



HHS Public Access

Author manuscript

Nat Methods. Author manuscript; available in PMC 2022 December 01.

Published in final edited form as:

Nat Methods. 2022 June ; 19(6): 671–674. doi:10.1038/s41592-022-01478-3.

Metagenome assembly of high-fidelity long reads with hifiasm-meta

Xiaowen Feng^{1,2}, Haoyu Cheng^{1,2}, Daniel Portik³, Heng Li^{1,2,*}

¹Department of Data Sciences, Dana-Farber Cancer Institute, Boston, MA, USA

²Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA

³Pacific Biosciences, Menlo Park, CA, USA

Abstract

De novo assembly of metagenome samples is a common approach to the study of microbial communities. Current metagenome assemblers developed for short sequence reads or noisy long reads were not optimized for accurate long reads. We thus developed hifiasm-meta, a new metagenome assembler that exploits the high accuracy of recent data. Evaluated on seven empirical datasets, hifiasm-meta reconstructed tens to hundreds of complete circular bacterial genomes per dataset, consistently outperforming other metagenome assemblers.

A short-read metagenome assembly¹ often results in contigs of tens of kilobases (kb) in length², ~1% of a bacterial genome. After years of metagenome sequencing, there were only 62 complete genomes assembled from metagenome samples as of September 2019³. Although we can cluster short contigs into metagenome-assembled genomes (MAGs) with binning algorithms⁴, binning can be an important source of errors which complicate or mislead downstream analysis³. The limitation of short-read MAGs motivated the development of metaFlye⁵, the only published assembler specialized for long-read metagenome assembly. Initially developed for noisy long reads of error rate ~10%, Flye⁶, which metaFlye is based on, does not take advantage of PacBio's high-fidelity reads (HiFi) and is suboptimal for single-species HiFi assembly⁷. To leverage the full power of long accurate HiFi reads, we developed hifiasm-meta, extending our earlier work⁸ to metagenome samples.

In comparison to the assembly of a single species, metagenome assembly poses several unique challenges^{1,9}, such as a larger variance in read length distribution in PacBio HiFi data, and high ploidy combined with low coverage in certain haplotypes. We made several major changes in hifiasm-meta to address these challenges. First, hifiasm-meta has an optional read selection step that reduces the coverage of highly abundant strains without

*To whom correspondence should be addressed: hli@ds.dfci.harvard.edu.

Author Contributions Statement

X.F. and H.L. conceived the project, designed the algorithm and wrote the manuscript. X.F. implemented the algorithm and evaluated the metagenome assemblies. All co-authors helped the data analysis and revised the manuscript.

Competing Interests Statement

D.P. is an employee of Pacific Biosciences. H.L. is a consultant of Integrated DNA Technologies and on the Scientific Advisory Boards of Sentieon and Innozeon. The remaining authors declare no competing interests.

losing reads on low abundant strains. Second, during the construction of the assembly graph, hifiasm-meta tries to protect reads in genomes of low coverage, which may be treated as chimeric reads and dropped by the original hifiasm. Third, hifiasm-meta only drops a contained read if other reads exactly overlapping with the read are inferred to come from the same haplotype. This reduces contig breakpoints caused by contained reads¹⁰. Fourth, after the initial graph construction, hifiasm-meta uses the coverage information to prune unitig overlaps, assuming unitigs from the same strain tend to have similar coverage. It also tries to join unitigs from different haplotypes to patch the remaining assembly gaps. These strategies make hifiasm-meta more robust to features in metagenome datasets.

We first evaluated hifiasm-meta r58-31876a0, metaFlye⁵ v2.9 and HiCanu⁷ v2.2 on two mock communities ATCC and zymo (Table 1). ATCC consists of 20 distinct species, 15 of which are abundant at 0.18–18% and 5 are rare at 0.02% abundance. We were able to reconstruct 13 of the abundant species each as a complete circular contig, comparable to metaFlye and Hicanu (Table S1). All tools assembled *P. gingivalis*, at 18% abundance, into two contigs. No assemblers could fully reconstruct the five species of low abundance. We manually checked the read alignment of these species and found their assembly gaps are all caused by insufficient coverage. We would not be able to assemble these species in full with the current data. The zymo dataset features 21 strains of 17 species, including five strains of *E. coli* at 8% abundance each. A challenge of this dataset lies in the phasing of the *E. coli* strains. Hifiasm-meta assembled strain B766 into a complete circular contig, strain B3008 into 2 contigs and the rest as fragmented contigs. HiCanu assembled both B766 and B3008 into complete circular contigs; metaFlye failed to assemble all 5 strains as circular contigs. Hifiasm-meta produces a more contiguous assembly for *M. smithii* at 0.04% abundance (Table S1). Generally, all three assemblers have comparable accuracy on the two mock community datasets.

We then evaluated the three HiFi metagenome assemblers on real datasets (Table 1). Due to the lack of their true compositions, we used CheckM¹¹ to measure the completeness and the contamination level of each assembly. We define quality brackets in line with the minimum information requirement¹². From the sheepA gut sample, hifiasm-meta reconstructed 323 contigs longer than 1Mb (Fig. 1a; Extended Data Fig. 1) totaling 651Mb in length. 176 of them were near-complete according to CheckM (Fig. 1b). Most long contigs that failed to reach this category are due to incompleteness, not due to excessive contamination. Among the 176 near-complete hifiasm-meta contigs, 134 are circular (Fig. 1b), representing a significant improvement over HiCanu (71 circular near-complete contigs) and metaFlye (47). We aligned hifiasm-meta, HiCanu and metaFlye assemblies to each other and investigated the similarity between them. We found 86% and 94% of circular near-complete HiCanu and metaFlye contigs, respectively, are also circular in the hifiasm-meta assembly and are of similar lengths (Table S3). The remaining near-complete circular HiCanu and metaFlye contigs are assembled into either one linear contig or two linear contigs by hifiasm-meta. Hifiasm-meta can reconstruct most high-quality contigs found by other assemblers. Furthermore, the mash¹³ sequence divergence between hifiasm circular contigs is mostly above 1%, except for four pairs of contigs at 0.62–0.92% divergence. In general, strains of high divergence (more than a couple of percent) can be separated into disconnected contigs; a few strains of low divergence tend to be collapsed and represented

by mosaic contigs; many strains of mixed divergence may lead to complex assembly subgraphs and are the most challenging to assemble.

To reconstruct MAGs from non-circular contigs, we applied the MetaBAT2 binning algorithm⁴ to each assembly. Not optimized for long-read assemblies, MetaBAT2 may mistakenly group different strains of the same species into one MAG and even group two circular contigs occasionally. Such MAGs would be considered to be contaminated by CheckM. To improve binning, we separated circular contigs into individual bins. In the end, we identified more than 110 non-circular MAGs of medium or higher quality from each sheepA assembly (Fig. 1b). Hifiasm-meta still finds more quality MAGs in total.

We applied hifiasm-meta to the larger sheepB dataset¹⁴ (Table 1) and obtained 379 near complete MAGs and 279 circular contigs. Bickhart et al¹⁴ assembled the combined sheepA and sheepB datasets with metaFlye and clustered contigs into MAGs using additional Hi-C data. They reported 44 circular contigs and 428 near complete MAGs evaluated by DAS Tool. For a direct comparison, we ran CheckM on their assembly and identified 241 near complete MAGs instead. Hifiasm-meta produced a more contiguous assembly with HiFi data alone.

For the chicken and the four human gut metagenomes (Table 1), hifiasm-meta consistently produced more circular contigs and more total MAGs than HiCanu and metaFlye as well (Fig. 1b). Hifiasm-meta and metaFlye have comparable performance on the sludge dataset, both better than HiCanu. All assemblers produced fewer MAGs in comparison to the sheepA gut sample. To see how much this is caused by the higher data volume of sheepA, we randomly sampled sheepB, which represents the same specimen, but was sequenced in SequelII and had similar read length distribution to that of humanO1, to ~18Gb of sequences, comparable to the size of humanO1 and sludge. On the downsampled dataset, we could assemble 70 circular contigs, much more than the number of circular contigs in humanO1 and sludge. This suggests that data volume does affect the assembly quality but the more contiguous sheepA assembly is probably more related to the composition of the sample.

Among the four human gut datasets, two were collected from omnivore donors and the other two from vegan donors. Each dataset represents a pool of four individuals (Table 1). We further pooled the four datasets together and co-assembled them. With read names reported in the final hifiasm-meta assembly, we can identify the composition of each contig based on the sources of reads. We found the great majority of contigs of ~1Mb in size and almost all ~1Mb circular contigs are either omnivore-specific or vegan-specific (Fig. 1c), whereas the two omnivore samples are well mixed in long omnivore-specific contigs, so it is with the two vegan samples. We see more reads coming from the humanO1 and humanV1 datasets probably because humanO1 and humanV1 have more and longer reads than the other two human gut samples.

Omnivore and vegan samples are also well separated among co-assembled MAGs, though omnivore- and vegan-specific MAGs are mixed in the phylogenetic tree (Fig. 1d): in this tree, 20 genera consist of three or more MAGs, 17 of which contain both omnivore- and

vegan-specific MAGs. This suggests hifiasm-meta assembly is better at untangling subtle composition differences. Also notably, a clade of seven circular contigs (in the northeastern direction in Fig. 1d) only has 75–79% CheckM completeness but they all have 5S, 16S and 23S rRNA genes and 18 tRNA genes (Table S4). Two of them were assembled by HiCanu into circular contigs of near identical lengths, so they are less likely to be truncated misassembly. We speculate this clade may be underrepresented in CheckM.

On performance, hifiasm-meta took ~18 hours over 48 CPU threads to assemble the sheepA and the chicken datasets and took ~3 hours for the human gut samples (Table S3). On these datasets, it is as fast as metaFlye and is consistently faster than HiCanu by a few folds. Hifiasm-meta tends to use more memory than metaFlye and HiCanu, consuming ~200Gb RAM for the sheepA and chicken gut samples. Hifiasm-meta assembled the largest sheepB dataset in 8.9 days and used 724Gb RAM at the peak.

In the era of short-read sequencing, metagenome assembly was rarely considered as a method to reconstruct full genomes³. This view has been changed by recent progress in long-read assembly^{5,14–16}. Optimized for long accurate HiFi reads, hifiasm-meta moves metagenome assembly even further. It possibly assembles more circular MAGs from one deeply sequenced sample, without manual intervention, than all circular MAGs published in the past. Such high-quality metagenome assemblies may fundamentally change the practice in metagenome analysis and shed light on the biological and biomedical implications of microbial communities.

Methods

Overview of the hifiasm-meta algorithm.

The hifiasm-meta workflow consists of optional read selection, sequencing error correction, read overlapping, string graph construction and graph cleaning. The error correction and read overlapping steps are largely identical to the original hifiasm. We added optional read selection and revamped the rest of steps.

Optional downsampling of input reads.

If read selection is enabled, hifiasm-meta will first make a crude guess of whether there are too many alignments to be performed for the whole read set. This is done by examining anchors and is alignment free. We proceed to do the selection if 2/3 reads have more than 300 target reads. We start with an empty hash table which will record kmer counts, and go through reads in batches of 2000. In a batch, for each read encountered, we collect its canonical kmers and query the hash table for their occurrences. Three percentiles 3%, 5% and 10% are checked against the corresponding thresholds 10, 50 and 50 respectively. If any percentile is lower than the given threshold, the read is kept. The rationale is that we would like to keep a read when it has some rare kmers, i.e. when discarding it will lead to loss of information. Note that the “rare kmers” here are not necessarily rare globally, and the read selection result might change if the inputs are shuffled. We assume that the input is not particularly sorted. After all reads in the batch have been processed, we update the kmer counting hash table with them (kmers of discarded reads are also counted). The termination

criterion of the read selection is either the total number of reads being kept has exceeded the desired count, or all reads have been processed.

Modified chimera detection.

Before graph construction, the original hifiasm regards a read to be chimeric and discards it if a middle part of the read is not covered by other reads. A read from a genome of low abundance may have such an uncovered region due to statistical fluctuation. Hifiasm-meta disables the heuristic if both ends of the read overlap with five or fewer other reads. This extra threshold improves the contiguity of genomes of low abundance.

Treatment of contained reads.

The standard procedure to construct a string graph discards a read contained in a longer read. This may lead to an assembly gap if the contained read and the longer read actually reside on different haplotypes¹⁰. The original hifiasm patches such gaps by rescuing contained reads after graph construction. Hifiasm-meta tries to resolve the issue before graph construction instead. It retains a contained read if other reads exactly overlapping with the read are inferred to come from different haplotypes. In other words, hifiasm-meta only drops a contained read if there are no other similar haplotypes around it. This strategy often retains extra contained reads that are actually redundant. These extra reads usually lead to bubble-like subgraphs and are later removed by the bubble popping algorithm in the original hifiasm.

Changes to graph cleaning.

At the graph construction stage, the original hifiasm-meta rejects overlaps between unitigs inferred to come from different haplotypes. Hifiasm-meta may do this to patch remaining assembly gaps. Hifiasm-meta also uses the unitig coverage to prune overlaps. Suppose unitig A overlaps unitig B and C in the same orientation. Such a bifurcation is an ambiguity in the assembly graph. Let $r_{AB} = \min\{\text{cov}(A), \text{cov}(B)\}$, where $\text{cov}(A)$ is the coverage of A . Hifiasm-meta drops the overlap between A and C if $r_{AB} > 0.7$ and $r_{AC} < 0.7$. This strategy is only applied to unitigs longer than 100kb as it is difficult to accurately estimate coverage for short unitigs. In addition, attempt to resolve short unitigs would not greatly improve the assembly quality anyway in our testing.

Assembly of metagenome datasets.

We evaluated hifiasm-meta r58, HiCanu v2.2 and metaFlye v2.9 all with 48 CPU threads. We used “hifiasm-meta reads.fa” for the assembly of empirical gut samples and used “hifiasm-meta --force-rs reads.fa” to enable read selection for the two mock community datasets. We ran HiCanu with “canu maxInputCoverage=1000 genomeSize=100m batMemory=200 -pacbio-hifi reads.fa”. We tried to increase the “genomeSize” parameter to 1000m for sheepA and got identical results. We ran metaFlye with “flye --pacbio-hifi reads.fa --plasmids --meta”. Hifiasm-meta and metaFlye report assembly time and peak memory usage. We used a script (<https://gist.github.com/xfengnefx/d4abf19de8ebae9cc8ccd56e9898604d>) to check /proc/ID/status files to measure the performance of HiCanu. For general file manipulations, we used seqtk (<https://>

github.com/lh3/seqtk, 1.3-r107-dirty), readfq.py (<https://github.com/lh3/readfq>, 7c04ce7), GNU Parallel²⁰ and SAMtools²¹.

Metagenome binning.

We used MetaBAT2 for initial binning and then post-process MetaBAT2 results to get final MAGs. We aligned raw reads to an assembly with “minimap2 -ak19 -w10 -I10G -g5k -r2k --l-j-min-ratio 0.5 -A2 -B5 -O5,56 -E4,1 -z400,50 contigs.fa reads.fa”²², calculated the depth with “jgi_summa_rsize_bam_contig_depths --outputDepth depth.txt input.bam”, and ran MetaBAT2 with “metabat2 --seed 1 -i contigs.fa -a depth.txt”. We tried different random seeds or “-s 500000”, and got similar results. We only applied MetaBAT2 to the primary hifiasm-meta and HiCanu assemblies as including alternate assembly led to worse binning. After MetaBAT2 binning, we separate circular contigs of 1Mb or longer into a separate MAG if it is binned together with other contigs.

Evaluating assemblies of mock metagenome libraries.

To evaluate the quality of assemblies, we mapped contigs with “minimap2 -cxasm20” to the reference genomes and inspected structural changes in the alignment. Out of 22 circular hifiasm-meta contigs assembled from the two mock communities, 21 are consistent with the reference, except for the assembly of *S. mutans*. For this genome, hifiasm-meta introduced a 20kb deletion that is supported by a small fraction of reads in alignment, suggesting this is a real but rare allele in the community.

In the contig-to-reference alignment, we observed up to several thousand mismatches and gaps per genome (Table S1). The number of small differences is much lower between HiFi assemblies. For example, for *N. meningitidis*, there are 6,090 small differences between the hifiasm-meta contig and the reference genome, but there are only two small base-pair differences between the hifiasm-meta and HiCanu contigs. We suspect most of these 6,090 differences may be consensus errors in the reference genome.

Evaluating metagenome assemblies.

We ran CheckM v1.1.3 to measure the completeness and the contamination level of MAGs. The command line is “checkm lineage_wf -x fa input/wd/; checkm qa -o 2 wd/lineage.ms .”. We also tried DAS Tool²³ for evaluation on the sheepA dataset. DAS Tool is more optimistic, identifying 22% more near-complete MAGs in comparison to CheckM. As CheckM is more often used for evaluation, we only applied CheckM to all assemblies. For sheepB, additionally, yak QV (<https://github.com/lh3/yak>) was used to evaluate contig correctness (Extended Data Fig. 2).

We used GTDB-Tk v1.3.0 and its database version r95 for phylogenetic placement with command line “gtdbtk classify_wf”. We annotated the tree and used GraPhlAn for visualization.

We used INFERNAL²⁴ to identify rRNAs and tRNAs from contigs. The command line is “cmscan -cut_ga -rfam -nohmmonly -fmt 2 -tblout cmscan.table Rfam.cm in.fa”. Entries with E-value larger than 0.01 were dropped. 733 of the 738 long circular contigs assembled


by hifiasm-meta in this manuscript were RNA-complete (Table S5), i.e. having at least one full length copy for all three types of rRNAs, and at least 18 full length copies of tRNAs.

Data availability


HiFi data were obtained from NCBI Sequence Read Archive (SRA) with accession numbers shown in Table 1. All generated assemblies and underlying data of the figures are available at <https://zenodo.org/record/6330282>. ZymoBIOMICS mock reference genomes were downloaded from <https://s3.amazonaws.com/zymo-files/BioPool/D6331.refseq.zip>. The list of reference genomes in the ATCC mock community is available at <https://www.atcc.org/products/msa-1003>. CheckM database: https://data.ace.uq.edu.au/public/CheckM_databases/checkm_data_2015_01_16.tar.gz. GTDB-Tk database: https://data.ace.uq.edu.au/public/gtdb/data/releases/release95/95.0/auxillary_files/.

Code availability

Hifiasm-meta is available at <https://github.com/xfengnefx/hifiasm-meta>.

Extended Data**Extended Data Figure 1.**

The hifiasm-meta assembly graph of the sheepA dataset. Short disconnected contigs are not shown.



Extended Data Figure 2.

Yak QV score correlated with contig coverage. Plots showing >1Mb contigs in sheepB assemblies. Contig coverage was estimated by `jgi_summarize_bam_contig_depths` from `metabat2`, and alignment was done with `minimap2 -a -k 19 -w 10 -I 10G -g 5000 -r 2000 -lj -min-ratio 0.5 -A 2 -B 5 -O 5,56 -E 4,1 -z 400,50`. Hifiasm-meta assembled 37 contigs without k-mer errors. HiCanu and metaFlye each assembled 4 contigs without k-mer errors.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank Wei Fan from Agricultural Genomics Institute at Shenzhen, China, for sharing the chicken gut dataset. This study was supported by US National Institutes of Health (grant R01HG010040 and U01HG010971 to H.L.). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

References

1. Lapidus AL & Korobeynikov AI. Metagenomic data assembly - the way of decoding unknown microorganisms. *Front Microbiol* 12, 613791 (2021).
2. Almeida A. et al. A new genomic blueprint of the human gut microbiota. *Nature* 568, 499–504 (2019). [PubMed: 30745586]
3. Chen L-X, Anantharaman K, Shaiber A, Eren AM. & Banfield JF. Accurate and complete genomes from metagenomes. *Genome Res* 30, 315–333 (2020). [PubMed: 32188701]
4. Kang DD. et al. MetaBAT 2: an adaptive binning algorithm for robust and efficient genome reconstruction from metagenome assemblies. *PeerJ* 7, e7359 (2019). [PubMed: 31388474]
5. Kolmogorov M. et al. metaFlye: scalable long-read metagenome assembly using repeat graphs. *Nat. Methods* 17, 1103–1110 (2020). [PubMed: 33020656]
6. Kolmogorov M, Yuan J, Lin Y. & Pevzner PA. Assembly of long, error-prone reads using repeat graphs. *Nat Biotechnol* 37, 540–546 (2019). [PubMed: 30936562]
7. Nurk S. et al. HiCanu: accurate assembly of segmental duplications, satellites, and allelic variants from high-fidelity long reads. *Genome Res.* 30, 1291–1305 (2020). [PubMed: 32801147]
8. Cheng H, Concepcion GT, Feng X, Zhang H. & Li H. Haplotype-resolved de novo assembly using phased assembly graphs with hifiasm. *Nat. Methods* 18, 170–175 (2021). [PubMed: 33526886]

9. Cao C. et al. Reconstruction of microbial haplotypes by integration of statistical and physical linkage in scaffolding. *Molecular biology and evolution* 38, 2660–2672 (2021). [PubMed: 33547786]
10. Hui J, Shomorony I, Ramchandran K. & Courtade TA. Overlap-based genome assembly from variable-length reads. In *IEEE International Symposium on Information Theory, ISIT 2016, Barcelona, Spain, July 10–15, 2016*, 1018–1022 (IEEE, 2016).
11. Parks DH, Imelfort M, Skennerton CT, Hugenholtz P. & Tyson GW. CheckM: assessing the quality of microbial genomes recovered from isolates, single cells, and metagenomes. *Genome Res.* 25, 1043–1055 (2015). [PubMed: 25977477]
12. Bowers RM. et al. Minimum information about a single amplified genome (misag) and a metagenome-assembled genome (mimag) of bacteria and archaea. *Nature biotechnology* 35, 725–731 (2017).
13. Ondov BD. et al. Mash: fast genome and metagenome distance estimation using minhash. *Genome biology* 17, 1–14 (2016). [PubMed: 26753840]
14. Bickhart DM. et al. Generating lineage-resolved, complete metagenome-assembled genomes from complex microbial communities. *Nature biotechnology* 1–9 (2022).
15. Moss EL, Maghini DG. & Bhatt AS. Complete, closed bacterial genomes from microbiomes using nanopore sequencing. *Nat Biotechnol* 38, 701–707 (2020). [PubMed: 32042169]
16. Vicedomini R, Quince C, Darling AE. & Chikhi R. Strawberry: automated strain separation in low-complexity metagenomes using long reads. *Nat Commun* 12, 4485 (2021). [PubMed: 34301928]
17. Hon T. et al. Highly accurate long-read HiFi sequencing data for five complex genomes. *Sci Data* 7, 399 (2020). [PubMed: 33203859]
18. Chaumeil P-A, Mussig AJ, Hugenholtz P. & Parks DH. GTDB-Tk: a toolkit to classify genomes with the genome taxonomy database. *Bioinformatics* 36, 1925–1927 (2019).
19. Asnicar F, Weingart G, Tickle TL, Huttenhower C. & Segata N. Compact graphical representation of phylogenetic data and metadata with graphlan. *PeerJ* 3, e1029 (2015). [PubMed: 26157614]

References

20. Tange O. GNU Parallel - the command-line power tool. *The USENIX Magazine* 36, 42–47 (2011).
21. Li H. et al. The sequence alignment/map format and samtools. *Bioinformatics* 25, 2078–2079 (2009). [PubMed: 19505943]
22. Li H. Minimap2: pairwise alignment for nucleotide sequences. *Bioinformatics* 34, 3094–3100 (2018). [PubMed: 29750242]
23. Sieber CMK. et al. Recovery of genomes from metagenomes via a dereplication, aggregation and scoring strategy. *Nature microbiology* 3, 836–843 (2018).
24. Nawrocki EP. & Eddy SR. Infernal 1.1: 100-fold faster rna homology searches. *Bioinformatics* 29, 2933–2935 (2013). [PubMed: 24008419]

four human gut samples. Each bar represents a contig of 1Mb in length. It gives the number of reads used in the contig. Each color corresponds to a human gut sample. A cross at the top of a bar indicates the contig being circular. (d) Phylogeny of human gut MAGs from the co-assembly. A colored clade corresponds to a phylum inferred by GTDB-Tk^{18,19}. A MAG is omnivore/vegan-specific if 90% of reads in the MAG come from omnivore/vegan samples, respectively. Two adjacent leaves coming in the same genus have the same shade on the “Genus” outer ring.

Table 1.

Evaluated metagenome datasets

Sample	Accession	# bases (Gb)	N50 read length (kb)	Median read QV	Sample description
ATCC	SRR11606871	59.2	12.0	36	Mock community ATCC MSA-1003 ¹⁷
zymo	SRR13128014	18.0	10.6	40	Mock community ZymoBIOMICS D6331
sheepA	SRR10963010	51.9	14.3	25	Sheep gut microbiome ⁵
sheepB	SRR14289618	206.4	11.8	N/A	Sheep gut microbiome ¹⁴
chicken	SRR15214153	33.6	17.6	30	Chicken gut microbiome
sludge	ERR7015089	15.3	15.4	32	Anaerobic digester sludge
humanO1	SRR15275213	18.5	11.4	40	Human gut from a pool of four omnivore samples
humanO2	SRR15275212	15.5	10.3	41	Human gut from a pool of four omnivore samples
humanV1	SRR15275211	18.8	11.0	39	Human gut from a pool of four vegan samples
humanV2	SRR15275210	15.2	9.6	40	Human gut from a pool of four vegan samples

The N50 read length is the length of the shortest read at 50% of the total number of read bases. The quality value (QV) of a read is $-10\log_{10}e$, where e is the expected sequencing error rate of the read, assuming accurate base quality. No base quality is available for the sheepB dataset.