American Genetic Association

Journal of Heredity, 2018, 16–28
doi:10.1093/jhered/esx042
Symposium Article
Advance Access publication 8 May 2017

OXFORD

Symposium Article

Effects of Gene Action, Marker Density, and Timing of Selection on the Performance of Landscape Genomic Scans of Local Adaptation

Jeremy B. Yoder and Peter Tiffin

From the Department of Forest and Conservation Genetics, University of British Columbia, Vancouver, BC, Canada (Yoder); and Department of Plant Biology, University of Minnesota, Saint Paul, MN 55108 (Tiffin).

Address correspondence to P. Tiffin at the address above, or e-mail: ptiffin@umn.edu.

Received February 13, 2017; First decision March 27, 2017; Accepted May 5, 2017.

Corresponding Editor: Lynda Delph

Abstract

Genomic "scans" to identify loci that contribute to local adaptation are becoming increasingly common. Many methods used for such studies have assumed that local adaptation is created by loci experiencing antagonistic pleiotropy (AP) and that the selected locus itself is assayed, and few consider how signals of selection change through time. However, most empirical data sets have marker density too low to assume that a selected locus itself is assayed, researchers seldom know when selection was first imposed, and many locally adapted loci likely experience not AP but conditional neutrality (CN). We simulated data to evaluate how these factors affect the performance of tests for genotype–environment association (GEA). We found that 3 types of regression-based analyses (linear models, mixed linear models, and latent factor mixed models) and an implementation of BayEnv all performed well, with high rates of true positives and low rates of false positives, when the selected locus experienced AP, and when the selected locus was assayed directly. However, all tests had reduced power to detect loci experiencing CN, and the probability of detecting associations was sharply reduced when physically linked rather than causative loci were sampled. AP also maintained detectable GEAs much longer than CN. Our analyses suggest that if local adaptation is often driven by loci experiencing CN, genome-scan methods will have limited capacity to find loci responsible for local adaptation.

Subject area: Genomics and gene mapping
Key words: antagonistic pleiotropy, BayEnv, BayPass, conditional neutrality, genotype–environment association, latent-factor mixed models, selective sweeps

Genomic data enable studies that use population genetic analyses to identify genes responsible for local adaptation. Identifying the genetic basis of local adaptation can advance our understanding of selective constraints, the genetic basis of naturally occurring variation, how geographically variable selection has shaped genomic diversity, and can provide insight into the tempo and mode of local adaptation (Savolainen et al. 2013; Tiffin and Ross-Ibarra 2014).

The most commonly used statistical approaches for identifying locally adapted loci use F_{ST} or other variance partitioning methods to identify loci with elevated among-population divergence (Lewontin and Krakauer 1973; Beaumont and Nichols 1996; Vitalis et al. 2001; Excoffier et al. 2009b; Bonhomme et al. 2010; Günther and Coop 2013); or they test for strong genotype–environment associations (GEA) (Joost et al. 2007, 2008; Poncet et al. 2010).

© The American Genetic Association 2017. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

Evaluating the performance of these tests under biologically realistic scenarios is important for understanding the analysis of empirical datasets, and a number of recent studies have used simulated data to achieve this goal (De Mita et al. 2013; Jones et al. 2013; De Villemereuil et al. 2014; Lotterhos and Whitlock 2014, 2015; Frichot et al. 2015). DeMita et al. (2013) compared the statistical performance of regression models and F_{ST}-type tests to identify targets of local adaptation under various sampling schemes for both selfing and outcrossing species; Jones et al. (2013) compared different regression and classification-based tests for GEA under different strengths of selection; Lotterhos and Whitlock (2014) examined various F_{ST}-type tests for identifying targets of selection in stable and expanding populations with different population structures. Finally, De Villemereuil et al. (2014) tested the robustness of GEA and F_{ST}-like tests to hierarchical population structure, and selection acting on highly polygenic traits.

These simulation studies, together with an extensive earlier literature (e.g., Teshima et al. 2006; Jensen et al. 2007; Meirmans 2012), have identified the potential pitfalls and power of these tests. For example, when the strength of selection is weak or population structure co-varies with environmental variation that drives local adaptation, tests can be plagued by false positives (Meirmans 2012; De Mita et al. 2013; De Villemereuil et al. 2014). Similarly, the power of local adaptation tests can be highly dependent on demographic history, with both false positives and false negatives being more likely in expanding populations (Lotterhos and Whitlock 2014, 2015) or when the underlying assumptions of the test do not fit the population structure to which they are applied (Excoffier et al. 2009a; De Mita et al. 2013; Frichot et al. 2015; Lotterhos and Whitlock 2015). The power of landscape genomic tests can also depend on sampling scheme. De Mita et al. (2013) showed that when populations are distributed across a continuous environmental gradient, linear-model-based tests perform well when sampling a single individual from each of many populations; but can have very high false positive rates when sampling is concentrated in a few discrete populations. When sampling a few discrete populations, methods that model allelic diversity within populations can provide better performance. Moreover, Lotterhos and Whitlock (2015) found that GEA tests and F_{ST}-like tests had more power to detect weakly selected loci when samples were drawn from pairs of sites in contrasting environments, rather than collected along transects spanning environmental gradients.

The biological relevance of simulation studies depends on the assumptions made when simulating data. Here, we examine several assumptions made in previous analyses that may not correctly capture the reality of empirical data. The first of these assumptions is that, at a biallelic locus under spatially variable selection, one allele will be favored in one part of the range and the alternative allele will be favored in the other part of the range—that is, the phenotypic effects of the alleles exhibit antagonistic pleiotropy (AP) for fitness in different environments (e.g., De Mita et al. 2013; de Villemereuil et al. 2014; Lotterhos and Whitlock 2015). Empirical studies have found some evidence that alleles responsible for local adaptation show tradeoffs across contrasting environments (Anderson et al. 2011; Ågren et al. 2013). However, it also is clear that many alleles contributing to local adaptation experience conditional neutrality (CN), contributing to higher (or lower) fitness in one environment while having no fitness effect in other locations (Schnee and Thompson 1984; Thompson 2005; Gardner and Latta 2006; Verhoeven et al. 2008; Lowry et al. 2009; Fournier-Level et al. 2011; Anderson et al. 2013, 2014a; Yoder et al. 2014). We do not know how often local adaptation in natural populations evolves by AP or under CN. A review of phenotypic studies of local adaptation conducted by Anderson et al. (2011) found greater evidence for CN than AP, but also caution that there may be an ascertainment bias in favor of finding CN using paired tests of selection in contrasting environments. The reason for this bias is that experiments with limited statistical power are less likely to detect signals of selection in two environments (as is required to conclude AP) rather than only one environment. Loci identified as responsible for adaptation in only one environment may be CN alleles, or may be AP loci that have not been properly characterized.

Contrasting selection in different parts of a species range is expected to maintain variation at AP loci. By contrast, CN alleles are expected to be driven to fixation in one part of the range by selection but will eventually fix throughout the range through via gene flow and drift. For this reason, CN alleles may harbor weaker signals of local adaptation than AP loci. Moreover, the signals of local adaptation at CN loci may be expected to decay through time—as the allele that is favored in one part of the range drifts to fixation in the non-selected portion of the range.

Previous simulation studies of local adaptation scans also have assumed that the causative locus, or a marker in complete linkage disequilibrium (LD) with it, is among those sampled (e.g., De Mita et al. 2013; Lotterhos and Whitlock 2014). Whole-genome sequence data for samples from multiple populations or from across a species range, which are necessary to ensure that causal loci are likely to be sampled, are available for relatively few species. Even for these, the complexity of mapping sequence reads to reference genomes means that whole-genome "resequencing" is unlikely to produce complete coverage of all variable sites. Empirical genome scans rarely assume that they are sampling causal loci, but rather that they are at best able to find markers or regions in close linkage with causal variants (Savolainen et al. 2013). Marker density is important because as marker density decreases, LD between assayed markers and the true targets of selection will weaken, and with it the ability to identify those targets of selection (Caballero et al. 2012; Tiffin and Ross-Ibarra 2014). Therefore, simulations that assume sampling of causal variants will overestimate the sensitivity of landscape genomics methods.

Tests of species-wide adaptation have much less power to detect selection that has acted on standing variation than selection acting on a new mutation (Hermisson and Pennings 2005; Przeworski et al. 2005). Selection on standing variation may be particularly important in the evolution of CN alleles, because they may drift in the neutral part of the range before migrating into the selected part of the range. Selection on standing variation may also greatly weaken the power to detect selection acting on linked loci, given that the further away an assayed marker is from the selected variant, the more opportunity there will be for recombination and the breakdown of LD.

The objective of this study is to better understand what landscape genomic scans can reveal about the genetics of local adaptation under scenarios that likely reflect the biological reality of many systems and datasets. Specifically, we 1) evaluate the power of 4 methods commonly used for genomic scans of local adaptation to detect AP and CN loci, and when selection acted on either new mutations or standing variation; 2) examine the effects of time between the onset of selection and the time at which a population is sampled, and 3) look for a population-genetic signal in linked loci that might be used to differentiate AP from CN loci and selection acting on standing variation or new mutations.

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

Methods

We used QuantiNEMO (release 1.0.4; Neuenschwander et al. 2008) to simulate a landscape of 100 populations arrayed in a 10 × 10 grid, each containing 100 diploid, hermaphroditic individuals. Each simulated individual carries a number of biallelic loci, arranged along a single set of homologous chromosomes; one of these loci is under geographically varying selection as described below, while the others are selectively neutral. QuantiNEMO simulates nonoverlapping generations, with mating occurring within populations, between pairs of individuals drawn at random with replacement. The number of offspring produced by each mating was determined by a single biallelic locus. We set the strength and direction of selection on this locus to vary across the landscape according to 1 of 4 geographic patterns (Figure 1). After reproduction, offspring migrated to adjacent populations with a probability of 0.01, in a "stepping-stone" model of migration. We ran simulations for 10 000 generations and analyzed results at 1, 2, 4, 6, and 10 thousand generations after selection was imposed. The QuantiNEMO configuration files, scripts used for analysis, and the output of analyses are available in the Dryad Digital Repository, http://dx.doi.org/10.5061/dryad.p12q4).

Simulated Genotypes

We imposed selection on either new mutations or on standing variation. In stochastic forward simulations such as QuantiNEMO, creating a true hard sweep with an origin in a single sequence is challenging, because a single allele will often be lost to drift, even when it has a strong selective advantage. When variation at a locus is lost, QuantiNEMO reintroduces variation by forcing a single mutation event in the next generation. To achieve adaptation from new mutations, then, we started simulations with one allele at each selected locus having a frequency of 2×10^{-4} in the half of the landscape where it was selectively favored, and kept only simulations in which this allele was either not lost or lost and reintroduced as a single new mutation within the first 1000 generations of simulation. For selection on standing variation, we set global starting allele frequencies based on expectations from a standard neutral model at equilibrium (Ganapathy and Uyenoyama 2009). To evaluate the effects of adaptation on genetic diversity in regions linked to selected loci, we simulated 4 clusters with 20 neutral biallelic loci each, at recombination distances of 0, 0.01, 0.1, and 1 centimorgans (cM) away from the selected locus, plus another 20 neutral loci on a separate chromosome. For these linked loci, we drew starting

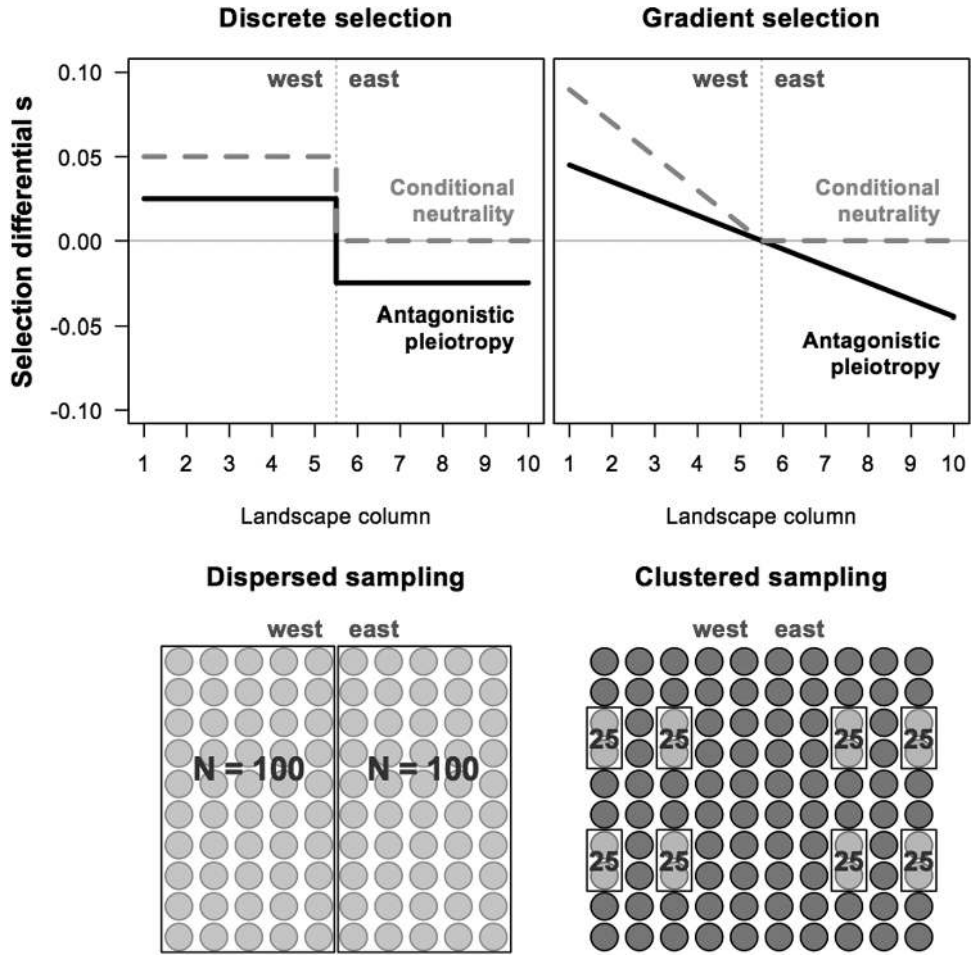

Figure 1. The models of geographically varying selection we simulated. Top panel: We conducted all simulations in a landscape consisting of a 10 × 10 array of populations linked by migration in a stepping stone pattern. The selection differential s acting on the A allele of the selected locus changed from the western half of the landscape to the eastern half. For discrete selection, the change was discontinuous (left graph); for gradient selection, s varied continuously from the westernmost column of populations to the easternmost column (right graph). Bottom panel: For all analyses, we used samples of genotypes drawn from the simulated populations in 1 of 2 ways. For "dispersed" samples, we drew 100 genotypes at random from across the entire western half of the landscape, and 100 from the east; in "clustered" samples, we drew 25 genotypes from each of 8 nonadjoining pairs of populations, highlighted.

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

allele frequencies from the equilibrium distribution of frequencies for a standard neutral model, as for the selected locus understanding variation. Because QuantiNEMO assembles starting populations by creating individuals' multi-locus genotypes at random based on the specified starting allele frequencies, these loci began each simulation in complete linkage equilibrium (i.e., LD, as measured by $r^2 = 0$) with the selected locus.

Selection

We ran 100 simulations each for each of 8 possible combinations of (1) new mutations or standing variation experiencing either (2) AP or CN that varied across the landscape as (3) either 2 discrete environments or a gradient (DS or GR; Figure 1). As a standard for comparison, we also ran 100 simulations each with no selection (neutral, or NT), starting from either new mutations or standing variation. We parameterized the directional selection option in QuantiNEMO to create selection differentials that conformed to each of these forms of selection. Given 3 possible genotypes at a biallelic focal locus, AA, AB, and BB, we defined the selection differential s so that the relative fitness of each genotype, w_{AA}, w_{AB}, and w_{BB} was equal to 1, $1 - s$, and $1 - 2s$, respectively. We chose selection differentials that resulted in clinal distributions of allele frequencies in each population after 2000 generations in trial simulations; QuantiNEMO's output of global F_{ST} in these trials also indicated that 5000 generations were usually sufficient for the landscape to achieve mutation-drift-selection-migration equilibrium.

To create discrete and gradient selective landscapes, we varied the selection differential s in different columns of the 10 × 10 landscape grid (Figure 1). For discrete (DS) selection on AP alleles, we set the selection differential (s) to 0.025 in the 5 "western" columns of populations in the 10 × 10 landscape grid, $s = -0.025$ in the 5 "eastern" columns. For DS on CN alleles, we set $s = 0.05$ in the 5 western columns, and $s = 0$ in the 5 eastern columns. For gradient (GR) selection on AP, we varied the value of s linearly from 0.05 in the westernmost column to −0.05 in the easternmost column; for GR selection acting on CN alleles, we varied s from 0.10 in the first column to 0 in the sixth column, and set $s = 0$ for the remaining 4 columns (Figure 1). The different maximum values of s in each of these 4 scenarios ensured that the selected alleles experienced roughly the same strength of selection, averaged across the entire landscape, in all scenarios.

Analysis

We analyzed all QuantiNEMO output using the R statistical computing environment (R Core Team 2014). For each time point (1, 2, 4, 6, and 10 thousand generations after selection commenced) we sampled 200 diploid individuals (400 alleles) according to 1 of 2 sampling schemes intended to capture the sampling used in typical landscape genomic studies. In the first, which we term "dispersed sampling," we drew 100 individuals at random from each side of the landscape, without respect to their population of origin (Figure 1). We also examined "clustered sampling" in which we drew 25 individuals from each of 4 pairs of adjacent populations in the center of the western and eastern halves of the landscape (Figure 1). This was intended to introduce cryptic population structure into the analysis, and it approximates sampling schemes used in studies that attempt to estimate local allele frequencies by sampling multiple individuals at each site (e.g., Eckert et al. 2010; Cheng et al. 2012; Jones et al. 2012; Pyhäjärvi et al. 2013). For all population sampling, we checked that the selected locus was polymorphic in the sample and took new samples if this was not the case. In all downstream analysis, we only tested for signals of local adaptation at loci for which the frequency of the less-common allele was greater than 0.05, following standard practice in empirical studies, which typically exclude loci with rare minor alleles as part of genotyping procedures or later analysis.

GEA Tests

We conducted 4 tests for GEA, each with different forms of control for the confounding effect of population structure. In all GEA tests, the "environment" was the landscape column number (Figure 1). First, we fitted a simple linear model (LM) using genotype as a predictor of the environment from which individuals were sampled (i.e., environment–genotype). This method had no control for population structure, and served primarily as a standard for comparison to the other methods.

The second GEA method was the mixed linear model (MLM) framework implemented in the programs TASSEL and GAPIT (Yu et al. 2006; Zhang et al. 2010; Lipka et al. 2012), which is widely used for genome-wide association studies, and which we applied in a previous landscape genomic analysis (Yoder et al. 2014). The MLM approach fits a model for each locus with the observed allele as the predictor variable and environment as the response (i.e., as in the LM, environment–genotype), with a pairwise kinship (K-) matrix as a covariate to control for confounding effects of population structure. We calculated the K-matrix as a similarity matrix, using genotypes at 1000 unlinked neutral loci for 200 individuals simulated separately under the same demographic history as the genotypes used for analysis (i.e., equivalent to loci from additional chromosomes beyond the one carrying the selected locus in the simulations of local adaptation).

We also tested for association using the latent factor mixed model (LFMM) framework of Frichot et al. (2013), as implemented in the R package LEA (http://membres-timc.imag.fr/Olivier.Francois/LEA). This method uses a Markov chain Monte Carlo algorithm to fit a LFMM, in which individual genotypes are predicted by a mixed model with an environmental variable as a fixed effect and population structure modeled as a set of K latent factors. That is, LFMMs fit a model of the form genotype–environment whereas LM and MLM methods fit a model of the form environment–genotype. Following Frichot et al. (2013) and the procedure outlined in the LEA manual, we wrote a script to run LFMM analysis using utilities in the LEA package. As with the MLM test, we used genotypes from 1000 unlinked neutral loci simulated independently under the same demographic conditions, and sampled in the same manner, as the genotypes used for analysis, to estimate appropriate values of K for each simulation scenario. Our script then fit LFMMs with $K - 1$ through $K + 2$ latent factors (or 1 through $K + 2$ if the initial approximation found $K = 1$), and selected the fitted LFMM with an optimal genomic inflation value (Devlin and Roeder 1999). The script fit LFMMs by running the MCMC algorithm for 10^{-4} iterations, with the first 5000 iterations discarded as burn-in. For each locus, the script recorded the value of K for the optimal LFMM, the raw z-score from the LFMM, a P-value adjusted for the genomic inflation value, and whether or not that P-value achieved statistical significance at $P < 0.05$ with correction for false discovery (Benjamini and Hochberg 1995).

Finally, we tested for association using the BayEnv method (Günther and Coop 2013) as implemented in the program BayPass (Gautier 2015). The BayEnv method calculates among-population

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

genetic differentiation from allele count data, and tests for association between genetic differentiation at each locus and between-site environmental differences, controlling for the effects of among-population covariance in allele frequencies, using an allele-frequency covariance matrix calculated from putatively neutral control loci (Günther and Coop 2013). We wrote a script in R to run the BayEnv analysis on our simulated data. The script first called utilities provided in BayPass to estimate an allele frequency covariance matrix based on genotypes from 1000 unlinked neutral loci simulated under the same demographic conditions, and sampled in the same manner, as the genotypes used for analysis. It then called BayPass to run the BayEnv analysis, using the estimated covariance matrix, to estimate the strength of association between allele frequencies at each locus and the environment.

We calculated the sensitivity of each GEA test as the proportion of simulations in which the selected locus was positively identified as having a greater GEA than expected by chance, and the false positive rate as the proportion of simulations in which a focal neutral locus was identified as having a greater GEA than expected by chance. For the LM, MLM, and LFMM tests, we expressed the probability of greater-than-expected GEA returned by each test as a LOD score (the $-\log_{10}(P)$) and assessed sensitivity at LOD cutoff values up to 9 (or $P \leq 10^{-9}$). The BayEnv test implemented in BayPass expresses confidence in GEA for each locus not as a *P*-value, but as a Bayes factor (BF) scaled in dB units equal to $10 \times \log_{10}(BF)$; we assessed BayEnv sensitivity using cutoff values up to 9 for this scaled Bayes factor as well.

Population Genetic Statistics

Because recombination and drift are expected to be differ between AP and CN loci and also between selection acting on new alleles and standing variation, we calculated 2 population genetic parameters that may provide some ability to differentiate between these types of alleles and differences in selection. First, we calculated the H12 index of haplotype frequency distortion (Messer and Petrov 2013; Garud et al. 2014). H12 is calculated from the ranked frequencies of the n unique haplotypes in a sample, as $(f_1 + f_2)^2 + f_3^2 + \ldots + f_n^2$, where f_i is the frequency of the ith-most-common haplotype (Garud et al. 2014). H12 varies between 0 and 1, where values approaching 1 indicate a population dominated by 1 or 2 haplotypes, and lower values indicate a population with more haplotypes at low to intermediate frequencies. H12 provides a haplotype-frequency-based test for a recent selective sweep that can be robust to scenarios in which a selectively favored allele occurs on multiple haplotype backgrounds, which may often be suitable for putative cases of local adaptation (Yoder et al. 2014). We calculated H12 on haplotypes assembled from each individual's genotype at the selected locus and 10 loci drawn randomly from the set of 20 neutral loci simulated at each recombination distance, to provide estimates of those statistics given different degrees of physical linkage to the selected locus. To directly evaluate LD between the selected locus and the linked loci, we calculated r^2, a measure of LD that is equivalent to the correlation coefficient between the allelic state at one locus and the allelic state at a second locus.

Results

In simulations of selection on a new mutation, AP alleles usually fixed within the part of the range where they were selectively favored (i.e., achieved a global frequency of 0.5) within 2000 generations, though there was heterogeneity due to the loss, through drift, of alleles with very low initial frequency (Figure 2). Meanwhile CN alleles that began as new mutations almost always fixed in the selected part of the range (the western half) within 1000 generations. The faster rate of fixation for CN alleles reflects the stronger local selection acting on CN alleles within the selected region (note, across the range the strength of selection acting on CN and AP alleles is equal; Figure 1). Following local fixation, CN alleles continued to increase in global frequency, as gene flow from the west increased allele frequency in the eastern part of the landscape (Figure 2). Evolution from standing variation resulted in more rapid fixation of the selected allele in all cases.

Detecting Targets of Local Adaptation

Under the conditions we simulated, the ability to identify an allele responsible for local adaptation was dependent on the statistical test applied, whether alleles at the gene were conditionally neutral or antagonistically pleiotropic, whether the selected SNP or a physically linked variant was assayed, and the time of sampling after selection was imposed. The ability to detect selection also was affected by whether selection acted on a new mutation or standing variation and whether the imposed selection was discrete or acted on a gradient, but for these factors the effects were either minor or very limited (Figure 3). The only case in which power to detect selection on a new mutation differed strongly from power to detect selection on standing variation was for BayPass analyses on clustered samples (Figure 3H); the only case in which power to detect local adaptation on a discrete versus a gradient landscape differed strongly was for LFMM analyses on dispersed samples (Figure 3C).

The 4 GEA tests we applied include 3 variations on linear model tests of association: a simple liner model with no correction for the confounding effect of population structure (LM); a mixed linear model (MLM) method implemented in GAPIT (Lipka et al. 2012), which uses a kinship matrix covariate to control for effects of population structure, and LFMM method (Frichot et al. 2013) that incorporates latent factors to model population structure. The fourth test we applied is the BayEnv method (Günther and Coop 2013), as implemented in the program BayPass (Gautier 2015), which identifies putative targets of local adaptation as those loci with unusually strong among-population covariance between allele frequency and environmental distance, given the genome-wide distribution of those covariances. Both the LM and MLM approaches had high power to detect true positives (at or near 100% for all selected loci in the LM tests and for all but selection acting on CN standing variation for the MLM tests; Figures 3 and 4). The LFMM test had high power to detect AP loci under all circumstances except for dispersed samples taken from simulations of selection on a new mutation; and it had consistently lower power to detect CN loci under all circumstances. From the perspective of true positives, the BayEnv/BayPass test had the poorest performance, presumably because it overcorrected for population structure. BayPass did, however, have similar true positive rates to the 3 linear model-based approaches (and with far lower false positive rates than LM) when selection acted on new mutations and the samples were clustered (i.e., 25 sequences sampled from each of 4 proximate subpopulations on each side of the landscape; Figure 1).

False Positive Rates

The performance of statistical tests is dependent not only on the power of those tests to identify true positives, but also their ability

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

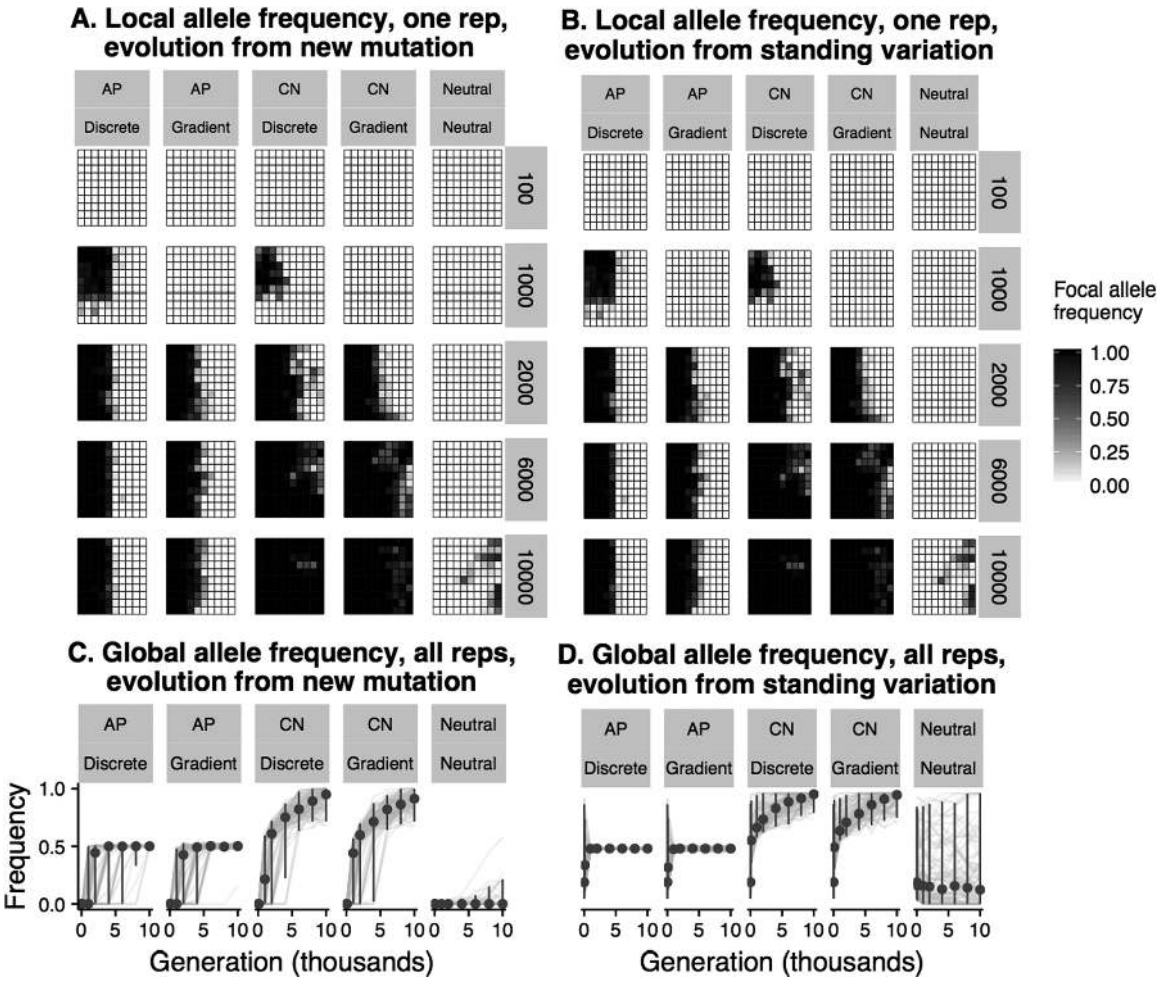

Figure 2. Evolution of allele frequency at the selected locus. (**A**, **B**) Image plots illustrating the frequency of the focal allele in each of the 10 × 10 populations in the simulated landscape (one cell per site) in one replicate simulation of each of the geographic selection scenarios outlined in Figure 1, with evolution from a new mutation or standing variation. (**C**, **D**) Plots of the landscape-wide frequency of the focal allele over time, for all replicates of each selection scenario (light gray lines), with the mean frequency and 95% density interval superimposed at each time point (dark points and error bars).

to minimize false positives. This is especially important when analyzing a large number of sites, most of which are not expected to have a direct role in adaptation, in a typical "genome scan" approach. We assessed false positives as the proportion of neutral simulations in which each of the 4 GEA tests identified a significant association. Because neutral simulations starting from new mutations rarely evolved sufficiently high allele frequency to pass our minor allele frequency filter (MAF > 0.05), we used results from simulations of neutral evolution starting from standing variation as a standard for comparison to all selection scenarios. These distributions should accurately reflect the range of environmental associations seen for neutral loci in empirical samples.

The 4 tests differed considerably in their false positive rate. The rate was highest for LM and very low for MLM, LFMM, and BayPass (Figure 3, dotted lines). It is not surprising that the LM test, which has no correction for population structure or demographic history, had the highest false positive rate; for all simulated scenarios more than 30% of neutral loci exceeded a LOD value = 5.0. The high false positive rates for the LM indicate that the LOD scores should not be directly interpreted as the probability of rejecting the null hypothesis that an allele does not contribute to local adaptation. However, LOD scores might still prove useful for ranking SNPs in a genome-wide sample, as is evident from the distributions of LOD scores for the selected locus in each simulation scenario (Figure 4). For both LM and MLM tests, the LOD scores for selected AP loci far exceed those from the neutral loci. For LM and MLM tests, the mean LOD scores for CN loci also clearly exceed those of neutral simulations, suggesting that using the upper tail of the distribution of LOD scores may be an effective strategy for identifying genetic variants that contribute to local adaptation.

Given that the MLM test performed well, with a high rate of true positives and a low rate of false positives, we focus the rest of the Results section on results from the MLM analyses.

Timing of Selection

Variation at AP loci is expected to be maintained in a population because different alleles are favored in different parts of the range. By contrast, variation at CN loci is transient—going rapidly to fixation in the part of the range where the allele is selectively favored and then drifting to fixation in the part of the range that the allele is neutral (Figure 2). This shapes the power of GEA tests to identify AP versus CN loci at different times after the onset of selection. At AP loci, the power to identify a selected locus is high regardless of the time of sampling (Figure 5). The one exception to this is for sampling

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

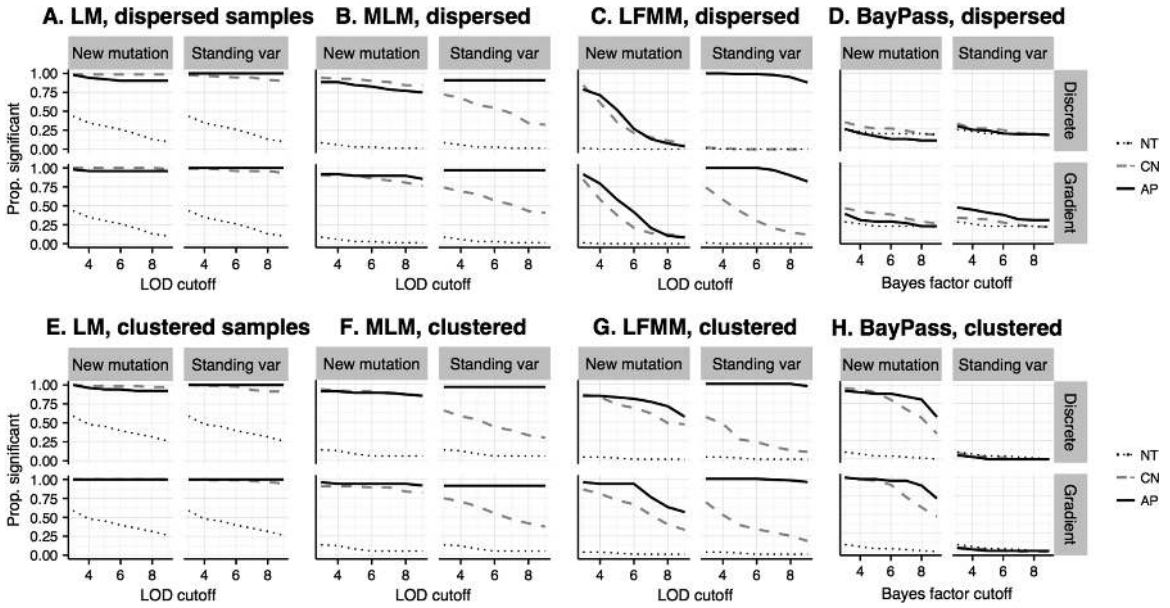

Figure 3. Detection of GEAs. Proportion of simulated datasets passing the MAF > 0.05 filter in which the selected locus had greater GEAs than expected by chance in 4 different GEA tests, under different a priori cutoffs for statistical significance. Solid lines are data from simulations of AP; dashed lines, simulations of conditional neutrality; dotted lines, simulations without selection acting. For A–D, tests were done on samples collected in a dispersed manner; for E–H, on samples collected in clusters (see Figure 1, Methods), all taken in generation 2000.

at 1000 generations after selection is imposed upon a new mutation, in which case the power of the MLM test to detect the target of selection is <80% (for LOD > 5, ~50% for LOD > 9) on a discrete selection landscape. This reduced power reflects the fact that selection was not great enough to always overcome the effects of drift, and in some simulations many generations and recurrent mutation occurred before the frequency of the locally adapted allele was high enough to be detected.

By contrast, the power to detect a CN allele responsible for local adaptation is greatest at the earliest sampling time and decays steadily with time. When samples are taken 1000 generations after selection is first imposed, 70% of selected loci (averaged across the selection scenarios) were identified as targets of selection by the MLM test at LOD > 5. By 10 000 generations, only 10% of selected CN alleles, averaged across all scenarios, were identified as targets of selection at LOD > 5 and only ~2% of the alleles were identified at LOD > 9. Not only is the frequency of CN alleles identified as targets of selection lower than the frequency of AP alleles identified as targets, but fewer CN alleles are polymorphic at MAF > 0.05, at 10 000 generations all of the AP loci meet the MAF > 0.05 filter, whereas <70% of CN alleles that began as new mutations and <40% of the CN alleles that began as standing variation are still segregating at MAF > 0.05.

Signal at Linked Loci

The four local adaptation tests performed well, for the most part, when the selected locus was analyzed. However, even with full genome sequence data uneven coverage will often mean that studies are testing for evidence of local adaptation using SNPs that are physically linked to the selected locus, but not themselves the target of selection. The power to detect targets of selection when sampling physically linked SNPs is clearly much weaker than when the selected locus itself is sampled. Because the relative performance of the 4 statistical tests on linked SNPs mirrored that of the performance when assaying the selected SNP, we focus our discussion of linked SNPs on results from the MLM test only.

When the selected locus was assayed, the MLM test with a LOD cutoff of >5 positively identified nearly 100% of all AP alleles and CN alleles that began as new mutations, and identified ~50% of CN alleles that began as standing variation alleles. By contrast, when physically linked neutral loci were assayed only ~40% of selected loci were identified under the best conditions—that is, when sampling linked loci that were in complete physical linkage to an AP allele that began as a new mutation and had been subject to selection for 10 000 generations. The power to detect loci linked to targets of selection was much lower when selection acted on standing variation or CN alleles, and when the selected locus was more than 0.01 cM away from the assayed locus (Figure 6).

Averaged across all of the sampling times, when AP loci were in complete physical linkage (i.e., at a linkage distance of 0 cM) with the selected locus, only 20% of loci had LOD > 5.0 in the MLM test when selection acted on a new mutation on a gradient landscape; and 8% when selection acted on standing variation. The power of the tests fell to 17% and 7% (new mutations and standing variation, respectively) when the assayed loci were 0.01 cM away from the target of selection, and to 10% and 4% when the assayed loci were 0.1 cM away from the target of selection. For comparison, approximately 2% of neutral mutations had LOD > 5—the loci just 0.01 cM from the selected locus were less than 10 times more likely to show positive signs of GEA. When assayed SNPs were 1 cM away from the selected locus, the true positive rates were approximately equal to the false positive rates from neutral loci (2.2% and 2.3% for AP and CN alleles).

Signals of Selection in Haplotype Diversity and LD

In addition to evaluating the performance of tests to identify targets of local adaptation using genome–environment association tests, we examined whether linkage information might provide a signature

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

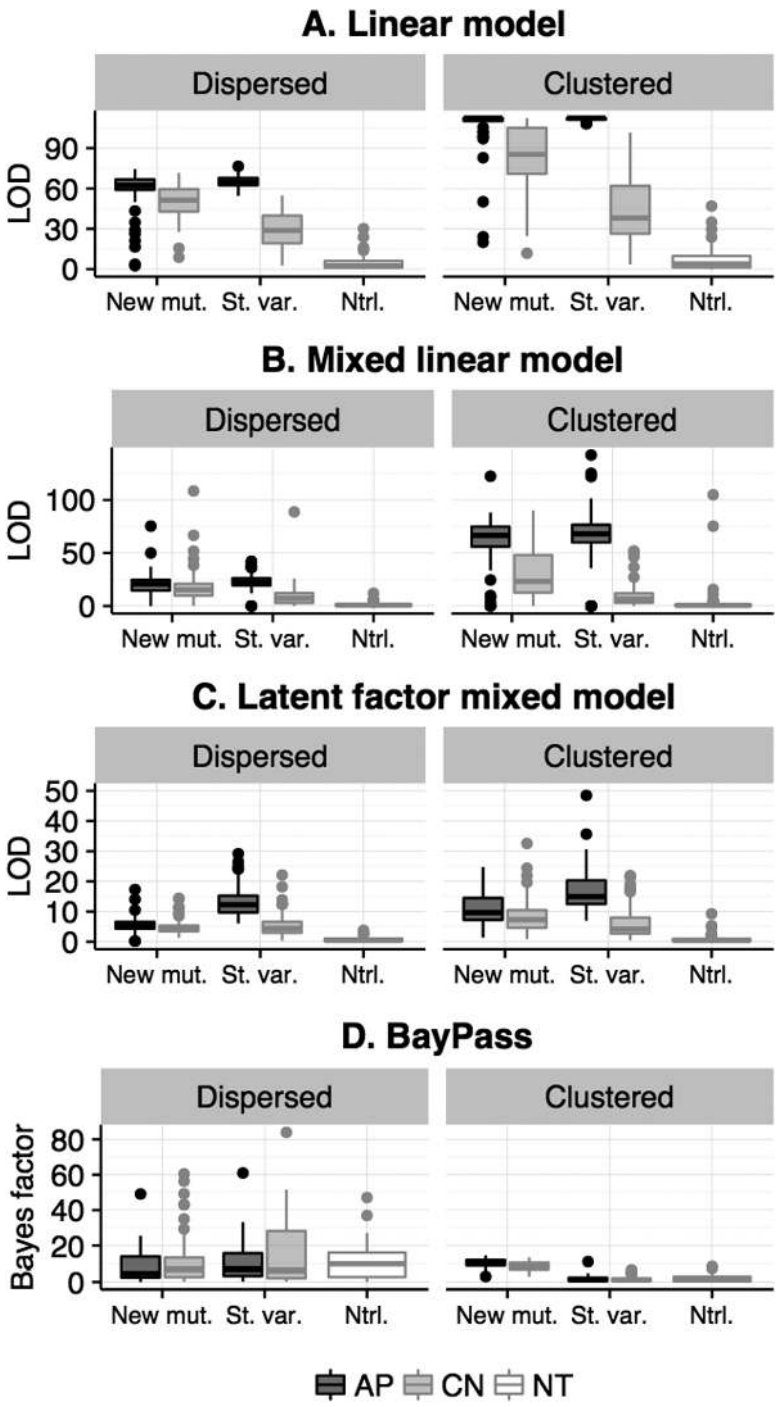

Figure 4. Distribution of test statistics for GEA tests. Distribution of LOD score or Bayes factor for the locus under selection in simulations of AP, CN, or neutral (NT) loci, starting from new mutations or standing variation, and with tests conducted on dispersed or clustered samples (Figure 1) taken in generation 2000.

that could be used to differentiate selection that had acted on new mutations versus standing variation and differentiate AP from CN alleles. To do this, we compared r^2 and H12 values calculated for loci that were in complete physical linkage (cM = 0) with selected loci that had been identified by the MLM as being a target of selection with LOD > 5. Unfortunately, the distributions of both H12 and r^2 values for regions linked to CN and AP loci are largely overlapping (Figure 7) and thus provide no clear signal for differentiating CN from AP loci. Selection that acted on new mutations resulted in r^2 and H12 values that are generally higher than selection that acted on standing variation (Figure 7). Moreover, the distributions of r^2 and H12 for loci linked to CN alleles are very similar to those obtained for neutral loci.

Haplotype diversity, as captured by H12, had generally lower power to positively identify targets of selection than did GEA tests, as demonstrated by comparing the distributions of MLM LOD and H12 (Figure 8). Except for clustered samples from simulations of CN loci, at least 90% of simulations on a gradient landscape had

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

MLM LOD greater than the 95th percentile from corresponding neutral simulations, across different sampling methods, gene actions, or whether selection acted on a new mutation or standing variation. By contrast, at loci in complete physical linkage with the selected locus (i.e., at a linkage distance of 0 cM), 48% to 79% of simulations of selection on a new mutation resulted in H12 greater than the 95th neutral percentile, and at most 6% of simulations of selection on standing variation resulted in such H12 outliers.

Discussion

Evaluating the performance of the statistical methods used to identify loci responsible for local adaptation helps in the design and interpretation of empirical searches for adaptive loci. Simulations are valuable for achieving these goals because one can control the scenarios in which selection is acting. The utility of simulation-based analyses is, however, dependent upon the degree to which the assumptions of the simulation capture biological reality. We compared the power of 4 tests used in genomic scans of local adaptation in natural populations to identify loci contributing to local adaptation via AP or CN, when selection acted on a new mutation or standing variation, and when sampling linked rather than causative loci.

Previous simulations studies of GEA test performance have assumed that the selected site is among those assayed. However, many empirical studies analyze relatively sparse genome-wide marker data and are unlikely to have captured more than a small minority of selected variants (Tiffin and Ross-Ibarra 2014; Lowry et al. 2016). In our analyses, all tests had high true positive rates when the selected site itself was assayed; however performance was much lower when physically linked sites were assayed, dropping from near 100% true-positive rates at selected sites to as low as 5% at a recombination distance of just 0.1 cM (Figure 6).

That the signal of adaptation is weaker at loci that are physically linked but not in complete LD with the variants on which selection is acting is unsurprising (e.g., Vilas et al. 2012; Hoban et al. 2016). Nevertheless, it is important to realize that when designing and interpreting local adaptation scans based on reduced representation data such as RAD-seq (reviewed by Davey et al. 2011), the ability to identify targets of selection depends upon more than tight physical linkage between the assayed and the causative SNPs. Rather, it depends on the nonrandom association of alleles, or LD. Even 2 alleles in close physical linkage may not have strong LD given that mutations do not enter the population in high LD (as measured by r^2) with adjacent variants already segregating in a population and LD can break down rapidly through recombination. For these reasons, even physically proximate SNPs segregating in genomic regions with high mean LD may be in low LD with one another (e.g., Remington et al. 2001; Branca et al. 2011).

Rapid decay in the signal of selection with increasing physical distance might be viewed as a weakness because it requires the assayed SNPs to be very close to the target of selection; but it can

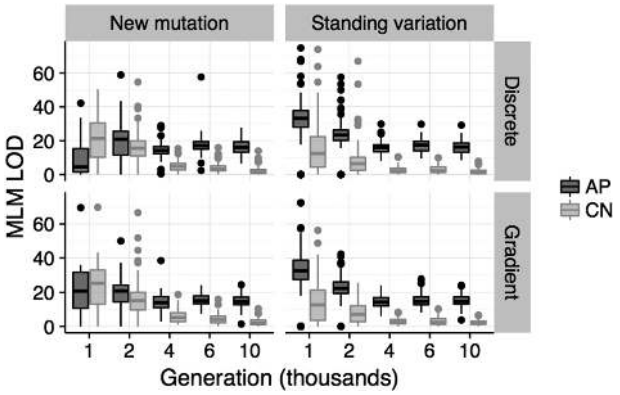

Figure 5. Signals of local adaptation over time. Boxplots of the LOD returned by the MLM test for AP (dark) or CN (lighter) loci based on dispersed samples (Figure 1) at each time point we examined, up to 10000 generations after the onset of selection.

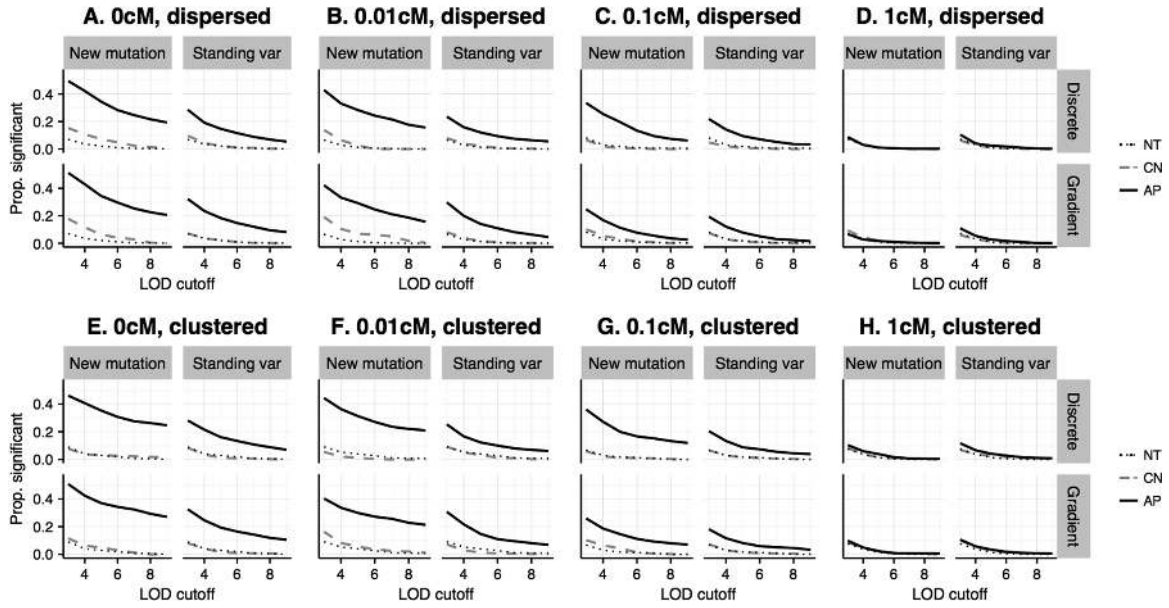

Figure 6. Signs of local adaptation at linked loci. Proportion of simulated datasets passing the MAF > 0.05 filter in which locus at different distances from the selected locus had greater genotype–environment associations than expected by chance in the MLM test. Solid lines are data from simulations of AP; dashed lines, simulations of conditional neutrality; dotted lines, simulations without selection acting. For A–D, tests were done on samples collected in a dispersed manner; for E–H, on samples collected in clusters (Figure 1), all taken in generation 10000.

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

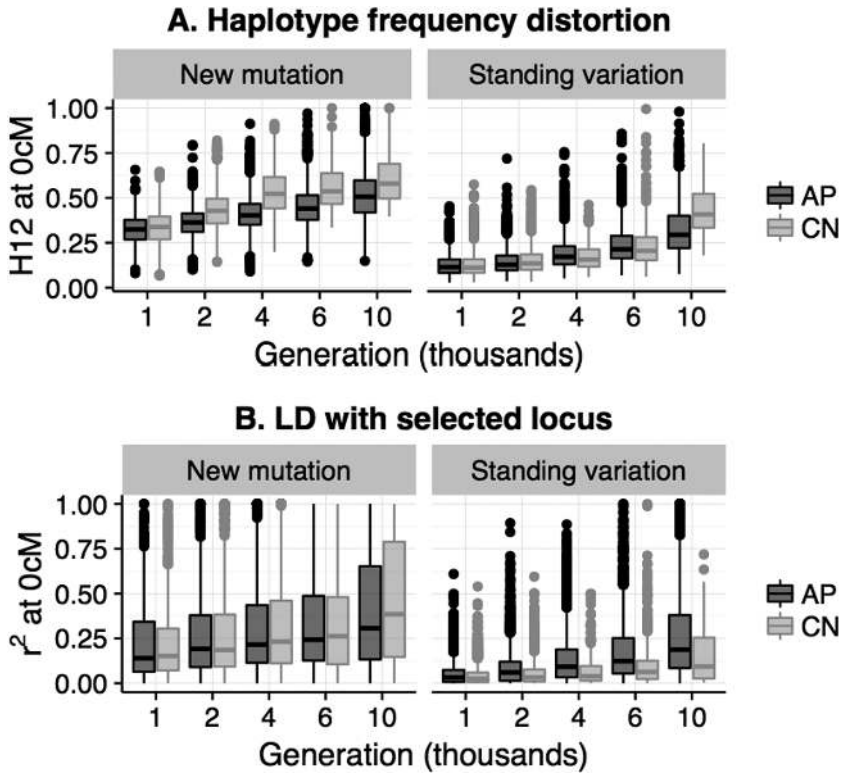

Figure 7. Haplotype diversity and LD with the selected locus. (**A**) Distributions of the H12 index of haplotype frequency distortion and (**B**) LD, as r^2, with AP the selected locus, calculated on data from neutral loci in complete physical linkage (i.e., at 0 cM linkage distance) to AP (dark) or CN (lighter) selected loci with MAF > 0.05 and MLM LOD > 5. Data are from dispersed samples collected at each of the time points indicated, on a discrete landscape (Figure 1).

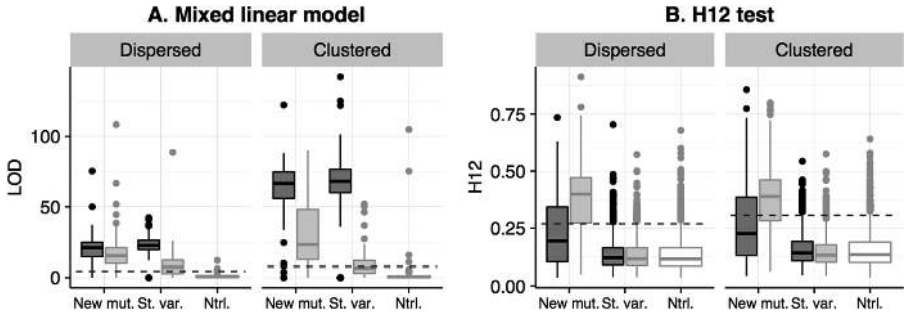

Figure 8. Relative power of the MLM test versus H12. Distribution of (**A**) MLM LOD score for the locus under selection and (**B**) H12 for loci at 0 cM from the selected locus, in simulations of AP (dark), CN (light), or neutral (white) loci on a gradient landscape, starting from new mutations or standing variation, and with tests conducted on dispersed or clustered samples (Figure 1) taken in generation 2000. In all panels, the dashed horizontal line marks the 95th percentile of the LOD or H12 distribution for neutral simulations.

also be viewed as a strength, because it means that if a signal of selection is detected one can have fairly high confidence that the actual target of selection is nearby. This proximity to the selected site is likely to be especially important if the goal is to identify the genes (or even causative variants) responsible for adaptation. Identification of these genes is necessary if researchers seek to characterize the functional basis of local adaptation or follow up with empirical work to provide more refined understanding of the process driving local adaptation (e.g., Anderson et al. 2013).

Although all of the GEA tests we examined had generally high power to detect a locally adapted locus when the selected locus was assayed, they differed widely in their vulnerability to false positives. In particular, the LM test had a very high rate of false positives. This high rate is not surprising given that the LM test includes no covariates that correct for population structure. Nevertheless, the difference in the magnitude of absolute LOD values between selected loci and neutral loci (Figure 4) suggests that even the simple LM test could reliably differentiate locally adapted loci from neutral loci if locally adapted loci are identified as outliers in the genome-wide GEA distribution rather than based on an a priori significance cutoff. The corrections imposed by the LFMM resulted in a low true positive rate when linked sites were assayed and a false positive rate that was higher than MLM. The MLM test corrects for population structure by including a kinship matrix as a covariate in the mixed linear model. This had 2 desirable effects. First, it resulted in a high ratio of true to false positives, and second the true positives it detected were in close physical linkage to the actual target of selection. If the goal of a GEA study is identify the genes responsible for local adaptation,

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

then close proximity of true positives to the actual target of selection is, of course, a desirable characteristic of a test for GEA.

Ideally, population genetic data could be used to differentiate whether local adaptation was due to genes with CN or AP and whether selection had acted on a new mutation or standing variation. Because of the expectation that drift and recombination would differentially affect the genomic variants in tight physical linkage with CN compared to AP loci, we examined H12 and r^2 among SNPs in the same genomic region as selected variants. H12 should be greater if a selected variant is on only 1 or 2 genetic backgrounds, and r^2 estimate linkage-disequilibrium between 2 variants. Unfortunately, neither H12 nor r^2 differ greatly between AP and CN loci and thus do not provide a strong signal to differentiate these types of loci. However, differences in H12 and r^2 together with the magnitude of the LOD scores that differentiate selected from neutral loci suggest that these 2 statistical signals might be useful in differentiating AP and CN alleles. In particular, alleles with very high LOD values but r^2 and H12 values that are similar to the genome wide averages might more likely be AP than CN alleles. H12 and r^2 might also by themselves provide direct insight into whether selection has acted on a new mutation of standing variation; r^2 and especially H12 in the regions of SNPs that the MLM test identifies as targets of selection show little overlap between new mutations and standing variation (Figures 7 and 8).

Although our simulated data and analyses did not provide clear signals that could be used to differentiate AP from CN loci, our analyses made it clear that CN loci will remain polymorphic in populations for far fewer generations than AP loci. The loss of variation at CN loci was particularly striking when selection acted on standing variation—by 10 000 generations after the onset of selection on a gradient-type landscape (Figure 1), only 37% of CN loci remained biallelic with a minor allele frequency of at least 0.05; on a discrete landscape only 31% did. The loss of diversity at CN loci is consistent with expectations that loci favored in part of the range and neutral in the other parts of the range will eventually go to fixation in all parts of the range due to gene flow from the regions where they are favored by selection. This transient nature of polymorphism at CN loci suggests that if local adaptation is often driven by CN loci (Gardner and Latta 2006; Verhoeven et al. 2008; Lowry et al. 2009; Fournier-Level et al. 2011; Anderson et al. 2013, 2014a, 2014b), genetic differences between populations connected by gene flow will be temporary, unless new selected alleles continue to enter populations or the selective environment changes frequently.

Of course, the results from our simulations come with the caveat that we analyzed populations at demographic equilibrium experiencing relatively strong selection on a simple landscape. Populations at nonequilibrium demographic conditions, experiencing more complex landscapes, experiencing weaker selection, or selection distributed across multiple loci (Kelly 2006; Pritchard et al. 2010; Savolainen et al. 2013; Anderson et al. 2014b), would all violate these assumptions and thereby both increase the risk of false positives and weaken the ability to identify loci responsible for local adaptation (Excoffier et al. 2009a; Fourcade et al. 2013; De Villemereuil et al. 2014; Lotterhos and Whitlock 2014, 2015). On the other hand, natural systems in which populations in contrasting environments are linked by lower rates of gene flow or experience stronger selection for local adaptation should see greater effects on neutral loci linked to locally adapted loci (Charlesworth et al. 1997; Storz and Kelly 2008). Finally, the relative performance of the individual GEA tests examined here are known to vary with sampling design and the spatial structure of geographically varying selection (De Mita et al. 2013; Frichot et al. 2015). It is important to understand that many natural populations and real-world sampling designs may behave differently than what we see from our simulations, and we suggest the use of customized simulations based on prior knowledge of individual study systems to inform sampling design, analysis choice, and interpretation of results from population genomic scans.

Conclusions

Landscape genomic scans have become a common method for trying to identify locally adapted loci and in order to understand how geographically variable selection shapes the genetic diversity within species (Savolainen 2011; Tiffin and Ross-Ibarra 2014). These scans provide a means to identify locally adapted genome regions without an a priori list of genes under selection, or even a strong prior hypothesis as to what environmental gradients drive local adaptation, or which traits are responsible for local adaptation. However, what population genomic scans can detect is limited by gene action, the time since selection began to act, and the density with which the genome is sampled. If local adaptation is often driven by CN, rather than AP, our results suggest that genome-scan methods might miss many loci that contribute to local adaptation.

Funding

This work was funded by the National Science Foundation (1237993 to Nevin D. Young and P.T.). J.B.Y. is currently supported as a postdoctoral fellow with the CoAdapTree Project (Genome Canada to Sally Aitken, Sam Yeaman, and Richard Hamelin).

Acknowledgments

Computational resources for simulation and analysis were provided by the University of Minnesota Supercomputing Institute. We thank Sam Yeaman for comments on preliminary results.

Data Availability

We have deposited QuantiNEMO simulation configuration files, analysis scripts, and GEA test results and population genetic statistics in the Dryad (doi: 10.5061/dryad.p12q4).

References

Ågren JA, Oakley CG, McKay JK, Lovell JT, Schemske DW. 2013. Genetic mapping of adaptation reveals fitness tradeoffs in *Arabidopsis thaliana*. *Proc Natl Acad Sci USA*. 110:21077–21082.

Anderson JT, Lee CR, Mitchell-Olds T. 2014a. Strong selection genome-wide enhances fitness trade-offs across environments and episodes of selection. *Evolution*. 68:16–31.

Anderson JT, Lee CR, Rushworth CA, Colautti RI, Mitchell-Olds T. 2013. Genetic trade-offs and conditional neutrality contribute to local adaptation. *Mol Ecol*. 22:699–708.

Anderson JT, Wagner MR, Rushworth CA, Prasad KVSK, Mitchell-Olds T. 2014b. The evolution of quantitative traits in complex environments. *Heredity*. 112:4–12.

Anderson JT, Willis JH, Mitchell-Olds T. 2011. Evolutionary genetics of plant adaptation. *Trends Genet*. 27:258–266.

Beaumont MA, Nichols RA. 1996. Evaluating loci for use in the genetic analysis of population structure. *Proc R Soc Lond B*. 263:1619–1626.

Benjamini Y, Hochberg Y. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Ser B*. 57:289–300.

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

Bonhomme M, Chevalet C, Servin B, Boitard S, Abdallah J, Blott S, Sancristobal M. 2010. Detecting selection in population trees: the Lewontin and Krakauer test extended. *Genetics*. 186:241–262.

Branca A, Paape TD, Zhou P, Briskine R, Farmer AD, Mudge J, Bharti AK, *et al*. 2011. Whole-genome nucleotide diversity, recombination, and linkage disequilibrium in the model legume *Medicago truncatula*. *Proc Natl Acad Sci USA*. 108:E864–E870.

Caballero A, Vilas A, Pe A, Pérez-Figueroa A, Caballero A. 2012. A simulation study on the performance of differentiation-based methods to detect selected loci using linked neutral markers. *J Evol Biol*. 25:1364–1376.

Charlesworth B, Nordborg M, Charlesworth D. 1997. The effects of local selection, balanced polymorphism and background selection on equilibrium patterns of genetic diversity in subdivided populations. *Genet Res*. 70:155–174.

Cheng C, White BJ, Kamdem C, Mockaitis K, Costantini C, Hahn MW, Besansky NJ. 2012. Ecological genomics of *Anopheles gambiae* along a latitudinal cline: a population-resequencing approach. *Genetics*. 190:1417–1432.

Davey JW, Hohenlohe PA, Etter PD, Boone JQ, Catchen JM, Blaxter ML. 2011. Genome-wide genetic marker discovery and genotyping using next-generation sequencing. *Nat Rev Genet*. 12:499–510.

De Mita S, Thuillet AC, Gay L, Ahmadi N, Manel S, Ronfort J, Vigouroux Y. 2013. Detecting selection along environmental gradients: analysis of eight methods and their effectiveness for outbreeding and selfing populations. *Mol Ecol*. 22:1383–1399.

de Villemereuil P, Frichot É, Bazin É, François O, Gaggiotti OE. 2014. Genome scan methods against more complex models: when and how much should we trust them? *Mol Ecol*. 23:2006–2019.

Devlin B, Roeder K. 1999. Genomic control for association studies. *Biometrics*. 55:997–1004.

Eckert AJ, van Heerwaarden J, Wegrzyn JL, Nelson CD, Ross-Ibarra J, González-Martínez SC, Neale DB. 2010. Patterns of population structure and environmental associations to aridity across the range of loblolly pine (*Pinus taeda* L., Pinaceae). *Genetics*. 185:969–982.

Excoffier L, Foll M, Petit RJ. 2009a. Genetic consequences of range expansions. *Annu Rev Ecol Evol Syst*. 40:481–501.

Excoffier L, Hofer T, Foll M. 2009b. Detecting loci under selection in a hierarchically structured population. *Heredity*. 103:285–298.

Fourcade Y, Chaput-Bardy A, Secondi J, Fleurant C, Lemaire C. 2013. Is local selection so widespread in river organisms? Fractal geometry of river networks leads to high bias in outlier detection. *Mol Ecol*. 22:2065–2073.

Fournier-Level A, Korte A, Cooper MD, Nordborg M, Schmitt J, Wilczek AM. 2011. A map of local adaptation in *Arabidopsis thaliana*. *Science*. 334:86–89.

Frichot E, Schoville SD, Bouchard G, François O. 2013. Testing for associations between loci and environmental gradients using latent factor mixed models. *Mol Biol Evol*. 30:1687–1699.

Frichot E, Schoville SD, de Villemereuil P, Gaggiotti OE, François O. 2015. Detecting adaptive evolution based on association with ecological gradients: orientation matters! *Heredity*. 115:22–28.

Ganapathy G, Uyenoyama MK. 2009. Site frequency spectra from genomic SNP surveys. *Theor Popul Biol*. 75:346–354.

Gardner KM, Latta RG. 2006. Identifying loci under selection across contrasting environments in *Avena barbata* using quantitative trait locus mapping. *Mol Ecol*. 15:1321–1333.

Garud NR, Messer PW, Buzbas EO, Petrov DA. 2014. *Soft selective sweeps are the primary mode of recent adaptation in Drosophila melanogaster*. ArXiv 1303.0906v.

Gautier M. 2015. Genome-wide scan for adaptive divergence and association with population-specific covariates. *Genetics*. 201:1555–1579.

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

Hermisson J, Pennings PS. 2005. Soft sweeps: molecular population genetics of adaptation from standing genetic variation. *Genetics*. 169:2335–2352.

Jensen JD, Thornton KR, Bustamante CD, Aquadro CF. 2007. On the utility of linkage disequilibrium as a statistic for identifying targets of positive selection in nonequilibrium populations. *Genetics*. 176:2371–2379.

Jones FC, Grabherr MG, Chan YF, Russell P, Mauceli E, Johnson J, Swofford R, Pirun M, Zody MC, White S, *et al*.; Broad Institute Genome Sequencing Platform & Whole Genome Assembly Team. 2012. The genomic basis of adaptive evolution in threespine sticklebacks. *Nature*. 484:55–61.

Jones MR, Forester BR, Teufel AI, Adams RV, Anstett DN, Goodrich BA, Landguth EL, Joost S, Manel S. 2013. Integrating landscape genomics and spatially explicit approaches to detect loci under selection in clinal populations. *Evolution*. 67:3455–3468.

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

Joost S, Kalbermatten M, Bonin A. 2008. Spatial analysis method (sam): a software tool combining molecular and environmental data to identify candidate loci for selection. *Mol Ecol Resour*. 8:957–960.

Kelly JK. 2006. Geographical variation in selection, from phenotypes to molecules. *Am Nat*. 167:481–495.

Lewontin RC, Krakauer J. 1973. Distribution of gene frequency as a test of the theory of the selective neutrality of polymorphisms. *Genetics*. 74:175–195.

Lipka AE, Tian F, Wang Q, Peiffer J, Li M, Bradbury PJ, Gore MA, Buckler ES, Zhang Z. 2012. GAPIT: genome association and prediction integrated tool. *Bioinformatics*. 28:2397–2399.

Lotterhos KE, Whitlock MC. 2014. Evaluation of demographic history and neutral parameterization on the performance of FST outlier tests. *Mol Ecol*. 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. *Mol Ecol*. 24:1031–1046.

Lowry DB, Hall MC, Salt DE, Willis JH. 2009. Genetic and physiological basis of adaptive salt tolerance divergence between coastal and inland *Mimulus guttatus*. *New Phytol*. 183:776–788.

Lowry DB, Hoban S, Kelley JL, Lotterhos KE, Reed LK, Antolin MF, Storfer A. 2016. Breaking RAD: an evaluation of the utility of restriction site-associated DNA sequencing for genome scans of adaptation. *Mol Ecol Resour*.

Meirmans PG. 2012. The trouble with isolation by distance. *Mol Ecol*. 21:2839–2846.

Messer PW, Petrov DA. 2013. Population genomics of rapid adaptation by soft selective sweeps. *Trends Ecol Evol*. 28:659–669.

Neuenschwander S, Hospital F, Guillaume F, Goudet J. 2008. quantiNemo: an individual-based program to simulate quantitative traits with explicit genetic architecture in a dynamic metapopulation. *Bioinformatics*. 24:1552–1553.

Poncet BN, Herrmann D, Gugerli F, Taberlet P, Holderegger R, Gielly L, Rioux D, Thuiller W, Aubert S, Manel S. 2010. Tracking genes of ecological relevance using a genome scan in two independent regional population samples of *Arabis alpina*. *Mol Ecol*. 19:2896–2907.

Pritchard JK, Pickrell JK, Coop G. 2010. The genetics of human adaptation: hard sweeps, soft sweeps, and polygenic adaptation. *Curr Biol*. 20:R208–R215.

Przeworski M, Coop G, Wall JD. 2005. The signature of positive selection on standing genetic variation. *Evolution*. 59:2312–2323.

Pyhäjärvi T, Hufford MB, Mezmouk S, Ross-Ibarra J. 2013. Complex patterns of local adaptation in teosinte. *Genome Biol Evol*. 5:1594–1609.

R Core Team. 2016. R: a language and environment for statistical computing. Vienna (Austria): R Foundation for Statistical Computing. Available from: https://www.R-project.org/

Remington DL, Thornsberry JM, Matsuoka Y, Wilson LM, Whitt SR, Doebley J, Kresovich S, *et al*. 2001. Structure of linkage disequilibrium and phenotypic associations in the maize genome. *Proc Natl Acad Sci USA*. 98:11479–11484.

Savolainen O. 2011. Evolution. The genomic basis of local climatic adaptation. *Science*. 334:49–50.

Savolainen O, Lascoux M, Merilä J. 2013. Ecological genomics of local adaptation. *Nat Rev Genet*. 14:807–820.

Schnee FB, Thompson JNJ. 1984. Conditional neutrality of polygene effects. *Evolution*. 38:42–46.

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022

Storz JF, Kelly JK. 2008. Effects of spatially varying selection on nucleotide diversity and linkage disequilibrium: insights from deer mouse globin genes. *Genetics*. 180:367–379.

Teshima KM, Coop G, Przeworski M. 2006. How reliable are empirical genomic scans for selective sweeps? *Genome Res*. 16:702–712.

Thompson JN. 2005. *The geographic mosaic of coevolution*. Chicago: University of Chicago Press.

Tiffin PL, Ross-Ibarra J. 2014. Advances and limits of using population genetics to understand local adaptation. *Trends Ecol Evol*. 29:673–680.

Verhoeven KJ, Poorter H, Nevo E, Biere A. 2008. Habitat-specific natural selection at a flowering-time QTL is a main driver of local adaptation in two wild barley populations. *Mol Ecol*. 17:3416–3424.

Vitalis R, Dawson K, Boursot P. 2001. Interpretation of variation across marker loci as evidence of selection. *Genetics*. 158:1811–1823.

Yoder JB, Stanton-Geddes J, Zhou P, Briskine R, Young ND, Tiffin P. 2014. Genomic signature of adaptation to climate in *Medicago truncatula*. *Genetics*. 196:1263–1275.

Yu J, Pressoir G, Briggs WH, Vroh Bi I, Yamasaki M, Doebley JF, McMullen MD, Gaut BS, Nielsen DM, Holland JB, *et al*. 2006. A unified mixed-model method for association mapping that accounts for multiple levels of relatedness. *Nat Genet*. 38:203–208.

Zhang Z, Ersoz E, Lai CQ, Todhunter RJ, Tiwari HK, Gore MA, Bradbury PJ, Yu J, Arnett DK, Ordovas JM, *et al*. 2010. Mixed linear model approach adapted for genome-wide association studies. *Nat Genet*. 42:355–360.

Downloaded from https://academic.oup.com/jhered/article/109/1/16/3803393 by guest on 20 August 2022