

# LEAP: Constructing gene-coexpression networks for single-cell sequencing data using pseudo-time ordering

Alicia T. Specht and Jun Li

October 18, 2016

## Abstract

**Summary:** To construct gene co-expression networks based on single-cell RNA-Sequencing data, we present an algorithm called LEAP, which utilizes the estimated pseudo-time information of the cells to find stronger associations between pairs of genes.

**Contact:** [aspect2@nd.edu](mailto:aspect2@nd.edu)

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Data format</b>	<b>2</b>
<b>3</b>	<b>Maximum Absolute Correlation (MAC) Counter Function</b>	<b>2</b>
<b>4</b>	<b>Permutation Analysis Function</b>	<b>3</b>
<b>5</b>	<b>Creating Table 1 from our paper</b>	<b>5</b>
<b>6</b>	<b>Generating matrices for use in WGCNA</b>	<b>6</b>
<b>7</b>	<b>Session information</b>	<b>7</b>
<b>8</b>	<b>Citation information</b>	<b>7</b>

## 1 Introduction

Advances in sequencing technology now allow researchers to capture the expression profiles of individual cells. Several algorithms have been developed to attempt to account for these effects by determining a cell's so-called 'pseudo-time', or relative biological state of transition.

By applying these algorithms to single-cell sequencing data, we can sort cells into their pseudotemporal ordering based on gene expression. LEAP (Lag-based Expression Association for Pseudotime-series) then applies a time-series inspired lag-based correlation analysis to reveal linearly dependent genetic associations.

## 2 Data format

LEAP takes a data matrix for which the rows are genes and the columns are experiments, sorted by their pseudo-time. For example, consider this dataset consisting 20 genes from a dataset of high throughput single-cell RNA sequencing counts of *Mus musculus* dendritic cells(Shalek *et al*, 2014):

```
> library("LEAP")
> example_data[,1:5]

  X..0.000000 X..1.056273 X..1.402318 X..1.607550 X..1.779713
1    5.445686    5.102184    0.000000    2.9839979    0.000000
2    4.526036    6.532284    6.0904955    5.6275842    6.1838837
3    6.528766    6.214955    0.000000    3.2003647    6.4477135
4    7.252316    6.627344    0.000000    1.6035663    0.000000
5    0.000000    5.209410    5.5418589    5.5792501    2.8199493
6    0.000000    0.000000    0.000000    5.8262664    3.4519732
7    3.253274    2.959225    3.9272823    6.3940181    6.2686141
8    0.000000    0.000000    0.9958159    1.3969442    5.0141277
9    6.685728    6.414861    6.6039581    6.0960811    5.8442669
10   0.000000    3.201347    5.9733666    6.4752538    5.1608972
11   0.000000    0.000000    2.8945397    5.3855163    0.000000
12   0.000000    0.000000    0.000000    5.9502785    0.000000
13   0.000000    5.413914    5.6621031    0.000000    5.2452502
14   0.000000    3.174056    0.000000    3.2505840    6.6724467
15   6.664693    7.092854    0.000000    0.000000    6.7583925
16   1.035870    7.706152    0.000000    0.9082936    6.6226587
17   5.773570    4.750586    0.000000    0.000000    0.9529196
18   3.884139    4.667165    0.000000    5.6001221    4.5408934
19   6.288354    6.437330    6.4375612    6.2336922    6.5298076
20   0.000000    0.000000    0.000000    0.000000    5.3029461
```

We've shown only the first 5 cells here. The column names are the pseudo-times that were generated for each sample using Monocle (Trapnell *et al.*, 2014). As you can see, the samples have been ordered from lowest to greatest pseudo-time. We've also applied a  $\log(x+1)$  transformation to the count data.

## 3 Maximum Absolute Correlation (MAC) Counter Function

Once your data is in the above format, you can use the `MAC_counter()` function to calculate the Max Absolute Correlation (MAC) matrix for you data. The output of this function is a matrix where **Row gene index** and **Column gene index** correspond to the indeces for the gene pair (i,j), **Correlation** is the maximum absolute correlation (MAC) achieved for the pair, and **Lag** is the lag at which the MAC occurred. Note that the pair (i,j) and (j,i) will both appear in the results, as they will potentially have different MACs. As can be seen below, setting `MAC_cutoff=0.2` restricts the output to only those pairs with an MAC of 0.2 or greater.

```
> MAC_results = MAC_counter(data=example_data, max_lag_prop=1/3, MAC_cutoff=0.2, file_name="example", lag_m
> MAC_results[99:119,]
```

	Correlation	Lag	Row gene index	Column gene index
[1,]	0.2635403	172	2	19
[2,]	0.2526990	0	4	1
[3,]	0.2526990	0	1	4
[4,]	-0.2496095	3	1	7
[5,]	0.2485226	0	15	9
[6,]	0.2485226	0	9	15
[7,]	0.2420226	93	2	5

```

[8,] -0.2322217  3           1           12
[9,]  0.2314823  0           16          13
[10,] 0.2314823  0           13          16
[11,]  0.2224057  0           4           3
[12,]  0.2224057  0           3           4
[13,] -0.2214821  0           16          15
[14,] -0.2214821  0           15          16
[15,]  0.2205179  0           17           4
[16,]  0.2205179  0           4           17
[17,] -0.2193975 114          15           8
[18,]  0.2161187  0           7           4
[19,]  0.2161187  0           4           7
[20,] -0.2137327  90          17           5
[21,] -0.2125380 102          17          18

```

Here, `max_lag_prop` is the largest proportion of your experiments that you want your lag to be. For this example, we have 512 experiments, so the largest lag we will try is 172. We recommend using at most a `max_lag_prop=1/3`. The variable `file_name` is the name you'd to associate with your files. Our example creates the file `MAC_example.csv`.

```
> MAC_example[1:4,1:4]
```

```

      [,1]      [,2]      [,3]      [,4]
[1,]      NA -0.1728618  0.4182790  0.2526990
[2,] 0.1414134      NA -0.1427626  0.1463172
[3,] 0.4182790  0.1583516      NA  0.2224057
[4,] 0.2526990 -0.1593688  0.2224057      NA

```

The variable `lag_matrix` decides whether you would like the associated matrix of lag values to be saved as well. For our example, setting `lag_matrix=T` creates the file `lag_example.csv`.

```
> lag_example[1:4,1:4]
```

```

  V1 V2 V3 V4
1 NA 172  0  0
2 121 NA 141 165
3  0  67 NA  0
4  0  47  0 NA

```

Again, note that the diagonal is set to NA. It is important to note that each of the values in the lag matrix correspond to the size of the lag used on the gene listed in the column. In our example, 172 corresponds to starting gene 1's expression at its first pseudo-time point and staggering the expression of gene 2 by 172 pseudo-time points (hence starting at 173).

## 4 Permutation Analysis Function

To determine a cutoff for significant MAC values, you can use the `MAC_perm()` function.

```
> MAC_perm(data=example_data, MACs_observ=MAC_example, num_perms=10, max_lag_prop=1/3,
+          FDR_cutoffs=101, perm_file_name="example")
```

The variable `num_perms` determines the number of permutations to use. Note we've only used 10 here to simplify our example. For larger datasets, using 100 is most likely appropriate. `FDR_cutoffs` determines the number of cutoffs you'd like to use to split the domain [0,1] for the correlation. `data`, `max_lag_prop` and `perm_file_name` follow the same use as described for `MAC_counter()`.

This returns the dataset below, where `cors` are the correlation cutoffs, `MACs_observed` are the number of observed correlations at that cutoff, `MACs_ave_perm` are the average number observed in the permuted datasets at that cutoff, and `fdr` is the false discovery rate (FDR) observed at that cutoff. We can see that for our example dataset, if we would like to control the FDR around 0.1, then a correlation cutoff of 0.18 would be appropriate. Below are shown the results with nonzero FDR:

```
> perm_example[74:101,]
```

	cors	MACs_observed	MACs_ave_perm	fdr
74	0.27	98	0.0	0.0000000000
75	0.26	99	0.0	0.0000000000
76	0.25	101	0.1	0.0009900990
77	0.24	104	0.1	0.0009615385
78	0.23	106	0.3	0.0028301887
79	0.22	110	0.7	0.0063636364
80	0.21	113	2.4	0.0212389381
81	0.20	120	4.7	0.0391666667
82	0.19	137	8.6	0.0627737226
83	0.18	152	17.7	0.1164473684
84	0.17	168	35.8	0.2130952381
85	0.16	188	59.8	0.3180851064
86	0.15	217	95.4	0.4396313364
87	0.14	235	135.8	0.5778723404
88	0.13	244	170.9	0.7004098361
89	0.12	248	186.7	0.7528225806
90	0.11	248	191.1	0.7705645161
91	0.10	248	191.7	0.7729838710
92	0.09	248	191.7	0.7729838710
93	0.08	248	191.7	0.7729838710
94	0.07	248	191.7	0.7729838710
95	0.06	248	191.7	0.7729838710
96	0.05	248	191.7	0.7729838710
97	0.04	248	191.7	0.7729838710
98	0.03	248	191.7	0.7729838710
99	0.02	248	191.7	0.7729838710
100	0.01	248	191.7	0.7729838710
101	0.00	248	191.7	0.7729838710

## 5 Creating Table 1 from our paper

To generate the table found in our paper, we first compared our resulting LEAP and simple correlation based-networks with the true network taken from FunCoup (Schmitt *et al.*, 2013). The results of this analysis were assembled into tables such as the one shown below, where each row shows the number of correctly identified known associations up to the current row ("Count"), the MAC value, and the row index and column indices for the current association. See `TP_counts` below:

```
> TP_counts[10:20,]
      Count      MAC Row Column
10      1 0.7696197  73    74
11      2 0.7665432  74   220
12      2 0.7660433 110   358
13      2 0.7632045  61   453
14      2 0.7596131 220   541
15      2 0.7595162  38   358
16      2 0.7514357 110   111
17      3 0.7486041 275   345
18      4 0.7476178  38   359
19      4 0.7468748 110   359
20      4 0.7448159  61    74
>
```

We then counted the number of correctly identified known associations for the cutoffs 0.18-0.23 for LEAP and simple correlation, resulting in the counts shown below (note only the code for LEAP's analysis is shown for simplicity):

```
> cutoffs= c(0.23,0.22,0.21,0.20,0.19,0.18)
> true.cor.leap=rep(NA, 6)
> for(i in (1:length(cutoffs))){
+
+   inds = which(TP_counts[,2]<= cutoffs[i])
+   max = max(TP_counts[inds,2])
+   ind = which(TP_counts[,2]==max)
+
+   true.cor.leap[i] = TP_counts[ind,1]
+
+ }
> true.cor.leap
[1] 2394 3315 4686 6767 9508 12989
> true.cor.sim
[1] 1313 1508 1751 2022 2367 2804
```

Next we counted for each cutoff the number of lags greater than 0 and 50. We do this by first finding the rows in our network comparison tables containing a correctly identified association. Note that `lags` here is the lag matrix outputted by the `MAC_counter()` function.

```
> # get rows where true association is found #
>
> inds.trues=c()
> for(n in(1:(nrow(TP_counts)-1))){
+
+   if(TP_counts[n,1]<TP_counts[n+1,1]){inds.trues = c(inds.trues, n+1)}
+
+ }
```

```

+ }
> lag.count.0 = rep(NA, 6)
> lag.count.50 = rep(NA, 6)
> TP_trues = TP_counts[inds.trues,]
> # count the number of lags at each cutoff #
>
> for(j in (1:length(cutoffs))) {
+
+   inds = which(TP_trues[,2] >= cutoffs[j])
+
+   row.inds = TP_trues[inds,3]
+   col.inds = TP_trues[inds,4]
+   lag.cutoff = rep(NA, length(row.inds))
+
+   for(k in (1:length(row.inds))) {
+
+     lag.cutoff[k] = lags[row.inds[k], col.inds[k]]
+
+   }
+
+   lag.count.0[j] = length(which(lag.cutoff > 0))
+   lag.count.50[j] = length(which(lag.cutoff > 50))
+
+ }
>

```

We then count the total number of observations for LEAP and simple correlation that were identified at each cutoff, and finally pull all of these results together to create Table 1. Note that the FDR values are pulled directly from the `MAC_perm()` function output.

```

> Table1 = cbind(cutoffs, num.cor.sim, true.cor.sim, num.cor.leap,
+               true.cor.leap, lag.count.0, lag.count.50, fdr)
> Table1

```

	cutoffs	num.cor.sim	true.cor.sim	num.cor.leap	true.cor.leap	lag.count.0
[1,]	0.23	8494	1313	14735	2394	911
[2,]	0.22	9640	1508	20405	3315	1661
[3,]	0.21	11101	1751	28843	4686	2818
[4,]	0.20	12942	2022	41778	6767	4691
[5,]	0.19	15142	2367	59556	9508	7204
[6,]	0.18	17761	2804	82424	12989	10446

```


```

	lag.count.50	fdr
[1,]	526	0.002373495
[2,]	1007	0.005096089
[3,]	1744	0.010747098
[4,]	2927	0.021010715
[5,]	4579	0.047542090
[6,]	6746	0.082107443

```

>

```

## 6 Generating matrices for use in WGCNA

If you intend to use the results from LEAP for further analysis with WGCNA (Langfelder and Horvath, 2008), then you will require a symmetric matrix of correlations. LEAP will compute this matrix by setting `symmetric=T`.

This creates two files, `lag_symmetric_example.csv` and `MAC_symmetric_example.csv`. LEAP finds this matrix by comparing the correlation of each (i,j) and (j,i) pair, and keeping the value with the maximum absolute correlation. In our example, the pair gene 1 and gene 2 have correlation -0.17 when gene 2 is the lagged gene, (1,2), and correlation 0.14 when gene 1 is the lagged gene, (2,1), then in the symmetric matrix (1,2)=(2,1) = -0.17. Below are the results when we find the symmetric matrix for our example dataset:

```
> output=MAC_counter(data=example_data, max_lag_prop=1/3, file_name="example", lag_matrix=T, symmetric=T)
> MAC_symmetric[1:5,1:5]
      V1      V2      V3      V4      V5
1      NA -0.1728618 0.4182790 0.2526990 -0.1765579
2 -0.1728618      NA 0.1583516 -0.1593688 0.2420226
3 0.4182790 0.1583516      NA 0.2224057 0.1622316
4 0.2526990 -0.1593688 0.2224057      NA 0.1699593
5 -0.1765579 0.2420226 0.1622316 0.1699593      NA
```

## 7 Session information

```
> sessionInfo()

R version 3.2.2 (2015-08-14)
Platform: x86_64-w64-mingw32/x64 (64-bit)
Running under: Windows 7 x64 (build 7601) Service Pack 1

locale:
 [1] LC_COLLATE=English_United States.1252
 [2] LC_CTYPE=English_United States.1252
 [3] LC_MONETARY=English_United States.1252
 [4] LC_NUMERIC=C
 [5] LC_TIME=English_United States.1252

attached base packages:
 [1] stats      graphics  grDevices  utils      datasets  methods   base

other attached packages:
 [1] LEAP_0.2

loaded via a namespace (and not attached):
 [1] tools_3.2.2
```

## 8 Citation information

### References

- [Langfelder and Horvath, 2008] Langfelder, Peter and Horvath, Steve (2008) WGCNA: an R package for weighted correlation network analysis. *BMC bioinformatics*, **9**(1), 1.
- [Schmitt *et al.*, 2013] Schmitt, T., Ogris, C., and Sonnhammer, E. L. (2013) FunCoup 3.0: database of genome-wide functional coupling networks. *Nucleic Acids Research*, **42(Database issue)**, D821-8.
- [Shalek *et al.*, 2014] Shalek AK, Satija R., Shuga J., Trombetta J.J. et al. (2014) Single-cell RNA-seq reveals dynamic paracrine control of cellular variation. *Nature*, **510(7505)**, 363-369.
- [Trapnell *et al.*, 2014] Trapnell, Cole and Cacchiarelli, Davide and Grimsby, Jonna and Pokharel, Prapti and Li, Shuqiang and Morse, Michael and Lennon, Niall J. and Livak, Kenneth J. and Mikkelsen, Tarjei S. and Rinn, John L. (2014) The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells, *Nature Biotechnology*, **32(4)**, 381-386.