

SCIENTIFIC REPORTS



OPEN

Change Point Detection in Correlation Networks

Ian Barnett & Jukka-Pekka Onnela

Received: 02 June 2015

Accepted: 21 October 2015

Published: 07 January 2016

Many systems of interacting elements can be conceptualized as networks, where network nodes represent the elements and network ties represent interactions between the elements. In systems where the underlying network evolves, it is useful to determine the points in time where the network structure changes significantly as these may correspond to functional change points. We propose a method for detecting change points in correlation networks that, unlike previous change point detection methods designed for time series data, requires minimal distributional assumptions. We investigate the difficulty of change point detection near the boundaries of the time series in correlation networks and study the power of our method and competing methods through simulation. We also show the generalizable nature of the method by applying it to stock price data as well as fMRI data.

Many systems of scientific and societal interest are composed of a large number of interacting elements, examples ranging from proteins interacting within each living cell to people interacting with one another within and across societies. These and many other systems can be conceptualized as networks, where network nodes represent the elements in a given system and network ties represent interactions between the elements. Network science and network analysis are used to analyze and model the structure of interactions in a network, an approach that is commonly motivated by the premise that network structure is associated with the dynamical behavior exhibited by the network, which in turn is expected to be associated with its function. In many cases, however, network structure is not static but instead evolves in time. This suggests that given a sequence of networks, it would be useful to determine points in time where the structure of the network changes in a non-trivial manner. Determining these points is known as the network change point detection problem. Given the connection between network structure and function, it seems reasonable to conjecture that a change in network structure may be coupled with a change in network function. Consequently, detecting structural change points for networks could be informative about functional change points as well.

In this paper, we consider the change point detection problem for correlation networks. These networks belong to a class of networks sometimes called similarity networks and they are obtained by defining the edges based on some form of similarity or correlation measure between each pair of nodes¹. Examples of correlation networks appear in many financial and biological contexts, such as stock market price and gene expression data^{2–6}. In general, when evaluating correlation networks, the full data is used to estimate the correlations between the nodes. When using this approach for longitudinal data, it is sometimes implied that the network structure is the same over time. This assumption may however be inaccurate in some cases. For example, in Onnela *et al.*² a stock market correlation network is created from almost two decades of stock prices. In reality the relationship between the stocks, and therefore the structure of the underlying network, likely change over such a long period of time, an issue that they addressed by dividing the data into shorter time windows. Similarly, in functional magnetic resonance imaging (fMRI) trials it is likely that the brain interacts differently during different tasks⁷, or possibly even within a given task, so it may be inaccurate to assume a constant brain activity correlation network in trials with multiple tasks.

Suppose that a network is constant or may be assumed so until a known point in time before undergoing sudden change. The alternative, where network change occurs gradually, is not considered here. In the case of sudden change, the underlying data should be split up at the change point into two parts, and two separate correlation networks should be constructed from the two subsets of the data. In reality the location of the change point, or possibly several change points, is not known *a priori* and must also be inferred from the data. This problem belongs to a wider class of so-called change point detection problems, which has been studied in the field of process control. When the observed node characteristics are independent and normally distributed, methods exist for general time series data to detect changes in the multivariate normal mean or covariance^{8–10}. Dropping the normality assumption, there are also univariate^{11–13} and multivariate^{11,14,15} change point detection methods for shifts in the

Harvard University, Department of Biostatistics, Boston, 02115, USA. Correspondence and requests for materials should be addressed to J.-P.O. (email: onnela@hsph.harvard.edu)

mean of time series. We also consider a non-parametric change point test statistic proposed for detecting changes in the covariance in the multivariate case¹⁶.

There have been some promising efforts at change point detection for structural networks, but in this case the actual network is observed over time rather than relying on correlations of node characteristics that are used to construct the network^{17–22}. If a network is first inferred from correlations, then these methods could be applied. However, inferring networks from correlations is not trivial. When thresholding correlations to determine the adjacency matrix, which is a common approach, the inferred networks tend to be highly sensitive to the chosen threshold. For this reason, these methods cannot be directly applied to correlation networks without first solving the problem of inferring the correlation networks themselves. Therefore, despite this large body of methods developed for change point detection for both time series data and for networks, there is a need for a change point detection method specifically for correlation networks that is not hampered by stringent distributional assumptions.

In this paper we propose a computational framework for change point detection in correlation networks that has minimal distributional assumptions. This framework offers a novel and flexible approach to change point detection. Change point detection methods suggested by Zamba *et al.*^{10,23} and Aue *et al.*¹⁶ are adapted to our framework and power to detect change points is compared to our method using simulation. Also, we investigate the general difficulty of change point detection near the boundaries of the data both analytically and through simulation. Finally, we apply our framework to both stock market and fMRI correlation network data and demonstrate its success and limitations for detecting functionally relevant change points.

Results

The relationship between T and n for statistical power. Consider a system with a fixed number of $i = 1, \dots, n$ nodes with characteristics, Y_{ij} , observed at $j = 1, \dots, T$ distinct time points. For example, nodes could be stocks with their daily log returns as node characteristics. If the characteristics of a pair of nodes covary, this may indicate the presence an edge between these two nodes in the underlying network. The n by n covariance matrix of all the nodes taken over the time interval from time t_1 to t_2 is defined as $S(t_1, t_2)$. Our approach is to detect a change point at time k if the covariance prior to that point, $S(1, k)$, is significantly different from the covariance after that point, $S(k + 1, T)$, as measured by the Frobenius norm: $d(k) = \|S(1, k) - S(k + 1, T)\|_F^2$. Significance is determined through simulation of the null distribution of $d(k)$ using a bootstrap resampling procedure (see Methods).

Estimation of the covariance matrix requires T to be large relative to n because the empirical covariance matrix has $n(n - 1)/2$ elements that need to be estimated, so there is high variability in estimates if T is small. If T is too small, then even if a change point exists, the empirical covariance matrix may be so variable that the change point is undetectable. This problem is exacerbated when trying to detect change points near the boundary (close to 1 or T). When a change point k is very close to 1, then the empirical covariance matrix $S(1, k)$ is constructed using a very small amount of data and its estimate is unstable with high variance. Similarly, when k is very close to T , $S(k + 1, T)$ suffers from the same problem. This makes change point detection hard: if the empirical covariance matrix is highly variable, the noise from the estimation of the covariance matrices can make any possible differences between Σ_1 and Σ_2 statistically difficult to detect.

In an attempt to quantify just how difficult of a problem change point detection is near the boundary, we find the analytic form of $E[d(k)]$ under H_0 in the case of normally distributed Y_{ij} .

Theorem 1. *Let $Y_j = [Y_{1j}, \dots, Y_{nj}]^T \sim MVN(0, \Sigma)$ for $j \in \{1, \dots, T\}$ all *i.i.d.*, then for $d(k) = \|S(1, k) - S(k + 1, T)\|_F^2$ for any $k \in \{2, \dots, T - 1\}$ we have*

$$E[d(k)] = \left(\frac{1}{k} + \frac{1}{T - k}\right)(\text{tr}(\Sigma^2) + \text{tr}(\Sigma)^2) \tag{1}$$

Proof. We have that $d(k) = \text{tr}\{[S(1, k) - S(k + 1, T)]^2\} = \text{tr}\{[Y^T Y(D(1, k) - D(k + 1, T))]^2\} = \sum_{i=1}^n \sum_{j=1}^n (Y_i^T Y_j)^2 C_{jj} C_{ii}$ where $C = D(1, k) - D(k + 1, T)$. From the variance of a Gaussian quadratic form we have that $E[(Y_i^T Y_j)^2] = 2 \text{tr}(\Sigma^2) + \text{tr}(\Sigma)^2$. Similarly for the covariance case when $i \neq j$, we have that $E[(Y_i^T Y_j)^2] = E[E[(Y_i^T Y_j)^2 | Y_j]] = E[Y_j^T \Sigma Y_j] = \text{tr}(\Sigma^2)$. These combine to give us the result. For the more detailed algebra expanded upon, see the Supplementary Information.

The implication of the Theorem (equation (1)) is that the expected difference asymptotes to infinity as k approaches 0 or T and is minimized when $k = \lfloor T/2 \rfloor$. A similar result was observed in univariate normal change point detection settings^{24,25}. Although this result assumes multivariate normal data, we expect that the qualitative nature of the result generalizes beyond the normal distribution. The increase in $E[d(k)]$ is confirmed through simulation under H_0 and demonstrated in Fig. 1. The implication is that the noise in the estimation of the covariance matrices on both sides of the change point is minimized when both $S(0, k)$ and $S(k + 1, T)$ have sufficient data for their estimation. When k is close to 0, then even though $S(k + 1, T)$ has low variability, the large increase in variability of $S(0, k)$ leads to an overall noisier outcome. This demonstrates that the strength of the method is only as strong as its weakest estimate. For the purposes of study design and data collection, if we suspect that a change point occurs at a certain location, perhaps for theoretical reasons or based on past studies, we need to ensure that there is sufficient data collected both before and after the suspected change point if we are to have any hope of detecting it.

While it is intuitive that T needs to grow as some function of n in order to maintain any reasonable statistical power to detect change points, it is unclear what that function of n is. We investigate here further, using simulation, at what rate the number of longitudinal observations T needs to grow with system size n in order to maintain the

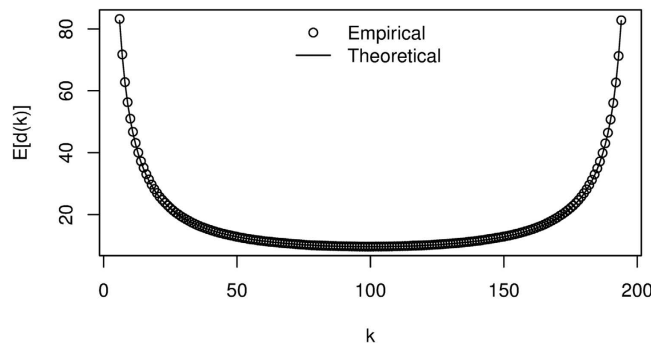


Figure 1. Difficulty of change point detection near the boundaries of data. With $T=200$, $n=20$, and $\Sigma = I_n$ (the identity matrix of order n), for each potential change point $1 < k < T$, we estimate $E[d(k)]$ by averaging $d(k)$ over 10000 simulations under H_0 and show the location of expected values with markers. These empirical estimates are contrasted with the theoretical expectation, shown as a solid line, given by Theorem 1, equation (1).

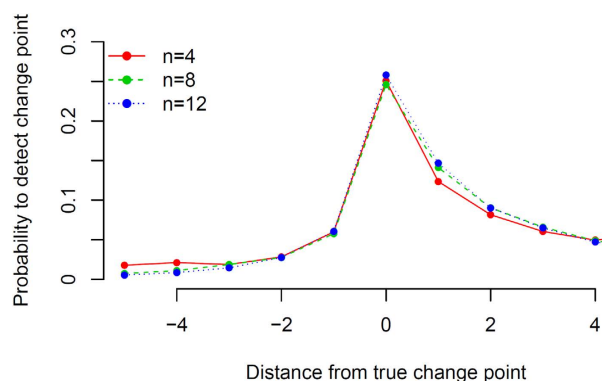


Figure 2. Change point detection statistical power as a function of n and T . The y -axis represents the probability that a change point is detected at a particular time point. The x -axis is the distance of a time point in either direction from the true change point. These probabilities are estimated based on 10000 iterations for each n .

same statistical power. We consider the case where a single change point occurs at the midpoint $\lfloor T/2 \rfloor$, and the node characteristics are $Y_j \sim \text{MVN}(0, \Sigma_1)$ for $j \leq T/2$ and $Y_j \sim \text{MVN}(0, \Sigma_2)$ for $j > T/2$ where:

$$\Sigma_1 = I_n \quad \Sigma_2 = \begin{pmatrix} \underbrace{(1-\rho)I + \rho 11^T}_{4 \times 4} & \underbrace{0}_{4 \times (n-4)} \\ \underbrace{0}_{(n-4) \times 4} & \underbrace{I}_{(n-4) \times (n-4)} \end{pmatrix} \quad (2)$$

where $\rho = 0.9$. In other words, Σ_2 is a block or partitioned matrix with exchangeable correlation within the blocks on the diagonal, and 0s in the off-diagonal blocks. We simulate instances of Y in this fashion 10000 times for each of $n = 4, 8, 12$.

In Fig. 2 we compare the performance of our method for change point detection, measured by the proportion of the 10000 iterations resulting in a statistically significant change point, for the three different values of n . The asymmetry in Fig. 2 around the true change point is caused by having Σ_1 first followed by Σ_2 . If the order of Σ_1 and Σ_2 is reversed, then the asymmetry will be reversed as well. We find that the probability of detecting the correct change point is the same for all n if we increase T by a quadratic rate in n as $T(n) = n(n-1) + C$ for the constant C , a functional form we discovered by numerical exploration. In our simulations we considered the $\alpha = 0.05$ significance level and $C = 30$. The intuition behind a quadratic rate is that as n increases, the number of entries in the empirical covariance matrix increases quadratically and therefore the noise in the Frobenius norm increases quadratically. Increasing T quadratically with n appears to balance out the added noise for increasing the dimensions of the correlation matrices, and stabilizes the statistical power to detect the change point. We would therefore recommend that if one wants to increase n , then there needs to be an associated increase in the number of observations that is quadratic in n in order to retain the ability to detect a change point with the same power.

Comparison of different matrix norms. Up until this point our proposed method has dealt with taking the Frobenius norm of the difference of empirical covariance matrices. The choice of the Frobenius norm was for

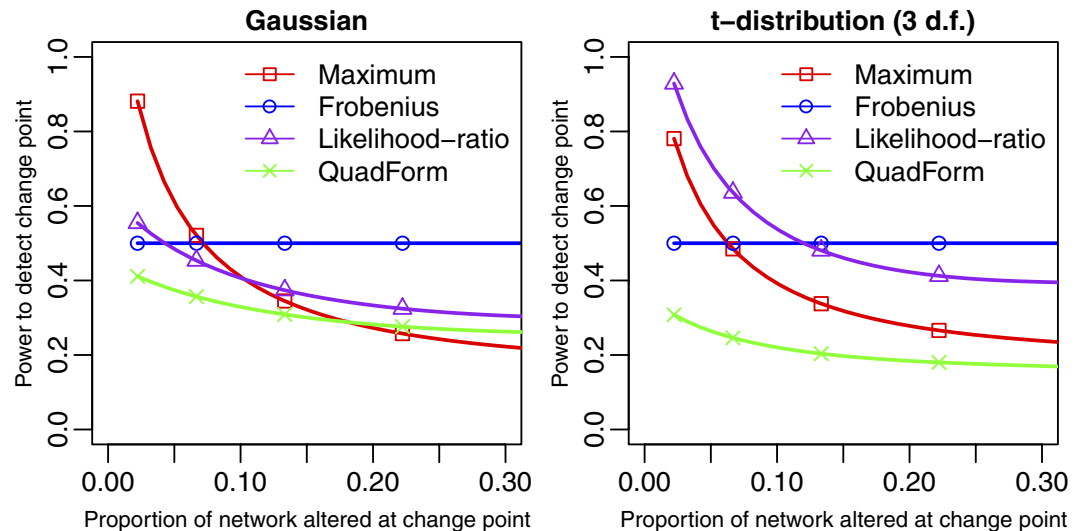


Figure 3. Power comparison for different change point test statistics. The QuadForm method is the test from Aue *et al.*¹⁶ For each point on the x -axis, a value of ρ in the definition of Σ_2 is selected such that the Frobenius norm has 50% power to detect a change point. As the proportion of the network altered at the change point increases, we adjust the value of ρ correspondingly. The left panel corresponds to characteristics with a multivariate normal distribution, while the right panel corresponds to characteristics with a multivariate t -distribution with 3 d.f.

algebraic simplicity of Theorem 1 (equation (1)). Though it is more simple than many other matrix norms for such calculations, there is no reason to believe that the Frobenius norm is uniformly the best choice of matrix norm if the objective is to maximize the statistical power of change point detection. There may be some types of change points that the Frobenius norm is good at detecting, but there may be other types of change points for which a different matrix norm or distance metric would be more suitable. We investigate this question more closely here.

Because the Frobenius norm sums the squared entries of a matrix, it is intuitive to expect that the Frobenius norm would be ideal for detecting change points in systems that demonstrate large-scale, network-wide changes in the correlation pattern. On the other hand, the Frobenius norm likely would not be very powerful in detecting small-scale local changes in correlation network structure. The rationale for this argument is that by summing over all the changes in the network structure, if there are very few changes relative to the entire network, then the Frobenius norm would be dominated by noise from the largely unchanged matrix elements.

We consider a different matrix norm, the Maximum norm, that is appealing for the case of small-scale, local changes. The Maximum norm of a matrix is simply the largest element of the matrix in absolute value. Intuitively, this norm would be ideal if there was just a single, but very large, change in the covariance matrix. If only one element of the covariance matrix changes, but the change is quite large, the Maximum norm would still be able to detect this change. Here the Frobenius norm would likely fail due to the sum of the all the changes being dominated by noise. The likelihood ratio test is more similar to the Frobenius norm than the Maximum norm in that it utilizes all entries in the covariance matrix rather than using only one element. More specifically, the likelihood ratio test statistic involves the product of the determinant of the covariance matrices before and after the change point instead of considering their difference, like the Frobenius norm does. However, the variance of the Frobenius norm is very sensitive to the variance of the correlation matrix, whereas the determinant, and therefore the likelihood ratio test statistic, is very robust to increased noise.

We also consider the nonparametric change point detection of Aue *et al.*¹⁶, which tests for the difference between the empirical covariance matrix before the change point compared to the empirical covariance matrix of the full data. This matrix is written as a vector containing all $n(n-1)/2$ off-diagonal elements, and is normalized by the inverse of its covariance through a quadratic form.

We compare the Frobenius norm, the Maximum norm, the likelihood-ratio test statistic, and the quadratic form of Aue *et al.*¹⁶ through simulation with varying proportion of the network altered at the change point. To do this, we generated Y_j from a multivariate normal distribution with $T = 400$ and a single change point occurring at $t = 200$. Prior to the change point $Y_j \sim MVN(0, \Sigma_1)$ for $j \leq t$ and after the change point $Y_j \sim MVN(0, \Sigma_2)$ for $j > t$, where we modify the dimension of the upper-left block of Σ_2 to change the proportion of the network that is altered at the change point. This simulation is also repeated for a multivariate t -distribution with 3 d.f. In each case ρ is selected such that the change point is detected with 50% power using the Frobenius norm. The results are displayed in Fig. 3, which confirms our intuition. As evidenced by its steep negative slope, the Maximum norm benefits most from having only a small proportion of the network altered. When a large proportion of the network is altered at the change point, then the Frobenius norm benefits the most. While the likelihood-ratio metric is more similar to the Frobenius norm than it is to the Maximum norm, it is still less sensitive to wide-spread subtle network changes than the Frobenius norm. As demonstrated by its gain in relative power, the likelihood-ratio test shows its resilience to increased noise in the case when characteristics are t -distributed. The performance of the Maximum

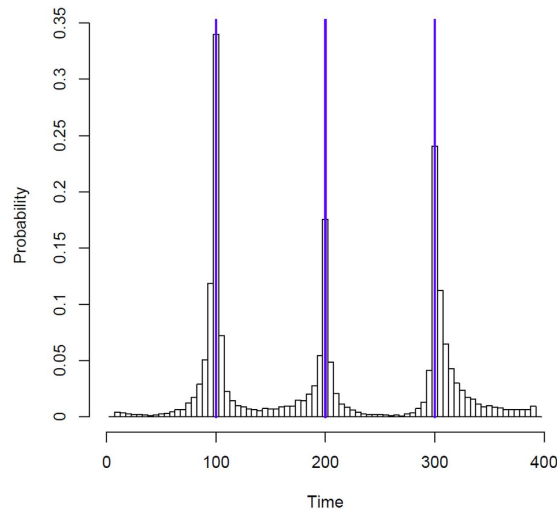


Figure 4. Detection of multiple change points. The y axis represents the probability a change point is detected in a bin, each containing five adjacent time points, based on 10000 simulations total. The true locations of the change points are marked with vertical blue lines.

norm suffers most under this increased noise. The quadratic form of Aue *et al.*¹⁶ has very poor performance in all settings due to the necessity of estimating a large $n(n-1)/2$ by $n(n-1)/2$ matrix, which requires a much larger T relative to n than do the other methods.

These considerations naturally lead to the following question: which norm, statistic, or metric should be used? The answer clearly depends on the anticipated nature of the change point, and is therefore difficult because often the nature of the change point is unknown. In fact, change point detection is used even when one is not sure a change point exists. If there is some *a priori* knowledge of a type of change point perhaps specific to the problem at hand, then that information could be used to select an appropriate norm. For example, suppose we investigate a network constructed from stock return correlations and the time period under investigation happens to encompass a sudden economic recession. The moment the recession strikes, it is likely that there will be large-scale changes in the underlying network, and therefore the Frobenius norm might be a good choice.

One important issue that deserves emphasis is *when* the choice of the norm to use should be made. It is very important that the choice of norm is made prior to looking at the data. If the analysis is performed multiple times repeatedly with different choices for the matrix norm, and the norm with the “best” results is selected, this would be deeply flawed and would invalidate the interpretation of the p -value. See Gelman and Loken²⁶ for an informative discussion of the problem of inflated false positive rates that result when the choice of the specific statistical procedure to use, in this case the norm, is not made prior to all data analysis.

Detecting multiple change points. It may be the case that more than one change point occurs in the data. In this case, our method can still be applied to search for additional change points by splitting the data into two segments at the first significant change point and then repeating the procedure on each segment separately. This process is repeated recursively on segments split around significant change points until no additional statistically significant change points remain. Each test is performed at the $\alpha = 0.05$ level (or at another user-specified level). Though multiple comparisons may seem like a potential problem here, in fact there is no problem because further tests are only performed conditional on the previous change points being elected as statistically significant. This prevents the false positive rate from being inflated. A sliding window approach can also be used to detect multiple change points as is done in Hawkins *et al.*²⁴ and Peel and Clauset¹⁹. Though we assume a constant structure between change points here, for univariate settings this assumption can be dropped and change points can be detected for arbitrary continuous functions between change points²⁷.

We investigate the performance of our method for detecting multiple change points through simulation. Consider the case where $T = 400$ time points are observed for $n = 10$ nodes in a network. Data follows a multivariate normal distribution with mean 0 and covariance Σ_1 for $1 \leq t \leq 100$ and for $201 \leq t \leq 300$, but has covariance Σ_2 for $101 \leq t \leq 200$ and for $301 \leq t \leq 400$. We define Σ_1 and Σ_2 as in (2) with $\rho = 0.9$, except that the upper left block of Σ_2 is 5×5 in this case. The probability that a change point is detected at each particular location is estimated from 10000 iterations and is shown in Fig. 4. Statistical significance is assessed at the usual $\alpha = 0.05$ significance level. We use the Frobenius norm here because changing between Σ_1 and Σ_2 constitutes large-scale change in the network.

As expected, the closer a location is to a change point, the more likely that location is found to be a statistically significant change point. However, there is an asymmetry in the ability to detect the different change points. The change point at $t = 200$ is more difficult to detect than the change points at $t = 100$ and $t = 300$. This is because if we consider all 400 data points and split them around $t = 200$, the two resulting covariance matrices on either side of $t = 200$ are expected to be the same because for both sets of 200 observations, half are from Σ_1 and half are from Σ_2 . This means we have almost no statistical power to detect a change point in the first iteration at $t = 200$. Instead, the change points at $t = 100$ and $t = 300$ are picked up first. The reason the change point at $t = 100$ is easier to detect than the one at $t = 300$, despite each being equally far from its respective boundary, is because

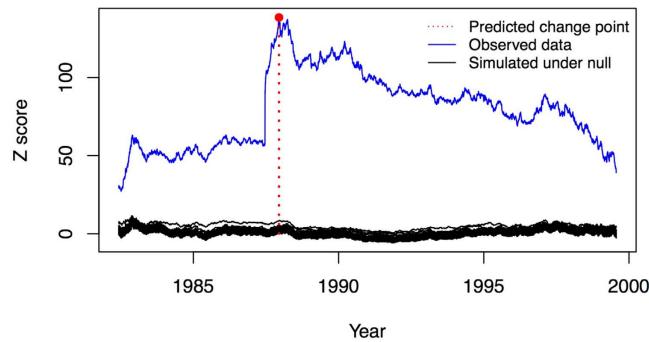


Figure 5. Change point detection in stock returns. With $n = 116$ stocks tracked over $T = 4786$ days (~ 19 years), the blue line is the empirical z-score $z(k)$ while the clustered black lines are the z-scores simulated under H_0 using bootstrap. A significant change point is detected near the end of the year 1987 corresponding to the well documented crash at the end of that year.

the data is ordered with the first 100 observations generated from Σ_1 and the last 100 from Σ_2 . If this is reversed, then $t = 300$ becomes the change point most likely to be detected. After the first change point is detected, power is reduced for the remaining change points due to the reduction in sample size that occurs due to dividing the data into smaller segments.

Correlation networks of stock returns. Our first data analysis example deals with networks constructed from correlations of stock returns. Networks constructed from correlations of stock returns have been used in the past to investigate the correlation structure of markets as well as to detect changes in their structure^{6,28,29}. Here we use a data set first analyzed in Onnela *et al.*² and apply our change point detection method to it.

A total of $n = 114$ S&P 500 stocks were followed from the beginning of 1982 to the end of 2000, keeping track of the stock price at closing for $T = 4786$ trading days over that time period. This data is publicly available and had been gathered for analysis previously where correlation networks were constructed based on the correlation between log returns in moving time windows²⁹. If the price of the i th stock on the j th day is P_{ij} , then the corresponding log return is $R_{ij} = \log(P_{ij}) - \log(P_{i,j-1})$. The log returns did not demonstrate statistically significant autocorrelation (Durbin-Watson p-value of 0.11) so the independent bootstrap was used.

Given that the stock market evolves constantly, and given the long time interval in the observed data, it may not be safe to assume that the correlation between the log returns of any two stocks stays fixed over time. If a correlation network were constructed by assuming edges between every two stocks with correlation greater than some threshold, and if all of the 19 years of data were used at once, the resulting network would likely be an inaccurate representation of the market if in fact the true underlying network changes with time³⁰. A more principled approach would be to first test for change points in the correlation network and then build multiple networks around those change points if necessary.

The stock price data is resampled 500 times assuming the null hypothesis of no change points, and the observed data is compared with these simulations to determine if and where a change point occurs. These simulation results are displayed in Fig. 5. There is strong statistical evidence of a change point at the end of the year 1987 evidenced by a p-value < 0.002 . The sieve bootstrap approach arrives at the same result, confirming the lack of temporal autocorrelation in log returns. We used the Frobenius norm as our goal was to find events that could lead to large-scale shocks to the correlation network. There were several other significant change points, but we focus on the first and most significant change point here. The stock market crash of October 1987, known as “Black Monday”, coincides with the first detected change point²⁸. The stock market crash evidently drastically changed the relationship between many of the stocks leading to a stark change in the correlation network. For this reason it is advisable to consider the network of stocks before and after the stock market crash separately, as well as splitting the data further around potential additional significant change points, rather than lumping all of the data together to construct a single correlation network.

Correlation networks of fMRI activity. Our second data analysis example deals with networks constructed from correlations in fMRI activity in the human brain. These represent functional activity networks as opposed to structural networks. The Center for Cognitive Brain Imaging at Carnegie Mellon University collected fMRI data as part of the star/plus experiment for six individuals as they each completed a set of 40 trials³¹. Each trial took approximately 27 seconds to complete. The subjects were positioned inside an MRI scanner, and at the start of a trial, each subject was shown a picture for four seconds before it was replaced by a blank screen for another four seconds. Then a sentence making a statement about the picture just shown was displayed, such as “The plus sign is above the star,” and the subject had four seconds to press a button “yes” or “no” depending on whether or not the sentence was in agreement with the picture. After this the subject had an interstimulus period of no activity for 15 seconds until the end of the trial. We avoid referring to this as “resting state” due the reserved meaning of that label for extended periods of brain inactivity. Trials were repeated with different variations, such as the picture being presented first before the sentence, or with the sentence contradicting the picture. MRI images were recorded

every 0.5 seconds, for a total of about 54 images over the course of a trial, corresponding to a total of $40 \times 54 = 2160$ images total. Each image was partitioned into 4698 voxels of width 3 mm. The study data are publicly available³².

If we were to analyze a single trial, change point detection would be quite difficult for the data in its raw form for $n = 4698$ voxels which is very large compared with the number of data points $T = 54$. Any empirical covariance matrix for these values of n and T would be too noisy to detect any statistically significant change point. We therefore combine our analysis on the eight trials where the picture is presented first and the sentence agrees with the picture for all six individuals. To accommodate the repeated trials in our correlation estimates, we define $S_k(i, j)$ to be the covariance estimator in equation (5) for the l th individual and k th trial only. The resulting covariance matrix averaged over all the trials and individuals is

$$S^*(i, j) = \sum_{l=1}^6 \sum_{k=1}^8 S_k(i, j) / 48 \quad (3)$$

Change point detection is then performed as before except using the $S^*(i, j)$ instead of the usual $S(i, j)$. Even though we have effectively increased the amount of data 48-fold by combining multiple trials and individuals together, the number of observed data points is still far fewer than the $n = 4698$ voxels. To reduce the number of nodes to a manageable size, we group the voxels into 24 distinct regions of interest (ROIs) in the brain following Hutchinson *et al.*³³, and we average the signals over all voxels within the same ROI. With 24 nodes and $54 \times 48 = 2592$ data points, empirical covariance matrices can be estimated with sufficient accuracy to detect change points in the network of ROIs so long as the change points occur sufficiently far from the beginning or end of the trial. Under the assumption that the network is drastically different when comparing the interstimulus state to the active state, we use the Frobenius norm. For each of the 6 individuals there was a significant presence of autocorrelation (Durbin-Watson p-value < 0.001), so the sieve bootstrap is used for inference. First order autocorrelation ($s = 1$) was used as it minimized mean squared error in cross-validation.

The most significant change point occurred $t = 12$ seconds into the trial, though it was not statistically significant (p-value of 0.18). This indicates that the network of interactions in the first part of the trial when the subject is actively reading, visualizing, responding, and connecting stimuli is most different from the interstimulus portion of the trial, though the difference is not statistically significant. If autocorrelation is ignored and instead independence is assumed, then the same change point appears significant with a p-value less than 0.001. This example demonstrates how ignoring autocorrelation in the data can lead to an inflation of the false positive rate.

Though we fail to reject the null hypothesis and find no statistically significant change point at the $\alpha = 0.05$ level, we examine the position of the most likely change point, if one exists, at $t = 12$. We construct two networks between the ROIs, one from $S^*(1, 24)$ corresponding to the first 12 seconds of the trial (recall that i and j in $S^*(i, j)$ index time points that are 0.5 seconds apart) and one from $S^*(25, 54)$ corresponding to the remaining 15 seconds of inactivity in the trial. The networks are constructed such that an edge is shown between two nodes if and only if their pairwise correlation is greater than 0.5 in absolute value. The two networks are displayed in Fig. 6. The correlation threshold to determine if an edge is present was selected such that the network after the change point had 20 edges, and the same threshold was used for both before and after networks. During the inactive period after $t = 12$ there is an increase in connectivity in the network. An explanation for why the most likely change point occurs at $t = 12$ could be because behavior after the change point corresponds with constant inactivity, whereas before the change point there is a mixture of inactivity, thinking, decisions, and other mental activity likely taking place. These different mental activities could dampen the observed correlations when averaged all together. Given the many different tasks that occur over the course of the trial, there almost certainly exist numerous changes in the correlation structure of brain activity. The lack of statistical significance in this example helps to illustrate the phenomenon discussed previously; even if change points exist, as is almost certainly the case here due to the numerous mental tasks required over the course of a trial, there is no hope of discovering them unless the number of observations is far greater than the size of the network ($T \gg n$).

Discussion

In this paper, new and existing change point detection methods were adapted to correlation networks using a computational framework. Many past treatments of change point detection make strict distributional assumptions on the observed characteristics, but our framework utilizes the bootstrap in order to avoid these restrictions. Traditional methods also assume independence between observations and upon first glance this assumption seems unreasonable. For instance, consider the stock market data. Stock prices are often modeled as a Markov process, which implies a strong autocorrelation between consecutive observations. For this reason the stock prices themselves cannot be used as input for our algorithm, but rather the log returns are used. Similar to the random noise in a Markov chain being independent, the assumption of the returns being independent is more reasonable, as was also demonstrated by the Watson-Durbin test. The fMRI voxel intensities, however, demonstrated significant autocorrelation and required the sieve bootstrap procedure to accommodate this autocorrelation.

We extended our framework to allow for multiple change points. If the first change point is found to be statistically significant, then the data is split into two parts on either side of the change point and the algorithm is repeated for each subset. This process of splitting the data around significant change points continues until there are no more significant change points. The fMRI data analysis found no significant change points but, due to the many changes in stimuli, there are likely multiple points in time where the structure of interaction between regions of interest in the brain changes. This negative result could likely be remedied by collecting higher temporal frequency imaging data. With each split of the data, T is approximately halved while n remains the same and this further lowers the power to detect change points. However, increasing temporal resolution alone is not sufficient because adjacent measurements can become redundant and highly correlated if the resolution is too high. Appropriately increased

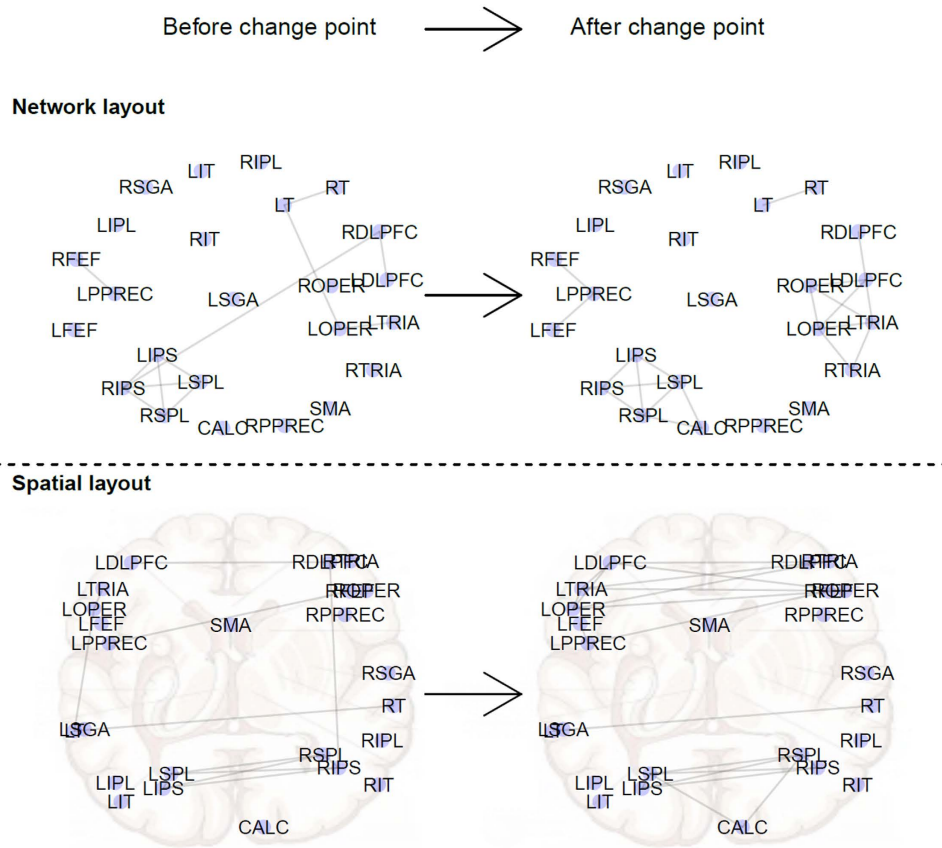


Figure 6. Brain region of interest (ROI) networks before and after the most likely change point. The network transition around the most likely change point are displayed with two different layouts. In the top panel, nodes are positioned to best display the pre-change point network topology. Those same node positions are used in the post-change point network in the top right. The networks on the bottom have nodes positioned according to their Talarach coordinates³⁶ that accurately represent their anatomical location in the brain.

temporal resolution coupled with longer trials or reduced network sizes should make it possible to use the proposed framework to detect multiple change points in correlation networks across different domains.

Methods

Assume that the system under investigation consists of a fixed number of n nodes with characteristics observed at T distinct time points, where the observed characteristics are $Y_{n \times T} = [Y_1, \dots, Y_T]$ where $Y_j = [Y_{1j}, \dots, Y_{nj}]^T$ is the j th n -dimensional column vector of Y and we assume that $Y_j \sim f(\cdot | \Sigma_j)$ is an unknown function with all columns of Y i.i.d. (independent and identically distributed) and where $\text{cov}(Y_j) = \Sigma_j$. We also assume that the rows of Y , corresponding to observations at individual nodes, are centered to have temporal mean 0 and scaled to have unit variance. Note that the centering and scaling, resulting in standardized observations for each node, can always be performed.

We define a set of diagonal matrices $D(i, j)_{T \times T}$ for $1 \leq i < j \leq T$ such that

$$D(i, j)_{kk} = \begin{cases} 1/(j - i + 1) & \text{if } i \leq k \leq j \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

We define the covariance matrix $S(i, j)_{n \times n}$ on the subset of the data ranging from the i th column to the j th column, i.e., from time point i to time point j ($1 \leq i < j \leq T$), to be:

$$S(i, j) = \sum_{k=i}^j Y_k Y_k^T / (j - i + 1) = YD(i, j)Y^T \tag{5}$$

In order to detect a change point, we wish to find the value of k in the range $(1 + \Delta, T - \Delta)$ that maximizes the differences between $S(1, k)$ and $S(k + 1, T)$, where Δ is picked large enough to avoid ill-conditioned covariance matrices ($\Delta > n$). The rationale for this approach is that if there were a change point in the data, the sample covariance matrices on each side of the change point ought to be different in structure. We choose the squared Frobenius norm as our metric for the distance between two matrices. Let our matrix distance metric be:

$$d(k) = \|S(1, k) - S(k + 1, T)\|_F^2 = \text{tr}\{[S(1, k) - S(k + 1, T)]^T [S(1, k) - S(k + 1, T)]\}, \tag{6}$$

where tr is the matrix trace operator. We wish to test the hypotheses:

$$\begin{aligned} H_0: \Sigma_j &= \Sigma \quad \forall j \\ H_A: \text{There exists } k \text{ such that } \Sigma_j &= \begin{cases} \Sigma_1 & \text{for } j \leq k \\ \Sigma_2 & \text{for } j > k \end{cases} \end{aligned} \tag{7}$$

where $S(1, k)$ represents the sample estimate of Σ_1 and $S(k + 1, T)$ represents the sample estimate of Σ_2 .

Existing methodology for change point detection. Consider for a moment the case where the vector Y_j is multivariate normal with expectation μ_1 and variance-covariance matrix Σ_1 before the change point and expectation μ_2 and variance-covariance Σ_2 after it. We denote this $Y_j \sim \text{MVN}(\mu_1, \Sigma_1)$ for $j \leq k$ and $Y_j \sim \text{MVN}(\mu_2, \Sigma_2)$ for $j > k$. A multivariate exponentially weighted moving average (EWMA) model has been developed for the detecting when μ_1 changes to μ_2 ^{9,10}. A likelihood ratio test for detecting change points in the covariance matrix Σ_1 at a known fixed point k was considered by Zamba²³ and Anderson³⁴. The likelihood ratio test statistic for detecting a change point at k is

$$\Lambda_k = \frac{|S(1, k)|^{\frac{k-1}{2}} \cdot |S(k + 1, T)|^{\frac{T-k-1}{2}}}{|S(1, T)|^{\frac{T-1}{2}}}, \tag{8}$$

where $|\cdot|$ is the matrix determinant operator.

This approach makes the assumption that the location of the change point is known to be at k . In reality however the location of the change point is unknown, and the method can be extended to allow an unknown change point location by considering $\max_{1+\Delta \leq k \leq T-\Delta} \{\Lambda_k\}$. When the Y_j are normally distributed then, for a fixed k , $-2 \log(\Lambda_k)$ follows a chi-square distribution for large T and for large $T - k$. Taking the maximum of Λ_k over all possible k results in a less tractable analytic distribution for the test statistic due to the necessity of correcting for multiple testing. For this reason, along with the fact that we do not wish to restrict ourselves to these distributional and asymptotic assumptions, we note that equation (8) can be easily adapted to our framework by defining $d(k) = \Lambda_k$ and proceeding as usual via simulation of the null distribution. Also, even if the characteristics are not normally distributed (like in the case of the t-distribution), the same normal likelihood-ratio test statistic will have the correct size due to this simulation of the null distribution.

We also consider the non-parametric approach of Aue *et al.*¹⁶ which defines

$$v_k = \frac{1}{\sqrt{T}} \left(\sum_{j=1}^k \text{vech}[Y_j Y_j^T] - \frac{k}{T} \sum_{j=1}^T \text{vech}[Y_j Y_j^T] \right) \tag{9}$$

where $\text{vech}(\cdot)$ is the operator that stacks the columns below the diagonal of a symmetric n by n matrix into a vector of length $n(n - 1)/2$. Then the test statistic for the existence of a change point is the sum of the quadratic forms $\frac{1}{n} \sum_{k=1}^T v_k^T \hat{C}_k^{-1} v_k$ where \hat{C}_k is an estimator for $\text{cov}(v_k)$.

Simulation based change point detection. It is of interest to establish a method of change point detection that does not require any distributional assumptions on Y_j , and we develop such a method in this section. Our approach is based on the bootstrap which offers a computational alternative that can well approximate the distribution of Y_j through resampling. Under H_0 , if the Y_j are all independent and come from the same distribution $Y_j \sim f(\cdot|\Sigma)$ for all $1 \leq j \leq T$, then bootstrapping the columns of Y is appropriate. Though $f(\cdot|\Sigma)$ is unknown, we approximate it with the empirical distribution $\hat{f}(\cdot|\Sigma)$ which gives each observed column vector Y_j an equal point mass of $1/T$. This is equivalent to resampling from the columns of Y with replacement.

For many time series applications there may be autocorrelation present between the columns of Y . In this case resampling the columns of Y would break the correlation structure and lead to bias in the approximation of the null distribution. To account for this autocorrelation in the resampling procedure, we use the sieve bootstrap³⁵. In particular, for correlated data the $Y^{(b)}$ are generated for autocorrelation of order s by fitting the model

$$Y_j = \sum_{k=1}^s \hat{\phi}_k Y_{j-k} + \hat{\varepsilon}_j \tag{10}$$

for each $j > s$. The $\hat{\phi}_k$ are estimated from the Yule-Walker equations, and used to solve for the $\hat{\varepsilon}_j$ s through equation (10). The bootstrap residuals, $\hat{\varepsilon}_j^{(b)}$, are resampled from all the $\hat{\varepsilon}_j$ s with replacement. This generates the bootstrapped $Y_j^{(b)}$ according to $Y_j^{(b)} = \sum_{k=1}^s \hat{\phi}_k Y_{j-k} + \hat{\varepsilon}_j^{(b)}$.

Let $Y^{(b)}$ be one of the bootstrap resamples from Y , where each $Y_j^{(b)}$ are generated by bootstrapping from $\hat{f}(\cdot, \Sigma)$ in the case of independence or from the sieve bootstrap for correlated data. This is repeated for $b \in \{1, \dots, B\}$ where B is the total number of bootstrap samples. Δ is a ‘‘buffer’’ that limits the change point detection from searching too close to the boundaries of data. We recommend $\Delta \approx n$. In the case where a change point location k is closer than Δ to either 1 or T , the change point will not be detected but these cases are near impossible to detect regardless of how small we make Δ . For each $k \in \{1 + \Delta, \dots, T - \Delta\}$, $S^{(b)}(1, k)$, $S^{(b)}(k + 1, T)$, and $d^{(b)}(k)$ are calculated where

$S^{(b)}(i, j) = Y^{(b)}D(i, j)Y^{(b)T}$ and $d^{(b)}(k) = \|S^{(b)}(1, k) - S^{(b)}(k + 1, T)\|_F$. Then $\hat{\mu}_0(k) = \frac{1}{B} \sum_{b=1}^B d^{(b)}(k)$ and $\hat{\sigma}_0^2(k) = \frac{1}{B-1} \sum_{b=1}^B (d^{(b)}(k) - \hat{\mu}_0(k))^2$ are calculated for each $k \in \{1 + \Delta, \dots, T - \Delta\}$.

A z-score is then calculated for each potential change point $k \in \{1 + \Delta, \dots, T - \Delta\}$ as

$$z^{(b)}(k) = \frac{d^{(b)}(k) - \hat{\mu}_0(k)}{\sqrt{\hat{\sigma}_0^2(k)}} \quad (11)$$

The change point occurs for the value of k for which the z-score is largest, so we let $Z^{(b)} = \max_k \{z^{(b)}(k)\}$ for each bootstrap sample b . This is also performed on the observed data, with $z(k) = [d(k) - \hat{\mu}_0(k)] / \sqrt{\hat{\sigma}_0^2(k)}$ and $Z = \max_k \{z(k)\}$ being the test statistic. The corresponding p-value obtained from bootstrapping is

$$p\text{-value} = \frac{1}{B} |\{b: Z^{(b)} \geq Z\}|, \quad (12)$$

where $|\cdot|$ is the cardinality of the set. If the p-value is significant, i.e., if sufficiently few bootstrap replicates $Z^{(b)}$ exceed Z , then we reject H_0 and declare a change point exists for the value of k with the highest z-score, i.e., at $\arg \max_k z(k)$.

It is also often the case that there exist more than one change point. In this case the data is split into two segments, one before the first change point and the other after it, and the bootstrap procedure is then repeated separately for each of the two segments. If a significant change point is found on a segment, then that segment is split in two again and this process is repeated until no more statistically significant change points are found. In practice, this procedure terminates after a small number of rounds because each iteration on average halves the amount of data which greatly reduces power to detect a change point after each subsequent iteration.

References

1. Onnela, J.-P. *et al.* Taxonomies of networks from community structure. *Physical Review E* **86**, 036104 (2012).
2. Onnela, J.-P., Kaski, K. & Kertész, J. Clustering and information in correlation based financial networks. *The European Physical Journal B-Condensed Matter and Complex Systems* **38**, 353–362 (2004).
3. Mizuno, T., Takayasu, H. & Takayasu, M. Correlation networks among currencies. *Physica A: Statistical Mechanics and its Applications* **364**, 336–342 (2006).
4. Bhan, A., Galas, D. J. & Dewey, T. G. A duplication growth model of gene expression networks. *Bioinformatics* **18**, 1486–1493 (2002).
5. Kose, F., Weckwerth, W., Linke, T. & Fiehn, O. Visualizing plant metabolomic correlation networks using clique–metabolite matrices. *Bioinformatics* **17**, 1198–1208 (2001).
6. Mantegna, R. N. Hierarchical structure in financial markets. *The European Physical Journal B-Condensed Matter and Complex Systems* **11**, 193–197 (1999).
7. Keightley, M. L. *et al.* An fmri study investigating cognitive modulation of brain regions associated with emotional processing of visual stimuli. *Neuropsychologia* **41**, 585–596 (2003).
8. Hawkins, D. M. & Zamba, K. A change-point model for a shift in variance. *Journal of Quality Technology* **37**, 21–31 (2005).
9. Zamba, K. & Hawkins, D. M. A multivariate change-point model for statistical process control. *Technometrics* **48**, 539–549 (2006).
10. Lowry, C. A., Woodall, W. H., Champ, C. W. & Rigdon, S. E. A multivariate exponentially weighted moving average control chart. *Technometrics* **34**, 46–53 (1992).
11. Qiu, P. *Introduction to statistical process control* (CRC Press, 2013).
12. Qiu, P. & Zhang, J. On phase ii spc in cases when normality is invalid. *Quality and Reliability Engineering International* **31**, 27–35 (2015).
13. Qiu, P. & Li, Z. On nonparametric statistical process control of univariate processes. *Technometrics* **53**, 390–405 (2012).
14. Qiu, P. Distribution-free multivariate process control based on log-linear modeling. *IIE Transactions* **40**, 664–677 (2008).
15. Qiu, P. & Hawkins, D. A rank-based multivariate cusum procedure. *Technometrics* **43**, 120–132 (2012).
16. Aue, A., Hörmann, S., Horváth, L., Reimherr, M. *et al.* Break detection in the covariance structure of multivariate time series models. *The Annals of Statistics* **37**, 4046–4087 (2009).
17. Lindquist, M. A., Waugh, C. & Wager, T. D. Modeling state-related fmri activity using change-point theory. *Neuroimage* **35**, 1125–1141 (2007).
18. Lindquist, M. A. *et al.* The statistical analysis of fmri data. *Statistical Science* **23**, 439–464 (2008).
19. Peel, L. & Clauset, A. Detecting change points in the large-scale structure of evolving networks. *arXiv preprint arXiv:1403.0989* (2014).
20. Akoglu, L. & Faloutsos, C. Event detection in time series of mobile communication graphs. Paper presented at *Army Science Conference: Orlando FL. Proceedings of the Army Science Conference*, 77–79 (2010).
21. McCulloh, I. & Carley, K. M. Detecting change in longitudinal social networks. *Journal of Social Structure* **12** (2011).
22. Tang, M., Park, Y., Lee, N. H. & Priebe, C. E. Attribute fusion in a latent process model for time series of graphs. *IEEE Transactions on Signal Processing* **61**, 1721–1732 (2013).
23. Zamba, K. A multivariate change point model for change in mean vector and/or covariance structure *Journal of Quality Technology* **41**, 285–303 (2009).
24. Hawkins, D. M., Qiu, P. & Kang, C. W. The changepoint model for statistical process control. *Journal of quality technology* **35**, 355–366 (2003).
25. Hawkins, D. M. Testing a sequence of observations for a shift in location. *Journal of the American Statistical Association* **72**, 180–186 (1977).
26. Gelman, A. & Loken, E. The garden of forking paths: Why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time. *Technical report*, Department of Statistics, Columbia University. (2013) (Date of access: 14/08/2014), from http://www.stat.columbia.edu/gelman/research/unpublished/p_hacking.pdf.
27. Xia, Z. & Qiu, P. Jump information criterion for statistical inference in estimating discontinuous curves. *Biometrika* **102**, 397–408 (2015).
28. Onnela, J.-P., Chakraborti, A., Kaski, K. & Kertész, J. Dynamic asset trees and black monday. *Physica A: Statistical Mechanics and its Applications* **324**, 247–252 (2003).
29. Onnela, J.-P., Chakraborti, A., Kaski, K., Kertész, J. & Kanto, A. Dynamics of market correlations: Taxonomy and portfolio analysis. *Physical Review E* **68**, 056110 (2003).
30. Onnela, J.-P. *Complex networks in the study of financial and social systems*. Phd Thesis (Helsinki University of Technology, 2006). (Date of access: 23/05/2014).

31. Mitchell, T. M. *et al.* Learning to decode cognitive states from brain images. *Machine Learning* **57**, 145–175 (2004).
32. Just, M. Starplus fmri data (2001). (Date of access: 03/06/2014) <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-81/www/>.
33. Hutchinson, R. A., Niculescu, R. S., Keller, T. A., Rustandi, I. & Mitchell, T. M. Modeling fmri data generated by overlapping cognitive processes with unknown onsets using hidden process models. *NeuroImage* **46**, 87–104 (2009).
34. Anderson, T. *An Introduction to Multivariate Statistical Analysis* (New York: Wiley, 1984).
35. Bühlmann, P. *et al.* Sieve bootstrap for time series. *Bernoulli* **3**, 123–148 (1997).
36. Lancaster, J. L. *et al.* Automated talairach atlas labels for functional brain mapping. *Human brain mapping* **10**, 120–131 (2000).

Acknowledgements

I.B. is supported by PHS Grant Number 2T32ES007142-32. I.B. and J.P.O. are supported by Harvard T.H. Chan School of Public Health Career Incubator Award to J.P.O. J.P.O. is further supported by NIH 4R37AI051164-14 (DeGruttola). The authors would like to thank Dr. Stephen Maher for his review of the fMRI analysis.

Author Contributions

I.B. and J.-P.O. conceived the project; I.B. performed the simulations and analysis; both authors interpreted the results and produced the final manuscript.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Barnett, I. and Onnela, J.-P. Change Point Detection in Correlation Networks. *Sci. Rep.* **6**, 18893; doi: 10.1038/srep18893 (2016).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>