



Published in final edited form as:

Nat Genet. 2015 March ; 47(3): 276–283. doi:10.1038/ng.3196.

Probabilities of Fitness Consequences for Point Mutations Across the Human Genome

Brad Gulko¹, Melissa J. Hubisz², Ilan Gronau^{2,3}, and Adam Siepel^{1,2,3}

¹Graduate Field of Computer Science, Cornell University, Ithaca, NY 14853, USA

²Department of Biological Statistics & Computational Biology, Cornell University, Ithaca, NY 14853, USA

Abstract

We describe a novel computational method for estimating the probability that a point mutation at each position in a genome will influence fitness. These *fitness consequence* (fitCons) scores serve as evolution-based measures of potential genomic function. Our approach is to cluster genomic positions into groups exhibiting distinct “fingerprints” based on high-throughput functional genomic data, then to estimate a probability of fitness consequences for each group from associated patterns of genetic polymorphism and divergence. We have generated fitCons scores for three human cell types based on public data from ENCODE. Compared with conventional conservation scores, fitCons scores show considerably improved prediction power for *cis*-regulatory elements. In addition, fitCons scores indicate that 4.2–7.5% of nucleotides in the human genome have influenced fitness since the human-chimpanzee divergence, and they suggest that recent evolutionary turnover has had limited impact on the functional content of the genome.

During the past decade, two major developments—the emergence of massively parallel, ultra-cheap DNA sequencing technologies and the use of these technologies as digital readouts for functional genomic assays—have led to a profusion of data describing various features of genomes, epigenomes, and transcriptomes^{1,2}. However, investigators still have

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

Correspondence and requests for materials should be addressed to A.S. (asiepel@cshl.edu) and I.G. (ilan.gronau@idc.ac.il).

³Present addresses:

Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, USA (A.S.); The Efi Arazi School for Computer Science, The Interdisciplinary Center, Herzliya, Israel. (A.G.)

URLs

Cold Spring Harbor Laboratory Mirror of UCSC Genome Browser: <http://genome-mirror.cshl.edu/>

UCSC Genome Browser: <http://genome.ucsc.edu>

INSIGHT: <http://compugen.cshl.edu/INSIGHT/>

GENCODE v15: <ftp://ftp.sanger.ac.uk/pub/genocode/release15/>

GERP: <http://mendel.stanford.edu/SidowLab/downloads/gerp/>

CADD: <http://cadd.gs.washington.edu/download/>

RegulomeDB: <http://regulome.stanford.edu/downloads/>

Gerstein Lab ENCODE-nets: <http://encodenets.gersteinlab.org/>

European Bioinformatics Institute's E-GEUV-1 data set: http://www.ebi.ac.uk/arrayexpress/files/E-GEUV-1/analysis_results/

Competing Interests The authors declare that they have no competing financial interests.

Author Contributions I.G. and A.S. conceived the study framework. B.G. and I.G. performed the experiments. All authors analyzed the data. B.G., M.J.H. and I.G. developed analysis tools. B.G., I.G. and A.S. wrote the paper. I.G. and A.S. supervised the research.

only rudimentary tools for integrating these diverse sources of information to obtain useful insights about genomic function and evolution. The limitations of current methods are particularly evident in the vast noncoding regions of eukaryotic genomes, which, despite recent progress^{3–6}, remain poorly annotated and understood. These limitations hamper progress in many areas, ranging from basic molecular genetics to disease association and personalized medicine⁷.

Many computational methods for gaining functional insights from sequence data are based on the simple but profound observation that functionally important nucleotides tend to remain unchanged over evolutionary time, because mutations at these sites generally reduce fitness and are therefore eliminated by natural selection^{7–15}. A major strength of these conservation- or constraint-based approaches is that they sidestep thorny questions about the relationship between the outcomes of biochemical experiments and fitness-influencing functional roles^{16–19} by getting at fitness directly through observations of evolutionary change. In essence, the “experiment” considered by these methods is the one conducted directly on genomes by nature over millennia, and the outcomes of interest are the presence or absence of fixed mutations.

These conservation-based methods, however, depend critically on the assumption that genomic elements are present at orthologous locations and maintain similar functional roles over relatively long evolutionary time periods. Evolutionary turnover may cause inconsistencies between sequence orthology and functional homology that substantially limit this type of analysis. Consequently, investigators have developed two major alternative strategies for the identification and characterization of functional elements. The first strategy is to augment information about interspecies conservation with information about genetic polymorphism^{20–28}. The shorter evolutionary time scales associated with intraspecies variation make this approach more robust to evolutionary turnover and less sensitive to errors in alignment and orthology detection. Polymorphic sites tend to be sparse along the genome, however, so this approach requires some type of pooling of information across genomic positions, which can be problematic in the absence high-quality genomic annotations. The second strategy is to forgo the use of evolutionary information and instead to predict functional roles from genomic data alone, typically with machine-learning methods for supervised classification^{29,30} or clustering followed by labeling based on known examples^{31–33}. This approach has the limitation that it depends strongly on previously characterized elements, which, in noncoding regions, are typically small in number and perhaps unrepresentative of the genome.

In this paper, we introduce a method for genomic analysis that combines many of the strengths of these polymorphism-based and functional genomic approaches. Like functional genomic methods, our approach groups genomic regions according to functional genomic “fingerprints” across multiple assays. Instead of relying on known examples for classification, however, we characterize each group by a probability of mutational fitness consequences—or fitCons score—inferred from patterns of genetic variation. These fitCons scores are estimated using a recently developed statistical method, called Inference of Natural Selection from Interspersed Genomically coHerent elemenTs (INSIGHT), that contrasts patterns of polymorphism and divergence in a collection of dispersed genomic

sites with those in nearby neutrally evolving sites, accounting for negative and positive selection³⁴. Thus, the method integrates both evolutionary and functional data in characterizing the potential functional importance of genomic regions. We demonstrate that these fitCons scores are useful for visualization, for prediction of *cis*-regulatory elements, and for measuring the global influence of recent natural selection across the genome.

Results

General Features of the Prediction Problem

Information about genetic variation can be used to estimate probabilities of fitness consequences for moderately large groups of genomic positions but not for individual loci, owing to the sparsity of informative sites along the genome. This property of “groupwise” but not “individual” predictivity is common to many statistical problems, but it is complicated in our case by two additional features. First, an appropriate scheme for grouping or stratification is not clear *a priori* here, because genomic correlates of fitness consequences are incompletely understood. Second, the outcomes of interest in our problem—the fitness consequences of point mutations—are not directly evident from the data. To highlight these challenges, consider the simpler problem of estimating the expected risk of an automobile accident. This problem must also be addressed at the level of groups (either explicitly through stratification of drivers, or implicitly through regression), but in this case, the relevant features—such as the age, sex, and number of traffic violations of the driver—are mostly plain to the analyst. In addition, the outcomes of interest—the occurrences and costs of accidents—are directly observed. In our problem, the genomic “risk factors” for fitness-influencing mutations, particularly in unannotated noncoding regions of the genome, are much less clear. Furthermore, once a grouping is determined, it is still not possible to read off the associated fitness consequences of mutations; instead they must be inferred from patterns of genetic variation using an evolutionary model.

Calculation of FitCons Scores

We have addressed these challenges using the following strategy. Beginning with genome-wide functional genomic data sets obtained from each cell type (Fig. 1A), we first cluster genomic positions by their joint functional genomic “fingerprints” (Fig. 1B). We focus on three highly informative and largely orthogonal functional genomic data types—DNase-seq data, RNA-seq data, and ChIP-seq data describing histone modifications—which describe DNA accessibility, transcription, and chromatin states, respectively. We divide genomic positions into three levels of DNase-seq signal, four levels of RNA-seq signal, and 26 distinct chromatin states based on the ChromHMM method^{31,33}. In addition, we distinguish between sites that fall outside (0) or within (1) annotated protein-coding sequences (CDSs). We then consider all possible combinations of these four types of assignments, obtaining $3 \times 4 \times 26 \times 2 = 624$ distinct functional genomic classes. We apply this clustering step separately to three karyotypically normal cell types: human umbilical vein epithelial cells (HUVEC), H1 human embryonic stem cells (H1 hESC), and lymphoblastoid cells (GM12878), resulting in 443–447 usable classes of sites, with median numbers of 165 to 224 thousand sites per class (see Supplementary Table 1 and Methods for details).

Next, we use INSIGHT to estimate the probabilities of mutational fitness consequences within each of these classes based on patterns of polymorphism and divergence (Fig. 1C). This step yields an estimate of the fraction of sites under selection (ρ) for each of the analyzed classes, which serves as the fitCons score for that class. Finally, we assign to each nucleotide position in the genome the score estimated for the corresponding functional genomic class (Fig. 1D). Each genomic position is thus assigned a value between 0 and 1, representing the probability that the nucleotide at that position influences fitness, as estimated from patterns of variation at all genomic sites displaying the same functional genomic fingerprint. A vital property of these fitCons scores is that they integrate information from both evolutionary data and cell-type-specific functional genomic data.

Genomic Distribution of FitCons Scores

To obtain a general overview of the genomic distribution of fitCons scores, we first considered the composition and coverage of nucleotide sites of various annotation types as a variable threshold S was applied to the fitCons score, focusing on HUVEC (see Discussion for other cell types). When S is zero, all sites are considered and the composition of annotations reflects the overall genomic distribution (Fig. 2A). As S increases, however, the sites in known functional classes become strongly enriched relative to the intergenic and intronic sites. Regions such as 5' and 3' UTRs, promoters, and introns are most enriched at intermediate scores, reflecting moderate levels of natural selection in these regions, while CDSs dominate at the highest scores. The coverage properties (Fig. 2B) are best for CDSs, 3' UTRs, and 5' UTRs (in that order), but they are also considerably elevated above the intergenic background for promoters, transcription factor binding sites (TFBS), long intergenic noncoding RNAs (lincRNAs), and small noncoding RNAs (snRNAs). Notably, the enrichment for functionally annotated genomic regions at high scores occurs despite no use of genomic annotations in the scoring scheme (except for CDS annotations). Instead, these elevated scores reflect differences in patterns of polymorphism and divergence that arise naturally from the fitness consequences of mutations in these regions and become evident after clustering based on functional genomic data. The fitCons scores for each cell type are displayed genome-wide as tracks in our public mirror of the UCSC Genome Browser (Fig. 3; Supplementary Fig. 1).

The fitCons scores generally depend in expected ways on the marginal signals of functional genomic covariates, but they are also capable of capturing complex, non-additive relationships among covariates. For example, the scores outside of CDSs increase with marginal DNase-seq (Fig. 4A) and RNA-seq (Fig. 4B) signals, as expected; yet a closer examination reveals that the scores actually decrease with DNase-seq intensity in the presence of a high RNA-seq intensity, apparently due to an implicit partitioning of 5' and 3' UTRs (Fig. 4C). This example demonstrates that our exhaustive partitioning scheme allows the method to capture unanticipated relationships between functional genomic covariates and natural selection.

Predictive Power for *Cis*-Regulatory Loci

We evaluated the predictive power of the fitCons scores for known cell-type-specific regulatory elements in comparison with three widely used phylogenetic conservation scoring

original curves (Supplementary Fig. 4), demonstrating that the apparent performance of the fitCons scores was not artificially inflated by the coarse-grained nature of our scores and TFBSs. Second, we examined an alternative set of predicted enhancers for GM12878 cells based on characteristic patterns of divergent transcription initiation³⁹. Unlike the chromatin-based enhancer predictions described above, these predictions are based on data completely independent from those underlying the fitCons scores. Nevertheless, the fitCons scores still displayed excellent predictive power for this set, better than all other methods besides EnhancerFinder (Supplementary Fig. 5).

Proportion of the Human Genome Under Selection

The proportion of nucleotides in the human genome that directly influence fitness—sometimes called the “share under selection” (SUS)—has primarily been estimated using methods that consider divergence patterns among mammals, for which turnover of functional elements may be an important confounding factor^{40–44}. In addition to being useful as predictors of function, the fitCons scores could be useful in obtaining estimates of the SUS that are less sensitive to turnover because they measure natural selection over much shorter time scales.

An initial estimate of the SUS can be obtained by simply averaging the fitCons scores across all nucleotide positions in the genome. Because each score represents a probability that an individual nucleotide influences fitness, their average represents an expected fraction of nucleotides in the genome having fitness-influencing functions, or an expected SUS. This approach yields an estimate of 7.5% ($\pm 0.1\%$) for HUVEC, or 7.5–7.8% across cell types. These estimates are largely consistent with, but on the high end, of those based on cross-species divergence, which generally have fallen between 3 and 8%^{12,40,44–46}. Among sites under selection, we estimate that 9.0% are in CDS, 2.2% in 3' UTRs, 35.2% in introns, 51.7% in intergenic regions, and <1% in each of several other noncoding annotation classes (Supplementary Table 2). Our estimates of the SUS are somewhat lower than previous estimates that have explicitly allowed for evolutionary turnover, most of which have been 2–3 times higher than the pan-mammalian estimates of ~5%^{26,44,46–48}. However, they are similar to a recent estimate of 7.1–9.2% based on improved alignments and a new model for turnover⁴⁹ (see Discussion).

Violations of modeling assumptions will tend to bias the fitCons scores upward, particularly for functional classes for which the true fraction is close to zero (see Supplementary Note). To address this problem, we performed a parallel calculation for “neutral” sites that intersect the large class of genomic positions having a “null” functional genomic fingerprint (i.e., no DNase-seq, RNA-seq, or histone modification signal). This calculation results in an estimate of 3.3%, which can be considered an upper bound on the contribution of error because these putatively neutral sites undoubtedly include some sites under selection. By subtracting this 3.3% from our naive estimate of 7.5%, we obtain an estimated lower bound for the SUS of 4.2%, with somewhat higher fractions of selected sites in CDSs and 3' UTRs (Supplementary Table 2). (These estimates are for HUVEC, but the results for the other cell types were very similar.) Overall, our analysis of the SUS suggests that between 4.2% and

7.5% of nucleotides in the genome have direct fitness-influencing functions, and that the ratio of noncoding to coding functional sites is between 5.4 and 10.1.

Implications for Evolutionary Turnover of Functional Elements

To better understand the differences between fitCons scores and conventional divergence-based scores, we devised an alternative scoring system (denoted “fitConsD”) based on the same site clusters but an estimator of the fraction of nucleotides under selection that instead considers nucleotide divergence patterns across primates (see Methods). Thus, the fitCons and fitConsD scores both represent probabilities of fitness consequences per nucleotide but over two different evolutionary time scales. Overall, these two measures are remarkably well correlated, with $R^2=0.88$ (Fig. 6A). Furthermore, a measure based on the difference between fitConsD and fitCons scores suggests relatively low amounts of turnover across annotation classes, accounting for no more than about 10% of all functional sites (Fig. 6B). These observations suggest that the main signal for selection has been maintained over long evolutionary time periods and that turnover has been modest during primate evolution, but that there are some classes of sites that show stronger recent than ancient natural selection.

Discussion

The essential idea of our approach is to use functional genomic data to group sites into classes that are relatively homogeneous in terms of their functional roles, then to characterize the bulk influence of natural selection on those classes based on their patterns of polymorphism and divergence. For our estimation of natural selection, we make use of a recently developed probabilistic model of evolution and efficient algorithms for genome-wide inference (INSIGHT). We interpret INSIGHT-based estimates of fractions of nucleotides under selection as probabilities that each nucleotide influences fitness, or *fitness consequence* (fitCons) scores. Even with a simple clustering scheme, these fitCons scores appear to be highly informative about genomic function.

Based on our experiments, the fitCons scores have excellent predictive performance for putative *cis*-regulatory elements, outperforming several divergence-based methods (phastCons, phyloP, GERP, and CADD) and one annotation-based method (RegulomeDB) by clear margins. They also performed slightly better in enhancer prediction than EnhancerFinder, a program specifically designed for this purpose, although it should be noted that EnhancerFinder was trained on other cell types. Notably, prediction performance does not appear to be sensitive to the choice of neutral sites used by INSIGHT (Supplementary Figs. 6 and 7). In part, the observed improvement in performance reflects the use of cell-type-specific data (Fig. 5 and Supplementary Fig. 2), but the fitCons scores also show a clear performance advantage when considering all annotated elements rather than just active ones (Supplementary Fig. 3). Thus, the approach of grouping genomic sites by functional genomic signatures and then measuring group wise fitness consequences based on patterns of genetic variation appears to offer real benefits for prediction of regulatory function, as compared with methods that consider either genetic divergence or functional genomic data alone.

Interestingly, the recently published CADD method performed no better on our tests than conventional conservation scores, despite reports by the authors of significant advantages over phyloP, phastCons, GERP, and other methods³⁵. This inconsistency appears to reflect several important differences between our validation experiments and those reported in reference [35]. First, our tests focused specifically on putative *cis*-regulatory elements, while many of their tests considered a mixture of coding and noncoding elements. In particular, the ClinVar database, which figured prominently in their experiments, includes very few noncoding variants (~5% of pathogenic variants). Second, when Kircher et al. did consider noncoding regions, they generally did not distinguish between *cis*-regulatory elements and sequences that more directly influence the structure and content of protein-coding transcripts, such as splice sites. CADD has a natural advantage with these variants due to its use of gene annotations, whereas the annotation-free fitCons scores may perform better in completely unannotated regions of the genome. Finally, the tests by Kircher et al. that explicitly considered putative *cis*-regulatory elements were limited to a few loci and considered only correlations with saturation mutagenesis experiments, irrespective of a prediction threshold. We view our ROC-type comparisons based on multiple independent genome-wide sets of elements as a more direct and comprehensive demonstration of predictive power for *cis*-regulatory elements. In any case, the comparison of these two closely related, yet distinct, approaches helps to reveal strengths and weaknesses of each, and may lead to new ideas for improved methodologies.

A side benefit of our model-based approach is that the basewise probabilities of fitness consequences lead in a straightforward manner to an estimate of the “share under selection” (SUS) in the human genome. This estimate of the SUS reflects time scales since the divergence of humans and chimpanzees, about 4–6 million years ago, unlike conventional estimates based on tens or hundreds of millions of years of mammalian evolution. Nevertheless, our estimate of the SUS—at 4.2–7.5%—ends up being remarkably similar to those based on longer time scales, which have generally fallen between 3 and 8%^{12,40–43,45,50}. It also overlaps with a recent estimate of 7.1–9.2% based on patterns of insertion and deletion and an explicit model of evolutionary turnover⁴⁹. We take the general concordance of these estimates, both with one another and with our fitCons- and fitConsD-based estimates, as a strong indication that the SUS has remained quite low (probably <10%) over various time-scales in mammalian evolution. This finding stands in contrast to estimates that ~80% of nucleotides may be functional, based on measures of “biochemical activity”³. However, it is important to bear in mind that these evolutionary and biochemical estimates reflect somewhat different definitions of “function,” and this may explain some of the difference between them^{16,18,19}. For example, the fitCons- and conservation-based estimates (excluding those based on indels) generally represent fractions of positions at which point mutations will have fitness consequences, but they do not account for sequences (such as spacer elements) that would have fitness consequences if deleted but not mutated (see Supplementary Note for discussion).

Apart from the absolute fraction of functional DNA in the human genome is the question of how much the functional content of the genome has changed over time through gains and losses of functional elements. Several studies have estimated that such “turnover” could

allow the current SUS in the human genome to be ~2–3 times larger than estimated from comparisons across mammals^{26,44,46–48}. Indeed, these findings have been proposed to explain, in part, the discordance between evolution-based and biochemical estimates of the functional fraction of the genome^{26,51,52}. However, most of these analyses have accounted for turnover using relatively crude methods, for example, by relying on an apparently near-linear relationship between pairwise divergence and the estimated SUS^{46,47} or by estimating functional content from mean SNP densities or derived allele frequencies in genomic regions not conserved across mammals²⁶ (but see reference [49] for an improved model). Our analysis is more direct, by comparing analogous divergence-based and polymorphism-based estimates of the SUS based on exactly the same clusters of nucleotide positions. In addition, our analysis focuses on primate evolution, rather than attempting to account for turnover across mammals, where factors such as alignment error, orthology detection, and genomic rearrangements can be problematic. The similarity between our estimates based on polymorphism (fitCons) and divergence (fitConsD) strongly suggests that evolutionary turnover has been modest during primate evolution, because massive turnover would be expected to lead to a substantial downward bias in the divergence-based estimates. Our power experiments indicate that this observation is not an artifact of reduced sensitivity in the fitCons scores. Nevertheless, we cannot rule out the possibility that compensating gains and losses on very recent time scales maintain a similar SUS while substantially altering the genomic composition of functional sequences.

We have focused on HUVEC in this paper, but we also generated fitCons scores for two other cell types (H1 hESC and GM12878). A comparison across cell types (see Supplementary Note) indicated that the genomic positions assigned to each functional class differed substantially across cell types, but equivalently defined clusters had concordant fitCons scores in the different cell types (Supplementary Fig. 8). When cell-type-specific scores were examined, elements “active” in that cell type displayed significantly higher scores than inactive elements. Moreover, particular elements had higher scores in cell types for which they were active than in cell types for which they were inactive (Supplementary Fig. 9). Notably, we found that a set of integrated scores based on a simple, heuristic procedure (see Methods) performed nearly as well as the cell-type-specific scores in the target cell types, but much better on elements from mismatched or pooled cell types (Supplementary Fig. 10). With more flexible and scalable clustering techniques, it may be possible to improve these methods by considering all cell types simultaneously, clustering sites by multi-cell-type functional genomic fingerprints, and then producing a single set of scores reflecting these joint patterns. Such improvements, together with increases in the resolution and quality of the available functional genomic data, should result in improved power for individual functional elements and refined estimates of the share under selection.

Online Methods

Functional Genomic Data

RNA-seq and DNase-seq data for HUVEC, H1 hESC, and GM12878 were downloaded from the University of California, Santa Cruz (UCSC) Genome Browser. Chromatin states for the same three cell types were downloaded from the European Bioinformatics Institute

ftp site (see Supplementary Table 3). For DNase-seq, we considered two replicate experiments from University of Washington (UW) data for each cell type. However, only one UW replicate was available for H1 hESC so additional DNase-seq data for this cell line was obtained from Duke University. For each replicate DNase-seq experiment, we downloaded broad and narrow peak calls. For RNA-seq, we selected a single replicate from the Caltech poly-A+ 75 bp paired-end read data, after examining several alternative data sets. For chromatin states, we used the 25-state ChromHMM segmentation generated in December, 2012³³.

Clustering Approach

We produced a separate partitioning for each cell type based on the functional genomic data. The broad and narrow DNase-seq peaks were used to partition sites in the genome into three mutually exclusive classes: sites that fall in a narrow peak in both replicate experiments (2); sites that fall in a broad peak in at least one replicate and do not fall in a narrow peak in both replicates (1); and sites that fall outside of all called peaks (0). This three-level scheme allows for both high sensitivity (class 1) and high specificity (class 2). For H1 hESC, only one set of broad peak calls was available to define class 1. For the RNA-seq data, we partitioned sites in the genome into four mutually exclusive classes (0–3) based on numbers of reads aligned at each position. Read depth thresholds were set separately for each cell type through a process that aims to minimize the conditional entropy of concentrations of predicted sites under selection (see Supplementary Note). Chromatin states were defined directly from the 25 states in ChromHMM, with a 26th state containing sites not assigned to any chromatin class. The Cartesian product of these partitions, together with the partition into coding and noncoding sequences (see below), resulted in $3 \times 4 \times 26 \times 2 = 624$ distinct functional classes.

Running INSIGHT

The INSIGHT method infers the fraction of nucleotide sites under selection (ρ) for a given collection of sites by comparing patterns of within-species polymorphism and between-species divergence within those sites and within putatively neutrally evolving sites nearby. A detailed description of the method, sequence data, and data-quality filters is given in reference [34]. For each non-empty fitCons site cluster, INSIGHT was used to estimate ρ , which was used as the fitCons score of all sites in that cluster. To reduce sensitivity to estimates with high uncertainty, we filtered out clusters for which the estimated standard error was larger than 40% of the estimated value of ρ . To increase computational efficiency, clusters larger than 20 Mb were partitioned into smaller subclasses, and estimates of ρ were computed as weighted averages (by numbers of informative sites) across subclasses.

Neutral Sites

The collection of sites predicted to be free from the influence of natural selection (“neutral” sites) was derived from a set identified previously^{28,34,53}. Briefly, this set is obtained by eliminating from the genome sites likely to be under direct natural selection, including (1) exons of annotated protein-coding genes and the 1000 bp flanking them on either side; (2) RNA genes from GENCODE v11 and 1000 bp flanks; and (3) conserved non-coding

elements (identified by phastCons) and 100 bp flanks. This set was used both by INSIGHT and for our power analysis (see below).

GENCODE Annotations

Transcript annotations from GENCODE v15⁵⁴ were downloaded from the Sanger Institute FTP server and used to define eight site classes: coding sequences (CDS), 5' untranslated regions (UTRs), 3' UTRs, promoters, introns, long intergenic noncoding RNAs (lincRNA), short noncoding RNAs (sncRNA), and intergenic (sites not falling within any protein-coding transcription unit). Transcripts annotated with feature type="CDS" and gene type="protein coding" were used to define the CDS set for fitCons. For subsequent analysis, we used a slightly more conservative set, obtained by additionally requiring feature type="gene", gene status="KNOWN", transcript status="KNOWN", and the identification of both start and stop codons within the transcript. UTRs were defined from transcripts having feature type="UTR" and gene type="protein coding" and designated as 5' or 3'. Introns were defined by positions that fall within a protein-coding transcript but outside of the CDS and UTR regions. Promoters were defined as the 1000 bp immediately upstream of the first (i.e., most upstream) transcription start site for each protein-coding gene. A similarly defined alternative set of 100 bp promoter regions was used in assessing differences between cell types (see Supplementary Fig. 9). LincRNAs were identified by transcripts with feature type="exon" and gene type="lincRNA". Similarly, sncRNAs consisted of transcripts with feature type="exon" and gene type ∈ {"miRNA", "snRNA", "snoRNA"}. Positions in the more inclusive CDS set were removed from all noncoding classes.

Cis-Regulatory Elements

Transcription factor binding sites (TFBSs) were drawn from a set for 78 transcription factors, based on ChIP-seq data from ENCODE²⁸ downloadable from our Genome Browser mirror. This set contains roughly 1.4 million binding sites of mean length 11 bp, each of which is associated with the cell types in which it was detected. For some tests, we considered only the subset of nucleotide positions inside these TFBSs that corresponded to motif positions with strong base preferences, defined as those positions at which the consensus allele appeared in at least 90% of all binding sites (according to the inferred motif model). For enhancers, we used the distal regulatory modules described in reference [38]. We downloaded the file enets4.Distal cell line.txt from the Gerstein Lab and extracted from it a total of 19,005 enhancer-transcript associations, covering 5,834 unique autosomal loci with a mean length of 888 bp, along with the cell types associated with each predicted enhancer. Expression quantitative trait loci (eQTL) described in reference [6] were downloaded from the European Bioinformatics Institute's E-GEUV-1 data set. We used the four files {EUR373, YRI89}.{exon, gene}.cis.FDR5.best.rs137.txt.gz to identify 6,760 distinct autosomal positions and the associated transcripts, removing all positions overlapping CDSs.

Identifying Active Elements per Cell Type

In several analyses, we considered the subset of elements in each annotation class for which we had evidence of activity in a given cell type. To identify the cell types in which TFBS and enhancers were active, we used the cell type designations provided in the corresponding

annotation files (see above). For other classes of elements, we defined the active elements using a set of GENCODE transcripts and genes that showed significantly elevated levels of RNA transcription in the Caltech RNA-seq data. Those were transcripts (or genes) for which the 95% confidence interval of the normalized read count in a given cell type fell within the top one third of normalized read counts for transcripts (or genes) across all three cell types (threshold 1.477 for transcripts and 4.966 for genes). Active eQTLs were identified via associated active genes using the GENCODE gene identifier specified for each eQTL. Active promoters, UTRs, CDS, and introns, were identified via associated active transcripts. For the comparison between cell types (Supplementary Fig. 9) we also used collections of eQTLs and promoters found to be inactive in a given cell type. Those were defined in a similar way, by using transcripts and genes falling in the bottom third of the distribution of normalized read counts.

Comparison with Other Scores

Base-wise scores from the Genomic Evolutionary Rate Profiling (GERP)¹³ method, the Combined Annotation Dependent Depletion (CADD) method, and phastCons¹² and phyloP¹⁵ methods were downloaded from the respective websites (see URLs; file hg19.GERP.scores.tar.gz generated in August, 2010 for GERP, file whole genome SNVs.tsv.gz, downloaded in September, 2013 for CADD, and the 46 placental mammals conservation tracks for phastCons and phyloP). CADD scores are specified for each genomic position and each of the three possible variant bases in that position. We took the maximum of these three scores, which yielded the best performance for CADD in our comparisons. We also used RegulomeDB³⁶ (downloaded in January, 2013) to rank single nucleotide polymorphisms (SNPs), such as eQTLs, into one of 13 categories according to evidence from functional genomic data. Finally, we obtained EnhancerFinder scores³⁷ for 1500 bp windows tiled across the genome directly from the authors. We used the general, non-tissue-specific scores, and averaged them at positions contained in multiple overlapping windows.

Receiver Operating Characteristic (ROC) Curves

We used ROC curves to measure the ability of each scoring scheme to discriminate between functional and nonfunctional regulatory elements. For TFBSs and enhancers, we used the annotations described above as “true positives” and defined “true negatives” from our filtered, putatively neutral sites (see above). For eQTLs, our negative set consisted of all 9.8 million variants tested in reference [6], excluding indels, non-simple variants, and positions that showed possible associations at a threshold of nominal $p < 0.05$ (7.6 million SNPs remained). In all three cases, we additionally removed any sites in the positive set from the negative set. A point on a ROC plot indicates the fraction of annotated genomic positions with scores higher than a given score (true positive rate) versus the fraction of control genomic positions with scores higher than that score (false positive rate). Positions with no scores are ignored when computing fractional coverage.

Integrating fitCons Scores across Cell Types

We generated a series of fitCons scores that integrate functional genomic data across the three cell types by using the original 624 fingerprints (see above) and altering the rule by which sites are assigned to clusters to reflect information from multiple cell types. Our approach attempts to select a fingerprint for each site that was likely to be most informative about the site's function, while avoiding a bias toward higher scores with increasing cell types. See Supplementary Note for details.

Share Under Selection

Assume a partitioning of the genome into K mutually exclusive and exhaustive clusters, C_1, \dots, C_K , and a corresponding set of fitCons scores, $\rho(C_1), \dots, \rho(C_K)$. Note that the expected number of genomic positions under selection in cluster C_i is given by $\rho(C_i)|C_i|$, because ρ is an estimate of the fraction of sites under selection. For an arbitrary collection of sites, S , the expected number of sites in S that are under selection is given by $\text{sel}(S) = \sum_i \rho(C_i) |C_i \cap S|$, and the average fitCons score for S is given by $\rho(S) = \text{sel}(S) / |S|$. To avoid under-estimation of $\rho(S)$, we do not filter out fitCons scores with high uncertainty in these calculations, as we do for other analyses (see above). In addition, to account for possible overestimation of ρ in very large clusters having low fractions of sites under selection, we ran INSIGHT on the intersection of our neutral sites (see above) and all non-coding sites in the 'Quiescent' chromatin state with no DNase-seq or RNA-seq signal. We then subtracted the estimated value of ρ , denoted ρ_{neut} , from the raw fitCons score to obtain a conservative lower bound, $\rho(S) - \rho_{\text{neut}}$, for the fraction of sites under selection in S .

FitConsD and Evolutionary Turnover

To make the comparison between fitCons and fitConsD as direct as possible, fitConsD scores were computed using the same pipeline we developed for fitCons (see Fig. 1), except that in step C we replaced the INSIGHT model with an evolutionary model that considers sequence divergence between the human, chimpanzee, orangutan, and rhesus macaque genomes. The fitConsD scores are based on an estimate s_i for the relative evolutionary rate of each cluster C_i compared with a neutral model globally estimated for the four-primate phylogeny. This relative rate is then compared with the relative rate s_i^{neut} estimated for the putative neutral regions flanking sites in C_i and a divergence-based estimate of the fraction of sites under selection in cluster C_i is given by $\rho_{\text{div}}(C_i) = 1 - s_i / s_i^{\text{neut}}$ (see Supplementary Note).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank Leonardo Arbiza for helpful discussions and assistance with early analyses, and Greg Cooper for constructive criticism of our validation experiments and comparisons with CADD. This research was supported by NIH grant GM102192, a David and Lucile Packard Fellowship for Science and Engineering (to AS), and a postdoctoral fellowship from the Cornell Center for Comparative and Population Genomics (to IG). The content is solely the responsibility of the authors and does not necessarily represent the official views of the US National Institutes of Health.

REFERENCES

1. Mardis ER. A decade's perspective on DNA sequencing technology. *Nature*. 2011; 470:198–203. [PubMed: 21307932]
2. Wold B, Myers RM. Sequence census methods for functional genomics. *Nat Methods*. 2008; 5:19–21. [PubMed: 18165803]
3. Consortium EP. An integrated encyclopedia of DNA elements in the human genome. *Nature*. 2012; 489:57–74. [PubMed: 22955616]
4. Shen Y, et al. A map of the cis-regulatory sequences in the mouse genome. *Nature*. 2012; 488:116–120. [PubMed: 22763441]
5. Neph S, et al. An expansive human regulatory lexicon encoded in transcription factor footprints. *Nature*. 2012; 489:83–90. [PubMed: 22955618]
6. Lappalainen T, et al. Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*. 2013; 501:506–511. [PubMed: 24037378]
7. Cooper GM, Shendure J. Needles in stacks of needles: finding disease-causal variants in a wealth of genomic data. *Nat Rev Genet*. 2011; 12:628–640. [PubMed: 21850043]
8. Mayor C, et al. VISTA : visualizing global DNA sequence alignments of arbitrary length. *Bioinformatics*. 2000; 16:1046–1047. [PubMed: 11159318]
9. Margulies EH, Blanchette M, Program NCS, Haussler D, Green ED. Identification and characterization of multi-species conserved sequences. *Genome Res*. 2003; 13:2507–2518. [PubMed: 14656959]
10. Boffelli D, et al. Phylogenetic shadowing of primate sequences to find functional regions of the human genome. *Science*. 2003; 299:1391–1394. [PubMed: 12610304]
11. Ovcharenko I, Boffelli D, Loots GG. eShadow: a tool for comparing closely related sequences. *Genome Res*. 2004; 14:1191–1198. [PubMed: 15173121]
12. Siepel A, et al. Evolutionarily conserved elements in vertebrate, insect, worm, and yeast genomes. *Genome Res*. 2005; 15:1034–1050. [PubMed: 16024819]
13. Cooper GM, et al. Distribution and intensity of constraint in mammalian genomic sequence. *Genome Res*. 2005; 15:901–913. [PubMed: 15965027]
14. Asthana S, Roytberg M, Stamatoyannopoulos J, Sunyaev S. Analysis of sequence conservation at nucleotide resolution. *Plos Computational Biology*. 2007; 3:2559–2568.
15. Pollard KS, Hubisz MJ, Rosenbloom KR, Siepel A. Detection of nonneutral substitution rates on mammalian phylogenies. *Genome Res*. 2010; 20:110–121. [PubMed: 19858363]
16. Graur D, et al. On the immortality of television sets: "function" in the human genome according to the evolution-free gospel of ENCODE. *Genome Biol Evol*. 2013; 5:578–590. [PubMed: 23431001]
17. Niu DK, Jiang L. Can ENCODE tell us how much junk DNA we carry in our genome? *Biochem Biophys Res Commun*. 2013; 430:1340–1343. [PubMed: 23268340]
18. Doolittle WF. Is junk DNA bunk? A critique of ENCODE. *Proceedings of the National Academy of Sciences of the United States of America*. 2013; 110:5294–5300. [PubMed: 23479647]
19. Eddy SR. The ENCODE project: Missteps overshadowing a success. *Current Biology*. 2013; 23:R259–R261. [PubMed: 23578867]
20. McDonald JH, Kreitman M. Adaptive Protein Evolution at the Adh Locus in *Drosophila*. *Nature*. 1991; 351:652–654. [PubMed: 1904993]
21. Fay JC, Wyckoff GJ, Wu CI. Positive and negative selection on the human genome. *Genetics*. 2001; 158:1227–1234. [PubMed: 11454770]
22. Andolfatto P. Adaptive evolution of non-coding DNA in *Drosophila*. *Nature*. 2005; 437:1149–1152. [PubMed: 16237443]
23. Eyre-Walker A, Woolfit M, Phelps T. The distribution of fitness effects of new deleterious amino acid mutations in humans. *Genetics*. 2006; 173:891–900. [PubMed: 16547091]
24. Boyko AR, et al. Assessing the evolutionary impact of amino acid mutations in the human genome. *PLoS Genet*. 2008; 4:e1000083. [PubMed: 18516229]

25. Wilson DJ, Hernandez RD, Andolfatto P, Przeworski M. A population genetics-phylogenetics approach to inferring natural selection in coding sequences. *PLoS Genet.* 2011; 7:e1002395. [PubMed: 22144911]
26. Ward LD, Kellis M. Evidence of abundant purifying selection in humans for recently acquired regulatory functions. *Science.* 2012; 337:1675–1678. [PubMed: 22956687]
27. Khurana E, et al. Integrative annotation of variants from 1092 humans: application to cancer genomics. *Science.* 2013; 342:1235587. [PubMed: 24092746]
28. Arbiza L, et al. Genome-wide inference of natural selection on human transcription factor binding sites. *Nat Genet.* 2013; 45:723–729. [PubMed: 23749186]
29. Narlikar L, et al. Genome-wide discovery of human heart enhancers. *Genome Res.* 2010; 20:381–392. [PubMed: 20075146]
30. Ritchie GR, Dunham I, Zeggini E, Flicek P. Functional annotation of noncoding sequence variants. *Nat Methods.* 2014; 11:294–296. [PubMed: 24487584]
31. Ernst J, Kellis M. Discovery and characterization of chromatin states for systematic annotation of the human genome. *Nat Biotechnol.* 2010; 28:817–825. [PubMed: 20657582]
32. Hoffman MM, et al. Unsupervised pattern discovery in human chromatin structure through genomic segmentation. *Nature Methods.* 2012; 9:473–U488. [PubMed: 22426492]
33. Hoffman MM, et al. Integrative annotation of chromatin elements from ENCODE data. *Nucleic Acids Research.* 2013; 41:827–841. [PubMed: 23221638]
34. Gronau I, Arbiza L, Mohammed J, Siepel A. Inference of Natural Selection from Interspersed Genomic Elements Based on Polymorphism and Divergence. *Molecular Biology and Evolution.* 2013; 30:1159–1171. [PubMed: 23386628]
35. Kircher M, et al. A general framework for estimating the relative pathogenicity of human genetic variants. *Nat Genet.* 2014; 46:310–315. [PubMed: 24487276]
36. Boyle AP, et al. Annotation of functional variation in personal genomes using RegulomeDB. *Genome Research.* 2012; 22:1790–1797. [PubMed: 22955989]
37. Erwin GD, et al. Integrating diverse datasets improves developmental enhancer prediction. *PLoS Comput Biol.* 2014; 10:e1003677. [PubMed: 24967590]
38. Gerstein MB, et al. Architecture of the human regulatory network derived from ENCODE data. *Nature.* 2012; 489:91–100. [PubMed: 22955619]
39. Core LJ, et al. Analysis of nascent RNA identifies a unified architecture of transcription initiation regions at mammalian promoters and enhancers. In press. *Nat Genet.* 2014 In Press.
40. Mouse Genome Sequencing, C. Initial sequencing and comparative analysis of the mouse genome. *Nature.* 2002; 420:520–562. [PubMed: 12466850]
41. Cooper GM, et al. Characterization of evolutionary rates and constraints in three Mammalian genomes. *Genome Res.* 2004; 14:539–548. [PubMed: 15059994]
42. Lindblad-Toh K, et al. Genome sequence, comparative analysis and haplotype structure of the domestic dog. *Nature.* 2005; 438:803–819. [PubMed: 16341006]
43. Lindblad-Toh K, et al. A high-resolution map of human evolutionary constraint using 29 mammals. *Nature.* 2011; 478:476–482. [PubMed: 21993624]
44. Ponting CP, Nellaker C, Meader S. Rapid Turnover of Functional Sequence in Human and Other Genomes. *Annual Review of Genomics and Human Genetics, Vol 12.* 2011; 12:275–299.
45. Chiaromonte F, et al. The share of human genomic DNA under selection estimated from human-mouse genomic alignments. *Cold Spring Harb Symp Quant Biol.* 2003; 68:245–254. [PubMed: 15338624]
46. Meader S, Ponting CP, Lunter G. Massive turnover of functional sequence in human and other mammalian genomes. *Genome Res.* 2010; 20:1335–1343. [PubMed: 20693480]
47. Smith NG, Brandstrom M, Ellegren H. Evidence for turnover of functional noncoding DNA in mammalian genome evolution. *Genomics.* 2004; 84:806–813. [PubMed: 15475259]
48. Ponting CP, Hardison RC. What fraction of the human genome is functional? *Genome Research.* 2011; 21:1769–1776. [PubMed: 21875934]

49. Rands CM, Meader S, Ponting CP, Lunter G. 8.2% of the Human Genome Is Constrained: Variation in Rates of Turnover across Functional Element Classes in the Human Lineage. *Plos Genetics*. 2014; 10
50. Lunter G, Ponting CP, Hein J. Genome-wide identification of human functional DNA using a neutral indel model. *Plos Computational Biology*. 2006; 2:2–12.
51. Kellis M, et al. Defining functional DNA elements in the human genome. *Proceedings of the National Academy of Sciences of the United States of America*. 2014; 111:6131–6138. [PubMed: 24753594]
52. Pheasant M, Mattick JS. Raising the estimate of functional human sequences. *Genome Res*. 2007; 17:1245–1253. [PubMed: 17690206]
53. Gronau I, Hubisz MJ, Gulko B, Danko CG, Siepel A. Bayesian inference of ancient human demography from individual genome sequences. *Nat Genet*. 2011; 43:1031–1034. [PubMed: 21926973]
54. Harrow J, et al. GENCODE: the reference human genome annotation for The ENCODE Project. *Genome Res*. 2012; 22:1760–1774. [PubMed: 22955987]

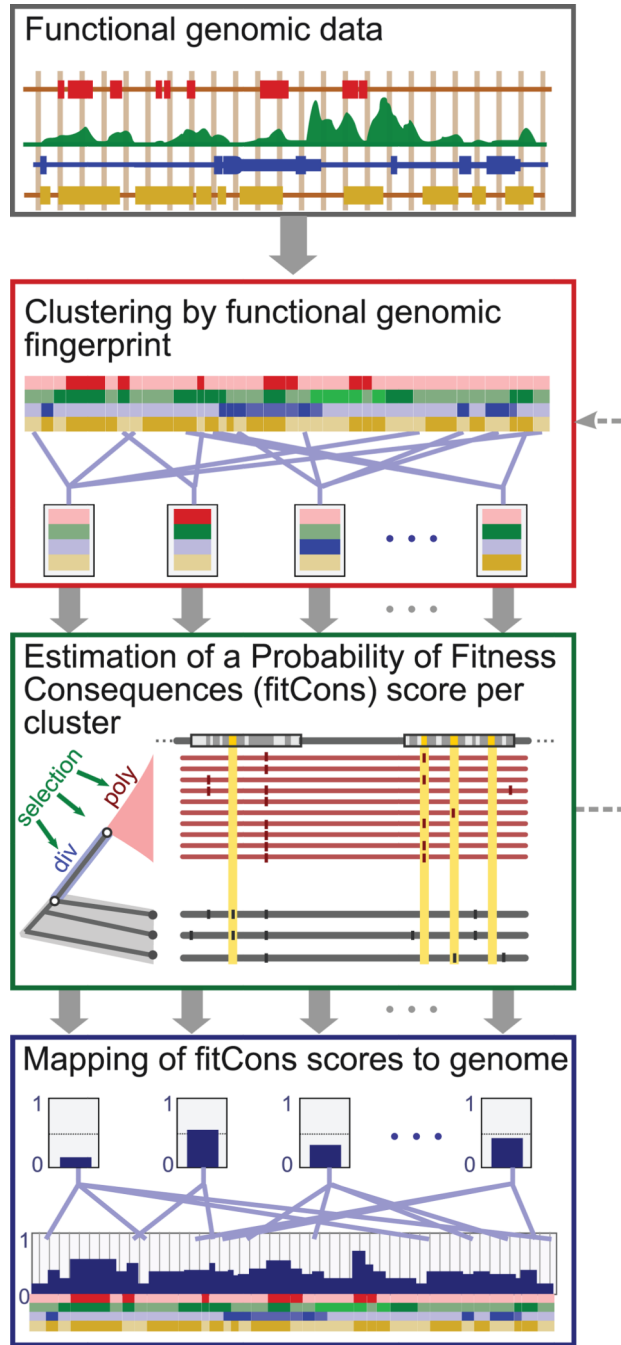
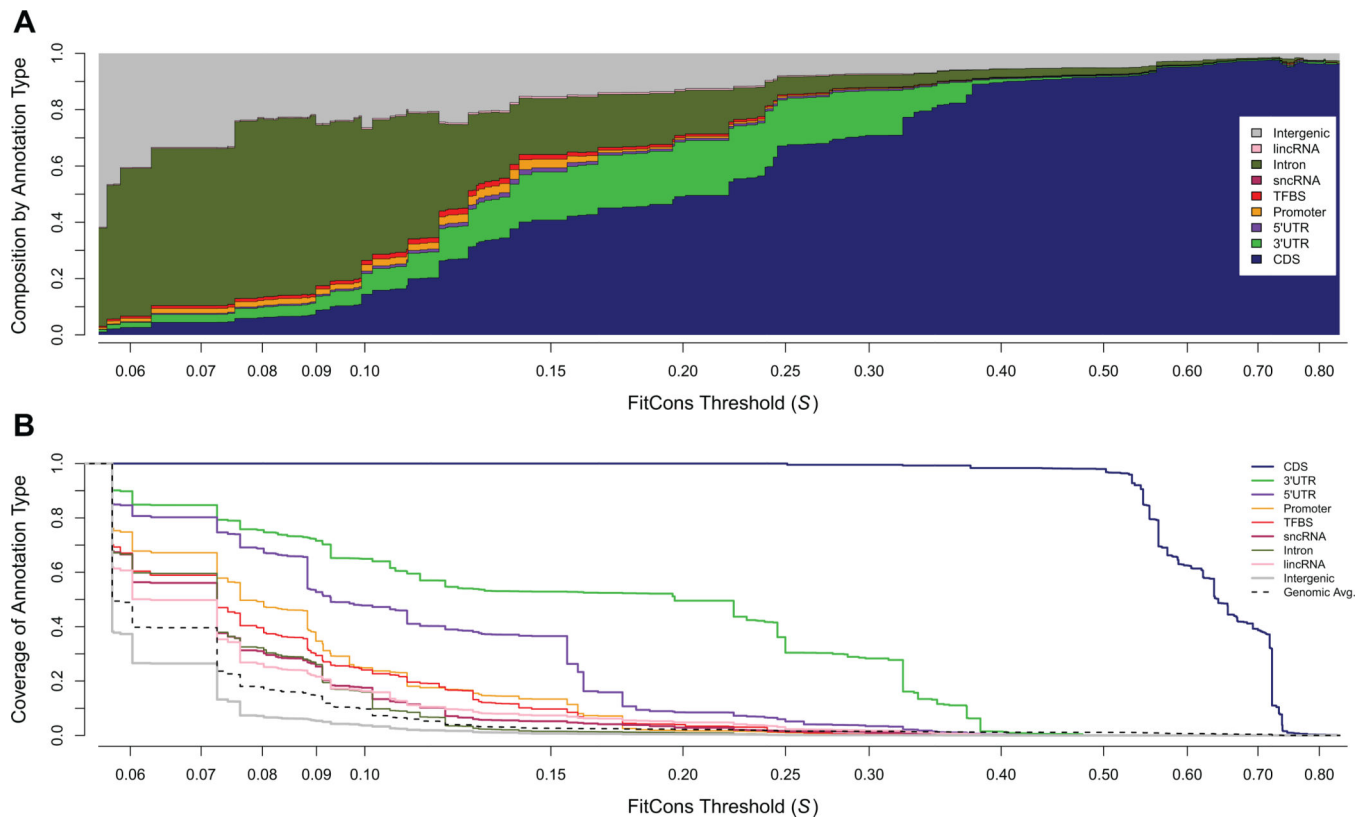


Figure 1. Illustration of procedure for calculating fitCons scores. (A) Functional genomic data, such as DNase-seq, RNA-seq and histone modification data, are arranged along the genome sequence in tracks. (B) Nucleotide positions in the genome are clustered by joint patterns across these functional genomic tracks. For example, one cluster might contain genomic positions with a high DNase-seq signal, a moderate RNA-seq signal, and high signals for H3K4me1 and H3K27ac, suggesting transcribed enhancers. Another might contain positions with a low DNase-seq signal, a high RNA-seq signal, and a signal for H3K36me3,

suggesting actively transcribed gene bodies. Notice that clusters will generally contain genomic positions dispersed along the genome sequence. (C) Patterns of polymorphism and divergence are analyzed using INSIGHT³⁴ to obtain an estimate of the fraction of nucleotides under natural selection (ρ) in each cluster. This quantity is interpreted as a probability that each nucleotide position influences the fitness of the organism that carries it, or a fitness consequence (fitCons) score. (D) The fitCons score for each cluster is assigned to all genomic positions that were included in the cluster. In this way, all nucleotide positions are assigned a score, but there can be no more distinct scores than there are clusters. Note that, in our initial work, the clustering of genomic positions is accomplished by a simple exhaustive partitioning scheme that produces 624 distinct clusters. In future work, however, it may be desirable to iterate between clustering and calculating scores (dashed line).

**Figure 2.**

Composition and coverage of high-scoring genomic regions according to fitCons. (A) Composition by annotation type in regions that exceed a fitCons score threshold of S , as S is varied across the range of possible scores. Each vertical cross-section of the plot can be thought of as a narrow “stacked bar” representation of the composition by annotation type of all genomic positions at which the fitCons score $> S$. At the left side of the plot, when S is small, the composition by annotation type is representative of the genome as a whole. As the threshold S increases, coding sequences (CDS; navy) are increasingly enriched and intergenic (gray) sequences are increasingly depleted. Regions experiencing moderate levels of selection, such as untranslated regions (UTRs; purple and green), promoters (orange), small noncoding RNAs (snRNAs; eggplant), and introns (olive), are most enriched at intermediate scores. Note the logarithmic scale for the x -axis. (B) Coverage of the same annotation types by genomic regions having fitCons score $> S$, with an x -axis matching that in (A). The dashed line indicates the genome-wide average. At each value of S , the relative height of a given curve in comparison to the dashed line indicates the enrichment (or depletion) of the corresponding annotation type in genomic regions having score $> S$. The legend at the right lists the annotation types in order of decreasing enrichment. When multiple annotations applied to a single nucleotide position, one was selected in the following order: CDS, TFBS, promoter, snRNA, lincRNA, 5' UTR, 3' UTR, intron, and intergenic. These figures summarize data at 2.9 billion genomic sites.

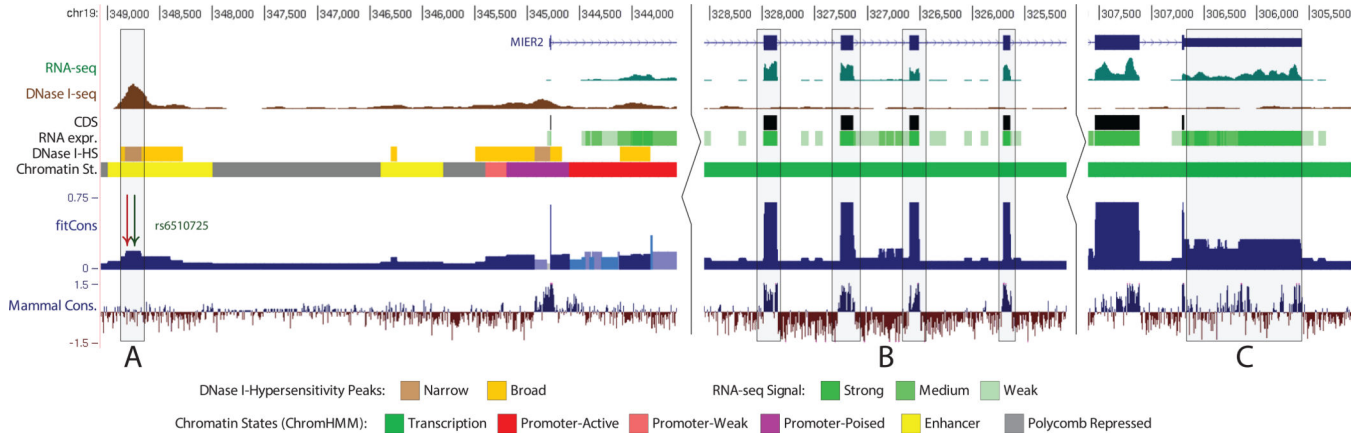


Figure 3. Genome browser display showing functional genomic fingerprints and fitCons scores. Shown, from top to bottom, are the exons of the MIER2 gene (blue); the raw RNA-seq (turquoise) and DNase-seq (brown) signals; the four discretized tracks used to define the 624 functional genomic fingerprints, including annotation-based CDSs (black), RNA-seq signal (green), DNase-seq signal (yellow and brown), and chromatin modifications (multiple colors; see key at bottom); the fitCons scores based on those fingerprints (dark blue, with lighter blues less statistically significant); and, for comparison, phyloP-based conservation scores for mammals. (A) An apparent enhancer, marked by a combination of enhancer-associated chromatin modifications and a strong DNase-seq signal, displays elevated fitCons scores but no elevation in conservation scores. Many regulatory elements display such a pattern, either because they have arisen recently in evolutionary time or because errors in orthology detection or alignment result in spuriously low conservation scores. Here, a ChIP-seq-supported TFBS for AP-1 (red arrow) and a lung-cancer-associated SNP (green arrow) are highlighted. (B) CDS exons show elevated scores according to both fitCons and phyloP. (C) The 3' UTR, marked by transcription-associated chromatin modifications, a high RNA-seq signal, and an absence of DNase-I hypersensitivity or CDS annotations, displays moderately elevated fitCons scores and patches of evolutionary conservation. The fitCons scores are fairly well correlated with phyloP conservation scores¹⁵ genome-wide, with some notable exceptions in noncoding regions (Supplementary Fig. 1). Browser tracks are publicly available on our genome browser mirror (hg19 assembly).

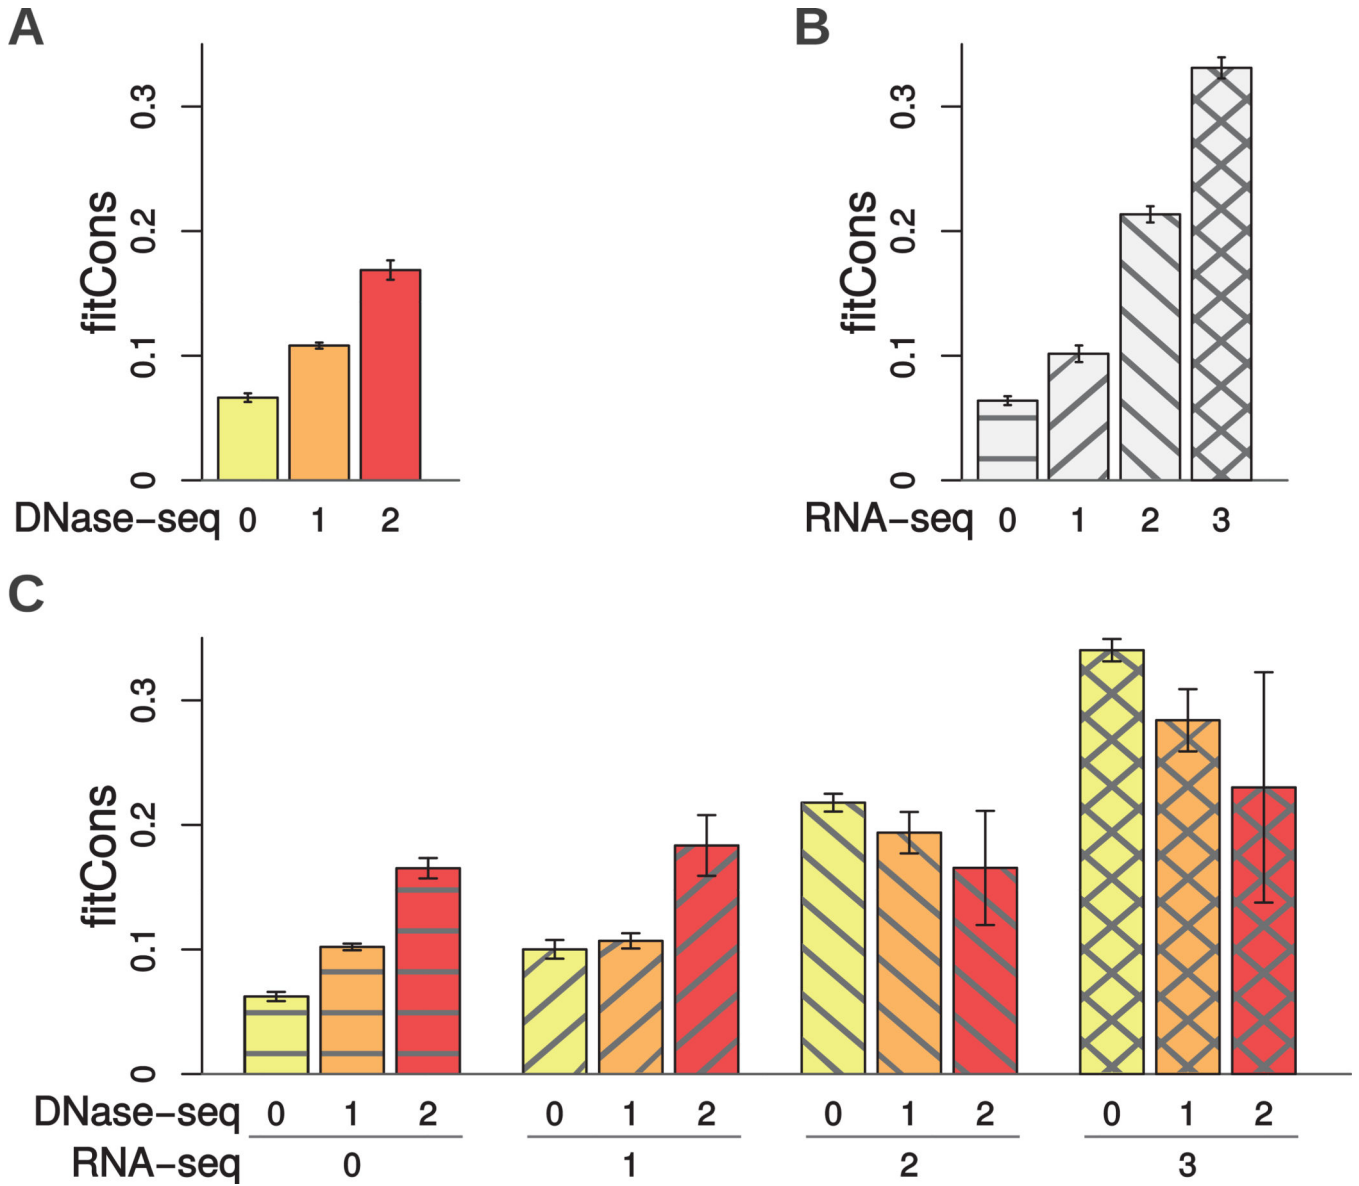


Figure 4. Average fitCons scores as a function of DNase-seq and RNA-seq intensity. Results represent averages across all non-CDS clusters having the marginal or joint property of interest. Error bars represent standard errors of the aggregated scores (see Methods). (A) FitCons scores increase with DNase-seq intensity, probably due to an increasing density of *cis*-regulatory elements. 0: No DNase-seq signal; 1: broad peaks; 2: narrow peaks. (B) FitCons scores increase with RNA-seq intensity. 0: no RNA-seq reads; 1–3: weak to strong RNA-seq signal (see Methods). (C) FitCons scores behave in a non-additive manner as joint combinations of DNase- and RNA-seq intensity are considered. In particular, at medium to high RNA-seq read depth (classes 2 and 3), fitCons scores decrease (rather than increase) with increasing DNase-seq signal. This unexpected pattern is explained by an enrichment for DNase-I hypersensitivity near the 5' ends of genes. Conditional on a high RNA-seq signal, a high DNase-seq signal tends to be associated with the 5' UTRs and upstream regions of genes,

which are under fairly weak selection, while a low DNase-seq signal is associated with 3' UTRs, which are under stronger selection. Each bar in (A) summarizes 104 clusters, each bar in (B) summarizes 78 clusters, and each bar in (C) summarizes 26 clusters.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

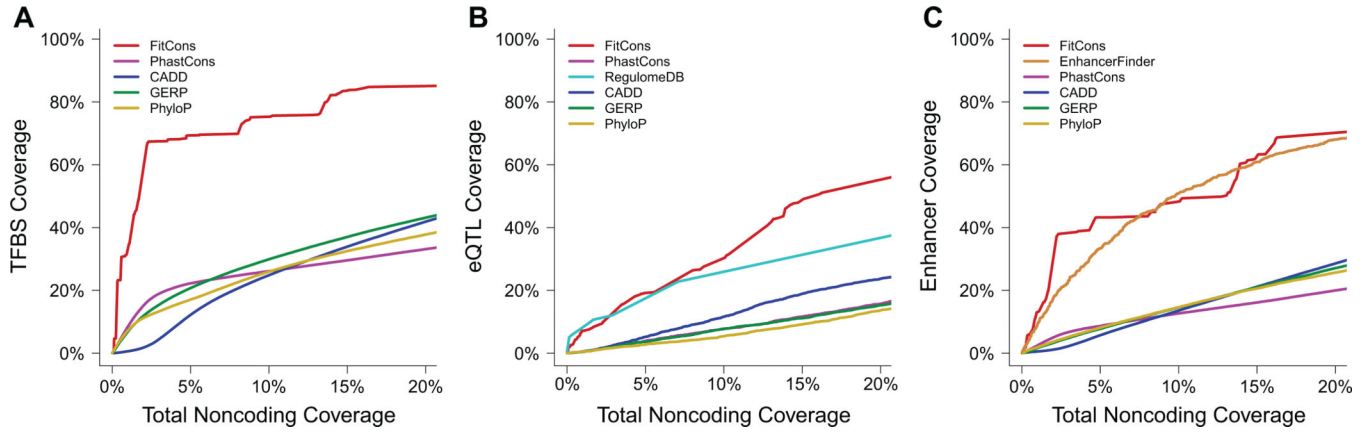


Figure 5.

Coverage of active *cis*-regulatory elements as a function of total coverage of the noncoding genome. Coverage of each type of element is shown as the score threshold is adjusted to alter the total coverage of noncoding sequences in the genome, excluding sites annotated as CDS or UTR. FitCons is compared with scores from the Combined Annotation Dependent Depletion (CADD)³⁵, Genomic Evolutionary Rate Profiling (GERP)¹³, phastCons¹², and phyloP¹⁵ programs (see Methods). (A) Coverage of 55,844 transcription factor binding sites detected by CHIP-seq in HUVEC²⁸. (B) Coverage of high-resolution eQTLs identified in a recent large-scale study⁶, restricted to 3,662 eQTLs associated with genes transcribed in HUVEC. Coverage of eQTLs is also shown for classification of single nucleotide variants by RegulomeDB³⁶. The divergence-based scores (phastCons, phyloP, GERP, and CADD) all perform poorly on the eQTL data set, probably because the ascertainment for segregating sites creates a bias against evolutionary conservation. Note also that the apparent performance of RegulomeDB, particularly at low total noncoding coverage, is somewhat influenced by consideration of eQTL data in its scoring scheme. (C) Coverage of 462 enhancers identified by characteristic chromatin marks³⁸ assayed in HUVEC. Coverage of these enhancers by EnhancerFinder³⁷ predictions is also shown. In all three plots, the *x*-axis represents coverage at 2.8 billion noncoding positions.

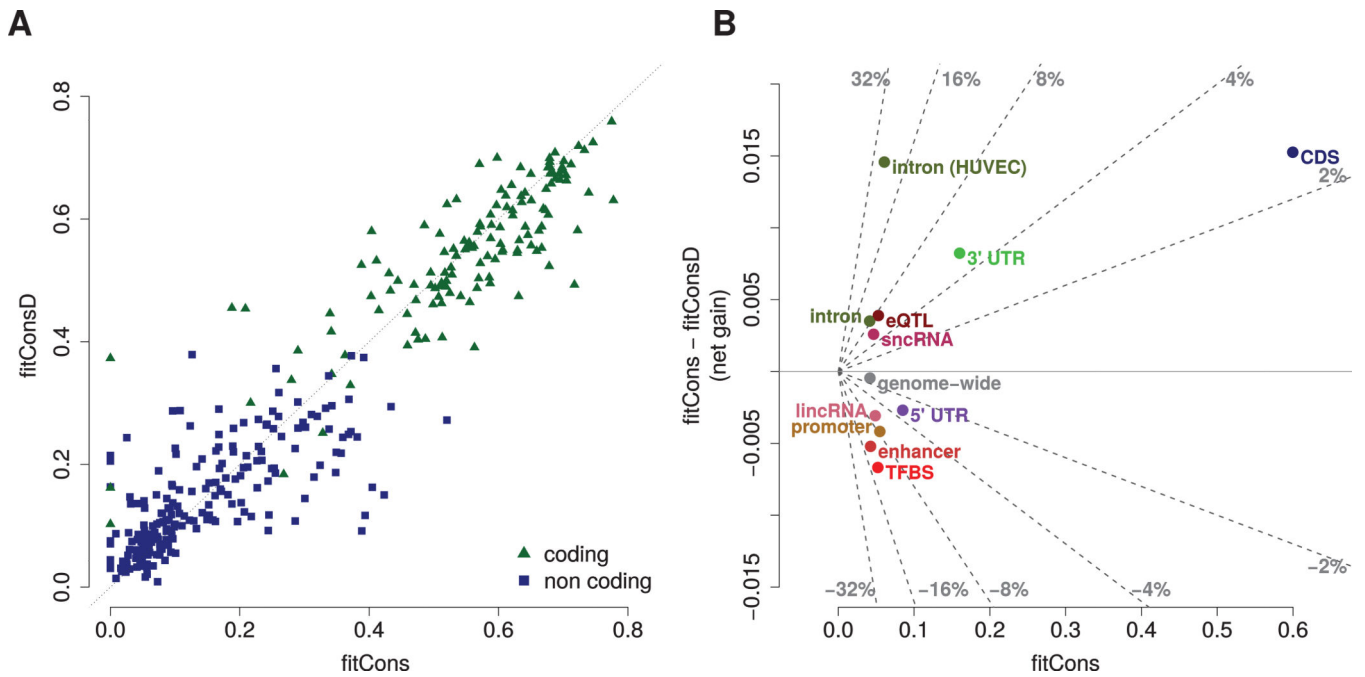


Figure 6.

Comparison between fitCons and fitConsD scores. FitConsD is an alternative estimate of fitness consequences, analogous to fitCons, but based on an estimator of the fraction of sites under natural selection that considers divergence patterns across four primate genomes (see Methods). (A) FitCons and FitConsD scores are shown for the clusters defined using functional genomic data from HUVEC. Scores are shown for the 348 clusters of size 10 kb or larger, distinguishing between coding clusters (green triangles) and non-coding clusters (blue squares). Both sets of scores are corrected by subtracting the possible contribution from model misspecification (see Methods). Correlation between the two sets of scores is high ($R^2 = 0.88$) overall, and is somewhat higher for coding ($R^2 = 0.69$) than for noncoding ($R^2 = 0.51$) clusters. (B) The net gain in the fraction of sites under selection on population genetic time scales relative to primate-divergence time scales, computed by subtracting average fitConsD scores from average fitCons score for different classes of functional elements (negative values imply net loss). Net gain is plotted against average fitCons score, and lines of constant slope radiating from the origin represent constant values of a “net gain rate” per functional site, computed as $NGR = (\text{fitCons} - \text{fitConsD})/\text{fitCons}$. The NGR is small (10%) for almost all annotation classes considered, with the main exception being the introns of active genes ($NGR > 20\%$; see “intron (HUVEC)”), which are enriched in clusters that exhibit an absence of DNase-seq or RNA-seq signal and chromatin modifications suggesting transcriptional elongation.