

---

# **Metagenomics Workshop SciLifeLab Documentation**

***Release 1.0***

**Johannes Alneberg, John Larsson, Ino de Bruijn, Luisa Hugerth, A**

November 09, 2016



<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Connecting to UPPMAX</b>	<b>5</b>
2.1	Connecting to UPPMAX . . . . .	5
2.2	Getting a node of your own . . . . .	5
2.3	Load virtual environment . . . . .	6
2.4	Set sample variables . . . . .	6
<b>3</b>	<b>Checking required software</b>	<b>7</b>
3.1	Programs used in this workshop . . . . .	7
3.2	Using which to locate a program . . . . .	7
3.3	Check all programs in one go with which . . . . .	8
3.4	(Optional exercise) Install sickle by yourself . . . . .	8
<b>4</b>	<b>Quality Control</b>	<b>11</b>
4.1	Quality Control with FastQC . . . . .	11
4.2	Optional: Quality trimming Illumina paired-end reads . . . . .	12
<b>5</b>	<b>Community analysis using 16S reads</b>	<b>15</b>
5.1	Community analysis using rRNA gene reads . . . . .	15
<b>6</b>	<b>Metagenomic Assembly</b>	<b>17</b>
6.1	Assembling reads with Velvet . . . . .	17
<b>7</b>	<b>Taxonomic Classification</b>	<b>21</b>
7.1	Phylogenetic Classification using Phylosift . . . . .	21
<b>8</b>	<b>Functional Annotation</b>	<b>23</b>
8.1	Annotating the assembly using the PROKKA pipeline . . . . .	23
8.2	Predicting metabolic pathways using MinPath . . . . .	24
8.3	Mapping reads and quantifying genes . . . . .	25
8.4	Summarize and explore the functional annotation . . . . .	27



This is a one day metagenomics workshop. We will discuss quality checking, assembly, taxonomic classification, binning and annotation of metagenomic samples. This workshop is developed by the 'Environmental Genomics group' at KTH / SciLifeLab.

Here is a [link](#) to the official homepage for the SciLifeLab workshop in metagenomics.

A presentation for this workshop is available [here](#).

Program:

- **Getting started**

- *Introduction*
- *Connecting to UPPMAX*
- *Checking required software*

- **Sessions**

- *Quality Control*
- *Community analysis using 16S reads*
- *Metagenomic Assembly*
- *Taxonomic Classification*
- *Functional Annotation*

Contents:



---

# Introduction

---

In this workshop we will be working with human associated metagenomes from the human microbiome project ([HMP](#)), following all the bioinformatic steps beginning with read sequence quality checking and ending with functional annotation of assembled contigs. You will be asked to choose one dataset to work with throughout the entire workshop and in the end we will compare results among different groups. The three datasets you can choose from is:

- Skin metagenome
- Tooth metagenome
- Gut metagenome





---

## Connecting to UPPMAX

---

### 2.1 Connecting to UPPMAX

The first step of this lab is to open a ssh connection to the computer cluster Milou on [UPPMAX](#). If you have a Mac or a PC running Linux, start the terminal (black screen icon). If you work on a PC running Windows, download and start MobaXterm (<http://mobaxterm.mobatek.net>). Now type (change username to your own username):

```
ssh -X username@milou.uppmax.uu.se
```

and give your password when prompted. As you type the password, nothing will show on screen. No stars, no dots. It is supposed to be that way. Just type the password and press enter, it will be fine. You should now get a welcoming message from Uppmax to show that you have successfully logged in.

### 2.2 Getting a node of your own

Usually you would do most of the work in this lab directly on one of the login nodes at uppmax, but we have arranged for you to have half of one node (=8 cores) each to avoid disturbances. To get this reservation you need to use the `salloc` command like this:

```
salloc -A g2015028 -t 08:00:00 -p core -n 8 --no-shell --reservation=g2015028_1 &
```

Now check which node you got (replace username with your uppmax user name) like this:

```
squeue -u username
```

The `nodelist` column gives you the name of the node that has been reserved for you (starts with “m”). Connect to that node using:

```
ssh -X nodename
```

Note: there is a uppmax specific tool called `jobinfo` that supplies the same kind of information as `squeue` that you can use as well ( `$ jobinfo -u username`). You are now logged in to your reserved node, and there is no need for you to use the SLURM queuing system. You can now continue with the specific exercise instructions.

**IMPORTANT:** If it happens that you are logged out from your Uppmax session during the course (for instance during lunch) you **should not** run the `salloc` command. Instead **just login to the same node** using `ssh`.

## 2.3 Load virtual environment

We have already installed all programs for you, all you have to do is load the virtual environment for this workshop. Once you are logged in to the server run:

```
source /proj/g2015028/metagenomics/virtenv/bin/activate
```

If you would have to, you deactivate the virtual environment with the command *deactivate*, but you don't have to do that yet.

NOTE: This is a [python virtual environment](#). The binary folder of the virtual environment has symbolic links to all programs used in this workshop so you should be able to run those without problems.

## 2.4 Set sample variables

You will now have to make your decision on which kind of dataset you want to work with during this workshop. The choices you have are three different sampling sites on or within the human body:

- Gut
- Skin
- Teeth

**Run only *\*one\** of the following commands in the terminal**

This will set the `SAMPLE` and `SAMPLE_ID` variables that will be used in the commands in the next steps of the tutorial. If for some reason you have to restart the terminal you will have to set these variable names again.

### 2.4.1 Gut

```
SAMPLE=gut  
SAMPLE_ID=SRS011405
```

### 2.4.2 Teeth

```
SAMPLE=teeth  
SAMPLE_ID=SRS014690
```

### 2.4.3 Skin

```
SAMPLE=skin  
SAMPLE_ID=SRS015381
```

After you have chosen a sample you will create the file structure continuously throughout the workshop. This will make it possible for us to only use '`$SAMPLE`' in the commands, and it will automatically be changed to the sample type that you chose. [Here](#) you can see an overview of what this structure should look like at the end of the day (the "results" part of this structure).

---

## Checking required software

---

An often occurring theme in bioinformatics is installing software. Here we will go over some steps to help you check whether you actually have the right software installed. There's an optional exercise on how to install the quality trimmer `sickle`.

### 3.1 Programs used in this workshop

The following programs are used in this workshop:

- Bowtie2
- Velvet
- samtools
- Picard
- Phylosift
- Fastqc
- Sortmerna
- Rdp\_Classifier
- Krona
- Prokka
- MinPath
- BedTools

### 3.2 Using which to locate a program

An easy way to determine whether you have a certain program installed is by typing:

```
which programname
```

where `programname` is the name of the program you want to use. The program `which` searches all directories in `$PATH` for the executable file `programname` and returns the path of the first found hit. This is exactly what happens when you would just type `programname` on the command line, but then `programname` is also executed. To see what your `$PATH` looks like, simply echo it:

```
echo $PATH
```

For more information on the `$PATH` variable see this link: [http://www.linfo.org/path\\_env\\_var.html](http://www.linfo.org/path_env_var.html).

### 3.3 Check all programs in one go with `which`

To check whether you have all programs installed in one go, you can use `which`. In order to do so we will iterate over all the programs and call `which` on each of them. First make a variable containing all programs separated by whitespace:

```
req_progs="bowtie2 bowtie2-build velveth velvetg parallel samtools interleave-reads.py phylosift fastq
echo $req_progs
```

Now iterate over the variable `req_progs` and call `which`:

```
for p in $req_progs; do which $p || echo $p not in PATH; done
```

In Unix-like systems a program that successfully completes its tasks should return a zero exit status. For the program `which` that is the case if the program is found. The `||` character does not mean *pipe the output onward* as you are probably familiar with (otherwise see <http://tldp.org/HOWTO/Bash-Prog-Intro-HOWTO-4.html>), but checks whether the program exists successfully and executes the part behind it if not.

If any of the installed programs is missing, try to install them yourself or ask. If you are having troubles following these examples, try to find some bash tutorials online next time you have some time to kill. Educating yourself on how to use the command line effectively increases your productivity immensely.

Some bash resources:

- Excellent bash tutorial <http://tldp.org/HOWTO/Bash-Prog-Intro-HOWTO.html>
- Blog post on pipes for NGS <http://www.vincebuffalo.com/2013/08/08/the-mighty-named-pipe.html>
- Using bash and GNU parallel for NGS <http://bit.ly/gwbash>

### 3.4 (Optional exercise) Install `sickle` by yourself

Follow these steps only if you want to install `sickle` by yourself.

From the `sickle` project description: “`Sickle` is a tool that uses sliding windows along with quality and length thresholds to determine when quality is sufficiently low to trim the 3'-end of reads and also determines when the quality is sufficiently high enough to trim the 5'-end of reads. It will also discard reads based upon the length threshold.”

Installation procedures of research software often follow the same pattern, so it's useful to learn how to do this. Download the code, *compile* it and copy the binary to a location in your `$PATH`. The code for `sickle` is on <https://github.com/najoshi/sickle>. I prefer *compiling* my programs in `~/src` and then copying the resulting program to my `~/bin` directory, which is in my `$PATH`. This should get you a long way:

```
mkdir -p ~/src

# Go to the source directory and clone the sickle repository
cd ~/src
git clone https://github.com/najoshi/sickle
cd sickle

# Compile the program
make
```

```
# Create a bin directory
mkdir -p ~/bin
cp sickle ~/bin
```



---

## Quality Control

---

The first step of any sequencing project is to do quality control of your reads and remove (trim) low quality bases from the end of the read. In this exercise, you will work with Illumina data from the Human Microbiome Project that has already been trimmed. We still want to check the quality of reads, though.

In this part of the metagenomics workshop we will learn how to:

- Check the quality of your raw sequencing data
- Perform quality trimming using sickle

The workshop has the following exercises:

### 4.1 Quality Control with FastQC

In this exercise you will use [FastQC](#) to investigate the quality of your sequences using a nice graphical summary output.

#### 4.1.1 Retrieving your data

For the first step, make a workshop folder in your home directory and move into it:

```
mkdir -p ~/mg-workshop
cd ~/mg-workshop
```

Inside it, make a folder for your input files:

```
mkdir -p ~/mg-workshop/data/$SAMPLE/reads/1M/
cd ~/mg-workshop/data/$SAMPLE/
```

Now make a copy of the files you want to work on: gut, skin or teeth datasets. These files were originally taken from the [Human Microbiome Project](#) and then subsampled to include only 1 million reads each. You can copy these files from the project directory:

```
cp /proj/g2015028/nobackup/metagenomics-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.1.fastq ~/mg-w
cp /proj/g2015028/nobackup/metagenomics-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.2.fastq ~/mg-w
```

You will now have two files in your reads directory: one for the forward reads \*\_1.fastq and one for the reverse reads \*\_2.fastq.

### 4.1.2 FastQC

We will now use FastQC to generate a report about the quality of our sequencing reads. For most programs and scripts in this workshop, you can see their instructions by typing their name in the terminal followed by the flag `-h`. There are many options available, and we'll show you only a few of those.

First, make a folder to keep your quality control results:

```
mkdir -p ~/mg-workshop/results/quality_check/$SAMPLE/
```

Now, run fastqc for each file:

```
fastqc -o ~/mg-workshop/results/quality_check/$SAMPLE/ --nogroup ~/mg-workshop/data/$SAMPLE/reads/1M
```

FastQC will generate two files for each input file, one compressed, and one not. To view your files, copy the html results into your local computer and open them with a browser.

From **your own shell (not inside Uppmax - open a new terminal window)**:

```
mkdir -p ~/mg-workshop/  
cd ~/mg-workshop/  
scp -r username@milou.uppmx.uu.se:~/mg-workshop/results/quality_check/*/*html .
```

Instead of username, type your own username!

Now open the reports. Make sure you understand the results. Do they look ok? Is there a difference between forward and reverse? Are there any warnings or errors? What do they mean? Do you have adapter sequences in your reads? The FastQC project includes an ugly, but useful, [tutorial](#).

## 4.2 Optional: Quality trimming Illumina paired-end reads

In this exercise you will learn how to quality trim Illumina paired-end reads. Illumina paired-end reads are by far the most common Next Generation Sequencing (NGS) approach for metagenomics. The reads downloaded from the HMP are already quality trimmed. However, if you have time and want to try it out for yourself, you can run some more stringent quality trimming on them and see what happens.

### 4.2.1 Running sickle on a paired-end library

For quality trimming Illumina paired end reads we use the library sickle which trims reads from 5' end to 3' end using a sliding window. If the mean quality of bases inside a window drops below a specified number, the remaining of the read will be trimmed.

As a default, sickle trims a read at the point needed to maintain its average quality over 20. It also discards reads that are shorter than 20 bp. These are very good default values, but in this extra exercise you're welcome to change the values of these parameters using the `-q` and `-l` flags.

You can use the same qc directory as before for this step, since these reads won't be further processed.

Make sure you understand the input parameters and then run sickle,:

```
mkdir -p ~/mg-workshop/results/quality_check/sickle/$SAMPLE  
sickle pe \  
  -f ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.1.fastq \  
  -r ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.2.fastq \  
  -t sanger \  
  -o ~/mg-workshop/results/quality_check/sickle/$SAMPLE/qtrim.1.fastq \  
  -p ~/mg-workshop/results/quality_check/sickle/$SAMPLE/qtrim.2.fastq \  
  -l 20 -q 20
```



```
-s ~/mg-workshop/results/quality_check/sickle/$SAMPLE/qtrim.unpaired.fastq \  
-q 20 -l 20
```

Chek what files have been generated. Do you understand each of them?

**Question:** How many paired reads are left after trimming? How many singletons?

**Question:** What are the different quality scores that sickle can handle? Why do we specify `-t sanger` here?

### 4.2.2 Run FastQC again

We would like to see if sickle has done a good job. We do so by verifying the quality of the reads again with fastqc. Please refer to the FastQC exercise for instructions on how to do this.

**Question:** Does the quality improve much?

### 4.2.3 Trimming adapter sequence

To remove adapter sequences from your reads you can use [cutadapt](#). This is a crucial step to guarantee the quality of your assembly, but we'll skip that in this workshop.

At least a basic knowledge of how to work with the command line is required otherwise it will be very difficult to follow some of the examples. Have fun!



---

## Community analysis using 16S reads

---

In this part of the metagenomics workshop we will learn how to analyse the taxonomic composition of a sample using reads containing parts of 16S rRNA genes.

Continue to the following exercise:

### 5.1 Community analysis using rRNA gene reads

In this exercise we will analyse the taxonomic composition of your sample by utilising reads containing parts of 16S rRNA genes. Partial 16S rRNA genes will be extracted from the reads using the program [sortmeRNA](#) and these will subsequently be classified using the [RDP](#) classifier. Finally, the results will be visualised with the interactive program [KRONA](#).

#### 5.1.1 SortmeRNA

We will extract 16S rRNA encoding reads using [sortmeRNA](#) which is one of the fastest software for this. We start by making the necessary folders and assigning all necessary databases to a variable called DB:

```
mkdir -p ~/mg-workshop/results/phylogeny/16S/$SAMPLE
cd ~/mg-workshop/results/phylogeny/16S/$SAMPLE
ln -s ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.1.fastq reads.1.fastq
ln -s ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.2.fastq reads.2.fastq
DB="/proj/g2015028/nobackup/metagenomics-workshop/reference_db/sortmerna/fasta/silva-arc-16s-database"
```

SortMeRNA has built-in multithreading support that we will use for parallelization (-a). We still have to launch one sample at a time, though:

```
for readfile in reads.*.fastq;
do sortmerna --reads $readfile --ref $DB --fastx --aligned ${readfile}_rrna -v -a 2;
done
```

#### 5.1.2 RDP classifier

sortmeRNA outputs the reads, or part of reads, that encode rRNA in a fasta file. These rRNA sequences can be classified in many ways. One option is blasting them against a suitable database. Here we use a simple and fast method, the classifier tool at [RDP](#) (the Ribosomal Database Project). This uses a naïve bayesian classifier trained on kmer frequencies of many sequences of defined taxonomies. It gives bootstrap support values for each taxonomic level - usually, the support gets lower the further down the hierarchy you go. Genus level is the lowest level provided. You

can use the web service if you prefer, and upload each file individually, or you can use the uppmx installation of RDP classifier like this:

```
for file in *_rrna*.fastq;
do name=$(basename $file);
java -Xmx1g -jar /proj/g2015028/metagenomics/virtenv/rdp_classifier_2.6/dist/classifier.jar classify
done
```

### 5.1.3 Krona

To get a graphical representation of the taxonomic classifications you can use [Krona](#), which is an excellent program for exploring data with hierarchical structures in general. The output file is an html file that can be viewed in a browser. Again make a directory for Krona and run it, specifying the name of the output file (-o), the minimum bootstrap support to use (-m) and that the two input files should be treated as only one (-c):

```
ktImportRDP -o 16S.tax.html -m 50 -c reads.1.fastq_rrna.fastq.class.tsv reads.2.fastq_rrna.fastq.class.tsv
```

Copy the resulting file 16S.tax.html to your local computer with scp and open it a browser, like you did for the FastQC output.

**Question: What's the dominant type of organisms found in your sample?**

Have fun!

---

## Metagenomic Assembly

---

In this part of the metagenomics workshop we will learn how to:

- Perform assemblies with velvet

The part has the following exercise:

### 6.1 Assembling reads with Velvet

In this exercise you will learn how to perform an assembly with [Velvet](#). Velvet takes your reads as input and assembles them into contigs. It consists of two steps. In the first step, `velveth`, the de Bruijn graph is created. In the second one, the graph is traversed and contigs are created with `velvetg`. When constructing the de Bruijn graph, a *kmer* has to be specified. Reads are cut up into pieces of length *k*, each representing a node in the graph, edges represent an overlap (some de Bruijn graph assemblers do this differently, but the idea is the same). The advantage of using kmer overlap instead of read overlap is that the computational requirements grow with the number of unique kmers instead of unique reads. A more detailed explanation can be found in [this paper](#).

You can test different kmer lengths, as long as they're odd numbers. A good margin is to have the kmer length between 21 and 51. We'll then look at a few statistics on the assembly; if your choice of kmer wasn't good, you might have to run another assembly (but this is very fast).

#### 6.1.1 Pick your kmer

Fill in which value for *k* you want to do in the [Google doc](#). The value should be odd and somewhere in the range between maybe 19 and 99. Later we will compare the results from the different kmers for each group.

#### 6.1.2 velveth

Create the graph data structure with `velveth`. First create a directory with symbolic links to the pairs that you want to use:

```
mkdir -p ~/mg-workshop/results/assembly/$SAMPLE/  
cd ~/mg-workshop/results/assembly/$SAMPLE/  
ln -s ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.1.fastq pair1.fastq  
ln -s ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.2.fastq pair2.fastq
```

Create a directory for the kmer of your choice. **Replace N with the kmer length below:**

```
mkdir ${SAMPLE}_N
```

The reads need to be interleaved (forward and reverse read from the same fragment following each other in one file) for `velveth`. There are many tools available for performing this simple task. We'll be using one borrowed from `khmer`, but really anything will do:

```
interleave-reads.py -o pair.fasta pair1.fastq pair2.fastq
```

Run `velveth`, replacing `N` with the `kmer length` you chose:

```
velveth ${SAMPLE}_N N -fasta -shortPaired pair.fasta
```

Check what directories have been created:

```
ls
```

### 6.1.3 velvetg

To get the actual contigs you will have to run `velvetg` on the created graph. You can vary options such as the expected coverage and the coverage cut-off if you want, but we do not do that in this tutorial. We only choose not to do scaffolding. Again **replace `N` for your current kmer length**:

```
velvetg ${SAMPLE}_N -scaffolding no
```

### 6.1.4 assemstats

After the assembly, one wants to look at the length distributions of the resulting assemblies. We have written the `assemstats` script for that:

```
assemstats 100 ${SAMPLE}_N/contigs.fa
```

Try to find out what each of the stats represent by trying other cut-off values than 100. One of the most often used statistics in assembly length distribution comparisons is the *N50 length*, a weighted median of the length, where you weigh each contig by its length. This way, you assign more weight to larger contigs. Fifty per cent of all the bases in the assembly are contained in contigs shorter or equal to N50 length. Add your results to the [Google doc](#).

**Question: What are the important length statistics? Do we prefer sum over length? Should it be a combination?**

### 6.1.5 (Optional) Megahit

The `Megahit` is a recent improvement to assembly algorithms that can assemble large and complex metagenomes in an efficient manner. It runs on a single node and runs multiple values for `k` in a predefined or custom sequence. The default sequence is 21, 41, 61, 81 and 99. Here is how to run `megahit` for a specified list of `kmer` lengths, using up to 8 cores (threads) and maximum half the available memory on the node.

```
mkdir -p ~/mg-workshop/results/assembly/megahit/${SAMPLE}/
rm -rf ~/mg-workshop/results/assembly/megahit/${SAMPLE}/megahit_output
time megahit -l ~/mg-workshop/data/${SAMPLE}/reads/1M/${SAMPLE_ID}_1M.1.fastq \
  -2 ~/mg-workshop/data/${SAMPLE}/reads/1M/${SAMPLE_ID}_1M.2.fastq -t 8 -m 0.5 \
  -o ~/mg-workshop/results/assembly/megahit/${SAMPLE}/megahit_output/ --k-list 21,41,61,81,99
```

There is another [sheet\\_megahit](#) where you can add the `Megahit` assembly results.

**Question: How do `Megahit`'s results compare to those from `Velvet`? When would you choose one assembler over the other?**

### 6.1.6 (Optional) Ray

The [Ray](#) assembler was made to play well with metagenomics. Furthermore, it uses [MPI](#) to distribute the computation over multiple computational nodes and/or cores. You can run Ray on 8 cores with the command:

```
mkdir -p ~/mg-workshop/results/assembly/ray/$SAMPLE/
module unload intel
module load gcc openmpi/1.7.5
rm -rf ~/mg-workshop/results/assembly/ray/$SAMPLE/${SAMPLE}_N
time mpiexec -n 8 Ray -k N -p ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.{1,2}.fastq \
    -o ~/mg-workshop/results/assembly/ray/$SAMPLE/${SAMPLE}_N
module unload gcc
module load intel
```

Replace N again with your chosen kmer. There is another [sheet\\_ray](#) where you can add the Ray assembly results.

**Question: How do Ray's results compare to those from Velvet? When would you choose one assembler over the other?**

Have fun!





---

## Taxonomic Classification

---

- Taxonomic annotation of contigs using Phylosift

Contents:

### 7.1 Phylogenetic Classification using Phylosift

In this section, we will investigate our contigs with Phylosift, to investigate from which species they originate.

#### 7.1.1 Phylosift

Phylosift is software created for the purpose of determining the phylogenetic composition of metagenomic data. It uses a defined set of genes to predict the taxonomy of each sequence in your dataset. You can read more about how this works here: <http://phylosift.wordpress.com>. Let's prepare for the phylosift run in the usual way:

```
mkdir -p ~/mg-workshop/results/phylogeny/phylosift/$SAMPLE
cd ~/mg-workshop/results/phylogeny/phylosift/$SAMPLE
ln -s ~/mg-workshop/results/assembly/$SAMPLE/${SAMPLE}_N/contigs.fa .
```

Phylosift can be run using raw sequencing reads directly, but it excels at classifying contigs. Classification is performed in several sequential [steps](#), starting with a search for conserved marker genes in the data. We will run phylosift in parallel (8 cores):

```
phylosift search --threads 8 --debug --output phylosift_output contigs.fa > phylosift.log 2> phylosift.log
```

The progress of the run is stored in the file **phylosift\_output/run\_info.txt**, but if you want more detailed info you can have a look at **phylosift.log**.

When the run finishes you will have a directory called 'blastDir' inside the main phylosift output directory. Phylosift uses a program called [LAST](#), which is similar to BLAST but much faster, to identify sequences matching a set of marker genes among your sequences.

In the subsequent steps these sequences are aligned to reference marker alignments using **hmmalign** of the [HMMER suite](#). Alignments are then used to place each identified sequence into a phylogenetic reference tree using the program [pplacer](#). Finally, all placements are evaluated and summarized.

Notice that we give the `--continue` flag to phylosift, telling it to continue from the previous step in the analysis.:

```
phylosift align --threads 8 --debug --continue --output phylosift_output contigs.fa >> phylosift.log
```

When this step is complete the main output directory will contain an ‘alignDir’ with alignments from **hmmalign** and a ‘treeDir’ with placement files from **pplacer**.

*Unfortunately, the phylosift takes a fairly long time.* So if you can’t afford to wait for it, you can choose to copy the results from the project directory:

```
cp -r /proj/g2015028/nobackup/metagenomics-workshop/results/phylogeny/phylosift/$SAMPLE/phylosift_out
```

When all phylosift runs are completed (or copied), browse the output directory:

```
ls ~/mg-workshop/results/phylogeny/phylosift/$SAMPLE/phylosift_output/
```

All of these files are interesting, but the most fun one is the html file, so let’s download this to your own computer and have a look. **Again, switch to a terminal where you’re not logged in to UPPMAX:**

```
mkdir ~/mg-workshop/  
scp username@milou.uppmx.uu.se:~/mg-workshop/results/phylogeny/phylosift/phylosift_output/*.html ~/mg-workshop/
```

**Question: Compare with the taxonomic results from the 16S analysis. Do the results match? If not, what could be the explanation for the differences?**

---

## Functional Annotation

---

This part of the workshop has the following exercises:

1. Gene annotation pipeline - PROKKA
2. Predict metabolic pathways using MinPath
3. Quantify genes by mapping reads to the assembly
4. Explore gene annotation using KRONA

Contents:

### 8.1 Annotating the assembly using the PROKKA pipeline

Now that you have assembled the data into contigs the next natural step to do is annotation of the data, i.e. finding the genes and doing functional annotation of those. A range of programs are available for these tasks but here we will use **PROKKA**, which is essentially a **pipeline** comprising several open source bioinformatic tools and databases.

PROKKA automates the process of locating open reading frames (ORFs) and RNA regions on contigs, translating ORFs to protein sequences, searching for protein homologs and producing standard output files. For gene finding and translation, PROKKA makes use of the program **Prodigal**. Homology searching (via BLAST and HMMER) is then performed using the translated protein sequences as queries against a set of public databases (**CDD**, **PFAM**, **TIGRFAM**) as well as custom databases that come with PROKKA.

Set up the necessary files and run PROKKA, replacing **N** below with the kmer you chose for the assembly step.:

```
mkdir -p ~/mg-workshop/results/functional_annotation/prokka/  
cd ~/mg-workshop/results/functional_annotation/prokka/  
ln -s ~/mg-workshop/results/assembly/$SAMPLE/${SAMPLE}_N/contigs.fa .  
prokka contigs.fa --outdir $SAMPLE --norrna --notrna --metagenome --cpus 8  
cd $SAMPLE
```

PROKKA produces several types of output, such as:

- a **GFF** file, which is a standardised, tab-delimited, format for genome annotations
- a **Genbank (GBK)** file, which is a more detailed description of nucleotide sequences and the genes encoded in these.

When your dataset has been annotated you can view the annotations directly in the GFF file. PROKKA names the resulting files according to the current date like so: PROKKA\_mmddyyyy. So if you're following this workshop on a different date than 11242015 or running PROKKA on your own later on you will have to modify the following commands to match.

Now, take a look at the GFF file by doing:

```
less -S PROKKA_11242015.gff
```

You will notice that the first lines in the GFF file show the annotated sequence regions. To skip these and get directly to the annotations you can do:

```
grep -v "^#" PROKKA_11242015.gff | less -S
```

The caret (^) symbol tells grep to match at the beginning of each line and the '-v' flag means that these lines are skipped. The remaining lines are then piped to less.

**Question: How many coding regions were found by Prodigal? Hint: use grep -c to count lines**

Some genes in your dataset should now contain annotations from several databases, such as [enzyme commission](#) and [COG](#) (Clusters of Orthologous Groups) identifiers.

**Question: How many of the coding regions were given an enzyme identifier? How many were given a COG identifier?**

In the downstream analyses we will quantify and compare the abundance of enzymes and metabolic pathways, as well as COGs in the different samples. To do this, we will first extract lists of the genes with enzyme and COG IDs from the GFF file that was produced by PROKKA. First we find lines containing enzyme numbers using pattern matching with grep and then reformat the output using a combination of cut and sed

```
grep "eC_number=" PROKKA_11242015.gff | cut -f9 | cut -f1,2 -d ';' | sed 's/ID=//g' | sed 's/;eC_number=/'
```

Then we extract COG identifiers:

```
egrep "COG[0-9]{4}" PROKKA_11242015.gff | cut -f9 | sed 's/.\.+COG\([0-9]\+\);locus_tag=\(PROKKA_[0-9]'
```

**Make sure you understand what the different parts of these lines of code does. Try removing parts between the pipe (|) symbols and see how this changes the output.**

The COG table we will save for later. Next up is to predict pathways in the sample based on the enzymes annotated by PROKKA.

## 8.2 Predicting metabolic pathways using MinPath

Metabolic pathways are made up of enzymes that catalyze various reactions. Depending on how pathways are defined, they may contain any number of enzymes. A single enzyme may also be part of one or several pathways. One way of predicting metabolic pathways in a sample is to simply consider all the pathways that a set of enzymes are involved in. This may however overestimate pathways, for instance if only a few of the enzymes required for a pathway are annotated in the sample.

Here we will predict pathways using the program [MinPath](#) to get conservative estimate of the pathways present. MinPath only considers the minimum number of pathways required to explain the set of enzymes in the sample. As input, MinPath requires 1) a file with gene identifiers and enzyme numbers, separated by tabs, and 2) a file that links each enzyme to one or several pathways. The first of these we produced above using pattern matching from the PROKKA gff file. The second file exist in two versions, one that links enzymes to pathways defined in the Metacyc database and one that links enzymes to pathways defined in the KEGG database.

First we make sure that all the required files are available:

```
mkdir -p ~/mg-workshop/results/functional_annotation/minpath/$SAMPLE/
cd ~/mg-workshop/results/functional_annotation/minpath/$SAMPLE/
mkdir -p ~/mg-workshop/reference_db/
cp -r /proj/g2015028/nobackup/metagenomics-workshop/reference_db/cog ~/mg-workshop/reference_db/
```

```
cp -r /proj/g2015028/nobackup/metagenomics-workshop/reference_db/kegg ~/mg-workshop/reference_db/
cp -r /proj/g2015028/nobackup/metagenomics-workshop/reference_db/metacyc ~/mg-workshop/reference_db/
ln -s ~/mg-workshop/results/functional_annotation/prokka/$SAMPLE/PROKKA.$SAMPLE.ec
```

Run MinPath with this command to predict Metacyc pathways:

```
MinPath1.2.py -any PROKKA.$SAMPLE.ec -map ~/mg-workshop/reference_db/metacyc/ec.to.pwy --report PROKKA.$SAMPLE.ec
```

And to predict KEGG pathways:

```
MinPath1.2.py -any PROKKA.$SAMPLE.ec -map ~/mg-workshop/reference_db/kegg/ec.to.pwy --report PROKKA.$SAMPLE.ec
```

Take a look at the report files:

```
less -S PROKKA.$SAMPLE.metacyc.minpath
```

**Question: How many Metacyc and KEGG pathways did MinPath predict in your sample? How many were predicted if you had counted all possible pathways as being present? (HINT: look for the ‘naive’ and ‘minpath’ tags)**

## 8.3 Mapping reads and quantifying genes

### 8.3.1 Overview

So far we have only got the number of genes and annotations in the sample. Because these annotations are predicted from assembled reads we have lost the quantitative information for the annotations. So to actually **quantify** the genes, we will map the input reads back to the assembly.

There are many different mappers available to map your reads back to the assemblies. Usually they result in a [SAM or BAM file](#). Those are formats that contain the alignment information, where BAM is the binary version of the plain text SAM format. In this tutorial we will be using [bowtie2](#). You can also take a look at the [Bowtie2 documentation](#).

The SAM/BAM file can afterwards be processed with [Picard](#) to remove duplicate reads. Those are likely to be reads that come from a [PCR duplicate](#).

[BEDTools](#) can then be used to retrieve coverage statistics.

### 8.3.2 Mapping reads with bowtie2

First set up the files needed for mapping. **Replace ‘N’ with the kmer you used for the velet assembly:**

```
mkdir -p ~/mg-workshop/results/functional_annotation/mapping/$SAMPLE/
cd ~/mg-workshop/results/functional_annotation/mapping/$SAMPLE/
ln -s ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.1.fastq pair1.fastq
ln -s ~/mg-workshop/data/$SAMPLE/reads/1M/${SAMPLE_ID}_1M.2.fastq pair2.fastq
ln -s ~/mg-workshop/results/assembly/$SAMPLE/${SAMPLE}_N/contigs.fa
```

Then run the bowtie2-build program on your assembly:

```
bowtie2-build contigs.fa contigs.fa
```

**Question: What does bowtie2-build do? (Refer to the documentation)**

Next we run the actual mapping using [bowtie2](#):

```
bowtie2 -p 8 -x contigs.fa -1 pair1.fastq -2 pair2.fastq -S $SAMPLE.map.sam
```

The output SAM file needs to be converted to BAM format. For this we will use [samtools](#). First we create an index of the assembly for samtools:

```
samtools faidx contigs.fa
```

Then the SAM file is converted to BAM format ([view](#)), sorted by left most alignment coordinate ([sort](#)) and indexed ([index](#)) for fast random access in these steps:

```
samtools view -bt contigs.fa.fai $SAMPLE.map.sam > $SAMPLE.map.bam
samtools sort $SAMPLE.map.bam $SAMPLE.map.sorted
samtools index $SAMPLE.map.sorted.bam
```

### 8.3.3 Removing duplicates

We will now use **MarkDuplicates** from the Picard tool kit to identify and remove duplicates in the sorted and indexed BAM file:

```
java -Xms2g -Xmx32g -jar /sw/apps/bioinfo/picard/1.92/milou/MarkDuplicates.jar INPUT=$SAMPLE.map.sorted.bam \
METRICS_FILE=$SAMPLE.map.markdup.metrics AS=TRUE VALIDATION_STRINGENCY=LENIENT \
MAX_FILE_HANDLES_FOR_READ_ENDS_MAP=1000 REMOVE_DUPLICATES=TRUE
```

Picard's documentation also exists! Two bioinformatics programs in a row with decent documentation! Take a moment to celebrate, then take a look [at it](#).

**Question: Why not just remove all identical pairs instead of mapping them and then removing them?**

**Question: What is the difference between samtools rmdup and Picard MarkDuplicates?**

### 8.3.4 Calculating coverage

We have now mapped reads back to the assembly and have information on how much of the assembly that is covered by the reads in the sample. We are interested in the coverage of each of the genes annotated in the previous steps by the PROKKA pipeline. To extract this information from the BAM file we first need to define the regions to calculate coverage for. This we will do by creating a custom BED file defining the regions of interest (the PROKKA genes). Here we use an in-house bash script called [prokkagff2bed.sh](#) that searches for the gene regions in the PROKKA output and then prints them in a suitable format:

```
prokkagff2bed.sh ~/mg-workshop/results/functional_annotation/prokka/$SAMPLE/PROKKA_11242015.gff > $SAMPLE.map.bed
```

We then use [bedtools](#) to extract coverage information from the BAM file for the regions defined in the BED file we just created:

```
bedtools coverage -hist -abam $SAMPLE.map.markdup.bam -b $SAMPLE.map.bed > $SAMPLE.map.hist
```

**\*Note:** When using bedtools 2.24.0 or later, the *A* and the *B* files are switched as follows:

```
bedtools coverage -hist -a $SAMPLE.map.bed -b $SAMPLE.map.markdup.bam > $SAMPLE.map.hist
```

Have a look at the output file with `less` again. The final four columns give you the histogram i.e. coverage, number of bases with that coverage, length of the contig/feature/gene, bases with that coverage expressed as a ratio of the length of the contig/feature/gene. For each gene, we calculate coverage as  $c_{\text{gene}} = \text{sum}(\text{depth} * \text{fraction\_at\_depth})$ .

This calculation is performed using the in-house script [get\\_coverage\\_for\\_genes.py](#)

```
get_coverage_for_genes.py -i <(echo $SAMPLE.map.hist) > $SAMPLE.coverage
```

We now have coverage values for all genes predicted and annotated by the PROKKA pipeline. Next, we will use the annotations and coverage values to summarize annotations for the sample and produce interactive plots.

**Question: Coverage can also be calculated for each contig. Do you expect the coverage to differ for a contig and for the genes encoded on the contig? When might it be a good idea to calculate the latter?**

## 8.4 Summarize and explore the functional annotation

Now that we have annotated genes and quantified them in the sample using read mapping we will summarize and explore the annotations. This we will do by producing interactive plots detailing the proportion of functional categories such as metabolic pathways and orthologous gene families.

### 8.4.1 KRONA interactive plots

KRONA is a tool that takes as input a table of abundance values and several hierarchical categories and produces HTML files that can be explored interactively. The enzyme annotations from PROKKA are particularly suited for this purpose because these annotations can be grouped into higher functional categories such as pathways (e.g. glycolysis) and pathway classes (e.g. energy metabolism) for enzymes. Similarly, COG annotations can be summed up into higher categories such as “Carbohydrate transport and metabolism” and “Metabolism”.

First we will create a new directory for the krona output and link to the necessary files:

```
mkdir -p ~/mg-workshop/results/functional_annotation/krona/$SAMPLE
cd ~/mg-workshop/results/functional_annotation/krona/$SAMPLE/
ln -s ~/mg-workshop/results/functional_annotation/mapping/$SAMPLE/$SAMPLE.coverage
ln -s ~/mg-workshop/results/functional_annotation/prokka/$SAMPLE/PROKKA.$SAMPLE.ec
ln -s ~/mg-workshop/results/functional_annotation/prokka/$SAMPLE/PROKKA.$SAMPLE.cog
ln -s ~/mg-workshop/results/functional_annotation/minpath/$SAMPLE/PROKKA.$SAMPLE.kegg.minpath
ln -s ~/mg-workshop/results/functional_annotation/minpath/$SAMPLE/PROKKA.$SAMPLE.metacyc.minpath
```

Next, use the in-house `genes.to.kronaTable.py` script to produce the tabular output needed for KRONA.

For Metacyc pathways (from enzymes, only considering pathways predicted by MinPath):

```
genes.to.kronaTable.py -i PROKKA.$SAMPLE.ec -m ~/mg-workshop/reference_db/metacyc/ec.to.pwy -H ~/mg-wor
```

For KEGG pathways (from enzymes, only considering pathways predicted by MinPath):

```
genes.to.kronaTable.py -i PROKKA.$SAMPLE.ec -m ~/mg-workshop/reference_db/kegg/ec.to.pwy -H ~/mg-wor
```

For COG annotations:

```
genes.to.kronaTable.py -i PROKKA.$SAMPLE.cog -m ~/mg-workshop/reference_db/cog/cog.to.cat -H ~/mg-wor
```

Finally, use Kronatools `ktImportText` script to generate the HTML files:

```
ktImportText -o $SAMPLE.krona.metacyc.minpath.html $SAMPLE.krona.metacyc.minpath.tab
ktImportText -o $SAMPLE.krona.kegg.minpath.html $SAMPLE.krona.kegg.minpath.tab
ktImportText -o $SAMPLE.krona.COG.html $SAMPLE.krona.COG.tab
```

Copy the resulting html files to your local computer with scp as before and open it a browser, like you did for the FastQC output.

**Question: What are the main differences between the databases you have worked with: COG, Metacyc and KEGG? Which one do you prefer and why?**

**Question:** What are the main differences between the different samples (gut, skin and teeth)? Compare with results from other groups. Can you, for instance, find differences in degradation of compounds?

Enjoy!