

LEA: An R Package for Landscape and Ecological Association Studies

Eric Frichot and Olivier François
Université Grenoble-Alpes,
Centre National de la Recherche Scientifique,
TIMC-IMAG UMR 5525, Grenoble, 38042, France.

Contents

1 Overview	1
2 Introduction	2
2.1 Input files	3
3 Analysis of population structure and imputation of missing data	4
3.1 Principal Component Analysis	4
3.2 Inference of individual admixture coefficients using <code>snmf</code> . . .	5
3.3 Population differentiation tests using <code>snmf()</code>	6
3.4 Missing genotype imputation using <code>snmf</code>	6
4 Ecological association tests using <code>lfmm</code>	9
5 Ecological association tests using <code>lfmm2</code>	13
6 Predictive Ecological Genomics: Genetic offsets	14

1 Overview

LEA is an R package dedicated to landscape genomics and ecological association tests (Frichot and François, 2015; Gain and François, 2021). LEA performs analyses of population structure and imputation of missing genotypes. It also runs genome scans for selection, tests association of genotypes with environmental variables, and computes predictive values of genetic offsets based on new or future environments. The package includes factor methods for estimating ancestry coefficients from large genotypic matrices and for evaluating the number of ancestral populations (`snmf`, `pca`). It implements

latent factor mixed models for identifying genetic polymorphisms that exhibit high correlation with environmental gradients (lfmm, lfmm2). LEA is mainly based on optimized programs that can scale with the dimension of large data sets.

2 Introduction

The goal of this tutorial is to give an overview of the main functionalities of the R package LEA. It will show the main steps of analysis, including 1) analysing population structure and preparing a genotypic matrix for genomewide association studies, 2) fitting GWAS latent factor mixed models to the data and extracting candidate regions of interest.

As some functions may take a few hours to analyse very large data sets, output files are written into text files that can be read by LEA after each batch of runs (called a 'project'). We advise creating working directories containing genotypic data and covariables when starting LEA. Note that two files with the same names but a different extension are assumed to contain the same data in distinct formats.

```
# creation of a directory for LEA analyses  
dir.create("LEA_analyses")  
# set the created directory as the working directory  
setwd("LEA_analyses")
```

This tutorial is based on a small dataset consisting of 400 SNPs genotyped for 50 diploid individuals. The last 50 SNPs are correlated with an environmental variable, and represent the target loci for an association analysis. Similar artificial data were analyzed in the computer note introducing the R package LEA (Frichot and François, 2015).

```
library(LEA)  
# Creation a the genotypic file: "genotypes.lfmm"  
# The data include 400 SNPs for 50 individuals.  
data("tutorial")  
# Write genotypes in the lfmm format  
write.lfmm(tutorial.R, "genotypes.lfmm")  
# Write genotypes in the geno format  
write.geno(tutorial.R, "genotypes.geno")  
# creation of an environment gradient file: gradient.env.  
# The .env file contains a single ecological variable  
# for each individual.  
write.env(tutorial.C, "gradients.env")
```

Note that the **LEA** package is to be able to handle very large population genetic data sets. Genomic data are processed using fast C codes wrapped into the R code. Most **LEA** functions use character strings containing paths to input files as arguments.

2.1 Input files

The R package **LEA** can handle several input file formats for genotypic matrices. More specifically, the package uses the **lfmm** and **geno** formats, and provides functions to convert from other formats such as **ped**, **vcf**, and **ancestrymap** formats. The program **VCFTOOLS** can be very useful in providing one of those format (**ped** is the closest to an **lfmm** matrix).

The **lfmm** and **geno** formats can also be used for coding multiallelic marker data (eg, microsatellites). For multiallelic marker data, the conversion function **struct2geno()** converts files in the **STRUCTURE** format in the **geno** or **lfmm** formats. **LEA** can also process allele frequency data if they are encoded in the **lfmm** format. In that case, the **lfmm()** and **lfmm2()** functions will use allele counts for populations in its model.

Ecological predictors or phenotypic traits must be formatted in the **env** format. This format corresponds to a matrix where each variable is represented as a column (Frichot et al., 2013). It uses the **.env** extension.

When using ecological data, we often need to decide which variables should be used among a large number of ecological indicators (e.g., climatic variables). Here, we suggest that users summarize their data using linear combinations of those indicators. Considering principal component analysis and using the first principal components as proxies for ecological gradients linked to selective forces can be useful in this context.

The **LEA** package can handle missing data in population structure analyses. In association analyses, missing genotypes must be replaced by imputed values using a missing data imputation method. We encourage users to remove their missing data by using the function **impute()**, which is based on population structure analysis and nonnegative matrix factorization (see next section). Note that specialized genotype imputation programs such as **BEAGLE**, **IMPUTE2** or **MENDEL-IMPUTE** could provide better imputation results than **LEA**, in particular for model species with a published reference genome. Filtering out rare variants – retaining minor allele frequency greater than 5 percent –, and pruning regions in strong LD may also result in better analyses with **LEA**.

3 Analysis of population structure and imputation of missing data

The R package LEA implements two classical approaches for the estimation of population genetic structure: principal component analysis (pca) and admixture analysis (Patterson et al., 2006; Pritchard et al., 2000) using sparse nonnegative matrix factorization (snmf). The algorithms programmed in LEA are improved versions of pca and admixture analysis, that are able to process large genotypic matrices efficiently.

3.1 Principal Component Analysis

The LEA function `pca()` computes the scores of a PCA for a genotypic matrix, and returns a screeplot for the eigenvalues of the sample covariance matrix. Using the function `pca()`, an object of class `pcaProject` is created. This object contains a path to the files storing eigenvectors, eigenvalues and projections.

```
# run of pca  
# Available options, K (the number of PCs),  
# center and scale.  
# Create files: genotypes.eigenvalues - eigenvalues,  
# genotypes.eigenvectors - eigenvectors,  
# genotypes.sdev - standard deviations,  
# genotypes.projections - projections,  
# Create a pcaProject object: pc.  
pc = pca("genotypes.lfmm", scale = TRUE)
```

The number of "significant" components can be evaluated using graphical methods based on the screeplot (Figure 1). The knee in the screeplot indicates that there are around $K = 4$ major components in the data (≈ 5 genetic clusters). Following (Patterson et al., 2006), the `tracy.widom()` function computes Tracy-Widom tests for each eigenvalue as follows.

```
# Perform Tracy-Widom tests on all eigenvalues.  
# create file: tuto.tracyWidom - tracy-widom test information.  
tw = tracy.widom(pc)
```

```
# display p-values for the Tracy-Widom tests (first 5 pcs).  
tw$pvalues[1:5]
```

```
[1] 8.000e-09 8.000e-09 8.000e-09 1.503e-04 3.152e-02
```

```
# plot the percentage of variance explained by each component
plot(tw$percentage)
```

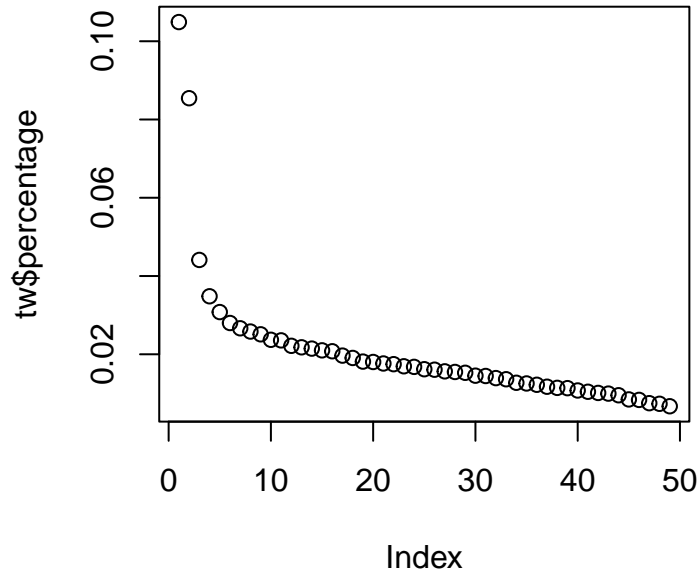


Figure 1: Screeplot for the percentage of variance explained by each component in a PCA of the genetic data. The knee at $K = 4$ indicates that there are 5 major genetic clusters in the data.

3.2 Inference of individual admixture coefficients using `snmf`

The package `LEA` includes the R function `snmf()` that estimates individual admixture coefficients from the genotypic matrix. The function provides results very close to Bayesian clustering programs such as `STRUCTURE` (Pritchard et al., 2000; François and Durand, 2010). Assuming K ancestral populations, the R function `snmf()` provides least-squares estimates of ancestry proportions rather than maximum likelihood estimates (Frichot et al., 2014).

```
# main options
# K = number of ancestral populations
# entropy = TRUE: computes the cross-entropy criterion,
# CPU = 4 the number of CPUs.
project = NULL
```

```
project = snmf("genotypes.geno",
               K = 1:10,
               entropy = TRUE,
               repetitions = 10,
               project = "new")
```

The `snmf()` function computes an entropy criterion that evaluates the quality of fit of the statistical model to the data using a cross-validation technique (Figure 2). The entropy criterion can help choosing the number of ancestral populations that best explains the genotypic data (Alexander and Lange, 2011; Frichot et al., 2014). Here we have a clear minimum at $K = 4$, suggesting 4 genetic clusters in the data. Often, the plot shows a less clear pattern, and choosing the "knee" point is a generally good approach. The number of ancestral populations is closely linked to the number of principal components that explain variation in the genomic data. Both numbers can help determining the number of latent factors when correcting for confounding effects due to population structure in ecological association tests with `lfmm()` and `lfmm2()`.

The next step is to display a barplot for the Q -matrix. In Figure 3, the `Q()` function of `LEA` is called and the output Q -matrix is converted into a `Qmatrix` object. The conversion of the Q -matrix as a `Qmatrix` object is also useful for running improved graphical functions from other packages such as `tess3r` (Caye et al., 2016, 2018).

3.3 Population differentiation tests using `snmf()`

The most common approaches to detecting outlier loci from a genomic background have focused on extreme values of the fixation index, F_{st} , across loci. The `snmf()` function can compute fixation indices when the population is genetically continuous, when predefining subpopulations is difficult, and in the presence of admixed individuals in the sample (Martins et al., 2016). In the `snmf` approach, population differentiation statistics are computed from ancestry coefficients obtained from an `snmf` object, and p -values are returned for all loci. Figure 4 is an example of outlier analysis with `snmf()`.

3.4 Missing genotype imputation using `snmf`

Missing genotypes are critical to genome-wide association studies. Before running an association study, an important step is to replace the missing data, represented as '9' in the `geno` and `lfmm` files, by better values. To provide an example of missing data imputation, let's start by removing 100 genotypes from the original data. The resulting matrix is saved in the file `genotypeM.geno`.

```
# plot cross-entropy criterion for all runs in the snmf project
plot(project, col = "blue", pch = 19, cex = 1.2)
```

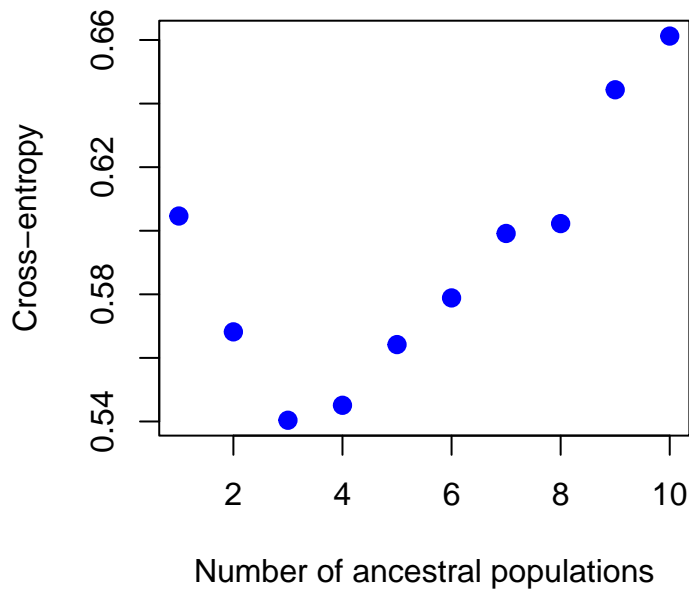


Figure 2: Value of the cross-entropy criterion as a function of the number of populations in **snmf**.

```
# creation of a genotypic matrix with missing genotypes
dat = as.numeric(tutorial.R)
dat[sample(1:length(dat), 100)] <- 9
dat <- matrix(dat, nrow = 50, ncol = 400)
write.lfmm(dat, "genoM.lfmm")

## [1] "genoM.lfmm"
```

Next, the function **snmf()** can be run on the data with missing genotypes as follows. The completion of the genotypic matrix is based on estimated ancestry coefficients and ancestral genotype frequencies.

```
project.missing = snmf("genoM.lfmm", K = 4,
  entropy = TRUE, repetitions = 10,
  project = "new")
```

```

# select the best run for K = 4
best = which.min(cross.entropy(project, K = 4))
my.colors <- c("tomato", "lightblue",
              "olivedrab", "gold")
barchart(project, K = 4, run = best,
          border = NA, space = 0,
          col = my.colors,
          xlab = "Individuals",
          ylab = "Ancestry proportions",
          main = "Ancestry matrix") -> bp
axis(1, at = 1:length(bp$order),
     labels = bp$order, las=1,
     cex.axis = .4)

```

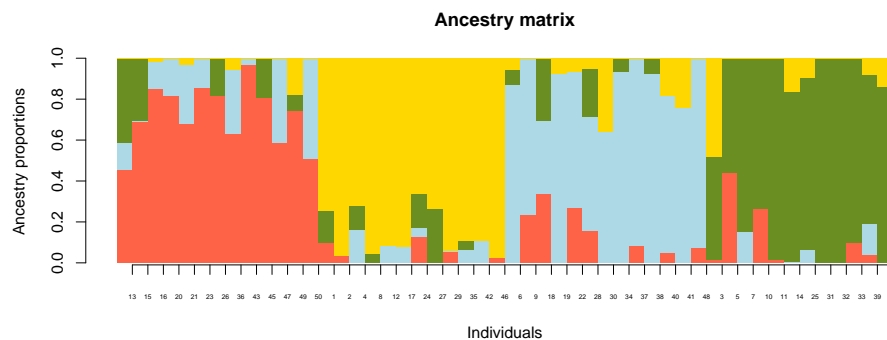


Figure 3: Ancestry coefficients obtained from `snmf()`.

The `snmf` project data can be used to impute the missing data as follows.

```

# select the run with the lowest cross-entropy value
best = which.min(cross.entropy(project.missing, K = 4))

# Impute the missing genotypes
impute(project.missing, "genoM.lfmm",
        method = 'mode', K = 4, run = best)

## Missing genotype imputation for K = 4
## Missing genotype imputation for run = 4
## Results are written in the file: genoM.lfmm_imputed.lfmm

# Proportion of correct imputation results
dat.imp = read.lfmm("genoM.lfmm_imputed.lfmm")
mean( tutorial.R[dat == 9] == dat.imp[dat == 9] )

```



```
# Population differentiation tests
p = snmf.pvalues(project,
  entropy = TRUE,
  ploidy = 2,
  K = 4)
pvalues = p$pvalues
par(mfrow = c(2,1))
hist(pvalues, col = "orange")
plot(-log10(pvalues), pch = 19, col = "blue", cex = .5)
```

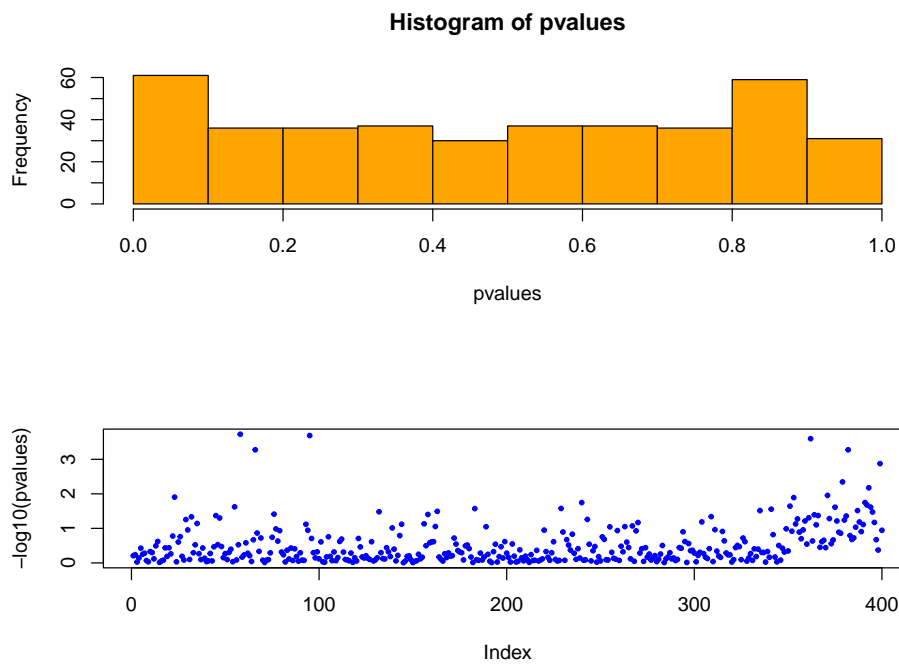


Figure 4: P -values for population differentiation tests with `snmf()`.

```
## [1] 0.75
```

The results are saved in an output file with the string "imputed" in its suffix name.

4 Ecological association tests using `lfmm`

The R package `LEA` performs genome-wide association analysis based on latent factor mixed models using the `lfmm()` function (Frichot et al., 2013).

To recall the model, let G denote the genotypic matrix, storing allele frequencies for each individual at each locus, and let X denote a set of d ecological predictors or phenotypic traits. LFMMs consider the genotypic matrix entries as response variables in a latent factor regression model

$$G_{i\ell} = \mu_\ell + \beta_\ell^T X_i + U_i^T V_\ell + \epsilon_{i\ell}, \quad (1)$$

where μ_ℓ is a locus specific effect, β_ℓ is a d -dimensional vector of regression coefficients, U_i contains K latent factors, and V_ℓ contains their corresponding loadings (i stands for an individual and ℓ for a locus). The residual terms, $\epsilon_{i\ell}$, are statistically independent Gaussian variables with mean zero and variance σ^2 .

In latent factor models, association between predictors and allele frequencies can be tested while estimating unobserved latent factors that model confounding effects. In principle, the latent factors include levels of population structure due to shared demographic history or background genetic variation. After correction for confounding effects, association between allele frequencies and an ecological predictor at a particular locus is often interpreted as a signature of natural selection.

Running LFMM. The `lfmm()` program is based on a stochastic algorithm (MCMC) which does not provide exact results. See the section on the `lfmm2()` for an alternative method which provides exact results under simplified assumptions. We recommend using large number of cycles (e.g., `-i 6000`) and the burnin period should set at least to one-half of the total number of cycles (`-b 3000`). We have noticed that the program results are sensitive to the run-length parameter when data sets have relatively small sizes (e.g., a few hundreds of individuals, a few thousands of loci). We recommend increasing the burnin period and the total number of cycles in this situation.

```
# main options:
# K the number of latent factors
# Runs with K = 6 and 5 repetitions.
project = NULL
project = lfmm("genotypes.lfmm",
               "gradients.env",
               K = 6,
               repetitions = 5,
               project = "new")

## lfmm uses a very naive imputation method which has low power
when genotypes are missing: See impute() for a better imputation
method.
```

```
## Note that lfmm has an improved estimation algorithm implemented  
in lfmm2, which should be the preferred option.
```

Deciding the number of latent factors. Deciding an appropriate value for the number of latent factors in the `lfmm()` call can be based on the analysis of histograms of test significance values. Ideally, histograms should be flat, with a peak close to zero.

Since the objective is to control the false discovery rate (FDR) while keeping reasonable power to reject the null hypothesis, we recommend using several runs for each value of K and combine p -values (use 5 to 10 runs, see our script below). Choosing values of K for which the histograms show their correct shape warrants that the FDR can be controlled efficiently.

Testing all K values in a large range, say from 1 to 20, is generally useless. A careful analysis of population structure and estimates of the number of ancestral populations contributing to the genetic data indicates the range of values to be explored. For example, if the `snmf()` command estimates 4 ancestral populations, then running `lfmm()` for $K = 3 - 6$ often provides good results.

Combining z -scores obtained from multiple runs. We use the Fisher-Stouffer method to combine z -scores from multiple runs. In practice, we found that using the median z -scores of 5-10 runs and re-adjusting the p -values afterwards can increase the power of `lfmm` tests. This procedure is implemented in LEA function `lfmm.pvalues()`.

```
# compute adjusted p-values  
p = lfmm.pvalues(project, K = 6)  
pvalues = p$pvalues
```

The results displayed in Figure 5 show that the null-hypothesis is correctly calibrated. The loci exhibiting significant associations are found at the right on the Manhattan plot.

To adjust p -values for multiple testing issues, we use the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). We set the expected levels of FDR to $q = 5\%$, 10% , 15% and 20% respectively. The lists of candidate loci are given by the following script. Since we the ground truth is known for the simulated data, we can compare the expected FDR levels to their observed levels, and compute the power (TPR) of the test.

```
for (alpha in c(.05,.1,.15,.2)) {  
  # expected FDR  
  print(paste("Expected FDR:", alpha))  
}
```

```
# GWAS significance test
par(mfrow = c(2,1))
hist(pvalues, col = "lightblue")
plot(-log10(pvalues), pch = 19, col = "blue", cex = .7)
```

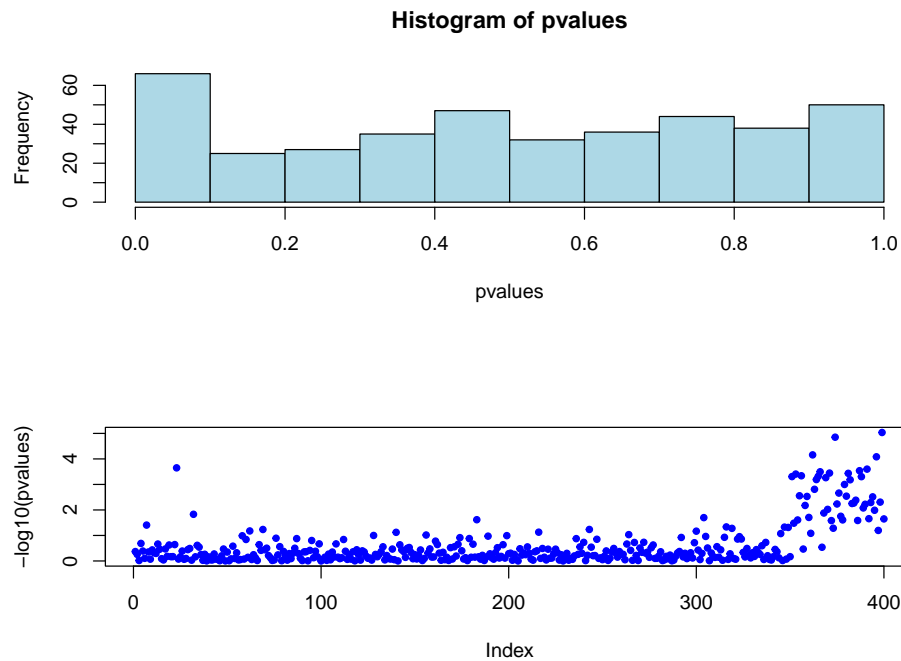


Figure 5: P -values for LFMM tests. The loci showing significant associations are at the right on the Manhattan plot.

```
L = length(pvalues)

# return a list of candidates with expected FDR alpha.
# Benjamini-Hochberg's algorithm:
w = which(sort(pvalues) < alpha * (1:L) / L)
candidates = order(pvalues)[w]

# estimated FDR and True Positive Rate
Lc = length(candidates)
estimated.FDR = sum(candidates <= 350)/Lc
print(paste("Observed FDR:",
            round(estimated.FDR, digits = 2)))
estimated.TPR = sum(candidates > 350)/50
```

```

    print(paste("Estimated TPR:",
                round(estimated.TPR, digits = 2)))
}

## [1] "Expected FDR: 0.05"
## [1] "Observed FDR: 0.04"
## [1] "Estimated TPR: 0.44"
## [1] "Expected FDR: 0.1"
## [1] "Observed FDR: 0.03"
## [1] "Estimated TPR: 0.66"
## [1] "Expected FDR: 0.15"
## [1] "Observed FDR: 0.03"
## [1] "Estimated TPR: 0.72"
## [1] "Expected FDR: 0.2"
## [1] "Observed FDR: 0.07"
## [1] "Estimated TPR: 0.76"

```

5 Ecological association tests using lfmm2

As an efficient alternative to the MCMC algorithm implemented in the `lfmm()`, genome-wide association analysis can be performed in LEA by using the `lfmm2()` command. This function allows estimating K latent factors and the effect sizes corresponding to environmental variables in the same statistical model as with `lfmm()`. For `lfmm2()`, the estimation algorithm is based on exact solutions of a least-squares minimization problem (Caye et al., 2019). For large data sets, `lfmm2()` is much faster than the MCMC version (Gain and François, 2021). The function also decouples the estimation of latent factors from the association tests, which allows implementing various types of tests including linear or generalized linear models. Let us consider an example simulated from the LFMM generative model. The simulated genotypic matrix contains $n = 100$ individuals genotyped for $L = 1000$ loci, ten of which are truly associated with an artificial environmental variable, X .

```

# Simulate non-null effect sizes for 10 target loci
# individuals
n = 100
# loci
L = 1000
# Environmental variable
X = as.matrix(rnorm(n))
# effect sizes

```

```
B = rep(0, L)
target = sample(1:L, 10)
B[target] = runif(10, -10, 10)
```

The latent factors, \mathbf{U} , contained in an $n \times 3$ matrix, are random vectors created as follows. Correlations between the factors and environmental predictor, X , are introduced in the simulation model.

```
# Create 3 hidden factors and their loadings
U = t(tcrossprod(as.matrix(c(-1,0.5,1.5)), X)) +
  matrix(rnorm(3*n), ncol = 3)

V <- matrix(rnorm(3*L), ncol = 3)
```

The genotypic matrix, \mathbf{Y} , is simulated according to an approximation of the LFMM generative model. This matrix has dimension $n \times L$.

```
# Simulate a matrix containing haploid genotypes
Y <- tcrossprod(as.matrix(X), B) +
  tcrossprod(U, V) +
  matrix(rnorm(n*L, sd = .5), nrow = n)

Y <- matrix(as.numeric(Y > 0), ncol = L)
```

We fit an LFMM by using $K = 3$ latent factors. This value corresponds to the true value in the model (note that $K = 3$ could be easily recovered from a PCA screeplot).

```
# Fitting an LFMM with K = 3 factors
mod <- lfmm2(input = Y, env = X, K = 3)
```

The `lfmm2()` command generates an object of class `lfmm2Class` which contains estimated factors (`mod2@U`) and loadings (`mod2@V`) for being introduced as correction factors in genome-wide association tests. Note here that LEA is using S4 objects rather than S3 objects. Getting the factors could be useful for implementing customized tests. For example, they could be used for computing a covariance matrix for random effects in a mixed linear model. The `lfmm2.test()` function implements simpler tests such as linear or generalized linear model tests. The latter are much slower. To adjust p -values for multiple testing issues, we can use the Benjamini-Hochberg procedure as with the `lfmm()` tests. The tests and a Manhattan plot can be performed as follows.

6 Predictive Ecological Genomics: Genetic offsets

to be developed, see example in `help(genetic.offset)`

```

# Computing P-values and plotting their minus log10 values
pv <- lfmm2.test(object = mod,
                 input = Y,
                 env = X,
                 linear = TRUE)

plot(-log10(pv$pvalues), col = "grey", cex = .6, pch = 19)
points(target, -log10(pv$pvalues[target]), col = "red")

```

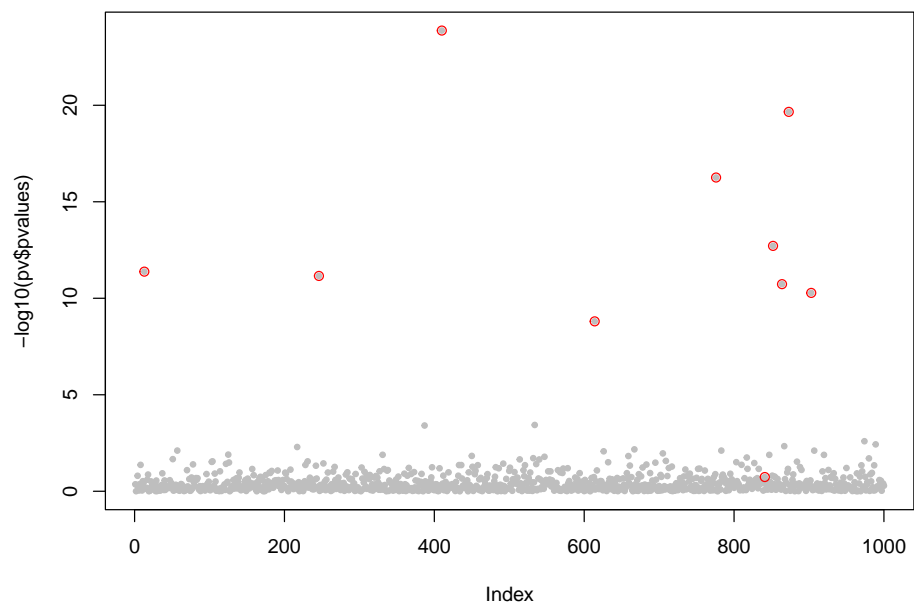


Figure 6: Manhattan plot of $\log_{10} p$ -values for LFMM2 tests. The loci showing real associations are circled in red.

References

- Alexander DH and Lange K. 2011. Enhancements to the ADMIXTURE algorithm for individual ancestry estimation. BMC Bioinformatics 12:246.
- Benjamini Y and Hochberg Y. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. J R Stat Soc B Met. pp. 289–300.
- Caye K, Jay F, Michel O and Francois O. 2016. Fast inference of individual admixture coefficients using geographic data. bioRxiv p. 080291.

- Caye K, Jay F, Michel O and Francois O. 2018. Fast inference of individual admixture coefficients using geographic data. *The Annals of Applied Statistics* 12:586–608.
- Caye K, Jumentier B, Lepeule J and François O. 2019. Lfmm 2: Fast and accurate inference of gene-environment associations in genome-wide studies. *Molecular Biology and Evolution* 36:852–860.
- François O and Durand E. 2010. Spatially explicit bayesian clustering models in population genetics. *Mol Ecol Resour.* 10:773–784.
- Frichot E and François O. 2015. LEA: an R package for Landscape and Ecological Association studies. *Methods in Ecology and Evolution* 6:925–929.
- Frichot E, Mathieu F, Trouillon T, Bouchard G and François O. 2014. Fast and efficient estimation of individual ancestry coefficients. *Genetics* 196:973–983.
- Frichot E, Schoville SD, Bouchard G and François O. 2013. Testing for associations between loci and environmental gradients using latent factor mixed models. *Mol Biol Evol.* 30:1687–1699.
- Gain C and François O. 2021. LEA 3: Factor models in population genetics and ecological genomics with r. *Molecular Ecology Ressources* 21.
- Martins H, Caye K, Luu K, Blum MGB and Francois O. 2016. Identifying outlier loci in admixed and in continuous populations using ancestral population differentiation statistics. *Molecular Ecology* 25:5029–5042.
- Patterson N, Price AL and Reich D 2006. Population structure and eigenanalysis. *PLoS Genet.* 2:20.
- Pritchard JK, Stephens M and Donnelly P. 2000. Inference of population structure using multilocus genotype data. *Genetics* 155:945–959.