Accounting for 16S rRNA copy number prediction uncertainty and its implications in microbial diversity analyses

Yingnan Gao¹, yg5ap@virginia.edu

Martin Wu^{1*}, mw4yv@virginia.edu

¹Department of Biology, University of Virginia, 485 McCormick Road, Charlottesville, VA,
22904, USA

*Corresponding author

Abstract

Background: 16S rRNA gene has been widely used in microbial diversity studies to determine the community composition and structure. 16S rRNA gene copy number (16S GCN) varies among microbial species and this variation introduces biases to the relative cell abundance estimated using 16S rRNA read counts. To correct the biases, methods (e.g., PICRUST2) have been developed to predict 16S GCN. 16S GCN predictions come with inherent uncertainty, which is often ignored in the downstream analyses. However, a recent study suggests that the uncertainty can be so great that copy number correction is not justified in practice. Despite the significant implications in 16S rRNA based microbial diversity studies, the uncertainty associated with 16S GCN predictions has not been well characterized and its impact on microbial diversity studies needs to be investigated.

Results: Here we develop RasperGade16S, a novel method and software to better model and capture the inherent uncertainty in 16S rRNA GCN prediction. RasperGade16S implements a maximum likelihood framework of pulsed evolution model and explicitly accounts for intraspecific GCN variation and heterogeneous GCN evolution rates among species. Using cross validation, we show that our method provides robust confidence estimates for the GCN predictions and outperforms PICRUST2 in both precision and recall. We have predicted GCN for 592605 OTUs in the SILVA database and tested 113842 bacterial communities that represent an exhaustive and diverse list of engineered and natural environments. We found that the prediction uncertainty is small enough for 99% of the communities that 16S GCN correction should improve their compositional and functional profiles estimated using 16S rRNA reads. On the other hand, we found that GCN variation has limited impacts on beta-diversity analyses such as PCoA, PERMANOVA and random forest test.

Conclusion: We have developed a method to accurately account for uncertainty in 16S rRNA GCN predictions and the downstream analyses. For almost all 16S rRNA surveyed bacterial communities, correction of 16S GCN should improve the results when estimating their compositional and functional profiles. However, such correction is not necessary for beta-diversity analyses.

Keywords: 16S rRNA gene copy number, confidence estimates, *RasperGade16S*, hidden state prediction, microbial composition

Background

The 16S ribosomal RNA (16S rRNA) gene is the gold standard for microbial diversity study and has been commonly used to estimate the composition of bacterial and archaeal communities through amplicon sequencing. Sequence reads are usually matched to reference databases like SILVA [1] and GreenGene [2] to determine the presence of taxa and their relative cell abundances. However, the 16S rRNA gene copy number (GCN) can vary from 1 to more than 15 [3, 4] and this large copy number variation introduces bias in the relative cell abundance estimated using the gene read counts (thereafter referred to as gene abundance) [5], and consequently it can skew the community profiles, diversity measures and lead to qualitatively incorrect interpretations [5–7]. As a result, it has been argued that the 16S rRNA GCN variation should be taken into account in 16S rRNA gene-based analyses [5].

The majority of bacteria species have not been cultured or sequenced and their 16S rRNA GCNs are unknown. Studies have shown that 16S rRNA GCN exhibits a strong phylogenetic signal [5,

7], and therefore 16S rRNA GCN can be inferred from closely related reference bacteria. Based on this principle, software has been developed to predict the 16S rRNA GCN [5, 7–9] in a process often referred to as hidden state prediction [10]. However, a recent study rightly points out that the accuracy of 16S rRNA GCN prediction deteriorates as the minimum phylogenetic distance from the query sequence to the reference sequences increases, and the prediction of 16S rRNA GCN is still an open question [11].

The increasing error of 16S rRNA GCN prediction with increasing phylogenetic distance roots from the stochastic nature of trait evolution, which leads to inherent uncertainty in the predicted trait values. One way of reducing the inherent uncertainty is to improve taxon sampling in the reference phylogeny to reduce the query’s phylogenetic distance to the reference [12]. Another way of addressing the inherent uncertainty is to model the uncertainty directly and have a confidence estimate. This will help us evaluate how confident we should be about a GCN prediction and therefore make meaningful interpretation of predictions. Unfortunately, few 16S rRNA GCN prediction tools provide a confidence estimation for the predicted 16S rRNA GCN, and uncertainty is mostly ignored when interpreting the results of downstream analyses [5, 7, 8]. For example, PICRUST2 predicts functional profiles of microbial communities from 16S rRNA sequence data. It predicts 16S rRNA GCN for each operational taxonomic unit (OTU) in the community and uses the predicted values (point estimates) to estimate “corrected” relative cell abundances and metagenomes, without accounting for the uncertainty of the predictions. As a result, the impact of uncertainty in 16S rRNA prediction on microbial diversity analyses remains unknown and needs to be investigated.

Several points need to be considered to properly model the prediction uncertainty. First, because the uncertainty roots from the stochastic nature of trait evolution, we need to develop a good model for 16S rRNA GCN evolution. Previously the evolution of the 16S rRNA GCN trait has been modeled as gradual evolution using the Brownian motion (BM) model [5, 7, 8]. However, alternative models exist and need to be considered [13–15]. For example, the Ornstein-Uhlenbeck model assumes a centralizing trend towards an optimum [13, 14]. As 16S rRNA GCN has been linked to the ecological strategy of bacterial species [16, 17] and bacteria diversify across all types of environments, a consistent trend in the evolution of 16S rRNA GCN is unlikely. Thus, a model without any trend like the BM model is preferred. Pulsed evolution (PE) is another model that assumes no trend in evolution. Unlike the BM model, where the trait evolves gradually as small trait changes accumulate overtime, the PE model postulates that evolution happens in pulses of “jumps” in the trait values followed by stasis [13, 18]. Previous studies have showed that pulsed evolution is prevalent in the evolution of mammalian body size [13, 19]. To develop a proper model for 16S rRNA GCN evolution, the tempo and mode of evolution need to be examined.

Secondly, 16S rRNA GCN can vary within the same species [20–23], which introduces uncertainty to GCN prediction that needs to be accounted for. It has been shown that modeling the intraspecific variation is essential for the analysis of comparative trait data and failing to account for this variation can result in model misspecification [13]. Because species are usually separated by zero branch length in the phylogeny of the 16S rRNA gene, the intraspecific variation can be modelled as time-independent variation, which can also account for measurement errors [19].

Thirdly, there is notable rate heterogeneity in 16S rRNA GCN evolution. For example, the obligately intracellular bacteria and free-living bacteria with streamlined genomes (e.g., *Rickettsia* and *Pelagibacter*) have elevated molecular evolutionary rates and therefore relatively long branches in the 16S rRNA gene phylogeny. Nevertheless, they have only one copy of 16S rRNA in their genomes and the GCNs rarely change. It is expected that the 16S rRNA GCN prediction for this group of bacteria should be accurate despite their large phylogenetic distances to the reference genomes. Such examples suggest that the rate heterogeneity of 16S rRNA GCN evolution should be systematically evaluated and modelled properly. However, no previous methods have evaluated and modeled such rate heterogeneity, leading to potential model misspecification in 16S rRNA GCN predictions.

Here, we develop a novel tool *RasperGade16S* that employs a heterogeneous pulsed evolution model for 16S rRNA GCN prediction. Using simulation and cross-validation, we show that *RasperGade16S* outperforms current methods such as those implemented in PICRUST2 in providing a much-improved confidence estimate. We show that even if we cannot eliminate inherent uncertainty in 16S rRNA GCN prediction, having an accurate confidence estimation allows us to incorporate the uncertainty in all downstream analyses and therefore to better interpret their results with confidence intervals. We show that correcting 16S rRNA GCN improves the relative cell abundance estimates of the microbial communities and will be beneficial for the vast majority of environmental samples we have analyzed. We also show that GCN correction is unnecessary for beta-diversity analyses because it has limited impact on the results.

Methods

Preparing reference genomes and phylogeny

We downloaded annotated RNA gene sequences from 21245 complete bacterial genomes in the NCBI RefSeq database on April 9, 2021. For each genome, we counted the 16S rRNA gene copy number and picked the 16S rRNA gene sequence that has the highest average similarity to other 16S rRNA gene sequences in the genome as the representative sequence. To remove potential errors introduced by genome mis-assemblies [24], we removed genomes whose 16S rRNA GCN differs from their 5S rRNA GCN by greater than 2 copies, whose 16S rRNA sequence contains ambiguous bases or genomes on the list of withheld genomes in the curated ribosomal RNA operon copy number database rrnDB [3]. We aligned the remaining representative 16S rRNA gene sequences using HMMER version 3.2 [25] with the hidden Markov model built from the GreenGene 13.8 16S rRNA gene alignment, and trimmed the alignment with a mask from the GreenGene database [2]. After collapsing identical 16S rRNA alignments, 6408 representative sequences remained and we built a phylogenetic tree from the trimmed alignment using RAxML version 8.2 [26] with options -f d -m GTRGAMMA. We used the *Deinococcus-Thermus* group to root the reference phylogeny.

Evaluating the rate heterogeneity of 16S rRNA GCN evolution

To estimate the degree of rate heterogeneity in 16S rRNA GCN evolution, we calculated the local average rate of evolution for each genus that contains at least 10 genomes in the reference phylogeny and examined the distribution of the average rates among genera. The average rate of

a genus is calculated as the variance of phylogenetically independent contrasts (PICs) [27] of GCN within the genus.

Modeling 16S rRNA GCN evolution with a heterogeneous pulsed evolution model

Using the R package *RasperGade*, we developed a heterogeneous pulsed evolution model of 16S rRNA GCN evolution with a two-step strategy. First, we fitted a homogeneous pulsed evolution model with the reference phylogeny and used the fitted parameters to guide the first round of binning of nodes into two groups: the slowly- and regularly-evolving groups. Preliminary estimation of the slowly- and regularly-evolving groups using the variance of the PICs shows that their jump frequencies differ by at least 100 folds (assuming the jump sizes are the same in these two groups, Figure S1). Therefore, in our first round of binning, we assumed that the regularly-evolving group evolves under the fitted homogeneous model, and that the slowly-evolving group has a jump frequency 100-fold lower than that of the regularly-evolving group. From the root to the tip, we classified each node into the two groups by the Akaike Information Criterion (AIC) value of the PIC under these two distinct pulsed evolution models. We joined connected nodes of the same group into a neighborhood and then flipped the assignment of a neighborhood if the flip resulted in a better overall AIC value. For the second round of binning, we fitted two pulsed evolution models, one with nodes that were classified as regularly-evolving and the other with slowly-evolving nodes. Using the newly fitted parameters, we repeated the procedure in the first round of binning. After the second round of binning, we calculated r , the rate of evolution in each group, as the process variance per unit branch length defined in a previous study [13]. We then rescaled the reference tree by multiplying the branches in the slowly-evolving group by the ratio $r_{slow}/r_{regular}$. To accommodate time-independent variation in

the tip trait values, we also added a branch length over which the process variance of the fitted pulsed evolution model is equal to the model's time-independent variation to each tip branch.

Predicting 16S rRNA GCN

We used the R package *RasperGade16S* to predict 16S rRNA GCN using the heterogeneous pulsed evolution model. *RasperGade16S* first assigns the query sequence to either the regularly-evolving or the slowly-evolving group based on where it is inserted in the reference phylogeny. For a query sequence inserted into the slowly-evolving group, its insertion branch length is scaled by the ratio $r_{slow}/r_{regular}$ and a small branch length is added to represent the estimated time-independent variation of the specific group the query sequence is inserted into. *RasperGade16S* then predicts the GCN of the query using the rescaled reference phylogeny. Because 16S rRNA GCN is an integer trait, the continuous prediction from hidden state prediction is rounded and a confidence (probability) that the prediction is equal to the truth is estimated by integrating the predicted uncertainty distribution. We marked the 16S rRNA GCN prediction with a confidence smaller than 95% as unreliable, and otherwise as reliable. As a comparison, we also predicted GCN using PICRUST2, which employs multiple hidden state prediction methods in the R package *castor* [28] for 16S rRNA GCN predictions. We selected three methods by which confidence can be estimated: the phylogenetically independent contrast (pic) method, the maximum parsimony (mp) method, and the empirical probability (emp) method. Otherwise, we run PICRUST2 using default options and the unscaled reference phylogeny. For the pic method, the confidence measure is not provided by the hidden state prediction function of the *castor* package, and thus we use a customized R script to reroot the tree at the query sequence and estimate the confidence of the prediction. We did not test the tool CopyRighter [7] in this study

because CopyRighter employs the pic method and we expect its performance to be highly similar to PICRUST2 running the pic method. For the same reason, we did not test the tool PAPRICA [9] because it employs the subtree average method, which is a continuous analogue to the emp method, and we expect its performance to be highly similar to PICRUST2 running the emp method.

Adjust NSTD and NSTI with rate heterogeneity

The adjusted nearest-sequenced-taxon-distances (NSTDs) [11] is calculated using the rescaled reference tree. The adjusted nearest-sequenced-taxon-index (NSTI) [8] is calculated as the weighted average of adjusted NSTDs of the community members.

Validating the quality of predicted 16S rRNA GCN and its confidence estimate

We used cross-validations to evaluate the quality of 16S rRNA GCN prediction and its confidence estimate, and how they vary with NSTD. We randomly selected 2% of the tips in the reference phylogeny as the test set, and filtered the remaining reference set by removing tips with a NSTD to any test sequence smaller than a threshold. We then predicted the 16S rRNA GCN for each tip in the test set using the filtered reference set. We conducted cross-validation within 9 bins delineated by 10 NSTD thresholds: 0, 0.002, 0.005, 0.010, 0.022, 0.046, 0.100, 0.215, 0.464 and 1.000 substitutions/site, and for each bin we repeated the cross-validation 50 times with non-overlapping test sets. We evaluated the quality of the 16S rRNA GCN prediction by the coefficient of determination (R^2), the fraction of variance in the true copy numbers explained by the prediction. We evaluated the quality of confidence estimate by precision and recall. Precision is defined as the proportion of accurately predicted 16S rRNA GCN in predictions considered as

reliable (with $\geq 95\%$ confidence), and recall is defined as the proportion of reliable predictions in the accurately predicted 16S rRNA GCNs. We averaged the R^2 , precision and recall for the 50 cross-validations in each bin.

Simulating bacterial communities with 16S rRNA GCN variation

To evaluate the effect of 16S rRNA GCN correction on microbial diversity analyses, we simulated two sets of bacterial communities using the reference genomes: one set for analyses on relative cell abundances (SC1) and the other set for beta-diversity analyses (SC2). We assumed that each reference genome represents one OTU. For a simulated community in SC1, we randomly selected 100 OTUs from the reference genomes in the test set of cross-validation experiments (around 129 genomes per experiment), and assigned their true cell abundances randomly from a log-series species abundance distribution. We repeated the simulation 20 times for each of the 50 cross-validation experiments at each NSTD threshold, resulting in 9000 simulated communities in SC1. For communities in SC2, simulation was performed similarly as in SC1 except that 1. we set two hypothetical environmental types, 2. randomly assigned 5 unique OTUs to each environmental type, and 3. the unique OTUs were twice more likely to be placed in top ranks of the log-series distribution (i.e., to be more abundant) than the non-unique OTUs in their enriched environment. We repeated this simulation to generate 10 communities for each environment and each of the 50 cross-validation experiments at the smallest NSTD threshold, resulting in 1000 simulated communities in SC2.

Evaluating the effect of 16S rRNA GCN correction on relative cell abundance estimation

We evaluated the effect of 16S rRNA GCN correction on the simulated bacterial communities (SC1). To estimate the confidence interval (CI) of the corrected relative cell abundance of each OTU in a community, we randomly drew 1000 sets of 16S rRNA GCNs from their predicted uncertainty distribution. For each set of 16S rRNA GCNs, we divided the gene read count of OTUs by their corresponding 16S rRNA GCNs to get the corrected cell counts. The median of the corrected cell count for each OTU in the 1000 sets is used as the point estimate of the corrected cell count, and the OTU's relative cell abundance is calculated by normalizing the corrected cell count with the sum of corrected cell counts of all OTUs in the community. The OTU with the highest corrected relative cell abundance is considered the most abundant taxon. We calculated the CIs for the corrected relative cell abundances and the support value for the most abundant OTU. The 95% CI for each OTU's relative cell abundance is determined using the 2.5% and 97.5% quantiles of the 1000 sets of corrected relative cell abundances. The support value for the most abundant OTU is calculated as the empirical probability that the OTU has the highest cell abundance in the 1000 sets of corrected cell abundances. We calculated the coverage probability of the CI as the empirical frequency that the relative gene abundance or true relative cell abundance is covered by the estimated CI.

Evaluating the effect of 16S rRNA GCN correction on beta-diversity analyses

We used the Bray-Curtis dissimilarity, a commonly used metric in beta-diversity analyses, for any beta-diversity analysis that requires a dissimilarity matrix, and evaluated the effect of 16S rRNA GCN correction on the simulated bacterial communities (SC2). To correct for 16S rRNA GCN variation and apply confidence estimation in beta-diversity analyses, for each community

we randomly drew 1000 sets of 16S rRNA GCNs from predicted uncertainty distribution of each OTU in the community. For each set of randomly drawn 16S rRNA GCNs, we calculated the corrected relative cell abundance table and the corresponding Bray-Curtis dissimilarity matrix. We used the corrected cell abundance table to generate the principal coordinates analysis (PCoA) plot and to conduct the permutational multivariate analysis of variance (PERMANOVA) and the random forest test with the R package *vegan* and *randomForest*, respectively.

Predicting 16S rRNA GCN for SILVA OTUs

We downloaded 592605 full length representative bacterial 16S rRNA sequences of non-redundant OTUs at 99% similarity (OTU99) in the SILVA release 132 [1]. We aligned and trimmed the sequences using the method described above. We then inserted the OTUs into the reference phylogeny using the evolutionary placement algorithm (EPA-ng) [29] with the model parameters estimated by RAxML when building the reference phylogeny. We limited the maximum number of placements per SILVA representative sequence to 1. We predicted the 16S rRNA GCN for each SILVA OTU99 as described above using the heterogeneous pulsed evolution model and calculated adjusted NSTDs.

Evaluating the effect of copy number correction in HMP1 dataset

To check the effect of 16S rRNA GCN correction in empirical data, we analyzed the 16S rRNA V1-V3 amplicon sequence data in the first phase of Human Microbiome Project (HMP1) [30]. The 16S rRNA GCN for each OTU in the HMP1 dataset was predicted using *RasperGade16S*. We picked 2560 samples in the HMP1 dataset with complete metadata associated with them and determined the adjusted NSTI and relative cell abundance in each community as described above.

For beta-diversity, we randomly picked 100 samples from each of the 5 body sites and analyzed their beta-diversity as described above.

Examining the adjusted NSTI of empirical bacterial communities

To check the predictability of 16S rRNA GCN in empirical data, we examined bacterial communities surveyed by 16S rRNA amplicon sequencing in the MGnify resource platform [31] that were processed with the latest two pipelines (4.1 and 5.0). The MGnify resource platform uses the SILVA database release 132 [1] for OTU-picking in their latest pipelines, and therefore predicted GCNs for OTUs in the SILVA database can be used directly. We filtered the surveyed communities from the MGnify platform so that only communities with greater than 80% of their gene reads mapped to the SILVA reference at a similarity of 97% or greater were included. This filtering yielded 113842 bacterial communities representing a broad range of environment types. We calculated the adjusted NSTI for each community and examined the adjusted NSTI distribution in various environmental types.

Results

Time-independent variation is present in 16S rRNA GCN evolution

To evaluate the extent of intraspecific variation in 16S rRNA GCN, we examined 5437 pairs of reference genomes with identical 16S rRNA gene alignments. The 16S rRNA GCN differs in 607 (11%) of them, suggesting the presence of significant intraspecific variation or time-independent variation. To formally test its significance, we modeled time-independent variation as a normal white noise, and fitted the Brownian motion (BM) model to the evolution of 16S rRNA GCN in 6408 high-quality reference genomes we compiled, with and without time-

independent variation. Using Akaike Information Criterion (AIC), we found that incorporating time-independent variation with the BM model greatly improves the model fit (Table 1), indicating the necessity to take time-independent variation into account in 16S rRNA GCN prediction. In addition, we observed that the rate of evolution in the fitted BM model is inflated by 1670 folds when time-independent variation is not included in the model. Such inflation in the estimated rate of evolution will lead to overestimation of uncertainty in the 16S rRNA GCN prediction.

Table 1. The AICs of Brownian motion model and pulsed evolution model.

Model	BM	BM (with time-independent variation)	PE (with time-independent variation)
Homogenous model	34338	18028	-7925
Heterogeneous model	29798	3040	-15395

Pulsed evolution model better explains the 16S rRNA GCN evolution than the Brownian motion model

When predicting traits using phylogenetic methods, the BM model is commonly assumed to be the model of evolution. Previous studies have shown that the pulsed evolution (PE) model is a better model for explaining the evolution of animal body size [13, 19], prompting us to test whether pulsed evolution can be applied to explain 16S rRNA GCN evolution as well. Using the R package *RasperGade* which implements the maximum likelihood framework of pulsed evolution described by Landis and Schraiber [13], we fitted the PE model with time-independent variation to the same dataset. Table 1 shows that the PE model provides a significantly better fit

than the BM model, indicating that 16S rRNA GCN prediction should assume the PE model instead of the BM model.

Substantial rate heterogeneity exists in 16S rRNA GCN evolution

To systematically examine the rate heterogeneity of 16S rRNA GCN evolution in the reference genomes, we first used the variance of phylogenetically independent contrasts (PICs) as an approximate estimate of the local evolution rate of 16S rRNA GCN. We found that the rate of evolution varies greatly among genera (Figure S1), but can be roughly divided into two groups with high and low rates of evolution. Therefore, we developed a heterogeneous pulsed evolution model where the frequency of jumps varies between two groups to accommodate the heterogeneity among different bacterial lineages. Using a likelihood framework and AIC, we classified 3049 and 3358 nodes and their descending branches into slowly-evolving and regularly-evolving groups respectively (Figure S2). The frequency of jumps in the regularly-evolving group is 145 folds of the frequency in the slowly-evolving group. All heterogeneous models provide significantly better fits than their homogeneous counterparts (Table 1), indicating that a heterogeneous model should be assumed in predicting 16S rRNA GCN. In addition, we found that PE provides a significantly better model fit in both groups, reinforcing that pulsed evolution should be assumed in the prediction of 16S rRNA GCN.

Apart from the rate of pulsed evolution, we also observed heterogeneity in time-independent variation: for the slowly-evolving group, the fitted model parameters indicate no time-independent variation, while for the regularly-evolving group, the magnitude of time-independent variation is approximately 40% of a jump in the pulsed evolution. The presence of

time-independent variation caps the confidence of prediction in the regularly-evolving group at 85%, which can only be achieved when the query has identical 16S rRNA gene alignment to one of the reference genomes.

The effect of NSTD on accuracy and uncertainty of 16S rRNA GCN predictions

Because of the stochastic nature of evolution, the inherent uncertainty in hidden state prediction accumulates over time, and consequently the accuracy of the prediction decreases as the phylogenetic distance to the reference increases [11]. To get a better understanding of the relationships between the nearest-sequenced-taxon-distances (NSTD) and metrics that measure the accuracy and uncertainty of the prediction, we performed a cross-validation experiment using simulated datasets. We simulated the evolution of 16S rRNA GCN under the fitted heterogeneous pulsed evolution model along the reference phylogeny, and predicted the simulated traits of the genomes using different methods: *RasperGade16S* with the pulsed evolution model and PICRUST2 with the BM model (pic), maximum parsimony (mp) and empirical probability (emp). As expected, we found that the true uncertainty of the prediction, as predicted by the pulsed evolution model under which the simulated GCN evolves, increases with the NSTD. The true confidence of the prediction, calculated as 1 - uncertainty, decreases with the NSTD (Figure 1A, purple line). PICRUST2 with the BM model predicts greatly inflated uncertainty (Figure 1A, red line), while the maximum parsimony method (the default method in PICRUST2) predicts no uncertainty at all (Figure 1A, blue line). The empirical probability method predicts intermediate uncertainty that is greater than the true uncertainty at small NSTDs, but smaller than the truth at large NSTDs (Figure 1A, green line). We used the coefficient of determination (R^2) of the predicted trait values to the true ones to evaluate the accuracy of the

prediction. Figure 1B shows *RasperGade16S* performs the best. It is followed by the maximum parsimony method and the BM model method. The empirical probability method performs the worst. As observed in previous research [11], the accuracy decreases as the NSTD increases (Figure 1B).

Because we will use the uncertainty measure to evaluate the reliability of the prediction, we tested whether the uncertainty predicted by the various methods we compared here captures the true reliability of the prediction. Specifically, we calculated the precision and recall of predictions with a confidence of 95% or greater to examine the recovery of highly reliable predictions. We define precision as the proportion of accurately predicted 16S rRNA GCN in predictions with $\geq 95\%$ confidence, and recall as the proportion of predictions with $\geq 95\%$ confidence in the accurately predicted 16S rRNA GCNs. Ideally, the precision should be greater than 95% throughout the NSTD spectrum, while the recall should gradually drop as the NSTD and the uncertainty of prediction increase. We found that *RasperGade16S* yields high recall at small NSTDs and it decreases as NSTD and uncertainty in the prediction increase (Figure 1C, purple line). In terms of precision, *RasperGade16S* yields high precision throughout the spectrum of NSTD (Figure 1D, purple line). PICRUST2 using the BM model (the pic method) has the lowest recall rate at the smallest NSTD and no recovery beyond as it overestimates the uncertainty (Figure 1C and D, red lines). On the contrary, the maximum parsimony (mp) method yields the highest recall regardless of NSTD as it predicts no uncertainty at all (Figure 1C, blue line), but it suffers from lower precision than *RasperGade16S* and the precision drops quickly as NSTD increases (Figure 1D, blue line). In essence, the pic method thinks the uncertainty is so great that few predictions can be trusted. On the other hand, the mp method thinks there is no

uncertainty and every prediction is reliable. The empirical probability (emp) method shows a similar trend in recall compared to *RasperGade16S*, but suffers from lower precision than *RasperGade16S*.

In summary, the uncertainty of 16S rRNA GCN prediction increases with the increase of NSTD, and as a result, the accuracy of prediction drops as the NSTD increases for all methods. The recall rate of highly reliable predictions also drops with the increasing NSTD and uncertainty, while the precision can remain high throughout the NSTD spectrum.

RasperGade16S improves confidence estimate for 16S rRNA GCN prediction in empirical data

Using empirical data and cross-validation, we compared the performance of various 16S rRNA prediction methods in accuracy and confidence estimates. In general, the trends of uncertainty, accuracy, precision and recall plotted against the NSTD (Figure 2) are very similar to those observed in the simulation study (Figure 1), indicating that *RasperGade16S* models the 16S rRNA GCN evolution reasonably well. As observed in the simulation, the pic and mp methods produce very large and zero uncertainty respectively (Figure 2A), leading to both poor precision and recall rates (Figure 2C and 2D). The emp method performs the worst in terms of accuracy. *RasperGade16S* produces the best overall precision, achieving an average precision rate of 0.96 throughout the NSTD spectrum. In summary, *RasperGade16S*, using the heterogeneous pulsed evolution model, provides one of the best accuracies and the best confidence estimate for 16S rRNA GCN prediction over the full spectrum of NSTD, and should be preferred when predicting 16S rRNA GCN.

As NSTD of the 16S rRNA gene also depends on the sequence alignment and how it is trimmed, it can vary from study to study using the same set of reference sequences. Therefore, to put NSTD values of this study in a taxonomic context, we calculated the NSTDs between taxa at different taxonomical levels. For example, at the species level, we calculated the NSTD of a species to another species within the same genus. We found that the median NSTD between congeneric species is around 0.01 substitutions/site and the maximum NSTD threshold (0.464 substitutions/site) in our cross-validation experiment roughly correspond to the median NSTD between classes (Figure S3).

Copy number correction improves relative cell abundance estimation

Because 16S rRNA GCN variation biases gene abundances disproportionately among the community members, it distorts the relative cell abundance estimated from the gene abundance [5]. In general, community members with lower relative cell abundances suffer from greater impacts by 16S rRNA GCN variation, while those with higher relative cell abundances appear to be less affected by it (Figure 3A). The impact of 16S rRNA GCN variation also depends on the deviation of a member's GCN from the average GCN of the community members, with larger deviations resulting in larger impacts (Figure 3A). When a member's 16S rRNA GCN is greater than the average GCN of the community, its relative abundance will be overestimated. On the other hand, when a member's 16S rRNA GCN is smaller than the average GCN, its relative abundance will be underestimated. In our simulated experiment, we found that 16S rRNA GCN variation has a large detrimental effect on the estimated relative cell abundance (Figure 3B). On average, the relative cell abundance estimated using the gene abundance increased or decreased by 1.8-fold compared to the true relative cell abundance, and the empirical probability of

correctly identifying the most abundant OTU based on the gene abundance is only around 45% (Figure 3C). Correcting for 16S rRNA GCN improves the estimated relative cell abundance (Figure 3B). As expected, the improvement is greatest when the adjusted nearest-sequenced-taxon-index (NSTI) is small (i.e., when there are closely related reference genomes), and it gradually diminishes when the adjusted NSTI increases. At the smallest adjusted NSTI, the average fold change of the estimated relative cell abundance decreases to 1.1-fold after 16S rRNA GCN correction and the empirical probability of correctly identifying the most abundant OTU increases to around 80% (Figure 3C).

Because we predict each OTU's 16S rRNA GCN with a confidence estimate, we can provide 95% confidence intervals (95% CIs) for their relative cell abundance as well. Ideally, 95% of the true relative cell abundances should be covered by the 95% CIs. Figure 3D shows that the average coverage probability of the true relative cell abundance is about 80%, indicating that our 95% CIs perform decently. Similarly, we can also calculate the coverage probability of our 95% CI to the relative gene abundance. As expected, when the coverage probability to the relative gene abundance increases, the improvement by GCN correction decreases (Figure 3E), and that when this coverage probability exceeds 95%, GCN correction starts to show no improvement in relative cell abundance estimates. In empirical studies when the true abundance is unknown, we can use the coverage probability to the relative gene abundance to decide if GCN correction for a community will likely improve the relative abundance estimation or not. For the most abundant OTU in the community, we can calculate its support value from the 16S rRNA GCN's confidence estimates. We found that the calculated support value closely matches the empirical probability that the most abundant OTU is correctly identified (Figure 3C).

To demonstrate the effect of 16S rRNA GCN correction, we analyzed the data from the first phase of the Human Microbiome Project (HMP1). We included 2560 samples with high-quality metadata from 18 body subsites in the analysis and we found that on average the relative cell abundance with and without 16S rRNA GCN correction changes around 1.25-fold among different communities. Using 95% CIs, we found that less than 3% of the HMP1 samples have a coverage probability to the relative gene abundance greater than 95% (Figure 3F), suggesting that 16S rRNA GCN correction will likely improve the estimated relative cell abundances in almost all the HMP1 samples. In terms of the most abundant OTU, we found that the identity of the most abundant OTU changes after copy number correction in around 20% of the communities. The support values for the most abundant OTUs are around 0.85 on average, indicating high confidence in the identification of the most abundant OTUs.

Copy number correction provides limited improvements on beta-diversity analyses

Because 16S rRNA GCN variation affects the estimated relative cell abundances, it may also affect the beta-diversity analyses such as principal coordinates analysis (PCoA), permutational multivariate analysis of variance (PERMANOVA), and the random forest test that use the relative cell abundance information. To examine the effect of 16S rRNA GCN variation on these analyses, we simulated communities in two types of environments where 5% of the OTUs are enriched in one environment compared to the other. We performed beta-diversity analyses on the simulated data and generated the PCoA plot (Figure 4A). We found that when the relative gene abundance is used to calculate the Bray-Curtis dissimilarity, the positions of the samples in the PCoA plot shift from their positions based on the true relative cell abundance (solid lines in

Figure 4A). Correcting for 16S rRNA GCN reduces more than 60% of the shift in the dissimilarity space ($P < 0.001$, Figure 4A, dotted lines). Despite the shift in the PCoA plot, we found that the clustering of communities does not seem to be affected by the 16S rRNA GCN variation.

In addition to the PCoA plot, we observed a limited effect of 16S rRNA GCN variation on other beta-diversity analyses. In PERMANOVA, the proportion of variance explained (PVE) by the environmental type is 10.8% on average using the true cell abundances. Using gene abundance, the average PVE increases slightly to 10.9% and the increase is not statistically significant ($P = 0.688$, Figure 4B), indicating that PERMANOVA is not very sensitive to 16S rRNA GCN variation. Nevertheless, correcting for 16S rRNA GCN reduces the deviation from the PVE estimated using the true cell abundances by 70% on average.

It is a common practice to compare the relative cell abundance of OTUs of interest between environments. We found that such comparison is also not sensitive to 16S rRNA GCN variation (Figure 4C), with the fold-change of relative cell abundance estimated using the gene abundance and the truth highly concordant ($R^2 > 0.98$). For the top 10 OTUs with the highest fold-change in true cell abundance, on average 93% of them are also the top 10 OTUs with highest fold-change in gene abundance. Alternatively, we can use the random forest test to identify OTUs that are differentially abundant between environments by their importance scores (defined as the mean decrease in classification accuracy if removed from the data). We found that for the top 10 OTUs ranked by the importance score using the true cell abundance, on average 80% of them are also recovered by the random forest test using gene abundance, indicating that similar to

PERMANOVA, random forest test is not very sensitive to the 16S rRNA GCN variation either.

Correcting for 16S rRNA GCN slightly increases the recovery rate to 84% on average.

To examine the effect of 16S rRNA GCN variation correction on beta-diversity in empirical data, we analyzed the beta-diversity using 500 samples from the HMP1 dataset. We found that correction of 16S rRNA GCN does not seem to affect the clustering of communities by body sites in the PCoA plot. Pairwise PERMANOVA shows that the mean PVE by the body site is 14.9% before 16S rRNA GCN correction and decreases marginally to 14.6% after correction, and the PVEs using the gene abundance and the corrected cell abundance are also highly concordant ($R^2 > 0.98$). On the other hand, pairwise random forest tests yield similar results before and after 16S rRNA GCN correction, with around 9 out of the top 10 features identified by the random forest test remaining unchanged before and after correction. In terms of the fold-change of relative cell abundances between body sites, we found that copy number correction has little impact as the estimated fold-change before and after correction are highly similar ($R^2 > 0.95$).

Predicting 16S rRNA GCNs for SILVA OTUs

Using *RasperGade16S*, we predicted the 16S rRNA GCN for 592605 bacterial OTUs (99% identity) in the release 132 of the SILVA database. Overall, the median adjusted NSTD for all bacterial OTUs is 0.070 substitutions/site, and 34.7% of the predictions have a high confidence of 95% or greater, and 74.9% of the predictions have a moderate confidence of 50% or greater (Table 2). In comparison, randomly guessing has a null confidence of around 6.7% (1 out of 15 possible GCNs). Among major phyla with more than 10000 OTUs, the proportion of highly

confident predictions varies greatly (Table 2), with Cyanobacteria having the lowest proportion of 19.1% and Acidobacteria having the highest proportion of 50.4%. Similarly, the proportion of moderately confident predictions varies from 58.3% to 89.5% among these phyla. Interestingly, the proportions of highly confident predictions closely match the proportions of slowly-evolving OTUs in each phylum (Table 2), suggesting a causal relationship between them.

Table 2. Summary of SILVA 16S rRNA GCN predictions.

Taxonomic group	Number of OTUs	Median adjusted NSTD (substitutions/site)	Proportion of OTUs in slowly-evolving group	Proportion of highly confident predictions	Proportion of moderately confident predictions
Bacteria	592605	0.070	34.9%	34.7%	74.9%
Proteobacteria	238929	0.062	43.4%	43.1%	85.4%
Firmicutes	149757	0.091	21.6%	21.5%	68.9%
Actinobacteria	60510	0.061	45.2%	45.2%	89.5%
Bacteroidetes	55663	0.117	29.9%	29.8%	58.3%
Acidobacteria	14534	0.006	50.4%	50.4%	82.7%
Cyanobacteria	13970	0.285	19.9%	19.1%	60.5%

Highly confident predictions are defined as predictions with a confidence of 95% or greater. Moderately confident predictions are defined as predictions with a confidence of 50% or greater.

Vast majorities of microbial communities should benefit from copy number correction

To examine if analysis of real communities would benefit from 16S rRNA GCN correction, we calculated the adjusted NSTI for 113842 communities in the microbiome resource platform MGnify (formerly known as EBI Metagenomics) [31] that passed our quality control. The adjusted NSTI varies greatly among samples and the median across all samples is 0.01 substitutions/site. For example, the smallest adjusted NSTI (8.2×10^{-7} substitutions/site) comes

from a human vaginal clinical sample that is dominated by one OTU (relative cell abundance >0.98). This OTU is closely related to *Lactobacillus iners* in the slowly-evolving group, which results in the extremely small adjusted NSTI. On the other hand, the largest adjusted NSTI (0.6 substitutions/site) belongs to a sample from the rumen of dairy cows. The majority of OTUs in this sample has an adjusted NSTD greater than 0.1 substitutions/site and account for more than 90% of the total cell abundance. In the simulated communities, we observed that GCN correction significantly improves the estimated relative cell abundances ($P < 0.001$) even when the adjusted NSTI reaches 0.3 substitutions/site. We found that the vast majority of the communities (>99%) from MGnify have an adjusted NSTI less than 0.3 substitutions/site, suggesting that they should benefit from 16S rRNA GCN correction when estimating the relative cell abundances. The distribution of adjusted NSTI varies among different environmental types (Figure 5), but the proportion of communities that will likely benefit from 16S rRNA GCN correction remains high, ranging from 98% to 100%.

Discussion

16S rRNA GCN variation skews microbial community composition estimated from the 16S rRNA read count. To correct for the bias introduced by the GCN variation, several tools have been developed to predict GCN from reference genomes. A recent study has pointed out that the GCN predictions come with inherent uncertainty, particularly for these taxa without closely related genomes [11]. The concern is that inaccurate predictions can introduce noise to community compositions that can be worse than the original GCN-related biases, thereby raising doubt about the usefulness of the 16S rRNA GCN correction in microbial diversity analyses.

We address the inherent uncertainty problem in 16S rRNA GCN prediction by directly measuring it with confidence estimates. Using simulations and cross-validation, we show that *RasperGade16S* outperforms PICRUST2 in both the precision and recall rates. *RasperGade16S*' strength comes from three features of its modeling of the 16S rRNA GCN evolution: implementation of a pulsed evolution model, accounting for the rate heterogeneity and time-independent trait variation. Pulsed evolution model expects no trait changes to occur over a short branch as jumps are not likely to happen on that branch. This leads to a higher confidence to 16S rRNA GCN prediction with a short NSTD, and thus improves the recall of the accurate predictions. By incorporating rate heterogeneity, we can make predictions in the slowly-evolving groups with high confidence, even when their NSTDs are large, thereby further improving the overall precision and recall rates. In the reference phylogeny, 48% of branches were estimated to fall within this slowly-evolving group, whose evolution rate is 145 times slower compared to that of the regularly-evolving group. The third source of improvement for *RasperGade16S* comes from accounting for time-independent variation, which can result from measurement error and intraspecific variation. We show that failing to account for time-independent variation results in model misspecification (Table 1) and overestimated rate of evolution for the pic method.

Having confidence estimates is critical in the presence of inherent uncertainty because they provide direct evaluation of the uncertainty associated with the predictions. Although such uncertainty is positively correlated with the NSTD (Figure 2), use of NSTD as a measure of uncertainty lacks a clear statistical meaning. Using cross-validation, we show that *RasperGade16S* has high precision (around 0.96), which means for predictions with high confidence ($\geq 95\%$), 96% of the predictions are accurate. Therefore, we can use the confidence

score provided by *RasperGade16S* to select high-quality predictions if necessary, or we can draw firm conclusions from the 16S rRNA data when the confidence is high. For example, 16S rRNA GCN has been linked to the ecological strategy of microbial species, with oligotrophs generally having low GCNs and copiotrophs having higher GCNs. To better understand the overall ecological strategy of a bacterial community, we can predict its members' GCNs and classify the community into either an oligotroph-dominant, copiotroph-dominant or a mixed community, and we can do this with a measure of confidence.

The application of confidence estimation extends beyond the prediction of 16S rRNA GCN. Because the uncertainty in the prediction is inherited by statistics derived from the predicted 16S rRNA GCN, we can estimate the uncertainty and confidence intervals of important parameters in downstream analyses, such as the relative cell abundance. With confidence intervals, we can draw more meaningful and sound conclusions, such as identifying the most abundant OTU in the community with a support value. Getting confidence estimates of the relative cell abundance is also important for predicting the functional profile of a community based on 16S rRNA sequences. Although PICRUST2 uses an extremely lenient NSTD cut-off to eliminate problematic sequences, it does not provide an accurate confidence measurement of its predictions. As shown in this study, the default maximum parsimony method used by PICRUST2 to predict 16S rRNA GCN essentially thinks there is no uncertainty in the predictions, which is unrealistic and leads to poor precision. Incorporation of a more meaningful confidence estimate of 16S rRNA GCN prediction in PICRUST2 should make its functional profile prediction more informative.

We predicted GCN for 592605 bacterial OTU99 in the SILVA database. Not surprisingly, we observed considerable uncertainty in the GCN predictions. This is because only a small fraction of bacterial diversity in the SILVA database has been captured by the fully sequenced genomes. In addition, 65.1% of OTUs in SILVA database belong to the regularly-evolving group and the confidence of predictions for these OTUs is capped at 85% because of the time-independent variation. However, we would like to point out that natural communities are not a random subsampling of the SILVA OTUs and the median NSTI (NSTD weighted by community members' relative abundance) of the 113842 microbial communities we examined is 0.01 substitutions/site, much lower than the median NSTD of SILVA OTUs. Strikingly, 99% of 113842 microbial communities we examined have an adjusted NSTI less than 0.3 substitutions/site, a range where we show that GCN correction improves the accuracy of the relative cell abundance estimation (Figure 3B). Because these communities represent a comprehensive and diverse list of natural and engineered environments, we believe 16S rRNA GCN correction should be applied systematically in practice if getting accurate estimates of the relative cell abundance is important in the study. Our results therefore affirm the conclusion of the previous studies based on analyses of a much smaller number of communities [5, 7].

McLaren et al. have shown that statistics that are functions of individual taxon's relative abundance will be sensitive to the systematic biases introduced in the sequencing and data analysis pipeline (including bias introduced by 16S rRNA GCN variation) [6]. Contrary to some popular belief, these biases will not cancel out when analyzing the differences between samples that have been measured by the same protocol [6]. Nevertheless, few studies have investigated to what extent the bias introduced by 16S rRNA GCN variation will have on the microbiome beta

diversity analyses. We show that the effect sizes of 16S rRNA bias on beta-diversity analyses are small. Correcting 16S rRNA GCN provides limited improvement on the beta-diversity analyses such as random forest analysis and PERMANOVA test. One possible reason is that for an OTU, the fold change in the relative cell abundance between samples remains more or less the same with or without correcting for the copy number (Figure 4C). For example, assuming the estimated relative cell abundances of an OTU in samples A and B are r_a and r_b respectively without copy number correction. When correcting for the copy number, its relative abundance is adjusted with the scaling factor ACN/GCN, where the GCN is the 16S rRNA copy number of the OTU and the ACN is the average copy number of the sample. Assuming the ACN does not vary much between samples, then the scaling factor for the OTU will be roughly the same in samples A and B. So even with copy number correction, the relative abundance change will still be close to r_a/r_b .

It should be noted that having a confidence associated with the 16S rRNA GCN prediction helps to estimate the uncertainty of the prediction, but it does not improve the accuracy of the prediction. Accuracy of the prediction is constrained by the inherent uncertainty, which can only be improved by better sampling the reference genomes. However, as our current sampling is inadequate for accurate 16S rRNA GCN prediction of all environmental bacteria, we believe that incorporating confidence estimates is the best practice to control for the uncertainty in the 16S rRNA based microbial diversity studies, as opposed to not correcting the GCN bias as previously suggested [11].

Conclusion

We have developed a robust model to estimate the confidence of 16S rRNA GCN predictions. 16S rRNA GCN correction should be beneficial for almost all 16S rRNA surveyed bacterial communities when estimating their compositional and functional profiles. However, the GCN correction is not necessary for microbial beta-diversity analyses.

List of abbreviations

16S rRNA: 16S ribosomal RNA. GCN: gene copy number. 16S GCN: 16S rRNA gene copy number. OUT: operational taxonomic unit. BM: Brownian motion. PE: pulsed evolution. PIC: phylogenetically independent contrast. AIC: Akaike information criterion. NSTI: nearest-sequenced-taxon-index. NSTD: nearest-sequenced-taxon-distance. CI: confidence interval. PCoA: principal coordinates analysis. PERMANOVA: permutational multivariate analysis of variance

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and material

The NCBI accession numbers of the reference genomes, the representative 16S rRNA sequences and alignments, the reference phylogeny, the predicted GCN for OTU99 in the SILVA database, and the simulated bacterial community data generated during the current study are available in the Dryad repository (https://datadryad.org/stash/share/OaS9BjM_kIVdJ3WkZRT7KO8fDr8D4k8jy3LsOtIYELM).

The R package *RasperGade16S* can be downloaded from <https://github.com/wu-lab-uva/RasperGade16S>. The original scripts to conduct the analyses in this study are available in the GitHub repository (<https://github.com/wu-lab-uva/16S-rRNA-GCN-Predcition>).

Competing interests

The authors declare that they have no competing interests.

Funding

Not applicable

Authors' contributions

YG developed the R package *RasperGade16S*, conducted statistical analyses in the manuscript and was a major contributor in writing the manuscript. MW conceptualized the rate heterogeneity model and was a major contributor in writing the manuscript. Both authors read and approved the final manuscript.

Acknowledgement

Not applicable.

References

1. Quast C, Pruesse E, Yilmaz P, Gerken J, Schweer T, Yarza P, et al. The SILVA ribosomal RNA gene database project: improved data processing and web-based tools. *Nucleic Acids Res* 2012; **41**: D590–D596.
2. DeSantis TZ, Hugenholtz P, Larsen N, Rojas M, Brodie EL, Keller K, et al. Greengenes, a Chimera-Checked 16S rRNA Gene Database and Workbench Compatible with ARB. *Appl Environ Microbiol* 2006; **72**: 5069–5072.
3. Klappenbach JA. rrndb: the Ribosomal RNA Operon Copy Number Database. *Nucleic Acids Res* 2001; **29**: 181–184.
4. Větrovský T, Baldrian P. The Variability of the 16S rRNA Gene in Bacterial Genomes and Its Consequences for Bacterial Community Analyses. *PLoS One* 2013; **8**: e57923.
5. Kembel SW, Wu M, Eisen JA, Green JL. Incorporating 16S gene copy number information improves estimates of microbial diversity and abundance. *PLoS Comput Biol* 2012; **8**: 16–18.
6. McLaren MR, Willis AD, Callahan BJ. Consistent and correctable bias in metagenomic sequencing experiments. *Elife* 2019; **8**.
7. Angly FE, Dennis PG, Skarshewski A, Vanwonterghem I, Hugenholtz P, Tyson GW. CopyRighter: a rapid tool for improving the accuracy of microbial community profiles through lineage-specific gene copy number correction. *Microbiome* 2014; **2**: 11.
8. Langille MGI, Zaneveld J, Caporaso JG, McDonald D, Knights D, Reyes JA, et al. Predictive functional profiling of microbial communities using 16S rRNA marker gene sequences. *Nat Biotechnol* 2013; **31**: 814–821.

9. Bowman JS, Ducklow HW. Microbial Communities Can Be Described by Metabolic Structure: A General Framework and Application to a Seasonally Variable, Depth-Stratified Microbial Community from the Coastal West Antarctic Peninsula. *PLoS One* 2015; **10**: e0135868.
10. Zaneveld JRR, Thurber RL V. Hidden state prediction: a modification of classic ancestral state reconstruction algorithms helps unravel complex symbioses. *Front Microbiol* 2014; **5**: 431.
11. Louca S, Doebeli M, Parfrey LW. Correcting for 16S rRNA gene copy numbers in microbiome surveys remains an unsolved problem. *Microbiome* 2018; **6**: 41.
12. Ané C. Analysis of comparative data with hierarchical autocorrelation. *Ann Appl Stat* 2008; **2**: 1078–1102.
13. Landis MJ, Schraiber JG. Pulsed evolution shaped modern vertebrate body sizes. *Proc Natl Acad Sci U S A* 2017; **114**: 13224–13229.
14. Joy JB, Liang RH, McCloskey RM, Nguyen T, Poon AFY. Ancestral Reconstruction. *PLOS Comput Biol* 2016; **12**: e1004763.
15. Elliot MG, Mooers AØ. Inferring ancestral states without assuming neutrality or gradualism using a stable model of continuous character evolution. *BMC Evol Biol* 2014; **14**: 226.
16. Roller BRK, Stoddard SF, Schmidt TM. Exploiting rRNA operon copy number to investigate bacterial reproductive strategies. *Nat Microbiol* 2016; **1**: 1–7.
17. Lauro FM, McDougald D, Thomas T, Williams TJ, Egan S, Rice S, et al. The genomic basis of trophic strategy in marine bacteria. *Proc Natl Acad Sci* 2009; **106**: 15527–15533.
18. Eldredge N, Gould SJ. Punctuated equilibria - an alternative to phyletic gradualism.

- 753 *Models in Paleobiology*. 1972. pp 82–115.
- 754 19. Uyeda JC, Hansen TF, Mcpeck A. The million-year wait for macroevolutionary bursts.
755 *Proc Natl Acad Sci* 2011; **108**: 15908–15913.
- 756 20. Sadeghifard N, Guertler V, Beer M, Seviour RJ. The Mosaic Nature of Intergenic 16S-
757 23S rRNA Spacer Regions Suggests rRNA Operon Copy Number Variation in
758 *Clostridium difficile* Strains. *Appl Environ Microbiol* 2006; **72**: 7311–7323.
- 759 21. Lee CM, Sieo CC, Abdullah N, Ho YW. Estimation of 16S rRNA gene copy number in
760 several probiotic *Lactobacillus* strains isolated from the gastrointestinal tract of chicken.
761 *FEMS Microbiol Lett* 2008; **287**: 136–141.
- 762 22. Bodilis J, Nsague-Meilo S, Besaury L, Quillet L. Variable Copy Number, Intra-Genomic
763 Heterogeneities and Lateral Transfers of the 16S rRNA Gene in *Pseudomonas*. *PLoS One*
764 2012; **7**: e35647.
- 765 23. Lavrinienko A, Jernfors T, Koskimäki JJ, Pirttilä AM, Watts PC. Does Intraspecific
766 Variation in rDNA Copy Number Affect Analysis of Microbial Communities? *Trends*
767 *Microbiol* 2021; **29**: 19–27.
- 768 24. Perisin M, Vetter M, Gilbert JA, Bergelson J. 16Stimator: statistical estimation of
769 ribosomal gene copy numbers from draft genome assemblies. *ISME J* 2016; **10**: 1020–
770 1024.
- 771 25. Eddy SR. Accelerated Profile HMM Searches. *PLoS Comput Biol* 2011; **7**: e1002195.
- 772 26. Stamatakis A. RAxML version 8: a tool for phylogenetic analysis and post-analysis of
773 large phylogenies. *Bioinformatics* 2014; **30**: 1312–1313.
- 774 27. Felsenstein J. Phylogenies and the comparative method. *Am Nat* 1985; **125**: 1–15.
- 775 28. Louca S, Doebeli M. Efficient comparative phylogenetics on large trees. *Bioinformatics*

2018; **34**: 1053–1055.

29. Barbera P, Kozlov AM, Czech L, Morel B, Darriba D, Flouri T, et al. EPA-ng: Massively Parallel Evolutionary Placement of Genetic Sequences. *Syst Biol* 2019; **68**: 365–369.
30. The Human Microbiome Project Consortium. A framework for human microbiome research. *Nature* 2012; **486**: 215–221.
31. Mitchell AL, Almeida A, Beracochea M, Boland M, Burgin J, Cochrane G, et al. MGnify: the microbiome analysis resource in 2020. *Nucleic Acids Res* 2019; **48**: D570–D578.

Figure legends

Figure 1. The performance of prediction on simulated 16S rRNA GCN. Using cross-validation of simulated data, the mean estimated uncertainty and confidence of predictions (A), the mean coefficient of determination R^2 of the predictions (B), and the recall (C) and precision (D) of classification of predictions by their associated confidence estimate, plotted against the mean NSTD. The error bars represent the standard error of the mean.

Figure 2. The performance of prediction on empirical 16S rRNA GCN. Using cross-validation of empirical data, the mean estimated uncertainty and confidence of predictions (A), the mean coefficient of determination R^2 of the predictions (B), and the recall (C) and precision (D) of classification of predictions by their associated confidence estimate, plotted against the mean NSTD. The error bars represent the standard error of the mean.

Figure 3. The impact of 16S rRNA GCN variation on estimated relative cell abundances. (A) The theoretical impact of GCN variation on estimated relative cell abundance. The color of the

lines denotes the ratio of an OTU's of GCN to the average GCN of the community. (B) The average fold-change to the true relative cell abundance. (C) The empirical probability of correctly identifying the most abundant OTU in the community and the support value for the most abundant OTU. (D) The coverage probability of relative cell abundances' estimated 95% CIs to the true relative cell abundance. (E) The correlation between the coverage probability to the relative gene abundance and the improvement by GCN correction. The horizontal red dashed line represents no improvement in relative cell abundance estimates; the vertical red dashed line represents 95% coverage probability to the relative gene abundance. (F) The empirical cumulative distribution of the coverage probability to the relative gene abundance in 2560 samples from the HMP1 dataset. All error bars represent standard error of the mean.

Figure 4. The impact of 16S rRNA GCN variation on beta-diversity. Examples of shift in the Bray-Curtis dissimilarity matrix and other beta-diversity measurements due to 16S rRNA GCN variation. (A) The shift in PCoA plot comparing 20 simulated samples from two hypothetical environments. Solid lines represent the shift of a sample from its true location when using the gene abundance. Dotted lines represent the shift of a sample from its true location when using corrected cell abundance. (B) The shift in the proportion of variance explained by the environment. The red dashed line represents identical outcome with and without 16S rRNA GCN correction. (C) The shift in the fold-change of relative abundances between environments. The red dashed line represents identical outcome with and without 16S rRNA GCN correction.

820 **Figure 5. The distribution of adjusted NSTI in empirical data.** The distribution of adjusted
821 NSTI of 113842 communities in the MGnify database representing various environmental types.
822 The red dashed line marks the adjusted NSTI of 0.3 substitutions/site.

A

GCN to average GCN ratio

B

Fold-change with correction
Fold-change without correction

C

Empirical probability without correction
Empirical probability with correction
Support value

D **E** **F**