

Quantile-based permutation thresholds for QTL hotspot analysis: a tutorial

Elias Chaibub Neto* and Brian S Yandell†

November 26, 2024

1 Motivation

QTL hotspots, groups of traits co-mapping to the same genomic location, are a common feature of genetical genomics studies. Genomic locations associated with many traits are biologically interesting since they may harbor influential regulators. Nonetheless, non-genetic mechanisms, uncontrolled environmental factors and unmeasured variables are capable of inducing a strong correlation structure among clusters of transcripts, and as a consequence, whenever a transcript shows a spurious linkage, many correlated transcripts will likely map to the same locus, creating a spurious QTL hotspot. Permutation approaches that do not take into account the phenotypic correlation tend to underestimate the size of the hotspots that might appear by change in these situations (Breitling et al. 2008).

This issue motivated the development of permutation tests that preserve the correlation structure of the phenotypes in order to determine the significance of QTL hotspots (Breitling et al. 2008, Chaibub Neto et al. 2012). In this tutorial we present software tools implementing the *NL*-method (Chaibub Neto et al. 2012), the *N*-method (Breitling et al. 2008), and the *Q*-method (West et al. 2007, Wu et al. 2008) permutation approaches.

2 Overview

This tutorial illustrates the application of the *NL*-, *N*- and *Q*-methods, implemented in the `qtlhot` R package, to a few toy examples. The `qtlhot` package is built over the `R/qtl` package (Broman et al. 2003), and we assume the reader is familiar with it.

3 Basic functionality with No Real Hotspots

In this section we consider two toy simulated examples. In the first we simulate highly correlated phenotypes. In the second, we simulate uncorrelated phenotypes.

*Department of Computational Biology, Sage Bionetworks, Seattle WA

†Department of Statistics, University of Wisconsin-Madison, Madison WI

```
> library(qtlhot)
```

We start by simulating a “null backcross” data set composed of 1,000 phenotypes, 4 chromosomes, 51 equally spaced genetic markers per chromosome, and 100 individuals, with the `sim.null.cross` function. The `latent.eff` parameter control the amount of correlation among the phenotypes. Each phenotype k is generated according to the model $Y_k = \theta L + \epsilon_k$, where $L \sim N(0, \sigma^2)$ is a latent variable, θ represents the effect (`latent.eff`) of the latent variable on the phenotype, and $\epsilon_k \sim N(0, \sigma^2)$ represents a residual error term with σ^2 set to `res.var`. Note that we do not simulate any QTLs in a “null cross” and any linkages we might detect in such a data set are due entirely to chance.

```
> ncross1 <- sim.null.cross(chr.len = rep(100, 4),
+                           n.mar = 51,
+                           n.ind = 100,
+                           type = "bc",
+                           n.phe = 1000,
+                           latent.eff = 3,
+                           res.var = 1,
+                           init.seed = 123457)
```

The function `include.hotspots` takes the “null cross” as an input and includes 3 hotspots of size `hsize` at position `hpos` of chromosome `hchr` into it. Explicitly, it simulates each one of the hotspots according to the model $Y_k^* = \beta M + Y_k$, where Y_k is the phenotype generated by the `generate.null.cross` function; $M = \gamma Q + \epsilon_M$ is a master regulator that affects all phenotypes in the hotspot; Q is a QTL located at position `hpos` of chromosome `hchr`; γ represents the QTL effect (`Q.eff`); $\epsilon_M \sim N(0, \sigma^2)$; and β is computed such that the association between Y_k^* and Q , measured by the LOD score, is given (theoretically) by a valued sampled from the user specified LOD score range (`lod.range.1` and etc).

```
> cross1 <- include.hotspots(cross = ncross1,
+                            hchr = c(2, 3, 4),
+                            hpos = c(25, 75, 50),
+                            hsize = c(100, 50, 20),
+                            Q.eff = 2,
+                            latent.eff = 3,
+                            lod.range.1 = c(2.5, 2.5),
+                            lod.range.2 = c(5, 8),
+                            lod.range.3 = c(10, 15),
+                            res.var = 1,
+                            n.phe = 1000,
+                            init.seed = 12345)
```

Note that by choosing `latent.eff = 3` we generate highly correlated phenotype data. The distribution of the correlation values for each pair of phenotypes is given below.

```

> ncor1 <- cor(cross1$pheno)
> summary(ncor1[lower.tri(ncor1)])

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.4343 0.8559 0.8929 0.8654 0.9063 0.9725

> rm(ncor1)

```

Next we obtain standard permutation thresholds (Churchill and Doerge 1994) for single trait QTL mapping analysis for the sequence `alphas`, representing target genome wide error rates (GWER).

```

> set.seed(123)
> pt <- scanone(ncross1, method = "hk", n.perm = 1000)

> alphas <- seq(0.01, 0.10, by=0.01)
> lod.thrs <- summary(pt, alphas)
> lod.thrs

```

LOD thresholds (1000 permutations)

```

      lod
1%  3.11
2%  2.89
3%  2.68
4%  2.57
5%  2.44
6%  2.34
7%  2.26
8%  2.20
9%  2.15
10% 2.11

```

```

> lod.thr <- lod.thrs[5]

```

We perform QTL mapping analysis for 1,000 phenotypes using Haley-Knott regression, and keep only the `drop.lod = 1.5` LOD support interval (Manichaikul et al. 2006) around significant peaks (above `lod.thr`) at each chromosome for each trait. LOD support intervals are the most commonly used interval estimate for the location of a QTL.

```

> scan1 <- qtl::scanone(cross1, pheno.col = 1:1000, method = "hk")

```

The routine `highlod` just saves the LOD support intervals for significant peaks, which dramatically reduces the object size from the result of `scanone`. We can examine the maximum hotspot size for each possible single trait threshold using the `max` method for `highlod` objects. The first column gives the counts associated with QTL mapping threshold of 3.11, whereas the last one shows the counts based on the more liberal threshold 2.11. Note how the counts increase as the QTL mapping thresholds decrease.

```
> high1 <- highlod(scan1, lod.thr = min(lod.thrs), drop.lod = 1.5)
> max(high1, lod.thr = lod.thrs)
```

	chr	pos	max.N	lod.thr
D3M36	3	70	50	3.11
D2M12	2	22	56	2.89
D2M13	2	24	67	2.68
D2M131	2	24	78	2.57
D2M121	2	22	83	2.44
D1M34	1	66	96	2.34
D1M341	1	66	123	2.26
D1M342	1	66	150	2.20
D1M343	1	66	169	2.15
D1M344	1	66	189	2.11

Next we infer the hotspot architecture at varying QTL mapping thresholds. In other words, for each genomic position, we count the number of traits that map to it with a LOD score equal or higher than the threshold in `lod.thr`. The `hots1` object is a `scanone` object with added attributes and specialized summary and plot methods. The summary shows the counts for the most significant loci per chromosome, essentially using the `summary` method for `scanone` objects.

```
> hots1 <- hotsize(high1, lod.thr = lod.thr)
> summary(hots1)
```

```
hotsize elements: chr pos max.N
LOD threshold: 2.438697
```

	chr	pos	max.N
D1M34	1	66	71
D2M12	2	22	83
D3M37	3	72	50
D4M27	4	52	20

We plot the hotspot architecture inferred using the single trait permutation threshold 2.44 ($\alpha = 0.05$). Figure 1 shows the counts across the genome. Recall that in the call of function `include.hotspots` we set to simulate 3 hotspots: (1) a hotspot of size 100 at position 25cM of chromosome 2 with LOD scores around 2.5; (2) a hotspot of size 50 at position 75cM of chromosome 3 with LOD scores ranging from 5 to 8; and (3) a hotspot of size 20 at position 50cM of chromosome 4 with LOD scores ranging from 10 to 15. Nonetheless, Figure 1 shows several spurious peaks on chromosome 1, that arise because of the high correlation of the phenotypes.

```
> plot(hots1, cex.lab = 1.5, cex.axis = 1.5)
```

Next, we perform permutation tests to assess the statistical significance of the hotspots detected on Figure 1. We consider the N - and NL -methods. [The Q method of West and Wu is

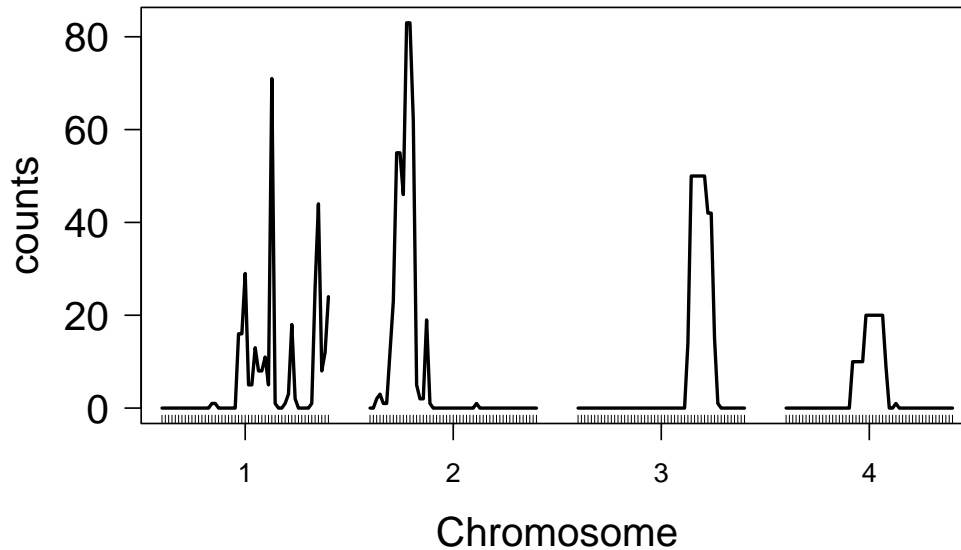


Figure 1: Hotspot architecture associated with QTL mapping threshold of 2.44 in example 1.

documented in the `ww.perm` manual page but not shown here.] The `hotperm` function implements the N - and NL -methods' permutation schemes (see Chaibub Neto et al. 2012, for details). The parameter `n.quant` sets the maximum hotspot size to be analyzed by the NL -method. The parameter `drop.lod` controls the magnitude of the LOD support interval computation during the LOD profile processing step. The function's output is a list with two elements: `max.lod.quant` and `max.N`.

```
> set.seed(12345)
> hotperm1 <- hotperm(cross = cross1,
+                   n.quant = 300,
+                   n.perm = 100,
+                   lod.thrs = lod.thrs,
+                   alpha.levels = alphas,
+                   drop.lod = 1.5,
+                   verbose = FALSE)

> names(hotperm1)

[1] "max.lod.quant" "max.N"

> summary(hotperm1)

max.N: hotspot threshold by single-trait LOD threshold and significance level
      99% 98% 97% 96% 95% 94% 93% 92% 91% 90%
3.11  52  43  39  30  28  10   8   7   7   6
```

```

2.89 80 72 62 59 51 22 19 17 16 13
2.68 156 151 134 96 59 51 44 41 28 21
2.57 209 191 188 120 77 71 67 60 35 31
2.44 294 275 270 177 119 110 92 79 63 60
2.34 375 345 332 227 163 154 127 95 88 87
2.26 441 407 379 270 203 199 162 116 112 111
2.2 485 452 426 306 240 233 201 145 139 132
2.15 516 497 472 334 267 257 234 169 160 153
2.11 553 532 504 373 306 287 266 195 181 178

```

```

max.lod.quant: LOD threshold by hotspot size quantile and significance level
      99% 98% 97% 96% 95% 94% 93% 92% 91% 90%
1    5.42 4.90 4.63 4.35 4.28 4.20 4.12 4.03 3.91 3.78
2    4.98 4.74 4.51 4.27 4.00 3.74 3.71 3.68 3.59 3.51
5    4.79 4.69 4.59 4.44 4.28 4.12 3.92 3.71 3.51 3.49
10   4.48 4.46 4.44 4.38 4.23 4.08 3.92 3.78 3.66 3.53
20   4.23 4.20 4.18 4.16 4.14 4.03 3.92 3.80 3.69 3.58
50   3.84 3.74 3.64 3.55 3.45 3.35 3.25 3.15 3.06 3.05
100  3.72 3.66 3.60 3.54 3.48 3.42 3.36 3.31 3.25 3.19
200  3.38 3.36 3.34 3.31 3.29 3.26 3.24 3.21 3.19 3.16
300  3.23 3.23 3.23 3.23 3.23 3.23 3.23 3.23 3.23 3.23

```

The `max.N` element of the `hotperm1` object stores the output of the N -method's permutations and is given by a matrix with 100 rows representing the permutations, and 10 columns representing the QTL mapping thresholds. Entry ij stores the maximum genome wide hotspot size detected at permutation i using the QTL mapping threshold j . The `max.N` element of the `summary` is a 10 by 10 matrix with rows indexing the QTL mapping thresholds and columns indexing the target genome wide error rates. Each entry ij shows the hotspot size above which a hotspot is considered significant at a GWER j using the QTL mapping threshold i . As before, our interest focus on the diagonal, and the N -method's threshold that controls the hotspot GWER at a 5% level when the QTL mapping was controlled at a GWER of 5% (single trait LOD threshold of 2.44) is 119 traits per hotspot (rounded up). Note that according to the N -method, none of the hotspots on Figure 1 is significant.

The `max.lod.quant` element of the `hotperm1` object stores the output of the NL -method's permutations and is given by a matrix with 100 rows representing the permutations, and 300 columns representing the hotspot sizes analyzed. Entry ij stores the maximum genome wide $qLOD(n)$ value—the n th LOD score in a sample ordered from highest to lowest—computed at permutation i using the QTL mapping threshold j . The `max.lod.quant` element of the `summary` shows just the extremes and quartiles.

The `max.lod.quant` element of the `quantile` contains the sliding LOD quantiles from 1 up to 300 for each single trait LOD threshold. This can be used to plot how many traits pass the sliding LOD threshold, and are significant, at each locus in the genome. Figure 2 shows the hotspot significance profile for the thresholds targeting GWER at a 5% level. Note that we only consider LODs above `lod.thr`, and chromosome 1 has none.

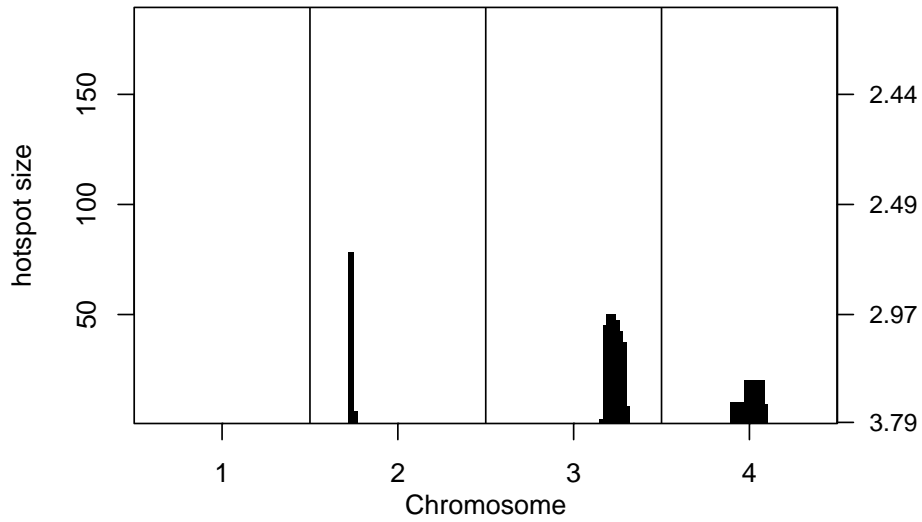


Figure 2: Hotspot size significance profile targeting GWER at a 5% level for example 1.

```
> quant1 <- quantile(hotperm1, 0.05, lod.thr = lod.thr)
> plot(high1, quant.level = quant1, sliding = TRUE)
```

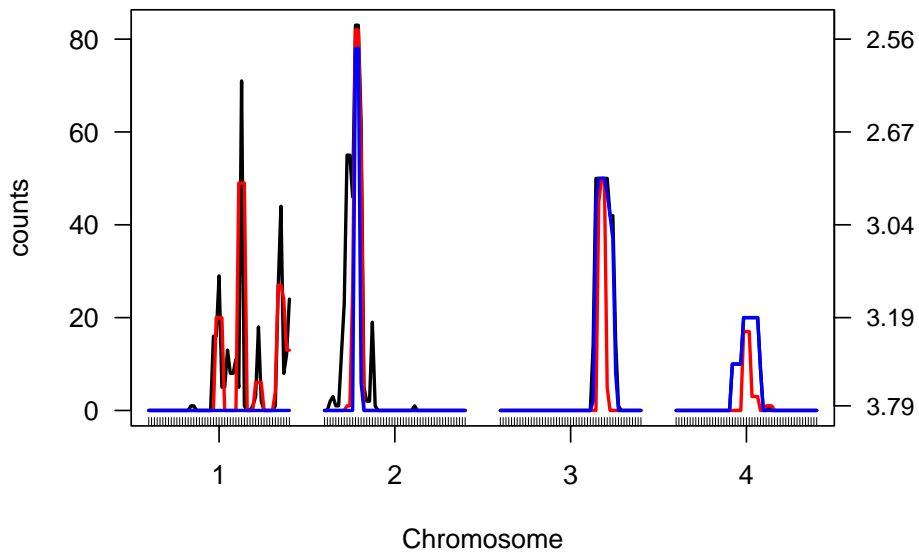


Figure 3: Hotspot size significance profile at 5% level for example 1. Black is raw counts, red is counts smoothed over a 5cM window, blue is the sliding LOD threshold approach.

Figure 2 depicts a sliding window of hotspot size thresholds ranging from $n = 1, \dots, N$, where $N = 100$ corresponds to the (approximate) hotspot size threshold derived from the N -method. For each genomic location this figure shows the hotspot sizes at which the hotspot was significant, that is, at which the hotspot locus had more traits than the hotspot size threshold on the left mapping to it with a LOD score higher than the threshold on the right than expected by chance. For example, the hotspot on chromosome 3 was significant at all sizes from 1 to 50, meaning that at least 1 trait mapped to the hotspot locus with LOD higher than 3.79, at least 20 traits mapped to the hotspot locus with LOD higher than 3.19, and so on up to hotspot size 50, all with LODs higher than 2.97. The N -method that did not detect any hotspots, but the NL -method's sliding window correctly detected the simulated hotspots and showed that the apparent hotspots on chromosome 1 were noisy artifacts.

Figure 3 is another way to visualize this, using the sliding LOD threshold (blue curve). We also show a 5-cM smoothing window applied to the counts (red curve). Note that with the sliding LOD threshold, no hotspots (or individual traits) are significant on chr 1, and hotspots of size 20 are found on chr 2 and 4, with a significant hotspot of size 50 on chr 3.

```
> hotsq1 <- hotsize(high1, lod = lod.thr, window = 5, quant.level = quant1)
> plot(hotsq1)

> summary(hotsq1)
```

```
hotsize elements:  chr pos max.N max.N.window quant
LOD threshold: 2.438697
smooth window: 5
quantile level summary:
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 2.439  2.528   2.673   2.826  3.077   3.790
```

```
      chr pos max.N max.N.window quant
D1M34  1  66    71             49    0
D2M13  2  24    83             82   78
D3M39  3  76    50              5   47
D4M28  4  54    20              3   20
```

4 Example with Uncorrelated Phenotypes

Next we consider a second toy example with uncorrelated phenotype data. We repeat the simulation and analysis steps presented previously changing `latent.eff` to zero.

```
> ncross2 <- sim.null.cross(chr.len = rep(100,4),
+                           n.mar = 51,
+                           n.ind = 100,
+                           type = "bc",
+                           n.phe = 1000,
```

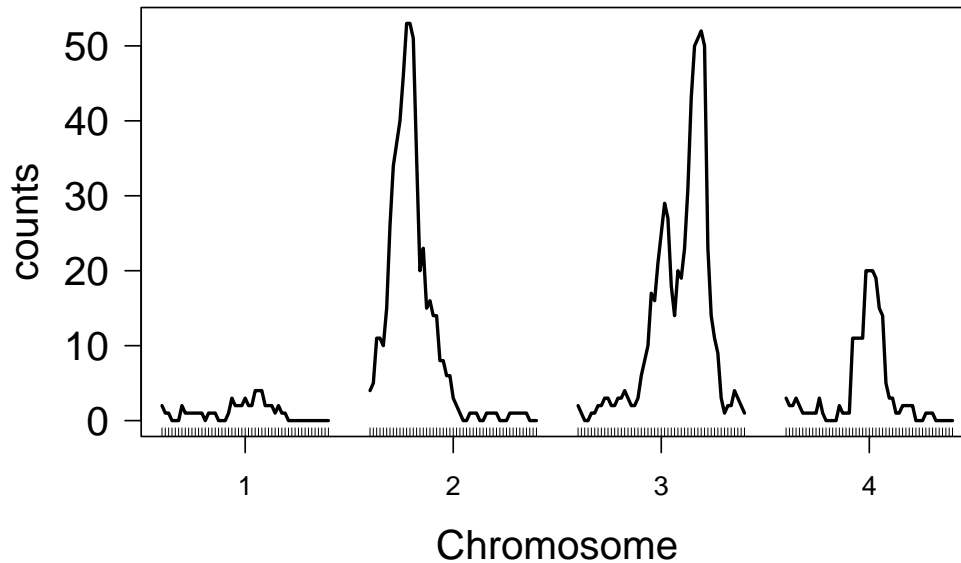



Figure 4: Hotspot architecture targeting 5% GWER for example 2.

```

+             latent.eff = 0,
+             res.var = 1,
+             init.seed = 123457)
> cross2 <- include.hotspots(cross = ncross2,
+             hchr = c(2, 3, 4),
+             hpos = c(25, 75, 50),
+             hsize = c(100, 50, 20),
+             Q.eff = 2,
+             latent.eff = 0,
+             lod.range.1 = c(2.5, 2.5),
+             lod.range.2 = c(5, 8),
+             lod.range.3 = c(10, 15),
+             res.var = 1,
+             n.phe = 1000,
+             init.seed = 12345)

> ncor2 <- cor(cross2$pheno)
> summary(ncor2[lower.tri(ncor2)])

      Min.   1st Qu.   Median     Mean   3rd Qu.     Max.
-0.471212 -0.067418  0.001111  0.002256  0.070412  0.724543

> rm(ncor2)

> scan2 <- scanone(cross2, pheno.col = 1:1000, method = "hk")
> high2 <- highlod(scan2, lod.thr = lod.thr, drop.lod = 1.5)

```

```
> hots2 <- hotsize(high2)
> plot(hots2, cex.lab = 1.5, cex.axis = 1.5)
```

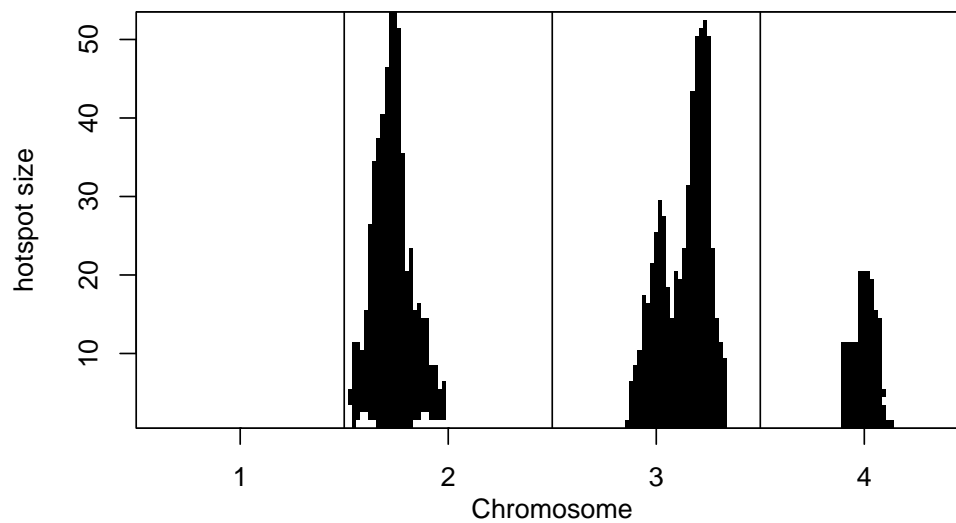


Figure 5: Hotspot significance profile targeting 5% GWER for example 2.

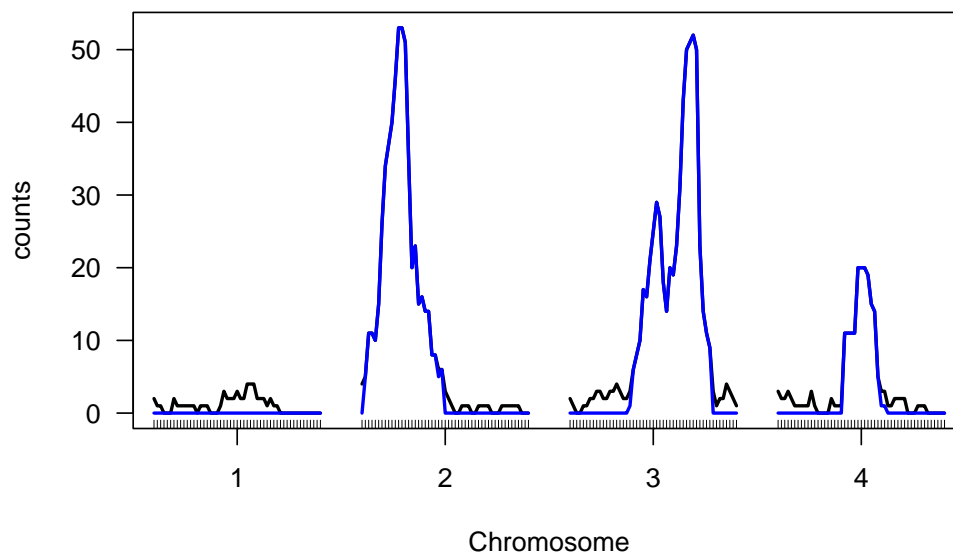


Figure 6: Hotspot significance scan targeting 5% GWER for example 2.

```

> set.seed(12345)
> hotperm2 <- hotperm(cross = cross2,
+                   n.quant = 300,
+                   n.perm = 100,
+                   lod.thrs = lod.thrs,
+                   alpha.levels = alphas,
+                   drop.lod = 1.5,
+                   verbose = FALSE)
> quant2 <- quantile(hotperm2, 0.05, lod.thr = lod.thr)

```

The N -method, as expected, gave rise to much smaller thresholds in this second example with uncorrelated phenotypes. Additionally, inspection of Figure 4 shows no spurious hotspots on chromosome 1.

```

> plot(high2, lod.thr = lod.thr, quant.level = quant2, sliding = TRUE)

```

Figure 5 presents the hotspot significance profile targeting 5% GWER. For this second example, all methods correctly detected the simulated hotspots.

```

> plot(high2, quant.level = quant2)

```

Figure 6 shows another way to represent significant hotspots. We overlay the largest significant hotspot counts using the sliding quantiles in red on top of the curve on Figure 4. Notice that the large sizes are all significant, but only small sizes corresponding to larger LOD scores are significant. We add a right axis with the sliding LOD thresholds.

5 References

1. Breitling R., Y. Li, B. M. Tesson, J. Fu, C. Wu, et al., 2008 Genetical genomics: spotlight on QTL hotspots. *PLoS Genetics* **4**: e1000232.
2. Broman K. W., W. Wu, S. Sen, G. A. Churchill, 2003 R/qtl: QTL mapping in experimental crosses. *Bioinformatics* **19**: 889-890.
3. Chaibub Neto et al., 2012 Quantile-based permutation thresholds for QTL hotspots. *Genetics* (under review).
4. Churchill G. A., and R. W. Doerge, 1994 Empirical threshold values for quantitative trait mapping. *Genetics* **138**: 963-971.
5. Manichaikul A., J. Dupuis, S. Sen, and K. W. Broman, 2006 Poor performance of bootstrap confidence intervals for the location of a quantitative trait locus. *Genetics* **174**: 481-489.
6. West M. A. L., K. Kim, D. J. Kliebenstein, H. van Leeuwen, R. W. Michelmore, R. W. Doerge, D. A. St. Clair 2007 Global eQTL mapping reveals the complex genetic architecture of transcript-level variation in *Arabidopsis*. *Genetics* **175**: 1441-1450.
7. Wu C., D. L. Delano, N. Mitro, S. V. Su, J. Janes, et al. 2008 Gene set enrichment in eQTL data identifies novel annotations and pathway regulators. *PLoS Genetics* **4**: e1000070.