High performance computation of landscape genomic models integrating local indices of spatial association

Sylvie Stucki[1], Pablo Orozco-terWengel[2], Michael W. Bruford[2], Licia Colli[3], Charles Masembe[4], Riccardo Negrini[3,5], Pierre Taberlet[5,6], Stéphane Joost[1,*], and the NEXTGEN Consortium[8]

[1]Laboratory of Geographic Information Systems (LASIG), School of Architecture, Civil and Environmental Engineering (ENAC), Ecole Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland
[2]School of Biosciences, Sir Martin Evans Building, Cardiff University, Cardiff, UK
[3]Istituto di Zootecnica and BioDNA - Centro di Ricerca sulla Biodiversità e sul DNA Antico, Università Cattolica del S. Cuore, via E. Parmense 84, 29100 Piacenza, Italy
[4]Makerere University, Institute of Environment and Natural Resources, Molecular Biology Laboratory, Kampala, Uganda
[5]Associazione Italiana Allevatori, Roma, Italy
[6]CNRS, Laboratoire d'Ecologie Alpine, Grenoble, France
[7]Univ. Grenoble Alpes, Laboratoire d'Ecologie Alpine, Grenoble, France
[8]http://nextgen.epfl.ch

May 30, 2014

Abstract

Motivation: The increasing availability of high-throughput datasets requires powerful methods to support the detection of signatures of selection in landscape genomics.
Results: We present an integrated approach to study signatures of local adaptation, providing rapid processing of whole genome data and enabling assessment of spatial association using molecular markers.
Availabilty: Samβada is an open source software written in C++ available at http:lasig.epfl.ch/sambada (under the license GNU GPL 3). Compiled versions are provided for Windows, Linux and MacOS X.
Contact: stephane.joost@epfl.ch, sylvie.stucki@a3.epfl.ch
Supplementary material is available online.

[*]to whom correspondence should be addressed

1 Introduction

The time interval between Mitton *et al.*'s (1977) first attempt to correlate allelic frequencies with environmental variables to look for signatures of selection in ponderosa pine, and Joost *et al.*'s (2007; 2008) application of this concept allowing parallel processing of large numbers of logistic regressions was otherwise marked by little developments. During this period correlative approaches were used in parallel with population genetics outlier-detection methods (e.g. Beaumont and Nichols, 1996; Vitalis *et al.*, 2003; Foll and Gaggiotti, 2008) as cross-validation (e.g. Jones *et al.*, 2013; Henry and Russello, 2013) to detect signatures of selection (see a review in Vitti *et al.*, 2013). However, while such methods are still in vogue (e.g. Colli *et al.*, 2014), there has been a revival in the interest of developing new statistical approaches for the emerging field of landscape genomics (e.g. Coop *et al.*, 2010; Günther and Coop, 2013; Frichot *et al.*, 2013; Guillot *et al.*, 2014). For example, BayEnv (Günther and Coop, 2013) implements a Bayesian method to compute correlations between allele frequencies and ecological variables taking into account differences in sample sizes and shared demographic history. LFMM (Frichot *et al.*, 2013) estimates the influence of population structure on allele frequencies by introducing unobserved variables as latent factors. Finally, SGLMM (Guillot *et al.*, 2014) uses a spatially-explicit computational framework including a random effect to quantify the correlation between genotypes and environmental variables. Yet, important functions are still lacking such as high performance computing capacity to process whole genome data, and the integration of spatial statistics to support a distinction between selection and demographic signals. Here we present the software Samβada, which aims at filling these gaps offering an open source multivariate analysis framework to detect signatures of selection. Samβada's use is illustrated with a case study dedicated to the detection of potentially adaptive loci in 813 *Bos taurus* and *Bos indicus* individuals in Uganda genotyped for $\sim 40,000$ SNP. Lastly, Samβada's performance is described with respect to other state of the art software to detect signatures of selection.

2 Methods

2.1 Samβada's approach

Samβada uses logistic regressions to model the probability of presence of an allelic variant for a polymorphic marker given the environmental conditions of the sampling locations (Joost *et al.*, 2007). Since each of the states of a given character is considered independently (i.e. as binary presence/absence in each sample), Samβadacan handle many types of molecular data (e.g. SNPs, indels, copy number variants and haplotypes), provided the user formats the input. Specifically, biallelic SNPs are recoded as three distinct genotypes. A maximum likelihood approach is used to fit the models (Dobson and Barnett, 2008).

2.1.1 Univariate analysis

In the univariate case, each model for a given genotype is compared to a constant model, where the probability of presence is the same at each location and is equal to the frequency of the genotype. Significance is assessed with both log-likelihood ratio (G) and Wald tests (Dobson and Barnett, 2008). Bonferroni correction is applied for multiple comparisons. In order to avoid numerous computations of p-values, the significance threshold α is converted to a minimum score threshold. G and Wald scores are used to compare models rather than Akaike or Bayesian information criterion in order to automate model selection.

2.2 Multivariate analysis

The model selection procedure is adapted to assess the significance of multivariate models. Both G and Wald tests refer to a null model to build the null hypothesis. The current model can be compared to the constant model (the same as in the univariate analysis) using multivariate χ^2 statistics. While rejecting the null hypothesis in this configuration indicates that at least one parameter is statistically significant, it does not provide information about which parameter(s) is relevant to the model. Therefore model selection is based on simpler models nested in the current one, and parameter significance is determined with either a Wald test applied to each parameter separately (except the constant parameter) or with G tests excluding a parameter at the time. For the latter, if a model A has q parameters, we define the parents of A as the q models with $q-1$ parameters obtained by dropping one parameter from A. For instance, if A models the occurrence of genotype X_i with 3 environmental variables E_2, E_3 and E_5,

$$A = (X_i | E_2, E_3, E_5),$$

then the *parents* of A are the three models

$$(X_i | E_2, E_3), (X_i | E_2, E_5) \text{ and } (X_i | E_3, E_5).$$

The parent of A with the highest log-likelihood is used as the reference model for the significance test. This way, the G score is the smallest possible among all parents, thus if the null hypothesis is rejected, it will also be rejected by comparing A with each of its parents. This method ensures that adding a new parameter leads to a better modelling of the presence of the genotype. The overall procedure for fitting and selecting models for each genotype begins with the computation of the constant model. Univariate models are built and tested against the constant one, followed by testing bivariate models against their parents, and so forth until the user-defined maximum number of parameters is reached.

2.3 Spatial autocorrelation

Beyond detection of selection signatures, Samβada quantifies the level of spatial dependence in the distribution of each genotype. This measure of spatial autocorrelation refers

to similarities or differences among neighbouring individuals that cannot be explained by chance. Assessing whether the geographic location has an effect on allele frequency is especially important in landscape genomics since statistical models assume independence between events. Thus if individuals with similar genotypes tend to concentrate in space, spurious correlations may co-occur with specific values of environmental variables. On the other hand, spatial independence of data strengthens the confidence in the detections.

Samβada measures the global spatial autocorrelation in the whole dataset with Moran's I, as well as the spatial dependency of each point with Local Indicators of Spatial Association (LISA) (Moran, 1950; Anselin, 1995). In practice, LISAs are computed by comparing the value of each point with the mean value of its neighbours as defined by a specific weighting scheme based on a kernel function (see supplementary material). Both a spatially fixed kernel type relying on distance only, and a varying kernel type considering point density can be used. Samβada includes three fixed kernels (moving window, Gaussian and bisquare) and a varying one (nearest neighbours). The sum of LISAs on the whole dataset is proportional to Moran's I (Anselin, 1995). Significance assessment relies on an empirical distribution of the indices. For Moran's I, values (genotype occurrences) are permuted among the locations of individuals of the whole dataset and a pseudo p-value is computed as the proportion of permutations for which I is equal to - or more extreme (higher for a positive Moran's I or lower for a negative Moran's I) - than the observed I. For LISA, the pseudo p-value is separately computed for each point (individual), by keeping the individual of interest's value fixed and permuting its neighbouring points with the rest of the dataset.

3 Samβada implementation

Samβada is a standalone application written in C++. The application was developed using the Scythe Statistical Library (Pemstein *et al.*, 2011) for matrix computation and probability distributions, which was also chosen for its straightforward application programming interface (API). Samβada is distributed under an open source GNU General Public License license in order to ease its use for research and teaching.

3.1 Desktop and High Performance Computing

Two complementary versions of the software were developed: a desktop option-rich program well suited to small to medium-sized datasets, and a High Performance Computing version dedicated to large datasets.

3.1.1 Desktop version (Samβada)

Samβada includes multivariate analysis and spatial autocorrelation. Many options are provided to facilitate the formatting of the data and to customise the analysis. For instance, the significance of models is assessed during the analysis and non-significant associations can be discarded. Moreover models can be sorted according to their scores

before writing the results in order to make it possible to directly be in a position to interpret them. All results presented in this paper were processed with Samβada Desktop.

3.1.2 Parallel computing version (CoreSAM)

When processing large datasets, primary analysis usually focuses on univariate models. Multivariate models and spatial autocorrelation may be considered as a second step, but are too computationally intensive to be applied to the whole dataset. In order to speed-up the process, CoreSAM is a light version of Samβada, written in C, which focuses on univariate analysis. Compared with Samβada, fewer options are available, but the computation is up to seven times faster.

Combining Samβada and CoreSAM, large datasets may be analysed by steps: Univariate logistic models identify candidate loci exhibiting selection signatures. These loci may be then investigated in the light of spatial autocorrelation measures and multivariate models. The former may point out whether the observed correlation is due to similarities between neighbours, while the latter allows including the population structure, if any, in the model in order to assess whether the environmental variable provides supplementary information on the marker frequency when taking the demography into account.

3.2 Modules

Samβada includes several modules that enhance interfacing with other programs.

3.2.1 Geovisualization of spatial statistics

Samβada provides an option to save the spatial autocorrelation results as a shapefile (.shp), a common format for storing vector information in Geographic Information Systems (GIS). This feature relies on the shplib open source library (shape-lib.maptools.org).

3.2.2 Recoding molecular data

Samβada is distributed with a utility for recoding molecular data into binary information. Currently RecodePlink handles ped/map files, a format for SNP data used by PLINK (Purcell *et al.*, 2007).

3.2.3 Supervision

For very large molecular datasets, Samβada provides a module to share workload between computers. "Supervision" splits the input data in several files that can be processed separately, even on heterogeneous computers. Lastly, Supervision merges the results to provide the same output as if the whole dataset had been processed at once.

4 Case study

4.1 Sampling design

This study addressed local adaptation in Ankole and Shorthorn zebu cattle in Uganda. Sampling was designed to cover the whole country, including each eco-geographic region, and to obtain a homogeneous distribution of individuals across the country. A regular grid made of 51 cells of 70 x 70 km was produced to this end. On average, four farms were visited in each cell and four unrelated individuals were selected from each farm, for a total of 917 biological samples retrieved from 202 farms. Recorded information also included the location of the farm, the name of the breed, a picture and morphological information on each individual. These elements were stored in a database accessible through a Web interface, enabling real-time monitoring of the sampling campaign.

4.2 Molecular data

Out of the 917 individuals, 813 samples were genotyped with a medium-density SNP chip (54,609 SNPs, BovineSNP50 BeadChip, Illumina Inc., San Diego, CA). Only markers located on the autosomal chromosomes were considered in the analyses. The dataset was filtered with PLINK (Purcell *et al.*, 2007) with a call rate set to 95% for both individuals and SNPs, and a minimum allele frequency (MAF) set to 1%. The resulting dataset contains 804 samples and 40,019 SNPs.

4.3 Environmental data

The geographical coordinates of the individuals sampled enabled the characterisation of their habitat with the help of the WorldClim dataset containing monthly values of precipitation, minimum, mean and maximum temperature as well as 19 derived variables, at 1km resolution (Hijmans *et al.*, 2005). These environmental variables were originally stored in four tiles (portions of map) which were pasted using the Geospatial Data Abstraction Library (GDAL Development Team, 2013) and a customized Python script. The topography is described by the 90m resolution SRTM3 (Shuttle Radar Topography Mission) digital elevation model (Farr *et al.*, 2007). SAGA GIS (www.saga-gis.org) was used to paste the 36 tiles covering the country and to derive slope and orientation from the altitude. Longitude and latitude were also included as a proxy of population structure. Finally the values of the 72 environmental variables were extracted for each animal using the "Point Sampling Tool" extension (http://hub.qgis.org/projects/pointsamplingtool) in QuantumGIS (www.qgis.org).

Considering all environmental variables in the computation of the multiple logistic regressions would have provided a comprehensive analysis with a low risk of missing detections. Nonetheless some variables are highly correlated; this implies dependence between models and increases the variance of parameters in multivariate analyses. Thus we used the Variance Inflation Factor (VIF) to control for multicollinearity (Dobson and Barnett, 2008). A maximum VIF of 5 corresponds to a coefficient of correlation of 0.9 between pairs of variables. The number of variables was reduced iteratively by randomly

removing one of the two most highly correlated variables until the maximum correlation was lower than the threshold. This procedure led to a set of 23 environmental variables that were used for univariate landscape genomic analysis (table S1). Bivariate models were also processed with Samβada to assess the effect of a combination of predictors, and to take the population structure into account. This information was constituted by the coefficient of membership of individuals to the two main populations of Ugandan cattle. As a single coefficient was sufficient to represent the origin of each individual, a new variable "population structure" was defined as the coefficient of membership of each individual to the population Ankole. This variable was added to the set of 23 variables and the correlation-based variable selection method was reapplied to limit the VIF to 5. On this basis, fifteen environmental variables were considered for Samβada bivariate analysis (table S1).

4.4   Population structure

Population structure was analysed with Admixture (Alexander *et al.*, 2009), which estimates maximum likelihood of individual ancestries from multilocus SNP genotype datasets. This approach assumes that samples descend from a predefined number of ancestor populations that are mixing. Admixture estimates both the fraction of each sample coming from each population and the marker frequencies in these populations. The optimal number of populations is assessed by a *k*-fold cross-validation procedure.

4.5   Alternative methods to detect selection

For comparison purpose, three alternative approaches to Samβada were used to detect signatures of selection in Ugandan cattle data. Two of these are correlative approaches (BayEnv and Latent Factor Mixed Models, Coop *et al.*, 2010; Frichot *et al.*, 2013), while the third is an outlier-detection population genetics approach (Beaumont and Nichols, 1996) included in Arlequin 3.5 (Excoffier and Lischer, 2010).

4.5.1   BayEnv

BayEnv uses a Bayesian approach to detect candidate SNPs under selection while accounting for the inherent correlation in allele frequencies across populations due to shared demographic history (Coop *et al.*, 2010). BayEnv first uses a set of neutral loci to build a null model robust to demographic history, against which an alternative model including an environmental variable is compared. Markers exhibiting the highest Bayes factors are potentially subject to selection. In this study the set of neutral loci was chosen at random among loci identified as neutral by Samβada.

4.5.2   Latent Factor Mixed Models

LFMM is a Bayesian individual-based approach to detect selection in landscape genomics (Frichot *et al.*, 2013). Population structure is added into the model via unobserved variables. Thus the significance of the association between genome and environment can

be assessed while taking into account the effect of the population structure. The number of latent factors (unobserved variables) must be specified for the analysis. Although this number is related to the population structure, it is usually lower than the number of populations (Frichot personal communication).

4.5.3 Outlier approach

Arlequin is a comprehensive software for population genetics analyses (Excoffier and Lischer, 2010) that includes an outlier-based method to detect signatures of selection (Beaumont and Nichols, 1996). This approach assumes an island model, where individuals are sampled in distinct populations that exchange migrants. Each locus is characterised by the fixation index F_{ST} (Wright, 1949) and its heterozygosity. Neutral coalescent simulations are used to estimate the distribution of F_{ST} conditional on heterozygosity. Loci exhibiting extreme values of F_{ST} are candidate targets of selection.

5   Results

5.1   Population structure

The best partition of the dataset consisted of four populations, although the vast majority of the samples were allocated to two clusters (almost 96%) on the basis of the ancestry coefficients (Figure S1). Mapping these coefficients revealed the two clusters (340 and 431 individuals out of 804) occurred in the South-West and North-East of Uganda respectively. Using pictures of sampled individuals, the first cluster was identified as Ankole cattle and the second one as zebu. The remaining two clusters (32 animals) possibly represent introgression from allochthonous gene pools.

5.2   Detection of selection signatures

Four approaches were applied to detect selection signatures. The statistical significance threshold for Samβada, LFMM and Arlequin was set to 1% before applying Bonferroni correction. For BayEnv, model selection was based on the distribution of Bayes Factors (BF) for neutral loci (Coop *et al.*, 2010). Results were analysed separately for each environmental variable and models showing a BF higher than the 1[st] percentile of the neutral distribution were detected as candidate loci. For BayEnv and Arlequin, samples were classified into populations using a threshold of 0.85 for the higher admixture coefficient, leading to three clusters of 162 Ankole cattle, 8 zebus and 10 cattle from the third population; samples from the fourth population were highly admixed and none satisfied the condition.

Using univariate models, Samβada identified 2,499 SNPs (6.2 %) potentially subject to selection, BayEnv 1977 (4.9 %), LFMM 280 (0.7 %) and Arlequin did not identify any loci as significant. The loci detected by Samβada with the highest G scores were compared among methods in table S2. Thirty-six loci were identified by the three cor-

relative methods and three of them were among the most significant models in Samβada (Table S3). These three SNPs occur close to each other in chromosome five.

Samβada's multivariate analysis identified 84 significant bivariate models, corresponding to 29 loci. In Samβada's framework, this means these models provided a significantly more accurate estimation of the genotype's frequency than their univariate parents, while at least one of their parents was also significant. Among those, three models that included the "population structure" variable also had a parent model showing a significant association with this variable. Therefore, although the population structure partly explains the distribution of these genotypes, adding an environmental variable provided a significantly more accurate estimation of their distribution ($p \leq 7.9 \cdot 10^{-10} \Leftrightarrow G \geq 37.8$). These models correspond to three loci that were detected by all correlative approaches.

5.3 Spatial autocorrelation

Global and local indicators of spatial autocorrelation were computed for two genotypes with a weighting scheme based on the 20 nearest neighbours: ARS-BFGL-NGS-113888 (hereon ARS-113) (allele GG), which was detected by Samβada with the highest G score and was also detected by BayEnv, was compared with Hapmap28985-BTA-73836 (hereon HM-28) (allele GG), which was detected by Samβada, BayEnv and LFMM. Logistic models significantly associated both genotypes with isothermality, which measures the stability of temperature during the year. Figure 1 shows local indices of spatial autocorrelation for these 2 genotypes: on the one hand, ARS-113 (GG) was positively autocorrelated for the majority of points and the indicator was significant for half of them. The distribution of this marker shows spatial dependence, non-significant associations were found at the edge of Lake Victoria and in a corridor in the North of the Lake with some occurrences in the West of Uganda. This widespread pattern of spatial autocorrelation could originate from the underlying population structure, since Ankole cattle are more common in the South-West while zebus are more common in the North-East of the country. Thus the correlation detected by logistic regressions between ARS-113 (GG) and environmental variables could be spurious and due to demographic factors. On the other hand, the local indicators of spatial association of HM-28 (GG) showed lower values in general and were only significant in the North-West of Uganda. This particular region also showed the lowest values of isothermality in Uganda, i.e. a high variability of temperatures. The low value of spatial autocorrelation for HM-28 (GG) implies that the distribution of this genotype was mostly independent from the location and this supports a possible adaptive origin of the observed correlation between HM-28 (GG) and isothermality with logistic models. This correlation between HM-28 (GG) and isothermality also appeared with bivariate LISAs, where the presence of the genotype was compared to the mean value of isothermality among neighbouring points (not shown).

9
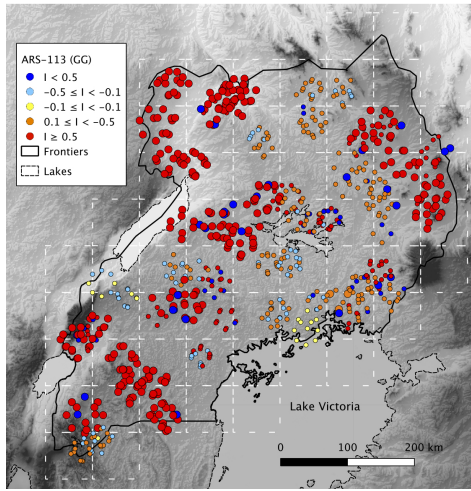
(a) ARS-113 (GG)
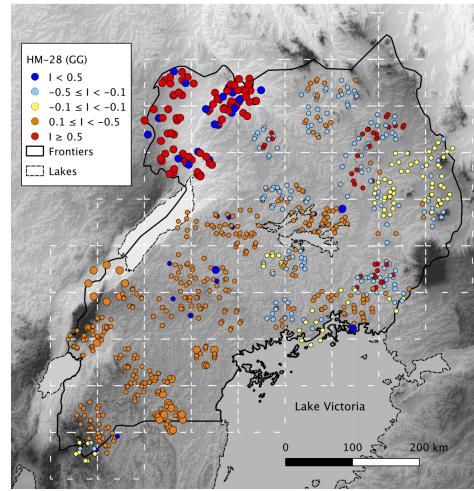
(b) HM-28 (GG)

Figure 1: Local indicators of spatial association of markers ARS-113 (allele GG) and HM-28 (allele GG). The weighting scheme is based on the first 20 nearest neighbours. Red points tend to be similar to their neighbours while blue points differ from them. Yellow points are independent from their neighbourhood. Small points indicate non-significant values ($p > 0.001$). The map in the background represents the relief, the darker the shade, the higher the altitude. Samples coming from the same farm have been spread on a circle around their actual location.

6 Discussion

The key features of Samβada are the multivariate modelling and the measure of spatial autocorrelation. Both can help the interpretation of results in the case that the dataset features population structure. Bivariate models may include the global ancestry coefficients provided by a preliminary analysis. This setup can detect which loci are correlated with the environment while taking demography into account. Additionally, the introduction of measurements of spatial autocorrelation into these analyses integrates spatial statistics with landscape genomics. Contrary to most current and non-spatial models (e.g. Frichot *et al.*, 2013; Coop *et al.*, 2010), this approach integrated in Samβada allows the determination of whether the observed data reflects independent samples, a requirement of the underlying modelling assumptions of such methodologies. Measuring spatial autocorrelation assesses whether the occurrence of a genotype is related to its frequency in the surrounding locations. More specifically, local indices of spatial autocorrelation allow the mapping of areas prone to spatial dependency. On the basis of the present analysis, using spatial statistics in conjunction with correlative models may lower the risk of false positives due to population structure in landscape genomics.

In the present study, Samβada detected the highest number of SNPs as potentially subject to selection among the four approaches. However when comparing the positions of these SNPs, 1,029 of them were less than 100,000 base pairs apart from another detected locus, thus some of these detections might refer to the same signature of selection. Samβada's results partially match with those of BayEnv with 435 common SNPs (i.e. 22% of BayEnv's detections). Concerning the third correlative approach, LFMM is more conservative than Samβada but the correspondence is better since 154 loci (out of 280, i.e. 55% of LFMM's detections) are detected by both methods. Moreover, 25 SNPs detected by LFMM only are less than 100,000 base pairs apart from a loci detected by Samβada, potentially identifying the same selection signature. The order of detections differed between the two methods, as the most significant loci detected by Samβada are ignored by LFMM. Lastly, Arlequin's best results involved 17 SNPs with p-values lower than 10^{-4} (significance threshold: $\alpha = 2.5 \cdot 10^{-7}$), out of which 2 were common with Samβada and 16 were common with BayEnv. This result suggests that population-based methods, whether using outliers or environmental correlations, tend to detect the same selection signatures. On the one hand, Samβada's detection rate may indicate the occurrence of some false positives due to population structure; on the other hand, the discrepancy between the results may indicate that the more conservative approaches have some false negatives. Thus the actual number of loci subject to selection is likely to lie in between. Comparing the results in the light of spatial dependence gives information about the differences between Samβada's and LFMM's detections. Maps of local spatial autocorrelation for ARS-113 (GG) and HM-28 (GG) illustrated a general trend: LFMM discarded SNPs showing significant local spatial autocorrelation for a large proportion of the sampling locations, while Samβada detected them. Thus measuring local autocorrelation of candidate genotypes may help distinguishing between the effects of local adaptation and those of population structure among Samβada detections.

Regarding common detections, the three SNPs identified by Samβada when population structure was included as a covariate were among the common detections of correlatives approaches. Thus pre-existent knowledge on demography may be built on to refine correlation-based detections of selection signatures. One possible approach could consist of computing population structure and then including one variable summarising this structure in the constant model used by Samβada. This way, only genotypes showing a significant correlation with the environment while taking the population structure into account would be detected. Concerning the biological function of the common detections, these three loci are located on chromosome 5, near the gene POLR3B whose mouse counterpart is involved in limiting infection by intracellular bacteria and DNA viruses (UniProt, www.uniprot.org). Moreover, genotype HM-28 (GG) shows spatial autocorrelation in the North-Western part of Uganda and this area overlaps with one of those where the higher load of tse-tse fly (*Glossina* spp.) occur in the country (Abila *et al.* (2008); MAAIF *et al.*, 2010). Hence the risk of cattle trypanosomiasis is high in this region and the detected mutations may be involved in parasite resistance.

The increasing availability of large molecular datasets raises challenges regarding their analysis. Correlative approaches in landscape genomics enable fast detection of candidate loci to local adaptation. However these methods must take into account the effect of population structure (Frichot *et al.*, 2013; Joost *et al.*, 2013; De Mita *et al.*, 2013). Limited dispersal of individuals leads to spatial autocorrelation of marker frequencies, which may cause spurious correlations with the environment. Samβada addresses the first topic by detecting rapidly selection signatures and the second one by measuring the level of spatial autocorrelation for candidate loci. The next methodological step involves developing spatially-explicit models that directly include autocorrelation. Guillot *et al.* (2014) provide such a model, however the current R-based implementation does not enable whole-genome analysis. Alternatively Geographically Weighted Regressions (GWR) measure the spatial stationarity of regression coefficients by fitting a distinct model for each sampling location. The number of neighbouring points considered for each sampling location is given by the weighting scheme. These models allow some "local" coefficients to differ between sampling points while some "global" coefficients are common to all points (Fotheringham *et al.*, 2002; Joost *et al.*, 2013). Thus GWR enables building a null model where the constant term may vary in space and then refining it by adding a global environmental effect for all locations. Comparing these two models would enable an assessment of whether the global environmental effect is needed to describe the distribution of the genotype. The key advantage of allowing the constant term to vary in space is to take spatial autocorrelation into account in the models. This way, GWR allows an investigation of the spatial behaviour of loci showing selection signature with standard logistic regressions and may help to distinguish between local adaptation and population structure in landscape genomics. However GWR models require a fine-tuning of the weighting scheme from the user, which restrains their application to very large datasets.

Computation time is critical when processing large datasets. In this context, Samβada is able to swiftly analyse high-density SNP-chips and variants from whole-genome se-

quencing (e.g. the case study presented in here is analysed within 69 minutes for univariates models alone and 8.5 hours for both univariate and bivariate models). When considering single-process computations, Samβada is approximately 4.5 times quicker than LFMM and 30 times than BayEnv. Both Samβada and LFMM enable parallelised processing. Samβada's processing speed, combined with its ability to analyse the spatial autocorrelation in molecular data and to incorporate prior knowledge on population structure, suits a wide range of applications, especially those involving whole genome sequence data.

Acknowledgement

References

Abila, P. P., Slotman, M. A., Parmakelis, A., Dion, K. B., Robinson, A. S., Muwanika, V. B., Enyaru, J. C. K., Lokedi, L. M., Aksoy, S., and Caccone, A. (2008). High levels of genetic differentiation between ugandan *Glossina fuscipes fuscipes* populations separated by lake kyoga. *PLOS Neglected Tropical Diseases*, **2**(5), e242.

Alexander, D. H., Novembre, J., and Lange, K. (2009). Fast model-based estimation of ancestry in unrelated individuals. *Genome Research*, **19**(9), 1655–1664.

Anselin, L. (1995). Local Indicators of Spatial Association - LISA. *Geographical Analysis*, **27**(2), 93–115. GISDATA (Geographic Information Systems Data) Specialist Meeting on GIS (Geographic Information Systems) and Spatial Analysis, Amsterdam, Netherlands, Dec 01-05, 1993.

Beaumont, M. A. and Nichols, R. A. (1996). Evaluating loci for use in the genetic analysis of population structure. *Proceedings Royal Society London B*, **263**, 1619–1626.

Colli, L., Joost, S., Negrini, R., Nicoloso, L., Crepaldi, P., Ajmone-Marsan, P., and Consortium, E. (2014). Assessing the spatial dependence of adaptive loci in 43 European and Western Asian goat breeds using AFLP markers. *PLOS ONE*, **9**(1).

Coop, G., Witonsky, D., Di Rienzo, A., and Pritchard, J. K. (2010). Using environmental correlations to identify loci underlying local adaptation. *Genetics*, **185**(4), 1411–1423.

De Mita, S., Thuillet, A.-C., Gay, L., Ahmadi, N., Manel, S., Ronfort, J., and Vigouroux, Y. (2013). Detecting selection along environmental gradients: analysis of eight methods and their effectiveness for outbreeding and selfing populations. *Molecular Ecology*, **22**(5), 1383–1399.

Dobson, A. J. and Barnett, A. G. (2008). *An Introduction to Generalized Linear Models*. Chapman & Hall, 3 edition.

Excoffier, L. and Lischer, H. E. L. (2010). Arlequin suite ver 3.5: a new series of programs to perform population genetics analyses under Linux and Windows. *Molecular Ecology Ressources*, **10**(3), 564–567.

Farr, T. G., Rosen, P. A., Caro, E., Crippen, R., Duren, R., Hensley, S., Kobrick, M., Paller, M., Rodriguez, E., Roth, L., Seal, D., Shaffer, S., Shimada, J., Umland, J., Werner, M., Oskin, M., Burbank, D., and Alsdorf, D. (2007). The shuttle radar topography mission. *Reviews of Geophysics*, **45**(2).

Foll, M. and Gaggiotti, O. (2008). A genome-scan method to identify selected loci appropriate for both dominant and codominant markers: A Bayesian perspective. *Genetics*, **180**, 977–993.

Fotheringham, A. S., Brunsdon, C., and Charlton, M. (2002). *Geographically Weighted Regression: the analysis of spatially varying relationships*. John Wiley & Sons, Chichester, 1 edition.

Frichot, E., Schoville, S. D., Bouchard, G., and François, O. (2013). Testing for associations between loci and environmental gradients using latent factor mixed models. *Molecular Biology and Evolution*, **30**(7), 1687–1699.

GDAL Development Team (2013). *GDAL - Geospatial Data Abstraction Library, Version 1.10*. Open Source Geospatial Foundation.

Guillot, G., Vitalis, R., le Rouzic, A., and Gautier, M. (2014). Detecting correlation between allele frequencies and environmental variables as a signature of selection. A fast computational approach for genome-wide studies. *Spatial Statistics*, **8**, 145–155.

Günther, T. and Coop, G. (2013). Robust identification of local adaptation from allele frequencies. *Genetics*, **195**(1), 205–220.

Henry, P. and Russello, M. A. (2013). Adaptive divergence along environmental gradients in a climate-change-sensitive mammal. *Ecology and Evolution*, **3**(11), 3906–3917.

Hijmans, R., Cameron, S., Parra, J., Jones, P., and Jarvis, A. (2005). Very high resolution interpolated climate surfaces for global land areas. *International Journal Of Climatology*, **25**(15), 1965–1978.

Jones, M., Forester, B., Teufel, A., Adams, R., Anstett, D., Goodrich, B., Joost, S., and Manel, S. (2013). Integrating spatially explicit approaches to detect adaptive loci in a landscape genomics context. *Evolution*, **67**, 3455–3468.

Joost, S., Bonin, A., Bruford, M. W., Després, L., Conord, C., Erhardt, G., and Taberlet, P. (2007). A spatial analysis method (SAM) to detect candidate loci for selection: towards a landscape genomics approach to adaptation. *Molecular Ecology*, **16**(18), 3955–3969.

Joost, S., Kalbermatten, M., and Bonin, A. (2008). Spatial Analysis Method (SAM): a software tool combining molecular and environmental data to identify candidate loci for selection. *Molecular Ecology Ressources*, **8**, 957–960.

Joost, S., Vuilleumier, S., Jensen, J. D., Schoville, S., Leempoel, K., Stucki, S., Widmer, I., Melodelima, C., Rolland, J., and Manel, S. (2013). Uncovering the genetic basis of adaptive change: on the intersection of landscape genomics and theoretical population genetics. *Molecular Ecology*, **22**(14), 3659–3665.

Ministry of Agriculture, Animal Industry and Fisheries, Uganda, Uganda Bureau of Statistics, Food and Agriculture Organization of the United Nations, International Livestock Research Institute, and World Resources Institute (2010). *Mapping a better future: Spatial Analysis and Pro-Poor Livestock Strategies in Uganda*, pages 30–37. World Resources Institute, Washington, DC and Kampala.

Mitton, J. B., Linhart, Y. B., Hamrick, J. L., and Beckman, J. S. (1977). Observations on genetic structure and mating system of ponderosa pine in Colorado front range. *Theoretical and Applied Genetics*, **51**(1), 5–13.

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, **37**(1/2), pp. 17–23.

Pemstein, D., Quinn, K. M., and Martin, A. D. (2011). The Scythe statistical library: An open source C++ library for statistical computation. *Journal of Statistical Software*, **42**(12), 1–26.

Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M. A., Bender, D., Maller, J., Sklar, P., de Bakker, P. I., Daly, M. J., and Sham, P. C. (2007). PLINK: A tool set for whole-genome association and population-based linkage analyses. *The American Journal of Human Genetics*, **81**(3), 559 – 575.

Vitalis, R., Dawson, K., Boursot, P., and Belkhir, K. (2003). DetSel 1.0: A computer program to detect markers responding to selection. *Journal of Heredity*, **94**(5), 429–431.

Vitti, J. J., Grossman, S. R., and Sabeti, P. C. (2013). Detecting Natural Selection in Genomic Data. In Bassler, BL and Lichten, M and Schupbach, G, editor, *Annual Review of Genetics, Vol 47*, volume 47 of *Annual Review of Genetics*, pages 97–120. Annual Reviews, 4139 EL CAMINO WAY, PO BOX 10139, PALO ALTO, CA 94303-0897 USA.

Wright, S. (1949). The genetical structure of populations. *Annals of Eugenics*, **15**(1), 323–354.

High performance computation of landscape genomic models integrating local indices of spatial association
Supplementary material

1 Spatial autocorrelation

Several indices are available for measuring the global spatial autocorrelation in a dataset. Samβada uses the Moran's I (Moran, 1950), defined as follows:

$$I = \frac{n}{S_0} \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij}(y_i - \bar{y})(y_j - \bar{y})}{\sum_{i=1}^{n}(y_i - \bar{y})^2} = \frac{n}{S_0} \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij} z_i z_j}{\sum_{i=1}^{n} z_i^2} \tag{1}$$

avec

n	number of sampling points;
w_{ij}	weight of point j in the neighbourhood of i, defined by the spatial kernel;
S_0	sum of all weights $\left(S_0 = \sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij}\right)$;
y_i, y_j	values for points i and j;
$\bar{y}$	mean value;
z_i, z_j	deviations from the mean.

Local indices of spatial association (LISA, Anselin, 1995) measure the local association between the value of a point and the neighbouring points. Samβada computes a local variant of the Moran's I for each point i:

$$I_i = \frac{n-1}{\sum_{i=1}^{n} z_i^2} \left(z_i \sum_{j=1}^{n} w_{ij} z_j \right) \tag{2}$$

$$\tag{3}$$

LISAs are defined in such a way that their sum over all points is proportional to a global measure of spatial autocorrelation, in this case the Moran's I:

$$I = \frac{1}{n-1} \sum_{i=1}^{n} I_i \tag{4}$$

2 Environmental variables

Variable	Description	Data source	Used for univariate analysis	Used for bivariate analysis
alt_SRTM	Altitude [m]	SRTM3	X	
aspect	Orientation of the relief [°]	Derived from SRTM3	X	X
BIO2	Mean Diurnal Range (Mean of monthly (max temp - min temp))	WorldClim	X	X
BIO3	Isothermality (BIO2/BIO7) (* 100)		X	X
BIO7	Temperature Annual Range (max temp - min temp)		X	
BIO9	Mean Temperature of Driest Quarter		X	
BIO12	Annual Precipitation		X	X
BIO15	Precipitation Seasonality (Coefficient of Variation)		X	X
BIO18	Precipitation of Warmest Quarter		X	X
latitude	Latitude	Sampling measurements	X	X
longitude	Longitude		X	X
prec2	Precipitation in February	WorldClim	X	
prec3	Precipitation in March		X	
prec4	Precipitation in April		X	X
prec5	Precipitation in May		X	X
prec6	Precipitation in June		X	
prec7	Precipitation in July		X	
prec9	Precipitation in September		X	
prec10	Precipitation in October		X	X
prec11	Precipitation in November		X	X
slope	Slope of the relief [%]	Derived from SRTM3	X	X
tmin10	Minimal temperature in October	WorldClim	X	X
tmax10	Maximal temperature in October		X	
Ankole	Coefficient of ancestry to the population Ankole	Analysis with Admixture		X
Number of variables			23	15

Table S1: Environmental variables used to detect selection signatures with correlative approaches. Univariate analyses were performed with Sambada, BayEnv and LFMM and bivariate analyses with Sambada
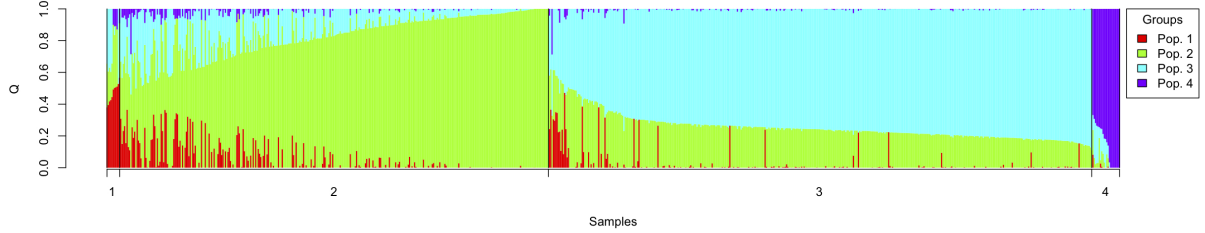
3 Population structure

Figure S1: Population structure computed with Admixture (Alexander *et al.*, 2009). Individuals are gathered together by populations, labeled horizontally. The assignation is based on the highest membership coefficient $Q_{\max}$ of the sample. Inside each population, individuals are ranked by increasing (or decreasing) value of $Q_{\max}$.

4 Computation times

	41,215 SNPs 804 samples	634,849 SNPs 102 samples
Samβada	1.2	2.9
Samβada biv.	8.7	18.4
BayEnv	41.3	62.,2
LFMM	3.2	16.0
LFMM (mono)	6.1	58.1
Arlequin	?	?

Table S2: Comparison of computation times among methods. The datasets used in this case include chromosome X. The first dataset refers to the data used in the article (SNPs and individual call rates=5%, MAF=1%, including chr. X). The second dataset refers to 102 samples of Ugandan cattle that were chosen among the 917 samples to be genotyped with a high-density SNP chip (BovineHD, Illumina Inc; SNPs and ind. call rates=5%, MAF=5%). Durations are expressed in hours. "LFMM (mono)" shows the durations if using a single thread. Arlequin was run on another computer, thus its performance is not comparable.

3

5 Comparison of detections

	Loci	Chr.	Pos. [Mbp]	Samβada	BayEnv	LFMM	Arlequin	Detections
1	ARS-BFGL-NGS-113888	5	48.32	1	1	0	0	2
2	Hapmap41074-BTA-73520	5	48.35	1	1	0	0	2
3	Hapmap41762-BTA-117570	5	18.94	1	1	0	0	2
4	ARS-BFGL-NGS-46098	20	2.95	1	1	0	0	2
5	Hapmap41813-BTA-27442	5	49.04	1	1	0	0	2
6	BTA-73516-no-rs	5	48.75	1	1	0	0	2
7	**Hapmap28985-BTA-73836**	**5**	**70.34**	**1**	**1**	**1**	**0**	**3**
8	Hapmap31863-BTA-27454	5	48.99	1	1	0	0	2
9	**ARS-BFGL-NGS-106520**	**5**	**70.20**	**1**	**1**	**1**	**0**	**3**
10	**BTA-73842-no-rs**	**5**	**70.18**	**1**	**1**	**1**	**0**	**3**
11	Hapmap50523-BTA-98407	5	46.74	1	1	0	0	2
12	BTB-01400776	20	2.70	1	1	0	0	2
13	Hapmap23956-BTA-36867	15	47.20	1	1	0	0	2
14	ARS-BFGL-NGS-10586	2	128.64	1	1	0	0	2
15	ARS-BFGL-NGS-43694	5	49.65	1	1	0	0	2
16	BTA-122374-no-rs	14	16.44	1	1	0	0	2
17	BTB-01356178	20	2.49	1	1	0	0	2
18	**ARS-BFGL-NGS-94862**	**11**	**103.53**	**1**	**1**	**1**	**0**	**3**
19	BTA-108359-no-rs	14	16.31	1	0	0	0	1
20	ARS-BFGL-NGS-15960	5	28.02	1	1	0	0	2
21	ARS-BFGL-NGS-116294	2	128.58	1	1	0	0	2
22	INRA-566	13	57.94	1	0	1	0	2
23	**BTA-49720-no-rs**	**5**	**69.66**	**1**	**1**	**1**	**0**	**3**
24	ARS-BFGL-NGS-56387	13	24.36	1	1	0	0	2
25	BTA-28185-no-rs	26	22.78	1	0	0	0	1

Table S3: List of SNPs detected by Samβada corresponding to the models with the highest G scores. Loci are identified by their name, their chromosome and their position in million base pairs. The following columns show which method detected them and the last one counts these detections. Loci in bold type are the commons discoveries of Samβada, LFMM and BayEnv. Local indices of spatial autocorrelation were computed for SNPs on lines 1 and 7.

References

Alexander, D. H., Novembre, J., and Lange, K. (2009). Fast model-based estimation of ancestry in unrelated individuals. *Genome Research*, **19**(9), 1655–1664.

Anselin, L. (1995). Local Indicators of Spatial Association - LISA. *Geographical Analysis*, **27**(2), 93–115. GISDATA (Geographic Information Systems Data) Specialist Meeting on GIS (Geographic Information Systems) and Spatial Analysis, Amsterdam, Netherlands, Dec 01-05, 1993.

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, **37**(1/2), pp. 17–23.