A sequence-based global map of regulatory activity for deciphering human genetics

Kathleen M. Chen[1,2], Aaron K. Wong[2], Olga G. Troyanskaya[1,2,3*], Jian Zhou[4,*]

[1]Department of Computer Science, Princeton University, Princeton, New Jersey, United States of America
[2]Flatiron Institute, Simons Foundation, New York, New York, United States of America
[3]Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, New Jersey, United States of America
[4]Lyda Hill Department of Bioinformatics, University of Texas Southwestern Medical Center, Dallas, Texas, United States of America

Correspondence to:
Olga G. Troyanskaya, email: ogt@cs.princeton.edu
Jian Zhou, email: jian.zhou@utsouthwestern.edu

Abstract

Sequence is at the basis of how the genome shapes chromatin organization, regulates gene expression, and impacts traits and diseases. Epigenomic profiling efforts have enabled large-scale identification of regulatory elements, yet we still lack a sequence-based map to systematically identify regulatory activities from any sequence, which is necessary for predicting the effects of any variant on these activities. We address this challenge with Sei, a new framework for integrating human genetics data with sequence information to discover the regulatory basis of traits and diseases. Our framework systematically learns a vocabulary for the regulatory activities of sequences, which we call sequence classes, using a new deep learning model that predicts a compendium of 21,907 chromatin profiles across >1,300 cell lines and tissues, the most comprehensive to-date. Sequence classes allow for a global view of sequence and variant effects by quantifying diverse regulatory activities, such as loss or gain of cell-type-specific enhancer function. We show that sequence class predictions are supported by experimental data, including tissue-specific gene expression, expression QTLs, and evolutionary constraints based on population allele frequencies. Finally, we applied our framework to human genetics data. Sequence classes uniquely provide a non-overlapping partitioning of GWAS heritability by tissue-specific regulatory activity categories, which we use to characterize the regulatory architecture of 47 traits and diseases from UK Biobank. Furthermore, the predicted loss or gain of sequence class activities suggest specific mechanistic hypotheses for individual regulatory pathogenic mutations. We provide this framework as a resource to further elucidate the sequence basis of human health and disease.

Introduction

Deciphering how regulatory functions are encoded in genomic sequences is a major challenge in understanding how genome variation links to phenotypic traits. Cell-type-specific regulatory activities encoded in elements such as promoters, enhancers, and chromatin insulators are critical to defining the complex expression programs essential for multicellular organisms, like those affecting cell lineage specificity and development. The majority of disease-associated variants from genome-wide association studies (GWAS) are located in noncoding regions[1] and may perturb regulatory elements, yet without knowing the sequence dependencies of these elements--that is, how changes in sequence affect regulatory activities--we cannot predict the impact of these variants and uncover the regulatory mechanisms contributing to complex diseases and traits. For instance, a variant in an embryonic stem cell-specific enhancer may turn off the expression of a gene critical for early development, while other variants in the same region may increase enhancer activity or have no effect at all.

Substantial progress has been made in the experimental profiling and integrative analysis of epigenomic marks, such as histone marks and DNA accessibility, across a wide range of tissues and cell types[2–4]. Histone marks are commonly used to identify regulatory elements; for example, H3K4me3 can indicate active promoter regions and H3K27ac/H3K4me1 can indicate

active enhancer regions. Moreover, histone marks and chromatin accessibility can be integrated with chromatin state models[5–9]. These works have been instrumental to annotating the genome with regulatory elements across many tissues.

At the same time, deep learning sequence modeling techniques have been successfully applied to learn sequence features that are predictive of transcription factor binding and histone modifications[10–15]. These models are powerful tools for inferring the impact of sequence variation at the chromatin level. However, each chromatin-level prediction can only inform a very specific aspect of sequence--for example, whether a variant causes an increase or decrease of C/EBP-β binding. We continue to lack a global, integrative view of sequence regulatory activities, including all major aspects of cis-regulatory functions, such as tissue-specific or broad enhancer and promoter activities. This limits our ability to interpret the integrated effects of all chromatin-level perturbations caused by genomic variants and determine their impact on human health and diseases.

We address this challenge by creating a global map for sequence regulatory activity based on a new deep-learning-based framework called Sei. This framework introduces a new sequence model that predicts a comprehensive compendium of 21,907 publicly available chromatin profiles--the broadest set to-date--and uses the model to quantitatively characterize regulatory activities for any sequence with a novel vocabulary of sequence classes. Sequence classes cover diverse types of regulatory activities, such as promoter or cell-type-specific enhancer activity, across the whole genome by integrating sequence-based predictions from histone marks, transcription factors, and chromatin accessibility across a wide range of cell types. For example, 'embryonic stem cell-specific enhancer' sequence class activity may be estimated from the predicted binding of multiple transcription factors including Pou5F1, Sox2, and Nanog, as well as various histone marks, on a sequence. Importantly, sequence classes can be used to both classify and quantify the regulatory activities of any sequence based on predictions made by the deep learning sequence model. Therefore, sequence classes allow for the quantitative mapping of any mutation to its impact on cell-type-specific regulatory activities.

The Sei framework thus provides an interpretable and systematic integration of sequence-based regulatory activity predictions (intrinsic information, based on sequence function) with human genetics data (extrinsic information, based on variant-phenotype association) for discovering the regulatory basis of human traits and disease. We applied our framework to characterize disease- and trait-associated regulatory disruptions by combining sequence class information and UK Biobank GWAS data. Sequence classes provide a non-overlapping partitioning of heritability in GWAS by regulatory activity, which we use to profile the regulatory architecture of all 47 diseases and traits in UK Biobank GWAS[16].

Moreover, variant effect prediction at the sequence-class-level newly enables the interpretation of regulatory mechanisms for individual disease mutations and can differentiate between gain-of-

function and loss-of-function regulatory mutations. The regulatory and tissue-specific view provided by sequence classes suggests potential new mechanisms for individual disease-associated variants: for example, we used sequence classes to link mutations in blood-related diseases with previously unknown mechanisms to the malfunctioning of cell-type-specific enhancers.

We provide the Sei framework as a resource for systematically classifying and scoring any sequence and variant with sequence classes, additionally providing the Sei model predictions for the 21,907 chromatin profiles underlying the sequence classes. The framework can be run using the code at https://github.com/FunctionLab/sei-framework and a user-friendly interface to view Sei sequence class predictions is available at hb.flatironinstitute.org/sei.

Results

Developing a comprehensive sequence model for 21,907 chromatin profiles

To capture the widest range of sequence features that are predictive of regulatory activities, we first developed a new deep learning sequence model, which we refer to as the Sei model, that enables the base-level interpretation of sequences by predicting 21,907 genome-wide cis-regulatory targets--including 9,471 transcription factor profiles, 10,064 histone mark profiles and 2,372 chromatin accessibility profiles--with single nucleotide sensitivity. The majority of this data (19,905 profiles) is from the Cistrome Project[4], a resource that uniformly processes and annotates public ChIP-, DNase-, and ATAC-seq datasets, and the remaining chromatin profiles were processed by the ENCODE[2] and Roadmap Epigenomics[3] projects. The Sei model encompasses an estimated ~1000 non-histone DNA-binding proteins (which we refer to as transcription factors), 77 histone marks, and chromatin accessibility across >1300 cell lines and tissues (Supplementary Files 1, 2).

To efficiently predict 21,907 chromatin profiles from sequence, we designed a novel model architecture (Supplementary Figure 1) and improved our training pipeline. The Sei model uses a new residual-block architecture with a dual linear and nonlinear path design: the linear path allows for fast and statistically efficient training, while the nonlinear path offers strong representation power and the capability to learn complex interactions. For scaling and performance, we introduced a layer of spatial basis functions, which integrates information across spatial locations with much higher memory efficiency than fully connected layers. The model takes as input a 4kb length sequence and predicts the probabilities of 21,907 targets at the center position. Our model training pipeline was updated to improve training speed and performance by using on-the-fly sampling, which reduces overfitting by generating new training samples for every training step.

The model achieved an average area under the receiver-operating characteristic (AUROC) of 0.972 and average area under the precision-recall curve (AUPRC) of 0.409 across all 21,907 chromatin

profiles (Supplementary Figure 2). In addition to accurately predicting individual profiles, the predictions also recapitulated the correlation structure of these profiles, which indicates that the Sei model is able to capture the co-localization patterns of chromatin profiles (Supplementary Figure 3). Furthermore, the Sei model also improved over our best previously published model, DeepSEA "Beluga"[12], on the 2002 chromatin profiles predicted by both models by 19% on average (as measured by AUROC/1-AUROC, Supplementary Figure 4).

Therefore, the Sei model is the most comprehensive chromatin-level sequence model to-date, and offers an expansive new resource for sequence and variant interpretation.

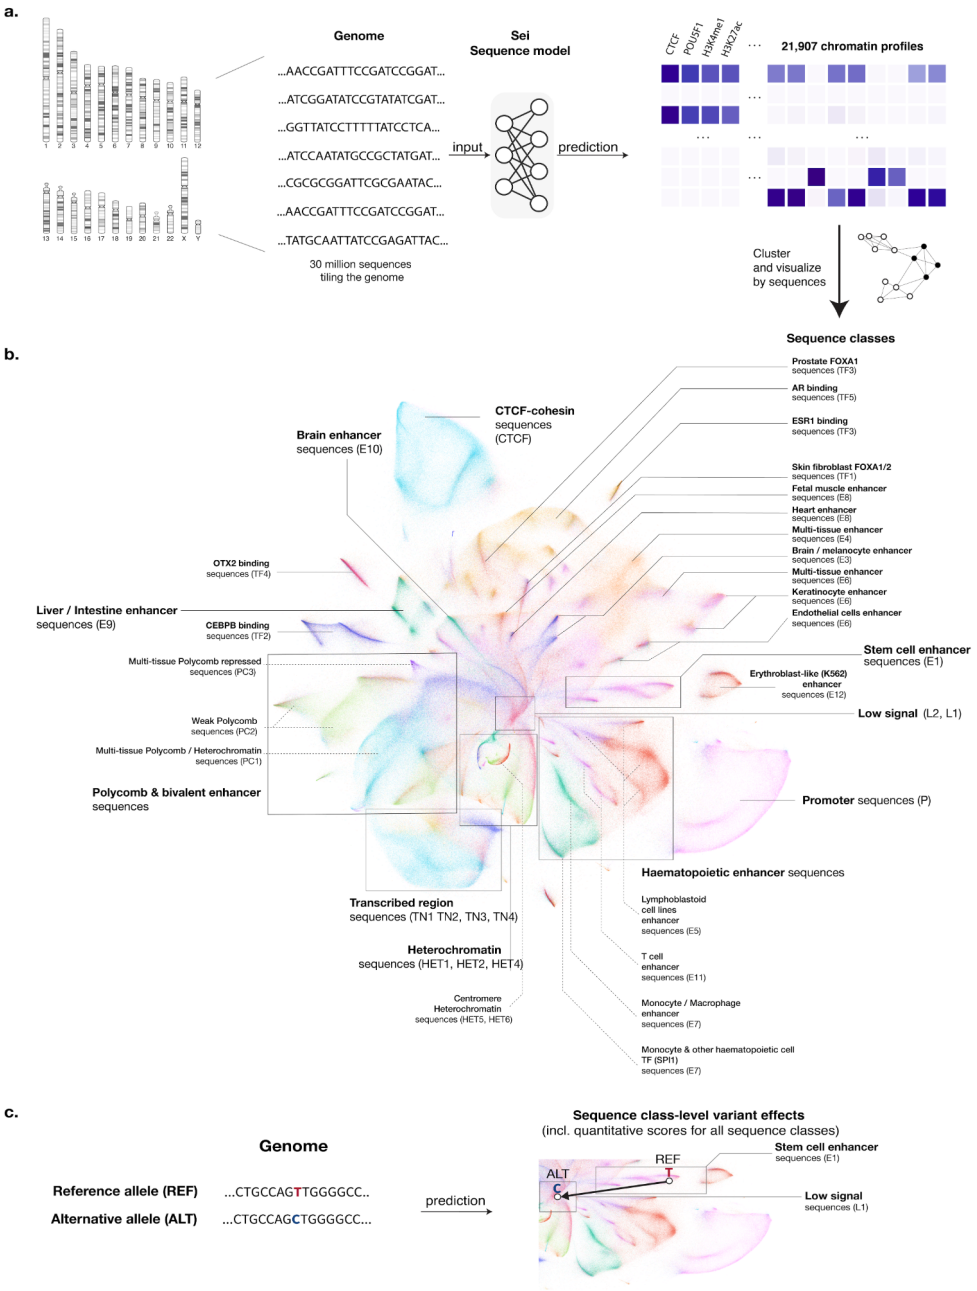

Figure 1. Mapping the global regulatory landscape of genomic sequences.
a, Overview of the Sei framework for systematic prediction of sequence regulatory activities. Sequence classes are extracted from the predicted chromatin profiles of 30 million sequences evenly tiling the genome. The predictions were made by Sei, a new deep convolutional network sequence model trained on 21,907 chromatin profiles. Specifically, classes are identified by applying Louvain community detection to the nearest-neighbor graph of 180 principal components extracted from the predictions data. **b,** Visualizing the global regulatory landscape of human genome sequences discovered by this approach with UMAP. Major sequence classes include cell-type-specific enhancer classes, CTCF-cohesin, promoter, TF-specific, and heterochromatin/centromere classes. **c,** This framework is further applied to predict sequence-class-level genome variant effects, quantified by changes in sequence class scores.

Defining sequence classes using a sequence model from whole genome sequences

Next, we applied the Sei model to develop a global, quantitative map from genomic sequences to specific classes of regulatory activities, which we term sequence classes, by integrating the wide range of chromatin profiles predicted by Sei. Sequence classes are therefore mapped directly from sequence, and each sequence class represents a distinct program of regulatory activities across tissues and cell types as covered by the Sei model. Furthermore, sequence classes allow for the mapping of any sequence to quantitative scores that represent a broad spectrum of regulatory activities.

To cover the whole spectrum of sequence activities, we identified sequence classes from Sei predictions for 30 million sequences uniformly tiling the whole genome (4kb windows with 100bp step size). We visualized the global structure of sequence regulatory signals as represented by the model's chromatin profile predictions with nonlinear dimensionality reduction techniques[17,18] (Figure 1) and applied Louvain community clustering[19] to these predictions to categorize the 30 million sequences into 40 sequence classes (Figure 1a).

This visualization of human genome sequences demonstrates the global organization of sequence regulatory activities (Figure 1b). The center of the visualization contains sequences with weak or no regulatory activity based on histone mark and TF enrichment, and sequences with specific regulatory activities radiate outwards, establishing a continuum from no activity to strong specific activity. Different branches of sequences are enriched in distinct chromatin modifications and transcription factors, and sequences with similar regulatory activities are grouped together. For example, tissue-specific enhancer sequences were predominantly grouped by tissue in the visualization (Figure 1b). In addition, sequences with repressive Polycomb marks were spatially adjacent to H3K9me3-marked heterochromatin sequences (Figure 1b), reflecting their extensive crosstalk in epigenetic silencing[20–22]. Notably, promoter-proximal and CTCF-cohesin binding sequences form two well-defined clusters that are separated from other sequences, which may reflect the distinct nature of these activities (Figure 1b).

The sequence classes identified from whole genome sequences recapitulate the sequence organization shown in the visualization and provide a basis for summarizing sequence activities globally. To facilitate intuitive interpretation of sequence classes, we named them based on the corresponding enrichment of cis-regulatory profiles (Figure 2a, Supplementary Figures 5-11); specifically, we label each sequence class with a functional group acronym and index denoting the rank of the sequence class within the group (Supplementary Figure 12, e.g. E1 encompasses a larger proportion of the genome than E2). Because genomic sequences encode their regulatory activity programs across all cell types, sequence classes also show distinct activity patterns across cell types and tissues. We label sequence classes primarily based on their active, cell-type-specific regulatory activities--in particular, promoter and enhancer activities. Therefore, sequence classes that are not labeled as enhancer ('E') or promoter ('P') generally lack enhancer or promoter activity in any cell type predicted by Sei.

In summary, sequence classes contain 1 'P' promoter class, which is most strongly enriched in the active promoter histone mark H3K4me3 across all cell types (Figure 2a, Supplementary Figure 5); 12 'E' enhancer classes, which are strongly enriched in enhancer histone marks, such as H3K4me1 and H3K27ac, and transcription factors relevant to their activities in select cell types (e.g. PU.1/Spi1 in the E7 monocyte/macrophage enhancer class, HNF4-α in E9 liver/intestine, and Sox2/Nanog/Pou5f1 in E1 stem cell), and often display repressive H3K27me3 marks in inactive cell types (Figure 2a, Supplementary Figures 6-8, 11); 1 'CTCF' sequence class, which is strongly enriched in CTCF and cohesin (Figure 2a, Supplementary Figure 11); 5 'TF' sequence classes, which are enriched in a few specific transcription factors (e.g. CEBPB sequence class) but have weak or no enhancer mark enrichment (Figure 2a, Supplementary Figure 11); 4 'PC' Polycomb classes, which are enriched in the Polycomb-repressed region mark H3K27me3 and generally not enriched in active promoter or enhancer marks (Figure 2a, Supplementary Figure 8); 6 'HET' heterochromatin classes, which are enriched in the heterochromatin mark H3K9me3 (Figure 2a, Supplementary Figure 9); 4 'TN' sequence classes, which are enriched in transcription elongation marks H3K36me3 or H3K79me2 (Figure 2a, Supplementary Figure 10); and finally, 7 'L' (low signal) sequence classes, which are not strongly enriched in any of the above marks (Figure 2a). As a whole, the 40 sequence classes cover >97.4% of the genome (Supplementary Figure 12).

Beyond classifying genomic sequences to sequence classes, we define sequence class scores to provide a global and quantitative representation of sequence regulatory activities. This for the first time allows us to (1) predict the regulatory activity for any sequence and (2) quantify the changes in regulatory activity caused by any sequence variant. Sequence class scores summarize predictions for all 21,907 chromatin profiles based on weights specific to each sequence class, which are computed by projecting Sei predictions onto unit-length vectors that point to the center of each sequence class. Sequences that score highly for a particular sequence class have high predictions for the chromatin profiles associated with that class. Sequence class scores thus allow for the quantification of the regulatory activity of any sequence, and the impact of a variant can

253 be represented by the difference between the sequence class scores for the reference and
254 alternative alleles.
255
256 **Enhancer sequence classes predict tissue-specific gene expression**
257 The group of sequences that are likely most impactful to tissue-specific gene expression
258 regulation are the enhancer ('E') sequence classes, thus here we assessed the association of
259 enhancer sequence class scores with tissue-specific gene expression.
260
261 In the visualization of sequence regulatory activities, sequence classes with different cell type-
262 and tissue-specific enhancer activities are localized to distinct subregions (Figure 1b). 'E'
263 sequence classes capture both specific and broad enhancer activities. Based on enhancer mark
264 enrichment (Supplementary Figures 6, 7), E7 is specific for monocyte/macrophage, E11 is
265 specific for T-cell, E5 is specific for lymphoblastoid/B-cell-like cell lines, E9 is specific for liver
266 and intestine, E1 is specific for embryonic stem cells & induced pluripotent stem cells, and E10
267 and E3 are specific for brain (Figure 1, 2; all enrichments stated are significant with $p<2.2e-16$,
268 Fisher's exact test, two-sided). In contrast, broad enhancer sequence classes can either
269 encompass enhancer activity in similar cell types across different tissues, such as fibroblast (E2)
270 and epithelial (E6) cell types (Supplementary Figures 6, 7), or encompass enhancer activity in
271 many different cell types; for example, E4 is enriched in fibroblast, muscle, astrocytes,
272 osteoblast, epithelial, and other cell types. Consistent with their predicted enhancer activities, the
273 coverage of 'E' sequence class annotations within a 10kb window to transcription start sites
274 (TSS) are correlated with the differential expression patterns of these genes in the corresponding
275 cell types over the tissue-average (Figure 2b).
276
277 Since sequence class scores allow us to systematically predict the effects of variants on higher-
278 level regulatory functions, we can estimate whether a given variant diminishes, maintains, or
279 increases the enhancer activity of a sequence based on the difference between the sequence class
280 scores for the reference and alternative alleles. Evaluated on GTEx eQTL data[23], we found that
281 variants predicted to increase 'E' sequence class activity were significantly positively correlated
282 with higher gene expression, whereas those predicted to increase 'PC' sequence class activity were
283 significantly negatively correlated with gene expression--consistent with the expected repressive
284 role of 'PC' sequence class activities (Figure 2c). Therefore, sequence classes can distinguish the
285 effects of variants on gene expression based on their consequences in regulatory activities.
286
287
288

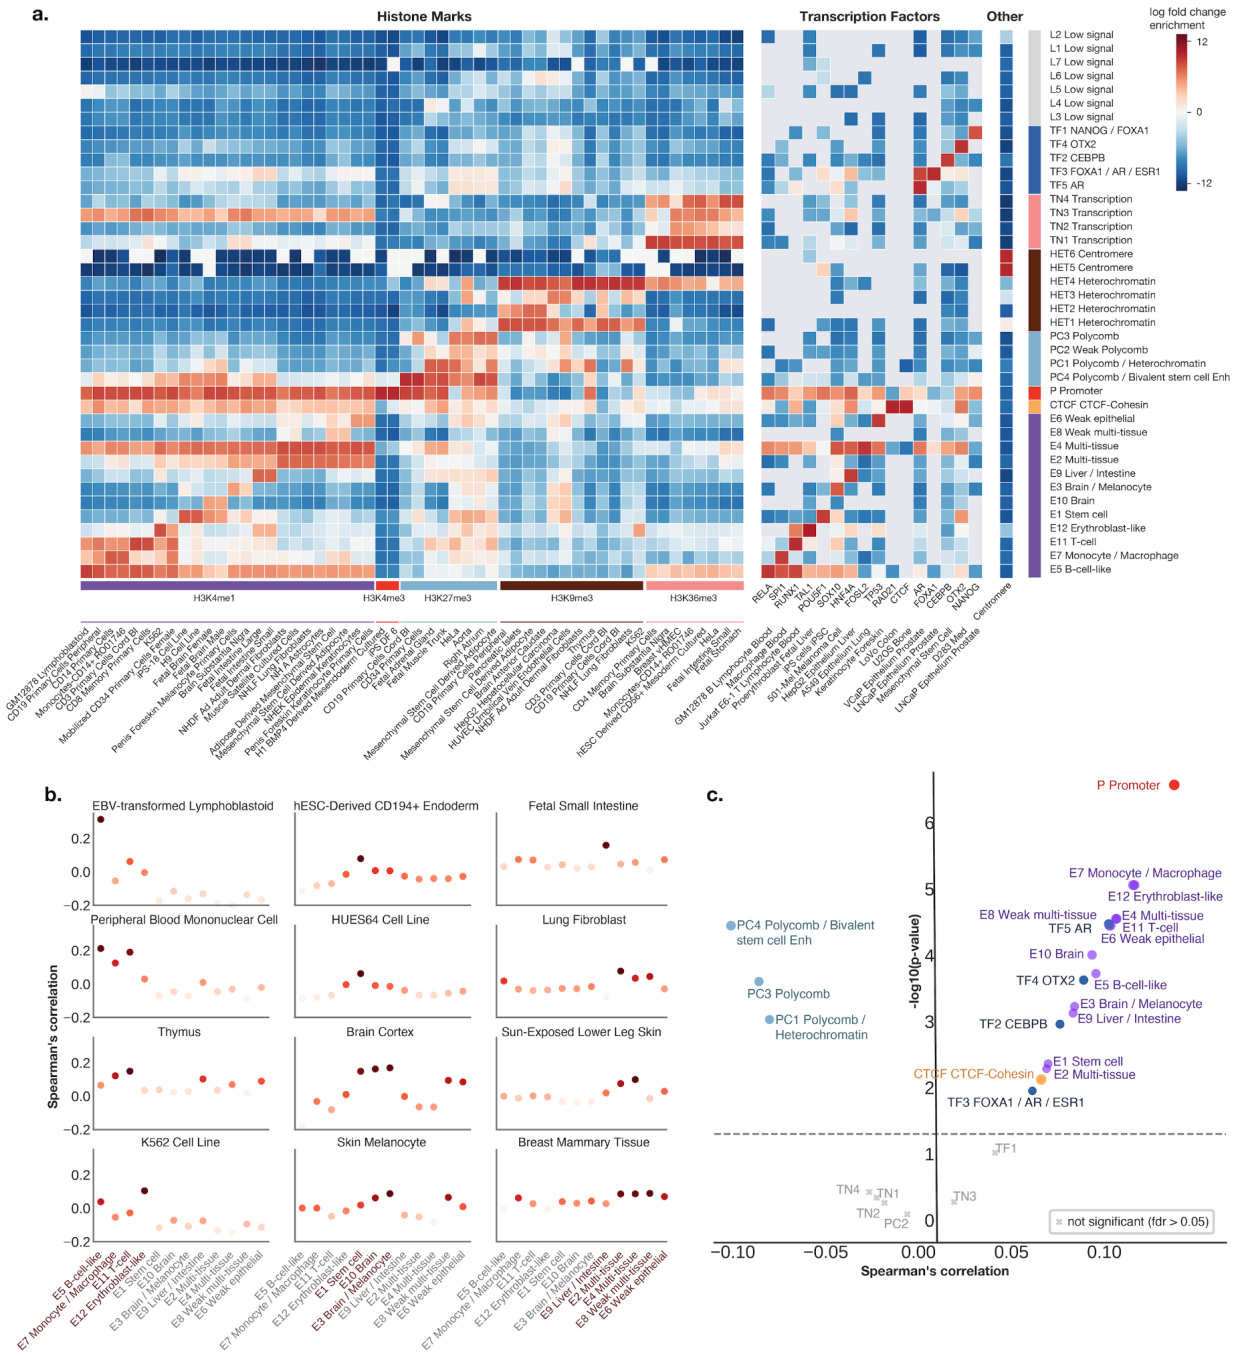

Figure 2. Sequence classes predict cell-type-specific regulatory activities and directional, expression-altering variant effects. a, Sequence-class-specific enrichment of histone marks, transcription factors, and repeat annotations. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap. **b,** Enhancer sequence classes near transcription start sites are correlated with cell-type-specific gene expression in the applicable tissue or cell types (see Methods). The y-axis shows the Spearman correlation between the proportion of each sequence class annotation within 10kb to TSS and the tissue-specific differential gene expression (fold over tissue-average). **c,** Regulatory sequence-class-level variant effects are predictive of directional GTEx variant gene expression effects. The x-axis shows Spearman

correlations between the predicted sequence-class-level variant effects and the signed GTEx variant effect sizes (slopes) for variants with strong predicted effects near transcription start sites (Methods) and the y-axis shows the corresponding log10 p-values. All colored dots are above the Benjamini-Hochberg FDR < 0.05 threshold.

Regulatory sequence classes are under evolutionary constraints

Variants that alter regulatory activities of sequences often disrupt gene regulation and are therefore expected to impact human health and disease. We tested this expectation by comparing human population genome variant allele frequencies[24] based on the sequence class in which each variant is located and the predicted variant effect on that sequence class. Indeed, we found that variants localized in E-, P-, and CTCF- sequence classes have lower population allele frequency than variants in other sequence classes (Figure 3a). More importantly, variants predicted to strongly perturb regulatory sequence classes had significantly lower allele frequencies than variants that weakly perturb these classes (Figure 3a), consistent with the hypothesis that disruption of regulatory sequence class activities has a major negative impact on fitness, which we refer to as a negative selection signature.

Specifically, we observed strong negative selection signatures for variants assigned to all E, CTCF and P sequence classes (Figure 3). Multi-tissue enhancer sequence classes E4 and E2 and the brain enhancer sequence class E10 show the strongest association of predicted sequence-class-level variant effects and allele frequencies. Notably, for the CTCF sequence class, only negative variant effects--decreasing sequence class activity--appear to be under very strong constraints, suggesting that CTCF sites are generally tolerant to positive effect mutations that further increase CTCF binding. This is in contrast to the generally deleterious impact of both increase and decrease of enhancer and promoter activities. As expected, TN sequence classes, which overlap with protein-coding regions, are among the sequence classes with the lowest allele frequency (Supplementary Figure 13).

In contrast, those assigned to HET, PC, TF, and L sequence classes generally did not show strong negative selection signatures and had higher overall allele frequencies (Supplementary Figure 13). Importantly, this does not suggest that Polycomb or transcription factors are inessential: the HET, PC, TF, and L classes generally do not show strong enhancer or promoter histone mark enrichment in any cell type (with the exception of bivalent marks in stem cells observed in PC4), and thus they are expected to play less major roles in gene expression regulation. However, Polycomb-related regulation is likely critical for ‘E’ and ‘P’ sequence classes, which are often Polycomb-repressed in some cell types but enhancers or promoters in other cell types (Supplementary Figures 5-8). Similarly, we expect that TF binding plays a central role in ‘E’ classes that are highly enriched in relevant TFs (Figure 2a, Supplementary Figure 11).

Therefore, sequence classes show distinct evolutionary constraints, and 'E' enhancer sequence classes show the strongest bidirectional constraints. This suggests that both increases and decreases of enhancer activity are expected to lead to deleterious effects on fitness, highlighting the importance of precisely controlling gene expression.

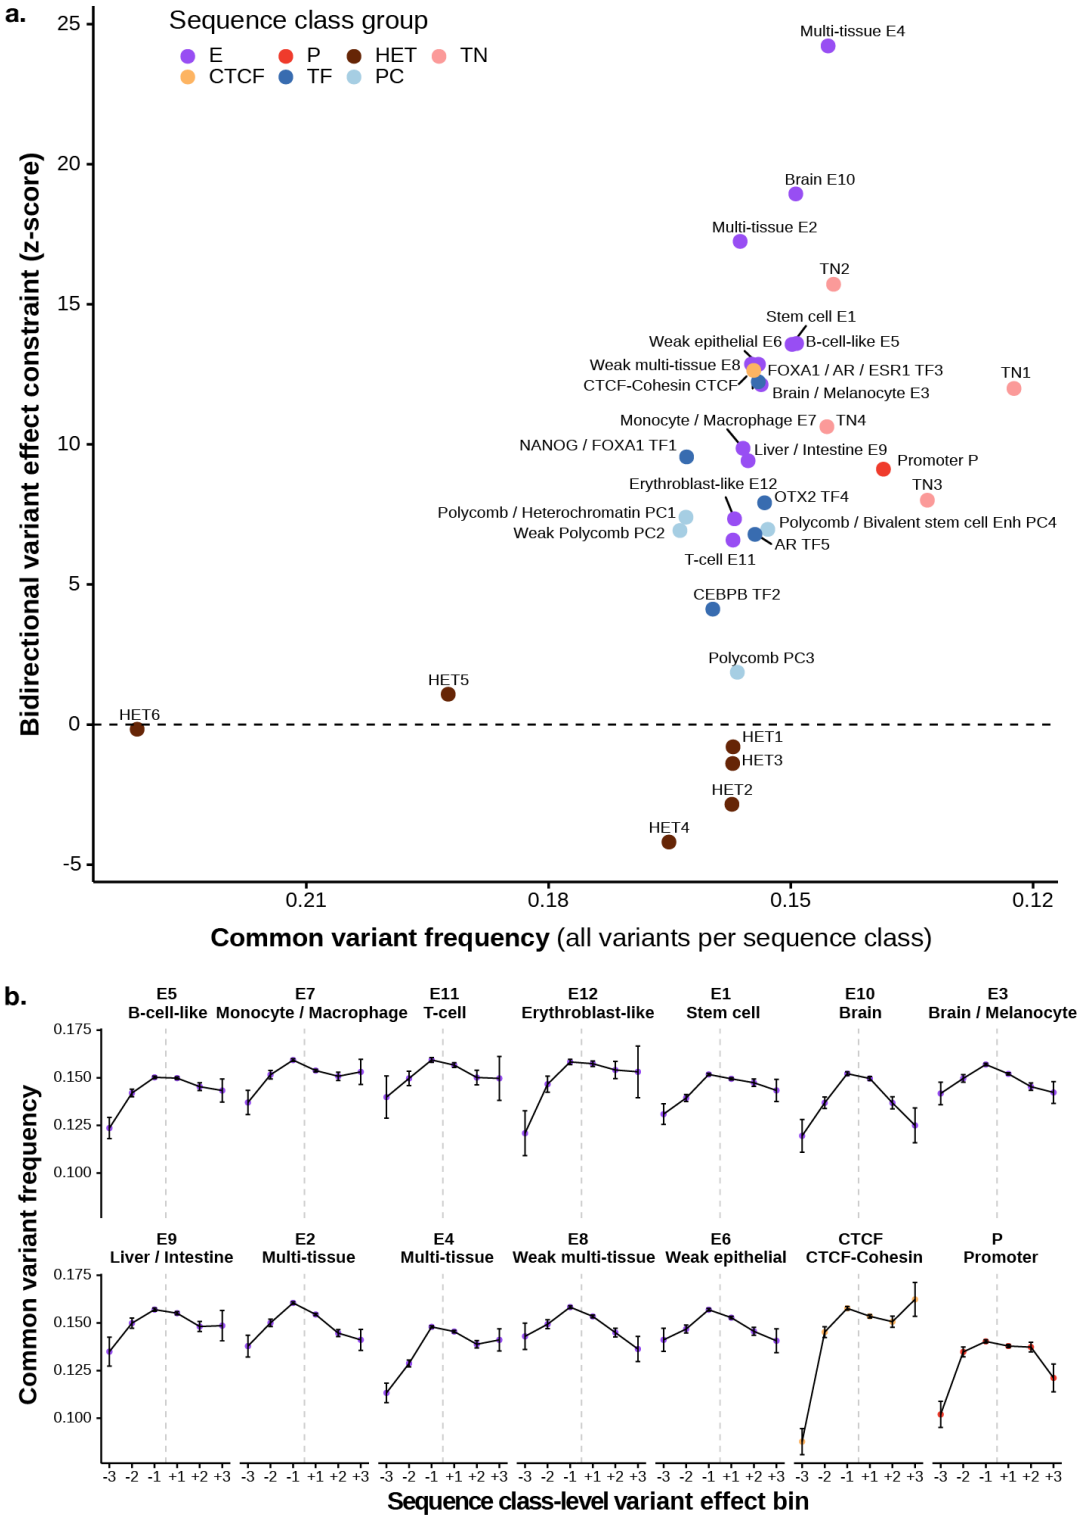

Figure 3. Variants with strong regulatory sequence class effects show negative selection signatures. a, Scatter plot for allele-frequency-based analysis of each sequence class. The x-axis shows the average variant allele frequency, and the y-axis shows the bidirectional variant effect constraint z-score, which is computed based on logistic regressions predicting common variant (allele frequency > 0.01) from sequence-class-level variant effect score for both positive and negative effects (Methods). 'L' sequence classes are excluded due to lack of interpretation for their sequence-class-level variant effect scores. **b,** Comparison of common variant frequencies for 1000 Genomes variants assigned to different sequence classes and variant effect bins. The common variant threshold is >0.01 allele frequency across the 1000 Genomes population. Error bars show +/- 1 standard error (SE). The sequence-class-level variant effects are assigned to 6 bins (+3: top 1% positive, +2: top 1%-10% positive, +1: top 10% -100% positive, -1: top 10% -100% negative, -2: top 1%-10% negative, -3: top 1% negative).

Sequence classes elucidate the tissue-specific regulatory architecture of GWAS traits

The population allele frequency analysis on sequence classes suggest that variants perturbing regulatory sequence class activities are likely involved in human health and disease. Therefore, to explore this hypothesis, we used GWAS data to delineate the genetic contribution of each sequence class to diseases and traits.

Partitioned heritability from LD score regression (LDSR) has been a powerful tool for understanding the genetic architecture of diseases and traits using GWAS summary statistics[25]. Previous applications of LDSR use overlapping annotations[25–27] and therefore cannot unambiguously partition heritability across annotations. Because sequence classes are both non-overlapping and cover nearly the entire genome, they provide a clear and more easily interpretable picture of the regulatory architecture of diseases and traits. To show this, we estimated the proportion of heritability explained by each sequence class for 47 GWAS traits in UK Biobank (UKBB)[16,28] (Methods). Specifically, we applied LDSR and used a conservative estimate of the proportion of heritability, subtracting one standard error and lower-bounding by 0. Our analysis of UKBB GWAS revealed genetic signatures of sequence-class-specific regulatory functions (Figure 4, Supplementary File 3).

Remarkably, 'E' and 'P' sequence classes cover almost all classes that explain a high proportion of heritability for GWAS traits and diseases--the same sequence classes inferred to be under strong evolutionary constraints (Figure 3a, Supplementary File 3). We observed three main groups of traits that share similar heritability composition signatures across sequence classes. The first group is blood-related traits, which contains two subgroups of immune-related and non-immune-related traits. The majority of heritability signals in blood-related traits are explained by enhancer classes for the relevant cell type(s), such as monocyte/macrophage enhancer (E7) for *Monocyte Count*, B-cell-like enhancer (E5) for *Auto Immune Traits*, and erythroblast-like (red blood cell progenitor) enhancer (E12) for *Red Blood Cell Distribution Width,* which measures the range of variation in red blood cell volume. Furthermore, autoimmune-related traits are

386 selectively associated with the immune cell type enhancer sequence classes E5 (B-cell like), E11
387 (T-cell), and E7 (monocyte/macrophage), while erythroblast-like enhancer E12 is specifically
388 linked to red-blood-cell-related traits. Therefore, sequence classes can dissect the cell-type-
389 specific regulatory architecture of traits and diseases with heritability decomposition, even
390 without relying on gene-level information.
391
392 Cognitive and mental traits (*Morning Person*, *Neuroticism*, *Smoking Status*, *Years of Education*,
393 *College Education*) have similar sequence-class-level heritability decompositions as well; for
394 this second group of traits, heritability was mostly explained by brain enhancer (E10 and E3) and
395 stem cell enhancer (E1) sequence classes. The link to E1 is consistent with our observation that
396 E1 was also moderately enriched for active enhancer mark H3K4me1 in brain cell types (Figure
397 2a, Supplementary Figure 5) and is positively correlated with gene expression in brain tissues
398 (Figure 2b).
399
400 The third group of traits is intriguingly diverse, including *Balding*, *Lung Forced Vital Capacity*,
401 *Waist-hip Ratio*, *Height*, and *Heel T-score.* The heritability of these traits are mostly explained
402 by multi-tissue enhancer classes (E4, E2, and E8), which show activity in epithelial cells,
403 fibroblast, muscle, and many other cell types. Enhancer activity across multiple tissues in the
404 body may explain the diverse phenotypes that are associated with these traits.
405
406 Beyond these three groups, there are a number of traits with unique heritability patterns that are
407 also linked to highly relevant sequence classes. For example, the *High Cholesterol* trait was most
408 associated with the liver and intestine enhancer sequence class (E9), which is consistent with the
409 physiology of cholesterol metabolism and known etiology of this condition[29]. E9 was also linked
410 to red-blood-cell-related traits, in line with the role of liver in erythropoiesis.
411
412 Finally, the promoter sequence class P uniquely explained a sizable proportion of heritability in
413 nearly all traits, suggesting a near-universal involvement of promoter sequence variations in all
414 traits and diseases.

Figure 4. Sequence-class-based partitioning of GWAS heritability shows trait associations with tissue-specific regulation. Partitioned genome-wide heritability in UKBB GWAS with all 40 sequence classes. The size of the dot indicates the proportion of heritability estimated from LDSR, which is conservatively estimated as one standard error below the estimated heritability proportion (bounded by 0). The color of the dot indicates the significance z-score of the fold enrichment of heritability relative to the proportion of all SNPs assigned to the sequence class (bounded by 0). Colored boxes indicate traits associated with blood (red), brain (green), multiple tissues (blue) and promoters (orange).

We next assessed whether our new sequence classes could explain GWAS heritability beyond that explained by annotations discovered in prior studies. To this end, we performed LDSR analysis with our whole genome annotations of sequence classes conditioned on an up-to-date set of previously identified baseline annotations (v2.2, https://alkesgroup.broadinstitute.org/LDSCORE/). We uncovered 83 significant sequence-class-

trait associations with a corrected p-value cutoff of <0.05 (Supplementary File 4). 70% of all UKBB GWAS traits and most of the E and P sequence classes (9/13) have at least one significant association after multiple hypothesis testing correction (Supplementary File 4). This finding suggests that sequence classes can identify extensive new regulatory signals that enrich GWAS interpretation.

Disease mutations are predicted to disrupt the activities of sequence classes

Sequence-class-level effects enable the prediction of specific regulatory mechanisms at the individual, pathogenic mutation level. To showcase our framework's capability to predict the mechanisms of individual mutations, we used Sei to predict the direction and magnitude of sequence-class-level mutation effects for all 853 regulatory disease mutations from the Human Gene Mutation Database (HGMD)[30]. For systematic classification and quantification of these mutations, we assign each mutation to an affected sequence class based on its mutation effects (the sequence class with the strongest score change) and the sequence that it alters (Methods).

Overall, the average variant effect score of disease mutations is 4.2x larger than the de novo mutations in healthy individuals (0.903 vs 0.217, $p<2.2e-16$, Wilcoxon rank-sum test two-sided, max absolute effect across sequence classes) and 6.5x larger than the 1000 Genomes common variants with AF>0.01 (0.903 vs 0.139, $p<2.2e-16$). Here we focus on analyzing the mutations with the strongest predicted effects (>1.1, n=138/853), where predicted effect refers to the variant effect of the assigned sequence class for each mutation (Figure 5, Supplementary Figure 14). Because sequence-class-level variant effects are directional--that is, predicting whether the alternative allele increases or decreases sequence-class level activity--we are able to discover that while the majority (~80%) of pathogenic mutations with strong predicted effects are predicted to decrease sequence class activity, the remaining 20% of HGMD pathogenic mutations are predicted to increase sequence class activity. Moreover, perturbations to E-, P-, and CTCF-classes make up >99% of the mutations with strong predicted effects on sequence class activity (Supplementary File 5): 44.9% are predicted to affect tissue-specific E sequence classes, 38.4% are predicted to affect the P promoter sequence class, and interestingly, 15.9% are predicted to affect the CTCF-cohesin sequence class (Methods).

We found that almost all mutations with strong predicted effects in cell-type-specific E sequence classes contributed to diseases relevant to that same cell type (Figure 5, Supplementary File 5)--for most of these mutations, the nearby gene is known to be relevant to the disease but the molecular mechanisms of regulatory disruption is unknown. For example, mutations causing Protein C deficiency and Hemophilia B, two diseases characterized by the deficiency of specific plasma proteins produced in the liver (protein C and coagulation factor IX, respectively), are predicted to decrease E9 liver/intestine sequence class activities. Blood cell-type-specific enhancer sequence classes are disrupted in distinct blood-related diseases and deficiencies relevant to the corresponding cell type: the E12 erythroblast-like enhancer sequence class is

disrupted in red blood cell-specific diseases such as pyruvate kinase deficiency, erythropoietic porphyria, delta-thalassemia, and beta-thalassemia; the E7 monocyte/macrophage-like sequence class is disrupted in monocyte and macrophage-related chronic granulomatous disease; and the E5 B-cell-like enhancer sequence class is disrupted in X-linked agammaglobulinemia, a functional deficiency of B-cell. For developmental diseases, such as preaxial polydactyly triphalangeal thumb and radial ray deficiency and triphalangeal thumb-polysyndactyly syndrome, the E1 embryonic stem cell-specific enhancer sequence class is predicted to be disrupted by mutations in a known distal enhancer of Sonic Hedgehog (SHH) (chr7:156583951 G>A[31], chr7:156583949 G>C[32]), a gene that plays a crucial role in the positioning and growth of limbs, fingers, and toes during development.

In addition, 38% of the regulatory mutations with strong predicted effects affect the activity of the promoter sequence class P, including a hypercholesterolemia mutation near the LDLR gene (chr19:11200089 C>T[33]), a microcephaly & developmental delay mutation near the PIGY gene (chr4:89444948 C>T[34]), and a retinoblastoma mutation near the RB1 gene (chr13:48877851 G>T[35]). The high proportion of mutations perturbing the P sequence class likely reflects both the critical role of promoters in diseases and the emphasis on promoter-proximal mutations in past studies.

While the mutations we've discussed thus far are negative effect mutations which decrease sequence class activity, 20% of HGMD pathogenic mutations are predicted to increase sequence class activity. Indeed, these mutations included many known gain-of-function mutations, which validated our predictions. The top positive effect was observed for a mutation (chrX:73072592 G>C) near the XIST gene that skews X-inactivation of the mutant chromosome in females[36]; this mutation was predicted to increase the activity of the CTCF sequence class and has been experimentally validated to increase CTCF binding[37]. Positive effect predictions for 'E' and 'P' sequence classes were also validated by previously studied mutations: an alpha-thalassemia mutation near the HBM gene (chr16:209709 T>C[38]) known to create a GATA1 binding site and increase intergenic transcription was predicted to increase the activity of the erythroblast-specific E12 sequence class, and a TERT gene mutation found in individuals with familial melanoma (chr5:1295161 T>G[39]) was predicted to increase the activity of P. Beyond this, many mutations predicted to have strong positive effects were not previously understood. For example, a mutation near the HBG1 gene (chr11:5271262 A>G[40]) that causes persistence of fetal hemoglobin is also predicted to increase the activity of the erythroblast-specific E12 sequence class. Previously, this mutation was known to create an ATGCAAAT octamer[40] that matches the POU family transcription factor motif, but its functional consequences were unclear.

Notably, even though pathogenic mutations from prior genetics studies are subjected to selection bias, the observation of pathogenic mutations with strong impacts on E-, P-, and CTCF-sequence classes are consistent with our estimation that regulatory sequence classes are under

strong evolutionary constraints, and strong disruptions to these classes are likely to cause health consequences. We also note that the pathogenic mutations with strong positive effects on the CTCF class do not contradict the population allele frequency analysis which inferred that further increase of activity on the CTCF sequence class is generally tolerated, because all these pathogenic mutations are located in sequences in other sequence classes, but their sequence class identities are altered to the CTCF class by these mutations. This is in contrast to the allele frequency analysis, which only focuses on variants located in sequences in the CTCF sequence class.

Therefore, sequence-class-level effects both corroborate existing regulatory mechanisms and propose new mechanisms for individual pathogenic mutations. We expect our framework to be a valuable tool in accelerating genetic discoveries of disease-causal mutations and their mechanisms in the regulatory genome.

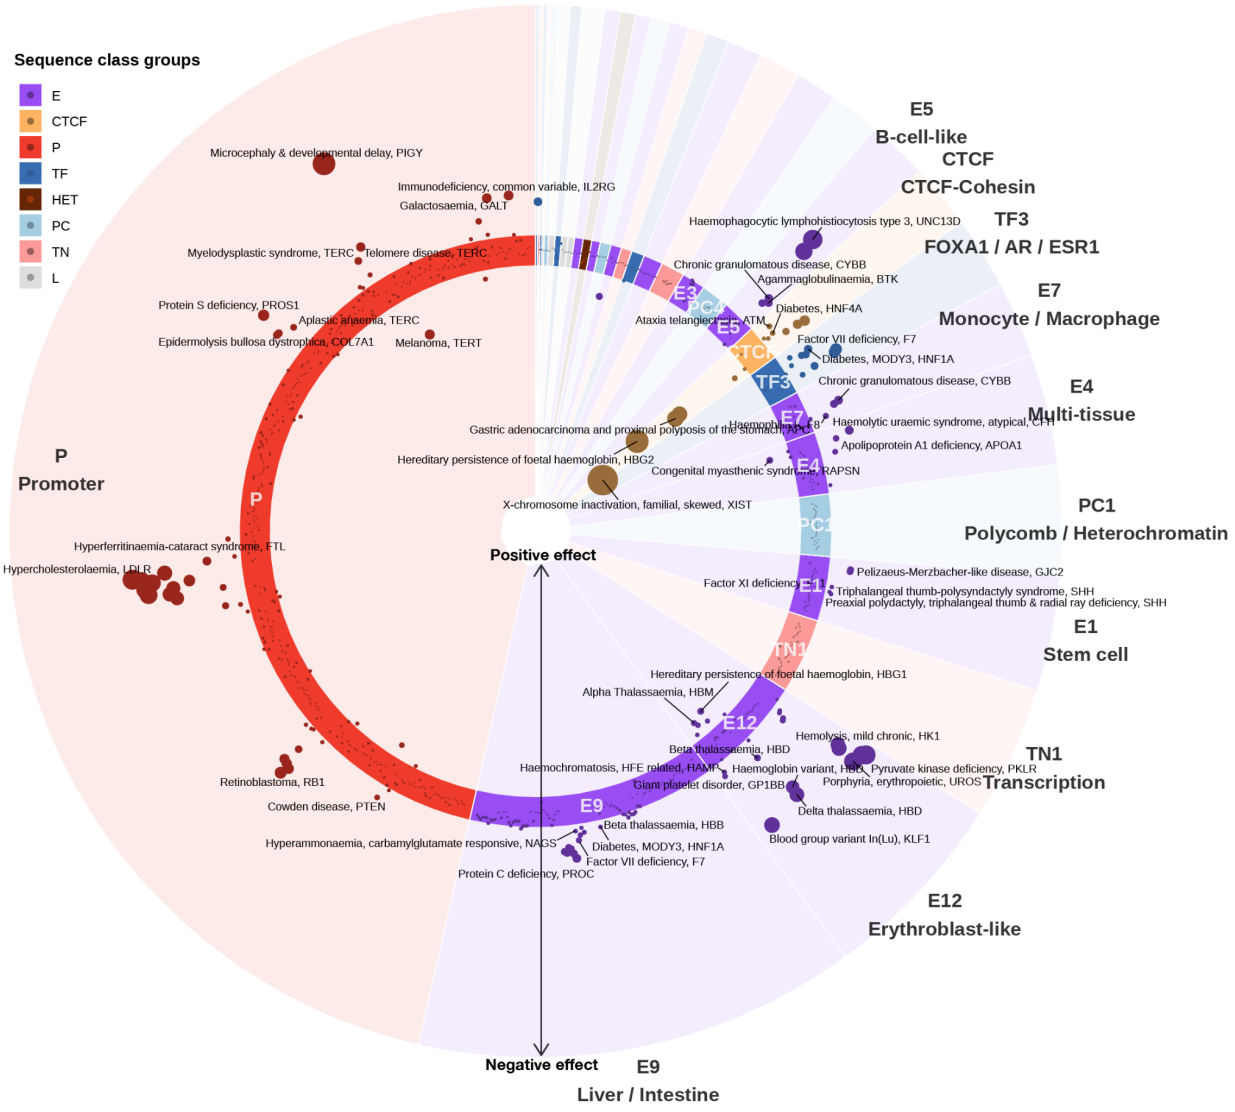

Figure 5. Disease regulatory mutations are predicted to disrupt promoter, CTCF, and tissue-specific enhancer sequence classes. Sequence-class-level mutation effects of pathogenic noncoding HGMD mutations are plotted. A polar coordinate system is used, where the radial coordinate indicates the sequence-class-level effects. Each dot represents a mutation, and mutations inside the circle are predicted to have positive effects (increased activity of sequence class), while mutations outside of the circle are predicted to have negative effects (decreased activity of sequence class). Dot size indicates the absolute value of the effect. Mutations are assigned to sequence classes based on their sequences and predicted effects (Methods). Within each sequence class, mutations are ordered by chromosomal coordinates. The associated disease and gene name are annotated for each mutation, and only the strongest mutation is annotated if there are multiple mutations associated with the same disease, gene, and sequence class.

Discussion

We developed a genome-wide sequence-based map of regulatory activities using sequence classes, a vocabulary for genomic sequence activities discovered using a data-driven, systematic method. Our deep-learning-based framework uses a compendium covering 21,907 publicly available cis-regulatory profiles and the whole genome sequence to create a mapping from any sequence to a comprehensive set of sequence classes. This provides a global sequence-based view of sequence regulatory activities and allows for the quantitative prediction of variant effects on sequence class activities. Sequence classes are a concise vocabulary of regulatory activities that is interpretable, quantifiable, and easily analyzed globally (across all sequence classes) and individually. To our knowledge, it is the first such attempt to systematically map regulatory activities from any sequence.

We demonstrated that E- and P- sequence classes are strongly enriched in trait and disease GWAS heritability and under evolutionary constraints. Importantly, sequence classes provide insights into the mechanisms of individual pathogenic mutations by predicting effects on the function of tissue-specific enhancers, promoter activity, and long-range genome interactions (e.g. CTCF-cohesin sequence class). Using sequence-class-level variant effect predictions, we linked many pathogenic mutations to tissue-specific regulatory changes in the relevant tissues. These predictions point to potential mechanisms that can be experimentally tested in the future.

Sequence classes leverage a sequence model trained on most publicly available cis-regulatory profile data; however, there remains substantial space for improvement as more data becomes available. For example, we are still lacking data for many cell types, developmental stages, transcription factors, and combinations of chromatin targets measured in new cell types or conditions. More data will likely enable the characterization of still more cell-type-specific and developmental stage-specific sequence classes, defining sequence classes with increasingly fine-grained regulatory resolution.

This work demonstrates the potential of sequence classes to discover regulatory disruptions in

human diseases, through both the aggregation of genome-wide variant association signals and prediction of the impact of individual mutations. We provide sequence classes and the Sei model as a resource for further research into understanding the regulatory genetic landscape of human health and diseases. Our framework is applicable to any variant, regardless of whether it is common, rare, or never previously observed, and we expect it to be a powerful tool for understanding the mechanistic effects of noncoding mutations in human health.

Methods

Training data

21,907 cis-regulatory profiles in peak format were compiled from the processed files of the Cistrome[4], ENCODE[2], and Roadmap Epigenomics projects[3]. The Cistrome Project, which systematically processed publicly available cis-regulatory profiles, contributed the majority of the profiles predicted in Sei (19,905). We excluded profiles from Cistrome with less than 1000 peaks. Genome sequences are from the GRCh38/hg38 human reference genome. The full list of cis-regulatory profiles is available in Supplementary File 1.

Deep learning sequence model training

The Sei model is trained to predict 21,907 transcription factor binding, histone marks, and DNA accessibility from cis-regulatory profile peaks at the center of 4kb sequences.

The model architecture is composed of three sequential sections: 1) a convolutional network with dual linear and nonlinear paths, 2) residual dilated convolution layers, 3) spatial basis function transformation and output layers. A detailed specification of the model is available in Supplementary File 6 and in the code repository (https://github.com/FunctionLab/sei-framework, downloadable from https://doi.org/10.5281/zenodo.4906996). In the convolutional architecture, we introduced a new design composed of both linear and nonlinear convolution blocks. The nonlinear blocks are composed of convolution layers, batch normalization layers and rectified linear activation functions (ReLU), similar to regular convolutional networks. The linear blocks have the same structure as the nonlinear blocks but do not include activation functions to facilitate learning of linear dependencies. Each nonlinear block is stacked on top of a linear block with a residual connection adding the input of the nonlinear block to the output, allowing the computation to go through either the linear or nonlinear path. Dilated convolutional layers with residual connections further expands the receptive fields without reducing spatial resolution. Finally, spatial basis functions are used to reduce dimensionality of the spatial dimension while preserving the capability to discriminate spatial patterns of sequence representations. Specifically, in the Sei model, a B-spline basis matrix (256x16) with 16 degrees-of-freedom across 256 uniformly-spaced spatial bins is generated and multiplied with the convolutional layers output to reduce the 256 spatial dimensions to 16 spline basis function dimensions. After the spline basis function transformation, a fully-connected layer and an output layer are used for

integrating information across the whole sequence and generating the final 21,907-dimensional predictions.

Training, validation, and testing datasets are specified by different sets of chromosomes in the hg38 genome (holding out chromosome 8 and 9 for the test set and chromosome 10 for the validation set), and samples drawn uniformly across the hg38 genome for these partitions, excluding regions specified in the ENCODE blacklist[41]. For training, we sampled training sequences and their labels on-the-fly from the training set of chromosomes using Selene[42]. As a result, almost all training samples are drawn from unique genomic intervals with distinct start and end positions to reduce overfitting during the training process. For each 4kb region, a 21,907-dimensional binary label vector is created for the 21,907 cis-regulatory profiles based on whether the center basepair overlaps with a peak in each of the profiles. The model is implemented in PyTorch and trained with Selene. A detailed training configuration file is available at https://github.com/FunctionLab/sei-framework/blob/main/train/train.yml.

Model performance

We computed the AUROC and AUPRC for all cis-regulatory profiles predicted by Sei on the test holdout dataset, excluding profiles that had fewer than 25 positive samples in the test set. Additionally, to assess the correlation structure of the predictions, we compared the rank-transformed pairwise Spearman's rank correlations for the predicted cis-regulatory profiles to the pairwise correlations for the true labels (peak calls provided in Cistrome DB).

The model performance comparison between DeepSEA and Sei is computed on the 2,002 cis-regulatory profiles from Roadmap and ENCODE that both DeepSEA and Sei predict. Because both models have the same chromosomal test holdout (chr8 and chr9), we use the regions specified in the DeepSEA test holdout set to create a common test dataset of sequences and labels on which to evaluate the models.

Sequence classes

We selected 30 million genomic positions that uniformly tile the genome with 100bp step size and then computed Sei predictions for 4kb sequences centered at each position. Sequences overlapping with ENCODE blacklist regions[41] or assembly gaps ("N"s) are removed. To process the 30 million x 21,907 predictions matrix, the dimensionality is first reduced with principal component analysis (PCA). The PCA transformations were fitted with incremental PCA using a batch size of 1,000,000 for one pass of the whole dataset, and genomic positions were randomly assigned to batches. The top 180 principal components, scaled to unit variance, were used for constructing a nearest neighbor graph where each node is connected to its k-nearest neighbors by Euclidean distance (k=14). Louvain community clustering with default parameters was applied to the nearest neighbor graph with the python-louvain package, which resulted in 61 clusters. We refer to the largest 40 clusters as sequence classes and exclude the remaining (smallest) 21

clusters, which constitute <2.6% of the genome, from our analyses due to their size. We refer to this cluster assignment to sequence classes at 100bp resolution as sequence class annotations. We visualized the genome-wide predictions by computing UMAP embedding with a subsample of PCA-transformed Sei predictions of 30 million sequences, and then fine-tuned the visualization with OpenTSNE. The detailed procedures will be available in our code repository.

Sequence class scores

Each sequence class is represented as a unit vector in the 21,907-dimensional cis-regulatory profile space, in the direction of the average prediction of all sequences assigned to this sequence class among the 30 million. In more formal notation, the vector for sequence class i is $v_i = \frac{\overline{p_{s \in Sequence\ class\ i}}}{||p_{s \in Sequence\ class\ i}||_2}$, where p_s represents the 21,907-dimensional Sei prediction for sequence s. Each Sei prediction can then be projected onto any sequence class vector to obtain a sequence class-level representation of the prediction, which we call sequence class score or$score_{s,i} = p_s \cdot {v_i}^T$. In addition, predicted sequence-class-level variant effects are represented by the difference between the sequence class scores of the sequences carrying the reference allele and the alternative allele, or $score_{v,i} = score_{alt,i} - score_{ref,i}$. To better represent predicted variant effects on histone marks, it is necessary to normalize for the nucleosome occupancy (e.g. loss-of-function mutation near TSS can decrease H3K4me3 modification level while increasing nucleosome occupancy, resulting in an overall increase in observed H3K4me3 quantity). Therefore, for variant effect computation, we use the sum of all histone profile predictions as an approximation to nucleosome occupancy and adjust all histone mark predictions to remove the impact of nucleosome occupancy change (non-histone mark predictions are unchanged):

$$p^{hm*} = {p^{hm}}_{ref} \frac{\sum_k {p^{hm^k}}_{ref} + \sum_k {p^{hm^k}}_{alt}}{\sum_k {p^{hm^k}}_{ref}}; p^{hm*} = {p^{hm}}_{alt} \frac{\sum_k {p^{hm^k}}_{ref} + \sum_k {p^{hm^k}}_{alt}}{\sum_k {p^{hm^k}}_{alt}}$$

where $\sum_k {p^{hm^k}}_{ref}$ represents the sum over all histone mark predictions (among 21907-dimensions of a prediction) for the reference allele. We generally exclude Low Signal sequence classes in sequence-class-level variant effect analyses because they lack an intuitive biological interpretation.

Sequence class enrichment of histone mark annotations, transcription factors, and repeats

We computed the log fold change enrichment of various chromatin profiles and genome annotations for each sequence class based on sequence class annotations (described above, see 'Sequence classes'). Log fold change enrichment is computed by taking the log ratio of the proportion of a sequence class intersecting with the annotation versus the background proportion of the annotation, where we consider all regions assigned to any sequence class. We computed enrichment for all 21,907 profiles predicted by Sei, filtered the chromatin profiles for each sequence class to only those having Benjamini-Hochberg corrected p-values (Fisher's exact test, two-sided) below 2.2e-16, and selected the top 25 profiles based on log fold change enrichment. Cistrome Project profile enrichment is computed over 2 million random genomic positions.

683
684 The annotation of centromere repeats is obtained from the UCSC RepeatMasker track, and
685 annotations of histone marks over multiple cell types are obtained from the Roadmap
686 Epigenomics project--enrichments for both of these sets of annotations are computed over the
687 entire genome.
688
689 **Enhancer sequence class correlations with cell-type-specific gene expression**
690 Tissue expression profiles are from GTEx[23], Roadmap Epigenomics[3], and ENCODE[2] and
691 transformed to log-scale RPKM (reads per kilobase per million reads mapped) scores as
692 previously described[12] and normalized by tissue-average. Specifically, a pseudocount was added
693 before log transformation (0.0001 for GTEx tissues, which are averaged across individuals, and
694 0.01 for Roadmap and ENCODE tissues). After log transformation, the average scores across
695 tissues were subtracted for each gene; as a result, the processed scores represent log fold change
696 relative to tissue-average.
697
698 Gene-wide expression prediction is evaluated on sequence class annotations (from Louvain
699 community clustering) for positions within +/-10kb of the TSSs for these genes. For each
700 enhancer sequence class and tissue, we compute the Spearman correlation between the sequence
701 class annotation coverage and gene expression.
702
703 **Correlation between regulatory sequence class variant effects and directional eQTL**
704 **variant effect sizes**
705 We collected the eQTLs within +/-5kb of gene TSSs from GTEx v8, combined across all GTEx
706 tissues, and computed the Spearman correlation between the top 15k variant effect predictions
707 for each sequence class and the eQTL variant effect sizes (averaged across multiple tissues if the
708 variant is an eQTL in multiple tissues). The p-values are derived from the Spearman's rank
709 correlation test (two-sided) and BH correction is applied. Low Signal and Heterochromatin
710 sequence classes are excluded from this analysis due to lack of interpretation for their variant
711 effect scores in this context.
712
713 **Evolutionary constraints on variant effects**
714 We computed sequence-class-level variant effects for all 1000 Genomes project phase 3
715 variants[24]. Variants are assigned to sequence classes based on the 100bp resolution genome-wide
716 assignment derived from Louvain community clustering as described above. For each sequence
717 class we divide variants into 6 bins based on their effects in the same sequence class as
718 illustrated in Figure 3, and summarize common variant (AF>0.01) frequencies in each bin by
719 mean and standard error of the mean. We also estimated statistical significance of allele
720 frequency dependency on sequence-class-level variant effects. For each sequence class, we
721 applied logistic regression separately for positive effect and negative effect variants, to predict
722 common variants (AF>0.01) from the absolute value of sequence-class-level variant effect score,

723 and obtained the significance z-score of the regression coefficient of variant effect. The
724 bidirectional evolutionary constraint z-score is defined as the negative value of the combined z-
725 scores from positive and negative effect variants with Stouffer's method.
726
727 **Partitioning GWAS heritability by sequence classes**
728 UKBB GWAS summary statistics were obtained from [16]. To study the association of sequence
729 class genome annotation and sequence class variant effects and trait heritability, we performed
730 partitioned heritability LD score regression (LDSR) as described in [25]. To partition the
731 heritability as sums of heritability explained by each sequence class, we run LDSR with only
732 sequence class annotations and a baseline all-ones annotation. We obtained the estimated
733 proportion of h^2 explained by each sequence class and its standard error with LDSR as
734 implemented in https://github.com/bulik/ldsc. As the estimated proportions can have high
735 variance or even be negative, we use a conservative estimator of the estimated proportion of
736 h^2 subtracted by 1 standard error, lower-bounded by 0. To assess the contribution of sequence
737 classes to explaining additional heritability when conditioned on known baseline annotations, we
738 also run LDSR with the baseline annotations (v2.2,
739 https://alkesgroup.broadinstitute.org/LDSCORE/). The p-values are derived from the coefficient
740 z-score, and BH correction is applied.
741
742 **Sequence class-level variant effect analysis of noncoding pathogenic mutations**
743 We obtained all mutations assigned "DM" and "regulatory" annotation in the Human Gene
744 Mutation Database (HGMD) database (2019.1 release). RMRP gene mutations are excluded
745 because they are likely pathogenic due to impacting RNA function instead of regulatory
746 perturbations, despite being annotated to the regulatory category in HGMD. For every mutation,
747 we predicted the sequence class scores for both the reference and the alternative allele and
748 computed the sequence-class-level variant effect as the predicted scores for the alternative allele
749 subtracting the scores for the reference allele. To provide an overview of sequence-class level
750 effects of human noncoding pathogenic mutations, mutations are first assigned to sequence
751 classes based on the sequence class annotations of the mutation position. For mutations with a
752 strong effect in a different sequence class than the originally assigned sequence class (absolute
753 value higher than the original sequence class by >1 absolute difference and >2.5 fold relative
754 difference), we reassign the mutation to the sequence class with the strongest effects.
755
756 **Code and data availability**
757 The Sei framework code is provided in https://github.com/FunctionLab/sei-framework, and the
758 model and associated data files downloadable by following the instructions in the GitHub
759 repository. Code and data for the manuscript results are available at
760 https://github.com/FunctionLab/sei-manuscript.
761
762 **Acknowledgements**

The authors acknowledge all members of the Troyanskaya lab for helpful discussions. This work was performed using the high-performance computing resources, supported by the Scientific Computing Core, at the Flatiron Institute and the Terascale Infrastructure for Groundbreaking Research in Science and Engineering high-performance computer center at Princeton University. K.M.C. is supported by the National Science Foundation Graduate Research Fellowship Program (NSF-GRFP). O.G.T. is supported by National Institutes of Health grant nos. R01HG005998, U54HL117798 and R01GM071966, U.S. Department of Health and Human Services grant no. HHSN272201000054C and Simons Foundation grant no. 395506. O.G.T. is a senior fellow of the Genetic Networks program of the Canadian Institute for Advanced Research. J.Z. is supported by the Cancer Prevention and Research Institute of Texas grant (RR190071) and the UT Southwestern Endowed Scholars program.

Author Contributions

K.M.C. and J.Z. conceived the Sei framework, developed the computational methods, and performed the analyses. A.K.W. developed the Sei web server. K.M.C., J.Z., and O.G.T. wrote the manuscript.

Supplementary Figures

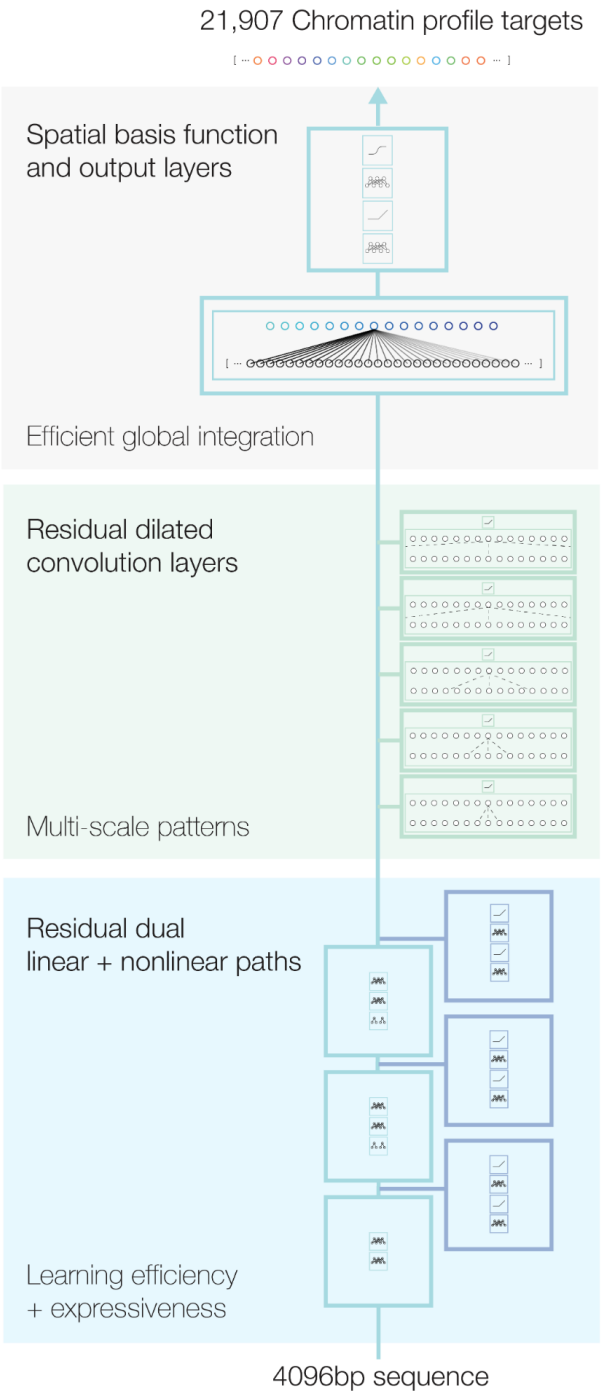

Supplementary Figure 1. Schematic overview of Sei model architecture. 4096bp sequences, one-hot encoded, are the input to the model (bottom) and the predicted 21,907 cis-regulatory profiles are the output (top).

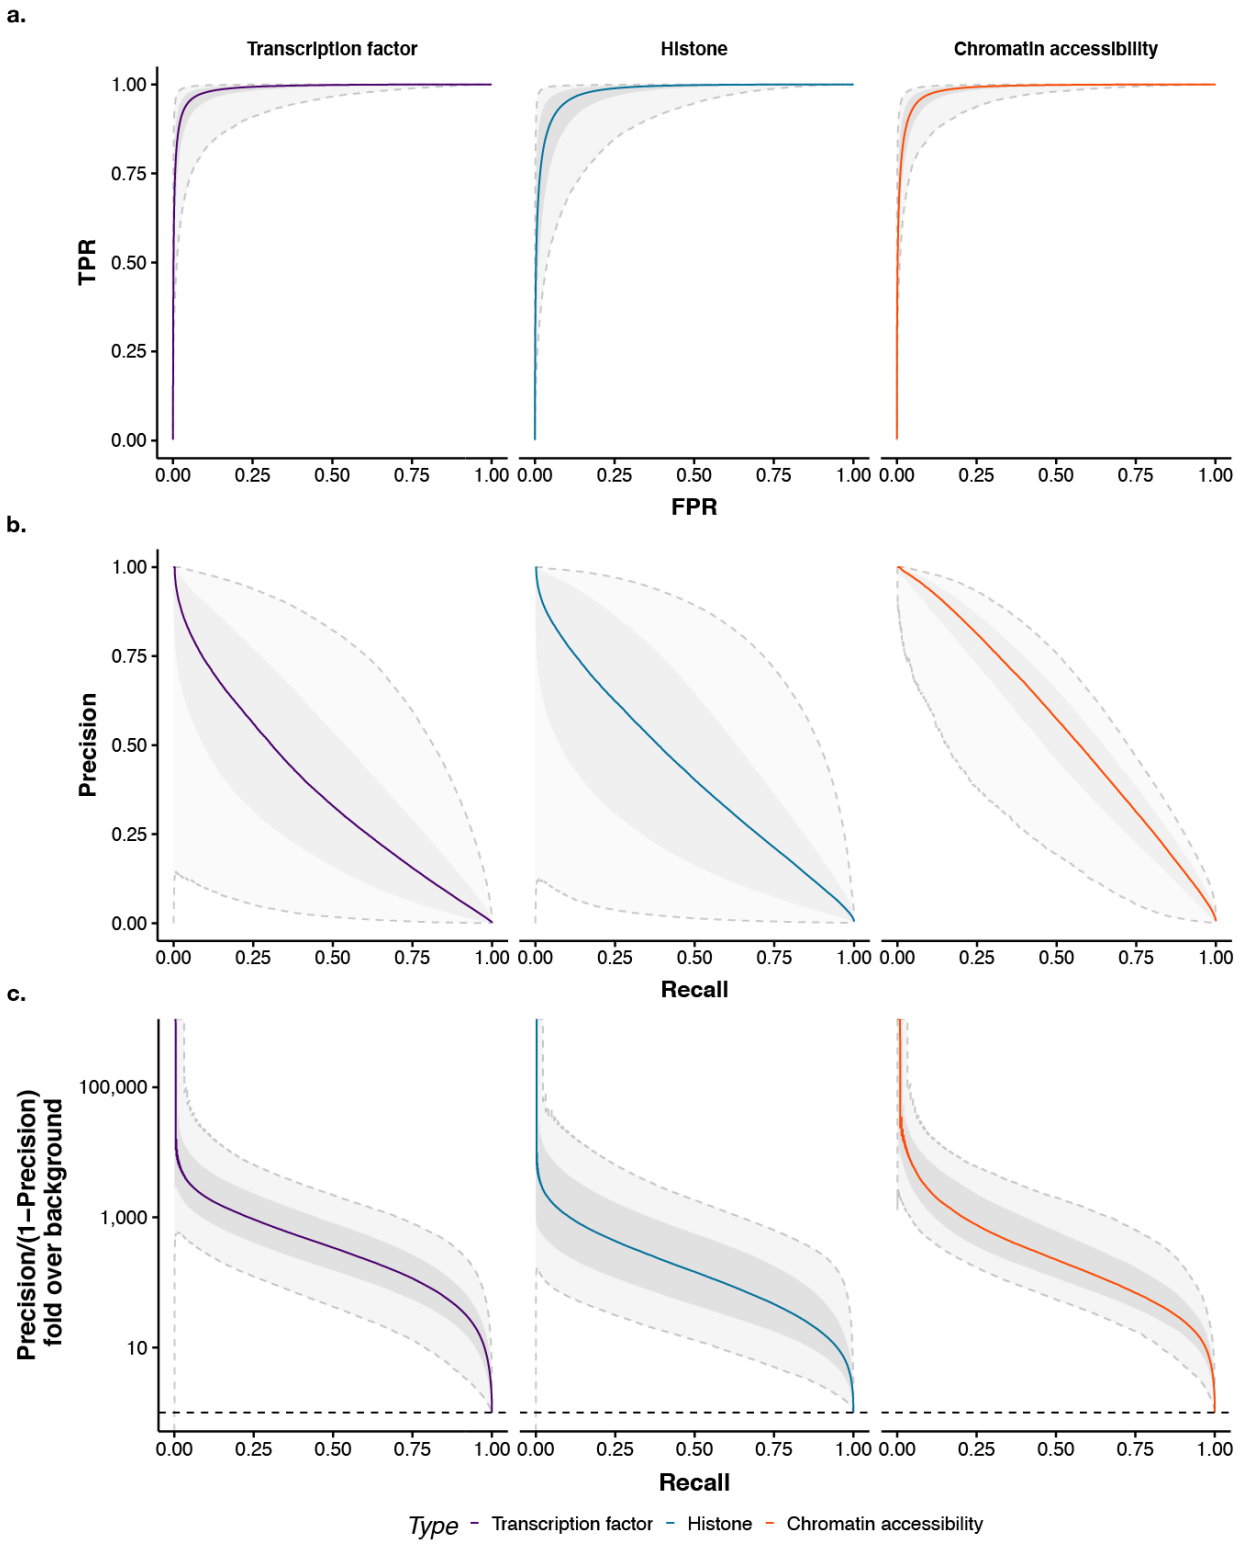

785
786 **Supplementary Figure 2. Sei model performance on predicting 21907 cis-regulatory profiles on**
787 **holdout chromosomes.**
788

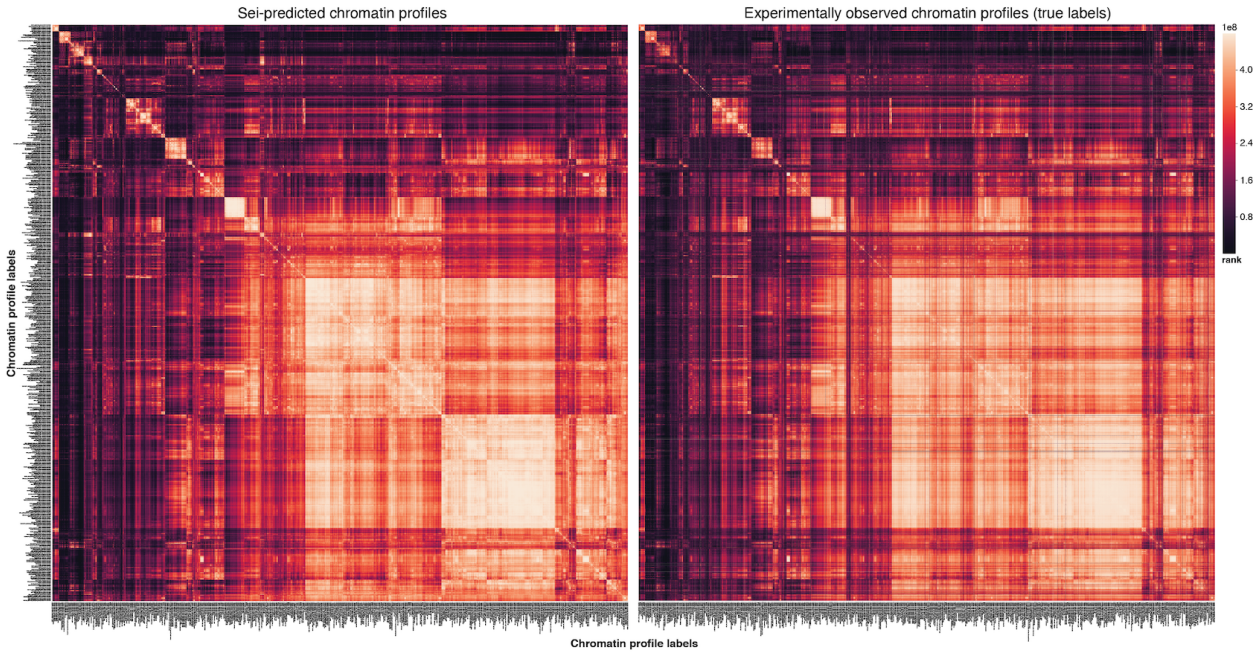

Supplementary Figure 3. Visualizing the rank-transform of pairwise Spearman correlations for the 21,907 cis-regulatory profiles in Sei. Sei model predictions share a highly similar correlation structure with the experimental observations.

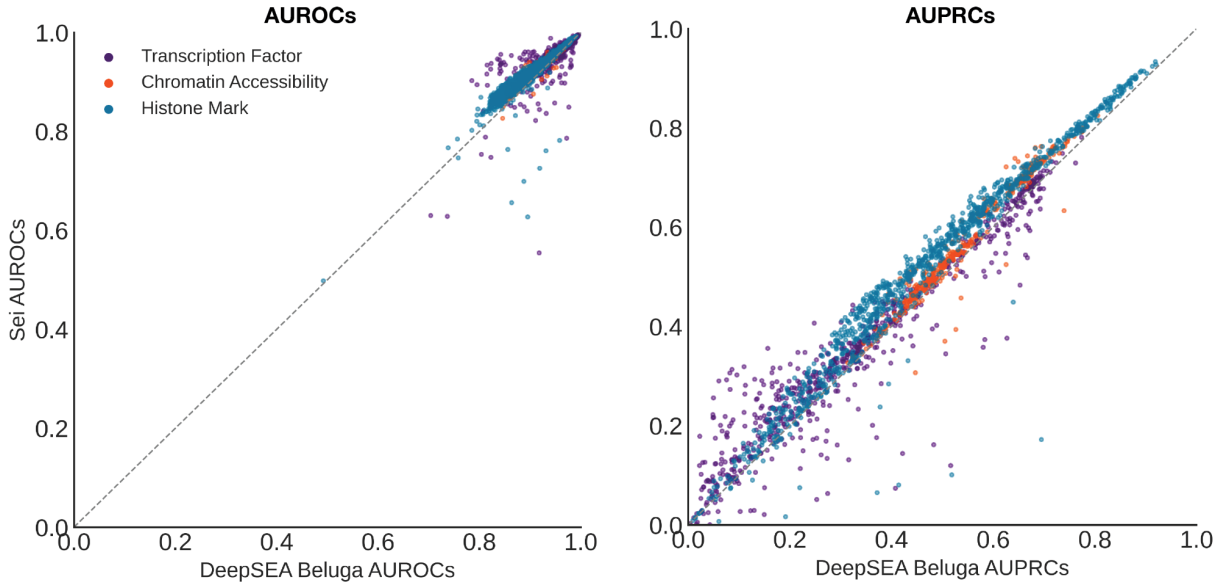

Supplementary Figure 4. Sei model performance comparison with DeepSEA. Performance on the shared 2002 DeepSEA "Beluga" (2018) cis-regulatory profiles are compared.

Supplementary Figure 5. Enrichment of tissue/cell type-specific H3K4me3 (promoter mark) profiles in sequence classes. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap.

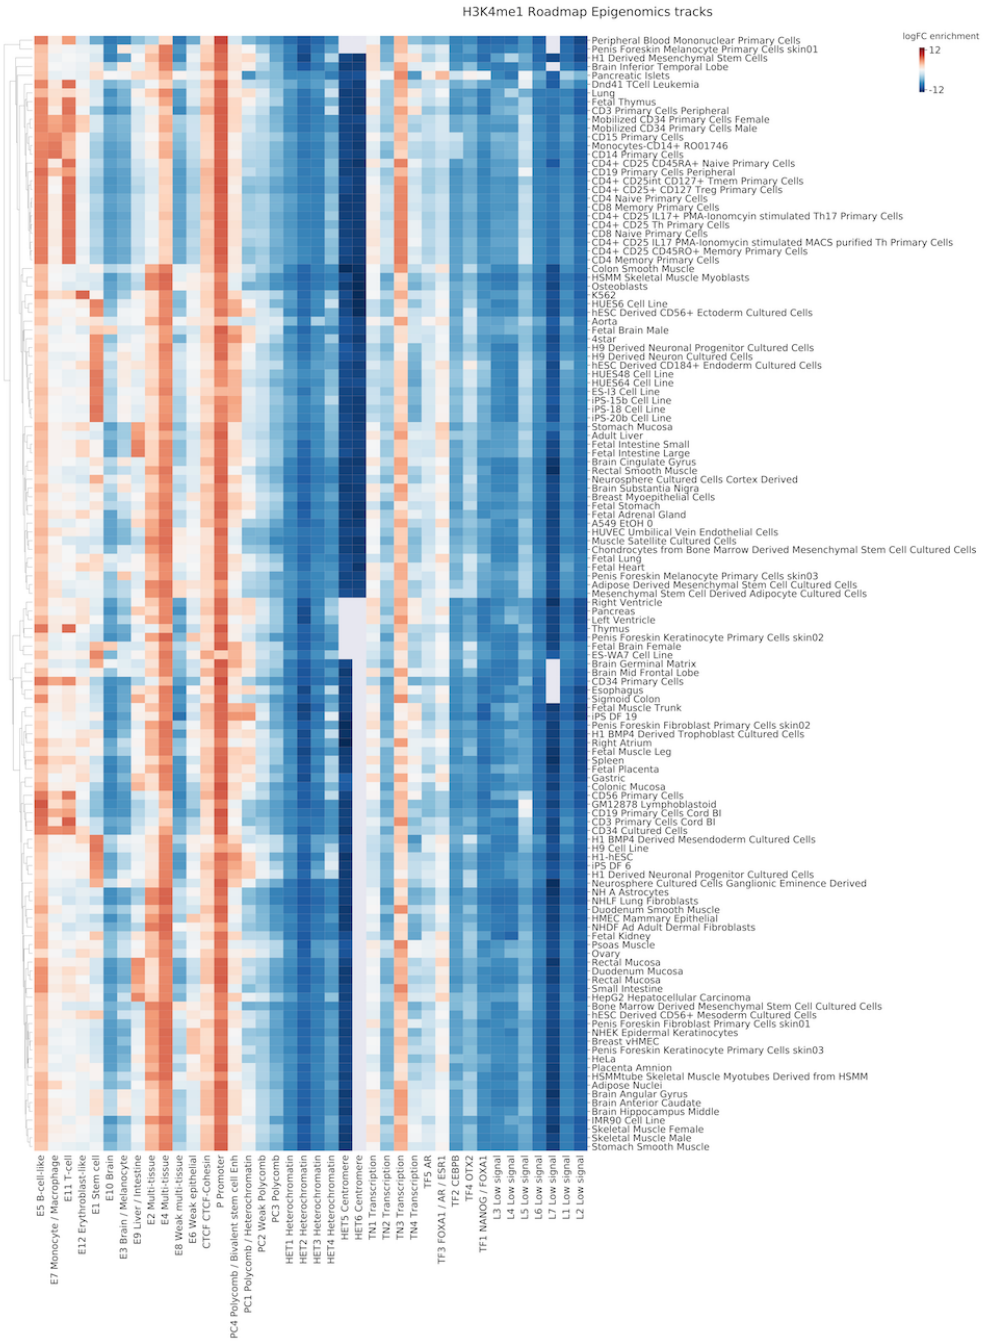

Supplementary Figure 6. Enrichment of tissue/cell type-specific H3K4me1 (enhancer mark) profiles in sequence classes. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap.

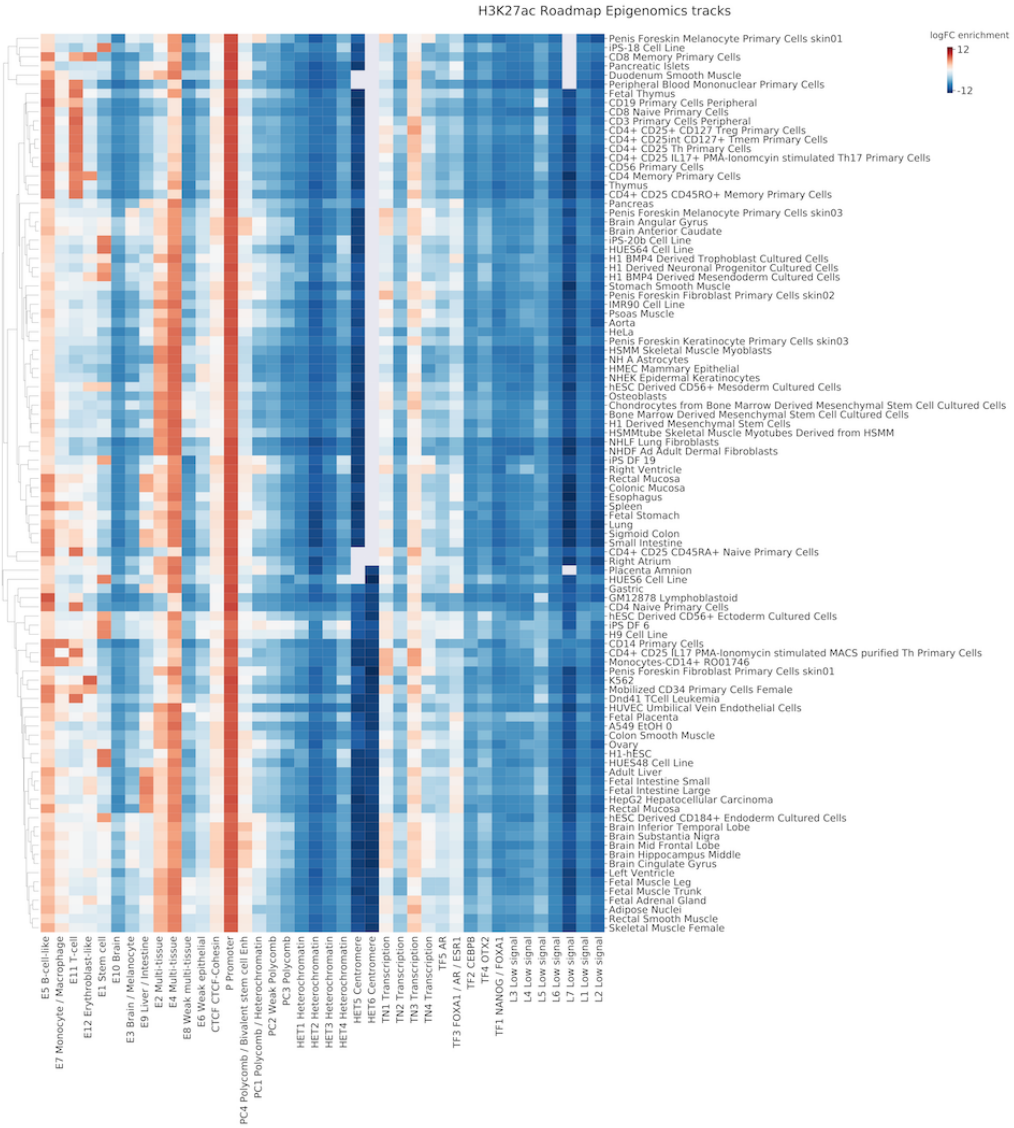

Supplementary Figure 7. Enrichment of tissue/cell type-specific H3K27ac (enhancer mark) profiles in sequence classes. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap.

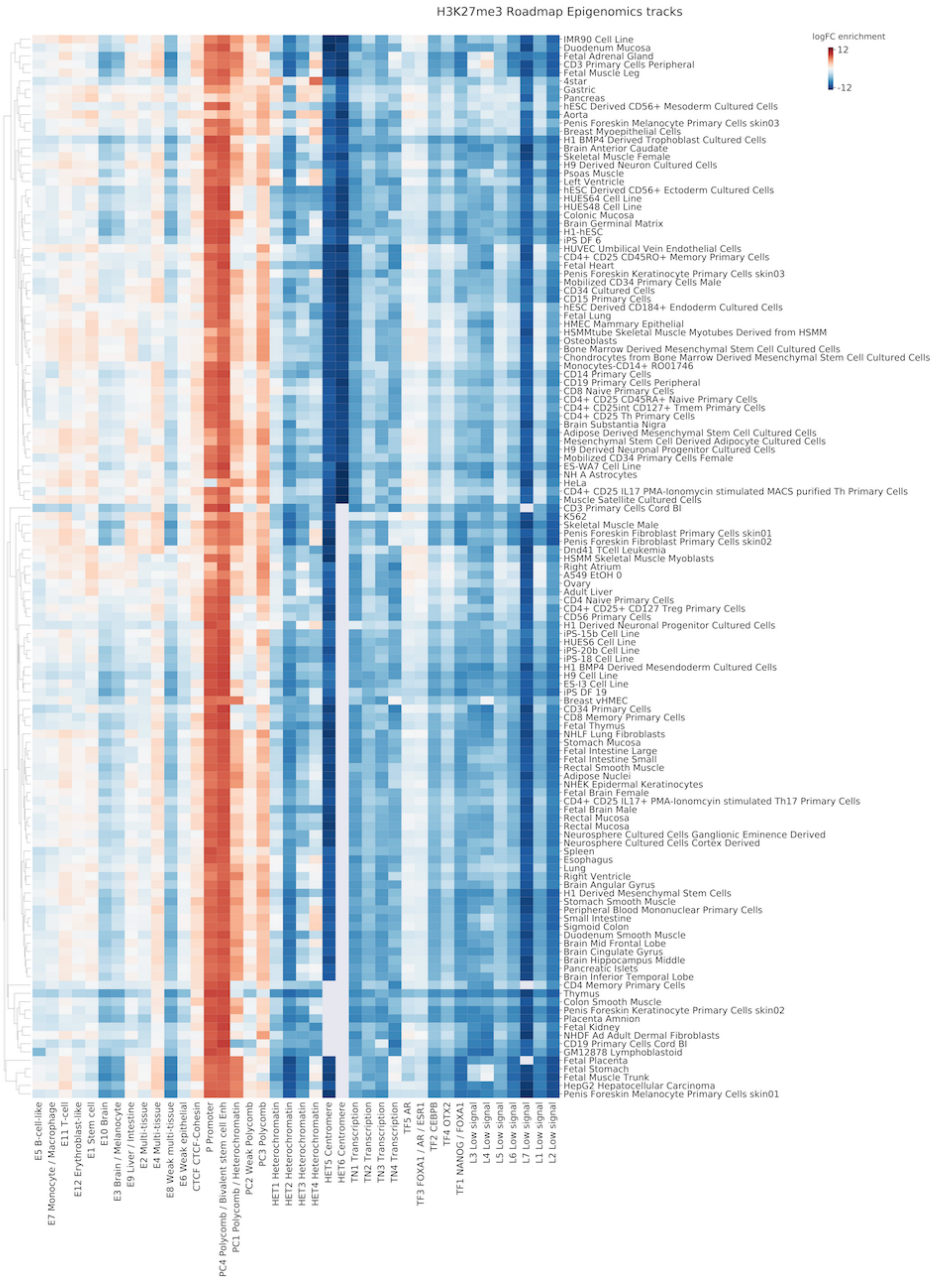

Supplementary Figure 8. Enrichment of tissue/cell type-specific H3K27me3 (Polycomb mark) profiles in sequence classes. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap.

Supplementary Figure 9. Enrichment of tissue/cell type-specific H3K9me3 (heterochromatin mark) profiles in sequence classes. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap.

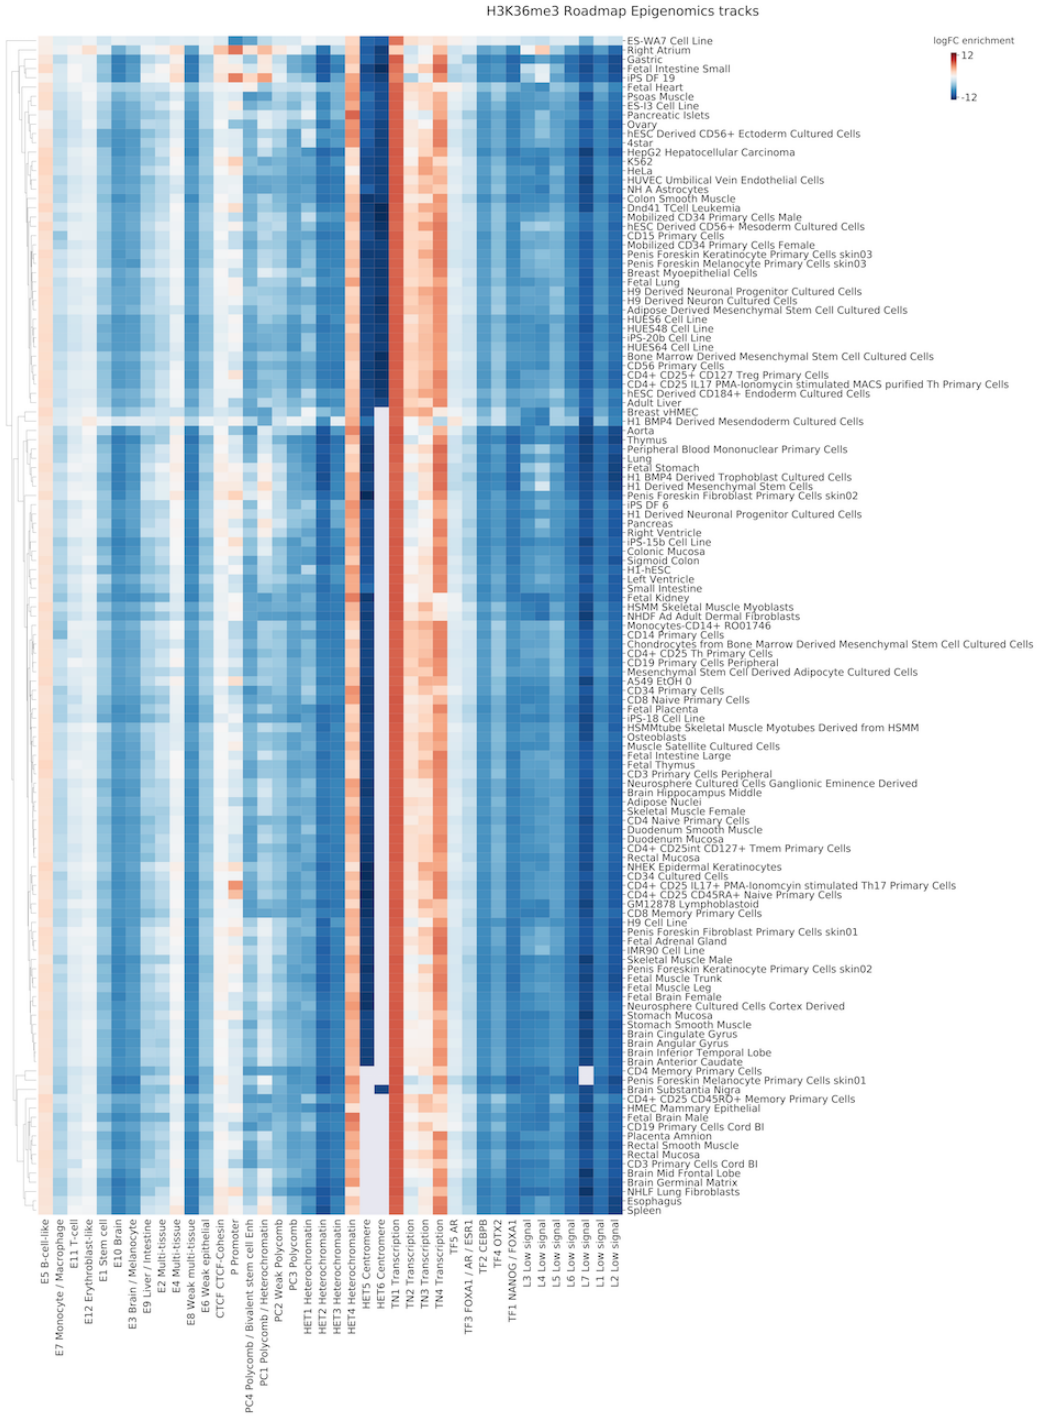

Supplementary Figure 10. Enrichment of tissue/cell type-specific H3K36me3 (transcription mark) profiles in sequence classes. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap.

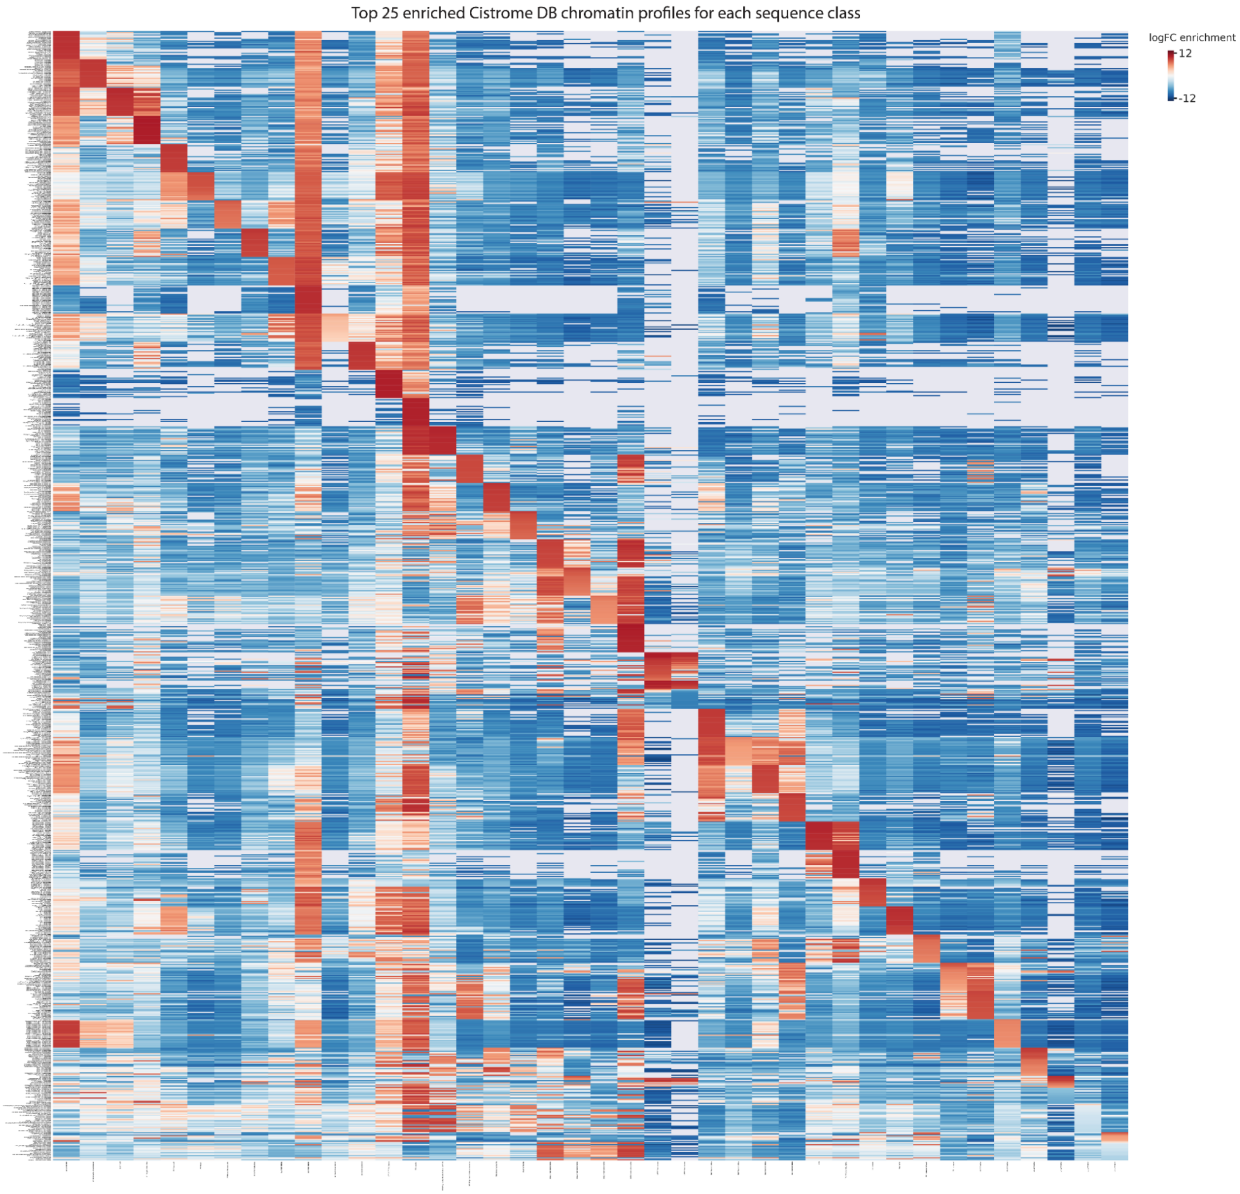

Supplementary Figure 11. Top 25 enriched Cistrome Project chromatin profiles for each sequence class. Log fold change enrichment over genome-average background is shown in the heatmap. No overlap is indicated by the gray color in the heatmap. We computed the enrichment for all 21,907 profiles predicted by Sei over 2 million random genomic positions. For each sequence class, the chromatin profiles are filtered to those having Benjamini-Hochberg corrected p-values (Fisher's exact test, two-sided) < 2.2e-16 selecting the top 25 profiles based on log fold change enrichment.

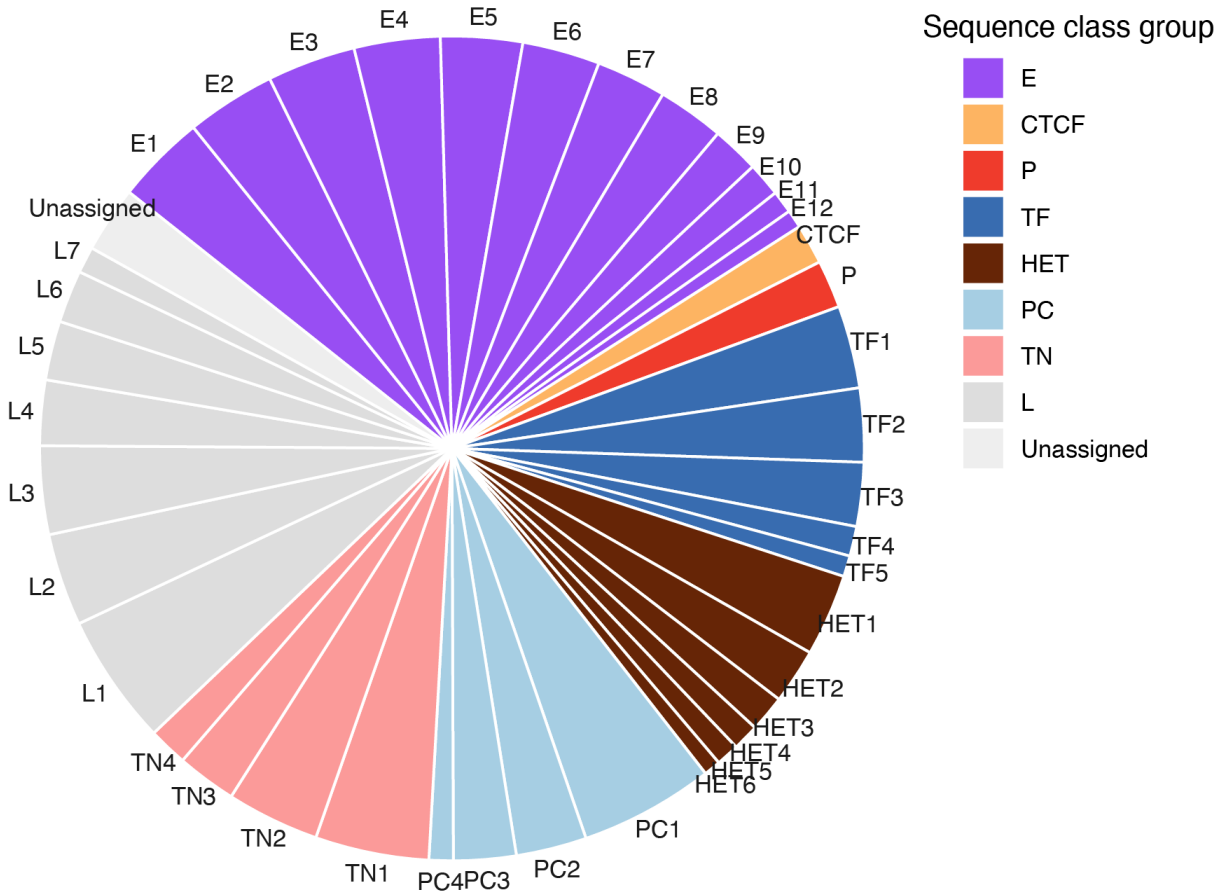

Supplementary Figure 12. Genome sequence proportion covered by each sequence class. The proportion of each sequence class is shown in the pie chart. Genome-wide sequence class assignments were based on Louvain clustering of Sei predictions of sequence tiling the genome with 100bp step size. The clusters unassigned to sequence classes due to the small size (below top 40 clusters) were categorized as "Unassigned".

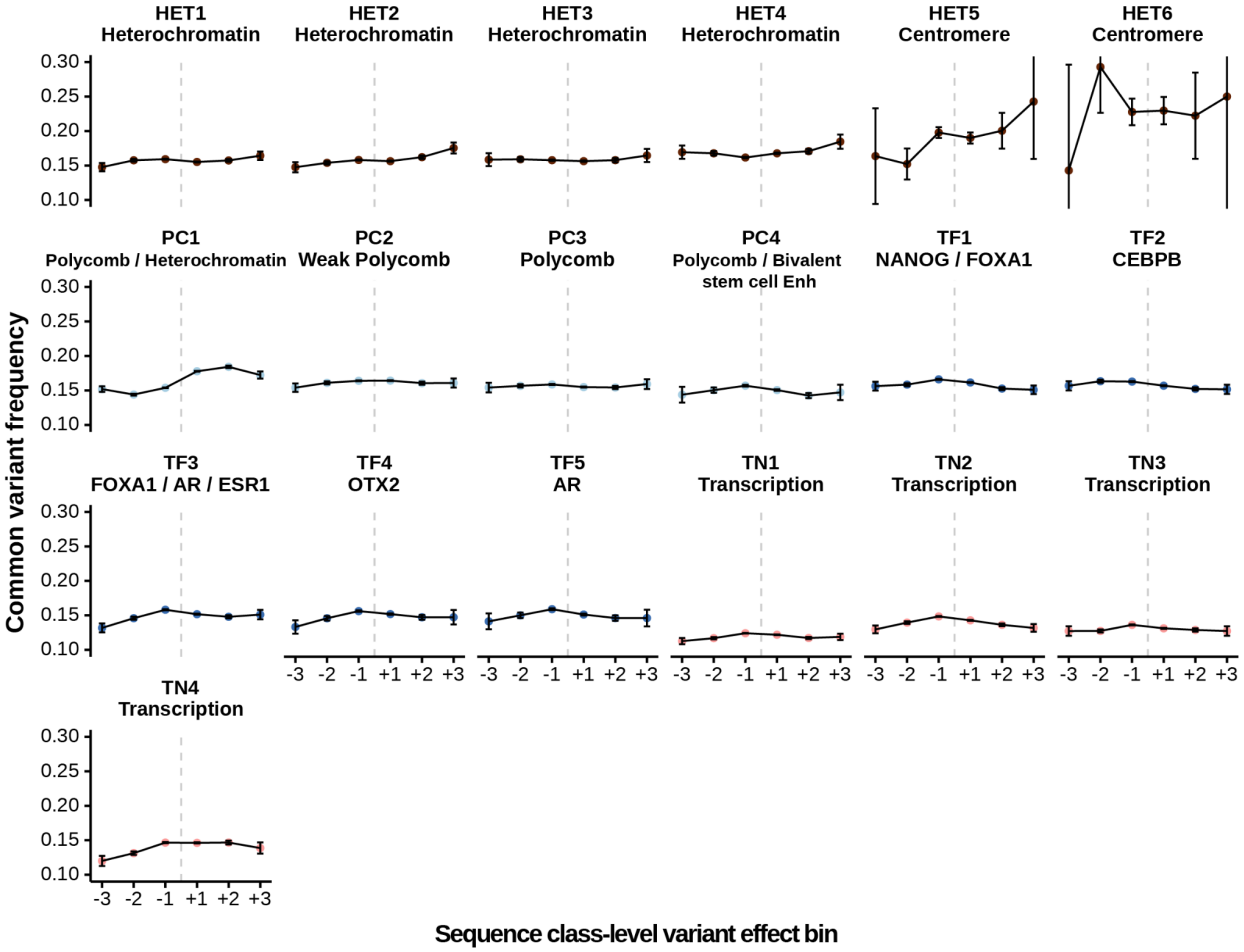

Supplementary Figure 13. Population allele frequency profiles for variants in heterochromatin, low signal, polycomb, and transcription sequence classes. Comparison of common variant frequencies of 1000 Genomes variants assigned to different sequence classes and variant effect bins. The common variant threshold is >0.01 allele frequency across the 1000 Genomes population. Error bars show +/- 1 standard error(SE). The sequence-class-level variant effects are assigned to 6 bins (+3: top 1% positive, +2: top 1%-10% positive, +1, top 10% -100% positive, -3: top 1% negative, -2: top 1%-10% negative, -1, top 10% -100% negative).

Absolute sequence-class-level variant effect score
5
10
Effect direction
Negative
Positive
E5 B cell-like
E7 Macrophage / Macrophage
E11 T cell
E12 Erythroblast-like
E13 Stem cell
E10 Brain
E9 Brain / Astrocyte
E3 Liver / Hepatocyte
E4 Muscle tissue
E8 Weak muscle tissue
E6 Weak epithelial
CTCF CTCF-binding
P Promoter
PC4 Postcortin / Brodmann's Mem cell Exh
PC5 Postcortin / Heterochromatin
PC2 Weak Postcortin
PC1 Postcortin
HE1 Hemostatic
HE2 Hemostatic
HE3 Hemostatic
HE4 Hemostatic
HE5 Cardiomyocytes
TN1 Neuroblastoma
TN2 Neuroblastoma
TN3 Trophoblast
TN4 Trophoblast
TE5 AR
TF3 FOXA1 AR / ESR1
TF2 CEBPB
TF4 OTX2
TF1 NANOG / FOXA1
L3 Low signal
L4 Low signal
L5 Low signal
L6 Low signal
L7 Low signal
L1 Low signal
L2 Low signal
Haemophagocytic lymphohistiocytosis type 3, UNC13D, 17:73839907 C>T
Haemophagocytic lymphohistiocytosis type 3, UNC13D, 17:73839908 G>A
Chronic granulomatous disease, CYBB, X:37639262 A>C
Chronic granulomatous disease, CYBB, X:37639264 T>C
Agammaglobulinaemia, BTK, X:100641212 T>C
Chronic granulomatous disease, CYBB, X:37639266 C>T
Chronic granulomatous disease, CYBB, X:37639267 C>T
Haemophilia A, F8, X:154251045 A>G
Haemophilia A, F8, X:154251048 A>T
Pyruvate kinase deficiency, PKLR, 1:155271258 T>C
Pyruvate kinase deficiency, PKLR, 1:155271259 C>G
Hemolysis, mild chronic, HK1, 10:71075518 A>G
Porphyria, erythropoietic, UROS, 10:127505271 A>T
Porphyria, erythropoietic, UROS, 10:127505271 A>G
Porphyria, erythropoietic, UROS, 10:127505277 C>T
Porphyria, erythropoietic, UROS, 10:127505287 G>T
Porphyria, erythropoietic, UROS, 10:127505287 G>C
Porphyria, erythropoietic, UROS, 10:127505291 G>T
Delta thalassaemia, HBD, 11:5255749 G>T
Delta thalassaemia, HBD, 11:5255781 G>A
Haemoglobin variant, HBD, 11:5255789 T>A
Delta thalassaemia, HBD, 11:5255790 A>G
Beta thalassaemia, HBD, 11:5255793 C>T
Hereditary persistence of foetal haemoglobin, HBG1, 11:5271262 A>G
Hereditary persistence of foetal haemoglobin, HBG1, 11:5271285 A>G
Hereditary persistence of foetal haemoglobin, HBG2, 11:5276188 A>G
Alpha Thalassaemia, HBM, 16:209709 T>C
Blood group variant In(Lu), KLF1, 19:12998078 A>G
Haemolytic anaemia, nonspherocytic, KLF1, 19:12998108 G>A
Haemochromatosis, HFE related, HAMP, 19:35773328 C>T
Giant platelet disorder, GP1BB, 22:19710933 C>G
Pelizaeus-Merzbacher-like disease, GJC2, 1:228337558 A>G
Pelizaeus-Merzbacher-like disease, GJC2, 1:228337561 A>G
Factor XI deficiency, F11, 4:187186995 G>A
Triphalangeal thumb-polysyndactyly syndrome, SHH, 7:156583949 G>C
Preaxial polydactyly, triphalangeal thumb & radial ray deficiency, SHH, 7:156583951 G>A
Protein C deficiency, PROC, 2:128175983 A>G
Protein C deficiency, PROC, 2:128175984 A>G
Protein C deficiency, PROC, 2:128175988 T>A
Protein C deficiency, type I, PROC, 2:128175994 T>G
Protein C deficiency, PROC, 2:128176001 T>C
Protein C deficiency, PROC, 2:128176005 C>T
Beta thalassaemia, HBB, 11:5248388 G>T
Beta thalassaemia, HBB, 11:5248389 G>C
Beta thalassaemia, HBB, 11:5248388 G>A
Beta thalassaemia, HBB, 11:5248389 G>A
Beta thalassaemia, HBB, 11:5248391 G>C
Beta thalassaemia, HBB, 11:5248391 G>A
Diabetes, MODY3, HNF1A, 12:121418289 A>C
Factor VII deficiency, F7, 13:113760091 G>C
Factor VII deficiency, F7, 13:113760095 T>G
Factor VII deficiency, F7, 13:113760096 T>G
Factor VII deficiency, F7, 13:113760097 T>G
Hyperammonaemia, carbamylglutamate responsive, NAGS, 17:42073968 C>A
Haemophilia B, F9, X:138612871 A>G
Haemolytic uraemic syndrome, atypical, CFH, 1:196620941 C>T
Muscular dystrophy with mental retardation, POMT1, 9:134379574 A>G
Delta thalassaemia, HBD, 11:5255778 T>C
Congenital myasthenic syndrome, RAPSN, 11:47470715 G>C
Congenital myasthenic syndrome, RAPSN, 11:47470726 T>C
Apolipoprotein A1 deficiency, APOA1, 11:116708365 T>G
Beta thalassaemia, HBB, 11:5248329 T>G
Dystonia, dopa-responsive, SPR, 2:73114549 G>A
Gastric adenocarcinoma and proximal polyposis of the stomach, APC, 5:112043220 A>C
Gastric adenocarcinoma and proximal polyposis of the stomach, APC, 5:112043223 A>G
Gastric adenocarcinoma and proximal polyposis of the stomach, APC, 5:112043223 A>T
Gastric adenocarcinoma and proximal polyposis of the stomach, APC, 5:112043224 T>C
Gastric adenocarcinoma and proximal polyposis of the stomach, APC, 5:112043225 G>A
Parkinson disease, early-onset, PRKN, 6:163148706 C>G
Melanoma, CDKN2A, 9:21974875 G>T
Hereditary persistence of foetal haemoglobin, HBG2, 11:5276213 G>C
Ataxia telangiectasia, ATM, 11:108093770 A>G
Diabetes, MODY1, HNF4A, 20:42984264 G>A
Diabetes, HNF4A, 20:42984278 C>T
X-chromosome inactivation, familial, skewed, XIST, X:73072592 G>C
Epidermolysis bullosa dystrophica, COL7A1, 3:48632779 G>A
Epidermolysis bullosa dystrophica, COL7A1, 3:48632780 G>A
Protein S deficiency, PROS1, 3:93692761 G>A
Aplastic anaemia, TERC, 3:169482399 C>T
Pulmonary fibrosis, TERC, 3:169482524 C>A
Aplastic anaemia, TERC, 3:169482777 G>C
Aplastic anaemia, TERC, 3:169482782 C>T
Myelodysplastic syndrome, TERC, 3:169482814 G>A
Pulmonary fibrosis, TERC, 3:169482838 C>A
Telomere disease, TERC, 3:169482906 G>C
Microcephaly & developmental delay, PIGY, 4:89444948 C>T
Melanoma, TERT, 5:1295161 T>G
Lipoprotein lipase deficiency, LPL, 8:19796725 T>C
Ataxia, isolated vitamin E deficiency, TTPA, 8:63998581 G>A
Melanoma, CDKN2A, 9:21974847 G>A
Melanoma, CDKN2A, 9:21974860 C>A
Galactosaemia, GALT, 9:34648606 T>G
Thrombocytopaenia 2, ANKRD26, 10:27389389 C>T
Cowden disease, PTEN, 10:89623056 C>T
Cowden disease, PTEN, 10:89623296 G>A
Cowden disease, PTEN, 10:89623365 G>T
Schizophrenia, NRGN, 11:124609907 C>G
Deafness, autosomal recessive 1, GJB2, 13:20767158 G>A
Retinoblastoma, RB1, 13:48877851 G>A
Retinoblastoma, RB1, 13:48877851 G>T
Retinoblastoma, RB1, 13:48877852 G>A
Retinoblastoma, RB1, 13:48877856 T>A
Retinoblastoma, RB1, 13:48877856 T>G
Retinoblastoma, RB1, 13:48877857 G>A
Retinoblastoma, RB1, 13:48877860 G>T
Retinoblastoma, RB1, 13:48878045 C>G
Autism spectrum disorder, GABRB3, 15:27018162 C>A
Protoporphyria, erythropoietic, FECH, 18:55254103 C>G
Hypercholesterolaemia, LDLR, 19:11200019 C>T
Hypercholesterolaemia, LDLR, 19:11200034 C>A
Hypercholesterolaemia, LDLR, 19:11200037 C>T
Hypercholesterolaemia, LDLR, 19:11200083 C>T
Hypercholesterolaemia, LDLR, 19:11200085 C>A
Hypercholesterolaemia, LDLR, 19:11200085 C>G
Hypercholesterolaemia, LDLR, 19:11200085 C>T
Hypercholesterolaemia, LDLR, 19:11200086 C>A
Hypercholesterolaemia, LDLR, 19:11200086 C>G
Hypercholesterolaemia, LDLR, 19:11200087 T>C
Hypercholesterolaemia, LDLR, 19:11200088 C>T
Hypercholesterolaemia, LDLR, 19:11200089 C>G
Hypercholesterolaemia, LDLR, 19:11200089 C>T
Hypercholesterolaemia, LDLR, 19:11200090 C>G
Hypercholesterolaemia, LDLR, 19:11200124 T>C
Hypercholesterolaemia, LDLR, 19:11200202 A>C
Hyperferritinaemia-cataract syndrome, FTL, 19:49468575 C>T
Late foetal loss, PROCR, 20:33759640 T>G
Immunodeficiency, common variable, IL2RG, X:70331494 G>A
Choroideraemia, CHM, X:85302634 G>T
Cleft palate, TBX22, X:79270181 G>A
Pancreatic agenesis, PTF1A, 10:23508365 A>G
Diabetes, MODY3, HNF1A, 12:121418354 T>C
Factor VII deficiency, F7, 13:113760060 C>T
Factor VII deficiency, F7, 13:113760062 C>G
Factor VII deficiency, F7, 13:113760094 C>T
Factor VII deficiency, F7, 13:113760117 A>G
Factor VII deficiency, F7, 13:113760124 A>C
Factor VII deficiency, F7, 13:113760126 A>C
Factor VII deficiency, F7, 13:113760155 C>T

Supplementary Figure 14. Predicted sequence class-level variant effects for HGMD regulatory disease mutations. HGMD regulatory disease mutations with sequence-class level variant effect score >1.1 are included.

Supplementary File 1. The list of 21907 cis-regulatory profiles and Sei prediction performance.

Supplementary File 2. Summary of tissues and cell types covered by the Sei model.

Supplementary File 3. Partitioned UKBB GWAS heritability by sequence classes using LDSR. The proportions of heritability are represented by the LDSR estimate - 1 standard error, lower bounded by 0. The LDSR enrichment z-scores are also lower bounded by 0.

Supplementary File 4. Significant UKBB GWAS trait - sequence class associations identified with LDSR conditioned on the baseline annotations.

Supplementary File 5. Predicted sequence class-level variant effects for HGMD regulatory disease mutations. HGMD regulatory disease mutations with sequence-class level variant effect score >1.1 are included.

Supplementary File 6. Detailed Sei model architecture specification.

References

1. Edwards, S. L., Beesley, J., French, J. D. & Dunning, M. Beyond GWASs: Illuminating the dark road from association to function. *American Journal of Human Genetics* vol. 93 779–797 (2013).
2. Bernstein, B. E. *et al.* An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**, 57–74 (2012).
3. Roadmap Epigenomics Consortium *et al.* Integrative analysis of 111 reference human epigenomes. *Nature* **518**, 317–330 (2015).
4. Zheng, R. *et al.* Cistrome Data Browser: expanded datasets and new tools for gene regulatory analysis. *Nucleic Acids Res.* **47**, D729–D735 (2019).
5. Filion, G. J. *et al.* Systematic protein location mapping reveals five principal chromatin types in Drosophila cells. *Cell* **143**, 212–224 (2010).
6. Ernst, J. *et al.* Mapping and analysis of chromatin state dynamics in nine human cell types. *Nature* **473**, 43–49 (2011).
7. Hoffman, M. M. *et al.* Unsupervised pattern discovery in human chromatin structure through genomic segmentation. *Nat. Methods* **9**, 473–476 (2012).
8. Zhou, J. & Troyanskaya, O. G. Probabilistic modelling of chromatin code landscape reveals functional diversity of enhancer-like chromatin states. *Nat. Commun.* (2016) doi:10.1038/ncomms10528.
9. Boix, C. A., James, B. T., Park, Y. P., Meuleman, W. & Kellis, M. Regulatory genomic

circuitry of human disease loci by integrative epigenomics. *Nature* **590**, 300–307 (2021).
10. Alipanahi, B., Delong, A., Weirauch, M. T. & Frey, B. J. Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. *Nat. Biotechnol.* **33**, 831–838 (2015).
11. Zhou, J. & Troyanskaya, O. G. Predicting effects of noncoding variants with deep learning-based sequence model. *Nat. Methods* **12**, 931–934 (2015).
12. Zhou, J. *et al.* Deep learning sequence-based ab initio prediction of variant effects on expression and disease risk. *Nat. Genet.* (2018) doi:10.1038/s41588-018-0160-6.
13. Kelley, D. R. *et al.* Sequential regulatory activity prediction across chromosomes with convolutional neural networks. *Genome Res.* (2018) doi:10.1101/gr.227819.117.
14. Kelley, D. R. Cross-species regulatory sequence activity prediction. *PLoS Comput. Biol.* **16**, e1008050 (2020).
15. Avsec, Ž. *et al.* Deep learning at base-resolution reveals motif syntax of the cis-regulatory code. *bioRxiv* 737981 (2019) doi:10.1101/737981.
16. Loh, P.-R., Kichaev, G., Gazal, S., Schoech, A. P. & Price, A. L. Mixed-model association for biobank-scale datasets. *Nat. Genet.* **50**, 906–908 (2018).
17. McInnes, L., Healy, J., Saul, N. & Großberger, L. UMAP: Uniform Manifold Approximation and Projection. *Journal of Open Source Software* (2018) doi:10.21105/joss.00861.
18. Poličar, P. G., Stražar, M. & Zupan, B. openTSNE: a modular Python library for t-SNE dimensionality reduction and embedding. *BioRxiv* (2019).
19. Blondel, V. D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *J. Stat. Mech.* **2008**, P10008 (2008).
20. Hawkins, R. D. *et al.* Distinct epigenomic landscapes of pluripotent and lineage-committed human cells. *Cell Stem Cell* **6**, 479–491 (2010).
21. Boros, J., Arnoult, N., Stroobant, V., Collet, J.-F. & Decottignies, A. Polycomb repressive complex 2 and H3K27me3 cooperate with H3K9 methylation to maintain heterochromatin protein 1α at chromatin. *Mol. Cell. Biol.* **34**, 3662–3674 (2014).
22. Schwämmle, V. *et al.* Systems Level Analysis of Histone H3 Post-translational Modifications (PTMs) Reveals Features of PTM Crosstalk in Chromatin Regulation. *Mol. Cell. Proteomics* **15**, 2715–2729 (2016).
23. The GTEx Consortium. The GTEx Consortium atlas of genetic regulatory effects across human tissues. *Science* **369**, 1318–1330 (2020).
24. 1000 Genomes Project Consortium, T. 1000 G. P. *et al.* An integrated map of genetic variation from 1,092 human genomes. *Nature* (2012) doi:10.1038/nature11632.
25. Finucane, H. K. *et al.* Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nat. Genet.* **47**, 1228–1235 (2015).
26. Finucane, H. K. *et al.* Heritability enrichment of specifically expressed genes identifies disease-relevant tissues and cell types. *Nat. Genet.* **50**, 621–629 (2018).
27. Reshef, Y. A. *et al.* Detecting genome-wide directional effects of transcription factor

938 binding on polygenic disease risk. doi:10.1101/204685.
939 28. Sudlow, C. *et al.* UK biobank: an open access resource for identifying the causes of a wide
940 range of complex diseases of middle and old age. *PLoS Med.* **12**, e1001779 (2015).
941 29. Paththinige, C. S., Sirisena, N. D. & Dissanayake, V. H. W. Genetic determinants of
942 inherited susceptibility to hypercholesterolemia – a comprehensive literature review. *Lipids*
943 *Health Dis.* **16**, 103 (2017).
944 30. Stenson, P. D. *et al.* The Human Gene Mutation Database: 2008 update. *Genome Med.* **1**, 13
945 (2009).
946 31. Al-Qattan, M. M., Al Abdulkareem, I., Al Haidan, Y. & Al Balwi, M. A novel mutation in
947 the SHH long-range regulator (ZRS) is associated with preaxial polydactyly, triphalangeal
948 thumb, and severe radial ray deficiency. *Am. J. Med. Genet. A* **158A**, 2610–2615 (2012).
949 32. Gurnett, C. A. *et al.* Two novel point mutations in the long-range SHH enhancer in three
950 families with triphalangeal thumb and preaxial polydactyly. *Am. J. Med. Genet. A* **143A**,
951 27–32 (2007).
952 33. De Castro-Orós, I. *et al.* Functional analysis of LDLR promoter and 5' UTR mutations in
953 subjects with clinical diagnosis of familial hypercholesterolemia. *Hum. Mutat.* **32**, 868–872
954 (2011).
955 34. Ilkovski, B. *et al.* Mutations in PIGY: expanding the phenotype of inherited
956 glycosylphosphatidylinositol deficiencies. *Hum. Mol. Genet.* **24**, 6146–6159 (2015).
957 35. Sakai, T., Ohtani, N., McGee, T. L., Robbins, P. D. & Dryja, T. P. Oncogenic germ-line
958 mutations in Sp1 and ATF sites in the human retinoblastoma gene. *Nature* **353**, 83–86
959 (1991).
960 36. Plenge, R. M. *et al.* A promoter mutation in the XIST gene in two unrelated families with
961 skewed X-chromosome inactivation. *Nat. Genet.* **17**, 353–356 (1997).
962 37. Pugacheva, E. M. *et al.* Familial cases of point mutations in the XIST promoter reveal a
963 correlation between CTCF binding and pre-emptive choices of X chromosome inactivation.
964 *Hum. Mol. Genet.* **14**, 953–965 (2005).
965 38. De Gobbi, M. *et al.* A regulatory SNP causes a human genetic disease by creating a new
966 transcriptional promoter. *Science* (2006) doi:10.1126/science.1126431.
967 39. Horn, S. *et al.* TERT promoter mutations in familial and sporadic melanoma. *Science* **339**,
968 959–961 (2013).
969 40. Surrey, S., Delgrosso, K., Malladi, P. & Schwartz, E. A single-base change at position -175
970 in the 5'-flanking region of the G gamma-globin gene from a black with G gamma-beta+
971 HPFH. *Blood* **71**, 807–810 (1988).
972 41. Amemiya, H. M., Kundaje, A. & Boyle, A. P. The ENCODE Blacklist: Identification of
973 Problematic Regions of the Genome. *Sci. Rep.* **9**, 9354 (2019).
974 42. Chen, K. M., Cofer, E. M., Zhou, J. & Troyanskaya, O. G. Selene: a PyTorch-based deep
975 learning library for sequence data. *Nat. Methods* (2019) doi:10.1038/s41592-019-0360-8.

976