

Pigengene: Computing and using eigengenes

Habil Zare

Modified: 26 April, 2016. Compiled: January 4, 2019

Contents

1	Introduction	1
2	How to run <i>Pigengene</i> ?	2
2.1	Installation	2
2.2	A quick overview	2
2.3	What is an eigengene?	3
2.4	A toy example	3
2.5	Citation	7
3	Session Information	8

1 Introduction

Gene expression profiling technologies such as microarray or RNA Seq provide valuable datasets, however, inferring biological information from these data remains cumbersome. *Pigengene* address two challenges:

1. **Curse of dimensionality:** The number of features in an expression profile is usually very high. For instance, there are about 20,000 genes in human. In contrast, the number of samples (patients) is often very limited in practice, and may not exceed a few hundreds. Yeung et al. have shown that standard data reduction methods such as principal component analysis (PCA) are not appropriate to directly apply on gene expression data [1]. Instead, *Pigengene* addresses this challenge by applying PCA on gene modules.

2. **Normalization:** Data produced using different technologies, or in different labs, are not easily comparable. *Pigengene* identifies *eigengenes*, informative biological signatures that are robust with respect to the profiling platform. For instance, it can identify the signatures (compute the eigengenes) on microarray data, and infer them on biologically-related RNA Seq data. The resulting signatures are directly comparable even if the set of samples (patients) are independent and disjoint in the two analyzed datasets.

2 How to run *Pigengene* ?

2.1 Installation

Pigengene is an *R* package that can be downloaded and installed from *Bioconductor* by the following commands in *R*:

```
if (!requireNamespace("BiocManager", quietly=TRUE)) install.packages("BiocManager")
BiocManager::install("Pigengene")
```

Alternatively, if the source code is already available, the package can be installed by the following command in Linux:

```
R CMD INSTALL Pigengene_x.y.z.tar.gz
```

where x.y.z. determines the version. The second approach requires all the dependencies be installed manually, therefore, the first approach is preferred.

2.2 A quick overview

Pigengene identifies gene modules (clusters), computes an eigengene for each module, and uses these biological signatures as features for classification. The main function is `one.step.pigengene` which requires a gene expression profile and the corresponding conditions (types). Individual functions are also provided to facilitate running the pipeline in a customized way. The inferred biological signatures (eigengenes) are useful for supervised or unsupervised analyses.

2.3 What is an eigengene?

In most functions of this package, eigenegenes are computed or used as robust biological signatures. Briefly, each eigengene is a weighted average of the expression of all genes in a given set of genes (also known as a gene module or a cluster of genes). The weights are adjusted in a way that the explained variance is maximized. This guarantees that the loss in the biological information is minimized.

2.4 A toy example


For a quick start, the application of [Pigengene](#) pipeline on some leukemia dataset is demonstrated below [2]. The first step is to load the package and data in R:

```
library(Pigengene)

## Loading required package: graph
## Loading required package: BiocGenerics
## Loading required package: parallel
##
## Attaching package: 'BiocGenerics'
## The following objects are masked from 'package:parallel':
##
##   clusterApply, clusterApplyLB, clusterCall, clusterEvalQ,
##   clusterExport, clusterMap, parApply, parCapply, parLapply,
##   parLapplyLB, parRapply, parSapply, parSapplyLB
## The following objects are masked from 'package:stats':
##
##   IQR, mad, sd, var, xtabs
## The following objects are masked from 'package:base':
##
##   Filter, Find, Map, Position, Reduce, anyDuplicated, append,
##   as.data.frame, basename, cbind, colMeans, colSums, colnames,
##   dirname, do.call, duplicated, eval, evalq, get, grep, grepl,
##   intersect, is.unsorted, lapply, lengths, mapapply, match, mget,
##   order, paste, pmax, pmax.int, pmin, pmin.int, rank, rbind,
##   rowMeans, rowSums, rownames, sapply, setdiff, sort, table,
##   tapply, union, unique, unsplit, which, which.max, which.min
```

Pigengene: Computing and using eigengenes

```
##  
data(aml)  
data(mds)  
d1 <- rbind(aml,mds)  
Labels <- c(rep("AML",nrow(aml)),rep("MDS",nrow(mds)))  
names(Labels) <- rownames(d1)  
Disease <- as.data.frame(Labels)  
h1 <- heatmap.type(d1[,1:20],annRow=Disease,show_rownames=FALSE)
```



Please note that the provided data in the package is sub-sampled for a quicker demonstration. For real applications, the expression of thousands of genes should be provided in order to the co-expression network analysis to be appropriate. It is common to first perform differential expression analysis, sort all the genes based on p-values, and use the top-third as the input [3]. Analyzing such input with *Pigengene* can take a few hours and may require 5-10 GB of memory. The following command runs *Pigengene* pipeline on the toy data:

Pigengene: Computing and using eigengenes

```
p1 <- one.step.pigengene(Data=d1,saveDir='pigengene', bnNum=0, verbose=1,
  seed=1, Labels=Labels, toCompact=FALSE, doHeat=FALSE)

## Pigengene started analizing 366 samples using 1000 genes...

## Warning:  executing %dopar% sequentially:  no parallel backend registered

## Pigengenes...

## Pigengene plots in:

## /private/tmp/RtmpzVRDfc/Rbuild5fa8538d591c/Pigengene/vignettes/pigengene/plots

## Making decision trees...

## minPerLeaf:  8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21,
22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37

## costs:

##      AML MDS
## AML   0   1
## MDS   1   0


## toCompact:  FALSE
```

Results and figures are saved in `pigengene` folder under the current directory. For more advanced applications, the user is encouraged to analyze the data step-by-step and customize the individual functions such as `compute.pigenegene` and `make.decision.tree`.

In addition to the provided decision trees, the user can also take alternative approaches to perform classification, clustering, survival analysis, etc. *using eigengenes as robust biological signatures (informative features)*. Eigengenes and other useful objects can be retrieved from the output. For instance, `c5treeRes` is a list containing the results of fitting decision trees to the eigengenes. As shown above, a couple of trees were fitted, one per a value for `minPerLeaf`. The following command plots the tree corresponding to 34, i.e., it was fitted requiring the minimum number of samples per every leaf to be at least 34.

```
plot(p1$c5treeRes$c5Trees[["34"]])
```


Pigengene: Computing and using eigengenes



The tree corresponding to other values are saved in `pigengene` folder. Of note, is the `pigenegene` object that contains the matrix of inferred eigengenes. Each row corresponds to a sample, and each column represents an eigen gene.

```
dim(p1$pigengene$eigenenes)
## [1] 366 4

p1 <- pheatmap.type(p1$pigengene$eigenenes,annRow=Disease,show_rownames=FALSE)
```



2.5 Citation

The methodology and an interesting application of *Pigengene* on studying hematological malignancies is presented in the following reference [4].

```
citation("Pigengene")
```

To cite package 'Pigengene' in publications use:

Amir Foroushani et al.(2016) Large-scale gene network analysis reveals the significance of extracellular matrix pathway and homeobox genes in acute myeloid leukemia: an introduction to the Pigengene package and its applications, Foroushani et al., BMC Medical Genomics. URL: <https://bmcmmedgenomics.biomedcentral.com/articles/10.1186/s12920-017-0253-6>.

A BibTeX entry for LaTeX users is

@Article, author = Amir Foroushani and et al., title = Large-scale gene network analysis reveals the significance of extracellular matrix pathway and homeobox genes in acute myeloid leukemia: an introduction to the Pigengene package and its applications, journal = BMC Medical Genomics, year = 2017, volume = 10, number = 1, pages = 16, month = 3,

3 Session Information

The output of `sessionInfo` on the system that compiled this document is as follows:

```
toLatex(sessionInfo())
```

- R version 3.5.2 Patched (2018-12-20 r75875),
x86_64-apple-darwin15.6.0
- Locale: C/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
- Running under: OS X El Capitan 10.11.6
- Matrix products: default
- BLAS:
/Library/Frameworks/R.framework/Versions/3.5/Resources/lib/libRblas.0.dylib
- LAPACK:
/Library/Frameworks/R.framework/Versions/3.5/Resources/lib/libRlapack.dylib
- Base packages: base, datasets, grDevices, graphics, methods, parallel,
stats, utils
- Other packages: BiocGenerics 0.28.0, Pigengene 1.8.1, graph 1.60.0
- Loaded via a namespace (and not attached): AnnotationDbi 1.44.0,
Biobase 2.42.0, BiocManager 1.30.4, BiocStyle 2.10.0, C50 0.1.2,
Cubist 0.2.2, DBI 1.0.0, DEoptimR 1.0-8, Formula 1.2-3, GO.db 3.7.0,
Hmisc 4.1-1, IRanges 2.16.0, MASS 7.3-51.1, Matrix 1.2-15, R6 2.3.0,
RColorBrewer 1.1-2, RSQLite 2.1.1, Rcpp 1.0.0, Rgraphviz 2.26.0,
S4Vectors 0.20.1, WGCNA 1.66, acepack 1.4.1, assertthat 0.2.0,
backports 1.1.3, base64enc 0.1-3, bindr 0.1.1, bindrcpp 0.2.2, bit 1.1-14,
bit64 0.9-7, blob 1.1.1, blearn 4.4, checkmate 1.8.5, cluster 2.0.7-1,
codetools 0.2-16, colorspace 1.3-2, compiler 3.5.2, crayon 1.3.4,
data.table 1.11.8, digest 0.6.18, doParallel 1.0.14, dplyr 0.7.8,
dynamicTreeCut 1.63-1, evaluate 0.12, fastcluster 1.1.25,
fit.models 0.5-14, foreach 1.4.4, foreign 0.8-71, ggplot2 3.1.0,

glue 1.3.0, grid 3.5.2, gridExtra 2.3, gtable 0.2.0, highr 0.7,
htmlTable 1.13, htmltools 0.3.6, htmlwidgets 1.3, impute 1.56.0,
inum 1.0-0, iterators 1.0.10, knitr 1.21, lattice 0.20-38,
latticeExtra 0.6-28, lazyeval 0.2.1, libcoin 1.0-2, magrittr 1.5,
matrixStats 0.54.0, memoise 1.1.0, munsell 0.5.0, mvtnorm 1.0-8,
nnet 7.3-12, partykit 1.2-2, pcaPP 1.9-73, pheatmap 1.0.12, pillar 1.3.1,
pkgconfig 2.0.2, plyr 1.8.4, preprocessCore 1.44.0, purrr 0.2.5,
reshape2 1.4.3, rlang 0.3.0.1, rmarkdown 1.11, robust 0.4-18,
robustbase 0.93-3, rpart 4.1-13, rrcov 1.4-7, rstudioapi 0.8, scales 1.0.0,
splines 3.5.2, stats4 3.5.2, stringi 1.2.4, stringr 1.3.1, survival 2.43-3,
tibble 2.0.0, tidyselect 0.2.5, tools 3.5.2, xfun 0.4, yaml 2.2.0

References

- [1] Ka Yee Yeung and Walter L. Ruzzo. Principal component analysis for clustering gene expression data. *Bioinformatics*, 17(9):763–774, 2001.
- [2] Ken I Mills, Alexander Kohlmann, P Mickey Williams, Lothar Wieczorek, Wei-min Liu, Rachel Li, Wen Wei, David T Bowen, Helmut Loeffler, Jesus M Hernandez, et al. Microarray-based classifiers and prognosis models identify subgroups with distinct clinical outcomes and high risk of aml transformation of myelodysplastic syndrome. *Blood*, 114(5):1063–1072, 2009.
- [3] Bin Zhang, Chris Gaiteri, Liviu-Gabriel Bodea, Zhi Wang, Joshua McElwee, Alexei A Podtelezhnikov, Chunsheng Zhang, Tao Xie, Linh Tran, Radu Dobrin, et al. Integrated systems approach identifies genetic nodes and networks in late-onset alzheimer’s disease. *Cell*, 153(3):707–720, 2013.
- [4] A Foroushani, R Agrahari, R Docking, A Karsan, and H Zare. Large-scale gene network analysis reveals the significance of extracellular matrix pathway and homeobox genes in acute myeloid leukemia. *In preparation*.