

RESEARCH

Ariadne: Synthetic Long Read Deconvolution Using Assembly Graphs

Lauren Mak^{1,2}, Dmitry Meleshko^{1,2}, David C. Danko^{1,2}, Waris N. Barakzai³, Natan Belchikov⁴ and Iman Hajirasouliha^{2,5*}

Abstract

Background: *De novo* assemblies are critical for capturing the genetic composition of complex samples. Synthetic Long Read sequencing techniques such as 10x Genomics' Linked-Reads, UST's TELL-Seq, Loop Genomics' LoopSeq, and BGI's Long Fragment Read combines 3' barcoding with standard short-read sequencing to expand the range of linkage resolution from hundreds to tens of thousands of base-pairs. The application of SLR sequencing to genome assembly has demonstrated that barcoding-based technologies balance the tradeoffs between long-range linkage, per-base coverage, and costs. However, multiple long fragments may be associated with the same 3' nucleotide-based unique molecular identifier (UMI). The lack of a 1:1 correspondence between a long fragment and a UMI, in conjunction with low sequencing depth, confounds the assignment of linkage between short-reads.

Results: We introduce Ariadne, a novel SLR deconvolution algorithm based on assembly graphs, that can be used to extract single-species read-sets from a large SLR dataset. Ariadne deconvolution of SLR clouds increases the proportion of read clouds containing only reads from a single fragment by up to 37.5-fold. Using these enhanced read clouds in *de novo* assembly significantly improves assembly contiguity and the size of the largest aligned blocks in comparison to the non-deconvolved read clouds.

Conclusions: Integrating Ariadne into the post-processing pipeline for SLR technologies increases the quality of *de novo* assembly for complex populations, such as microbiomes.

Keywords: Linked-read sequencing; Assembly graphs; Metagenomics; Barcode deconvolution

1 Background

Next generation sequencing technologies underpin the large-scale genetic analyses of complex genomic mixtures. However, reconstruction of multiple genomes or identification of structural variations is more computationally demanding than assembling a single haploid genome.

Though standard Illumina short-read sequencing is the most popular platform for genome-wide characterizations due to its sequencing depth to cost ratio, the length of Illumina short-reads limits the linkage information that can be extracted from sequencing libraries. To address this, researchers are increasingly using nucleotide barcode-based, chromatin conformation-based, or long-read sequencing technologies to approximate single-cell resolution from complex mixtures of microorganisms. Long-read sequenc-

ing technologies, such the range of options from Oxford Nanopore Technologies, are capable of resolving a limited number of high-quality circularized draft genomes from metagenomic sequencing data [1]. However, high-quality assemblies of long reads are reliant on large amounts of input DNA to adequately cover all of the species of the sample, which may not be available depending on its size and alpha diversity [2]. When sequencing depth is low, the assembly graph contains many ambiguous junctions comprised of sequence material from multiple genetically similar species and/or strains [1].

1.1 Overview of Synthetic Long Read (SLR) Technology
Alternative approaches that rely on short-read sequencing while still generating long-range linkage information include Hi-C and synthetic long-read (SLR) technologies [3]. While similar ambiguous junctions are inherent in short-read assemblies due to limitations in sequencing insert size, *de novo* assemblies generated

*Correspondence: imh2003@med.cornell.edu

²Institute for Computational Biomedicine, Department of Physiology and Biophysics, Weill Cornell Medicine of Cornell University, New York, USA
Full list of author information is available at the end of the article

by short-read and synthetic long read (SLR) are inherently larger and more comprehensive representations of the sample composition due to the relative volume of genomic coverage. Since SLR library preparations require less input material than long reads, it is much more suitable for extracting sequencing information from low-volume samples [1, 4, 5, 6].

SLR technologies, starting with Illumina TruSeq Synthetic Long reads, [7], associate short reads originating from the same extracted genomic fragment with some type of nucleotide-based unique molecular identifier (UMI). UMIs are colloquially referred to as 'barcodes' in the context of SLR sequencing. Briefly, input (meta)genomic DNA is sheared into long fragments of 5-100 kbp. After shearing, a UMI (usually 16 - 20 base-pairs long) is ligated to short-reads from the fragments such that short-reads from the same fragment share the same UMI. Linked-read UMIs are unrelated to the 5' UMI used for sample multiplexing. We refer to the set of reads that share a UMI as a read cloud. Finally, the short-reads are sequenced using standard sequencing technologies (e.g. Illumina HiSeq). SLRs offer additional long-range information over standard short-reads. Reads with matching UMIs are more likely to have emerged from the same long fragment of DNA than two randomly sampled reads, which extends the relative positional information encoded in the read's short 50 - 250 bp sequence past the standard limitations of a short-read insert, which are typically several hundreds of base-pairs. There is a slight coverage tradeoff due to the size of the UMIs and associated library preparation costs relative to standard short-read sequencing.

SLR techniques are differentiated by the biochemical mechanism that associates UMIs with genomic fragments or short reads. 10x linked-read sequencing uses oil-based droplets to encapsulate a few genomic fragments and a UMI, which are subsequently splintered into short reads which are amplified with the UMI [8]. While 10x Genomics' linked reads has been discontinued after nearly half a decade of widespread usage, a variety of other UMI-based SLR methods such as TELL-Seq, LoopSeq, and BGI's Long Fragment Reads (stLFR) have been commercialized recently with the promise of (near-)single-molecule resolution along with simplified library preparation procedures and compatibility with standard Illumina sequencing machines. TELL-Seq uses proprietary TELL beads, which both capture genomic fragments and use transposon-based reactions to insert the UMI sequence throughout the fragment [9]. TELL fragments are approximately 20 - 40 kbp long [9], similar to the (on average) 10 kbp fragments generated by 10x [10]. LoopSeq similarly uses an intramolecular enzyme-based distribution method, but

does not use beads to partition genomic fragments at all [11]. One major advantage of both TELL-Seq and LoopSeq over 10x and BGI's Long Fragment Reads is their single-tube reaction chemistry, and compatibility with standard Illumina sequencing machines. 10x library preparation depends on costly instruments for droplet management [8], and BGI's Long Fragment Reads are incompatible with Illumina sequencing. For detailed explanations of each method's capabilities and library preparation procedures, we direct the reader to [8, 9, 11, 12]. In this study, we also provide a first head-to-head comparison of multiple SLR technologies to profile their strengths and weaknesses on metagenomics datasets of similar complexity.

The central drawback of SLR technologies is that the UMI-based linkage information must be efficiently interpreted from the sequencing data to simulate long read resolution and increase the average contiguity (i.e.: NA50) of *de novo* assembly. The need for novel algorithms to leverage this information has been partially met by the proliferation of SLR-based assemblers, such as Athena, Supernova, and cloudSPAdes [13, 14, 15]. However, additional novel algorithmic efforts are needed to achieve the desired contiguity and reduce the number of observed assembly errors [15]. Many taxonomic lineages identified in large-scale studies are not represented in reference sequence databases, and are not associated with isolated cultures [16]. Metagenomic analyses that rely on existing reference genomes, such as read-based taxonomic classification, will inherently bias the genomic reconstruction of a mixed population, generally towards the most common and already well-characterized species within the sample [17].

1.2 Applications of SLR Sequencing in Metagenomics

SLR sequencing has been shown to resolve species compositions of metagenomics samples in both 16S-based and shotgun whole-genome-based analyses. Because of its low UMI multiplicity, LoopSeq was able to identify multiple copies of the 16S rRNA gene as belonging to a single strain [18, 19]. When paired with additional sources of information, such as targeted amplification or longitudinal sequencing, SLRs are capable of resolving metagenomics to the strain level. In conjunction with high-throughput qPCR, LoopSeq has been used to associate antimicrobial resistance genes with specific species in environmental samples [20]. The long-read-like linkage information encoded in read clouds has been used to track single nucleotide variants in the human gut microbiome longitudinally, demonstrating that prioritizing depth of coverage over strict read length can be an optimal analysis strategy especially when the number of strains/haplotypes is an

unknown variable [10]. In another study, SLRs have also been used to identify the presence of structural variation on bacterial chromosomes [21].

1.3 Challenges Posed by SLR Sequencing

Despite the additional linkage information offered by SLR sequencing, there are new computational challenges involved in applying barcoded reads to *de novo* assembly. Because metagenomic samples are intrinsically multiplexed samples of multiple species, long-range linkage information is confounded by the multiplicity of fragments assigned to 3' UMIs. Existing systems employ on the order of 10^{6-7} 3' UMIs [22]. In previous studies with the 10x Genomics system, it was observed that there were 2–20 long fragments of DNA per 3' UMI [23]. The larger the barcoded read cloud, the more likely that reads tended to originate from multiple fragments. Our analyses suggest that at least 97% of read clouds are composed of ≥ 2 fragments, with the exception of a LoopSeq dataset. In the absence of additional information about the sample, it is difficult to distinguish the genomic origin of a random assortment of reads with the same UMI. Furthermore, each fragment of DNA is only fractionally covered by reads (typically 10–20%). Because overall coverage is reduced, SLRs provide long-range information at the expense of short-range blocks of contiguous sequence.

1.4 The Barcode Deconvolution Problem

The barcode deconvolution problem, previously described in [24, 23], is defined as the assignment of each read with a given 3' UMI to a subgroup such that every read in the subgroup originates from the same fragment or contiguous genomic segment. A solution to the barcode deconvolution problem for a set of read clouds would be a map from each read cloud to a function which solves the UMI deconvolution problem for that read cloud. Reads with the same 3' UMI that are highly likely to have originated from the same metagenomic fragment are more likely to co-occur in one another's sequence space within the assembly graph than reads from different fragments.

Though the multiplicity of genomic fragments to 3' UMIs has been problematic since the inception of SLR sequencing, the UMI deconvolution problem has only been addressed recently by two computational tools: EMA [24] and Minerva [23]. The EMA approach augments read alignment for barcoded reads based on alignment to a reference sequence. In the process of generating probabilistic alignments, reads from a single read cloud are sub-grouped into sets of reads that map close to each other. EMA is particularly applicable for highly repetitive regions where a barcoded read can align to multiple locations within the genome

[24]. However, the EMA approach relies on the user to supply reference sequence(s) as *a priori* information about the input sample. We also demonstrate that EMA (as of the date of publication) is unable to recognize > 99.9% of SLR UMIs, even with UMI whitelists tailored for each dataset. Additionally, is an extension of the EMA approach that makes a graph of UMIs using the degree of read connectivity between UMIs as edges to deconvolve read clouds [25]. While the species composition of popular sampling sites such as the human gut are well-characterized, such reference-based methods are not designed for microbial samples from under-studied environments such as urban landscapes [26]. Minerva does not require a reference genome, using instead *k*-mer similarities between read clouds to approximately solve the UMI deconvolution problem for metagenomic samples [23]. However, Minerva is memory-inefficient and requires extensive parameter optimization to deconvolve all of the read clouds in a dataset.

In this paper, we present Ariadne as an advanced approach to tackle UMI deconvolution. Instead of using read alignments to reference sequences, which are unknown for poorly characterized environmental microbial samples [26], or relying on computationally expensive string-based graphs, Ariadne leverages the linkage information encoded in the full de Bruijn-based assembly graph generated by a *de novo* assembly tool such as cloudSPAdes [15] to generate up to 37.5-fold more read clouds containing only reads from a single fragment, improve the summed NA50 by up to 500 kbp, and maintain a proportional rate of mis-assembly relative to *de novo* assembly without prior UMI deconvolution. Searching through the pre-made assembly graph for reads in the neighboring sequence space makes the search for co-occurring reads computationally tractable, generalizable, and scalable to large datasets.

2 Results

2.1 Benchmarking Datasets

We used real data sets from four microbial mock communities which we refer as MOCK5 10x, MOCK5 LoopSeq, MOCK20 10x, MOCK20 TellSeq throughout this manuscript. See Table 1 for an overview of the mock microbiome datasets and Table 2 for the relative species abundances. Two of the datasets, MOCK5 LoopSeq and MOCK20 TELL-Seq, are analysed for the first time in this work. Note that while both comprised of 5 species, the MOCK5 10x and MOCK5 LoopSeq datasets share only one species—*Escherichia coli*—and thus cannot be considered as a comparison of 10x and LoopSeq technology. The original MOCK20 10x dataset is approximately 330 million reads, but

for efficiency a random subset of 100 million reads was used in the subsequent analyses. Similarly, the original MOCK20 TELL-Seq dataset is approximately 210 million reads large, but was also subsetted to 100 million reads. Reference genome sequences can be downloaded [here](#).

2.2 Gold-Standard ('Reference') Cloud Deconvolution

The use of mock communities allowed us to infer the genomic fragment of origin of the barcoded reads. These gold-standard fragment assignments served as a benchmark to compare read clouds and assemblies with and without deconvolution. We mapped the reads to the reference sequences of the species that were known to form the mock communities using Bowtie2 [27] (version 2.3.4.1). Using the read mapping positions along the reference genome, we further subsetted the reads into gold-standard read clouds such that the left- and right-most starting positions of reads in the same cloud were no further than 200 kbp apart, in case multiple fragments from the same genome were present in the same read cloud. We termed this method 'reference deconvolution' as it represents the maximum likelihood inference of the genomic fragments that originated reads tagged with the same 3' UMI.

2.3 Ariadne Generates a Large Number of High-Quality Enhanced Read Clouds

Ariadne generated enhanced read clouds, or subgroups of original read clouds, that largely corresponded to individual fragments of DNA. We measured the quality of each enhanced read cloud using two metrics: Shannon entropy index $H = -\sum p_i \log p_i$ and purity $P = \max(\vec{p})$ where p_i indicates the proportion of reads in an (enhanced) read cloud that originates from the same most prevalent 200 kbp region in a single reference sequence. Prior to these quality checks, we excluded reads from standard and enhanced read groups of size 2 or smaller (i.e. consisting of a single read-pair or smaller), which are trivially pure.

Without deconvolution, there is a large spread of P at each read cloud size, and nearly 100% of the reads are in mixed-origin read clouds, or clouds that are comprised of reads that have likely originated from different SLR fragments (Table 3). Larger clouds are more likely to contain reads of multiple species origins and thus lower purity. The exception to this trend is the MOCK5 LoopSeq dataset (Figure 1 top row). Since the goal of linked reads is to approximate the linkage range of long fragments, having mixed-origin clouds as a result of fragment-to-UMI multiplicity confounds downstream applications such as taxonomic classification. Ariadne deconvolution reduces the proportion of multi-origin read clouds by at least 2-fold

in the MOCK5 LoopSeq dataset, up to 7.5-fold in the MOCK20 10x dataset (Table 3 column 7). Since Ariadne deconvolution generally decreases this trend (except in the case of LoopSeq), it is unlikely that our results have been inflated by a large number of small and trivially pure clouds. Furthermore, reference-based deconvolution generates a similar number of clouds of a similar size to Ariadne, indicating that search-distance based subgrouping models genomic fragment boundaries sufficiently.

Ariadne at least doubles the proportion of completely deconvolved read clouds, which represent the entirety of the reads from a single inferred genomic fragment, except in the case of the MOCK5 LoopSeq dataset (Table 3 column 8). In comparison, under-deconvolved clouds are clouds comprised of reads that originate from multiple fragments, whereas over-deconvolved clouds are comprised of single-origin reads that are a subset of all of the reads from that fragment. The proportion of single-origin read clouds, the sum of complete and over-deconvolved read clouds, has increased between 2- and 37.5-fold.

While the 10x and MOCK20 TELL-Seq size-to-purity distributions are similar, the MOCK5 LoopSeq distribution peaks at $P = 0.96$, which explains why Ariadne deconvolution has minimal effect on read cloud quality. For the rest of the datasets, with respect to size, the relative purity of Ariadne-enhanced read clouds is significantly larger than that of the original read clouds (Figure 1 bottom row). The ideal deconvolution based on reference-mapped positions is shown in the Figure 1 middle row, where there are extremely few clouds with $P < 1$. Even with deconvolution, there are still large read clouds (≥ 100 bp) in the deconvolved set, indicating that Ariadne is capable of maintaining the integrity of existing single-origin read clouds through a limited search of the assembly graph (Figure 1). Overall, applying Ariadne deconvolution to SLR datasets generates enhanced read clouds that closely resemble the size-to-purity distribution of the reference deconvolution, thereby approximating the ideal deconvolution without *a priori* information about the microbial composition of the originating data.

We have additionally demonstrated the effect of increased search distance on overall dataset quality metrics (Supp. Table 1). The number of read clouds and the average entropy increase as the search distance increases, with minimal increases in the read cloud purity along with the proportion of single-origin (complete and over-deconvolved read clouds). Given with the large increases in computational runtime, only results with the deconvolution search distance of 5 kbp are included going forward in the main text. Similar results- smaller and on average fewer-origin read

clouds- can be observed with the Shannon entropy measure (Supp. Fig. 1).

2.4 Enhanced Read Clouds Improve Metagenomic Assembly

Since Ariadne was able to deconvolve the original SLR dataset into high-quality and non-trivial enhanced read clouds, we applied Ariadne to the full 97M-read MOCK5 10x dataset, the full 75M-read MOCK5 LoopSeq dataset, and 100M randomly sampled reads from each of the MOCK20 10x and MOCK20 TELL-Seq datasets to generate enhanced read clouds as input to cloudSPAdes [15] in metagenomics mode. The assembly quality of the resulting scaffolds with reference and Ariadne deconvolution was compared to that without prior deconvolution. cloudSPAdes generates *de novo* assemblies a wide variety of sequencing data into contigs and scaffolds, and has been benchmarked on the MOCK5 10x and MOCK20 10x datasets previously [15]. As such, we have used similar metaQUAST metrics to evaluate and compare the quality of the assemblies [28].

The largest improvements are in the overall assembly contiguity and the largest alignment, demonstrating that enhanced read clouds with increased fragment specificity generate higher-quality assemblies (Figure 2). NA50 is another measure of assembly contiguity, reporting the length of the aligned block such that using longer or equal-length contigs produces half of the bases of the assembly. *De novo* assemblies generated from reference- or Ariadne-deconvolved read clouds are significantly more contiguous than those generated from the original, non-deconvolved read clouds (Figure 2). To calculate the overall performance improvements when assembling each dataset, the assembly statistic for each species was summed and the relative difference between reference- or Ariadne-enhanced scaffolds and no-deconvolution scaffolds calculated. For NA50 and largest alignment, this means the no-deconvolution summed NA50 were subtracted from that of the reference- or Ariadne-enhanced scaffolds. For the rate of misassembled bases, the number of misassembled bases was divided by the total assembly length. Then, the reference- or Ariadne-enhanced scaled misassembly rates were divided by that of the no-deconvolution scaffolds.

While the differences between the deconvolution methods and no deconvolution were minimal in the MOCK5 10x dataset (Figure 2 first column), there were significantly positive differences in NA50 and largest alignment in all other datasets. While Ariadne generally produces shorter NA50 and largest alignments than the ideal genomic deconvolution, Ariadne assemblies significantly outperform no-deconvolution

assemblies in terms of contiguity statistics (Figure 2 top and middle). However, in the case of the two MOCK5 datasets, the Ariadne assembly NA50 gains were larger than the reference deconvolution assembly gains, suggesting that the most ideal read cloud partition does not always generate the most contiguous assemblies (Figure 2 top first and second panels).

In terms of misassembly rate, Ariadne assemblies largely match no-deconvolution assemblies and significantly outperform reference deconvolution. With the MOCK20 10x dataset, the third quartile of misassembly rate by reference deconvolution is 8-fold larger than that of no deconvolution, whereas Ariadne assemblies contained up to 2-fold more at maximum (Figure 2 bottom third panel). Ariadne underperforms assembly without deconvolution in terms of largest alignment and misassembly rate in MOCK20 TELL-Seq (Figure 2 bottom last panel). However, both Ariadne- and reference-enhanced assembly obtain NA50 for 4 species that the no-deconvolution assembly failed to capture, as well as some substantial differences between Ariadne and reference-based deconvolution with respect to no deconvolution (Table 4). In other datasets, there may be species that are more easily reconstructed with read cloud deconvolution that cloudSPAdes would not otherwise find long, contiguous reference subpaths for through the assembly graph. Outlier values of NA50, largest alignment, and misassembly rates were omitted from Figure 2 for visual clarity were all from the reference deconvolution scaffolds and can be found in Supp. Table 2. There were no major changes in the fraction of reference bases that were reconstructed, the total aligned length, the mean number of mismatches, and the number of contigs.

2.5 Comparison of SLR Sequencing Technologies

While both comprised of 5 species, the only species common to the MOCK5 10x and MOCK5 LoopSeq datasets is *Escherichia coli*. As such, while they are comparable in terms of input community complexity, they cannot be treated as a comparison of SLR technologies with respect to *de novo* assembly. However, the two MOCK20 datasets were generated from the same 20-species mock community product from Zymo, and as such, it is possible to compare assembly quality metrics (Table 5).

The difference between the accurately recovered fraction of each species' genome is small (10x is on average 8% larger). There were some species (e.g. *Bacteroides vulgatus*, approx. abundance is 0.03%) that both no-deconvolution and reference-deconvolved 10x and TELL-Seq assemblies struggled to reconstruct because of their low abundances. It is noteworthy that the Ariadne-deconvolved 10x assembly managed

to reconstruct a small contig from *Schaalia odontolytica*, which both no-deconvolution and reference-deconvolved assemblies completely missed. The difference between the total alignable lengths of the assemblies is larger (13.3%), with all 10x contig-to-reference genome alignments roughly the same or significantly longer than TELL-Seq alignments. At the species level, however, there was some variability in terms of the SLR technology that reconstructed the larger alignment. This variation was probably in part due to fluctuations in coverage as well as downsampling the full sequencing read dataset for efficient comparison. For example, in the case of *Bacillus cereus*, the longest alignable 10x contig was 166,772 bp and the longest alignable TELL-Seq contig was 2,079,803 bp. However, the recovered genome fraction for both technologies was 97.03% and 98.92% respectively, indicating that while the 10x assembly of *B. cereus* was more fragmented, it was still by and large complete and gaps were likely due to coverage variability. Similar reversals can be found where the TELL-Seq assemblies of a species are more fragmented but similarly complete. Nonetheless, the 10x assembly contains far fewer contigs than the TELL-Seq assembly- 2,481 to 7,729- which is also reflected in its smaller number of larger read clouds (Table 3).

As expected with larger assemblies, the 10x assembly has significantly more misassemblies and misassembled content than the TELL-Seq data, which is probably due to the fact that more assembled bases are contained in fewer contigs. The amount of unalignable sequence content was quite small in both, and comprised < 0.2% of both assemblies.

2.6 Runtime and Performance

Ariadne consumes approximately as much memory as the initial assembly graph generation process, with the exceptions of the MOCK5 datasets (Supp. Table 3). Any computing cluster with enough RAM to conduct *de novo* assembly will have sufficient resources to also run Ariadne. The primary technical cost of running Ariadne is compute time. The search distance-based deconvolution process scales with assembly graph connectivity, which is a function of i) sequencing coverage, ii) genomic fragment length, and iii) the genetic content of the originating metagenome. When the search distance is 5 kbp or less, Ariadne's runtime is similar to that of the entire *de novo* assembly process for MOCK5 10x and MOCK20 10x datasets (Supp. Table 3). At search distances larger than 5 kbp and for MOCK5 LoopSeq or MOCK20 TELL-Seq datasets, Ariadne's runtime becomes significantly larger than the time needed to generate the assembly graph, likely due to the increased number of nodes

considered per read cloud deconvolution (Supp. Table 3). For instance, when comparing MOCK20 10x and MOCK20 TELL-Seq datasets, Ariadne deconvolution with a search distance of 5 kbp takes nearly 5 times as long to complete for MOCK20 TELL-Seq, despite the number of reads being identical and *de novo* assembly taking nearly the same amount of time.

2.7 EMA, Longranger and Lariat, and Minerva

EMA takes SLR UMIs into account to align linked reads to a reference sequence(s) using a latent variable model, thereby serving as an alternative to Bowtie2-based read alignment to generate gold-standard read cloud deconvolution. However, the first step of the EMA pipeline (last downloaded: May 7, 2021), `ema count`, which counts UMIs and partitions the original FastQ file into a number of deconvolution bins, is unable to recognize most to all of the UMIs, even when custom whitelists with the exact list of UMIs in the dataset are directly provided (Supp. Table 4). In one case, `ema count` failed to detect any UMIs in the MOCK5 LoopSeq dataset altogether. In the best-case scenario with the MOCK20 10x dataset, `ema count` recognized 199,488 barcoded reads. The core deconvolution module in the pipeline, `ema align`, would have been able to deconvolve at maximum 0.002% of a 94-million read dataset (Supp. Table 4). To deconvolve reads that are not recognized as barcoded, the EMA pipeline uses the same procedure as our reference deconvolution pipeline, with `bwa` as the read aligner instead of Bowtie2. Due to the paucity of aligned reads and the subsequent lack of deconvolution, EMA-enhanced read clouds were not featured in this analysis. For ease of use, we decided to use our own reference-based deconvolution pipeline, which automates all of the read alignment, read-subgrouping, and FastQ generation steps with a single submission command. Similar UMI recognition issues were encountered with the Longranger align pipeline and the Lariat aligner it was based on. Similar to EMA, Lariat incorporates UMI information to align linked reads to a reference sequence(s). However, the available FastQ files for all four of the datasets were not comprised of raw Illumina BCL files or the raw output of `longranger mkfastq`. As such, it was not possible to apply Longranger align or Lariat to the four datasets. Minerva was similarly unable to deconvolve a sufficient number of read clouds (0.002% at best with the MOCK5 LoopSeq dataset), and it too was not featured in this full analysis. However, both Ariadne and Minerva completed a deconvolution test-run of 20-million reads from the MOCK5 10x dataset. In summary, while Minerva generated nearly 100% single-species read clouds, it was only able to deconvolve 4% of the reads in total (Supp.

Figs. 2 and 3, Supp. Table 4). Performance-wise, the Minerva runs were 4 times as long and consumed 3 times as much memory. The results are explained in greater detail in the Supplementary Materials (pg. 6 and 7).

3 Discussion

There have been other recent developments in the SLR space, intended to extract as much linkage information from barcoded reads as possible in order to maximize recovery of input genomic information. Instead of tackling the UMI deconvolution problem, Guo *et al.* [29] and Weng *et al.* [30] innovate downstream of the *de novo* assembly problem. By using SLRs in combination with graph- or *k*-mer-based methods, SLRsuper-scaffolder and IterCluster attempt to generate longer and higher-quality assemblies. Either of these, paired with the largely single-origin read clouds generated by Ariadne, could potentially improve the NA50 and average alignment size generated by cloudSPAdes. Hybrid sequencing and analysis strategies in the future may still take advantage of the coverage and linkage depth of SLR datasets, even when used in combination with long reads [31, 32].

Ariadne deconvolution generates enhanced read clouds that are up to 37.5-fold more likely to be single-origin, which will improve downstream applications that depend on approximating the long-range linkage information from a single species, such as taxonomic classification. *De novo* assemblies of mock metagenome communities are significantly more contiguous with Ariadne-enhanced read clouds, without the outsized increase in misassembly rate as observed with the ideal deconvolution strategy. In terms of the dataset-specific results, increasing the number of species did not change the shape of the size-to-purity distribution greatly, reflecting the inter-technology similarities in UMI-to-fragment multiplicity. As such, Ariadne is capable of improving assembly results and read cloud composition across all of the SLR strategies, and unlike EMA or the Longranger pipeline, easily facilitates re-analyses of existing SLR datasets for higher-resolution *de novo* assembly and taxonomic classification without pre-existing knowledge of the originating species. As with all other assembly-based algorithms, the degree of assembly contiguity will large depend on i) the true sample composition, which determines the intrinsic genetic heterogeneity to be resolved, and ii) the amount of raw sequencing data generated, which is a function of DNA extraction efficiency (i.e.: genomic fragment size) and sequencing coverage [33].

Though Ariadne relies on cloudSPAdes parameters to generate the assembly graph (e.g., iterative *k*-mer

sizes), the program by itself only has two: search distance and size cutoff. The maximum search distance determines the maximum path length of the Dijkstra graphs surrounding the focal read. Since each read is modelled as the center of a genomic fragment, the search distance can be thought of as the width of the fragment. As such, it should be set as the mean estimated fragment length, as determined using other means such as *a priori* knowledge of shearing duration and intensity or mapping reads with the same UMI to known reference genomes. While the estimated mean length of a metagenomic fragment according to [15] is approximately 40 kbp, the most balanced results were obtained with a significantly shorter search distance. Of the distances tried in this analysis, to balance the highest-quality assembly results, single-origin read clouds, and computational efficiency, we recommend that the user set a search distance 5 kbp or shorter to generate their own enhanced read clouds. If the user has additional information that their SLR data is comprised of larger read clouds that are nearly single-origin, such as in the MOCK5 LoopSeq dataset, then larger size cutoffs should be tried. Similarly, if the goal is to generate as much alignable genetic material as possible without as much concern for synteny or the relative spacing of genes along the genome, then larger search distances can be tried. Reference-based deconvolution, which directly approximates genomic fragments by mapping reads to the the known reference sequence composition of the sample, is better for large contiguous aligned blocks. However, the high misassembly rate is detrimental to the overall integrity of the assembly and should be carefully applied even in cases where the species composition of the sample is known. For some SLR technologies, the average genomic fragment size must be estimated prior to library preparation to ensure the correct balance of DNA and reagent molarity [9, 8, 11]. This average fragment size serves as a useful upper bound of an appropriate search distance.

This study is the first to compare the performances of multiple SLR technologies on well-characterized mock microbiomes. While 10x and TELL-Seq assemblies on a 20-species community are comparable in terms of the fraction of recovered species genomes, the 10x assembly was consistently larger, more accurate, and more contiguous than the TELL-Seq assembly while only slightly more prone to misassembly. LoopSeq is an interesting case where the overall fraction of recovered species genomes was low but the overall read cloud purity was extremely high, leading to smaller but extremely contiguous (i.e.: large alignable contigs) assemblies. While 10x is a well-characterized and thoroughly validated all-around choice, it may be

advantageous to use LoopSeq on microbiomes composed of a few bacteria with small genomes. As with setting search distances, we advise researchers comparing SLR methods to consider the known complexity of their samples and to choose accordingly.

4 Conclusion

We have developed Ariadne, a novel SLR deconvolution algorithm based on assembly graphs that addresses the UMI deconvolution problem, and enables the complete usage of the linkage information in SLR technology. Ariadne deconvolution has the largest impact when the input microbial community is large and complex, which is ideal for environmental microbial samples with minimal prior characterization [26]. In its current form, Ariadne is implemented as a module of cloudSPAdes to directly ingest the generated assembly graph and generate enhanced read clouds. Further algorithmic improvements to maximize the correspondence between 3' UMIs and fragments may bring about significant improvements in assembly quality, and increase its scalability to other large-scale SLR problems, such as haplotype phasing.

5 Methods

5.1 Algorithm Overview

We have developed Ariadne, an algorithm that approximately solves the UMI deconvolution problem for SLR sequencing datasets. Ariadne deconvolves read clouds by positioning each read on a de Bruijn-based assembly graph, and grouping reads within a read cloud that are located on nearby edges of the assembly graph. The grouped reads are termed enhanced read clouds. Currently, Ariadne is implemented as a module of cloudSPAdes version 3.12.0.

5.1.1 The Usage of UMI Information in De Novo Assembly

The following is a summary of cloudSPAdes' usage of UMIs to identify contigs from assembly graphs. For a more detailed description, see the Materials and Methods section of [15]. cloudSPAdes constructs and iteratively simplifies a de Bruijn assembly graph using sequencing reads and several sizes of k-mers (by default, 21, 33, and 55 bp). The goal is to recover genomic cycles through the assembly graph that correspond to whole chromosomes. Due to read error, incomplete sequencing coverage, limited sequencing depth, and biological phenomena such as repetitive regions, de novo assemblers employ a number of heuristics to find the optimal unbranching paths through the graph [13, 14, 15]. In the case of SLRs and cloudSPAdes, UMIs provide one mechanism of identifying edges that are likely part of the same genomic cycle [15]. If read

i carrying UMI b aligns to edge i , then edge i is said to be associated with UMI b . The UMI similarity between two edges is the proportion of associated UMIs in common. By using UMI similarities to identify edges that likely originated from the same genomic region, cloudSPAdes connects edges within the assembly graph to approximate chromosomes.

Due to fragment-to-UMI multiplicity, barcoded metagenomic read clouds are likely comprised of reads from a few species. Since these reads are all carrying the same UMI b , the long edges associated with UMI b may be erroneously connected to form contigs with genomic material from multiple species. Chimeric contigs are hazardous for downstream analyses, since as the binning of contigs into metagenomic-assembled genomes [6, 34].

5.1.2 The Assembly Graph Conception of a Fragment

Searching through the cloudSPAdes-generated assembly graph for reads in the neighboring sequence space is equivalent to approximating the sequence content of a genomic fragment given the entirety of the sequence content in the input read dataset. By necessity, reads originating from the same fragment must i) have the same 3' UMI and ii) be no more further apart than the total length of the fragment. The user can specify the search distance according to *a priori* knowledge about the median size of genomic fragments generated in the first step of SLR sequencing.

This process is diagrammed in Figure 3. Reads can be mapped to a de Bruijn graph. A read (Figure 3A blue and red bars) is said to coincide with an edge of the assembly graph (Figure 3B) if the read's sequence aligns to the edge's sequence. The read can also be said to coincide with the vertices bordering the edge it coincides with. For example, in Figure 3C, read i , which consists of string 'TGACTGC' coincides with an edge i that also contains the string 'TGACTGC'.

5.1.3 The Dijkstra Graph Conception of a Fragment

The Dijkstra graphs (Figure 3D) comprising the nearby assembly graph for each read represent the potential sequence space that the read-originating fragment occupies on that assembly graph. A Dijkstra graph contains the shortest paths from source node to all vertices in the given graph. The Dijkstra graphs for a read are comprised of the vertices bordering assembly graph edges that are reachable within the maximal search distance. The maximum search distance is a user-provided parameter limiting the size of the

Dijkstra graph, or the search space to be considered in the assembly graph. The maximum search distance reflects the user's *a priori* knowledge of the size of a fragment.

5.2 The Ariadne Algorithm

Ariadne requires a prior step to locate reads within the (meta)genome of the sample. Where EMA requires an alignment step using the bwa read mapper to generate initial mappings for its barcoded reads [24, 35], Ariadne uses the raw assembly graph generated by any of the SPAdes family of *de novo* assemblers to identify the locations of reads relative to one another in the sample's sequence space. Though Ariadne is intended to be a standalone tool, in the future, it will be possible to integrate Ariadne as an intermediary step in the *de novo* assembly procedure to improve the assembly graph *in situ*.

The assembly graph is first generated by applying cloudSPAdes to the raw SLR metagenomics reads as described in [36]. For a more detailed explanation of the process of forming the assembly graph, we refer the reader to [36, 37]. Subsequently, the Ariadne deconvolution algorithm is applied to the raw reads, with the cloudSPAdes-generated assembly graph supplying potential long-range linkage connections derived from the dataset.

5.2.1 Step 1: Extract the assembly graph from a cloudSPAdes run.

A simple representation of an assembly graph is depicted in (Figure 3B). In the cloudSPAdes assembly procedure, the assembly graph is obtained from the raw assembly graph by condensing each non-branching path into a single edge and closing gaps, removing loops, bulges, and redundant contigs. In the process of constructing the assembly graph, cloudSPAdes also maps each read to the assembly graph, generating the mapping path of a read, or the edges in the assembly graph either partially or fully spanned by the read (Figure 3C read i). The mapping path P_i , or finite walk, of read i is comprised of the set of edges e from graph G that read i covers (Equation 1).

$$P_i = \{e_1, e_2, \dots, e_n\} \quad (1)$$

The mapping path can equivalently be described as a vertex sequence V_i , which is composed of the set of vertices that border each of the edges in P_i (Equation 2).

$$V_i = \{v_1, v_2, \dots, v_{n+1} : \phi(e_i) = v_i, v_{i+1}\} \quad (2)$$

Importantly, the vertices and edges in the assembly graph are unique numeric indices replacing their sequence content. Instead of having to compare the sequences comprising the assembly graph edges, reads sharing edges can be identified by the indices in their mapping paths. As such, for the purposes of UMI deconvolution, reads do not need to be re-mapped along the assembly graph. Instead, the index-based mapping paths and vertex sequences of each read are used to locate the read. This represents a considerable speed-up over the Minerva procedure, which relies on hashing string-based sequence comparisons between reads and read clouds. Steps 2 and 3 are trivially parallelized such that the deconvolution procedure processes as many read clouds as there are threads available.

5.2.2 Step 2: Generating connected read-sets for each read i .

If the number of reads in the read cloud is greater than the user-set size cutoff, the read cloud (i.e.: tagged with the same 3' UMI) is loaded into memory (Figure 3A). For each read i , the following steps are conducted to identify other reads with the same 3' UMI that potentially originated from the same fragment.

Step 2A: Generating forward and reverse Dijkstra graphs. In Figure 3, read i aligns to an edge in the de Bruijn assembly graph. Assuming that read i is oriented in the 5' \rightarrow 3' direction, the 3'-terminal vertex is the 3'-most sequence that read i is contiguous with. By finding edges reachable within search distance d of the 3'-terminal vertex, we will be able to all reads with the same UMI that are likely to originate from the same fragment. Reads that are oriented in the 3' \rightarrow 5' direction can be reverse-complemented to apply this same procedure.

To facilitate this search, a forward Dijkstra graph (Figure 3D) is constructed starting at read i 's 3'-terminal vertex. The forward Dijkstra graph $D_{f,i}$ comprises the set of vertices v_k , which are all vertices in the assembly graph reachable within the maximum search distance d the 3'-terminal vertex (Equation 3).

$$D_{f,i} = \{v_k : \min(\text{distance}(v_j, v_k)) \leq d\} \quad (3)$$

The process of constructing the Dijkstra graph, which represents the nearby sequence space of read i , is as follows. The tentative distances to all downstream vertices in the assembly graph are set to the maximum search distance d . From the 3'-terminal vertex, the tentative distances to these vertices are calculated from lengths of the edges connecting the vertices. This value is compared to the current distance value, and the smaller of the two is assigned as the actual

distance. The vertex with the smallest actual distance from the current node is selected as the next node from which to find minimal paths, and this process is repeated. This process concludes for any path from v_j when the smallest tentative distance to vertices v_k is the maximum search distance.

Correspondingly, a reverse Dijkstra graph is constructed with the goal node set as read i 's 5'-terminal vertex, the start-vertex of the edge that the start of the read coincides with. For the reverse Dijkstra $D_{r,i}$, the 5'-terminal vertex is treated as the terminal node of Dijkstra graph construction, and the set of distances in $D_{r,i}$ is comprised of the vertices v_g such that the minimal edge-length distance between the 5'-terminal vertex and v_g is less than the maximum search distance.

Step 2B: Identifying other reads potentially originating from the same (meta)genomic fragment. Due to long-range linkage, reads from the same genomic fragment should occur in both the nearby assembly graph and each others' Dijkstra graphs. As such, all other reads j in the same read cloud are compared to read i and its forward and reverse Dijkstra graphs. For every read j in the read cloud, if one of the following conditions holds, the read j is added to the connected read-set of read i . If any vertex in the vertex sequence of j , V_j (or its reverse complement sequence) is found in the forward or reverse Dijkstra graphs (Figure 3D). This process is conducted via binary search of sorted vertex names in the forward and reverse Dijkstra graphs, eliminating the need to test for sequence string equivalency as in Minerva. Note that edge lengths can range from hundreds to tens of thousands of base-pairs in length. If read i aligned to an edge that was 10 kbp long, many reads from the same originating fragment may in fact align to the same edge as read i . Ariadne thus deconvolves read clouds by dynamically identifying the maximal sets of connected reads within each cloud.

Else, if neither of the above is true, the read is likely to have originated from a different fragment than read i (Figure 3E) despite being tagged with the same 3' UMI, and thus no connecting component is built.

The complexity of the overall read cloud deconvolution, which consists of Dijkstra graph construction and binary searches through the assembly graph, is

$$\sum_i O(E_{G_{i,d}} \log(V_{G_{i,d}})) \quad (4)$$

where G is the Dijkstra subgraph induced on the overall assembly graph with read i as the focal read, d as the maximum search distance, E is the number of

edges in the Dijkstra subgraph G , and V is the number of vertices in the Dijkstra subgraph G .

The number of vertices searched for connectivity (Step 2B) depends on technical properties intrinsic to the dataset—such as the number of reads, the read coverage, error rate, number of repeats in the metagenome, and number of chromosomes—and the user-set search distance parameter.

5.2.3 Step 3: Output maximal sets of connected reads, or enhanced read clouds.

Each enhanced read cloud is the set of reads from the same read cloud that are part of the same search-distance-limited connected subgraph, as identified in step 2B. These enhanced read clouds are output as solutions to the UMI deconvolution problem (Figure 3F). Each original read cloud is either subdivided into two or more enhanced read clouds, or kept intact if all of the reads were found within the search distance d of each other. To avoid a preponderance of trivial-sized enhanced read clouds, if there is an enhanced read cloud generated that is composed solely of a pair of reads, the two reads are instead added to the smallest enhanced read cloud generated from the same original cloud.

5.3 Selection of Search Distances

The maximum search distance was selected two quantities: i) the expected size of the physical genomic fragment that generates the sequencing reads and the ii) the likelihood of observing reads some N base-pairs (bp) away from a focal read, where N is an integer. [15] had previously estimated several dataset parameters to accurately model metagenomic fragments as paths through the assembly graph. To estimate the average fragment length, a method termed single linkage clustering was used to partition reads into clusters corresponding to alignments of likely fragments to the known reference genomes of the species comprising the sample. For example, the expected fragment size of the MOCK5 10x dataset was estimated to be 39,139 bp. For a complete analysis of the genomic fragment length estimates, the number of genomic fragments per 3' UMI, and overall coverage of MOCK5 10x, see the Supplementary Materials of [15].

Setting a maximum search distance of $d = 5$ kbp seemed to model linked-read genomic fragments with an expected length of 40 kbp reasonably. To limit the occurrence of under-deconvolution, the search distance is set conservatively, as a fraction of the expected genomic fragment length. Each read i is used as a focal read to search the assembly graph. If read j truly originated from the same genomic fragment of read i , it is significantly more likely for read j to occur within d

of read i than a full 40 kbp away. Several search distances smaller than the estimated fragment length—5, 10, 15, and 20 kbp—were tried for this study, and it was found that search distance $d = 5$ kbp provided the best balance of deconvolution accuracy, assembly quality, and computational efficiency (Supp. Table 1).

5.4 Mathematical Justification of Fragment Dissimilarity

We provide a short summary of the mathematical model that justifies the modelling of fragments as search-distance limited connected subgraphs within a de Bruijn assembly graph. The details of the mathematical model for drawing fragments of DNA from a SLR sequencing sample are fully described in [23], the article describing Minerva, the previous iteration of Ariadne. In brief, random fragments of genomic sequence on the scale of kilo-base-pairs from a large number of species with genomes that are many mega-base-pairs in length are unlikely to be similar in terms of sequence to overlap on the assembly graph, even with k -mers as small as 22 bp. For metagenomic datasets, at least 95% of read clouds consist of fragments that do not overlap in this model. This is important because it means that overlapping fragments, and thus spurious connections, are uncommon and will not hinder deconvolution in most read clouds. The algorithm is theoretically capable of uniquely deconvolving 95% of read clouds by searching the assembly graph surrounding the read for connecting reads. The only drawback of this model is that it does not account for the fact that individual fragments may have similar sequences, such as those resulting from repetitive regions or highly conserved genes.

Appendix

Acknowledgements

We would like to thank Indira Wu and Tuval Ben-Yehezkel at Loop Genomics for providing the MOCK5 LoopSeq dataset, as well as Yu Xia and Tom Chen at Universal Sequencing for providing the MOCK20 TELL-Seq dataset.

Funding

This work was supported by an NIGMS Maximizing Investigators' Research Award (MIRA)[grant number R35 GM138152-01 to I.H.]. L.M., D.M. and D.C.D. were also supported by the Tri-Institutional Training Program in Computational Biology and Medicine (CBM) funded by the NIH grant 1T32GM083937. L.M. was additionally supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Postgraduate Scholarship-Doctoral (PGS D). W.N.B. was supported by the Tri-Institutional Computational Biology Summer Program.

Abbreviations

Not applicable.

Availability of data and materials

5.5 Data
The MOCK5 and MOCK20 10x datasets originally used in [15] are available at cloudSPAdes.tar.gz. The MOCK5 LoopSeq and MOCK20 TELL-Seq datasets generated in this study have been submitted to the NCBI BioProject database under the accession number [PRJNA728470](https://www.ncbi.nlm.nih.gov/bioproject/PRJNA728470).

5.6 Code

Project name: Ariadne
Project home page: <https://github.com/lauren-mak/Ariadne>
Archived version: NA
Operating system(s): Linux
Programming language: C++, Python
Other requirements: g++ (version 5.3.1 or higher), cmake (version 2.8.12 or higher), zlib, libbz2
License: MIT License
Other: Pipelines and scripts used to automate *de novo* assembly, summary statistics, and figure generation are available at <https://github.com/lauren-mak/ariadne-assembly-support>.

Ethics approval and consent to participate

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Consent for publication

Not applicable.

Authors' contributions

L.M. implemented Ariadne, and conducted all benchmarking tests. W.N.B. help with an early prototype of Ariadne. N.B. assisted with analyses. D.M., D.C.D., and I.H. conceived of the initial idea, D.M. and D.C.D. provided direction and datasets to the study. All authors contributed to the writing of the manuscript. I.H. supervised the study.

Authors' information

Not applicable.

Additional Files

Additional file 1 — Supplementary Materials
Supplementary tables and figures with an extended comparison of Ariadne and Minerva on a subset of the full MOCK5 10x dataset demonstrating Ariadne's improvements in deconvolution.

Author details

¹Tri-Institutional Computational Biology & Medicine Program, Weill Cornell Medicine of Cornell University, New York, USA. ²Institute for Computational Biomedicine, Department of Physiology and Biophysics, Weill Cornell Medicine of Cornell University, New York, USA. ³Department of Computer Science, New York University, New York, USA. ⁴Physiology, Biophysics & Systems Biology Program, Weill Cornell Medicine of Cornell University, New York, USA. ⁵Institute of Biology, Weill Cornell Medicine of Cornell University, New York, USA.

References

1. Moss, E.L., Maghini, D.G., Bhatt, A.S.: Complete, closed bacterial genomes from microbiomes using nanopore sequencing. *Nature biotechnology* **38**(6), 701–707 (2020)
2. De Coster, W., Weissensteiner, M.H., Sedlazeck, F.J.: Towards population-scale long-read sequencing. *Nature Reviews Genetics* **22**(9), 572–587 (2021)
3. Sedlazeck, F.J., Lee, H., Darby, C.A., Schatz, M.C.: Piercing the dark matter: Bioinformatics of long-range sequencing and mapping. *Nature Reviews Genetics* **19**(6), 329–346 (2018). doi:[10.1038/s41576-018-0003-4](https://doi.org/10.1038/s41576-018-0003-4)
4. Bertrand, D., Shaw, J., Kalathiyappan, M., Ng, A.H.Q., Kumar, M.S., Li, C., Dvornic, M., Soldo, J.P., Koh, J.Y., Tong, C., *et al.*: Hybrid metagenomic assembly enables high-resolution analysis of resistance determinants and mobile elements in human microbiomes. *Nature biotechnology* **37**(8), 937–944 (2019)
5. Zlitni, S., Bishara, A., Moss, E.L., Tkachenko, E., Kang, J.B., Culver, R.N., Andermann, T.M., Weng, Z., Wood, C., Handy, C., *et al.*: Strain-resolved microbiome sequencing reveals mobile elements that drive bacterial competition on a clinical timescale. *Genome medicine* **12**(1), 1–17 (2020)
6. Brown, C.L., Keenum, I.M., Dai, D., Zhang, L., Vikesland, P.J., Pruden, A.: Critical evaluation of short, long, and hybrid assembly for contextual analysis of antibiotic resistance genes in complex environmental metagenomes. *Scientific Reports* **11**(1) (2021). doi:[10.1038/s41598-021-83081-8](https://doi.org/10.1038/s41598-021-83081-8)

7. Voskoboynik, A., Neff, N.F., Sahoo, D., Newman, A.M., Pushkarev, D., Koh, W., Passarelli, B., Fan, H.C., Mantalas, G.L., Palmeri, K.J., *et al.*: The genome sequence of the colonial chordate, botryllus schlosseri. *life* **2**, 00569 (2013)
8. Zheng, G.X.Y., Lau, B.T., Schnall-Levin, M., Jarosz, M., Bell, J.M., Hindson, C.M., Kyriazopoulou-Panagiotopoulou, S., Masquelier, D.A., Merrill, L., Terry, J.M., Mudivarti, P.A., Wyatt, P.W., Bharadwaj, R., Makarewicz, A.J., Li, Y., Belgrader, P., Price, A.D., Lowe, A.J., Marks, P., Vurens, G.M., Hardenbol, P., Montesclaros, L., Luo, M., Greenfield, L., Wong, A., Birch, D.E., Short, S.W., Bjornson, K.P., Patel, P., Hopmans, E.S., Wood, C., Kaur, S., Lockwood, G.K., Stafford, D., Delaney, J.P., Wu, I., Ordonez, H.S., Grimes, S.M., Greer, S., Lee, J.Y., Belhocine, K., Giorda, K.M., Heaton, W.H., McDermott, G.P., Bent, Z.W., Meschi, F., Kondov, N.O., Wilson, R., Bernate, J.A., Gauby, S., Kindwall, A., Bermejo, C., Fehr, A.N., Chan, A., Saxonov, S., Ness, K.D., Hindson, B.J., Ji, H.P.: Haplotyping germline and cancer genomes with high-throughput linked-read sequencing. *Nature Biotechnology* **34**(3), 303–311 (2016). doi:[10.1038/nbt.3432](https://doi.org/10.1038/nbt.3432)
9. Chen, Z., Pham, L., Wu, T.-C., Mo, G., Xia, Y., Chang, P.L., Porter, D., Phan, T., Che, H., Tran, H., *et al.*: Ultralow-input single-tube linked-read library method enables short-read second-generation sequencing systems to routinely generate highly accurate and economical long-range sequencing information. *Genome research* **30**(6), 898–909 (2020)
10. Roodgar, M., Good, B.H., Garud, N.R., Martis, S., Avula, M., Zhou, W., Lancaster, S.M., Lee, H., Babveyh, A., Nesamoney, S., *et al.*: Longitudinal linked-read sequencing reveals ecological and evolutionary responses of a human gut microbiome during antibiotic treatment. *Genome research* **31**(8), 1433–1446 (2021)
11. Callahan, B.J., Grinevich, D., Thakur, S., Balamotis, M.A., Yehezkel, T.B.: Ultra-accurate microbial amplicon sequencing with synthetic long reads. *Microbiome* **9**(1), 1–13 (2021)
12. Wang, O., Chin, R., Cheng, X., Yan Wu, M.K., Mao, Q., Tang, J., Sun, Y., Anderson, E., Lam, H.K., Chen, D., Zhou, Y., Wang, L., Fan, F., Zou, Y., Xie, Y., Zhang, R.Y., Drmanac, S., Nguyen, D., Xu, C., Villarosa, C., Gablenz, S., Barua, N., Nguyen, S., Tian, W., Liu, J.S., Wang, J., Liu, X., Qi, X., Chen, A., Wang, H., Dong, Y., Zhang, W., Alexeev, A., Yang, H., Wang, J., Kristiansen, K., Xu, X., Drmanac, R., Peters, B.A.: Efficient and unique cobarcoding of second-generation sequencing reads from long DNA molecules enabling cost-effective and accurate sequencing, haplotyping, and de novo assembly. *Genome Research* **29**(5), 798–808 (2019). doi:[10.1101/gr.245126.118](https://doi.org/10.1101/gr.245126.118)
13. Weisenfeld, N.I., Kumar, V., Shah, P., Church, D.M., Jaffe, D.B.: Direct determination of diploid genome sequences. *Genome Research* **27**(5), 757–767 (2017). doi:[10.1101/gr.214874.116](https://doi.org/10.1101/gr.214874.116)
14. Bishara, A., Moss, E.L., Kolmogorov, M., Parada, A.E., Weng, Z., Sidow, A., Dekas, A.E., Batzoglou, S., Bhatt, A.S.: High-quality genome sequences of uncultured microbes by assembly of read clouds. *Nature Biotechnology* **36**(11), 1067–1080 (2018). doi:[10.1038/nbt.4266](https://doi.org/10.1038/nbt.4266)
15. Tolstoganov, I., Bankevich, A., Chen, Z., Pevzner, P.A.: CloudSPAdes: Assembly of synthetic long reads using de Bruijn graphs. *Bioinformatics* **35**(14), 61–70 (2019). doi:[10.1093/bioinformatics/btz349](https://doi.org/10.1093/bioinformatics/btz349)
16. Hug, L.A., Baker, B.J., Anantharaman, K., Brown, C.T., Probst, A.J., Castelle, C.J., Butterfield, C.N., HERNSDORF, A.W., Amano, Y., Ise, K., Suzuki, Y., Dudek, N., Relman, D.A., Finstad, K.M., Amundson, R., Thomas, B.C., Banfield, J.F.: A new view of the tree of life. *Nature Microbiology* **1**(5), 1–6 (2016). doi:[10.1038/nmicrobiol.2016.48](https://doi.org/10.1038/nmicrobiol.2016.48)
17. Ayling, M., Clark, M.D., Leggett, R.M.: New approaches for metagenome assembly with short reads. *Briefings in Bioinformatics* **00**(January), 1–11 (2019). doi:[10.1093/bib/bbz020](https://doi.org/10.1093/bib/bbz020)
18. Chung, N., Van Goethem, M.W., Preston, M.A., Lhota, F., Cerna, L., Garcia-Pichel, F., Fernandes, V., Giraldo-Silva, A., Kim, H.S., Hurowitz, E., *et al.*: Accurate microbiome sequencing with synthetic long read sequencing. *bioRxiv* (2020)
19. Jeong, J., Yun, K., Mun, S., Chung, W.-H., Choi, S.-Y., Nam, Y.-d., Lim, M.Y., Hong, C.P., Park, C., Ahn, Y.J., *et al.*: The effect of taxonomic classification by full-length 16s rRNA sequencing with a synthetic long-read technology. *Scientific reports* **11**(1), 1–12 (2021)
20. Gupta, C.L., Avidov, R., Kattusamy, K., Saadi, I., Varma, V.S., Blum, S.E., Zhu, Y.-G., Zhou, X.-Y., Su, J.-Q., Laor, Y., *et al.*: Spatial and temporal dynamics of microbiomes and resistomes in broiler litter stockpiles. *Computational and structural biotechnology journal* **19**, 6201–6211 (2021)
21. Yilmaz, B., Mooser, C., Keller, I., Li, H., Zimmermann, J., Bosshard, L., Fuhrer, T., de Agüero, M.G., Trigo, N.F., Tschanz-Lischer, H., *et al.*: Long-term evolution and short-term adaptation of microbiota strains and sub-strains in mice. *Cell host & microbe* **29**(4), 650–663 (2021)
22. Marks, P., Garcia, S., Barrio, A.M., Belhocine, K., Bernate, J., Bharadwaj, R., Bjornson, K., Catalanotti, C., Delaney, J., Fehr, A., Fiddes, I.T., Galvin, B., Heaton, H., Herschleb, J., Hindson, C., Holt, E., Jabara, C.B., Jett, S., Keivanfar, N., Kyriazopoulou-Panagiotopoulou, S., Lek, M., Lin, B., Lowe, A., Mahamdallie, S., Maheshwari, S., Makarewicz, T., Marshall, J., Meschi, F., O’Keefe, C.J., Ordonez, H., Patel, P., Price, A., Royall, A., Ruark, E., Seal, S., Schnall-Levin, M., Shah, P., Stafford, D., Williams, S., Wu, I., Xu, A.W., Rahman, N., MacArthur, D., Church, D.M.: Resolving the full spectrum of human genome variation using Linked-Reads. *Genome Research* **29**(4), 635–645 (2019). doi:[10.1101/gr.234443.118](https://doi.org/10.1101/gr.234443.118)
23. Danko, D.C., Meleshko, D., Bezdán, D., Mason, C., Hajirasouliha, I.: Minerva: an alignment- and reference-free approach to deconvolve Linked-Reads for metagenomics. *Genome research* **29**(1), 116–124 (2019). doi:[10.1101/gr.235499.118](https://doi.org/10.1101/gr.235499.118)
24. Shajji, A., Numanagić, I., Whelan, C., Berger, B.: Statistical Binning for Barcoded Reads Improves Downstream Analyses. *Cell Systems* **7**(2), 219–2265 (2018). doi:[10.1016/j.cels.2018.07.005](https://doi.org/10.1016/j.cels.2018.07.005)
25. Dufresne, Y., Sun, C., Marijon, P., Lavenier, D., Chauve, C., Chikhi, R.: A Graph-Theoretic Barcode Ordering Model for Linked-Reads. *Drops-Iln/12800* (11), 1–11 (2020). doi:[10.4230/LIPICS.WABI.2020.11](https://doi.org/10.4230/LIPICS.WABI.2020.11)
26. Danko, D., Bezdán, D., Afshinnekoo, E., Ahsanuddin, S., Bhattacharya, C., Butler, D.J., Chng, K.R., De Filippis, F., Hecht, J., Kahles, A., Karasikov, M., Kyrpides, N.C., Leung, M.H.Y., Meleshko, D., Mustafa, H., Mutai, B., Neches, R.Y., Ng, A., Nieto-Caballero, M., Nikolayeva, O., Nikolayeva, T., Png, E., Sanchez, J.L., Shaaban, H., Sierra, M.A., Tong, X., Young, B., Alicea, J., Bhattacharyya, M., Blekhan, R., Castro-Nallar, E., Cañas, A.M., Chatziefthimiou, A.D., Crawford, R.W., Deng, Y., Desnues, C., Dias-Neto, E., Donnellan, D., Dybdal, M., Elhaik, E., Ercolini, D., Frolova, A., Graf, A.B., Green, D.C., Hajirasouliha, I., Hernandez, M., Iraola, G., Jang, S., Jones, A., Kelly, F.J., Knights, K., Łabaj, P.P., Lee, P.K.H., Shawn, L., Ljungdahl, P., Lyons, A., Mason-Buck, G., McGrath, K., Mongodin, E.F., Moraes, M.O., Nagarajan, N., Noushmehr, H., Oliveira, M., Ossowski, S., Osuolale, O.O., Özcan, O., Paez-Espino, D., Rascovan, N., Richard, H., Rättsch, G., Schriml, L.M., Semmler, T., Sezerman, O.U., Shi, L., Song, L.H., Suzuki, H., Court, D.S., Thomas, D., Tighe, S.W., Udekwi, K.I., Ugalde, J.A., Valentine, B., Vassilev, D.I., Vayndorf, E., Velavan, T.P., Zambrano, M.M., Zhu, J., Zhu, S., Mason, C.E.: Global Genetic Cartography of Urban Metagenomes and Anti-Microbial Resistance. *bioRxiv*, 724526 (2019). doi:[10.1101/724526](https://doi.org/10.1101/724526)
27. Langmead, B., Salzberg, S.L.: Fast gapped-read alignment with Bowtie 2. *Nature Methods* **9**(4), 357–359 (2012). doi:[10.1038/nmeth.1923](https://doi.org/10.1038/nmeth.1923)
28. Mikheenko, A., Saveliev, V., Gurevich, A.: Metaquast: evaluation of metagenome assemblies. *Bioinformatics* **32**(7), 1088–1090 (2016)
29. Guo, L., Xu, M., Wang, W., Gu, S., Zhao, X., Chen, F., Wang, O., Xu, X., Fan, G., Deng, L., Liu, X.: SLR-superscaffolder: A de novo scaffolding tool for synthetic long reads using a top-to-bottom scheme. *bioRxiv* (2019). doi:[10.1101/762385](https://doi.org/10.1101/762385)
30. Weng, J., Chen, T., Xie, Y., Xu, X., Zhang, G., Peters, B.A., Drmanac, R.: IterCluster: a barcode clustering algorithm for long fragment read analysis. *PeerJ* **8**, 8431 (2020). doi:[10.7717/peerj.8431](https://doi.org/10.7717/peerj.8431)
31. Ma, Z.S., Li, L., Ye, C., Peng, M., Zhang, Y.-P.: Hybrid assembly of ultra-long nanopore reads augmented with 10x-genomics contigs: Demonstrated with a human genome. *Genomics* **111**(6), 1896–1901 (2019)
32. Coombe, L., Zhang, J., Vandervalk, B.P., Chu, J., Jackman, S.D., Birol, I., Warren, R.L.: Arks: chromosome-scale scaffolding of human genome drafts with linked read kmers. *BMC bioinformatics* **19**(1), 1–10 (2018)
33. Zhang, L., Fang, X., Liao, H., Zhang, Z., Zhou, X., Han, L., Chen, Y.,

- Qiu, Q., Li, S.C.: A comprehensive investigation of metagenome assembly by linked-read sequencing. *Microbiome* **8**(1), 1–11 (2020)
34. Sangwan, N., Xia, F., Gilbert, J.A.: Recovering complete and draft population genomes from metagenome datasets. *Microbiome* **4**(1) (2016). doi:[10.1186/s40168-016-0154-5](https://doi.org/10.1186/s40168-016-0154-5)
35. Li, H., Durbin, R.: Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* **25**(14), 1754–1760 (2009). doi:[10.1093/bioinformatics/btp324](https://doi.org/10.1093/bioinformatics/btp324)
36. Nurk, S., Meleshko, D., Korobeynikov, A., Pevzner, P.A.: MetaSPAdes: A new versatile metagenomic assembler. *Genome Research* **27**(5), 824–834 (2017). doi:[10.1101/gr.213959.116](https://doi.org/10.1101/gr.213959.116). [1604.03071](https://pubmed.ncbi.nlm.nih.gov/28011116/)
37. Bankevich, A., Nurk, S., Antipov, D., Gurevich, A.A., Dvorkin, M., Kulikov, A.S., Lesin, V.M., Nikolenko, S.I., Pham, S., Prjibelski, A.D., Pyshkin, A.V., Sirotkin, A.V., Vyahhi, N., Tesler, G., Alekseyev, M.A., Pevzner, P.A.: SPAdes: A new genome assembly algorithm and its applications to single-cell sequencing. *Journal of Computational Biology* **19**(5), 455–477 (2012). doi:[10.1089/cmb.2012.0021](https://doi.org/10.1089/cmb.2012.0021)
38. Woodcroft, B.: CoverM. <https://github.com/wood/CoverM> (2022)

Figures

Figure 1 The size-weighted purity of SLR read clouds increases after applying deconvolution methods. All graphs were generated from 40,000 randomly sampled clouds from the dataset. **Top row:** No deconvolution. **Middle row:** Reference deconvolution based on read alignment to species and then grouping reads in 200 kbp regions. **Bottom row:** Ariadne deconvolution with a search distance of 5 kbp and a minimum cloud size cutoff of 6.

Figure 2 Read cloud deconvolution improves metagenomic assembly compared to raw SLR data. We compared assemblies built from raw linked reads (no deconvolution) to assemblies built from reads deconvolved using two methods: reference deconvolution, which maps reads to reference genomes, and deconvolution using Ariadne. **Top Row:** The NA50 of assemblies for each species in each sample between deconvolved reads and raw reads. Larger numbers indicate better performance. **Middle Row:** The largest alignment of assemblies for each species in each sample between deconvolved reads and raw reads. Larger numbers indicate better performance. **Bottom row:** The proportion of misassembled bases p_{miss} is the number of bases in misassembled contigs over the total number of assembled bases.

Figure 3 Graphical description of the Ariadne deconvolution process. **(A)** Reads with the same 3' UMI are in a read cloud. Blue and red reads originate from different fragments. The **(B)** de Bruijn assembly graph is generated by cloudSPAdes, and **(C)** a focal read is mapped to one of its edges. From a read's 3'-terminal vertex, **(D)** a Dijkstra graph (indicated by a large black circle) is created from all edges and vertices within the maximum search distance from the 3'-terminal vertex. These vertices and edges (within the black circle) comprise read i 's search-distance-limited connected subgraph within the whole assembly graph. Reads aligning to edges in this connected subgraph are added to read i 's connected set. **(E)** Reads originating from different fragments likely coincide with non-included vertices. **(F)** Connected read-sets with at least one intersection (i.e.: one read in common) are output together as an enhanced read cloud.

Tables

Table 1 Overview of mock microbiome linked-read datasets. * indicates that MOCK20 10x and TELL-Seq generated from the same mock microbiome community product

Dataset	Num. Reads	Num. Barcoded Reads	Prop. Barcoded Reads
MOCK5 10x	97,491,080	91,101,472	0.93446
MOCK5 LoopSeq	75,107,814	75,107,814	1
MOCK20 10x*	100,000,000	94,151,528	0.94152
MOCK20 TELLSeq*	100,000,000	100,000,000	1

Table 2 Relative species abundance in each mock microbiome community as calculated by the tool CoverM [38]. Abundance based on read coverage of reference genome sequence adjusted for genome size.

Dataset	Species	Relative Abundance
MOCK5 10x	<i>Enterobacter cloacae</i>	7.129894635
	<i>Pseudomonas fluorescens</i>	11.63274034
	<i>Micrococcus luteus</i>	16.54618068
	<i>Escherichia coli</i>	22.04082093
	<i>Staphylococcus epidermidis</i>	42.6503661
MOCK5 LoopSeq	<i>Pseudomonas aeruginosa</i>	2.2308808
	<i>Streptococcus mutans</i>	5.312954304
	<i>Porphyromonas gingivalis</i>	26.40226344
	<i>Escherichia coli</i>	28.9818004
	<i>Rhodobacter sphaeroides</i>	37.07210106
MOCK20 10x and TELLSeq	<i>Schaalia odontolytica</i>	0.005069583
	<i>Bifidobacterium adolescentis</i>	0.005194844
	<i>Deinococcus radiodurans</i>	0.014251377
	<i>Enterococcus faecalis</i>	0.015894382
	<i>Bacteroides vulgatus</i>	0.027645357
	<i>Cutibacterium acnes</i>	0.184333912
	<i>Lactobacillus gasseri</i>	0.189651678
	<i>Neisseria meningitidis</i>	0.229268868
	<i>Helicobacter pylori</i>	0.241180209
	<i>Acinetobacter baumannii</i>	0.250956306
	<i>Bacillus cereus</i>	1.214779563
	<i>Clostridium beijerinckii</i>	1.220119982
	<i>Pseudomonas aeruginosa</i>	1.366135524
	<i>Staphylococcus aureus</i>	1.816064537
	<i>Streptococcus agalactiae</i>	1.914207527
	<i>Rhodobacter sphaeroides</i>	13.2715074
	<i>Escherichia coli</i>	16.14207583
	<i>Staphylococcus epidermidis</i>	19.55632719
	<i>Porphyromonas gingivalis</i>	21.07923273
	<i>Streptococcus mutans</i>	21.25610476

Table 3 Read cloud summary statistics. For Ariadne deconvolution, we used a search distance of 5 kbp. The 3-5th columns contain the average and standard deviations of read cloud statistics. Prop. Under, Comp., and Over refer to the proportion of total original or deconvolved read clouds that were over- or under-deconvolved, or completely and exactly comprised of all of the reads from a single inferred genomic fragment.

Dataset	Deconv. Method	Avg. Purity	Avg. Entropy	Avg. Size	Num. Clouds	Prop. Under	Prop. Comp.	Prop. Over
MOCK5 10x	None	0.53±0.16	1.61±0.52	63.51 ±51.64	1 425 430	0.98	0.02	0
MOCK5 10x	Reference	0.99±0.05	0.03±0.14	18.47 ±19.26	4 733 136	0.06	0.91	0.03
MOCK5 10x	Ariadne	0.89±0.22	0.37±0.7	12.18 ±15.26	4 753 473	0.25	0.04	0.71
MOCK5 LoopSeq	None	0.93±0.11	0.31±0.31	269.18 ±356.67	276 665	0.68	0.32	0
MOCK5 LoopSeq	Reference	0.98±0.07	0.06±0.21	116.19 ±252.59	637 012	0.14	0.82	0.04
MOCK5 LoopSeq	Ariadne	0.88±0.21	0.4±0.7	101.44 ±225.71	734 143	0.35	0.14	0.51
MOCK20 10x	None	0.38±0.15	2.28±0.62	186.49 ±129.32	503 205	0.97	0.03	0
MOCK20 10x	Reference	0.99±0.06	0.05±0.18	29.31 ±29.26	3 146 784	0.13	0.83	0.04
MOCK20 10x	Ariadne	0.94±0.18	0.22±0.64	13.56 ±22.69	6 920 196	0.13	0.05	0.83
MOCK20 TELL-Seq	None	0.39±0.15	2.37±0.59	160.45 ±110.51	623 026	0.98	0.02	0
MOCK20 TELL-Seq	Reference	1±0.02	0±0.05	23.83 ±27.65	4 067 621	0.01	0.99	0
MOCK20 TELL-Seq	Ariadne	0.87±0.24	0.47±0.89	21.96 ±30.74	4 074 321	0.25	0.04	0.71





Table 4 Both Ariadne and reference deconvolution increase the summed NA50 of *de novo* assembled metagenomes. The rightmost column shows the difference between the summed NA50 of assemblies obtained from deconvolved reads and the summed NA50 of non-deconvolved assembly. The asterisk (*) indicates species that had no associated NA50 in the non-deconvolved assembly.

Dataset	Species	Deconv. Method	Difference in NA50 (bp)
MOCK5 10x	Enterobacter cloacae*	Ariadne	16 655
MOCK20 TELL-Seq	Streptococcus agalactiae*	Ariadne	464 957
MOCK20 TELL-Seq	Staphylococcus epidermidis*	Ariadne	489 177
MOCK20 TELL-Seq	Streptococcus mutans*	Ariadne	906 518
MOCK20 TELL-Seq	Staphylococcus aureus*	Ariadne	106 598
MOCK5 10x	Enterobacter cloacae*	Reference	13 628
MOCK20 10x	Rhodobacter sphaeroides	Reference	-580 531
MOCK20 10x	Staphylococcus epidermidis	Reference	639 974
MOCK20 TELL-Seq	Staphylococcus epidermidis*	Reference	1 503 747
MOCK20 TELL-Seq	Streptococcus agalactiae*	Reference	2 047 925
MOCK20 TELL-Seq	Streptococcus mutans*	Reference	2 012 488
MOCK20 TELL-Seq	Staphylococcus aureus*	Reference	134 663



Table 5 Comparison of 10x vs. TELL-Seq *de novo* assembly statistics as calculated by MetaQUAST. MOCK20 datasets were made using 10x and TELL-seq library preparation and sequencing protocols, reference-deconvolved, and *de novo* assembled using cloudSPAdes.

Summary statistic	MOCK20 10x	MOCK20 TELL-Seq
Genome fraction (%)	75.827	67.719
Total aligned length (bp)	52,363,841	45,407,218
Longest single-contig alignments (bp)	6,234,231	2,079,899
Num. contigs	2481	7729
Num. misassemblies	127	59
Total size of misassembled contigs (bp)	23,155,091	15,651,764
Proportion of assembled sequence in misassembled contigs	44.22%	34.47%
Num. unaligned contigs	1	45
Size of unaligned contigs (bp)	599	105,526




No Deconvolutions






Reference



Arriane





PROVIDER

PROVIDER

PROVIDER

PROVIDER

PROVIDER

PROVIDER



PROVIDER

PROVIDER

PROVIDER

