

Expression reflects population structure

Brielin C Brown^{1,†}, Nicolas L. Bray^{+,2,3} and Lior Pachter^{+,4}

1. Verily Life Sciences
2. Institute for Innovative Genomics, UC Berkeley
3. Department of Molecular & Cell Biology, UC Berkeley
4. Division of Biology and Biological Engineering, California Institute of Technology

+ Correspondence should be addressed to: nicolas.bray@berkeley.edu, lpachter@caltech.edu.

† Work performed while at UC Berkeley

Abstract

Population structure in genotype data has been extensively studied, and is revealed by looking at the principal components of the genotype matrix. However, no similar analysis of population structure in gene expression data has been conducted, in part because a naïve principal components analysis of the gene expression matrix does not cluster by population. We identify a linear projection that reveals population structure in gene expression data. Our approach relies on the coupling of the principal components of genotype to the principal components of gene expression via canonical correlation analysis. Furthermore, we analyze the variance of each gene within the projection matrix to determine which genes significantly influence the projection. We identify thousands of significant genes, and show that a number of the top genes have been implicated in diseases that disproportionately impact African Americans.

Author Summary

High dimensional, multi-modal genomics datasets are becoming increasingly common, which warrants investigation into analysis techniques that can reveal structure in the data without overfitting. Here, we show that the coupling of principal component analysis to canonical correlation analysis offers an efficient approach to exploratory analysis of this kind of data. We apply this method to the GEUVADIS dataset of genotype and gene expression values of European and Yoruban individuals, finding as-of-yet unstudied population structure in the gene expression values. Moreover, many of the top genes identified by our method have been previously implicated in diseases that disproportionately impact African Americans.

Introduction

Genes mirror geography to the extent that in homogeneous populations, individuals can be localized to within hundreds of kilometers purely on the basis of their genotype [1,2,3]. Population structure in genotypes is revealed via projection of single nucleotide polymorphism (SNP) data onto the first few principal components of the population-genotype matrix. The principal components space, which is a lower-dimensional distinguished subspace of the high-dimensional data, is computed by a procedure called principal components analysis (PCA). While PCA has been successful in revealing population structure from SNP data, it does not identify such structure in some other genomic data types. For example, in the case of gene expression data, PCA has not revealed obvious population signatures (Figure 1A, [4]). Here we show that although the first two principal components of expression data do not capture population structure, there are other projections that do. One approach to finding such projections is the coupling of dimension reduction to correlation maximization after adjustment for confounding. This approach, utilizing PCA and canonical correlation analysis (CCA), has been used to effectively analyze the relationship between gene expression and copy number variation [5]. The method is implementable via singular value decomposition and is therefore also efficient. We apply it to finding population structure in expression data, thereby further highlighting the combination of PCA and CCA as a powerful approach to integrative analysis of genomics data.

As an optimization procedure, PCA can be viewed as the projection of data onto the lower-

dimension subspace that minimizes the average distance of the data to its projection. This is algebraically equivalent to finding the lower-dimensional subspace that maximizes the variance of the projected data. This statistical view of PCA helps to explain why PCA of expression data might not reveal population structure: even if such structure is present in the data, it may not lie on the directions of maximal variance. CCA is a widely used method for joint analysis of heterogeneous data and provides a linear-algebraic mechanism for identifying shared structure among a pair of datasets. Given a pair of data matrices, CCA finds maximally correlated linear combinations of the columns of each matrix [6]. We show that CCA applied to the PCA projections of expression and genotype data identifies a projection of the expression data that reveals population structure, which we validate via a leave-one-out cross validation experiment.

To validate our method, we examined population structure in expression data from the Genetic European Variation in Health and Disease (GEUVADIS) project [7], which consists of RNA-seq data obtained from lymphoblastoid cell lines derived from whole-genome sequenced individuals belonging to five distinct populations. The GEUVADIS data has been extensively studied [7], yet our analysis reveals structure not previously examined in this well-characterized dataset.

Results

A naïve PCA analysis of the GEUVADIS expression data (Figure 1A) shows that unlike genotype data (Figure 1B), there is no clear clustering of individuals by population. This result is consistent with other analyses of expression data, in which population structure is not detected by PCA [4]. To understand the sources of variation that could explain the first and second principal component axes, we labeled the individuals according to the lab where they were sequenced (Figure 1C). This provides some insight into the sources of variation. For example, samples from Lab 3 are distinctly separated from Lab 1. We therefore proceeded to correct for confounding by regressing the gene expression matrix on a matrix of potentially confounding variables and taking the residual (Figure 1D—F, see Methods). We note that it is also possible to correct for confounding using CCA by exploiting the relationship between CCA with categorical data and linear discriminant analysis [8] (Supplementary Methods, Supplementary Figure 1).

Next, we examined whether coupling of expression data to genotype data could identify a projection that reveals population structure. To do this we performed PCA followed by CCA on the batch-corrected expression matrix and the genotype matrix (see Methods). The resulting CCA projection of expression data (Figure 2A), reveals distinct population patterns in the data, although not as clearly as the resulting PCA of the genotype data (Figure 1B). We validate this observation by performing a leave-one-out cross-validation experiment, where we remove each individual from the dataset and show that the reconstruction error of the model on the held out point is close to the error in the training set, and that the principal components of the reconstructed gene expression matrix show similar population patterns (Figure 2B). Moreover, the CCA projection is indexed by linear combinations of genes, which can be understood to discriminate individuals based on expression signatures. That is, genes with high variance in the CCA expression projection (see Supplemental Methods) have expression distributions that segregate based on patterns in the genotype PCs, which we interpret to represent population structure [1,2,3]. After correction for correlated multiple testing using the Benjamini–Hochberg–Yekutieli procedure [9], we identified 3,581 genes with significant scores at FDR 5%, indicating that population structure within gene expression data is pervasive.

The three genes with largest z-score in this analysis were LATS-2, EIF4EBP2 and STX7 (Figure 3). All three genes display increased expression in the YRI population, highlighting the possibility of relating population expression differences directly to disease phenotypes such as discussed in [10], where EIF4B is shown to be associated with vascular disease in African Americans. In addition, LATS-2 is a known tumor suppressor gene, and it is well known that African Americans have significantly worse outcomes at every stage of cancer treatment [11]. Other notable genes identified with this method include PSPH, NEDD4-2 and ARHGEF11 (all with $p < 1e-07$). PSPH

was examined in [12] where it was found to be the gene with the highest degree of differential allelic expression. The eQTL rs6700, which is associated with expression of that gene, is an ancestry informative marker [13]. NEDD4-2, which is associated with hypertension [14], contains SNPs such as rs945508 which are similarly associated with ancestry. Allelic variants in ARHGEF11 are associated with kidney disease in mouse models [15], which has higher prevalence in the African-American community. While we view the identification of such genes as important, we caution that African-Americans also experience substantial structural inequality in healthcare, which confounds this analysis [16]. We also note that while genes such as NEDD4-2 must also have substantial genetic/epigenetic regulation that is linked to population differences, the projection-associated genes identified by our method does not produce that information. Indeed, its power to detect genes associated with population structure comes by virtue of requiring only one test per gene and being agnostic to the source of regulation. While a complete analysis of population-associated expression differences is beyond the scope of this paper, these examples suggest that our method should be a powerful approach for directly identifying genes whose expression associates with population.

Conclusion

The identification of population structure in expression data suggests that it should be interesting to extend population genetic methods such as [17] to population transcriptomics. The example of joint analysis of expression and genotype data can be extended to include other datatypes via an extension of CCA to more than two matrices [8,18,19,20], and the coupling of PCA to CCA could also be extended to a hierarchical factor analysis method. Importantly, the coupling of PCA and CCA is not the only projection that reveals population structure. For example, connecting the principal components using linear regression gives similar visualizations (Supplementary Methods, Supplementary Figure 2). The choice of model should reflect the variance structure of the data, which here we have deliberately remained agnostic to. Ultimately, it is important to identify the optimal model for inference.

While we believe the extensions describe above will be interesting to pursue, our analysis and that in [5] shows that PCA+CCA is a useful and rapid approach to exploratory analysis of heterogeneous data. As the generation of large-scale, high-dimensional, multi-modal genomics datasets becomes more commonplace [20, 21, 22], we expect the combination of PCA and CCA to become as common as PCA is today.

Methods

We obtained genotype data of the Phase 1 1000 genomes individuals in PLINK format [23] from cog-genomics [See Data and Software Availability]. GEUVADIS project RNA-seq reads were downloaded from the European Nucleotide Archive (accession number ENA: ERP001942). In the analyses performed we omitted the CEU population because it has been previously found to display biased expression patterns due to the age of the cell line [24]. Importantly, this bias affects every CEU sample and therefore cannot be corrected for traditional methods of handling confounding.

There are 343 individuals with genotype data from 1000 genomes phase 1 and corresponding RNA-seq data from GEUVADIS in the FIN, GBR, TSI and YRI populations. We quantified the transcript abundances of these individuals using kallisto [25] with the GENCODE v27 protein coding transcript sequences and annotation. The GENCODE v27 annotation contains 95,659 transcripts. We omitted all transcripts with mean transcripts per million (TPM) less than 0.1 across the quantified samples, leaving 58,012 transcripts. We then used the GENCODE v27 annotation to obtain gene level quantifications by summing transcript quantifications in TPM units. Finally, we removed genes in the MHC region and on non-autosomal chromosomes. This left 14,070 genes for analysis. The Phase 1 1000 genomes genotypes contain 39,728,178 variants. We filtered indels, variants with minor allele frequency (MAF) less than 5%, and non-biallelic

SNPs leaving 6,785,201 SNPs for analysis. Finally, we quantile-normalized the expression matrix, and centered and scaled each gene quantification vector to have mean 0 and variance 1. In the following analyses, we chose to keep 27 principal components of expression and 11 principal components of genotype, while analyzing the first two canonical components. However we note that our results are stable under different choices of numbers of components (Supplementary Figure 3, Supplementary Table 1).

To remove batch effects from the expression matrix, we one-hot encoded the lab identification vector, and then added a column for sample gender [26], resulting in a 343 x 7 matrix of potentially confounding variables. We then regressed each gene expression vector on the confounding matrix and used the residual expression vector for all further analysis. Next, we computed principal components of the genotype matrix using PLINK and principal components of the corrected expression matrix using the eigendecomposition of the sample correlation matrix. Finally, we computed the canonical variables between the top principal components of the genotype and corrected expression matrices (see the Supplementary Methods for details on the linear algebra).

To verify that we did not over-fit in estimating coefficients using CCA, we performed leave-one-out cross validation. We removed each of the 343 individuals one-by-one from the dataset, re-calculated the principal components of the genotype and expression matrices, and re-estimated the canonical variables and bases. We then projected each held out individual into the resulting CCA gene expression subspace. After this process, for each individual, we plotted the first two principal components of the re-constructed expression matrix to verify the individual clusters by population (see the Supplementary Methods for details of the how the projection was performed). Furthermore, we calculated the in-sample and out-of-sample reconstruction error as the squared Frobenius norm of the original and reconstructed data points, and verified that it was similar for both left-in and held-out samples.

Finally, we asked which genes had significant variance for the CCA gene expression projection. We computed the variance of each gene in the projection, and calculated significance via a permutation test with 10 million permutations. In each iteration, we shuffled the genotype principal components and recomputed the variance explained. The p -value derived from this test is the number of times the permuted score is greater than the observed score, divided by the number of permutations (see the Supplementary Methods for details of how the variance is computed). We further estimated a z-score for each gene as the difference between the estimated and mean permutation variance divided by the variance of the permuted variance.

Data and software availability

The software used to produce the analyses is on GitHub. We provide a package of tools for computing the projections and estimating gene significance, as well as a Snakemake file [27] that can be used to completely reproduce the analysis, from data acquisition to figure generation.

- Analysis software:
<https://github.com/pachterlab/PCACCA/>
- Gencode v27 transcripts:
ftp://ftp.sanger.ac.uk/pub/gencode/Gencode_human/release_27/gencode.v27.pc_transcripts.fa.gz
- Gencode v27 GTF:
ftp://ftp.sanger.ac.uk/pub/gencode/Gencode_human/release_27/gencode.v27.annotation.gtf.gz
- GEUVADIS RNA-seq reads:
<ftp://ftp.sra.ebi.ac.uk/vol1/ERA169/ERA169774/fastq>
- 1000 genomes genotypes:
https://www.dropbox.com/s/k9ptc4kep9hmvz5/1kg_phase1_all.tar.gz

Acknowledgements

The authors would like to thank Shannon McCurdy for invaluable feedback on this manuscript.

References

1. Cavalli-Sforza LL, Menozzi P, Piazza A. The history and geography of human genes. Princeton university press; 1994.
2. Novembre, John, et al. "Genes mirror geography within Europe." *Nature* 456.7218 (2008): 98-101.
3. Yang, Wen-Yun, et al. "A model-based approach for analysis of spatial structure in genetic data." *Nature genetics* 44.6 (2012): 725-731.
4. Stranger, Barbara E., et al. "Patterns of cis regulatory variation in diverse human populations." *PLoS genetics* 8.4 (2012): e1002639.
5. Sonesson, C et al. "Integrative analysis of gene expression and copy number alterations using canonical correlation analysis." *BMC Bioinformatics* 11 (2010):191.
6. Hotelling, Harold. "Relations between two sets of variates." *Biometrika* 28.3-4 (1936): 321-377.
7. Lappalainen, Tuuli, et al. "Transcriptome and genome sequencing uncovers functional variation in humans." *Nature* 501.7468 (2013): 506-511.
8. Bach, Francis and Michael Jordan. "A Probabilistic Interpretation of Canonical Correlation Analysis." *TR 688, University of California, Berkeley* (2005)
9. Benjamini, Yoav and Daniel Yekutieli. "The control of the false discovery rate in multiple testing under dependency." *Annals of Statistics*. 29.4 (2001): 1165-1188.
10. Wei, Peng et al. "Differential endothelial cell gene expression by African Americans versus Caucasian Americans: a possible contribution to health disparity in vascular disease and cancer." *BMC Medicine* 9:2 (2011).
11. Long, B., F. W. Liu, and R. E. Bristow. "Disparities in uterine cancer epidemiology, treatment, and survival among African Americans in the United States." *Gynecologic oncology* 130.3 (2013): 652-659.
12. Morley, Michael, et al. "Genetic analysis of genome-wide variation in human gene expression." *Nature* 430.7001 (2004): 743-747
13. Marcus, Joseph H., and John Novembre. "Visualizing the geography of genetic variants." *Bioinformatics* 33.4 (2017): 594-595.
14. Shi, Peijun P., et al. "Salt-sensitive hypertension and cardiac hypertrophy in mice deficient in the ubiquitin ligase Nedd4-2." *American Journal of Physiology-Renal Physiology* 295.2 (2008): F462-F470.
15. Williams, Jan M., et al. "Genetic variants in Arhgef11 are associated with kidney injury in the Dahl salt-sensitive rat." *Hypertension* (2012): HYPERTENSIONAHA-112.
16. Hill, Shirley A. *Inequality and African-American health: How racial disparities create sickness*. Policy Press, 2016.
17. Novembre, John, and Benjamin M. Peter. "Recent advances in the study of fine-scale population structure in humans." *Current opinion in genetics & development* 41 (2016): 98-105.
18. Tenenhaus, Arthur, and Michel Tenenhaus. "Regularized generalized canonical correlation analysis." *Psychometrika* 76.2 (2011): 257-284.
19. Witten, Daniela and Robert J. Tibshirania. "Extensions of Sparse Canonical Correlation Analysis with Applications to Genomic Data." *Statistical Applications in Genetics and Molecular Biology* 8.1 (2009): 1-27.
20. Butler, Andrew et al. "Integrating single-cell transcriptomic data across different conditions, technologies and species." *Nature Biotechnology* (2018).
21. Todorovic, Vesna "Gene expression: Single-cell RNA-seq—now with protein." *Nature Methods* volume 14, pages 1028–1029 (2017)
22. Chen, Xingqi et al. "Joint single-cell DNA accessibility and protein epitope profiling reveals environmental regulation of epigenomic heterogeneity" *bioRxiv* 310359

23. Purcell, Shaun, et al. "PLINK: a tool set for whole-genome association and population-based linkage analyses." *The American Journal of Human Genetics* 81.3 (2007): 559-575. Yuan, Yuan, Lei Tian, Dongsheng Lu, and Shuhua Xu. "Analysis of genome-wide RNA-sequencing data suggests age of the CEPH/Utah (CEU) lymphoblastoid Cell lines systematically biases gene expression profiles." *Scientific reports* 5 (2015).
24. Yuan, Yuan, et al. "Analysis of genome-wide RNA-sequencing data suggests age of the CEPH/Utah (CEU) lymphoblastoid cell lines systematically biases gene expression profiles." *Scientific reports* 5 (2015): 7960.
25. Bray, Nicolas L., et al. "Near-optimal probabilistic RNA-seq quantification." *Nature biotechnology* 34.5 (2016): 525-527.
26. Whitney, Adeline et al. "Individuality and variation in gene expression patterns in human blood" *Proceedings of the National Academy of Sciences* 100 (4) (2003): 1896-1901
27. Köster, Johannes, and Sven Rahmann. "Snakemake—a scalable bioinformatics workflow engine." *Bioinformatics* 28.19 (2012): 2520-2522.

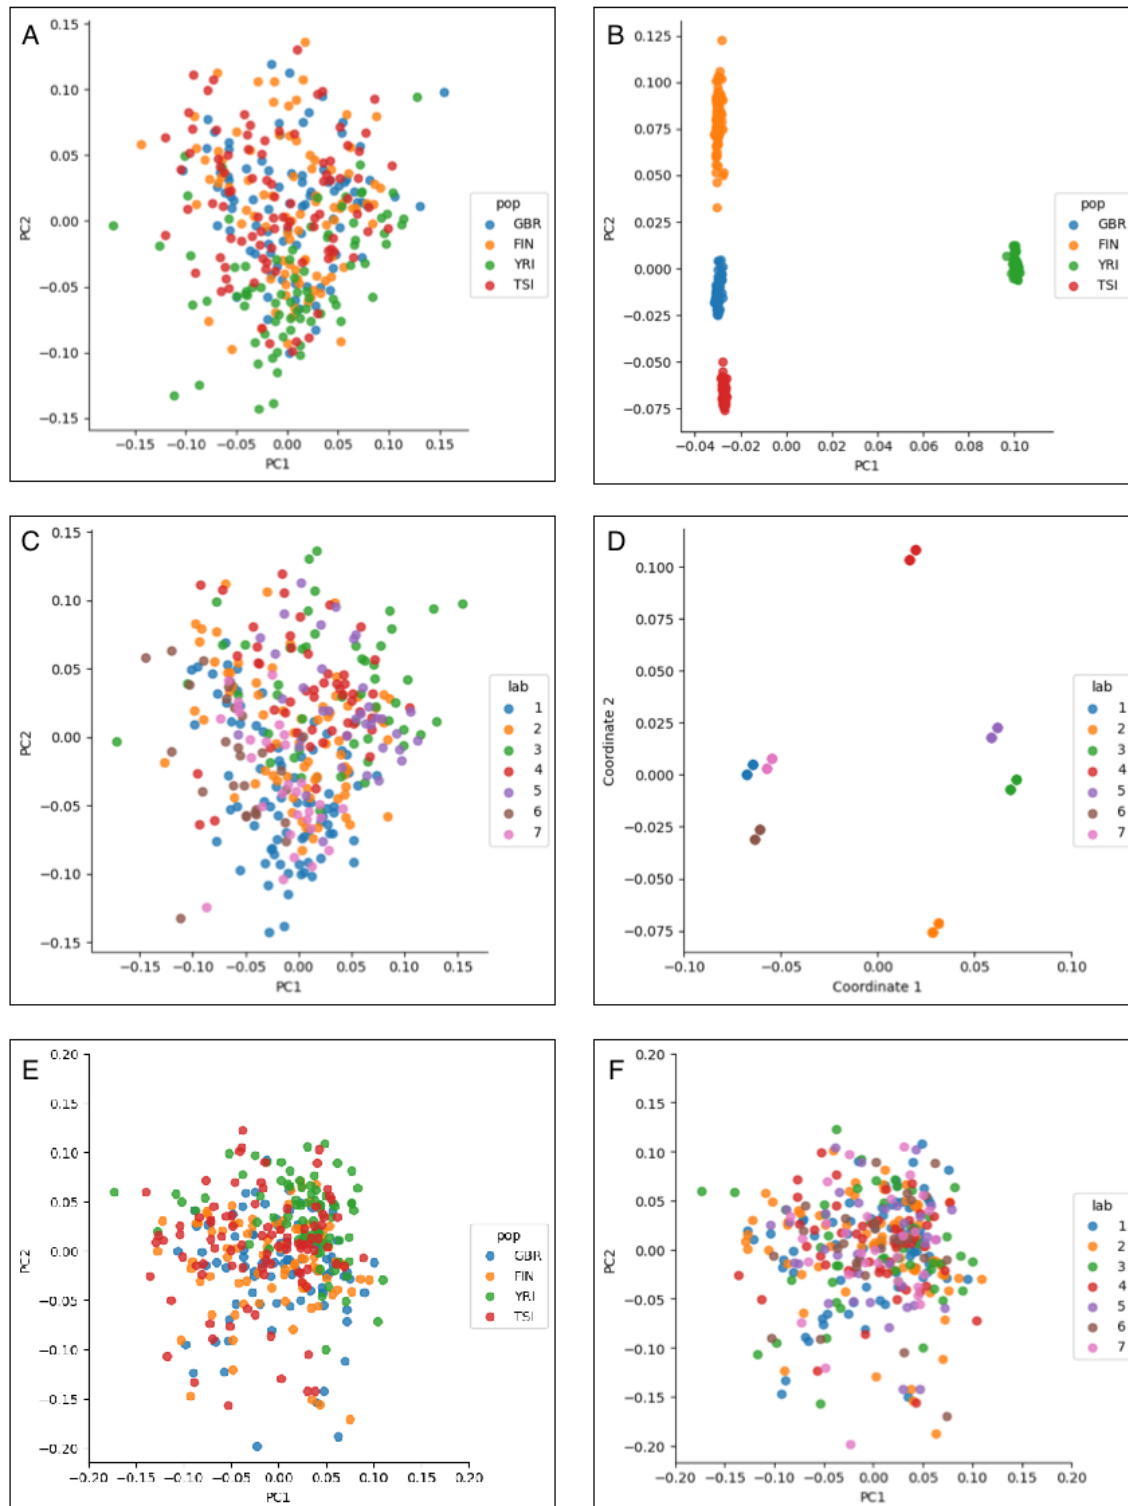


Figure 1: (A) PCA of the expression matrix fails to reveal clustering by population, whereas (B) PCA of the genotype matrix reveals clear clustering by population. (C) Coloring of samples by batch reveals that PC1 and PC2 are being partly defined by batch source. (D) Projection into the space learned by a regression model of gene expression from batch and gender identity reveals strong clustering in the data. (E) PCA of the batch-corrected expression matrix. (F) The corrected data no longer cluster by batch.

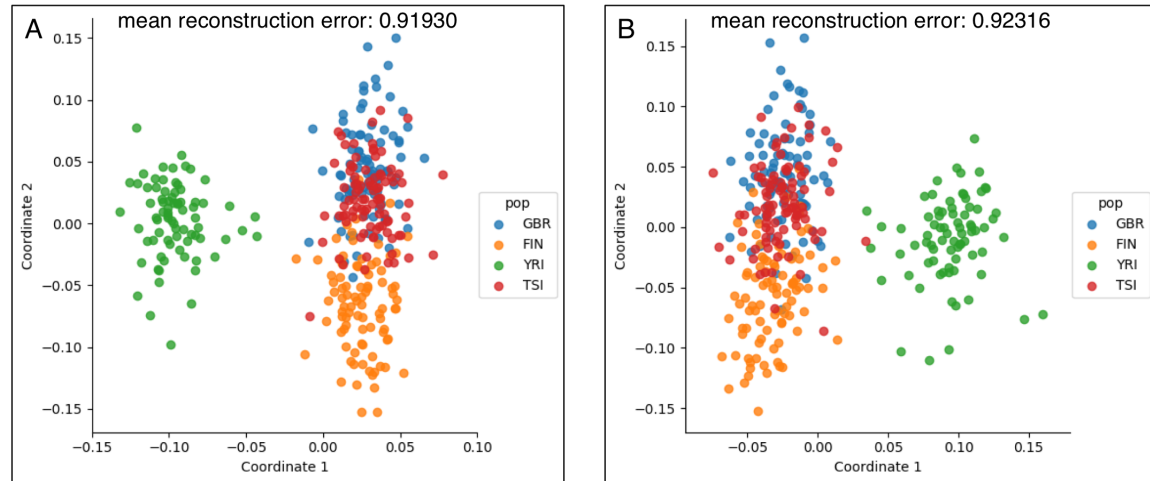


Figure 2: (A) A CCA projection of the batch-corrected expression matrix that shows that expression reflects population structure. While the individuals, labeled according to their population, do not cluster as clearly as with genotype data (Figure 1B), there is clear population structure in the CCA projection of the batch-corrected expression data. (B) A leave-one-out cross-validation experiment showing that individuals are approximately projected to their populations of origin even when the projection matrix is learned without their expression or genotype data. The mean re-construction errors in (A) the left-in samples and (B) the held-out samples are similar and overlaid on top of the Figure. The first two canonical correlations are 0.963 and 0.766.

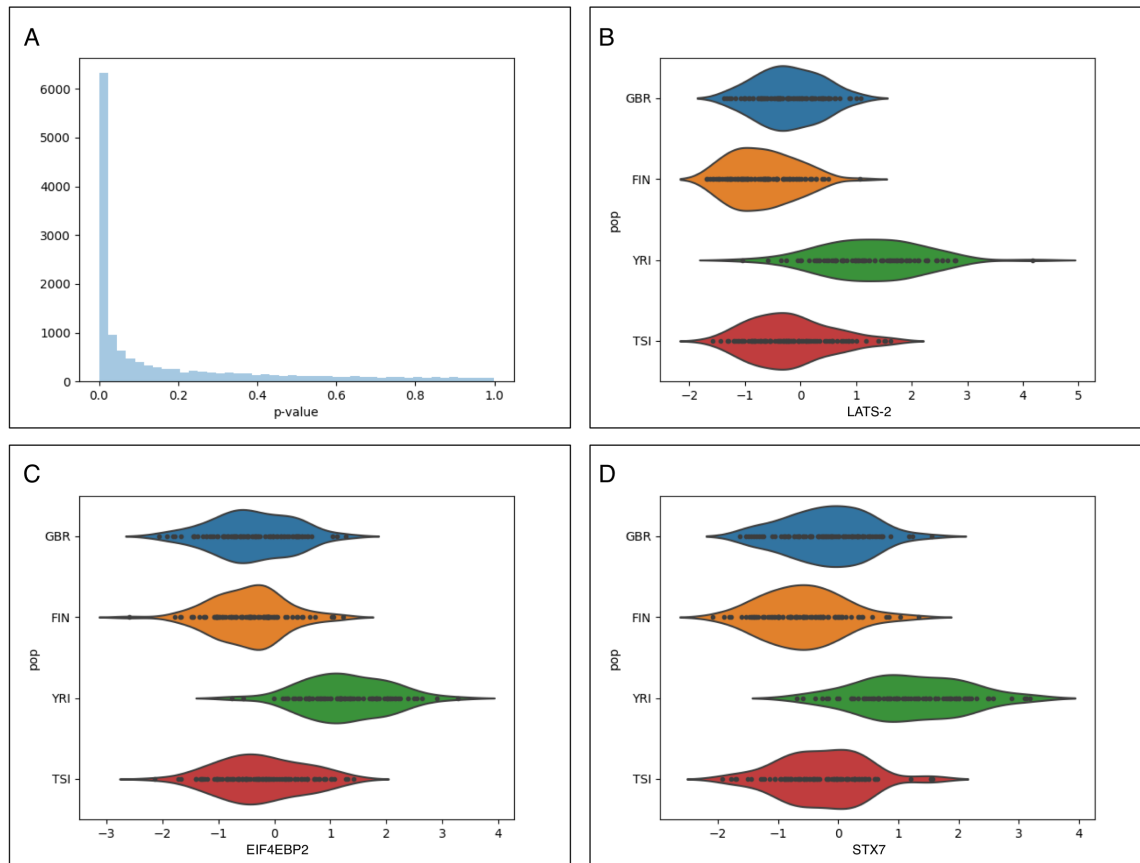


Figure 3: (A) The p -value distribution for tests that the variance of each gene in the projection is greater than the null shows a large number of genes with significant scores in the PCA+CCA projection. The expression distributions by population for the three genes with highest z-scores are shown in (B) the LATS-2 gene, (C) the EIF4EBP2 gene, (D) the STX7 gene.

Supplementary Methods

Expression reflects population structure

Brielin C Brown, Nicolas Bray and Lior Pachter

July 6, 2018

1 Principal Component Analysis

Let X be an $N \times D$ matrix of rank $r \leq N < D$. The singular value decomposition of X is a factorization of the form $X = U\Lambda V^\top$, where U and V are $N \times r$ and $D \times r$ unitary matrices, respectively, and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_r)$ is a diagonal matrix containing the singular values of X . The columns of U and V are called the left and right singular vectors, respectively [?]. While keeping the full spectrum of singular values gives an exact decomposition of X , one can instead keep only the top $l < r$ singular values, using U_l and V_l to denote the first l columns of U and V , respectively, and Λ_l to denote the first l singular values. This gives an approximation to X which we call $\tilde{X} = U_l \Lambda_l V_l^\top \approx X$.

Let $\Sigma = \frac{X^\top X}{N-1}$ be the covariance of X . We use $A \cdot_\Sigma B = A\Sigma^+ B$ to indicate matrix multiplication with respect to the Mahalanobis metric with covariance Σ . When X is column-mean subtracted, the principal components of X are the l -dimensional coordinates defined by the rows of U_l . These coordinates have basis $F = \frac{V_l \Lambda_l}{\sqrt{N-1}}$ which are orthogonal with respect to the Mahalanobis metric on X . We note that this definition of principal component analysis (PCA) is slightly different than the classical definition. Specifically, the classical formulation takes $U_l \Lambda_l$ as its coordinates and V_l as the orthonormal Euclidean basis [?]. However in genetics the basis scaling we use here is more common. Specifically, if X is a genotype matrix this is equivalent to defining the principal components as the eigenvectors of the realized relationship matrix as in [?, ?].

Then projection of X into the space spanned by F and its complement are given by

$$\begin{aligned} \tilde{X} &= X \cdot_\Sigma F F^\top \\ &= X \left(\frac{X^\top X}{N-1} \right)^+ F F^\top \\ &= U \Lambda V^\top V \Lambda^{-2} V^\top V_l \Lambda_l^2 V_l^\top \\ &= U_l \Lambda_l V_l^\top. \\ \tilde{X}^\perp &= X - \tilde{X}. \end{aligned}$$

Next, we extend these concepts to canonical correlation analysis before combining them to define explicitly the method we use.

2 Canonical Correlation Analysis

Given two data matrices X and Y with the same number of rows representing distinct but related data, canonical correlation analysis (CCA) finds maximally correlated linear combinations of the columns of X and Y . CCA identifies matrices A and B such that the sequence of correlations $\rho_{XY,i} = \text{corr}(XA^{(i)}, YB^{(i)})$, where $A^{(i)}$ is the i th column of matrix A , is successively maximized so that the correlation matrix $\rho_{XY} = \text{corr}(XA, YB)$ is diagonal [?].

Specifically, let X and Y be $N \times D_X$ and $N \times D_Y$ data matrices of rank r_X and r_Y respectively. As in the previous section, we denote the singular value decompositions of X and Y by $X = U_X \Lambda_X V_X^\top$ and $Y = U_Y \Lambda_Y V_Y^\top$. The canonical correlations of X and Y are given by the singular values of the matrix $M = U_X^\top U_Y = U_M \rho_{XY} V_M^\top$. The canonical bases are given by $F_X = \frac{V_X \Lambda_X U_M}{\sqrt{N-1}}$ and $F_Y = \frac{V_Y \Lambda_Y V_M}{\sqrt{N-1}}$, again orthonormal with respect to the Mahalanobis metric on Σ_X and Σ_Y respectively [?]. As with PCA, one can instead keep only the top k singular vectors, which we denote by $U_{M,k}$ and $V_{M,k}$ to indicate the first k columns of U_M and V_M respectively. In this case, the bases are $F_{X,k} = \frac{V_X \Lambda_X U_{M,k}}{\sqrt{N-1}}$ and $F_{Y,k} = \frac{V_Y \Lambda_Y V_{M,k}}{\sqrt{N-1}}$.

The coordinates of the data matrices in k -dimensional CCA space, which we refer to as the canonical variables, are given by

$$\begin{aligned} C_{X,k} &= X \cdot_{\Sigma_X} F_{X,k} = U_X U_{M,k}, \\ C_{Y,k} &= Y \cdot_{\Sigma_Y} F_{Y,k} = U_Y V_{M,k}. \end{aligned}$$

Similarly, the projections of the data matrices onto the canonical bases are given by

$$\begin{aligned} \tilde{X}_{C,k} &= X \cdot_{\Sigma_X} F_{X,k} F_{X,k}^\top = U_X U_{M,k} U_{M,k}^\top \Lambda_X V_X^\top, \\ \tilde{Y}_{C,k} &= Y \cdot_{\Sigma_Y} F_{Y,k} F_{Y,k}^\top = U_Y V_{M,k} V_{M,k}^\top \Lambda_Y V_Y^\top. \end{aligned}$$

As with PCA the complement is given by

$$\begin{aligned} \tilde{X}_{C,k}^\perp &= X - \tilde{X}_{C,k} = U_X (\mathbb{I} - U_{M,k} U_{M,k}^\top) \Lambda_X V_X^\top, \\ \tilde{Y}_{C,k}^\perp &= Y - \tilde{Y}_{C,k} = U_Y (\mathbb{I} - V_{M,k} V_{M,k}^\top) \Lambda_Y V_Y^\top. \end{aligned}$$

From this, it is straightforward to interpret linear discriminant analysis (LDA) as a special case of CCA [?]. Let Y be a matrix of data observations and let L be a length N categorical data vector with K categories. Consider the $N \times K - 1$ indicator matrix X with the i th column the indicator vector $\mathcal{I}[L = i]$, sometimes called the one-hot encoding of the data. Then LDA between X and Y is equivalent to learning the CCA projection $Y_{C,k}$, and one can project out effects correlated with the categories in L by computing $Y_{C,k}^\perp$. Indeed, we can replace the regression correction for batch and gender in the main text with this approach and obtain similar visualizations. See Supplementary Figure 1.

We combine PCA with CCA as in [?] to arrive at our final projection. Specifically, we work with the first l_X PCA components of X and the first l_Y PCA components of Y . In this case, $M = U_{X,l_X}^\top U_{Y,l_Y}$ such that the coordinates in k -dimensional CCA space are

$$\begin{aligned} C'_X &= U_{X,l_X} U_{M,k}, \\ C'_Y &= U_{Y,l_Y} V_{M,k}. \end{aligned}$$

with bases

$$\begin{aligned} F'_X &= \frac{V_{X,l_X} \Lambda_{X,l_X} U_{M,k}}{\sqrt{N-1}}, \\ F'_Y &= \frac{V_{Y,l_Y} \Lambda_{Y,l_Y} V_{M,k}}{\sqrt{N-1}}. \end{aligned}$$

As before, the projection of the data matrices onto the canonical bases and their complements are given by

$$\begin{aligned} \tilde{X}_{C'} &= X \cdot_{\Sigma_X} F'_X F'^\top_X = U_{X,l_X} U_{M,k} U_{M,k}^\top \Lambda_{X,l_X} V_{X,l_X}^\top, \\ \tilde{Y}_{C'} &= X \cdot_{\Sigma_Y} F'_Y F'^\top_Y = U_{Y,l_Y} V_{M,k} V_{M,k}^\top \Lambda_{Y,l_Y} V_{Y,l_Y}^\top, \\ \tilde{X}_{C'}^\perp &= X - \tilde{X}_{C'} \\ \tilde{Y}_{C'}^\perp &= Y - \tilde{Y}_{C'} \end{aligned}$$

3 Leave-One-Out Cross-Validation

Let the genotype matrix of the sample be X and the expression matrix of the sample be Y . For individual i , the full matrices X and Y are reduced by removing row i to create X_{-i} and Y_{-i} . We then learn the canonical bases F'_X , F'_Y of the data matrices using the combination of PCA and CCA described above. Next we project the held out individual gene vector y^i into this space to get $y^i_C = y^i \cdot_{\Sigma} F'_{Y,-i} F'^{\top}_{Y,-i}$. After repeating this for all individuals, we form the data matrix Y'_C where each row i is the projection of individual i 's gene expression vector into the CCA-gene space learned without using i . A plot of the first two principal components of this data matrix shows that the population structure learned by this method is valid with respect to held-out samples (Figure 2, Supplementary Figure 3).

4 Determining the Significance of a Gene for the CCA Projection

Let $F'_{Y,(j)}$ be the j 'th row in the basis matrix F'_Y . The variance v_j for gene j is given by $v_j = \|F'_{Y,(j)}\|_2^2$. Therefore, to determine whether the variance of a gene is non-null, we can perform a permutation test. Specifically, in each permutation p we shuffle the genotype principal components and compute the permuted variance for each gene v_j^p . The p -value for the test that the variance is greater than the null score is then the number of times the permuted variance is greater than the true variance. That is, for N_p permutations, the p -value is $p_j = \sum_p \mathbb{1}[v_j^p > v_j] / N_p$.

5 A regression based approach

While performing CCA between the projected data matrices is one approach to visualizing the common structure underlying the principal components of two data matrices, there are possibly many projections that yield such a result. One involves assuming that the principal components of one matrix have a linear relationship to the principal components of the other with additive noise. That is if, $U_Y = U_X \beta + \epsilon$, the least squares solution is

$$\hat{\beta} = U_X^{\top} U_Y.$$

This time, the coordinates of the data matrix in the projection are

$$\begin{aligned} C_{Y,k} &= U_X \hat{\beta} = U_X U_X^{\top} U_Y, \\ C_{X,k} &= U_X. \end{aligned}$$

Compared to the CCA approach, the regression approach presumes a different noise model for the data. Rather than modeling the principal components of both gene expression and genotype as linearly related to an underlying hidden factor with additive noise, the regression approach implicitly models the genotype principal components as the underlying factor that influences gene expression principal components with additive noise. Ultimately, the choice of model should reflect the ‘‘noise’’ structure emerging from the underlying biology and the nature of the measurements. One notable drawback of the regression based approach is that it uses the location of the genotype in genotype-PCA space to approximate the location of the expression vector in expression-PCA space, as opposed to the CCA projection which only needs to know the global relationship between the points in expression-PCA and genotype-PCA space. A consequence of this is that any held-out individual needs to have their genotype projected into genotype-PCA space, which is computationally infeasible when working with millions of SNPs. Therefore, we have omitted the leave-one-out cross-validation for this model. However the visualization obtained when applying this model is depicted in Supplementary Figure 2.

6 Supplementary Figures and Tables

	27, 11	13, 5	42, 15	100, 2
27, 11	1	0.691	0.925	0.890
13, 5	0.691	1	0.666	0.753
42, 15	0.925	0.666	1	0.823
100, 2	0.890	0.753	0.823	1

Table 1: Cross-correlation matrix of the Z-scores obtained for each gene across the different choices of components presented in Figure S3.

References

- [1] Roger A Horn, Charles R Johnson Matrix analysis, 2nd ed. ed. Cambridge University Press, New York, 2013.
- [2] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. The elements of statistical learning. Vol. 1. New York: Springer series in statistics, 2001.
- [3] Alkes L Price, Nick J Patterson, Robert M Plenge, Michael E Weinblatt, Nancy A Shadick, and David Reich. Principal components analysis corrects for stratification in genome-wide association studies. *Nature genetics*, 38(8):904–909, 2006.
- [4] Jian Yang, S Hong Lee, Michael E Goddard, and Peter M Visscher. GCTA: a tool for genome-wide complex trait analysis. *The American Journal of Human Genetics*, 88(1):76–82, 2011.
- [5] Michael J Greenacre. Theory and applications of correspondence analysis. 1984
- [6] Francis R Bach and Michael I Jordan A probabilistic interpretation of canonical correlation analysis. 2005
- [7] Charlotte Sonesson, Henrik Lilljebjörn, Thoas Fioretos, and Magnus Fontes Integrative analysis of gene expression and copy number alterations using canonical correlation analysis. *BMC Bioinformatics*, 11:191, 2010.

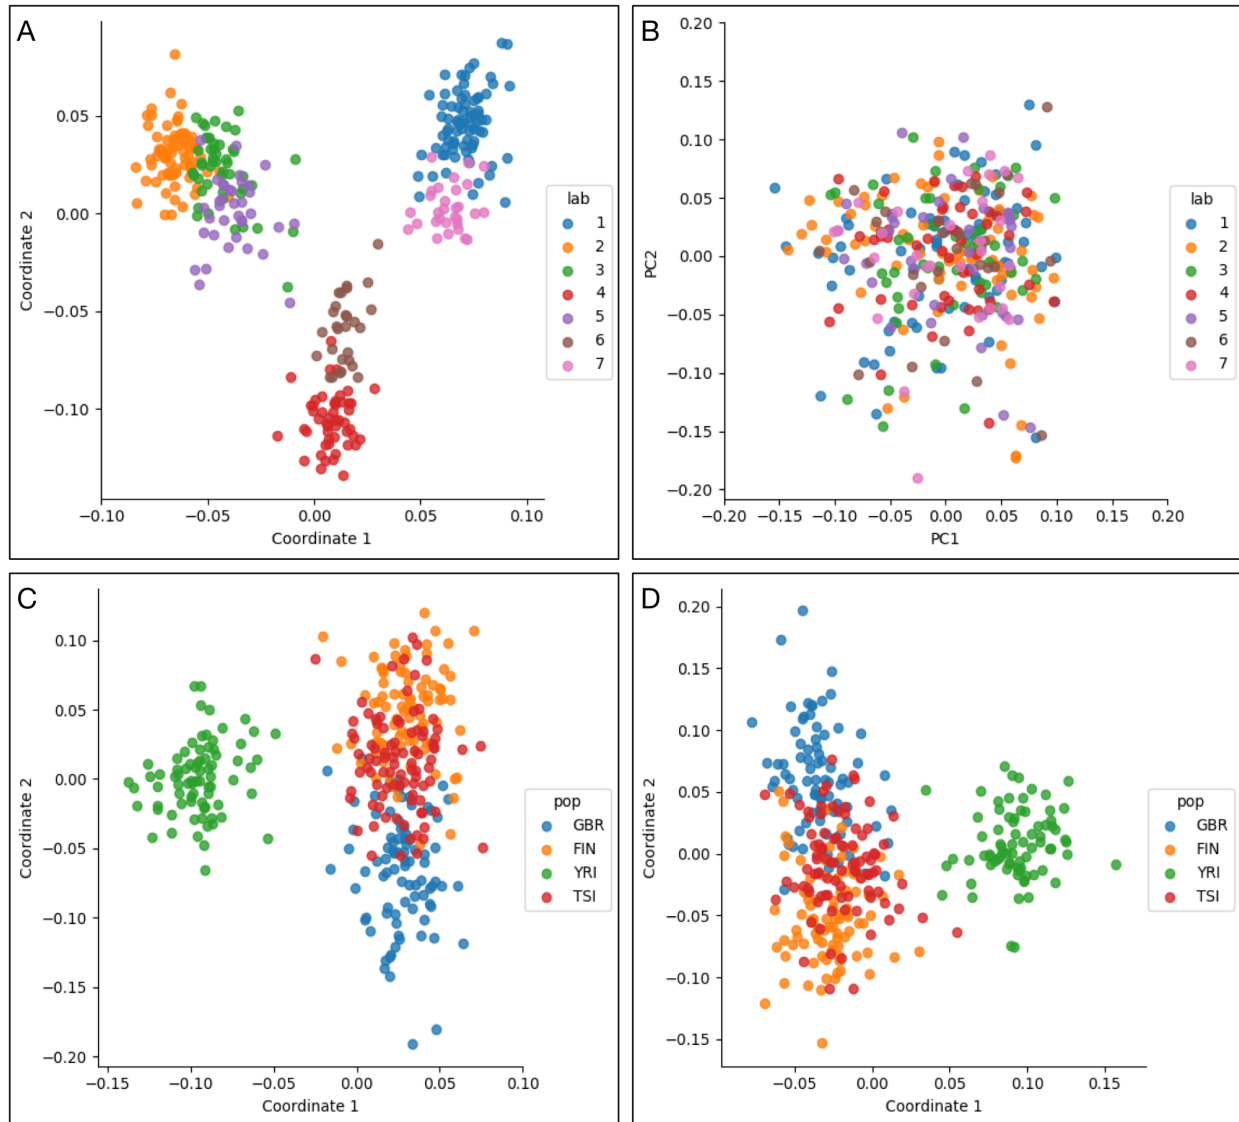


Figure 1: (A) CCA between PCs of expression and a confounding matrix is related to LDA, and projection into the learned space reveals strong clustering by batch within the data. (B) Projecting orthogonally to this space leaves samples scrambled by batch in the first PCs. (C) Using this correction instead of regression gives similar results for CCA between PCs of expression and genotype. (D) As in the main text, structure is maintained during a cross-validation experiment.




Figure 2: Using regression rather than CCA to relate the principal components of the two data matrices also yields a projection that reveals population structure within the expression data.

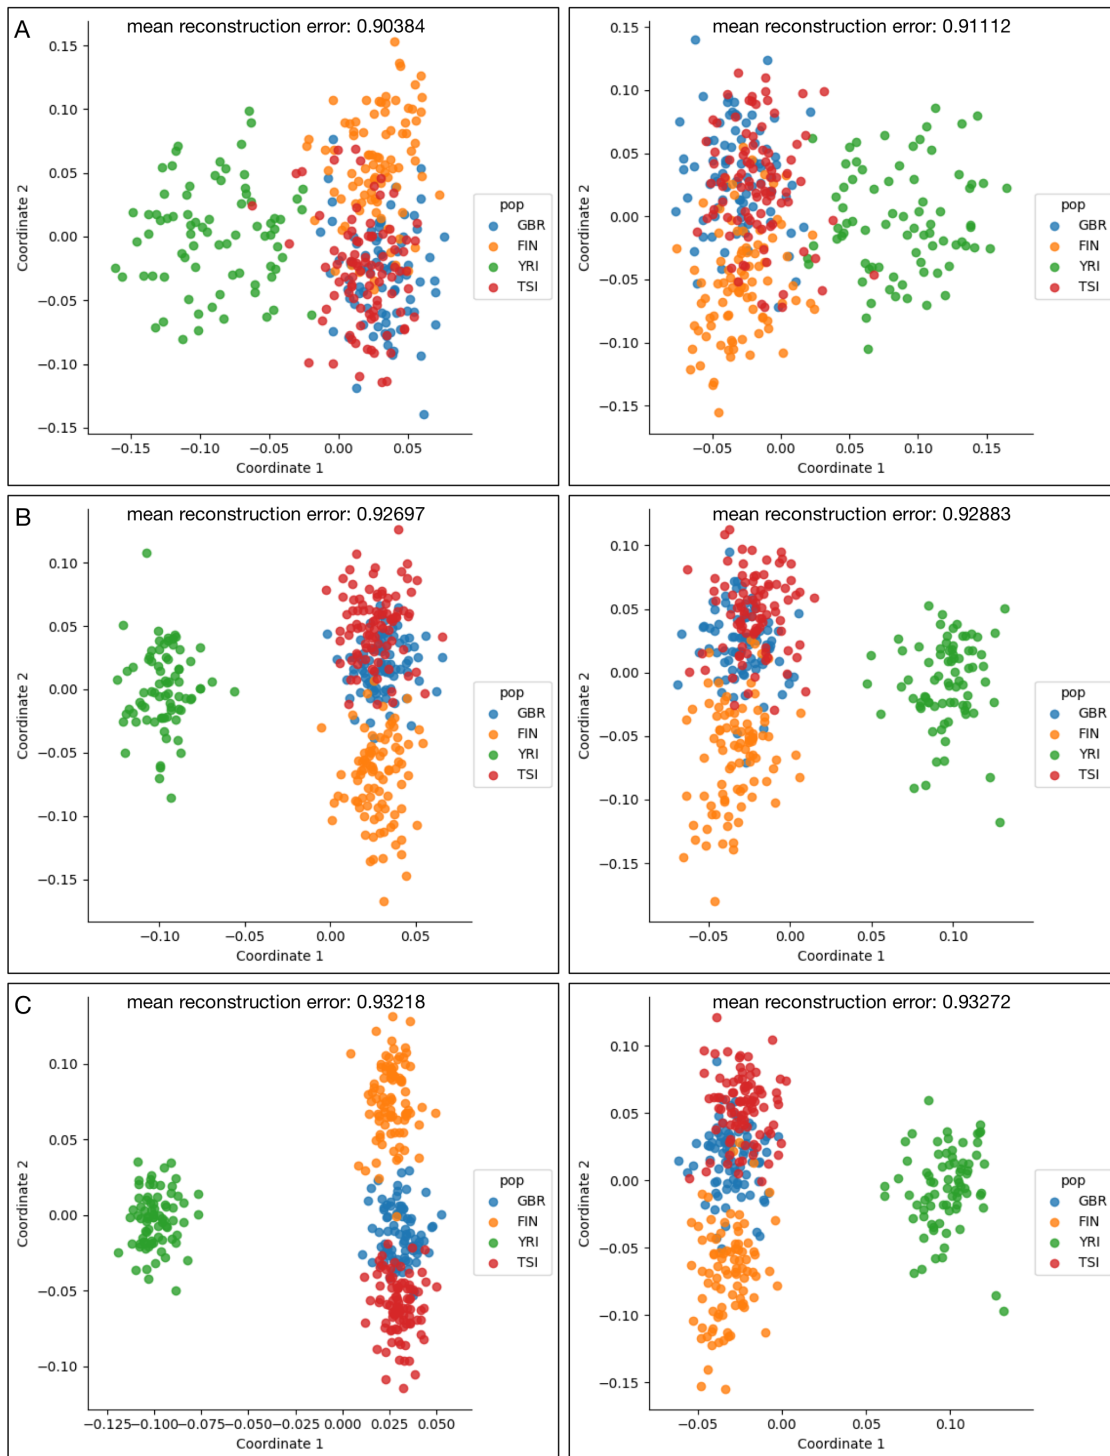


Figure 3: PCA and CCA of genotype and expression data with different numbers of principal components. (A) 13 expression and 5 genotype components. (B) 42 gene expression and 15 genotype components. (C) 100 expression and 2 genotype components.