RegEvol: detection of directional selection in regulatory sequences through phenotypic predictions and phenotype-to-fitness functions

Alexandre Laverré[1,2], Thibault Latrille[3], Marc Robinson-Rechavi[1,2]

1. Department of Ecology and Evolution, University of Lausanne, Lausanne, Switzerland.

2. Swiss Institute of Bioinformatics (SIB), Lausanne, Switzerland.

3. Department of Computational Biology, University of Lausanne, Lausanne, Switzerland.

Abstract

Regulatory DNA controls when and where genes are expressed, making it a key driver of phenotypic evolution. Yet detecting selection in non-coding regions remains difficult, as most approaches rely on sequence conservation or changes in substitution rate rather than molecular effects. RegEvol bridges this gap by linking machine learning-based predictions of transcription factor binding to explicit evolutionary models. It uses the distribution of predicted mutational effects to infer fitness functions under different evolutionary scenarios (random drift, stabilising selection, or directional selection). Through maximum-likelihood estimation, it identifies the regime that best explains observed changes along a lineage from an ancestral sequence. RegEvol corrects biases that affected previous tests based on machine learning of transcription factor binding, while remaining conservative across different levels of divergence. Applied to over 3 million *Drosophila melanogaster* regulatory regions, we identify 5.1% of them under directional selection, enriched near reproductive and immune genes. The framework is readily applicable to experimentally detected regulatory elements with alignable ancestral sequences and is flexible to future advances in understanding regulatory function, providing a powerful basis for investigating adaptation in non-coding regions.

Key-words: regulatory evolution, directional section, fitness function, machine learning

Introduction

The evolution of *cis*-regulatory sequences plays a central role in shaping phenotypic diversity and adaptation across species. Regulatory evolution has been proposed as a significant factor in morphological and physiological change since King and Wilson's well-known suggestion that many phenotypic divergences between closely related species may be caused by regulatory changes rather than differences in protein coding (King and Wilson 1975). This view is now supported by abundant evidence, with the majority of variants associated with phenotypic traits and human diseases falling within non-coding regulatory elements, such as enhancers and promoters (Maurano et al. 2012; Zhang and Lupski 2015; Lee et al. 2018). These sequences attract transcription factors (TFs) that bind to DNA by recognising specific sequence motifs, enabling precise gene expression patterns. This motif-based binding provides a direct link between genotype and phenotype, serving as the basis for sequence-based models of regulatory function (Lai et al. 2019; Sokolova et al. 2024).

Despite their functional importance, studying the evolutionary forces on regulatory sequences remains a major challenge. Unlike protein-coding sequences, where the ratio of nonsynonymous to synonymous substitutions (dN/dS) provides a framework to detect selection (Tanaka and Nei 1989; Yang 1998; Yang 2014), regulatory sequences lack a comparable metric. This is partly due to the absence of a universal “regulatory code” that would allow for a systematic interpretation of mutations' functionnal effects in non-coding DNA (Kim and Wysocka 2023). As a result, it remains difficult to distinguish between neutral drift, stabilising selection, and adaptive evolution in regulatory regions. The majority of commonly applied approaches to detect selection in non-coding DNA are based on substitution rate variation among phylogenies. Tools such as PhastCons (Siepel et al. 2005) and PhyloP (Pollard et al. 2010) identify conserved and accelerated regions by comparing observed substitution rates to neutral expectations. These approaches have been widely used to identify deeply conserved elements that are likely to be functionally important, and can allow for branch-specific rate shifts. However, they remain limited in several significant aspects. First, substitution rate-based methods are inherently indirect: they infer selection from sequence conservation or acceleration patterns, without modelling the functional consequences of mutations. As a result, they can be confounded by non-adaptive processes such as biased gene conversion, variation in mutation rates, or demographic history (Galtier and Duret 2007). Second, because they are based on conservation across lineages, they can be poorly suited for detecting selection in fast-evolving or recently gained regulatory elements. Indeed, most cis-regulatory elements are not conserved at large evolutionary

distances (Villar et al. 2015), and can maintain conserved regulatory activity despite substantial sequence divergence (Wong et al. 2020; Phan et al. 2025). Finally, they often rely on predefined neutral proxies, such as 4-fold degenerate sites or repetitive elements. These may not accurately reflect local mutational or selective backgrounds, especially in non-model species with limited annotations. Recent studies have attempted to enhance these models by incorporating phylogenetic structure and functional annotations. For example, probabilistic framework has been developed to model the turnover of cis-regulatory elements using ChIP-seq data and phylogenetic hidden Markov models (Dukler et al. 2020). Similarly, Yan et al. introduced PhyloAcc-GT, a Bayesian method that detects substitution rate shifts in conserved noncoding elements while accounting for gene tree discordance (Yan et al. 2023). These approaches represent important advances in modeling substitution rate variation, but they still rely on the underlying premise that conservation of function is coupled to conservation of sequence.

To overcome these limitations, recent efforts have focused on building genotype-to-phenotype maps that directly link regulatory sequence variation to functional output. Massively parallel reporter assays have enabled the systematic measurement of the effects of thousands of mutations on enhancer activity, providing empirical fitness landscapes for regulatory elements (Patwardhan et al. 2012; Smith et al. 2013; Gallego Romero and Lea 2023). These data have been used to train machine learning models that predict TF binding, chromatin accessibility, and gene expression from DNA sequence alone (Vaishnav et al. 2022; Krieger et al. 2022). These approaches demonstrate that the functional impact of mutations in regulatory DNA can be inferred directly from sequence, providing a powerful alternative to conservation-based methods. Building on this conceptual foundation, Liu and Robinson-Rechavi (2020) previously developed a machine learning-based method to detect positive selection on regulatory sequences by comparing observed changes in predicted TF binding affinity to a null distribution of random mutations. This approach leveraged gapped k-mer support vector machines (gkm-SVMs) trained on ChIP-seq data to predict the impact of substitutions on TF binding. It can infer whether a regulatory sequence has evolved under positive selection by comparing the predicted change in binding affinity (ΔSVM) to a null distribution generated by in silico mutagenesis. This method offers clear advantages over traditional rate-based approaches, as it avoids reliance on predefined neutral sites, can be applied to individual regulatory elements, and models the functional impact of mutations directly from sequence. Yet despite these strengths, several important limitations have emerged. One is ascertainment bias in ChIP-seq data, where enrichment for high-affinity binding sites can inflate false-positive rates in selection tests (Jiang and Zhang 2024). Another is that the method's accuracy declines with

divergence, since its non-parametric null model becomes unreliable over larger evolutionary distances.

Here we present RegEvol, a new framework building on our previous work by combining machine learning-based predictions of mutational effects with explicit population-genetic models of selection. Rather than relying on substitution rate shifts or permutation-based tests, we directly model the fitness effects of mutations in regulatory sequences. Using a trained sequence-to-function model, our approach predicts the functional impact of all possible point mutations in ChIP-seq-defined regulatory regions, resulting in genotype-to-phenotype maps. These predictions are then incorporated into a population-genetic framework that uses phenotype-to-fitness functions to determine the fitness effect of phenotypic changes. We evaluate the performance of this method through simulations and compare it with our previous one. Finally, we apply it to diverse empirical datasets in *Drosophila melanogaster* to study the evolutionary dynamics of regulatory sequences.

Results

1 - SVM predictions capture functional signals beyond sequence conservation

We trained gapped k-mer support vector machine (gkm-SVM) models to distinguish transcription factor (TF)-bound regions identified by ChIP-seq from random genomic sequences matched for length, GC content, and repeat composition (Lee 2016). These models learn to recognize specific combinations of short DNA motifs (k-mers) enriched in TF-bound regions, thereby capturing the sequence context and features that define functional regulatory elements. Despite the rise of deep learning frameworks such as DeepSEA (Zhou and Troyanskaya 2015) and BPNet (Avsec et al. 2021), gkm-SVM remains among the most interpretable and TF-specific approaches, with performance comparable to deep models when trained on individual TF datasets (Tognon et al. 2025; Vorontsov et al. 2025). Moreover, gkm-SVM directly quantifies sequence-level contributions, enabling efficient in silico mutagenesis at base-pair resolution. Once trained, the model assigns an SVM score to any DNA sequence, reflecting its predicted similarity to the ChIP-seq peak class based on k-mer composition.

We first confirmed the biological relevance of these predictions by analysing how SVM scores vary along individual ChIP-seq peaks. Using a sliding window approach, we computed position-wise SVM scores across *Homo sapiens* peaks bound by CEBPA. A representative example is shown in Figure 1A, which reveals two distinct regions of elevated scores, each approximately 10 bp in length, corresponding to predicted TF binding sites.

These high-scoring regions matched the canonical CEBPA motif from the JASPAR database (Castro-Mondragon et al. 2022), suggesting that the model accurately identifies functional binding sites within regulatory elements. To generalise this observation, we extracted the top 100 10-mers with the highest SVM weights from the CEBPA model and compared them to known motifs using the TomTom tool from the MEME Suite (Gupta et al. 2007; Bailey et al. 2009). The top-scoring matches corresponded to the CEBPA motif and related forkhead TFs (Supplementary Table 1), confirming that the model captures biologically meaningful sequence features associated with TF binding.

We then quantified the predicted impact of point mutations using ΔSVM, defined as the change in SVM score resulting from a single-nucleotide substitution. For each position in the sequence, we introduced all possible point mutations in silico and computed the resulting change in predicted binding affinity. The absolute ΔSVM values reflect the magnitude of predicted functional disruption caused by each mutation. To assess how well these metrics reflect TF binding intensity, we computed normalised read coverage for each base within CEBPA ChIP-seq peaks, using it as a proxy for TF abundance. All peaks were centred on their summits (i.e, the position of maximum read coverage) and aligned to produce an average signal profile. We then calculated SVM scores, ΔSVM values, and well-known conservation metrics, phastCons and phyloP scores, at each aligned position. Both SVM and ΔSVM scores exhibit a sharp peak at the summit, closely mirroring the ChIP-seq signal (Figure 1B). This indicates that mutations in these regions are predicted to most strongly affect TF binding, either directly or indirectly. On the other hand, phastCons and phyloP scores lack such enrichment and have a lower dynamic range, suggesting that the conservation scores are less sensitive to fine-scale variation of TF occupancy and thus regulatory region function. Indeed, both SVM and ΔSVM scores are positively correlated with read coverage over individual peaks (median Pearson's correlation coefficient = 0.31 and 0.22 respectively; Fig 1C), while phastCons and phyloP scores are largely uncorrelated (median Pearson's correlation coefficient = 1.8e-2 and 2.7e-3). This confirms that SVM-based metrics capture TF binding intensity more accurately than conservation-based scores. These patterns were consistent across multiple TFs and species (Supplementary Figure 1).

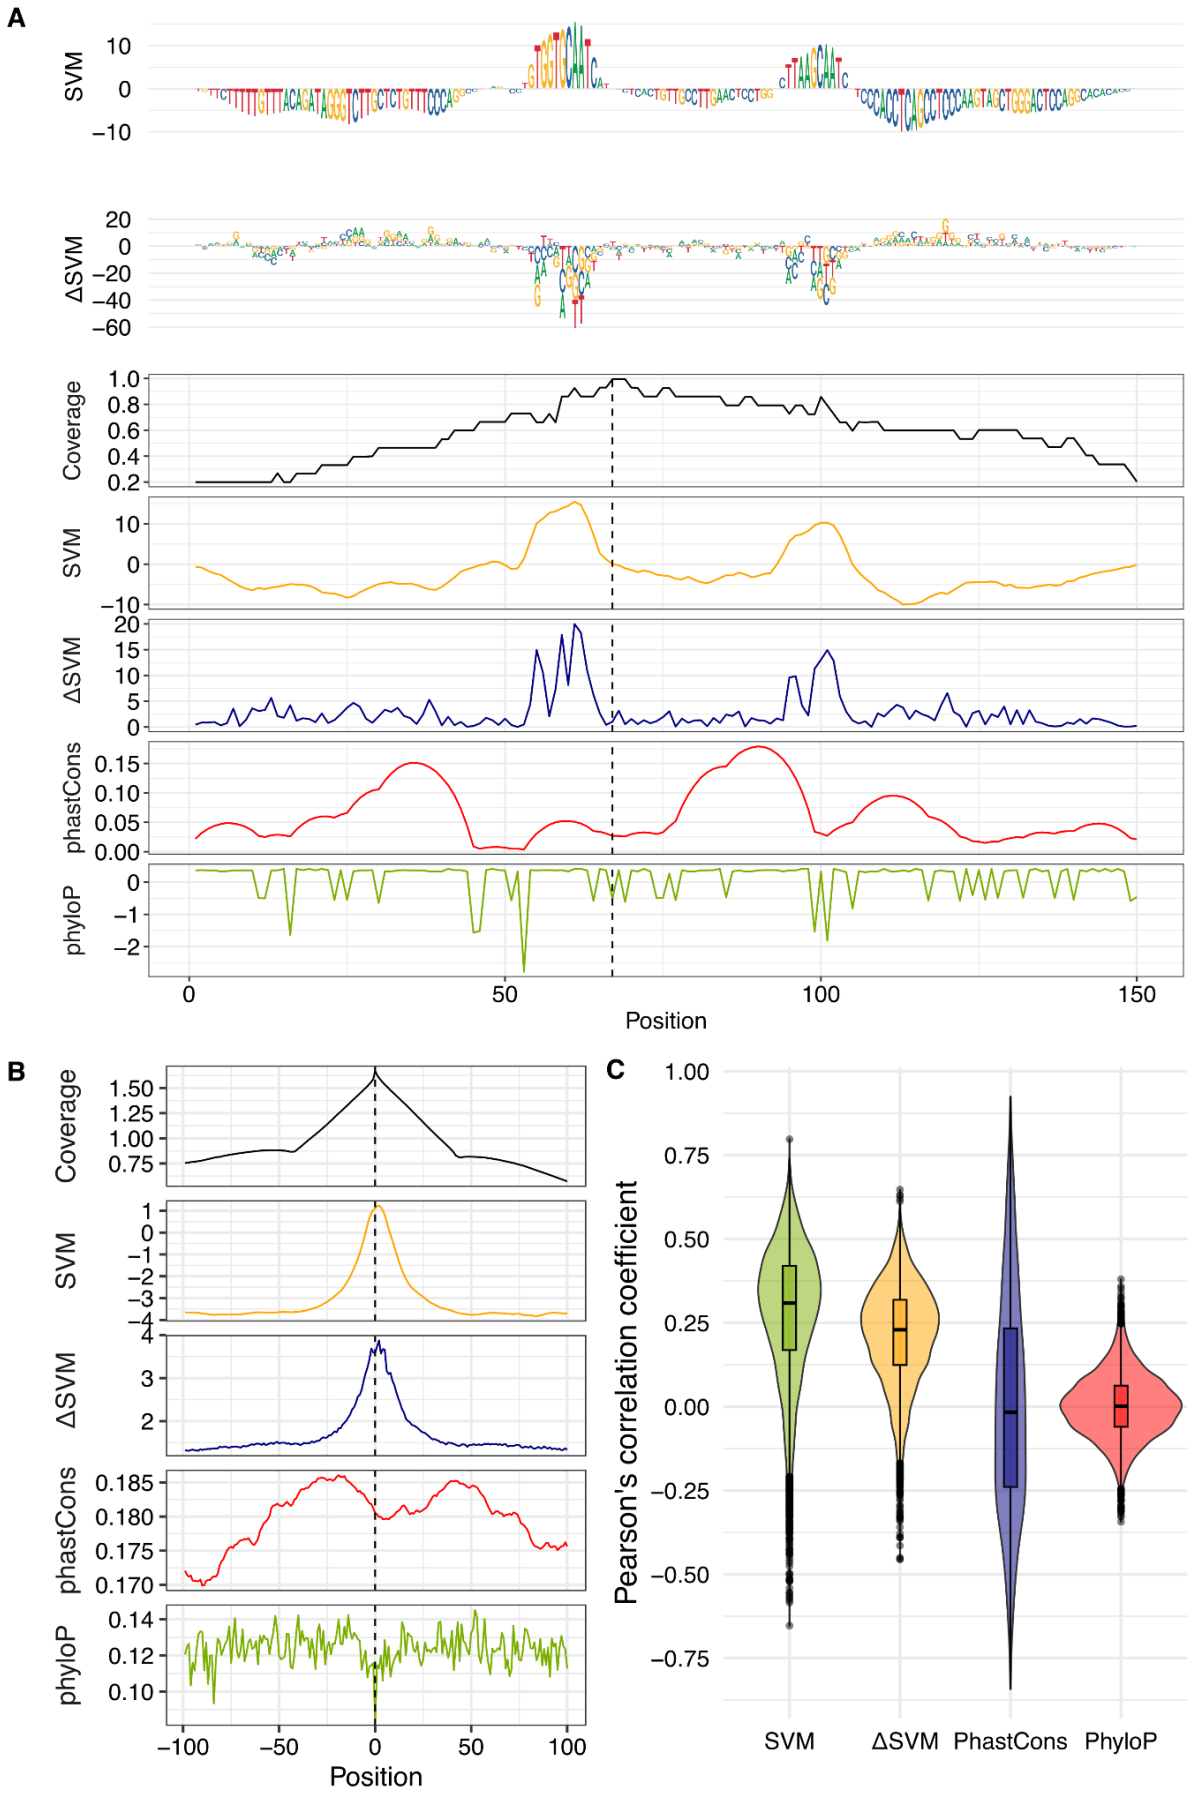
A
SVM
ΔSVM
Coverage
SVM
ΔSVM
phastCons
phyloP
Position
B
Coverage
SVM
ΔSVM
phastCons
phyloP
Position
C
Pearson's correlation coefficient
SVM
ΔSVM
PhastCons
PhyloP

Figure 1: SVM-based predictions capture functional signals of transcription factor binding. A) SVM scores computed using a sliding window across a CEBPA ChIP-seq peak in *Homo sapiens*. Each letter represents a nucleotide, with font size proportional to the SVM score of the overlapping window. Two regions of elevated scores match the canonical CEBPA binding motifs (JASPAR). The summit of the peak is shown by a dotted line. **B)** Average profiles of normalised read coverage (red), SVM prediction scores (yellow), absolute ΔSVM values (green), phastCons score (blue), and phyloP score (purple), centred on the summits of CEBPA ChIP-seq peaks. **C)** Distribution of Pearson correlation coefficients between each metric (SVM, ΔSVM, phastCons, phyloP) and the read coverage profile across individual CEBPA peaks.

Finally, we examined the consistency of model predictions across species and TFs using mammalian liver ChIP-seq datasets (Supplementary Figure 2-3). Clustering based on k-mer scores revealed that peaks grouped primarily by TF identity rather than species, indicating that the models capture conserved binding preferences. An exception was observed for HNF4 and FOXA1, which clustered by species rather than individual TF. These two transcription factors are known to co-occupy regulatory regions in the genome, acting cooperatively and sequentially to activate liver-specific gene expression (Horisawa et al. 2020), and thus there is little sequence signal differentiating their individual binding from each other.

2 - Defining genotype-to-fitness maps for TF-ChIP-seq peaks

To leverage this functional prediction by SVMs and investigate the evolutionary forces acting on ChIP-seq peaks, we developed RegEvol. This framework quantifies selection on regulatory sequences by integrating predictive models of binding affinity with evolutionary theory. This approach is illustrated in Figure 2 and detailed mathematical formalism is described in Supplementary Materials.

The first step involves constructing a genotype-to-phenotype map for each ChIP-seq peak. This is achieved by performing in silico mutagenesis on all possible point mutations within the peak and computing the change in predicted binding affinity for each mutation using ΔSVM. The resulting ΔSVM values define a Distribution of Phenotypic Effects (DPE), which captures the range of potential binding changes accessible from the ancestral sequence in one point mutation (Figure 2.A, left). Each DPE is specific to a given individual peak and provides a quantitative representation of its local genotype-to-phenotype landscape. Substitutions that occurred along the focal lineage are identified by comparing the reference sequence to its inferred ancestral state, using whole-genome alignments (Figure 2.A, right). The corresponding ΔSVM is computed for each observed substitution, representing the phenotypic effect of the mutations that were actually fixed in the lineage. These observed substitutions can be viewed as a subsample of the whole mutational landscape defined by the DPE.

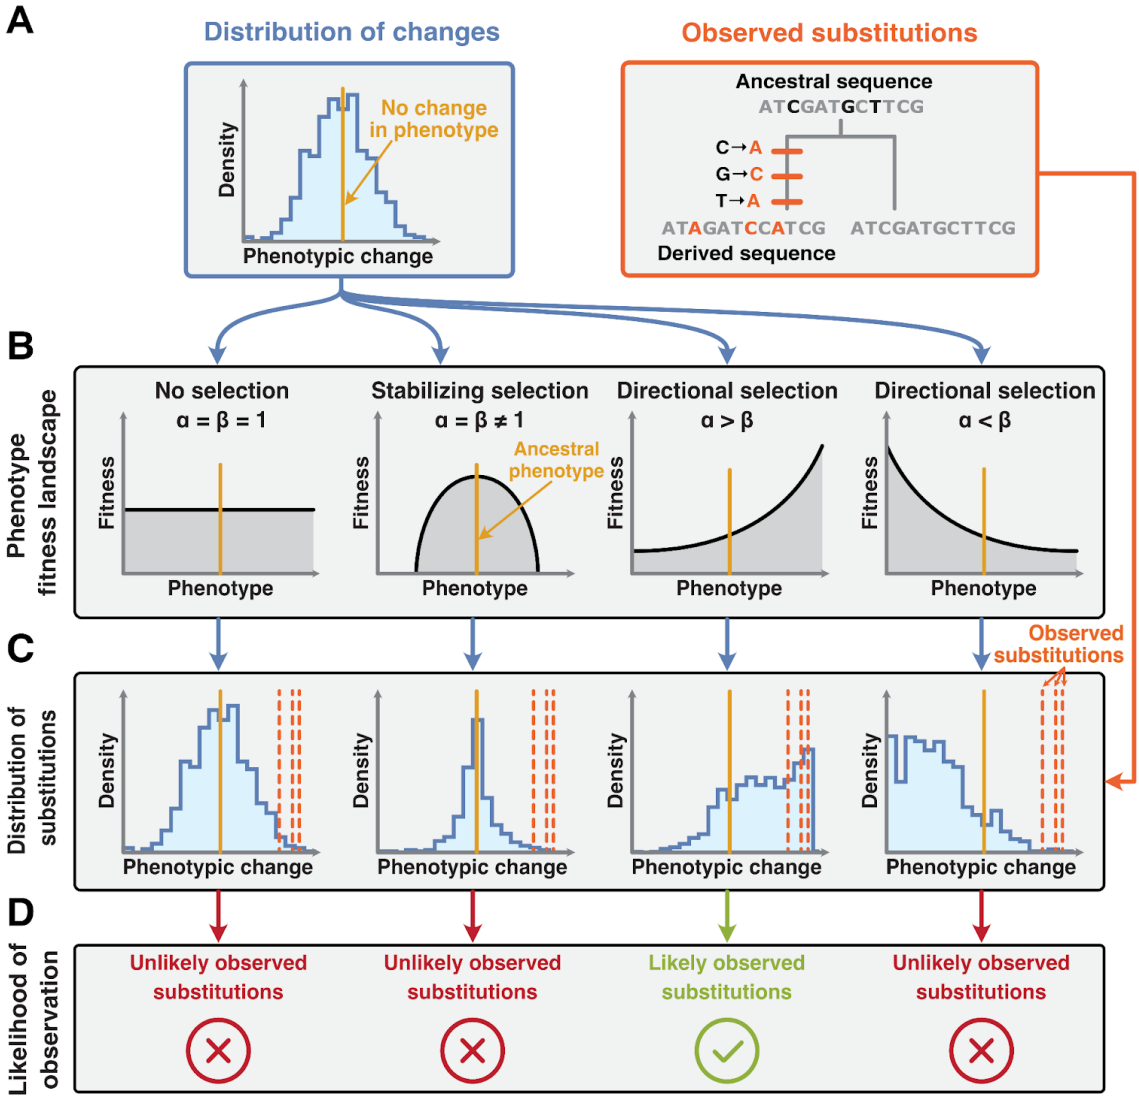

Figure 2. Outline of RegEvol: from genotype to fitness. A) Distribution of phenotypic effects (DPE) for all possible point mutations within a ChIP-seq peak (left). The observed substitutions relative to the ancestral sequence represent a subset of this distribution and are used to infer the peak's evolutionary regime (right). **B)** Definition of nested selective regimes using Beta distributions parameterized by α and β, which define the shape of the underlying fitness landscape. **C)** Expected distributions of all possible substitutions under each phenotypic fitness landscape. Vertical orange lines indicate the observed substitutions. **D)** Likelihood of the observed substitutions under each selective regime. In this example, because the substitutions are biased toward positive phenotypic change, the model of positive directional selection is the most likely.

Three nested evolutionary scenarios are defined to model how selection may have shaped these substitutions, each corresponding to a different phenotype-to-fitness map (Figure 2.B). These maps transform the predicted phenotypic effects into fitness consequences, effectively modelling a distribution of fitness effects (DFE) specific to each regulatory element. Rather than assuming a fixed DFE, we derive it from a mechanistic model of molecular function, consistent with the population genetic framework proposed by Eyre-Walker and Keightley (2007). These models are parameterised using Beta distributions.

The neutral model assumes that there is no selection on binding affinity, with mutations having a probability of fixation that does not depend on ΔSVM. This is captured by a flat fitness landscape, corresponding to a uniform Beta distribution with $\alpha = \beta = 1$. Importantly, variation in mutation and substitution rates across the genome is accounted for in all evolutionary scenarios (Materials & Methods). The stabilising selection model assumes that the ancestral binding affinity is optimal, disfavoring mutations that deviate from this optimum. This is modelled by a symmetric Beta distribution centred at ΔSVM = 0, with $\alpha = \beta \neq 1$. The shape of the distribution reflects the strength of selection, with higher values of α and β indicating stronger selection. The directional selection model allows for asymmetric fitness landscapes, favouring mutations that increase or decrease binding affinity. This is modelled by an asymmetric Beta distribution with $\alpha \neq \beta$, allowing the fitness peak to shift away from the ancestral state. Each fitness function is combined with the DPE to generate an expected distribution of fixed mutations under the corresponding selection regime. The observed substitutions are then compared to these expectations to assess which model best explains the data.

To infer the most likely evolutionary scenario for each peak, model parameters are estimated by maximising the likelihood of the observed substitutions, given the DPE (Figure 2.C). The selection coefficient for each mutation is computed as the log-ratio of fitness between the derived and ancestral states. Fixation probabilities are calculated accordingly, and the likelihood of the observed substitutions is computed under each model. Likelihood ratio tests are then used to compare models and determine the best-fitting scenario for each peak. The model with the highest likelihood, adjusted for complexity, is selected as the most probable evolutionary explanation: neutral evolution, stabilising selection, or directional selection (Figure 2.D).

3 - Detection of Directional selection on simulated peaks evolution

To evaluate the performance of our likelihood-based method, we simulated the evolution of ChIP-seq peaks under three evolutionary scenarios: random sampling (i.e., drift), skewed towards mutations with low effect (i.e., stabilising selection), and skewed towards a directional change (i.e., directional selection). For each scenario, we generated synthetic sequences by sampling substitutions according to their mutation probability, predicted phenotypic effect (ΔSVM), and fixation probability under a specified selection model. Selection strength was controlled by adjusting the parameters of the Beta distribution used to define the phenotype-to-fitness map.

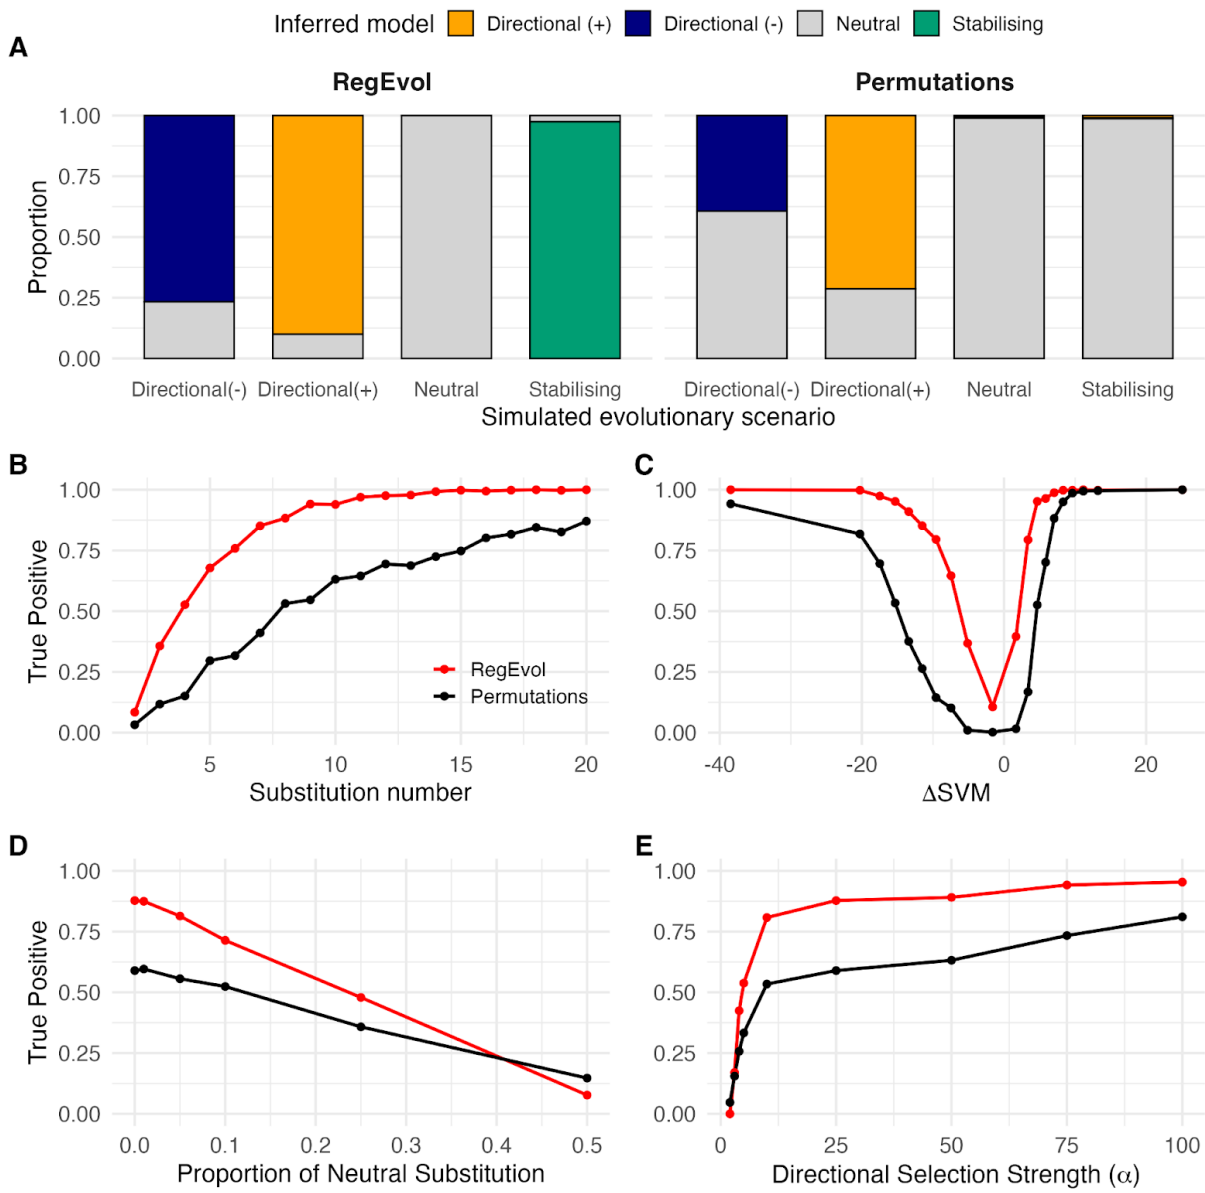

Figure 3: Detection of directional selection on simulated evolution of drosophila CTCF peaks. A) Proportion of peaks simulated under Directional, Stabilising and Random scenarios and detected by RegEvol and the Permutations Test as evolving under positive directional selection (orange), negative directional selection (blue), stabilising selection (green) and random drift (grey). **B-E)** True Positive rate of peaks simulated under directional selection (toward positive (+) or negative (-) changes) detected by RegEvol (red) or the Permutation Test (black), as a function of **B)** number of substitutions per peak, **C)** ΔSVM (quantile 5%), **D)** proportion of random substitutions, and **E)** selection strength (α parameter)**.** *N* = 10,000 simulated peaks for each evolutionary scenario; FDR ≤ 0.01.

RegEvol produced well-calibrated p-values and controlled false discovery rates across scenario comparisons, demonstrating that likelihood-based inference accurately reflects statistical confidence (Supplementary Figure 4). Across all scenarios, the method showed high accuracy in identifying the correct evolutionary model (Figure 3.A). The method showed strong power to detect directional selection, with true positive rates (TPR) of 0.90 for Positive

Directional Selection and 0.77 for Negative Directional Selection under moderate selection strength (α=10) Notably, the test was conservative: false positives (i.e cases where directional selection was inferred under neutrality or stabilising selection) remained at 1e-4 at a FDR threshold of 1%.

We quantified the factors influencing detection power by simulating the evolution of sequences under directional selection from *D. melanogaster* CTCF peaks (Figure 3.B-E), and contrasted RegEvol to our previous permutation test (Liu and Robinson-Rechavi 2020). RegEvol outperformed the permutation test across all ranges of conditions, with both higher sensitivity and specificity. As expected, detection rates increased sharply with substitution counts (Figure 3.B). Peaks with fewer than five substitutions rarely reached significance in the permutation test, showing a TPR below 0.3, and remained below 0.75 for RegEvol. The TPR for RegEvol plateaued near its maximum after about ten substitutions, whereas the permutation test increased more gradually and never reached full detection, even at 20 substitutions per peak. With only two substitutions per peak, RegEvol's TPR remained low even under strong directional selection (TPR=0.24 at α=100) but reached its maximum plateau at four substitutions (Supplementary Figure 5). This reflects the conservative nature of our likelihood framework, which requires consistent phenotypic shifts across multiple substitutions to confidently reject the null model.

We observed an asymmetry in detection power for directional selection: peaks evolving toward decreased binding affinity (i.e., negative ΔSVM) were less frequently detected (Figure 3.A and C). This reflects the typical asymmetry of the DPE across species and transcription factors, with a higher proportion of mutations showing negative effects (Supplementary Figure 6; mean ΔSVM in DPE for Drosophila CTCF peaks=-0.33, sd=1.44). Moreover, as expected, detection rates were low for simulated peaks with the smallest ΔSVM (*i.e.*, those causing little to no change in predicted binding affinity). However, while the permutation test performs poorly in this regime, RegEvol remains capable of detecting consistent directional shifts even when individual substitutions have modest effects. This is because the likelihood framework integrates directional consistency across multiple substitutions, rather than relying solely on the sum of their effects. This distinction highlights the strength of RegEvol in identifying subtle but systematic evolutionary trends.

To test the robustness of our method, we introduced a proportion of randomly fixed substitutions (i.e., mutations sampled without regard to their ΔSVM) to simulate evolutionary noise. This reflects realistic scenarios where drift is strong or multiple factors beyond the focal TF impact the evolution of regulatory elements. As expected, detection power declined

with increasing proportions of such substitutions, dropping below 0.5 when one-quarter were random (Figure 3.D). RegEvol outperformed the permutation test until half of the substitutions were random, where the two methods have TPR below 0.25. A similar pattern was observed when varying directional selection strength: both methods improved with stronger selection and plateaued at comparable levels, with RegEvol maintaining superior performance (Figure 3.E).

4 - Specificity and Ascertainment Bias

We also evaluated the sensitivity of both methods to extreme-effect substitutions, sequence divergence, and potential ascertainment bias. The permutation test, by design, is highly responsive to large ΔSVM values. While this can be advantageous for detecting isolated, high-impact substitutions, it also makes the test more prone to false positives when such mutations are fixed by chance. In simulations where substitutions were drawn independently of their ΔSVM (Figure 4.A), the permutation test showed an overall low false positive rate (FPR=0.01) but a sharp increase for extreme ΔSVM values, particularly positive ones (quantile 1% FPR=0.15; quantile 99% FPR=0.39). A similar pattern was observed in simulations under stabilising selection, where substitutions are constrained to low ΔSVM values (Figure 4.B, quantile 99% FPR=0.26). In contrast, RegEvol remained stable and conservative across all simulations, as it integrates information across all substitutions and emphasizes consistent directional trends rather than outliers (maximum FPR across quantile=1e-4).

Distinguishing stabilising selection from random drift is inherently difficult, as most substitutions have low predicted effects on TF binding affinity (Supplementary Figure 6). We therefore assessed RegEvol's performance in detecting stabilising selection across a range of simulation parameters (Supplementary Figure 7). Moreover, because of the asymmetry of the DPE toward negative changes, the distribution of ΔSVM in random permutations becomes increasingly negative with the number of substitutions. Consequently, under stabilising selection, the false discovery rate of the permutation test positively correlates with sequence divergence and can reach a FPR of 0.05 with high number of substitutions (Figure 4.C, Spearman's rho=0.96, p=7.4e-16). By integrating the ΔSVM distribution as a DFE for each peak and explicitly modeling stabilising selection around ΔSVM = 0, RegEvol effectively controls for this bias (FPR=0).

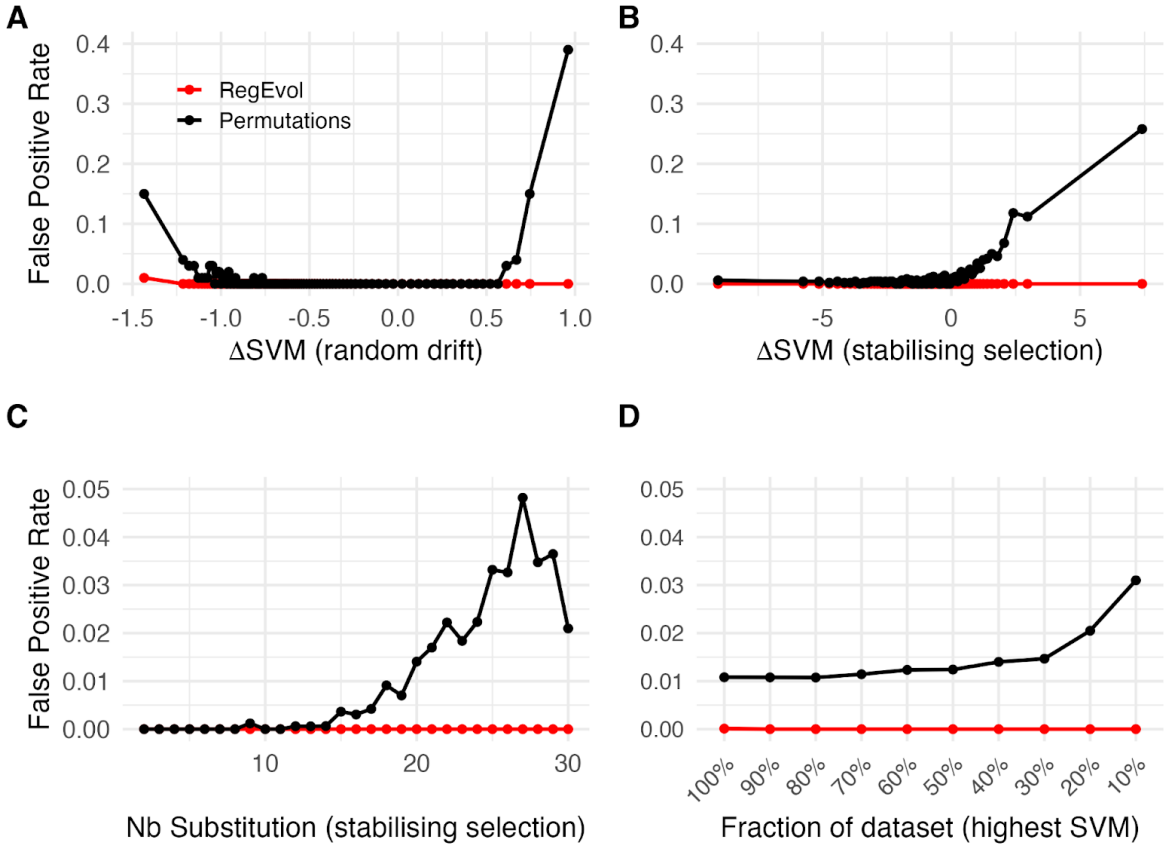

Figure 4. Sensitivity to extreme substitutions and robustness to ascertainment bias. Proportion of simulated *Drosophila* CTCF peaks detected under directional selection using the Permutation Test (black lines) and RegEvol (red lines). Substitutions were either drawn randomly (**A, D**) or under stabilising selection (**B, C**; mean(α=β)=45). **A-B)** False positive rate across 1% quantiles of ΔSVM; **C)** false positive rate as a function of the number of substitutions; **D)** fraction of datasets stratified by the highest derived SVM scores. *N* = 10,000 simulated peaks; FDR ≤ 0.01.

An ascertainment bias on ChIP-seq peaks has been suggested to affect permutation tests, as higher SVM score peaks are both more likely to be called and to show high positive ΔSVM (Jiang and Zhang 2024). To test this, we subsampled peaks by derived ΔSVM in random simulations and assessed the proportion detected under directional selection. RegEvol maintained a FPR close to 0 across peak categories, while the permutation test showed a positive association with derived ΔSVM reaching FPR above 0.03 (Figure 4.D). Applying the same subsampling to the previously analysed human CEBPA dataset (Liu and Robinson-Rechavi 2020), the correlation for RegEvol remained low, with only the highest SVM category showing a slight increase, potentially reflecting a true biological signal (Supplementary Figure 8). These results demonstrate that RegEvol is robust to peak strength, supporting its reliability in real datasets where ChIP-seq signal intensity may influence peak calls.

To evaluate practical limitations of RegEvol, we examined how detection power is affected by the number of substitutions per regulatory element and multiple testing correction (Supplementary Figures 9-10). In empirical datasets, many peaks do not have enough substitutions, particularly in closely related species. For example, in mammalian datasets, the median number of substitutions per peak can be low depending on the divergence since the closest ancestral state (human median substitution number = 2.8), which reduces statistical power and leads to a scarcity of low p-values after multiple testing correction (Supplementary Figure 10). In Drosophila, higher divergence results in more substitutions per peak, providing a more stable detection power across regulatory elements (Supplementary Figure 9). Given these observations and the low true positive rate for peaks with few substitution even under strong simulated directional selection (Supplementary Figure 5), we restricted our analyses to peaks with at least four substitutions in vertebrates datasets (Supplementary Figure 11). This correspond to an average of 1.7% sequence divergence (median length=291bp; Supplementary Figure 12-13).

4 - Analysis of *Drosophila melanogaster* ChIP-seq peaks

To illustrate the utility of RegEvol, we applied it to a large-scale dataset of transcription factor binding sites in *Drosophila melanogaster*. Using ChIP-seq peaks from the modERN consortium (Kudron et al. 2024), we analysed regulatory elements across multiple developmental stages and transcription factors (Supplementary Data Table 2). For each peak, we computed the DPE using gkm-SVM models trained on the corresponding TF and inferred the most likely selection regime since the divergence from *D. simulans*, based on ancestral sequences reconstructed from whole-genome alignments.

We inferred a total proportion of 5.1% of fruit fly peaks under directional selection (Figure 5A), out of 2.8M tested. This is not due to high permissiveness of the test, since on mammalian ChIP-seq peaks we find much lower proportions of directional selection (Supplementary Figure 11). This is consistent with reports of a higher proportion of amino acid substitutions fixed by directional selection, and generally higher efficiency of selection, in the fruit fly than in mammals (Eyre-Walker 2006; Sella et al. 2009; Lin et al. 2025). Among the fly peaks, there is considerable variation between TFs and conditions, from almost none to 50% of peaks inferred under directional selection. With the low developmental resolution of those data, there is no strong trend of differences over ontogeny but there is less signal of directional selection in prepupal and pupal stages (Supplementary Figure 14).

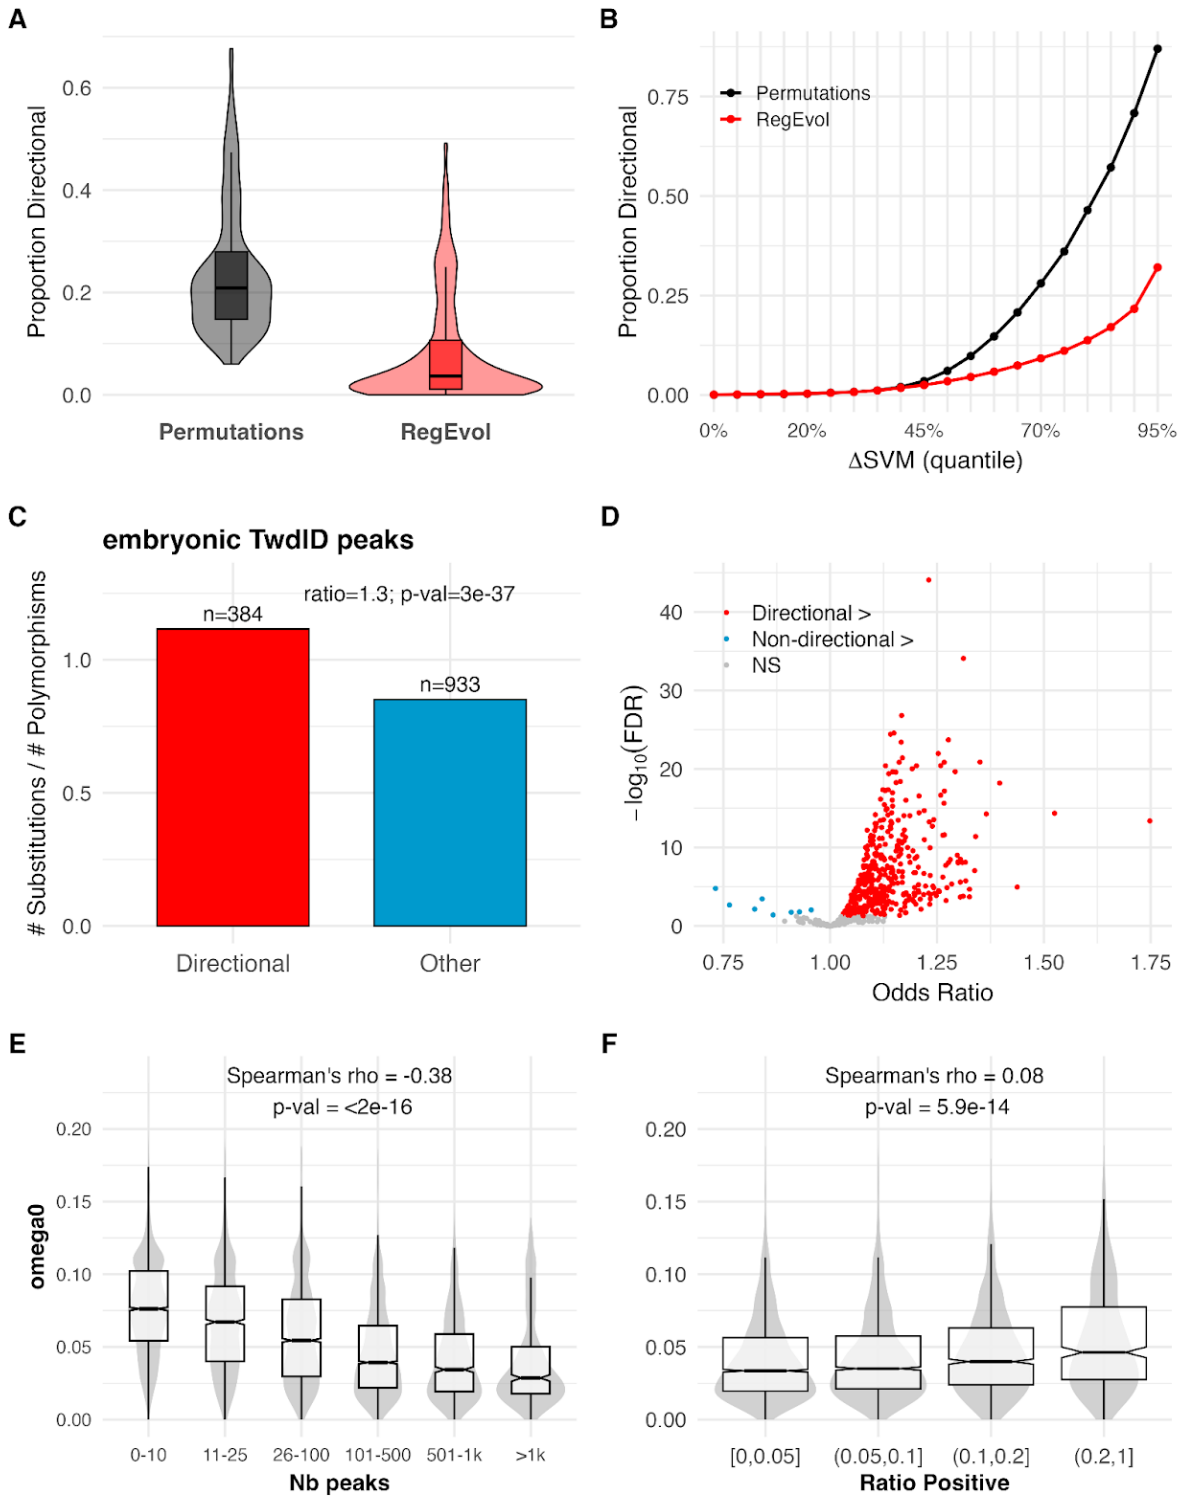

Figure 5. Application of RegEvol to Drosophila melanogaster transcription factor binding sites. A. Proportion of peaks inferred to be under directional selection using the original permutation test (black; $p \leq 0.01$) and RegEvol (red; FDR ≤ 0.05). **B.** Proportion of peaks under directional selection across the 5% quantiles of derived ΔSVM values for all peaks. **C.** Ratio of substitutions to single-nucleotide polymorphisms (SNPs) in embryonic *TwdID* peaks inferred by RegEvol under directional selection (red; $N = 385$) or not (blue; $N = 877$). The odds ratio and Fisher's exact test between the two categories are indicated. **D.** Odds ratios between directional and non-directional peaks across datasets (restricted to experiments with ≥10 directional peaks). Experiments where directional peaks show higher

397 ratios (red; odds ratio > 1), lower ratios (blue; odds ratio ≤ 1), or no significant difference
398 (grey; FDR > 0.05) are shown. **E-F.** Relationship between purifying selection on
399 protein-coding genes (ω_0) and the total number of associated peaks (**E**) or the proportion of
400 their peaks inferred under directional selection (**F**). Spearman's correlation coefficients were
401 computed across all genes.

402 As expected, the signal for directional selection increases with the ΔSVM (Figure 5B): peaks
403 with larger predicted phenotypic effects are more often inferred to have evolved under
404 directional selection. We observe much more positive than negative directional selection, a
405 pattern likely biological rather than methodological, since it is absent from simulations. We
406 detect few peaks under stabilising selection, which can be explained by the framing of our
407 test relative to the nature of purifying selection: stabilising selection tends to eliminate
408 mutations, leading to few substitutions and a low power of the likelihood test. The role of
409 stabilising selection in our model is essentially to avoid false positives in the test of
410 directional selection; however, our approach is poorly adapted to inferring the proportion of
411 peaks evolving under stabilising selection *per se*.

412 A key prediction of recent directional selection is a local reduction in polymorphism due to
413 selective sweeps, together with increased inter-species divergence from fixation of mutations
414 (McDonald and Kreitman 1991; Fu and Akey 2013). Accordingly, we expect a higher
415 substitution-to-SNP ratio in peaks inferred to have evolved under directional selection. This
416 expectation is confirmed in *D. melanogaster*, with almost all experiments supporting this
417 pattern (74.8% experiment, odds ratio > 1 and Fisher's exact test FDR ≤ 0.05) (Figure 5 C,
418 D). Because our model gains power with increasing numbers of substitutions, this pattern
419 could in principle reflect a bias. To evaluate this, we compared substitutions and SNPs
420 separately and found that peaks under directional selection have both significantly more
421 substitutions and fewer SNPs (Wilcoxon test p-values ≤ $5e10^{-4}$; Supplementary Figure 15),
422 consistent with recent selective sweeps rather than a power artefact. This provides
423 independent validation, as SNP frequencies are not used by the model, and shows that our
424 k-mer-based DFE inference aligns with classical expectations of natural selection.

425 While it is difficult to reliably associate TF ChIP-seq peaks with the genes they regulate, the
426 closest genes provide a good approximation, especially in compact genomes such as that of
427 *D. melanogaster* (Kudron et al. 2024). As expected, protein-coding genes under stronger
428 purifying selection (ω_0) are associated with more TF peaks (Figure 5E), thus presumably
429 more complex regulation and more pleiotropy. There is also an association, although weaker,
430 between selection on protein-coding genes and on regulatory sequences (Spearman's
431 rho=0.08, p-values=8e-16, Figure 5F): proteins under weaker purifying selection are
432 associated with a higher proportion of peaks under directional selection. We do not find an

association between adaptive selection on protein-coding genes (ω2) and the proportion of directional selection on their associated peaks.

Genes associated with peaks under directional selection are enriched for morphogenesis and developmental processes (Supplementary Figure 16). However, genes linked to a larger total number of peaks are also overrepresented in these functional categories. When considering the ratio of peaks under directional selection, no significant functional enrichment was observed (Supplementary Figure 16). The lack of functional enrichment may reflect a limitation of this approach. By aggregating TF peaks from 740 different experiments and contexts, the resulting gene set may be too large and functionally heterogeneous for conventional functional tests. To explore whether signals of directional selection could be detected at a broader scale, we applied TopAnat (Komljenovic et al. 2018a; Bastian et al. 2021), which identifies anatomical structures where a gene set is preferentially expressed. Genes with a high proportion of selected TF peaks (≥0.2) were enriched in reproductive tissues (testis, seminal fluid glands, ovary) and immune-associated organs (fat body, Malpighian tubule) (Supplementary Figure 17). These tissues are major sites of sexual selection and immune processes (Buchon et al. 2014; Vaibhvi et al. 2022), which are known to evolve rapidly, supporting the biological relevance of our directional peaks.

Discussion

Understanding how regulatory sequences evolve is critical for linking genotype to phenotype, yet most evolutionary analyses of non-coding DNA still rely on substitution rates or sequence conservation (Adam Siepel et al. 2005; Pollard et al. 2010). These approaches are powerful for identifying constraints but often lack sensitivity to detect selection acting on specific molecular functions (Ward and Kellis 2012; Villar et al. 2015), especially directional selection. RegEvol addresses this limitation by testing how predicted mutational effects on transcription factor (TF) binding have been shaped by selection. By linking quantitative predictions of regulatory activity to patterns of sequence divergence, RegEvol provides a functionally informed and mechanistically interpretable view of regulatory evolution. In contrast to motif-centric models, the predictive component of RegEvol is trained on entire ChIP-seq peaks, capturing local sequence context and potential co-factors that influence TF binding (Yan et al. 2021). This shift from rate-based to phenotype-based inference enables the detection of subtle selection signals that may not produce strong shifts in sequence conservation but shape regulatory activity.

RegEvol builds on our previous permutation-based framework (Liu and Robinson-Rechavi 2020), which evaluated selection by comparing the summed ΔSVM of observed substitutions to a null distribution derived from random permutations. While this approach provided a

useful baseline, it was sensitive to extreme substitution effects, lacked robust multiple-test correction, and was affected by ascertainment bias in the choice of regulatory peaks (Jiang and Zhang 2024). RegEvol overcomes these limitations by explicitly modelling three selective regimes: neutral, stabilising, and directional. Each regime is defined by a distinct fitness function that transforms the distribution of predicted phenotypic effects on TF binding into an expected substitution pattern. Using a maximum-likelihood framework, our method identifies the selective regime that best explains the observed substitutions. Moreover, because this inference is analytical rather than simulation-based, p-values remain well calibrated and support rigorous multiple-test correction, avoiding one of the key weaknesses of the permutation approach (Jiang and Zhang 2024).

A key advantage of RegEvol is its nested modeling of selective regimes, allowing for a likelihood-based comparison of directional and stabilising selection. This is important because the expectation under stabilising selection is centred on no change in the phenotype ($\Delta SVM \approx 0$), whereas the null distribution in the permutation test is typically slightly negatively shifted. This asymmetry, has little impact when few substitutions are available but leads to a marked inflation of false positives with increasing divergence (Figure 4). Including a stabilising regime corrects this bias and improves robustness by providing an intermediate model between neutrality and directional selection.

RegEvol remains robust across a wide range of divergence levels: it is conservative at short timescales, minimising false positives, and gains statistical power while maintaining specificity on longer branches. Even when the overall signal is modest, RegEvol can still detect directional selection when substitution effects are consistent. This focus on the consistency of substitution effects rather than their magnitude reduces false positives driven by a few extreme mutations and mitigates biases due to peak ascertainment. This leads to more conservative selection inference, an advantage when detecting positive selection, where minimising false positives is critical (Jordan and Goldman 2012; Venkat et al. 2018; Wisotsky et al. 2020; Soni et al. 2023). Conversely, a few strong-effect mutations may not be sufficient to support a confident inference, illustrating that the method prioritises robustness over sensitivity. This trade-off contrasts with the permutation test, which interprets a handful of large-effect substitutions as evidence of directional selection. While the latter approach may be more powerful when adaptation truly depends on such rare events, it does so at the cost of an elevated false-positive rate. The reduced power of RegEvol in cases with few substitutions is not a methodological bias but an inherent limitation of statistical inference, common to all selection tests. Because the statistical power of RegEvol necessarily depends on the number of substitutions within a peak, very low divergence limits the ability to reliably

infer a fitness function, leading to reduced sensitivity across selective regimes. In practice, this means that RegEvol is most informative for elements with sufficient evolutionary change, while low divergent regions remain better characterised by conservation-based methods, which complement rather than compete with functionally informed approaches.

By evaluating the full distribution of substitution effects rather than their sum, RegEvol can also detect compensatory changes in which opposing substitutions collectively maintain overall function, corresponding to a disruptive or balancing pattern of evolution. This is especially important in regulatory regions, which can evolve conservatively under turn-over of binding sites (Moses et al. 2006; Krieger et al. 2022). However, because it assumes additive contributions of individual substitutions, the model does not capture non-additive epistatic effects, such as synergistic interactions between mutations, which remain a limitation of the method.

RegEvol also enables branch-specific analyses, providing a framework to identify adaptive regulatory changes that occurred along particular evolutionary lineages. In practice, the method requires both alignable sequence data between the focal species and an ancestral sequence reconstruction, as well as lineage-specific functional data to train predictive models of TF binding. This dependence on empirical data currently limits large-scale comparative scans, but recent advances in deep learning models trained on extensive functional genomics datasets are beginning to mitigate this constraint. Such models can transfer regulatory knowledge across species to infer changes in TF binding affinity or even in the underlying motif architecture, providing increasingly accurate functional priors for non-model organisms (Avsec et al. 2021; Benegas et al. 2023; Karollus et al. 2024). By integrating these predictive advances, RegEvol could eventually support systematic analyses of lineage-specific selection on regulatory elements, connecting molecular adaptation to phenotypic evolution with greater resolution.

By decoupling substitution rate from phenotypic effect, RegEvol enables a more functionally informed view of regulatory evolution, analogous to how coding sequence evolution is interpreted through the nature of amino acid changes rather than their frequency (Halpern and Bruno 1998; Rodrigue et al. 2021). While coding sequences are typically constrained by protein structure and pleiotropy, regulatory elements evolve under different pressures, including redundancy, turnover, and modularity, which enhance evolvability but complicate evolutionary analysis (Wittkopp and Kalay 2012; Kim and Wysocka 2023). Application of RegEvol to Drosophila ChIP-seq data illustrates its ability to detect consistent directional trends in regulatory sequences across varying substitution numbers, distinguishing phenotypic shifts from random drift. Approximately 5.1% of fly peaks were inferred under

directional selection, a pattern consistent with previous reports of adaptive evolution in Drosophila TFBS (He et al. 2011) and with the higher efficiency of selection in flies compared to mammals (Eyre-Walker 2006; Sella et al. 2009; Lin et al. 2025). Peaks under directional selection showed elevated substitution-to-SNP ratios, consistent with recent selective sweeps, and genes associated with these peaks were enriched in reproductive and immune tissues, supporting the biological relevance of the inferred selection. This framework thus provides a mechanistic link between regulatory sequence evolution and downstream phenotypic effects, offering new perspectives on the co-evolution of regulatory elements and associated coding sequences, and more broadly connecting regulatory variation to organismal adaptation.

Improving the genotype-to-phenotype map will be essential to extend the reach of RegEvol. While gkm-SVM models remain interpretable and TF-specific, capturing the additive effects of individual substitutions, they have limited capacity to represent complex regulatory features such as motif syntax, long-range dependencies, or non-additive interactions (de Boer and Taipale 2024). Recent deep learning models trained on high-resolution functional genomics data have shown the ability to learn these features directly from sequence, providing more accurate, context-aware predictions of mutational impact (Avsec et al. 2021; Benegas et al. 2023; Karollus et al. 2024). Additionally, integrating massively parallel reporter assays and other synthetic functional assays can empirically validate predicted effects and uncover subtle regulatory principles not captured by sequence-based models alone. Because RegEvol is agnostic to the underlying predictive model, it can incorporate these advances, producing more informative distributions of phenotypic effects and improving sensitivity for detecting selection in complex regulatory landscapes. Regulatory elements often integrate inputs from multiple TFs, chromatin states, and cellular contexts, and the functional impact of a mutation may vary across these layers (Hu and Tee 2017; Kurafeiski et al. 2018). Probabilistic frameworks, such as Bayesian or mixed-effect models, could capture heterogeneous and overlapping selective pressures within a single element, capturing this complexity more realistically than a single fitness landscape per element.

Together, these developments position RegEvol as a flexible and extensible framework for studying the evolution of regulatory sequences in a functionally informed manner. By integrating predictive models of molecular phenotypes, which link genotype to quantitative regulatory effects, with well-established evolutionary equations that map phenotypes to fitness, RegEvol provides a principled approach to detect directional selection, complementing traditional conservation-based analyses. As functional genomics, predictive modelling, and comparative data continue to advance, RegEvol can help bridge the gap

between regulatory sequence variation and phenotypic evolution, offering new insights into the molecular mechanisms underlying adaptation across evolutionary timescales.

Materials and Methods

Catalogue of Drosophila melanogaster Transcription Factor Binding Sites

We used pre-processed ChIP-seq peak data for *Drosophila melanogaster* TFBS from the modENCODE (model organism Encyclopedia of DNA Elements) and modERN (model organism Encyclopedia of Regulatory Networks) consortia (Kudron et al. 2024). It provides a harmonised dataset across 740 experiments covering 645 TFs and five developmental stages (Supplementary Data Table 2).

ChiP-seq peaks detection and coverage quantification

We collected publicly available ChIP-seq data for five transcription factors (TF) across ten species: *Homo sapiens*, *Macaca macaca*, *Mus musculus*, *Mus spretus*, *Mus caroli*, *Rattus norvegicus*, *Canis lupus familiaris*, *Felis catus*, *Oryctolagus cuniculus, Gallus gallus*, *Drosophila melanogaster* (Ballester et al. 2014; Schmidt et al. 2010; 2012; Rensch et al. 2016; Stefflova et al. 2013) (Supplementary Data Table 1). To ensure consistency across datasets, we re-processed all raw sequencing data using the NextFlow ChIP-seq pipeline v2.0 (Ewels et al. 2022), specifying Bowtie2 (Langmead and Salzberg 2012) as the aligner and using the narrow_peak option for peak calling with MACS2 (Zhang et al. 2008). For downstream analyses, we retained only peaks located on complete chromosomes, excluding unplaced scaffolds.

Read coverage was computed at each base within ChIP-seq peaks using bamCoverage, normalised by counts per million (Ramírez et al. 2014). For each experiment, we calculated the consensus peak coverage by summing normalised read counts across replicates. Peak summits were defined as the position with the highest normalised read coverage within each peak. In cases where multiple positions shared the maximum value, we selected the position with the highest mean coverage across a 10 bp window centred on the candidate site.

Whole genome alignment and ancestral state

We retrieved publicly available whole-genome alignments generated with Progressive Cactus. For *Drosophila melanogaster*, we used a 20-species alignment (Peng and Zhao 2024), and for vertebrates, we used the 241-species alignment from the Zoonomia Consortium (Armstrong et al. 2020) . These alignments include both reference genomes and inferred ancestral genomes at each internal node of the species tree.

For each analysed species, we used the reference genome as the query and extracted alignments with the closest related species and their most recent common ancestor. Aligned sequences were retrieved using the hal2maf tool from the Comparative Genomics Toolkit (Hickey et al. 2013), with the noDupes option enabled to exclude putative duplicated regions. We extracted the alignment of each set of ChIP-seq peaks using their genomic coordinates as input for the mafsInRegion tool from UCSC (Nassar et al. 2023).

gkm-SVM model training and validation

To model the relationship between DNA sequence and TF binding affinity, we trained gapped k-mer support vector machine (gkm-SVM; Ghandi et al. 2014) models for each TF and tissue. The positive training set consisted of genomic regions identified as ChIP-seq peaks in the reference genome. The negative set was generated by randomly sampling genomic sequences matched for length, GC content, and repeat content using the genNullSeqs function from the gkmSVM R package (Ghandi et al. 2016). When pre-built BSgenome objects were not available for the genome assemblies analysed, we constructed them following the BSgenome R package guidelines (Pagès 2024), using annotations from UCSC and GenBank (Supplementary Data Table 1).

Each model was trained using the gkmtrain function from the LS-GKM software (Lee 2016), with default parameters except for a k-mer length of 10 bp. Model performance was evaluated using 5-fold cross-validation using the -x 5 option. After training, we used the gkmpredict function to compute SVM weights for all possible 10-mers. Experiments with fewer than 1,000 peaks or with an associated model AUC below 0.8 were excluded from downstream analyses (Supplementary Figure 18).

SVM Score Computation for Peaks and Positions

We computed SVM scores for each ChIP-seq peak at two levels. At the peak level, the total SVM score was calculated as the sum of SVM weights from a 10 bp sliding window with a 1 bp step size. At the positional level, the SVM score for a specific nucleotide position was defined as the sum of SVM weights for all overlapping windows that include that position. These two levels of scoring were used for downstream analyses of both peak-level and site-specific effects.

In Silico Mutagenesis and ΔSVM Computation

For each ChIP-seq peak, we performed in silico mutagenesis by introducing all possible single-nucleotide mutations. The change in predicted binding affinity for each mutation was computed as ΔSVM, defined as the difference in SVM score between the mutated and

ancestral sequence. This yielded a Distribution of Phenotypic Effects (DPE) for each peak, representing the genotype-to-phenotype landscape.

Observed substitutions were identified by comparing the reference sequence to its inferred ancestral state. Indels were removed using TrimAl with the “nogaps” option (Capella-Gutiérrez et al. 2009). SVM scores were computed for both the reference and ancestral sequences using the same sliding window approach, and ΔSVM values were calculated for each observed substitution.

Substitution matrix inference

To account for variation in mutation and substitution rates across the genome, we estimated substitution matrices for each chromosome. Whole-genome alignments were split by chromosome using the mafSplit tool, and exonic regions were masked using GFF annotations. We then ran PhyML (Guindon et al. 2010) on each chromosome alignment using a Generalised Time-Reversible model with the params=r option to optimise substitution rate parameters. These matrices were used to weight mutation probabilities in the evolutionary model.

Evolutionary Model Inference via Likelihood Optimisation

To infer the selection regime acting on each ChIP-seq peak, we modeled the evolutionary process using three nested phenotype-to-fitness maps parameterized by Beta distributions: (i) a neutral model with a flat fitness landscape ($\alpha = \beta = 1$), (ii) a stabilising selection model with maximal fitness at the ancestral phenotype ($\alpha = \beta \neq 1$), and (iii) a directional selection model allowing asymmetric fitness landscapes ($\alpha \neq \beta$).

For each possible single-nucleotide mutation within a peak, we computed the predicted phenotypic effect (ΔSVM), the corresponding selection coefficient (as the log-ratio of fitness between derived and ancestral phenotypes), and the fixation probability under an origin-fixation model. These were combined with mutation probabilities derived from chromosome-specific substitution matrices to calculate a substitution probability for every in silico mutation.

The likelihood of the observed substitutions was then computed by summing the log-probabilities of the corresponding mutations under each model. Model parameters (α, β) were optimised using the scipy.optimize.minimize function (Virtanen et al. 2020). To compare the performance of each model and determine the most likely evolutionary scenario for each peak, we performed likelihood ratio tests using a threshold value of 1%. A detailed mathematical formalism of this framework is provided in Supplementary Materials.

Improvement on the Permutation Test

We refined the previously published Permutation Test to improve accuracy and interpretability (Liu and Robinson-Rechavi 2020). Mutation and substitution rates are now incorporated using chromosome-specific substitution matrices. Substitution positions are sampled without replacement to match the observed count, avoiding undersampling artefacts. We also implemented a two-sided test to detect both increases and decreases in binding affinity, and applied false discovery rate (FDR) correction to account for multiple testing.

Transcription factor binding motif recovery

To assess whether the gkm-SVM models captured known TF binding preferences, we extracted the SVM weights for all possible 10-mers from the trained models for human and mouse datasets. We selected the top 100 scoring 10-mers for each model and aligned them using Clustal Omega with default parameters (Sievers et al. 2011). The resulting alignments were submitted to the TomTom motif comparison tool (Gupta et al. 2007) in the MEME Suite website (Bailey et al. 2015) to identify matches against known TF motifs in the JASPAR CORE 2022 database (Castro-Mondragon et al. 2022).

Sequence conservation metrics

To compare model-based predictions with evolutionary conservation, we retrieved per-base conservation scores from the UCSC Genome Browser (Perez et al. 2025). Specifically, we used phastCons (Siepel et al. 2005) and phyloP scores (Pollard et al. 2010) from the 17-way vertebrate alignment for human, the 60-way glire subset for mouse, and the 27-way alignment for *Drosophila melanogaster*. For each ChIP-seq peak, we computed the conservation score at every nucleotide position and summarised the signal by calculating the average score across the peak.

Simulation of sequence evolution

To assess the accuracy and robustness of our selection inference framework, we simulated the evolution of regulatory sequences under controlled evolutionary scenarios. From randomly sampled drosophila CTCF ChIP-seq peak, we computed the predicted phenotypic effect (ΔSVM values) for all possible single-nucleotide mutations, along with mutation probabilities derived from chromosome-specific substitution matrices.

Substitutions in these peaks were sampled probabilistically according to their mutation rate, predicted phenotypic effect, and fixation probability under a specified evolutionary model. These models included neutral evolution, stabilising selection around the ancestral phenotype, and directional selection favouring increased or decreased binding affinity. Each

model was parameterised using Beta distributions to define the shape and strength of the phenotype-to-fitness map.

Simulations were performed across a range of conditions, varying the number of substitutions per sequence, the strength of selection and the proportion of neutral mutations. This allowed us to generate synthetic datasets reflecting different evolutionary regimes, which were then used to benchmark the performance of RegEvol and the Permutations Test.

Drosophila divergence and polymorphism

Polymorphism data from the *Drosophila melanogaster* Genetic Reference Panel (DGRP) (Mackay et al. 2012) were retrieved and lifted over from the dm3 to the dm6 genome assembly using LiftOver (Hinrichs 2006). Overlap with TF peaks was determined using bedtools intersect. Fixed substitutions between *D. melanogaster* and *D. simulans* were counted within each peak, and a substitution-to-SNP ratio was calculated for every peak. Fisher's exact tests were then applied for each experiment to compare peaks inferred to have evolved under directional selection against all other peaks.

Gene Ontology and Anatomical Enrichments

Associations between ChIP-seq peaks and genes were obtained from the modERN consortium based on genomic proximity (Kudron et al. 2024). Gene Ontology enrichment analysis was performed using the *clusterProfiler* R package (version 4.16) with gene annotations from *org.Dm.eg.db* (version 3.21.0) for *Drosophila melanogaster* (Wu et al. 2021). Enrichment was computed against a background of all genes with at least one associated peak. We tested genes with at least one peak under directional selection, ranked gene lists based on either the total number of peaks or the number of peaks under directional selection, and genes with a directional-to-total peak ratio above 0.25 compared to all genes with at least two associated peaks.

Anatomical enrichment was carried out using TopAnat in the BgeeDB R package (version 2.34.0) with Bgee release 15.2 (Komljenovic et al. 2018; Bastian et al. 2025). TopAnat identifies anatomical entities where a gene set is preferentially expressed based on expression calls integrated by Bgee across anatomical structures. The analysis was restricted to RNA-seq expression calls and limited to *Drosophila melanogaster* fully formed stage (UBERON:0000066). Only the gene set with a directional-to-total peak ratio above 0.25 and the corresponding background with at least two peaks was tested. Enrichment was assessed using the *weight01* algorithm (Alexa et al. 2006), and anatomical entities with FDR ≤ 0.01 were considered significant.

Drosophila coding sequences evolution

To assess whether genes associated with TF peaks also show signatures of coding sequence evolution, we retrieved ω_0 and ω_2 estimates from Selectome (Moretti et al. 2014). Selectome applies branch-site codon models (Davydov et al. 2019) to orthologous gene alignments to estimate the strength of purifying selection (ω_0) and positive selection (ω_2) on specific phylogenetic branches. These values were extracted for *D. melanogaster* genes associated with TF peaks and used to test for correlations with the number of peaks and the directional-to-total peak ratio.

Computation and Statistics

All data analyses were performed using R (version 4.5.0), Python (version 3.10), and Snakemake (version 7.24.0) for workflow management. The analyses were conducted on the high-performance computing infrastructure of the DCSR at the University of Lausanne. Multiple testing correction was applied using the Benjamini-Hochberg procedure to control the false discovery rate (FDR) across genomic regions. To assess potential ascertainment bias, we stratified ChIP-seq peaks by their SVM scores and evaluated detection rates across quantiles.

Data Availability

The RegEvol pipeline, along with all scripts used to generate the analyses and figures, is available at the following stable GitHub URL: https://github.com/mrrlab/RegEvol. Additionally, all RegEvol results and data generated for this study are publicly available at the Zenodo repository: https://zenodo.org/records/17867860.

Author Contributions

M.R-R and A.L. originally conceived the project, A.L. and T.L. designed and implemented the method and computational framework. A.L. conducted all analyses and wrote the initial draft of the manuscript. All authors discussed the results and contributed to the final version. M.R-R supervised the project.

Funding

This work was supported by the Swiss National Science Foundation grant 207853 to Marc Robinson-Rechavi and grant 219757 to Nicolas Salamin.

Competing Interests

The authors declare no competing interests.

References

Alexa, Adrian, Jörg Rahnenführer, and Thomas Lengauer. 2006. "Improved Scoring of Functional Groups from Gene Expression Data by Decorrelating GO Graph Structure." *Bioinformatics* 22 (13): 1600–1607. https://doi.org/10.1093/bioinformatics/btl140.

Armstrong, Joel, Glenn Hickey, Mark Diekhans, et al. 2020. "Progressive Cactus Is a Multiple-Genome Aligner for the Thousand-Genome Era." *Nature* 587 (7833): 246–51. https://doi.org/10.1038/s41586-020-2871-y.

Avsec, Žiga, Vikram Agarwal, Daniel Visentin, et al. 2021. "Effective Gene Expression Prediction from Sequence by Integrating Long-Range Interactions." *Nature Methods* 18 (10): 10. https://doi.org/10.1038/s41592-021-01252-x.

Bailey, T. L., M. Boden, F. A. Buske, et al. 2009. "MEME SUITE: Tools for Motif Discovery and Searching." *Nucleic Acids Research* 37 (Web Server): W202–8. https://doi.org/10.1093/nar/gkp335.

Bailey, Timothy L., James Johnson, Charles E. Grant, and William S. Noble. 2015. "The MEME Suite." *Nucleic Acids Research* 43 (W1): W39–49. https://doi.org/10.1093/nar/gkv416.

Ballester, Benoit, Alejandra Medina-Rivera, Dominic Schmidt, et al. 2014. "Multi-Species, Multi-Transcription Factor Binding Highlights Conserved Control of Tissue-Specific Biological Pathways." *Elife* 3: e02626.

Bastian, Frederic B, Julien Roux, Anne Niknejad, et al. 2021. "The Bgee Suite: Integrated Curated Expression Atlas and Comparative Transcriptomics in Animals." *Nucleic Acids Research* 49 (D1): D831–47. https://doi.org/10.1093/nar/gkaa793.

Bastian, Frederic B, Alessandro Brandulas Cammarata, Sara Carsanaro, et al. 2025. "Bgee in 2024: Focus on Curated Single-Cell RNA-Seq Datasets, and Query Tools." *Nucleic Acids Research* 53 (D1): D878–85. https://doi.org/10.1093/nar/gkae1118.

Benegas, Gonzalo, Sanjit Singh Batra, and Yun S. Song. 2023. "DNA Language Models Are Powerful Predictors of Genome-Wide Variant Effects." *Proceedings of the National Academy of Sciences* 120 (44): e2311219120. https://doi.org/10.1073/pnas.2311219120.

Boer, Carl G. de, and Jussi Taipale. 2024. "Hold out the Genome: A Roadmap to Solving the Cis-Regulatory Code." *Nature* 625 (7993): 41–50. https://doi.org/10.1038/s41586-023-06661-w.

Buchon, Nicolas, Neal Silverman, and Sara Cherry. 2014. "Immunity in Drosophila Melanogaster — from Microbial Recognition to Whole-Organism Physiology." *Nature Reviews Immunology* 14 (12): 796–810. https://doi.org/10.1038/nri3763.

Capella-Gutiérrez, Salvador, José M. Silla-Martínez, and Toni Gabaldón. 2009. "trimAl: A Tool for Automated Alignment Trimming in Large-Scale Phylogenetic Analyses." *Bioinformatics* 25 (15): 1972–73. https://doi.org/10.1093/bioinformatics/btp348.

Castro-Mondragon, Jaime A, Rafael Riudavets-Puig, Ieva Rauluseviciute, et al. 2022. "JASPAR 2022: The 9th Release of the Open-Access Database of Transcription

Factor Binding Profiles." *Nucleic Acids Research* 50 (D1): D165–73. https://doi.org/10.1093/nar/gkab1113.

Davydov, Iakov I, Nicolas Salamin, and Marc Robinson-Rechavi. 2019. "Large-Scale Comparative Analysis of Codon Models Accounting for Protein and Nucleotide Selection." *Molecular Biology and Evolution* 36 (6): 1316–32. https://doi.org/10.1093/molbev/msz048.

Dukler, Noah, Yi-Fei Huang, and Adam Siepel. 2020. "Phylogenetic Modeling of Regulatory Element Turnover Based on Epigenomic Data." *Molecular Biology and Evolution* 37 (7): 2137–52. https://doi.org/10.1093/molbev/msaa073.

Ewels, Philip, Alexander Peltzer, Sven Fillinger, et al. 2022. *The Nf-Core Framework for Community-Curated Bioinformatics Pipelines.* V. 2.4.1. Zenodo, released May 16. https://doi.org/10.5281/zenodo.7139814.

Eyre-Walker, Adam. 2006. "The Genomic Rate of Adaptive Evolution." *Trends in Ecology & Evolution* 21 (10): 569–75. https://doi.org/10.1016/j.tree.2006.06.015.

Eyre-Walker, Adam, and Peter D. Keightley. 2007. "The Distribution of Fitness Effects of New Mutations." *Nature Reviews Genetics* 8 (8): 610–18. https://doi.org/10.1038/nrg2146.

Fu, Wenqing, and Joshua M. Akey. 2013. "Selection and Adaptation in the Human Genome." *Annual Review of Genomics and Human Genetics* 14 (Volume 14, 2013): 467–89. https://doi.org/10.1146/annurev-genom-091212-153509.

Gallego Romero, Irene, and Amanda J. Lea. 2023. "Leveraging Massively Parallel Reporter Assays for Evolutionary Questions." *Genome Biology* 24 (1): 26. https://doi.org/10.1186/s13059-023-02856-6.

Galtier, Nicolas, and Laurent Duret. 2007. "Adaptation or Biased Gene Conversion? Extending the Null Hypothesis of Molecular Evolution." *Trends in Genetics* 23 (6): 273–77. https://doi.org/10.1016/j.tig.2007.03.011.

Ghandi, Mahmoud, Dongwon Lee, Morteza Mohammad-Noori, and Michael A. Beer. 2014. "Enhanced Regulatory Sequence Prediction Using Gapped K-Mer Features." *PLOS Computational Biology* 10 (7): e1003711. https://doi.org/10.1371/journal.pcbi.1003711.

Ghandi, Mahmoud, Morteza Mohammad-Noori, Narges Ghareghani, Dongwon Lee, Levi Garraway, and Michael A. Beer. 2016. "gkmSVM: An R Package for Gapped-Kmer SVM." *Bioinformatics* 32 (14): 2205–7. https://doi.org/10.1093/bioinformatics/btw203.

Guindon, Stéphane, Jean-François Dufayard, Vincent Lefort, Maria Anisimova, Wim Hordijk, and Olivier Gascuel. 2010. "New Algorithms and Methods to Estimate Maximum-Likelihood Phylogenies: Assessing the Performance of PhyML 3.0." *Systematic Biology* 59 (3): 307–21. https://doi.org/10.1093/sysbio/syq010.

Gupta, Shobhit, John A. Stamatoyannopoulos, Timothy L. Bailey, and William Stafford Noble. 2007. "Quantifying Similarity between Motifs." *Genome Biology* 8 (2): R24. https://doi.org/10.1186/gb-2007-8-2-r24.

Halpern, A L, and W J Bruno. 1998. "Evolutionary Distances for Protein-Coding Sequences: Modeling Site-Specific Residue Frequencies." *Molecular Biology and Evolution* 15 (7): 910–17. https://doi.org/10.1093/oxfordjournals.molbev.a025995.

He, Bin Z., Alisha K. Holloway, Sebastian J. Maerkl, and Martin Kreitman. 2011. “Does Positive Selection Drive Transcription Factor Binding Site Turnover? A Test with Drosophila Cis-Regulatory Modules.” *PLOS Genetics* 7 (4): e1002053. https://doi.org/10.1371/journal.pgen.1002053.

Hickey, Glenn, Benedict Paten, Dent Earl, Daniel Zerbino, and David Haussler. 2013. “HAL: A Hierarchical Format for Storing and Analyzing Multiple Genome Alignments.” *Bioinformatics* 29 (10): 1341–42. https://doi.org/10.1093/bioinformatics/btt128.

Hinrichs, A. S. 2006. “The UCSC Genome Browser Database: Update 2006.” *Nucleic Acids Research* 34 (90001): D590–98. https://doi.org/10.1093/nar/gkj144.

Horisawa, Kenichi, Miyako Udono, Kazuko Ueno, et al. 2020. “The Dynamics of Transcriptional Activation by Hepatic Reprogramming Factors.” *Molecular Cell* 79 (4): 660-676.e8. https://doi.org/10.1016/j.molcel.2020.07.012.

Hu, Zhenhua, and Wee-Wei Tee. 2017. “Enhancers and Chromatin Structures: Regulatory Hubs in Gene Expression and Diseases.” *Bioscience Reports* 37 (2): BSR20160183. https://doi.org/10.1042/BSR20160183.

Jiang, Daohan, and Jianzhi Zhang. 2024. “Ascertainment Bias in the Genomic Test of Positive Selection on Regulatory Sequences.” *Molecular Biology and Evolution* 41 (1): msad284. https://doi.org/10.1093/molbev/msad284.

Jordan, Gregory, and Nick Goldman. 2012. “The Effects of Alignment Error and Alignment Filtering on the Sitewise Detection of Positive Selection.” *Molecular Biology and Evolution* 29 (4): 1125–39. https://doi.org/10.1093/molbev/msr272.

Karollus, Alexander, Johannes Hingerl, Dennis Gankin, Martin Grosshauser, Kristian Klemon, and Julien Gagneur. 2024. “Species-Aware DNA Language Models Capture Regulatory Elements and Their Evolution.” *Genome Biology* 25 (1): 83. https://doi.org/10.1186/s13059-024-03221-x.

Kim, Seungsoo, and Joanna Wysocka. 2023. “Deciphering the Multi-Scale, Quantitative *Cis*-Regulatory Code.” *Molecular Cell*, Reimagining the Central Dogma, vol. 83 (3): 373–92. https://doi.org/10.1016/j.molcel.2022.12.032.

King, Mary-Claire, and Allan C. Wilson. 1975. “Evolution at Two Levels in Humans and Chimpanzees.” *Science* 188 (4184): 107–16.

Komljenovic, Andrea, Julien Roux, Julien Wollbrett, Marc Robinson-Rechavi, and Frederic B. Bastian. 2018a. “BgeeDB, an R Package for Retrieval of Curated Expression Datasets and for Gene List Expression Localization Enrichment Tests.” *F1000Research* 5 (August): 2748. https://doi.org/10.12688/f1000research.9973.2.

Komljenovic, Andrea, Julien Roux, Julien Wollbrett, Marc Robinson-Rechavi, and Frederic B. Bastian. 2018b. “BgeeDB, an R Package for Retrieval of Curated Expression Datasets and for Gene List Expression Localization Enrichment Tests.” Preprint, F1000Research, August 7. https://doi.org/10.12688/f1000research.9973.2.

Krieger, Gat, Offir Lupo, Patricia Wittkopp, and Naama Barkai. 2022. “Evolution of Transcription Factor Binding through Sequence Variations and Turnover of Binding Sites.” *Genome Research* 32 (6): 1099–111. https://doi.org/10.1101/gr.276715.122.

Kudron, Michelle, Louis Gevirtzman, Alec Victorsen, et al. 2024. “Binding Profiles for 961 Drosophila and C. Elegans Transcription Factors Reveal Tissue-Specific Regulatory Relationships.” *Genome Research* 34 (12): 2319–34. https://doi.org/10.1101/gr.279037.124.

Kurafeiski, Jasmin D, Paulo Pinto, and Erich Bornberg-Bauer. 2018. “Evolutionary Potential of Cis-Regulatory Mutations to Cause Rapid Changes in Transcription Factor Binding.” *Genome Biology and Evolution* 11 (2): 406–14. https://doi.org/10.1093/gbe/evy269.

Lai, Xuelei, Arnaud Stigliani, Gilles Vachon, et al. 2019. “Building Transcription Factor Binding Site Models to Understand Gene Regulation in Plants.” *Molecular Plant* 12 (6): 743–63. https://doi.org/10.1016/j.molp.2018.10.010.

Langmead, Ben, and Steven L Salzberg. 2012. “Fast Gapped-Read Alignment with Bowtie 2.” *Nature Methods* 9 (4): 357–59. https://doi.org/10.1038/nmeth.1923.

Lee, Dongwon. 2016. “LS-GKM: A New Gkm-SVM for Large-Scale Datasets.” *Bioinformatics* 32 (14): 2196–98. https://doi.org/10.1093/bioinformatics/btw142.

Lee, Phil H., Christian Lee, Xihao Li, Brian Wee, Tushar Dwivedi, and Mark Daly. 2018. “Principles and Methods of In-Silico Prioritization of Non-Coding Regulatory Variants.” *Human Genetics* 137 (1): 15–30. https://doi.org/10.1007/s00439-017-1861-0.

Lin, Meixi, Sneha Chakraborty, Carlos Eduardo G. Amorim, et al. 2025. “The Distribution of Fitness Effects Varies Phylogenetically across Animals.” Preprint, bioRxiv, May 13. https://doi.org/10.1101/2025.05.13.653358.

Liu, Jialin, and Marc Robinson-Rechavi. 2020. “Robust Inference of Positive Selection on Regulatory Sequences in the Human Brain.” *Science Advances* 6 (48): eabc9863. https://doi.org/10.1126/sciadv.abc9863.

Mackay, Trudy F. C., Stephen Richards, Eric A. Stone, et al. 2012. “The Drosophila Melanogaster Genetic Reference Panel.” *Nature* 482 (7384): 173–78. https://doi.org/10.1038/nature10811.

Maurano, Matthew T, Richard Humbert, Eric Rynes, et al. 2012. *Systematic Localization of Common Disease-Associated Variation in Regulatory DNA*. 337: 7.

McDonald, J. H., and M. Kreitman. 1991. “Adaptive Protein Evolution at the Adh Locus in Drosophila.” *Nature* 351 (6328): 652–54. https://doi.org/10.1038/351652a0.

Moretti, Sébastien, Balazs Laurenczy, Walid H. Gharib, et al. 2014. “Selectome Update: Quality Control and Computational Improvements to a Database of Positive Selection.” *Nucleic Acids Research* 42 (D1): D917–21. https://doi.org/10.1093/nar/gkt1065.

Moses, Alan M., Daniel A. Pollard, David A. Nix, et al. 2006. “Large-Scale Turnover of Functional Transcription Factor Binding Sites in Drosophila.” *PLOS Computational Biology* 2 (10): e130. https://doi.org/10.1371/journal.pcbi.0020130.

Nassar, Luis R., Galt P. Barber, Anna Benet-Pagès, et al. 2023. “The UCSC Genome Browser Database: 2023 Update.” *Nucleic Acids Research* 51 (D1): D1188–95. https://doi.org/10.1093/nar/gkac1072.

Pagès, Hervé. 2024. “BSgenome: Software Infrastructure for Efficient Representation of Full Genomes and Their SNPs. R Package Version Version 1.72.0, Https://Bioconductor.Org/Packages/BSgenome.” https://rdrr.io/bioc/BSgenome/.

Patwardhan, Rupali P., Joseph B. Hiatt, Daniela M. Witten, et al. 2012. “Massively Parallel Functional Dissection of Mammalian Enhancers in Vivo.” *Nature Biotechnology* 30 (3): 265–70. https://doi.org/10.1038/nbt.2136.

Peng, Junhui, and Li Zhao. 2024. “The Origin and Structural Evolution of de Novo Genes in Drosophila.” *Nature Communications* 15 (1): 810. https://doi.org/10.1038/s41467-024-45028-1.

Perez, Gerardo, Galt P Barber, Anna Benet-Pages, et al. 2025. “The UCSC Genome Browser Database: 2025 Update.” *Nucleic Acids Research* 53 (D1): D1243–49. https://doi.org/10.1093/nar/gkae974.

Phan, Mai H. Q., Tobias Zehnder, Fiona Puntieri, et al. 2025. “Conservation of Regulatory Elements with Highly Diverged Sequences across Large Evolutionary Distances.” *Nature Genetics* 57 (6): 1524–34. https://doi.org/10.1038/s41588-025-02202-5.

Pollard, K. S., M. J. Hubisz, K. R. Rosenbloom, and A. Siepel. 2010. “Detection of Nonneutral Substitution Rates on Mammalian Phylogenies.” *Genome Research* 20 (1): 110–21. https://doi.org/10.1101/gr.097857.109.

Ramírez, Fidel, Friederike Dündar, Sarah Diehl, Björn A. Grüning, and Thomas Manke. 2014. “deepTools: A Flexible Platform for Exploring Deep-Sequencing Data.” *Nucleic Acids Research* 42 (Web Server issue): W187–91. https://doi.org/10.1093/nar/gku365.

Rensch, Thomas, Diego Villar, Julie Horvath, Duncan T. Odom, and Paul Flicek. 2016. “Mitochondrial Heteroplasmy in Vertebrates Using ChIP-Sequencing Data.” *Genome Biology* 17 (1): 139. https://doi.org/10.1186/s13059-016-0996-y.

Rodrigue, Nicolas, Thibault Latrille, and Nicolas Lartillot. 2021. “A Bayesian Mutation–Selection Framework for Detecting Site-Specific Adaptive Evolution in Protein-Coding Genes.” *Molecular Biology and Evolution* 38 (3): 1199–208. https://doi.org/10.1093/molbev/msaa265.

Schmidt, D., M. D. Wilson, B. Ballester, et al. 2010. “Five-Vertebrate ChIP-Seq Reveals the Evolutionary Dynamics of Transcription Factor Binding.” *Science* 328 (5981): 1036–40. https://doi.org/10.1126/science.1186176.

Schmidt, Dominic, Petra C. Schwalie, Michael D. Wilson, et al. 2012. “Waves of Retrotransposon Expansion Remodel Genome Organization and CTCF Binding in Multiple Mammalian Lineages.” *Cell* 148 (1–2): 335–48. https://doi.org/10.1016/j.cell.2011.11.058.

Sella, Guy, Dmitri A. Petrov, Molly Przeworski, and Peter Andolfatto. 2009. “Pervasive Natural Selection in the Drosophila Genome?” *PLOS Genetics* 5 (6): e1000495. https://doi.org/10.1371/journal.pgen.1000495.

Siepel, A, G Bejerano, J.S Pedersen, et al. 2005. “Evolutionarily Conserved Elements in Vertebrate, Insect, Worm, and Yeast Genomes.” *Genome Research* 15 (8): 1034–50. https://doi.org/10.1101/gr.3715005.

Siepel, Adam, Gill Bejerano, Jakob S. Pedersen, et al. 2005. “Evolutionarily Conserved Elements in Vertebrate, Insect, Worm, and Yeast Genomes.” *Genome Research* 15 (8): 1034–50. https://doi.org/10.1101/gr.3715005.

Sievers, Fabian, Andreas Wilm, David Dineen, et al. 2011. “Fast, Scalable Generation of High-Quality Protein Multiple Sequence Alignments Using Clustal Omega.” *Molecular Systems Biology* 7 (October): 539. https://doi.org/10.1038/msb.2011.75.

Smith, Justin D., Kimberly F. McManus, and Hunter B. Fraser. 2013. “A Novel Test for Selection on Cis-Regulatory Elements Reveals Positive and Negative Selection Acting on Mammalian Transcriptional Enhancers.” *Molecular Biology and Evolution* 30 (11): 2509–18. https://doi.org/10.1093/molbev/mst134.

Sokolova, Ksenia, Kathleen M. Chen, Yun Hao, Jian Zhou, and Olga G. Troyanskaya. 2024. “Deep Learning Sequence Models for Transcriptional Regulation.” *Annual Review of Genomics and Human Genetics* 25 (Volume 25, 2024): 105–22. https://doi.org/10.1146/annurev-genom-021623-024727.

Soni, Vivak, Parul Johri, and Jeffrey D Jensen. 2023. “Evaluating Power to Detect Recurrent Selective Sweeps under Increasingly Realistic Evolutionary Null Models.” *Evolution* 77 (10): 2113–27. https://doi.org/10.1093/evolut/qpad120.

Stefflova, Klara, David Thybert, Michael D. Wilson, et al. 2013. “Cooperativity and Rapid Evolution of Cobound Transcription Factors in Closely Related Mammals.” *Cell* 154 (3): 530–40. https://doi.org/10.1016/j.cell.2013.07.007.

Tanaka, T, and M Nei. 1989. “Positive Darwinian Selection Observed at the Variable-Region Genes of Immunoglobulins.” *Molecular Biology and Evolution* 6 (5): 447–59. https://doi.org/10.1093/oxfordjournals.molbev.a040569.

Tognon, Manuel, Alisa Kumbara, Andrea Betti, Lorenzo Ruggeri, and Rosalba Giugno. 2025. “Benchmarking Transcription Factor Binding Site Prediction Models: A Comparative Analysis on Synthetic and Biological Data.” *Briefings in Bioinformatics* 26 (4): bbaf363. https://doi.org/10.1093/bib/bbaf363.

Vaibhvi, Vaibhvi, Sven Künzel, and Thomas Roeder. 2022. “Hemocytes and Fat Body Cells, the Only Professional Immune Cell Types in Drosophila, Show Strikingly Different Responses to Systemic Infections.” *Frontiers in Immunology* 13 (November). https://doi.org/10.3389/fimmu.2022.1040510.

Vaishnav, Eeshit Dhaval, Carl G. de Boer, Jennifer Molinet, et al. 2022. “The Evolution, Evolvability and Engineering of Gene Regulatory DNA.” *Nature* 603 (7901): 7901. https://doi.org/10.1038/s41586-022-04506-6.

Venkat, Aarti, Matthew W. Hahn, and Joseph W. Thornton. 2018. “Multinucleotide Mutations Cause False Inferences of Lineage-Specific Positive Selection.” *Nature Ecology & Evolution* 2 (8): 1280–88. https://doi.org/10.1038/s41559-018-0584-5.

Villar, Diego, Camille Berthelot, Sarah Aldridge, et al. 2015. “Enhancer Evolution across 20 Mammalian Species.” *Cell* 160 (3): 554–66. https://doi.org/10.1016/j.cell.2015.01.006.

Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, et al. 2020. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python.” *Nature Methods* 17 (3): 261–72. https://doi.org/10.1038/s41592-019-0686-2.

Vorontsov, Ilya E., Ivan Kozin, Sergey Abramov, et al. 2025. “Cross-Platform Motif Discovery and Benchmarking to Explore Binding Specificities of Poorly Studied Human Transcription Factors.” *Communications Biology* 8 (1): 1545. https://doi.org/10.1038/s42003-025-08909-9.

Ward, Lucas D., and Manolis Kellis. 2012. “Interpreting Noncoding Genetic Variation in Complex Traits and Human Disease.” *Nature Biotechnology* 30 (11): 1095–106. https://doi.org/10.1038/nbt.2422.

Wisotsky, Sadie R, Sergei L Kosakovsky Pond, Stephen D Shank, and Spencer V Muse. 2020. “Synonymous Site-to-Site Substitution Rate Variation Dramatically Inflates False Positive Rates of Selection Analyses: Ignore at Your Own Peril.” *Molecular Biology and Evolution* 37 (8): 2430–39. https://doi.org/10.1093/molbev/msaa037.

Wittkopp, Patricia J., and Gizem Kalay. 2012. “Cis-Regulatory Elements: Molecular Mechanisms and Evolutionary Processes Underlying Divergence.” *Nature Reviews Genetics* 13 (1): 59–69. https://doi.org/10.1038/nrg3095.

Wong, Emily S., Dawei Zheng, Siew Z. Tan, et al. 2020. “Deep Conservation of the Enhancer Regulatory Code in Animals.” *Science* 370 (6517): eaax8137. https://doi.org/10.1126/science.aax8137.

Wu, Tianzhi, Erqiang Hu, Shuangbin Xu, et al. 2021. “clusterProfiler 4.0: A Universal Enrichment Tool for Interpreting Omics Data.” *Innovation (Cambridge (Mass.))* 2 (3): 100141. https://doi.org/10.1016/j.xinn.2021.100141.

Yan, Han, Zhirui Hu, Gregg W C Thomas, Scott V Edwards, Timothy B Sackton, and Jun S Liu. 2023. “PhyloAcc-GT: A Bayesian Method for Inferring Patterns of Substitution Rate Shifts on Targeted Lineages Accounting for Gene Tree Discordance.” *Molecular Biology and Evolution* 40 (9): msad195. https://doi.org/10.1093/molbev/msad195.

Yan, Jian, Yunjiang Qiu, André M. Ribeiro dos Santos, et al. 2021. “Systematic Analysis of Binding of Transcription Factors to Noncoding Variants.” *Nature* 591 (7848): 147–51. https://doi.org/10.1038/s41586-021-03211-0.

Yang, Z. 1998. “Likelihood Ratio Tests for Detecting Positive Selection and Application to Primate Lysozyme Evolution.” *Molecular Biology and Evolution* 15 (5): 568–73. https://doi.org/10.1093/oxfordjournals.molbev.a025957.

Yang, Ziheng. 2014. *Molecular Evolution: A Statistical Approach*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199602605.001.0001.

Zhang, Feng, and James R. Lupski. 2015. “Non-Coding Genetic Variants in Human Disease.” *Human Molecular Genetics* 24 (R1): R102–10. https://doi.org/10.1093/hmg/ddv259.

Zhang, Yong, Tao Liu, Clifford A. Meyer, et al. 2008. “Model-Based Analysis of ChIP-Seq (MACS).” *Genome Biology* 9 (9): R137. https://doi.org/10.1186/gb-2008-9-9-r137.

Zhou, Jian, and Olga G. Troyanskaya. 2015. “Predicting Effects of Noncoding Variants with Deep Learning–Based Sequence Model.” *Nature Methods* 12 (10): 10. https://doi.org/10.1038/nmeth.3547.