CoSMIC - A hybrid approach for large-scale, high-resolution microbial profiling of novel niches

Maor Knafo*[1], Shahar Rezenman*[1], Tal Idan[1], Michael Elgart[2], Shlomi Dagan[3], Ziv Reich[1], Ruti Kapon[1], Dagan Sade^[1] & Noam Shental^[4]

[1]Department of Biomolecular Sciences, Weizmann Institute of Science Rehovot, Rehovot, Israel.

[2] Metabolomic Discovery Unit, Metabolic Center, Sheba Medical Center.

[3] Fisher Drug Discovery Resource Center, The Rockefeller University, New York USA.

[4]Department of Mathematics and Computer Science, The Open University of Israel.

^Corresponding authors:

Dagan Sade - dagan.sade@weizmann.ac.il &

Noam Shental - shental@openu.ac.il

*These authors contributed equally to this work

The authors declare they have no competing interests.

CoSMIC - A hybrid approach for large-scale, high-resolution microbial profiling of novel niches

Abstract

Standard microbial profiling based on 16S rRNA (16S) sequencing suffers from a lack of primer universality, primer biases and often yields low resolution. We introduce 'Comprehensive Small Ribosomal Subunit Mapping and Identification of Communities' (CoSMIC), addressing these challenges, especially in unexplored niches. CoSMIC begins with long-read sequencing of the full-length 16S gene, amplified by generic Locked Nucleic Acid primers over pooled samples, thus augmenting reference databases with novel niche-specific gene sequences. Subsequently, CoSMIC amplifies multiple non-consecutive variable regions along the gene, followed by short-read sequencing of each sample. Data from the different regions are integrated using the SMURF framework, alleviating primer biases and providing de-facto full gene resolution. Using a mock community, CoSMIC identified full-length 16S genes with significantly higher specificity and sensitivity while dramatically increasing resolution compared to standard methods. Evaluating CoSMIC across environmental samples provided higher accuracy and unprecedented resolution while detecting thousands of novel full-length 16S sequences.

Introduction

During the last decade, with the advent of massively parallel sequencing, the study and exploration of microbiomes have become an integrative science that spans multiple disciplines and research fields. Medicine, environmental studies, agriculture, and even behavioral studies have associated specific phenotypes with bacteria(1–4). Most of these studies rely on sequencing predefined variable regions of the 16S rRNA gene using short reads. However, it has been widely shown that the results of such studies depend on the amplified region due to biases introduced by the specific

primers applied(5,6), hence limiting the ability to harmonize findings across studies(7).

Moreover, due to short read length, standard methods amplify one or two adjacent

variable regions; therefore, many bacteria may share their sequence over the

amplicon yet be very different along the rest of the 16S rRNA gene. Such ambiguity in

assigning reads to bacteria often results in lower resolution profiling. However, current

bioinformatic pipelines do not adequately address either problem, regardless of

whether they use amplicon sequence variants (ASVs) or operational taxonomic units

(OTUs).

To overcome these problems, we have previously presented the Short MUlti-Region

Framework (SMURF) approach(8). In SMURF, each sample is amplified by any

number of primer pairs, and reads from different regions are computationally combined

to provide coherent, high-resolution microbial profiling. For such integration, SMURF

relies on available databases of full-length 16S rRNA sequences (e.g., SILVA(9) or

Greengenes(10)) containing millions of sequences. These databases are almost

comprehensive with respect to human-associated bacteria, while other habitats may

be severely underrepresented(11,12), rendering the approach suboptimal for less

explored niches.

Here, we present Comprehensive 16S rRNA gene Mapping and Identification of

Communities (CoSMIC), which builds on SMURF to provide a framework for high-

resolution profiling of novel niches. Before applying SMURF in a new niche, we

suggest amplifying a pool of samples using generic locked nucleic acid (LNA) primers,

which provide an almost universal amplification of full-length 16S rRNA sequences.

Sequencing these amplicons by long-read technologies (e.g., PacBio) enriches 16S

rRNA gene databases with relevant sequences, thus allowing SMURF to provide

efficient and high-resolution profiling of these niches. We provide a workflow for

applying CoSMIC detailing the selection of specific subsets of primers to enable cost

savings while not compromising resolution. The advantages of CoSMIC over standard

methods are demonstrated by sequencing a mock mixture and over a variety of

environmental samples.

CoSMIC is best suited for profiling novel ecological niches. It provides de facto full-

length 16S rRNA resolution at about 10% of the cost of per-sample PacBio

sequencing. This cost-effectiveness, combined with CoSMIC's flexibility in primer

selection, allows researchers to optimize their experimental designs for both scope

and budget. In addition, the full-length 16S rRNA sequences detected in CoSMIC's

initial step may be uploaded to public 16S rRNA databases, thus allowing other
researchers to perform even more cost-effective profiling studies over these niches.

Results

The outcome of microbial profiling highly depends on amplified regions

To demonstrate the effect of preselecting a single variable region (formerly shown,
e.g., by (5,6)), we performed 16S rRNA profiling of leaves, roots, and their surrounding
soil from seven different plants (see Methods). Profiling was conducted using the
commercial Qiagen QIAseq 16S/ITS Region Panel(13), which amplifies six regions
along the gene (Fig. 1A), resulting in 200K-600K reads per sample (Supplementary
Figure 1). We used Qiagen's CLC Genomics Microbiome package, which applies a
closed-reference Operational Taxonomic Units (OTU) picking strategy using a
similarity threshold of 97% (see Methods). The package independently analyzes each
region by assigning taxonomy to OTUs detected in a region and outputs the
taxonomies detected across all regions.
Fig. 1B illustrates the number of OTUs detected in leaf, soil, and root samples of *Aster*
poliothamnus using each of the six amplified regions. For example, the same root
sample results in 600-1100 OTUs, depending on the amplified region (Fig. 1B, right).
Moreover, there is a low overlap in assigned taxonomies detected by different regions,
as evident by the beta diversity among all pairs of amplified regions (Fig. 1C). Beta
diversity was calculated by a Jaccard distance and is higher than zero for any pair of
amplified regions (median Jaccard distances of 0.27, 0.30 and 0.28 for leaf, root, and
soil, respectively).
To evaluate the contribution of each amplicon to the analysis of a given sample type,
we identified the amplicon having the largest number of unique OTUs and sequentially
added regions according to their additional unique OTUs with respect to regions
selected thus far (Fig.1D-F). The "top" region varied across samples, e.g., the most
"informative" regions for leaf, soil, and root in the case of *Aster poliothamnus* were
V3V4, V1V2, and V4V5, respectively. In addition, it required four regions to exhaust
the number of unique OTUs. Yet, the identity of these regions varied among leaf, soil,

and root samples, thus highlighting the inherent biases in region selection.
Comparable observations were made when profiling the other plants - *Brandisia*
hancei, *Lepidium apetalum*, *Lathyrus sativus, Angelica hirsutiflora,* and *Tetraena*
mongolica (Supplementary Figure 2-6), where the most "informative" amplicon was
inconsistent both across plants for a given niche (e.g., leaves across all plants), and
over niches for each plant (Supplementary Figure 7-9).

Figure 1. Microbial profiling of *Aster altaicus* samples depends on the applied amplicon. Leaf, soil, and root
samples of *Aster altaicus* were profiled by the Qiagen kit that provides six amplicons along the 16S rRNA gene,
and their reads were subsequently analyzed using the CLC Genomics platform. **a.** The six amplicons of the Qiagen
kit, each amplifying two or three adjacent variable regions. Taken together, these amplicons cover all variable
regions. **b.** The number of OTUs detected for soil, leaf, and root samples of *Aster altaicus* for each of the six
amplicons. **c.** Distribution of Jaccard distances among pairs of amplicons. In each pair, e.g., V1V2 vs. V3V4, a
Jaccard distance is calculated between the lists of taxons detected by CLC Genomics in each region. **d-f.** Decay
graph of the number of unique taxons identified by the amplicons analyzed. Amplicons are ordered by the number
of unique taxons a region contributes to the former set of regions. D-Leaf, E-Soil, and F-Root.

Combining results from several variable regions partially overcomes problems in profiling

We previously suggested employing the Sample MUlti-Region Framework (SMURF) approach(8) to avoid issues related to pre-selecting a single region. In SMURF, any combination of regions is amplified, and their sequencing results are computationally combined, resulting in a single coherent profiling solution. SMURF relies on a database of full-length 16S rRNA sequences to solve an optimization problem, seeking the gene sequences and their relative frequencies that have given rise to the set of reads from the relevant regions. The output of SMURF is a set of "groups" and their relative frequencies, where each group contains full-length 16S rRNA genes having an identical sequence over the amplified regions. For example, suppose V1V2 and V7V9 regions are being amplified. In that case, full-length 16S rRNA genes having the same sequence over these two regions are indistinguishable and, if detected by SMURF, would be part of the same group. Apart from circumventing the issue of selecting a single region with its associated inherent complications, SMURF provides higher profiling resolution. The larger the number of regions, the smaller the size of a characteristic group, typically allowing the identification of a *single* full-length gene sequence that unambiguously appears in the sample.

We applied SMURF to analyze all eight plants across leaf, soil, and root samples. Every sample underwent analysis through all 63 region combinations of the amplified six regions. For example, an analysis was performed using solely V1V2, by V1V2 combined with V3V4, followed by V1V2 combined with V3V4 and V7V9, and so forth. Our assessment of each region's contribution was based on three parameters: *(a)* ambiguity, i.e., the average group size detected by SMURF. The lower the ambiguity, the higher the resolution as could be expected from the added regions; *(b)* relative detection, representing the fraction of identified 16S rRNA sequences out of those identified when all six regions are used (results based on all six regions are the best possible in this case; hence, all subsets are compared to them); and *(c)* Pearson correlation between the frequencies of 16S rRNA sequences detected by a set of regions vs. those detected by the six-region combination.

The average performance across all three parameters improved when more regions were incorporated into the SMURF analysis. The average ambiguity across all 24 samples was 20,000 sequences when using a single region. Yet, by adding only two

more regions to the analysis, ambiguity was reduced to less than 2,000 (Fig. 2A).
Secondly, the capacity to accurately detect a substantial proportion of the mutual
species increased linearly (Fig. 2B). Since sequences belonging to the same group
differ by un-sequenced regions, including each additional region removed sequences
that were inaccurately declared by the smaller number of regions. Thirdly, as more
amplicons are added, the correlation improves significantly, with many samples
achieving a correlation coefficient (r) higher than 0.9 by adding just two or three regions
(Fig. 2C). Figs. 2A-C displays a significant variance in performance for a given number
of regions, depending on their identity. Assuming that costs scale with the number of
amplified regions, the optimal selection of primers may be crucial. Such selection can
be tailored to suit the specific aims and budget. One can identify an optimal selection
by visualizing mutual taxonomies and Pearson correlation along the horizontal and
vertical axes, respectively, and ambiguity designated by the marker size (Fig. 2D-E).
For example, combining three primer pairs, V1V2, V2V3, and V4V5, yields better
results in soil samples than some sets of four primers (Fig. 2D, red cross vs. all other
pink circles). However, the same combination performs poorly for root (Fig. 2E, red
cross) and leaf (Supplementary Figure 10, red cross).
184

Figure 2. Evaluating the performance of sets of primer pairs across niches. Twenty-four samples

corresponding to the leaf, soil, and root of eight plants were profiled by SMURF, using all 63 combinations of the

six amplicons. **a.** Ambiguity, i.e., the number of 16S rRNA gene sequences per group as detected by SMURF, as

a function of the number of amplicons. Each dot corresponds to the average ambiguity in one of the 63 distinct

cases. The blue line and gray background correspond to the average and standard deviation of the ambiguity

across combinations of a specific number of regions. **b.** We consider SMURF's solution based on all six regions

the gold standard and display a fraction of species identified by subsets of the six regions. Shown is the fraction of

species whose taxonomy matches the six-region solution. **c.** Pearson correlation between profiling based on all six

regions and each set of regions. **d.** A comparison of amplicon sets for soil samples. Each dot corresponds to a

specific region combination averaged across 24 samples, where the marker size scales with the ambiguity and

axes correspond to Pearson correlation (horizontal) and mutual species (vertical). The marker's color indicates the

number of amplicons used for SMURF analysis. Hence, a single marker corresponds to six regions (yellow), and

six dark green dots correspond to the different groups of five regions. The marker "x" corresponds to a specific set

of 3 regions comprising V1V2, V2V3, and V4V5. **e.** The same for root samples.

Enriching the database using long-read sequencing

Although SMURF provides improved profiling, its results depend on the

comprehensiveness of the available full-length 16S rRNA sequence databases.

However, these databases are highly enriched for human-associated bacteria, while

other habitats are often poorly covered.

To bridge this gap and harness the advantages of SMURF across habitats, we
propose a simple preceding step in a microbiome sampling project designated to
provide niche-specific full-length 16S rRNA sequences. The procedure entails a PCR
step, utilizing a generic pre-designed multiplexed set of locked nucleic acid (LNA)
primers aimed at effectively amplifying the whole 16S rRNA gene. LNA primers(14,15)
are shorter than regular primers (10-13nt instead of ~20nt), thus allowing an increased
universality compared to regular primers while maintaining primer specificity and high
Tm (see Methods). We, therefore, propose to pool samples from the relevant niche,
apply PCR amplification using the pre-designed LNA primers, and sequence by long-
reads. To account for potential chimeras, the resulting reads undergo several quality
assurance steps and are clustered at 99% identity to remove redundancy. A
representative sequence from each cluster is aligned to the SILVA NR99 database(9),
and sequences whose alignment to SILVA is lower than 99% are considered novel
(for a detailed description of the abovementioned steps, see Methods). We refer to the
database containing SILVA and the novel sequences as "augmented SILVA".

To illustrate this idea and estimate the number of novel sequences that can be
detected, we used this approach over pools of samples, each from a different niche:
(a) sponges of the species *Spongia officinalis* collected from the Mediterranean coast
of France; (b) *Spongia officinalis* collected from the Mediterranean coast of Israel; (c)
plants from the Arava desert in Israel (climate type, BWh); and (d) plants from the
temperate lowland in Israel (climate type, Csa). These pools were sequenced using
PacBio long reads amplicon sequencing(16), resulting in ~200K high-quality, full-
length 16S rRNA sequences (Supplementary Figure 11) of average length 1600 bp
(Supplementary Figure 12).

Figure 3A presents the percentage of novel sequences as a function of the identity
threshold when aligning to SILVA. Sponges seem to have significantly higher
percentages of novel sequences(17), irrespective of identity threshold. Using the
SILVA definition of a novel full-length 16S rRNA sequence, i.e., a threshold of 99%
identity (vertical dashed line in Fig. 3A), more than 95% of detected sequences were
novel for samples of sponges and about 85% for samples taken from plants. These
novel sequences include many whose relative abundance in the pool is high,
especially in sponges. Fig. 3B shows histograms of the number of PacBio reads across
detected 16S rRNA sequences for the pools of plants and sponges (left and right
columns, respectively) when split into known SILVA sequences and novel sequences

(upper and lower rows, respectively). For sponges, 98% of PacBio sequences are

novel, compared to 90% in plants. Moreover, 17 novel sponge sequences appear in

more than 1000 reads, potentially corresponding to abundant sequences, compared

to zero such sequences in plants (analogously, 214 and 25 sequences have more

than 100 reads for sponges and plants, respectively). Hence, LNA amplification and

sequencing detect novel sequences, some of potentially significant frequency,

especially in the less-explored sponge niche. When subsequently performing short-

read sequencing over individual sponge samples and performing SMURF using the

augmented SILVA, more than 93% of the relative frequency in each sample was

attributed to novel sequences (Fig. 3C, two left columns). Therefore, standard SMURF

is inadequate for these sponge samples. In contrast, on average, 90% of our plant

samples were attributed to known SILVA sequences (Fig. 3C, two right columns).

Another way to evaluate the significance of novel sequences is to calculate the

percentage of short reads assigned to each database, thus estimating its "utility" (Fig.

3D). More than 82% of short reads in each sponge sample align to augmented SILVA

compared to 58% that align to SILVA (Fig. 3D upper panels). We note that each short

read may align to SILVA and novel sequences, yet the additional 24% of the reads

aligned to only novel sequences "drive" SMURF's optimization to reject SILVA

sequences and select novel sequences as part of the profiling solution. The case for

plant samples differs as more than 80% of the reads are aligned to SILVA, and novel

sequences contribute only an additional 8% and 4% for BWh and Csa, respectively

(Fig. 3D lower panels). Consequently, SMURF's solution attributes most of the relative

abundance to SILVA sequences.

Another observation is that 91% and 97% of the novel sequences for sponges and

plants, respectively, appear in less than 10 PacBio reads (Fig. 3B). Also, a rarefaction

curve plotted over these pools does not reach a plateau for either sponges or plants

(Supplementary Figure 13). Although part of these sequences may correspond to

chimeric sequences undetected by the applied algorithms, they may also, in part, be

a manifestation of underexplored communities. Hence, additional long-read

sequencing may detect novel sequences that may close the remaining gap in

unaligned reads.

Figure 3. Enriching the database with novel full-length 16S rRNA sequences. We amplified the full-length 16S rRNA gene using pre-designed LNA primers and sequenced the resulting amplicons by PacBio long-read sequencing. Four samples were processed, each containing a pool of samples from a different habitat, i.e., plants from the Arava desert in Israel, characterized by a climate type BWh, according to Koppen climate classification; plants from the lowland region in Israel, characterized by a climate type Csa; and pools of samples from the marine sponge *Spongia officinalis* collected in the Mediterranean Sea along the coast of either France or Israel. **a.** All non-redundant PacBio sequences (Plants BWh - 14852, Plants Csa - 14589, France - 6507, Israel - 5468) were aligned to SILVA using the CD-HIT algorithm(30). The percentage of sequences that did not appear in SILVA is shown as a function of the identity threshold. The vertical dashed line corresponds to SILVA's standard threshold for calling novel full-length 16S rRNA sequences (SILVA NR99). Colors indicate different biomes sampled. **b.** Histograms of read counts per 16S rRNA sequence for the plants' pool (left panels) and the sponges' pool (right panels). The upper and lower panels show histograms of SILVA and novel sequences, respectively. An identity threshold of 99% was used to call novel sequences with respect to SILVA. **c.** Individual samples from the abovementioned pools were sequenced via short reads and analyzed by SMURF using the augmented database to determine their microbial composition (France – 6 samples, Israel – 3 samples, Plants (BWh) – 10 samples, Plants (Csa) – 10 samples). The total relative frequency assigned by SMURF to novel 16S rRNA sequences per sample (dots) and across types (columns) is shown. **d.** The fraction of reads aligned to either SILVA or augmented SILVA for each

291   sample depicted in C. Each panel corresponds to a different sample type. To evaluate the statistical significance,

292   a paired t-test was performed for each sample type; * P<0.05, *** P<0.001, **** P<0.0001.

294   The CoSMIC approach

When considering the problems and solutions presented in the former sections, we

propose guidelines for large-scale microbiome studies that would provide high-

resolution profiling in a new or known niche while accommodating different budget

tiers (Fig. 4). We recommend treating each research project as a new ecological niche

(orange, right-hand side in the flowchart) and enrich current databases with relevant

full-length 16S rRNA gene sequences, i.e., pool multiple samples of different origins

of the ecological niche, apply LNA amplification and perform long reads sequencing.

Next, using the newly enriched database, we propose to sequence a representative

set of samples using the most extensive set of primers available (e.g., all nine Swift

pairs or all six Qiagen pairs), apply SMURF, and select the subset of primer pairs that

provide sufficient resolution per niche (or select all primer pairs in case costs do not

scale considerably with the number of primers). Once these steps are finalized, the

project transitions to a "familiar niche" case (purple, left-hand side), which performs

routine amplification of the optimal primer set followed by short reads sequencing,

subsequently analyzed by augmented-database SMURF.

Figure 4. The suggested CoSMIC flowchart. Gray - Perform standard DNA extraction from environmental
samples. **Orange -** Establish a niche-relevant full-length 16S rRNA gene database by LNA-based amplification and
long-read sequencing. **Orange-purple -** Test multiple primer sets to find an optimal primer pair combination. **Purple**
- Perform high throughput profiling in a large-scale project. Samples are amplified using the optimal primer
combination and analyzed by SMURF based on the augmented database.

Experimental evaluation of CoSMIC

To evaluate the advantages of CoSMIC, we sampled *Hymelaea hirsuta*, *Solanum*
lycopersicum, and *Triticum aestivum* at their leaves, roots, and surrounding soil. In
addition, we created a mock community of 12 known species. To construct our ground
truth, we performed standard shotgun metagenomic sequencing of each of the ten
samples and used phyloFlash(18) to extract 16S rRNA sequences and their relative
frequencies in each sample (see Methods). We refer to these sequences as ground
truth (gt16S), and their relative frequencies are called ground truth frequencies. An
optimal performance corresponds to detecting the exact full-length sequences of
gt16S while adding no other sequence. Note that each species in the community may
contain several different 16S rRNA sequences, thus the task is to detect all these
sequences.

As the first step in CoSMIC, we pooled DNA from all ten samples and used our LNA
primers to amplify the full-length16S rRNA gene. We sequenced the amplicons using
synthetic long reads technology (Loop Genomics) to achieve full-length 16S rRNA
gene sequences. Novel sequences, with respect to the SILVA NR99 standard, were
added to SILVA to create our reference database.

Each sample was then sequenced using Illumina short reads using the Swift pipeline,
which amplifies several regions along the 16S rRNA gene (see Methods section for
detailed information about obtained reads for Loop, Swift, and shotgun
metagenomics). To apply SMURF, we used our reference database while considering
nine amplicons spanning the 16S rRNA gene based on the abovementioned primers.
For comparison, we provide a standard analysis provided by Swift utilizing the
QIIME2(19) package to generate tables of operational taxonomic units (OTUs) or
amplicon sequence variants (ASVs) utilizing DADA2(20), each based on either all
amplified regions or only the widely used V3V4 amplicon.

Comparison to the ground truth of each sample was based on four criteria: (a)
Alignment identity - the average alignment score between detected 16S rRNA
sequences and gt16S; (b) *Explained frequency* - the sum of ground truth relative
frequencies for gt16S that were also detected by each method; (c) *Erroneous 16S*
rRNA genes sequences - the number of detected sequences that were not present in
gt16S; and (d) *Ambiguity* - the number of sequences in the reference database that
match detected sequences by each method (for details regarding a-d see Methods).

CoSMIC outperforms traditional methods on all parameters (Fig 5).

Alignment identity - Mean and overall alignment identities are higher for CoSMIC

across all samples (Fig. 5A). The distribution varies according to the tissue sampled,

with mock, root, and soil consistently producing better scores than leaf tissue. This is

likely related to the low number of bacteria detected due to chloroplasts occupying

most reads. While the difference between analysis methods is evident across all

tissues, the most striking example is the mock community scores. Swift-based

analysis, based on all regions or only on V3V4, had a median alignment identity of

about 90% even for this simple, well-defined mixture, compared to an almost 100%

alignment identity for CoSMIC.

Explained frequency - Significant improvements are also apparent in the accuracy of

community composition. Fig. 5B presents the total explained frequency as a function

of the identity threshold averaged per sample type. Using CoSMIC, more sequences

are accurately reported, accounting for a higher proportion of the ground-truth sample.

The difference between the methods is particularly evident when only high-quality

alignments are considered, resulting in almost zero frequency explained by traditional

methods.

In three of the four cases, such reduction also seems to affect CoSMIC, yet to a lesser

degree. To understand this phenomenon, we considered the mock mixture whose

species composition is known. The significant drop in explained frequency occurs at

an identity threshold of 98.4% and mainly originates from a mismatch between an *E.*

coli sequence in gt16S and the detected CoSMIC sequence. Aligning Illumina short

reads of the Swift pipeline to both sequences showed a perfect match to the detected

CoSMIC sequence, while the gt16S *E. coli* sequence contained 12 SNPs relative to

the reads. Therefore, the actual value of CoSMIC's explained frequency at identity

thresholds higher than 99% is at least 88%. This indicates that phyloFlash-derived

16S rRNA sequences may not be entirely accurate, demonstrating the inherent

problems in achieving highly accurate assembly of 16S rRNA sequences from shotgun

metagenomics. This problem is not restricted to phyloFlash and also characterizes

other state-of-the-art assemblers(21,22). To alleviate this problem yet provide an

estimate solely based on shotgun results, we used MarkerMAG(23) to improve the

gt16S sequences (see Methods), yielding an explained frequency of 89% for identity

thresholds higher than 99% (Fig. 5B left panel, grey dashed line). This also indicates

that most of the improvement in gt16S is attributed to the errors in the abovementioned

E. coli sequence.

Erroneous 16S rRNA sequences - The percentage of erroneous sequences, i.e., those

reported yet not present in the ground truth sample, is significantly lower for CoSMIC

than the other methods (Fig. 5C). Moreover, in three of the four cases, the number of

erroneous sequences using standard methods grows linearly with the identity

threshold, while CoSMIC displays an increase starting from a threshold of 95%. To

understand this phenomenon in the case of the mock mixture, we used the

MarkerMAG-improved sequences as ground truth, yielding a reduction of 9.5% (Fig.

5C left panel, grey dashed line). Although the percentage of erroneous sequences at

a high identity threshold is still significant, we note that since MarkerMAG-based

sequences explain 89% of the mock mixture sequences at an identity threshold of

99%, the erroneous sequences correspond to only ~11% of the population. Moreover,

the absolute number of these erroneous sequences is merely 12 (Supplementary

Figure 14). In addition, these errors may also be another manifestation of the

limitations of shotgun-based ground truth. Although more than 60 million reads were

allocated to shotgun sequencing of the mock mixture (Supplementary Figure 15), only

~88,000 were assigned to 16S rRNA sequences. Hence, the ground truth may be

missing some sequences, and the actual erroneous sequences at 99% identity may

be lower than 12.

Ambiguity - Fig. 5D presents a histogram of the ambiguity across all 16S rRNA

sequences detected in the ten samples. Most OTUs or ASVs detected by Swift

correspond to hundreds or thousands of the database's 16S rRNA sequences, while

CoSMIC detects a *single* sequence, yielding a more precise, actionable result.

Figure 5E combines all criteria using radar plots for each method, clearly showing

CoSMIC's improved performance, especially in addressing ambiguity (Amb) and

explained frequency at high alignment identity scores (Cov-99). CoSMIC outperforms

standard pipelines regardless of methodology (OTU or ASV) and whether multi-region

or single-region analysis is used. While standard methods can identify species (Id) at

low sequence resolution (Cov-80), they often result in many erroneous species (FP)

not present in the sample and have very high ambiguity (Amb). As expected, adding

more regions to the standard analysis does little to improve results since these

methods do not combine inputs from the different regions.

Figure 5. Experimental evaluation of standard methods vs. CoSMIC. Ten samples were sequenced - leaf, soil, and root samples of *Hymelaea hirsuta*, *Solanum lycopersicum,* and *Triticum aestivum*, and a mock mixture of 12 known species (see Methods). Each sample underwent standard shotgun metagenomic sequencing, from which we detected its full-length 16S rRNA sequences to be considered ground truth. In addition, as part of CoSMIC, we pooled extracted DNA from all ten samples, amplified by LNA primers, and sequenced via synthetic long reads (Loop Genomics) to create full-length 16S rRNA sequences that were added to the SILVA database. Standard amplicon sequencing was performed by a commercially available kit by Swift, which uses primers spanning the 16S rRNA gene to amplify several variable regions (see Methods). Analysis was then performed by CoSMIC using

nine primer pairs and by a standard QIIME2 pipeline. The latter used either nine primer pairs (All) or the V3V4
region only (V3V4) and was performed using OTUs or by ASV calling. Hence, four QIIME2-based results are shown
– OTU-All, OTU-V3V4, ASV-All, and ASV-V3V4. **a.** Alignment identity, i.e., the average alignment score between
detected 16S rRNA sequences and ground truth 16S rRNA genes (gt16S) detected by shotgun metagenomics.
Boxplots present the distribution of alignment identities for each tissue and analysis method. Leaf, soil, and root
panels correspond to an average of over three samples (one from each plant). A Mann-Whitney U test was
performed for all comparisons, which resulted in very low p-values, with the highest being 0.0064. **b.** The total
frequency of correctly detected gt16S by each method as a function identity threshold. The dotted horizontal line
in the left panel (mock mixture) marks the results if we improve the ground truth sequences by MarkerMag. Insets
present the total frequency achieved per sample at a cutoff of 99%; the dotted horizontal line is their average. The
y-axis, shown for the leftmost inset, is relevant for all insets. **c.** Percentage of erroneous 16S rRNA sequences, i.e.,
those identified by an analysis method yet not detected by shotgun metagenomics. The percentage is a function
of the identity threshold, where alignments below the threshold are considered erroneous. The dotted line marks
the MarkerMag-based improvement for the mock community using CoSMIC for up to a 99% identity threshold. **d.**
Histogram of the number of taxa matching an identification by each analysis type. **e.** Radar plot combining all
samples over several criteria (normalized on a 0-1 scale, where large values correspond to better performance):
Id, mean alignment identity for each analysis method as described in A; **Cov-80,** mean total ground truth frequency
at 80% identity threshold as described in B; **Cov-99,** the same for 99% identity threshold; **FP,** One minus the mean
false positive rate as described in C; **Amb,** mean ambiguity as described in d.

Discussion

The diverse applications of microbiome studies have underscored the need for robust and standardized methodologies(24). However, primer selection and arbitrary analysis decisions made by standard methods lead to lower profiling resolution (Figs. 5A, D) and fall short due to the omission of species (Fig. 5B) or the inclusion of species that are not present in the sample (Fig. 5C).

To address these limitations, we formerly proposed applying sets of primers from different variable regions and computationally combining their results via the SMURF approach(8). To do so, SMURF requires a comprehensive database of full-length 16S rRNA sequences. Yet, such comprehensiveness holds mainly for human-associated samples, and SMURF's performance is sub-optimal for other niches (e.g., Fig. 3C-D). CoSMIC aims to bridge this gap in a cost-effective way by introducing a preliminary step of proactively enriching the database with niche-relevant full-length 16S rRNA sequences. We pool samples from the novel niche and apply amplification by our generic LNA-based (almost) universal primers followed by long-read sequencing. Subsequently, one can apply short-read sequencing over each sample and apply SMURF using the augmented database to achieve high-resolution, more accurate

profiling results (Fig. 3C-D, Fig. 5). Applying CoSMIC is more important in less
explored environments, e.g., marine sponges, where close to 100% of the reads align
to novel 16S rRNA gene sequences detected by CoSMIC (Fig. 3C). Therefore, we
hope CoSMIC will be adopted by researchers seeking efficient and accurate large-
scale profiling of less explored niches.

To our surprise, CoSMIC also outperformed standard methods when profiling a mock
community of known species. The average alignment score between 16S rRNA
sequences detected by CoSMIC and the ground truth sequences was 97.2% (Fig. 5A
left panel), while the score of standard methods was lower by ~8%. This corresponds
to a ~120bp difference along the gene, thus potentially leading to a lower resolution.
Also observed was an overestimation of the diversity of 16S rRNA sequences over the
mock mixture, evident in both ASV or OTU-based standard analysis (Fig. 5C left
panel). Hence, CoSMIC provides a more robust estimation of this important ecological
property.

Overall costs per sample in CoSMIC (including all expenses - from DNA extraction to
those related to short- and long-read sequencing) currently range between $25 to $30,
depending on the number of primer pairs applied. Hence, overall costs only marginally
depend on the number of primer pairs since the expensive stages are DNA extraction
and other sample preparation steps, which do not depend on the number of applied
primer pairs. Also, amplification is often applied in multiplex; hence, sample
preparation and labor may be comparable to using a single primer pair. Moreover,
large-scale projects may benefit from a 10-15% reduction in costs associated with the
number of applied primer pairs. Therefore, we advise researchers to estimate the
minimal subset of primer pairs that provide adequate results for their samples (e.g.,
Fig. 2).

The straightforward approach of using our generic LNA primers and applying PacBio
long-read sequencing to individual samples rather than pools of samples is possible
and would achieve equivalent results to CoSMIC. However, its costs become
prohibitive for large projects. For example, the estimated cost for PacBio sequencing
is about $300 per sample for tens of samples and even higher costs for larger cohorts
due to the price of additional indices. Hence, CoSMIC provides a means of
circumventing such high costs while providing analogous results.

Researchers may also turn to shotgun metagenomics techniques to retrieve 16S rRNA
genes, yet the aims of shotgun sequencing differ from mere profiling, and their costs
are much higher. Moreover, state-of-the-art assemblers often fail to provide highly
accurate 16S rRNA gene sequences of a sample's organisms (e.g., the *E. coli*
sequence in our mock mixture of Fig. 5). The gene's conserved regions across species
and often its almost identical copies within the same organism hampers accurate
assembly. As a result, relying solely on metagenomics-based approaches for 16S
rRNA profiling may provide less accurate results than CoSMIC.

While the CoSMIC approach represents an advancement in microbiome analysis, it
requires careful quality assurance by researchers. In CoSMIC, a novel niche refers to
cases where a significant fraction of the reads cannot be matched to the available
database of full-length 16S rRNA sequences. Hence, researchers are expected to
monitor this fraction to ensure comprehensive results and re-apply the LNA step upon
encountering new communities.

Adding irrelevant or incorrect full-length 16S rRNA sequences to an ad-hoc CoSMIC
database does not affect its reconstruction results. For example, thousands of 16S
rRNA gene sequences that appeared as single PacBio reads and were part of the
augmented database did not participate in any of CoSMIC's reconstructions. Also,
chimeric reads that end up in an augmented database cannot be part of CoSMIC's
reconstructions since there will not be a set of short reads that supports such an
erroneous sequence. However, as performed here, we advise researchers to apply
chimera-detection (e.g., intrinsic algorithms by PacBio, VSEARCH
uchime_denovo(25), etc.) prior to adding sequences to the ad-hoc database to keep
the database "clear" of irrelevant sequences. Moreover, when compiling a set of novel
sequences to be added to a public general-purpose 16S rRNA database, or a niche-
specific pool of sequences, researchers should only consider sequences detected by
SMURF in at least one sample. By this, databases would only contain "validated"
sequences, i.e., those that passed both chimera-detection algorithms and were
detected by SMURF. This would potentially allow researchers of the same niche to
directly apply the short-read part of CoSMIC and skip long-read sequencing once an
adequate "coverage" of the environment has been achieved.

Comparisons between CoSMIC and other methods were based on detecting the

correct bacteria yet ignoring their relative frequencies. This omission is due to the fact

that frequency estimates are not robust enough in all primer-based methods, including

CoSMIC. These inaccuracies mainly result from noise and biases in PCR

amplification, yet they also stem from different copy numbers of the 16S rRNA

genes or different sequences of the gene in the same cell. To overcome this problem

and allow accurate estimates of relative frequencies, we aim to utilize unique

molecular identifiers (UMIs) as part of the CoSMIC sample preparation, similar to our

former quantification method(26).

In conclusion, employing CoSMIC will circumvent the problems caused by standard

profiling methods, allow harmonizing results across studies from the same niche or

from different niches, and potentially aid in inferring species functions and biological

importance. We hope that this framework, allowing low-cost, high-resolution microbial

profiling of any niche, would provide a modular and optimized solution for the evolving

landscape of microbiome research.

Methods

Sample collection and preservation

Overall, collected samples consisted of 33 individual plants and their surroundings, one mock community sample, two sample pools derived from multiple plants and their surroundings (collected in distinct locations in Israel), two sample pools derived from sponges collected in Israel and France, and another pool composed of all 33 individual samples used in the study (Supplementary Table 1). Leaf, root, and soil samples were collected in the field and kept in a nucleic acid protection buffer(27). Sponge samples were collected and immediately preserved in ethanol absolute. Upon arrival at the lab, samples were kept at -80°C until DNA extraction.

DNA extraction

Genomic DNA was extracted from 100mg of the root and leaf samples, 250mg of soil samples, and an approximately 1cm x 0.5cm piece from the sponges using the ZymoBIOMICS DNA Kit (Cat. D4300), following the manufacturer's protocol.

For the mock community, DNA was extracted from each isolate using GenElute Bacterial Genomic DNA Kits (Cat. NA2110-1KT), and 20ng from each species was used to create a mock community.

Pools were created by adding 50ng from each amplicon post LNA amplification and cleaning up to generate a 500ng pool.

The following subsections describe the specific samples and all steps per processing pipeline (the order of the presentation corresponds to the Results sections).

The Qiagen processing pipeline as used in Figures 1 and 2

Samples

Leaf, soil, and root samples of the following plants were collected: *Aster poliothamnus*, *Brandisia hancei*, *Tetraena mongolica*, *Lathyrus sativus,* and *Angelica hirsutiflora*. In addition, two *Lepidium apetalum* plants were collected. In total, 24 samples were processed.

Library preparation and sequencing

Libraries for deep sequencing were prepared using the QIAseq 16S/ITS Region Panel protocol (Cat. 333842), which amplifies the V1V2, V2V3, V3V4, V4V5, V5V7, and V7V9 regions (primer sequences may be available upon request from Qiagen). Sequencing was performed on an Illumina MiSeq using 250|8|8|250 cycles.

Qiagen-based analysis in Figure 1

The resulting FASTQ files were analyzed using the QIAGEN CLC Genomics Workbench software with the CLC Microbial Genomics module, which performs closed-reference Operational Taxonomic Unit (OTU) analysis. More specifically, the software preprocesses the raw sequencing data, aligns the reads against a reference database (Greengenes(10) clustered at 99%), and clusters the sequences into OTUs based on a similarity threshold of 97%. It then assigns taxonomy to each OTU and generates an OTU table containing the frequency of each OTU in each sample.

Combinatorial estimation of sets of variable regions in Figure 2

The same FASTQ files were provided to SMURF which was independently applied to the 63 combinations of the six amplified regions (Fig. 2). The performance of each combination was evaluated by: (a) ambiguity - the number of sequences in each SMURF-detected group, i.e., in each set of 16S rRNA genes having the same sequence over the amplified regions; (b) mutual species - the proportion of correctly identified species, i.e., the number of headers that appear in both groups (six regions-full, and region combination-partial) this parameter would be 1 if the partial combination has all the headers and 0 if they share none with the full six-region analysis; and (c) Pearson correlation between the correctly identified mutual species. A Jupyter Notebook tutorial is available (GitHub NoamShental/CoSMIC).

Enriching the database with novel 16S rRNA gene sequences

Samples for pooled LNA amplification and PacBio long-read sequencing in Figure 3

Four pools were created, each by combining extracted DNA from relevant samples:

Spongia officinalis France pool: Samples from seven different individuals were collected along the coast of Marseille, France.

Spongia officinalis Israel pool: Samples from three different individuals were collected along the coast of Haifa, Israel.

Arava desert Israel pool: Leaf, root and surrounding soil samples from the following
Arava desert species were pooled: *Moringa oleifera, Pulicaria incisa, Ochradenus*
baccatus, Anabasis articulata, Salsola vermiculata, Commiphora gileadensis,
Ochradenus baccatus, Fagonia mollis, Zilla spinosa, and *Astersicus gravelonis.*

Temperate lowlands Israel pool: Leaf, root, and surrounding soil samples from the
following lowlands region species were collected near Rehovot, Israel: *Retama*
raetam, Ferula communis, Vicia villosa, Heterotheca subaxillaris, Sarcopoterium
spinosum, Prasium majus and *Withania somnifera.*

Samples for pooled LNA amplification and Loop Genomics synthetic long reads
sequencing in Figure 5
A pool was created for LNA amplification and Loop Genomics synthetic long-read
sequencing(28) by combining 11 different libraries: a mixture of the 24 samples
prepared using the Qiagen QIAseq kit; one library per each of the nine samples
prepared using the Swift 16S+ITS PANEL kit; and one library for the mock community.

Designing LNA primers for full-length 16S rRNA sequences
Using a refined version of the Greengenes database containing 1.4M full-length 16S
rRNA gene sequences(10) we sought short primer pairs (of length 10-13nt) that
maximize the number of amplified genes. Considering only perfect matches while
ignoring thermodynamic considerations(14), we identified three candidate pairs that
potentially amplify 1.2M out of the 1.4M sequences (Table 1). Candidate pairs were
provided to Qiagen to replace specific nucleotides with locked nucleic acids, resulting
in our LNA primer set.

A.

primer	sequence	Qiagen catalog number
fwd_pair1 (LNA-F-52)	CCACATGCAAG	YCO0190926 (339406)
rev_pair1 (LNA-R-1382)	GGAACGTATTCAC	YCO0190932 (339406)
fwd_pair2 (LNA-F-32)	AACGCTGGCG	YCO0190928 (339406)
rev_pair2 (LNA-R-1376)	CCTATTCACCG	YCO0190929 (339406)
fwd_pair3 (LNA-F-57)	CCGCAAGTCG	YCO0190930 (339406)
rev_pair3 (LNA-R-1406)	GACGGGCGGT	YCO0190931 (339406)

B.

	pair1	pair2	pair3
pair1	1008809	700830	570143

pair2	700830	751633	379021
pair3	570143	379021	712836

Table 1. LNA primers. A. Three primer pairs (pairs 1-3) and their location along the *E. coli* 16S rRNA gene. Primers
are available via the Qiagen catalog number. **B.** The number of amplified full-length 16S rRNA genes by each pair
and by their intersection. About 360K are amplified by all three pairs, 570K by two pairs, and 250K by just a single
pair.

Sample preparation

The LNA primer set was applied to augment the database with novel full-length 16S

rRNA gene sequences to be sequenced by either PacBio or Loop Genomics pipelines.

The forward and reverse primers were combined into two separate mixes at a stock

concentration of 20µM.

Samples were individually amplified at an input concentration of 10-50ng. Each

sample was amplified using the following volumes: 25µl Kapa HiFi HotStart ReadyMix

(cat. KK2602), 1.25µl forward primer pool, and 1.25µl reverse primer pool (each to a

final concentration of 0.5µM), 50ng template, and DDW up to 50µl. PCR plan used:

95°C - 3 minutes, 2 cycles: 98°C - 20 seconds, 55°C - 15 seconds, 72°C - 1 minute;

24 cycles: 98°C - 20 seconds, 58°C - 15 seconds, 72°C - 1 minute; 72°C - 5 minutes,

hold at 4°C. Following amplification, the products were cleaned using Ampure XP

beads (Cat. A63881) at 0.6X concentration. The relevant samples from each niche

were pooled at 100 ng/µl each (for a full list of samples for each pool, see

Supplementary Table 1).

PacBio sequencing

16S rRNA full-length amplicons, obtained from the LNA amplification of the *Spongia*

officinalis microbial community (France and Israel) and of microbial communities in

plants from Israel's lowland and Arava, were used as a starting template for the

SMRTbell Express template Prep Kit 2.0 kit (Product PN. 100-938-900). Sequencing

was done on the PacBio Sequel sequencer. An initial analysis step of demultiplexing

and filtering was performed using the SMRT(28) analysis software. The software also

detects and discards chimeric reads (reads are eliminated if two SMRTbell adaptors

are sequenced).

Loop Genomics sequencing

Samples used to validate the CoSMIC method (Fig. 5) were sequenced using Loop Genomics(29). 16S rRNA full-length amplicons were used as a template for the LoopSeq PCR Amplicon Kit. Sequencing was done on the Illumina NovaSeq sequencer using the following cycles 151|8|8|151, resulting in 2.5-17 Mbp reads per sample (Supplementary Figure 16). Analysis was performed on the company's remote server, yielding fasta files containing assembled long reads.

Filtering potential chimeras

Although the SMRT analysis software already filters out potential chimeric reads, and to support chimera detection for Loop Genomics reads, the uchime_denovo of VSEARCH(25) was applied to detect and remove chimeric sequences. While no chimeric reads were found among full-length 16S rRNA sequences that were to be "augmented" to our database, we acknowledge the possibility that such sequences may still be present. The formation of chimeric sequences during 16S rRNA amplification, especially for full-length genes, has been well-documented in the literature (29). Combining chimera-detection algorithms and the fact that CoSMIC's profiling is robust to chimeric sequences provides significant assurances regarding the validity of detected sequences.

Calling novel 16S rRNA genes following the sequencing of LNA amplicons

Candidate 16S rRNA gene sequences were clustered within each experiment using CD-HIT(30) at 99% identity to remove redundant sequences. Non-redundant sequences were compared to the SILVA database using the cd-hit-est-2d function to determine their alignment to SILVA at different identity levels (Fig. 3A). Script for this analysis appears in the GitHub repository "NoamShental/CoSMIC". This script was used for sequences of both PacBio and Loop Genomics.

The Swift processing pipeline as used in Figure 5

Samples

Leaf, soil, and root samples of *Hymelaea hirsuta, Solanum lycopersicum,* and *Triticum aestivum* were collected.

A mock community was formed using 12 species obtained from environmental isolates
- *Acidovorax citrulli*, *Bacillus subtilis*, *Paenibacillus dendritiformis*, *Clavibacter*
michiganensis, *Escherichia coli*, *Zymomonas mobilis*, *Proteus mirabilis*,
Pseudomonas syringae, *Xanthomonas campestris*, *Agrobacterium tumefaciens*,
Sphingomonas yanoikuyae, and *Bacillus pumilus*. Species identification was
performed by Sanger sequencing using E8f-939r as primers.

Library preparation and sequencing
Libraries for deep sequencing were prepared with the Swift AMPLICON™ 16S+ITS
PANEL (Cat. AL-51696), using five forward and four reverse primers pooled in a single
reaction (Table 2)(31). Sequencing was performed on the Illumina NovaSeq
sequencer with the following cycles: 151|8|8|151, yielding 2.5-15 Mbp reads per
sample (Fig. S15 lower panel). Base calling and demultiplexing were performed using
Illumina's bcl2fastq command-line interface.

Forward primers	Reverse primers
V1_f: GAGTTTGATCMTGGCTCAG	V4_r: CTACCAGGGTATCTAATCC
V3_f: CCTACGGGAGGCAGCAG	V5_r: CCGTCAATTCMTTTGAGTTT
V4_f: GCCAGCAGCCGCGGTAA	V8_r: GACGGGCGGTGTGTACAA
V6_f: ATGGCTGTCGTCAGCT	V9_r: TACCTTGTTACGACTT
V7_f: GYAACGAGCGCAACCC	

701 **Table 2. Primers used in the Swift AMPLICON™ 16S+ITS PANEL.** Primers were applied to amplify the variable
702 regions along the 16S rRNA gene as described in the kit's protocol.

Data analysis
Standard analysis: Microbial profiling was performed using the standard Swift analysis
scripts to create Amplicon Sequence Variant (ASV) or OTU tables. These tables were
based on the QIIME2(19) package, where paired R1 and R2 reads were filtered and
analyzed using R's DADA2(20) package to account for sequencing errors.
CoSMIC analysis: Each sample was demultiplexed to its variable regions using
Cutadapt(32). In principle, since five forward primers were mixed with four reverse
primers, SMURF can be applied to up to 20 specific regions. We selected nine
potential amplicons (Table 3) while ignoring long amplicons (e.g., the pair V1_f and
V9_r was omitted) since the number of such reads was negligible. Paired R1 and R2
reads were filtered, analyzed using R's DADA2 package to account for sequencing
errors, and submitted to SMURF.

Forward primer	Reverse primer
V1_f	V4_r
V1_f	V5_r
V3_f	V4_r
V3_f	V5_r
V4_f	V4_r
V4_f	V5_r
V6_f	V8_r
V6_f	V9_r
V7_f	V8_r
V7_f	V9_r

715 **Table 3. Swift primer pairs applied in SMURF-based reconstruction.** Multiple amplicon types are generated in
716 a multiplex amplification of the Swift nine primer pairs. SMURF analysis considered nine forward-reverse pairs of
717 adjacent regions. For example, two amplicon types starting with V3_f were considered, i.e., those that end with
718 V4_r or V5_r.

Generating ground truth 16S rRNA sequences in Figure 5

Samples

721 The same samples used for the Swift processing pipeline were used to generate
722 ground truth data.

Library preparation and sequencing

724 Extracted DNA from each sample was used with the Nextera DNA Flex Library Prep
725 (Cat. 20018704) to prepare the libraries for metagenomic sequencing. Libraries were
726 sequenced on the Illumina NovaSeq sequencer using the following cycles
727 151|8|8|151, resulting in 20-80 Mbp reads per sample (Fig. S15 upper panel).

Data analysis – generating ground truth 16S rRNA sequences

Ground truth sequences were extracted using phyloFlash(18), which employs a multi-
step approach to identify full-length 16S rRNA sequences from metagenomic data.
First, phyloFlash assembles the metagenomic reads using metaSPAdes(33) to
generate metagenome-assembled genomes (MAGs). Next, it aligns the metagenomic
reads to the SILVA database using a low identity threshold and retaining reads that
align to 16S rRNA gene sequences. These reads are then aligned to three reference
sets: (a) Trusted contigs, i.e., full-length 16S rRNA gene sequences obtained by either
Loop Genomics sequencing or PacBio; b) 16S rRNA genes extracted from the

metaSPAdes assembly of each sample using Barrnap(34); and (c) The SILVA
database of 16S rRNA genes, for reads that were unassigned to the two previous
sources.
This alignment step allows phyloFlash to determine the composition of the bacterial
community in the sample. The list of full-length 16S rRNA sequences from the trusted
contigs, assembled MAGs, and SILVA database alignments served as the ground
truth 16S rRNA sequences (gt16S).

Data analysis – improving candidate 16S rRNA sequences by MarkerMAG
To improve the ground truth sequences for the mock mixture further, MarkerMAG(23)
was applied, which uses a set of marker genes to identify and validate the presence
of specific taxons in metagenomic data. The *matam_16s* script in MarkerMAG, which
is based on the SILVA database, was utilized to generate a list of 16S rRNA marker
sequences for the taxons present in our sample. Next, the marker sequences were
linked to contigs assembled by metaSPAdes using the *link* script in MarkerMAG. This
step allowed us to validate the marker list by confirming that the marker sequences
were linked to the assembled contigs.
To evaluate alignment identity and the erroneous frequency, the validated marker
sequences of the mock mixture were compared against the output of CoSMIC.

Calculating alignment identity, explained frequency, erroneous
755
16S rRNA sequences, and ambiguity in Figure 5

The output of SMURF is a set of groups, where each group includes 16S rRNA genes
whose sequence is identical over the amplified regions yet varies over the rest of the
gene. An analogous definition of a group for standard analysis was created based on
the SILVA database, in which we extracted all relevant 16S rRNA gene sequences
that share the same taxonomy set by the QIIME2 pipeline. For example, if QIIME2
provides a genus-level taxonomy for a certain ASV or OTU, its "group" comprises all
SILVA sequences of this genus. This definition is equivalent to SMURF's groups as
all 16S rRNA sequences sharing the same taxonomy are indistinguishable based on
the reads provided.
Groups of QIIME2 or CoSMIC were used for the comparisons below:
Ambiguity (Fig. 5D): Ambiguity is the number of full-length sequences in a group. Fig.
5D displays a histogram of group sizes for both CoSMIC and QIIME2-based analyses.

The other measures in Fig. 5 required processing the groups. The sequences in each group were clustered using CD-HIT, and the most frequent sequence was selected as the group's representative.

Alignment identity (Fig 5A): Each representative sequence was aligned to its corresponding gt16S list using Blast, and the gt16S sequence with the highest alignment was selected. The unweighted average of Blast's identity score served as the alignment identity.

Explained frequency (Fig. 5B): Overall relative frequencies of groups whose representative 16S rRNA gene sequences aligned to gt16S for a given alignment score.

Erroneous 16S rRNA gene sequences (Fig. 5C): Percentage of 16S rRNA sequences that did not appear in gt16S out of the total number of sequences detected by each method.

The CoSMIC software package

The software package (REF) provides a comprehensive set of scripts facilitating all steps of CoSMIC. The package includes scripts augmenting an existing database by full-length 16S rRNA sequences obtained from long-read sequencing (e.g., PacBio). It then provides code that formats the augmented database in a SMURF-like format per the specific primers' set to be used in the short-read sequencing experiments. The package includes scripts for denoising and demultiplexing paired-end Illumina short reads using the DADA2 algorithm and running SMURF. The package supports SMURF profiling based on any subset of the original set of primers.

Hardware requirements

The CoSMIC package was optimized for efficiency, requiring a 20-core machine with 16GB of RAM per core, a specification within reach for most academic settings. With adjustments, our processes can run on standard PCs, albeit more slowly, making our pipeline accessible and practical for a wide range of research environments without the need for exceptional computational resources.

Documentation for code and processing requirements is provided at the GitHub repository.

References

1. Hamdi C, Balloi A, Essanaa J, Crotti E, Gonella E, Raddadi N, et al. Gut microbiome dysbiosis and honeybee health. Journal of Applied Entomology. 2011 Aug;135(7):524–33.

2. Johnson KVA, Foster KR. Why does the microbiome affect behaviour? Nat Rev Microbiol. 2018 Oct;16(10):647–55.

3. Nejman D, Livyatan I, Fuks G, Gavert N, Zwang Y, Geller LT, et al. The human tumor microbiome is composed of tumor type-specific intracellular bacteria. Science. 2020 May 29;368(6494):973–80.

4. Livyatan I, Straussman R. A spatial perspective on bacteria in tumours. Nature. 2022 Nov;611(7937):674–5.

5. Parada AE, Needham DM, Fuhrman JA. Every base matters: assessing small subunit rRNA primers for marine microbiomes with mock communities, time series and global field samples. Environ Microbiol. 2016 May;18(5):1403–14.

6. Klindworth A, Pruesse E, Schweer T, Peplies J, Quast C, Horn M, et al. Evaluation of general 16S ribosomal RNA gene PCR primers for classical and next-generation sequencing-based diversity studies. Nucleic Acids Res. 2013 Jan;41(1):e1.

7. Schloss PD. Amplicon Sequence Variants Artificially Split Bacterial Genomes into Separate Clusters. mSphere. 2021 Aug 25;6(4):e0019121.

8. Fuks G, Elgart M, Amir A, Zeisel A, Turnbaugh PJ, Soen Y, et al. Combining 16S rRNA gene variable regions enables high-resolution microbial community profiling. Microbiome. 2018 Jan 26;6(1):17.

9. SILVA ribosomal RNA gene database project: improved data processing and web-based tools | Nucleic Acids Research | Oxford Academic [Internet]. [cited 2023 Jun 27]. Available from: https://academic.oup.com/nar/article/41/D1/D590/1069277

10. Greengenes, a Chimera-Checked 16S rRNA Gene Database and Workbench Compatible with ARB - PMC [Internet]. [cited 2023 Sep 3]. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1489311/

11. Prodan A, Tremaroli V, Brolin H, Zwinderman AH, Nieuwdorp M, Levin E. Comparing bioinformatic pipelines for microbial 16S rRNA amplicon sequencing. PLoS One. 2020;15(1):e0227434.

12. Castelle CJ, Banfield JF. Major New Microbial Groups Expand Diversity and Alter our Understanding of the Tree of Life. Cell. 2018 Mar 8;172(6):1181–97.

13. Bradshaw MJ, Aime MC, Rokas A, Maust A, Moparthi S, Jellings K, et al. Extensive intragenomic variation in the internal transcribed spacer region of fungi. iScience. 2023 Jul 10;26(8):107317.

14. Kurreck J, Wyszko E, Gillen C, Erdmann VA. Design of antisense
oligonucleotides stabilized by locked nucleic acids. Nucleic Acids Res. 2002 May
1;30(9):1911–8.

15. Owczarzy R, You Y, Moreira BG, Manthey JA, Huang L, Behlke MA, et al. Effects
of sodium ions on DNA duplex oligomers: improved predictions of melting
temperatures. Biochemistry. 2004 Mar 30;43(12):3537–54.

16. Wenger AM, Peluso P, Rowell WJ, Chang PC, Hall RJ, Concepcion GT, et al.
Accurate circular consensus long-read sequencing improves variant detection
and assembly of a human genome. Nat Biotechnol. 2019 Oct;37(10):1155–62.

17. Thomas T, Moitinho-Silva L, Lurgi M, Björk JR, Easson C, Astudillo-García C, et
al. Diversity, structure and convergent evolution of the global sponge
microbiome. Nat Commun. 2016 Jun 16;7(1):11870.

18. Gruber-Vodicka HR, Seah BKB, Pruesse E. phyloFlash: Rapid Small-Subunit
rRNA Profiling and Targeted Assembly from Metagenomes. mSystems. 2020 Oct
27;5(5):10.1128/msystems.00920-20.

19. Bolyen E, Rideout JR, Dillon MR, Bokulich NA, Abnet CC, Al-Ghalith GA, et al.
Reproducible, interactive, scalable and extensible microbiome data science using
QIIME 2. Nat Biotechnol. 2019 Aug;37(8):852–7.

20. 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.

21. Pericard P, Dufresne Y, Couderc L, Blanquart S, Touzet H. MATAM:
reconstruction of phylogenetic marker genes from short sequencing reads in
metagenomes. Bioinformatics. 2018 Feb 15;34(4):585–91.

22. Yuan C, Lei J, Cole J, Sun Y. Reconstructing 16S rRNA genes in metagenomic
data. Bioinformatics. 2015 Jun 15;31(12):i35–43.

23. Song W, Zhang S, Thomas T. MarkerMAG: linking metagenome-assembled
genomes (MAGs) with 16S rRNA marker genes using paired-end short reads.
Bioinformatics. 2022 Aug 2;38(15):3684–8.

24. Thomas T, Gilbert J, Meyer F. Metagenomics - a guide from sampling to data
analysis. Microb Informatics Exp. 2012 Feb 9;2(1):3.

25. Edgar RC. UCHIME2: improved chimera prediction for amplicon sequencing
[Internet]. bioRxiv; 2016 [cited 2024 May 21]. p. 074252. Available from:
https://www.biorxiv.org/content/10.1101/074252v1

26. Rezenman S, Knafo M, Tsigalnitski I, Barad S, Jona G, Levi D, et al. gUMI-
BEAR, a modular, unsupervised population barcoding method to track variants
and evolution at high resolution. PLOS ONE. 2023 Jun 7;18(6):e0286696.

27. Camacho-Sanchez M, Burraco P, Gomez-Mestre I, Leonard JA. Preservation of
RNA and DNA from mammal samples under field conditions. Mol Ecol Resour.
2013 Jul;13(4):663–73.

28. Picelli S, Faridani OR, Björklund ÅK, Winberg G, Sagasser S, Sandberg R. Full-
length RNA-seq from single cells using Smart-seq2. Nat Protoc. 2014
Jan;9(1):171–81.

29. Callahan BJ, Grinevich D, Thakur S, Balamotis MA, Yehezkel TB. Ultra-accurate
microbial amplicon sequencing with synthetic long reads. Microbiome. 2021 Jun
5;9(1):130.

30. Fu L, Niu B, Zhu Z, Wu S, Li W. CD-HIT: Accelerated for clustering the next-
generation sequencing data. Bioinformatics. 2012 Dec;28(23):3150–2.

31. GitHub - swiftbiosciences/16S-SNAPP-py3: 16S SNAPP on Python 3 [Internet].
[cited 2024 Feb 14]. Available from: https://github.com/swiftbiosciences/16S-
SNAPP-py3

32. Martin M. Cutadapt removes adapter sequences from high-throughput
sequencing reads. EMBnet.journal. 2011 May 2;17(1):10–2.

33. Nurk S, Meleshko D, Korobeynikov A, Pevzner PA. metaSPAdes: a new versatile
metagenomic assembler. Genome Res. 2017 May;27(5):824–34.

34. Seemann T. tseemann/barrnap [Internet]. 2024 [cited 2024 Jun 6]. Available
from: https://github.com/tseemann/barrnap

895
896
897
898

Declarations

Ethics approval and Consent to participate

901 Not applicable

Consent for publication

904 Not applicable

Availability of data and materials

Acknowledgments

Shira Holand for scientific illustrations and graphic design; Micha Ilan for the collections of *Spongis officinallis* from Israel; Theirry Perez and Marie Grenier for the collections of *Spongis officinallis* from France. David Pilzer for PacBio sequencing support. Jonathan Nutkiewitz and Elad Jacobi-Carmeli for helpful thoughts and discussions. Funding was provided by Sustainability and Energy Research Initiative by TI; The Yotam project and the Weizmann Institute sustainability and energy research initiative to DS; Combined Weizmann-Abroad Postdoctoral Grant for Advancing Women in Science and the Emergency Fellowships for Israeli Postdocs Affected by COVID-19 to TI. The Open University of Israel internal grant 10024 to NS.

Authors' contributions

MK, SR, DS, and NS conceived and initiated the project; TI and DS collected most of the samples; MK, SR, and DS designed and carried out the experiments; MK, SR, and NS wrote the data analysis code and the software package; ME, SD, RK, and ZR were involved in the work design and have substantively revised it. MK, SR, TI, and NS prepared the manuscript.

All authors have approved the submitted version (and any substantially modified version that involves the author's contribution to the study), read and approved the final manuscript, and have agreed both to be personally accountable for the author's contributions and to ensure that questions related to the accuracy or integrity of any part of the work, even ones in which the author was not personally involved, are appropriately investigated, resolved, and the resolution documented in the literature.

Competing interests

The authors declare that they have no competing interests.

944