

SPECIAL ISSUE: POPULATION GENOMICS WITH R

MINOTAUR: A platform for the analysis and visualization of multivariate results from genome scans with R Shiny

ROBERT VERITY,^{*,†} CAITLIN COLLINS,^{*,†} DAREN C. CARD,[†] SARA M. SCHAAAL,[‡] LIUYANG WANG[§] and KATIE E. LOTTERHOS[‡]

^{*}MRC Centre for Outbreak Analysis and Modelling, Department of Infectious Disease Epidemiology, Imperial College London, St. Mary's Campus, Norfolk Place, London W21PG, UK, [†]Department of Biology, The University of Texas at Arlington, 501 S. Nedderman Drive, Arlington, TX 76019, USA, [‡]Department of Marine and Environmental Sciences, Northeastern University, 430 Nahant Road, Nahant, MA 01908, USA, [§]Department of Molecular Genetics and Microbiology, School of Medicine, Duke University, 213 Research Drive, Durham, NC 27710, USA

Abstract

Genome scans are widely used to identify 'outliers' in genomic data: loci with different patterns compared with the rest of the genome due to the action of selection or other nonadaptive forces of evolution. These genomic data sets are often high dimensional, with complex correlation structures among variables, making it a challenge to identify outliers in a robust way. The Mahalanobis distance has been widely used, but has the major limitation of assuming that data follow a simple parametric distribution. Here, we develop three new metrics that can be used to identify outliers in multivariate space, while making no strong assumptions about the distribution of the data. These metrics are implemented in the R package MINOTAUR, which also includes an interactive web-based application for visualizing outliers in high-dimensional data sets. We illustrate how these metrics can be used to identify outliers from simulated genetic data and discuss some of the limitations they may face in application.

Keywords: genomic scans, kernel density, Mahalanobis

Received 20 May 2016; accepted 26 July 2016

Introduction

Knowledge of the genetic architecture of biological traits—the number of loci that affect a phenotype, the magnitude of their effect and their distribution across the genome—not only illuminates the evolutionary processes that shape genomes, but also has important implications for complex diseases (McCarthy & Hirschhorn 2008), conservation (Kohn *et al.* 2006; Allendorf *et al.* 2010; Funk *et al.* 2012) and breeding programmes (Goddard & Hayes 2009; Varshney *et al.* 2009). With the advent of next-generation sequencing, we now have the ability to examine genomes at a fine scale, and, as a result, we have identified a large number of genomic variants that are implicated in complex diseases (Carlson *et al.* 2004; Hindorff *et al.* 2009) and adaptation to the local environment (Savolainen *et al.* 2013). This wealth of data is likely to yield new insights, but it also brings with it the challenge of extracting the relevant signal from noisy, complex,

multidimensional data sets. This is perhaps one reason why most of the variants detected so far have only managed to explain a very small proportion of the observable phenotypic variation (Yang *et al.* 2010; Brachi *et al.* 2011).

The preferred method for detecting genomic variants is via genome scans. There are many different approaches towards scanning genomes, but all are based on the same premise: that the loci of interest to the investigator are likely to be statistical outliers when compared with the rest of the genome. The particular choice of statistic will depend on the question being asked and the experimental design and may include one or more statistics from the following categories: tests for genetic differentiation (Lotterhos & Whitlock 2014; Hoban *et al.* in press), scans for strong positive selection and/or selective sweeps (Hohenlohe *et al.* 2010; Vatsiou *et al.* 2016), genomewide association studies for phenotype-associated loci (GWAS, reviewed in Carlson *et al.* 2004 and McCarthy *et al.* 2008), linkage mapping for quantitative trait loci (QTL, Savolainen *et al.* 2013), genetic-environment associations (reviewed in Rellstab *et al.* 2015) and scans for differentially expressed genes (Wang *et al.* 2009). A number of different genome scan test

Correspondence: Robert Verity, Fax: 020 7402 3927; E-mail: r.verity@imperial.ac.uk

¹These authors contributed equally to this work.

statistics may be calculated for a single genomic data set, and these are usually examined one at a time (i.e. in univariate analyses). Some test statistics may be highly correlated, while the power of other test statistics may vary for different regions of the genome depending on the details of selection, recombination, mutation and migration rates (Tiffin & Ross-Ibarra 2014). Additionally, the power of different approaches may vary among species because of demographic history, and within a species because of sampling design (De Mita *et al.* 2013; De Villemerueil *et al.* 2014; Lotterhos & Whitlock 2015). Finally, loci with intermediate probabilities of detection will often exhibit the highest variance in results from genome scans.

Given the complex evolutionary histories of most species, it is doubtful whether any single statistic can fully capture the genomic signal of interest in the majority of cases (Verity & Nichols 2014). Furthermore, the uncertainty in demographic history, coupled with the variation in statistical outcomes in different scenarios, makes it difficult to know which statistics have the greatest power to detect selection and which have the highest false-positive rates. These issues point to a need for composite, multivariate outlier methods that integrate information across multiple test statistics.

Multivariate methods have been utilized extensively in many biological applications, although in application to genome scans, the power of the multivariate approach for detecting outliers has not yet been fully evaluated. Because some dimension reduction methods such as principal component analysis rely on assumptions about the data that may be unjustifiable in the context of genome scans (O'Reilly *et al.* 2012), these methods are not ideally designed for the identification of multivariate outliers (Patterson *et al.* 2006). Some GWAS analyses have successfully employed multivariate approaches to identify genetic associations with multiple phenotypes (O'Reilly *et al.* 2012; Galesloot *et al.* 2014). Additionally, multivariate approaches have also been used in GWAS meta-analysis to simultaneously consider multiple genetic or phenotypic variables (reviewed in Evangelou & Ioannidis 2013). It is evident, however, that more opportunities exist for the use of multivariate approaches in outlier detection than are currently being capitalized on.

While there are dedicated software tools for calculating a variety of test statistics, there does not currently exist a unified platform for the filtering, visualization and integration of test statistics in multivariate space. Here, we describe a new R package called **MINOTAUR** (**M**ultivariate **V**isualisatio**N** and **O**u**T**lier **A**nalysis **U**sing **R**) built specifically for this purpose. This software package—initiated during a hackathon for population genetics in R (<https://github.com/NESCent/r-popgen>

hackathon)—provides functions for detecting outliers in multivariate space alongside procedures to manipulate, summarize and visualize these data. The R software environment (R Core Team 2015) is free, open source and hosts a large collection of tools for statistical analysis, making it the ideal host for the development and uptake of such a platform. Furthermore, because data visualization is an important part of verifying and identifying outliers, the R Shiny and Shiny Dashboard environments (Chang 2015; Chang *et al.* 2016) have been employed to provide **MINOTAUR** users with an interactive interface that streamlines the process of data input, statistical analysis and graphical exploration. Together, these tools have the potential to increase the efficiency with which the results of genome scans are interrogated.

Approaches to identifying multivariate outliers

In the **MINOTAUR** package, we implement four composite measures that can be used to integrate information over multiple univariate statistics: the Mahalanobis distance, harmonic mean distance, nearest neighbour distance and kernel density deviance. We developed the latter three measures, which are related to Mahalanobis distance, but make no strong assumptions about the parametric form of the data, meaning they can be applied to multivariate statistics that have complex correlated or even multimodal distributions. Some of these measures are heavily influenced by the distance of points from the multivariate centroid (Mahalanobis and harmonic mean distance), while others are mainly influenced by the sparseness of points in the local vicinity (nearest neighbour distance and kernel density deviance), and so we would expect the measures to behave differently from one another and to vary in their behaviour depending on the data at hand.

The calculation of these composite measures has been optimized for genome-scale data using precompiled routines, written in C++ and integrated into R using the package **Rcpp** (Eddelbuettel & François 2011; Eddelbuettel 2013). Several packages devoted to multivariate statistics that may be appropriate for genome-scale data already exist in R (see Table S1, Supporting information), and thus, users are free to utilize both existing statistical methods and the more targeted functions included within the **MINOTAUR** package.

Mahalanobis distance

The Mahalanobis distance is a multidimensional measure of the number of standard deviations that a point lies from the mean of a distribution. The Mahalanobis distance of a d -dimensional observation $x_i = (x_{i1}, x_{i2}, \dots, x_{id})^T$ from a distribution of N variables with mean

$\bar{x} = (\bar{x}_1, \bar{x}_2, \dots, \bar{x}_d)^T$ and covariance matrix S is defined as follows (Mahalanobis 1936):

$$D_M(x_i) = \sqrt{(x_i - \bar{x})^T S^{-1} (x_i - \bar{x})} \quad (\text{eqn 1})$$

This distance differs from the ordinary Euclidean distance due to the correction for covariance among observations, making it a better distance measure for genome scan summary statistics because it does not assume that statistics are independent (i.e. Euclidean distance equals Mahalanobis distance when S is a diagonal matrix). However, this distance does make the assumption that points disperse smoothly from a single multivariate centroid, and so it will tend to perform poorly when observations have a complex or multimodal distribution.

Harmonic mean distance

In this context, the 'harmonic mean distance' of an observation x_i refers to the harmonic mean of the distances between this point and all other points. The distance measure used here is the Euclidian distance normalized by multiplying by the inverse covariance matrix. This ensures that results are not dominated by a few statistics with a large spread and also accounts for any potential correlation between statistics, analogously to the Mahalanobis distance. Mathematically, we can define the harmonic mean distance as follows:

$$D_H(x_i) = N \left[\sum_{j \neq i} [(x_i - x_j) S^{-1} (x_i - x_j)]^{-1/2} \right]^{-1} \quad (\text{eqn 2})$$

The harmonic mean is heavily influenced by small values, which in this context means local effects are amplified. However, more distant points also have some effect on the final value (unlike the nearest neighbour distance described below), and so the harmonic mean strikes a balance between local and global effects. This has some advantages in outlier detection, as observations that are both distant from the main mass of the data and have few neighbours in the local vicinity will tend to be outliers.

Nearest neighbour distance

The nearest neighbour distance of the observation x_i gives the minimum distance between this point and any other point. As with the harmonic mean distance, we use the Euclidian distance normalized by the inverse covariance matrix. Mathematically, we can write

$$D_N(x_i) = \min_{j \neq i} \left(\sqrt{(x_i - x_j)^T S^{-1} (x_i - x_j)} \right) \quad (\text{eqn 3})$$

This statistic exclusively measures local effects, being largest when an observation is a long way from any other point. Because this distance is only based on two points (the focal point and its nearest neighbour), it is not influenced by the global distribution of the data, unlike the harmonic mean distance.

Kernel density deviance

Kernel density-based methods attempt to capture mathematically the distribution of the data as the sum of a number of simple parametric distributions. Here, we apply these methods to identifying multivariate outliers, defined as those points with a low density of data around them in multivariate space. We assume a multivariate normal kernel $G(x_i|x_j, \lambda^2 S)$ centred at the point x_j , where λ is the bandwidth of the kernel, which is scaled in each dimension by the covariance matrix of the data. We then calculate the leave-one-out log-likelihood (Leiva-Murillo & Artés-Rodríguez 2012) of the point x_i as follows:

$$L(x_i|\lambda) = \log \left(\frac{1}{N-1} \sum_{j \neq i} G(x_i|x_j, \lambda^2 S) \right) \quad (\text{eqn 4})$$

In other words, this is equal to the log-probability density of the point x_i under the kernel density distribution constructed from all points *apart from* x_i . Our final density-based measure is defined as follows:

$$D_K(x_i) = -2L(x_i|\lambda), \quad (\text{eqn 5})$$

which is sometimes referred to as the Bayesian deviance. This will be large whenever the density of the point x_i is low, and so the kernel density deviance can be thought of as a measure of the sparseness of points around the focal point.

One challenge when using kernel density methods is choosing an appropriate value for the bandwidth. Here, we simply use the bandwidth for which the total deviance of all points is minimized, that is

$$\lambda^* = \operatorname{argmin}_{\lambda} \left(\sum_{i=1}^N -2L(x_i|\lambda) \right) \quad (\text{eqn 6})$$

It can be shown that this is equivalent to the maximum-likelihood value of λ under the leave-one-out criterion. The value λ^* can be found using the MINOTAUR function `kernelDeviance()`, which takes a vector of bandwidths as input and returns the total deviance of

each. This function can be used to search for the minimum value of λ manually, or via an optimization routine such as `optim()`. Users are also free to use any other bandwidth, entered manually, or in the absence of a user-defined bandwidth, a simple method based on Silverman's rule is implemented as a default (this assumes that data are normally distributed and are a simple function of the standard deviation of the samples (Silverman 1986)).

The MINOTAUR R package—an R Shiny graphical user interface for multivariate outlier analysis and visualization

The MINOTAUR package performs two main functions: (i) it calculates the compound multivariate outlier statistics described above and (ii) it enables users to harness the interactive graphical power of the R Shiny environment to manipulate and visualize their data within the MINOTAUR graphical user interface (GUI). The GUI allows users to perform the former task with the click of a button; however, outlier identification can also be performed on the R command line using stand-alone functions available in MINOTAUR, if preferred. Directions for downloading and installing the package can be found at the end of this manuscript.

The MINOTAUR GUI is designed to streamline the process of genomic data analysis and outlier identification, taking users from data input to graphical output within a single platform. Distinct panels are used for each stage of the analysis, including data input and filtering, outlier detection via the methods described above and plotting results (e.g. histograms, scatter plots and Manhattan plots). An overview of the MINOTAUR GUI workflow is shown in Fig. 1.

In the *Data* panel, the MINOTAUR GUI allows users to either upload their own data sets or select among a set of four in-built example data sets. Data can be uploaded in a number of file formats, including comma- or tab-separated text files, and Rdata. Regardless of the file format, MINOTAUR expects all incoming data sets to be arranged in data frames, with each row representing a different genetic locus and each column representing a different univariate genome scan statistic (e.g. F_{ST} , Tajima's D) or other piece of locus-specific metadata (e.g. SNP identifiers, chromosomes/scaffolds and positions). Raw data objects can be filtered within the GUI, meaning, for example, that columns not related to outlier analysis can be dropped at an early stage.

Four example data sets are made available to users within the MINOTAUR package and GUI. The 'HumanGWAS' data set contains example output from an unpublished human genomewide association study.

The simulated 'NonParametricInverse' and 'NonParametricMultimodal' data sets each contain an example of nonparametric data, one with an inverse relationship and one that is highly multimodal. The 'TwoRef-Sim' data set contains population genetic data simulated under a model of expansion from two refugia (Lotterhos & Whitlock 2015). Note that the example data sets can also be accessed outside the GUI by running the `data()` command with the appropriate data set name. For example, to load the 'HumanGWAS' data set, type `data(HumanGWAS)` and hit ENTER. To learn more about a data set while in the R terminal, add a question mark before the data set name to load the relevant Help page; for example, type `?HumanGWAS` and hit ENTER.

In the *Multivariate Measures* panel, multiple univariate statistics can be integrated to produce the compound distance measures described above. These measures can be appended to the data frame and visualized interactively in the *Produce Plots* panel, which includes several sub-menus with useful plots for visualizing high-dimensional data sets, including Manhattan plots, 1D histograms and density-based 2D scatter plots. The plotting methods are designed with large genomic data sets in mind; for example, the `plot2d()` function included with the package calculates the density of points for a given bin size and shades bins according to the density of points within them and then optionally adds user-supplied points (ideally a small subset of points, for example the outliers only) to the plot. Additional options allow users to log-scale statistics and control various other visual settings commonly used when plotting data in R (Fig. 2).

Example applications of multivariate outliers

Evaluation of computational speed

First, we evaluated the speed of calculating the four compound distance measures for data sets with increasing numbers of loci (rows) and univariate statistics (columns). For this example, variables were randomly generated from a multivariate normal distribution. Table 1 gives the 'order' of complexity of these algorithms, together with measured run times for a data set composed of 50 000 loci and 10 variables (see Table S2, Supporting information). Overall, the Mahalanobis distance is calculated in a matter of seconds, even with particularly large data sets. The harmonic mean distance, nearest neighbour distance and kernel density deviance each scale approximately equally with increasing data set sizes, although the maximum-likelihood estimate of the ideal bandwidth for the latter measure can add significant computation time.

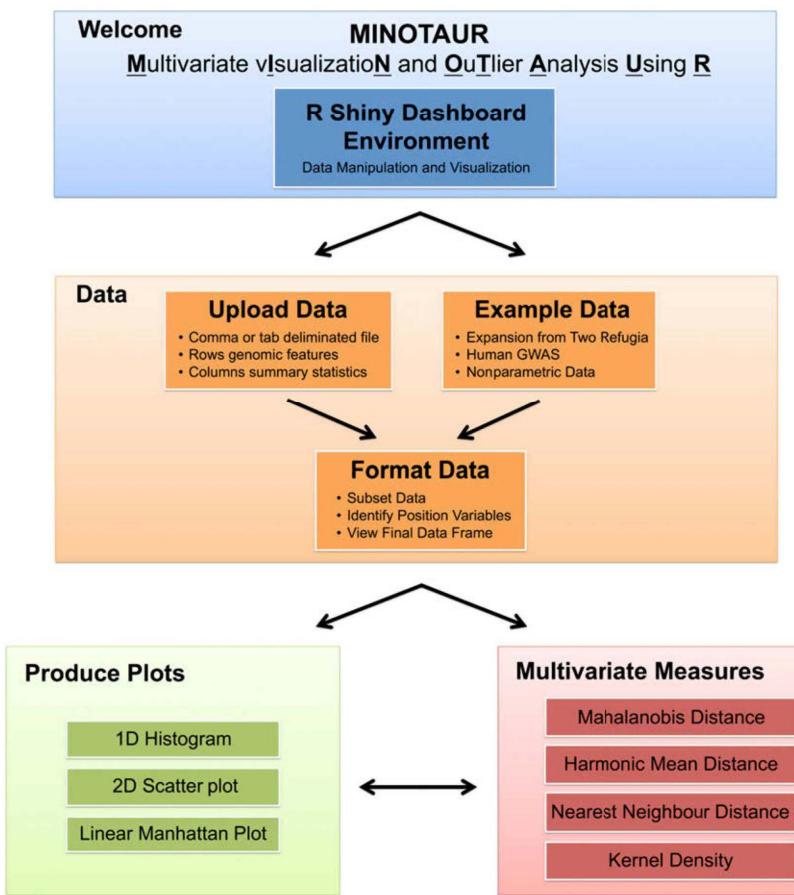


Fig. 1 Graphical overview of the MINOTAUR GUI workflow.

Example on simulated nonparametric distributions

Some kinds of genomic data—for example gene expression data—may generate complex nonparametric distributions. Genes that have high expression in one environment may have low expression in another environment, while investigators may be interested in identifying genes that have moderate expression in both environments. To test the performance of the multivariate outlier statistics in nonparametric situations, we simulated two examples of nonparametric distributions.

In the first example, we simulated a distribution of two variables that follow an inverse relationship, with some additional noise. We used contour plots to visualize the different ways in which each of the compound distance measures changes over the two-dimensional plane (Fig. 3). In these plots, the darker red lines indicate less-significant values of the test statistic and lighter yellow lines indicate more-significant values of the test statistic. We also looked at two manually chosen points on the plane—indicated by a square and a triangle—chosen to represent different sorts of outliers. The triangle would not be considered an outlier from the perspective of either one-dimensional distribution despite being a clear outlier from the two-dimensional

distribution, while the square would be considered an outlier in the first dimension but not the second. In this example, the nonparametric distribution affects the relative ability of the four statistics to identify each of these outliers (Fig. 4). The blue triangle would not have the largest value (i.e. not be the most outlying point) by the Mahalanobis or the harmonic mean distance, while it would have the largest value by nearest neighbour distance or kernel density deviance. In contrast, the blue square has the largest value of the test statistic by all four methods.

In the second example, we simulated a highly multimodal distribution from a normal mixture model. In this example, it can be seen how the parametric assumption of the Mahalanobis distance fails to capture the complexity of the data (Fig. S1, Supporting information). In contrast to the previous example, the harmonic mean distance behaves similarly to the kernel density deviance, and nearest neighbour distance has the most complex contour landscape.

Example on simulated genomic data

To test the power of multivariate statistics for genome scans, we applied them to a published simulated data set

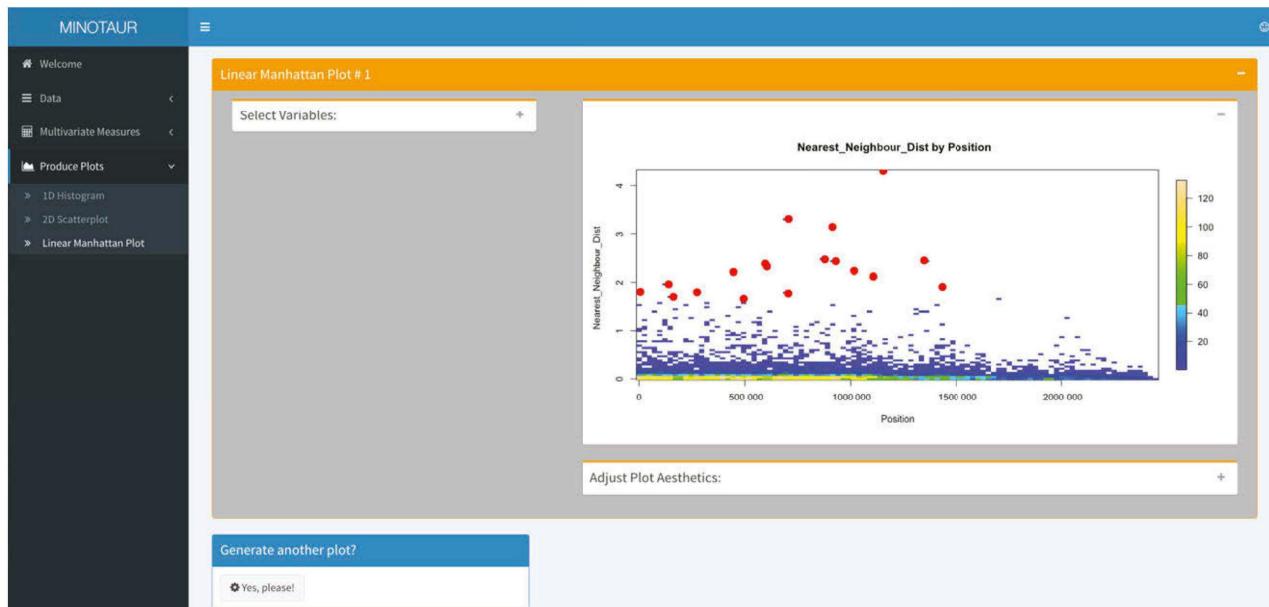


Fig. 2 Screenshot of MINOTAUR GUI highlighting the overall interface and the ability to visualize multivariate distributions. The plot is a Manhattan plot of the nearest neighbour distance across loci for all traits in the 'HumanGWAS' example data set provided as part of MINOTAUR. The base scatter plot demonstrates the binned visualization, where the density of data in an area is apparent from the colour. About 99.5 percentile outliers are indicated with solid orange circles. Visualization menus have been collapsed to simplify the image. Additional plots can also be stacked below to enable comparisons across multiple plots (not shown).

Table 1 Multivariate outlier detection methods implemented in MINOTAUR and associated computational run times

| Compound measure | Description | R Function | Computational complexity (big O notation) | Computation Elapsed Time for 50 000 loci and 10 variables (hh:mm:ss.ms) |
|----------------------------|---|----------------|---|---|
| Mahalanobis distance | Distance from multivariate centroid | Mahalanobis() | $O(Nk^2)$ | 00:00:00.095 |
| Harmonic mean distance | Inverse-weighted distance from all other points | harmonicDist() | $O(Nk^2)$ | 00:04:13.620 |
| Nearest neighbour distance | Distance to nearest neighbour | neighborDist() | $O(Nk^2)$ | 00:04:07.020 |
| Kernel density deviance | Local density of points | kernelDist() | $O(Nk^2)$ | 01:40:03.600 |

Computational complexity is given in 'big O' notation, with N referring to the number of observations and k the number of statistics (dimensions). Run times were determined using an Apple iMac with a 3.1 GHz Intel Core i5 processor and 32 GB of RAM running Apple OSX 10.9.5 and R version 3.2.3. Note that for computation time, the kernel density deviance includes both the maximum-likelihood estimation of the optimal bandwidth and the density calculations based on the optimal bandwidth

that was used to test different genome scan methods (Lotterhos & Whitlock 2014, 2015). Briefly, a landscape simulator was used to simulate haploid neutral and selected loci that adapted to an environmental cline (Lotterhos & Whitlock 2015). The landscape consisted of 360×360 demes, and the allele frequency of each deme changed each generation according to recurrence equations for mutation, migration, selection (if applicable) and drift (Lotterhos & Whitlock 2015). For the data set

used in this example, a total of 9900 neutral and 100 selected loci (simulated under varying strengths of selection: 12 loci with $s = 0.1$, 38 loci with $s = 0.01$ and 50 loci with $s = 0.005$) were simulated under a two-refuge demographic expansion. Individuals were then sampled from the landscape according to the allele frequency in each deme at 30 randomly chosen locations on the landscape at 20 individuals per location. For additional details, see Lotterhos & Whitlock (2014, 2015).

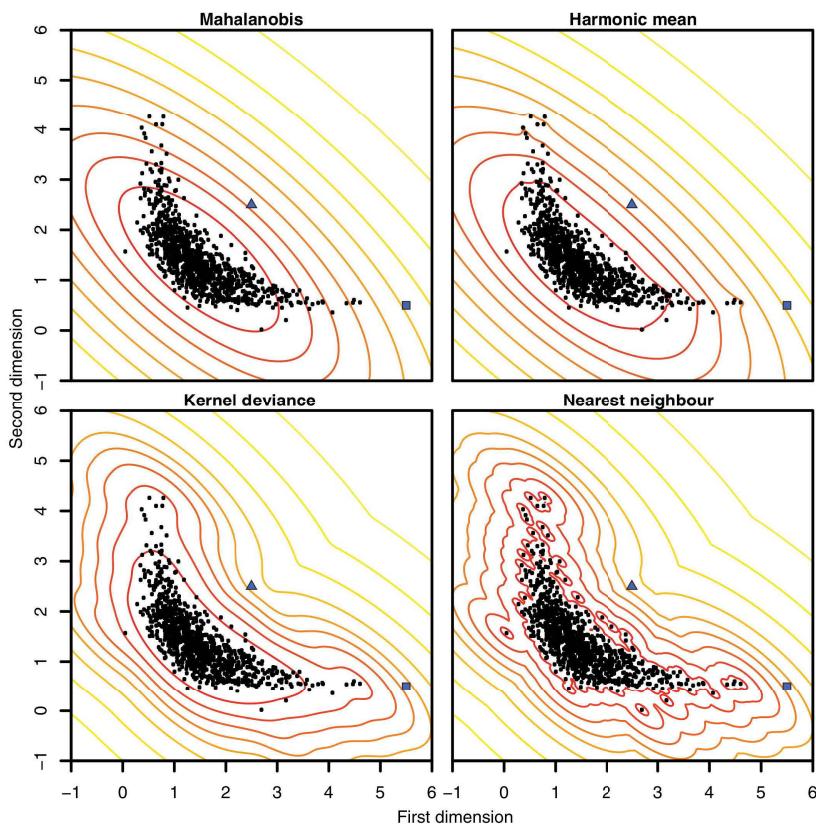


Fig. 3 Comparison of multivariate distance measures for nonparametric example data. Black dots show the simulated data, in which the two statistics (dimensions) are assumed to follow an inverse relationship with some additional noise. Solid lines show the distance measure computed at each point in the plane, arranged in 10% quantiles (e.g. the inner ring shows the 10% of locations with the smallest distance). The blue square and triangle show particular outlier points referred to in the main text.

The simulated data were used to create a single nucleotide polymorphism (SNP) table, and these data were used to perform genome scans in the programs BAYENV2 (Günther and Coop 2013) and LFMM (Frichot *et al.* 2013; now implemented in the R package LEA: Frichot & François 2015). A total of four univariate statistics from these two programs were used in the search for multivariate outliers: (i) log-Bayes factor (log-BF, a measure of the association between allele frequency and the environment in BAYENV2), (ii) Spearman's rho (a measure of the association between allele frequency and the environment in BAYENV2), (iii) $X^T X$ (a measure of genetic differentiation among populations in BAYENV2) and (iv) Z-score (a measure of the association between genotype and the environment in LFMM). These four univariate statistics, plotted in Fig. 4, were previously shown to have different strengths and weaknesses depending on sampling design and demographic history (Lotterhos & Whitlock 2015).

To illustrate the flexibility of the outlier functions implemented in MINOTAUR, we calculated multivariate outliers in two ways, corresponding to two different ways of calculating the covariance matrix S in equations (1) to (4). First, we used the traditional method of calculating the covariance matrix based on all the data. For high-dimensional data, estimation of the multivariate

mean and covariance (location and scatter) is expected to be robust to outliers as long as the proportion of outliers in the data is less than $1/(k + 1)$, where k is the number variables in the data frame (Rousseeuw & Van Driessen 1999). However, we found that even in this small data frame of only four variables and 10 000 loci, the 1% of selected loci (a fraction of which were true outliers) affected the estimation of the covariance matrix. For this reason, our MINOTAUR functions are designed to allow the user to input their own covariance matrix. To illustrate this use of the function, we also calculated a robust multivariate location and scatter estimate with a high breakdown point, using the 'Fast MCD' (minimum covariance determinant) estimator with the function CovNAMcd in the R package rrcovNA (Rousseeuw & Van Driessen 1999; Todorov *et al.* 2011).

To compare the ability of the univariate statistics and the multivariate statistics to separate neutral from selected loci, we calculated the empirical power. The empirical power is based on using all known neutral loci to generate a null distribution, and then, for each locus, an empirical P -value is calculated based on its cumulative frequency in this null distribution. To control for false discovery rate, empirical p -values were converted to q -values (in the R package qvalue: Dabney & Storey 2014) and loci with a q -value less than 0.05 were retained

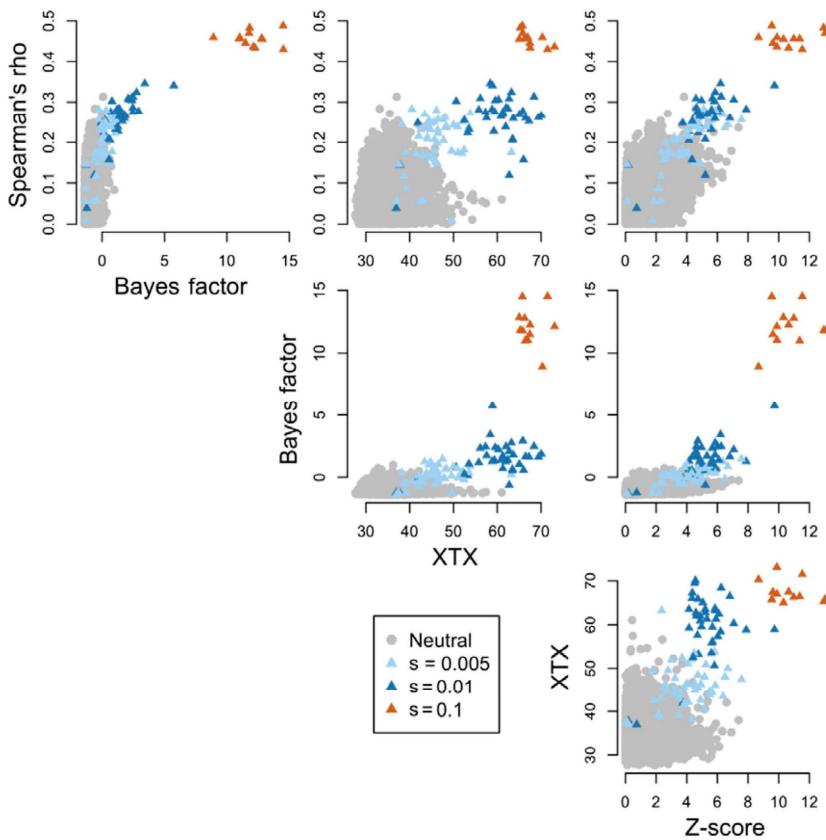


Fig. 4 Distributions of four univariate statistics from the two-refuge data set from Lotterhos & Whitlock (2015).

as positive hits (a q -value of 0.05 has a desired rate of five false positives of 100 positive hits).

For the univariate statistics, the empirical power was highest for log-BF (0.54) and lowest for Z-score (0.15), with Spearman's rho (0.46) and $X^T X$ (0.42) intermediate. For the multivariate statistics with the default covariance estimation, the empirical power was high for harmonic mean distance and Mahalanobis distance (0.41 for both), with kernel density and nearest neighbour distance performing poorly in this case (0.09 for both) (Fig. S2, Supporting information). For the user-input covariance matrix estimated with a high breakdown point (i.e. less influenced by outliers), the empirical power was highest for harmonic mean distance and Mahalanobis distance (0.58 for both), with kernel density and nearest neighbour distance still performing poorly (Fig. 5). This final example illustrates the potential of Mahalanobis and harmonic mean distance to improve the signal-to-noise ratio in genome scans, because the empirical power in this case was higher than any univariate statistic alone.

Discussion

Although the number of packages for population genetic data analysis in the R software is rapidly increasing (<http://popgen.nescent.org/PACKAGES.html>), basic

tools for manipulating and visualizing genome-scale data sets have so far been lacking. MINOTAUR fills this gap using the R Shiny Dashboard package to implement a GUI that makes it easy to upload, manipulate, analyse and visualize genomic data.

The multivariate metrics calculated in MINOTAUR contribute to a growing number of multivariate tools implemented in the R environment (see Table S1, Supporting information). Methods that are influenced heavily by the distance of a point from the centroid in multivariate space (such as Mahalanobis and the harmonic mean distance) will perform differently compared with methods that are influenced mainly by the sparseness of points in multivariate space (such as nearest neighbour distance and kernel density), as illustrated in the examples here. However, depending on how the data are distributed, the harmonic mean distance may be influenced by both these factors. For a single simulated data set, we found that robust use of the Mahalanobis or harmonic mean distance (i.e. when the covariance matrix used was estimated with a high breakdown point) could have higher power than any single univariate statistic alone. Although nearest neighbour distance and kernel density deviance performed poorly on the simulated genomic data, they may be useful in application to other kinds of nonparametric data, as illustrated in our examples

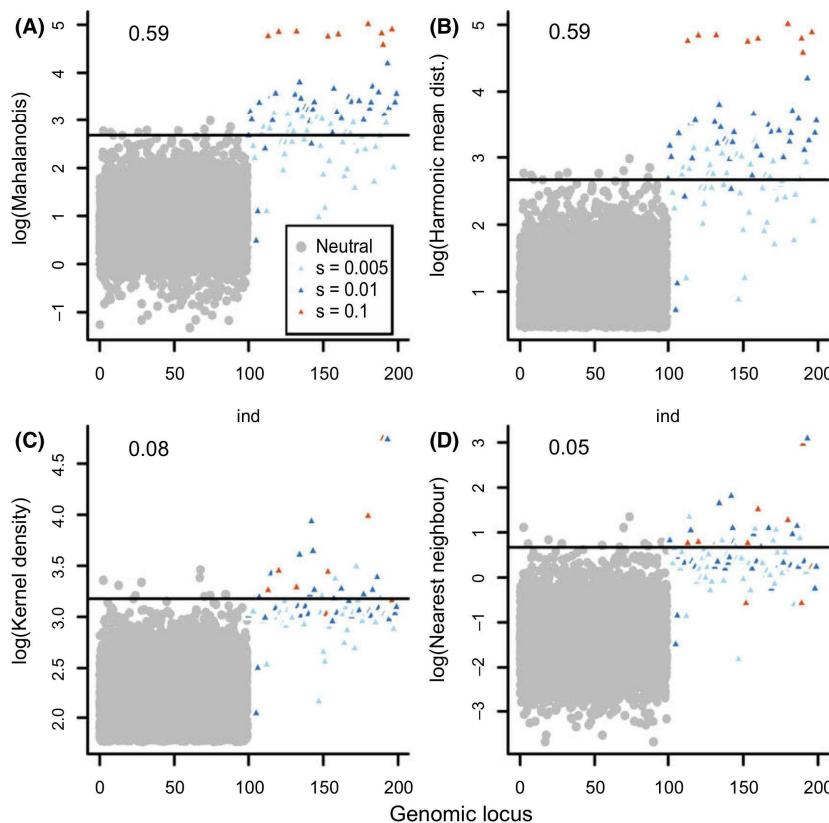


Fig. 5 Distributions of the four multivariate compound statistics applied to the four univariate statistics shown in Fig. 2. The MCD calculation of the covariance matrix was used. All 9900 neutral loci are plotted on indexes 0–100, and the selected loci are plotted on indexes 100–200. Note log transformation of each variable on the y-axis for (A) Mahalanobis distance, (B) harmonic mean distance, (C) kernel density and (D) nearest neighbour distance. The empirical power of the statistic to discriminate neutral from selected loci (see main text for details) is shown in the upper left-hand corner.

(Figs 3 and S1, Supporting information). Further evaluation, however, will be needed on both simulated and empirical data to determine whether multivariate outlier approaches will improve the signal-to-noise ratio in genome scans.

The MINOTAUR package is designed to complement existing tools for the analysis and integration of genome scan data. Thus, in addition to providing its own tools for genome-scale analyses, MINOTAUR can serve as a platform for the further analysis and visualization of results generated by other R packages. Examples include results from differential gene expression (LIMMA: Ritchie *et al.* 2015; DESEQ: Anders & Huber 2010; SEQGSEA: Wang & Cairns 2014), outliers for genetic differentiation (OUTFLANK: Whitlock & Lotterhos 2015; PCADAPT: Luu & Blum 2015), genetic–environment associations (LEA: Fritchot & François 2015) or genomewide association studies (e.g. GENABEL: Aulchenko *et al.* 2007; BLUESNP: Huang *et al.* 2013).

Recent developments such as the R Shiny and Shiny Dashboard environments (Chang 2015; Chang *et al.* 2016) dramatically aid in the development of R-based user-friendly web interfaces. Taking advantage of these tools, MINOTAUR is able to offer a new platform for visualizing and integrating genomic data that may appeal to molecular ecologists, modellers, statisticians and public health agencies.

Resources

Availability: Upon acceptance for publication, MINOTAUR will be distributed on CRAN (<http://cran.r-project.org/>) and be available for R on Windows, MAC OSX and LINUX platforms. Currently, MINOTAUR can be accessed via the following steps:

- `install.packages("devtools", dependencies =TRUE)`
- `library(devtools)`
- `install_github("NESCent/MINOTAUR", build_vignettes=TRUE)`
- `library(MINOTAUR)`
- `MINOTAUR()`

Note to reviewers: If you are facing issues with installation, try updating to the newest version of R and re-installing devtools from source. MINOTAUR has been tested on R version 3.3.0.

Licence: GNU General Public Licence (GPL) ≥ 2 .

Documentation: Besides the usual package documentation, MINOTAUR is released with a tutorial which can be opened by typing `VIGNETTE("MINOTAUR")`.

Development: The development of MINOTAUR is hosted on GitHub: (<https://github.com/NESCent/MINOTAUR>).

Acknowledgments

The resource reported in this study began at the Population Genetics in R Hackathon, which was held in March 2015 at the National Evolutionary Synthesis Center (NESCent) in Durham, NC, with the goal of addressing interoperability, scalability and workflow building challenges for the population genetics package ecosystem in R. The authors were participants in the hackathon and are indebted to NESCent (NSF #EF-0905606) for hosting and supporting the event. RV, CC, DCC, SMS benefited from travel support to attend this Hackathon.

References

Allendorf FW, Hohenlohe PA, Luikart G (2010) Genomics and the future of conservation genetics. *Nature Reviews Genetics*, **11**, 697–709.

Anders S, Huber W (2010) Differential expression analysis for sequence count data. *Genome Biology*, **11**, 1–12.

Aulchenko YS, Ripke S, Isaacs A, Duijn CMV (2007) GENABEL: an R library for genome-wide association analysis. *Bioinformatics*, **23**, 1294–1296.

Brachi B, Morris GP, Borevitz JO (2011) Genome-wide association studies in plants: the missing heritability is in the field. *Genome Biology*, **12**, 232.

Carlson CS, Eberle MA, Kruglyak L, Nickerson DA (2004) Mapping complex disease loci in whole-genome association studies. *Nature*, **429**, 446–452.

Chang W (2015) shinydashboard: Create Dashboards with ‘Shiny’. R package version 0.5.1. <https://CRAN.R-project.org/package=shinydashboard>

Chang W, Cheng J, Allaire JJ, Xie Y, McPherson J (2016) shiny: Web Application Framework for R. R package version 0.13.1. <https://CRAN.R-project.org/package=shiny>

Dabney A, Storey JD (2014) qvalue: Q-value estimation for false discovery rate control. R package, version 1.38.0. <https://github.com/jd-storey/qvalue>.

De Mita S, Thuillet AC, Gay L et al. (2013) Detecting selection along environmental gradients: analysis of eight methods and their effectiveness for outbreeding and selfing populations. *Molecular Ecology*, **22**, 1383–1399.

De Villemereuil PD, Frichot É, Bazin É, François O, Gaggiotti OE (2014) Genome scan methods against more complex models: when and how much should we trust them? *Molecular Ecology*, **23**, 2006–2019.

Eddelbuettel D (2013) *Seamless R and C++ Integration with Rcpp*. Springer, New York. ISBN 978-1-4614-6867-7.

Eddelbuettel D, François R (2011) Rcpp: Seamless R and C Integration. *Journal of Statistical Software*, **40**, 1–18.

Evangelou E, Ioannidis JP (2013) Meta-analysis methods for genome-wide association studies and beyond. *Nature Reviews Genetics*, **14**, 379–389.

Frichot E, François O (2015) LEA: An R package for landscape and ecological association studies. *Methods in Ecology and Evolution*, **6**, 925–929.

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

Funk WC, McKay JK, Hohenlohe PA, Allendorf FW (2012) Harnessing genomics for delineating conservation units. *Trends in Ecology & Evolution*, **27**, 489–496.

Galesloot TE, van Steen K, Kiemeney LALM, Janss LL, Vermeulen SH (2014) A comparison of multivariate genome-wide association methods. *PLoS ONE*, **9**, 1–8.

Goddard ME, Hayes BJ (2009) Mapping genes for complex traits in domestic animals and their use in breeding programmes. *Nature Reviews Genetics*, **10**, 381–391.

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

Hindorff LA, Sethupathy P, Junkins HA et al. (2009) Potential etiologic and functional implications of genome-wide association loci for human diseases and traits. *Proceedings of the National Academy of Sciences*, **106**, 9362–9367.

Hoban S, Kelley JK, Lotterhos KE et al. (in press) Finding the genomic basis of local adaptation: pitfalls, practical solutions and future directions. *American Naturalist*.

Hohenlohe PA, Phillips PC, Cresko WA (2010) Using population genomics to detect selection in natural populations: key concepts and methodological considerations. *International Journal of Plant Sciences*, **171**, 1059–1071.

Huang H, Tata S, Prill RJ (2013) BLUESNP: R package for highly scalable genome-wide association studies using Hadoop clusters. *Bioinformatics*, **29**, 135–136.

Kohn MH, Murphy WJ, Ostrander EA, Wayne RK (2006) Genomics and conservation genetics. *Trends in Ecology & Evolution*, **21**, 629–637.

Leiva-Murillo JM, Artés-Rodríguez A (2012) Algorithms for maximum-likelihood bandwidth selection in kernel density estimators. *Pattern Recognition Letters*, **33**, 1717–1724.

Lotterhos KE, Whitlock MC (2014) Evaluation of demographic history and neutral parameterization on the performance of FST outlier tests. *Molecular Ecology*, **23**, 2178–2192.

Lotterhos KE, Whitlock MC (2015) The relative power of genome scans to detect local adaptation depends on sampling design and statistical method. *Molecular Ecology*, **24**, 1031–1046.

Luu K, Blum MGB (2015) pcadapt: Principal component analysis for outlier detection. R package version 2.1. <https://CRAN.R-project.org/package=pcadapt>

Mahalanobis PC (1936) On the generalised distance in statistics. *Proceedings of the National Institute of Sciences of India*, **2**, 49–55.

McCarthy MI, Hirschhorn JN (2008) Genome-wide association studies: past, present and future. *Human Molecular Genetics*, **17**, R100–101.

McCarthy MI, Abecasis GR, Cardon LR et al. (2008) Genome-wide association studies for complex traits: consensus, uncertainty and challenges. *Nature Reviews Genetics*, **9**, 356–369.

O'Reilly PF, Hoggart CJ, Pomyen Y et al. (2012) MULTIPHEN: joint model of multiple phenotypes can increase discovery in GWAS. *PLoS ONE*, **7**, 1–12.

Patterson N, Price AL, Reich D (2006) Population structure and eigen-analysis. *PLoS Genetics*, **2**, 2074–2093.

R Core Team (2015) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>

Rellstab C, Gugerli F, Eckert AJ, Hancock AM, Holderegger R (2015) A practical guide to environmental association analysis in landscape genomics. *Molecular Ecology*, **24**, 4348–4370.

Ritchie ME, Phipson B, Wu D et al. (2015) LIMMA powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Research*, **43**, 1–13.

Rousseeuw PJ, Van Driessche KV (1999) A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, **41**, 212–223.

Savolainen O, Lascoux M, Merilä J (2013) Ecological genomics of local adaptation. *Nature Reviews Genetics*, **14**, 807–820.

Silverman BW (1986) *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London.

Tiffin P, Ross-Ibarra J (2014) Advances and limits of using population genetics to understand local adaptation. *Trends in Ecology & Evolution*, **29**, 673–680.

Todorov V, Templ M, Filzmoser P (2011) Detection of multivariate outliers in business survey data with incomplete information. *Advances in Data Analysis and Classification*, **5**, 37–56.

Varshney RK, Nayak SN, May GD, Jackson SA (2009) Next-generation sequencing technologies and their implications for crop genetics and breeding. *Trends in Biotechnology*, **27**, 522–530.

Vatsiou AI, Bazin E, Gaggiotti OE (2016) Detection of selective sweeps in structured populations: a comparison of recent methods. *Molecular Ecology*, **25**, 89–103.

Verity R, Nichols RA (2014) What is genetic differentiation, and how should we measure it—GST, D, neither or both? *Molecular Ecology*, **23**, 4216–4225.

Wang X, Cairns MJ (2014) SeqGSEA: a Bioconductor package for gene set enrichment analysis of RNA-Seq data integrating differential expression and splicing. *Bioinformatics*, **30**, 1777–1779.

Wang Z, Gerstein M, Snyder M (2009) RNA-Seq: a revolutionary tool for transcriptomics. *Nature Reviews Genetics*, **10**, 57–63.

Whitlock MC, Lotterhos KE (2015) Reliable detection of loci responsible for local adaptation: inference of a null model through trimming the distribution of FST. *The American Naturalist*, **186**, S24–36.

Yang J, Benyamin B, McEvoy BP *et al.* (2010) Common SNPs explain a large proportion of the heritability for human height. *Nature Genetics*, **42**, 565–569.

R.V. and K.E.L. conceptualized the study. R.V. derived the compound outlier measures and implemented them in R. R.V. and C.C. conceptualized and implemented the Shiny Dashboard GUI. C.C. managed R package development, package structure and documentation. All authors contributed R code for plotting in the Shiny Dashboard. D.C.C. wrote the vignette for the package and estimated function computational time. K.E.L. performed analysis of the performance of compound measures on simulated data. All authors contributed to writing this manuscript.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Table S1 Table of multivariate outlier statistics in other R packages that could be used in the context of genomic scans.

Table S2 Computation times for the four multivariate outlier detection methods in MINOTAUR for datasets up to 100 000 loci (rows) and 20 variables (columns) in hh:mm:ss.ms format. Run times were determined using an Apple iMac with a 3.1 GHz Intel Core i5 processor and 32 GB of RAM running Apple OSX 10.9.5 and R version 3.2.3. Note that the kernel density deviance includes both the maximum likelihood estimation of the optimal bandwidth and the density calculations based on the optimal bandwidth.

Fig. S1 Comparison of multivariate distance measures for multi-modal example data.

Fig. S2 Analogue to Fig. 5 in the main paper, but with a default estimate of covariance using all the data.