Optimizing taxonomic classification of marker gene sequences

Nicholas A. Bokulich[1#*], Benjamin D. Kaehler[2#*], Jai Ram Rideout[1], Matthew Dillon[1], Evan Bolyen[1], Rob Knight[3], Gavin A. Huttley[2#], J. Gregory Caporaso[1,4,#]

[1]The Pathogen and Microbiome Institute, Northern Arizona University, Flagstaff, AZ, USA

[2]Research School of Biology, Australian National University, Canberra, Australia

[3]Departments of Pediatrics and Computer Science & Engineering, and Center for Microbiome Innovation, University of California San Diego, La Jolla, CA, USA

[4]Department of Biological Sciences, Northern Arizona University, Flagstaff, AZ, USA

*These authors contributed equally

#Corresponding authors

Gregory Caporaso
Department of Biological Sciences
1298 S Knoles Drive
Building 56, 3rd Floor
Northern Arizona University
Flagstaff, AZ, USA
(303) 523-5485
(303) 523-4015 (fax)
Email: gregcaporaso@gmail.com

Nicholas Bokulich
The Pathogen & Microbiome Institute
PO Box 4073
Flagstaff, Arizona 86011-4073, USA
Email: nicholas.bokulich@nau.edu

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Benjamin Kaehler
Research School of Biology
46 Sullivans Creek Road,
The Australian National University,
Acton ACT 2601, Australia
Email: benjamin.kaehler@anu.edu.au

Gavin Huttley
Research School of Biology
46 Sullivans Creek Road,
The Australian National University,
Acton ACT 2601, Australia
Email: gavin.huttley@anu.edu.au

Abstract

Background: Taxonomic classification of marker-gene sequences is an important step in microbiome analysis. **Results**: We present q2-feature-classifier (https://github.com/qiime2/q2-feature-classifier), a QIIME 2 plugin containing several novel machine-learning and alignment-based taxonomy classifiers that meet or exceed classification accuracy of existing methods. We evaluated and optimized several commonly used taxonomic classification methods (RDP, BLAST, BLAST+, UCLUST) and several new methods (a scikit-learn naive Bayes machine-learning classifier, and VSEARCH and SortMeRNA alignment-based methods). **Conclusions**: Our results illustrate the importance of parameter tuning for optimizing classifier performance, and we make explicit

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

recommendations regarding parameter choices for a range of standard operating conditions. q2-feature-classifier and our evaluation framework, tax-credit, are both free, open-source, BSD-licensed packages available on GitHub.

Background

High-throughput sequencing technologies have transformed our ability to explore complex microbial communities, offering insight into microbial impacts on human health [1] and global ecosystems [2]. This is achieved most commonly by sequencing short, conserved marker genes amplified with 'universal' PCR primers, such as 16S rRNA for bacteria and archaea, or internal transcribed spacer (ITS) regions for fungi. Targeted marker-gene primers can also be used to profile specific taxa or functional groups, such as nifH genes [3]. These sequences often are compared against an annotated reference sequence database to determine the likely taxonomic origin of each sequence with as much specificity as possible. Accurate and specific taxonomic information is a crucial component of many experimental designs.

Challenges in this process include the short length of typical sequencing reads with current technology, sequencing and PCR errors [4], selection of appropriate marker genes that contain sufficient heterogeneity to differentiate target species but that are homogeneous enough in some regions to design broad-spectrum primers, quality of

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

reference sequence annotations [5], and selection of a method that accurately predicts the taxonomic affiliation of millions of sequences at minimal computational cost. Numerous methods have been developed for taxonomy classification of DNA sequences, but few have been directly compared in the specific case of short marker-gene sequences.

We introduce q2-feature-classifier, a QIIME 2 (https://qiime2.org/) plugin for taxonomy classification of marker-gene sequences. QIIME 2 is the successor to the QIIME [6] microbiome analysis package. The q2-feature-classifier plugin allows users to use any of the numerous machine-learning classifiers available in scikit-learn [7][8] for marker gene taxonomy classification, and currently provides two alignment-based taxonomy consensus classifiers based on BLAST+ [9] and vsearch [10]. We evaluate the latter two methods and the scikit-learn multinomial naive Bayes classifier (labelled “Naive Bayes” in the Results section) for the first time. We show that the classifiers provided in q2-feature-classifier match or outperform the classification accuracy of several widely-used methods for sequence classification, and that performance of the naive Bayes classifier can be significantly increased by providing it with information regarding expected taxonomic composition.

We also developed tax-credit (https://github.com/caporaso-lab/tax-credit/), an extensible computational framework for evaluating taxonomy classification accuracy. This framework streamlines the process of methods benchmarking by compiling multiple different test data sets, including mock communities [11] and simulated sequence reads. It additionally stores pre-computed results from previously evaluated methods, including the results presented here, and provides a framework for parameter sweeps and method

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

optimization. tax-credit could be used as an evaluation framework by other research groups in the future, or its raw data could be easily extracted for integration in another evaluation framework.

Results

We used tax-credit to optimize and compare multiple taxonomy classifiers. We evaluated two commonly used, pre-existing classifiers that are wrapped in QIIME 1 (RDP Classifier (version 2.2) [12], legacy BLAST (version 2.2.22) [13]), two QIIME 1 alignment-based consensus taxonomy classifiers (the default UCLUST classifier available in QIIME 1 (based on version 1.2.22q) [14], and SortMeRNA (version 2.0 29/11/2014) [15]), two alignment-based consensus taxonomy classifiers newly released in q2-feature-classifier (based on BLAST+ (version 2.6.0) [9] and vsearch (version 2.0.3) [10]), and a new multinomial naive Bayes machine-learning classifier in q2-feature-classifier (see materials and methods for information about q2-feature-classifier methods and source code availability). We performed parameter sweeps to determine optimal parameter configurations for each method.

Mock community evaluations

We first benchmarked classifier performance on mock communities, which are artificially constructed mixtures of microbial cells or DNA combined at known ratios [11].

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

We utilized 15 bacterial 16S rRNA mock communities and 4 fungal internal transcribed spacer (ITS) mock communities (Table 1) sourced from mockrobiota [11], a public repository for mock community data. Mock communities are useful for method benchmarking because: 1) unlike for simulated communities, they allow quantitative assessments of method performance under actual operating conditions, i.e., incorporating real sequencing errors that can be difficult to model accurately; and 2) unlike for natural community samples, the actual composition of a mock community is known in advance, allowing quantitative assessments of community profiling accuracy.

An additional priority was to test the effect of setting class weights on classification accuracy for the naive Bayes classifier implemented in q2-feature-classifier. In machine learning, class weights or prior probabilities are vectors of weights that specify the frequency at which each class is expected to be observed (and should be distinguished from the use of this term under Bayesian inference as a probability distribution of weights vectors). An alternative to setting class weights is to assume that each query sequence is equally likely to belong to any of the taxa that are present in the reference sequence database. This assumption, known as uniform class priors in the context of a naive Bayes classifier, is made by the RDP classifier [12], and its impact on marker-gene classification accuracy has yet to be validated. Making either assumption, that the class weights are uniform or known to some extent, will affect results and cannot be avoided. The mock communities have taxonomic abundances that are far from uniform over the set of reference taxonomies, as any real data set must. We can therefore use them to assess the impact of making assumptions regarding class weights. Where we have set the class

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

139 weights to the known taxonomic composition of a sample, we have labelled the results
140 "bespoke".

141 We evaluated classifier performance accuracy on mock community sequences
142 classified at taxonomic levels from class through species. Mock community sequences were
143 classified using the Greengenes 99% OTUs 16S rRNA or UNITE 99% OTUs ITS reference
144 sequences for bacterial and fungal mock communities, respectively. As expected,
145 classification accuracy decreased as classification depth increased, and all methods could
146 predict the taxonomic affiliation of mock community sequences down to genus level with
147 median F-measures exceeding 0.8 across all parameter sets (minimum: UCLUST F=0.81,
148 maximum: Naive Bayes Bespoke F=1.00) (Figure 1A). However, species affiliation was
149 predicted with much lower and more variable accuracy among method configurations
150 (median F-measure minimum: UCLUST F=0.42, maximum: Naive Bayes Bespoke F=0.95),
151 highlighting the importance of parameter optimization (discussed in more detail below).
152 Figure 1A illustrates line plots of mean F-measure at each taxonomic level, averaged across
153 all classifier configurations; hence, classifier performance is underestimated for some
154 classifiers that are strongly affected by parameter configurations or for which a wider
155 range of parameters were tested (e.g., Naive Bayes). Comparing only optimized methods
156 (i.e., the top-performing parameter configurations for each method), Naive Bayes Bespoke
157 achieved significantly higher F-measure (paired t-test $P < 0.05$) (Figure 1B), recall, taxon
158 detection rate, and taxon accuracy rate scores (Figure 1C) and lower Bray-Curtis
159 dissimilarity than all other methods (Figure 1D).

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Mock communities are necessarily simplistic, and cannot assess method performance across a diverse range of taxa. Sequences matching the expected mock community sequences are not removed from the reference database prior to classification, in order to replicate normal operating conditions and assess recovery of expected sequences. However, this approach may implicitly bias toward methods that find an exact match to the query sequences, and does not approximate well natural microbial communities in which few or no detected sequences exactly match the reference sequences. Hence, we performed simulated sequence read classifications (described below) to further test classifier performance.

Cross-validated taxonomy classification

Simulated sequence reads, derived from reference databases, allow us to assess method performance across a greater diversity of sequences than a single mock community generally encompasses. We first evaluated classifier performance using stratified k-fold cross-validation of taxonomy classification to simulated reads. The k-fold cross-validation strategy is modified slightly to account for the hierarchical nature of taxonomic classifications, which all of the classifiers in this study (with the exception of legacy BLAST) handle by assigning the lowest (i.e., most specific) taxonomic level where the classification surpasses some user-defined “confidence” or “consensus” threshold (see materials and methods). The modification is to truncate any expected taxonomy in each test set to the maximum level at which an instance of that taxonomy exists in the training set. Simulated reads were generated from Greengenes 99% OTUs 16S rRNA or UNITE 99% OTUs ITS

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

reference sequences with species-level annotations, and do not incorporate artificial sequencing errors (see materials and methods for more details). In this set of tests and below for novel taxa, the "bespoke" classifier had prior probabilities that were inferred from the training set each time it was trained.

Classification of cross-validated reads yielded similar results to mock community classification tests. For bacterial sequences, average classification accuracy for all methods declined from near-perfect scores at family level (median F-measure minimum: BLAST+ F=0.92, maximum legacy BLAST F=0.99), but still retained accurate scores at species level (median minimum: BLAST+ F=0.76, maximum SortMeRNA F=0.84), relative to some mock community data sets (Figure 2A). Fungal sequences exhibited similar performance, with the exception that mean BLAST+ and vsearch performance was markedly lower at all taxonomic levels, indicating high sensitivity to parameter configurations, and species-level F-measures were in general much lower (median minimum: BLAST+ F=0.17, maximum UCLUST F=0.45) than those of bacterial sequence classifications (Figure 2A). At species level, optimized UCLUST and SortMeRNA configurations achieved the highest F-measures for 16S rRNA simulated sequences (Figure 2B). UCLUST achieved the highest F-measure for ITS classification (F = 0.51). However, all optimized classifiers achieved similar F-measure ranges, with the exception of legacy BLAST for ITS sequences (Figure 2B).

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Novel taxa evaluations

Novel taxa classification offers a unique perspective on classifier behavior, assessing how classifiers perform when challenged with a "novel" clade that is not represented in the reference database. An ideal classifier should identify the nearest taxonomic lineage to which this taxon belongs, but no further. In this evaluation, a reference database is subsampled k times to generate query and reference sequence sets, as for cross-validated classification, but two important distinctions exist: 1) the reference database used for classification excludes any sequence that matches the taxonomic affiliation of the query sequences at taxonomic level *L*, the taxonomic rank at which classification is being attempted; and 2) this is performed at each taxonomic level, in order to assess classification performance when each method encounters a "novel" species, genus, family, et cetera.

Due to these differences, interpretation of novel taxa evaluation results is different from that of mock community and cross-validated classifications. For the latter, classification accuracy may be assessed at each taxonomic level for each classification result: mean classification accuracy at family level and species level evaluate the same results but focus on different taxonomic levels of classification. For novel taxa, however, different query and reference sequences are compiled for classification at each taxonomic level and separate classifications are performed for each. Hence, classifications at family and species level are independent events — one assesses how accurately each method

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

221 performs when it encounters a "novel" family that is not represented in the reference
222 database, the other when a "novel" species is encountered.

223 Novel taxa evaluations employ a suite of modified metrics, to provide more
224 information on what types of classification errors occur. Precision, recall, and F-measure
225 calculations at each taxonomic level *L* assess whether an accurate taxonomy classification
226 was made at level *L*-1: for example, a "novel" species should be assigned a genus, because
227 the correct species class is not represented within the reference database. Any species-
228 level classification in this scenario is an *overclassification* (affecting both recall and
229 precision). Overclassification is one of the key metrics for novel taxa evaluation, indicating
230 the degree to which novel sequences will be interpreted as known organisms. This
231 overclassification is often highly undesirable because it leads, for example, to the incorrect
232 classification of unknown but harmless environmental sequences as known pathogens.
233 Novel sequences that are classified within the correct clade, but to a less specific level than
234 *L,* are *underclassified* (affecting recall but not precision). Sequences that are classified into a
235 completely different clade are *misclassified* (affecting both recall and precision).

236 Precision, recall, and F-measure all gradually increase from average scores near 0.0
237 at class level, reaching peak scores at genus level for bacteria and species level for fungi
238 (Figure 3A-C). These trends are paired with gradual decreases in underclassification and
239 misclassification rates for all classification methods, indicating that all classifiers perform
240 poorly when they encounter sequences with no known match at the class, order, or family
241 levels (Figure 3D-F). At species level, UCLUST, BLAST+, and vsearch achieved significantly
242 better F-measures than all other methods for 16S rRNA classifications ($P < 0.05$) (Figure

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

3G). UCLUST achieved significantly better F-measures than all other methods for ITS classifications (Figure 3G). Over-, under-, and misclassification scores are less informative for optimizing classifiers for real use cases, as most methods could be optimized to yield near-zero scores for each of these metrics separately, but only through extreme configurations, leading to F-measures that would be unacceptable under any scenario. Note that all comparisons were made between methods optimized to maximize (or minimize) a single metric, and hence the configurations that maximize precision are frequently different from those that maximize recall or other metrics. This trade-off between different metrics is discussed in more detail below.

The novel taxa evaluation provides an estimate of classifier performance given a specific reference database, but its generalization is limited by the quality of the reference databases available and by the label-based approach used for partitioning and evaluation. Mislabeled and polyphyletic clades in the database, e.g. Clostridium group, increase the probability of (potentially erroneous) misclassification. A complementary analysis based on sequence similarity between a novel query and top reference hit could mitigate this issue. However, we choose to apply a label-based approach, as it better reflects the biological problem that users can expect to encounter; i.e., using a particular reference sequence database (which will contain some quantity of mislabeled and polyphyletic taxa inherent to currently available resources), how likely is a classifier to misclassify a taxonomic label?

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Multi-evaluation method optimization

The mock community and cross-validation classification evaluations yielded similar trends in configuration performance, but optimizing parameters choices for the novel taxa generally lead to suboptimal choices for the mock community and cross-validation tests (Figure 4). We sought to determine the relationship between method configuration performance for each evaluation, and use this information to select configurations that perform best across all evaluations. For 16S rRNA sequence species-level classification, method configurations that achieve maximum F-measures for mock and cross-validated sequences perform poorly for novel taxa classification (Figure 4B). Optimization is more straightforward for genus-level classification of 16S rRNA sequences (Figure 4A) and for fungal sequences (Figure 4C-D), for which configuration performance (measured as mean F-measure) is maximized by similar configurations among all three evaluations.

To identify optimal method configurations, we set accuracy score minimum thresholds for each evaluation by identifying natural breaks in the range of quality scores, selecting methods and parameter ranges that meet these criteria. Table 2 lists method configurations that maximize species-level classification accuracy scores for mock community, cross-validated, and novel taxa evaluations under several common operating conditions. “Balanced” configurations are recommended for general use, and are methods that maximize F-measure scores. “Precision” and “Recall” configurations maximize precision and recall scores, respectively, for mock, cross-validated, and novel-taxa classifications (Table 2). “Novel” configurations optimize F-measure scores for novel taxa

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

285 classification, and secondarily for mock and cross-validated performance (Table 2). These
286 configurations are recommended for use with sample types that are expected to contain
287 large proportions of unidentified species, for which overclassification is undesirable.
288 However, these configurations may not perform optimally for classification of known
289 species (i.e., underclassification rates will be higher). For fungi, the same configurations
290 recommended for "Precision" perform well for novel taxa classification (Table 2). For 16S
291 rRNA sequences, BLAST+, UCLUST, and vsearch consensus classifiers perform best for
292 novel taxa classification (Table 2).
293

294 ## Computational runtime

295 High-throughput sequencing platforms (and experiments) continue to yield increasing
296 sequence counts, which — even after quality filtering and dereplication or operational
297 taxonomic unit clustering steps common to most microbiome analysis pipelines — may
298 exceed thousands of unique sequences that need classification. Increasing numbers of
299 query sequences and references sequences may lead to unacceptable runtimes, and under
300 some experimental conditions the top-performing method (based on precision, recall, or
301 some other metric) may be insufficient to handle large numbers of sequences within an
302 acceptable time frame. For example, quick turnarounds may be vital under clinical
303 scenarios as microbiome evaluation becomes common clinical practice, or commercial
304 scenarios, when large sample volumes and client expectations may constrain turnaround
305 times and method selection.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

306 We assessed computational runtime as a linear function of 1) the number of query
307 sequences and 2) the number of reference sequences. Linear dependence is empirically
308 evident in Figure 5. For both of these metrics, the slope is the most important measure of
309 performance. The intercept indicates the amount of time taken to train the reference
310 sequences, load environmental variables, or other "setup" steps that will diminish in
311 significance as sequence counts grow, and hence are negligible.
312 UCLUST (0.000028 s/sequence), vsearch (0.000072 s/sequence), BLAST+
313 (0.000080 s/sequence), and legacy BLAST (0.000100 s/sequence) all exhibit shallow
314 slopes with increasing numbers of reference sequences. Naive Bayes (0.000483
315 s/sequence) and SortMeRNA (0.000543 s/sequence) yield moderately higher slopes, and
316 RDP (0.001696 s/sequence) demonstrates the steepest slope (Figure 5A). For runtime as a
317 function of query sequence count, UCLUST (0.002248 s/sequence), RDP (0.002920
318 s/sequence), and SortMeRNA (0.003819 s/sequence) have relatively shallow slopes
319 (Figure 5B). Naive Bayes (0.022984 s/sequence), BLAST+ (0.026222s/sequence) , and
320 vsearch (0.030190 s/sequence) exhibit greater slopes. Legacy BLAST (0.133292
321 s/sequence) yielded a slope magnitudes higher than other methods, rendering this method
322 impractical for large data sets.
323
324

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Discussion

We have developed and validated several machine-learning and alignment-based classifiers provided in q2-feature-classifier and benchmarked these classifiers, as well as other common classification methods, to evaluate their strengths and weaknesses across a range of parameter settings for each (Table 2).

Each classifier required some degree of optimization to define top-performing parameter configurations, with the sole exception of QIIME 1's legacy BLAST wrapper, which was unaffected by its only user-defined parameter, e-value, over a range of 10^{-10} to 1000. For all other methods, performance varied widely depending on parameter settings, and a single method could achieve among the worst performance with one configuration but among the best performance with another. Configurations greatly affected accuracy with mock community, cross-validated, and novel taxa evaluations, indicating that optimization is necessary under a variety of performance conditions, and optimization for one condition may not necessarily translate to another. Mock community and cross-validated evaluations exhibited similar results, but novel taxa evaluations selected different optimal configurations for most methods (Figure 4), indicating that configurations optimized to one condition, e.g., high-recall classification of known sequences, may be less suited for other conditions, e.g., classification of novel sequences. Table 2 lists the top-performing configuration for each method for several standard performance conditions.

Optimal configurations also varied among different evaluation metrics. Precision and recall, in particular, exhibited some mutual opposition, such that methods increasing

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

precision reduced recall. For this reason, F-measure, the harmonic mean of precision and recall, is a useful metric for choosing configurations that are well balanced for average performance. "Balanced" method configurations — which maximize F-measure scores for mock, cross-validated, and novel taxa evaluations (Table 2) — are best suited for a wide range of user conditions. The naive Bayes classifier with kmer lengths of 6 or 7 and confidence = 0.7 (or confidence ≥ 0.9 if using bespoke class weights), RDP with confidence = 0.6-0.7, and UCLUST (minimum consensus = 0.51, minimum similarity = 0.9, max accepts = 3) perform best under these conditions (Table 2). Performance is dramatically improved using bespoke class weights for 16S rRNA sequences (Figure 4A-B), though this approach is developmental and only applicable when the expected composition of samples is known in advance (a scenario that is becoming increasingly common with the increasing quantity of public microbiome data, and which could be aided by microbiome data sharing resources such as Qiita (http://qiita.microbio.me)). For ITS sequences, the naive Bayes classifier with kmer lengths of 6 or 7 and confidence ≥ 0.9, or RDP with confidence = 0.7-0.9, perform best, and the effects of bespoke class weights are less pronounced (Figure 4C-D).

However, some users may require high-precision classifiers when false-positives may be more damaging to the outcome, e.g., for detection of pathogens in a sample. Precision scores are maximized by naive Bayes and RDP classifiers with high confidence settings (Table 2). Optimizing for precision will significantly damage recall by yielding a high number of false negatives.

Other users may require high-recall classifiers when false-negatives and underclassification hinder interpretation, but false positives (mostly overclassification to a

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

closely related species) are less damaging. For example, in environments with high numbers of unidentified species, a high-precision classifier may yield large numbers of unclassified sequences; in such cases, a second pass with a high-recall configuration (Table 2) may provide useful inference of what taxa are most similar to these unclassified sequences. When recall is optimized, precision tends to suffer slightly (leading to similar F-measure scores to "balanced" configurations) but novel taxa classification accuracy is minimized, as these configurations tend to overclassify (Table 2). Any user prioritizing recall ought to be aware of and acknowledge these risks, e.g., when sharing or publishing their results, and understand that many of the species-level classifications may be wrong, particularly if the samples are expected to contain many uncharacterized species. For 16S rRNA sequences, naive Bayes bespoke classifiers with kmer lengths between 12-32 and confidence = 0.5 yield maximal recall scores, but RDP (confidence = 0.5) and naive Bayes (uniform class weights, confidence = 0.5, kmer length = 11, 12, or 18) also perform well (Table 2). Fungal recall scores are maximized by the same configurations recommended for "Balanced" classification, i.e., naive Bayes classifiers with kmer lengths between 6-7 and confidence between 0.92-0.98, or RDP with confidence between 0.7-0.9 (Table 2).

Runtime requirements may also be the chief concern dictating method selection for some users. QIIME 1's UCLUST wrapper provides the fastest runtime while still achieving reasonably good performance for most evaluations; Naive Bayes, RDP, and BLAST+ also delivered reasonably low runtime requirements, and outperform UCLUST on most other evaluation metrics.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Conclusions

The classification methods provided in q2-feature-classifier will support improved taxonomy classification of marker-gene sequences, and are released as a free, open-source plugin for use with QIIME 2. We demonstrate that these methods perform as well as or better than other leading taxonomy classification methods on a number of performance metrics. The naive Bayes, vsearch, and BLAST+ consensus classifiers described here are released for the first time in QIIME 2, with optimized "balanced" configurations (Table 2) set as defaults.

We also present the results of a benchmark of several widely used taxonomy classifiers, and recommend the top-performing methods and configurations for the most common user scenarios. Our recommendations for "balanced" methods (Table 2) will be appropriate for most users who are classifying 16S rRNA or fungal ITS sequences, but other users may prioritize high-precision (low false-positive) or high-recall (low false-negative) methods.

We have also shown that great potential exists for improving the accuracy of taxonomy classifications by appropriately setting class weights for the machine learning classifiers. Currently, no tools exist that allow users to generate appropriate values for these class weights in real applications. Compiling appropriate class weights for different sample types could be a promising approach to further improve taxonomic classification of marker gene sequence reads.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Methods

Mock communities

All mock communities were sourced from mockrobiota [11]. Raw fastq files were demultiplexed and processed using tools available in QIIME 2 (version 2017.4) (https://qiime2.org/). Reads were demultiplexed with q2-demux (https://github.com/qiime2/q2-demux) and quality filtered and dereplicated with q2-dada2 [4]. Representative sequence sets for each dada2 sequence variant were used for taxonomy classification with each classification method.

The inclusion of multiple mock community samples is important to avoid overfitting; optimizing method performance to a small set of data could result in overfitting to the specific community compositions or conditions under which those data were generated, which reduces the robustness of the classifier.

Cross-validated simulated reads

The simulated reads used here were derived from the reference databases using the "Cross-validated classification performance" notebooks in our project repository. The reference databases where either Greengenes or UNITE (99% OTUs) that were cleaned according to taxonomic label to remove sequences with ambiguous or null labels.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Reference sequences were trimmed to simulate amplification using standard PCR primers and slice out the first 250 bases downstream (3') of the forward primer. The exact sequences were used for cross validation, and were not altered to simulate any sequencing error. The bacterial primers used were 515F/806R [17], and the fungal primers used were BITSf/B58S3r [18]. Each database was stratified by taxonomy and 10-fold randomised cross-validation data sets were generated using scikit-learn's library functions. Where a taxonomic label had less than 10 instances, taxonomies were amalgamated to make sufficiently large strata. If, as a result, a taxonomy in any test set was not present in the corresponding training set, the expected taxonomy label was truncated to the nearest common taxonomic rank observed in the training set (e.g., *Lactobacillus casei* would become *Lactobacillus*). The notebook detailing simulated read generation (for both cross-validated and novel taxa reads) prior to taxonomy classification is available at https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/novel-taxa/dataset-generation.ipynb.

Classification performance was also slightly modified from a standard machine-learning scenario as the classifiers in this study are able to refuse classification if they are not confident above a taxonomic level for a given sample. This also accommodates the taxonomy truncation that we performed for this test. The methodology was consistent with that used below for novel taxa, but we defer this description to the next section.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

"Novel taxa" simulation analysis

"Novel taxa" analysis was performed to test the performance of classifiers when assigning taxonomy to sequences that are not represented in a reference database, e.g., as a simulation of what occurs when a method encounters an undocumented species. In this analysis, simulated amplicons were filtered from those used for the cross-validation analysis. For all sequences present in each test set, sequences sharing taxonomic affiliation at a given taxonomic level L (e.g., to species level) in the corresponding training set were removed. Taxa are stratified among query and test sets such that for each query taxonomy at level L, no reference sequences match that taxonomy, but at least one reference sequence will match the taxonomic lineage at level L-1 (e.g., same genus but different species). An ideal classifier would assign taxonomy to the nearest common taxonomic lineage (e.g., genus), but would not "overclassify" to near neighbors (e.g., assign species-level taxonomy when species X is removed from the reference database). For example, a "novel" sequence representing the species *Lactobacillus brevis* should be classified as "*Lactobacillus*", without species-level annotation, in order to be considered a true positive in this analysis. As described above for cross-validated reads, these novel taxa simulated communities were also tested in both bacterial (B) and fungal (F) databases on simulated amplicons trimmed to simulate 250-nt sequencing reads.

Novel taxa classification performance is evaluated using precision, recall, F-measure, overclassification rates, underclassification rates, and misclassification rates for each taxonomic level (phylum to species), computed with the following modified

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

definitions (see below, *Performance analyses using simulated reads*, for full description of precision, recall, and F-measure calculations; these calculations use the modified definitions of true positive, false positive, and false negative as described here):

1) A true positive is considered the nearest correct lineage contained in the reference database. For example, if *Lactobacillus brevis* is removed from the reference database and used as a query sequence, the only correct taxonomy classification would be "*Lactobacillus*", without species-level classification.
2) A false positive would be either an classification to a different *Lactobacillus* species (*Overclassification*), or any genus other than *Lactobacillus* (*Misclassification*).
3) A false negative occurs if an expected taxonomy classification (e.g., "*Lactobacillus*") is not observed in the results. Note that this will be the modified taxonomy expected when using a naive reference database, and is not the same as the true taxonomic affiliation of a query sequence in the novel taxa analysis. A false negative results from misclassification, overclassification, or when the classification contains the correct basal lineage, but does not assign a taxonomy label at level L (*Underclassification)*. E.g., classification as "*Lactobacillaceae*", but no genus-level classification.

Taxonomy classification

Representative sequences for all analyses (mock community, cross-validated, and novel taxa) were classified taxonomically using the following taxonomy classifiers and setting sweeps:

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

1. q2-feature-classifier multinomial naive Bayes classifier. Varied k-mer length in {4, 6, 7, 8, 9, 10, 11, 12, 14, 16, 18, 32} and confidence threshold in {0, 0.5, 0.7, 0.9, 0.92, 0.94, 0.96, 0.98, 1}.

2. BLAST+ [9] local sequence alignment, followed by consensus taxonomy classification implemented in q2-feature-classifier. Varied max accepts from 1 to 100; percent identity from 0.80 to 0.99; and minimum consensus from 0.51 to 0.99. See description below.

3. vsearch [10] global sequence alignment, followed by consensus taxonomy classification implemented in q2-feature-classifier. Varied max accepts from 1 to 100; percent identity from 0.80 to 0.99; and minimum consensus from 0.51 to 0.99. See description below.

4. Ribosomal Database Project (RDP) naïve Bayesian classifier [12] (QIIME1 wrapper), with confidence thresholds between 0.0 to 1.0 in steps of 0.1.

5. Legacy BLAST [13] (QIIME1 wrapper) varying e-value thresholds from 1e-9 to 1000.

6. SortMeRNA [15] (QIIME1 wrapper) varying minimum consensus fraction from 0.51 to 0.99; similarity from 0.8 to 0.9; max accepts from 1 to 10; and coverage from 0.8 to 0.9.

7. UCLUST [14] (QIIME1 wrapper) varying minimum consensus fraction from 0.51 to 0.99; similarity from 0.8 to 0.9; and max accepts from 1 to 10.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

With the exception of the UCLUST classifier, we have only benchmarked the performance of open-source, free, marker-gene-agnostic classifiers, i.e., those that can be trained/aligned on a reference database of *any* marker gene. Hence, we excluded classifiers that can only assign taxonomy to a particular marker gene (e.g., only bacterial 16S rRNA genes) and those that rely on specialized or unavailable reference databases and cannot be trained on other databases, effectively restricting their use for other marker genes and custom databases.

Classification of bacterial/archaeal 16S rRNA sequences was made using the Greengenes reference sequence database (13_8 release) [5] preclustered at 99% ID, with V4 domain amplicons extracted using primers 515f/806r with q2-feature-classifier's `extract_reads` method. Classification of fungal ITS sequences was made using the UNITE database (version 7.1 QIIME developer release) [19] preclustered at 99% ID. For the cross validation and novel taxa tests we prefiltered to remove sequences with incomplete or ambiguous taxonomies (containing the substrings 'unknown', 'unidentified', or '_sp' or terminating at any level with '__').

The notebooks detailing taxonomy classification sweeps of mock communities are available at https://github.com/caporaso-lab/tax-credit/tree/0.2.2/ipynb/mock-community. Cross-validated read classification sweeps are available at https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/cross-validated/taxonomy-assignment.ipynb. Novel taxa classification sweeps are available at https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/novel-taxa/taxonomy-assignment.ipynb.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Runtime analyses

The tax-credit framework employs two different runtime metrics: as a function of 1) the number of query sequences or 2) the number of reference sequences. Taxonomy classifier runtimes were logged while performing classifications of pseudorandom subsets of 1, 2,000, 4,000, 6,000, 8,000, and 10,000 sequences from the Greengenes 99% OTU database. Each subset was drawn once then used for all of the tests as appropriate. All runtimes were computed on the same Linux workstation (Ubuntu 16.04.2 LTS, Intel Xeon CPU E7-4850 v3 @ 2.20GHz, 1TB memory). The exact commands used for runtime analysis are presented in the "Runtime analyses" notebook in the project repository (https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/runtime/analysis.ipynb).

Performance analyses using simulated reads

Cross-validated and novel taxa reads are evaluated using the classic precision, recall, and F-measure metrics [5] (novel taxa use the standard calculations as described below, but modified definitions for true positive (*TP*), false positive (*FP*), and false negative (*FN*), as described above for novel taxa).

Precision, recall, and F-measure are calculated as follows:

- *Precision = TP/(TP+FP)* or the fraction of sequences that were classified correctly at level L.
- *Recall = TP/(TP+FN)* or the fraction of expected taxonomic labels that were predicted at level L.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

- *F-measure* = 2 × Precision × Recall / (Precision + Recall), or the harmonic mean of precision and recall.

The Jupyter notebook detailing commands used for evaluation of cross-validated read classifications is available at https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/cross-validated/evaluate-classification.ipynb. The notebook for evaluation of novel taxa classifications is available at https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/novel-taxa/evaluate-classification.ipynb.

Performance analyses using mock communities

The Jupyter notebook detailing commands used for evaluation of mock communities, including the three evaluation types described below, is available at https://github.com/caporaso-lab/tax-credit/blob/0.2.2/ipynb/mock-community/evaluate-classification-accuracy.ipynb.

Precision and Recall

Classic precision, recall, and F-measure are used to calculate mock community classification accuracy, using the definitions given above for simulated reads. These metrics require knowing the expected classification of each sequence, which we determine by performing a gapless alignment between each representative sequence in the mock community and the marker-gene sequences of each microbial strain added to the mock community. These "expected sequences" are provided for the mock communities in mockrobiota [11]. Representative sequences are assigned the taxonomy of the best

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

alignment, and any representative sequence with more than 3 mismatches to the expected sequences are excluded from precision/recall calculations. If a representative sequence aligns to more than one expected sequence equally well, all top hits are accepted as the "correct" classification. This scenario is rare and typically only occurred when different strains of the same species were added to the same mock community to intentionally produce this challenge (e.g., for mock-12 as described by [4]). Precision, recall, and F-measure are then calculated by comparing the "expected" classification for each mock community sequence to the classifications predicted by each taxonomy classifier using the full reference databases, as described above.

Taxon accuracy rate and taxon detection rate

Taxon accuracy rate (TAR) and taxon detection rate (TDR) are used for qualitative compositional analyses of mock communities. As the true taxonomy labels for each sequence in a mock community are not known with absolute certainty, TAR and TDR are useful alternatives to precision and recall that instead rely on the presence/absence of expected taxa, or microbiota that are intentionally added to the mock community. In practice, TAR/TDR are complementary metrics to precision/recall and should provide similar results if the expected classifications for mock community representative sequences are accurate.

At a given taxonomic level, a classification is a:

- true positive (*TP*), if that taxon is both observed and expected.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

- false positive (*FP*), if that taxon is observed but not expected.
- false negative (*FN*), if a taxon is expected but not observed.

These are used to calculate TAR and TDR as:

- *TAR = TP/(TP+FP)* or the fraction of observed taxa that were expected at level L.
- *TDR = TP/(TP+FN)* or the fraction of expected taxa that are observed at level L.

Bray-Curtis Dissimilarity

Bray-Curtis dissimilarity [20] is used to measure the degree of dissimilarity between two samples as a function of the abundance of each species label present in each sample, treating each species as equally related. This is a useful metric for evaluating classifier performance by assessing the relative distance between each predicted mock community composition (abundance of taxa in a sample based on results of a single classifier) and the expected composition of that sample. For each classifier, Bray-Curtis distances between the expected and observed taxonomic compositions are calculated for each sample in each mock community dataset; this yields a single expected-observed distance for each individual observation. The distance distributions for each method are then compared statistically using paired or unpaired t-tests to assess whether one method (or configuration) performs consistently better than another.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

New taxonomy classifiers

We describe q2-feature-classifier (https://github.com/qiime2/q2-feature-classifier), a plugin for QIIME 2 (https://qiime2.org/) that performs multi-class taxonomy classification of marker-gene sequence reads. In this work we compare the consensus BLAST+ and vsearch methods and the naive Bayes scikit-learn classifier. The software is free and open-source.

Machine learning taxonomy classifiers

The q2-feature-classifier plugin allows users to apply any of the suite of machine learning classifiers available in scikit-learn (http://scikit-learn.org) to the problem of taxonomy classification of marker-gene sequences. It functions as a lightweight wrapper that transforms the problem into a standard document classification problem. Advanced users can input any appropriate scikit-learn classifier pipeline, which can include a range of feature extraction and transformation steps as well as specifying a machine learning algorithm.

The plugin provides a default method which is to extract k-mer counts from reference sequences and train the scikit-learn multinomial naive Bayes classifier, and it is this method that we test extensively here. Specifically, the pipeline consists of a sklearn.feature_extraction.text.HashingVectorizer feature extraction step followed by a sklearn.naive_bayes.MultinomialNB classification step. The use of a hashing feature extractor allows the use of significantly longer k-mers than the 8-mers that are used by

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

RDP Classifier, and we tested up to 32-mers. Like most scikit-learn classifiers, we are able to set class weights when training the multinomial naive Bayes classifiers. In the naive Bayes setting, setting class weights means that class priors are not derived from the training data or set to be uniform, as they are for the RDP Classifier. For more detail on how class weights enter the calculations please refer to the scikit-learn User Guide (http://scikit-learn.org).

In most settings, it is highly unlikely that the assumption of uniform weights is correct. That assumption is that each of the taxa in the reference database is equally likely to appear in each sample. Setting class weights to more realistic values can greatly aid the classifier in making more accurate predictions, as we show in this work. When testing the mock communities we made use of the fact that the sequence compositions were known *a priori* for the bespoke classifier. For the simulated reads studies, we allowed the classifier to set the class weights from the class frequencies observed in each training set for the bespoke classifier.

For this study, we performed two parameter sweeps on the mock communities: an initial broad sweep to optimize feature extraction parameters and then a more focussed sweep to optimise k-mer length and confidence parameter settings. These sweeps included varying the assumptions regarding class weights. The focussed sweeps were also performed for the cross-validated and novel taxa evaluations, but only for the assumption of uniform class

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

653 priors. The results for the focussed sweeps across all data sets are those which are
654 compared against the other classifiers in this work.
655
656 The broad sweeps used a modified scikit-learn pipeline which consisted of the
657 sklearn.feature_extraction.text.HashingVectorizer, followed by the
658 sklearn.feature_extraction.text.TfidfTransformer, then the
659 sklearn.naive_bayes.MultinomialNB. We performed a full grid search over the parameters
660 shown in Table 3. The conclusion from the initial sweep was that the TfidfTransformer step
661 did not significantly improve classification, that n_features should be set to 8192, feature
662 vectors should be normalised using L2 normalisation and that the alpha parameter for the
663 naive Bayes classifier should be set to 0.001. Please see https://github.com/caporaso-
664 lab/tax-credit/blob/0.2.2/ipynb/mock-community/evaluate-classification-accuracy-nb-
665 extra.ipynb for details.

666 **Consensus taxonomy alignment-based classifiers**

667
668 Two new classifiers implemented in q2-feature-classifier perform consensus taxonomy
669 classification based on alignment of a query sequence to a reference sequence. The
670 methods `classify_consensus_vsearch` and `classify_consensus_blast` use the global aligner
671 vsearch [10] or the local aligner BLAST+ [9], respectively, to return up to `maxaccepts`
672 reference sequences that align to the query with at least `perc_identity` similarity. A
673 consensus taxonomy is then assigned to the query sequence by determining the taxonomic

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

674 lineage on which at least `min_consensus` of the aligned sequences agree. This consensus
675 taxonomy is truncated at the taxonomic level at which less than `min_consensus` of
676 taxonomies agree. For example, if a query sequence is classified with `maxaccepts=3`,
677 `min_consensus=0.51`, and the following top hits:
678
679 `k__Bacteria; p__Firmicutes; c__Bacilli; o__Lactobacillales; f__Lactobacillaceae;`
680 `g__Lactobacillus; s__brevis`
681 `k__Bacteria; p__Firmicutes; c__Bacilli; o__Lactobacillales; f__Lactobacillaceae;`
682 `g__Lactobacillus; s__brevis`
683 `k__Bacteria; p__Firmicutes; c__Bacilli; o__Lactobacillales; f__Lactobacillaceae;`
684 `g__Lactobacillus; s__delbrueckii`
685
686 The taxonomy label assigned will be `k__Bacteria; p__Firmicutes; c__Bacilli;`
687 `o__Lactobacillales; f__Lactobacillaceae; g__Lactobacillus; s__brevis`. However, if
688 `min_consensus=0.99`, the taxonomy label assigned will be `k__Bacteria; p__Firmicutes;`
689 `c__Bacilli; o__Lactobacillales; f__Lactobacillaceae; g__Lactobacillus`.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

Declarations

Ethics approval and consent to participate

Not applicable

Consent for publication

Not applicable

Availability of data and materials

Mock community sequence data used in this study are publicly available in mockrobiota [11] under the study identities listed in Table 1. All other data generated in this study, and all new software, is available in our GitHub repositories under the BSD license. The tax-credit repository can be found at: https://github.com/caporaso-lab/tax-credit, and static versions of all analysis notebooks, which contain all code and analysis results, can be viewed there. The q2-feature-classifier repository can be accessed at https://github.com/qiime2/q2-feature-classifier; as a QIIME2 core plugin, it is automatically installed any time QIIME2 (https://qiime2.org/) is installed.

Project name: q2-feature-classifier
Project home page: https://github.com/qiime2/q2-feature-classifier
Operating system(s): macOS X, Linux
Programming language: Python
Other requirements: QIIME2
License: BSD-3-Clause
Any restrictions to use by non-academics: None

Project name: tax-credit

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

716 **Project home page**: https://github.com/caporaso-lab/tax-credit
717 **Operating system(s)**: macOS X, Linux
718 **Programming language**: Python
719 **Other requirements**: None (QIIME2 required for some optional functions)
720 **License**: BSD-3-Clause
721 **Any restrictions to use by non-academics**: None
722
723

724 Funding

725 This work was funded in part by National Science Foundation award 1565100 to JGC and
726 RK, awards from the Alfred P. Sloan Foundation to JGC and RK, and National Health and
727 Medical Research Council of Australia award APP1085372 to GAH and JGC. These funding
728 bodies had no role in the design of the study, the collection, analysis, or interpretation of
729 data, or in writing the manuscript.

730 Acknowledgments

731 The authors thank Stephen Gould and Cheng Soon Ong for advice on machine learning
732 optimisation.

733 Authors' Contributions

734 NAB, RK, and JGC conceived and designed tax-credit. NAB, BDK, JGC, and JRR contributed
735 to tax-credit. BDK, MD, JGC, and NAB contributed to q2-feature-classifier. BDK, JGC, MD,
736 JRR, and EB provided QIIME 2 integration with q2-feature-classifier. JGC and GAH provided
737 materials and support. NAB, BDK, JGC, and GAH wrote the manuscript with input from all
738 co-authors.

Competing Interests

The authors declare that they have no competing interests.

Tables and Figures

Table 1. Mock communities currently integrated in tax-credit.

Study ID*	target-gene	Species	Strains	Citation
mock-1	16S	46	48	[21]
mock-2	16S	46	48	[21]
mock-3	16S	21	21	[21]
mock-4	16S	21	21	[21]
mock-5	16S	21	21	[21]
mock-7	16S	67	67	[22]
mock-8	16S	67	67	[11]
mock-9	ITS	13	16	[11]
mock-10	ITS	13	16	[11]
mock-12	16S	26	27	[4]
mock-16	16S	56	59	[23]
mock-18	16S	15	15	[24]
mock-19	16S	15	27	[24]
mock-20	16S	20	20	[25]
mock-21	16S	20	20	[25]
mock-22	16S	20	20	[25]
mock-23	16S	20	20	[25]
mock-24	ITS	8	8	[26]
mock-26	ITS	11	11	[27]

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

744 *All studies are available on mockrobiota [11] at https://github.com/caporaso-

745 lab/mockrobiota/tree/master/data/[studyID]

746

747 Table 2. Optimized methods configurations for standard operating conditions.

				Mock			Cross-validated			Novel taxa			
Target	Condition	Method	Parameters	F	P	R	F	P	R	F	P	R	Threshold
16S rRNA	Balanced	NB-bespoke	[6,6]:0.9	0.705	0.98	0.582	0.827	0.931	0.744	0.165	0.243	0.125	F = (0.49, 0.8, 0.1)
			[6,6]:0.92	0.705	0.98	0.581	0.825	0.936	0.737	0.165	0.251	0.123	F = (0.7, 0.8, 0.15)
			[6,6]:0.94	0.703	0.98	0.579	0.822	0.942	0.729	0.162	0.259	0.118	
			[7,7]:0.92	0.712	0.978	0.592	0.831	0.931	0.751	0.151	0.221	0.115	
			[7,7]:0.94	0.708	0.978	0.586	0.829	0.936	0.743	0.157	0.239	0.117	
		naive-bayes	[7,7]:0.7	0.495	0.797	0.38	0.819	0.886	0.761	0.115	0.138	0.099	
		rdp	0.6	0.564	0.798	0.457	0.815	0.868	0.768	0.102	0.128	0.084	
			0.7	0.55	0.799	0.438	0.812	0.892	0.746	0.124	0.173	0.096	
		uclust	0.51:0.9:3	0.498	0.746	0.392	0.846	0.876	0.817	0.154	0.201	0.126	
	Precision	NB-bespoke	[6,6]:0.98	0.676	0.987	0.537	0.803	0.956	0.692	0.163	0.303	0.111	P = (0.94, 0.95, 0.25)
			[7,7]:0.98	0.687	0.98	0.551	0.815	0.951	0.713	0.164	0.283	0.115	
		rdp	1	0.239	0.941	0.16	0.632	0.968	0.469	0.12	0.457	0.069	
	Recall	NB-bespoke	[12,12]:0.5	0.754	0.8	0.721	0.815	0.83	0.801	0.053	0.058	0.049	R = (0.47, 0.75, 0.04)
			[14,14]:0.5	0.758	0.802	0.726	0.811	0.826	0.797	0.052	0.057	0.048	R = (0.7, 0.75, 0.04)
			[16,16]:0.5	0.755	0.785	0.732	0.808	0.825	0.792	0.052	0.058	0.047	
			[18,18]:0.5	0.772	0.803	0.748	0.805	0.823	0.789	0.055	0.061	0.05	
			[32,32]:0.5	0.937	0.966	0.913	0.788	0.818	0.76	0.054	0.067	0.045	
		naive-bayes	[11,11]:0.5	0.567	0.77	0.479	0.793	0.82	0.768	0.059	0.065	0.055	
			[12,12]:0.5	0.567	0.769	0.479	0.79	0.816	0.765	0.059	0.064	0.055	
			[18,18]:0.5	0.564	0.764	0.477	0.779	0.807	0.753	0.057	0.063	0.051	
		rdp	0.5	0.577	0.791	0.48	0.816	0.848	0.787	0.068	0.079	0.06	
	Novel	blast+	10:0.51:0.8	0.436	0.723	0.325	0.816	0.896	0.749	0.225	0.332	0.171	F = (0.4, 0.8, 0.2)
		uclust	0.76:0.9:5	0.467	0.775	0.348	0.84	0.938	0.76	0.219	0.358	0.158	
		vsearch	10:0.51:0.8	0.45	0.74	0.342	0.814	0.891	0.75	0.226	0.333	0.171	
			10:0.51:0.9	0.45	0.74	0.342	0.82	0.896	0.755	0.219	0.338	0.162	
Fungi	Balanced	naive-bayes	[6,6]:0.94	0.874	0.935	0.827	0.481	0.57	0.416	0.374	0.438	0.327	F = (0.85, 0.45, 0.37)

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

			[6,6]:0.96	0.874	0.935	0.827	0.495	0.597	0.423	0.399	0.473	0.344	
			[6,6]:0.98	0.874	0.935	0.827	0.505	0.629	0.423	0.426	0.52	0.361	
			[7,7]:0.98	0.874	0.935	0.827	0.485	0.596	0.409	0.388	0.47	0.33	
		NB-bespoke	[6,6]:0.94	0.928	0.968	0.915	0.48	0.567	0.416	0.371	0.433	0.325	
			[6,6]:0.96	0.928	0.968	0.915	0.491	0.59	0.42	0.393	0.466	0.34	
			[6,6]:0.98	0.927	0.97	0.913	0.504	0.624	0.422	0.421	0.512	0.358	
			[7,7]:0.98	0.935	0.97	0.921	0.487	0.596	0.412	0.386	0.466	0.329	
		rdp	0.7	0.929	0.939	0.922	0.479	0.572	0.413	0.382	0.451	0.332	
			0.8	0.924	0.939	0.915	0.507	0.633	0.422	0.434	0.534	0.366	
			0.9	0.922	0.937	0.913	0.517	0.698	0.411	0.47	0.617	0.379	
	Precision	naive-bayes	[6,6]:0.98	0.874	0.935	0.827	0.505	0.629	0.423	0.426	0.52	0.361	P = (0.92, 0.6, 0.3)
		NB-bespoke	[6,6]:0.98	0.927	0.97	0.913	0.504	0.624	0.422	0.421	0.512	0.358	
		rdp	0.8	0.924	0.939	0.915	0.507	0.633	0.422	0.434	0.534	0.366	
			0.9	0.922	0.937	0.913	0.517	0.698	0.411	0.47	0.617	0.379	
			1	0.821	0.943	0.742	0.461	0.81	0.322	0.459	0.774	0.327	
	Recall	NB-bespoke	[6,6]:0.92	0.938	0.971	0.924	0.467	0.544	0.409	0.353	0.407	0.312	R = (0.9, 0.4, 0.3)
			[6,6]:0.94	0.928	0.968	0.915	0.48	0.567	0.416	0.371	0.433	0.325	
			[6,6]:0.96	0.928	0.968	0.915	0.491	0.59	0.42	0.393	0.466	0.34	
			[6,6]:0.98	0.927	0.97	0.913	0.504	0.624	0.422	0.421	0.512	0.358	
			[7,7]:0.96	0.935	0.969	0.921	0.47	0.56	0.404	0.357	0.422	0.31	
			[7,7]:0.98	0.935	0.97	0.921	0.487	0.596	0.412	0.386	0.466	0.329	
		rdp	0.7	0.929	0.939	0.922	0.479	0.572	0.413	0.382	0.451	0.332	
			0.8	0.924	0.939	0.915	0.507	0.633	0.422	0.434	0.534	0.366	
			0.9	0.922	0.937	0.913	0.517	0.698	0.411	0.47	0.617	0.379	
	Novel	naive-bayes	[6,6]:0.98	0.874	0.935	0.827	0.505	0.629	0.423	0.426	0.52	0.361	F = (0.85, 0.45, 0.4)
		NB-bespoke	[6,6]:0.98	0.927	0.97	0.913	0.504	0.624	0.422	0.421	0.512	0.358	
		rdp	0.8	0.923	0.939	0.915	0.507	0.633	0.422	0.434	0.534	0.366	
			0.9	0.921	0.937	0.913	0.517	0.698	0.411	0.47	0.617	0.379	

748

749 [a]F = F-measure, P = precision, R = recall
750 [b]Naive Bayes parameters: k-mer range, confidence
751 [c]RDP parameters: confidence
752 [d]BLAST+/vsearch parameters: max accepts, minimum consensus, minimum percent
753 identity
754 [e]UCLUST parameters: minimum consensus, similarity, max accepts

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

[f]Threshold describes the score cutoffs used to define optimal method ranges, in the format: [metric = (mock score, cross-validated score, novel-taxa score)]. If two cutoffs are given, the second indicates a higher cutoff used to select parameters for the developmental NB-bespoke method, and the configurations listed are the union of the two cutoffs: the second cutoff for selecting NB-bespoke, the first for selecting all other methods.

Table 3. Naive Bayes broad grid search parameters

Step	Parameter	Values
sklearn.feature_extraction.text.HashingVectorizer	n_features	1024, 8192, 65536
	ngram_range	[4,4], [8, 8], [16, 16], [4,16]
sklearn.feature_extraction.text.TfidfTransformer	norm	l1', 'l2', None
	usd_idf	True, False
sklearn.naive_bayes.MultinomialNB	alpha	0.001, 0.01, 0.1
	class_prior	None, array of class weights
post processing	confidence	0, 0.2, 0.4, 0.6, 0.8

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

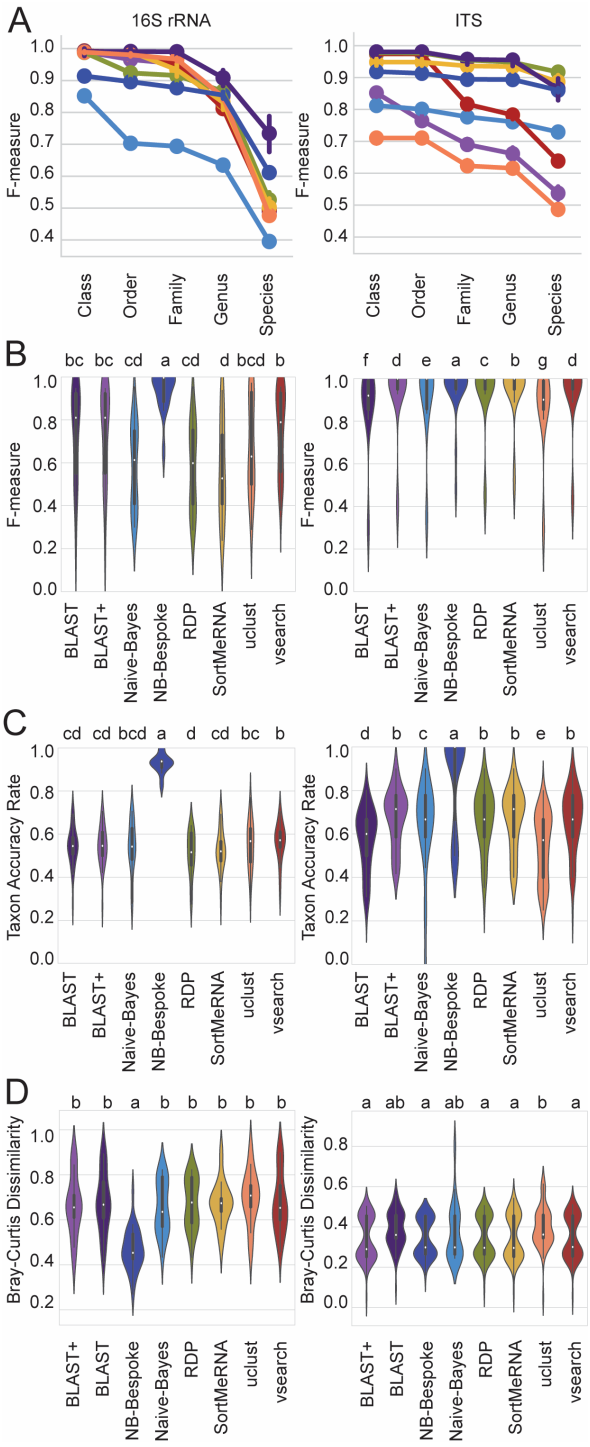

766

767

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

768 Figure 1. Classifier performance on mock community datasets for 16S rRNA sequences (left
769 column) and fungal ITS sequences (right column). A, Average F-measure for each taxonomy
770 classification method (averaged across all configurations and all mock community
771 datasets) from class to species level. Error bars = 95% confidence intervals. B, Average F-
772 measure for each optimized classifier (averaged across all mock communities) at species
773 level. C, Average taxon accuracy rate for each optimized classifier (averaged across all mock
774 communities) at species level. D, Average Bray-Curtis distance between the expected mock
775 community composition and its composition as predicted by each optimized classifier
776 (averaged across all mock communities) at species level. Violin plots show median (white
777 point), quartiles (black bars), and kernel density estimation (violin) for each score
778 distribution. Violins with different lower-case letters have significantly different means
779 (paired t-test false detection rate-corrected $P < 0.05$).

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

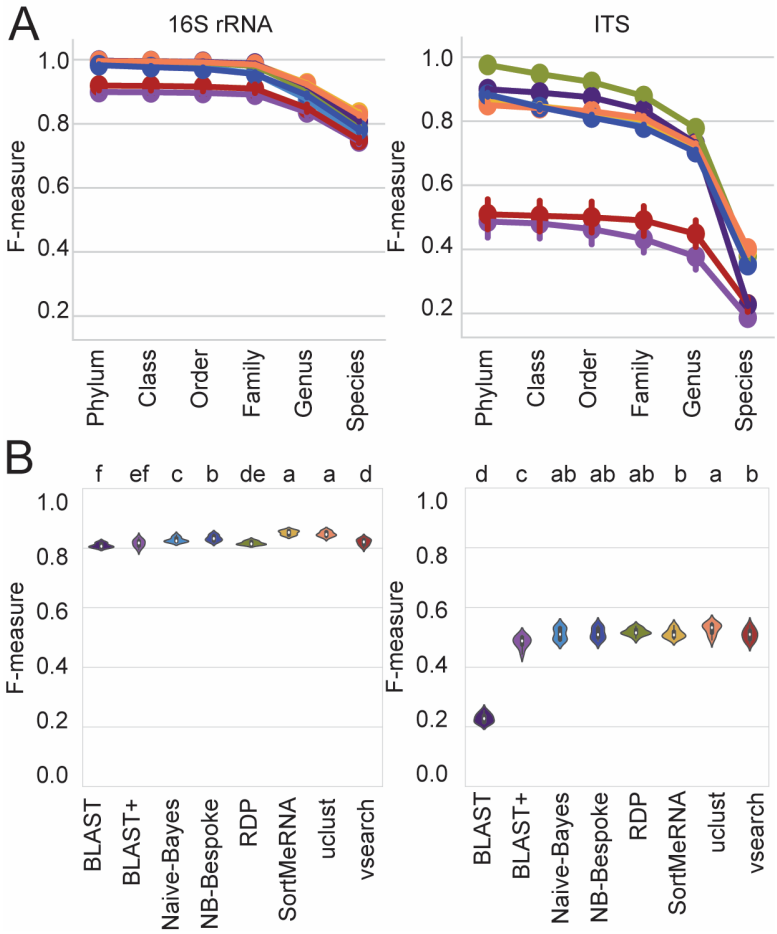

780

781 Figure 2. Classifier performance on cross-validated sequence datasets for 16S rRNA
782 sequences (left column) and fungal ITS sequences (right column). A, Average F-measure for
783 each taxonomy classification method (averaged across all configurations and all cross-
784 validated sequence datasets) from class to species level. Error bars = 95% confidence
785 intervals. B, Average F-measure for each optimized classifier (averaged across all cross-
786 validated sequence datasets) at species level. Violins with different lower-case letters have
787 significantly different means (paired t-test false detection rate-corrected $P < 0.05$).

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

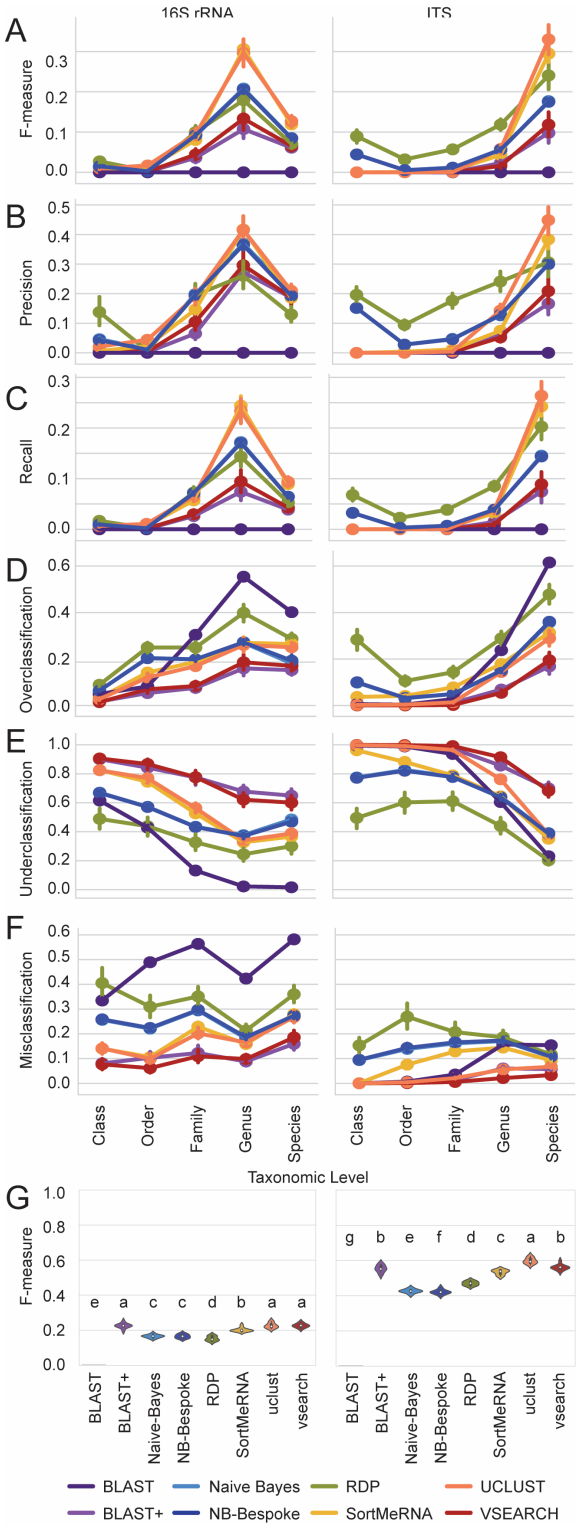

788

789 Figure 3. Classifier performance on novel-taxa simulated sequence datasets for 16S rRNA
790 sequences (left column) and fungal ITS sequences (right column). A-F, Average F-measure
791 (A), precision (B), recall (C), overclassification (D), underclassification (E), and
792 misclassification (F) for each taxonomy classification method (averaged across all
793 configurations and all novel taxa sequence datasets) from phylum to species level. Error
794 bars = 95% confidence intervals. B, Average F-measure for each optimized classifier
795 (averaged across all novel taxa sequence datasets) at species level. Violins with different
796 lower-case letters have significantly different means (paired t-test false detection rate-
797 corrected $P < 0.05$).

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

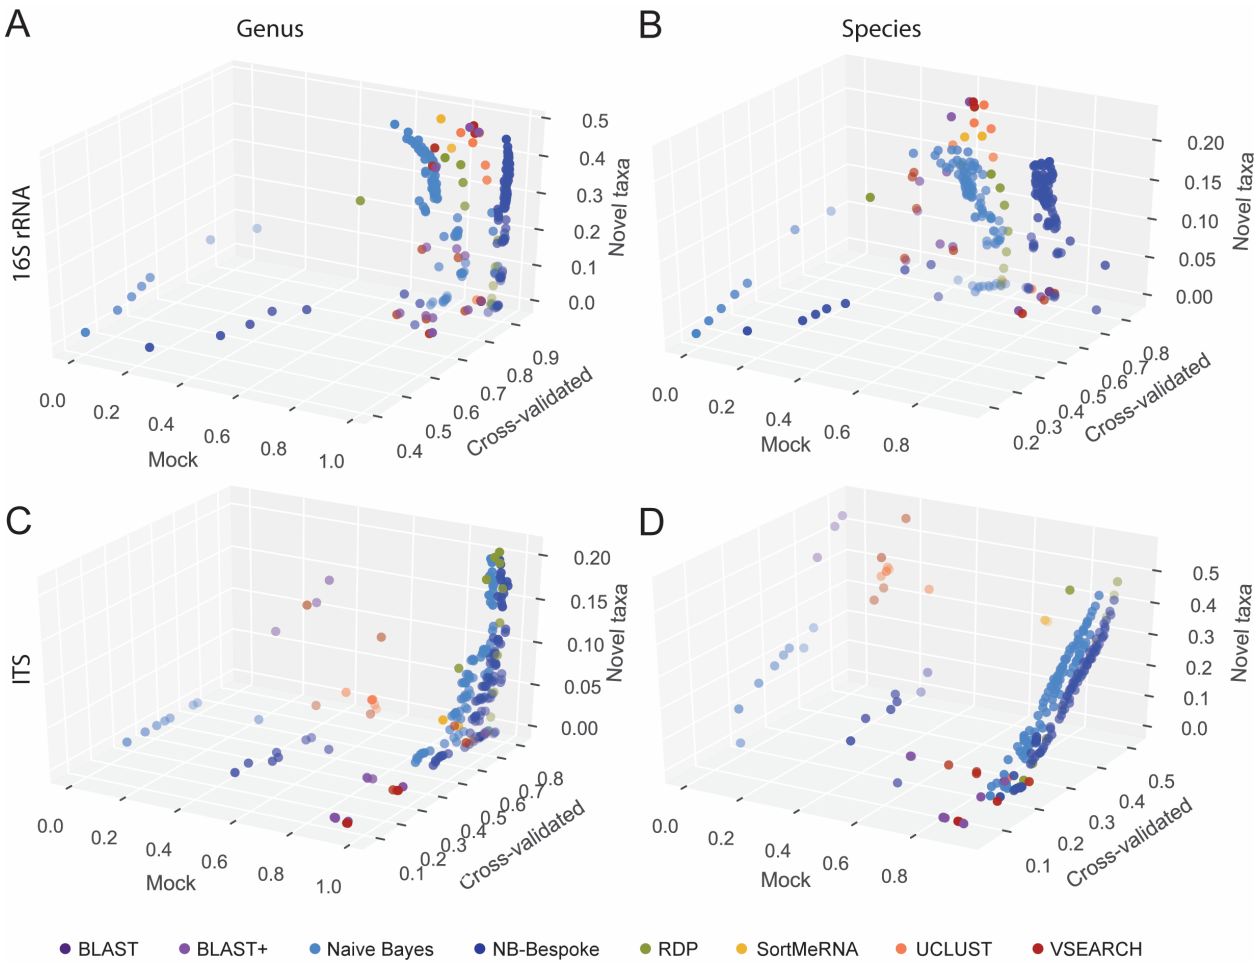

Figure 4. Classification accuracy comparison between mock community, cross-validated, and novel taxa evaluations. Scatterplots show mean F-measure scores for each method configuration, averaged across all samples, for classification of 16S rRNA at genus level (A) and species level (B), and fungal ITS sequences at genus level (C) and species level (D).

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

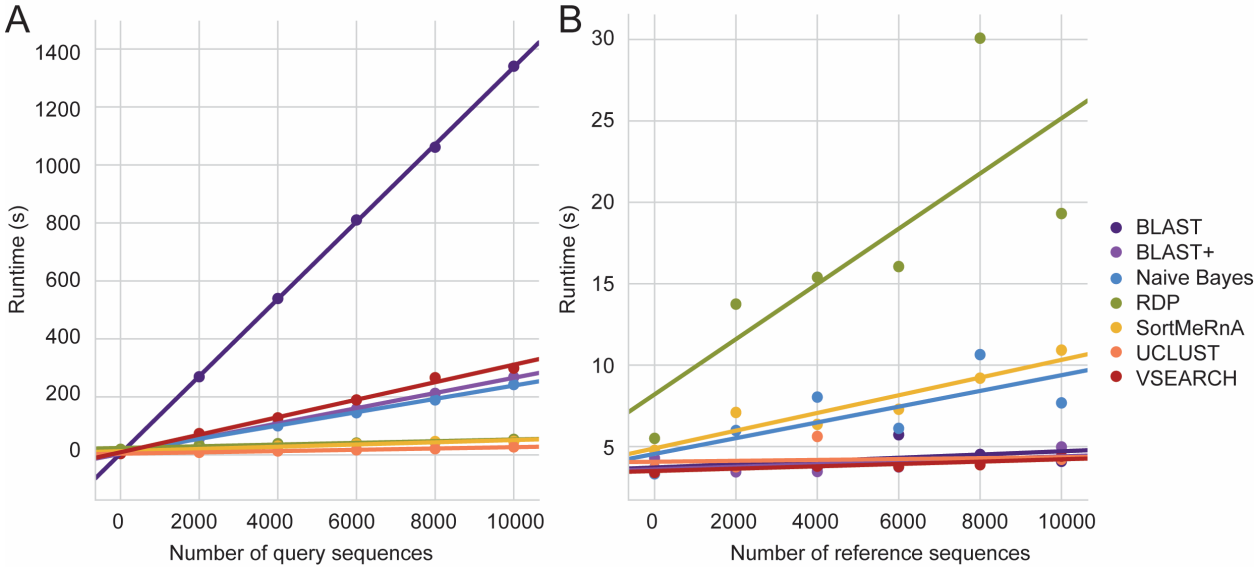

Figure 5. Runtime performance comparison of taxonomy classifiers. Runtime (s) for each taxonomy classifier either varying the number of query sequences and keeping a constant 10000 reference sequences (A) or varying the number of reference sequences and keeping a constant 1 query sequence (B).

References

1. Human Microbiome Project Consortium. A framework for human microbiome research. Nature. 2012 Jun 13;486(7402):215–21.

2. Gilbert JA, Jansson JK, Knight R. The Earth Microbiome project: successes and aspirations. BMC Biol [Internet]. 2014;12(1). Available from: http://dx.doi.org/10.1186/s12915-014-0069-1

3. Wang Q, Quensen JF 3rd, Fish JA, Lee TK, Sun Y, Tiedje JM, et al. Ecological patterns of nifH genes in four terrestrial climatic zones explored with targeted metagenomics using FrameBot,

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

a new informatics tool. MBio. 2013 Sep 17;4(5):e00592–13.

4. Callahan BJ, McMurdie PJ, Rosen MJ, Han AW, Johnson AJA, Holmes SP. DADA2: High-resolution sample inference from Illumina amplicon data. Nat Methods. 2016 Jul;13(7):581–3.

5. McDonald D, Price MN, Goodrich J, Nawrocki EP, DeSantis TZ, Probst A, et al. An improved Greengenes taxonomy with explicit ranks for ecological and evolutionary analyses of bacteria and archaea. ISME J. 2012 Mar;6(3):610–8.

6. Caporaso JG, Kuczynski J, Stombaugh J, Bittinger K, Bushman FD, Costello EK, et al. QIIME allows analysis of high-throughput community sequencing data. Nat Methods. 2010 May;7(5):335–6.

7. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B, Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E. Scikit-learn: Machine Learning in Python. J Mach Learn Res. 2011;12:2825–30.

8. Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F., Mueller, A., Grisel, O., Niculae, V., Prettenhofer, P., Gramfort, A., Grobler, J., Layton, R., VanderPlas, J., Joly, A., Holt, B., Varoquaux G. API design for machine learning software: experiences from the scikit-learn project. In: ECML PKDD Workshop: Languages for Data Mining and Machine Learning. 2013. p. 108–22.

9. Camacho C, Coulouris G, Avagyan V, Ma N, Papadopoulos J, Bealer K, et al. BLAST : architecture and applications. BMC Bioinformatics. 2009;10(1):421.

10. Rognes T, Flouri T, Nichols B, Quince C, Mahé F. VSEARCH: a versatile open source tool for metagenomics. PeerJ. 2016 Oct 18;4:e2584.

11. Bokulich NA, Rideout JR, Mercurio WG, Shiffer A, Wolfe B, Maurice CF, et al. mockrobiota: a Public Resource for Microbiome Bioinformatics Benchmarking. mSystems [Internet]. 2016 Sep;1(5). Available from: http://dx.doi.org/10.1128/mSystems.00062-16

12. Wang Q, Garrity GM, Tiedje JM, Cole JR. Naive Bayesian classifier for rapid assignment of rRNA sequences into the new bacterial taxonomy. Appl Environ Microbiol. 2007 Aug;73(16):5261–7.

13. Altschul SF, Gish W, Miller W, Myers EW, Lipman DJ. Basic local alignment search tool. J Mol Biol. 1990 Oct 5;215(3):403–10.

14. Edgar RC. Search and clustering orders of magnitude faster than BLAST. Bioinformatics. 2010 Oct 1;26(19):2460–1.

15. Kopylova E, Noé L, Touzet H. SortMeRNA: fast and accurate filtering of ribosomal RNAs in metatranscriptomic data. Bioinformatics. 2012 Dec 15;28(24):3211–7.

16. Müller AC, Behnke S. pystruct - Learning Structured Prediction in Python. J Mach Learn Res. 2014;15:2055–60.

17. Caporaso JG, Lauber CL, Walters WA, Berg-Lyons D, Huntley J, Fierer N, et al. Ultra-high-

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017

852 throughput microbial community analysis on the Illumina HiSeq and MiSeq platforms. ISME J.
853 2012 Aug;6(8):1621–4.

854 18. Bokulich NA, Mills DA. Improved Selection of Internal Transcribed Spacer-Specific Primers
855 Enables Quantitative, Ultra-High-Throughput Profiling of Fungal Communities. Appl Environ
856 Microbiol. 2013;79(8):2519–26.

857 19. Kõljalg U, Nilsson RH, Abarenkov K, Tedersoo L, Taylor AFS, Bahram M, et al. Towards a unified
858 paradigm for sequence-based identification of fungi. Mol Ecol. 2013 Nov;22(21):5271–7.

859 20. Bray JR, Curtis JT. An Ordination of the Upland Forest Communities of Southern Wisconsin.
860 Ecol Monogr. 1957;27(4):325–49.

861 21. Bokulich NA, Subramanian S, Faith JJ, Gevers D, Gordon JI, Knight R, et al. Quality-filtering
862 vastly improves diversity estimates from Illumina amplicon sequencing. Nat Methods. 2013
863 Jan;10(1):57–9.

864 22. Maurice CF, Haiser HJ, Turnbaugh PJ. Xenobiotics Shape the Physiology and Gene Expression of
865 the Active Human Gut Microbiome. Cell. 2013;152(1-2):39–50.

866 23. Schirmer M, Ijaz UZ, D'Amore R, Hall N, Sloan WT, Quince C. Insight into biases and sequencing
867 errors for amplicon sequencing with the Illumina MiSeq platform. Nucleic Acids Res. 2015 Mar
868 31;43(6):e37.

869 24. Tourlousse DM, Yoshiike S, Ohashi A, Matsukura S, Noda N, Sekiguchi Y. Synthetic spike-in
870 standards for high-throughput 16S rRNA gene amplicon sequencing. Nucleic Acids Res.
871 2016;gkw984.

872 25. Gohl DM, Vangay P, Garbe J, MacLean A, Hauge A, Becker A, et al. Systematic improvement of
873 amplicon marker gene methods for increased accuracy in microbiome studies. Nat Biotechnol.
874 2016 Sep;34(9):942–9.

875 26. Taylor DL, Walters WA, Lennon NJ, Bochicchio J, Krohn A, Caporaso JG, et al. Accurate
876 Estimation of Fungal Diversity and Abundance through Improved Lineage-Specific Primers
877 Optimized for Illumina Amplicon Sequencing. Appl Environ Microbiol. 2016 Dec
878 15;82(24):7217–26.

879 27. Ihrmark K, Bödeker ITM, Cruz-Martinez K, Friberg H, Kubartova A, Schenck J, et al. New
880 primers to amplify the fungal ITS2 region--evaluation by 454-sequencing of artificial and
881 natural communities. FEMS Microbiol Ecol. 2012 Dec;82(3):666–77.

882

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.3208v1 | CC BY 4.0 Open Access | rec: 30 Aug 2017, publ: 30 Aug 2017