

Supplementary Materials

The Genetic Architecture of Cutaneous Melanoma Across the Variant Spectrum

Shiyu Zhang, Deze Zhao, Ziyu Guo, Daishi Li, Zehao Luo, Yuming Sun, Yu Meng, Danyao Chen, Yating Dian, Zhiqin Zhang, Yao Yu, Zihua Wu, Jiayuan Le, Yangqian Chen, Can Gong, Yaosheng Fu, Linfeng Li, Yixiao Xiong, Lei Yao, Qian Zhou, Hui Su, Yi He, Guanxiong Zhang, Shuang Zhao, Furong Zeng, Zheng Wang, Jinchen Li, Chao Chen, Juan Su, Jian Yang, Guangtong Deng*, Xiang Chen*

*Corresponding Authors:

Guangtong Deng, dengguangtong@outlook.com; Xiang Chen, chenxiangck@126.com.

Table of Contents

- Supplementary Methods (pp. 2–19)
- Supplementary Notes (pp. 20–24)
- Supplementary References (pp. 25-26)
- Supplementary Figures (pp. 27–49; Supplementary Fig. 1–23)
- Supplementary Tables (Supplementary Tables 1–39; provided as separate Excel files)
- Supplementary Datasets (available via the Zenodo repository: 10.5281/zenodo.19562796)

Supplementary Methods

GWAS cohorts and summary statistics included in the meta-analysis

(1) UK Biobank

The UK Biobank study was approved by the North West Centre for Research Ethics Committee (11/NW/0382)[1]. Genotyping was performed using two closely related arrays (UK BiLEVE Axiom and UK Biobank Axiom), followed by imputation with the Haplotype Reference Consortium panel.

GWAS summary statistics for cutaneous melanoma were obtained from the GWAS Catalog (GCST90435596), with the primary analysis conducted by Wei Zhou et al.[2]. Summary statistics were harmonized and lifted over to the GRCh38 reference genome. The phenotype was defined using PheCode 172.11, with cases comprising individuals assigned to the corresponding phecode and controls defined as individuals without the phecode of interest or related phecodes.

The original analysis included White British participants who passed UK Biobank quality control, comprising 2,691 cases and 395,071 controls. Approximately 28 million variants with minor allele count ≥ 20 and imputation INFO score ≥ 0.3 were included.

Association testing was performed using SAIGE (v0.13) under a logistic mixed model, adjusting for sex, birth year, and the first four genetic principal components. A genetic relationship matrix was included as a random effect to account for sample relatedness, and saddlepoint approximation was applied to control for case–control imbalance.

(2) FinnGen Study

FinnGen is a large-scale research project integrating genomic and health registry data from Finnish biobank participants to investigate the genetic basis of disease[3]. Genotyping was performed using Illumina and Affymetrix arrays, followed by imputation with the population-specific SISu v4.2 reference panel using Beagle.

Sample-level quality control excluded individuals with ambiguous sex, genotype missingness $>5\%$, excess heterozygosity (± 4 SD), or non-Finnish ancestry. Variant-level quality control excluded variants with missingness $>2\%$, Hardy–Weinberg equilibrium $P < 1 \times 10^{-6}$, or minor allele count < 3 .

GWAS summary statistics were obtained from the FinnGen Data Freeze 12 release (November 4, 2024), comprising 500,348 participants of genetically inferred Finnish ancestry and 21,311,644 variants (GRCh38). Cutaneous melanoma was defined using registry-based endpoints, including ICD-10 (C43), ICD-9/8 (172), and cancer registry codes (ICD-O-3 topography C44; morphology codes 8720, 8721, 8726, 8730, 8742–8745;

behaviour code 3). The analysis included 5,753 cases and 378,749 controls, with controls defined as individuals without any cancer diagnosis.

Association testing was performed using REGENIE (v2.2.4). In step 1, leave-one-chromosome-out (LOCO) predictions were computed adjusting for age, sex, genotyping batch, and the first 10 genetic principal components. In step 2, association tests were performed for variants with minor allele count ≥ 5 . Approximate Firth correction was applied for variants with $P < 0.01$ (--firth --approx --pThresh 0.01), and standard errors were derived from likelihood ratio test statistics.

(3) Million Veteran Program (MVP) study

The VA Million Veteran Program (MVP) is a national cohort initiated in 2011 to investigate the contributions of genetic, environmental, and lifestyle factors to health and disease among U.S. veterans[4]. Genotyping was performed using a custom ThermoFisher Axiom MVP 1.0 array, followed by imputation to a hybrid reference panel combining the African Genome Resources panel and the 1000 Genomes Project Phase 3 (version 5).

Variant-level quality control excluded variants with imputation quality < 0.3 , minor allele count < 20 , call rate $< 97.5\%$ for common variants ($MAF > 1\%$), and call rate $< 99\%$ for rare variants ($MAF < 1\%$). Variants deviating by more than 10% from expected allele frequencies based on 1000 Genomes reference data were also removed.

Participants were assigned to genetically inferred ancestry groups (AFR, AMR, EAS, EUR, and SAS). Within each ancestry group, GWAS was performed using SAIGE under a generalized linear mixed model to account for relatedness and case–control imbalance. Only variants with imputation quality > 0.3 and minor allele count > 40 within each ancestry group were included. Analyses were adjusted for age, sex, and the first 10 ancestry-specific genetic principal components. Results were combined across ancestry groups using fixed-effect, inverse-variance weighted (IVW) meta-analysis implemented in GWAMA. Genomic control was implemented for each ancestry group but not applied to the resulting meta-analysis summary statistics.

Cutaneous melanoma was defined using PheCode 172.1 (“melanomas of skin”). GWAS summary statistics for this phenotype were obtained from the GWAS Catalog (GCST90479794), representing meta-analyzed results across ancestry groups. The dataset comprised 45,274 European ancestry cases and 389,597 controls, 256 African ancestry cases and 121,283 controls, and 1,032 Hispanic/Latino cases and 58,328 controls.

Definition of significant loci

Independently associated loci were defined using a two-stage framework integrating linkage disequilibrium (LD)–based clumping and conditional analysis. First, LD clumping was performed using PLINK2 ($P < 5 \times 10^{-8}$, $r^2 < 0.1$, ± 1 Mb window) to identify index variants. Regions spanning ± 1 Mb around each index variant were subsequently merged to define non-overlapping primary loci, yielding 117 loci.

To identify secondary independent association signals beyond LD-based clumping, we applied GCTA-COJO (Genome-wide Complex Trait Analysis—conditional and joint association analysis), which performs stepwise model selection to detect conditionally independent variants[5]. Whole-genome sequencing (WGS) data from 5,000 randomly selected unrelated UK Biobank participants were used as the LD reference panel, and variants with allele frequency discrepancies greater than 0.3 between the GWAS meta-analysis and the LD reference were excluded. This procedure identified 247 independently significant variants.

These variants were mapped to the primary loci, with 27 located outside any defined locus. Variants within 250 kb of the nearest locus boundary were assigned to that locus, whereas those more than 250 kb away were used to define secondary loci by constructing ± 500 kb intervals. These intervals were subsequently merged with existing loci to maintain non-overlapping boundaries, yielding a final set of 126 independent loci.

Fine-mapping

Fine-mapping was performed for each locus using CARMA (Causal Robust Mapping for Meta-analysis)[6], a Bayesian framework specifically designed for meta-analytic GWAS settings. In GWAS meta-analysis settings, heterogeneity arising from uneven sample size coverage across variants and discrepancies in LD structure between GWAS summary statistics and reference panels can bias posterior inference. CARMA addresses these challenges by explicitly modeling inconsistencies in LD structure between summary statistics and the reference panel, thereby improving robustness to LD misspecification. A heavy-tailed Cauchy prior on effect sizes enables adaptive modeling of heterogeneous signal-to-noise ratios across loci arising from uneven variant-level sample size in meta-analytic settings[6].

Whole-genome sequencing data from 490,541 UK Biobank participants were used as the LD reference panel. For each locus, symmetric LD matrices were computed using PLINK v1.9, with allele coding aligned to the effect allele orientation in the GWAS summary statistics. Summary statistics were reordered to match the LD matrix variant order, and Z-scores were computed as $Z = \beta/SE$.

Fine-mapping was performed using the CARMA function with outlier detection enabled (outlier.switch = TRUE) to account for inconsistencies between GWAS summary statistics and the LD reference, as well as heterogeneity in effective sample size across variants in meta-analysis. The maximum number of causal variants per locus was set to 10, with all other parameters left at default. Credible sets were defined at a posterior inclusion probability (PIP) threshold of 0.99.

For three loci exhibiting extensive allelic heterogeneity (loci 94, 103 and 121; >20 conditionally independent variants identified by GCTA-COJO), the maximum number of causal variants was increased in proportion to the number of independent signals at each locus to accommodate their increased complexity. The 46th locus (chr6:29,456,029 – 33,713,500), overlapping the major histocompatibility complex (MHC) region (chr6:28,510,120–33,480,577), was excluded due to its complex and extended LD structure.

Credible sets were restricted to those containing at least one variant with a genome-wide significant association ($P \leq 5 \times 10^{-8}$) in the meta-analysis. Candidate causal variants were then functionally annotated using the Ensembl Variant Effect Predictor (VEP, v115) against the GRCh38.p14 reference genome (annotations provided in the [Supplementary datasets](#)).

Epistatic effects between fine-mapped *MC1R* missense variants

Fine-mapping of the 103rd locus identified two independent missense variants in *MC1R* with PIP = 1.0 (chr16:89919709:C:T and chr16:89919510:C:A). Functional annotations curated in the Open Targets platform indicated that chr16:89919709:C:T is a fine-mapped trans-pQTL for circulating KIT protein levels (PIP = 1.0), whereas chr16:89919510:C:A was not included in the pQTL analysis due to its low allele frequency (MAF < 1%). Although a direct interaction between *MC1R* and *KIT* has not been established, both genes have well-characterized roles in melanocyte biology and melanoma pathogenesis[7]. The presence of two independent causal signals at *MC1R*, together with evidence linking one variant to circulating KIT levels, prompted us to examine the role of the second variant and the joint effects of both variants on circulating KIT levels.

To this end, we evaluated the association between genotypes at these two *MC1R* missense variants and plasma KIT concentrations using whole-genome sequencing data from 490,541 UK Biobank participants and Olink proteomic measurements from a subset of 53,013 individuals.

KIT normalized protein expression (NPX) values were transformed using a rank-based inverse normal transformation to approximate normality, calculated as $Z = \Phi^{-1}((r -$

0.5)/ n), where r is the rank of the NPX value, n is the sample size, and Φ^{-1} denotes the inverse standard normal cumulative distribution function.

We fitted linear regression models with transformed KIT levels as the outcome and variant dosage, modeled additively (0–2 copies of the effect allele), as predictors. Four models were evaluated: (i) each variant separately (two single-variant models), (ii) both variants jointly under an additive framework, and (iii) a joint model including an interaction term to assess deviation from additivity. All models were adjusted for age, sex, age², sex-by-age and sex-by-age² interaction terms, fasting time, and the first 20 genetic principal components. Analyses were restricted to individuals with complete data for KIT protein levels, both variant dosages, and covariates (N = 51,637).

For an intuitive assessment of joint allelic effects, participants were additionally classified by carrier status (dosage > 0) into four mutually exclusive groups: non-carriers, carriers of chr16:89919709:C:T only, carriers of chr16:89919510:C:A only, and carriers of both variants. Group means and 95% confidence intervals of transformed KIT levels were estimated. Differences across groups were assessed using one-way ANOVA. The expected KIT level under an additive model for double carriers was defined as $E_{both} = \mu_{v1-only} + \mu_{v2-only} - \mu_{none}$, and deviation from additivity was quantified as $\Delta = \mu_{none} - E_{both}$.

Gene prioritization

We performed functional annotation and gene prioritization for genome-wide significant loci using FUMA (v1.5.2)[8]. Meta-analytic GWAS summary statistics were converted from GRCh38 to GRCh37 genomic coordinates prior to analysis. Variants located at conversion-unstable positions were pre-excluded before liftover to ensure coordinate accuracy and to prevent build-conversion artefacts[9]. Candidate target genes were prioritized using three complementary strategies: (i) positional mapping, (ii) expression quantitative trait locus (eQTL) mapping and (iii) chromatin interaction mapping.

(1) Positional mapping

Variants were mapped to genes based on functional annotation and genomic proximity. SNPs annotated as exonic, splicing, 3' untranslated region (UTR) or 5' UTR were assigned to their corresponding genes as candidate targets.

(2) eQTL mapping

To link associated variants to putative target genes via regulatory effects on gene expression, we performed eQTL mapping using the FUMA framework across skin- and blood-relevant tissues, as well as immune cell-specific datasets.

- Bulk tissue eQTL resources

We included large-scale *cis*-eQTL datasets from multiple skin-/blood-relevant tissues: GTEx v8 (whole blood, cultured fibroblasts, skin not sun-exposed suprapubic, skin sun-exposed lower leg), eQTLGen Consortium (blood *cis*-eQTLs), BIOS Consortium (whole-blood eQTLs), MuTHER (skin) and the TwinsUK skin dataset from the eQTL Catalogue.

- Immune cell-specific eQTL resources

To extend eQTL mapping from bulk tissues to cell-type-resolved contexts, we incorporated datasets spanning single-cell and purified immune populations. Single-cell eQTL (scRNA-eQTL) datasets included B cells, dendritic cells, natural killer (NK) cells, monocytes (classical and non-classical), CD4⁺ T cells, CD8⁺ T cells, and peripheral blood mononuclear cells (PBMCs). In parallel, we included data from the DICE (Database of Immune Cell Expression, eQTLs and Epigenomics), comprising purified immune subsets such as naive and activated CD4⁺ and CD8⁺ T cells, B cells, NK cells, monocytes, T helper subsets (TH1, TH2, TH17, TFH), and regulatory T cells (naive and memory).

(3) Chromatin interaction mapping

To link variants to candidate target genes via long-range regulatory interactions, we performed chromatin interaction mapping using promoter-enhancer correlation data from the FANTOM5 consortium, incorporating both cell type-specific and organ-level datasets. Significant enhancer-promoter links were defined at a false discovery rate (FDR) $< 1 \times 10^{-6}$. Gene promoters were defined as regions spanning 250 bp upstream and 500 bp downstream of transcription start sites (TSS). Chromatin regulatory context was further supported using epigenomic annotations from the Roadmap Epigenomics Project. Variants located within enhancer regions showing significant correlation with gene promoters were assigned to those genes, enabling prioritization of candidate targets supported by long-range regulatory evidence.

For all prioritized candidate genes identified using the strategies above, gene set and pathway enrichment were evaluated using the GENE2FUNC module in FUMA.

Drug target enrichment of GWAS prioritized genes

To evaluate whether GWAS-prioritized genes are enriched among known drug targets, we integrated gene-drug annotations from three curated resources: Drug Repurposing Hub[10], DrugBank[11], and Open Targets Platform[12]. Gene-drug links were standardized by harmonizing gene symbols (HGNC) and drug names, and duplicate

entries were removed. The union of all unique target genes across the three databases was defined as the drug target gene set.

GWAS-prioritized genes were obtained from FUMA and restricted to protein-coding genes (N = 526). Ensembl gene identifiers were mapped to HGNC symbols using Ensembl BioMart, and only genes with valid annotations were retained. The background gene universe was defined as all human protein-coding genes with valid HGNC symbols retrieved from BioMart. Both GWAS-prioritized genes and drug target genes were restricted to this shared background to ensure consistency in enrichment testing.

Enrichment was assessed using a two-sided Fisher's exact test based on a 2×2 contingency table comparing the overlap between GWAS-prioritized genes and drug target genes relative to the background gene universe. Odds ratios and 95% confidence intervals were reported to quantify enrichment effect size, and corresponding P values were used to assess statistical significance.

Cell type deconvolution of cutaneous melanoma spatial transcriptomic data

Cellular composition across spatial transcriptomic (ST) spots was inferred using SpaCET (Spatial Cellular Estimate for Tumors, v1.4.0), a computational framework developed for tumor spatial transcriptomic deconvolution[13]. The method is designed to address key challenges inherent to tumor tissues, including intratumoral heterogeneity, variable tissue density across spots, immune cell integrity and collinearity among closely related immune sublineages.

Malignant melanocyte fractions were first inferred using a melanoma-specific copy number alteration (CNA) gene-pattern dictionary derived from skin cutaneous melanoma. Following malignant fraction estimation, non-malignant cellular components were resolved using a hierarchical cell reference atlas derived from publicly available single-cell RNA sequencing datasets. SpaCET applies constrained linear regression to estimate cell-type proportions within a two-level hierarchical framework:

(1) At the first level, non-malignant fractions were decomposed into 12 major lineages, comprising 10 immune populations (plasma cells, B cells, CD4⁺ T cells, CD8⁺ T cells, natural killer cells, conventional dendritic cells, plasmacytoid dendritic cells, macrophages, mast cells and neutrophils), two stromal populations (cancer-associated fibroblasts and endothelial cells), and an unassigned component capturing residual signal not explained by the reference lineages.

(2) At the second level, major lineages were further deconvolved into functionally defined sublineages. B cells were partitioned into naïve, non-switched memory, switched

memory and exhausted subsets. CD4⁺ T cells were resolved into naïve, T helper 1 (Th1), T helper 2 (Th2), T helper 17 (Th17), T follicular helper (Tfh) and regulatory T (Treg) populations. CD8⁺ T cells were classified into naïve, central memory, effector memory, effector and exhausted subsets. Conventional dendritic cells were subdivided into cDC1 (CLEC9A⁺), cDC2 (CD1C⁺) and cDC3 (LAMP3⁺) populations. Macrophages were resolved into M1-like and M2-like states.

Based on spot-level deconvolved cell-type proportions, tumor–stroma interface regions were defined as follows. Spots with a malignant cell fraction ≥ 0.5 were classified as “Tumor”, whereas all remaining spots were initially labeled as “Stroma”. For each spot initially labeled as “Stroma”, its six adjacent neighbors on the Visium hexagonal grid were examined; a spot was retained as “Stroma” only if all neighboring spots were also labeled as “Stroma”, otherwise it was reclassified as “Interface”. These “Tumor”, “Stroma”, and “Interface” annotations were used for downstream spatial analyses.

All analyses were performed using SpaCET (v1.4.0) with default parameters unless otherwise specified.

Genetically informed spatial mapping of cells (gsMap)

To identify cellular and spatial contexts enriched for the heritability of cutaneous melanoma, we applied gsMap(v1.73.6), which integrates spatial transcriptomic profiles with genome-wide association study (GWAS) summary statistics to map trait-associated genetic signals to relevant cell types and their spatial organization[14].

We followed the gsMap analytical framework as described by the authors. First, latent embeddings were learned for each spatial spot by jointly modeling the standardized gene expression matrix and the spatial adjacency matrix, thereby capturing both transcriptional similarity and spatial structure. Second, spots with similar latent representations were grouped into homogeneous microdomains, and gene specificity scores (GSS) were computed for each spot by aggregating gene expression patterns across its corresponding microdomain. Third, gene-level GSS values were linked to genetic variants by assigning SNPs located within ± 50 kb of each gene’s transcribed region. These SNP sets were used to construct spot-specific annotations for stratified linkage disequilibrium score regression (S-LDSC). Linkage disequilibrium (LD) scores were calculated using genotype data from 10,000 randomly selected unrelated UK Biobank participants with whole-genome sequencing data as the LD reference panel. Fourth, S-LDSC was used to test whether SNPs with higher GSS values were enriched for the heritability of cutaneous melanoma, conditional on the baseline SNP annotation model. Statistical significance was evaluated using a one-sided z-test assessing whether the regression coefficient associated with each spot-specific annotation was greater than zero. To evaluate the

association between defined spatial regions (for example, Tumor, Interface, and Stromal regions) and trait-relevant genetic risk, spot-level P values within each region were aggregated using the Cauchy combination test.

gsMap signal heterogeneity across tumor spots

To quantify heterogeneity across tumor spots, we analyzed spot-level gsMap signals in 14 cutaneous melanoma spatial transcriptomic samples, defining the signal for each spot as the $-\log_{10}(P)$ value from spatial S-LDSC analysis.

To account for variation in library size/sequencing depth across spots, we performed within-sample adjustment using a linear model as $\text{gsMap}_{-\log_{10} P} \sim \log(1 + \text{UMI})$. Model residuals were used as UMI-adjusted gsMap signals, representing the component of the gsMap signal independent of sequencing depth within each sample.

We assessed heterogeneity of gsMap signals across two types of tumor spots: (i) all Tumor-classified spots (malignant cell fraction ≥ 0.5 , as defined by SpaCET), and (ii) high-purity tumor spots (malignant cell fraction ≥ 0.99). For each sample, heterogeneity within tumor spots was quantified using dispersion metrics—including the mean, median, standard deviation, variance, interquartile range (IQR), and median absolute deviation (MAD)—calculated for both raw gsMap signals and UMI-adjusted residuals.

Spatial autocorrelation of UMI-adjusted gsMap signals within tumor spots was assessed using Moran's I . Spatial coordinates were defined by Visium grid positions (`array_row`, `array_col`). For both tumor spot definitions—(i) all Tumor-classified spots (malignant cell fraction ≥ 0.5) and (ii) high-purity tumor spots (malignant cell fraction ≥ 0.99)—a k -nearest-neighbor graph ($k = 6$) was constructed within each sample, and row-standardized spatial weights (w_{ij}) were derived from the resulting adjacency structure. Moran's I was computed as:

$$I = \frac{n}{\sum_i \sum_j w_{ij}} \cdot \frac{\sum_i \sum_j w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{\sum_i (x_i - \bar{x})^2}$$

where x_i denotes the UMI-adjusted gsMap residual at spot i , \bar{x} is the mean across tested spots and n is the number of tumor spots. Positive values of I indicate spatial clustering of similar residual gsMap signals, whereas values near zero are consistent with spatial randomness. Statistical significance was assessed using a permutation-based Moran's I test (999 permutations; one-sided alternative for positive autocorrelation), and P values were adjusted across samples and tumor definitions using the Benjamini–Hochberg procedure.

Functional state heterogeneity of tumor spots

Given the consistently observed heterogeneity of gsMap signals across tumor spots, we considered two potential contributors: variation in tumor cell functional states and differences in their spatial context. To evaluate the contribution of functional heterogeneity, we quantified activity scores for 50 hallmark pathways for each tumor spot. Within each sample, Spearman rank correlations were computed between gsMap signals ($-\log_{10}P$) and pathway activity scores across tumor spots.

P values from all pairwise correlations between pathway scores and gsMap signals across samples were adjusted using the Benjamini–Hochberg procedure, and associations with a false discovery rate (FDR) < 0.05 were considered significant. Significant positive associations ($\rho > 0$) were summarized across samples to assess reproducibility.

Spatial-context heterogeneity of tumor spots

To quantify spatial-context heterogeneity among tumor spots, we evaluated the association between gsMap signals ($-\log_{10}P$) and proximity to the tumor–immune interface. For each tumor spot, the proximity was defined as the minimum Euclidean distance (in Visium grid units) to the nearest tumor–immune interface.

Within each sample, associations between gsMap signals and distance to the interface were assessed across tumor spots using Spearman rank correlation. To account for variation in tumor purity, we fitted a linear model including malignant cell fraction as a covariate: $\text{gsMap}_{-\log_{10} p} = \beta_0 + \beta_1 \cdot \text{Dist}_{\text{interface}} + \beta_2 \cdot \text{Malignant fraction} + \varepsilon$. P values were adjusted across samples using the Benjamini–Hochberg procedure.

Colocalization with tumor cells and enrichment of T cells in boundary-adjacent regions

Tumor–immune boundary-adjacent regions were defined as the union of tumor–immune interface spots (classified by SpaCET) and their six immediate neighbors on the Visium hexagonal lattice. Immune cell abundance in these regions was quantified using inferred spot-level cell-type proportions. For each sample, we computed Spearman rank correlations between malignant cell proportions and the proportions of each major immune lineage (Plasma, B cell, CD4⁺ T, CD8⁺ T, NK, cDC, pDC, macrophage, mast and neutrophil) across boundary-adjacent spots.

P values were adjusted across all malignant–immune lineage pairs and samples using the Benjamini–Hochberg procedure, and associations with $FDR < 0.05$ and positive correlation coefficients were considered evidence of colocalization.

Among all major immune lineages, T cells exhibited reproducible positive colocalization with malignant cells across samples and were therefore prioritized for downstream analysis. For each spot, overall T cell proportion was defined as the sum of $CD4^+$ and $CD8^+$ T cell proportions.

To quantify T cell enrichment in boundary-adjacent regions, we compared overall T cell abundance in boundary-adjacent spots with that in non-boundary spots from the same sample. Enrichment was defined as the ratio of the median T cell proportion in boundary-adjacent spots to that in non-boundary spots, with a small constant ($\epsilon = 1 \times 10^{-4}$) added to both medians to stabilize estimates. Ninety-five percent confidence intervals were estimated using non-parametric bootstrap with 2,000 resamples. Two-sided empirical bootstrap P values were calculated against the null ratio of 1.

Spatial ligand–receptor co-expression analysis

Spatial ligand–receptor (L–R) co-expression patterns were quantified using bivariate Moran’s I, which measures spatial coupling between ligand expression at a given location and receptor expression in neighboring locations. For each sample, L–R pairs with positive spatial association (Moran’s I > 0) and false discovery rate (FDR) < 0.05 were retained. To identify conserved high-confidence interactions, we further restricted to L–R pairs that were consistently detected across all samples (N = 14) with positive and significant associations, and that exhibited strong spatial coupling (mean Moran’s I > 0.2).

UK Biobank whole-genome sequencing data and quality control

This study used the UK Biobank 500K WGS DRAGEN aggregated dataset (release 2, May 2024). Whole genomes from 490,640 UK Biobank participants were sequenced to a mean coverage of 32.5× using Illumina NovaSeq 6000 platforms. Details of the sequencing and processing pipeline have been described previously[15]. Multi-allelic variants were decomposed into bi-allelic sites, followed by cohort-level filtering using the DRAGEN machine learning–based quality control framework. For each variant allele, a site-level quality score (MLSQ) was computed and recorded as the QUAL value; alleles with $QUAL \geq 0.1$ were assigned FILTER = PASS.

In this release, 1,081,661,407 PASS single-nucleotide polymorphisms (SNPs) and 129,273,976 PASS insertions and deletions (indels) were identified across autosomes,

sex chromosomes, mitochondrial DNA, and alternative contigs. Downstream analyses were restricted to high-quality (FILTER = PASS) autosomal variants. Additional site-level filtering excluded variants with genotype quality (GQ) ≤ 10 , local allele depth (LAD) < 8 , or missing genotype rate > 0.1 .

Variant annotation

Variants were functionally annotated using the Ensembl Variant Effect Predictor (VEP, v114) based on the GRCh38 reference genome (GRCh38.p14). Transcript annotations were derived from GENCODE v48. Predicted loss-of-function variants were annotated using the LOFTEE plugin to identify high-confidence protein-truncating variants. Additional functional impact scores, including REVEL and CADD Phred-scaled scores, were incorporated from dbNSFP (v5.1a). Population allele frequencies were obtained from gnomAD (v4.1). These annotations were used to prioritize rare variants with large predicted functional impact for downstream collapsing analyses.

Construction of rare-variant masks for protein-coding genes

VEP-derived functional annotations were integrated with splice-disruption predictions from Illumina precomputed SpliceAI scores for both SNVs and indels to inform gene-based rare-variant collapsing analyses. Analyses were restricted to variants within protein-coding genes, retaining MANE Select transcripts when available and otherwise the canonical transcript. For each variant, the maximum SpliceAI delta score across acceptor and donor gain/loss categories (DS_max) was computed and incorporated into downstream annotation.

Specifically, variants were classified into functional groups based on predicted coding impact. Predicted loss-of-function (pLoF) variants were defined as LOFTEE high-confidence annotations, including transcript ablation, splice acceptor or donor variants, stop-gained, and frameshift variants. Damaging missense or protein-altering variants were defined as missense or inframe insertion/deletion variants with either REVEL ≥ 0.773 , CADD Phred-scaled score ≥ 28.1 , or SpliceAI DS_max ≥ 0.20 . Missense or inframe variants not meeting these thresholds were classified as other missense or protein-altering variants. Synonymous variants were defined based on VEP consequence terms, including synonymous and stop-retained variants.

Gene-based rare-variant masks were constructed by combining prioritized coding variant classes (for example, pLoF alone or pLoF plus damaging missense variants) across multiple minor allele frequency (MAF) thresholds ([Supplementary Note 3](#)).

Construction of rare-variant masks for noncoding RNAs

Rare-variant masks for noncoding RNAs (ncRNAs) were constructed based on Ensembl annotations (Ensembl.Hsapiens.v86; GRCh38). Non-protein-coding gene biotypes were retained to define ncRNA transcripts, and exonic intervals were extracted accordingly.

To prioritize non-coding variants with potential functional relevance, three complementary metrics capturing predicted pathogenicity, constraint, and evolutionary conservation respectively, were applied: CADD (Combined Annotation Dependent Depletion; ≥ 20)[16], JARVIS (Junk Annotation genome-wide Residual Variation Intolerance Score; ≥ 0.99)[17], and GERP (Genomic Evolutionary Rate Profiling; ≥ 2)[18]. Variants meeting these criteria were grouped into masks (for example, CADD- or GERP-prioritized) for downstream rare-variant collapsing analyses ([Supplementary Note 3](#)).

Construction of rare-variant masks for epidermal melanocyte-specific candidate *cis*-regulatory elements

Candidate *cis*-regulatory element (cCRE) annotations were obtained from the ENCODE consortium, which integrates biochemical signatures across diverse human cell types to define regulatory elements with putative functional activity[19]. We used the ENCODE v4 release (~2.37 million cCREs across 1,888 cell types) and restricted analyses to epidermal melanocyte-specific cCREs to prioritize non-coding regulatory elements associated with genetic risk of cutaneous melanoma. In total, 2,272,843 melanocyte-specific cCREs were included.

These cCREs were classified into eight functional categories based on genomic proximity to transcription start sites (TSSs) and combinations of biochemical signatures (including DNase/ATAC accessibility, H3K4me3, H3K27ac, CTCF binding, and transcription factor occupancy). Specifically, these categories comprised promoter-like signatures (PLS), proximal enhancer-like signatures (pELS), distal enhancer-like signatures (dELS), chromatin accessibility with H3K4me3 (CA-H3K4me3), chromatin accessibility with CTCF (CA-CTCF), chromatin accessibility with transcription factor binding (CA-TF), chromatin accessibility only (CA), and transcription factor-associated elements (TF).

Rare-variant masks for cCREs were constructed using the same prioritization framework applied to ncRNA genes, with variants prioritized based on CADD, GERP, and JARVIS scores ([Supplementary Note 3](#)).

Phenotype definition and sample inclusion in the UKB WGS cohort

Participants from the UK Biobank with available whole-genome sequencing data were included in rare-variant association analyses. Cases were defined as individuals with a registry-confirmed diagnosis of cutaneous melanoma, identified using ICD-10 codes C43.0–C43.9 (malignant melanoma of skin) or D03.0–D03.9 (melanoma in situ), and without any record of other primary cancers. Controls were defined as participants without any cancer diagnosis in the UK Biobank cancer registry. Individuals with discordance between self-reported and genetically inferred sex, or evidence of sex chromosome aneuploidy, were excluded.

Among 3,848 cases with WGS data, 3,819 (99.2%) were of European ancestry ([Supplementary Fig. 20](#)). Given the limited representation of other ancestry groups and to minimize confounding from population stratification in rare-variant analyses, subsequent analyses were restricted to individuals of European ancestry.

Cutaneous melanoma diagnoses were further ascertained using linked national cancer registry data. Tumor histology (UK Biobank field 40011) and behavior codes (field 40012) were combined to derive morphology–behavior annotations according to the International Classification of Diseases for Oncology, 3rd Edition (ICD-O-3). In the study cohort, superficial spreading melanoma (ICD-O-3 code 8743/3) was the predominant histological subtype, accounting for 1,596 cases (41.5%) ([Supplementary Fig. 21](#)).

Genome-wide rare-variant collapsing analysis

Genome-wide rare-variant collapsing analyses for protein-coding genes and non-coding regions were performed using SAIGE-GENE+ (v1.4.0)[20]. To account for cryptic relatedness and population structure, a sparse genetic relatedness matrix (GRM) was constructed from linkage disequilibrium (LD)-pruned hard-called genotypes using PLINK2 (--indep-pairwise 50 5 0.05). This approach provides a computationally efficient approximation of the full GRM while preserving local kinship structure required for mixed-model association testing.

A generalized linear mixed model (GLMM) was fitted to model cutaneous melanoma case–control status as a binary outcome. Fixed-effect covariates included genetically inferred sex, age, age², sex-by-age and sex-by-age² interaction terms, WGS sequencing batch, and the first 20 genetic principal components, with the sparse GRM included as a random effect. The fitted null model was used for all downstream region-based rare-variant association analyses. To improve calibration under case–control imbalance, saddlepoint approximation (SPA) was applied to the score test statistics.

Genome-wide rare-variant collapsing analyses were conducted separately for 18,343 protein-coding genes, 34,829 ncRNAs, and 2,272,843 cCREs using their respective

prioritized variant masks and across multiple minor allele frequency thresholds (1×10^{-4} , 1×10^{-3} and 1×10^{-2}). Statistical significance was evaluated within each category based on two criteria: Bonferroni correction ($0.05/N$, where N denotes the number of tested genes or regions) and Benjamini–Hochberg false discovery rate (FDR) control at 0.05.

Conditional analysis for rare-variant associations

For rare-variant associations passing $FDR < 0.05$, collapsing analyses were repeated while conditioning on independently significant variants identified by GCTA-COJO from the GWAS meta-analysis, applied consistently across protein-coding genes and noncoding regions (cCREs and ncRNAs).

To further evaluate whether noncoding signals could be explained by LD with nearby coding variants, additional conditioning analyses were performed. For cCREs, conditional models additionally included all tested coding variants within candidate target protein-coding genes linked to each element based on ENCODE cCRE–gene annotations, as well as the nearest protein-coding gene when not already included. For ncRNAs, conditioning was performed on all tested coding variants within the nearest protein-coding gene.

Three independent replication cohorts for rare-variant associations

Genomics England (GEL) 100,000 Genomes Project

The 100,000 Genomes Project (100KGP) received ethical approval from the Health Research Authority Research Ethics Committee East of England—Cambridge South (Ref. 14/EE/111). Participants, including cancer patients, rare disease probands and their relatives, were recruited through 13 NHS Genomic Medicine Centres across the UK, with informed consent obtained from all individuals[21].

We analyzed the aggregated release of 78,195 whole-genome sequencing germline genomes (AggV2), provided as a multi-sample VCF. All samples were sequenced using 150 bp paired-end reads on Illumina HiSeq X instruments and processed uniformly using the Illumina North Star Version 4 whole-genome sequencing workflow (NSV4, v2.6.53.23). Reads were aligned to the GRCh38 reference genome with decoy sequences.

All samples passed stringent quality control filters, including low contamination (freemix < 0.03), a balanced heterozygous-to-homozygous SNV ratio (< 3), and an expected number of SNVs (3.2–4.7 million). Additional criteria included $>90\%$ concordance with genotyping array data, median fragment size >250 bp, $<5\%$ chimeric reads, $>60\%$ mapped reads, and $<10\%$ AT dropout.

Variants were normalized by decomposing multiallelic sites, enforcing parsimony, and left-aligning indels. Autosomal variants were labeled as “PASS” only if they satisfied all site-level filters, including low missingness ($\leq 5\%$), sufficient depth (median depth ≥ 10), high genotype quality (median GQ ≥ 15), balanced allelic representation in heterozygous calls (ABratio $\geq 25\%$), adequate call rate (completeGTRatio $\geq 50\%$), and no significant deviation from Hardy–Weinberg equilibrium in unrelated individuals of genetically inferred European ancestry (mid-P $\geq 1 \times 10^{-5}$).

Overall, 540,098,760 of 722,342,407 autosomal variants (74.8%) passed quality control. Downstream analyses were performed using the masked PLINK2 PGEN genotype files.

For the replication cohort, cases were defined as individuals recruited to the cancer programme with a diagnosis of malignant cutaneous melanoma, as well as individuals recruited through the rare disease programme who had a diagnosis of cutaneous melanoma (ICD-10: C43.x or D03.x) as their only recorded cancer. Controls were selected from the rare disease programme and had no cancer registry record. Analyses were restricted to unrelated individuals of genetically inferred European ancestry, excluding samples with discordant genetic and reported sex. The final replication cohort comprised 441 cases and 28,069 controls.

We sought to replicate rare-variant associations identified in the UK Biobank discovery cohort (FDR < 0.05), including 128 protein-coding genes, 3 ncRNAs, and 7 melanocyte-specific cCREs. Variant annotation pipeline and rare-variant mask definitions were applied consistently between discovery and replication analyses. Variants with minor allele frequency $> 1\%$ in the UK Biobank were excluded, and masks with a minor allele count < 5 in the replication cohort were not tested.

Rare-variant associations were performed using Firth logistic regression, adjusting for genetically inferred sex, age, age², sex-by-age and sex-by-age² interaction terms, and the first 20 genetic principal components.

All of Us and Mass General Brigham Biobank

Exome-wide rare-variant collapsing analyses across 601 disease phenotypes in the All of Us (AoU) and Mass General Brigham Biobank (MGB) cohorts have been reported previously[22]. Summary statistics for cutaneous melanoma (phecode 172.1) were obtained from their publicly available portal (https://hugeamp.org:8000/research.html?pageid=600_traits_app_home).

The AoU cohort included whole-genome sequencing data from 2,014 cases and 119,504 controls of genetically inferred European ancestry, whereas MGB cohort comprised whole-exome sequencing data from 2,110 cases and 41,143 controls of genetically inferred European ancestry. Detailed descriptions of sequencing, quality control, and phenotype definition are provided in the original study[22].

In their exome-wide rare-variant collapsing analyses, they defined nine rare-variant masks based on allele frequency and predicted functional impact, including: (1) LOF variants (MAF < 1%); (2) LOF + missense (score > 0.8; MAF < 1%); (3) LOF + missense (score > 0.5; MAF < 1%); (4) LOF variants (MAF < 0.1%); (5) LOF + missense (score > 0.8; MAF < 0.1%); (6) LOF + missense (score > 0.5; MAF < 0.1%); (7) LOF + missense (score > 0.5; MAF < 0.001% or MAC < 5); (8) missense-only variants (score > 0.5; MAF < 0.001% or MAC < 5); and (9) missense-only variants (score > 0.2; MAF < 0.001% or MAC < 5). Association testing was performed using two-sided logistic mixed-effects score tests, adjusting for age, age², sex, sequencing batch, and cohort-specific genetic principal components. Relatedness was modelled using a sparse kinship matrix as a random effect, and P values were computed using saddle-point approximation to account for case–control imbalance.

Definition of replication criteria

Replication of rare-variant associations was evaluated across independent cohorts. For protein-coding genes with FDR < 0.05 in the discovery analysis (N = 128), replication was assessed using results from the AoU, MGB, and GEL cohorts. For noncoding associations (3 ncRNAs and 7 cCREs), replication was assessed in the GEL cohort only.

Replication was defined by (i) concordant direction of burden effect estimates with the UK Biobank discovery dataset and (ii) nominal significance (two-sided $P < 0.05$) in at least one replication dataset. For protein-coding genes, robust replication required evidence from at least two independent cohorts. For noncoding associations, replication was defined as concordant effect direction with $P < 0.05$ in the GEL cohort.

Common-variant polygenic risk scores (cvPRS)

GWAS meta-analysis summary statistics were lifted over from GRCh38 to GRCh37. Variants located in conversion-unstable regions were excluded prior to liftover following recommended guidelines (<https://github.com/cathaloruaidh/genomeBuildConversion>). Polygenic risk scores were constructed using PRS-CS[23]. The UK Biobank European-ancestry LD reference panel was used, restricting to non-ambiguous HapMap3 variants with imputation INFO > 0.8 and minor allele frequency > 1%. The global shrinkage parameter (ϕ) was inferred from the data under a fully Bayesian framework (PRS-CS-auto). Markov chain Monte Carlo sampling was performed with 10,000 iterations, including 5,000 burn-in iterations.

The GEL cohort was used as an out-of-sample evaluation dataset, and individual-level cvPRS were calculated using posterior effect size estimates from PRS-CS.

Construction of an integrative network linking genes prioritized by GWAS and rare-variant analyses

To contextualize the interaction landscape underlying cutaneous melanoma associated genes, we constructed an integrated network comprising 654 entries (647 unique genes) prioritized by GWAS meta-analysis (N = 526) and rare-variant association analyses (N = 128). Molecular interaction data were curated from the OmniPath database using the OmnipathR package (v3.15.1)[24].

Two main classes of interactions were considered: (i) post-translational interactions, including protein–protein interactions (PPIs), and (ii) transcriptional regulatory (TR) interactions. PPI data were obtained from the omnipath, kinaseextra, pathwayextra and ligrecextra resources, whereas TR interactions were curated from DoRothEA, CollecTRI and related transcription factor–target databases. Only human interactions (NCBI taxonomy ID 9606) were retained.

Interactions involving protein complexes were decomposed into their constituent proteins, and all pairwise interactions between components were enumerated, with self-interaction loops removed. Interaction endpoints were mapped from UniProt identifiers to Ensembl gene IDs.

For each directed interaction (source–target pair), duplicate records were resolved by retaining the entry with the strongest supporting evidence, prioritizing the number of literature references, curation level, and supporting databases. To construct an undirected interaction network, reciprocal interactions (A→B and B→A) were collapsed into a single edge, with evidence metrics summarized conservatively by taking the maximum value across directions.

To quantify evidence support, interaction weights were derived from the number of literature references annotated in OmniPath. Reference counts were transformed as $\log(1+n_{references})$, capped at the 99th percentile to limit the influence of outliers, and rescaled to the [0,1] interval using min-max normalization. A small constant ($\varepsilon = 0.01$) was added to avoid zero weights, yielding final weights defined as $\omega = \varepsilon + (1 - \varepsilon) \times \frac{x - x_{min}}{x_{max} - x_{min}}$, where x denotes the capped log-transformed reference count. For gene pairs supported by both PPI and TR evidence, interaction weights were integrated using a probabilistic union model, $\omega = 1 - \prod_i(1 - \omega_i)$, where ω_i denotes the normalized weight of each supporting interaction. This formulation assumes independent sources of evidence and increases the combined weight as multiple lines of support accumulate.

The final network comprised 427 genes connected by 1,063 edges. Genes without any curated interaction with other prioritized genes were excluded, as they formed isolated nodes outside the connected network.

Supplementary Notes

Supplementary Note 1 | Summary of GWAS datasets included in the meta-analysis

	UK Biobank	FinnGen (R12)	MVP
Source of summary statistics	GWAS Catalog (GCST9043559 5)	Release DF12 (Nov 4, 2024) finngen_R12_C3_MELANOMA_SKIN_EXA LLC	GWAS Catalog (GCST9047979 4)
Phenotype definition	Melanoma of skin (PheCode 172.1) ICD-9: 172.x, V10.82; ICD-10: C43.x, D03.x	Malignant melanoma of skin: - ICD-8/9/10: 172, C43; - cancer registry: Topography ICD-O-3: C44; Morphology ICD-O-3: 8720 8721 8726 8730 8742 8743 8744 8745 Behavior codes: 3 (Controls: no cancer diagnosis)	Melanoma of skin (PheCode 172.1) ICD-9: 172.x, V10.82; ICD-10: C43.x, D03.x
Number of cases	2,691	5,753	AFR: 256 EUR: 45,274 AMR: 1,032 (Total: 46,562)
Number of controls	395,071	378,749	AFR: 121,283 EUR: 389,597 AMR: 58,328 (Total: 569,208)
Genome build	GRCh38	GRCh38	GRCh38
Genotype array	UK BiLEVE Axiom array and UKB Axiom array	Illumina & Affymetrix chip arrays	Thermo Fisher Axiom MVP 1.0
Imputation panel	HRC + UK10K	SISu v4.2 (8,554 WGS)	AGR + 1000G Phase 3
Genotype QC	MAC ≥ 20 ; INFO ≥ 0.3	Genotype missingness $\leq 2\%$; HWE $P \geq 1e-6$; MAC ≥ 3	INFO ≥ 0.3 ; MAC ≥ 20 ; call rate $\geq 97.5\%$ for common variants and

			≥99% for rare variants
Software used	SAIGE v0.13 Covariates: PCs1–4, sex, birth year	REGENIE v2.2.4 Covariates: sex, age, 10 PCs, genotyping batch	SAIGE (GPU) Covariates: age, sex, 10 ancestry-specific PCs
Ancestry	White British (EUR)	Finnish (EUR)	Meta-analysis of AFR, AMR, EUR

Supplementary Note 2 | Fine-mapped causal variants with candidate epistatic effects in *MC1R*

	chr16:89919709:C:T	chr16:89919510:C:A
Marginal p-value	1.55×10 ⁻³⁰²	8.95×10 ⁻²³
Conditional p-value	1.00×10 ⁻²⁶⁰	4.55×10 ⁻⁴⁷
Beta	0.44	0.32
Conditional beta	0.46	0.48
Effect allele	T	A
PIP	1.00	1.00
Functional consequence	Missense variant	Missense variant
Amino acid change	R151C	D84E
CADD Phred score	24.5	24.1
dbSNP ID	rs1805007	rs1805006
gnomAD AF – African/African-American	0.01314	0.00184
gnomAD AF – Latino/Admixed American	0.01260	0.00202
gnomAD AF – Amish	0.05044	0.00658
gnomAD AF – Ashkenazi Jewish	0.06848	0.00179
gnomAD AF – East Asian	0.00074	0
gnomAD AF – Finnish	0.06711	0.00205
gnomAD AF – Middle Eastern	0.00874	0.00050
gnomAD AF – Non-Finnish European	0.09177	0.01209
gnomAD AF – South Asian	0.00404	0.00012
Meta-analysis AF	0.066	0.0083

Supplementary Note 3 | Definition of rare-variant masks

Category	Mask
Protein-coding genes	<p>pLoF</p> <p>damaging_missense_or_protein_altering</p> <p>other_missense_or_protein_altering</p> <p>synonymous</p> <p>damaging_missense_or_protein_altering + other_missense_or_protein_altering (All missense)</p> <p>pLoF + damaging_missense_or_protein_altering</p> <p>pLoF + damaging_missense_or_protein_altering + other_missense_or_protein_altering</p> <p>pLoF + damaging_missense_or_protein_altering + other_missense_or_protein_altering + synonymous</p>
Noncoding RNAs	<p>CADD-prioritized (CADD_PHRED \geq 20)</p> <p>GERP-prioritized (GERP \geq 2)</p> <p>JARVIS-prioritized (JARVIS \geq 0.99)</p> <p>CADD–GERP prioritized (CADD_PHRED \geq 20 or GERP \geq 2)</p> <p>CADD–JARVIS prioritized (CADD_PHRED \geq 20 or JARVIS \geq 0.99)</p> <p>All-score prioritized (CADD_PHRED \geq 20 or GERP \geq 2 or JARVIS \geq 0.99)</p> <p>All ncRNA exon variants (All variants within ncRNA exons irrespective of score support)</p>
Melanocyte-specific cCREs	<p>CADD-prioritized (CADD_PHRED \geq 20)</p> <p>GERP-prioritized (GERP \geq 2)</p> <p>JARVIS-prioritized (JARVIS \geq 0.99)</p> <p>CADD–GERP prioritized (CADD_PHRED \geq 20 or GERP \geq 2)</p> <p>CADD–JARVIS prioritized (CADD_PHRED \geq 20 or JARVIS \geq 0.99)</p> <p>All-score prioritized (CADD_PHRED \geq 20 or GERP \geq 2 or JARVIS \geq 0.99)</p> <p>All cCRE variants (All variants within melanocyte-specific cCREs irrespective of score support)</p>

Supplementary Note 4 | Discussion on distinct genes prioritized by GWAS and rare-variant analyses and their network topology

To provide a unified view of genetic architecture across the allele-frequency spectrum, we constructed a biological network to contextualize interactions and regulatory relationships among 647 CM-associated genes prioritized by GWAS or rare-variant collapsing analyses. The two approaches prioritized largely distinct gene sets with limited overlap (seven genes identified by both, and 519 and 121 genes uniquely prioritized by GWAS and rare-variant analyses, respectively), consistent with previous observations[25]. This discrepancy arises because rare-variant burden tests primarily prioritize trait-specific genes, whereas GWAS preferentially identifies trait-specific variants[25]. Variants can be trait-specific in two ways: they may directly act within trait-specific genes (e.g., coding variants) or exert trait-specific regulatory effects on pleiotropic genes (e.g., noncoding variants), for example by regulating gene expression only in trait-relevant cell types, developmental stages or biological contexts. GWAS can therefore prioritize those pleiotropic and highly constrained genes through trait-specific regulatory variants in the noncoding genome, where regulatory effects can remain highly context-dependent despite the broad functional importance of the target gene. Moreover, these regulatory variants are generally under weaker negative selection than damaging coding variants disrupting the same genes. In contrast, rare-variant collapsing analyses are constrained by the depletion of deleterious coding variants in pleiotropic genes under purifying selection. This depletion results in low carrier counts and consequently larger standard errors, limiting statistical power even in the presence of substantial effect sizes and partially decoupling association strength ($Z^2 = (\beta / \text{s.e.})^2$) from the biological importance of the gene to the trait.

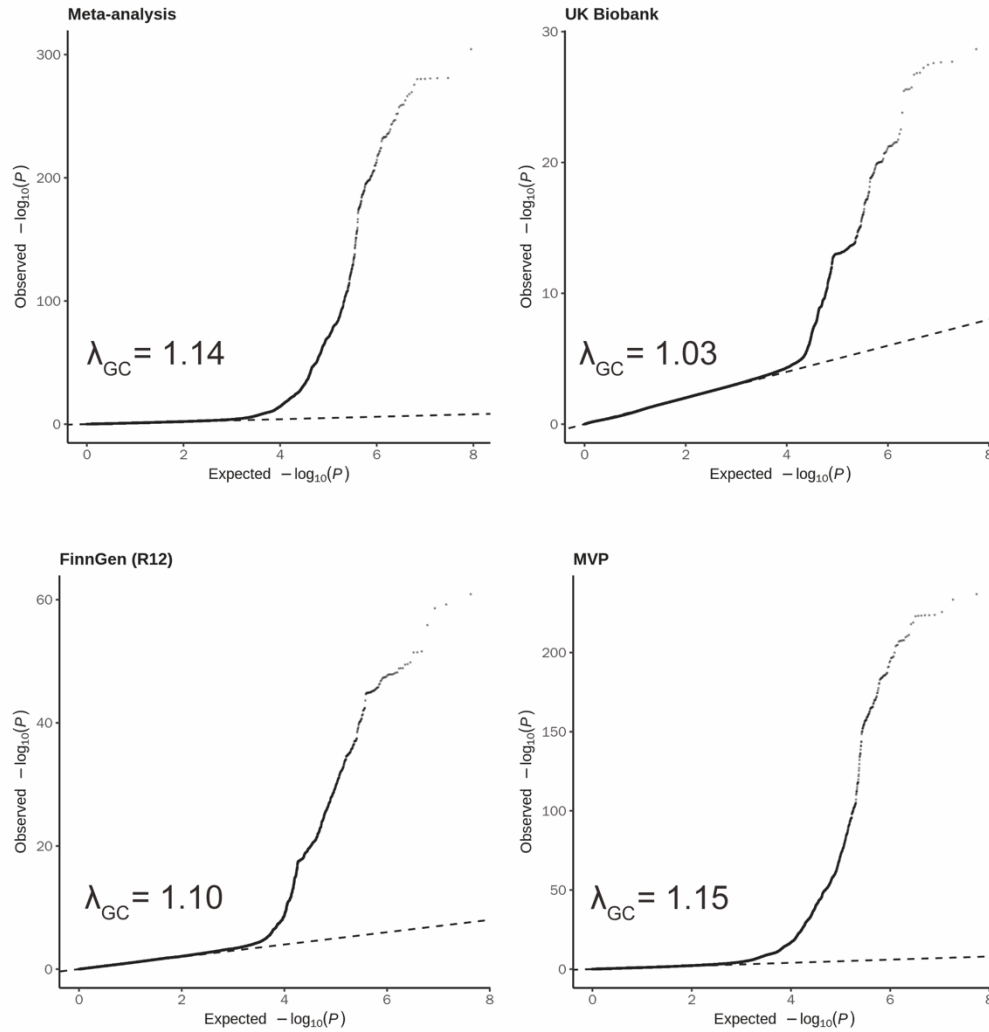
Consistent with this framework, genes located in the innermost core of the network were enriched for highly constrained and pleiotropic genes and were predominantly prioritized by GWAS rather than rare-variant analyses. Together, these findings suggest that GWAS and rare-variant collapsing analyses capture distinct yet complementary aspects of genetic architecture and should therefore be jointly considered to achieve a more complete view of disease biology.

Supplementary References





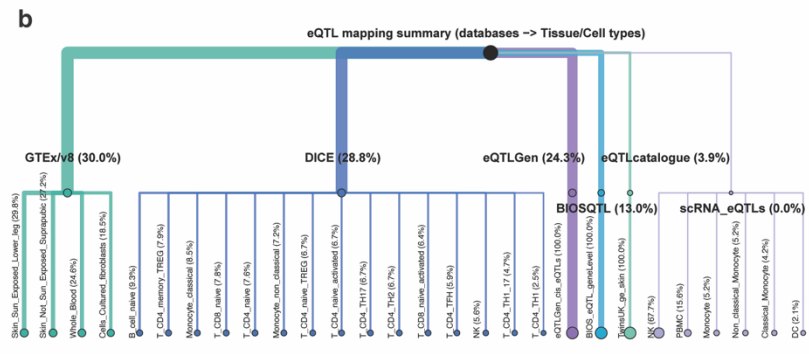
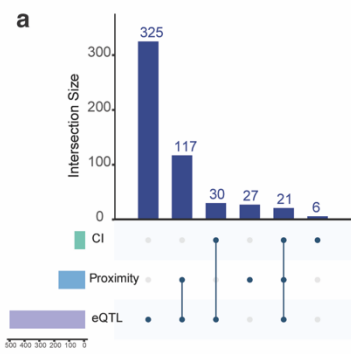
1. Bycroft, C., et al., *The UK Biobank resource with deep phenotyping and genomic data*. Nature, 2018. **562**(7726): p. 203-209.
2. Zhou, W., et al., *Efficiently controlling for case-control imbalance and sample relatedness in large-scale genetic association studies*. Nature Genetics, 2018. **50**(9): p. 1335-1341.
3. Kurki, M.I., et al., *FinnGen provides genetic insights from a well-phenotyped isolated population*. Nature, 2023. **613**(7944): p. 508-518.
4. Verma, A., et al., *Diversity and scale: Genetic architecture of 2068 traits in the VA Million Veteran Program*. Science, 2024. **385**(6706): p. eadj1182.
5. Yang, J., et al., *Conditional and joint multiple-SNP analysis of GWAS summary statistics identifies additional variants influencing complex traits*. Nat Genet, 2012. **44**(4): p. 369-75, s1-3.
6. Yang, Z., et al., *CARMA is a new Bayesian model for fine-mapping in genome-wide association meta-analyses*. Nat Genet, 2023. **55**(6): p. 1057-1065.
7. Flaherty, K.T., F.S. Hodi, and D.E. Fisher, *From genes to drugs: targeted strategies for melanoma*. Nature Reviews Cancer, 2012. **12**(5): p. 349-361.
8. Watanabe, K., et al., *Functional mapping and annotation of genetic associations with FUMA*. Nature Communications, 2017. **8**(1): p. 1826.
9. Ormond, C., et al., *Converting single nucleotide variants between genome builds: from cautionary tale to solution*. Briefings in Bioinformatics, 2021. **22**(5).
10. Corsello, S.M., et al., *The Drug Repurposing Hub: a next-generation drug library and information resource*. Nature Medicine, 2017. **23**(4): p. 405-408.
11. Knox, C., et al., *DrugBank 6.0: the DrugBank Knowledgebase for 2024*. Nucleic Acids Research, 2023. **52**(D1): p. D1265-D1275.
12. Buniello, A., et al., *Open Targets Platform: facilitating therapeutic hypotheses building in drug discovery*. Nucleic Acids Research, 2024. **53**(D1): p. D1467-D1475.
13. Ru, B., et al., *Estimation of cell lineages in tumors from spatial transcriptomics data*. Nature Communications, 2023. **14**(1): p. 568.
14. Song, L., et al., *Spatially resolved mapping of cells associated with human complex traits*. Nature, 2025. **641**(8064): p. 932-941.
15. Carss, K., et al., *Whole-genome sequencing of 490,640 UK Biobank participants*. Nature, 2025. **645**(8081): p. 692-701.
16. Rentzsch, P., et al., *CADD: predicting the deleteriousness of variants throughout the human genome*. Nucleic Acids Res, 2019. **47**(D1): p. D886-d894.
17. Vitsios, D., et al., *Prioritizing non-coding regions based on human genomic constraint and sequence context with deep learning*. Nature Communications, 2021. **12**(1): p. 1504.
18. Huber, C.D., B.Y. Kim, and K.E. Lohmueller, *Population genetic models of GERP scores suggest pervasive turnover of constrained sites across mammalian evolution*. PLoS Genet, 2020. **16**(5): p. e1008827.
19. Moore, J.E., et al., *An expanded registry of candidate cis-regulatory elements*. Nature, 2026.

20. Zhou, W., et al., *SAIGE-GENE+ improves the efficiency and accuracy of set-based rare variant association tests*. *Nature Genetics*, 2022. **54**(10): p. 1466-1469.
21. Cipriani, V., et al., *Rare disease gene association discovery in the 100,000 GenomesProject*. *Nature*, 2025.
22. Jurgens, S.J., et al., *Rare coding variant analysis for human diseases across biobanks and ancestries*. *Nature Genetics*, 2024. **56**(9): p. 1811-1820.
23. Ge, T., et al., *Polygenic prediction via Bayesian regression and continuous shrinkage priors*. *Nature Communications*, 2019. **10**(1): p. 1776.
24. Türei, D., et al., *OmniPath: integrated knowledgebase for multi-omics analysis*. *Nucleic Acids Research*, 2025. **54**(D1): p. D652-D660.
25. Spence, J.P., et al., *Specificity, length and luck drive gene rankings in association studies*. *Nature*, 2026. **649**(8098): p. 918-925.

Supplementary Figures 1-23



Supplementary Fig. 1 | Quantile–quantile plots of genome-wide association study (GWAS) summary statistics. Quantile–quantile (Q–Q) plots of observed versus expected $-\log_{10}(P)$ values for GWAS summary statistics from UK Biobank (UKB; $N = 397,762$), FinnGen (R12; $N = 384,502$), and the Million Veteran Program (MVP; $N = 615,770$), as well as the meta-analysis ($N = 1,398,034$). Genomic inflation factors (λ_{GC}) are shown for each analysis.

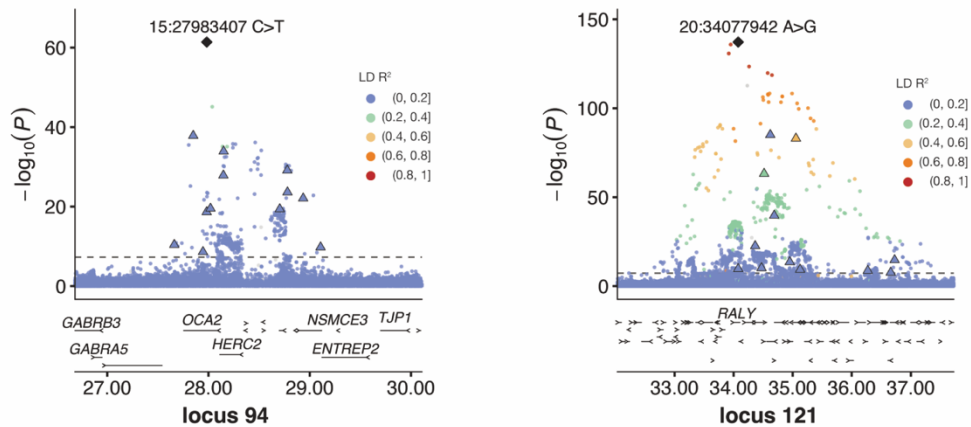


Supplementary Fig. 2 | Gene prioritization and functional enrichment analyses.

(a) UpSet plot summarizing 526 candidate genes prioritized by complementary strategies, including genomic proximity, expression quantitative trait locus (eQTL) mapping, and enhancer–promoter chromatin interaction (CI) links.

(b) Summary of eQTL mappings linking genome-wide association study (GWAS) loci to genes across datasets and tissue or cell types. eQTL associations were aggregated from multiple resources, including Genotype-Tissue Expression (GTEx v8), Database of Immune Cell Expression (DICE), eQTLGen Consortium, BIOSQTL, eQTL Catalogue, and single-cell eQTL datasets. Connections represent mappings from eQTL resources to tissue or cell types, with edge width proportional to the number of associations and node size scaled by relative contribution. Percentages indicate the proportion of mappings contributed by each dataset and, within each dataset, by individual tissue or cell types.


(c–g) Gene-set enrichment analysis of prioritized genes. Enrichment signals are shown for cancer-related copy-number variation gene sets from the Chemical and Genetic Perturbation collection (c), representative Gene Ontology Biological Process (GO:BP) pathways (d), representative Gene Ontology Cellular Component (GO:CC) pathways (e), Gene Ontology Molecular Function (GO:MF) pathways (f), and curated traits from the GWAS Catalog (g). The x axis shows $-\log_{10}(P)$, and point size indicates the proportion of overlapping genes relative to gene-set size.




Supplementary Fig. 3 | Allelic heterogeneity at locus 94 and locus 121. Association signals are shown across each locus, with variants coloured by linkage disequilibrium (LD; r^2) with the lead variant (diamond). Candidate causal variants with posterior inclusion probability (PIP) ≥ 0.99 are highlighted as triangles. The dashed horizontal line denotes the genome-wide significance threshold ($P = 5 \times 10^{-8}$). The lower panels show the positions and orientations of protein-coding genes within each plotted interval. Locus identifiers are indicated for each panel.

Supplementary Figs. 4–17 | Spatial distribution of gsMap signals and malignant melanocyte proportions across samples. For each spatial transcriptomics sample, spot-level gsMap signals ($-\log_{10}(P)$, bottom) and malignant melanocyte proportions inferred by SpaCET (top) are shown across the tissue section. Black outlines indicate tumor–stroma interface regions. gsMap signals exhibit marked spatial heterogeneity within tumor regions, with localized enrichment that is not uniformly distributed across tumor spots.



MEL74





MEL74




MEL75



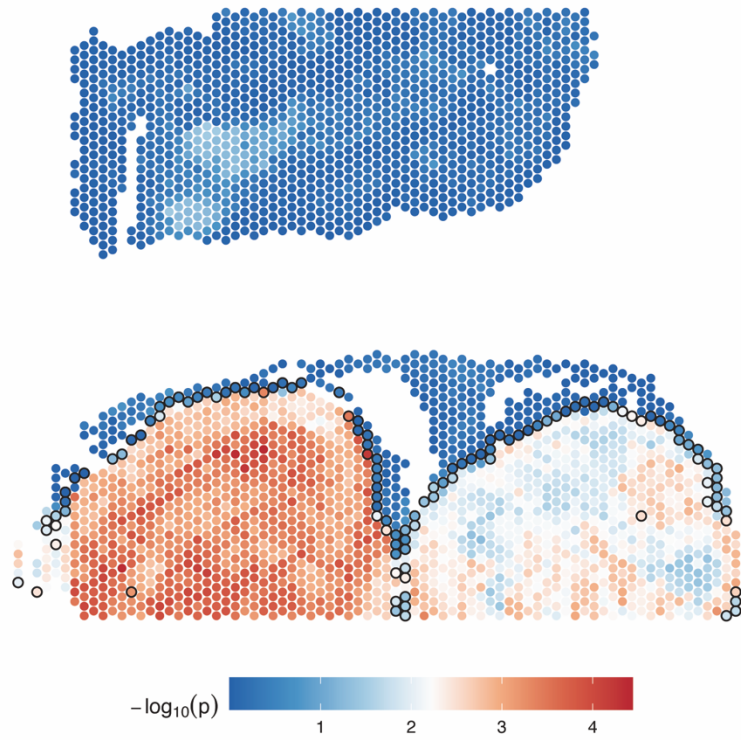
MEL75




MEL99




MEL99





MEL109





MEL109





MEL126




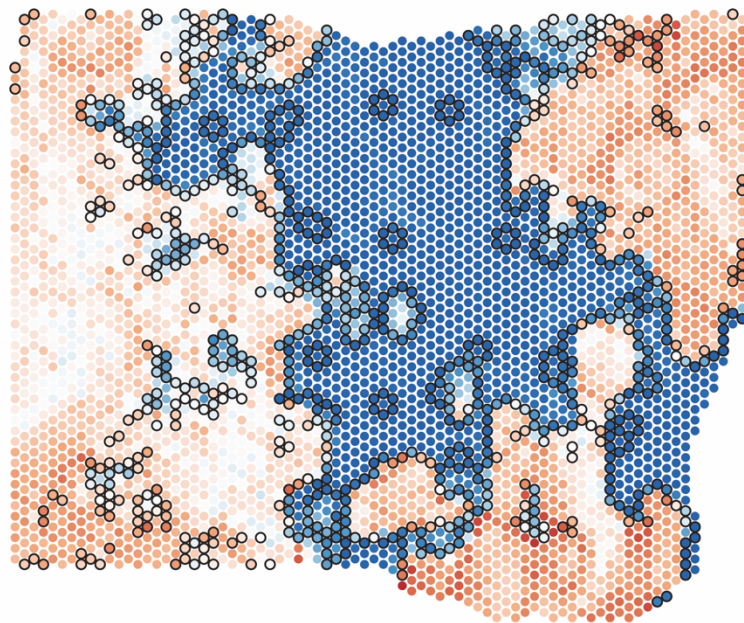
MEL126





MEL131





MEL131





MEL162





MEL162




MEL170





MEL170



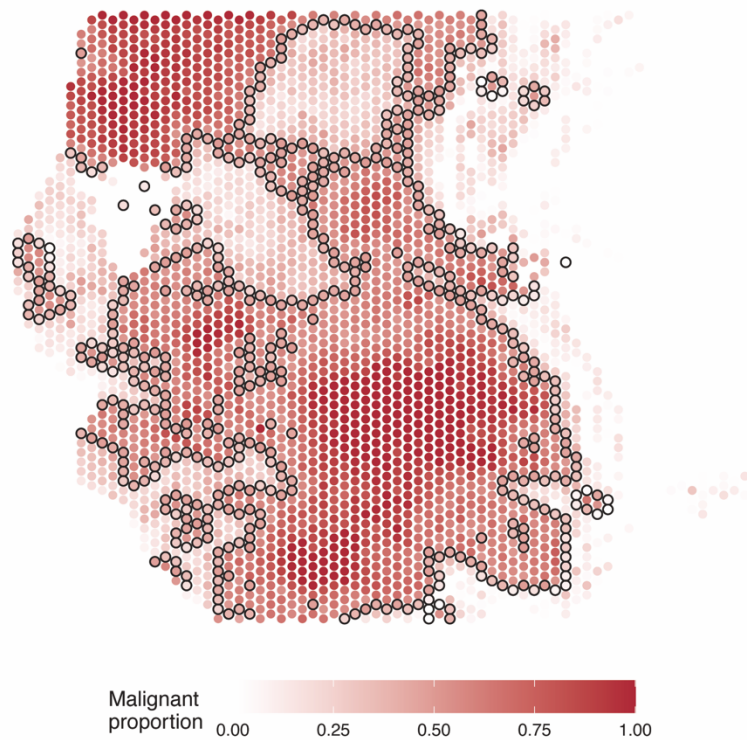
MEL176




MEL176





MEL201





MEL201




MEL222



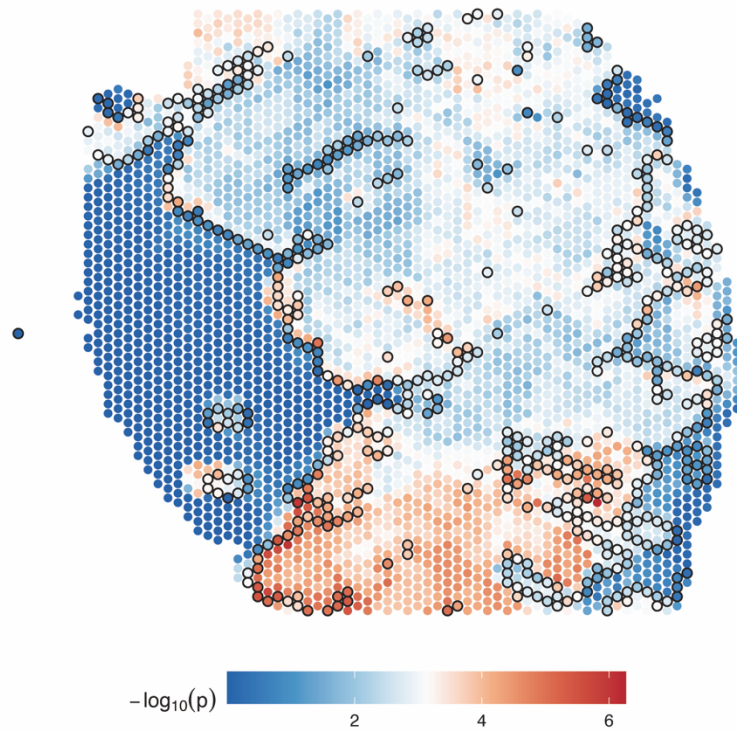
MEL222




MEL239




MEL239





MEL256





MEL256

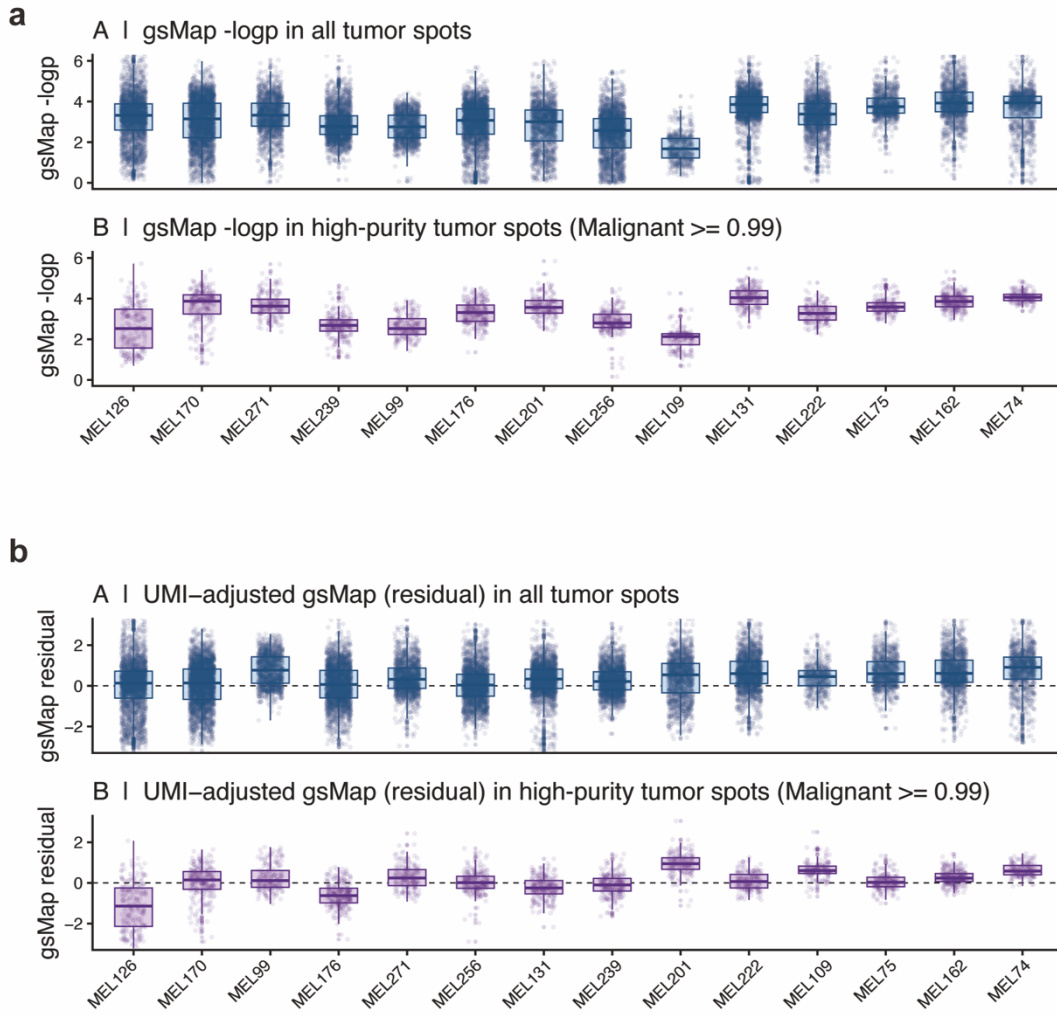


MEL271



MEL271





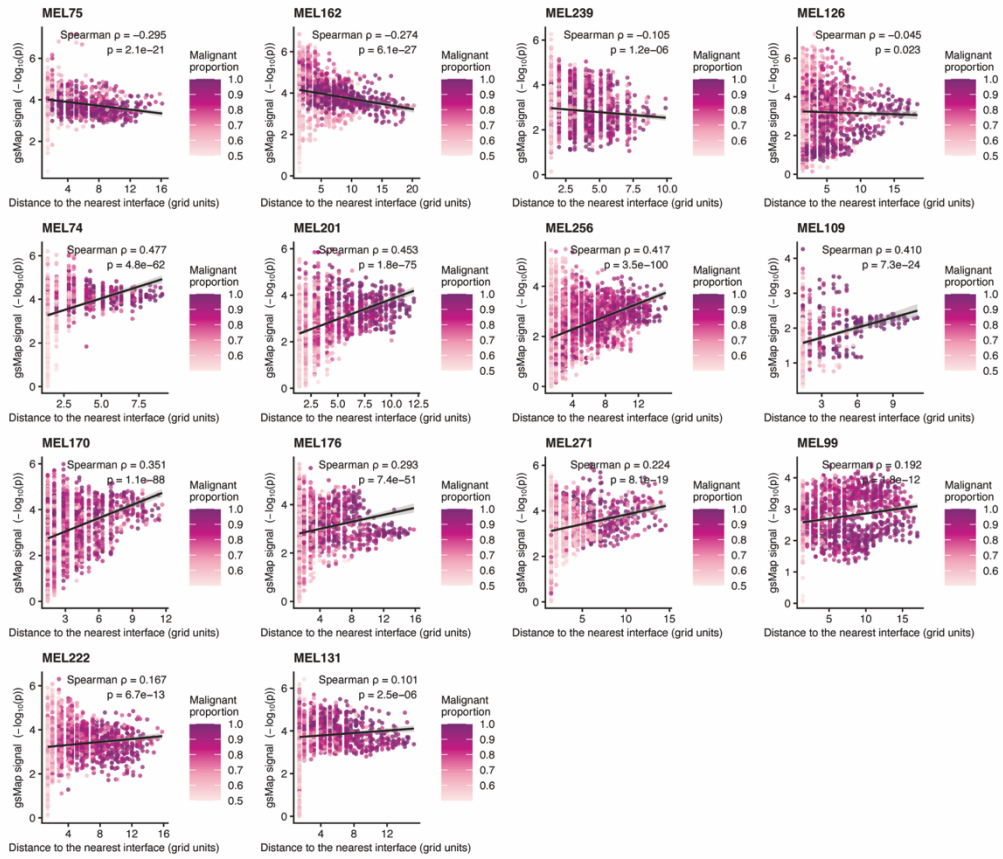
Supplementary Fig. 18 | Intratumoral heterogeneity of gsMap signals before and after adjustment for library size.

(a) Distribution of gsMap signals across tumor spots in each sample, shown for all tumor spots (top) and for high-purity tumor spots with malignant melanocyte proportion ≥ 0.99 (bottom).

(b) Distribution of UMI-adjusted gsMap signals, represented by residuals from per-sample linear regression of gsMap signal on spot-level library size (total UMI counts per spot), shown for all tumor spots (top) and high-purity tumor spots (bottom). Persistent between-spot variability after UMI adjustment indicates that intratumoral heterogeneity of gsMap signals is not explained solely by differences in sequencing depth.

Correlation between gsMap signal and the nearest distance to the interface (Tumor spots)

Negatives (most negative→0) followed by positives (largest→smallest)




Supplementary Fig. 19 | Correlations between gsMap signals and local tumor-immune spatial context across samples. For each sample, tumor-spot gsMap signals are plotted against the distance to the nearest tumor-immune interface. Points are coloured by malignant cell proportion, with black lines indicating linear fits (95% confidence intervals). Spearman correlation coefficients and P values are shown for each sample.

UK Biobank WGS dataset				
Characteristic	Control N = 364,581 ¹	Case N = 3,848 ¹	p-value ²	Overall N = 368,429 ¹
Sex			2.40e-06	
Female	198,018 (54%)	2,237 (58%)		200,255 (54%)
Male	166,563 (46%)	1,611 (42%)		168,174 (46%)
Age (years)	55.5 ± 8.1	57.5 ± 7.7	2.10e-54	55.5 ± 8.1
Ethnic background			7.15e-45	
Africa	6,531 (1.8%)	1 (<0.1%)		6,532 (1.8%)
Chinese	1,321 (0.4%)	4 (0.1%)		1,325 (0.4%)
Mixed	2,348 (0.6%)	4 (0.1%)		2,352 (0.6%)
Other	3,707 (1.0%)	4 (0.1%)		3,711 (1.0%)
South Asian	8,330 (2.3%)	4 (0.1%)		8,334 (2.3%)
White	340,050 (94%)	3,819 (100%)		343,869 (94%)
Unknown	2,294	12		2,306
Skin color			5.69e-67	
Black	3,193 (0.9%)	0 (0%)		3,193 (0.9%)
Brown	12,174 (3.4%)	14 (0.4%)		12,188 (3.4%)
Dark olive	7,061 (2.0%)	41 (1.1%)		7,102 (2.0%)
Fair	240,719 (67%)	2,831 (74%)		243,550 (67%)
Light olive	68,634 (19%)	499 (13%)		69,133 (19%)
Very fair	26,985 (7.5%)	426 (11%)		27,411 (7.6%)
Unknown	5,815	37		5,852
Hair color			4.98e-82	
Black	31,623 (8.7%)	139 (3.6%)		31,762 (8.7%)
Blonde	37,432 (10%)	575 (15%)		38,007 (10%)
Dark brown	137,472 (38%)	1,183 (31%)		138,655 (38%)
Light brown	137,750 (38%)	1,602 (42%)		139,352 (38%)
Other	4,272 (1.2%)	44 (1.1%)		4,316 (1.2%)
Red	14,692 (4.0%)	300 (7.8%)		14,992 (4.1%)
Unknown	1,340	5		1,345


¹ n (%); Mean ± SD
² Pearson's Chi-squared test; Welch Two Sample t-test

Supplementary Fig. 20 | Characteristics of the UK Biobank WGS cutaneous melanoma cohort. Demographic and phenotypic characteristics of individuals with whole-genome sequencing data, stratified by cutaneous melanoma status (364,581 controls and 3,848 cases; total N = 368,429). Continuous variables are presented as mean ± s.d., and categorical variables as number (percentage). P values were calculated using Welch's two-sample t-test (continuous variables) and Pearson's chi-squared test (categorical variables).

A



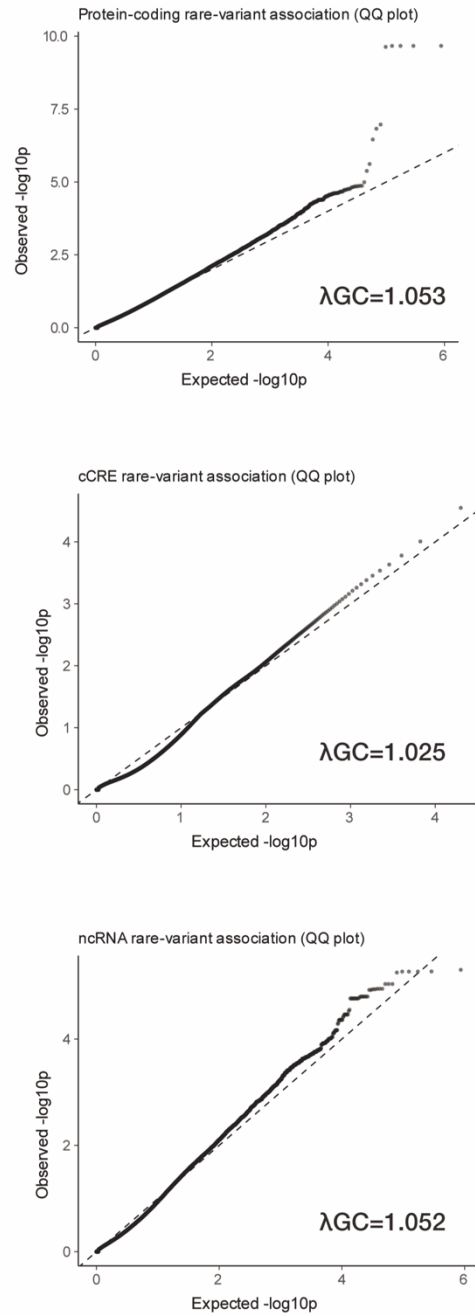
B



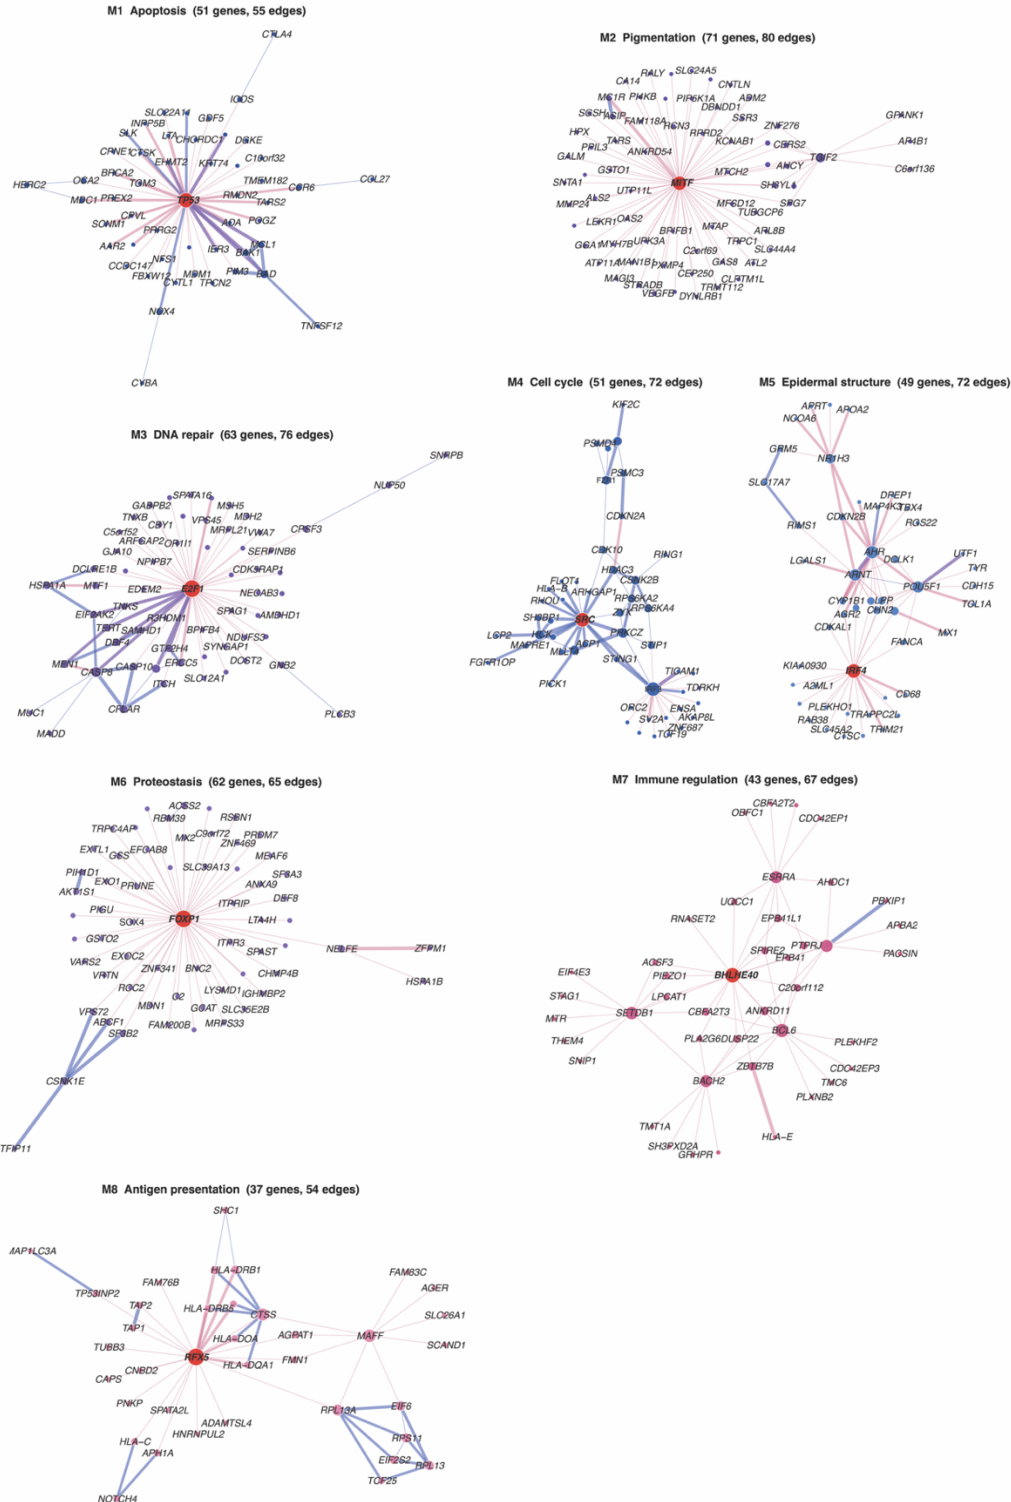
Supplementary Fig. 21 | Distribution of diagnostic codes among cutaneous melanoma cases in the UK Biobank WGS cohort.

(a) Distribution of International Classification of Diseases, Tenth Revision (ICD-10) codes among cutaneous melanoma cases.

(b) Distribution of International Classification of Diseases for Oncology, Third Edition (ICD-O-3) morphology/behaviour codes among cases. Bars indicate the number of cases in each category, and total numbers of coded cases are shown for each panel.



Supplementary Fig. 22 | Calibration of rare-variant association tests across annotation classes. Quantile–quantile (QQ) plots comparing observed versus expected $-\log_{10}(P)$ values for rare-variant association tests in protein-coding genes (top), candidate cis-regulatory elements (cCREs; middle) and ncRNAs (bottom). The dashed line indicates the null expectation. Genomic inflation factors (λ_{GC}), estimated using the 95th percentile of the null χ^2 distribution, are shown for each annotation class. Overall, minimal inflation was observed (protein-coding $\lambda_{GC} = 1.053$; cCRE $\lambda_{GC} = 1.025$; ncRNA $\lambda_{GC} = 1.052$), indicating well-calibrated association statistics.



Supplementary Fig. 23 | Within-module network structures and hub genes. Gene-level subnetworks for each of the eight modules identified by consensus Louvain clustering (M1–M8). Nodes represent genes, with size proportional to within-module degree. The top hub gene in each module (highest degree) is highlighted in red. Edges denote regulatory interactions, coloured by type—transcriptional (pink), post-translational (blue) or both (purple)—with width proportional to the strength of supporting evidence. Module identity and size (number of genes and edges) are indicated in each panel.