

Imputed SNP analyses and meta-analysis with `snpStats`

David Clayton

November 18, 2024

Getting started

The need for imputation in SNP analysis studies occurs when we have a smaller set of samples in which a large number of SNPs have been typed, and a larger set of samples typed in only a subset of the SNPs. We use the smaller, complete dataset (which will be termed the *training dataset*) to impute the missing SNPs in the larger, incomplete dataset (the *target dataset*). Examples of such applications include:

- use of HapMap data to impute association tests for a large number of SNPs, given data from genome-wide studies using, for example, a 500K SNP array, and
- meta-analyses which seek to combine results from two platforms such as the Affymetrix 500K and Illumina 550K platforms.

Here we will not use a real example such as the above to explore the use of `snpStats` for imputation, but generate a fictitious example using the data analysed in earlier exercises. This is particularly artificial in that we have seen that these data suffer from extreme heterogeneity of population structure.

We start by attaching the required libraries and accessing the data used in the exercises:

```
> library(snpStats)
> library(hexbin)
> data(for.exercise)
```

We shall sample 200 subjects in our fictitious study as the training data set, select every second SNP to be missing in the target dataset, and split the training set into two parts accordingly:

```
> training <- sample(1000, 200)
> select <- seq(1, ncol(snps.10), 2)
> missing <- snps.10[training, select]
> present <- snps.10[training, -select]
> missing
```

```
A SnpMatrix with 200 rows and 14251 columns
Row names:  jpt.360 ... jpt.832
Col names:  rs7909677 ... rs12218790
```

```
> present
```

```
A SnpMatrix with 200 rows and 14250 columns
Row names:  jpt.360 ... jpt.832
Col names:  rs7093061 ... rs7899159
```

Thus the training dataset consists of the objects `missing` and `present`. The target dataset holds a subset of the SNPs for the remaining 800 subjects.

```
> target <- snps.10[-training, -select]
> target
```

```
A SnpMatrix with 800 rows and 14250 columns
Row names:  jpt.869 ... ceu.464
Col names:  rs7093061 ... rs7899159
```

But, in order to see how successful we have been with imputation, we will also save the SNPs we have removed from the target dataset

```
> lost <- snps.10[-training, select]
> lost
```

```
A SnpMatrix with 800 rows and 14251 columns
Row names:  jpt.869 ... ceu.464
Col names:  rs7909677 ... rs12218790
```

We also need to know where the SNPs are on the chromosome in order to avoid having to search the entire chromosome for suitable predictors of a missing SNP:

```
> pos.miss <- snp.support$position[select]
> pos.pres <- snp.support$position[-select]
```

Calculating the imputation rules

The next step is to calculate a set of rules which for imputing the missing SNPs from the present SNPs. This is carried out by the function `snp.imputation`¹:

```
> rules <- snp.imputation(present, missing, pos.pres, pos.miss)
```

¹Sometimes this command generates a warning message concerning the maximum number of EM iterations. If this only concerns a small proportion of the SNPs to be imputed it can be ignored.

SNPs tagged by a single SNP: 4862

SNPs tagged by multiple tag haplotypes (saturated model): 9210

This step executes remarkably quickly when we consider what the function has done. For each of the 14251 SNPs in the “missing” set, the function has performed a forward step-wise regression on the 50 nearest SNPs in the “present” set, stopping each search either when the R^2 for prediction exceeds 0.95, or after including 4 SNPs in the regression, or until R^2 is not improved by at least 0.05. The figure 50 is the default value of the `try` argument of the function, while the values 0.95, 4 and 0.05 together make up the default value of the `stopping` argument. After the predictor, or “tag” SNPs have been chosen, the haplotypes of the target SNP plus tags was phased and haplotype frequencies calculated using the EM algorithm. These frequencies were then stored in the `rules` object.²

A short listing of the first 10 rules follows:

```
> rules[1:10]
```

```
rs7909677 ~ rs2496276+rs2288680+rs7475011+rs1616133 (MAF = 0.05583756, R-squared = 0.6)
rs12773042 ~ rs2496276+rs2288680+rs4881551+rs4880750 (MAF = 0.04822335, R-squared = 0.9)
rs11253563 ~ rs7093061+rs9419496+rs4390277+rs4881551 (MAF = 0.2309645, R-squared = 0.9)
rs4881552 ~ rs7475011+rs4881551+rs7093061+rs9419496 (MAF = 0.4107143, R-squared = 0.97)
rs10904596 ~ rs7093061+rs9419496+rs7081782+rs4881551 (MAF = 0.2233503, R-squared = 0.9)
rs4880781 ~ rs4880983+rs4881551+rs10736957+rs4390277 (MAF = 0.1700508, R-squared = 0.9)
rs7910845 ~ rs3132006 (MAF = 0.08585859, R-squared = 1)
rs6560730 ~ rs9419496 (MAF = 0.3115578, R-squared = 0.9644185)
rs9329280 ~ rs2436024 (MAF = 0.08673469, R-squared = 0.968183)
rs4880517 ~ rs9419498 (MAF = 0.07286432, R-squared = 0.9637778)
```

The rules are also selectable by SNP name for detailed examination:

```
> rules[c('rs11253563', 'rs2379080')]
```

```
rs11253563 ~ rs7093061+rs9419496+rs4390277+rs4881551 (MAF = 0.2309645, R-squared = 0.9)
rs2379080 ~ rs4880983+rs2288680+rs2379078+rs10794885 (MAF = 0.1809045, R-squared = 0.9)
```

Rules are shown with a + symbol separating predictor SNPs. (It is important to know which SNPs were used for each imputation when checking imputed test results for artifacts.)

A summary table of all the 14,251 rules is generated by


```
> summary(rules)
```

²For imputation from small samples, some smoothing of these haplotype frequencies would be advantageous and some ability to do this has been included. The `use.haps` argument to `snp.imputation` controls this. But invoking this option slows down the algorithm and it is not advised other than for very small sample sizes.

| R-squared | SNPs used | | | | |
|-------------|-----------|--------|--------|--------|------|
| | 1 tags | 2 tags | 3 tags | 4 tags | <NA> |
| [0,0.1) | 98 | 38 | 3 | 0 | 0 |
| [0.1,0.2) | 0 | 87 | 38 | 4 | 0 |
| [0.2,0.3) | 0 | 45 | 80 | 38 | 0 |
| [0.3,0.4) | 0 | 34 | 79 | 101 | 0 |
| [0.4,0.5) | 0 | 30 | 78 | 219 | 0 |
| [0.5,0.6) | 0 | 31 | 79 | 336 | 0 |
| [0.6,0.7) | 0 | 37 | 86 | 498 | 0 |
| [0.7,0.8) | 0 | 110 | 137 | 752 | 0 |
| [0.8,0.9) | 0 | 247 | 304 | 1354 | 0 |
| [0.9,0.95) | 297 | 337 | 362 | 1141 | 0 |
| [0.95,0.99) | 2168 | 397 | 356 | 1053 | 0 |
| [0.99,1] | 2299 | 139 | 166 | 414 | 0 |
| <NA> | 0 | 0 | 0 | 0 | 179 |

Columns represent the number of tag SNPs while rows represent grouping on R^2 . The last column (headed <NA>) represents SNPs for which an imputation rule could not be computed, either because they were monomorphic or because there was insufficient data (as determined by the `minA` optional argument in the call to `snp.imputation`). The same information may be displayed graphically by

```
> plot(rules)
```



Carrying out the association tests

The association tests for imputed SNPs can be carried out using the function `single.snp.tests`.

```
> imp <- single.snp.tests(cc, stratum, data=subject.support,
+                          snp.data=target, rules=rules)
```


Using the observed data in the matrix `target` and the set of imputation rules stored in `rules`, the above command imputes each of the imputed SNPs, carries out 1- and 2-df

single locus tests for association, returns the results in the object `imp`. To see how successful imputation has been, we can carry out the same tests using the *true* data in `missing`:

```
> obs <- single.snp.tests(cc, stratum, data=subject.support, snp.data=lost)
```


The next commands extract the p -values for the 1-df tests, using both the imputed and the true “missing” data, and plot one against the other (using the `hexbin` plotting package for clarity):

```
> logP.imp <- -log10(p.value(imp, df=1))
> logP.obs <- -log10(p.value(obs, df=1))
> hb <- hexbin(logP.obs, logP.imp, xbin=50)
> sp <- plot(hb)
> hexVP.abline(sp$plot.vp, 0, 1, col="black")
```




As might be expected, the agreement is rather better if we only compare the results for SNPs that can be computed with high R^2 . The R^2 value is extracted from the `rules` object, using the function `imputation.r2` and used to select a subset of rules:

```
> use <- imputation.r2(rules)>0.9
> hb <- hexbin(logP.obs[use], logP.imp[use], xbin=50)
> sp <- plot(hb)
> hexVP.abline(sp$plot.vp, 0, 1, col="black")
```



Similarly, the function `imputation.maf` can be used to extract the minor allele frequencies of the imputed SNP from the `rules` object. Note that there is a tendency for SNPs with a high minor allele frequency to be imputed rather more successfully:

```
> hb <- hexbin(imputation.maf(rules), imputation.r2(rules), xbin=50)
> sp <- plot(hb)
```

The function `snp.rhs.glm` also allows testing imputed SNPs. In its simplest form, it can be used to calculate essentially the same tests as carried out with `single.snp.tests`³ (although, being a more flexible function, this will run somewhat slower). The next commands recalculate the 1 df tests for the imputed SNPs using `snp.rhs.tests`, and plot the results against those obtained when values are observed.


```
> imp2 <- snp.rhs.tests(cc~strata(stratum), family="binomial",
+                       data=subject.support, snp.data=target, rules=rules)
> logP.imp2 <- -log10(p.value(imp2))
```

³There is a small discrepancy, of the order of $(N - 1) : N$.

```

> hb <- hexbin(logP.obs, logP.imp2, xbin=50)
> sp <- plot(hb)
> hexVP.abline(sp$plot.vp, 0, 1, col="black")

```



Storing imputed genotypes

In the previous two sections we have seen how to (a) generate imputation rules and, (b) carry out tests on SNPs imputed according to these rules, but without storing the imputed genotypes. It is also possible to store imputed SNPs in an object of class `SnpMatrix` (or


`XSnpmatrix`). The posterior probabilities of assignment of each individual to the three possible genotypes are stored within a one byte variable, although obviously not to full accuracy.

The following command imputes the “missing” SNPs using the “target” dataset and stores the imputed values in an object of class `SnpMatrix`:

```
> imputed <- impute.snps(rules, target, as.numeric=FALSE)
```


(If `as.numeric` were set to `TRUE`, the default, the resulting object would be a simple numeric matrix containing posterior expectations of the 0, 1, 2 genotype.) A nice graphical description of how `snpStats` stores uncertain genotypes is provided by the function `plotUncertainty`. This plots the frequency of the stored posterior probabilities on an equilateral triangle. The posterior probabilities are represented by the perpendicular distances from each side, the vertices of the triangle corresponding to certain assignments. Thus, the SNP `rs4880568` is accurately imputed ($R^2 = 0.94$)

```
> plotUncertainty(imputed[, "rs4880568"])
```



while rs2050968 is rather less so ($R^2 = 0.77$)

```
> plotUncertainty(imputed[, "rs2050968"])
```



Tests can be carried out on these uncertainly assigned genotypes. For example

```
> imp3 <- single.snp.tests(cc, stratum, data=subject.support,
+                           snp.data=imputed, uncertain=TRUE)
```

The `uncertain=TRUE` argument ensures that uncertainly assigned genotypes are used in the computations. This should yield nearly the same result as before. For the first five SNPs we have

```
> imp3[1:5]
```

| | N | Chi.squared.1.df | Chi.squared.2.df | P.1df | P.2df |
|-----------|-----|------------------|------------------|------------|-----------|
| rs7909677 | 800 | 0.32501973 | 0.5810182 | 0.56860647 | 0.7478827 |

```
rs12773042 800      0.25120021      1.0717618 0.61623124 0.5851536
rs11253563 800      2.73636355      3.5829634 0.09808775 0.1667130
rs4881552  800      0.07183608      0.2371646 0.78868197 0.8881787
rs10904596 800      2.94350804      3.4098248 0.08622329 0.1817883
```

```
> imp[1:5]
```

| | N | N.r2 | Chi.squared.1.df | Chi.squared.2.df | P.1df | P.2df |
|------------|-----|----------|------------------|------------------|------------|-----------|
| rs7909677 | 800 | 521.4821 | 0.40434993 | 0.7280041 | 0.52485128 | 0.6948898 |
| rs12773042 | 800 | 707.7510 | 0.23991679 | 1.0018354 | 0.62426622 | 0.6059743 |
| rs11253563 | 800 | 774.5257 | 2.80117306 | 3.7008386 | 0.09419537 | 0.1571712 |
| rs4881552 | 800 | 782.4351 | 0.08976712 | 0.2341699 | 0.76447342 | 0.8895096 |
| rs10904596 | 800 | 744.8878 | 3.07814923 | 3.5827782 | 0.07935077 | 0.1667284 |


There are small discrepancies due to the genotype assignment probabilities not being stored to full accuracy. However these should have little effect on power of the tests and no effect on the type 1 error rate.

Note that the ability of `snpStats` to store imputed genotypes in this way allows alternative programs to be used to generate the imputed genotypes. For example, the file “`mach1.out.mlprob.gz`” (which is stored in the `extdata` sub-directory of the `snpStats` package) contains imputed SNPs generated by the MACH program, using the `-mle` and `-mldetails` options. In the following commands, we find the full path to this file, read it, and inspect one the imputed SNP in column 50:

```
> path <- system.file("extdata/mach1.out.mlprob.gz", package="snpStats")
> mach <- read.mach(path)
```

```
Reading MACH data from file /tmp/RtmpZ0GdXZ/Rinst56055af7de0/snpStats/extdata/mach1.out.
Reading SnpMatrix with 500 rows and 178 columns
```

```
> plotUncertainty(mach[,50])
```



Meta-analysis

As stated at the beginning of this document, one of the main reasons that we need imputation is to perform meta-analyses which bring together data from genome-wide studies which use different platforms. The `snpStats` package includes a number of tools to facilitate this. All the tests implemented in `snpStats` are “score” tests. In the 1 df case we calculate a score defined by the first derivative of the log likelihood function with respect to the association parameter of interest at the parameter value corresponding to the null hypothesis of no

association. Denote this by U . We also calculate an estimate of its variance, also under the null hypothesis — V say. Then U^2/V provides the chi-squared test on 1 df. This procedure extends easily to meta-analysis; given two independent studies of the same hypothesis, we simply add together the two values of U and the two values of V , and then calculate U^2/V as before. These ideas also extend naturally to tests of several parameters (2 or more df tests).

In `snpStats`, the statistical testing functions can be called with the option `score=TRUE`, causing an extended object to be saved. The extended object contains the U and V values, thus allowing later combination of the evidence from different studies. We shall first see what sort of object we have calculated previously using `single.snp.tests` *without* the `score=TRUE` argument.

```
> class(imp)

[1] "SingleSnpTests"
attr(,"package")
[1] "snpStats"
```

This object contains the imputed SNP tests in our target set. However, these SNPs were observed in our training set, so we can test them. We will also recalculate the imputed tests. In both cases we will save the score information:

```
> obs <- single.snp.tests(cc, stratum, data=subject.support, snp.data=missing,
+                          score=TRUE)
> imp <- single.snp.tests(cc, stratum, data=subject.support,
+                          snp.data=target, rules=rules, score=TRUE)
```

The extended objects have been returned:

```
> class(obs)

[1] "SingleSnpTestsScore"
attr(,"package")
[1] "snpStats"

> class(imp)

[1] "SingleSnpTestsScore"
attr(,"package")
[1] "snpStats"
```

These extended objects behave in the same way as the original objects, so that the same functions can be used to extract chi-squared values, p -values etc., but several additional functions, or methods, are now available. Chief amongst these is `pool`, which combines evidence across independent studies as described at the beginning of this section. Although `obs` and `imp` are *not* from independent studies, so that the resulting test would not be valid, we can use them to demonstrate this:


```

> both <- pool(obs, imp)
> class(both)

[1] "SingleSnpTests"
attr(,"package")
[1] "snpStats"

> both[1:5]

      N      N.r2 Chi.squared.1.df Chi.squared.2.df      P.1df      P.2df
rs7909677  997 718.4821      0.1659395      0.2951012 0.6837460 0.8628188
rs12773042  997 904.7510      0.1754096      0.9817742 0.6753482 0.6120832
rs11253563  997 971.5257      1.4415384      2.0224763 0.2298906 0.3637683
rs4881552   996 978.4351      0.3159415      0.4064953 0.5740569 0.8160761
rs10904596  997 941.8878      1.7552949      1.9367421 0.1852125 0.3797010

```

Note that if we wished at some later stage to combine the results in `both` with a further study, we would also need to specify `score=TRUE` in the call to `pool`:

```

> both <- pool(obs, imp, score=TRUE)
> class(both)

[1] "SingleSnpTestsScore"
attr(,"package")
[1] "snpStats"

```

Another reason to save the score statistics is that this allows us to investigate the *direction* of findings. These can be extracted from the extended objects using the function `effect.sign`. For example, this command tabulates the signs of the associations in `obs`:

```

> table(effect.sign(obs))

  -1    0    1
7065  38 7148

```

In this table, -1 corresponds to tests in which effect sizes were negative (corresponding to an odds ratio less than one), while +1 indicates positive effect sizes (odds ratio greater than one). Zero sign indicates that the effect was NA (for example because the SNP was monomorphic). Reversal of sign can be the explanation of a puzzling phenomenon when two studies give significant results individually, but no significant association when pooled. Although it is not impossible that such results are genuine, a more usual explanation is that the two alleles have been coded differently in the two studies: allele 1 in the first study is allele 2 in the second study and vice versa. To allow for this, `snpStats` provides the `switch.alleles` function, which reverses the coding of specified SNPs. It can be applied to `SnpMatrix` objects but, because allele switches are often discovered quite late on in the

analysis and recoding the original data matrices could have unforeseen consequences, the `switch.alleles` function can also be applied to the extended test output objects. This modifies the saved scores *as if* the allele coding had been switched in the original data. The use of this is demonstrated below.

```
> effect.sign(obs)[1:6]
```

```
rs7909677 rs12773042 rs11253563 rs4881552 rs10904596 rs4880781
      -1           1           -1           -1           1           1
```

```
> sw.obs <- switch.alleles(obs, 1:3)
```

```
> class(sw.obs)
```

```
[1] "SingleSnpTestsScore"
```

```
attr(,"package")
```

```
[1] "snpStats"
```

```
> effect.sign(sw.obs)[1:6]
```

```
rs7909677 rs12773042 rs11253563 rs4881552 rs10904596 rs4880781
      1           -1           1           -1           1           1
```