
Difference-Based Causal Models: Bridging the gap between Granger causality and DCMs

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

Denver Dash*
Intel Labs Pittsburgh
Pittsburgh, PA, 15213, USA

Marek J. Druzzel†
Decision Systems Laboratory
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA, 15260, USA

Dean Pomerleau‡
Intel Labs Pittsburgh
Pittsburgh, PA, 15213, USA

Gustavo Sudre
Center for the Neural Basis of Cognition
Carnegie Mellon University
Pittsburgh, PA, 15213, USA

In recent years, the study of causality in the brain, or *effective connectivity*, has become increasingly important in the field of computational neuroscience. Neuroimaging modalities such as EEG, fMRI and MEG provide an immense and rich source of data. Causal analysis has the potential to help make sense of this complex data, and provide important insights about the neural basis of cognition. Two of the predominant methods for this task are Granger causality [Granger, 1969] and Dynamic Causal Modeling (DCM) [Friston et al., 2003]. Stated simply, DCM is the practice of modeling neurological activation with a set of continuous-time differential equations and using data to fit the coupling parameters of the model. By contrast, Granger causality is a more exploratory method which seeks to explain the variance of variables at time t by conditioning on other variables at times $\tau < t$.

Because of their detailed nature, DCMs make strong theoretical assumptions about how macroscopic interactions in the brain impact sensor measurements, and can therefore account for complex phenomena which would otherwise confound the measurement of causal connectivity, such as the haemodynamic response function causing random delays in fMRI measurements [see Friston, 2009, for a discussion of how Granger causality and DCMs compare in this context]. Such models typically require extensive background knowledge; however, even when some underlying mechanisms are not known for certain, DCMs can be considered as plausible hypotheses whose validity can be tested through aptly designed empirical studies. Granger causality (which is a specific instance of vector autoregression (VAR)), on the other hand uses temporal precedence and linear correlation to provide a data-driven method to search for qualitative causal relations. This approach assumes linearity, which is a strong assumption, but is often a good approximation even when the underlying system is nonlinear. Furthermore, for the purposes of detecting causality, having the correct form of the functional relationship between two variables is perhaps not critical. The ability to discover effective connectivity with little or no background knowledge about the underlying mechanisms is appealing, especially in light of the massive amounts of data available and the constant innovation in new sensing modalities which are less understood than more established techniques such as fMRI or EEG. Furthermore, this technique has proven useful for studying the neural correlates of cognitive states [c.f., Gaillard et al., 2009].

In this paper we argue that, despite some positive features, Granger causality and VAR methods are not well-suited for modeling systems which are based on differential equations, due to the way causality across time is represented. We show that, in principle, even the simplest physical dynamical systems such as the simple harmonic oscillator can result in infinite-order Markov models when Granger causality is applied without regard for the underlying dynamics. We present an alternative method, called *Difference-Based Causal Models (DBCMS)* that are based on a formalism introduced by Iwasaki and Simon [1994] and recently proposed by Dash [2003, 2005] for analyzing causality of dynamic systems in equilibrium. The name “DBCMS” was coined recently by Voortman et al. [2008], who also presented a first algorithm for discovering them from data. DBCMS have many of the desirable properties of Granger causality models, but are based around the assumption that the underlying system is governed by differential equations, and so all causation across time is due to integration of a variable whose derivative is being caused instantaneously by some other variables. Because of this, we show that DBCMS can result in much more parsimonious (and accurate) models when applied to systems which are in fact driven by differential equations.

Stated briefly, a DBCMS is a discrete-time model consisting of a set of variables and a set of equations, with each variable being specified by an equation. The defining characteristic of a DBCMS is that all causation across time is due to a derivative (e.g., \dot{x}) causing a change in its integral (e.g., x). Equations describing this relationship are called *integral equations* and are deterministic. All other causation is *contemporaneous*, occurring between two variables on a time-scale that is smaller than the time-step of the model. In this paper, we perform tests of independence that assume all non-integral equations are linear and contain independent normal error terms, although in principle those could be replaced with nonlinear or distribution-free tests to generalize the applicability of our method. As with Granger causality and VAR models, DBCMS are nearly a subset of causal models as defined by Pearl [2000] and structural equation models similar to those discussed 50 years ago by Strotz and Wold [1960]. Granger causality and VAR models allow arbitrary edges to exist across time; however, we argue that for many real physical systems this representation is too general. DBCMS, by contrast, assume that all causation works in the same way as in mechanical systems, i.e., all causation *across time* is due to integration. This restriction represents a tradeoff between expressibility and tractability. On the one hand, DBCMS are able to represent all systems of differential equations up to arbitrary order, including all systems governed by Newtonian mechanics. On the other hand, at least mathematically, there exist systems that cannot be cast into the DBCMS representation.

* Also adjunct faculty in the Department of Biomedical Informatics, School of Medicine, University of Pittsburgh.

† Also faculty of Computer Science, Bialystok Technical University, Wiejska 45A, 15-351 Bialystok, Poland.

‡ Also adjunct faculty in the Robotics Institute, Carnegie Mellon University.

Our algorithm for learning DBCMs, which we call *DBCML Learner (DBCML)* is in essence a constraint-based causal discovery algorithm (although in principle any structure discovery method could be used as the core; our implementation is based around the PC algorithm of Spirtes et al. [2000]). However, DBCML is novel in that it takes advantage of many unique features of DBCMs to make the structure search more efficient and accurate. First, DBCML treats differences (discrete analogues of derivatives) as latent variables and conducts an efficient search to find them. DBCML thus exploits the well-known fact that accurately detecting the presence of latent variables can lead to much more parsimonious models. However, unlike other algorithms that attempt to detect general latent variables, when detecting latent differences, we have much a priori knowledge about the hidden variables we are searching for: we can estimate their values from the original data, and the structural relationship of the hypothesized latents are known with respect to their corresponding integral. Also important is the fact that once all the relevant differences have been discovered, DBCMs are guaranteed to be first-order Markovian models *and* all cross-temporal edges are known a priori; thus the search space for structure is contained within a given time-slice, dramatically simplifying the search for edges. We prove that DBCML is correct given a conditional independence oracle, and show empirically that it is also efficient and robust in the sense that it avoids calculating higher-order derivatives unless they are required by the model, thus avoiding unnecessary computation and mistakes due to numerical errors that arise with higher-order differences.

As an example motivating the advantages of DBCMs, consider the set of equations describing the motion of a damped simple harmonic oscillator shown on the left of Figure 1: a block of mass m is suspended from a spring in a viscous fluid. The harmonic oscillator is an archetypal dynamic system, ubiquitous in nature and, although linear, is often a good approximation to nonlinear systems close to equilibrium. Like all mechanical systems, the equations of motion for the harmonic oscillator are given by Newton’s 2nd law describing the acceleration a of the mass under the forces (due to the force of gravity $F_g = mg$, due to the spring, F_x , and due to viscosity, F_v) acting on the block. These forces *instantaneously* determine a ; furthermore, they indirectly determine the values of all integrals of a . Writing this continuous time system as a discrete time model, v and x are approximately determined by the difference equations that represent integration across time: $v^{t+1} = v^t + a^t \Delta t$ and $x^{t+1} = x^t + v^t \Delta t$. Treating these equations as a structural equation model across time results in the causal graphs shown in the center of Figure 1. When confronted with data that has not made all relevant derivatives explicit, the distinction between DBCMs




Figure 1: Left: The damped simple harmonic oscillator. Center: The DBCM for this system is always first-order Markovian. Right: The causal graph learned by the perfect Granger causality learner (with all v ’s and a ’s marginalized out) is infinite-order Markovian.

and Granger causality becomes glaring. Rather than trying to search for the derivatives that make the model first-order Markovian, Granger causality would simply learn the structure of the system with those latent variables marginalized out. One might suspect that there is not much difference. For example, one might expect that a second order differential equation would simply result in a second-order Markov model when the derivatives are marginalized out. Unfortunately, that is not the case, because the causation among the derivatives forms an infinite chain into the past resulting in the infinite-order Markov model on the right of Figure 1. Thus, the benefits of using the DBCM representation are not merely computational, but in fact without learning the derivatives directly, the system does not even have a finite representation.

We conjecture that most brain sensing modalities are indeed governed by differential equations, and thus DBCMs are more appropriate than Granger causality or VAR models. This is a conjecture that is supported by the success of DCMs, which are also based on differential equations. At the same time, as with DCMs, it should be relatively straightforward to incorporate background information into the DBCML procedure by adding known equations to the system prior to learning structure. Thus DBCMs in principle provide a “dial” by which we can smoothly transition between the exploratory nature of Granger causal models to the more knowledge-rich DCMs. We have begun to test the applicability of DBCMs on alpha wave detection in EEG data. Although preliminary, these results agree with the well-known fact that alpha waves are generated in the occipital lobes. Additionally, our results seem to imply that alpha wave activity under waking conditions with eyes closed are transmitted nearly instantaneously to other regions of the brain. Our hypothesis is that these signals are primarily due to conductivity effects of the skull. We have also begun exploratory analysis of effective connectivity in MEG data during simple visual and motor tasks, with encouraging results. While preliminary, these results provide evidence that DBCMs may be useful for modeling causality in the brain. More work needs to be done to compare and contrast their benefits relative to Granger causality and DCMs.

References

Denver Dash. *Caveats for Causal Reasoning with Equilibrium Models*. 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/>).

Karl Friston. Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS Biology*, 7(2):220–225, February 2009.

K.J. Friston, Harrison L., and Penny W. Dynamic causal modeling. *NeuroImage*, 19:1273–1302, 2003.

Raphaël Gaillard, Stanislas Dahene, Claude Adam, Cléopâtre Clémenceau, Dominique Hasboun, Michel Baulac, Laurent Cohen, and Lionel Naccache. Converging intracranial markers of conscious access. *PLoS Biology*, 7(3):472–492, 2009.

Clive W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, July 1969.

Yumi Iwasaki and Herbert A. Simon. Causality and model abstraction. *Artificial Intelligence*, 67(1):143–194, May 1994.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2000.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. The MIT Press, Cambridge, MA, second edition, 2000.

Robert H. Strotz and H.O.A. Wold. Recursive vs. nonrecursive systems: An attempt at synthesis; Part I of a triptych on causal chain systems. *Econometrica*, 28(2):417–427, April 1960.

Mark Voortman, Denver Dash, and Marek Druzdzal. Learning causal models that make correct manipulation predictions with time series data. In *Workshop on Causality, in conjunction with the 22nd annual conference on Neural Information Processing Systems (NIPS)*, 2008.