

SCIENTIFIC REPORTS



OPEN

Powerful and Efficient Strategies for Genetic Association Testing of Symptom and Questionnaire Data in Psychiatric Genetic Studies

Aaron M. Holleman^{1,2}, K. Alaine Broadaway³, Richard Duncan³, Andrei Todor^{2,3}, Lynn M. Almli⁴, Bekh Bradley^{4,5}, Kerry J. Ressler⁶, Debashis Ghosh⁷, Jennifer G. Mulle^{2,3} & Michael P. Epstein^{2,3}

Genetic studies of psychiatric disorders often deal with phenotypes that are not directly measurable. Instead, researchers rely on multivariate symptom data from questionnaires and surveys like the PTSD Symptom Scale (PSS) and Beck Depression Inventory (BDI) to indirectly assess a latent phenotype of interest. Researchers subsequently collapse such multivariate questionnaire data into a univariate outcome to represent a surrogate for the latent phenotype. However, when a causal variant is only associated with a subset of collapsed symptoms, the effect will be challenging to detect using the univariate outcome. We describe a more powerful strategy for genetic association testing in this situation that jointly analyzes the original multivariate symptom data collectively using a statistical framework that compares similarity in multivariate symptom-scale data from questionnaires to similarity in common genetic variants across a gene. We use simulated data to demonstrate this strategy provides substantially increased power over standard approaches that collapse questionnaire data into a single surrogate outcome. We also illustrate our approach using GWAS data from the Grady Trauma Project and identify genes associated with BDI not identified using standard univariate techniques. The approach is computationally efficient, scales to genome-wide studies, and is applicable to correlated symptom data of arbitrary dimension.

Evidence indicates that common genetic variants should explain a sizeable role of the variation in many psychiatric disorders. For example, common variants are estimated to explain 40% of the heritability for bipolar disorder¹, 21% of the heritability of depression², 30–45% of the heritability of post-traumatic stress disorder (PTSD)^{3–6}, and 50% of the heritability of autism spectrum disorder⁷. However, even in studies involving thousands of subjects, identification of specific common trait-influencing polymorphisms remains a challenge. To discover new associations, much attention has been spent on improving genotyping and sequencing technologies to interrogate more genetic variation; however, comparatively less attention has been afforded to thorough characterization of the underlying psychiatric phenotypes that are considered for genetic analysis.

In genetic analyses of a psychiatric phenotype, we often envision our outcome of interest as a single, measurable entity. In practice, we are rarely able to measure the outcome of interest directly and instead attempt to capture the true, latent phenotype via several connected but discrete measurements. As an example, psychiatric genetic studies attempt to account for the heterogeneity of symptoms found in a single psychiatric disorder by measuring the symptoms from several angles via a questionnaire or exam. For example, in studies of post-traumatic stress disorder (PTSD), researchers often measure the outcome using the PTSD Symptom Scale (PSS), which is a

¹Department of Epidemiology, Emory University, Atlanta, GA, USA. ²Center for Computational and Quantitative Genetics, Emory University, Atlanta, GA, USA. ³Department of Human Genetics, Emory University, Atlanta, GA, USA. ⁴Department of Psychiatry and Behavioral Sciences, Emory University, Atlanta, GA, USA. ⁵Clinical Psychologist, Mental Health Service Line, Department of Veterans Affairs Medical Center, Atlanta, GA, USA. ⁶Department of Psychiatry, McLean Hospital, Harvard Medical School, Belmont, MA, USA. ⁷Department of Biostatistics and Informatics, Colorado School of Public Health, Aurora, CO, USA. Aaron M. Holleman and K. Alaine Broadaway contributed equally. Correspondence and requests for materials should be addressed to M.P.E. (email: mpepste@emory.edu)

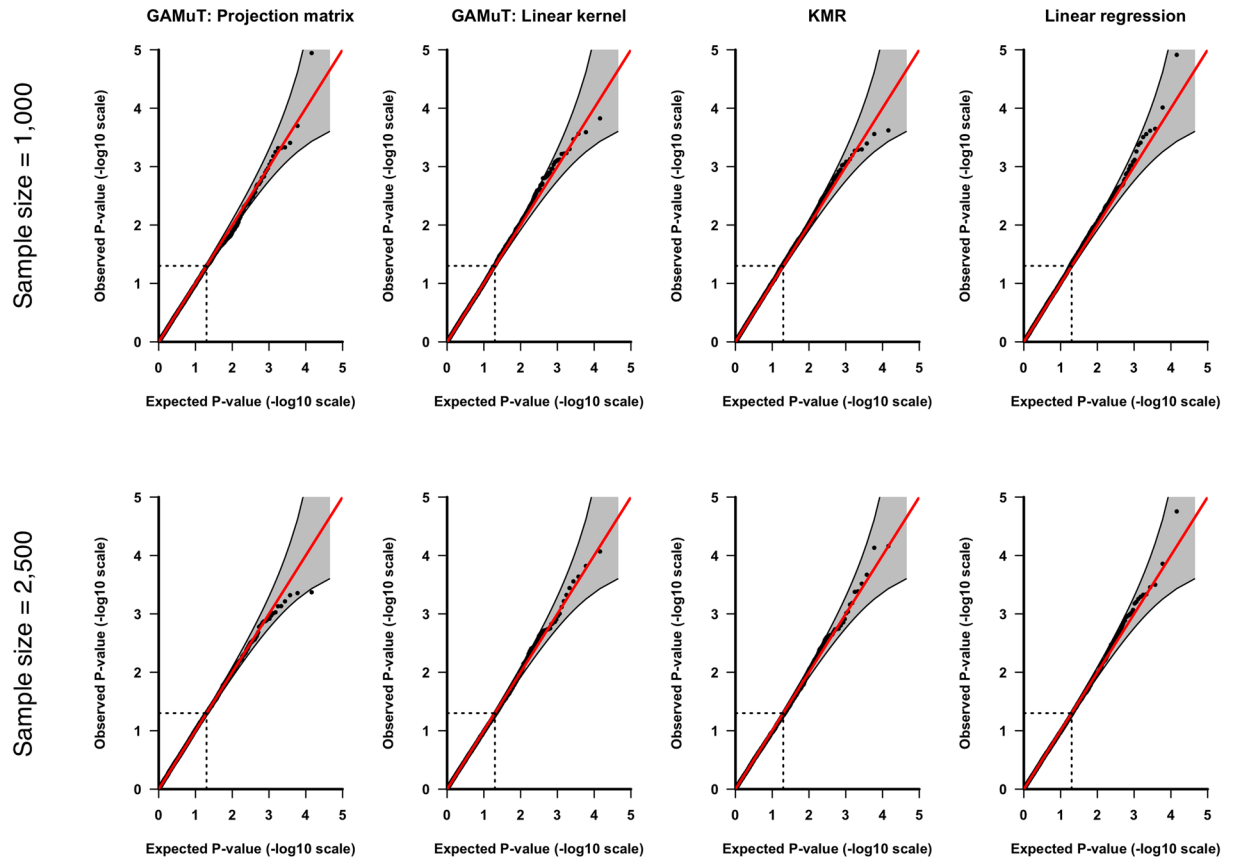


Figure 1. QQ plots applying GAMuT, KMR, and linear regression to 10,000 simulated null data sets assuming sample sizes of 1,000 (top row) or 2,500 (bottom row), for *STAT3*. For each simulation, 21 ordinal questionnaire items were generated for each subject. For KMR and linear regression, the 21 questionnaire items were summed together to yield a single cumulative score. Results are shown applying GAMuT using a projection matrix and a linear kernel for modeling phenotypic similarity. GAMuT and KMR analyses used a weighted linear genotype kernel, with weights based on sample MAF.

17-item questionnaire for assessing and diagnosing PTSD according to the DSM-IV. Each item corresponds to a PTSD symptom and is rated from 0 to 3, with higher scores indicating greater symptom frequency/intensity^{8,9}. Meanwhile, in studies of depression, many studies attempt to measure the phenotype using multiple symptom measurements from the Beck Depression Inventory-II (BDI). The BDI is a 21-item questionnaire with each question developed to correspond to DSM-IV diagnostic criteria for major depressive disorder. The answers to each question are scored from 0 to 3, with higher scores indicating more severe depressive symptoms¹⁰.

Data captured by the PSS, BDI or other questionnaires can actually be considered a collection of interrelated multivariate phenotypes that, in the case of symptom scales, are usually ordinal in nature. The view of a mental disorder as a constellation of multiple correlated symptoms is aligned with the National Institute of Mental Health's (NIMH) Research Domain Criteria (RDoC), which emphasize basic functional dimensions or mechanisms involved in psychopathology (e.g., fear, reward-seeking, attention, perception, arousal) rather than DSM or ICD diagnostic categories¹¹. Nevertheless, practical use of such multivariate symptom data in genetic analysis is complicated by the fact that standard statistical techniques for genetic analysis are generally univariate and designed to handle a single outcome at a time. To improve analytical utility, many questionnaires like the BDI and PSS were designed so that the multivariate symptoms are collapsed into a univariate phenotype for analysis. The simplest and most common collapsing method is unweighted summation of each question's score^{10,12–14} into a univariate cumulative score. The cumulative score can then be treated either as a continuous outcome, or cutoffs can be applied to indicate presence/absence of disease symptoms.

An important issue with applying a univariate cumulative score in genetic analysis is that reducing multivariate information to univariate data nearly always comes at a cost. Carefully defining a phenotype is as vital in a GWAS as reliable genotyping; any association between gene and trait may be diluted by phenotypic heterogeneity. For example, if a gene were associated with a subset of the BDI questionnaire outcomes (e.g., a somatic symptom of depression-like changes in sleep patterns) but not other subsets (e.g., affective symptoms like mood or attitude), the magnitude of the overall effect size of the gene would be attenuated if the two subsets were combined into a univariate outcome measure.

A few key assumptions must be met in order for a univariate cumulative score to sufficiently summarize multivariate ordinal data. As noted by van der Sluis *et al.*^{15–17}, the three primary assumptions that must be met

		Sample Size = 1,000			Sample Size = 2,500		
		$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.001$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.001$
STAT3	GAMuT: Projection Matrix	0.0492	0.0073	0.0010	0.0503	0.0099	0.0008
	GAMuT: Linear Kernel	0.0515	0.0110	0.0013	0.0466	0.0100	0.0009
	KMR	0.0533	0.0098	0.0013	0.0455	0.0108	0.0010
	Linear Regression	0.0541	0.0112	0.0013	0.0479	0.0093	0.0013
LRFN5	GAMuT: Projection Matrix	0.0506	0.0095	0.0012	0.0445	0.0080	0.0008
	GAMuT: Linear Kernel	0.0480	0.0096	0.0014	0.0492	0.0105	0.0010
	KMR	0.0491	0.0102	0.0010	0.0500	0.0107	0.0006
	Linear Regression	0.0508	0.0111	0.0011	0.0533	0.0104	0.0007

Table 1. Empirical Type-I Error Rates. Empirical type-I error rates are presented for GAMuT, univariate KMR and linear regression, considering two genes (*STAT3* on chromosome 17 and *LRFN5* on chromosome 14) and different combinations of sample size and significance (α) level. Error rates are calculated as the proportion of *P*-values less than the specified significance threshold given 10,000 null simulations. All GAMuT and KMR analyses used a weighted linear genotype kernel, with weights based on sample MAF. GAMuT, KMR and linear regression properly control for type-I error across all scenarios tested.

are: (1) the correlation between all questions in the questionnaire must be explained by a single (latent) phenotype; (2) the genetic effect must be on the latent phenotype; (3) the genetic effect—acting through the latent phenotype—must have identical effects on all of the questions in the questionnaire. For applied psychiatric phenotypes, it is more plausible that the assumptions are violated than maintained, a perspective that is supported by NIMH's focus on RDoC. Depressive symptoms identified by the BDI might come from multiple sources (e.g., major depressive disorder, bereavement, post-traumatic stress disorder), violating the first assumption. The causal genetic effect might directly increase somatic symptoms of depression such as changes in appetite and sleep, but not impact mood, violating the second assumption. Alternatively, a variant might in fact affect each trait identified by every question, but have slightly different effect sizes on different questions. If any of these assumptions are not met, association analysis using the cumulative score will result in a substantial loss of power^{15,17–19}.

A few alternatives have been presented to model the complex multivariate data captured within questionnaires. A popular type of approach is a data reduction method like principal component analysis (PCA), which relies on identifying a linear combination of the set of questionnaire responses that maximizes response variance across questions. Once the top few principal components are identified (i.e., those principal components that explain most of the questionnaire variance), association testing is performed between those top principal components and genotype^{20,21}. However, PCA-based strategies that consider only high-variance principal components were recently shown to be generally suboptimal²². As an alternative, van der Sluis *et al.*¹⁷ presented a multivariate gene-based association test by extended Simes procedure (MGAS) that combines the *P*-values obtained from standard, single-SNP association tests for each outcome to produce a single multivariate gene-based *P*-value. However, MGAS relies on permutations to establish significance, which make genome-wide analyses of psychiatric phenotypes cumbersome. Alternatively, Basu *et al.*²³ introduced a rapid multivariate multiple linear regression method (RMMLR), which operates on a MANOVA-based platform. However, while RMMLR establishes significance analytically, it cannot incorporate the important ordinal outcomes commonly measured in questionnaires and surveys.

To allow computationally-efficient and powerful genetic analysis of multivariate symptom data, we show in this paper that we can use a kernel distance-covariance (KDC)^{24–28} method called the Gene Association with Multiple Traits (GAMuT) test²⁹, to assess association between high-dimensional symptom data and multiple variants (common or rare) in a gene. The framework is designed to test whether pairwise similarity in questionnaire responses is independent of pairwise genotypic similarity in a region of interest. The framework allows for an arbitrary number of categorical questions within the questionnaire as well as an arbitrary number of genotypes, thereby permitting gene-based or pathway-based testing of genetic variants. The method allows for covariate adjustment and is a closed-form test that yields analytic *P*-values, thus scaling easily to genome-wide analysis. GAMuT is therefore well-suited to facilitate research that is directly aligned with RDoC's goal of encouraging investigation of biological, cognitive-behavioral, and self-report data using multivariate methods¹¹.

The remainder of this manuscript is organized as follows. We first provide a short overview of the GAMuT method and its features. We then present simulation work to demonstrate that the framework can be considerably more powerful than the standard univariate test based on a cumulative score derived from a questionnaire. We then illustrate the approach using a GWAS study of BDI scores collected as part of the Grady Trauma Project^{30–32}. We finish with concluding remarks and discuss potential extensions to our approach.

Methods

Overview of GAMuT. We provide a brief overview of the GAMuT method²⁹ here and relegate the technical details of the procedure to the Supplementary Methods section online. For a sample of *N* unrelated subjects, GAMuT examines the relationship between a set of *Q* questions (each question assumed to be an ordinal categorical variable with an arbitrary number of levels) and a set of *V* genetic variants within a gene or pathway of interest.

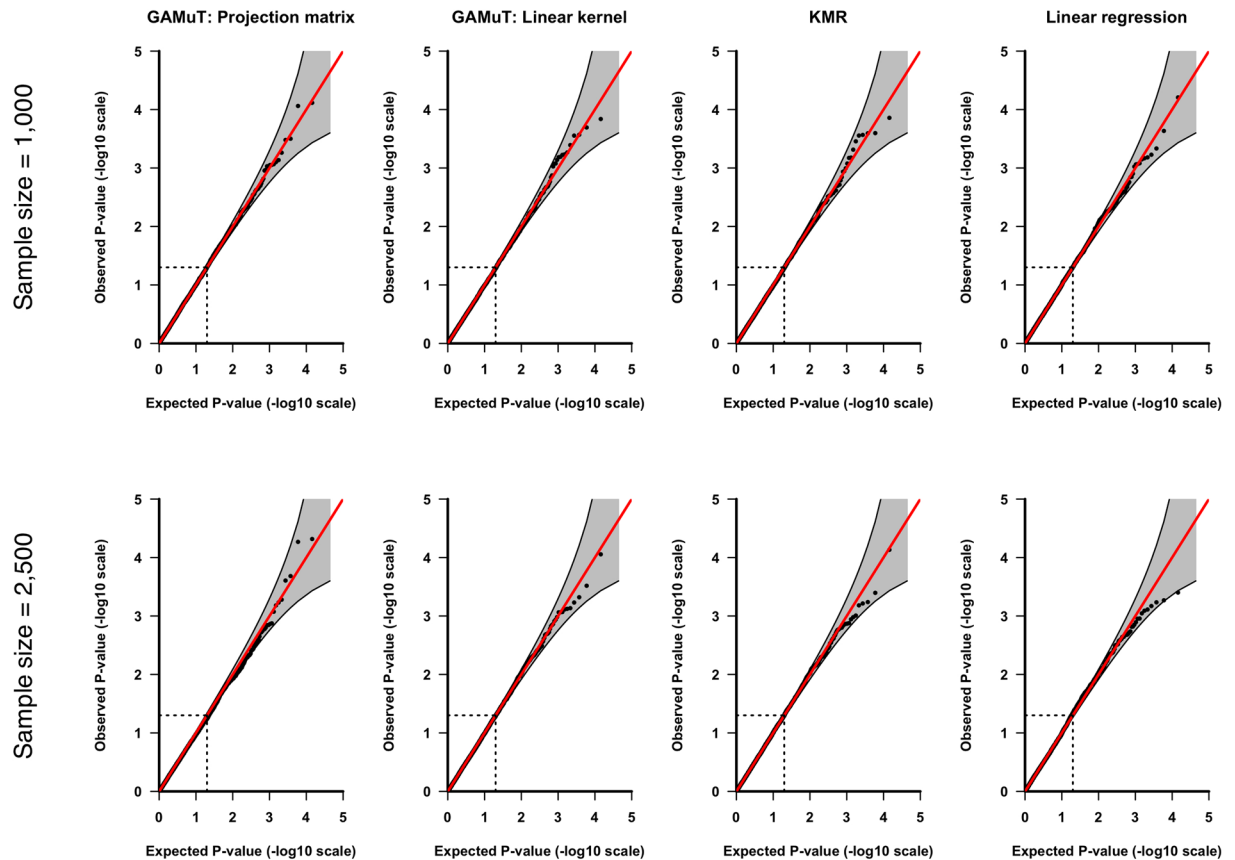


Figure 2. QQ plots applying GAMuT, KMR, and linear regression to 10,000 simulated null data sets assuming sample sizes of 1,000 (top row) or 2,500 (bottom row), for *LRFN5*. For each simulation, 21 ordinal questionnaire items were generated for each subject. For KMR and linear regression, the 21 questionnaire items were summed together to yield a single cumulative score. Results are shown applying GAMuT using a projection matrix and a linear kernel for modeling phenotypic similarity. GAMuT and KMR analyses used a weighted linear genotype kernel, with weights based on sample MAF.

GAMuT is motivated by the idea that, for a pair of individuals, increased genetic similarity at trait-influencing loci across a gene should lead to increased similarity of their questionnaire outcome data. Consequently, GAMuT constructs two different similarity matrices; one similarity matrix for the questionnaire outcomes and the other similarity matrix for the genetic variation within a gene. Each similarity matrix has N rows and N columns with individual elements of the matrix denoting the similarity (phenotypic or genetic) among different pairs of subjects. GAMuT creates a test statistic that evaluates whether the pairwise elements in the similarity matrix of questionnaire outcomes is independent of the pairwise elements in the genetic similarity matrix. The resulting test follows a known asymptotic distribution, which leads to easy and rapid calculation of P -values. GAMuT allows for questionnaire outcomes of arbitrary dimension and can further adjust for covariates.

Simulations. We conducted simulations to verify that GAMuT properly preserves type-I error (i.e., empirical size) and to assess power of GAMuT relative to standard association tests that treat questionnaire responses as a univariate outcome variable resulting from summing the responses into a continuous score. We briefly summarize the simulation design here and provide more comprehensive details in the Supplementary Methods section online. We considered sample sizes of either 1,000 or 2,500 independent subjects. We performed simulations based on SNPs and LD patterns located within 2 kb up- and down-stream from two genes: *signal transducer and activator of transcription 3 (STAT3)*, a gene on chromosome 17q21.31, and *leucine rich repeat and fibronectin type III domain containing 5 (LRFN5)*, a gene on chromosome 14q21.1 recently identified as possibly involved in major depressive disorder³³ (see Supplementary Figs S1 and S2 for the MAF and pairwise LD structure of SNPs in *STAT3* and *LRFN5*). We generated simulated genotypes for all SNPs identified in HapMap within these genes (27 SNPs for *STAT3*; 127 SNPs for *LRFN5*), but applied the testing approaches only to those SNPs that would be typed on standard genotyping arrays (14 SNPs for *STAT3*; 50 SNPs for *LRFN5*).

We simulated multivariate questionnaire data to mimic the BDI questionnaire results obtained from Grady Trauma Project participants. The BDI consists of 21 groups of statements that reflect various symptoms and attitudes associated with depression. Each group includes 4 statements, which correspond to a scale of 0 to 3 in terms of intensity. The BDI is scored by summing the ratings given to each of the 21 items, yielding a cumulative

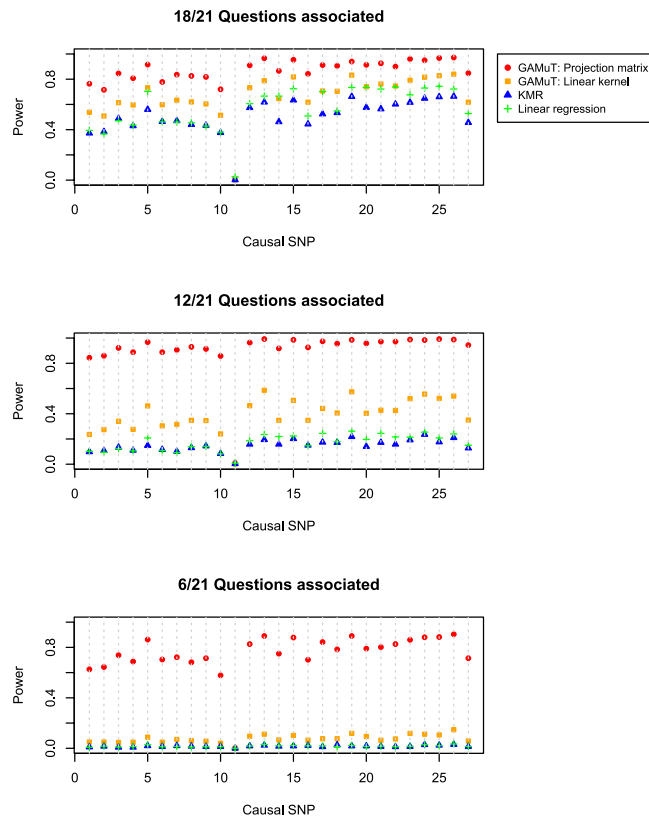


Figure 3. Power for GAMuT, KMR, and linear regression is plotted as a function of causal SNP, for *STAT3*. Top plot assumes the causal SNP is associated with 18 of the 21 BDI questions. Middle plot assumes 12 of 21 questions are associated with causal SNP. Bottom plot assumes only 6 of 21 questions are associated with the causal SNP. Results are shown applying GAMuT using a projection matrix and a linear kernel for modeling phenotypic similarity. GAMuT and KMR analyses used a weighted linear genotype kernel, with weights based on sample MAF. Sample size is 2,500. Power is calculated assuming a significance (α) level of 0.001.

score ranging from 0–63. To mimic the BDI, we generated 21 ordinal responses using the observed distributions and correlations of these responses within the GTP BDI dataset. We show the correlation matrix among ordinal responses in Supplementary Fig. S3 and the distribution of observations for each of the 21 ordinal responses in Supplementary Fig. S4.

For both *STAT3* and *LRFN5*, we applied GAMuT to 10,000 null simulated datasets to estimate empirical size. To investigate the performance of GAMuT under confounding and to assess whether the approach can successfully adjust for relevant covariates in this setting, we also tested empirical size by simulating questions under a confounding model where question responses were independent of genotype, but both questions and genotype were associated with a continuous covariate. For power models, we simulated data sets in which each of the SNPs within the particular gene was modeled as being causal (i.e., each of the 27 SNPs within *STAT3* was modeled as causal, and each of the 127 SNPs within *LRFN5* was set as causal) with effect size of the causal SNP on each question resulting in mean effect sizes with modest effect on the overall cumulative score. We varied the number of questions associated with the causal SNP, considering situations where 18/21, 12/21, and 6/21 questions were actually associated with the causal SNP.

Using the simulated data, we evaluated GAMuT using either projection matrices or linear kernels to model phenotypic similarity and using weighted linear kernels to model genotypic similarity (with weights based on sample MAF). We compared GAMuT to two standard approaches that use the univariate cumulative questionnaire score for inference: standard linear regression and kernel machine regression (KMR)³⁴. Standard linear regression considers individual SNPs for analysis. KMR tests, on the other hand, jointly model multiple SNPs within a gene. KMR can be thought of as a specialized version of GAMuT that considers only one phenotype (the univariate cumulative sum of symptoms/questions) rather than the observed multivariate phenotypes. For KMR, we modeled genotypic similarity in a fashion analogous to GAMuT by using a weighted linear kernel with weights based on sample MAF. Thus, comparison of GAMuT to KMR helps highlight the benefit of considering a multivariate questionnaire phenotype over a traditional cumulative-based score for gene-based analysis.

Analysis of the Grady Trauma Project. Data used in our analyses were collected as part of the Grady Trauma Project (GTP), which investigates the role of genetic risk factors for psychiatric disorders such as PTSD and depression^{32,35}. This study has been carried out according to protocols approved by the IRBs of Emory University School of Medicine and Grady Memorial Hospital. Participants in the GTP are served by the Grady

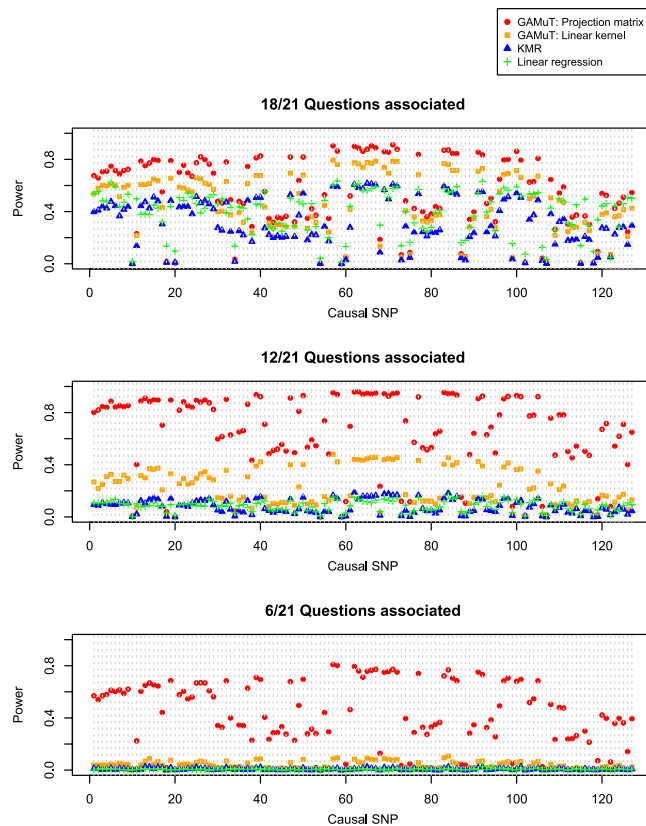


Figure 4. Power for GAMuT, KMR, and linear regression is plotted as a function of causal SNP, for *LRFN5*. Top plot assumes the causal SNP is associated with 18 of the 21 BDI questions. Middle plot assumes 12 of 21 questions are associated with causal SNP. Bottom plot assumes only 6 of 21 questions are associated with the causal SNP. Results are shown applying GAMuT using a projection matrix and a linear kernel for modeling phenotypic similarity. GAMuT and KMR analyses used a weighted linear genotype kernel, with weights based on sample MAF. Sample size is 2,500. Power is calculated assuming a significance (α) level of 0.001.

Hospital in Atlanta, Georgia, and are predominantly urban, African American, and of low socioeconomic status. GTP staff approach subjects in the waiting rooms of Grady Primary Care, Obstetrics and Gynecology, and other clinics, obtaining their written informed consent to participate. In addition to collecting an Oragene salivary sample for DNA extraction, GTP staff conduct an extensive verbal interview, which includes demographic information, a history of stressful life events, and several psychological surveys, including the BDI.

The GTP initially genotyped participants on the Illumina HumanOmni1-Quad array to permit GWAS analyses. Applying standard GWAS quality control filters left 4,607 African-American subjects with good quality genotype data. Further removal of subjects who did not report at least one past trauma, subjects with missing BDI scores, or subjects with incomplete covariate data (age, gender, and the top ten principal components to account for ancestry) yielded a final sample size of 3,520 subjects.

For our sample, we used the support files provided by Illumina to identify 765,580 common genetic variants (MAF > 5%) that fell within 19,609 autosomal genes. We applied GAMuT to the BDI data using a linear kernel to measure pairwise phenotypic similarity in multivariate symptom scores. To measure genetic similarity, we used a linear genotype kernel within GAMuT and performed unweighted analyses as well as weighted analyses, with weights based on variants' MAF or the variants' estimated log odds ratios derived from external and independent GWAS studies of MDD, bipolar disorder, and schizophrenia that are available from the Psychiatric Genomics Consortium^{36–38}. For comparison, we also applied SNP-based linear regression and gene-based univariate KMR to the cumulative BDI score. For KMR, we applied the same genotype weighting schemes as used for the GAMuT analyses.

Results

Type-I error simulations. Fig. 1 shows the quantile-quantile (QQ) plots based on application of GAMuT, KMR, and linear regression to null datasets consisting of 1,000 or 2,500 subjects assayed for 21 BDI questions and SNPs within *STAT3*. We also present empirical type-I error rates for the *STAT3* gene at three nominal significance thresholds in the top half of Table 1. The results show that, for both sample sizes tested, GAMuT properly controls for type-I error, even at the extreme tails of the test. KMR and linear regression, using the cumulative score approach, also demonstrated appropriate empirical size. Null simulations for *LRFN5* likewise produced results indicating that GAMuT, KMR and linear regression all properly control for type-I error (see QQ plots in Fig. 2

Gene	Chr	Number of variants	Genotype weights	GAMuT	KMR	Linear regression (minimum <i>P</i> -value of SNP in gene)
ZHX2	8	97	PGC SZ	2.73×10^{-6}	4.42×10^{-4}	1.00×10^{-3}
		76	PGC MDD	8.59×10^{-6}	1.36×10^{-3}	
FAM43A	3	115	PGC BPD	7.35×10^{-5}	4.15×10^{-4}	2.27×10^{-5}
NUP214	9	19	MAF-based	7.97×10^{-5}	5.16×10^{-3}	5.57×10^{-3}
E2F6	2	29	PGC MDD	9.45×10^{-5}	4.91×10^{-2}	3.32×10^{-2}
GUK1	1	6	PGC SZ	9.54×10^{-5}	5.52×10^{-3}	2.34×10^{-3}
SLC22A5	5	38	PGC SZ	9.69×10^{-5}	3.20×10^{-3}	3.98×10^{-4}

Table 2. Full GAMuT Results for BDI (21 items). Genes with $P < 1 \times 10^{-4}$ identified in the GAMuT analyses are shown. For all analyses, GAMuT employed a linear kernel to model phenotypic similarity, and both GAMuT and KMR utilized a linear genotype kernel (possibly weighted). Of the genes listed, univariate KMR identified none using these same criteria, and standard linear regression identified no SNPs of suggestive significance (suggestive significance threshold for single SNPs: $P < 1 \times 10^{-6}$) within the gene. PGC MDD, PGC BPD, PGC SZ denote weights based on log odds ratios from the Psychiatric Genomics Consortium GWAS of major depressive disorder, bipolar disorder, and schizophrenia, respectively; MAF-based = weights based on minor allele frequencies of variants calculated using the Grady Trauma Project genotype data.

and empirical type-I error rates in the bottom half of Table 1). Supplementary Fig. S5 shows that residualization of questionnaire data prior to GAMuT analysis effectively controls for confounding that, unadjusted, would yield inflated results.

Power simulations. Next we compared the power of GAMuT with univariate KMR and linear regression analyses in a series of simulation studies. For these power simulations, we set sample size to 2,500. For all approaches, power was estimated as the proportion of *P*-values < 0.001 and was evaluated based on 500 replicates of the data per model. To calculate the power of linear regression on the gene-level, we conducted inference on the smallest SNP-level *P*-value in the gene. We adjusted this minimum *P*-value for the multiple correlated SNP-based tests performed in the gene using P_{ACT}^{39} and used the resulting adjusted *P*-value as a gene-level *P*-value.

Figures 3 and 4 show the power results for *STAT3* and *LRFN5*, respectively. We plot power as a function of the causal SNP, where the causal SNPs are ordered by genomic location. The genotyped SNPs (14 for *STAT3* and 50 for *LRFN5*; denoted by 'x' on the bottom of Supplementary Figs S1 and S2) were used to calculate test statistics, but all SNPs (27 for *STAT3* and 127 for *LRFN5*) were treated as causal in turn. Therefore, in situations where the causal SNP is not typed, we rely on correlation of the causal SNP with observed typed SNPs in the gene to gain statistical power. GAMuT offers considerably more power than the two competing univariate methods using cumulative scores for each of the three simulation models considered. For *STAT3*, when approximately half of the questions (12/21) are associated with the causal SNP, both KMR and linear regression observe less than 30% power to detect all 27 causal SNPs; by comparison, GAMuT using a projection matrix to model phenotypic similarity maintains power greater than 80% for 26 of the 27 causal SNPs, while GAMuT using a linear kernel for modeling phenotypic similarity maintains power greater than 50% for 7 of the causal SNPs. When only six questions are associated with the causal SNP, the univariate methods using cumulative scores show nearly zero power for detecting effects across all SNPs, whereas GAMuT employing a projection matrix maintains power greater than 60% for 25 of the 27 causal SNPs.

Similarly, for *LRFN5*, when 12 of 21 questions are associated with the causal SNP, both KMR and linear regression show less than 20% power for detecting all 127 causal SNPs; whereas GAMuT employing a projection matrix observes greater than 80% power to detect 55 of the causal SNPs and greater than 50% power to detect 96 causal SNPs. For the *LRFN5* simulations, when only 6 of the 21 questions have an association with the causal SNP, the univariate methods have almost zero power for detecting all 127 causal SNPs, while GAMuT with a projection matrix maintains greater than 50% power for detecting 57 of the causal SNPs.

We observe a slight drop in power using GAMuT when nearly all of the questions are associated with the causal variant (top row Figs 3 and 4) compared with when approximately half of questions are associated (middle row Figs 3 and 4). This pattern of decreased power when the proportion of associated phenotypes is close to 1 has been observed in other multivariate approaches, including multivariate analysis of variance (MANOVA)^{40,41}. Regardless, our power results demonstrate the benefits of modeling the questionnaire data in a multivariate framework like that employed by GAMuT rather than using a traditional cumulative score.

Application to Grady Trauma Project. We used the GTP dataset to test for associations between the BDI questionnaire and common variants in up to 19,609 genes. Prior to analyses, we controlled for gender, age, and ancestry in the 3,520 unrelated subjects. We applied GAMuT using a linear kernel to measure pairwise phenotypic similarity. We used several approaches for weighting SNPs, including MAF-based weights as well as external weights based on log odds ratio estimates from the PGC GWAS of MDD, bipolar disorder, and schizophrenia. For external weights, we note that not all GTP variants were present within the PGC GWAS results, and therefore the GAMuT analyses utilizing PGC-based genotype weights necessarily included fewer SNPs and corresponding genes than the analyses using MAF-based weights or no weights. Specifically, GAMuT analyses using PGC MDD weights involved 16,716 genes containing 469,582 SNPs. Meanwhile, GAMuT analyses using PGC bipolar disorder weights involved 16,761 genes containing 586,505 SNPs while analyses using PGC schizophrenia weights

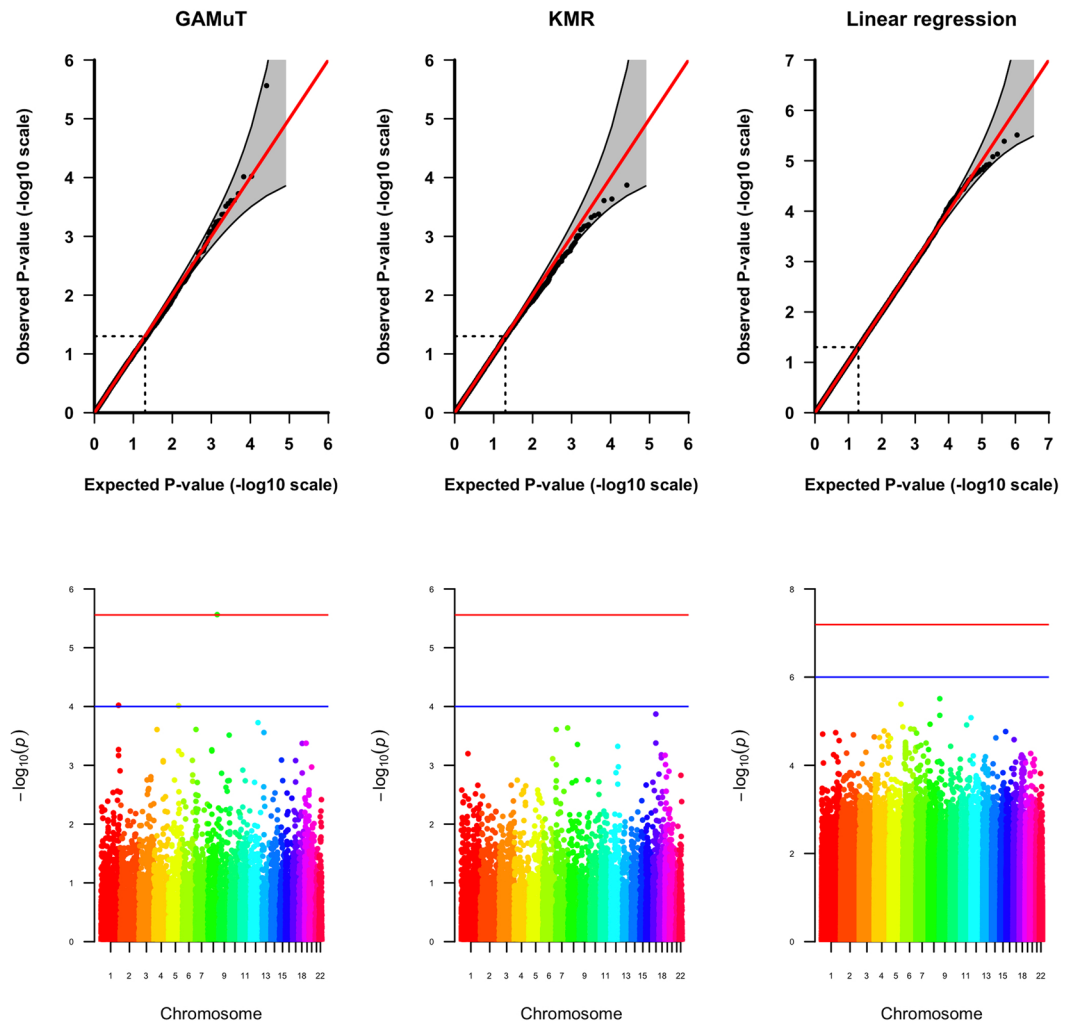


Figure 5. QQ and Manhattan plots for GAMuT, KMR, and linear regression analyses of BDI. The GAMuT analysis used a linear kernel to model phenotypic similarity and genotype weights derived from results of the PGC GWAS for schizophrenia. The KMR analysis also used weights based on the PGC GWAS for schizophrenia. In the Manhattan plots, the red line represents the study-wide significance threshold and the blue line represents the suggestive significance threshold. The study-wide significance thresholds for the GAMuT and KMR analyses are based on a Bonferroni correction for 18,067 genes tested, while the study-wide significance threshold for the linear regression analysis is based on a Bonferroni correction for 775,255 SNPs tested. In the Manhattan plot for the GAMuT results, the point exceeding the study-wide significance threshold is the $-\log_{10}(P\text{-value})$ for *ZHX2*, a gene on chromosome 8. These analyses used a sample of $N = 3,520$.

involved 18,067 genes containing 661,879 SNPs. For comparison with the GAMuT results, we ran univariate KMR using the cumulative BDI. For these KMR analyses, we employed the same genotype weighting schemes as used for GAMuT, and tested the exact same genes as tested in the GAMuT analyses. We also performed standard univariate linear regression of each of 775,255 common variants (SNP-level analyses) on the cumulative BDI score.

Since GAMuT and KMR analyze genes whereas linear regression analyzes SNPs, the multiple-testing adjusted significance thresholds differ between the former tests and the latter test. For each GAMuT and KMR analysis, we set a stringent study-wide significance threshold corresponding to a Bonferroni correction based on the number of genes tested (e.g., $0.05/19,609 = 2.55 \times 10^{-6}$). Thus, the study-wide significance threshold differed depending on the particular genotype weights used, ranging from a threshold of $0.05/16,716 = 2.99 \times 10^{-6}$ for PGC MDD weights to $0.05/19,609 = 2.55 \times 10^{-6}$ for MAF-based weights and no weights. For all GAMuT and KMR analyses we considered P -values less than $P < 1 \times 10^{-4}$ as suggestive. For SNP-based linear regression, we tested 775,255 SNPs across the genome. While we could apply the standard GWAS significance threshold of 5×10^{-8} , we note that this threshold is more conservative than a Bonferroni correction based on the number of SNPs tested. Thus, for linear regression, we instead used a study-wide significance threshold of $0.05/775,255 = 6.45 \times 10^{-8}$, and we considered P -values less than $P < 1 \times 10^{-6}$ as suggestive.

We provide QQ and Manhattan plots for all GAMuT, KMR, and linear regression analyses of BDI in Supplementary Fig. S6. We also list genes identified by GAMuT to be associated with BDI at study-wide or suggestive

significance levels within Table 2. The GAMuT analyses of BDI identified one gene exceeding study-wise significance, while univariate KMR and linear regression of BDI did not detect any study-wise associated genes or SNPs. GAMuT found *ZHX2*, on chromosome 8, to be strongly associated with BDI ($P = 2.73 \times 10^{-6}$), when using genotype weights based on estimated log odds ratios from the PGC GWAS for schizophrenia. We present QQ and Manhattan plots for this particular analysis in the first column of Fig. 5. As noted in Table 2, *ZHX2* was also found to be highly suggestively associated with BDI when employing genotype weights based on the PGC GWAS of MDD ($P = 8.59 \times 10^{-6}$). Previous research suggests a possible link between *ZHX2* and autism spectrum disorder⁴². In comparison with the GAMuT analyses, KMR of cumulative BDI did not identify *ZHX2* as having even suggestive association (Table 2; Fig. 5, middle column; Supplementary Fig. S6b), and univariate linear regression revealed no SNPs suggestively associated with BDI within *ZHX2* or anywhere else across the genome (Table 2; Fig. 5, last column; Supplementary Fig. S6c).

Discussion

As genetic studies of mental-health and psychiatric disorders increasingly shift to the study of high-dimensional symptom and questionnaire data (in greater alignment with NIMH RDoC), it is imperative to employ powerful statistical tests that maximize the possibility of novel genetic discoveries. Here, we have shown that multivariate methods like GAMuT are substantially more powerful for gene mapping of multivariate symptom data than standard methods that typically summarize such symptoms into a single univariate cumulative score for analysis. Methods like GAMuT that jointly model individual questionnaire outcomes are robust to phenotypic heterogeneity, in which a genetic risk factor only affects a subcategory within the questionnaire. In standard cumulative approaches, including KMR and linear regression, phenotypic heterogeneity can dilute the association between gene and trait, making the association extremely difficult to detect. While we focused here on gene-based studies of common variants, we note that our findings are generalizable to studies of rare genetic variation as well as studies of methylation patterns throughout the genome.

We applied GAMuT to the GTP dataset to test for associations between the BDI questionnaire and up to 19,609 genes. After controlling for important covariates, GAMuT found a strong association between BDI and *ZHX2* ($P = 2.73 \times 10^{-6}$), which previous research suggests might be associated with autism spectrum disorder⁴². In comparison, univariate KMR and linear regression did not identify *ZHX2* or SNPs within it to be associated with BDI, at even suggestive levels. This demonstrates through use of real-world data the capacity for multivariate methods like GAMuT to detect genotype-phenotype associations that would be missed using standard cumulative univariate approaches.

GAMuT derives analytic *P*-values based on Davies' exact method, thereby improving computational efficiency and permitting application of the approach on a genome-wide scale. Like the popular KMR framework for univariate analysis, our approach allows for inclusion of prior information, such as biological plausibility of the SNPs under study. We provide R⁴³ software implementing the approach on our website (see Web Resources). Computation run times for GAMuT are primarily dependent on sample size. Using an R script running single-threaded on a 1.7 GHz Intel Core i7 CPU processor, the time required for GAMuT to analyze 10 phenotypes for 1,000, 5,000, and 10,000 subjects is 0.52 seconds/gene, 13.2 seconds/gene, and 68.6 seconds/gene. Thus, genome-wide implementation is feasible particularly when high-performance cluster services are available.

Web resources. Epstein Software: <https://github.com/epstein-software>.
OMIM: <http://www.omim.org>.

Data Availability

The Grady Trauma Project dataset is available from KJR on reasonable request.

References

- Lee, S. H., Wray, N. R., Goddard, M. E. & Visscher, P. M. Estimating missing heritability for disease from genome-wide association studies. *Am J Hum Genet* **88**, 294–305, <https://doi.org/10.1016/j.ajhg.2011.02.002> (2011).
- Lee, S. H. *et al.* Genetic relationship between five psychiatric disorders estimated from genome-wide SNPs. *Nat Genet* **45**, 984–994, <https://doi.org/10.1038/ng.2711> (2013).
- Lyons, M. J. *et al.* Do genes influence exposure to trauma? A twin study of combat. *Am J Med Genet* **48**, 22–27, <https://doi.org/10.1002/ajmg.1320480107> (1993).
- Sartor, C. E. *et al.* Common heritable contributions to low-risk trauma, high-risk trauma, posttraumatic stress disorder, and major depression. *Arch Gen Psychiatry* **69**, 293–299, <https://doi.org/10.1001/archgenpsychiatry.2011.1385> (2012).
- Stein, M. B., Jang, K. L., Taylor, S., Vernon, P. A. & Livesley, W. J. Genetic and environmental influences on trauma exposure and posttraumatic stress disorder symptoms: a twin study. *Am J Psychiatry* **159**, 1675–1681, <https://doi.org/10.1176/appi.ajp.159.10.1675> (2002).
- True, W. R. *et al.* A twin study of genetic and environmental contributions to liability for posttraumatic stress symptoms. *Arch Gen Psychiatry* **50**, 257–264 (1993).
- Gaugler, T. *et al.* Most genetic risk for autism resides with common variation. *Nat Genet* **46**, 881–885, <https://doi.org/10.1038/ng.3039> (2014).
- Foa, E. B., Riggs, D. S., Dancu, C. V. & Rothbaum, B. O. Reliability and validity of a brief instrument for assessing posttraumatic-stress-disorder. *J Trauma Stress* **6**, 459–473, <https://doi.org/10.1007/Bf00974317> (1993).
- Foa, E. B. & Tolin, D. F. Comparison of the PTSD Symptom Scale-Interview Version and the Clinician-Administered PTSD scale. *J Trauma Stress* **13**, 181–191, <https://doi.org/10.1023/A:1007781909213> (2000).
- Beck, A. T., Steer, R. A., Ball, R. & Ranieri, W. Comparison of Beck Depression Inventories -IA and -II in psychiatric outpatients. *J Pers Assess* **67**, 588–597, https://doi.org/10.1207/s15327752jpa6703_13 (1996).
- Clark, L. A., Cuthbert, B., Lewis-Fernandez, R., Narrow, W. E. & Reed, G. M. Three approaches to understanding and classifying mental disorder: ICD-11, DSM-5, and the National Institute of Mental Health's Research Domain Criteria (RDoC). *Psychol Sci Public Interest* **18**, 72–145, <https://doi.org/10.1177/1529100617727266> (2017).
- Wasserman, D., Wasserman, J., Rozanov, V. & Sokolowski, M. Depression in suicidal males: genetic risk variants in the CRHR1 gene. *Genes Brain Behav* **8**, 72–79, <https://doi.org/10.1111/j.1601-183X.2008.00446.x> (2009).

13. McCaffery, J. M. *et al.* Genetic predictors of depressive symptoms in the Look AHEAD trial. *Psychosom Med* **77**, 982–992, <https://doi.org/10.1097/PSY.0000000000000242> (2015).
14. Almlil, L. M. *et al.* ADCYAP1R1 genotype associates with post-traumatic stress symptoms in highly traumatized African-American females. *Am J Med Genet B Neuropsychiatr Genet* **162B**, 262–272, <https://doi.org/10.1002/ajmg.b.32145> (2013).
15. van der Sluis, S., Verhage, M., Posthuma, D. & Dolan, C. V. Phenotypic complexity, measurement bias, and poor phenotypic resolution contribute to the missing heritability problem in genetic association studies. *PLoS One* **5**, e13929, <https://doi.org/10.1371/journal.pone.0013929> (2010).
16. van der Sluis, S., Posthuma, D. & Dolan, C. V. TATES: efficient multivariate genotype-phenotype analysis for genome-wide association studies. *PLoS Genet* **9**, e1003235, <https://doi.org/10.1371/journal.pgen.1003235> (2013).
17. van der Sluis, S. *et al.* MGAS: a powerful tool for multivariate gene-based genome-wide association analysis. *Bioinformatics* **31**, 1007–1015, <https://doi.org/10.1093/bioinformatics/btu783> (2015).
18. Minica, C. C., Boomsma, D. I., van der Sluis, S. & Dolan, C. V. Genetic association in multivariate phenotypic data: power in five models. *Twin Res Hum Genet* **13**, 525–543, <https://doi.org/10.1375/twin.13.6.525> (2010).
19. Medland, S. E. & Neale, M. C. An integrated phenomic approach to multivariate allelic association. *Eur J Hum Genet* **18**, 233–239, <https://doi.org/10.1038/ejhg.2009.133> (2010).
20. He, L. N. *et al.* Genomewide linkage scan for combined obesity phenotypes using principal component analysis. *Ann Hum Genet* **72**, 319–326, <https://doi.org/10.1111/j.1469-1809.2007.00423.x> (2008).
21. Zhang, F. *et al.* Genome-wide pathway association studies of multiple correlated quantitative phenotypes using principle component analyses. *PLoS One* **7**, e53320, <https://doi.org/10.1371/journal.pone.0053320> (2012).
22. Aschard, H. *et al.* Maximizing the power of principal-component analysis of correlated phenotypes in genome-wide association studies. *Am J Hum Genet* **94**, 662–676, <https://doi.org/10.1016/j.ajhg.2014.03.016> (2014).
23. Basu, S. *et al.* A rapid gene-based genome-wide association test with multivariate traits. *Hum Hered* **76**, 53–63, <https://doi.org/10.1159/000356016> (2013).
24. Gretton, A. *et al.* A kernel statistical test of independence. *Advances in Neural Information Processing Systems*, pp. 585–592 (2008).
25. Hua, W. Y. & Ghosh, D. Equivalence of kernel machine regression and kernel distance covariance for multidimensional phenotype association studies. *Biometrics*, <https://doi.org/10.1111/biom.12314> (2015).
26. Szekely, G. J. & Rizzo, M. L. Brownian distance covariance. *Ann Appl Stat* **3**, 1236–1265, <https://doi.org/10.1214/09-Aoas312> (2009).
27. Szekely, G. J., Rizzo, M. L. & Bakirov, N. K. Measuring and testing dependence by correlation of distances. *Ann Stat* **35**, 2769–2794, <https://doi.org/10.1214/009053607000000505> (2007).
28. Zhang, K., Peters, J., Janzing, D. & Schölkopf, B. Kernel-based conditional independence test and application in causal discovery. *arXiv:1202.3775v1 [cs.LG]* (2012).
29. Broadaway, K. A. *et al.* A Statistical Approach for Testing Cross-Phenotype Effects of Rare Variants. *Am J Hum Genet* **98**, 525–540, <https://doi.org/10.1016/j.ajhg.2016.01.017> (2016).
30. Almlil, L. M. *et al.* Correcting systematic inflation in genetic association tests that consider interaction effects: application to a genome-wide association study of posttraumatic stress disorder. *JAMA Psychiatry* **71**, 1392–1399, <https://doi.org/10.1001/jamapsychiatry.2014.1339> (2014).
31. Binder, E. B. *et al.* Association of FKBP5 polymorphisms and childhood abuse with risk of posttraumatic stress disorder symptoms in adults. *JAMA* **299**, 1291–1305, <https://doi.org/10.1001/jama.299.11.1291> (2008).
32. Bradley, R. G. *et al.* Influence of child abuse on adult depression: moderation by the corticotropin-releasing hormone receptor gene. *Arch Gen Psychiatry* **65**, 190–200, <https://doi.org/10.1001/archgenpsychiatry.2007.26> (2008).
33. Wray, N. R. *et al.* Genome-wide association analyses identify 44 risk variants and refine the genetic architecture of major depression. *Nat Genet* **50**, 668–681, <https://doi.org/10.1038/s41588-018-0090-3> (2018).
34. Kwee, L. C., Liu, D., Lin, X., Ghosh, D. & Epstein, M. P. A powerful and flexible multilocus association test for quantitative traits. *Am J Hum Genet* **82**, 386–397, <https://doi.org/10.1016/j.ajhg.2007.10.010> (2008).
35. Ressler, K. J. *et al.* Post-traumatic stress disorder is associated with PACAP and the PAC1 receptor. *Nature* **470**, 492–497, <https://doi.org/10.1038/nature09856> (2011).
36. Major Depressive Disorder Working Group of the Psychiatric GWAS Consortium. A mega-analysis of genome-wide association studies for major depressive disorder. *Mol Psychiatry* **18**, 497–511, <https://doi.org/10.1038/mp.2012.21> (2013).
37. Psychiatric GWAS Consortium Bipolar Disorder Working Group. Large-scale genome-wide association analysis of bipolar disorder identifies a new susceptibility locus near ODZ4. *Nat Genet* **43**, 977–983, <https://doi.org/10.1038/ng.943> (2011).
38. Schizophrenia Working Group of the Psychiatric Genomics Consortium. Biological insights from 108 schizophrenia-associated genetic loci. *Nature* **511**, 421–427, <https://doi.org/10.1038/nature13595> (2014).
39. Conneely, K. N. & Boehnke, M. So many correlated tests, so little time! Rapid adjustment of P values for multiple correlated tests. *Am J Hum Genet* **81**, 1158–1168, <https://doi.org/10.1086/522036> (2007).
40. Cole, D. A., Maxwell, S. E., Arvey, R. & Salas, E. How the power of MANOVA can both increase and decrease as a function of the intercorrelations among dependent variables. *Psychological Bulletin* **115**, 465–474 (1994).
41. Ray, D., Pankow, J. S. & Basu, S. USAT: A Unified Score-Based Association Test for Multiple Phenotype-Genotype Analysis. *Genet Epidemiol* **40**, 20–34, <https://doi.org/10.1002/gepi.21937> (2016).
42. Walker, S. & Scherer, S. W. Identification of candidate intergenic risk loci in autism spectrum disorder. *BMC Genomics* **14**, 499, <https://doi.org/10.1186/1471-2164-14-499> (2013).
43. R Core Team R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/> (2017).

Acknowledgements

This work was supported by NIH grants GM117946, HG007508, MH071537, and AR060893. We thank Drs. David Cutler and Nigel Williams for their comments on a previous version of this manuscript. We appreciate the technical support of all of the staff and volunteers of the Grady Trauma Project. Most importantly, we are extremely indebted to and appreciative of the time and effort given from all of the participants of the Grady Trauma Project.

Author Contributions

K.A.B. and M.P.E. developed the original research concept. K.A.B. and R.D. performed simulation-based analyses for *STAT3*; A.M.H. and A.T. performed simulation-based analyses for *LRFN5*. A.M.H. performed all applied analyses of the Grady Trauma Project. L.M.A., B.B. and K.J.R. provided genetic and questionnaire data from the Grady Trauma Project. A.M.H., K.A.B. and M.P.E. drafted the manuscript and continuously improved the writing with help from L.M.A., B.B., K.J.R., D.G. and J.G.M. All authors reviewed the manuscript.

Additional Information

Supplementary information accompanies this paper at <https://doi.org/10.1038/s41598-019-44046-0>.

Competing Interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019