DEPP: Deep Learning Enables Extending Species Trees using Single Genes

Yueyu Jiang,^{1*} Metin Balaban,² Qiyun Zhu,³ and Siavash Mirarab¹

¹*Department of Electrical and Computer Engineering, UC San Diego, CA 92093, USA*

²*Bioinformatics and Systems Biology Graduate Program, UC San Diego, CA 92093, USA*

³*Center for Fundamental and Applied Microbiomics, Arizona State University, Tempe, AZ 85281, USA*

*Corresponding author: y5jiang@ucsd.edu

Abstract

Placing new sequences onto reference phylogenies is increasingly used for analyzing environmental samples, especially microbiomes. However, existing placement methods have a fundamental limitation: they assume that query sequences have evolved using specific models directly on the reference phylogeny. Thus, they can place single-gene data (e.g., 16S rRNA amplicons) onto their own gene tree. This practice is a proxy for a more ambitious goal: extending a (genome-wide) species tree given data from individual genes. No algorithm currently addresses this challenging problem. Here, we introduce Deep-learning Enabled Phylogenetic Placement (DEPP), an algorithm that learns to extend species trees using single genes without pre-specified models. We show that DEPP updates the multi-locus microbial tree-of-life with single genes with high accuracy. We further demonstrate that DEPP can achieve the long-standing goal of combining 16S and metagenomic data onto a single tree, enabling community structure analyses that were previously impossible and producing robust patterns.

Keywords: Phylogenetic placement, Neural networks, Deep learning, Gene tree discordance, Microbiome analyses, Metagenomics

Introduction

In recent years, phylogenetic inference has found widespread application in identifying organisms that make up a biological sample (Hebert et al., 2003; Seifert et al., 2007; Munch et al., 2008). Microbiome analyses often rely on phylogenetic analyses to identify species present in an environment sampled using the 16S rRNA gene amplicon sequencing or whole metagenome shotgun sequencing data (Handelsman, 2004; Langille et al., 2013; Sunagawa et al., 2013; Matsen, 2015; N.-p. Nguyen et al., 2014; Truong et al., 2015; Asnicar et al., 2020). Using the phylogenetic context, we can identify species even when exact matches to the reference datasets are not present. Similarly, outside microbiome analyses, identifying new and known species using (meta)barcoding and genome skimming data rely on phylogenetic analyses (Kress et al., 2009; QUICKE et al., 2012; Ballesteros and Hormiga, 2018; Bohmann et al., 2020).

In these high-throughput applications, phylogenetic placement – adding a new sequence onto an existing reference tree – is sufficient, and the more challenging *de novo* reconstruction is neither necessary nor always more accurate (Janssen et al., 2018). Phylogenetic placement has a long history of method development (Felsenstein, 1981; Desper and Gascuel, 2002; Mirarab, N. Nguyen, et al., 2012; Matsen et al., 2010; Berger et al., 2011; Barbera et al., 2019; Balaban, Sarmashghi, et al., 2020; Rabiee and Mirarab, 2020). However, these algorithms add sequences from a single gene family onto a tree showing its evolutionary history (i.e., the gene tree). Phylogenetic relationships change across the genome (Maddison, 1997; Degnan and Salter, 2005) due to processes such as horizontal gene transfer (HGT) (Ochman et al., 2000), and accounting for such discordance is a subject of much recent method development (Warnow, 2017). Given data from individual genes, it must be assumed that they have evolved on a gene tree, not the species tree. Thus, existing methods place on gene trees but use the gene tree as a proxy for the species tree, hoping that their differences are not consequential. For example, species identification using marker genes such as 16S or COI (e.g., Konstantinidis and Tiedje, 2005; Zaneveld et al., 2010) implicitly assumes that the gene tree is close enough to the species tree to allow accurate identification of the *species* by placement on the gene tree. More recently, Rabiee and Mirarab, 2020 enabled inserting a genome onto a species tree by minimizing quartet distance, but this approach requires genome-wide data.

DEPP LEARNING PHYLOGENETIC PLACEMENT

54 Users of phylogenetic placement often face a question: given query sequence data from a
55 single gene (or a handful of genes), is it possible to place the query onto the species tree (a
56 goal we name *discordant placement*). The correct position of a query on the species tree is not
57 always determinable from a single gene. Nevertheless, the gene tree is related to the species
58 tree, and gene data have *some* information about the correct placement on the species tree,
59 giving us hope that discordant placement can be achieved with sufficient if imperfect accuracy.

60 The ability to extend a species tree using single-gene data would be tremendously useful
61 in studying ecology. Microbiome analyses are split between metagenomic and 16S data, which
62 remain mostly disconnected. Accurate discordant placement methods would enable researchers
63 to add 16S samples onto species trees and treat them as if they were metagenomic samples,
64 albeit with less signal and more uncertainty. Thus, it would become possible to combine 16S and
65 metagenomic data (Fig. 1) in downstream analyses such as sample differentiation (Matsen IV et
66 al., 2013) using methods such as UniFrac (Lozupone and Knight, 2005) and taxonomic profiling.
67 In particular, the ability to add *all* types of data, including 16S, metagenome-assembled genomes
68 (MAGs), and marker genes on the same backbone tree using the same algorithm will be helpful
69 in ensuring consistency for downstream applications. For example, we often have a large
70 number of 16S samples and a smaller number of metagenomic samples (presumably with better
71 differentiation ability) available to study the impact of microbiome on a phenotype. Inserting
72 all samples onto the same tree will enable probing the associations between composition of
73 samples and phenotypes using a unified analysis. Similarly, the ability to place eukaryotic
74 species sampled only through their barcode genes (e.g., COI) on species trees will greatly
75 benefit ecology research.

76 How can we approach the discordant placement problem? One solution is to misuse the
77 existing methods and simply assume the species tree and the gene tree are the same. However,
78 it is unclear how these methods are expected to behave when the sequence data has not directly
79 evolved on the tree. The standard phylogenetics methods are model-based: they assume data
80 has evolved on a tree according to a model of sequence evolution (e.g., GTR Tavaré, 1986)
81 and infer the tree using maximum likelihood, Bayesian sampling, or model-corrected distances
82 (Warnow, 2017). While this paradigm has enjoyed much success, its reliance on explicit models

can make it less accurate in the face of deviations from assumptions of the model (Jermin et al., 2020; Naser-Khdour et al., 2019; Lartillot et al., 2007; Sullivan and Swofford, 2001; Halpern and Bruno, 1998). Using existing methods for the discordant placement problem would lead to a misspecified model, which can result in low accuracy, as we show. We suggest that instead of relying on the dominant paradigm, discordant placement can be solved using machine learning.

As early as the 1990s, researchers attempted phylogenetic inference using general purpose machine learning models, such as neural networks. Dopazo and Carazo, 1997 formulated the problem as unsupervised learning and designed a neural network that reflected the tree shape. More recently, the success of deep neural networks (DNNs) in solving other challenging problems has motivated efforts to adopt DNNs in phylogenetics. Zou et al., 2020a and Suvorov et al., 2020 have used a semi-supervised approach with two steps: for every four species (a quartet), classify the input data to one of the three possible quartet topologies, then combine these quartet trees. This formulation raises a question: where are we to find labeled training data in the high volume needed by DNNs? Both papers turn to simulations for an answer: use complex models to simulate data on known trees, from which we then train the model. The point of these methods is to use complex generative models that can be sampled but do not avail themselves to scalable inference. However, learning from simulated data runs the risk of missing relevant parts of the huge parameter space and model misspecification. As Zaharias et al., 2021 recently showed, these methods can have lower accuracy than standard methods in careful benchmarking.

Phylogenetic placement offers a way to use general purpose models without simulations. Given a reference tree, however computed, and sequence data which are a function of the tree, we can use the reference data to train a machine learning model (such as a DNN) that can place a query sequence onto the reference tree. The reference tree may be a species tree inferred using large numbers of genes, using complex models, and perhaps after spending much computational resources. Such reference data are increasingly available. For example, several comprehensive trees were published in the past year on tens of thousands of microbial species (Zhu, Mai, et al., 2019; Parks et al., 2020; Asnicar et al., 2020) using 120 to 400 genes, with analyses that took > 200,000 hours of CPU and GPU time in one case. These available trees

are excellent candidates for providing the training data.

Why would we need to turn to black-box machine learning models for discordant placement? Note that discordant placement only poses a weak connection between sequence data and the tree: the sequence data should be a function of the tree. In contrast, traditional models require that sequence data have evolved *directly* on the tree according to a specific Markov model, which is a much stronger assumption. A compelling reason to use machine learning is that it provides ways to learn general functions (here from sequences to positions on trees). By avoiding explicit mechanistic models, machine learning has the potential to build general models that map sequences onto trees even when the tree and the sequences are not fully compatible, enabling placement onto species trees using single genes. Such a model would, in effect, simultaneously place sequences onto a gene tree and *reconcile* (Doyon et al., 2011) the differences with the species tree.

In this paper, we introduce the Deep-learning Enabled Phylogenetic Placement (DEPP) framework (Fig. 1). Given a reference tree, inferred in any way, and some sequence data labeled by leaves of the tree but potentially evolved on a tree incongruent with the reference tree, DEPP learns a neural network to embed sequences in a high dimensional Euclidean space, such that pairwise distances in the new space correspond to the square root of tree distances. Given such a model, the placement of new sequences can proceed by computing the embedding, computing distances, and using distance-based phylogenetic placement.

Materials & Methods

Discordant phylogenetic placement

The standard phylogenetic placement takes as input a *reference* tree T , its associated sequences S , and a query sequence q . The output is the best placement of q on T , which consists of a specific position of a particular edge of T and the length of a new terminal branch. In contrast to the standard placement, we assume the relationship between the reference tree and sequences is indirect. Thus, S (typically sequences from a single gene) does not directly evolve on T (typically a species tree) but is influenced by T . Also, T can be inferred from any source of data (not just S). We define the discordant phylogenetic placement as the problem

Figure 1: **Overview of the method.** Top: an illustration of discordant placement and the overall process. A neural network model is trained from a given reference tree, however inferred, and reference sequences labeled by the leaves of the tree. The model maps sequences to embeddings in k dimensions. Once sequences are embedded, phylogenetic placement is performed using distances between embeddings of the query and reference species. Middle: the structure of the CNN model and both training and placement processes. The training process minimizes the difference between Euclidean distances in the $\mathbb{R}^k$ space and the square root of tree distances across all pairs of species. The placement process computes Euclidean distances between embeddings of the query and references. The resulting distance vector is placed using weighted least square error minimization (APPLES). Bottom: The process for combining 16S and metagenomic samples using DEPP.

of finding the best placement of the query on T using S despite disagreements between S and T . What “best placement” depends on the context. When the discordance between data and the reference tree is because T is a species tree but the sequence data come from a single gene, we define the ideal placement as the true placement of the species on the species tree. This definition is meant to enable the application laid out before; namely, updating the species tree and identifying samples using a single gene. If this problem could be solved completely accurately, we could build species trees using single genes, and we could fully identify samples using their marker genes. Alternative definition of “best” could be imagined. For example, we could seek the place of queries that would minimize the distance between the updated tree and

the true gene tree; such definitions may be appropriate for other applications.

Background on Distance-based Estimation

Let T denote a rooted phylogenetic tree on a set of n taxa $\{t_i\}_{i=1}^n$ represented as leaves and each branch labeled with its length. The tree T defines a distance matrix where each entry $d_{ij} \in \mathbb{R}^+$ is the path length between leaves i and j . A distance matrix may or may not be equal to that of some tree, but when it does match a tree, it matches a *unique* tree and is called additive (Buneman, 1974). Assume we are given a set of sequences S , and each $s_i \in S$ corresponds to a leaf t_i . Computing distance between sequences produces a sequence distance matrix. These distances can converge to additivity if computed under the correct statistical model. For example, under the Jukes and Cantor, 1969 (JC) model, $\frac{3}{4} \ln(1 - \frac{4}{3}h)$ asymptotically converges to additivity where h is the hamming distance between sequences. Note that in discordant phylogenetic placement, traditional methods for obtaining distances from S would not match d_{ij} as the sequences have not evolved on the reference tree directly.

Placement as supervised learning

We approach discordant placement using supervised learning (Fig. 1). The training data are T and S and our model is a convolutional neural network (CNN). We use a distance-based approach and express the training data as $\{(s_i, s_j), d_{ij}\}$ where $s_i, s_j \in S$ are for pairs of reference sequences and d_{ij} is the distance between taxa i and j on the tree T . Due to the discordance, which leads to model misspecification, such using JC or similar distances may not be accurate. Instead, we use machine learning to compute distances. In addition, since real datasets almost always include missing data, we use machine learning to also reconstruct the missing parts of a sequence to obtain more accurate distances.

Learning Objective

To compute distances, we build a CNN that embeds sequences in the $\mathbb{R}^k$ space; we then use distances between embeddings as estimates of phylogenetic distances. The use of embeddings enjoys a theoretical justification. As Vienne et al., 2012 showed and Layer and Rhodes, 2017

elaborated, for any tree T , there *exists* a collection of points $P = \{\Phi(t_i)\}_{i=1}^n$ in the $\mathbb{R}^{n-1}$ Euclidean space such that the distance between the points $\Phi(t_i)$ and $\Phi(t_j)$ equals to $\sqrt{d_{ij}}$. Thus, if sequences are a function of the tree, there must exist an embedding that corresponds to the tree, and we use machine learning to find an embedding that minimizes the divergence between embedding distances and (the square root of) distances on the given reference tree.

More precisely, we treat the reference tree T and sequences S as training data and seek a model that maximizes the match of Euclidean distances between embeddings and the square root of phylogenetic distances in the reference tree (i.e., $\sqrt{d_{ij}}$). The square root is to match the theory by Vienne et al., 2012 and Layer and Rhodes, 2017. To make our goal precise, we need to define a measure of matrix similarity. While any metric can be used (and measures such as log-determinant divergence have shown promise (e.g., Xie et al., 2018)), here, we simply use mean squared error, seeking:

$$\arg \min_{\Phi} \sum_{i,j} (\|\Phi(s_i) - \Phi(s_j)\|_2 - \sqrt{d_{ij}})^2$$

where $\Phi : \{A, C, G, T\}^L \rightarrow \mathbb{R}^k$ is an embedding of sequences in the Euclidean space, and d_{ij} give pairwise path distances in the reference tree T . While we focus on nucleotides here, a similar formulation can be used for amino acid sequences or any type of character data. Note that the estimated distance of i and j is $(\|\Phi(s_i) - \Phi(s_j)\|_2)^2$.

Motivated by strong evidence in distance-based phylogenetics that weighting down long distances improves accuracy (Gascuel, 2000; Desper and Gascuel, 2002; Balaban, Sarmashghi, et al., 2020), we define a weighted version of the objective function:

$$\arg \min_{\Phi} \sum_{i,j} \frac{1}{d_{ij}^2} (\|\Phi(s_i) - \Phi(s_j)\|_2 - \sqrt{d_{ij}})^2 \quad (1)$$

which follows the Fitch and Margoliash, 1967 (FM) principle. Note that other weighting schemes (e.g., Beyer et al., 1974; Desper and Gascuel, 2002) can be used and will be tested.

Embedding size. The Layer and Rhodes, 2017 (LR) formulation requires $n - 1$ dimensions, which introduces some challenges. According to the theory, the number of dimensions needs

to increase by one after inserting the query. Our supervised learning formulation does not allow that (the embedding size is fixed after training). Thus, there is no guarantee that the embeddings remain correct after addition, even if they are before addition. However, we note that, in LR embeddings, adding a leaf would require simply dividing one of the $n - 1$ dimensions into two dimensions, leaving the rest of the embeddings intact. Thus, one can hope that having one less dimension has a minimal practical impact. More broadly, for large n , training models with n -dimensional embedding is impractical. Thus, we often set $k < n - 1$, and the gap can be more than an order of magnitude for some of our tests described below. In practice, we use a rule of thumb to select the default k (which the user can change), setting $k = 10\sqrt{n}$, rounded to the nearest power of 2 from below (i.e., $2^{\lfloor \frac{1}{2} \lfloor \log_2(100n) \rfloor \rfloor}$).

Model structure and training

CNN. We use a convolutional neural network (Fig. 1). Nucleotide sequences are encoded using 4-bit one-hot binary vectors; we refer to each bit as a channel. Gaps can be encoded as all zeros (DEPP version $\leq 0.1.13$) or by setting all four channels to $\frac{1}{4}$ ($v < 0.2.2$). Moreover, as detailed below, we can use a separate model to guess the best values to represent a gap ($v \geq 0.2.0$). To accommodate reference species that have multiple copies of a gene, we change the encoding so that instead of a binary vector, it includes the frequencies of the nucleotide characters among the gene copies. These encodings provide the input “features” that are processed through three linear convolutional layers, each followed by a nonlinear layer and a fully-connected linear layer.

As detailed in Appendix C, a convolutional layer applies a set of parameterized kernels by convolving them across its input (i.e., using the dot product of the kernel entries and the input). Convolutional layers are usually used as feature extractors, and multiple layers are used to detect high-level abstraction from the input. Here, we use them to enable the model to go beyond the traditional i.i.d models of sequence evolution and capture k -mer signatures. The first convolutional layer takes as input L features, each encoded as four channels, and outputs an $8 \times L$ matrix by applying a kernel size of 1 (but applied to all four input channels). The next two convolutional layers each have a kernel size of 5 (i.e., operating on 5-mers). The

input of each layer is padded with zeros on both sides so that the output has the same length as the input. The input to the second convolutional layer is added to the output of the third convolutional layer, forming a residual block, which is a cornerstone of deep learning (He et al., 2016). Using residual blocks can help solve the vanishing gradient problem, which is why they are commonly used, including for sequence data analyses (Killoran et al., 2017; Zou et al., 2020b). To enable the model to capture nonlinear relations, after each convolutional layer, we use a nonlinear layer built using the continuously differentiable exponential linear unit, which has performed well in other contexts (Barron, 2017). The last layers are fully connected, taking features from outputs of convolutional layers and producing the final embeddings or the probability vectors. This layer aggregates the signal from convolutional layers. Each activation in the output is connected to all the inputs, and the output is a weighted sum of all the inputs.

Handling missing data. Multiple sequence alignments almost always include gaps, which may represent missing data or indels. While indels may represent real signal, here, just like traditional maximum likelihood phylogenetic models, we treat all gaps as missing data. As mentioned earlier, we can deal with missing data by leaving the one-hot encoding ambiguous (e.g., setting all four channels to $\frac{1}{4}$). An alternative is to amputate the missing data using. Motivated by the BERT (Devlin et al., 2018) model used extensively in NLP, DEPP includes a reconstruction neural network to guess the best encoding for missing data (implemented since v0.2.0?). The input of the model is a sequence with missing data (gaps) encoded as a one-hot vector, and the output is the reconstruction of the sequence where the sites with gaps are probability vectors inferred by the model. During training, we randomly select sites and label them as gaps to provide supervised signal for learning the reconstruction objective. This learning objective is the Kullback-Leibler divergence between the one-hot encoding of the letter and the output probability vector:

$$\arg \min_{\Theta} \sum_{i \in D_{mask}} D_{KL}(I_i || O_i) \quad (2)$$

where D_{mask} are randomly chosen masked sites, and I_i and O_i are the one-hot encoding of the letter and the output probability vector for the site i . The reconstruction model, which

consists of three convolutional layers and one fully-connected layer, is trained separately from the DEPP encoder. At the time of testing/placement, a query sequence is first run through the reconstruction model to fill in the gaps, and the output of the reconstruction network is fed into the DEPP encoder to generate embeddings.

Training. We trained the model with Eq. 1 as the loss function for DEPP encoder and Eq. 2 as the loss function for reconstruction network using the stochastic gradient descent algorithm RMSProp, which divides the gradient by a running average of its recent magnitude to speed up training (Tieleman and Hinton, 2012). The batch size is fixed to 32. We check the training loss every 50 epochs and stop the training when the value of the loss function fails to decrease in two consecutive checks. The model with the optimal objective function value is chosen.

Placement

Once the CNN model is trained, we use it to map a given query sequence q to a vector of distances $D_1 \dots D_n$. For datasets with missing data (gaps), we compute two sets of distances, $\{D_i\}_{i=1}^n$ and $\{D_i^R\}_{i=1}^n$, using the models with and without gap reconstruction, respectively. The final distances is set to the weighted sum of the distances, i.e. $(1 - \alpha)D_i + \alpha D_i^R$ where α is the proportion of the sites with gaps in the query sequences. The weighted sum is used to reduce the impact of reconstructed bases (which are guessed, as opposed to being observed) on the final distance and will be empirically tested. Given these distances, we then place q onto T using distance-based placement (Balaban, Sarmashghi, et al., 2020), which uses dynamic programming to find the placement with the minimum $\sum_{i=1}^n D_i^{-2} (D_i - d_{qi}(T))^2$ where $d_{qi}(T)$ represents the tree-based distance between the query and each taxon i (Fig. 1).

Uncertainty calculation

In principle, bootstrapping, the dominant method used in phylogenetics, can be used to estimate uncertainty around distances and thus placements. However, bootstrapping assumes i.i.d models, and convolutional networks like DEPP do not treat input as i.i.d. Moreover, bootstrapping would require retraining our model on each replicate bootstrap, which we do not afford. Instead, we measure the uncertainty using a subsampling procedure based on solid grounds

from the non parametric support estimation literature. For each query, we randomly select m sites and mask them as gaps. The masked sequences are input to the pretrained DEPP model, which produces a distance vector corresponding to the distances from the query to backbone species. We repeat this step r times and get r distance vectors $\hat{D}_1 \dots \hat{D}_r$. A correction is then applied to the r distance vectors as

$$\tilde{D}_i = \sqrt{\frac{m}{n}}(\hat{D}_i - \bar{D}) + D$$

where $\bar{D}$ is the average over $\hat{D}_i$, D is the distance vector corresponds to the sequence with no site masked, and n is the length of the sequence. We use $\tilde{D}_i$ for placement and get r placements for each query, which are then used to calculate support of the placements by counting the number of times each edge is chosen. In our experiment, we choose $r = 200$ and set $m = n/\log^{0.1}(n)$.

The $\sqrt{\frac{m}{n}}$ term adjusts for the increased variance of estimates obtained from fewer data points, and is justified as follows. For a statistically consistent estimator $\hat{\theta}_n$ of a parameter θ based on n data points, if there is some rate of convergence τ_n such that $\tau_n(\hat{\theta}_n - \theta)$ weakly converges to *some* distribution as $n \rightarrow \infty$ (assumption 2.2.1 of Politis et al., 1999), then, under forgiving conditions, the distribution of $\tau_b(\hat{\theta}_b - \theta_n)$ converges to that of $\tau_n(\hat{\theta}_n - \theta)$ as $n \rightarrow \infty$ as long as $b \rightarrow \infty$ and $b/n \rightarrow 0$ (see Theorem 2.2.1 in Politis et al., 1999). While the choice of τ_n is not obvious in general, by central theorem limit, for any estimator that can be described as the sum of independent random variables, $\tau_n = \sqrt{n}$ is the correct choice. Motivated by this observation, we set $\tau_n = \sqrt{n}$ which gives us the $\sqrt{\frac{m}{n}}$ correction term, noting that we have no proof that the rate of convergence of our estimator is proportional to $\sqrt{n}$ (or for that matter, that our estimator is consistent). We will evaluate the support values empirically.

DEPP implementation Details

We implemented DEPP using PyTorch and treeswift python packages (Moshiri, 2020) and trained the models on 2080Ti NVIDIA GPUs. The embedding size k is set to 128 for datasets with 200 taxon and 512 for larger datasets (including the real WoL dataset). Other hyperpa-

rameters are fixed to their defaults (Table S2) unless otherwise specified. DEPP is trained on the reference tree and is used to compute distances that are then fed to APPLES-II (Balaban, Jiang, et al., 2021). APPLES-II is used identically to APPLES-II+JC (see below). Branch lengths of the backbone tree provided to DEPP are reestimated using RAxML-8 (Stamatakis, 2014) under the GTR+CAT model. Given more than one gene, DEPP has two options: concatenating genes or computing a summary of distances. For each query, we can compute the distance between a query and backbone species j according to each of the N genes, obtaining $D_j^1 \dots D_j^N$ (ignoring missing genes). We summarize all D_j^i values by setting D_j to the average of all D_j^i values that fall between 25 and 75 percentiles of all D_j^i values (to remove the impact of outlier genes).

Methods compared

We compared DEPP v.0.2.2 (unless otherwise specified) to three methods.

EPA-ng. (Barbera et al., 2019) This maximum likelihood method is widely used but is *not* designed for discordant placement. Nevertheless, we test it for placing gene data on the species tree, with branch lengths re-estimated using RAxML-ng, under the GTR + Γ model.

INSTRAL. (Rabiee and Mirarab, 2020) This method updates a species tree given input gene trees (already updated to include the query) and accounts for discordance by maximizing the quartet score. Here, input gene trees are inferred using FastTree-II (Price et al., 2010) or RAxML. Results with FastTree or RAxML give similar results (Fig. S16), so in the main paper we use the results with input tree built by FastTree-II. While INSTRAL does account for discordance, it requires *at least two* genes and is designed for cases with many genes.

APPLES+JC. We use APPLES in its default settings where it computes distances using Jukes Cantor (JC) model (Jukes and Cantor, 1969), chosen because Balaban, Sarmashghi, et al., 2020 found no evidence that more complex models improve accuracy. The branch lengths of reference trees are re-estimated based on the JC model using RAxML-ng (Kozlov et al., 2019). APPLES also includes two options d_m and b_m to set $w_{qi} = 0$ (i.e., ignore distances) for at most $n - b_m$ reference taxa per query when $\delta_{qi} > d_m$. We use $d_m = 0, b_m = 5$ for all datasets.

Datasets

Simulated dataset. For studying the gene tree and species tree discordance due to incomplete lineage sorting (ILS), we use a published simulated dataset (Mirarab and Warnow, 2015), with 200 ingroup species and gene trees that are discordant with the species tree due to ILS. The dataset contains model conditions corresponding to high, medium, and low ILS, each with 50 replicates. We arbitrarily selected the first $2^0, \dots, 2^5$ genes for each replicate.

We also simulated a second dataset that consists of gene trees and species trees discordant due to horizontal gene transfer (HGT) in addition to low levels of ILS. We used Simphy 1.0 (Mallo et al., 2016) to simulate 10 replicates, each with 10,000 ingroup species and 500 genes. Species trees are simulated using 10^8 -generation birth-death process (Kendall, 1948) with birth and death rates set to 5×10^{-7} and 4.167×10^{-7} , respectively. This would lead to low levels of ILS —average normalized RF distance between species and gene trees due to ILS is 0.03. In addition to ILS, each gene goes through HGT at a rate drawn from lognormal distribution with $\mu \sim \mathcal{N}(-18, 0.4)$, $\sigma^2 = 0.75$. Simphy uses a HGT model that reduces the probability of transfer proportionally to the distance between source and recipient. The combined effect of ILS and HGT leads to 0.44 (normalized RF) median gene tree discordance. Given true gene trees simulated using Simphy, we simulate alignments with length ranging between 231 and 2054 using Indelible (Fletcher and Yang, 2009). Since training a model on 10,000 species takes considerable computational resources, we then pick five genes per each replicate. To do so, we ordered the genes by the number of species that are horizontally transferred and pick 50th, 150th, 250th, 350th, and 450th genes in that the ordered gene list to ensure our test cases include genes with different level of HGT.

Web-of-life (WoL) marker genes. Zhu, Mai, et al., 2019 built a species tree of 10,575 prokaryotic genomes using ASTRAL from 381 marker genes and computed mutation unit branch lengths using 100 sites randomly selected from each of the 381 marker genes. We categorized the marker genes into three equal-sized groups based on the rank of quartet-distance (Sand et al., 2013) between their gene tree and the species tree. For each discordance category, we selected the 10 most commonly present genes among all species (Table S3). The mean gene

tree discordance with the species tree, measured using the quartet distance, is 0.18, 0.34, 0.50 in low, medium, and high discordant groups. Gene alignments are available from Zhu, Mai, et al., 2019.

WoL rRNA genes. 16S and 5S rRNA genes were predicted using RNAmmer (Lagesen et al., 2007) and aligned using UPP (N.-p. D. Nguyen et al., 2015). In genomes with multiple copies of 16/5S, we train and test DEPP on all copies and re-estimate backbone tree branch lengths for APPLES+JC using an arbitrary copy. Due to their wide usage in microbiome analyses, we also perform analyses with three regions of 16S commonly used in amplicon sequencing: V3+V4 ($\approx$ 400bp), V4 (100bp), and V4 (150bp). We removed from 16S datasets any predicted sequence output by RNAmmer that was shorter than half of the average sequence length (removing less than 1% of species; see Table S4).

Traveler’s Diarrhea microbiomes. Quality-controlled 16S rRNA gene amplicons and manually curated MAGs were derived from a study by Zhu, Dupont, et al., 2018, which identified novel pathogenic profiles from the fecal samples of 22 Traveler’s Diarrhea (TD) patients as compared with seven healthy traveler (HT) controls. The 16S rRNA amplicon sequence variants (ASVs) were generated using Deblur from QIIME 2 and are 250bp long. The 381 marker genes were identified using PhyloPhlAn on the translated protein sequences inferred by Prodigal from the contigs included in each MAG. This protocol is identical to that used in the WoL study.

Evaluation procedure and leave-out experiments

To ensure query sets (i.e., testing data) are separate from the training data, we removed 5% of species (10 and 500 for simulations with 200 and 10,000 backbone taxon and $\approx$ 445 for real data) from the species tree to obtain reference trees. We did not re-estimate species trees after removing queries. These left-out species are used as the query. The reference tree is the true species tree for the simulated dataset and the ASTRAL tree from the WoL dataset (Zhu, Mai, et al., 2019). For the simulated dataset, branch lengths of the backbone species tree are estimated using sites randomly selected from the genes we used in the experiments

(32 genes for ILS data and 5 genes for HGT data) with each gene providing 500 sites. For WoL dataset, branch lengths of the species tree are available from the original study (estimated under GTR+ Γ from 100 randomly selected sites from 381 marker genes). Training is done using DEPP v0.2.2. For testing, each query taxon is placed independently, and the result is compared against the full reference tree before pruning the query (i.e., the true tree for simulations and the ASTRAL tree for WoL). The error metric we report is the number of edges between the position on the reference tree and the inferred placement. In total, we have 8,934 and 25,000 test cases for the ILS simulated and HGT simulated data respectively and 14,616 test cases for the WoL dataset. INSTRAL fails in 66/9000 tests; we exclude these cases for all the methods.

In addition, we categorize test cases by their level of ILS, level of HGT, and phylogenetic signal. We compute the level of ILS by measuring the Robinson and Foulds, 1981 (RF) distance between true gene trees and the species tree. The phylogenetic signal is a function of many factors, including sequence length, tree height, and the rate of evolution. Here, to quantify the lack of signal, we use the RF distance between true gene trees and those estimated using FastTree-II (Price et al., 2010). These two measurements are per backbone tree. In contrast, we measure HGT levels on a per query basis by inspecting species close to the query species in the true gene tree and their placement in the species tree. Specifically, for the five nearest species in the gene tree to the query q (denote them by N_5), we compute the sum of their path length (number of branches) to q in the species tree. Note that this sum can never be less than 17, which is the value obtained if N_5 are the five closest leaves to q in the species tree and the topology is identical and balanced. We measure HGT as the average path length of N_5 above 17; i.e., $(-17 + \sum_{i \in N_5} e_{q,i})/5$ where $e_{q,i}$ is the number of branches from the query q to species i in the species tree. Thus, 0 means the query is placed in a similar context in the gene tree and thus there is no HGT *close* to that leaf, while a high value indicates that the species close to the query in the gene tree are far away from it in the true species tree, indicating HGT.

Case study on Traveler’s Diarrhea

For ASV placement, a single model is trained using DEPP v0.2.2 with the reference tree set to the WoL species tree (backbone tree) and backbone sequences coming from V4 region of

DEPP LEARNING PHYLOGENETIC PLACEMENT

414 16S (≈ 250 bp). The trained model is applied on the ASV in the Traveler’s Diarrhea dataset to
 415 calculate the distance matrix between the sequences in the studied dataset and the backbone
 416 sequences. We removed one gene (p0150), where all backbone sequences were gapped for at
 417 least half of the sites. Finally, we train a model using DEPP v.0.2.2 on the remaining sequences,
 418 giving us 380 models in total. We release and maintain these reference DEPP models for public
 419 use (see Data and Code Availability); v1.0.0 of the database is used in these analyses.

420 For placing MAGs, first, we used UPP to extend the alignments of all 380 marker genes
 421 to include the markers identified in the TD dataset. Given these alignments, we used DEPP
 422 and the 380 trained models to compute distances between query marker genes and backbone
 423 WoL species, resulting in 380 distance matrices. We then use the distance-summary option
 424 of DEPP (i.e. mean distance in the interquartile range) to summarize all distances and use
 425 APPLES-II for placement.

426 To compare pairs of samples, we use weighted UniFrac (Lozupone and Knight, 2005; Mc-
 427 Donald et al., 2018). A feature table containing the frequency of ASVs or the number of reads
 428 matching a MAG in each sample is available. We use the feature table and the placement tree
 429 to calculate weighted UniFrac between each pair of the samples using QIIME 2 (Bolyen et al.,
 430 2019). We then use the PERMANOVA (Anderson, 2001), as implemented in QIIME 2, to
 431 compare the HT group and the TD group (the number of permutation is set to be $10^6 - 1$). To
 432 visualize the correlation between samples, we apply Principal Coordinates Analysis (PCoA)
 433 on the weighted UniFrac distance matrix using QIIME-2 (P. Legendre and L. Legendre, 2012;
 434 Halko et al., 2010), picking the top 3 coordinates for visualization.

435 We calculated a *MAG coverage* metric for each sample to represent the proportion of se-
 436 quencing data covered by a MAG. It equals $(\sum_{i \in M} L_i C_i) / (\sum_{i \in N} L_i C_i)$, where M is the set
 437 of contigs constituting MAGs, N is the set of all contigs in a metagenome assembly, L is the
 438 length (bp) of a contig, C is the coverage of a contig as determined by the average number of
 439 times each nucleotide of the contig is included in any sequencing read recruited to the contig.

Results

Evaluation on simulated datasets

DEPP training and parameter sensitivity

We start by evaluating DEPP on simulated datasets, testing the ability to train the CNN model in reasonable times. As the training epochs advance, the loss function (1) drops rapidly and stabilizes after around 500 epochs in a typical case (Fig. S1). Here, training, which is a one-time process for each reference tree, finished in around 20 minutes for the 200-taxon dataset and 260 minutes for 10000-taxon dataset, on a machine with one 2080Ti NVIDIA GPU and 8 CPU cores; placement of 1000 queries took 4 seconds for the 200-taxon and 30 seconds for the 10000-taxon datasets using a single CPU core.

DEPP is mostly robust to weighing schema, with all four schemes tested resulting in statistically indistinguishable performance (Table S2). Models with more parameters, i.e., deeper network or larger embedding size, tend to have better performance but also longer training time. For example, for the 200-taxon tree, the time for training a model with one residual block is around 15 minutes while this number goes up to 20 minutes when the model has 5 residual blocks. Reducing the number of residual blocks from five to one or reducing the embedding size to 32 reduce the accuracy significantly (Table S2). In our final model, we use five residual blocks for backbone tree with 200 taxon and one residual block for the rest of dataset to trade-off performance and training time. Note that while the preliminary results motivated the choice of default settings (Table S2) used in the rest of analyses, we did not select the optimal settings for this dataset and have not tested various settings on other datasets (thus, hyperparameters are not overfit to the data). The use of the gap reconstruction model dramatically improves accuracy when the query has 40% or more gaps, and the use of the weighted approach results in further reductions in error (Fig. S17).

Comparison to other methods

We now compare accuracy of DEPP to distance-based APPLES-II (Balaban, Jiang, et al., 2021) used with the standard Jukes Cantor (JC) model, maximum likelihood method EPA-ng

(Barbera et al., 2019), and the quartet-based discordant-aware method INSTRAL (Rabiee and Mirarab, 2020). Note that APPLES-JC and EPA-ng are not designed for discordant placement using a single gene, and INSTRAL is designed only for datasets with many genes (at least two but ideally many more). However, since no existing method is designed for discordant placement, we had to compare it to these existing methods.

ILS discordance. On the 200-taxon dataset, DEPP is comparable to EPA-ng and outperforms APPLES+JC and INSTRAL when given a single gene (Fig. 2ab). While DEPP and EPA-ng have similar average error rates overall, DEPP has fewer cases with very high errors (Fig. 2a). When the gene tree discordance level is medium (or low), DEPP, EPA-ng and APPLES+JC have similar average error, which is as low as 1.5 edges (or 0.9 edge) on average but DEPP has a shorter error tail (Fig. 2a). DEPP has the lowest mean error among the methods when discordance is high (3.38 edges for DEPP versus 3.50 for EPA-ng and 4.31 for APPLES+JC). For context, random placements on these species trees would give a placement error of 14 edges on average (empirically computed). Furthermore, DEPP outperforms other methods in difficult cases when the phylogenetic signal is weak (Fig. 2b).

All methods experience a sharp rise in the error when the phylogenetic signal weakens (Fig. S5a) or discordance increases (Fig. S5bc). According to an ANOVA test (Table S1), both gene tree discordance and signal have a significant impact on the placement error (p -value $< 10^{-20}$). However, these two factors combined explain only around 15% of the variance in error for DEPP and EPA-ng.

As the number of concatenated genes increases, unsurprisingly, the mean errors of all methods reduce (Fig. 2c). Computing per-gene distances and summarizing them instead of concatenating them increases accuracy under some conditions but reduces accuracy under others (Fig. S2). DEPP, EPA-ng as well as APPLES+JC are more accurate than INSTRAL for low numbers of genes (≤ 4) but not for more genes. In fact, as the number of genes increases to 32, INSTRAL starts to have the best accuracy, a result consistent with the theory as INSTRAL is statistically consistent under ILS.

We further examine the example cases where DEPP performs well or poorly. We compare

Figure 2: Results on simulated datasets. a) Empirical cumulative distribution function (ECDF) of the placement error on a single gene for high, medium, and low discordance (69% 34%, and 21% mean RF). INSTRAL needs at least 2 genes. b) Mean and standard error of placement error versus the number of genes on ILS data. c) Sensitivity to gene properties on ILS data. Error comparison between DEPP and other method using a single gene on (left) different level of true gene tree discordance (RF distance between true gene trees and the species tree) and (right) different level of gene signal missing (RF distance between true gene trees and estimated gene trees) combining all discordance levels. y-axis: the error of DEPP minus error an alternative method. White squares: mean error difference. d) Empirical cumulative distribution function (ECDF) of the placement error on HGT data. e) Error comparison between DEPP and other method on different level of HGT. HGT is measured by $\sum_{i \in N_5} e_{q,i}$, where N_5 is the 5 closest species on the gene tree and $e_{q,n}$ is the number of branches between queries and species i on species tree.

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure 3: Examples of phylogenetic distance calculation using DEPP. Phylogenetic distance on the true species tree (in units of substitution per site) versus the distance calculated by DEPP for (a) examples with high error (11 edges of error) and (b) examples with zero error. Dashed gray line: identity line. High errors correspond to long terminal and short internal branches; see Fig. S3 for trees.

the ML phylogenetic distances computed using RAxML on the true species tree versus distances computed by DEPP for low and high error cases (Fig. 3). Both high and low error cases seem to result in unbiased distances computed in the training set; however, high error examples have higher variance. The high variance can be attributed to a lack of signal: two identical sequences in the gene may belong to different parts of the species tree, a problem that the model cannot overcome. At the time of testing (e.g., for a query), distances are systematically overestimated in cases with high error, and a large range of values are estimated for pairs with equal ML distances. In some cases, DEPP assigns small distances to some reference species that have high distances to the query (Fig. 3; bottom left).

Examining example trees shows that these cases of high error tend to correspond to novel query taxa; i.e., those on long branches on sparsely sampled clades (Fig. S3). Queries with higher terminal branch lengths lead to higher error (Fig. S4); however, it appears that the shortest of branches also have a higher error, perhaps because distinguishing very similar taxa requires strong signal. Similarly, queries with a large clade as the sister tend to be more difficult for DEPP (Fig. S4). Trees that lead to high error tend to have long terminal branches and short

branches close to the root, a condition that corresponds to rapid radiations; in contrast, easy cases are those with shorter terminal branches and long branches closer to the root (Fig. S3).

HGT discordance. On the 10000-taxon dataset, which includes HGT, EPA-ng has the best overall accuracy, with DEPP coming as a close second (mean error: 2.43 and 2.80, respectively). However, these average performances mask larger differences as HGT levels change. Breaking down the dataset by the HGT level per query taxon, we observe that DEPP has slightly worse accuracy than EPA-ng for the queries with low or medium HGT level but better accuracy with high levels of HGT. DEPP has the largest advantage with the most challenging cases when gene sequences moved far away; e.g., error for DEPP is 3.5 edges better than EPA-ng and 0.9 edges better than APPLES+JC on average with the highest level of HGT.

The real WoL dataset

We then tested DEPP on the real WoL dataset using 30 marker genes, pre-selected to represent of the range of discordance among all 381 genes, in addition to 16S and 5S. We tested DEPP, EPA-ng and APPLES+JC using both novel queries (left-out) and observed queries (training data). Despite the size of the datasets, neither DEPP training nor placement was prohibitively slow. For example, on the 16S gene with 7407 species, training takes 240 minutes and uses 5GB of memory using a 2080Ti NVIDIA GPU. Placement of 800 queries on the backbone takes less than 40 seconds using a single core.

Given large backbone trees such as WoL, a query sequence will have a considerable chance of matching a species present in the reference tree. Under such conditions, DEPP has better performance over all the genes compared with EPA-ng and APPLES+JC (Fig. S6). DEPP has close to perfect accuracy for all genes except the multi-copy genes 16S and 5S (Figs. 4a and S6). For these genes, the error slightly increases; DEPP finds the optimal placement in 91% (16S) and 73% (5S) of cases but, on average, has an error of 0.32 (16S) and 1.13 (5S) edges. Patterns of error do not change whether we train DEPP using the weighting scheme described before or by simply selecting an arbitrary copy (Fig. S7). While APPLES+JC and EPA-ng also had low error levels for many genes, they had considerable error levels on the training

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure 4: **Results on the WoL dataset.** a) Mean and standard error of placement error for APPLES+JC, EPA-ng and DEPP, applied to novel queries and known queries, on the species tree from WoL dataset, which is treated as the ground truth. The x-axis shows genes, ordered by their quartet distance to the reference species tree (shown parenthetically). For novel queries, the placement is performed in a leave-out fashion. For Empirical Cumulative Distribution (ECDF) of the placement error, see Fig. S8. b) ECDF of placement error on the 16S gene with full length or individual regions of 16S used in amplicon sequencing. V3+V4 region is $\approx$ 400bp. Dashed lines show the mean error. c) Methods comparison on rRNA genes. We cut the x-axis at 15 for better visualization; see Fig. S11 for the full range.

datasets for eight and three genes, respectively (Fig. S6).

In the most interesting case, when the query sequences are novel (i.e., are not in the training set), both DEPP and EPA-ng greatly outperform APPLES+JC (Figs. 4a). On average, the placement error of DEPP (2.17 edges) and EPA-ng (2.15 edges) is much lower than APPLES+JC (3.34 edges). Moreover, EPA-ng and DEPP have low error about the same number of times (respectively, 89% and 88% of DEPP and EPA-ng placements have four edges or less error), DEPP is less often far away from the optimal placement. For example, on average, the maximum error of DEPP for each gene is 7 edges lower than EPA-ng; or, the placement error of EPA-ng is larger than 15 edges in 3.1% of cases compared to 2.4% for DEPP. Thus, just like the simulated dataset, DEPP has fewer cases of complete failure.

Across all 32 genes, in 85% of tests, DEPP and EPA-ng finds a placement within 3 edges of the optimal placement; for the full-length 16S gene, this value is 91% for DEPP and 87% for EPA-ng. A random placement on a tree with 10,575 leaves is, on average, 26 edges away from the optimal placement. Our results are also consistent with using 16S as a marker gene, which among the 32 genes had one of the lowest mean error rates (1.43 edges). Finally, note that similar to simulated datasets, accuracy of DEPP is a function of the accuracy of its distances. Calculated distances have very little bias and high variance where DEPP works well (Fig. S9) but high variance and bias when it works poorly (Fig. S10).

Going beyond a single gene, given 50 randomly selected genes, DEPP used with distance summary strategy (see Materials & Methods) was able to place within three branches of the optimal placement in 94% of cases, with a mean error of only 0.98 edges (Fig. S12). Accuracy slightly degrades if we concatenate genes instead of summarizing distances among them (e.g., mean error increases to 1.6 edges). Thus, DEPP can not only place using single genes, but it can also place with high accuracy given data from multiple genes.

Testing performance on rRNA genes 16S and 5S, average error for DEPP is lower than EPA-ng or APPLES+JC for both genes (Fig. 4c). Mean error for DEPP over all rRNA data is 2.48 edges compared to 2.81 for EPA-ng and 3.47 for APPLES+JC. For full-length 16S, DEPP and EPA-ng find the optimal placements about the same number of times; however DEPP is more often *close* to the optimal placements (e.g., 91% of queries are within three edges for DEPP

566 compared to 87% for EPA-ng and 80% for APPLES+JC). When given short amplicon-length
 567 sequences, the advantage of DEPP over alternatives becomes more substantial. For example,
 568 given 100bp amplicons from the V4 region, DEPP has an average error of 3.4 edges while the
 569 average error of EPA-ng and APPLES+JC are 4.12 and 4.94 edges, respectively. While 3.4
 570 edges of error may sound high, note that insertion of a 100bp read into a *species tree* is clearly a
 571 difficult task. Comparing performance on different rRNA genes, for 16S data, longer sequences
 572 give better performance for all the methods (Fig. 4). Interestingly, 5S sequences (which are
 573 ≈ 100 bp) have better accuracy than 16S V4 region sequences with the similar length, possibly
 574 indicating that 5S carries more phylogenetic signal or less discordance with the species tree.

575 **Support accuracy.** We test our method of measuring support on 30 marker genes of
 576 WoL data versus the support values generated by EPA-ng. Support values are vastly different
 577 between EPA-ng and DEPP. While EPA-ng tends to produce 100% support for a single place-
 578 ment for most queries, DEPP generate far more placements, most with low support (Fig 5a).
 579 Across all 12,727 queries, EPA-ng generates only 14,355 placements versus 163,110 produced
 580 by DEPP. DEPP estimates full support for a single placement for only 9.8% of queries whereas
 581 EPA-NG produces full support for 85% of queries. Because of its high confidence in its unique
 582 placement, no threshold of EPA-NG support can produce low levels of false positive detec-
 583 tion (Fig 5b), in contrast to DEPP, which can produce FPRs close to zero. Comparing FPR
 584 and recall, DEPP can achieve the same recall level as EPA-NG with far lower levels of FPR
 585 (Fig 5b).

586 Both EPA-NG and DEPP support values tend to be higher in distribution for correct
 587 placements than incorrect placements (Fig 5ac). However, DEPP support values show a much
 588 larger gap between correct and incorrect branches and are more predictive of accuracy compared
 589 to EPA-ng (Fig 5ac). Both methods clearly over-estimate support so that even branches with
 590 100% support are often not incorrect; for example, only 66%, 71% and 76% of the DEPP
 591 placements with support ≥ 0.9 , ≥ 0.95 and ≥ 0.99 are correct. However, the over-estimation
 592 problem is worse for EPA-ng, which gives high support in the vast majority of cases (Fig 5cd).
 593 Overall, 62% of wrong placements with EPA-ng have 100% support, compared with only 0.2%

Figure 5: Support of the placements. a) Empirical cumulative distribution function (ECDF) of the supports. b) Receiver operating characteristic (ROC) curve. TP, TN, FP, FN is defined by correct placement with support above the threshold, incorrect placement with support below the threshold, incorrect placement with support above the threshold, correct placement with support below the threshold respectively (support threshold step size: 0.5%). c) Correlation between support and the correctness of the placements (area of each dot is proportional to the number of placements it contains). d) Correlation between support and placement accuracy. x-axis: cumulative support of the top placements (placements with highest supports); y-axis: proportion of the cases when the correct placements are among the top placements. These figures are using 163110 placements for DEPP and 14355 placements for EPA-ng in total across 12727 queries on 30 marker genes of the WoL dataset.

for DEPP.

Case study on combined 16S rRNA and shotgun metagenomic data

We next studied how adding the the metagenome-assembled genomes (MAGs) and 16S rRNA amplicon sequence variants (ASVs) onto the same tree enables new analyses. We used the dataset by Zhu, Dupont, et al., 2018 with gut microbiomes from seven healthy controls (HT) and 22 patients with Traveler's Diarrhea (TD), all of whom were sampled using both 16S amplicon sequencing and metagenomics with available MAGs and ASVs. We added the ASVs and MAGs onto the same WoL backbone tree using DEPP (see Materials & Methods) obtaining two placement profiles for each subject. We compare pairs of profiles using the weighted UniFrac (Lozupone and Knight, 2005) distance.

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure 6: **Combined MAG and ASV results on the TD dataset.** a) For every sample, the UniFrac distance (y -axis) is shown between its MAG placements and ASV placements of the same sample (orange) and ASV placements of other samples (blue) versus the MAG coverage (proportion of sequencing data covered by the MAGs). For inter-sample comparisons, we show all pairs of comparisons (small dots), the mean and standard error (large sold dots and bars), and the median (empty circles). b) PCoA analyses of combined ASV and MAG placements based on weighted UniFrac Distances. Each sample is represented with two connected dots, one for MAG and one for ASV. Examples of samples where ASV and MAG have very low distances are highlighted in green. Nine samples dominated by *E. coli* are highlighted with a red border. c) Distribution of UniFrac distances among pairs of samples within HT or TD group and across the groups, using MAGs, ASVs or the two data combined. F-statistic and p-values are calculated using the PERMANOVA test. d) Distance of the placement position (geometric mean of tMRCA of the species in the closest sibling clade to the query) using ASV and MAG data; each dot represents one bin or ASV. e) The ASV and MAG placements on the WoL phylogenetic tree (grey: backbone taxa).

The UniFrac distances of the MAG profile of a sample to the ASV profiles of other samples were higher on average than its distance to the ASV profile of the same sample in all except one

case (Fig. 6a). In nine samples, the MAG profile had a lower distance to its own ASV than *any* other sample. The intra-sample distances between the ASV and MAG profiles substantially reduced as the MAGs represented a larger proportion of the sequencing data of each sample (Fig. 6a; p -value = 0.001), whereas distances across samples slightly increase (p -value = 0.02). As a result, the gap between inter- and intra-sample distances grew substantially with higher MAG coverage (Fig. 6a).

Placements on a single tree enabled us to visualize all ASV- and MAG-informed community structures using a unified Principal Coordinates Analysis (PCoA) (Fig. 6b). This results show that for some samples, ASV and MAG placements indicate extremely similar community structures while for others, there is a substantial disagreement between the two. The intra-sample distances in the PCoA plots tend to be shorter for high coverage MAGs (Fig. S13a).

MAGs and ASVs can both distinguish healthy and diseased samples (p -values: 0.019 and 0.024 using the standard PERMANOVA method), but MAGs provide a higher statistical power that implies a larger effect size (Fig. 6c). While the median ASV distances between pairs of TD samples are similar to distances between TD and HD samples, both intra-group median MAG distances are lower than inter-group MAG median. Furthermore, combining two type of data, 16S and MAG, which is enabled by DEPP, provides a large separation with much larger F statistics and increased statistical significant (p -values: 0.002) (Fig. 6c).

Despite the substantial agreement between MAG and ASV placements, there are also differences. For three samples (78, 10, 80152), the ASV/MAG agreement is low compared to the background distance levels. Zhu, Dupont, et al., 2018 characterized these samples as suffering from the co-infection of multiple Enterobacteriaceae organisms (*Escherichia*, *Enterobacter*, *Klebsiella*, and *Citrobacter*), which increased the challenge of accurately binning contigs from these closely related microbes – a possible explanation for the relatively low congruence. In addition, there are several groups that are found by MAGs or ASVs but not the other (Figs. 6e, S15). The MAG placements include a clade representing Saccharibacteria which is not represented in the ASV placements. This group of bacteria is classified under the candidate phyla radiation (CPR), which has distinct physiological and genetic characteristics from all remaining bacteria. Commonly used 16S rRNA primers have reduced sensitivity in capturing the Sac-

charibacteria group (Castelle and Banfield, 2018) that, according to Zhu, Dupont, et al., 2018, may be responsible for the disease status. On the other hand, the ASV placements include several clades under Cyanobacteria not found by MAGs. These may represent the commonly seen contamination from chloroplasts of dietary plants in 16S analyses (Di Rienzi et al., 2013), a problem that does not afflict metagenomic assembly.

Beyond the detected groups, branch lengths also reveal interesting patterns. While MAGs tend to be placed close to the tips, a majority of ASV placements are deep in the tree (Fig. 6d). These more basal placements of ASVs are consistent with the lower phylogenetic signal included in short sequences, which can lead to less specific characterizations. Shorter terminal branches highlight the advantage of MAGs versus 16S rRNA amplicons in understanding microbiome compositions.

We observed substantial correlation (r^2 : 0.57; p -value $< 10^{-5}$) between the Faith's phylogenetic (alpha) diversity computed using ASV and MAG placements (Fig. S14). These strong but imperfect correlations once again show the two sources of data capture similar but subtly different patterns. Overall, due to the more basal placements discussed earlier, alpha diversity measured using ASV tends to be higher. Several low alpha diversity cases are related to nine TD samples reported by Zhu, Dupont, et al., 2018 to be dominated by *E. coli*. The first axis in the combined PCoA analysis clearly separates these samples from others using MAGs, but the separation is less strong using ASVs, a pattern observed if PCoA analyses are performed separately for ASV and MAGs (Fig. S13bc).

Discussion

We introduced a deep learning approach for extending an existing phylogenetic tree without a need for pre-specified models of sequence evolution or gene tree discordance. Our approach learns how to add new taxa by capturing patterns in an existing reference set. Thus, it uses the backbone alignment and tree to learn a model that maps sequences onto the tree. This automatic learning of the model eliminates the need for assuming rigid models. Given a correct sequence evolution model and no discordance, we see no reason machine learning should be

more accurate than traditional phylogenetics. The power of DEPP is in its ability to learn from data without knowing the model. For example, in our simulations where the GTR+Gamma model of sequence evolution and MSC or HGT models of gene tree discordance were used, DEPP was able to match the accuracy of maximum likelihood placement (and surpassed it when discordance was high) without any prior knowledge of the underlying model.

The model misspecification in our analyses came from gene tree discordance, but other forms of model misspecification exist and should be explored in future. Moreover, other reasons for discordance (e.g., the reference tree may be the taxonomic tree or built from morphological data) can also be imagined and provide potential use cases of DEPP. Finally, beyond accuracy, saving computational effort can provide a compelling reason to use black-box methods. If an expensive model is used to infer a tree, perhaps a black-box model can learn its essential features so that the tree can be extended further without repeating the effort. Such an approach would be presumably cheaper than applying the expensive model to the entire dataset. We leave the exploration of such applications to future work.

The specific formulation that we chose, embedding sequences in high dimensional spaces, allowed us to define a loss function that can be easily optimized using back propagation. One can argue that the ideal loss function for placement would be one that evaluates the accuracy of the final placement, not the distances. Designing such loss functions would be easy enough. However, optimizing a loss function with discrete components (e.g., the placement branch) will lose the differentiability of the loss function, which is necessary for backpropagation using standard methods. Finding ways to perform backpropagation in partially differentiable spaces like phylogenetic placement is an interesting topic, which we leave to future work.

Here, we trained our model on backbone trees that ranged in size from 200 to 10,000 species. While deep learning is believed to require extremely large labeled datasets for training, we were able to train DEPP, which has a moderate number of layers, with only 200 species because we use pairwise information. Thus, with 200 species, we have $\binom{200}{2} = 19900$ observed pairwise distances for training. Nevertheless, it is reasonable to expect that as reference trees become more densely sampled, the accuracy of DEPP would increase. Moreover, recall that our embedding is much smaller than size of the tree. Our results indicate that while theory

suggests we need $n - 1$ dimensions for Euclidean embedding of trees, far fewer dimensions suffice in practice; we had to reduce k from 128 to 32 to observe substantial drops in accuracy on simulated data (Table S2). Note that LR embedding states that $n - 1$ dimensions are sufficient, but it does not state that $n - 1$ dimensions are necessary. Future work should explore recent advances in hyperbolic neural networks (Ganea et al., 2018) and hyperbolic distances (Tabaghi and Dokmanić, 2020) to overcome limitations of Euclidean distances.

Our choices of hyper-parameters such as k and the training parameters such as the stopping criterion were based on pre-selected values that were not fine-tuned on any dataset. In extensive simulations, this simple procedure was necessary due to computational reasons. When applying in practice on real data, it is possible to fine-tune all the hyperparameters using a validation set. We can first randomly select a subset of the reference species as the validation set, use these as testing data to tune the parameters for the given dataset, and then train one last time with all the data with the fine-tuned parameters. Note that such a procedure will require repeated training and can be slow, and our preliminary results (Table S2) show that the gain in accuracy obtained can be small.

Our results on the real prokaryotic dataset indicate that DEPP *can* overcome horizontal transfer to find the right species tree placement, *as long as* similar patterns of HGT are observed in the training data. Note that HGT is the main cause of gene tree discordance in prokaryotes, even for the marker genes such as 16S (Gogarten et al., 2002; Stepanauskas et al., 2020). However, despite relatively low levels of error overall, a small tail of placements far away from the optimal species placement with errors >20 edges is observed for many genes (Fig. S8). This long tail may be a signature of horizontally transferred genes pointing to a very different position on the species tree than the genome-wide position. When HGT events are observed in the training data, DEPP has a chance to learn them and account for them. However, when an HGT event is novel (not seen among training data), DEPP has no way of recovering the correct position given just that one gene. These novel HGT events are a likely source of those infrequent cases of large error.

Similar to unobserved HGT events, unobserved sequence patterns can negatively impact the accuracy of DEPP. For the cases with high errors, we observe distances being overestimated,

especially when the ML distances are low. We attribute this bias to the inability of the model to easily take advantage of novel data (e.g., changes in sites that are invariable in training data). We hope to remedy this limitation of our model in future work by more explicit modeling of unseen data, data augmentation, or changing the loss function.

The most immediate use of DEPP is in connecting 16S and metagenomics, as we demonstrated in our case study. DEPP often places 16S within three branches of the correct edge in the species tree; thus, while some errors remain, DEPP results enable combined 16S and metagenomic analyses with high accuracy. Currently, the main method used to combine results from 16S and metagenomic data is to use each data type to perform *taxonomic* identification and use the resulting classification in downstream analyses. Taxonomic classification clearly has less resolution than phylogenetic placement. DEPP (and more broadly, discordant placement) allows a phylogenetic, instead of taxonomic, approach for combining data. Once sequences from all the data types are added to the same tree, many downstream measurements, such as UniFrac distances and beta diversity, can be performed on the combined data. Our case study demonstrated that DEPP is capable of resolving real microbial community structures using either 16S amplicon or metagenomic data. Moreover, on this dataset, we observed improved ability to distinguish healthy and diseased samples by combining 16S and MAG data. Thus, combining both sources of data can reveal patterns that are relevant to the pathogenic profiles of the samples and the clinical status of the subjects. We saw remarkable levels of agreement between MAGs and ASV data in our case study, but also disagreements. While our data tended to provide evidence supporting the advantages of metagenomics over 16S amplicons, some limitations of MAGs (such as assembling similar species) were also revealed. Thus, due to pros and cons of each data type, the two sources of information are likely to remain complementary. Given 16S and metagenomic data, DEPP can add *all* samples to the same underlying tree, and this unified view of multiple data types enables downstream statistical analyses (e.g., Unifrac) to analyze both sets of samples jointly. By unifying the backbone tree, DEPP eliminates one source of bias, which is the use of separate reference libraries.

747 **Data and Code Availability.**

748 DEPP is publicly available at <https://github.com/yueyujiang/DEPP>.

749 Data from this paper are available at <https://tera-trees.com/data/depp/>.

750 **Funding.**

751 This work was supported by the National Science Foundation (NSF) grant IIS-1845967 to Y.J.
 752 and S.M., NSF NSF-1815485 to M.B. and S.M., and an Arizona State University start-up
 753 grant to Q.Z.; Computations were performed on the San Diego Supercomputer Center (SDSC)
 754 through XSEDE allocations, which is supported by the NSF [ACI- 1053575].

References

- Anderson, Marti J (2001). “A new method for non-parametric multivariate analysis of variance”. In: *Austral ecology* 26.1, pp. 32–46. DOI: <https://doi.org/10.1111/j.1442-9993.2001.01070.pp.x>.
- Asnicar, Francesco et al. (Dec. 2020). “Precise phylogenetic analysis of microbial isolates and genomes from metagenomes using PhyloPhlAn 3.0”. In: *Nature Communications* 11.1, p. 2500. DOI: [10.1038/s41467-020-16366-7](https://doi.org/10.1038/s41467-020-16366-7).
- Balaban, Metin, Yueyu Jiang, et al. (2021). “Fast and Accurate Distance-based Phylogenetic Placement using Divide and Conquer”. In: *Molecular Ecology Resources* March, pp. 1–15. DOI: [10.1111/1755-0998.13527](https://doi.org/10.1111/1755-0998.13527).
- Balaban, Metin, Shahab Sarmashghi, and Siavash Mirarab (May 2020). “APPLES: Scalable Distance-Based Phylogenetic Placement with or without Alignments”. In: *Systematic Biology* 69.3. Ed. by David Posada, pp. 566–578. DOI: [10.1093/sysbio/syz063](https://doi.org/10.1093/sysbio/syz063).
- Ballesteros, Jesús A. and Gustavo Hormiga (Apr. 2018). “Species delimitation of the North American orchard-spider *Leucauge venusta* (Walckenaer, 1841) (Araneae, Tetragnathidae)”. In: *Molecular Phylogenetics and Evolution* 121. December 2017, pp. 183–197. DOI: [10.1016/j.ympev.2018.01.002](https://doi.org/10.1016/j.ympev.2018.01.002).
- Barbera, Pierre et al. (Mar. 2019). “EPA-ng: Massively Parallel Evolutionary Placement of Genetic Sequences”. In: *Systematic Biology* 68.2. Ed. by David Posada, pp. 365–369. DOI: [10.1093/sysbio/syy054](https://doi.org/10.1093/sysbio/syy054).
- Barron, Jonathan T (2017). “Continuously differentiable exponential linear units”. In: *arXiv*, arXiv–1704.
- Berger, Simon A., Denis Krompass, and Alexandros Stamatakis (May 2011). “Performance, accuracy, and Web server for evolutionary placement of short sequence reads under maximum likelihood.” In: *Systematic biology* 60.3, pp. 291–302. DOI: [10.1093/sysbio/syr010](https://doi.org/10.1093/sysbio/syr010).
- Beyer, William A. et al. (1974). “A molecular sequence metric and evolutionary trees”. In: *Mathematical Biosciences*. DOI: [10.1016/0025-5564\(74\)90028-5](https://doi.org/10.1016/0025-5564(74)90028-5).

- 782 Bohmann, Kristine et al. (July 2020). “Beyond DNA barcoding: The unrealized potential of
783 genome skim data in sample identification”. In: *Molecular Ecology* 29.14, pp. 2521–2534.
784 DOI: [10.1111/mec.15507](https://doi.org/10.1111/mec.15507).
- 785 Bolyen, Evan et al. (2019). “Reproducible, interactive, scalable and extensible microbiome data
786 science using QIIME 2”. In: *Nature Biotechnology* 37.8, pp. 852–857. DOI: [10.1038/s41587-
787 019-0209-9](https://doi.org/10.1038/s41587-019-0209-9).
- 788 Buneman, Peter (1974). “A note on the metric properties of trees”. In: *Journal of Combinatorial
789 Theory, Series B* 17.1, pp. 48–50. DOI: [10.1016/0095-8956\(74\)90047-1](https://doi.org/10.1016/0095-8956(74)90047-1).
- 790 Castelle, Cindy J. and Jillian F. Banfield (Mar. 2018). “Major New Microbial Groups Expand
791 Diversity and Alter our Understanding of the Tree of Life”. In: *Cell* 172.6, pp. 1181–1197.
792 DOI: [10.1016/j.cell.2018.02.016](https://doi.org/10.1016/j.cell.2018.02.016).
- 793 Degnan, James H. and Laura A Salter (Jan. 2005). “Gene tree distributions under the coalescent
794 process”. In: *Evolution* 59.1, pp. 24–37. DOI: [10.1111/j.0014-3820.2005.tb00891.x](https://doi.org/10.1111/j.0014-3820.2005.tb00891.x).
- 795 Desper, Richard and Olivier Gascuel (2002). “Fast and accurate phylogeny reconstruction
796 algorithms based on the minimum-evolution principle”. In: *International Workshop on Al-
797 gorithms in Bioinformatics*. Springer, pp. 357–374.
- 798 Devlin, Jacob et al. (2018). “Bert: Pre-training of deep bidirectional transformers for language
799 understanding”. In: *arXiv preprint arXiv:1810.04805*.
- 800 Di Rienzi, Sara C et al. (Oct. 2013). “The human gut and groundwater harbor non-photosynthetic
801 bacteria belonging to a new candidate phylum sibling to Cyanobacteria”. In: *eLife* 2. DOI:
802 [10.7554/eLife.01102](https://doi.org/10.7554/eLife.01102).
- 803 Dopazo, Joaquín and José María Carazo (1997). “Phylogenetic reconstruction using an un-
804 supervised growing neural network that adopts the topology of a phylogenetic tree”. In:
805 *Journal of Molecular Evolution* 44.2, pp. 226–233. DOI: [10.1007/PL00006139](https://doi.org/10.1007/PL00006139).
- 806 Doyon, J.-P. JP et al. (Sept. 2011). “Models, algorithms and programs for phylogeny reconcil-
807 iation”. In: *Briefings in Bioinformatics* 12.5, pp. 392–400. DOI: [10.1093/bib/bbr045](https://doi.org/10.1093/bib/bbr045).
- 808 Felsenstein, Joseph (1981). “Evolutionary trees from DNA sequences: a maximum likelihood
809 approach”. In: *Journal of molecular evolution* 17.6, pp. 368–376.

- 810 Fitch, Walter M. and Emanuel Margoliash (Jan. 1967). “Construction of Phylogenetic Trees”.
811 In: *Science* 155.3760, pp. 279–284. DOI: [10.1126/science.155.3760.279](https://doi.org/10.1126/science.155.3760.279).
- 812 Fletcher, W. and Z. Yang (Aug. 2009). “INDELible: A Flexible Simulator of Biological Se-
813 quence Evolution”. In: *Molecular Biology and Evolution* 26.8, pp. 1879–1888. DOI: [10.1093/molbev/msp098](https://doi.org/10.1093/molbev/msp098).
814
- 815 Ganea, Octavian Eugen, Gary Bécigneul, and Thomas Hofmann (May 2018). “Hyperbolic
816 neural networks”. In: *Advances in Neural Information Processing Systems*. Vol. 2018-Decem,
817 pp. 5345–5355.
- 818 Gascuel, Olivier (Mar. 2000). “On the Optimization Principle in Phylogenetic Analysis and
819 the Minimum-Evolution Criterion”. In: *Molecular Biology and Evolution* 17.3, pp. 401–405.
820 DOI: [10.1093/oxfordjournals.molbev.a026319](https://doi.org/10.1093/oxfordjournals.molbev.a026319).
- 821 Gogarten, J. Peter, W. Ford Doolittle, and Jeffrey G. Lawrence (Dec. 2002). “Prokaryotic
822 Evolution in Light of Gene Transfer”. In: *Molecular Biology and Evolution* 19.12, pp. 2226–
823 2238. DOI: [10.1093/oxfordjournals.molbev.a004046](https://doi.org/10.1093/oxfordjournals.molbev.a004046).
- 824 Halko, Nathan et al. (2010). “An algorithm for the principal component analysis of large data
825 sets”. In: *arXiv e-prints*, arXiv:1007.5510. eprint: [1007.5510](https://arxiv.org/abs/1007.5510).
- 826 Halpern, Aaron L. and William J. Bruno (July 1998). “Evolutionary distances for protein-
827 coding sequences: modeling site- specific residue frequencies”. In: *Molecular Biology and*
828 *Evolution* 15.7, pp. 910–917. DOI: [10.1093/oxfordjournals.molbev.a025995](https://doi.org/10.1093/oxfordjournals.molbev.a025995).
- 829 Handelsman, Jo (Dec. 2004). “Metagenomics: application of genomics to uncultured microor-
830 ganisms.” In: *Microbiology and molecular biology reviews : MMBR* 68.4, pp. 669–85. DOI:
831 [10.1128/MMBR.68.4.669-685.2004](https://doi.org/10.1128/MMBR.68.4.669-685.2004).
- 832 He, Kaiming et al. (2016). “Deep residual learning for image recognition”. In: *Proceedings of*
833 *the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- 834 Hebert, Paul D. N. et al. (Feb. 2003). “Biological identifications through DNA barcodes”. In:
835 *Proceedings of the Royal Society of London. Series B: Biological Sciences* 270.1512, pp. 313–
836 321. DOI: [10.1098/rspb.2002.2218](https://doi.org/10.1098/rspb.2002.2218).

DEPP LEARNING PHYLOGENETIC PLACEMENT

- 837 Janssen, Stefan et al. (Apr. 2018). “Phylogenetic Placement of Exact Amplicon Sequences
838 Improves Associations with Clinical Information”. In: *mSystems* 3.3. Ed. by Nicholas Chia,
839 pp. 00021–18. DOI: [10.1128/mSystems.00021-18](https://doi.org/10.1128/mSystems.00021-18).
- 840 Jermini, Lars S, Renee A Catullo, and Barbara R Holland (June 2020). “A new phylogenetic
841 protocol: dealing with model misspecification and confirmation bias in molecular phyloge-
842 netics”. In: *NAR Genomics and Bioinformatics* 2.2. DOI: [10.1093/nargab/lqaa041](https://doi.org/10.1093/nargab/lqaa041).
- 843 Jukes, T H and C R Cantor (1969). “Evolution of protein molecules”. In: *Mammalian protein*
844 *metabolism, Vol. III (1969), pp. 21-132*. Vol. III, pp. 21–132.
- 845 Kendall, David G. (1948). “On the Generalized ”Birth-and-Death” Process”. In: *The Annals*
846 *of Mathematical Statistics* 19.1, pp. 1–15. DOI: [10.1214/aoms/1177730285](https://doi.org/10.1214/aoms/1177730285).
- 847 Killoran, Nathan et al. (2017). “Generating and designing DNA with deep generative models”.
848 In: *arXiv preprint arXiv:1712.06148*.
- 849 Konstantinidis, Konstantinos T. and James M. Tiedje (Feb. 2005). “Genomic insights that
850 advance the species definition for prokaryotes”. In: *Proceedings of the National Academy of*
851 *Sciences* 102.7, pp. 2567–2572. DOI: [10.1073/pnas.0409727102](https://doi.org/10.1073/pnas.0409727102).
- 852 Kozlov, Alexey M et al. (2019). “RAxML-NG: a fast, scalable and user-friendly tool for maxi-
853 mum likelihood phylogenetic inference”. In: *Bioinformatics* 35.21, pp. 4453–4455.
- 854 Kress, W. John et al. (Nov. 2009). “Plant DNA barcodes and a community phylogeny of
855 a tropical forest dynamics plot in Panama”. In: *Proceedings of the National Academy of*
856 *Sciences* 106.44, pp. 18621–18626. DOI: [10.1073/pnas.0909820106](https://doi.org/10.1073/pnas.0909820106).
- 857 Lagesen, Karin et al. (May 2007). “RNAmmer: consistent and rapid annotation of ribosomal
858 RNA genes”. In: *Nucleic Acids Research* 35.9, pp. 3100–3108. DOI: [10.1093/nar/gkm160](https://doi.org/10.1093/nar/gkm160).
- 859 Langille, Mgi et al. (2013). “Predictive functional profiling of microbial communities using 16S
860 rRNA marker gene sequences.” In: *Nature biotechnology* 31.9, pp. 814–21. DOI: [10.1038/
861 nbt.2676](https://doi.org/10.1038/nbt.2676).
- 862 Lartillot, Nicolas, Henner Brinkmann, and Hervé Philippe (2007). “Suppression of long-branch
863 attraction artefacts in the animal phylogeny using a site-heterogeneous model”. In: *BMC*
864 *Evolutionary Biology* 7.Suppl 1, S4. DOI: [10.1186/1471-2148-7-S1-S4](https://doi.org/10.1186/1471-2148-7-S1-S4).

- 865 Layer, Mark and John A. Rhodes (2017). “Phylogenetic trees and Euclidean embeddings”. In:
866 *Journal of Mathematical Biology* 74.1-2, pp. 99–111. DOI: [10.1007/s00285-016-1018-0](https://doi.org/10.1007/s00285-016-1018-0).
- 867 Legendre, Pierre and Louis Legendre (2012). *Numerical ecology*. Elsevier.
- 868 Lozupone, Catherine and Rob Knight (2005). “UniFrac : a New Phylogenetic Method for
869 Comparing Microbial Communities UniFrac : a New Phylogenetic Method for Comparing
870 Microbial Communities”. In: *Applied and environmental microbiology* 71.12, pp. 8228–8235.
871 DOI: [10.1128/AEM.71.12.8228](https://doi.org/10.1128/AEM.71.12.8228).
- 872 Maddison, Wayne P. (Sept. 1997). “Gene Trees in Species Trees”. In: *Systematic Biology* 46.3,
873 pp. 523–536. DOI: [10.2307/2413694](https://doi.org/10.2307/2413694).
- 874 Mallo, Diego, Leonardo De Oliveira Martins, and David Posada (Mar. 2016). “SimPhy : Phy-
875 logenomic Simulation of Gene, Locus, and Species Trees”. In: *Systematic Biology* 65.2,
876 pp. 334–344. DOI: [10.1093/sysbio/syv082](https://doi.org/10.1093/sysbio/syv082).
- 877 Matsen, Frederick A (Jan. 2015). “Phylogenetics and the Human Microbiome”. In: *Systematic*
878 *Biology* 64.1, e26–e41. DOI: [10.1093/sysbio/syu053](https://doi.org/10.1093/sysbio/syu053).
- 879 Matsen, Frederick A, Robin B Kodner, and E Virginia Armbrust (Oct. 2010). “pplacer: linear
880 time maximum-likelihood and Bayesian phylogenetic placement of sequences onto a fixed
881 reference tree.” In: *BMC bioinformatics* 11.1, p. 538. DOI: [10.1186/1471-2105-11-538](https://doi.org/10.1186/1471-2105-11-538).
- 882 Matsen IV, Frederick A. et al. (2013). “Edge Principal Components and Squash Clustering:
883 Using the Special Structure of Phylogenetic Placement Data for Sample Comparison”. In:
884 *PLoS ONE* 8.3, pp. 1–17. DOI: [10.1371/journal.pone.0056859](https://doi.org/10.1371/journal.pone.0056859).
- 885 McDonald, Daniel et al. (2018). “Striped UniFrac: enabling microbiome analysis at unprece-
886 dented scale”. In: *Nature Methods* 15 (11), pp. 847–848. DOI: [10.1038/s41592-018-0187-](https://doi.org/10.1038/s41592-018-0187-8)
887 [8](https://doi.org/10.1038/s41592-018-0187-8).
- 888 Mirarab, Siavash, Nam Nguyen, and Tandy Warnow (Dec. 2012). “SEPP: SATé-Enabled Phy-
889 logenetic Placement.” In: *Pacific Symposium On Biocomputing*. WORLD SCIENTIFIC,
890 pp. 247–58. DOI: [10.1142/9789814366496_{_}0024](https://doi.org/10.1142/9789814366496_{_}0024).
- 891 Mirarab, Siavash and Tandy Warnow (June 2015). “ASTRAL-II: Coalescent-based species tree
892 estimation with many hundreds of taxa and thousands of genes”. In: *Bioinformatics* 31.12,
893 pp. i44–i52. DOI: [10.1093/bioinformatics/btv234](https://doi.org/10.1093/bioinformatics/btv234).

DEPP LEARNING PHYLOGENETIC PLACEMENT

- 894 Moshiri, N. (Jan. 2020). “TreeSwift: A massively scalable Python tree package”. In: *SoftwareX*
895 11, p. 100436. DOI: [10.1016/j.softx.2020.100436](https://doi.org/10.1016/j.softx.2020.100436).
- 896 Munch, Kasper et al. (Dec. 2008). “Fast phylogenetic DNA barcoding”. In: *Philosophical*
897 *Transactions of the Royal Society B: Biological Sciences* 363.1512, pp. 3997–4002. DOI:
898 [10.1098/rstb.2008.0169](https://doi.org/10.1098/rstb.2008.0169).
- 899 Naser-Khdour, Suha et al. (Dec. 2019). “The Prevalence and Impact of Model Violations in
900 Phylogenetic Analysis”. In: *Genome Biology and Evolution* 11.12. Ed. by David Bryant,
901 pp. 3341–3352. DOI: [10.1093/gbe/evz193](https://doi.org/10.1093/gbe/evz193).
- 902 Nguyen, Nam-phuong et al. (Dec. 2014). “TIPP: taxonomic identification and phylogenetic pro-
903 filing”. In: *Bioinformatics* 30.24, pp. 3548–3555. DOI: [10.1093/bioinformatics/btu721](https://doi.org/10.1093/bioinformatics/btu721).
- 904 Nguyen, Nam-phuong D. et al. (Dec. 2015). “Ultra-large alignments using phylogeny-aware
905 profiles”. In: *Genome Biology* 16.1, p. 124. DOI: [10.1186/s13059-015-0688-z](https://doi.org/10.1186/s13059-015-0688-z).
- 906 Ochman, Howard, Jeffrey G Lawrence, and Eduardo A Groisman (May 2000). “Lateral gene
907 transfer and the nature of bacterial innovation”. In: *Nature* 405.6784, pp. 299–304. DOI:
908 [10.1038/35012500](https://doi.org/10.1038/35012500).
- 909 Parks, Donovan H. et al. (Sept. 2020). “A complete domain-to-species taxonomy for Bacteria
910 and Archaea”. In: *Nature Biotechnology* 38.9, pp. 1079–1086. DOI: [10.1038/s41587-020-](https://doi.org/10.1038/s41587-020-0501-8)
911 [0501-8](https://doi.org/10.1038/s41587-020-0501-8).
- 912 Politis, Dimitris N, Joseph P Romano, and Michael Wolf (1999). *Subsampling*. Springer Science
913 & Business Media.
- 914 Price, Morgan N., Paramvir S. Dehal, and Adam P. Arkin (Mar. 2010). “FastTree-2 – Approx-
915 imately Maximum-Likelihood Trees for Large Alignments”. In: *PLoS ONE* 5.3, e9490. DOI:
916 [10.1371/journal.pone.0009490](https://doi.org/10.1371/journal.pone.0009490).
- 917 QUICKE, DONALD L. J. et al. (July 2012). “Utility of the DNA barcoding gene fragment for
918 parasitic wasp phylogeny (Hymenoptera: Ichneumonoidea): data release and new measure
919 of taxonomic congruence”. In: *Molecular Ecology Resources* 12.4, pp. 676–685. DOI: [10.](https://doi.org/10.1111/j.1755-0998.2012.03143.x)
920 [1111/j.1755-0998.2012.03143.x](https://doi.org/10.1111/j.1755-0998.2012.03143.x).

- 921 Rabiee, Maryam and Siavash Mirarab (Aug. 2020). “INSTRAL: Discordance-Aware Phyloge-
922 netic Placement Using Quartet Scores”. In: *Systematic Biology* 69.2. Ed. by David Posada,
923 pp. 384–391. DOI: [10.1093/sysbio/syz045](https://doi.org/10.1093/sysbio/syz045).
- 924 Robinson, DF and LR Foulds (1981). “Comparison of phylogenetic trees”. In: *Mathematical*
925 *Biosciences* 53.1-2, pp. 131–147.
- 926 Sand, Andreas et al. (Sept. 2013). “Algorithms for Computing the Triplet and Quartet Dis-
927 tances for Binary and General Trees”. In: *Biology* 2.4, pp. 1189–1209. DOI: [10.3390 /
928 biology2041189](https://doi.org/10.3390/biology2041189).
- 929 Seifert, K. A. et al. (Mar. 2007). “Prospects for fungus identification using CO1 DNA barcodes,
930 with Penicillium as a test case”. In: *Proceedings of the National Academy of Sciences* 104.10,
931 pp. 3901–3906. DOI: [10.1073/pnas.0611691104](https://doi.org/10.1073/pnas.0611691104).
- 932 Stamatakis, Alexandros (2014). “RAxML version 8: a tool for phylogenetic analysis and post-
933 analysis of large phylogenies”. In: *Bioinformatics* 30.9, pp. 1312–1313.
- 934 Stepanauskas, Ramunas et al. (Jan. 2020). “Gene exchange networks define species-like units in
935 marine prokaryotes”. In: *bioRxiv*, p. 2020.09.10.291518. DOI: [10.1101/2020.09.10.291518](https://doi.org/10.1101/2020.09.10.291518).
- 936 Sullivan, Jack and David L. Swofford (Sept. 2001). “Should We Use Model-Based Methods
937 for Phylogenetic Inference When We Know That Assumptions About Among-Site Rate
938 Variation and Nucleotide Substitution Pattern Are Violated?” In: *Systematic Biology* 50.5.
939 Ed. by R. Olmstead, pp. 723–729. DOI: [10.1080/106351501753328848](https://doi.org/10.1080/106351501753328848).
- 940 Sunagawa, Shinichi et al. (Dec. 2013). “Metagenomic species profiling using universal phyloge-
941 netic marker genes”. In: *Nature Methods* 10.12, pp. 1196–1199. DOI: [10.1038/nmeth.2693](https://doi.org/10.1038/nmeth.2693).
- 942 Suvorov, Anton, Joshua Hochuli, and Daniel R. Schrider (Sept. 2020). “Accurate Inference of
943 Tree Topologies from Multiple Sequence Alignments Using Deep Learning”. In: *Systematic*
944 *Biology* 69.2. Ed. by Laura Kubatko, pp. 221–233. DOI: [10.1093/sysbio/syz060](https://doi.org/10.1093/sysbio/syz060).
- 945 Tabaghi, Puoya and Ivan Dokmanić (2020). “Hyperbolic Distance Matrices”. In: *Proceedings*
946 *of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*,
947 pp. 1728–1738. DOI: [10.1145/3394486.3403224](https://doi.org/10.1145/3394486.3403224).
- 948 Tavaré, Simon (1986). “Some Probabilistic and Statistical Problems in the Analysis of DNA
949 Sequences”. In: *Lectures on Mathematics in the Life Sciences* 17, pp. 57–86.

- 950 Tieleman, Tijmen and Geoffrey Hinton (2012). “Lecture 6.5-rmsprop: Divide the gradient by
951 a running average of its recent magnitude”. In: *COURSERA: Neural networks for machine*
952 *learning* 4.2, pp. 26–31.
- 953 Truong, Duy Tin et al. (Oct. 2015). “MetaPhlAn2 for enhanced metagenomic taxonomic pro-
954 filing”. In: *Nature Methods* 12.10, pp. 902–903. DOI: [10.1038/nmeth.3589](https://doi.org/10.1038/nmeth.3589).
- 955 Vienne, D. M. de, Sébastien Ollier, and Gabriela Aguilera (June 2012). “Phylo-MCOA: A
956 Fast and Efficient Method to Detect Outlier Genes and Species in Phylogenomics Using
957 Multiple Co-inertia Analysis”. In: *Molecular Biology and Evolution* 29.6, pp. 1587–1598.
958 DOI: [10.1093/molbev/msr317](https://doi.org/10.1093/molbev/msr317).
- 959 Warnow, Tandy (2017). *Computational phylogenetics: An introduction to designing methods*
960 *for phylogeny estimation*. Cambridge University Press.
- 961 Xie, Pengtao et al. (2018). *Orthogonality-Promoting Distance Metric Learning: Convex Relax-*
962 *ation and Theoretical Analysis*. arXiv: [1802.06014 \[cs.LG\]](https://arxiv.org/abs/1802.06014).
- 963 Zaharias, Paul, Martin Grosshauser, and Tandy Warnow (2021). “Re-evaluating Deep Neu-
964 ral Networks for Phylogeny Estimation: The issue of taxon sampling”. In: *RECOMB (To*
965 *appear)*.
- 966 Zaneveld, Jesse R. et al. (July 2010). “Ribosomal RNA diversity predicts genome diversity in
967 gut bacteria and their relatives”. In: *Nucleic Acids Research* 38.12, pp. 3869–3879. DOI:
968 [10.1093/nar/gkq066](https://doi.org/10.1093/nar/gkq066).
- 969 Zhu, Qiyun, Christopher L. Dupont, et al. (Dec. 2018). “Visualization-assisted binning of
970 metagenome assemblies reveals potential new pathogenic profiles in idiopathic travelers’
971 diarrhea”. In: *Microbiome* 6.1, p. 201. DOI: [10.1186/s40168-018-0579-0](https://doi.org/10.1186/s40168-018-0579-0).
- 972 Zhu, Qiyun, Uyen Mai, et al. (Dec. 2019). “Phylogenomics of 10,575 genomes reveals evolution-
973 ary proximity between domains Bacteria and Archaea”. In: *Nature Communications* 10.1,
974 p. 5477. DOI: [10.1038/s41467-019-13443-4](https://doi.org/10.1038/s41467-019-13443-4).
- 975 Zou, Zhengting et al. (May 2020a). “Deep Residual Neural Networks Resolve Quartet Molecular
976 Phylogenies”. In: *Molecular Biology and Evolution* 37.5. Ed. by Li Liu, pp. 1495–1507. DOI:
977 [10.1093/molbev/msz307](https://doi.org/10.1093/molbev/msz307).

978 Zou, Zhengting et al. (2020b). “Deep residual neural networks resolve quartet molecular phy-
979 logenies”. In: *Molecular Biology and Evolution* 37.5, pp. 1495–1507.

Supplementary Material

A Supplementary Figures

Figure S1: Train loss

Figure S2: **Mean and standard error of placement error versus the number of genes for simulated data.** Blue: Results of single DEPP model trained by concatenating genes; Orange: Results of aggregating distance matrix from multiple DEPP models (one model per gene) by taking the mean distance among the interquartile range. Note that these analyses are using the older version of DEPP (v.0.1.56), and an older version of the dataset with ultrametric species trees used as the backbone (data release, v1.1.0).

Figure S3: Examples of phylogenetic trees with high (top) and low (bottom) placement errors by DEPP. For each tree, we show the (correctly placed) query in red and show the error of DEPP on top in red.

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure S4: **Impact of query properties on error.** For each mode condition, we show the impact of the (a) length of the terminal branch of the query and (b) the number of leaves in the sibling clade of the query on the accuracy of all models when run with a single gene. The boxes correspond to levels of discordance, showing high, medium, and low from left to right.

Figure S5: **The impact of true gene tree discordance and lack of signal** left: Impact of lacking signal measured as the RF distance between true gene trees and the estimated gene trees. Middle: Impact of true gene tree discordance measured as the RF distance between true gene trees and the species tree. Right Impact of true gene tree discordance measured as the quartet score between true gene trees and the species tree, combining all discordance levels. Dots represent the mean of 10 queries for each of 148 replicate datasets.

Figure S6: The mean and standard error of errors of APPLES+JC and EPA-ng versus DEPP when tested on known sequences (i.e., those in the reference set). Note that, unlike other genes, here, 16S and 5s have multiple copies (1.9 and 3.8 copies per species, respectively), causing the spike in the error.

Figure S7: **Impact of weighted encoding on WoL 16S data.** The figure shows the comparison between training the model using all the available copies (using weighted encoding described in the text) versus selecting one arbitrary copy to represent each species. With multiple copies, sequenced are encoded by the frequency of the nucleotide character at each site. The results show that using more sequences per species has a very small impact on accuracy. Nevertheless, we use all the available sequences for training in all other experiments under the assumption that using all the sequences can reduce the variance of the DNN model. Results are based on version 1.0.0 of the DEPP WoL reference library.

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure S8: **Placement error on real microbial data in leave-out-experiments** We show the Empirical Cumulative Distribution (ECDF) of the placement error for JC and DEPP per gene case. The top two rows correspond to genes with high discordance, the next two rows medium discordance, and the next two rows low discordance.

Figure S9: **Examples of distance calculation using DEPP with errors equal to zero** x-axis: the true phylogenetic distance between the query; y-axis: distance calculated by DEPP. Each panel is a query arbitrarily chosen from the experiment corresponding to the gene indicated by the panel's title. See Table S3 for gene properties.

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure S10: **Examples of distance calculation using DEPP with errors larger than 15** x-axis: the true phylogenetic distance between the query; y-axis: distance calculated by DEPP. Each panel is a query arbitrarily chosen from the experiment corresponding to the gene indicated by the panel's tile. See Table S3 for gene properties.

Figure S11: **Complete Empirical CDF on the WoL dataset.** In this figure we show the complete ECDF figure on WoL data; a truncated version is shown Fig. 4b.

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure S12: **Results of multiple genes on WoL dataset summarized using two different ways.** We show results of using 50 randomly selected genes in a leave-out experiment. Blue: We train 50 models which gives us 50 distance matrices. We then aggregate the 50 distance matrices by taking the mean distances in the interquartile range. Orange: We concatenate the sequences from 50 genes and train a model on the concatenated sequences. Note that these analyses are using the older version of DEPP (v.0.1.56), and an older version of the dataset with ultrametric species trees used as the backbone (data release, v1.1.0).

Figure S13: **Visualization of beta diversity of Traveler's Diarrhea (TD) gut microbiomes.** We plot samples using PCoA run on weighted UniFrac distances. a) The PCoA visualization of both MAG and ASV profiles. Profiles from the same sample are linked by an edge, and the color of the edge shows the percentage of sequencing data covered by MAGs. Note that lighter edges, corresponding to samples with a higher MAG coverage, tend to have shorter lengths. b,c) The PCoA visualization of MAG profiles (b) and ASV profiles (c) separately. *E. coli*-dominated samples are more easily delineated using MAGs.

Figure S14: **Correlation of Alpha diversity measured using ASVs and MAGs on Traveler's Diarrhea dataset.** We show the alpha diversity of MAGs versus ASVs, measured using Faith's phylogenetic diversity (PD). Since this measure ignores abundance, to remove the impact of noise, we only consider the most abundant species that cumulatively cover at least 95% of the placed sequences for each sample. The dotted line shown in the figure is fitted using alpha diversity from all 29 samples.

Figure S15: ASV and MAG placement on the WoL phylogenetic tree (no backbone species included)

DEPP LEARNING PHYLOGENETIC PLACEMENT

Figure S16: **Mean and standard error of placement error versus the number of genes**
Input tree to INSTRAL is estimated using FastTree and RAXML

Figure S17: **Impact of reconstruction network** The x-axis is the proportion of gaps in the alignments of the queries; Data is from 30 marker genes in the WoL dataset

982 B Supplementary Tables

Table S1: ANOVA statistical tests of the impact of gene tree discordance (RF between true gene trees and the species tree) and lack of gene signal (RF between true gene trees and estimated gene trees), both log transformed for better linearization, on the placement error of DEPP and APPLES+JC.

Method	Variable	Df	Sum Sq	Mean Sq	F value	p-value	Explained %
DEPP	log(1 - gt. err)	1	618.71	618.71	85.00	0	4.8
	log(1- discord)	1	1325.80	1325.80	182.14	0	10.4
	Residuals	1486.0	10816.49	7.27			84.8
JC	log(1 - gt. err)	1	4252.44	4252.44	384.01	0	18.7
	log(1- discord)	1	1973.77	1973.77	178.24	0	8.7
	Residuals	1486.0	16455.38	11.07			72.6
EPA-ng	log(1 - gt. err)	1	1596.27	1596.27	140.09	0	8.0
	log(1- discord)	1	1328.98	1328.98	116.63	0	6.7
	Residuals	1486.0	16932.23	11.39			85.3
INSTRAL	log(1 - gt. err)	1	1541.83	1541.83	99.43	0	5.1
	log(1- discord)	1	5658.22	5658.22	364.92	0	18.7
	Residuals	1486.0	23040.67	15.50			76.2
Using quartet score (for true discordance)							
DEPP	log(1 - gt. err)	1	618.71	618.71	81.40	0	4.9
	log(1- discord)	1	848.12	848.12	111.59	0	6.6
	Residuals	1486.0	11294.16	7.60			88.5

Table S2: Impact of hyperparameters on the error of DEPP. We test on the single-gene high discordance 200-taxon simulated dataset. Bold: the default setup, used for the results reported in the paper. k : embedding size. To change the architecture, we change the number of resblocks. Various weighting schemes are incorporated in (1). Error: mean placement error (edges). Statistical tests of significance compare each condition to the default condition using a two-sided paired Student's t test (p -values).

	32	64	128	128	128	128	128	128	256
resblocks	5	5	1	3	5	5	5	5	5
weighting	$\frac{1}{\delta^2}$	$\frac{1}{\delta^2}$	$\frac{1}{\delta^2}$	$\frac{1}{\delta^2}$	$\frac{1}{\delta^2}$	1	$\frac{1}{\delta}$	$\frac{1}{\delta^4}$	$\frac{1}{\delta^2}$
error	3.836	3.380	3.734	3.534	3.442	3.374	3.516	3.520	3.302
p -values	0.002	0.544	0.023	0.359	N/A	0.472	0.454	0.442	0.161

Observations:

- Increasing embedding size (k) tend to reduce the error.
- More residual blocks lead to better performance.
- DEPP is robust to weighting schema. Different weighting schema do not have statistically significant difference in error.

Table S3: Properties of selected genes from the WoL dataset Zhu, Mai, et al., 2019. Length indicates the length of the gene’s protein alignment. Occupancy shows the number of species that has the each gene in their genome. QD: Quartet distance between the published gene tree and the species tree. Finally, group shows the assignment of genes to the three groups of low, medium, and high discordance based on the QD.

marker	UniProt ID	gene	length	occupancy	QD	group
p0010	C9RQN5	<i>rpoB</i>	1429	9911	0.09324	low
p0003	A0A0T5XBS3	<i>rpoC</i>	1652	10249	0.11948	low
p0067	P62663	<i>rpsC</i>	239	10295	0.16015	low
p0011	Q1IJG7	<i>alaS</i>	1098	10253	0.16170	low
p0038	K9TTL6	<i>rplB</i>	287	10208	0.16852	low
p0023	D3R080	<i>infB</i>	1155	10197	0.17352	low
p0002	B1AI94	<i>ftsH</i>	721	10243	0.22639	low
p0047	D2PUL3	<i>Kfla (0407)</i>	1113	10201	0.20781	low
p0018	B5I9V0	<i>pgk</i>	409	9969	0.23642	low
p0004	P19486	<i>tuf</i>	424	10054	0.25889	low
p0056	B9CN99	<i>rimI</i>	816	10047	0.28959	mid
p0000	B3PLT3	<i>oppF-valS</i>	1631	9985	0.29736	mid
p0006	Q7US70	<i>cysS</i>	537	9977	0.31252	mid
p0008	D5WTU4	<i>serS</i>	428	9944	0.32622	mid
p0132	A0A0P0KSS3	<i>metAP1b</i>	333	9745	0.33467	mid
p0001	A0A173SLF7	<i>fusA1</i>	880	10237	0.34126	mid
p0033	D3LW72	<i>HMPREF0889 (1213)</i>	550	10315	0.35079	mid
p0005	E1GXM3	<i>ileS</i>	1220	10187	0.37819	mid
p0075	E8X3N9	<i>AciX9 (1167)</i>	848	10135	0.38996	mid
p0020	D8PAC1	<i>thrS</i>	651	10063	0.39864	mid
p0032	Q9X3X7	<i>topA</i>	1212	10101	0.41792	high
p0024	A5F9G1	<i>epd</i>	341	9337	0.42137	high
p0069	B1GZ24	<i>fold</i>	296	9389	0.43362	high
p0258	E1WYY2	<i>tyrS</i>	423	9238	0.46740	high
p0243	D5E5H2	<i>hisRS</i>	478	9769	0.48789	high
p0074	A0A173ZQ45	<i>rhIE</i>	747	8977	0.50476	high
p0310	P75510	<i>trpS</i>	346	9368	0.53021	high
p0350	B8J115	<i>metN</i>	348	9536	0.56286	high
p0131	A0A1C6GRF3	<i>potA1</i>	525	9110	0.58579	high
p0151	I5AQ83	<i>EubceDRAFT1 (0090)</i>	1320	10070	0.59077	high

Table S4: Number of species left after filtering out sequences that is shorten than roughly half the average length.

Dataset	Unfiltered	full length	V3+V4	V4 150bp	V4 100bp
# of species	7813	7797	7730	7731	7740
Threshold	N/A	800	200	75	50

C Background on CNNs

Convolutional layers. Denote an input feature map of the intermediate layer of a model as $X \in \mathbb{R}^{L \times c_{in}}$, where L is the length of the feature map (here, sequence length), and c_{in} is the number of input channels (here, 4 for the initial layer). Also given is a filter $F \in \mathbb{R}^{k \times c_{in} \times c_{out}}$ where k is the kernel size and c_{out} is the number of output channels. The convolutional operation is defined as:

$$Y(l, t) = \sum_{i=0}^k \sum_{c=0}^{c_{in}} X(l + i, c) \times F(i, c, t)$$

One-dimensional convolutional layer use sliding windows to process the input sequence. Each site in the output is the weighted sum of the fragment in the corresponding windows plus the bias.

In a neural network, convolutional layers are usually used as feature extractors. Multiple convolutional layers are able to detect high-level abstraction from the input. One-dimensional convolutional layer use sliding windows to process the input sequence. Each site in the output is the weighted sum of the fragment in the corresponding windows plus the bias. The weights are determined by a learnable kernel.

Fully Connected Layer. In a neural network, fully connected layers are usually used as the last few layers to aggregate information and get the final output. In the fully connected layer, each activation in the output has connections with all the input. Specifically, each output dimension is a weighted sum of all the input, and the weights are optimized during training.

Nonlinear layer. Convolutional layers and fully connected layers consist of only linear operations. Introducing nonlinear layers can give the model the ability to capture nonlinear relations. While many nonlinear kernels exist, in this work, we use Continuously Differentiable Exponential Linear Units (CELU) as the nonlinear layer applied at the end of each layer.

1000 CELU has the following form:

$$\text{CELU}(x, \alpha) = \begin{cases} x, & x \leq 0, \\ \alpha(e^{\frac{x}{\alpha}} - 1), & \text{otherwise} \end{cases} \quad (\text{S1})$$

1001 **D Exact commands and versions**

1002 **Software version**

- 1003 • APPLES 2.0.0
- 1004 • DEPP 0.2.2
- 1005 • EPA-ng 0.3.8
- 1006 • INSTRAL 5.13.5
- 1007 • RAxML-ng 1.0.1
- 1008 • RAxML 8.2.12
- 1009 • UPP 4.3.10
- 1010 • Qiime2 2020.11.1
- 1011 • Prodigal v2.6.3
- 1012 • PhyloPhlAn commit 2c0e61a
- 1013 • Simphy 1.0
- 1014 • Indelible V1.03

1015 **Branch length reestimation**

- 1016 • `raxml-ng --evaluate --msa $sequence_file --tree $tree --model JC --threads 2 --blopt`
- 1017 `nr_safe`
- 1018 • `raxmlHPC-PTHREADS -s $sequence_file -w $outdir -n run -p 12345 -T 32 -m GTRCAT`
- 1019 `-g $tree`

1020 **Detecting Genes**

- 1021 • Predict open reading frames (ORFs) and translate DNA sequences into amino acid sequences
- 1022 `– prodigal -p $mode -f gff -g 11 -i $g.fna -o $g.gff -a $g.faa -d $g.ffn`
- 1023 `$g`: genome ID, `$mode`: "single" for single genomes (e.g., WoL) and "meta" for metagenomes
- 1024 (e.g., TD)
- 1025 • Identify the 400 marker genes from ORFs
- 1026 `– phylophlan.py`
- 1027 `--c_dat $tmpdir --c_in $indir`

1028 --c_out \$outdir --nproc \$cpus -u \$rundir

1029 **Gene alignment**

1030 Genes were aligned using UPP using the following commands for aligning the backbone se-
1031 quences(first command), and aligning the novel queries to the existing backbone alignments
1032 (second command).

1033 • `run_upp.py -s $input_seq -B 2000 -M -1 -T 0.33 -A 200`

1034 • `run_upp.py -s $query_seq -a $backbone_seq -t $backbone_tree -A 100 -d $outdir`

1035 **Placement**

1036 • APPLES+DEPP

1037 – Training

1038 * Simulated dataset

1039 `train.depp.py`

1040 `backbone_tree_file=$backbone_tree`

1041 `backbone_seq_file=$backbone_seq`

1042 `patience=5 lr=1e-4 embedding_size=128`

1043 * Marker genes in WoL dataset

1044 `train.depp.py`

1045 `backbone_tree_file=$backbone_tree`

1046 `backbone_seq_file=$backbone_seq`

1047 `patience=5 lr=1e-4 embedding_size=512`

1048 * 5S/16S in WoL dataset

1049 `train.depp.py`

1050 `backbone_tree_file=$backbone_tree`

1051 `backbone_seq_file=$backbone_seq`

1052 `patience=5 lr=1e-4 embedding_size=512 replicate_seq=True`

1053 * Note: for version 1.0.0 of this reference DB, we used DEPP 0.1.13, for version 1.1.0
1054 (used everywhere unless otherwise specified), we use DEPP 0.1.54 for version 1.2.0, we

DEPP LEARNING PHYLOGENETIC PLACEMENT

```

1055     use DEPP 0.2.2.

1056 - Query time

1057 * Calculating distance matrix

1058   · depp_distance.py

1059     query_seq_file=$query_seq

1060     backbone_seq_file=$backbone_seq

1061     model_path=$model_path

1062     outdir=$out_dir

1063 * Placement using APPLES

1064   · run_apples.py -d $distance_file -t $backbone_tree -o -f 0 -b 5

1065 • APPLES+JC

1066   run_apples.py -s $backbone_tree -q $query_seq -t $backbone_tree -f 0 -b 5

1067 • EPA-ng

1068 - raxml-ng

1069   --evaluate

1070   --msa $backbone_seq

1071   --tree $backbone_tree

1072   --prefix info

1073   --model GTR+G+F

1074   --threads 2 --blopt nr_safe

1075 - epa-ng

1076   --ref-msa $backbone_seq

1077   --tree $backbone_tree

1078   --query $query_seq

1079   --model $GTR_info

1080 • INSTRAL

1081 - java -Djava.library.path=./ -jar __instral.jar__

1082   -i $gene_trees

1083   -f $backbone_tree

```

1084 `--placement $query`

1085 `-o $placement_tree -C`

1086 **Weighted UniFrac**

1087 • `qiime diversity beta-phylogenetic`

1088 `--p-metric weighted_unifrac`

1089 `--i-table $feature_table`

1090 `--i-phylogeny $placement_tree`

1091 `--o-distance-matrix $unifrac_distance_matrix`

1092 **PERMANOVA**

1093 • `qiime diversity beta-group-significance`

1094 `--i-distance-matrix $unifrac_distance_matrix`

1095 `--m-metadata-file $metadata`

1096 `--o-visualization $permanova_output`

1097 `--m-metadata-column group`

1098 `--p-permutations 999999`

1099 **PCoA**

1100 • `qiime diversity pcoa`

1101 `--i-distance-matrix $unifrac_distance_matrix --o-pcoa $pcoa_matrix`

1102 D.1 HGT simulation

```
1103 • simphy -rs 10 -rl f:500 -rg 1 -sb f:0.0000005 -sd f:0.000000416666667
1104 -st f:100000000 -sl f:10000 -si f:1 -sp f:500000 -su f:4e-08 -hh f:1
1105 -hs ln:1.5,1 -hl ln:1.3692114,0.6931472 -hg ln:1.5,1 -cs 14907 -gt n:-18,0.4
1106 -lt ln:gt,0.75 -lk 1 -lb f:0 -ld f:0 -v 3
1107 -o model.10000.100000000.0.0000005.norm-lognormal -ot 0 -op 1 -od 1 > log.txt
```