



MetaPalette: a k -mer Painting Approach for Metagenomic Taxonomic Profiling and Quantification of Novel Strain Variation

David Koslicki^a, Daniel Falush^b

Mathematics Department, Oregon State University, Corvallis, Oregon, USA^a; Institute of Life Sciences, University of Swansea, Singleton Park, Swansea, United Kingdom^b

ABSTRACT Metagenomic profiling is challenging in part because of the highly uneven sampling of the tree of life by genome sequencing projects and the limitations imposed by performing phylogenetic inference at fixed taxonomic ranks. We present the algorithm MetaPalette, which uses long k -mer sizes ($k = 30, 50$) to fit a k -mer “palette” of a given sample to the k -mer palette of reference organisms. By modeling the k -mer palettes of unknown organisms, the method also gives an indication of the presence, abundance, and evolutionary relatedness of novel organisms present in the sample. The method returns a traditional, fixed-rank taxonomic profile which is shown on independently simulated data to be one of the most accurate to date. Tree figures are also returned that quantify the relatedness of novel organisms to reference sequences, and the accuracy of such figures is demonstrated on simulated spike-ins and a metagenomic soil sample. The software implementing MetaPalette is available at: <https://github.com/dkoslicki/MetaPalette>. Pretrained databases are included for *Archaea*, *Bacteria*, *Eukaryota*, and viruses.

IMPORTANCE Taxonomic profiling is a challenging first step when analyzing a metagenomic sample. This work presents a method that facilitates fine-scale characterization of the presence, abundance, and evolutionary relatedness of organisms present in a given sample but absent from the training database. We calculate a “ k -mer palette” which summarizes the information from all reads, not just those in conserved genes or containing taxon-specific markers. The compositions of palettes are easy to model, allowing rapid inference of community composition. In addition to providing strain-level information where applicable, our approach provides taxonomic profiles that are more accurate than those of competing methods.

KEYWORDS: taxonomic profiling, metagenomics, quantitative methods

Metagenomics is a developing field used to characterize the organismal composition of microbial communities in environmental or clinical samples (1). A key step in most metagenomic analyses is to identify the organisms in the sample and their relative frequencies. A wide variety of different algorithms have been developed for this purpose.

Most approaches, including the one described here, are based on relating sequenced reads to reference organism genome sequences. Conceptually, the aim of these approaches is to place the organisms in the sample on a “tree of life” that has been defined in advance. In practice, the available reference organisms are extremely unevenly scattered through the true tree of life. Many medically important branches, such as enterobacteria, are relatively well sampled, with many strains from the same species, while there are entire phyla of unculturable organisms that are unrepresented (2–5).

Received 29 February 2016 Accepted 28 March 2016 Published 7 June 2016

Citation Koslicki D, Falush D. 2016. MetaPalette: a k -mer painting approach for metagenomic taxonomic profiling and quantification of novel strain variation. *mSystems* 1(3):e00020-16. doi:10.1128/mSystems.00020-16.

Editor Janet K. Jansson, Pacific Northwest National Laboratory

Copyright © 2016 Koslicki and Falush. This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/).

Address correspondence to David Koslicki, david.koslicki@math.oregonstate.edu.

A further difficulty, both in theory and in practice, is that a fully resolved tree of life cannot be established, even from complete reference genomes. At the scale of individual species, homologous recombination scrambles variation so that a tree is not necessarily an appropriate representation of organismal relationships, while more distant phylogenetic relationships can be difficult to estimate due to the various technical challenges of reconstructing ancient evolutionary events (6–10).

Based on these practical considerations, an effective metagenomic method should both identify the closest organism or set of organisms in the reference set and also estimate the genetic difference between the closest reference(s) and the organism present in the sample. The method should work both if the closest neighbor is a distant member of the same phylum and if there are multiple strains within the species in question. Fine-scale classification is important because the detailed knowledge we have, of *E. coli*, for example, shows that organisms from the same species can have entirely different ecologies and phenotypic effects on their host (11).

Given these difficulties, a number of different approaches are taken to characterize metagenomic samples. A commonly used approach is to first place individual reads onto a tree constructed for a particular set of genes and then attempt to sum the phylogenetic information across the reads (12–15). Phylogenetic analysis of each read can be computationally challenging for large datasets, and individual reads can often only be placed inaccurately. It is challenging to appropriately represent this uncertainty in later stages of the analysis. These approaches also break down if a tree is not a good representation of relationships among organisms, e.g., within species. Furthermore, while utilizing specific genes (so-called *marker genes*) can increase computational efficiency, this approach throws away a considerable amount of information from sequences that do not align to the marker genes. As a result of these issues, these methods are typically accurate for genus level or higher classification but not for fine-scale classification.

Another approach identifies features that are characteristic of particular organisms, such as the frequency of *k*-mers (16–18). These features are used either for taxonomic binning of individual reads or in order to compute the overall composition. Depending on the *k*-mer size utilized, these methods either are suitable only for higher-level phylogenetic analysis (for small *k*-mers) or are highly dependent on the training database utilized (for larger *k*-mers). In either case, no existing method using this approach can accurately detect and classify organisms that are highly diverged from ones in the training database, and the existing methods still struggle with quantifying strain-level variation. Using longer *k*-mers allows for higher specificity, but using *k*-mers that are unique to specific taxa in the reference data set (as in references 16 and 17) ignores a great deal of information about evolutionary relatedness provided by other *k*-mers. It also makes the approach highly dependent on the specific composition of the reference data set. We argue that utilizing all *k*-mers in a reference database and multiple *k*-mer sizes allows the modeling of the *k*-mer signatures of organisms absent from a given training database.

In this paper, we present an approach based on defining a “palette” for each reference organism. Specifically, we count the number of *k*-mers found in the sample DNA that are present in each reference organism. Our approach thus uses all *k*-mers of a particular length in the reference data set, while discarding the specific information provided by matches of individual *k*-mers. This is similar in spirit to the so-called pseudoalignment approach in reference 19, except that here, we use *k*-mer counts of the entire sample, not of individual reads whose origins may be ambiguous. We model these palettes using a simple linear mixture model which includes both the reference organisms and “hypothetical organisms” of different degrees of genetic relatedness to the reference organisms. The algorithm is called MetaPalette, and the outputs of the algorithm are demonstrated in Fig. 1.

We first introduce the concept of a “common *k*-mer matrix” and demonstrate how utilizing multiple *k*-mer sizes allows accurate quantification of evolutionary relatedness. We then develop a mixture modeling procedure that utilizes this information to

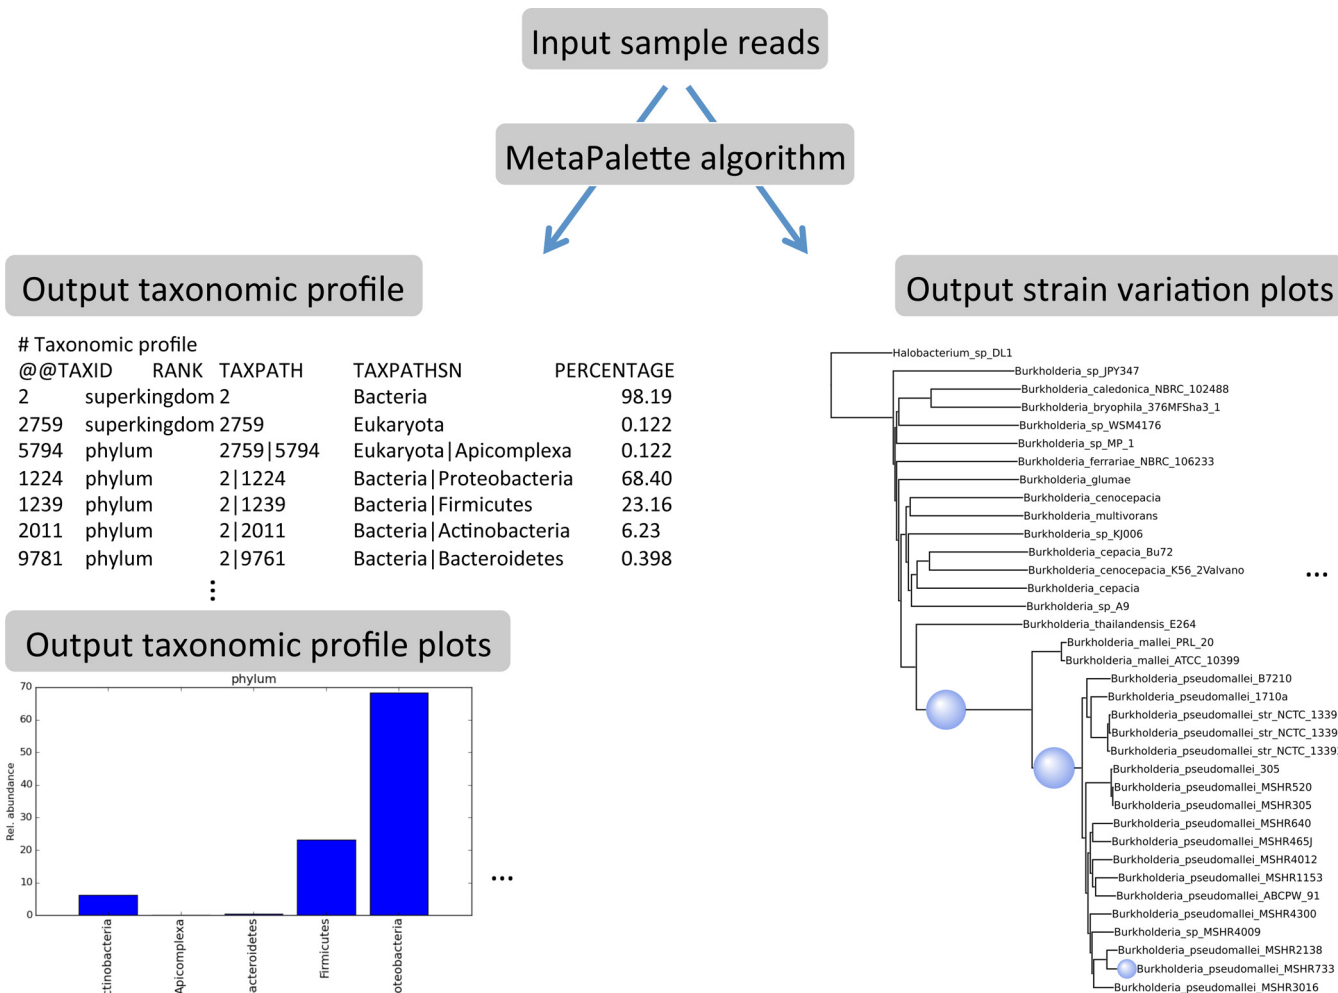


FIG 1 Illustration of the MetaPalette algorithm. Along with an output taxonomic profile and bar chart plots at all inferred taxonomic ranks, figures of strain-level variation for each inferred genus and/or species are also output.

taxonomically profile a metagenomic sample and indicate the evolutionary relatedness of novel organisms. Evidence on simulated and real data is given that demonstrates that this approach can accurately capture strain-level variation, and we then benchmark this approach against other, commonly utilized metagenomic-profiling techniques.

RESULTS

MetaPalette design. We begin by introducing the basic analytic approach. For a more detailed description, see Materials and Methods.

(i) Common *k*-mer training matrix. To quantify the similarity of two genomes, we count (with multiplicity) the fraction of each genome’s *k*-mers that are in common with the other. Rigorous mathematical definitions of this and other quantities are contained in Materials and Methods. This quantity, denoted by $pckm_k(\cdot, \cdot)$ “percentage of common *k*-mers,” is similar to the well-known Jaccard index (20) except that, among other differences, $pckm_k(\cdot, \cdot)$ is not symmetric but does incorporate the counts of *k*-mers, not just their occurrence.

When given a set of genomes (i.e., a training database), a pairwise similarity matrix can be formed using the equation $A_{i,j}^{(k)} = pckm_k(g_i, g_j)$ for g_i and g_j training genomes. The column vector $pckm_k(\cdot, g_j)$ can be thought of as a palette, representing the particular *k*-mer profile of g_j in relation to those of other genomes. We call each of these matrices a “common *k*-mer matrix.” These matrices reflect the relatedness of the training

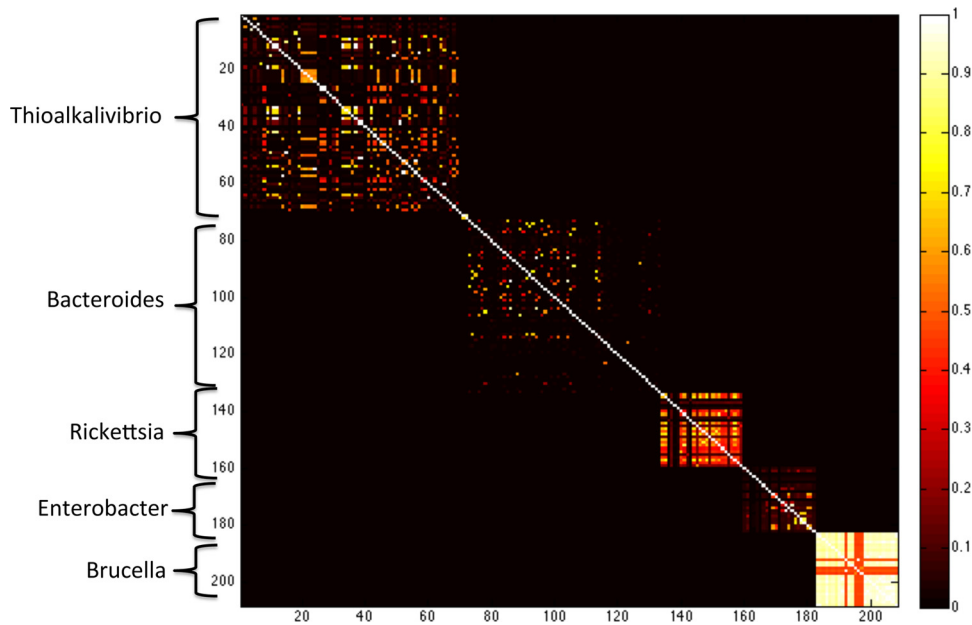


FIG 2 Heatmap of the common k -mer matrix $A^{(40)}$ for $k = 40$ using a subset of the NCBI bacterial genome database. Delineations between genera can clearly be seen. In a given genus, differing similarities of species are also visible.

genomes based on k -mer similarity. For larger k -mer sizes, one can clearly extract taxonomic information from these matrices, as shown in Fig. 2.

Beyond genus-level variation, strain-level variation can be captured through these common k -mer matrices. For example, using all the strains of the species *Burkholderia multivorans* accessible via NCBI, we formed a neighbor-joining tree using the average of the 30-mer and 50-mer common k -mer matrices. This tree, shown in Fig. 3, demonstrates how the common k -mer matrices can capture variations among these strains.




FIG 3 Neighbor-joining tree for the species *Burkholderia multivorans* based on the average of the common 30-mer and 50-mer matrices (shown in heat map to the right), depicting the ability of the common k -mer matrices to capture strain-level variation.

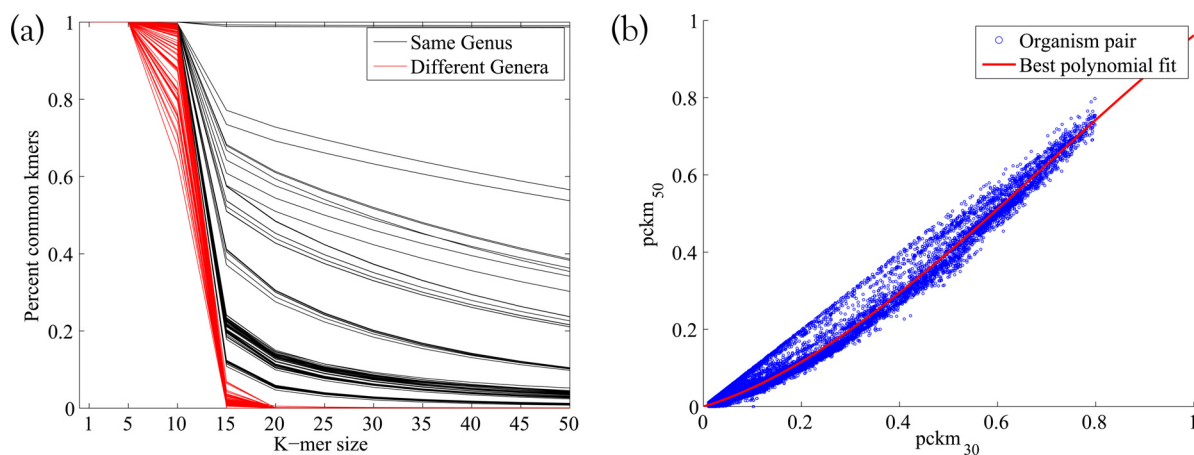


FIG 4 (a) Plot of k -mer similarity $pckm_k(g_i, g_j)$ as a function of k for 100 organism pairs of the same genus and 100 of different genera. (b) Scatterplot of the $6,914^2$ pairs of entries of the common 30-mer and 50-mer matrices shown with the best-fit polynomial.

The entries of $A^{(k)}$ can be calculated in a computationally efficient manner. We take the approach of forming bloom count filters (using Jellyfish [21]) for each of the training genomes and then counting the common k -mers using a simple C++ program based on heap data structures.

(ii) Modeling related organisms. To model the k -mer counts for organisms with various degrees of relation from the training database, we take advantage of the differing behavior of $pckm_k(\cdot, \cdot)$ as a function of k for closely related organisms and distantly related organisms. In particular, the percentage of common k -mers, $pckm_k(\cdot, \cdot)$, decays more slowly as a function of k for closely related organisms than for distantly related ones. This is consistent with the intuitive idea that, for example, two organisms from different phyla will have similar percentages of shared 1-mers but very few common 50-mers. Conversely, two closely related strains will have both high percentages of shared 1-mers and high percentages of shared 50-mers. This is demonstrated in Fig. 4a. This property means that using more than one k -mer length should in principle allow us to distinguish between having an organism that is identical to a training organism at a low frequency and having an organism that is distantly related to all training organisms but present in the sample at a higher frequency.

We focus on two particular k -mer sizes, $k = 30$ and $k = 50$, due to the predictability of $pckm_k$ for these k -mer sizes. Indeed, using 6,914 whole bacterial genomes downloaded from a variety of publicly accessible repositories (via RepoPhlAn [<https://bitbucket.org/nsegata/repophlan>]), we observed that the percentage of shared 30-mers can be predicted from the percentage of shared 50-mers (Fig. 4b). A degree 3 polynomial was used (as it resulted in the lowest root mean square error [RMSE] and R^2 values, which did not improve for higher-degree polynomials). Namely, we observed that for the polynomial $p(x) = -.5141x^3 + 1.0932x^2 + 0.3824x$, $pckm_{50}(g_i, g_j) \approx p(pckm_{30}(g_i, g_j))$.

For k -mer lengths substantially shorter than 30, the behavior of $pckm_k$ is more variable—for example, because of convergence of sequence composition between distantly related organisms. On the other hand, k -mers much larger than 50 are increasingly time consuming to compute and are likely to be more sensitive to sequencing error and other technical artifacts.

We can augment the matrices $A^{(k)}$ with columns that represent hypothetical organisms which are related by different degrees to the reference organism. For a given organism with genome g_i , if we wish to include a hypothetical organism h that is 90% similar to genome g_i in its 30-mers, we can round down each entry of the column vector $pckm_{30}(\cdot, g_i)$ to be no more than 0.90. Call this vector $pckm_{30}(\cdot, h)$. The entries below 90% do not need to be changed, since we assume that the hypothetical organism has the same patterns of k -mer sharing to more distantly related “outgroup” taxa as to the reference organism.

We model the 50-mer similarity by setting $pckm_{50}(\cdot, h) = p(pckm_{30}(\dots, g_i))$ for the previously defined polynomial p . Adding these vectors to $A^{(30)}$ and $A^{(50)}$ effectively adds a hypothetical organism that has a common k -mer signature 90% similar to that of genome g_i . We then repeat this procedure for all training genomes g_i and for similarities ranging from 90%, 80%, ..., 10% and append these columns to $A^{(30)}$ and $A^{(50)}$.

(iii) Sample k -mer signature. Given a metagenomic sample, we form two vectors, $y^{(30)}$ and $y^{(50)}$, consisting of the total counts in the sample of the 30-mers and 50-mers shared with the training organisms. In “Mathematical formulation” in Materials and Methods, we show that these vectors are linearly related to the organism abundances via the common k -mer matrices $A^{(30)}$ and $A^{(50)}$.

Note that in forming $y^{(k)}$, we count the k -mers in the entire sample, not of the individual reads. This allows for a very computationally efficient approach: as the training genomes typically have low error, their k -mers can be efficiently stored in de Bruijn graphs (formed using Bcalm [22]). We can then query the bloom count filter formed from the sample in a highly parallel fashion.

(iv) Sparsity-promoting optimization procedure. After forming $y^{(k)}$, we note that some of the entries $y_i^{(k)}$ may be nonzero due not to the presence of organism i in the sample but to the fact that there exists an organism j that shares portions of its genome with organism i . Since $A_{ij}^{(k)}$ represents the “overlap” of these two organisms, we can deconvolute this linear mixture relationship by solving the equation $A^{(k)}x = y^{(k)}$ for x , the vector of organism abundances. However, after having augmented $A^{(k)}$ with the hypothetical organisms, this system of equations is underdetermined (10 times more columns than rows). We can employ a sparsity-promoting optimization procedure to infer the most parsimonious x consistent with the equation $A^{(k)}x = y^{(k)}$ for $k = 30, 50$. This procedure, first introduced in reference 23 and proven correct in reference 24, is detailed in “Optimization procedure” in Materials and Methods.

(v) Inferring taxonomy. The abundances of the hypothetical organisms are then mapped back onto the taxonomy (for the output taxonomic profile) or the neighbor-joining tree formed from $A^{(k)}$ (for the output strain variation figures), utilizing a least-common-ancestor approach detailed in “Inferring taxonomy” in Materials and Methods.

Quantification of strain-level variation. We demonstrate in two ways that the inclusion of the hypothetical organisms allows the inference of strain-level variation. First, we spike novel organisms into a mock metagenomic community and show that MetaPalette can accurately predict their presence. Second, we utilize a real metagenomic soil sample to give evidence for a novel strain that MetaPalette predicts.

(i) HMP mock community. We first formed the common k -mer matrices $A^{(k)}$ using 31 strains of *Lysinibacillus sphaericus*. We then used Grinder (25) to simulate a data set consisting of two novel strains (not included in the training database). These reads were then spiked into the Human Microbiome Project (HMP) mock even community (an ~6.6 million-read metagenome consisting of 22 select organisms sampled using an Illumina GA-II sequencer; NCBI accession number [SRR172902](https://www.ncbi.nlm.nih.gov/nuccore/SRR172902)). The output of MetaPalette is shown in Fig. 5, demonstrating the ability of the method to correctly infer the presence of organisms absent from the training data.

Decreasing the number or changing the identity of the training organisms does not impede the method. As shown in Fig. 6, 50,000 simulated reads from the species *Providencia alcalifaciens* were again spiked into the HMP mock even community, and the inferred abundance was again placed optimally on the neighbor-joining tree. See Fig. 11 to 14 for a variety of such figures spanning all domains of life. These results provide evidence that MetaPalette can correctly infer the presence of organisms related to but absent from the training database.

(ii) Metagenomic soil sample. To assess MetaPalette on a real metagenomic sample, we utilized the Iowa prairie metagenomic sample from reference 26 (corresponding to MG-RAST project accession number 6377). After running MetaPalette on a subset of these data (metagenome accession number 4539594.3), the taxonomic profile

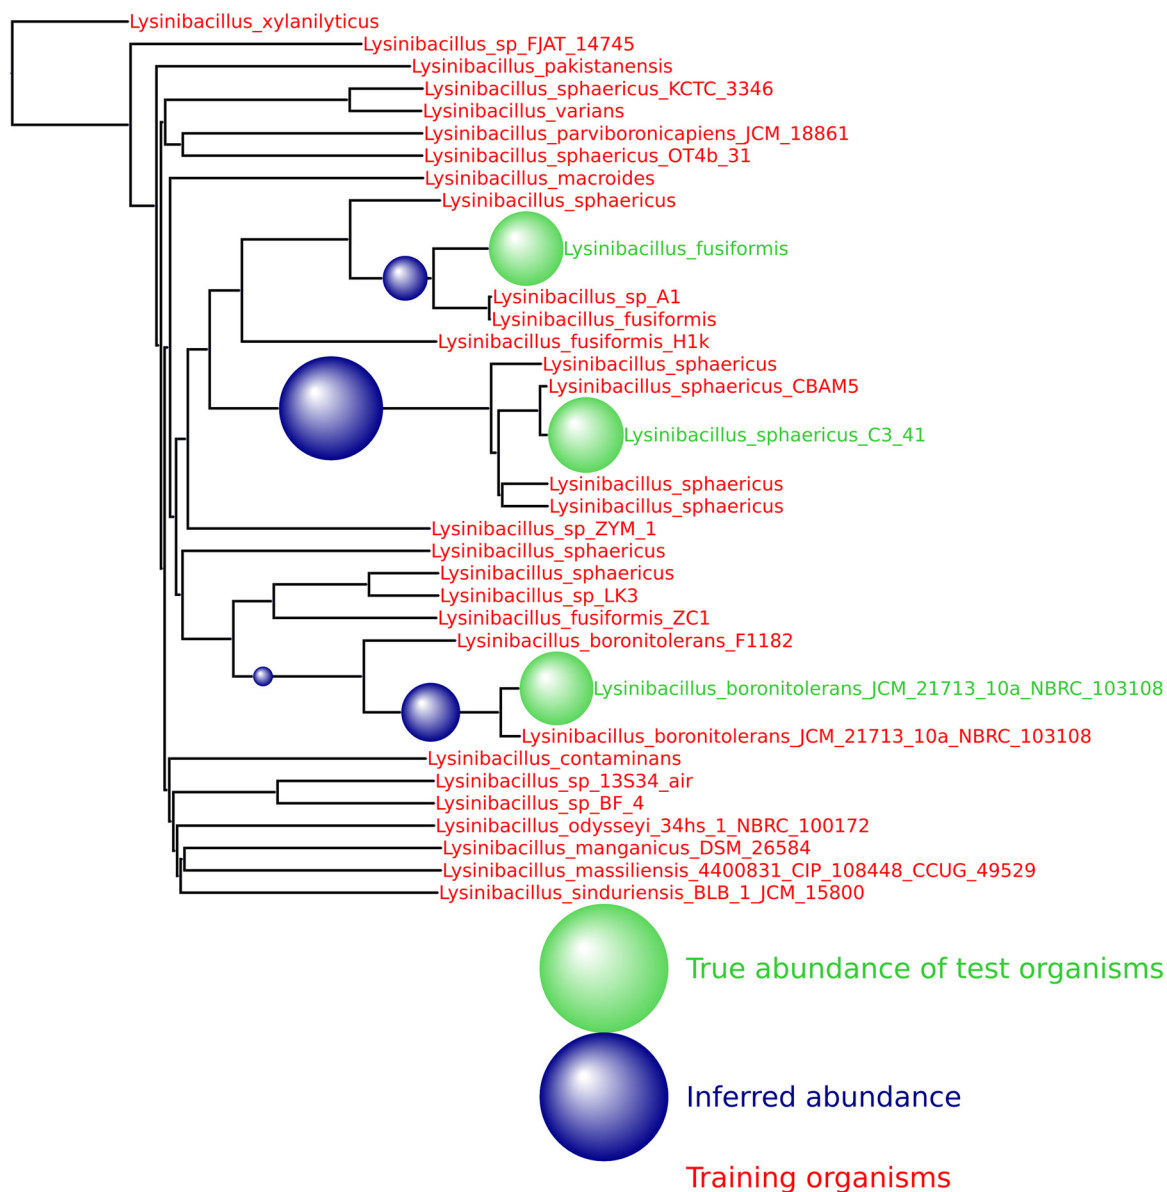


FIG 5 Result from training on 30 strains of *L. sphaericus* and testing on three novel strains. A total of 50,000 reads from the novel strains were spiked into the HMP mock even community. The training organisms are denoted with red font, and the names of the testing organisms are in green font.

that was returned predicted the presence of the genus *Bradyrhizobium*. Generating the tree plot on a subset of this genus resulted in, among others, a prediction of a novel organism in the clade defined by strains of *Bradyrhizobium valentinum* (Fig. 7a). To verify this, we aligned the entire soil metagenome to the reference genome of the strain *B. valentinum* LmjM3 using Bowtie2 with -very-sensitive- local settings (27) and extracted the aligned reads. Interestingly, 0.29% of the reads aligned, while the MetaPalette-predicted abundance for this putative novel organism of interest was 0.33%. The depths of coverage of the extracted reads are pictured in Fig. 7b; the mean depth was 74.3×.

To assess the evolutionary relatedness of this predicted organism, we utilized the *B. valentinum* LmjM3 *nifH* gene sequence (NCBI accession number [KF806461](#)), which was used in the work reported in reference 28, along with other genes, to determine the taxonomy of *B. valentinum* LmjM3. Aligning the extracted reads to *nifH* resulted in a mean depth of coverage of 22×. We collapsed the aligned reads (via a majority vote)

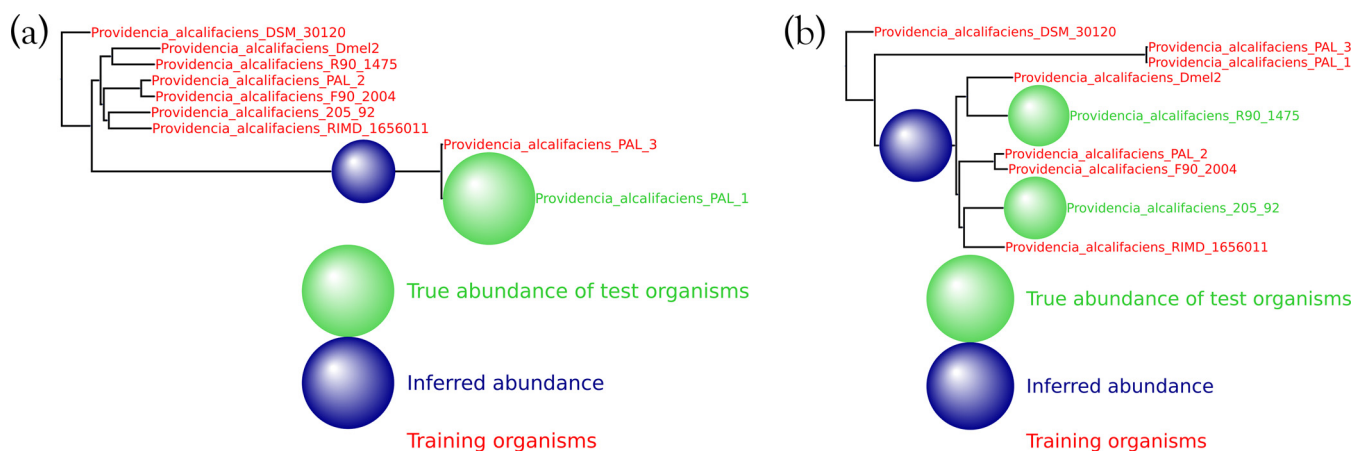


FIG 6 For each of the samples, a total of 50,000 reads from novel strains of *P. alcalifaciens* were spiked into the HMP mock even community. (a) Result from training on 8 strains and testing on 1 novel strain. (b) Result from training on 7 strains and testing on 2 novel strains.

in regions of coverage of at least $22\times$ and called this the maximum-likelihood sequence. We then performed a multiple-sequence alignment of this sequence along with the *nifH* sequences of 20 other organisms closely related to *B. valentinum*. The topology of the bootstrap consensus neighbor-joining tree is pictured in Fig. 7c and shows that the maximum-likelihood sequence is placed at the same location as was predicted by MetaPalette. While this is not enough evidence to unequivocally claim the existence of a novel strain in this sample, this gives support that MetaPalette correctly inferred the abundance and placement, as shown in Fig. 7a, of a potentially novel strain in the clade defined by strains of *B. valentinum*.

Comparison to other metagenomic profiling methods. To facilitate an objective comparison with other methods with minimal “author bias,” we utilized the same data and metrics used by other authors in a recent paper evaluating metagenomics methods (29). This allowed comparison to the following algorithms: CLARK (17), Kraken (16), OneCodex (30), LMAT (31), MG-RAST (32), MetaPhlAn (33), mOTU (14), Genometa (34), QIIME (35), EBI (36), MetaPhyler (15), MEGAN (37), taxator-tk (38), and GOTTCHA (39).

(i) Training data. Each of the methods was trained using the default recommended databases. We trained our method using 6,914 whole-genome sequences and assemblies obtained from various public repositories via RepoPhlAn (<https://bitbucket.org/nsegata/repophlan>). The training procedure for MetaPalette on these 6,914 organisms took a total of approximately 7 h on a 48-core server.

(ii) Testing data. The testing data consisted of 6 samples and are fully explained in reference 29 (in Methods), but we briefly summarize them here. Three replicates were formed from two different distributions of over 900 different genomes spanning the tree of life (including *Eukaryote* genomes). Included in each test sample were shuffled/randomized genomes (not meant to be assigned to any known taxon), as well as sequences from the genome of *Leptospira interrogans* that were evolved using Rose (40) to simulate novelty. Error profiles were based on those of 6 real soil metagenomic samples sequenced using an Illumina HiSeq 2000. Each of the resulting test samples contains between 27 and 37 million read pairs.

(iii) Error metrics. We utilized the same divergence error metric as was used in reference 29; that is, for x_i , representing the true frequency of taxon i in the sample, and x_i^* , representing the predicted frequency of taxon i for a given method,

$$\text{divergence} = \sum_i \log_2 \left(\frac{x_i^*}{x_i} \right)$$

where the summation is over those indices such that $X_i > 0$ and $x_i^* > 0$. Since this error metric does not take into consideration the number of spurious assignments (that is,

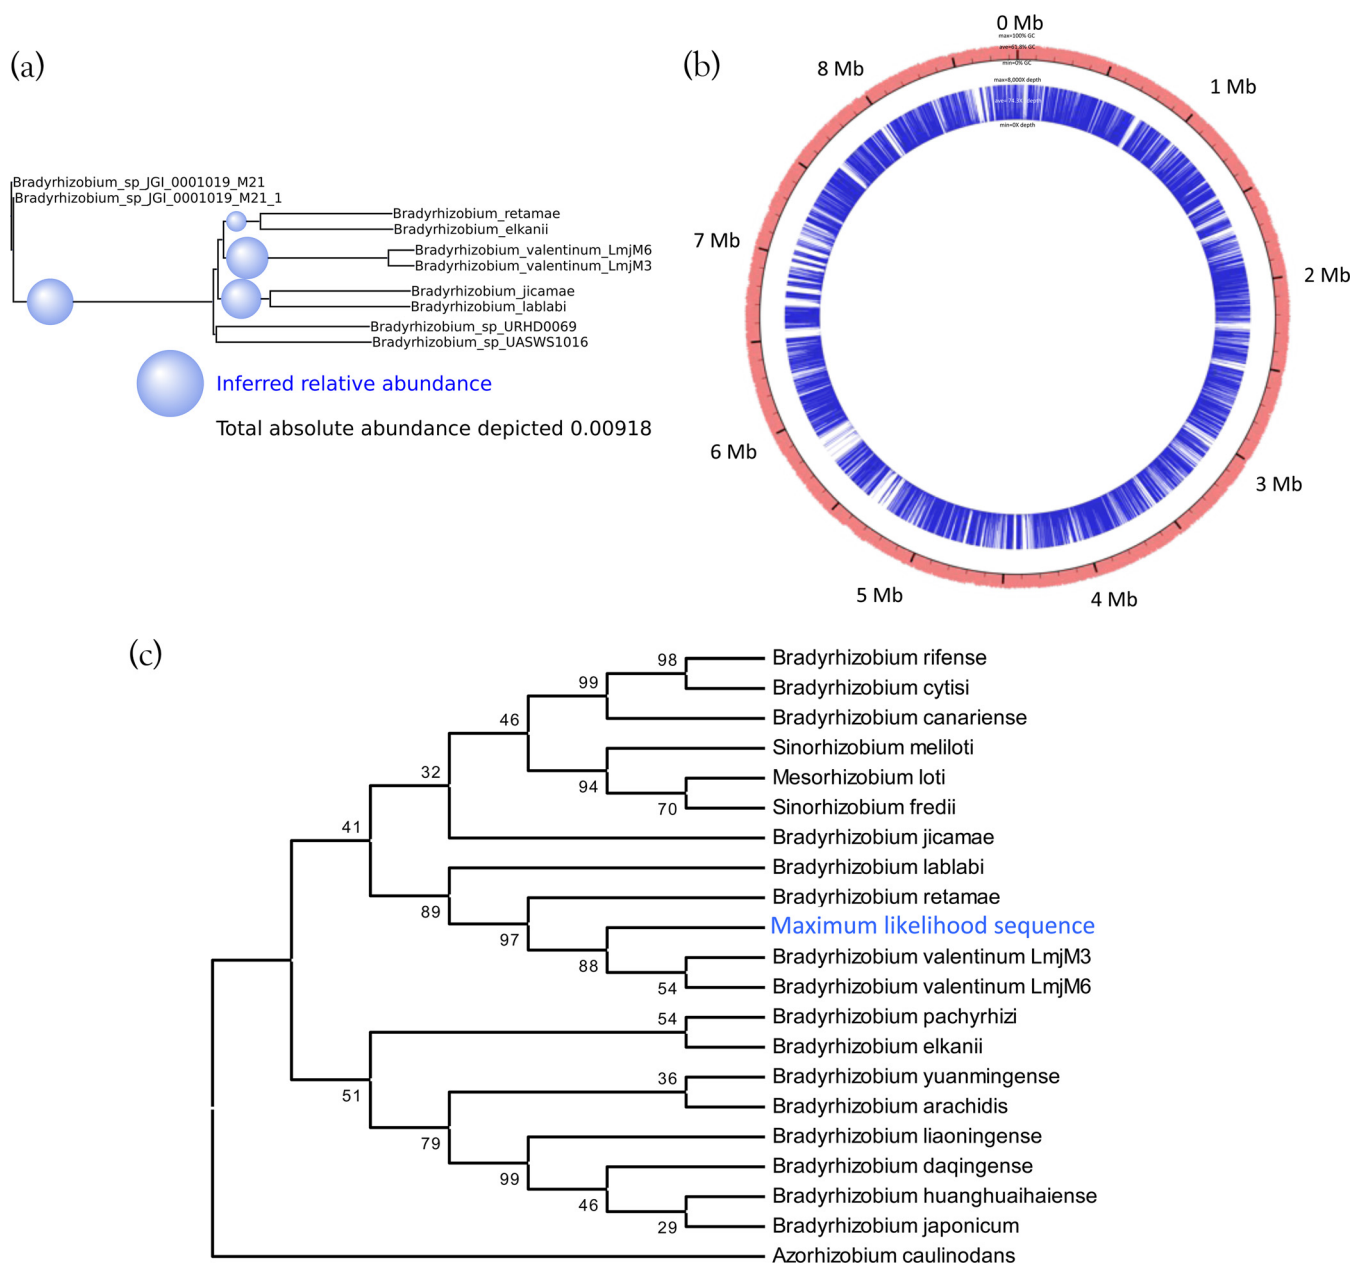


FIG 7 (a) Subtree of the MetaPalette output tree for the Iowa prairie metagenome using organisms from the genus *Bradyrhizobium*. (b) Depths of alignment for reads from the soil metagenome that aligned to *B. valentinum* LmjM3. The outer red ring shows the %GC for *B. valentinum* LmjM3, and the inner blue ring shows the alignment depth (truncated to 8,000× for ease of viewing). All contigs of the reference strain *B. valentinum* LmjM3 were concatenated in this figure. (c) Bootstrap consensus tree topology based on *nifH* for 20 organisms, along with the maximum-likelihood sequence obtained from aligning the soil metagenome to the *nifH* gene sequence of *B. valentinum* LmjM3. Bootstrap values (500 replicates) are shown next to the branches. Full details regarding the formation of the tree are given in “Sequence analysis details” in Materials and Methods.

taxa predicted by a method to be in a sample but not actually present), we also use the number of false positives (FP) at a given taxonomic rank, as follows:

$$FP = |\{i : x_i^* > 0 \text{ and } x_i = 0\}|.$$

(iv) Comparison results. Each method was run using the default parameters. For each method, we averaged the divergence error metric over all the test samples at the genus level (Fig. 8a). Furthermore, we selected a number of the more accurate methods and averaged the number of false positives over all the test samples at the phylum level (Fig. 8b). The results in Fig. 8a and b clearly show the competitive nature of MetaPalette, as it has the lowest error using both metrics. However, when comparing to other

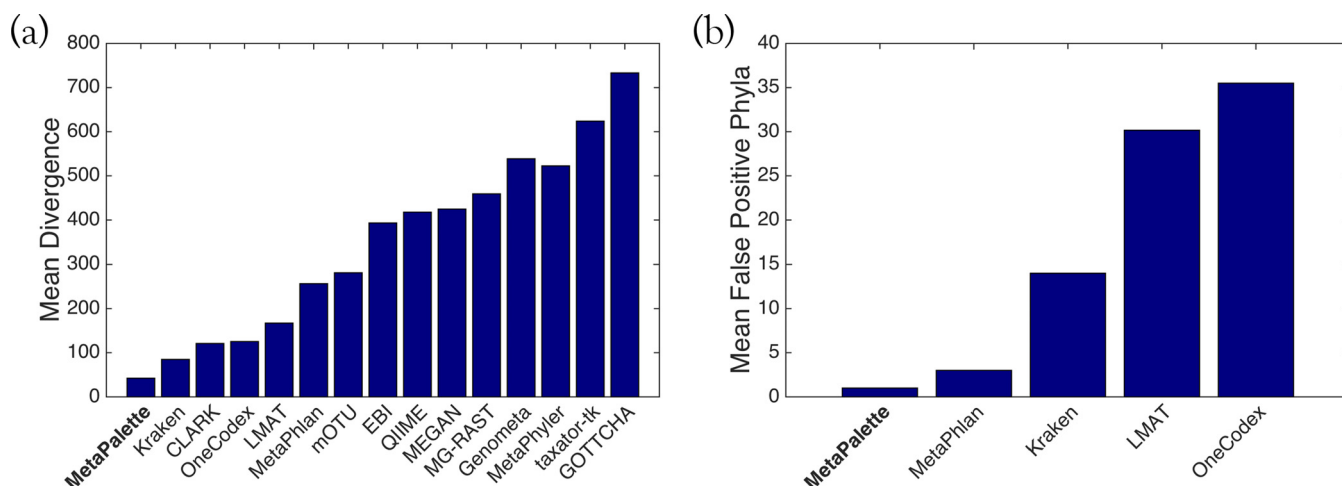


FIG 8 Plot of performance metrics for all metagenomics methods averaged over all test samples. Smaller values indicate better performance. (a) Divergence error metric at the genus level. (b) Number of false-positive phyla.

methods, one should be careful of their intended use. For example, taxator-tk is intended to be used on an assembled metagenome (and here unassembled reads were used), and QIIME only uses the 16S rRNA sequences in a sample. Furthermore, most of these methods assign individual reads and then summarize them to obtain a taxonomic profile, while our method only profiles the entire sample and returns relative proportions of organisms.

Figure 9 shows the execution time of each of the methods (on a log scale, obtained from reference 29), further showing the competitive nature of MetaPalette.

Software and pretrained data. (i) Software. The source code for MetaPalette, along with installation instructions and directions, is accessible at <https://github.com/dkoslicki/MetaPalette>. MetaPalette is written primarily in python and accepts input reads in uncompressed fasta or fastq format, as well as compressed fasta/fastq using bzip2 and gzip. For fastq input, optional parameters can be given to specify counting *k*-mers only above a certain quality score (Phred), thereby attenuating the negative impact of sequencing error in the correct inference of relative abundances. The output taxonomic profile is compliant with the Bioboxes profiling format, version 0.9, found at <https://github.com/bioboxes/rfc/tree/master/data-format>. Python scripts are also in-

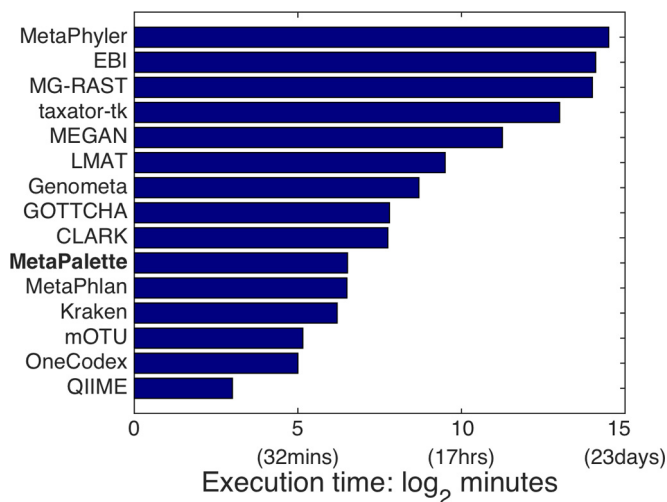


FIG 9 Mean execution time of each method averaged over all 6 test samples.

cluded to aid in downloading data, forming custom databases, and creating the appropriate taxonomy files.

To facilitate cross-platform usability, a Docker (41) container has been created and is accessible at: <https://hub.docker.com/r/dkoslicki/metapalette>, with an accompanying docker file at: <https://github.com/dkoslicki/MetaPalette/blob/master/Docker/Dockerfile>.

If users wish to use MetaPalette but lack computational resources, they may utilize the Galaxy (42–44) server located at: <http://math-galaxy.cgrb.oregonstate.edu/>.

A preliminary version of this software was submitted to Critical Assessment of Metagenomic Interpretation (CAMI: <http://www.cami-challenge.org/>) under the name CommonKmers. However, since significant changes have been made since that point, we strongly recommend using the current MetaPalette software instead.

(ii) Pretrained data. To decrease the computational burden, pretrained databases are accessible at http://files.cgrb.oregonstate.edu/Koslicki_Lab/MetaPalette. Databases and accompanying taxonomies have been included for *Archaea* (666 organisms; 1.7 GB uncompressed), *Bacteria* (15,147 organisms; 60 GB), *Eukaryota* (1,307 organisms; 41 GB), and viruses (4,798 organisms; 0.6 GB). All organisms were obtained via RepoPhlAn.

The 6,914-organism database used for the comparison to other profiling methods is accessible at http://files.cgrb.oregonstate.edu/Koslicki_Lab/MetaPalette (Comparison.tar.gz).

DISCUSSION

We have described a fast, flexible, and accurate method for estimating the taxonomic composition of organisms which is based on reconstructing a k -mer-based profile of a sample. Each reference organism has a k -mer “palette,” and we fit the sample as a mixture of different palettes, both of the reference organisms and organisms absent from the training data at various degrees of relatedness to the training database. Our approach is in part inspired by the chromosome-painting method used to deduce fine-scale population structure in human genetics (45, 46), which is also based on mixture modeling of palettes. A particular advantage of MetaPalette over other metagenomic profiling methods is that MetaPalette provides an indication of how related the organisms in a given sample are to the closest matching organisms of the training database, whether they are within the same species or distantly related organisms from the same phyla.

Furthermore, the standard approach to summarizing composition information has been to place organisms at different taxonomic levels. We produce a standard taxonomic profile which we have shown to be more accurate than that produced by other methods. This fixed-rank approach is sensible at the genus level and above but omits fine-scale information. Hence, for branches of the tree of life that are well represented in the training database, we can also output a phylogenetic tree giving detailed information on how the sampled taxa relate to the organisms in the training database (Fig. 5 to 7; see also Fig. 11 to 14 in Materials and Methods).

For many applications, it is of interest to understand which individual reads belong to which organisms (1). A principled approach to this problem is to first estimate the overall composition of the sample, using MetaPalette or an equivalent, and then to assign individual reads conditional on the overall assignment. This represents a promising avenue for future methodological development.

MATERIALS AND METHODS

Mathematical formulation. We include here all rigorous mathematical definitions of the quantities discussed in the main text.

Given the alphabet $\mathcal{A} = \{A, C, T, G\}$, let \mathcal{A}^n denote the set of all words v of length $|v| = n$ on \mathcal{A} , and let $\mathcal{A}^* = \cup_{n \geq 0} \mathcal{A}^n$ be the set of all finite words on \mathcal{A} . Hence, words containing non-ACTG characters are ignored. Let $D = \{g_1, \dots, g_M\}$ be a database of genomic sequences $g_i \in \mathcal{A}^*$ and let $S = \{s_1, \dots, s_N\}$ be a set of sample sequences (the reads to be classified). For notational simplicity, assume that the read length is fixed; i.e., for all $t = 1, \dots, N$ $|s_t| = r$. Fix a k -mer size and endow $\mathcal{A}^k = \{v_1, \dots, v_{4^k}\}$ with the lexicographic order. Let $\text{occ}_v(w)$ represent the number of occurrences (with overlap) of the subword v in the word w . That is, for $w, v \in \mathcal{A}^n$, let

$$\text{occ}_v(w) = |\{j : w_j w_{j+1} \cdots w_{j+k-1} = v\}| \quad (1)$$

For a fixed k -mer size and two genomes, g_i and g_j , we calculate the number of k -mers in genome j common to both g_i and g_j . That is, the (i,j) th entry of the common k -mer training matrix $A^{(k)}$ is as follows:

$$A_{i,j}^{(k)} = \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)}} \frac{\text{occ}_w(g_j)}{|g_j| - k + 1} \quad (2)$$

Refer to the entries of the common k -mer matrix as $\text{pckm}_k(g_i, g_j) = A_{i,j}^{(k)}$. Let $s_t \subset g_j$ denote the relationship that read s_t was derived from genome g_j . We represent the taxonomic profile of the sample S by the probability vector x as follows:

$$x_j = \frac{1}{N} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (3)$$

where $\mathbb{1}$ is the indicator function. Now let the measurement vector y be given by the probability vector

$$y_i^{(k)} = \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_S(k)} \text{occ}_w(S) / \sum_{w \in \text{SW}_S(k)} \text{occ}_w(S) \quad (4)$$

We assume that the reads s_t are uniformly randomly selected from the genome g_j . Then, for $w \in \mathcal{A}^k$, let $\mathbb{P}(w|g_j)$ be the probability that k -mer w is found in genome g_j . Then, we have that the proportion of k -mers w in the sample is similar to the proportion of the appearance of w in the genome g_j when weighted by the relative abundance of the genome g_j in the sample, as follows:

$$\frac{\sum_{t=1}^N \text{occ}_w(s_t)}{\sum_{w \in \text{SW}_S(k)} \text{occ}_w(S)} \approx \frac{1}{N} \sum_{j=1}^M \mathbb{P}(w|g_j) \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (5)$$

$$= \frac{1}{N} \sum_{j=1}^M \frac{\text{occ}_{g_j}(w)}{|g_j| - k + 1} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (6)$$

We then calculate

$$y_i^{(k)} = \frac{1}{\sum_{w \in \text{SW}_S(k)} \text{occ}_w(S)} \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_S(k)} \text{occ}_w(S) \quad (7)$$

$$= \frac{1}{\sum_{w \in \text{SW}_S(k)} \text{occ}_w(S)} \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_S(k)} \sum_{t=1}^N \text{occ}_w(s_t) \quad (8)$$

$$\approx \frac{1}{N} \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_S(k)} \sum_{j=1}^M \frac{\text{occ}_{g_j}(w)}{|g_j| - k + 1} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (9)$$

$$= \frac{1}{N} \sum_{j=1}^M \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)} \cap \text{SW}_S(k)} \frac{\text{occ}_{g_j}(w)}{|g_j| - k + 1} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (10)$$

$$\approx \frac{1}{N} \sum_{j=1}^M \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)}} \frac{\text{occ}_{g_j}(w)}{|g_j| - k + 1} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (11)$$

$$= \sum_{j=1}^M A_{i,j}^{(k)} x_j = (A^{(k)} x)_i \quad (12)$$

Equation 10 is justified since if $w \notin \text{SW}_{g_j(k)}$, then $\text{occ}_{g_j}(w) = 0$. For computational reasons, we make the assumption in equation 11 that $\text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)} \cap \text{SW}_S(k) = \text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)}$. However, this assumption can be mitigated by adding hypothetical organisms (see ‘‘Hypothetical organisms’’ below). Our assumptions imply that

$$A^{(k)} x \approx y^{(k)} \quad (13)$$

We will try to recover the vector x satisfying $x_j \geq 0$ for all $j = 1, \dots, M$ from equation 13.

Further improvements. A few further improvements are possible, but not pursued here. Namely, we could use just the k -mers that are actually in the sample to form the training matrix. That is, use $\text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)} \cap \text{SW}_S(k)$ in the formation of $A^{(k)}$, as follows:

$$A_{i,j}^{(k)} = \sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_{g_j(k)} \cap \text{SW}_S(k)} \frac{\text{occ}_w(g_j)}{|g_j| - k + 1}$$

The disadvantage of this is that the (slow) training step would need to be rerun for each sample.

For a second improvement, we could make the approximation in equation 5 more delicate by incorporating the coverage

$$\frac{\sum_{t=1}^N \text{occ}_w(s_t)}{\sum_{w \in \text{SW}_S(k)} \text{occ}_w(S)} \approx \frac{1}{N} \sum_{j=1}^M \mathbb{P}(w|g_j) \frac{\sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_S(k)} \text{occ}_w(g_j)}{\sum_{w \in \text{SW}_{g_j(k)}} \text{occ}_w(g_j)} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (14)$$

$$= \frac{1}{N} \sum_{j=1}^M \frac{\text{occ}_{g_j}(w)}{|g_j| - n + 1} \frac{\sum_{w \in \text{SW}_{g_i(k)} \cap \text{SW}_S(k)} \text{occ}_w(g_j)}{|g_j| - k + 1} \sum_{t=1}^N \mathbb{1}_{\{s_t \subset g_j\}} \quad (15)$$




FIG 10 Neighbor-joining tree based on *nifH* for 20 organisms along with the maximum-likelihood sequence obtained from aligning the soil data to the gene sequence for *nifH* of *B. valentinum* LmjM3. Bootstrap values are shown next to the branches, and the bar indicates 0.5 nucleotide substitutions per site.

So $A^{(k)}$ would have the form:

$$A_{ij}^{(k)} = \frac{\sum_{w \in SW_{g_i}(k) \cap SW_{g_j}(k)} \text{occ}_w(g_j)}{|g_j| - k + 1} \sum_{w \in SW_{g_i}(k) \cap SW_{g_j}(k) \cap SW_{g_s}(k)} \frac{\text{occ}_w(g_j)}{|g_j| - k + 1} \quad (16)$$

This effectively multiplies column j of $A^{(k)}$ by the percent coverage of genome j . Finally, in equation 16, we could put a weighting factor that represents how unique a k -mer is to the genome in question. This would down-weight k -mers shared among many diverse genomes and up-weight those unique to certain strains/species/genera/etc.

Hypothetical organisms. To simulate an organism that is, say, 90% related to a database genome g_j , we augment the common k -mer matrix $A^{(k)}$ with a column derived by rounding down the entries of the column vector $\{A_{ij}^{(k)}\}_{i=1, \dots, M}$ that are above 90%. Two k -mer sizes are needed to form the hypothetical organism's common k -mer matrices. For the first k -mer size, k_1 , we define $A^{(k_1),h}$ for a fixed number of hypothetical bins $h \in \{0.9, 0.8, \dots, 0.1\}$, where

$$A_{ij}^{(k_1),h} = \max \left[h, \sum_{w \in SW_{g_i}(k_1) \cap SW_{g_j}(k_1)} \frac{\text{occ}_w(g_j)}{|g_j| - n + 1} \right] \quad (17)$$

For the second k -mer size, k_2 , using the polynomial $p(x) = -.5141x^3 + 1.0932x^2 + 0.3824x$, we define

$$A_{ij}^{(k_2),h} = \max \left[h, p \left(\sum_{w \in SW_{g_i}(k_2) \cap SW_{g_j}(k_2)} \frac{\text{occ}_w(g_j)}{|g_j| - n + 1} \right) \right] \quad (18)$$

Instead of thresholding, as we did here, one can imagine other scalings obtained from studying the relationship between a given taxonomy and the common k -mer matrix $A^{(k)}$. In particular, to deal with differing rates of evolution in the tree of life, a fruitful area of future investigation would be to modify the polynomial p depending on the taxonomy of the organisms under consideration.

Optimization procedure. We choose two k -mer sizes to be $k_1 = 30$ and $k_2 = 50$, as this seems to give a good trade-off between reconstruction fidelity and computational performance. We then collect the common k -mer matrix and hypothetical matrices blockwise into the $2|D| \times 10|D|$ size matrix as follows:

$$A = \begin{bmatrix} A^{(30)} & A^{(30),0.9} & \dots & A^{(30),0.1} \\ A^{(50)} & A^{(50),0.9} & \dots & A^{(50),0.1} \end{bmatrix} \quad (19)$$

Collect also the k -mer sample vectors $y^{(k)}$:

$$y = \begin{bmatrix} y^{(30)} \\ y^{(50)} \end{bmatrix} \quad (20)$$

The problem at hand is then to reconstruct the phylogenetic profile x by solving the linear system

$$Ax = y \tag{21}$$

Equation 21 is solved by using a sparsity-promoting optimization procedure motivated by techniques used in the compressive sensing literature. Sparsity is emphasized due to the inclusion of the hypothetical organisms, as well as the reasonable assumption that relatively few organisms from the database D are actually present in the given sample. We use a variant of nonnegative basis pursuit denoising which reduces to a nonnegative least-squares problem (24, 47). We aim to solve

$$x^* = \arg \min_z \|z\|_1 \text{ subject to } Az = y, z \geq 0 \tag{(\ell_1\text{-min})}$$

This optimization procedure has the advantage of being transformed into a nonnegative least-squares problem. Indeed, as $\lambda \rightarrow \infty$, we can regularize $(\ell_1\text{-min})$ as

$$x^* = \arg \min_z \|z\|_1^2 + \lambda^2 \|Az - y\|_2^2 \text{ subject to } z \geq 0 \tag{(NNREG)}$$

This reduces to a nonnegative least-squares problem by defining

$$\tilde{A} = \begin{bmatrix} 1 & \dots & 1 \\ \lambda A \end{bmatrix}, \tilde{y} = \begin{bmatrix} 0 \\ \lambda y \end{bmatrix}$$

So (NNREG) is equivalent to the nonnegative least-squares problem

$$x^* = \arg \min_z \|\tilde{A}z - \tilde{y}\|_2^2 \text{ subject to } z \geq 0$$

This can be solved efficiently by using the LawsonHanson algorithm (48). We use the value $\lambda = 200$ throughout, as this value gives a good trade-off between sparsity and accuracy of fit of the k -mer counts.

Inferring taxonomy. Since the reconstructed vector x may have nonzero entries corresponding to a hypothetical bin, we need to develop a method to map from a hypothetical bin to a specific taxonomic rank. A naive approach would be to assign a fixed taxonomic rank to each hypothetical bin (call this the “fixed rank” method). For example, all nonzero entries of x corresponding to $A^{(k)}$ would be assigned to the strain level, all nonzero entries of x corresponding to $A^{(k),0.9}$ would be assigned to the species level, etc.

We take a more biologically informed approach: we take the least common ancestor (LCA) taxon between a hypothetical organism and a nearby organism in the database D : if $x_i > 0$ corresponds to the hypothetical bin h , find an organism g_j such that $|A_{ij}^{(k)} - h| < \delta$ for some threshold δ . In the output taxonomic profile, we assign x_i to the lowest taxonomic rank common to the organisms with genomes g_i and g_j . For the output strain variation figures, we assign the abundance x_i relative to the least common ancestor of g_i and g_j (above the LCA if $h < A_{ij}^{(k)}$ and below the LCA if $h > A_{ij}^{(k)}$).

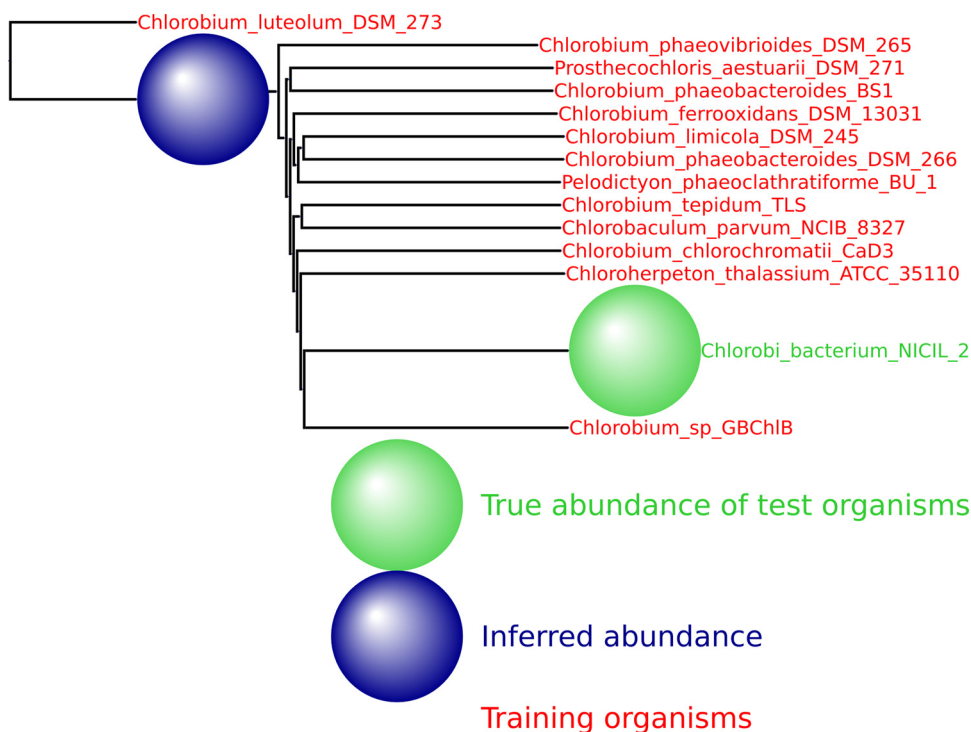


FIG 11 HMP spiking results for the bacterial phylum *Chlorobi* with 13 training organisms and 1 novel testing organism.

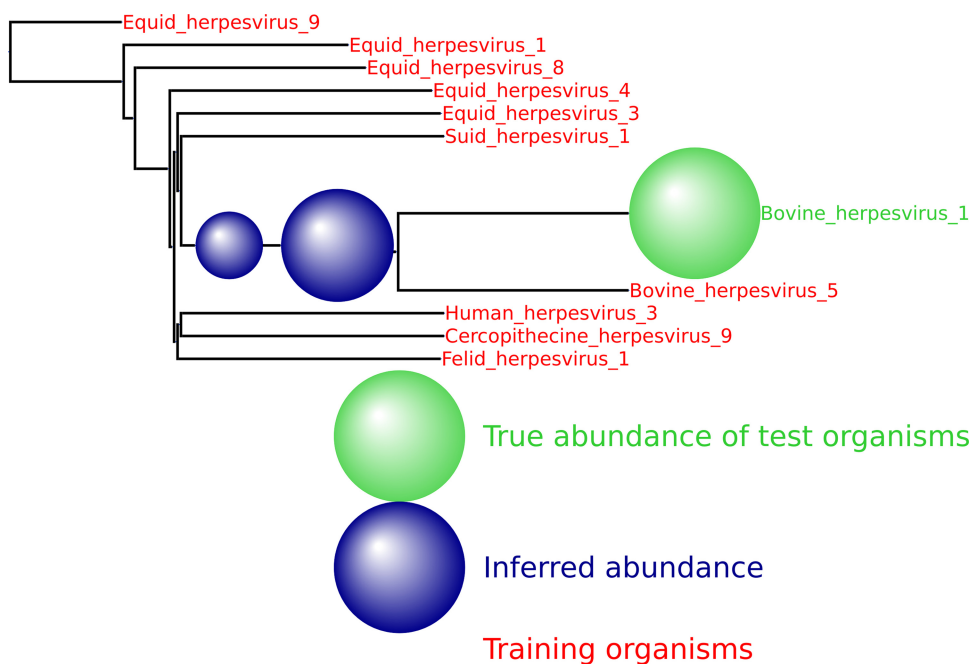


FIG 12 HMP spiking results for the viral genus *Varicellovirus* with 10 training organism and 1 novel testing organism. Only 5,000 simulated reads were spiked into the HMP sample.

For the output taxonomic profile, a hybrid of the fixed rank and LCA approaches can increase sensitivity or specificity. We thus include three options: the default option is the LCA approach, while the sensitive and specific options are various hybrids of the two methods.

Sequence analysis details. For the tree shown in Fig. 7c, the evolutionary history was inferred using the neighbor-joining method (49). The bootstrap consensus tree, inferred from 500 replicates, is taken to represent the evolutionary history of the taxa analyzed (50). Branches corresponding to partitions reproduced in less than 50% of the bootstrap replicates are collapsed. The percentages of replicate trees in which the associated taxa clustered together in the bootstrap test (500 replicates) are shown next to the branches. The analysis involved 21 nucleotide sequences. All positions with less than 95% site coverage were eliminated. That is, less than 5% alignment gaps, missing data, and ambiguous bases were allowed at any position. There were a total of 652 positions in the final data set. Evolutionary analyses were conducted in MEGA6 (51).

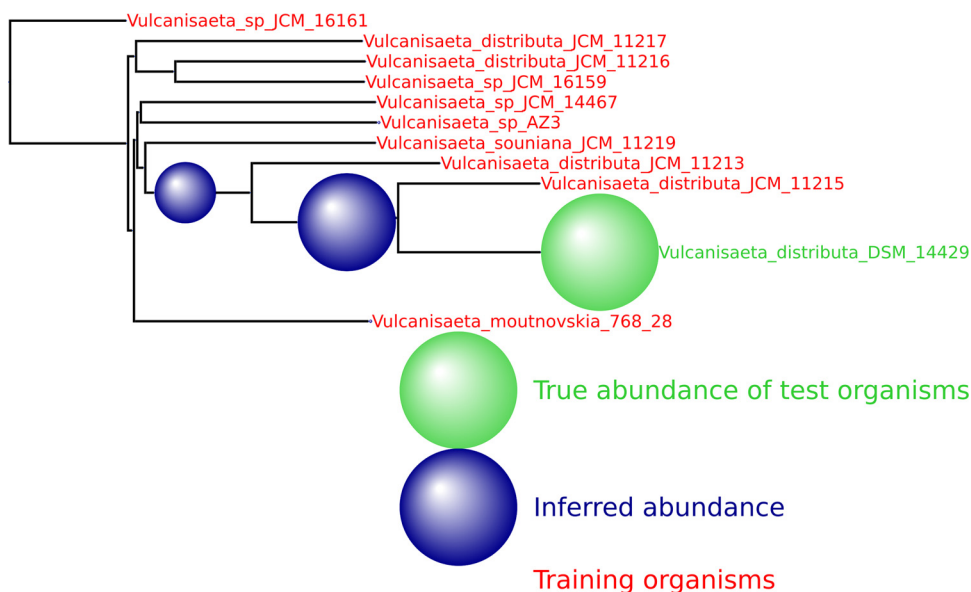


FIG 13 HMP spiking results for the archaeal genus *Vulcanisaeta* with 10 training organism and 1 novel testing organism. Only 5,000 simulated reads were spiked into the HMP sample.

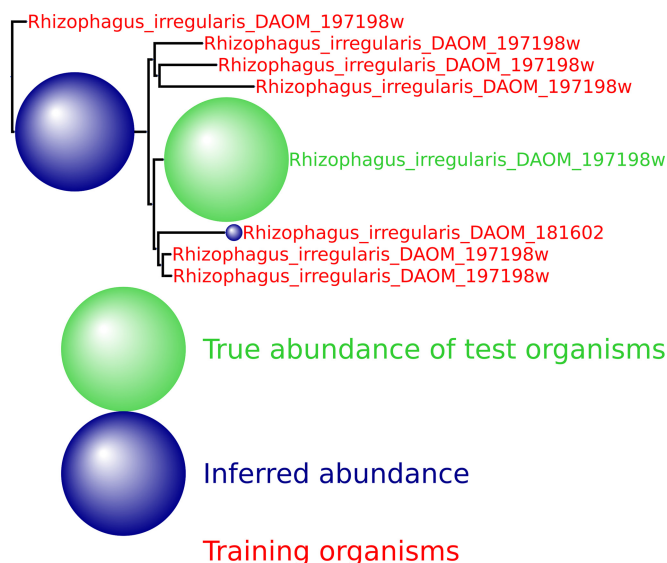


FIG 14 HMP spiking results for the *Eukaryota* genus *Rhizophagus* with 7 training organism and 1 novel testing organism.

Figure 10 depicts a tree constructed using the same method as just described but with evolutionary distances computed using the maximum-composite-likelihood method (52). The unit is the number of base substitutions per site.

Additional figures. We provide here a number of additional output figures from MetaPalette (Fig. 11 to 14) to demonstrate that the ability to correctly infer the presence of organisms related to but absent from the training database is not dependent on the particular kingdom/phyla/etc. used. Unless otherwise noted, a total of 50,000 simulated reads from the novel organisms were spiked into the HMP mock even community. Figures are included for *Bacteria*, *Archaea*, *Eukaryota*, and viruses.

ACKNOWLEDGMENTS

We thank the Isaac Newton Institute at Cambridge University for their hospitality during the program on metagenomics. This project was conceived while both authors were attending this program. Daniel Falush is supported by a Medical Research Fellow fellowship as part of the Climb Consortium for Medical Microbiology.

We also thank Nam Nguyen and Daniel Alemany, who both contributed to this project in its preliminary stages.

FUNDING INFORMATION

Daniel Falush is supported by a Medical Research Fellow fellowship as part of the CLIMB Consortium for Medical Microbiology.

REFERENCES

1. **Sharpton TJ.** 2014. An introduction to the analysis of shotgun metagenomic data. *Front Plant Sci* **5**:209. <http://dx.doi.org/10.3389/fpls.2014.00209>.
2. **Marcy Y, Ouverney C, Bik EM, Lösekann T, Ivanova N, Martin HG, Szeto E, Platt D, Hugenholtz P, Relman DA, Quake SR.** 2007. Dissecting biological “dark matter” with single-cell genetic analysis of rare and uncultivated tm7 microbes from the human mouth. *Proc Natl Acad Sci U S A* **104**:11889–11894. <http://dx.doi.org/10.1073/pnas.0704662104>.
3. **Baker BJ, Comolli LR, Dick GJ, Hauser LJ, Hyatt D, Dill BD, Land ML, VerBerkmoes NC, Hettich RL, Banfield JF.** 2010. Enigmatic, ultrasmall, uncultivated Archaea. *Proc Natl Acad Sci* **107**:8806–8811. <http://dx.doi.org/10.1073/pnas.0914470107>.
4. **Lan Y, Wang Q, Cole JR, Rosen GL.** 2012. Using the rdp classifier to predict taxonomic novelty and reduce the search space for finding novel organisms. *PLoS One* **7**:e32491. <http://dx.doi.org/10.1371/journal.pone.0032491>.
5. **Rinke C, Schwientek P, Sczyrba A, Ivanova NN, Anderson IJ, Cheng JF, Darling A, Malfatti S, Swan BK, Gies EA, Dodsworth JA, Hedlund BP, Tsiamis G, Sievert SM, Liu WT, Eisen JA, Hallam SJ, Kyrpides NC, Stepanauskas R, Rubin EM.** 2013. Insights into the phylogeny and coding potential of microbial dark matter. *Nature* **499**:431–437. <http://dx.doi.org/10.1038/nature12352>.
6. **Puigbò P, Wolf YI, Koonin EV.** 2009. Search for a “tree of life” in the thicket of the phylogenetic forest. *J Biol* **8**:59. <http://dx.doi.org/10.1186/jbiol159>.
7. **Wang B.** 2001. Limitations of compositional approach to identifying horizontally transferred genes. *J Mol Evol* **53**:244–250. <http://dx.doi.org/10.1007/s002390010214>.
8. **Rhille H, Brinkmann H, Lavrov DV, Littlewood DTJ, Manuel M, Wörheide G, Baurain D.** 2011. Resolving difficult phylogenetic questions: why more sequences are not enough. *PLoS Biol* **9**:e1000602. <http://dx.doi.org/10.1371/journal.pbio.1000602>.
9. **Salichos L, Rokas A.** 2013. Inferring ancient divergences requires genes

- with strong phylogenetic signals. *Nature* **497**:327–331. <http://dx.doi.org/10.1038/nature12130>.
10. **Roure B, Baurain D, Philippe H.** 2013. Impact of missing data on phylogenies inferred from empirical phylogenomic data sets. *Mol Biol Evol* **30**:197–214. <http://dx.doi.org/10.1093/molbev/mss208>.
 11. **Kaper JB, Nataro JP, Mobley HL.** 2004. Pathogenic *Escherichia coli*. *Nat Rev Microbiol* **2**:123–140. <http://dx.doi.org/10.1038/nrmicro818>.
 12. **Haft DH, Tovchigrechko A.** 2012. High-speed microbial community profiling. *Nat Methods* **9**:793–794. <http://dx.doi.org/10.1038/nmeth.2080>.
 13. **Nguyen NP, Mirarab S, Liu B, Pop M, Warnow T.** 2014. TIPP: taxonomic identification and phylogenetic profiling. *Bioinformatics* **30**:3548–3555. <http://dx.doi.org/10.1093/bioinformatics/btu721>.
 14. **Sunagawa S, Mende DR, Zeller G, Izquierdo-Carrasco F, Berger SA, Kultima JR, Coelho LP, Arumugam M, Tap J, Nielsen HB, Rasmussen S, Brunk S, Pedersen O, Guarner F, de Vos WM, Wang J, Li J, Doré J, Ehrlich SD, Stamatakis A, Bork P.** 2013. Metagenomic species profiling using universal phylogenetic marker genes. *Nat Methods*. <http://dx.doi.org/10.1038/nmeth.2693>.
 15. **Liu B, Gibbons T, Ghodsi M, Pop M.** 2010. MetaPhyler: taxonomic profiling for metagenomic sequences, p 95–100. 2010 IEEE International Conference on IEEE Bioinformatics and Biomedicine (BIBM), 18 to 20 December 2010, Hong Kong, China. Institute of Electrical and Electronics Engineers, New York, NY. <http://dx.doi.org/10.1109/BIBM.2010.5706544>.
 16. **Wood DE, Salzberg SL.** 2014. Kraken: ultrafast metagenomic sequence classification using exact alignments. *Genome Biol* **15**:R46. <http://dx.doi.org/10.1186/gb-2014-15-3-r46>.
 17. **Ounit R, Wanamaker S, Close TJ, Lonardi S.** 2015. Fast and accurate classification of metagenomic and genomic sequences using discriminative k-mers. *BMC Genomics* **16**:236. <http://dx.doi.org/10.1186/s12864-015-1419-2>.
 18. **Koslicki D, Foucart S, Rosen G.** 2014. WGSQuikr: fast whole-genome shotgun metagenomic classification. *PLoS One* **9**:1784. <http://dx.doi.org/10.1371/journal.pone.0091784>.
 19. **Schaeffer L, Pimentel H, Bray N, Melsted P, Pachter L.** 2015. Pseudo-alignment for metagenomic read assignment. arXiv arXiv:1510.07371. [q-bio.QM] <http://arxiv.org/abs/1510.07371v2>.
 20. **Jaccard P.** 1901. Etude comparative de la distribution florale dans une portion des Alpes et du Jura. *Bull Soc Vaudoise Sci Nat* **37**:547–579.
 21. **Marçais G, Kingsford C.** 2011. A fast, lock-free approach for efficient parallel counting of occurrences of k-mers. *Bioinformatics* **27**:764–770. <http://dx.doi.org/10.1093/bioinformatics/btr011>.
 22. **Chikhi R, Limasset A, Jackman S, Simpson JT, Medvedev P.** 2014. On the representation of de Bruijn graphs, p 35–55. In Sharan R (ed), *Research in computational molecular biology. Proceedings of the 18th Annual International Conference, RECOMB 2014, Pittsburgh, PA, 2 to 5 April 2014. Lecture notes in computer science 8394*. SpringerVerlag, Cham, Switzerland.
 23. **Koslicki D, Foucart S, Rosen G.** 2013. Quikr: a method for rapid reconstruction of bacterial communities via compressive sensing. *Bioinformatics* **29**:2096–2102. <http://dx.doi.org/10.1093/bioinformatics/btt336>.
 24. **Foucart S, Koslicki D.** 2014. Sparse recovery by means of nonnegative least squares. *IEEE Signal Process Lett* **21**:498–502. <http://dx.doi.org/10.1109/LSP.2014.2307064>.
 25. **Angly FE, Willner D, Rohwer F, Hugenholz P, Tyson GW.** 2012. Grinder: a versatile amplicon and shotgun sequence simulator. *Nucleic Acids Res* **40**:e94. <http://dx.doi.org/10.1093/nar/gks251>.
 26. **Howe AC, Jansson JK, Malfatti SA, Tringe SG, Tiedje JM, Brown CT.** 2014. Tackling soil diversity with the assembly of large, complex metagenomes. *Proc Natl Acad Sci U S A* **111**:4904–4909. <http://dx.doi.org/10.1073/pnas.1402564111>.
 27. **Langmead B, Salzberg SL.** 2012. Fast gapped-read alignment with bowtie2. *Nat Methods* **9**:357–359. <http://dx.doi.org/10.1038/nmeth.1923>.
 28. **Durán D, Rey L, Navarro A, Busquets A, Imperial J, Ruiz-Argüeso T.** 2014. *Bradyrhizobium valentinum* sp. nov., isolated from effective nodules of *Lupinus mariae-josephae*, a lupine endemic of basic-lime soils in Eastern Spain. *Syst Appl Microbiol* **37**:336–341. <http://dx.doi.org/10.1016/j.syapm.2014.05.002>.
 29. **Lindgreen S, Adair KL, Gardner PP.** 2016. An evaluation of the accuracy and speed of metagenome analysis tools. *Sci Rep* **6**:19233. <http://dx.doi.org/10.1038/srep19233>.
 30. **Minot SS, Krumm N, Greenfield NB.** 2015. One codex: a sensitive and accurate data platform for genomic microbial identification. *bioRxiv* <http://dx.doi.org/10.1101/027607>.
 31. **Ames SK, Hysom DA, Gardner SN, Lloyd GS, Gokhale MB, Allen JE.** 2013. Scalable metagenomic taxonomy classification using a reference genome database. *Bioinformatics* **29**:2253–2260. <http://dx.doi.org/10.1093/bioinformatics/btt389>.
 32. **Meyer F, Paarmann D, D'Souza M, Olson R, Glass E, Kubal M, Paczian T, Rodriguez A, Stevens R, Wilke A, Wilkening J, Edwards R.** 2008. The metagenomics RAST server—a public resource for the automatic phylogenetic and functional analysis of metagenomes. *BMC Bioinformatics* **9**:386. <http://dx.doi.org/10.1186/1471-2105-9-386>.
 33. **Truong DT, Franzosa EA, Tickle TL, Scholz M, Weingart G, Pasoli E, Tett A, Huttenhower C, Segata N.** 2015. Metaphlan2 for enhanced metagenomic taxonomic profiling. *Nat Methods* **12**:902–903. <http://dx.doi.org/10.1038/nmeth.3589>.
 34. **Davenport CF, Neugebauer J, Beckmann N, Friedrich B, Kameri B, Kokott S, Paetow M, Siekmann B, Wieding-Drewes M, Wienhöfer M, Wolf S, Tümmler B, Ahlers V, Sprengel F.** 2012. Genometa—a fast and accurate classifier for short metagenomic shotgun reads. *PLoS One* **7**:e41224. <http://dx.doi.org/10.1371/journal.pone.0041224>.
 35. **Caporaso JG, Kuczynski J, Stombaugh J, Bittinger K, Bushman FD, Costello EK, Fierer N, Peña AG, Goodrich JK, Gordon JI, Huttley GA, Kelley ST, Knights D, Koenig JE, Ley RE, Lozupone CA, McDonald D, Muegge BD, Pirrung M, Reeder J.** 2010. QIIME allows analysis of high-throughput community sequencing data. *Nat Methods* **7**:335–336. <http://dx.doi.org/10.1038/nmeth.f303>.
 36. **Hunter S, Corbett M, Denise H, Fraser M, Gonzalez-Beltran A, Hunter C, Jones P, Leinonen R, McAnulla C, Maguire E, Maslen J, Mitchell A, Nuka G, Oisel A, Pesseat S, Radhakrishnan R, Rocca-Serra P, Scheremetjew M, Sterk P, Vaughan D, Cochrane G, Field D, Sansone SA.** 2014. EBI metagenomics—a new resource for the analysis and archiving of metagenomic data. *Nucleic Acids Res* **42**:D600–D606. <http://dx.doi.org/10.1093/nar/gkt961>.
 37. **Huson DH, Mitra S, Ruscheweyh HJ, Weber N, Schuster SC.** 2011. Integrative analysis of environmental sequences using MEGAN4. *Genome Res* **21**:1552–1560. <http://dx.doi.org/10.1101/gr.120618.111>.
 38. **Dröge J, Gregor I, McHardy AC.** 2015. Taxator-tk: precise taxonomic assignment of metagenomes by fast approximation of evolutionary neighborhoods. *Bioinformatics* **31**:817–824. <http://dx.doi.org/10.1093/bioinformatics/btu745>.
 39. **Freitas TA, Li PE, Scholz MB, Chain PS.** 2015. Accurate read-based metagenome characterization using a hierarchical suite of unique signatures. *Nucleic Acids Res* **43**:e69. <http://dx.doi.org/10.1093/nar/gkv180>.
 40. **Stoye J, Evers D, Meyer F.** 1998. Rose: generating sequence families. *Bioinformatics* **14**:157–163. <http://dx.doi.org/10.1093/bioinformatics/14.2.157>.
 41. **Merkel D.** 2014. Docker: lightweight Linux containers for consistent development and deployment. *Linux J* **2014**:2.
 42. **Giardine B, Riemer C, Hardison RC, Burhans R, Elitski L, Shah P, Zhang Y, Blankenberg D, Albert I, Taylor J, Miller W, Kent WJ, Nekrutenko A.** 2005. Galaxy: a platform for interactive large-scale genome analysis. *Genome Res* **15**:1451–1455. <http://dx.doi.org/10.1101/gr.4086505>.
 43. **Goecks J, Nekrutenko A, Taylor J, Galaxy Team.** 2010. Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biol* **11**:R86. <http://dx.doi.org/10.1186/gb-2010-11-8-r86>.
 44. **Blankenberg D, Kuster GV, Coraor N, Ananda G, Lazarus R, Mangan M, Nekrutenko A, Taylor J.** 2010. Galaxy: a web-based genome analysis tool for experimentalists. *Curr Protoc Mol Biol* Chapter 19:Unit 19.10.1. <http://dx.doi.org/10.1002/0471142727.mb1910s89>.
 45. **Leslie S, Winney B, Hellenenthal G, Davison D, Boumertit A, Day T, Hutnik K, Royrvik EC, Cunliffe B, Lawson DJ, Falush D, Lawson DJ, Falush D, Freeman C, Pirinen M, Myers S, Robinson M, Donnelly P, Bodmer W.** 2015. The fine-scale genetic structure of the British population. *Nature* **519**:309–314. <http://dx.doi.org/10.1038/nature14230>.
 46. **Hellenenthal G, Busby GB, Band G, Wilson JF, Capelli C, Falush D, Myers S.** 2014. A genetic atlas of human admixture history. *Science* **343**:747–751. <http://dx.doi.org/10.1126/science.1243518>.
 47. **Chen SS, Donoho DL, Saunders MA.** 1998. Atomic decomposition by basis pursuit. *SIAM J Sci Comput* **20**:33–61. <http://dx.doi.org/10.1137/S1064827596304010>.
 48. **Lawson CL, Hanson RJ.** 1974. Solving least squares problems, vol. 161. Society for Industrial and Applied Mathematics, Philadelphia, PA.

49. **Saitou N, Nei M.** 1987. The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Mol Biol Evol* **4**:406–425.
50. **Felsenstein J.** 1985. Confidence limits on phylogenies: an approach using the bootstrap. *Evolution* **39**:783–791. <http://dx.doi.org/10.2307/2408678>.
51. **Tamura K, Stecher G, Peterson D, Filipinski A, Kumar S.** 2013. MEGA6: molecular evolutionary genetics analysis version 6.0. *Mol Biol Evol* **30**: 2725–2729. <http://dx.doi.org/10.1093/molbev/mst197>.
52. **Tamura K, Nei M, Kumar S.** 2004. Prospects for inferring very large phylogenies by using the neighbor-joining method. *Proc Natl Acad Sci U S A* **101**:11030–11035. <http://dx.doi.org/10.1073/pnas.0404206101>.