SINGLE Package Vignette

Ricardo Pio Monti, Christoforos Anagnostopoulos and Giovanni Montana

January 11, 2014

Given multiple longitudinal time series we are often interested in quantifying the relationship between the time series over time. For example, given fMRI data corresponding to various regions of interest in the brain we may be interested in measuring their statistical dependencies over time. These relationships can subsequently be summarised as graphs or networks where each node corresponds to a time series (e.g., a BOLD time series for a region of interest) and edges between nodes (or their absence) provide information regarding the nature of the pairwise relationship. Here we present a short tutorial and introduction to using the R package SINGLE to estimate such dynamic graphs over time from noisy time series data. The SINGLE package provides an implementation of the Smooth Incremental Graphical Lasso Estimation algorithm, full details of which can be found in [Monti et al., 2013].

1   Introduction

There is an increasing interest in summarising the relationships between time series using undirected graphs. Here each node represents a time series (e.g., this can be the BOLD time series for a brain region) and the edge structure serves to summarise the statistical relationship between nodes. Edge structure is commonly estimated using partial correlations. In this case, the absence of an edge between two nodes implies that the two nodes are conditionally independent given all other nodes.

However it is often the case that only a global graph is estimated using the entire time series. Whilst this may be appropriate in some scenarios, it is often the case that we expect the statistical dependencies between nodes to change over time. A clear example of this can be seen by considering fMRI time series of a subject performing alternating tasks: we naturally expect the relationship between brain regions to change depending on task.

In this vignette, we introduce the `SINGLE` package which can be used to estimate dynamic graphs from noisy time series data. The remainder of this vignette is organised as follows: in Section 2 we give a brief motivating example to show the capabilities of the SINGLE algorithm. In Section 3 we give a brief background description of the SINGLE algorithm. Finally, in Section 4 we give a more detailed description of each of the functions in the `SINGLE` package and their usage.

2   Motivating example

Here we present a brief motivational example to show the capabilities of the SINGLE algorithm and the functionality of the `SINGLE` package. The main function is the `SINGLE` function which provides an implementation of the Smooth Incremental Graphical Lasso Estimation (SINGLE) algorithm. Here we give a brief illustration of how the SINGLE algorithm can be used to estimate non-stationary networks.

We begin by simulating non-stationary data in order to test the performance of the SINGLE algorithm. We simulate the data using the `generate_random_data` function:

```
> library('SINGLE')
> set.seed(1)
> sim = generate_random_data(ROI=5, length_=50, seg=3, sparsity=.1, str=-.6)
```

Here `sim` is a list of length two. The first entry contains a matrix of simulated non-stationary data and the second contains an array which summaries the true correlation structure over time. Thus `sim$data` is a matrix with $p = 5$ columns and $n = 150$ rows. For example, this could correspond to fMRI data from 5 regions of interest over a time period of length 150. As mentioned previously each column represents a node in our graph and each row is a chronologically ordered observation. Thus `sim$data[i,j]` is the ith observation of the jth node.

The true partial correlation structure over time is stored in the array `sim$true_cov`. We also note that this dataset contains two changepoints, one at the 50th observation and the other at the 100th. We note that `sim$true_cov[,,i]` contains the correlation structure at the ith observation. Thus we can view the correlation structure for at the first observation as follows:

```
> sim$true_cov[,,1]
```

```
      [,1] [,2] [,3] [,4] [,5]
[1,]   1.0    0    0    0 -0.6
[2,]   0.0    1    0    0  0.0
[3,]   0.0    0    1    0  0.0
[4,]   0.0    0    0    1  0.0
[5,]  -0.6    0    0    0  1.0
```

We can now run the SINGLE algorithm to estimate the network structure over time. The SINGLE algorithm requires the input of three intuitive parameters. Parameters λ_1 and λ_2 control the level of sparsity and temporal homogeneity respectively and h represents the width of the kernel used to estimate sample covariance matrices. Optimal values for each parameter can be estimated efficiently and this functionality is incorporated into the SINGLE function. As a result, in order to implement the SINGLE algorithm we simply run:

```
> S = SINGLE(data=sim$data, verbose=TRUE)

Covariance matrices not provided... will be estimated
Kernel width, h, not provided... will be estimated
l1 parameter not provided... will be estimated
l2 parameter not provided... will be estimated
Estimating optimal l1 and l2 parameters...
```

We note from the output that each of the parameters λ_1, λ_2 and h are estimated directly from the data. Moreover, the SINGLE function is flexible allowing the user to provide specific parameters if necessary. The user can either specify a given value for λ_1, λ_2 and h as shown below:

```
> S = SINGLE(data=sim$data, l1=.75, l2=.5, h=40)
```

Alternatively, a range of possible choices can be provided for any of the three parameters as follows:

```
> S = SINGLE(data=sim$data, l1=seq(.25,1, .25), l2=seq(.25,1, .25), h=c(30,40,50))
```

Additionally, kernel width parameter, h, can be estimated using the function `choose_h`. This function estimates the optimal choice of h by calculating the look-ahead log-likelihood (see section 4.5 for full details). Below we provide an example

```
> data = sim$data
> h_G = choose_h(data=data, sample_size=30, kernel="gaussian",
+               h=seq(10,100,10))
> h_W = choose_h(data=data, sample_size=30, kernel="window",
+               h=seq(10,100,10))
```

Here the choice of `kernel` allows the user to alternate between a Gaussian kernel and a sliding window. The user can also estimate sample covariance matrices $S_1, \ldots, S_T$ using the `get_kern_cov` function. The estimated sample covariance matrices can then be provided directly to the SINGLE function, therefore reducing the computational cost if several iterations need to be run on the same dataset. This is demonstrated below:

```
> C_gaus = get_kern_cov(data=data, h=h_G, kernel="gaussian")
> C_slid = get_kern_cov(data=data, h=h_W, kernel="window")
> S_gaus = SINGLE(data=data, C=C_gaus, l1=.75, l2=0.5)
> S_slid = SINGLE(data=data, C=C_slid, l1=.75, l2=0.5)
```

This allows the user to run the SINGLE algorithm on sample covariance matrices that have been estimated using any method of their choice. The `plotSINGLE` function can then be used to visualise partial correlations over time.

```
> plotSINGLE(object=S_gaus, index=c(1,2,3,4,5), x.axis = seq(1,150),
+            n.row=2,
+            col.names=seq(1,5), fix.axis=TRUE)
```
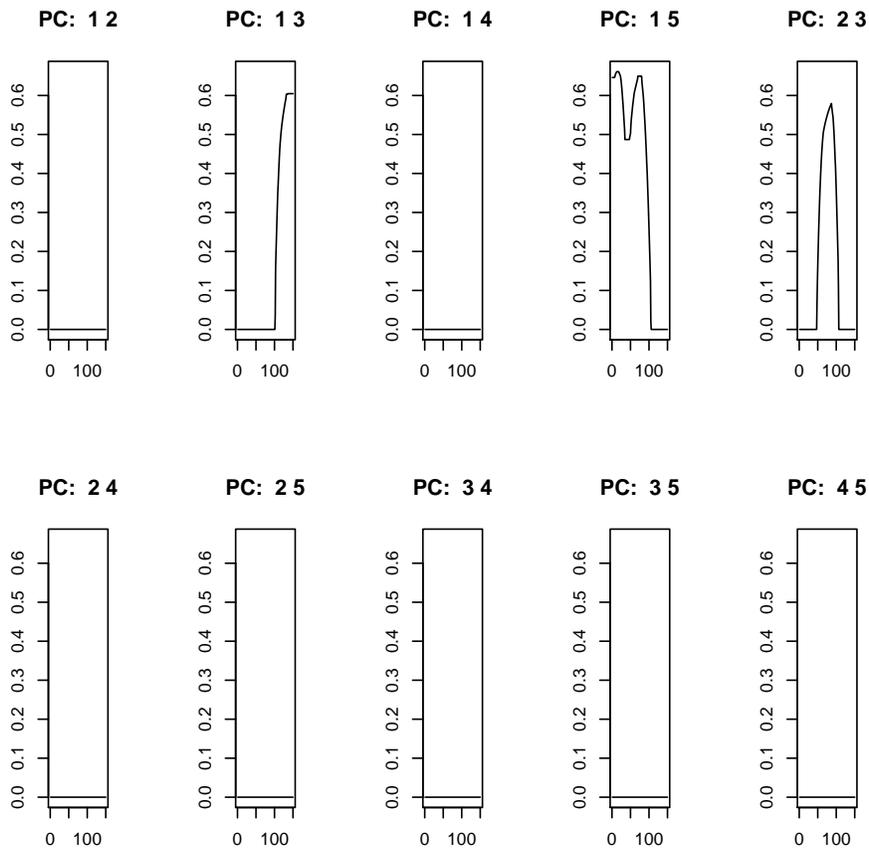

Figure 1: Estimated partial correlations for simulated example

We can see that the estimated partial correlations accurately reflect the true network structure (stored in `sim$true_cov`). Given this is a simulated example we can assess the performance of the SINGLE algorithm using precision, recall and F scores over time using the `precision_recall` function:

```
> result = precision_recall(true_cov=sim$true_cov, estimated_cov=S_gaus$P_)
> plot(result$F1, type='l', ylim=c(0,1), ylab='',
+       main='F Score', xlab='Time')
```
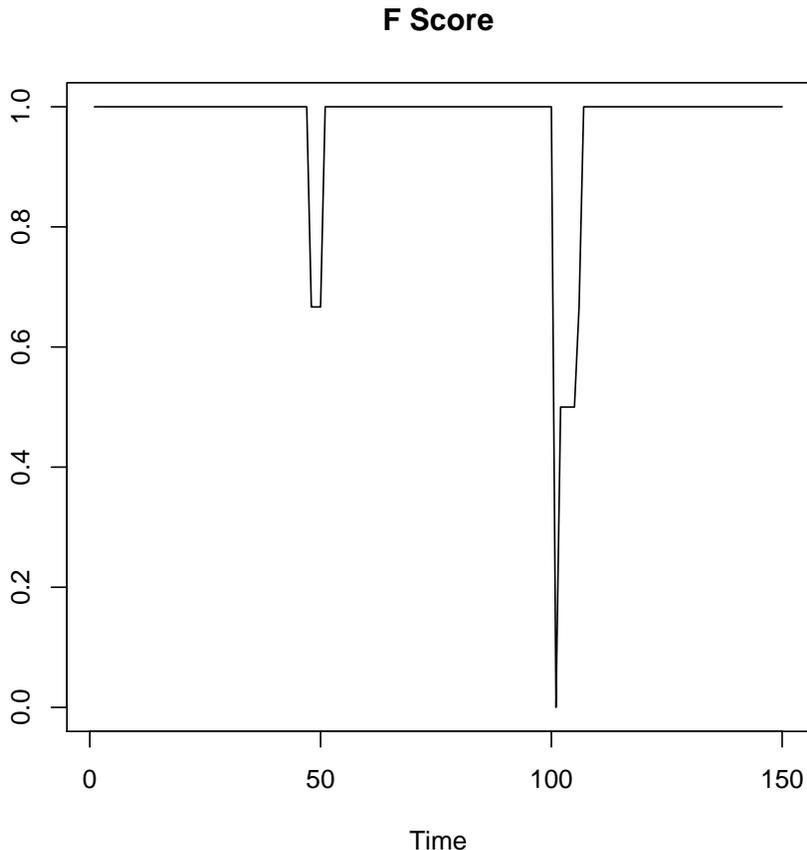
F Score

Figure 2: Plot of the F score over time for the SINGLE algorithm on simulated data

3 Background

We assume to have obtained time series denoted by $X_1, \ldots, X_T$, where each vector $X_i \in \mathbb{R}^{1 \times p}$ contains the measurements for each of the p variables of interest at the ith time point. We are interested in inferring a sequence of graphs $\{G_1, G_2, \ldots, G_T\}$ where each $G_i = (V, E_i)$ corresponds to the functional connectivity between nodes, V, at time i. The edge structure, E, is determined using partial correlations.

The SINGLE algorithm has the following two desirable properties: (a) each graph G_i is sparse thus allowing for accurate and interpretable estimates of the underlying graphs, and (b) the structure of the estimated graphs is temporally homogeneous, implying that there are sparse innovations in network structure over time.

At any given time point, we assume that the random vector X_i follows a multivariate Gaussian distribution, however both the mean and the covariance of this distribution are assumed to be dependent on the time index. We write S_i to denote the estimated covariance

5

matrix at time i.

The set of partial correlations is summarised in the precision (inverse covariance) matrix. Thus our objective is equivalent to obtaining a sequence of time-dependent estimates of precision matrices, $\hat{\Theta}_1, \ldots, \hat{\Theta}_T$. The Smooth Incremental Graphical Lasso (SINGLE) obtains these estimates by solving a constrained optimisation problem which balances goodness-of-fit, sparsity and temporal smoothness. This is achieved by formulating the following objective function:

$$f(\{\hat{\Theta}\}) = \sum_{i=1}^{T} \left[-\log \det \hat{\Theta}_i + \text{trace}\,(S_i \hat{\Theta}_i) \right] + \lambda_1 \sum_{i=1}^{T} ||\hat{\Theta}_i||_1 + \lambda_2 \sum_{i=2}^{T} ||\hat{\Theta}_i - \hat{\Theta}_{i-1}||_1, \quad (1)$$

where $\{\hat{\Theta}\} = \{\hat{\Theta}_1, \ldots, \hat{\Theta}_T\}$ contains all the precision matrices indexed by time. By taking a closer look at the objective function we can gain a clear understanding of what the SINGLE algorithm is looking to achieve. The first sum is proportional to the sum of negative log likelihoods of the estimated precision matrices. The first penalty term regularised by λ_1 corresponds to the Graphical Lasso penalty and ensures sparsity in the estimated graphical structure by imposing a penalty on the sum of absolute entries in each $\hat{\Theta}_i$. On the other hand, the second penalty function regularised by λ_2, ensures smoothness by penalising the differences between temporally adjacent networks. This penalty can be seen as an extension of the Fused Lasso penalty from the context of penalised regression (i.e., in the Fused Lasso we penalise regression coefficients $|\beta_i - \beta_{i+1}|$ and here this is extended to the difference over graphs).

The SINGLE estimation procedure consists of two independent steps performed in sequence: initially, recursive estimates of the covariance matrices $S_1, \ldots, S_T$ are obtained using a Gaussian kernel; then an iterative optimisation algorithm is run until convergence in order to produce a sequence of estimated graphs. Full details of the SINGLE algorithm are given in [Monti et al., 2013].

4   The `SINGLE` package

In this section we give a more detailed description of each of the functions contained in the `SINGLE` package.

4.1   `generate_random_data`

This function allows for the generation of data with a random correlation structure that is piecewise continuous. This is achieved using Vector Autoregressive Processes (VAR). The choice of VAR processes here is motivated by their ability to encode autocorrelation within a time series as well as cross-correlations across time series. Network structure can be simulated according to either Erdos-Renyi random graphs or Barabasi-Albert scale-free networks. This choice is specified by `mode` input in the `generate_random_data` function.

As shown in the motivational example, we can easily simulate data using this function:

```
> sim_ER = generate_random_data(ROI=5, length_=50, mode='ER', seg=3, sparsity=.1)
> sim_BA = generate_random_data(ROI=5, length_=50, mode='BA', seg=3, sparsity=.1)
```

We note that the number of piecewise continuous segments of data and their lengths can be specified using `seg` and `length_` respectively. Random data with segments of different lengths can be generated using the `generate_random_data` command several times as follows:

```
> sim1 = generate_random_data(ROI=5, length_=25, mode='BA', seg=1, sparsity=.15)
> sim2 = generate_random_data(ROI=5, length_=75, mode='BA', seg=1, sparsity=.15)
> data = rbind(sim1$data, sim2$data)
```

This function depends on both the `igraph` [Csardi and Nepusz, 2006] and `dse` packages [Gilbert, 2006] for simulating random networks and VAR processes respectively.

4.2 SINGLE

The `SINGLE` function is the main function within the `SINGLE` package and provides an implementation of the SINGLE algorithm. The SINGLE algorithm involves three parameters, all of which can be efficiently estimated. These parameters can either be specified by the user or alternatively estimated from the data directly. In order to provide a greater insight to the user, we detail each of these parameters below:

1. `h`: This is kernel width to be used when estimating sample covariance matrices. If `h` is not provided it is estimated using the look-ahead log-likelihood, provided by the `choose_h` function (full details provided in section 4.5).

2. `l1`: This penalty measures the regularisation of the sparsity penalty in the SINGLE objective function. Increasing the value of λ_1 will increase the sparsity of the estimated dynamic networks.

3. `l2`: This is the value of λ_2 in the SINGLE objective function. Increasing the value of λ_2 will encourage sparse innovations in the estimated network structures by penalising the difference between consecutive networks.

As shown in Section 2 the SINGLE algorithm can be run as follows:

```
> set.seed(1)
> sim = generate_random_data(ROI=5, length_=50, seg=3, sparsity=.1, str=-.6)
> data = sim$data
> S = SINGLE(data=data)
```

Since the SINGLE algorithm relies on the Fused Lasso for one step of its iterative optimisation this function depends on the `flsa` package [Hoefling, 2013].

4.3 precision_recall

The `precision_recall` function allows us to measure the performance of the SINGLE algorithm. Precision is defined as the fraction of reported edges which are true edges while recall is defined as the fraction of true edges which are correctly reported. It follows that the closer precision and recall are to 1 the better the performance. Finally we also define the F score as follows:

$$F = 2 \cdot \frac{P \cdot R}{P + R},$$

where P and R refer to the precision and recall respectively.

This function calculates the precision and recall for an iteration of the SINGLE algorithm at each observation. That is, at the ith observation the precision, recall and F score is calculated by comparing the true correlation structure at the time and the correlation structure estimated by the SINGLE algorithm.

7

The `precision_recall` function is run directly on the array of estimated precision matrices (`S$P_` in the example above) as follows:

```
> result = precision_recall(true_cov=sim$true_cov, estimated_cov=S$P_)
```

4.4 plotSINGLE

The `plotSINGLE` function allows for the visualisation of the estimated correlation structures over time. This is achieved by plotting the pairwise partial correlations between each of the nodes. Since there may be a potentially large number of nodes (resulting in an even larger number of pairwise correlations to plot) the `plotSINGLE` function allows the user to specify a subset of nodes to plot using its `index` command. In the example from before we plot the partial correlations for all nodes. Alternatively, we could plot only the partial correlations between the 1st, 3rd and 5th nodes as follows:

```
> plotSINGLE(object=S, index=c(1,2,3,5), x.axis = seq(1,150), n.row=2,
+           col.names=c(1,2,3,5), fix.axis=TRUE)
```
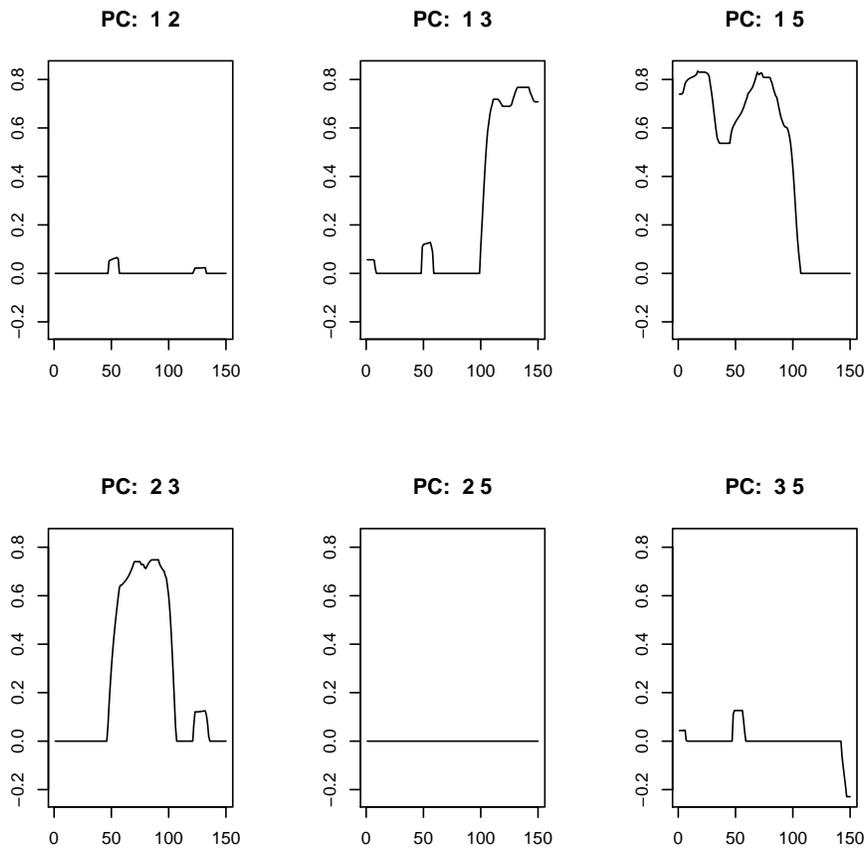
Figure 3: Estimated partial correlations for simulated example discussed previously

4.5 `choose_h`

The `choose_h` function allows the user to estimate the optimal kernel width using look-ahead log-likelihood. We define

$$\mathcal{L}_{-1}(i; h) = -\frac{1}{2} \log \det (S_{i-1}) - \frac{1}{2}(X_i - \mu_{i-1})^T S_{i-1}^{-1}(X_i - \mu_{i-1}) \tag{2}$$

as the estimated look-ahead log-likelihood for the ith observation for some fixed choice of h. We note that both μ_{i-1} and S_{i-1} are estimated with the ith observation removed. It follows that calculating $\mathcal{L}_{-1}(h)$ for all observations is computationally expensive, thus in order to save computational effort we sample a subset of observations $R \subseteq \{2, \ldots, T\}$. The optimal value of h can then be estimated by choosing h to maximise the following score function:

$$CV(h) = \sum_{i \in R} \mathcal{L}_{-1}(i; h), \tag{3}$$

where in order to reduce the variability of estimated look-ahead log-likelihoods we use the same sample R to calculate each $CV(h)$.

The `sample_size` input in the `choose_h` function determines the size of subset R while the input `h` is a vector specifying the values of h at which we should calculate $CV(h)$....

4.6 `get_kern_cov`

The `get_kern_cov` function allows the user to estimate sample covariance matrices for a given data using either a Gaussian kernel or sliding window of any given width. The following formula is used to estimate sample covariance matrices:

$$\mu_i = \frac{\sum_{j=1}^{T} K_h(i, j) \cdot X_j}{\sum_{j=1}^{T} K_h(i, j)}, \tag{4}$$

$$S_i = \frac{\sum_{j=1}^{T} K_h(i, j) \cdot (X_j - \mu_j)^T (X_j - \mu_j)}{\sum_{j=1}^{T} K_h(i, j)}. \tag{5}$$

where $K_h(\cdot, \cdot)$ is a kernel function.

References

G. Csardi and T. Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006. URL `http://igraph.sf.net`.

P. D. Gilbert. *Brief User's Guide: Dynamic Systems Estimation distributed with the dse package*, 2006. URL `http://cran.r-project.org/web/packages/dse/vignettes/dse-guide.pdf`.

H. Hoefling. *flsa: Path algorithm for the general Fused Lasso Signal Approximator*, 2013. URL `http://CRAN.R-project.org/package=flsa`. R package version 1.05.

R. P. Monti, P. Hellyer, D. Sharp, R. Leech, C. Anagnostopoulos, and G. Montana. Estimating Time-varying Brain Connectivity Networks from Functional MRI Time Series. *Preprint*, 2013.