

Using Categories defined by Chromosome Bands

D. Sarkar, S. Falcon, R. Gentleman

August 7, 2008

1 Chromosome bands

Typically, in higher organisms, each chromosome has a centromere and two arms. The short arm is called the p arm and the longer arm the q arm. Chromosome bands (see Figure 1) are identified by differential staining, usually with Giemsa-based stains, and many disease-related defects have been mapped to these bands; such mappings have played an important role in classical cytogenetics. With the availability of complete sequences for several genomes, there have been efforts to link these bands with specific sequence locations (Furey and Haussler, 2003). The estimated location of the bands in the reference genomes can be obtained from the UCSC genome browser, and data linking genes to particular bands can be obtained from a variety of sources such as the NCBI. This vignette demonstrates tools that allow the use of categories derived from chromosome bands, that is, the relevant categories are determined *a priori* by a mapping of genes to chromosome bands.

Figure 1 shows an ideogram of human chromosome 12, with the band 12q21 shaded. As shown in the figure, 12q21 can be divided into more granular levels 12q21.1, 12q21.2, and 12q21.3. 12q21.3 can itself be divided at an even finer level of resolution into 12q21.31, 12q21.32, and 12q21.33. Moving towards less granular bands, 12q21 is a part of 12q2 which is again a part of 12q. We take advantage of this nested structure of the bands in our analysis.

```
> library("Category")
> library("ALL")
> library("hgu95av2.db")
> library("annotate")
> library("genefilter")
> library("SNPchip")
> library("geneplotter")
> library("limma")
> library("lattice")
> library("graph")
```

2 Data Preparation

For illustration, we use a microarray dataset (Chiaretti et al., 2005) from a clinical trial in acute lymphoblastic leukemia (ALL). The data are described in Chapter 2 of Hahne et al.

NULL

NULL




Figure 1: Ideogram for human chromosome 12. The p arm is on the left, the q arm is on the right, and the centromere is indicated by a notch. The shaded bands together represent 12q21. This band is composed of three sub-bands: 12q21.1, 12q21.2, and 12q21.3. The last of these is composed of sub-sub-bands 12q21.31, 12q21.32, and 12q21.33.

(2008). For the analysis presented here, we consider the comparison of two subsets: patients identified as having a BCR/ABL gene fusion present, typically as a result of a translocation of chromosomes 9 and 22 (labeled BCR/ABL), and those that have no observed cytogenetic abnormalities (labeled NEG). The full dataset is available in the ALL package, and the relevant subset of the data can be obtained by

```
> data(ALL, package="ALL")
> subsetType <- "BCR/ABL"
> Bcell <- grep("^B", as.character(ALL$BT))
> bcrAblOrNegIdx <- which(as.character(ALL$mol.biol) %in% c("NEG", subsetType))
> bcrAblOrNeg <- ALL[, intersect(Bcell, bcrAblOrNegIdx)]
> bcrAblOrNeg$mol.biol <- factor(bcrAblOrNeg$mol.biol)
```

We also create relevant annotation maps to go from feature names to Entrez ID, gene symbol, and chromosome band.

```
> annType <- c("db", "env")
> entrezMap <- getAnnMap("ENTREZID", annotation(bcrAblOrNeg),
+                         type=annType, load=TRUE)
> symbolMap <- getAnnMap("SYMBOL", annotation(bcrAblOrNeg),
+                         type=annType, load=TRUE)
> bandMap <- getAnnMap("MAP", annotation(bcrAblOrNeg),
+                       type=annType, load=TRUE)
```

We applied a non-specific filter to the dataset to remove probesets lacking the desired annotation as well as those with an interquartile range (IQR) below the median IQR, as probesets with little variation across samples are uninformative. We also ensured that each Entrez Gene identifier maps to exactly one probeset by selecting the probeset with the largest IQR when two or more probesets map to the same Entrez Gene ID.

```
> filterAns <- nsFilter(bcrAblOrNeg,
+                         require.entrez = TRUE,
```

```

+               remove.dupEntrez = TRUE,
+               var.func = IQR, var.cutoff = 0.5)
> nsFiltered <- filterAns$eset

```

We also remove probesets with no gene symbol, as well as those with no mapping to a chromosome band.

```

> hasSYM <- sapply(mget(featureNames(nsFiltered), symbolMap, ifnotfound=NA),
+               function(x) length(x) > 0 && !is.na(x[1]))
> hasMAP <- sapply(mget(featureNames(nsFiltered), bandMap, ifnotfound=NA),
+               function(x) length(x) > 0 && !is.na(x[1]))
> nsFiltered <- nsFiltered[hasSYM & hasMAP, ]

```

We define the *gene universe* to be the subset of genes that remain after this filtering.

```

> affyUniverse <- featureNames(nsFiltered)
> entrezUniverse <- unlist(mget(affyUniverse, entrezMap))
> names(affyUniverse) <- entrezUniverse
> if (any(duplicated(entrezUniverse)))
+   stop("error in gene universe: can't have duplicate Entrez Gene Ids")

```

We assessed differential expression between the BCR/ABL and NEG groups using an empirical Bayes approach, as implemented in the software package limma (Smyth, 2005), yielding an attenuated *t*-statistic for each gene.

```

> design <- model.matrix(~ 0 + nsFiltered$mol.biol)
> colnames(design) <- c("BCR/ABL", "NEG")
> contr <- c(1, -1) ## NOTE: we thus have BCR/ABL w.r.t NEG
> fm1 <- lmFit(nsFiltered, design)
> fm2 <- contrasts.fit(fm1, contr)
> fm3 <- eBayes(fm2)
> ttestLimma <- topTable(fm3, number = nrow(fm3), adjust.method = "none")
> ttestLimma <- ttestLimma[featureNames(nsFiltered), ]
> tstats <- ttestLimma$t
> names(tstats) <- entrezUniverse[rownames(ttestLimma)]
> ##

```

We used a *p*-value cutoff of 0.01 to identify a list of potentially differentially expressed genes.

```

> ttestCutoff <- 0.01
> smPV <- ttestLimma$P.Value < ttestCutoff
> pvalFiltered <- nsFiltered[smPV, ]
> selectedEntrezIds <- unlist(mget(featureNames(pvalFiltered), entrezMap))
> ##

```

3 Methods

There are two important features of gene sets based on chromosome bands: (1) the bands are nested hierarchically, and (2) they almost form a partition (for most species almost all genes

appear in only one location in the genome). This naturally leads to two different dichotomies in approaches to testing: one is a top-down versus a bottom-up approach, and the other contrasts local tests with global tests. We use human chromosome 12 as an example to describe these approaches. Users will need to identify which of the approaches best align with their objectives and then make use of the software appropriately.

Conceptually one can start a sequential testing approach either at the most coarse level of organization (probably the level of the arm: p or q), or the most specific level (that of sub-sub-bands). In the top-down approach, one first tests the hypothesis of interest on a coarse level of organization, and if rejected, the next level of organization is considered. For example, we might first consider the band 12q21, and if that hypothesis is rejected, test each of the sub-bands 12q21.1, 12q21.2, and 12q21.3. If the hypothesis for 12q21.3 is rejected, then we may subsequently examine each of 12q21.31, 12q21.32, and 12q21.33.

The bottom-up approach differs in that we begin at the most granular level and move upward, amalgamating adjacent bands at each level. The bottom-up approach is easier to put into a conditional hypothesis testing framework. Our initial null hypotheses involve the smallest, or most granular bands, and if there is evidence that these are unusual (i.e., we reject the null hypothesis) then moving to a larger, or less granular, band requires additional information to declare it significant, over and above what we have used to identify the smaller band. In our example, we would first test 12q21.31, 12q21.32, and 12q21.33, and then move up and test 12q21.3. If one or more of the three sub-bands had been declared significant, we would exclude the evidence from genes annotated in those sub-bands when testing the coarser band.

It is important to note that the top-down versus bottom-up approaches represent a fundamental trade-off between false positive and false negative errors. The bottom-up approach necessarily involves performing a larger number of tests, yielding a correspondingly larger absolute number of false positives for a given false positive rate at which each individual test is controlled. The top-down approach cuts down on the number of false positives by starting with fewer top-level tests, and performing further tests at sublevels only when a top-level test is rejected. The disadvantage to this approach is loss of power to detect real departures that are localized to a sub-level, a phenomenon commonly illustrated using Simpson’s paradox (see, e.g., Wagner, 1982).

Whether a test is local or global is a different question, orthogonal to that of top-down or bottom-up. There are two distinct but potentially relevant questions that may be of interest. The first is whether genes in a particular gene set are “different” relative to all other genes under consideration. For a Hypergeometric test, this question may be formalized as whether the proportion of interesting genes in 12q21 is different from the proportion of interesting genes in the rest of the genome, or equivalently, whether membership in 12q21 is independent of being selected. Such tests are *global* in the sense that all genes in the gene universe are used to determine whether or not genes at a location are unusual or not. An alternative is to ask whether genes in a genomic location are different relative to other genes in some meaningfully defined neighbourhood. Such a test can be performed simply by restricting the gene universe to a suitable subset; for example, when testing 12q21, we may only consider genes in 12q. A more natural approach is to use a 2×3 contingency table to test the hypothesis that the proportion of interesting genes is the same in 12q21, 12q22, and 12q23. Both these tests are *local* in the sense that only nearby genes are used.

Contingency table tests are inherently local and although they do not naturally extend to conditional testing, we can use a top-down approach to test at various resolutions. Such tests can be performed by the `cb_contingency()` function, which we do not discuss in this vignette. Instead, we focus on the bottom-up approach, which allows for conditional testing.

4 Utility functions

We first define a few utility functions that we subsequently use in presentation. The `chrSortOrder()` function reorders rows of data frame for display in a natural order.

```
> chrSortOrder <- function(df) {
+   chrs <- sub("(\\p{q}+).*$", "\\1", rownames(df))
+   xyIdx <- chrs %in% c("X", "Y")
+   xydf <- NULL
+   if (any(xyIdx)) {
+     chrs <- chrs[!xyIdx]
+     xydf <- df[xyIdx, ]
+     df <- df[!xyIdx, ]
+   }
+   ord <- order(as.integer(chrs), rownames(df))
+   df <- df[ord, ]
+   if (!is.null(xydf))
+     df <- rbind(df, xydf)
+   df
+ }
```

The `gseaTstatStripplot()` function creates a comparative strip plot of the t -statistics for specified bands.

```
> gseaTstatStripplot <- function(bands, g, ..., include.all = FALSE)
+ {
+   chroms <- c(1:22, "X", "Y")
+   chromArms <- c(paste(chroms, "p", sep=""), paste(chroms, "q", sep=""))
+   egid <- lapply(nodeData(g, bands), "[", "geneIds")
+   if (include.all) {
+     egid$All <-
+       unique(unlist(lapply(nodeData(g)[chromArms], "[", "geneIds")))
+   }
+   tdf <- do.call(make.groups, lapply(egid, function(x) tstats[x]))
+   stripplot(which ~ data, tdf, jitter = TRUE, ...)
+ }
>
>
```

The `esetBWPlot()` function creates box-and-whisker plots for every gene in an “*ExpressionSet*”.

```

> esetBWPlot <- function(tmpSet, ..., layout=c(1, nrow(emat)))
+ {
+   emat <- exprs(tmpSet)
+   pd <- pData(tmpSet)
+   probes <- rownames(emat)
+   syms <-
+     sapply(mget(probes, hgu95av2SYMBOL, ifnotfound=NA),
+       function(x) if (all(is.na(x))) "NA" else as.character(x)[1])
+   selectedAffy <-
+     probes %in% affyUniverse[selectedEntrezIds]
+   symsSelected <- syms[selectedAffy]
+   symsWithStatus <-
+     paste(syms,
+       ifelse(selectedAffy, "*", ""),
+       sep = "")
+   pdat <-
+     cbind(exprs=as.vector(emat),
+       genes=factor(probes, levels = probes, labels = syms),
+       pd[rep(seq_len(nrow(pd)), each=nrow(emat)), ])
+   pdat <- transform(pdat, genes = reorder(genes, exprs))
+   panels.to.shade <- levels(pdat$genes) %in% symsSelected
+   bwplot(mol.biol ~ exprs | genes, data=pdat,
+     layout = layout,
+     auto.key=TRUE,
+     scales=list(x=list(log=2L)),
+     xlab="Log2 Expression",
+     panels.to.shade = panels.to.shade,
+     panel = function(..., panels.to.shade) {
+       if (panels.to.shade[packet.number()])
+         panel.fill(col = "lightgrey")
+       panel.bwplot(...)
+     },
+     strip=FALSE,
+     strip.left=TRUE, ...)
+ }
> g1 <- makeChrBandGraph(annotation(nsFiltered), univ=entrezUniverse)
> ct <- ChrBandTreeFromGraph(g1)
> subsetByBand <- function(eset, ct, band) {
+   egIDs <- unlist(nodeData(ct@toChildGraph, n=band,
+     attr="geneIds"), use.names=FALSE)
+   wantedProbes <- affyUniverse[as.character(egIDs)]
+   eset[intersect(wantedProbes, featureNames(eset)), ]
+ }
>

```

5 Hypergeometric Testing

We use a method similar to that described in Falcon and Gentleman (2007) to conditionally test for over-representation of chromosome bands in the selected gene list. A test is set up by creating a suitable object of class “*ChrMapHyperGParams*”. We first create an object to perform a standard Hypergeometric analysis, treating each chromosome band independently, and then modify a copy to represent a conditional Hypergeometric computation.

```
> params <- new("ChrMapHyperGParams",
+               conditional=FALSE,
+               testDirection="over",
+               universeGeneIds=entrezUniverse,
+               geneIds=selectedEntrezIds,
+               annotation="hgu95av2",
+               pvalueCutoff=0.05)
> paramsCond <- params
> paramsCond@conditional <- TRUE
```

The test computations are performed by

```
> hgans <- hyperGTest(params)
> hgansCond <- hyperGTest(paramsCond)
```

The results can be summarized by

```
> sumUn <- summary(hgans, categorySize=1)
> chrSortOrder(sumUn)
```

	ChrMapID	Pvalue	OddsRatio	ExpCount	Count	Size
1	14q22.2	0.0006148332	22.305949	0.4986034	4	6
2	7q31.2	0.0050133381	16.682203	0.4155028	3	5
3	9q21.1	0.0065743613	7.427762	0.8310056	4	10
4	13q31	0.0068879625	Inf	0.1662011	2	2
5	13q31.1	0.0068879625	Inf	0.1662011	2	2
6	6p23	0.0068879625	Inf	0.1662011	2	2
7	9q33.2	0.0094118429	11.118644	0.4986034	3	6
8	2p22.2	0.0154651699	8.336864	0.5817039	3	7
9	12q21.33	0.0195249836	22.185915	0.2493017	2	3
10	2q31.3	0.0195249836	22.185915	0.2493017	2	3
11	4q22.1	0.0195249836	22.185915	0.2493017	2	3
12	14q22	0.0210288739	3.715909	1.6620112	5	20
13	6q2	0.0230390078	2.165196	5.6508380	11	68
14	12q14	0.0232403572	6.667797	0.6648045	3	8
15	1p36.11	0.0306574202	4.046356	1.2465084	4	15
16	12q21.3	0.0369130300	11.090141	0.3324022	2	4
17	4q22	0.0369130300	11.090141	0.3324022	2	4
18	10p11.22	0.0369130300	11.090141	0.3324022	2	4

19	12q13.11	0.0369130300	11.090141	0.3324022	2	4
20	6q24.2	0.0369130300	11.090141	0.3324022	2	4
21	1p21	0.0439701634	4.760291	0.8310056	3	10
22	20p12	0.0439701634	4.760291	0.8310056	3	10
23	22q11.23	0.0439701634	4.760291	0.8310056	3	10
24	10p	0.0440389351	2.367273	3.3240223	7	40
25	10p1	0.0440389351	2.367273	3.3240223	7	40
26	9	0.0453714799	1.608622	12.6312849	19	152

```
> sumCond <- summary(hgansCond, categorySize=1)
> chrSortOrder(sumCond)
```

	ChrMapID	Pvalue	OddsRatio	ExpCount	Count	Size
1	14q22.2	0.0006148332	22.305949	0.4986034	4	6
2	7q31.2	0.0050133381	16.682203	0.4155028	3	5
3	9q21.1	0.0065743613	7.427762	0.8310056	4	10
4	13q31.1	0.0068879625	Inf	0.1662011	2	2
5	6p23	0.0068879625	Inf	0.1662011	2	2
6	9q33.2	0.0094118429	11.118644	0.4986034	3	6
7	2p22.2	0.0154651699	8.336864	0.5817039	3	7
8	12q21.33	0.0195249836	22.185915	0.2493017	2	3
9	2q31.3	0.0195249836	22.185915	0.2493017	2	3
10	4q22.1	0.0195249836	22.185915	0.2493017	2	3
11	6q2	0.0230390078	2.165196	5.6508380	11	68
12	12q14	0.0232403572	6.667797	0.6648045	3	8
13	1p36.11	0.0306574202	4.046356	1.2465084	4	15
14	10p11.22	0.0369130300	11.090141	0.3324022	2	4
15	12q13.11	0.0369130300	11.090141	0.3324022	2	4
16	6q24.2	0.0369130300	11.090141	0.3324022	2	4
17	1p21	0.0439701634	4.760291	0.8310056	3	10
18	20p12	0.0439701634	4.760291	0.8310056	3	10
19	22q11.23	0.0439701634	4.760291	0.8310056	3	10
20	10p1	0.0440389351	2.367273	3.3240223	7	40
21	9	0.0453714799	1.608622	12.6312849	19	152

For the standard test, the structure of the chromosome band graph is ignored and a Hypergeometric test is applied to each band independently. For the conditional test, the hierarchical relationship among the bands as represented by the graph is used in the computation. The highest-resolution bands (those with no children in the graph) are tested first. Testing proceeds with the bands whose children (sub-bands) have already been tested. For these bands, the gene annotations that are inherited from significant child nodes (children with p -value smaller than the specified cutoff) are removed prior to testing to yield a conditional test.

The effect of the conditional test is illustrated by examining the results for 14q and its sub-bands. In the standard test, we see that 14q22 and 14q22.2 both have a significant p -value. In the conditional test, only 14q22.2 remains. The conclusion is that there is not enough additional evidence beyond that provided by 14q22.2 to mark 14q22 as significant.

6 GSEA using linear models

GSEA is a popular method that can be used to assess whether or not particular gene sets are associated with a phenotype of interest (Subramanian et al., 2005; Tian et al., 2005; Jiang and Gentleman, 2007). Most applications of this method do not explicitly deal with structure of the gene sets, but when analyzing chromosomal location such methods are desirable. We present a simple approach that is similar in spirit to traditional GSEA, and generalizes nicely to accommodate nested categories. Consider the situation where gene i has an associated measure of differential expression y_i ; for example, an attenuated t -statistic derived from a differential expression analysis such as *limma* (Smyth, 2005). Given a particular category, GSEA asks whether the distribution of y_i -s restricted to the category of interest is “unusual”. Thus, in Figure 2, we might be interested in knowing whether the distribution of y_i values for genes in 12q21 is different from that of the genes in 12q (for a local test) or of all genes in the gene universe (for a global test). Figure 2 is produced by

```
> gseaTstatStripplot(c("12q21.1", "12q21", "12q2", "12q"),
+                     include.all = TRUE,
+                     g = g1,
+                     xlab = "Per-gene t-statistics",
+                     panel = function(...) {
+                         require(grid, quietly = TRUE)
+                         grid.rect(y = unit(2, "native"),
+                                   height = unit(1, "native"),
+                                   gp =
+                                   gpar(fill = "lightgrey",
+                                         col = "transparent"))
+                         panel.grid(v = -1, h = 0)
+                         panel.stripplot(...)
+                         panel.average(..., fun = mean, lwd = 3)
+                     })
```

We fit a factorial model to see whether the distribution of y_i is associated with category membership. Specifically, for category j , we fit the model

$$y_i = \beta_0 + \beta_1 a_{ij} + \varepsilon_i \quad (1)$$

where $a_{ij} = 1$ if gene i belongs to category j , and 0 otherwise. The index i may range over the full gene universe, or a subset, depending on whether one wishes to perform global or local tests. The null hypothesis of no association is represented by $H_0 : \beta_1 = 0$. The model nominally assumes that the y_i -s are Normally distributed with equal variance, but in practice the results are robust against mild deviations. The presence of an intercept term allows nonzero overall mean, which can be important in many situations, especially for local tests. We expect the test to be fairly insensitive to distributional assumptions on the y_i -s.

We can fit (1) by least squares and test $H_0 : \beta_1 = 0$ to obtain a marginal test for each category j ; in this case, each chromosome band. The procedure also generalizes to incorporate the nesting structure of chromosome bands. Specifically, if band j_2 (e.g., 12q21.1) is nested within a coarser band j_1 (e.g., 12q21) and the more granular band j_2 is significant, then the




Figure 2: Per-gene t -statistics, as computed by *limma*, for selected chromosome bands. The points are jittered vertically to alleviate overlap. A thick grey line joins the mean value within each group. A GSEA test would compare the genes in 12q21 with those in 12q, or the entire gene universe, by fitting a linear model with per-gene t -statistics as the response, and a term indicating membership in 12q21 as the predictor. If 12q21.1 is declared as significant, a conditional GSEA test would include an additional term for 12q21.1.

effect of membership in j_1 over and above the effect attributable to membership in j_2 can be tested by fitting the model

$$y_i = \beta_0 + \beta_1 a_{ij_1} + \beta_2 a_{ij_2} + \varepsilon_i \quad (2)$$

and testing the null hypothesis $H_0 : \beta_1 = 0$. Multiple significant sub-bands and multiple levels of nesting can be incorporated by including further terms in the model. The complete process can be summarized as follows: Start by marginally testing each band which has no sub-bands. For all other bands, first test all sub-bands, then test the current band using a linear model that includes a term for each *significant* sub-band.

We apply this procedure to perform global tests using per-gene t -statistics as a measure of differential expression in BCR/ABL relative to NEG samples. As with the Hypergeometric tests, we start by creating objects of class “*ChrMapLinearMParams*”.

```
> params <- new("ChrMapLinearMParams",
+               conditional = FALSE,
+               testDirection = "up",
+               universeGeneIds = entrezUniverse,
+               geneStats = tstats,
+               annotation = "hgu95av2",
+               pvalueCutoff = 0.01,
+               minSize = 4L)
> params@graph <- makeChrBandGraph(params@annotation, params@universeGeneIds)
> params@gsc <- makeChrBandGSC(params@graph)
```

```
> paramsCond <- params
> paramsCond@conditional <- TRUE
```

The tests are performed, and the results summarized by

```
> lmans <- linearMTest(params)
> lmansCond <- linearMTest(paramsCond)
> chrSortOrder(summary(lmans))
```

	ChrMapID	Pvalue	Effect	Size	Conditional	TestDirection
53	12q21	1.386566e-03	1.1132172	17	FALSE	up
54	12q21.3	8.112109e-03	1.8418037	4	FALSE	up
76	14	6.492349e-03	0.3112771	155	FALSE	up
77	14q	6.492349e-03	0.3112771	155	FALSE	up
81	14q2	2.440494e-04	0.6177206	76	FALSE	up
83	14q22	1.353609e-04	1.2491242	20	FALSE	up
133	18p	1.206333e-03	0.8807705	28	FALSE	up
134	18p1	1.206333e-03	0.8807705	28	FALSE	up
135	18p11	1.206333e-03	0.8807705	28	FALSE	up
137	18p11.3	6.233312e-03	1.0630065	13	FALSE	up
161	1p2	3.077861e-03	0.7561752	31	FALSE	up
162	1p21	9.256086e-04	1.5089370	10	FALSE	up
216	2	2.273852e-04	0.3257970	291	FALSE	up
227	2q	7.597359e-05	0.4601438	165	FALSE	up
228	2q1	6.433980e-03	0.6207694	38	FALSE	up
229	2q11	5.758756e-03	1.1186350	12	FALSE	up
237	2q3	3.364310e-03	0.4283786	96	FALSE	up
240	2q33	6.867392e-03	0.9452907	16	FALSE	up
243	3	4.451464e-03	0.2671785	238	FALSE	up
256	3q	6.295149e-04	0.4608853	118	FALSE	up
263	3q2	2.769151e-03	0.4340754	98	FALSE	up
266	3q25	4.381852e-04	1.2030677	18	FALSE	up
271	4	2.645253e-03	0.3461723	158	FALSE	up
277	4q	3.649205e-03	0.3968617	110	FALSE	up
279	4q13	6.759492e-03	0.8934856	18	FALSE	up
291	5	3.549008e-04	0.3752981	200	FALSE	up
297	5q	2.463796e-03	0.3379143	169	FALSE	up
303	5q2	1.702037e-03	1.0897674	17	FALSE	up
306	5q23	3.818815e-03	1.6690195	6	FALSE	up
311	6	2.316308e-04	0.3404264	264	FALSE	up
322	6q	8.401000e-06	0.6794107	96	FALSE	up
326	6q2	6.899269e-06	0.8126623	68	FALSE	up
330	6q24	3.457299e-03	1.5646194	7	FALSE	up
348	7q21.1	9.883068e-03	1.1302842	10	FALSE	up
385	9q	5.937007e-03	0.3721895	110	FALSE	up
387	9q21	3.237225e-04	1.2331567	18	FALSE	up

388	9q21.1	4.994149e-06	2.1386990	10	FALSE	up
857	10p11.22	8.509458e-03	1.8283876	4	FALSE	up
988	14q22.2	1.022892e-05	2.6617280	6	FALSE	up
1019	15q24.3	9.012446e-03	1.6210440	5	FALSE	up
1033	16q12.1	4.493977e-03	1.6345576	6	FALSE	up
1067	18p11.22	1.098996e-03	1.6590621	8	FALSE	up
1106	1p21.3	2.696967e-03	1.7407370	6	FALSE	up
1218	2q11.2	9.261727e-03	1.0889984	11	FALSE	up
1316	4q13.3	9.132299e-03	1.0450689	12	FALSE	up
1390	6p22.3	3.301504e-03	1.1554181	13	FALSE	up
1407	6q21	1.461539e-03	1.1072339	17	FALSE	up
1411	6q23.3	3.448624e-03	1.5651026	7	FALSE	up
1413	6q24.2	7.851576e-03	1.8509171	4	FALSE	up
1444	7q31.2	4.484345e-05	2.6816192	5	FALSE	up
1499	9q21.13	8.252375e-04	1.9681613	6	FALSE	up
1513	9q33.2	4.188486e-04	2.0886878	6	FALSE	up

> chrSortOrder(summary(lmansCond))

	ChrMapID	Pvalue	Effect	Size	Conditional	TestDirection
54	12q21.3	8.112109e-03	1.8418037	4	TRUE	up
81	14q2	8.223851e-03	0.4419018	76	TRUE	up
137	18p11.3	6.233312e-03	1.0630065	13	TRUE	up
227	2q	4.011884e-03	0.3509407	165	TRUE	up
240	2q33	6.867392e-03	0.9452907	16	TRUE	up
266	3q25	4.381852e-04	1.2030677	18	TRUE	up
297	5q	9.128698e-03	0.2885130	169	TRUE	up
306	5q23	3.818815e-03	1.6690195	6	TRUE	up
348	7q21.1	9.883068e-03	1.1302842	10	TRUE	up
388	9q21.1	8.874923e-04	2.3911605	10	TRUE	up
857	10p11.22	8.509458e-03	1.8283876	4	TRUE	up
988	14q22.2	1.022892e-05	2.6617280	6	TRUE	up
1019	15q24.3	9.012446e-03	1.6210440	5	TRUE	up
1033	16q12.1	4.493977e-03	1.6345576	6	TRUE	up
1067	18p11.22	1.098996e-03	1.6590621	8	TRUE	up
1106	1p21.3	2.696967e-03	1.7407370	6	TRUE	up
1218	2q11.2	9.261727e-03	1.0889984	11	TRUE	up
1316	4q13.3	9.132299e-03	1.0450689	12	TRUE	up
1390	6p22.3	3.301504e-03	1.1554181	13	TRUE	up
1407	6q21	1.461539e-03	1.1072339	17	TRUE	up
1411	6q23.3	3.448624e-03	1.5651026	7	TRUE	up
1413	6q24.2	7.851576e-03	1.8509171	4	TRUE	up
1444	7q31.2	4.484345e-05	2.6816192	5	TRUE	up
1499	9q21.13	8.252375e-04	1.9681613	6	TRUE	up
1513	9q33.2	4.188486e-04	2.0886878	6	TRUE	up




Figure 3: Expression values for the BCR/ABL and NEG samples for the genes located within 1p36.2, which is declared as significant by the Hypergeometric test, but not by GSEA. Genes in the selected list are highlighted with a shaded background. For these genes, The NEG samples, those with no known cytogenetic abnormalities, have significantly lower expression than the BCR/ABL samples. However, the direction is reversed (albeit mildly) for many of the other genes.

```
>
> ##
```

These examples only test for consistently upregulated categories; similar calls with `testDirection = "down"` can be used to test for downregulation. As we see, the GSEA approach picks out many more bands as significant, but there is some concordance with the Hypergeometric approach. For example, 7q31, 8p22, and 14q22.2 come up in both analyses. Figure 3 shows box-and-whisker plots of genes in one category (1p36.2) that is declared as significant by the Hypergeometric test, but not by GSEA. It is produced by

```
> tmpSet <- subsetByBand(nsFiltered, ct, "1p36.2")
> esetBWPlot(tmpSet, ylab="1p36.2", layout = c(2, 8),
+           par.strip.text = list(cex = 0.8))
```

References

- S. Chiaretti, X. Li, R. Gentleman, A. Vitale, K. S. Wang, F. Mandelli, R. Foa, and J. Ritz. Gene expression profiles of b-lineage adult acute lymphocytic leukemia reveal genetic patterns that identify lineage derivation and distinct mechanisms of transformation. *Clinical Cancer Research*, 11:7209–7219, 2005.
- S Falcon and R Gentleman. Using GOstats to test gene lists for GO term association. *Bioinformatics*, 23(2):257–8, 2007.
- T. S. Furey and D. Haussler. Integration of the cytogenetic map with the draft human genome sequence. *Human Molecular Genetics*, 12:1037–1044, 2003. URL DOI:10.1093/hmg/ddg113.
- F. Hahne, W. Huber, R. Gentleman, and S. Falcon. *Bioconductor Case Studies*. Springer, New York, 2008.
- Z. Jiang and R. Gentleman. Extensions to gene set enrichment. *Bioinformatics*, 23:306–313, 2007.
- Gordon K. Smyth. Limma: linear models for microarray data. In R. Gentleman, V. Carey, S. Dudoit, R. Irizarry, and W. Huber, editors, *Bioinformatics and Computational Biology Solutions using R and Bioconductor*, pages 397–420. Springer, New York, 2005.
- A. Subramanian, P. Tamayo, V. K. Mootha, et al. Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences*, 102(43):15545–15550, 2005.
- L. Tian, S. A. Greenberg, S. W. Kong, et al. Discovering statistically significant pathways in expression profiling studies. *Proceedings of the National Academy of Sciences*, 102(38):13544–13549, 2005.
- Clifford H. Wagner. Simpson’s paradox in real life. *The American Statistician*, 36(1):46–48, 1982. ISSN 00031305.