NeuroImage 84 (2014) 986–1006

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Non-Gaussian methods and high-pass filters in the estimation of effective connections

Joseph D. Ramsey *, Ruben Sanchez-Romero, Clark Glymour

Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:
Accepted 25 September 2013
Available online 5 October 2013

Keywords:
Non-Gaussianity
Orientation
Causal networks
fMRI
High-pass filters

ABSTRACT

We consider several alternative ways of exploiting non-Gaussian distributional features, including some that can in principle identify direct, positive feedback relations (graphically, 2-cycles) and combinations of methods that can identify high dimensional graphs. All of the procedures are implemented in the TETRAD freeware (Ramsey et al., 2013). We show that in most cases the limited accuracy of the several non-Gaussian methods in the Smith et al. (2011) simulations can be attributed to the high-pass Butterworth filter used in that study. Without that filter, or with the filter in the widely used FSL program (Jenkinson et al., 2012), the directional accuracies of several of the non-Gaussian methods are at or near ceiling in many conditions of the Smith et al. simulation. We show that the improvement of an apparently Gaussian method (Patel et al., 2006) when filtering is removed is due to non-Gaussian features of that method introduced by the Smith et al. implementation. We also investigate some conditions in which multi-subject data help with causal structure identification using higher moments, notably with non-stationary time series or with 2-cycles. We illustrate the accuracy of the methods with more complex graphs with and without 2-cycles, and with a 500 node graph; to illustrate applicability and provide a further test we apply the methods to an empirical case for which aspects of the causal structure are known. Finally, we note a number of cautions and issues that remain to be investigated, and some outstanding problems for determining the structure of effective connections from fMRI data.

Introduction

The simulation study of Smith et al. (2011) using Dynamic Causal Models (DCMs, Friston et al., 2003) has become a test bed and frequently cited reference for studies of the accuracy of automated methods for estimating "effective connections" between brain regions of interest (ROIs). Testing 35 search methods in 28 simulated conditions each with 50 simulated subjects, Smith et al. found that adjacencies—which regions are causally connected with which others, without specification of directions of influence—can be identified with good to excellent accuracy by several algorithms, a result that has been confirmed by comparison of inferences from human fMRI with experimental animal studies (Dawson et al., 2013). The Smith et al. description of the 28 conditions is given in Appendix A. Correct estimation of the direction of influence is more challenging. Smith et al. found only one method, Patel's tau (Patel et al., 2006), that identified directions with even fair accuracy. None of the methods were able to identify direct feedback cycles (graphically, 2-cycles) or the causal structure of a DCM model generating non-stationary time series.

Since the Smith et al. paper several search methods have been published that exploit the use of multiple data sets (Gates and Molenaar, 2012; Iyer et al., 2013; Ramsey et al., 2011) and/or non-Gaussian features of the BOLD signal (Hyvärinen and Smith, 2013) to determine directions of influence between pairs of variables assumed to be directly connected (relative to the set of regions of interest (ROIs) used in a data analysis). Ramsey et al. (2011) estimated directions of influence with two heuristics based on the assumption that sums of variables should be closer to a Gaussian distribution than the distributions of the summands. Hyvärinen and Smith (2013) introduced several methods based on the LiNGAM model. Using data from 10 simulated subjects at a time, after estimating connections (without direction) with a multisubject Bayesian algorithm, IMaGES, Ramsey et al's (2010) non-Gaussian orientation methods were substantially more accurate than the methods Smith et al. tested with realistic sample sizes and stationary time series, typically finding edges and directing them with both 85% precision and recall for data generated from DCM models whose effective connections specified directed acyclic graphs (DAGs). With their methods, Hyvärinen and Smith found comparable accuracies on the Smith et al. simulations using data from one subject at a time, given the true undirected edges of the generating graph. All of these the Hyvärinen and Smith methods can in principle identify graphs of effective connections with a cyclic structure, provided that the cycles are not direct feedbacks—i.e., not 2-cycles.

These results invite several questions. Is 85% the best that non-Gaussian methods can do with the Smith et al. data, or can the

* Corresponding author.

1053-8119/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.neuroimage.2013.09.062

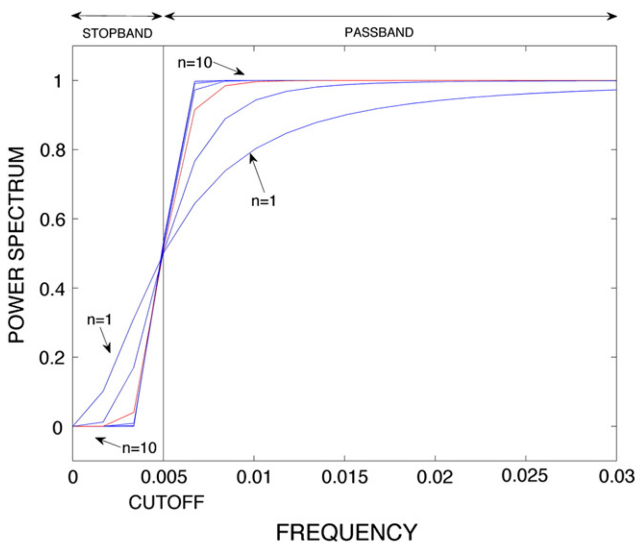

Fig. 1. Power spectrum for various high-pass filters with different orders n and the same cutoff frequency of 1/200 s indicated with a black vertical line. The filter with an order n = 4 is shown in red.

accuracies be improved? Can the methods identify effective connections represented in denser or higher dimensional graphs? For which methods in which conditions does the use of multiple subject data improve accuracy? Can the published or other non-Gaussian methods identify 2-cycles? Can non-Gaussian methods identify the generating effective connection structure from non-stationary time series? These are the issues we address. We show that the Smith et al. simulations contain a filtering step that handicaps non-Gaussian methods. Once that filter is removed, or replaced by a less stringent filter, with appropriate measures of non-Gaussianity several of the published non-Gaussian methods are at or near ceiling for both precision and recall in data from DCM models essentially identical to that of many conditions in Smith et al. We offer further simulations showing that the methods are robust with denser graphical generating structures. We show by simulation that combining second moment methods for determining causal connections (graphically: adjacencies between variables) with non-Gaussian methods for determining directions of influence allows excellent precision and mediocre recall for high dimensional (500 ROIs) models. We show that there are other methods inspired by the LiNGAM model that give good, although less than perfect, results with data from the Smith et al. stimulations that are non-stationary or from structures with 2-cycles.

A puzzle

A note in a webpage ancillary to the Smith et al. paper contains the following remark:

> "We (Smith et al.) recently (24/08/2012) found that the highpass temporal filtering originally applied to Sim19 and Sim20 was more aggressive than intended. Whereas, for all other simulations, the simulated data was conservatively highpass filtered at a cutoff frequency of 1/200 s (to simulate preprocessing of fMRI data), this cutoff had originally been set at a fixed fraction of Nyquist, and hence for Sim19/20 (with the lower TR = 0.25 s) was unrealistically aggressive, at 1/17 s. We have now re-run Sim19/20, with the highpass cutoff now set to the intended 1/200 s. We found no significant differences in the bottom-line results, except for somewhat improved performance (in directionality estimation) by Patel's tau and LiNGAM."[1]

[1] http://www.fmrib.ox.ac.uk/analysis/netsim/correction.html.

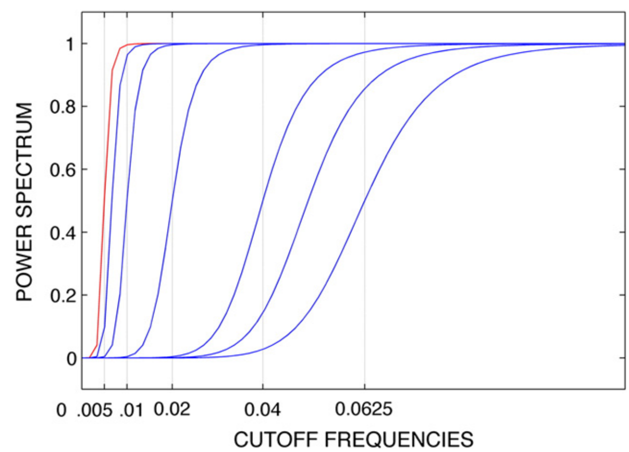

Fig. 2. Power spectrum for various high-pass filters with different cutoff frequencies and the same order (n = 4). The filter with a cutoff of 1/200 s is shown in red.

High-pass filters are used to remove biological and behavioral contributions of low frequency to the fMRI time series (Kruggel et al., 1999; Power et al., 2012). Here is the puzzle: Why did changing the high-pass filter improve the accuracy of LiNGAM and Patel's tau but not the accuracies of the other methods? In what follows we propose an answer to the question and show that the answer allows for considerably increased accuracy in estimating directions of effective connections from fMRI data.

Except for comparisons in which simulated data is not high-pass filtered, all of our reanalyses of the Smith et al. simulations are carried out with exactly their specifications of graphical structure, linear coefficient parameter ranges, sample sizes, sampling rates (TR values) and disturbance terms. The only difference, explained in Appendix B, is that we have allowed more random variation in coefficient parameter values across subjects.[2]

What high-pass filters do

All conditions in the Smith et al. simulations used a Butterworth high-pass temporal filter to emulate preprocessing of fMRI data to reduce low frequency artifacts in the signals. The behavior of a Butterworth filter can be controlled using a cutoff frequency, which sets the division between the passband (the frequencies that will be allowed) and the stopband (the frequencies that will be suppressed by the filter); and by changing the order of the filter, which controls the sharpness of the separation of the passband and stopband.

Butterworth filters are characterized by being maximally flat in the passband (i.e. almost all the frequencies in the passband are affected by the same factor) and monotonic in the passband and stopband. In frequency space, the power spectrum of the Butterworth high-pass filter can be expressed as a function of the frequency,

$$|H(\Omega)|^2 = 1 - \left[1/\left(1 + (\Omega/\Omega_c)^{2n}\right)\right]$$

where Ω is the frequency, Ω_c is the cutoff frequency and n is the order of the filter. Except as noted in the passage quoted above, in the Smith et al. simulations a Butterworth high-pass filter with an order of n = 4 and a

[2] The Smith et al. study also shows evidence of an additional influence on non-Gaussianity. Conditions 2 and 17 are exactly the same except from the level of Gaussian noise: ~N(0,1) and ~N(0,0.1) respectively. This implies that the BOLD signals of condition 17 are more non-Gaussian than those of condition 2. The impact of the differences in Gaussianity of both signals can be seen in the Smith et al. results showing better accuracy for condition 17 for LiNGAM and Patel's tau.