

ilifu Online Training

Dane Kennedy <dane@idia.ac.za>

User Training Workshop – Advanced Training

13 April 2021

Context

- Environment modules (Imod)
- Python virtual environments
- R and RStudio

http://docs.ilifu.ac.za/#/tech_docs/software_environments?id=environment-modules

- Configure a working environment
- Setup PATH, LD_LIBRARY_PATH, alias etc.

```
$ module avail
----- /software/modules/common -----
R/3.6.3      openmpi/4.0.3      python/3.8.2

----- /usr/share/lmod/lmod/modulefiles -----
Core/lmod/6.6  Core/settarg/6.6

$ R
Command 'R' not found, but can be installed with:
...
$ module add R/3.6.3
$ R
R version 3.6.3 (2020-02-29) -- "Holding the
Windsock"
Copyright (C) 2020 The R Foundation for Statistical
Computing
Platform: x86_64-pc-linux-gnu (64-bit)
...
```

- virtualenv
- <https://virtualenv.pypa.io/en/latest/>
- Isolated Python environment
- Less risk of conflicts occurring with pip install --user
- Similar to venv (python -m venv)
- Can customize which os python is used: python2.7, python3+
- Limited by os libraries

```
virtualenv --help
```

```
virtualenv /path/to/virtual_environment
```

```
--python           The Python interpreter to use
```

```
--system-site-packages Give the virtual environment access to  
the global site-packages
```

```
virtualenv /path/to/virtual_environment
```

Example:

```
virtualenv ~/.jupenv  
source ~/.jupenv/bin/activate  
which python  
  □ ~/.jupenv/bin/bin/python  
pip install scikit-learn  
deactivate (to exit)
```


Python virtualenv as a Jupyter kernel

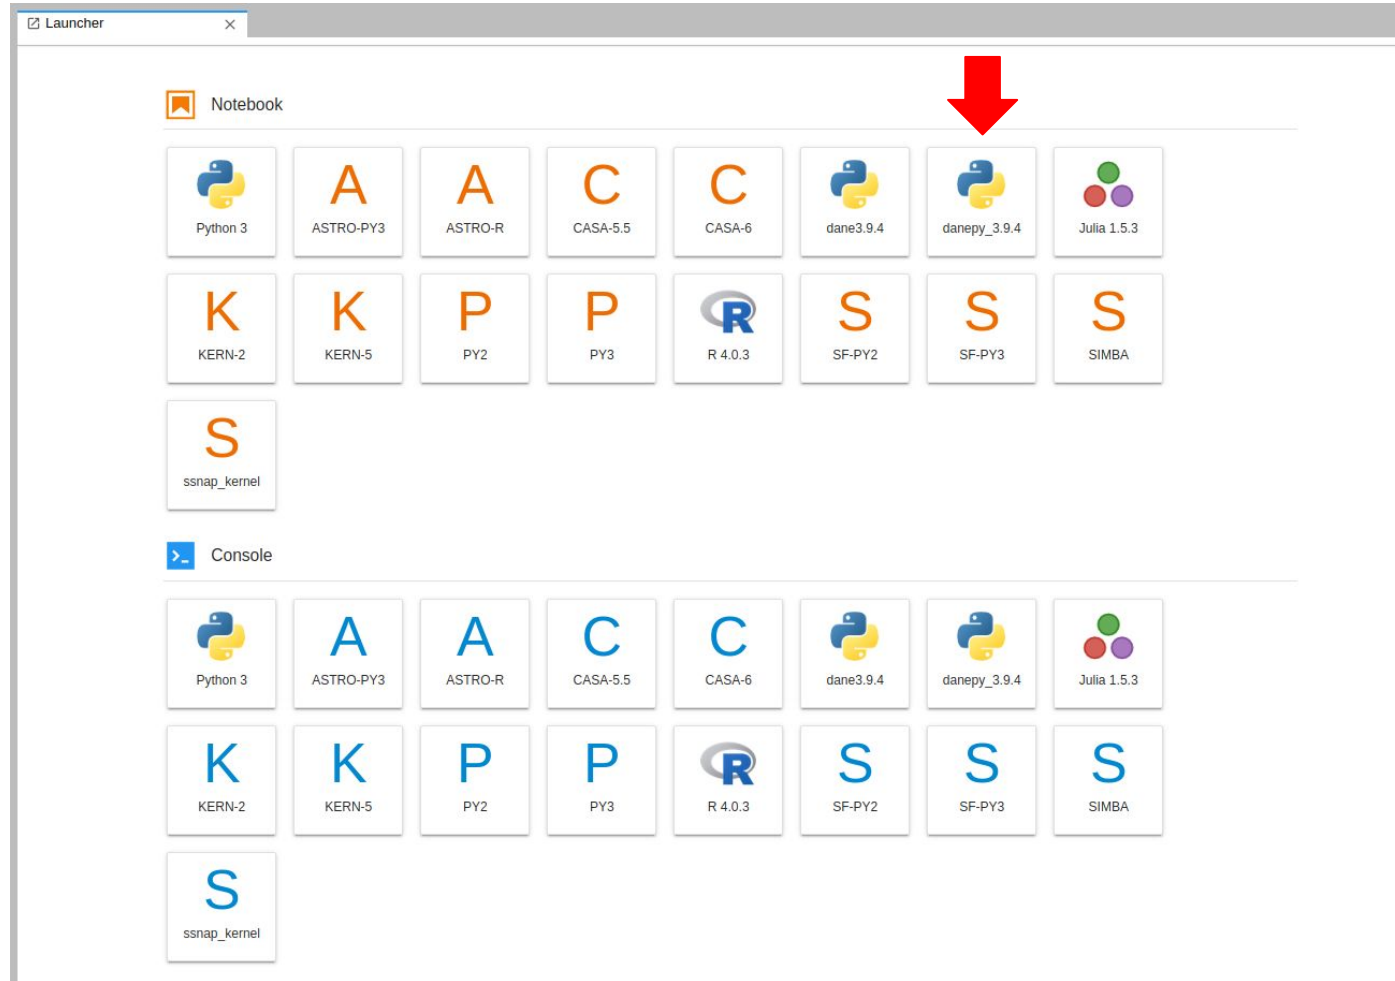
Once the virtual environment is active:

```
ipython kernel install --name "<kernel_name>" --user
```

Example:

```
. ./jupyter/bin/activate  
pip install jupyter  
ipython kernel install --name "danepy_3.9.4" --user  
Installed kernelspec danepy_3.9.4 in /users/dane/.local/share/jupyter/kernels/dane_py3.9.4
```

Creates the kernel.json file at /users/dane/.local/share/jupyter/kernels/dane_py3.9.4/kernel.json



R in a slurm job

http://docs.ilifu.ac.za/#/tech_docs/software_environments?id=running-r

```
dane@slurm-login:~$ srun --nodes=1
--tasks=2 --mem=4g --time 08:00:00 --pty
bash
dane@compute-003:~$ module add R/4.0.3
dane@compute-003:~$ R
R version 4.0.3 (2020-10-10) --
"Bunny-Wunnies Freak Out"
...
```

```
#!/bin/bash
#SBATCH --nodes=1
#SBATCH --cpus-per-task=1
#SBATCH --mem=4GB
#SBATCH --job-name=covid19_analysis
#SBATCH --time=00:10:00
```

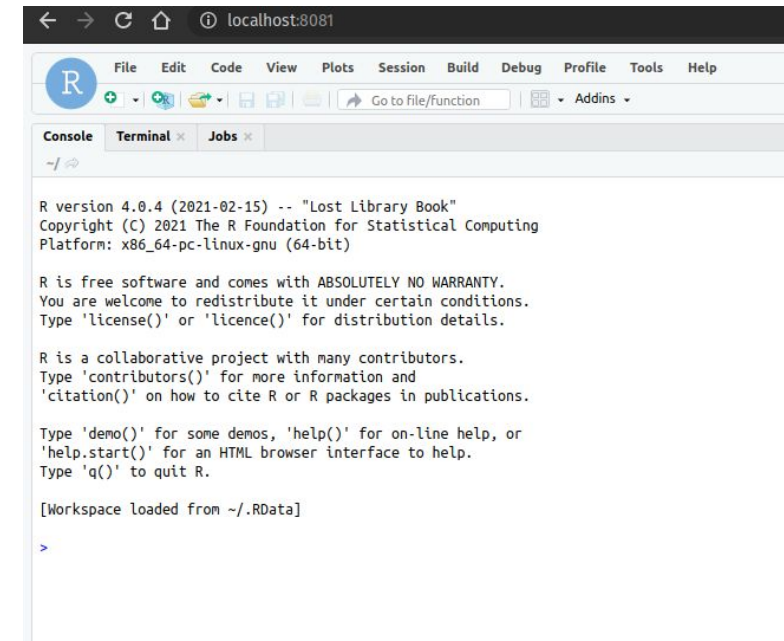
```
module add R/4.0.3
Rscript covid19.R
```

```
library('covid19.analytics')
write.csv(covid19.data(), 'covid19.csv')
```

http://docs.ilifu.ac.za/#/tech_docs/software_environments?id=running-rstudio

```
dane@slurm-login:~$ srun --nodes=1
--tasks=2 --mem=4g --time 08:00:00 --pty
bash
dane@compute-003:~$ module add
R/RStudio1.2.5042-R4.0.4
dane@compute-003:~$ rstudio
The environment variable RSTUDIO_PASSWORD
was not set, so your password has been
chosen for you. It's: ****
Running rserver on port 59447
To connect to this server run this on your
local machine:
    ssh -A dane@compute-003 -o ...
then vist http://localhost:8081 in your
browser and use the username "dane" to
login with the password "****"
```

```
dane@localhost:~$ ssh -A dane@compute-003
-o ...
```

A screenshot of the RStudio web interface running in a browser. The browser address bar shows 'localhost:8081'. The RStudio interface includes a menu bar (File, Edit, Code, View, Plots, Session, Build, Debug, Profile, Tools, Help), a toolbar with icons for file operations, and a main console area. The console displays the R startup message for version 4.0.4 (2021-02-15), including copyright information and usage instructions. The workspace is loaded from '~/.RData'.

```
localhost:8081
R
File Edit Code View Plots Session Build Debug Profile Tools Help
Go to file/function Addins
Console Terminal Jobs
~/
R version 4.0.4 (2021-02-15) -- "Lost Library Book"
Copyright (C) 2021 The R Foundation for Statistical Computing
Platform: x86_64-pc-linux-gnu (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.

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.

[Workspace loaded from ~/.RData]
>
```

Installing R packages

http://docs.ilifu.ac.za/#/tech_docs/software_environments?id=running-r

```
install.packages()
```

versus

```
BiocManager::install()
```

