

BAYESIAN HIDDEN MARKOV TREE MODELS FOR CLUSTERING GENES WITH SHARED EVOLUTIONARY HISTORY¹

BY YANG LI^{*,§,2}, SHAOYANG NING^{*,2}, SARAH E. CALVO^{†,‡,§},
VAMSI K. MOOHA^{¶,†,‡,§} AND JUN S. LIU^{*}

Harvard University^{*}, Broad Institute[†], Harvard Medical School[‡], Massachusetts General Hospital[§], and Howard Hughes Medical Institute[¶]

Determination of functions for poorly characterized genes is crucial for understanding biological processes and studying human diseases. Functionally associated genes are often gained and lost together through evolution. Therefore identifying co-evolution of genes can predict functional gene-gene associations. We describe here the full statistical model and computational strategies underlying the original algorithm *CLustering by Inferred Models of Evolution* (CLIME 1.0) recently reported by us (*Cell* **158** (2014) 213–225). CLIME 1.0 employs a mixture of tree-structured hidden Markov models for gene evolution process, and a Bayesian model-based clustering algorithm to detect gene modules with shared evolutionary histories (termed evolutionary conserved modules, or ECMs). A Dirichlet process prior was adopted for estimating the number of gene clusters and a Gibbs sampler was developed for posterior sampling. We further developed an extended version, CLIME 1.1, to incorporate the uncertainty on the evolutionary tree structure. By simulation studies and benchmarks on real data sets, we show that CLIME 1.0 and CLIME 1.1 outperform traditional methods that use simple metrics (e.g., the Hamming distance or Pearson correlation) to measure co-evolution between pairs of genes.

1. Introduction. The human genome encodes more than 20,000 protein-coding genes, of which a large fraction do not have annotated function to date (Galperin and Koonin (2010)). Predicting unknown member genes to biological pathways/complexes and the determination of function for poorly characterized genes are crucial for understanding biological processes and human diseases. It has been observed that functionally associated genes tend to be gained and lost together during evolution (Kensche et al. (2008), Pellegrini et al. (1999)). Identifying shared evolutionary history (aka, co-evolution) of genes can help predict functions for unstudied genes, reveal alternative functions for genes considered to be well characterized, propose new members of biological pathways, and provide new insights into human diseases.

Received June 2018; revised August 2018.

¹Supported in part by NSF Grant DMS-1613035, NIGMS Grant R01GM122080, and NIH Grant R35 GM122455-02.

²These authors contributed equally to the manuscript.

Key words and phrases. Co-evolution, Dirichlet process mixture model, evolutionary history, gene function prediction, tree-structured hidden Markov model.

Phylogenetic profile matrix

	S1	S2	S3	S4	S5	S6	S7	S8	
G1	1	1	1	1	0	0	0	0	distance = 1
G2	1	1	0	1	0	0	0	0	
G3	1	1	0	0	1	1	1	1	distance = 0
G4	1	1	1	0	1	0	1	1	
G5	1	1	1	1	0	0	0	0	
G6	1	1	1	0	1	0	1	1	

FIG. 1. A toy example of phylogenetic profile matrix for $N = 6$ genes ($G1, \dots, G6$) and $S = 8$ species ($S1, \dots, S8$). Blue and white squares respectively denote presence or absence of genes in corresponding genomes. $G1$ and $G2$ have Hamming distance 1, while $G4$ and $G6$ have Hamming distance 0.

The concept of “phylogenetic profiling” was first introduced by Pellegrini et al. (1999) to characterize phylogenetic distributions of genes. One can predict a gene’s function based on its phylogenetic similarity to those with known functions. Let the binary phylogenetic profile matrix $\mathbf{X}_{N \times S}$ denote the presence/absence of N genes across S species. Pellegrini et al. (1999) proposed to measure the “degree” of co-evolution of a pair of genes i and j as the Hamming distance (Hamming (1950)) between the i th and j th rows of \mathbf{X} . A toy example is shown in Figure 1. Various methods have since been developed (see, Kensche et al. (2008), for a review) and applied with success in predicting components for prokaryotic protein complexes (Pellegrini et al. (1999)); phenotypic traits such as pili, thermophily, and respiratory tract tropism (Jim et al. (2004)); cilia (Li et al. (2004)); mitochondrial complex I (Ogilvie, Kennaway and Shoubridge (2005), Pagliarini et al. (2008)); and small RNA pathways (Tabach et al. (2013)).

Currently there are more than 200 eukaryotic species with their genomes completely sequenced and about 2000 species with full genomes being sequenced (JGI GOLD³). The growing availability of genome sequences from diverse species provides us unprecedented opportunities to chart the evolutionary history of human genes. However, existing phylogenetic profiling methods still suffer from some limitations (Kensche et al. (2008)). First, most available methods perform only pairwise comparison between an input query gene and a candidate, and are thus unable to discover subtle patterns that show up only after aligning multiple input query genes. Such methods also cannot handle cases where members in the query gene set exhibit different phylogenetic profiles. Second, most methods ignore errors in phylogenetic profiles, which are often caused by inaccuracies in genome assembly, gene annotation, and detection of distant homologs (Trachana et al. (2011)). Third, most methods (with exceptions of Barker and Pagel (2005),

³JGI Genome Online Database: <https://gold.jgi.doe.gov/>.

Vert (2002), Von Mering et al. (2003), Zhou et al. (2006)) assume independence across input species, ignoring their phylogenetic relationships, for example, the tree structure of their evolutionary history. These methods are rather sensitive to the organisms' selection in the analysis. Currently available tree-based methods, however, are computationally cumbersome and hardly scalable for analyzing large input sets, let alone entire genomes (Barker, Meade and Pagel (2007), Barker and Pagel (2005)).

To cope with the aforementioned limitations, Li et al. (2014) introduced the two-step procedure *CLustering by Inferred Models of Evolution* (denoted by CLIME 1.0). In its Partition step, CLIME 1.0 clusters the input gene set \mathcal{G} into disjoint evolutionarily conserved modules (ECMs), simultaneously inferring the number of ECMs and each gene's ECM membership. In the Expansion step, CLIME 1.0 scores and ranks other genes not in \mathcal{G} according to a log-likelihood-ratio (LLR) statistic for their likelihood of being new members of an inferred ECM. Li et al. (2014) systematically applied CLIME 1.0 to over 1000 human canonical complexes and pathways, resulting in a discovery of unanticipated co-evolving components and new members of important gene sets.

We here provide a full statistical account of CLIME 1.0 and its computational strategies, evaluate CLIME 1.0's performances with extensive simulations, extend it to incorporate uncertainties in the phylogenetic tree structure, and compare CLIME 1.0 with existing methods such as BayesTraits. Finally we apply CLIME 1.0 to gene sets in OMIM (Online Mendelian Inheritance in Man) to reveal new insights on human genetic disorders. Compared with existing methods, by incorporating a coherent statistical model, CLIME 1.0 (1) takes proper account of the dependency between species; (2) automatically learns the number of distinct evolutionary modules in the input gene set \mathcal{G} ; (3) leverages information from the entire input gene set to more reliably predict new genes that have arisen with a shared pattern of evolutionary gains and losses; (4) uses the LLR statistic as a principled measure of co-evolution compared to naive metrics (e.g., Hamming distance, Pearson correlation).

Complementary to the original CLIME 1.0, we further provide an extended version, named CLIME 1.1, which inherits the Bayesian hidden Markov tree model from CLIME 1.0, but further accounts for the uncertainty of the input phylogenetic tree structure by incorporating a prior on the evolutionary tree. Instead of a single, fixed tree as by CLIME 1.0, CLIME 1.1 takes an empirical distribution of tree structures, in addition to the phylogenetic profiles of a given gene set, as input; infers the posterior of the hidden evolutionary histories, hidden cluster (ECM) labels and parameters, as well as the posterior of evolutionary tree structure through Gibbs sampling; eventually outputs the ECMs of input gene set in the Partition step, and then classify novel genes into inferred ECMs in the Expansion step.

Rather than using only a point tree estimate, CLIME 1.1 adds to the original CLIME 1.0 by allowing the estimation error in the tree-building process as well as the variability of phylogenetic trees among genes, and thus alleviating the risk

of misspecification in the tree structure. In practice, popular tree-building methods and softwares such as PhyML (Guindon et al. (2010)) and MrBayes (Ronquist and Huelsenbeck (2003)) characterize the uncertainty in the estimation with bootstrap or posterior tree samples. CLIME 1.1 can readily utilize such output samples as empirical approximation for tree prior distribution. We also compare CLIME 1.1 with CLIME 1.0 and other benchmark methods in extensive simulations and real data to showcase its features and strengths. We find that CLIME 1.1 is more robust and accurate when there is high uncertainty in tree estimation or gene-wise variability in the evolutionary tree structures.

The rest of this article is organized as follows. In Section 2, we introduce the tree-structured hidden Markov model (HMM) for genes' stochastic gain/loss events on a given phylogenetic tree, and the Dirichlet process mixture (DPM) model for clustering genes into modules with shared history. The Partition step of CLIME 1.0, which implements the Gibbs sampler to sample from the posterior distribution of the DPM model, is described in Section 3. The Expansion step is introduced in Section 4. In Section 5, we briefly introduce the pre-processing of CLIME 1.0. The extended model and inference procedure of CLIME 1.1 are described in Section 6. Simulation studies that compare CLIME 1.0 and CLIME 1.1 with hierarchical clustering are presented in Section 7. In Section 8, we apply CLIME 1.0 and 1.1 on real data, and use leave-one-out cross-validation to compare the performance of CLIME 1.0 with hierarchical clustering on gene sets from GO (Gene Ontology) and KEGG (Kyoto Encyclopedia of Genes and Genomes) databases. We conclude this paper with a discussion in Section 9.

2. Bayesian mixture of HMM on a phylogenetic tree.

2.1. *Notation.* Let \mathcal{G} denote the input gene set with n genes, and N be the total number of genes in the reference genome. Let \mathbf{X}_i be the phylogenetic profile of gene i , $i = 1, \dots, N$, and specifically, let \mathbf{X} denote the phylogenetic profile of the input gene set. For example, \mathcal{G} can be the set of 44 subunit genes of human mitochondrial complex I, and \mathbf{X} is their phylogenetic profile matrix; for reference genome, we have $N = 20,834$ human genes with their phylogenetic profile matrix denoted by $\mathbf{X}_{1:N}$. For notational simplicity, we let $1, \dots, n$ index the n genes in \mathcal{G} and let $n + 1, \dots, N$ index the rest in the genome. The input phylogenetic tree has S living species indexed by $1, \dots, S$, and $S - 1$ ancestral extinct species indexed by $S + 1, \dots, 2S - 1$. The $2S - 1$ living and extinct species are connected by the $2S - 2$ branches on the tree. For simplicity, we assume that the phylogenetic tree is binary, while the model and algorithm can be easily modified for non-binary input trees. For each gene $i = 1, \dots, N$, its *phylogenetic profile* is defined as the observed vector $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,S})$ with $X_{i,j} = 1$ or 0 denoting the presence or absence of gene i across the S extant species. Let $\mathbf{H}_i = (H_{i,1}, \dots, H_{i,2S-1})$ denote gene i th ancestral (unobserved) and extant presence/absence states in the $2S - 1$ species.

We call a cluster of genes with shared evolutionary history an evolutionarily conserved module (ECM). Let $\mathbf{I} = (I_1, \dots, I_n)$ denote the ECM assignment indicators of genes, where $I_i = k$ indicates that gene i is assigned to ECM k . We assume that each gene can only be “gained” once throughout the entire evolutionary history, which happens at branch λ_i , $i = 1, \dots, N$. Let $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_N)$ denote the gain nodes of the N genes, where $\lambda_i = s$ indicates that gene i was gained at tree node s . With the available data, we can estimate $\boldsymbol{\lambda}$ in the preprocessing stage as described in Section 5 with very small estimation error. We thus assume that $\boldsymbol{\lambda}$ is a known parameter throughout the main algorithm.

2.2. Tree-structured HMM for phylogenetic profiles. We introduce here a tree-structured HMM to model the presence/absence history and phylogenetic profile of genes. For each gene i , its complete evolutionary history $\mathbf{H}_i = (H_{i,1}, \dots, H_{i,2S-1})$ is only partially observed at the bottom level, that is, the phylogenetic profile vector $\mathbf{X}_i = (X_{i,1}, \dots, X_{i,S})$ is the observation of presence/absence states for only the living species, $H_{i,1}, \dots, H_{i,S}$. Due to sequencing and genome annotation errors, there are also observation errors on the presence/absence of genes. In other words, $X_{i,1}, \dots, X_{i,S}$ are noisy observations on $H_{i,1}, \dots, H_{i,S}$. We assume that genes in ECM k share the same set of branch-specific probabilities of gene loss for the $2S - 2$ branches, denoted by $\boldsymbol{\theta}_k = (\theta_{k,1}, \dots, \theta_{k,2S-2})$. For genes in ECM k , the transition of absence/presence states from its direct ancestor to species s is specified by transition matrix $\mathbf{Q}_{k,s}$,

$$\mathbf{Q}_{k,s} = \begin{matrix} & \begin{matrix} 0 & 1 \end{matrix} \\ \begin{matrix} 0 \\ 1 \end{matrix} & \begin{bmatrix} 1 & 0 \\ \theta_{k,s} & 1 - \theta_{k,s} \end{bmatrix} \end{matrix}.$$

Thus, for every evolutionary branch (after the gain branches $\boldsymbol{\lambda}$), there is a \mathbf{Q} matrix. We assume that once a gene got lost, it cannot be re-gained, which is realistic for eukaryotic species. Therefore the first row of $\mathbf{Q}_{k,s}$ indicates that the transition probability from absence to presence (re-gain) is 0. The second row shows our parameterization that the transition probability from presence to absence (gene loss) is $\theta_{k,s}$, and presence to presence is $1 - \theta_{k,s}$.

Let $\sigma(s)$ denote the direct ancestor species of s , and let set $\mathcal{T}(s)$ include all of the offspring species in the sub-tree rooted at node s . Obviously $H_{i,s} = 0$ if species s is not in $\mathcal{T}(\lambda_i)$. The likelihood function of evolutionary history \mathbf{H}_i conditional on gene i in ECM k is

$$\begin{aligned} & \Pr(\mathbf{H}_i \mid \boldsymbol{\theta}_k, I_i = k) \\ &= \begin{cases} \prod_{s \in \mathcal{T}(\lambda_i) \setminus \lambda_i} \mathbf{Q}_{k,s}(H_{i,\sigma(s)}, H_{i,s}) & \text{if } H_{i,s} = 0 \forall s \notin \mathcal{T}(\lambda_i), \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

To account for errors in determining the presence/absence of a gene, we allow each component of the observed phylogenetic profile, $X_{i,s}$, to have an independent

probability q to be erroneous (i.e., different from the true state $H_{i,s}$). The error probability q is low and assumed to be known. By default, we set $q = 0.01$ based on our communication with biologists with expertise in genome sequencing and annotation. We note that estimating it in the MCMC procedure is straightforward, but a strong prior on q is needed for its proper convergence and identifiability. For each gene i , the likelihood function of \mathbf{X}_i given \mathbf{H}_i is

$$(1) \quad \Pr(\mathbf{X}_i | \mathbf{H}_i) = \prod_{s=1}^S \Pr(X_{i,s} | H_{i,s}) = \prod_{s=1}^S (1 - q)^{\mathbb{I}\{X_{i,s}=H_{i,s}\}} (q)^{\mathbb{I}\{X_{i,s} \neq H_{i,s}\}},$$

where $\mathbb{I}\{\cdot\}$ is the indicator function that is equal to 1 if the statement is true, and 0 otherwise. The complete likelihood for gene i is

$$(2) \quad \Pr(\mathbf{X}_i, \mathbf{H}_i | \boldsymbol{\theta}, I_i) = \left[\prod_{s \in \mathcal{T}(\lambda_i) \setminus \lambda_i} \mathbf{Q}_{i,s}(H_{i,\sigma(s)}, H_{i,s}) \right] \left[\prod_{s=1}^S (1 - q)^{\mathbb{I}\{X_{i,s}=H_{i,s}\}} (q)^{\mathbb{I}\{X_{i,s} \neq H_{i,s}\}} \right]$$

and the complete likelihood for all the genes is

$$(3) \quad \Pr(\mathbf{X}, \mathbf{H} | \boldsymbol{\theta}, \mathbf{I}) = \prod_{i=1}^n \Pr(\mathbf{X}_i, \mathbf{H}_i | \boldsymbol{\theta}, I_i).$$

2.3. *Dirichlet process mixture of tree hidden Markov models.* The number of ECMs K may be specified by users reflecting their prior knowledge on the data set. When the prior information about the data set is not available, we can estimate K from data by MCMC sampling with a Dirichlet process prior on $\boldsymbol{\theta}$ (Ferguson (1973), Neal (2000)). For each gene $i \in \{1, \dots, n\}$, we let the prior distribution of $\boldsymbol{\theta}_i$ follow Dirichlet process with concentration parameter α and base distribution \mathcal{F}_0 , denoted by $\text{DP}(\mathcal{F}_0, \alpha)$. This gives us the following Bayesian hierarchical model. For each gene $i = 1, \dots, n$,

$$(4) \quad \begin{aligned} \mathbf{X}_i | \mathbf{H}_i &\sim P(\mathbf{X}_i | \mathbf{H}_i), \\ \mathbf{H}_i | \boldsymbol{\theta}_i &\sim P(\mathbf{H}_i | \boldsymbol{\theta}_i), \\ \boldsymbol{\theta}_i | \mathcal{F} &\sim \mathcal{F}, \\ \mathcal{F} &\sim \text{DP}(\mathcal{F}_0, \alpha), \\ \mathcal{F}_0 &= \prod_{s=1}^{2S-2} \text{Beta}(a, b). \end{aligned}$$

The base distribution \mathcal{F}_0 is set as the product of a set of Beta distributions for branch-specific gene loss probabilities.

We use the Chinese restaurant process representation (Aldous (1985), Pitman (1996)) of the Dirichlet process and implement a Gibbs sampler (Gelfand and

Smith (1990), Liu (2008)) to draw from the posterior distribution of ECM assignments $\mathbf{I} = (I_1, \dots, I_n)$. The Chinese restaurant process prior for cluster assignments is exchangeable (Aldous (1985)), therefore the prior distribution for \mathbf{I} is invariant to the order of n genes. More precisely, the mixture model in equation (4) can be formulated as follows:

$$\begin{aligned}
 X_i | \mathbf{H}_i &\sim P(X_i | \mathbf{H}_i), & i = 1, 2, \dots, n, \\
 \mathbf{H}_i | \theta_{I_i} &\sim P(\mathbf{H}_i | \theta_{I_i}), & i = 1, 2, \dots, n, \\
 \theta_k &\sim \prod_{s=1}^{2S-2} \text{Beta}(a, b), & k = 1, 2, \dots,
 \end{aligned}
 \tag{5}$$

$$\begin{aligned}
 \Pr(I_i = I_j, j < i | I_1, \dots, I_{i-1}) &= n_{i,j} / (i - 1 + \alpha), & i = 1, 2, \dots, n, \\
 \Pr(I_i \neq I_j, \forall j < i | I_1, \dots, I_{i-1}) &= \alpha / (i - 1 + \alpha), & i = 1, 2, \dots, n,
 \end{aligned}$$

where $n_{i,j} = \sum_{l=1}^{i-1} \mathbb{I}\{I_l = I_j\}$.

2.4. *Dynamic programming for integrating out H.* In Section 3.3, we will introduce the Gibbs sampler to sample from the posterior distribution of \mathbf{I} . In the Gibbs sampler, we need to calculate the marginal probability of \mathbf{X}_i given the HMM parameter θ , with gene i 's evolutionary history \mathbf{H}_i integrated out. Suppose gene i is in ECM k , then

$$\Pr(X_i | \theta_k) = \sum_{\mathbf{H}_i} \Pr(X_i, \mathbf{H}_i | \theta_k).$$

We use the following tree-version of the backward procedure to calculate this marginal probability. For gene i , define X_i^s as its phylogenetic profile in the subtree rooted at species s (obviously $X_i^{2S-1} = X_i$). We calculate the marginal probability by recursively computing factors $\beta_{i,s}(h)$, defined as

$$\beta_{i,s}(h) \equiv \Pr(X_i^s | \theta_k, H_{i,s} = h).$$

For a living species s , which is a leaf of the tree,

$$\beta_{i,s}(h) = \Pr(X_i^s | \theta_k, H_{i,s} = h) = (1 - q)^{\mathbb{I}\{X_i^s = h\}} (q)^{\mathbb{I}\{X_i^s \neq h\}}.$$

Let $\delta_1(s)$ and $\delta_2(s)$ denote those two children species of s . For an inner tree species s , we can factorize $\beta_{i,s}(t)$ as

$$\begin{aligned}
 \beta_{i,s}(h) &= \sum_{h_1, h_2 \in \{0,1\}} \Pr(X_i^s, H_{i,\delta_1(s)} = h_1, H_{i,\delta_2(s)} = h_2 | \theta_k, H_{i,s} = h) \\
 &= \sum_{h_1, h_2 \in \{0,1\}} \Pr(X_i^s | \theta_k, H_{i,\delta_1(s)} = h_1) \cdot \Pr(H_{i,\delta_1(s)} = h_1 | \theta_k, H_{i,s} = h) \\
 &\quad \cdot \Pr(X_i^s | \theta_k, H_{i,\delta_2(s)} = h_2) \cdot \Pr(H_{i,\delta_2(s)} = h_2 | \theta_k, H_{i,s} = h) \\
 &= \left[\sum_{h_1 \in \{0,1\}} \beta_{i,\delta_1(s)}(h_1) \mathbf{Q}_{k,\delta_1(s)}(h, h_1) \right] \left[\sum_{h_2 \in \{0,1\}} \beta_{i,\delta_2(s)}(h_2) \mathbf{Q}_{k,\delta_2(s)}(h, h_2) \right].
 \end{aligned}$$

For each gene i , we calculate the β 's recursively bottom-up along the tree, until the gain branch λ_i , resulting in the marginal probability:

$$\begin{aligned}
 \Pr(X_i | \theta_k) &= \sum_{h \in \{0,1\}} \Pr(X_i^{\lambda_i} | \theta_k, H_{i,\lambda_i} = h) \Pr(H_{i,\lambda_i} = h | \theta_k) \\
 (6) \qquad &= 0 + \Pr(X_i^{\lambda_i} | \theta_k, H_{i,\lambda_i} = 1) \stackrel{\text{def}}{=} \beta_{i,\lambda_i}(1).
 \end{aligned}$$

2.5. *Dynamic programming for integrating out θ .* In each step of the Gibbs sampler, we pull out each gene from its current ECM and either re-assign it to an existing ECM or create a new singleton ECM for it according to the calculated conditional probability $\Pr(I_i | X_i, \mathbf{H}_i, \theta)$. For each ECM k , its parameter $\theta_k = \{\theta_{k,s}\}_{s=1}^{2S-2}$ is a vector containing $2S - 2$ loss probabilities. Our real data has $S = 139$, which makes each θ_k a 276-dimensional vector. The high dimensionality of $\theta_1, \dots, \theta_K$ adds heavy computational burden and dramatically slows down the convergence rate of the Gibbs sampler. To overcome this difficulty, we develop a collapsed Gibbs sampler (Liu (1994)) by applying the predictive updating technique (Chen and Liu (1996)) to improve the MCMC sampling efficiency. In particular, we integrate θ_k out from the conditional probability $\Pr(I_i = k | X_i, \mathbf{H}_i, \theta_k)$, so that

$$\begin{aligned}
 \Pr(I_i = k | X_i, \mathbf{H}, \mathbf{I}_{-i}) &= \int \Pr(I_i = k | X_i, \mathbf{H}, \theta_k) \Pr(\theta_k | X_i, \mathbf{H}, \mathbf{I}_{-i}) d\theta_k \\
 &\propto \Pr(X_i | \mathbf{H}_{-i}^k, I_i = k) \Pr(I_i = k | \mathbf{I}_{-i}),
 \end{aligned}$$

where $\mathbf{H}^k = \{\mathbf{H}_j : I_j = k, j = 1, \dots, n\}$ denotes the evolutionary histories for genes in ECM k , and $\mathbf{H}_{-i}^k = \mathbf{H}^k \setminus \{\mathbf{H}_i\}$. $\Pr(I_i = k | \mathbf{I}_{-i}) = \sum_{j \neq i} \mathbb{I}\{I_j = k\} / (n - 1 + \alpha)$ is the Chinese restaurant prior on \mathbf{I} , and $\Pr(X_i | \mathbf{H}_{-i}^k, I_i = k)$ is the marginal likelihood of X_i conditional on gene i is in ECM k with θ_k integrated out. We calculate $\Pr(X_i | \mathbf{H}_{-i}^k, I_i = k)$ as follows.

Conditional on \mathbf{H}_{-i}^k , the distribution of $\theta_{k,s}$, $s = 1, \dots, 2S - 2$, is simply a conjugate Beta posterior distribution,

$$\begin{aligned}
 \theta_{k,s} | \mathbf{H}_{-i}^k &\sim \text{Beta} \left(a + \sum_{j \neq i, I_j = k} \mathbb{I}\{H_{j,\sigma(s)} = 1, H_{j,s} = 0\}, \right. \\
 &\qquad \left. b + \sum_{j \neq i, I_j = k} \mathbb{I}\{H_{j,\sigma(s)} = 1, H_{j,s} = 1\} \right).
 \end{aligned}$$

Integrating out θ_k with respect to this distribution, we obtain the likelihood of X_i conditional on \mathbf{H}_{-i}^k :

$$\begin{aligned}
 \Pr(X_i | \mathbf{H}_{-i}^k, I_i = k) &= \int \Pr(X_i | \theta_k, I_i = k) \Pr(\theta_k | \mathbf{H}_{-i}^k) d\theta_k \\
 (7) \qquad &= \int \beta_{i,\lambda_i}(1) \Pr(\theta_k | \mathbf{H}_{-i}^k) d\theta_k = \bar{\beta}_{i,\lambda_i}(1),
 \end{aligned}$$

where $\bar{\beta}$ is defined as

$$\bar{\beta}_{i,s}(h) \equiv \mathbb{E}[\beta_{i,s}(h) \mid \mathbf{H}_{-i}^k] = \mathbb{E}[\Pr(\mathbf{X}_i^s \mid \boldsymbol{\theta}_k, H_{i,s} = h) \mid \mathbf{H}_{-i}^k].$$

For a leaf species s , $\bar{\beta}_{i,s}(h) = \beta_{i,s}(h)$. For an inner tree species s , $\bar{\beta}_{i,s}(h)$ can be calculated recursively from bottom of the tree to the top as

$$\begin{aligned} \bar{\beta}_{i,s}(h) &= \mathbb{E}[\beta_{i,s}(h) \mid \mathbf{H}_{-i}^k] \\ &= \left[\sum_{h_1=0,1} \bar{\beta}_{i,\delta_1(s)}(h_1) \bar{\mathbf{Q}}_{k,\delta_1(s)}(h, h_1) \right] \left[\sum_{h_2=0,1} \bar{\beta}_{i,\delta_2(s)}(h_2) \bar{\mathbf{Q}}_{k,\delta_2(s)}(h, h_2) \right], \end{aligned}$$

where $\bar{\mathbf{Q}}_{k,s}$ is the expectation of transition probability matrix $\mathbf{Q}_{k,s}$ conditional on \mathbf{H}_{-i}^k ,

$$(8) \quad \bar{\mathbf{Q}}_{k,s} = \mathbb{E}[\mathbf{Q}_{k,s} \mid \mathbf{H}_{-i}^k] = \begin{bmatrix} 1 & 0 \\ \mathbb{E}[\theta_{k,s} \mid \mathbf{H}_{-i}^k] & 1 - \mathbb{E}[\theta_{k,s} \mid \mathbf{H}_{-i}^k] \end{bmatrix},$$

and $\mathbb{E}[\theta_{k,s} \mid \mathbf{H}_{-i}^k]$ is simply the expectation of a Beta conjugate posterior distribution.

$$\mathbb{E}[\theta_{k,s} \mid \mathbf{H}_{-i}^k] = \frac{a + \sum_{j:I_j=k, j \neq i} \mathbb{I}\{H_{j,\delta(s)} = 1, H_{j,s} = 0\}}{a + b + \sum_{j:I_j=k, j \neq i} \mathbb{I}\{H_{j,\delta(s)} = 1\}}.$$

In the Gibbs sampler, we also need to compute the marginal probability that gene i is in its own singleton group, that is, $\Pr(\mathbf{X}_i \mid I_i \neq I_j, \forall j \neq i)$. By integrating out \mathbf{H}_i and $\boldsymbol{\theta}_i$, we have

$$\begin{aligned} \Pr(\mathbf{X}_i \mid I_i \neq I_j, \forall j \neq i) &= \int \sum_{\mathbf{H}_i} \Pr(\mathbf{X}_i, \boldsymbol{\theta}_i, \mathbf{H}_i \mid I_i \neq I_j, \forall j \neq i) d\boldsymbol{\theta}_i \\ (9) \quad &= \int \Pr(\mathbf{X}_i \mid \boldsymbol{\theta}_i, I_i \neq I_j, \forall j \neq i) d\mathcal{F}_0(\boldsymbol{\theta}_i) \\ &= \int \beta_{i,\lambda_i}(1) \Pr(\boldsymbol{\theta}_i) d\boldsymbol{\theta}_i. \end{aligned}$$

Note that (9) is a special case of (7) with $\mathbf{H}_{-i}^k = \emptyset$, thus it can be calculated in the same recursive way with

$$\bar{\mathbf{Q}}_{k,s} = \mathbb{E}[\mathbf{Q}_{k,s} \mid \mathbf{H}_{-i}^k = \emptyset] = \begin{bmatrix} 1 & 0 \\ a/(a+b) & b/(a+b) \end{bmatrix}.$$

2.6. ECM strength measurement. After partitioning the input gene set \mathcal{G} into ECMs, it is of great interest to determine which of the ECMs share more informative and coherent evolutionary histories than others, since the ranking of ECMs leads to different priorities for further low-throughput experimental investigations. In our Bayesian model-based framework, the strength of ECM k , denoted by ϕ_k ,

is defined as the logarithm of the Bayes Factor between two models normalized by the number of genes in that ECM. The first model is under the assumption that these genes have co-evolved in the same ECM and share the same θ parameter, and the second model is under the assumption that each gene has evolved independently in its own singleton ECM with different θ s. Specifically, with a partitioning configuration I , the strength for ECM k is defined as

$$(10) \quad \phi_k = \left\{ \log \left[\frac{\int \prod_{i:I_i=k} \Pr(X_i | \theta) \Pr(\theta) d\theta}{\prod_{i:I_i=k} \int \Pr(X_i | \theta) \Pr(\theta) d\theta} \right] \right\} / \sum_{i=1}^n \mathbb{I}\{I_i = k\}.$$

This strength measurement reflects the level of homogeneity among the evolutionary histories of genes in this ECM. A larger ϕ_k indicates that genes in ECM k share more similar and informative evolutionary history with more branches having high loss probabilities.

3. Partition step: MCMC sampling and point estimators.

3.1. *Choice of hyper-parameters.* Several hyper-parameters need to be specified, including the concentration parameter α in the Dirichlet process prior and hyper-parameters a, b for the Beta prior of θ s. Concentration parameter α controls the prior belief for the number of components in the mixture model, as larger α makes it easier to create a new ECM in each step of the Gibbs sampling. We set Dirichlet process concentration parameter as widely used $\alpha = 1$. To test the method’s robustness on α , we applied the algorithm to simulated and real data with $\alpha = 1, \alpha = \log(n)$, and $\alpha = \sqrt{n}$ respectively, and observed no significant changes on the posterior distribution of K . The reason is that histories of ECMs are often so different from each other that the likelihood function dominates the prior on determining K .

We set hyper-parameters $\alpha = 0.03, \beta = 0.97$ to make the prior have mean 0.03, which reflects our belief that overall 3% of times a gene gets lost when evolving from one species to another on a branch of the tree. The 3% average loss probability was determined based on the genome-wide average loss rate observed in our data.

3.2. *Forward-backward sampling for H .* In the Gibbs sampler, we apply a tree-version of forward-summation-backward-sampling method (Liu (2008), Section 2.4) to sample/impute the hidden evolutionary history states in H . Conditional on gene i is in ECM k , we want to sample H_i from the conditional distribution $\Pr(H_i | X_i, \theta_k)$. Note that, by the Markovian structure of tree HMM, $\Pr(H_i | X_i, \theta_k)$ can be written as

$$\Pr(H_i | X_i, \theta_k) = \begin{cases} \prod_{s \in \mathcal{T}(\lambda_i) \setminus \lambda_i} \Pr(H_{i,s} | H_{i,\sigma(s)}, X_i, \theta_k) & \text{if } H_{i,s} = 0 \forall s \notin \mathcal{T}(\lambda_i), \\ 0 & \text{otherwise,} \end{cases}$$

which suggests a sequential sampling procedure: draw $H_{i,s}$ for each species $s \in \mathcal{T}(\lambda_i) \setminus \lambda_i$ top-down along the tree from $\Pr(H_{i,s} | H_{i,\sigma(s)}, \mathbf{X}_i, \boldsymbol{\theta}_k)$ conditional on the previously drawn state $H_{i,\sigma(s)}$ of its ancestral species $\sigma(s)$.

We first use the backward procedure described in Section (2.4) to calculate the $\beta_{i,s}$ for all species $s \in \mathcal{T}(\lambda_i) \setminus \lambda_i$ bottom-up along the tree, then we have

$$\begin{aligned} \Pr(H_{i,s} | H_{i,\sigma(s)}, \mathbf{X}_i, \boldsymbol{\theta}_k) &\propto \Pr(H_{i,s}, \mathbf{X}_i^s | H_{i,\sigma(s)}, \boldsymbol{\theta}_k) \\ &= \Pr(\mathbf{X}_i^s | H_{i,s}, \boldsymbol{\theta}_k) \cdot \Pr(H_{i,s} | H_{i,\sigma(s)}, \boldsymbol{\theta}_k) \\ &= \beta_{i,s}(H_{i,s}) \cdot \mathbf{Q}_{k,s}(H_{i,\sigma(s)}, H_{i,s}). \end{aligned}$$

Similar to Section 2.5, we integrate out $\boldsymbol{\theta}_k$ to derive that

$$\begin{aligned} &\Pr(\mathbf{X}_i, \mathbf{H}_i | \mathbf{H}_{-i}, I_i = k) \\ (11) \quad &= \int \Pr(\mathbf{X}_i, \mathbf{H}_i | \boldsymbol{\theta}_k) \Pr(\boldsymbol{\theta}_k | \mathbf{H}_{-i}, I_i = k) d\boldsymbol{\theta}_k \\ &= \left[\prod_{s \in \mathcal{T}(\lambda_i) \setminus \lambda_i} \bar{\mathbf{Q}}_{k,s}(H_{i,\sigma(s)}, H_{i,s}) \right] \left[\prod_{s=1}^S (1-q)^{\mathbb{I}\{X_{i,s}=H_{i,s}\}} q^{\mathbb{I}\{X_{i,s} \neq H_{i,s}\}} \right], \end{aligned}$$

where $\bar{\mathbf{Q}}_{k,s}$ was defined in equation (8). Obviously, equation (11) is in the same form as the complete likelihood in equation (2) with transition probabilities matrix $\mathbf{Q}_{k,s}$ replaced by $\bar{\mathbf{Q}}_{k,s}$. The sequential sampling strategy for \mathbf{H}_i from conditional distribution $\Pr(\mathbf{H}_i | \mathbf{X}_i, \mathbf{H}_{-i}, I_i = k)$ is to start with $H_{i,\lambda_i} = 1$ and draw $H_{i,s}$ for each species $s \in \mathcal{T}(\lambda_i) \setminus \lambda_i$ top-down along the tree from distribution $\Pr(H_{i,s} | H_{i,\sigma(s)}, \mathbf{X}_i, \mathbf{H}_{-i}, I_i = k)$ conditional on the sampled state $H_{i,\sigma(s)}$ of its ancestral species $\sigma(s)$, with matrices $\mathbf{Q}_{k,s}$ replaced by $\bar{\mathbf{Q}}_{k,s}$.

3.3. Gibbs sampling implementation. In each step of Gibbs sampling, we pull out each gene from its current ECM and assign it to an existing ECM or create a new singleton ECM for it with respect to the calculated conditional distribution $\Pr(I_i | \mathbf{X}_i, \mathbf{H}, \mathbf{I}_{-i})$, which is calculated as

$$(12) \quad \begin{aligned} &\Pr(I_i = k | \mathbf{X}_i, \mathbf{H}, \mathbf{I}_{-i}) \\ &\propto \begin{cases} \frac{\sum_{j:j \neq i} \mathbb{I}\{I_j = k\}}{n_\alpha - 1 + \alpha} \cdot \Pr(\mathbf{X}_i | \mathbf{H}_{-i}, I_i = k) & \exists j \neq i, \text{ s.t. } I_j = k, \\ \frac{1}{n - 1 + \alpha} \cdot \Pr(\mathbf{X}_i | I_i \neq I_j, \forall j \neq i) & \text{otherwise,} \end{cases} \end{aligned}$$

where $\Pr(\mathbf{X}_i | \mathbf{H}_{-i}, I_i = k)$ and $\Pr(\mathbf{X}_i | I_i \neq I_j, \forall j \neq i)$ are respectively calculated in equations (7) and (9).

We implement the collapsed Gibbs sampler to calculate the posterior distribution of \mathbf{I} and \mathbf{H} . In each Gibbs sampler iteration, we conduct the following two steps:

1. Draw $\mathbf{H}_i \sim \Pr(\mathbf{H}_i \mid \mathbf{X}_i, \mathbf{H}_{-i}, \mathbf{I})$, $i = 1, \dots, n$ by the procedure in Section 3.2.
2. Draw $I_i \sim \Pr(I_i \mid \mathbf{X}_i, \mathbf{H}, \mathbf{I}_{-i})$, $i = 1, \dots, n$ as calculated in equation (12).

By using this Gibbs sampling scheme, genes with similar evolutionary history will be clustered to the same ECM, and genes without any close neighbor will stay in their own singleton ECMs. This automatically estimates the number of ECMs K .

We implemented this Gibbs sampler in C++, and tested its computational efficiency. On a typical input gene set with ~ 100 genes across 139 species, the Gibbs sampler takes about 30 minutes to finish 1000 iterations on a standard Linux server using a single CPU. For input gene sets of size 5000, the Gibbs sampler takes less than 24 hours to finish 1000 iterations.

3.4. *Point estimator for ECM assignments \mathbf{I} .* While the posterior distribution of \mathbf{I} is calculated by the Gibbs sampler, users may prefer a single optimal solution for \mathbf{I} as it is easier to interpret and proceed to further experimental investigations. To obtain a point estimator of \mathbf{I} , we calculate the posterior probability $\Pr(\mathbf{I} \mid \mathbf{X})$ at the end of each Gibbs sampling iteration. The *maximum a posteriori* (MAP) assignment, $\arg \max_{\mathbf{I}} \Pr(\mathbf{I} \mid \mathbf{X})$, will be reported as the final MAP estimation. Suppose we have M MCMC samples, denoted by $\mathbf{I}^{(1)}, \dots, \mathbf{I}^{(M)}$, then the MAP assignment can be approximated by

$$\hat{\mathbf{I}} = \arg \max_{\mathbf{I}^{(m)}; m=1, \dots, M} \Pr(\mathbf{I}^{(m)} \mid \mathbf{X}).$$

We know that

$$\Pr(\mathbf{I} \mid \mathbf{X}) \propto \Pr(\mathbf{X} \mid \mathbf{I}) \Pr(\mathbf{I}),$$

where $\Pr(\mathbf{I})$ is the Chinese restaurant process prior,

$$\Pr(\mathbf{I}) = \frac{\prod_{k=1}^K (n_k - 1)!}{n!}, \quad \text{where } n_k = \sum_{i=1}^n \mathbb{I}\{I_i = k\},$$

and $\Pr(\mathbf{X} \mid \mathbf{I}) = \prod_{k=1}^K \Pr(\mathbf{X}_k \mid \mathbf{I})$, where $\mathbf{X}_k = \{\mathbf{X}_i : I_i = k, i = 1, \dots, n\}$ and $\Pr(\mathbf{X}_k \mid \mathbf{I})$ is the marginal probability for phylogenetic profiles of genes in ECM k , that is,

$$\Pr(\mathbf{X}_k \mid \mathbf{I}) = \int \left[\prod_{i: I_i=k} \Pr(\mathbf{X}_i \mid \boldsymbol{\theta}_k) \right] \Pr(\boldsymbol{\theta}_k) d\boldsymbol{\theta}_k.$$

This integral has no closed-form solution, but we can approximate this marginal likelihood by the method in Chib (1995) using samples obtained by the Gibbs sampler. In particular, we have the following equation holds for any $\boldsymbol{\theta}_k^* = (\theta_{k,1}^*, \dots, \theta_{k,2S-1}^*)$:

$$(13) \quad \log \Pr(\mathbf{X}_k \mid \mathbf{I}) = \sum_{i: I_i=k} \log \Pr(\mathbf{X}_i \mid \boldsymbol{\theta}_k^*) + \log \Pr(\boldsymbol{\theta}_k^*) - \log \Pr(\boldsymbol{\theta}_k^* \mid \mathbf{X}_k, \mathbf{I}).$$

In the equation above, prior probability $\Pr(\theta_k^*)$ can be calculated directly and the likelihood $\Pr(X_i | \theta_k^*)$ can be calculated by dynamic programming with computational complexity $O(S)$. We approximate $\Pr(\theta_k^* | X_k, \mathbf{I})$ by running additional Gibbs sampling. Let $\mathbf{H}_k = \{H_i : I_i = k, i = 1, \dots, n\}$. We fix ECM assignments at \mathbf{I} and re-run Gibbs sampler for T iterations to draw samples $\{\mathbf{H}_k^{(1)}, \dots, \mathbf{H}_k^{(M)}\}$ from $\Pr(\mathbf{H}_k | X_k, \mathbf{I})$, and then $\Pr(\theta_k^* | X_k, \mathbf{I})$ can be approximated as

$$(14) \quad \Pr(\theta_k^* | X_k, \mathbf{I}) = \sum_{\mathbf{H}_k} \Pr(\theta_k^* | \mathbf{H}_k) \Pr(\mathbf{H}_k | X_k, \mathbf{I}) \approx \frac{1}{M} \sum_{m=1}^M \Pr(\theta_k^* | \mathbf{H}_k^{(m)}),$$

where

$$\Pr(\theta_k^* | \mathbf{H}_k^{(m)}) = \prod_{s=1}^{2S-2} \text{Be} \left(\theta_{k,s}^* | a + \sum_{i:I_i=k} \mathbb{I}\{H_{i,\delta(s)}^{(m)} = 1, H_{i,s}^{(m)} = 0\}, \right. \\ \left. b + \sum_{i:I_i=k} \mathbb{I}\{H_{i,\delta(s)}^{(m)} = 1, H_{i,s}^{(m)} = 1\} \right).$$

$\text{Be}(\theta | \alpha, \beta)$ is the Beta density function. Plug equation (14) in equation (13), we get the approximation for marginal likelihood $\Pr(X_k | \mathbf{I})$.

Though the approximation is consistent for any θ_k^* , as pointed out by Chib (1995), the choice of θ_k^* determines the efficiency of approximation. The approximation is likely to be more precise with a θ_k^* that is close to the true θ_k . A natural choice for θ_k^* is the posterior mean estimator of θ_k as calculated in equation (15).

3.5. Point estimator for loss probabilities θ . In the implementation of the Gibbs sampler, we integrate out the θ 's from the model and run the collapsed Gibbs sampler, which improves the MCMC sampling efficiency. After obtaining the final partitioning $\hat{\mathbf{I}}$, we want to calculate the point estimators for the θ 's for the K ECMs defined in $\hat{\mathbf{I}}$, denoted by $\{\hat{\theta}_1, \dots, \hat{\theta}_K\}$. For each ECM k , those branches with estimated high loss probabilities $\hat{\theta}_{k,s}$ are evolutionary signature of ECM k and distinguish it from other ECMs. In Section 4, we plug the estimated parameters $\{\hat{\theta}_1, \dots, \hat{\theta}_K\}$ into the likelihood ratio statistics to identify novel genes that are not in \mathcal{G} but share close history with any of the K ECMs. The point estimator of $\theta_{k,s}$ is defined as the posterior mean of $\theta_{k,s}$ conditional on \mathbf{X} and $\hat{\mathbf{I}}$, that is, $\hat{\theta}_{k,s} = \mathbb{E}[\theta_{k,s} | \mathbf{X}, \hat{\mathbf{I}}]$. To compute $\hat{\theta}_{k,s}$, we re-run the Gibbs sampler conditional on $\hat{\mathbf{I}}$ to draw $M = 1000$ samples $\mathbf{H}_k^{(1)}, \dots, \mathbf{H}_k^{(M)}$ from $\Pr(\mathbf{H}_k | \mathbf{X}, \hat{\mathbf{I}})$, where $\mathbf{H}_k = \{H_i : I_i = k, i = 1, \dots, n\}$. $\hat{\theta}_{k,s}$ is approximated by the following Rao–Blackwellized estimator (Liu, Wong and Kong (1994)):

$$(15) \quad \hat{\theta}_{k,s} \approx \frac{1}{M} \sum_{m=1}^M \mathbb{E}[\theta_{k,s} | \mathbf{H}_k^{(m)}, \hat{\mathbf{I}}],$$

where

$$(16) \quad \mathbb{E}[\theta_{k,s} | \mathbf{H}_k^{(m)}, \hat{\mathbf{I}}] = \frac{a + \sum_{i:\hat{I}_i=k} \mathbb{I}\{H_{i,\delta(s)}^{(m)} = 1, H_{i,s}^{(m)} = 0\}}{a + b + \sum_{i:\hat{I}_i=k} \mathbb{I}\{H_{i,\delta(s)}^{(m)} = 1\}}.$$

4. Expansion step: Identifying novel genes co-evolved with each ECM. In the Partition step, CLIME 1.0 clusters the input set \mathcal{G} into disjoint evolutionarily conserved modules (ECMs), simultaneously inferring the number of ECMs and each gene's ECM membership. The second step of CLIME 1.0, the Expansion step, identifies novel genes that are not in the input gene set \mathcal{G} but share evolutionary history with any ECM k identified in the Partition step. The Expansion step is essential to CLIME 1.0 as the main goal of it is to identify novel genes that are co-evolved with a subset of \mathcal{G} . The underlying logic is that if a ECM k consists of a large number of genes of \mathcal{G} , then the other genes not in \mathcal{G} but share history with ECM k are likely functionally associated with \mathcal{G} .

For each candidate gene g and ECM k , $g = 1, \dots, N$ and $k = 1, \dots, K$, we calculate the log-likelihood ratio (LLR),

$$\text{LLR}_{g,k} = \log \Pr(\mathbf{X}_g | \hat{\boldsymbol{\theta}}_k) - \log \Pr(\mathbf{X}_g | \hat{\boldsymbol{\theta}}_0),$$

where the background null model $\hat{\boldsymbol{\theta}}_0$ is defined as the estimated genome-wide average loss probabilities over all $N = 20,834$ human genes. The estimation of $\hat{\boldsymbol{\theta}}_0$ is straightforward and described in Section 5. In the LLR, the first term $\log \Pr(\mathbf{X}_g | \hat{\boldsymbol{\theta}}_k)$ quantifies the likelihood that \mathbf{X}_g was generated from the HMM of ECM k , and the second term $\log \Pr(\mathbf{X}_g | \hat{\boldsymbol{\theta}}_0)$ quantifies the likelihood that \mathbf{X}_g was generated from the background null HMM. High value of $\text{LLR}_{g,k}$ indicates that the HMM of ECM k explains the phylogenetic profile \mathbf{X}_g much better than the background null model, which suggests that gene g is more probable to share the same evolutionary history with the genes in ECM k , than a randomly selected gene in human genome.

For each ECM, CLIME 1.0 scores all $N - n$ human genes, ranks them by LLR scores, and reports the list of genes with $\text{LLR} > 0$ (denoted by ECM+). Compared to naïve metrics (e.g., Hamming distance, Pearson correlation between phylogenetic profiles), this LLR statistic measures co-evolution more appropriately and achieves substantially higher prediction sensitivity and specificity (see Section 8.2).

5. Preprocessing: Estimation of gain branches λ and background null model $\boldsymbol{\theta}_0$. In the preprocessing stage, CLIME 1.0 infers the gain branch λ_i for each gene i and estimates the background null model parameter $\hat{\boldsymbol{\theta}}_0$ for gene loss events from phylogenetic profiles of all human genes in the input matrix. The null model is an ECM-independent HMM whose branch-specific loss probabilities are averaged over all genes in the human genome.

We estimate $\boldsymbol{\theta}_0$ under the model that all $N = 20,834$ human genes share the same loss probability vector $\boldsymbol{\theta}_0$, that is, $\boldsymbol{\theta}_1 = \boldsymbol{\theta}_2 = \dots = \boldsymbol{\theta}_N = \boldsymbol{\theta}_0$, and implement

a Gibbs sampler to sample from the posterior distribution of $\Pr(\boldsymbol{\theta}_0, \boldsymbol{\lambda} \mid \mathbf{X}_{1:N})$. We start the Gibbs sampler from the initial state with $\boldsymbol{\theta}_0 = (0.03, \dots, 0.03)$ and $\boldsymbol{\lambda} = (2S - 1, \dots, 2S - 1)$. In each step of the Gibbs sampler, we conduct the following steps:

1. Draw $\lambda_i \sim \Pr(\lambda_i \mid \mathbf{X}_i, \boldsymbol{\theta}_0)$, $i = 1, \dots, N$.
2. Draw $\mathbf{H}_i \sim \Pr(\mathbf{H}_i \mid \mathbf{X}_i, \boldsymbol{\lambda}, \boldsymbol{\theta}_0)$, $i = 1, \dots, N$ by the forward-backward procedure.
3. Draw $\boldsymbol{\theta}_0 \sim \Pr(\boldsymbol{\theta}_0 \mid \mathbf{H}_{1:N}, \boldsymbol{\lambda})$, $i = 1, \dots, N$.

Both conditional distributions $\Pr(\lambda_i \mid \mathbf{X}_i, \boldsymbol{\theta}_0)$ and $\Pr(\boldsymbol{\theta}_0 \mid \mathbf{H}_{1:N}, \boldsymbol{\lambda})$ are straightforward to sample from. $\Pr(\lambda_i \mid \mathbf{X}_i, \boldsymbol{\theta}_0)$ is a discrete distribution and for $s = 1, \dots, 2S - 1$,

$$\Pr(\lambda_i = s \mid \mathbf{X}_i, \boldsymbol{\theta}_0) \propto \Pr(\mathbf{X}_i \mid \lambda_i = s, \boldsymbol{\theta}_0) \Pr(\lambda_i = s).$$

We adopt a uniform prior on $\Pr(\lambda_i = s) = 1/(2S - 1)$ and calculate likelihood function $\Pr(\mathbf{X}_i \mid \lambda_i = s, \boldsymbol{\theta}_0)$ with dynamic programming outlined in equation (6). $\Pr(\boldsymbol{\theta}_0 \mid \mathbf{H}_{1:N}, \boldsymbol{\lambda})$ is simply a product of Beta distributions, and each $\theta_{0,s}$, $s = 1, \dots, 2S - 2$, can be drawn independently. Similar to equation (15), we define the point estimator of $\boldsymbol{\theta}_0$ as $\hat{\boldsymbol{\theta}}_0 = \mathbb{E}[\boldsymbol{\theta}_0 \mid \mathbf{X}_{1:N}]$ and approximate it with MCMC samples. Suppose we have M MCMC samples on $\boldsymbol{\lambda}$, denoted by $\boldsymbol{\lambda}^{(1)}, \dots, \boldsymbol{\lambda}^{(M)}$. For each gene $i = 1, \dots, N$, we define $\hat{\lambda}_i$ as the *maximum a posteriori* (MAP) estimator approximated by MCMC samples,

$$\hat{\lambda}_i = \arg \max_s \sum_{m=1}^M \mathbb{I}\{\lambda_i^{(m)} = s\}.$$

In both the Partition and the Expansion steps of CLIME 1.0, gain branches $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_N)$ are considered as known and fixed. An alternative way for estimating the gain branch for each gene i in \mathcal{G} is to update λ_i in the Gibbs sampler of the Partitioning step and calculate their posterior distributions. There are two reasons why we chose to estimate the gain branch for each gene in the Preprocessing step and kept it fixed in the later two steps. First, the gain branches can usually be reliably estimated with little uncertainty. For example, if a gene i was truly gained at node s , then most likely we will observe its presences only in \mathbf{X}_i^s , which informs us that the gain event happened at node s . Second, by estimating the gain branches at the Preprocessing step, we reduce the computation complexity compared to a full model that updates $\boldsymbol{\lambda}$ at each MCMC iteration of the Partition step.

6. The extended model with uncertainty of phylogenetic tree.

6.1. *The extended model of CLIME 1.1.* Here we introduce the model of CLIME 1.1, which extends CLIME 1.0 by incorporating the uncertainty in phylogenetic trees. We keep the same notation as in the original CLIME 1.0. Conditioning on the tree structure T , we follow the same specification as in equation (5).

Additionally, we assume that the tree structure follows a prior $T \sim \mathcal{F}_T$, so that jointly we have:

$$\begin{aligned} X_i | \mathbf{H}_i^T, T &\sim P(X | \mathbf{H}_i^T), & i = 1, 2, \dots, n, \\ \mathbf{H}_i^T | \boldsymbol{\theta}_k^T, I_i = k, T &\sim P(\mathbf{H}_i^T | \boldsymbol{\theta}_k^T), & i = 1, 2, \dots, n, \\ \boldsymbol{\theta}_k^T &\sim \prod_{s=1}^{2S-2} \text{Beta}(a, b), & k = 1, 2, \dots, \\ I_i &\sim \text{CRP}(\alpha), & i = 1, 2, \dots, n, \\ T &\sim \mathcal{F}_T. \end{aligned}$$

Here, the superscript T indicates dependency on the tree structure, which will be suppressed in the following derivations for simplicity. In practice, we utilize the bootstrap samples or posterior draws of trees from the output of tree-constructing softwares to approximate the prior distribution \mathcal{F}_T . That is, suppose we have N_T sampled tree structures $\{T_1, \dots, T_{N_T}\}$, we assume that $\mathcal{F}_T = \frac{1}{N_T} \delta_{T_i}(T)$, where δ is the Dirac point mass. This distribution is derived based on a probabilistic model of evolution and can well characterize the variability in the estimation of the evolutionary tree.

6.2. Posterior inference of CLIME 1.1 with Gibbs sampler. We implement a collapsed Gibbs sampler (Liu (1994)) to draw from the posterior distribution, which cycles through the samplings of the hidden evolutionary history \mathbf{H} , the tree structure T , and the ECM label \mathbf{I} . The high-dimensional parameter vector $\boldsymbol{\theta}$ is integrated out throughout the process similarly as what we did for CLIME 1.0 to improve the sampling efficiency.

1. Sampling $[\mathbf{H} | \mathbf{X}, \mathbf{I}, T]$: For each gene i , we sample its evolutionary history \mathbf{H}_i from $\Pr(\mathbf{H}_i | \mathbf{X}, \mathbf{H}_{-i}, T, \mathbf{I})$, which can be achieved by the same procedure described in Section 3.2 to sample \mathbf{H}_i , conditioning on tree structure T .

2. Sampling $[\mathbf{I} | \mathbf{X}, \mathbf{H}, T]$: For each gene i , we sample its cluster label I_i from $\Pr(I_i | \mathbf{I}_{-i}, \mathbf{X}, \mathbf{H}_{-i}, T)$, which, conditioning on tree structure T , can be similarly calculated as in equation (12).

3. Sampling $[T | \mathbf{X}, \mathbf{I}]$: We sample T based on posterior

$$\Pr(T | \mathbf{X}, \mathbf{I}) \propto \mathcal{F}_T(T) \Pr(\mathbf{X} | T, \mathbf{I}).$$

Since the prior \mathcal{F}_T is taken as the empirical distribution $\frac{1}{N_T} \delta_{T_i}(T)$, we sample $T = T_i$ with probability proportional to $\Pr(\mathbf{X} | T_i, \mathbf{I})$, where $\Pr(\mathbf{X} | T_i, \mathbf{I})$ can be approximated by the method of Chib (1995) as in equation (13). Note that the conditional distribution $\Pr(\mathbf{X} | T_i, \mathbf{I})$ will be used again in the Partition step for calculating $\arg \max_{\mathbf{I}} \Pr(\mathbf{I} | \mathbf{X})$, and the Expansion step for calculating the LLR of novel genes.

6.3. *Partition step of CLIME 1.1.* We are mainly interested in estimating the ECM clustering labels of all input genes. Similar to CLIME 1.0, we adopt the MAP estimator $\hat{\mathbf{I}} = \arg \max_{\mathbf{I}} \Pr(\mathbf{I}|\mathbf{X})$, approximated by searching through all MCMC samples of \mathbf{I} , that is,

$$\hat{\mathbf{I}} = \arg \max_{\mathbf{I}^{(m)}:m=1,\dots,M} \Pr(\mathbf{I}^{(m)} | \mathbf{X}).$$

Specifically,

$$\Pr(\mathbf{I}|\mathbf{X}) \propto \int \Pr(\mathbf{X}, \mathbf{I}|T)\mathcal{F}_T(T) dT = \Pr(\mathbf{I}) \sum_{T_i} \frac{1}{N_T} \Pr(\mathbf{X}|\mathbf{I}, T_i),$$

where the conditional distribution $\Pr(\mathbf{X}|\mathbf{I}, T_i)$ has been calculated in Step 3 of the Gibbs sampler in Section 6.2, and the prior $\Pr(\mathbf{I})$ is assumed to be the Chinese restaurant process.

6.4. *Expansion step of CLIME 1.1.* Suppose a gene g 's phylogenetic profile being \mathbf{X}_g ($g = 1, \dots, N$). We calculate its LLR for all ECMs, $k = 1, \dots, K$, similarly as for CLIME 1.0, that is,

$$\text{LLR}_{g,k} = \log \Pr(\mathbf{X}_g|I_g = k, \mathbf{X}, \hat{\mathbf{I}}) - \log \Pr(\mathbf{X}_g|I_g = 0, \mathbf{X}, \hat{\mathbf{I}}),$$

where $I_g = 0$ indicates the background null model.

We calculate the predictive likelihood by integrating out θ_k and T :

$$\begin{aligned} &\Pr(\mathbf{X}_g|I_g = k, \mathbf{X}, \hat{\mathbf{I}}) \\ &= \int \Pr(\mathbf{X}_g|I_g = k, \theta_k, T) \Pr(\theta_k, T|\mathbf{X}, \hat{\mathbf{I}}) dT d\theta_k \\ &= \int \Pr(\mathbf{X}_g|I_g = k, \theta_k, T) \Pr(\theta_k|T, \hat{\mathbf{I}}, \mathbf{X}) \Pr(T|\mathbf{X}, \hat{\mathbf{I}}) dT d\theta_k \\ &\propto \int \Pr(\mathbf{X}_g|I_g = k, \theta_k, T) \Pr(\theta_k|T, \hat{\mathbf{I}}, \mathbf{X}) \Pr(\mathbf{X}|T, \hat{\mathbf{I}})\mathcal{F}_T(T) dT d\theta_k. \end{aligned}$$

Note that $\mathcal{F}_T = \frac{1}{N_T} \delta_{T_i}(T)$, and

$$\Pr(\theta_k|T, \hat{\mathbf{I}}, \mathbf{X}) = \int \Pr(\theta_k|T, \hat{\mathbf{I}}, \mathbf{H}) \Pr(\mathbf{H}|\mathbf{X}, T, \hat{\mathbf{I}}) d\mathbf{H},$$

which can be approximated using the Gibbs sampling draws as

$$\Pr(\theta_k|T, \hat{\mathbf{I}}, \mathbf{X}) \approx \frac{1}{M} \sum_{i=1}^M \Pr(\theta_k|T, \hat{\mathbf{I}}, \mathbf{H}^{(i)}).$$

Plugging in the foregoing integral, we have the following approximation

$$\Pr(\mathbf{X}_g|I_g = k, \mathbf{X}, \hat{\mathbf{I}}) \approx \frac{1}{N_T} \sum_{i=1}^{N_T} \left[\frac{1}{M} \sum_{m=1}^M \Pr(\mathbf{X}_g|I_g = k, \bar{\theta}_k^{(i,m)}, T_i) \Pr(\mathbf{X}|T_i, \hat{\mathbf{I}}) \right],$$

where $\bar{\theta}_k^{(i,m)} = E(\theta_k | T_i, \hat{\mathbf{I}}, \mathbf{H}^{(m)})$ can be calculated by conjugate Beta distribution as in equation (16); the predictive likelihood $\Pr(X_g | I_g = k, \bar{\theta}_k^{(i,m)}, T_i)$ can then be calculated by dynamic programming introduced in Section 2.4; and the likelihood of input gene set $\Pr(\mathbf{X} | T_i, \hat{\mathbf{I}})$ has been previously calculated in the step 3 of Gibbs sampler in Section 6.2.

7. Simulation studies. We simulated the phylogenetic profile data from two models: a tree-based hidden Markov model and a tree-independent model where CLIME 1.0 and CLIME 1.1's model is mis-specified. The simulated input gene sets contain 50 genes, comprising a mixture of 5 ECMs, each with 10 genes, whose phylogenetic profiles were generated using the tree-based and tree-independent models. We analyzed the data with four methods: (1) CLIME 1.0; (2) CLIME 1.1; (3) hierarchical clustering based on Hamming distance (Pellegrini et al. (1999)); (4) hierarchical clustering based on squared anti-correlation distance (Glazko and Mushegian (2004)), where the distance between gene i and j is defined as $d_{i,j} = 1 - [\text{corr}(\mathbf{X}_i, \mathbf{X}_j)]^2$.

For the tree-based hidden Markov model, we first used MrBayes (Ronquist and Huelsenbeck (2003)) to obtain 100 phylogenetic trees generated from the posterior distribution of the tree structure model based on 16 highly reserved proteins of 138 eukaryotic species (Bick, Calvo and Mootha (2012)) and an additional prokaryote outgroup (139 species in total). For each simulation, we randomly picked one of the 100 tree structures, and generated the phylogenetic profiles and ECM assignments based on the tree-based HMM and this picked tree structure. Note that here we simulated uncertainties in the tree structure. Thus, the original CLIME 1.0 with a single phylogenetic tree (the consensus) input runs the risk of tree misspecification for these simulated data. For each ECM, we first randomly selected one branch in the evolution tree to be the gain branch, and then, along its sub-tree, selected N_L branches to be the potential gene loss branches and assign P_L to be their gene loss probability to generate the phylogenetic profile of each gene. A higher P_L leads to a more similar evolutionary history among the simulated genes in the same ECM, and a lower P_L makes the underlying histories of genes less similar and adds more difficulty to the algorithms. We simulated observation error with rate $q = 0.02$, which is different from $q = 0.01$ as pre-specified in CLIME 1.0 and CLIME 1.1's model. In addition, we simulated $N_S \in \{0, 10, 20, 50\}$ singleton ECMs with one gene in each ECM as the background noise. Eventually, each input dataset contains a $(50 + N_S) \times 139$ binary matrix indicating the presence or absence of each gene in each species.

For the tree-independent generating model in comparison, N_L potential gene losses were randomly selected from 139 species without any reference to their evolutionary relations. Note that such a tree-independent model is equivalent to a tree-based model when all the losses are constrained to happen exclusively on leaf branches. We range $P_L \in \{0.6, 0.7, 0.8, 0.9\}$ and $N_L \in \{4, 6, 8, 10\}$ for both

the tree-based model and the tree-independent model. Higher N_S gives more noise and higher P_L and N_L indicate more independent loss events across various ECMs thus stronger signal.

For each set of parameters, we simulated phylogenetic profile matrices for 20 times, applied all four methods, and adopted the average adjusted Rand index (ARI) (Hubert and Arabie (1985)) between the estimated and true partitioning for these 20 simulated datasets to evaluate clustering accuracy. For CLIME 1.0, to be consistent with the online software, we used the consensus phylogenetic tree built from 16 highly reserved proteins of 138 species (Bick, Calvo and Mootha (2012)) with one outgroup prokaryote species as the single input tree structure, shown in Figure 4. For CLIME 1.1, we included the 100 MrBayes samples described above as the input for empirical prior of the tree structure to account for estimation uncertainty. For hierarchical clustering, we used 10% singleton genes as cutoff for clustering as adopted in (Glazko and Mushegian (2004)). The complete simulation results for tree-based model and tree-independent model are reported in Figures 2 and 3, respectively.

As shown in Figure 2, when phylogenetic profiles were generated from a tree-based model of evolution with the risk of tree misspecification, CLIME 1.1 dominates all other clustering methods in terms of accuracy with the tree-uncertainty taken into account. CLIME 1.0, in general, also holds the lead over hierarchical clustering methods. The advantages of our tree-based Markov model are even more significant in scenarios where more loss events are present along the evolutionary history, that is, more loss branches ($N_L \geq 6$) with higher ($P_L \geq 0.7$), to provide stronger signals for our tree-based model. Another feature of our methods is the robustness against the varying number of singleton ECMs, or the noise in clustering. As the noise level (N_S) increases, both CLIME 1.0 and CLIME 1.1 demonstrate consistent clustering accuracy, while hierarchical clustering methods show severe decay in performance. Notably, by incorporating the uncertainty of tree structure and weighting the clustering on the more probable tree structures, CLIME 1.1 further boosts the clustering accuracy of CLIME 1.0, where the latter draws inference based solely on a single possibly incorrect tree input.

Simulations under the tree-independent model give all four methods a more even ground. Yet still, both CLIME 1.0 and CLIME 1.1 outperformed other benchmark methods in most of the simulation settings. Specifically, CLIME 1.1 maintained its domination over all other methods, sustaining the benefit of incorporating of tree structure viability. With a distribution of possible evolutionary trees to integrate, CLIME 1.1 takes advantage of the effect of model averaging through posterior updates of tree structure, and adapts more successfully to the change of the generative model. Both CLIME 1.0 and CLIME 1.1 maintained consistency in performance across varying simulation setting, while hierarchical methods, especially the one with Hamming distance, is very sensitive to the noise level ($N_S > 0$).




FIG. 2. Simulation study results under tree-based model. Comparison of clustering accuracy (ARI) between CLIME 1.0 (black solid line), CLIME 1.1 (red dash), hierarchical clustering by Hamming distance (green dot), and hierarchical clustering by anti-correlation (blue dot-dash). N_L : number of tree branches for each ECM to have nonzero loss probability. P_L : loss probability for the N_L branches. N_s : number of singleton ECMs for each dataset.

8. Application to real data. We next apply both CLIME 1.0 and CLIME 1.1 to several real datasets, including two selected gene sets (mitochondrial complex I and proteinaceous extracellular matrix), as well as 409 manually curated gene sets from OMIM (Online Mendelian Inheritance in Man) (Hamosh et al. (2005)), where each gene set consists of genes known to be associated to a specific genetic disease. We show that CLIME 1.0 and CLIME 1.1 enjoy advantages in clustering accuracy over existing methods. Furthermore, the results of clustering and expansion analysis by CLIME 1.0 and CLIME 1.1 on these gene sets agree with estab-




FIG. 3. Simulation study results under tree-independent model. Comparison of clustering accuracy (ARI) between CLIME 1.0 (black solid line), CLIME 1.1 (red dash), hierarchical clustering by Hamming distance (green dot), and hierarchical clustering by anti-correlation (blue dot-dash). N_L : number of tree branches for each ECM to have nonzero loss probability. P_L : loss probability for the N_L branches. N_S : number of singleton ECMs for each dataset.

lished biological findings and also shed lights on potential biological discovery on gene functions and pathway compositions.

8.1. *Phylogenetic tree and matrix.* To facilitate the following analyses by CLIME 1.0, we used a single, consensus species tree that was published in Bick, Calvo and Mootha (2012) consisting of 138 diverse, sequenced eukaryotes with an additional prokaryote outgroup. For the analyses by CLIME 1.1, we used 100 posterior samples obtained by MrBayes (Ronquist and Huelsenbeck (2003))




FIG. 4. *Phylogenetic tree in use with 138 eukaryotic species (Bick, Calvo and Mootha (2012)). The tree consists of species in four different eukaryotic kingdoms (Protists, Plants, Fungi and Animals), labeled in four different colors. Human is the rightmost species on the tree.*

based on the 16 highly reserved proteins of 138 species used by Bick, Calvo and Mootha (2012). We used the phylogenetic profile matrix in Li et al. (2014) for all $N = 20,834$ human genes across the 139 species. A greater diversity of the organisms in the input tree often leads to a greater power for CLIME 1.0 and CLIME 1.1, through the increased opportunity for independent loss events.

8.2. Leave-one-out cross validation. We compared CLIME 1.0 with Hamming distance and BayesTraits (BT) (Barker and Pagel (2005), Pagel and Meade (2007)), another phylogenetic-tree-based method for gene co-evolution analysis. We conducted leave-one-out cross-validation analysis on two selected pathway/gene sets (mitochondrial complex I and proteinaceous extracellular matrix) to evaluate the clustering accuracy of the three methods. Note that we here focus on the performance of CLIME 1.0, considering the computational demands of CLIME 1.1. In Section 8.3 and 8.4, we show that CLIME 1.0 and CLIME 1.1 give relatively consistent results in real pathway-based data analysis.

For each gene set, we applied CLIME 1.0 to all but one gene within a specific pathway as test set for ECM identification and then expand the identified ECMs with the remaining human genes ($\sim 20,000$ candidate genes). We varied the LLR threshold in the expansion step of CLIME 1.0 and repeated this leave-one-out procedure for all genes in the gene set to calculate the average sensitivity and specificity of the algorithm. Note that the true positive calls (sensitivity) are made when the left-out gene is included in the expansion list and false positive calls are made when genes outside the pathway are included in the expansion list of any established ECM. For comparison, we also conducted the same experiment with the Hamming distance method (Pellegrini et al. (1999)) and BayesTraits.




FIG. 5. *Real data leave-one-out cross-validation on gene set: mitochondrial respiratory chain complex I. Comparison of ROC curves between CLIME 1.0, BayesTrait, and Hamming distance.*

BayesTraits is computationally intensive as it evaluates genetic profiles in a pairwise manner (estimated ~ 244 -hour CPU time for $44 \times 20,000$ pairwise calculation, one leave-one-out experiment for a 44-gene pathway test set; versus CLIME 1.0's ~ 2 -hour CPU time). For efficiency in computation, we only subsampled 500 genes from remaining ($\sim 20,000$) human genes as the candidate set for gene set expansion. We calculated the pairwise co-evolution p -values by BayesTraits between all genes in the leave-one-out test set and the candidate set, and made a positive call if the minimal p -value between the candidate gene and each gene in the test set is below certain threshold. Similarly, we varied the threshold to obtain the sensitivity and specificity of the algorithm.

We applied all these methods to two gene sets, mitochondrial respiratory chain complex I (44 genes), and proteinaceous extracellular matrix (194 genes) and report the receiver operating characteristic curves (ROC, true positive rate (TPR) versus false positive rate (FPR)) of all methods in Figure 5 and 6 respectively.




FIG. 6. *Real data leave-one-out cross-validation on gene set: proteinaceous extracellular matrix. Comparison of ROC curves between CLIME 1.0, BayesTrait, and Hamming distance.*




FIG. 7. Partition of 44 human Complex I genes by CLIME 1.0 and CLIME 1.1. Genes with same colored blocks are included in the same nonsingleton ECMs. Grey color indicates singleton genes.

Both CLIME 1.0 and BayesTraits dominated the Hamming-distance-based method, showing the strong advantage of incorporating the information from phylogenetic trees for functional pathway analysis based on genetic profiles. Compared with BayesTraits, CLIME 1.0 performed slightly better than BayesTraits in majority of the evaluation range of the ROC curve. Particularly, CLIME 1.0 dominated BayesTraits in both experiments when false positive rates are under 0.2%, indicating CLIME 1.0's strength in providing accurate gene clustering with controlled mis-classification errors.

8.3. *Human complex I.* We compared CLIME 1.0 and CLIME 1.1 on a set of 44 human genes encoding complex I, the largest enzyme complex of the mitochondrial respiratory chain essential for the production of ATP (Balsa et al. (2012)). CLIME 1.0 partitioned the 44 genes into five nonsingleton ECMs, and CLIME 1.1 gave nearly identical clustering (ARI: 0.962), as shown in Figure 7, except that CLIME 1.1 combines the two ECMs by CLIME 1.0 that are related to nuclear DNA encoded subunits of the alpha subcomplex (with prefix NDUF) (Mimaki et al. (2012)). Both CLIME 1.0 and CLIME 1.1 identified an ECM containing only the subunits encoded by mitochondrial DNA (ECM1: ND1, ND4, and ND5, ECM strength by CLIME 1.0: $\phi = 30.1$, CLIME 1.1: $\phi = 30.1$), and an ECM comprising solely the core components of the N module in complex I (ECM2: NDUFV1 and NDUFV2, ECM strength by CLIME 1.0: $\phi = 6.2$, CLIME 1.1: $\phi = 6.7$) (Mimaki et al. (2012)). A detailed report on the largest ECM (indexed ECM3, ECM strength by CLIME 1.0: $\phi = 5.0$, CLIME 1.1: $\phi = 5.8$) by both methods and their respective top extended gene sets (ECM3+) is shown in Table 1. ECM3 mainly contains the nuclear-DNA-encoded subunits of complex I, including all four core subunits in the module Q of complex I (marked by asterisk). Among the top extended genes in ECM3+, multiple complex I assembly factors and core subunits are identified (marked by boldface).

8.4. *Gene sets related to human genetic diseases.* We performed the analysis by CLIME 1.0 on 409 manually curated gene sets from OMIM (Online Mendelian Inheritance in Man) (Hamosh et al. (2005)), where each gene set consists of genes known to be associated with a specific genetic disease. CLIME 1.0 identified non-singleton ECMs in 52 of these 409 gene sets (check <http://www.people.fas.harvard.edu/~junliu/CLIME/> for complete results). Figure 8 shows the top 20

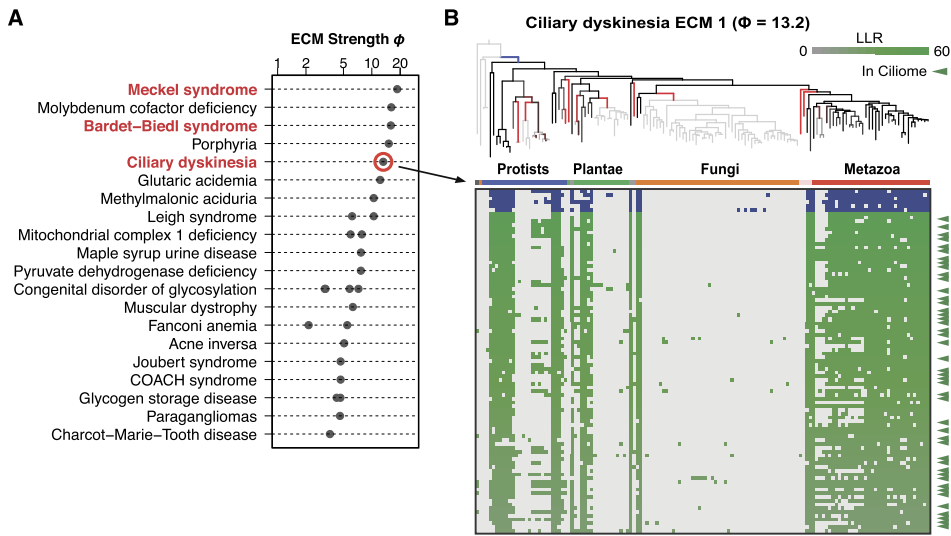


FIG. 8. (A) Top 20 OMIM gene sets with highly informative ECMs by CLIME 1.0, ranked by strength of the top ECM. All nonsingleton ECMs are shown as separate dots. Three gene sets related to human ciliary dysfunction are highlighted in red. (B) ECM 1 for ciliary dyskinesia gene set. The inferred gain/loss events are indicated by blue and red tree branches. Blue/white and green/white matrices show phylogenetic profiles of ECM and expanded genes, respectively. Green arrows indicate predicted new genes that are supported by Ciliome database.

disease-associated gene sets with the highest strength ECMs. For gene sets related to diseases such as Leigh syndrome, mitochondrial complex I deficiency, and congenital disorder of glycosylation, multiple high-strength ECMs were identified by CLIME 1.0, which suggests that functionally distinct sub-groups may exist in these gene sets. We note that among the top five gene sets, three are related to the human ciliary disease (highlighted in red). Specifically, the only nonsingleton ECM ($\phi = 13.2$) for ciliary dyskinesia, defined by having more than 15 independent loss events, is fully displayed in Figure 8(B). The expansion list contains 100 novel genes with LLR > 0. As illustrated by the heat map in Figure 8(B), all genes in the ciliary dyskinesia ECM and its expansion list share a remarkable consensus in their phylogenetic profiles. Furthermore, about 50 of the 100 expansion genes belong to the Ciliome database (Inglis, Boroevich and Leroux (2006)), an aggregation of data from seven large-scale experimental and computational studies, showing strong functional relevance of CLIME 1.0's expansion prediction.

We next compared CLIME 1.1 with CLIME 1.0 on this ciliary dyskinesia gene set. The ECM partition by CLIME 1.1 is identical to CLIME 1.0, providing a strong support of such a subgroup structure among the ciliary-dyskinesia-related genes. We further compared the extended gene sets (ECM+) obtained by CLIME 1.0 and CLIME 1.1. Among the top 100 predicted genes, 89 are shared by CLIME 1.0 and CLIME 1.1, with the top 20 reported in Table 2. The major-

TABLE 2

The nonsingleton ECM and its extension ECM+ of the ciliary dyskinesia gene set by CLIME 1.0 and CLIME 1.1. For ECM+, boldface indicates predictions for functional association with the input set; asterisk indicates direct association with ciliary dyskinesia disease by recent experimental or human genetic support

	CLIME 1.0				CLIME 1.1			
ECM	RSPH4A CCDC40	HEATR2 DNAAF2	RSPH9	CCDC39	RSPH4A CCDC40	HEATR2 DNAAF2	RSPH9	CCDC39
ECM+	RSPH6A* DRC1* MAATS1 C21orf59* CCDC146	CCDC65* SPEF1 CCDC113 SPAG16 CCDC135	RSPH3* PIBF1 CCDC147 IQUB CCDC63	C6orf165* SPATA4 ODF3 RIBC2 PIH1D3*	RSPH6A* CCDC113 SPATA4 ODF3 TTC26	CCDC65* DRC1* MAATS1 IQUB SPAG16	C6orf165* SPEF1 CCDC147 CCDC135 CEP164	RSPH3* PIH1D3* PIBF1 CCDC146 CCDC13

ity of the new members predicted by CLIME 1.0 and CLIME 1.1 can be validated as having functional association with cilia (cross-referenced by GeneCards: <https://www.genenames.org/>). In addition, the top four predicted genes have been found related to the primary ciliary dyskinesia (Horani et al. (2016)), further demonstrating the promising power of CLIME 1.0 and CLIME 1.1 in the prediction of functional relevance.

9. Discussion. Instead of integrating the pairwise co-evolution information of the genes in the input gene set in an *ad hoc* way, CLIME 1.0 explicitly models multiple genes in a functional gene set as a set of disjoint gene modules, each with its own evolutionary history. Leveraging information from multiple genes and modeling profile errors are critical because phylogenetic profiles are often noisy due to incomplete assemblies/annotations and errors in detecting distant homologs. Furthermore, CLIME 1.0 automatically infers the number of modules and gene assignments to each module. As an extension, CLIME 1.1 inherits these strengths of CLIME 1.0 and enhances its robustness and accuracy by incorporating uncertainty of evolutionary trees. CLIME 1.1, thereby, takes into account the estimation error in the tree estimation process, as well as the variability of phylogenetic relationships among genes. Simulation studies and leave-one-out cross-validations on real data showed that CLIME 1.0 achieved a significantly improved accuracy in detecting shared evolution compared with benchmark methods we tested. CLIME 1.1 further adds to CLIME 1.0 with improved robustness and precision.

Applications of CLIME 1.0 and CLIME 1.1 to real data testified the algorithms' excellent power in predicting functional association between genes and in providing guidance for further biological studies (see Li et al. (2014) for more details). Based on our exemplary pathway/gene set data, CLIME 1.0 and 1.1 showed a great consistency in identifying evolutionarily conserved subsets of genes, and demonstrated high accuracy in recovering and predicting functionally connected gene groups. CLIME 1.1 further added in with discoveries of improved robustness and relevance.

Specifically, in our investigations of the 44 complex-I-encoding genes, both CLIME 1.0 and 1.1 were able to identify subgroups of genes encoding different functional modules of complex I, and connect assembly factors associated with each module. CLIME 1.1 added to CLIME 1.0 by combining the two subgroups with nuclear DNA encoded subunits, further improving the biological interpretation of the clusters. This helps provide insights on the complete picture of complex I's catalyzing process and mechanism. We also applied our methods to more than 400 gene sets related to human genetic diseases, where CLIME 1.0 and 1.1 showed great potentials in predicting genes' functional associations with human genetic diseases. Focusing on the ciliary dyskinesia, both CLIME methods established novel connections between classic disease-driven genes and other cilia-related genes from the human genome. CLIME 1.1 furthered prediction relevance with 5% more cilia-related genes among the top predictions. Most notably, the

top four predicted genes by both CLIME methods have been validated by recent studies on primary ciliary dyskinesia. This prompts biologists with a great confidence in using CLIME as a powerful tool and in following up CLIME's findings for further experimental validations and studies on such human genetic diseases.

To trade for a gain in predication accuracy, CLIME 1.0 demands a comparatively high computational capacity. The computational complexity is about $O(Sn^2)$ per MCMC iteration in the Partition step. For CLIME 1.1, with incorporation of tree uncertainty, the step-wise computational complexity is about $O(N_T Sn^2)$. In practice, to ensure computational efficiency, we recommend implementing CLIME 1.0 firstly for a general, large-scale exploration and CLIME 1.1 for more focused, follow-up analyses and validations, as demonstrated in the Sections 8.3 and 8.4.

As shown in simulation studies, CLIME 1.0 and CLIME 1.1 gain most of its prediction power from the abundance of independent gene loss events through the evolutionary process. In fact, independent gene losses create variability of phylogenetic profiles between distinctive gene clusters, thus providing a strong signal for CLIME 1.0 and 1.1 to make inference on. Similarly, in real data we observe that CLIME 1.0 and 1.1's power is derived from the diversity of species genomes, as it provides us opportunity to observe more shared loss events. In recent years, the availability of completely sequenced eukaryotic genomes is dramatically increasing. With growing abundance and quality of eukaryotic genome sequences, the power of CLIME 1.0 and 1.1 will increase as evolutionarily distant species are more likely to possess abundant gene loss events, and thus stronger signals for CLIME 1.0 and 1.1 to extract.

Further improvements of the model are possible. Currently, we do not estimate q but set it as 0.01 based on our prior knowledge on the observation error rate. Though we observe that the model is robust to q , it is more statistically rigorous to estimate q from data. Furthermore, as there is variation between the quality of sequenced genomes, we can further assume that different species have different mis-observation rates with independent priors, which can be estimated through posterior updating. Admittedly, point estimates for cluster labels by MAP provide an interpretable representation of the posterior results, especially convenient for scientists to conduct follow-up analysis or experiment. We may also consider alternative representation of the posterior on the cluster assignment, for example, the co-assignment probability for genes.

The results, a C++ software implementing the proposed method, and an online analysis portal are freely available at <http://www.gene-clime.org>. The website was previously introduced in Li et al. (2014).

Acknowledgments. V. K. Mootha is an Investigator of the Howard Hughes Medical Institute.

REFERENCES

- ALDOUS, D. J. (1985). Exchangeability and related topics. In *École D'été de Probabilités de Saint-Flour, XIII—1983. Lecture Notes in Math.* **1117** 1–198. Springer, Berlin. [MR0883646](#)
- BALSA, E., MARCO, R., PERALES-CLEMENTE, E., SZKLARCZYK, R., CALVO, E., LANDÁZURI, M. O. and ENRÍQUEZ, J. A. (2012). NDUFA4 is a subunit of complex IV of the mammalian electron transport chain. *Cell Metab.* **16** 378–386.
- BARKER, D., MEADE, A. and PAGEL, M. (2007). Constrained models of evolution lead to improved prediction of functional linkage from correlated gain and loss of genes. *Bioinformatics* **23** 14–20.
- BARKER, D. and PAGEL, M. (2005). Predicting functional gene links from phylogenetic-statistical analyses of whole genomes. *PLoS Comput. Biol.* **1** e3.
- BICK, A. G., CALVO, S. E. and MOOHA, V. K. (2012). Evolutionary diversity of the mitochondrial calcium uniporter. *Science* **336** 886.
- CHEN, R. and LIU, J. S. (1996). Predictive updating methods with application to Bayesian classification. *J. Roy. Statist. Soc. Ser. B* **58** 397–415. [MR1377840](#)
- CHIB, S. (1995). Marginal likelihood from the Gibbs output. *J. Amer. Statist. Assoc.* **90** 1313–1321. [MR1379473](#)
- FERGUSON, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1** 209–230. [MR0350949](#)
- GALPERIN, M. Y. and KOONIN, E. V. (2010). From complete genome sequence to “complete” understanding? *Trends Biotechnol.* **28** 398–406.
- GELFAND, A. E. and SMITH, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409. [MR1141740](#)
- GLAZKO, G. V. and MUSHEGIAN, A. R. (2004). Detection of evolutionarily stable fragments of cellular pathways by hierarchical clustering of phyletic patterns. *Genome Biol.* **5** R32.
- GUINDON, S., DUFAYARD, J.-F., LEFORT, V., ANISIMOVA, M., HORDIJK, W. and GASCUEL, O. (2010). New algorithms and methods to estimate maximum-likelihood phylogenies: Assessing the performance of PhyML 3.0. *Syst. Biol.* **59** 307–321.
- HAMMING, R. W. (1950). Error detecting and error correcting codes. *Bell Syst. Tech. J.* **29** 147–160. [MR0035935](#)
- HAMOSH, A., SCOTT, A. F., AMBERGER, J. S., BOCCHINI, C. A. and MCKUSICK, V. A. (2005). Online Mendelian Inheritance in Man (OMIM), a knowledgebase of human genes and genetic disorders. *Nucleic Acids Res.* **33** D514–D517.
- HORANI, A., FERKOL, T. W., DUTCHER, S. K. and BRODY, S. L. (2016). Genetics and biology of primary ciliary dyskinesia. *Paediatr. Respir. Rev.* **18** 18–24.
- HUBERT, L. and ARABIE, P. (1985). Comparing partitions. *J. Classification* **2** 193–218.
- INGLIS, P. N., BOROEVIKH, K. A. and LEROUX, M. R. (2006). Piecing together a ciliome. *Trends Genet.* **22** 491–500.
- JIM, K., PARMAR, K., SINGH, M. and TAVAZOIE, S. (2004). A cross-genomic approach for systematic mapping of phenotypic traits to genes. *Genome Res.* **14** 109–115.
- KENSCHKE, P. R., VAN NOORT, V., DUTILH, B. E. and HUYNEN, M. A. (2008). Practical and theoretical advances in predicting the function of a protein by its phylogenetic distribution. *J. R. Soc. Interface* **5** 151–170.
- LI, J. B., GERDES, J. M., HAYCRAFT, C. J., FAN, Y., TESLOVICH, T. M., MAY-SIMERA, H., LI, H., BLACQUE, O. E., LI, L., LEITCH, C. C. et al. (2004). Comparative genomics identifies a flagellar and basal body proteome that includes the BBS5 human disease gene. *Cell* **117** 541–552.
- LI, Y., CALVO, S. E., GUTMAN, R., LIU, J. S. and MOOHA, V. K. (2014). Expansion of biological pathways based on evolutionary inference. *Cell* **158** 213–225.
- LIU, J. S. (1994). The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem. *J. Amer. Statist. Assoc.* **89** 958–966. [MR1294740](#)

- LIU, J. S. (2008). *Monte Carlo Strategies in Scientific Computing. Springer Series in Statistics.* Springer, New York. MR2401592
- LIU, J. S., WONG, W. H. and KONG, A. (1994). Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes. *Biometrika* **81** 27–40. MR1279653
- MIMAKI, M., WANG, X., MCKENZIE, M., THORBURN, D. R. and RYAN, M. T. (2012). Understanding mitochondrial complex I assembly in health and disease. *Biochim. Biophys. Acta, Bioenerg.* **1817** 851–862.
- NEAL, R. M. (2000). Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Statist.* **9** 249–265. MR1823804
- Ogilvie, I., Kennaway, N. G. and Shoubridge, E. A. (2005). A molecular chaperone for mitochondrial complex I assembly is mutated in a progressive encephalopathy. *J. Clin. Invest.* **115** 2784–2792.
- PAGEL, M. and MEADE, A. (2007). BayesTraits. Computer program and documentation. Available at <http://www.evolution.rdg.ac.uk/bayestraits.html>.
- PAGLIARINI, D. J., CALVO, S. E., CHANG, B., SHETH, S. A., VAFAI, S. B., ONG, S. E., WALFORD, G. A. et al. (2008). A mitochondrial protein compendium elucidates complex I disease biology. *Cell* **134** 112–123.
- PELLEGRINI, M., MARCOTTE, E. M., THOMPSON, M. J., EISENBERG, D. and YEATES, T. O. (1999). Assigning protein functions by comparative genome analysis: Protein phylogenetic profiles. *Proc. Natl. Acad. Sci. USA* **96** 4285–4288.
- PITMAN, J. (1996). Some developments of the Blackwell–MacQueen urn scheme. In *Statistics, Probability and Game Theory. Institute of Mathematical Statistics Lecture Notes—Monograph Series* **30** 245–267. IMS, Hayward, CA. MR1481784
- RONQUIST, F. and HUELSENBECK, J. P. (2003). MrBayes 3: Bayesian phylogenetic inference under mixed models. *Bioinformatics* **19** 1572–1574.
- TABACH, Y., BILLI, A. C., HAYES, G. D., NEWMAN, M. A., ZUK, O., GABEL, H., KAMATH, R., YACOBY, K., CHAPMAN, B., GARCIA, S. M. et al. (2013). Identification of small RNA pathway genes using patterns of phylogenetic conservation and divergence. *Nature* **493** 694–698.
- TRACHANA, K., LARSSON, T. A., POWELL, S., CHEN, W.-H., DOERKS, T., MULLER, J. and BORK, P. (2011). Orthology prediction methods: A quality assessment using curated protein families. *BioEssays* **33** 769–780.
- VERT, J.-P. (2002). A tree kernel to analyse phylogenetic profiles. *Bioinformatics* **18** S276–S284.
- VON MERING, C., HUYNEN, M., JAEGGI, D., SCHMIDT, S., BORK, P. and SNEL, B. (2003). STRING: A database of predicted functional associations between proteins. *Nucleic Acids Res.* **31** 258–261.
- ZHOU, Y., WANG, R., LI, L., XIA, X. and SUN, Z. (2006). Inferring functional linkages between proteins from evolutionary scenarios. *J. Mol. Biol.* **359** 1150–1159.

Y. LI
S. NING
J. S. LIU
DEPARTMENT OF STATISTICS
HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS 02138
USA
E-MAIL: yangli.stat@gmail.com
shaoyangning@fas.harvard.edu
jliu@stat.harvard.edu

S. E. CALVO
BROAD INSTITUTE
CAMBRIDGE, MASSACHUSETTS 02142
USA
E-MAIL: scalvo@broadinstitute.org

V. K. MOOTHA
HOWARD HUGHES MEDICAL INSTITUTE
DEPARTMENT OF SYSTEMS BIOLOGY
HARVARD MEDICAL SCHOOL
BOSTON, MASSACHUSETTS 02115
USA
AND
MASSACHUSETTS GENERAL HOSPITAL
BOSTON, MASSACHUSETTS 02114
USA
E-MAIL: vamsi@hms.harvard.edu