

Using *GGtools* for eQTL discovery and interpretation

VJ Carey `stvjc` at `channing.harvard.edu`

April 10, 2013

Contents

| | | |
|----------|---|-----------|
| 1 | Overview and installation | 1 |
| 2 | Data structures | 2 |
| 2.1 | Reference data supplied with Bioconductor | 2 |
| 2.2 | Working with your own data | 4 |
| 2.3 | Filters and permutation support | 5 |
| 2.3.1 | Bracket operations | 5 |
| 2.3.2 | Large scale filters | 5 |
| 2.3.3 | Permutation of expression against genotype | 6 |
| 2.4 | Post-analysis data structures | 6 |
| 3 | Focused analyses | 7 |
| 4 | Comprehensive surveys | 10 |
| 4.1 | A set of genes vs. all SNP on a chromosome | 10 |
| 4.2 | Tabulating best associated cis-rSNP with permutation-based FDR: small example | 11 |
| 4.3 | Tabulating all cis-rSNP with association score exceeding a given threshold | 13 |
| 4.4 | Removal of empirically identified expression heterogeneity | 14 |
| 5 | Exercises | 17 |
| 6 | Appendix: building your own structures | 18 |
| 7 | Session information | 20 |

1 Overview and installation

This document addresses data structure and analytic workflow for investigations of genetic sources of expression variation. Key background references are Williams et al.

(2007) for general biologic overview, Cheung et al. (2005) and Stranger et al. (2007) for key applications, and Stegle et al. (2010), Petretto et al. (2010), and Leek et al. (2010) for various methodological issues. Majewski and Pastinen (2011) reviews potentials of eQTL investigations with expression measures based on RNA sequencing.

This document is constructed using R version 3.0.0. See the session information at the end of the document for full details. Using a comparable version of R, you can obtain the software needed for the production of this document using

```
source("http://www.bioconductor.org/biocLite.R")
biocLite("GGtools", dependencies=TRUE)
```

2 Data structures

2.1 Reference data supplied with Bioconductor

A collection of 30 trios of central European ancestry was genotyped for 4 million SNP loci in HapMap phase II. Immortalized B-cell lines were assayed for gene expression using Illumina's HumanWG6v1 bead array. Digital data on expression and genotype for the 90 CEU individuals is distributed in Bioconductor package *GGdata*; the expression data were retrieved from the GENEVAR website of Wellcome Trust, e.g.,

```
ftp://ftp.sanger.ac.uk/pub/genevar/CEU_parents_norm_march2007.zip
```

and the genotype data were obtained directly from hapmap.org at build 36:

```
ftp://ftp.ncbi.nlm.nih.gov/hapmap/genotypes/2008-03/forward/non-redundant/
```

The data in *GGdata* are likely derived from r23, while r23a is now distributed. Some effort at updating genotypes may be supplied in the future.

Acquire the genome-wide expression data and the genotype data for chromosome 20 as follows:

```
> suppressPackageStartupMessages(library(GGtools))
> library(parallel)

> g20 = getSS("GGdata", "20")

> g20
```

```
SnpMatrix-based genotype set:
number of samples: 90
number of chromosomes present: 1
annotation: illuminaHumanv1.db
Expression data dims: 47293 x 90
Total number of SNP: 119921
```

```
Phenodata: An object of class 'AnnotatedDataFrame'
  sampleNames: NA06985 NA06991 ... NA12892 (90 total)
  varLabels: famid persid ... male (7 total)
  varMetadata: labelDescription
```

```
> class(g20)
```

```
[1] "smlSet"
attr(,"package")
[1] "GGBase"
```

The `smlSet` class was designed in 2006 as an experiment in unifying high-throughput expression and genotype data. A key resource was the *snpMatrix* (now *snpStats*) package of David Clayton, which defined an 8-bit representation of genotype calls (including uncertain calls obtained by statistical imputation), import and coercion infrastructure allowing use of the 8-bit representation with popular genetic data formats (pedfiles, mach and beagle outputs, etc.), and statistical testing infrastructure employing this representation.

The expression and sample-level data are handled just as with familiar `ExpressionSet` instances:

```
> exprs(g20)[1:5,1:5]
```

| | NA06985 | NA06991 | NA06993 | NA06994 | NA07000 |
|---------------|-----------|-----------|-----------|-----------|-----------|
| GI_10047089-S | 5.983962 | 5.939529 | 5.912270 | 5.891347 | 5.906675 |
| GI_10047091-S | 6.544493 | 6.286516 | 6.244446 | 6.277397 | 6.330893 |
| GI_10047093-S | 9.905235 | 10.353804 | 10.380972 | 9.889223 | 10.155686 |
| GI_10047099-S | 7.993935 | 7.593970 | 8.261215 | 6.598430 | 6.728085 |
| GI_10047103-S | 11.882265 | 12.204753 | 12.249708 | 11.798415 | 12.015252 |

```
> pData(g20)[1:4,]
```

| | famid | persid | mothid | fathid | sampid | isFounder | male |
|---------|-------|--------|--------|--------|---------|-----------|-------|
| NA06985 | 1341 | 14 | 0 | 0 | NA06985 | TRUE | FALSE |
| NA06991 | 1341 | 2 | 14 | 13 | NA06991 | FALSE | FALSE |
| NA06993 | 1341 | 13 | 0 | 0 | NA06993 | TRUE | TRUE |
| NA06994 | 1340 | 9 | 0 | 0 | NA06994 | TRUE | TRUE |

The genotype data are held in a list with elements intended to represent chromosomes, and the list is stored in an environment to reduce copying efforts.

```
> smList(g20)
```

```
$`20`
A SnpMatrix with 90 rows and 119921 columns
Row names: NA06985 ... NA12892
Col names: rs4814683 ... rs6090120
```

```
> as(smList(g20)[[1]][1:5,1:5], "matrix")
```

| | rs4814683 | rs6076506 | rs6139074 | rs1418258 | rs7274499 |
|---------|-----------|-----------|-----------|-----------|-----------|
| NA06985 | 03 | 03 | 03 | 03 | 03 |
| NA06991 | 02 | 03 | 02 | 02 | 03 |
| NA06993 | 01 | 03 | 01 | 01 | 03 |
| NA06994 | 01 | 03 | 01 | 01 | 03 |
| NA07000 | 03 | 03 | 03 | 03 | 03 |

The leading zeroes in the display above indicate that raw bytes are used to represent the genotypes per sample (rows) per SNP (columns). Coercions:

```
> as(smList(g20)[[1]][1:5,1:5], "numeric")
```

| | rs4814683 | rs6076506 | rs6139074 | rs1418258 | rs7274499 |
|---------|-----------|-----------|-----------|-----------|-----------|
| NA06985 | 2 | 2 | 2 | 2 | 2 |
| NA06991 | 1 | 2 | 1 | 1 | 2 |
| NA06993 | 0 | 2 | 0 | 0 | 2 |
| NA06994 | 0 | 2 | 0 | 0 | 2 |
| NA07000 | 2 | 2 | 2 | 2 | 2 |

```
> as(smList(g20)[[1]][1:5,1:5], "character")
```

| | rs4814683 | rs6076506 | rs6139074 | rs1418258 | rs7274499 |
|---------|-----------|-----------|-----------|-----------|-----------|
| NA06985 | "B/B" | "B/B" | "B/B" | "B/B" | "B/B" |
| NA06991 | "A/B" | "B/B" | "A/B" | "A/B" | "B/B" |
| NA06993 | "A/A" | "B/B" | "A/A" | "A/A" | "B/B" |
| NA06994 | "A/A" | "B/B" | "A/A" | "A/A" | "B/B" |
| NA07000 | "B/B" | "B/B" | "B/B" | "B/B" | "B/B" |

Any number of chromosomes can be held in the genotype list component, but the design allows working with only one chromosome at a time and examples emphasize this approach. Amalgamation of results across chromosomes is generally straightforward.

2.2 Working with your own data

The `make_smlSet` function can be used to bind a list of suitably named `SnpMatrix` instances with an `ExpressionSet` instance to create a single `smlSet` instance covering multiple chromosomes.

The **externalize** function can be applied to such an **smlSet** instance to create a new *package* which can be installed for use with **getSS**. This is the preferred way of managing work with large genotyping panels (such as the 10 million locus panel achievable with “thousand genomes imputation”).

Briefly, **externalize** arranges a DESCRIPTION file and system of folders that can be installed as an R package. The expression data are stored as object **ex** in `data/eset.rda`, and the **SnpMatrix** instances are stored separately as `.rda` files in `inst/parts`. The **getSS** function will create **smlSet** instances on the fly from the externalize-generated package.

2.3 Filters and permutation support

Here “filter” is used to refer to any function that processes an **smlSet** instance and returns an **smlSet** instance after altering the contents, which may involve eliminating probes or SNPs, performing numerical transformations to expression or genotype measures, transforming sample data or eliminating samples.

2.3.1 Bracket operations

Coordinated manipulations of genotype, expression, and phenotype information on samples can be accomplished with the second subscript to the bracket operator. Thus `g20[,1:5]` is the restriction of **g20** to the first five samples. Sample names may also be used for such manipulations.

Reduction of the expression component can be accomplished with the first subscript to the bracket operator. Thus `g20[1:5,]` is the restriction of **g20** to five expression probes; all other data are unaltered. If it is desired to use feature names for such manipulations, the character vector of feature names must be cast to class **probeId** with the **probeId()** method. At present no such operations can be used to alter the genotype data contents.

2.3.2 Large scale filters

It is known that non-specific filtering (removal of probes with low variation across samples, without regard to sample phenotype or class information) can increase sensitivity and specificity of certain differential expression test procedures (Bourgon et al., 2010). The **nsFilter** function of the *genefilter* package has been adapted to work with **smlSet** instances.

SNPs can be filtered out of the **smlSet** instance on the basis of observed minor allele or genotype frequencies using **MAFfilter** and **GTFfilter** respectively.

Various approaches to reduction of “expression heterogeneity” can be examined. `clipPCs(x, vec)` will form the singular value decomposition of the expression matrix and remove principal components enumerated in **vec** by reassembling the expression matrix after setting eigenvalues in **vec** to zero. It is also possible to employ any computed

quantities such as principal components or surrogate variables identified in SVA (Leek and Storey, 2007) as covariates in the formula element of analysis functions in **GGtools**, but note that simple permutations do not lead to valid permutation tests of SNP effects in the presence of covariates (see (Bůžková et al., 2011) who focus on interaction, but describe the problem for main effects models, with references, early in the paper.)

The introduction of novel approaches to expression transformation can be accomplished using code similar to the following, illustrating of the use of PEER (Stegle et al., 2010):

```
> library(peer)
> model = PEER()
> PEER_setPhenoMean(model, t(exprs(g20)))
> PEER_setNk(model, 10)
> PEER_setCovariates(model, matrix(1*g20$male,nc=1))
> PEER_update(model)
> resid=t(PEER_getResiduals(model))
> rownames(resid) = featureNames(g20)
> colnames(resid) = sampleNames(g20)
> g20peer10 = g20
> g20peer10@assayData = assayDataNew("lockedEnvironment", exprs=resid)
```

At this point, `g20peer10` holds expression data with 10 latent factors removed after adjustment for gender.

2.3.3 Permutation of expression against genotype

Because the *snpStats* testing procedures defensively match predictor to response variable orderings using sample labels, special steps must be taken to ensure that tests use properly permuted responses. The `permEx` function takes care of this, using the current state of the random number generator.

2.4 Post-analysis data structures

While it is possible to construe the results of an eQTL search as a static report, it is more productive to conceptualize the result as a data object for further analysis. Unfortunately, the number of tests to be managed can be very large – at least hundreds of millions, and these must be joinable with location metadata to be maximally useful.

Several data structures for managing post-analysis results have emerged as this package has matured. Of particular concern are those that use *ff* out-of-memory archiving for test statistic values or effect estimates and those that use the **GRanges** infrastructure to facilitate efficient query resolution in genomic coordinates. These will be described along with the related analytic workflow steps.

3 Focused analyses

A specific gene can be checked for eQTL on a given chromosome or set of chromosomes with `gwSnpTests`. There are various convenience facilities. In the call to `gwSnpTests` below, a gene symbol is used to pick out an expression element, and adjustment for gender is commodated in an additive genetic model for effects of B allele copy number on expression of CPNE1. One degree of freedom chi-squared tests are computed.

```
> t1 = gwSnpTests(genesym("CPNE1")~male, g20, chrnum("20"))
> t1
```

```
gwSnpScreenResult for gene  CPNE1  [probe  GI_23397697-A ]
```

```
> topSnps(t1)
```


```
$`20`
```

| | p.val |
|------------|--------------|
| rs17093026 | 6.485612e-14 |
| rs1118233 | 1.897898e-13 |
| rs2425078 | 2.426168e-13 |
| rs1970357 | 2.426168e-13 |
| rs12480408 | 2.426168e-13 |
| rs6060535 | 2.426168e-13 |
| rs11696527 | 2.426168e-13 |
| rs6058303 | 2.426168e-13 |
| rs6060578 | 2.426168e-13 |
| rs7273815 | 2.544058e-13 |

It is possible to compute tests for this specific gene for association with SNP across several chromosomes if desired; change the value of the third argument to a vector.

There are a few approaches to visualization of the results that are relevant, but complications arise in relation to choice of genomic coordinates.

```
> plot(t1, snplocsDefault())
> plot_EvG(genesym("CPNE1"), rsid("rs17093026"), g20)
```



Code like the following can be used to display scores ($-\log_{10} p$) on the genome browser, here with hg19 locations.

```
> library(snplocsDefault(), character.only=TRUE)
> S20 = getSNPlocs("ch20", as.GRanges=TRUE)
> GR20 = makeGRanges(t1, S20)
> library(rtracklayer)
> export(GR20, "~/cpne1new.wig")
```

With this code, it will be necessary to manually alter the chr assignment in the wig file, and place an informative title to get the following display.

4 Comprehensive surveys

4.1 A set of genes vs. all SNP on a chromosome

The performance of *snpStats* `snp.rhs.tests` is very good and so our principle for large-scale searches is to compute all association statistics, save them in truncated form, and filter results later. This is carried out with the `eqtlTests` function. To illustrate, the expression data is sharply filtered to the 50 most variable genes on chromosome 20 as measured by cross-sample median absolute deviation, SNP with $MAF < 0.05$ are removed, and then all SNP-gene association tests are executed.

```
> g20 = GGtools:::restrictProbesToChrom(g20, "20")
> mads = apply(exprs(g20),1,mad)
> oo = order(mads, decreasing=TRUE)
> g20 = g20[oo[1:50],]
> tf = tempfile()
> dir.create(tf)
> e1 = eqtlTests(MAFfilter(g20, lower=0.05), ~male,
+   geneApply=mclapply, targdir=tf)
> e1
```

```
eqtlTestsManager  computed Wed Apr 10 20:11:07 2013
gene annotation: illuminaHumanv1.db
some genes (out of 50): hmm9615-S GI_31077201-S ... GI_34335276-S GI_32483384-A
some snps (out of 55039): rs4814683 rs6139074 ... rs6062363 rs4809418
```

On a two-core macbook pro, this computation takes well under a minute. The details of the underlying data structure are involved. Briefly, a short integer is used to represent each chi-squared statistic obtained in the 1 tests computed, in an `ff` archive. Use `topFeats` to manually harvest this.

```
> pm1 = colnames(e1@fffile)
> tops = sapply(pm1, function(x) topFeats(probeId(x), mgr=e1, n=1))
> top6 = sort(tops, decreasing=TRUE)[1:6]
```

```
> print(top6)
```

| | | |
|-------------------------|-------------------------|--------------------------|
| GI_17149835-I.rs6041750 | GI_34147330-S.rs6037097 | GI_41327717-S.rs2223883 |
| 40.30 | 36.49 | 24.87 |
| GI_31563517-A.rs6085541 | GI_29826316-A.rs6135980 | GI_13540544-S.rs11905405 |
| 21.38 | 19.70 | 19.20 |

R has propagated the names of probes and SNPs with the scores so that a table can be created as follows:

```

> nms = strsplit(names(top6), "\\.")
> gn = sapply(nms, "[", 1)
> sn = sapply(nms, "[", 2)
> tab = data.frame(snp=sn, score=as.numeric(top6))
> rownames(tab) = gn
> tab

```

| | snp | score |
|---------------|------------|-------|
| GI_17149835-I | rs6041750 | 40.30 |
| GI_34147330-S | rs6037097 | 36.49 |
| GI_41327717-S | rs2223883 | 24.87 |
| GI_31563517-A | rs6085541 | 21.38 |
| GI_29826316-A | rs6135980 | 19.70 |
| GI_13540544-S | rs11905405 | 19.20 |

Statistical interpretation of the scores in this table is not clear as the data structure includes familial aggregation in trios and extended pedigrees, and may include population stratification. Nevertheless, consistency of these findings with other published results involving multiple populations can be checked. *GGtools* includes a table published by Stranger and colleagues in 2007 enumerating multipopulation eQTL (Stranger et al., 2007).

```

> data(strMultiPop)
> strMultiPop[ strMultiPop$rsid %in% tab$snp, ]

```

| | rsid | genesym | illv1pid | snpChr | snpCoordB35 | probeMidCoorB35 |
|-------|-----------|----------|---------------|--------|-------------|-----------------|
| 11306 | rs6041750 | FKBP1A | GI_17149835-I | 20 | 1298709 | 1300503 |
| 11398 | rs6037097 | PYGB | GI_34577123-S | 20 | 25295221 | 25226352 |
| 25370 | rs6037097 | C20orf22 | GI_34147330-S | 20 | 25295221 | 25223609 |

| | snp2probe | minuslog10p | adjR2 | assocGrad | permThresh | popSet |
|-------|-----------|-------------|--------|-----------|------------|-----------------|
| 11306 | 1794 | 21.6949 | 0.3798 | 0.4708 | 12.4815 | CEU-CHB-JPT-YRI |
| 11398 | 68869 | 8.9349 | 0.1645 | 0.1324 | 6.5410 | CEU-CHB-JPT-YRI |
| 25370 | 71612 | 8.9684 | 0.2257 | 0.3396 | 6.5535 | CEU-CHB-JPT |

Thus the top two SNP in the table computed here are identified as multipopulation eQTL by Stranger. The other association scores are not very strong and likely do not correspond to genuine associations.

4.2 Tabulating best associated cis-rSNP with permutation-based FDR: small example

The workhorse for identifying genes to which can be associated putatively regulating SNP (rSNP) is `best.cis.eQTLs`. This can be used for genome-wide analysis, but here

an alternative table is created for the sharply filtered chromosome 20 data given above. This call says that gene location information will be acquired from the Bioconductor TxDb.Hsapiens annotation package for hg19 UCSC known genes, and that tests for association within 1 Mbase of the coding region for each gene will be considered. The expression data will be permuted against genotype data in two independent draws to assemble the null reference distribution for association scores; these are used to enumerate false significance claims at various magnitudes of the distribution of association scores. The plug-in procedure for estimating FDR XXX cite Hastie Tibs Friedman is used.

```
> if (file.exists("db2")) unlink("db2", recursive=TRUE)
> fn = probeId(featureNames(g20))
> exTx = function(x) MAFfilter( x[fn, ], lower=0.05)
> b1 = best.cis.eQTLs("GGdata", ~male, radius=1e6,
+   folderstem="db2", nperm=2, geneApply=mclapply,
+   smFilter= exTx, chrnames="20", snpannopk=snplocsDefault())
```

```
get data...build map...run smFilter...filter probes in map...tests...filter...done.
get data...build map...run smFilter...filter probes in map...tests...filter...done.
get data...build map...run smFilter...filter probes in map...tests...filter...done.
```

```
> b1
```

GGtools mcwBestCis instance. The call was:

```
best.cis.eQTLs(smpack = "GGdata", rhs = ~male, folderstem = "db2",
  radius = 1e+06, chrnames = "20", geneApply = mclapply, snpannopk = snplocsDefault(),
  smFilter = exTx, nperm = 2)
```

Best loci for 50 probes are recorded.

There were 83969 gene:snp tests.

Top 4 probe:SNP combinations:

GRanges with 4 ranges and 6 metadata columns:

| | seqnames | ranges | strand | score | snpid |
|---------------|----------|----------------------|--------|-----------|-------------|
| | <Rle> | <IRanges> | <Rle> | <numeric> | <character> |
| GI_34147330-S | 20 | [24280834, 26371618] | * | 36.49 | rs6037097 |
| GI_31077201-S | 20 | [61679079, 63680979] | * | 14.78 | rs13044229 |
| GI_38327656-I | 20 | [45286150, 47415360] | * | 9.80 | rs6094800 |
| GI_33598945-A | 20 | [38766161, 40804357] | * | 9.85 | rs16988882 |

| | snploc | radiusUsed | nsnp | fdr |
|---------------|-----------|------------|-----------|-----------|
| | <integer> | <numeric> | <integer> | <numeric> |
| GI_34147330-S | 25347221 | 1e+06 | 1387 | 0.0000000 |
| GI_31077201-S | 61738288 | 1e+06 | 766 | 0.7500000 |
| GI_38327656-I | 46369894 | 1e+06 | 2157 | 0.7666667 |
| GI_33598945-A | 38879237 | 1e+06 | 1658 | 0.7931034 |

```

seqlengths:
  20
  NA
====
use chromsUsed(), fullreport(), etc. for additional information.

```

4.3 Tabulating all cis-rSNP with association score exceeding a given threshold

The gene-centered analysis described in the previous subsection yields a set of thresholds corresponding to various FDRs. It may be of interest to enumerate all SNP with association scores exceeding some threshold. All.cis.eQTLs can be used for this.

```

> args(GGtools:::All.cis.eQTLs)

function (maxfdr = 0.05, inbestcis = NULL, smpack = "GGdata",
  rhs = ~1, folderstem = "cisScratch", radius = 50000, shortfac = 100,
  chrnames = as.character(1:22), smchrpref = "", gchrpref = "",
  schrpref = "ch", geneApply = lapply, geneannopk = "illuminaHumanv1.db",
  snpannopk = snplocsDefault(), smFilter4cis = function(x) nsFilter(MAFfilter(clipPCs
    1:10), lower = 0.05), var.cutoff = 0.85), smFilter4all = function(x) MAFfilter(
    1:10), lower = 0.05), nperm = 2, excludeRadius = NULL,
  exFilter = function(x) x, SSgen = GGBase::getSS)
NULL

```

The computation can be done *de novo* on the basis of an `smpack` argument, or can be regarded as a followup on a completed `best.cis.eQTLs` call.

```

> b1.all = GGtools:::All.cis.eQTLs( maxfdr = 0.05, inbestcis = b1, smpack="GGdata", r
+ chrnames = "20",
+ smFilter4all = function(x) MAFfilter(x, lower=0.05))

```

PHASE 1: determining cis threshold...

PHASE 2: extracting associations passing cis threshold...

build map...SSgen/filter the probes ...test...done.

Here a second SNP is identified for the one gene exhibiting an eQTL at $FDR \leq 0.05$.

Nota bene. Many users have attempted to work with All.cis.eQTLs with `maxfdr` set to 1 in order to obtain all cis tests. This is not the intended use. eQTLtests or a new function called All.cis can be used for exhaustive testing tasks.

4.4 Removal of empirically identified expression heterogeneity

Leek et al. (2010) describe implications of batch effects in high-throughput experimental contexts. Various methods for adjustment of responses have been proposed; a very simple but evidently risky approach is nonspecific removal of principal components of variation. To maximize information on possible batch effects, the entire expression matrix is restored and decomposed for principal component removal.

```
> if (file.exists("db2")) unlink("db2", recursive=TRUE)
> g20 = getSS("GGdata", "20")
> exTx = function(x) MAFfilter( clipPCs(x,1:10)[fn, ], lower=0.05)
> g20f = exTx(g20)

> b2 = best.cis.eQTLs("GGdata", ~male, radius=1e6,
+   folderstem="db2", nperm=2, geneApply=mclapply,
+   smFilter= exTx, chrnames="20", snpannopk=snplocsDefault())

get data...build map...run smFilter...filter probes in map...tests...filter...done.
get data...build map...run smFilter...filter probes in map...tests...filter...done.
get data...build map...run smFilter...filter probes in map...tests...filter...done.

> b2
```

GGtools mcwBestCis instance. The call was:

```
best.cis.eQTLs(smpack = "GGdata", rhs = ~male, folderstem = "db2",
  radius = 1e+06, chrnames = "20", geneApply = mclapply, snpannopk = snplocsDefault(),
  smFilter = exTx, nperm = 2)
```

Best loci for 50 probes are recorded.
 There were 83969 gene:snp tests.
 Top 4 probe:SNP combinations:
 GRanges with 4 ranges and 6 metadata columns:

| | seqnames | ranges | strand | score | snpid |
|---------------|----------|----------------------|--------|-----------|-------------|
| | <Rle> | <IRanges> | <Rle> | <numeric> | <character> |
| GI_34147330-S | 20 | [24280834, 26371618] | * | 54.90 | rs2261747 |
| GI_29029560-A | 20 | [46662783, 48713497] | * | 30.96 | rs4810909 |
| GI_17149835-I | 20 | [352356, 2373816] | * | 26.03 | rs6041749 |
| GI_31652260-S | 20 | [41295709, 43345122] | * | 25.90 | rs619289 |

| | snploc | radiusUsed | nsnp | fdr |
|---------------|-----------|------------|-----------|-----------|
| | <integer> | <numeric> | <integer> | <numeric> |
| GI_34147330-S | 25261784 | 1e+06 | 1387 | 0 |
| GI_29029560-A | 47621372 | 1e+06 | 1952 | 0 |
| GI_17149835-I | 1350672 | 1e+06 | 2382 | 0 |
| GI_31652260-S | 42294613 | 1e+06 | 1664 | 0 |

```

seqlengths:
  20
  NA
====
use chromsUsed(), fullreport(), etc. for additional information.

```

An improvement in sensitivity is seen after this adjustment. The probes with FDR at 0.001 or lower are identified by the helper function

```
> goodProbes = function(x) names(x@scoregr[elementMetadata(x@scoregr)$fdr<0.001])
```

All probes identified as significant (at $FDR \leq 0.001$) before the PCA adjustment are identified as such after it:





```
> setdiff(goodProbes(b2), goodProbes(b1))

[1] "GI_29029560-A" "GI_17149835-I" "GI_31652260-S" "GI_4502372-S"
[5] "GI_21264601-S"
```

The effects of the adjustment for genes that were significant only in the adjusted analysis can be visualized:

```
> newp = setdiff(goodProbes(b2), goodProbes(b1))
> np = length(newp)
> bestSnp = function(pn, esm) elementMetadata(esm@scoregr[pn])$snpid
> par(mfrow=c(2,2))
> plot_EvG(probeId(newp[1]), rsid(bestSnp(newp[1], b2)), g20, main="raw")
> plot_EvG(probeId(newp[1]), rsid(bestSnp(newp[1], b2)), g20f, main="PC-adjusted")
> plot_EvG(probeId(newp[np]), rsid(bestSnp(newp[np], b2)), g20, main="raw")
> plot_EvG(probeId(newp[np]), rsid(bestSnp(newp[np], b2)), g20f, main="PC-adjusted")

```



5 Exercises

1. All computations performed above ignore familial structure in the data that can be determined using the `famid`, `mothid`, `fathid` variables in the `pData(g20)`. Reduce the `smlSet` instance used for eQTL testing to parents only, who have parent identifiers equal to zero, and recompute the main tables.
2. For selected eQTLs that are were significant with low FDR in the full data ignoring, but are not significant in the analysis of the reduced data, use a reasonably specified variance components model on the full data with familial structure and compute a third test statistic. Is the restriction to parents only a good policy for eQTL discovery? Is there evidence of substantial familial aggregation in expression after heterogeneity reduction?
3. How can we select in a principled way the number of principal components to be removed for heterogeneity reduction?

6 Appendix: building your own structures

The following code illustrates construction of a minimal `smlSet` instance.

```
> library(Biobase)
> suppressPackageStartupMessages(library(GGtools))
> ex = matrix(0, nr=5, nc=3)
> pd = data.frame(v1 = 1:3, v2=5:7)
> colnames(ex) = rownames(pd) = LETTERS[1:3]
> adf = AnnotatedDataFrame(pd)
> rownames(ex) = letters[1:5]
> es = ExpressionSet(ex, phenoData=adf)
> exprs(es)

  A B C
a 0 0 0
b 0 0 0
c 0 0 0
d 0 0 0
e 0 0 0

> pData(es)

  v1 v2
A  1  5
B  2  6
C  3  7

> library(snpStats)
> mysnp = matrix(rep(1:3, 10), nr=3) # note 1=A/A, ... 0 = NA
> rownames(mysnp) = colnames(ex)
> mysm = new("SnpMatrix", mysnp)

coercing object of mode numeric to SnpMatrix

> as(mysm, "character")

[,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10]
A "A/A" "A/A" "A/A" "A/A" "A/A" "A/A" "A/A" "A/A" "A/A" "A/A"
B "A/B" "A/B" "A/B" "A/B" "A/B" "A/B" "A/B" "A/B" "A/B" "A/B"
C "B/B" "B/B" "B/B" "B/B" "B/B" "B/B" "B/B" "B/B" "B/B" "B/B"

> as(mysm, "numeric")
```

| | [,1] | [,2] | [,3] | [,4] | [,5] | [,6] | [,7] | [,8] | [,9] | [,10] |
|---|------|------|------|------|------|------|------|------|------|-------|
| A | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| B | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| C | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 |

```
> sml = make_smlSet(es, list(c1=mysm))
> annotation(sml)
```

```
character(0)
```

```
> colnames(smList(sml)[[1]])
```

```
NULL
```

7 Session information

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

- R version 3.0.0 (2013-04-03), x86_64-apple-darwin10.8.0
- Locale: C
- Base packages: base, datasets, grDevices, graphics, methods, parallel, splines, stats, stats4, utils
- Other packages: AnnotationDbi 1.22.1, BSgenome 1.28.0, Biobase 2.20.0, BiocGenerics 0.6.0, Biostrings 2.28.0, DBI 0.2-5, GGBase 3.22.0, GGtools 4.8.0, GenomicRanges 1.12.1, IRanges 1.18.0, Matrix 1.0-12, RSQLite 0.11.2, Rsamtools 1.12.0, SNPlocs.Hsapiens.dbSNP.20120608 0.99.9, illuminaHumanv1.db 1.18.0, lattice 0.20-15, org.Hs.eg.db 2.9.0, snpStats 1.10.0, survival 2.37-4
- Loaded via a namespace (and not attached): AnnotationForge 1.2.0, GenomicFeatures 1.12.0, KernSmooth 2.23-10, RCurl 1.95-4.1, VariantAnnotation 1.6.1, XML 3.96-1.1, annotate 1.38.0, biomaRt 2.16.0, bit 1.1-10, bitops 1.0-5, ff 2.2-11, genefilter 1.42.0, grid 3.0.0, rtracklayer 1.20.0, tools 3.0.0, xtable 1.7-1, zlibbioc 1.6.0

References

- Richard Bourgon, Robert Gentleman, and Wolfgang Huber. Independent filtering increases detection power for high-throughput experiments. *Proc Natl Acad Sci USA*, 107(21):9546–51, May 2010. doi: 10.1073/pnas.0914005107.
- Petra Bůžková, Thomas Lumley, and Kenneth Rice. Permutation and parametric bootstrap tests for gene-gene and gene-environment interactions. *Ann Hum Genet*, 75(1): 36–45, Jan 2011. doi: 10.1111/j.1469-1809.2010.00572.x.
- Vivian G Cheung, Richard S Spielman, Kathryn G Ewens, Teresa M Weber, Michael Morley, and Joshua T Burdick. Mapping determinants of human gene expression by regional and genome-wide association. *Nature*, 437(7063):1365–9, Oct 2005. doi: 10.1038/nature04244. URL <http://192.168.100.1:5280/?redirect=http%253A/eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi>.
- Jeffrey T Leek and John D Storey. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genet*, 3(9):1724–35, Sep 2007. doi: 10.1371/journal.pgen.0030161. URL <http://www.plosgenetics.org/article/info%253Adoi%252F10.1371%252Fjournal.pgen.0030161>.

- Jeffrey T Leek, Robert B Scharpf, Héctor Corrada Bravo, David Simcha, Benjamin Langmead, W Evan Johnson, Donald Geman, Keith Baggerly, and Rafael A Irizarry. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat Rev Genet*, 11(10):733–9, Oct 2010. doi: 10.1038/nrg2825. URL <http://www.nature.com/nrg/journal/v11/n10/full/nrg2825.html>.
- Jacek Majewski and Tomi Pastinen. The study of eqtl variations by rna-seq: from snps to phenotypes. *Trends Genet*, 27(2):72–9, Feb 2011. doi: 10.1016/j.tig.2010.10.006.
- Enrico Petretto, Leonardo Bottolo, Sarah R Langley, Matthias Heinig, Chris McDermott-Roe, Rizwan Sarwar, Michal Pravenec, Norbert Hübner, Timothy J Aitman, Stuart A Cook, and Sylvia Richardson. New insights into the genetic control of gene expression using a bayesian multi-tissue approach. *PLoS Comput Biol*, 6(4):e1000737, Apr 2010. doi: 10.1371/journal.pcbi.1000737. URL <http://www.ploscompbiol.org/article/info%253Adoi%252F10.1371%252Fjournal.pcbi.1000737>.
- Oliver Stegle, Leopold Parts, Richard Durbin, and John Winn. A bayesian framework to account for complex non-genetic factors in gene expression levels greatly increases power in eqtl studies. *PLoS Comput Biol*, 6(5):e1000770, May 2010. doi: 10.1371/journal.pcbi.1000770. URL <http://www.ploscompbiol.org/article/info%253Adoi%252F10.1371%252Fjournal.pcbi.1000770>.
- Barbara E Stranger, Alexandra C Nica, Matthew S Forrest, Antigone Dimas, Christine P Bird, Claude Beazley, Catherine E Ingle, Mark Dunning, Paul Flicek, Daphne Koller, Stephen Montgomery, Simon Tavaré, Panos Deloukas, and Emmanouil T Dermitzakis. Population genomics of human gene expression. *Nat Genet*, 39(10):1217–1224, Oct 2007. doi: 10.1038/ng2142.
- R. B.H Williams, E. K.F Chan, M. J Cowley, and P. F.R Little. The influence of genetic variation on gene expression. *Genome Research*, 17(12):1707–1716, Dec 2007. doi: 10.1101/gr.6981507.