Supplemental Method 1: Alzheimer's MR score

Our goal is to develop a score to quantify the magnitude of causal evidence from Mendelian Randomization which supports a nonzero causal effect when multiple, potentially non-independent, causal effect estimates are made. Such a score is intended to be largest for genes with the strongest causal effects and for which we have the highest degree of confidence that they are so. We additionally want the score to capture evidence aggregated across different biologically related gene phenotypes (e.g., gene expression, protein abundance).

For a single gene let $\hat{\theta}_{sl_sf}$ represent the estimated marginal causal effect of the gene phenotype s on Alzheimer's disease risk using gene phenotype cohort l_s and AD GWAS cohort f, for $l_s = 1, \dots, L_s$, $f = 1, \dots, F$, $s = 1, \dots, S$. We seek some linear function of $\hat{\boldsymbol{\theta}} = (\hat{\theta}_{sl_sf})$ that can be used to measure the degree to which the evidence collectively suggests that $\boldsymbol{\theta} \neq \mathbf{0}$ and that one gene may have stronger evidence than another. This may be considered from an inferential point of view as estimating the magnitude and number of nonzero elements in $\boldsymbol{\theta}$, but here we will not provide such a test. A standard choice for a score of this type is a weighted sum of the chi-square statistics corresponding to the tests of $H_0^{sl_sf}: \theta_{sl_sf} = 0$. One potential challenge with this score is that the l_s cohorts are not guaranteed to contain different subjects, producing non-zero correlation between the summands of the score. This is potentially problematic because we may on average expect to observe *larger* scores from causal estimates made using overlapping GWAS subjects compared to those made using independent cohorts. In general, the score we seek is only intended to be used to compare evidence between genes and prioritize those with larger scores. If the score for each gene is equally likely to be inflated to the same extent on average, then this problem will have no impact on our prioritization of one gene over another using this score. We now derive the score as a sum of Z-statistics, to retain the sign of effect, and show its properties.

Signed score

In what follows, we will drop the subscripts (s, l_s, f) for notational simplicity and will refer to a pair of causal estimates $(\hat{\theta}_1, \hat{\theta}_2)$ made for potentially different gene phenotypes (e.g., gene expression and protein abundance) and from potentially different gene and/or AD GWAS cohorts, but each corresponding to gene G. It is well known that the IVW estimator for θ_k is $\hat{\theta}_k = \widehat{\boldsymbol{\beta}}_k^{\top} \widehat{\boldsymbol{\alpha}}_k \hat{h}_k^{-1}$ for $k \in \{1,2\}$, where $\widehat{\boldsymbol{\beta}}_k$ is the m_k-length vector of estimated IV associations with the exposure, $\widehat{\boldsymbol{\alpha}}_k$ is correspondingly the same for AD, and $\hat{h}_k \coloneqq \widehat{\boldsymbol{\beta}}_k^{\top} \widehat{\boldsymbol{\beta}}_k$. Let $\hat{s}_k^2$ denote the estimated variance of $\hat{\theta}_k$, which under the causal null hypothesis is equal to $\hat{h}_k^{-1}$ when $\hat{\alpha}_j$ is standardized to have unit variance. We make no use of weights in this IVW estimator because we assume $(\widehat{\boldsymbol{\beta}}_k, \widehat{\boldsymbol{\alpha}}_k)$ are pre-transformed to accomplish weighting in IVW, losing no generality. The score we propose to use for gene G is:

$$S_G = x\left(\sum_k \hat{\theta}_k \hat{s}_k^{-1}\right) = x\left(\sum_k T_k\right) = \mathbf{x}^{\top}\mathbf{t}$$

where x is some fixed linear function which may induce a weighting scheme and $\mathbf{x}$ is its vector representation. We now derive the expectation and variance of S_G. Recall that IVW estimation conditions on $\widehat{\boldsymbol{\beta}}_k$ and so treats it as fixed, while $\widehat{\boldsymbol{\alpha}}_k$ is random, and implicitly assumes that $E\left(\hat{\beta}_{kj_k}\right) = \beta_{kj_k}$ and $E\left(\hat{\alpha}_{kj_k}\right) = \alpha_{kj_k} = \hat{\beta}_{kj_k}\theta_k$. It has been shown that for non-overlapping exposure and outcome GWAS participants, $E\left(\hat{\theta}_k\right) = \eta\theta_k$, where $0 < \eta < 1$ is the so-called reliability ratio of $\eta \approx \beta_{jk}^2\left(\mathrm{Var}\left(\hat{\beta}_{kj_k}\right) + \beta_{jk}^2\right)^{-1}$. The value η can be conceptually considered as measuring the degree of weak instrument bias. See Lorincz-Comi et al. (2024) for a more complete discussion. In our analyses, we restricted the set of selected IVs to those for which the corresponding F-statistic for instrument strength was greater than 10, indicating that $\eta > 0.9$ for all genes. This ensures that $E(\hat{\beta}_{jk} - \beta_{jk})$ is relatively small and so in what follows we assume $E\left(\hat{\beta}_{jk} - \beta_{jk}\right) \approx 0$ using IVW. Using this property, it follows that:

$$E\left(S_G \middle| \hat{h}_k^{-1/2}\right) \approx x\left(\sum_k \theta_k \hat{s}_k^{-1}\right) = x\left(\sum_k T_k^0\right) = \mathbf{x}^\top \mathbf{t}_0$$

where the only approximation that is made is for $E\left(\hat{\beta}_{jk} - \beta_{jk}\right) \approx 0$ as stated above. In words, the score S_G is the weighted sum of causal effect sizes divided by their standard errors. This score will be largest for causal effect sizes that are large and precisely estimated, and smallest for those which are either small or imprecisely estimated. It is not necessary to include weights via a nonzero vector $\mathbf{x}$. We will show later that a suitable choice for $\mathbf{x}$ is $(K^{-1}, \ldots, K^{-1})^\top$ such that $\sqrt{K}S_G$ is asymptotically normal. Note that the score S_G will be largest in absolute value when the signs of all causal effects are the same, whereas it may be made smaller by causal effects of different sign.

Since our goal is to compare the score from one gene to that of another, we now turn to estimation of the variance of S_G. It is immediately clear that $Var(S_G)$ will involve covariance terms, representing scaled covariance between $\hat{\theta}_k$ and $\hat{\theta}_{k'}$ for genes k and k' when $k \neq k'$. Using $(T_k, T_{k'})$, it follows that:

$$\begin{aligned} Cov(T_k, T_{k'}) &= \frac{\widehat{\boldsymbol{\beta}}_k^\top}{\hat{h}_k^{3/2}} Cov\left(\widehat{\boldsymbol{\alpha}}_k, \widehat{\boldsymbol{\alpha}}_{k'} \middle| \widehat{\boldsymbol{\beta}}_k, \widehat{\boldsymbol{\beta}}_{k'}\right) \frac{\widehat{\boldsymbol{\beta}}_{k'}}{\hat{h}_{k'}^{3/2}} \\ &\leq \frac{\widehat{\boldsymbol{\beta}}_k^\top}{\hat{h}_k^{3/2}} \boldsymbol{D}_k Corr(\widehat{\boldsymbol{\alpha}}_k, \widehat{\boldsymbol{\alpha}}_{k'}) \boldsymbol{D}_{k'} \frac{\widehat{\boldsymbol{\beta}}_{k'}}{\hat{h}_{k'}^{3/2}} \\ &\approx \frac{n_0^{ff'}}{n_f n_{f'}} \frac{\widehat{\boldsymbol{\beta}}_k^\top \widehat{\boldsymbol{\beta}}_{k'}}{\left(\hat{h}_k \hat{h}_{k'}\right)^{3/2}} \end{aligned}$$

where

$$\boldsymbol{D}_k = \mathrm{diag}\left(\mathrm{Var}\left(\hat{\alpha}_{kj}\right)^{\frac{1}{2}}\right) \approx \frac{1}{\sqrt{n_f}} \boldsymbol{I}$$

and $n_0^{ff'}$ is the number of subjects overlapping in the two AD GWAS cohorts of sizes n_f and $n_{f'}$. We used the property that $Cov(\widehat{\boldsymbol{\alpha}}_k, \widehat{\boldsymbol{\alpha}}_{k'} | \widehat{\boldsymbol{\beta}}_k, \widehat{\boldsymbol{\beta}}_{k'}) \leq \boldsymbol{D}_k Corr(\widehat{\boldsymbol{\alpha}}_k, \widehat{\boldsymbol{\alpha}}_{k'}) \boldsymbol{D}_{k'}$ since the former represents conditional covariance and the latter marginal covariance. For reasons that will become clearer later, we are interested in the largest value which $Cov(T_k, T_{k'})$ could take and so we assume throughout that the marginal covariance is used instead of the conditional covariance. When no causal effects of either gene k or k' exist, these two terms will be equal, indicating that our calculations above and below are performed under the causal null hypothesis.

It is observed above that $E(S_G)$ is not related to sample overlap but $Var(S_G)$ is. However, also note that $Cov(T_k, T_{k'}) = O\left(\max\left[n_f, n_{f'}\right]\right)$. A key observation here is that:

$$\lim_{\max\left(n_f, n_{f'}\right) \to \infty, m \to M} \hat{h}_k^{-1} = \sigma_k^2,$$

which is the heritability of the gene phenotype where M is the finite number of truly causal gene phenotype SNPs and m is the number of IVs used in MR. In our analyses, m was often less than 10. The variance of S_G including its covariance terms is:

$$\begin{aligned} Var(S_G) &= \mathbf{x}^\top \left(\sum_k Var\left(\hat{\theta}\right) \hat{h}_k^{-1} + 2Cov(T_k, T_{k'}) \right) \mathbf{x} \\ &= \mathbf{x}^\top \left(\sum_k 1 + 2Cov(T_k, T_{k'}) \right) \mathbf{x} \\ &= \mathbf{x}^\top K \mathbf{x} + O\left(\max\left[n_f, n_f'\right]\right). \end{aligned}$$

When equal weights of $K^{-1/2}$ are applied to each summand, i.e. when $\mathbf{x} = (K^{-1/2})$, $Var(S_G) = K + O\left(\max\left[n_f, n_f'\right]\right)$. We now want to fully specify the limiting distribution of S_G. Recall that $\hat{\theta}_k$ is asymptotically normal and denote $E(S_G)$ as S_G^0 and $Var(S_G)$ as $(1 + \kappa_G^{01})^2$, where κ_G^{01} represents all covariance terms in $Var(S_G)$ and we use $\mathbf{x} = (K^{-1})$ where K is the number of phenotype-cohort pairs contributing to the score S_G, it follows that as $K \to \infty$:

$$\sqrt{K} S_G \xrightarrow{D} Normal(S_G^0, (1 + \kappa_G^{01})^2).$$

Additionally, it was implied earlier that as $n_0^{ff'} \left(n_f n_{f'}\right)^{-1} \to 0$, S_G is the weighted sum of K independent terms and so under this condition $\sqrt{K} S_G \xrightarrow{D} Normal(S_G^0, 1)$. The purpose of a score such as S_G is to compare it to scores for $S_{G'}$ when $G \neq G'$, i.e. to compare the strength of evidence between pairs of genes.

Effect of sample overlap

In our data, we used summary statistics from 5 separate AD GWAS, all of which shared at least some overlapping participants. In the table below, we show the factor $n_0^{ff'}/n_f n_{f'}$ for every pair of AD GWAS study samples f and f'.

Factors $n_0^{ff'}/n_f n_{f'}$ for pairs of AD GWAS studies f and f'					
	[1]	[2]	[3]	[4]	[5]
[1] Bellenguez et al. (2022)	-	-	-	-	-
[2] Jansen et al. (2019)	1.87E-6	-	-	-	-
[3] Kunkle et al. (2019)	1.36E-6	1.43E-7	-	-	-
[4] Schwartzentruber et al. (2021)	1.69E-6	1.91E-6	1.56E-6	-	-
[5] Wightman et al. (2022)	3.18E-7	3.95E-7	1.94E-6	3.24E-7	-

These results suggest that if $\widehat{\boldsymbol{\beta}}_k^{\top}\widehat{\boldsymbol{\beta}}_{k'}\left(\hat{h}_k\hat{h}_{k'}\right)^{-3/2}$ is small relative to the $n_0^{ff'}/n_f n_{f'}$ values in the table above, the covariance terms in $Var(S_G)$ will be negligibly small compared to the variance terms. We performed the following calculations to estimate the maximum value any of these covariance terms could take. First, we re-write the covariance terms as

$$\begin{aligned} Cov\left(T_k, T_{k'}|\hat{h}_k, \hat{h}_{k'}\right) &= \frac{n_0^{ff'}}{n_f n_{f'}}\,\widehat{\boldsymbol{\beta}}_k^{\top}\widehat{\boldsymbol{\beta}}_{k'}\left(\hat{h}_k\hat{h}_{k'}\right)^{-\frac{3}{2}} \\ &= \frac{n_0^{ff'}}{n_f n_{f'}}\varrho_{kk'}\frac{1}{\hat{h}_k\hat{h}_{k'}} \\ &= \sigma_{ff'}^{kk'} \\ &< 1.94\times 10^{-6}\frac{1}{\hat{h}_k\hat{h}_k'} \end{aligned} \tag{1}$$

where $\varrho_{kk'}$ is the genetic correlation between gene phenotypes (e.g., expression, protein abundance) k and k' and $\hat{h}_k\hat{h}_{k'}$ is the product of heritabilities explained by the IVs in MR performed on gene phenotypes k and k' separately. It is immediately clear that the only term in Equation (1) which may be greater than 1 is the final one. The maximum values that $\sigma_{ff'}^{kk'}$ could achieve is therefore only a function of $\left(\hat{h}_k\hat{h}_{k'}\right)^{-1}$ since the maximum possible value of $n_0^{ff'}\left(n_f n_{f'}\right)^{-1}\varrho_{kk'}$ in our sample is 1.94E-6. It is clear that $\sigma_{ff'}^{kk'}$ is maximized when $\hat{h}_k\hat{h}_{k'}$ is minimized. Assuming the minimum heritability in the gene phenotype(s) explained by the MR SNP instruments was 1%, the maximum value that $\sigma_{ff'}^{kk'}$ could take is 0.02. Since this is much less than the variance terms in $Var(S_G)$, which are each 1, we can conclude that 1 is a close approximation to $(1+\kappa_G^{01})^2$ and so we may ignore $\sqrt{\sigma_{ff'}^{kk'}}$ when making inference about S_G and $S_{G'}$ for genes G and G'.

The advantage of being able to approximate $Var(S_G)$ by 1 in its asymptotic distribution is that the distribution of S_G and $S_{G'}$ will only differ by their asymptotic expectations and

not their asymptotic variances. This makes direct comparison of S_G and $S_{G'}$ by ranking them appropriate since an estimator of the gene with the largest score is the gene with the largest observed S value (notice the subscript G has been dropped for generality), an estimator of the gene with the second largest score is the gene with the second largest S value, and so on.

Supplemental Method 2: Non-sparse joint estimator

Consider the following model:

$$\mathbf{x} = \mathbf{G}^\top \mathbf{b} + \boldsymbol{\epsilon} \tag{1}$$

where $\mathbf{G} = \left(\mathbf{g}_j\right)_{j=1}^{m} = \left(G_{ij}\right)_{i=1}^{n}$ is an $n \times m$ matrix of genotypes for m SNPs and n unrelated individuals, $\mathbf{b} = \left(b_j\right)_{j=1}^{m}$ is the $m-$ length vector of their joint effects on the phenotype X, and $\boldsymbol{\epsilon}$ is random error uncorrelated with $\mathbf{G}$ that has expectation $\mathbf{0}$ and variance $1 - h^2$, implying each G_{ij} is standardized such that $E(\mathbf{G}^\top \mathbf{G}) = \mathbf{I}$ and h^2 is the heritability of the phenotype X. The model in Equation (1) can be considered a *joint* model since the associations $(b_1, \ldots, b_m)^\top$ are in the presence of all G_{ij} simultaneously. Consider now two *marginal* models for G_{is} and G_{ik}, where $s, k \in \{1, \ldots, m\}$:

$$\mathbf{x} = \mathbf{g}_s \beta_s + \boldsymbol{\epsilon}_s, \qquad \mathbf{x} = \mathbf{g}_k \beta_k + \boldsymbol{\epsilon}_k. \tag{2}$$

In these models (β_s, β_k) represent *total* associations between (G_{is}, G_{ik}) and x_i. In most publicly available GWAS data, we only have access to estimates of (β_s, β_k) and not (b_s, b_k). However, it is generally (b_1, b_2) that we are most interested in making inferences about in practice. As such, we desire some way estimate $\mathbf{b}$ from $\widehat{\boldsymbol{\beta}}$, which is the $m-$ length vector of estimated total associations between the m SNPs and phenotype X.

We begin by deriving covariance between the MLEs $\hat{\beta}_s$ and $\hat{\beta}_k$ under the assumptions of model (2) above. The MLEs are:

$$\hat{\beta}_s = \frac{\mathbf{g}_s^\top \mathbf{x}}{n}, \qquad \hat{\beta}_k = \frac{\mathbf{g}_k^\top \mathbf{x}}{n},$$

and their covariance is

$$Cov\left(\hat{\beta}_s, \hat{\beta}_k\right) \approx \frac{1}{n^2} \mathbf{g}_s^\top Var(\mathbf{x} | \mathbf{g}_s, \mathbf{g}_k) \mathbf{g}_k$$

$$= \frac{1}{n^2} \mathbf{g}_s^\top \mathbf{g}_k \sigma_s \sigma_k \xrightarrow{P} \frac{\sigma_s \sigma_k}{n} \rho_{sk}$$

where ρ_{sk} is the LD correlation between G_{is} and G_{ik} and the convergence statement holds since $n^{-1}\mathbf{g}_s^\top \mathbf{g}_k \xrightarrow{P} \rho_{sk}$. This result has been generalized to further show that $\mathbf{R}^{-1}\widehat{\boldsymbol{\beta}}$ is a close approximation of the joint effects $\mathbf{b}$ (Zhu & Stephens, 2017), where $\mathbf{R}$ is the $m \times m$ matrix of LD correlations between the m SNPs.

Problem

One challenge frequently encountered in practice is that $\mathbf{R}$ is unknown and can only be estimated from reference panels. Denote an unbiased estimator of $\mathbf{R}$ as $\widehat{\mathbf{R}}$. In practice, we often use individuals from a reference panel to estimate LD in our study population, and we often assume individuals in the reference are ancestrally representative of those in the study population. Consider a reference panel of N individuals with the same genetic ancestry as that of the individuals from which $\widehat{\boldsymbol{\beta}}$ was calculated. Let their standardized genotypes be contained in the $N \times m$ matrix $\mathbf{L}$ and $Y \approx F$ indicate that Y is approximately distributed as F. It follows that:

$$\widehat{\mathbf{R}} = N^{-1}\mathbf{L}^\top\mathbf{L} \approx \text{Wishart}(N, N^{-1}\mathbf{R}).$$

Using the properties of the Wishart distribution, $E(\widehat{\mathbf{R}}) = \mathbf{R}$ and $E(\widehat{\mathbf{R}}^{-1}) = N(N-m-1)^{-1}\mathbf{R}^{-1}$. In other words, $\widehat{\mathbf{R}}$ is an unbiased estimator of $\mathbf{R}$ but $\widehat{\mathbf{R}}^{-1}$ is a biased estimator of $\mathbf{R}^{-1}$. Since it is $\widehat{\mathbf{R}}^{-1}$ that we intend to use when estimating $\mathbf{b}$ from $\widehat{\boldsymbol{\beta}}$, it is possible that either estimation of $\mathbf{b}$ or $Var(\hat{\mathbf{b}})$ may be affected by the bias in $\widehat{\mathbf{R}}^{-1}$ for the estimator $\widehat{\mathbf{R}}^{-1}\widehat{\boldsymbol{\beta}}$ of $\mathbf{b}$.

Standard approaches to estimating $\mathbf{b}$ use $\tilde{\mathbf{b}}$:

$$\begin{aligned}
\tilde{\mathbf{b}} &= \widehat{\mathbf{R}}^{-1}\widehat{\boldsymbol{\beta}}, \\
E(\tilde{\mathbf{b}}) &= N(N-m-1)^{-1}\mathbf{b} \\
\widehat{Var}(\hat{\mathbf{b}}) &= \widehat{\mathbf{R}}^{-1}\mathbf{D}\widehat{\mathbf{R}}\mathbf{D}\widehat{\mathbf{R}}^{-1} \\
&\approx \bar{\sigma}^2\widehat{\mathbf{R}}^{-1} \\
E[\widehat{Var}(\tilde{\mathbf{b}})] &= \frac{N}{N-m-1}\bar{\sigma}^2\mathbf{R}^{-1} \qquad (3)
\end{aligned}$$

where

$$\mathbf{D} = \text{diag}\left(SE[\hat{\beta}_j]_{j=1}^{m}\right), \qquad \bar{\sigma}^2 = \mathbf{1}^\top\mathbf{D}\mathbf{D}\mathbf{1}/\text{m}.$$

The approximation is made by assuming $SE(\hat{\beta}_j) \approx SE(\hat{\beta}_{j'})$ for $j \neq j'$; in general, GWAS sample sizes are large enough that $\widehat{SE}(\hat{\beta}_j) = SE(\hat{\beta}_j)$ can be safely assumed. The primary issue with this approach to estimating $\mathbf{b}$ is that $\widehat{Var}(\tilde{\mathbf{b}})$ will underestimate

$Var(\tilde{\mathbf{b}})$ when $N - m < \infty$. To see this, we derive $Var(\tilde{\mathbf{b}})$ using properties of the Wishart distribution and treating $\widehat{\mathbf{R}}^{-1}$ as random instead of fixed. It follows that:

$$Var(\tilde{\mathbf{b}}) = E\left[Var\left(\tilde{\mathbf{b}}\middle|\widehat{\mathbf{R}}^{-1}\right)\right] + Var\left[E\left(\tilde{\mathbf{b}}\middle|\widehat{\mathbf{R}}^{-1}\right)\right] = I_1 + I_2$$

where

$$I_1 = \bar{\sigma}^2 E\left(\widehat{\mathbf{R}}^{-1}\mathbf{R}\widehat{\mathbf{R}}^{-1}\right). \tag{4}$$

$$I_2 = (\boldsymbol{\beta}^\top \otimes \mathbf{I}_m) Var\left(\operatorname{vec} \widehat{\mathbf{R}}^{-1}\right)(\mathbf{I}_m \otimes \boldsymbol{\beta}), \tag{5}$$

Beginning with I_1, we can use the Taylor expansion to find that:

$$I_1 = \frac{N^2}{(N-m-1)^2}\,\bar{\sigma}^2\mathbf{R}^{-1}. \tag{6}$$

Let $c = N(N-m-1)^{-1}$ and note that that c^2 is only approaching 1 (from above) as $\frac{N}{m} \to \infty$. In all other cases, $c^2\bar{\sigma}^2\mathbf{R}^{-1} > c\bar{\sigma}^2\mathbf{R}^{-1} \coloneqq E[\widehat{Var}(\tilde{\mathbf{b}})]$. By comparing Equation (3) to Equation (6), it is established that $E\left[\widehat{Var}\left(\tilde{\mathbf{b}}\right)\right] \leq Var\left(\tilde{\mathbf{b}}\right)$ since $c^2 > c$ in small samples and I_2 will simultaneously be non-negative. Note that Equation (6) is exactly true and not approximated since the first derivative of $c^2\bar{\sigma}^2\mathbf{R}^{-1}\mathbf{R}\mathbf{R}^{-1}$ with respect to $c\mathbf{R}^{-1}$ is $c\bar{\sigma}^2$ and so no higher-order terms in the expansion are nonzero.

Correction

We now propose an alternative estimator of $\mathbf{b}$ which accounts for $c \neq 1$. We then derive its non-asymptotic expectation and variance and propose an estimator of its variance. Asymptotically as $N/m \to \infty$, $c \to 1$, but we are primarily interested in the case when $N \sim m$. Let the alternative estimator be:

$$\ddot{\mathbf{b}} = c^{-1}\widehat{\mathbf{R}}^{-1}\widehat{\boldsymbol{\beta}}, \tag{7}$$

which can be seen to be unbiased when compared to $E(\tilde{\mathbf{b}})$ from above. The expression for $Var(\ddot{\mathbf{b}})$ becomes $Var(\ddot{\mathbf{b}}) \coloneqq \ddot{I}_1 + \ddot{I}_2$, where:

$$\ddot{I}_1 = \bar{\sigma}^2\mathbf{R}^{-1},$$

$$\ddot{I}_2 = c^{-2}\mathbf{W}\left(d_1\mathbf{P} + d_2\left[\mathbf{I}_{m^2} + \mathbf{K}_{m,m}\right]\mathbf{T}\right)\mathbf{W}^\top,$$

$$\mathbf{W} = (\boldsymbol{\beta}^\top \otimes \mathbf{I}_m), \qquad \mathbf{P} = \operatorname{vec}(\mathbf{R}^{-1})\operatorname{vec}(\mathbf{R}^{-1})^\top, \qquad \mathbf{T} = \mathbf{R}^{-1} \otimes \mathbf{R}^{-1},$$

$$d_1 = \frac{N^2}{(N-m)(N-m-1)^2(N-m-3)},$$

$$d_2 = \frac{N-m-1}{2}d_1,$$

and $\mathbf{K}_{m,m}$ is a $m^2 \times m^2$ commutation matrix. We propose the estimator of $Var(\ddot{\mathbf{b}})$ which simply replaces $\mathbf{R}^{-1}$ by $c^{-1}\widehat{\mathbf{R}}^{-1}$ and $\boldsymbol{\beta}$ by $\widehat{\boldsymbol{\beta}}$. One potential limitation of this estimator of $Var(\ddot{\mathbf{b}})$ is that the estimation error of $(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ is assumed to be absent when we know it generally is not. In the following section, we can observe a small upward bias a result of this assumption, though its effect is not large.

Simulation

We performed a simulation in which we generated individual-level data for 200 individuals and 20 SNPs, 3 of which explained 5% heritability in the quantitative phenotype X and the rest of which had no association with X. We set the true LD structure between each of the 20 SNPs to be first-order autoregressive with correlation parameter 0.5 and randomly drew genotypes from a multivariate normal distribution. This choice was made to examine the statistical properties described above without encountering computational issues related to estimating a covariance matrix from multivariate binomial (i.e., correlated genotype) data. We estimated $\mathbf{b}$ using four methods: (i) marginal association estimates (termed hereafter as 'marginal'), (ii) joint association estimates using individual-level data and including all SNPs in the same model together ('true joint'), (iii) standard joint estimates from marginal summary statistics ('standard joint'), corresponding to $\tilde{\mathbf{b}}$ above, and (iv) alternative joint estimates from marginal summary statistics ('alternative joint'), corresponding to $\ddot{\mathbf{b}}$ above.

Figure 1

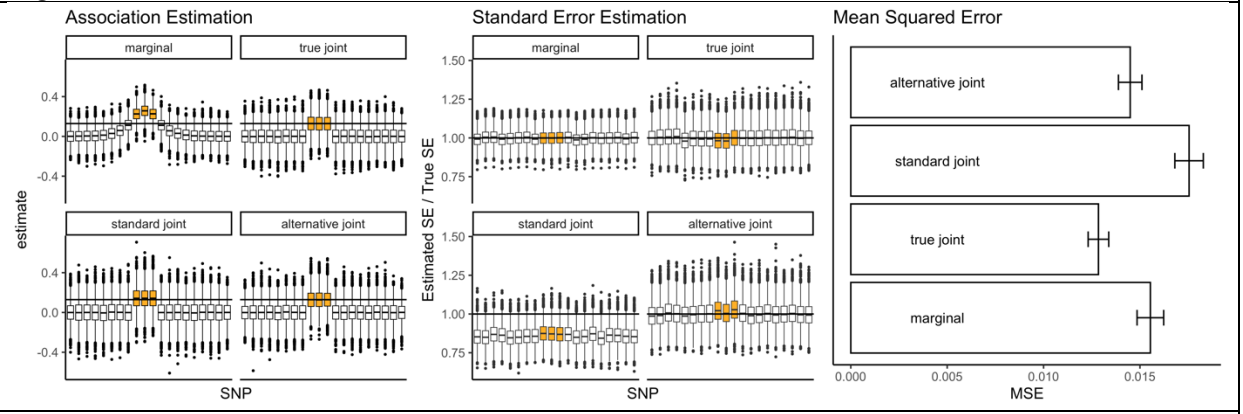

Displayed in the left panel are estimates of association between each of 20 correlated SNPs and a quantitative phenotype using four methods. Displayed in the middle panel are estimated standard errors divided by true standard errors, equal here to the empirical standard deviations of association estimates across all 1,000 simulations, for each of four methods. Displayed in the right panel are values of average mean squared error (MSE) across the 5,000 simulations, where in each simulation MSE was proportional to $\left\|\hat{\mathbf{b}} - \mathbf{b}\right\|_2^2$ for a generic estimator $\hat{\mathbf{b}}$ of $\mathbf{b}$.

Figure 1 displays the simulation results which demonstrate that $E\left(\tilde{b}_j - b_j | b_j \neq 0\right) > 0$ and $E\left(\widehat{Var}\left[\tilde{b}_j\right]\right) < Var\left(\tilde{b}_j\right)$. In other words, the standard estimator of joint effects from marginal GWAS summary statistics is biased away from 0 and the corresponding estimated standard errors are biased towards 0. Using Equations (3) and (6) above, we

can expect under these simulation conditions that the proportion of bias in the estimated standard errors of $\tilde{\mathbf{b}}$ will be equal to $c = (N - m - 1)/N$, which is what we approximately observe in Figure 1. Conversely, the alternative estimator $\ddot{\mathbf{b}}$ that we proposed above appears to estimate $\mathbf{b}$ without bias. There is a very small upward bias of $\widehat{Var}(\ddot{b}_j)$, making tests using these standard errors relatively conservative, but this bias is small and so is not expected to have a substantial impact on inference. Additionally, this upward bias is significantly smaller than the downward bias of $\widehat{Var}(\tilde{\mathbf{b}})$ using the standard approach. This small bias likely arises from the source previously mentioned earlier: Calculation of $\widehat{Var}(\ddot{\mathbf{b}})$ is effectively conditioned on $\widehat{\boldsymbol{\beta}}$ when in fact $\widehat{\boldsymbol{\beta}}$ should be treated as random. Such an estimator of $Var(\ddot{\mathbf{b}})$ would correct for GWAS estimation error of $(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$. We also observe that marginal estimates of association are generally biased away from 0, even for SNPs with no true association with the phenotype. Finally, our method using $\ddot{\mathbf{b}}$ has lower mean squared error (MSE) (i.e., $\left\|\hat{\mathbf{b}} - \mathbf{b}\right\|_2^2$ for a generic estimator $\hat{\mathbf{b}}$ of $\mathbf{b}$) than the standard method using $\tilde{\mathbf{b}}$. Smaller MSE using $\ddot{\mathbf{b}}$ may be explained by the small upward bias of $\tilde{\mathbf{b}}$ or by its larger variance. Note that $\tilde{\mathbf{b}}$ will necessarily have a larger variance than $\ddot{\mathbf{b}}$ since $\left|\ddot{b}_j\right| < \left|\tilde{b}_j\right|$ for all $j = 1, \ldots, m$. As a result, $\ddot{\mathbf{b}}$ is both a less biased and more powerful estimator of $\mathbf{b}$ than $\tilde{\mathbf{b}}$.

Supplementary References

1. Bellenguez, C. et al. New insights into the genetic etiology of Alzheimer's disease and related dementias. Nature Genetics 54, 412-436 (2022)

2. Jansen, I.E. et al. Genome-wide meta-analysis identifies new loci and functional pathways influencing Alzheimer's disease risk. Nat Genet 51, 404-413 (2019).

3. Kunkle, B.W. *et al.* Novel Alzheimer Disease Risk Loci and Pathways in African American Individuals Using the African Genome Resources Panel: A Meta-analysis. *JAMA Neurol* **78**, 102-113 (2021).

4. Lorincz-Comi, N., Yang, Y., Li, G., & Zhu, X. (2024). MRBEE: A bias-corrected multivariable Mendelian randomization method. *Human Genetics and Genomics Advances*, *5*(3).

5. Schwartzentruber, J. *et al.* Genome-wide meta-analysis, fine-mapping and integrative prioritization implicate new Alzheimer's disease risk genes. *Nature Genetics* **53**, 392-402 (2021).

6. Wightman, D.P. *et al.* A genome-wide association study with 1,126,563 individuals identifies new risk loci for Alzheimer's disease. *Nature Genetics* **53**, 1276-1282 (2021).

7. Zhu, X., & Stephens, M. (2017). Bayesian large-scale multiple regression with summary statistics from genome-wide association studies. *The annals of applied statistics*, *11*(3), 1561.