Reference Trait Analysis Reveals Correlations Between Gene Expression and Quantitative Traits in Disjoint Samples

Daniel A. Skelly,* Narayanan Raghupathy,* Raymond F. Robledo,* Joel H. Graber,*[†] and Elissa J. Chesler*¹

*The Jackson Laboratory, Bar Harbor, Maine 04609 and [†]MDI Biological Laboratory, Bar Harbor, Maine 04609

ORCID IDs: 0000-0002-2329-2216 (D.A.S.); 0000-0002-5642-5062 (E.J.C.)

ABSTRACT Systems genetic analysis of complex traits involves the integrated analysis of genetic, genomic, and disease-related measures. However, these data are often collected separately across multiple study populations, rendering direct correlation of molecular features to complex traits impossible. Recent transcriptome-wide association studies (TWAS) have harnessed gene expression quantitative trait loci (eQTL) to associate unmeasured gene expression with a complex trait in genotyped individuals, but this approach relies primarily on strong eQTL. We propose a simple and powerful alternative strategy for correlating independently obtained sets of complex traits and molecular features. In contrast to TWAS, our approach gains precision by correlating complex traits through a common set of continuous phenotypes instead of genetic predictors, and can identify transcript–trait correlations for which the regulation is not genetic. In our approach, a set of multiple quantitative “reference” traits is measured across all individuals, while measures of the complex trait of interest and transcriptional profiles are obtained in disjoint subsamples. A conventional multivariate statistical method, canonical correlation analysis, is used to relate the reference traits and traits of interest to identify gene expression correlates. We evaluate power and sample size requirements of this methodology, as well as performance relative to other methods, via extensive simulation and analysis of a behavioral genetics experiment in 258 Diversity Outbred mice involving two independent sets of anxiety-related behaviors and hippocampal gene expression. After splitting the data set and hiding one set of anxiety-related traits in half the samples, we identified transcripts correlated with the hidden traits using the other set of anxiety-related traits and exploiting the highest canonical correlation ($R = 0.69$) between the trait data sets. We demonstrate that this approach outperforms TWAS in identifying associated transcripts. Together, these results demonstrate the validity, reliability, and power of reference trait analysis for identifying relations between complex traits and their molecular substrates.

KEYWORDS systems genetics; RNA-seq; transcriptome-wide association study; statistical genetics; behavioral genetics; complex traits; functional genomics

A major goal of complex trait analysis is to discover pathways and mechanisms associated with disease. By definition, these traits exhibit hallmarks of genetic complexity including pleiotropy, epistasis, and gene–environment interaction. Genetic mapping is a powerful approach for detecting QTL that influence complex trait variation, but it has limited power for detecting small effect loci and can suffer from poor

mapping resolution, hindering the identification of causal genes. Moreover, these causal genetic variants do not always reside in relevant therapeutic targets. Therefore, many systems genetic strategies have emerged to correlate complex traits directly with molecular phenotypic variation, with the goal of constructing molecular networks that are correlated with trait variation from a trait-relevant tissue or cell population.

Ideally, trait correlation networks are constructed using direct phenotypic measurements for each member of a population. However, there are wide-ranging circumstances in which this approach is infeasible or impossible because it is physically, technically, or financially difficult to obtain all of the measures in the same individuals. To refer to phenotypes

Copyright © 2019 by the Genetics Society of America

doi: <https://doi.org/10.1534/genetics.118.301865>

Manuscript received December 11, 2018; accepted for publication May 14, 2019; published Early Online May 21, 2019.

Supplemental material available at FigShare: <https://doi.org/10.25386/genetics.7970984>.

¹Corresponding author: The Jackson Laboratory, 600 Main St., Bar Harbor, ME 04609. E-mail: elissa.chesler@jax.org

whose measurement on the same individual is infeasible or impossible, we will use the term incompatible phenotypes. Incompatible phenotypes arise in common experimental designs such as studies of susceptibility to exposure effects where the exposure affects physiology (e.g., predisposition to psychostimulant addiction) or studies of disease that relate early stage changes to late-stage outcomes (e.g., early molecular correlates predictive of Alzheimer's disease risk). Moreover, incompatible phenotypes arise when the original study population is no longer available but there is a desire to extend the study to a new set of traits, a situation that is common in human genetic analyses. Finally, phenotypes could be incompatible for strictly financial or logistical reasons, for example due to prohibitively high costs of genomic assays in large cohorts, leading to fractional collection of data on some samples and more thorough characterization of others.

One emerging approach for relating gene expression and complex traits measured in different cohorts of genetically diverse individuals is to exploit genetic variants that affect gene expression (gene expression QTL; eQTL) to impute transcript abundance from genotypes alone (Gamazon *et al.* 2015; Gusev *et al.* 2016, 2018; Barbeira *et al.* 2018; Mancuso *et al.* 2017). This enables estimation of the association between imputed gene expression and complex traits, an approach that has been called a transcriptome-wide association study (TWAS) (Gusev *et al.* 2016). However, the TWAS approach suffers from several limitations, most notably a reliance on strong local (presumably *cis*-acting) eQTL and consequent inability to impute transcript abundance for genes without detected eQTL. In contrast to using sparse, discrete *genotypes* to impute per-individual gene expression and infer correlation to complex traits; our approach uses shared variation across a rich set of quantitative, multidimensional *phenotypes* to infer gene expression correlates of phenotypic variability.

Rather than impute gene expression from genetic data, another strategy is to predict phenotypic data from other phenotypes. Hormozdiari *et al.* (2016a) used this approach to impute unmeasured phenotypes in the context of genome-wide association studies (GWAS). Specifically, the method of Hormozdiari *et al.* (2016a) uses the correlation structure in one set of traits to predict a single unmeasured target trait in a second cohort using only phenotypic data. In the present study, we extend this strategy to multivariate phenotyping and apply it to transcriptomics, providing a precise transcript-to-trait correlation approach that can be compared to the TWAS method.

We outline a simple method, reference trait analysis, to study relations between a set of complex traits of interest (*target traits*) and a set of high-dimensional molecular traits obtained in disjoint subsets of individuals. Reference trait analysis relates these two incompatible, multidimensional sets of phenotypes indirectly through the use of a shared set of *reference traits* measured in all individuals. Since target and molecular traits are not measured in the same individu-

als, direct comparisons are impossible. Instead, we relate these traits through reference traits. Reference traits are best chosen with *a priori* knowledge that they share biological underpinnings with target traits. This relationship between reference and target traits is exploited to compute scores from reference traits that capture variation in unmeasured target traits and can be directly related to transcriptional profiles. By design, our method is robust to the detection of transcript-trait associations for which the regulation is not genetic or is characterized by multiple weak, indirect genetic effects. Therefore, it captures biological variability associated with both genetic and environmental sources of vulnerability and has the potential to identify molecular networks of complex trait variation, even when there is insufficient power to detect a QTL or genome-wide significant SNP association.

In this study we develop and evaluate the reference trait analysis method using data from a previously published behavioral study of Diversity Outbred mice (Logan *et al.* 2013). Diversity Outbred mice are genetically unique; consequently, per-subject terminal traits such as brain gene expression can only be obtained in a single exposure condition. However, the approach we propose can be useful in any heterogeneous population for which a common reference set of traits is assessed. Our assessment data set consists of multiple measures of anxiety-related traits in a sample of Diversity Outbred mice, all of whom have been subjected to brain transcriptional profiling as well as measurements of two sets of related behaviors. We present an overview of our method, use these data to assess sample-size requirements, and quantify the method's reliability across a range of target-reference trait correlations. Finally, we test whether the reference trait method more faithfully recovers trait-gene expression correlations than the TWAS approach.

Materials and Methods

Mouse rearing and phenotyping

Diversity Outbred mice (J:DO; The Jackson Laboratory) are a heterogeneous stock derived from the same eight founder strains as the Collaborative Cross (Churchill *et al.* 2012; Svenson *et al.* 2012; Gatti *et al.* 2014; Chesler *et al.* 2016). In this study we used a subset ($N = 258$) of the 283 Diversity Outbred mice studied by Logan *et al.* (2013) with hippocampal gene expression profiled by RNA sequencing (RNA-seq) (see below). Mice in this study were from generations four to five (G4–G5) of the Diversity Outbred population. Briefly, each mouse was acclimated to the housing area and subjected to a brief testing battery which included a 20-min novel open-field test and a 10-min light–dark test, among other common behavioral tasks. The open-field and light–dark tests are used to measure exploratory activity and approach–avoidance behavior. Many complex trait measures can be extracted from these tasks. For this analysis, we chose two sets of informative measures (Supplemental Material,

Table S1). Complete details of animal rearing, husbandry, and phenotyping are presented in Logan *et al.* (2013). Mice were sacrificed using decapitation which was necessary to preserve fresh brain tissue in the absence of drug or asphyxiation. All procedures and protocols were approved by The Jackson Laboratory Animal Care and Use Committee, and were conducted in compliance with the National Institutes of Health Guidelines for the Care and Use of Laboratory Animals.

Genotyping

DNA was prepared from tail biopsies and samples were genotyped using the Mouse Universal Genotyping Array (Morgan *et al.* 2016). We obtained genotypes at 7802 markers from arrays processed by GeneSeek (Lincoln, NE). We used intensities from each array to infer the haplotype blocks in each individual Diversity Outbred genome using a hidden Markov model (Gatti *et al.* 2014).

Gene expression profiling

Total hippocampal RNA was isolated using the TRIzol Plus RNA purification kit (Life Technologies, Carlsbad, CA) with on-column DNase digestion. Samples for RNA-seq analysis were prepared using the TruSeq Kit (Illumina, San Diego, CA) according to the manufacturer's protocols and subjected to paired-end, 100-bp sequencing on the HiSeq 2000 (Illumina) per manufacturer's recommendations. RNA-seq was performed in nine sequencing runs with two technical replicates for each sample, resulting in an average sequencing depth of ~ 24 million reads per sample after pooling technical replicates. To obtain estimates of gene expression, we aligned reads to individualized diploid genomes using the bowtie aligner (Langmead *et al.* 2009) and quantified transcript abundance by allocating multimapping reads using the EM algorithm with RSEM (Li and Dewey 2011) as described in Munger *et al.* (2014). Raw counts in each sample were normalized to the upper quartile value and transformed to normal scores.

Reference trait analysis

We conducted reference trait analysis using R version 3.3.2 (R Core Team 2016). Canonical correlation analysis was carried out using the `cancor` function in base R. We regressed out the effect of sex on each phenotype because it is not of primary interest in this study. An example walk-through of a reference trait analysis and code to carry out the analyses described in this article are available at https://daskelly.github.io/reference_traits/reference_trait_analysis_walkthrough.html.

To examine the power and robustness of reference trait analysis, we simulated data with varying sample sizes and canonical correlation coefficients. We based our simulations on the anxiety phenotype data, consisting of open-field exploration and light-dark box behavioral measures. Specifically, for each of 1000 simulations, we started with the covariance matrix computed from 5 open-field and 5 light-dark box traits and randomly increased or decreased each of

the $5 \times 5 = 25$ interdata-set covariances by 20%. We then simulated multivariate normal phenotype data with the specified covariance matrix. This procedure resulted in two multivariate data sets (simulated open-field and light-dark box traits), where the covariance structure *within* each data set was similar to that in the real data but with different covariances *between* data sets. When a canonical correlation analysis was carried out on each pair of simulated data sets, the magnitude of the first canonical correlation coefficient varied between $R = 0.1$ and $R = 0.98$ due to the variation in interdata-set covariances.

We simulated gene expression traits with exact correlation to the first target trait canonical variable $\bar{v}_1$ in the simulated data set. To simulate a random vector of observations with defined correlation to an existing vector, we took advantage of the geometric property that the cosine between two mean-centered vectors equals their correlation. Therefore, a random vector with defined correlation to an existing vector can be computed by starting with random draws from a normal distribution, mean centering, and applying standard linear algebra operations.

After hiding target traits in half of the individuals and gene expression data in the other half, we conducted reference trait analysis and quantified performance as the fraction of the time true trait-gene expression correlations were detected using a 10% false discovery rate (FDR) threshold. *P*-values for trait-gene expression correlations were calculated using a two-sided *T* statistic and correlations deemed significant at a 10% FDR were identified using *q*-values (Storey and Tibshirani 2003).

Imputing gene expression using TWAS

We divided the anxiety data set in half and considered open-field measurements as target traits, hiding gene expression measurements for the animals where we did not hide open-field traits. For the TWAS strategy, our training cohort consisted of animals with genotypes and gene expression data, and our testing cohort consisted of animals with genotypes and open-field traits (*i.e.*, training/testing labels are reversed from reference trait analysis, see Figure S1). Diversity Outbred mice are an outbred population with genomic ancestry derived from eight inbred founder strains. We used methods implemented in R/qtl2 software (<http://kbroman.org/qtl2/>) to impute SNP variation in each mouse from array-based genotypes obtained at coarser resolution (see above) using known SNP genotypes present in founder haplotypes. This resulted in genotypes for ~ 30 million SNPs. Given the limited number of generations of outbreeding, haplotype blocks in Diversity Outbred mice typically stretch for megabases (Svenson *et al.* 2012), leading to strong local linkage disequilibrium (LD). As such, we used PLINK version 1.9 (Purcell *et al.* 2007) to prune variants in very strong LD in the eight founder strains, using the following parameters: indep-pairwise 200 kb 40 kb 0.95. This procedure reduced the number of SNPs to 235,335 with minimal loss of information.

To impute gene expression, we used the PredictDB module of PrediXcan (Gamazon *et al.* 2015) to build predictive models of gene expression from local genotypes within 10 Mb of each gene, with sex included as a covariate. We conducted 1000 permutations with random 50:50 divisions of the anxiety data set to account for stochastic sampling effects. For each replicate we obtained predictive models of gene expression by running PredictDB on the training cohort and applied them to the testing cohort to impute gene expression. Following Gamazon *et al.* (2015; <https://github.com/hakym-lab/PrediXcan>), we considered only genes with models that were significantly predictive of gene expression ($\text{FDR} \leq 5\%$). Finally, we calculated correlations between imputed gene expression and a summary measure of the target traits (first canonical variable) in the testing cohort. Results were nearly identical whether we correlated to the first canonical variable or first principal component of the target traits (median transitive reliability 45% vs. 44%), but correlations to first canonical variable allow for direct comparison with results from reference trait analysis.

Scoring gene sets to assess retrieval of known anxiety-related genes

To score gene sets for relevance to anxiety, a rater with expertise in behavioral neuroscience who was blind to the analysis methods scored a combined list of all gene sets obtained herein. We assigned a score of zero to irrelevant data sets; a score of two to gene sets with general brain or behavior relevance; and a score of four to anxiety-relevant data sets in which either the gene set was generated in an anxiety-relevant experiment, the gene set consisted of genes interacting with a compound known to be anxiolytic or anxiogenic, or the gene set was a Gene Ontology annotation set with direct biological relevance to anxiety. For compounds, a single MEDLINE query of the compound name and the term anxiety was performed and the results of the query were examined for overall conceptual relevance.

Data availability

Raw RNA-seq gene expression data from the hippocampus of 258 Diversity Outbred mice are available from ArrayExpress (accession number E-MTAB-7637). Genotypes of the 258 Diversity Outbred mice are available from The Jackson Laboratory Diversity Outbred database (<https://www.jax.org/research-and-faculty/genetic-diversity-initiative/tools-data/diversity-outbred-database>). A processed and normalized gene expression matrix is available as Supplemental Dataset 1. Phenotype data acquired via the open-field and light–dark box paradigms are available as Supplemental Datasets 2 and 3. An example walk-through of a reference trait analysis and code to carry out the analyses described in this article are available at https://daskelly.github.io/reference_traits/reference_trait_analysis_walkthrough.html. Supplemental material available at FigShare: <https://doi.org/10.25386/genetics.7970984>.

Figure 1 Overview of reference trait analysis. Target and reference traits are measured in a set of training individuals (top plots, gray triangles), while reference traits and gene expression are measured in test individuals (bottom plots and $\times$ symbols). (1) Canonical correlation is used to identify a linear combination of reference traits (top blue $\rightarrow$) that best captures variation in the traits of interest (red $\rightarrow$; yellow curve connecting arrows represents canonical correlation analysis). (2) The weights derived from canonical correlation analysis are applied to reference traits in the testing population to derive reference trait scores for each individual (projected reference traits; bottom blue $\rightarrow$). (3) Projected reference traits are correlated with molecular phenotypes.

Results and Discussion

Outline of approach

The reference trait analysis procedure is straightforward and relies on the well-characterized statistical technique of canonical correlation analysis. Beginning with a population of individuals, reference traits (labeled using the variable U) are measured on all individuals, target traits (labeled with V) on the training cohort, and high-dimensional molecular traits (labeled with M) on the testing cohort (Figure 1). Although target traits and their molecular correlates are of primary interest, the choice of reference traits is an important aspect of the method. First, as we will show, the strength of the multivariate relationship between target and reference traits is a key parameter determining the power to detect trait–transcript correlations. Second, because our method leverages shared variation between target and reference traits, it identifies trait–transcript correlations driven by the portion of target trait variation that is shared with reference traits. For example, studying addiction-related traits using novelty behaviors as reference traits would be expected to uncover transcripts associated with addiction behaviors through biological pathways that also contribute to the etiology of novelty-seeking behaviors.

To conduct reference trait analysis, we employ canonical correlation (Hotelling 1936), which can be thought of as a parent analysis of the more familiar multiple regression. A multiple regression of Y on X models the relationships between multiple X measures $X_1, X_2, \dots, X_p$ and univariate Y .

In contrast, canonical correlation reveals the magnitude and nature of relationships between multivariate U and V , e.g., $U_1, U_2, \dots, U_p$ and $V_1, V_2, \dots, V_q$. Specifically, canonical correlation identifies linear combinations of two multivariate measures U and V such that the (univariate) linear combinations of each measure $\vec{u}$ and $\vec{v}$, known as canonical variables, are maximally correlated. In this study we use canonical correlation to build linear combinations of reference traits (transforming U to $\vec{u}$) that maximize shared variance with target traits (V) in the set of training individuals. The possible number of canonical variables is limited to the size of the smaller of U and V , and each successive covariate captures a diminishing proportion of the shared variance between the traits. In this study we focus on the first canonical variable, $\vec{u}_1$ or $\vec{v}_1$, which explains the largest fraction of shared variance between U and V . This quantity can be thought of as a summary of each set of traits analogous to their first principal component, but rather than being aligned with the axis of maximal variation among a single set of variables, it is aligned in the direction of maximal shared variation between the two sets of traits U and V . For data sets with a very large number of reference and/or target traits (i.e., $P \gg n$), sparse canonical correlation analysis (Witten and Tibshirani 2009; Wilms and Croux 2016) may reduce overfitting, but this situation is not common when relating two sets of traits U and V that contain organism-level phenotypes as opposed to molecular features.

The analysis of training data defines canonical coefficients that can be used to compute first canonical variables from individual-level trait data (i.e., transform U to $\vec{u}_1$ or V to $\vec{v}_1$). We use these coefficients learned from the training data (Figure 1, top) to transform reference trait data from the testing cohort U' , which projects these data in the direction of maximal shared variation with target traits. Multiple “projected” traits could be calculated using canonical coefficients corresponding to the first, second, or subsequent canonical variables; here we focus on the first canonical variable as explained above. The projected trait $\vec{u}'_1$ represents a linear combination of reference traits that optimally captures the largest portion of variation shared between reference and target traits due to their underlying genetic and environmental covariation. The univariate projected trait vector is then compared to high-dimensional genomic measurements to extract molecular phenotypes in one sample set that covary with target traits from another group (Figure 1, bottom right). For interpretability, it may be desirable to assess a single specific target trait in the canonical correlation, rather than a derived linear combination of a set of targets.

Transitive reliability captures global patterns of covariation between incompatible traits

Reference trait analysis reveals covariation between molecular phenotypes and target trait variation. There are many possible applications of this strategy. For example, in addiction research, many studies evaluate transcriptional response to drug exposure but are unable to evaluate the predisposing

characteristics of a drug-naïve brain that associate with addiction-related behaviors. Using a reference trait strategy, one can evaluate the transcriptomes of drug-naïve brains and relate them to the response to drug self-administration through a set of reference traits that do not involve drug exposure. We have previously estimated the association of novelty seeking and drug self-administration in mice, revealing a canonical correlation of 0.61 among these sets of traits (Dickson *et al.* 2015).

To evaluate whether the reference-traits strategy could be applied to find transcriptional correlates of drug self-administration, we used a data set where reference, target, and molecular trait profiling were performed on the same individuals to allow for assessment of the robustness and sample-size requirements of the method. Specifically, we studied relationships between two distinct sets of anxiety-related traits and hippocampal gene expression, where all traits were measured in each of the 258 Diversity Outbred mice (Logan *et al.* 2013). The anxiety-related traits consisted of 11 measurements of open-field arena exploration behaviors and 5 measurements of light–dark box behaviors (Table S1). A canonical correlation analysis of these two sets of traits yielded a statistically significant model ($F_{55,1123.75} = 4.48$, $P < 2 \times 10^{-16}$, Wilk’s $\lambda = 0.400$) that had a first canonical correlation coefficient of magnitude 0.69. This was higher than all univariate correlations between open-field and light–dark box traits (median 0.11, maximum 0.65), and similar in magnitude to the shared variation revealed by the first canonical variable in the motivating analysis of novelty-related behaviors and cocaine self-administration (Dickson *et al.* 2015). We arbitrarily designated the open-field traits as target traits and light–dark box traits as reference traits. For reference trait analysis, we hid gene expression data for half of the mice (training set) and open-field data for the remaining mice (testing set). As described below, we varied simulated total sample size to explore the performance of our method.

In this evaluation of reference trait analysis, we know the true values of all hidden data and can directly evaluate the power of the method to reveal gene expression patterns associated with target trait variation. Specifically, we estimate canonical coefficients (weights to calculate canonical variables) from the training set and use them to calculate projected traits $\vec{u}'_1$ in the testing set. To quantify the performance of reference trait analysis when the true answer is known, we computed correlations in testing-set animals between gene expression E and either (1) the first projected trait $\vec{u}'_1$, $\text{cor}(E, \vec{u}'_1)$, or (2) the first canonical variable computed using hidden target traits $\vec{v}'_1$, $\text{cor}(E, \vec{v}'_1)$. The latter quantity, the “truth,” is unavailable in a real application of reference trait analysis. A set of reference traits that perfectly captures all variation in target traits would result in a vector of gene expression–trait correlations that is identical whether the target traits were known or projected from reference traits (i.e., the reference traits serve as a perfect surrogate for target traits). We define transitive reliability as the correlation between these vectors, i.e., $\text{cor}[\text{cor}(E, \vec{u}'_1), \text{cor}(E, \vec{v}'_1)]$. High

on x-axis indicates the number of individuals used in each of the training and testing groups (the number of individuals phenotyped for target traits and the number with high-dimensional molecular phenotypes, respectively). Black line indicates magnitude of first canonical correlation calculated from full data set. Color indicates whether training and testing groups were fully or partially independent.

Figure 2 Reference trait analysis reveals overall patterns of covariation between incompatible traits. (A) Relationship between canonical correlation and transitive reliability. To evaluate the mathematical relationship between these quantities, we simulated two vectors with known correlation to represent the canonical variables and calculated transitive reliability with real gene expression data. Canonical correlation shown is absolute value and transitive reliability is sign matched. (B) Sample-size increases lead to higher and more precise transitive reliability. Plot shows transitive reliability estimated using anxiety data with animals subsampled as described in the main text. Sample size

transitive reliability would indicate that strong correlations between gene expression and target traits are likely to be identified using projected traits.

Transitive reliability, estimated using real gene expression data and simulated canonical variables with known correlation, scales linearly with the magnitude of the canonical correlation coefficient (Figure 2A), confirming our intuition that greater sharing of variation between target and reference traits increases the utility of leveraging reference traits to understand target trait variation. We divided the anxiety data set into equally sized subsets (partially overlapping for larger sample sizes) to examine the dependence of transitive reliability on sample size. The canonical correlation was upwardly biased for small sample sizes ($N < 90$; data not shown), as has previously been recognized (e.g., Thompson 1990). When we used Wherry's correction as suggested by Thompson (1990), canonical correlations no longer depended on sample size (linear model; $P > 0.8$). Overall, transitive reliability asymptotically approached the magnitude of the canonical correlation coefficient calculated from the full data set (Figure 2B, black line), demonstrating that global patterns of trait–gene expression correlation can be recovered with relatively modest sample sizes using the reference trait approach. In contrast, weights from the smallest (fifth) canonical variable, which captures little shared variation between data sets, produced low transitive reliabilities (median 0.11).

Reference trait analysis successfully identifies known trait correlations

Ultimately, the primary goal of reference trait analysis is to identify molecular correlates of unmeasured phenotypes. To discover these correlates, individual gene expression levels are correlated to projected traits. To test this strategy, we first employed reference trait analysis on the anxiety-related phenotype data described above. After randomly splitting the data set and withholding open-field data (arbitrarily designated as target traits) in half of the individuals, we identified gene expression levels correlated to projected reference traits. We found high overlap between the genes most strongly correlated with hidden target trait canonical variable 1 and those

most strongly correlated with projected traits (23% overlap among genes with the top 5% of correlations to each trait, compared to 2.5% expected overlap; $P < 1 \times 10^{-15}$, Fisher's exact test). Across all genes, including those with weaker correlations, we found that the vector of trait–gene expression correlations computed using reference trait analysis showed significant similarity to the true correlations ($P < 0.001$, permutation test using generalized Jaccard similarity statistic). Moreover, in contrast to the alternative methods for identifying trait–gene expression correlations discussed above, some correlations detected using reference trait analysis involved genes with no significant eQTL (e.g., 42% of the top 50 correlations). These genes, which are demonstrably associated with trait variation, would not be detectable using TWAS-type approaches.

To examine the power and robustness of reference trait analysis across a wide range of biologically plausible parameter values, we conducted extensive simulations. We simulated data across a range of sample sizes (100, 200, 300, ..., 1000, 1200, 1400, ..., 2000) and enforced a similar covariance structure to the observed data. Specifically, data were simulated using observed covariances within each set of anxiety traits, but we perturbed covariances between the two sets of traits to generate data sets with varying canonical correlations. We then simulated gene expression levels with known correlation to the first target trait canonical variable, $\bar{v}_1$ ($\rho = 0.2, 0.225, 0.25, \dots, 0.9$ with 20 genes each). We simulated trait data and gene expression data at random for each of 1000 simulations for each sample size.

For each simulation, after hiding target traits in half the individuals and gene expression data in the other half, we conducted reference trait analysis. We computed projected reference traits, correlated to gene expression, and quantified performance as the fraction of true trait–gene expression correlations that were detected using a 10% FDR threshold. For high trait–gene correlations ($\rho > 0.6$) and strong target–reference trait canonical correlations ($R = 0.7$ or 0.9), the correlation of interest was essentially always detected (Figure 3). For lower target–reference trait canonical correlations ($R = 0.5$), even relatively modest true trait–gene expression correlations (e.g., $\rho = 0.3$) were often detected with sample

Figure 3 Reference trait analysis identifies simulated trait–gene expression correlations across a wide variety of parameter values. Sample size plotted along x-axis is the number of individuals used in each of the training and testing groups (equal sample size for the two groups, where the training group consists of individuals phenotyped for target traits and the testing group consists of those with high-dimensional molecular phenotypes). True correlation (y-axis) indicates correlation between first target trait canonical variable ($\bar{v}_1$) and simulated gene expression. Facets indicate magnitude of canonical correlation coefficient between reference and target traits (R listed along gray strips, ± 0.02). Pseudocoloring indicates the fraction of trait–gene expression correlations detected by reference trait analysis for a variety of sample size and true magnitudes of correlation. A 10% FDR threshold implies that 10% of trait–gene expression correlations deemed significant are false positives. Navy and magenta contour lines depict regions above/below which true positive trait–gene expression correlations are detected >80% and <20% of the time, respectively.

sizes above ~ 300 individuals (Figure 3). Thus, reference trait analysis was a highly effective means for identifying trait–gene expression correlations across a diverse range of practical sample sizes, typical values for trait-to-gene expression correlation, and canonical correlation parameters.

Comparison of reference trait analysis to related approaches

An alternative approach to identifying genes associated with complex traits is to make use of known genetic variation that regulates gene expression (eQTL). There has been considerable recent interest in methods that integrate complex trait associations and gene expression genetics to identify genes whose expression is associated with trait variation (Nica *et al.* 2010; Wallace *et al.* 2012; He *et al.* 2013; Gamazon *et al.* 2015; Gusev *et al.* 2016; Hormozdiari *et al.* 2016b; Zhu *et al.* 2016; Hauberg *et al.* 2017; Wen *et al.* 2017). Several methods perform tests of the hypothesis that genome-wide association (GWA) signals and eQTL are truly colocated *vs.* independent but appearing colocated due to LD (Nica *et al.* 2010; Wallace *et al.* 2012; Giambartolomei *et al.* 2014; Fortune *et al.* 2015; Hormozdiari *et al.* 2016b; Zhu *et al.* 2016; Hauberg *et al.* 2017; Wen *et al.* 2017). Another approach that is more directly applicable to the experimental designs studied herein is to harness strong genetic predictors of gene expression variation (eQTL) to impute transcriptomes in genotyped and phenotyped cohorts, which allows detection of trait–expression correlations (the TWAS approach; Gamazon *et al.* 2015; Gusev *et al.* 2016, 2018; Mancuso *et al.* 2017; Barbeira *et al.* 2018). TWAS is an approach that is complementary to reference trait analysis and has been a particularly powerful method for discovery of candidate genes driving GWA signals detected in very large

human cohorts (tens or hundreds of thousands of individuals). Figure S1 provides a comparison of genotype, phenotype, and gene expression data in the reference traits and TWAS strategies. One weakness of the TWAS approach is that it hinges on the presence of detectable eQTL (typically local, presumably *cis*-acting eQTL; but see He *et al.* 2013 and Vervier and Michaelson 2016). In humans, even panels of 1000 individuals with gene expression measurements only result in a modest number of genes (500–4000) with significant *cis*-heritability that can be imputed in the cohort lacking gene expression data (Gusev *et al.* 2016). In contrast, reference trait analysis has no requirement for detection of eQTL. Therefore, it is amenable to detect correlation between transcripts with complex expression regulatory mechanisms and traits of similarly complex regulation.

Although TWAS and reference trait analysis use different data types, both are tools for inferring relations among complex traits and transcript abundance, so we sought to compare their performance on the same data set. For TWAS, we used methods implemented in the software suite PrediXcan (Gamazon *et al.* 2015). We randomly divided our anxiety data set in half and considered open-field measurements as target traits. We withheld gene expression measurements in half of the animals; therefore, only genotype and reference trait data were visible for all animals. We built predictive models of gene expression from the training cohort of mice, applied these models to impute gene expression in the testing cohort, and calculated correlations between imputed gene expression and a summary measure of the target traits (first canonical variable). We conducted 1000 permutations with random 50:50 divisions of the anxiety data set to account for stochastic sampling effects. For each replicate, we compared global trait–gene expression correlations for PredictDB-imputed

gene expression vs. those computed using projected traits obtained with our new method. In the former case, trait data are available and gene expression data are imputed, while in the latter case gene expression data are available and trait data are imputed.

For direct comparisons between reference trait analysis and TWAS, we ran reference trait analysis using only genes that were significantly predicted by the PredictDB module of PrediXcan (FDR < 5%; see *Materials and Methods*). Across the 1000 permutations, we imputed gene expression for a mean 12,250 genes (range 11,640–12,750; this mean represents ~70% of total 17,539 genes measured), indicating that a substantial fraction of genes has insufficient local genetic signal for accurate imputation. An advantage of reference trait analysis is that it is not limited by the presence of strong eQTL and all genes can be tested for association with projected reference traits. For each of the 1000 permutations, we computed the transitive reliability of TWAS and of reference trait analysis. Reference trait analysis more accurately captured global patterns of trait–transcript correlation than TWAS (Figure 4). Specifically, transitive reliability for target trait first canonical variable–gene expression correlations was higher using the reference trait approach (measured gene expression and projected reference traits) compared to the TWAS approach (imputed gene expression and measured traits) for 92.7% of simulations (Figure 4; Figure S2 shows an example of results from one permutation). Thus, we show empirically that reference trait analysis outperforms TWAS in the mouse anxiety data set.

In addition to the quantitative comparison of the methods, we sought to determine which approach provided the best retrieval of known anxiety-related genes. To perform this analysis we made use of GeneWeaver’s database of gene sets curated from multiple sources (Baker *et al.* 2016). The top 400 genes identified using each analysis method were entered as three gene lists, containing results of (1) direct correlations of target traits (first canonical variable) to transcripts, (2) correlations of target traits to TWAS-imputed gene expression, and (3) correlations of gene expression to projected traits calculated using reference trait analysis. Each gene list was compared to every gene set in the GeneWeaver database via Jaccard similarity. For each of the three gene lists, the name, ID, and description for the top 249 similar gene sets were exported and combined to form a single list of results. A rater with expertise in behavioral neuroscience who was blind to the analysis methods scored the combined list of all similar gene sets obtained in these three analyses. Gene sets were scored discretely based on relevance of their semantic descriptions to anxiety, with scores assigned based on three levels (irrelevant, generally relevant to brain or behavior, and specifically relevant to anxiety). We found that true open-field first canonical variable–gene expression correlations had highest relevance to anxiety. The top truly correlated genes were similar to gene sets more relevant to anxiety than those genes identified using reference traits or those using TWAS ($P = 0.0065$ and $P = 1.5 \times 10^{-14}$, re-

Figure 4 Reference trait analysis recovers true trait–gene expression correlations more accurately than TWAS. Binned hexagon plot shows the results of 1000 random samples where the anxiety data set was split into two halves randomly designated the training and testing groups. Reference trait analysis and TWAS were used to recover trait–gene expression correlations (correlating gene expression to projected traits in the case of reference trait analysis, and correlating imputed gene expression to the trait first canonical variable in the case of TWAS). The true values of both the traits and gene expression are known in this data set, but were hidden when running reference trait analysis or TWAS. For each method, the transitive reliability was computed based on comparing projected/imputed data to true values.

spectively; two-sided Fisher’s exact test). Nevertheless, reference trait analysis performed significantly better than TWAS at identifying genes with similarity to anxiety-relevant gene sets ($P = 7.3 \times 10^{-6}$).

Finally, another alternative to relating traits and transcripts between population cohorts is to make use of polygenic risk predictors trained using genome-wide genotypes and phenotypes, and applied to individuals with genotypes but missing phenotypes (in this case, samples with only transcriptional profiles available) (Makowsky *et al.* 2011; Dudbridge 2013; Wray *et al.* 2013). However, theoretical considerations and empirical results suggest that this approach generally requires sample sizes much larger than 1000 individuals to obtain accurate predictions (Dudbridge 2013). In the context of reference trait analysis, relating complex reference and target traits that share high canonical correlation implicitly leverages the common polygenic or omnigenic (Boyle *et al.* 2017) basis of these traits by making use of all of the information contained in continuous quantitative variation.

Conclusions

We have described a general method for exploring trait covariation among incompatible and independently collected phenotypes studied in disjoint samples of genetically diverse

individuals to extract molecular networks associated with disease. Our method uses canonical correlation analysis, a standard multivariate statistical method, to relate incompatible phenotypes using a set of reference traits measured on all individuals. Our analyses demonstrate that this approach performs well over a range of parameters encountered in the study of trait correlations (Dickson *et al.* 2015) and under sample-size requirements that are practical to obtain. This approach can be useful both for capturing global patterns of covariation between target traits and high-dimensional molecular phenotypes, as well as for identifying specific molecular correlates to target traits. Our method identifies trait–gene expression associations and we do not assert that these associations are necessarily causal, as has been recognized by studies relating GWAS results and eQTL (Gamazon *et al.* 2015; Gusev *et al.* 2016; Hauberg *et al.* 2017). Transcripts that do casually affect target traits may do so directly (via pleiotropic effects on both reference and target traits) or indirectly (through reference traits). Statistical methods such as structural equation modeling, path analysis, or mediation analysis (*e.g.*, Chick *et al.* 2016) may provide insight into the nature of the relationship between these traits.

When will reference trait analysis be a useful tool? Intuitively, and as demonstrated in Figure 3, large sample sizes, precise trait measurements, and high shared variance between reference and target traits would allow for the most accurate estimation of canonical correlation coefficients and high power to detect correlations to molecular phenotypes. Although our method could be applied in a wide variety of scenarios, it is likely to be particularly useful for studies of highly complex, polygenic, multidimensional traits (*e.g.*, behavior, physiology, and morphology) in cohorts of modest size. As with any method that applies information learned from one cohort to biological measures from another cohort, reference trait analysis performs best in the absence of systematic differences (*i.e.*, heterogeneity in population characteristics or unequal genetic composition caused by population structure) between the training and testing cohorts. This requirement can equivalently be formulated in terms of exchangeability of individuals in the training and testing cohorts. Applying reference trait analysis when individuals are not exchangeable (for example, where training and testing individuals are sampled from different genetic reference populations) is not likely to perform as well, although the method may still be useful if the genetic architecture of target and reference traits is similar in both populations. For example, we might expect reference trait analysis to enable extrapolation from Diversity Outbred mice to Collaborative Cross mice derived from the same founders but, depending on the nature of trait regulation, it may not enable extrapolation to the BXD population which derives from different founders. Likewise, in human studies drawn from similar populations, we would anticipate the approach to perform better than it would across distinct subpopulations. For very large cohorts of individuals, where obtaining suitable reference traits may be difficult, polygenic scores

based on either genetic predictors alone or on a combination of genetic and environmental risk factors (Dudbridge *et al.* 2018) may be a valuable approach for predicting phenotypic variation in a test cohort that can then be correlated with molecular networks.

Although our application of reference trait analysis involves correlations to high-dimensional molecular phenotypes, the method could, in principle, be applied to any sets of phenotypes that are multivariate in nature. Moreover, the high relative performance of our method underscores the importance of extensive phenotyping using quantitative traits rather than relying on binary indicators of disease and disease-related phenotypes that may mask complex underlying etiologies. We anticipate that the framework outlined in this study will be increasingly useful as studies of diverse, genetically unique populations become more widespread. A useful future extension to this approach would incorporate statistical techniques such as sparse canonical correlation analysis (Witten and Tibshirani 2009; Wilms and Croux 2016), which could permit inference in phenome-level studies where the target or reference traits are high dimensional. Overall, our approach is likely to be particularly important in functional genomics studies, those using postmortem subjects, and large population studies in which individuals are unavailable for further characterization.

Acknowledgments

We thank Daniel M. Gatti for assistance with processing mouse genotypes and obtaining genotype probabilities for mapping, Juliet Ndukum for implementing an early prototype of the reference trait analysis methodology, and Timothy Reynolds for assistance with the GeneWeaver platform. We thank Jackson Laboratory Genome Technologies for assistance with library preparation and sequencing of RNA-seq samples. This study was supported by National Institutes of Health grants R01 DA-037927 (E.J.C.), R01 AA-018776 (E.J.C.), and P50 DA-039841 (E.J.C.) and by program funds to E.J.C. from The Jackson Laboratory.

Author contributions: D.A.S.: conceptualization, data curation, formal analysis, methodology, software, visualization, and writing. N.R.: data curation, formal analysis, and software. R.F.R.: investigation, methodology. J.H.G.: formal analysis. E.J.C.: conceptualization, funding acquisition, methodology, project administration, supervision, and writing.

Literature Cited

- Baker, E., J. A. Bubier, T. Reynolds, M. A. Langston, and E. J. Chesler, 2016 GeneWeaver: data driven alignment of cross-species genomics in biology and disease. *Nucleic Acids Res.* 44: D555–D559. <https://doi.org/10.1093/nar/gkv1329>
- Barbeira, A. N., S. P. Dickinson, J. M. Torres, R. Bonazzola, J. Zheng *et al.*, 2018 Exploring the phenotypic consequences of tissue specific gene expression variation inferred from GWAS summary statistics. *Nat. Commun.* 9: 1825.

- Boyle, E. A., Y. I. Li, and J. K. Pritchard, 2017 An expanded view of complex traits: from polygenic to omnigenic. *Cell* 169: 1177–1186. <https://doi.org/10.1016/j.cell.2017.05.038>
- Chesler, E. J., D. M. Gatti, A. P. Morgan, M. Strobel, L. Trepanier *et al.*, 2016 Diversity outbred mice at 21: maintaining allelic variation in the face of selection. *G3 (Bethesda)* 6: 3893–3902. <https://doi.org/10.1534/g3.116.035527>
- Chick, J. M., S. C. Munger, P. Simecek, E. L. Huttlin, K. Choi *et al.*, 2016 Defining the consequences of genetic variation on a proteome-wide scale. *Nature* 534: 500–505. <https://doi.org/10.1038/nature18270>
- Churchill, G. A., D. M. Gatti, S. C. Munger, and K. L. Svenson, 2012 The Diversity Outbred mouse population. *Mamm. Genome Off. J. Int. Mamm. Genome Soc.* 23: 713–718. <https://doi.org/10.1007/s00335-012-9414-2>
- Dickson, P. E., J. Ndukum, T. Wilcox, J. Clark, B. Roy *et al.*, 2015 Association of novelty-related behaviors and intravenous cocaine self-administration in Diversity Outbred mice. *Psychopharmacology (Berl.)* 232: 1011–1024. <https://doi.org/10.1007/s00213-014-3737-5>
- Dudbridge, F., 2013 Power and predictive accuracy of polygenic risk scores. *PLoS Genet.* 9: e1003348 (erratum: *PLoS Genet.* 9: 10.1371/annotation/b91ba224-10be-409d-93f4-7423d502cba0). <https://doi.org/10.1371/journal.pgen.1003348>
- Dudbridge, F., N. Pashayan, and J. Yang, 2018 Predictive accuracy of combined genetic and environmental risk scores. *Genet. Epidemiol.* 42: 4–19. <https://doi.org/10.1002/gepi.22092>
- Fortune, M. D., H. Guo, O. Burren, E. Schofield, N. M. Walker *et al.*, 2015 Statistical colocalization of genetic risk variants for related autoimmune diseases in the context of common controls. *Nat. Genet.* 47: 839–846 (erratum: *Nat. Genet.* 47: 962). <https://doi.org/10.1038/ng.3330>
- Gamazon, E. R., H. E. Wheeler, K. P. Shah, S. V. Mozaffari, K. Aquino-Michaels *et al.*, 2015 A gene-based association method for mapping traits using reference transcriptome data. *Nat. Genet.* 47: 1091–1098. <https://doi.org/10.1038/ng.3367>
- Gatti, D. M., K. L. Svenson, A. Shabalin, L.-Y. Wu, W. Valdar *et al.*, 2014 Quantitative trait locus mapping methods for diversity outbred mice. *G3 (Bethesda)* 4: 1623–1633. <https://doi.org/10.1534/g3.114.013748>
- Giambartolomei, C., D. Vukcevic, E. E. Schadt, L. Franke, A. D. Hingorani *et al.*, 2014 Bayesian test for colocalisation between pairs of genetic association studies using summary statistics. *PLoS Genet.* 10: e1004383. <https://doi.org/10.1371/journal.pgen.1004383>
- Gusev, A., A. Ko, H. Shi, G. Bhatia, W. Chung *et al.*, 2016 Integrative approaches for large-scale transcriptome-wide association studies. *Nat. Genet.* 48: 245–252. <https://doi.org/10.1038/ng.3506>
- Gusev, A., N. Mancuso, H. K. Finucane, Y. Reshef, L. Song *et al.*, 2018 Transcriptome-wide association study of schizophrenia and chromatin activity yields mechanistic disease insights. *Nat. Genet.* 50: 538–548.
- Haugberg, M. E., W. Zhang, C. Giambartolomei, O. Franzén, D. L. Morris *et al.*, 2017 Large-scale identification of common trait and disease variants affecting gene expression. *Am. J. Hum. Genet.* 100: 885–894 (erratum: *Am. J. Hum. Genet.* 101: 157). <https://doi.org/10.1016/j.ajhg.2017.04.016>
- He, X., C. K. Fuller, Y. Song, Q. Meng, B. Zhang *et al.*, 2013 Sherlock: detecting gene-disease associations by matching patterns of expression QTL and GWAS. *Am. J. Hum. Genet.* 92: 667–680. <https://doi.org/10.1016/j.ajhg.2013.03.022>
- Hormozdiari, F., E. Y. Kang, M. Bilow, E. Ben-David, C. Vulpe *et al.*, 2016a Imputing phenotypes for genome-wide association studies. *Am. J. Hum. Genet.* 99: 89–103. <https://doi.org/10.1016/j.ajhg.2016.04.013>
- Hormozdiari, F., M. van de Bunt, A. V. Segrè, X. Li, J. W. J. Joo *et al.*, 2016b Colocalization of GWAS and eQTL signals detects target genes. *Am. J. Hum. Genet.* 99: 1245–1260. <https://doi.org/10.1016/j.ajhg.2016.10.003>
- Hotelling, H., 1936 Relations between two sets of variates. *Biometrika* 28: 321–377. <https://doi.org/10.1093/biomet/28.3-4.321>
- Langmead, B., C. Trapnell, M. Pop, and S. L. Salzberg, 2009 Ultrafast and memory-efficient alignment of short DNA sequences to the human genome. *Genome Biol.* 10: R25. <https://doi.org/10.1186/gb-2009-10-3-r25>
- Li, B., and C. N. Dewey, 2011 RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinformatics* 12: 323. <https://doi.org/10.1186/1471-2105-12-323>
- Logan, R. W., R. F. Robledo, J. M. Recla, V. M. Philip, J. A. Bubier *et al.*, 2013 High-precision genetic mapping of behavioral traits in the diversity outbred mouse population. *Genes Brain Behav.* 12: 424–437. <https://doi.org/10.1111/gbb.12029>
- Makowsky, R., N. M. Pajewski, Y. C. Klimentidis, A. I. Vazquez, C. W. Duarte *et al.*, 2011 Beyond missing heritability: prediction of complex traits. *PLoS Genet.* 7: e1002051. <https://doi.org/10.1371/journal.pgen.1002051>
- Mancuso, N., H. Shi, P. Goddard, G. Kichaev, A. Gusev *et al.*, 2017 Integrating gene expression with summary association statistics to identify genes associated with 30 complex traits. *Am. J. Hum. Genet.* 100: 473–487. <https://doi.org/10.1016/j.ajhg.2017.01.031>
- Morgan, A. P., C. P. Fu, C. Y. Kao, C. E. Welsh, J. P. Didion *et al.*, 2016 The mouse universal genotyping array: from substrains to subspecies. *G3 (Bethesda)* 6: 263–279. <https://doi.org/10.1534/g3.115.022087>
- Munger, S. C., N. Raghupathy, K. Choi, A. K. Simons, D. M. Gatti *et al.*, 2014 RNA-Seq alignment to individualized genomes improves transcript abundance estimates in multiparent populations. *Genetics* 198: 59–73. <https://doi.org/10.1534/genetics.114.165886>
- Nica, A. C., S. B. Montgomery, A. S. Dimas, B. E. Stranger, C. Beazley *et al.*, 2010 Candidate causal regulatory effects by integration of expression QTLs with complex trait genetic associations. *PLoS Genet.* 6: e1000895. <https://doi.org/10.1371/journal.pgen.1000895>
- Purcell, S., B. Neale, K. Todd-Brown, L. Thomas, M. A. R. Ferreira *et al.*, 2007 PLINK: a tool set for whole-genome association and population-based linkage analyses. *Am. J. Hum. Genet.* 81: 559–575. <https://doi.org/10.1086/519795>
- R Core Team, 2016 *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna.
- Storey, J. D., and R. Tibshirani, 2003 Statistical significance for genomewide studies. *Proc. Natl. Acad. Sci. USA* 100: 9440–9445. <https://doi.org/10.1073/pnas.1530509100>
- Svenson, K. L., D. M. Gatti, W. Valdar, C. E. Welsh, R. Cheng *et al.*, 2012 High-resolution genetic mapping using the Mouse Diversity outbred population. *Genetics* 190: 437–447. <https://doi.org/10.1534/genetics.111.132597>
- Thompson, B., 1990 Finding a correction for the sampling error in multivariate measures of relationship: a Monte Carlo study. *Educ. Psychol. Meas.* 50: 15–31. <https://doi.org/10.1177/0013164490501003>
- Vervier, K., and J. J. Michaelson, 2016 SLINGER: large-scale learning for predicting gene expression. *Sci. Rep.* 6: 39360. <https://doi.org/10.1038/srep39360>
- Wallace, C., M. Rotival, J. D. Cooper, C. M. Rice, J. H. M. Yang *et al.*, 2012 Statistical colocalization of monocyte gene expression and genetic risk variants for type 1 diabetes. *Hum. Mol. Genet.* 21: 2815–2824. <https://doi.org/10.1093/hmg/ddc098>
- Wen, X., R. Pique-Regi, and F. Luca, 2017 Integrating molecular QTL data into genome-wide genetic association analysis:

- probabilistic assessment of enrichment and colocalization. *PLoS Genet.* 13: e1006646. <https://doi.org/10.1371/journal.pgen.1006646>
- Wilms, I., and C. Croux, 2016 Robust sparse canonical correlation analysis. *BMC Syst. Biol.* 10: 72. <https://doi.org/10.1186/s12918-016-0317-9>
- Witten, D. M., and R. J. Tibshirani, 2009 Extensions of sparse canonical correlation analysis with applications to genomic data. *Stat. Appl. Genet. Mol. Biol.* 8: 1–27. <https://doi.org/10.2202/1544-6115.1470>
- Wray, N. R., J. Yang, B. J. Hayes, A. L. Price, M. E. Goddard *et al.*, 2013 Pitfalls of predicting complex traits from SNPs. *Nat. Rev. Genet.* 14: 507–515. <https://doi.org/10.1038/nrg3457>
- Zhu, Z., F. Zhang, H. Hu, A. Bakshi, M. R. Robinson *et al.*, 2016 Integration of summary data from GWAS and eQTL studies predicts complex trait gene targets. *Nat. Genet.* 48: 481–487. <https://doi.org/10.1038/ng.3538>

Communicating editor: J. Wolf