

Searching the DCM Model Space, and Some Generalizations

Michael Freenor and Clark Glymour

April 30, 2010

Technical Report No. CMU-PHIL-185

Philosophy

Methodology

Logic

Carnegie Mellon

Pittsburgh, Pennsylvania 15213

Searching the DCM Model Space, and Some Generalizations

Michael Freenor and Clark Glymour

Department of Philosophy, Carnegie Mellon University

Address for correspondence:

Michael Freenor

Baker Hall 135

Carnegie Mellon University

Pittsburgh, PA 15213-3890

Fax number: 412.268.1440

Telephone: 412.268.8568

Email address: mfreenor@andrew.cmu.edu

Abstract

We describe the (enormous) size of the search space for Dynamic Causal Models and generalizations of them. We describe a greedy heuristic search for modulation relations, and give evidence of its reliability by showing it gives the same results as does an exhaustive search using an SPM8 example. Using that example and the SPM8 Bayes factor model comparison criteria, we compare “conventional” DCM models with models in which input variables can also have modulation effects, and give another empirical illustration. Using the SPM8 example, we also consider the statistical result of allowing models in which any variable, whether designated “input” or “ancillary” or “contextual” can both directly influence regions of interest and modulate effective connections. We conclude by considering what these results may mean for finding and interpreting the best fitting DCM model or generalization.

Introduction

The Dynamic Causal Model (DCM) framework for interpreting fMRI contrasts requires the specification of a detailed model, which is aided by SPM8 software. Penny, et al. (2004) have called for machine learning approaches to DCM model specification, and the SPM8 manual lists comparison of alternative models by Bayesian model selection in its counsels for using DCM models in fMRI studies. Specification of DCM models has at least two aspects: the specification of a directed graphical model describing the causal connectivity relations between inputs and regions of interest and between regions of interest, and the specification of which of these connections are modulated by ancillary or contextual variables. While the SPM8 manual and supporting literature (Friston, *et al.*, 2003) on DCM model specification separates, in practice, exogenous variables that directly affect regions of interest (“input”) from exogenous variables that modulate the strengths of effective connections (“ancillary,” “contextual”), that separation is sometimes violated in applications that posit models in which an input variable influences a region of interest directly and also modulates one or more effective connections. To our knowledge, the statistical advantages or disadvantages of allowing such multiple roles for exogenous variables is unexplored. Ramsey, et al., (2010) have proposed and validated on simulated data a procedure (“IMAGES”) for identifying feed-forward graphical structure, but search for feedback relations remains a matter of trial and error using Bayes factor comparisons. To our knowledge, no search method for modulation relations has been proposed and validated except trial and error search using Bayes factors.

In what follows, we describe the (enormous) size of the search space for DCM models and generalizations of them. We describe a greedy heuristic search for modulation relations, and, using an SPM8 example, we give evidence of its reliability by showing it gives the same results as does an exhaustive search. Using that example and the SPM8 Bayes factor model comparison criteria, we compare “conventional” DCM models with models in which input variables can also have modulation effects, and give another empirical illustration. Finally, using the SPM8 example, we consider the statistical result of allowing models in which any variable, whether designated “input” or “ancillary” or “contextual” can both directly influence regions of interest and modulate effective connections. We conclude by considering what these results may mean for interpreting the best fitting DCM model.

The Structure of Dynamic Causal Models

Dynamic causal models represent a system of ordinary differential equations. The neural state equation expresses the rate of change of the system in terms of the current state and its interaction with the inputs.

$$\dot{z} = (A + \sum u_j B^j)z + Cu$$

A , B , and C are matrices representing intrinsic connectivity (effective connections between ROIs), modulation of this connectivity due to different inputs (B^j is the modulation matrix for the j 'th input), and the main-effects from inputs respectively. These matrices represent the rate of change of the neural state with respect to the state itself (A), the rate of change of the neural state with respect to the j 'th input and the state (B), and the rate of change of the neural state with respect to the inputs directly (C), as shown in the above diagram from the SPM8 manual.

Specifying a DCM requires the specification of the A , B , and C matrices. An illustration from the SPM8 manual based on data from Buchel and Friston (1997) is shown in Figure 1.




Figure 1: "mod_bwd", the first a priori specified model examined in the SPM8 manual for the Friston & Buchel data.

If we allow the rows and columns of the A and B matrices to denote V1, V5, and SPC in order, and if we interpret a "1" at column i , row j as "there is an edge from i to j ", then we can write out Figure 1 in matrix form. Furthermore, number the inputs as follows: 1 – photic, 2—motion, 3—attention.

$$\begin{array}{cccc}
A = \begin{array}{ccc} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{array} & B^1 = \begin{array}{ccc} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{array} & B^2 = \begin{array}{ccc} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{array} & B^3 = \begin{array}{ccc} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{array} \\
C = \begin{array}{ccc} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{array} & & &
\end{array}$$

The columns in the C matrix are to be interpreted as the inputs, in the order specified above, while the rows remain the same ROIs. Each B^l entry specifies a modulating edge (or absence) from input i . The DCM framework assumes a Gaussian shrinkage prior distribution over these and other parameters representing the hemodynamic response, and uses an iterative expectation-maximization scheme for estimation.

Evaluation and Comparison of Dynamic Causal Models

Let Θ denote the parameters of model m with respect to data y . Following the SPM manual, we can write the likelihood of these parameters as:

$$p(\theta|y, m) = \frac{p(y|\theta, m)p(\theta|m)}{p(y|m)}$$

$p(y|m)$ is the likelihood of the data given the model, which acts as a normalization term. This is called the “model evidence”, and can be written as:

$$p(y|m) = \int p(\theta|y, m)p(\theta|m)d\theta$$

This integral is typically not solvable analytically, and a variety of approximations to the model evidence exist, including the BIC, AIC, and F (negative free-energy) approximations (Penny et al. 2004). We will use F, which is expressed as the following balance between accuracy and complexity:

$$\ln p(y|m) = \text{accuracy}(m) - \text{complexity}(m)$$

where

$$\text{accuracy}(m) = \left[-\frac{1}{2} \ln |C_e| - \frac{1}{2} (y - h(u, \eta_{\theta|y}))^T C_e^{-1} (y - h(u, \eta_{\theta|y})) \right]$$

$$\text{complexity}(m) = \left[\frac{1}{2} \ln |C_p| - \frac{1}{2} \ln |C_{\theta|y}| - \frac{1}{2} (\eta_{\theta|y} - \theta_p)^T C_p^{-1} (\eta_{\theta|y} - \theta_p) \right]$$

Θ is the model parameters, y the data, h the prediction function, $\eta_{\theta|y}$ the conditional probability of the parameters given the data, $C_{\theta|y}$ the conditional covariance of the parameters given the data, and C_{θ} , C_p , and θ_p the error covariance, prior covariance, and prior mean of the parameters respectively.

The Bayes factor between models m_i and m_j , is:

$$BF_{ij} = \frac{p(y|m_i)}{p(y|m_j)}$$

Assuming the models are exclusive and exhaustive, the Bayes factor can be mapped to relative posterior probability estimate of the models. The model with higher posterior is of course preferred, other things equal, but Penny, et al., (2004) suggest preferring the more probable model only if the ratio of probabilities is equal to or greater than the base of the natural logarithms.

In the SPM8 manual, model comparison is illustrated with the DCMs in Figures 1 and 2.




Figure 2: “mod_fwd”, the preferred a priori specified model discussed in the SPM8 manual, selected by comparing its model evidence against mod_bwd.

The ratio of the probabilities of the model of Figure 2 to the model of Figure 1 given in the SPM8 manual is roughly 0.8/0.2.

DCM Search Spaces

Given a set of ROIs, stimulus inputs and ancillary variables in an experimental setup, possible DCMs can differ in their intrinsic connections (A -matrix), their modulatory connections (B -matrix), or their main-effects from inputs to ROIs (C -matrix). In what follows we will consider violating the separation of

variables into C matrix (input) variables and B matrix (ancillary) variables. We describe the search space of these generalized DCM models as a function of the number of regions of interest.

Let there be n ROIs and k inputs, not distinguishing between experimentally controlled inputs and ancillary variables. There are n^2 possible directed edges, hence $\binom{n^2}{i}$ directed graphs with i edges, or equivalently non-zero entries in the A matrix for n ROIs. For any particular configuration of the A-matrix with i edges there are 2^{ki} possible sets of modulation relations given by the B-matrix. There are 2^{kn} possible configurations of the C-matrix.

This yields the following measure of generalized DCM possibility spaces,

$$2^{kn} \sum_{i=0}^{n(n-1)} \binom{n(n-1)}{i} 2^{ki}$$

For many interesting problems, the number of mathematically possible generalized DCMs is enormous, and exhaustive searches are obviously not feasible. This space of possibilities for the B-matrix is reduced by adhering to the conventional separation of input and ancillary variables. With r input variables and a disjoint set of s modulation variables, the function of n changes to:

$$2^{rn} \sum_{i=0}^{n(n-1)} \binom{n(n-1)}{i} 2^{si}$$

Substantive prior knowledge may of course reduce these dimensions of search. In the case of the Buchel and Friston, (1997) data that is analyzed in the SPM8 manual $n = 3$ and $k = 3$, meaning there are 272,097,792 hypothetical models for this data expressible in the framework of DCM.

Modulation Search

Figure 2 shows the DCM model preferred to the model of Figure 1 in the SPM8 manual, assuming the directed graph and applying a trial and error search for modulation effects using the SPM8 model selection criteria. An exhaustive search over configurations of possible modulation relations, under the SPM8 separation of input (photic) and modulation (motion, attention) variables, finds that the most probable model is that in Figure 3.




Figure 3: result of an exhaustive search over B-matrix entries from motion and attention, starting from the A and C-matrix of mod_fwd. The posterior probability of this model versus mod_fwd is 1.0/0.0 according to SPM8 BMS.

The model of Figure 3 is almost saturated: only one possible modification relation is omitted. In fact, starting with no modulation effects, the addition of further modifications increases the posterior probability estimate almost monotonically until the most probable model is found. The histogram of models as a function of differences of each possible model's posterior probability to mod_fwd's (Figure 2) in direct pair-wise comparison is plotted in Figure 4.




Figure 4: Histogram of the posterior probabilities of modulation models. 0 marks the SPM8 preferred model of figure 2, mod_fwd. All models are compared to this model. Models to the right of 0 are deemed superior to mod_fwd by SPM8.

Exhaustive search for modulation effects is generally infeasible even under the SPM8 restriction of possible modulators. Consequently, we propose the following simple heuristic search for modulation relations: the first stage considers each possible one edge addition from possible modulators to edges of the effective connectivity graph, and adds the modulation edge that has the largest probability difference with the original graph. At stage n , the single edge that most improves the probability compared to the graph obtained at stage $n - 1$ is added. The search stops when no further improvement in posterior probability is obtained. The result of this procedure for the SPM8 example, beginning with the SPM8 A and C matrices, is the same as the result of the exhaustive search in figure 3. We find in almost all of the cases examined in this paper that the most probable model instantiates more than half of the possible modulation relations allowed, suggesting that it would be more efficient

to run the heuristic search in reverse, starting with all possible modulation relations and eliminating them one by one in a greedy search.

The heuristic search for modulation effects could conceivably find a sub-optimal model by the SPM8 model comparison criterion if, for example, edge α is preferred to edge β , and with edge α added, edge γ is preferred at the next (or later) step to edge δ , but the addition of edges β and δ yields a more probable model than the addition of edges α and γ . We have no proof that, within the DCM framework which uses Gaussian shrinkage priors on parameters, no such errors can occur, but they clearly do not in the SPM8 example.

Violating B-Matrix Conventions

In SPM8 and associated literature, input variables are not afforded B-matrix entries for modulator relations. Schuyler, et al. (2010) violate this convention but give no comparison with models that accord with it. Ramsey, et al., have provided a feed-forward directed graphical model for data from an experiment (Xue and Poldrack, 2007) in which subjects are shown rhyming and non-rhyming words and non-words in 4 blocks. We add feedback for each feed-forward edge, as shown in Figure 5, and compare DCM models based on that graph in which the input does, and alternatively does not, also modulate effective connections. We consider only a single subject (subject 1) from among thirteen.

The input variable for this example was constructed from the onset vector of the experimental condition. It is binary – “on” when the experimental condition is present, “off” when it’s not, and lagged by two time-steps to synchronize it roughly with the hemodynamic response. Our iMAGES result, run across all thirteen subjects as in Ramsey, *et al.* (2010), is identical to Ramsey’s result, which uses a convolved input.

There are 2^{14} possible modulation models, and exhaustive search is impossible. The best modulation model obtained with our heuristic search is shown in Figure 6. The relative posterior probabilities for the comparison of the models in Figure 5 and Figure 6 are 0 and 1, respectively.




Figure 5: DCM of the Xue & Poldrack data, respecting the notion that driving inputs do not modulate.




Figure 6: result of the greedy modulation search applied to figure 5m allowing the input to modulate. The posterior probability is 1.0 versus figure 5.

In the best model found, the input modulates 9 of the 14 postulated effective connections.

The same method can be applied to the 3 variable SPM8 model, where, allowing photic, motion and attention to modulate freely, the space of possible modulation models is $2^{12} = 4098$ and can be searched exhaustively and compared with the heuristic search. The two searches agree on the best model, shown in Figure 7. Figure 7, when compared to Figure 3, yields a posterior probability of 1.




Figure 7: result of both an exhaustive search and a greedy modulation search, starting from the A and C-matrices of mod_fwd, allowing photic to have B-matrix entries as well.

Violating B and C-Matrix Conventions

In the conventional DCM framework, ancillary variables that have non-zero entries in the B matrix are not (typically) allowed entries in the C matrix of inputs. That restriction can be violated as well. In the SPM8 example, motion and attention are logically related to photic input. If the Ramsey, et al. search is run with all variables, motion and attention are screened off from the regions of interest by photic. Alternatively, the Ramsey search can be run separately with photic as the only input, with motion as the only input, and with attention as the only input. Motion and attention are then found to directly influence V5.

Combining the edges of the three graphs thus obtained, supplemented with feedback for the effective connections, and then carrying out an exhaustive, or heuristic, search for modulation relations, results in the model of Figure 8.




Figure 8: result of unioning iMAGES searches with each controlled variable sequentially, allowing ancillary variables to have main-effects, followed by an exhaustive or heuristic search for modulation relations. The iMAGES searches establish the feed-forward structure; back-projections are added to each edge.

The model of Figure 8 has posterior probability essentially 1 compared to the model of Figure 7, and is the most probable model we have found using the effective connections postulated in the SPM8 models.

The models of Figure 8 is not plausible: we expect influences on V5 to pass through V1. That it is nonetheless more probable than any models that prohibit main influences from ancillary variables is curious, and therefore potentially troublesome.

Violating A, B, and C-Matrix Conventions

An IMAGES search yields the feed-forward structure of the SPM8 model effective connections, but does not yield feedback relations. In the general case, an exhaustive search for feedback models using posterior probability connections would be daunting, but perhaps feasible with some parallelization. In the SPM8 example, however, only one further feedback edge is possible, from SPC to V1. Adding that edge to the effective connections, repeating the three IMAGES searches separately with photic, motion and attention as unique inputs, combining the results, and carrying out an exhaustive and heuristic search for modulation relations yields the model of Figure 9.




Figure 9: result of searching across possible feedback structures. Each possible feedback structure was investigated with a greedy modulation search. These results were compared to find the best, resulting in the above model. Compared to figure 8, the posterior probabilities are 0.8242/0.1758.

Figure 9 and Figure 8, in a heads-up comparison using SPM8, yield posterior probabilities of 0.8242 and 0.1758 respectively. Figure 9 is the most probable model we have found for the Buchel and Friston, (1997) data.

Discussion

Our exploration of the search space of DCM models and generalizations raises a number of fundamental questions to which we have no answers: Are there compelling neurophysiological reasons for strictly separating modulators from inputs, B matrix variables from C matrix variables? By what neurophysiological mechanisms can contextual variables modulate effective connections? Are there neurochemical reasons for thinking that external stimulations may stimulate modulation of many effective connections nearly simultaneously? Why are the most probable models we find so heavily parameterized?

Breaking the conventional separation of roles for ancillary and input variables improves the posterior probability considerably, but does not necessarily result in a plausible model. That may be because in the SPM example, attention and motion are logically connected to photic. When ancillary variables are independent of regions of interest conditional on input variables, we do not recommend allowing the ancillary variables entry into the DCM C matrix, a procedure we used only to illustrate the behavior of the SPM8 model selection criterion. We suggest a useful and applicable criterion for separating C matrix from B matrix variables is this: find the smallest set, I , of exogenous variables conditional on which the remaining exogenous variables are independent of the ROIs. I is the set of variables that influence ROIs; the remaining exogenous variables do not. This separation is easily found with Ramsey's procedure or many other.

In our examples, the model selection criteria favor heavily parameterized models. The most probable model, shown in Figure 9, omits only one possible effective connection, and only 4 of 15 possible modulation connections. Abiding by the separation of variable roles in the conventional DCM formalism yields a most probable model, shown in Figure 3, in which all but one of the possible modulation connections is realized. In the DCM model for the Xue and Poldrack data, the most probable model has the input modulating the majority of effective connections.

The heavy parameterization of a most probable model within some search space may be because the model is true, or truer than the many alternatives examined, but it may also be because the posterior probabilities calculated are an unreliable guide to truth. Posterior probabilities may fail as a guide either because of biases in the prior probability distribution, or because of unsatisfactory approximations in computing posterior probabilities, or because the model search space is inappropriate. In other contexts, it is a familiar fact that misspecification of the model search space often leads to over-parameterization. Merely for one kind of example, if in truth a collection of variables are all confounded by a common cause, and the postulated search space excludes that possibility, the measured variables will typically all be dependent; each pair of measured variables will typically be dependent even conditionally on any other set of measured variables. A complete graph and complete parameterization will result. That is why substantive assumptions about the space of possible models are invaluable—if correct.

Acknowledgments

We thank Russell Poldrack and Stephen Hanson for many helpful conversations, and Karl Friston for valuable comments and suggestions on a draft of this paper. Joseph Ramsey provided both helpful comments and computational assistance. We also thank Guillaume Flandin and Maria Joao for their corrections to the SPM8 code. Research was supported by a grant from the James S. McDonnell Foundation.

References

- C. Buchel and K.J. Friston, (1997). Modulation of connectivity in visual pathways by attention: Cortical interactions evaluated with structural equation modelling and fMRI. *Cerebral Cortex*, 7:768-778.
- Buxton, R.B., Wong, E.C., Frank, L.R., (1998). Dynamics of blood flow and oxygenation changes during brain activation: the Balloon model. *MRM* 39, 855–864.
- D.M. Chickering and C. Boutilier, (2002). Optimal search structure identification with greedy search. *Journal of Machine Learning Research*, 3:507-554.

K.J. Friston, L. Harrison, and W.D. Penny, (2003). Dynamic Causal Modelling. *NeuroImage*, 19(4):1273-1302.

Mandeville, J.B., Marota, J.J., Ayata, C., Zararchuk, G., Moskowitz, M.A., Rosen, B., Weisskoff, R.M., (1999). Evidence of a cerebrovascular postarteriole Windkessel with delayed compliance. *J. Cereb. Blood Flow Metab* 19, 679–689.

Meek, C. (1997). *Graphical Models: Selecting causal and statistical models*. PhD thesis, Carnegie Mellon University.

W.D. Penny, K.E. Stephan, A. Mechelli, and K.J. Friston, (2004). Comparing Dynamic Causal Models. *NeuroImage*, 22(3):1157-1172.

J.D. Ramsey, S.J. Hanson, C. Hanson, Y.O. Halchenko, R.A. Poldrack, and C. Glymour, (2010). Six problems for causal inference from fMRI. *NeuroImage*, 49(2): 1545-1558.

Schuyler B, Ollinger JM, Oakes TR, Johnstone T, and Davidson RJ, (2010). Dynamic Causal Modeling applied to fMRI data shows high reliability. *Neuroimage*, 49(1):603-11.

SPM8 Manual (<http://www.fil.ion.ucl.ac.uk/spm/doc/manual.pdf>)

Xue, G., Poldrack, R., (2007). The neural substrates of visual perceptual learning of words: implications for the visual word form area hypothesis. *J. Cogn. Neurosci.* 19, 1643–1655.