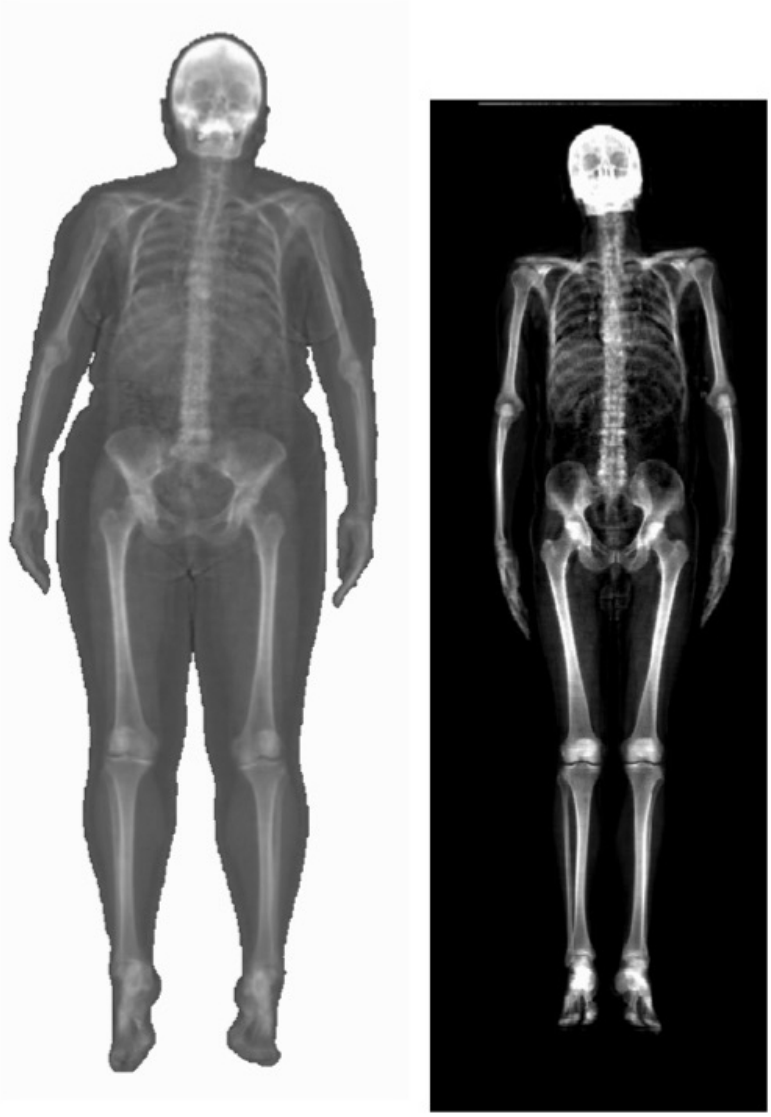

Fig. S1. Types of DXA images acquired from the UKB. (Left) Image of patient imaged on white background. (Right) Image of patient imaged on black background. Sizes of images are true to scale.

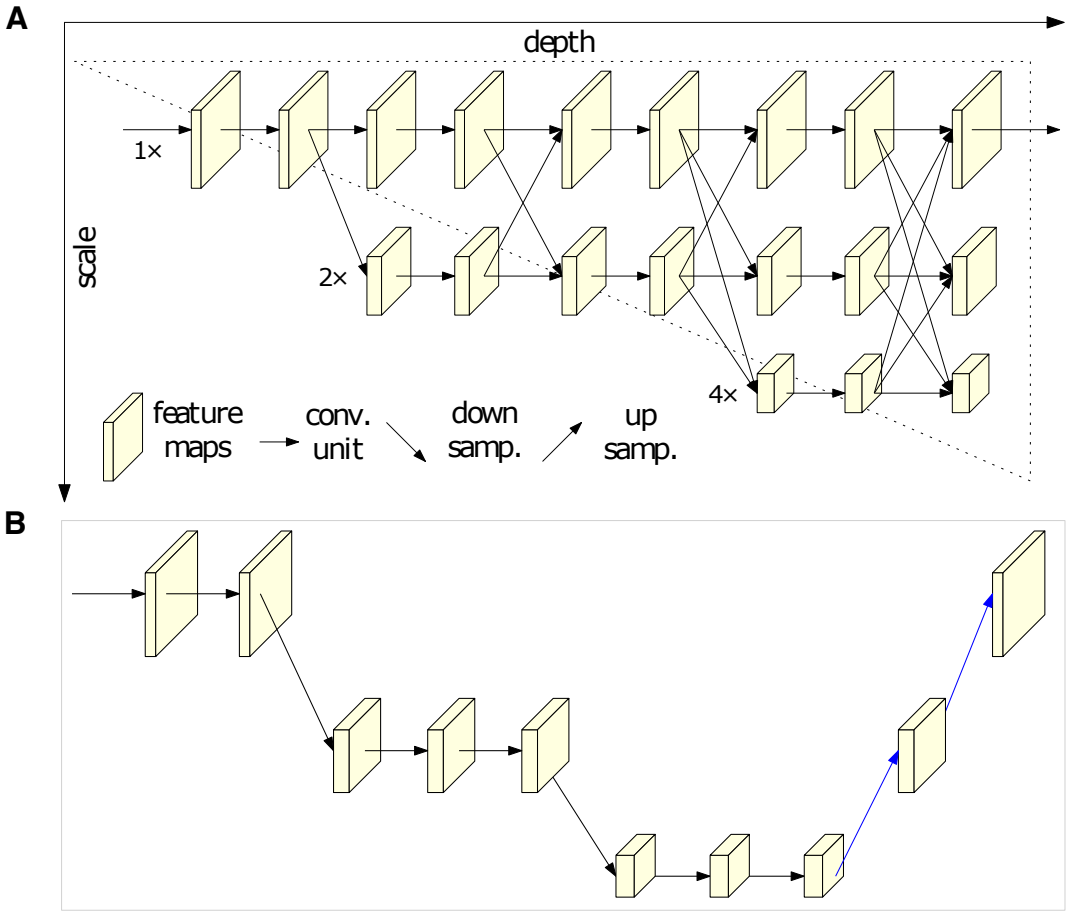

741
742 **Fig. S2. A comparison of HRNet and ResNet deep learning architectures**. (**A)** High-
743 Resolution Network (HRNet) architecture maintains parallel high to low resolution subnetworks.
744 (**B)** Simple Baseline deep learning architecture (ResNet) which relies on a high-to-low and low-
745 to-high framework. Both images are taken directly from Sun et al., (*35*) to illustrate the
746 architectural differences between HRNet and a standard architecture for this prediction task.
747

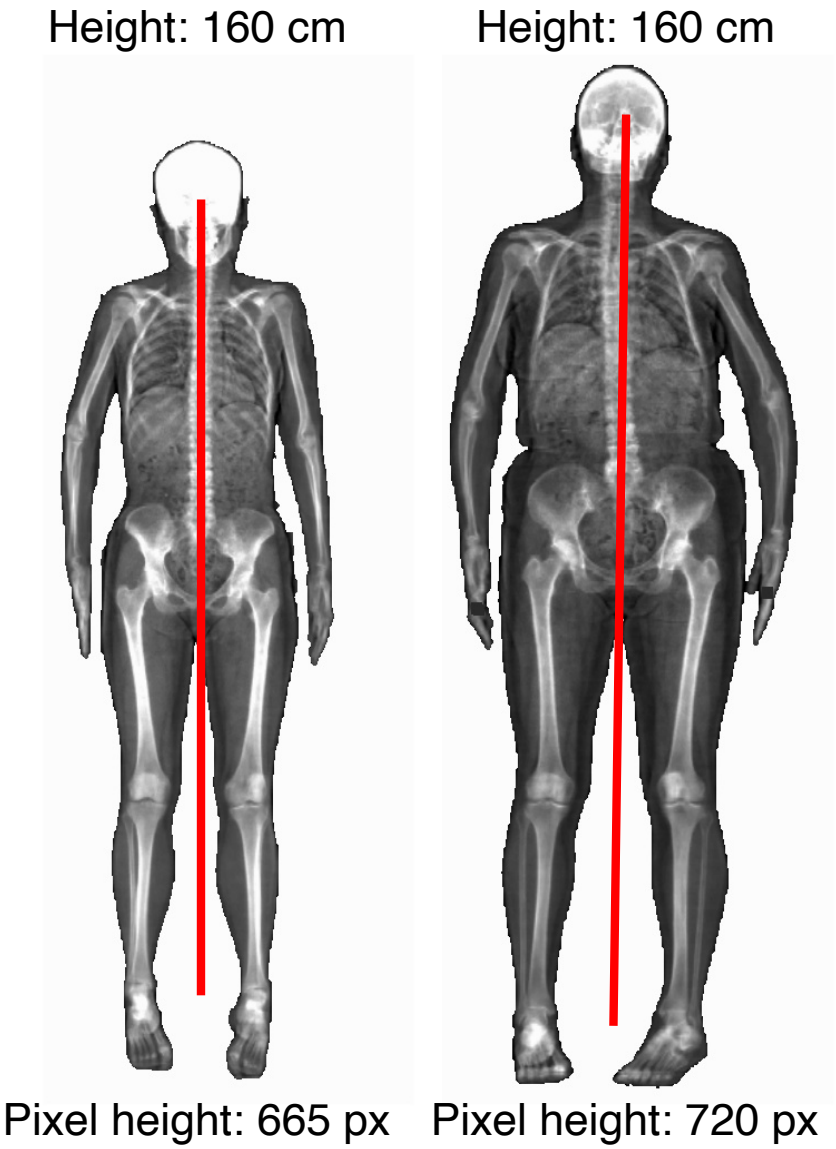

Fig. S3. DXA images from the UKB that have undergone different image scaling. Example of two individuals who were measured to be the same height in the FID 50 in the UKB (overall height) but pixel-based measurements of one image were considerably smaller than the other due to image scaling/resolution differences.

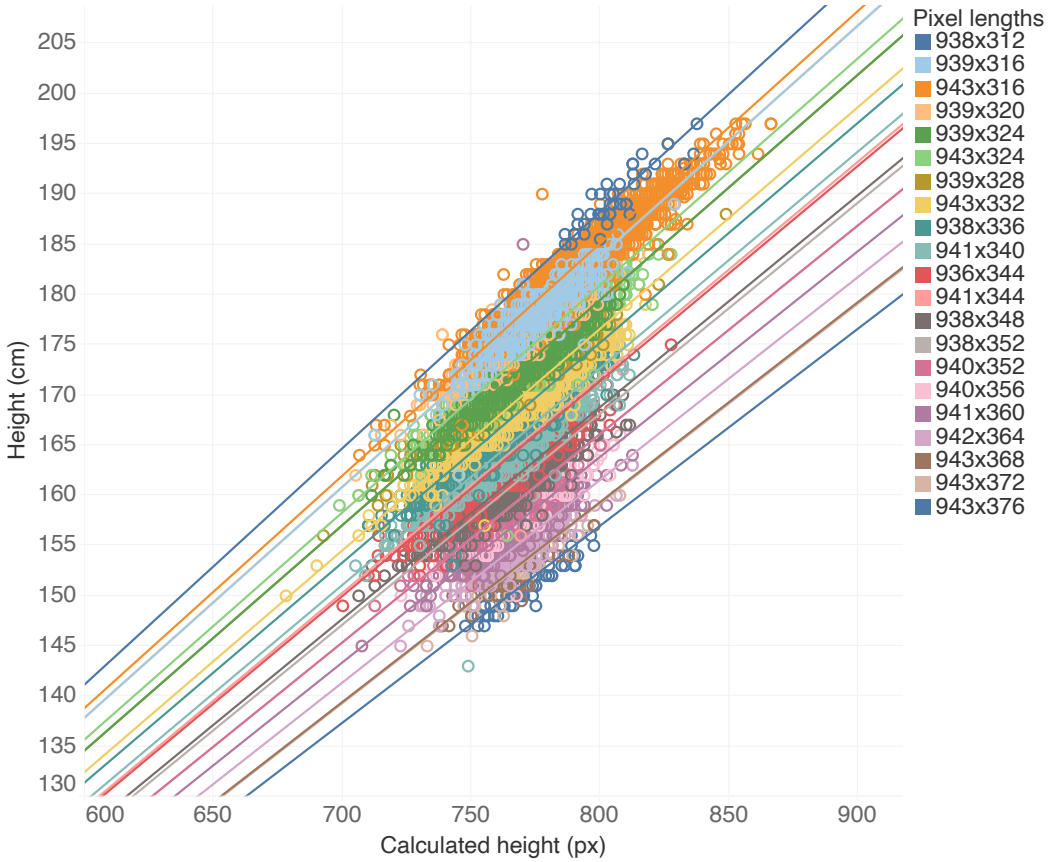

Fig. S4. A linear regression of image-measured height against UKB-measured height. For each image pixel-ratio, we regressed height measured in the UKB with height we calculated in pixels from the DXA scan. This provided a conversion from pixels to cm that we used as a normalization factor to correct for differences in resolution.

Fig. S5. Examples of individuals who were outliers on our measurement and were removed from analysis. (Left) Individual with femur deformity and metal implants. (Right) Individual with missing forearm.

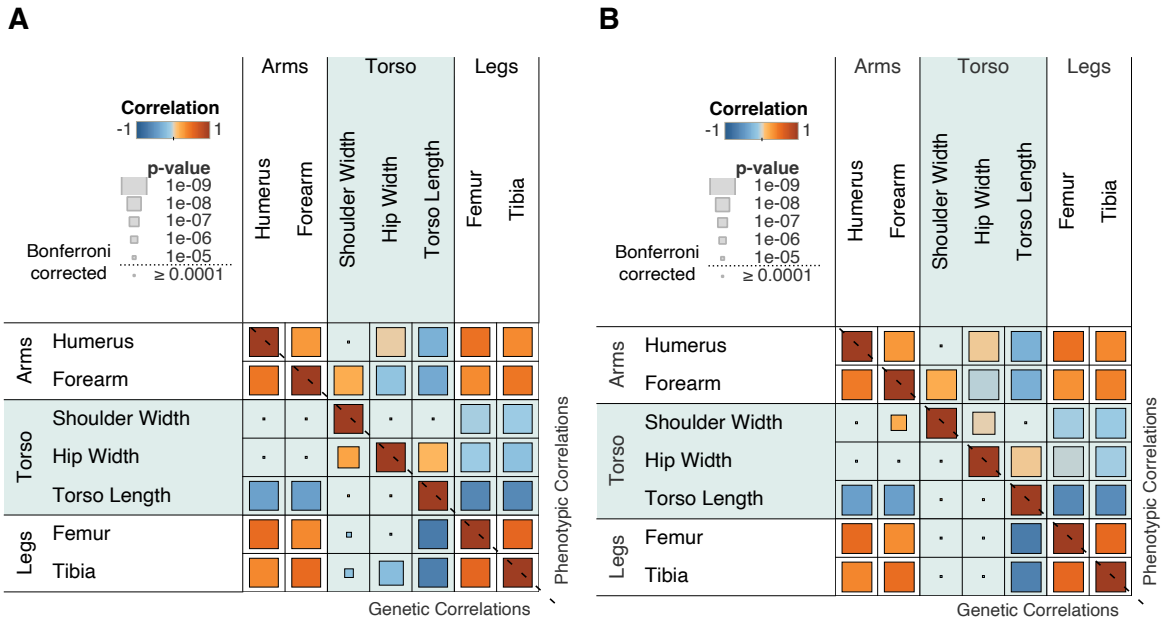

Fig. S6. A heatmap comparison of genotype and phenotype correlations between ratios and residuals. (A) Matrix of genotype and phenotype correlation with each phenotype computed as a ratio of height. (**B)** Matrix of genotype and phenotype correlation with each phenotype computed by regressing the phenotype with height and then obtaining residuals.

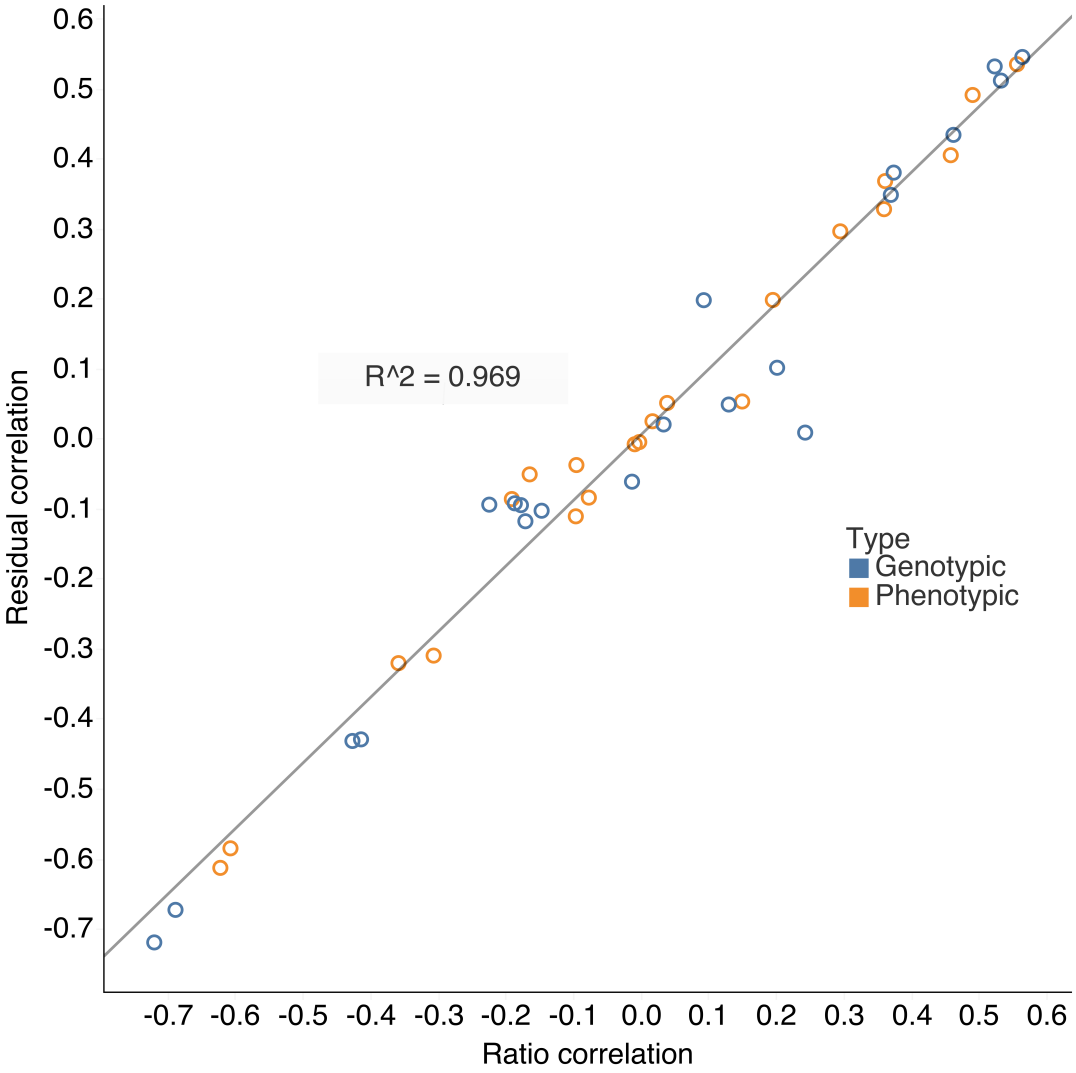

Fig. S7. Correlation of genotype and phenotype correlations across skeletal traits, computed using ratios with height and second residualizing for height

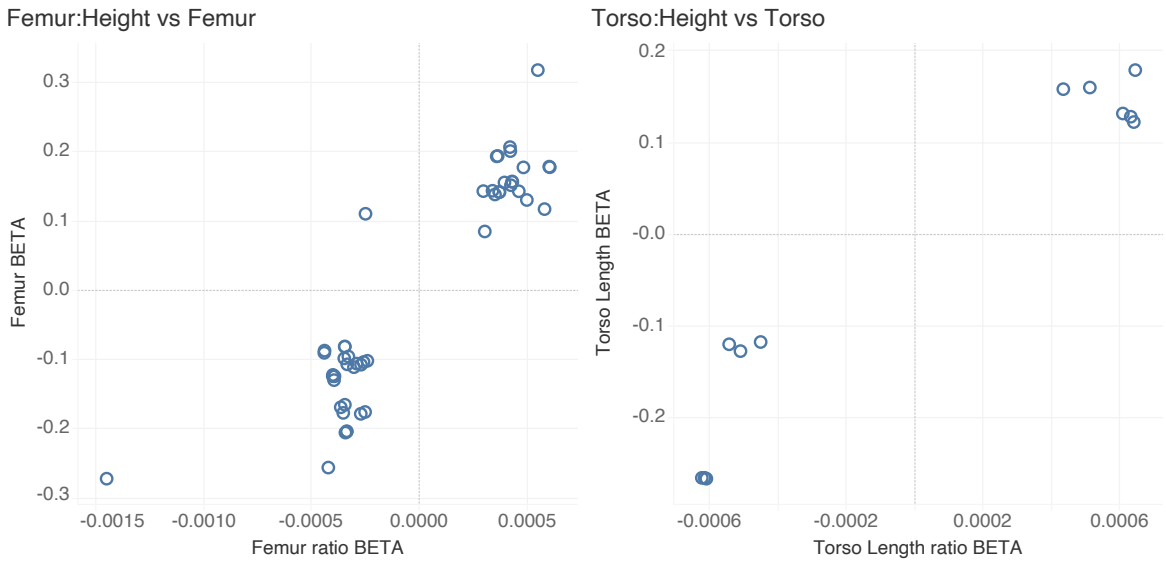

Fig. S8. Comparison of effect estimates of independent genome-wide significant SNPs across different phenotypes. Effect estimates of genome-wide significant SNPs for each phenotype ($p < 5e-08$) showing same effect directionality for skeletal proportions and raw measurements.

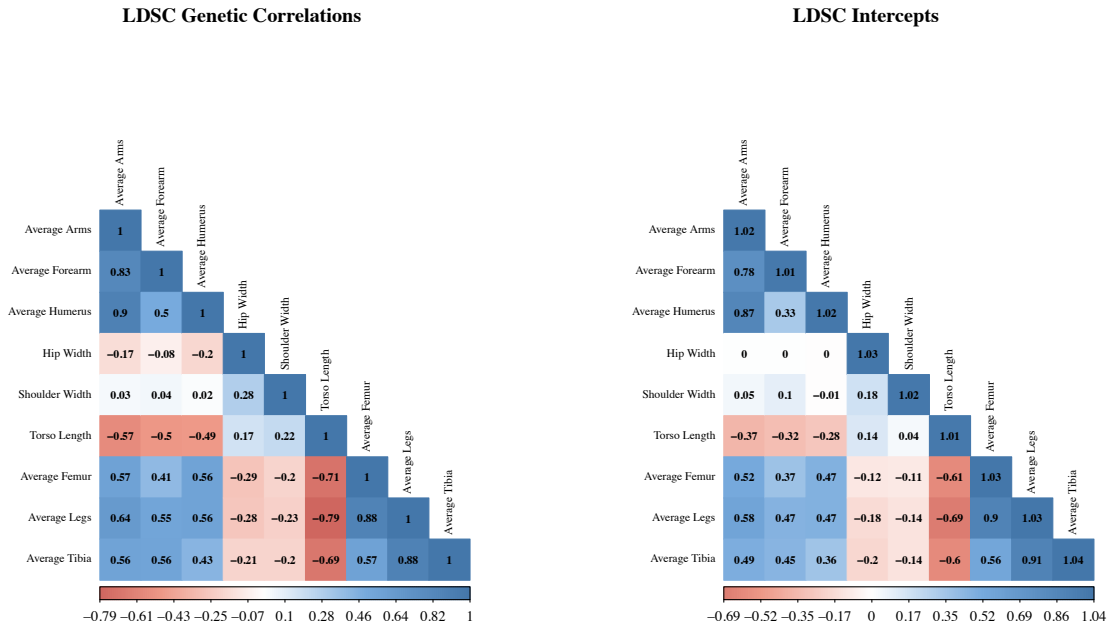

Fig. S9. Heatmap of genetic correlations and LDSC cross-trait intercepts across skeletal proportion phenotypes within odd-numbered chromosomes

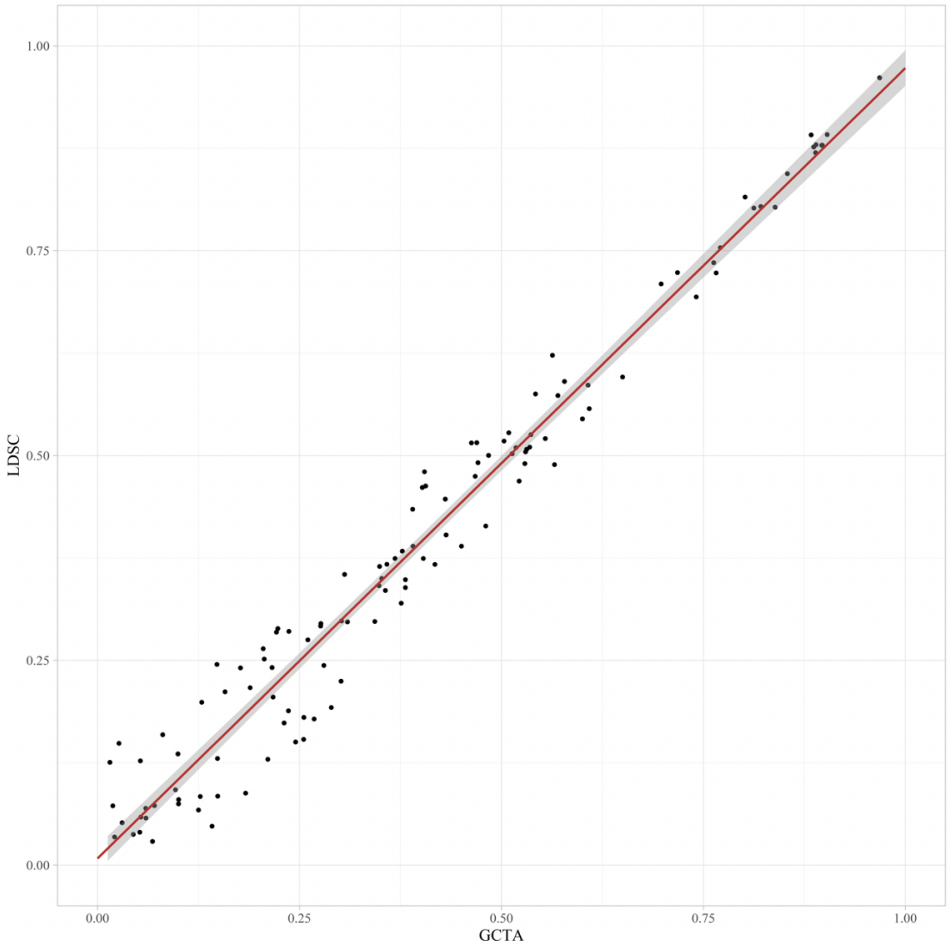

Fig. S10. Scatterplot of GCTA and LDSC genetic correlation estimates across skeletal ratios.

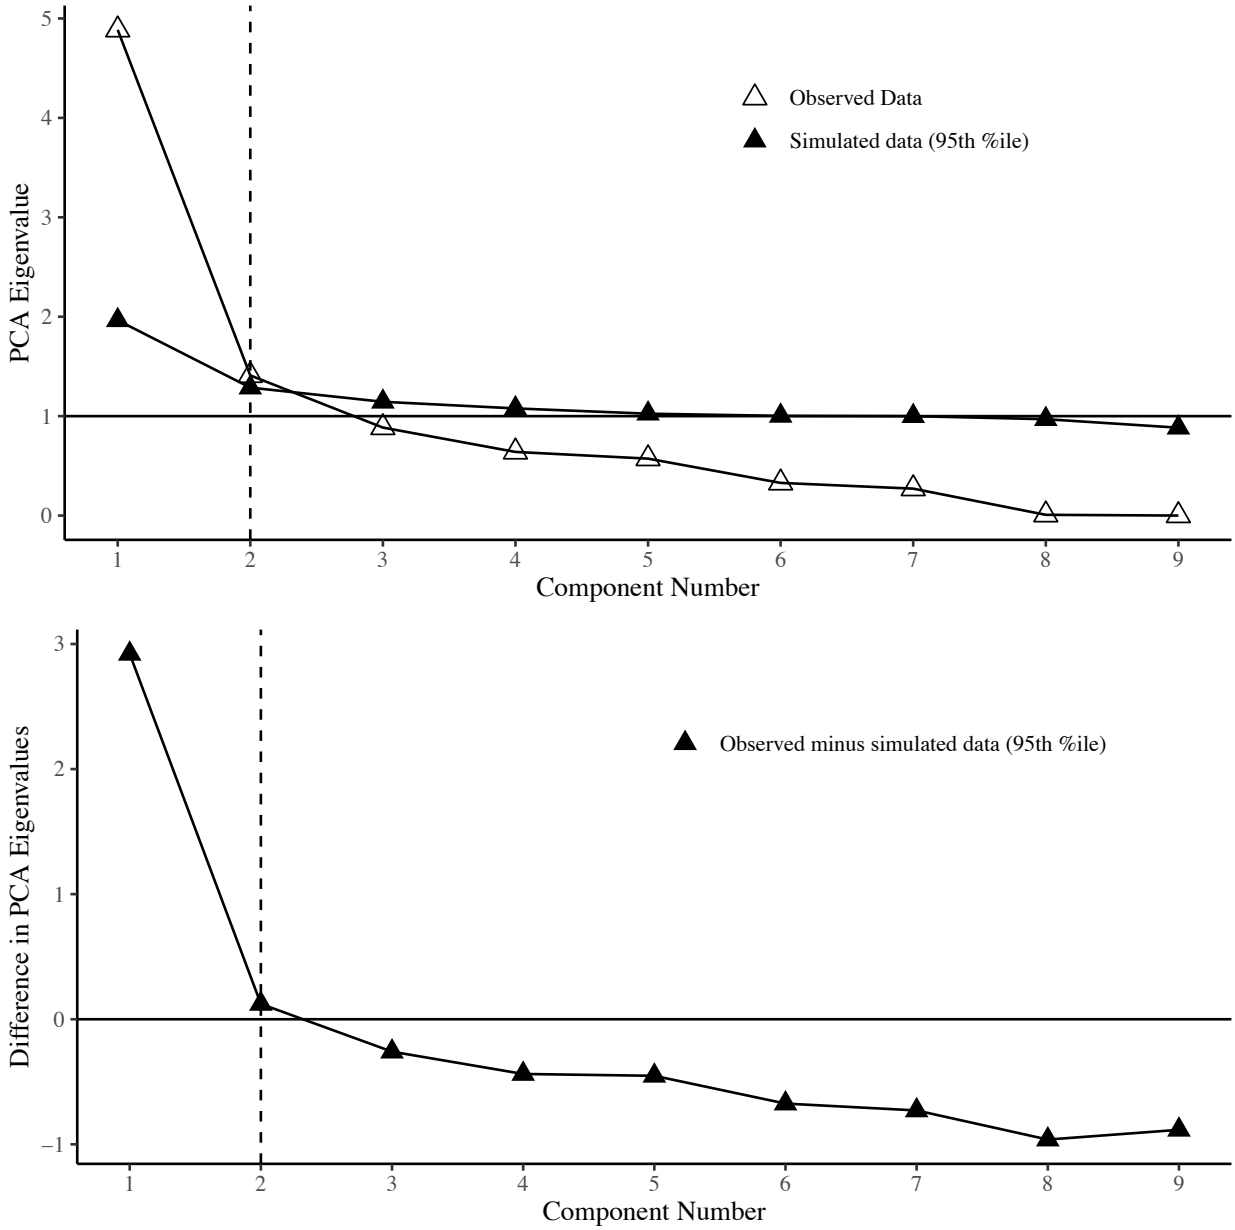

Fig. S11. Screeplots of PCA and difference in PCA from LDSC Parallel Analysis

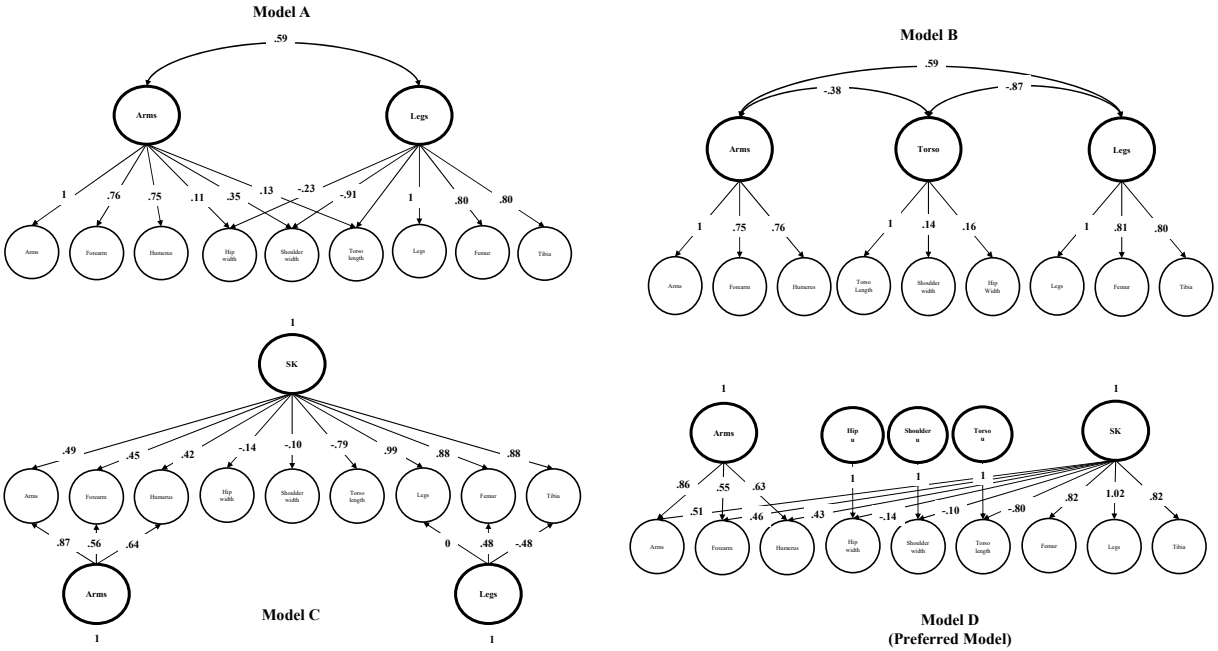

Fig. S12. Confirmatory factor models of skeletal traits (even autosomes)

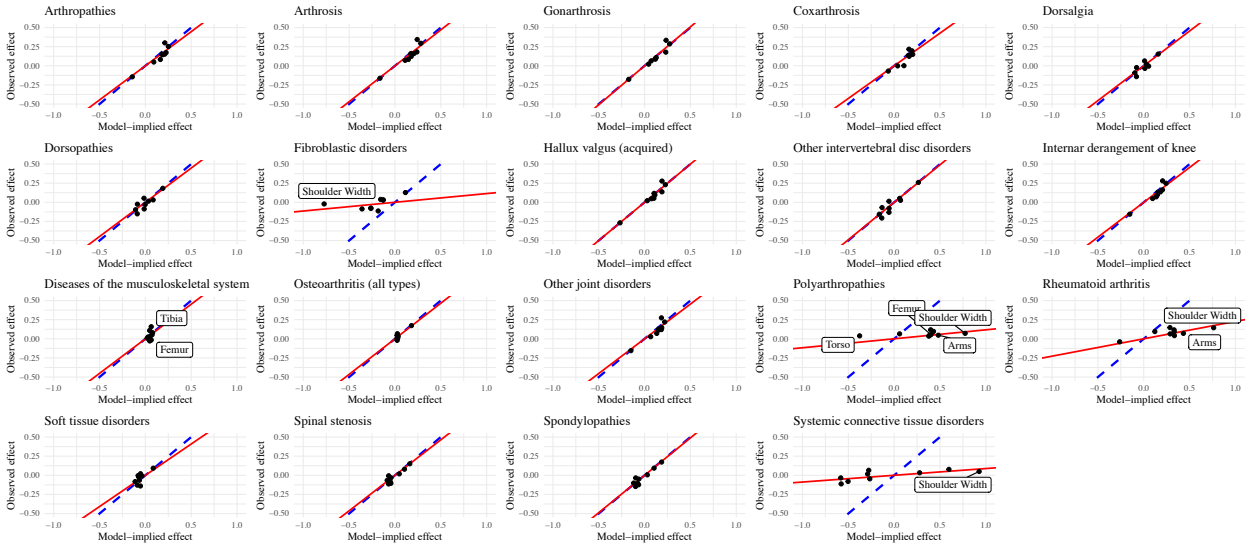

Fig. S13. Scatterplots of observed and model-implied effects between musculoskeletal diseases and skeletal endophenotypes. Red lines represent best fitting regression lines. Blue dashed lines represent perfect fit (Observed effects = Model-implied effects). Labeled traits are outliers detected based on standardized differences between the observed and the common factor model-implied effects for the skeletal traits > 2.

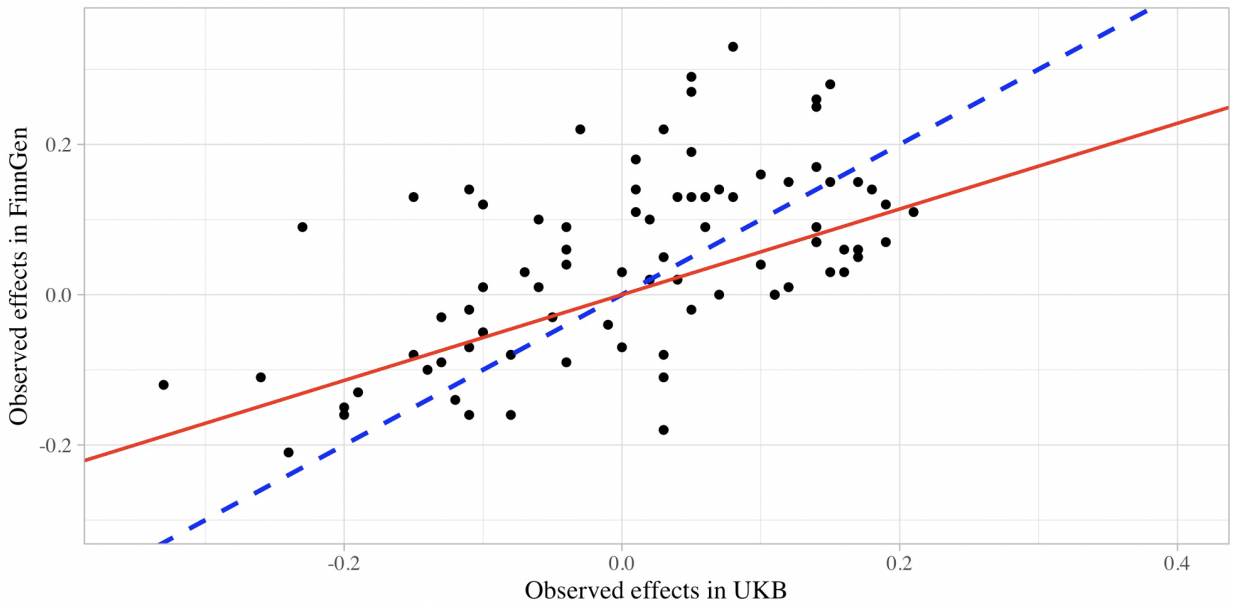

Fig. S14. Observed effects of skeletal traits on musculoskeletal diseases common in UKB and FinnGen. Red lines represent best fitting regression lines. Blue dashed lines represent perfect correspondence (intercept = 0, slope = 1).

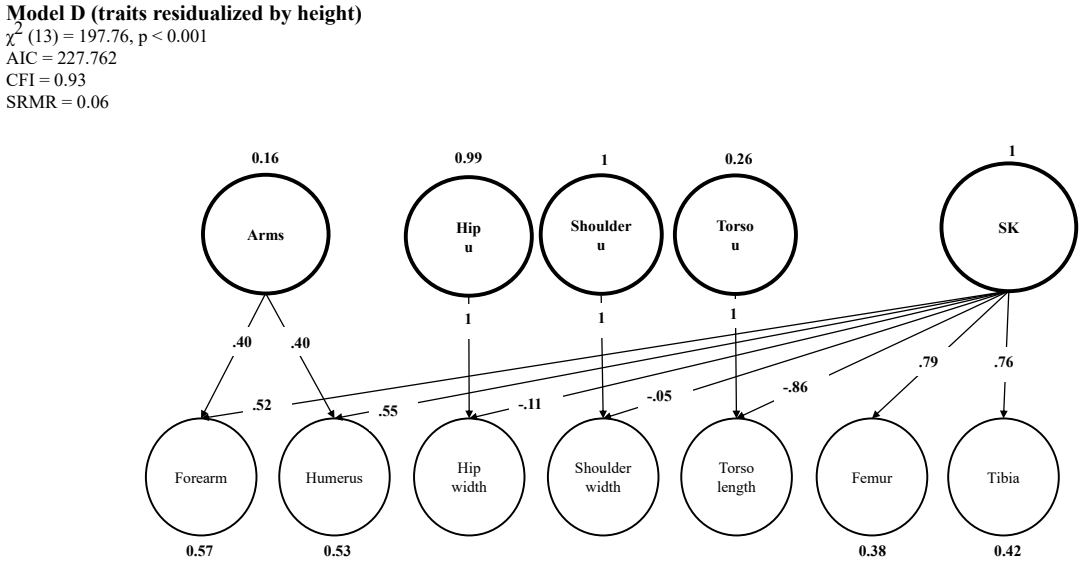

Fig. S15. Confirmatory Factor Model D applied to residuals. Preferred model D fully standardized parameter estimates fitted on height-residualized skeletal traits as well as excluding overall arms and leg length

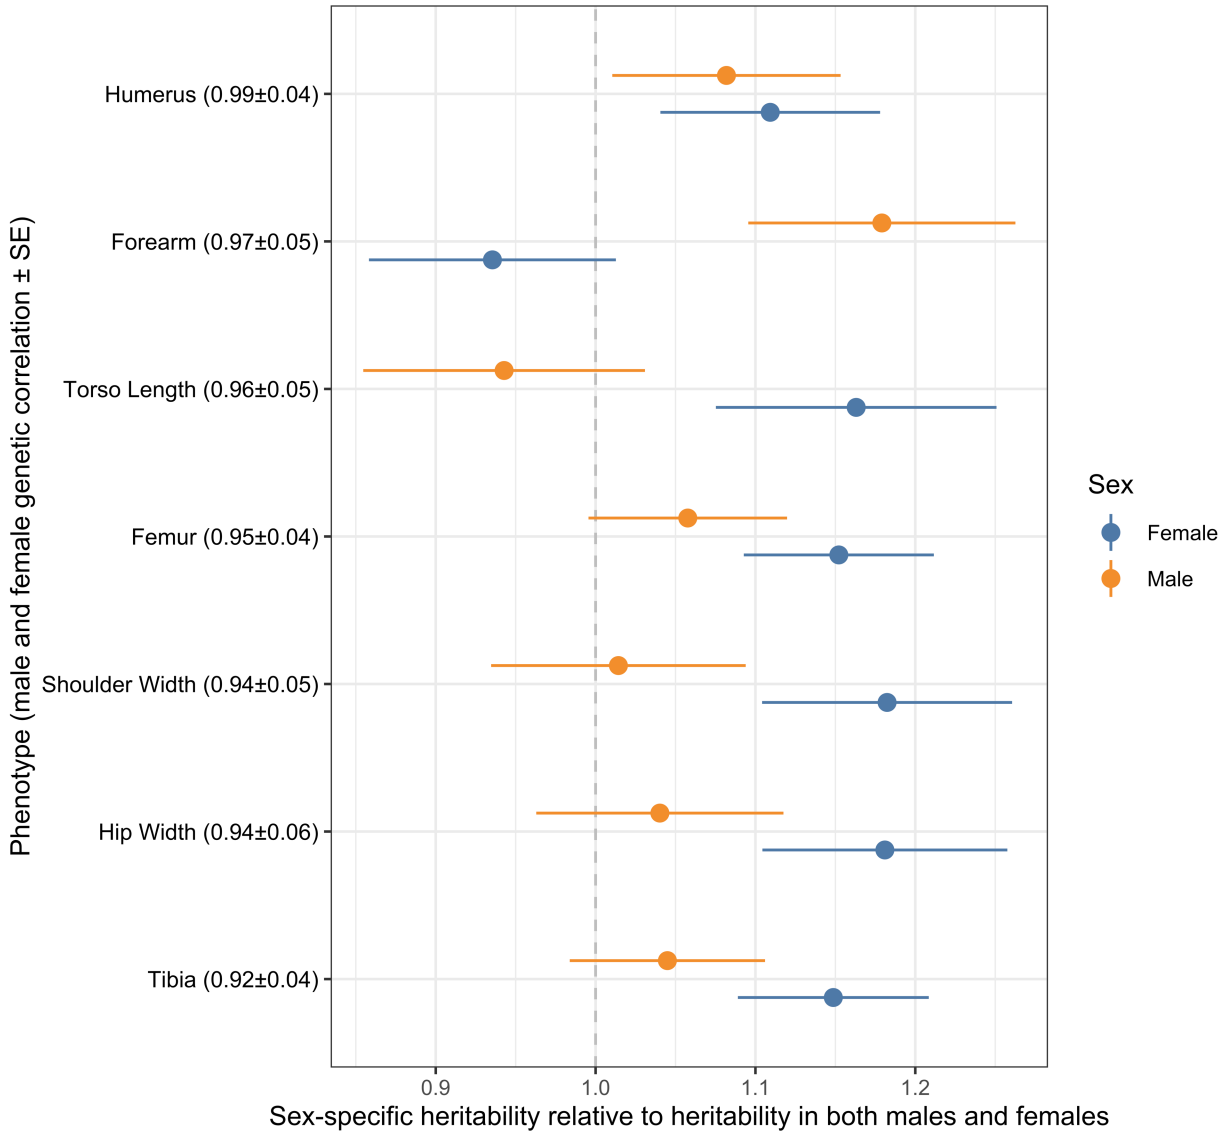

Fig. S16. Genetic correlations between males and females, estimated using bi-variate LD Score Regression shown for each trait (y-axis). SNP heritability divided by the SNP heritability estimated in the sample with both sexes combined (x-axis) for all traits

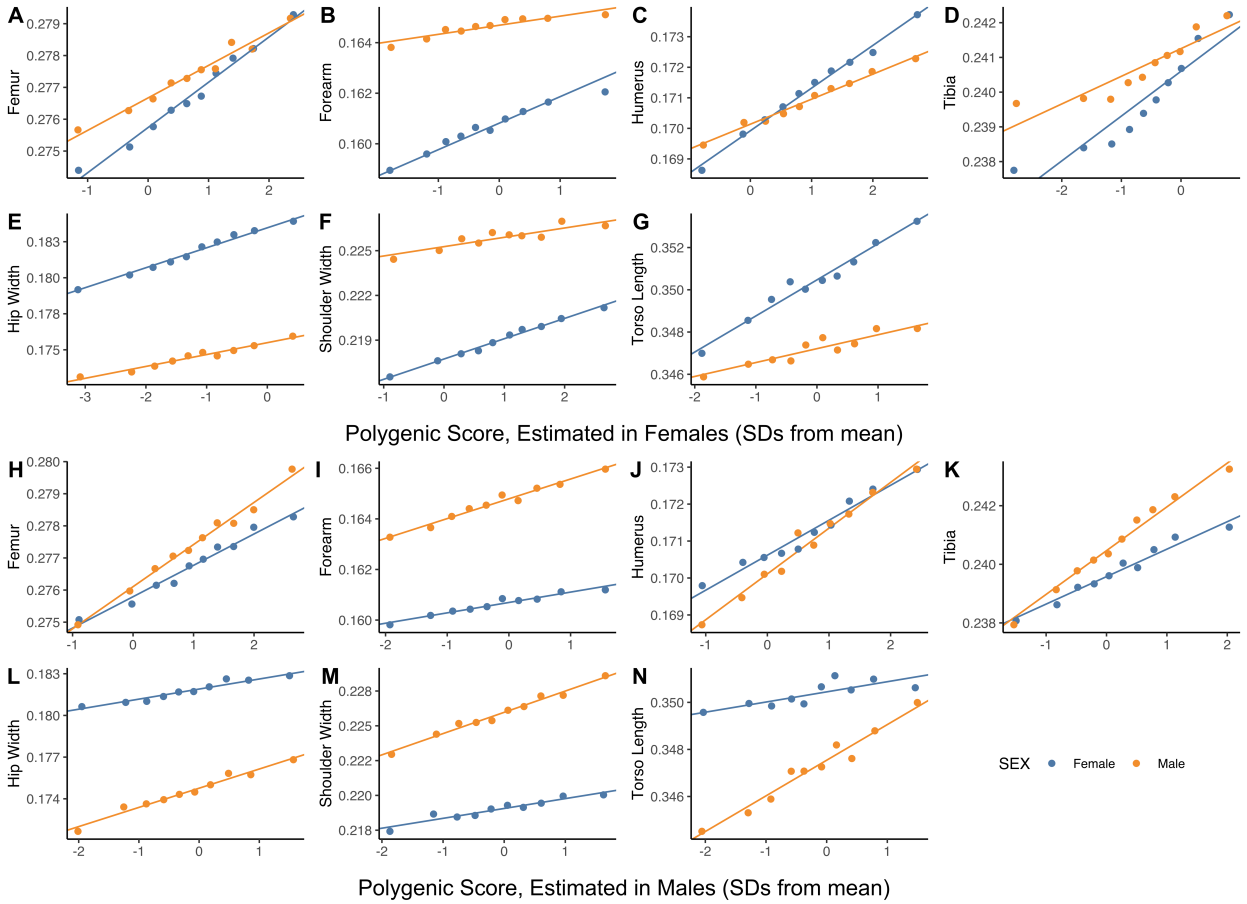

Fig. S17. Regression of trait values in males (orange) and separately in females (blue) to a polygenic score estimated in an independent sample of females. Points show mean values in one decile of the polygenic score; the fitted line and associated effect estimate and R^2 correspond to regressions on the raw, non-binned data.

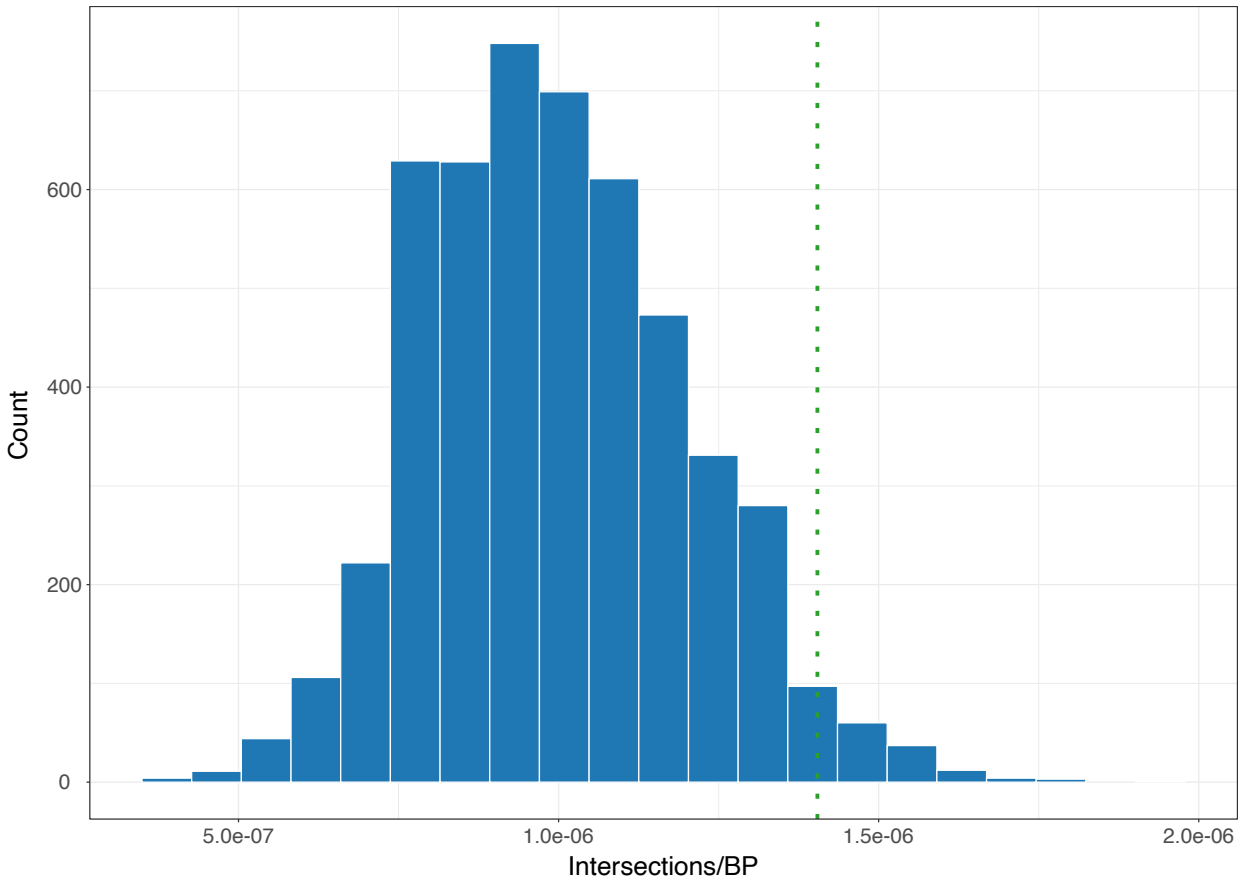

Fig. S18. HAR background distributions. Intersections per base pair occurring between human accelerated regions (HARs) and phenotype-associated genes. Blue bars are background distributions generated from 5,000 simulations of matched element sets. An example is shown here for skin pigmentation.