

globalSeq: testing for association between RNA-Seq and high-dimensional data

A Rauschenberger, MA Jonker, MA van de Wiel
and RX Menezes

May 15, 2016

This vignette explains how to use the R package **globalSeq**. The function **omnibus** tests for associations between an overdispersed count variable and a high-dimensional covariate set. The function **proprius** decomposes the test statistic to show the contributions of individual samples or covariates. And the function **cursus** performs genome-wide analyses.

1 Initialisation

Start with installing the R package **globalSeq** from Bioconductor:

```
source("https://bioconductor.org/biocLite.R")  
BiocInstaller::biocLite("globalSeq")
```

Please type the following command to load and attach the package:

```
library(globalSeq)
```

If you want to reproduce the examples, you should attach the toy database:

```
attach(toydata)
```

The following commands access the R documentation:

```
utils::help(globalSeq)  
utils::vignette("globalSeq")
```

2 Test of association

2.1 Data

Data is available for 10 individuals and 16 variables.

```
cbind(y,X)

##      y  X1 X2 X3 X4 X5 X6 X7 X8 X9 X10 X11 X12 X13 X14 X15
## ind1  7  0  1  0  0  0  0  1  1  1  0  1  0  1  1  0
## ind2  1  1  1  0 37 27 37 26 26 50  0  1  0  1  0  0
## ind3  1  0  1  0  0  0  1  1  1  0  0  0  0  0  0  1
## ind4 15 40 19 33 34 23 33 38 21 20 19 28  1  1  1 32
## ind5  6  1  1  0  0  0  1  1  1  1  1  1  1  1  0  1
## ind6 16 26 40 25 23 25 25 20 34 38 19 25  1  0  1 36
## ind7  5  0  0  1  1  0  0  1  0  0  0  0  1  1  1  1
## ind8  2  0  1  1  0  0  0  1  0  0  0  0  1  0  1  1
## ind9  2  1  1  1  1  0  0  0  1  1  0  0  1  1  1  1
## ind10 5  1  0  1  0  1  0  0  1  1  1  0  1  1  1  0
```

2.2 Minimal example

We are interested whether the response variable \mathbf{y} is associated with the covariate matrix \mathbf{X} . Looking at the data might already lead to an answer.¹

Because the number of covariates exceeds the sample size, classical tests cannot test their joint significance. But the function **omnibus** also works in high-dimensional settings:

```
set.seed(1)
omnibus(y,X)

##      pvalue teststat covs
## 1  0.025 2234.206    15
```

¹Note that the response variable takes higher values for individuals 4 and 6 than for the other individuals. Looking at the covariate matrix, we observe that individuals 2, 4 and 6 are peculiar. We conclude: The data on individuals 4 and 6 speak for an association, but the data on individual 2 speaks against an association. Covariates 12, 13 and 14 are uninformative, and the role of other covariates is less clear.

2.3 Offset

Suppose that an offset is available. Relative to the offset, the response \mathbf{y} is more or less constant across samples:

```
rbind(y,offset)

##          ind1 ind2 ind3 ind4 ind5 ind6 ind7 ind8 ind9 ind10
## y           7   1   1  15   6  16   5   2   2   5
## offset      15   1   2  28  10  28   9   2   5   7
```

If we account for this offset, there is no evidence for an association between the response \mathbf{y} and the covariate matrix \mathbf{X} :

```
set.seed(1)
omnibus(y,X,offset=offset)

##    pvalue  teststat covs
## 1  0.967 -10693.18   15
```

2.4 Confounding variable

Suppose that each sample belongs either to group 1 or to group 2. We can observe that \mathbf{y} tends to take small values in one group, and large values in the other:

```
rbind(y,group)

##          ind1 ind2 ind3 ind4 ind5 ind6 ind7 ind8 ind9 ind10
## y           7   1   1  15   6  16   5   2   2   5
## group        1   1   1   2   1   2   1   1   1   1
```

We suspect that the group membership explains some variation of the response \mathbf{y} or the covariate matrix \mathbf{X} . Therefore we account for this confounding variable by using stratified permutations:

```
set.seed(1)
omnibus(y,X,group=group)

##    pvalue  teststat covs
## 1 0.8714575 2234.206   15
```

2.5 Overdispersion

Setting the dispersion parameter of the negative binomial distribution equal to zero is equivalent to using the Poisson distribution:

```
set.seed(1)
omnibus(y,X,phi=0)

##      pvalue teststat covs
## 1  0.025 77991.33   15
```

2.6 Multiple covariate sets

Suppose that two covariate sets are available:

```
X1 <- X[,c(1:11,15)]
X2 <- X[,12:14]
```

We are interested in testing for associations between y on one hand, and $X1$ or $X2$ on the other:

```
set.seed(1)
omnibus(y,list(X1,X2))

##      joint teststat single.1 single.2 covs.1 covs.2
## 1 0.063 2.347201   0.025   0.561   12   3
```

The output includes the p -value and the test statistic for the joint test, the p -values for the individual tests, and the numbers of tested covariates.

2.7 P-values

When testing single covariate sets, the user can choose between three different types of p -values. By default p -values are calculated by permutation without repetitions ($kind = 1$). Alternatively, permutation can be interrupted when it becomes impossible to reach a predefined significance level ($0 < kind < 1$), or the method of control variables can be used ($kind = 0$).


```
omnibus(y,X,kind=1) # crude permutation test
omnibus(y,X,kind=0.05) # interrupting permutation
omnibus(y,X,kind=0) # method of control variables
```

3 Decomposition

3.1 Minimal example

Even though a single hypothesis is tested on the covariate set, the function **proprius** can obtain the contributions of individual covariates or samples to the test statistic.

```
proprius(y,X,type="samples")  
proprius(y,X,type="covariates")
```




We observe that individual 2 contributes negatively to the test statistic, whereas the contributions of individuals 4 and 6 are positive. We also observe that several covariates have large positive contributions. Summing over the individual contributions gives back the test statistic.

3.2 Null distribution

If a significance level α is specified, then the $1 - \alpha$ lower quantile under the null hypothesis is plotted:

```
proprius(y,X,type="covariates",alpha=0.05)
```



3.3 Further arguments

Offsets are included as in section 2.3, confounding variables are taken into account as in section 2.4, and overdispersion is treated as in section 2.5. The decompositions have not been implemented for multiple covariate sets.

4 Genome-wide analysis

4.1 Data

Suppose the matrix Y contains expression levels of 7 genes and 5 individuals, and the matrices V and W represent two different molecular profiles. Suppose that we furthermore know the locations of the genes ($Yloc$), the locations of the genetic or epigenetic alterations ($Vloc$ and $Wloc$), and the chromosome indicators $Ychr$, $Vchr$ and $Wchr$.

```
cbind(Yloc,Ychr,Y)
```

```
##      Yloc Ychr ind1 ind2 ind3 ind4 ind5
## gene1  19   1  37  61  21  24   8
## gene2  52   1  21  10  36   6  14
## gene3  53   1   6  23  19  14  13
## gene4  65   1  35  40  15  50  30
## gene5  80   1   8  24   6  33  18
## gene6  83   1  27  11  16  14  18
## gene7  84   1  54  55  16  32  15
```

```
cbind(Vloc,Vchr,V)
```

```
##      Vloc Vchr ind1 ind2 ind3 ind4 ind5
## V1      4   1  0.4  0.0  0.7 -1.3  0.3
## V2     14   1 -0.1 -0.2 -0.6  0.8  0.7
## V3     19   1 -0.4 -2.2 -0.7 -0.8  0.8
## V4     34   1  1.3  0.5  0.4 -0.4 -2.2
## V5     36   1 -0.7  1.6  1.3 -1.4 -0.9
## V6     51   1  0.6 -1.2  0.3  3.5 -0.2
## V7     55   1  0.3  1.7 -0.5 -1.2 -0.1
## V8     81   1  0.2  0.2  1.3  0.2 -0.5
## V9     83   1  0.9  1.4 -0.3  0.1  0.5
## V10    91   1 -0.9 -0.9  0.8 -0.9  0.0
```

```
cbind(Wloc,Wchr,W)
```

```
##      Wloc Wchr ind1 ind2 ind3 ind4 ind5
## W1     47   1   1   0   0   2   2
## W2     53   1   0   2   1   1   2
## W3     54   1   0   0   2   1   2
## W4     61   1   1   2   1   1   1
## W5     69   1   1   1   2   1   1
## W6     90   1   1   2   1   2   1
## W7     91   1   1   1   1   1   1
## W8     93   1   0   2   0   2   1
```

4.2 Minimal example

We are interested in testing the expression of each gene for associations with *local* genetic or epigenetic alterations.

Setting the argument `window` to 5 entails that each gene is tested for association with all covariates that are located no further than 5 units from the gene. In this example the covariates of interest are:

```
## gene1: V3
## gene2: V6 V7
## gene3: V6 V7
## gene4:
## gene5: V8 V9
## gene6: V8 V9
## gene7: V8 V9
```

The chromosome-wide analysis will not return any p -value for gene 4, because no covariates are in its vicinity.

```
set.seed(1)
cursus(Y,Yloc,V,Vloc,window=5)
```

4.3 Multiple chromosomes

If multiple chromosomes are analysed, it is important to include the chromosome indicators. They make sure that genes are mapped to covariates on the same chromosome:

```
set.seed(1)
cursus(Y,Yloc,V,Vloc,window=5,Ychr,Vchr)
```

4.4 Different library sizes

To account for different sequencing depths, an offset can be included:

```
offset <- colSums(Y) # library sizes
set.seed(1)
cursus(Y,Yloc,V,Vloc,window=5,offset=offset)
```

Otherwise the offset is calculated based on \mathbf{Y} .

4.5 Multiple molecular profiles

For the simultaneous analysis of multiple molecular profiles, the function **cursus** expects one covariate matrix, one location vector and one window size for each profile.

```
set.seed(1)
cursus(Y,Yloc,list(V,W),list(Vloc,Wloc),list(5,50))

## Analysing a single chromosome.
##           joint    teststat  single.1  single.2 covs.1 covs.2
## gene1 0.01666667  3.03041885 0.1333333 0.05000000      2      5
## gene2 0.28333333  1.02593007 0.5500000 0.08333333      2      8
## gene3 0.77500000 -1.00030785 0.7583333 0.62500000      2      8
## gene4          NA          NaN          NA 0.27500000      0      8
## gene5 0.54166667 -0.03586313 0.5250000 0.47500000      2      8
## gene6 0.58333333 -0.54128922 0.8416667 0.23333333      2      8
## gene7 0.22500000  0.82337451 0.5083333 0.17500000      2      8
```

4.6 Further arguments

Confounding variables are taken into account as in section 2.4. By setting `phi=rep(0,q)` where `q` is the number of genes, the user can restrict the negative binomial distribution to the Poisson distribution. By default, offsets and dispersion parameters are calculated internally, but they can also be inserted. The following example uses the R package **edgeR** from Bioconductor:

```
list <- edgeR::DGEList(counts=Y)
list <- edgeR::calcNormFactors(list)
list <- edgeR::estimateDisp(list)
offset <- list$samples$norm.factors
phi <- list$tagwise.dispersion

cursus(Y,Yloc,V,Vloc>window=5,offset=offset,phi=phi)
```

References

The R package **globalSeq** is based on Rauschenberger et al. [5], where detailed references to previous work are given. If you use **globalSeq** for publications, please cite Rauschenberger et al. [5].

Based on calculations from le Cessie and van Houwelingen [2] and using the generalised linear modelling framework from McCullagh et al. [3], Goeman et al. [1] showed how to test for association between a response variable from the exponential family of distributions and a high-dimensional covariate set, and derived the contributions of samples and covariates to the test statistic. Menezes et al. [4] extended this test to multiple covariate sets. Rauschenberger et al. [5] adapted the methods from Goeman et al. [1] and Menezes et al. [4] to the negative binomial distribution with an unknown dispersion parameter. This permutation test was made computationally efficient using ideas from Wieringen et al. [7] and Senchaudhuri et al. [6].

- [1] Jelle J Goeman, Sara A van de Geer, Floor de Kort, and Hans C van Houwelingen. A global test for groups of genes: testing association with a clinical outcome. *Bioinformatics*, 20:93–99, 2004.
- [2] Saskia le Cessie and Hans C van Houwelingen. Testing the fit of a regression model via score tests in random effects models. *Biometrics*, 51:600–614, 1995.
- [3] Peter McCullagh and John A Nelder. *Generalized linear models. 2nd ed.* Chapman and Hall, London, 1989.
- [4] Renée X Menezes, Leila Mohammadi, Jelle J Goeman, and Judith Boer. Analysing multiple types of molecular profiles simultaneously: connecting the needles in the haystack. *BMC Bioinformatics*, 17:77, 2015.
- [5] Armin Rauschenberger, Marianne A Jonker, Mark van de Wiel, and Renée X Menezes. Testing for association between RNA-Seq and high-dimensional data. *BMC Bioinformatics*, 17:118, 2016.
- [6] Pralay Senchaudhuri, Cyrus R Mehta, and Nitin R Patel. Estimating exact p values by the method of control variates or monte carlo rescue. *Journal of the American Statistical Association*, 90:640–648, 1995.
- [7] Wessel N van Wieringen, Mark A van de Wiel, and Aad W van der Vaart. A test for partial differential expression. *Journal of the American Statistical Association*, 103:1039–1049, 2008.