Installing Jupyter Notebooks:

Install the latest version of the anaconda python distribution:
https://www.anaconda.com/distribution/

Open Anaconda Navigator and you should notice that Jupyter is already installed. If not, click Install in the Home menu.

Opening Jupyter Notebooks:

To open Jupyter, you can use the Anaconda Navigator GUI or call “jupyter notebook” at the command line.

Jupyter will open as a new tab in a browser window, which displays http://localhost:8888/tree.

To create a new notebook, select a folder and create a notebook via the “New” dropdown on the top right of the screen.

Here’s an example notebook you could create:

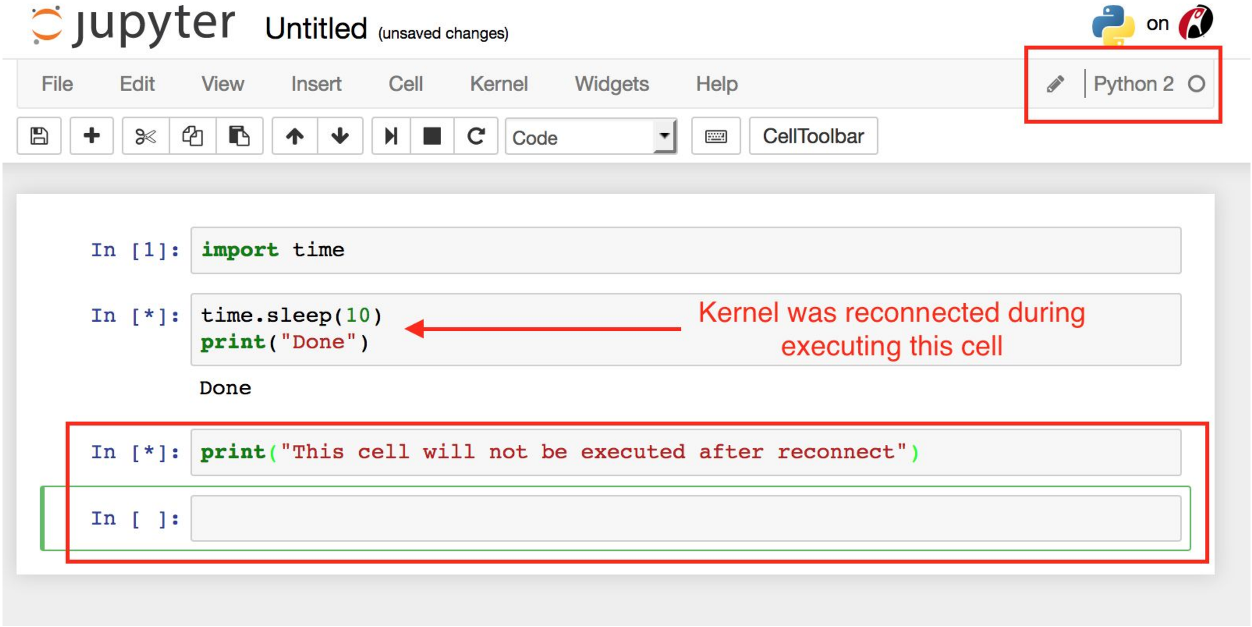

```
In [1]: import time

In [*]: time.sleep(10)
        print("Done")

        Done

In [*]: print("This cell will not be executed after reconnect")

In [ ]:
```

Kernel was reconnected during executing this cell

How Jupyter notebook works:

Think of Jupyter like a text editor that keeps track of your code. When you “start the kernel”, Jupyter starts up a python process, called the kernel, that runs in the background. When you execute a cell, Jupyter sends your code to the python process, runs it, and displays the results below the cell. Hence you can run and execute your cells in any order.

This can get confusing, since if you run cell 2 and then cell 1, then variables accessible in cell 1 can be set by cell 2. The python process only knows about the order of execution, not the order of cells in the notebook. After running a cell, a number in brackets appears to the left of the cell indicating the order in which cells were run to help you remember.

If you restart the Kernel, it restarts the python process, deleting all variables in memory (but not your code!) so you can run it again afresh. It is always a good idea to restart your kernel and rerun all cells before submitting homework or when you encounter a bug; this can help you check that you haven’t executed cells in a weird order that introduces bugs.