Genome-wide mapping of global-to-local genetic effects on human facial shape

Peter Claes [1,2,3]*, Jasmien Roosenboom[4], Julie D. White[5], Tomek Swigut[6], Dzemila Sero [1,2],
Jiarui Li[1,2], Myoung Keun Lee[4], Arslan Zaidi[5], Brooke C. Mattern[5], Corey Liebowitz[5], Laurel Pearson [5],
Tomás González[5], Elizabeth J. Leslie[4], Jenna C. Carlson [7], Ekaterina Orlova[8], Paul Suetens[1,2],
Dirk Vandermeulen[1,2], Eleanor Feingold[7,8], Mary L. Marazita[4,8], John R. Shaffer[8], Joanna Wysocka[6,9]*,
Mark D. Shriver[5]* and Seth M. Weinberg [4,10]*

Genome-wide association scans of complex multipartite traits like the human face typically use preselected phenotypic measures. Here we report a data-driven approach to phenotyping facial shape at multiple levels of organization, allowing for an open-ended description of facial variation while preserving statistical power. In a sample of 2,329 persons of European ancestry, we identified 38 loci, 15 of which replicated in an independent European sample (*n* = 1,719). Four loci were completely new. For the others, additional support (*n* = 9) or pleiotropic effects (*n* = 2) were found in the literature, but the results reported here were further refined. All 15 replicated loci highlighted distinctive patterns of global-to-local genetic effects on facial shape and showed enrichment for active chromatin elements in human cranial neural crest cells, suggesting an early developmental origin of the facial variation captured. These results have implications for studies of facial genetics and other complex morphological traits.

The human face is a multipartite trait composed of distinct features (eyes, nose, chin and mouth), whose size, shape and composition are clearly heritable. However, knowledge of which genetic variants are responsible for human facial variation is still lacking[1]. Several genome-wide association studies (GWAS) have each identified a handful of loci associated with a small number of facial traits, with few of these loci having been replicated[2–7]. While a GWAS benefits from being an unbiased approach to gene mapping, the phenotypic descriptions in such studies are typically preselected and used to classify individuals in a 'phenotype-first' way of thinking. This approach may be appropriate in certain instances (for example, affected and unaffected disease status) but is less so for complex, multipartite traits like the human face. The result is that facial shape has been reduced to a limited series of measurements (for example, linear distances) that are analyzed individually, resulting in a loss of information.

In this study, we present a data-driven approach to facial phenotyping that exploits both the partable and integrated information contained in 3D facial images, allowing for the identification of genetic effects on facial shape at multiple levels of organization—from global to local. This approach generates a nested series of multivariate GWAS, with a low computational burden and, more importantly, controlled multiple-testing burden. We applied this new approach to a European-derived discovery cohort and then tested significantly associated variants for replication in an independent European-derived cohort. In an effort to provide additional validation, we integrated our work with previously published human

facial GWAS. We show a number of newly associated genetic loci supported by strong statistical evidence, identifying unreported patterns in global-to-local genetic effects on facial shape. Furthermore, these loci are preferentially marked by active chromatin signatures in human cranial neural crest cells (CNCCs), an embryonic cell type that gives rise to most of the craniofacial structures. This suggests a developmental origin of much of the facial variation uncovered by our study. These results offer new insights on the genetic basis of human facial shape with potentially far-reaching implications. More generally, the results present an alternative to the prevailing phenotype-first mindset and our approach is widely applicable to any GWAS on complex, quantitative and multipartite traits, especially those captured thoroughly using images.

Results

Study samples. A study sample of 2,329 unrelated participants of European ancestry made up the discovery cohort for our analysis (the Pittsburgh sample, PITT). These participants had a median age of 23 years and were recruited from several US sites. An additional, independently collected and genotyped sample of 1,719 adult participants of European ancestry made up our replication cohort (the Penn State sample, PSU). These participants had a median age of 22 years and were recruited at several sites in the United States and Europe. For both cohorts, imputation of unobserved genetic variants and sporadic missing genotype calls for assayed SNPs was performed using the 1000 Genomes Project[8] Phase 3 reference panel. Basic demographic descriptors and general physical characteristics

(for example, sex, age, height and weight) were available from participants in both cohorts.

Global-to-local facial segmentations. We used digital stereophotogrammetry to obtain 3D facial images from all participants. After trimming and cleaning, 3D facial images were aligned in dense correspondence[9], ensuring that homology was established among the roughly 10,000 3D points (considered quasi-landmarks[10]) making up an individual's facial shape. Facial shape, in both cohorts separately, was adjusted for potential confounding variables (age, sex, height, weight, facial size and cohort-specific population structure) and then partitioned in an unsupervised manner into a series of global-to-local facial segments. To accomplish this, we applied hierarchical spectral clustering[11] to the facial quasi-landmark configurations of the PITT sample. First, the quasi-landmarks were listed in a squared ($\sim$10,000 $\times$ $\sim$10,000) similarity matrix using pairwise 3D correlations. Second, a Laplacian transformation was applied to enhance similarities before an Eigen decomposition of the squared matrix. Finally, within the Eigen spectral map, k-means++ clustering was used to group highly correlated quasi-landmarks, which, when mapped back to the facial surface, result in a segmentation of the face. This was done in a bifurcating hierarchical manner using five levels, such that facial segments with closer relationships were located nearby one another and each facial segment was split in two toward the next level, resulting in a total of 63 segments, as depicted in Fig. 1 using a polar dendrogram layout. The hierarchical design gradually focused on more local shape aspects that were otherwise overlooked, without ignoring the integration of facial parts more globally at higher levels. The same pipeline was applied to the PSU cohort (Supplementary Fig. 1), and the resulting segmentation was compared to the PITT segmentation using the normalized

mutual information ($0 < \mathrm{NMI} < 1$; 0, no overlap; 1, perfect overlap) at each of the five levels ($\mathrm{NMI_{L0}} = 1$, $\mathrm{NMI_{L1}} = 0.90$, $\mathrm{NMI_{L2}} = 0.80$, $\mathrm{NMI_{L3}} = 0.72$, $\mathrm{NMI_{L4}} = 0.75$, $\mathrm{NMI_{L5}} = 0.78$). These high NMI values indicate substantial overlap and, therefore, good reproducibility of the global-to-local facial segmentations across similar but independent samples. The hierarchical spectral clustering (Fig. 1; using the PITT cohort as the reference for the remainder of the study) subdivided facial shape into meaningful and recognizable segments. Globally, the midface was first separated from the rest of the face and was further partitioned into the region of the mouth (quadrant 1, starting at segment 4) and nose (quadrant 2, starting at segment 5). The remainder of the face was further partitioned into the lower facial area (quadrant 3, starting at segment 6) and the upper facial area (quadrant 4, starting at segment 7). Each quadrant was repeatedly partitioned, increasingly focusing on smaller facial parts. This provided an efficient and objective means for subdividing facial shape into global-to-local parts with internally well-correlated shape variations.

To generate variables that captured biological shape, we applied generalized Procrustes analysis (GPA) separately to the quasi-landmarks making up each facial segment. As such, a shape space for each facial segment was constructed that was independent from the other segments and its relative positioning in lower-level (larger) segments. Following GPA, we applied principal-component analysis (PCA) to extract the major factors of shape variation characterizing each facial segment. We used parallel analysis[12] to determine the number of principal components (PCs) needed to adequately summarize shape variation for the given segment. The number of PCs retained for each facial segment after parallel analysis and the percentage of total variation they explained in the PITT cohort are depicted in Supplementary Fig. 2. As expected, we found that lower-level facial segments required more PCs and that the retained PCs for all facial segments explained most of the total shape variance (median of 95%, min = 89%, max = 97%) present within the respective segments. The result was nearly complete coverage of all facial shape variation at five different levels of detail.

Genetic mapping of global-to-local shape. We performed a series of GWAS to test the genetic association of a total of 9,478,608 SNPs with the shape information contained in each of the 63 facial segments using the PITT discovery cohort. As each of the facial segments was represented by multiple dimensions of variation (PCs), we used a multivariate canonical correlation analysis (CCA)[13]. In brief, CCA extracts the linear combination of PCs from a facial segment that has maximal correlation with the SNP being tested, under the additive genetic model, after correcting for confounding variables. Therefore, CCA avoids the preselection of individual PCs and reveals the associated facial effect as a linear direction in a multidimensional shape space.

We found a total of 1,932 SNPs among 38 separate loci that reached nominal genome-wide significance ($P \le 5 \times 10^{-8}$). Quantile–quantile, regional association and detailed supplementary illustrations of all 38 loci are provided online (see URLs). To accommodate the multiple-testing burden present in performing separate genome scans for 63 facial segments, we determined the false discovery rate for any test dependency structure[14] (FDRd) P-value threshold at 2.82×10^{-8} and a more conservative study-wide Bonferroni-corrected significance threshold at 1.28×10^{-9}. Following imputation in the PSU cohort, 1,931 of these 1,932 SNPs were available for replication testing. Because it is possible for a SNP to affect different aspects of variation in the same facial segment across different datasets, replicating an observed genetic association necessitated both replicating a significant effect on a given facial segment and a determination that the same aspects of variation were affected. Therefore, we first used the discovery cohort as a phenotyping reference for the replication cohort (see the Methods for additional

Fig. 1 | Hierarchical spectral clustering of facial shape. Global-to-local facial segmentation of the PITT cohort (n = 2,329) obtained using hierarchical spectral clustering. Segments are colored in green. Segment 1, representing the entire face and shown in the center, is partitioned into segments 2 and 3, shown in the innermost concentric circle. Moving outward across the five concentric circles, corresponding to the five hierarchical levels, each facial segment is further partitioned into two more localized segments.

details). Then, the multidimensional PC scores of the PSU cohort were projected onto the CCA loadings of the discovered effect. In doing so, the associated facial trait, once identified using CCA in the PITT cohort, was kept fixed and thus consistent between the PITT discovery cohort and the PSU replication cohort. This resulted in a univariate facial measure that was modeled using linear regression, to test for replication. This was done separately for each SNP and facial segment combination that achieved genome-wide significance in the discovery phase. From all replication efforts combined, we computed the FDRd adjusted significance threshold at 4.20×10^{-3} and a Bonferroni-corrected significance threshold at 3.28×10^{-4}. A total of 1,821 SNPs across 15 loci replicated with P values well below the FDRd threshold (Table 1).

We focused on the 15 replicated loci to investigate patterns of genome to facial shape associations. These 15 loci involved a variety of facial segments, and several loci affected segments in more than one facial region, as illustrated in Figs. 2 and 3. Several interesting patterns were observed. First, the majority of affected segments were located in the nose and the lower facial quadrant, primarily the chin. Second, some associations were involved in high-level segments, often at the outer layer(s) of the polar dendrogram, suggesting a localized phenotypic effect (for example, 1p32.1). In contrast, we observed other associations involving numerous linked facial segments from the edge of the dendrogram (level$_5$) to the center (level$_0$). Interestingly, the evidence of association (obtained by tracking the $-\log_{10} P$ values in Fig. 3) in these instances may increase from center to edge (global to local; for example, 7q21.3) or from edge to center (local to global; for example, 1p12), reflecting a localized versus global phenotypic effect, respectively. Still others, for example, 1q31.3, reached a maximum association partway up the dendrogram, reflecting an effect on an intermediate global-to-local level. Interestingly, both of the loci on chromosome 17 had the same pattern of association across nose segments, but 17q24.3 (rs11655006[T>C]) showed the stronger association at the more global levels and 17q24.3 (rs5821892[C>CG]) showed strongest association at the most local level. Third, most associations were limited to linked segments within one facial quadrant of the polar dendrogram, but in a few cases genetic associations involved two distinct quadrants, for example, 2q31.1 and 3q21.3. Finally, in the online illustrations, we observed that facial effects propagated consistently in linked facial segments across different levels in the hierarchical design.

For completeness, we present an overview of all 38 loci identified in the discovery cohort with their peak SNP statistical details and imputation quality scores in Supplementary Tables 1 and 2, respectively. For 14 of the non-replicating loci, association was represented by a single SNP, and for an additional 5 loci the peak SNP had a minor allele frequency (MAF) less than 2%. Another locus that did not replicate showed a significant discrepancy in peak SNP MAF between the discovery and replication cohorts. Of the remaining 18 loci (also given in Table 1), 16 loci replicated at a nominal level of significance ($P \leq 0.05$), 15 of which replicated below the FDRd threshold and 12 of which replicated below the Bonferroni threshold. Two loci with robust associations in the discovery cohort failed to replicate. Additional illustrations on the discovery and replication results of all 38 loci are presented online (see URLs).

Association between 15 genomic regions and signatures of active regulatory elements in cranial neural crest cells. To explore biological processes and phenotypes associated with the identified GWAS loci, we performed Genomic Regions Enrichment of Annotations (GREAT) analysis[15]. Remarkably, even though only 15 loci were used in this analysis, we detected significant associations with craniofacial development, including categories such as chondrocyte differentiation, first and second branchial arch mesenchyme, development of facial bones and cleft palate (Supplementary

Fig. 3). Although craniofacial development is complex and involves interactions between multiple cell types, CNCCs, an embryonic cell group that arises at 3–6 weeks of gestation, have a central role in the formation of the facial plan and in determining its species-specific and individual variation[16–18]. During embryogenesis, CNCCs migrate away from the neural tube along stereotypical paths and form the majority of the cranial mesenchyme, which then differentiates into the bone, cartilage and connective tissue of the face and head[19,20]. We reasoned that, if divergence of the facial shape within the human population captured by our GWAS indeed originates early during embryogenesis, then the 15 loci should be preferentially active in CNCCs. To test this hypothesis, we took advantage of the epigenomic mapping datasets generated from human CNCCs, which were derived in vitro from human embryonic stem cells (ESCs) or induced pluripotent stem cells (iPSCs)[21,22]. Acetylation of histone H3 on lysine 27 (H3K27ac) has been associated with the promoters of transcriptionally active genes and with active distal enhancers, and thus is a useful mark to consider when exploring cell-type-specific activity of both coding and noncoding regions of the genome[23,24]. We quantified H3K27ac ChIP–seq signals in the vicinity (for example, within 10 kb) of the 15 peak SNPs in CNCCs of different genetic backgrounds and compared them to the H3K27ac signals over the same regions in over 30 other cell types, representing distinct adult, embryonic and in vitro–derived cell types. We observed higher H3K27ac signals in the vicinity of the peak SNPs in CNCCs than in any other cell type examined (Fig. 4). These observations are consistent with the preferential activity of the identified loci in CNCCs and with an embryonic origin of the human facial variation captured by our study.

A large majority of the genetic variants associated with human trait variation map to the noncoding parts of the genome. Many such variants are thought to reside within cis-regulatory elements[25]. However, cis-regulatory elements, especially distal enhancers, are often characterized by highly cell-type-specific activity patterns (reflected in their cell-type-specific chromatin marking patterns), underscoring the importance of analyzing relevant cell types for GWAS interpretation (reviewed in refs. [26,27]). To explore the overlap between the 15 loci and regulatory regions active in CNCCs, we used our epigenomic datasets (which included p300, TFAP2A, NR2F1, H3K4me1, H3K4me3, H3K27me3 and H3K27ac ChIP–seq data and nucleosomal depletion analysis by ATAC–seq) to identify and classify cis-regulatory elements on the basis of the type of element and activity. Specifically, transcription factor binding and transposase hypersensitivity maps were used to identify the genomic positions of candidate regulatory regions, whereas relative enrichments of H3K4me1 to H3K4me3 were used to distinguish enhancers from promoters and H3K27ac versus H3K27me3 signals were used to infer active, premarked or poised/repressed activity states (depicted in different colors in Fig. 5a). We observed that the peak SNPs were more likely to be located within or in the immediate vicinity of (<1 kb away from) a detectable regulatory chromatin feature in CNCCs ($P = 1.58 \times 10^{-3}$). Furthermore, these SNPs were enriched in enhancer elements, especially strong active CNCC enhancers ($P = 1.19 \times 10^{-2}$, odds ratio = 7.7, Fisher's exact test). An example of this is shown in Fig. 5b: peak SNP rs6740960[A>T] (affecting the lower facial quadrant) is located within an active CNCC enhancer marked by p300, H3K27ac and H3K4me1. Interestingly, the same variant has also recently been associated with an elevated risk of cleft lip and/or palate (CL/P) in a European cohort[28]. Furthermore, comparisons with our epigenomic datasets previously obtained in chimpanzee CNCCs[21] showed that the overlapping enhancer is biased toward chimpanzees in activity (as evidenced by the elevated levels of H3K27ac in chimpanzee as compared to human CNCCs). This example and others of loci (for example, *PAX3*) that showed association with facial variation, disease and enrichment for regulatory elements divergent between humans and chimpanzees raise

Table 1 | Genome-wide significant loci and replication results

Annotation					Effect		Discovery					Replication			
Chromosomal position	SNP	Candidate gene	Alleles	Lit. int.	Quadrant	Module	CC	CC2	P value	MAF	No. SNPs	β	P value	MAF	No. SNPs
Discovered and replicated results															
1p32.1 (60,997,570)	rs4916068	Intergenic	C<T		2	41	0.16	0.03	4.81×10^{-10}	0.500	28	0.25	1.39×10^{-12}	0.494	28
1p12 (119,762,175)	rs72691108	TBX15	A<G	ST5	4	7	0.25	0.06	5.81×10^{-14}	0.236	224	0.34	1.01×10^{-15}	0.243	224
1q31.3 (197,329,041)	rs2821116	**ASPM**	A<T	ST5[a]	3	13	0.24	0.06	2.63×10^{-19}	0.210	342	0.32	2.29×10^{-15}	0.209	342
2p21 (42,181,679)	rs6740960	PKDCC	T<A	ST4[a], ST5[a]	3	26	0.21	0.05	3.03×10^{-14}	0.497	2	0.17	3.44×10^{-6}	0.499	2
2q31.1 (177,111,819)	rs970797	HOXD cluster	T<G	ST4, ST5	2,3	38	0.17	0.03	6.17×10^{-11}	0.434	15	0.11	1.12×10^{-3}	0.428	4
2q36.1 (223,039,052)	rs10176525	PAX3	A<T	ST3[a], ST4[a], ST5[a]	2	10	0.21	0.05	3.76×10^{-14}	0.224	38	0.26	4.39×10^{-11}	0.233	38
3q21.3 (128,106,267)	rs2977562	**RAB7A, ACAD9**	G<A	ST4[a]	2,3	16	0.22	0.05	1.39×10^{-17}	0.244	389	0.20	8.63×10^{-7}	0.248	374
3q27.1 (184,333,169)	rs58022575	EPHB3, DVL3	G<GA	ST4[a], ST5[a]	2	20	0.17	0.03	7.99×10^{-10}	0.473	16	0.21	2.39×10^{-9}	0.467	16
4q31.3 (154,828,366)	rs9995821	DCHS2	C<T	ST3[a], ST4[a]	2	45	0.22	0.05	3.67×10^{-18}	0.200	10	0.39	7.33×10^{-19}	0.204	10
6p21.1 (44,681,840)	rs227833	SUPT3H	C<G	ST3[a], ST4[a], ST5[a]	2	10	0.21	0.05	3.38×10^{-14}	0.261	164	0.14	5.63×10^{-4}	0.254	150
6q23.2 (133,615,646)	rs5880172	**RPS12, EYA4**	A<ACT		4	30	0.21	0.04	6.10×10^{-12}	0.312	33	0.26	6.91×10^{-13}	0.297	33
7q21.3 (96,124,975)	rs10238953	**DLX6, DYNC1L1**	G<A	ST5[a]	3	55	0.24	0.06	1.06×10^{-22}	0.140	256	0.55	1.29×10^{-26}	0.137	256
17q24.3 (69,139,583)	rs11655006	**BC039327**	G<A	ST5[a]	2	2	0.27	0.07	5.24×10^{-21}	0.376	311	0.15	1.77×10^{-5}	0.364	308
17q24.3 (70,036,479)	rs5821892	SOX9	C<CG	ST5[a]	2	10	0.19	0.03	2.30×10^{-9}	0.427	31	0.13	3.01×10^{-4}	0.453	26
19q13.11 (34,290,995)	rs287104	**KCTD15**	G<A		2	41	0.17	0.03	1.26×10^{-10}	0.349	8	0.16	2.86×10^{-5}	0.329	8
Discovered results with weak evidence of replication															
5p14.3 (20,558,014)	rs62354288	**CDH18**	T<A		2	21	0.17	0.03	8.05×10^{-9}	0.355	4	0.02	7.04×10^{-1}	0.357	0
5q35.3 (177,760,576)	rs2913791	**NHP2, ZNF354A**	A<G		3	55	0.16	0.03	4.96×10^{-9}	0.429	5	0.00	9.14×10^{-1}	0.417	0
20p11.22 (21,626,627)	rs2424390	PAX1	A<T	ST3[a], ST4[a]	2	11	0.17	0.03	1.50×10^{-8}	0.233	15	0.09	2.47×10^{-2}	0.233	0

Properties of the top SNPs in 18 chromosomal regions showing genome-wide significance. The candidate genes for intragenic SNPs are indicated in bold. CC, canonical correlation; CC2, canonical correlation squared; β, regression coefficient; MAF, minor allele frequency; No. SNPs (discovery), the number of SNPs reaching genome-wide significance within the same locus; No. SNPs (replication), the number of SNPs reaching FDRd replication significance within the same locus; Lit. int., literature integration providing a link to supplementary tables reported in the manuscript (Supplementary Tables 3–5, as indicated). Additional details, such as the degrees of freedom for CC and the standard error for β, can be found in Supplementary Table 1.
The alleles are presented as minor allele < major allele.
aThe associated trait as described in the literature involved the facial segments identified, for example, nose protrusion–nose segments.

Fig. 2 | Fifteen replicating loci. Top, a polar dendrogram structurally corresponding to the polar dendrogram displaying the facial segments in Fig. 1; concentric circles represent loci reaching genome-wide significance for each segment. Bottom, Manhattan plot of all 63 facial segments combined illustrating the chromosomal position of the associated loci. The bottom (dotted) horizontal line represents the genome-wide significance threshold ($P = 5 \times 10^{-8}$), the middle (dashed) horizontal line represents the FDRd significance threshold ($P = 2.82 \times 10^{-8}$) and the top (solid) horizontal line represents study-wide significance ($P = 1.28 \times 10^{-9}$). The 15 loci are ordered by chromosome number, increasing from left to right, and are colored according to corresponding chromosome.

an intriguing possibility that the genetic variation within an overlapping set of loci/regulatory elements may influence both species-specific and individual facial shape variation in humans as well as determine predisposition to craniofacial malformations.

Candidate genes and integration with facial GWAS literature. A literature-based annotation of genes was performed, and their involvement in facial variation is reported in Supplementary Note 1. Additionally, in Supplementary Note 2, Supplementary Tables 3–6 and Supplementary Figs. 4 and 5, we first report statistical replication for 10 of 16 SNPs from six previous facial GWAS based on 2D or 3D images[2–7] (Supplementary Table 3). Second, we consulted GWAS Central and Phenoscanner[29] and report additional evidence of similar effects on factial shape for nine SNPs in Table 1 (Supplementary Table 4). Third, we cross-referenced ten of our loci with a recent

study[30] based on self-reported information on chin dimple and nose size (Supplementary Table 5). Table 1 indicates the involvement of each locus in any of these three integration efforts. In summary, for the 15 replicated loci in this work, 4 were completely new, 9 were previously discovered showing consistent facial effects and 2 were previously discovered with pleiotropic facial effects.

Discussion
A deeper understanding of the genetic basis of human facial traits may provide insights into the mechanisms of craniofacial morphogenesis, improve knowledge of the complex relationship between genotype and phenotype in craniofacial syndromes and birth defects, and eventually provide a basis for predicting facial features in numerous applications, ranging from early diagnosis to personalized medicine, treatment planning in craniofacial surgery and

Fig. 3 | Facial shape effects. Outer layer, for each locus, the $-\log_{10} P$ value of the lead SNP is shown for all 63 facial segments plotted on a polar dendrogram (ranging from 0 to the maximum $-\log_{10} P$ value as reported in Table 1). Facial segments with a P value lower than the genome-wide threshold ($P = 5 \times 10^{-8}$) are encircled. Inner layer, for each locus, its effect on a representative segment is illustrated using the normal displacement (displacement in the direction locally perpendicular to the facial surface). Blue and red indicate a local shape depression and protrusion, respectively, due to the minor allele at the SNP. Two loci (2q31.1 and 3q21.3) were associated with segments from separate quadrants of the polar dendrogram; hence, normal displacement is illustrated for representative segments in two areas of the face. Loci are ordered counterclockwise by chromosomal location.

orthodontics, and biometrics and forensics. In mice, whole-genome quantitative trait locus (QTL) studies have identified numerous genetic regions associated with craniofacial traits[31–33]. In humans, both candidate gene approaches[10,34,35] and GWAS[2–7] have implicated specific variants. While the highly complex nature of craniofacial morphogenesis suggests that many genes are likely to influence facial morphology, the number of associated loci uncovered so far using GWAS has been limited, with few independent replications.

Although traditional anthropometric facial measures remain a valuable resource for anthropologists and clinicians, because of their widespread use and simplicity, limited GWAS results have been obtained with such measurements so far. One reason for this is the difficulty inherent in defining proper facial types or relevant facial measurements a priori. As noted by Adhikari et al.[4], quantitative measures, on the basis of which all previous facial GWAS efforts from 2D or 3D images have been conducted, are expected to yield higher statistical power than categorical ones, but they have not yet resulted in many robust findings. Using well-defined categorical scale ratings instead, they discovered and replicated several associated loci, which we also replicate here. Similarly, but using self-reported categorizations on chin dimple and nose size, Pickrell et al.[30] identified a number of associated loci in an impressively sized sample ($n > 70,000$). However, the relatively simple categorical

phenotypes used in these studies preclude a comprehensive and precise description of the effects of genetic variants on facial morphology. We argue that these approaches are perfectly in line with the traditional phenotype-first thinking common to GWAS and remain a good paradigm for case–control designs and simple quantitative or qualitative traits, but perhaps are less suitable for multipartite quantitative traits like the human face.

The craniofacial complex is initially modulated by precisely timed embryonic gene expression and molecular and cellular processes mediated through complex signaling pathways[36]. As humans grow, hormones, nutritional status and biomechanical factors affect the face[37,38]. A natural consequence of these forces and constraints during facial morphogenesis and growth is that the face exhibits a modular organization, with suites of facial features at different scales that show internal integration but remain relatively independent from other features[39]. Therefore, we expect the human face to be influenced by many genes that exhibit a range of effects, with some influencing only localized parts of the face and others influencing more global aspects of morphology. As a result, the proposed facial segmentation is related to the concept of modularity and integration in morphological studies[40], with the main addition that the partitioning was done hierarchically, moving from globally integrated to locally focused modules. This allows for an investigation

Fig. 4 | Preferential activity in CNCCs. SNPs for facial variation are associated with signatures of active regulatory regions in human CNCCs. Shown are box plots of the normalized distribution of H3K27ac ChIP–seq signals in regions immediately surrounding (within 10 kb of) each of the 15 top replicated SNPs (or, as a control, randomized SNP positions; bottom) in various adult, embryonic and in vitro–derived cell types, as indicated. Samples corresponding to in vitro–derived CNCCs are highlighted in red and separated from other cell types. The box plot central line represents the median, the hinges represent the first and third quartiles, and the whiskers extend to 1.5 times the interquartile distance.

of facial shape effects propagating at different scales. On the basis of structural correlations between a vast number of individual 3D points on facial surfaces, we were able to hierarchically cluster the human face into segments. These were learned from the data and thus are unsupervised and data derived instead of candidate driven, and they could be duplicated across independent cohorts. This allowed for an efficient divide-and-conquer strategy enabling the screening of a huge set of facial variations without compromising statistical power.

Efficient use of phenotypic data from modest-sized cohorts is necessary for investigating 3D facial shape. In contrast to a recent GWAS on human body height[41], which, owing to the widespread availability of data on the phenotype, was able to identify 423 loci in combined samples over two orders of magnitude larger than what was used here, the specialized phenotype data needed for investigating 3D facial shape are limited to a few, smaller cohorts. Using the PITT cohort from Shaffer et al.[6], we managed to substantially increase the number of associated loci identified.

The global and local facial patterns as depicted in Fig. 3 for the discovered loci may help elucidate their roles during craniofacial development. It is clear that many of the genes at these loci are expressed in relevant tissues during embryonic development. Connecting these patterns of early expression and function to eventual morphology is a major challenge. The kind of highly refined facial effects presented here can help provide a roadmap to clarify the connections between molecules and morphology. Consider, for instance, two loci associated with different aspects of nasal shape, which are depicted in Supplementary Fig. 6. 19q13.11 (*KCTD15*) showed a highly focal effect limited to the nasal tip. In contrast, 6p21.1 (*SUPT3H*) affected primarily the nasal root and lateral parts of the nasal bridge, with its effects effectively sparing the nasal tip. A similar nasal phenotype was observed by Adhikari et al.[4] at 6p21.1. These highly specific phenotypic effects might provide important insights about the role of these two genes during human facial morphogenesis and growth. *KCTD15* has been shown to regulate *TFAP2A*, which has a critical role in neural
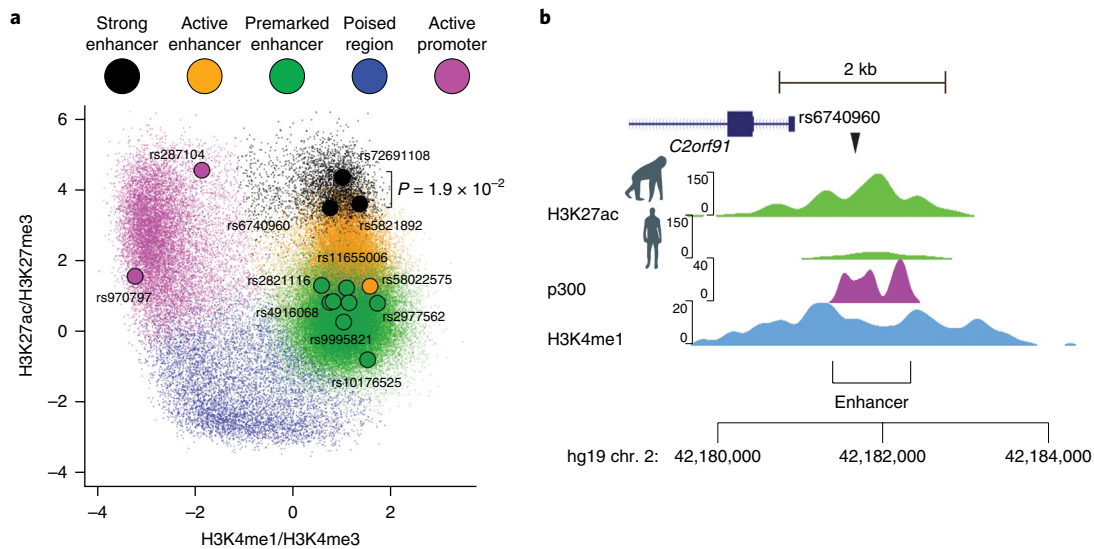
Fig. 5 | Regulatory regions active in CNCCs. a, Top replicated SNPs are enriched in the vicinity of active CNCC enhancer regions. Plotted is the promoter/enhancer score, corresponding to the ratio of H3K4me1 to H3K4me3 ChIP–seq signal (ordinate), versus the activity score, corresponding to the ratio of H3K27ac to H3K27me3 ChIP–seq signal (abscissa). Small dots represent all putative regulatory regions identified in CNCCs, and large dots highlight the chromatin regions closest to each of the significant SNPs. Colors represent the classes and activity states of regulatory elements: black, strong active enhancers; orange, weakly active enhancers; green, primed enhancers; magenta, promoters; blue, poised/repressed regulatory regions. Peak SNPs are enriched in strong active enhancers ($P = 1.19 \times 10^{-2}$, odds ratio = 7.7, Fisher's exact test). In addition, SNPs that retested with combined FDR < 0.001 are more likely to be close to (within 1 kb of) a detectable regulatory chromatin feature in CNCCs ($P = 1.58 \times 10^{-3}$, Fisher's exact test). **b,** A significant SNP, rs6740960[A>T], associated with lower-face variation is located within an active CNCC enhancer. Plotted are ChIP–seq tracks for H3K27ac (green), the p300 acetyltransferase (purple) and H3K4me1 (blue). Human and chimpanzee H3K27ac tracks are indicated with silhouette icons. The position of the rs6740960[A>T] SNP is indicated by an arrowhead and that of the enhancer with a bracket. Track heights are normalized enrichment scores from QuEST 2.0.

crest formation[42] and, when mutated, results in reduced snout length in mice, among other defects[43]. Perhaps *KCTD15* affects nasal tip shape in humans by influencing chondrocyte proliferation in the nasal septum, whereas *SUPT3H* exerts its influence on nasal shape by affecting portions of the maxilla and nasal bones. The precision of the phenotypic effects identified here can form the basis of such testable hypotheses in the laboratory. As another example, previous research extensively described the role of *SOX9* in the development of a broad range of tissues[44]. For facial development specifically, *SOX9* is expressed in CNCCs populating the pharyngeal arches in the head region. Previous animal and developmental studies described that the first pharyngeal arch gives rise to the maxillary and mandibular prominences[45]. However, the results of our study show that the locus associated with *SOX9* (17q24.3) and the one located 1 Mb away (*CASC17*, which is known to functionally interact with *SOX9* through transcription factors) both influence nasal shape, although the nose is known to be formed from the frontonasal prominence. This raises a fascinating question about the involvement of *SOX9* during the development of the nose and is an interesting subject for functional follow-up studies.

In our search through web-based GWAS repositories, we did find associations ($P < 1 \times 10^{-7}$) with non-facial traits for the loci in Table 1, including aspects of body size and composition (waist–hip ratio and height, *TBX15–WARS2* and *SUPT3H*, respectively), developmental traits like age at menarche (*RAB7A*) and brain DNA methylation levels (*RPS12–EVA4*), and risk of biliary atresia (*PAX3*). That genes and, more specifically, even relatively small gene regions, such as those identified in GWAS, have multiple functions in the body, or pleiotropy, has been appreciated for some time[46]. However, the extent to which pleiotropy has a role in trait variation and disease risk is perhaps much greater than what has previously been appreciated[47].

In conclusion, we proposed a data-driven global-to-local facial phenotyping approach, well suited for a genome-wide association scan. Using this approach, we have substantially advanced the literature on facial genetics on several fronts. First, we identified and replicated a number of new associated genetic loci using modest-sized cohorts. Second, we provided additional support for numerous previously identified loci and showed a strong integration of our results with the facial GWAS literature. Third, we demonstrated the preferential activity of the replicated loci in CNCCs, consistent with an embryonic origin of the human facial variation captured by our study. Lastly, we identified patterns of global-to-local genetic effects on facial shape, supporting the genetic organization of facial features at different scales.

URLs. National Institute for Dental and Craniofacial Research, http://www.nidcr.nih.gov/; National Human Genome Research Institute, https://www.genome.gov/; National Institute of Justice, https://www.nij.gov/funding/Pages/welcome.aspx; US Department of Defense, http://www.defense.gov/; MeshMonk facial mapping software, https://github.com/TheWebMonks/meshmonk; GREAT analysis, http://great.stanford.edu/; database of Genotypes and Phenotypes (dbGaP), http://www.ncbi.nlm.nih.gov/gap; facial images for the PITT cohort, https://www.facebase.org/data/record/#1/isa:dataset/RID=14283; GWAS Central, http://www.gwascentral.org/; Table Browser of the UCSC Genomes Browser, http://genome.ucsc.edu/cgi-bin/hgTables; OMIM database, https://www.omim.org/. All relevant study and results data to run future replications and meta-analysis efforts are provided at https://www.esat.kuleuven.be/psi/research/global-to-local-facial-phenotyping.

Methods

Methods, including statements of data availability and any associated accession codes and references, are available at https://doi.org/10.1038/s41588-018-0057-4.

References

1. Roosenboom, J., Hens, G., Mattern, B. C., Shriver, M. D. & Claes, P. Exploring the underlying genetics of craniofacial morphology through various sources of knowledge. *BioMed Res. Int.* **2016**, 3054578 (2016).
2. Paternoster, L. et al. Genome-wide association study of three-dimensional facial morphology identifies a variant in *PAX3* associated with nasion position. *Am. J. Hum. Genet.* **90**, 478–485 (2012).
3. Liu, F. et al. A genome-wide association study identifies five loci influencing facial morphology in Europeans. *PLoS Genet.* **8**, e1002932 (2012).
4. Adhikari, K. et al. A genome-wide association scan implicates *DCHS2*, *RUNX2*, *GLI3*, *PAX1* and *EDAR* in human facial variation. *Nat. Commun.* **7**, 11616 (2016).
5. Cole, J. B. et al. Genomewide association study of African children identifies association of *SCHIP1* and *PDE8A* with facial size and shape. *PLoS Genet.* **12**, e1006174 (2016).
6. Shaffer, J. R. et al. Genome-wide association study reveals multiple loci influencing normal human facial morphology. *PLoS Genet.* **12**, e1006149 (2016).
7. Lee, M. K. et al. Genome-wide association study of facial morphology reveals novel associations with *FREM1* and *PARK2*. *PLoS One* **12**, e0176566 (2017).
8. Abecasis, G. R. et al. An integrated map of genetic variation from 1,092 human genomes. *Nature* **491**, 56–65 (2012).
9. Snyders, J., Claes, P., Vandermeulen, D. & Suetens, P. Development and comparison of non-rigid surface registration and extensions (technical report KUL/ESAT/PSI/1401) 1–55 (2014).
10. Claes, P. et al. Modeling 3D facial shape from DNA. *PLoS Genet.* **10**, e1004224 (2014).
11. Ng, A. Y., Jordan, M. I. & Weiss, Y. On spectral clustering: analysis and an algorithm. in *Proc. 14th International Conference Neural Information Processing Systems: Natural and Synthetic* 849–856 (MIT Press, Cambridge, MA, 2001).
12. Hayton, J. C., Allen, D. G. & Scarpello, V. Factor retention decisions in exploratory factor analysis: a tutorial on parallel analysis. *Organ. Res. Methods* **7**, 191–205 (2004).
13. Thompson, B. Canonical correlation analysis. in *Encyclopedia of Statistics in Behavioral Science* (eds. Everitt, B. & Howell, D.) (John Wiley and Sons, Hoboken, NJ, 2005).
14. Benjamini, Y. & Yekutieli, D. The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.* **29**, 1165–1188 (2001).
15. McLean, C. Y. et al. GREAT improves functional interpretation of cis-regulatory regions. *Nat. Biotechnol.* **28**, 495–501 (2010).
16. Bronner, M. E. & LeDouarin, N. M. Development and evolution of the neural crest: an overview. *Dev. Biol.* **366**, 2–9 (2012).
17. Green, S. A., Simoes-Costa, M. & Bronner, M. E. Evolution of vertebrates as viewed from the crest. *Nature* **520**, 474–482 (2015).
18. Schneider, R. A. Regulation of jaw length during development, disease, and evolution. *Curr. Top. Dev. Biol.* **115**, 271–298 (2015).
19. Cordero, D. R. et al. Cranial neural crest cells on the move: their roles in craniofacial development. *Am. J. Med. Genet. A* **155A**, 270–279 (2011).
20. Kulesa, P. M. & McLennan, R. Neural crest migration: trailblazing ahead. *F1000Prime Rep.* **7**, 02 (2015).
21. Prescott, S. L. et al. Enhancer divergence and cis-regulatory evolution in the human and chimp neural crest. *Cell* **163**, 68–83 (2015).
22. Rada-Iglesias, A. et al. Epigenomic annotation of enhancers predicts transcriptional regulators of human neural crest. *Cell Stem Cell* **11**, 633–648 (2012).
23. Rada-Iglesias, A. et al. A unique chromatin signature uncovers early developmental enhancers in humans. *Nature* **470**, 279–283 (2011).
24. Creyghton, M. P. et al. Histone H3K27ac separates active from poised enhancers and predicts developmental state. *Proc. Natl. Acad. Sci. USA* **107**, 21931–21936 (2010).
25. Tak, Y. G. & Farnham, P. J. Making sense of GWAS: using epigenomics and genome engineering to understand the functional relevance of SNPs in non-coding regions of the human genome. *Epigenetics Chromatin* **8**, 57 (2015).
26. Long, H. K., Prescott, S. L. & Wysocka, J. Ever-changing landscapes: transcriptional enhancers in development and evolution. *Cell* **167**, 1170–1187 (2016).
27. Corradin, O. & Scacheri, P. C. Enhancer variants: evaluating functions in common disease. *Genome Med.* **6**, 85 (2014).
28. Ludwig, K. U. et al. Imputation of orofacial clefting data identifies novel risk loci and sheds light on the genetic background of cleft lip ± cleft palate and cleft palate only. *Hum. Mol. Genet.* **26**, 829–842 (2017).
29. Staley, J. R. et al. PhenoScanner: a database of human genotype–phenotype associations. *Bioinformatics* **32**, 3207–3209 (2016).
30. Pickrell, J. K. et al. Detection and interpretation of shared genetic influences on 42 human traits. *Nat. Genet.* **48**, 709–717 (2016).
31. Klingenberg, C. P. Quantitative genetics of geometric shape: heritability and the pitfalls of the univariate approach. *Evolution* **57**, 191–195 (2003).
32. Pallares, L. F. et al. Mapping of craniofacial traits in outbred mice identifies major developmental genes involved in shape determination. *PLoS Genet.* **11**, e1005607 (2015).
33. Maga, A. M., Navarro, N., Cunningham, M. L. & Cox, T. C. Quantitative trait loci affecting the 3D skull shape and size in mouse and prioritization of candidate genes in-silico. *Front. Physiol.* **6**, 92 (2015).
34. Coussens, A. K. & van Daal, A. Linkage disequilibrium analysis identifies an *FGFR1* haplotype-tag SNP associated with normal variation in craniofacial shape. *Genomics* **85**, 563–573 (2005).
35. Peng, S. et al. Detecting genetic association of common human facial morphological variation using high density 3D image registration. *PLoS Comput. Biol.* **9**, e1003375 (2013).
36. Carlson, D. S. Theories of craniofacial growth in the postgenomic era. *Semin. Orthod.* **11**, 172–183 (2005).
37. Sperber, G. H., Guttmann, G. D. & Sperber, S. M. *Craniofacial Development (Book for Windows & Macintosh)* (PMPH-USA, Hamilton-London, 2001).
38. Williams, S. E. & Slice, D. E. Regional shape change in adult facial bone curvature with age. *Am. J. Phys. Anthropol.* **143**, 437–447 (2010).
39. Mitteroecker, P. & Bookstein, F. The evolutionary role of modularity and integration in the hominoid cranium. *Evolution* **62**, 943–958 (2008).
40. Klingenberg, C. P. Morphological integration and developmental modularity. *Annu. Rev. Ecol. Evol. Syst.* **39**, 115–132 (2008).
41. Wood, A. R. et al. Defining the role of common variation in the genomic and biological architecture of adult human height. *Nat. Genet.* **46**, 1173–1186 (2014).
42. Zarelli, V. E. & Dawid, I. B. Inhibition of neural crest formation by Kctd15 involves regulation of transcription factor AP-2. *Proc. Natl. Acad. Sci. USA* **110**, 2870–2875 (2013).
43. Ahituv, N. et al. An ENU-induced mutation in *AP-2α* leads to middle ear and ocular defects in Doarad mice. *Mamm. Genome* **15**, 424–432 (2004).
44. Lee, Y. H. & Saint-Jeannet, J. P. *Sox9* function in craniofacial development and disease. *Genesis* **49**, 200–208 (2011).
45. Francis-West, P. H., Robson, L. & Evans, D. J. *Craniofacial Development: The Tissue and Molecular Interactions That Control Development of the Head* (Springer Science & Business Media, Berlin-Heidelberg, 2012).
46. Stearns, F. W. One hundred years of pleiotropy: a retrospective. *Genetics* **186**, 767–773 (2010).
47. Boyle, E. A., Li, Y. I. & Pritchard, J. K. An expanded view of complex traits: from polygenic to omnigenic. *Cell* **169**, 1177–1186 (2017).

Author contributions

This work is the result of a four-center collaboration between KU Leuven, the University of Pittsburgh, Penn State University and the Stanford University School of Medicine, led by P.C., S.M.W., M.D.S. and J.W., respectively. All four centers contributed equally to this work. P.C. and D.S. with input from P.S. and D.V. conceptualized and implemented the global-to-local facial segmentations. J.L. provided genomic data support and analysis at the KU Leuven site and together with P.C. ran the GWAS. M.K.L., J.R., E.J.L., J.C.C., E.O.,

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at https://doi.org/10.1038/s41588-018-0057-4.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to P.C. or J.W. or M.D.S. or S.M.W.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Methods

Sample and recruitment details. For the discovery cohort (the Pittsburgh (PITT) sample), data from 2,449 participants were obtained from the 3D Facial Norms repository[48]. The repository includes 3D facial surface images and self-reported demographic descriptors (for example, age and ancestry) as well as basic physical characteristics (for example, height and weight) from individuals recruited at four US sites: Pittsburgh, PA; Seattle, WA; Houston, TX; and Iowa City, IA. Recruitment was limited to individuals aged 3 to 40 years and of self-reported European ancestry. Individuals were excluded if they reported a personal or family history of any birth defect or syndrome affecting the head or face, a personal history of any significant facial trauma or facial surgery, or any medical condition that might alter the structure of the face. A total of 2,329 participants were retained for analysis, after removing 120 participants, with missing information on sex, age, height, weight or with 3D image mapping artifacts ($n = 22$).

For the replication cohort (the Penn State (PSU) sample), participants were recruited through several studies at the Pennsylvania State University and sampled in the following locations: State College, PA; New York, NY; Urbana-Champaign, IL; Twinsburg, OH; Dublin, Ireland; Rome, Italy; Warsaw, Poland; and Porto, Portugal. Participants self-reported information on age, ethnicity, ancestry and body characteristics, and data were gathered on height, weight and pigmentation of the hair and skin. Individuals were excluded from the analysis if they were below 18 years of age (range 18–88) and if they reported a personal history of significant facial trauma or facial surgery, or any medical condition that might alter the structure of the face. No restriction on ancestry or ethnicity was imposed during recruitment, but only individuals of European descent were used in this study ($n = 2,059$). Participants were removed because of missing sex, age, height or weight information ($n = 81$) or the presence of 3D image mapping artifacts ($n = 33$). A further reduction to $n = 1,719$ was done by excluding participants who were not included in the genotype imputation efforts.

Facial phenotyping, 3D imaging quality control and shape variables. *3D facial surface imaging.* Digital facial stereophotogrammetry was used to capture 3D facial surfaces in both samples. This well-established approach uses digital photography to generate a dense 3D point cloud representing the surface geometry of the face from multiple 2D images with overlapping fields of view. For the Pittsburgh sample, facial surfaces were acquired using 3dMDface camera systems. For the Penn State sample, 3D images were obtained with 3dMDface or Vectra H1 (Canfield Scientific) systems. Applying standard facial image acquisition protocols[49], participants were asked to close their mouths and hold their faces with a neutral expression.

Spatially dense facial quasi-landmarking. 3D images in wavefront.obj file format were imported into an in-house 3D image-cleaning program for cropping and trimming, removing hair, ears and any dissociated polygons. Five crude positioning landmarks were placed on the face to establish a rough facial orientation. An anthropometric mask (AM)[50] was non-rigidly mapped[9] onto all 3D images and their reflections, which were constructed by changing the sign of the x coordinate[51]. The AM is essentially a surface template covering the facial area of interest, and the mapping thereof onto the facial images is a process equivalent to the indication of traditional landmarks. This establishes homologous spatially dense (~10,000) quasi-landmark (QL) configurations for all 3D images and their reflections. In other words, image data from different individuals became standardized, enabling a spatially dense analysis.

Facial size was obtained as the centroid size of the quasi-landmark configurations. Facial shape was symmetrized using GPA[52] to eliminate differences in position, orientation and size of both original and reflected quasi-landmark configurations. The average of an original and its reflected quasi-landmark configuration constitutes the symmetric component, while the difference between the two configurations constitutes the asymmetric component. Although of interest, in this work we currently ignore variations in facial asymmetry. Therefore, unless otherwise mentioned, when discussing facial shape we always refer to the symmetric quasi-landmark configuration.

Facial quality control. Outlier faces, due to quasi-landmark mapping errors, were detected by measuring the Mahalanobis distance for each face to the overall average face in the symmetrized shape space spanned by an orthogonal basis of principal components that captures 98% of the total variation. From the distribution of Mahalanobis distances, a z score for each facial shape was established, and each face with a z score equal to or larger than 2 was manually inspected for quasi-landmark errors. Identified erroneous faces were removed, and the whole process starting from the generalized Procrustes superimposition of both original and reflected quasi-landmark configurations was repeated. The PSU and PITT datasets were processed separately with the same AM, resulting in error-free symmetric and compatible (homologous) quasi-landmark configurations across both datasets.

Phenotyping the discovery panel. The PITT cohort served as the discovery panel. First, the superimposed and symmetrized facial shapes were corrected using a partial least-squares regression (PLSR; function plsregress from Matlab 2016b)

for the confounders of sex, age, age[2], weight, height, facial size and the first four genetic PCA axes to correct for population stratification. Of note, while correcting for genetic background and other confounders is a required step in genetic mapping efforts, it additionally ensures that structural facial variations due to the confounders do not influence the global-to-local segmentation of faces. Second, a 3D correlation, using the RV coefficient[53], between each pair of corrected quasi-landmarks was computed to construct the squared similarity matrix (~10,000 × ~10,000). Subsequently, a hierarchical spectral clustering with five levels was performed, resulting in a total of 63 hierarchically linked facial segments with 1, 2, 4, 8, 16 and 32 non-overlapping modules at levels 0, 1, 2, 3, 4 and 5 (Fig. 1). The hierarchical design provides a stepwise focused shape analysis at different scales, by gradually zooming in to more local shape variations without ignoring the integration of facial parts at previous levels. Furthermore, specific shape effects should propagate consistently in facial segments linked across different levels, lending additional support that the discovered signals are robust. At each level, for each segment, all quasi-landmarks in the segment are subjected to a new GPA. As such, a shape space for each facial module is constructed independently of the other modules and its relative positioning within the full face. This is particularly interesting for smaller shape variations (for example, nose tip) that when superimposed using the full facial surface are overlooked and become undetectable. Finally, after GPA, each modular shape space is spanned by an orthogonal basis using PCA combined with parallel analysis[12] to determine the number of significant components contributing to facial shape.

This application of PCA had several advantages. First, PCs form an orthogonal basis, such that shape variations could be described as linear combinations of PCs. Second, by selecting significant PCs, it was possible to eliminate (via parallel analysis[12]) noisy or meaningless shape variations that result from various sources of error, such as 3D image acquisition and/or quasi-landmark mapping. Third, PCA is an excellent tool to reduce the dimensions of high-dimensional data when strong correlations between the individual data elements are present. The construction of facial modules, as performed here, resulted in the grouping of highly correlated quasi-landmarks, such that a substantial reduction of dimensions was obtained. This in turn enabled the use of multivariate association techniques, such as CCA, that are constrained in the number of variables tested as a function of the sample size. For almost all segments and especially for the larger segments containing up to ~10,000 quasi-landmarks, this constraint would otherwise have been violated. Furthermore, the multivariate association was more computationally tractable given the lower number of variables tested. Finally, an effect identified in PCA space could be transformed back into the quasi-landmark shape space.

Note that some prior studies[2,3,5] also analyzed facial shape data using PCs, but they did so by analyzing them separately. Basically, in these studies, full facial multivariate sparse landmark configurations were projected onto each PC and thus the phenotypic data were preselected along the PCs individually and one by one. Any preselection of measurements, either distances or PCs that are subsequently analyzed individually results in a loss of information, as combinations of these measurements are not considered. In contrast, our work uses all PCs together in a single multivariate association effort based on CCA. Therefore, any possible linear combination of PCs is investigated simultaneously and information is not lost. In CCA, 'canonical' is the statistical term for analyzing latent variables (which are not directly observed) that represent multiple variables (which are directly observed). Such an optimization and construction of latent variables is absent in a linear regression on an individual PC.

Phenotyping the replication panel. The PSU cohort was used as the replication panel. First, independently of the PITT dataset, the superimposed and symmetrized facial shapes were corrected, using a PLSR, for the confounders of sex, age, age[2], weight, height, facial size and the first three PCA axes of genetic background. The reason for this separate and independent correction was twofold: (i) both cohorts have independent axes of genetic background and (ii) this workflow shows how future replication efforts can be done without the need to merge, and therefore have access to, all the genetic and covariate data from the discovery cohort. The establishment of spatially dense quasi-landmarks (AM mapping) allowed for consistent phenotyping. After AM mapping, the PSU faces were segmented using the quasi-landmark clustering from the discovery PITT cohort. Finally, for each facial segment, all PSU shape instances were superimposed onto the segmental average shape and PCs of the discovery panel. In this way, the replication effort only required the discovery quasi-landmark cluster labels, average shapes and PCs for each of the facial segments.

Genotyping, quality checks, imputation, population structure and annotation. *The Pittsburgh sample.* Genotyping was performed at the Center for Inherited Disease Research (CIDR) at Johns Hopkins University. Participants, including 70 duplicate samples and 72 HapMap control samples, were genotyped on the Illumina OmniExpress + Exome v1.2 array plus 4,322 investigator-chosen SNPs included to capture variation in specific regions of interest based on previous studies of the genetics of facial variation. Standard data cleaning and quality assurance procedures[54] were performed in collaboration with the University of Washington Genetics Coordinating Center. Specifically, samples were evaluated for concordance of genetic and reported sex, evidence of chromosomal aberrations,

biological relatedness across study participants, ancestry, genotype call rate and batch effects. SNPs were evaluated for call rate, discordant genotype calls between duplicate samples, Mendelian errors in HapMap control parent–offspring trios, deviation from Hardy–Weinberg genotype proportions and sex differences in allele frequency and heterozygosity.

Imputation of unobserved genetic variants and sporadic missing genotype calls for assayed SNPs was performed using the 1000 Genomes Project[8] Phase 3 reference panel. SHAPEIT2[55] was used for prephasing of haplotypes, and IMPUTE2[56,57] was used to impute nearly 35 million variants. SNP-level (INFO score > 0.5) and genotype-per-participant-level (genotype probability > 0.9) filters were used to omit poorly imputed variants. Masked variant analyses, in which genotyped SNPs were imputed as though they had not been assayed, indicated high concordance between imputed and observed genotypes: 0.982 for SNPs with MAF ≥0.05 and 0.998 for SNPs with MAF <0.05.

Population structure was assessed using PCA of approximately 97,000 autosomal genotyped SNPs chosen for call rate (>95%), MAF (>0.05) and LD (pairwise $r^2 < 0.1$ across variants in a sliding window of 10 Mb). Tests of genetic association between the first 20 PCs and all SNPs confirmed that PCs did not represent local variation at specific genetic loci. On the basis of the scree plot and joint distributions, four PCs are sufficient to capture population structure within the PITT sample.

The Penn State sample. Participants sampled from 2006–2012 (IRB 32341) were genotyped on the Illumina Human Hp200c1 BeadChip. Participants sampled from 2013–2016 (IRB 44929, 13103, 2503 and 4320) were genotyped on the 23andMe v3 and v4 arrays. Samples were evaluated for concordance of genetic and reported sex, evidence of chromosomal aberrations, biological relatedness across study participants, ancestry, genotype call rate and batch effects. SNPs were evaluated for call rate, discordant genotype calls between duplicate samples, Mendelian errors in HapMap control parent–offspring trios, deviation from Hardy–Weinberg genotype proportions and sex differences in allele frequency and heterozygosity.

Using the 1000 Genomes Project[8] Phase 3 reference panel, samples with > 500,000 variants were imputed, as fewer variants result in uncertain imputation probabilities. SHAPEIT2[55] was used for prephasing of haplotypes, and IMPUTE2[56,57] was used to impute nearly 35 million variants. SNP-level (INFO score > 0.5) and genotype-per-participant-level (genotype probability > 0.9) filters were used to omit poorly imputed variants.

To select the Europeans used in this study, the HapMap 3 dataset ($n=998$) was merged with the non-imputed genotypes, and genetic ancestry was estimated using the ADMIXTURE program[58] assuming $k=3$ to 9 ancestral populations. On the basis of the cross-validation (CV) error for each k value, we selected a value of $k=6$ as appropriate for the dataset. These results were then used to provide a matrix of genetic ancestry axes, which were used to select samples most closely related to the CEU and TSI population references.

Additional population structure was assessed using PCA of approximately 100,000 autosomal genotyped SNPs chosen for call rate (>95%), MAF (>0.05) and LD (pairwise $r^2 < 0.1$ across variants in a sliding window of 10 Mb). Tests of genetic association between the first 20 PCs and all SNPs confirmed that PCs did not represent local variation at specific genetic loci. on the basis of the scree plot and joint distributions, three PCs are sufficient to capture population structure within the European-only PSU sample.

Top SNPs and annotation. We observed 1,932 genome-wide significant SNPs across 38 loci using a 500-kb window. For each locus, a top SNP was defined as the SNP generating the highest association (lowest P value) in any of the 63 facial segments. Genes 500 kb up- and downstream of the top SNPs were identified using the Table Browser of the UCSC Genome Browser (see URLs). Each gene was annotated on the basis of the literature found in PubMed. As a search term, only the gene name was used, and relevant articles were selected based on their title and abstract. The OMIM database (see URLs) was searched for syndromes associated with the annotated genes.

Statistics and epigenomic analysis. *SNPs—statistical association.* Our global-to-local facial phenotyping partitioned facial shape into 63 facial segments, each of which was represented by multiple dimensions of variation (PCs). CCA[13], as implemented in the function canoncorr from Matlab 2016b, was used as a straightforward multivariate testing framework (note that CCA is also implemented in PLINK 1.9[59]). CCA extracts the linear combination of PCs from a facial segment that has maximal correlation with the SNP being tested. The correlation is tested for significance based on Rao's F-test approximation[60] (right-tail, one-sided test). Using CCA, we tested each SNP ($n=9,478,608$) individually under the additive genetic model in the PITT cohort ($n=2,329$). Note that CCA does not accommodate adjustments for covariates, but effects of important variables such as sex, age, height, weight, facial size and genetic ancestry were corrected for (using PLSR) at the phenotyping stage. Additionally, we applied a similar correction on each SNP, again using PLSR. Therefore, the CCA analysis was performed under the reduced model, which was obtained after removing the effects of covariates on both the independent (SNP) as well as the dependent (facial shape) variables. The minimum MAF cutoff for SNPs to test

was 1%. Quantile–quantile plots, provided online, indicate that the population stratification under this reduced model was dealt with properly.

Multiple-testing corrections. Given the burden of multiple comparisons, a strict significance threshold of $P \leq 5 \times 10^{-8}$ was used to declare 'genome-wide significance', which corresponds to a Bonferroni correction for 1 million independent tests and mostly applicable in a European GWAS cohort[61]. Given that we tested facial variation as one of many facial segments separately, the multiple-comparisons burden was magnified. Therefore, we also determined a more stringent threshold for declaring 'study-wide significance' corresponding to an additional adjustment for the effective number of independent tests[62]. The eigenvalues of pairwise multivariate correlations of all 63 modules determined a total of 39 effective independent tests. This reduction in effective tests was expected because of the hierarchical and overlapping construction of the facial segments. Therefore, the study-wide significance threshold was determined to be 1.2821×10^{-9} (i.e., $5 \times 10^{-8}/39$). Additionally, we computed an FDR-adjusted significance threshold of Benjamini and Yekutieli[14] that is accurate for any test dependency structure.

SNPs—statistical replication. On the basis of previous work[10], the effect found in the discovery cohort was measured directly in the replication cohort ($n=1,719$), which was followed by a univariate analysis. First, the discovery cohort was used as a phenotyping reference for the replication cohort. Facial shapes of the replication cohort were projected onto the PCs of each of the 63 facial segments. Subsequently, they were further projected onto the SNP-related CCA loadings, constructing specific genetic effect scores[10]. In doing so, the phenotypic trait, once discovered, was fixed and explicitly measured in the replication cohort. This resulted in a univariate phenotypic measure that was tested for association using a standard linear regression. We report the regression coefficient, standard error and a P value (two-sided) based on a t statistic for regression coefficients using the function regstats from Matlab 2016b. This was done separately for each combination of SNP ($n=1,931$) and facial segment that achieved genome-wide significance. From all replication efforts combined ($n=7,467$), we computed an FDR-adjusted significance threshold[14].

SNPs—tissue-specific enhancer association. ChIP–seq datasets from human and chimpanzee CNCCs were obtained from Prescott et al.[21] and compared to 72 publicly available H3K27ac ChIP–seq datasets corresponding to 31 cell types/tissues (as indicated in Fig. 4a) and downloaded from the Sequence Read Archive (SRA) repository. Normalized H3K27ac ChIP coverage was calculated for 800,000 genomic regions corresponding to a superset of ENCODE DNase-hypersensitive sites (DHSs) and transcription factor–binding sites. Next, all intervals in the 20-kb windows centered on the 15 top replicated SNPs were selected (intersectBed), and a box plot (using the boxplot function in R with default interval settings) of the logarithm of normalized ChIP coverage was plotted. As a control, a five times larger set of randomized SNP positions was used in the same procedure.

SNPs—CNCC chromatin state association and statistics. Identification of different classes of CNCC regulatory regions based on chromatin signatures was performed as described in Prescott et al.[21]. In brief, 106,000 human CNCC candidate regulatory regions were called on the basis of combined p300, AP2α and NR2F1 ChIP–seq and ATAC–seq data by combining MACS2-called peaks into one set using mean-shift clustering. Normalized coverage around these regions was obtained for H3K4me1, H3K4me3, H3K27ac and H3K27me3 histone modifications. To visualize chromatin states at the candidate regions, the log of the H3K4me1 to H3K4me3 signal ratio was plotted versus the log of the ratio of H3K27ac to H3K27me3 as a scatterplot and colored by k-means ($k=5$) cluster identity, corresponding to promoters, active enhancers, poised regions and premarked regions. Regulatory regions closest to the top replicated GWAS (closestBed function) SNPs were identified and are highlighted by circles. We observed 3 of 15 lead SNPs to be closest to a region with an active chromatin state (distance less than 9 kb in each case). This represents a significant enrichment for active chromatin–lead SNP association under Fisher's exact test (contingency table: 3 active regions proximal to SNP, 12 other regions proximal to SNP, 3,622 active regions not proximal to SNP, 102,694 other regions not proximal to SNP), exact $P=1.19 \times 10^{-2}$. We also compared replicated SNPs to non-replicated SNPs from the screen in terms of direct overlap (1-kb distance) with CNCC regulatory elements (contingency table: 26 non-replicated SNPs with no overlap, 0 non-replicated SNPs with overlap, 7 replicated SNPs with no overlap, 5 replicated SNPs with overlap), Fisher's exact test $P=1.58 \times 10^{-3}$.

Chromatin modification visualization in the genome browser. Kernel smoothed coverage files in wig format were generated with QuEST2.0[63] from datasets from Prescott et al.[21] and uploaded to the UCSC browser. Representative tracks were selected to make Fig. 5b concise and instructive.

Gene Ontology term enrichment analysis. Analysis was performed with GREAT 3.0.0[15]. Genomic positions for replicated SNPs were uploaded to http://great.

stanford.edu/, and the results of the analysis in terms of FDR and enrichments for each significant category are summarized in Supplementary Fig. 4.

Cell line sources. Cells from WiCell (H9 ESC) and the laboratory of F. Gage (iPSC; Salk Institute) were authenticated by analyzing genomic sequence data from the cell lines. PCR tests were done to test for mycoplasma contamination.

Ethics statement. Institutional review board (IRB) approval was obtained at each recruitment site, and all participants gave their written informed consent before participation; for children, written consent was obtained from a parent or legal guardian. For the Pittsburgh sample, the following local ethics approvals were obtained: University of Pittsburgh IRB PRO09060553 and RB0405013; UT Health Committee for the Protection of Human Subjects HSC-DB-09-0508; Seattle Children's IRB 12107; University of Iowa Human Subjects Office/IRB 200912764 and 200710721. For the Penn State sample, the following local ethics approvals were obtained: State College, PA (IRB 44929 and 4320); New York, NY (IRB 45727); Urbana-Champaign, IL (IRB 13103); Dublin, Ireland; Rome, Italy; Warsaw, Poland; and Porto, Portugal (IRB 32341); Twinsburg, OH (IRB 2503).

Life Sciences Reporting Summary. Further information on experimental design is available in the Life Sciences Reporting Summary.

Data availability. All of the genotypic markers for the Pittsburgh dataset are available to the research community through the dbGaP controlled-access repository (see URLs) at accession phs000949.v1.p1. The raw source data for the phenotypes—the 3D facial surface models in .obj format—are available through the FaceBase Consortium (see URLs). Access to these 3D facial surface models requires proper institutional ethics approval and approval from the FaceBase data access committee. Additional details can be requested from S.M.W.

The participants making up the Penn State University dataset were not collected with broad data sharing consent. Given the highly identifiable nature of both facial and genomic information and unresolved issues regarding risk to participants, we opted for a more conservative approach to participant recruitment. Broad data sharing of these collections would thus be in legal and ethical violation of the informed consent obtained from the participants. This restriction is not because of any personal or commercial interests. Additional details can be requested from M.D.S.

KU Leuven provides the spatially dense facial mapping software, free to use for academic purposes: MeshMonk (see URLs). Matlab implementations of the hierarchical spectral clustering to obtain facial segmentations are available upon request from P.C. The statistical analyses in this work were based on functions of the statistical toolbox in Matlab 2016b as mentioned throughout the Methods.

All relevant data to run future replications and meta-analysis efforts are provided in Matlab format online (see URLs). This includes the anthropometric mask used, facial segmentation cluster labels, PCA shape spaces for all 63 facial segments in the PITT cohort, CCA loadings and all association statistics for the peak SNPs in Table 1. An example Matlab script to explore the data is provided as well.

References

48. Weinberg, S. M. et al. The 3D Facial Norms Database: Part 1. A web-based craniofacial anthropometric and image repository for the clinical and research community. *Cleft Palate Craniofac. J.* **53**, e185–e197 (2016).
49. Heike, C. L., Upson, K., Stuhaug, E. & Weinberg, S. M. 3D digital stereophotogrammetry: a practical guide to facial image acquisition. *Head Face Med.* **6**, 18 (2010).
50. Claes, P., Walters, M. & Clement, J. Improved facial outcome assessment using a 3D anthropometric mask. *Int. J. Oral Maxillofac. Surg.* **41**, 324–330 (2012).
51. Claes, P., Walters, M., Vandermeulen, D. & Clement, J. G. Spatially-dense 3D facial asymmetry assessment in both typical and disordered growth. *J. Anat.* **219**, 444–455 (2011).
52. Rohlf, F. & Slice, D. Extensions of the procrustus method for the optimal superimposition of landmarks. *Syst. Zool.* **39**, 40–59 (1990).
53. Robert, P. & Escoufier, Y. A unifying tool for linear multivariate statistical methods: the RV-coefficient. *Appl. Stat.* **25**, 257–265 (1976).
54. Laurie, C. C. et al. Quality control and quality assurance in genotypic data for genome-wide association studies. *Genet. Epidemiol.* **34**, 591–602 (2010).
55. Delaneau, O., Zagury, J.-F. & Marchini, J. Improved whole-chromosome phasing for disease and population genetic studies. *Nat. Methods* **10**, 5–6 (2013).
56. Howie, B. N., Donnelly, P. & Marchini, J. A flexible and accurate genotype imputation method for the next generation of genome-wide association studies. *PLoS Genet.* **5**, e1000529 (2009).
57. Howie, B., Marchini, J. & Stephens, M. Genotype imputation with thousands of genomes. *G3* **1**, 457–470 (2011).
58. Alexander, D. H., Novembre, J. & Lange, K. Fast model-based estimation of ancestry in unrelated individuals. *Genome Res.* **19**, 1655–1664 (2009).
59. Purcell, S. et al. PLINK: a tool set for whole-genome association and population-based linkage analyses. *Am. J. Hum. Genet.* **81**, 559–575 (2007).
60. Olson, C. L. On choosing a test statistic in multivariate analysis of variance. *Psychol. Bull.* **83**, 579 (1976).
61. Pe'er, I., Yelensky, R., Altshuler, D. & Daly, M. J. Estimation of the multiple testing burden for genomewide association studies of nearly all common variants. *Genet. Epidemiol.* **32**, 381–385 (2008).
62. Li, J. & Ji, L. Adjusting multiple testing in multilocus analyses using the eigenvalues of a correlation matrix. *Heredity* **95**, 221–227 (2005).
63. Valouev, A. et al. Genome-wide analysis of transcription factor binding sites based on ChIP-Seq data. *Nat. Methods* **5**, 829–834 (2008).

nature research

Corresponding author(s):   Peter Claes

☐ Initial submission ☐ Revised version ☒ Final submission

Life Sciences Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form is intended for publication with all accepted life science papers and provides structure for consistency and transparency in reporting. Every life science submission will use this form; some list items might not apply to an individual manuscript, but all fields must be completed for clarity.

For further information on the points included in this form, see Reporting Life Sciences Research. For further information on Nature Research policies, including our data availability policy, see Authors & Referees and the Editorial Policy Checklist.

▶ Experimental design

1. **Sample size**

   Describe how sample size was determined.

   > Sample size was determined by the amount of data available in the public facial data repository (Facebase.org). This sample was previously used, and this work shows a strong increase in power from this data set. More information is found in M&M.

2. **Data exclusions**

   Describe any data exclusions.

   > Data exclusions were pre-established based on image quality, facial interventions (e.g. facial surgery), and missing data. More information is found in M&M.

3. **Replication**

   Describe whether the experimental findings were reliably reproduced.

   > A completely independently sampled, imaged, genotyped, imputed... cohort was used to replicate any statistical findings. More information is found in M&M.

4. **Randomization**

   Describe how samples/organisms/participants were allocated into experimental groups.

   > No randomization took place, group membership of discovery and replication was determined by the separately obtained cohorts available. Important covariates were conditioned for during the analysis. More information is found in M&M.

5. **Blinding**

   Describe whether the investigators were blinded to group allocation during data collection and/or analysis.

   > Blinding was not relevant to this study. Two independent cohorts were used, one for discovery and one for replication. These cohorts were sampled by different research centra, and the analysis was done, by yet another research center. Each cohort was analyzed as a whole and no further grouping was applied.

   Note: all studies involving animals and/or human research participants must disclose whether blinding and randomization were used.

6. Statistical parameters

For all figures and tables that use statistical methods, confirm that the following items are present in relevant figure legends (or in the Methods section if additional space is needed).

n/a	Confirmed	
☐	☒	The <u>exact sample size</u> (*n*) for each experimental group/condition, given as a discrete number and unit of measurement (animals, litters, cultures, etc.)
☐	☒	A description of how samples were collected, noting whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	A statement indicating how many times each experiment was replicated
☐	☒	The statistical test(s) used and whether they are one- or two-sided (note: only common tests should be described solely by name; more complex techniques should be described in the Methods section)
☐	☒	A description of any assumptions or corrections, such as an adjustment for multiple comparisons
☐	☒	The test results (e.g. *P* values) given as exact values whenever possible and with confidence intervals noted
☐	☒	A clear description of statistics including <u>central tendency</u> (e.g. median, mean) and <u>variation</u> (e.g. standard deviation, interquartile range)
☐	☒	Clearly defined error bars

See the web collection on statistics for biologists for further resources and guidance.

▶ Software

Policy information about availability of computer code

7. Software

Describe the software used to analyze the data in this study.	Custom build facial mapping software was used and is made available on Github as stated in the manuscript. The rest of the work and analysis was performed using Matlab 2016b or R, all relevant functions used are quoted in the M&M, other more custom functions are available on request as stated in the Manuscript.

For manuscripts utilizing custom algorithms or software that are central to the paper but not yet described in the published literature, software must be made available to editors and reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). *Nature Methods* guidance for providing algorithms and software for publication provides further information on this topic.

▶ Materials and reagents

Policy information about availability of materials

8. Materials availability

Indicate whether there are restrictions on availability of unique materials or if these materials are only available for distribution by a for-profit company.	No unique materials were used

9. Antibodies

Describe the antibodies used and how they were validated for use in the system under study (i.e. assay and species).	No antibodies were used.

10. Eukaryotic cell lines

a. State the source of each eukaryotic cell line used.	WiCell (H9 ESC) and Fred Gage laboratory (iPSC)
b. Describe the method of cell line authentication used.	We have analyzed genomic sequences data from the lines
c. Report whether the cell lines were tested for mycoplasma contamination.	Yes, by PCR tests. M&M, Epigenomic analyses of CNCC regulatory chromatin signatures
d. If any of the cell lines used are listed in the database of commonly misidentified cell lines maintained by ICLAC, provide a scientific rationale for their use.	No commonly misidentified cell lines were used

▸ Animals and human research participants

Policy information about studies involving animals; when reporting animal research, follow the ARRIVE guidelines

11. Description of research animals

Provide details on animals and/or animal-derived materials used in the study.

No animals were used

Policy information about studies involving human research participants

12. Description of human research participants

Describe the covariate-relevant population characteristics of the human research participants.

For the discovery cohort (the Pittsburgh (PITT) sample), data from 2,449 participants were obtained from the 3D Facial Norms repository. The repository includes 3D facial surface images and self-reported demographic descriptors (e.g., age, ancestry) as well as basic physical characteristics (e.g. height and weight) from individuals recruited at four U.S. sites: Pittsburgh, PA; Seattle, WA; Houston, TX; and Iowa City, IA. Recruitment was limited to individual's age three to 40 years and of self-reported European ancestry. Individuals were excluded if they reported a personal or family history of any birth defect or syndrome affecting the head or face, a personal history of any significant facial trauma or facial surgery, or any medical condition that might alter the structure of the face.

For the replication cohort (the Penn State (PSU) sample), participants were recruited through several studies at The Pennsylvania State University and sampled in the following locations: State College, PA; New York, NY; Urbana-Champaign, IL; Twinsburg, OH; Dublin, Ireland; Rome, Italy; Warsaw, Poland, and Porto, Portugal. Participants self-reported information on age, ethnicity, ancestry, and body characteristics, and data were gathered on height, weight, and pigmentation of the hair and skin. Individuals were excluded from the analysis if they were below 18 years of age (range 18 – 88) and if they reported a personal history of significant facial trauma or facial surgery, or any medical condition that might alter the structure of the face. No restriction on ancestry or ethnicity was imposed during recruitment, but only individuals of European descent were used in this study (N = 2,059).

For both cohorts, there was no imbalance in gender, and weight and height distributions follow typical distributions seen in a European-derived population. Population stratification was dealt with using genetic PCs from full genome data.