

---

# **labibi Documentation**

***Release 5.0***

**C. Titus Brown**

August 14, 2014



<b>1</b>	<b>Dramatis personae</b>	<b>3</b>
<b>2</b>	<b>Papers and References</b>	<b>5</b>
2.1	Books . . . . .	5
2.2	RNAseq . . . . .	5
2.3	Computing and Data . . . . .	5
<b>3</b>	<b>Links</b>	<b>7</b>
3.1	Humor . . . . .	7
3.2	Resources . . . . .	7
3.3	Blogs . . . . .	7
<b>4</b>	<b>Complete table of contents</b>	<b>9</b>
4.1	Day 1 - Getting started with Amazon . . . . .	9
4.2	Day 2 – Running BLAST and other things at the command line . . . . .	18
4.3	Variant calling . . . . .	26
4.4	Assembling E. coli sequences with Velvet . . . . .	29
4.5	Interval Analysis and Visualization . . . . .	32
4.6	Running bedtools . . . . .	33
4.7	Understanding the SAM format . . . . .	33
4.8	R Tutorial for NGS2014 . . . . .	34
4.9	Control Flow . . . . .	55
4.10	ifelse() . . . . .	55
4.11	Other vectorized ways of control flow. . . . .	56
4.12	Simple loops . . . . .	56
4.13	for loop . . . . .	57
4.14	So for the for loop we would do the following: . . . . .	58
4.15	More avoiding loops . . . . .	58
4.16	The step above creates a vector of n NA's. They will be replaced sequentially with the random numbers as we generate them (using a function like the above one). . . . .	60
4.17	Variant calling and exploration of polymorphisms . . . . .	61
4.18	A complete de novo assembly and annotation protocol for mRNASeq . . . . .	64
4.19	Amazon Web Services instructions . . . . .	65
4.20	Instructor's Guide to ANGUS Materials . . . . .	86
4.21	Workshop Code of Conduct . . . . .	86



This is the schedule for the [2014 MSU NGS course](#).

This workshop has a [Workshop Code of Conduct](#).

[Assessment questionnaire](#)

Day	Schedule
Monday 8/4	<ul style="list-style-type: none"> <li>• 1:30pm lecture: Welcome! (Titus)</li> <li>• Tutorial: <i>Day 1 - Getting started with Amazon</i></li> <li>• 7pm: research presentations</li> </ul>
Tuesday 8/5	<ul style="list-style-type: none"> <li>• <i>Day 2 – Running BLAST and other things at the command line</i></li> <li>• 9:15am lecture: Sequencing considerations (Titus)</li> <li>• 10:30am: tutorial, <i>Running command-line BLAST</i> (Titus)</li> <li>• Afternoon: assessment</li> <li>• 1:15pm: tutorial, <i>Short Read Quality Control</i> (Elijah and Istvan)</li> <li>• Evening: <i>firepit social</i></li> </ul>
Wednesday 8/6	<ul style="list-style-type: none"> <li>• 9:15am lecture: Mapping and Assembly (Titus)</li> <li>• 10:30am: tutorial, <i>Variant calling</i> (Titus)</li> <li>• 1:15pm: <i>Understanding the SAM format</i> (Istvan)</li> <li>• 7:15pm: tutorial, UNIX command line (Elijah)</li> </ul>
Thursday 8/7	<ul style="list-style-type: none"> <li>• 9:15am lecture: Genomic Intervals (Istvan)</li> <li>• 10:30am mini-diversion: The Bioinformatics Skill System (Istvan)</li> <li>• 10:45am: tutorial, <i>Interval Analysis and Visualization</i> (Istvan)</li> <li>• 1:15pm: tutorial, <i>Assembling E. coli sequences with Velvet</i> (Titus)</li> <li>• 5:30pm: leave for Kalamazoo</li> </ul>
Friday 8/8	<ul style="list-style-type: none"> <li>• 9:15am-noon lecture/tutorial, <i>R Tutorial for NGS2014</i> R etc. (Ian Dworkin and Martin Schilling)</li> <li>• 1:15pm: tutorial, <i>Variant calling and exploration of polymorphisms</i></li> <li>• 1:15pm: lecture, more variant calling (Martin Schilling)</li> <li>• 7pm: lecture, Gene and genome annotation: PowerPoint   PDF (Daniel Standage)</li> </ul>
Saturday 8/9	<ul style="list-style-type: none"> <li>• 9:15am-noon: lecture/tutorial, <i>A complete de novo assembly and annotation protocol for mRNASeq</i> (Titus)</li> <li>• 1:15pm: lecture/discussion, mRNAseq assembly with Trinity (Meg Staton)</li> </ul>
Monday 8/11	<ul style="list-style-type: none"> <li>• 9:15am lecture, mRNAseq and counting PDF (Ian Dworkin)</li> <li>• 10:30am tutorial, <code>drosophila_rnaseq1</code> (Chris Chandler)</li> <li>• 10:45am tutorial, <code>mount_chris_snapshot</code> (Meg)</li> </ul>
2	<ul style="list-style-type: none"> <li>• 2:15pm <b>Contents</b> SOAPdeNovoTrans_count_eXpress (Matt)</li> <li>• 7:15pm tutorial, Mapping reads to transcriptomes</li> </ul>

---

## Dramatis personae

---

Instructors:

- Istvan Albert
- C Titus Brown
- Ian Dworkin

TAs:

- Amanda Charbonneau
- Elijah Lowe
- Will Pitchers
- Aswathy Sebastian
- Qingpeng Zhang

Lecturers:

- Chris Chandler
- Adina Chuang Howe
- Matt MacManes
- Martin Schilling
- Daniel Standage
- Meg Staton

He Who Drives Many Places:

- Cody Nicks





---

## Papers and References

---

### 2.1 Books

- [Practical Computing for Biologists](#)

This is a highly recommended book for people looking for a systematic presentation on shell scripting, programming, UNIX, etc.

### 2.2 RNAseq

- [Differential gene and transcript expression analysis of RNA-seq experiments with TopHat and Cufflinks](#), Trapnell et al., Nat. Protocols.
- One paper that outlines a pipeline with the tophat, cufflinks, cuffdiffs and some associated R scripts.
- [Statistical design and analysis of RNA sequencing data.](#), Auer and Doerge, Genetics, 2010.
  - [A comprehensive comparison of RNA-Seq-based transcriptome analysis from reads to differential gene expression and cross-comparison with microarrays: a case study in \*Saccharomyces cerevisiae\*](#). Nookaew et al., Nucleic Acids Res. 2012.
  - [Challenges and strategies in transcriptome assembly and differential gene expression quantification. A comprehensive in silico assessment of RNA-seq experiments](#) Vijay et al., 2012.
  - [Computational methods for transcriptome annotation and quantification using RNA-seq](#), Garber et al., Nat. Methods, 2011.
  - [Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments.](#), Bullard et al., 2010.
  - [A comparison of methods for differential expression analysis of RNA-seq data](#), Sonesson and Delorenzi, BMC Bioinformatics, 2013.
  - [Measurement of mRNA abundance using RNA-seq data: RPKM measure is inconsistent among samples.](#), Wagner et al., Theory Biosci, 2012. Also see [this blog post](#) explaining the paper in detail.

### 2.3 Computing and Data

- [A Quick Guide to Organizing Computational Biology Projects](#), Noble, PLoS Comp Biology, 2009.

- [Willingness to Share Research Data Is Related to the Strength of the Evidence and the Quality of Reporting of Statistical Results](#), Wicherts et al., PLoS One, 2011.
- [Got replicability?](#), McCullough, Economics in Practice, 2007.

Also see this great pair of blog posts on [organizing projects](#) and [research workflow](#).

## **3.1 Humor**

- [Data Sharing and Management Snafu in 3 Short Acts](#)

## **3.2 Resources**

- [Biostar](#)  
A high quality question & answer Web site.
- [SEQanswers](#)  
A discussion and information site for next-generation sequencing.
- [Software Carpentry lessons](#)  
A large number of open and reusable tutorials on the shell, programming, version control, etc.

## **3.3 Blogs**

- <http://www.genomesunzipped.org/>  
Genomes Unzipped.
- <http://ivory.idyll.org/blog/>  
Titus's blog.
- <http://bcbio.wordpress.com/>  
Blue Collar Bioinformatics
- <http://massgenomics.org/>  
Mass Genomics
- <http://blog.nextgenetics.net/>  
Next Genetics
- <http://gettinggeneticsdone.blogspot.com/>  
Getting Genetics Done

- <http://omicsomics.blogspot.com/>  
Omics! Omics!

---

## Complete table of contents

---

### 4.1 Day 1 - Getting started with Amazon

We're going to start by getting you set up on Amazon Web Services. For the duration of the course, we'll be running analyses on computers we rent from Amazon; this has a number of benefits that we'll discuss in the lecture.

#### 4.1.1 Start up an EC2 instance

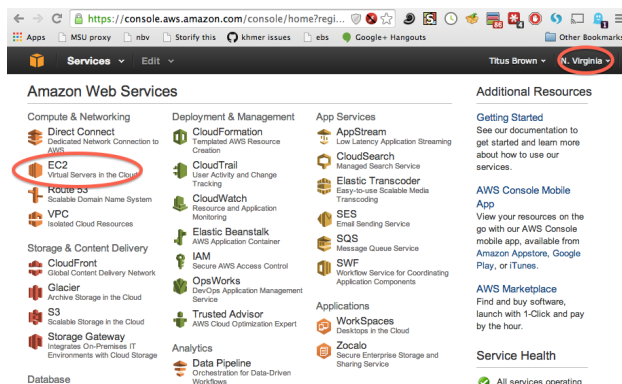
Here, we're going to startup an Amazon Web Services (AWS) Elastic Cloud Computing (EC2) "instance", or computer.

Go to '<https://aws.amazon.com>' in a Web browser.


Select 'My Account/Console' menu option 'AWS Management Console.'

Log in with your username & password.

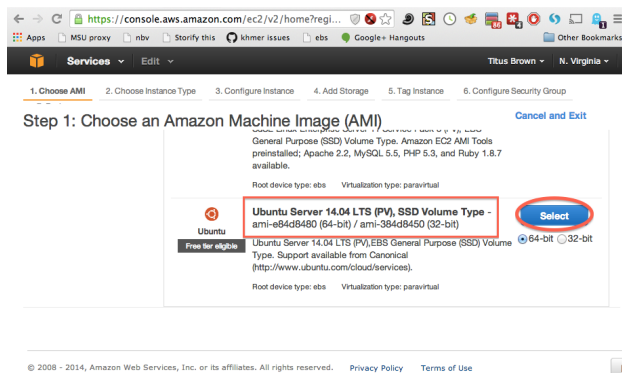
Make sure it says North Virginia in the upper right, then select EC2 (upper left).



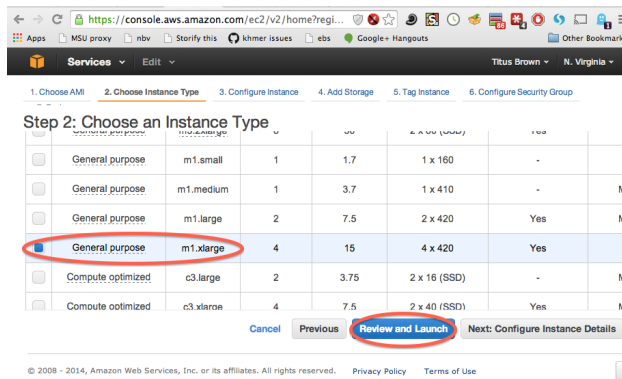
Select "Launch Instance" (midway down the page).




Next, scroll down the list of operating system types until you find Ubuntu 14.04 LTS (PV) – it should be at the very bottom. Click ‘select’. (See *Starting up a custom operating system* if you want to start up a custom operating system instead of Ubuntu 14.04.)



Scroll down the list of instance types until you find “m1.xlarge”. Select the box to the left, and then click “Review and Launch.”

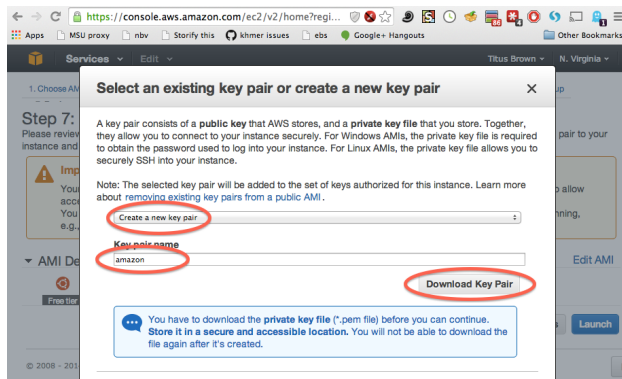


Ignore the warning, check that it says “Ubuntu 14.04 LTS (PV)”, and click “Launch”.

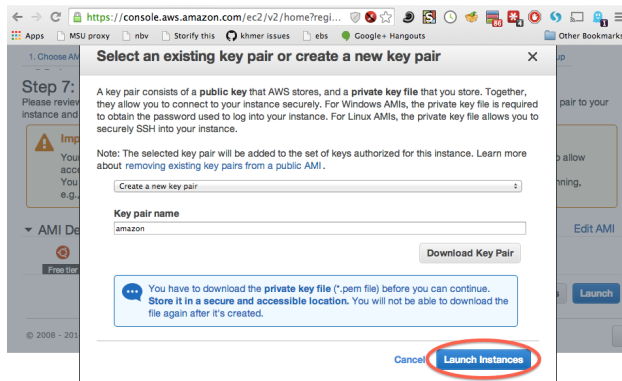


The *first* time through, you will have to “create a new key pair”, which you must then name (something like ‘amazon’) and download.

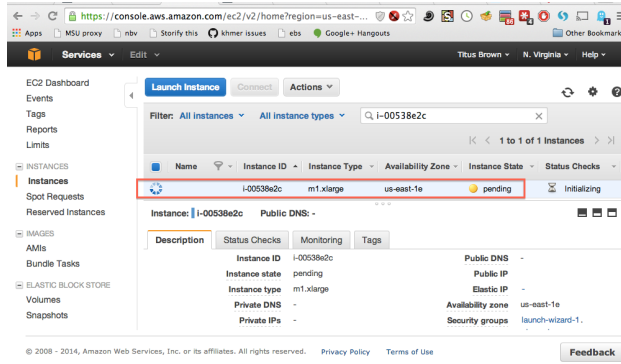
After this first time, you will be able to select an existing key pair.



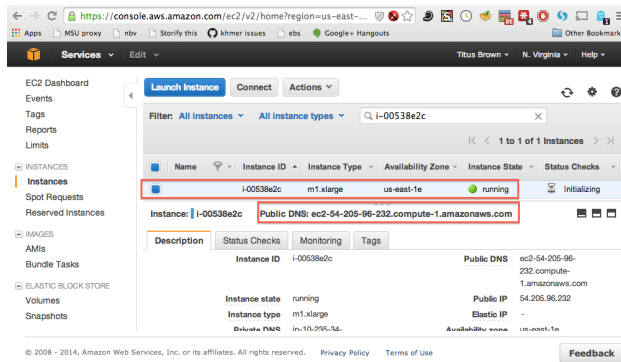
Select “Launch Instance.”



Select “view instance” and you should see a “pending” line in the menu.



Wait until it turns green, then make a note of the “Public DNS” (we suggest copying and pasting it into a text notepad somewhere). This is your machine name, which you will need for logging in.



Then, go to *Logging into your new instance “in the cloud” (Windows version)* or *Logging into your new instance “in the cloud” (Mac version)*

You might also want to read about *Terminating (shutting down) your EC2 instance*.

### 4.1.2 Logging into your new instance “in the cloud” (Windows version)


Download Putty and Puttygen from here: <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

#### Generate a ppk file from your pem file


(You only need to do this once!)

Open puttygen; select “Load”.

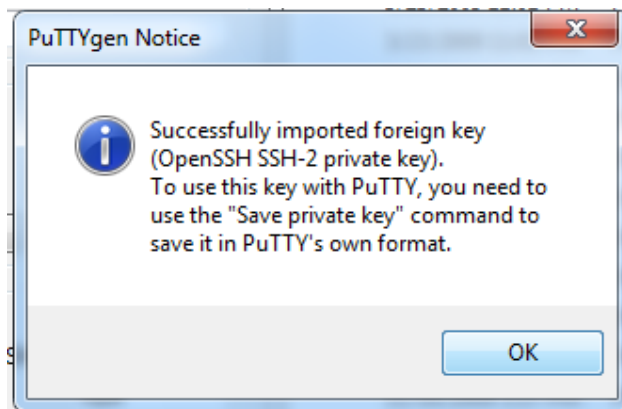




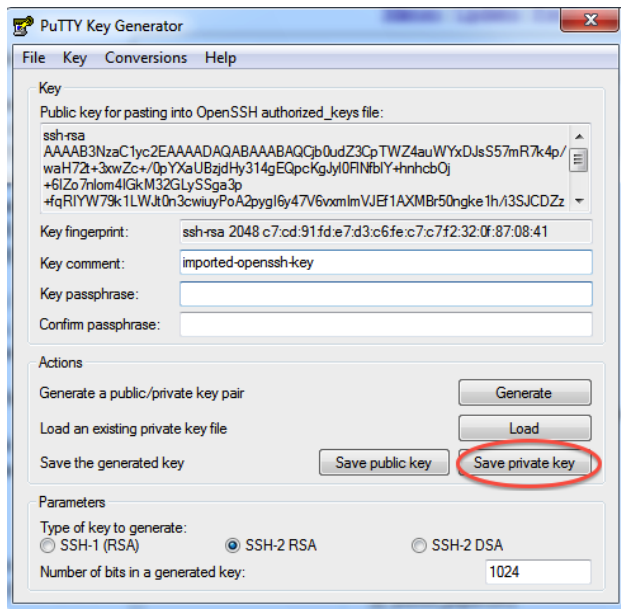
Find and load your '.pem' file; it's probably in your Downloads folder. Note, you have to select 'All files' on the bottom.



Load it.




Now, "save private key". Put it somewhere easy to find.

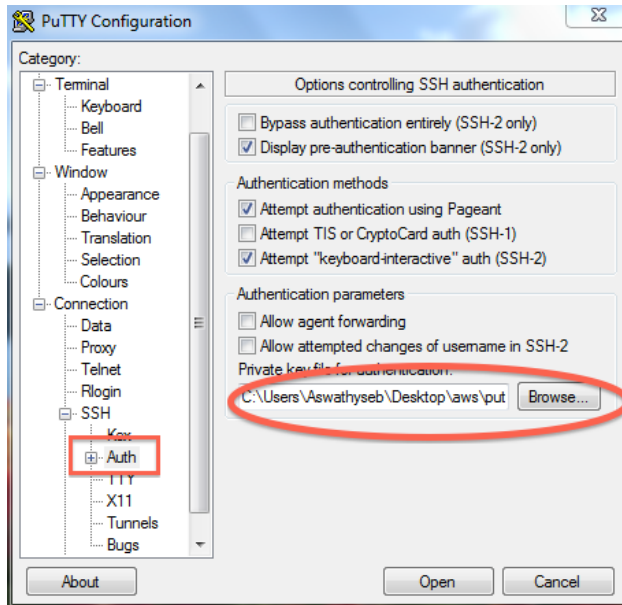


## Logging into your EC2 instance with Putty


Open up putty, and enter your hostname into the Host Name box.




Now, go find the 'SSH' section and enter your ppk file (generated above by puttygen). Then select 'Open'.



Log in as “ubuntu”.



Declare victory!



Here, you’re logging in as user ‘ubuntu’ to the machine ‘ec2-174-129-122-189.compute-1.amazonaws.com’ using the authentication key located in ‘amazon.pem’ on your Desktop.

You should now see a text line that starts with something like `ubuntu@ip-10-235-34-223:~$`. You’re in! Now type:

```
sudo bash
cd /root
```

to switch into superuser mode (see: <http://xkcd.com/149/>) and go to your home directory.

This is where the rest of the tutorials will start!

If you have Dropbox, you should now visit *Installing Dropbox on your EC2 machine*.

You might also want to read about *Terminating (shutting down) your EC2 instance*.

To log out, type:

```
exit
logout
```

or just close the window.

### 4.1.3 Logging into your new instance “in the cloud” (Mac version)

OK, so you’ve created a running computer. How do you get to it?

The main thing you’ll need is the network name of your new computer. To retrieve this, go to the instance view and click on the instance, and find the “Public DNS”. This is the public name of your computer on the Internet.

Copy this name, and connect to that computer with ssh under the username ‘root’, as follows.

First, find your private key file; it’s the .pem file you downloaded when starting up your EC2 instance. It should be in your Downloads folder. Move it onto your desktop and rename it to ‘amazon.pem’.

Next, start Terminal (in Applications... Utilities...) and type:

```
chmod og-rwx ~/Desktop/amazon.pem
```

to set the permissions on the private key file to “closed to all evildoers”.

Then type:

```
ssh -i ~/Desktop/amazon.pem ubuntu@ec2-???-???-???-???.compute-1.amazonaws.com
```

(but you have to replace the stuff after the ‘@’ sign with the name of the host).

Here, you’re logging in as user ‘ubuntu’ to the machine ‘ec2-174-129-122-189.compute-1.amazonaws.com’ using the authentication key located in ‘amazon.pem’ on your Desktop.

You should now see a text line that starts with something like ubuntu@ip-10-235-34-223:~\$. You’re in! Now type:

```
sudo bash
cd /root
```

to switch into superuser mode (see: <http://xkcd.com/149/>) and go to your home directory.

This is where the rest of the tutorials will start!

If you have Dropbox, you should now visit *Installing Dropbox on your EC2 machine*.

You might also want to read about *Terminating (shutting down) your EC2 instance*.

To log out, type:

```
exit
logout
```

or just close the window.

### 4.1.4 Installing Dropbox on your EC2 machine

**IMPORTANT:** Dropbox will sync everything you have to your EC2 machine, so if you are already using Dropbox for a lot of stuff, you might want to create a separate Dropbox account just for the course.

Start at the login prompt on your EC2 machine:

```
cd /root
```

If you can not do this(“Permission denied”), make sure you are in superuser mode. (You should see a text line that starts with something like

```
root@ip-10-235-34-223:~$. If not, use “sudo bash” to switch.)
```

Then, grab the latest dropbox installation package for Linux:

```
wget -O dropbox.tar.gz "http://www.dropbox.com/download/?plat=lnx.x86_64"
```

Unpack it:

```
tar -xvzf dropbox.tar.gz
```

Make the Dropbox directory on /mnt and link it in:

```
mkdir /mnt/Dropbox
ln -fs /mnt/Dropbox /root
```

and then run it, configuring it to put stuff in /mnt:

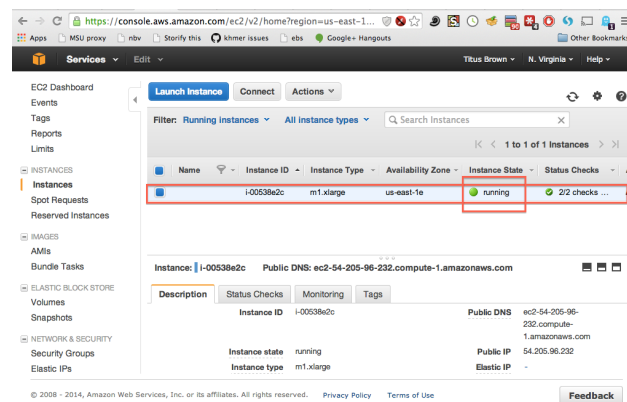
```
HOME=/mnt /root/.dropbox-dist/dropboxd &
```

When you get a message saying “this client is not linked to any account”, copy/paste the URL into browser and go log in. Voila!

Your directory ‘/root/Dropbox’, or, equivalently, ‘/mnt/Dropbox’, is now linked to your Dropbox account!

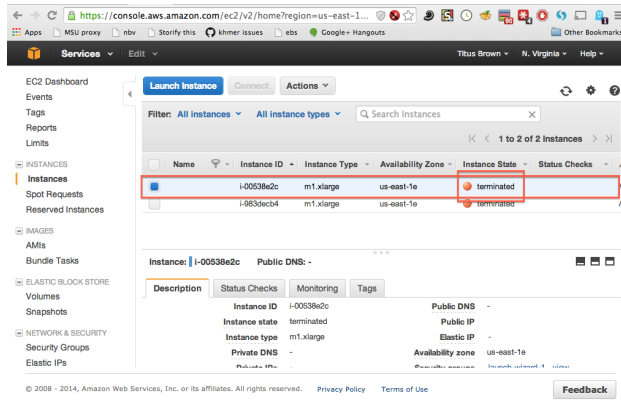
### 4.1.5 Terminating (shutting down) your EC2 instance

While your instance is running, Amazon will happily charge you on a per-hour basis – [check out the pricing](#) for more information. In general, you will want to shut down your instance when you’re done with it; to do that, go to your EC2 console and find your running instances (in green).



Then, select one or all of them, and go to the ‘Actions...’ menu, and then select ‘Terminate’. Agree.

After a minute or two, the console should show the instance as “terminated”.



## 4.2 Day 2 – Running BLAST and other things at the command line

Before following the procedures below, go through the process of starting up an ec2 instance and logging in – see [Day 1 - Getting started with Amazon](#) for details. Make sure you follow the Dropbox instructions, too!

The lecture will start at 9:15, the first tutorial ([Running command-line BLAST](#)) will start at 10:30, and the second tutorial will start at 1:30.

### 4.2.1 Running command-line BLAST

The goal of this tutorial is to run you through a demonstration of the command line, which you may not have seen or used much before.

Prepare for this tutorial by working through [Start up an EC2 instance](#), but follow the instructions to start up [Starting up a custom operating system](#) instead; use AMI ami-7606d01e.

All of the commands below can and should be copy/pasted rather than re-typed.

Note: on Windows using TeraTerm, you can select the commands in the Web browser, then go to TeraTerm and click your right mouse button to paste. On Mac OS X using Terminal, you can select the commands in the Web browser, use Command-C to copy, and then go the terminal and use Command-V to paste.

#### Switching to root

Start by making sure you’re the superuser, root:

```
sudo bash
```

#### Updating the software on the machine

Copy and paste the following two commands

```
apt-get update
apt-get -y install screen git curl gcc make g++ python-dev unzip \
    default-jre pkg-config libncurses5-dev r-base-core \
    r-cran-gplots python-matplotlib sysstat
```

(make sure to hit enter after the paste – sometimes the last line doesn’t paste completely.)

If you started up a custom operating system, then this should finish quickly; if instead you started up Ubuntu 14.04 blank, then this will take a minute or two.

## Install BLAST

Here, we’re using curl to download the BLAST distribution from NCBI; then we’re using ‘tar’ to unpack it into the current directory; and then we’re copying the program files into the directory /usr/local/bin, where we can run them from anywhere.

```
cd /root

curl -O ftp://ftp.ncbi.nih.gov/blast/executables/release/2.2.26/blast-2.2.26-x64-linux.tar.gz
tar xzf blast-2.2.26-x64-linux.tar.gz
cp blast-2.2.26/bin/* /usr/local/bin
cp -r blast-2.2.26/data /usr/local/blast-data
```

OK – now you can run BLAST from anywhere!

Again, this is basically what “installing software” means – it just means copying around files so that they can be run, and (in some cases) setting up resources so that the software knows where specific data files are.

## Running BLAST

Try typing:

```
blastall
```

You’ll get a long laundry list of output, with all sorts of options and arguments. Let’s play with some of them.

First! We need some data. Let’s grab the mouse and zebrafish RefSeq protein data sets from NCBI, and put them in /mnt, which is the scratch disk space for Amazon machines

```
cd /mnt

curl -O ftp://ftp.ncbi.nih.gov/refseq/M_musculus/mRNA_Prot/mouse.protein.faa.gz
curl -O ftp://ftp.ncbi.nih.gov/refseq/D_rerio/mRNA_Prot/zebrafish.protein.faa.gz
```

If you look at the files in the current directory, you should see both files, along with a directory called lost+found which is for system information:

```
ls -l
```

should show you:

```
drwx----- 2 root root    16384 2013-01-08 00:14 lost+found
-rw-r--r--  1 root root 9454271 2013-06-11 02:29 mouse.protein.faa.gz
-rw-r--r--  1 root root 8958096 2013-06-11 02:29 zebrafish.protein.faa.gz
```

Both of these files are FASTA protein files (that’s what the .faa suggests) that are compressed by gzip (that’s what the .gz suggests).

Uncompress them

```
gunzip *.faa.gz
```

and let’s look at the first few sequences:

```
head -11 mouse.protein.faa
```

These are protein sequences in FASTA format. FASTA format is something many of you have probably seen in one form or another – it’s pretty ubiquitous. It’s just a text file, containing records; each record starts with a line beginning with a ‘>’, and then contains one or more lines of sequence text.

Let’s take those first two sequences and save them to a file. We’ll do this using output redirection with ‘>’, which says “take all the output and put it into this file here.”

```
head -11 mouse.protein.faa > mm-first.fa
```

So now, for example, you can do ‘cat mm-first.fa’ to see the contents of that file (or ‘less mm-first.fa’).

Now let’s BLAST these two sequences against the entire zebrafish protein data set. First, we need to tell BLAST that the zebrafish sequences are (a) a database, and (b) a protein database. That’s done by calling ‘formatdb’

```
formatdb -i zebrafish.protein.faa -o T -p T
```

Next, we call BLAST to do the search

```
blastall -i mm-first.fa -d zebrafish.protein.faa -p blastp
```

This should run pretty quickly, but you’re going to get a LOT of output!! What’s going on? A few things –

- if you BLAST a sequence against a large database, odds are it will turn up a lot of spurious matches. By default, blastall uses an e-value cutoff of 10, which is very relaxed.
- blastall also reports the first 100 matches, which is usually more than you want.
- a lot of proteins also have trace similarity to other proteins!

For all of these reasons, generally you only want the first few BLAST matches, and/or the ones with a “good” e-value. We do that by adding ‘-b 2 -v 2’ (which says, report only two matches and alignments); and by adding ‘-e 1e-6’, which says, report only matches with an e-value of 1e-6 or better

```
blastall -i mm-first.fa -d zebrafish.protein.faa -p blastp -b 2 -v 2 -e 1e-6
```

Now you should get a lot less text! (And indeed you do...) Let’s put it in an output file, ‘out.txt’

```
blastall -i mm-first.fa -d zebrafish.protein.faa -p blastp -b 2 -v 2 -o out.txt
```

The contents of the output file should look exactly like the output before you saved it into the file – check it out:

```
cat out.txt
```

## Converting BLAST output into CSV

Suppose we wanted to do something with all this BLAST output. Generally, that’s the case - you want to retrieve all matches, or do a reciprocal BLAST, or something.

As with most programs that run on UNIX, the text output is in some specific format. If the program is popular enough, there will be one or more parsers written for that format – these are just utilities written to help you retrieve whatever information you are interested in from the output.

Let’s conclude this tutorial by converting the BLAST output in out.txt into a spreadsheet format, using a Python script. (We’re not doing this just to confuse you; this is really how we do things around here.)

First, we need to get the script. We’ll do that using the ‘git’ program

```
git clone https://github.com/ngs-docs/ngs-scripts.git /root/ngs-scripts
```



We'll discuss 'git' more later; for now, just think of it as a way to get ahold of a particular set of files. In this case, we've placed the files in `/root/ngs-scripts/`, and you're looking to run the script `blast/blast-to-csv.py` using Python

```
python /root/ngs-scripts/blast/blast-to-csv.py out.txt
```

This outputs a spread-sheet like list of names and e-values. To save this to a file, do:

```
python /root/ngs-scripts/blast/blast-to-csv.py out.txt > /root/Dropbox/out.csv
```

The end file, 'out.csv', should soon be in your Dropbox on your local computer. If you have Excel installed, try double clicking on it.

---

And that's the kind of basic workflow we'll be teaching you:

1. Download program
2. Download data
3. Run program on data
4. Look at results

...but in many cases more complicated :).

---

Note that there's no limit on the number of sequences you BLAST, etc. It's just sheer compute speed and disk space that you need to worry about, and if you look at the files, it turns out they're not that big – so it's mostly your time and energy.

This will also maybe help you understand why UNIX programs are so powerful – each program comes with several, or several dozen, little command line “flags” (parameters), that help control how it does its work; then the output is fed into another such program, etc. The possibilities are literally combinatorial.

---

We're running a Python program 'blast-to-csv.py' above – if you're interested in what the Python program does, take a look at the source code:

<https://github.com/ngs-docs/ngs-scripts/blob/master/blast/blast-to-csv.py>

## Summing up

Command-line BLAST lets you do BLAST searches of any sequences you have, quickly and easily. It's probably the single most useful skill a biologist can learn if they're doing anything genomics-y ;).

Its main computational drawback is that it's not fast enough to deal with some of the truly massive databases we now have, but that's generally not a problem for individual users. That's because they just run it and “walk away” until it's done!

The main practical issues you will confront in making use of BLAST:

- getting your sequence(s) into the right place.
  - formatting the database.
  - configuring the BLAST parameters properly.
  - doing what you want after BLAST!
- 

Other questions to ponder:

- if we're using a pre-configured operating system, why did we have to install BLAST?

## 4.2.2 Short Read Quality Control

As useful as BLAST is, we really want to get into sequencing data, right? One of the first steps you must do with your data is evaluate its quality and try to improve it.

**Summary:** a sequencing run may contain data of low reliability. It may also contain various contaminants and artificial sequence fragments. Some (but not all) of these problem can be corrected.

**Caution:** Don't apply any type of correction without evaluating the results it produces. In general it is best to be conservative with QC. We are altering the data based on our expectations of what it should be like! The process may also introduce its own biases into the dataset.

### Biostar QoD (questions of the day)

QC is one of the most *mis-under-estimated* tasks of NGS data analysis. People assume there is very little to it once they know how to run the tool. The reality is a more complicated than that.

QC also happens to be a pet peeve of mine (Istvan) as demonstrated below in the following Biostar threads (and others):

1. [FastQC quality control shootout](#)
2. [So What Does The Sequence Duplication Rate Really Mean In A Fastqc Report](#)
3. [Revisiting the FastQC read duplication report](#)

### Quick Start

The first part of this tutorial will run on your own computer. It assumes that you have Java installed. Download both FastQC and two smaller datasets onto your system

1. <http://www.bioinformatics.babraham.ac.uk/projects/fastqc/>
2. [http://apollo.huck.psu.edu/data/SRR519926\\_1.fastq.zip](http://apollo.huck.psu.edu/data/SRR519926_1.fastq.zip)
3. [http://apollo.huck.psu.edu/data/SRR447649\\_1.fastq.zip](http://apollo.huck.psu.edu/data/SRR447649_1.fastq.zip)

Run FastQC on each the files from the graphical interface. Let's discuss the output in more detail.

### FastQC at the command line

Before you can do that, though, you need to install a bunch o' software.

We will use a so called `Makefile`, a simple text file that can contains a series of commands that you could otherwise type in yourself. There will be more information on shell programming and automation later. For now think of a `Makefile` as a simple way to pick which commands you can execute yourself. Let's get started. Install `make`:

```
sudo apt-get install make -y
```

You can also investigate the `Makefile` yourself: <https://github.com/ngs-docs/angus/blob/2014/files/Makefile-short-read-quality>

This tutorial will download datasets. You may want to create a directory (folder) that stores this data:

```
mkdir qc
cd qc
```

We assume that you are running the subsequent scripts from this folder.

---

**Important:** Obtain the *Makefile* and save it onto your cloud system:

```
# This is where you get the Makefile
wget https://raw.githubusercontent.com/ngs-docs/angus/2014/files/Makefile-short-read-quality -O Makefile
```

---

You can investigate the file:

```
# Look at the Makefile
# more Makefile or pico Makefile
```

So we now have a *Makefile* and our system can execute this *Makefile* via the *make* command.:

```
make
```

## Setup

In our case you have to always specify which section of the *Makefile* do you wish to execute. For example you can type:

```
make setup
```

This will execute the parts of the *Makefile* that is listed below:

```
#
# Run initial setup tasks
#

# This directory will contain the executables
mkdir -p ~/bin

# Add the ~/bin to the user path
echo 'export PATH=$PATH:~/bin' >> ~/.bashrc

# Install the unzip library
sudo apt-get install -y unzip

# Change the mount point so it is user writeable
sudo chmod 777 /mnt

# Update the installation sources
sudo apt-get update

# Install java
sudo apt-get install -y default-jdk
```

Note that you could also just type in these commands yourself for the same results. The *Makefile* just automates this.

## Software Install

The next step is installing FastQC and Trimmomatic on your instance:

```
make install
```

command will execute the following lines.

```
#
# The src folder will contain the downloaded software
#
mkdir -p ~/src

#
# The bin folder will contain the binaries that you have downloaded
#
mkdir -p ~/bin

#
# Install Trimmomatic
#
curl http://www.usadellab.org/cms/uploads/supplementary/Trimmomatic/Trimmomatic-0.32.zip -o ~/src/Trimmomatic-0.32.zip
unzip -o ~/src/Trimmomatic-0.32.zip -d ~/src
ln -fs ~/src/Trimmomatic-0.32/trimmomatic-0.32.jar ~/bin

#
# Install FastQC
#
curl http://www.bioinformatics.babraham.ac.uk/projects/fastqc/fastqc_v0.11.2.zip -o ~/src/fastqc_v0.11.2.zip
unzip -o ~/src/fastqc_v0.11.2.zip -d ~/src
chmod +x ~/src/FastQC/fastqc
ln -fs ~/src/FastQC/fastqc ~/bin/fastqc
```

Where did stuff go? The downloaded code went into `~/src` the binaries are linked into `~/bin` To test that everything works well type:

```
fastqc -v
```

This will print the version number of *fastqc*

## Download Data

Start at your EC2 prompt, then type

```
make download
```

This will execute the following lines. Remember that you could also type these in yourself.

```
#
# Data download
#
curl apollo.huck.psu.edu/data/SRR.tar.gz -o SRR.tar.gz

# Unpack the gzipped data
tar xzvf SRR.tar.gz
```

## The FASTQ Format

In class explanation of the format. See a good description at [http://en.wikipedia.org/wiki/FASTQ\\_format](http://en.wikipedia.org/wiki/FASTQ_format)

If you don't understand the format, you don't understand the basic premise of the data generation!

Run a FastQC analysis on each dataset:

```
make fastqc
```

would run the commands:

```
#
# Run FastQC on every file that we have downloaded.
#
fastqc *.fastq
```

This command will generate an HTML file for each file. Copy these files to your dropbox and look at them (a short walkthrough on what each plot means).

Alternatively you can generate the fastqc output directly to your Dropbox like so:

```
fastqc *.fastq -o /mnt/Dropbox
```

## Pattern Matching

We can also investigate what the files contain by matching:

```
# Find start codons
grep ATG SRR519926_1.fastq --color=always | head

# Find a subsequence
grep AGATCGGAAG SRR519926_1.fastq --color=always | head
```

Pattern matching via expressions is an extremely powerful concept. We'll revisit them later.

## Trimming

Note that there are vary large number of tools that perform quality/adaptor trimming.

Now, run [Trimmomatic](#) to eliminate Illumina adapters from your sequences. First we need to find the adapter sequences:

```
ln -s ~/src/Trimmomatic-0.32/adapters/TruSeq3-SE.fa
```

---

**Tip:** You may also want to shorten the command line like so:

```
alias trim='java -jar ~/src/Trimmomatic-0.32/trimmomatic-0.32.jar'
```

You can now invoke the tool just by typing:

```
trim
```

---

Among the (many) agonizing decisions that you will have to make is what parameters to pick: how big should be my window be, how long should the reads be, what should be the average quality be? What kinds of contaminants do I have. Run, rerun and evaluate. Err on the side of caution.

Trim by quality alone:

```
#
# Run the quality trimming.
#
java -jar ~/src/Trimmomatic-0.32/trimmomatic-0.32.jar SE SRR447649_1.fastq good.fq SLIDINGWIN
fastqc good.fq -o /mnt/Dropbox
```

Quality and clipping:

```
#
# Run quality trimming and clipping at the same time.
#
java -jar ~/src/Trimmomatic-0.32/trimmomatic-0.32.jar PE SRR519926_1.fastq good1.fq bad1.fq
fastqc good1.fq -o /mnt/Dropbox
```

Now a homework:

---

**Note:** Read the manual for [Trimmomatic](#). Trim the reads in parallel for both readfiles in a sample.

---

**Note:** BTW: cool kids have pretty prompts, but you too can be cool, all you need to do is:

```
echo "export PS1='\[\e]0;\w\a\]\n\[\e[32m\]\u@h \[\e[33m\]\w\[\e[0m\]\n\$ '" >> ~/.bashrc
```

Then relog. Don't ask why this works, it is one of those things that is best left undisturbed.

---

## 4.3 Variant calling

The goal of this tutorial is to show you the basics of variant calling using [Samtools](#).

We'll be using data from one of Rich Lenski's LTEE papers, the one on [the evolution of citrate consumption in LTEE](#).

### 4.3.1 Booting an Amazon AMI

Start up an Amazon computer (m1.large or m1.xlarge) using AMI ami-7607d01e (see [Start up an EC2 instance](#) and [Starting up a custom operating system](#)).

Log in with Windows or from Mac OS X.

### 4.3.2 Updating the operating system

Copy and paste the following two commands

```
apt-get update
apt-get -y install screen git curl gcc make g++ python-dev unzip \
    default-jre pkg-config libncurses5-dev r-base-core \
    r-cran-gplots python-matplotlib sysstat
```

to update the computer with all the bundled software you'll need.

### 4.3.3 Install software

First, we need to install the [BWA aligner](#):

```
cd /root
wget -O bwa-0.7.10.tar.bz2 http://sourceforge.net/projects/bio-bwa/files/bwa-0.7.10.tar.bz2/download
tar xvfj bwa-0.7.10.tar.bz2
cd bwa-0.7.10
make
```

```
cp bwa /usr/local/bin
```

Also install samtools:

```
apt-get -y install samtools
```

#### 4.3.4 Download data

Download the reference genome and the resequencing reads:

```
cd /mnt
```

```
curl -O http://athyra.idyll.org/~t/REL606.fa.gz  
gunzip REL606.fa.gz
```

```
curl -O ftp://ftp.sra.ebi.ac.uk/vol1/fastq/SRR098/SRR098038/SRR098038.fastq.gz
```

Note, this last URL is the “Fastq files (FTP)” link from the European Nucleotide Archive (ENA) for this sample: <http://www.ebi.ac.uk/ena/data/view/SRR098042>.

#### 4.3.5 Do the mapping

Now let’s map all of the reads to the reference. Start by indexing the reference genome:

```
cd /mnt
```

```
bwa index REL606.fa
```

Now, do the mapping of the raw reads to the reference genome:

```
bwa aln REL606.fa SRR098038.fastq.gz > SRR098038.sai
```

Make a SAM file (this would be done with ‘sampe’ if these were paired-end reads):

```
bwa samse REL606.fa SRR098038.sai SRR098038.fastq.gz > SRR098038.sam
```

This file contains all of the information about where each read hits on the reference.

Next, index the reference genome with samtools:

```
samtools faidx REL606.fa
```

Convert the SAM into a BAM file:

```
samtools import REL606.fa.fai SRR098038.sam SRR098038.bam
```

Sort the BAM file:

```
samtools sort SRR098038.bam SRR098038.sorted
```

And index the sorted BAM file:

```
samtools index SRR098038.sorted.bam
```

### 4.3.6 Visualizing alignments

At this point you can visualize with samtools tview or [Tablet](#).

‘samtools tview’ is a text interface that you use from the command line; run it like so:

```
samtools tview SRR098038.sorted.bam REL606.fa
```

The ‘.’s are places where the reads align perfectly in the forward direction, and the ‘,’s are places where the reads align perfectly in the reverse direction. Mismatches are indicated as A, T, C, G, etc.

You can scroll around using left and right arrows; to go to a specific coordinate, use ‘g’ and then type in the contig name and the position. For example, type ‘g’ and then ‘rel606:553093<ENTER>’ to go to position 553093 in the BAM file.

Use ‘q’ to quit.

For the [Tablet viewer](#), click on the link and get it installed on your local computer. Then, start it up as an application. To open your alignments in Tablet, you’ll need three files on your local computer: REL606.fa, SRR098042.sorted.bam, and SRR098042.sorted.bam.bai. You can copy them over using Dropbox, for example.

### 4.3.7 Counting alignments

This command:

```
samtools view -c -f 4 SRR098038.bam
```

will count how many reads DID NOT align to the reference (214518).

This command:

```
samtools view -c -F 4 SRR098038.bam
```

will count how many reads DID align to the reference (6832113).

And this command:

```
gunzip -c SRR098038.fastq.gz | wc
```

will tell you how many lines there are in the FASTQ file (28186524). Reminder: there are four lines for each sequence.

### 4.3.8 Calling SNPs

You can use samtools to call SNPs like so:

```
samtools mpileup -uD -f REL606.fa SRR098038.sorted.bam | bcftools view -bvcg - > SRR098038.raw.bcf
```

(See the ‘mpileup’ docs [here](#).)

Now convert the BCF into VCF:

```
bcftools view SRR098038.raw.bcf > SRR098038.vcf
```

You can check out the VCF file by using ‘tail’ to look at the bottom:

```
tail *.vcf
```



Each variant call line consists of the chromosome name (for E. coli REL606, there's only one chromosome - rel606); the position within the reference; an ID (here always '.'); the reference call; the variant call; and a bunch of additional information about

Again, you can use 'samtools tview' and then type (for example) 'g' 'rel606:4616538' to go visit one of the positions. The format for the address to go to with 'g' is 'chr:position'.

You can read more about [the VCF file format here](#).

### 4.3.9 Questions/discussion items

Why so many steps?

## 4.4 Assembling E. coli sequences with Velvet

The goal of this tutorial is to show you the basics of assembly using [the Velvet assembler](#).

We'll be using data from [Efficient de novo assembly of single-cell bacterial genomes from short-read data sets](#), Chitsaz et al., 2011.

### 4.4.1 Booting an Amazon AMI

Start up an Amazon computer (m1.large or m1.xlarge) using AMI ami-7607d01e (see [Start up an EC2 instance](#) and [Starting up a custom operating system](#)).

Log in with Windows or from Mac OS X.

#### Logging in

Log in and type:

```
sudo bash
```

to change into superuser mode.

### 4.4.2 Updating the operating system

Copy and paste the following two commands

```
apt-get update
apt-get -y install screen git curl gcc make g++ python-dev unzip \
    default-jre pkg-config libncurses5-dev r-base-core \
    r-cran-gplots python-matplotlib sysstat
```

to update the computer with all the bundled software you'll need.

#### Packages to install

Install [khmer](#):

```
cd /usr/local/share
git clone https://github.com/ged-lab/khmer.git
cd khmer
git checkout v1.1
make install
```

and install the Velvet assembler:

```
cd /root
curl -O http://www.ebi.ac.uk/~zerbino/velvet/velvet_1.2.10.tgz
tar xzf velvet_1.2.10.tgz
cd velvet_1.2.10
make MAXKMERLENGTH=51
cp velvet? /usr/local/bin
```

as well as [Quast](#), software for evaluating the assembly against the known reference:

```
cd /root
curl -O -L https://downloads.sourceforge.net/project/quast/quast-2.3.tar.gz
tar xzf quast-2.3.tar.gz
```

## Getting the data

Now, let's create a working directory:

```
cd /mnt
mkdir assembly
cd assembly
```

Download some E. coli data. The first data set (ecoli\_ref-5m-trim.fastq.gz) is the trimmed PE data sets from the other day (see [Short Read Quality Control](#)), and the second data set is a specially processed data set using [digital normalization](#) that will assemble quickly.

```
curl -O https://s3.amazonaws.com/public.ged.msu.edu/ecoli_ref-5m-trim.fastq.gz
curl -O https://s3.amazonaws.com/public.ged.msu.edu/ecoli-reads-5m-dn-paired.fa.gz
```

## Running an assembly

Now... assemble the small, fast data sets, using the Velvet assembler. Here we will set the required parameter k=21:

```
velveth ecoli.21 21 -shortPaired -fasta.gz ecoli-reads-5m-dn-paired.fa.gz
velvetg ecoli.21 -exp_cov auto
```

Check out the stats for the assembled contigs for a cutoff of 1000:

```
python /usr/local/share/khmer/sandbox/assemstats3.py 1000 ecoli.*/contigs.fa
```

Also try assembling with k=23 and k=25:

```
velveth ecoli.23 23 -shortPaired -fasta.gz ecoli-reads-5m-dn-paired.fa.gz
velvetg ecoli.23 -exp_cov auto
```

```
velveth ecoli.25 25 -shortPaired -fasta.gz ecoli-reads-5m-dn-paired.fa.gz
velvetg ecoli.25 -exp_cov auto
```

Now check out the stats for the assembled contigs for a cutoff of 1000:

```
python /usr/local/share/khmer/sandbox/assemstats3.py 1000 ecoli.*/contigs.fa
```

(Also read: [What does k control in de Bruijn graph assemblers?](#).)

## Comparing and evaluating assemblies - QUAST

Download the true reference genome:

```
cd /mnt/assembly
curl -O https://s3.amazonaws.com/public.ged.msu.edu/ecoliMG1655.fa.gz
gunzip ecoliMG1655.fa.gz
```

and run QUAST:

```
/root/quast-2.3/quast.py -R ecoliMG1655.fa ecoli.*/contigs.fa
```

Note that here we're looking at *all* the assemblies we've generated.

Now look at the results:

```
more quast_results/latest/report.txt
```

The first bits to look at are Genome fraction (%) and # misassembled contigs, I think.

## Searching assemblies – BLAST

Install BLAST:

```
cd /root

curl -O ftp://ftp.ncbi.nih.gov/blast/executables/release/2.2.24/blast-2.2.24-x64-linux.tar.gz
tar xzf blast-2.2.24-x64-linux.tar.gz
cp blast-2.2.24/bin/* /usr/local/bin
cp -r blast-2.2.24/data /usr/local/blast-data
```

Build BLAST databases for the assemblies you've done:

```
cd /mnt/assembly

for i in 21 23 25
do
    extract-long-sequences.py -o ecoli-$i.fa -l 500 ecoli.$i/contigs.fa
    formatdb -i ecoli-$i.fa -o T -p F
done
```

and then let's search for a specific gene – first, download a file containing your protein sequence of interest:

```
curl -O http://athyra.idyll.org/~t/crp.fa
```

and now search:

```
blastall -i crp.fa -d ecoli-21.fa -p tblastn -b 1 -v 1
```

## Questions and Discussion Points

Why do we use a lower cutoff of 1kb for the assemstats3 links, above? Why not 0?

## Followup work

Try running an assembly of the larger read data set:

```
velveth ecoli-full.31 31 -short -fastq.gz ecoli_ref-5m-trim.fastq.gz
velvetg ecoli-full.31 -exp_cov auto
```

## 4.5 Interval Analysis and Visualization

The results generate below are based on a question posed by a participant in the course. She wanted to know how well contigs of an unfinished genomic build of and ecoli strain match the common (K-12 strain MG1655) genomic build.

Download the results from:

<http://apollo.huck.psu.edu/data/ms115.zip>

How did we get the results in the file above? A short description follows:

### 4.5.1 Data collection

The partial genomic build is located at:

[http://www.ncbi.nlm.nih.gov/genome/167?genome\\_assembly\\_id=161608](http://www.ncbi.nlm.nih.gov/genome/167?genome_assembly_id=161608)

From this we downloaded the summary file `code/ADTL01.txt` that happens to be a tab delimited file that lists accession numbers. We then wrote a very simple program `code/getdata.py` to parse the list of accessions and download the data like so

```
# requires BioPython
from Bio import Entrez
Entrez.email = "A.N.Other@example.com"
stream = file("ADTL01.txt")
stream.next()

for line in stream:
    elems = line.strip().split()
    val = elems[1]
    handle = Entrez.efetch(db="nucleotide", id=val, rettype="fasta", retmode="text")
    fp = file("data/%s.fa" % val, 'wt')
    fp.write(handle.read())
    fp.close()
```

Finally we merged all data with:

```
cat *.fa > MS115.fa
```

Then we went hunting for the EColi genome, we found it here:

<http://www.genome.wisc.edu/sequencing/updating.htm>

Turns out that this site only distributes a GBK (Genbank file). We now need to extract the information from the GBK file to FASTA (genome) and GFF (genomic feature) file. For this we need to install the ReadSeq program:

<http://iubio.bio.indiana.edu/soft/molbio/readseq/java/>

Once we have this we typed:

```
# GBK to GFF format
java -jar readseq.jar -inform=2 -f 24 U00096.gbk

# GBK to FASTA
java -jar readseq.jar -inform=2 -f 24 U00096.gbk
```

This will create the files U00096.gbk.fasta and U00096.gbk.gff

Now lets map the ms115.fa contigs to the U00096.fa reference:

```
bwa index U00096.fa
bwa mem U00096.fa ms115.fa | samtools view -bS - | samtools sort - U00096
```

will produce the U00096.bam file. We have converted the U00096.bam to BED format via the:

```
bedtools bamtobed -i U00096.bam > U00096.bed
```

Visualizing the data

Download and run IGV

<http://www.broadinstitute.org/igv/>

Create a custom genome via *Genomes* -> *Create .genome* options

We will now visualize the BAM, GFF and BED files and discuss the various aspects of it.

## 4.6 Running bedtools

Install bedtools:

```
sudo apt-get bedtools
```

This works best if you store your files in Dropbox, that way you can edit the file on your computer then load them up on your IGV instance.

## 4.7 Understanding the SAM format

Log into your instance, create a new directory, navigate to that directory:

```
cd /mnt
  mkdir sam
  cd sam

# Get the makefile.
wget https://raw.githubusercontent.com/ngs-docs/angus/2014/files/Makefile-samtools -O Makefile
```

A series of exercises will show what the SAM format is and how it changes when the query sequence is altered and how that reflects in the output.

Also, for the speed of result generation here is a one liner to generate a bamfile:

```
# One line bamfile generation.
bwa mem index/sc.fa query.fa | samtools view -bS - | samtools sort - results
```

This will produce the results.bam output.

This tutorial may be clearer to view it in markdown [https://github.com/ngs-docs/angus/blob/2014/R\\_Introductory\\_tutorial\\_2014.md](https://github.com/ngs-docs/angus/blob/2014/R_Introductory_tutorial_2014.md)>

## 4.8 R Tutorial for NGS2014

In this tutorial I will introduce R, a programming language that is currently the *lingua franca* for data analysis (although python has many powerful data analysis tools through `numpy`, `scipy` and other libraries).

### 4.8.1 What is R?

R is a bit of funny programming language, with a [funny history](#). While many researchers with a statistical bent really appreciate aspects of the syntax, and how it works, there are a great number of idiosyncracies, and “features” of the language that [drive many programmers crazy](#). It is also known to be a bit of a memory hog, so big data sets **can** be an issue to work with if you are not aware of the appropriate approaches to deal with. Despite this, there are some features of R that make it great for doing data analyses of all kinds, and there are a huge number of libraries of tools covering all aspects of statistics, genomics, bioinformatics, etc... In many cases new methods are associated with at least a basic implementation in R if not a full blown library.

### 4.8.2 Installing R

Many of you may already have R installed on your local machine. If you would like to install R on your computer just go to <http://www.r-project.org/> click download and follow for your OS. For linux machines, it is best to instead use your package manager (`yum` or `apt-get` for instance).

If you would like to run this on an Amazon EC2 instance, [set up and log into your instance](#) as you did in the earlier tutorial.

R generally does not come pre-installed on most Linux distributions including Ubuntu (Debian) Linux, which we are using on our Amazon EC2 instance but it is very easy to install:

```
apt-get install r-base
```

You may wish to [link and mount your instance to dropbox](#) as we did in the earlier tutorials.

It is also important to point out that R has been packaged in a number of different easy to use ways. For instance, many scientists really like [R-studio](#) which is also a free easy to use IDE that makes it easier for you to implement *version control* and do reproducible research. I personally do not use R via R-studio, but feel free to try it out on your local machine if you want to.

For the remainder of the tutorial I am assuming that we will all be running R on your Amazon EC2 instance though.

## R

### 4.8.3 What is R, really....

While many of you are used to seeing R through some GUI (whether on windows, OSX,etc).. It is fundamentally just a command line program like what you have been using at the command line or when you call `python`, etc.. The GUI versions just add some additional functionality.

At your shell prompt just type:

```
R
```

You now have a new command prompt (likely `>`) which means you are in R. Indeed you probably see something that looks like:

```
## R version 3.0.1 (2013-05-16) -- "Good Sport"
## Copyright (C) 2013 The R Foundation for Statistical Computing
## Platform: x86_64-apple-darwin10.8.0 (64-bit)

## R is free software and comes with ABSOLUTELY NO WARRANTY.
## You are welcome to redistribute it under certain conditions.
## Type 'license()' or 'licence()' for distribution details.
##
## Natural language support but running in an English locale
##
## R is a collaborative project with many contributors.
## Type 'contributors()' for more information and
## 'citation()' on how to cite R or R packages in publications.

## Type 'demo()' for some demos, 'help()' for on-line help, or
## 'help.start()' for an HTML browser interface to help.
## Type 'q()' to quit R.
##
## >
```

I will quickly (on the Mac OSX version) show you some of navigating the GUI.

## 4.8.4 How to close R

To quit R:

```
q()
```

**NOTE to MAC (OSX) users:** This may no longer work (R V2.11.+ ) from the Mac R GUI... See below in OSX specific notes..

**NOTE:** For the moment when it asks you to save workspace image, say no.

## 4.8.5 R Basics

Ok, return to R (type `R` at unix command prompt) so we can begin the tutorial.

I stole this table from *Josh Herr's* “R” tutorial <[https://github.com/jrherr/quick\\_basic\\_R\\_tutorial/blob/master/R\\_tutorial.md](https://github.com/jrherr/quick_basic_R_tutorial/blob/master/R_tutorial.md)>‘\_\_:

Table 1 - Important R Symbols

Symbol	Meaning / Use
<code> </code>	Prompt for a new command line, you do NOT type this, it appears in the console
<code>   </code>	Continuation of a previously started command, you do NOT type this
<code>    #</code>	Comment in code; R does not “see” this text – important to explain your computational train of thought
<code>    &lt;-</code> or <code>    =</code>	set a value to a name

**Note:** Comments in R are performed by using `#`. Anything followed by the number sign is ignored by R.

For the purposes of this tutorial we will be typing most things at the prompt. However this is annoying and difficult to save. So it is generally best practice to write things in a script editor. The R GUIs have some (OSX version with syntax highlighting as does gedit). So does R Studio. If we have time I will show you how to set up syntax highlighting on nano as well.

### 4.8.6 R as a calculator

Let's start by adding 2+2.

```
2 + 2
```

This will produce the output:

```
## [1] 4
```

Since R is a programming language we can easily generate variables to store all sorts of things.

```
y = 2
```

When you create a variable like this, it does not provide any immediate output.

However when you type `y` and press return:

```
y
```

You will see the output:

```
## [1] 2
```

Note that the `[1]` is just an index for keeping track where the answer was put. It actually means that it is the first element in a vector (more on this in a few minutes).

R is case SENSITIVE. That is `y` & `Y` are not the same variable.

```
Y
```

```
## [1] 2
```

```
Y
```

```
## Error: object 'Y' not found
```

We can generate variables and do computations with them.

```
x = 3
```

```
x + y
```

```
## [1] 5
```

```
z <- x + y
```

You will notice that sometimes I am using `'='` and sometimes `'<-'`. These are called *assignment operators*. In most instances they are equivalent. The `'<-'` is preferred in R, and can be used anywhere. You can look at the help file (more on this in a second) to try to parse the difference

```
`?`("=")
```

We may want to ask whether a variable that we have computed equals something in particular for this we need to use `'=='` not `'='` (one equals is an assignment, two means 'equal to')

```
x == 3
```

```
## [1] TRUE
```

```
x == 4
```



```
## [1] FALSE
```

```
x == y
```

```
## [1] FALSE
```

What happens if we write?

```
x = y
```

We have now assigned the current value of `y` (2) to `x`. This also shows you that you can overwrite a variable assignment. This is powerful, but also means you need to be very careful that you are actually using the value you think you are.

Operator `*` for multiplication.

```
2 * 3
```

```
## [1] 6
```

For division `/`.

```
6/3
```

```
## [1] 2
```

Operator for exponents `^`. Can also use `**`

```
3^2
```

```
## [1] 9
```

```
3^2 # same as above
```

```
## [1] 9
```

You can use `^0.5` or `sqrt()` function for square roots.

```
9^0.5
```

```
## [1] 3
```

```
sqrt(9)
```

```
## [1] 3
```

**to raise something to  $e^{\text{some exponent}}$**

```
exp(2) # this is the performing  $e^2$ 
```

```
## [1] 7.389
```

For natural log (base  $e$ ), use `log()`.

```
log(2.7)
```

```
## [1] 0.9933
```

To raise to an arbitrary base use `log( , base)` like the following:

```
log(2.7, 10) # base 10
## [1] 0.4314
```

You can also use `log10()` or `log2()` for base 10 or base 2.

While these are all standard operators (except `<-`) common to most programming languages, it is a lot to remember.

### A bit on data structures in R

R is **vectorized**. It can do its operations on vectors. Indeed there is no data structure in R for scalar values at all. This means when you assign `y <- 2` you are in fact generating a vector of length 1.

```
a <- c(2, 6, 4, 5)
b <- c(2, 2, 2, 1)
```

The `c` is short for concatenate you can add the elements of the vectors together. Keep in mind *how* it is doing its addition this way (element-by-element).

```
a + b
## [1] 4 8 6 6
```

Or multiply the elements of the vector together (this is **NOT** vector multiplication, but element-by-element multiplication. For vector multiplication (inner & outer products) there are specific operators, i.e. `%*%`).

```
a * b
## [1] 4 12 8 5
```

If you want to take vectors `a` and `b` (imagine these are each columns of read counts for RNAseq from different samples) and put them together in a single vector you use the `c()` (concatenate) function. I just realized that I am using `c` as a variable name (for the new vector), and the function is `c()`. This is entirely by accident and they have no relation to one another.

```
c <- c(a, b)
```

How might you make a vector that repeats vector `a` 3 times?

MANY MANY operations can be vectorized, and R is really good at it!!! Vectors are very useful as way of storing data relevant for all sorts of data analysis.

## 4.8.7 GETTING HELP in R

There are a number of places where you can get help with R directly from the console. For instance, what if we want to get help on the function `lm()` which is what is used to fit all sorts of *linear models* like regression, ANOVA, ANCOVA etc..

```
?lm
```

This brings up a description of the function ‘lm’

For operators we need to use quotes

```
? '*' # for help for operators use quotes
```

sometimes you will need to use `help.search('lm')` This brings up all references to the `lm()` function in packages and commands in R. We will talk about packages later.

You can also use `RSiteSearch('lm')`. This is quite a comprehensive search that covers R functions, contributed packages and R-help postings. It is very useful but uses the web (so you need to be online).

You can also use the html version of help using `help.start()`.

Or if using the GUI, just go to the help menu!

## 4.8.8 Simple functions in base R

R has many functions, and indeed everything (including the operators) are functions (even though the details of the function call are sometimes hidden). As we will see it is very easy to also write new functions.

Here are some examples of the useful functions in base R.

You can find out the length of the new vector using `length()`:

```
length(c)
```

```
## [1] 8
```

`length()` is an example of a pre-built function in R. Most things in R revolve around using functions to do something, or extract something. We will write our own simple functions soon.

Here are some more common ones that you may use

```
mean(c)
```

```
## [1] 3
```

```
sum(c)
```

```
## [1] 24
```

standard deviation

```
sd(c)
```

```
## [1] 1.773
```

variance

```
var(c)
```

```
## [1] 3.143
```

Correlation coefficient.

```
cor(a, b)
```

This gets the Pearson correlation (there are options to change this to other types of correlations, among the arguments for this function.....).

```
## [1] -0.2928
```

Say we want to keep the mean of `c` for later computation we can assign it to a variable

```
mean_c <- mean(c)
```

We can look at the underlying code of the function (although some times it is buried, in these cases).

We can use a function like `mean()` to add up all of the elements of the vector.

We can also join the two vectors together to make a matrix of numbers.

```
d <- cbind(a, b)
d

##           a b
## [1,]  2  2
## [2,]  6  2
## [3,]  4  2
## [4,]  5  1
```

We can double check that we really made a matrix:

```
is.matrix(d)
```

This sets up a ‘Boolean’. In other words when we ask ‘is d a matrix?’ it answers TRUE or FALSE.

```
## [1] TRUE
```

```
mode(d)
```

```
## [1] "numeric"
```

```
class(d)
```

```
## [1] "matrix"
```

While the mode of d is still numeric (the most basic “atomic” type of the data), the class of the object we have created is a matrix.

Exercise: Make a new vector q that goes a,b,a\*b

### 4.8.9 Objects in R, classes of objects, mode of objects.

R is an object-oriented language. Everything in R is considered an object. Each object has one or more attributes (which we do not generally need to worry about, but useful for programming.) Most objects in R have an attribute which is the ‘class’ of the object, which is what we will usually care about. R has a bunch of useful classes for statistical programming. `bioconductor` also has expanded the classes for objects of use for bioinformatics and genomics. We will see these more this afternoon and in the tutorials next week.

The most basic (atomic) feature of an object. NOTE this does not mean the ‘mode’ of a distribution.

```
mode(c)
```

```
## [1] "numeric"
```

```
class(c) # class of the object
```

```
## [1] "numeric"
```

In this case for the vector c the mode and class of c is the same. This is not always going to be the case as we see below.

```
# typeof(c) # internal representation of type, rarely of interest
```

Let’s look at some other objects we have generated.

```
mode(mean_c)

## [1] "numeric"

class(mean_c) #

## [1] "numeric"
```

Despite what we have seen up to now, mode and class are not the same thing. mode tells the basic ‘structures’ for the objects. integer, numeric (vector), character, logical (TRUE,FALSE) are examples of the atomic structures.

There are many different classes of objects each with their own attributes. The basic ones that we will see are numeric, character, matrix, data.frame, formula, array, list & factor. It is relatively straightforward (but we will not discuss it here) to extend classes as you need.

We can also make vectors that store values that are not numeric.

```
cities <- c("Okemos", "E.Lansing", "Toronto", "Montreal")
class(cities)

## [1] "character"

mode(cities)

## [1] "character"
```

Let’s use one of the built-in functions that we used before to look at the “length” of cities.

```
length(cities)

## [1] 4
```

This tells us how many strings we have in the object ‘cities’ not the length of the string! To get the number of characters in each string we use nchar().

```
nchar(cities) # This tells us how many characters we have for each string.

## [1] 6 9 7 8
```

So if we just do this

```
q = "okemos"
length(q)

## [1] 1

nchar(q)

## [1] 6
```

Exercise: How would you compute the total number of characters for all of cities in the object cities?

Let’s create a second vector storing the names of rivers in each city.

```
rivers <- c("Red Cedar", "Red Cedar", "Don Valley", "Sainte-Laurent")
cities_rivers <- cbind(cities, rivers)
cities_rivers
```

```
##      cities      rivers
## [1,] "Okemos"    "Red Cedar"
## [2,] "E.Lansing" "Red Cedar"
## [3,] "Toronto"   "Don Valley"
## [4,] "Montreal"  "Sainte-Laurent"
```

```
class(cities_rivers)
```

```
## [1] "matrix"
```

```
mode(cities_rivers)
```

```
## [1] "character"
```

In this above example we have made a matrix, but filled with characters, not numerical values.

Another type of object we will need for this workshop (well a variant of it) is called formula Not surprisingly this is used generally to generate a formula for a statistical model we want to fit.

```
model_1 <- y ~ x1 + x2 + x1:x2
model_1
```

This is just the model formula, and we HAVE NOT FIT ANY MODEL YET!!!!!! It just tells us the model we want to fit. That is the object `model_1` has not yet been ‘evaluated’.

```
## y ~ x1 + x2 + x1:x2
## <environment: 0x100b504b0>
```

```
# typeof(model_1)
class(model_1)
```

```
## [1] "formula"
```

```
terms(model_1) # also see all.names() and all.vars
```

```
## y ~ x1 + x2 + x1:x2
## attr(,"variables")
## list(y, x1, x2)
## attr(,"factors")
##      x1 x2 x1:x2
## y      0  0      0
## x1     1  0      1
## x2     0  1      1
## attr(,"term.labels")
## [1] "x1"      "x2"      "x1:x2"
## attr(,"order")
## [1] 1 1 2
## attr(,"intercept")
## [1] 1
## attr(,"response")
## [1] 1
## attr(,".Environment")
## <environment: 0x100b504b0>
```

Let’s make a new vector that will store some read counts of transcript expression values (from a single transcript). The first four are from samples of a “wild type” genotype and the second four samples from a “mutant” genotype.

```
counts_transcript_a <- c(250, 157, 155, 300, 125, 100, 153, 175)
```

We may be interested in comparing and contrasting these. So we need to make a variable that stores the information on the two different genotypes. We do this using one of the underlying classes in R, called `factors`.

We will make a new variable for the mutant and wild type using the `gl()` (generate levels) function. There are other ways of generating levels for a categorical treatment variable (and R usually can figure this out), but this will work for here.

```
genotype <- gl(n=2, k=4, labels = c("wild_type", "mutant"))
```

This is our first time using arguments within a function. for the `gl()` the `n=2` means we have two treatment levels. `k=4` means we 4 samples within each.

One obvious thing we might want to do is make a single object with both variables (genotype, and counts\_transcript\_a).

Your first thought (given that we just did this a few minutes ago) might be to make a matrix from this.

```
expression_data <- matrix(counts_transcript_a, genotype)
```

But as you see you gget an error message

```
## Error in matrix(counts_transcript_a, genotype) :  
## non-numeric matrix extent
```

Why? Because for an object of class `matrix` all of the atomic types in it need to be the same. Basically we can not combine numeric and factors.

The way we do this is to instead create a `data.frame`. A `data.frame` is a particular representation of another type of object (`list`) that allows for the storage of heterogeneous data types. Effectively you can think about it like a spreadsheet for purposes of analysis (for now anywaus).

```
expression_data <- data.frame(counts_transcript_a, genotype)  
expression_data
```

When you import most of your data, and do most analyses, you will more often than not be using a data frame.

## 4.8.10 Workspaces, and objects in them

R stores variables, datafiles, functions, vectors, etc in what is called the Workspace. This contains all of the items that you can access directly within your R session. You can list all of the objects in your # workspace using:

```
ls()
```

```
## [1] "a"           "b"           "c"           "cities"  
## [5] "cities_rivers" "d"           "mean_c"      "model_1"  
## [9] "q"           "rivers"      "x"           "y"  
## [13] "z"
```

If you want to remove a particular variable (say `x`) use the `rm()` function

```
rm(x)
```

you could remove multiple objects

```
rm(x, y, z)
```

```
## Warning: object 'x' not found
```

If you want to remove all of the objects in your workspace `rm(list = ls())`. We will learn what this means later, but basically we are making a list that contains all of the objects found by performing `ls()`, and then removing everything in that list.

Some people like to save their workspaces, not only because it contains all of the commands they have written, but also all of the objects they have created during that session. I personally do not do this unless I have created objects that have taken a long time to compute. Instead I just save the scripts I write.

However if you write your commands directly at the console (like we have been doing, but really you should not do) without a script editor, you should save your workspaces.

```
save.image('file_name')
```

Which will save it to your current working directory (you can find that using `getwd()`).

If you want to load it again

```
load('file_name.RData')
```

You will need to have the correct working directory set, which I will show you how to do shortly.

### 4.8.11 SCRIPT!

Writing everything at the console can be a bit annoying, so we will use a script editor.

In Mac OS X I personally find the built-in script editor useful. You can highlight the text in the script editor and press command (apple) + return to send it to the R console. Or place the cursor at the end of the line that you want to submit to R with command+ return. It also provides syntax highlighting, and shows the syntax & options for functions.

However, for those of you are under the spell of Bill Gates..... While the basic script editor for windows does not have much functionality, many people have written excellent script editors. The base script editor in windows will submit a line with ctrl-R(???). There are many windows script editors with syntax highlighting (such as Tinn-R).

For a list of some [http://www.sciviews.org/\\_rgui/](http://www.sciviews.org/_rgui/) In general we will save R scripts with the extension .R

Also there is R-studio.

let's type something into our new script

```
x <- c(3, 6, 6, 7)
```

now highlight that line and press ctrl+r (windows), or apple key + return (mac). This should send the highlighted portion to R.

```
x <- c(2, 2, 2, 2)
y <- c(3, 3, 3, 3)
z <- cbind(x, y)
z
```

```
##      x y
## [1,] 2 3
## [2,] 2 3
## [3,] 2 3
## [4,] 2 3
```

We have also just used a new function, `cbind`, which stands for column bind. This will create a new object stitching them together as columns.

### 4.8.12 Writing our own functions in R

We have now used a few built in functions in R (there are many). Anything where you use `()` is a function in R. Like I mentioned, pretty much everything you do in R is actually a call to a function.



However, we will often want to compute something for which there is no pre-built function. Thankfully it is very easy to write our own functions in R. You should definitely get in the habit of doing so.

Functions have the following format:

```
aFunction <- function(input variable 1, input variable 2, argument, etc...) {expressions to calculate}
```

This is abstract so let me give you a real example. For our read counts, we want to compute the standard error of the mean (a measure of sampling uncertainty), which is  $\sim$ equal to the  $\text{sd}/\sqrt{\text{sample size}}$ . How might we do it?

We want to compute it for the numeric vector of read counts.

```
counts_transcript_a
```

We could do it by hand

```
sd_a <- sd(counts_transcript_a)
sample_a <- length(counts_transcript_a)
sd_a/sqrt(sample_a)
```

```
## [1] 23.35857
```

(notice that we the last value was printed, as we did not store it in a variable).

Or we could do it in one line

```
sd(a)/sqrt(length(a)) # notice the function within a function
```

```
## [1] 23.35857
```

But we can also do it so that we can use any vector input we wanted by writing a function.

```
StdErr <- function(vector) {
  sd(vector)/sqrt(length(vector))
}
```

Now type StdErr

```
StdErr
```

```
## function(vector) {
##   sd(vector)/sqrt(length(vector))
## }
## <environment: 0x100b504b0>
```

This just repeats the function. If you want to edit the function just type `edit(StdErr)`.

Let's use our new function

```
StdErr(counts_transcript_a)
```

```
## [1] 23.35857
```

But now we can use this for any set of samples that we need to calculate the SEM. In this case transcript 'b'.

```
counts_transcript_b <- c(75, 85, 82, 79, 77, 83, 96, 62)
StdErr(counts_transcript_b)
```

```
## [1] 3.414452
```

Exercise: Write your own function to do something simple, like calculate the co-efficient of variation (CV) which is the  $\text{sd}/\text{mean}$ .

It takes some practice but learning to write small discrete functions can be extremely helpful for R.

One thing to keep in mind, is that it is very easy to call one function from within another. It is generally considered good practice to write functions that do one thing, and one thing only. It is way easier to find problems (debug).

One of the great things that we can (and need to do) often is to compute the mean (or SEM, etc..) for each transcript by sample. We can do this for the data set we made

```
expression_data
```

I will not explain it at the moment, but we can use one of the apply like functions to compute the mean and SEM for each genotype (works the same whether it is 2 or 2000).

```
with(expression_data, tapply(X=counts_transcript_a, INDEX=genotype, FUN=mean))
```

```
## wild_type    mutant
## 215.50      138.25
```

And then for the SEM

```
with(expression_data, tapply(X=counts_transcript_a, INDEX=genotype, FUN=StdErr))
```

```
# wild_type    mutant
# 35.83876    16.34715
```

The `with()` just makes it a bit easier to utilize the variables within the `expression_data` object. There are a number of other ways of doing it, primarily using the `$` (extract) operator (for lists including data frames). You will also see people use the `attach()`. Avoid using `attach()` at all costs.

### 4.8.13 Using `source()` to load your functions

One of the great things about writing simple functions, is that once you have them working, you can keep using them over and over. However, it is generally a pain (and bad practice) to have to include the text of the function in every script you write (what if there is a bug...). Instead, R has a function `source()` which allows you to 'load' a script that contains functions you have written (and other options you may want), so that you can use them.

We will likely see `source()` in later tutorials, so watch for it.

### 4.8.14 Regular Sequences

Sometimes we want regular sequences or to create objects of repeated numbers or characters. R makes this easy.

If you want to create regular sequences of integers by units of 1

```
one_to_20 <- 1:20
one_to_20
```

```
## [1] 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

```
twenty_to_1 <- 20:1
twenty_to_1
```

```
## [1] 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1
```

for other more complicated sequences, use the `seq()` function

```
seq1 <- seq(from = 1, to = 20, by = 0.5)
seq1
```

```
## [1] 1.0 1.5 2.0 2.5 3.0 3.5 4.0 4.5 5.0 5.5 6.0 6.5 7.0 7.5
## [15] 8.0 8.5 9.0 9.5 10.0 10.5 11.0 11.5 12.0 12.5 13.0 13.5 14.0 14.5
## [29] 15.0 15.5 16.0 16.5 17.0 17.5 18.0 18.5 19.0 19.5 20.0
```

or

```
seq1 <- seq(1, 20, 0.5)
seq1
```

```
## [1] 1.0 1.5 2.0 2.5 3.0 3.5 4.0 4.5 5.0 5.5 6.0 6.5 7.0 7.5
## [15] 8.0 8.5 9.0 9.5 10.0 10.5 11.0 11.5 12.0 12.5 13.0 13.5 14.0 14.5
## [29] 15.0 15.5 16.0 16.5 17.0 17.5 18.0 18.5 19.0 19.5 20.0
```

This shows that for default options (in the correct order) you do not need to specify things like ‘from’ or ‘by’

Exercise: Make a sequence from -10 to 10 by units of 2

What if you want to repeat a number or character a set number of times?

```
many_2 <- rep(2, times = 20)
```

Works for characters as well

```
many_a <- rep("a", times = 10)
```

We can even use this to combine vectors

```
seq_rep <- rep(20:1, times = 2)
seq_rep
```

```
## [1] 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 20 19 18
## [24] 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1
```

What if you wanted to repeat a sequence of numbers (1,2,3) 3 times?

```
rep_3_times <- rep(c(1, 2, 3), times = 3)
# or
rep(1:3, times = 3)
```

```
## [1] 1 2 3 1 2 3 1 2 3
```

What if we wanted to perform this to create a matrix

```
matrix(rep(20:1, 4), 20, 4)
```

```
##      [,1] [,2] [,3] [,4]
## [1,] 20  20  20  20
## [2,] 19  19  19  19
## [3,] 18  18  18  18
## [4,] 17  17  17  17
## [5,] 16  16  16  16
## [6,] 15  15  15  15
## [7,] 14  14  14  14
## [8,] 13  13  13  13
## [9,] 12  12  12  12
## [10,] 11  11  11  11
## [11,] 10  10  10  10
## [12,] 9   9   9   9
## [13,] 8   8   8   8
## [14,] 7   7   7   7
## [15,] 6   6   6   6
```

```
## [16,] 5 5 5 5
## [17,] 4 4 4 4
## [18,] 3 3 3 3
## [19,] 2 2 2 2
## [20,] 1 1 1 1
```

#### 4.8.15 Indexing, extracting values and subsetting from the objects we have created

Often we will want to extract certain elements from a vector, list or matrix. Sometimes this will be a single number, sometimes a whole row or column.

We index in R using `[ ]` (square brackets).

**NOTE** R indexes starting with 1, not 0!

```
a <- 1:20
b <- 5 * a
a

## [1] 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20

b

## [1] 5 10 15 20 25 30 35 40 45 50 55 60 65 70 75 80 85
## [18] 90 95 100

length(a)

## [1] 20

length(b)

## [1] 20
```

If we want to extract the 5th element from `a`.

```
a[5]

## [1] 5
```

If we want to extract the 5th and 7th element from `'b'`

```
b[c(5, 7)]

## [1] 25 35
```

If we want to extract the fifth through 10th element from `'b'`

```
b[5:10]

## [1] 25 30 35 40 45 50
```

How about if we want all but the 20th element of `'a'`?

```
a[-20]

## [1] 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19
```

Indexing can also be used when we want all elements greater than (less than etc...) a certain value. Under the hood this is generating a logical/boolean (T vs. F).

```
b[b > 20]

## [1] 25 30 35 40 45 50 55 60 65 70 75 80 85 90 95 100
```

Or between certain numbers

```
b[b > 20 & b < 80]

## [1] 25 30 35 40 45 50 55 60 65 70 75
```

Exercise: generate a vector with 20 elements create a 'sub' vector that has elements 1:5, 16:20 create a 'sub' vector with odd elements 1,3,5,...,19.

```
c <- a + b
q_matrix <- cbind(a, b, c)
q_matrix
```

`cbind()` 'binds' column vectors together into a matrix (also see `rbind()`).

```
##      a      b      c
## [1,]  1      5      6
## [2,]  2     10     12
## [3,]  3     15     18
## [4,]  4     20     24
## [5,]  5     25     30
## [6,]  6     30     36
## [7,]  7     35     42
## [8,]  8     40     48
## [9,]  9     45     54
## [10,] 10     50     60
## [11,] 11     55     66
## [12,] 12     60     72
## [13,] 13     65     78
## [14,] 14     70     84
## [15,] 15     75     90
## [16,] 16     80     96
## [17,] 17     85    102
## [18,] 18     90    108
## [19,] 19     95    114
## [20,] 20    100    120
```

(What happens if we ask for the length of `q.matrix`?...)

```
length(q_matrix)
```

We can instead ask for number of rows or columns

```
nrow(q_matrix)
```

```
## [1] 20
```

```
ncol(q_matrix)
```

```
## [1] 3
```

Or just use `dim(q_matrix)` (dimensions) to get both # rows and # columns. R always specifies in row by column format.

```
## [1] 20 3
```

Say we want to extract the element from the 3rd row of the second column (b)?

```
q_matrix[3, 2]
```

```
## b
## 15
```

How about if we want to extract the entire third row?

```
q_matrix[3, ]
```

```
## a b c
## 3 15 18
```

We can also pull things out by name

```
q_matrix[, "c"]
```

```
## [1] 6 12 18 24 30 36 42 48 54 60 66 72 78 84 90 96 102
## [18] 108 114 120
```

This is an example of indexing via ‘key’ instead of numerical order.

The at @ is used to extract the contents of a slot in an object.. We will not use it much for this class, but it is essential for object oriented programming in R (S4) objects. `objectName@slotName`.

More often we will use the dollar sign \$, which is used to extract elements of an object of class list (including data frames).. We will use this a lot to extract information from objects (such as information from our models, like coefficients) `object.name$element.name`.

For more information ? ‘\$’.

#### 4.8.16 Where to go from here?

There are a huge number of resources for R. Everyone has favorite tutorials and books. Here are but a few.

I have a few [screencasts](#) that you can access. I also have a number of [tutorials](#). I have way more R resources for a graduate class I teach in computational statistical approaches which I am happy to share as well.

The R site also has access to [numerous tutorials and books](#) or [other documents](#).

For more advanced programming check out [Hadley Wickham’s online book](#).

Here is a reasonably decent [R wikibook](#)

I really like the book [art of R programming](#).

#### 4.8.17 A few advanced topics... For your own amusement (not necessary for now, but helps for more advanced R programming).

Objects have attributes. The one we have thought about most is the class of the object, which tells us (and R) how to think about the object, and how it can be used or manipulated (methods). We have also looked at `dim()` which is another attribute Here is a list of common ones: `class`, `comment`, `dim`, `dimnames`, `names`, `row.names` and `tsp`.

We can set attributes of objects in easy ways like

```
x <- 4:6
names(x) <- c("observation_1", "observation_2", "observation_3")
x
```

```
## observation_1 observation_2 observation_3
##              4              5              6
```

You can see the attributes in a bunch of ways

```
str(x)

## Named int [1:3] 4 5 6
## - attr(*, "names")= chr [1:3] "observation_1" "observation_2" "observation_3"
```

```
attributes(x)
```

```
## $names
## [1] "observation_1" "observation_2" "observation_3"
```

Same as above, but we will be able to use this to set attributes of the object x as well

```
attr(x, "names")
```

```
## [1] "observation_1" "observation_2" "observation_3"
```

```
y <- cbind(1:5, 11:15)
attributes(y)
```

```
## $dim
## [1] 5 2
```

```
colnames(y) <- c("vec1", "vec2")
comment(y) <- c("the first column is pretend data", "the second column is yet more pretend data ")
str(y)
```

```
## int [1:5, 1:2] 1 2 3 4 5 11 12 13 14 15
## - attr(*, "dimnames")=List of 2
## ..$ : NULL
## ..$ : chr [1:2] "vec1" "vec2"
## - attr(*, "comment")= chr [1:2] "the first column is pretend data" "the second column is yet more pretend data "
```

```
attributes(y)
```

```
## $dim
## [1] 5 2
##
## $dimnames
## $dimnames[[1]]
## NULL
##
## $dimnames[[2]]
## [1] "vec1" "vec2"
##
##
## $comment
## [1] "the first column is pretend data"
## [2] "the second column is yet more pretend data "
```

```
““T
```

Calling a function like `summary()` will do very different things for different object classes. We will use this call alot for data frames and output from statistical models, etc..

```
summary(x)  # numeric vector

##      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##      4.0     4.5     5.0     5.0     5.5     6.0

summary(string_1)  # character string

##      Length      Class      Mode
##           1 character character
```

The call to `summary()` is generic, which first looks at the class of the object, and then uses a class specific method to generate a summary of `x`.

```
summary(x)

##      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##      4.0     4.5     5.0     5.0     5.5     6.0

summary.default(x)

##      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##      4.0     4.5     5.0     5.0     5.5     6.0
```

but..

```
summary.lm(x)  # Since this was looking for an object of class lm

## Error: $ operator is invalid for atomic vectors
```

To see all of the methods used when you call the generic `summary()` for S3 classes.

```
methods(summary)

## [1] summary.aov                summary.aovlist
## [3] summary.aspell*            summary.connection
## [5] summary.data.frame         summary.Date
## [7] summary.default            summary.ecdf*
## [9] summary.factor              summary.glm
## [11] summary.infl                summary.lm
## [13] summary.loess*              summary.manova
## [15] summary.matrix              summary.mlm
## [17] summary.nls*                summary.packageStatus*
## [19] summary.PDF_Dictionary*    summary.PDF_Stream*
## [21] summary.POSIXct             summary.POSIXlt
## [23] summary.ppr*                summary.prcomp*
## [25] summary.princomp*           summary.proc_time
## [27] summary.srcfile              summary.scref
## [29] summary.stepfun              summary.stl*
## [31] summary.table                summary.tukeysmooth*
##
##      Non-visible functions are asterisked
```

## 4.8.18 Syntax style guide

Generally it is advisable to use a consistent way of scripting. For any given programming language there are syntax style guide. The [Style guide for my class](#). You can also check out the [R style guide from Google](#).



### 4.8.19 Random bits

Note about using `q()` on the Mac R GUI in v2.11.+ The programming team decided the default behaviour was potentially ‘dangerous’, and people may lose their files, so they have changed it to `command + q` to quit instead. If you are an old-fogey like me and like to use `q()`, you have a couple of options. `base::q()` # This will work, but it is annoying.

you can set your `.Rprofile` to have the following line.

```
options(RGUI.base.quit=T)
```

and the next time you run R the old `q()` will work.

If you do not know how to create or edit `.Rprofile`, come speak with me...

### 4.8.20 session info

The R session information (including the OS info, R version and all packages used):

```
sessionInfo()

## R version 3.0.1 (2013-05-16)
## Platform: x86_64-apple-darwin10.8.0 (64-bit)
##
## locale:
## [1] en_US.UTF-8/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
##
## attached base packages:
## [1] stats      graphics  grDevices  utils      datasets  methods   base
##
## other attached packages:
## [1] knitr_1.5
##
## loaded via a namespace (and not attached):
## [1] evaluate_0.5.1 formatR_0.10  stringr_0.6.2 tools_3.0.1

Sys.time()

## [1] "2014-08-07 10:39:49 EDT"
```

### R indexing begins at 1 (not 0 like Python) Negative values of indexes in R

mean something very different. for instance

`a[-1]` # this removes the first element of `a`, and prints out all of the remaining elements.

```
[1] 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
```

---

```
““
```

```
# As far as I know all classes of objects are mutable, which means you can
# write over the name of the objects, values within the objects, and
# slots....
```

```
# Indexing on a character string does not work in R
string_1 <- "hello world"
string_1[1]
```

```
## [1] "hello world"

# instead you need to use the substr() function
substr(x = string_1, start = 1, stop = 1)

## [1] "h"

# similarly
length(string_1) # this gives an output of 1

## [1] 1

nchar(string_1) # this gives the 11 characters

## [1] 11
```

## **TOC**

**Section 1: What is R; R at the console; quitting R**

**Section 2: R basics; R as a calculator; assigning variables; vectorized computation in R**

**Section 3: pre-built functions in R**

**Section 4: Objects, classes, modes - Note: should I add attributes?**

**Section 5: The R workspace; listing objects, removing objects (should I add attach and detach?)**

**Section 6: Getting Help in R**

**Section 7: Using A script editor for R**

**Section 8: Writing simple functions in R**

**Section 8b: Using source() to call a set of functions**

**Section 9: Regular sequences in R**

**Section 10: Extracting (and replacing), indexing & subsetting (using the index). Can also be used for sorting.**

**Advanced stuff to learn on your own...**

**..... setting attributes of objects.... (names, class, dim )**

**..... environments (see ?environments)**

**Control Flow and loops in R #####**

## 4.9 Control Flow

### 4.9.1 The standard if else

```
p.test <- function(p) {
  if (p <= 0.05)
    print("yeah!!!!") else if (p >= 0.9)
    print("high!!!!") else print("somewhere in the middle")
}
```

Now pick a number and put it in `p.test`

```
p.test(0.5)

## [1] "somewhere in the middle"
```

### 4.10 ifelse()

A better and vectorized way of doing this is `ifelse(test, yes, no)` function. `ifelse()` is far more useful as it is vectorized.

```
p.test.2 <- function(p) {
  ifelse(p <= 0.05, print("yippee"), print("bummer, man"))
}
```

Test this with the following sequence. See what happens if you use `if` vs. `ifelse()`.

```
x <- runif(10, 0, 1)
x

## [1] 0.27332 0.14155 0.89000 0.07041 0.79419 0.25013 0.02324 0.86766
## [9] 0.41114 0.56165
```

Now try it with `p.test()` (uses `if`).

```
p.test(x)

## Warning: the condition has length > 1 and only the first element will be used
## Warning: the condition has length > 1 and only the first element will be used

## [1] "somewhere in the middle"
```

Now try it with `p.test.2()`

```
p.test.2(x)

## [1] "yippee"
## [1] "bummer, man"

## [1] "bummer, man" "bummer, man" "bummer, man" "bummer, man" "bummer, man"
## [6] "bummer, man" "yippee"          "bummer, man" "bummer, man" "bummer, man"
```

## 4.11 Other vectorized ways of control flow.

There are many times that you may think you need to use an if with (iterating with a for loop... see below), or ifelse, but there may be far better ways.

For instance, say you are doing some simulations for a power analysis, and you want to know how often your simulation gives you a p-value less than 0.05.

```
p.1000 <- runif(n = 1000, min = 0, max = 1)
```

The line above generates 1000 random values between 0-1, which we will pretend are our p-values for differential expression from our simulation.

You may try and count how often it less than 0.05

```
p.ifelse <- ifelse(p.1000 < 0.05, 1, 0) # If it is less than 0.05, then you get a 1, otherwise 0.
```

Our approximate false positives. Should be close to 0.05

```
sum(p.ifelse)/length(p.1000)
```

```
## [1] 0.059
```

However the best and fastest way to accomplish this is to use the index, by setting up the Boolean (TRUE/FALSE) in the index of the vector.

```
length(p.1000[p.1000 < 0.05])/length(p.1000)
```

```
## [1] 0.059
```

Same number, faster and simpler computation.

## 4.12 Simple loops

### 4.12.1 while() function..

I tend to avoid these, so you will not see them much here

```
i <- 1
while (i <= 10) {
  print(i)
  i <- i + 0.5
}
```

```
## [1] 1
## [1] 1.5
## [1] 2
## [1] 2.5
## [1] 3
## [1] 3.5
## [1] 4
## [1] 4.5
## [1] 5
## [1] 5.5
## [1] 6
## [1] 6.5
## [1] 7
```

```
## [1] 7.5
## [1] 8
## [1] 8.5
## [1] 9
## [1] 9.5
## [1] 10
```

## 4.13 for loop

If I run a loop I most often use `for () {}` automatically iterates across a list (in this case the sequence from 1:10).

```
for (i in 1:10) {
  print(i)
}
```

```
## [1] 1
## [1] 2
## [1] 3
## [1] 4
## [1] 5
## [1] 6
## [1] 7
## [1] 8
## [1] 9
## [1] 10
```

If you do not want to use integers, how might you do it using the `for()`?

```
for (i in seq(from = 1, to = 5, by = 0.5)) {
  print(i)
}
```

```
## [1] 1
## [1] 1.5
## [1] 2
## [1] 2.5
## [1] 3
## [1] 3.5
## [1] 4
## [1] 4.5
## [1] 5
```

Using strings is a bit more involved in R, compared to other languages. For instance the following does not do what you want:

```
.. code:: r

  for (letter in "word") { print(letter)
  }

## [1] "word"
```

(try letters for a hoot.)

Instead in R, we have to split the word “word” into single characters using `strsplit()`, i.e:

```
.. code:: r

    strsplit("word", split = "")

## [[1]]
## [1] "w" "o" "r" "d"
```

## 4.14 So for the for loop we would do the following:

```
for (letter in strsplit("word", split = "")) {
  print(letter)
}

## [1] "w" "o" "r" "d"
```

## 4.15 More avoiding loops

Many would generate random numbers like so.

```
for (i in 1:100) {
  print(rnorm(n = 1, mean = 0, sd = 1))
}

## [1] -0.1837
## [1] -0.9313
## [1] 1.648
## [1] -0.6964
## [1] 0.2112
## [1] 0.3441
## [1] 1.036
## [1] 0.7439
## [1] 0.5859
## [1] -0.6087
## [1] -0.4014
## [1] 1.44
## [1] -0.3906
## [1] -1.861
## [1] -0.739
## [1] -1.204
## [1] 0.07794
## [1] -1.65
## [1] 1.261
## [1] 0.6753
## [1] 0.6736
## [1] 0.3238
## [1] -1.316
## [1] 0.2965
## [1] 1.499
## [1] 0.4326
## [1] 0.4488
## [1] 0.8873
## [1] -1.304
## [1] -0.347
## [1] 0.3491
```

```
## [1] 0.24
## [1] 0.1425
## [1] -0.2785
## [1] -0.5072
## [1] -1.775
## [1] -0.04051
## [1] 0.9452
## [1] 0.3322
## [1] -0.01994
## [1] -0.2308
## [1] -0.4053
## [1] -0.5685
## [1] -1.631
## [1] -0.1484
## [1] 0.434
## [1] 1.653
## [1] 1.57
## [1] 0.1308
## [1] -1.059
## [1] -0.7157
## [1] -0.8316
## [1] 0.06561
## [1] 0.8243
## [1] 0.1841
## [1] 1.048
## [1] 0.1612
## [1] -0.9553
## [1] -0.7569
## [1] -0.288
## [1] -1.837
## [1] 0.7301
## [1] -2.103
## [1] -1.869
## [1] -1.298
## [1] -1.077
## [1] -0.2139
## [1] -0.9419
## [1] 0.4694
## [1] -1.344
## [1] -0.08514
## [1] -2.055
## [1] -0.803
## [1] -0.7281
## [1] 1.778
## [1] -1.116
## [1] 1.33
## [1] 0.1535
## [1] -2.897
## [1] 0.7305
## [1] 1.228
## [1] 1.697
## [1] -0.8183
## [1] -1.013
## [1] -0.634
## [1] -0.942
## [1] -0.3395
## [1] 0.1396
## [1] 1.022
```

```
## [1] 0.9868
## [1] -0.7778
## [1] 1.075
## [1] -0.1029
## [1] 0.2644
## [1] 0.01165
## [1] 0.8025
## [1] -1.24
## [1] -0.8865
## [1] 0.981
## [1] 0.5333
```

We are cycling through and generating one random number at each iteration. Look at the indices, and you can see we keep generating vectors of length 1.

better/cleaner/faster to generate them all at one time

```
rnorm(n = 100, mean = 0, sd = 1)
```

```
## [1] -0.08683 -1.55262 -1.16909 0.30451 -1.14555 0.76682 0.12643
## [8] -0.61174 -0.29103 -0.10707 -0.03397 -0.05926 0.27294 1.32693
## [15] -0.53284 1.83234 0.43959 -0.88991 0.25383 0.96709 -0.23210
## [22] -1.00190 -1.32289 1.80030 1.15272 -1.82907 0.75989 1.35966
## [29] 0.53943 0.01429 -0.58707 -0.11886 -0.70367 -2.38988 0.08033
## [36] -0.22795 -0.62166 -0.19832 -1.95990 -0.85127 0.94236 0.37771
## [43] 0.32617 -0.08393 -0.54506 -2.58781 -0.58433 0.20985 -0.41613
## [50] 0.60527 0.51713 1.57950 -0.61079 -0.28564 -0.16444 0.55007
## [57] 0.57258 0.58513 -0.86728 -0.81185 -0.29333 -1.23935 0.46169
## [64] -1.53586 -0.32583 0.17629 -0.85579 1.04989 1.22120 1.53359
## [71] -2.37276 1.44393 1.47506 0.40110 -0.10157 0.35485 -0.72068
## [78] -1.27910 0.63152 -0.65216 1.60160 0.27109 0.50904 -1.00531
## [85] 0.76743 -0.78954 -0.01159 1.06944 1.15661 -0.91031 1.54919
## [92] -0.84334 2.19994 0.26716 0.02081 0.53577 0.07840 -0.79387
## [99] -1.18941 1.24745
```

### The not advisable approach

First we initialize a vector to store all of the numbers. Why do we initialize this vector first?

```
n <- 1e+05
x <- rep(NA, n)
```

## 4.16 The step above creates a vector of n NA's. They will be replaced sequentially with the random numbers as we generate them (using a function like the above one).

```
head(x)
```

```
## [1] NA NA NA NA NA NA
```

Now we run the for loop.

```
for (i in 1:n) {
  x[i] <- rnorm(n = 1, mean = 0, sd = 1)
}
```



for each `i` in the index, one number is generated, and placed in `x`

```
head(x)
```

```
## [1]  0.2848 -0.5432  1.1391 -1.0901  0.8515  0.5490
```

However this is computationally inefficient in R. Which has vectorized operations.

```
system.time(
```

```
  for (i in 1:n){
    x[i] <- rnorm(n=1, mean=0, sd=1)})
```

```
##      user  system elapsed
##  0.562    0.023    0.584
```

We can also use the `replicate` function to do the same thing. Easier syntax to write.

```
system.time(z <- replicate(n, rnorm(n = 1, mean = 0, sd = 1)))
```

```
##      user  system elapsed
##  0.561    0.035    0.841
```

This is ~20% faster.

However, since R is vectorized, both of the will be far slower than:

```
system.time(y <- rnorm(n, 0, 1))
```

```
##      user  system elapsed
##  0.010    0.000    0.011
```

About 65 times faster than the for loop

The general rule in R is that loops are slower than the `apply` family of functions (for small to medium data sets, not true for very large data) which are slower than vectorized computations.

## 4.17 Variant calling and exploration of polymorphisms

Now that we have some experience in R, we will check out a `vcf` file with polymorphisms from

### 4.17.1 ## Getting the data and installing extra packages

Installing a bunch of stuff:

```
get bwa
```

```
cd /root
```

```
wget -O bwa-0.7.10.tar.bz2 http://sourceforge.net/projects/bio-bwa/files/bwa-0.7.10.tar.bz2/download
```

untar and compile (via make) bwa

```
tar xvfj bwa-0.7.10.tar.bz2
```

```
cd bwa-0.7.10
```

```
make
```

```
cp bwa /usr/local/bin
```

install some tools

```
apt-get update
apt-get -y install samtools screen git curl gcc make g++ python-dev unzip \
    default-jre pkg-config libncurses5-dev r-base-core \
    r-cran-gplots python-matplotlib sysstat libcurl4-openssl-dev libxml2-dev

git clone https://github.com/schimar/ngs2014_popGen.git

cd ngs2014_popGen/var_call2/
```

### 4.17.2 Let's do another round of variant calling

index the reference genome

```
bwa index ref_genome.fna
```

map our reads to the indexed reference genome

```
bwa aln ref_genome.fna read_file.fq > mapped_reads.sai
```

Create the SAM file

```
bwa samse ref_genome.fna mapped_reads.sai read_file.fq > mapped_reads.sam
```

Index the reference genome

```
samtools faidx ref_genome.fna
```

Convert from SAM to BAM

```
samtools view -b -S -o mapped_reads.bam mapped_reads.sam
```

Sort the BAM

```
samtools sort mapped_reads.bam mapped_reads.sorted
```

And index again, but now the sorted BAM file

```
samtools index mapped_reads.sorted.bam
```

Visualize the alignment

```
samtools tview mapped_reads.sorted.bam ref_genome.fna
```

### 4.17.3 Variant exploration with Bioconductor

Now simply type R in the shell and:

```
source("http://bioconductor.org/biocLite.R")
biocLite()
biocLite("VariantAnnotation")
biocLite("SNPlocs.Hsapiens.dbSNP.20101109")
biocLite("BSgenome.Hsapiens.UCSC.hg19_1.3.1000")
```

#### 4.17.4 Quality control

Now we load the VariantAnnotation package as well as the data. The objective of this exercise is to compare the quality of called SNPs that are located in dbSNP, versus those that are novel.

```
library(VariantAnnotation)
fl <- system.file("extdata", "chr22.vcf.gz", package="VariantAnnotation")
```

Locate the sample data in the file system. Explore the metadata (information about the content of the file) using scanVcfHeader. Discover the ‘info’ fields VT (variant type), and RSQ (genotype imputation quality).

```
(hdr <- scanVcfHeader(fl))
info(hdr)[c("VT", "RSQ"),]
```

Input the data and peak at their locations:

```
(vcf <- readVcf(fl, "hg19"))
head(rowData(vcf), 3)
```

SNPs were called with MaCH/thunder (part of GotCloud), for more info, see <http://genome.sph.umich.edu/wiki/Thunder> and [http://genome.sph.umich.edu/wiki/MaCH\\_FAQ](http://genome.sph.umich.edu/wiki/MaCH_FAQ). Notice that the seqnames (chromosome levels) are set to ‘22’, we want to rename those

```
rowData(vcf) <- renameSeqlevels(rowData(vcf), c("22"="ch22"))
```

We now load the SNP database and discover whether our SNPs are in dbSNP

```
library(SNPlocs.Hsapiens.dbSNP.20101109)

destination <- tempfile()
pre <- FilterRules(list(isLowCoverageExomeSnp = function(x) {
  grepl("LOWCOV,EXOME", x, fixed=TRUE)
}))
filt <- FilterRules(list(isSNP = function(x) info(x)$VT == "SNP"))
snpFilt <- filterVcf(fl, "hg19", destination, prefilters=pre, filters= filt)
vcf_filt <- readVcf(snpFilt, "hg19")

rowData(vcf)
rowData(vcf_filt)
```

If we compare vcf and vcf\_filt, we see that of the 10376 SNPs in our initial vcf file, 794 are in the database.

```
inDbSNP <- rownames(vcf) %in% rownames(vcf_filt)
table(inDbSNP)
metrics <- data.frame(inDbSNP = inDbSNP, RSQ = info(vcf)$RSQ)
```

Let’s finally visualize it:

```
library(ggplot2)
ggplot(metrics, aes(RSQ, fill=inDbSNP)) +
  geom_density(alpha=0.5) +
  scale_x_continuous(name="MaCH / Thunder Imputation Quality") +
  scale_y_continuous(name="Density") +
  theme(legend.position="top")
```

(This won’t work in R on EC2, simply because we can’t run X11 through an ssh connection)

## 4.18 A complete de novo assembly and annotation protocol for mRNASeq

The goal of this tutorial is to run you through (part of) a real mRNAseq analysis protocol, using a small data set that will complete quickly.

Prepare for this tutorial by working through *Start up an EC2 instance*, but follow the instructions to start up *Starting up a custom operating system* instead; use AMI ami-7607d01e.

### 4.18.1 Switching to root

Start by making sure you're the superuser, root:

```
sudo bash
```

### 4.18.2 Updating the software on the machine

Copy and paste the following two commands

```
apt-get update
apt-get -y install screen git curl gcc make g++ python-dev unzip \
    default-jre pkg-config libncurses5-dev r-base-core \
    r-cran-gplots python-matplotlib sysstat samtools python-pip
```

If you started up a custom operating system, then this should finish quickly; if instead you started up Ubuntu 14.04 blank, then this will take a minute or two.

### 4.18.3 Downloading the sample data

The mRNAseq protocol works with the data set that you put in '/data'. Here, we will download a small data set (a subset of the data from [this paper](#), data from embryonic *Nematostella*>'), and put it in /data

```
mkdir /mnt/data
ln -fs /mnt/data /data
cd /data
curl -O http://athyra.idyll.org/~t/mrnaseq-subset.tar
tar xvf mrnaseq-subset.tar
```

Check it out:

```
ls
```

You'll see a bunch of different files – these are the kinds of files you'll get from your sequencing facility.

### 4.18.4 Starting on the protocols

We're going to work with a special version of the protocols today, one that we adapted specifically for this course.

**In general**, you should use the latest version, which will be at <https://khmer-protocols.readthedocs.org/>.

For today, we'll be using <http://khmer-protocols.readthedocs.org/en/ngs2014/> instead.

Work through the following:

1. [Quality trimming](#)

2. Applying digital normalization
3. Running the actual assembly
4. BLASTing your assembly

#### 4.18.5 Actually using the BLAST Web server

To connect to your BLAST Web server, you need to enable inbound traffic on your computer. Briefly:

- go to your instance and look at what security group you're using

(should be 'launch-wizard-' something). On the left panel, under Network and Security, go into Security Groups. Select your security group, and select Inbound, and Edit. Click "Add rule", and change "Custom TCP rule" to "http". Then click "save". Done!

You can try pasting this into your BLAST server:

```
MDRSVNVIQCAAPTRIQCEEINAKMLLGVGVFGLCMNIVLAVIMSFGAAPHSHGMLSSVEFDHVDYH
SRDNHHGHSHLHHEHQHRDGCSSHSHGNGGADMQRLECASPESEMMEVVVETTSNAESICSHERGSQSM
NLRAAVLHVFGDCLQSLGVVLAACVIWAGNNSVGVPSAHSYYNLADPLLSVLFVITVYTTLNLFKEV
IVILLEQVPPAVEYTVARDALLSVEKVQAVDDLHIWAVGPGFVLSAHLCTNGCATTSEANAVVEDAECR
CRQLGIVHTTIQLKHAADVRNTGA
```

## 4.19 Amazon Web Services instructions

### 4.19.1 Start up an EC2 instance

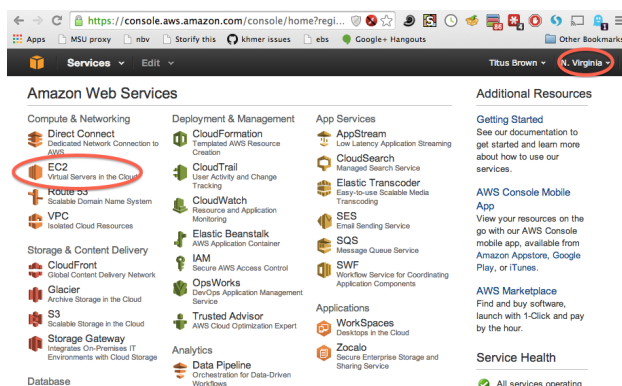
Here, we're going to startup an Amazon Web Services (AWS) Elastic Cloud Computing (EC2) "instance", or computer.

Go to '<https://aws.amazon.com>' in a Web browser.


Select 'My Account/Console' menu option 'AWS Management Console.'

Log in with your username & password.

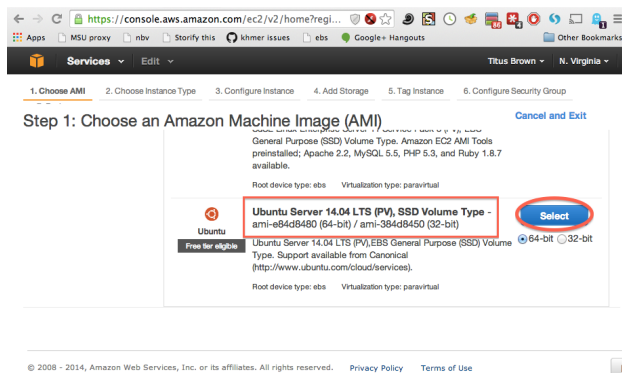
Make sure it says North Virginia in the upper right, then select EC2 (upper left).



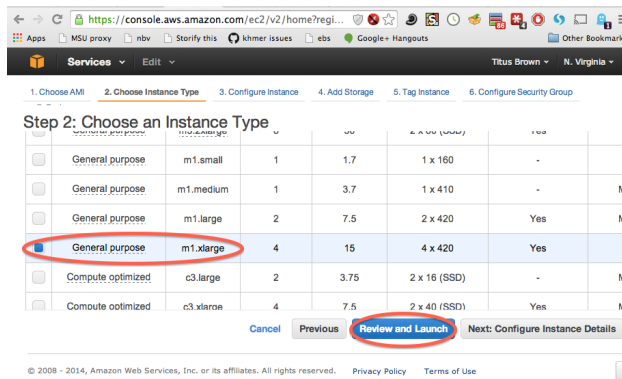
Select "Launch Instance" (midway down the page).




Next, scroll down the list of operating system types until you find Ubuntu 14.04 LTS (PV) – it should be at the very bottom. Click ‘select’. (See *Starting up a custom operating system* if you want to start up a custom operating system instead of Ubuntu 14.04.)



Scroll down the list of instance types until you find “m1.xlarge”. Select the box to the left, and then click “Review and Launch.”




Ignore the warning, check that it says “Ubuntu 14.04 LTS (PV)”, and click “Launch”.

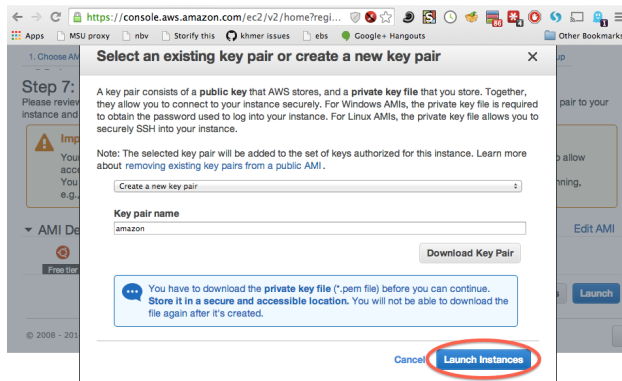


The *first* time through, you will have to “create a new key pair”, which you must then name (something like ‘amazon’) and download.

After this first time, you will be able to select an existing key pair.



Select “Launch Instance.”



Select “view instance” and you should see a “pending” line in the menu.





You should now see a text line that starts with something like `ubuntu@ip-10-235-34-223:~$`. You're in! Now type:

```
sudo bash
cd /root
```

to switch into superuser mode (see: <http://xkcd.com/149/>) and go to your home directory.

This is where the rest of the tutorials will start!

If you have Dropbox, you should now visit *Installing Dropbox on your EC2 machine*.

You might also want to read about *Terminating (shutting down) your EC2 instance*.

To log out, type:

```
exit
logout
```

or just close the window.

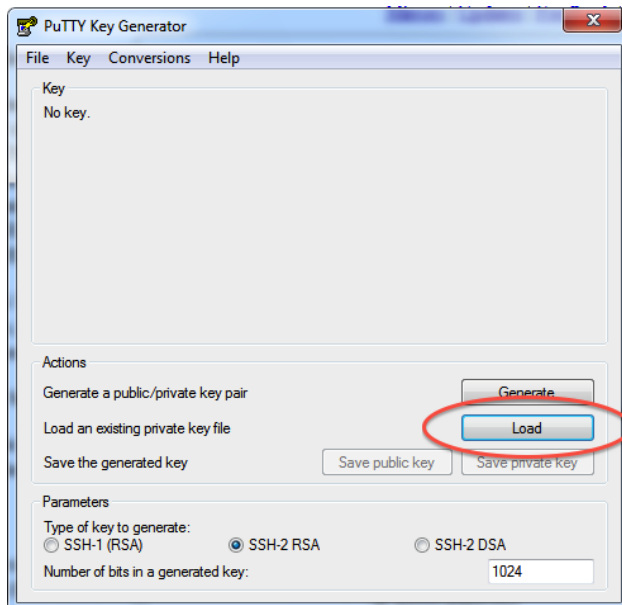
### 4.19.3 Logging into your new instance “in the cloud” (Windows version)

Download Putty and Puttygen from here: <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

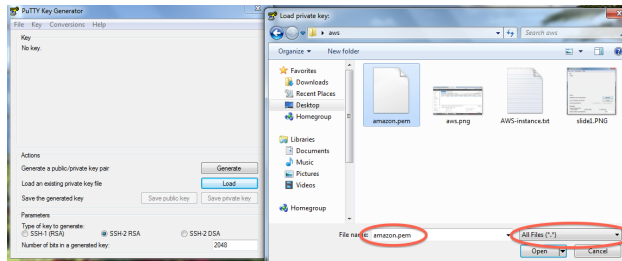
#### Generate a ppk file from your pem file

(You only need to do this once!)

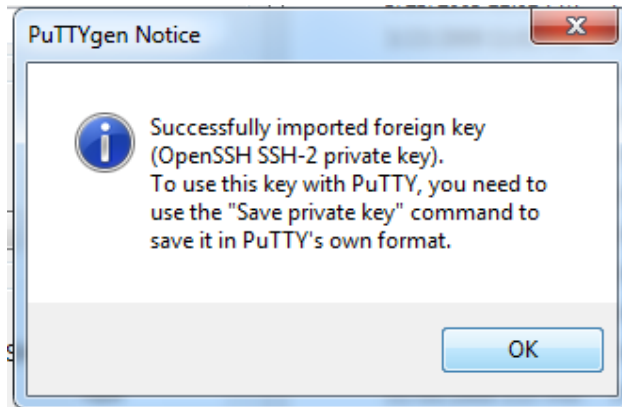
Open puttygen; select “Load”.



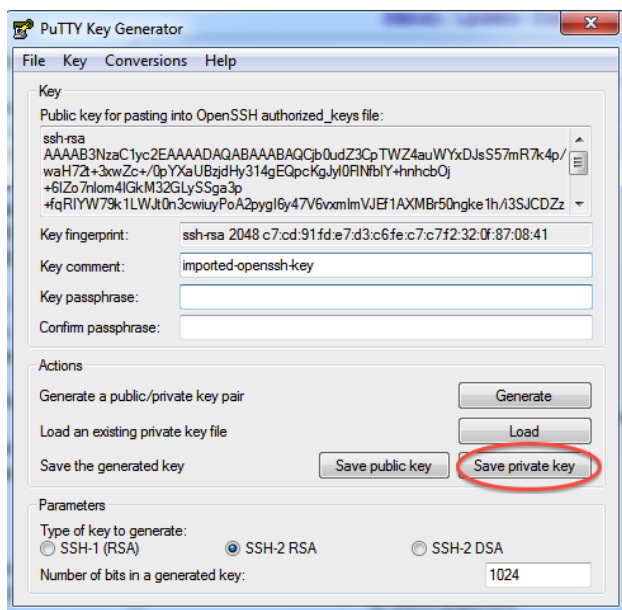
Find and load your ‘.pem’ file; it’s probably in your Downloads folder. Note, you have to select ‘All files’ on the bottom.



Load it.

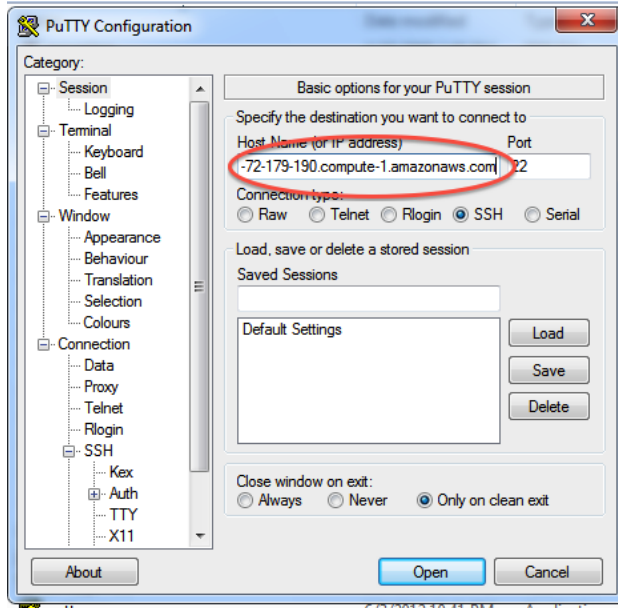


Now, "save private key". Put it somewhere easy to find.

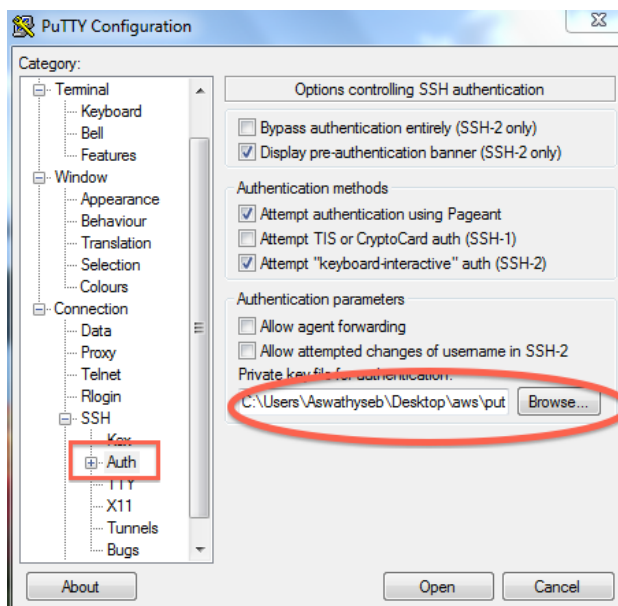


## Logging into your EC2 instance with Putty


Open up putty, and enter your hostname into the Host Name box.




Now, go find the 'SSH' section and enter your ppk file (generated above by puttygen). Then select 'Open'.



Log in as "ubuntu".



Declare victory!



Here, you're logging in as user 'ubuntu' to the machine 'ec2-174-129-122-189.compute-1.amazonaws.com' using the authentication key located in 'amazon.pem' on your Desktop.

You should now see a text line that starts with something like `ubuntu@ip-10-235-34-223:~$`. You're in! Now type:

```
sudo bash
cd /root
```

to switch into superuser mode (see: <http://xkcd.com/149/>) and go to your home directory.

This is where the rest of the tutorials will start!

If you have Dropbox, you should now visit *Installing Dropbox on your EC2 machine*.

You might also want to read about *Terminating (shutting down) your EC2 instance*.

To log out, type:

```
exit
logout
```

or just close the window.

## 4.19.4 Installing Dropbox on your EC2 machine

**IMPORTANT:** Dropbox will sync everything you have to your EC2 machine, so if you are already using Dropbox for a lot of stuff, you might want to create a separate Dropbox account just for the course.

Start at the login prompt on your EC2 machine:

```
cd /root
```

If you can not do this (“Permission denied”), make sure you are in superuser mode. (You should see a text line that starts with something like

```
root@ip-10-235-34-223:~$. If not, use “sudo bash” to switch.)
```

Then, grab the latest dropbox installation package for Linux:

```
wget -O dropbox.tar.gz "http://www.dropbox.com/download/?plat=lnx.x86_64"
```

Unpack it:

```
tar -xvzf dropbox.tar.gz
```

Make the Dropbox directory on /mnt and link it in:

```
mkdir /mnt/Dropbox
ln -fs /mnt/Dropbox /root
```

and then run it, configuring it to put stuff in /mnt:

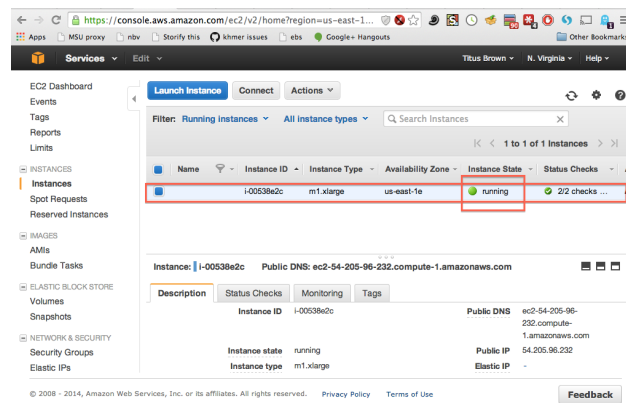
```
HOME=/mnt /root/.dropbox-dist/dropboxd &
```

When you get a message saying “this client is not linked to any account”, copy/paste the URL into browser and go log in. Voila!

Your directory ‘/root/Dropbox’, or, equivalently, ‘/mnt/Dropbox’, is now linked to your Dropbox account!

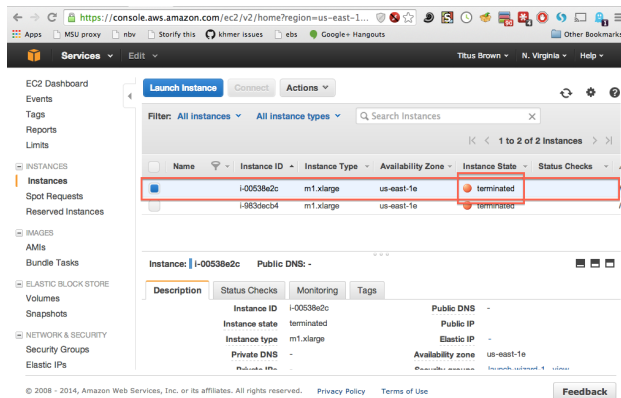
### 4.19.5 Terminating (shutting down) your EC2 instance

While your instance is running, Amazon will happily charge you on a per-hour basis – [check out the pricing](#) for more information. In general, you will want to shut down your instance when you’re done with it; to do that, go to your EC2 console and find your running instances (in green).



Then, select one or all of them, and go to the ‘Actions...’ menu, and then select ‘Terminate’. Agree.

After a minute or two, the console should show the instance as “terminated”.



### 4.19.6 Storing data persistently with Amazon EBS Volumes

If you want to save your data across instances – that is, if you want to have persistent data – Amazon can do that for you, too. Amazon is happy to rent disk space to you, in addition to compute time. They’ll rent you disk space in a few different ways, but the way that’s most useful for us is through what’s called Elastic Block Store (EBS). This is essentially a hard-disk rental service.

There are two basic concepts – “volume” and “snapshot”. A “volume” can be thought of as a pluggable-in hard drive: you create an empty volume of a given size, attach it to a running instance, and voila! You have extra hard disk space. Volume-based hard disks have two problems, however: first, they cannot be used outside of the “availability zone” they’ve been created in, which means that you need to be careful to put them in the same zone that your instance is running in; and they can’t be shared amongst people.

Snapshots, the second concept, are the solution to transporting and sharing the data on volumes. A “snapshot” is essentially a frozen copy of your volume; you can copy a volume into a snapshot, and a snapshot into a volume.

(Learn more from <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>)

In this and the following tutorials, you will be instructed to create new Amazon EBS Volume, to create Amazon EBS Snapshot from EBS Volume and to restore EBS Volume from EBS Snapshot.

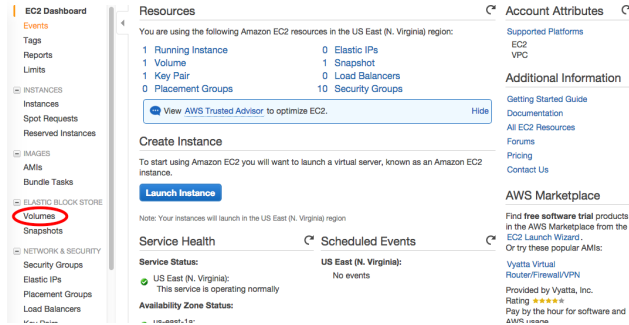
### Prerequisites

This tutorial assumes you’ve already set up an account on Amazon Web Services, and that you’ve completed the EC2 tutorial to set up an Amazon instances.


### Creating an Amazon EBS Volume

Firstly open the Amazon EC2 console at ‘<https://console.aws.amazon.com/ec2>’ and make sure it says North Virginia in the upper right.

At the AWS Management Console, on the left menu bar, click “Volumes”.

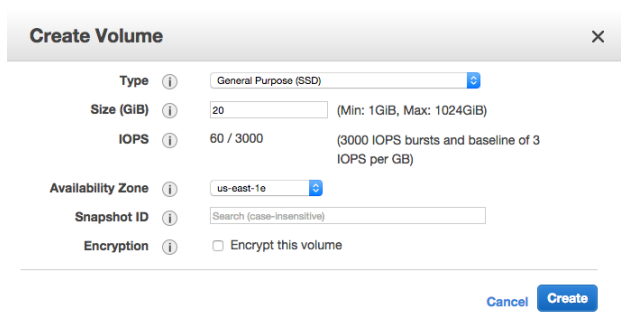


Click “Create Volume”.




Enter the desired size, and select the zone in which your instance is running. The volume and instance must be in the same zone. Otherwise, the volume cannot be attached to your instance.

Then click “Create”.

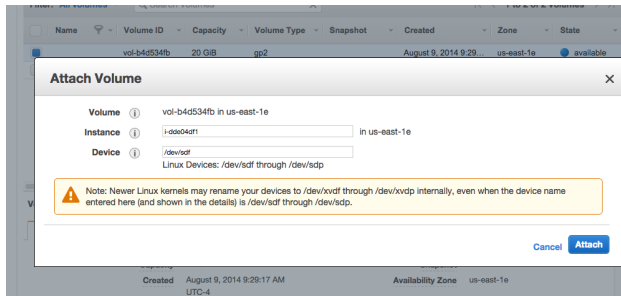


Wait for your volume to finish being created, then click “Attach Volume”.



Select the desired running instance. You may leave the Device at the default: /dev/sdf. This is the name that your EC2 instance can use to connect to the disk space. (If this is the second volume you want to attach to this instance, this may be different, like /dev/sdg.)

Click “Attach”.



When attachment is complete, connect to your instances via SSH.

If the volume is newly created, you must format the volume. **\*\*WARNING: ONLY DO THIS ONCE, WHEN YOU FIRST CREATE THE VOLUME. OTHERWISE, YOU WILL LOSE ALL YOUR DATA. YOU SHOULD NOT DO THIS IF THE VOLUME IS RESTORED FROM A SNAPSHOT AND YOU WANT TO USE THE DATA ON IT \*\***.

```
ubuntu@ip-10-31-232-122:~$ sudo mkfs -t ext4 /dev/xvdf
```

Note1: here we use “xvdf” instead of “sdf”.

Note2: The device may be different. If this is the second volume you attached to the same instance, the device ID may be /dev/xdg, so here just use “xvdg” instead of “sdg”. In this situation, for all the commands below, replace “/dev/xvdf” by “/dev/xvdg”.

Then, mount the volume. You’ll do this every time you attach the volume to an instance:

```
ubuntu@ip-10-31-232-122:~$ sudo mkdir /data
ubuntu@ip-10-31-232-122:~$ sudo mount /dev/xvdf /data
```

Your drive is now ready to use – it will be available under /data/.

## Detaching an Amazon EBS Volume

Any volumes you have attached will automatically detach when you shut down the instance. You can also stop all processes that are using the volume, change out of the directory, and type

```
ubuntu@ip-10-31-232-122:~$ sudo umount -d /dev/xvdf
```

and then detach the volume via the AWS Web site.

### 4.19.7 Using Amazon EBS Snapshots for sharing and backing up data

Now you have the Amazon EBS Volume to store the data for your instance to use. But it can only be attached to the EC2 instance you created. If you want to share the EBS Volume with the data with other colleagues so they can use the data on the EC2 instance they created, you need to create an Amazon EBS Snapshot and share it with other Amazon EC2 user. You can also create Amazon EBS Snapshots periodically to backup your Amazon EBS Volume.

In this tutorial, you will be instructed to create an Amazon EBS Snapshot from an Volume, to share the EBS Snapshot with others and to restore an Volume from an Snapshot.

## Prerequisites

This tutorial assumes you’ve completed the EC2 tutorial to set up an Amazon EBS Volume.




## Creating an Amazon EBS Snapshot from a Volume

Firstly open the Amazon EC2 console at [‘https://console.aws.amazon.com/ec2’](https://console.aws.amazon.com/ec2) and make sure it says North Virginia in the upper right.

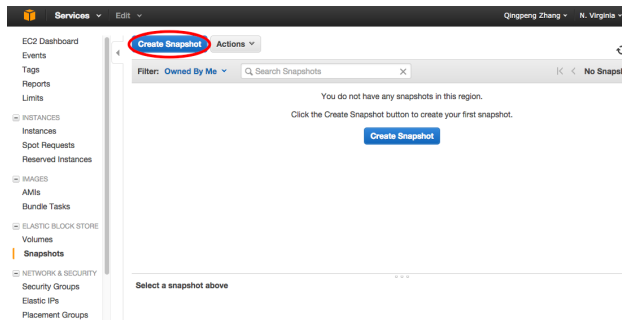
At the AWS Management Console, on the left menu bar, click “Volumes”.

Here you can see the 20GiB volume you created in the tutorial “Storing data persistently with Amazon EBS Volumes”.



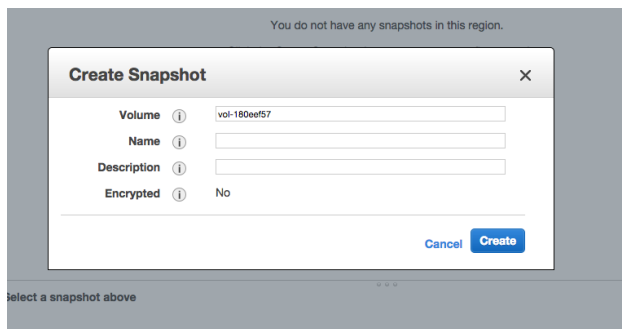
This tutorial will guide you to create a Amazon EBS Snapshot for this volume. In this example, the “Volume ID” of that volume is “vol-180eef57”. Record this ID, we will use it later.

Next, on the left menu bar, click “Snapshots” and click “Create Snapshot”.



Choose the Volume we want to create Snapshot for. (The ID is vol-180eef57 for this example, as we recored in last step.) You can enter the information for “Name” and “Description” as you like or just leave them blank.

Then click “Create”.



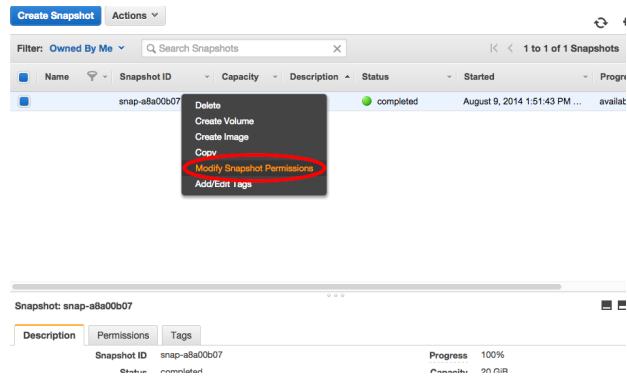
Ok. Now you have created a Snapshot from that 20G Volume.

## Sharing an Amazon EBS Snapshot

For now, the snapshot we just created is private, which can only be viewed and used by yourself.

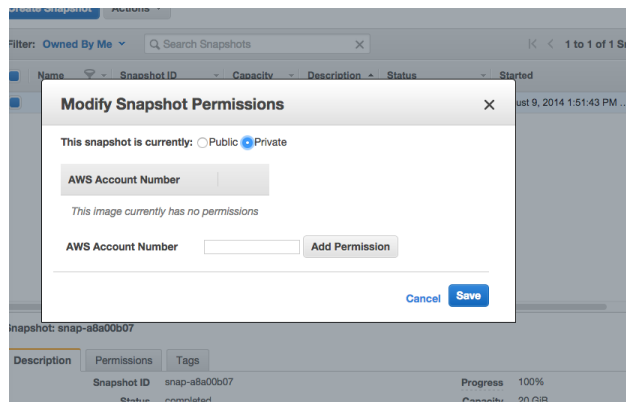
To make it public or share it with other Amazon EC2 user:

Right click the snapshot and click “Modify Snapshot Permissions”.



If you want to make this snapshot public, which means any Amazon EC2 user can have access to this snapshot and all the data in it, click “Public”.

If you just want to share this snapshot with specific person, like your colleague, keep the option as “Private” but put the AWS Account Number (which can be acquired by checking “Account Settings”) of the person in the box and click “Add Permission”.




Now you can share the “Snapshot ID” (snap-a8a00b07 in this example) to your colleague so they can restore an EBS Volume from this snapshot and use the data in it.

## Restoring an Amazon EBS Volume from a Snapshot

Your colleague shares an Amazon EBS Snapshot with you and you want to use the data on it. You did something terribly wrong to the data on your Volume and you want to restore the data from the backup Snapshot of that Volume. Under these circumstances, you want to restore an EBS Volume from a Snapshot.

It is similar to how you create a new EBS Volume. The only difference is that in the dialog box after you click “Create Volume” button, input the “Snapshot ID” of the snapshot you want to restore. Similarly, also select the zone in which your instance is running.



Ok, now you have the volume available to attach to your running instance.

### 4.19.8 Transferring Files between your laptop and Amazon instance

For linux/Unix/Mac system, we can use a command-line tool “scp” to transfer files between your laptop and Amazon instance. Also we can use a GUI tool “FileZilla” to do the transfer, which is more user-friendly.

#### Using scp to transfer data

“scp” means “secure copy”, which can copy files between computers on a network. You can use this tool in a Terminal on a Unix/Linux/Mac system.

To upload a file from your laptop to Amazon instance:

```
$scp -i ~/Desktop/amazon.pem ~/Desktop/MS115.fa ubuntu@ec2-54-166-128-20.compute-1.amazonaws.com:~/o
```

This command will upload a file - MS115.fa in your ~/Desktop/ folder of your laptop to folder ~/data/ on an Amazon instance. Note you still need to use the private key you used to connect to the Amazon instance with ssh. (In this example, it is the amazon.pem file in ~/Desktop/.

Note: You need to make sure that the user “ubuntu” has the permission to write in the target directory. In this example, if ~/data/ was created by user “ubuntu”, it should be fine.

Similarly, to download a file from Amazon instance to your laptop:

```
$scp -i ~/Desktop/amazon.pem ubuntu@ec2-54-166-128-20.compute-1.amazonaws.com:/data/ecoli_ref-5m-trim
```

This command will download a file /data/ecoli\_ref-5m-trim.fastq.gz from Amazon instance to your ~/Download folder in your laptop.


Note: You can use asterisk(\*) to download multiple files, like \*.fasta.gz.

#### Using FileZilla to transfer data

If you want a more user-friendly tool to transfer data, FileZilla is a good choice. It is free, it supports Windows/Linux/Mac systems, and it has a good user interface. It supports FTP, SFTP and other file transfer protocols.


Firstly, go to ‘<https://filezilla-project.org/>’ and click “Download FileZilla Client” button to download it.

The interface of FileZilla is like this:

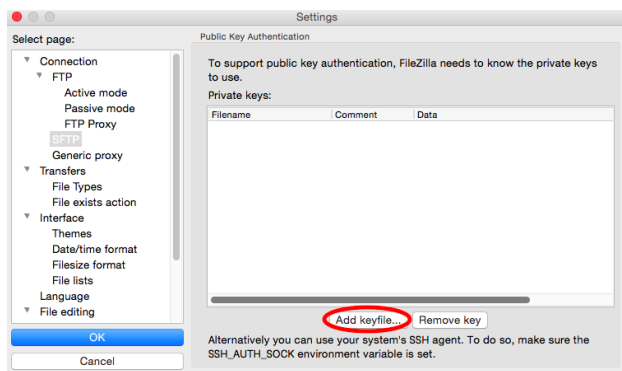


If you want to use FileZilla to upload to or download data from a normal FTP server if you have the user and password, just put the information in the “Host”, “Username”, “Password” box and connect. However for Amazon instance, we use key-pair to log in instead of password for better safety. So it is a little bit more complicated to configure.


Open “Settings” and click “SFTP”:



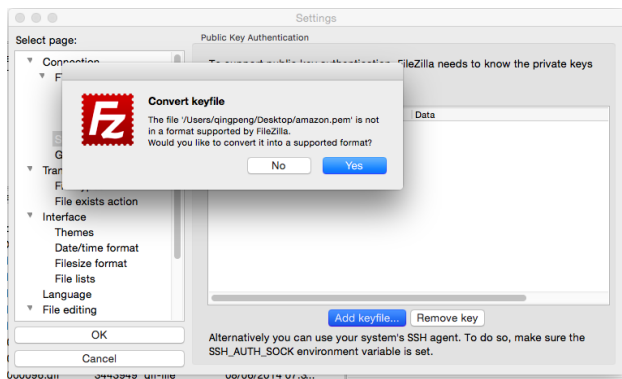
Click “Add keyfile...”:



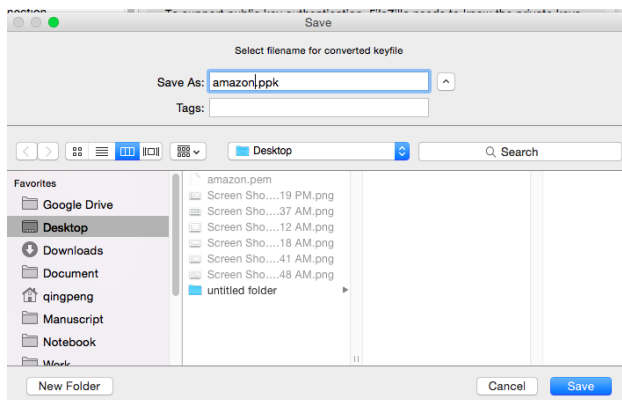
Then select the “.pem” file you used to connect to Amazon instance with ssh.



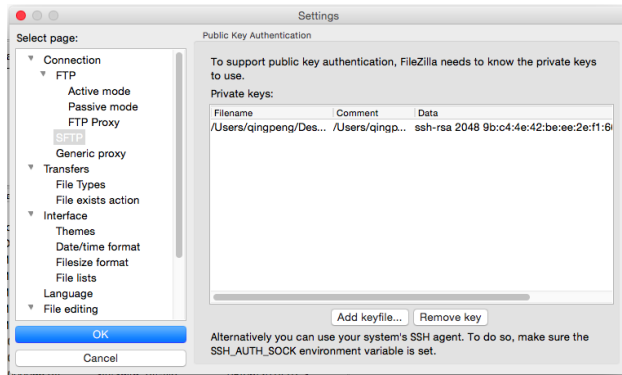
There is a dialog box to ask you if you want to convert the “.pem” file into a supported format. Click “Yes”.



Name it with extension as “.ppk” and save it.

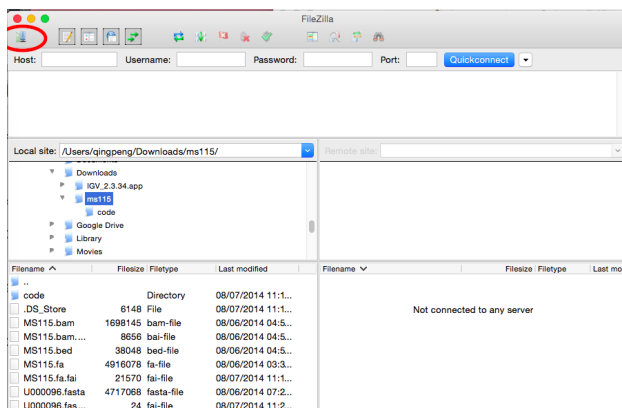


You will see the a private key has been added.

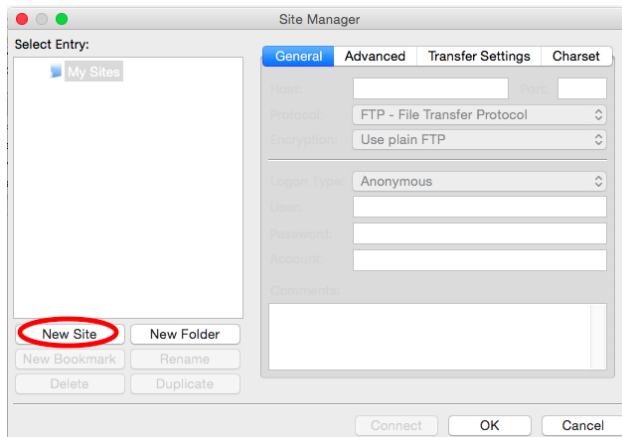


Close “Settings” and go back to the main interface.

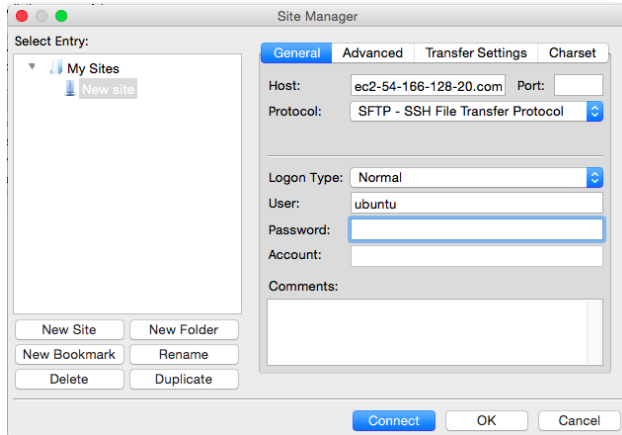
Click button to open the site manager.




Click “New Site”.



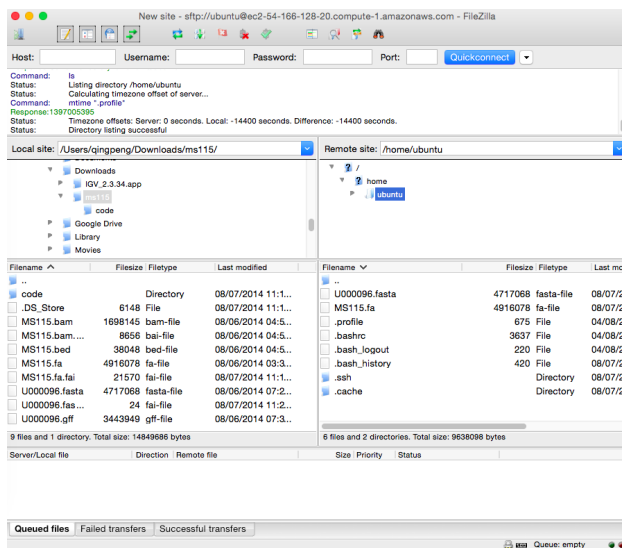
Put the Amazon instance URL like ec2-54-166-128-20.compute-1.amazonaws.com in the “Host” box. Set “Protocol” as “SFTP”, “Logon Type” as “Normal”, “User” as “ubuntu” and leave “Password” as blank. Then click “Connect”.



There will be a dialogue box to ask you about “Unknown host key”, just click “Ok”.



All right. Now you have logged in the Amazon instance. You can drag and drop to transfer the files between the remote machine and your local laptop.



#### 4.19.9 Uploading files to Amazon S3 to share

“Amazon Simple Storage Service(S3) is storage for the Internet.” It is designed to make web-scale computing easier for developers with a highly scalable, reliable, fast, inexpensive data storage infrastructure. Companies like Dropbox, Pinterest, Tumblr store their data on the S3 servers.

(Learn more from <http://docs.aws.amazon.com/AmazonS3/latest/dev/Welcome.html>)

For personal users like us, S3 is also a reliable, inexpensive service to store data or share data around the world.

This tutorial will instruct you to upload files to Amazon S3 and share them.

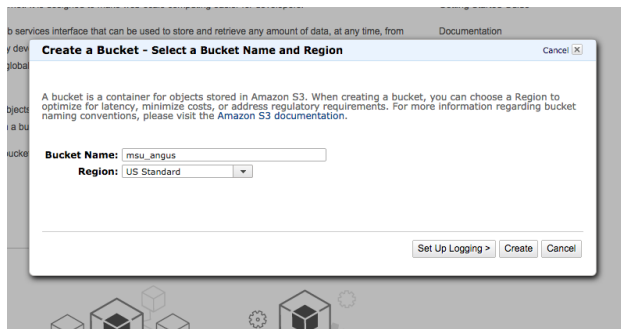
### Uploading files to Amazon S3

Go to Amazon S3 Console: <https://console.aws.amazon.com/s3/>

Click “Create Bucket”.

In the dialog box, in the Bucket Name box, enter a bucket name. Here a “bucket” is like a “folder” in concept. The bucket name must be unique across all existing bucket names in Amazon S3. You can not change the name after you create a bucket. Also, the bucket name you chose here will be visible in the URL that points to the objects(files) stored in the bucket. So make sure the bucket name you choose is appropriate.

Leave the “Region” as default and click “Create”.



Next you can click the bucket we just created, and it is an empty one.


Now we can add files into this bucket(folder) by clicking “Upload”.



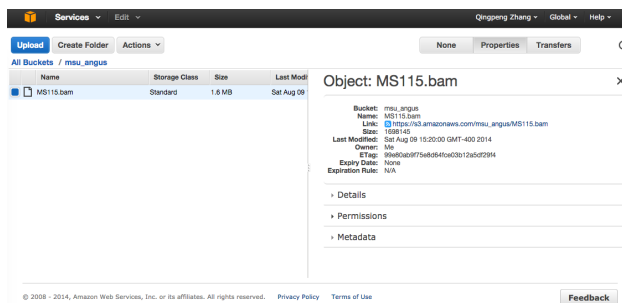
After you select the files and upload them successfully, you can see them in the current bucket. Right click the file, you can manipulate it in many ways you like, like “Rename”, “Delete”, “Download”, and others.

Out of them, “Make Public” is what you want to click to share the file.





When it is done, highlight the file you just shared and click “Properties” on the upper right. You will see the link of this file, like [https://s3.amazonaws.com/msu\\_angus/MS115.bam](https://s3.amazonaws.com/msu_angus/MS115.bam) in this example.



So that's the public URL of that file. You can share this link to the person you want to share this file with.

## Downloading files from Amazon S3

You can use any internet browser to download the file from Amazon S3 directly with the public URL.

In command-line environment, you can use curl to download the file:

```
$curl -O https://s3.amazonaws.com/msu_angus/MS115.bam
```

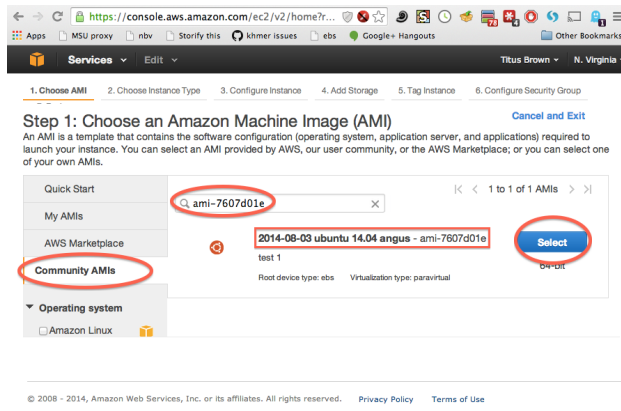
This command will download the file into current directory.

### 4.19.10 Starting up a custom operating system

The instructions in *Start up an EC2 instance* tell you how to start up a machine with Ubuntu Linux version 14.04 on it, but that machine comes with very little software installed. For anything where you are executing actual analyses, you're going to want to have a bunch of basic software installed.

Therefore, we make custom versions of Ubuntu available as well, that come with some software pre-installed. (See *Technical guide to the ANGUS course* for a list of the packages installed on the ANGUS custom AMI.)

To boot these, go to EC2/Launch and select “Community AMIs” instead of the default Quick Start; then type in the AMI number or name given to you in the tutorial. Below is a screenshot of an example for ‘ami-7606d01e’. Then proceed with the rest of *Start up an EC2 instance*.



### 4.19.11 Technical guide to the ANGUS course

#### Packages we install

Install:

```
apt-get update
apt-get -y install screen git curl gcc make g++ python-dev unzip \
    default-jre pkg-config libncurses5-dev r-base-core \
    r-cran-gplots python-matplotlib sysstat
```

## 4.20 Instructor's Guide to ANGUS Materials

The main repository is here: <https://github.com/ngs-docs/angus>. Please try to keep everything in there as much as possible.

For 2014, contribute to the branch '2014'.

We use [Sphinx](#) to build the site from multiple files, and each file is written in [reStructuredText](#).

Merges to the '2014' branch are automatically built and posted by [readthedocs.org](#) at <http://angus.readthedocs.org/en/2014/>

You can use pull requests OR you can just send Titus your github ID and he will give you merge privileges. For the first few modifications we would still suggest using pull requests just so you can get the hang of reST.

Put static files that you do not want interpreted by Sphinx (e.g. presentation PDFs) in the `files/` directory.

### 4.20.1 Licensing

Everything you do must be releasable under CC0 except for your presentation slides, which must be accessible and posted somewhere reasonably permanent (in our repo, on slideshare, etc) but can be under whatever license you choose.

## 4.21 Workshop Code of Conduct

All attendees, speakers, sponsors and volunteers at our workshop are required to agree with the following code of conduct. Organisers will enforce this code throughout the event. We are expecting cooperation from all participants to

help ensuring a safe environment for everybody.

**tl; dr: Don't be a jerk.**

### 4.21.1 Need Help?

You can reach the course director, Titus Brown, at [titus@idyll.org](mailto:titus@idyll.org) or via the cell phone number on the course info page. You can also talk to any of the instructors or TAs if you need immediate help.

Judi Brown Clarke, [jbc@msu.edu](mailto:jbc@msu.edu), is the person to contact if Titus is not available or there are larger problems; she is available via phone at 517.353.5985.

### 4.21.2 The Quick Version

Our workshop is dedicated to providing a harassment-free workshop experience for everyone, regardless of gender, age, sexual orientation, disability, physical appearance, body size, race, or religion (or lack thereof). We do not tolerate harassment of workshop participants in any form. Sexual language and imagery is not appropriate for any workshop venue, including talks, workshops, parties, Twitter and other online media. Workshop participants violating these rules may be sanctioned or expelled from the workshop *without a refund* at the discretion of the workshop organisers.

### 4.21.3 The Less Quick Version

Harassment includes offensive verbal comments related to gender, age, sexual orientation, disability, physical appearance, body size, race, religion, sexual images in public spaces, deliberate intimidation, stalking, following, harassing photography or recording, sustained disruption of talks or other events, inappropriate physical contact, and unwelcome sexual attention.

Participants asked to stop any harassing behavior are expected to comply immediately.

If a participant engages in harassing behavior, the workshop organisers may take any action they deem appropriate, including warning the offender or expulsion from the workshop with no refund.

If you are being harassed, notice that someone else is being harassed, or have any other concerns, please contact a member of workshop staff immediately.

Workshop instructors and TAs will be happy to help participants contact KBS security or local law enforcement, provide escorts, or otherwise assist those experiencing harassment to feel safe for the duration of the workshop. We value your attendance.

We expect participants to follow these rules at workshop and workshop venues and workshop-related social events.

This work is licensed under a [Creative Commons Attribution 3.0 Unported License](https://creativecommons.org/licenses/by/3.0/).

---

This Code of Conduct taken from <http://confcodeofconduct.com/>. See

<http://www.ashedryden.com/blog/codes-of-conduct-101-faq>

for more information on codes of conduct.