
Metagenomics Workshop Documentation

Release 1

Alex Sczyrba

Oct 05, 2021

Contents

1	Setting up your de.NBI Cloud SimpleVM instance	3
2	The Tutorial Data Set	5
3	FastQC Quality Control	7
4	Assembly	9
4.1	Velvet Assembly	9
4.2	MEGAHIT Assembly	10
4.3	metaSPAdes Assembly	10
4.4	IDBA-UD Assembly	11
4.5	Ray Assembly	11
5	Gene Prediction	13
6	Assembly Evaluation	15
6.1	Read Mapping	15
6.2	MetaQUAST	16
7	Binning	19
7.1	MaxBin Binning	19
7.2	MetaBAT Binning	20
8	Classification	21
8.1	Genome Taxonomy Database (GTDB)	21
8.2	Kraken Taxonomic Sequence Classification System	22

Welcome to the one-day metagenomics assembly workshop. This tutorial will guide you through the typical steps of metagenome assembly and binning.

CHAPTER 1

Setting up your de.NBI Cloud SimpleVM instance

As metagenome assemblies require a lot of compute resources, we will run the tutorial in the *de.NBI Cloud* <<http://cloud.denbi.de>>. Each workshop participant has access to a virtual machine and run all jobs on it. All necessary software tools have been pre-installed on this VM.

Use the link you received via e-Mail to connect to your instance and open a terminal.

CHAPTER 2

The Tutorial Data Set

We have prepared a small toy data set for this tutorial. Please use the following commands to download the data to your VM:

```
sudo chown ubuntu:ubuntu /mnt
cd /mnt
wget https://openstack.cebitec.uni-bielefeld.de:8080/swift/v1/denbi-mg-course/WGS-
↪data.tar
tar xvf WGS-data.tar
```

The */mnt/WGS-data* directory has the following content:

File	Content
genomes/	Directory containing the reference genomes
gold_std/	Gold Standard assemblies
read1.fq	Read 1 of paired reads (FASTQ)
read2.fq	Read 2 of paired reads (FASTQ)
reads.fas	Shuffled reads (FASTA)

CHAPTER 3

FastQC Quality Control

FastQC aims to provide a simple way to do some quality control checks on raw sequence data coming from high throughput sequencing pipelines. It provides a modular set of analyses which you can use to give a quick impression of whether your data has any problems of which you should be aware before doing any further analysis.

The main functions of FastQC are

- Import of data from BAM, SAM or FastQ files (any variant)
- Providing a quick overview to tell you in which areas there may be problems
- Summary graphs and tables to quickly assess your data
- Export of results to an HTML based permanent report
- Offline operation to allow automated generation of reports without running the interactive application

See the [FastQC home page](#) for more info.

To run FastQC on our data, simply type:

```
cd /mnt/WGS-data
fastqc read1.fq read2.fq
```

After FastQC finished running you can access the report using a web browser:

```
firefox *.html
```

Check out the [FastQC home page](#) for examples of reports including bad data.

We are going to use different assemblers and compare the results.

4.1 Velvet Assembly

Velvet was one of the first de novo genomic assemblers specially designed for short read sequencing technologies. It was developed by Daniel Zerbino and Ewan Birney at the European Bioinformatics Institute (EMBL-EBI). Velvet currently takes in short read sequences, removes errors then produces high quality unique contigs. It then uses paired-end read and long read information, when available, to retrieve the repeated areas between contigs. See the [Velvet GitHub page](#) for more info.

4.1.1 Step 1: velveth

velveth takes in a number of sequence files, produces a hashtable, then outputs two files in an output directory (creating it if necessary), Sequences and Roadmaps, which are necessary for running velvetg in the next step.

Let's create multiple hashtables using kmer-lengths of 31 and 51. We are going to run two jobs in parallel:

```
cd /mnt/WGS-data

velveth velvet_31 31 -shortPaired -fastq -separate read1.fq read2.fq &
velveth velvet_51 51 -shortPaired -fastq -separate read1.fq read2.fq &
```

Once the two jobs are finished (use *top* to monitor your jobs), you should have two output directories for the two different kmer-lengths: *velvet_31* and *velvet_51*.

4.1.2 Step 2: velvetg

Now we have to start the actual assembly using velvetg. velvetg is the core of Velvet where the de Bruijn graph is built then manipulated. Let's run assemblies for both kmer-lengths. See the [Velvet manual](#) for more info about parameter settings. Again, we submit the job to the compute cluster:

```
velvetg velvet_31 -cov_cutoff auto -ins_length 270 -min_contig_lgth 500 -exp_cov auto_
↪&
velvetg velvet_51 -cov_cutoff auto -ins_length 270 -min_contig_lgth 500 -exp_cov auto_
↪&
```

The contig sequences are located in the *velvet_31* and *velvet_51* directories in file *contigs.fa*. Let's get some very basic statistics on the contigs. The script `getN50.pl` reads the contig file and computes the total length of the assembly, number of contigs, N50 and largest contig size. In our example we will exclude contigs shorter than 500bp (option `-s 500`):

```
getN50.pl -s 500 -f velvet_31/contigs.fa
getN50.pl -s 500 -f velvet_51/contigs.fa
```

4.2 MEGAHIT Assembly

MEGAHIT is a single node assembler for large and complex metagenomics NGS reads, such as soil. It makes use of succinct de Bruijn graph (SDBG) to achieve low memory assembly. MEGAHIT can optionally utilize a CUDA-enabled GPU to accelerate its SDBG construction. See the [MEGAHIT home page](#) for more info.

MEGAHIT can be run by the following command. As our compute instance have multiple cores, we use the option `-t 14` to tell MEGAHIT it should use 14 parallel threads. The output will be redirected to file *megahit.log*:

```
cd /mnt/WGS-data
megahit -1 read1.fq -2 read2.fq -t 14 -o megahit_out
```

The contig sequences are located in the *megahit_out* directory in file *final.contigs.fa*. Again, let's get some basic statistics on the contigs:

```
getN50.pl -s 500 -f megahit_out/final.contigs.fa
```

4.3 metaSPAdes Assembly

SPAdes – St. Petersburg genome assembler – is an assembly toolkit containing various assembly pipelines. See the [SPAdes home page](#) for more info.

metaSPAdes can be run by the following command:

```
cd /mnt/WGS-data
metaspades.py -o metaspades_out --pe1-1 read1.fq --pe1-2 read2.fq
```

The contig sequences are located in the *metaspades_out* directory in file *contigs.fasta*. Again, let's get some basic statistics on the contigs:

```
getN50.pl -s 500 -f metaspades_out/contigs.fasta
```

4.4 IDBA-UD Assembly

IDBA is the basic iterative de Bruijn graph assembler for second-generation sequencing reads. IDBA-UD, an extension of IDBA, is designed to utilize paired-end reads to assemble low-depth regions and use progressive depth on contigs to reduce errors in high-depth regions. It is a generic purpose assembler and especially good for single-cell and metagenomic sequencing data. See the [IDBA home page](#) for more info.

IDBA-UD requires paired-end reads stored in single FastA file and a pair of reads is in consecutive two lines. You can use *fq2fa* (part of the IDBA repository) to merge two FastQ read files to a single file. The following command will generate a FASTA formatted file called *reads12.fas* by “shuffling” the reads from FASTQ files *read1.fq* and *read2.fq*:

```
cd /mnt/WGS-data
fq2fa --merge read1.fq read2.fq reads12.fas
```

IDBA-UD can be run by the following command. As our compute instances have multiple cores, we use the option *--num_threads 14* to tell IDBA-UD it should use 14 parallel threads:

```
cd /mnt/WGS-data
idba_ud -r reads12.fas --num_threads 14 -o idba_ud_out
```

The contig sequences are located in the *idba_ud_out* directory in file *contig.fa*. Again, let’s get some basic statistics on the contigs:

```
getN50.pl -s 500 -f idba_ud_out/contig.fa
```

4.5 Ray Assembly

Ray is a parallel software that computes de novo genome assemblies with next-generation sequencing data. Ray is written in C++ and can run in parallel on numerous interconnected computers using the message-passing interface (MPI) standard. See the [Ray home page](#) for more info.

Ray can be run by the following command using a kmer-length of 51 and 31, respectively. As our compute instance have multiple cores, we specify this in the ‘*mpirun -n 14*’ command to let Ray know it should use 14 parallel MPI processes:

```
cd /mnt/WGS-data
mpirun -n 14 /usr/local/bin/Ray -k 51 -p read1.fq read2.fq -o ray_51
```

If there is enough time, you can run another Ray assembly using a smaller kmer size:

```
mpirun -n 14 /usr/local/bin/Ray -k 31 -p read1.fq read2.fq -o ray_31
```

This will create the output directory *ray_51* (and *ray_31*), the final contigs are located in *ray_51/Contigs.fasta* (and *ray_31/Contigs.fasta*). Again, let’s get some basic statistics on the contigs:

```
getN50.pl -s 500 -f ray_51/Contigs.fasta
getN50.pl -s 500 -f ray_31/Contigs.fasta
```

Now that you have run assemblies using Velvet, MEGAHIT, IDBA-UD and Ray, let’s have a quick look at the assembly statistics of all of them:

```
cd /mnt/WGS-data
sh ./get_assembly_stats.sh
```


CHAPTER 5

Gene Prediction

Prodigal (Prokaryotic Dynamic Programming Genefinding Algorithm) is a microbial (bacterial and archaeal) gene finding program developed at Oak Ridge National Laboratory and the University of Tennessee. See the [Prodigal home page](#) for more info.

To run `prodigal` on our data, simply type:

```
cd /mnt/WGS-data/megahit_out

prodigal -p meta -a final.contigs.genes.faa -d final.contigs.genes.fna -f gff -o_
↪final.contigs.genes.gff -i final.contigs.fa
```

Output files:

<code>final.contigs.genes.gff</code>	positions of predicted genes in GFF format
<code>final.contigs.genes.faa</code>	protein translations of predicted genes
<code>final.contigs.genes.fna</code>	nucleotide sequences of predicted genes

Assembly Evaluation

We are going to evaluate our assemblies using the reference genomes.

6.1 Read Mapping

In this part of the tutorial we will look at the assemblies by mapping the reads to the assembled contigs. Different tools exist for mapping reads to genomic sequences such as [bowtie](#) or [bwa](#). Today, we will use the tool BBMap.

BBMap: Short read aligner for DNA and RNA-seq data. Capable of handling arbitrarily large genomes with millions of scaffolds. Handles Illumina, PacBio, 454, and other reads; very high sensitivity and tolerant of errors and numerous large indels. Very fast. See the [BBTools home page](#) for more info.

bbmap needs to build an index for the contigs sequences before it can map the reads onto them. Here is an example command line for mapping the reads back to the MEGAHIT assembly:

```
cd /mnt/WGS-data/megahit_out  
bbmap.sh ref=final.contigs.fa
```

Now that we have an index, we can map the reads:

```
bbmap.sh in=../read1.fq in2=../read2.fq out=megahit.bam threads=14
```

bbmap produces output in BAM format (the binary version of the [SAM format](#)). BAM files can be viewed and manipulated with [SAMtools](#). Let's first build an index for the FASTA file:

```
samtools faidx final.contigs.fa
```

We have to sort the BAM file by starting position of the alignments. This can be done using samtools again:

```
samtools sort -o megahit_sorted.bam -@ 14 megahit.bam
```

Now we have to index the sorted BAM file:

```
samtools index megahit_sorted.bam
```

To look at the BAM file use:

```
samtools view megahit_sorted.bam | less
```

We will use the IGV genome browser to look at the mappings:

```
igv
```

Now let's look at the mapped reads:

1. Load the contig sequences into IGV. Use the menu Genomes->Load Genome from File...
2. Load the BAM file into IGV. Use menu File->Load from File...
3. Load the predicted genes as another track. Use menu File->Load from File... to load the GFF file.

6.2 MetaQUAST

QUAST stands for Quality ASsessment Tool. The tool evaluates genome assemblies by computing various metrics. You can find all project news and the latest version of the tool at [sourceforge](https://openstack.cebitec.uni-bielefeld.de:8080/swift/v1/denbi-mg-course/). QUAST utilizes MUMmer, GeneMarkS, GeneMark-ES, GlimmerHMM, and GAGE. In addition, MetaQUAST uses MetaGeneMark, Krona tools, BLAST, and SILVA 16S rRNA database. See the [metaQuast home page](#) for more info.

In case not all assemblies finished so far, copy the pre-computed assembly results to your local directory:

```
cd /mnt/WGS-data
wget https://openstack.cebitec.uni-bielefeld.de:8080/swift/v1/denbi-mg-course/
↪assembly_results.tar
tar xvf assembly_results.tar
cp -v -r /mnt/WGS-data/assembly_results/* .
```

To call `metaquast.py` we have to provide reference genomes which are used to calculate a number of different metrics for evaluation of the assembly. In real-world metagenomics, these references are usually not available, of course:

```
cd /mnt/WGS-data

metaquast.py --threads 14 --gene-finding \
-R /mnt/WGS-data/genomes/Aquifex_aeolicus_VF5.fna,\
/mnt/WGS-data/genomes/Bdellovibrio_bacteriovorus_HD100.fna,\
/mnt/WGS-data/genomes/Chlamydia_psittaci_MN.fna,\
/mnt/WGS-data/genomes/Chlamydophila_pneumoniae_CWL029.fna,\
/mnt/WGS-data/genomes/Chlamydophila_pneumoniae_J138.fna,\
/mnt/WGS-data/genomes/Chlamydophila_pneumoniae_LPCoLN.fna,\
/mnt/WGS-data/genomes/Chlamydophila_pneumoniae_TW_183.fna,\
/mnt/WGS-data/genomes/Chlamydophila_psittaci_C19_98.fna,\
/mnt/WGS-data/genomes/Finegoldia_magna_ATCC_29328.fna,\
/mnt/WGS-data/genomes/Fusobacterium_nucleatum_ATCC_25586.fna,\
/mnt/WGS-data/genomes/Helicobacter_pylori_26695.fna,\
/mnt/WGS-data/genomes/Lawsonia_intracellularis_PHE_MN1_00.fna,\
/mnt/WGS-data/genomes/Mycobacterium_leprae_TN.fna,\
/mnt/WGS-data/genomes/Porphyromonas_gingivalis_W83.fna,\
/mnt/WGS-data/genomes/Wigglesworthia_glossinidia.fna \
-o quast \
```

(continues on next page)

(continued from previous page)

```
-l MegaHit,metaSPAdes,Ray_31,Ray_51,velvet_31,velvet_51,idba_ud \  
megahit_out/final.contigs.fa \  
metaspades_out/contigs.fasta \  
ray_31/Contigs.fasta \  
ray_51/Contigs.fasta \  
velvet_31/contigs.fa \  
velvet_51/contigs.fa \  
idba_ud_out/contig.fa
```

QUAST generates HTML reports including a number of interactive graphics. To access these reports, load the reports in your web browser:

```
firefox quast/report.html
```


After the assembly of metagenomic sequencing reads into contigs, binning algorithms try to recover individual genomes to allow access to uncultivated microbial populations that may have important roles in the samples community.

7.1 MaxBin Binning

MaxBin is a software that is capable of clustering metagenomic contigs into different bins, each consists of contigs from one species. MaxBin uses the nucleotide composition information and contig abundance information to do achieve binning through an Expectation-Maximization algorithm. For users' convenience MaxBin will report genome-related statistics, including estimated completeness, GC content and genome size in the binning summary page. See the [MaxBin home page](#) for more info.

Let's run a MaxBin binning on the MEGAHIT assembly. First, we need to generate an abundance file from the mapped reads:

```
cd /mnt/WGS-data/megahit_out
mkdir maxbin
cd maxbin
```

Activate the conda base environment to get MaxBin into the PATH:

```
conda activate base
```

Then:

```
pileup.sh in=../megahit_sorted.bam out=cov.txt
awk '{print $1"\t"$5}' cov.txt | grep -v '^#' > abundance.txt
```

Next, we can run MaxBin:

```
run_MaxBin.pl -thread 14 -contig ../final.contigs.fa -out maxbin -abund abundance.txt
```

Assume your output file prefix is (out). MaxBin will generate information using this file header as follows.

(out).0XX.fasta	the XX bin. XX are numbers, e.g. out.001.fasta
(out).summary	summary file describing which contigs are being classified into which bin.
(out).log	log file recording the core steps of MaxBin algorithm
(out).marker	marker gene presence numbers for each bin. This table is ready to be plotted by R or other 3rd-party software.
(out).marker.pdf	visualization of the marker gene presence numbers using R
(out).noclass	all sequences that pass the minimum length threshold but are not classified successfully.
(out).tooshort	all sequences that do not meet the minimum length threshold.

Now you can run a gene prediction on each genome bin and BLAST one sequence for each bin for a (very crude!) classification:

```
for i in max*fasta; do prodigal -p meta -a $i.genes.faa -d $i.genes.fna -f gff -o $i.  
→genes.gff -i $i& done
```

Does the abundance of the bins match the 16S profile of the community?

7.2 MetaBAT Binning

MetaBAT, An Efficient Tool for Accurately Reconstructing Single Genomes from Complex Microbial Communities.

Grouping large genomic fragments assembled from shotgun metagenomic sequences to deconvolute complex microbial communities, or metagenome binning, enables the study of individual organisms and their interactions. MetaBAT is an automated metagenome binning software which integrates empirical probabilistic distances of genome abundance and tetranucleotide frequency. See the [MetaBAT home page](#) for more info.

Let's run a MetaBAT binning on the MEGAHIT assembly:

```
cd /mnt/WGS-data/megahit_out  
mkdir metabat  
cd metabat  
  
runMetaBat.sh ../megahit_out/final.contigs.fa ../megahit_out/megahit_sorted.bam
```

MetaBAT will generate 11 bins from our assembly:

```
ls final.contigs.fa.metabat-bins  
  
bin.1.fa  
bin.2.fa  
bin.3.fa  
bin.4.fa  
bin.5.fa  
bin.6.fa  
bin.7.fa  
bin.8.fa  
bin.9.fa  
bin.10.fa  
bin.11.fa
```


Taxonomic classification tools assign taxonomic labels to reads or assembled contigs of metagenomic datasets.

8.1 Genome Taxonomy Database (GTDB)

GTDB-Tk is a software toolkit for assigning objective taxonomic classifications to bacterial and archaeal genomes based on the [Genome Database Taxonomy GTDB](#). It is designed to work with recent advances that allow hundreds or thousands of metagenome-assembled genomes (MAGs) to be obtained directly from environmental samples. It can also be applied to isolate and single-cell genomes.

See the [GTDBTk homepage](#) for more info.

First, we need to download the GTDB database files. The database is pretty big (33 Gb), so even downloading it from our local copy in the de.NBI Cloud will take a couple of minutes. After downloading, we need to extract the tar archive (please be patient ;):

```
cd /mnt
wget https://openstack.cebitec.uni-bielefeld.de:8080/swift/v1/denbi-mg-course/gtdbtk_
↪r95_data.tar.gz
tar xvzf gtdbtk_r95_data.tar.gz
```

Now we need to set an environment variable that stores the path to the database:

```
export GTDBTK_DATA_PATH=/mnt/release95
```

Next, let's assign taxonomic labels to our binning results using GTDB-Tk:

```
cd /mnt/WGS-data/megahit_out/maxbin
gtdbtk classify_wf --extension fasta --cpus 14 --genome_dir . --out_dir gtdbtk_out
```

8.2 Kraken Taxonomic Sequence Classification System

Kraken is a system for assigning taxonomic labels to short DNA sequences, usually obtained through metagenomic studies. Kraken aims to achieve high sensitivity and high speed by utilizing exact alignments of k-mers and a novel classification algorithm.

In its fastest mode of operation, for a simulated metagenome of 100 bp reads, Kraken processed over 4 million reads per minute on a single core, over 900 times faster than Megablast and over 11 times faster than the abundance estimation program MetaPhlAn. Kraken's accuracy is comparable with Megablast, with slightly lower sensitivity and very high precision.

See the [Kraken home page](#) for more info.

Let's assign taxonomic labels to our binning results using Kraken. First, we need to compare the genome bins against the Kraken database:

```
cd /vol/spool/workdir/assembly/megahit_out/maxbin
mkdir kraken

for i in maxbin./*.fasta
do
qsub -cwd -N kraken_${i} -b y \
/usr/local/bin/kraken --db /usr/local/share/krakendb --threads 1 --fasta-input ${i} --
↪output kraken/${i}.kraken
done
```

If you need the full taxonomic name associated with each input sequence, Kraken provides a script named `kraken-translate` that produces two different output formats for classified sequences. The script operates on the output of `kraken`:

```
cd kraken
for i in *.kraken
do
qsub -cwd -b y -o ${i}.labels \
/usr/local/bin/kraken-translate --db /usr/local/share/krakendb ${i}
done
```

Does the abundance of the bins match the 16S profile of the community?