

AN ORFOME ASSEMBLY APPROACH TO METAGENOMICS SEQUENCES ANALYSIS

Yuzhen Ye* and Haixu Tang

*School of Informatics, Indiana University
Bloomington, Indiana 47408, USA*

**Email: yye@indiana.edu*

Metagenomics is an emerging methodology for the direct genomic analysis of a mixed community of uncultured microorganisms. The current analyses of metagenomics data largely rely on the computational tools originally designed for microbial genomics projects. The challenge of assembling metagenomic sequences arises mainly from the short reads and the high species complexity of the community. Alternatively, individual (short) reads will be searched directly against databases of known genes (or proteins) to identify homologous sequences. The latter approach may have low sensitivity and specificity in identifying homologous sequences, which may further bias the subsequent diversity analysis. In this paper, we present a novel approach to metagenomic data analysis, called Metagenomic ORFome Assembly (MetaORFA). The whole computational framework consists of three steps. Each read from a metagenomics project will first be annotated with putative open reading frames (ORFs) that likely encode proteins. Next, the predicted ORFs are assembled into a collection of peptides using an EULER assembly method. Finally, the assembled peptides (i.e., ORFome) are used for database searching of homologs and subsequent diversity analysis. We applied MetaORFA approach to several metagenomics datasets with low coverage short reads. The results show that MetaORFA can produce long peptides even when the sequence coverage of reads is extremely low. Hence, the ORFome assembly significantly increased the sensitivity of homology searching, and may potentially improve the diversity analysis of the metagenomic data. This improvement is especially useful for the metagenomic projects when the genome assembly does not work because of the low sequence coverage.

Keywords: Metagenomics, ORFome, ORFome assembly, Function annotation.

1. INTRODUCTION

Owing to the rapid advancement of the ultra-high throughput DNA sequencing technologies¹, the genomic studies of microorganisms in environmental samples have recently shifted from the focused sequencing of 16sRNA sequences² to the shotgun sequencing of the whole DNAs in the sample. This new methodology, now called *metagenomics* or *environmental genomics*, has opened a door for biologists to assess the unknown world of the uncultured microorganisms that are believed to be the majority in any environmental sample. The early attempts of this kind can be traced back to a report published in 2002, in which extremely high diversity of uncultured marine viral communities were revealed through genome sequencing³. However, the most important progress in shotgun metagenomics happened in 2004⁴⁻⁷, when two research groups published results from their large-scale environmental sequencing projects. The first project studied the sample from the Sargasso Sea, and revealed ~ 2000 dis-

tinct species of microorganisms, including 148 types of bacteria that have never been observed before⁸. In the second project, a handful of genomes of bacteria and archaea that had previously resisted attempts to culture them were revealed based on the analysis of the sample from the acid mine drainage⁹. Since then, many more metagenomics projects have been conducted, involving broadened applications from ecology and environmental sciences to chemical industry¹⁰ and human health, e.g., the human gut microbiome projects^{11, 12}.

The rapid growth of metagenomic data has posed great challenges to the computational analysis^{13, 14}. Some metagenomics projects applied directly the data analysis pipeline that includes the whole genome assemblers¹⁵⁻¹⁸ and gene finding programs¹⁹—originally designed for the conventional Whole Genome Shotgun (WGS) sequencing projects—with only some small parameter modifications^{8, 9, 12, 20}. However, it is unclear how accurate these existing tools for fragment assembly and genome annotation are when applied to metagenomic

*Corresponding author.

data. Mavromatis and colleagues have conducted a valuable benchmarking experiment to evaluate the performance of conventional genome assembly and annotation pipeline on simulated metagenomic data²¹. In this experiment, sequencing reads were randomly collected from 113 assembled genomes that are mixed at various complexities. Afterwards, the quality of the results from each processing step (i.e., assembly, gene prediction, and phylogenetic binning) was assessed separately by comparison to the corresponding genomes used in the simulation. This experiment delivered an encouraging message that the number of errors made at each step overall is not high, and some errors (e.g., the chimeric contigs) would not be propagated into the subsequent steps (e.g., binning). Nevertheless, we argue that this experiment may not completely reflect the challenge of metagenomic data analysis, especially the difference between metagenomic data and the data from conventional genome sequencing. Conventional genome projects deal with only one or sometimes a few individual genomes from the same species that are isolated prior to sequencing, whereas metagenomics attempts to analyze simultaneously a huge amount of genomes not only from hundreds of different microorganisms, but also from many individuals of each organism. As a result, even the reads from the same species might be quite different from each other since they might be sampled from different individuals' genomes. Furthermore, those microbial species may exist in the sample at a wide range of abundances. Hence, typically, only a few dominant species can receive good sequence coverage for their genomes, whereas the sequence coverage for the remaining species is low.

More and more metagenomic projects have applied Next-Generation Sequencing (NGS) technologies that produce massive but shorter reads (e.g., ~ 200 bps for 454 pyrosequencing machines) than those from the Sanger sequencing methods. Therefore, many metagenomic sequencing projects that acquired a merely small number of short sequencing reads often skipped the step of fragment assembly, and directly used the short reads for downstream analysis^{3, 22, 23}. For instance, short reads

can be used to search against protein database using TBLASTX to identify homologous proteins, in which an arbitrary E-value (e.g., $\leq 1e-5$) was chosen as a cutoff²². This direct search approach, however, often missed many homologous genes (or proteins)²⁴, and resulted in a very low false positive rate^a but high false negative rate. This drawback may bias the further analysis of species diversity (i.e., how many different species are present in the sample) and functional coverage (i.e., how many functional categories of proteins are present in the sample).

In this paper, we present a novel *ORFome assembly* approach to assembling metagenomic sequencing reads. Different from the conventional genome analysis pipeline that first assembles sequencing reads into contigs (or scaffolds) and then predicts protein coding regions within the contigs, our method first identifies putative protein coding regions (i.e., open reading frames, or ORFs) within unassembled reads, and then focuses on the assembly of only these sequences (i.e., *ORFome*). The *ORFome assembly* approach has several advantages. First, it significantly simplifies the task of fragment assembly that is often complicated by the repetitive sequences present mainly in non-coding regions²⁵. Meanwhile, we argue that *ORFome assembly* does not lose much useful information by neglecting the non-coding sequences due to several reasons: (1) the set of proteins (or the *ORFome* that encodes them) carry the most important information for the downstream analysis; (2) the microbial genomes are often very compact and protein coding regions comprise a major fraction of them; and (3) microbial proteins are mainly encoded by continuous non-split open reading frames (ORFs), thus the prediction of coding sequences prior to assembly is relatively straightforward. Second, from *ORFome assembly*, complete proteins (or long peptides) may be derived, thus higher sensitivity and specificity can be achieved in the step of database searching for homologs²⁴. Furthermore, most single nucleotide polymorphisms are synonymous mutations that do not change the encoding amino acids so that *ORFome assembly* does not even feel them. So by working on the peptide sequences (translated from sequencing reads *in silico*)

^aFor example, the MEGAN analysis based on the direct BLAST search method has achieved a 0 false positive rate²³!

instead of the raw DNA sequences, the ORFome assembly alleviates the assembly difficulty caused by the differences among individual genomes at polymorphic sites. We used four marine viral metagenomic datasets of short reads, acquired using 454 sequencing technique, to test our ORFome assembly method—no genome assemblies are available for these metagenomic datasets because the reads are extremely short and the sequence coverage is low.

2. METHODS

The computational framework of ORFome assembly consists of three steps (Fig. 1 (e-f)): (1) each read is assessed individually and the putative open reading frames (ORFs) that likely encode proteins are annotated; (2) the annotated ORFs are assembled into a collection of peptides using a modified EULER assembly method²⁶; and (3) the assembled peptides are used for the database searching of homologs.

A major difference between the ORFome assembly approach and the conventional whole genome assembly is that the former approach conducts gene annotation before assembly, whereas the latter approach conducts gene annotation after assembly. Conventional fragment assembly algorithms are mostly based on the analysis of overlap graph, in which the reads are represented by vertices and the overlaps between reads are represented by edges²⁷. The presence of repeats in the genomes often induce many spurious edges in the overlap graph, which is a major challenge in fragment assembly. There are two additional aspects in the metagenomic data that make fragment assembly even more challenging. First, metagenomics projects often apply NGS technique, and produce shorter reads (~ 200 bps) than Sanger sequencing methods (500-1000 bps). As a result, many short repeats (with lengths between 200 bps and 500 bps) may increase the complexity of the overlap graph, and cause many more mis-assemblies²⁸. Second, unlike the conventional genome shotgun sequencing, which handles a single species, metagenomics sequencing reads are collected from a large amount of different genomes. Hence, we anticipate these reads should be assembled into not one but many sequences that may even share high similarity on multiple regions. Therefore, the straightforward application of conventional fragment

assemblers may encounter difficulties. In contrast, the ORFome assembly approach attempts to assemble only the most important portions of the target genomes, i.e., the protein coding regions, which can highly reduce the complexity of the overlap graph and thus improve the assembly quality.

It is worth pointing out the idea of ORFome assembly can be viewed as an extension of the repeat masking approach used in whole genome assembly of large eukaryotic (including human) genomes. To avoid the complication induced by the many interspersed repeat copies present in most eukaryotic genomes, Celera Assembler first masked out putative repeats in the unassembled reads, and then focused on the assembly of the remaining reads from non-repetitive regions^{29, 30}. The resulting overlap graph, which consists of a number of connected components each representing reads from continuous non-repetitive regions, is much simpler and easy to be analyzed. Similarly, the ORFome assembly approach divides the complex overlap graph into a number of components each representing reads from a single gene or several highly similar genes from the same family.

We applied the ORFome assembly approach to several metagenomics datasets from Ocean samples with low coverage and short reads²². The results show that MetaORFA can produce long peptides even when the sequence coverage of reads is extremely low. Hence, further analysis of assembled peptides significantly increased the sensitivity for subsequent homology searching, and may potentially improve the diversity analysis of the metagenomic data.

2.1. ORFome Assembly Algorithm

We implemented a tool called MetaORFA in C/C++ under linux platforms for the ORFome assembly. MetaORFA consists of two programs. One program takes as input a set of reads and predicts a number of putative ORFs; and the other program takes as input the set of putative ORFs, and reports a set of peptides corresponding to the assembled ORFs. Prior to be supplied to MetaORFA, the original reads were first processed by MDUST (a popular tool for autonomous masking from TIGR, which implements the DUST algorithm³¹) to mask out low-

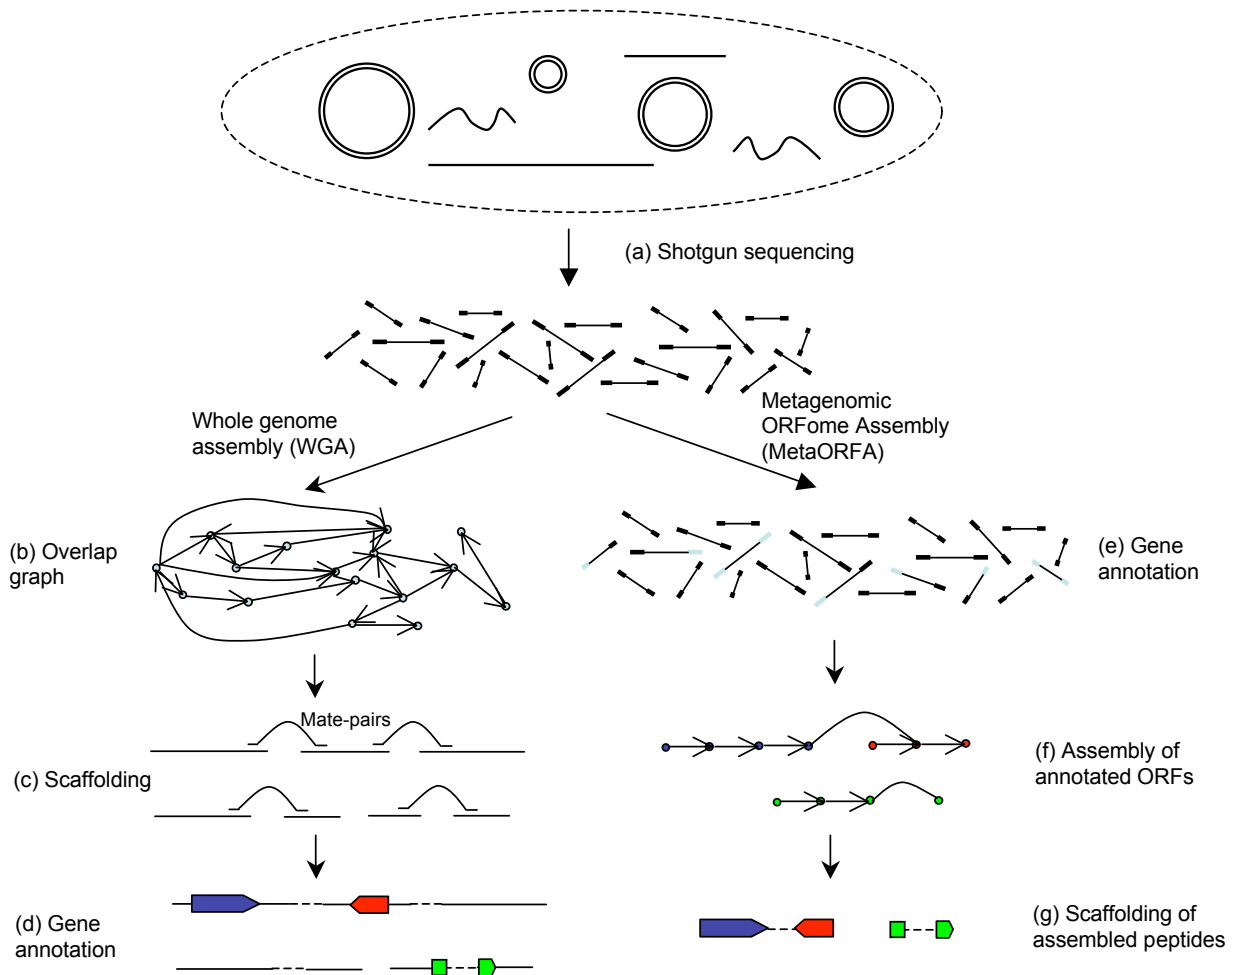


Fig. 1. A schematic comparison of the ORFome assembly approach with the whole genome assembly (WGA) pipeline for the metagenomic sequence analysis. Both approaches attempt to characterize the protein coding genes in the shotgun sequencing reads from the metagenomic analysis of an environmental sample containing a number of different microorganisms (the reads are shown as double-barreled, as currently several NGS techniques are capable of generating such data; however, some early metagenomics projects, including the datasets used in this paper, did not produce double-barreled sequencing reads, and thus the scaffolding step is not feasible) (a). The whole genome assembly (WGS) pipeline (b-d) first assembles the reads into contigs and scaffolds, and then annotates the genes in the assembled sequences. In comparison, ORFome assembly approach (e-g) first applies gene finding in the unassembled reads, and then assembles only those annotated (partial) ORFs into peptides. These peptides may be further connected to form scaffolds if there are mate-pairs available from double-barreled sequencing (g).

complexity regions, and then processed by Tandem Repeat Finder (TRF V4.0)³² to mask out short tandem repeats.

In this preliminary study, we adopted a very simple method for ORF prediction. For each read (and its reverse complement), a region from the beginning (i.e., position 1, 2, or 3, depending on the frame) or a start codon to the end of the read or a stop codon is considered as a potential ORF. Only ORFs with

more than a threshold K (default $K = 30$) codons were reported. These ORFs will be then transformed into peptide sequences, and subsequently assembled using a modified EULER algorithm originally designed for DNA fragment assembly²⁶. In this process, we first build a de Bruijn graph using all k -mers (default $k = 10$) in the putative peptides from previous step, and then apply the equivalent transformations as described in Ref. 26 to resolve short

repeats among peptides. Unlike many other genome assemblers that assemble reads into linear contigs, EULER aims at constructing from the reads a *repeat graph* that represents not only the unique regions but also the repeat structures³³. Although we anticipate there are not many repeats in the coding sequences, the similar parts of homologous proteins from the same family may act like repeats during the ORFome assembly. EULER assembly method can generate a compact graph structure representing the architecture of domain combinations, including domain recurrences and shuffling³⁴. We note further analysis of the ORFome assembly results, as described below, has not fully taken advantage of the information embodied in the repeat graph. Ideally, one can adopt a network matching approach to identify a path in the repeat graph representing a peptide sequence that is most similar to a protein in databases of known proteins. Nevertheless, our analysis has demonstrated that even the simple analysis of individual assembled peptides (corresponding to the edges in the repeat graph) revealed more proteins in the sample.

2.2. Functional Coverage Assessment

The ORFome, i.e., the set of assembled peptides, is ready for further computational analysis with different purposes, e.g., searching against database for homologous sequences, or mapping to biological pathways to study metabolic diversity³⁵. Here we show that we can improve the functional coverage of metagenomics sequences by using assembled peptides instead of unassembled reads. There are various ways to estimate functional coverage of a sample. In this study we used PANTHER (Protein ANalysis THrough Evolutionary Relationships) protein family classification³⁶ for such assessment. The comparison of the functional coverage between different ORFomes is then straightforward, and we can simply count the number of families (subfamilies) found in each assembled ORFome and calculate their differences.

In the PANTHER classification system, proteins are classified into families and subfamilies of shared function by experts. Families and subfamilies are presented as Hidden Markov Models (HMMs). We downloaded the PANTHER HMM library Version 6.1 (release date December 17, 2007) from

<ftp://ftp.pantherdb.org>, which contains 5547 protein family HMMs, divided into 24,582 functionally distinct protein subfamily HMMs. We also downloaded the HMM searching tool (`pantherScore.pl`, version 1.02), which utilized fast BLAST search prior to the more sensitive but time-consuming HMM matching procedure to speed up the process. The query protein sequence will first be blasted against the consensus sequences of each PANTHER HMMs, and then based on the results, some heuristics are applied to determine which HMMs (i.e., protein families or subfamilies) that the query should be compared with using `hmmsearch` from the `hmmer` package (<http://hmmer.janelia.org>).

2.3. Metagenomic Sequences Datasets

We tested our algorithm on four datasets each containing metagenomics sequences of a major oceanic region community (the four regions are Sargasso Sea, Coast of British Columbia, Gulf of Mexico, and Arctic Ocean) (referred to as Ocean Virus datasets)²². The reads were acquired by 454 sequencing machine, and they are typically very short. All the metagenomic sequences were downloaded from CAMERA website (<http://camera.calit2.net/>)³⁷.

3. RESULTS

We applied our ORFome assembly tool MetaORFA to assemble the four Ocean Virus datasets. The assembly of a dataset took about from several minutes to half an hour for the four datasets we used here (on a linux machine with Intel(R) Core(TM)2 CPU@ 2.40GHz). The unassembled reads and assembled peptides were searched against Integrated Microbial Genomics (IMG) database³⁸ using BLASTP to identify known homologous proteins in pre-sequenced microbial genomes. To show the improvement of functional coverage after the ORFome assembly, we also searched both sets of sequences against PANTHER families and subfamilies. Below we first report the basic statistics of the assembled peptides as compared to the unassembled reads, and then show the annotation of the ORFs by BLAST search and PANTHER family annotation.

Table 1. Statistics of the ORFs for Ocean Virus datasets

Sample		Num	Min	Max	Ave	Num60
Arctic Ocean	Reads	688590	35	370	99	-
	UA-ORF	1015432	30	58	33	0
	A-Pep	368278	30	175	37	12428
Sargasso Sea	Reads	399343	36	282	104	-
	UA-ORF	345411	30	49	33	0
	A-Pep	214330	30	162	34	1530
Coast of British Columbia	Reads	16456	37	254	102	
	UA-ORF	426666	30	61	33	1
Gulf of Mexico	A-Pep	304106	30	196	34	2559
	Reads	771849	38	246	95	-
	UA-ORF	467085	30	54	33	0
	A-Pep	206111	30	157	34	1703

Num, Min, Max and Ave represent the total number, the minimum, maximum, and average length of the reads (in nucleotides), unassembled ORFs (UA-ORF, in amino acid residues) and assembled peptides (A-Pep, in amino acid residues), respectively. Num60 represents the total number of unassembled ORFs and assembled peptides of length ≥ 60 . We note that the minimum length of unassembled ORFs is 30 because we used a default cutoff 30 to detect ORFs in original short reads, which are used as inputs for MetaORFA.

3.1. Assembled Peptides from the ORFome Assembly

Table 1 shows the statistics of the reads, unassembled putative ORFs and assembled peptides for the four Ocean Viruses datasets. For all four datasets, the ORFome assembly successfully produced long peptides (≥ 60) that are not present in the unassembled reads. However, the number and the length of long peptides are different from one dataset to another. For example, the ORFome assembly produced the largest number (12,428) of long peptides with longest average length (37 aa) in the Arctic Ocean dataset, even though comparable number of sequencing reads were acquired in each of these four datasets. This may indicate either the diversity of the microorganisms in Arctic Ocean sample is lower than the diversity in the other samples, or the microorganism genomes in this sample are more compact than the genomes in the other samples.

We use the second longest peptide assembled from the Gulf of Mexico dataset as an example to illustrate the advantages of the ORFome assembly.^b Fig. 2 shows that 12 putative ORFs detected from different short reads were assembled into a long peptide (144 aa) by the ORFome assembly, which shows strong similarity across the entire peptide with an

annotated protein in IMG database.

3.2. Homology Search of Assembled Peptides

One of the commonly used analysis of metagenomic data is the searching of the unassembled reads against databases of known microbial proteins in an attempt to use the identified homologous proteins to assess the function and species diversity in the sample^{39, 23}. In this type of analysis, a quite high cutoff is often chosen for the BLAST E-values (i.e., less significant) because the query sequences (i.e., reads) are quite short. As a result, there may be many false hits included in the final list of homologous proteins, which can mislead the diversity analysis. Comparing with this straightforward approach, we anticipate the homology search using the assembled peptides from the ORFome assembly can achieve higher sensitivity and result in more hits with higher significance (i.e., lower E-values).

We compared the results of homology searches using assembled peptides with the results using unassembled reads. The four Ocean Virus datasets were tested separately against IMG database. As reported in Ref. 22^c, only few reads hit proteins in the database. We emphasize that the assembled

^bThe longest peptide has 157 amino acids, which only has very weak similarity to the protein sequences collected in IMG; it is hard to identify this peptide based on the similarity search result.

^cWe note that a direct comparison is not feasible since different databases were used for homology searching in these two studies.

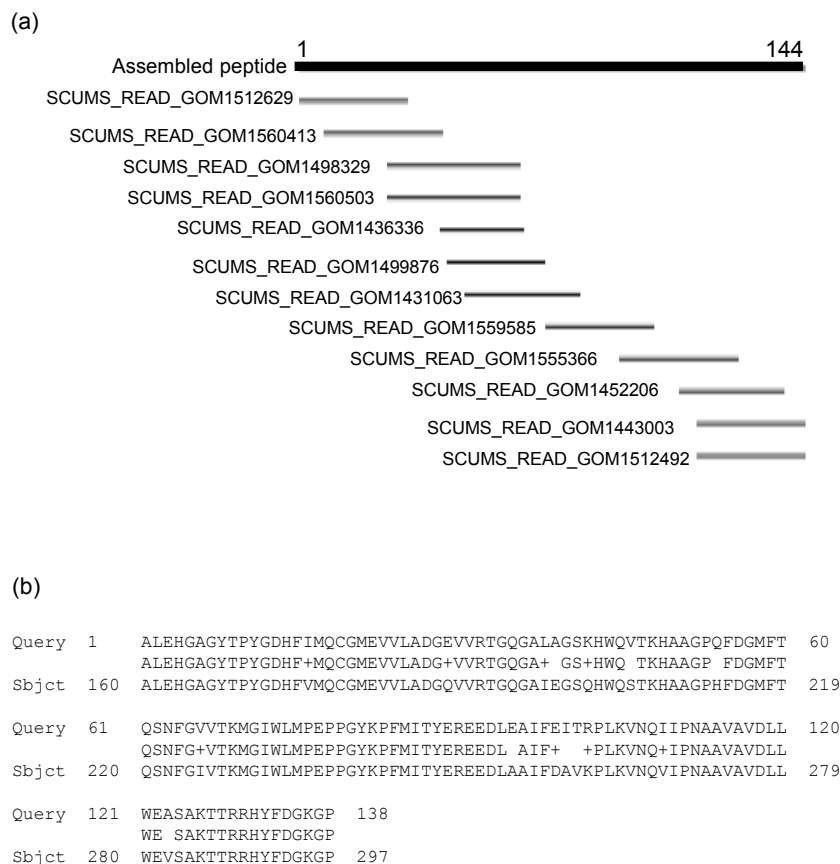


Fig. 2. A long peptide with 144 aa (contig11088, highlighted in bold line) assembled from 12 putative ORFs (represented as thin lines below the contig) in the Gulf of Mexico dataset shows strong similarity with proteins in IMG database with known function (a). (b) shows the BLAST alignment between the peptide and the flavoprotein subunit p-cresol methylhydroxylase from *Novosphingobium aromaticivorans* in IMG database.

peptides increase the number of significant hits (i.e., $E\text{-value} \leq 1e-5$) in all four datasets, from 26% in the Gulf of Mexico dataset (i.e., 2489 read hits were added to 9517 read hits received from the searching using unassembled reads) to 43% in the Arctic Ocean dataset (26,578 read hits were added to 61,578 original read hits). Fig. 3 shows the detailed comparison of the added number of read hits when various E-value cutoffs were applied. For all four datasets, a nearly constant number of read hits can be added by using assembled peptides. In comparison, a majority of read hits from the similarity searching using unassembled reads received high E-values. For in-

stance, there are only 11,144 read hits in the Arctic Ocean dataset with $E\text{-values} \leq 1e-10$, whereas 30,151 additional read hits (i.e., 270% more!) can be added from the similarity searching using assembled peptides. Among four datasets, the improvement of read hits is most significant in the Arctic Ocean dataset, which is consistent with the result that this dataset also achieved the best assembly performance.

3.3. Novel Assignments of Functional Categories by Assembled Peptides

We further assessed the performance of the ORFome assembly in improving the function annotation on

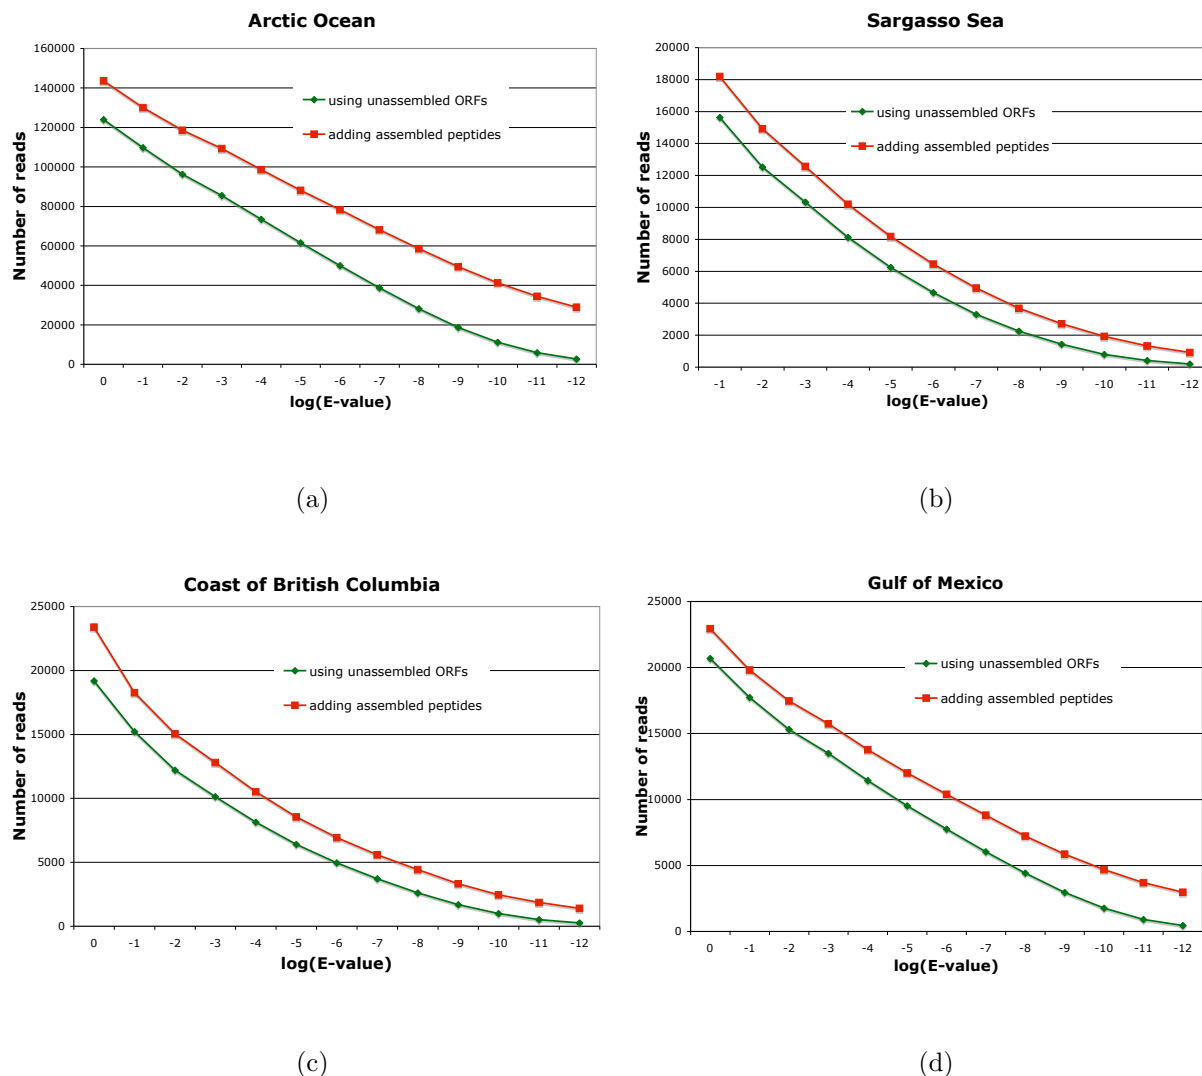


Fig. 3. Detailed comparison of the total number of read hits in IMG database using unassembled and the total number of read hits including those read hits belonging to the assembled peptides at different BLAST E-value cutoffs. The deviation between the two lines indicates the gain of read hits by using assembled peptides from the ORFome assembly.

the Ocean Virus datasets. Table 2 summarizes the statistics of the number of matched families in PANTHER database for all four datasets. Both the number of hits from the searching of unassembled reads as well as the additional number of hits from the searching of assembled peptide are listed. Although the additional numbers of families detected by using assembled peptides are relatively low for all datasets, there are still some new protein families (or novel protein functions) that can be annotated when assembled peptides were used. This suggests that we

may be able to improve the protein function annotation using assembled peptides.

In the Gulf of Mexico dataset, the assembled peptides hit additional 20 PANTHER protein families. One of them is ATP synthase mitochondrial F1 complex assembly factor 2 (Panther family ID PTHR21013). The peptide corresponding to this hit is assembled from two ORFs from different reads (see Fig. 4).

SCUMS_READ_GOM1547922: 1-99
 RTPQRAPLIVASAALAETIAAEW***QDQD******TV******DP******PA***
 RTPQRAPLIVASAALAETIAAEW***QDQD******TV******DP******PA***AAMPITGLTNAAILDLALPDPLGFAE
QDQD***TV******DP******PA***AAMPITGLTNAAILDLALPDPLGFAE
 SCUMS_READ_GOM1406881: 2-103

Fig. 4. An example of assembled peptide in the Gulf of Mexico dataset, which is assembled from two reads. The sequence of the assembled ORF is shown in the box; the overlap of the two putative ORFs from two reads are highlighted by italic and bold fonts. This assembled peptide hits the protein family of ATP synthase mitochondrial F1 complex assembly factor 2 in PANTHER database.

Table 2. Summary of the family annotation of assembled peptides versus unassembled reads for the four ocean virus datasets

Sample	Family	Add-on	Example
Arctic Ocean	590	33	PTHR22748
Sargasso Sea	265	4	PTHR11527
Coast of British Columbia	352	7	PTHR10566
Gulf of Mexico	413	21	PTHR17630

The Family column lists the total number of protein families that are found from unassembled reads. The "Add-on" column lists the additional panther protein families that are detected by using assembled peptides. The last column gives a few examples of the additional protein families (or functions) that are annotated based the assembled peptides only: PTHR22748, AP endonuclease (E-value = 5.4e-12); PTHR11527 (subfamily SF15), heat shock protein 16 (E-value = 1.5e-07); PTHR10566 (subfamily SF7), ubiquinone biosynthesis protein AARF (E.coli)/ABC (Yeast)-related (E-value = 7.3e-11); PTHR17630 (subfamily SF20), carboxymethylenebutenolidase (Evalue = 4.7e-08) .

4. DISCUSSION

One of the main issues in whole genome assembly is the chimeric contigs that are resulted from mis-assemblies. Tremendous finishing efforts have to be invested in order to identify and correct these errors. This issue is expected to be more serious in metagenomics data analysis because of the higher complexity of metagenomics sample and the introduction of short reads. Although it remains unclear whether the mis-assemblies will dramatically influence the conclusion on the principal aims of metagenomics, such as the assessment of species diversity

in the sample, many metagenomics project avoided assembling sequencing reads, and analyzed the original reads directly. The ORFome assembly provides a simple solution to conduct a small-scale but accurate assembly of protein coding regions that can improve the sensitivity of homology search. In this study, although we showed the homology searching was improved after the ORFome assembly, we have not systematically evaluated the influence of these improvements on the diversity analysis. We will apply the ORFome assembly approach to more datasets with various sequence coverage and sample complexities (i.e., the approximate number of species and the range of abundances among these species). Our intention is to estimate the sequencing efforts that are required to get a good assessment of species diversity for samples with different complexities.

There are several ways to further improve the ORFome assembly algorithm described here. First, the current method for predicting putative ORFs in sequencing reads can be improved by incorporating additional features of gene coding sequences (e.g., the codon usages) and utilizing sophisticated probabilistic models. Second, a few parameters (e.g., the length cutoff of putative ORFs) used in the ORFome assembly should be optimized. This indicates that there is still room for the further improvement of the ORFome assembly approach by selecting more appropriate parameters. Finally, as we mentioned in the METHODS section, the advantages of the

ORFome assembly have not been fully taken in the downstream data analysis in this study. The EULER method used here can assemble putative ORFs into a repeat graph, in addition to the peptides represented by edges in the graph. Therefore, one can adopt a network matching approach as used in Ref. 40 to achieve a more sensible database searching.

Finally we point out that the basic method we adopted for ORF prediction may generate some spurious peptides, and some of the assembled ORFs may be not real proteins. Those spurious peptides may not cause serious problems in applications such as the homology search based annotations as used in this paper. However, we should not neglect them in other types of applications, such as comparison of the number of protein clusters (families) among different metagenomic datasets.

5. CONCLUSION

We present a novel ORFome assembly approach to metagenomics data analysis. The application of this method on four metagenomics datasets achieved promising results. Even with low coverage short reads from these datasets, our method has assembled many long peptides, which can hit on annotated proteins in the database that are not detectable otherwise. The ORFome assembly provides a useful tool to retrieve rich information from metagenomic sequencing reads, and it shows potential to facilitate an accurate assessment of the species and functional diversity in metagenomics.

Acknowledgement

This research was supported by the Indiana METACyt Initiative of Indiana University, funded in part through a major grant from the Lilly Endowment, Inc, and School of Informatics. The authors thank the University Information Technology Services team in Indiana University for their help with high-performance computing (for BLAST search).

References

- Mardis E. Anticipating the 1,000 dollar genome. *Genome Biol*, 2006; **7**:112.
- Lane D, Pace B, Olsen G, et al. Rapid determination of 16S ribosomal RNA sequences for phylogenetic analyses. *Proc Natl Acad Sci USA*, 1985; **82**:6955–6959.
- Breitbart M, Salamon P, Andresen B, et al. Genomic analysis of uncultured marine viral communities. *Proc Natl Acad Sci USA*, 2002; **99**:14250–14255.
- Galperin M. Metagenomics: from acid mine to shining sea. *Environ Microbiol*, 2004; **6**:543–545.
- Eyers L, George I, Schuler L, et al. Environmental genomics: exploring the unmined richness of microbes to degrade xenobiotics. *Appl Microbiol Biotechnol*, 2004; **66**:123–130.
- Streit W, Schmitz R. Metagenomics—the key to the uncultured microbes. *Curr Opin Microbiol*, 2004; **7**:492–498.
- Riesenfeld C, Schloss P, Handelsman J. Metagenomics: genomic analysis of microbial communities. *Annu Rev Genet*, 2004; **38**:525–552.
- Venter J, Remington K, Heidelberg J, et al. Environmental genome shotgun sequencing of the Sargasso Sea. *Science*, 2004; **304**:66–74.
- Tyson G, Chapman J, Hugenholtz P, et al. Community structure and metabolism through reconstruction of microbial genomes from the environment. *Nature*, 2004; **428**:37–43.
- Lorenz P, Eck J. Metagenomics and industrial applications. *Nat Rev Microbiol*, 2005; **3**:510–516.
- Turnbaugh P, Ley R, Mahowald M, et al. An obesity-associated gut microbiome with increased capacity for energy harvest. *Nature*, 2006; **444**:1027–1031.
- Gill S, Pop M, Deboy R, et al. Metagenomic analysis of the human distal gut microbiome. *Science*, 2006; **312**:1355–1359.
- Chen K, Pachter L. Bioinformatics for whole-genome shotgun sequencing of microbial communities. *PLoS Comput Biol*, 2005; **1**:106–112.
- Foerstner K, von Mering C, Bork P. Comparative analysis of environmental sequences: potential and challenges. *Philos Trans R Soc Lond B Biol Sci*, 2006; **361**:519–523.
- Batzoglou S, Jaffe D, Stanley K, et al. ARACHNE: a whole-genome shotgun assembler. *Genome Res*, 2002; **12**:177–189.
- Jaffe D, Butler J, Gnerre S, et al. Whole-genome sequence assembly for mammalian genomes: Arachne 2. *Genome Res*, 2003; **13**:91–96.
- Huson D, Reinert K, Kravitz S, et al. Design of a compartmentalized shotgun assembler for the human genome. *Bioinformatics*, 2001; **17 Suppl 1**:S132–139.
- Aparicio S, Chapman J, Stupka E, et al. Whole-genome shotgun assembly and analysis of the genome of *Fugu rubripes*. *Science*, 2002; **297**:1301–1310.
- Azad R, Borodovsky M. Probabilistic methods of identifying genes in prokaryotic genomes: connections to the HMM theory. *Brief Bioinformatics*, 2004; **5**:118–130.
- Yooseph S, Sutton G, Rusch D, et al. The Sorcerer II Global Ocean Sampling expedition: expanding the universe of protein families. *PLoS Biol*, 2007; **5**:e16.

21. Mavromatis K, Ivanova N, Barry K, et al. Use of simulated data sets to evaluate the fidelity of metagenomic processing methods. *Nat Methods*, 2007; **4**:495–500.
22. Angly F, Felts B, Breitbart M, et al. The marine viromes of four oceanic regions. *PLoS Biol*, 2006; **4**:e368.
23. Huson D, Auch A, Qi J, et al. MEGAN analysis of metagenomic data. *Genome Res*, 2007; **17**:377–386.
24. Wommack K, Bhavsar J, Ravel J. Metagenomics: read length matters. *Appl Environ Microbiol*, 2008; **74**:1453–1463.
25. Tang H. Genome assembly, rearrangement, and repeats. *Chem Rev*, 2007; **107**:3391–3406.
26. Pevzner P, Tang H, Waterman M. An Eulerian path approach to DNA fragment assembly. *Proc Natl Acad Sci USA*, 2001; **98**:9748–9753.
27. Myers E. Toward simplifying and accurately formulating fragment assembly. *J Comput Biol*, 1995; **2**:275–290.
28. Chaisson M, Pevzner P, Tang H. Fragment assembly with short reads. *Bioinformatics*, 2004; **20**:2067–2074.
29. Myers E, Sutton G, Delcher A, et al. A whole-genome assembly of *Drosophila*. *Science*, 2000; **287**:2196–2204.
30. Venter J, Adams M, Myers E, et al. The sequence of the human genome. *Science*, 2001; **291**:1304–1351.
31. Morgulis A, Gertz E, Scher A, et al. A fast and symmetric DUST implementation to mask low-complexity DNA sequences. *J Comput Biol*, 2006; **13**:1028–1040.
32. Benson G. Tandem repeats finder: a program to analyze DNA sequences. *Nucleic Acids Res*, 1999; **27**:573–580.
33. Pevzner PA, Tang H, Tesler G. De novo repeat classification and fragment assembly. *Genome Res*, 2004; **14**(9):1786–1796.
34. Raphael B, Zhi D, Tang H, et al. A novel method for multiple alignment of sequences with repeated and shuffled elements. *Genome Res*, 2004; **14**:2336–2346.
35. Rodriguez-Brito B, Rohwer F, Edwards R. An application of statistics to comparative metagenomics. *BMC Bioinformatics*, 2006; **7**:162.
36. Thomas P, Campbell M, Kejariwal A, et al. PANTHER: a library of protein families and subfamilies indexed by function. *Genome Res*, 2003; **13**:2129–2141.
37. Seshadri R, Kravitz S, Smarr L, et al. CAMERA: a community resource for metagenomics. *PLoS Biol*, 2007; **5**:e75.
38. Markowitz V, Szeto E, Palaniappan K, et al. The integrated microbial genomes (IMG) system in 2007: data content and analysis tool extensions. *Nucleic Acids Res*, 2007; .
39. Edwards R, Rohwer F. Viral metagenomics. *Nat Rev Microbiol*, 2005; **3**:504–510.
40. Ye Y, Jaroszewski L, Li W, et al. A segment alignment approach to protein comparison. *Bioinformatics*, 2003; **19**:742–749.