

PhyloCSF: a comparative genomics method to distinguish protein-coding and non-coding regions

Michael F. Lin^{*†} Irwin Jungreis^{*} Manolis Kellis^{*†‡}

Abstract

As high-throughput transcriptome sequencing provides evidence for novel transcripts in many species, there is a renewed need for accurate methods to classify small genomic regions as protein-coding or non-coding. We present PhyloCSF, a novel comparative genomics method that analyzes a multi-species nucleotide sequence alignment to determine whether it is likely to represent a conserved protein-coding region, based on a formal statistical comparison of phylogenetic codon models. We show that PhyloCSF’s classification performance in 12-species *Drosophila* genome alignments exceeds all other methods we compared in a previous study, and we provide a software implementation for use by the community. We anticipate that this method will be widely applicable as the transcriptomes of many additional species, tissues, and subcellular compartments are sequenced, particularly in the context of ENCODE and modENCODE, and as interest shifts to long non-coding RNAs, defined primarily by their lack of protein-coding potential rather than conserved RNA structures.

Introduction

High-throughput transcriptome sequencing (mRNA-Seq) is yielding precise structures for novel transcripts in many species, including mammals [1]. Accurate computational methods are needed to classify these transcripts and the corresponding genomic exons as protein-coding or non-coding, even if the transcript models are incomplete or if they only reveal novel exons of already-known genes. In addition to classifying novel transcript models, such methods also have applications in evaluating and revising existing gene annotations [2, 3, 4, 5, 6], and as input features for *de novo* gene structure predictors [7, 8]. We have previously [9] compared numerous methods for determining whether an exon-length nucleotide sequence is likely to be protein-coding or non-coding, including single-sequence metrics that analyze the genome of interest only, and comparative genomics metrics that use alignments of orthologous regions in the genomes of related species.

Among the comparative methods benchmarked in our previous study, one of our original contributions was the Codon Substitution Frequencies (CSF) metric, which assigns a score to each codon substitution observed in the input alignment based on the relative frequency of that substitution in known coding and non-coding regions. We showed that CSF is highly effective, performing competitively with a phylogenetic modeling approach with much less computational expense, and indeed we have applied it successfully in flies [10, 3], fungi [5], and mammals [4, 11, 1]. However,

^{*}Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology. 32 Vassar St. 32-D510, Cambridge, MA 02139

[†]The Broad Institute. 7 Cambridge Center, Cambridge, MA 02142

[‡]Correspondence: {mikelin,manoli}@mit.edu

as discussed in our previous work, CSF has certain drawbacks arising from its *ad hoc* scheme for combining evidence from multiple species. For example, it makes only partial use of the evidence available in a multi-species alignment, and it produces a score lacking a precise theoretical interpretation, meaningful only relative to its empirical distributions in known coding and non-coding regions.

This paper introduces a rigorous reformulation of CSF, which frames the evaluation of a given alignment as a statistical model comparison problem, choosing between phylogenetic models estimated from known coding and non-coding regions as the best explanation for the alignment. This new “PhyloCSF” method fully leverages multiple alignments in a phylogenetic framework, produces meaningful likelihood ratios as its output, and rests upon the sweeping theoretical foundation for statistical model comparison. Benchmarking on the classification datasets from our original study, we show that PhyloCSF outperforms all of the other methods we had previously considered. Lastly, we briefly describe our software implementation, which we make publicly available, as well as several recent applications in fungi, flies, and mammals.

Phylogenetic tests for distinguishing coding and non-coding regions

PhyloCSF is based on the well-established theoretical framework for statistical phylogenetic model comparison. In this context, phylogenetic models are generative probabilistic models that produce alignments of molecular sequences, based on a prior distribution over a common ancestral sequence, the topology and branch lengths of a phylogenetic tree relating the descendants, and a substitution process along each branch giving the rates (per unit branch length) at which each character changes to any other. In phylogenetic model comparison, we wish to choose between two competing models as the better explanation for a given alignment. A standard approach is to decide based on the *likelihood ratio* between the two models, which quantifies how much more probable the alignment is under one model than under the other. This general approach has been used to explore many different aspects of the evolution of protein-coding genes, as recently reviewed in [12] and [13].

To distinguish coding and non-coding regions, we design one phylogenetic model to represent the evolution of codons in protein-coding genes, and another to represent the evolution of nucleotide triplet sites in non-coding regions. These models may have one or more parameters θ that adjust them to the genomic region of interest, e.g. the neutral substitution rate or G+C content. To analyze a given alignment A of extant sequences, we first determine the probability of the alignment under the maximum likelihood estimate (MLE) of the parameters for the coding model, $p_C = \max_{\theta_C} \Pr(A | \text{Coding}, \theta_C)$. We similarly estimate the alignment’s probability under the non-coding model, $p_N = \max_{\theta_N} \Pr(A | \text{Noncoding}, \theta_N)$. Finally, we decide if the alignment is more likely to represent a protein-coding region or a non-coding region based on the log-likelihood ratio $\Lambda = \log \frac{p_C}{p_N}$. The cutoff can be chosen to achieve a certain level of statistical significance, based on known asymptotic convergence properties of the log-likelihood ratio statistic [14, 15, 16], or it can be chosen empirically based on classification performance in a test set; we use the latter strategy in this work.

The d_N/d_S test

A standard method for detecting purifying selection on protein-coding sequences is to test for evidence that non-synonymous substitutions occur at a slower rate than synonymous substitutions. In the widely used PAML implementation of this test [17, 18], the codon frequencies π and the ratio of transition to transversion rates κ are the only parameters used to determine all triplet substitution rates in the background/non-coding model, while the coding model additionally supposes that non-

synonymous codon substitution rates are reduced relative to synonymous rates by a scale factor ω (also called d_N/d_S). PAML takes the phylogenetic tree topology as input, and estimates the branch lengths, π , κ , and ω for each alignment. The log-likelihood ratio between the coding and non-coding models can then be obtained from PAML’s output. (By convention, the log-likelihood ratio is set to zero if the estimated $\omega \geq 1$.)

Our previous work [9] showed this to be one of the best comparative methods for distinguishing coding and non-coding regions, outperforming our CSF metric according to standard classification error measures. Notably however, the d_N/d_S test performed worse than CSF for short regions (≤ 180 nt). This is not surprising since PAML was designed for evolutionary analysis of complete open reading frames, not short exon-length regions, which probably provide too little information to reliably estimate both the branch lengths and codon frequencies in addition to the two rate parameters.

PhyloCSF

PhyloCSF differs from the standard d_N/d_S test in two main ways. First, it takes advantage of recent advances in phylogenetic codon models that enable much more detailed representations of coding and non-coding sequence evolution. Specifically, while the d_N/d_S test uses only a few parameters to model the rates of all possible codon substitutions [17], PhyloCSF uses *empirical codon models* (ECMs) based on several thousand parameters modeling these rates [19] – one ECM estimated from alignments of many known coding regions, and another ECM from non-coding regions. By comparing these two rich evolutionary models, PhyloCSF can observe many additional informative features of a given alignment compared to the d_N/d_S test. For example, the coding ECM captures not only the decreased overall rate of non-synonymous substitutions, but also the different rates of specific non-synonymous substitutions reflecting the chemical properties of the amino acids. (Earlier codon modeling approaches also incorporate amino acid distances, e.g. [20], but to our knowledge, these are not widely used for distinguishing coding and non-coding regions.) Also, our ECMs explicitly model the very low rates of nonsense substitutions (giving rise to stop codons) in coding regions compared to non-coding regions.

Second, PhyloCSF also takes advantage of genome-wide training data to provide prior information about the branch lengths in the phylogenetic tree and the codon frequencies, rather than attempting to re-estimate these *a priori* in each individual alignment. PhyloCSF assumes a fixed tree “shape” based on the genome-wide MLEs of the branch lengths, and estimates only two scale factors ρ_C and ρ_N , applied uniformly to all of the branch lengths in the coding and non-coding models respectively, for each individual region analyzed. This allows PhyloCSF to accommodate some region-specific rate variation and reduces its sensitivity to the absolute degree of conservation – without needing to estimate a large number of parameters for each alignment, which may be difficult for short regions.

In summary (Figure 1), PhyloCSF relies on two ECMs fit to genome-wide training data, which include estimates for the branch lengths, codon frequencies, and codon substitution rates for known coding and non-coding regions. To evaluate a given nucleotide sequence alignment, PhyloCSF (1) determines the MLE of the scale factor ρ on the branch lengths for each of these models, (2) computes the likelihood of each model (the probability of the alignment under the model) using the MLE of the scale factor, and (3) reports the log-likelihood ratio between the coding and non-coding models.




Figure 1: PhyloCSF method overview. (A) PhyloCSF uses phylogenetic codon models estimated from genome-wide training data based on known coding and non-coding regions. These models include a phylogenetic tree and codon substitution rate matrices Q_C and Q_N for coding and non-coding regions, respectively, shown here for 12 *Drosophila* species. Q_C captures the characteristic evolutionary signatures of codon substitutions in conserved coding regions, while Q_N captures the typical evolutionary rates of triplet sites in non-coding regions. (B) PhyloCSF applied to a short region from the first exon of the *D. melanogaster* homeobox gene *Dfd*. The alignment of this region shows only synonymous substitutions compared to the inferred ancestral sequence (green). Using the maximum likelihood estimate of a scale factor ρ applied to the assumed branch lengths, the alignment has higher probability under the coding model than the non-coding model, resulting in a positive log-likelihood ratio Λ . (C) PhyloCSF applied to a conserved region within a *Dfd* intron. In contrast to the exonic alignment, this region shows many non-synonymous substitutions (red), nonsense substitutions (blue, purple), and frameshifts (orange). The alignment has lower probability under the coding model, resulting in a negative score.

PhyloCSF outperforms other methods

We used the datasets and benchmarks from our previous study [9] to evaluate our new method. Briefly, the datasets consist of known protein-coding regions and randomly selected non-coding regions (about 50,000 total regions) in the genome of the fruitfly *Drosophila melanogaster*, aligned with eleven other *Drosophila* species using MULTIZ [21, 22, 10]. The lengths of the regions in both the coding and non-coding sets match the length distribution of fly coding exons. Consistent with our previous work, we trained and applied PhyloCSF on this dataset using four-fold cross-validation, to ensure that any observed performance differences are not due to overfitting. We assessed the results by examining ROC curves and computing the minimum average error (MAE), the average false positive and false negative rates at the cutoff that minimizes this average. To compare the power of the methods for short exons specifically, we additionally computed these benchmarks only for the 37% of examples from 30 to 180 nucleotides in length (a range which includes three-quarters of mammalian coding exons).

These benchmarks show that PhyloCSF outperforms the other comparative methods we previously benchmarked (Figure 2, left column), effectively dominating them all at good sensitivity/specificity tradeoffs. PhyloCSF’s overall MAE was 19% lower than that of the Reading Frame Conservation metric, 15% lower than that of our older CSF method, and 8% lower than the d_N/d_S test’s. PhyloCSF also clearly outperformed the other methods for short exons (Figure 2B), with an MAE 11% lower than the d_N/d_S test’s. Our previous study [9] showed that these comparative methods in turn outperform single-species metrics based on sequence composition (e.g. interpolated Markov models and the Z curve discriminant).

We also compared PhyloCSF to the other methods using only pairwise alignments between *D. melanogaster* and *D. ananassae*, which we previously showed to be the best single informant for this purpose. PhyloCSF also dominates other methods in these pairwise alignments (Figure 2, right column), with MAE at least 24% lower for all aligned regions and at least 21% lower for the short regions. Some of PhyloCSF’s greater relative advantage in the pairwise case arises from its ability to produce an informative score even for regions lacking alignment with *D. ananassae*, based on the composition of the *D. melanogaster* region and the codon frequencies included in PhyloCSF’s generative ECMs.

Overall, these benchmarks show that PhyloCSF provides superior power to distinguish fly coding and non-coding regions based on either multi-species or pairwise genome alignments.

Using non-independence of codon sites to increase accuracy

Like most codon modeling approaches, PhyloCSF assumes statistical independence of each codon site from all the others in the alignment (conditional on the model parameters). We next studied the extent to which this assumption is violated in the real alignments in our dataset, and sought to use this to further improve our ability to distinguish the coding and non-coding regions.

One simple way of assessing the dependence of codon sites in a set of alignments is to measure the relationship between the PhyloCSF log-likelihood ratio, Λ , and the number of sites in each alignment, n . It can be shown that, for alignments consisting of n independent sites all generated from either the coding or non-coding ECM, the standard deviation $SD(\Lambda) \propto n^{1/2}$ [16]. On the other hand, $SD(\Lambda) \propto n$ under complete dependence of the sites in each alignment – that is, where each alignment consists of a single randomly sampled site, repeated n times. We applied a numerical procedure to estimate coefficients A and B in $SD(\Lambda) \sim An^B$ based on the real alignments in our dataset. For coding regions, we found $B_C = 0.91$ and for non-coding regions, we found $B_N = 0.75$. Compared to the expected values of $B = 0.5$ under complete independence and $B = 1$ under

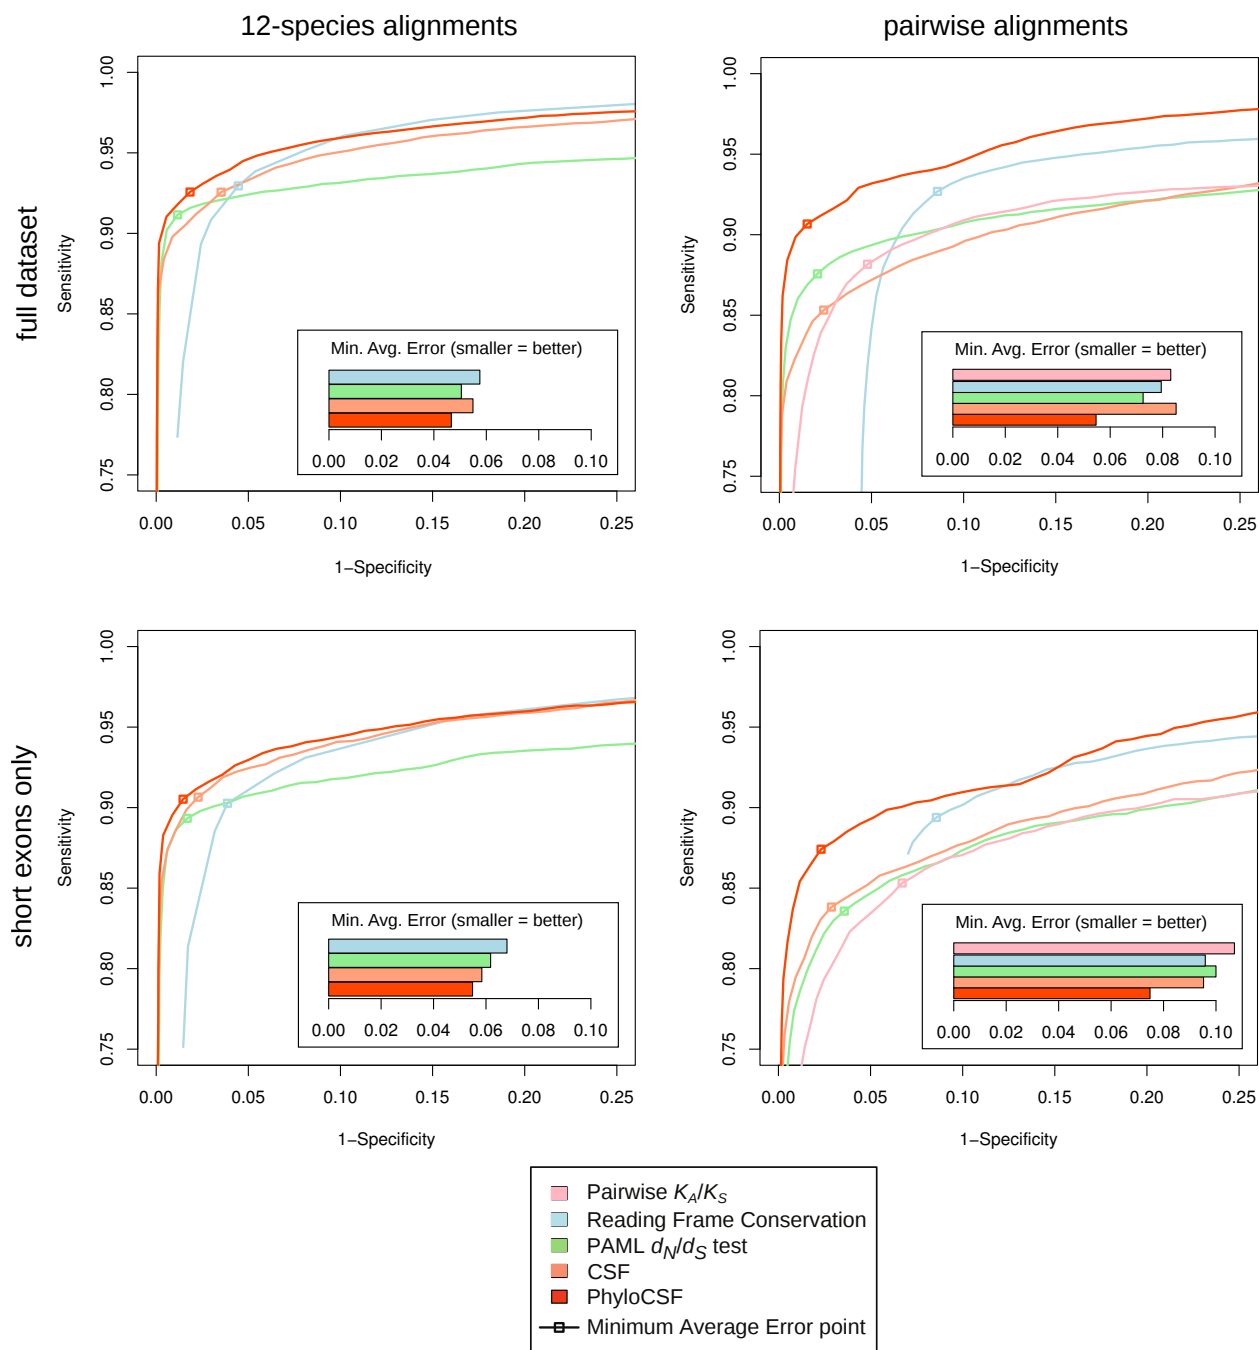


Figure 2: PhyloCSF performance benchmarks. ROC plots and error measures for several methods to distinguish known protein-coding and randomly selected non-coding regions in *D. melanogaster*. The top row of plots shows the results for our full dataset of approximately 50,000 regions matching the fly exon length distribution, while the bottom row of plots is based only the 37% of these regions between 30 and 180 nucleotides in length. The left-hand plots show the performance of the methods applied to multiple alignments of 12 fly genomes, while the right-hand plots use pairwise alignments between *D. melanogaster* and *D. ananassae*. PhyloCSF effectively dominates the other methods.




Figure 3: Exploiting non-independence of codon sites. The PhyloCSF log-likelihood ratio Λ is transformed based on alignment length into a new log-odds score Ψ (see main text). Ψ provides a superior discriminant in the full 12 fly dataset.

complete dependence, these results suggest strong site-to-site dependencies in the real alignments – considerably stronger in coding regions than non-coding regions.

The differing length-dependent dispersions of Λ for coding and non-coding regions suggest that a better classification rule could be obtained by interpreting the log-likelihood ratio for each example in the context of the alignment’s length. This contrasts with the suggestion of the “law of likelihood” that, under the model’s independence assumptions, any information pertinent to the decision would already be captured in the likelihood (Hacking 1965).

We defined a new log-odds score Ψ for each alignment by assuming that Λ normally distributes with mean and variance as different functions of n for coding and non-coding regions:

$$\Psi(\Lambda, n) = \log \frac{\mathcal{N}(\Lambda; n\mu_C, (A_C n^{B_C})^2)}{\mathcal{N}(\Lambda; n\mu_N, (A_N n^{B_N})^2)}$$

where $\mathcal{N}(x; \mu, \sigma^2)$ is the normal density and the six parameters μ_C , A_C , B_C , μ_N , A_N , and B_N are estimated numerically from the training dataset (see Methods).

We calculated Ψ for each of the examples in our *Drosophila* dataset (with cross-validation) and found that it indeed provides a superior overall discriminant between the coding and non-coding regions (Figure 3), with an MAE 17% lower than that of Λ and 23% lower than the d_N/d_S test’s. While Ψ somewhat distorts the formal inferential meaning captured by Λ , this approach exploits information from site-to-site dependencies in a principled way at negligible additional computational cost – in contrast to attempts to explicitly capture such dependencies in the generative codon models [12, 13].

Implementation and availability

To facilitate the use of PhyloCSF by the community, we provide an implementation that evaluates an input sequence alignment in Multi-FASTA format and reports the resulting log-likelihood ratio

in units of decibans. The program supports batch processing of many alignments, and can either evaluate the alignments as-is or search for high-scoring open reading frames within them. We also provide the ECMs and other parameter settings for several phylogenies. The Objective Caml source code and executables for several platforms are available at: <http://compbio.mit.edu/PhyloCSF>

Applying PhyloCSF in fungi, flies, and mammals

We now briefly discuss several recent applications of PhyloCSF in three different phylogenies. While biological results of each of these applications are presented elsewhere in greater detail, together they illustrate PhyloCSF’s usefulness in practice.

Recent efforts to reconstruct the transcriptome of the fission yeast *Schizosaccharomyces pombe* based on mRNA-Seq experiments suggested numerous novel transcript models. We generated whole-genome alignments of *S. pombe* with *S. octosporus*, *S. japonicus* and *S. cryophilus*, and identified 89 novel protein-coding genes that are highly conserved among these species, even though they do not show primary sequence similarity to known proteins. The mRNA-Seq reconstruction also suggested several hundred revisions to the exon-intron structures of existing gene annotations. PhyloCSF showed that the revised transcript models strongly tend to have higher coding potential than the originals, confirming the quality of the proposed revisions (Rhind *et al.*, “Comparative functional genomics reveals fission yeast innovation in centromere structure, transcriptional control, carbon metabolism and antisense transcription of meiotic genes”, submitted).

The modENCODE project undertook a comprehensive effort to map the transcriptome of *D. melanogaster*, identifying 1938 novel transcribed loci in this well-studied genome. Using PhyloCSF with the 12-species whole-genome alignments, we identified 138 of these likely to represent novel coding genes (The modENCODE Consortium, 2010). We also used PhyloCSF in concert with the modENCODE data to identify and characterize 238 known *D. melanogaster* genes that show high coding potential immediately downstream of their stop codon, suggesting a surprisingly widespread mechanism of stop codon readthrough (Jungreis *et al.*, “Evidence of abundant stop codon readthrough in *Drosophila* and other metazoa”, submitted). PhyloCSF’s ability to systematically resolve small regions was especially important in this application, as many of these putative readthrough regions are quite short (mean 67 codons).

Lastly, we have also used PhyloCSF in the human and mouse genomes, using alignments of 29 placental mammals (2X Mammals Sequencing and Analysis Consortium, in prep.). We applied PhyloCSF to transcript models reconstructed from mRNA-Seq on 16 human tissues to identify several candidate novel coding genes (Figure 4), despite the extensive decade-long efforts to annotate the human genome. We also used CSF and PhyloCSF to evaluate coding potential in novel mouse transcripts, helping to identify thousands of long *non*-coding RNAs with diverse functional roles [11] (Hung *et al.*, “Extensive and choreographed transcription of noncoding RNAs within cell cycle promoters”, submitted).

Discussion

We have introduced PhyloCSF, a comparative genomics method for distinguishing protein-coding and non-coding regions, and shown that it outperforms all previous methods. In addition to its superior discriminatory power, PhyloCSF is far more theoretically attractive than our older CSF and other *ad hoc* metrics, relying on a formal statistical comparison of phylogenetic codon models. However, we note that PhyloCSF and CSF produce highly correlated scores (Pearson coefficient 0.95 in our dataset), and the new method is much more computationally demanding. Also, unlike many

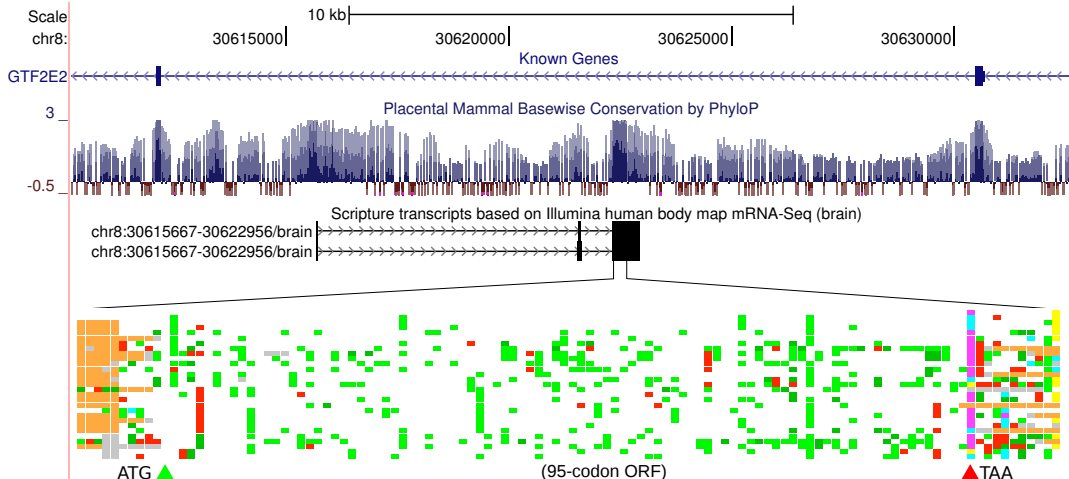


Figure 4: A novel human coding gene found with the help of PhyloCSF. Transcriptome reconstruction by Scripture [1] based on brain mRNA-Seq data provided by Illumina, Inc. produced two alternative transcript models lying antisense to an intron of *GTF2E2*, a known protein-coding gene. PhyloCSF identified a 95-codon ORF in the third exon of this transcript, highly conserved across placental mammals. The color schematic illustrates conservation (white), synonymous and conservative codon substitutions (green), other non-synonymous codon substitutions (red), stop codons (blue/magenta/yellow), and frame-shifted regions (orange) in the genome alignment of 29 mammals. Despite its unmistakable protein-coding evolutionary signatures, the ORF shows no sequence similarity to known proteins.

of the other methods compared in our previous study, CSF and PhyloCSF require extensive genome-wide training data from known coding and non-coding regions, which can present an obstacle to their use outside of well-studied phylogenies such as mammals and flies.

The classification approach described here, in which we are given individual sequence alignments and must decide whether each one represents a coding or non-coding region, is complementary to *de novo* comparative gene predictors that attempt to parse the complete primary genome sequence into exon-intron structures and intergenic regions. In particular, we suggest that the classification approach is more naturally applicable to transcript models reconstructed from expression evidence such as mRNA-Seq, where the exon-intron structure is essentially known. Moreover, the rapidly decreasing cost of such transcriptome sequencing is arguably reducing the need for *de novo* structure gene predictors, although they will remain useful for predicting possible lowly- or rarely-expressed genes.

Methods

PhyloCSF

PhyloCSF’s trained parameters include a phylogenetic tree (with branch lengths) and two 64-by-64 codon rate matrices \mathbf{Q}_C and \mathbf{Q}_N representing coding and non-coding sequence evolution, respectively, as reversible, homogeneous, continuous-time Markov processes. To evaluate a given alignment, we first evaluate the likelihood of the coding model as follows. We define an alignment-specific parameter ρ_C that operates as a scale factor applied to all of the branch lengths in the predefined tree. Given a setting of ρ_C , the substitution probability matrix along any branch with

length t is given by $\mathbf{P} = \exp(t\rho_C \mathbf{Q}_C)$. We can then compute the probability of the full alignment using Felsenstein’s algorithm [23] – assuming independence of the codon sites, using the equilibrium frequencies implicit in \mathbf{Q}_C as the prior distribution over the common ancestral sequence, and marginalizing out any gapped or ambiguous codons. We numerically maximize this probability over ρ_C to obtain the likelihood of the coding model p_C . We then evaluate the likelihood of the non-coding model p_N in the same way, using \mathbf{Q}_N and an independent scale factor ρ_N , and report the log-likelihood ratio $\log \frac{p_C}{p_N}$ as the result.

Estimation of empirical codon models

To estimate the phylogenetic tree and the empirical rate matrices \mathbf{Q}_C and \mathbf{Q}_N for the species of interest, we rely on sequence alignments of many known coding and random non-coding regions. Given this genome-wide training data, we optimize the parameters for the coding and non-coding models using an expectation-maximization approach. The E-step is carried out as previously described [24, 25]. In each M-step, we update the ECM exchangeability parameters using a spectral approximation method [26] and the branch lengths by numerical optimization of the expected log-likelihood function [25]. Meanwhile, the codon/triplet frequencies are fixed to their empirical averages in the training examples, and we assume a fixed species tree topology.

Parameter estimation for Ψ transformation

We estimated the parameters needed for the Ψ transformation by first defining a Z-score for each training example i ,

$$Z_i = \frac{\Lambda_i - n_i \mu}{A n_i^B}$$

where Λ_i is the PhyloCSF log-likelihood ratio for the i ’th example and n_i is the number of codon sites in this example. We then numerically solve these three equations for the three parameters μ , A , and B :

$$\begin{aligned} \sum_i Z_i &= 0 \\ \sum_i Z_i^2 &= 1 \\ \text{Cor}(\mathbf{Z}, \mathbf{\Lambda}) &= 0 \end{aligned}$$

leading to Z-scores which have mean 0 and variance 1, and do not correlate with length. We apply this procedure separately for all the coding and all the non-coding examples in the training dataset to obtain the six parameters μ_C , A_C , B_C , μ_N , A_N , and B_N .

Acknowledgments

The authors thank Matthew D. Rasmussen and Manuel Garber for helpful comments and discussions. Funding for this work was provided by the National Institutes of Health (U54 HG004555-01) and the National Science Foundation (DBI 0644282).

References

- [1] Guttman M, Garber M, Levin JZ, Donaghey J, Robinson J, et al. (2010) Ab initio reconstruction of cell type-specific transcriptomes in mouse reveals the conserved multi-exonic structure of lincrnas. *Nat Biotech* 28: 503-510.
- [2] Kellis M, Patterson N, Endrizzi M, Birren B, Lander ES (2003) Sequencing and comparison of yeast species to identify genes and regulatory elements. *Nature* 423: 241-254.
- [3] Lin MF, Carlson JW, Crosby MA, Matthews BB, Yu C, et al. (2007) Revisiting the protein-coding gene catalog of *Drosophila melanogaster* using 12 fly genomes. *Genome research* 17: 1823-1836.
- [4] Clamp M, Fry B, Kamal M, Xie X, Cuff J, et al. (2007) Distinguishing protein-coding and noncoding genes in the human genome. *Proceedings of the National Academy of Sciences* 104: 19428-19433.
- [5] Butler G, Rasmussen MD, Lin MF, Santos MAS, Sakthikumar S, et al. (2009) Evolution of pathogenicity and sexual reproduction in eight *Candida* genomes. *Nature* 459: 657-662.
- [6] Pruitt KD, Harrow J, Harte RA, Wallin C, Diekhans M, et al. (2009) The consensus coding sequence (CCDS) project: Identifying a common protein-coding gene set for the human and mouse genomes. *Genome research* 19: 1316-1323.
- [7] Brent MR (2008) Steady progress and recent breakthroughs in the accuracy of automated genome annotation. *Nature reviews Genetics* 9: 62-73.
- [8] Alioto T, Guig R (2008). State of the art in eukaryotic gene prediction.
- [9] Lin MF, Deoras AN, Rasmussen MD, Kellis M (2008) Performance and scalability of discriminative metrics for comparative gene identification in 12 *Drosophila* genomes. *PLoS Comput Biol* 4: e1000067.
- [10] Stark A, Lin MF, Kheradpour P, Pedersen JS, Parts L, et al. (2007) Discovery of functional elements in 12 *Drosophila* genomes using evolutionary signatures. *Nature* 450: 219-232.
- [11] Guttman M, Amit I, Garber M, French C, Lin MF, et al. (2009) Chromatin signature reveals over a thousand highly conserved large non-coding RNAs in mammals. *Nature* 458: 223-227.
- [12] Anisimova M, Kosiol C (2009) Investigating protein-coding sequence evolution with probabilistic codon substitution models. *Molecular biology and evolution* 26: 255-271.
- [13] Delport W, Scheffler K, Seoighe C (2009) Models of coding sequence evolution. *Briefings in bioinformatics* 10: 97-109.
- [14] Whelan S, Goldman N (1999) Distributions of statistics used for the comparison of models of sequence evolution in phylogenetics. *Mol Biol Evol* 16: 1292-1299.
- [15] Ota R, Waddell PJ, Hasegawa M, Shimodaira H, Kishino H (2000) Appropriate likelihood ratio tests and marginal distributions for evolutionary tree models with constraints on parameters. *Molecular biology and evolution* 17: 798-803.
- [16] Vuong QH (1989) Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57: 307-333.
- [17] Yang Z, Nielsen R (1998) Synonymous and nonsynonymous rate variation in nuclear genes of mammals. *Journal of Molecular Evolution* 46: 409-418.
- [18] Yang Z (2007) Paml 4: Phylogenetic analysis by maximum likelihood. *Molecular biology and evolution* 24: 1586-1591.
- [19] Kosiol C, Holmes I, Goldman N (2007) An empirical codon model for protein sequence evolution. *Molecular biology and evolution* 24: 1464-1479.
- [20] Goldman N, Yang Z (1994) A codon-based model of nucleotide substitution for protein-coding DNA sequences. *Molecular biology and evolution* 11: 725-736.

- [21] Blanchette M, Kent WJ, Riemer C, Elnitski L, Smit AFA, et al. Aligning multiple genomic sequences with the threaded blockset aligner. *Genome research* 14: 708-715.
- [22] Drosophila 12 Genomes Consortium (2007) Evolution of genes and genomes on the drosophila phylogeny. *Nature* 450: 203-218.
- [23] Felsenstein J (2004) *Inferring Phylogenies*. Sunderland, MA: Sinauer Associates, Inc.
- [24] Holmes I, Rubin GM (2002) An expectation maximization algorithm for training hidden substitution models. *Journal of Molecular Biology* 317: 753-764.
- [25] Siepel A, Haussler D (2004) Phylogenetic estimation of context-dependent substitution rates by maximum likelihood. *Mol Biol Evol* 21: 468-488.
- [26] Arvestad L, Bruno WJ (1997) Estimation of reversible substitution matrices from multiple pairs of sequences. *Journal of Molecular Evolution* 45: 696-703.