

Phylogenetic Modeling of Regulatory Element Turnover Based on Epigenomic Data

Noah Dukler ^{1,2}, Yi-Fei Huang,³ and Adam Siepel ^{*}1

¹Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY

²Physiology, Biophysics, and Systems Biology, Weill Cornell Medical College, New York, NY

³Department of Biology and Huck Institute of Life Sciences, Pennsylvania State University, University Park, PA

^{*}Corresponding author: E-mail: asiepel@cshl.edu.

Associate editor: Sergei Kosakovsky Pond

Abstract

Evolutionary changes in gene expression are often driven by gains and losses of *cis*-regulatory elements (CREs). The dynamics of CRE evolution can be examined using multispecies epigenomic data, but so far such analyses have generally been descriptive and model-free. Here, we introduce a probabilistic modeling framework for the evolution of CREs that operates directly on raw chromatin immunoprecipitation and sequencing (ChIP-seq) data and fully considers the phylogenetic relationships among species. Our framework includes a phylogenetic hidden Markov model, called epiPhyloHMM, for identifying the locations of multiply aligned CREs, and a combined phylogenetic and generalized linear model, called phyloGLM, for accounting for the influence of a rich set of genomic features in describing their evolutionary dynamics. We apply these methods to previously published ChIP-seq data for the H3K4me3 and H3K27ac histone modifications in liver tissue from nine mammals. We find that enhancers are gained and lost during mammalian evolution at about twice the rate of promoters, and that turnover rates are negatively correlated with DNA sequence conservation, expression level, and tissue breadth, and positively correlated with distance from the transcription start site, consistent with previous findings. In addition, we find that the predicted dosage sensitivity of target genes positively correlates with DNA sequence constraint in CREs but not with turnover rates, perhaps owing to differences in the effect sizes of the relevant mutations. Altogether, our probabilistic modeling framework enables a variety of powerful new analyses.

Key words: phylogenetics, *cis*-regulation, evolution.

Introduction

It is now well established that the evolution of form and function is often driven by mutations in *cis*-regulatory elements (CREs), particularly in multicellular eukaryotes having complex programs for regulating gene expression (Wittkopp and Kalay 2012; Siepel and Arbiza 2014; Villar et al. 2014; Reilly and Noonan 2016). In humans, patterns of genetic polymorphism, patterns of interspecies divergence, and the results of genome-wide association studies all indicate that a majority of phenotype- or fitness-influencing nucleotides fall in non-coding sequences and likely function in gene regulation (Waterston et al. 2002; Siepel et al. 2005; Hindorff et al. 2009; Arbiza et al. 2013; Gusev et al. 2014; Battle et al. 2017). Although the mutations that underlie regulatory evolution sometimes have subtle effects on, say, protein-DNA binding or chromatin accessibility, in many of the best-known cases, they instead alter gene expression through the gain or loss in activity of a whole CRE (Tournamille et al. 1995; Prabhakar et al. 2008; McLean et al. 2011). A number of lines of evidence indicate that this gain-and-loss process—sometimes called “turnover”—occurs at substantial rates over evolutionary time (Dermitzakis and Clark 2002; Moses et al. 2006; Doniger and Fay 2007; Wang et al. 2007; Bradley et al. 2010; Schmidt et al. 2010; Jones et al. 2012;

Cotney et al. 2013; Kasowski et al. 2013). Indeed, the evolutionary dynamics of this process appear to play out over considerably shorter time periods than those for other critical functional elements, such as protein-coding genes, microRNAs, or long noncoding RNAs (Weirauch and Hughes 2010; Danko et al. 2018).

There have been numerous attempts to model the evolutionary dynamics of CRE turnover at the level of the primary DNA sequence (Dermitzakis and Clark 2002; Moses et al. 2006; Siepel et al. 2006; Bullaughey 2011; Tugrul et al. 2015). However, characterizing this process at the sequence level is fundamentally challenging owing to limitations in the inference of regulatory function from the DNA sequence alone. During the past 15 years, new technologies for collecting high-throughput epigenomic data—such as chromatin immunoprecipitation and sequencing (ChIP-seq) data for transcription factors or histone modifications—have provided a path forward, by more directly indicating similarities and differences across species in molecular phenotypes that are closely related to *cis*-regulatory activity. A considerable number of comparative epigenomic studies have now been carried out in a variety of organisms, including studies based on transcription factor binding (Odom et al. 2007; Bradley et al. 2010; Schmidt et al. 2010; Paris et al. 2013; Wong et al. 2015;




FIG. 1. Illustration of modeling framework. (A) ChIP-seq data are aligned separately to the reference genome for each species then converted to the coordinate system of the human (hg38) genome using the liftOver program (see Materials and Methods). Only regions of apparent one-to-one orthology are considered, based on synteny. Cis-regulatory elements (CREs) that are active in one or more species are then identified using epiPhyloHMM. Finally, the dynamics of CRE turnover within these elements are modeled using phyloGLM, which accounts for the associations between various genomic features and local rates of gain and loss. (B) Both epiPhyloHMM and phyloGLM use a core “two-state” phylogenetic model in which the presence ($s_i = 1$) and absence ($s_i = 0$) of CREs are allowed to change in a branch length-dependent manner along a fixed phylogeny, according to a continuous-time Markov model. The model is defined by an instantaneous rate matrix Q (dashes indicate values required for rows to sum to zero). The conditional probabilities of the raw ChIP-seq read counts at the tips of the tree (x_i) given the corresponding state (s_i) are modeled using negative binomial (NB) distributions. The color intensities for the “0” and “1” boxes are proportional to the probability of each state. π , stationary frequency of CRE presence; γ , gain/loss rate; t_i , length of branch i .

Jubb et al. 2016), particular histone modifications (Bernstein et al. 2005; Mikkelsen et al. 2010; Xiao et al. 2012; Prescott et al. 2015; Villar et al. 2015), chromatin accessibility or chromatin contacts (Shibata et al. 2012; Vietri Rudan et al. 2015; Maher et al. 2018), DNA methylation (Qu et al. 2018), and nascent transcription (Danko et al. 2018) (partially reviewed in Marinov and Kundaje 2018). Among other findings, these studies have confirmed generally rapid rates of CRE gain and loss, and demonstrated that turnover rates are substantially higher in enhancers than in promoters, that depth of conservation correlates with various measures of functional impact, and that the evolutionary stability of gene expression correlates with the complexity and conservation of the local CRE architecture. However, with rare exceptions (Qu et al. 2018; Yang et al. 2018) (see Discussion), the available comparative epigenomic data sets have been analyzed using heuristic, model-free methods that do not consider the phylogenetic relationships of the species under study or the uncertainty in epigenomic data.

In this article, we introduce new model-based inference methods that address these deficiencies by fully accounting for the species phylogeny as well as the relationship between element activity and raw ChIP-seq read counts. Our methods can also account for correlations of turnover rates with local features along the genome sequence—such as gene expression patterns across tissues, distance to the transcription start site (TSS), or DNA sequence conservation—and they are efficient enough to be applied to genome-wide data sets. As a proof of concept, we apply these methods to previously published ChIP-seq data for the H3K4me3 and H3K27ac histone modifications in liver tissue across a phylogeny of nine

mammals (Villar et al. 2015). As described in detail below, we confirm several previous findings regarding relative rates of turnover in enhancers and promoters, and correlations with gene expression patterns and local regulatory architecture. In addition, we examine differences between patterns of constraint at the DNA sequence and CRE turnover levels, and find evidence suggesting that they reflect differences in the effect sizes of the relevant mutations.

Results

General Approach

Our approach for analyzing multispecies epigenomic data consists of three major stages (fig. 1A). First, we carry out a series of preprocessing steps to summarize the ChIP-seq read counts for each species in a common coordinate system (based on version hg38 of the human reference genome), excluding genomic regions where we could not establish clear one-to-one orthology based on genomic synteny (see Materials and Methods for details). Second, we apply a newly developed probabilistic inference method, called epiPhyloHMM, to identify “active” regions based on the ChIP-seq read counts, working in the common coordinate system. At this stage, an “active” region is one containing CREs in any one or more species. This method accounts for the phylogenetic gain and loss process, as well as noise in the ChIP-seq data, at the same time as it predicts the locations of the elements. Third, we apply a new probabilistic modeling program, called phyloGLM, to describe the process of phylogenetic gain and loss in more detail, within the “active” regions identified by epiPhyloHMM. PhyloGLM conditions

on a rich set of genomic features, capturing their correlations with local rates of gain and loss.

Shared Phylogenetic Model

The epiPhyloHMM and phyloGLM programs both make use of the same core “two-state” probabilistic model for the gain and loss of CREs along the branches of a phylogeny. Moreover, in both cases, this model also describes the generation of read counts that are reflective of CRE presence or absence at the tips of the tree. Thus, it serves as a generative model for multispecies read counts that can be fitted to the raw data by maximum likelihood (fig. 1B). In this article, we focus on read counts from ChIP-seq experiments, but the model can easily be extended to other data types, such as those arising from DNase-seq, ATAC-seq, or PRO-seq experiments (see Discussion).

The phylogenetic component of the model is a straightforward “two-state” presence/absence model (with state variables $s \in \{0, 1\}$) for CREs along the branches of a phylogeny. It assumes a tree topology is given together with nonnegative real-valued branch lengths. In practice, the tree and branch lengths can be obtained from the literature or estimated from sequence data (see Materials and Methods). The stochastic process for gains and losses is defined, in the usual manner, by a continuous-time Markov model with an instantaneous rate matrix Q , from which branch length-dependent turnover (gain/loss) probabilities can be obtained as $P(t) = \exp(Qt)$ for each branch length t (see Felsenstein 2004). The model has two free parameters: the stationary probability of CRE presence (π) and a single turnover rate parameter (γ), which together define a 2×2 reversible rate matrix Q (fig. 1B). Given data at the leaves of the tree, phylogenetic inference with this model can be accomplished using Felsenstein’s pruning algorithm (Felsenstein 1973, 1981) (see Materials and Methods).

Unlike with standard phylogenetic models for DNA sequences, however, the observed data here consists of epigenomic (typically ChIP-seq) read counts, which provide only an approximate indication of whether or not an active CRE exists in each species. We accommodated the uncertainty in read counts by borrowing from the literature on statistical peak-calling for ChIP-seq data (Zhang et al. 2008). In particular, we described both the probability of the observed read counts x_i in species i given an active CRE in that species, $P(x_i | s_i = 1)$, and the probability of the observed read counts given no active CRE, $P(x_i | s_i = 0)$, using negative binomial distributions (fig. 1B). Moreover, for the “active” model, we used a mixture of three negative binomial distributions to accommodate peaks of various heights (see Materials and Methods) (Anders et al. 2013; Love et al. 2014). We also adapted these emission distributions to accommodate missing data due to alignment gaps (see Materials and Methods). Altogether, this modeling approach allows us to perform multispecies peak-calling and phylogenetic inference simultaneously, accounting for uncertainty in both the locations of present day CREs along the genome and their presence/absence over evolutionary time.

epiPhyloHMM: Prediction of Multispecies CREs from Epigenomic Data

To address the problem of predicting “active” CREs, we made use of the framework of phylogenetic hidden Markov models, or phylo-HMMs (Felsenstein and Churchill 1996; Siepel and Haussler 2005). Phylo-HMMs are hidden Markov models whose hidden states are associated with phylogenetic models, which in turn, define distributions over columns in multiply aligned sequences of observations. Phylo-HMMs are sometimes called “space-time” models (Yang 1995) because they describe stochastic processes in both a spatial dimension, along the genome sequence, and a temporal dimension, along the branches of a phylogeny. In this case, the temporal (phylogenetic) models describe distributions over aligned ChIP-seq readcounts from multiple species, as described in the previous section. The spatial (hidden Markov) model, in turn, is designed to allow the identification of CREs with various patterns of presence/absence at the tips of the tree.

This hidden Markov model consists of a single “inactive” state and a set of states representing each possible presence/absence pattern (fig. 2A). Assuming most of the genome will be inactive, the transition model is sparse, with each active state being accessible only from the inactive state, and not from other active states (see Siepel et al. 2006 for a similar approach). It is completely defined by two free parameters, ρ_0 and ρ_1 (fig. 2A and Materials and Methods).

In practice, we constrain the complexity of the model by including states only for presence/absence patterns that are achievable by at most three gain/loss events along the branches of the phylogeny (see Materials and Methods and supplementary fig. S1, Supplementary Material online). The free parameters of both the phylogenetic model (π, γ) and the HMM (ρ_0, ρ_1) are fitted to aligned epigenomic data by maximum likelihood, and then active elements are called in the standard way, using the Viterbi algorithm (fig. 2B). The method generally performs well on simulated data (supplementary Methods and figs. S2–S6, Supplementary Material online).

Application of epiPhyloHMM to Histone-Modification Data for Nine Mammals

We applied epiPhyloHMM to recently published H3K4me3 and H3K27ac ChIP-seq data for liver tissue from mammals (Villar et al. 2015), preprocessing and aligning the data as outlined above (see also Materials and Methods). These data have the benefit of serving as highly general, if imperfect (Benton et al. 2019), indicators of regulatory function and of having been generated uniformly across species. We used data for 9 of the 20 species examined in Villar et al. (2015). We chose these nine species to provide good coverage of the placental mammals across a variety of timescales, by representing three separate mammalian clades (primates, rodentia, and carnivora) and an outgroup (opossum). The ingroup species were selected to include good-quality assemblies. We mitigated potential problems stemming from the lesser




Fig. 2. Prediction of multispecies *cis*-regulatory elements (CREs) using epiPhyloHMM. (A) State-transition diagram for the phylogenetic hidden Markov model. Each state represents a different combination of active and inactive elements in the observed species, corresponding to a different gain/loss scenario along the branches of the tree (gold: active; black: inactive). The fully inactive (background) state is shown on the “left” and the states representing various possible presence/absence patterns for active elements are grouped on the “right.” The red star indicates the state associated with the focal site in the cartoon data at top. Notice that each active state is accessible only from the inactive state, leading to a sparse transition matrix for the hidden Markov model. (B) Example of predictions obtained by applying epiPhyloHMM to H3K4me3 ChIP-seq data from Villar et al. (2015) in a region along human chromosome 1. Shown are the nine-species phylogeny (left) and the corresponding ChIP-seq read counts along the chromosome (right), with predicted elements highlighted in red boxes. For reference, annotated protein-coding genes and an antisense transcription unit are also shown (top).

quality of the opossum assembly by assigning a separate set of parameters to the outgroup branch.

This analysis produced an average of $\sim 16,000$ and $\sim 47,000$ elements per species for the H3K4me3 and H3K27ac marks, respectively (supplementary fig. S7, Supplementary Material online), with some variation across species owing to differences in data quality and alignability. The substantially greater abundance of H3K27ac elements was expected because the H3K27ac mark is associated with both active promoters and active enhancers, whereas the H3K4me3 mark is more specific to promoters. The two types of elements also had highly distinct distributions of state assignments, with the fully conserved state being the most frequent for the H3K4me3 mark but ranking much lower for the H3K27ac mark, beneath most single-species states (supplementary fig. S8, Supplementary Material online). This difference suggests substantially lower rates of turnover in promoters than enhancers (see below).

For comparison, we analyzed the same data with a heuristic pipeline similar to those employed by other researchers (Villar et al. 2015; Danko et al. 2018), which applied a peak-calling method (MACS2) separately in each species, and then mapped all elements to the human genome using liftOver. We found that epiPhyloHMM found the vast majority of regions identified by this alternative pipeline, as well as a number of new regions, which did indeed exhibit hallmarks of real regulatory function (supplementary figs. S9 and S10A and B, Supplementary Material online). Moreover, epiPhyloHMM tended to call fewer regions as active in a single species and estimated greater numbers of multispecies regions, by effectively sharing information across the phylogeny (supplementary fig. S10C, Supplementary Material online).

phyloGLM: Modeling of CRE Turnover Conditional on Local Genomic Features

We addressed the third stage in our pipeline—modeling of gain/loss dynamics conditional on genomic features such as nearby gene expression or sequence conservation—with a new program, called phyloGLM, that allows the free parameters of the our phylogenetic model (π and γ) to be determined by a function of genomic features through a generalized linear model (GLM; fig. 3). These genomic features are treated as being constant across the phylogeny, but they can be computed in any manner from any combination of species.

As shown below, this GLM-based design provides a rigorous framework for measuring the strength of association of individual genomic features with turnover rate, and for testing for differences in turnover rate between distinct groups of CREs (see Discussion).

Application of phyloGLM to Real Data

We applied phyloGLM to the genome-wide predictions from epiPhyloHMM, separately analyzing the H3K4me3 promoter data set and two subsets of the H3K27ac data that correspond to likely promoters and likely enhancers. Thus, we were able to compare the enhancer data set with two distinct promoter data sets, one of which (H3K27ac) included more abundant but less precise predictions than the other (H3K4me3). To set up the analysis, we first assigned each CRE a putative target gene from Ensembl (Zerbino et al. 2018) using simple distance-based rules, which essentially associated each CRE with the closest TSS of a gene but discarded elements that could plausibly be associated with more than one gene (see Materials and Methods and




Fig. 3. The phyloGLM model. phyloGLM combines the general modeling framework of [figure 1](#) with a generalized linear model (GLM) to account for the influence of local genomic covariates (such as gene expression levels or sequence conservation scores) on the turnover process. In each genomic bin j , the turnover rate γ_j and the equilibrium frequency for active elements π_j are determined by a logistic function f (see Materials and Methods) applied to a linear combination of covariate values C_j with weight vectors θ_γ and θ_π , respectively. The phylogeny and hypothetical read counts x_j for bin j are shown at the left, and hypothetical covariate values are represented by colored squares to the right. The shades of gray to the right of the phylogeny illustrate relative values of the conditional likelihoods $P(x_j|s)$.

[supplementary fig. S11, Supplementary Material](#) online). In addition, based on proximity to the nearest TSS, we classified H3K27ac CREs as likely promoter (within 1.5 kb) or enhancer (within 100 kb) elements, and we similarly classified H3K4me elements as likely promoters (within 1.5 kb) or discarded them. Finally, we associated each CRE with a collection of genic and *cis*-regulatory features that could potentially impact its turnover rate ([supplementary fig. S12, Supplementary Material](#) online). Broadly, these features described the numbers of CREs associated with the target gene, the expression patterns and annotated function of that gene, and measures of evolutionary constraint on the local DNA sequence. After removing all elements with incomplete covariate data, we were left with 5,368 H3K4me3 promoters, 7,220 H3K27ac promoters, and 25,673 H3K27ac enhancers for further analysis. These features were somewhat correlated with one another, but most correlations were weak ([supplementary fig. S13, Supplementary Material](#) online). We fitted phyloGLM separately to these three sets of elements, estimating all free parameters by maximum likelihood and conditioning on a phylogeny with branch lengths based on published estimates of divergence times in millions of years ([Kumar et al. 2017](#)). In all cases, we separately parameterized the branch to the outgroup (opossum), on which gains and losses are difficult to distinguish, to avoid skewing the other parameter estimates.

This model permitted us to compute the expected total numbers of gains and losses along each branch of the phylogeny, conditional on the data and the fitted model. Similar patterns of gain and loss were observed for the H3K4me3 and H3K27ac promoters, with fewer total events in the H3K4me3 elements (0.97 vs. 1.22 events per element on an average; [supplementary fig. S15, Supplementary Material](#) online). For H3K27ac elements, we found that the overall rates of gains and losses are fairly similar, with

somewhat more gains than losses, both within promoters and within enhancers ([supplementary fig. S16, Supplementary Material](#) online). However, the total rate of turnover for enhancers appears to be approximately twice that of promoters, with 2.47 events per element compared with 1.22 for the H3K27ac elements. The numbers of expected events per branch were roughly proportional to the branch lengths, with long branches tending to be assigned more events than short branches. The gain/loss proportions were somewhat variable across branches, but this variation likely reflects a combination of true differences and biases from human-referenced alignments and differences in ChIP-seq data abundance and quality across species (see Discussion).

A comparison of the distributions of numbers of events per CRE provided further support for a roughly 2-fold higher rate of turnover at enhancers than promoters, with median values of 0.0075 and 0.0031 events per million years (My), respectively, for the H3K27ac data ([fig. 4A](#); $P < 2.2 \times 10^{-16}$; Wilcoxon signed-rank test). From these distributions and the estimated phylogeny, it was also possible to estimate a distribution of the “half-life” (time required for half of active elements to be lost) for each type of CRE. For the H3K27ac data, the median half-life for enhancers is 130 My and that for promoters is 552 My ([fig. 4B](#); see Materials and Methods). These estimates are substantially lower than previous estimates of 296 and 939 My, respectively ([Villar et al. 2015](#)). However, our estimate of the median turnover rate for promoters based on the less noisy H3K4me3 data set was $\sim 30\%$ lower, at 0.0022 events per million years, corresponding to a half-life estimate of 937 My, in much better agreement with the corresponding previous estimate. Thus, it seems likely that our H3K27ac turnover rate estimates are substantially inflated by the lower resolution ChIP-seq data (see Discussion).




Fig. 4. Enhancers show greater rates of turnover than promoters. (A) Distributions of estimated turnover rates per bin (γ) across all elements, based on the full phyloGLM model. Results are shown for separate analyses of promoters based on the H3K4me3 mark, promoters based on the H3K27ac mark, and enhancers based on the H3K27ac mark (see Materials and Methods). The phylogeny and branch lengths (in millions of years) were obtained from Kumar et al. (2017). The branch to the opossum outgroup was excluded in rate estimation. To simplify the visualization, a small number of elements (41 of 38,170) with turnover rates > 0.044 were omitted. The bar in each violin plot represents the median. (B) Similar distributions of times for half of all active elements to decay to an inactive state (half-life).




Fig. 5. Marginal effects of gene expression-related features on turnover rates. Estimates of the coefficients in the generalized linear model (θ_γ) for four gene expression-related features, for promoters based on H3K4me3 (left) and both promoters and enhancers based on H3K27ac (right). A positive value indicates that a feature increases the turnover rate and a negative value indicates that it decreases the turnover rate. The branch to the outgroup is excluded from rate estimation. Error bars represent approximate 95% confidence intervals (see Materials and Methods). “*” indicates a significant difference from zero ($P < 0.05$ after Benjamini–Hochberg correction) based on a likelihood ratio test.

Gene Expression-Related Features Associated with Turnover Rates

In addition to allowing us to characterize overall rates of turnover, our GLM-based framework enabled us to examine the strength and directionality of the association between each of the genomic features we considered and the rates of turnover at enhancers and promoters. We begin by considering these relationships for gene expression-related

features, and examine the remaining features in subsequent sections.

For promoters, three gene expression-related covariates had statistically significant associations with turnover rate: the level of expression of the target gene in the liver (the assayed tissue here), the number of tissues in which the target gene was expressed, and the cross-tissue expression dispersion (fig. 5). For liver expression and the number of tissues,




Fig. 6. Marginal effects of remaining features on turnover rates. Estimates of the coefficients in the generalized linear model (θ_i) for four additional features, for promoters based on H3K4me3 (left) and both promoters and enhancers based on H3K27ac (right). A positive value indicates that a feature increases the turnover rate and a negative value indicates that it decreases the turnover rate. The branch to the outgroup is excluded from rate estimation. Error bars represent approximate 95% confidence intervals (see Materials and Methods). “*” indicates a significant difference from zero ($P < 0.05$ after Benjamini–Hochberg correction) based on a likelihood ratio test.

increased values of the covariate were associated with significantly decreased turnover rates. These observations are broadly consistent with a variety of previous analyses that have indicated that CREs associated with high levels of expression in the tissue of interest or with broad expression patterns across tissues tend to experience elevated levels of constraint (Arbiza et al. 2013; Khan et al. 2013; Villar et al. 2015; Berthelot et al. 2018). The positive correlation with cross-tissue expression dispersion (significant for H3K4me3 only) appears to reflect a similar trend.

The results for enhancers were generally similar to those for promoters, but tended to be somewhat weaker, likely in part owing to the difficulty of correctly linking enhancers with target genes. One notable difference for enhancers was that the number of tissues in which the target gene was expressed was not significantly associated with turnover rate. This difference might result in part from tissue-general “housekeeping” tending to have fewer enhancers than tissue-specific genes (Zabidi et al. 2015). In addition, housekeeping genes are likely enriched for proximal enhancers, which tend to be excluded by our filters.

Additional Features Associated with Turnover Rates

The remaining genomic features describe either measures of sequence constraint of the CRE (phastCons; Siepel et al. 2005) or the gene (pLI; Lek et al. 2016), or aspects of the local regulatory “architecture” of each target gene, including the numbers of enhancers and promoters and the distance of each enhancer from the TSS (fig. 6). Both enhancers and promoters displayed a negative correlation between CRE sequence conservation, as measured by phastCons, and turnover rate. As

previously noted (Villar et al. 2015; Berthelot et al. 2018; Danko et al. 2018), this observation indicates that elements that are more constrained at the DNA sequence level are also more resistant to evolutionary gain and loss. Interestingly, however, we observed no significant correlation between constraint against loss-of-function variants in the gene, as measured by pLI scores, and turnover rates of associated enhancers or promoters (see next section).

Among the architectural features, the strongest correlate at the enhancer level is the distance to the TSS, a quantity that is positively associated with turnover rate. As has been noted in several recent studies (Villar et al. 2015; Berthelot et al. 2018; Danko et al. 2018), this increased constraint against turnover on enhancers close to the TSS likely reflects an enrichment for genuine enhancer–gene interactions and direct influence on the expression of the target gene. Another observation that echoes a previous finding is that the number of enhancers per gene is positively correlated with the enhancer turnover rate but negatively correlated with the promoter turnover rate. As previously noted (Danko et al. 2018), this observation suggests that larger ensembles of enhancers associated with the same target gene tend to impose additional constraints against promoter turnover, but nevertheless to relax constraint on each of the enhancers themselves, perhaps because each enhancer is less essential to the overall regulatory architecture of the locus (see also Frankel et al. 2010; Perry et al. 2010; Osterwalder et al. 2018). We also observed significant effects for several top-level Reactome categories, suggesting that biological function may provide additional information about regulatory constraint (supplementary fig. S17, Supplementary Material online). Together, these




FIG. 7. Turnover rates and sequence conservation at CREs for dosage-sensitive and insensitive genes. (A) Estimated coefficients θ_γ for indicators for the Reactome (Fabregat et al. 2016) categories “Metabolism” (less dosage-sensitive) and “Gene expression (Transcription)” (more dosage-sensitive) at associated promoters and enhancers (both based on H3K27ac). Error bars represent approximate 95% confidence intervals (see Materials and Methods). “*” indicates a statistically significant difference from zero ($P < 0.05$ after Benjamini–Hochberg correction) based on a likelihood ratio test. The differences between the two categories of genes are not significant (ns). (B) Mean phastCons scores for vertebrates (100-way alignment in UCSC Genome Browser) in H3K27ac-based promoters of the same two categories of genes, stratified by the expected number of turnover events per CRE. Notches correspond to $1.58 \cdot \text{IQR} / \sqrt{n}$ indicating an approximate 95% confidence interval for the median. “*” indicates a significant difference between gene categories ($P < 0.05$, t -test with Benjamini–Hochberg correction). (C) Same as (B) but for H3K27ac-based enhancers. For visual clarity, outliers ($> 1.5 \cdot \text{IQR}$ from the hinge) are not shown in (B) and (C).

observations suggest that CREs evolve in a manner that is strongly dependent on the local regulatory context in which they appear.

Differences between DNA Sequence and Epigenetic Conservation in CREs

The correlation of turnover rate with the phastCons scores of CREs but not with the pLI scores of target genes is curious, but what does it signify? The pLI score for a gene measures the probability of intolerance to (generally heterozygous) loss-of-function mutations in that gene, as inferred from patterns of variation in ultradeep human exome-sequencing data (Lek et al. 2016). pLI scores can be used to differentiate between haploinsufficient and haplosufficient genes or, similarly, between dosage-sensitive and insensitive genes (Lek et al. 2016) (although, strictly speaking, the scores are directly informative only about the strength of selection acting on heterozygotes; Fuller et al. 2019). By contrast, phastCons scores simply measure a reduction in fixed derived alleles, and do not effectively differentiate among various forms of negative selection. Therefore, the observed difference in correlation suggests that gains and losses of CREs are generally deleterious but perhaps do not depend strongly on the dominance or dosage properties of target genes.

To investigate this issue further, we examined CRE turnover rates at the promoters of two classes of genes that serve as proxies for dosage insensitivity and sensitivity, respectively: genes that encode proteins involved in metabolism such as enzymes (whose action tends to be relatively insensitive to protein abundance) and genes that encode proteins that regulate gene expression such as transcription factors (which tend to be more sensitive to abundance) (Wilkie 1994; Kondrashov and Koonin 2004). We found no significant difference between these classes of genes in the turnover rates of CREs in either promoters or enhancers (fig. 7A), further

supporting the idea that turnover rates have little dependency on dosage or dominance. Interestingly, however, when we condition on the expected number of turnover events at each CRE and focus on the CREs that have undergone the fewest events, we do observe increased sequence conservation at the more dosage-sensitive regulatory genes. This difference is observed both for promoters (fig. 7B) and enhancers (fig. 7C), although it is statistically significant for promoters only. It suggests that, although both classes of genes are similarly resistant to mutations that result in the complete gain and loss of elements (fig. 7A), dosage-insensitive genes are more tolerant of nucleotide substitutions that do not result in complete gain or loss events (fig. 7B and C). These observations illustrate how the evolutionary dynamics of CRE gain and loss may differ from those for nucleotide substitutions owing to the pronounced effects of turnover events on gene expression, and, more generally, how patterns of evolutionary constraint across the genome may depend on the effect sizes of mutations.

Discussion

In this article, we have introduced a new probabilistic modeling framework for inferring the dynamics of CRE gain and loss, which accounts for phylogenetic correlations among species, uncertainty in peak calls from ChIP-seq data, and the influence of local genomic features on turnover rates. We have applied our methods to H3K4me3 and H3K27ac ChIP-seq data (Villar et al. 2015). These histone modifications are somewhat noisy and imperfect indicators of regulatory function (Benton et al. 2019), but they have the important advantages of being highly general (i.e., not depending on any particular regulatory role or transcription factor) and of having been assayed uniformly across species, making them a natural starting point for an evolutionary analysis. We find support for a number of previously reported results, including a

substantially higher rate of turnover in enhancers than promoters, negative correlations of turnover rate with DNA sequence conservation, expression level, and tissue breadth, positive correlations with distance from the TSS, and a strong dependency on features of the local regulatory architecture such as number of enhancers per gene. Overall, we find that enhancers are gained and lost at about twice the rate of promoters during mammalian evolution, with median rates of 0.0075 and 0.0031 events per element per million years, respectively, based on the H3K27ac data.

In addition, we made use of our modeling framework to examine an apparent lack of correlation between turnover rates at CREs and the haploinsufficiency or dosage sensitivity of target genes, as measured approximately using pLI scores (Lek et al. 2016). We found that turnover rates were significantly negatively correlated with DNA sequence conservation in CREs, suggesting that both whole gain/loss events and nucleotide substitutions are deleterious, but that turnover rates were not correlated with pLI scores, suggesting that gain/loss events are no more deleterious at haploinsufficient/dosage-sensitive genes than at haplosufficient/dosage-insensitive genes. However, when we conditioned on the expected number of turnover events at each CRE, a positive correlation became evident between sequence conservation and dosage sensitivity at low-turnover CREs (fig. 7). We interpret this result as indicating that DNA substitutions that do not cause complete gain or loss events are more easily tolerated at dosage-insensitive genes than at dosage-sensitive genes. These mutations of small effect at dosage-insensitive genes may be allowed to accumulate and perhaps compensate for one another, permitting drift in CRE sequences as long as it does not cause the gain or loss of a whole element. By contrast, gains and losses of entire CREs have sufficiently large effects that they are deleterious at both dosage-sensitive and -insensitive genes. Thus, the divergence between epigenetic and sequence constraint is potentially informative about the mode of selection at each locus. These observations may help to explain previous reports of CREs that display conservation of epigenetic marks but not the DNA sequence (Ludwig et al. 2000; Fisher et al. 2006; Hare et al. 2008; Yang et al. 2015; see also Duque et al. 2014; Khoueiry et al. 2017).

For various reasons, we have approached the problems of identifying regulatory elements and modeling their evolution separately, using the epiPhyloHMM and phyloGLM programs, respectively. This strategy allows us to address the problem of segmenting the genome into “active” and “inactive” regions in a relatively efficient manner, using a simpler model, and then characterize the turnover process using a richer model that conditions on a diverse collection of genomic features. It also has practical advantages in terms of modularity of software development and efficient processing of genome-wide data. At the same time, this strategy has the limitation that the richer evolutionary model implemented in phyloGLM is not exploited in element identification, which in principle, could result in loss of power. Still, this limitation does not appear to be of major practical importance because identifying regions containing active elements turns out to be fairly

straightforward. A related limitation is that we analyze the genome in bins of fixed size, which occasionally results in spurious inferences of turnover when the boundaries of the bins align poorly with the locations of peaks. Nevertheless, this simple approach is generally fairly effective for pooling read counts along the genome and accommodating limitations in the genomic resolution of peak locations (see below).

Notably, our flexible phyloGLM design allows not only for improved modeling of evolutionary dynamics but also direct assessment of hypotheses about how these dynamics depend on various aspects of genomic context, such as gene expression, local regulatory architecture, and sequence conservation. As phyloGLM considers all covariates together in a single model, we avoid the need for complex post hoc analyses, for example, that require matching of foreground and background sets of elements in terms of relevant covariates. Importantly, phyloGLM is not designed to identify individual elements that have undergone gain and loss events, a task for which it would have weak power owing to noisy per-site data. Instead, its purpose is to pool information across the genome in order to enable hypotheses about groups of elements to be tested. Similar approaches have been used for the estimation of dN/dS rates (Meyer and Wilke 2013) and probabilities of fitness consequences for new mutations (Huang et al. 2017; Huang and Siepel 2019), but to our knowledge, this approach has not been previously employed in the study of CRE evolution.

Another strength of phyloGLM and epiPhyloHMM is that both programs use the same modular emission model to describe the relationships between epigenetic data and CRE state. Thus, it would be straightforward to adapt them to consider alternative epigenomic data types, possibly in combination. As long as the phylogenetic model has two states (active and inactive), the remaining portions of the existing programs could be used essentially without change. They could further be extended to consider additional states—for example, to allow for various combinations or various levels of epigenomic signals—but such an approach may be prohibitive in terms of runtime, unless heuristic methods were used to limit the portion of the genome under consideration.

Two other recently published methods (Qu et al. 2018; Yang et al. 2018) have addressed the problem of inferring evolutionary dynamics from multispecies epigenomic data using strategies that are similar to ours, but are also different in key respects. Importantly, both of these methods avoid separating the element identification and evolutionary inference problems—a decision that has potential advantages, provided the data have sufficiently high resolution to avoid overfitting, but that is also costly in computational efficiency.

The first method (Yang et al. 2018) directly models an evolving continuous signal (replication timing, in their application) along a collection of aligned genomes using an elegant combination of a branching Orstein–Uhlenbeck process along a phylogeny and a hidden Markov model along the genome sequence, which is fitted to genome-wide data by expectation maximization. This approach appears to be quite powerful but it differs from ours in that it focuses on direct

modeling of a continuous molecular phenotype, rather than on describing a relationship between a discrete property of interest (such as transcription factor binding or CRE activity) and functional genomic data describing that property (such as ChIP-seq read counts). The strategy has potential advantages when the property of interest truly is continuous, but could also tend to overfit a complex signal. The second method (Qu et al. 2018) is more similar to ours in that it does distinguish between discrete “positive” and “negative” states (in this case, reflecting DNA methylation status), again using a combination of a hidden Markov model and a phylogenetic model. The structure of this HMM is more complex than ours, however, and requires Monte Carlo methods for fitting. The setting for this method is also different from ours in that the whole genome bisulfite sequencing data being analyzed appear to provide a higher resolution, more precise readout of the feature of interest, avoiding some of the uncertainty of peak-calling from ChIP-seq data. More generally, modeling of comparative epigenomic data remains an active area of research, with a number of newly developed methods that make similar but complementary modeling assumptions, and more work will be needed to find which approaches are best suited for various applications of interest.

A major problem faced by all current modeling approaches—and a reason why model-free methods have often been used instead—is the fundamental imprecision of comparative epigenomic data. In most data sets, there is considerable uncertainty not only in the strength of the signal along the genome but also in the precise genomic position and breadth of that signal (i.e., in peak height, location, and width, in the context of ChIP-seq data). This uncertainty is compounded by errors in alignment and orthology identification between species. Evolutionary models must therefore strike a balance between getting the most out of the available data, and avoiding biases that come from assuming unrealistic levels of precision or resolution. Indeed, in practice, all of the current modeling approaches have required the use of extensive heuristics and filters before and/or after they are applied. In our case, the imprecision of the data resulted in a tendency to fragment individual elements into multiple predicted segments, for example, because peaks did not align well in width or position across species. We attempted to mitigate this problem by postprocessing the epiPhyloHMM predictions with heuristic rules that joined and filtered elements (see Materials and Methods). Nevertheless, these misalignment and fragmentation issues undoubtedly produced some upward bias in our estimates of turnover rate. This effect most likely drives the substantial elevation in rate for the promoters based on the noisier H3K27ac data as compared with the more precise H3K4me3 data (supplementary fig. S14, Supplementary Material online). More work will be needed both to improve the precision of comparative epigenomic data and to accommodate uncertainty in models of evolutionary dynamics.

Two other limitations of our analysis that are broadly shared with other comparative genomics studies concern the use of reference-based multiple alignments and proximity-based rules for associating CREs with target genes.

Our use of human-referenced multiple alignments (which are also used in Qu et al. 2018; Yang et al. 2018) prevents us from analyzing portions of the other genomes that do not align to the human genome, and therefore creates a general bias toward “gains” in human and closely related species, and toward “losses” in more distant parts of the tree, as is evident in a close inspection of supplementary figures S15 and S16, Supplementary Material online. This same limitation makes it infeasible to test for differences in gain or loss rates across branches or clades of the phylogeny, because alignment-induced biases would likely overwhelm real biological differences. There has been some progress in recent years toward generalized reference-free multiple alignment methods (Paten et al. 2011; Armstrong et al. 2019) but much more work is needed on this important problem.

Similarly, the linking of CREs, particularly enhancers, with target genes is another fundamental unsolved problem that pervades many genomic analyses. In our case, it is likely that a substantial fraction of enhancers are mis-assigned a target gene, with downstream effects on a number of our analyses (e.g., figs. 5–7). Experimental work to link enhancers to the correct target genes, either via 3D-chromatin capture (Sanyal et al. 2012; Jin et al. 2013; Mifsud et al. 2015), or large-scale genome editing (Fulco et al. 2016), will help to improve this issue over time.

Materials and Methods

ChIP-seq Data Preparation

All ChIP-seq data were obtained from Villar et al. (2015). Reads were aligned to the reference genome for each species (Waterston et al. 2002; Rat Genome Sequencing Project Consortium 2004; Lindblad-Toh et al. 2005, 2011; Mikkelsen et al. 2007; Yan et al. 2011; Marmoset Genome Sequencing and Analysis Consortium 2014; Peng et al. 2014) (obtained from the UCSC genome browser) using bowtie2 (v2.2.9) (Langmead and Salzberg 2012). The reference genomes were hg38 (human), rheMac3 (macaque), callJac3 (marmoset), mm10 (mouse), rn6 (rat), musFur1 (ferret), canFam3 (dog), felCat8 (cat), and monDom5 (opossum). Each read was summarized by the single base at the center of the read and converted to the human reference genome using the liftOver utility and the best reciprocal chains supplied by the UCSC genome browser (Hinrichs et al. 2006) (hg38). A coverage map was computed from the converted reads and summed to get the total number of reads per 250-bp bin. Regions that could be aligned to the human genome from one or more other species were combined if they were less than 50 kb apart and were expanded by 5 kb to either side to create genomic blocks within which to run the phylo-HMM. Sections of the human genome that were not covered by one of these blocks were excluded from the analysis, leaving 2.83 Gb for further analysis.

Model for Peak-Calling

Our peak-calling model has two versions: the full model used in epiPhyloHMM and a simpler two-state model that is applied separately to the data for each species in a preprocessing

step (as detailed below). The two versions are the same except for the state space. In the two-state model, the probability of the data in bin j , for presence/absence state $p \in \{0, 1\}$ is given by a negative binomial mixture model:

$$P(\tilde{\mathbf{x}}_j | \tilde{\boldsymbol{\mu}}_p, \tilde{\mathbf{w}}_p, \gamma_j, \tilde{\boldsymbol{\lambda}}) = \sum_m w_{p,m} \cdot \prod_r \text{nbinom}(x_{j,r} | \mu_{p,m} \gamma_j \lambda_r, \sigma_{p,m,j,r}).$$

Here, $\tilde{\mathbf{x}}_j$ is a vector of the read counts for each replicate r , $w_{p,m}$ is the weight of mixture component m for state p (such that $\sum_m w_{p,m} = 1$), $\mu_{p,m}$ is the mean count for state p and component m , $\sigma_{p,m,j,r}$ represents the dispersion parameter for the negative binomial distribution (see below), γ_j is the fraction of bases in bin j that are aligned to the human reference genome, and $\tilde{\boldsymbol{\lambda}} = \{\lambda_r\}$ is a set of scaling factors that account for the sequencing depth of each replicate r . As noted in the Results section, we use a three-component mixture model for the “presence” state ($p = 1$) and a single component for the “absence” state ($p = 0$). The scaling factor λ_r is calculated as:

$$\lambda_r = \frac{1}{\alpha} \sum_j x_{j,r},$$

where $\alpha = \max_r(\sum_j x_{j,r})$. To account for differences in the dispersion of read counts for different mean depths, we make use of the model from DESeq2 (Love et al. 2014), where the dispersion of the distribution σ is defined as a function of the mean μ and two free parameters, θ_1 and θ_2 :

$$\sigma(\mu, \theta_1, \theta_2) = \frac{\theta_1}{\mu} + \theta_2.$$

We use the DESeq2 software to estimate θ_1 and θ_2 , after subsampling the genome to obtain roughly similar numbers of low-, medium-, and high-coverage sites. We then calculate the dispersion per data-point as,

$$\sigma_{p,m,j,r} = \frac{\theta_1}{\mu_{p,m} \gamma_j \lambda_r} + \theta_2.$$

This strategy accounts for effects of sample library depth and cross-species alignability on the expected read-count depth for each state p , mixture component m , replicate r , and bin j .

The transition model is a simple two-state model with auto-correlation parameters ρ_1 for the peak state and ρ_0 for the background state (fig. 2).

Hidden Markov Model

The hidden Markov model used by epiPhyloHMM includes a set of states $\{s_1, \dots, s_E\}$ representing all patterns of CRE presence and absence at the tips of the tree, up to a maximum number of gain/loss events (three, in our application; supplementary fig. S1, Supplementary Material online). Conditional on the state (and, implicitly, on the gain/loss history), the read-counts at the tips of the tree are independent. Thus, the emission probability for the observed data in

state s_e at site j is given by a product over tips (species) t and presence/absence states p :

$$P(\tilde{\mathbf{x}}_j | s_e, \tilde{\boldsymbol{\mu}}, \tilde{\mathbf{w}}, \gamma_j, \tilde{\boldsymbol{\lambda}}) = \prod_t \prod_{p \in \{0,1\}} P(\tilde{\mathbf{x}}_{jt} | \tilde{\boldsymbol{\mu}}_{pt}, \tilde{\mathbf{w}}_{pt}, \gamma_{jt}, \tilde{\boldsymbol{\lambda}}_t)^{\delta(s_e \approx \{t,p\})},$$

where $\delta(s_e \approx \{t,p\})$ takes a value of one if HMM state s_e is consistent with species t having presence/absence state p and a value of zero otherwise. Where there is a large alignment gap in a species (with size $\geq 5\text{kb}$), we force the emission probability for “presence” ($p=1$) in that species to zero, presuming a deletion.

The matrix of transition probabilities consists of three different types of transitions. First, self-transitions for all active and inactive states have probabilities ρ_1 and ρ_0 , respectively (fig. 2). Second, all transitions from active states to the inactive state have probability $1 - \rho_1$. Third, each transition from the inactive state to any active state s_e has probability $(1 - \rho_0) \cdot \frac{1}{Z} P_{\text{Fels}}(s_e)$. Here, $P_{\text{Fels}}(s_e)$ is the probability of the presence/absence pattern at the tips of the tree consistent with state s_e , as computed using Felsenstein’s pruning algorithm, under our phylogenetic turnover model (implicitly conditioned on the given phylogeny and the parameters π and γ). Because not all presence/absence patterns are possible, this probability must explicitly be normalized by the sum across all allowable states, $Z = \sum_{\text{active } s_e} P_{\text{Fels}}(s_e)$. Thus, the relative frequencies of the active states are proportional to their equilibrium probabilities under the specified phylogenetic process. All other elements of the transition matrix are fixed at zero, preventing direct transitions between active states.

epiPhyloHMM Model Fitting

For reasons of efficiency, we fit the epiPhyloHMM model to the data approximately in several successive steps. We first fit the peak-calling models separately to the data for each species, using a subset (125 Mb) of the mapped reads and estimating all free parameters by maximum likelihood with the L-BFGS-B algorithm (Zhu et al. 1997). As noted earlier, this calculation made use of the dispersion model that was pre-estimated using DESeq2. We then split the multispecies data into 20 partitions of similar size, with break-points in 50-kb long regions lacking alignment to the human genome by any other species. Separately in each species, we converted bins with $< 15\%$ of bases aligning to the human genome to missing data. We then estimated the ρ_1 , ρ_0 , π , and γ parameters of the epiPhyloHMM model by maximum likelihood using the L-BFGS-B algorithm (Zhu et al. 1997), keeping the species-specific peak-calling parameters—namely, the mixture coefficients $\tilde{\mathbf{w}}$ and mean counts $\tilde{\boldsymbol{\mu}}$ —at their previously estimated values. We then obtained an initial set of element calls using the Viterbi algorithm, and filtered them by the following heuristics:

- (1) We grouped maximal sets of elements that were separated by at most one “background” bin (as sometimes occurs due to the sparse design of the transition matrix; fig. 2).

- (2) For each grouped element, we computed an alignment “score” equal to the sum of alignment scaling factors γ_j across all bins j in the element and across all species.
- (3) We retained the element in the group having the highest alignment score.
- (4) In addition, we retained any elements having a score that exceeded a designated threshold T ($T = 16$ in our analysis).
- (5) We masked any remaining elements from the data by resetting their alignment scale factors γ_j to 0.

After this masking step, we re-estimated the ρ_1 , ρ_0 , π , and γ parameters of the epiPhyloHMM model separately for each partition of the data. Then we obtained our final set of predictions by running the Viterbi algorithm genome-wide using the median values of these per-block estimates.

The Running Time of epiPhyloHMM

The asymptotic running time of the forward algorithm for computing the likelihood for a fully general HMM is $O(Lk^2)$, where L is the number of sites being analyzed and k is the number of states in the model. In our case, however, we can reduce the complexity to roughly $O(Lk)$ by taking advantage of our sparse transition matrix (i.e., such that the “active” states have nonzero probability of transitioning only to themselves and the “inactive” state; see [fig. 2](#)). Furthermore, the number of possible states scales as $k = 2^S$, where S is the number of species. As noted earlier, we substantially reduce the state space by only allowing for at most three mutations on the tree and combining mutational patterns that result in the same active/inactive patterns at the tips of the tree.

To test the effects of these innovations in practice, we fit the phylogenetic parameters of epiPhyloHMM, using pre-fit peak-calling emission parameters (described above), for models having 9, 64, and 416 states (corresponding to 3, 6, and 9 species trees, respectively). We obtained run-times of approximately 6, 30, and 188 min, respectively. Thus, in practice, the run-time increases approximately linearly with the number of states.

CRE-Gene Association and Annotation as Enhancers and Promoters

All assignments were based on distances to genes obtained from Ensembl build 93 ([Zerbino et al. 2018](#)) via BiomaRt ([Durinck et al. 2005; 2009](#)). Promoter regions were defined per transcript as the interval ± 1.5 kb of the annotated TSS. Each promoter region was associated with the gene linked to the TSS in question, but multiple promoter regions were allowed per gene. H3K4me3 elements that overlapped (by at least one nucleotide) with a single promoter region were annotated as “promoter” and associated with the corresponding gene. H3K4me3 elements that overlapped with multiple promoter regions were annotated as an “unassociated promoter” (promoter-UA). H3K4me elements that did not overlap with any promoter regions were annotated as “unknown” (unk). For H3K27ac marks, the same rules were used to label an element as promoter or promoter-UA.

To classify enhancers, we first defined “expanded promoter regions” as intervals ± 10 kb of the TSS, again merging them across transcripts of the same gene. H3K27ac elements that overlapped with a single expanded promoter region were annotated as a “proximal enhancer” (enhancer_proximal) and associated with the corresponding gene. H3K27ac elements that overlapped with multiple expanded promoter regions were annotated as an “unassociated proximal enhancer” (enhancer_proximal-UA). H3K27ac elements that did not overlap with any expanded promoter regions but still fell within 100 kb of a TSS were annotated as a “distal enhancer” and associated with the closest gene (enhancer_distal). H3K27ac elements that met none of these criteria were labeled as “unknown.” This scheme is represented in [supplementary figure S11, Supplementary Material](#) online.

Linking Phylogenetic Parameters to Genomic Features for phyloGLM

As described in the Results section, the phylogenetic model for CRE turnover is defined by two free parameters: the turnover rate, γ , and the equilibrium frequency of element “presence” distribution, π . These parameters are defined as generalized linear functions of a vector of genomic features, or covariates, that are assumed to be available for each bin j .

Specifically, the turnover rate, γ_j , for a given bin j with covariate vector C_j is defined by passing a linear combination of C_j and a vector of coefficients θ_γ through the logistic function,

$$w(\theta_\gamma, C_j) = \frac{e^{C_j \cdot \theta_\gamma}}{1 + e^{C_j \cdot \theta_\gamma}},$$

with bounds of $\gamma_{\min} \leq \gamma_j \leq \gamma_{\max}$ imposed as follows,

$$\begin{aligned} \gamma_j &= \gamma_{\min} \cdot (1 - w(\theta_\gamma, C_j)) + \gamma_{\max} \cdot w(\theta_\gamma, C_j) \\ &= \gamma_{\min} + \frac{(\gamma_{\max} - \gamma_{\min})e^{C_j \cdot \theta_\gamma}}{1 + e^{C_j \cdot \theta_\gamma}} \\ &= \gamma_{\min} + \frac{\gamma_{\max} - \gamma_{\min}}{1 + e^{-C_j \cdot \theta_\gamma}}. \end{aligned}$$

Our implementation allows for a separate set of coefficients $\theta_{\gamma,e}$ for each edge e of the phylogeny, but in practice, we separate only the branch to the outgroup and apply the same coefficients to all other branches of the tree. This strategy ensures that weak power to distinguish gains and losses on this branch, and poor alignability to the outgroup, do not drive the maximum likelihood estimates for other parts of the tree.

Similarly, the stationary distribution of element presence π_j is defined as

$$\pi_j = \frac{e^{C_j \cdot \theta_\pi}}{1 + e^{C_j \cdot \theta_\pi}} = \frac{1}{1 + e^{-C_j \cdot \theta_\pi}}.$$

In order to prevent numerical underflow, if $\pi_j < 10^{-3}$, we reset $\pi_j = 10^{-3}$; similarly, if $1 - \pi_j < 10^{-3}$, we reset $\pi_j = 1 - 10^{-3}$.

Fitting a phyloGLM Model

The phyloGLM model was fit to the data using the L-BFGS-B algorithm (Zhu et al. 1997) as implemented in the “optim” function in R (R Core Team 2018). Gradients with respect to the log likelihood $\mathcal{L}(\bar{\theta}; \tilde{\mathbf{x}})$ (where $\bar{\theta}$ denotes the entire parameter set) for elements of the coefficient vectors θ_γ and θ_π were computed using the chain rule. For example, the gradient for an element i of the vector θ_γ , denoted $\theta_{\gamma,i}$, is given by,

$$\frac{\delta \mathcal{L}(\bar{\theta}; \tilde{\mathbf{x}})}{\delta \theta_{\gamma,i}} = \sum_j \frac{\delta \mathcal{L}(\bar{\theta}; \tilde{\mathbf{x}}_j)}{\delta U_j} \cdot \frac{\delta U_j}{\delta \theta_{\gamma,i}},$$

where the sum is across all bins, $\mathcal{L}(\bar{\theta}; \tilde{\mathbf{x}}_j)$ represents the contribution of bin j to the log likelihood function, and $U_j = \theta_\gamma \cdot C_j$ is the linear combination of features for bin j . The partial derivative $\frac{\delta \mathcal{L}(\bar{\theta}; \tilde{\mathbf{x}}_j)}{\delta U_j}$ was computed numerically. However, the final term can easily be computed analytically as,

$$\begin{aligned} \frac{\delta U_j}{\delta \theta_{\gamma,i}} &= \frac{\delta}{\delta \theta_{\gamma,i}} (\theta_{\gamma,0} + \theta_{\gamma,1} C_{j,1} + \dots + \theta_{\gamma,i} C_{j,i} + \dots + \theta_{\gamma,n} C_{j,n}) \\ &= C_{j,i}. \end{aligned}$$

Notice that the use of the chain rule considerably accelerates these calculations because $\frac{\delta \mathcal{L}(\bar{\theta}; \tilde{\mathbf{x}}_j)}{\delta U_j}$ only needs to be computed once per phylogenetic parameter (γ and π), and then can be propagated efficiently to all of the individual coefficients. The same approach works for elements of θ_π .

Preparation of Genomic Features

Distances to the TSS were based on annotations from Ensembl build 93 (Zerbino et al. 2018), accessed via BiomaRt (Durinck et al. 2005; 2009). The distance was computed from the nearest boundary of the CRE to the annotated TSS position. Gene expression features were based on data downloaded from the GTEx web portal (<https://gtexportal.org/>; last accessed April 2, 2020). Mean phastCons-100way scores were calculated using the GenomicScores package (Puigdevall and Castelo 2018). pLI scores were collected from ftp://ftp.broadinstitute.org/pub/ExAC_release/release1/manuscript_data/ (last accessed April 2, 2020). Functional annotations of genes were obtained from Reactome 2018 (Fabregat et al. 2016). For the fitting of the phyloGLM model, only bins that were unambiguously associated with a single gene and had no missing covariate values were used. After this filtering, 7,220 promoters and 25,990 enhancers associated with 5,307 and 5,552 unique genes remained. All noncategorical CRE covariates were scaled to have mean 0 and SD 1 to improve model-fitting performance and produce comparable coefficient values across covariates.

Expected Numbers of Gains and Losses per Branch

We calculated the expected numbers of gains and losses per branch (supplementary figs. S15 and S16, Supplementary Material online) using the standard message-passing algorithm on the phylogeny, followed by estimation of the

probabilities of each state transition per branch (see Siepel and Haussler 2003 for details).

Half-Life Estimation

Under our model, the instantaneous rate of transition from the active state to the inactive state is given by $\gamma(1 - \pi)$. Thus, the half-life $t_{\frac{1}{2}}$, or time required for half of active elements to become inactive, is given by:

$$\frac{1}{2} = e^{-\gamma(1-\pi)t_{\frac{1}{2}}} \Rightarrow t_{\frac{1}{2}} = \frac{\ln(2)}{\gamma(1-\pi)}$$

Dosage-Sensitivity Analysis

Annotations of genes as “Metabolic” or “Generic Transcription Pathway” were obtained from Reactome 2018 (Fabregat et al. 2016). Only genes having one of the functional annotations “R-HSA-1430728” or “R-HSA-212436” were selected for analysis; genes with both labels were omitted. To test for significant differences between the two gene categories, we compared the log likelihoods of phyloGLM models having 1) separate coefficients for the two gene categories and 2) one coefficient that applied to both categories, computing P values for one degree of freedom. Mean 100way phastCons scores were calculated using bwtool (Pohl and Beato 2014).

Software Availability

epiPhyloHMM is available as an R package at <https://github.com/ndukler/flexPhyloHMM> (last accessed April 2, 2020). phyloGLM is also available as an R package at <https://github.com/ndukler/phyloGLM> (last accessed April 2, 2020).

Supplementary Material

Supplementary data are available at *Molecular Biology and Evolution* online.

Acknowledgments

This research was supported by US National Institutes of Health (Grant Nos R01-HG010346 and R35-GM127070). 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

- Anders S, McCarthy DJ, Chen Y, Okoniewski M, Smyth GK, Huber W, Robinson MD. 2013. Count-based differential expression analysis of RNA sequencing data using R and Bioconductor. *Nat Protoc*. 8(9):1765–1786.
- Arbiza L, Gronau I, Aksoy BA, Hubisz MJ, Gulko B, Keinan A, Siepel A. 2013. Genome-wide inference of natural selection on human transcription factor binding sites. *Nat Genet*. 45(7):723–729.
- Armstrong J, Hickey G, Diekhans M, Deran A, Fang Q, Xie D, Feng S, Stiller J, Genereux D, Johnson J, et al. 2019. Progressive alignment with cactus: a multiple-genome aligner for the thousand-genome era. *bioRxiv*.
- Battle A, Brown CD, Engelhardt BE, Montgomery SB. 2017. Genetic effects on gene expression across human tissues. *Nature* 550(7675):204–213.

- Benton ML, Talipineni SC, Kostka D, Capra JA. 2019. Genome-wide enhancer annotations differ significantly in genomic distribution, evolution, and function. *BMC Genomics* 20(1):511.
- Bernstein BE, Kamal M, Lindblad-Toh K, Bekiranov S, Bailey DK, Huebert DJ, McMahon S, Karlsson EK, Kulbokas EJ, Gingeras TR, et al. 2005. Genomic maps and comparative analysis of histone modifications in human and mouse. *Cell* 120(2):169–181.
- Berthelot C, Villar D, Horvath JE, Odom DT, Flicek P. 2018. Complexity and conservation of regulatory landscapes underlie evolutionary resilience of mammalian gene expression. *Nat Ecol Evol.* 2:152.
- Bradley RK, Li X-Y, Trapnell C, Davidson S, Pachter L, Chu HC, Tonkin LA, Biggin MD, Eisen MB. 2010. Binding site turnover produces pervasive quantitative changes in transcription factor binding between closely related *Drosophila* species. *PLoS Biol.* 8(3):e1000343.
- Bullaughhey K. 2011. Changes in selective effects over time facilitate turnover of enhancer sequences. *Genetics* 187(2):567–582.
- Cotney J, Leng J, Yin J, Reilly SK, DeMare LE, Emera D, Ayoub AE, Rakic P, Noonan JP. 2013. The evolution of lineage-specific regulatory activities in the human embryonic limb. *Cell* 154(1):185–196.
- Danko CG, Choate LA, Marks BA, Rice EJ, Wang Z, Chu T, Martins AL, Dukler N, Coonrod SA, Tait Wojno ED, et al. 2018. Dynamic evolution of regulatory element ensembles in primate CD4 + T cells. *Nat Ecol Evol.* 2:537.
- Dermitzakis ET, Clark AG. 2002. Evolution of transcription factor binding sites in Mammalian gene regulatory regions: conservation and turnover. *Mol Biol Evol.* 19(7):1114–1121.
- Doniger SW, Fay JC. 2007. Frequent gain and loss of functional transcription factor binding sites. *PLoS Comput Biol.* 3(5):e99.
- Duque T, Samee MAH, Kazemian M, Pham HN, Brodsky MH, Sinha S. 2014. Simulations of enhancer evolution provide mechanistic insights into gene regulation. *Mol Biol Evol.* 31(1):184–200.
- Durinck S, Moreau Y, Kasprzyk A, Davis S, De Moor B, Brazma A, Huber W. 2005. BioMart and Bioconductor: a powerful link between biological databases and microarray data analysis. *Bioinformatics* 21(16):3439–3440.
- Durinck S, Spellman PT, Birney E, Huber W. 2009. Mapping identifiers for the integration of genomic datasets with the R/Bioconductor package biomaRt. *Nat Protoc.* 4(8):1184–1191.
- Fabregat A, Sidiropoulos K, Garapati P, Gillespie M, Hausmann K, Haw R, Jassal B, Jupe S, Korninger F, McKay S, et al. 2016. The Reactome pathway Knowledgebase. *Nucleic Acids Res.* 44(D1):D481–D487.
- Felsenstein J. 1973. Maximum likelihood and minimum-steps methods for estimating evolutionary trees from data on discrete characters. *Syst Biol.* 22(3):240–249.
- Felsenstein J. 1981. Evolutionary trees from DNA sequences: a maximum likelihood approach. *J Mol Evol.* 17(6):368–376.
- Felsenstein J. 2004. Inferring phylogenies. Sunderland (MA): Sinauer Associates.
- Felsenstein J, Churchill GA. 1996. A hidden Markov model approach to variation among sites in rate of evolution. *Mol Biol Evol.* 13(1):93–104.
- Fisher S, Grice EA, Vinton RM, Bessling SL, McCallion AS. 2006. Conservation of RET regulatory function from human to zebrafish without sequence similarity. *Science* 312(5771):276–279.
- Frankel N, Davis GK, Vargas D, Wang S, Payre F, Stern DL. 2010. Phenotypic robustness conferred by apparently redundant transcriptional enhancers. *Nature* 466(7305):490–493.
- Fulco CP, Munschauer M, Anyoha R, Munson G, Grossman SR, Perez EM, Kane M, Cleary B, Lander ES, Engreitz JM, et al. 2016. Systematic mapping of functional enhancer-promoter connections with CRISPR interference. *Science* 354(6313):769–773.
- Fuller ZL, Berg JJ, Mostafavi H, Sella G, Przeworski M. 2019. Measuring intolerance to mutation in human genetics. *Nat Genet.* 51(5):772–776.
- Gusev A, Lee SH, Trynka G, Finucane H, Vilhjálmsson BJ, Xu H, Zang C, Ripke S, Bulik-Sullivan B, Stahl E, et al. 2014. Partitioning heritability of regulatory and cell-type-specific variants across 11 common diseases. *Am J Hum Genet.* 95(5):535–552.
- Hare EE, Peterson BK, Iyer VN, Meier R, Eisen MB. 2008. Sepsid even-skipped enhancers are functionally conserved in *Drosophila* despite lack of sequence conservation. *PLoS Genet.* 4(6):e1000106.
- Hindorf LA, Sethupathy P, Junkins HA, Ramos EM, Mehta JP, Collins FS, Manolio TA. 2009. Potential etiologic and functional implications of genome-wide association loci for human diseases and traits. *Proc Natl Acad Sci U S A.* 106(23):9362–9367.
- Hinrichs AS, Karolchik D, Baertsch R, Barber GP, Bejerano G, Clawson H, Diekhans M, Furey TS, Harte RA, Hsu F, et al. 2006. The UCSC Genome Browser Database: update 2006. *Nucleic Acids Res.* 34(90001):D590–D598.
- Huang Y-F, Gulko B, Siepel A. 2017. Fast, scalable prediction of deleterious noncoding variants from functional and population genomic data. *Nat Genet.* 49(4):618–624.
- Huang Y-F, Siepel A. 2019. Estimation of allele-specific fitness effects across human protein-coding sequences and implications for disease. *Genome Res.* 245522:118.
- Jin F, Li Y, Dixon JR, Selvaraj S, Ye Z, Lee AY, Yen C-A, Schmitt AD, Espinoza CA, Ren B, et al. 2013. A high-resolution map of the three-dimensional chromatin interactome in human cells. *Nature* 503(7475):290–294.
- Jones FC, Grabherr MG, Chan YF, Russell P, Mauceci E, Johnson J, Swofford R, Pirun M, Zody MC, White S, et al. 2012. The genomic basis of adaptive evolution in threespine sticklebacks. *Nature* 484(7392):55–61.
- Jubb AW, Young RS, Hume DA, Bickmore WA. 2016. Enhancer turnover is associated with a divergent transcriptional response to glucocorticoid in mouse and human macrophages. *J Immunol.* 196(2):813–822.
- Kasowski M, Kyriazopoulou-Panagiotopoulou S, Grubert F, Zaugg JB, Kundaje A, Liu Y, Boyle AP, Zhang QC, Zakharia F, Spacek DV, et al. 2013. Extensive variation in chromatin states across humans. *Science* 342(6159):750–752.
- Khan Z, Ford MJ, Cusanovich DA, Mitrano A, Pritchard JK, Gilad Y. 2013. Primate transcript and protein expression levels evolve under compensatory selection pressures. *Science* 342(6162):1100–1104.
- Khoueiry P, Girardot C, Ciglar L, Peng P-C, Gustafson EH, Sinha S, Furlong EE. 2017. Uncoupling evolutionary changes in DNA sequence, transcription factor occupancy and enhancer activity. *eLife.* 6:1–29.
- Kondrashov FA, Koonin EV. 2004. A common framework for understanding the origin of genetic dominance and evolutionary fates of gene duplications. *Trends Genet.* 20(7):287–290.
- Kumar S, Stecher G, Suleski M, Hedges SB. 2017. TimeTree: a resource for timelines, timetrees, and divergence times. *Mol Biol Evol.* 34(7):1812–1819.
- Langmead B, Salzberg SL. 2012. Fast gapped-read alignment with Bowtie 2. *Nat Methods.* 9(4):357–359.
- Lek M, Karczewski KJ, Minikel EV, Samocha KE, Banks E, Fennell T, O'Donnell-Luria AH, Ware JS, Hill AJ, Cummings BB, et al. 2016. Analysis of protein-coding genetic variation in 60,706 humans. *Nature* 536(7616):285–291.
- Lindblad-Toh K, Garber M, Zuk O, Lin MF, Parker BJ, Washietl S, Kheradpour P, Ernst J, Jordan G, Mauceci E, et al. 2011. A high-resolution map of human evolutionary constraint using 29 mammals. *Nature* 478(7370):476–482.
- Lindblad-Toh K, Wade CM, Mikkelsen TS, Karlsson EK, Jaffe DB, Kamal M, Clamp M, Chang JL, Kulbokas EJ, Zody MC, et al. 2005. Genome sequence, comparative analysis and haplotype structure of the domestic dog. *Nature* 438(7069):803–819.
- Love MI, Huber W, Anders S. 2014. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biol.* 15(12):550.
- Ludwig MZ, Bergman C, Patel NH, Kreitman M. 2000. Evidence for stabilizing selection in a eukaryotic enhancer element. *Nature* 403(6769):564–567.
- Maher KA, Bajic M, Kajala K, Reynoso M, Pauluzzi G, West DA, Zumstein K, Woodhouse M, Bubbs K, Dorrity MW, et al. 2018. Profiling of accessible chromatin regions across multiple plant species and cell

- types reveals common gene regulatory principles and new control modules. *Plant Cell* 30(1):15–36.
- Marinov GK, Kundaje A. 2018. ChIP-ping the branches of the tree: functional genomics and the evolution of eukaryotic gene regulation. *Brief Funct Genomics*. 17(2):116–137.
- Marmoset Genome Sequencing and Analysis Consortium. 2014. The common marmoset genome provides insight into primate biology and evolution. *Nat Genet*. 46:850–857.
- McLean CY, Reno PL, Pollen AA, Bassan AI, Capellini TD, Guenther C, Indjeian VB, Lim X, Menke DB, Schaar BT, et al. 2011. Human-specific loss of regulatory DNA and the evolution of human-specific traits. *Nature* 471(7337):216–219.
- Meyer AG, Wilke CO. 2013. Integrating sequence variation and protein structure to identify sites under selection. *Mol Biol Evol*. 30(1):36–44.
- Mifsud B, Tavares-Cadete F, Young AN, Sugar R, Schoenfelder S, Ferreira L, Wingett SW, Andrews S, Grey W, Ewels PA, et al. 2015. Mapping long-range promoter contacts in human cells with high-resolution capture Hi-C. *Nat Genet*. 47(6):598–606.
- Mikkelsen TS, Wakefield MJ, Aken B, Amemiya CT, Chang JL, Duke S, Garber M, Gentles AJ, Goodstadt L, Heger A, et al. 2007. Genome of the marsupial *Monodelphis domestica* reveals innovation in non-coding sequences. *Nature* 447(7141):167–177.[10.1038/nature05805]
- Mikkelsen TS, Xu Z, Zhang X, Wang L, Gimble JM, Lander ES, Rosen ED. 2010. Comparative epigenomic analysis of murine and human adipogenesis. *Cell* 143(1):156–169.
- Moses AM, Pollard DA, Nix DA, Iyer VN, Li X-Y, Biggin MD, Eisen MB. 2006. Large-scale turnover of functional transcription factor binding sites in *Drosophila*. *PLoS Comp Biol*. 2(10):e130.
- Odom DT, Dowell RD, Jacobsen ES, Gordon W, Danford TW, Maclsaac KD, Rolfe PA, Conboy CM, Gifford DK, Fraenkel E, et al. 2007. Tissue-specific transcriptional regulation has diverged significantly between human and mouse. *Nat Genet*. 39(6):730–732.
- Osterwalder M, Barozzi I, Tissières V, Fukuda-Yuzawa Y, Mannion BJ, Afzal SY, Lee EA, Zhu Y, Plajzer-Frick I, Pickle CS, et al. 2018. Enhancer redundancy provides phenotypic robustness in mammalian development. *Nature* 554(7691):239–243.
- Paris M, Kaplan T, Li XY, Villalta JE, Lott SE, Eisen MB. 2013. Extensive divergence of transcription factor binding in *Drosophila* embryos with highly conserved gene expression. *PLoS Genet*. 9(9):e1003748.
- Paten B, Earl D, Nguyen N, Diekhans M, Zerbino D, Haussler D. 2011. Cactus: algorithms for genome multiple sequence alignment. *Genome Res*. 21(9):1512–1528.
- Peng X, Alféldi J, Gori K, Eisfeld AJ, Tyler SR, Tisoncik-Go J, Brawand D, Law GL, Skunca N, Hattar M, et al. 2014. The draft genome sequence of the ferret (*Mustela putorius furo*) facilitates study of human respiratory disease. *Nat Biotechnol*. 32(12):1250–1255.
- Perry MW, Boettiger AN, Bothma JP, Levine M. 2010. Shadow enhancers foster robustness of *Drosophila* gastrulation. *Curr Biol*. 20(17):1562–1567.
- Pohl A, Beato M. 2014. Bwtool: a tool for bigWig files. *Bioinformatics* 30(11):1618–1619.
- Prabhakar S, Visel A, Akiyama JA, Shoukry M, Lewis KD, Holt A, Plajzer-Frick I, Morrison H, FitzPatrick DR, Afzal V, et al. 2008. Human-specific gain of function in a developmental enhancer. *Science* 321(5894):1346–1350.
- Prescott SL, Srinivasan R, Marchetto MC, Grishina I, Narvaiza I, Sella L, Gage FH, Swigut T, Wysocka J. 2015. Enhancer divergence and cis-regulatory evolution in the human and chimp neural crest. *Cell* 163(1):68–83.
- Puigdevall P, Castelo R. 2018. GenomicScores: seamless access to genome-wide position-specific scores from R and Bioconductor. *Bioinformatics* 34(18):3208–3210.
- Qu J, Hodges E, Molaro A, Gagneux P, Dean MD, Hannon GJ, Smith AD. 2018. Evolutionary expansion of DNA hypomethylation in the mammalian germline genome. *Genome Res*. 28(2):145–158.
- R Core Team. 2018. R: a language and environment for statistical computing. Vienna (Austria): R Foundation for Statistical Computing. Available from: <https://www.R-project.org>.
- Rat Genome Sequencing Project Consortium. 2004. Genome sequence of the Brown Norway rat yields insights into mammalian evolution. *Nature* 428:493–521.
- Reilly SK, Noonan JP. 2016. Evolution of gene regulation in humans. *Annu Rev Genomics Hum Genet*. 17(1):45–67.
- Sanyal A, Lajoie BR, Jain G, Dekker J. 2012. The long-range interaction landscape of gene promoters. *Nature* 489(7414):109–113.
- Schmidt D, Wilson MD, Ballester B, Schwalie PC, Brown GD, Marshall A, Kutter C, Watt S, Martinez-Jimenez CP, Mackay S, et al. 2010. Five-vertebrate ChIP-seq reveals the evolutionary dynamics of transcription factor binding. *Science* 328(5981):1036–1040.
- Shibata Y, Sheffield NC, Fedrigo O, Babbitt CC, Wortham M, Tewari AK, London D, Song L, Lee B-K, Iyer VR, et al. 2012. Extensive evolutionary changes in regulatory element activity during human origins are associated with altered gene expression and positive selection. *PLoS Genet*. 8(6):e1002789.
- Siepel A, Bejerano G, Pedersen JS, Hinrichs AS, Hou M, Rosenbloom K, Clawson H, Spieth J, Hillier LW, Richards S, et al. 2005. Evolutionarily conserved elements in vertebrate, insect, worm, and yeast genomes. *Genome Res*. 15(8):1034–1050.
- Siepel A, Arbiza L. 2014. Cis-regulatory elements and human evolution. *Curr Opin Genet Dev*. 29:81–89.
- Siepel A, Haussler D. 2003. Phylogenetic estimation of context-dependent substitution rates by maximum likelihood. *Mol Biol Evol*. 21(3):468–488.
- Siepel A, Haussler D. 2005. Phylogenetic hidden Markov models. In: Nielsen R, editor. *Statistical methods in molecular evolution*. New York: Springer-Verlag. p. 325–351.
- Siepel A, Pollard KS, Haussler D. 2006. New methods for detecting lineage-specific selection. In: Hutchison D, editor. *Research in computational molecular biology*. Vol. 3909. Heidelberg, Berlin: Springer. p. 190–205.
- Tournamille C, Colin Y, Cartron JP, Le Van Kim C. 1995. Disruption of a GATA motif in the Duffy gene promoter abolishes erythroid gene expression in Duffy-negative individuals. *Nat Genet*. 10(2):224–228.
- Tugrul M, Paixao T, Barton NH, Tkacik G. 2015. Dynamics of transcription factor binding site evolution. *PLoS Genet*. 11:e1005639.
- Vietri Rudan M, Barrington C, Henderson S, Ernst C, Odom DT, Tanay A, Hadjir S. 2015. Comparative Hi-C reveals that CTCF underlies evolution of chromosomal domain architecture. *Cell Rep*. 10(8):1297–1309.
- Villar D, Berthelot C, Aldridge S, Rayner TF, Lukk M, Pignatelli M, Park TJ, Deaville R, Erichsen JT, Jasinska AJ, et al. 2015. Enhancer evolution across 20 mammalian species. *Cell* 160(3):554–566.
- Villar D, Flicek P, Odom DT. 2014. Evolution of transcription factor binding in metazoans – mechanisms and functional implications. *Nat Rev Genet*. 15(4):221–233.
- Wang T, Zeng J, Lowe CB, Sellers RC, Salama SR, Yang M, Burgess SM, Brachmann RK, Haussler D. 2007. Species-specific endogenous retroviruses shape the transcriptional network of the human tumor suppressor protein p53. *Proc Natl Acad Sci U S A*. 104(47):18613–18618.
- Waterston RH, Lindblad-Toh K, Birney E, Rogers J, Abril JF, Agarwal P, Agarwala R, Ainscough R, Alexandersson M, An P, et al. 2002. Initial sequencing and comparative analysis of the mouse genome. *Nature* 420(6915):520–562.
- Weirauch MT, Hughes TR. 2010. Conserved expression without conserved regulatory sequence: the more things change, the more they stay the same. *Trends Genet*. 26(2):66–74.
- Wilkie AO. 1994. The molecular basis of genetic dominance. *J Med Genet*. 31(2):89–98.
- Wittkopp PJ, Kalay G. 2012. Cis-regulatory elements: molecular mechanisms and evolutionary processes underlying divergence. *Nat Rev Genet*. 13(1):59–69.

- Wong ES, Thybert D, Schmitt BM, Stefflova K, Odom DT, Flicek P. 2015. Decoupling of evolutionary changes in transcription factor binding and gene expression in mammals. *Genome Res.* 25(2):167–178.
- Xiao S, Xie D, Cao X, Yu P, Xing X, Chen C-C, Musselman M, Xie M, West FD, Lewin HA, et al. 2012. Comparative epigenomic annotation of regulatory DNA. *Cell* 149(6):1381–1392.
- Yan G, Zhang G, Fang X, Zhang Y, Li C, Ling F, Cooper DN, Li Q, Li Y, van Gool AJ, et al. 2011. Genome sequencing and comparison of two nonhuman primate animal models, the cynomolgus and Chinese rhesus macaques. *Nat Biotechnol.* 29(11):1019–1023.
- Yang S, Oksenberg N, Takayama S, Heo S-J, Poliakov A, Ahituv N, Dubchak I, Boffelli D. 2015. Functionally conserved enhancers with divergent sequences in distant vertebrates. *BMC Genomics* 16(1):882.
- Yang Y, Gu Q, Zhang Y, Sasaki T, Crivello J, O'Neill RJ, Gilbert DM, Ma J. 2018. Continuous-Trait Probabilistic Model for Comparing Multi-species Functional Genomic Data. *Cell Syst.* 7:208–218.e11.
- Yang Z. 1995. A space-time process model for the evolution of DNA sequences. *Genetics* 139(2):993–1005.
- Zabidi MA, Arnold CD, Schernhuber K, Pagani M, Rath M, Frank O, Stark A. 2015. Enhancer-core-promoter specificity separates developmental and housekeeping gene regulation. *Nature* 518(7540):556–559.
- Zerbino DR, Achuthan P, Akanni W, Amode MR, Barrell D, Bhai J, Billis K, Cummins C, Gall A, Girón CG, et al. 2018. Ensembl 2018. *Nucleic Acids Res.* 46(D1):D754–D761.
- Zhang Y, Liu T, Meyer CA, Eeckhoute J, Johnson DS, Bernstein BE, Nussbaum C, Myers RM, Brown M, Li W, et al. 2008. Model-based analysis of ChIP-Seq (MACS). *Genome Biol.* 9(9):R137.
- Zhu C, Byrd RH, Lu P, Nocedal J. 1997. Algorithm 778: LBFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans Math Softw.* 23(4):550–560.