Genome-wide association mapping of transcriptome variation in *Mimulus guttatus* indicates differing patterns of selection on cis- versus trans-acting mutations

Keely E. Brown*[,1,2] and John K. Kelly*,

*Department of Ecology and Evolutionary Biology, University of Kansas, Lawrence, Kansas 66045, USA
[1]Current address: Department of Botany and Plant Sciences, University of California Riverside, Riverside, California 92521, USA
[2]Correspondence: 900 University Ave, Riverside, CA 92521, keelyb@ucr.edu

Running head: Selection on gene expression in monkeyflower

Keywords: eGWAS; coexpression modules; selection; floral traits; transcriptome; regulatory; monkeyflower

ABSTRACT

We measured the floral bud transcriptome of 151 fully sequenced lines of *Mimulus guttatus* from one natural population. Thousands of single nucleotide polymorphisms (SNPs) are implicated as transcription regulators, but there is a striking difference in the Allele Frequency Spectrum (AFS) of cis-acting and trans-acting mutations. Cis-SNPs have intermediate frequencies (consistent with balancing selection) while trans-SNPs exhibit a rare-alleles model (consistent with purifying selection). This pattern only becomes clear when transcript variation is normalized on a gene-to-gene basis. If a global normalization is applied, as is typically in RNAseq experiments, asymmetric transcript distributions combined with “rarity disequilibrium” produce a super-abundance of false positives for trans-acting SNPs. To explore the cause of purifying selection on trans-acting mutations, we identified gene expression modules as sets of co-expressed genes. The extent to which trans-acting mutations influence modules is a strong predictor of allele frequency. Mutations altering expression of genes with high “connectedness” (those that are highly predictive of the representative module expression value) have the lowest allele frequency. The expression modules can also predict whole-plant traits such as flower size. We find that a substantial portion of the genetic (co)variance among traits can be described as an emergent property of genetic effects on expression modules.

INTRODUCTION

Genetically controlled gene expression variation is prevalent within and between species, across different tissues, environments, and treatment contexts (Harding et al. 1989, Whitehead and Crawford 2006, McManus et al. 2010, Meiklejohn et al. 2014, Signor and Nuzhdin 2018). Changes in gene expression can facilitate divergence between species (Johnson and Porter 2000, Tulchinsky et al. 2014, Mack and Nachman 2017, McGirr and Martin 2020), and provide a mechanism for a population to rapidly adapt to a new environment (Morris et al. 2014, Ghalambor et al. 2015, Campbell-Staton et al. 2017, Margres et al. 2017, Mack et al. 2018, Hamann et al. 2020). Standing genetic variation and plasticity in gene expression can buffer a population against environmental fluctuations (Podrabsky and Somero 2004, Stern et al. 2007, Acar et al. 2008, López-Maury et al. 2008). While much has been learned about the regulation of particular genes, genome-wide patterns in the evolutionary dynamics of gene expression are just beginning to be explored (e.g. Josephs et al. 2020). It also remains unclear how gene expression, as a molecular phenotype, might mediate the genetic underpinnings of quantitative trait variation, and ultimately fitness.

Evolutionary dynamics of transcriptional effectors

Mutations can alter gene expression in many ways, and we expect selection to act differently on different types of variants (Lawrence et al. 2016, Bewick and Schmitz

2017, Duren et al. 2017). Gene expression can be affected by mutations acting either in cis or in trans. Cis-acting variants affect a closely linked gene directly, perhaps by altering sequences normally bound by transcription factors or other regulatory machinery. In contrast, trans-acting regulatory variants change the cellular environment in which transcription happens, say by altering diffusible products like the transcription factors themselves (Wittkopp et al 2004, Emerson and Li 2010).

Natural selection could differ systematically between cis- and trans-acting variants for several reasons. First, the mutational target for trans-acting effectors of a gene could be a substantial fraction of the genome (Boyle et al. 2017) while a more limited set of sites are available for cis-acting mutations (Gruber et al. 2012, Metzger et al. 2016). Second, trans-acting variants have the potential to affect multiple genes, and may thus have negative consequences on finely-tuned pathways (Stern and Orgogozo 2008; but see Hoekstra and Coyne, 2007). If trans-mutations have opposing pleiotropic effects on many genes (antagonism), they may still increase in frequency in a conditional manner (Hall et al. 2010, Anderson et al. 2011). Third, if a trans-acting variant with weakly deleterious effects on the expression of one or more target genes increases substantially in frequency (due to drift or selection), cis-variants specific to each affected gene might then act as a buffer, leading to positive directional selection on cis-compensatory mutations. This often occurs, for example, with pleiotropic mutations associated with antibiotic resistance (Maisnier-Patin and Andersson 2004, Brandis et al. 2012) and compensatory pairs of cis- and trans-effectors have been documented in several systems (Coolon et al. 2014, Wang et al. 2015, Fear et al. 2016, Mack et al. 2016, Verta et al. 2016, Metzger et al. 2017). These theories generally suggest that trans-acting variants should be under stronger negative selection than cis-acting variants based on the premise that gene expression should usually experience strong stabilizing selection (Denver et al. 2005, Rifkin et al. 2005, Whitehead and Crawford 2006, Hodgins-Davis et al. 2015). If this is correct, then any mutation with broad effects on expression, regardless of cis- or trans-effect, will more likely be deleterious, perhaps through cascading effects on connected pathways or networks (Fisher 1930).

Broad patterns of selection can be inferred from the allele frequency spectrum (AFS) of cis- and trans-acting variants. When compared to the neutral expectation, an excess of intermediate frequency variants suggests balancing selection while an excess of rare variants suggests purifying selection (Tajima, 1989). Demographic events, such as population expansions or contractions, can perturb the AFS away from the neutral expectation (Hartl and Clark, 1997). However, since demographic effects are genome-wide, we can make inferences about selection by comparing the AFS for a particular class of polymorphism (e.g. cis-effectors of gene expression) to that of the entire genome. Of course, this is just a first step; inferences about selection require corroboration from multiple lines of evidence (Beaumont and Balding 2004, Bigham et al. 2010). In this study, we find an AFS consistent with purifying selection for trans-acting expression variants and corroborate this pattern by showing that the skew towards extreme allele frequencies is greatest at loci with the broadest effects on expression. In contrast, cis-acting SNPs exhibit an AFS suggesting balancing selection. The processes most likely to generate balancing selection on cis-SNPs depend on the

specific ways that these mutations affect whole organism phenotypes, and also on the complicated and variable mapping from phenotype to fitness in nature.

Transcriptional mutations generating genetic (co)variation in traits

How important are transcriptional regulators in modifying fitness-related traits? Case studies of specific genes with known mutant phenotypes provide many examples where gene expression influences fitness relevant traits of plants (Streisfeld and Rausher 2009, Sobel and Streisfeld 2013, Ning et al. 2017, Kremling et al. 2018, Alonge et al. 2020). The quantitative importance of transcriptional mutations relative to those that effect enzymatic or structural protein function remains a point of contention (Hoekstra and Coyne 2007, Stern and Orgogozo 2008), but a steady increase of evidence from human eQTL/eGWAS research suggests a predominant role for gene expression variation in generating quantitative trait variation (Nicolae et al. 2010, Maurano et al. 2012, Torres et al. 2014, Farh et al. 2015, Boyle et al. 2017). As a first step to understanding the relationships between mutations affecting transcription, quantitative trait variation, and fitness, we here use observed gene expression variation to predict variation and co-variation among a set of quantitative traits. These traits correlate with field fitness components in yellow monkeyflower (*Mimulus guttatus)* and were previously analyzed as part of a GWAS that predicted trait and fitness measures directly from SNPs (Troth et al 2018).

In this study, we associate SNPs segregating within inbred lines derived from the Iron Mountain population with gene expression variation in flower buds. Allele frequencies in the inbred lines accurately represent those in the natural population (Troth et al. 2018). We first document strikingly different patterns of apparent selection from the AFS of cis- and trans-acting regulatory SNPs. We then show that modules of co-expressed genes predict the trait means of the inbred lines, despite that we measured gene expression and traits on different plants grown in different places. The stability of the relationship between transcriptome and trait is surprising, given the frequently cited “noisiness” of transcriptome data (Arias and Hayward 2006, Raj and Oudenaarden 2008). Finally, we demonstrate that correlations, including tradeoffs between fitness-related traits can be predicted from gene expression variation.

METHODS

Study system- We used randomly derived inbred lines of the yellow monkeyflower, *Mimulus guttatus (*syn *Erythranthe guttata,* Phrymaceae) from the Iron Mountain (IM) population in the Cascade Mountains of Oregon (44.402217N, 122.153317W; Willis 1999, Kelly 2003). This population is predominantly outcrossing with little internal population structure (Sweigart et al. 1999, Willis 1993). Due to its annual/winter annual lifespan and short growing season, the IM population experiences a fitness tradeoff caused by variation in flower size and life-history phenotypes (Mojica et al. 2012). In 2018, Troth et al. sequenced whole genomes of 187 IM inbred lines and phenotyped

them for 13 flower size and developmental timing traits known to influence fitness in the field.

RNAseq- We grew plants from 151 of the genome-sequenced inbred lines in the University of Kansas greenhouse under standard conditions (Monnahan and Kelly 2015) in three different cohorts. For each cohort, we grew more plants than needed for tissue collection and randomly selected plants for sampling soon after germination. We chose a recognizable and consistent stage at which to collect tissue, which we call the late floral bud stage. These are unopened flower buds (approximately 2-6 mm in length) on the first flowering node (so the corolla is presumably not fully expanded), but are advanced enough that buds on the second flowering node are visible. We chose bud tissue to enrich for transcripts related to flower size. When beginning the first cohort, it was unclear if this tissue type/amount would yield enough RNA for adequate sequencing. We thus pooled bud tissue from three plants of the same line in each tube prior to RNA extraction. Biological replicates were then multiple tubes of pooled tissue, all from the same line. Pooling was done randomly with regard to flowering time (ie. if 6 plants per line were sequenced, 3 in each of 2 tubes, one tube was not all three earliest flowering plants). This process was not repeated in cohorts 2 and 3, for which 1-3 biological replicates (plants) of each line were collected and sequenced separately. We collected tissue into liquid nitrogen at the same time of day (with regard to both actual time and hours after greenhouse lights turn on) within a two-hour window that was consistent between cohorts.

We ground the collected tissue finely with a plastic micropestle and extracted RNA using the Qiagen RNeasy Plant Mini Kit (Hilden, Germany). We generated sequencing libraries using the QuantSeq 3'mRNA-Seq Library Prep Kit for Illumina (Lexogen, Vienna, Austria) per protocol, modified to perform half reactions, and we sequenced the libraries using NextSeq HO-SR75bp (Illumina, San Diego CA, USA) at the University of Kansas Genome Sequencing Core. Each cohort was sequenced separately with a maximum of 96 samples per flow cell (a total of 4 flow cells and 281 individual samples).

To calculate read counts, we implemented the programs in Lexogen's BlueBee pipeline. First, we trimmed reads with bbduk (k=13, ktrim=r, useshortkmers=t, mink=5, qtrim=r, trimq=10, minlength=20) from BBTools 38.86 (Bushnell 2014) and aligned reads to the *M. guttatus* V2.0 reference genome (Phytozome, Hellsten et al., 2013) with STAR 2.5.0a (Dobin et al. 2013) using Lexogen's recommended parameters (outFilterMultimapNmax 20, alignSJoverhangMin 8, alignSJDBoverhangMin 1, outFilterMismatchNmax 999, outFilterMismatchNoverLmax 0.1, alignIntronMin 20, alignIntronMax 1000000, alignMatesGapMax 1000000). Finally, we counted transcript copies using htseq-count 0.11.2 and the genome annotation (Anders et al. 2014). The output is a table of read counts for each transcript. We then removed 5 samples that had fewer than 250k mapped reads (mean for remaining samples of 3,877,524 mapped reads) and normalized the counts for each sample (to account for variable library quality and sequencing depth) using the estimateSizeFactors function in DESeq2 1.28.1 (Love et al. 2014).

Predicting transcript levels from SNPs- Across all samples, 28,615 of 33,573 total annotated transcripts had at least one mapped read. We kept each gene isoform as a separate transcript. We first filtered out any transcripts which had mapped reads in fewer than 5% of samples, and which did not have 10 or more mapped reads in at least one sample. This left 20,463 transcripts for association mapping. The vast majority of genes (19,721 out of 20,463), had only one isoform with mapped reads. To account for the effect of cohort, we fit a linear model using lm in base R (R Core Team, 2013) to each transcript with cohort as a categorical predictor and then subtracted the estimated effect from each read count. We then transformed each transcript's expression in each sample by two methods: 1) log(expression + 1), and 2) Box-Cox transformation (Box and Cox, 1964) using the boxcox() function in the R package *EnvStats* 2.3.1 with a range for λ between -5 and 5 (Millard, 2014). Because some counts were negative after factoring out the effect of cohort, we shifted the distributions of all gene counts such that the minimum value was 0 for both types of transformation. Additionally, because Box-Cox transformation cannot accommodate zeros, we added a small value to each count that was equal to 10% of the minimum difference between any two samples (such that the difference between that value and zero was essentially undetectable in the original counts). Finally, we averaged every gene's expression across plants with each inbred line.

We obtained a filtered set of polymorphisms of the sequenced lines by starting with sites called by Troth et al. (2018). We kept only biallelic SNPs with a minor allele frequency above 2.5% that were called in at least half of the sequenced lines. We then pruned these sites for local LD using PLINK 1.90b3.38 (Purcell et al. 2007) with a window size of 50 SNPs, a step size of 10 SNPs, and an R^2 threshold of 0.9. This left 2,952,894 SNPs for downstream analysis. We performed the GWAS using GEMMA 0.98.1 (Zhou and Stephens 2014) by first constructing a centered relatedness matrix using all filtered, but unpruned sites. Finally, we used the univariate linear mixed model (-lmm) in GEMMA, which in the case of no covariates takes the form: expression = SNP genotype effect + random effect of relatedness + error. As part of the model, GEMMA outputs the 'chip heritability' for each gene; an estimate of the proportion of transcription variation that can be explained by all genetic causes. We used p-values taken from the likelihood ratio test to find associations between the levels of 20,463 transcripts and each of the 2,953,894 SNPs.

We classified the associations as cis-acting if the site was within 25kb of any part of the transcribed gene and trans-acting otherwise. This distance-based approach for calling cis-effectors can be undermined by long-distance LD. A physically distant SNP (which we would classify as trans) might be associated with expression simply because it is in LD with a cis-acting SNP. For this reason, we excluded data from the meiotic drive locus on chromosome 11 (a known region of extended LD, Fishman and Willis 2005, Fishman and Saunders 2008, Fishman and Kelly 2015) from genome-wide summaries. Overall, the sequenced lines from IM show a rapid decay of LD as inter-SNP distances exceed 10kb (Puzey et al. 2017) which makes our 25kb cutoff conservative. Cis- and trans-acting mutations can be distinguished more directly using allele-specific

expression data (Wittkopp et al. 2004, Springer and Stupar 2007, Tirosh et al. 2009, Shi et al. 2012, Osada et al. 2017, Signor and Nuzhdin 2018), but only in heterozygous individuals and here we are measuring expression in highly homozygous inbred lines.

To interpret our estimates for SNP effects on transcription, we permuted the vector of gene expression data (all genes) against line genotypes 100 times. For each replicate, we applied the GEMMA -lmm for all cis-tests (all cases where a SNP was within 25kb of a gene). Across permutation replicates, we obtained ca. 1.2 billion tests to relate SNP allele frequency to significance levels under the null hypothesis of no SNP effect on expression (as in Josephs et al. 2015).

We next performed simulations allowing SNP effects on gene expression. As previously, we first permuted expression values against genotypes. This simulates the "environmental variance" in expression. We then added $2 * \beta$ to all lines carrying the homozygous alternate (non-reference) genotype, where the genotypic effect (β) was determined separately for each gene. Given the large number of lines in our panel, the variance in expression attributable to SNP genotype, V_{snp}, is:

$$V_{snp} = 4\,q(1-q)\beta^2 \qquad \text{(Equation 1)}$$

where q is the frequency of the reference base in the line panel. Based on results from our significant cis-effect tests (described below), we set $V_{snp} = V_{residual}$ for the first set of simulations. In other words, SNP explains half the variance in expression. We calibrate the simulations with SNP effects in two different ways: 1) We consider the case where the proportion of variance due to SNP is held constant at 0.5. This implies:

$$\beta = \sqrt{\frac{V_{residual}}{4q(1-q)}} \qquad \text{(Equation 2)}$$

Since all variation is environmental after permutation, $V_{residual}$ is simply the variance of expression in the gene before adding effects to genotypes. 2) We set V_{snp} for each gene assuming that q = 0.5. This implies $\beta = \sqrt{V_{residual}}$. For this case, V_{snp} obtained after adding genotypic effects will vary with q (and be lower for SNPs with lower minor allele frequencies). These simulation schemes were considered by Tung et al (2015) in analyzing expression data, although these authors simulated the residual variance from a normal distribution while we use permutation of the observed expression levels. For both schemes, we performed tests on a random selection of 20 of the cis-SNPs for each gene with a distinct permutation of expression values versus line for each test.

Predicting phenotypes from gene coexpression modules- To identify sets of coexpressed genes, we used WGCNA 1.69 in R (Langfelder and Horvath 2008) using the sample normalized and cohort factored, but untransformed, counts as input with a power of 3, max block size of 21000, minimum module size of 30, dynamic tree cut method, correlation using dissimilarity, and merge cut height of 0.25. During co-expression analysis, one sample was removed as an outlier. WGCNA identified 37 modules of coexpressed transcripts (Table S1, Fig S1). Next, we extracted the line

means for 13 traits from Troth et al. (2018) for the 151 lines used in this study (germination date, days to flower, corolla width, corolla length, floral tube length, throat width, stigma length, anther length, height at flowering, first flowering node, width of widest leaf, and the first two principal components calculated from all floral dimensions. To look for associations between gene expression modules and measured traits, we Box-Cox transformed each module's eigen expression value (which is the first principal component of a PCA for expression of all member genes in a module), as well as every trait and fit a linear model (all in R, R Core Team, 2013). We used the eigen gene expression for each module as a predictor in a simple regression, as well as fitting the multiple regression for each trait using all 37 modules simultaneously. We also used the program stepAIC from the R package MASS 7.3-52 (Venables and Ripley 2013) to choose a lowest AIC (Akaike Information Criterion) regression model including some but not all modules as predictors.

For each best-fit multiple regression model, we used permutation to test for significance. We treated the set of traits as one block and the set of modules as another block and permuted which line had which of each block. This preserved the correlations between modules and between traits, but changed which sets of trait values and module values went together. To elaborate, imagine a line has a set of trait values X and module values K, and another line has a set of trait values Y and module values L. Then, a permuted data set might combine traits X with modules L and traits Y with modules K. This is referred to as Permutation 1 in the Results. We next used the coefficients from the best-fit model to predict trait variances and covariances. For each trait, we estimated the effect of each module included in the best fit model from a multiple linear regression and constructed an equation to predict trait value for each line:

$$Y_{jz}{}' = b + \sum_i m_{ij} x_{iz} \qquad \text{(Equation 3)}$$

where $Y_{jz}{}'$ is the predicted value of trait j for line z, m_{ij} is the estimated effect of module i on trait j, x_{iz} is the eigen expression of line z for module i, and the sum is taken over all modules in the model. The covariance of predicted values for traits j and k is:

$$Cov(Y_j', Y_k') = \frac{1}{n-1}\sum_z (Y_{jz}{}' - \mu_j) * (Y_{kz}{}' - \mu_k) \qquad \text{(Equation 4)}$$

where μ_j is the mean of trait j, μ_k is the mean of trait k, and the sum is taken over all n lines. Calculations were done using a custom python script (Supplemental File 1). We permuted traits against module values for testing. For each permuted set, we again found a best-fit model with a subset of gene expression modules and asked how much trait covariation we could predict (using the above method) to generate a distribution. We determined if the amount of trait covariation explained by the real gene expression data, as represented by modules, was significant using alpha levels calculated from the permuted distribution.

For each trait, we randomly sorted genes into modules with the same number of genes, calculated the eigen gene expression value (PC1) for each module for each line, Box-Cox transformed the module eigenvalues, and then included them in multiple

regression. In each case, we fit two models, one with all permuted modules and one with the stepAIC chosen set. We permuted the module composition 1000 times to generate distributions of R^2 for each trait. Reported p-values for Permutation 2 in the Results are calculated from those distributions.

RESULTS

Genetic effects on transcription levels of individual genes

The number of detected associations between genotype and expression, as well as the putative type of regulatory association (cis vs. trans), are contingent on how we transform transcript counts. Before testing for genetic effects on gene expression, we transformed expression read count data using two methods, log(count + 1) and Box-Cox with λ estimated for each gene separately. Using a rounded p-value cutoff of 1e-12 (Bonferroni = 8.27e-13), we identified 106,585 significant SNP/transcript associations using log(counts+1) transformed counts (10,087 cis associations and 96,498 trans, Fig 1A). Using Box-Cox transformed counts, over 90% of the significant trans-effects evaporate and we find only 8,088 cis and 7,685 trans associations (Fig 1B).

A careful inspection of the differences between the two methods indicates that the Box-Cox results are more reliable. With log(counts+1) transformation, individual genes often have skewed distributions with a small number of lines exhibiting atypically high or low expression (Fig1C, Fig S2). The few outlier lines with extreme expression will harbor the same minor (and in most cases very rare) allele at many loci. In fact, the majority of the 96,498 trans-regulatory associations (Fig 1A) involve SNPs with a minor allele frequency (MAF) between 2.5-5% (Fig 2A). When unlinked but rare alleles occur together in the same lines, LD is high owing to “rarity disequilibrium” (Houle and Márquez 2015, Lappalainen et al. 2013). If those same lines have extreme expression, all of the linked SNPs will show a strong association with expression.

The Box-Cox transformation provides a scale adjustment specific to each gene. In transcripts with the largest number of genetic associations in the log(counts+1) analysis, Box-Cox more completely “normalizes” expression reducing the effect of outliers (Fig 1C-D, Fig S2). Estimates from the Box-Cox are less affected by the pull of extreme values and we will subsequently limit attention to these tests. The mean estimated heritability of gene expression was 0.31 (see Fig S3 for the full distribution). For genome-wide analyses, we removed SNPs on chromosome 11 because the large block of apparent trans-effects on chromosome 11 are within the meiotic drive locus (Fishman and Willis 2005, Fishman and Saunders 2008, Fishman and Kelly 2015). The Drive allele is essentially a single DNA sequence over >5Mb of DNA segregating in the inbred lines at ~30%. As a consequence, it is impossible to distinguish trans-acting SNPs from those that are simply in linkage disequilibrium with cis-acting SNPs. This leaves 7,832 cis and 4,626 trans associations.

SNPs with significant cis-effects explain between 26-74% of the total variance in expression of a gene (Table S2). Significant trans-effectors have a larger average effect size than cis-SNPs (cis=0.54, trans=1.16, F-value = 9.07, p-value =2.6e-3), and effect size is negatively correlated with MAF (both log-transformed; estimated effect of cis-SNP effect size on MAF = -0.029, t = -8.347, p < 2e-16; estimated effect of trans-SNP effect size on MAF = -0.025, t = -4.351, p = 1.39e-05; Fig S4). Since we find trans-acting SNPs have a distribution skewed toward low frequency, it follows that such mutations would also have larger effect sizes, as has been reported in other systems (Josephs et al. 2020).

Across the genome, the distribution of minor allele frequency differs greatly between cis- and trans-acting mutations (Fig 2A). We find many rare alleles responsible for trans-regulatory effects on gene expression, and an increasing number of cis effects at higher minor allele frequency (MAF). We find a mean MAF for cis and trans associations of 0.342 and 0.215, respectively (F-value = 2197, p-value < 2.2e-16). The MAF distributions for cis- and trans-acting sites are both different from the MAF distribution of the entire genome (Fig S5, two-sample Kolmogorov-Smirnov tests: cis-to-all comparison: D = 0.51866, p < 2.2e-16, trans-to-all comparison: D = 0.1661, p < 2.2e-16), and are different from each other (D = 0.43651, p < 2.2e-16). The difference in MAF could be due to differential power to detect cis- vs. trans-acting loci with different effect sizes, since we found larger effect sizes for trans effectors. To test whether differences in effect size were driving the MAF pattern, we took only the top quartile of effect sizes (after normalizing by mean expression level) in each regulatory class. The pattern remains the same (Fig 2B) – an excess of common cis-acting alleles and an excess of rare trans-acting alleles. Finally, we established that the pattern is insensitive to our distance cutoff (25kb) for trans-effectors. If we limit trans to SNPs that affect expression on different chromosomes (Fig S6), the cis/trans difference remains.

Permutation tests indicated that our significance threshold for cis- tests is quite stringent. Across 100 whole genome permutations, only three of out of 1.2 billion SNP tests passed our $p < 10^{-12}$ threshold (Table S3A). Following Josephs et al. (2015), we considered significance levels across allele frequency categories of SNPs to establish a null distribution for the AFS of tests. For a given minor allele frequency, the fraction of tests that yield p-values less than a specified threshold (say 10^{-5}) are reported in Table S3A. Because permutation reiterates the null hypothesis, we expect the fraction to equal the threshold, e.g. about 1 test in a million would have $p < 10^{-6}$. In fact, we find that tests on intermediate allele frequency SNPs (minor allele > 20%) tend to be conservative (low p-values under-represented) while rare allele SNPs (minor allele < 10%) tend to be anti-conservative. These results imply a pull towards more minor frequencies in the null distribution for AFS. However, it is noteworthy that an extremely small number of tests approach our actual threshold.

The simulations allowing SNP effects on expression routinely yield significant results. The fraction of tests passing various thresholds for our two simulation schemes (constant V_{snp} and constant β) are reported in Table S3B,C. With V_{snp} held constant (effect size varies with allele frequency), about 90% of tests pass our $p < 10^{-12}$ threshold

regardless of allele frequency. Essentially all tests pass for lower thresholds. With fixed β (where the proportion of variance explained by a SNP varies with AF), there is lower power for rare alleles than intermediate frequency SNPs, which is expected given that rare alleles generate less variation. Most relevant to the results, we consider the relative proportion of significant tests that fall into each allele frequency class for each significance level, and how this compares to the observed AFS of significant tests. This is depicted for three thresholds in Figure 3. The AFS of significant tests with V_{snp} held constant is unaffected by threshold and matches the AFS of all tested SNPs. With V_{snp} constant, allele frequency has no effect on ascertainment. With fixed β, a smaller fraction of tests are significant for rare alleles (contrast orange to grey bars in Fig 3). However, this skew towards intermediacy is not sufficient to explain the data – the AFS of real tests is substantially more intermediate than predicted by ascertainment with fixed β (contrast orange to blue bars) at all significance thresholds.

We find no evidence for "trans-eQTL hotspots", single SNPs affecting many genes, similar to *Populus tremula* winter buds (Mähler et al. 2017). In fact, there are more genes with transcript levels that are affected by many trans-SNPs than SNPs with more than 2 trans-associations (Fig S7). The three genes with the most trans-acting SNPs are Migut.D00926 (160 SNPs) annotated as a jasmonate ZIM domain-containing protein (JAZ); Migut.M00568 (114 SNPs) annotated as a chlorophyll A/B binding protein; and N01403 (105 SNPs) annotated as an auxin-responsive F-box transport inhibitor response protein. For genes M00568 and N01403, the trans-associations are concentrated on the same chromosome as the gene (Fig 1B). There is an apparent association between the Chr11 Drive Locus and one gene on chromosome 8 (Fig 1A). This gene (Migut.H01175) has three putative homologs in the *Mimulus* genome, only one of which was expressed in our samples (Migut.K01148). We found that all of our samples had high expression for only one of the two genes (Fig S8), and low to no expression for the other, which could indicate mis-assembly. Indeed, when we mapped reads from two samples with expression of either gene to the newer reference genome build (*Mimulus guttatus* TOL v5.0, DOE-JGI, http://phytozome.jgi.doe.gov/) they all mapped to the same region on chromosome 11, which supports that Migut.H01175 is a mis-assembled isoform of Migut.K01148. Finally, the number of cis-associations for a gene is positively correlated with gene size (effect estimate for log(number of associations+1) ~ log(gene length) = 0.1797, t-value = 6.04, p = 1.87e-9 Fig S9).

Rare allele load refers to the proportion of segregating loci at which an individual carries the minor allele, if the population frequency of that allele is very low. It is similar to the concept of deleterious mutation load, but assumes nothing about the fitness effect of individual rare alleles. Instead, it is usually used to test whether or not there is a cumulative fitness effect of harboring many rare variants. This load predicts dysregulation of gene expression in maize (Kremling et al. 2018) and the severity of inbreeding depression in *M. guttatus* (Brown and Kelly 2020). We tested whether lines with an excess of rare alleles exhibit differing patterns of expression, but found no correlation between load and the number of genes showing extreme expression (plus/minus two standard deviations from the mean, Fig S10). We also find no clustering of lines by rare allele load in gene expression principal component (PC)

space (Fig S11). The many associations between gene expression and rare variants suggested by Fig 1A (and by the MAF of associations removed by Box-Cox transformation) is thus likely not a real cumulative effect of many rare alleles generating extreme gene expression genome-wide.

Construction of gene co-expression networks

We next sought to establish sets of genes that co-vary in expression across inbred lines. Using the cohort and individual normalized gene expression counts, WGCNA identified 37 modules of co-expressed transcripts. Each module includes between 37 and 5767 genes (mean 553, median 231) and each transcript (gene) belongs to only one module (Table S1, Fig S1). WGCNA groups genes with correlated expression and further collapses groups such that gene expression between modules should not be highly correlated ($R^2 > 0.8$). However, eigengene expression (principal component 1 for the PCA of all genes in a module) of a few pairs of modules remain moderately correlated (20 of 666 pairwise comparisons with 0.74 > R^2 > 0.5) (Fig S12).

To determine if the apparent purifying selection on trans-effecting sites (Fig 2A) is due to their impact on regulatory networks, we calculated the "connectedness" of each gene by correlating the gene's expression with the eigengene expression value of its module. This measures how predictive a gene's expression is of the expression of all genes in the module. Note that we are not calculating the number of edges a gene has in a regulatory or interaction network, which is sometimes called connectivity. Hence, the use of a nonstandard term. The distribution of correlation coefficients (R^2 with module PC1) is highly right-skewed (Fig 4A). For this reason, we grouped genes by "connectedness" quartile and then calculated the average MAF of sites affecting each gene either in cis or in trans. We find a consistent difference in MAF of cis- and trans-effectors, with trans having lower MAF, especially in the highest quartile for "connectedness," for which MAF is significantly lower than in all other categories (effect estimate for quartile 4 on trans-MAF = -0.143, p = 1.53e-14, effect estimate for quartile 4 on cis-MAF = -0.022, p = 0.00872) (Fig 4B). We did not find enrichment for any GO terms in the set of genes in connectedness quartile 4, using the closest *Arabidopsis thaliana* putative homologs.

Predicting phenotypes from expression

We next tested whether floral bud gene expression affects quantitative traits (Line means from Troth et al. 2018). We use modules of co-expressed genes as predictors of phenotype because this provides a tractable way to incorporate the whole transcriptome. We used multiple linear regression including all 37 modules as predictors of trait, and then chose the AIC-best model for each trait. This selected model included from 6 (widest leaf) to 20 (flower size PC1) of the 37 total expression modules (Table 1). Parameter estimates for the best-fit models, including effect sizes for each included module, are reported in Table S4. The best-fitting model for each trait explained from 23% to 47% of trait variation, with the strongest prediction being for overall flower size

(PC1 in Table 1). To establish statistical significance for prediction of trait variation, we permuted the data in two ways:

Permutation 1: Does gene expression predict trait variation? To test the hypothesis that a model using gene expression explains no more trait variation than by chance, we permuted modules by line. Correlations among traits and among modules were preserved (see methods), but randomly associated with each other across lines. This tests whether the transcriptome (as collapsed into coexpression modules) is a significant predictor of traits in a linear model. Using this method, trait variation predicted by the real gene expression modules is highly significant ($p < 0.01$) for 7 of 8 flower-size measurements (except flower size PC2), and marginally significant ($0.01 < p < 0.05$) for height and node (Table 1). These 9 traits are significantly correlated with each other, except for throat width with node (Fig 5, Table S5).

Permutation 2: Do gene co-expression modules better predict traits than random groups of genes? The significant prediction of traits by modules does not imply that modules are necessarily the best summary of gene expression for trait prediction. In order to test the hypothesis that the predicted trait variation is just a function of including the whole transcriptome (by creating groups of genes as predictors), we permuted module membership by shuffling genes into random groups of the same size as the real modules. These groups contain the same amount of information in terms of fraction of transcriptome included, but eliminate clustering of genes based on co-expression that defines the real modules. Essentially, we are asking if co-expression networks are a better way of decreasing parameter space than grouping genes randomly when the goal is to predicts trait values. By permuting gene module membership, we find that only 4 traits (corolla width and length, anther length, and flower size PC1) are significantly better predicted by co-expression modules than by random assortment of genes (Table 1). For all other traits, the amount of trait variation explained is attributable to the inclusion of the whole transcriptome, not variation in co-expressed groups of genes. However, while most traits are not significantly better predicted by real modules than scrambled sets of genes, real modules better predict traits than the *average* permuted data set for all but four traits (days to flower, node, widest leaf, and flower size PC2). Quantiles for the distribution of permuted R^2 for both permutations are presented in Table S6).

Overlapping sets of gene expression modules are included in the best-fit model for the four traits where expression modules are significant by both permutation tests (corolla width and length, anther length, and flower size PC1). All four are significantly predicted (in their own best-fit regression models) by 14 common modules. These traits are all positively correlated (Fig 5). The 14 modules are not correlated (Fig S12), but they affect all 4 traits in the same direction. As a consequence, trait correlations can emerge from the joint effects of uncorrelated modules.

Prior studies indicate that tradeoffs between fitness components (and associated traits) are central to the maintenance of variation in this population (Mojica and Kelly 2010, Scoville et al. 2011, Mojica et al. 2012, Monnahan and Kelly 2015, Monnahan and Kelly

2017, Brown and Kelly 2018). For this reason, we estimated the extent to which gene expression modules generate trait covariances. Among pairwise comparisons between the 9 traits that are significantly predicted by gene expression (Table 1 column 3, and see Permutation 1 above), 35 of 36 pairs are significantly correlated (R^2 between 0.07-0.91, $p < 0.05$ for all but node by throat width). We used estimates for the effect of each gene expression module on each trait from the best-fit multiple linear model (Table 1, column 1) to predict trait covariances using equations 1-2. If a module affects two traits, some fraction of the covariance between the traits can be attributed to the shared effect of that module. We find that 26-54% of the covariance between traits is attributable to this module-predicted covariance (35 pairwise comparisons). 33 of 36 covariances are significantly predicted by gene expression modules (Same permuted datasets as above, see equations 1 and 2 in Methods, $p < 0.05$, Fig 5 upper triangle). Modules are most strongly predictive of trait covariances among the four traits that are better predicted by modules than by the randomly grouped whole transcriptome (corolla width, corolla length, stigma length, anther length, and flower size PC1, see Table 1 column 4).

A large fraction of module variation is genetic. For each individual, we calculated the eigen expression (PC1) for each module and tested for an effect of inbred line using an ANOVA. Line explains 57-91% of variance in gene module expression (Table S7). Of the 37 modules, 29 are significantly affected by line ($p < 0.05$). There is no correlation between the estimated genetic control of a module and the number of traits for which a module is a significant predictor. However, all modules that significantly predict at least half of our measured traits (save one, “brown”) are significantly affected by genotype (p from 0.046 to 8.18e-15, F from 1.34 to 4.04). That is to say, modules that significantly predict many traits exhibit genetic variation among lines.

DISCUSSION

Natural selection on regulatory variants- Using a collection of sequenced inbred lines derived from a single natural population of yellow monkeyflower (*Mimulus guttatus*), we have dissected the genetic variation in the floral bud transcriptome. We found 12,458 SNPs with genome-wide significant associations with expression, 62% of which act in cis. Striking differences in the allele frequency spectrum (AFS) suggest differing selection regimes on cis- and trans-acting regulatory SNPs. Sites proximal to the affected gene are enriched for intermediate frequency variants. SNPs distant from target genes are enriched for rare variants. Hodgins-Davis et al. (2015) argue that gene expression should evolve according to a “house of cards” model, characterized by few mutations with large effects and moderate stabilizing selection (as opposed to a Gaussian model of evolution with many mutations of small effect and weak selection). Stabilizing selection on a quantitative trait with a fixed optimum predicts that minor alleles should be less common than under neutral evolution. Trans-acting mutations are more likely to be deleterious than cis-acting mutations if they have more pronounced effects (see introduction). The results of this study suggest that different selective pressures operate on cis and trans variation, consistent with previous work in a natural population of *Capsella grandiflora* (Josephs et al. 2020). The distribution of MAF for

trans-effecting SNPs in Capsella was similar to the Mimulus estimate (Fig 2). However, Capsella exhibits a nearly uniform distribution of MAF for cis-SNPs, while there is a definite inflation of intermediate frequency SNPs in Mimulus. The skew of cis-acting SNPs towards intermediate frequency, relative not only to trans-acting but also the genome as a whole, suggests balancing selection.

Ascertainment is a central concern for inference in QTL and association mapping studies (Beavis 1994). For loci with no effect on expression, our permutation study indicates that rare-allele SNPs are more likely to yield very low p-values than intermediate frequency SNPs. Thus, false positives are more likely to come from rare alleles than common, although permutation almost never produced p-values that pass the thresholds imposed on the real data (Table S3A). Considering SNPs with effects on expression, ascertainment depends on how we measure the 'importance' of a SNP. In simulated data where SNPs explain the same amount of expression variation as observed in the real data (constant V_{snp}), we find no effect of allele frequency on the probability that a SNP is detected (Fig 3, gray bars). In contrast, if we hold the effect of alleles constant (constant β), then V_{snp} is lower with extreme than intermediate allele frequencies. For fixed β, the simulations indicate that a SNP with equally frequent alleles (q = 0.5) is very likely to prove significant (90% of cases) while the detection probability falls to below 25% if $q \leq 0.15$ (Table S3C). However, this sieve does not explain the intermediacy of q for significant cis-SNPs in the data. First, the number of significant tests in the highest MAF categories significantly exceed the predicted number under the constant β simulations (Fig 3C). Second, the constant β simulations predict that tests on SNPs with q<<0.5 will 'fill in' the lower portion of the distribution when the significance threshold is reduced (Orange bars in Fig 3A,B). In other words, SNPs that do not pass the stringent 10^{-12} threshold simply because the minor allele is present in fewer lines should still routinely yield $p < 10^{-9}$ or $p < 10^{-6}$. The real data provide no indication of these "almost significant" SNPs in the range if $0.05 \leq q \leq 0.25$. The distribution is skewed intermediate across significance thresholds.

Figures 2 and 4 support the hypothesis that trans-effectors are routinely subject to purifying selection, at least for mutations with large enough effects to be detected in this study. Loci influencing expression in trans can affect multiple components of finely-tuned networks simultaneously. Here, we show that the minor allele frequency of SNPs affecting a gene's expression is correlated with how well that gene predicts the expression of many other genes (those in the same coexpression module), what we call "connectedness." Genes that are well-connected in this sense are likely to be the hub of a regulatory network, a role commonly filled by transcription factors (Babu et al. 2004), although we do not detect an enrichment for any particular type of gene in this set. We find that SNPs affecting well-connected genes tend to be lower in frequency and that the magnitude of decrease in MAF is stronger for trans-acting SNPs than cis-acting SNPs (Fig 4). This difference supports the idea that trans-effectors with broad pleiotropic effects on many genes are more likely to affect regulatory hubs and therefore be routinely subjected to purifying selection. Previous studies suggest that genes with high network connectivity are constrained by selection (Hahn and Kern 2005, Ramsay

et al. 2009, Josephs et al. 2017), which could explain why their expression would also be stabilized.

Connectedness of genes affected by trans-SNPs might explain the pattern of purifying selection, but it does not explain why cis-acting variants exhibit an MAF distribution suggestive of balancing selection. Cis-acting variants did have smaller effect sizes, which would explain a difference in severity of purifying selection, but not that allele frequencies at cis- SNPs are more intermediate than the genome-wide average. One potential explanation is that cis-acting variants may evolve on a gene-by-gene basis to counter the pleiotropic effects that trans-acting loci have on many genes. The hypothesis that cis-acting variants might evolve to mitigate trans-pleiotropy is supported by many studies finding opposing cis- and trans effects on the same gene (Coolon et al. 2014, Wang et al. 2015, Mack et al. 2016, Metzger et al. 2017). In this study, we find no such preponderance of compensatory cis/trans pairs. Using a conservative set of 24 genes with both cis-SNPs and inter-chromosome trans-SNPs, we find only one example of cis/trans compensation.

Genetic effect on traits mediated through gene expression- Understanding selection requires that we look at how genetic effects on gene expression translate to effects on whole-organism phenotypes and, ultimately, to fitness in the natural environment. Table 1 shows that floral and plant height measures can be significantly predicted by the flower bud transcriptome when abstracted into coexpression modules. The most precise prediction is for overall flower size (flower size PC1) and for the component measurements that jointly determine flower size (corolla width and length, stigma and anther lengths). The strength of prediction (nearly 50% of variation explained) is notable given that flower traits are likely established early in development (Krizek and Anderson, 2013). Accurate *prediction* of flower size from the RNAseq data does not imply that the bud mRNA from the exact time of sampling were *causal* to trait variation or covariation. Measured transcript levels might simply be strongly correlated through development time, which might suggest that trait variation is continually reinforced through development.

Prediction precision was likely reduced by the fact that modules were estimated from RNAseq performed on one set of plants, while the mean phenotypes were estimated from different plants of the same inbred lines. Plants from the two experiments almost certainly experienced subtle environmental differences (different greenhouses, growth at different times of the year, different years). The high R^2 for flower size despite these limitations suggests that stable relationships between genotypes and traits are mediated through transcriptome variation. Additionally, the separation of experiments avoids a subtle but potentially important bias. When phenotypes and gene expression levels are measured on the same plants, the two can become associated owing to confounding factors, even if there is no effect of expression on phenotype. Imagine that plants differ randomly in receipt of a resource such as soil nitrogen. If nitrogen affects both gene expression and phenotype, expression and phenotype will be correlated even if there is no inherent relationship. Establishing the mean phenotype of each line prior to measuring expression eliminates this bias (Rausher 1992).

Gene expression modules predict not only trait variation but also the covariances between traits (Fig 5). Trait correlations emerge when the same module influences multiple traits (Eqs 3-4). We find that module predictions can explain up to 54% of the observed covariance between traits (throat width and height). We further show that a substantial fraction of the variation in expression modules has a genetic basis (Table S7), which suggests variation in gene expression as a potential cause of genetic correlations between whole-plant traits. Understanding trait covariances is essential when natural selection involves trade-offs between traits. Such trade-offs can provide the mechanistic basis of balancing selection (Mérot et al. 2020), which is suggested in our data by the intermediacy of the AFS for cis-acting variants.

In many annual plants, suites of correlated life-history traits related to rate of development (progression to flowering) are subject to a tradeoff between flowering time and fecundity. In *Mimulus* specifically, variation in life-history traits is maintained by opposing selective pressures on survival to flower and seed set (Kelly 2008, Mojica and Kelly 2010, Mojica et al. 2012, Monnahan and Kelly 2015, Troth et al. 2018; Monnahan et al. 2021). As with many other species, small, fast-growing plants survive to flower but make fewer seeds. Large, slow-growing plants have the capacity to make more seeds and perhaps disperse more pollen, but risk not reaching maturity before the end of the growing season. This type of tradeoff can maintain polymorphism through balancing selection on loci affecting the underlying traits such as days to flower or flower size (Austen et al. 2017, Brown and Kelly 2018, Exposito-Alonso et al. 2018). Our bud transcriptome modules predict floral dimensions, but not development rate under greenhouse conditions (days to germination or days to flower; Table 1). However, we suggest that future studies measuring gene modules from a range of tissues at different time points, coupled with the statistical methods that we employ here (Eqs 3-4), might determine whether the survival/fecundity tradeoff in Mimulus contributes to the intermediate allele frequency pattern evident for cis-acting transcriptional mutations (Fig 2).

Scale of measurement for gene expression- Fig 1 contrasts two different ways to normalize read counts, Box-Cox and log(count+1). The latter is most similar to models typically applied in RNAseq studies, such as generalized linear models that use the log-link function (e.g. DESeq2; Love et al. 2014). When expression is normalized in the same way across all genes (such as with the log(count+1) method), rare alleles occurring in lines with extreme expression produce many false positives as a result of "rarity disequilibrium" (Houle and Márquez 2015, Lappalainen et al. 2013). When counts are instead power transformed using an exponent (λ) estimated for each gene separately (Box-Cox), samples with extreme expression are pulled closer to the mean of the resulting distribution (compare Figs 1C,D). This decreases the occurrence of spurious associations due to rare alleles. We retained the log(count+1) analysis in Fig 1 as a caution for future studies. This issue is likely to emerge in any situation where the absolute count of individuals carrying the rare allele is small (say less than 5).

Conclusion

The two major findings from this study are connected through our summarization of the transcriptome in terms of gene expression modules. The first result is that cis-acting SNPs tend to have intermediate allele frequencies (relative to the genome as a whole), while trans-SNPs exhibit a rare-alleles model consistent with purifying selection. Trans-acting mutations are most rare if they have broad effects, with the latter measured by how strongly a trans-affected gene predicts the overall expression of its module. The second result is that expression modules predict flower size with a surprising degree of precision. As a consequence, we can attribute substantial fractions of the variance in flower size measures to variation in expression modules (R^2 values in Table 1). Despite that expression levels of different modules are largely uncorrelated (across lines), they can generate covariances among traits because individual modules influence multiple traits. This 'transcriptome-explained' covariance can be a substantial portion of the total covariance across lines (up to 54%, Fig 5). Our results do not provide a clear explanation for why cis-acting SNPs exhibit allele frequencies consistent with balancing selection, but the prediction of trait covariances suggests how future studies that may address this question. Specifically, experiments that determine the nature and extent of transcriptional control of development rate could provide a more mechanistic understanding of balancing selection.

Data Availability

Gene expression data has been submitted to NCBI's SRA (project number PRJNA736440). The Python scripts used to generate trait covariances as well as those used for permutation and simulations are available as Supplemental File 1.

Acknowledgements

We thank S. Wright, S. Macdonald, P. Veltsos, L. Hileman, E. Everman, R. Unckless, C. Wessinger, C. Fiscus, and two anonymous reviewers for comment on the manuscript, and J. Willis, B. Sikes, the Koenig and Seymour labs at UCR for comments on early stages of analysis.

Funding

JKK acknowledges support from NSF grants DEB-1753630 and MCB-1940785. KEB was supported by NSF PRFB IOS-1907061 and a KU Genome Sequencing Core Award Voucher. This work was supported by the Center for Research Computing at the University of Kansas and the KU Genome Sequencing Core.

References

Acar, M., Mettetal, J. T., & van Oudenaarden, A. (2008). Stochastic switching as a survival strategy in fluctuating environments. *Nature Genetics*, *40*(4), 471–475.

Alonge, M., Wang, X., Benoit, M., Soyk, S., Pereira, L., Zhang, L., Suresh, H., Ramakrishnan, S., Maumus, F., Ciren, D., Levy, Y., Harel, T. H., Shalev-Schlosser, G., Amsellem, Z., Razifard, H., Caicedo, A. L., Tieman, D. M., Klee, H., Kirsche, M., … Lippman, Z. B. (2020). Major Impacts of Widespread Structural Variation on Gene Expression and Crop Improvement in Tomato. *Cell*, *182*(1), 145–161.e23.

Anderson, J. T., Willis, J. H., & Mitchell-Olds, T. (2011). Evolutionary genetics of plant adaptation. *Trends in Genetics: TIG*, *27*(7), 258–266.

Anders, S., Pyl, P. T., & Huber, W. (2014). HTSeq—a Python framework to work with high-throughput sequencing data. *Bioinformatics* , *31*(2), 166–169.

Arias, A. M., & Hayward, P. (2006). Filtering transcriptional noise during development: concepts and mechanisms. *Nature Reviews. Genetics*, *7*(1), 34–44.

Austen, E. J., Rowe, L., Stinchcombe, J. R., & Forrest, J. R. K. (2017). Explaining the apparent paradox of persistent selection for early flowering. *The New Phytologist*, *215*(3), 929–934.

Babu, M. M., Luscombe, N. M., Aravind, L., Gerstein, M., & Teichmann, S. A. (2004). Structure and evolution of transcriptional regulatory networks. *Current Opinion in Structural Biology*, *14*(3), 283–291.

Beaumont, M. A., & Balding, D. J. (2004). Identifying adaptive genetic divergence among populations from genome scans. *Molecular Ecology*, *13*(4), 969–980.

Beavis, W. D. (1994, December). The power and deceit of QTL experiments: lessons from comparative QTL studies. In *Proceedings of the forty-ninth annual corn and sorghum industry research conference* (Vol. 250, p. 266).

Bewick, A. J., & Schmitz, R. J. (2017). Gene body DNA methylation in plants. *Current Opinion in Plant Biology*, *36*, 103–110.

Bigham, A., Bauchet, M., Pinto, D., Mao, X., Akey, J. M., Mei, R., Scherer, S. W., Julian, C. G., Wilson, M. J., López Herráez, D., Brutsaert, T., Parra, E. J., Moore, L. G., & Shriver, M. D. (2010). Identifying signatures of natural selection in Tibetan and Andean populations using dense genome scan data. *PLoS Genetics*, *6*(9), e1001116.

Box, G. E. P., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society*, *26*(2), 211–243.

Boyle, E. A., Li, Y. I., & Pritchard, J. K. (2017). An Expanded View of Complex Traits: From Polygenic to Omnigenic. *Cell*, *169*(7), 1177–1186.

Brandis, G., Wrande, M., Liljas, L., & Hughes, D. (2012). Fitness-compensatory mutations in rifampicin-resistant RNA polymerase. *Molecular Microbiology*, *85*(1), 142–151.

Brown, K. E., & Kelly, J. K. (2018). Antagonistic pleiotropy can maintain fitness variation in annual plants. *Journal of Evolutionary Biology*, *31*(1), 46–56.

Brown, K. E., & Kelly, J. K. (2020). Severe inbreeding depression is predicted by the “rare allele load” in Mimulus guttatus. *Evolution; International Journal of Organic Evolution*, *74*(3), 587–596.

Bushnell, B. (2014). BBTools software package. *URL Http://sourceforge. Net/projects/bbmap*.

Campbell-Staton, S. C., Cheviron, Z. A., Rochette, N., Catchen, J., Losos, J. B., & Edwards, S. V. (2017). Winter storms drive rapid phenotypic, regulatory, and genomic shifts in the green anole lizard. *Science*, *357*(6350), 495–498.

Coolon, J. D., McManus, C. J., Stevenson, K. R., Graveley, B. R., & Wittkopp, P. J. (2014). Tempo and mode of regulatory evolution in Drosophila. *Genome Research*, *24*(5), 797–808.

Denver, D. R., Morris, K., Streelman, J. T., Kim, S. K., Lynch, M., & Thomas, W. K. (2005). The transcriptional consequences of mutation and natural selection in Caenorhabditis elegans. *Nature Genetics*, *37*(5), 544–548.

Dobin, A., Davis, C. A., Schlesinger, F., Drenkow, J., Zaleski, C., Jha, S., Batut, P., Chaisson, M., & Gingeras, T. R. (2013). STAR: ultrafast universal RNA-seq aligner. *Bioinformatics* , *29*(1), 15–21.

Duren, Z., Chen, X., Jiang, R., Wang, Y., & Wong, W. H. (2017). Modeling gene regulation from paired expression and chromatin accessibility data. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(25), E4914–E4923.

Emerson, J. J., & Li, W.-H. (2010). The genetic basis of evolutionary change in gene expression levels. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *365*(1552), 2581–2590.

Farh, K. K.-H., Marson, A., Zhu, J., Kleinewietfeld, M., Housley, W. J., Beik, S., Shoresh, N., Whitton, H., Ryan, R. J. H., Shishkin, A. A., Hatan, M., Carrasco-Alfonso, M. J., Mayer, D., Luckey, C. J., Patsopoulos, N. A., De Jager, P. L., Kuchroo, V. K., Epstein, C. B., Daly, M. J., … Bernstein, B. E. (2015). Genetic and epigenetic fine mapping of causal autoimmune disease variants. *Nature*, *518*(7539), 337–343.

Fear, J. M., León-Novelo, L. G., Morse, A. M., Gerken, A. R., Van Lehmann, K., Tower, J., Nuzhdin, S. V., & McIntyre, L. M. (2016). Buffering of Genetic Regulatory Networks in Drosophila melanogaster. *Genetics*, *203*(3), 1177–1190.

Fisher, R. A. (1930). *The genetical theory of natural selection*. Clarendon Press, Oxford, U.K.

Fishman, L., & Kelly, J. K. (2015). Centromere-associated meiotic drive and female fitness variation in Mimulus. *Evolution; International Journal of Organic Evolution*, *69*(5), 1208–1218.

Fishman, L., & Saunders, A. (2008). Centromere-associated female meiotic drive entails male fitness costs in monkeyflowers. *Science*, *322*(5907), 1559–1562.

Fishman, L., & Willis, J. H. (2005). A novel meiotic drive locus almost completely distorts segregation in mimulus (monkeyflower) hybrids. *Genetics*, *169*(1), 347–353.

Ghalambor, C. K., Hoke, K. L., Ruell, E. W., Fischer, E. K., Reznick, D. N., & Hughes, K. A. (2015). Non-adaptive plasticity potentiates rapid adaptive evolution of gene expression in nature. *Nature*, *525*(7569), 372–375.

Gruber, J. D., Vogel, K., Kalay, G., & Wittkopp, P. J. (2012). Contrasting properties of gene-specific regulatory, coding, and copy number mutations in Saccharomyces cerevisiae: frequency, effects, and dominance. *PLoS Genetics*, *8*(2), e1002497.

Hahn, M. W., & Kern, A. D. (2005). Comparative genomics of centrality and essentiality in three eukaryotic protein-interaction networks. *Molecular biology and evolution*, *22*(4), 803-806.

Hall, M. C., Lowry, D. B., & Willis, J. H. (2010). Is local adaptation in Mimulus guttatus caused by trade-offs at individual loci?: TRADE-OFFS CAN CONTRIBUTE TO ADAPTATION IN MIMULUS. *Molecular Ecology*, *19*(13), 2739–2753.

Hamann, E., Pauli, C. S., Joly-Lopez, Z., Groen, S. C., Rest, J. S., Kane, N. C., Purugganan, M. D., & Franks, S. J. (2020). Rapid evolutionary changes in gene expression in response to climate fluctuations. *Molecular Ecology*. https://doi.org/10.1111/mec.15583

Harding, K., Hoey, T., Warrior, R., & Levine, M. (1989). Autoregulatory and gap gene response elements of the even-skipped promoter of Drosophila. *The EMBO Journal*, *8*(4), 1205–1212.

Hartl, D. L. & Clark, A. G. (1997). *Principles of population genetics* (Vol. 116). Sinauer associates Sunderland, MA.

Hodgins-Davis, A., Rice, D. P., & Townsend, J. P. (2015). Gene Expression Evolves under a House-of-Cards Model of Stabilizing Selection. *Molecular Biology and Evolution*, *32*(8), 2130–2140.

Hoekstra, H. E., & Coyne, J. A. (2007). The locus of evolution: evo devo and the genetics of adaptation. *Evolution; International Journal of Organic Evolution*, *61*(5), 995–1016.

Houle, D., & Márquez, E. J. (2015). Linkage Disequilibrium and Inversion-Typing of the Drosophila melanogaster Genome Reference Panel. *G3* , *5*(8), 1695–1701.

Johnson, N. A., & Porter, A. H. (2000). Rapid speciation via parallel, directional selection on regulatory genetic pathways. *Journal of Theoretical Biology*, *205*(4), 527–542.

Josephs, E. B., Lee, Y. W., Stinchcombe, J. R., & Wright, S. I. (2015). Association mapping reveals the role of purifying selection in the maintenance of genomic variation in gene expression. *Proceedings of the National Academy of Sciences*, *112*(50), 15390-15395.

Josephs, E. B., Lee, Y. W., Wood, C. W., Schoen, D. J., Wright, S. I., & Stinchcombe, J. R. (2020). The Evolutionary Forces Shaping Cis- and Trans-Regulation of Gene Expression within a Population of Outcrossing Plants. *Molecular Biology and Evolution*, *37*(8), 2386–2393.

Josephs, E. B., Wright, S. I., Stinchcombe, J. R., & Schoen, D. J. (2017). The relationship between selection, network connectivity, and regulatory variation within a population of Capsella grandiflora. *Genome biology and evolution*, *9*(4), 1099-1109.

Kelly, J. K. (2003). Deleterious mutations and the genetic variance of male fitness components in Mimulus guttatus. *Genetics*, *164*(3), 1071–1085.

Kelly, J. K. (2008). Testing the rare-alleles model of quantitative variation by artificial selection. *Genetica*, *132*(2), 187–198.

Kremling, K. A. G., Chen, S.-Y., Su, M.-H., Lepak, N. K., Romay, M. C., Swarts, K. L., Lu, F., Lorant, A., Bradbury, P. J., & Buckler, E. S. (2018). Dysregulation of expression correlates with rare-allele burden and fitness loss in maize. *Nature*, *555*(7697), 520–523.

Krizek, B. A., & Anderson, J. T. (2013). Control of flower size. *Journal of Experimental Botany*, *64*(6), 1427–1437.

Langfelder, P., & Horvath, S. (2008). WGCNA: an R package for weighted correlation network analysis. *BMC Bioinformatics*, *9*, 559.

Lappalainen, T., Sammeth, M., Friedländer, M. R., Ac't Hoen, P., Monlong, J., Rivas, M. A., ... & Dermitzakis, E. T. (2013). Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*, *501*(7468), 506-511.

Lawrence, M., Daujat, S., & Schneider, R. (2016). Lateral Thinking: How Histone Modifications Regulate Gene Expression. *Trends in Genetics: TIG*, *32*(1), 42–56.

López-Maury, L., Marguerat, S., & Bähler, J. (2008). Tuning gene expression to changing environments: from rapid responses to evolutionary adaptation. *Nature Reviews. Genetics*, *9*(8), 583–593.

Love, M. I., Huber, W., & Anders, S. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, *15*(12), 550.

Mack, K. L., Ballinger, M. A., Phifer-Rixey, M., & Nachman, M. W. (2018). Gene regulation underlies environmental adaptation in house mice. *Genome Research*, *28*(11), 1636–1645.

Mack, K. L., Campbell, P., & Nachman, M. W. (2016). Gene regulation and speciation in house mice. *Genome Research*, *26*(4), 451–461.

Mack, K. L., & Nachman, M. W. (2017). Gene Regulation and Speciation. *Trends in Genetics: TIG*, *33*(1), 68–80.

Mähler, N., Wang, J., Terebieniec, B. K., Ingvarsson, P. K., Street, N. R., & Hvidsten, T. R. (2017). Gene co-expression network connectivity is an important determinant of selective constraint. *PLoS genetics*, *13*(4), e1006402.

Maisnier-Patin, S., & Andersson, D. I. (2004). Adaptation to the deleterious effects of antimicrobial drug resistance mutations by compensatory evolution. *Research in Microbiology*, *155*(5), 360–369.

Margres, M. J., Wray, K. P., Hassinger, A. T. B., Ward, M. J., McGivern, J. J., Moriarty Lemmon, E., Lemmon, A. R., & Rokyta, D. R. (2017). Quantity, Not Quality: Rapid Adaptation in a Polygenic Trait Proceeded Exclusively through Expression Differentiation. *Molecular Biology and Evolution*, *34*(12), 3099–3110.

Maurano, M. T., Humbert, R., Rynes, E., Thurman, R. E., Haugen, E., Wang, H., Reynolds, A. P., Sandstrom, R., Qu, H., Brody, J., Shafer, A., Neri, F., Lee, K., Kutyavin, T., Stehling-Sun, S., Johnson, A. K., Canfield, T. K., Giste, E., Diegel, M., … Stamatoyannopoulos, J. A. (2012).

Systematic localization of common disease-associated variation in regulatory DNA. *Science*, *337*(6099), 1190–1195.

McGirr, J. A., & Martin, C. H. (2020). Ecological divergence in sympatry causes gene misexpression in hybrids. *Molecular Ecology*, *29*(14), 2707–2721.

McManus, C. J., Coolon, J. D., Duff, M. O., Eipper-Mains, J., Graveley, B. R., & Wittkopp, P. J. (2010). Regulatory divergence in Drosophila revealed by mRNA-seq. *Genome Research*, *20*(6), 816–825.

Meiklejohn, C. D., Coolon, J. D., Hartl, D. L., & Wittkopp, P. J. (2014). The roles of cis- and trans-regulation in the evolution of regulatory incompatibilities and sexually dimorphic gene expression. *Genome Research*, *24*(1), 84–95.

Mérot, C., Llaurens, V., Normandeau, E., Bernatchez, L., & Wellenreuther, M. (2020). Balancing selection via life-history trade-offs maintains an inversion polymorphism in a seaweed fly. *Nature Communications*, *11*(1), 670.

Metzger, B. P. H., Duveau, F., Yuan, D. C., Tryban, S., Yang, B., & Wittkopp, P. J. (2016). Contrasting Frequencies and Effects of cis- and trans-Regulatory Mutations Affecting Gene Expression. *Molecular Biology and Evolution*, *33*(5), 1131–1146.

Metzger, B. P. H., Wittkopp, P. J., & Coolon, J. D. (2017). Evolutionary Dynamics of Regulatory Changes Underlying Gene Expression Divergence among Saccharomyces Species. *Genome Biology and Evolution*, *9*(4), 843–854.

Millard, S. P. (2014). E nv S tats, an RP ackage for E nvironmental S tatistics. *Wiley StatsRef: Statistics Reference Online*. https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118445112.stat07181

Mojica, J. P., & Kelly, J. K. (2010). Viability selection prior to trait expression is an essential component of natural selection. *Proceedings. Biological Sciences / The Royal Society*, *277*(1696), 2945–2950.

Mojica, J. P., Lee, Y. W., Willis, J. H., & Kelly, J. K. (2012). Spatially and temporally varying selection on intrapopulation quantitative trait loci for a life history trade-off in Mimulus guttatus. *Molecular Ecology*, *21*(15), 3718–3728.

Monnahan, P. J., Colicchio, J., Fishman, L., Macdonald, S. J., & Kelly, J. K. (2021). Predicting evolutionary change at the DNA level in a natural Mimulus population. *PLoS Genetics*, *17*(1), e1008945.

Monnahan, P. J., & Kelly, J. K. (2015). Naturally segregating loci exhibit epistasis for fitness. *Biology Letters*, *11*(8). https://doi.org/10.1098/rsbl.2015.0498

Monnahan, P. J., & Kelly, J. K. (2017). The Genomic Architecture of Flowering Time Varies Across Space and Time in Mimulus guttatus. *Genetics*, *206*(3), 1621–1635.

Morris, M. R. J., Richard, R., Leder, E. H., Barrett, R. D. H., Aubin-Horth, N., & Rogers, S. M. (2014). Gene expression plasticity evolves in response to colonization of freshwater lakes in threespine stickleback. *Molecular Ecology*, *23*(13), 3226–3240.

Nicolae, D. L., Gamazon, E., Zhang, W., Duan, S., Dolan, M. E., & Cox, N. J. (2010). Trait-associated SNPs are more likely to be eQTLs: annotation to enhance discovery from GWAS. *PLoS Genetics*, *6*(4), e1000888.

Ning, Y., Liu, W., & Wang, G.-L. (2017). Balancing Immunity and Yield in Crop Plants. *Trends in Plant Science*, *22*(12), 1069–1079.

Osada, N., Miyagi, R., & Takahashi, A. (2017). Cis- and Trans-regulatory Effects on Gene Expression in a Natural Population of Drosophila melanogaster. *Genetics*, *206*(4), 2139–2148.

Podrabsky, J. E., & Somero, G. N. (2004). Changes in gene expression associated with acclimation to constant temperatures and fluctuating daily temperatures in an annual killifish Austrofundulus limnaeus. *The Journal of Experimental Biology*, *207*(Pt 13), 2237–2254.

Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M. A. R., Bender, D., Maller, J., Sklar, P., de Bakker, P. I. W., Daly, M. J., & Sham, P. C. (2007). PLINK: a tool set for whole-genome

association and population-based linkage analyses. *American Journal of Human Genetics*, *81*(3), 559–575.

Puzey, J. R., Willis, J. H., & Kelly, J. K. (2017). Population structure and local selection yield high genomic variation in Mimulus guttatus. *Molecular Ecology*, *26*(2), 519–535.

Raj, A., & van Oudenaarden, A. (2008). Nature, nurture, or chance: stochastic gene expression and its consequences. *Cell*, *135*(2), 216–226.

Ramsay, H., Rieseberg, L. H., & Ritland, K. (2009). The correlation of evolutionary rate with pathway position in plant terpenoid biosynthesis. *Molecular biology and evolution*, *26*(5), 1045-1053.

Rausher, M. D. (1992). THE MEASUREMENT OF SELECTION ON QUANTITATIVE TRAITS: BIASES DUE TO ENVIRONMENTAL COVARIANCES BETWEEN TRAITS AND FITNESS. *Evolution; International Journal of Organic Evolution*, *46*(3), 616–626.

Rifkin, S. A., Houle, D., Kim, J., & White, K. P. (2005). A mutation accumulation assay reveals a broad capacity for rapid evolution of gene expression. *Nature*, *438*(7065), 220–223.

Scoville, A. G., Lee, Y. W., Willis, J. H., & Kelly, J. K. (2011). Explaining the heritability of an ecologically significant trait in terms of individual quantitative trait loci. *Biology Letters*, *7*(6), 896–898.

Shi, X., Ng, D. W.-K., Zhang, C., Comai, L., Ye, W., & Chen, Z. J. (2012). Cis- and trans-regulatory divergence between progenitor species determines gene-expression novelty in Arabidopsis allopolyploids. *Nature Communications*, *3*, 950.

Signor, S. A., & Nuzhdin, S. V. (2018). The Evolution of Gene Expression in cis and trans. *Trends in Genetics: TIG*, *34*(7), 532–544.

Sobel, J. M., & Streisfeld, M. A. (2013). Flower color as a model system for studies of plant evo-devo. *Frontiers in Plant Science*, *4*, 321.

Springer, N. M., & Stupar, R. M. (2007). Allele-specific expression patterns reveal biases and embryo-specific parent-of-origin effects in hybrid maize. *The Plant Cell*, *19*(8), 2391–2402.

Stern, D. L., & Orgogozo, V. (2008). The loci of evolution: how predictable is genetic evolution? *Evolution; International Journal of Organic Evolution*, *62*(9), 2155–2177.

Stern, S., Dror, T., Stolovicki, E., Brenner, N., & Braun, E. (2007). Genome-wide transcriptional plasticity underlies cellular adaptation to novel challenge. *Molecular Systems Biology*, *3*(1), 106.

Streisfeld, M. A., & Rausher, M. D. (2009). Altered trans-regulatory control of gene expression in multiple anthocyanin genes contributes to adaptive flower color evolution in Mimulus aurantiacus. *Molecular Biology and Evolution*, *26*(2), 433–444.

Sweigart, A., Karoly, K., Jones, A., & Willis, J. H. (1999). The distribution of individual inbreeding coefficients and pairwise relatedness in a population of mimulus guttatus. *Heredity*, *83 (Pt 5)*, 625–632.

Tajima, F. (1989). Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics*, *123*(3), 585–595.

Team, R. C., & Others. (2013). *R: A language and environment for statistical computing*. https://cran.microsoft.com/snapshot/2014-09-08/web/packages/dplR/vignettes/xdate-dplR.pdf

Tirosh, I., Reikhav, S., Levy, A. A., & Barkai, N. (2009). A yeast hybrid provides insight into the evolution of gene expression regulation. *Science*, *324*(5927), 659–662.

Torres, J. M., Gamazon, E. R., Parra, E. J., Below, J. E., Valladares-Salgado, A., Wacher, N., Cruz, M., Hanis, C. L., & Cox, N. J. (2014). Cross-tissue and tissue-specific eQTLs: partitioning the heritability of a complex trait. *American Journal of Human Genetics*, *95*(5), 521–534.

Troth, A., Puzey, J. R., Kim, R. S., Willis, J. H., & Kelly, J. K. (2018). Selective trade-offs maintain alleles underpinning complex trait variation in plants. *Science*, *361*(6401), 475–478.

Tulchinsky, A. Y., Johnson, N. A., & Porter, A. H. (2014). Hybrid incompatibility despite pleiotropic constraint in a sequence-based bioenergetic model of transcription factor binding. *Genetics*, *198*(4), 1645–1654.

Tung J., Zhou X., Alberts S.C., & Stephens M, Gilad Y. (2015). The genetic architecture of gene expression levels in wild baboons. *eLife* 4:e04729.

Venables, W. N., & Ripley, B. D. (2013). *Modern Applied Statistics with S-PLUS*. Springer Science & Business Media.

Verta, J.-P., Landry, C. R., & MacKay, J. (2016). Dissection of expression-quantitative trait locus and allele specificity using a haploid/diploid plant system--insights into compensatory evolution of transcriptional regulation within populations. *The New Phytologist*, *211*(1), 159–171.

Wang, Z., Sun, X., Zhao, Y., Guo, X., Jiang, H., Li, H., & Gu, Z. (2015). Evolution of gene regulation during transcription and translation. *Genome Biology and Evolution*, *7*(4), 1155–1167.

Whitehead, A., & Crawford, D. L. (2006). Variation within and among species in gene expression: raw material for evolution. *Molecular Ecology*, *15*(5), 1197–1211.

Willis, J. H. (1993). Partial self-fertilization and inbreeding depression in two populations of Mimulus guttatus. *Heredity*, *71*(2), 145–154.

Willis, J. H. (1999). THE ROLE OF GENES OF LARGE EFFECT ON INBREEDING DEPRESSION IN MIMULUS GUTTATUS. *Evolution; International Journal of Organic Evolution*, *53*(6), 1678–1691.

Wittkopp, P. J., Haerum, B. K., & Clark, A. G. (2004). Evolutionary changes in cis and trans gene regulation. *Nature*, *430*(6995), 85–88.

Zhou, X., & Stephens, M. (2014). Efficient multivariate linear mixed model algorithms for genome-wide association studies. *Nature Methods*, *11*(4), 407–409.

Figure Legends

Figure 1. Genome-wide associations of gene expression. Above: Transcript levels were normalized using (A) log(count+1) or (B) Box-Cox. Associations are designated as cis (pink) or trans (blue). Chromosomes are numbered along both axes. Points are larger in panel B to aid visualization. The putative misassembly is indicated with a red asterisk. The three genes with the most trans-associations are indicated with blue asterisks. Distributions of read counts for a representative gene, Migut.D00004, which had 3,665 SNP associations with Log(counts+1) transformation (C), but none with Box-Cox transformation (D).

Figure 2. (A) Minor allele frequency distribution for all associations excepting Chromosome 11. (B) The subset of associations in the top quartile of effect sizes for each regulatory category. Cis-acting variants in pink and trans-acting variants in blue. Lowest MAF bin is 0.025 < x < 0.035.

Figure 3. The proportion of significant cis-tests from the real data (blue), simulations with constant β (orange), and constant V_{snp} (grey) are reported for 10 AFS categories (minor allele q=0.0-0.05, 0.05-0.1, etc). The panels indicate the proportions obtained by imposing different significance thresholds to call significance.

Figure 4. Purifying selection on trans-effectors of highly connected genes. (A) The distribution of connectedness (as measured by R^2 between a gene and its module expression) for genes with associated cis (pink) and trans (blue) acting variants. (B) The average minor allele frequency of sites affecting each gene in a given connectedness quartile, separated by cis- and trans-acting variants. Each data point in (B) is a gene, which is assigned a quartile and the MAF of sites affecting it is calculated and plotted on the Y-axis.

Figure 5. Gene expression predicts trait covariances. The bottom triangle shows trait correlations. A line denotes a significant correlation at p < 0.05. The diagonal displays normalized trait histograms. Fraction covariance explained by gene expression, using the best-fit model coefficients for prediction, is shown in the top diagonal. Any displayed number is significant, asterisks denote levels of significance (determined by permutation, p < 0.05, p* < 0.025, p < 0.01, p*** < 0.001).**

Tables

Table 1. Variation in traits explained by the AIC best-fit model of gene expression modules. Significance was established by permuting either the module eigen values across lines, or the module gene composition. Bold indicates significance with *: p < 0.05, **: p < 0.01, ***: p<0.001.

Trait	R^2	Modules	H_0: Gene expression does not predict trait variation	H_0: Modules do not predict traits better than random groups of genes
Days to germination	0.289	17	0.094	0.203
Days to flower	0.246	12	0.253	0.516
Corolla width	0.392	19	**0.005****	**0.040***
Corolla length	0.435	17	**0.000*****	**0.015***
Tube length	0.389	19	**0.002****	0.078
Throat width	0.309	11	**0.044***	0.190
Stigma length	0.382	18	**0.005****	0.071
Anther length	0.400	19	**0.001****	**0.032***
Height	0.333	11	**0.014***	0.341
Node	0.314	8	**0.036***	0.639
Widest leaf	0.231	6	0.294	0.891
Flower size PC1	0.472	20	**0.000*****	**0.006****
Flower size PC2	0.270	9	0.149	0.501