DOI: 10.1002/gepi.22077

RESEARCH ARTICLE

Genetic WILEY Epidemiology



Mendelian randomization with fine-mapped genetic data: Choosing from large numbers of correlated instrumental variables

Stephen Burgess^{1,2} \bigcirc | Verena Zuber^{1,3} | Elsa Valdes-Marquez⁴ | Benjamin B Sun² |

Jemma C Hopewell⁴

¹MRC Biostatistics Unit, Cambridge, United Kingdom

²Department of Public Health and Primary Care, University of Cambridge, Cambridge, United Kingdom

³European Bioinformatics Institute, Hinxton, nr Duxford, United Kingdom

⁴Clinical Trial Service Unit and Epidemiological Studies Unit, Nuffield Department of Population Health, University of Oxford, Oxford, United Kingdom

Correspondence

Stephen Burgess, MRC Biostatistics Unit, Cambridge Institute of Public Health, Robinson Way, Cambridge CB2 0SR, UK. Email: sb452@medschl.cam.ac.uk

Funding information:

Grant sponsor: Wellcome Trust; Grant sponsor: Royal Society; Grant number: 204623/Z/16/Z; Grant sponsor: British Heart Foundation Basic Science Research Fellowship; Grant number: FS/14/55/30806. ABSTRACT

Mendelian randomization uses genetic variants to make causal inferences about the effect of a risk factor on an outcome. With fine-mapped genetic data, there may be hundreds of genetic variants in a single gene region any of which could be used to assess this causal relationship. However, using too many genetic variants in the analysis can lead to spurious estimates and inflated Type 1 error rates. But if only a few genetic variants are used, then the majority of the data is ignored and estimates are highly sensitive to the particular choice of variants. We propose an approach based on summarized data only (genetic association and correlation estimates) that uses principal components analysis to form instruments. This approach has desirable theoretical properties: it takes the totality of data into account and does not suffer from numerical instabilities. It also has good properties in simulation studies: it is not particularly sensitive to varying the genetic variants included in the analysis or the genetic correlation matrix, and it does not have greatly inflated Type 1 error rates. Overall, the method gives estimates that are less precise than those from variable selection approaches (such as using a conditional analysis or pruning approach to select variants), but are more robust to seemingly arbitrary choices in the variable selection step. Methods are illustrated by an example using genetic associations with testosterone for 320 genetic variants to assess the effect of sex hormone related pathways on coronary artery disease risk, in which variable selection approaches give inconsistent inferences.

KEYWORDS

Mendelian randomization, allele score, correlated variants, summarized data, conditional analysis

1 | BACKGROUND

In a Mendelian randomization investigation, genetic variants that are instrumental variables for a given risk factor are used to assess the causal effect of the risk factor on an outcome (Burgess & Thompson, 2015; Davey Smith & Ebrahim, 2003). An association between such a genetic variant and the outcome is indicative of a causal effect of the risk factor on the outcome (Didelez & Sheehan, 2007; Lawlor, Harbord, Sterne, Timpson, & Davey Smith, 2008). When there are multiple uncorrelated genetic variants that are instrumental variables for the same risk factor, power to detect

© 2017 The Authors Genetic Epidemiology Published by Wiley Periodicals, Inc.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

a causal effect is maximized by including all such genetic variants in a single analysis (Pierce, Ahsan, & VanderWeele, 2011). However, when genetic variants are correlated, it is not clear how to choose which variants to include in the analysis to obtain the most efficient estimate possible without the analysis suffering from numerical instabilities when there are large numbers of highly correlated candidate variants (such as with fine-mapped genetic data).

1.1 | Theoretical viewpoint

2

-WILEY

If individual-level data are available on the genetic variants (potentially correlated), risk factor, and outcome for the same participants, then the two-stage least squares (2SLS) method provides the most efficient estimate of the causal effect (among all instrumental variable estimators using linear combinations of the instruments and under conditional homoscedasticity—the error term in the model relating the outcome to the risk factor has constant variance conditional on the instruments) (Wooldridge, 2009). Use of the 2SLS method for estimating a causal effect is discussed at length elsewhere in the literature; see Angrist and Imbens (1995) for a theoretical introduction, and Didelez, Meng, and Sheehan (2010) for a discussion in the context of Mendelian randomization.

The first stage of the 2SLS method regresses the risk factor on all the genetic variants. As the sample size increases, the coefficient of any variant that does not explain independent variation in the risk factor will tend to zero, and so its contribution to the analysis decreases to zero. This implies that an optimally efficient Mendelian randomization analysis should include all genetic variants associated with the risk factor in a conditional analysis. The inclusion of additional variants not independently associated with the risk factor will not have a negative impact on the analysis asymptotically (as their coefficient for contribution to the analysis will tend to zero), but

will not add to the precision of the causal estimate either. As an aside, fitted values from the first-stage of the 2SLS method are equivalent (up to an additive constant) to values of an allele score (also called a genetic risk score). This implies that the optimal weights in an allele score with correlated variants are the conditional (multivariable) associations of the variants with the risk factor.

1.2 | Estimating a causal effect using summarized data

The 2SLS estimate can also be obtained using summarized data on genetic associations with the risk factor and with the outcome from univariable regression analyses of the risk factor or outcome on each genetic variant in turn. This is important as such summarized data from large consortia are often publicly available, enabling Mendelian randomization investigations to be performed on large sample sizes without the need for costly and time-consuming data-sharing arrangements (Burgess et al., 2015). This estimate can also be calculated in a two-sample setting, in which genetic associations with the risk factor and with the outcome are estimated in different samples (Inoue & Solon, 2010).

If the genetic association with the risk factor for genetic variant *j* is $\hat{\beta}_{Xj}$ with standard error (SE) se($\hat{\beta}_{Xj}$), and with the outcome is $\hat{\beta}_{Yj}$ with SE se($\hat{\beta}_{Yj}$), and assuming that genetic variants are uncorrelated, then the causal estimate is (Johnson, 2013):

Inverse variance weighted estimate (uncorrelated variants)

$$= \frac{\sum_{j} \hat{\beta}_{Yj} \hat{\beta}_{Xj} \operatorname{se}(\hat{\beta}_{Yj})^{-2}}{\sum_{j} \hat{\beta}_{Xj}^{2} \operatorname{se}(\hat{\beta}_{Yj})^{-2}}.$$
 (1)

This is referred to as the inverse variance weighted (IVW) estimate (Burgess, Butterworth, & Thompson, 2013). It is the weighted mean of the 2SLS estimates using each genetic variant individually $(\hat{\beta}_{Y_i}/\hat{\beta}_{X_i})$ with the inverse-variance weights $[\operatorname{se}(\hat{\beta}_{Y_i})/\hat{\beta}_{X_i}]^{-2}$. The variant-specific estimates are combined using the standard formula for a fixed-effect meta-analysis (Borenstein, Hedges, Higgins, & Rothstein, 2009). This same estimate can be obtained by weighted regression of the genetic associations with the outcome $\hat{\beta}_{Yi}$ on the genetic associations with the risk factor $\hat{\beta}_{X_i}$ using weights se $(\hat{\beta}_{Y_i})^{-2}$ and with the intercept term set to zero. The IVW estimate is equivalent to the 2SLS estimate when the genetic variants are uncorrelated (Burgess, Dudbridge, & Thompson, 2015). This formula does not take into account uncertainty in the genetic associations with the risk factor; however, these associations are typically more precisely estimated than those with the outcome, and ignoring this uncertainty does not lead to inflated Type 1 error rates for the IVW estimate in realistic scenarios (Burgess et al., 2013).

When genetic variants are correlated, the IVW method can be extended to account for the correlations between genetic variants (Burgess, Dudbridge, & Thompson, 2016). This can be motivated by considering generalized weighted linear regression of the genetic associations with the outcome on the genetic associations with the risk factor using the weighting matrix Ω , where $\Omega_{j_1j_2} = \operatorname{se}(\hat{\beta}_{Yj1}) \operatorname{se}(\hat{\beta}_{Yj2}) \rho_{j_1j_2}$, and $\rho_{j_1j_2}$ is the correlation between genetic variants j_1 and j_2 . The causal estimate is:

IVW estimate (correlated variants)

$$= (\hat{\boldsymbol{\beta}}_X^T \Omega^{-1} \hat{\boldsymbol{\beta}}_X)^{-1} \hat{\boldsymbol{\beta}}_X^T \Omega^{-1} \hat{\boldsymbol{\beta}}_Y, \qquad (2)$$

where $\hat{\boldsymbol{\beta}}_X$, $\hat{\boldsymbol{\beta}}_Y$ are vectors of the genetic associations, and T is a vector transpose. Again, this estimate is equivalent to the 2SLS estimate that is obtained using individual-level data (see Appendix for proof). It therefore inherits the efficiency

property of the 2SLS estimate as the optimally efficient causal estimate based on all the genetic variants.

1.3 | Scope of paper

In this paper, we illustrate and provide guidance on choosing variants to include in a Mendelian randomization with finemapped genetic data. We first provide a motivating example analysis based on summarized genetic associations for hundreds of correlated genetic variants in a single gene region. We demonstrate and explain why including too many genetic variants in such an analysis can lead to numerical instabilities and inflated Type 1 error rates. We also show that estimates based on a few variants can be highly sensitive to the choice of these variants. A novel approach is presented using principal components analysis (PCA) to ensure that all variants contribute to the analysis, but without introducing numerical instabilities. We discuss practical implications of these findings for applied Mendelian randomization investigations. Software code in the R programming language for implementing the analyses discussed in the paper is provided in the Appendix.

2 | MOTIVATING EXAMPLE: SERUM TESTOSTERONE AND CORONARY HEART DISEASE RISK

We consider an example of Mendelian randomization analysis with serum testosterone as the risk factor and coronary artery disease (CAD) risk as the outcome using genetic variants in the SHBG gene region. The associations of 325 individual SNPs with testosterone in 3,225 men of European ancestry are reported by Jin et al. (2012); associations of 322 of these variants with CAD risk in 60,801 CAD cases and 123,504 controls are reported by the CARDIoGRAMplusC4D Consortium (2015). Previously, in an independent dataset, Coviello et al. (2012) demonstrated at least six separate signals in the SHBG gene region at a genome-wide level of significance in 21,791 individuals from 10 cohort studies, plus three more variants associated with sex hormone binding globulin (SHBG) on adjustment for these six variants. In all analyses, correlations between variants are estimated using 1,000 Genomes Phase 3 data on 503 individuals of European descent as reference data. A further two variants were monomorphic in the reference data; analyses are conducted using the remaining 320 variants. As variants in the SHBG gene region are associated with circulating levels of both testosterone and SHBG, a positive Mendelian randomization finding would not distinguish which of these is a causal risk factor, but would suggest that sex hormone related mechanisms have a causal role in cardiovascular disease.

Three approaches are taken here to choose which variants to include in a Mendelian randomization analysis. First, we take eight variants from the conditional analysis in the independent dataset reported by Coviello et al. (the association with testosterone in the data under analysis was not available for one variant). Second, we perform a stepwise conditional approach using the summarized associations reported by Jin et al., selecting at each step of the analysis the variant having the lowest P-value for association with the risk factor in the conditional analysis. We proceed until no additional variants are associated with the risk factor at P < 0.0001 or P <0.001. This approach is implemented using the GCTA software. Third, we perform a stepwise pruning approach (Yang et al., 2012), selecting at each step of the analysis the variant having the lowest P-value for association with the risk factor in a marginal (univariable) analysis. Once a variant is selected, all other variants whose correlation with the selected variant is greater in magnitude than a given correlation threshold (taken as 0.2, 0.4, 0.6, 0.8, 0.9, and 0.95; equivalent to an r^2 threshold of 0.04, 0.16, 0.36, 0.64, 0.81, and 0.9025) are removed from the analysis. We continue until each variant is either selected or removed. This ensures that a set of variants is chosen for each threshold correlation such that the variants are each marginally associated with the risk factor, and the pairwise correlations are all below the threshold correlation. Although a data-driven approach to selecting variants to include in a Mendelian randomization investigation is often unwise (Burgess, Thompson, & CRP CHD Genetics Collaboration, 2011), in this case the associations with the risk factor and with the outcome are estimated in nonoverlapping samples, and so "winner's curse" bias in the genetic associations with the outcome should not arise.

The Mendelian randomization estimates are presented in Table 1. Fixed-effect analysis models that account for correlations between variants are used throughout. A fixed-effect model assumes that all genetic variants are targeting the same causal effect parameter. This is reasonable when all the genetic variants are in the same gene region and so are likely to affect the risk factor in the same way. If genetic variants in different gene regions are used in a Mendelian randomization investigation, then a random-effects model should be preferred, particularly if the risk factor is a complex phenotype such as blood pressure, as different genetic variants influencing blood pressure via different biological mechanisms may lead to different magnitudes of change in the outcome (Bowden, Davey Smith, Haycock, & Burgess, 2016). Despite the two approaches using a conditional analysis and the pruning approach at a threshold correlation of 0.2 including similar numbers of variants in the analysis, the causal estimates in these three analyses differed substantially-by over two SEs, and gave opposing substantive conclusions. In the pruning approach, as the threshold correlation increased, more variants were included in the Mendelian randomization analysis, and the precision of the causal estimate increased. However, for very large values of the threshold correlation, the SE of

4 WILEY

	Threshold	Correlation			
Selection Approach	$\overline{\rho}$	r^2	Number of Variants	Estimate (SE)	
Conditional analysis in independent dataset (Coviello)	-	_	8	-0.258 (0.097)	
GCTA at $P < 0.0001$	-	-	6	-0.009 (0.058)	
GCTA at $P < 0.001$	-	-	19	-0.068 (0.042)	
Pruning	0.2	0.04	8	-0.110 (0.094)	
Pruning	0.4	0.16	20	-0.085 (0.067)	
Pruning	0.6	0.36	39	-0.017 (0.051)	
Pruning	0.8	0.64	62	-0.137 (0.031)	
Pruning	0.9	0.81	85	−0.537 (-) ^a	
Pruning	0.95	0.9025	104	-1.099(0.001)	

Estimates (SE) of causal effect of testosterone on CAD risk (estimates are log odds ratios per unit increase in log-transformed testosterone) from IVW method (accounting for correlation) with variants selected using three different approaches and (for the pruning method) six different threshold correlations (measured by ρ and by r^2). ^aThe variance estimate was negative, indicating that the weighting matrix was not positive definite, meaning that either the standard errors in the weighting matrix were imprecisely estimated, or else were not compatible with the correlation matrix.



FIGURE 1 Estimated genetic associations and 95% confidence intervals with testosterone (nmol/L, then log-transformed) and with coronary artery disease risk (log odds ratios): (left) for 104 genetic variants included in Mendelian randomization analysis with threshold correlation 0.95 ($r^2 = 0.9025$); (right) for 62 genetic variants with threshold correlation 0.8 ($r^2 = 0.64$)

Note: The heavy dashed line is the IVW estimate (accounting for correlation between variants).

the causal estimate is implausibly small. With a threshold correlation of 0.9, the SE of the causal estimate was not defined due to the variance estimate being negative. With a threshold correlation of 0.95, the causal estimate is clearly spurious, as can be seen by visual inspection of the data (Fig. 1, left panel). The result with a correlation of 0.8 is also suspect (Fig. 1, right panel), as the variants having the greatest associations with testosterone all lie above the causal effect estimate. Even at lower threshold correlations of 0.4 and 0.6, the SEs of the causal estimate are substantially lower than those calculated using the conditional approach. This may be due to the extra variants explaining additional variability in the risk factor; the reduction in SE corresponds to a 97% relative increase in variance explained by the variants at a threshold of 0.4 compared with at 0.2, and a 240% increase

at a threshold of 0.6. It is unclear which of the estimates in Table 1 are reliable, and therefore whether evidence supports testosterone as a causal risk factor for coronary heart disease risk or not.

3 | CHOOSING THE RIGHT NUMBER OF VARIANTS

To resolve the question of how to choose which variants to include in a Mendelian randomization analysis, we explore reasons why analyses that include too many or too few genetic variants may go wrong, and propose a solution that incorporates associations on large numbers of genetic variants into the analysis, but does not suffer from numerical instabilities.

3.1 | Too many variants: Near-singular genetic correlation matrix

A matrix is singular if it cannot be inverted-formally, if the determinant of the matrix is zero. This occurs when the rows or columns of a matrix are linearly dependent; that is, at least one column (or row) can be calculated as a linear sum of multiples of the other columns (known as multicollinearity). This will occur for the genetic correlation matrix when two genetic variants are in perfect linkage disequilibrium, or alternatively if a small number of haplotypes are present in the data (perfect multicollinearity can occur even if no pair of variants is highly correlated). In contrast, a near-singular matrix can be inverted, but its determinant is close to zero. This occurs in a regression model when there is substantial, but not perfect, multicollinearity. As sample sizes for estimating genetic correlations increase, singular matrices will become less common, but near-singular genetic correlation matrices are likely to become more common. This is because a discrepant allele count in a single individual (which could represent a genotyping error) can lead to a singular matrix becoming nonsingular. A near-singular matrix is problematic as elements of its inverse can be very large. In the motivating example with correlation thresholds of 0.9 and 0.95, the maximal element of the inverse of the correlation matrix is over 10 million.

If a matrix is exactly singular, then it cannot be inverted, and the analysis will report an error. If a matrix is nearsingular, then the analysis may report an estimate without giving any indication that the estimate may be misleading (as observed in Fig. 1). In conjunction with discrepancies in the genetic association estimates, near-singular behavior can lead to overly precise as well as highly misleading estimates. Discrepancies may occur because of the rounding of association estimates (particularly for summarized genetic associations taken from the literature), inaccuracy and uncertainty in correlation estimates, and genetic association estimates and/or correlation estimates being estimated in different samples. When multiplied by the large numbers in the inverse of a near-singular genetic correlation matrix, small discrepancies in association estimates are magnified. Overprecision in the causal estimate will occur when genetic association estimates that should be similar based on the correlation matrix are more dissimilar than expected.

3.2 | Too few variants: Unstable estimates

Although theoretical considerations suggest that a Mendelian randomization analysis should be based on only variants associated with the risk factor in a conditional analysis, in practice this results in a Mendelian randomization estimate that only uses data on a small number of variants. In the motivating example, the conditional analyses suggest that less than 10 variants should be included in the analysis; associations with the remaining over 300 variants are ignored. In some cases and in particular in the motivating example, the causal estimate is highly sensitive to the choice of which variants are included in the analysis. This leads to unstable Mendelian randomization estimates—if one of the selected variants in the conditional analysis happened not to be measured, or failed quality control (QC) criteria, then a different set of variants would have been obtained from the conditional analysis, resulting in a different Mendelian randomization estimate.

3.3 | Just right? Principal components analysis

One potential solution for resolving the problem of multiple correlated variants is PCA. The use of PCA has been previously suggested for reducing the dimensionality of the instrumental variable space to resolve issues of weak instrument bias (Winkelried & Smith, 2011), and as a tool for grouping variants in a fine-mapped gene region (Cai et al., 2013). We perform unscaled PCA on a weighted version of the genetic correlation matrix $\Psi_{j_1j_2} =$ $\hat{\beta}_{Xj1}\hat{\beta}_{Xj2} \operatorname{se}(\hat{\beta}_{Yj1})^{-1} \operatorname{se}(\hat{\beta}_{Yj2})^{-1}\rho_{j_1j_2}$. The diagonal elements of this matrix are the inverse-variance weights, and so each is equal to the precision of the causal estimate based on that variant alone.

Assuming that associations for all variants are estimated in the same sample size, these diagonal elements are proportional to the amount of variance in the risk factor explained by the genetic variant. This can be seen as the SEs of the associations with the outcome will be directly proportional to the SEs of the associations with the risk factor, which in turn relate to the minor allele frequencies MAF_i : if the variant is a diallelic SNP, then se $(\hat{\beta}_{X_i})^{-2} \propto MAF_i(1 - MAF_i)$ (Burgess et al., 2016). (The proportion of variance in the risk factor explained by genetic variant j is $\beta_{X_i}^2 \times MAF_i(1 - MAF_i)$, where β_{X_i} is measured in standard deviation [SD] units.) Hence, if the variants were uncorrelated, then the first principal component would be the genetic variant that explained the largest proportion of variance in the risk factor, and so on. For correlated variants, the first principal component represents a linear combination of variants that explains the largest proportion of variance in the risk factor, and each subsequent principal component is the linear combination of variants that explains the next largest proportion of variance while being orthogonal to the previous principal components.

This choice of matrix should be advantageous for Mendelian randomization investigations over PCA approaches on the unweighted matrix of genetic correlations. If two variants are perfectly correlated, but the estimates for one are measured in a larger sample size, then the precision of the association with the outcome (se($\hat{\beta}_{Yj}$)⁻¹) will be greater for this variant, and so it will (correctly) be preferentially selected. The number of principal components to be included in the analysis can be chosen based on a

threshold of variance in the weighted genetic correlation matrix. Once the principal components have been selected, we multiply the vector of genetic associations with the risk factor by the matrix of principal components, we multiply the vector of genetic associations with the outcome by the matrix of principal components, and pre- and postmultiply the genetic correlation matrix by the matrix of principal components. The IVW method is then performed on the transformed vectors of genetic associations and the transformed correlation matrix. If the matrix $\Psi = W \Lambda W^T$, where W is the matrix of eigenvectors (or loadings), and Λ is the diagonal matrix with the eigenvalues $\lambda_1 > ... > \lambda_J$ on the diagonal, then let W_k be the matrix constructed of the first kcolumns of W. Then we define:

6

-WILEY

 $\tilde{\boldsymbol{\beta}}_X = \boldsymbol{W}_k^T \hat{\boldsymbol{\beta}}_X$ as the transformed genetic associations with the risk factor:

 $\tilde{\boldsymbol{\beta}}_{Y} = \boldsymbol{W}_{k}^{T} \hat{\boldsymbol{\beta}}_{Y}$ as the transformed genetic associations with the outcome; AND

 $\tilde{\Omega} = W_k^T \Omega W_k$ as the transformed weighting matrix.

Then, the PCA-IVW estimate is given by:

$$(\tilde{\boldsymbol{\beta}}_X^T \tilde{\boldsymbol{\Omega}}^{-1} \tilde{\boldsymbol{\beta}}_X)^{-1} \tilde{\boldsymbol{\beta}}_X^T \tilde{\boldsymbol{\Omega}}^{-1} \tilde{\boldsymbol{\beta}}_Y.$$
(3)

For the example of testosterone and CAD risk, 99% of the variance in this matrix was explained by the first eight principal components, and 99.9% by the first 17 principal components. The corresponding estimates using these principal components as instruments were -0.065 (SE 0.099) and -0.045 (0.083), respectively. These estimates are similar in precision to that using the previous conditional analysis for variable selection, but less precise than those calculated using the GCTA method on the data under analysis or a liberal correlation threshold in the pruning method.

Overall, the conclusion from this motivating example is that there is no strong evidence of a causal relationship between sex hormone related pathways and coronary heart disease risk on the basis of the genetic evidence presented here. The more extreme estimates suggesting a causal relationship come from the less reliable methodological approaches, and these estimates should not be trusted. A more detailed analysis could be performed using genetic variants previously associated with either SHBG or testosterone from other gene regions, although the specific relevance of other variants to sex hormone related pathways is not always clear. Additionally, it could be argued that these analyses should be performed in men and women separately. An authoritative analysis conclusively judging the causal relevance of sex hormone related pathways to coronary heart disease risk is beyond the scope of this methodologically focused paper.

4 | SIMULATION STUDY

We illustrate statistical issues arising from using too many and too few variants in a series of simulation studies based on the motivating example. Again, fixed-effect analysis models are used throughout.

4.1 | Sensitivity to choice of genetic variants

First, we repeated the analyses of the motivating example except using only 180 of the 360 genetic variants at a time. This represents a scenario in which only a subset of the genetic variants in the analysis were measured. Sets of 180 variants were chosen at random 10,000 times.

4.2 | Sensitivity to correlation matrix

Second, we repeated the analyses of the motivating example except varying the correlation matrix. We took a bootstrap sample of the reference data (same size sample as the original data, sampled with replacement), and calculated a correlation matrix based on this sample. This procedure was performed 10,000 times.

For each of these simulation analyses, we performed the pruning method for selecting genetic variants at a threshold correlation of 0.2, 0.4, 0.6 and 0.8, and the PCA method using components that explained 99% and 99.9% of the variance in the summarized association matrix. Results are presented in Table 2. In both simulation studies, as the threshold in the pruning approaches increased, the mean SE of the causal estimates decreased, and the mean causal estimate also changed substantially. For a threshold correlation of $\rho = 0.8$, causal estimates were unstable, and were particularly sensitive to changes in the correlation matrix. In contrast, estimates using the PCA approach were not so precise, but they were far less variable between iterations.

4.3 | Rounding of association estimates

Finally, we simulated genetic associations with the risk factor and with the outcome directly. Genetic associations with the risk factor were drawn for 320 variants from a multivariable normal distribution with mean vector the measured genetic associations with testosterone from the motivating example and variance-covariance matrix Ω_X , where $\Omega_{Xj_1j_2} = \operatorname{se}(\hat{\beta}_{Xj1})\operatorname{se}(\hat{\beta}_{Xj2})\rho_{j_1j_2}$. The associations with the outcome are drawn from a multivariate normal distribution with mean zero and variance-covariance matrix Ω , where $\Omega_{j_1j_2} = \operatorname{se}(\hat{\beta}_{Yj1})\operatorname{se}(\hat{\beta}_{Yj2})\rho_{j_1j_2}$ as defined above. This

TABLE 2 Simulations Varying Choice of Variants and Correlation Matrix

	Varying Choice of Variants			Varying Correlation Matrix			
Selection Approach	Mean Estimate	SD	Mean SE	Mean Estimate	SD	Mean SE	
Pruning at $\rho = 0.2$	-0.100	0.044	0.094	-0.114	0.035	0.090	
Pruning at $\rho = 0.4$	-0.093	0.032	0.078	-0.074	0.027	0.065	
Pruning at $\rho = 0.6$	-0.009	0.049	0.060	-0.018	0.052	0.046	
Pruning at $\rho = 0.8$	-0.024	0.402	0.048 ^a	_b	-	-	
PCA at 99% of variance	-0.053	0.028	0.098	-0.051	0.027	0.096	
PCA at 99.9% of variance	-0.045	0.025	0.084	-0.047	0.017	0.083	

Means of estimates, SDs of estimates, and mean SEs for 10,000 iterations based on motivating example: (i) varying the choice of variants and (ii) varying the correlation matrix. Six approaches for selecting genetic variants are performed: four based on pruning at different correlation thresholds (ρ) and two based on PCA. ^aExcluding 536 iterations in which the standard error was not defined.

^bEstimates were highly variable and the standard error was not defined for a large proportion of iterations.

TABLE 3 Simulation Rounding Association Estimates

	Unround	ed		Three De	cimal Places		Two Deci	mal Places	
Selection Approach	SD	Mean SE	Power	SD	Mean SE	Power	SD	Mean SE	Power
Null causal effect									
Pruning at $\rho = 0.2$	0.080	0.079	5.0	0.080	0.080	4.9	0.086	0.077	7.3
Pruning at $\rho = 0.4$	0.067	0.066	5.0	0.067	0.066	5.1	0.073	0.063	9.2
Pruning at $\rho = 0.6$	0.049	0.049	5.0	0.050	0.050	4.9	0.066	0.047	16.5
Pruning at $\rho = 0.8$	0.027	0.022	10.5	0.175	0.022	40.8	0.418	0.020	62.2
PCA at 99% of variance	0.089	0.090	4.6	0.090	0.090	4.6	0.094	0.083	8.0
PCA at 99.9% of variance	0.075	0.075	4.6	0.075	0.076	4.5	0.079	0.069	9.0
Positive causal effect of 0.1									
Pruning at $\rho = 0.2$	0.080	0.079	24.8	0.080	0.080	24.6	0.086	0.077	27.9
Pruning at $\rho = 0.4$	0.067	0.066	33.6	0.067	0.066	33.2	0.073	0.063	37.0
Pruning at $\rho = 0.6$	0.049	0.049	54.3	0.050	0.050	51.9	0.066	0.047	53.1
Pruning at $\rho = 0.8$	0.027	0.022	88.8	0.172	0.022	86.7	0.644	0.020	79.3
PCA at 99% of variance	0.089	0.090	19.6	0.090	0.090	19.5	0.095	0.083	25.1
PCA at 99.9% of variance	0.075	0.075	26.1	0.075	0.076	25.6	0.079	0.069	32.6

SD of estimates, mean SEs, and empirical power based on the 95% confidence interval for 10,000 simulated datasets using six approaches for selecting genetic variants. Results are also given on rounding the association estimates to a fixed number of decimal places.

represents a null causal effect. We also set the mean of the distribution of the associations with the outcome as 0.1 times the associations with the risk factor, representing a causal effect of 0.1. We simulated 10,000 datasets for each value of the causal effect, and calculated the Mendelian randomization estimate using the same six approaches for variant selection as above. Additionally, we repeated the analyses but first rounding the genetic associations (and their SEs) to three and two decimal places.

4.4 | Results

Results are presented in Table 3 for the SD of estimates, the mean SE, and the empirical power of the 95% confidence interval (the proportion of datasets in which the confidence interval excluded the null; this is the Type 1 error rate for a null causal effect). The mean estimates (not presented) were

close to the true causal effect throughout for all approaches. As in the previous simulations, estimates from the pruning approach became more precise as the threshold correlation increased, although Type 1 error rates were above nominal levels for $\rho = 0.8$ even when the association estimates were not rounded. Rounding exacerbated false-positive findings, and inflated Type 1 error rates were present in all methods when associations were rounded to two decimal places. Coverage rates were least affected when pruning at a threshold correlation of $\rho = 0.2$ or 0.4 and for the PCA approaches. With a positive causal effect, power increased as the threshold increased, although judging estimators by power estimates is misleading when Type 1 error rates are inflated. Power of the PCA approaches was similar to that using a pruning threshold of $\rho = 0.2$, but lower than that at a threshold of $\rho = 0.4$, and was greater using principal components that explained a greater proportion of the weighted correlation matrix.

WILEY-

$\frac{8}{1}$ Wiley 5 | discussion

As the cost of high-density genome sequencing continues to fall, additional signals are likely to be identified within known loci. There will be growing demand for methods to exploit correlated instruments in Mendelian randomization, as the addition of correlated variants can improve power to detect a causal effect. In this paper, we first connected previously known results together to show from theoretical arguments that genetic variants included in a Mendelian randomization analysis should be those that are associated with the risk factor in a conditional analysis. If the variants are combined in an allele score, then the conditional (multivariable) associations with the risk factor should be used as weights in the allele score to obtain the most efficient analysis. If only summarized data are available, then the same analysis can be replicated with the marginal (univariable) associations using an extension to the IVW method to account for correlations between variants.

However, difficulties arise when there are many correlated genetic variants in a single gene region that are associated with the risk factor (fine-mapping genetic data). Including too few genetic variants in an analysis means that estimates are less precise, but also highly variable, in that different approaches to choosing variants can lead to markedly different estimates. However, including too many variants can lead to numerical instabilities and overly precise estimates with inflated Type 1 error rates. These numerical instabilities are not computational issues, but arise due to inconsistencies in the data: for example, if association estimates are rounded to a fixed number of decimal places, or if association or correlation estimates are obtained in different samples. It is difficult in practice to judge at what threshold these numerical issues begin to occur, although in the simulation examples considered, problems regularly occurred when pruning variants at a threshold correlation of 0.8 ($r^2 = 0.64$), and occasionally occurred at a threshold correlation of 0.6 ($r^2 = 0.36$). We note as well that r^2 is not always a good measure of correlation between genetic variants; near-singular matrices can occur when the pairwise correlations measured by r^2 are low, but there are haplotypes represented in the data, or when the minor allele frequencies of variants differ, but a common variant "tags" a rare variant (high D-prime, but low r^2).

As an alternative approach, we have proposed a method for selecting instruments based on PCA of a weighted version of the genetic correlation matrix. This approach constructs instruments as linear combinations of genetic variants. As the linear combinations are orthogonal, the approach does not suffer as much with respect to numerical instabilities. Additionally, the method incorporates data on all the genetic variants into the analysis, and consequently causal estimates from the approach are less variable. Estimates from the PCA approach are less precise than those from the variable selection approaches considered here (GCTA and pruning); however, they are less variable with respect to choices of how to implement the analysis (in particular the choice of variants).

5.1 | Comparison with previous work

The IVW method presented here is a simple application of generalized weighted linear regression, and is not unique to Mendelian randomization. The same method has been used in a variety of contexts including discovery genetics (Zhu et al., 2016), and prediction and model selection (Chen et al., 2015; Benner et al., 2016; Newcombe, Conti, & Richardson, 2016). A number of different solutions have been proposed to the problem of highly correlated variants, including pruning and clumping at a threshold correlation, and adding a small positive number to the diagonal of the correlation matrix (Gusev et al., 2016). In the applied example of the paper at a correlation threshold of $\rho = 0.8$, adding 0.1 to the diagonal of the correlation matrix changed the causal estimate from -0.137 (SE (0.031) to -0.065 (0.057). Although the substantial change in the causal estimate is indicative of near-singular behavior, it would seem preferable for estimation to simply use a stricter correlation threshold rather than misspecifying the correlation matrix (and better still to use the principal component approach presented in this manuscript).

We believe that Mendelian randomization differs somewhat from other analysis contexts, as an instrumental variable analysis relies on inferences from a single-gene region (e.g., for a protein risk factor where the gene region is the coding region for the risk factor) or a small number of gene regions. Another feature of Mendelian randomization is the prevalence of the summarized data and two-sample settings, in which discrepancies in genetic associations are likely to arise.

PCAs have been suggested before for fine-mapping data, with Wallace demonstrating that 70% of the variance in the genetic correlation matrix could be explained by an average of seven components for 49 test gene regions (Wallace, 2013). A key innovation here is weighting the genetic correlation matrix, meaning that principal components with the greatest eigenvalues will be those that explain the most variance in the risk factor. This means that it is more likely that an analysis based on a small number of principal components will have reasonable power to detect a causal effect. For example, if there is only one causal variant in the gene region, then 100% of the variance would be explained by one principal component, even if there were other uncorrelated variants in the gene region.

We advocate the PCA method proposed in this paper as a worthwhile approach to analyze fine-mapped genetic data for Mendelian randomization. It provides estimates that may be less precise compared with those from variable selection approaches such as GCTA, but are more robust to seemingly arbitrary choices in the variable selection step.

ACKNOWLEDGMENTS

We would like to thank Paul Newcombe and Chris Wallace (both MRC Biostatistics Unit, Cambridge) for comments on an earlier version of this manuscript. Stephen Burgess and Verena Zuber are supported by Sir Henry Dale Fellowship jointly funded by the Wellcome Trust and the Royal Society (grant number 204623/Z/16/Z). Jemma C Hopewell is supported by a British Heart Foundation Basic Science Research Fellowship (grant number FS/14/55/ 30806).

ORCID

Stephen Burgess (b) http://orcid.org/0000-0001-5365-8760

REFERENCES

- Angrist, J. D., & Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association*, 90(430), 431–442.
- Benner, C., Spencer, C. C., Havulinna, A. S., Salomaa, V., Ripatti, S., & Pirinen, M. (2016). FINEMAP: Efficient variable selection using summary data from genome-wide association studies. *Bioinformatics*, 32(10), 1493–1501.
- Borenstein, M., Hedges, L., Higgins, J., & Rothstein, H. (2009). Generality of the basic inverse-variance method. Introduction to metaanalysis (chap. 34). Chichester, West Sussex, UK: Wiley, 311–319.
- Bowden, J., Davey Smith, G., Haycock, P. C., & Burgess, S. (2016). Consistent estimation in Mendelian randomization with some invalid instruments using a weighted median estimator. *Genetic Epidemiol*ogy, 40(4), 304–314.
- Burgess, S., Butterworth, A. S., & Thompson, S. G. (2013). Mendelian randomization analysis with multiple genetic variants using summarized data. *Genetic Epidemiology*, 37(7), 658–665.
- Burgess, S., Dudbridge, F., & Thompson, S. G. (2015). Re: "Multivariable Mendelian randomization: The use of pleiotropic genetic variants to estimate causal effects." *American Journal of Epidemiology*, 181(4), 290–291.
- Burgess, S., Dudbridge, F., & Thompson, S. G. (2016). Combining information on multiple instrumental variables in Mendelian randomization: Comparison of allele score and summarized data methods. *Statistics in Medicine*, 35(11), 1880–1906.
- Burgess, S., Scott, R., Timpson, N., Davey Smith, G., Thompson, S. G., & EPIC-InterAct Consortium. (2015). Using published data in Mendelian randomization: A blueprint for efficient identification of causal risk factors. *European Journal of Epidemiology*, 30(7), 543– 552.
- Burgess, S., & Thompson, S. G. (2015). Mendelian randomization: Methods for using genetic variants in causal estimation. Boca Raton, FL: Chapman & Hall.
- Burgess, S., Thompson, S. G., & CRP CHD Genetics Collaboration. (2011). Avoiding bias from weak instruments in Mendelian randomization studies. *International Journal of Epidemiology*, 40(3), 755– 764.

- Cai, M., Dai, H., Qiu, Y., Zhao, Y., Zhang, R., Chu, M., ... Chen, F. (2013). SNP set association analysis for genome-wide association studies. *PLOS One*, 8(5), e62495.
- CARDIoGRAMplusC4D Consortium. (2015). A comprehensive 1000 genomes-based genome-wide association meta-analysis of coronary artery disease. *Nature Genetics*, *47*, 1121–1130.
- Chen, W., Larrabee, B. R., Ovsyannikova, I. G., Kennedy, R. B., Haralambieva, I. H., Poland, G. A., & Schaid, D. J. (2015). Fine mapping causal variants with an approximate Bayesian method using marginal test statistics. *Genetics*, 200(3), 719–736.
- Coviello, A. D., Haring, R., Wellons, M., Vaidya, D., Lehtimäki, T., Keildson, S., ... Perry, J. R. (2012). A genome-wide association meta-analysis of circulating sex hormone–binding globulin reveals multiple Loci implicated in sex steroid hormone regulation. *PLoS Genetics*, 8(7), e1002805.
- Davey Smith, G., & Ebrahim, S. (2003). "Mendelian randomization": Can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology*, *32*(1), 1–22.
- Didelez, V., Meng, S., & Sheehan, N. (2010). Assumptions of IV methods for observational epidemiology. *Statistical Science*, 25(1), 22–40.
- Didelez, V., & Sheehan, N. (2007). Mendelian randomization as an instrumental variable approach to causal inference. *Statistical Meth*ods in Medical Research, 16(4), 309–330.
- Gusev, A., Ko, A., Shi, H., Bhatia, G., Chung, W., Penninx, B. W., ... Pasaniuc, B. (2016). Integrative approaches for large-scale transcriptome-wide association studies. *Nature Genetics*, 48(3), 245– 252.
- Inoue, A., & Solon, G. (2010). Two-sample instrumental variables estimators. *Review of Economics and Statistics*, 92(3), 557– 561.
- Jin, G., Sun, J., Kim, S.-T., Feng, J., Wang, Z., Tao, S., ... Xu, J. (2012). Genome-wide association study identifies a new locus JMJD1C at 10q21 that may influence serum androgen levels in men. *Human Molecular Genetics*, 21(23), 5222–5228.
- Johnson, T. (2013). Efficient calculation for multi-SNP genetic risk scores. Technical report, The Comprehensive R Archive Network. Retrieved from http://cran.r-project.org/web/packages/ gtx/vignettes/ashg2012.pdf
- Lawlor, D., Harbord, R., Sterne, J., Timpson, N., & Davey Smith, G. (2008). Mendelian randomization: Using genes as instruments for making causal inferences in epidemiology. *Statistics in Medicine*, 27(8), 1133–1163.
- Newcombe, P. J., Conti, D. V., & Richardson, S. (2016). JAM: A scalable Bayesian framework for joint analysis of marginal SNP effects. *Genetic Epidemiology*, 40(3), 188–201.
- Pierce, B., Ahsan, H., & Vander Weele, T. (2011). Power and instrument strength requirements for Mendelian randomization studies using multiple genetic variants. *International Journal of Epidemiology*, 40(3), 740–752.
- Wallace, C. (2013). Statistical testing of shared genetic control for potentially related traits. *Genetic Epidemiology*, 37(8), 802–813.
- Winkelried, D., & Smith, R. J. (2011). Principal components instrumental variable estimation. Cambridge Working Papers in Economics 1119. Technical report, Faculty of Economics, University of Cambridge.

WILEV

- Wooldridge, J. (2009). Instrumental variables estimation and two stage least squares. *Introductory econometrics: A modern approach (chap.* 15). Nashville, TN: South-Western Publishing, 512–552.
- Yang, J., Ferreira, T., Morris, A., Medland, S. E., Genetic Investigation of ANthropometric Traits (GIANT) Consortium, DIAbetes Genetics Replication And Meta-analysis (DIAGRAM) Consortium, ... Visscher, P. M. (2012). Conditional and joint multiple-SNP analysis of GWAS summary statistics identifies additional variants influencing complex traits. *Nature Genetics*, 44(4), 369–375.
- Zhu, Z., Zhang, F., Hu, H., Bakshi, A., Robinson, M. R., Powell, J. E., ... Yang, J. (2016). Integration of summary data from GWAS and eQTL studies predicts complex trait gene targets. *Nature Genetics*, 48, 481–487.

How to cite this article: Burgess S, Zuber V, Valdes-Marquez E, Sun B, Hopewell JC. Mendelian randomization with fine-mapped genetic data: Choosing from large numbers of correlated instrumental variables. *Genet Epidemiol.* 2017; 00:1–12. https://doi.org/10.1002/gepi.22077

APPENDIX A

A.1 Software code

We provide R code to implement the methods discussed in this paper. The genetic variants are represented by g (a matrix of allele counts for the genetic variants), the risk factor by x, and the outcome by y. Weights for the allele score are represented by wts. The associations of the candidate instruments with the risk factor are denoted betaXG with SEs sebetaXG. The associations of the candidate instruments with the outcome are denoted betaYG with SEs sebetaXG. With a continuous outcome, these associations are usually estimated using linear regression; with a binary outcome, using logistic regression.

The two-stage least squares (2SLS) method can be implemented using the *sem* package:

library(sem)

Genetic variants can be collapsed into an allele score, and the score can be used in the 2SLS method:

If the genetic variants are perfectly uncorrelated, and the weights are the coefficients from univariable regression analyses of the risk factor on each of the genetic variants in turn, then these two analyses are equivalent.

IVW estimate (ignoring correlation):

Although fixed-effects models are used throughout this paper, the (multiplicative) random-effects analysis is preferred when heterogeneity between the causal estimates from each genetic variant is expected (provided that there are enough genetic variants in the model to obtain a reasonable estimate of the heterogeneity). Heterogeneity would generally be expected when using genetic variants from multiple gene regions that may have different mechanisms of influencing the risk factor, but not when using multiple variants in the same gene region that should have similar mechanisms of effect.

IVW estimate (accounting for correlation):

Omega	= sebetaYG%o%sebetaYG*rho
beta_IVWcorrel	<pre>= solve(t(betaXG)%*%solve(Omega)%*%betaXG)*t(betaXG)%*%solve(Omega)%*%betaYG</pre>
<pre>se_IVWcorrel.fixed</pre>	<pre>= sqrt(solve(t(betaXG)%*%solve(Omega)%*%betaXG))</pre>
resid	<pre>= betaYG-beta_IVWcorrel*betaXG</pre>
<pre>se_IVWcorrel.random</pre>	<pre>= sqrt(solve(t(betaXG)%*%solve(Omega)%*%betaXG))*</pre>
	<pre>max(sqrt(t(resid)%*%solve(Omega)%*%resid/(length(betaXG)-1)),1)</pre>

The matrix rho comprises the pairwise correlations between the genetic associations (in particular, the genetic associations with the outcome). Provided that are genetic associations estimated in the same participants, these are equal to the correlations between the genetic variants themselves.

IVW estimate (accounting for correlation) using principal components:

Phi = (betaXG/sebetaYG)%o%(betaXG/sebetaYG)*rho

```
summary(prcomp(Phi, scale=FALSE))
```

```
K = which(cumsum(prcomp(Phi, scale=FALSE)$sdev^2/sum((prcomp(Phi, scale=FALSE)$sdev^2))>0.99)[1]
```

K is number of principal components to include in analysis

```
# this code includes principal components to explain 99% of variance in the risk factor
```

```
betaXG0 = as.numeric(betaXG%*%prcomp(Phi, scale=FALSE)$rotation[,1:K])
```

```
betaYG0 = as.numeric(betaYG%*%prcomp(Phi, scale=FALSE)$rotation[,1:K])
```

```
Omega = sebetaYG%o%sebetaYG*rho
```

```
pcOmega = t(prcomp(Phi, scale=FALSE)$rotation[,1:K])%*%Omega%*%prcomp(Phi, scale=FALSE)$rotation[,1:K]
beta_IVWcorrel.pc = solve(t(betaXG0)%*%solve(pcOmega)%*%betaXG0)*t(betaXG0)%*%solve(pcOmega)%*%betaYG0
se_IVWcorrel.fixed.pc = sqrt(solve(t(betaXG0)%*%solve(pcOmega)%*%betaXG0))
```

The IVW method accounting for correlation can also be performed using the standard linear regression command after weighting the data by the Cholesky decomposition:

Omega = sebetaYG%o%sebetaYG*rho c_betaXG = solve(t(chol(Omega)))%*%betaXG c_betaYG = solve(t(chol(Omega)))%*%betaYG

```
beta_IVWcorrel = lm(c_betaYG~c_betaXG-1)$coef[1]
se_IVWcorrel.fixed = sqrt(1/(t(betaXG)%*%solve(Omega)%*%betaXG))
se_IVWcorrel.random = sqrt(1/(t(betaXG)%*%solve(Omega)%*%betaXG))*max(summary(lm(c_betaYG~c_betaXG-1))$sigma,1)
```

A.2 Proof of equality of 2SLS and IVW estimates

A.2.1 Variants uncorrelated

If the we write the risk factor as X (usually an $N \times 1$ matrix, although the result can be generalized for multiple risk factors), the outcome as Y (an $N \times 1$ matrix), and the instrumental variables as Z (an $N \times J$ matrix), then the two-stage least squares estimate of causal effects is:

$$\hat{\beta}_{2SLS} = [X^T Z (Z^T Z)^{-1} Z^T X]^{-1} X^T Z (Z^T Z)^{-1} Z^T Y.$$

This estimate can be obtained by sequential regression of the risk factor on the instrumental variables, and then the outcome on fitted values of the risk factor from the first-stage regression.

Regression of Y on Z gives beta-coefficients $\hat{\beta}_Y = (Z^T Z)^{-1} Z^T Y$ with SEs the square roots of the diagonal elements of the matrix $(Z^T Z)^{-1} \sigma^2$ where σ is the residual SE. If the instrumental variables are perfectly uncorrelated, then the off-diagonal

11

WILEN

elements of $(Z^T Z)^{-1} \sigma^2$ are all equal to zero. Regression of X on Z gives beta-coefficients $\hat{\beta}_X = (Z^T Z)^{-1} Z^T X$. Weighted linear regression of the beta-coefficients $\hat{\beta}_Y$ on the beta-coefficients $\hat{\beta}_X$ using the inverse-variance weights $(Z^T Z)\sigma^{-2}$ gives an estimate:

$$\begin{split} & [\hat{\beta}_X^T(Z^T Z)\hat{\beta}_X]^{-1} \sigma^{-2} \hat{\beta}_X^T(Z^T Z) \sigma^2 \hat{\beta}_Y \\ = & [X^T Z(Z^T Z)^{-1}(Z^T Z)(Z^T Z)^{-1} Z^T X]^{-1} X^T Z(Z^T Z)^{-1} (Z^T Z)(Z^T Z)^{-1} Z^T Y \\ = & [X^T Z(Z^T Z)^{-1} Z^T X]^{-1} X^T Z(Z^T Z)^{-1} Z^T Y \\ = & \hat{\beta}_{2SLS} \end{split}$$

The assumption of uncorrelated instrumental variables ensures that the regression coefficients from univariate regressions (as in the regression-based methods) equal those from multivariable regression (as in the two-stage least squares method).

A.2.2 Variants correlated

WILEY

12

If the variants are correlated, then the same argument holds, except that the weights in the weighted linear regression of the beta-coefficients $\hat{\beta}_X$ on the beta-coefficients $\hat{\beta}_X$ are $(Z^T Z)P\sigma^{-2}$, where *P* is the (symmetric) correlation matrix.