

Estimating Gene-Specific Phenotypes with `gespeR`

Fabian Schmich

April 21, 2024

Contents

| | | |
|----------|--------------------------------------|----------|
| 1 | The <code>gespeR</code> Model | 1 |
| 2 | Working Example | 1 |
| 3 | <code>sessionInfo()</code> | 7 |

1 The `gespeR` Model

This package provides algorithms for deconvoluting off-target confounded phenotypes from RNA interference screens. The package uses (predicted) siRNA-to-gene target relations in a regularised linear regression model, in order to infer individual gene-specific phenotype (GSP) contributions. The observed siRNA-specific phenotypes (SSPs) for reagent $i = 1, \dots, n$ as the weighted linear sum of GSPs of all targeted genes $j = 1, \dots, p$

$$Y_i = x_{i1}\beta_j + \dots + x_{ip}\beta_p + \epsilon_i, \quad (1)$$

where x_{ij} represents the strength of knockdown of reagent i on gene j , β_j corresponds to the GSP of gene j and ϵ_i is the error term for SSP i . The linear regression model is fit using elastic net regularization:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \left(y_i - \sum_{j=1}^p x_{ij}\beta_j \right)^2 + \lambda \sum_{j=1}^p \left(\alpha\beta_j^2 + (1 - \alpha)|\beta_j| \right) \right\}. \quad (2)$$

Here λ determines the amount of regularization and α is the mixing parameter between the ridge and lasso penalty with $0 \leq \alpha \leq 1$. The elastic net penalty selects variables like the lasso and shrinks together the coefficients of correlated predictors like ridge. This allows for a sparse solution of nonzero GSPs, while retaining simultaneous selection of genes with similar RNAi reagent binding patterns in their respective 3' UTRs. For more information and for citing the `gespeR` package please refer to:

Schmich F (2024). *gespeR: Gene-Specific Phenotype Estimator*. R package version 1.35.0, <http://www.cbg.ethz.ch/software/gespeR>.

2 Working Example

In this example, we load first load simulated phenotypic readout and siRNA-to-gene target relations. The toy data consists of four screens (A, B, C, D) of 1,000 siRNAs and a limited gene universe of 1,500 genes. Detailed description of how the data was simulated can be accessed using `?simData`. First, we load the package:

```
library(gesper)
```

Now the phenotypes and target relations can be initialised using the `Phenotypes` and `TargetRelations` commands. First, we load the four phenotypes vectors:

```
phenos <- lapply(LETTERS[1:4], function(x) {
  sprintf("Phenotypes_screen_%s.txt", x)
})
phenos <- lapply(phenos, function(x) {
  Phenotypes(system.file("extdata", x, package="gesper"),
    type = "SSP",
    col.id = 1,
    col.score = 2)
})
show(phenos[[1]])

## 1000 SSP Phenotypes

## Warning: 'tbl_df()' was deprecated in dplyr 1.0.0.
## i Please use 'tibble::as_tibble()' instead.
## i The deprecated feature was likely used in the gesper package.
## Please report the issue to the authors.
## This warning is displayed once every 8 hours.
## Call 'lifecycle::last_lifecycle_warnings()' to see where this warning was
## generated.


## # A tibble: 1,000 x 2
##   ID          Scores
##   <chr>         <dbl>
## 1 siRNAID_0001 -0.930
## 2 siRNAID_0002 -1.13
## 3 siRNAID_0003 -1.05
## 4 siRNAID_0004  0.808
## 5 siRNAID_0005 -1.42
## 6 siRNAID_0006  1.64
## 7 siRNAID_0007 -0.157
## 8 siRNAID_0008  0.748
## 9 siRNAID_0009 -0.959
## 10 siRNAID_0010 -0.0440
## # i 990 more rows
```

A visual representation of the phenotypes can be obtained with the `plot` method:

```
plot(phenos[[1]])

## 'stat_bin()' using 'bins = 30'. Pick better value with 'binwidth'.
```

SSP Phenotypes



Now, we load the target relations for all four screens using the constructor of the `TargetRelations` class:

```
tr <- lapply(LETTERS[1:4], function(x) {
  sprintf("TR_screen_%s.rds", x)
})
tr <- lapply(tr, function(x) {
  TargetRelations(system.file("extdata", x, package="gesper"))
})
show(tr[[2]])

## 1000 x 1500 siRNA-to-gene relations.
## 10 x 5 sparse Matrix of class "dgCMatrix"
##           colnames
## rownames  geneID_0001 geneID_0002 geneID_0003 geneID_0004 geneID_0005
## siRNAID_0001      .      .      .      .      .
## siRNAID_0002      .      .      .      .      .
## siRNAID_0003      . 0.4385757 .      .      .
## siRNAID_0004      .      .      .      .      .
## siRNAID_0005      .      .      .      .      .
## siRNAID_0006      .      .      .      .      .
## siRNAID_0007      .      .      .      .      .
## siRNAID_0008      .      .      .      .      .
## siRNAID_0009      .      .      .      .      .
## siRNAID_0010      .      .      .      .      .
## ...
```

For large data sets, e.g. genome-wide screens, target relations objects can become very big and the user may not want to keep all values in the RAM. For this purpose, we can use the `unloadValues` method. In this example, we write the values to a temp-file, i.e. not the original source file, which may be required, when we do not want to overwrite existing data, after, for instance, subsetting the target relations object.

```
# Size of object with loaded values
format(object.size(tr[[1]]), units = "Kb")

## [1] "717.7 Kb"

tempfile <- paste(tempfile(pattern = "file", tmpdir = tempdir()), ".rds", sep="")
tr[[1]] <- unloadValues(tr[[1]], writeValues = TRUE, path = tempfile)

# Size of object after unloading
format(object.size(tr[[1]]), units = "Kb")

## [1] "178.3 Kb"

# Reload values
tr[[1]] <- loadValues(tr[[1]])
```

In order to obtain deconvoluted gene-specific phenotypes (GSPs), we fit four models on the four separate data sets using cross validation by setting `mode = "cv"`. We set the elastic net mixing parameter to 0.5 and use only one core in this example:

```
res.cv <- lapply(1:length(phenos), function(i) {
  gesper(phenotypes = phenos[[i]],
        target.relations = tr[[i]],
        mode = "cv",
        alpha = 0.5,
        ncores = 1)
})
```

The `ssp` and `gsp` methods can be used to obtain SSP and GSP scores from a `gesper` object:

```
ssp(res.cv[[1]])

## 1000 SSP Phenotypes
##
## # A tibble: 1,000 x 2
##   ID           Scores
##   <chr>         <dbl>
## 1 siRNAID_0001 -0.930
## 2 siRNAID_0002 -1.13
## 3 siRNAID_0003 -1.05
## 4 siRNAID_0004  0.808
## 5 siRNAID_0005 -1.42
## 6 siRNAID_0006  1.64
## 7 siRNAID_0007 -0.157
## 8 siRNAID_0008  0.748
## 9 siRNAID_0009 -0.959
## 10 siRNAID_0010 -0.0440
## # i 990 more rows

gsp(res.cv[[1]])
```

```
## 1500 GSP Phenotypes
##
## # A tibble: 1,500 x 2
##   ID           Scores
##   <chr>        <dbl>
## 1 geneID_0001     NA
## 2 geneID_0002     NA
## 3 geneID_0003     NA
## 4 geneID_0004     NA
## 5 geneID_0005     NA
## 6 geneID_0006     NA
## 7 geneID_0007     NA
## 8 geneID_0008     NA
## 9 geneID_0009     NA
## 10 geneID_0010    NA
## # i 1,490 more rows

head(scores(res.cv[[1]]))


## # A tibble: 6 x 2
##   ID           Scores
##   <chr>        <dbl>
## 1 geneID_0001     NA
## 2 geneID_0002     NA
## 3 geneID_0003     NA
## 4 geneID_0004     NA
## 5 geneID_0005     NA
## 6 geneID_0006     NA
```

The fitted models can also be visualised using the `plot` method:

```
plot(res.cv[[1]])

## 'stat_bin()' using 'bins = 30'. Pick better value with 'binwidth'.
## Warning: Removed 1169 rows containing non-finite outside the scale range ('stat_bin()').
```

GSP Phenotypes




The `concordance` method can be used to compute the concordance between ranked lists of phenotypes. Here we compute concordance between all pairs of GSPs, as well as between all pairs of SSPs, from all four data sets:

```
conc.gsp <- concordance(lapply(res.cv, gsp))
conc.ssp <- concordance(lapply(res.cv, ssp))
```


We can visualise the `concordance` objects using the `plot` method:

```
plot(conc.gsp) + ggtitle("GSPs\n")
plot(conc.ssp) + ggtitle("SSPs\n")
```

GSPs



SSPs



3 sessionInfo()

- R version 4.4.0 beta (2024-04-15 r86425), x86_64-pc-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_GB, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Time zone: America/New_York
- TZcode source: system (glibc)
- Running under: Ubuntu 22.04.4 LTS
- Matrix products: default
- BLAS: /home/biocbuild/bbs-3.19-bioc/R/lib/libRblas.so
- LAPACK: /usr/lib/x86_64-linux-gnu/lapack/liblapack.so.3.10.0
- Base packages: base, datasets, grDevices, graphics, methods, stats, utils
- Other packages: gespeR 1.35.0, ggplot2 3.5.0, knitr 1.46
- Loaded via a namespace (and not attached): AnnotationDbi 1.65.2, Biobase 2.63.1, BiocFileCache 2.11.2, BiocGenerics 0.49.1, BiocManager 1.30.22, Biostrings 2.71.5, Category 2.69.0, DBI 1.2.2, GSEABase 1.65.1, GenomeInfoDb 1.39.14, GenomeInfoDbData 1.2.12, IRanges 2.37.1, KEGGREST 1.43.0, Matrix 1.7-0, MatrixGenerics 1.15.1, R6 2.5.1, RBGL 1.79.0, RColorBrewer 1.1-3, RSQLite 2.3.6, Rcpp 1.0.12, S4Vectors 0.41.6, UCSC.utils 0.99.7, XML 3.99-0.16.1, XVector 0.43.1, affy 1.81.0, affyio 1.73.0, annotate 1.81.2, biomaRt 2.59.1, bit 4.0.5, bit64 4.0.5, blob 1.2.4, cachem 1.0.8, cellHTS2 2.67.4, cli 3.6.2, codetools 0.2-20, colorspace 2.1-0, compiler 4.4.0, crayon 1.5.2, curl 5.2.1, dbplyr 2.5.0, digest 0.6.35, doParallel 1.0.17, dplyr 1.1.4, evaluate 0.23, fansi 1.0.6, farver 2.1.1, fastmap 1.1.1, filelock 1.0.3, foreach 1.5.2, genefilter 1.85.1, generics 0.1.3, glmnet 4.1-8, glue 1.7.0, graph 1.81.1, grid 4.4.0, gtable 0.3.4, highr 0.10, hms 1.1.3, httr 1.4.7, httr2 1.0.1, hwriter 1.3.2.1, iterators 1.0.14, jsonlite 1.8.8, labeling 0.4.3, lattice 0.22-6, lifecycle 1.0.4, limma 3.59.8, locfit 1.5-9.9, magrittr 2.0.3, matrixStats 1.3.0, memoise 2.0.1, munsell 0.5.1, parallel 4.4.0, pillar 1.9.0, pkgconfig 2.0.3, plyr 1.8.9, png 0.1-8, preprocessCore 1.65.0, prettyunits 1.2.0, progress 1.2.3, rappdirs 0.3.3, reshape2 1.4.4, rlang 1.1.3, scales 1.3.0, shape 1.4.6.1,

splines 4.4.0, splots 1.69.0, statmod 1.5.0, stats4 4.4.0, stringi 1.8.3, stringr 1.5.1, survival 3.5-8, tibble 3.2.1, tidyselect 1.2.1, tools 4.4.0, utf8 1.2.4, vctrs 0.6.5, vsn 3.71.1, withr 3.0.0, xfun 0.43, xml2 1.3.6, xtable 1.8-4, zlibbioc 1.49.3