

The fastLiquidAssociation Package

Tina Gunderson

April 27, 2020

1 Introduction

The fastLiquidAssociation package is intended to provide more efficient tools for genome-wide application of liquid association. The term liquid association (LA) was coined by K-C. Li in 2002 to describe a theory of co-expression dynamics which “conceptualize[s] the internal evolution of co-expression pattern for a pair of genes” [3]. He used the term “liquid” here (as opposed to “solid”) to convey the idea of a quantity which examined the extent to which the co-expression or correlation of two genes X_1 and X_2 might be mediated by the expression of a third gene X_3 , i.e. the degree to which the correlation of the two genes X_1 and X_2 might vary (e.g. go from contra- to co-expressed or vice-versa) based on the expression level of the third gene (with the third gene’s expression functioning as a representation of the cellular state).

Building on Li’s work, Ho et al. proposed a modification to the liquid association statistic in order to better capture more complex interdependencies between the variables, noting that, “[w]hen the conditional means and variances [of X_1 and X_2] also depend on X_3 , then the three-product-moment measure, as derived by Li for standardized variables, no longer precisely estimates the expected change of co-expression of X_1 and X_2 with respect to X_3 , but captures a larger set of interdependencies among the three variables” [1].

This software package provides an expansion of the functions available in the LiquidAssociation package. While the LiquidAssociation package refers to modified liquid association (MLA) as generalized liquid association (GLA), the 2011 publication refers to it as modified liquid association. For the purposes of this vignette, we will use the two terms interchangeably. The main function (**fastMLA**) reduces the processing power and memory needed to calculate MLA values for a genome ($\binom{N}{3} * 3$, with N being the number of genes in the genome, e.g. 1.079×10^{11} for 6000 genes) by using a pre-screening method to reduce the candidate pool to triplets likely to have a high MLA value. It does this using matrix algebra to create an approximation to the direct MLA estimate for all possible pairs of $X_1, X_2 | X_3$. Intuitively, we can picture MLA as the change in correlation between X_1 and X_2 when X_3 is at a high level of expression versus their correlation when X_3 is at a low level of expression. This difference, ρ_{diff} ($\rho_{diff} = \rho_{high} - \rho_{low}$), gives an approximation of the relative magnitude of the MLA value for most triplet combinations (ρ_{diff} values range from -2 to 2 compared to MLA values which range from $-\sqrt{\pi/2}$ to $\sqrt{\pi/2}$). Figure 1 shows a comparison using the first 50 and the first 250 genes of the Spellman et al. data set (available in the yeastCC package) [4] for all possible triplets of $|\widehat{MLA}|$ vs. $|\rho_{diff}|$ with loess curves.

The package also incorporates functions to test for significance of the triplets returning a high MLA value and provides a faster, more robust version of the MLA bootstrapping estimation procedure for use with multicore systems.

2 Simple Usage

Here we present an example of package usage on a smaller data set. We start by loading the R package and the example data. In this package, we use the yeast cell cycle gene expression data by Spellman [4]. The data can be obtained through package yeastCC. The annotation package for the yeast experiment *org.Sc.sgd.db* can be obtained through Bioconductor.



Figure 1: Comparison for all triplets possible using the first 50 and 250 genes of $|\widehat{MLA}|$ vs. $|\rho_{diff}|$ with lowess curves

2.1 fastMLA

```
> library(fastLiquidAssociation)
> suppressMessages(library(yeastCC))
> suppressMessages(library(org.Sc.sgd.db))
> data(spyCCES)
> lae <- spyCCES[, -(1:4)]
> ### get rid of samples with high % NA elements
> lae <- lae[apply(is.na(exprs(lae)), 1, sum) < ncol(lae)*0.3, ]
> data <- t(exprs(lae))
> data <- data[, 1:50]
> dim(data)
```

```
[1] 73 50
```

After removing genes with high missing percentage, for speed in this example, we reduce the data set to the first 50 genes. A normal score transformation of the data as described in both Li and Ho et al.'s work [3] [1] is performed within the **fastMLA** function. The top values are returned sorted by $|\widehat{MLA}|$.

```
> detectCores()
```

```
[1] 40
```

```
> example <- fastMLA(data=data, topn=50, nvec=1:5, rvalue=0.5, cut=4, threads=detectCores())
```

Allowing parallel execution with up to 40 working processes.

```
> example[1:5, ]
```

	X1 or X2	X2 or X1	X3	rhodiff	MLA value
1	YAL041W	YAL046C	YAL004W	-1.14880	-0.42498
2	YAL041W	YAL042W	YAL004W	-1.22538	-0.40982
3	YAL014C	YAL015C	YAL004W	-1.09253	-0.40199
4	YAL039C	YAL045C	YAL004W	-1.03147	-0.39181
5	YAL025C	YAL035W	YAL002W	0.94254	0.39086

The arguments to the **fastMLA** function are a matrix of numerical data with genes as columns (*data*), the number of results to return (*topn*), the column numbers of the genes to test (*nvec*), the $|\rho_{diff}|$ cutoff to use

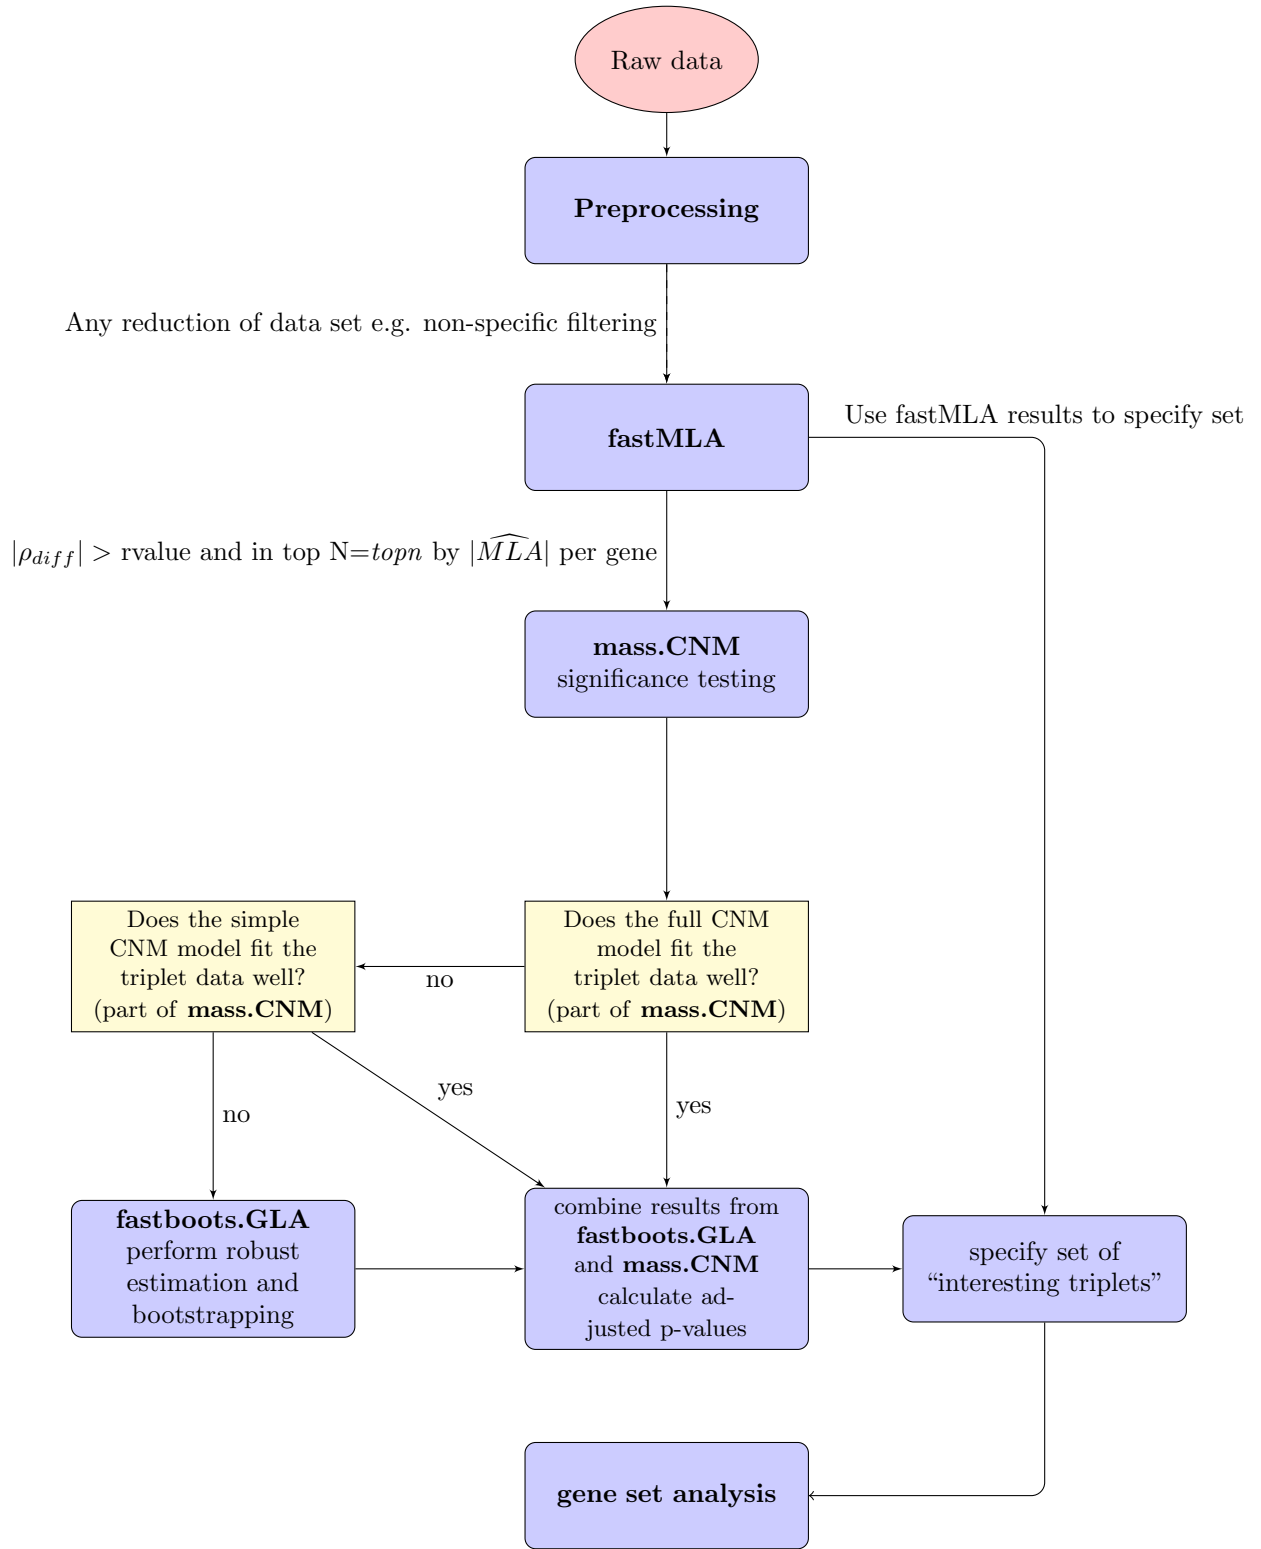


Figure 2: Process map for data processing, testing, and gene set analysis using all components of the *fastLiquidAssociation* package

(*rvalue*), the *cut* argument to pass to the GLA function from the LiquidAssociation package, and the number of processors available to use for increasing the speed of calculation for correlation values (*threads*) (for a fuller discussion of this argument please see the WGCNA package documentation, for a discussion of parallelization please see the section in this vignette entitled "Parallelization of the Direct Estimate"). The *cut* argument will depend on the number of observations in the data. Based on data from Ho et al. [1], the number of observations per bin should be in the 15-30 range for maximum specificity. The *cut* argument is equal to the number of bins plus one. Hence in this example, as we have 73 observations, we chose to set the *cut* value at 4. The maximum *rvalue* is theoretically 2 (as $\rho_{diff} = |\rho_{high} - \rho_{low}|$ and $-1 \leq \rho_{X_1, X_2} \leq 1$). The default value is set at 0.5 or 25% of the realizable correlation difference, however $|\rho_{diff}|$ a.k.a. *rvalue* was set at 1.1 in our example above for speed purposes. Caution should be used when determining the *rvalue*. Too high a value risks missing those triplets whose MLA values are not fully reflected by the more simplistic correlation, while too low a value approaches testing all possible triplets and forfeits any increase in testing efficiency.¹

2.2 Functions of the Conditional Normal Model

Two options for calculating significance are included in this package, **mass.CNM** which relies on estimation of parameters from a conditional normal model and **fastboots.GLA** which uses a more robust direct estimate method. Both functions return unadjusted p-values. The CNM model is written as:

$$\begin{aligned} X_3 &\sim N(\mu_3, \sigma_3^2) \\ X_1, X_2 | X_3 &\sim N\left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \Sigma\right). \end{aligned}$$

where $\Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}$. The mean vector (μ_1, μ_2) and variance matrix Σ depend on the level of X_3 as written below:

$$\begin{aligned} \mu_1 &= \beta_1 X_3, \\ \mu_2 &= \beta_2 X_3, \\ \log \sigma_1^2 &= \alpha_3 + \beta_3 X_3, \\ \log \sigma_2^2 &= \alpha_4 + \beta_4 X_3, \\ \log \left[\frac{1+\rho}{1-\rho} \right] &= \alpha_5 + \beta_5 X_3. \end{aligned}$$

```
> #from our example with fastMLA
> CNMcalc <- mass.CNM(data=data, GLA.mat=example, nback=5)
> CNMcalc
```

```
$`top p-values`
```

	X1 or X2	X2 or X1	X3	rhodiff	MLA value	estimates	san.se	wald	p value
1	YAL011W	YAL017W	YAL005C	0.79	0.30	0.70	0.17	17.14	0
2	YAL040C	YAL048C	YAL005C	0.87	0.31	0.67	0.20	10.87	0
3	YAL031C	YAL036C	YAL002W	-0.78	-0.31	-0.79	0.24	10.42	0
4	YAL028W	YAL041W	YAL004W	0.94	0.30	1.22	0.38	10.03	0
5	YAL041W	YAL046C	YAL004W	-1.15	-0.42	-1.55	0.50	9.51	0
model									
1	F								
2	F								
3	F								
4	F								
5	F								

¹The fastMLA function incorporates the more efficient correlation calculation of the **cor** function in the WGCNA package [2].

```
$`bootstrap triplets`
NULL
```

In our original subset of the data, there are no triplets not returning sensible values, so to demonstrate what is returned when sensible values are not returned for either model, we change the dataset used to the full 5721 genes in the *yeastCC* data set and use a matrix which specifically contains triplets not returning sensible values.

```
> fulldata <- t(exprs(lae))
> load(system.file('data', 'testmat.RData', package='fastLiquidAssociation'))
> notsense <- testmat
> CNMother <- mass.CNM(data=fulldata, GLA.mat=notsense, nback=5)
> CNMother

$`top p-values`
  X1 or X2 X2 or X1      X3 rhodiff MLA value estimates san.se    wald p value
1  YOL120C YOL143C YDR495C  1.2710   0.5376    1.5562 0.4030 14.9116 0.0001
2  YHR129C YPR045C YDR495C  1.4467   0.5474    2.4887 0.6495 14.6822 0.0001
3  YIR043C YLR375W YDR495C  1.3415   0.5358    1.7832 0.5269 11.4531 0.0007
4  YIR004W YOR291W YDR495C -1.4788  -0.5662   -2.1373 0.6774  9.9541 0.0016
5  YCRX15W YDR257C YDR495C -1.3919  -0.5435   -2.1025 0.6693  9.8679 0.0017

model
1      F
2      F
3      F
4      F
5      F

$`bootstrap triplets`
  X1 or X2 X2 or X1  X3      rhodiff GLA value estimates
[1,] "YKR032W" "YNR021W" "YDR294C" "1.2808" "0.4567" "22.1364"
[2,] "YBR012C" "YDL191W" "YDR346C" "1.4023" "0.4298" "-10211.647"

  san.se      wald    p value model
[1,] "20.0702"    "1.2165" "0.27"   "S"
[2,] "15288623.7058" "0"    "0.9995" "S"
```

In the above example, we apply the **mass.CNM** function to the results of **fastMLA** to test for significance using both of the conditional normal model-based approaches laid out in Ho et al. [1]. The **mass.CNM** function accepts three arguments: *data*, *GLA.mat*, and *nback*. *data* is the same matrix used as *data* in the **fastMLA** function, the *GLA.mat* is the returned matrix from a call to the **fastMLA** function, and *nback* is the number of results to return, sorted by p-value, then by size of the Wald statistic. The main parameter of interest in the CNM for significance testing using liquid association is b_5 .

The **mass.CNM** incorporates a two step process, first testing the triplets against the full CNM model and then for any triplets which do not return sensible values (sensible being defined as any triplet for which the p-value or SE estimate=0, SE estimate >10, or any β_5 statistics returned “NaN”, “NA”, or “Inf”), it tests them against the simple model. The full CNM model uses all parameter value (i.e. estimates the parameters of the model above using all five equations), while the simple model only uses the last three.

The **mass.CNM** output is a list with two components, one a data frame with *nback* rows which list the information from fastMLA (the triplet, the triplet’s GLA and ρ_{diff} values), the β_5 values estimated in the model (the estimate for β_5 , its SE, the Wald statistic, and p-value), and the model (full or simple) which was used to calculate the values and the second a list of any triplets who did not return sensible values for either model. For any triplets that do not return a sensible value using one of the two models, the more robust direct estimate method is available as the **fastboots.GLA** function. ²

²The above example may produce a warning message “In sqrt(diag(object\$alpha)) : NaNs produced”. This is normal behavior when one or both of the CNM models cannot fit at least one of the triplets.

2.3 Parallelization of the Direct Estimate

As mentioned before, for triplets that when tested with the full or simple models did not return sensible values, a direct estimator is available. The **fastboots.GLA** function makes use of the increased processing power that has become more widely available, both on standard desktops and servers. It is an updated and parallelized version of the **getsGLA** function available in the **LiquidAssociation** package. As the standard error in the direct estimate relies on bootstrapping and the significance calculations can require a large number of iterations, we sought a more efficient method to increase efficiency.

Parallelization has as its basis the idea that large problems (e.g. a high number of calculations) can be split into smaller problems that are being run at the same time, rather than sequentially. Parallelization of processes makes sense when each process itself takes significantly longer to complete than to allocate. In the below example, we use the **detectCores** function from the **parallel** package to detect the number of CPU cores. On systems where more processors are available, e.g. servers for supercomputing which normally have multiple nodes of several processors each, it is possible to set a larger number of cores than the **detectCores** function may detect.

fastboots.GLA is intended to be used with the results of a call to **mass.CNM**, specifically the output labelled bootstrap triplets. The **fastboots.GLA** function is intended for use on multicore systems in order to make use of possible parallelization. While it can be used on a single core system, there would be no speed benefit over the **getsGLA** function available in the **LiquidAssociation** package, though the **fastboots.GLA** function does resolve an error in the **getsGLA** function that can be caused by data redundancy. In the below example, we are using a six core system.

```
> #determine number of processors for multicore systems
> clust <- makeCluster(detectCores())
> registerDoParallel(clust)
> boottrips <- CNMother[[2]]
> dim(boottrips)
```

```
[1] 2 10
```

Below we an example of processing speed with the **fastboots.GLA** function. Using the **getsGLA** function from the **LiquidAssociation** package, the triplet had the following values: ³

user	system	elapsed	sGLA	p value
30.85	0.01	30.92	5.25	0.00000

```
> #We take the results for the single triplet and put it in matrix format
> example.boots <- boottrips[1, , drop = FALSE]
> dim(example.boots)
```

```
[1] 1 10
```

```
> set.seed(1)
> system.time(GLAnew <- fastboots.GLA(tripmat=example.boots,data=fulldata,
+ clust=clust, boots=30, perm=500, cut=4))
```

user	system	elapsed
1.64	0.00	253.32

```
> GLAnew
```

	X1 or X2	X2 or X1	X3 rhodiff	MLA value	MLA stat	boots	p-value
1	YKR032W	YNR021W	YDR294C	1.2808	0.4567	5.180175	0

³The original **getsGLA** code is not called here as it sometimes produces an error due to non-unique cut values with higher numbers of permutations.

```
> #the matrix conversion is not needed for the 2 line result
> set.seed(1)
> system.time(GLAtwo <- fastboots.GLA(tripmat=boottrips,data=fulldata,
+ clust=clust, boots=30, perm=500, cut=4))
```

```
      user  system elapsed
      1.90    0.12    6.71
```

```
> GLAtwo
```

```
      X1 or X2 X2 or X1      X3 rhodiff MLA value MLA stat boots p-value
1  YKR032W  YNR021W YDR294C  1.2808    0.4567 5.180175    0.00000000
2  YBR012C  YDL191W YDR346C  1.4023    0.4298 2.387118    0.05252525
```

```
> #close the cluster
> stopCluster(clust)
```

tripmat here is the list of triplets where a direct estimate is needed. The *data* is the numeric matrix of data used in both the **fastMLA** and **mass.CNM** functions. *clust* is a cluster of CPU cores to use in parallelization, created by a call to the **makeCluster** function from package **parallel**. The *boots* argument specifies the number of bootstrap iterations for estimating the bootstrap standard error. The *perm* argument specifies the number of iterations to calculate the permuted p-value. Given that the results of the permutations depend on random number generation, for reproducibility here we use the *set.seed* option. The *cut* option is used to determine the number of “buckets” to used in the estimation procedure.

Because the normalization of the data matrix is incorporated into the **fastboots.GLA** function, the full time decrease is not apparent until the number of iterations significantly exceeds the time to normalize the data matrix or the number of triplets to test increases. This is partially displayed in the example, where even the calculation of two triplets has a lower processing time than the single triplet run serially.

3 Further Applications

In this section, we cover further applications available with the results of the **fastLiquidAssociation** package. Using the **GOstats** package in conjunction with the results from either the **fastMLA**, the **mass.CNM**, or **fastboots.GLA** we can specify our set of interesting genes and perform GO analysis. We show an example here using the genes in the X_3 position of the results from **fastMLA** to define our list of interesting genes and a p-value cutoff of 0.05 for the BP ontology. Only categories containing a minimum of five genes are displayed.

```
> library(GOstats)
> library("org.Sc.sgd.db")
> ##X3 genes
> topX3 <- unique(example[,3])
> hyp.cutoff <- 0.05
> ####
> params <- new("GOHyperGParams", geneIds=topX3, universeGeneIds=colnames(data),
+ annotation="org.Sc.sgd.db", ontology="BP", pvalueCutoff=hyp.cutoff, conditional=TRUE,
+ testDirection="over")
> GOout <- hyperGTest(params)
> summary(GOout, categorySize=5)
```

```
      GOBPID      Pvalue OddsRatio ExpCount Count Size
1  GO:0034613 0.02451434  15.85714  0.8333333    3   10
2  GO:0043933 0.03307123  13.50000  0.9166667    3   11
3  GO:0033036 0.04324699  11.66667  1.0000000    3   12
```

	Term
1	cellular protein localization
2	protein-containing complex subunit organization
3	macromolecule localization

Because we used open reading frame (ORF) IDs in order to use the GOstats package with yeast data, we can the following to unpack the associated gene names.

```
> ###extracts GO list elements of summary(hyperGtestobj)<cutoff
> ###converts ORFids to Gene names and returns under GO list element
> ##for ontology BP
> GOids <- summary(GOout,categorySize=5)$GOBPID
> check <- GOout@goDag@nodeData@data
> subset <- check[GOids]
> terms <- summary(GOout,categorySize=5)$Term
> test <- sapply(subset,function(m) m[1])
> orflist <- lapply(test,function(x) intersect(x,topX3))
> ##creates mapping of ORFids to gene names
> x <- org.Sc.sgdGENENAME
> mappedgenes <- mappedkeys(x)
> xx <- as.list(x[mappedgenes])
> mapid <- names(xx)
> ##creates list of GO ids
> genename1 <- lapply(orflist,function(x) xx[match(x,mapid)])
> ###
> ### if reduced num of terms desired run for statement below ##
> # for(i in 1:length(genename1)){
> # if (length(genename1[[i]])>10) genename1[[i]] <- genename1[[i]][1:10]
> # }
> ## for full list use below
> genelist <- lapply(genename1,function(x) paste(x,collapse=" ", ""))
> ugenes <- unlist(genelist)
> names <- sapply(names(ugenes), function(x) unlist(strsplit(x, split='.', fixed=TRUE))[1])
> umat <- matrix(ugenes)
> umat <- cbind(terms,umat)
> rownames(umat) <- names
> colnames(umat) <- c("GO description","Associated genes")
> umat
```

	GO description	Associated genes
G0:0034613	"cellular protein localization"	"TFC3, VPS8, SSA1"
G0:0043933	"protein-containing complex subunit organization"	"VPS8, EFB1, SSA1"
G0:0033036	"macromolecule localization"	"TFC3, VPS8, SSA1"

4 References

References

- [1] Yen-Yi Ho, Giovanni Parmigiani, Thomas A Louis, and Leslie M Cope. Modeling liquid association. *Biometrics*, 67(1):133–141, 2011.
- [2] Peter Langfelder and Steve Horvath. Wgcna: an r package for weighted correlation network analysis. *BMC bioinformatics*, 9(1):559, 2008.
- [3] Ker-Chau Li. Genome-wide coexpression dynamics: theory and application. *Proceedings of the National Academy of Sciences*, 99(26):16875–16880, 2002.

- [4] Paul T Spellman, Gavin Sherlock, Michael Q Zhang, Vishwanath R Iyer, Kirk Anders, Michael B Eisen, Patrick O Brown, David Botstein, and Bruce Futcher. Comprehensive identification of cell cycle-regulated genes of the yeast *saccharomyces cerevisiae* by microarray hybridization. *Molecular biology of the cell*, 9(12):3273–3297, 1998.

5 Session Information

R version 4.0.0 (2020-04-24)

Platform: x86_64-w64-mingw32/x64 (64-bit)

Running under: Windows Server 2012 R2 x64 (build 9600)

Matrix products: default

locale:

- [1] LC_COLLATE=C
- [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] stats4 parallel stats graphics grDevices utils datasets
- [8] methods base

other attached packages:

- [1] GO.db_3.10.0 GOstats_2.54.0
- [3] graph_1.66.0 Category_2.54.0
- [5] Matrix_1.2-18 fastLiquidAssociation_1.24.0
- [7] Hmisc_4.4-0 ggplot2_3.3.0
- [9] Formula_1.2-3 survival_3.1-12
- [11] lattice_0.20-41 doParallel_1.0.15
- [13] iterators_1.0.12 foreach_1.5.0
- [15] LiquidAssociation_1.42.0 org.Sc.sgd.db_3.10.0
- [17] AnnotationDbi_1.50.0 IRanges_2.22.0
- [19] S4Vectors_0.26.0 yeastCC_1.27.0
- [21] Biobase_2.48.0 BiocGenerics_0.34.0
- [23] geepack_1.3-1 preprocessCore_1.50.0
- [25] impute_1.62.0 WGCNA_1.69
- [27] fastcluster_1.1.25 dynamicTreeCut_1.63-1

loaded via a namespace (and not attached):

- [1] nlme_3.1-147 bitops_1.0-6 matrixStats_0.56.0
- [4] bit64_0.9-7 RColorBrewer_1.1-2 Rgraphviz_2.32.0
- [7] tools_4.0.0 backports_1.1.6 R6_2.4.1
- [10] rpart_4.1-15 DBI_1.1.0 colorspace_1.4-1
- [13] nnet_7.3-14 withr_2.2.0 tidyselect_1.0.0
- [16] gridExtra_2.3 bit_1.1-15.2 compiler_4.0.0
- [19] htmlTable_1.13.3 scales_1.1.0 checkmate_2.0.0
- [22] genefilter_1.70.0 RBGL_1.64.0 stringr_1.4.0
- [25] digest_0.6.25 foreign_0.8-79 AnnotationForge_1.30.0
- [28] base64enc_0.1-3 jpeg_0.1-8.1 pkgconfig_2.0.3
- [31] htmltools_0.4.0 htmlwidgets_1.5.1 rlang_0.4.5
- [34] rstudioapi_0.11 RSQLite_2.2.0 generics_0.0.2

[37]	acepack_1.4.1	dplyr_0.8.5	RCurl_1.98-1.2
[40]	magrittr_1.5	Rcpp_1.0.4.6	munsell_0.5.0
[43]	lifecycle_0.2.0	stringi_1.4.6	MASS_7.3-51.6
[46]	grid_4.0.0	blob_1.2.1	crayon_1.3.4
[49]	splines_4.0.0	annotate_1.66.0	knitr_1.28
[52]	pillar_1.4.3	codetools_0.2-16	XML_3.99-0.3
[55]	glue_1.4.0	latticeExtra_0.6-29	data.table_1.12.8
[58]	png_0.1-7	vctrs_0.2.4	gtable_0.3.0
[61]	purrr_0.3.4	tidyr_1.0.2	assertthat_0.2.1
[64]	xfun_0.13	xtable_1.8-4	broom_0.5.6
[67]	tibble_3.0.1	memoise_1.1.0	cluster_2.1.0
[70]	ellipsis_0.3.0	GSEABase_1.50.0	