Chapter 2: Designing a GPU Computing Strategy

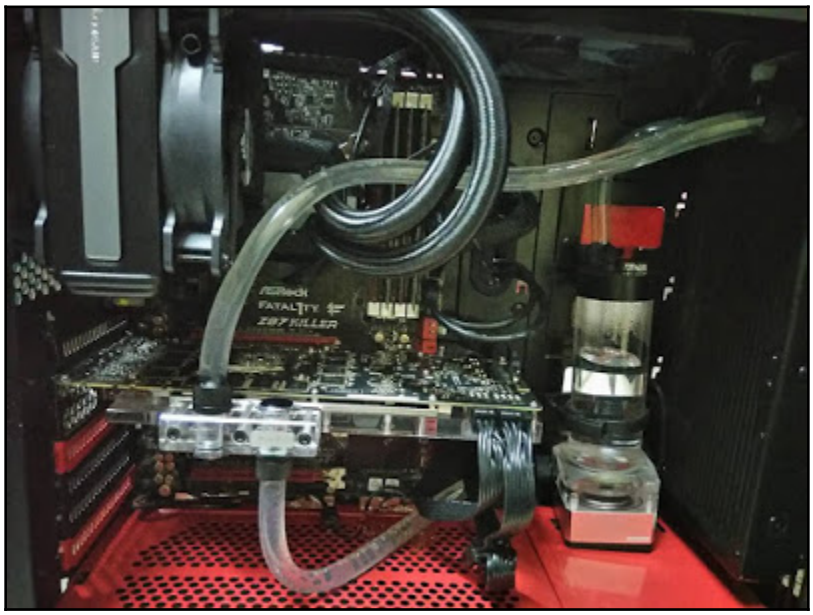

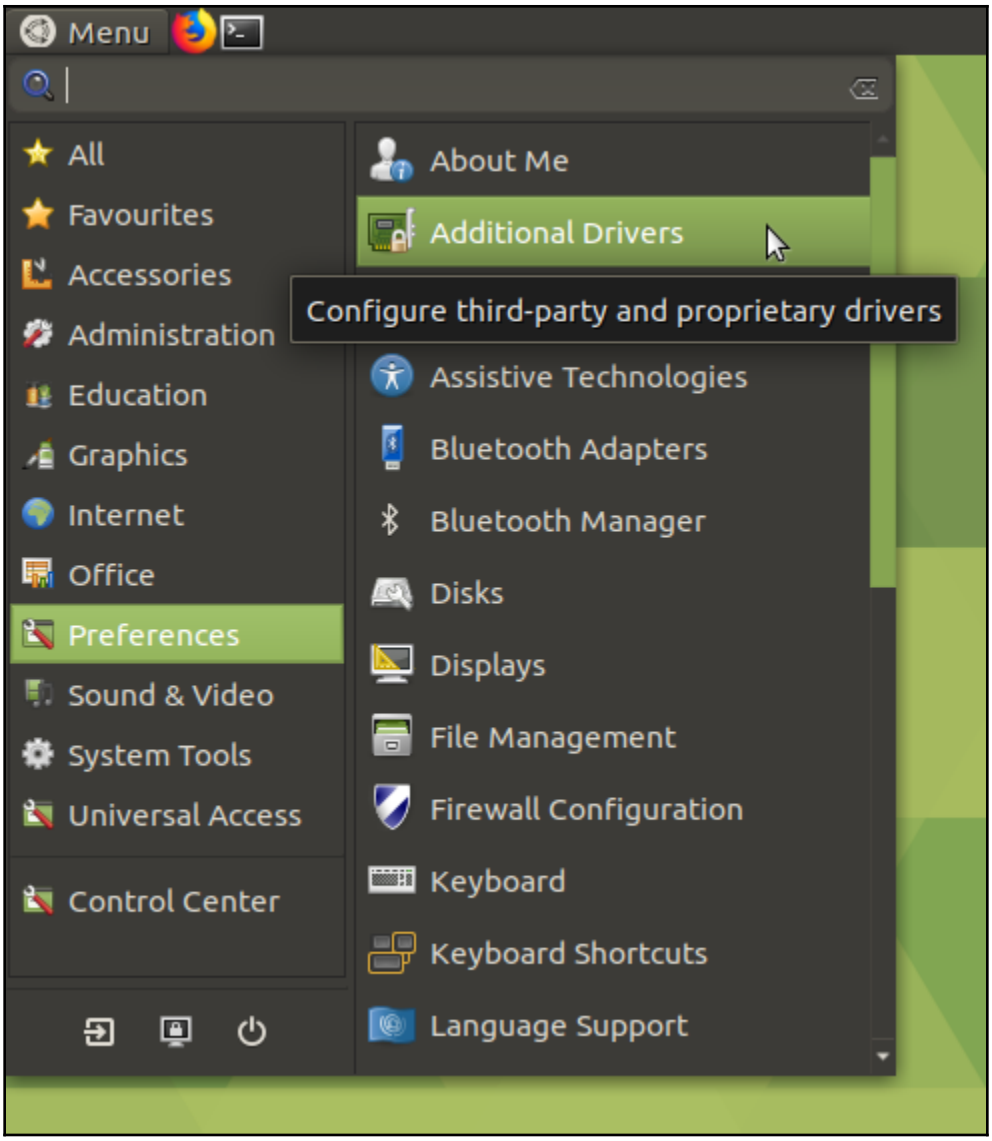
Menu
All
Favourites
Accessories
Administration
Education
Graphics
Internet
Office
Preferences
Sound & Video
System Tools
Universal Access
Control Center
About Me
Additional Drivers
Configure third-party and proprietary drivers
Assistive Technologies
Bluetooth Adapters
Bluetooth Manager
Disks
Displays
File Management
Firewall Configuration
Keyboard
Keyboard Shortcuts
Language Support

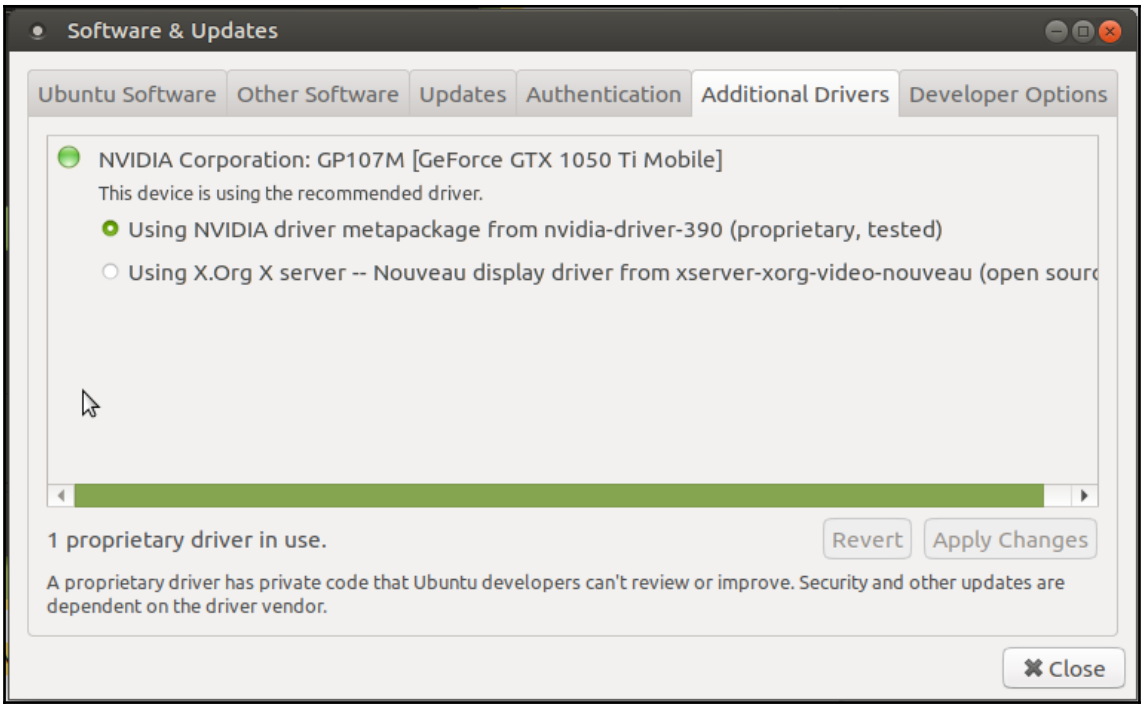
Software & Updates
Ubuntu Software
Other Software
Updates
Authentication
Additional Drivers
Developer Options
NVIDIA Corporation: GP107M [GeForce GTX 1050 Ti Mobile]
This device is using the recommended driver.
Using NVIDIA driver metapackage from nvidia-driver-390 (proprietary, tested)
Using X.Org X server -- Nouveau display driver from xserver-xorg-video-nouveau (open sourc
1 proprietary driver in use.
Revert
Apply Changes
A proprietary driver has private code that Ubuntu developers can't review or improve. Security and other updates are dependent on the driver vendor.
Close

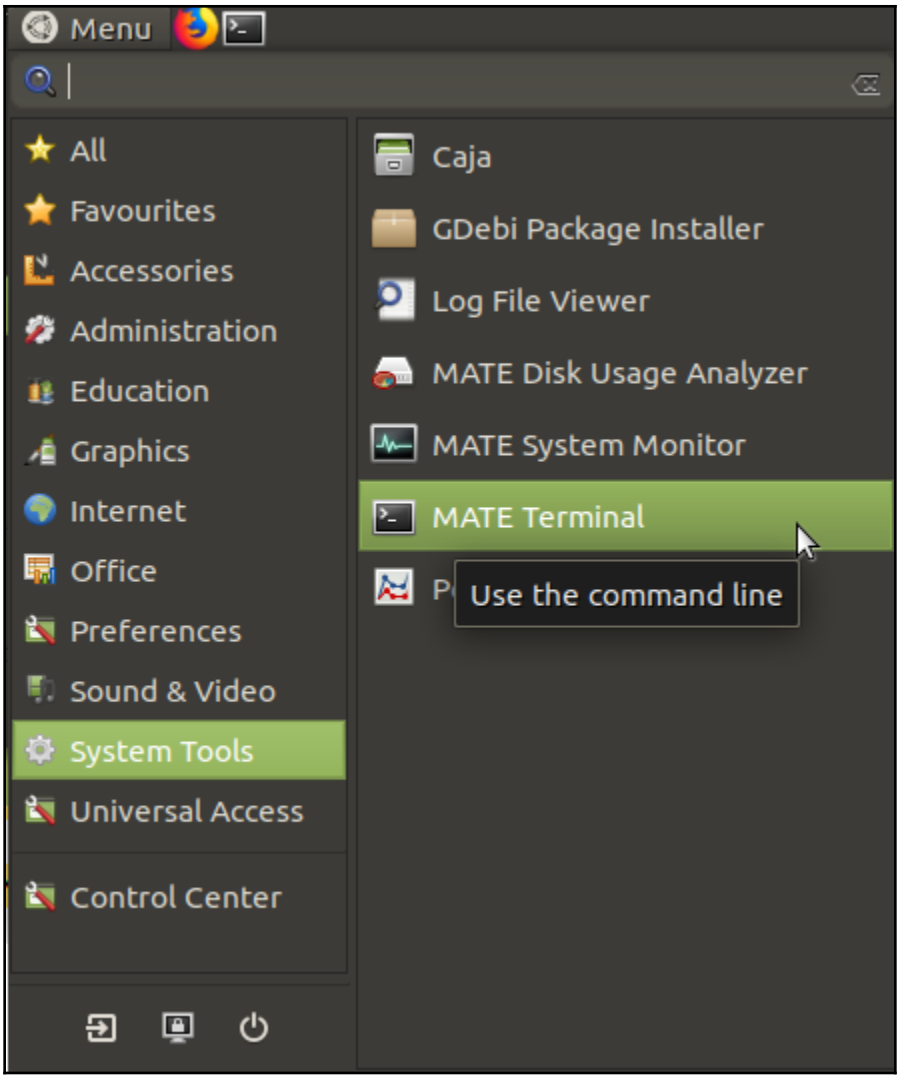
Menu
All
Favourites
Accessories
Administration
Education
Graphics
Internet
Office
Preferences
Sound & Video
System Tools
Universal Access
Control Center
Caja
GDebi Package Installer
Log File Viewer
MATE Disk Usage Analyzer
MATE System Monitor
MATE Terminal
Use the command line

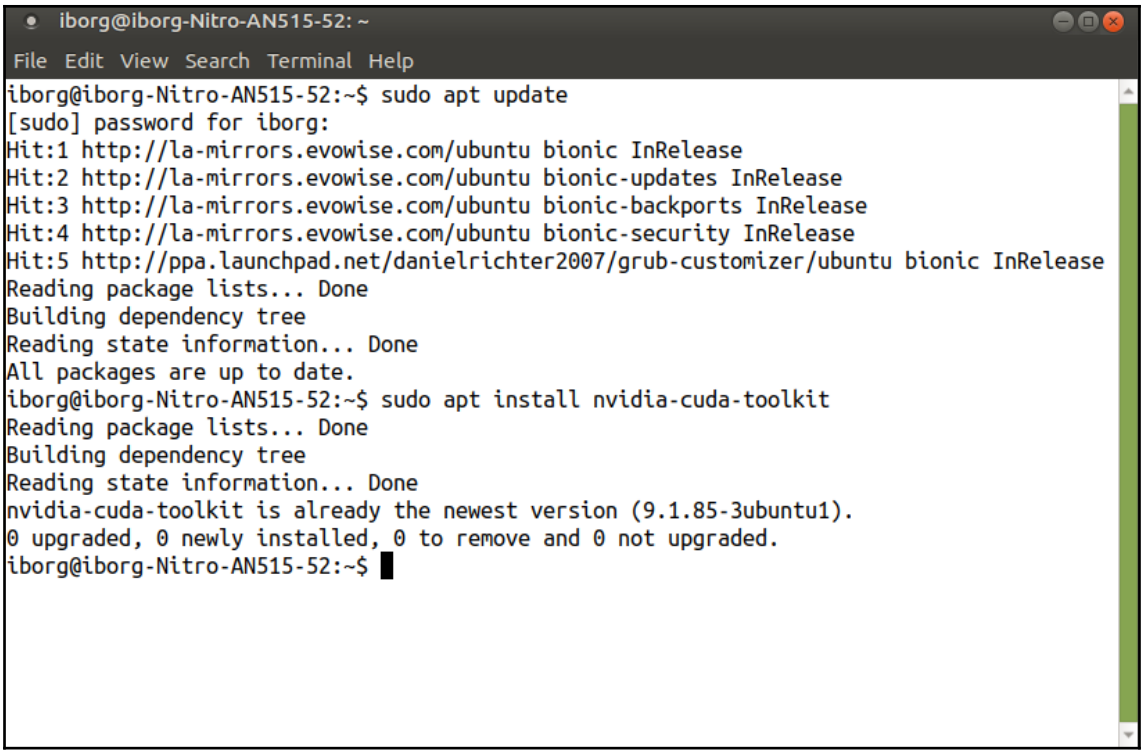
iborg@iborg-Nitro-AN515-52: ~
File Edit View Search Terminal Help
iborg@iborg-Nitro-AN515-52:~$ sudo apt update
[sudo] password for iborg:
Hit:1 http://la-mirrors.evowise.com/ubuntu bionic InRelease
Hit:2 http://la-mirrors.evowise.com/ubuntu bionic-updates InRelease
Hit:3 http://la-mirrors.evowise.com/ubuntu bionic-backports InRelease
Hit:4 http://la-mirrors.evowise.com/ubuntu bionic-security InRelease
Hit:5 http://ppa.launchpad.net/danielrichter2007/grub-customizer/ubuntu bionic InRelease
Reading package lists... Done
Building dependency tree
Reading state information... Done
All packages are up to date.
iborg@iborg-Nitro-AN515-52:~$ sudo apt install nvidia-cuda-toolkit
Reading package lists... Done
Building dependency tree
Reading state information... Done
nvidia-cuda-toolkit is already the newest version (9.1.85-3ubuntu1).
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
iborg@iborg-Nitro-AN515-52:~$

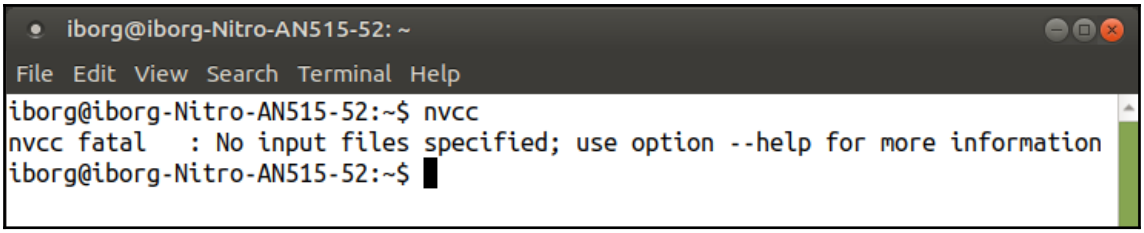
iborg@iborg-Nitro-AN515-52: ~
File Edit View Search Terminal Help
iborg@iborg-Nitro-AN515-52:~$ nvcc
nvcc fatal : No input files specified; use option --help for more information
iborg@iborg-Nitro-AN515-52:~$

Chapter 3: Setting Up a GPU Computing Platform with NVIDIA and AMD

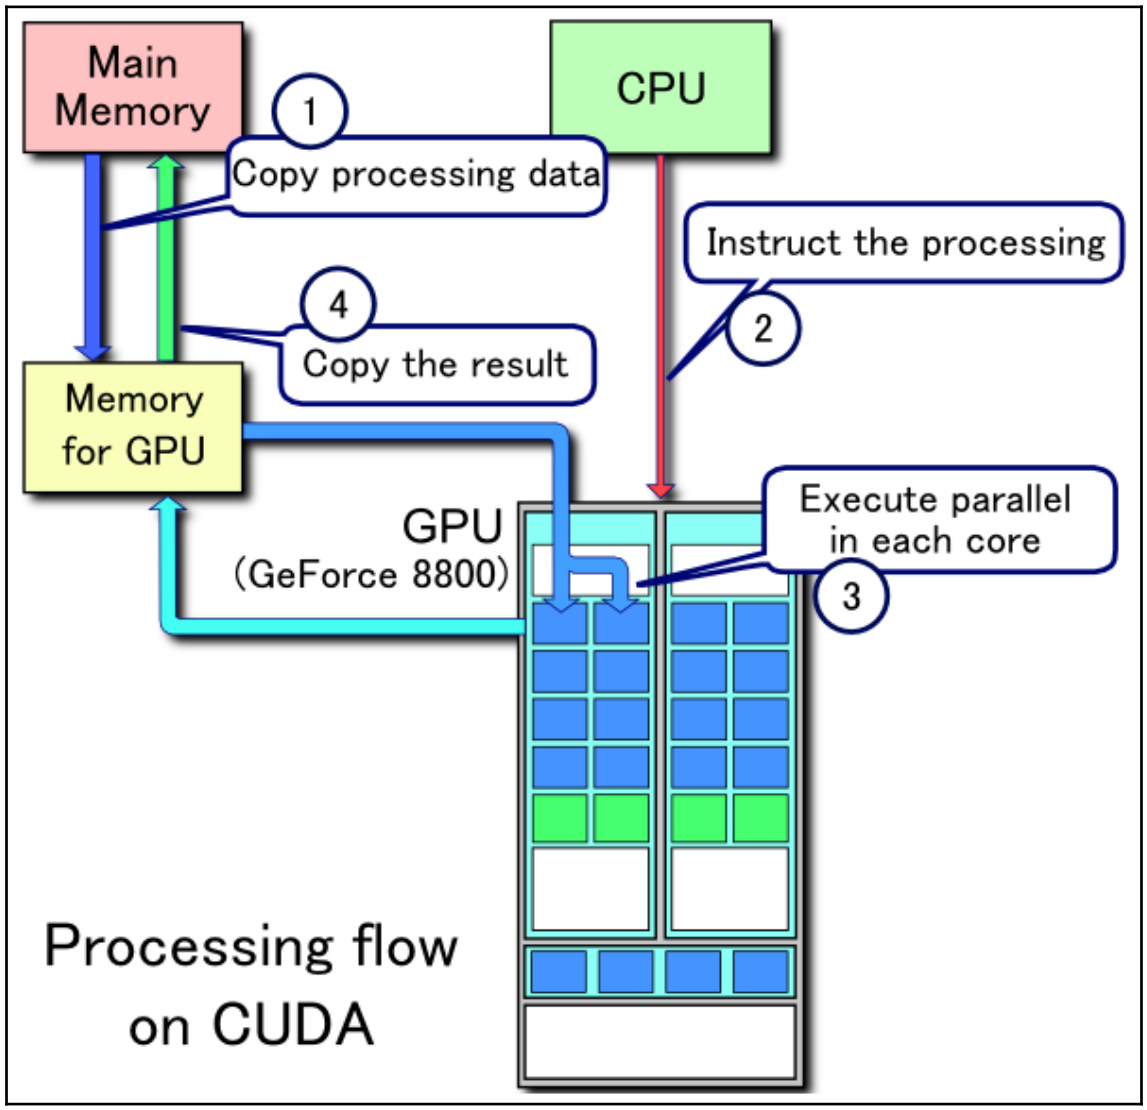

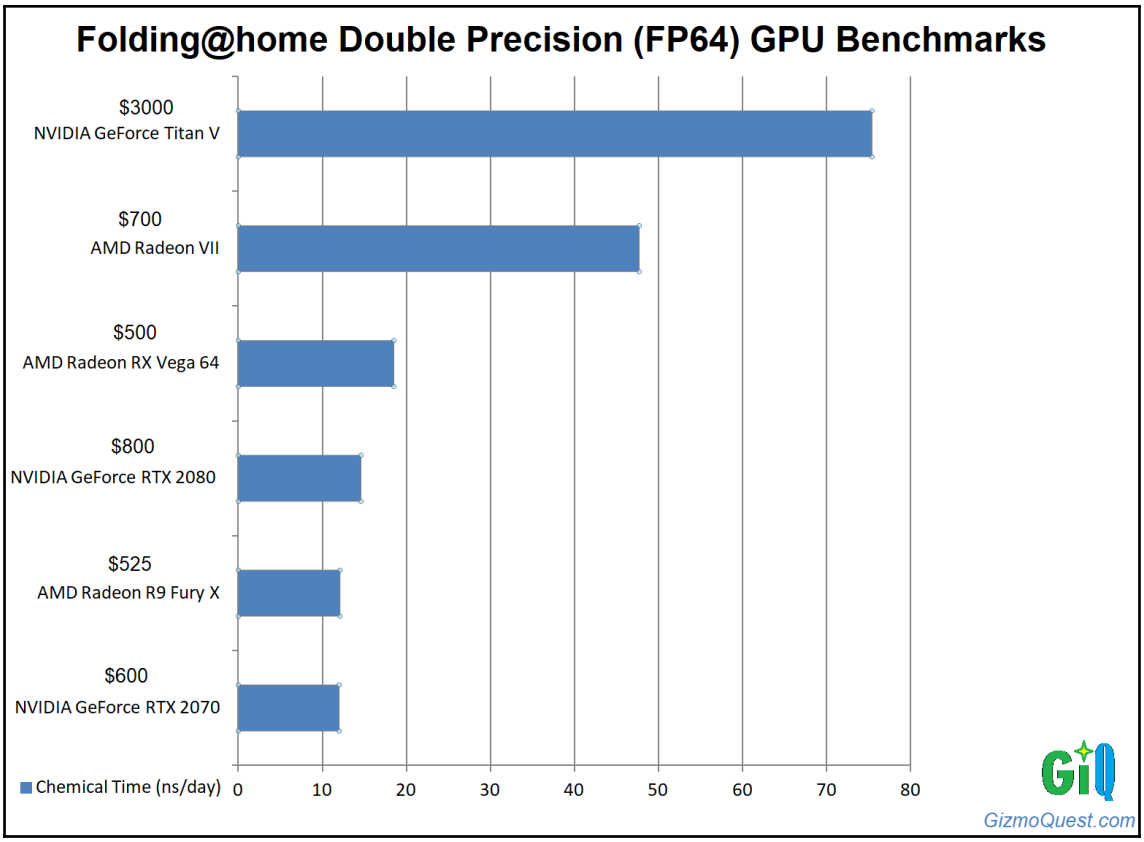
Folding@home Double Precision (FP64) GPU Benchmarks
$3000
NVIDIA GeForce Titan V
$700
AMD Radeon VII
$500
AMD Radeon RX Vega 64
$800
NVIDIA GeForce RTX 2080
$525
AMD Radeon R9 Fury X
$600
NVIDIA GeForce RTX 2070
Chemical Time (ns/day)
0
10
20
30
40
50
60
70
80
GizmoQuest.com

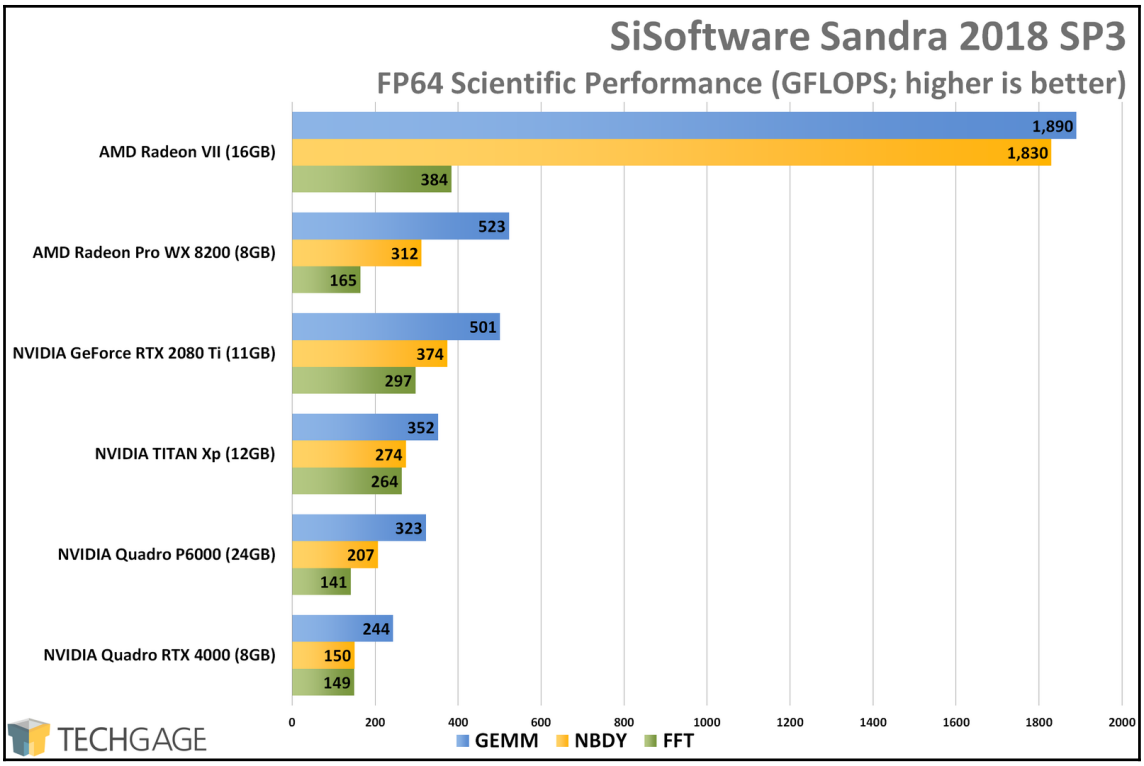
SiSoftware Sandra 2018 SP3
FP64 Scientific Performance (GFLOPS; higher is better)
AMD Radeon VII (16GB)
1,890
1,830
384
AMD Radeon Pro WX 8200 (8GB)
523
312
165
NVIDIA GeForce RTX 2080 Ti (11GB)
501
374
297
NVIDIA TITAN Xp (12GB)
352
274
264
NVIDIA Quadro P6000 (24GB)
323
207
141
NVIDIA Quadro RTX 4000 (8GB)
244
150
149
0
200
400
600
800
1000
1200
1400
1600
1800
2000
GEMM
NBDY
FFT
TECHGAGE

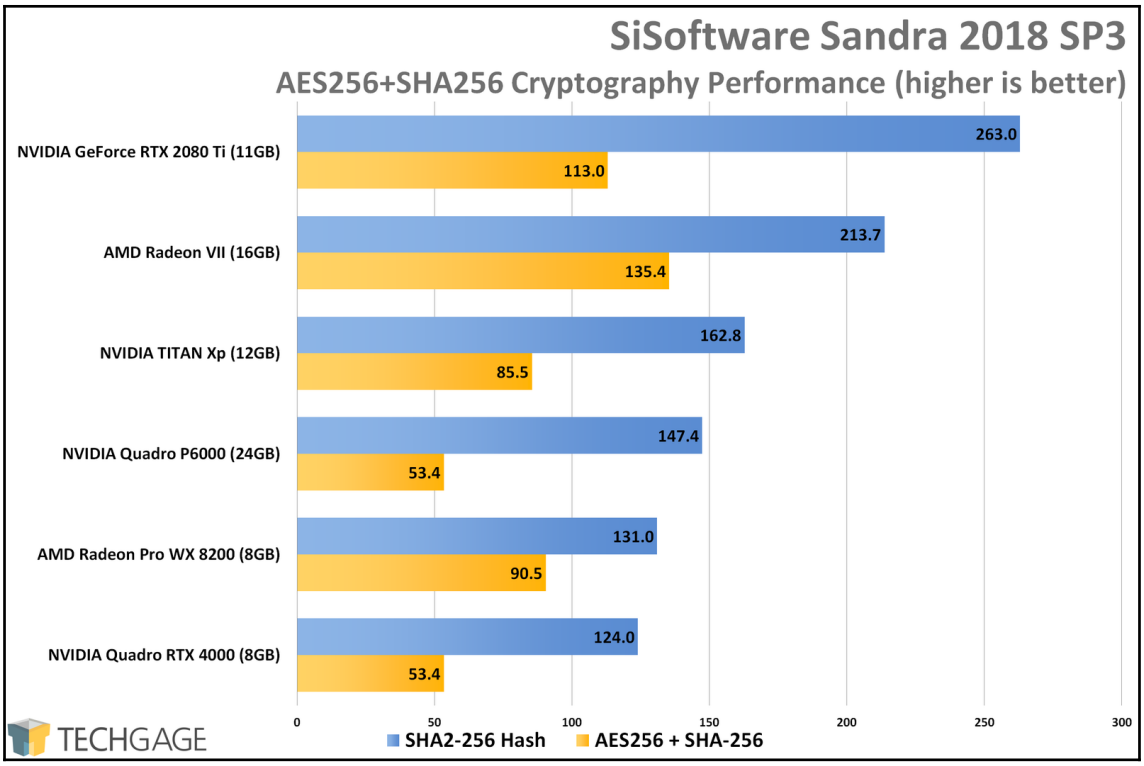
SiSoftware Sandra 2018 SP3
AES256+SHA256 Cryptography Performance (higher is better)
NVIDIA GeForce RTX 2080 Ti (11GB)
263.0
113.0
AMD Radeon VII (16GB)
213.7
135.4
NVIDIA TITAN Xp (12GB)
162.8
85.5
NVIDIA Quadro P6000 (24GB)
147.4
53.4
AMD Radeon Pro WX 8200 (8GB)
131.0
90.5
NVIDIA Quadro RTX 4000 (8GB)
124.0
53.4
0
50
100
150
200
250
300
TECHGAGE
SHA2-256 Hash
AES256 + SHA-256

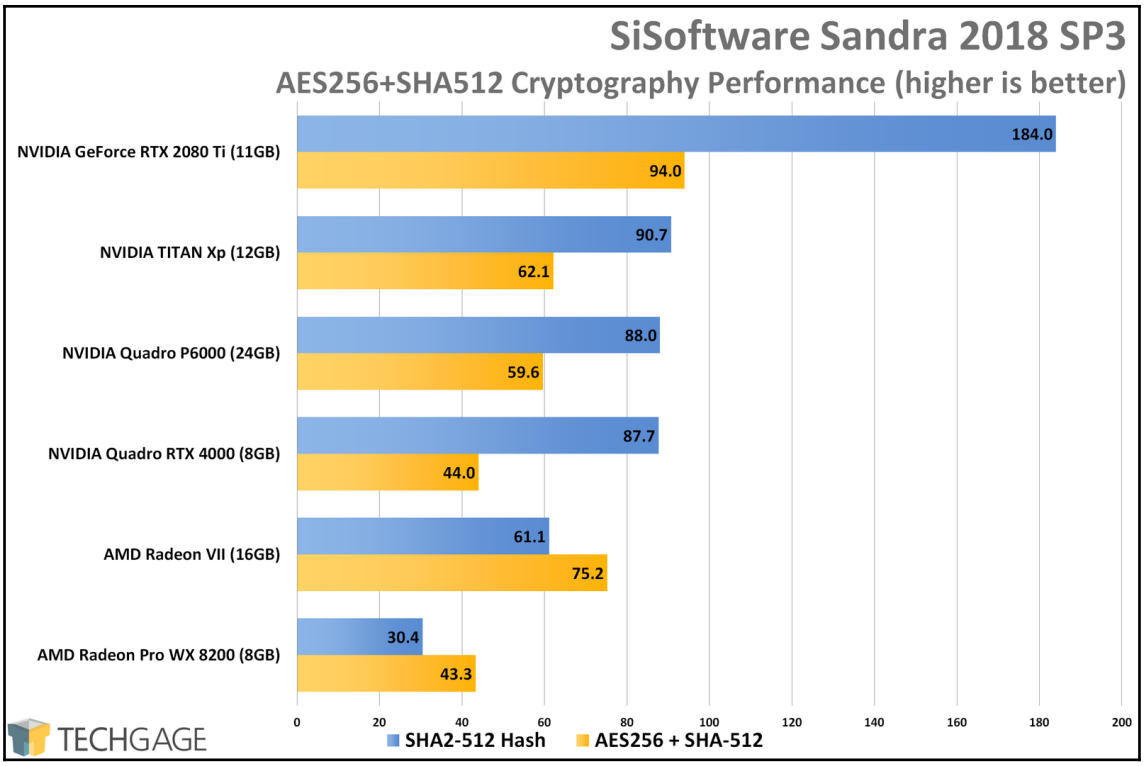
SiSoftware Sandra 2018 SP3
AES256+SHA512 Cryptography Performance (higher is better)
NVIDIA GeForce RTX 2080 Ti (11GB)
184.0
94.0
NVIDIA TITAN Xp (12GB)
90.7
62.1
NVIDIA Quadro P6000 (24GB)
88.0
59.6
NVIDIA Quadro RTX 4000 (8GB)
87.7
44.0
AMD Radeon VII (16GB)
61.1
75.2
AMD Radeon Pro WX 8200 (8GB)
30.4
43.3
0
20
40
60
80
100
120
140
160
180
200
TECHGAGE
SHA2-512 Hash
AES256 + SHA-512

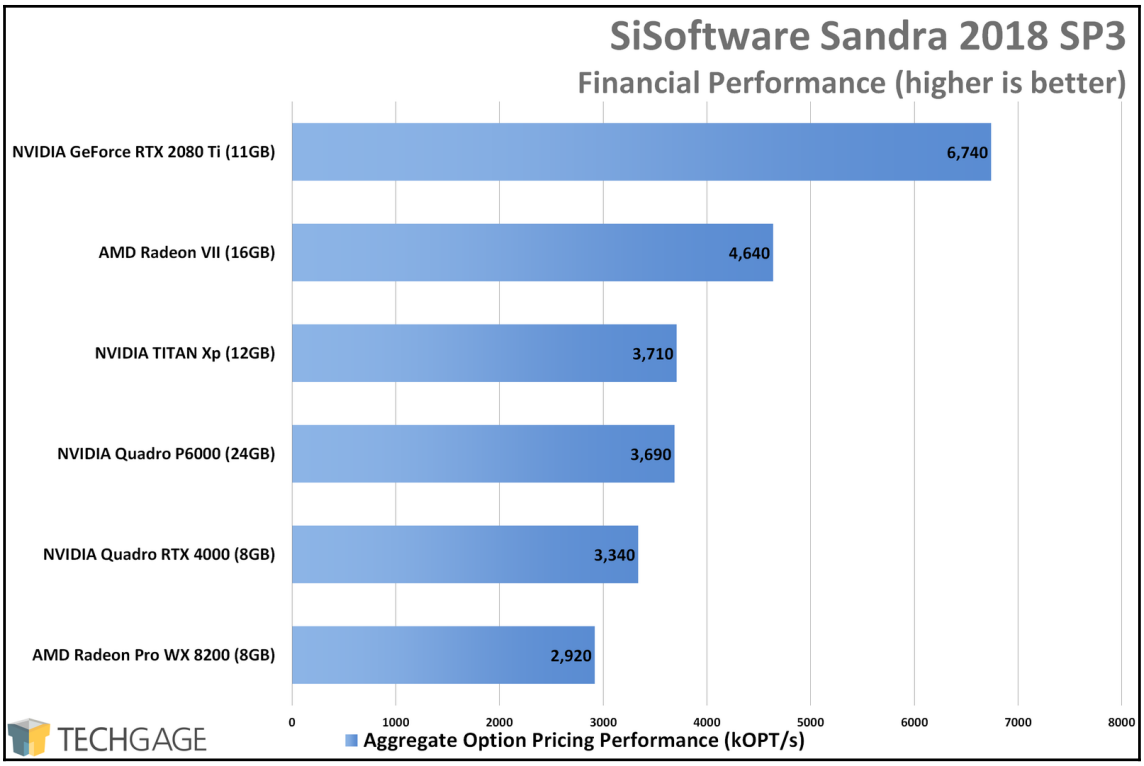
SiSoftware Sandra 2018 SP3
Financial Performance (higher is better)
NVIDIA GeForce RTX 2080 Ti (11GB)
6,740
AMD Radeon VII (16GB)
4,640
NVIDIA TITAN Xp (12GB)
3,710
NVIDIA Quadro P6000 (24GB)
3,690
NVIDIA Quadro RTX 4000 (8GB)
3,340
AMD Radeon Pro WX 8200 (8GB)
2,920
0
1000
2000
3000
4000
5000
6000
7000
8000
TECHGAGE
Aggregate Option Pricing Performance (kOPT/s)

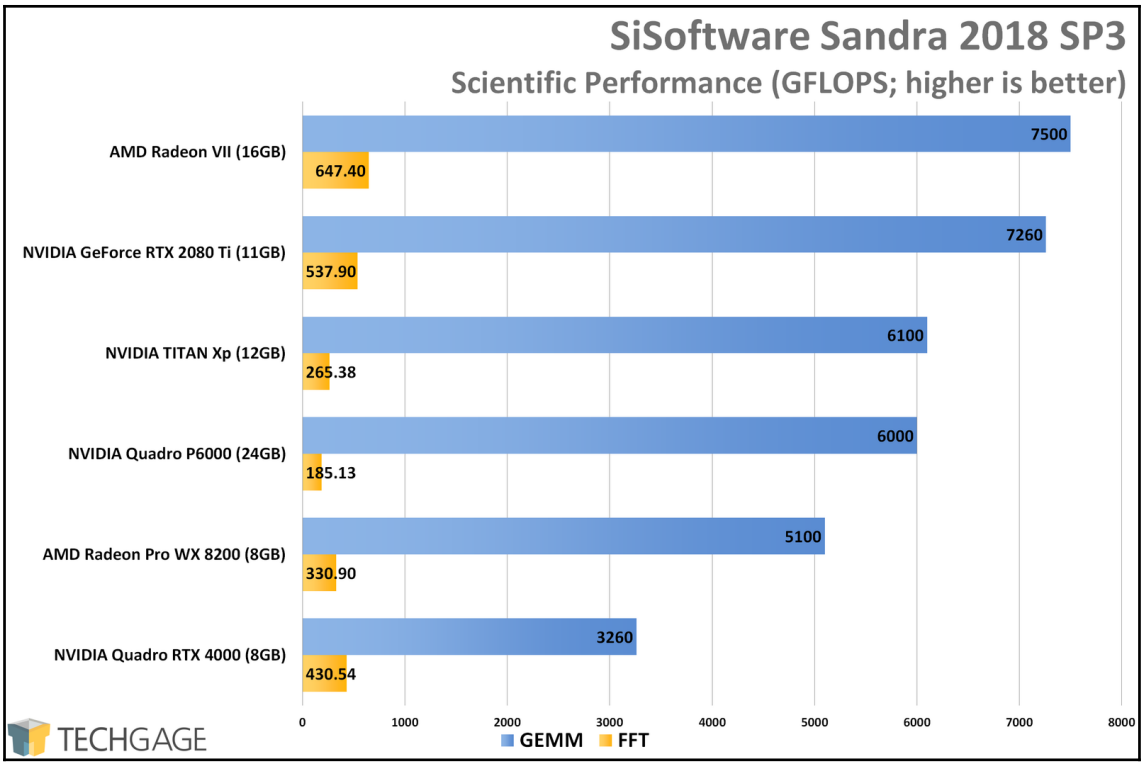
SiSoftware Sandra 2018 SP3
Scientific Performance (GFLOPS; higher is better)
AMD Radeon VII (16GB)
7500
647.40
NVIDIA GeForce RTX 2080 Ti (11GB)
7260
537.90
NVIDIA TITAN Xp (12GB)
6100
265.38
NVIDIA Quadro P6000 (24GB)
6000
185.13
AMD Radeon Pro WX 8200 (8GB)
5100
330.90
NVIDIA Quadro RTX 4000 (8GB)
3260
430.54
0
1000
2000
3000
4000
5000
6000
7000
8000
GEMM
FFT
TECHGAGE

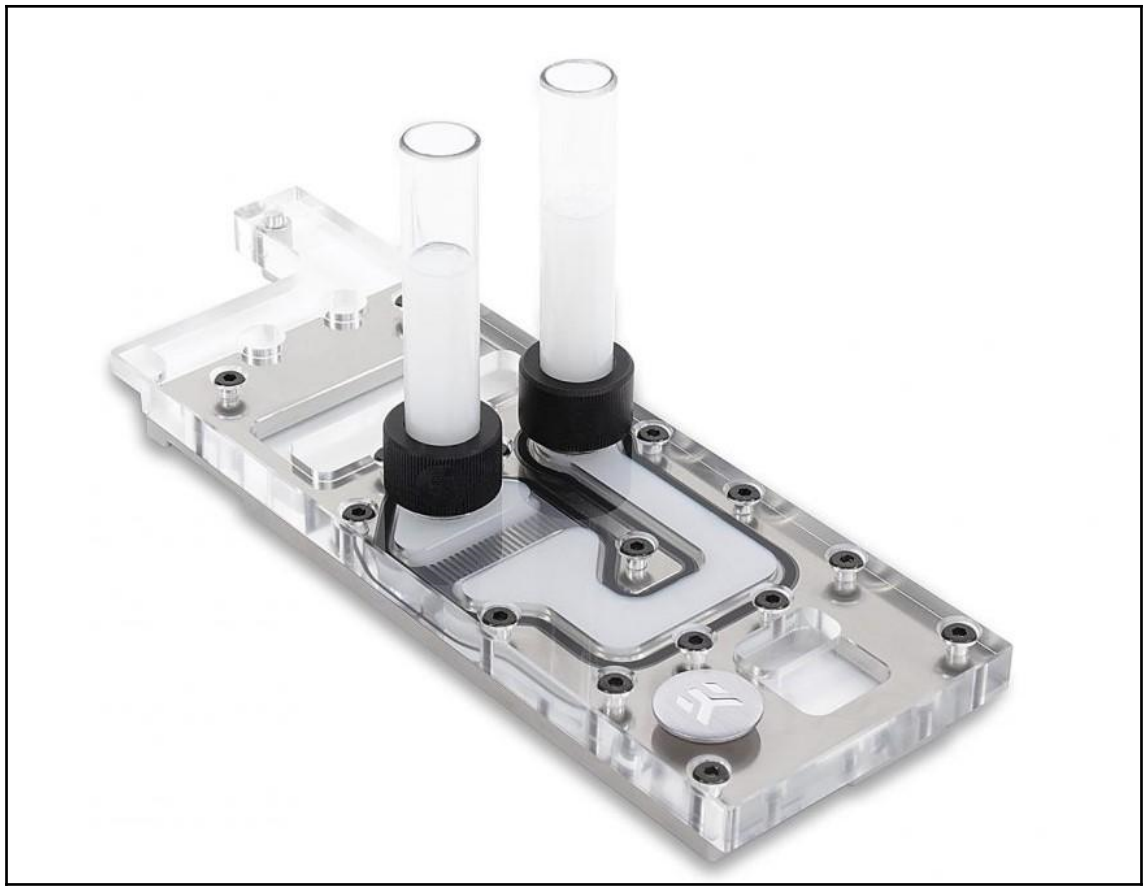

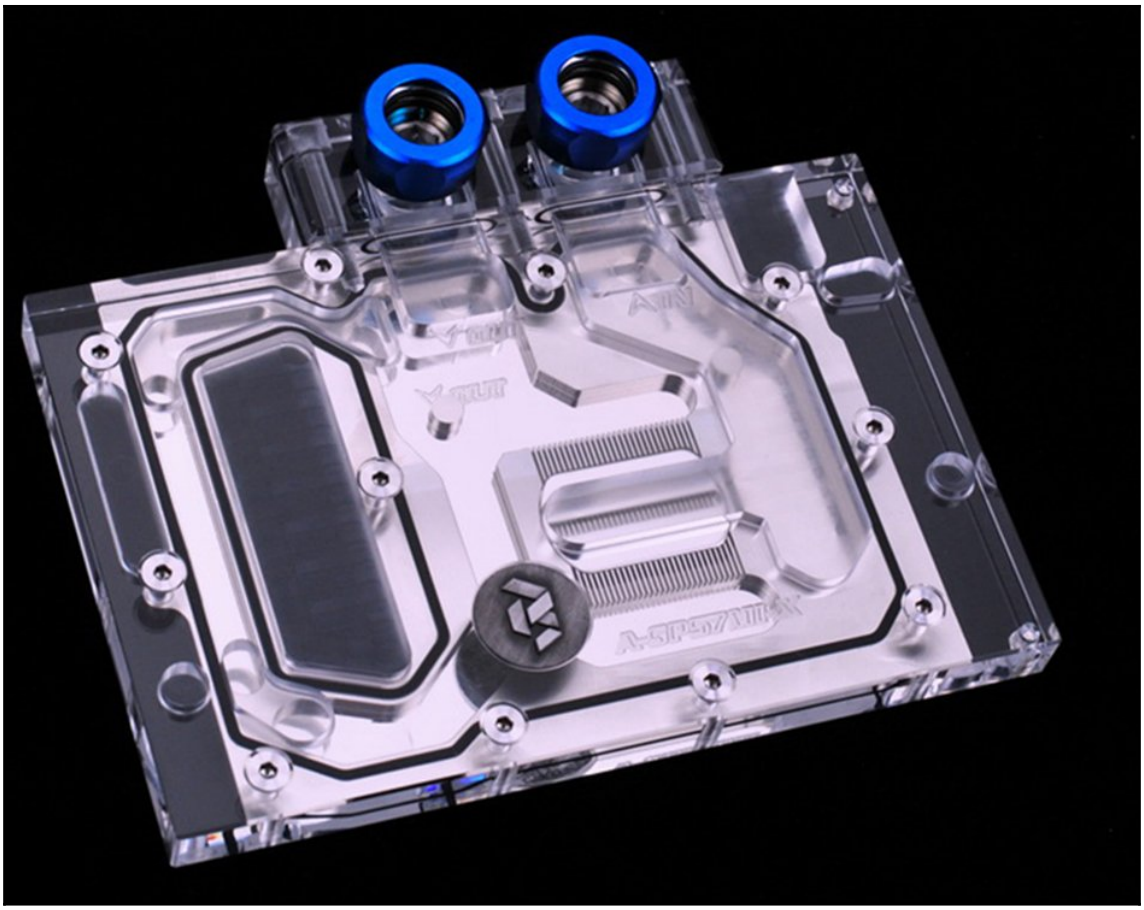


AMD
AMD
RADEON
R
STRIX
CORSAIR
TX850M

Chapter 4: Fundamentals of GPU Programming

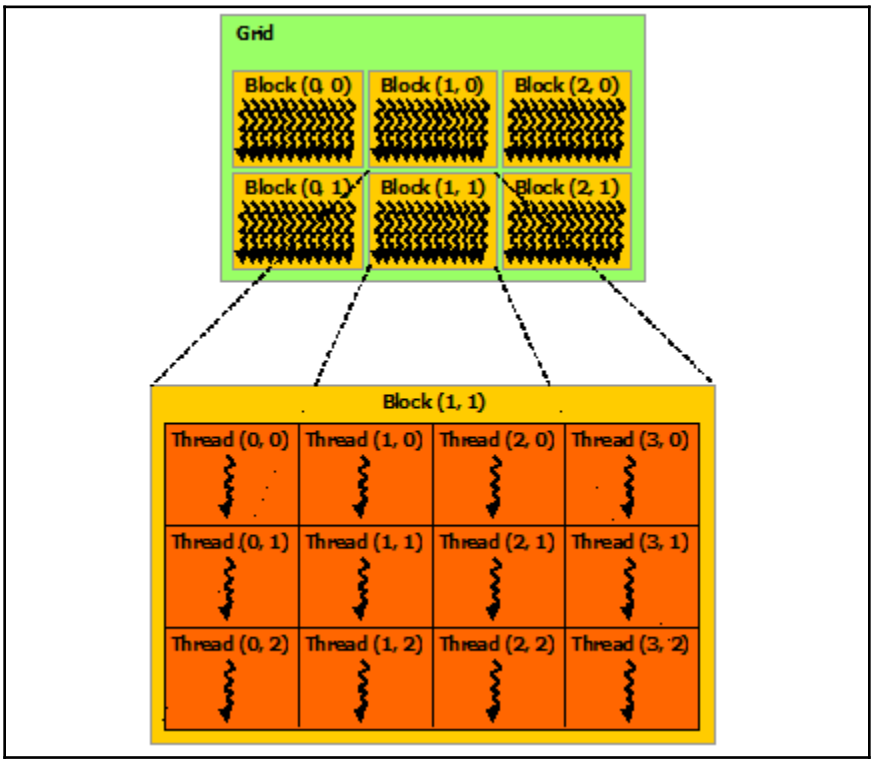

Chapter 5: Setting Up Your Environment for GPU Programming

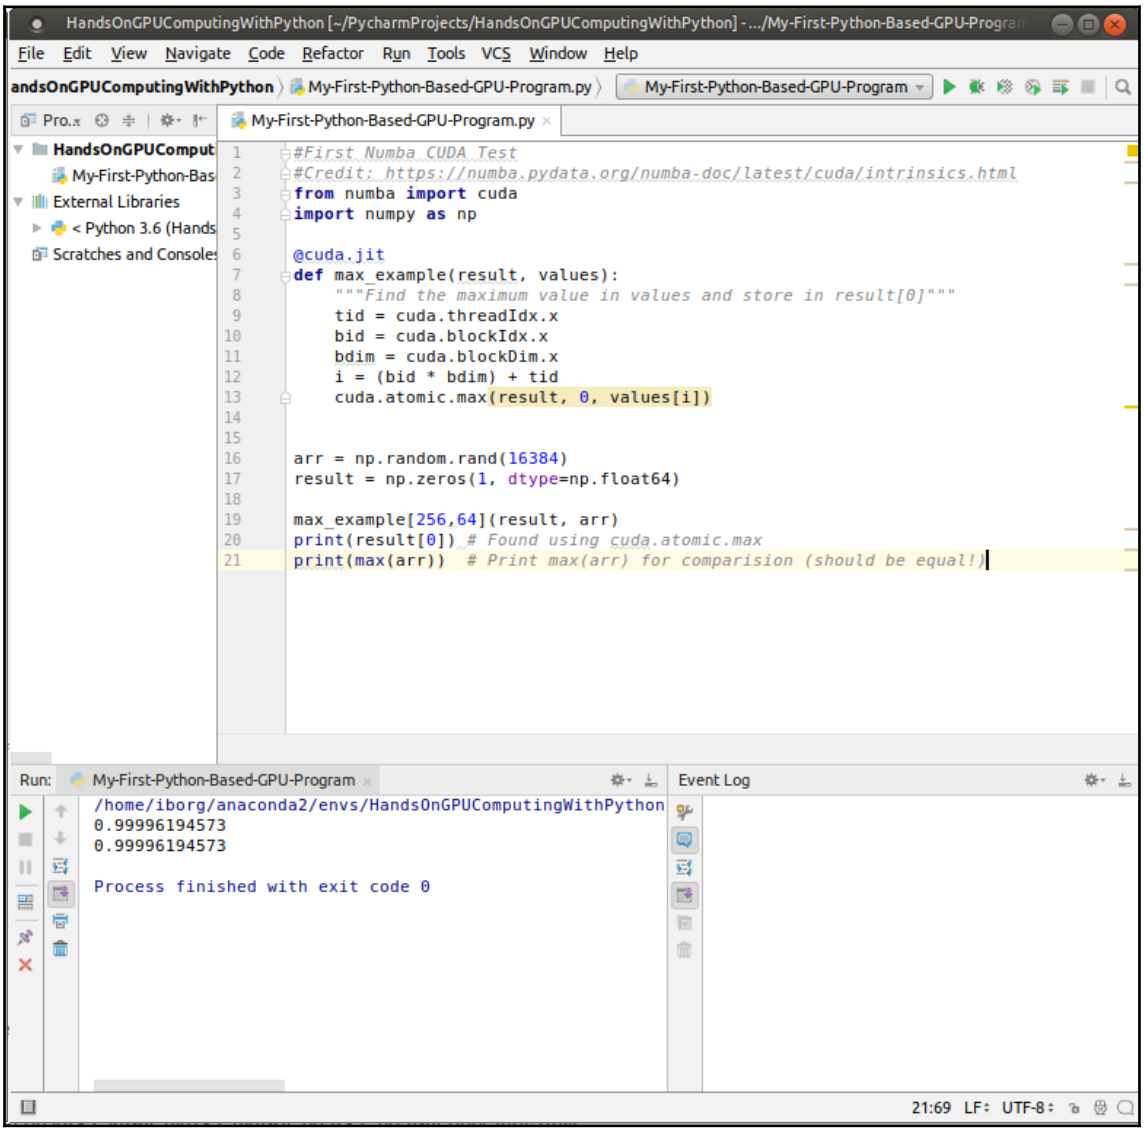

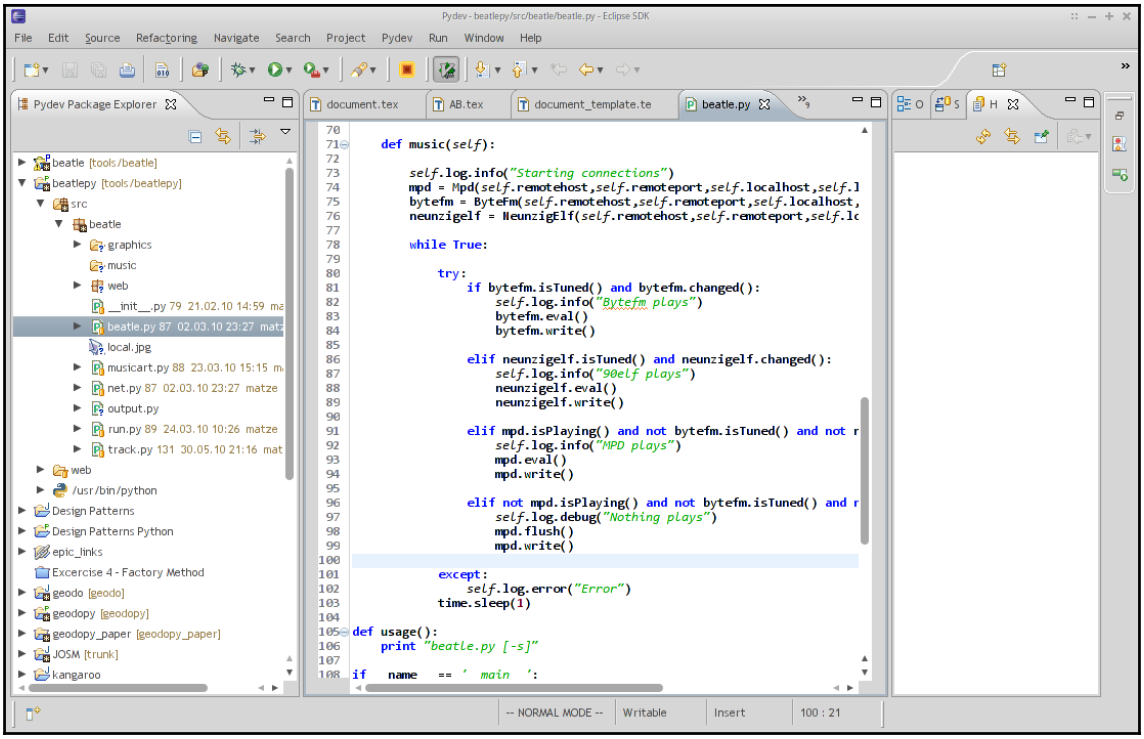
Pydev - beatlepy/src/beatle/beatle.py - Eclipse SDK
File Edit Source Refactoring Navigate Search Project Pydev Run Window Help
Pydev Package Explorer
beatle [tools/beatle]
beatlepy [tools/beatlepy]
src
beatle
graphics
music
web
__init__.py 79 21.02.10 14:59 ma
beatle.py 87 02.03.10 23:27 matz
local.jpg
musicart.py 88 23.03.10 15:15 m
net.py 87 02.03.10 23:27 matze
output.py
run.py 89 24.03.10 10:26 matze
track.py 131 30.05.10 21:16 mat
web
/usr/bin/python
Design Patterns
Design Patterns Python
epic_links
Excercise 4 - Factory Method
geodo [geodo]
geodopy [geodopy]
geodopy_paper [geodopy_paper]
JOSM [trunk]
kangaroo
document.tex
AB.tex
document_template.te
beatle.py
70
71 def music(self):
72
73 self.log.info("Starting connections")
74 mpd = Mpd(self.remotehost,self.remoteport,self.localhost,self.l
75 bytefm = ByteFm(self.remotehost,self.remoteport,self.localhost,
76 neunzigelf = NeunzigElf(self.remotehost,self.remoteport,self.lo
77
78 while True:
79
80 try:
81 if bytefm.isTuned() and bytefm.changed():
82 self.log.info("Bytefm plays")
83 bytefm.eval()
84 bytefm.write()
85
86 elif neunzigelf.isTuned() and neunzigelf.changed():
87 self.log.info("90elf plays")
88 neunzigelf.eval()
89 neunzigelf.write()
90
91 elif mpd.isPlaying() and not bytefm.isTuned() and not r
92 self.log.info("MPD plays")
93 mpd.eval()
94 mpd.write()
95
96 elif not mpd.isPlaying() and not bytefm.isTuned() and r
97 self.log.debug("Nothing plays")
98 mpd.flush()
99 mpd.write()
100
101 except:
102 self.log.error("Error")
103 time.sleep(1)
104
105 def usage():
106 print "beatle.py [-s]"
107
108 if __name__ == '__main__':
-- NORMAL MODE --
Writable
Insert
100 : 21

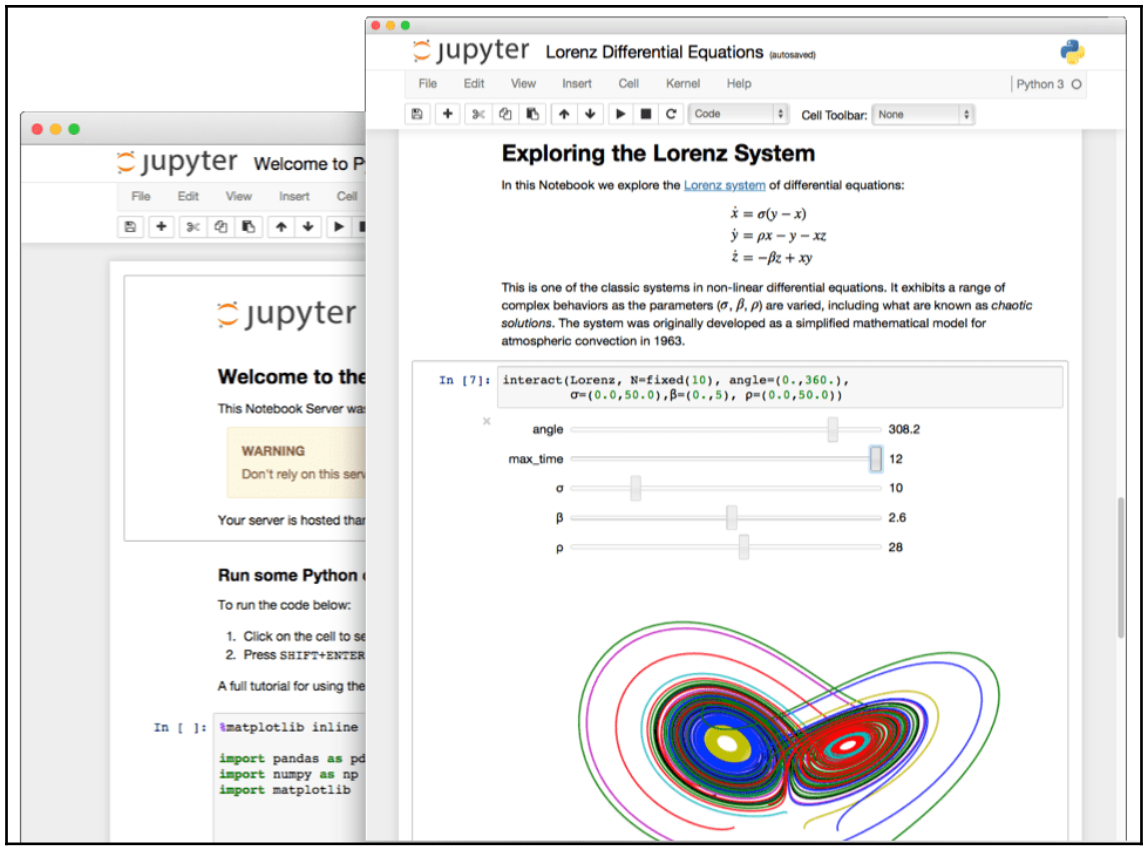
jupyter Welcome to P
File Edit View Insert Cell
jupyter
Welcome to the
This Notebook Server wa
WARNING
Don't rely on this serv
Your server is hosted thar
Run some Python
To run the code below:
1. Click on the cell to se
2. Press SHIFT+ENTER
A full tutorial for using the
In [ ]:
%matplotlib inline
import pandas as pd
import numpy as np
import matplotlib
jupyter Lorenz Differential Equations (autosaved)
File Edit View Insert Cell Kernel Help
Python 3
Code
Cell Toolbar: None
Exploring the Lorenz System
In this Notebook we explore the Lorenz system of differential equations:
$\dot{x} = \sigma(y - x)$
$\dot{y} = \rho x - y - xz$
$\dot{z} = -\beta z + xy$
This is one of the classic systems in non-linear differential equations. It exhibits a range of complex behaviors as the parameters (σ, β, ρ) are varied, including what are known as chaotic solutions. The system was originally developed as a simplified mathematical model for atmospheric convection in 1963.
In [7]:
interact(Lorenz, N=fixed(10), angle=(0.,360.),
σ=(0.0,50.0),β=(0.,5), ρ=(0.0,50.0))
angle 308.2
max_time 12
σ 10
β 2.6
ρ 28

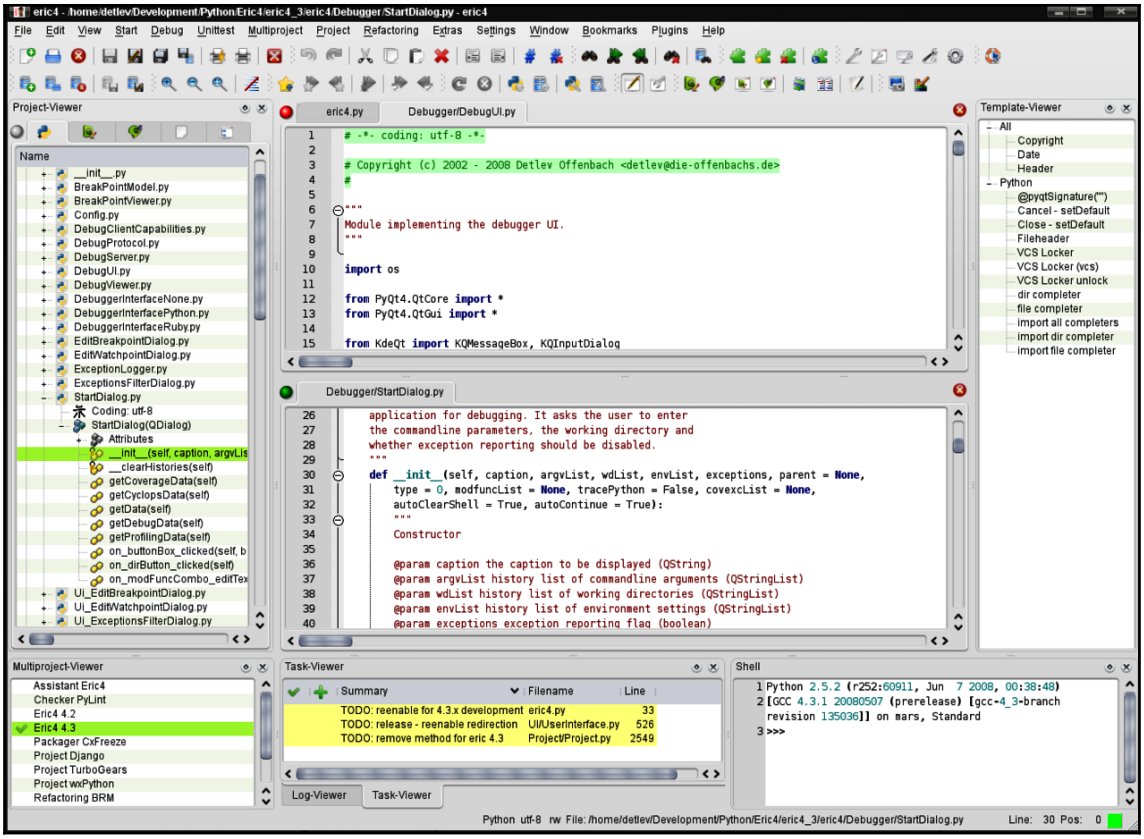
eric4 - /home/detlev/Development/Python/Eric4/eric4_3/eric4/Debugger/StartDialog.py - eric4
File Edit View Start Debug Unittest Multiproject Project Refactoring Extras Settings Window Bookmarks Plugins Help
Project-Viewer
Name
__init__.py
BreakPointModel.py
BreakPointViewer.py
Config.py
DebugClientCapabilities.py
DebugProtocol.py
DebugServer.py
DebugUI.py
DebugViewer.py
DebuggerInterfaceNone.py
DebuggerInterfacePython.py
DebuggerInterfaceRuby.py
EditBreakpointDialog.py
EditWatchpointDialog.py
ExceptionLogger.py
ExceptionsFilterDialog.py
StartDialog.py
Coding: utf-8
StartDialog(QDialog)
Attributes
__init__(self, caption, argvLis
__clearHistories(self)
getCoverageData(self)
getCyclopsData(self)
getData(self)
getDebugData(self)
getProfilingData(self)
on_buttonBox_clicked(self, b
on_dirButton_clicked(self)
on_modFuncCombo_editTex
Ui_EditBreakpointDialog.py
Ui_EditWatchpointDialog.py
Ui_ExceptionsFilterDialog.py
eric4.py
Debugger/DebugUI.py
-*- coding: utf-8 -*-
Copyright (c) 2002 - 2008 Detlev Offenbach <detlev@die-offenbachs.de>
#
"""
Module implementing the debugger UI.
"""
import os
from PyQt4.QtCore import *
from PyQt4.QtGui import *
from KdeQt import KQMessageBox, KQInputDialog
Debugger/StartDialog.py
application for debugging. It asks the user to enter
the commandline parameters, the working directory and
whether exception reporting should be disabled.
"""
def __init__(self, caption, argvList, wdList, envList, exceptions, parent = None,
type = 0, modfuncList = None, tracePython = False, covexcList = None,
autoClearShell = True, autoContinue = True):
"""
Constructor
@param caption the caption to be displayed (QString)
@param argvList history list of commandline arguments (QStringList)
@param wdList history list of working directories (QStringList)
@param envList history list of environment settings (QStringList)
@param exceptions exception reporting flag (boolean)
Template-Viewer
All
Copyright
Date
Header
Python
@pyqtSignature("")
Cancel - setDefault
Close - setDefault
Fileheader
VCS Locker
VCS Locker (vcs)
VCS Locker unlock
dir completer
file completer
import all completers
import dir completer
import file completer
Multiproject-Viewer
Assistant Eric4
Checker PyLint
Eric4 4.2
Eric4 4.3
Packager CxFreeze
Project Django
Project TurboGears
Project wxPython
Refactoring BRM
Task-Viewer
Summary
Filename
Line
TODO: reenable for 4.3.x development eric4.py 33
TODO: release - reenable redirection UI/UserInterface.py 526
TODO: remove method for eric 4.3 Project/Project.py 2549
Log-Viewer
Task-Viewer
Shell
Python 2.5.2 (r252:60911, Jun 7 2008, 00:38:48)
[GCC 4.3.1 20080507 (prerelease) [gcc-4_3-branch
revision 135036]] on mars, Standard
>>>
Python utf-8 rw File: /home/detlev/Development/Python/Eric4/eric4_3/eric4/Debugger/StartDialog.py
Line: 30 Pos: 0

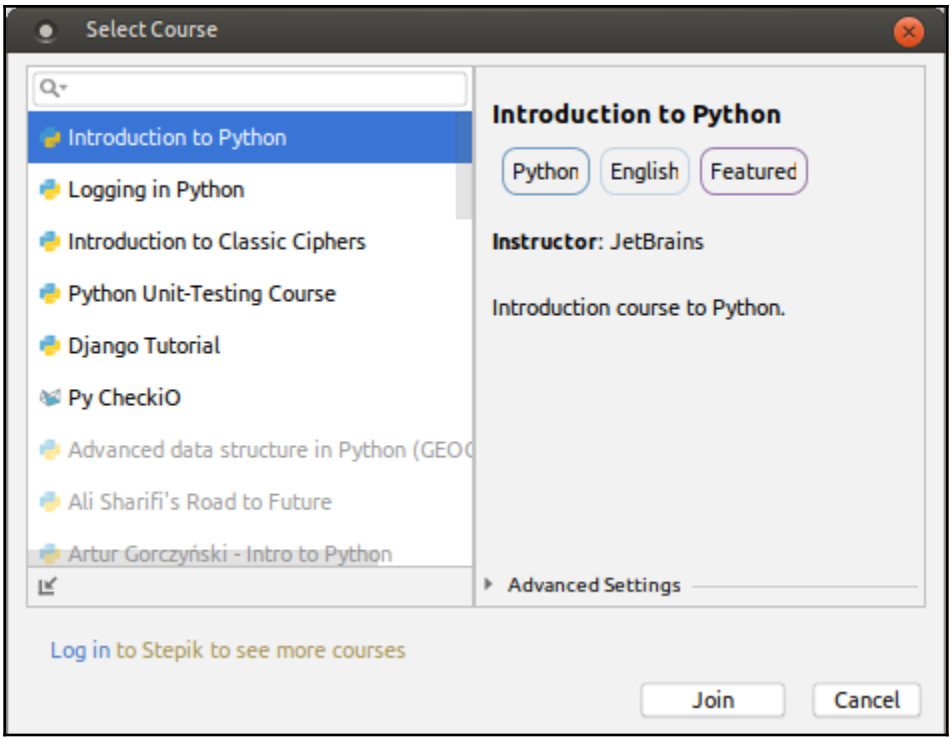
Select Course
Introduction to Python
Logging in Python
Introduction to Classic Ciphers
Python Unit-Testing Course
Django Tutorial
Py CheckiO
Advanced data structure in Python (GEO
Ali Sharifi's Road to Future
Artur Gorczyński - Intro to Python
Introduction to Python
Python
English
Featured
Instructor: JetBrains
Introduction course to Python.
Advanced Settings
Log in to Stepik to see more courses
Join
Cancel

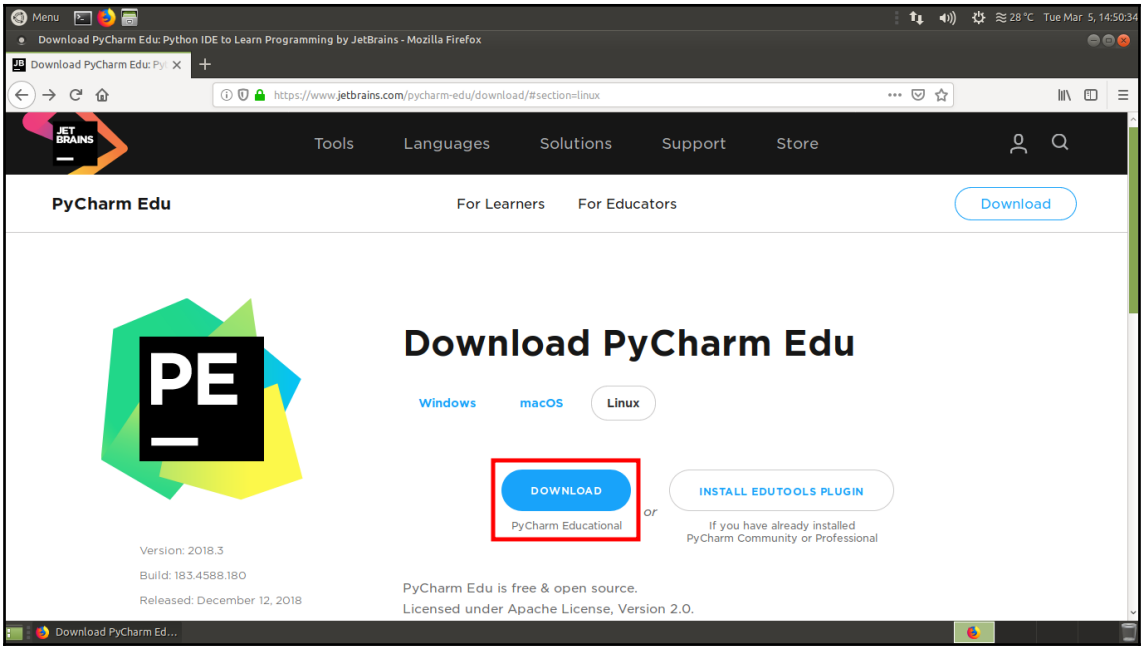
Download PyCharm Edu: Python IDE to Learn Programming by JetBrains - Mozilla Firefox
https://www.jetbrains.com/pycharm-edu/download/#section=linux
Tools
Languages
Solutions
Support
Store
PyCharm Edu
For Learners
For Educators
Download
Download PyCharm Edu
Windows
macOS
Linux
DOWNLOAD
PyCharm Educational
or
INSTALL EDUTOOLS PLUGIN
If you have already installed
PyCharm Community or Professional
Version: 2018.3
Build: 183.4588.180
Released: December 12, 2018
PyCharm Edu is free & open source.
Licensed under Apache License, Version 2.0.

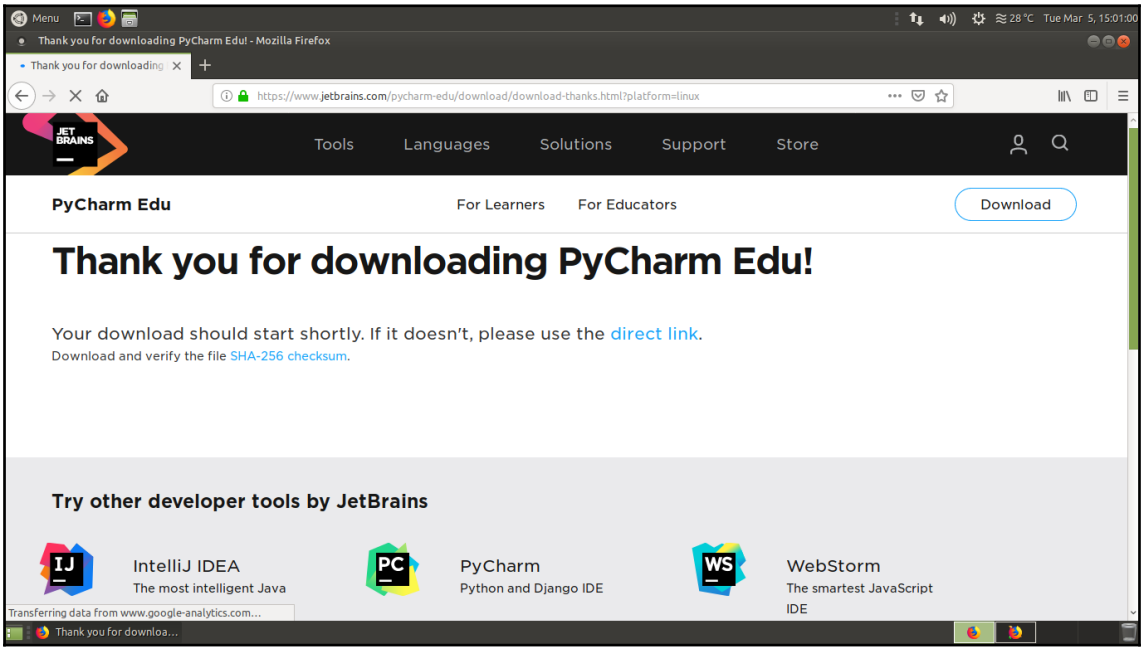
Thank you for downloading PyCharm Edu! - Mozilla Firefox
https://www.jetbrains.com/pycharm-edu/download/download-thanks.html?platform=linux
Tools
Languages
Solutions
Support
Store
PyCharm Edu
For Learners
For Educators
Download
Thank you for downloading PyCharm Edu!
Your download should start shortly. If it doesn't, please use the direct link.
Download and verify the file SHA-256 checksum.
Try other developer tools by JetBrains
IntelliJ IDEA
The most intelligent Java
PyCharm
Python and Django IDE
WebStorm
The smartest JavaScript
IDE

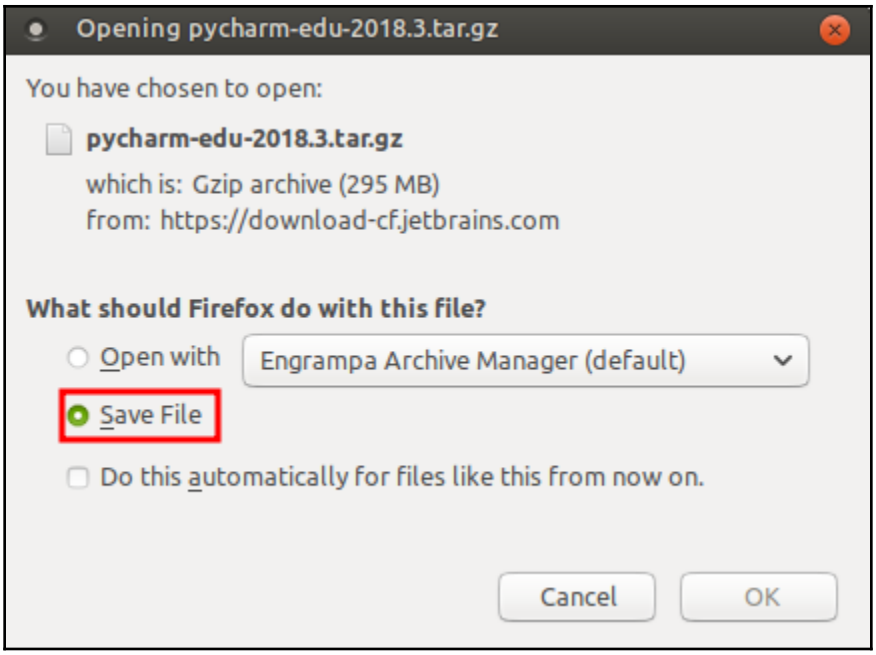
Opening pycharm-edu-2018.3.tar.gz
You have chosen to open:
pycharm-edu-2018.3.tar.gz
which is: Gzip archive (295 MB)
from: https://download-cf.jetbrains.com
What should Firefox do with this file?
Open with
Engrampa Archive Manager (default)
Save File
Do this automatically for files like this from now on.
Cancel
OK

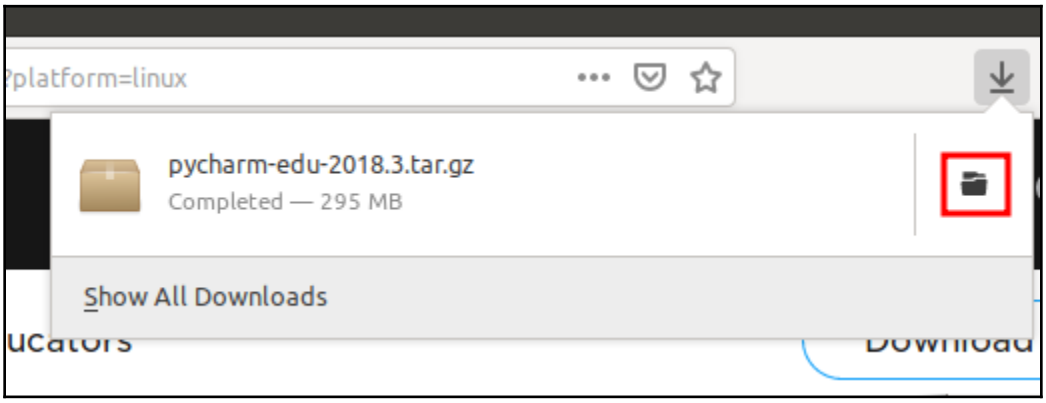
?platform=linux
pycharm-edu-2018.3.tar.gz
Completed — 295 MB
Show All Downloads

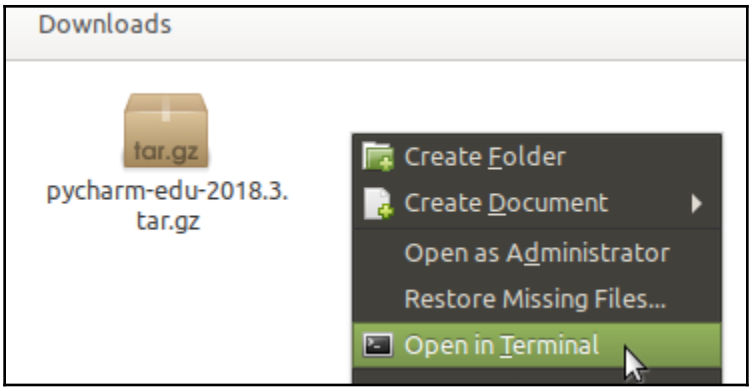
Downloads
tar.gz
pycharm-edu-2018.3.
tar.gz
Create Folder
Create Document
Open as Administrator
Restore Missing Files...
Open in Terminal

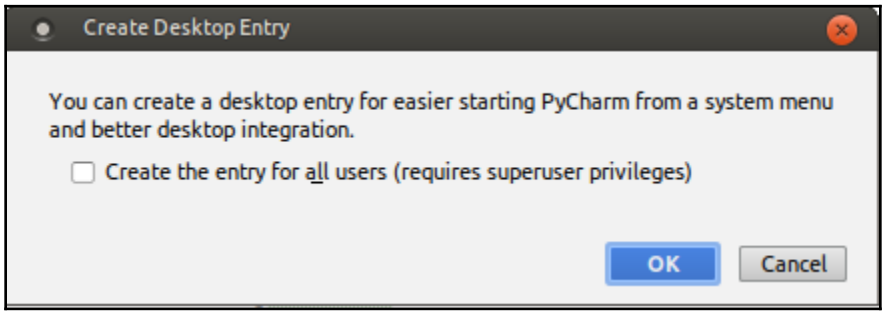
Create Desktop Entry
You can create a desktop entry for easier starting PyCharm from a system menu and better desktop integration.
Create the entry for all users (requires superuser privileges)
OK
Cancel

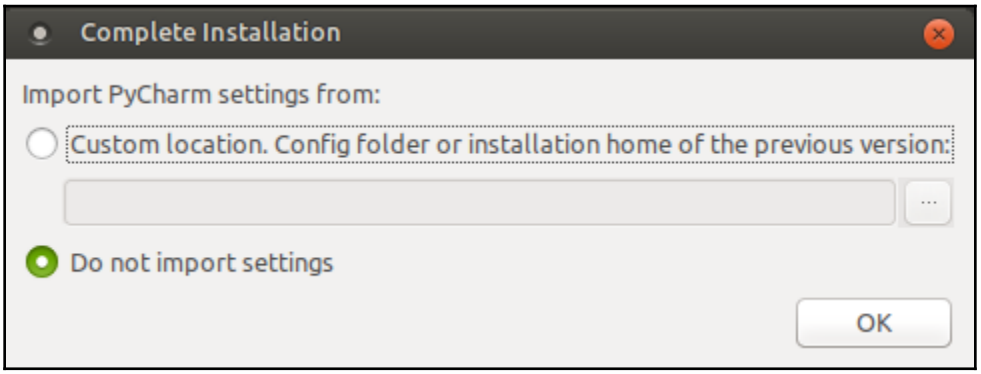
Complete Installation
Import PyCharm settings from:
Custom location. Config folder or installation home of the previous version:
Do not import settings
OK

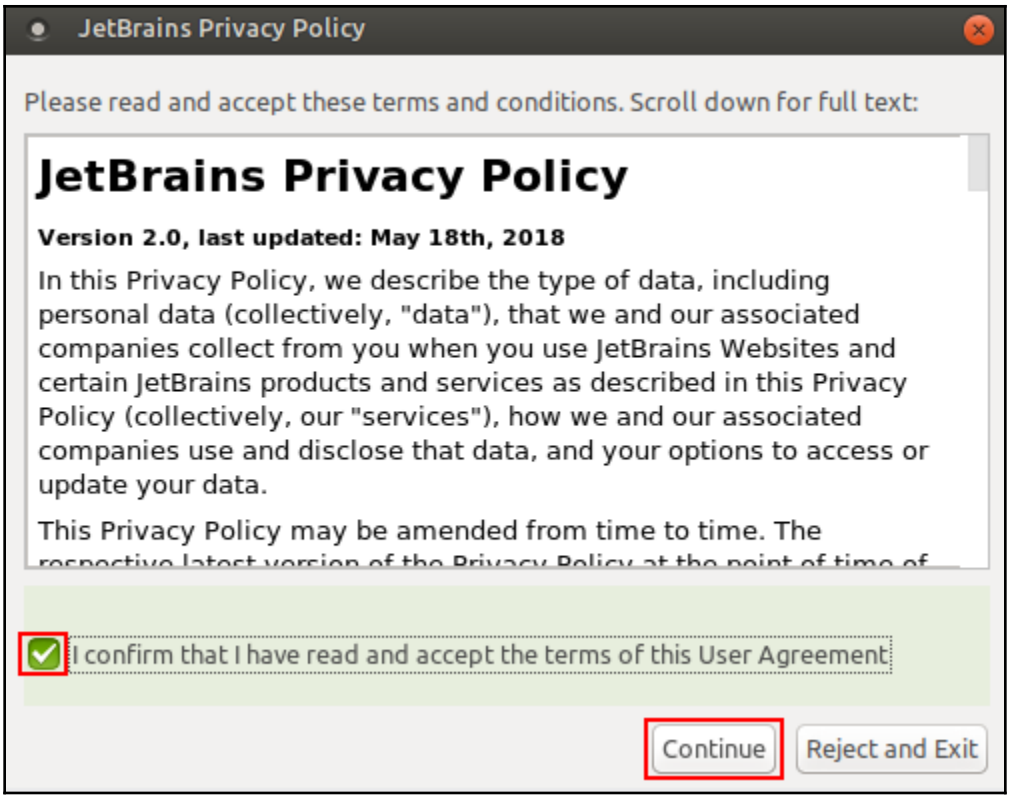
JetBrains Privacy Policy
Please read and accept these terms and conditions. Scroll down for full text:
JetBrains Privacy Policy
Version 2.0, last updated: May 18th, 2018
In this Privacy Policy, we describe the type of data, including personal data (collectively, "data"), that we and our associated companies collect from you when you use JetBrains Websites and certain JetBrains products and services as described in this Privacy Policy (collectively, our "services"), how we and our associated companies use and disclose that data, and your options to access or update your data.
This Privacy Policy may be amended from time to time. The
I confirm that I have read and accept the terms of this User Agreement
Continue
Reject and Exit

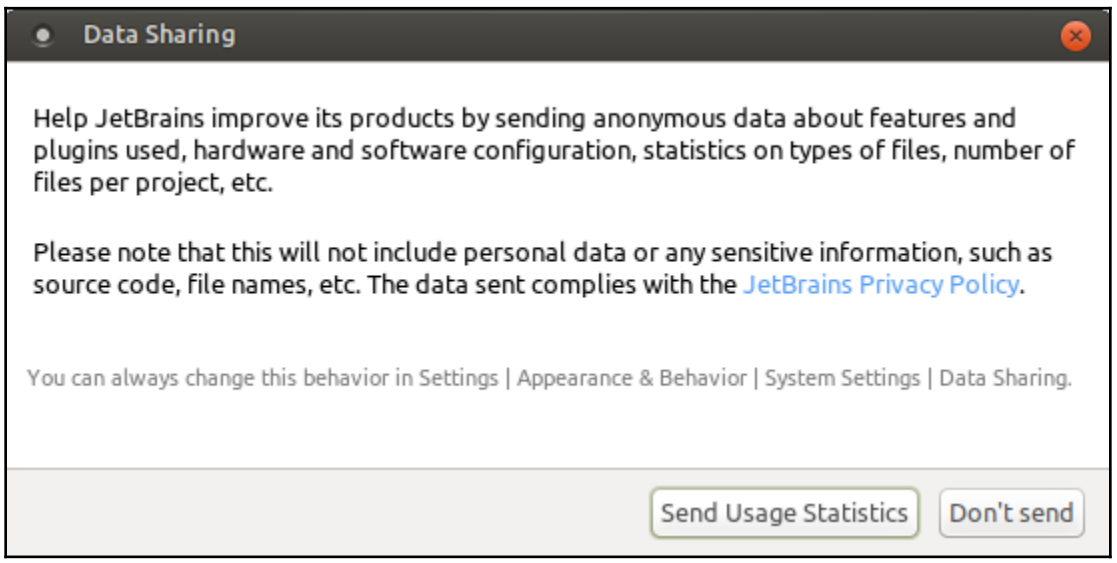
Data Sharing
Help JetBrains improve its products by sending anonymous data about features and plugins used, hardware and software configuration, statistics on types of files, number of files per project, etc.
Please note that this will not include personal data or any sensitive information, such as source code, file names, etc. The data sent complies with the JetBrains Privacy Policy.
You can always change this behavior in Settings | Appearance & Behavior | System Settings | Data Sharing.
Send Usage Statistics
Don't send

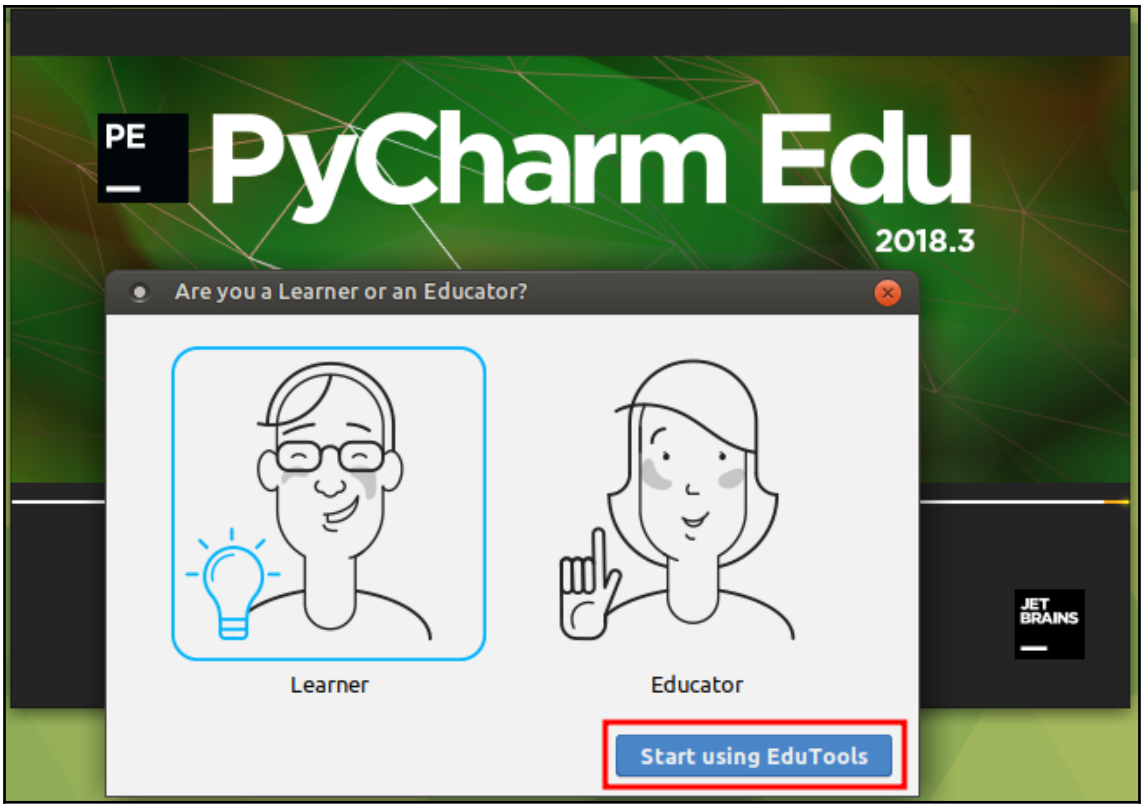
PE
PyCharm Edu
2018.3
Are you a Learner or an Educator?
Learner
Educator
Start using EduTools
JET BRAINS

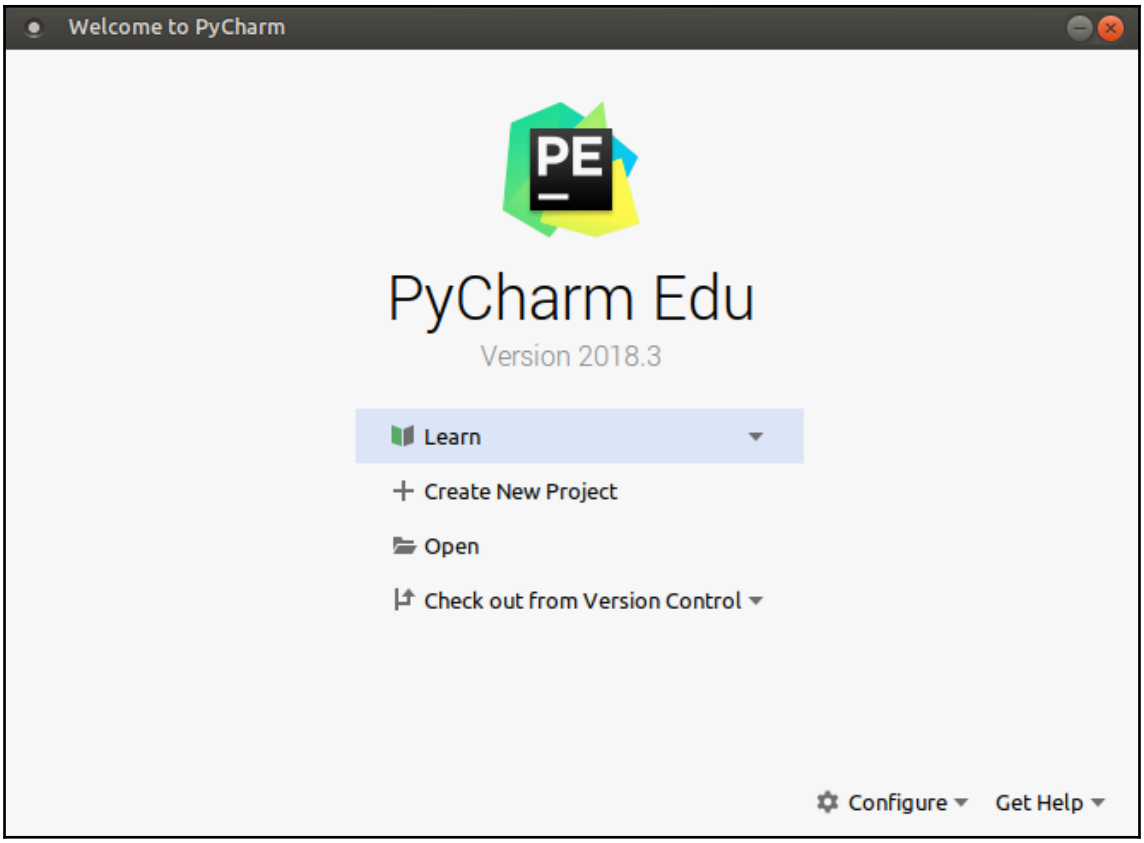
Welcome to PyCharm
PyCharm Edu
Version 2018.3
Learn
Create New Project
Open
Check out from Version Control
Configure
Get Help

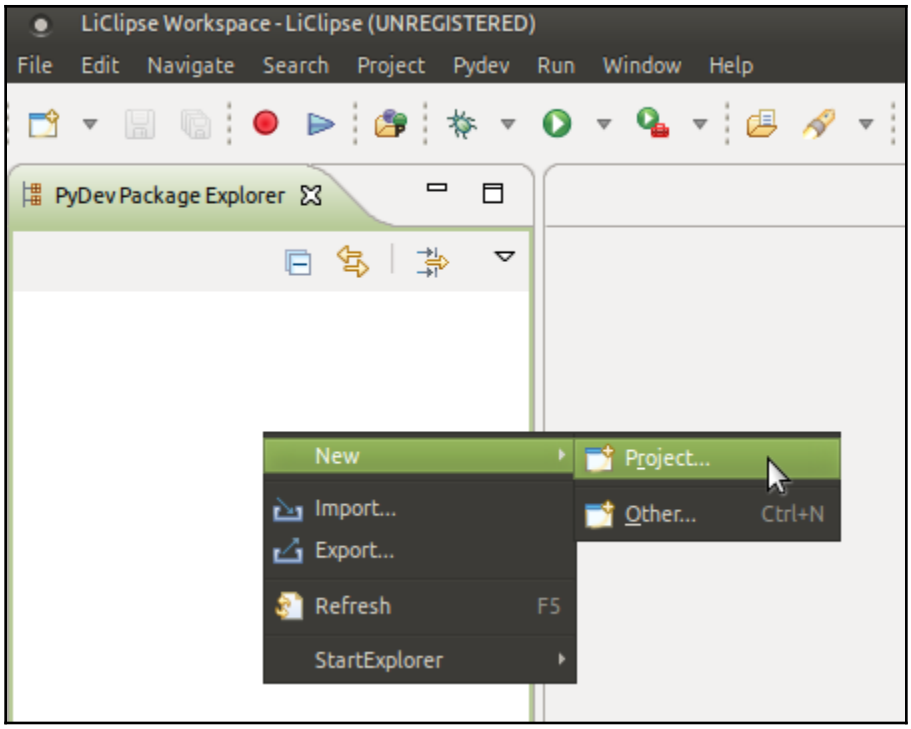
LiClipse Workspace - LiClipse (UNREGISTERED)
File
Edit
Navigate
Search
Project
Pydev
Run
Window
Help
PyDev Package Explorer
New
Import...
Export...
Refresh
F5
StartExplorer
Project...
Other...
Ctrl+N

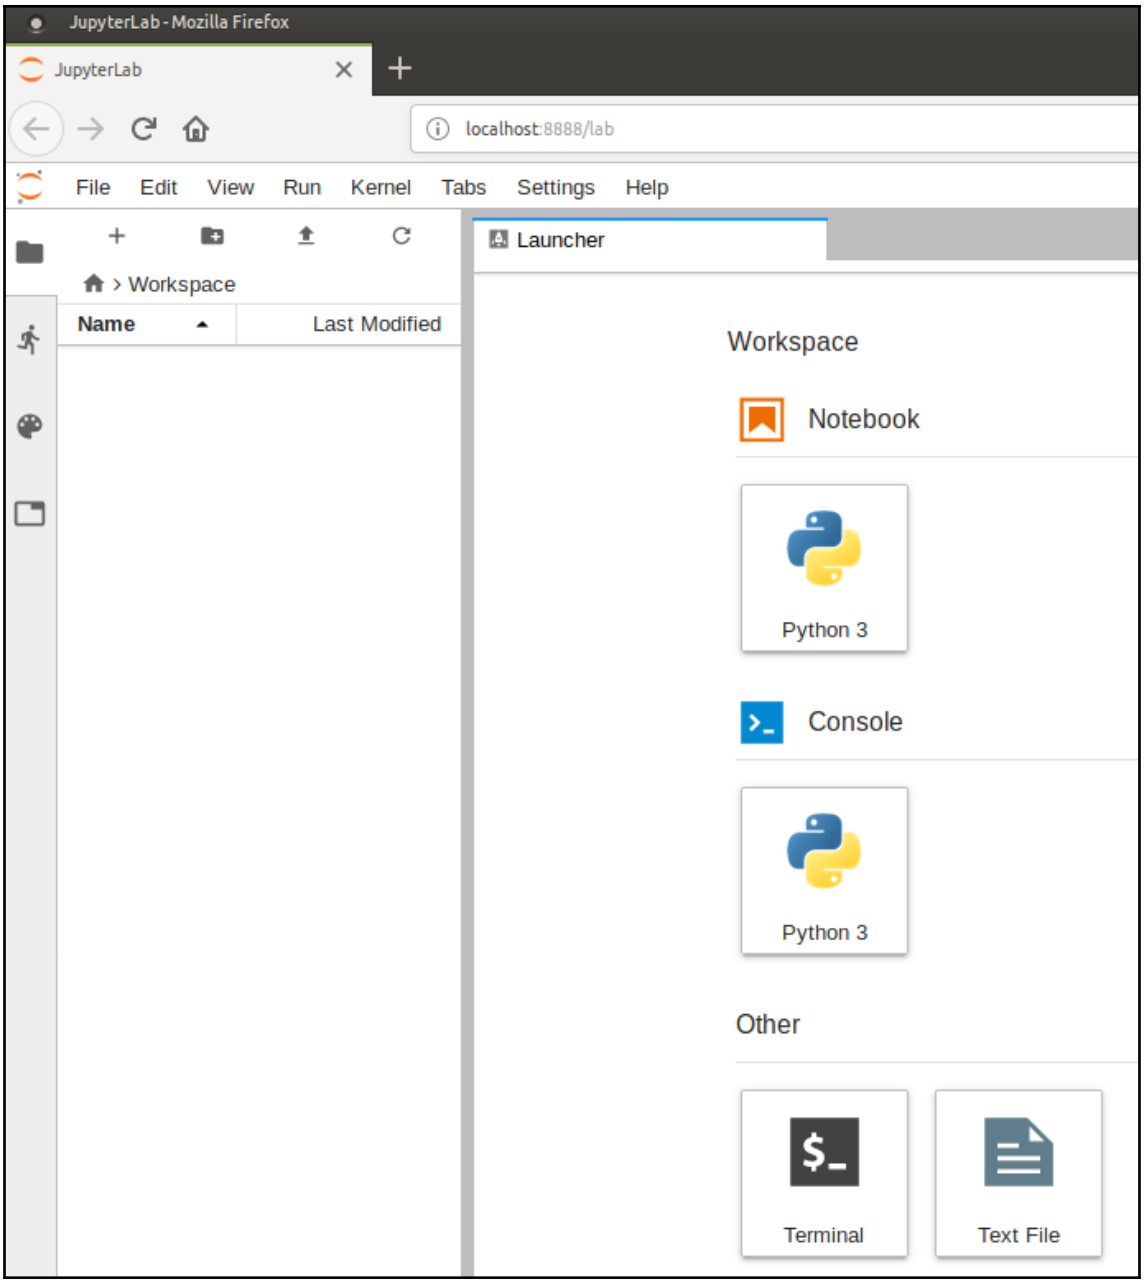
JupyterLab - Mozilla Firefox
JupyterLab
localhost:8888/lab
File
Edit
View
Run
Kernel
Tabs
Settings
Help
Launcher
Workspace
Name
Last Modified
Workspace
Notebook
Python 3
Console
Python 3
Other
Terminal
Text File

Installing LiClipse

The recommended way of using LiClipse is by downloading the native installer for your platform, which provides a pre-configured standalone for Windows (32 bits or 64 bits), Mac OS (64 bits) or Linux (32 bits or 64 bits).

However, it's also possible to install it through the update site:
http://update.liclipse.com/latest
(note that you must have Java 8 installed in this case -- see details on the Update Site Install section).

The latest versions may be downloaded through the links below:

- LiClipse 5.1.3 Windows (64 bits)
- LiClipse 5.1.3 Windows (32 bits)
- LiClipse 5.1.3 Mac OS (64 bits)
- LiClipse 5.1.3 Linux (64 bits)
- LiClipse 5.1.3 Linux (32 bits)
- LiClipse 5.1.3 Local Update Site

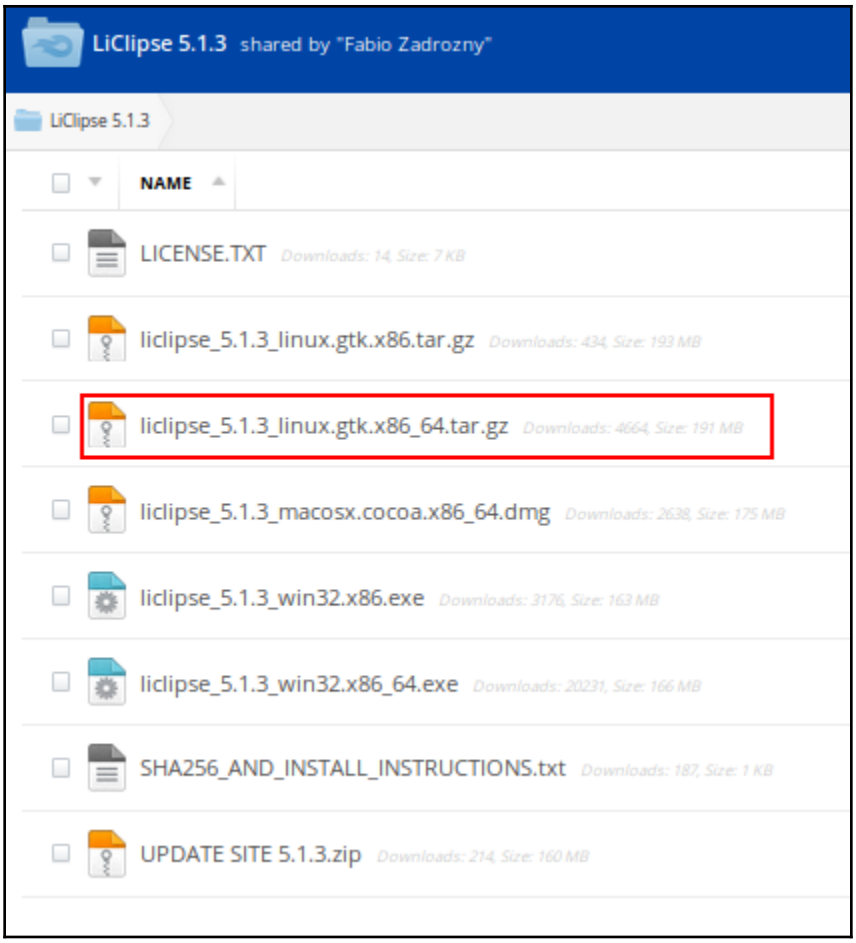
LiClipse 5.1.3 shared by "Fabio Zadrozny"
LiClipse 5.1.3
NAME
LICENSE.TXT
liclipse_5.1.3_linux.gtk.x86.tar.gz
liclipse_5.1.3_linux.gtk.x86_64.tar.gz
liclipse_5.1.3_macosx.cocoa.x86_64.dmg
liclipse_5.1.3_win32.x86.exe
liclipse_5.1.3_win32.x86_64.exe
SHA256_AND_INSTALL_INSTRUCTIONS.txt
UPDATE SITE 5.1.3.zip

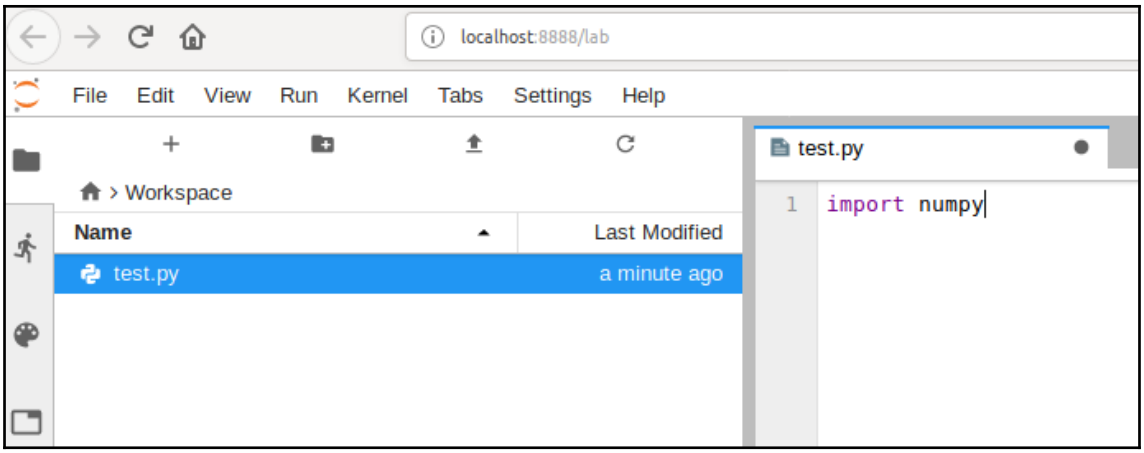
localhost:8888/lab
File
Edit
View
Run
Kernel
Tabs
Settings
Help
Workspace
Name
Last Modified
test.py
a minute ago
test.py
1
import numpy

Chapter 6: Working with CUDA and PyCUDA

```
$ nvprof ./gpu_multiply
==23253== NVPROF is profiling process 23253, command: ./gpu_multiply

GPU Device used for computation: GeForce GTX TITAN X

Multiplication on GPU computed in: 45.111168 milliseconds
Max error: 0
==23253== Profiling application: ./gpu_multiply
==23253== Profiling result:
            Type  Time(%)      Time     Calls       Avg       Min       Max  Name
 GPU activities:   64.48%  863.43ms         2  431.71ms  429.96ms  433.47ms  [CUDA memcpy HtoD]
                   32.16%  430.61ms         1  430.61ms  430.61ms  430.61ms  [CUDA memcpy DtoH]
                    3.37%  45.061ms         1  45.061ms  45.061ms  45.061ms  multiply(double*, double*, unsigned long)
      API calls:   90.15%  1.29438s         3  431.46ms  430.04ms  433.56ms  cudaMemcpy
                    5.78%  83.029ms         2  41.515ms  3.3217ms  79.708ms  cudaMalloc
                    3.14%  45.126ms         1  45.126ms  45.126ms  45.126ms  cudaEventSynchronize
                    0.87%  12.533ms         2  6.2664ms  2.7815ms  9.7513ms  cudaFree
                    0.02%  263.33us        94  2.8010us     300ns  106.54us  cuDeviceGetAttribute
                    0.02%  261.07us         1  261.07us  261.07us  261.07us  cudaGetDeviceProperties
                    0.01%  145.22us         1  145.22us  145.22us  145.22us  cuDeviceTotalMem
                    0.00%  28.937us         1  28.937us  28.937us  28.937us  cudaLaunch
                    0.00%  27.919us         1  27.919us  27.919us  27.919us  cuDeviceGetName
                    0.00%  9.4900us         2  4.7450us  1.3550us  8.1350us  cudaEventCreate
                    0.00%  7.2200us         2  3.6100us  3.0590us  4.1610us  cudaEventRecord
                    0.00%  5.6390us         1  5.6390us  5.6390us  5.6390us  cudaEventElapsedTime
                    0.00%  1.8280us         3     609ns     329ns  1.1010us  cuDeviceGetCount
                    0.00%  1.0300us         3     343ns     122ns     490ns  cudaSetupArgument
                    0.00%     992ns         2     496ns     324ns     668ns  cuDeviceGet
                    0.00%     871ns         1     871ns     871ns     871ns  cudaConfigureCall
                    0.00%     207ns         1     207ns     207ns     207ns  cudaGetDeviceCount
```

```
$ nvcc gpu_multiply_revised.cu -o gpu_multiply_revised
$ nvprof ./gpu_multiply_revised
==24085== NVPROF is profiling process 24085, command: ./gpu_multiply_revised

GPU Device used for computation: GeForce GTX TITAN X

Multiplication on GPU computed in: 707.247864 milliseconds
Max error: 0
==24085== Profiling application: ./gpu_multiply_revised
==24085== Profiling result:
            Type  Time(%)      Time     Calls       Avg       Min       Max  Name
 GPU activities:  100.00%  44.840ms         1  44.840ms  44.840ms  44.840ms  multiply(double*, double*, unsigned long)
      API calls:   35.14%  714.82ms         2  357.41ms  286.24ms  428.58ms  cudaMallocManaged
                   32.56%  662.39ms         1  662.39ms  662.39ms  662.39ms  cudaLaunch
                   30.04%  611.13ms         2  305.57ms  282.21ms  328.92ms  cudaFree
                    2.21%  44.925ms         2  22.462ms  47.642us  44.877ms  cudaDeviceSynchronize
                    0.03%  677.37us        94  7.2060us     374ns  304.22us  cuDeviceGetAttribute
                    0.01%  258.02us         1  258.02us  258.02us  258.02us  cudaGetDeviceProperties
                    0.01%  194.39us         1  194.39us  194.39us  194.39us  cuDeviceTotalMem
                    0.00%  45.546us         1  45.546us  45.546us  45.546us  cuDeviceGetName
                    0.00%  17.407us         2  8.7030us  8.3380us  9.0690us  cudaEventRecord
                    0.00%  16.356us         2  8.1780us  1.2330us  15.123us  cudaEventCreate
                    0.00%  4.2530us         2  2.1260us  1.1520us  3.1010us  cudaMemPrefetchAsync
                    0.00%  2.9140us         1  2.9140us  2.9140us  2.9140us  cudaEventElapsedTime
                    0.00%  2.8480us         1  2.8480us  2.8480us  2.8480us  cudaEventSynchronize
                    0.00%  2.5410us         3     847ns     411ns  1.4760us  cuDeviceGetCount
                    0.00%  1.4210us         2     710ns     489ns     932ns  cuDeviceGet
                    0.00%     763ns         3     254ns     123ns     327ns  cudaSetupArgument
                    0.00%     599ns         1     599ns     599ns     599ns  cudaConfigureCall
                    0.00%     153ns         1     153ns     153ns     153ns  cudaGetDeviceCount

==24085== Unified Memory profiling result:
Device "GeForce GTX TITAN X (0)"
   Count  Avg Size  Min Size  Max Size  Total Size  Total Time  Name
    3817  1.9988MB  256.00KB  2.0000MB  7.450584GB  662.1459ms  Host To Device
   68690  170.61KB  4.0000KB  0.9961MB  11.17603GB  983.5813ms  Device To Host
Total CPU Page faults: 34345
```

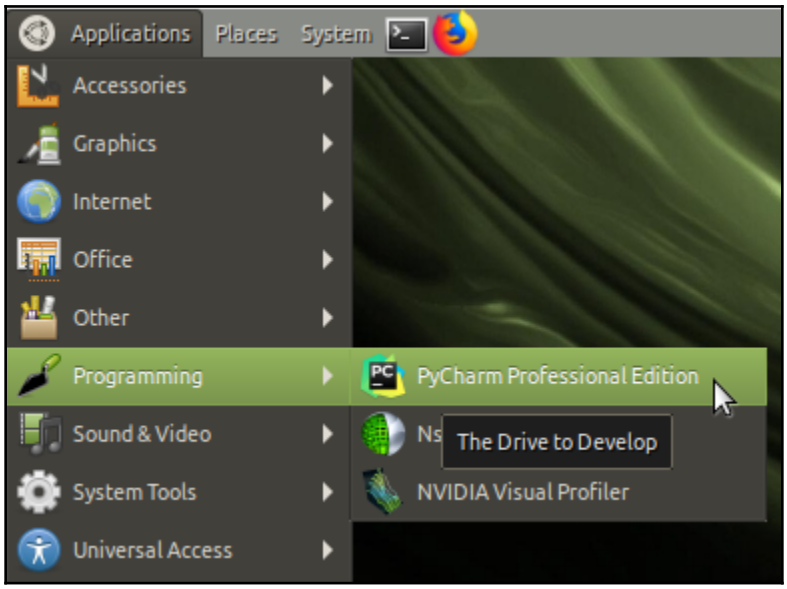

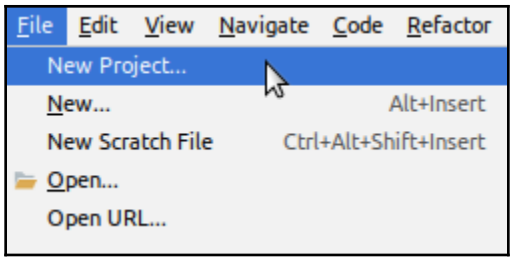
File
Edit
View
Navigate
Code
Refactor
New Project...
New...
Alt+Insert
New Scratch File
Ctrl+Alt+Shift+Insert
Open...
Open URL...

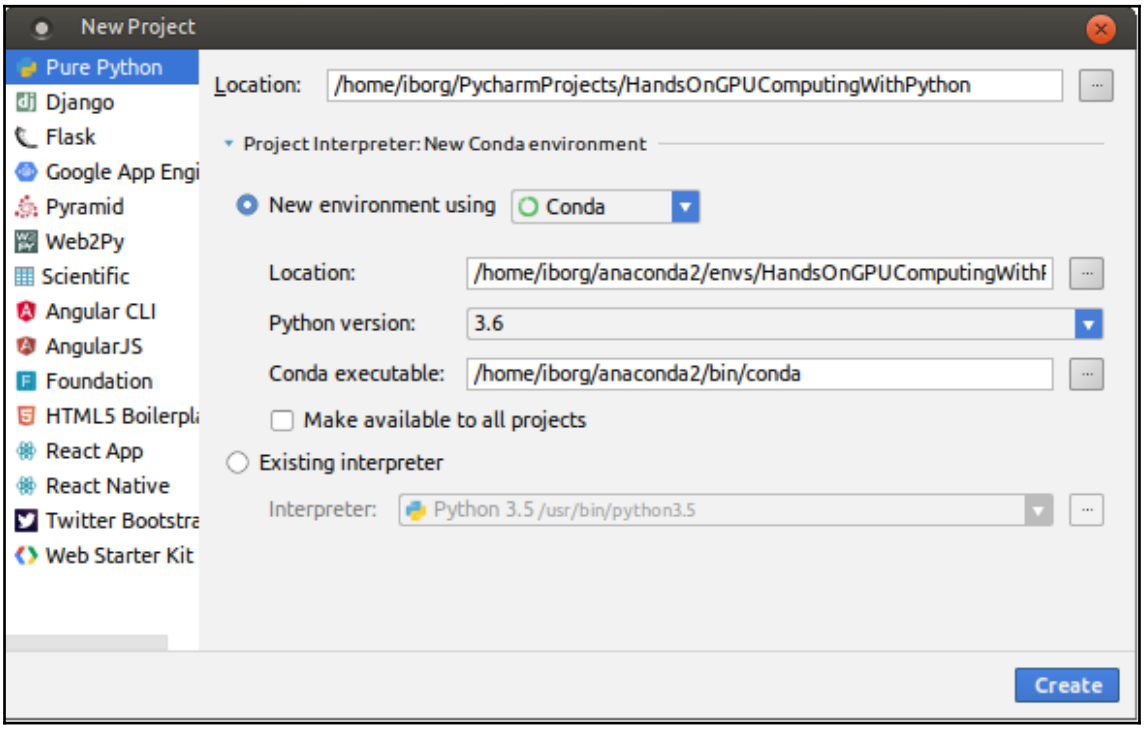
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location:
/home/iborg/PycharmProjects/HandsOnGPUComputingWithPython
Project Interpreter: New Conda environment
New environment using
Conda
Location:
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithF
Python version:
3.6
Conda executable:
/home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter:
Python 3.5 /usr/bin/python3.5
Create


import p
pycuda
<built-in>

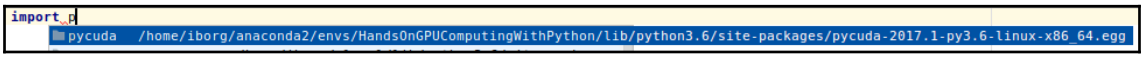
import p
pycuda /home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages/pycuda-2017.1-py3.6-linux-x86_64.egg

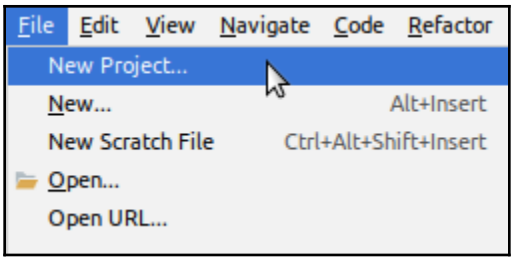
File
Edit
View
Navigate
Code
Refactor
New Project...
New...
Alt+Insert
New Scratch File
Ctrl+Alt+Shift+Insert
Open...
Open URL...

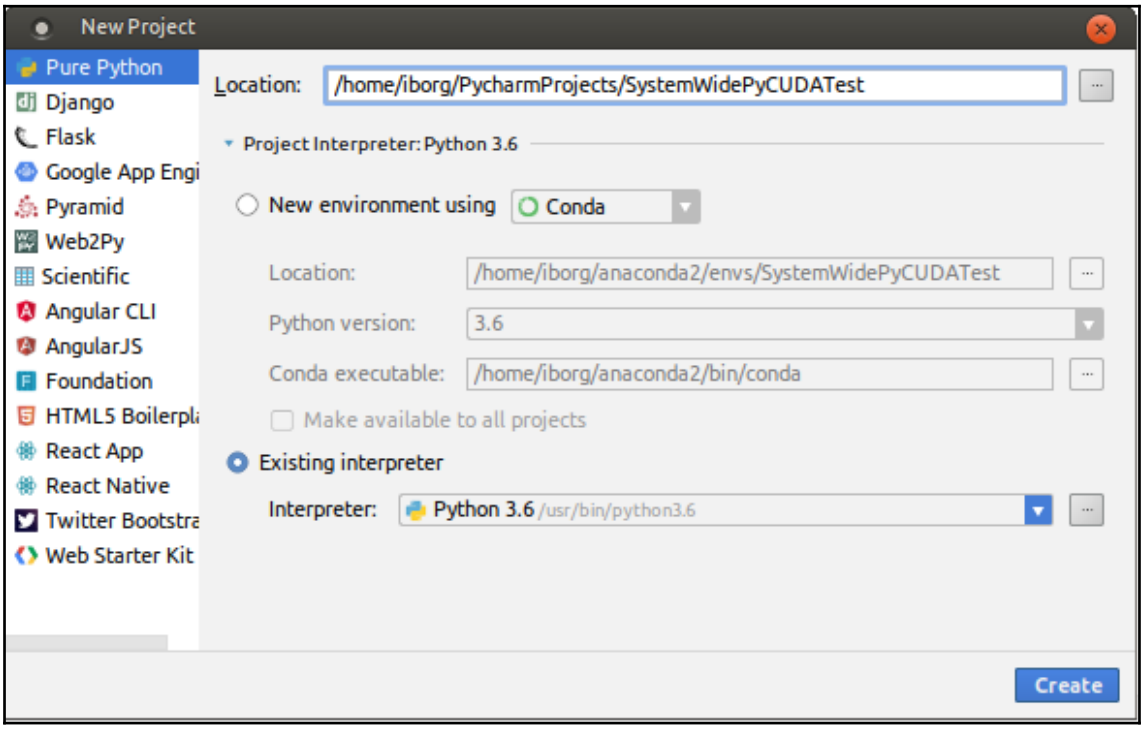
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location:
/home/iborg/PycharmProjects/SystemWidePyCUDATest
Project Interpreter: Python 3.6
New environment using
Conda
Location:
/home/iborg/anaconda2/envs/SystemWidePyCUDATest
Python version:
3.6
Conda executable:
/home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter:
Python 3.6 /usr/bin/python3.6
Create

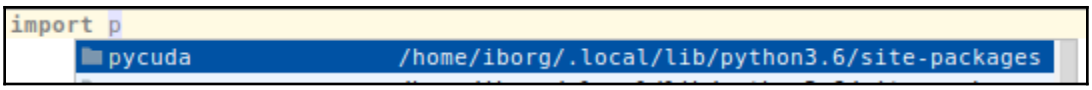
import p
pycuda
/home/iborg/.local/lib/python3.6/site-packages

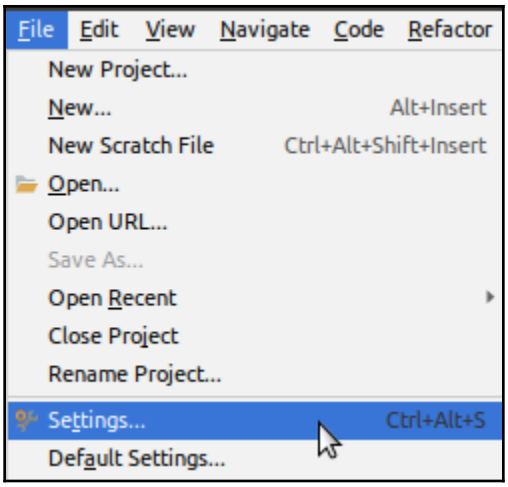
File
Edit
View
Navigate
Code
Refactor
New Project...
New... Alt+Insert
New Scratch File Ctrl+Alt+Shift+Insert
Open...
Open URL...
Save As...
Open Recent
Close Project
Rename Project...
Settings... Ctrl+Alt+S
Default Settings...

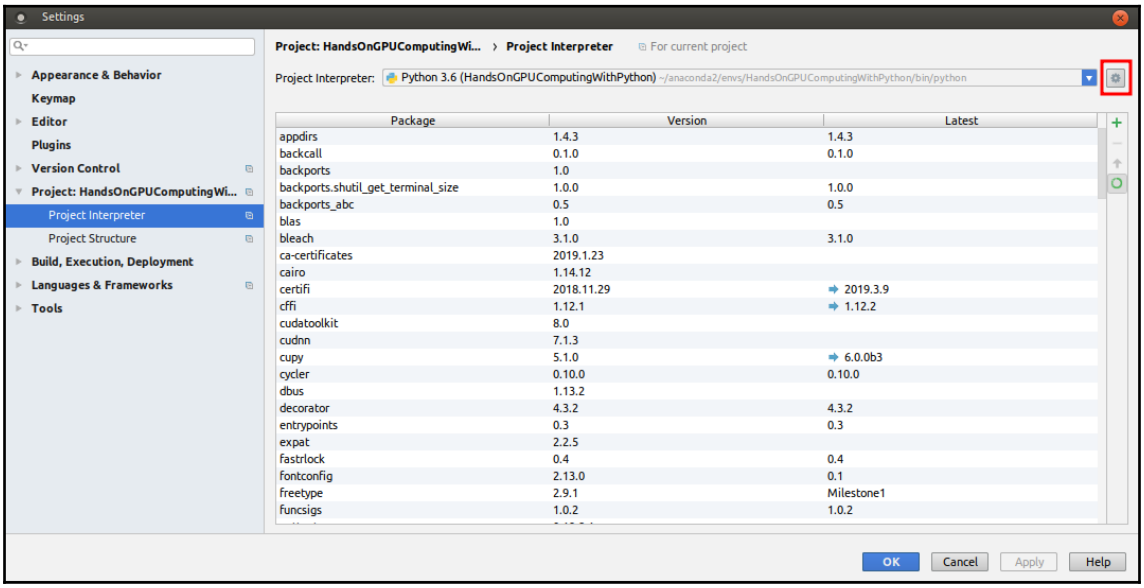
Settings
Project: HandsOnGPUComputingWi... › Project Interpreter
For current project
Appearance & Behavior
Keymap
Editor
Plugins
Version Control
Project: HandsOnGPUComputingWi...
Project Interpreter
Project Structure
Build, Execution, Deployment
Languages & Frameworks
Tools
Project Interpreter: Python 3.6 (HandsOnGPUComputingWithPython) ~/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python
Package Version Latest
appdirs 1.4.3 1.4.3
backcall 0.1.0 0.1.0
backports 1.0
backports.shutil_get_terminal_size 1.0.0 1.0.0
backports_abc 0.5 0.5
blas 1.0
bleach 3.1.0 3.1.0
ca-certificates 2019.1.23
cairo 1.14.12
certifi 2018.11.29 2019.3.9
cffi 1.12.1 1.12.2
cudatoolkit 8.0
cudnn 7.1.3
cupy 5.1.0 6.0.0b3
cycler 0.10.0 0.10.0
dbus 1.13.2
decorator 4.3.2 4.3.2
entrypoints 0.3 0.3
expat 2.2.5
fastrlock 0.4 0.4
fontconfig 2.13.0 0.1
freetype 2.9.1 Milestone1
funcsigs 1.0.2 1.0.2
OK
Cancel
Apply
Help

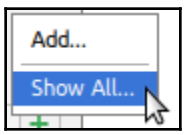
Add...
Show All...

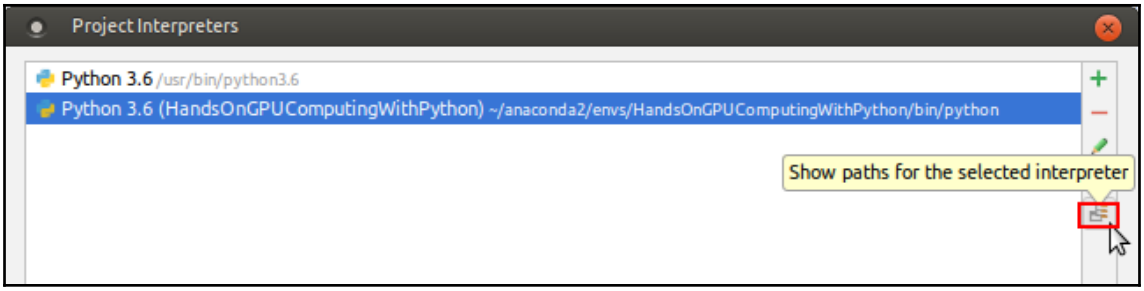
Project Interpreters
Python 3.6 /usr/bin/python3.6
Python 3.6 (HandsOnGPUComputingWithPython) ~/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python
Show paths for the selected interpreter

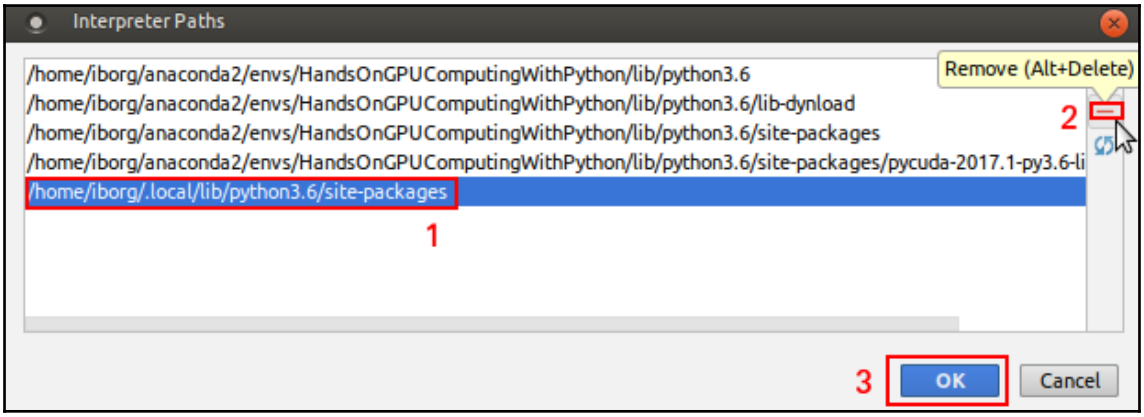
Interpreter Paths
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/lib-dynload
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages/pycuda-2017.1-py3.6-li
/home/iborg/.local/lib/python3.6/site-packages
Remove (Alt+Delete)
1
2
3
OK
Cancel

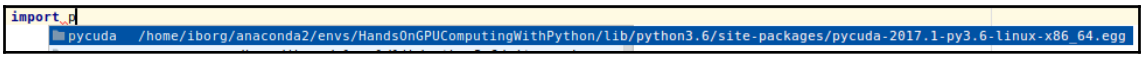
import p
pycuda /home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages/pycuda-2017.1-py3.6-linux-x86_64.egg

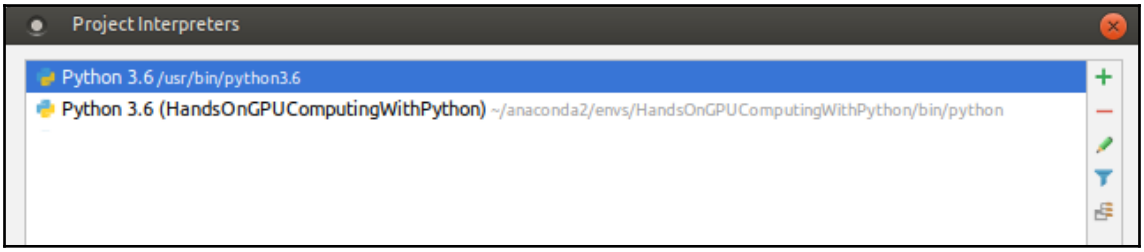
Project Interpreters
Python 3.6 /usr/bin/python3.6
Python 3.6 (HandsOnGPUComputingWithPython) ~/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python

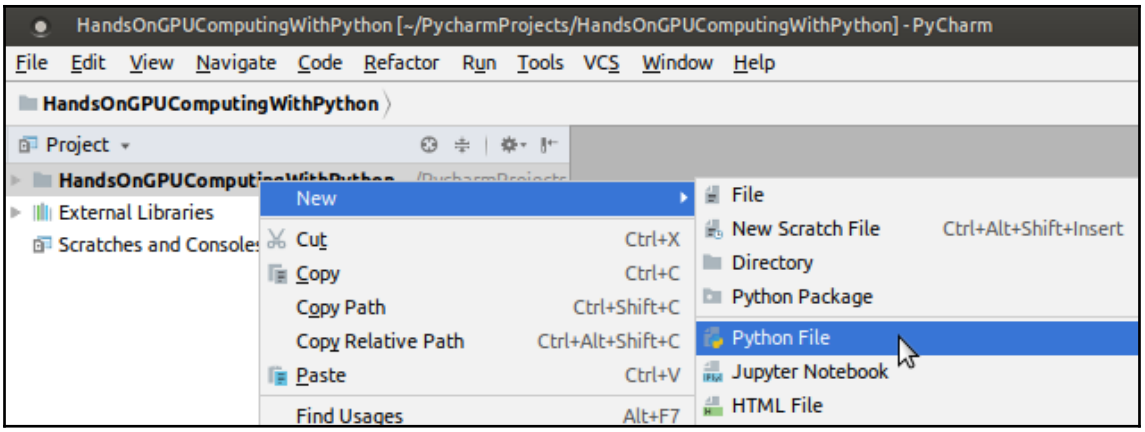
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File

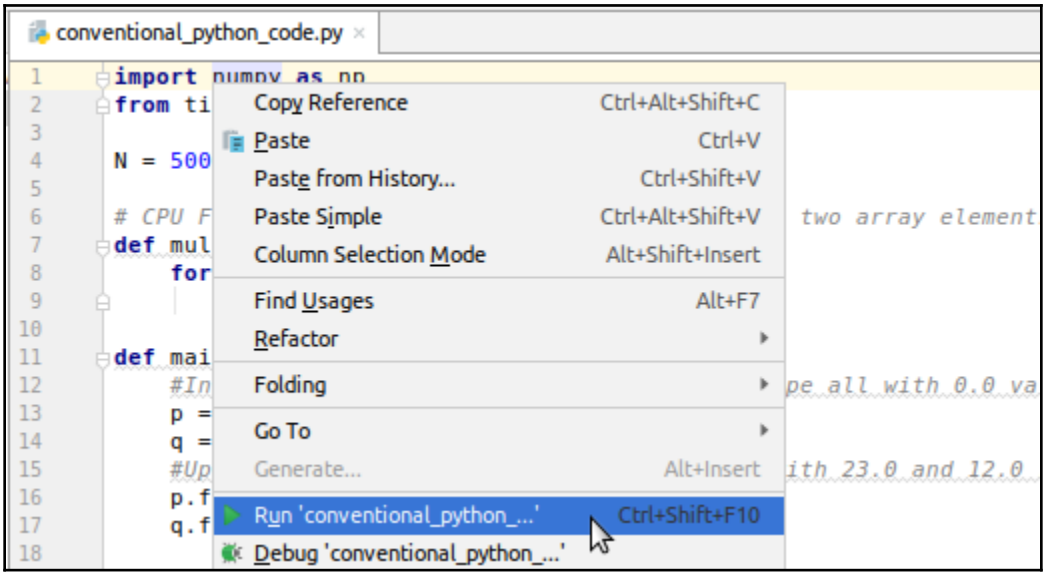
conventional_python_code.py
Copy Reference Ctrl+Alt+Shift+C
Paste Ctrl+V
Paste from History... Ctrl+Shift+V
Paste Simple Ctrl+Alt+Shift+V
Column Selection Mode Alt+Shift+Insert
Find Usages Alt+F7
Refactor
Folding
Go To
Generate... Alt+Insert
Run 'conventional_python_...' Ctrl+Shift+F10
Debug 'conventional_python_...'

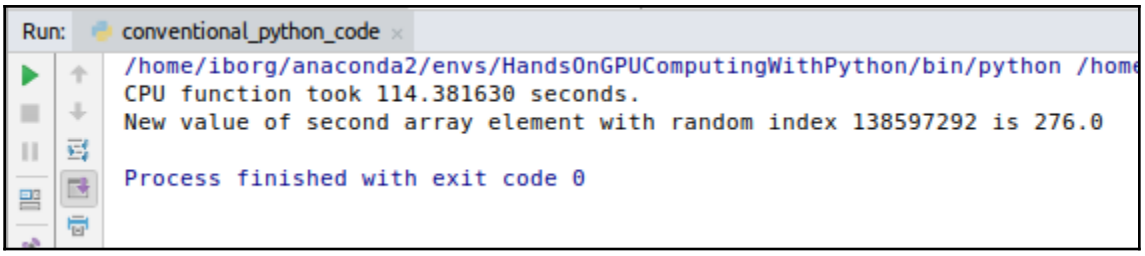
Run: conventional_python_code
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python /home
CPU function took 114.381630 seconds.
New value of second array element with random index 138597292 is 276.0
Process finished with exit code 0

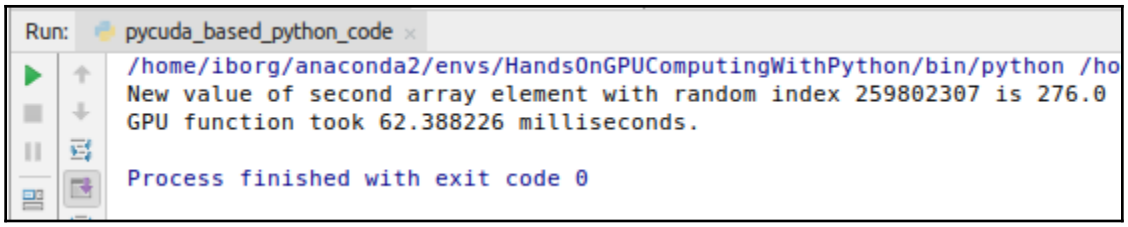
Run:
pycuda_based_python_code
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python /ho
New value of second array element with random index 259802307 is 276.0
GPU function took 62.388226 milliseconds.
Process finished with exit code 0

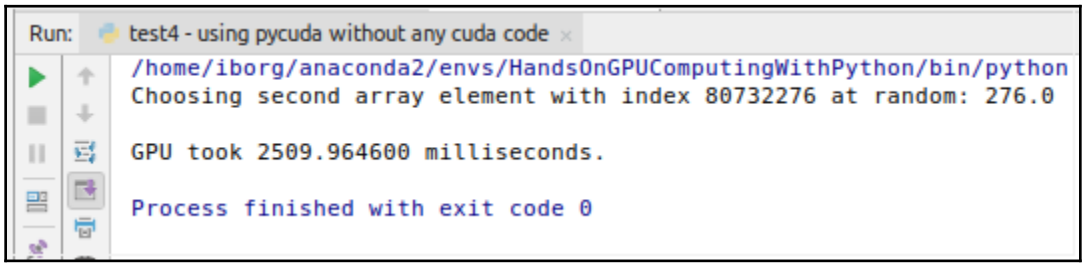
Run:
test4 - using pycuda without any cuda code
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python
Choosing second array element with index 80732276 at random: 276.0
GPU took 2509.964600 milliseconds.
Process finished with exit code 0

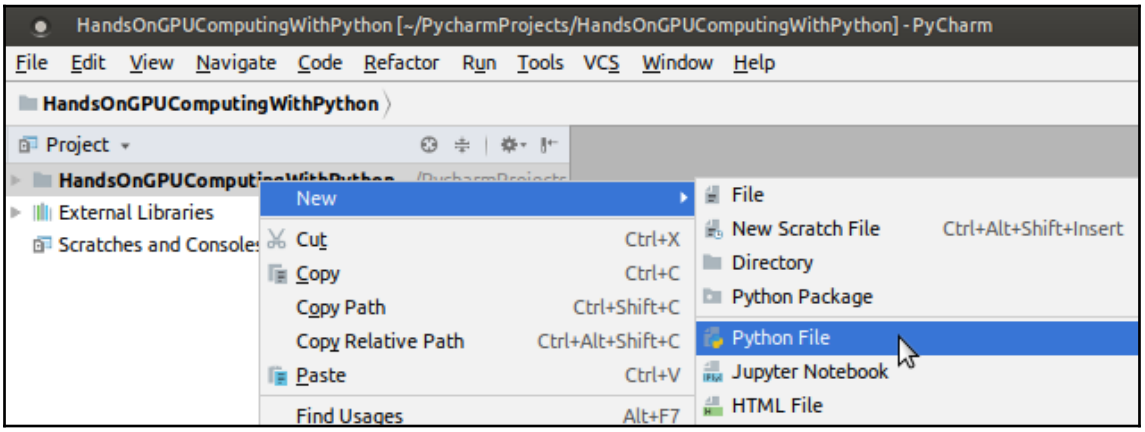
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File

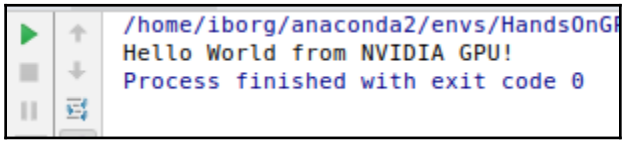
/home/iborg/anaconda2/envs/HandsOnG
Hello World from NVIDIA GPU!
Process finished with exit code 0

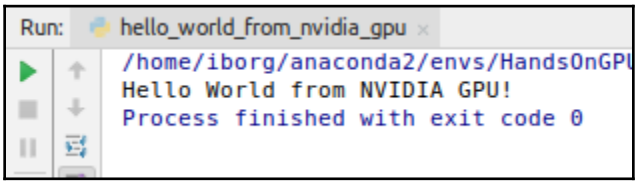
Run:
hello_world_from_nvidia_gpu
/home/iborg/anaconda2/envs/HandsOnGPU
Hello World from NVIDIA GPU!
Process finished with exit code 0

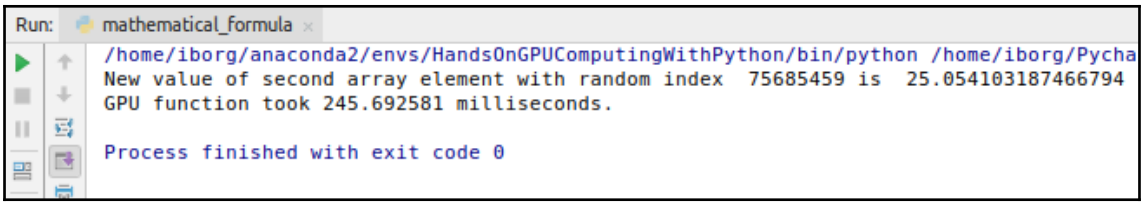
Run:
mathematical_formula
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python /home/iborg/Pycha
New value of second array element with random index 75685459 is 25.054103187466794
GPU function took 245.692581 milliseconds.
Process finished with exit code 0

Chapter 7: Working with ROCm and PyOpenCL

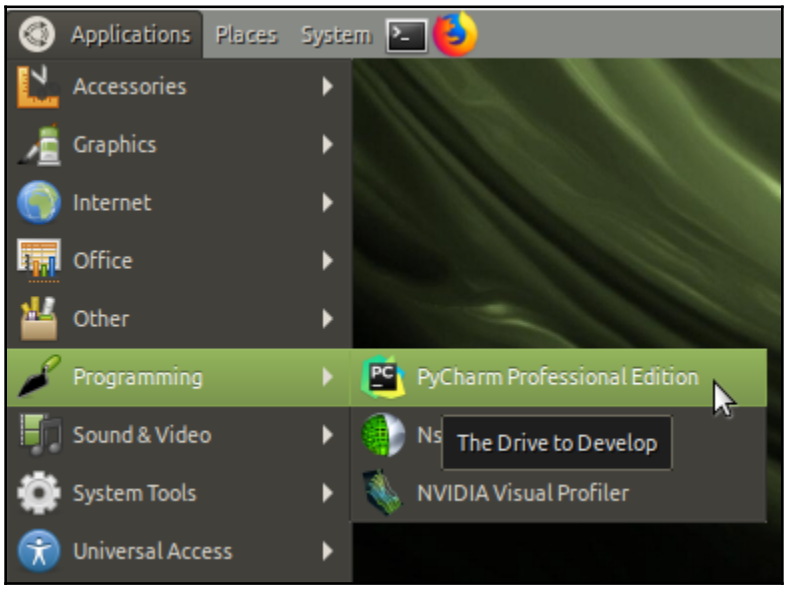

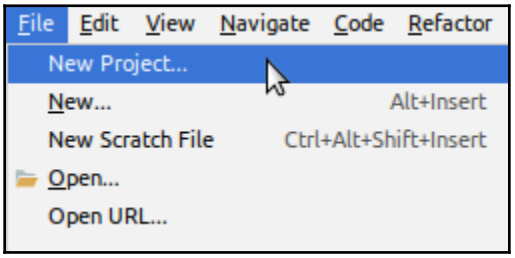

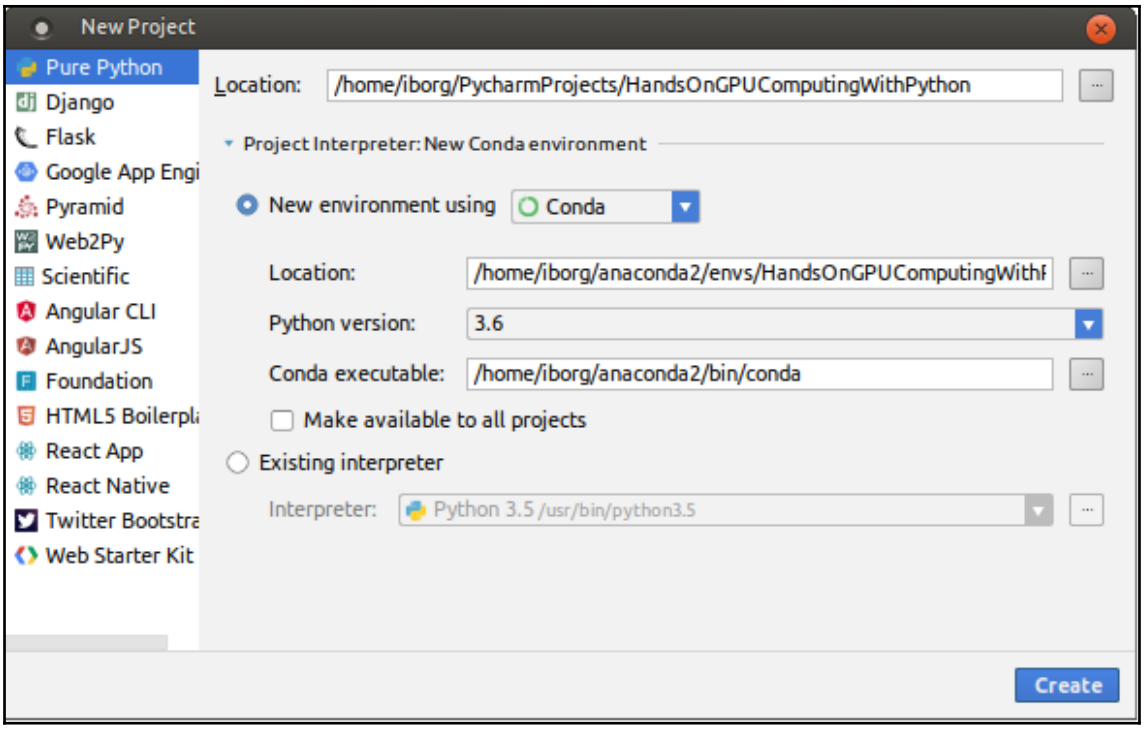
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location: /home/iborg/PycharmProjects/HandsOnGPUComputingWithPython
Project Interpreter: New Conda environment
New environment using Conda
Location: /home/iborg/anaconda2/envs/HandsOnGPUComputingWithF
Python version: 3.6
Conda executable: /home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter: Python 3.5 /usr/bin/python3.5
Create

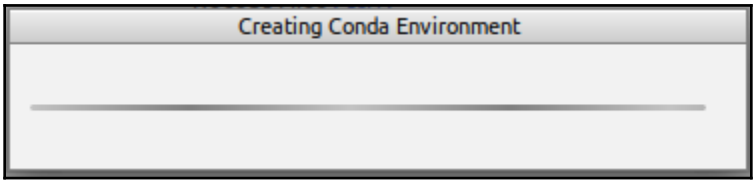
Creating Conda Environment

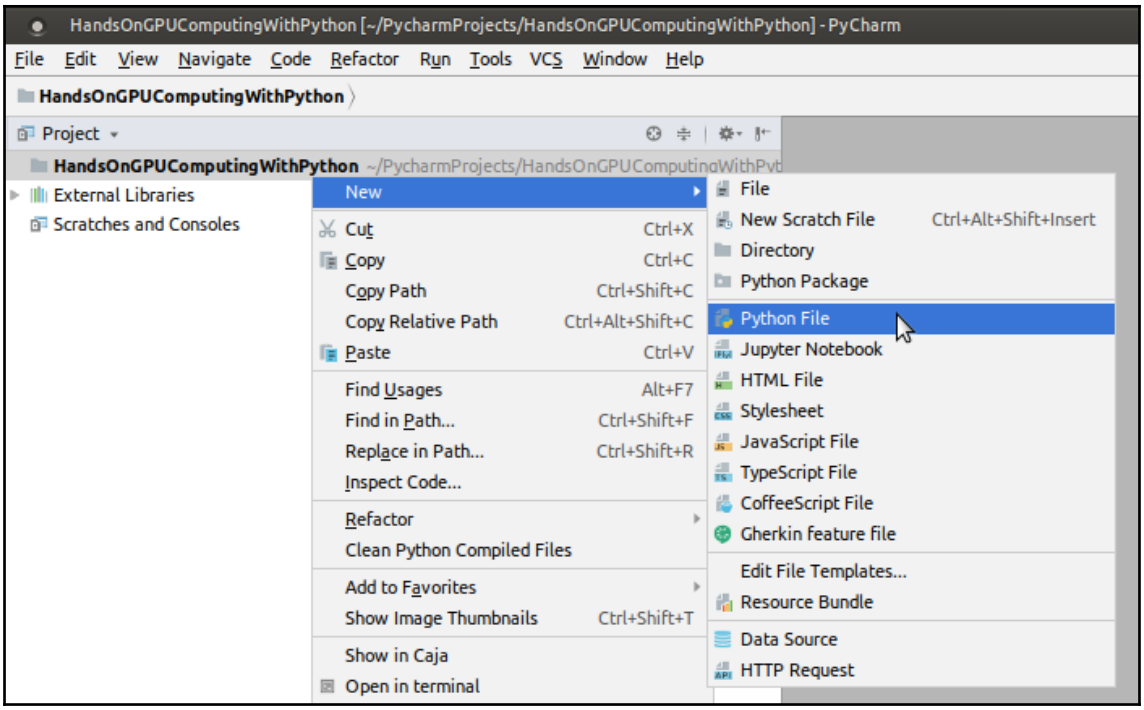
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
HandsOnGPUComputingWithPython ~/PycharmProjects/HandsOnGPUComputingWithPyt
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
Find in Path... Ctrl+Shift+F
Replace in Path... Ctrl+Shift+R
Inspect Code...
Refactor
Clean Python Compiled Files
Add to Favorites
Show Image Thumbnails Ctrl+Shift+T
Show in Caja
Open in terminal
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File
Stylesheet
JavaScript File
TypeScript File
CoffeeScript File
Gherkin feature file
Edit File Templates...
Resource Bundle
Data Source
HTTP Request


import pyo
pyopencl <built-in>

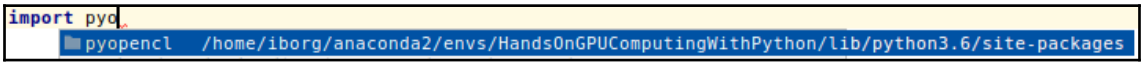
import pyo
pyopencl /home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages

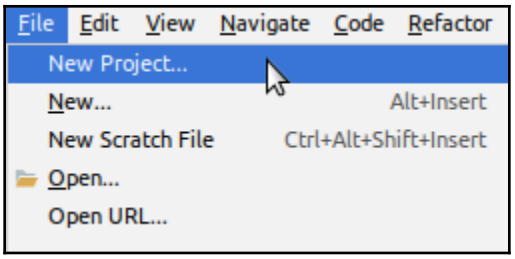
File Edit View Navigate Code Refactor
New Project...
New... Alt+Insert
New Scratch File Ctrl+Alt+Shift+Insert
Open...
Open URL...

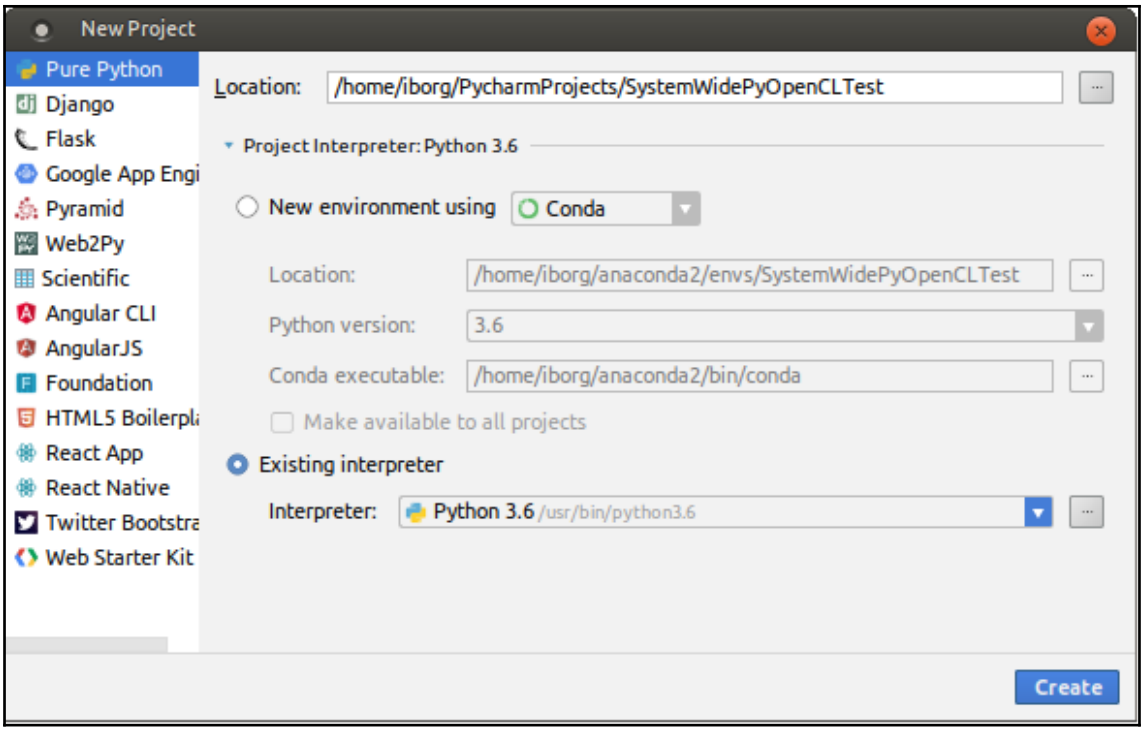
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location: /home/iborg/PycharmProjects/SystemWidePyOpenCLTest
Project Interpreter: Python 3.6
New environment using Conda
Location: /home/iborg/anaconda2/envs/SystemWidePyOpenCLTest
Python version: 3.6
Conda executable: /home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter: Python 3.6 /usr/bin/python3.6
Create

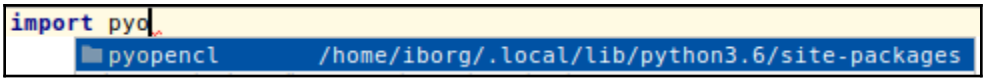
import pyo
pyopencl /home/iborg/.local/lib/python3.6/site-packages

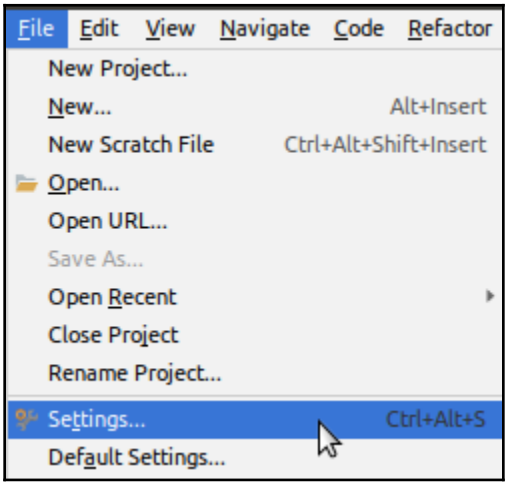
File Edit View Navigate Code Refactor
New Project...
New... Alt+Insert
New Scratch File Ctrl+Alt+Shift+Insert
Open...
Open URL...
Save As...
Open Recent
Close Project
Rename Project...
Settings... Ctrl+Alt+S
Default Settings...

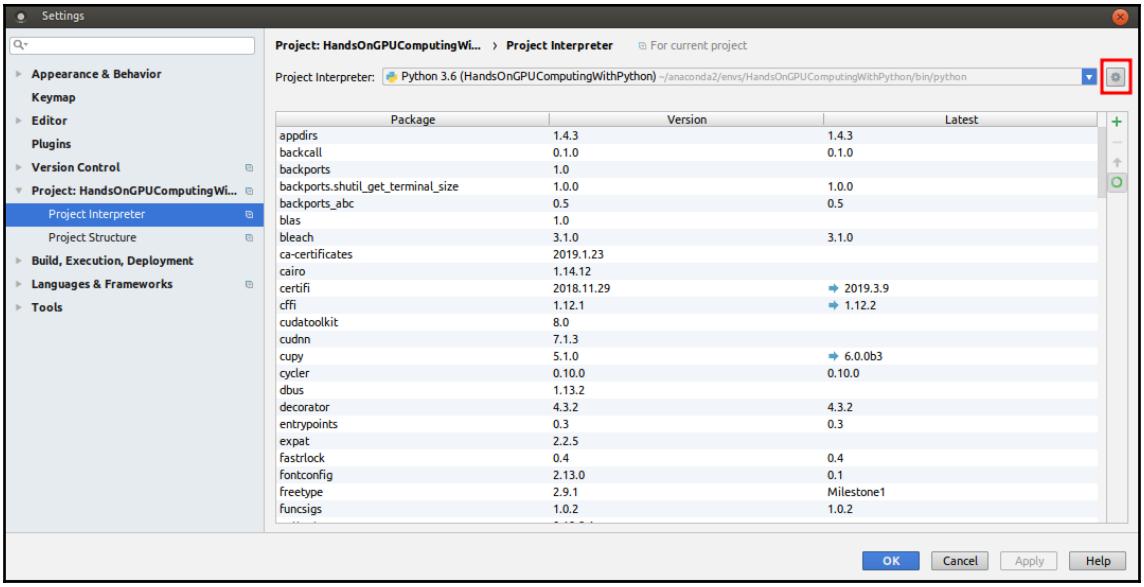
Settings
Appearance & Behavior
Keymap
Editor
Plugins
Version Control
Project: HandsOnGPUComputingWi...
Project Interpreter
Project Structure
Build, Execution, Deployment
Languages & Frameworks
Tools
Project: HandsOnGPUComputingWi... › Project Interpreter
For current project
Project Interpreter: Python 3.6 (HandsOnGPUComputingWithPython) ~/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python
Package
Version
Latest
appdirs 1.4.3 1.4.3
backcall 0.1.0 0.1.0
backports 1.0
backports.shutil_get_terminal_size 1.0.0 1.0.0
backports_abc 0.5 0.5
blas 1.0
bleach 3.1.0 3.1.0
ca-certificates 2019.1.23
cairo 1.14.12
certifi 2018.11.29 2019.3.9
cffi 1.12.1 1.12.2
cudatoolkit 8.0
cudnn 7.1.3
cupy 5.1.0 6.0.0b3
cycler 0.10.0 0.10.0
dbus 1.13.2
decorator 4.3.2 4.3.2
entrypoints 0.3 0.3
expat 2.2.5
fastrlock 0.4 0.4
fontconfig 2.13.0 0.1
freetype 2.9.1 Milestone1
funcsigs 1.0.2 1.0.2
OK
Cancel
Apply
Help

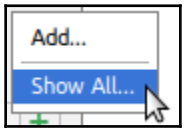
Add...
Show All...

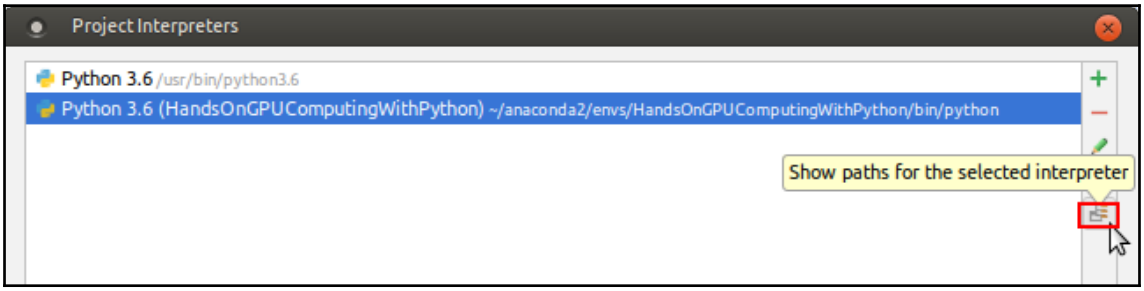
Project Interpreters
Python 3.6 /usr/bin/python3.6
Python 3.6 (HandsOnGPUComputingWithPython) ~/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python
Show paths for the selected interpreter

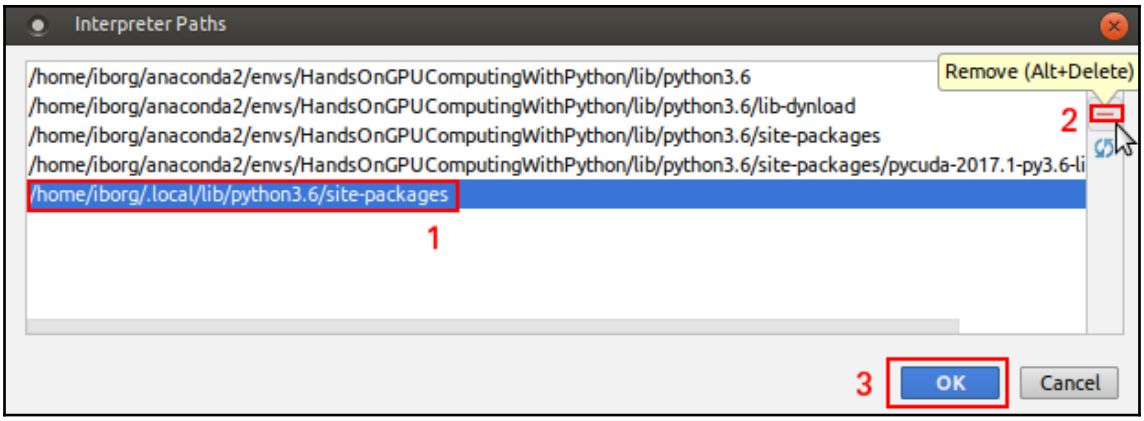
Interpreter Paths
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/lib-dynload
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages
/home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages/pycuda-2017.1-py3.6-li
/home/iborg/.local/lib/python3.6/site-packages
Remove (Alt+Delete)
1
2
3
OK
Cancel

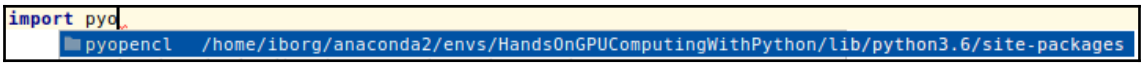
import pyo
pyopencl /home/iborg/anaconda2/envs/HandsOnGPUComputingWithPython/lib/python3.6/site-packages

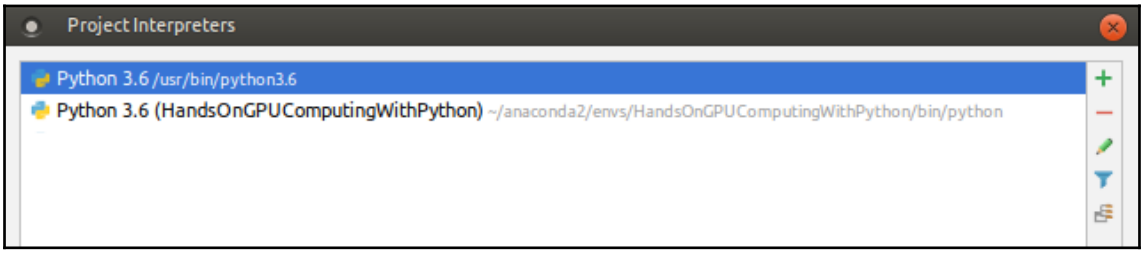
Project Interpreters
Python 3.6 /usr/bin/python3.6
Python 3.6 (HandsOnGPUComputingWithPython) ~/anaconda2/envs/HandsOnGPUComputingWithPython/bin/python

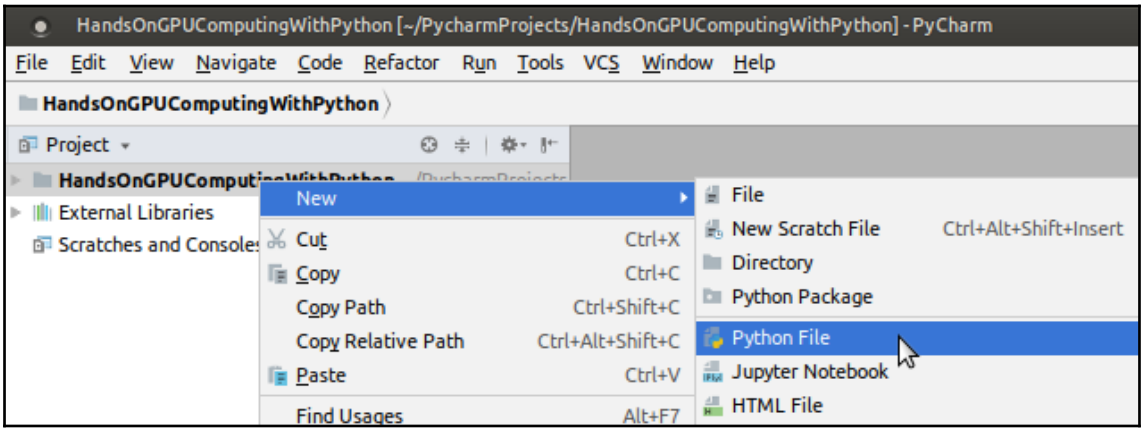
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File

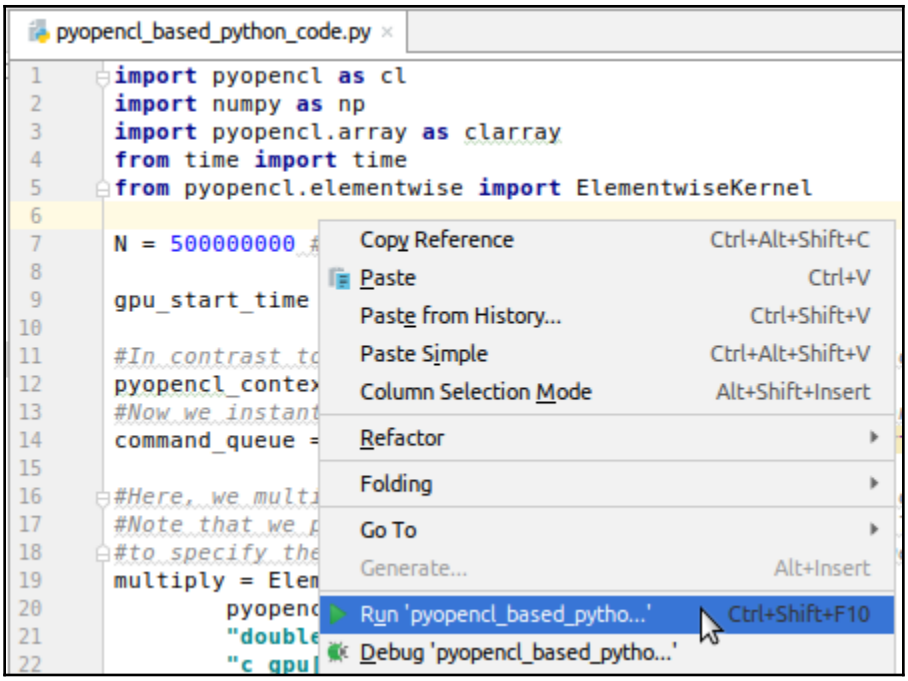


```
Run: pyopencl_based_python_code
GPU Kernel Function took 22.959746 milliseconds
GPU Time(Inclusive of memory transfer between host and device (GPU)): 6.434994697570801 seconds
New value of third array element with random index 340795000 is 276.0

Process finished with exit code 0
```

```
Run: pyopencl_version_of_pycuda_exam...
GPU Kernel Function took 17.676814 milliseconds
GPU Time(Inclusive of memory transfer between host and device (GPU)): 4.429342031478882 seconds
New value of third array element with random index 287997311 is 276.0

Process finished with exit code 0
```

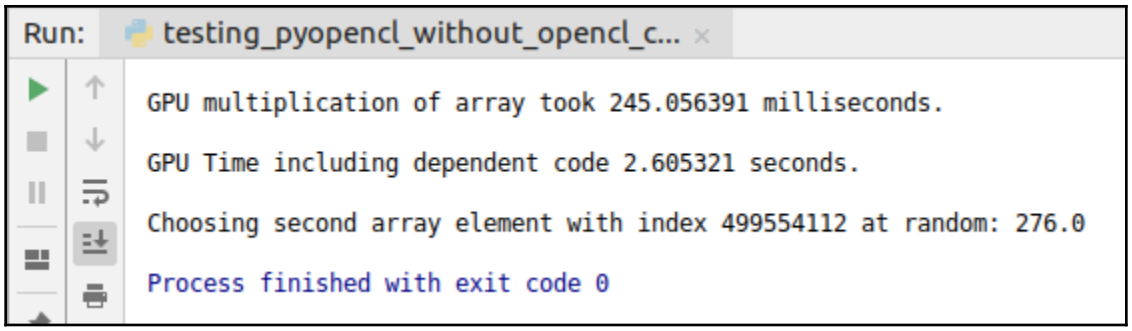
Run:
testing_pyopencl_without_opencl_c...
GPU multiplication of array took 245.056391 milliseconds.
GPU Time including dependent code 2.605321 seconds.
Choosing second array element with index 499554112 at random: 276.0
Process finished with exit code 0

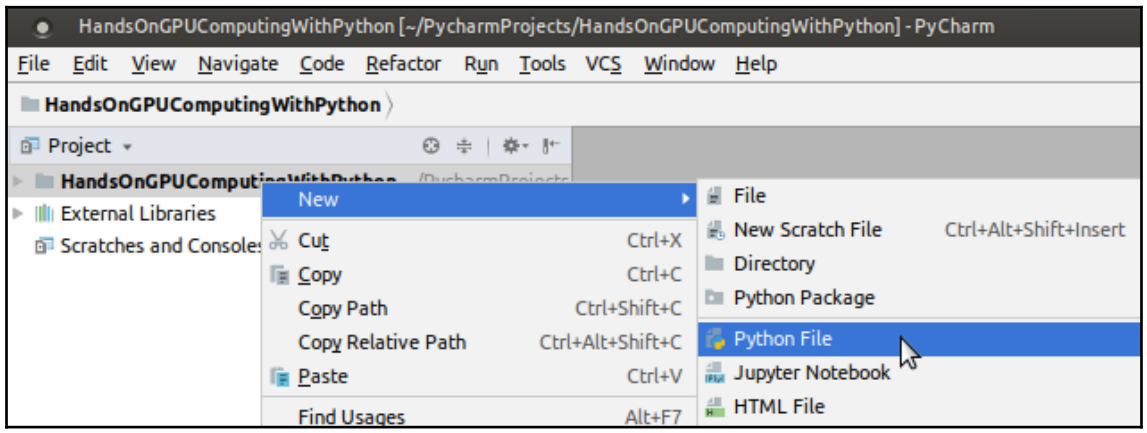
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File

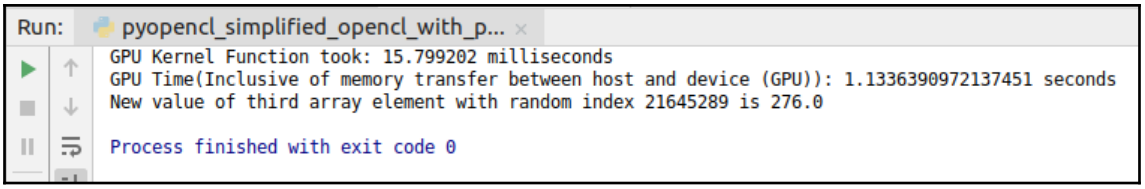
Run:
pyopencl_simplified_opencl_with_p...
GPU Kernel Function took: 15.799202 milliseconds
GPU Time(Inclusive of memory transfer between host and device (GPU)): 1.1336390972137451 seconds
New value of third array element with random index 21645289 is 276.0
Process finished with exit code 0

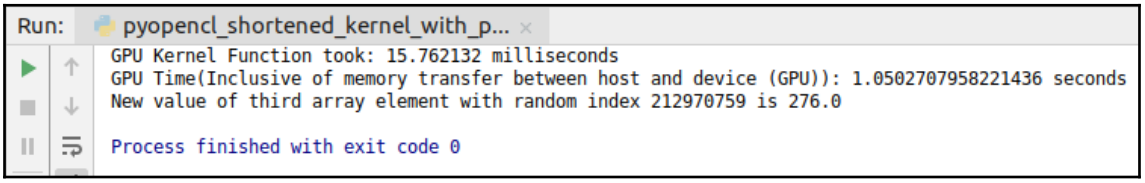
Run:
pyopencl_shortened_kernel_with_p...
GPU Kernel Function took: 15.762132 milliseconds
GPU Time(Inclusive of memory transfer between host and device (GPU)): 1.0502707958221436 seconds
New value of third array element with random index 212970759 is 276.0
Process finished with exit code 0

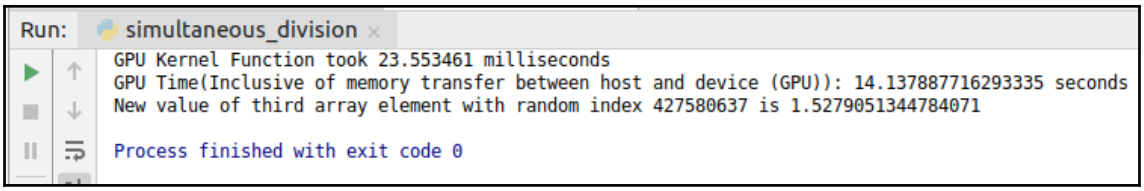
Run:
simultaneous_division
GPU Kernel Function took 23.553461 milliseconds
GPU Time(Inclusive of memory transfer between host and device (GPU)): 14.137887716293335 seconds
New value of third array element with random index 427580637 is 1.5279051344784071
Process finished with exit code 0

Chapter 8: Working with Anaconda, CuPy, and Numba for GPUs

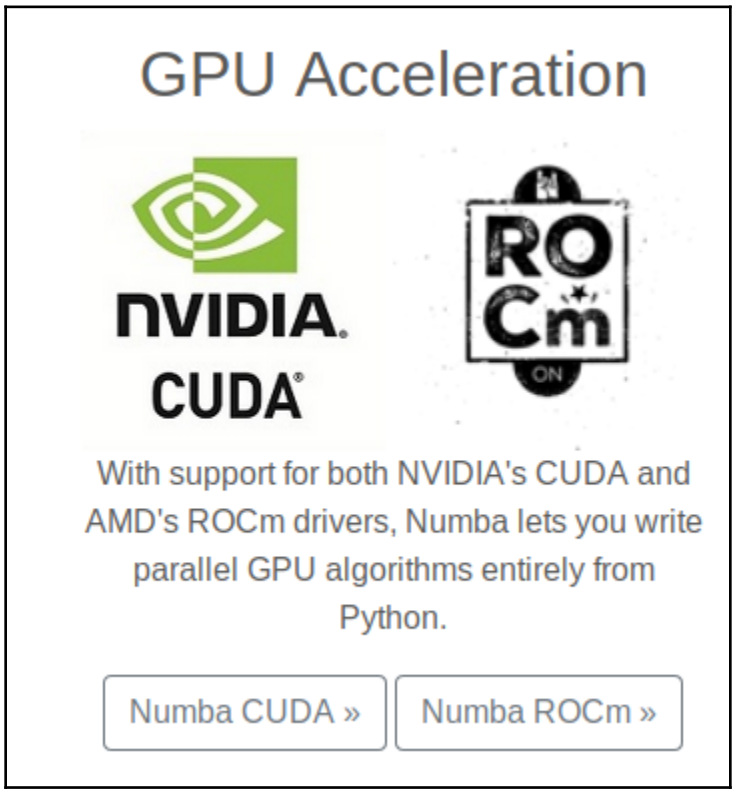

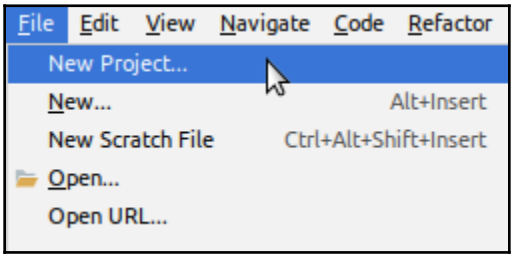

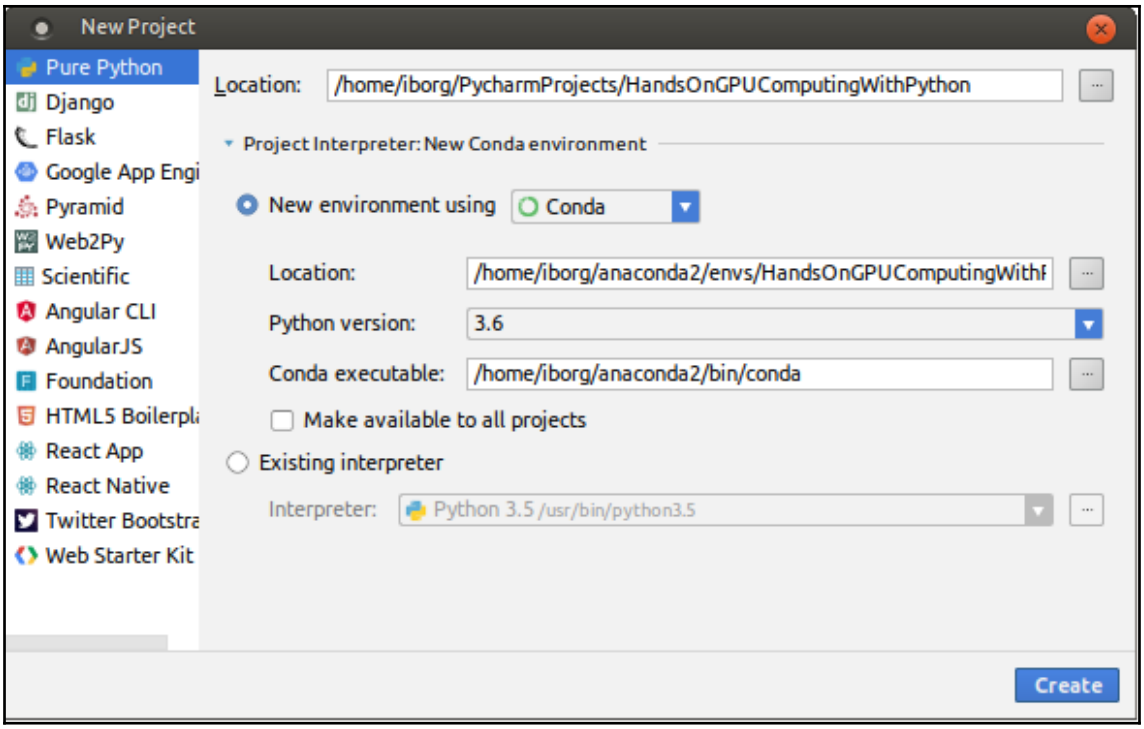
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location: /home/iborg/PycharmProjects/HandsOnGPUComputingWithPython
Project Interpreter: New Conda environment
New environment using Conda
Location: /home/iborg/anaconda2/envs/HandsOnGPUComputingWithF
Python version: 3.6
Conda executable: /home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter: Python 3.5 /usr/bin/python3.5
Create

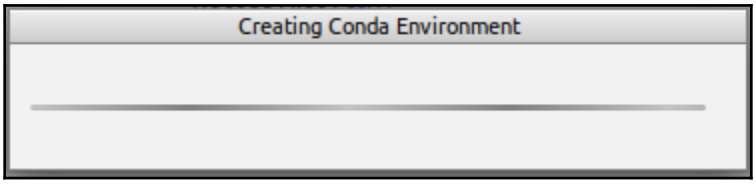
Creating Conda Environment

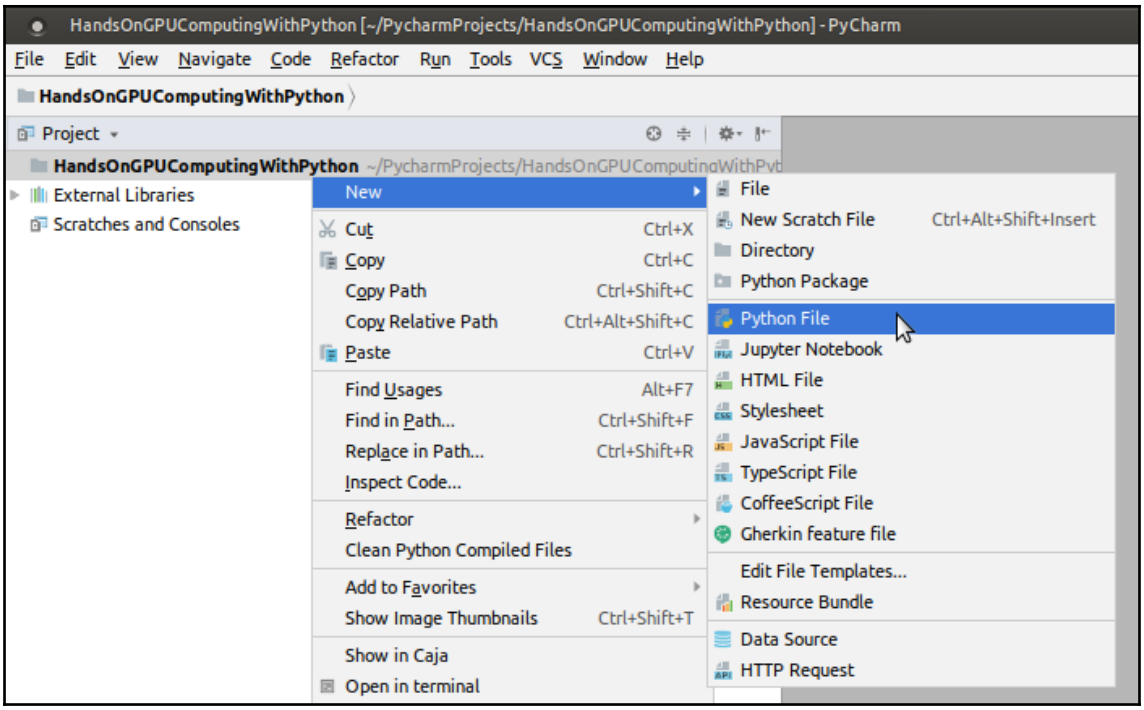
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
HandsOnGPUComputingWithPython ~/PycharmProjects/HandsOnGPUComputingWithPyt
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
Find in Path... Ctrl+Shift+F
Replace in Path... Ctrl+Shift+R
Inspect Code...
Refactor
Clean Python Compiled Files
Add to Favorites
Show Image Thumbnails Ctrl+Shift+T
Show in Caja
Open in terminal
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File
Stylesheet
JavaScript File
TypeScript File
CoffeeScript File
Gherkin feature file
Edit File Templates...
Resource Bundle
Data Source
HTTP Request


import cu
cupy


from cupy import cu
cuda cupy.cuda

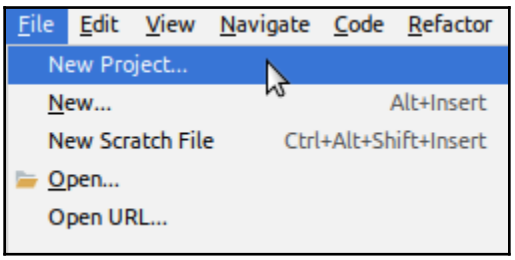
File Edit View Navigate Code Refactor
New Project...
New... Alt+Insert
New Scratch File Ctrl+Alt+Shift+Insert
Open...
Open URL...

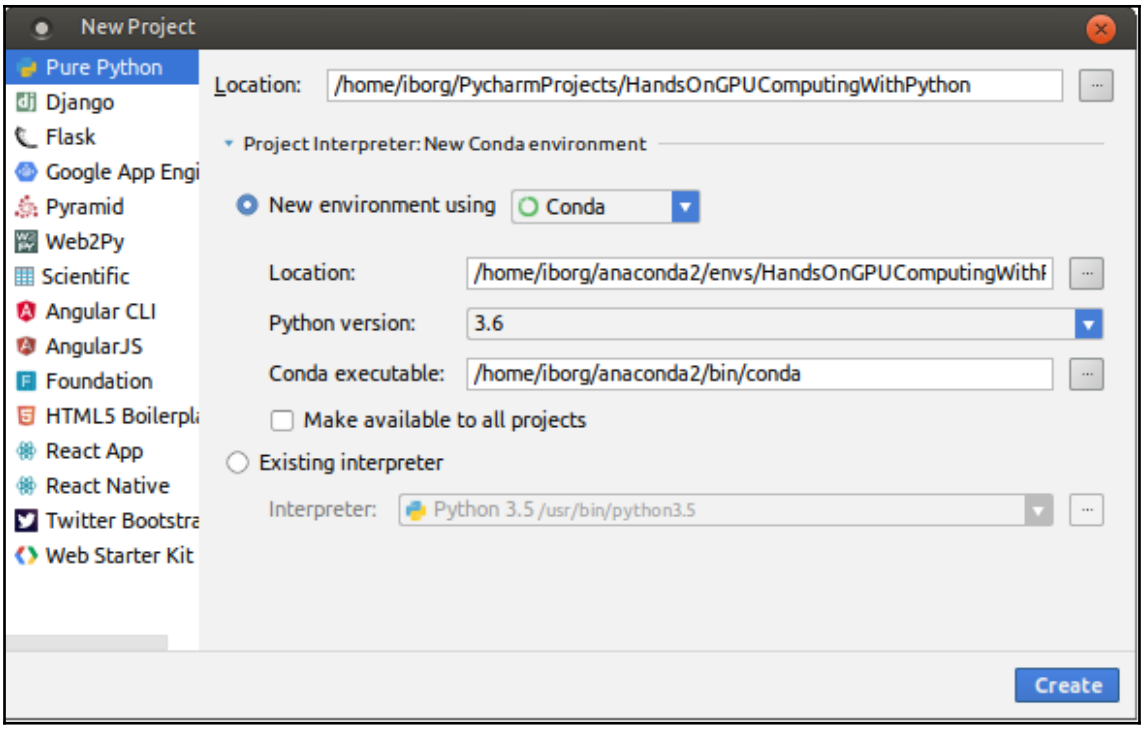
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location: /home/iborg/PycharmProjects/HandsOnGPUComputingWithPython
Project Interpreter: New Conda environment
New environment using Conda
Location: /home/iborg/anaconda2/envs/HandsOnGPUComputingWithF
Python version: 3.6
Conda executable: /home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter: Python 3.5 /usr/bin/python3.5
Create

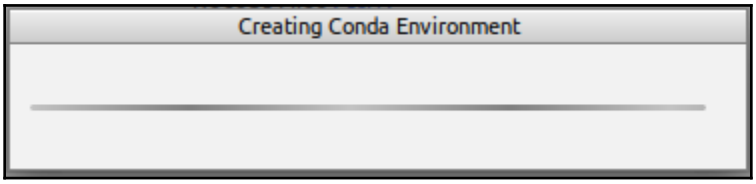
Creating Conda Environment

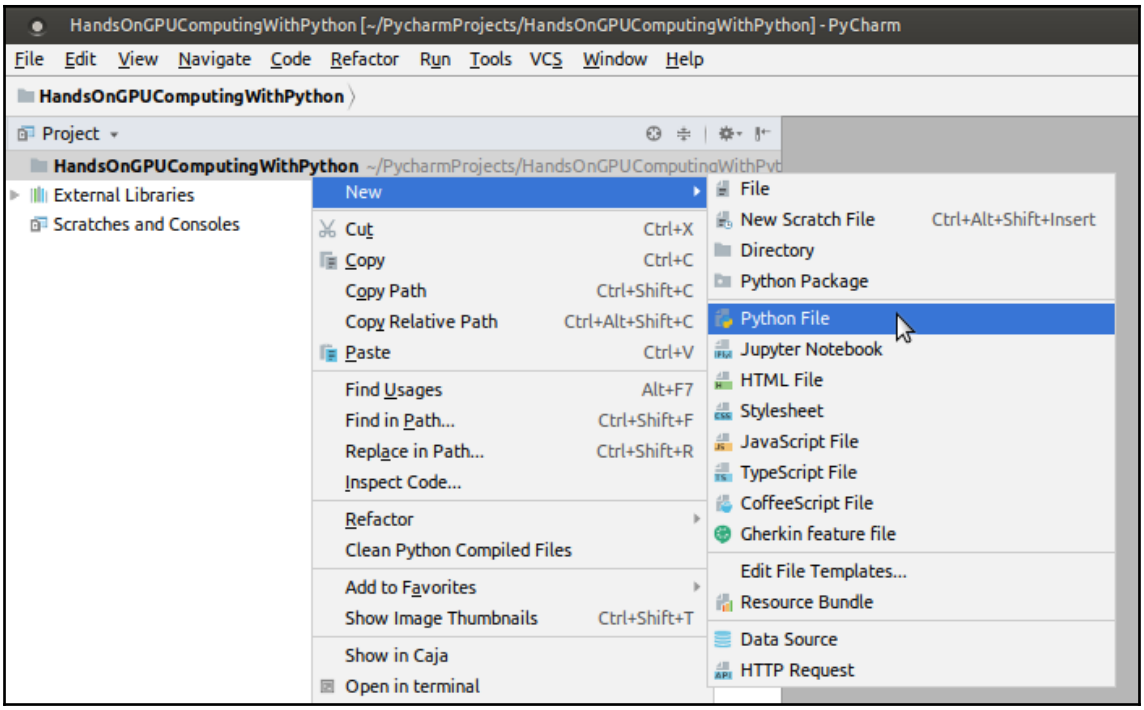
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
HandsOnGPUComputingWithPython ~/PycharmProjects/HandsOnGPUComputingWithPyt
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
Find in Path... Ctrl+Shift+F
Replace in Path... Ctrl+Shift+R
Inspect Code...
Refactor
Clean Python Compiled Files
Add to Favorites
Show Image Thumbnails Ctrl+Shift+T
Show in Caja
Open in terminal
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File
Stylesheet
JavaScript File
TypeScript File
CoffeeScript File
Gherkin feature file
Edit File Templates...
Resource Bundle
Data Source
HTTP Request


import nu
numba


from numba import cu
cuda (numba)


from numba import r
roc (numba)

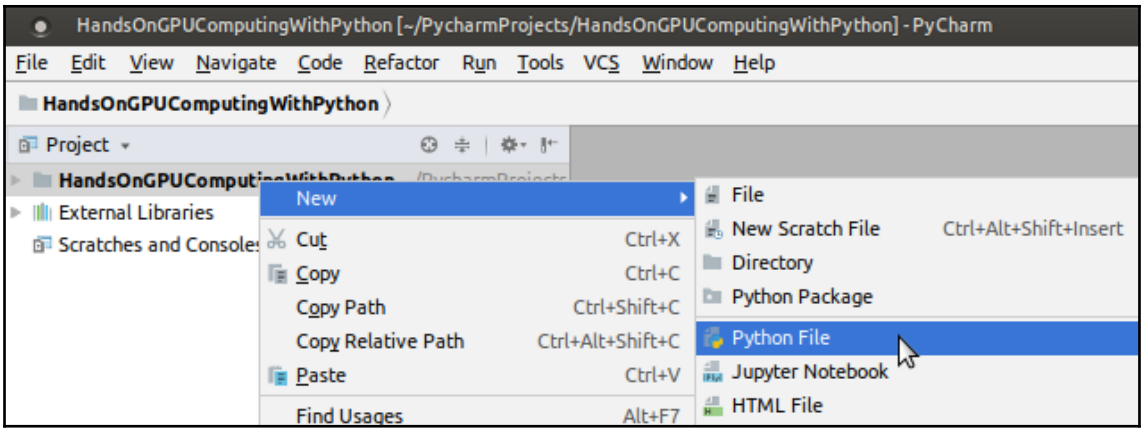


```
Run: replacing_numpy_code_with_cupy
Choosing array element 132060554 at random:
Random array element value is 276.000000
<CUDA Device 0>
CuPy on GPU took 0.433645 seconds.

Process finished with exit code 0
```

```
Run: raw_cuda_kernel_on_cupy
[    0     1     4     9    16    25    36    49    64    81   100   121
   144   169   196   225   256   289   324   361   400   441   484   529
   576   625   676   729   784   841   900   961  1024  1089  1156  1225
  1296  1369  1444  1521  1600  1681  1764  1849  1936  2025  2116  2209
  2304  2401  2500  2601  2704  2809  2916  3025  3136  3249  3364  3481
  3600  3721  3844  3969  4096  4225  4356  4489  4624  4761  4900  5041
  5184  5329  5476  5625  5776  5929  6084  6241  6400  6561  6724  6889
  7056  7225  7396  7569  7744  7921  8100  8281  8464  8649  8836  9025
  9216  9409  9604  9801 10000 10201 10404 10609 10816 11025 11236 11449
 11664 11881 12100 12321 12544 12769 12996 13225 13456 13689 13924 14161
 14400 14641 14884 15129 15376 15625 15876 16129 16384 16641 16900 17161
 17424 17689 17956 18225 18496 18769 19044 19321 19600 19881 20164 20449
 20736 21025 21316 21609 21904 22201 22500 22801 23104 23409 23716 24025
 24336 24649 24964 25281 25600 25921 26244 26569 26896 27225 27556 27889
 28224 28561 28900 29241 29584 29929 30276 30625 30976 31329 31684 32041
 32400 32761 33124 33489 33856 34225 34596 34969 35344 35721 36100 36481
 36864 37249 37636 38025 38416 38809 39204 39601 40000 40401 40804 41209
 41616 42025 42436 42849 43264 43681 44100 44521 44944 45369 45796 46225
 46656 47089 47524 47961 48400 48841 49284 49729 50176 50625 51076 51529
 51984 52441 52900 53361 53824 54289 54756 55225 55696 56169 56644 57121
 57600 58081 58564 59049 59536 60025 60516 61009 61504 62001 62500 63001
 63504 64009 64516 65025 65536 66049 66564 67081 67600 68121 68644 69169
 69696 70225 70756 71289 71824 72361 72900 73441 73984 74529 75076 75625
 76176 76729 77284 77841 78400 78961 79524 80089 80656 81225 81796 82369
 82944 83521 84100 84681 85264 85849 86436 87025 87616 88209 88804 89401]

Process finished with exit code 0
```

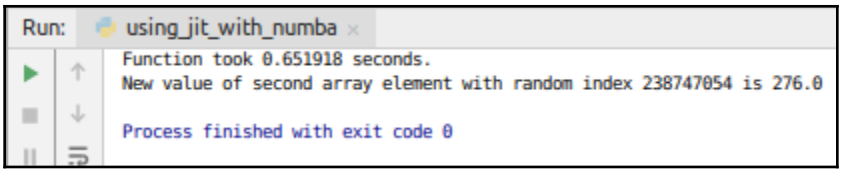
Run:
using_jit_with_numba
Function took 0.651918 seconds.
New value of second array element with random index 238747054 is 276.0
Process finished with exit code 0

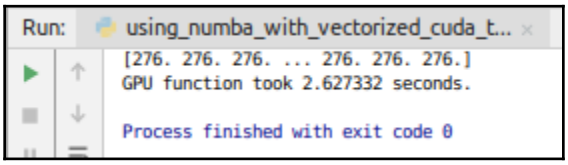
Run:
using_numba_with_vectorized_cuda_t...
[276. 276. 276. ... 276. 276. 276.]
GPU function took 2.627332 seconds.
Process finished with exit code 0

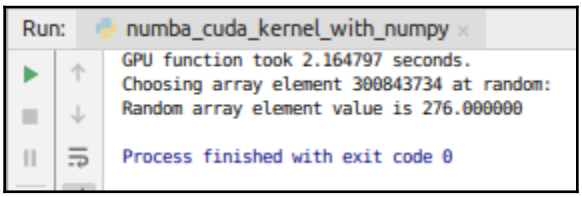
Run:
numba_cuda_kernel_with_numpy
GPU function took 2.164797 seconds.
Choosing array element 300843734 at random:
Random array element value is 276.000000
Process finished with exit code 0

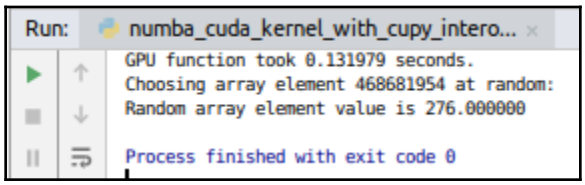
Run:
numba_cuda_kernel_with_cupy_intero...
GPU function took 0.131979 seconds.
Choosing array element 468681954 at random:
Random array element value is 276.000000
Process finished with exit code 0

Chapter 9: Containerization on GPU-Enabled Platforms

```
activate        activate_this.py   pip*       python3*
activate.csh    activate.xsh       pip3*      python3.6@
activate.fish   easy_install*      pip3.6*    python-config*
activate.ps1    easy_install-3.6*  python@    wheel*
```

```
abc.py@                          heapq.py@                    re.py@
base64.py@                       hmac.py@                     rlcompleter.py@
bisect.py@                       importlib@                   shutil.py@
_bootlocale.py@                  imp.py@                      site-packages/
codecs.py@                       io.py@                       site.py
collections@                     keyword.py@                  sre_compile.py@
_collections_abc.py@             lib-dynload@                 sre_constants.py@
config-3.6m-x86_64-linux-gnu@    LICENSE.txt@                 sre_parse.py@
copy.py@                         linecache.py@                stat.py@
copyreg.py@                      locale.py@                   struct.py@
distutils/                       no-global-site-packages.txt  tarfile.py@
_dummy_thread.py@                ntpath.py@                   tempfile.py@
encodings@                       operator.py@                 tokenize.py@
enum.py@                         orig-prefix.txt              token.py@
fnmatch.py@                      os.py@                       types.py@
functools.py@                    posixpath.py@                warnings.py@
__future__.py@                   __pycache__/                 weakref.py@
genericpath.py@                  random.py@                   _weakrefset.py@
hashlib.py@                      reprlib.py@
```

```
abc.cpython-36.pyc                  os.cpython-36.pyc
base64.cpython-36.pyc               posixpath.cpython-36.pyc
bisect.cpython-36.pyc               random.cpython-36.pyc
_bootlocale.cpython-36.pyc          re.cpython-36.pyc
codecs.cpython-36.pyc               reprlib.cpython-36.pyc
_collections_abc.cpython-36.pyc     shutil.cpython-36.pyc
copy.cpython-36.pyc                 site.cpython-36.pyc
copyreg.cpython-36.pyc              sre_compile.cpython-36.pyc
enum.cpython-36.pyc                 sre_constants.cpython-36.pyc
fnmatch.cpython-36.pyc              sre_parse.cpython-36.pyc
functools.cpython-36.pyc            stat.cpython-36.pyc
__future__.cpython-36.pyc           struct.cpython-36.pyc
genericpath.cpython-36.pyc          tarfile.cpython-36.pyc
hashlib.cpython-36.pyc              tempfile.cpython-36.pyc
heapq.cpython-36.pyc                token.cpython-36.pyc
hmac.cpython-36.pyc                 tokenize.cpython-36.pyc
io.cpython-36.pyc                   types.cpython-36.pyc
keyword.cpython-36.pyc              warnings.cpython-36.pyc
linecache.cpython-36.pyc            weakref.cpython-36.pyc
locale.cpython-36.pyc               _weakrefset.cpython-36.pyc
operator.cpython-36.pyc
```

```
easy_install.py         pkg_resources/   setuptools-41.0.0.dist-info/
pip/                    __pycache__/     wheel/
pip-19.0.3.dist-info/   setuptools/      wheel-0.33.1.dist-info/
```

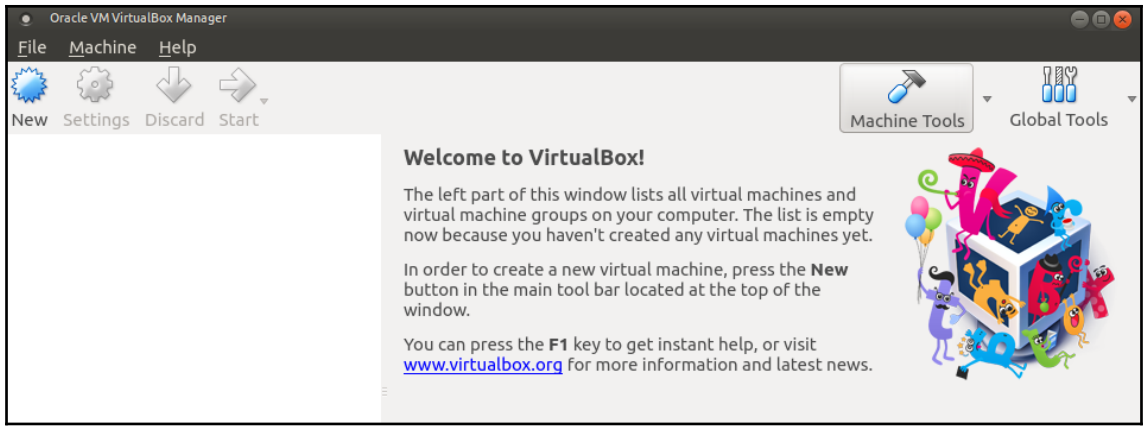
Oracle VM VirtualBox Manager
File Machine Help
New Settings Discard Start
Machine Tools
Global Tools
Welcome to VirtualBox!
The left part of this window lists all virtual machines and virtual machine groups on your computer. The list is empty now because you haven't created any virtual machines yet.
In order to create a new virtual machine, press the New button in the main tool bar located at the top of the window.
You can press the F1 key to get instant help, or visit www.virtualbox.org for more information and latest news.

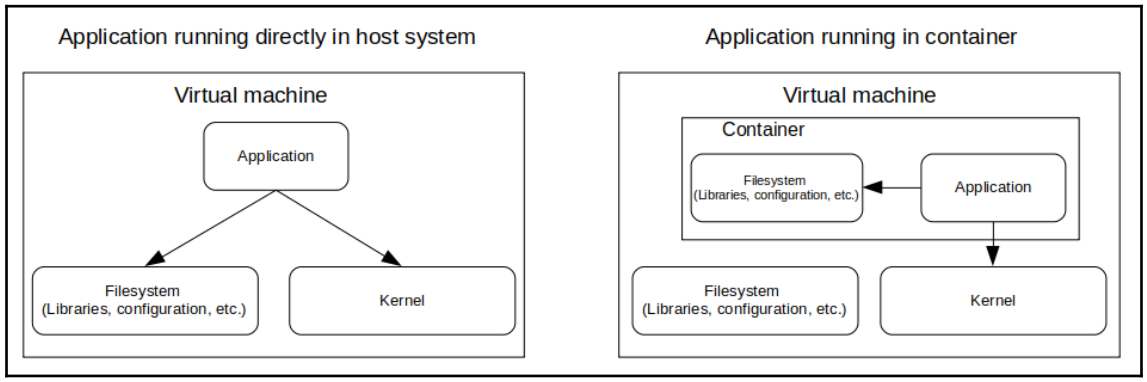
Application running directly in host system
Virtual machine
Application
Filesystem
(Libraries, configuration, etc.)
Kernel
Application running in container
Virtual machine
Container
Filesystem
(Libraries, configuration, etc.)
Application
Filesystem
(Libraries, configuration, etc.)
Kernel

```
from tensorflow.python.client import device_lib
device_lib.list_local_devices()
```

```
[name: "/device:CPU:0"
 device_type: "CPU"
 memory_limit: 268435456
 locality {
 }
 incarnation: 8097481313537670403, name: "/device:XLA_CPU:0"
 device_type: "XLA_CPU"
 memory_limit: 17179869184
 locality {
 }
 incarnation: 3509578143695564125
 physical_device_desc: "device: XLA_CPU device", name: "/device:XLA_GPU:0"
 device_type: "XLA_GPU"
 memory_limit: 17179869184
 locality {
 }
 incarnation: 8079698124687756107
 physical_device_desc: "device: XLA_GPU device", name: "/device:GPU:0"
 device_type: "GPU"
 memory_limit: 14800692839
 locality {
   bus_id: 1
   links {
   }
 }
 incarnation: 1379948111412660623
 physical_device_desc: "device: 0, name: Tesla T4, pci bus id: 0000:00:04.0, compute capability: 7.5"]
```

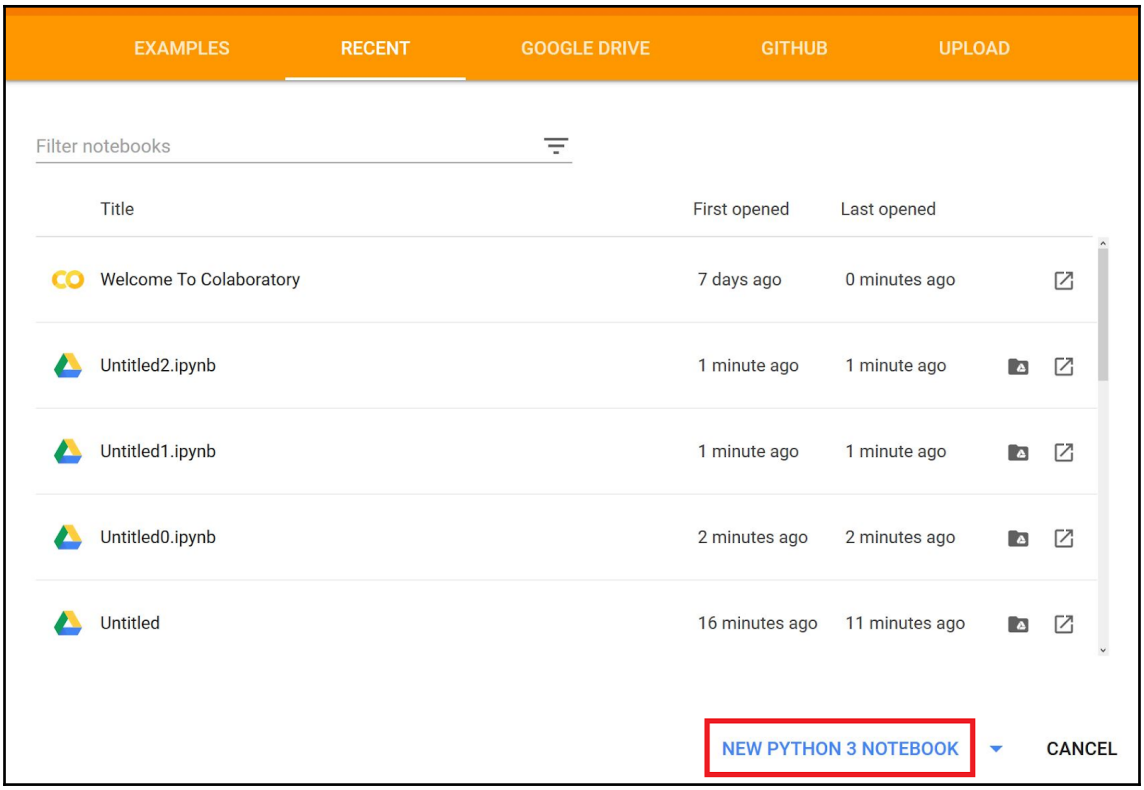
EXAMPLES
RECENT
GOOGLE DRIVE
GITHUB
UPLOAD
Filter notebooks
Title
First opened
Last opened
Welcome To Colaboratory
7 days ago
0 minutes ago
Untitled2.ipynb
1 minute ago
1 minute ago
Untitled1.ipynb
1 minute ago
1 minute ago
Untitled0.ipynb
2 minutes ago
2 minutes ago
Untitled
16 minutes ago
11 minutes ago
NEW PYTHON 3 NOTEBOOK
CANCEL

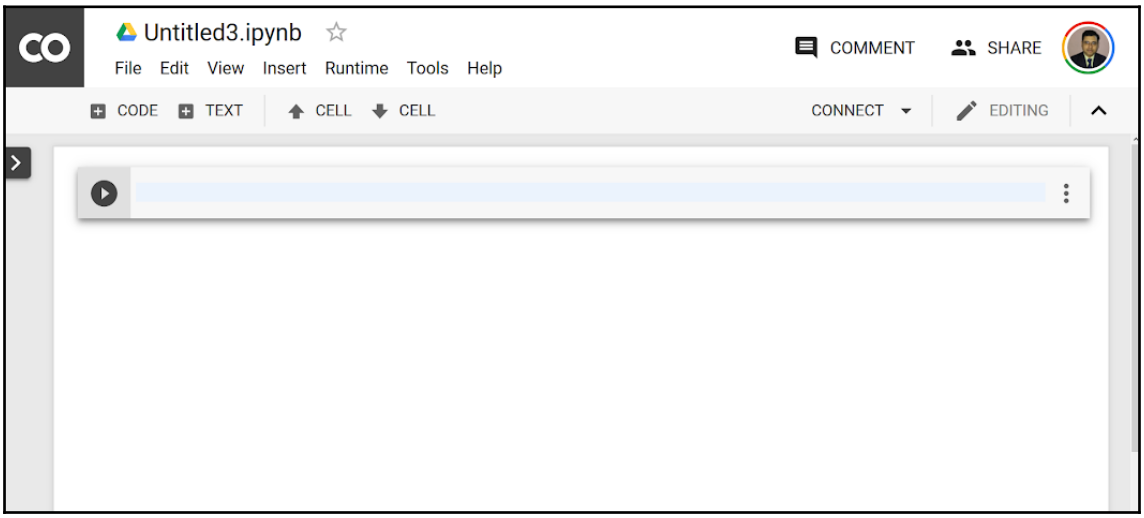
Untitled3.ipynb
File Edit View Insert Runtime Tools Help
COMMENT
SHARE
CODE
TEXT
CELL
CELL
CONNECT
EDITING

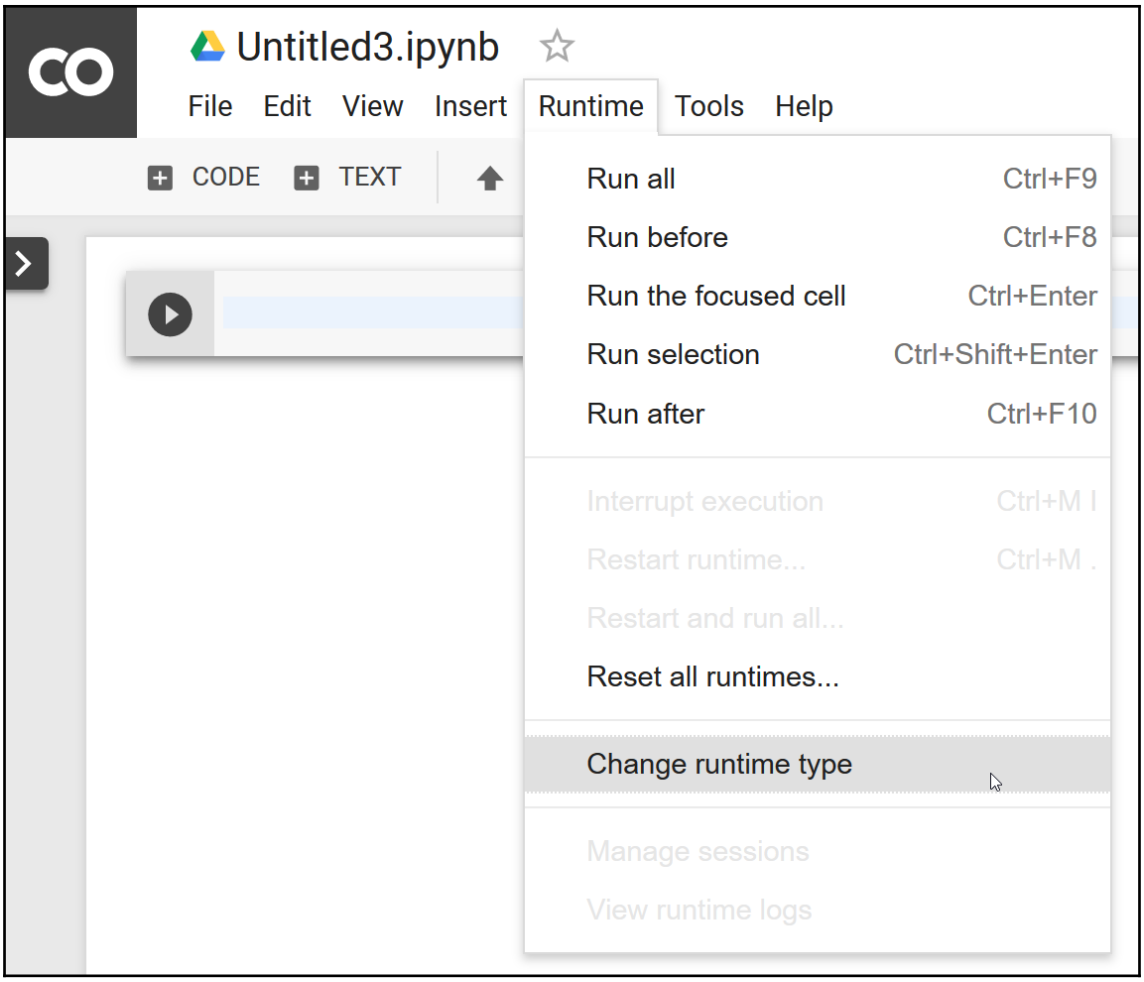
Untitled3.ipynb
File
Edit
View
Insert
Runtime
Tools
Help
CODE
TEXT
Run all
Ctrl+F9
Run before
Ctrl+F8
Run the focused cell
Ctrl+Enter
Run selection
Ctrl+Shift+Enter
Run after
Ctrl+F10
Interrupt execution
Ctrl+M I
Restart runtime...
Ctrl+M .
Restart and run all...
Reset all runtimes...
Change runtime type
Manage sessions
View runtime logs

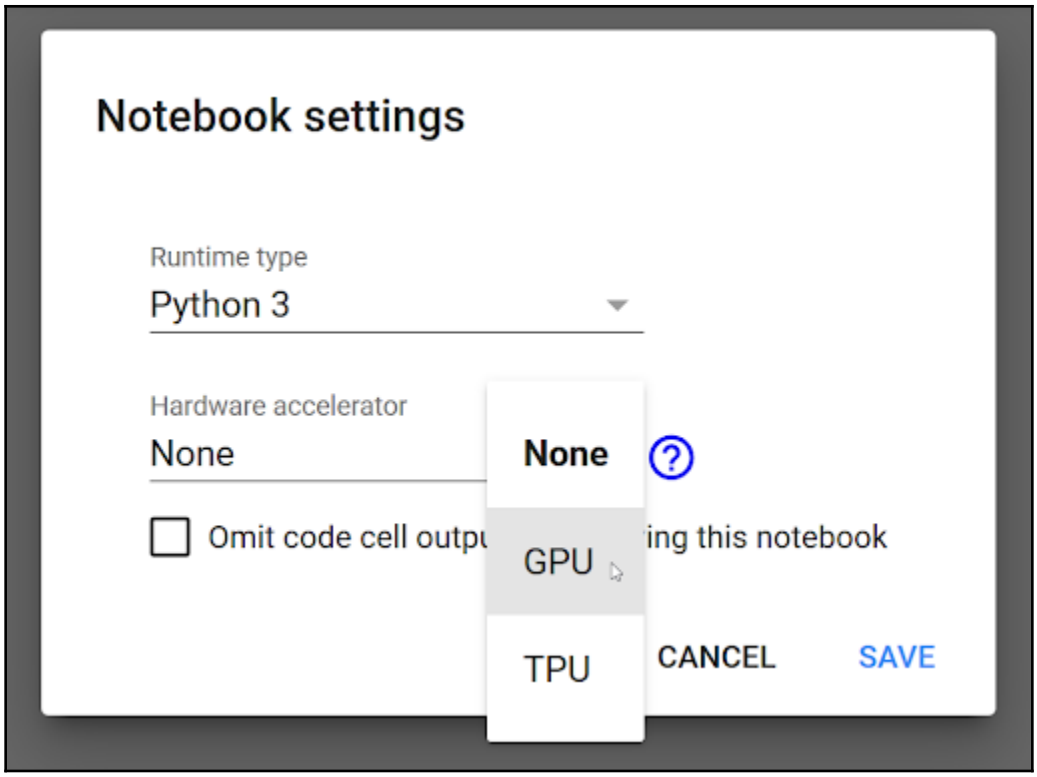


```
import numpy
```

```
import numpy
```

```
[4]  !ls

     sample_data
```

```
[10] !nvcc -V

nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2019 NVIDIA Corporation
Built on Fri_Feb__8_19:08:17_PST_2019
Cuda compilation tools, release 10.1, V10.1.105
```

```
!pip install pycuda

Collecting pycuda
  Downloading https://files.pythonhosted.org/packages/58/33/cced4891eddd1a3ac561ff99081019fddc7838a07cace272c941e3c2f915/pycuda-2018.1.1.tar.gz (1.6MB)
    100% |████████████████████████████████| 1.6MB 16.1MB/s
Collecting pytools>=2011.2 (from pycuda)
  Downloading https://files.pythonhosted.org/packages/ac/a3/f54f7190315ad41b7334d8733350e7fcefded8f25e0b45e2329b80279921/pytools-2019.1.tar.gz (57kB)
    100% |████████████████████████████████| 61kB 18.9MB/s
Requirement already satisfied: pytest>=2 in /usr/local/lib/python3.6/dist-packages (from pycuda) (3.6.4)
Requirement already satisfied: decorator>=3.2.0 in /usr/local/lib/python3.6/dist-packages (from pycuda) (4.4.0)
Collecting appdirs>=1.4.0 (from pycuda)
  Downloading https://files.pythonhosted.org/packages/56/eb/810e700ed1349edde4cbdc1b2a21e28cdf115f9faf263f6bbf8447c1abf3/appdirs-1.4.3-py2.py3-none-any.whl
Collecting mako (from pycuda)
  Downloading https://files.pythonhosted.org/packages/eb/69/6137c60cae2ab8c911bff510bb6d1d23a0189f75d114bb277606c6486b5f/Mako-1.0.8.tar.gz (468kB)
    100% |████████████████████████████████| 471kB 25.0MB/s
Requirement already satisfied: six>=1.8.0 in /usr/local/lib/python3.6/dist-packages (from pytools>=2011.2->pycuda) (1.11.0)
Requirement already satisfied: numpy>=1.6.0 in /usr/local/lib/python3.6/dist-packages (from pytools>=2011.2->pycuda) (1.14.6)
Requirement already satisfied: attrs>=17.4.0 in /usr/local/lib/python3.6/dist-packages (from pytest>=2->pycuda) (19.1.0)
Requirement already satisfied: atomicwrites>=1.0 in /usr/local/lib/python3.6/dist-packages (from pytest>=2->pycuda) (1.3.0)
Requirement already satisfied: py>=1.5.0 in /usr/local/lib/python3.6/dist-packages (from pytest>=2->pycuda) (1.8.0)
Requirement already satisfied: more-itertools>=4.0.0 in /usr/local/lib/python3.6/dist-packages (from pytest>=2->pycuda) (7.0.0)
Requirement already satisfied: pluggy<0.8,>=0.5 in /usr/local/lib/python3.6/dist-packages (from pytest>=2->pycuda) (0.7.1)
Requirement already satisfied: setuptools in /usr/local/lib/python3.6/dist-packages (from pytest>=2->pycuda) (40.9.0)
Requirement already satisfied: MarkupSafe>=0.9.2 in /usr/local/lib/python3.6/dist-packages (from mako->pycuda) (1.1.1)
Building wheels for collected packages: pycuda, pytools, mako
  Building wheel for pycuda (setup.py) ... done
  Stored in directory: /root/.cache/pip/wheels/a5/17/ac/99922221c732eeece43529d3e0f9d441f7301c75990b2cdbff
  Building wheel for pytools (setup.py) ... done
  Stored in directory: /root/.cache/pip/wheels/1b/f4/26/46b6cf949b3cccefcc41e6b526f7c16351a4a5c124fc6f6eaa
  Building wheel for mako (setup.py) ... done
  Stored in directory: /root/.cache/pip/wheels/b9/42/3c/79032e33bb2dfae60a10a6133f828a4a99ae46d33a2d323944
Successfully built pycuda pytools mako
Installing collected packages: appdirs, pytools, mako, pycuda
Successfully installed appdirs-1.4.3 mako-1.0.8 pycuda-2018.1.1 pytools-2019.1
```

Chapter 10: Accelerated Machine Learning on GPUs

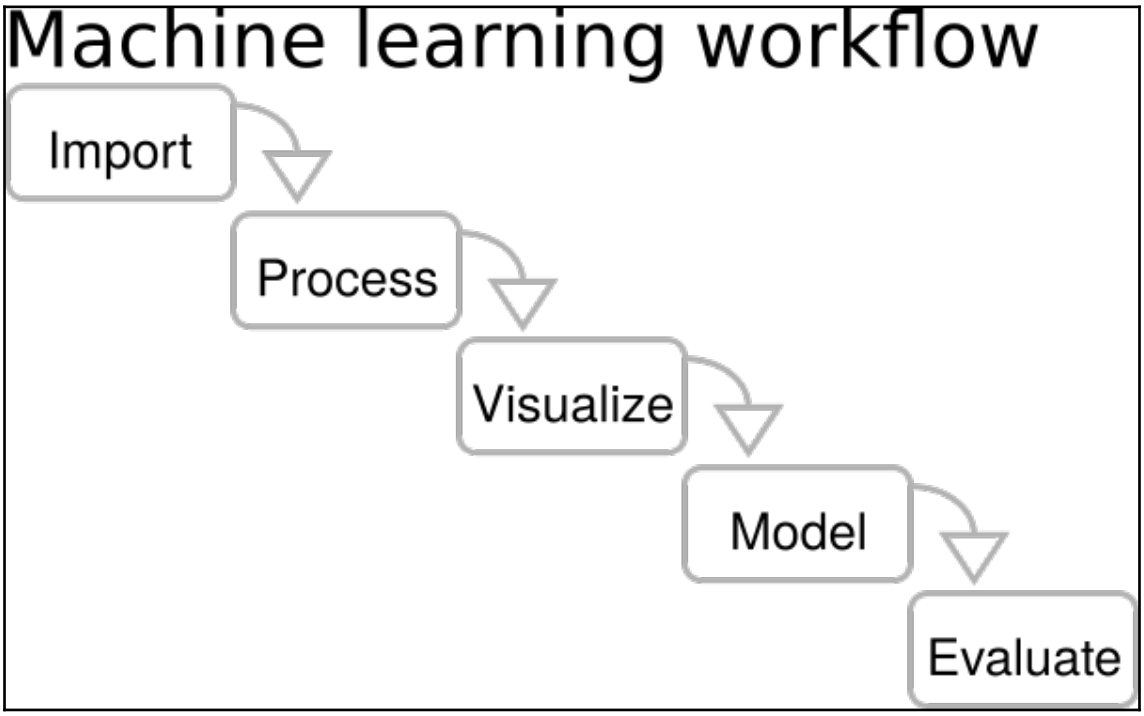

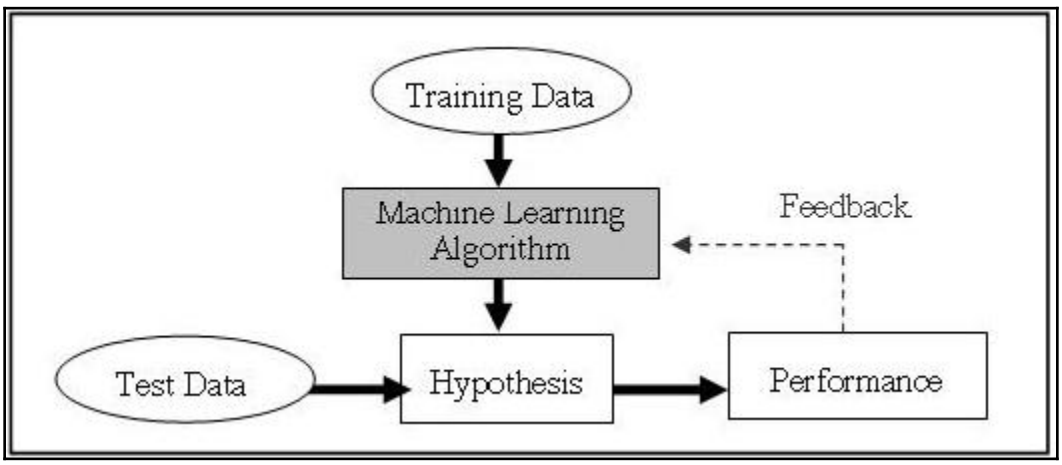

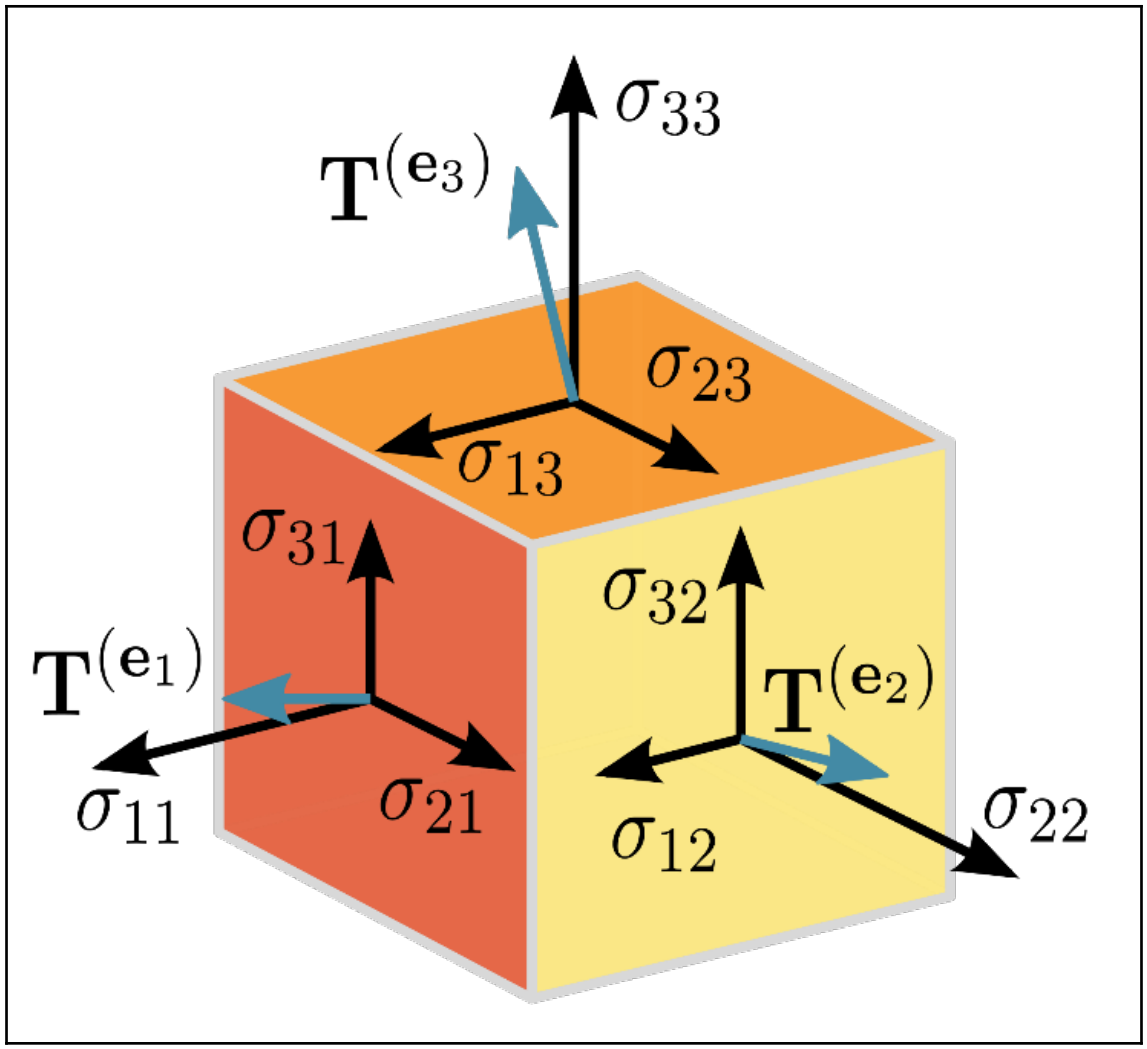
σ_{33}
$\mathbf{T}^{(\mathbf{e}_3)}$
σ_{23}
σ_{13}
σ_{31}
σ_{32}
$\mathbf{T}^{(\mathbf{e}_1)}$
$\mathbf{T}^{(\mathbf{e}_2)}$
σ_{11}
σ_{21}
σ_{12}
σ_{22}

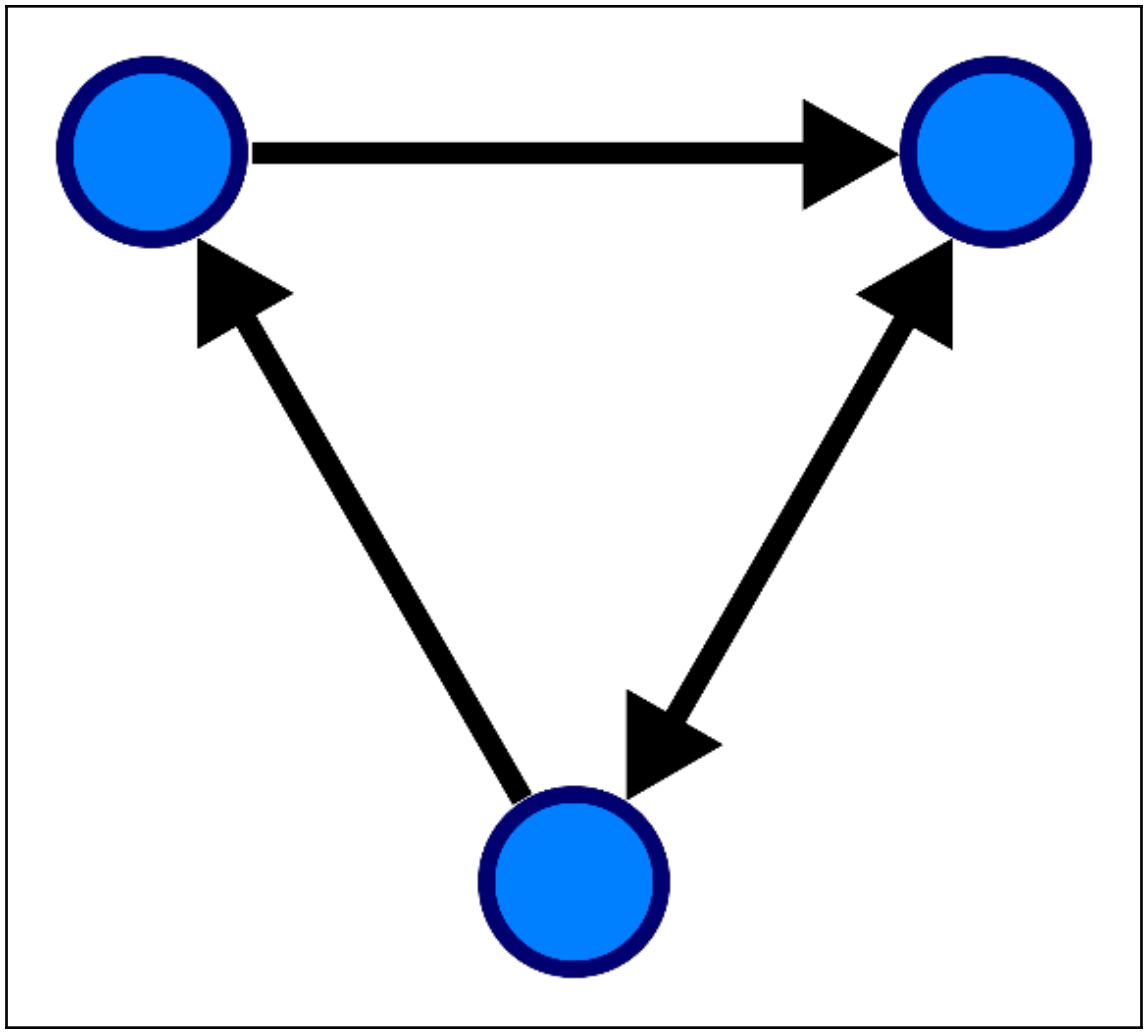

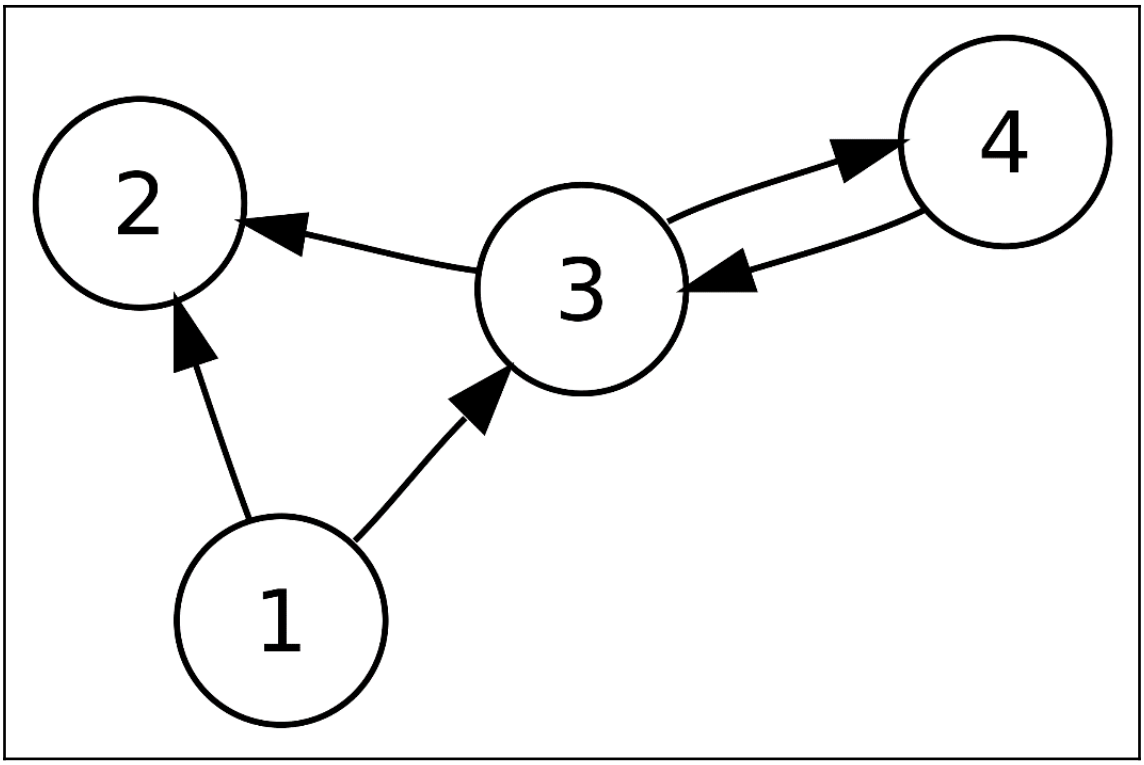
2
3
4
1

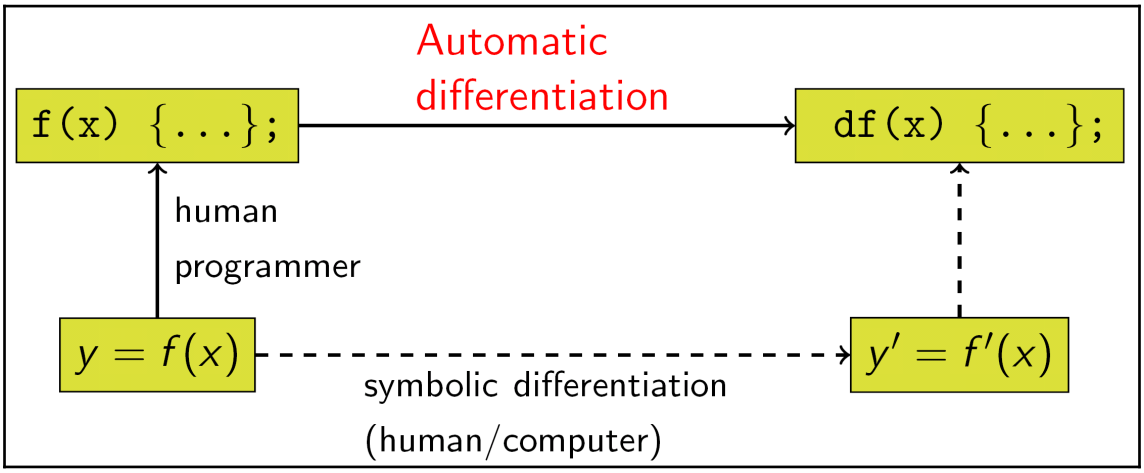
Automatic
differentiation
f(x) {...};
df(x) {...};
human
programmer
$y = f(x)$
$y' = f'(x)$
symbolic differentiation
(human/computer)

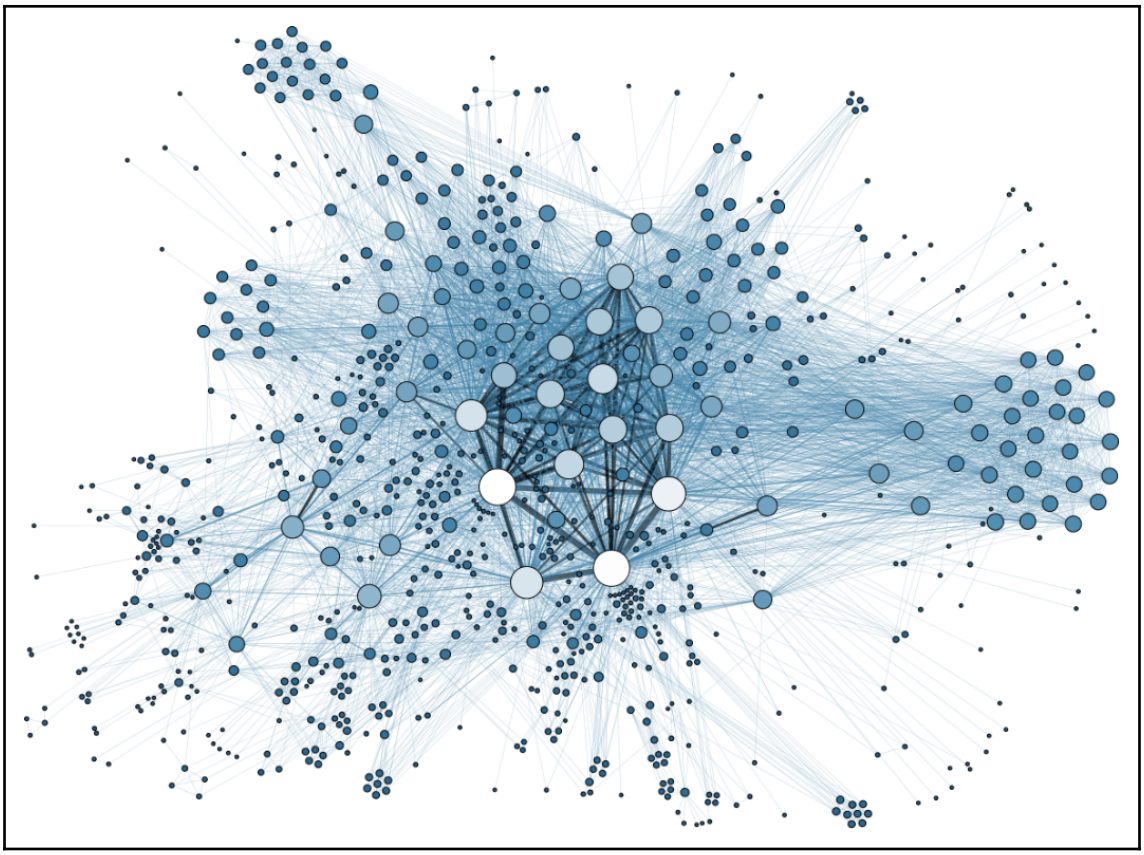

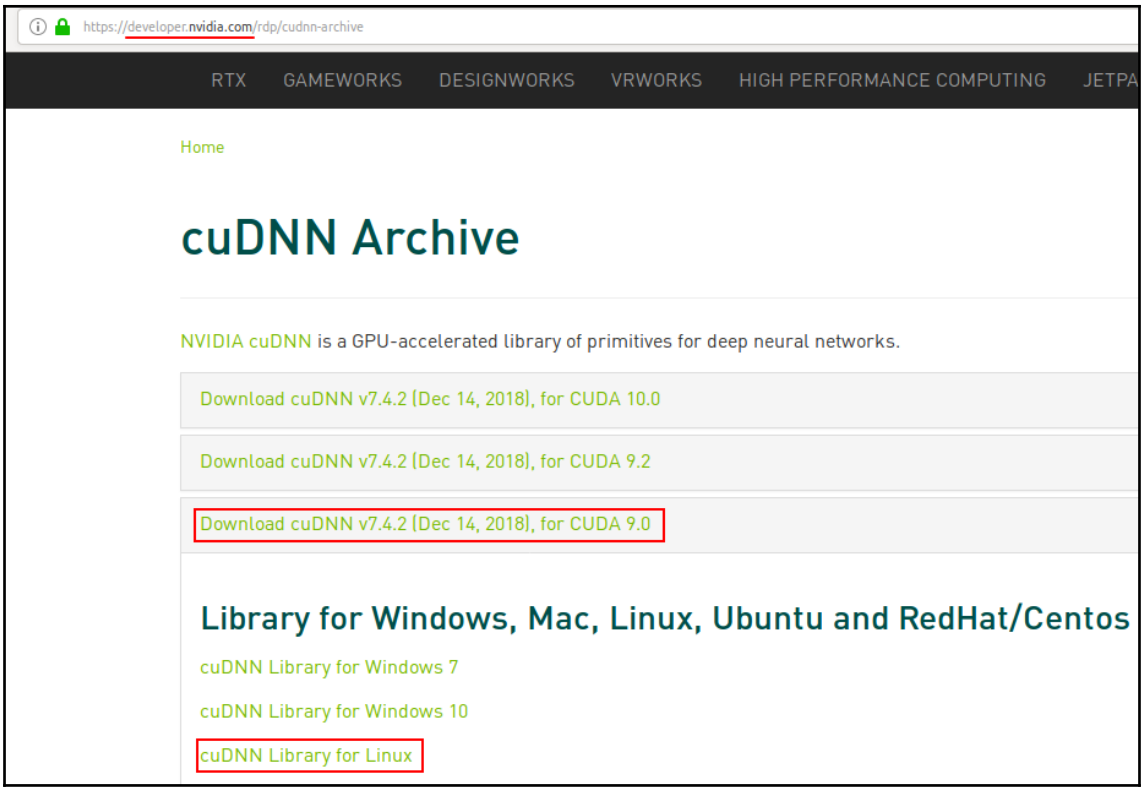
https://developer.nvidia.com/rdp/cudnn-archive
RTX
GAMEWORKS
DESIGNWORKS
VRWORKS
HIGH PERFORMANCE COMPUTING
JETPA
Home
cuDNN Archive
NVIDIA cuDNN is a GPU-accelerated library of primitives for deep neural networks.
Download cuDNN v7.4.2 (Dec 14, 2018), for CUDA 10.0
Download cuDNN v7.4.2 (Dec 14, 2018), for CUDA 9.2
Download cuDNN v7.4.2 (Dec 14, 2018), for CUDA 9.0
Library for Windows, Mac, Linux, Ubuntu and RedHat/Centos
cuDNN Library for Windows 7
cuDNN Library for Windows 10
cuDNN Library for Linux

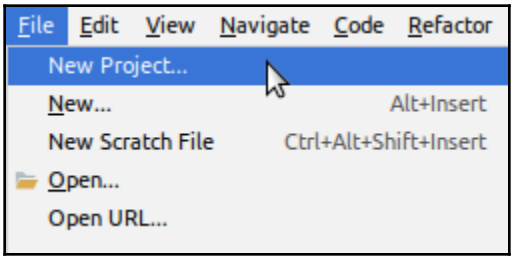
File
Edit
View
Navigate
Code
Refactor
New Project...
New...
Alt+Insert
New Scratch File
Ctrl+Alt+Shift+Insert
Open...
Open URL...

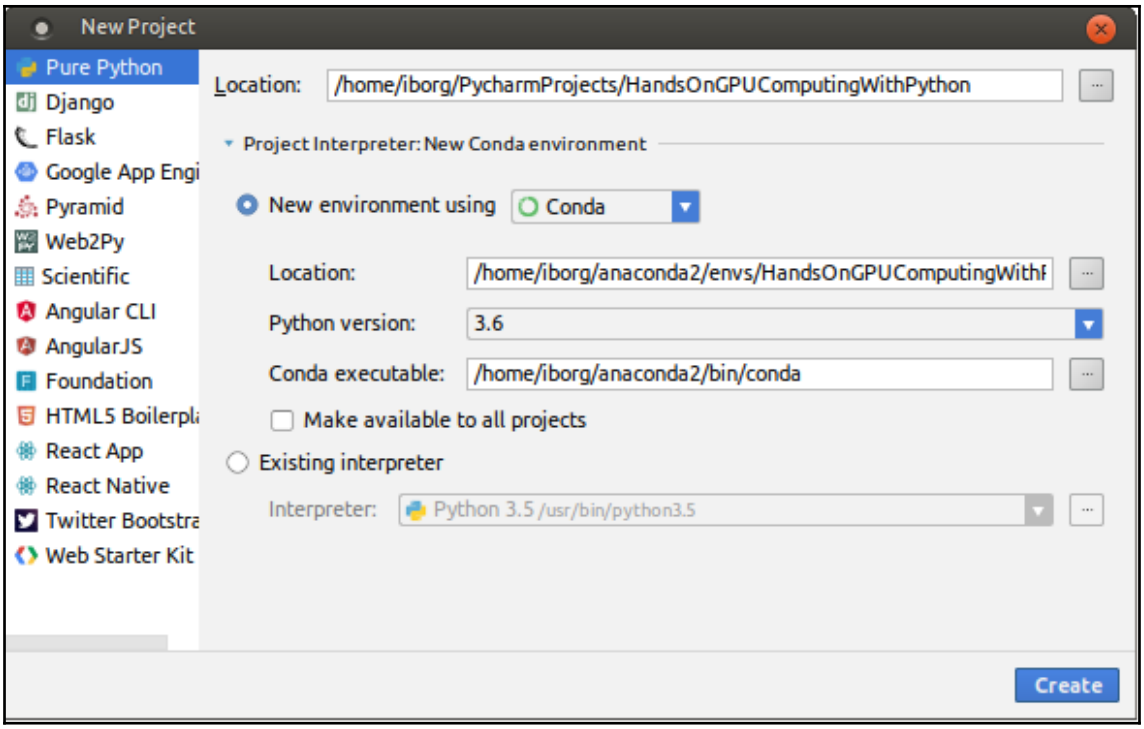
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location: /home/iborg/PycharmProjects/HandsOnGPUComputingWithPython
Project Interpreter: New Conda environment
New environment using Conda
Location: /home/iborg/anaconda2/envs/HandsOnGPUComputingWithF
Python version: 3.6
Conda executable: /home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter: Python 3.5 /usr/bin/python3.5
Create

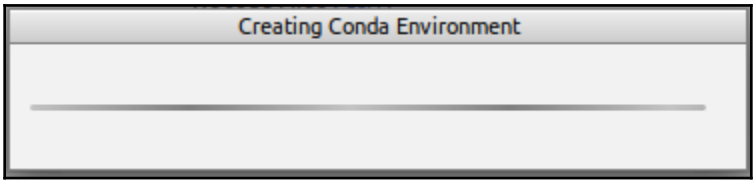
Creating Conda Environment

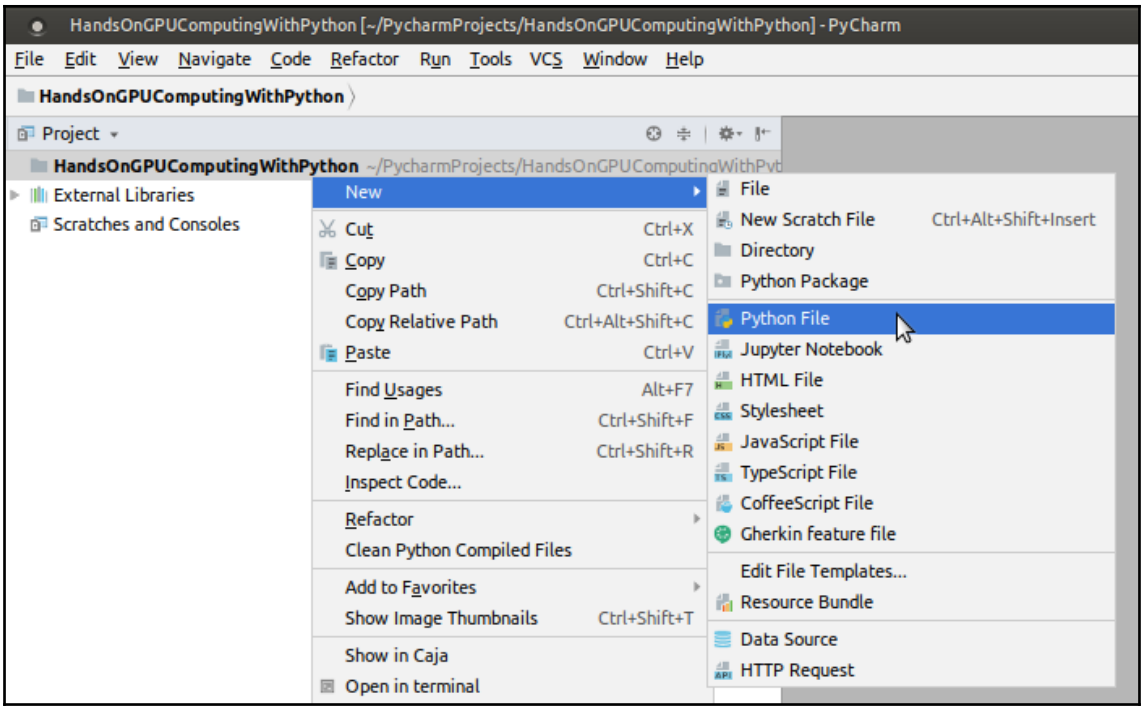
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
HandsOnGPUComputingWithPython ~/PycharmProjects/HandsOnGPUComputingWithPyt
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
Find in Path... Ctrl+Shift+F
Replace in Path... Ctrl+Shift+R
Inspect Code...
Refactor
Clean Python Compiled Files
Add to Favorites
Show Image Thumbnails Ctrl+Shift+T
Show in Caja
Open in terminal
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File
Stylesheet
JavaScript File
TypeScript File
CoffeeScript File
Gherkin feature file
Edit File Templates...
Resource Bundle
Data Source
HTTP Request


import tensorf
tensorflow

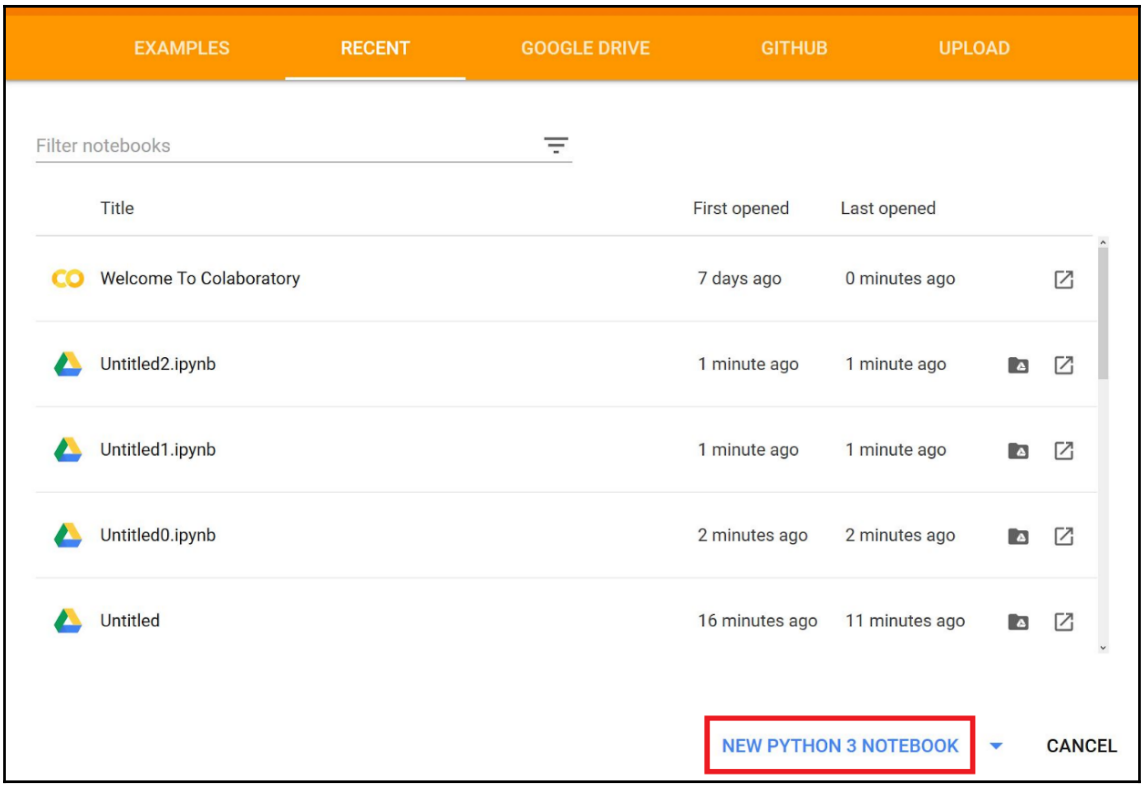
EXAMPLES
RECENT
GOOGLE DRIVE
GITHUB
UPLOAD
Filter notebooks
Title
First opened
Last opened
Welcome To Colaboratory
7 days ago
0 minutes ago
Untitled2.ipynb
1 minute ago
1 minute ago
Untitled1.ipynb
1 minute ago
1 minute ago
Untitled0.ipynb
2 minutes ago
2 minutes ago
Untitled
16 minutes ago
11 minutes ago
NEW PYTHON 3 NOTEBOOK
CANCEL

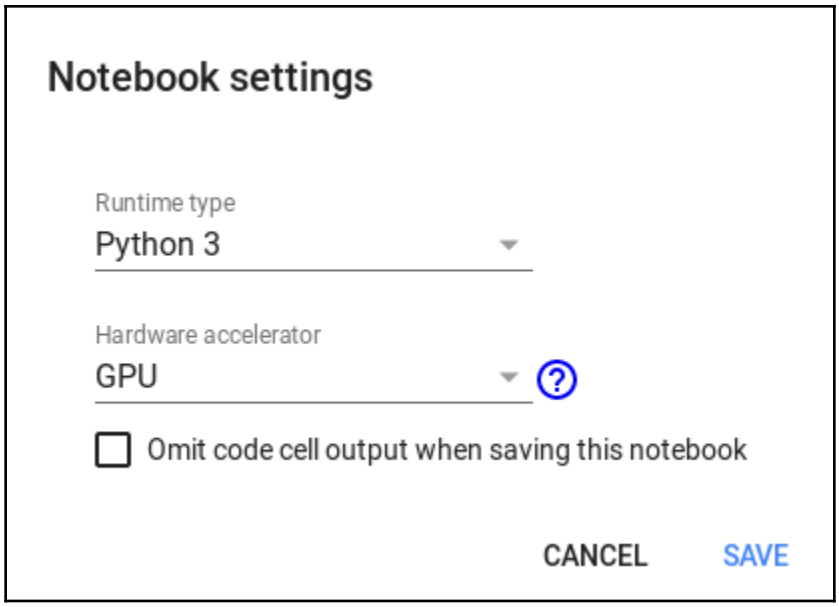
Notebook settings
Runtime type
Python 3
Hardware accelerator
GPU
Omit code cell output when saving this notebook
CANCEL
SAVE

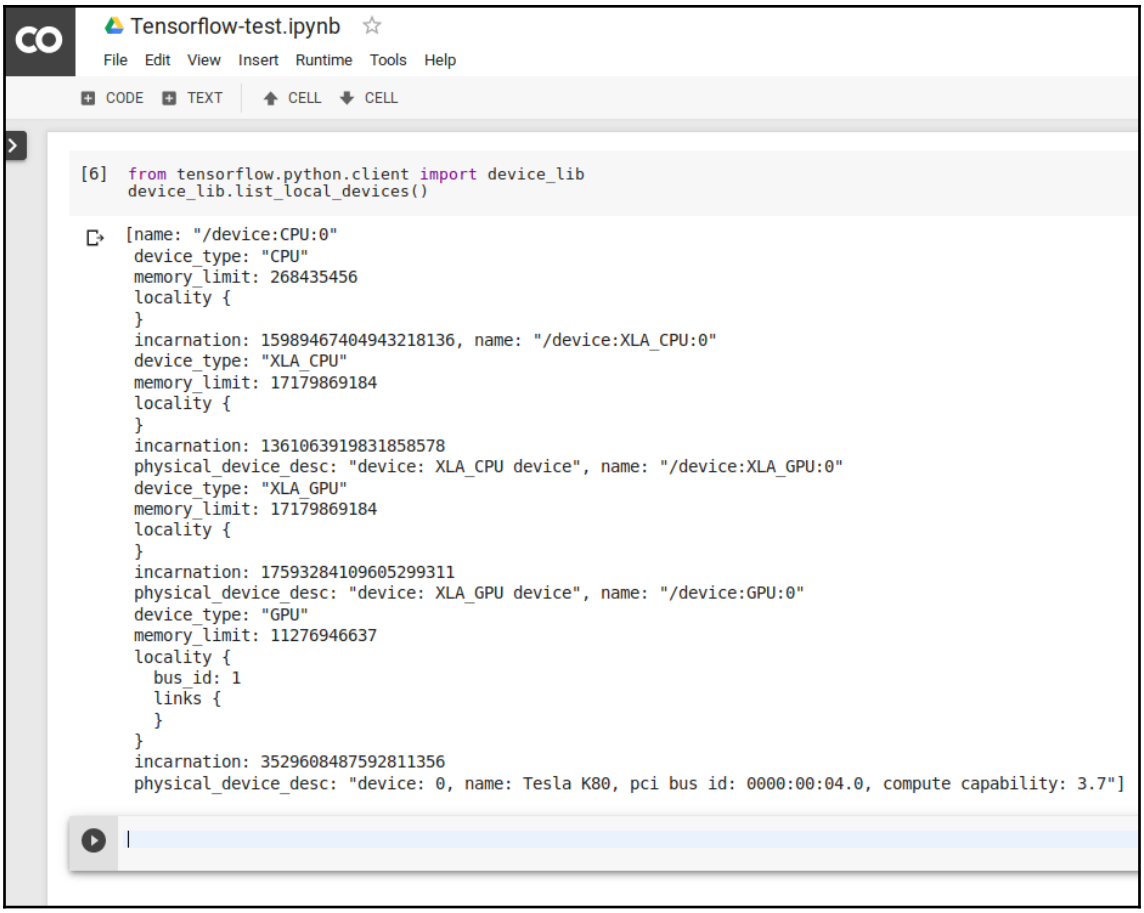
Tensorflow-test.ipynb
File Edit View Insert Runtime Tools Help
CODE TEXT CELL CELL
[6] from tensorflow.python.client import device_lib
device_lib.list_local_devices()
[name: "/device:CPU:0"
device_type: "CPU"
memory_limit: 268435456
locality {
}
incarnation: 15989467404943218136, name: "/device:XLA_CPU:0"
device_type: "XLA_CPU"
memory_limit: 17179869184
locality {
}
incarnation: 1361063919831858578
physical_device_desc: "device: XLA_CPU device", name: "/device:XLA_GPU:0"
device_type: "XLA_GPU"
memory_limit: 17179869184
locality {
}
incarnation: 17593284109605299311
physical_device_desc: "device: XLA_GPU device", name: "/device:GPU:0"
device_type: "GPU"
memory_limit: 11276946637
locality {
bus_id: 1
links {
}
}
incarnation: 3529608487592811356
physical_device_desc: "device: 0, name: Tesla K80, pci bus id: 0000:00:04.0, compute capability: 3.7"]

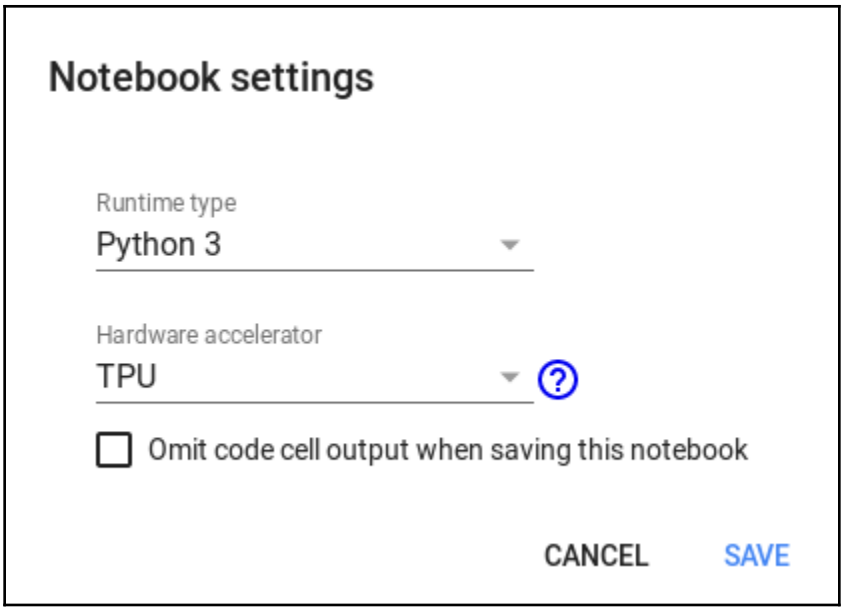


```
import tensorflow as tf
import os

tpu_path = 'grpc://' + os.environ['COLAB_TPU_ADDR']
print ('TPU Address:', tpu_path)
print ('TPU Devices:')
tf.Session(tpu_path).list_devices()
```

```
TPU Address: grpc://10.28.7.170:8470
TPU Devices:
[_DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:CPU:0, CPU, -1, 4777662762040939562),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:XLA_CPU:0, XLA_CPU, 17179869184, 18125897417461199249),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:0, TPU, 17179869184, 11197279434657959742),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:1, TPU, 17179869184, 2363331106931353136),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:2, TPU, 17179869184, 9315033542070624465),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:3, TPU, 17179869184, 13775977546227054969),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:4, TPU, 17179869184, 1688936288431572532),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:5, TPU, 17179869184, 2505726060068495099),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:6, TPU, 17179869184, 3340525144834310058),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU:7, TPU, 17179869184, 18112389880637035600),
 _DeviceAttributes(/job:tpu_worker/replica:0/task:0/device:TPU_SYSTEM:0, TPU_SYSTEM, 17179869184, 1563169429893506137)]
```

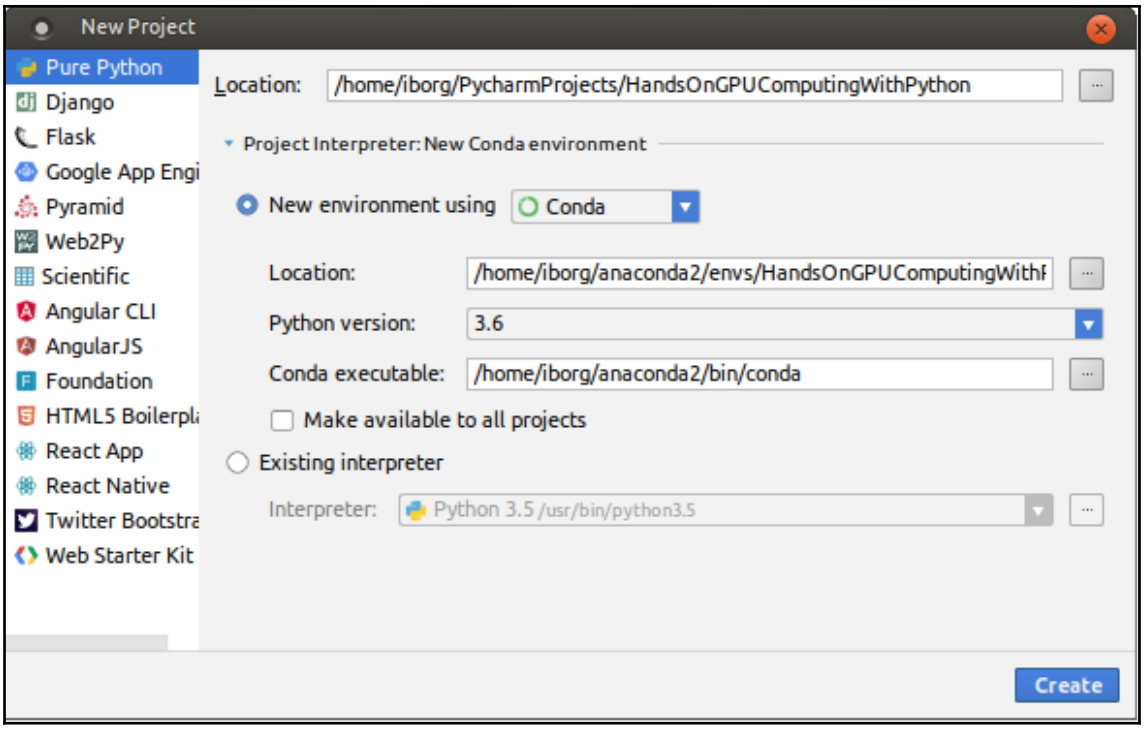
New Project
Pure Python
Django
Flask
Google App Engi
Pyramid
Web2Py
Scientific
Angular CLI
AngularJS
Foundation
HTML5 Boilerpl
React App
React Native
Twitter Bootstra
Web Starter Kit
Location: /home/iborg/PycharmProjects/HandsOnGPUComputingWithPython
Project Interpreter: New Conda environment
New environment using Conda
Location: /home/iborg/anaconda2/envs/HandsOnGPUComputingWithF
Python version: 3.6
Conda executable: /home/iborg/anaconda2/bin/conda
Make available to all projects
Existing interpreter
Interpreter: Python 3.5 /usr/bin/python3.5
Create

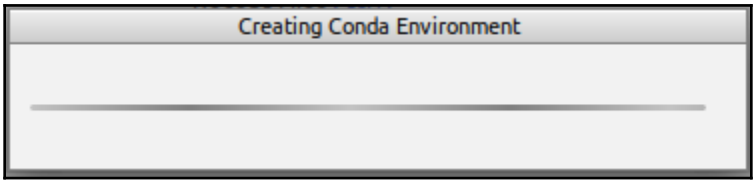
Creating Conda Environment

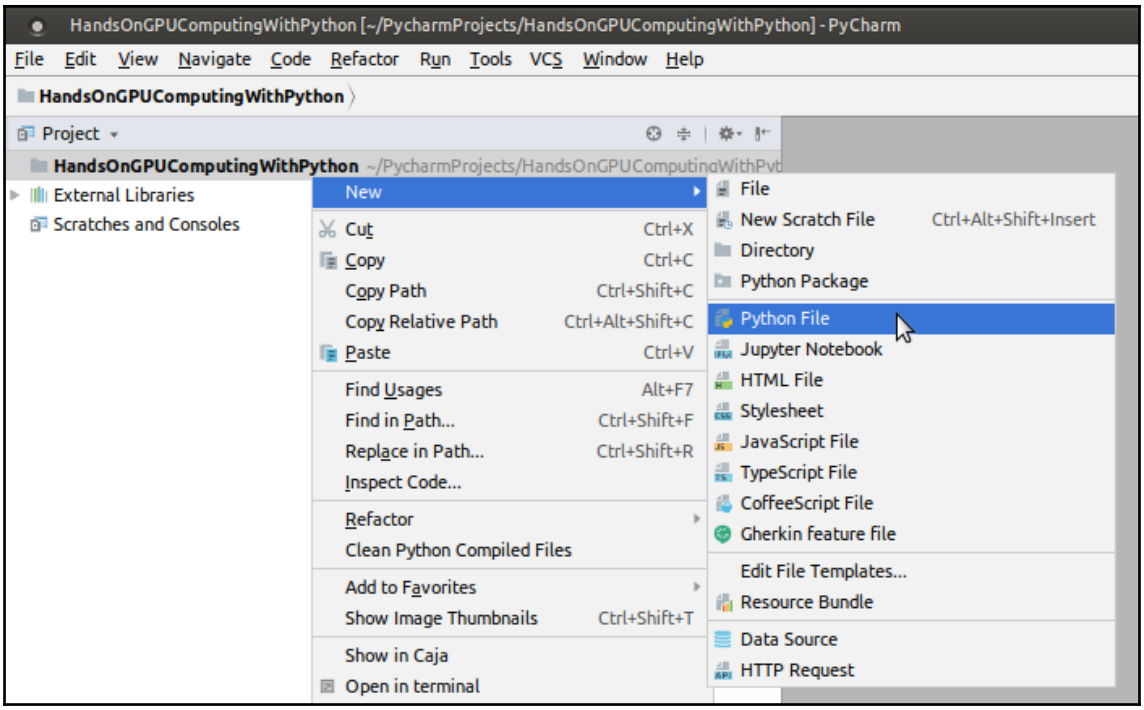


```
0
1
GeForce GTX TITAN X
(5, 2)

Process finished with exit code 0
```

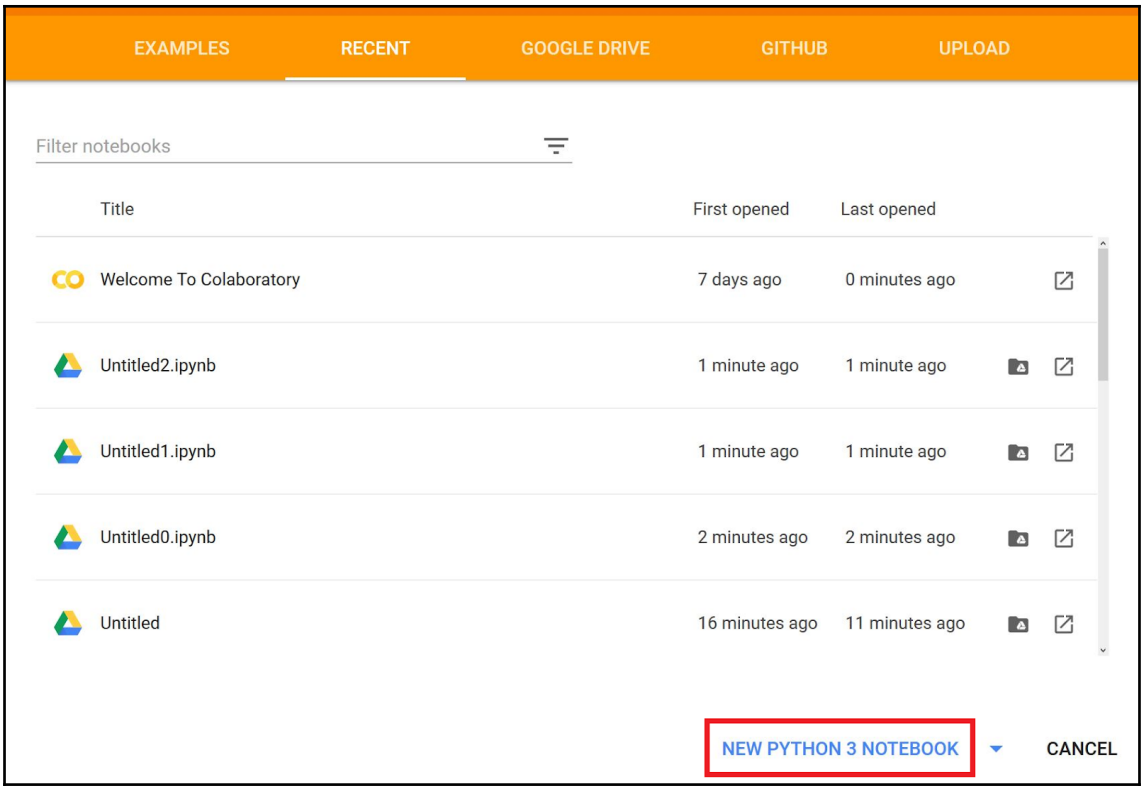
EXAMPLES
RECENT
GOOGLE DRIVE
GITHUB
UPLOAD
Filter notebooks
Title
First opened
Last opened
Welcome To Colaboratory
7 days ago
0 minutes ago
Untitled2.ipynb
1 minute ago
1 minute ago
Untitled1.ipynb
1 minute ago
1 minute ago
Untitled0.ipynb
2 minutes ago
2 minutes ago
Untitled
16 minutes ago
11 minutes ago
NEW PYTHON 3 NOTEBOOK
CANCEL

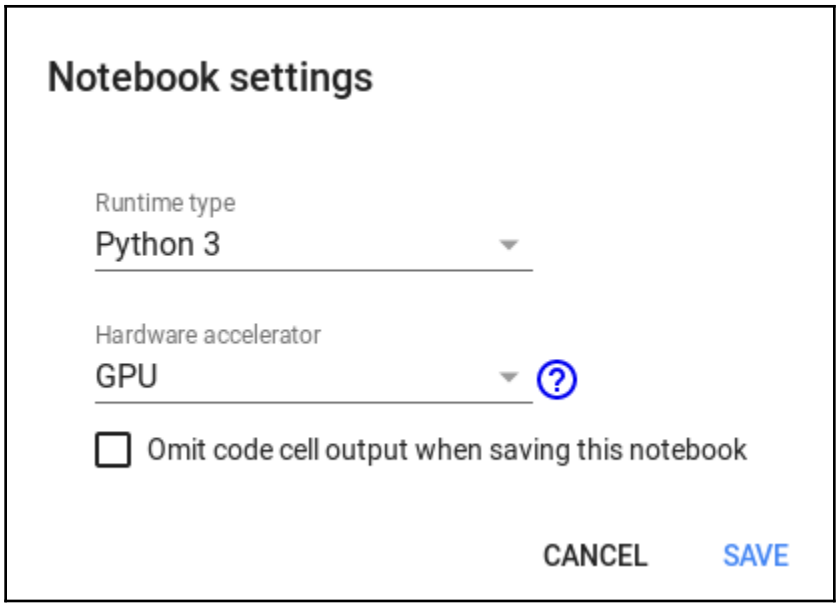
Notebook settings
Runtime type
Python 3
Hardware accelerator
GPU
Omit code cell output when saving this notebook
CANCEL
SAVE

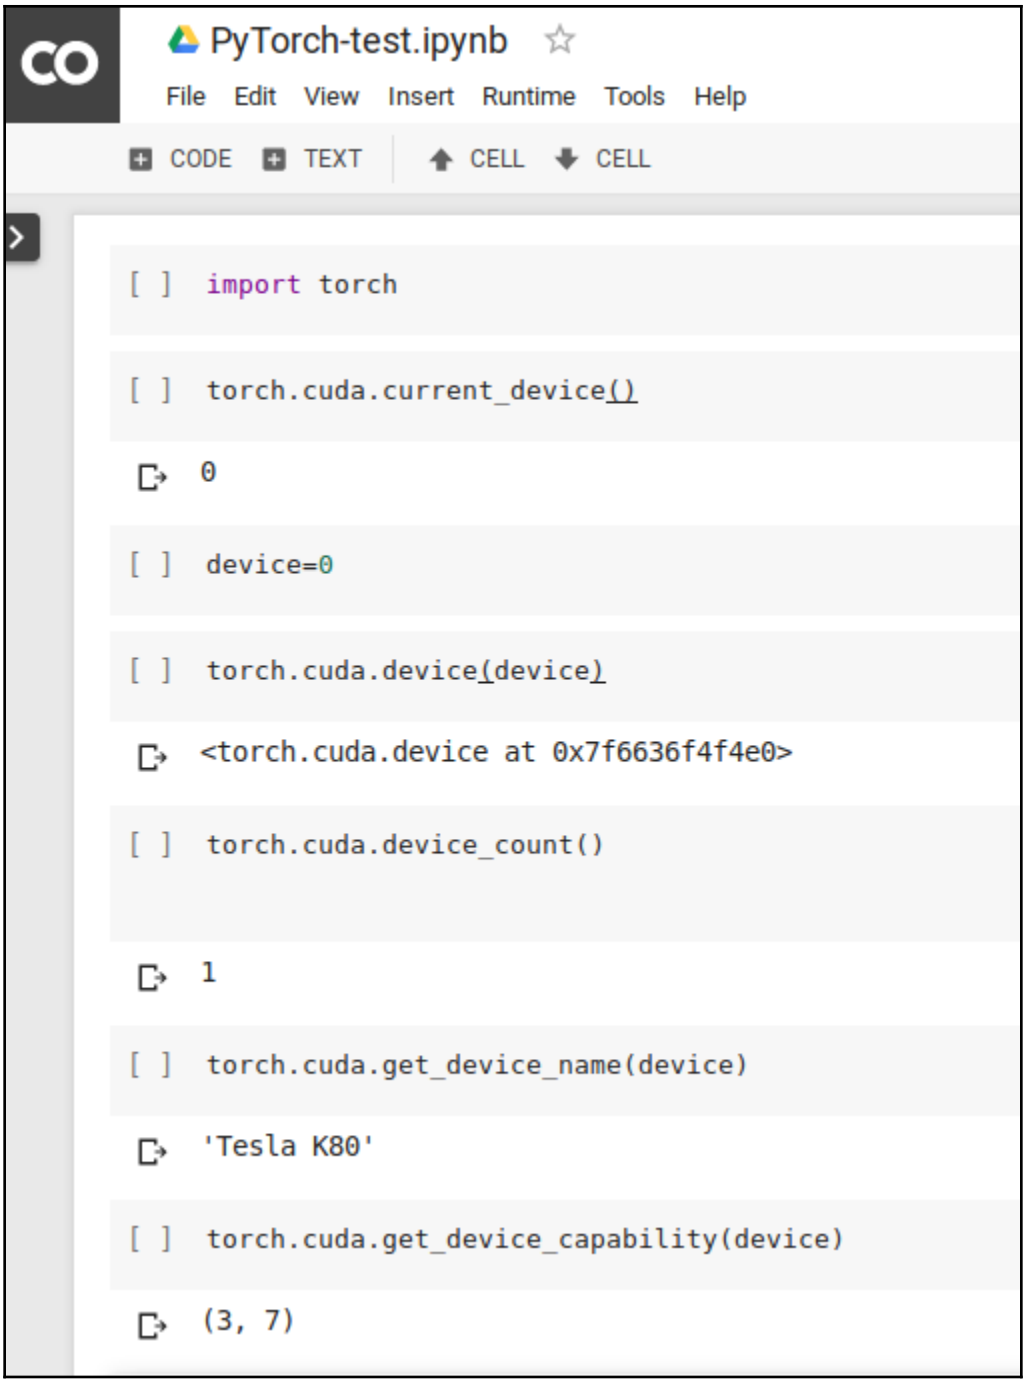
PyTorch-test.ipynb
File Edit View Insert Runtime Tools Help
CODE TEXT CELL CELL
[ ] import torch
[ ] torch.cuda.current_device()
0
[ ] device=0
[ ] torch.cuda.device(device)
<torch.cuda.device at 0x7f6636f4f4e0>
[ ] torch.cuda.device_count()
1
[ ] torch.cuda.get_device_name(device)
'Tesla K80'
[ ] torch.cuda.get_device_capability(device)
(3, 7)

```
import os

tpu_path = 'grpc://' + os.environ['COLAB_TPU_ADDR']
print ('TPU Address:', tpu_path)
```

```
TPU Address: grpc://10.28.7.170:8470
```

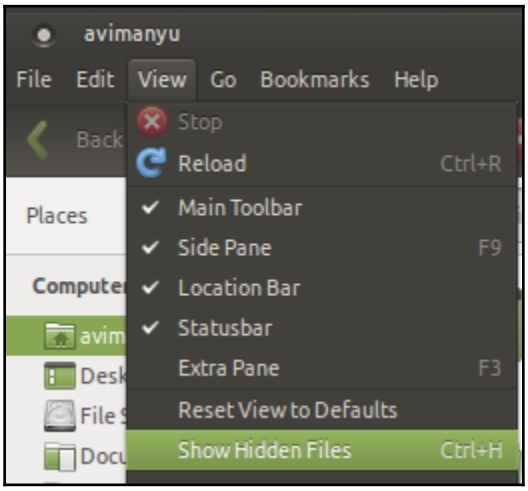

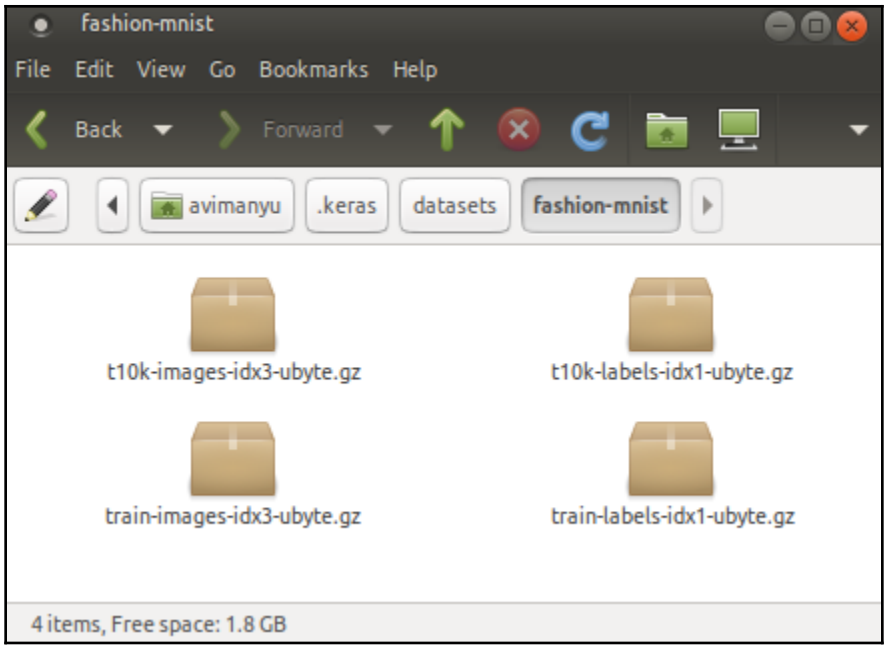
fashion-mnist
File Edit View Go Bookmarks Help
Back Forward
avimanyu .keras datasets fashion-mnist
t10k-images-idx3-ubyte.gz
t10k-labels-idx1-ubyte.gz
train-images-idx3-ubyte.gz
train-labels-idx1-ubyte.gz
4 items, Free space: 1.8 GB

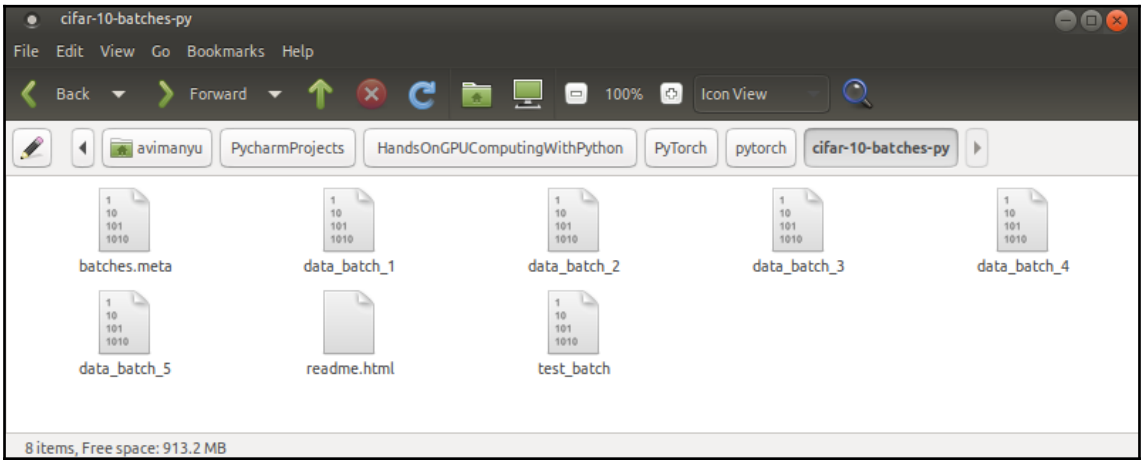
cifar-10-batches-py
File Edit View Go Bookmarks Help
Back Forward 100% Icon View
avimanyu PycharmProjects HandsOnGPUComputingWithPython PyTorch pytorch cifar-10-batches-py
batches.meta
data_batch_1
data_batch_2
data_batch_3
data_batch_4
data_batch_5
readme.html
test_batch
8 items, Free space: 913.2 MB

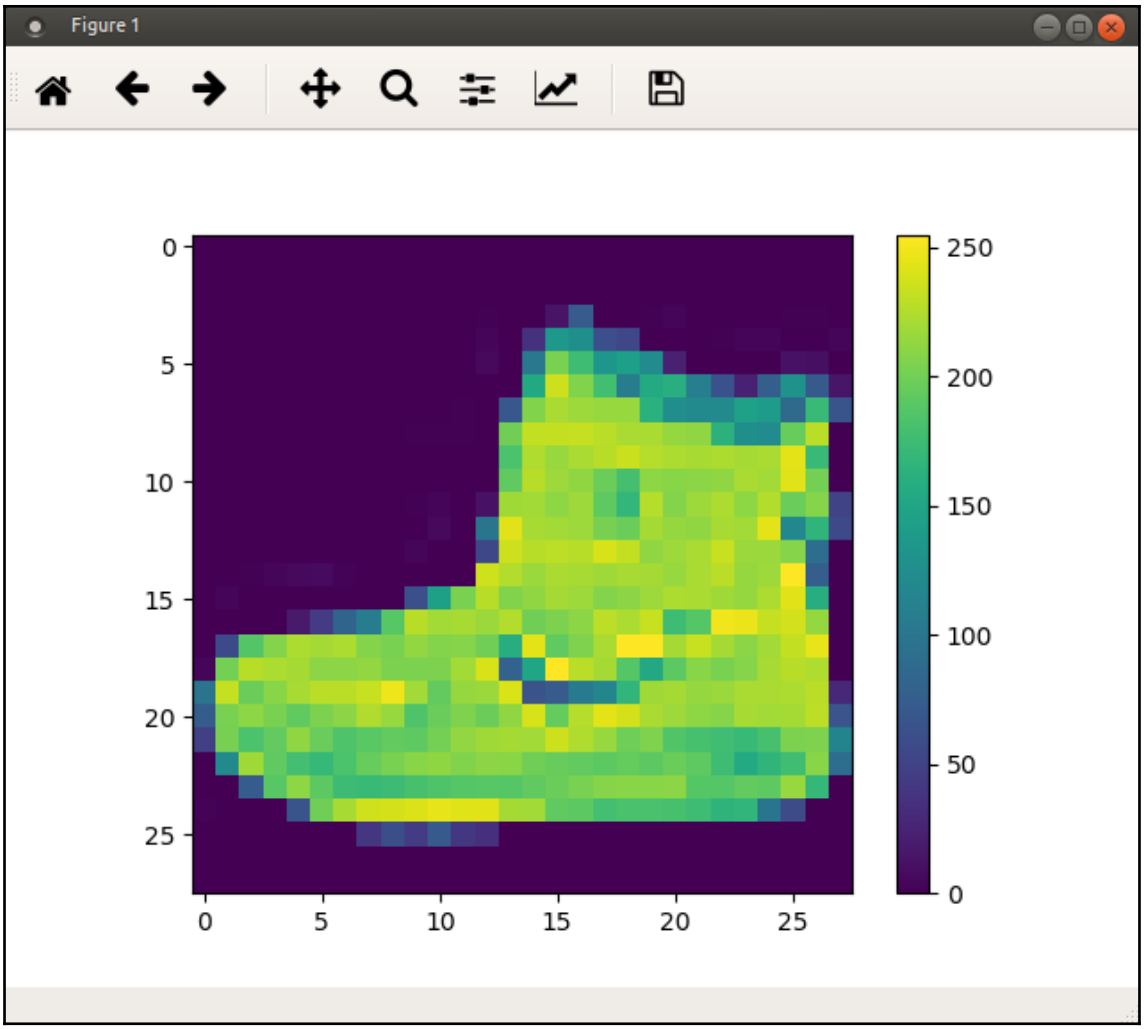
Figure 1
0
5
10
15
20
25
0
5
10
15
20
25
250
200
150
100
50
0

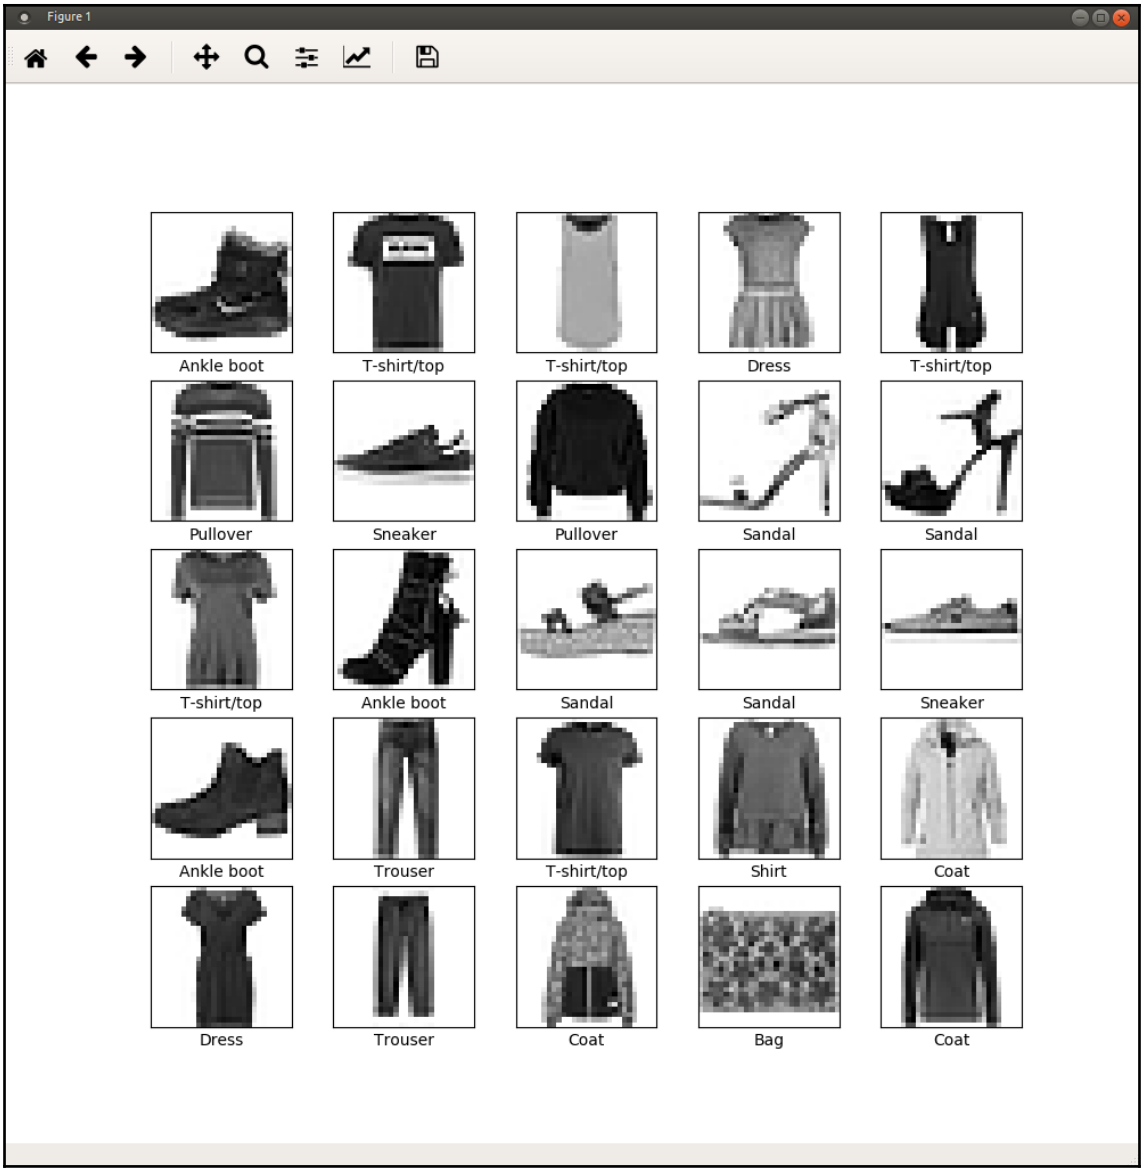
Figure 1
Ankle boot
T-shirt/top
T-shirt/top
Dress
T-shirt/top
Pullover
Sneaker
Pullover
Sandal
Sandal
T-shirt/top
Ankle boot
Sandal
Sandal
Sneaker
Ankle boot
Trouser
T-shirt/top
Shirt
Coat
Dress
Trouser
Coat
Bag
Coat

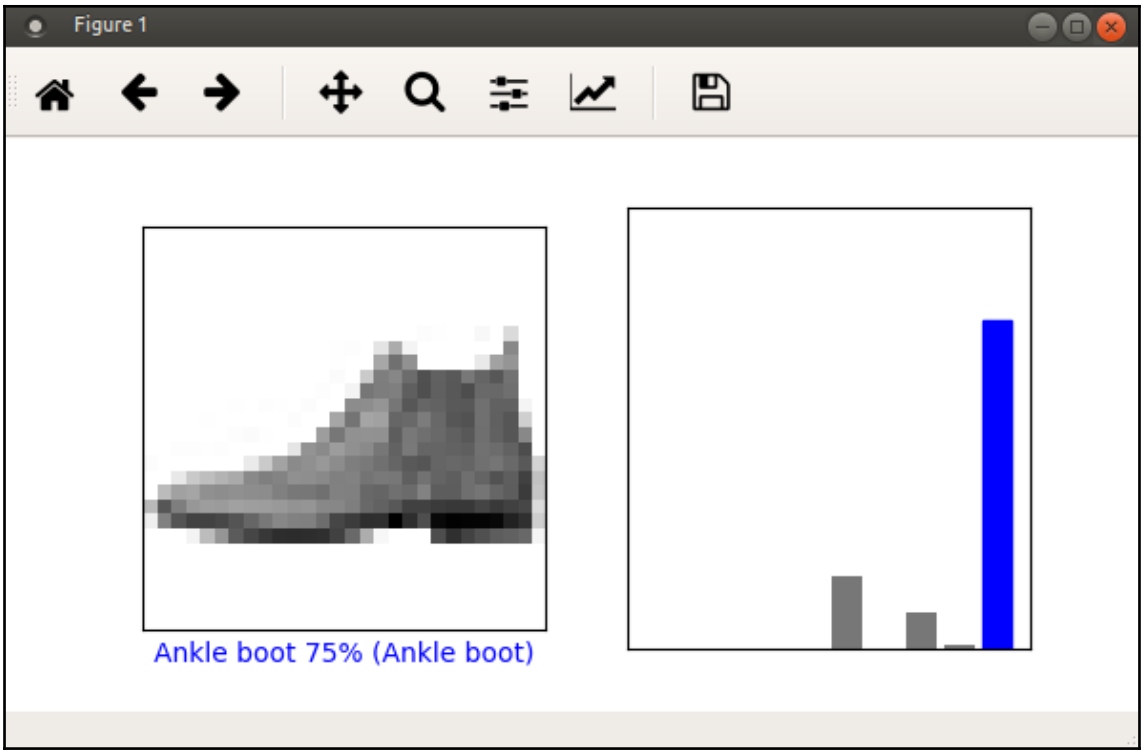
Figure 1
Ankle boot 75% (Ankle boot)

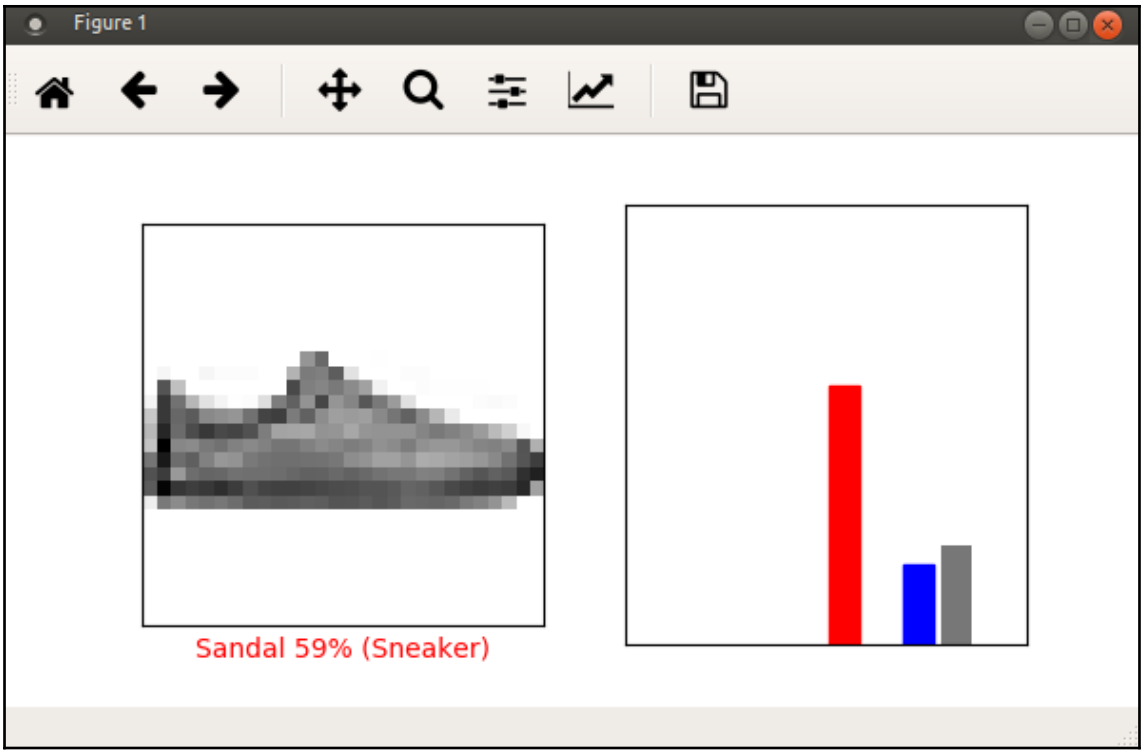
Figure 1
Sandal 59% (Sneaker)

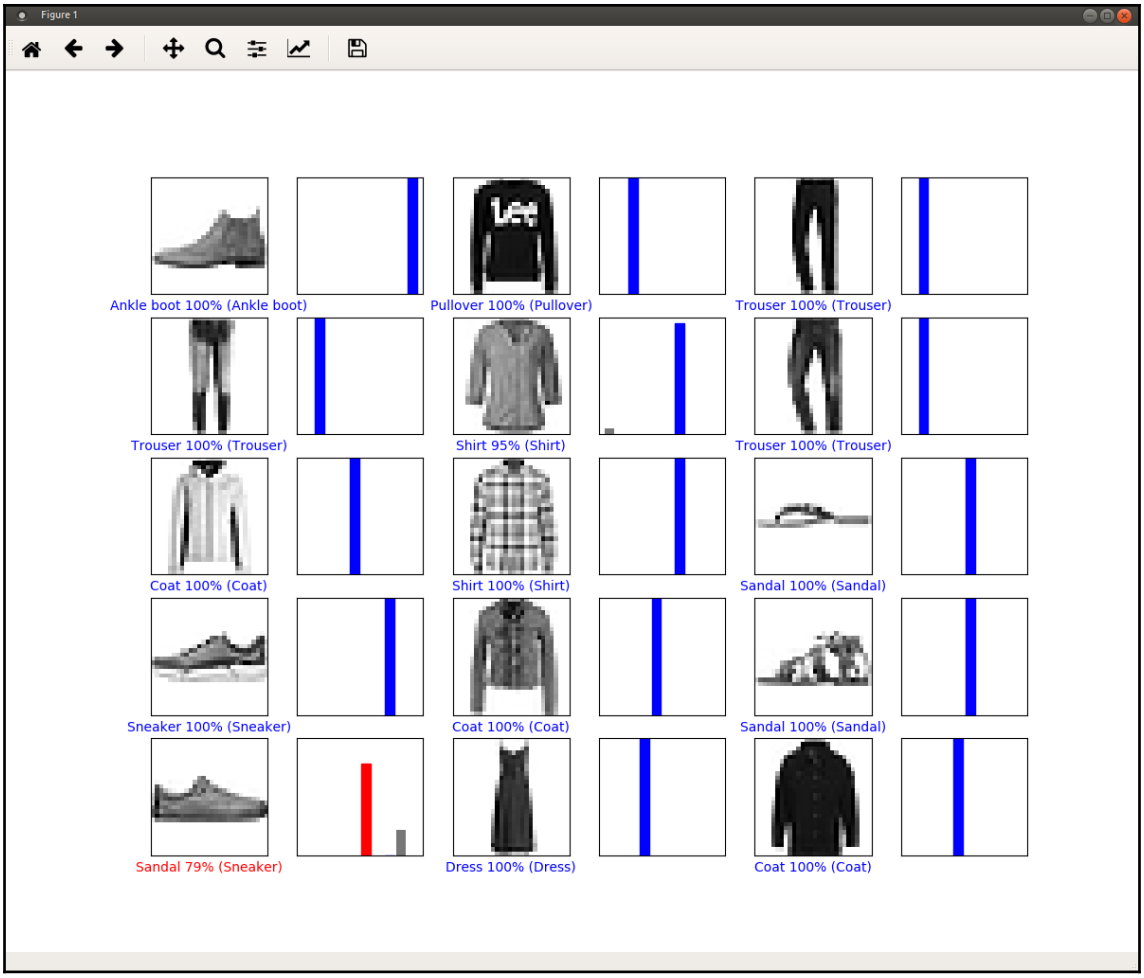
Figure 1
Ankle boot 100% (Ankle boot)
Pullover 100% (Pullover)
Trouser 100% (Trouser)
Trouser 100% (Trouser)
Shirt 95% (Shirt)
Trouser 100% (Trouser)
Coat 100% (Coat)
Shirt 100% (Shirt)
Sandal 100% (Sandal)
Sneaker 100% (Sneaker)
Coat 100% (Coat)
Sandal 100% (Sandal)
Sandal 79% (Sneaker)
Dress 100% (Dress)
Coat 100% (Coat)

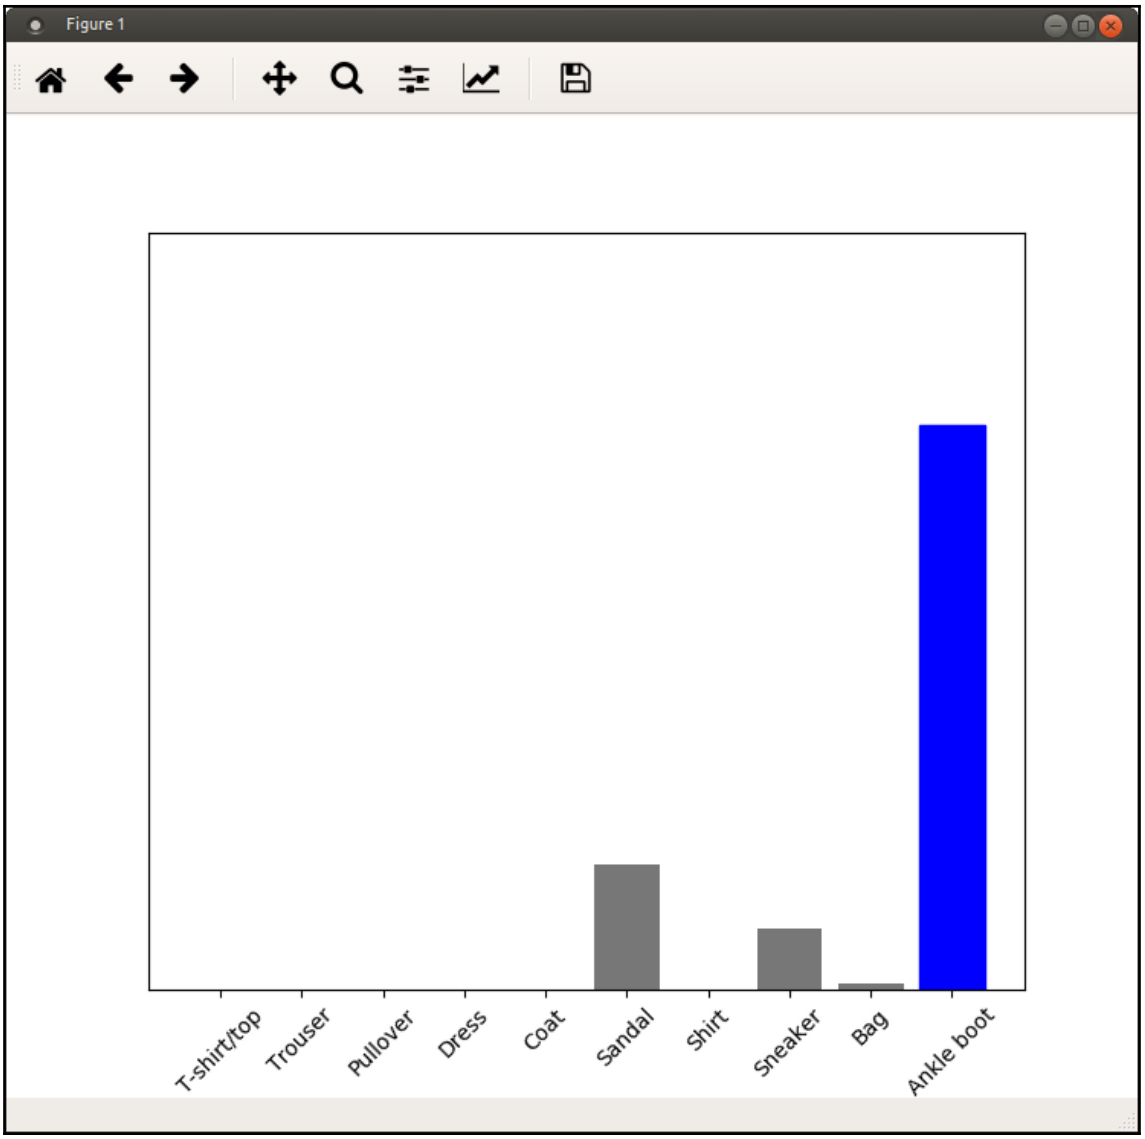
Figure 1
T-shirt/top
Trouser
Pullover
Dress
Coat
Sandal
Shirt
Sneaker
Bag
Ankle boot

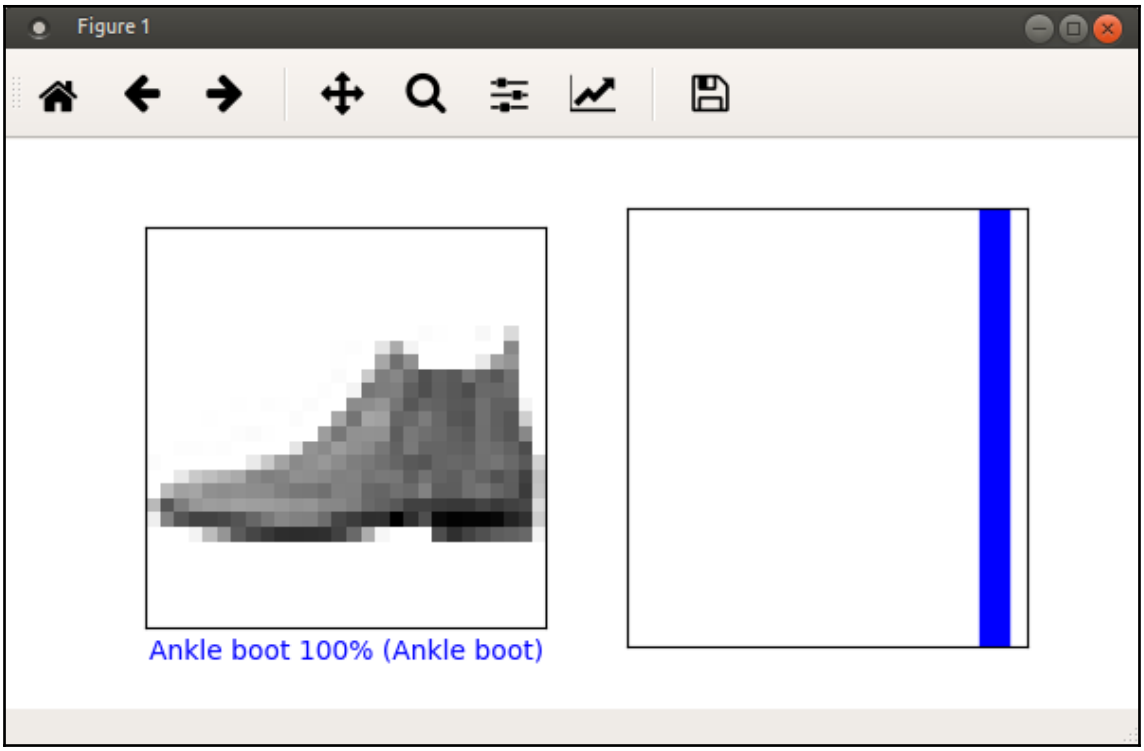
Figure 1
Ankle boot 100% (Ankle boot)

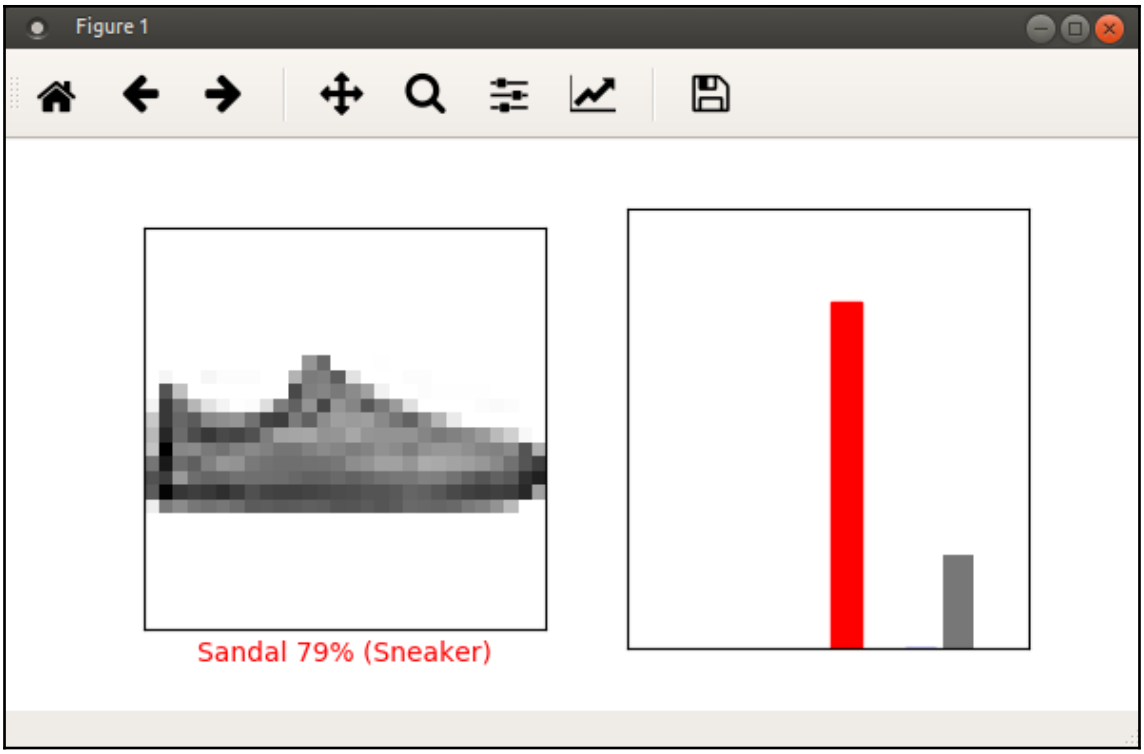
Figure 1
Sandal 79% (Sneaker)

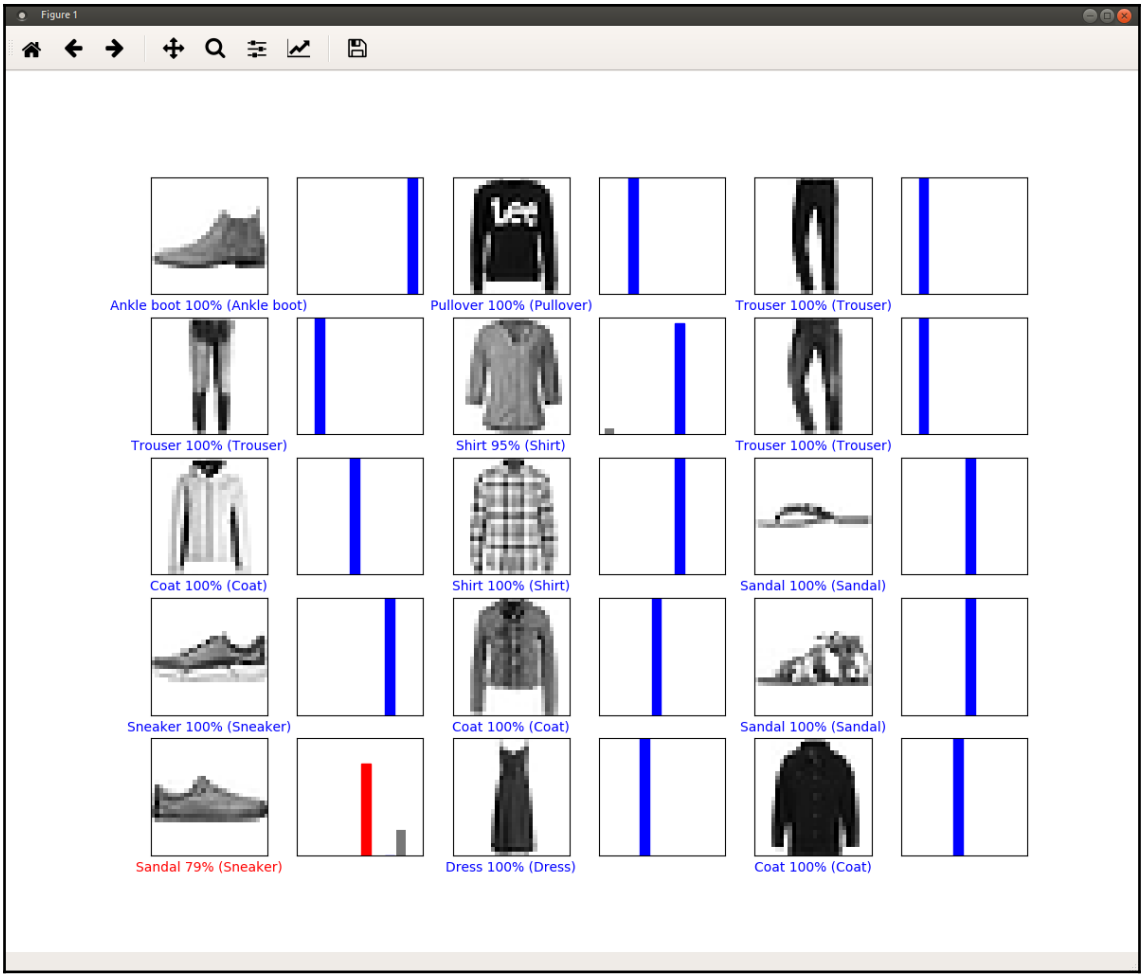
Figure 1
Ankle boot 100% (Ankle boot)
Pullover 100% (Pullover)
Trouser 100% (Trouser)
Trouser 100% (Trouser)
Shirt 95% (Shirt)
Trouser 100% (Trouser)
Coat 100% (Coat)
Shirt 100% (Shirt)
Sandal 100% (Sandal)
Sneaker 100% (Sneaker)
Coat 100% (Coat)
Sandal 100% (Sandal)
Sandal 79% (Sneaker)
Dress 100% (Dress)
Coat 100% (Coat)

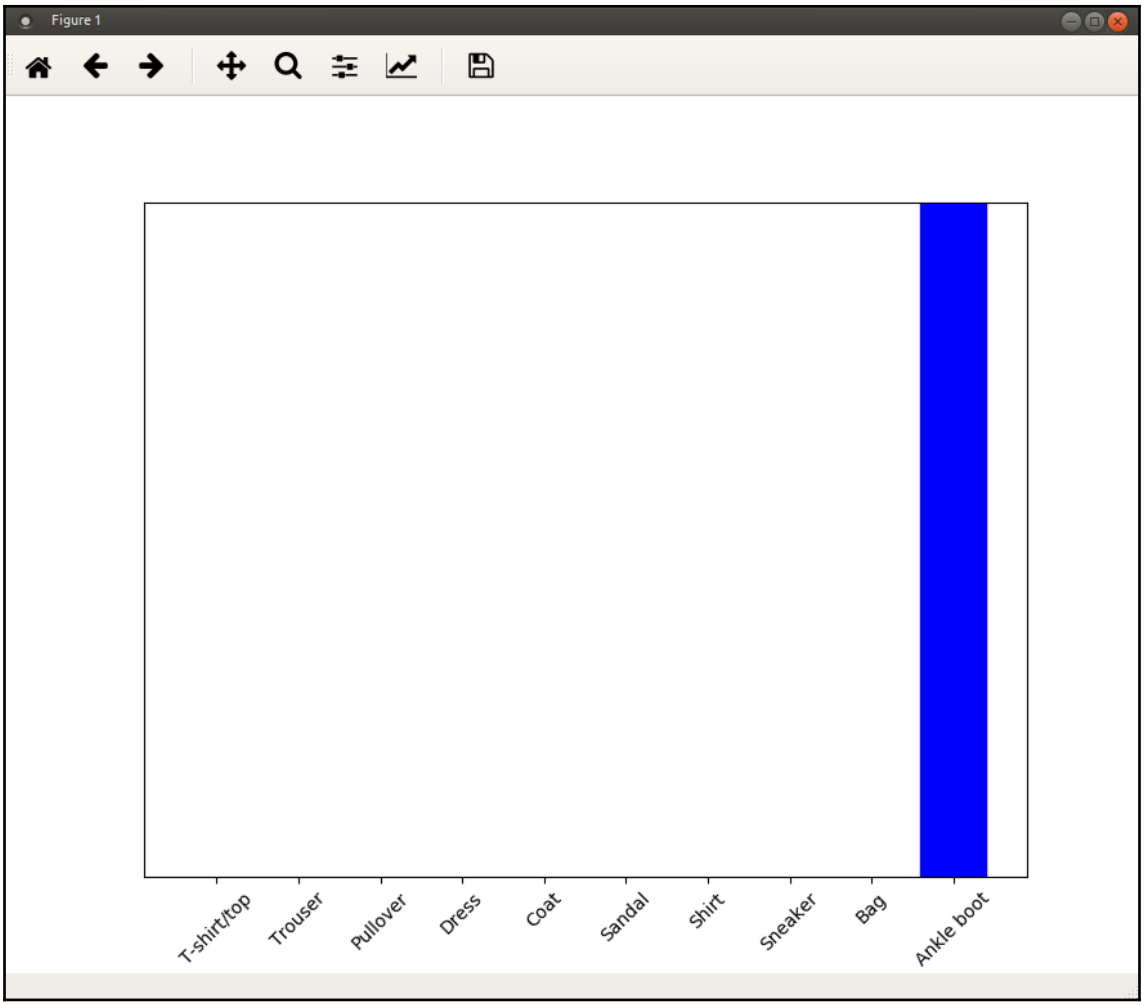
Figure 1
T-shirt/top
Trouser
Pullover
Dress
Coat
Sandal
Shirt
Sneaker
Bag
Ankle boot

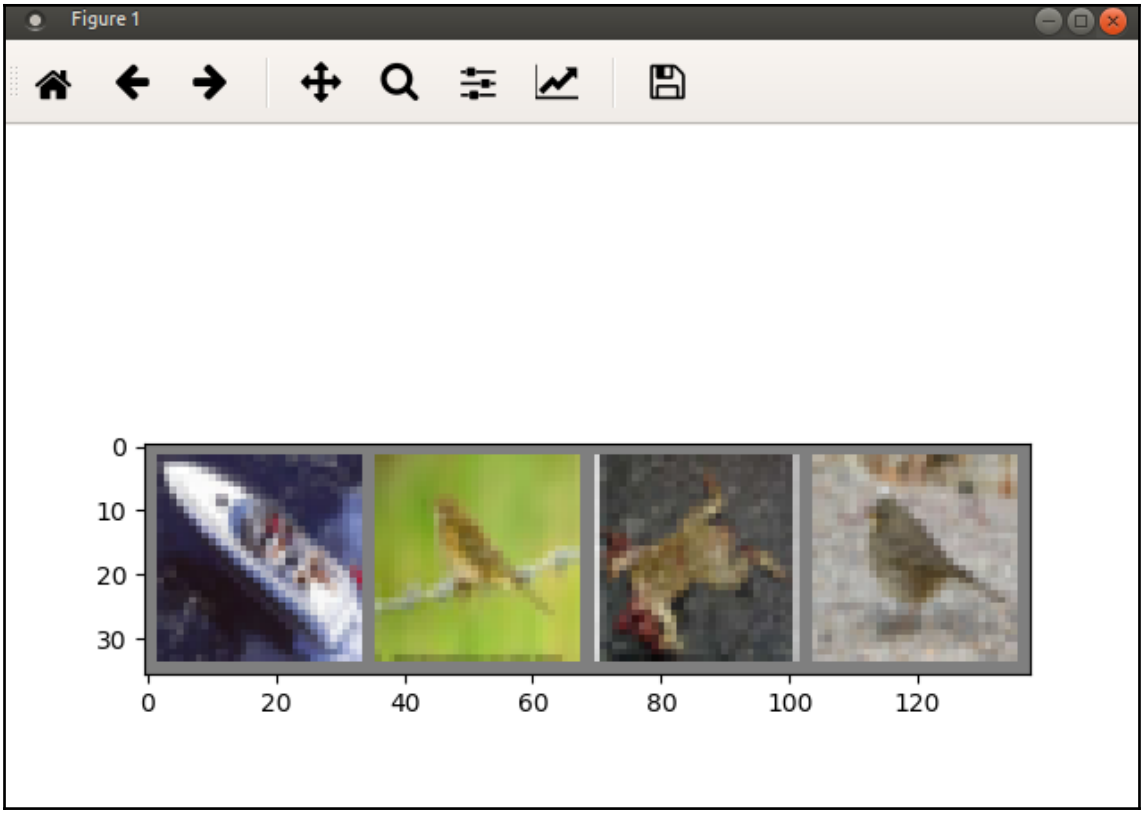
Figure 1
0
10
20
30
0
20
40
60
80
100
120

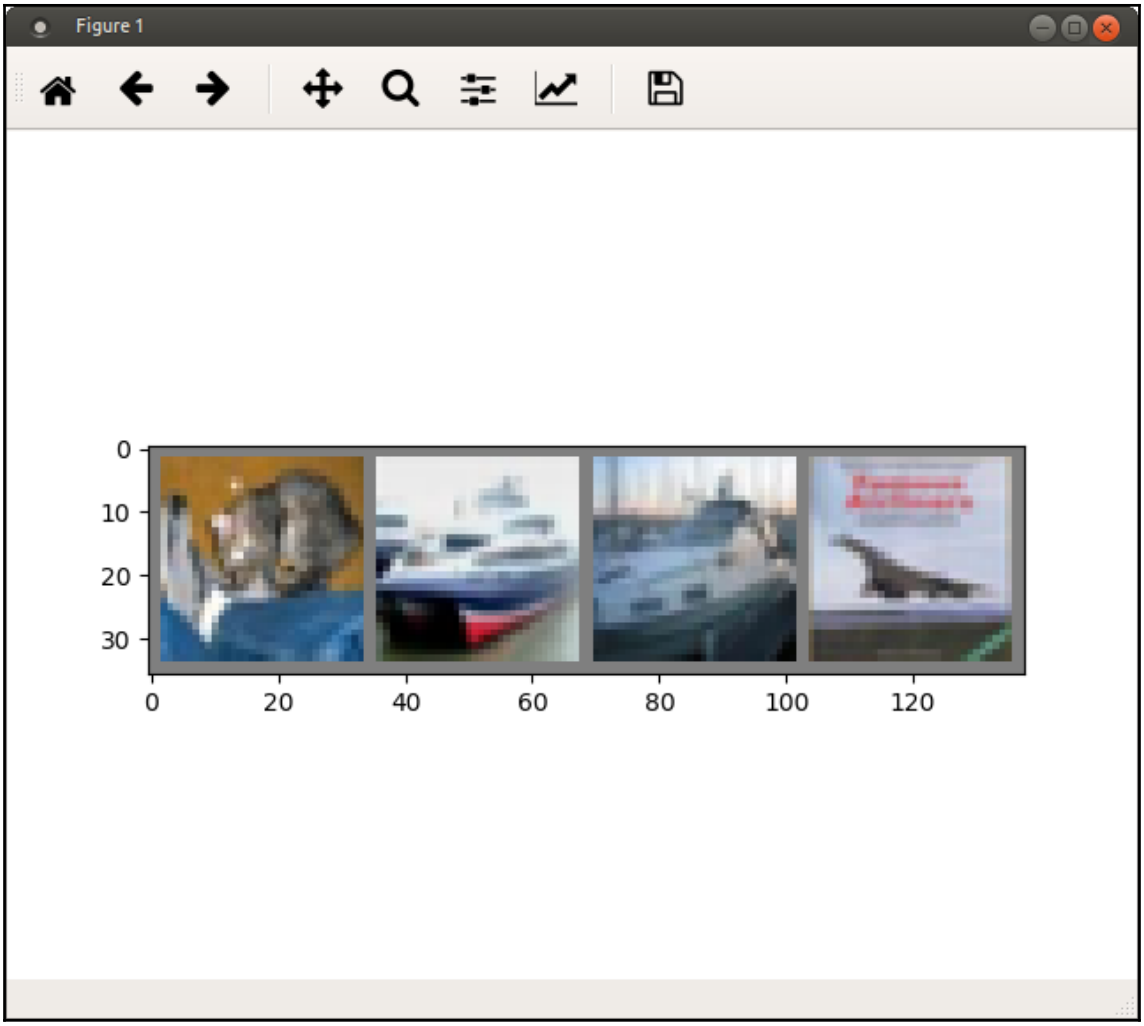
Figure 1
0
10
20
30
0
20
40
60
80
100
120

Chapter 11: GPU Acceleration for Scientific Applications Using DeepChem

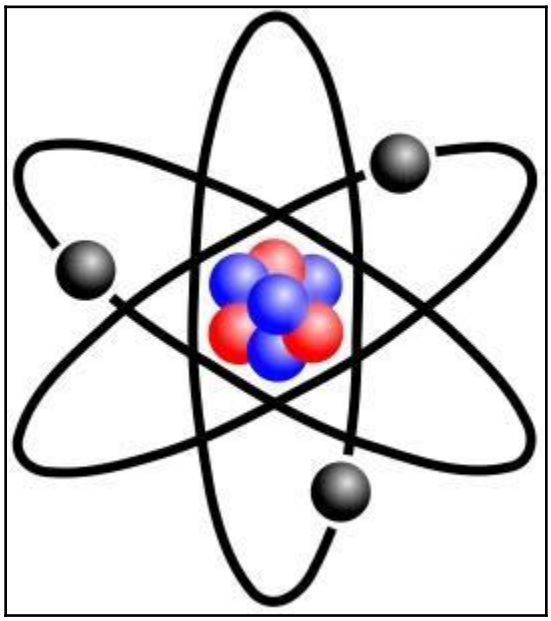

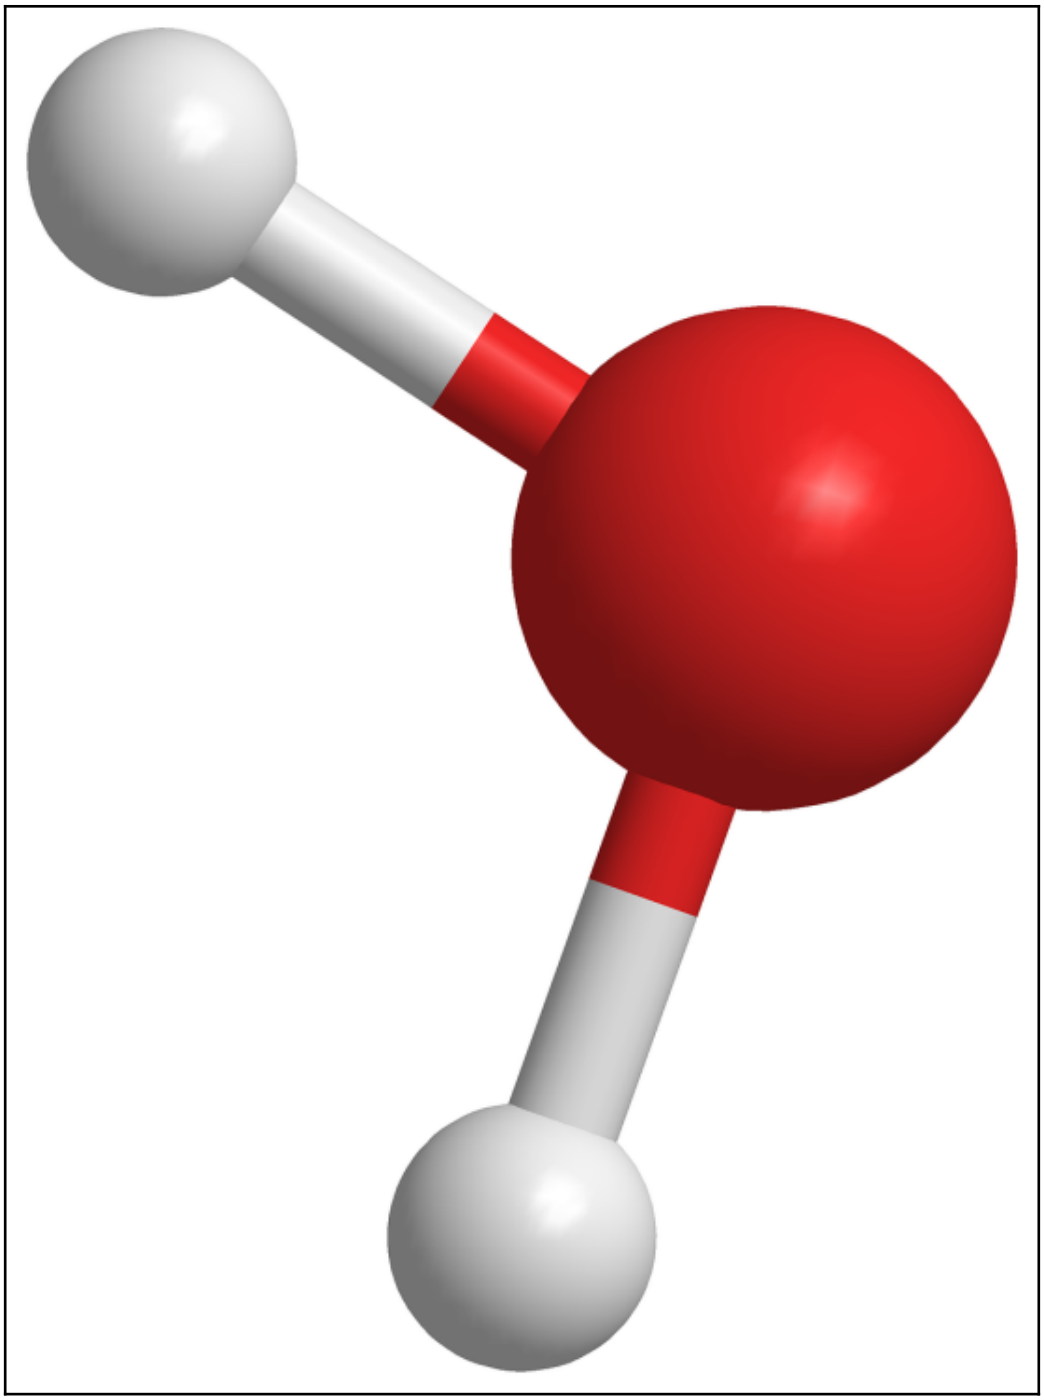

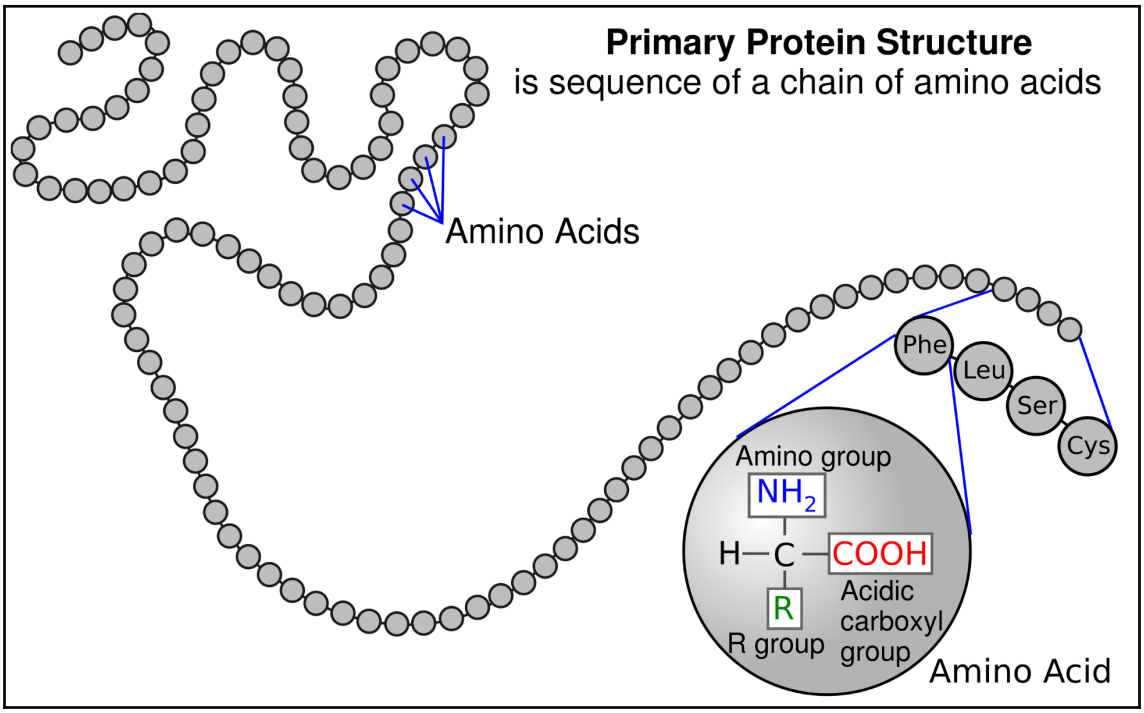
Primary Protein Structure
is sequence of a chain of amino acids
Amino Acids
Phe
Leu
Ser
Cys
Amino group
NH2
H—C—COOH
R
R group
Acidic carboxyl group
Amino Acid

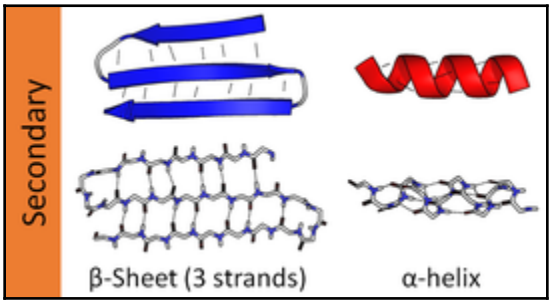
Secondary
β-Sheet (3 strands)
α-helix

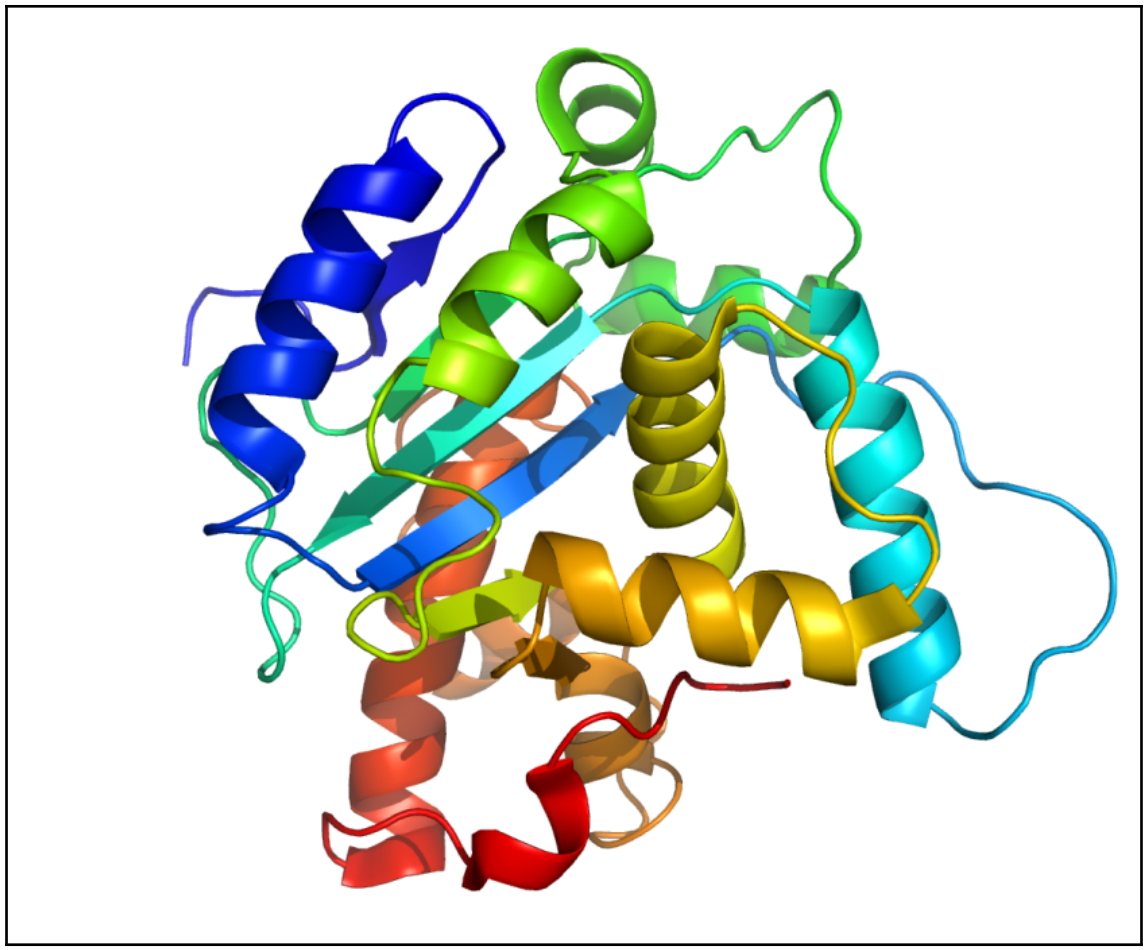

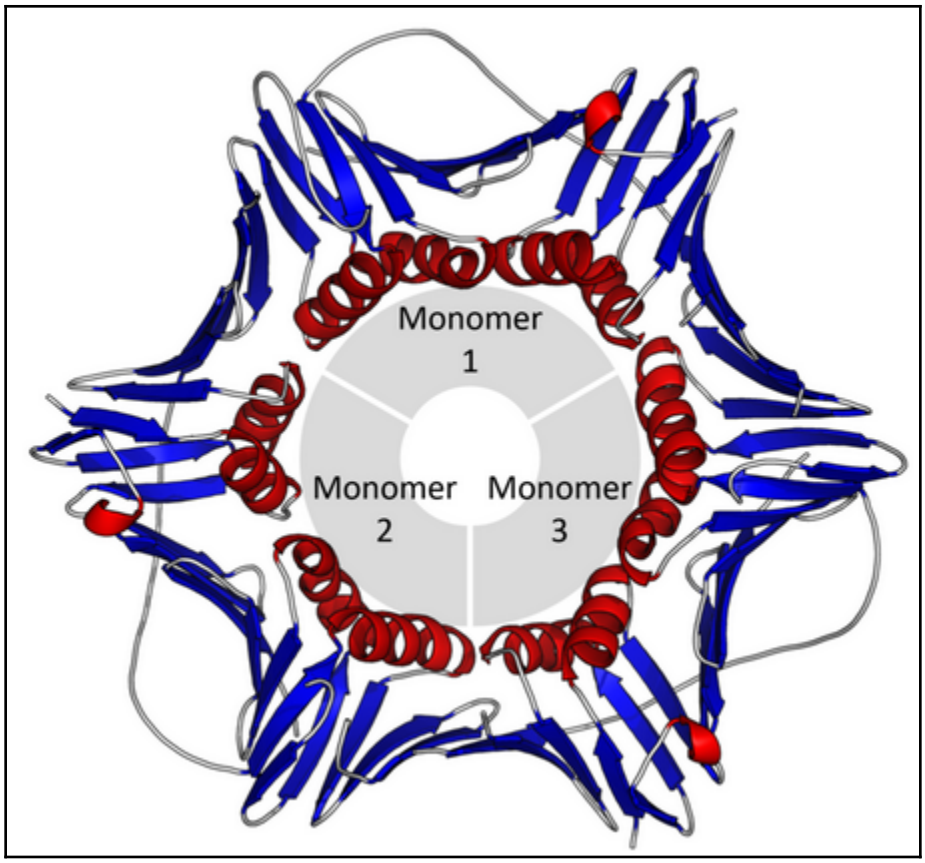
Monomer
1
Monomer
2
Monomer
3

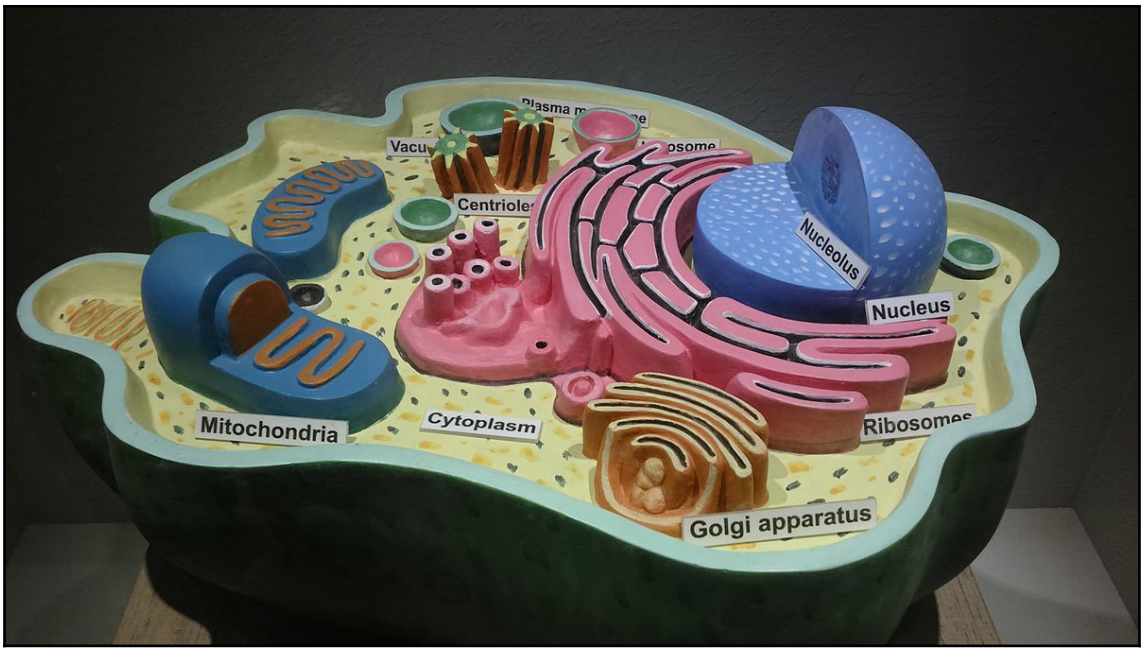
Nucleolus
Nucleus
Mitochondria
Cytoplasm
Golgi apparatus

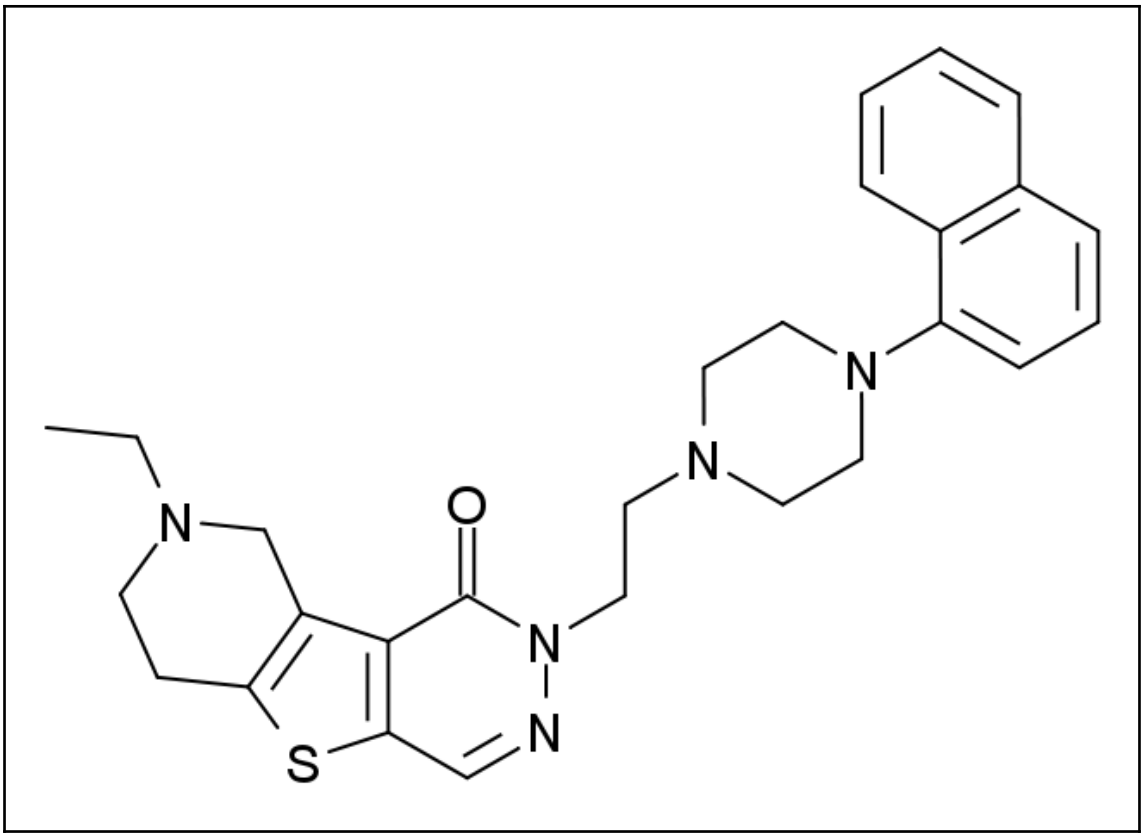
N
N
O
N
N
S

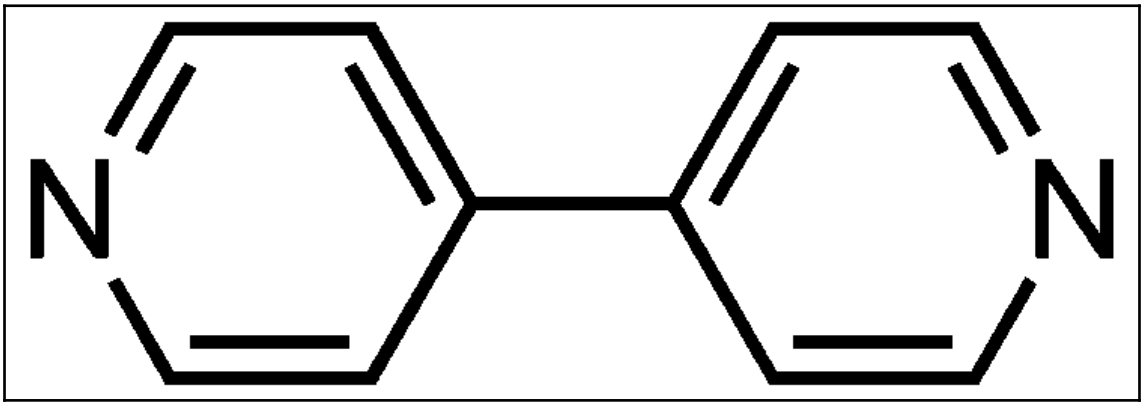
N
N

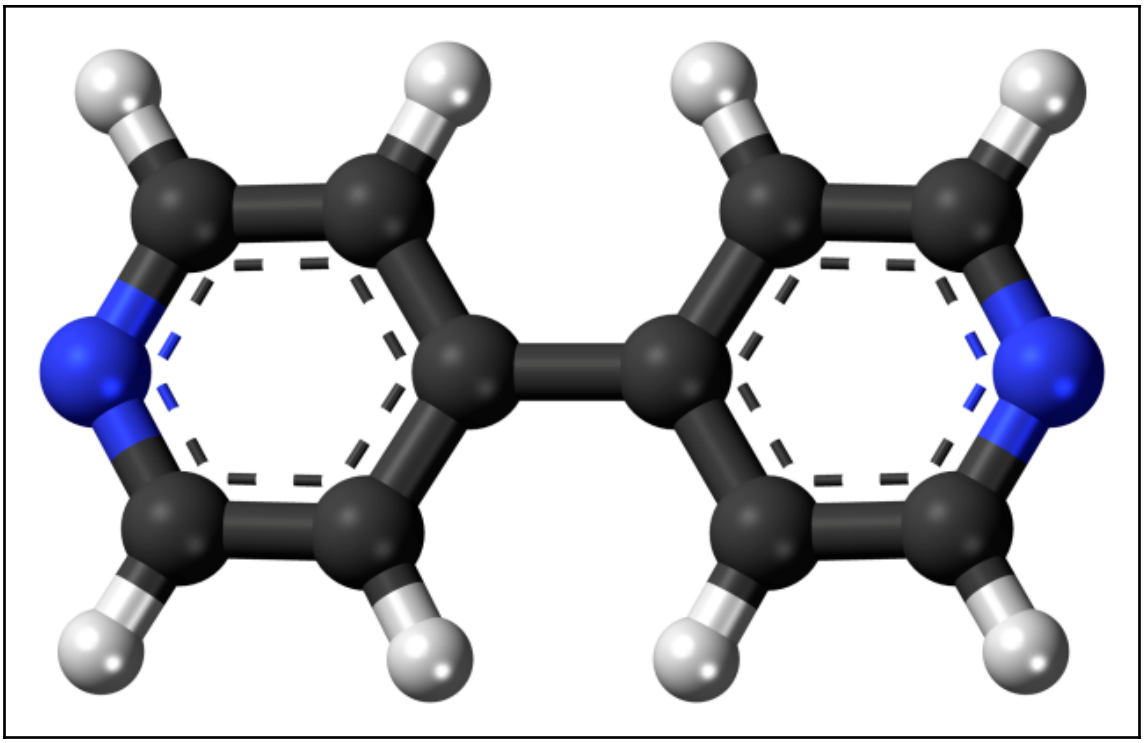

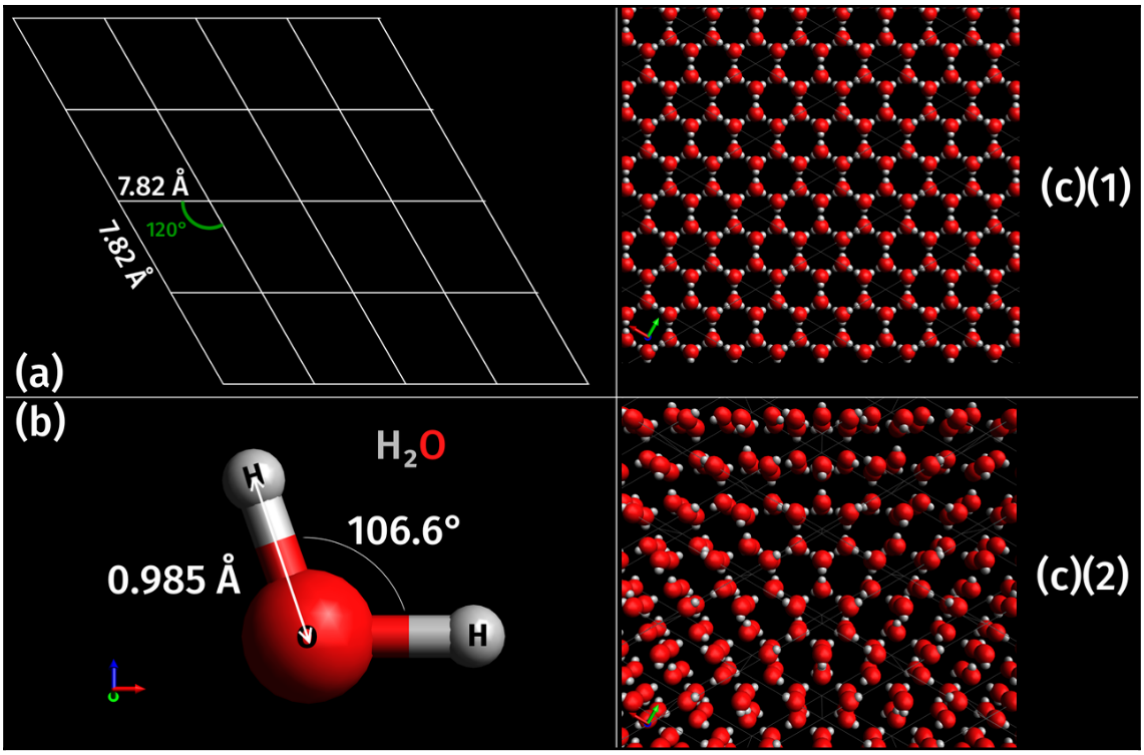
7.82 Å
120°
7.82 Å
(a)
(b)
H_2O
106.6°
0.985 Å
H
O
H
(c)(1)
(c)(2)

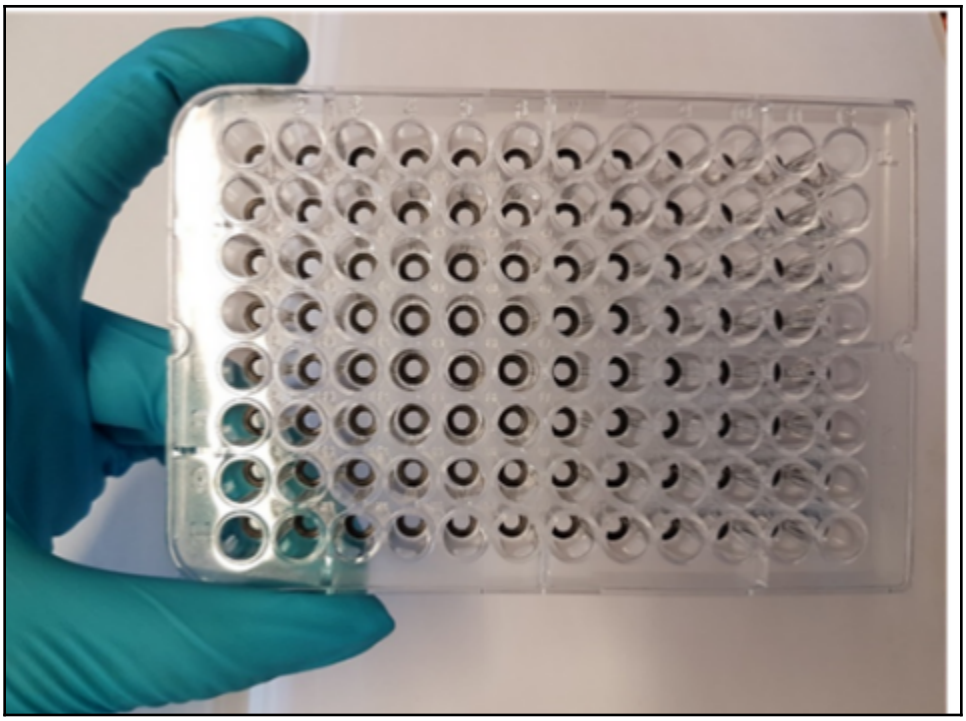

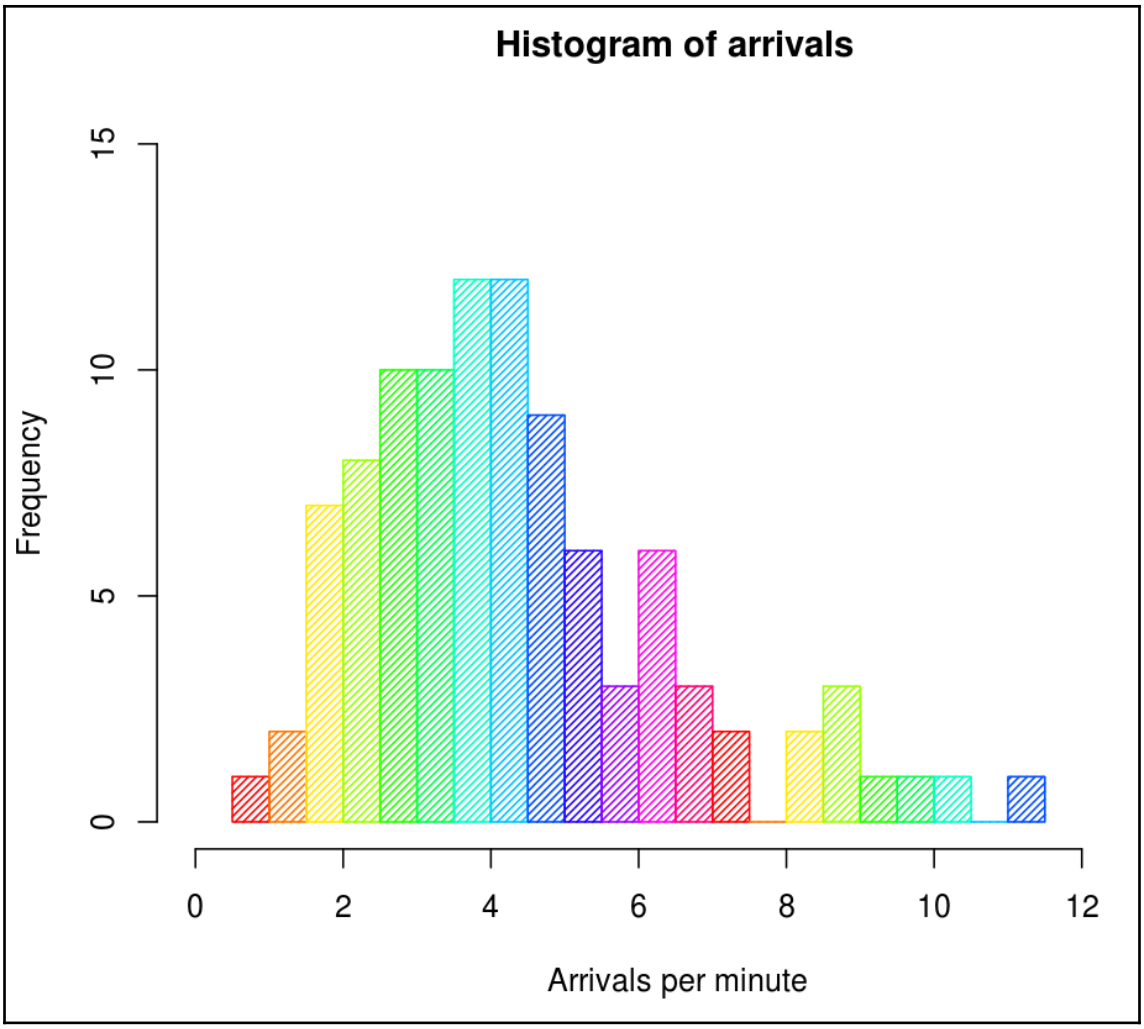
Histogram of arrivals
Frequency
0
5
10
15
0
2
4
6
8
10
12
Arrivals per minute

```
from tensorflow.python.client import device_lib
device_lib.list_local_devices()
```

```
[name: "/device:CPU:0"
 device_type: "CPU"
 memory_limit: 268435456
 locality {
 }
 incarnation: 16861077266969563597, name: "/device:XLA_CPU:0"
 device_type: "XLA_CPU"
 memory_limit: 17179869184
 locality {
 }
 incarnation: 13289663721869503433
 physical_device_desc: "device: XLA_CPU device", name: "/device:XLA_GPU:0"
 device_type: "XLA_GPU"
 memory_limit: 17179869184
 locality {
 }
 incarnation: 15512174984466094799
 physical_device_desc: "device: XLA_GPU device", name: "/device:GPU:0"
 device_type: "GPU"
 memory_limit: 14800692839
 locality {
   bus_id: 1
   links {
   }
 }
 incarnation: 158464110753286612
 physical_device_desc: "device: 0, name: Tesla T4, pci bus id: 0000:00:04.0, compute capability: 7.5"]
```

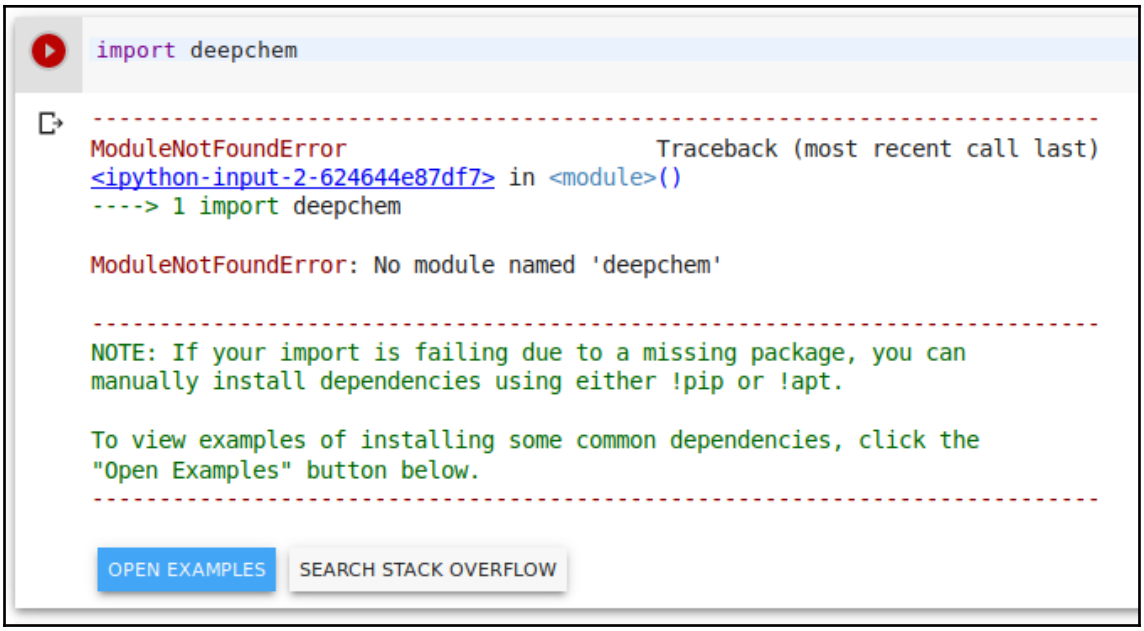
import deepchem

ModuleNotFoundError Traceback (most recent call last)
<ipython-input-2-624644e87df7> in <module>()
----> 1 import deepchem
ModuleNotFoundError: No module named 'deepchem'

NOTE: If your import is failing due to a missing package, you can
manually install dependencies using either !pip or !apt.
To view examples of installing some common dependencies, click the
"Open Examples" button below.

OPEN EXAMPLES
SEARCH STACK OVERFLOW

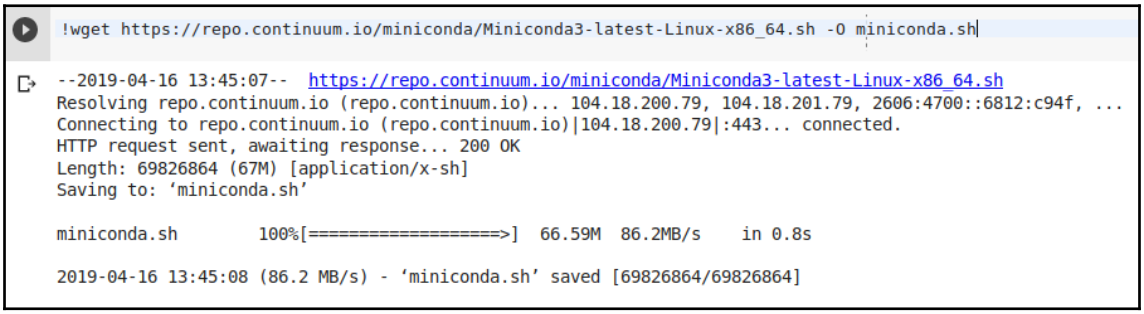
!wget https://repo.continuum.io/miniconda/Miniconda3-latest-Linux-x86_64.sh -O miniconda.sh
--2019-04-16 13:45:07-- https://repo.continuum.io/miniconda/Miniconda3-latest-Linux-x86_64.sh
Resolving repo.continuum.io (repo.continuum.io)... 104.18.200.79, 104.18.201.79, 2606:4700::6812:c94f, ...
Connecting to repo.continuum.io (repo.continuum.io)|104.18.200.79|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 69826864 (67M) [application/x-sh]
Saving to: 'miniconda.sh'
miniconda.sh 100%[==========>] 66.59M 86.2MB/s in 0.8s
2019-04-16 13:45:08 (86.2 MB/s) - 'miniconda.sh' saved [69826864/69826864]

```
!chmod +x miniconda.sh
!bash ./miniconda.sh -b -f -p /usr/local
!conda install -y --prefix /usr/local -c deepchem -c rdkit -c conda-forge -c omnia deepchem-gpu=2.1.0
```

```
PREFIX=/usr/local
reinstalling: python-3.7.1-h0371630_7 ...
Python 3.7.1
reinstalling: ca-certificates-2018.03.07-0 ...
reinstalling: conda-env-2.6.0-1 ...
reinstalling: libgcc-ng-8.2.0-hdf63c60_1 ...
reinstalling: libstdcxx-ng-8.2.0-hdf63c60_1 ...
reinstalling: libffi-3.2.1-hd88cf55_4 ...
reinstalling: ncurses-6.1-he6710b0_1 ...
reinstalling: openssl-1.1.1a-h7b6447c_0 ...
reinstalling: xz-5.2.4-h14c3975_4 ...
reinstalling: yaml-0.1.7-had09818_2 ...
reinstalling: zlib-1.2.11-h7b6447c_3 ...
reinstalling: libedit-3.1.20170329-h6b74fdf_2 ...
reinstalling: readline-7.0-h7b6447c_5 ...
reinstalling: tk-8.6.8-hbc83047_0 ...
reinstalling: sqlite-3.26.0-h7b6447c_0 ...
reinstalling: asn1crypto-0.24.0-py37_0 ...
reinstalling: certifi-2018.11.29-py37_0 ...
reinstalling: chardet-3.0.4-py37_1 ...
reinstalling: idna-2.8-py37_0 ...
```

```
import sys
sys.path
```

```
['',
 '/env/python',
 '/usr/lib/python36.zip',
 '/usr/lib/python3.6',
 '/usr/lib/python3.6/lib-dynload',
 '/usr/local/lib/python3.6/dist-packages',
 '/usr/lib/python3/dist-packages',
 '/usr/local/lib/python3.6/dist-packages/IPython/extensions',
 '/root/.ipython']
```

```
sys.path.append('/usr/local/lib/python3.6/site-packages/')
```

```
sys.path
```

```
['',
 '/env/python',
 '/usr/lib/python36.zip',
 '/usr/lib/python3.6',
 '/usr/lib/python3.6/lib-dynload',
 '/usr/local/lib/python3.6/dist-packages',
 '/usr/lib/python3/dist-packages',
 '/usr/local/lib/python3.6/dist-packages/IPython/extensions',
 '/root/.ipython',
 '/usr/local/lib/python3.6/site-packages/']
```

```
import deepchem
```

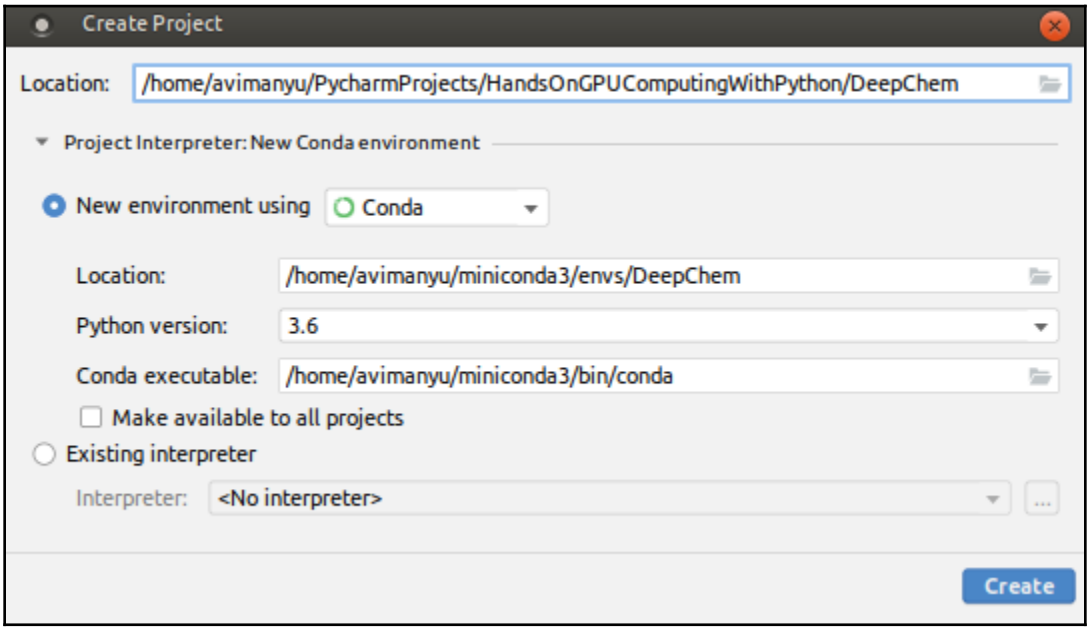

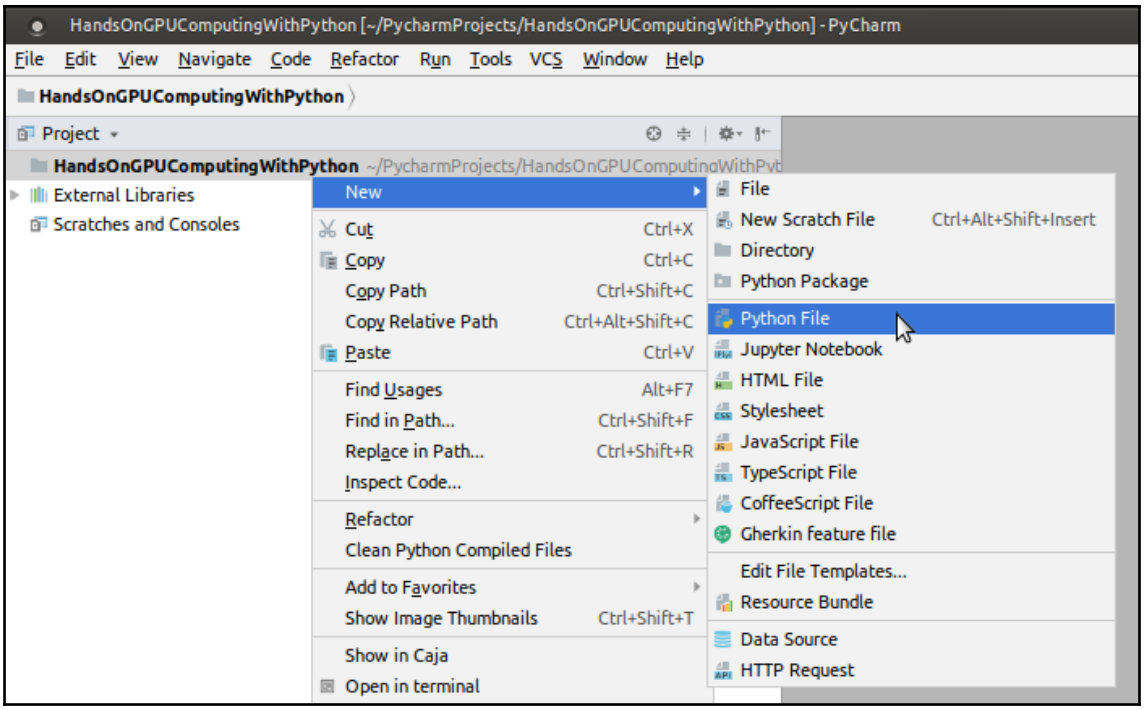
HandsOnGPUComputingWithPython [~/PycharmProjects/HandsOnGPUComputingWithPython] - PyCharm
File Edit View Navigate Code Refactor Run Tools VCS Window Help
HandsOnGPUComputingWithPython
Project
HandsOnGPUComputingWithPython ~/PycharmProjects/HandsOnGPUComputingWithPyt
External Libraries
Scratches and Consoles
New
Cut Ctrl+X
Copy Ctrl+C
Copy Path Ctrl+Shift+C
Copy Relative Path Ctrl+Alt+Shift+C
Paste Ctrl+V
Find Usages Alt+F7
Find in Path... Ctrl+Shift+F
Replace in Path... Ctrl+Shift+R
Inspect Code...
Refactor
Clean Python Compiled Files
Add to Favorites
Show Image Thumbnails Ctrl+Shift+T
Show in Caja
Open in terminal
File
New Scratch File Ctrl+Alt+Shift+Insert
Directory
Python Package
Python File
Jupyter Notebook
HTML File
Stylesheet
JavaScript File
TypeScript File
CoffeeScript File
Gherkin feature file
Edit File Templates...
Resource Bundle
Data Source
HTTP Request


import dee
deepchem

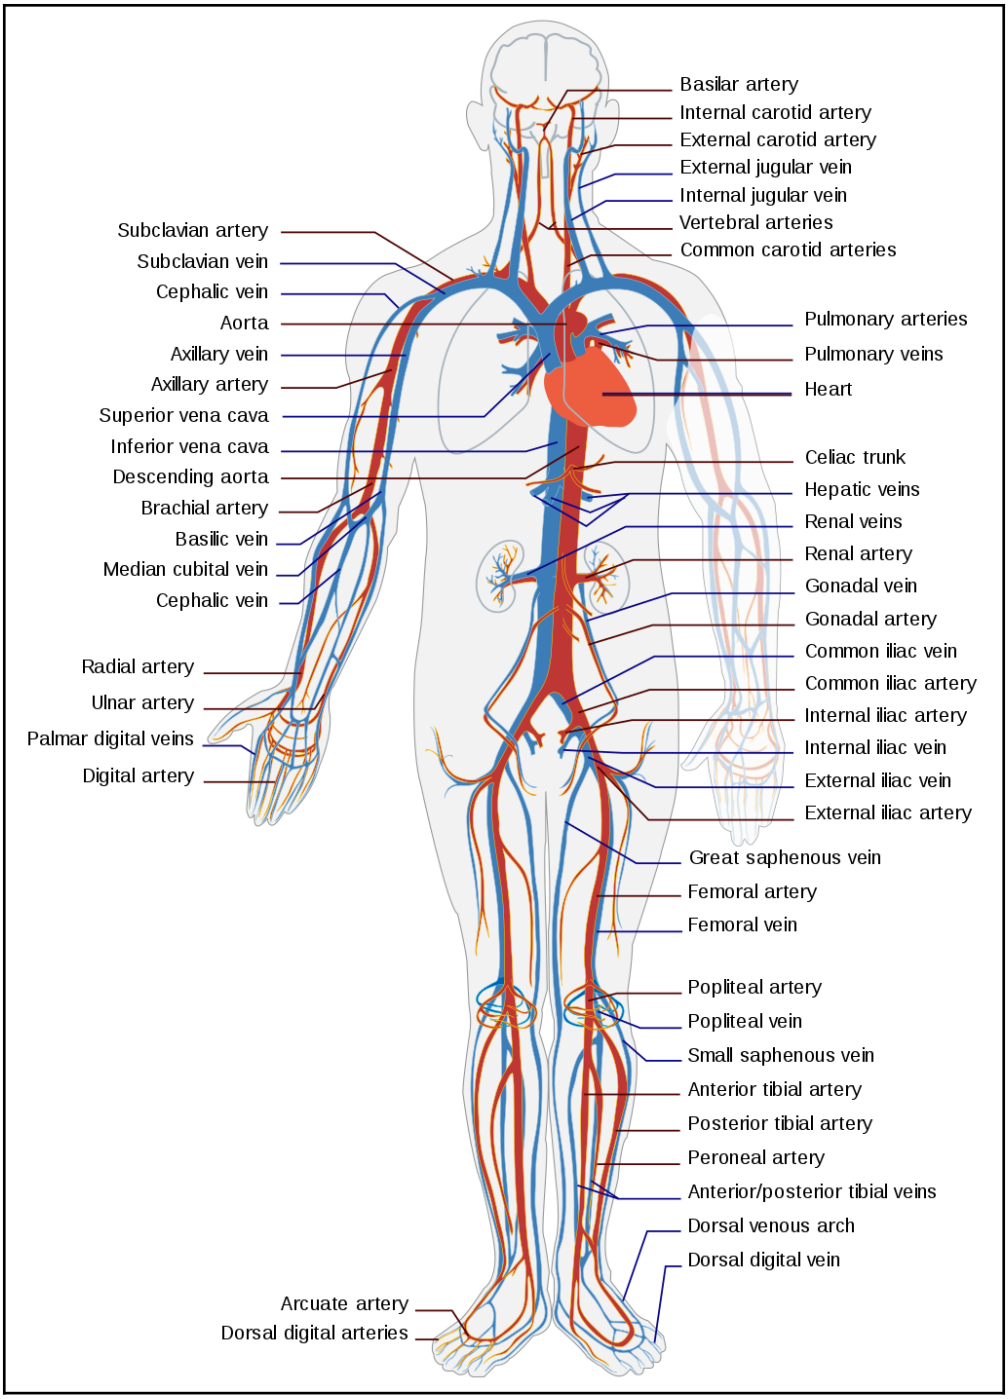
Basilar artery
Internal carotid artery
External carotid artery
External jugular vein
Internal jugular vein
Vertebral arteries
Common carotid arteries
Subclavian artery
Subclavian vein
Cephalic vein
Aorta
Axillary vein
Axillary artery
Superior vena cava
Inferior vena cava
Descending aorta
Brachial artery
Basilic vein
Median cubital vein
Cephalic vein
Radial artery
Ulnar artery
Palmar digital veins
Digital artery
Pulmonary arteries
Pulmonary veins
Heart
Celiac trunk
Hepatic veins
Renal veins
Renal artery
Gonadal vein
Gonadal artery
Common iliac vein
Common iliac artery
Internal iliac artery
Internal iliac vein
External iliac vein
External iliac artery
Great saphenous vein
Femoral artery
Femoral vein
Popliteal artery
Popliteal vein
Small saphenous vein
Anterior tibial artery
Posterior tibial artery
Peroneal artery
Anterior/posterior tibial veins
Dorsal venous arch
Dorsal digital vein
Arcuate artery
Dorsal digital arteries

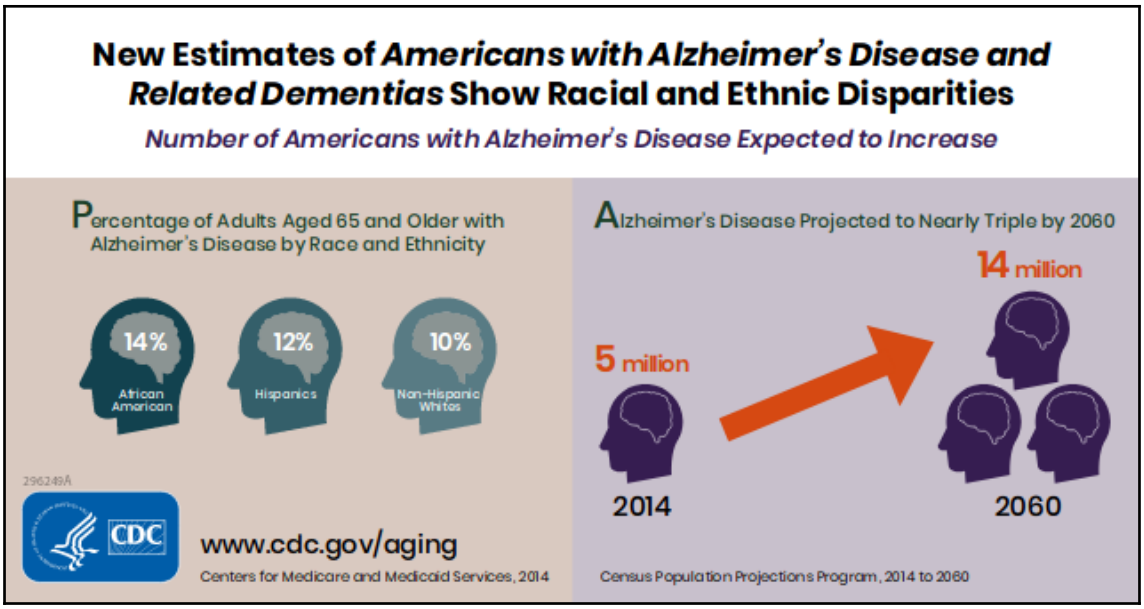
New Estimates of Americans with Alzheimer's Disease and Related Dementias Show Racial and Ethnic Disparities
Number of Americans with Alzheimer's Disease Expected to Increase
Percentage of Adults Aged 65 and Older with Alzheimer's Disease by Race and Ethnicity
14%
African American
12%
Hispanics
10%
Non-Hispanic Whites
296249A
CDC
www.cdc.gov/aging
Centers for Medicare and Medicaid Services, 2014
Alzheimer's Disease Projected to Nearly Triple by 2060
14 million
5 million
2014
2060
Census Population Projections Program, 2014 to 2060

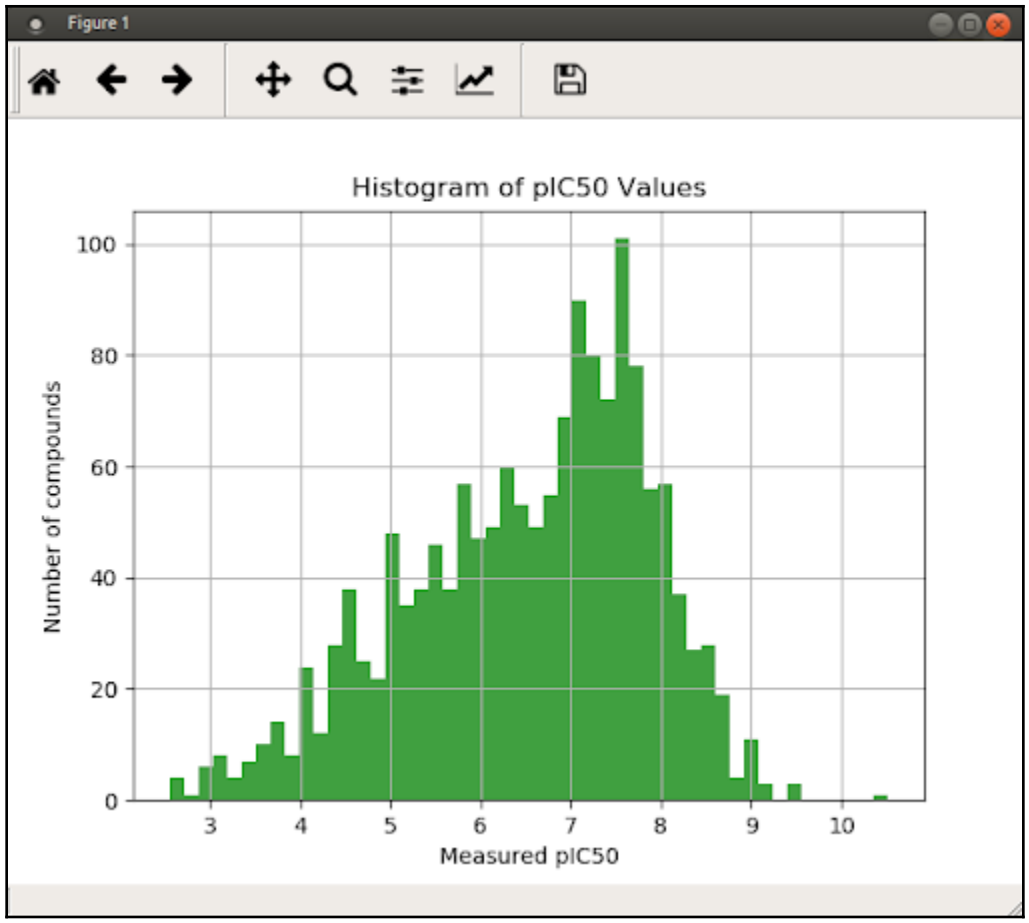
Figure 1
Histogram of pIC50 Values
Number of compounds
Measured pIC50
0
20
40
60
80
100
3
4
5
6
7
8
9
10

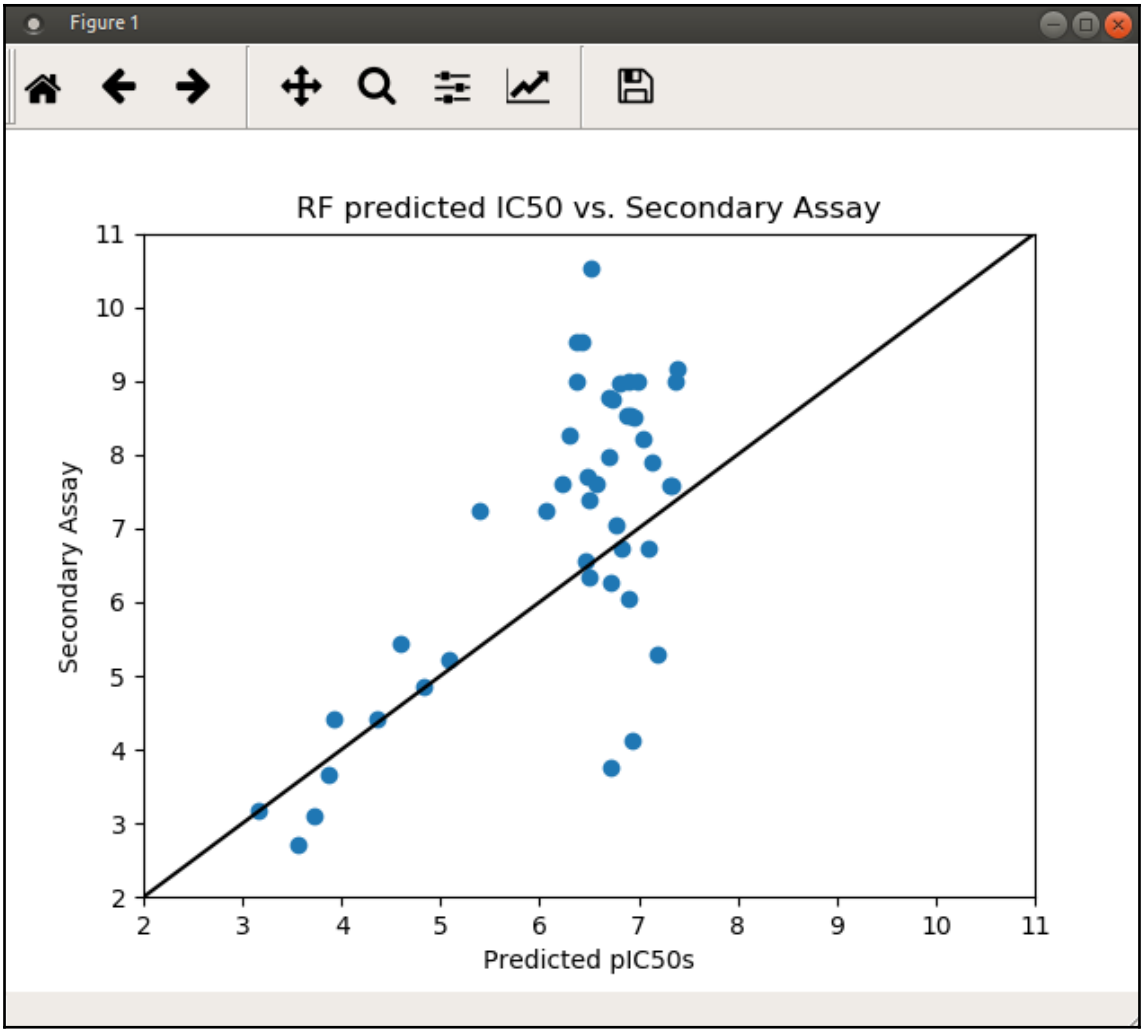
Figure 1
RF predicted IC50 vs. Secondary Assay
Secondary Assay
Predicted pIC50s
2
3
4
5
6
7
8
9
10
11

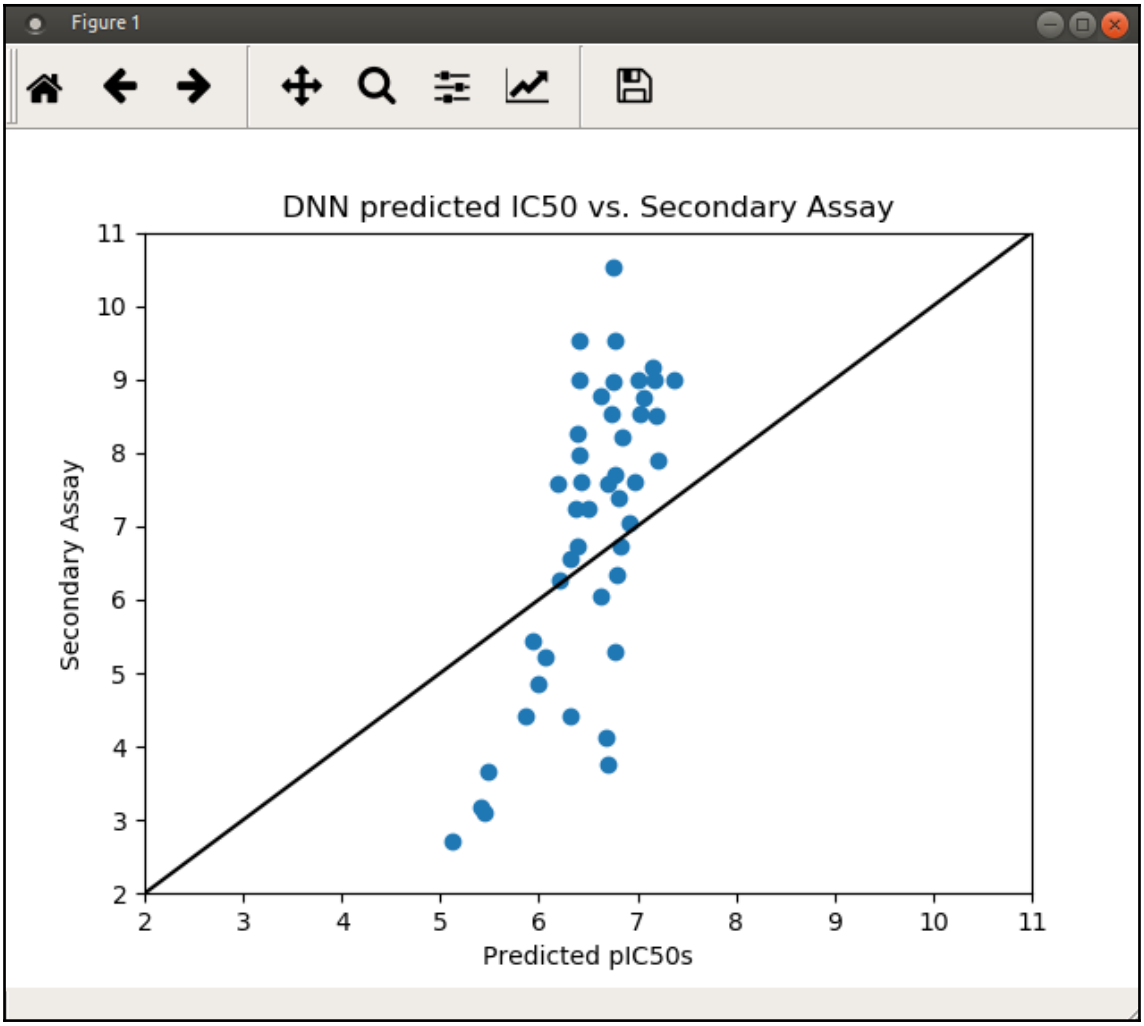
Figure 1
DNN predicted IC50 vs. Secondary Assay
Secondary Assay
Predicted pIC50s
2
3
4
5
6
7
8
9
10
11

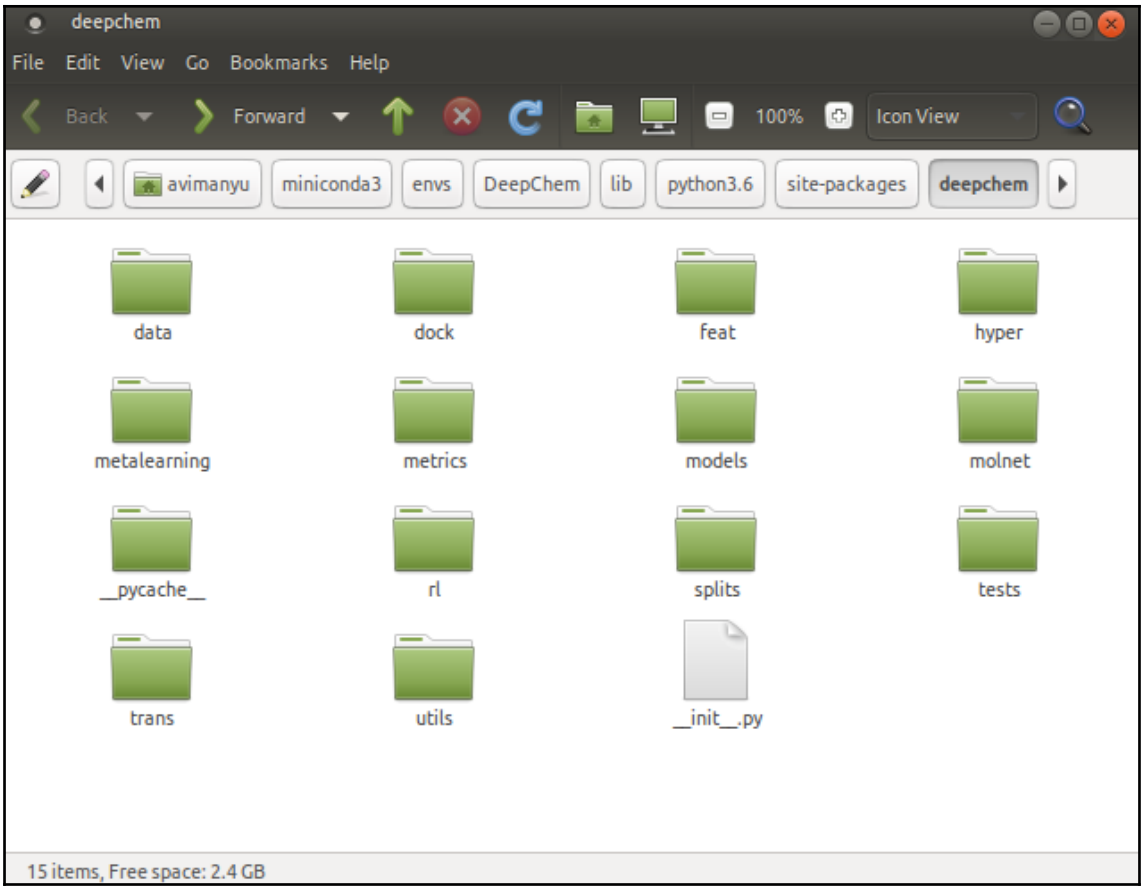
deepchem
File Edit View Go Bookmarks Help
Back Forward 100% Icon View
avimanyu miniconda3 envs DeepChem lib python3.6 site-packages deepchem
data
dock
feat
hyper
metalearning
metrics
models
molnet
__pycache__
rl
splits
tests
trans
utils
__init__.py
15 items, Free space: 2.4 GB

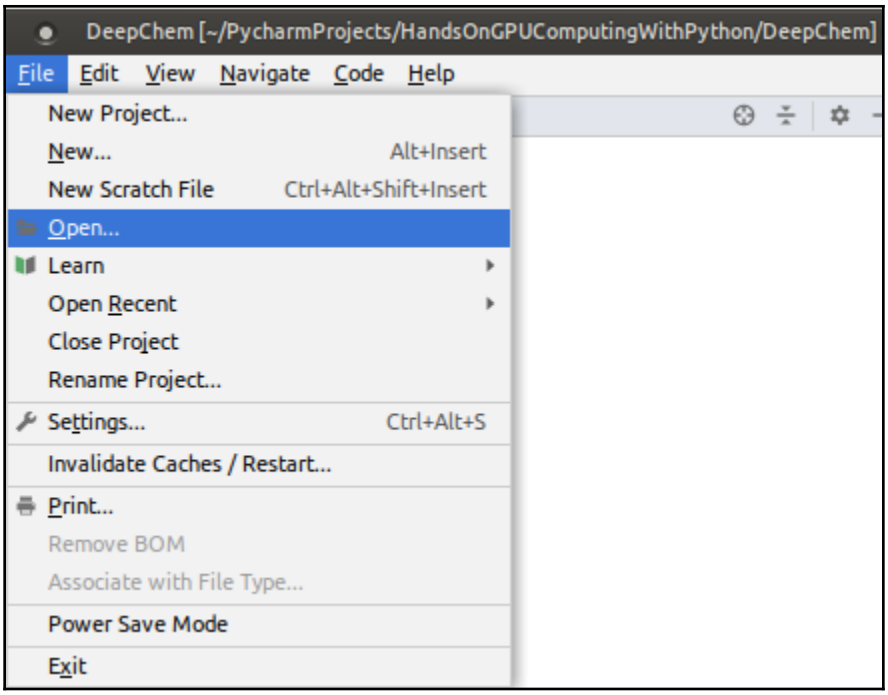
DeepChem [~/PycharmProjects/HandsOnGPUComputingWithPython/DeepChem]
File
Edit
View
Navigate
Code
Help
New Project...
New... Alt+Insert
New Scratch File Ctrl+Alt+Shift+Insert
Open...
Learn
Open Recent
Close Project
Rename Project...
Settings... Ctrl+Alt+S
Invalidate Caches / Restart...
Print...
Remove BOM
Associate with File Type...
Power Save Mode
Exit

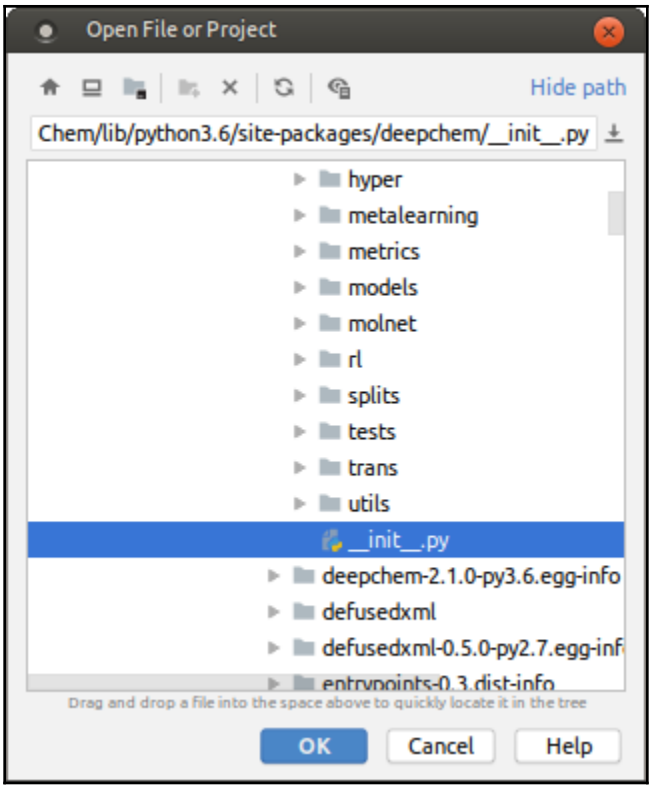
Open File or Project
Hide path
Chem/lib/python3.6/site-packages/deepchem/__init__.py
hyper
metalearning
metrics
models
molnet
rl
splits
tests
trans
utils
__init__.py
deepchem-2.1.0-py3.6.egg-info
defusedxml
defusedxml-0.5.0-py2.7.egg-inf
entrypoints-0.3.dist-info
Drag and drop a file into the space above to quickly locate it in the tree
OK
Cancel
Help

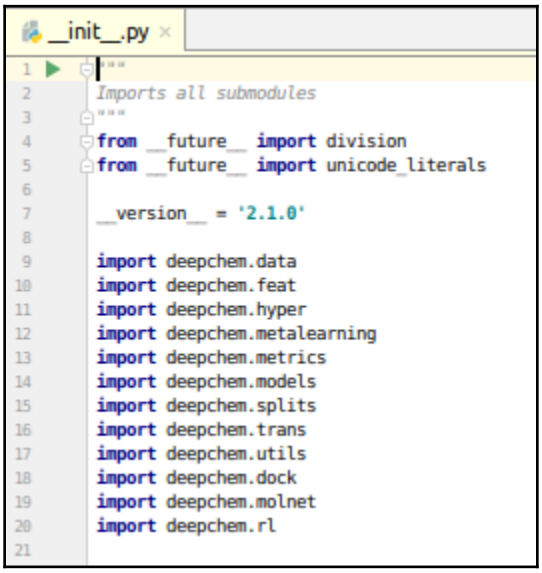
__init__.py
"""
Imports all submodules
"""
from __future__ import division
from __future__ import unicode_literals
__version__ = '2.1.0'
import deepchem.data
import deepchem.feat
import deepchem.hyper
import deepchem.metalearning
import deepchem.metrics
import deepchem.models
import deepchem.splits
import deepchem.trans
import deepchem.utils
import deepchem.dock
import deepchem.molnet
import deepchem.rl