Published in final edited form as:

Stat Biosci. 2017 June ; 9(1): 50–72. doi:10.1007/s12561-016-9154-z.

Annotation Regression for Genome-Wide Association Studies with an Application to Psychiatric Genomic Consortium Data

Sunyoung Shin¹ and Sündüz Kele¹

¹Department of Statistics, Department of Biostatistics and Medical Informatics, University of Wisconsin-Madison, Madison, USA

Abstract

Although genome-wide association studies (GWAS) have been successful at finding thousands of disease-associated genetic variants (GVs), identifying causal variants and elucidating the mechanisms by which genotypes influence phenotypes are critical open questions. A key challenge is that a large percentage of disease-associated GV are potential regulatory variants located in noncoding regions, making them difficult to interpret. Recent research efforts focus on going beyond annotating GV by integrating functional annotation data with GWAS to prioritize GV. However, applicability of these approaches is challenged by high dimensionality and heterogeneity of functional annotation data. Furthermore, existing methods often assume global associations of GV with annotation data. This strong assumption is susceptible to violations for GV involved in many complex diseases. To address these issues, we develop a general regression framework, named **Annotation Regression for GWAS (ARoG)**. ARoG is based on a finite mixture of linear regressions model where GWAS association measures are viewed as responses and functional annotations as predictors. This mixture framework addresses heterogeneity of effects of GV by grouping them into clusters and high dimensionality of the functional annotations by enabling annotation selection within each cluster. ARoG further employs permutation testing to evaluate the significance of selected annotations. Computational experiments indicate that ARoG can discover distinct associations between disease risk and functional annotations. Application of ARoG to autism and schizophrenia data from Psychiatric Genomics Consortium led to identification of GV that significantly affect interactions of several transcription factors with DNA as potential mechanisms contributing to these disorders.

Keywords

Finite mixture of regressions; Functional genomic data; Genome-wide; association studies; Integrative analysis; Regularized variable selection

[✉]Sündüz Kele, keles@stat.wisc.edu, Sunyoung Shin, shin@stat.wisc.edu.

Electronic supplementary material The online version of this article (doi:10.1007/s12561-016-9154-z) contains supplementary material, which is available to authorized users.

1 Introduction

Although genome-wide association studies (GWAS) have successfully identified thousands of genetic loci associated with human diseases, the design and analysis of these studies are challenged in two critical aspects. First, existing GWAS have revealed that, for many common disorders, the typical genetic architecture encompasses many genetic variants (GVs) with individually small effects on phenotypes [1], indicating the need for larger sample sizes to reliably identify them. Second, the roles of a large proportion of identified GV remain elusive, since they reside in noncoding sequences within introns or in regions between genes. Today, with the availability of affordable whole-genome sequencing, our ability to elucidate the mechanisms by which genotypes influence phenotypes is far behind to identify phenotype-associated variants.

In parallel to the rapid developments in the design and analysis of GWAS, large consortia projects such as the Encyclopedia of DNA Elements (ENCODE [2,3]), the Roadmap Epigenomics Mapping Consortium (REMC [4]), the Genotype-Tissue Expression Project (GTEx [5]), and the International Human Epigenome Consortium (IHEC [6]) as well as many investigator-driven projects are generating diverse data types of RNA transcription (RNA-seq), DNA accessibility (DNase-seq), DNA methylation (Methyl-seq), protein–DNA interactions (ChIP-seq/exo), protein–RNA interactions (CLIP-seq), and chromatin state (Histone ChIP-seq) across diverse cell/tissue types. There is a growing literature on methods for utilizing one or more classes of these functional annotation data to support GWAS results. In contrast to methods requiring individual-level GWAS data which are neither available publicly nor immediately [7,8], many existing methods have a useful key feature of using population-level data in the form of summary statistics of GV [9–12]. However, annotation data have rarely played more than an indirect role in assessing evidence for association in both types of approaches. Specifically, Iversen et al. [7], Pickrell [13], and Thompson et al. [12] used annotations to model prior probabilities of association of GV under Bayesian framework or hierarchical modeling. Chung et al. [9] and Kichaev et al. [11] developed models to integrate binarized functional annotation data and GWAS summary statistics such as p values and z scores of GV. Gagliano et al. [10] correlated annotations with GWAS association status and used Bayes factor for the annotations to estimate the posterior odds of association of GV for prioritization. A significant shortcoming of these methods is that they aim to globally relate the associations of GV to functional annotation data, despite the fact that the same disease mechanism might be governed by distinct functional annotations. For example, disruption of an important pathway may arise by GV in coding regions of the genes and/or in their regulatory mechanisms. Regulatory GV may have a variety of mechanisms such as transcription factor (TF) binding, histone modifications, enhancer activity through chromatin architecture, DNA methylation, and alternative splicing [14]. Another key shortcoming is that they preselect relevant annotations or use conveniently available annotations ignoring disease etiology, because they are not equipped to automatically select important annotations in a data-adaptive manner. Furthermore, several of them can only use annotation data in specific formats (e.g., most recent genetic analysis incorporating pleiotropy and annotation (GPA) [9] requires binary annotation variables).

To overcome these challenges, we develop a regression framework named **Annotation Regression for GWAS (ARoG)** and integrate GWAS and functional annotation data. ARoG models GWAS association measures, e.g., z scores from univariate analysis of GWAS, as a linear function of functional annotations. It employs a mixture of linear regressions framework to accommodate the heterogeneity of associations between GWAS association measures and functional annotations. It aims to capture locally distinct associations that would not be revealed with an analysis that assumes homogeneity of these associations. A critical aspect of ARoG is that it can automatically select relevant annotations among a large number of annotations with penalization techniques. ARoG accommodates categorical or continuous annotation types, both of which are commonly available. The rest of the paper is organized as follows: Section 2 presents empirical observations regarding GWAS association measures and functional annotations using Psychiatric Genomics Consortium (PGC) data. Section 3 develops ARoG and discusses implementation details. In Sect. 4, we analyze PGC autism and schizophrenia data and identify GVs that might influence these diseases with the potential to modulate TF-DNA interactions. Section 5 presents computational experiments with a wide variety of settings including PGC analysis-driven simulations. In Sect. 6, we provide concluding remarks and discuss extensions.

2 Exploring Psychiatric Genomics Consortium Data with Functional Annotations

PGC has conducted mega analysis of GWAS data for five psychiatric disorders: attention deficit/hyperactivity disorder (ADHD), autism spectrum disorder (AUT), bipolar disorder (BIP), major depressive disorder (MDD), and schizophrenia (SCZ) [15,16]. They identified four genome-wide significant loci for BIP [17], and more than 100 genome-wide significant loci for SCZ [18,19]. However, their analysis did not lead to any reproducible genome-wide significant loci for ADHD and MDD, and the analysis on AUT is in progress. Their GWAS summary datasets are publicly available at <http://www.med.unc.edu/pgc/downloads>. In what follows, we focus on AUT and SCZ data.

2.1 Autism GWAS

We generated a set of candidate SNPs by starting with the intersection of SNPs genotyped in all five disorder datasets from the PGC cross-disorder study [16]. After lifting the original genomic coordinates from hg18 to hg19, we obtained 1,219,561 SNPs common to all five disorders. We next selected a subset of the SNPs with a Benjamini-Hochberg (BH)-adjusted association p value smaller than or equal to 0.1 in any of the five disorders [20]. This led to a total of 1430 SNPs. Next, we included 761 linkage disequilibrium (LD) partners of these SNPs as identified by the SNAP tool [21] with an $r^2 \geq 0.8$ to one or more of the 1430 SNPs. As part of preprocessing, we discarded SNPs with more than two alleles, SNPs with nucleotide mismatches between the PGC dataset and the SNP database dbSNP [22], and SNPs not listed in dbSNP [22]. As a result, we obtained a total of 2191 SNPs for analysis. Supplementary Fig. 1a displays the histogram of the autism z scores for these sets of SNPs and illustrates that, as expected, LD partners tend to contribute z scores around zero to the overall distribution, since they had BH-adjusted p values larger than 0.1 in the initial selection step. The Manhattan plot of the p values in Fig. 1a indicates that the SNPs with the

strongest association are on chr 5 (6 of them) and chr 6 (3 of them). All of these have raw p values less than 10^{-6} ; however, they make neither the conventional GWAS p value cut-off of 5×10^{-8} nor the Bonferroni cut-off of 4.1×10^{-8} specific for this study, indicating that common practice for GWAS analysis would not confidently identify significant SNPs from this study.

Currently, most integrative analysis methods consider functional annotations enabled by the large-scale analysis results of consortia projects such as ENCODE (e.g., [9,11,13]). In our exposition, we consider a class of functional annotation which computationally quantifies the effects of SNPs on TF binding. We used atSNP [23] for this quantification and created an annotation score matrix for the 2191 SNPs [23]. Specifically, atSNP computes the likelihood that a given SNP disrupts or enhances the binding sites from a given set of position weight matrices (PWMs) characterizing the class of sequences which TFs recognize [24]. atSNP scans through subsequences overlapping with the SNP position with reference and SNP alleles for the best matches of both to a given PWM and quantifies the significance of the best matches with both alleles by p values. The natural logarithm of the ratio of the two p values is defined as the atSNP annotation score, which empirically reflects the change in the ranks of the PWM matches of the alleles. SNPs likely to enhance or disrupt binding of given TF have large absolute atSNP scores for the corresponding PWM, while SNPs with little potential impact on binding have scores close to zero. We refer to Zuo et al. [23] for further computational details. We considered the JASPAR CORE database [25] for vertebrates with 205 PWMs as our motif library and scored the SNP set. Supplementary Fig. 1b displays the heatmap of the resulting annotation score matrix along with the z scores. Here, most SNPs have relatively weak annotation scores and only SNPs colored as yellow or red are likely to lead to significant changes (as assessed by atSNP p values) in TF binding. As part of our exploratory analysis, we first regressed z scores from each of the five disorders on each annotation score separately. Figure 1b displays the $-\log_{10}$ transformed p values from these marginal regressions. We note that the overall association of the z scores and functional annotations for some diseases are apparent (e.g., MDD). However, for autism, none of the annotations can be deemed as contributing to the variation in the autism z scores based on this global marginal analysis as all BH-adjusted p values are greater than 0.1 (Supplementary Fig. 1c).

2.2 Schizophrenia GWAS

PGC provides analyses of two schizophrenia GWAS: a SCZ study from [18] and a more comprehensive SCZ study from [19]. We refer to the first study as SCZ1 and the second study as SCZ2. SCZ1 and SCZ2 have genotypes for 1,252,901 and 9,444,230 SNPs, respectively. We considered 1,179,262 SNPs common to SCZ1 and SCZ2, filtered out SNPs with BH-adjusted p values larger than 0.01 for both studies, and retained the remaining 8029 SNPs. Similar to the autism analysis, we excluded SNPs with more than two alleles or mismatches of alleles between the PGC datasets and dbSNP, and SNPs that are not in dbSNP. We next extended this set by including their LD partners with $r^2 \geq 0.8$. Our final set of SNPs for the analysis included 11,386 SNPs. z scores of these SNPs have bimodal distributions in both SCZ1 and SCZ2 (Supplementary Fig. 1d). Long tails of the z score distribution of SCZ2 indicate many more statistically significant SNPs from this study. This

may imply increased precision of SCZ2 over SCZ1, which is attributable to the seven-fold increase in the sample size. The Manhattan plot in Fig. 1c indicates that genome-wide significant SNPs from SCZ2 spread throughout the genome. In what follows, we used SCZ2 as the main analysis dataset for both marginal and ARoG analyses and SCZ1 as the validation dataset.

We generated an $11,386 \times 205$ annotation score matrix using atSNP with the JASPAR PWM library. Supplementary Fig. 1e displays the heatmap of the resulting annotation score matrix along with the SCZ1 and SCZ2 z scores and illustrates that only a small proportion of SNPs might impact binding of a small subset of TFs. Marginal regressions of SCZ2 z scores on annotation scores identify 40 annotations which are significant when adjusted for multiple testing by the BH procedure at level 0.1 (Fig. 1d). However, given the large sample size, i.e., the number of SNPs, we view the associations from this marginal analysis as suggestive and turn our attention to developing a framework to identify subgroups of SNPs whose association measures can be explained by a subgroup of functional annotations.

3 A Mixture of Linear Regressions Framework for Incorporating Functional Annotations into GWAS Analysis

ARoG utilizes association measures of SNPs with disease/phenotype as response and potential effects of SNPs on TF binding as predictors as in Sect. 2. It exploits these functional annotations in a regression framework and aims to simultaneously boost detection power of SNPs and select relevant annotations.

3.1 Basic Annotation Regression for GWAS (ARoG(I))

Let $z_i \in \mathbb{R}$, $i = 1, \dots, n$ denote z scores for n SNPs from a GWAS. Let $\mathbf{X} = [\mathbf{X}_i]_{i=1, \dots, n} \in \mathbb{R}^{n \times (p+1)}$ denote an annotation score matrix for the SNPs, where $\mathbf{X}_i = (X_{i0}, \dots, X_{ip}) \in \mathbb{R}^{p+1}$ is a vector of p functional annotations for the i th SNP with the first element of 1 as the intercept term. ARoG assumes that n SNPs are partitioned into K clusters and use a finite mixture of linear regressions model (FMR) to relate the response z_i and the predictor vector $\mathbf{X}_i$ for SNP i , $i = 1, \dots, n$. Basic ARoG, denoted by ARoG(I), performs l_1 -norm penalized maximum likelihood estimation for FMR (FMRLasso) proposed in Städler et al. [26]. Following the notation of FMRLasso, we denote the prior probability of the k th cluster as π_k , its regression parameters as $\boldsymbol{\beta}_k = (\beta_{k0}, \dots, \beta_{kp})^T$, and its variance as σ_k^2 . The conditional density function of z scores, that is z given a functional annotation vector, $\mathbf{X}$ is then

$$f_{\xi}(z|\mathbf{X}) = \sum_{k=1}^K \pi_k \frac{1}{\sqrt{2\pi}\sigma_k} \exp\left(-\frac{(z - \mathbf{X}^T \boldsymbol{\beta}_k)^2}{2\sigma_k^2}\right), \quad (1)$$

where $\xi = (\beta_1, \dots, \beta_K, \sigma_1, \dots, \sigma_K, \pi_1, \dots, \pi_{K-1}) \in \mathbb{R}^{K \cdot (p+1)} \times \mathbb{R}_{>0}^K \times \Pi$, and $\Pi = \{\boldsymbol{\pi} : \pi_k > 0 \text{ for } k = 1, \dots, K-1, \text{ and } \sum_{k=1}^{K-1} \pi_k < 1\}$ is a parameter space of $\boldsymbol{\pi}$ with $\pi_K = 1 - \sum_{k=1}^{K-1} \pi_k$. Here, $\boldsymbol{\beta}_k$ is a cluster-specific regression parameter specifying how z scores relate to the annotation scores within cluster k . Städler et al. [26] considered a reparametrized form of

this density for scale-invariant estimation and efficient computation. Specifically, they reparametrized the regression parameters and the variances as follows:

$$\phi_k = \beta_k / \sigma_k, \quad \rho_k = \sigma_k^{-1}, \quad k=1, \dots, K.$$

We can rewrite Eq. (1) with the new parameters as

$$f_{\theta}(z|X) = \sum_{k=1}^K \pi_k \frac{\rho_k}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(\rho_k z - X^T \phi_k)^2\right),$$

where $\theta = (\phi_1, \dots, \phi_K, \rho_1, \dots, \rho_K, \pi_1, \dots, \pi_{K-1}) \in \mathbb{R}^{K \cdot (p+1)} \times \mathbb{R}_{>0}^K \times \Pi$ is a new parameter vector and Π is the same set as above with $\pi_K = 1 - \sum_{k=1}^{K-1} \pi_k$.

FMRLasso penalizes the negative log-likelihood with an l_1 norm penalty [27]:

$$-\frac{1}{n} l_{\text{pen}, \lambda}(\theta) = -\frac{1}{n} l(\theta) + \lambda \sum_{k=1}^K \pi_k \|\phi_k\|_1 = -\frac{1}{n} \sum_{i=1}^n \log \left(\sum_{k=1}^K \pi_k \frac{\rho_k}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(\rho_k z_i - X_i^T \phi_k)^2\right) \right) + \lambda \sum_{k=1}^K \pi_k \|\phi_k\|_1,$$

where λ is a tuning parameter and $\|\phi_k\|_1 = \sum_{j=1}^p |\phi_{kj}|$. The penalty term weighs the contributions from each cluster by the corresponding prior probability. For a given number of clusters, K , and a given tuning parameter, λ , we define a minimizer of the penalized negative log-likelihood as the FMRLasso estimator, denoted by $\tilde{\theta}_{\lambda, K} = (\tilde{\phi}_1, \dots, \tilde{\phi}_K, \tilde{\rho}_1, \dots, \tilde{\rho}_K, \tilde{\pi}_1, \dots, \tilde{\pi}_{K-1})$. Städler et al. [26] also suggested an unweighted penalty term $\sum_k \|\phi_k\|_1$ and another weighted penalty term of the form $\sum_k \pi_k^{0.5} \|\phi_k\|_1$. The unweighted penalty term tends to perform poorly in unbalanced cases, where the numbers of SNPs across clusters differ significantly [26]. Therefore, ARoG utilizes the weighted penalty term with π_k , which performs well in both the balanced and unbalanced cases.

A key issue in the mixture linear regression models is the selection of the optimal number of clusters and the optimal tuning parameter. We use a modified Bayesian Information Criteria (BIC), defined by Städler et al. [26] as

$$BIC = -2l(\hat{\theta}_{\lambda, K}) + \log(n) d_e,$$

where $d_e = K + (K-1) + \sum_{r=1, \dots, K; j=0, \dots, p} 1_{\{\hat{\phi}_{r,j} \neq 0\}}$ is the effective number of parameters. We perform a grid search over a set of (λ, K) and find the optimal combinations, $(\hat{\lambda}, \hat{K})$, achieving the smallest modified BIC. Städler et al. [26] showed that a single cluster model

selects no variables with $\lambda_{\max} = \max_{j=0, \dots, p} \left| \frac{\langle z, X^j \rangle}{\sqrt{n} \|z\|} \right|$, where z is the response vector and X^j is the $(j+1)$ th column of X , and suggested $\lambda_{\max}$ as the upper bound for the value of the tuning parameter. ARoG increases this upper bound three to six times, since multiple clusters may

require a larger tuning parameter to avoid selecting false positive annotations. With a slight abuse of notation, we denote the ARoG(I) parameter estimates as $\tilde{\boldsymbol{\theta}} \equiv \tilde{\boldsymbol{\theta}}_{\hat{\lambda}, \hat{K}}$. The annotation coefficient and the variance for the k th cluster are estimated by $\tilde{\boldsymbol{\beta}}_k = \tilde{\boldsymbol{\phi}}_k / \tilde{\rho}_k$ and $\tilde{\sigma}_k = 1 / \tilde{\rho}_k$, respectively. The posterior probability that SNP i belongs to the k th cluster is given by

$$\tilde{\gamma}_{ik} \equiv P(k|X_i, z_i, \tilde{\boldsymbol{\theta}}) = \frac{\tilde{\pi}_k \frac{\tilde{\rho}_k}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(\tilde{\rho}_k z_i - X_i^T \tilde{\boldsymbol{\phi}}_k)^2\right)}{\sum_{k=1}^K \tilde{\pi}_k \frac{\tilde{\rho}_k}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(\tilde{\rho}_k z_i - X_i^T \tilde{\boldsymbol{\phi}}_k)^2\right)}.$$

ARoG assigns the SNPs to the clusters for which they have the largest posterior

probabilities. This generates K SNP sets with members $\mathcal{C}_k = \{i \in \{1, \dots, n\} : k = \arg\max_m \tilde{\gamma}_{im}\}$ and set sizes $|\mathcal{C}_k| = n_k$, $k = 1, \dots, K$. We denote the selected annotation set of each cluster as $S_k = \{j \in \{1, \dots, p\} : \tilde{\beta}_{kj} > 0\}$ with numbers of the selected annotations, $|S_k| = p_k$, and define the entire annotations selected by the model as the union of the selected annotations across all clusters, $S = \bigcup_k S_k \subseteq \{1, \dots, p\}$.

3.2 Permutation Testing for ARoG

The ARoG framework follows up the penalized likelihood-based selection with a permutation testing to evaluate the significance of the selected annotations, which typically have small effect sizes based on our data analysis results in Sect. 4. We specifically test whether the maximum absolute value of the coefficients of each functional annotation across all clusters can arise by chance. We randomly permute the z -scores of the SNPs a large number of times (at least 1000 times), and refit ARoG to each permuted dataset. At each fit, we record the maximum absolute value of the estimated coefficients for each annotation across all clusters. This collection generates functional annotation-specific null distributions. Then, the p value for the j th annotation is computed as the proportion of datasets with the maximum absolute value of the estimated coefficients for the annotation larger than

$\max_k |\tilde{\beta}_{kj}|$. We utilize the BH procedure [20] at level 0.1 to account for the multiplicity of the annotations.

3.3 Two-Step Annotation Regression for GWAS (ARoG(II))

Basic ARoG filters false positive annotations with a global penalization across all clusters; however, it is still prone for selecting a nonignorable number of false positive variables as both the simulations of Städler et al. [26] and our computational experiments in Sect. 5 illustrate. To reduce this effect, and thereby increasing specificity, we propose and study two-step ARoG, denoted by ARoG(II). ARoG(II) implements a cluster-level penalization and a refit estimation after the initial global penalization by ARoG(I). The additional cluster-level penalization is similar to relaxed Lasso of Meinshausen [28] which employs another level of Lasso in the context of standard multivariate linear regression model. Both ARoG(II) and relaxed Lasso aim to filter out false positive variables resulting from the initial penalization and whereby lead to better or comparable prediction with more accurate variable selection. Refitting has been widely used as a simple but practical tool to overcome

the biased estimation of Lasso [29]. The refit step aims to improve regression parameter estimation by alleviating shrinkage effects toward zero due to penalization.

ARoG(II) starts with the optimal clusters selected by the initial FMRLasso of ARoG(I) and considers a standard multivariate linear regression model for each cluster with its FMRLasso selected annotation set. Specifically, ARoG(II) adds an l_1 penalty term to the residual sum of squares within each cluster, and thus obtains the following preliminary estimators:

$$\hat{\beta}_{\lambda_k, k}^{S_k} = \underset{\beta_k^{S_k}}{\operatorname{argmin}} \frac{1}{n_k} \sum_{i \in \mathcal{C}_k} \left(z_i - X_i^{S_k T} \beta_k^{S_k} \right)^2 + \lambda_k \|\beta_k^{S_k}\|_1, \quad (2)$$

where $\beta_k^{S_k} \in \mathbb{R}^{p_k+1}$, $X_i^{S_k} = [1; X_{ij; j \in S_k}] \in \mathbb{R}^{p_k+1}$, $k = 1, \dots, K$. The tuning parameter selection for each cluster is through BIC [30] and the coefficients estimated with the optimal tuning parameter are denoted as $\hat{\beta}_k^{S_k} = \hat{\beta}_{\lambda_k, k}^{S_k}$. Next, based on the clusterwise Lasso, we obtain a smaller annotation set, $M_k = \{j \in S_k : \hat{\beta}_{kj}^{S_k} \neq 0\}$, with size $|M_k| = d_k \leq p_k$, and refit a least squares regression with this annotation set:

$$\hat{\beta}_k^{M_k} = \underset{\beta_k^{M_k}}{\operatorname{argmin}} \frac{1}{n_k} \sum_{i \in \mathcal{C}_k} \left(z_i - X_i^{M_k T} \beta_k^{M_k} \right)^2, \quad (3)$$

where $\beta_k^{M_k} \in \mathbb{R}^{d_k+1}$, $X_i^{M_k} = [1; X_{ij; j \in M_k}] \in \mathbb{R}^{d_k+1}$. We then have ARoG(II) annotation score coefficients for the k th cluster as

$$\hat{\beta}_{kj} = \begin{cases} \hat{\beta}_{kj}^{M_k} & j \in M_k \cup 0, \\ 0 & j \in \{1, \dots, p\} \setminus M_k. \end{cases} \quad (4)$$

Similar to ARoG(I), we define ARoG(II) annotations as the union of selected annotations over the clusters, $M = \bigcup_k M_k \subseteq \mathcal{S}$. There is a trade-off between ARoG(I) and ARoG(II), since cluster-level Lasso tends to gain specificity and lose sensitivity with more aggressive annotation screening. The level of the trade-off varies on a case-by-case basis. We further discuss this issue with computational experiments in Sect. 5.1. The permutation testing described in Sect. 3.2 is also part of ARoG(II).

3.4 Numerical Implementation

We implement ARoG with publicly available R packages `fmrlasso` and `glmnet`. The `fmrlasso` package fits FMRLasso with a block coordinate descent generalized expectation–maximization algorithm (BCD-GEM) proposed by Städler et al. [26]. It alternates between an expectation step (E-step) and a generalized maximization step (generalized M-step), which updates the prior probabilities, $\boldsymbol{\pi}$ at once, then updates the reparametrized regression coefficients, $\boldsymbol{\phi}$ and standard deviations, $\boldsymbol{\rho}$. In the M-step, coordinate updates for $\boldsymbol{\phi}$ and $\boldsymbol{\rho}$ are

performed on decoupled K optimization problems for each cluster separately. We use a hierarchical clustering strategy to initialize the BCD-GEM algorithm. This clustering operates on a distance matrix where the distance between any two SNPs is calculated as l_1 distance of their z scores and summary annotation scores, which is defined as the l_2 norm on the annotation score vector. This distance criterion ensures that SNPs with similar z scores and similar variability in the functional annotations are more close to each other. As a result of this clustering, each SNP has initial membership probabilities for the cluster it belongs to and the remaining clusters with a ratio of 9 to 1. For the M-step, we initialize

$\phi_{kj}^{(0)}=0, \rho_k^{(0)}=2, \pi_k^{(0)}=1/K, k=1, \dots, K, \text{ and } j=1 \dots, p$. Finally, we implement the cluster-level Lasso of ARoG(II) with a coordinate descent algorithm using `glmnet`.

4 ARoG Analysis of PGC Data

4.1 PGC Autism GWAS

We fitted ARoG(I) and ARoG(II) to the autism dataset described in Sect. 2.1 and varied the number of clusters as $K=1, \dots, 10$. Table 1 presents the parameter estimates from both ARoGs with $\hat{K}=3$ as the optimal number of clusters. Refitting for ARoG(II) is performed after each SNP is assigned to the cluster for which the SNP has the highest posterior probability based on ARoG(I) to reestimate both the regression parameters and the cluster-specific variances. Both ARoGs have the first and the second clusters as intercept-only models and select FOXL1 and Nkx2-5 TFs for the third cluster. We kept the ARoG(I) intercept estimate for the first cluster, since no SNPs were assigned to this cluster. Estimated coefficients for both TFs of the cluster 3 indicate that the SNP-driven increase in binding affinities for FOXL1 and Nkx2-5 associate with the increase in autism risk in cluster 3. We further support the significance of these associations with a permutation test described in Sect. 3.2 (Supplementary Fig. 2a). The third cluster has a total of 13 SNPs, nine of which constitute the most genome-wide significant SNPs depicted in the Manhattan plot of Fig. 1a. As an alternative multivariate approach to ARoG, we also used ordinary least squares (OLS) and Lasso regression to select the most relevant annotations from the set of 205. OLS did not select any annotations with a BH adjustment on the OLS p values at level 0.1 and had unadjusted p values of 0.332 and 0.014 for FOXL1 and Nkx2-5, respectively. We obtained 11 Lasso-selected annotations including Nkx2-5 and FOXL1 with 5-fold cross-validation to tune the l_1 penalty parameter. However, neither of these survived the permutation testing implemented in a way similar to that of ARoG's (Supplementary Fig. 2a). This analysis suggests that ARoG is indeed exploiting associations detectable only when appropriate subgroups of SNPs are considered.

We investigated the effects of the selected annotations, FOXL1 and Nkx2-5, on autism z scores of cluster 3. Figure 2a highlights the significant positive associations of the z scores with FOXL1 and Nkx2-5 annotation scores, within cluster 3. Considering the whole set of SNPs leads to weak positive associations without statistical support from marginal regressions for both TFs. Figure 2b displays the heatmap of FOXL1 and Nkx2-5 annotation scores of cluster 3 SNPs organized by hierarchical clustering along with their z scores on top and supports that the variation in z scores is well explained by these two annotation scores.

Next, we define AROG-driven candidate causal SNPs, i.e., *ARoG SNPs*, as SNPs leading to significant TF-binding affinity changes and having marginal association with the disorder. atSNP reports p values assessing whether the observed change in TF binding affinity due to SNP is significant. We use these p values along with raw GWAS association p values to refine the SNPs in cluster 3 and create a set of *ARoG SNPs*. For the autism application, we considered a subset of the cluster 3 SNPs with raw GWAS p value of at most 0.005 and atSNP p value of at most 0.01 for FOXL1 or Nkx2-5, resulting in a single *ARoG SNP*, rs17597926. The unadjusted p value of this SNP from autism GWAS is 0.0017 and the resulting atSNP FOXL1 p value is 0.0008. The composite logo plot in Fig. 2c confirms that rs17597926 is creating a potential FOXL1-binding site. rs17597926 is located within the 5th intron of the TCF4 gene, known to interact with helix-loop-helix proteins and regulate neurodevelopment [31]. Furthermore, this SNP has been identified as a *cis*-eQTL for TCF4 in a recent brain expression GWAS [32]. This is an additional support for a potential regulatory role of rs17597926 as a mediator of TCF4 gene in psychiatric disorders.

4.2 Comparison with GPA on PGC Autism GWAS

In addition to the AROG analysis, we also applied the GPA approach of Chung et al. [9] to the autism data. We emphasize that GPA and AROG approaches utilize functional annotations from different angles: GPA goes after global signals using all SNPs genotyped, whereas AROG aims to identify local signals by focusing on a smaller set of signals with potential significance. GPA is based on a joint generative model of association p values of the SNPs and annotation data and identifies annotations that the disease-associated SNPs are enriched for. It aims to simultaneously identify null (SNPs not associated with the phenotype) and nonnull (SNPs associated with the phenotype) and quantify the enrichment of a given annotation within these SNP sets. It specifically tests whether equal proportions of nonnull and null SNPs carry the annotation. Although it can handle multiple annotations simultaneously, our results from two application schemes of “one annotation at a time” and “all annotations simultaneously” showed extreme differences which could potentially be attributable to the violation of the GPA independence assumption of the annotations given the SNP’s null versus nonnull status. As a result, we focus on applying GPA one annotation at a time. Of the 1,219,561 genotyped SNPs in the PGC cross-disorder study, we considered a subset of 1,210,570 SNPs that were also in the dbSNP database. GPA works with binary annotations; therefore, we created a binary annotation score matrix by running atSNP [23] on the SNPs and thresholding atSNP p values at 0.05. We applied GPA to each of the 205 annotations separately and estimated the proportions of null and nonnull SNPs associated with each annotation. Results of GPA hypothesis testing for annotation enrichment did not identify any annotation as significantly enriched for autism-associated SNPs (Supplementary Fig. 3). This is consistent with our marginal analysis in Sect. 2 where none of the annotations exhibited significant marginal associations with the autism z scores. The estimated fold enrichments of FOXL1 and Nkx2-5 in the GPA analysis were 1.003 (s.e. 0.113) and 0.912 (s.e. 0.146), respectively. Both of these levels were too small to be detected with this analysis that considered only two global classes of SNPs (null and nonnull).

4.3 PGC Schizophrenia GWAS

We applied AROG to the SCZ2 dataset described in Sect. 2 with numbers of clusters $K = 1, \dots, 10$. Best BIC values were achieved at $\hat{K} = 6$ and $\hat{K} = 7$ with only a 0.01% difference between the two. We carried out the rest of the analysis with $\hat{K} = 6$.

Neither AROG(I) nor AROG(II) selected any annotations for clusters 1–4. AROG(I) selected FOXL1, Klf4, Prrx2, and NKX3-1 annotations for cluster 5, and Arnt::Ahr, E2F1, FOXL1, Klf4, Foxq1, Prrx2, ARID3A, and E2F4 annotations for cluster 6. Among the selected annotations, the pair of E2F1 and E2F4 and the pair of Prrx2 and ARID3A share similar sequence logos, respectively; thus, both pairs have relatively high correlations of 0.8096 and 0.6234 in the annotation score matrix. AROG(II) retained FOXL1 and Klf4 for cluster 5 and Arnt::Ahr and FOXL1 for cluster 6 (Table 2). Overall, both clusters 5 and 6 are populated with the most genome-wide significant SNPs depicted in the Manhattan plot of Fig. 1c. Permutation testing results for AROG(I) and AROG(II) support significance of the selected annotations with BH adjustment at level 0.1 (Supplementary Fig. 2b). OLS analysis of this dataset selected 15 annotations with unadjusted permutation p values smaller than 0.05; however, none of these survived the multiple testing correction with the BH adjustment at level of 0.1. In contrast, Lasso with 5-fold cross-validation tuning selected 17 annotations, of which only two (Foxq and Zfx annotations) survived the same multiple testing adjustment. The scatter plots of the SCZ2 z scores against the AROG(II) selected annotations exhibit associations in clusters 5 and 6 with a similar global trend across the whole SNP set (Supplementary Fig. 4).

We next created a set of *ARoG SNPs* with Bonferroni-corrected GWAS association p values less than 0.05 and atSNP p values less than 0.01. While the autism dataset suffers from low power, SCZ2 dataset has many SNPs reaching genome-wide significance. Thus, we used the more stringent rule based on Bonferroni cut-off of 0.05 on the SCZ2 GWAS association p values. *ARoG SNPs* included 14 SNPs from cluster 5 and 30 SNPs from cluster 6. Supplementary Table 1 presents the genomic locations, GWAS p values, and RegulomeDB scores [33] of these SNPs. RegulomeDB scores range from 1 for SNPs likely to affect TF binding and be linked to expression of a gene target to 7 for SNPs with no supporting data. For details on RegulomeDB scoring scheme, we refer Table 2 of Boyle et al. [33]. Of the 44 *ARoG SNPs*, 15 have RegulomeDB score of 1, providing evidence for potential importance of these SNPs to schizophrenia. We next compared the SCZ1 association measures (p values) of *ARoG SNPs* to those of other SNP sets one could have identified from the initial set of 11,386 SNPs without using additional functional annotation (Fig. 3a). The other SNP sets one could define are BH0.01 (SNPs defined by BH correction at level 0.01 on the SCZ2 p values), BF0.05 (SNPs defined by Bonferroni correction at level 0.05 on the SCZ2 p values), and BF0.01 (SNPs defined by Bonferroni correction at level 0.01 on the SCZ2 p values). The initial SNP set was added as a baseline. *ARoG SNPs* are on average more significant and reproducible in the SCZ1 than the other SNP sets. Comparison of *ARoG SNPs* with a randomly selected SNP set of the same size from BF0.05 also indicated that *ARoG SNPs* are on average more significant illustrating that the use of the functional annotation information is biasing the selection toward SNPs with reproducible associations. Figure 3b displays the ranking of SCZ1 p values of all 11,386 SNPs and illustrates that most

ARoG SNPs are among the most significant SNPs with respect to SCZ1. Four of these SNPs reach genome-wide significance with Bonferroni adjustment at level 0.05 in the SCZ1 study.

Next, we assessed whether any of the *ARoG* SNPs were among the schizophrenia-associated SNPs from dbGaP [34]. dbGaP harbors 249 SNPs associated with schizophrenia and 42 of these are among the 11,386 SNPs we utilized. Two of the *ARoG* SNPs (rs11191580 and rs10224497, located at chr10:104,906,211 and chr7:2,149,967) are among the dbGaP SNPs. SNP rs11191580 leads to the enhancement of Arnt::Ahr binding, while rs10224497 seems to disrupt Arnt::Ahr binding. Since we observed that enhanced Arnt::Ahr binding overall associated with increased schizophrenia risk in clusters 5 and 6 of the *ARoG* results, we further investigated rs11191580. rs11191580 is located within the 3rd intron of Nt5C2 and has rs732998, located within the 4th intron of Nt5C2, as a perfect LD partner. Their composite logo plots support that these SNPs might indeed enhance the binding of Arnt::Ahr (Fig. 3c, d). Furthermore, the association of rs11191580 is also validated in SCZ1 with p value of 2.23×10^{-8} . Although rs732998 does not quite make the genome-wide significance cut-off, it also exhibits a significant association in SCZ1 with p value of 9.50×10^{-8} . In summary, these two SNPs in perfect LD lead to sequence changes that are likely to improve the binding of the Arnt::Ahr complex. This complex regulates genes in response to the carcinogenic environmental contaminant 2,3,7,8-tetrachlorodibenzo-*p*-dioxin (TCDD). Akahoshi et al. [35] showed that Ahr TF is frequently detected in brain, and overexpression of AhR causes neural differentiation of Neuro2a cells. Furthermore, recent studies support that dioxins and related chemicals influence neural development, and the AhR-signaling pathway might mediate the impact of dioxins on the nervous system [36].

Finally, as we have done for the autism dataset in Sect. 4.2, we applied GPA to SCZ2 dataset using all 1,175,307 SNPs that were in the dbSNP database out of the 1,179,262 genotyped SNPs. This analysis identified nonnull SNPs as significantly depleted for Zfp423 annotation (Supplementary Fig. 5) under Bonferroni-adjusted significance level of 0.1. However, the probability that a SNP is nonnull is estimated as 0.3157; thus, the estimated nonnull set is likely to falsely include a large number of SNPs unassociated with schizophrenia. This implies that the Zfp423 depletion is likely to be a false positive finding.

5 Simulation Studies

We evaluated *ARoG*(I) and *ARoG*(II) with synthetic datasets and PGC data-driven simulated datasets. As alternative methods operating on multiple annotations, we included Lasso regression, and OLS with BH correction on the regression coefficient p values at level of 0.05. Both *ARoG*s enable clustering of SNPs and heterogeneous annotation coefficients across the clusters, while both Lasso regression and OLS assume the homogeneity of the annotation effects. Our simulation-based evaluations focus on the detection of relevant annotations from a large set of annotations with weak effect sizes. GPA under the settings enabled by binarizing the annotation scores behaved unstably and the GPA model fitting sometimes failed. This was mostly due to the severe sparsity of the binary annotations making their estimators be near or on the boundary of the parameter space. Therefore, we did not include GPA as an alternative method in these simulation studies.

We generated 100 simulated datasets under each scenario, where each simulated dataset consisted of training data, validation data, and test data. The validation dataset was used for selection of the optimal tuning parameter, and its sample size was increased 100 times compared to that of the training dataset. The test error was calculated as the negative log-likelihood on the test dataset with the same sample size as the training dataset. We report for each method the test error, numbers of true (TPs) and false positives (FPs), adjusted rand index (ARI) [37], receiver operating characteristic (ROC) curve, and precision-recall curve. TPs and FPs for ARoG are defined by pooling selected annotations across the identified clusters. We use ARI to measure the similarity between the true SNP clusters and estimated SNP clusters. ROC curves and precision-recall curves present the performance of annotation selection in a threshold (cut-off for BH-adjusted p values for OLS, tuning parameters for Lasso and both ARoGs) agnostic manner. In these curves, we plot averages of the true positive rate (TPR), false positive rate (FPR), and precision across the 100 simulation replications.

5.1 Synthetic Data

We generated data from several Gaussian finite mixture regressions model varying the sparsity level of annotation signals as sparse, intermediate, and dense models (Table 3). The columns of the predictor matrix X are generated from an independent standard normal distribution. Supplementary Figs. 6 and 7 and Fig. 4 present the results of these simulations.

The sparse setting is a small n , large p setting; hence, our comparisons only include ARoG(I), ARoG(II), and Lasso. Supplementary Fig. 6a, b shows that ARoG(I) has the smallest test error and both ARoGs have a median ARI of about 0.7. ARoG(I) and ARoG(II) have the same ARI by design, since they share the same clustering assignment, and Lasso has ARI of 0, since it does not perform clustering. Both versions of ARoG have high TPR with the optimal tuning parameter (Supplementary Fig. 6c); however, ARoG(I) has an inflated FPR with an average of 10 more FPs compared to ARoG(II) (Supplementary Fig. 6d). Lasso underselects annotations and on average has four false negatives. ROC and precision-recall curves (Supplementary Fig. 6e, f) indicate that ARoGs outperform Lasso significantly. We remark that the top left corner of the ROC curves for ARoGs roughly corresponds to the results with the optimal tuning parameters presented in the boxplots of TP and FP. We also note that Lasso, ARoG(I), and ARoG(II) do not select annotations in a sequentially augmenting way as tuning parameters decrease; thus, the ROC curves are not monotonically increasing. We also investigated this sparse setting by increasing the sample size to 1000 and observed almost perfect performance by all methods with an area under the ROC curve of 1.

Supplementary Fig. 7 presents the simulation results from the intermediate model of Table 3, where almost half of the regression parameters are set as zero. The test error evaluation, ROC curves, and precision-recall curves clearly indicate that ARoGs outperform OLS and Lasso. Both ARoG(I) and ARoG(II) have the optimal number of TPs, 55, with the optimal tuning parameters (Supplementary Fig. 7c); however, their numbers of FPs are substantially different (Supplementary Fig. 7d). ARoG(II), on average, selects 2 FPs, whereas ARoG(I) selects more than 45 FPs. This emphasizes the significance of the cluster-level Lasso step of

ARoG(II) for reducing the numbers of FPs. ROC curves (Supplementary Fig. 7e) reveal that although ARoG(I) and ARoG(II) overall have comparable performances, ARoG(II) performs marginally better in the top left corner, where both accurately identify all TPs, but only ARoG(II) succeeds in filtering out many FPs.

Figure 4 presents the results for the dense setting. These results agree with the superior performances of both versions of ARoG in the previous settings in terms of prediction error (Figure 4a). This setting also elucidates the contrast between the two ARoGs: ARoG(I) tends to select all annotations, essentially failing to achieve any variable selection, whereas ARoG(II) is able to filter out FPs (Fig. 4c, d). OLS tends to underselect annotations and Lasso tends to select more compared to OLS, however, still partially recovers the TPs. ARoG(II) has the best ROC and precision-recall curve performances (Fig. 4e, f). This setting includes 60 annotations that are not shared between clusters in contrast to the previous settings where the annotations were shared by multiple clusters. The selection performances on these cluster-specific variables are less stable; as a result, the TPRs of both ARoGs heavily fluctuate between 0.4 and 1 in the top left corner of the ROC curves where both have small FPR. Overall, we conclude that the differences between ARoG(I) and ARoG(II) become more pronounced as the sparsity level decreases, and in such dense settings, ARoG(I) tends to have much higher FPR than ARoG(II).

These computational experiments involved completely simulated datasets where the predictor matrix had independent columns and was not designed to have many weak signals. Supplementary Fig. 8 presents the results from a weak signal setting where the actual annotation predictor matrix from the PGC autism GWAS is used to simulate data with the parameters of Supplementary Table 2. This predictor matrix has many more scores close to zero compared to the randomly generated predictor matrix in the above simulations. The overall conclusions from this setting agree well with the sparse setting.

5.2 PGC Analysis-Driven Data

We next evaluated the performance of ARoG in two simulation settings based on the autism and SCZ2 data analyses of Sects. 4.1 and 4.3. The data were simulated based on the actual annotation score matrices from each application and the ARoG(II) parameter estimates of regression slopes, standard deviations, and prior probabilities. These data-driven simulation studies aim to capture the typical signal to noise levels observed in these type of studies. Figure 5 displays the results of the autism simulation setting. Both ARoGs reduce the prediction error compared to OLS by about 8% (Fig. 5a). Both OLS and Lasso tend to miss TPs and thus fail to recover the underlying associations (Fig. 5c). ARoG(I) and ARoG(II) tend to select at least one correct annotation with 86 and 78 times out of 100 repetitions, respectively. Specifically, FOXL1 is selected 71 and 61 times and Nkx2-5 is selected 49 and 42 times by ARoG(I) and ARoG(II), respectively. ARoG(II) on average filters out 2 more FPs than ARoG(I) (Fig. 5d). Based on the ROC and precision-recall curves, ARoG(I) has the best tuning parameter-free performance followed by ARoG(II) (Fig. 5e, f). OLS performs almost the same as random guess with an ROC curve on the 45 degree line.

Supplementary Fig. 9 presents the results for the SCZ2 simulation setting. Both ARoGs perform very well in terms of prediction error (Supplementary Fig. 9a). Similar to the autism

simulations, both OLS and Lasso fail to select any annotations (Supplementary Fig. 9c, d). ARoG(I) and ARoG(II) tend to recover the TPs to some extent by selecting at least one correct annotation except in one simulated dataset with ARoG(II). In this setting, a trade-off between two ARoGs is clear since ARoG(I) seems better at identifying the TPs, namely, Arnt::Ahr, FOXL1, and Klf4, whereas ARoG(II) is able to more aggressively eliminate FPs. ARoG(I) has a median of 11 FPs, with more than 25 FPs in 13 of the simulated datasets. In contrast, ARoG(II) has a median of 2 FPs, with less than 10 FP annotations in almost all simulated datasets. In terms of the ROC and precision-recall curve comparisons (Supplementary Fig. 9e, f), ARoG(I) exhibits a better tuning parameter-free performance compared to ARoG(II). Both OLS and Lasso perform similar to random guesses.

6 Discussion

We presented an integrative framework, named ARoG, for incorporating functional annotation data into GWAS analysis. The key idea behind ARoG is that even when a set of SNPs disrupts a global mechanism, e.g., pathway, that leads to disease, they might be achieving this by disrupting various submechanisms. Some might be disrupting coding sequences, some binding sites of TFs, some methylation profiles or chromatin accessibility. ARoG capitalizes on this idea and aims to identify clusters of SNPs for which GWAS association measures can be explained by a subset of functional annotations. ARoG utilizes FMRLasso [26] which enables selection among large numbers of functional annotations. We illustrated ARoG with an application to PGC data on autism and schizophrenia disorders by utilizing the impact of SNPs on TF-binding affinities as functional annotations. Our analyses led to identification of SNPs which do not necessarily make the genome-wide significance cut-offs, however, are potentially worthy of following up since their GWAS associations are supplemented by their significant effects on TF-binding affinities. This versatile framework provides many directions for useful extensions. First, its annotation selection capability makes it applicable with larger sets of functional annotations including TF ChIP-seq, DNase I accessibility, Histone ChIP-seq, and DNA methylation. Second, we focused our analysis on one disorder at a time; however, ARoG framework can be easily extended to simultaneously consider multiple related GWAS.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported by the National Institutes of Health Grants HG007019, HG003747, and U54AI117924. The authors thank the editor and two referees for their helpful comments.

References

1. Stranger BE, Stahl EA, Raj T. Progress and promise of genome-wide association studies for human complex trait genetics. *Genetics*. 2011; 187(2):367–383. [PubMed: 21115973]
2. Gerstein MB, Kundaje A, Hariharan M, Landt SG, Yan K-K, Cheng C, Mu XJ, Khurana E, Rozowsky J, Alexander R, Min R, Alves P, Abyzov A, Addleman N, Bhardwaj N, Boyle AP, Cayting P, Charos A, Chen DZ, Cheng Y, Clarke D, Eastman C, Euskirchen G, Fritze S, Fu Y,

Gertz J, Grubert F, Harmanci A, Jain P, Kasowski M, Lacroute P, Leng J, Lian J, Monahan H, O'Geen H, Ouyang Z, Partridge EC, Patacsil D, Pauli F, Raha D, Ramirez L, Reddy TE, Reed B, Shi M, Slifer T, Wang J, Wu L, Yang X, Yip KY, Zilberman-Schapira G, Batzoglu S, Sidow A, Farnham PJ, Myers RM, Weissman SM, Snyder M. Architecture of the human regulatory network derived from ENCODE data. *Nature*. 2012; 489(7414):91–100. [PubMed: 22955619]

3. Yue F, Cheng Y, Breschi A, Vierstra J, Wu W, Ryba T, Sandstrom R, Ma Z, Davis C, Pope BD, Shen Y, Pervouchine DD, Djebali S, Thurman RE, Kaul R, Rynes E, Kirilusha A, Marinov GK, Williams BA, Trout D, Amrhein H, Fisher-Aylor K, Antoshechkin I, DeSalvo G, See L-H, Fastuca M, Drenkow J, Zaleski C, Dobin A, Prieto P, Lagarde J, Bussotti G, Tanzer A, Denas O, Li K, Bender MA, Zhang M, Byron R, Groudine MT, McCleary D, Pham L, Ye Z, Kuan S, Edsall L, Wu Y-C, Rasmussen MD, Bansal MS, Kellis M, Keller CA, Morrissey CS, Mishra T, Jain D, Dogan N, Harris RS, Cayting P, Kawli T, Boyle AP, Euskirchen G, Kundaje A, Lin S, Lin Y, Jansen C, Malladi VS, Cline MS, Erickson DT, Kirkup VM, Learned K, Sloan CA, Rosenbloom KR, Lacerda de Sousa B, Beal K, Pignatelli M, Flicek P, Lian J, Kahveci T, Lee D, Kent JW, Ramalho Santos M, Herrero J, Notredame C, Johnson A, Vong S, Lee K, Bates D, Neri F, Diegel M, Canfield T, Sabo PJ, Wilken MS, Reh TA, Giste E, Shafer A, Kutayin T, Haugen E, Dunn D, Reynolds AP, Neph S, Humbert R, Hansen RS, De Bruijn M, Selleri L, Rudensky A, Josefowicz S, Samstein R, Eichler EE, Orkin SH, Levasseur D, Papayannopoulou T, Chang K-H, Skoultschi A, Gosh S, Disteche C, Treuting P, Wang Y, Weiss MJ, Blobel GA, Cao X, Zhong S, Wang T, Good PJ, Lowdon RF, Adams LB, Zhou X-Q, Pazin MJ, Feingold EA, Wold B, Taylor J, Mortazavi A, Weissman SM, Stamatoyannopoulos JA, Snyder MP, Guigo R, Gingeras TR, Gilbert DM, Hardison RC, Beer MA, Ren B. The Mouse ENCODE Consortium. A comparative encyclopedia of DNA elements in the mouse genome. *Nature*. 2014; 515(7527):355–364. <http://dx.doi.org/10.1038/nature13992>. [PubMed: 25409824]
4. Roadmap Epigenomics Consortium. Integrative analysis of 111 reference human epigenomes. *Nature*. 2015; 518(7539):317–330. <http://view.ncbi.nlm.nih.gov/pubmed/25693563>. [PubMed: 25693563]
5. The GTEx Consortium. The genotype-tissue expression (GTEx) pilot analysis: multi-tissue gene regulation in humans. *Science*. 2015; 348(6235):648–660. [PubMed: 25954001]
6. International Human Epigenome Consortium. <http://ihc-epigenomes.org/research/projects/>.
7. Iversen ES, Lipton G, Clyde MA, Monteiro AN. Functional annotation signatures of disease susceptibility loci improve SNP association analysis. *BMC Genom*. 2014; 15:398.
8. Wasserman WW, Long N, Dickson SP, Maia JM, Kim HS, Zhu Q, Allen AS. Leveraging prior information to detect causal variants via multi-variant regression. *PLoS Comput Biol*. 2013; 9(6):e1003093. [PubMed: 23762022]
9. Chung D, Yang C, Li C, Gelernter J, Zhao H. GPA: a statistical approach to prioritizing GWAS results by integrating pleiotropy and annotation. *PLoS Genet*. 2014; 10(11):e1004787. [PubMed: 25393678]
10. Gagliano SA, Barnes MR, Weale ME, Knight J. A Bayesian method to incorporate hundreds of functional characteristics with association evidence to improve variant prioritization. *PLoS ONE*. 2014; 9(5):e98122. [PubMed: 24844982]
11. Kichaev G, Yang WY, Lindstrom S, Hormozdiari F, Eskin E, Price AL, Kraft P, Pasaniuc B. Integrating functional data to prioritize causal variants in statistical fine-mapping studies. *PLoS Genet*. 2014; 10(10):e1004722. [PubMed: 25357204]
12. Thompson JR, Gögele M, Weichenberger CX, Modenese M, Attia J, Barrett JH, Boehnke M, DeGrandi A, Domingues FS, Hicks AA, Marroni F, Pattaro C, Ruggeri F, Borsani G, Casari G, Parmigiani G, Pastore A, Pfeufer A, Schwienbacher C, Taliun D, CKD Gen Consortium. Fox CS, Pramstaller PP, Minelli C. SNP prioritization using a Bayesian probability of association. *Genet Epidemiol*. 2013; 37(2):214–221. [PubMed: 23280596]
13. Pickrell JK. Joint analysis of functional genomic data and genome-wide association studies of 18 human traits. *Am J Hum Genet*. 2014; 94(4):559–573. [PubMed: 24702953]
14. Pai AA, Pritchard JK, Gilad Y. The genetic and mechanistic basis for variation in gene regulation. *PLoS Genet*. 2015; 11(1):e1004857. [PubMed: 25569255]
15. Psychiatric Genomics Consortium. <http://www.med.unc.edu/pgc>.

16. Cross-Disorder Group of the Psychiatric Genomics Consortium. Identification of risk loci with shared effects on five major psychiatric disorders: a genome-wide analysis. *Lancet*. 2013; 381(9875):1371–1379. [PubMed: 23453885]
17. 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*. 2011; 43(10):977–983. [PubMed: 21926972]
18. Schizophrenia Working Group of the Psychiatric Genomics Consortium. Genome-wide association study identifies five new schizophrenia loci. *Nat Genet*. 2011; 43(10):969–976. [PubMed: 21926974]
19. Schizophrenia Working Group of the Psychiatric Genomics Consortium. Biological insights from 108 schizophrenia-associated genetic loci. *Nature*. 2014; 511:421–427. [PubMed: 25056061]
20. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc B*. 1995; 57:289–300.
21. Johnson AD, Handsaker RE, Pulit SL, Nizzari MM, O'Donnell CJ, de Bakker PIW. SNAP: a web-based tool for identification and annotation of proxy SNPs using HapMap. *Bioinformatics*. 2008; 24(24):2938–2939. [PubMed: 18974171]
22. dbSNP: Short Genetic Variations. <http://www.ncbi.nlm.nih.gov/SNP/>
23. Zuo C, Shin S, Kele S. atSNP: transcription factor binding affinity testing for regulatory SNP detection. *Bioinformatics*. 2015; 31(20):3353–3355. [PubMed: 26092860]
24. Stormo GD, Shneider TD, Gold L, Ehrenfeucht A. Use of 'perceptron' algorithm to distinguish translational initiation sites in *E. coli*. *Nucleic Acids Res*. 1982; 10(9):2997–3010. [PubMed: 7048259]
25. Mathelier A, Zhao X, Zhang AW, Parcy F, Worsley-Hunt R, Arenillas DJ, Buchman S, Chen C, Chou A, Ienasescu H, Lim J, Shyr C, Tan G, Zhou M, Lenhard B, Sandelin A, Wasserman WW. JASPAR 2014: an extensively expanded and updated open-access database of transcription factor binding profiles. *Nucleic Acids Res*. 2014; 42(D1):D142–D147. [PubMed: 24194598]
26. Städler N, Bühlmann P, van de Geer S. l_1 -penalization for mixture regression models. *TEST*. 2010; 19(2):209–256.
27. Tibshirani R. Regression selection and shrinkage via the Lasso. *J R Stat Soc B*. 1994; 58:267–288. <http://citeseer.ist.psu.edu/viewdoc/summary?doi=10.1.1.35.7574>.
28. Meinshausen N. Relaxed Lasso. *Comput Stat Data Anal*. 2007; 52(1):374–393.
29. Candès E, Tao T. The Dantzig selector: statistical estimation when p is much larger than n . *Ann Stat*. 2007; 35(6):2313–2351.
30. Schwarz G. Estimating the dimension of a model. *Ann Stat*. 1978; 6(2):461–464.
31. Forrest MP, Hill MJ, Quantock AJ, Martin-Rendon E, Blake DJ. The emerging roles of TCF4 in disease and development. *Trends Mol Med*. 2014; 20(6):322–331. [PubMed: 24594265]
32. Zou F, Chai HS, Younkin CS, Allen M, Crook J, Pankratz VS, Carrasquillo MM, Rowley CN, Nair AA, Middha S, Maharjan S, Nguyen T, Ma L, Malphrus KG, Palusak R, Lincoln S, Bisceglia G, Georgescu C, Kouri N, Kolbert CP, Jen J, Haines JL, Mayeux R, Pericak-Vance MA, Farrer LA, Schellenberg GD, Petersen RC, Graff-Radford NR, Dickson DW, Younkin SG, Ertekin-Taner N. Brain expression genome-wide association study (eGWAS) identifies human disease-associated variants. *PLoS Genet*. 2012; 8(6):e1002707. [PubMed: 22685416]
33. Boyle AP, Hong EL, Hariharan M, Cheng Y, Schaub MA, Kasowski M, Karczewski KJ, Park J, Hitz BC, Weng S, Cherry JM, Snyder M. Annotation of functional variation in personal genomes using regulomedb. *Genome Res*. 2012; 22(9):1790–1797. [PubMed: 22955989]
34. dbGaP: The Database of Genotypes and Phenotypes. <http://www.ncbi.nlm.nih.gov/gap>
35. Akahoshi E, Yoshimura S, Ishihara-Sugano M. Over-expression of AhR(aryl hydrocarbon receptor) induces neural differentiation of Neuro2a cells: neurotoxicology study. *Environ Health*. 2006; 5(1):24. [PubMed: 16956419]
36. Xie HQ, Xu H-M, Fu H-L, Hu Q, Tian W-J, Pei X-H, Zhao B. AhR-mediated effects of dioxin on neuronal acetylcholinesterase expression in vitro. *Environ Health Perspect*. 2013; 121(5):613–618. [PubMed: 23426015]
37. Rand WM. Objective criteria for the evaluation of clustering methods. *J Am Stat Assoc*. 1971; 66(336):846–850.

Fig. 1.

a Manhattan plot for autism association p values across all 2191 SNPs. *ARoG SNP* identified in Sect. 4.1 is marked with a *purple diamond*. *Blue dashed* and *red dotted horizontal lines* depict the Bonferroni cut-off at a significance level of 0.05 and the conventional p value cut-off of 5×10^{-8} . **b** Heatmap of $-\log_{10} p$ values from marginal regressions of z scores for five psychiatric disorders on annotation scores. **c** Manhattan plot for SCZ2 association p values across all 11,386 SNPs. *ARoG SNPs* identified in Sect. 4.3 are marked with *purple diamonds*. *Blue dashed* and *red dotted horizontal lines* depict the

Bonferroni cut-off at a significance level of 0.05 and the conventional p value cut-off of 5×10^{-8} . **d** Ranking of annotations based on the marginal regressions of the SCZ2 z scores on annotation scores for each TF. FOXL1, Klf4, and Arnt::Ahr TFs that are identified as associated with the z scores in Sect. 4.3 are labeled. The *dashed vertical line* depicts the BH cut-off at a significance level of 0.1 (Color figure online)

Fig. 2.

a Autism z scores versus annotation scores for FOXL1 and Nkx2-5 TFs selected for cluster 3 along with the marginal linear regression line fit and ARoG estimates (intercept and slopes for FOXL1 and Nkx2-5). **b** Hierarchical clustering of SNPs in cluster 3 based on ARoG selected annotations with their AUT z scores. **c** Composite sequence logo of SNP rs17597926 with the FOXL1 PWM. The *middle two rows* represent best matching genomic subsequences to the FOXL1 PWM with the reference and SNP rs17597926 alleles, respectively. The *dashed boxes* mark the SNP location. *Top and bottom rows* display FOXL1

PWM sequence logos aligned to the best reference and SNP allele matches (Color figure online)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

(a)

(b)

(c)

(d)

Fig. 3.

a SCZ1 p values for multiple SNP sets generated based on the SCZ2 data: The SNP sets include ARoG: *ARoG* SNPs; All: All 11, 386 SNPs; BH0.01: SNPs with SCZ2 BH-adjusted p values less than or equal to 0.01; BF0.05/BF0.01: SNPs with SCZ2 Bonferroni-adjusted p values less than or equal to 0.05/0.01. **b** SNPs ranked based on their SCZ1 significance levels. *ARoG* SNPs are marked with asterisks. The vertical dashed line depicts Bonferroni cut-off of SCZ1 analysis under significance level of 0.05. **c** Composite sequence logo of rs11191580 with the Arnt::Ahr PWM: the SNP enhances the binding of Arnt::Ahr. **d** Composite sequence logo of SNP rs732998 with the Arnt::Ahr PWM: the SNP also enhances the binding of Arnt::Ahr

Fig. 4.
Simulation results for the dense model of Table 3 (Color figure online)

Fig. 5.
Simulation results for autism data analysis-driven setting (Color figure online)

Table 1

ARoG parameter estimates with PGC autism data

Cluster	1	2	3
Estimated prior prob. ($\tilde{\pi}_k$)	0.0929	0.8724	0.0347
ARoG(I)			
SSD ($\hat{\sigma}_k$)	0.1739	1.0153	2.3258
(Intercept)	-0.0686	0.0253	-0.0580
FOXL1	0	0	0.2639
Nkx2-5	0	0	0.1171
ARoG(II)			
SSD ($\hat{\sigma}_k$)	0.1739	0.9926	3.7191
(Intercept)	-0.0686	0.0183	-0.5724
FOXL1	0	0	1.2075
Nkx2-5	0	0	1.3215

Table 2

ARoG parameter estimates with PGC SCZ2 data

Cluster	1	2	3	4	5	6
Estimated prior prob. ($\tilde{\pi}_k$)	0.2335	0.2173	0.2202	0.2084	0.0615	0.0591
ARoG(I)						
SSD ($\hat{\sigma}_k$)	0.4122	0.3825	0.8805	0.9001	1.7423	1.8453
(Intercept)	4.1206	-4.1117	-4.5537	4.6279	-5.7453	5.8446
Amt::Ahr	0	0	0	0	0	0.0859
FOXL1	0	0	0	0	-0.2335	-0.1394
Klf4	0	0	0	0	0.1027	0.0653
ARoG(II)						
SSD ($\hat{\sigma}_k$)	0.3336	0.2941	0.9605	0.9427	1.5425	1.6721
(Intercept)	4.0791	-4.0776	-4.8039	4.9482	-7.0816	7.1786
Amt::Ahr	0	0	0	0	0	0.1814
FOXL1	0	0	0	0	-0.2879	-0.4295
Klf4	0	0	0	0	0.2654	0

Table 3

Simulation settings for Sect. 5.1

		β	σ	π
Sparse				
n		100		
Cluster	1	$(0, \underbrace{3, \dots, 3}_{5 \text{ repetitions}}, \underbrace{0, \dots, 0}_{100 \text{ repetitions}})$	0.5	0.5
	2	$(0, \underbrace{-1, \dots, -1}_{5 \text{ repetitions}}, \underbrace{0, \dots, 0}_{100 \text{ repetitions}})$	0.5	0.5
Intermediate				
n		1000		
Cluster	1	$(0, \underbrace{1.5, \dots, 1.5}_{55 \text{ repetitions}}, \underbrace{0, \dots, 0}_{50 \text{ repetitions}})$	1	0.5
	2	$(0, \underbrace{-0.5, \dots, -0.5}_{55 \text{ repetitions}}, \underbrace{0, \dots, 0}_{50 \text{ repetitions}})$	1	0.5
Dense				
n		1000		
Cluster	1	$(0, \underbrace{-1, \dots, -1}_{30 \text{ repetitions}}, \underbrace{1, \dots, 1}_{30 \text{ repetitions}}, \underbrace{0, \dots, 0}_{45 \text{ repetitions}})$	0.5	0.5
	2	$(0, \underbrace{-1, \dots, -1}_{30 \text{ repetitions}}, \underbrace{0, \dots, 0}_{30 \text{ repetitions}}, \underbrace{1, \dots, 1}_{30 \text{ repetitions}}, \underbrace{0, \dots, 0}_{15 \text{ repetitions}})$	0.5	0.5