X}_{t+1}^- (truncated at some maximal lag) would severely suffer from the curse of dimensionality and *overfitting* [13], which means that many variables do not actually carry useful information but merely fit the noise in the time series, and the prediction—trained on a learning set—would perform poorly on a test set. The goal is,

thus, to select a small subset of predictors from \mathbf{X}_{t+1}^- that carries a maximum of information about Y_{t+1} and still generalizes well on new data. However, for an N -variate process \mathbf{X} truncated to a maximum delay τ_{\max} the number of possible subsets grows exponentially in N and τ_{\max} .

To avoid this combinatorial explosion, simple search strategies such as ranking the individual predictors by their mutual information (MI scheme) with the target variable can be used which, however, is prone to include redundant variables that do not improve a prediction. Forward selection, a more advanced technique, iteratively determines predictors based on how much information they contain *additionally* to the already chosen variables using conditional mutual information (CMI scheme) [16] leading to a polynomial computational complexity. These strategies will be discussed in Sec. III. But forward selection is not a globally optimal strategy, one reason being that it might select variables that are not deterministically driving Y_{t+1} , and the other reason is that it fails to detect synergetic cases as demonstrated in our model example in Sec. VI.

The unknown deterministic drivers $\mathcal{P}_{Y_{t+1}}$ in Eq. (1) are, however, key to arrive at optimal predictors as can be shown by decomposing the MMI in Eq. (2) using the chain rule [14]

$$I(Y_{t+1}; \mathbf{X}_{t+1}^-) = I(Y_{t+1}; \mathcal{P}_{Y_{t+1}}) + \underbrace{I(Y_{t+1}; \mathbf{X}_{t+1}^- \setminus \mathcal{P}_{Y_{t+1}} | \mathcal{P}_{Y_{t+1}})}_{=0}, \quad (3)$$

where “ \setminus ” denotes the set subtraction. The second term is zero for processes satisfying the *Markov property* which states that Y_{t+1} is independent of the remaining past given its *causal parents* $\mathcal{P}_{Y_{t+1}}$, a term that originates from the theory of *graphical models* [20]. For multivariate time series these are called *time-series graphs* [17,21,22]. This proves that, theoretically, all information is already contained in the causal parents. Adding a variable from \mathbf{X}_{t+1}^- would not increase the information, but removing a parent from $\mathcal{P}_{Y_{t+1}}$ would decrease the MMI in Eq. (3). The causal parents can be efficiently estimated by the algorithm described in Sec. IV A. The Markovity rests on the assumption that the noise term η_{t+1}^Y driving Y is independent of the noise terms driving the other variables. While this assumption is not strictly fulfilled in many real-world systems, it often at least approximately holds. This assumption is also not as crucial for the prediction task as it is for the causal inference problem.

However, for finite time series, some predictors in $\mathcal{P}_{Y_{t+1}}$, even though causal, could only be weakly driving and lead to overfitting since they do not generalize well on new data. It is, therefore, crucial to optimize the selection of a minimal subset of causal parents. Since the set of causal parents $\mathcal{P}_{Y_{t+1}}$ is much smaller than \mathbf{X}_{t+1}^- , this can now be done using a global optimization strategy. In Sec. IV B we present such a strategy to select the *optimal subset* $\mathcal{P}_{Y_{t+1}}^{(p)} \subseteq \mathcal{P}_{Y_{t+1}}$ of causal predictors.

For predictions $h > 1$ steps into the future, the set of causal predictors $\mathcal{P}_{Y_{t+h}}$ for Y_{t+h} is no longer identical with the parents as for $h = 1$ because predictors can only be chosen among the

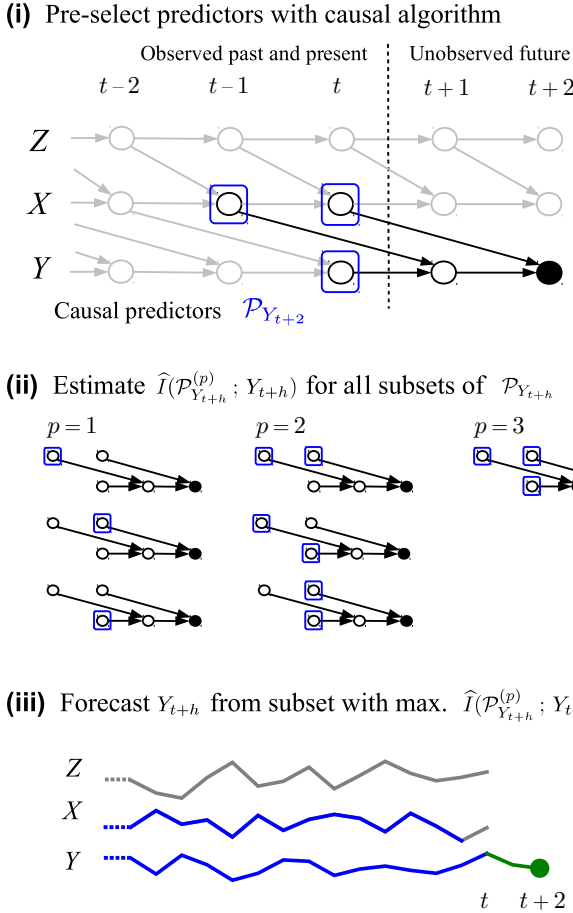


FIG. 1. (Color online) Optimal prediction scheme. (i) Causal preselection in an example time-series graph (see text) for predicting Y_{t+2} . Even though Z could have a high mutual information with Y , its influence is only indirect through the parents X_t and Y_{t+1} of Y_{t+2} . Among these the latter already lies in the unobserved future, but part of its information can still be recovered by measuring Y_t and X_{t-1} which share information along the paths marked with black arrows. These variables form the causal predictors $\mathcal{P}_{Y_{t+h}}$ (blue boxes) for $h = 2$, which can be found by determining the Markov set. A suitable algorithm for this task will be discussed in Sec. IV A. All paths from nodes further in the past to Y_{t+2} have to pass through this set. (ii) Selection of optimal subset from causal predictors. For all numbers of predictors p , the multivariate mutual information $\hat{I}(\mathcal{P}_{Y_{t+h}}^{(p)}; Y_{t+h})$ of all possible subsets is estimated. (iii) The optimal predictors are those where the estimated MMI takes its maximum or can be obtained from cross-validation. For example, if the optimal predictors were (Y_t, X_{t-1}) , only time series samples (blue shaded) of these predictors are used with a nearest-neighbor (Sec. IV C) or model-based (Sec. VIII) prediction scheme to forecast the future value $t + h$ of the time series of Y .

observed variables $\mathbf{X}_{t+1}^- = (\mathbf{X}_t, \mathbf{X}_{t-1}, \dots)$ prior to $t + 1$. Still the same algorithm as for the causal parents can be used to obtain the set of variables that separates Y_{t+h} from $\mathbf{X}_{t+1}^- \setminus \mathcal{P}_{Y_{t+h}}$ in the time-series graph. In Fig. 1(i) an example of such a graph is given. As defined in Ref. [17], each node in that graph represents a subprocess of a multivariate discrete time process \mathbf{X} at a certain time t . Nodes are connected by a directed link

if they are not independent conditionally on the past of the whole process, which implies a lag-specific *Granger causality* with respect to \mathbf{X} [21]. Using these causal predictors, the only uncertainty left comes from the source entropy of Y_{t+h} and the entropy from the unobserved ancestors of Y_{t+h} between $t+1$ and $t+h-1$ [see Fig. 1(i)].

In the following sections we discuss and numerically compare the four prediction schemes mentioned above: (1) MI selection, (2) CMI forward selection, (3) CMI forward selection of only causal predictors, and (4) our optimal scheme.

III. MI AND CMI PREDICTION SCHEMES

In the first prediction scheme, MI selection, the respective MI of each variable in \mathbf{X}_{t+1}^- up to a maximum lag τ_{\max} with the target variable Y_{t+h} is estimated. Then the predictors $X^{(i)} \in \mathbf{X}_{t+1}^-$ are ranked by their MI: $I(X^{(1)}; Y_{t+h}) > I(X^{(2)}; Y_{t+h}) > I(X^{(3)}; Y_{t+h}) > \dots$. To determine the best number p of the ranked predictors that should be used, one can apply either a heuristic criterion or *cross-validation* [13]. In the model experiments in Secs. VI and VII we evaluate both approaches, employing the heuristic criterion that the MI of the ranked predictor $X^{(p)}$ should be at least a fraction λ of the MI of the previous predictor $X^{(p-1)}$, i.e., $I(X^{(p)}; Y_{t+h}) > \lambda I(X^{(p-1)}; Y_{t+h})$ with $\lambda \in [0, 1)$. The last of the ranked predictors that satisfies this criterion determines the selected number \hat{p} of predictors. This scheme has the drawback that two or more predictors with high MI with the target variable might contain highly redundant information. Then the inclusion of redundant predictors leads to overfitting which will be discussed in Sec. VI.

The second prediction scheme, CMI forward selection, overcomes this drawback by excluding information already contained in the previous predictors [16]: First, the MIs $I(X_{t-\tau}; Y_{t+h})$ for all $X_{t-\tau} \in \mathbf{X}_{t+1}^-$ are estimated. The first predictor $X^{(1)}$ is the one that maximizes the MI with the target variable (i.e., the same one as in the MI-selection scheme). The next predictor $X^{(2)}$, however, is chosen according to the maximal CMI $I(X_{t-\tau}; Y_{t+h} | X^{(1)})$ among all remaining predictors, the third predictor is the maximum CMI conditional on the two previously selected predictors, and so on. In each step p , the CMI gives the gain to the MMI if this predictor is included because

$$\begin{aligned} & I[(X^{(1)}, \dots, X^{(p)}); Y_{t+h}] \\ &= \underbrace{I[(X^{(1)}, \dots, X^{(p-1)}); Y_{t+h}]}_{\text{MMI without } X^{(p)}} \\ &+ \underbrace{I[X^{(p)}; Y_{t+h} | (X^{(1)}, \dots, X^{(p-1)})]}_{\text{gain due to } X^{(p)}}. \end{aligned} \quad (4)$$

Here the heuristic criterion [16] is to select as the best \hat{p} the last p with at least a gain of

$$\begin{aligned} & I[X^{(p)}; Y_{t+h} | (X^{(1)}, \dots, X^{(p-1)})] \\ & > \lambda I[(X^{(1)}, \dots, X^{(p-1)}); Y_{t+h}], \end{aligned} \quad (5)$$

where $\lambda \in [0, 1)$ as before. In Ref. [16] also an adaptive choice of \hat{p} using a shuffle test is discussed. This scheme has been

proposed to infer causal drivers in Ref. [16]. However, it can be shown to fail for this task already in simple cases which will be shown in Sec. VI.

Rather than with a heuristic criterion, at the cost of additional computation time, the best number p of predictors can also be chosen by cross-validation. Here we use an m -fold cross-validation where the available observed set of time indices $\mathcal{T} = \{1, \dots, T\}$ is partitioned into m complementary segments. For each validation round, a fold m is retained as the *validation set* \mathcal{T}_m on which the prediction performance is evaluated. The nearest-neighbors are searched for in the complementary set $\mathcal{T} \setminus \mathcal{T}_m$ from which also the prediction estimate is generated. Then the number \hat{p} of predictors where the average prediction error across all m folds is minimal is chosen.

IV. OPTIMAL PREDICTION SCHEME

For the prediction of Y_{t+h} given the multivariate time series \mathbf{X} , our proposed optimal prediction scheme consists of the following steps (Fig. 1): (i) Estimate the causal predictors $\mathcal{P}_{Y_{t+h}} \subset \mathbf{X}_{t+1}^-$ using the causal algorithm described in Sec. IV A, (ii) check all subsets (except the empty one) and select the p causal predictors $\mathcal{P}_{Y_{t+h}}^{(p)} \subseteq \mathcal{P}_{Y_{t+h}}$ with the maximal *estimate* of the MMI $\hat{I}(\mathcal{P}_{Y_{t+h}}^{(p)}; Y_{t+h})$ with the target variable as the optimal ones (Sec. IV B), and (iii) use these predictors to forecast the target variable with nearest-neighbor prediction (Sec. IV C). In the following, we explain the causal preselection algorithm and discuss criteria for selecting the optimal subset. While here the actual prediction is conducted using a nearest-neighbor scheme [2], in Sec. VIII we will also discuss how a model-based prediction based on the inferred optimal predictors can further improve a forecast.

A. Causal preselection algorithm

The causal preselection algorithm is a modification of the algorithm introduced in Ref. [17], which is based on the PC algorithm [23,24] (named after its inventors Peter Spirtes and Clark Glymour). The main idea is to iteratively unveil the causal predictors by testing for independence between pairs of nodes in the time-series graph conditional on a subset of the remaining nodes. Since these conditions are efficiently chosen, the dimension stays as low as possible in every iteration step. This important feature helps to alleviate the curse of dimensionality in estimating CMIs [17] affecting the computation time as well as the reliability of conditional independence tests. Under some mild assumptions discussed below, the algorithm yields a consistent estimate of $\mathcal{P}_{Y_{t+h}}$. Instead of the commonly used binning estimators where the curse of dimensionality is especially severe, here we utilize an advanced nearest-neighbor estimator [25] that is most suitable for variables taking on a continuous range of values. This estimator has as a free parameter the number of nearest neighbors k which determines the size of hypercubes around each (high-dimensional) sample point. Small values of k lead

to a lower estimation bias but higher variance and vice versa. Therefore, we choose different values in the algorithm (k^{algo}) and the subsequent selection schemes ($k^{\text{CMI/MMI}}$). Note that for an estimation from (multivariate) time-series stationarity is required.

The algorithm starts with no *a priori* knowledge about the drivers and iteratively learns the set of predictors of Y : First, estimate unconditional dependencies $I(X_{t-\tau}; Y_{t+h})$ and initialize the preliminary predictors $\mathcal{P}_{Y_{t+h}} = \{X_{t-\tau} \in \mathbf{X}_{t+1}^- : I(X_{t-\tau}; Y_{t+h}) > 0\}$. This set contains also noncausal predictors which are now iteratively removed by testing whether the dependence between Y_{t+h} and each $X_{t-\tau} \in \mathcal{P}_{Y_{t+h}}$ conditioned on the incrementally increased subset $\mathcal{P}_{Y_{t+h}}^{n,i} \subseteq \mathcal{P}_{Y_{t+h}}$ vanishes:

- (n.) Iterate n over increasing number of conditions, starting with some $n_0 > 0$:
- (n.i) Iterate i through all combinations of picking n nodes from $\mathcal{P}_{Y_{t+h}}$ to define the conditions $\mathcal{P}_{Y_{t+h}}^{n,i}$ in this step, and estimate $I(X_{t-\tau}; Y_{t+h} | \mathcal{P}_{Y_{t+h}}^{n,i})$ for all $X_{t-\tau} \in \mathcal{P}_{Y_{t+h}}$. After each step the nodes $X_{t-\tau}$ with $I(X_{t-\tau}; Y_{t+h} | \mathcal{P}_{Y_{t+h}}^{n,i}) = 0$ are removed from $\mathcal{P}_{Y_{t+h}}$ and the iteration over i stops if all possible combinations have been tested. (In the implementation we limit the number n_i of combinations and check relevant combinations first, see Refs. [17,26] for details.)

If the cardinality $|\mathcal{P}_{Y_{t+h}}| \leq n$, then the algorithm converges, else increase n by one and iterate again. (In the implementation we limit the dimensionality up to some n_{max} . If the initial number of conditions is $n_0 > 1$ to speed up the algorithm, then also previously skipped combinations with $n < |\mathcal{P}_{Y_{t+h}}|$ need to be checked before convergence can be assessed.)

The main assumptions underlying the identification of the conditional independence structure with the PC algorithm are the *causal Markov condition*, i.e., Markovity of the process of any finite order, and *faithfulness*, which guarantees that the graph entails all conditional independence relations true for the underlying process and can be violated only in certain rather pathological cases [24]. If these assumptions are fulfilled, the causal algorithm is universally consistent, implying that the algorithm will converge to the true causal predictors with probability 1 for infinite sample size [24]. On the other hand, the CMI forward selection algorithm proposed for causal inference in Ref. [16] is not consistent since it yields noncausal drivers already in simple cases which will be analyzed in Sec. VI. But the CMI forward-selection scheme can be “repaired” by a further backward elimination step [27].

Practically, the causal algorithm involves conditional independence tests for $I(X_{t-\tau}; Y_{t+h} | \mathcal{P}_{Y_{t+h}}^{n,i}) = 0$. In Ref. [17] a shuffle test is proposed for testing whether $I(X_{t-\tau}; Y_{t+h} | \mathcal{P}_{Y_{t+h}}^{n,i}) > 0$: An ensemble of M values of $I(X_{t-\tau}^*; Y_{t+h} | \mathcal{P}_{Y_{t+h}}^{n,i})$ is generated where $X_{t-\tau}^*$ is a shuffled sample of $X_{t-\tau}$, i.e., with the indices permuted. Then the CMI values are sorted and for a test at a given α level, the αM -th value is taken as a significance threshold. In Ref. [17] a numerical study on the detection and false positive rates of the algorithm are given. The shuffle test comes at the additional cost that for each conditional independence test M surrogates of CMI have to

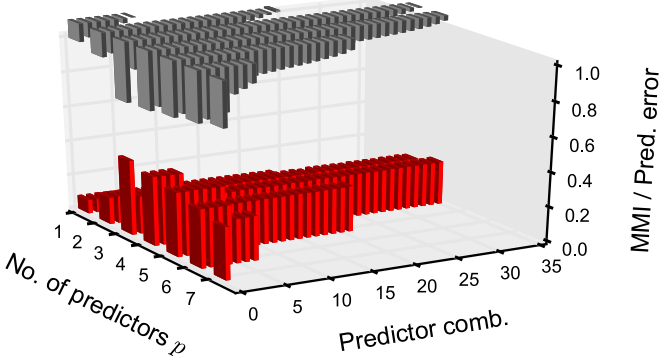


FIG. 2. (Color online) Multivariate mutual information (MMI, bottom red bars) and standardized root mean-squared prediction error (value proportional to lower end of gray bars at top) for all subsets of causal predictors for model (10), for details see Sec. VI. For each number of predictors p , the number of possible combinations varies according to the binomial coefficient $\binom{7}{p}$. The predictor combinations are sorted by their prediction error. The maximum of MMI and also smaller values match the minimum prediction error very well. Note that MMI is estimated on the learning set while the prediction error is evaluated out-of-sample on the test set indicating that the bias of higher-dimensional MMIs here serves as a good proxy for overfitting as discussed in the text.

be estimated. An alternative is to apply a fixed threshold I^* , which has the drawback that it does not adapt to the negative bias for higher-dimensional CMI [26,28]. The algorithm yields different numbers of predictors for different chosen fixed thresholds or significance levels and the value should be low enough to include weak but possibly synergetic causal predictors but high enough to limit computational complexity in the optimization step (Sec. V).

B. Optimal selection criteria

Once the causal predictors are determined, the optimal subset needs to be chosen (possibly containing all causal predictors). Here we also discuss a scheme using forward selection on the causal drivers (causal CMI scheme).

For the third prediction scheme, causal CMI selection, the forward-selection ranking discussed in Sec. III is applied not to the entire set \mathbf{X}_{t+1}^- but only to the preselected causal predictors $\mathcal{P}_{Y_{t+h}}$. Also here the same heuristic criterion as for the noncausal forward selection or cross-validation can be used.

For the optimal scheme [Fig. 1(ii)], the MMI $\widehat{I}(\mathcal{P}_{t+h}^{(p)}; Y_{t+h})$ for all subsets of the causal predictors from $\mathcal{P}_{Y_{t+h}}$ is estimated. In Fig. 2 the MMI values (ensemble average) are plotted for the model example discussed later in Sec. VI with $|\mathcal{P}_{Y_{t+h}}| = 7$ causal predictors. Even though all seven predictors are causal drivers, the *estimated* MMI is highest for just three of them and even decreases for more predictors [according to Eq. (3), the theoretical MMI should be maximal for seven predictors]. The reason is that the estimated MMI is negatively biased for higher dimensions [26,28,29] and if the additional information contained in the predictors does not outweigh this bias, the MMI decreases. The bias of the MMI estimator, therefore,

implies a penalty that avoids overfitting. In our heuristic criterion we exploit this property and simply select the subset $\mathcal{P}_{t+h}^{(p)} \subseteq \mathcal{P}_{t+h}$ with maximal estimated $\widehat{I}(\mathcal{P}_{t+h}^{(p)}; Y_{t+h})$. This model-free data-based criterion could be seen as an analog to model-based criteria like the Akaike information criterion (AIC) [30] where the penalty is derived from some measure of model complexity, but our criterion needs to be further investigated. We choose as the nearest-neighbor parameter in the MMI estimator $k^{\text{MMI}} = k$, where k is the nearest-neighbor parameter in the prediction (see next Sec. IV C).

In our numerical experiments in Sec. VII, we additionally test a combination of the MMI criterion with cross-validation: For each number of predictors p we check the MMI criterion to select the optimal combination and then use cross-validation to pick the optimal p . This approach gives always a slightly better prediction performance than using the maximum criterion alone—at the cost of much longer computation time. Asymptotically, for model-based predictions the AIC criterion is equivalent to cross-validation and in our numerical experiments (see Fig. 9) we also find that our heuristic MMI criterion well matches the cross-validation choice for longer time series. Note that our use of cross-validation treats p as a tuning parameter as is typically done [13]. One could also treat the choice of a subset $\mathcal{P}_{t+h}^{(p)}$ as a tuning parameter and run steps (i)–(iii) of the prediction scheme for every fold. $\mathcal{P}_{t+h}^{(p)}$ is, however, not a numeric “tuning parameter” and different folds in the cross-validation might not contain the same subsets. As a result the variance across the folds is considerable and it is hard to find the subset with minimal cross-validation error.

C. Nearest-neighbor prediction

Once the optimal predictors are selected, the actual prediction is conducted here using a scheme with a fixed number of nearest neighbors k [2]: For the optimal set of predictors $\mathcal{P}_{t+h}^{(p)}$, we first determine the distances

$$d_{t,s} = \|\mathcal{P}_{t+h}^{(p)} - \mathcal{P}_s^{(p)}\| \quad \text{for all } s \in \mathcal{T} \quad \text{with } s > \tau_{\max} + h, \quad (6)$$

where $\|\cdot\|$ denotes some norm. Here we apply the maximum norm as in the nearest-neighbor estimator of the (conditional) mutual information [25]. τ_{\max} is the maximum lag used to estimate $\mathcal{P}_{t+h}^{(p)}$. Next, we sort the distances in increasing order $d_{t,s_1} < d_{t,s_2} < \dots$, yielding an index sequence s_1, s_2, \dots . Now there are two approaches to use these distances: (i) A fixed distance ε is chosen and all points s with distance smaller than ε are taken into account to predict Y_{t+h} . Then the coarse-graining level is consistent for all sample points, but sometimes there might not be any point within a distance ε [6,15]. (ii) Here we use a fixed number of nearest neighbors which has the advantage that the same number of points contributes to a prediction making the estimate more reliable. For a chosen fixed number of nearest neighbors k the future value Y_{t+h} is then estimated by the conditional expectation and its prediction

interval by its standard deviation:

$$\hat{Y}_{t+h} = \frac{1}{k} \sum_{j=1}^k Y_{s_j}, \quad \hat{\sigma}(\hat{Y}_{t+h}) = \sqrt{\frac{1}{k} \sum_{j=1}^k (Y_{s_j} - \hat{Y}_{t+h})^2}. \quad (7)$$

Another option, instead of the expectation, is to fit an autoregressive model giving a *local-linear prediction* [2]. The summation can also be weighted with a function of the distance of the neighbors, different norms, or kernel-based methods can be chosen [13].

For the number of nearest neighbors k , we use $k = k^{\text{CMI/MMI}}$, where $k^{\text{CMI/MMI}}$ is the nearest-neighbor parameter in the estimation of CMI or MMI in the selection schemes. This choice approximately yields a consistent level of coarse-graining in the information-theoretic selection step and the nearest-neighbor prediction. Alternatively, at the cost of computation power, one can also utilize cross-validation for this choice [13]. The number of nearest neighbors needs to be balanced to guarantee that only nearby values are used as predictors, but still enough values are available to confidently estimate Y_{t+h} and possibly the prediction interval. The value will typically strongly depend on the data. As a skill metric we compute the standardized root-mean-squared error

$$\text{SRMSE} = \sqrt{\frac{\frac{1}{|\mathcal{T}_{\text{test}}|} \sum_{t \in \mathcal{T}_{\text{test}}} (Y_{t+h} - \hat{Y}_{t+h})^2}{\sigma_Y^2}}, \quad (8)$$

where σ_Y^2 is the variance of Y in the testing set $\mathcal{T}_{\text{test}}$.

V. COMPUTATIONAL COMPLEXITY

An important criterion for the practical applicability of the different prediction schemes discussed in Secs. III and IV is their respective computational complexity. The estimator for the (C)MI $I(X; Y|Z)$ employed here (nearest-neighbor technique with maximum norm [25]) has a computational complexity of $\mathcal{O}[T^2(D_X + D_Y + D_Z)]$ [29], where T is the time-series length and D the dimensionality of the respective variable. Fast neighbor searching algorithms can further reduce the dependency on T , but here we are interested in the relative complexity of the different predictor selection schemes and, therefore, only consider the linear scaling with the number of dimensions. In this case the first prediction scheme, MI selection, clearly is the cheapest option. For a N -dimensional process \mathbf{X} , this procedure involves just $N(\tau_{\text{max}} + 1)$ estimates of MIs with a dimensionality of $D_X = 1, D_Y = 1, D_Z = 0$.

The second scheme, CMI forward selection, is more demanding the more possible predictors are included. For cross-validation, a maximum number of p_{max} predictors has to be selected, increasing the dimension to maximally $D_X = 1, D_Y = 1, D_Z = p_{\text{max}} - 1$ due to more conditions in Eq. (4). The CMI forward selection then involves

$$\begin{aligned} & \sum_{p=0}^{p_{\text{max}}-1} [N(\tau_{\text{max}} + 1) - p] \\ &= \frac{1}{2} p_{\text{max}} [1 - p_{\text{max}} + 2N(1 + \tau_{\text{max}})] \end{aligned}$$

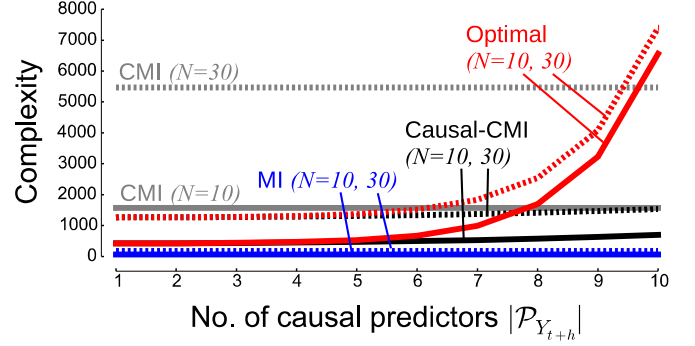


FIG. 3. (Color online) Scaling of computational complexity with the cardinality of the estimated causal predictors $|\mathcal{P}_{Y_{t+h}}|$ for $\tau_{\text{max}} = 2$. Solid lines will be used for $N = 10$, whereas dashed lines denote $N = 30$. The curve for the causal CMI-selection scheme (black) differs only slightly from the constant curve for the MI-selection scheme (blue). The causal schemes have the advantage that their complexity rather depends on how many causal drivers are found, while the complexity of the MI- and CMI-selection schemes depends on the number of processes and the maximum lag. For the noncausal CMI-selection scheme (gray lines) we fixed the maximum number of predictors to $p_{\text{max}} = 10$. To include the complexity due to the preselection step in the causal schemes we added 420 (1260) for $N = 10$ ($N = 30$) for these schemes according to Eq. (9) for $n_0 = 2$ and $n_i = 3$ as in our numerical experiments (see Sec. VII). The plot shows that if the number of causal predictors can be reduced below 8 (10) here, the optimal scheme even takes less computation time than the noncausal CMI-selection scheme.

estimates of CMI with iteratively increasing dimensionality. Then the complexity of the CMI-selection scheme scales as

$$\begin{aligned} & \sum_{p=0}^{p_{\text{max}}-1} [N(\tau_{\text{max}} + 1) - p](p + 2) \\ &= \frac{1}{6} p_{\text{max}} [5 - 3p_{\text{max}} - 2p_{\text{max}}^2 + 3N(3 + p_{\text{max}})(1 + \tau_{\text{max}})] \end{aligned}$$

for $p_{\text{max}} < N(\tau_{\text{max}} + 1) + 1$, i.e., with a high linear dependency on $N(\tau_{\text{max}} + 1)$ and a polynomial dependency $\sim p_{\text{max}}^3$. Note that for an m -fold cross-validation step using nearest-neighbor prediction an additional computational complexity of

$$m \cdot T \frac{T}{m} \sum_{p=1}^{p_{\text{max}}} (1 + p) = \frac{1}{2} T^2 p_{\text{max}} (3 + p_{\text{max}})$$

has to be added.

In Fig. 3 we compare the complexity of the different prediction schemes for $N = 10$ and $\tau_{\text{max}} = 2$ as in our numerical experiments in Sec. VII. While one can fix p_{max} to a small number for which nearest-neighbor predictions yield acceptable results, the main problem here is that the computation time quickly increases with $N(\tau_{\text{max}} + 1)$ (linear but with a large prefactor). One advantage of a causal preselection step is to reduce this number before the computationally expensive selection procedure is invoked.

Since typically the set of causal predictors is small, i.e., $|\mathcal{P}_{Y_{t+h}}| \ll N(\tau_{\text{max}} + 1)$, the third prediction scheme, causal CMI selection, has a drastically smaller computational

complexity than the noncausal CMI-selection scheme if the additional complexity due to the causal algorithm is not that large. If we rank all causal predictors (not limiting to some p_{\max} as in the noncausal scheme), the complexity scales as $\frac{1}{6}|\mathcal{P}_{Y_{t+h}}|(|\mathcal{P}_{Y_{t+h}}| + 1)(|\mathcal{P}_{Y_{t+h}}| + 5)$, where $|\mathcal{P}_{Y_{t+h}}|$ denotes the number of causal predictors. In Fig. 3, the complexity (black lines) shows only a very moderate increase with $|\mathcal{P}_{Y_{t+h}}|$.

For the optimal scheme the computational complexity grows exponentially as $2^{|\mathcal{P}_{Y_{t+h}}|-1}(2 + |\mathcal{P}_{Y_{t+h}}|) - 1$. However, Fig. 3 shows that if the number of causal predictors can be restricted, the optimal scheme even takes less computation time than the noncausal CMI-selection scheme. The number of causal predictors can be reduced by adjusting the significance level α or fixed threshold I^* in the conditional independence tests of the causal preselection algorithm. Most important, the causal scheme's complexity only scales with the number of causal drivers and *not* directly with the number of processes N or the maximal lag τ_{\max} as the noncausal schemes (dashed lines in Fig. 3). The dependence of the causal schemes on N and the maximal lag τ_{\max} is only via the algorithm.

The additional time complexity of the causal algorithm varies with the graph structure. The number of iterations can be limited by starting with a higher number of initial conditions n_0 and limiting the maximum dimensionality n_{\max} . This number determines up to which dimension of $\mathcal{P}_{Y_{t+h}}^{n,i}$ the conditional independence is checked. Also the number of combinations n_i in the i loop can be restricted. In a worst-case scenario where the spurious links only vanish if the maximum number of conditions is used, the computational complexity scales as

$$N(\tau_{\max} + 1) \left[2 + \sum_{n=n_0}^{n_{\max}} \sum_{i=0}^{n-1} (2 + n) \right] \\ = N(\tau_{\max} + 1) \left[2 + \frac{n_i}{2} (1 + n_{\max} - n_0)(4 + n_{\max} + n_0) \right].$$

However, for sparse graphs and the conditional set being efficiently chosen [17], typically links get removed already for an n_0 -dimensional condition with a complexity of

$$N(\tau_{\max} + 1)[2 + n_i(2 + n_0)], \quad (9)$$

and $n_i = 1$ or 2. This is also confirmed in numerical experiments in Ref. [17] and Sec. VII. Often the complexity is even lower because the MI value in the first iteration is already nonsignificant.

VI. MODEL EXAMPLE

The following nonlinear discrete-time stochastic delay equation provides an illustrative example where a simple MI or CMI forward selection procedure yields noncausal variables that deteriorate a prediction. Consider

$$Y_{t+1} = c \sum_{i=1}^4 W_{t-1}^{(i)} + b \prod_{i=1}^3 Z_{t-1}^{(i)} + \sigma \eta_{t+1}^Y \\ X_t^{(1)} = a[W_{t-1}^{(1)} + W_{t-1}^{(3)}] + \eta_t^{X^{(1)}} \\ X_t^{(2)} = a[W_{t-1}^{(2)} + W_{t-1}^{(4)}] + \eta_t^{X^{(2)}}, \quad (10)$$

where the causal drivers $W^{(\cdot)}$, $Z^{(\cdot)}$, and η are independent Gaussian processes with zero mean and unit variance $[\sim \mathcal{N}(0, 1)]$. This model illustrates why the MI and CMI prediction schemes fail to yield good predictions due to (1) selecting noncausal predictors and (2) missing out synergetic predictors. Here the $Z_{t-1}^{(\cdot)}$ are synergetic causal drivers, which—for certain parameters b and c —are *individually* less predictive than the drivers $W_{t-1}^{(\cdot)}$. But selected all *together*, their combined information $I[(Z_{t-1}^{(1)}, Z_{t-1}^{(2)}, Z_{t-1}^{(3)}); Y_{t+1}]$ is much larger than the single mutual informations. This synergetic effect can only be detected if all subsets of causal drivers are tested in a globally optimal scheme. In the following we analyze why the MI- and CMI-selection schemes fail to provide good predictions for such cases.

Regarding the problem of selecting noncausal drivers, for certain parameter combinations of (a, b, c) the mutual information $I(X_t^{(i)}; Y_{t+1})$ is larger than any $I(W_{t-1}^{(\cdot)}; Y_{t+1})$ or $I(Z_{t-1}^{(\cdot)}; Y_{t+1})$ for all i . The noncausal schemes based on iteratively selecting predictors with maximal MI or CMI (blue and gray box plots in Fig. 4) will, therefore, choose a noncausal $X_t^{(\cdot)}$ prior to the true causal predictors $W_{t-1}^{(\cdot)}$ and $Z_{t-1}^{(\cdot)}$. Since the predictors $W_{t-1}^{(\cdot)}$ have the largest MI after the $X_t^{(\cdot)}$, these are included next in the MI-selection scheme. In the CMI forward selection scheme, on the other hand, the synergetic variables $Z_{t-1}^{(\cdot)}$ are selected after the second iteration step. This leads to a drastic decrease in the prediction error at $p = 5$ (gray box plot in Fig. 4). The problem is that now the dimension of the predictors is already 5 and the two spuriously causal variables $X_t^{(\cdot)}$ deteriorate the prediction.

The causal preselection avoids this pitfall. The black and red box plots in Fig. 4 denote the schemes based on causal predictors using forward selection (black) and the optimal scheme (red). Also for causal drivers the forward selection scheme is not optimal because the selection of one predictor at a time fails for synergetic cases and selects the weak drivers $W^{(\cdot)}$ prior to the synergetic drivers $Z^{(\cdot)}$. Only the optimal scheme correctly identifies these drivers for the dimension $p = 3$ and reaches the minimal prediction error possible for this model (black line) for each p . In Fig. 2 we show that the three synergetic drivers $Z^{(\cdot)}$ are chosen with our heuristic optimal criterion since they have the largest MMI with the target variable.

VII. NUMERICAL EXPERIMENTS

We next compare the four schemes including the causal preselection algorithm on a larger class of synergetic nonlinear discrete-time stochastic processes with different coupling configurations:

$$\mathbf{X}_{t+1}^{(1)} = c g_{\text{lin}}^{(1)}(\mathbf{X}_{t+1}^-) + b g_{\text{syn}}(\mathbf{X}_{t+1}^-) + \eta_{t+1}^{(1)} \\ \mathbf{X}_{t+1}^{(j)} = a g_{\text{lin}}^{(j)}(\mathbf{X}_{t+1}^-) + \eta_{t+1}^{(j)} \quad \text{for } j = 2, \dots, N, \quad (11)$$

where we are interested in predicting $\mathbf{X}_{t+1}^{(1)}$ (i.e., $h = 1$). The linear function $g_{\text{lin}}^{(1)}$ is simply the sum of five randomly chosen subprocesses $\mathbf{X}_{t-\tau}^{(\cdot)}$ [excluding process $\mathbf{X}^{(1)}$] at random lags $0 \leq \tau \leq 2$. The nonlinear function g_{syn} , on the other

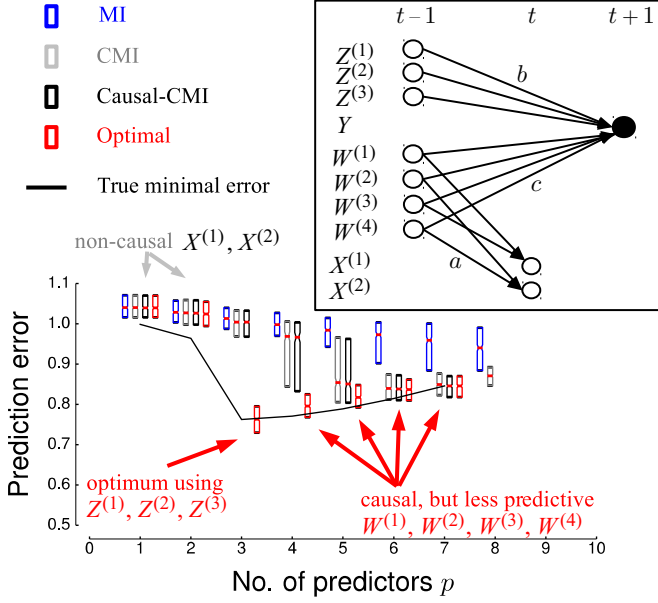


FIG. 4. (Color online) Comparison of prediction schemes for an ensemble of 500 realizations of model (10) for $a = 0.4, b = 2, c = 0.4, \sigma = 0.5$ with a time-series graph given in the inset [learning set length 500, test set length 125, nearest-neighbor prediction with $k = 10$ neighbors, (C)MIs estimated from the learning set with parameter $k^{\text{algo}} = 50$ in the algorithm and $k^{\text{CMI/MMI}} = 10$ for the optimization]. The box and whiskers plots give the ensemble median and the interquartile range of the standardized root-mean-squared prediction errors in the test set for each iteration step p in the four schemes (from left to right: MI, CMI, causal-CMI, optimal scheme). The black line gives the median of the true minimal prediction error obtained by minimizing the out-of-sample prediction error for each number of predictors p taken from the true causal drivers. For $p = 3$, only the optimal scheme (red) selects the best (synergetic) predictors $Z_{t-1}^{(1)}, Z_{t-1}^{(2)}, Z_{t-1}^{(3)}$ and reaches the minimum possible error while the causal forward selection (black) first picks one of the less predictive $W_{t-1}^{(1)}$ and the pure forward selection (gray) and MI-based schemes first pick the two noncausal drivers $X_t^{(1)}, X_t^{(2)}$. The nonoptimal schemes include the synergetic predictors only when the higher dimensionality is already worsening the prediction due to overfitting.

hand, is the product of three randomly chosen subprocesses [excluding process $\mathbf{X}^{(1)}$ and the ones already included in the linear term]. The other subprocesses for $j > 1$ are linearly driven by two other randomly chosen subprocesses, also at random lags. The coefficients are fixed to $a = 0.4, b = 2, c = 0.4$, and $N = 10$. With this setup we generate an ensemble of 500 realizations. We run the four schemes at different choices of the heuristic parameter λ and using cross-validation checking up to $p_{\text{max}} = 8$ predictors in the noncausal schemes. For the causal schemes, we use cross-validation for all estimated causal predictors (up to maximally $p_{\text{max}} = 8$, ranked by their CMI value in the algorithm [17]). The causal drivers are estimated using the algorithm parameters $n_0 = 2, n_{\text{max}} = 3, n_i = 3$ and $\tau_{\text{max}} = 2$ with a fixed significance threshold $I^* = 0.004$ (analyses for other thresholds are shown in the appendix).

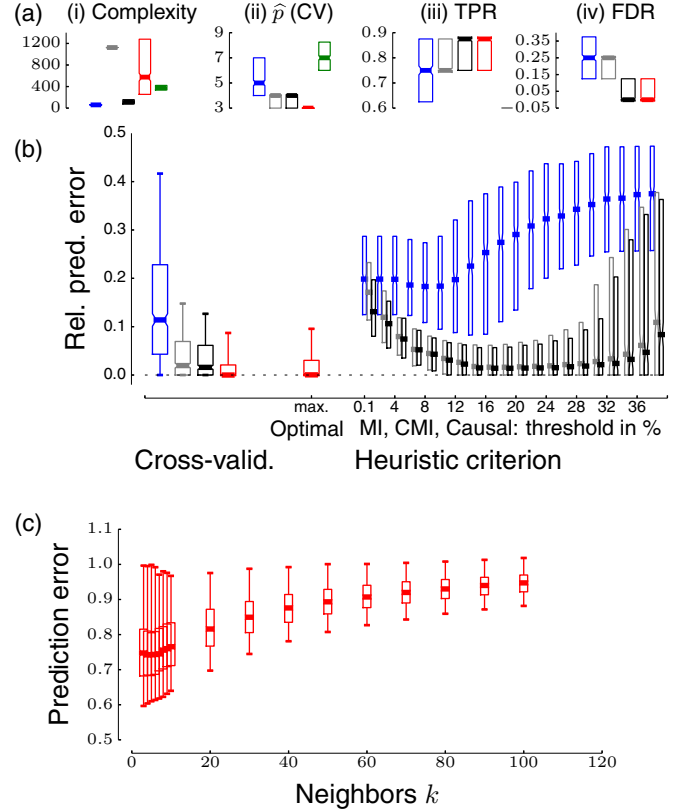


FIG. 5. (Color online) Results of numerical experiments for an ensemble of 500 trials of the synergetic model class (11) with time-series length $T = 500$ (125 in the test set) for the four different prediction schemes (from left to right in the panels: MI in blue, CMI in gray, causal-CMI in black, optimal in red). (a) The four box plots show the ensemble interquartile range of (i) computational complexity due to the causal algorithm, (ii) the range of numbers of predictors \hat{p} selected by cross-validation (CV, here the green bar shows the number of causal predictors $|\mathcal{P}_{Y_{t+h}}|$ in the preselection step), (iii) the true positive rate (TPR), and (iv) false discovery rate (FDR). The latter are evaluated for each scheme at $p = 8$, corresponding to the true number of causal drivers (or less if fewer causal drivers are estimated in the preselection step). (b) Box plots showing the median and interquartile range of the prediction error *relative* to the true minimal error obtained by minimizing the out-of-sample prediction error over all subsets of true causal drivers. On the left are the results if cross-validation is used to optimize p for each scheme (whiskers show the 5% and 95% quantiles). The red box in the center shows the result for the heuristic optimality criterion. The range of boxes on the right shows the results for different thresholds λ for the other schemes (only interquartile range). (c) Box and whiskers plots (showing the 5% and 95% quantiles) for the absolute prediction error of the optimal scheme at the cross-validated choice of p for different numbers of nearest neighbors k .

In Fig. 5, we show various statistics comparing the computational complexities and the prediction errors. Here the number of possible predictors is $N(\tau_{\text{max}} + 1) = 30$ yielding a computational complexity of 60 for the MI-selection scheme (blue) and 1124 for the CMI-selection scheme (gray) using $p_{\text{max}} = 8$. The causal algorithm reduces the number of

possible predictors to about 7 (median). This corresponds to a true positive rate (TPR) of roughly 0.9 [there are eight true drivers in model (11), but several are only weakly driving] and a zero false discovery rate (FDR), while the MI- and CMI-selection schemes detect fewer causal drivers and much more false positives. Fewer predictors result in a lower computational complexity for the causal prediction schemes. The causal CMI-selection scheme runs extremely fast (black) and the complexity of the optimal scheme (red) strongly varies among the different realizations, since it depends exponentially on how many causal predictors are preselected (Fig. 3) but still typically even stays below the noncausal CMI-selection scheme. Using cross-validation, the MI-selection scheme uses typically (median) $p = 5$, the CMI-selection schemes both $p = 4$, and the optimal scheme only $p = 3$ from the 8 true causal drivers for this model.

Finally, the relative prediction errors show that only the optimal scheme reaches the lowest possible errors with a median of zero relative error and even 90% of the ensemble below an error of 0.1. This demonstrates the large improvements due to the global optimization scheme that is only possible after reducing the set of variables to the few causal predictors.

The aforementioned results have been obtained using cross-validation to select the optimal p . The computationally cheaper alternative using a heuristic criterion here yields drastic differences in the prediction performance depending on the choice of λ . While here values in the range $\lambda = 20\% \dots 30\%$ give good results, in another experiment (Fig. 6) we found good predictions only for $\lambda = 10\% \dots 15\%$, making it hard to provide rules of thumb in practical applications. In the appendix we show that also the length of the time series results in different optimal ranges for λ . On the other hand, for the optimal scheme the heuristic choice leads to almost the same minimal errors as in cross-validation.

To test the robustness of our results, we also compare the prediction schemes on a class of nonsynergetic, but still nonlinearly coupled, models (generalized additive models [13] with $N = 4$ processes and polynomials of linear and quadratic degree) as analyzed in the supplement of Ref. [17]. For each ensemble member, we choose as a target variable the one with the largest sum of “incoming” coefficients (absolute values). The results shown in Fig. 6 demonstrate that for this case also the causal CMI-selection scheme reaches optimal prediction errors. In the appendix, we show that the optimality is robust also for different significance thresholds and other time-series lengths.

In Figs. 5(c) and 6(c) we evaluate the prediction for different phase-space resolutions. To this end we use the causal predictors and run steps (ii) and (iii) of the prediction scheme from Fig. 1 (using cross-validation to choose \hat{p}) for varying nearest-neighbor parameters k and MMI estimation parameter $k^{\text{MMI}} = k$, both set to the same value for consistency. While in the first ensemble [Fig. 5(c)] the error is minimal for very few neighbors and sharply rises if too many neighbors are used, in the second ensemble [Fig. 6(c)] too few neighbors yield worse results. In practice, the choice will very much depend on the process under study, but here we use a

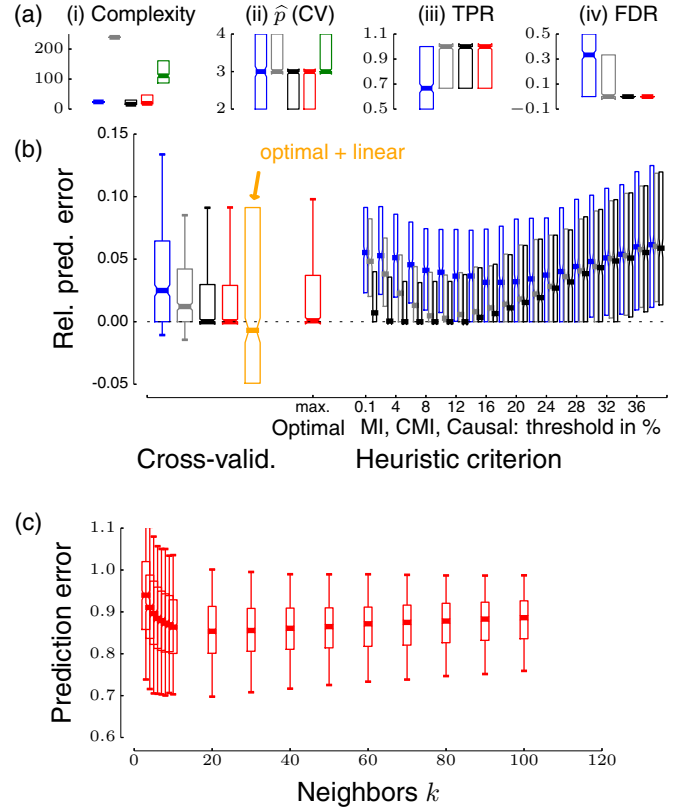


FIG. 6. (Color online) Numerical experiments on a class of non-synergetic but still nonlinearly coupled generalized additive models as analyzed in the supplement of Ref. [17]. Parameters as in Fig. 5 but with $N = 4$ processes and $p_{\text{max}} = 6$. The orange box plot in (b) shows the relative prediction error if the optimal predictors from the model-free selection scheme are used in conjunction with the linear autoregressive prediction model (12) (only the interquartile range shown).

value in the range $k = 5 \dots 10$ which constitutes a balance between local information and enough neighbors to reliably estimate \hat{Y}_{t+h} .

VIII. MODEL-FREE SELECTION COMBINED WITH MODEL-BASED PREDICTION

Up to now we have stayed in a model-free framework with information-theoretic optimal selection of predictors and a nearest-neighbor prediction. While nearest-neighbor prediction is a flexible method that will adapt to any function f in Eq. (1), it will in many cases be outperformed by a model-based prediction—if the right model class is chosen. If a misspecified model is chosen for variable selection and fitting, it might miss out nonlinear combinations of predictors. For example, in our synergetic model (10) a linear selection method would only include the weakly predictive variables $W^{(\cdot)}$ and largely miss out the highly predictive variables $Z^{(\cdot)}$. The functional dependency on the $W^{(\cdot)}$ is, on the other hand, much better fitted with a linear model than with nearest neighbors.

To take advantage of improved model-based predictions and at the same time not miss out synergetic predictor combinations, we propose to apply our optimal predictor selection scheme and conduct the final prediction step by fitting a model on the optimal set of predictors. Here we demonstrate this approach on the nonsynergetic model ensemble from Ref. [17]. To predict Y_{t+h} from the optimal subset of predictors $\mathcal{P}_{Y_{t+h}}^{(\hat{p})}$ (chosen by cross-validation or the heuristic criterion), we use the ordinary least-squares regression technique. Then the prediction interval is given by the variance $\hat{\sigma}_\varepsilon^2$ of the regression residual plus the errors in the estimated regression coefficients $\hat{\mathbf{B}}$ [1]:

$$\hat{Y}_{t+h} = \mathcal{P}_{t+h}^{(\hat{p})} \hat{\mathbf{B}}, \quad \hat{\sigma}(\hat{Y}_{t+h}) = \sqrt{\hat{\sigma}_\varepsilon^2 + \sum_{i=1}^{\hat{p}} \hat{\sigma}(\hat{\mathbf{B}}_i)^2 (X^{(i)})^2}. \quad (12)$$

In Fig. 6(b) the results for this approach are shown as the orange box plot. The prediction improvement varies strongly for the different realizations which include non-linear and linear drivers. About half of the realizations are better fitted using the optimized linear approach with prediction improvements of up to 10% compared to the nearest-neighbors prediction. More advanced techniques such as generalized additive models can further improve a prediction [13]. In addition to facilitating the prediction task, the knowledge of the functional forms of dependencies can also help to better understand coupling mechanisms.

IX. PREDICTING ENSO

The combined framework developed in the last section is now illustrated on a sea-surface temperature index of the ENSO in the tropical Pacific, which has been the focus of prediction research for many decades due to its far-reaching climatic and economic impacts [31,32]. The Nino3.4 index is defined as the average sea-surface temperature over the region 5°N - 5°S , 170° - 120°W [33]. As another possible predictor variable, we use an atmospheric index based on sea-level pressure, the Southern Oscillation Index (SOI), which is computed from the surface air pressure difference between Tahiti and Darwin, Australia.

In Fig. 7(a), we employ the model-free causal algorithm and predictor selection [steps (i) and (ii) in Fig. 1, here optimized using cross-validation] to obtain the optimal predictors and compare the skill of the nearest-neighbor and the linear prediction using the autoregressive model (12) fitted on the optimal predictors. Trained on the period 1951–2002, we test the prediction on the last decade, 2003–2014. From the 24 possible predictors, the optimal predictor for $h = 1$ month is only Nino3.4 at a 1-month lag, while for $h = 2$ months the three predictors (Nino3.4_t, SOI_t, SOI_{t-9}) are relevant, indicating that the atmospheric coupling, including a long memory, constitutes an important predictive mechanism. Here the linear autoregressive model significantly reduces the prediction error by about 0.05–0.1 compared to the nearest-neighbor approach using the same predictors, at least for a few months ahead. For steps larger than 5

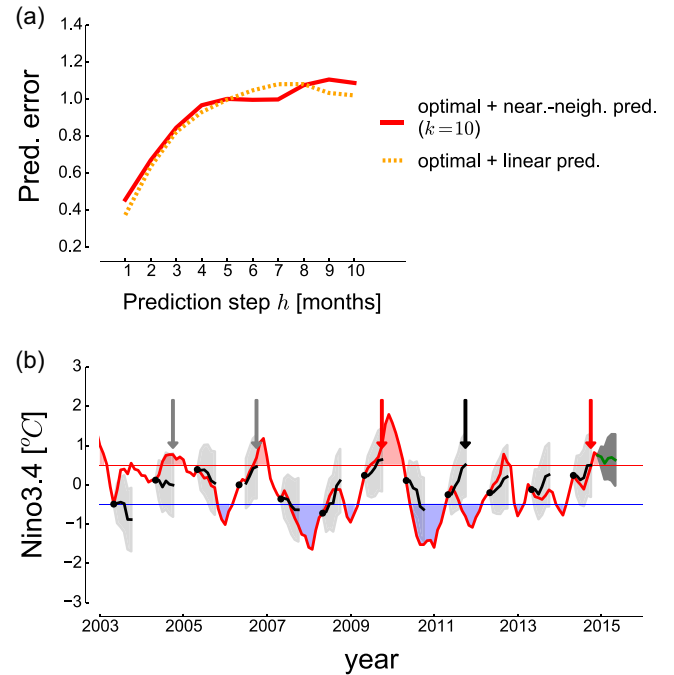


FIG. 7. (Color online) Prediction of the El Niño Southern Oscillation (ENSO) index Nino3.4 in the period 2003–2014 (to December) using 1951–2002 as a learning set, causal algorithm run with significance threshold $I^* = 0.03$ testing up to $\tau_{\max} = 12$ months. (a) Prediction error using nearest-neighbor prediction (solid red line) and linear prediction (dotted orange line) versus prediction step h . For both approaches the same optimal predictors obtained from the model-free scheme with cross-validated (fivefold within the learning set) choice of predictors are used. (b) Nino3.4 index with El Niño and La Niña events marked in red and blue, respectively. The black lines denote selected hindcasts and their 1σ -prediction intervals (gray) using the linear prediction. The dots mark the starting time t in May of each year, and the predicted values range from June ($h = 1$) to October ($h = 5$). The arrows mark correct (red), missed (gray), and false (black) hindcasts of the Nino3.4 index exceeding 0.5°C with more than 49% probability. The green line marks a real forecast starting in December 2014 giving a probability for the Nino3.4 index to stay above 0.5°C of 55–70% for the months until May 2015.

months, the error in both approaches quickly reaches 1, which implies that the prediction is merely a persistence forecast. The better linear prediction is a sign that exploiting the nonlinearities in Nino3.4 [34] does not improve the prediction much while the linear fit using the optimal predictors better harnesses the linear drivers of ENSO, at least on these time scales [35].

To give an impression of selected predictions from the linear model (actually hindcasts), we show in Fig. 7(b) the predictions up to 5 months ahead starting from May in each year. The important onsets of El Niño events are determined by expert assessment, but one definition is the 3-month-running-mean smoothed Nino3.4 index exceeding 0.5°C , here marked by a red line (La Niñas, where the index decreases below -0.5°C are marked in blue). With our hindcasts starting in May of each year, one can compute the probability of

an El Niño event as the part of the prediction distribution exceeding 0.5°C (assuming a Gaussian distribution with mean and standard deviation given by Eq. (12)). If this probability is larger than 49% for any of the 5 months ahead (until October), we predict an El Niño event. With this scheme we would have correctly predicted the moderate El Niño event in 2009 and the onset of the weak El Niño in current 2014–2015 season (red arrows) but missed the weak events of 2004 and 2006 (gray arrows, the latter being almost predicted with a probability of 48%). On the other hand, in 2011 (black arrow) a false alarm is given. The overall weak predictability of the recent El Niño events is also found in other studies using statistical as well as physical model predictions [32] and suggests that the mechanism of ENSO could be changing. Finally, our real forecast (green line) starting from December 2014 suggests that the weak current El Niño condition persists with a probability of 55–70% for the months until May 2015.

X. DISCUSSION AND CONCLUSIONS

In this article we have shown that the combinatorial explosion to search for globally optimal subsets of predictors can be overcome by restricting the search to causal drivers. Globally optimal predictors detect also synergetic mechanisms where the combination of multiple predictors strongly improves a prediction. Analytical considerations and numerical experiments indicate that such an approach is superior to schemes using MI ranking or forward selection with conditional mutual information. Another advantage is that the computational complexity only scales with the number of causal predictors and not directly with the number of processes included in the analysis. If the set of causal predictors is not that large, the optimal scheme is even computationally less expensive than the noncausal CMI-selection scheme. To determine the optimal size of this set, we have found that a parameter-free heuristic criterion performs almost as good as a computationally much more demanding cross-validation.

Note that even though, theoretically, only causal drivers can yield optimal predictions, noncausal variables could still be better predictors. Consider the case where a very high-dimensional process \mathbf{W} drives Y and X . Then the prediction of Y from the causal drivers \mathbf{W} is deteriorated due to the curse of dimensionality for finite samples, while the noncausal process X could potentially better aggregate this information. The same effect also explains why in Fig. 4 the CMI prediction using the noncausal $X_t^{(c)}$ together with the synergetic drivers has a slightly smaller prediction error than the causal CMI-selection scheme for $p = 5$ (gray box plot).

While we propose the model-free selection of predictors for processes where the underlying mechanisms are poorly understood, the actual prediction can be much improved using suitable model-based techniques compared to a pure model-free nearest-neighbor prediction. This approach combines the advantage of a model-free approach

to detect relevant variables with the smaller prediction variance of model-based methods and can also be used to better understand coupling mechanisms. The application of this combined approach significantly improves the prediction of an ENSO index compared to a nearest-neighbor scheme. The combined approach can be further improved by optimizing the number of predictors \hat{p} with a different criterion than the model-free criteria discussed in Sec. IV B. Linear models especially can harness much more predictors before the problem of overfitting becomes severe.

Here the scope of application was the prediction of future values of a time series. In a forthcoming paper we will investigate how the scheme can be adapted if, for example, only forecasts for the emergence of extreme events like El Niños [36] are needed. A *Python* script to estimate the causal predictors can be obtained from the author’s website at Ref. [37].

ACKNOWLEDGMENTS

We acknowledge the financial support by the German National Merit Foundation (Studienstiftung des deutschen Volkes), the Humboldt Graduate School, and the German Federal Ministry of Education and Research (Young Investigators Group CoSy-CC², Grant No. 01LN1306A).

APPENDIX

Robustness

In Fig. 8 we show the results as in Fig. 6 for different values of the significance threshold I^* . Obviously, this threshold affects the true positive and false discovery rate, which are, however, not directly of interest for the prediction task (as opposed to the causal inference problem). But a too-low significance level in the causal preselection algorithm leads to a high computational complexity and also increases the variance in the optimal subset selection step, which results in higher prediction errors. If, on the other hand, the significance level is too high, too few predictors are available to optimize the prediction such that the resulting optimal predictors equal the pre-selected causal predictors $\mathcal{P}_{Y_{t+h}}$. If the significance level is adjusted to yield just a few predictors more than the number of optimal predictors \hat{p} (obtained through cross-validation or the optimal heuristic criterion), the prediction error is minimal and also the computational complexity is lower than for the noncausal CMI forward selection scheme.

We also evaluate the prediction schemes for time-series lengths $T = 300$ and $T = 800$. The results shown in Fig. 9 demonstrate that the optimal scheme also works for very short time series and is even better for longer time series. For $T = 800$ and the synergetic model (11) the optimal scheme even results in 75% of the realizations reaching the true minimal prediction error.

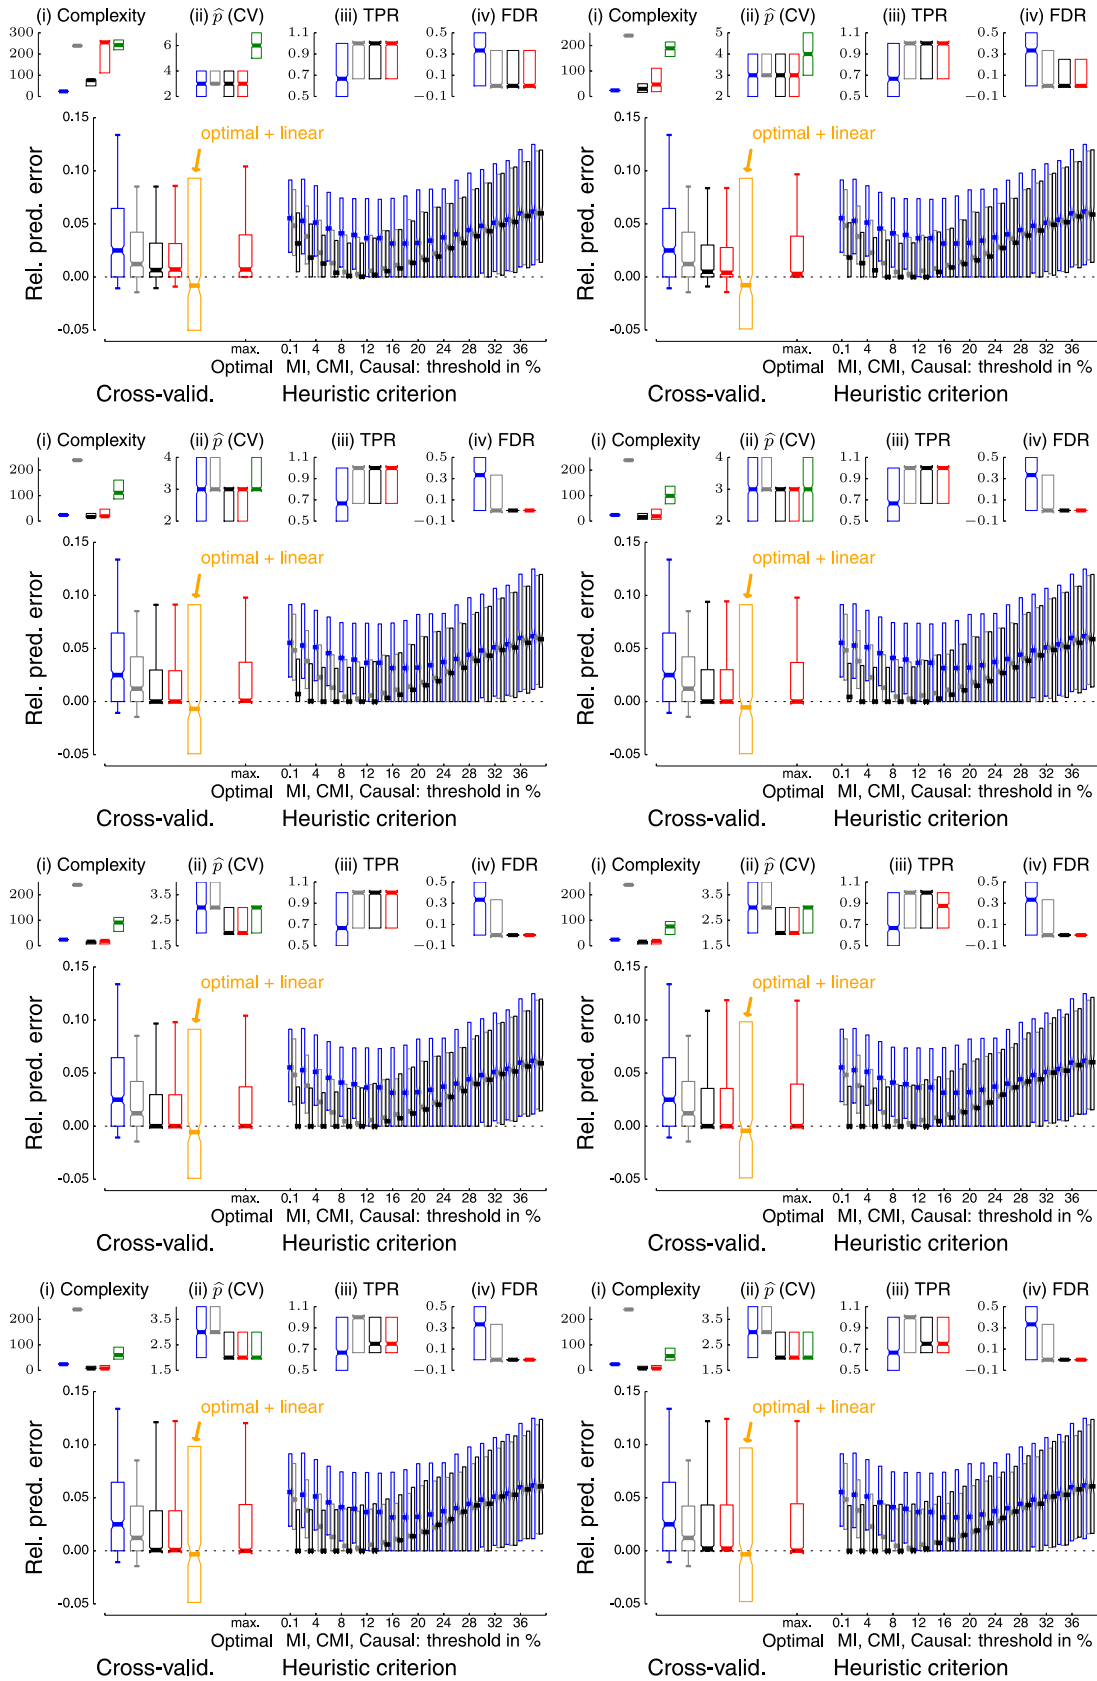


FIG. 8. (Color online) As in Figs. 6(a) and 6(b) but for a larger range of significance thresholds $I^* = 0.001, 0.002, 0.004, 0.005, 0.006, 0.008, 0.009, 0.01$ (row-wise from top left to bottom right).

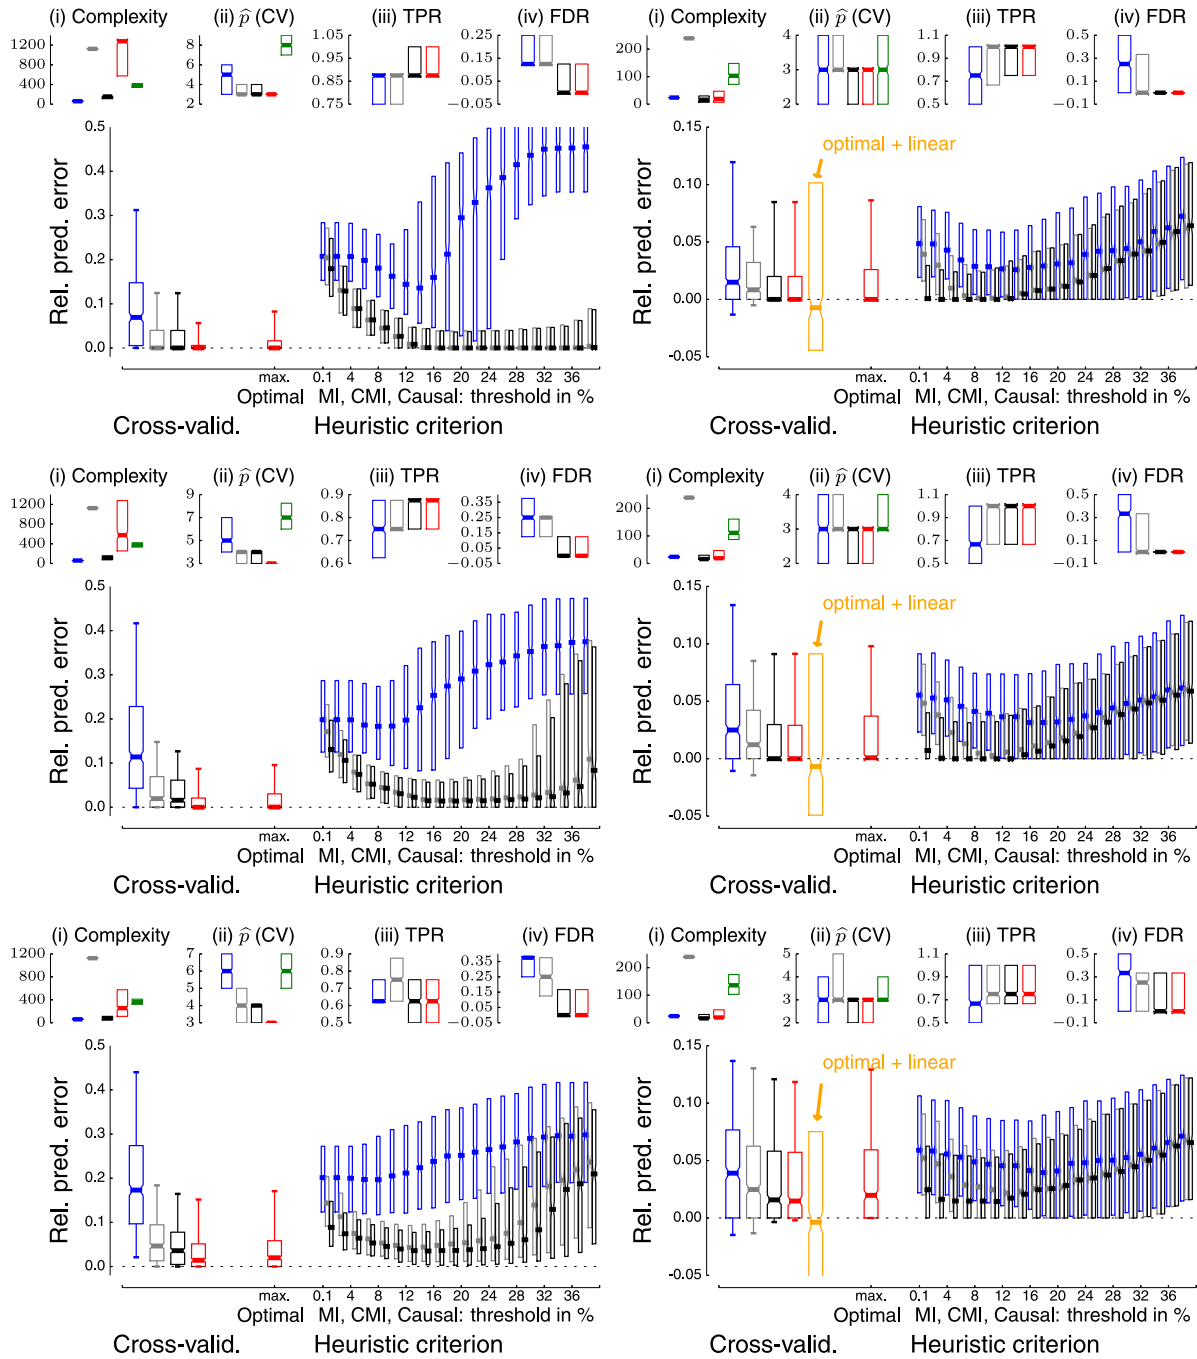


FIG. 9. (Color online) As in Fig. 5 (left column) and Fig. 6 (right column) but for a larger range of time-series lengths $T = 800, 500, 300$ (top to bottom, test set lengths 200, 125, 75).

[1] P. J. Brockwell and R. A. Davis, *Introduction to Time Series and Forecasting*, 2nd ed. (Springer, New York, 2002).
 [2] J. D. Farmer and J. J. Sidorowich, Predicting chaotic time series, *Phys. Rev. Lett.* **59**, 845 (1987).
 [3] H. D. I. Abarbanel, R. Brown, and J. B. Kadtke, Prediction in chaotic nonlinear systems: Methods for time series with broadband Fourier spectra, *Phys. Rev. A* **41**, 1782 (1990).
 [4] M. Giona, F. Lentini, and V. Cimagalli, Functional reconstruction and local prediction of chaotic time series, *Phys. Rev. A* **44**, 3496 (1991).
 [5] A. K. Alparslan, M. Sayar, and A. R. Atilgan, State-space prediction model for chaotic time series, *Phys. Rev. E* **58**, 2640 (1998).
 [6] B. Pompe, Mutual Information and Relevant Variables for Predictions, in *Modelling and Forecasting Financial Data* (Springer, New York, 2002), Vol. 2, p. 61.
 [7] M. Ragwitz and H. Kantz, Markov models from data by simple nonlinear time series predictors in delay embedding spaces, *Phys. Rev. E* **65**, 056201 (2002).

- [8] E. Eisenstein, I. Kanter, D. A. Kessler, and W. Kinzel, Generation and prediction of time series by a neural network, *Phys. Rev. Lett.* **74**, 6 (1995).
- [9] G. G. Szpiro, Forecasting chaotic time series with genetic algorithms, *Phys. Rev. E* **55**, 2557 (1997).
- [10] M. Small and C. Tse, Minimum description length neural networks for time series prediction, *Phys. Rev. E* **66**, 066701 (2002).
- [11] F. Takens, Detecting strange attractors in turbulence, in *Dynamical Systems and Turbulence, Warwick 1980: Proceedings of a Symposium Held at the University of Warwick 1979–80*, edited by D. A. Rand and L.-S. Young, Volume 898 of Lecture Notes in Mathematics (Springer, New York, 1981), p. 366.
- [12] T. Sauer, J. A. Yorke, and M. Casdagli, Embedology, *J. Stat. Phys.* **65**, 579 (1991).
- [13] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. (Springer, New York, 2009).
- [14] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, Hoboken, 2006).
- [15] A. Groth, Das Prinzip der maximalen Transinformation zur statistischen Modellierung, Diploma thesis, Ernst-Moritz-Arndt-Universität Greifswald, 2001.
- [16] D. Kugiumtzis, Direct-coupling information measure from nonuniform embedding, *Phys. Rev. E* **87**, 062918 (2013).
- [17] J. Runge, J. Heitzig, V. Petoukhov, and J. Kurths, Escaping the curse of dimensionality in estimating multivariate transfer entropy, *Phys. Rev. Lett.* **108**, 258701 (2012).
- [18] C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* **27**, 379 (1948).
- [19] B. Pompe and J. Runge, Momentary information transfer as a coupling measure of time series, *Phys. Rev. E* **83**, 051122 (2011).
- [20] S. L. Lauritzen, *Graphical Models* (Clarendon Press, Oxford, 1996), Vol. 16.
- [21] M. Eichler, Graphical modeling of multivariate time series, *Probab. Theor. Relat. Fields* **153**, 233 (2012).
- [22] J. Runge, V. Petoukhov, and J. Kurths, Quantifying the strength and delay of climatic interactions: The ambiguities of cross correlation and a novel measure based on graphical models, *J. Clim.* **27**, 720 (2014).
- [23] P. Spirtes and C. Glymour, An algorithm for fast recovery of sparse causal graphs, *Soc. Sci. Comput. Rev.* **9**, 62 (1991).
- [24] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search* (The MIT Press, Boston, 2000), Vol. 81.
- [25] S. Frenzel and B. Pompe, Partial mutual information for coupling analysis of multivariate time series, *Phys. Rev. Lett.* **99**, 204101 (2007).
- [26] J. Runge, Detecting and quantifying causality from time series of complex systems, Ph.D. thesis, Humboldt University, 2014.
- [27] J. Sun and E. M. Bollt, Causation entropy identifies indirect influences, dominance of neighbors and anticipatory couplings, *Physica D* **267**, 49 (2014).
- [28] J. Runge, J. Heitzig, N. Marwan, and J. Kurths, Quantifying causal coupling strength: A lag-specific measure for multivariate time series related to transfer entropy, *Phys. Rev. E* **86**, 061121 (2012).
- [29] A. Kraskov, H. Stögbauer, and P. Grassberger, Estimating mutual information, *Phys. Rev. E* **69**, 066138 (2004).
- [30] H. Akaike, A new look at the statistical model identification, *IEEE Trans. Aut. Control* **19**, 716 (1974).
- [31] M. A. Cane, S. E. Zebiak, and S. C. Dolan, Experimental forecasts of El Niño, *Nature* **321**, 827 (1986).
- [32] A. G. Barnston, M. K. Tippett, M. L. L'Heureux, S. Li, and D. G. DeWitt, Skill of real-time seasonal ENSO model predictions during 2002–11: Is our capability increasing?, *Bull. Am. Meteorol. Soc.* **93**, 631 (2012).
- [33] T. M. Smith, R. W. Reynolds, T. C. Peterson, and J. Lawrimore, Improvements to NOAA's historical merged land-ocean surface temperature analysis (1880–2006), *J. Clim.* **21**, 2283 (2008).
- [34] D. Dommenges, T. Bayr, and C. Frauen, Analysis of the non-linearity in the pattern and time evolution of El Niño southern oscillation, *Clim. Dynam.* **40**, 2825 (2013).
- [35] A. J. Gámez, C. S. Zhou, A. Timmermann, and J. Kurths, Nonlinear dimensionality reduction in climate data, *Nonlin. Process. Geophys.* **11**, 393 (2004).
- [36] J. Ludescher, A. Gozolchiani, M. I. Bogachev, A. Bunde, S. Havlin, and H. J. Schellnhuber, Improved El Niño forecasting by cooperativity detection, *Proc. Natl. Acad. Sci. USA* **110**, 11742 (2013).
- [37] www.pik-potsdam.de/members/jakrunge