Supplemental Methods

Genotyping and imputation

We extracted DNA from the two skeletal muscle biopsy samples. DNA samples were genotyped at the University of Michigan DNA sequencing core using the Infinium multi-ethnic global-8 v1 kit (Illumina, San Diego, CA, USA), along with 62 samples used in other projects. We downloaded the Illumina support files for the kit from https://support.illumina.com/downloads/infinium-multi-ethnic-global-8-v1-support-files.html on March 16 2020. We excluded multi mapping sites noted in the file Multi-EthnicGlobal_D1_MappingComment.txt and added variant rsIDs from the file Multi-EthnicGlobal_D1_b150_rsids.txt. We used the pre-imputation check toolkit developed by the McCarthy lab (script HRC-1000G-check-bim.pl, version 4.2.9 downloaded from https://www.well.ox.ac.uk/~wrayner/tools/, along with the 1000g legend file https://www.well.ox.ac.uk/~wrayner/tools/1000GP_Phase3_combined.legend.gz) to check strand, alleles, position and Ref/Alt assignments relative to the 1000g reference. We removed variants a) on ChrM, ChrXY, and ChrY, b) missing from 1000g or monomorphic in our set of samples or alleles mismatching with 1000g, c) palindromic SNPs with 1000g EUR MAF > 40%, d) MAF < 0.5% in 1000g EUR samples, e) with P-value for deviation from Hardy-Weinberg Equilibrium (HWE) < 1×10^{-6} (exact test implemented in VCFtools --hwe) in 1000g EUR samples and f) genotype missingness > 2.5% resulting in 409,966 variants. Ref and Alt alleles were reoriented relative to 1000g. We performed pre-phasing and imputation using the Michigan Imputation Server (Das et al., 2016). The standard pipeline (https://imputationserver.readthedocs.io/en/latest/pipeline/) included pre-phasing using Eagle2 (Loh et al., 2016) and genotype dosage imputation using Minimac4 (https://github.com/statgen/Minimac4) and the 1000g phase 3 v5 (build GRCh37/hg19) reference panel (The 1000 Genomes Project Consortium, 2015). Post-imputation, we selected biallelic variants with estimated imputation accuracy (r2) > 0.3, resulting in 7,764,973 variants.

snATAC-seq barcode correction

Barcodes were corrected in a similar manner as in the 10x Genomics Cell Ranger ATAC v. 1.0 software. In brief, barcodes were checked against the 10x Genomics whitelist. If a barcode was not on the whitelist, then we found all whitelisted barcodes within a hamming distance of two from the bad barcode. For each of these whitelisted barcodes, we calculated the probability that the bad barcode should be assigned to the whitelisted barcode using the Phred scores of the mismatched base(s) and the prior probability of a read coming from the whitelisted barcode (based on the whitelisted barcode's abundance in the rest of the data). If there was at least a 97.5% probability that the bad barcode was derived from one specific whitelisted barcode, it was corrected to the whitelisted barcode.

Nucleus sample assignments in libraries with two human samples

In snATAC-seq/snRNA-seq libraries containing nuclei from two human samples, we removed doublets and determined the sample of origin for singlets using demuxlet with the imputed

genotypes (demuxlet options '--field PL'; git commit b7453fc, modified as described in GitHub issue #15).

snRNA-seq datasets often show substantial levels of ambient RNA contamination, and this ambient RNA will reflect the most highly expressed genes in the common cell types and could impact demuxlet assignments; we therefore ran demuxlet twice for the snRNA-seq libraries, once using all genotyped positions and once masking those genotypes that overlap the top 30% of genes (by UMI count, using the 40k nuclei snRNA-seq library and gene coordinates from the hg19 GENCODE GTF). A nucleus was considered a singlet if it was called as a singlet (demuxlet 'BEST' assignment) in either of the two demuxlet runs.

Removal of ambient RNA counts from single nucleus gene expression matrices

Prior to clustering and downstream analysis, we used DecontX (Yang et al., 2020) (celda v. 1.6.1, in R v. 4.0.3 (R Core Team, 2020)) to adjust the nucleus x gene expression count matrices for ambient RNA. To generate initial cluster assignments for DecontX, we ran Seurat (Butler et al., 2018; Stuart et al., 2019) (v. 3.9.9.9010, in R v. 3.6.3) on the human and rat gene expression matrices separately, including all nuclei with gene expression (snRNA-seq or multiome libraries). For clustering of human nuclei, we split nuclei into two groups according to the individual from which they were derived; we then normalized and found the top 2000 variable features for nuclei from each individual, before integrating them using the FindIntegrationAnchors() and IntegrateData() functions (setting scale = False for FindIntegrationAnchors(), and dims = 1:20 for both functions). After integration the data was scaled and Louvain clustering was performed using the top 20 PCs, 10 nearest neighbors, Louvain resolution = 0.1, and 100 random starts. For rat clustering we scaled the data, computed the top 2000 variable features, normalized, and then clustered with the same parameters used for the human data. These clusters were passed to the decontX() function, and (rounded) decontaminated counts were used for clustering and all downstream analyses.

Clustering with LIGER

For LIGER clustering (LIGER v. 0.5.0) (Welch et al., 2019), we removed genes without signal in at least 10 nuclei from each input matrix (input matrices were the same as for Seurat, i.e. one RNA and one ATAC matrix per biological individual, with multiome library nuclei in the RNA matrix). We selected variable genes using the rat RNA and the HSM2 RNA matrices. For factorization, we used k = 10, lambda = 10, and nrep=4. For quantile normalization we set min_cells = 5, and knn_k = 20 for the human datasets and knn_k = 2 for the rat datasets (as fewer nuclei were present for rat; the quantile_norm function was customized for this purpose). We then determined clusters using the Louvain algorithm with resolution = 0.04.

Motif scores

Motif scores were generated using chromVAR (v. 1.8.0) (Schep et al., 2017) in R v. 3.6.0. To generate the peaks used as input, we merged the per-cluster broad peaks, then computed the center of those merged peaks and extended +/- 250 bps. To generate the motifs used as input, we scanned hg19 using FIMO (v. 4.11.2; default parameters, with background expectation calculated using the MEME suite's fasta-get-markov script on the hg19 FASTA file) (Grant et al.,

2011) with the human Cis-BP motif library (downloaded on August 6, 2018) (Weirauch et al., 2014). We ran the addGCBias() and computeDeviations() functions with default parameters. To determine which motifs to include in the heatmap displayed in Fig. S15, for each motif and each cell type cluster we computed the motif's mean deviation z-score across nuclei in that cluster. We then normalized this value across the cell type clusters (dividing by the maximum value across the clusters, and setting negative values to 0), such that each motif had a score between 0 and 1 in each cluster. With this value, we additionally calculated a "cell type specificity index" for each motif using the tissue specificity index equation from (Yanai et al., 2005) (resulting in a value between 0 and 1, where 0 indicates the motif mean deviation z-score is identical across all clusters and 1 indicates the motif mean deviation is positive only in one cluster and negative or 0 in the other clusters). We included in the heatmap all motifs with mean deviation z-score > 2 in at least one cluster, and cell type specificity > 0.85; these motifs should be those that show a considerably higher mean deviation in nuclei from a single cluster (or small subset of clusters) than in the remaining clusters.

UK Biobank GWAS

We downloaded UK Biobank GWAS summary statistics made available by the Benjamin Neale lab (v2 of their analysis, initially made public on August 1, 2018; http://www.nealelab.is/uk-biobank/) (Sudlow et al., 2015). Specifically, we downloaded the 'both sex' GWAS summary statistic files listed in the 'UKBB GWAS Imputed v3 - File Manifest Release 20180731' spreadsheet available at https://docs.google.com/spreadsheets/d/1kvPoupSzsSFBNSztMzl04xMoSC3Kcx3CrjVf4yBmESU/edit#gid=178908679 (downloaded on April 9, 2020). Because some traits may not be appropriate for such an enrichment analysis (because they are not strongly polygenic, because the phenotypes are untrustworthy, etc.), we kept only traits deemed as 'high confidence' and with estimated heritability > 0.01 (and z-score > 7) based on the Neale Lab's own LD score regression heritability analysis of the GWAS results. Their rating criteria are described on their UKBB LDSC GitHub page (https://nealelab.github.io/UKBB_ldsc/confidence.html) and their LD score regression results (with confidence ratings) were downloaded from https://www.dropbox.com/s/ipeqyhrpdqav5uh/ukb31063_h2_all.02Oct2019.tsv.gz?dl=1. For each trait, we used the 'primary' GWAS result, as indicated in that file. Any traits that did not have a combined male and female GWAS analysis were dropped.

Non-muscle cell type open chromatin annotations used in GWAS enrichments

To create the adipose open chromatin regions, we processed the three adipose ATAC-seq libraries from (Cannon et al., 2019). Adapter sequences were removed using cutadapt (v. 1.12) (Martin, 2011) before mapping to hg19 with BWA-MEM (-I 200,200,5000 -M). Duplicates were marked using Picard MarkDuplicates and BAM files were filtered using SAMtools view (-F 4 -F 256 -F 1024 -F 2048 -q 30) before converting to BED format (bamtools bamtobed) and calling peaks with MACS2 (--nomodel --shift -100 --seed 2018 --extsize 200 --broad --keep-dup all --SPMR). We then took the union of peaks across the three samples, keeping those merged peaks that appeared in at least two samples.

The beta cell ATAC-seq peaks were taken from (Rai et al., 2020) (peaks called using all beta cell nuclei).

Common open chromatin regions were derived from the DNase I hypersensitive sites from (Meuleman et al., 2020). The DHS index was downloaded from https://www.meuleman.org/DHS_Index_and_Vocabulary_hg38_WM20190703.txt.gz on March 21, 2020. We lifted open chromatin regions from hg38 to hg19 using liftOver with the chain file from http://hgdownload.cse.ucsc.edu/goldenpath/hg38/liftOver/hg38ToHg19.over.chain.gz (Hinrichs et al., 2006). We then kept those that were labeled as 'tissue invariant' and that appeared in at least 500 of the 733 samples.

We also used open chromatin regions from (Meuleman et al., 2020) for adrenal gland, bone, brain, eye, gonad, gum, H9 embryonic stem cells, heart, kidney, large intestine, liver, lung, mammary, mesoderm, ovary, placenta, prostate, skin, small intestine, spinal cord, spleen, stomach, and umbilical cord. For each tissue except the stem cells, we took the non-cancerous samples labeled 'Primary' from that tissue and kept those DNase I hypersensitive sites that appeared in at least 50% of the samples from that tissue. For the stem cells, we kept those DNase I hypersensitive sites that appeared in either of the two replicates from biosample 'H9_hESC'.

Re-processing of snATAC-seq data prior to allelic bias analysis

We used the WASP software package (commit 36c0e5f8b5) (van de Geijn et al., 2015) to re-process ATAC-seq data prior to the allelic bias analysis. Nuclei from each library, cluster, and each human sample were processed separately. Reads were filtered, before and after re-mapping, using SAMtools view with flags -f 3 -F 4 -F 8 -F 256 -F 2048 -q 30. Duplicates were removed using WASP's rmdup_pe.py script, modified to take into account the 10x nucleus barcode (such that only duplicates within a nucleus would be removed). Prior to gathering allele counts, overlapping read pairs were clipped using bamUtil's clipOverlap (v. 1.0.14) (Jun et al., 2015). Allele counts were computed using pysam's (https://github.com/pysam-developers/pysam; v. 0.16.0.1) pileup function with parameters (max_depth=9999999, min_base_quality=20, min_mapping_quality=30, ignore_overlaps=True). Allele counts were summed across all nuclei within a cluster (i.e., across human samples and libraries).

Potential of re-alignment to GRCh38 to affect conclusions

We believe that re-aligning the reads to GRCh38 would not significantly affect the conclusions presented in this paper for the following reasons:

The great majority of the analyses presented here are genome-wide in nature (e.g., clustering with gene read counts and examining GWAS enrichments in open chromatin regions) and are therefore unlikely to substantially change due to a small number of genomic regions that change between the hg19 and hg38 reference genomes, particularly as we perform stringent read mapping quality filtering and remove ATAC-seq peaks from ENCODE blacklisted regions prior to

downstream analyses (thereby removing the genomic regions most likely to be problematic and shift between reference genome versions).

In the cases where we do focus on single genomic loci, we also cite previously-published data that broadly support our conclusions, and/or perform a follow-up assay ourselves for validation. At the *ITPR2* locus, we cite previously-published luciferase and EMSA data supporting our claim that our SNP of interest lies in an enhancer and may be bound by AP-1; at the *ARL15* locus we perform an allelic luciferase assay ourselves validating the claimed enhancer activity and allelic effect.

Anonymization of shared data

BAM files were anonymized using anonymizeBAM (with --strict option; v. 0.4.5) (Ziegenhain and Sandberg, 2021), and following anonymization all BAM tags other than those related to the nucleus barcodes/UMIs were removed using SAMtools. BAM files were then converted to CRAM format (Hsi-Yang Fritz et al., 2011) for sharing using SAMtools, and were converted to single-ended BED format using BEDTools bamtobed. Shared ATAC-seq peaks and ATAC-seq coverage files were generated from the BED files using MACS2 callpeak (options --SPMR --nomodel --shift -100 --extsize 200 -B --broad --keep-dup all). Shared RNA-seq coverage files were generated using pybedtools (Dale et al., 2011) (a python wrapper around BEDTools; BedTool().genome_coverage() function, with arguments bg=True, split=True, and strand and scale set appropriately to produce both reverse and forward strand bedGraph files scaled to 1 million reads) followed by conversion of bedGraph files to bigWig files using bedGraphToBigWig (v. 4) (Kent et al., 2010).

References

Butler, A., Hoffman, P., Smibert, P., Papalexi, E., and Satija, R. (2018). Integrating single-cell transcriptomic data across different conditions, technologies, and species. Nat. Biotechnol. *36*, 411–420.

Cannon, M.E., Currin, K.W., Young, K.L., Perrin, H.J., Vadlamudi, S., Safi, A., Song, L., Wu, Y., Wabitsch, M., Laakso, M., et al. (2019). Open Chromatin Profiling in Adipose Tissue Marks Genomic Regions with Functional Roles in Cardiometabolic Traits. G3 GenesGenomesGenetics *9*, 2521–2533.

Dale, R.K., Pedersen, B.S., and Quinlan, A.R. (2011). Pybedtools: a flexible Python library for manipulating genomic datasets and annotations. Bioinformatics *27*, 3423–3424.

Das, S., Forer, L., Schönherr, S., Sidore, C., Locke, A.E., Kwong, A., Vrieze, S.I., Chew, E.Y., Levy, S., McGue, M., et al. (2016). Next-generation genotype imputation service and methods. Nat. Genet. *48*, 1284–1287.

van de Geijn, B., McVicker, G., Gilad, Y., and Pritchard, J.K. (2015). WASP: allele-specific software for robust molecular quantitative trait locus discovery. Nat. Methods *12*, 1061–1063.

Grant, C.E., Bailey, T.L., and Noble, W.S. (2011). FIMO: scanning for occurrences of a given motif. Bioinformatics *27*, 1017–1018.

Hinrichs, A.S., Karolchik, D., Baertsch, R., Barber, G.P., Bejerano, G., Clawson, H., Diekhans, M., Furey, T.S., Harte, R.A., Hsu, F., et al. (2006). The UCSC Genome Browser Database: update 2006. Nucleic Acids Res. *34*, D590-598.

Hsi-Yang Fritz, M., Leinonen, R., Cochrane, G., and Birney, E. (2011). Efficient storage of high throughput DNA sequencing data using reference-based compression. Genome Res. *21*,

734–740.
Jun, G., Wing, M.K., Abecasis, G.R., and Kang, H.M. (2015). An efficient and scalable analysis framework for variant extraction and refinement from population scale DNA sequence data. Genome Res. gr.176552.114.
Kent, W.J., Zweig, A.S., Barber, G., Hinrichs, A.S., and Karolchik, D. (2010). BigWig and BigBed: enabling browsing of large distributed datasets. Bioinformatics *26*, 2204–2207.
Loh, P.-R., Danecek, P., Palamara, P.F., Fuchsberger, C., A Reshef, Y., K Finucane, H., Schoenherr, S., Forer, L., McCarthy, S., Abecasis, G.R., et al. (2016). Reference-based phasing using the Haplotype Reference Consortium panel. Nat. Genet. *48*, 1443–1448.
Martin, M. (2011). Cutadapt removes adapter sequences from high-throughput sequencing reads. EMBnet.Journal *17*, 10–12.
Meuleman, W., Muratov, A., Rynes, E., Halow, J., Lee, K., Bates, D., Diegel, M., Dunn, D., Neri, F., Teodosiadis, A., et al. (2020). Index and biological spectrum of human DNase I hypersensitive sites. Nature *584*, 244–251.
R Core Team (2020). R: A Language and Environment for Statistical Computing (Vienna, Austria: R Foundation for Statistical Computing).
Rai, V., Quang, D.X., Erdos, M.R., Cusanovich, D.A., Daza, R.M., Narisu, N., Zou, L.S., Didion, J.P., Guan, Y., Shendure, J., et al. (2020). Single-cell ATAC-Seq in human pancreatic islets and deep learning upscaling of rare cells reveals cell-specific type 2 diabetes regulatory signatures. Mol. Metab. *32*, 109–121.
Schep, A.N., Wu, B., Buenrostro, J.D., and Greenleaf, W.J. (2017). chromVAR: inferring transcription-factor-associated accessibility from single-cell epigenomic data. Nat. Methods.
Stuart, T., Butler, A., Hoffman, P., Hafemeister, C., Papalexi, E., Mauck, W.M., Hao, Y., Stoeckius, M., Smibert, P., and Satija, R. (2019). Comprehensive Integration of Single-Cell Data. Cell *177*, 1888-1902.e21.
Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., et al. (2015). UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. PLoS Med. *12*.
The 1000 Genomes Project Consortium (2015). A global reference for human genetic variation. Nature *526*, 68–74.
Weirauch, M.T., Yang, A., Albu, M., Cote, A.G., Montenegro-Montero, A., Drewe, P., Najafabadi, H.S., Lambert, S.A., Mann, I., Cook, K., et al. (2014). Determination and inference of eukaryotic transcription factor sequence specificity. Cell *158*, 1431–1443.
Welch, J.D., Kozareva, V., Ferreira, A., Vanderburg, C., Martin, C., and Macosko, E.Z. (2019). Single-Cell Multi-omic Integration Compares and Contrasts Features of Brain Cell Identity. Cell *177*, 1873-1887.e17.
Yanai, I., Benjamin, H., Shmoish, M., Chalifa-Caspi, V., Shklar, M., Ophir, R., Bar-Even, A., Horn-Saban, S., Safran, M., Domany, E., et al. (2005). Genome-wide midrange transcription profiles reveal expression level relationships in human tissue specification. Bioinformatics *21*, 650–659.
Yang, S., Corbett, S.E., Koga, Y., Wang, Z., Johnson, W.E., Yajima, M., and Campbell, J.D. (2020). Decontamination of ambient RNA in single-cell RNA-seq with DecontX. Genome Biol. *21*, 57.
Ziegenhain, C., and Sandberg, R. (2021). anonymizeBAM: Versatile anonymization of human sequence data for open data sharing. BioRxiv 2021.01.11.426206.