Genome-phenome wide association in maize and Arabidopsis identifies a common molecular and evolutionary signature

Zhikai Liang[1,2], Yumou Qiu[3] and James C. Schnable[1,2*]

1. Department of Agronomy and Horticulture, University of Nebraska-Lincoln, Lincoln, NE, USA
2. Plant Science Innovation Center, University of Nebraska-Lincoln, Lincoln, NE, USA
3. Department of Statistics, Iowa State University, Ames, IA, USA

* Corresponding author: James C. Schnable; Email: schnable@unl.edu

Running title

Genome-phenome wide association in multiple plants

Abstract

Linking natural genetic variation to trait variation can help determine the functional roles different genes play. Often variation of one or several traits are assessed separately. High throughput phenotyping and data mining can capture dozens or hundreds of traits from the same individuals. Here we test the association between markers within a gene and many traits simultaneously. This Genome-Phenome Wide Association Study (GPWAS) is both a multi-marker and multi-trait test. Genes identified using GPWAS with 260 phenotypic traits in maize were enriched for genes independently linked to phenotypic variation. Traits associated with classical mutants were consistent with reported phenotypes for mutant alleles. Genes linked to phenomic variation in maize using GPWAS shared molecular, population genetic, and evolutionary features with classical mutants in maize. Genes linked to phenomic variation in Arabidopsis using GPWAS are significantly enriched in genes with known loss of function phenotypes. GPWAS may be an effective strategy to identify genes where loss of function alleles will produce mutant phenotypes. The shared signatures present in classical mutants and

genes identified using GPWAS may be a marker for genes with a role in specifying plant phenotypes generally, or pleiotropy specifically.

Key words

Quantitative genetics; phenomics; genome wide association studies; mutant phenotypes; maize; Arabidopsis

Introduction

In multicellular eukaryotes, only a small proportion of all annotated genes have yet to be linked to loss of function phenotypes (Schnable and Freeling, 2011; Schofield et al., 2012; Lamesch et al., 2012; Rhee and Mutwil, 2014; Chong et al., 2015; Schnable, 2019). Even in extensively studied single celled eukaryotes such as fission yeast (*Schizosaccharomyces pombe*) where the number of annotated genes is less and it is easier to screen for fitness effects under many environmental conditions there are thousands of genes where predicted gene functions have yet to be supported by a loss of function phenotype (Lock et al., 2019). The functions of many other genes have been inferred from quantitative genetic analyses. Arguably, the first such quantitative genetic association was the identification of a seed size QTL in dry bean (*Phaseolus vulgaris*) in 1923. This study used a single genetic marker, which was a qualitative trait controlled by a single gene (Sax, 1923). Soon after, quantitative trait variation could be linked directly to chromosome structural markers (Sprague, 1941). Technology for scoring genetic markers continued to advance, making it possible to genotype markers covering the entire genome across a population. This enabled Genome Wide Association Studies (GWAS) employing the linkage disequilibrium (LD) present in natural populations to identify functionally variable alleles of a gene influencing variation in a target trait (Eagle, 2006; Dewan et al., 2006; Atwell et al., 2010). The vast majority of loss of function mutations affect multiple traits. However, the majority of current quantitative genetics approaches seek to identify either genetic markers or genes linked to single phenotypes, with a subset considering data from multiple correlated phenotypes (Zhou and Stephens, 2012; O'Reilly et al., 2012; Korte et al., 2012; van der Sluis et al., 2013; Stephens, 2013; Wang et al., 2015; Turley et al., 2018; Pitchers et al., 2019). It is now feasible to collect data for thousands of intermediate molecular phenotypes, such as

transcript, protein, or metabolite abundance, from entire association populations and incorporate these data into quantitative genetics models as either explanatory (Lin et al., 2017; Kremling et al., 2019) or response variables (Wen et al., 2014; Matsuda et al., 2015; Diepenbrock et al., 2017; Kremling et al., 2018). Advances in high-throughput plant phenotyping have expanded the capacity of these techniques to score dozens or hundreds of whole-organism phenotypes across multiple time points and environments (Walter et al., 2015; Araus et al., 2018). Incorporating data on large sets of phenotypes scored in the same populations, including both correlated and uncorrelated traits, may aid in the identification of genes that, like the vast majority of classical loss of function mutants, play roles in controlling variation in multiple traits within an organism.

Here we employ a published dataset of 260 distinctly scored traits for 277 resequenced maize inbred lines (Flint-Garcia et al., 2005; Bukowski et al., 2018) to develop and evaluate a novel approach to identify the links between genes and quantitative phenotypic variation *per se* using a multi-trait multi-SNP framework. We demonstrate that the genes identified using this method, which we call Genome-Phenome Wide Association Study (GPWAS), show substantially greater cross-validation in an independent study using data from approximately 20 times as many individuals (Wallace et al., 2014) than do genes identified using conventional GWAS analysis of the same dataset. For a wide range of features, including expression level and breadth, syntenic conservation, purifying selection in related species, and the prevalence of presence-absence variation (PAV) across diverse maize lines, the genes identified using this multi-trait multi-SNP approach appear more similar to genes identified using forward mutagenesis, and less similar to the overall population of annotated maize gene models.

Results

Conceptual design of the GPWAS model

Scoring plant phenotypes from genetically identical populations across multiple environments and/or extracting plant phenotypes from different types of sensor data can produce high(er) dimensional trait datasets. We sought to develop and evaluate a model which could incorporate data from many traits, including uncorrelated traits, and

identify associations between genes and trait variation in practical amounts of computational time. Briefly our model requires two data matrices: one containing allele calls for many genetic markers across individuals in a population, and a second containing observed values for the same individuals across many traits. One or more genetic markers are assigned to a given gene or other genomic interval of interest. These markers are then treated as a response variable. The significance of the association between the gene or other genomic interval of interest and trait variation (considered as a whole) is determined by the comparison of two models. The first model seeks to predict genotype values for genetic markers in the target interval using solely information on population structure. The second model seeks to predict genotype values for genetic markers in the target interval using both trait data and population structure data using a stepwise selection procedure (see Methods and Figure 1).

Application of both GPWAS and GWAS to data from a maize diversity panel

Being able to employ and evaluate GPWAS required high density genotype and phenotype data collected from a common population. For this initial application and evaluation, data was collected from public sources. Genetic marker data were obtained from resequencing data of 277 inbred lines from the Buckler-Goodman maize association panel (Flint-Garcia et al., 2005). These lines are part of Maize HapMap3, which contains data for a total of 81,687,392 SNPs (Bukowski et al., 2018). After removing the SNPs with high levels of missing data, those that were not polymorphic among the 277 individuals employed here, and several other quality filtering parameters, 12,411,408 SNPs remained. Of these, 1,904,057 SNPs were assigned to 32,084 annotated gene models from the B73 RefGenV4 genome release. Filtering to eliminate redundancy between SNPs assigned to the same gene in high LD with each other reduced this number to 557,968 highly informative SNPs. A phenotypic dataset consisting of 57 specific traits scored for the Buckler-Goodman maize association panel across 1 to 16 distinct environments for a total of 285 unique phenotypic datasets was obtained from Panzea (Zhao et al., 2006). Removing datasets with extremely high levels of missing data resulted in 260 trait datasets with a median missing data rate of 18%. Of the total 72,020 potential trait datapoints (277 inbred lines $\times$ 260 traits), 23.6% or 16,963 trait datapoints were missing. Missing trait datapoints were imputed

using PHENIX a method which employs a combination of both kinship and between trait correlations to impute missing values in large-dimensional phenotypic datasets collected from populations (Dahl et al., 2016), and the estimated imputation accuracies for the individual traits are reported in Supplemental Data 1A. A conventional GWAS analysis generally employs either empirically determined statistical significance cutoffs (Wallace et al., 2014), or a Bonferroni correction based on the total number of hypothesis tests (International HapMap Consortium and others, 2005) or the number of "effective number" M_{eff} of independent hypothesis tests conducted (Li and Ji, 2005). For the above dataset, employing a naive Bonferroni correction would mean each individual analysis would be conducted using a multiple-testing corrected p-value cutoff of $8.96e^{-08}$, while a sequential analysis of all 260 traits should employ a multiple-testing corrected p-value of $3.45e^{-10}$. As shown in Figure 2A, a given gene might be identified in multiple independent GWAS analyses for individual traits but not be considered significantly associated with any traits when correcting for the total number of traits analyzed. In the example given, Zm00001d002175 shows a statistically significant association with flowering time in multiple environments, yet none of these associations are individually significant enough to meet the threshold for the full multiple testing correction.

Bonferroni multiple testing correction assumes that each test is independent of all other tests, however, the different trait datasets collected from the Buckler-Goodman association panel exhibited significant correlation (Figure 2C), including three large blocks of traits related to flowering time, plant architectural traits, and tassel structure traits respectively. To address the challenges of partially correlated traits and partially correlated genotype matrices, we developed an approach based upon a stepwise regression model fitting.

Multiple testing was corrected using a permutation-based method (see Methods), which controls for the complexities introduced by iterative model selection. It must be noted that this procedure controls the overall false discovery rate (FDR) among the population of genes identified as showing a statistically significant link between genotypic and phenotypic variation. It is not an estimate of Family Wise Error Rate (FWER) which would be obtained by only considering the single most significant p-value obtained for any gene in each permutation. Although computationally expensive, permutation has been shown to be robust for controlling false positives in both GWAS

and PheWAS studies (Tian et al., 2011; Namjou, B et al., 2014). Based on the permutation analysis, a p-value cutoff of $1.00e^{-23}$ resulted in the classification of 1,776 genes as being significantly associated with phenomic variation in the Buckler-Goodman association panel, resulting in an estimated FDR < $1.00e^{-3}$. Imputation accuracy varied significantly among individual traits. To test whether it might be beneficial to exclude traits with high missing data rates and low imputation accuracy, the effect of adding additional simulated zero-heritability phenotypes to the trait matrix was evaluated. Log transformed p-values for individual genes were highly correlated (R^2 = 0.96) before and after the addition of simulated zero heritability traits. All 260 traits were retained for downstream analysis as low information content traits may still provide some value and, based on the test above, appear to be, at worst, benign. For comparison purposes, the same set of traits and genotypes was also tested for associations using three conventional GWAS algorithms: a general linear model (GLM GWAS) (Price et al., 2006), a mixed linear model (MLM GWAS) (Yu et al., 2006; Price et al., 2006), and FarmCPU GWAS (Liu et al., 2016) (See Methods). Applying an equivalent permutation based FDR threshold to each conventional GWAS algorithm removed the vast majority of positive signals (Supplemental Figure 1). Therefore, for GWAS models, a conventionally multiple testing corrected p-value cutoff was employed (Supplemental Figure 1).

Validation of Gene-Phenome Associations

A second published dataset of genes identified as being associated with variation in trait values in the maize nested association mapping (NAM) population, which includes approximately 5,000 lines (McMullen et al., 2009), was employed to assess the relative power and accuracy of three conventional GWAS algorithms as well as the GPWAS algorithm (McMullen et al., 2009; Wallace et al., 2014). As the published data for the NAM population used B73 RefGenV2, all comparisons employed only the subset of 29,372 gene models with a clear 1:1 correspondence between gene models included in the B73 RefGenV2 and B73 RefGenV4 annotation versions. Of these, 4,227 genes were identified as being associated with at least one trait in the NAM dataset (Wallace et al., 2014). Genes identified using GPWAS showed significantly higher cross-validation in the NAM dataset than the sets of genes identified using GLM GWAS ($p = 2.05e^{-5}$; Chi-squared test; two-sided), MLM GWAS (p = 0.010; Chi-squared test; two-

sided), or FarmCPU GWAS (p = 0.013; Chi-squared test; two-sided) (Figure 3A and 3B; Supplemental Data 1B). Filtering to remove signals from rare SNPs where the minor allele was present in only one or two of the NAM population founder lines reduced the total number of genes identified in that study to 3,621. However, the overall trend observed remained consistent and statistically significant, with the genes identified using the GPWAS algorithm continuing to show statistically significantly higher rates of identification in the reduced NAM dataset (GLM GWAS, $p=1.63e^{-4}$; MLM GWAS, p=0.002; FarmCPU GWAS, p=0.025; Chi-squared test; two-sided) (Supplemental Data 1B). Analyses with two smaller real-world datasets for biochemical traits related to vitamin A (24 traits) and vitamin E (20 traits) metabolism (Owens et al., 2014; Diepenbrock et al., 2017) did not reveal any significant increase in the number of *a priori* gene candidates identified as showing a link to phenotypic variation relative to conventional GWAS approaches (Figure 3C). Conventional GWAS showed substantial advantages in power at low false discovery thresholds when compared to GPWAS for a single trait in simulation studies, while GPWAS showed significant advantages in power FDR trade offs when data on multiple traits (e.g more than 5 traits) was integrated into the analysis (Figure 3D). These simulations likely overstate GPWAS's performance advantage as phenotypes are likely to exhibit different degrees of genetic architecture complexity and differing degrees of shared v.s. unique causal loci.

Our GPWAS algorithm also produces a list of the specific traits included in the model for a given gene (Supplemental Data 1C). For example, in Figure 2B, the overall association between Zm00001d002175 and the trait dataset was statistically significant. The 11 individual traits included in the Zm00001d002175 model included both flowering time measured in multiple locations, as well as additional traits with indirect links to flowering time (e.g. number of leaves, Summer 2008, Cayuga, NY), and others with no obvious links to flowering time. These included the total kernel volume in one year in one location and kernel proteins as estimated using near infrared imaging in another year in a different location.

GPWAS Accurately Predicts Pleiotropic Consequences of Gene Knockouts

It is important to keep in mind that the associations of individual phenotypes identified

within the model are not rigorously controlled for false discovery. We therefore sought to qualitatively evaluate whether traits included in the model for an individual gene make sense in the context of existing detailed biological knowledge about the function of a given gene. One such gene was *anther ear1* (*an1*), a classical maize gene encoding an ent-copalyl diphosphate synthase involved in gibberellic acid biosynthesis, for which knockout alleles have been shown to reduce or abolish tassel branching, reduce plant height, delay growth, and delay flowering (Bensen et al., 1995). In a separate analysis of the 5,000 individual maize NAM lines, *an1* was identified as being associated with one trait, tassel spike length (Brown et al., 2011), however, it was not found to be associated with any individual traits through a conventional GWAS analysis of the Buckler-Goodman 282 dataset. GPWAS identified a statistically significant link between *an1* and a model incorporating multiple phenotypes including flowering time, plant height, and tassel branch number, all consistent with the known mutant phenotypes (Figure 4A and 4C). At least one additional phenotype included in the GPWAS model -- germination count (Summer 2006, Johnston, NC) -- was not supported by direct reports of characterization of the *an1* knockout allele, but is consistent with the role of *an1* in gibberellic acid metabolism (Peng and Harberd, 2002; Landoni et al., 2007). Overall, the set of phenotypes identified using GPWAS for the *an1* gene appeared to be consistent with previously reports based on either the characterization of the knockout allele or quantitative genetic analyses of natural populations. To disambiguate the effects of the multiple SNP and multiple traits portions of this analysis, the same multiple marker analysis for *an1* was conducted separately for each individual trait from the set of 260. This approach identified a larger total number of traits than a joint analysis of all 260 traits together. However, individually significantly associated traits tended to represent a smaller number of phenotype groups (i.e. multiple correlated measurements of the same phenotype in different environments) and failure to capture some of the traits consistent with the known function and mutant phenotypes of *an1* (Supplemental Figure 2).

The GPWAS model also identified *liguleless2* (*lg2*), another classical maize mutant with a well characterized knockout mutant phenotype (Brink, 1935). The *lg2* encodes a bZIP transcription factor (Walsh et al., 1998). The loss of *lg2* function disrupts the establishment of the ligule and auricle of the maize leaf and results in plants with extremely erect leaves (Brink, 1935; Harper and Freeling, 1996). Lines carrying *lg2*

knockout alleles have been reported to exhibit substantially (10-50%) higher grain yield than otherwise isogenic hybrids (Pendleton et al., 1968; Lambert and Johnson, 1978), reduced tassel branch numbers (Lambert and Johnson, 1978; Walsh and Freeling, 1999), and moderately increased central spike length (Walsh and Freeling, 1999). Quantitative genetic analyses have identified signals for leaf angle, tassel branch number, and kernel row number associated with the *lg2* locus (Walsh and Freeling, 1999; Tian et al., 2011; Brown et al., 2011; Li et al., 2018), although the effect on kernel row number was not significant in at least one study utilizing null alleles of *lg2* (Walsh and Freeling, 1999). In our study, GPWAS identified a statistically significant link between *lg2* and a model incorporating multiple phenotypes including upper leaf angle, leaf length, central spike length, kernel weight (a yield component trait), and cob diameter. Cob diameter exhibits substantial correlation and overlapping genetic architecture with kernel row number (Liu et al., 2015) (Figure 4B). The GPWAS model for *lg2* also incorporated a number of flowering-time related traits, which do not have consistent support in either the characterization of *lg2* knockout mutants, or previous quantitative genetic analyses of flowering time in maize. Despite this, knockout alleles of *lg2* have been reported to alter the vegetative-to-reproductive phase transition in maize and produce increased numbers of leaves on the main stalk, which would be consistent with its altered flowering time (Walsh and Freeling, 1999; Liu et al., 2015). As in the case of *an1*, the traits identified as being associated with *lg2* using GPWAS appear to be largely consistent with previous characterization of the functional roles of *lg2* in maize (Figure 4B and 4D).

Individual case studies such as the ones presented above can be misleading. As a control, we set out to identify similar case studies for genes identified by conventional GWAS. A total of seven classical maize mutants from the list in (Schnable and Freeling, 2011) were identified as linked to one or more traits by one of the three GWAS models tested: GLM 5, MLM 1, FarmCPU 1 (Supplemental Data 1D). In many cases there was no apparent link between the known gene function and the trait where a significant association was identified. For example, a gene -- *glossy8* -- involved in the reduction of ketones as part of the biosynthesis of cuticle waxes showed a statistical association with plant tillering, and a transcription factor regulating the production of anthocyanin -- *colored alurone1* -- showed a statistical association with cob diameter (Supplemental Data 1D) (Paz-Ares et al., 1987; Xu et al., 1997). A mutant involved in sex determination

-- *indeterminate spikelet1/Tasselseed6* -- showed a link to flowering time and a second gene involved in cuticular wax production -- *glossy1* -- showed a link to a disease resistance phenotype (Supplemental Data 1D) (Chuck et al., 1998; Sturaro et al., 2005). However, we were not able to identify any stronger links between GWAS trait associations and classical maize mutants in this dataset.

Greater Functional Specificity of Genes Identified Using GPWAS

Genes identified using GPWAS appear to be a significantly less random sample of total gene models than the set of genes identified using GLM GWAS. A set of 1,406 genes were uniquely identified using GPWAS but not GLM GWAS. An equivalent set of 1,630 genes were identified using GLM GWAS but not GPWAS. In the larger unique-to-GLM GWAS gene set, a single Gene Ontology (GO) term showed a statistically significant bias towards being associated with phenotypic variation (GO:0046034: ATP metabolic process), and two GO terms with nearly identical gene assignments showed a statistically significant bias towards not being associated with phenotypic variation (GO:0000723: Telomere maintenance and GO:003220 Telomere organization). However, the moderately smaller set of genes uniquely identified using GPWAS was enriched or purified for the presence of many more GO terms. A total of 71 GO terms were overrepresented in the unique-to-GPWAS (relative to GLM GWAS) gene set to a statistically significant degree, including numerous terms linked to development, hormone signalling, response to different stimuli, and cell growth (Supplemental Data 1E). The 13 GO terms that were underrepresented among genes uniquely identified using the GPWAS algorithm were generally associated with DNA conformation and replication (Supplemental Data 1E). A similar comparison was made between genes uniquely identified using GPWAS and FarmCPU GWAS. In this case only 706 genes were uniquely identified using FarmCPU. As it is more likely for an enrichment or purification to be statistically significant in larger populations, only the 706 most significant unique-to-GPWAS (relative to FarmCPU GWAS) genes were evaluated in this comparison to eliminate any potential bias. Among the unique-to-FarmCPU GWAS gene set, only a single GO term was overrepresented to a statistically significant degree (GO:0051707: Response to other organism). However, among the the unique-to-GPWAS (relative to FarmCPU GWAS) gene set of equal size, 39 GO terms showed a statistically significant overrepresentation, while another 4 were statistically

underrepresented (Supplemental Data 1E).

Several potential factors could explain the large difference in GO enrichment purification we observed between genes identified solely using GWAS and genes identified solely using GPWAS. A number of factors, including the number of GO terms per gene and the proportion of genes with no assigned GO term, differed modestly between the different populations of genes (Supplemental Data 1F). The specificity of GO terms was higher for genes identified using only GPWAS than for genes identified using only GWAS. GO terms assigned to genes identified using only GPWAS were assigned to a median of only 430 other genes. GO terms assigned to genes identified using only GLM GWAS were assigned to a median of 514 other genes. This difference in the number of genes that a given GO term is assigned does not appear to explain the differences observed in the enrichment or purification (Supplemental Figure 3). Rather, the large differences observed here are consistent with GWAS identifying a more random subset of annotated genes as being associated with phenotypic variation than did GPWAS.

Molecular, Structural, and Evolutionary Features of Genes Identified Using GPWAS

Genes identified using the GPWAS algorithm differed from the overall population of annotated maize gene models in a number of characteristics, as well as from the populations of genes identified using conventional GWAS. In many cases, the properties of genes identified using GPWAS appeared more similar to the population of genes with validated loss-of-function phenotypes (Schnable and Freeling, 2011). Slightly less than half of all annotated maize genes were expressed to a level above 1 fragment per kilobase of transcript per million mapped reads (FPKM) on the average of the 92 tissues/time points assayed (Schnable and Freeling, 2011; Stelpflug et al., 2016). Among genes identified using any of the three conventional GWAS algorithms, more than 2/3rds had average expression levels >1 FPKM. Among genes identified using the GPWAS algorithm and among genes with validated loss of function phenotypes, more than 3/4th had average expression levels >1 FPKM (Figure 5A; Supplemental Data 1B). Genes identified using GLM GWAS, MLM GWAS, FarmCPU GWAS, GPWAS, and the classical mutants all exhibited greater breadths of expression across tissues, larger

numbers of genes with observed evidence of translation, and greater gene lengths than the population of annotated genes as a whole (Supplemental Data 1B). The number of associated SNPs was positively correlated with the log-transformed inverse p-value assigned to genes using both GWAS (r = 0.566) and GPWAS (r = 0.625) (Supplemental Figure 4; Supplemental Data 1G). However, this association declined dramatically in the permuted data for GPWAS (median permuted r = 0.155), but remained high for GWAS (median permuted r = 0.626) (Supplemental Data 1H). This suggests that the high number of SNPs per gene for GPWAS (median: 43 SNPs, mean: 47.3 SNPs) relative to the overall gene set (median: 12 SNPs, mean: 17.4 SNPs) is a biological property of the genes controlling phenotypic variation in this population, rather than reflecting a bias in the GPWAS algorithm.

On a population and comparative genomics level, genes identified using the GPWAS algorithm also differed from the overall population of annotated maize gene models, and looked more like genes with validated loss-of-function phenotypes. Genes identified using both the conventional GWAS and GPWAS algorithms were significantly less likely to exhibit PAV in the maize populations (Figure 5B) than the overall population of maize gene models. The reduction in PAV frequency for genes identified using GPWAS (7.0%) was significantly greater than for genes identified only using GWAS (10.4%) (p=0.0015; Chi-squared test; two-sided), and not statistically significantly different from low level of presence absence variation observed for maize genes with validated loss of function phenotypes genes (4.1%) (p=0.36; Chi-squared test; two-sided) (Supplemental Data 1I). Genes identified using either conventional GWAS and GPWAS algorithms were significantly more likely to be conserved at syntenic orthologous locations in sorghum than the overall set of maize gene models (Figure 5C). Genes uniquely identified using GPWAS were more likely to be conserved at syntenic locations in the genome of sorghum (*Sorghum bicolor*) (91.8%) than those uniquely identified using GWAS (74-85%; see Supplemental Data 1I). This difference was statistically significant in comparison to all three GWAS algorithms tested and was comparable to the likelihood of syntenic conservation for maize genes with known loss of function mutant phenotypes (93.9%) (Supplemental Data 1I).

The genes identified as being associated with phenotypic variation using GPWAS also appeared to be under stronger purifying selection than either the overall population

of maize gene models or those identified using any of the three conventional GWAS algorithms (Figure 5D; Supplemental Data 1I). This analysis was constrained to the subset of gene models with conserved orthologs in sorghum (*Sorghum bicolor*), and foxtail millet (*Setaria italica*). Among these genes, those uniquely identified using GPWAS showed a reduced ratio of nonsynonymous substitution rate to synonymous substitution rate (Ka/Ks) (median: 0.168-0.169; mean 0.208-0.210), relative to the overall population of syntenically conserved maize gene models (median: 0.200; mean: 0.246), while those uniquely identified using GWAS showed elevated rates (median: 0.202-0.233; mean: 0.251-0.261) relative to the same overall population (Supplemental Data 1I). Among the maize genes with characterized loss-of-function phenotypes, this ratio declined even further (median: 0.144; mean: 0.177). In short, the typical annotated gene appears to experience notably less purifying selection than those associated with organismal-level phenotypic variation based on either characterized loss-of-function mutant phenotypes or those identified using the GPWAS, but not a GWAS, algorithm.

Genes identified using GPWAS in *A. thaliana* are enriched in reported loss of function phenotypes

Genotype and phenotype data were curated from a subset of the Arabidopsis 1001 Genomes Project (https://1001genomes.org/data/GMI-MPI/releases/v3.1/) and four independent publications reporting multiple phenotypic datasets (Atwell et al., 2010; Exposito-Alonso et al., 2019; Julkowska et al., 2017; Li et al., 2010) deposited in AraPheno, respectively. After excluding lines absent from the resequencing dataset and those phenotyped for an insufficient number of traits, a final dataset of 158 traits scored for 164 *A. thaliana* genotypes was obtained, alongside a set of 208,236 SNP markers assigned to *A. thaliana* gene models. Data were analyzed as described above, resulting in the identification of 131 genes associated with trait variation in *A. thaliana* using the GPWAS algorithm and the same permutation based statistical significance cutoff employed for maize data. The difference in the number of genes confidently identified using GPWAS algorithm at a permutation estimated FDR < $1.00e^{-3}$ between maize and *A. thaliana* may reflect differences in the phenotypes scored in each dataset or differences genetic architecture, but the simplest and most likely explanation is that the large difference simply results from the smaller number of individuals and traits available in the *A. thaliana* dataset. Within the set of 131 genes identified as showing

significant association with trait variation in *A. thaliana* using the GPWAS algorithm, 21 genes (16%) belonged to a set of *A. thaliana* genes with reported loss of function phenotypes (Lloyd and Meinke, 2012). This is approximately 2x the frequency of genes with reported loss of function phenotypes in the *A. thaliana* genome as a whole (7%). It is also higher than the frequency of genes with loss of function phenotypes in the 131 most significantly associated genes identified using either GLM based GWAS (10%) or FarmCPU based GWAS (11%). The difference in the frequency of genes with reported loss of function phenotypes in the set of genes identified using the GPWAS algorithm and the background set is statistically significant (p=0.00026; Chi-Square Test).

The total number of genes identified using these *A. thaliana* datasets was much smaller than the number of genes identified using the same algorithm in maize (approximately 1/10 as many genes). Similarly, the number of genes identified using conventional GWAS algorithms was also lower in the Arabidopsis dataset than the maize dataset. Eight-hundred-thirty-five genes showed statistically significant associations with at least one phenotype using the GLM GWAS model, 326 genes using the FarmCPU model, and eight genes using the MLM GWAS model. We speculate that this is a result of only having data for a smaller number of phenotypes scored across a smaller number of unique genotypes. Neither GLM GWAS nor FarmCPU GWAS model was significantly more likely to identify genes with validated loss of function phenotypes phenotypes than the background population. The MLM GWAS gene set was too small to test. Genes with validated loss of function phenotypes share many properties between maize and Arabidopsis, including higher expression, greater syntenic conservation, and lower Ka/Ks ratios (Schnable, 2019). More detailed comparisons between maize and *A. thaliana* must await datasets with more phenotypes scored over more *A. thaliana* genotypes, and the identification of larger sets of *A. thaliana* genes associated with phenomic variation.

Discussion

Complex datasets can contain scores for dozens or hundreds of traits across the same populations. The prevalence of these datasets and the challenges and opportunities they present is expected to grow in the coming years. Here, we developed an approach for identifying genotype-phenotype associations that can scale to the analysis of

datasets containing hundreds, or potentially even thousands, of traits. The set of genes identified by this GPWAS algorithm showed significant differences from the overall population of gene models for a range of features. However, these differences do not appear to be results of biases within the GPWAS algorithm itself. They may instead reflect biological differences between genes associated with phenotypic variation and other annotated gene models in the genome.

The statistical tests upon which the GPWAS approach is built become unstable once the number of traits exceeds the number of individuals scored, therefore, scaling to high numbers of traits would require the use of larger association populations than many of the most general used plant populations today (Flint-Garcia et al., 2005; Atwell et al., 2010; Huang et al., 2010; Morris et al., 2013). Multicollinearity in either the predictor or response variables can make the statistical estimation and inference procedures we employed unstable (Rencher and Schaalje, 2008). In cases where the number of measured traits exceeds the number of environments, it would be advisable to employ alternative approaches to reduce the dimensionality of the trait dataset, whether that be an *ad hoc* approach such as selecting a subset of representative traits from highly correlated blocks, or dimensional reduction analyses such as principal component analysis or multidimensional scaling. The automatic application of variable selection and/or dimensional reduction in such scenarios could be incorporated into future GPWAS implementations. The selection of the number of principal components to include in quantitative genetic tests is a matter of ongoing debate. Including too few PCs results in false positive associations driven by population structure, while including too many overcontrols population structure and dramatically reduces the power. A plot of the cumulative variance of the proportion of genetic variation in the maize 282 panel did not show a clear "elbow" and less than 40% of total genetic variation was explained when 30 PCs were included (Supplemental Figure 5A). While we chose to employ three PCs in the primary results presented here, a parallel analyses employing 20 PCs to control population structure for both GPWAS and GWAS algorithms recovered the same distinguishing signatures of genes associated with phenotypic variation relative to the overall set of gene models (Supplemental Figure 5B-E), although with substantially reduced numbers of total positive genes for all algorithms. This suggests that the pattern reported here is not an artifact of particular parameter selection decisions. Both multivariate tests of genotype phenotype association and multi-marker based tests have

been previously explored (Trégouët et al., 2009). Conventional approaches to multivariate GWAS analysis require that the set of traits being tested simultaneously be correlated (Turley et al., 2018; Pitchers et al., 2019). These algorithms cannot scale efficiently to testing associations with hundreds of traits simultaneously on a genome-wide scale, as adding more traits in multivariate GWAS leads to exponential increases in computational cost (Korte et al., 2012). GPWAS can iteratively search through large sets of correlated and uncorrelated traits, allowing it to utilize datasets, such as the one discussed here, which it is not possible to analyze using conventional multi-trait GWAS methods. While this provides advantages in terms of scalability, it does come at some cost. Firstly, as presently implemented, GPWAS cannot estimate marker/gene effects as is possible with many conventional GWAS algorithms. Secondly, as a result of the stepwise selection procedure, the set of traits identified as associated with a gene in a GPWAS analysis is unlikely to be exhaustive, particularly when multiple traits are closely correlated. For example, tassel spike length and tassel length are generally correlated including in this dataset. GPWAS identified variation in tassel length but not spike length as associated with *an1*. However, if tassel length is removed from the dataset, spike length becomes one of the phenotypes selected by the GPWAS model for *an1*. This manuscript is not intended to be an exhaustive comparison of methods for linking genotype and phenotype, but rather to highlight features associated with genes linked to phenomic variation which were either indetectable or less obvious using existing widely adopted methods.

A challenge for the present implementation of GPWAS is that it requires regions of interest to be defined across the genome. In this study, annotated gene models were used to define these regions, however, approximately 40% of the phenotypic variation in maize has been estimated to be explained by noncoding regulatory regions (Rodgers-Melnick et al., 2016). These regions can be separated from the genes whose expression they control by many kilobases (Studer et al., 2011; Castelletti et al., 2014), while LD in maize generally decays within one to several kilobases (Remington et al., 2001; Romay et al., 2013). Both sequence conservation and chromatin mark data could be used to define additional regions of interest likely to represent regulatory sequences (Zhang et al., 2012; Turco et al., 2013; Oka et al., 2017). Similar approaches could also be employed to identify currently unannotated regions of the genome with a high potential for containing cryptic genes, including functional long noncoding RNAs (Lloyd

et al., 2018). Finally, as presently implemented, GPWAS assumes the availability of multiple markers per gene -- or other region of interest -- genotyped across an association population. Presently this limits its applicability to a small number of species such as rice, Arabidopsis and maize. However the declining cost of obtaining high density genotypic data, via methods such as whole genome resequencing, genotyping by RNA-seq, or exom-seq means this barrier is likely to become less significant in the future.

The genes identified as linked to phenotypic variation in maize using the GPWAS algorithm shared many common features with the smaller set of maize genes with validated loss of function phenotypes. The smaller get of genes identified as linked to phenotypic variation in *A. thaliana* showed a statistically significant enrichment among the larger set of *A. thaliana* genes with validated loss of function phenotypes. Genes linked to variation individual phenotypes using GWAS in maize exhibited intermediate values between all gene models and GPWAS/classical mutants. And genes linked to variation in individual phenotypes using GWAS in *A. thaliana* exhibited an enrichment in genes with known loss of function phenotypes which was not statistically significant but was intermediate between the overall set of annotated genes and the significant enrichment among genes identified using the GPWAS algorithm. One potential explanation for this result is that mutant analysis and GPWAS may be more likely to identify genes with pleiotropic effects while conventional GWAS has greater strength to identify loci where functional variation influences only a small number of phenotypic outcomes. In this scenario the molecular and evolutionary signatures shared by classical mutants and genes identified using GPWAS may be associated with genes where functional variation is likely to have pleiotropic consequences. However, many of the signatures associated with reported loss of function phenotypes in both maize and Arabidopsis are consistently found at about the same strength in genes associated with lethal, non-lethal but constitutive, or environmentally conditional phenotypes (Schnable and Freeling, 2011; Schofield et al., 2012; Lamesch et al., 2012; Rhee and Mutwil, 2014; Chong et al., 2015; Schnable, 2019). This would not be consistent with the signatures we identify among genes identified as linked to phenotypic variation in maize using the GPWAS algorithm being a marker for essentiality. A second potential explanation is that, because the GPWAS algorithm incorporates data from many traits collected in differing environments it may simply have a lower false positive rate than is

possible through the analyses of any single phenotype. In this scenario the molecular and evolutionary signatures shared by classical mutants and genes identified using GPWAS may be a marker of which gene models play a notable role in influencing plant phenotype.

Methods

Genotype and Phenotype Sources, Filtering, and Imputation

Raw genotype calls from the resequencing of the maize 282 association panel (Bukowski et al., 2018) were retrieved from Panzea in AGPv4 coordinates. Missing genotypes were imputed using Beagle (version: 2018-06-10) (Browning and Browning, 2016; Bukowski et al., 2018). Only biallelic SNPs with fewer than 20% missing data points were subjected to imputation. After imputation, SNPs with a minor allele frequency (MAF) of less than 0.05 or which were scored as heterozygous in more than 10% of samples were discarded. A phenotype file (traitMatrix_maize282NAM_v15-130212.txt) containing a total of 285 traits, corresponding to 57 unique types of phenotypes scored in 1 to 16 environments was downloaded from Panzea. A set of 277 accessions with identical names in the HapMap3 data release and the Panzea trait data were employed for all downstream analyses.

Maize gene regions were extracted from AGPv4.39, which was downloaded from Ensembl. SNPs were clustered based on $R^2 > 0.8$ and only one randomly selected SNP per cluster was retained. If, after collapsing the highly correlated clusters, the number of SNPs exceeded 138 (50% of the number of inbred lines scored), a random subsample of 138 SNPs was employed for the downstream analyses. Identical final SNP sets were employed for the GPWAS and GWAS analyses.

Of the 285 initial trait datasets, 25 were removed because the data file contained a recorded trait value for only one individual, leaving a total of 260 trait datasets. Using a Bayesian multiple-phenotype mixed model (Dahl et al., 2016), missing phenotypes were imputed based on a kinship matrix calculated from 1.24 million SNPs generated using GEMMA (version: 0.94.1) (Zhou and Stephens, 2012; Dahl et al., 2016). For those traits with a sufficient number of real observations to enable evaluation, the accuracy of the phenotypic imputation was assessed independently by masking 1% available records

for each trait and comparing the imputed and masked values. This process was repeated 10x for each trait.

Calculating Principal Component Scores Used in Both GPWAS and GWAS

A subset of 1.24 million SNPs distributed across both intragenic and intergenic regions on all 10 chromosomes was used to perform PCA using the R prcomp function to provide controls for population structure. As the maize 282 panel exhibits relatively low population structure (Flint-Garcia et al., 2005), only the first three PCs were included in both GWAS and GPWAS analyses, however, comparable analyses could be run with different numbers of PCs included. In GPWAS, for analysis of the given gene on each chromosome, markers solely from the other 9 chromosomes were used to reduce the endogenous correlations between genes and principal components (Listgarten et al., 2012; Rincent et al., 2014). In GWAS, principal components were calculated using all of 1.24 million SNPs on 10 chromosomes.

GPWAS Analysis

All the operations for the GPWAS analyses are detailed in the R source code used to conduct the analysis -- and associated documentation -- which has been made available online (https://github.com/shanwai1234/GPWAS).

Briefly, we employed a model selection approach to adaptively select the most significant phenotypes associated with each gene. A F-test (one-sided) was used to compare a model to explain variation in SNPs based solely on population and a model which incorporated both population structure and trait data. The significance in the difference of the goodness of fit between these two models was used to determine the significance of the association of individual genes with phenotypic variation in the dataset, the Reduced Model (RM) based solely on population structure and the Phenotype Incorporating Model (PM) which used stepwise selection to include both population structure and trait data within its model. It should be noted that the direction of a regression does not require an assertion about the direction of causality. A regression model can be used to predict the value of a cause from a measured effect, or to predict the outcome from a known cause.

Reduced Model:

Let the subscripts *k* and *i* represent the *k*th individual and the *i*th gene. Let $g_{k,i}$ be the corresponding SNP values in that gene (*g*).

Here, $g_{k,i}$ is an *m*-dimensional vector of genotypes for SNPs where *m* is the number of distinct genetic markers associated with the gene after removal of SNPs in high linkage disequilibrium with each other.

We considered the multiple responses regression model:

$$(1)\, g_{k,i} = \sum_{q=1}^{v_{pc}} PC_{k,q}\beta_{iq} + \varepsilon_{k,i}$$

where *q* denotes the *q*th principal component, and v_{pc} is the number of the included principal components in (1). The regression coefficients β_{iq} and the errors $\varepsilon_{k,j}$ are all *m* dimensional vectors.

Stepwise Selection:

The final model begins with the initial model (1) and uses a stepwise selection procedure (Draper, N. R. and Smith, H., 1998) to add additional traits as explanatory variables. In each iteration, we consider the model:

$$(2) g_{k,i} = \sum_{q=1}^{v_{pc}} PC_{k,q}\beta_{iq} + \sum_{j}^{\square} Phe^{c}_{k,j}\tau_{ij} + \varepsilon_{k,i}$$

where { $Phe^{c}_{k,j}$ } are the currently selected, and τ_{ij} is the corresponding coefficient for the phenotype $Phe^{c}_{k,j}$ of the *i*th gene. The goodness of fit for each trait in the model (2) was assessed for all SNP markers in the *i*th gene jointly.

In detail, least square estimation was applied to fit the regression (2) and to obtain the residuals { $\widehat{\varepsilon_{k,i}}$ }. Let $\widehat{\Sigma_{e,i}} = n^{-1}\Sigma_{k=1}^{n} \widehat{\varepsilon_{k,i}} \widehat{\varepsilon'_{k,i}}$ be the residual sample covariance. The association between the *k*th trait and all of the evaluated SNPs were jointly evaluated by comparing the determinants of residual sample covariances $\widehat{\Sigma_{e.i}}$ with and without that trait included in the model. The p-value of the association was

obtained from a F-test built on a likelihood ratio statistic widely used in multi-response regression (see Section 7.7 in (Draper, N. R. and Smith, H., 1998; Johnson, R. A. et al., 2002). This F-test incorporated the dependence among the SNPs, and thus provided a more powerful test than the individual test of association for each single SNP by combining multiple signals across different SNPs.

If at least one trait passed a set significance threshold -- here $p < 0.01$ was selected -- the single most significant trait among all traits significantly associated at $p < 0.01$ was added to the model (2); see Figure 1. The model itself was then rerun using all traits selected to that point. If any of the traits already incorporated into the model failed to meet the original cut-off value of $p < 0.01$ after the incorporation of the newest trait, the single least significantly associated trait was removed from the model Figure 1. The cut-off value 0.01 is widely used for stepwise regression (Draper, N. R. and Smith, H., 1998). Looser thresholds (.05 and 0.1) required significantly more iterations to converge, dramatically increasing computational cost. The protocol employed here focused on controlling false discovery at the level of the final phenome-wide association test (as described below), while accepting that, regardless of the p-value threshold within the stepwise regression both false positives and false negatives for individual phenotypes will occur.

The process above constituted one iteration of the stepwise selection procedure. In the analyses presented in this manuscript, 35 sequential iterations of the stepwise selection procedure were performed per gene. With this dataset of genotypes and traits, every gene tested converged to a single stable model within less than 35 iterations (Supplemental Figure 6), however this assumption would need to be revisited when employing GPWAS on other datasets which might include either more individuals, more traits, or fewer traits with high pairwise correlation.

Phenotype Incorporating Model:

The final phenotype incorporating model (PM) can be represented as:

$$(3) g_{k,i} = \sum_{q=1}^{v_{pc}} PC_{k,q}\beta_{iq} + \sum_{j=1}^{v_i} Phe^{\iota}_{k,j}\tau^{\iota}_{ij} + \varepsilon_{k,i}$$

In the final model (3), there were v_i selected phenotypes for the *i*th gene, where

$v_i \leq 260$. The selected phenotypes { $Phe^{¿}_{k,j}$ } were a subset of the collection of all the phenotypes { $Phe_{k,1}, Phe_{k,2}, ..., Phe_{k,260}$ }, and $\tau^{¿}_{ij}$ was the corresponding coefficients for the selected phenotype $Phe^{¿}_{k,j}$ of the *i*th gene. Note that $g_{k,i}$, β_{iq}, and $\tau^{¿}_{ij}$ can be vectors corresponding to the multiple SNPs within the *i*th gene.

Model Comparison:

The final step was to evaluate how much the inclusion of trait data improved model fit (PM) relative to a purely population structure based model (RM). The statistical significance of the increase in goodness of fit of the two models was compared using a F-test (Johnson, R. A. and Wichern, D. W. et al., 2002) via comparing the residuals covariances of those models. The F-test takes into account all of the SNPs included from the target interval, as well as the degree of correlation between these SNPs. The filtering of highly linked SNPs described above satisfies the criteria of the F-tests that multiple response variables should not exhibit strong correlations with each other.

As adding more explanatory variables to a model will always tend to improve the goodness of fit, permutation based analyses were used to determine a threshold for a statistically significant increase in goodness of fit between the reduced model (RM) and the phenotype incorporating model (PM). Twenty permutations of the trait/genotype associations were conducted and GPWAS was run independently on each of these twenty permutations. Distributions of F-test PM/RM model comparisons from the permuted and unpermuted data were used to estimate false discovery rates at different cut off thresholds (Supplemental Figure 1).

GWAS Analysis

GLM GWAS and MLM GWAS analyses were conducted using the algorithm defined by Price and coworkers (Price et al., 2006). The FarmCPU GWAS with conducted using the algorithm defined by Liu and colleagues (Price et al., 2006; Liu et al., 2016). All algorithms were run using the R-based software rMVP (MVP version 1.0.1) (A Memory-efficient, Visualization-enhanced, and Parallel-accelerated Tool For Genome-Wide Association Study) (https://github.com/XiaoleiLiuBio/rMVP). FarmCPU analysis method was run using maxLoop = 10 and method.bin = "FaST-LMM" (Lippert et al., 2011). The first three principal components were considered to be additional covariates for the

population structure control in all analyses. The same kinship matrix used in the phenotype imputation was also used for controlling the genotype relationship in the MLM GWAS model, while the method for analyzing variance components (vc.method) was set to GEMMA (Zhou, 2017). To enable a comparison with the GPWAS results, each gene was assigned the p-value of the single most significant SNP among all the SNPs assigned to that gene across the 260 analyzed phenotypes in the GWAS model.

Comparison to Maize Classical Mutants

Maize classical loss of function mutant identities were taken from a previous study (Schnable and Freeling, 2011). To obtain an exhaustive list of reported mutant phenotypes, papers were mined from MaizeGDB loci pages, where both papers and conference abstracts that reference studies on individual maize genes, both cloned and uncloned, are captured by manual data curation (Schaeffer et al., 2011).

Nested Association Mapping Comparison in Maize

Published associations identified for 41 phenotypes scored across $\sim$ 5,000 maize recombinant inbred lines were retrieved from Panzea (http://cbsusrv04.tc.cornell.edu/users/panzea/download.aspx?filegroupid=14) (Schaeffer et al., 2011; Wallace et al., 2014). Following the thresholding proposed in that paper, a SNP and CNV (copy number variant) hits with a resample model inclusion probability $\geq$ 0.05, which were either within the longest annotated transcript for each gene (AGPv2.16) or within 15kb upstream or downstream of the annotated transcription start or stop sites were assigned to that gene respectively. Gene models were converted from the B73 RefGenV2 to B73 RefGenV4 using a conversion list published on MaizeGDB (https://www.maizegdb.org/search/gene/download_gene_xrefs.php?relative=v4).

Maize Gene Expression Analysis

Raw reads from a published maize expression atlas generated for the inbred line B73 were downloaded from the NCBI Sequence Read Archive PRJNA171684 (Stelpflug et al., 2016). Reads were trimmed using Trimmomatic-0.38 with default setting parameters (Bolger et al., 2014; Stelpflug et al., 2016). Trimmed reads were aligned to the maize

B73 RefGenV4 reference genome using GSNAP version 2018-03-25 (Wu and Nacu, 2010). Alignment results were converted to a sorted BAM file format using SAMtools 1.6 (Li et al., 2009; Wu and Nacu, 2010), and the FPKM values were calculated for each gene in the AGPv4.39 maize gene models in each sample using Cufflinks v2.2 (Trapnell et al., 2012). Only annotated genes located on 10 maize pseudomolecules were used for downstream analyses and the visualization of the FPKM distribution.

Calculating Ka/Ks ratios for Maize Gene Models

For each gene listed in a public syntenic gene list (Schnable J.C., 2018), the coding sequence for the single longest transcript per locus was downloaded from Ensembl Plants. Their sequences were each aligned to the single longest transcript of genes annotated as syntenic orthologs in *Sorghum bicolor* v3.1 (McCormick et al., 2018) and *Setaria italica* v2.2 (Bennetzen et al., 2012), retrieved from Phytozome v12.0 using a codon-based alignment as described previously (Zhang et al., 2017). The calculation of the ratio of the number of nonsynonymous substitutions per non-synonymous site (Ka) to the number of synonymous substitutions per synonymous site (Ks) was automatically calculated using scripts which are provided on github (https://github.com/shanwai1234/Grass-KaKs). Genes with a synonymous substitution rate less than 0.05 on the branch leading to maize after the maize/sorghum split were excluded from the analyses, as these atypically low Ks values tended to produce extreme Ka/Ks ratios. Genes with multiple tandem duplicates were also excluded from the Ka/Ks calculations. The calculated Ka/Ks ratios of maize genes are provided in Supplemental Data 1J.

Analysis of Presence/Absence Variation (PAV) Patterns in Maize

PAV data were downloaded from a published data file (Brohammer et al., 2018). Following the thresholding proposed in that paper, a gene was considered to exhibit presence absence variance if at least one inbred line had a coverage of less than 0.2.

Gene Ontology Enrichment Analysis in Maize

All GO analyses used the maize-GAMER GO annotations for B73 RefGenV4 gene models (Brohammer et al., 2018; Wimalanathan et al., 2018b). Statistical tests for GO

term enrichment and purification were performed using the goatools software package (v0.8.12) (Klopfenstein et al., 2018), with support for a two-sided Fisher's exact test provided by the fisher_exact function in SciPy. To determine the median information content of the GO term, each was assigned a score based on the total number of gene models to which this GO term was assigned to in the maize-GAMER dataset. This analysis considered only gene models to which a GO term was specifically applied to in the dataset, but not gene models where the assignment of the GO term may have been implied by the assignment of a child GO term. Genes in B73 RefGenV4 Zm00001d.2 that employed in maize-GAMER GO annotations (～ 40,000 genes) were used as the background population.

Evaluation of GPWAS and GWAS Power and FDR Using Simulated Data

SNP calls for the entire set of 1,210 individuals included in Maize HapMap3 were retrieved from Panzea (Bukowski et al., 2018), filtered, and assigned to genes as described above resulting in 1,648,398 SNPs assigned to annotated gene body regions in B73 RefGenV4. Two thousand genes, associated with 30,547 SNP markers were randomly sampled for downstream simulation. Independent phenotypes with known causal QTNs (Quantitative Trait Nucleotides) were simulated using the additive model in GCTA (v1.91.6) (Yang et al., 2011). Effect sizes for each QTN for each simulated phenotype in each permutation were drawn from a normal distribution centered on zero.

The resulting simulated trait data and genuine genotype calls were analyzed using GLM GWAS, FarmCPU GWAS, and GPWAS as described above, with the exception that the population structure PCs were calculated using a sample (1% or 191,856 SNPs) of the total SNPs remaining after filtering, rather than only using the subset of SNPs assigned to the 2,000 randomly selected genes included in this analysis.

For each analysis, the set of 2,000 genes was ranked from most to least statistically significant based on the significance of the most significantly associated SNP (for GLM and FarmCPU GWAS) or the significance of the overall model fit relative to a population structure only model (for GPWAS). The power evaluation for GPWAS was defined as the number of true positive genes relative to the total number of causal genes, and FDR was defined as the number of false positive genes relative to the total number of

positive genes. Power and FDR were calculated in a step of five genes starting with the five most significant genes and continuing to the 500 most significant genes (i.e. {5,10,...,495,500}).

In each phenotype simulation, 100 genes (5%) were randomly selected as causal genes to simulate one phenotype. For each causal gene in each simulation, a causal SNP was selected to simulate the phenotypic effect. Each of 100 phenotypic traits with heritability as 0.5 were simulated using the same set of 100 genes. Total 100 simulated phenotypes were split into 1, 5, 10, 20, 50 and 100 subgroups for running GPWAS.

Author Contributions

Z.L., Y.Q. and J.C.S conceived of the presented idea, Y.Q. designed the model, Z.L. designed the computational framework, Z.L. analyzed the data, Z.L. and J.C.S contributed to the implementation of the research, Z.L., Y.Q. and J.C.S wrote the manuscript. All authors discussed the results and contributed to the final manuscript.

Acknowledgements

This work is supported by National Science Foundation Awards MCB-1838307 and OIA-1826781 to JCS. In addition, we received support from the Quantitative Life Sciences Initiative at the University of Nebraska-Lincoln, which in turn received support from the University of Nebraska Program of Excellence. This work was completed utilizing the Holland Computing Center of the University of Nebraska, which receives support from the Nebraska Research Initiative. The authors thank Andy Dahl for the advice and instruction in the use of phenotype imputation software, Zheng Xu and Wenlong Ren for their consultation on the design of the association study, Sunil K. Kenchanmane Raju for a critical reading of the manuscript and valuable feedback, and the Panzea project (http://www.panzea.org) for gathering the phenotypic and genotypic data employed in this study. All the authors declare that they have no competing interests.

Conflict of Interest

The authors do not have any conflict of interest to declare.

Figure Legends

Figure 1. GPWAS algorithm implementation. (A) Example of trait and genotype matrices employed for GPWAS. (B) Flow chart showing initial data processing and the forward selection process within the GPWAS algorithm.

Figure 2. **Statistical association between the maize gene Zm00001d002175 and 260 distinct phenotypes.** Each diamond or triangle represents one specific phenotypic dataset. Symbol colors indicate the broad categories into which each specific phenotype falls. The specific identities of each phenotype ordered from left to right are given in Supplemental Data 1A. (A) The position of each diamond on the y-axis indicates the negative $\log_{10}$ p-value of the most statistically significant SNP assigned to that gene in a GLM GWAS analysis for that single trait. The dashed blue line indicates a p = 0.05 cutoff after Bonferroni correction for multiple testing based on the number of statistical tests in a single GWAS analysis ($8.96e^{-8}$). The solid line indicates a p = 0.05 cutoff after Bonferroni correction for multiple testing based on the number of statistical tests in GWAS for all 260 traits ($3.45e^{-10}$). (B) The placement of each triangle on the y-axis indicates whether a given phenotype was included in (Sel.) or excluded from (Uns.) the final GPWAS model constructed for this gene. The complete list of phenotypes incorporated into the GPWAS model for Zm00001d002175 is as follows: days to silk (Summer 2006, Cayuga, NY; Summer 2007, Johnston, NC), days to tassel (Summer 2007, Johnston, NC; Summer 2008, Cayuga, NY), GDD (Growing Degree Days) day to silk (Summer 2006, Cayuga, NY; Summer 2007, Johnston, NC), main spike length (Summer 2006, Johnston, NC), number of leaves (Summer 2008, Cayuga, NY), leaf width (Summer 2006, Champaign, IL), NIR (Near InfraRed)-measured protein (Summer 2006, Johnston, NC) and ear weight (Summer 2006, Champaign, IL). (C) The panel indicates the pairwise Pearson correlation coefficient between each pair of measured phenotypes. Clustering based on phenotypic correlation was used to determine the ordering of phenotypes along the x-axis. Each tick mark on the x-axes of the top and middle panels indicates a distance of five phenotype datasets.

Figure 3. Comparison of power to detect causal genes using either GWAS or GPWAS and either real and simulated data. (A) Proportion of genes linked to phenotypic variation in the Buckler-Goodman association panel using different statistical

methods which were also identified as linked to trait variation in a separate analysis of the maize Nested Association Mapping population (Wallace et al). The leftmost bar indicates the background rate among all maize gene models with 1:1 relationships between B73 RefGenV2 (used by Wallace et al) and B73 RefGenV4 (used in this study). *: $p \leq 0.05$; ***: $p \leq 1e^{-3}$ (Chi-Square Test). (B) Relationship between the total number of positive genes selected by each of the four quantitative genetics methods and the total number of positive genes which were also identified by Wallace et al. (C) Comparison of the performance of GPWAS and conventional GWAS methods in the identification of *a prior* candidate genes involved in vitamin A and E biosynthesis. Phenotypic data and published *a priori* candidate gene lists for vitamin A and vitamin E were taken from (Owens et al., 2014); (Diepenbrock et al., 2017). The methodology used here was otherwise identical to that employed for panel B. (D) Trade offs between power and FDR for GPWAS and FarmCPU GWAS when working with simulated data and different numbers of phenotypes. Each point indicates power and FDR calculated at a different step between the 5 and 500 most significantly associated genes, and derived from the 10 independent simulations. Error bars indicate standard errors for both FDR and power. The suffix -p# within the legend indicates the number of distinct phenotypes included in the analysis. FarmCPU provided the most favorable trade offs between power and FDR than the other two GWAS models at all steps. The results of the two other methods are omitted for readability. For simulation details see methods.

Figure 4. Evaluation of GLM GWAS, MLM GWAS, FarmCPU GWAS, single iteration single trait gene-trait association and GPWAS using known maize genes *Anther ear1* (*an1*) (Zm00001d032961) and *liguleless2* (*lg2*) (Zm00001d042777). (A) The dashed lines indicates a p-value corresponding to 0.05 after a Bonferroni correction for independent tests on 557,968 (SNPs). Solid lines indicate the stricter multiple testing corrected threshold, which considers both the number of SNPs and the number of phenotypes tested. In the GPWAS panel, Sel. and Uns. indicate traits that were selected and unselected respectively, in the model GPWAS fit for this particular gene. Phenotypes are ordered along the x-axis in the same order used for Figure 2, with each tick mark indicating a distance of five phenotypes. Phenotypes incorporated in the GPWAS model for *an1* were as follows: germination count (Summer 2006, Johnston, NC), days to tassel (Summer 2007, Cayuga, NY), GDD days to silk (Summer 2007, Johnston, NC; Summer 2007, Champaign, IL; Winter 2006, Miami-Dade, FL), tassel

length (Summer 2007, Cayuga, NY), spikelets primary branch (Summer 2006, Champaign, IL), secondary branch number (Summer 2006, Boone, MO), plant height (Summer 2006, Cayuga, NY), NIR-measured protein (Summer 2006, Johnston, NC), NIR-measured oil (Summer 2006, Johnston, NC; Winter 2006, Miami-Dade, FL), cob weight (Summer 2007, Johnston, NC), ear diameter (Summer 2007, Johnston, NC) and total kernel volume (Summer 2006, Cayuga, NY). (B) Phenotypes incorporated in the GPWAS model for *lg2* were as follows: days to silk (Summer 2006, Johnston, NC), days to tassel (Winter 2006, Ponce, PR), GDD days to tassel (Summer 2007, Champaign, IL), GDD anthesis-silking interval (Winter 2007, Miami-Dade, FL), main spike length (Summer 2006, Johnston, NC), leaf length (Summer 2006, Boone, MO), upper leaf angle (Summer 2006, Cayuga, NY), number of tillering plants (Summer 2007, Cayuga, NY), cob diameter (Winter 2006, Ponce, PR) and kernel weight (Summer 2007, Cayuga, NY). (C) The potential correspondence between phenotypes selected using the GPWAS model for *an1* using the GPWAS model and phenotypes either reported for loss of function *an1* mutants or previous quantitative genetic analyses (Bensen et al., 1995; Peng and Harberd, 2002; Landoni et al., 2007; Brown et al., 2011). (D) The potential correspondence between phenotypes selected by the GPWAS model for *lg2*, and phenotypes either reported for loss of function *lg2* mutants or previous quantitative genetic analyses (Brink, 1933; Pendleton et al., 1968; Lambert and Johnson, 1978; Harper and Freeling, 1996; Walsh and Freeling, 1999; Tian et al., 2011).

Figure 5. Comparisons among four gene populations: Background, GWAS (Genes linked to phenotype by GWAS (GLM) but not by GPWAS), GPWAS (Genes linked to phenotype by GPWAS but not by GWAS (GLM)) and classical mutants with known loss of function phenotypes (Schnable and Freeling, 2011). (A) Proportion of genes within each of the four populations which express with average FPKM > 1 of 92 assayed tissues/time points. (B) Proportion of genes within each of the four populations which exhibit presence absence variation (PAV) in maize. (C) Proportion of genes within each of the four populations which are conserved at syntenic orthologous locations in sorghum (*Sorghum bicolor*). (D) Distribution of Non-synonymous substitution rate/Synonymous substitution rate (Ka/Ks) for the subset of genes from each of the four populations with syntenic orthologs in both sorghum and foxtail millet (*Setaria italica*). A small number of individual genes with Ka/Ks ratios >1 are not show to aid readability (one gene in the GLM GWAS population, two in the GPWAS population). *:p value <

0.05; **: p value < 0.01; ***: p value < $1e^{-3}$ (Chi-Squared test employed for panels A-C; Mann–Whitney U employed for panel D.). GWAS, GPWAS, and classical mutant populations exhibited statistically significant differences from the background gene population for all four features shown. Comparable data for the other two GWAS algorithms evaluated as part of this study (MLM and FarmCPU), are provided in Supplemental Data.

Figure S1. Permutation testing based estimation of false discovery rates for GLM GWAS, FarmCPU, and GPWAS. For each panel, the dark curve shows the distribution of per gene p-values obtained from 20 permutations of genotype and trait data (see Methods), while the light curve indicates the distribution of per gene p-values obtained from the analysis of the non-permuted dataset. Red lines indicate the p-value analyses employed in these analysis, corresponding top p-value = $8.96e^{-8}$ for GLM and FarmCPU and an estimated FDR < 0.001 for GPWAS. Genes assigned p-values on the right side of each red line were employed for all downstream analyses in the main text. Panels A-C show the entirety of the distributions, while panels D-F display a magnified view of the regions of the curve where the p-value threshold is employed. When these data were used to estimate the p-value cut off corresponding to an estimated FDR < 0.001 for GLM GWAS, this was found to correspond to an uncorrected p-value of approximately 1e-14, resulting in 31 genes would remain statistically significantly associated with traits. For FarmCPU GWAS, the minimum FDR achieved was FDR <0.029 at a p-value threshold of $1e^{-15}$, resulting in 38 genes remaining statistically significantly associated with traits.

Figure S2. P-values assigned to the link between *an1* and individual phenotypes when the multi-SNP GPWAS algorithm was run with a single iteration and provided with only a single trait per analysis. Essentially this eliminates the effects of the forward selection and the "multi-trait" components of the GPWAS algorithm, retaining only the "multi-marker" components. Traits are sorted from most significant single-trait GPWAS association to least. Stars indicate those phenotypes which were associated with *an1* when all 260 phenotypes were analyzed together using the full GPWAS model.

Figure S3. Comparison of GO enrichment/purification among genes uniquely identified as being associated with phenotypic variation using different statistical

approaches. Each circle represents a single GO term in a single analysis. The position of each circle on the x axis indicates the total number of maize gene models which were assigned to this GO term in the maize GAMER dataset (Wimalanathan et al., 2018a). The position of each circle on the y-axis indicates the statistical significance of the enrichment or purification of this GO term in the given gene population relative to the background set of all annotated maize gene models. Red lines indicate the threshold for determining a significant GO term after a Bonferroni correction. (A) Comparison of the patterns of GO term enrichment/purification among genes either uniquely identified as being associated with phenotypic variation using a GLM GWAS analysis or uniquely identified as being associated with phenotypic variation in a GPWAS analysis. (B) As in panel A, but the comparison is between genes uniquely identified as being associated with phenotypic variation using a FarmCPU analysis or uniquely identified as being associated with phenotypic variation in a GPWAS analysis. Only the 706 genes uniquely identified using GPWAS with the strongest statistical signal were employed in panel b, to prevent any bias towards more significant p-values resulting from an analysis using a larger population of genes identified using GPWAS than those identified using FarmCPU.

Figure S4. Number of SNPs identified per gene and the p-value of genes identified using different models. (A) The number of SNPs assigned to genes uniquely identified using either GPWAS or GLM GWAS, as well as the total number of genes with identified SNPs. SNPs assigned to gene regions were filtered and employed in all analyses. The maximum remaining number of SNPs per gene was 138. The distributions of the genes uniquely identified using GLM GWAS or GPWAS were statistically significantly different, $p < 2.2e^{-16}$ (Mann-Whitney U test; two-sided). (B) Correlations between the SNP number per gene and the $-\log_{10}$ p-value of the total number of genes identified using GPWAS on real phenotype data. (C) Correlations between the SNP number per gene and the $-\log_{10}$ p-value of the total genes identified using GPWAS on randomly selected phenotype data from 20 permutations. (D) Correlations between the SNP number per gene and the $-\log_{10}$ p-value of the total genes identified using GLM GWAS on real phenotype data. (E) Correlations between the SNP number per gene and $-\log_{10}$ p-value of total genes identified using GLM GWAS on randomly selected phenotype data from 20 permutations. Spearman correlation methods were employed for the correlation test between SNP number and $-\log_{10}$ transformed p-value for each gene. Full statistical

reports are presented in Supplemental Data 1H.

Figure S5. Variance explained by principal components and genes detected by association models with top 20 PCs as covariates. (A) Cumulative variance explained by principal components in maize 282 association panel. (B) Proportion of genes within each of the four populations which express with average FPKM > 1 of 92 assayed tissues/time points. (C) Proportion of genes within each of the four populations which exhibit presence absence variation (PAV) in maize. (D) Proportion of genes within each of the four populations which are conserved at syntenic orthologous locations in sorghum (*Sorghum bicolor*). (E) Distribution of Non-synonymous substitution rate/Synonymous substitution rate (Ka/Ks) for the subset of genes from each of the four populations with syntenic orthologs in both sorghum and foxtail millet (*Setaria italica*). A small number of individual genes with Ka/Ks ratios > 1 are not show to aid readability (two genes in the GLM GWAS population). *:p value < 0.05; **: p value < 0.01 (Chi-Squared test employed for panels B-D; Mann–Whitney U employed for panel E.). GLM model represents GWAS model in this figure. GWAS, GPWAS, and classical mutant populations exhibited statistically significant differences from the background gene population for all four features shown, except for the insignificant difference between GWAS and background in Ka/Ks ratio.

Figure S6. Estimating the number of iterations required for convergence of the forward selection model for GPWAS using the set of phenotypes and genotypic data for the maize 282 buckler goodman association panel employed here. (A) Distribution of the number of phenotypes incorporated into GPWAS models for the 1,776 genes identified as significantly linked to maize phenomic variation in this study; (B) Change in the number of phenotypes incorporated into genetic models as the number of iterations employed for GPWAS increases. Data shown for the 10 genes with the largest total number of phenotypes incorporated into their models among the 1,776 genes identified in this study.

Supplemental Information

Supplemental Data 1.

References

Araus, J. L., Kefauver, S. C., Zaman-Allah, M., Olsen, M. S., and Cairns, J. E. (2018). Translating High-Throughput Phenotyping into Genetic Gain. *Trends Plant Sci.* **23**:451–466.

Atwell, S., Huang, Y. S., Vilhjálmsson, B. J., Willems, G., Horton, M., Li, Y., Meng, D., Platt, A., Tarone, A. M., Hu, T. T., et al. (2010). Genome-wide association study of 107 phenotypes in Arabidopsis thaliana inbred lines. *Nature* **465**:627–631.

Bennetzen, J. L., Schmutz, J., Wang, H., Percifield, R., Hawkins, J., Pontaroli, A. C., Estep, M., Feng, L., Vaughn, J. N., Grimwood, J., et al. (2012). Reference genome sequence of the model plant Setaria. *Nat. Biotechnol.* **30**:555–561.

Bensen, R. J., Johal, G. S., Crane, V. C., Tossberg, J. T., Schnable, P. S., Meeley, R. B., and Briggs, S. P. (1995). Cloning and characterization of the maize An1 gene. *Plant Cell* **7**:75–84.

Bolger, A. M., Lohse, M., and Usadel, B. (2014). Trimmomatic: a flexible trimmer for Illumina sequence data. *Bioinformatics* **30**:2114–2120.

Brink, R. A. (1933). HERITABLE CHARACTERS IN MAIZE: XLVI—Ligu1eless-2. *J. Hered.* **24**:325–326.

Brink, R. A. (1935). HERITABLE CHARACTERS IN MAIZE. *Journal of Heredity* **26**:249–251.

Brohammer, A. B., Kono, T. J. Y., Springer, N. M., McGaugh, S. E., and Hirsch, C. N. (2018). The limited role of differential fractionation in genome content variation and function in maize (Zea mays L.) inbred lines. *Plant J.* **93**:131–141.

Brown, P. J., Upadyayula, N., Mahone, G. S., Tian, F., Bradbury, P. J., Myles, S., Holland, J. B., Flint-Garcia, S., McMullen, M. D., Buckler, E. S., et al. (2011). Distinct genetic architectures for male and female inflorescence traits of maize. *PLoS Genet.* **7**:e1002383.

Browning, B. L., and Browning, S. R. (2016). Genotype Imputation with Millions of Reference Samples. *Am. J. Hum. Genet.* **98**:116–126.

Bukowski, R., Guo, X., Lu, Y., Zou, C., He, B., Rong, Z., Wang, B., Xu, D., Yang, B., Xie, C., et al. (2018). Construction of the third-generation Zea mays haplotype map. *Gigascience* **7**:1–12.

Castelletti, S., Tuberosa, R., Pindo, M., and Salvi, S. (2014). A MITE transposon insertion is associated with differential methylation at the maize flowering time QTL Vgt1. *G3* **4**:805–812.

Chong, J. X., Buckingham, K. J., Jhangiani, S. N., Boehm, C., Sobreira, N., Smith, J. D., Harrell, T. M., McMillin, M. J., Wiszniewski, W., Gambin, T., et al. (2015). The Genetic Basis of Mendelian Phenotypes: Discoveries, Challenges, and Opportunities. *Am. J. Hum. Genet.* **97**:199–215.

Chuck, G., Meeley, R. B., and Hake, S. (1998). The control of maize spikelet meristem

fate by the APETALA2-like gene indeterminate spikelet1. *Genes Dev.* **12**:1145–1154.

Dahl, A., Iotchkova, V., Baud, A., Johansson, Å., Gyllensten, U., Soranzo, N., Mott, R., Kranis, A., and Marchini, J. (2016). A multiple-phenotype imputation method for genetic studies. *Nat. Genet.* **48**:466–472.

Dewan, A., Liu, M., Hartman, S., Zhang, S. S.-M., Liu, D. T. L., Zhao, C., Tam, P. O. S., Chan, W. M., Lam, D. S. C., Snyder, M., et al. (2006). HTRA1 promoter polymorphism in wet age-related macular degeneration. *Science* **314**:989–992.

Diepenbrock, C. H., Kandianis, C. B., Lipka, A. E., Magallanes-Lundback, M., Vaillancourt, B., Góngora-Castillo, E., Wallace, J. G., Cepela, J., Mesberg, A., Bradbury, P. J., et al. (2017). Novel Loci Underlie Natural Variation in Vitamin E Levels in Maize Grain. *Plant Cell* **29**:2374–2392.

Draper, N. R. and Smith, H. (1998). *Applied regression analysis*. John Wiley & Sons.

Eagle, R. C. (2006). Complement Factor H Polymorphism in Age-Related Macular Degeneration. *Yearbook of Ophthalmology* **2006**:245–248.

Exposito-Alonso, M., 500 Genomes Field Experiment Team, Burbano, H. A., Bossdorf, O., Nielsen, R., and Weigel, D. (2019). Natural selection on the Arabidopsis thaliana genome in present and future climates. *Nature* **573**:126–129.

Flint-Garcia, S. A., Thuillet, A.-C., Yu, J., Pressoir, G., Romero, S. M., Mitchell, S. E., Doebley, J., Kresovich, S., Goodman, M. M., and Buckler, E. S. (2005). Maize association population: a high-resolution platform for quantitative trait locus dissection. *Plant J.* **44**:1054–1064.

Harper, L., and Freeling, M. (1996). Interactions of liguleless1 and liguleless2 function during ligule induction in maize. *Genetics* **144**:1871–1882.

Huang, X., Wei, X., Sang, T., Zhao, Q., Feng, Q., Zhao, Y., Li, C., Zhu, C., Lu, T., Zhang, Z., et al. (2010). Genome-wide association studies of 14 agronomic traits in rice landraces. *Nat. Genet.* **42**:961–967.

International HapMap Consortium and others (2005). A haplotype map of the human genome. *Nature* **437**:1299.

Johnson, R. A. and Wichern, D. W. et al. (2002). *Applied multivariate statistical analysis*. Prentice hall Upper Saddle River, NJ.

Julkowska, M. M., Koevoets, I. T., Mol, S., Hoefsloot, H., Feron, R., Tester, M. A., Keurentjes, J. J. B., Korte, A., Haring, M. A., de Boer, G.-J., et al. (2017). Genetic Components of Root Architecture Remodeling in Response to Salt Stress. *Plant Cell* **29**:3198–3213.

Klopfenstein, D. V., Zhang, L., Pedersen, B. S., Ramírez, F., Warwick Vesztrocy, A., Naldi, A., Mungall, C. J., Yunes, J. M., Botvinnik, O., Weigel, M., et al. (2018). GOATOOLS: A Python library for Gene Ontology analyses. *Sci. Rep.* **8**:10872.

Korte, A., Vilhjálmsson, B. J., Segura, V., Platt, A., Long, Q., and Nordborg, M. (2012). A mixed-model approach for genome-wide association studies of correlated traits in structured populations. *Nat. Genet.* **44**:1066–1071.

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., and Buckler, E. S. (2018). Dysregulation of expression correlates with rare-allele burden and fitness loss in maize. *Nature* **555**:520–523.

Kremling, K. A. G., Diepenbrock, C. H., Gore, M. A., Buckler, E. S., and Bandillo, N. B. (2019). Transcriptome-Wide Association Supplements Genome-Wide Association in *Zea mays*. *G3: Genes|Genomes|Genetics* **9**:3023–3033.

Lambert, R. J., and Johnson, R. R. (1978). Leaf Angle, Tassel Morphology, and the Performance of Maize Hybrids 1. *Crop Science* **18**:499–502.

Lamesch, P., Berardini, T. Z., Li, D., Swarbreck, D., Wilks, C., Sasidharan, R., Muller, R., Dreher, K., Alexander, D. L., Garcia-Hernandez, M., et al. (2012). The Arabidopsis Information Resource (TAIR): improved gene annotation and new tools. *Nucleic Acids Res.* **40**:D1202–10.

Landoni, M., Vecchia, F. D., Gavazzi, G., Giulini, A., La Rocca, N., Rascio, N., Colombo, M., Bononi, M., and Consonni, G. (2007). The an1-4736 mutation of anther ear1 in maize alters scotomorphogenesis and the light response. *Plant Science* **172**:172–180.

Li, J., and Ji, L. (2005). Adjusting multiple testing in multilocus analyses using the eigenvalues of a correlation matrix. *Heredity* **95**:221–227.

Li, H., Handsaker, B., Wysoker, A., Fennell, T., Ruan, J., Homer, N., Marth, G., Abecasis, G., Durbin, R., and 1000 Genome Project Data Processing Subgroup (2009). The Sequence Alignment/Map format and SAMtools. *Bioinformatics* **25**:2078–2079.

Li, Y., Huang, Y., Bergelson, J., Nordborg, M., and Borevitz, J. O. (2010). Association mapping of local climate-sensitive quantitative trait loci in Arabidopsis thaliana. *Proc. Natl. Acad. Sci. U. S. A.* **107**:21199–21204.

Li, M., Zhong, W., Yang, F., and Zhang, Z. (2018). Genetic and Molecular Mechanisms of Quantitative Trait Loci Controlling Maize Inflorescence Architecture. *Plant Cell Physiol.* **59**:448–457.

Lin, H.-Y., Liu, Q., Li, X., Yang, J., Liu, S., Huang, Y., Scanlon, M. J., Nettleton, D., and Schnable, P. S. (2017). Substantial contribution of genetic variation in the expression of transcription factors to phenotypic variation revealed by eRD-GWAS. *Genome Biol.* **18**:192.

Lippert, C., Listgarten, J., Liu, Y., Kadie, C. M., Davidson, R. I., and Heckerman, D. (2011). FaST linear mixed models for genome-wide association studies. *Nature Methods* **8**:833–835.

Listgarten, J., Lippert, C., Kadie, C. M., Davidson, R. I., Eskin, E., and Heckerman, D. (2012). Improved linear mixed models for genome-wide association studies. *Nature Methods* **9**:525–526.

Liu, L., Du, Y., Shen, X., Li, M., Sun, W., Huang, J., Liu, Z., Tao, Y., Zheng, Y., Yan, J., et al. (2015). KRN4 Controls Quantitative Variation in Maize Kernel Row Number. *PLoS Genet.* **11**:e1005670.

Liu, X., Huang, M., Fan, B., Buckler, E. S., and Zhang, Z. (2016). Iterative Usage of Fixed and Random Effect Models for Powerful and Efficient Genome-Wide Association Studies. *PLoS Genet.* **12**:e1005767.

Lloyd, J. P., Tsai, Z. T.-Y., Sowers, R. P., Panchy, N. L., and Shiu, S.-H. (2018). A Model-Based Approach for Identifying Functional Intergenic Transcribed Regions and Noncoding RNAs. *Mol. Biol. Evol.* **35**:1422–1436.

Lock, A., Rutherford, K., Harris, M. A., Hayles, J., Oliver, S. G., Bähler, J., and Wood, V. (2019). PomBase 2018: user-driven reimplementation of the fission yeast database provides rapid and intuitive access to diverse, interconnected information. *Nucleic Acids Res.* **47**:D821–D827.

Matsuda, F., Nakabayashi, R., Yang, Z., Okazaki, Y., Yonemaru, J.-I., Ebana, K., Yano, M., and Saito, K. (2015). Metabolome-genome-wide association study dissects genetic architecture for generating natural variation in rice secondary metabolism. *Plant J.* **81**:13–23.

McCormick, R. F., Truong, S. K., Sreedasyam, A., Jenkins, J., Shu, S., Sims, D., Kennedy, M., Amirebrahimi, M., Weers, B. D., McKinley, B., et al. (2018). The Sorghum bicolor reference genome: improved assembly, gene annotations, a transcriptome atlas, and signatures of genome organization. *Plant J.* **93**:338–354.

McMullen, M. D., Kresovich, S., Villeda, H. S., Bradbury, P., Li, H., Sun, Q., Flint-Garcia, S., Thornsberry, J., Acharya, C., Bottoms, C., et al. (2009). Genetic properties of the maize nested association mapping population. *Science* **325**:737–740.

Morris, G. P., Ramu, P., Deshpande, S. P., Hash, C. T., Shah, T., Upadhyaya, H. D., Riera-Lizarazu, O., Brown, P. J., Acharya, C. B., Mitchell, S. E., et al. (2013). Population genomic and genome-wide association studies of agroclimatic traits in sorghum. *Proc. Natl. Acad. Sci. U. S. A.* **110**:453–458.

Namjou, B., Marsolo, K., Carroll, R., Denny, J., Ritchie, M. D., Lingren, T., Porollo, A., Perry, C., Kottyan, L. C., Holm, I. A., et al. (2014). Phenome-wide association study (PheWAS) in EMR-linked pediatric cohorts. *Front. Genet.* **5**:401.

Oka, R., Zicola, J., Weber, B., Anderson, S. N., Hodgman, C., Gent, J. I., Wesselink, J.-J., Springer, N. M., Hoefsloot, H. C. J., Turck, F., et al. (2017). Genome-wide mapping of transcriptional enhancer candidates using DNA and chromatin features in maize. *Genome Biol.* **18**:137.

O'Reilly, P. F., Hoggart, C. J., Pomyen, Y., Calboli, F. C. F., Elliott, P., Jarvelin, M.-R.,

and Coin, L. J. M. (2012). MultiPhen: joint model of multiple phenotypes can increase discovery in GWAS. *PLoS One* **7**:e34861.

Owens, B. F., Lipka, A. E., Magallanes-Lundback, M., Tiede, T., Diepenbrock, C. H., Kandianis, C. B., Kim, E., Cepela, J., Mateos-Hernandez, M., Buell, C. R., et al. (2014). A foundation for provitamin A biofortification of maize: genome-wide association and genomic prediction models of carotenoid levels. *Genetics* **198**:1699–1716.

Paz-Ares, J., Ghosal, D., Wienand, U., Peterson, P. A., and Saedler, H. (1987). The regulatory c1 locus of Zea mays encodes a protein with homology to myb proto-oncogene products and with structural similarities to transcriptional activators. *EMBO J.* **6**:3553–3558.

Pendleton, J. W., Smith, G. E., Winter, S. R., and Johnston, T. J. (1968). Field Investigations of the Relationships of Leaf Angle in Corn (Zea mays L.) to Grain Yield and Apparent Photosynthesis1. *Agronomy Journal* **60**:422–424.

Peng, J., and Harberd, N. P. (2002). The role of GA-mediated signalling in the control of seed germination. *Curr. Opin. Plant Biol.* **5**:376–381.

Pitchers, W., Nye, J., Márquez, E. J., Kowalski, A., Dworkin, I., and Houle, D. (2019). A Multivariate Genome-Wide Association Study of Wing Shape inDrosophila melanogaster. *Genetics* **211**:1429–1447.

Price, A. L., Patterson, N. J., Plenge, R. M., Weinblatt, M. E., Shadick, N. A., and Reich, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.* **38**:904–909.

Remington, D. L., Thornsberry, J. M., Matsuoka, Y., Wilson, L. M., Whitt, S. R., Doebley, J., Kresovich, S., Goodman, M. M., and Buckler, E. S., 4th (2001). Structure of linkage disequilibrium and phenotypic associations in the maize genome. *Proc. Natl. Acad. Sci. U. S. A.* **98**:11479–11484.

Rencher, A. C., and Schaalje, G. B. (2008). *Linear models in statistics*. John Wiley & Sons.

Rhee, S. Y., and Mutwil, M. (2014). Towards revealing the functions of all genes in plants. *Trends Plant Sci.* **19**:212–221.

Rincent, R., Moreau, L., Monod, H., Kuhn, E., Melchinger, A. E., Malvar, R. A., Moreno-Gonzalez, J., Nicolas, S., Madur, D., Combes, V., et al. (2014). Recovering power in association mapping panels with variable levels of linkage disequilibrium. *Genetics* **197**:375–387.

Rodgers-Melnick, E., Vera, D. L., Bass, H. W., and Buckler, E. S. (2016). Open chromatin reveals the functional maize genome. *Proc. Natl. Acad. Sci. U. S. A.* **113**:E3177–84.

Romay, M. C., Millard, M. J., Glaubitz, J. C., Peiffer, J. A., Swarts, K. L., Casstevens, T. M., Elshire, R. J., Acharya, C. B., Mitchell, S. E., Flint-Garcia, S.

A., et al. (2013). Comprehensive genotyping of the USA national maize inbred seed bank. *Genome Biol.* **14**:R55.

Sax, K. (1923). The Association of Size Differences with Seed-Coat Pattern and Pigmentation in PHASEOLUS VULGARIS. *Genetics* **8**:552–560.

Schaeffer, M. L., Harper, L. C., Gardiner, J. M., Andorf, C. M., Campbell, D. A., Cannon, E. K. S., Sen, T. Z., and Lawrence, C. J. (2011). MaizeGDB: curation and outreach go hand-in-hand. *Database* **2011**:bar022.

Schnable, J. C. (2019). Genes and gene models, an important distinction. *New Phytol.* Advance Access published June 26, 2019, doi:10.1111/nph.16011.

Schnable, J. C., and Freeling, M. (2011). Genes identified by visible mutant phenotypes show increased bias toward one of two subgenomes of maize. *PLoS One* **6**:e17855.

Schnable J.C. (2018). Sorghum version 3, maize versions 3 and 4 syntenic gene list. *FigShare* Advance Access published 2018.

Schofield, P. N., Hoehndorf, R., and Gkoutos, G. V. (2012). Mouse genetic and phenotypic resources for human genetics. *Human Mutation* **33**:826–836.

Sprague, G. F. (1941). The location of dominant favorable genes in maize by means of an inversion. *Genetics* **26**:143–149.

Stelpflug, S. C., Sekhon, R. S., Vaillancourt, B., Hirsch, C. N., Buell, C. R., de Leon, N., and Kaeppler, S. M. (2016). An Expanded Maize Gene Expression Atlas based on RNA Sequencing and its Use to Explore Root Development. *Plant Genome* **9**.

Stephens, M. (2013). A unified framework for association analysis with multiple related phenotypes. *PLoS One* **8**:e65245.

Studer, A., Zhao, Q., Ross-Ibarra, J., and Doebley, J. (2011). Identification of a functional transposon insertion in the maize domestication gene tb1. *Nat. Genet.* **43**:1160–1163.

Sturaro, M., Hartings, H., Schmelzer, E., Velasco, R., Salamini, F., and Motto, M. (2005). Cloning and characterization of GLOSSY1, a maize gene involved in cuticle membrane and wax production. *Plant Physiol.* **138**:478–489.

Tian, F., Bradbury, P. J., Brown, P. J., Hung, H., Sun, Q., Flint-Garcia, S., Rocheford, T. R., McMullen, M. D., Holland, J. B., and Buckler, E. S. (2011). Genome-wide association study of leaf architecture in the maize nested association mapping population. *Nat. Genet.* **43**:159–162.

Trapnell, C., Roberts, A., Goff, L., Pertea, G., Kim, D., Kelley, D. R., Pimentel, H., Salzberg, S. L., Rinn, J. L., and Pachter, L. (2012). Differential gene and transcript expression analysis of RNA-seq experiments with TopHat and Cufflinks. *Nat. Protoc.* **7**:562–578.

Trégouët, D.-A., König, I. R., Erdmann, J., Munteanu, A., Braund, P. S., Hall, A. S.,

Grosshennig, A., Linsel-Nitschke, P., Perret, C., DeSuremain, M., et al. (2009). Genome-wide haplotype association study identifies the SLC22A3-LPAL2-LPA gene cluster as a risk locus for coronary artery disease. *Nat. Genet.* **41**:283–285.

Turco, G., Schnable, J. C., Pedersen, B., and Freeling, M. (2013). Automated conserved non-coding sequence (CNS) discovery reveals differences in gene content and promoter evolution among grasses. *Front. Plant Sci.* **4**:170.

Turley, P., Walters, R. K., Maghzian, O., Okbay, A., Lee, J. J., Fontana, M. A., Nguyen-Viet, T. A., Wedow, R., Zacher, M., Furlotte, N. A., et al. (2018). Multi-trait analysis of genome-wide association summary statistics using MTAG. *Nat. Genet.* **50**:229–237.

van der Sluis, S., Posthuma, D., and Dolan, C. V. (2013). TATES: efficient multivariate genotype-phenotype analysis for genome-wide association studies. *PLoS Genet.* **9**:e1003235.

Wallace, J. G., Bradbury, P. J., Zhang, N., Gibon, Y., Stitt, M., and Buckler, E. S. (2014). Association mapping across numerous traits reveals patterns of functional variation in maize. *PLoS Genet.* **10**:e1004845.

Walsh, J., and Freeling, M. (1999). The liguleless2 gene of maize functions during the transition from the vegetative to the reproductive shoot apex. *Plant J.* **19**:489–495.

Walsh, J., Waters, C. A., and Freeling, M. (1998). The maize gene liguleless2 encodes a basic leucine zipper protein involved in the establishment of the leaf blade-sheath boundary. *Genes Dev.* **12**:208–218.

Walter, A., Liebisch, F., and Hund, A. (2015). Plant phenotyping: from bean weighing to image analysis. *Plant Methods* **11**:14.

Wang, Y., Liu, A., Mills, J. L., Boehnke, M., Wilson, A. F., Bailey-Wilson, J. E., Xiong, M., Wu, C. O., and Fan, R. (2015). Pleiotropy analysis of quantitative traits at gene level by multivariate functional linear models. *Genet. Epidemiol.* **39**:259–275.

Wen, W., Li, D., Li, X., Gao, Y., Li, W., Li, H., Liu, J., Liu, H., Chen, W., Luo, J., et al. (2014). Metabolome-based genome-wide association study of maize kernel leads to novel biochemical insights. *Nat. Commun.* **5**:3438.

Wimalanathan, K., Friedberg, I., Andorf, C. M., and Lawrence-Dill, C. J. (2018a). Maize GO Annotation—Methods, Evaluation, and Review (maize-GAMER). *Plant Direct* **2**:e00052.

Wimalanathan, K., Friedberg, I., Andorf, C. M., and Lawrence-Dill, C. J. (2018b). Maize GO Annotation-Methods, Evaluation, and Review (maize-GAMER). *Plant Direct* **2**:e00052.

Wu, T. D., and Nacu, S. (2010). Fast and SNP-tolerant detection of complex variants and splicing in short reads. *Bioinformatics* **26**:873–881.

Xu, X., Dietrich, C. R., Delledonne, M., Xia, Y., Wen, T. J., Robertson, D. S., Nikolau, B. J., and Schnable, P. S. (1997). Sequence Analysis of the Cloned glossy8 Gene of Maize Suggests That It May Code for a [beta]-Ketoacyl Reductase Required for the Biosynthesis of Cuticular Waxes. *Plant Physiology* **115**:501–510.

Yang, J., Lee, S. H., Goddard, M. E., and Visscher, P. M. (2011). GCTA: a tool for genome-wide complex trait analysis. *Am. J. Hum. Genet.* **88**:76–82.

Yu, J., Pressoir, G., Briggs, W. H., Vroh Bi, I., Yamasaki, M., Doebley, J. F., McMullen, M. D., Gaut, B. S., Nielsen, D. M., Holland, J. B., et al. (2006). A unified mixed-model method for association mapping that accounts for multiple levels of relatedness. *Nat. Genet.* **38**:203–208.

Zhang, W., Wu, Y., Schnable, J. C., Zeng, Z., Freeling, M., Crawford, G. E., and Jiang, J. (2012). High-resolution mapping of open chromatin in the rice genome. *Genome Res.* **22**:151–162.

Zhang, Y., Ngu, D. W., Carvalho, D., Liang, Z., Qiu, Y., Roston, R. L., and Schnable, J. C. (2017). Differentially Regulated Orthologs in Sorghum and the Subgenomes of Maize. *Plant Cell* **29**:1938–1951.

Zhao, W., Canaran, P., Jurkuta, R., Fulton, T., Glaubitz, J., Buckler, E., Doebley, J., Gaut, B., Goodman, M., Holland, J., et al. (2006). Panzea: a database and resource for molecular and functional diversity in the maize genome. *Nucleic Acids Res.* **34**:D752–7.

Zhou, X. (2017). A UNIFIED FRAMEWORK FOR VARIANCE COMPONENT ESTIMATION WITH SUMMARY STATISTICS IN GENOME-WIDE ASSOCIATION STUDIES. *Ann. Appl. Stat.* **11**:2027–2051.

Zhou, X., and Stephens, M. (2012). Genome-wide efficient mixed-model analysis for association studies. *Nature Genetics* **44**:821–824.