

Learning Causal Models That Make Correct Manipulation Predictions With Time Series Data

Mark Voortman

*Decision Systems Laboratory
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA, 15260, USA*

VOORTMAN@SIS.PITT.EDU

Denver Dash

*Intel Research
Pittsburgh, PA, 15213, USA*

DENVER.H.DASH@INTEL.COM

Marek J. Druzdzel

*Decision Systems Laboratory
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA, 15260, USA*

MAREK@SIS.PITT.EDU

Editors: Isabelle Guyon, Dominik Janzing and Bernhard Schölkopf.

Abstract

One of the fundamental purposes of causal models is using them to predict the effects of manipulating various components of a system. It has been argued by Dash (2005, 2003) that the *Do* operator will fail when applied to an equilibrium model unless the underlying dynamic system obeys what he calls *Equilibration-Manipulation Commutability (EMC)*. Unfortunately, this fact renders most existing causal discovery algorithms unreliable for reasoning about manipulations, unless the details of the underlying dynamics of the system are explicitly represented in the model. Motivated by this caveat, in this paper we present a novel approach to causal discovery of dynamic models from time series. The approach uses a representation of dynamic causal models motivated by Iwasaki and Simon (1994), which asserts that all “causation across time” occurs because a variable’s derivative has been affected instantaneously. We present an algorithm that exploits this representation within a constraint-based learning framework by numerically calculating derivatives and learning instantaneous relationships. We argue that due to numerical errors in higher order derivatives, care must be taken when learning causal structure, but we show that the Iwasaki-Simon representation reduces the search space considerably, allowing us to forego calculating many high-order derivatives. Finally, we show that our approach can correctly recover the structure of a fairly complex dynamic system, and can predict the effect of manipulations accurately when a manipulation does not cause an instability. To our knowledge, this is the first causal discovery algorithm that has demonstrated that it can correctly predict the effects of manipulations for a system that does not obey the EMC condition.

Keywords: Causal discovery, dynamic systems, manipulations.

1. Introduction

One of the fundamental purposes of causal models is to use them for predicting the effects of manipulating various components of a system. It has been argued by Dash (2005, 2003) that the *Do* operator will fail when applied to an equilibrium model unless the underlying dynamic system obeys what he calls *Equilibration-Manipulation Commutability*, a principle which is illustrated by the graph in Figure 1. Unfortunately, this fact renders most existing

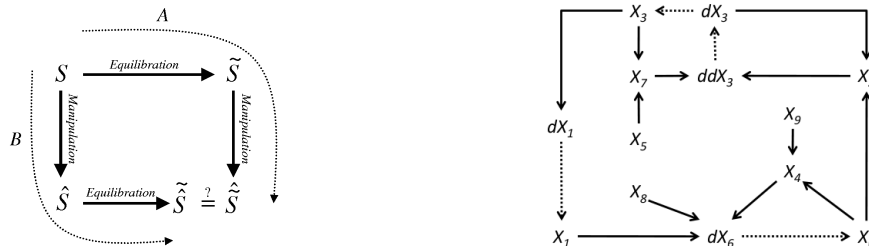


Figure 1: Equilibration-Manipulation Commutability provides sufficient conditions for an equilibrium causal graph to correctly predict the effect of manipulations (left). The dynamic system we used to simulate data (right).

causal discovery algorithms unreliable for reasoning about manipulations, unless the details of the underlying dynamics of the system are explicitly represented in the model.

Most classical causal discovery algorithms in AI make use of the class of independence constraints found in the data to infer causality between variables, assuming the faithfulness assumption (e.g., Spirtes et al., 2000; Pearl and Verma, 1991; Cooper and Herskovits, 1992). These methods will not be guaranteed to obey EMC if the observation time-scale of the data is long enough for some process in the underlying dynamic system to go through equilibrium. On the other hand, there have been previous approaches for learning dynamic causal models and Bayesian networks. Friedman et al. (1998) learn the structure of first-order Markov model by using time series data, and it would be straightforward to extend these approaches to higher-order Markovian models. However, the search space rapidly gets very large when searching for arbitrary dependencies across time.

Our approach, by contrast, uses an alternative representation of a dynamic system, explicitly modeling derivatives (or differences) of variables. It is beyond the scope of this paper to perform a quantitative comparison of prediction to these other approaches, however, we argue here that the representation that we learn helps us constrain the search space, and we expect that this reduction in complexity will make our algorithm perform better in practice and be more efficient than methods that try to learn fixed-order Markov structures for all variables.

2. Representation and Assumptions

Our approach uses a representation of dynamic causal models inspired by Iwasaki and Simon (1994), which asserts that all “causation across time” occurs because a variable’s derivative has been affected instantaneously. Iwasaki and Simon called these models “mixed causal

structures”. We use a slightly modified version of them and we call them “differential-based dynamic causal models” (DBD causal models, for short).

We use the notation $X^{(n)}$ to denote the n -th order derivative (or discrete version thereof) of variable X .

Definition 1 (DBD graphs) *Differential-based dynamic causal graphs over a set of time-dependent variables \mathbf{X} are discrete-time directed acyclic causal graphs, in which all “change across time” of a variable X occurs because there exists some n such that $X^{(n)}$ is being caused contemporaneously. That is, an edge exists from variable $Y_t \rightarrow X_{t+1}$ only if $Y_t = X_t^{(1)}$ or $Y_t = X_t$, in which case the parent set of X_{t+1} is $\{X_t, X_t^{(1)}\}$.*

The reason we constrain the parent set of X_{t+1} to be $\{X_t, X_t^{(1)}\}$ when X_t^1 is determined, is simply that, by definition,

$$X_{t+1} = X_t + X_t^{(1)}dt.$$

DBD models are unique in that they focus on uncovering *contemporaneous* causal relations that impact *derivatives* of some variables. They are motivated by real physical systems based on classical mechanics. For example, systems governed by Newton’s 2nd Law are archetypical causal systems: some “force” acts on a body, “causing” it to accelerate. The acceleration of the object, in turn, causes it to change velocity, which can cause the object to change position. The DBD representation assumes that all causation can be described in terms of “forces” causing a variable to change by impacting a derivative of some order instantaneously.

We show an example DBD graph in Figure 1. We will also use this example to illustrate the algorithm in one of the next sections. In the graph, two kinds of arcs are used: *solid arcs* that denote instantaneous causation, and *dashed arcs* that denote causation across time. The dashed arcs are also called integration links and are implicit, because they always point from a higher order derivative to a lower order derivative or the variable itself. The variables that have derivatives in the model are called dynamic, the remaining variables are called static. The learning algorithm has to find out which variables are dynamic and find the instantaneous causal arcs between variables and derivatives.

We now state two conjectures concerning DBD models that are useful in the causal discovery algorithm.

Conjecture 2 *Let A , B and C be different variables in a DBD model, and \dot{A} any order derivative of A . If the model contains an arc $A \leftarrow B$, then it does not contain the arc $\dot{A} \leftarrow C$.*

Conjecture 3 *Let A , B and C be different variables in a DBD model, and \dot{A} any order derivative of A . Let $A - B$ denote that there is an edge between A and B , that is directed as $A \rightarrow B$ or $A \leftarrow B$. If the model contains edges $A - B$ and $\dot{A} - C$, then the edge $A - B$ must be oriented $A \rightarrow B$.*

Our algorithm makes use of the PC algorithm (Spirtes et al., 2000), although we could have used any other causal discovery algorithm as well. Besides the assumptions required for the PC algorithm, we make two additional assumptions. First, we assume that the

system is stable, i.e., every dynamic variable must be part of a feedback loop. This implies that the highest order derivative of each dynamic variable must have at least one incoming arc. Second, exogenous variables are held constant over time and thus easily detectable in a data set.

3. The Algorithm

We present an algorithm that exploits the DBD representation within a constraint-based learning framework. The aim is to learn DBD models like the one given in Figure 1 directly. The input data¹ consisted of multiple time series that were generated first by parametrizing the model of Figure 1 with linear equations with independent Gaussian error terms, then by choosing different initial conditions for exogenous and dynamic variables and simulating 10000 discrete time steps. All derivatives of variables have been omitted from the data; thus, part of the challenge was that our method had to infer from the data which variables were changing due to the presence of derivatives and which were changing due to contemporaneous causation. Since calculating higher order derivatives using differences is sensitive to numerical errors, we opt for an incremental approach that gradually adds derivatives to the data set only when necessary, and exploits constraints given by our conjectures about feasible structures in these DBD graphs.

Our algorithm starts with only the original variables given in the data set and tries to learn structure between variables. Sometimes the presence of structure (plus our conjectures above) allow us to eliminate the possibility that some variables have derivatives. We can then calculate all feasible next-order derivatives, assuming the original edges learned are correct, and try to learn additional structure with these new variables. The process continues recursively until we reach a maximum-order derivative which is an input into the algorithm.

To illustrate the algorithm, we will use the example model from Figure 1. To find out which variables are dynamic, we run the PC algorithm on a data set containing only the variables. The resulting structure will be the graph in Figure 1 with all the derivatives removed. The following four disconnected graphs will be discovered:

- X_1 ; this variable has to be dynamic, because it is not exogenous.
- X_8 ; this variable is exogenous and, therefore, not dynamic.
- $X_5 \rightarrow X_7 \leftarrow X_3$; X_5 is exogenous, X_7 is instantaneously caused and not dynamic, X_3 is not exogenous and, therefore, dynamic.
- $X_2 - X_6 \rightarrow X_4 \leftarrow X_9$; either X_2 or X_6 is dynamic, X_4 is not, and X_9 is exogenous.

Summarizing, X_1 , X_3 , and either X_2 or X_6 are dynamic variables so there are only two competing models.

In the second step, for each of the competing models, the first order derivatives of the dynamic variables are added to the data set and the PC algorithm is executed again. The competing model in which X_2 is a dynamic variable will lead to an inconsistent structure, because there will be a v -structure into X_2 , namely $dX_3 \rightarrow X_2 \leftarrow X_6$. This violates

1. Downloadable from <http://www.causality.inf.ethz.ch/repository.php?id=16>

Conjecture 2 and so the structure is inconsistent. The other competing model is consistent, although no derivative of X_3 has an incoming edge. Therefore, as the last step, we add the second derivative of X_3 to the data set and run PC again to retrieve the original structure. We used Conjecture 3 as an extra rule to orient edges.

4. Prediction of Manipulations

The following results² were obtained by using the data and applying the instructions described in Appendix A. After running our algorithm on the data to obtain a causal structure, we estimated the coefficients in the equations in order to be able to make quantitative predictions. In the next step, we used the model and the values of the first four time steps in the data set to make predictions for time steps $\{5, 50, 100, 500, 1000, 2000, 4000, 10000\}$. We do not attempt to correct our predictions by using the data at times $t > 4$ when predicting later times, although doing this is possible and should improve our results.

The results are shown in Figure 2. Due to space constraints we chose not to present six tables, but instead calculated the average RMSE per time step for each manipulated variable. The graph shows that the error for the first few time steps is relatively small, but for all variables (except X_1) grows large in later times. Three variables in particular (X_2 , X_7 and X_4) had astronomical errors in later times. These huge rms errors are not indicative that our model was poor. In fact, in our case, since we generated the model, we could verify that the structure was exactly correct and the linear Gaussian parameters were very well identified. The reason for the unstable errors is that in the model of Figure 1, manipulating any variable except X_1 will break the feedback loop of a dynamic variable and thus will in general result in an instability (Dash, 2003). Thus our absolute rms error is expected to also be unstable in those manipulations, simply because we are predicting such large values.

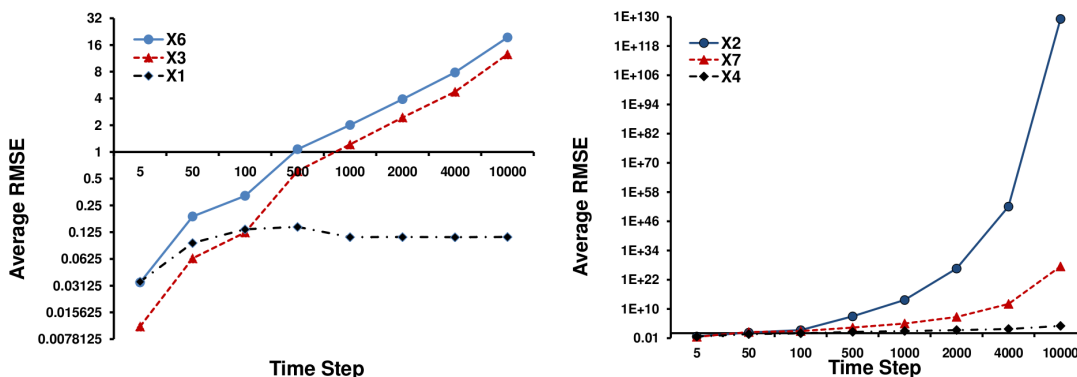


Figure 2: Average RMSE for each manipulated variable.

More important than getting the correct RMS error for these manipulations is the fact that our learned model correctly predicts that an instability will occur when any variable except X_1 is manipulated. In the absence of instability, our method has very low RMS error, as indicated by the curve of variable X_1 in Figure 2. This fact is significant, because the

2. Results can be downloaded from <http://pittsburgh.intel-research.net/~dhdash/rmse.zip>.

model retrieved from our system when variable X_1 is allowed to come to equilibrium will not obey the EMC condition (Dash, 2005). Thus, to our knowledge, we have presented the first algorithm that has demonstrated that it can correctly predict the effects of manipulations on systems that do not obey this condition.

5. Conclusions

We have described a first effort to construct an algorithm that can predict the effects of manipulations on systems that do not obey the EMC condition. We accomplish this by learning dynamic causal graphs in a representation very similar to that of Iwasaki and Simon. We have proposed a set of conjectures which are effective at constraining the search space for high-order Markovian relationships, which we expect will make this method more reliable and more efficient than other methods for learning temporal models, especially when higher-order relationships are present. We have shown that on a benchmark dataset generated from a fairly sophisticated dynamic system having multiple inter-related processes operating at widely varying time-scales, we were able to correctly learn the structure of the underlying system, and were able to predict that manipulating some variables in that system would result in an instability. Finally, in the absence of instabilities, we were able to predict with high accuracy the results of manipulating a variable, even far into the future. Future work will involve performing quantitative comparisons to other time-series methods. Also, although our conjectures formed useful heuristics for this method, we have been able to construct counter-examples where at least one of them is incorrect, so more work is needed to prove our existing conjectures and finding additional constraints on the search for derivatives.

References

- Gregory F. Cooper and Edward Herskovits. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9(4):309–347, 1992.
- Denver Dash. *Caveats for Causal Reasoning*. PhD thesis, Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, April 2003. <http://etd.library.pitt.edu/ETD/available/etd-05072003-102145/>.
- Denver Dash. Restructuring dynamic causal systems in equilibrium. In Robert G. Cowell and Zoubin Ghahramani, editors, *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTats 2005)*, pages 81–88. Society for Artificial Intelligence and Statistics, 2005. (Available electronically at <http://www.gatsby.ucl.ac.uk/aistats/>).
- Nir Friedman, Kevin Murphy, and Stuart Russell. Learning the structure of dynamic probabilistic networks. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 139–147. Morgan Kaufmann, 1998.
- Yumi Iwasaki and Herbert A. Simon. Causality and model abstraction. *Artificial Intelligence*, 67(1):143–194, May 1994.
- Judea Pearl and Thomas S. Verma. A theory of inferred causation. In J.A. Allen, R. Fikes, and E. Sandewall, editors, *KR-91, Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, pages 441–452, Cambridge, MA, 1991. Morgan Kaufmann Publishers, Inc., San Mateo, CA.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. Springer Verlag, New York, NY, USA, second edition, 2000.

Appendix A. Pot-luck challenge: FACT SHEET .

Repository URL: <http://www.causality.inf.ethz.ch/repository.php?id=16>

Title: Mixed Dynamic Systems

Authors: Denver Dash, Mark Voortman, Marek Druzdzal

Contact name, address, email and website: Denver Dash, denver.h.dash@intel.com, <http://pittsburgh.intel-research.net/~dhdash>

Key facts:

Simulated time series data of 9 variables based on linear Gaussian models with no latent common causes, but with multiple dynamic processes at varying time-scales.

Training Data: 9 Variables, 10000 time series, each time series sampled at 12 distinct times (relative to when exogenous variables were first manipulated).

Testing Data: Manipulation data. Each of the 6 non-exogenous variables is manipulated and held fixed for the duration of the time series. This is repeated 100 times for each of the 6 variables.

Summary: A Mixed Dynamic System is one that consists of multiple dynamic processes operating at widely different time-scales. This data represents a 9 variable (labeled $X_1 \dots X_9$) dynamic system with several dynamic processes acting on qualitatively different time scales from one another. The goal is to learn a causal model of the system with the training data, and then correctly predict the effects of various manipulations on the system (using the testing data for a quantitative measure of performance). This dataset was meant to be both simple and extremely challenging. All relations are linear with independent Gaussian error terms. There are no hidden confounders. However, we believe the inter-related dynamic processes will make prediction of manipulations challenging.

Training Data: The training data consists of 9 tab-separated text files (labelled X_1 .tsv, X_2 .tsv, etc.) one for each variable, and is arranged so that the rows in each file represent distinct time series for each variable (there are 10000 of these). That time series has been sampled at a few points in time after the exogenous variables of the system have been manipulated (all exogenous variables are held fixed for the duration of the time series). Specifically, the variables have been measured at the following discrete time intervals: $t = \{1, 2, 3, 4, 5, 50, 100, 500, 1000, 2000, 4000, 10000\}$, so there are 12 columns in each data file. Variables X_8 , X_5 and X_9 are all exogenous as can be verified by looking at X_9 .tsv, etc.

Test Data: The test data is organized into several ($6 \times 9 = 54$) data files labeled X_i -manipj.tsv (For example X_2 -manip3.tsv shows the values of variable X_2 when X_3 has been manipulated and held fixed). Each variable in the set of endogenous variables $\{X_1, X_2, X_3, X_4, X_6, X_7\}$ is manipulated 100 times for the entire 10000 time-step duration of each time series while the remaining variables are measured once at each of the 12 predetermined time-intervals. Thus each X_i -manipj.tsv file has 100 rows and 12 columns, and there are 9 files for each variable manipulated from the set $\{X_1, X_2, X_3, X_4, X_6, X_7\}$.

Evaluation: The objective of this problem is to use the first set of data labeled X^* .tsv to build a model which is then able to predict the effects of manipulation on the system as given by the X^* -manipN.tsv files. When predicting the effect of the manipulations, the goal is to predict the values of non-manipulated variables at times 5–10000 (columns 5 – 12) using the values of the previous times as input. For example, when predicting time 100 (column 7), you could use times 1, 2, 3, 4, 5, 50 (columns 1-6) as input. The output of the evaluation should be one table for each variable in the set $\{X_1, X_2, X_3, X_4, X_6, X_7\}$ of manipulated variables. Each table should have 5 rows and 8 columns, one row for each variable in $\{X_1, X_2, X_3, X_4, X_6, X_7\} \setminus X_j$, (where X_j is the manipulated variable), and one column for each time in the set $\{5, 50, 100, 500, 1000, 2000, 4000, 10000\}$. The entry of the table is the RMS error (over the 100 runs) between the predicted value of the variable at that time and the actual value in the test data.