

# How do I install PyTorch and TensorFlow with Python to use ccln2p3 GPU?

Thomas BAUDIER

Research engineer

[thomas.baudier@creatis.insa-lyon.fr](mailto:thomas.baudier@creatis.insa-lyon.fr)

- PyTorch and TensorFlow available without GPU
- How I did -> Probably not the best way to do
- For people who want to start quickly
- PyTorch 1.0.1.post2 and TensorFlow 1.13.0 (latest version on 28/03/2019)

 PyTorch

# Outline

- Requirements
  - Python version
  - Cuda version
- Installation
- Use
- Application example

# Requirements - Python

- Python versions available at ccin2p3

```

/pbs/software/centos-7-x86_64/python(0)>ll
total 394
drwxrwxr-x 9 dbouvet ccin2p3 163 11 avril 2018 2.7.14
drwxrwxr-x 9 dbouvet ccin2p3 163 4 juin 2018 2.7.15
drwxr-xr-x 8 dbouvet ccin2p3 142 5 oct. 17:22 2.7.3
drwxrwxr-x 9 gadrat ccin2p3 163 26 sept. 2018 3.4.8
drwxrwxr-x 9 gadrat ccin2p3 163 26 sept. 2018 3.6.5
drwxr-xr-x 9 gadrat ccin2p3 163 5 déc. 10:45 3.6.7
-rwxrwxr-x 1 dbouvet ccin2p3 331 20 sept. 2018 ccenv.csh
-rwxrwxr-x 1 gadrat ccin2p3 230 20 sept. 2018 ccenv.sh
-rw-rw-r-- 1 gadrat ccin2p3 40 20 août 2018 ccenv.yaml
drwxr-xr-x 4 gadrat ccin2p3 79 4 mars 13:28 modules
-rw-rw-r-- 1 gadrat ccin2p3 2853 19 févr. 15:52 VERSIONS
    
```

- Mandatory versions

<https://pytorch.org/get-started/locally/>

PyTorch Build	Stable (1.0)		Preview (Nightly)	
Your OS	Linux	Mac	Windows	
Package	Conda	Pip	LibTorch	Source
Language	Python 2.7	Python 3.5	Python 3.6	Python 3.7, C++
CUDA	8.0	9.0	10.0	None

Run this Command:

```

pip3 install https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl
pip3 install torchvision
    
```

<https://www.tensorflow.org/install/pip?lang=python3>

Python 3
  Python 2.7

Check if your Python environment is already configured:

★ Requires Python 3.4, 3.5, or 3.6

# Requirements - Python

- Change python version

```
/pbs/software/centos-7-x86_64/python(0)>cat ~/.profile | grep python -A 3  
. python_env.sh 3.6.5  
export LC_ALL=fr_FR.utf-8  
export LANG=fr_FR.utf-8
```

- Creation of virtual environment

```
/pbs/throng/creatis/users/tbaudier(0)>cd  
/pbs/home/t/tbaudier(0)>  
/pbs/home/t/tbaudier(0)>  
/pbs/home/t/tbaudier(0)>cd /pbs/throng/creatis/users/tbaudier/  
/pbs/throng/creatis/users/tbaudier(0)>mkdir venvAI  
/pbs/throng/creatis/users/tbaudier(0)>python_env.sh 3.6.7  
You are using version 3.6.7 of python.  
/pbs/throng/creatis/users/tbaudier(0)>pip install --user virtualenv  
Collecting virtualenv  
Using cached https://files.pythonhosted.org/packages/33/5d/314c760d4204f64e4a968275182b7751bd5c3249094757b39ba987dcfb5a/virtualenv-16.4.3-py2.py3-none-any.whl  
Installing collected packages: virtualenv  
The script virtualenv is installed in '/pbs/home/t/tbaudier/.local/bin' which is not on PATH.  
Consider adding this directory to PATH or, if you prefer to suppress this warning, use --no-warn-script-location.  
Successfully installed virtualenv-16.4.3  
You are using pip version 19.0.2, however version 19.0.3 is available.  
You should consider upgrading via the 'pip install --upgrade pip' command.  
/pbs/throng/creatis/users/tbaudier(0)>virtualenv venvAI/  
Using base prefix '/pbs/software/centos-7-x86_64/python/3.6.7'  
New python executable in /pbs/throng/creatis/users/tbaudier/venvAI/bin/python3.6  
Also creating executable in /pbs/throng/creatis/users/tbaudier/venvAI/bin/python  
Installing setuptools, pip, wheel...  
done.  
/pbs/throng/creatis/users/tbaudier(0)>□
```

# Requirements - Cuda

- Connect to interactive GPU jobs

```

/pbs/home/t/tbaudier(0)>qlogin -l GPU=1 -l GPUtype=K80 -q mc_gpu_interactive -pe multicores_gpu 4
JSV "/opt/sge/util/resources/jsv/corebinding.jsv" has been started
JSV "/opt/sge/util/resources/jsv/corebinding.jsv" has been stopped
Your job 2206803 ("QLOGIN") has been submitted
waiting for interactive job to be scheduled ...
Your interactive job 2206803 has been successfully scheduled.
Establishing /usr/bin/qlogin_wrapper session to host ccwgige010.in2p3.fr ...
The authenticity of host '[ccwgige010.in2p3.fr]:41865 ([134.158.48.126]:41865)' can't be established.
ECDSA key fingerprint is SHA256:wmXwcX0cJrRE/D5kyxYRVq2hYpF7PTYHud6X4/Ygmh0.
ECDSA key fingerprint is MD5:93:0b:6c:35:dc:cd:6f:eb:9b:39:68:03:44:d7:a5:dc.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '[ccwgige010.in2p3.fr]:41865,[134.158.48.126]:41865' (ECDSA) to the list of known hosts.
tbaudier@ccwgige010.in2p3.fr's password:
Last login: Thu Mar 28 10:10:16 2019 from cca004.in2p3.fr
  
```

```

  _ _ _  _ _ _  _ _  _ _ _  _ _ _  _ _ _  _ _ _  _ _ _
 / _ _ / _ _  _ _ ( ) ( ( \ _ _ \ ( _ _ \ ( _ _ \
 ( _ _ ( _ _ ( _ _ ) ( / _ _ / / _ _ / ) _ _ / ( _ _ (
 \ _ _ \ _ _ ) ( _ ) \ _ _ ) ( _ _ ) ( _ _ ) ( _ _ /
  
```

```

Platform: CentOS 7.6.1810
Architecture: x86_64
  
```

```


You are using version 3.6.5 of python.
[tbaudier@ccwgige010 ~]$ █
  
```

# Requirements - Cuda

- Cuda versions available at ccln2p3

```
[tbaudier@ccwige010 ~]$ ll -d /opt/cuda-*
drwxr-xr-x 11 root root 4096  8 mars  14:07 /opt/cuda-10.1
drwxr-xr-x  3 root root  20 18 sept. 2018 /opt/cuda-7.5
drwxr-xr-x 11 root root 4096 26 sept. 2017 /opt/cuda-8.0
drwxr-xr-x 12 root root 4096 15 oct.  17:31 /opt/cuda-9.2
```

- Mandatory versions

 PyTorch <https://pytorch.org/get-started/locally/>

PyTorch Build	Stable (1.0)		Preview (Nightly)		
Your OS	Linux		Mac	Windows	
Package	Conda	Pip	LibTorch	Source	
Language	Python 2.7	Python 3.5	Python 3.6	Python 3.7	C++
CUDA	8.0	9.0	10.0	None	
Run this Command:	<pre>pip3 install https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl pip3 install torchvision</pre>				



<https://www.tensorflow.org/install/source>

Version	Python version	Compiler	Build tools	cuDNN	CUDA
tensorflow_gpu-1.13.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.19.2	7.4	10.0
tensorflow_gpu-1.12.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.15.0	7	9
tensorflow_gpu-1.11.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.15.0	7	9
tensorflow_gpu-1.10.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.15.0	7	9
tensorflow_gpu-1.9.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.11.0	7	9
tensorflow_gpu-1.8.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.10.0	7	9
tensorflow_gpu-1.7.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.9.0	7	9
tensorflow_gpu-1.6.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.9.0	7	9
tensorflow_gpu-1.5.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.8.0	7	9
tensorflow_gpu-1.4.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.5.4	6	8
tensorflow_gpu-1.3.0	2.7, 3.3-3.6	GCC 4.8	Bazel 0.4.5	6	8

# Requirements - Cuda

- Download and install Cuda 10.0

<https://developer.nvidia.com/cuda-10.0-download-archive>

## CUDA Toolkit 10.0 Archive

**Select Target Platform**

Click on the green buttons that describe your target platform. Only supported platforms will be shown.

<b>Operating System</b>	<div style="display: flex; gap: 5px;"> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">Windows</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">Linux</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">Mac OSX</div> </div>
<b>Architecture</b>	<div style="display: flex; gap: 5px;"> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">x86_64</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">ppc64le</div> </div>
<b>Distribution</b>	<div style="display: flex; gap: 5px;"> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">Fedora</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">OpenSUSE</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">RHEL</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">CentOS</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">SLES</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">Ubuntu</div> </div>
<b>Version</b>	<div style="display: flex; gap: 5px;"> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">7</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">6</div> </div>
<b>Installer Type</b>	<div style="display: flex; gap: 5px;"> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">runfile [local]</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">rpm [local]</div> <div style="background-color: #76b82a; padding: 2px 5px; border: 1px solid #ccc;">rpm [network]</div> </div>

**Download Installer for Linux CentOS 7 x86\_64**

The base installer is available for download below.

Base Installer
Download (2.0 GB)

Installation Instructions:

1. Run `sudo sh cuda_10.0.130_410.48_linux.run``
2. Follow the command-line prompts

The CUDA Toolkit contains Open-Source Software. The source code can be found [here](#).  
 The checksums for the installer and patches can be found in [Installer Checksums](#).  
 For further information, see the [Installation Guide for Linux](#) and the [CUDA Quick Start Guide](#).



# Requirements - Cuda

- Download and install Cuda 10.0 - in an interactive GPU job

```
[tbaudier@ccwige010 ~]$ cd /pbs/throng/creatis/users/tbaudier/
[tbaudier@ccwige010 tbaudier]$ mkdir cuda100
[tbaudier@ccwige010 tbaudier]$ cd cuda100/
[tbaudier@ccwige010 cuda100]$ wget https://developer.nvidia.com/compute/cuda/10.0/Prod/local_installers/cuda_10.0.130_410.48_linux
--2019-03-28 14:54:53-- https://developer.nvidia.com/compute/cuda/10.0/Prod/local_installers/cuda_10.0.130_410.48_linux
Résolution de developer.nvidia.com (developer.nvidia.com)... 192.229.182.215
Connexion vers developer.nvidia.com (developer.nvidia.com)|192.229.182.215|:443...connecté.
requête HTTP transmise, en attente de la réponse...302 Found
Emplacement: https://developer.download.nvidia.com/compute/cuda/10.0/secure/Prod/local_installers/cuda_10.0.130_410.48_linux.run?8af_
iNNAU_o7b6HV6EnMCIj9BUxyPLjy9DV9rdyTLrZ3FzXgZoRDEtBXPY6PEOXGDgXJTvs6hUyLSyrvEVcdmq4WkYuM1_b6BAFXa0Vi1L6hs [suivant]
--2019-03-28 14:54:54-- https://developer.download.nvidia.com/compute/cuda/10.0/secure/Prod/local_installers/cuda_10.0.130_410.48_li
q5bLa-6KEN-CiNNAU_o7b6HV6EnMCIj9BUxyPLjy9DV9rdyTLrZ3FzXgZoRDEtBXPY6PEOXGDgXJTvs6hUyLSyrvEVcdmq4WkYuM1_b6BAFXa0Vi1L6hs
Résolution de developer.download.nvidia.com (developer.download.nvidia.com)... 192.229.221.58, 2606:2800:233:ef6:15dd:1ece:1d50:1e1
Connexion vers developer.download.nvidia.com (developer.download.nvidia.com)|192.229.221.58|:443...connecté.
requête HTTP transmise, en attente de la réponse...200 OK
Longueur: 2020126691 (1,9G) [application/octet-stream]
Sauvegarde en : «cuda_10.0.130_410.48_linux»

100%[=====
2019-03-28 14:55:27 (76,1 MB/s) - «cuda_10.0.130_410.48_linux» sauvegardé [2020126691/2020126691]

[tbaudier@ccwige010 cuda100]$ sh cuda_10.0.130_410.48_linux
Extraction failed.
Ensure there is enough space in /tmp and that the installation package is not corrupt
Signal caught, cleaning up
[tbaudier@ccwige010 cuda100]$ □
```

# Requirements - Cuda

- Download and install Cuda 10.0 - in an interactive GPU job

```
[tbaudier@ccwige010 ~]$ cd /pbs/throng/creatis/users/tbaudier/
[tbaudier@ccwige010 tbaudier]$ mkdir cuda100
[tbaudier@ccwige010 tbaudier]$ cd cuda100/
[tbaudier@ccwige010 cuda100]$ wget https://developer.nvidia.com/compute/cuda/10.0/Prod/local_installers/cuda_10.0.130_410.48_linux
--2019-03-28 14:54:53-- https://developer.nvidia.com/compute/cuda/10.0/Prod/local_installers/cuda_10.0.130_410.48_linux
Résolution de developer.nvidia.com (developer.nvidia.com)... 192.229.182.215
Connexion vers developer.nvidia.com (developer.nvidia.com)|192.229.182.215|:443...connecté.
requête HTTP transmise, en attente de la réponse...302 Found
Emplacement: https://developer.download.nvidia.com/compute/cuda/10.0/secure/Prod/local_installers/cuda_10.0.130_410.48_linux.run?8af_
iNNAU_o7b6HV6EnMCIj9BUxyPLjy9DV9rdyTLrZ3FzXgZoRDEtBXPY6PEOXGDgXJTvs6hUyLSyrvEVcdmq4WkYuM1_b6BAFXa0Vi1L6hs [suivant]
--2019-03-28 14:54:54-- https://developer.download.nvidia.com/compute/cuda/10.0/secure/Prod/local_installers/cuda_10.0.130_410.48_li
q5bLa-6KEN-CiNNAU_o7b6HV6EnMCIj9BUxyPLjy9DV9rdyTLrZ3FzXgZoRDEtBXPY6PEOXGDgXJTvs6hUyLSyrvEVcdmq4WkYuM1_b6BAFXa0Vi1L6hs
Résolution de developer.download.nvidia.com (developer.download.nvidia.com)... 192.229.221.58, 2606:2800:233:ef6:15dd:1ece:1d50:1e1
Connexion vers developer.download.nvidia.com (developer.download.nvidia.com)|192.229.221.58|:443...connecté.
requête HTTP transmise, en attente de la réponse...200 OK
Longueur: 2020126691 (1,9G) [application/octet-stream]
Sauvegarde en : «cuda_10.0.130_410.48_linux»

100%[=====
2019-03-28 14:55:27 (76,1 MB/s) - «cuda_10.0.130_410.48_linux» sauvegardé [2020126691/2020126691]

[tbaudier@ccwige010 cuda100]$ sh cuda_10.0.130_410.48_linux
Extraction failed.
Ensure there is enough space in /tmp and that the installation package is not corrupt
Signal caught, cleaning up
[tbaudier@ccwige010 cuda100]$ □
```

**ERROR**

# Requirements - Cuda

- Download and install Cuda 10.0 - without interactive job

```
/pbs/home/t/tbaudier(0)>cd /pbs/throng/creatis/users/tbaudier/  
/pbs/throng/creatis/users/tbaudier(0)>mkdir cuda100  
/pbs/throng/creatis/users/tbaudier(0)>cd cuda100/  
/pbs/throng/creatis/users/tbaudier/cuda100(0)>wget https://developer.nvidia.com/compute/cuda/10.0/Prod/local_installers/cuda_10.0.130_410.48_linux  
--2019-03-28 15:02:31-- https://developer.nvidia.com/compute/cuda/10.0/Prod/local_installers/cuda_10.0.130_410.48_linux  
Résolution de developer.nvidia.com (developer.nvidia.com)... 192.229.182.215  
Connexion vers developer.nvidia.com (developer.nvidia.com)|192.229.182.215|:443...connecté.  
requête HTTP transmise, en attente de la réponse...302 Found  
Emplacement: https://developer.download.nvidia.com/compute/cuda/10.0/secure/Prod/local_installers/cuda_10.0.130_410.48_linux.run?vtXso4kli0_mNlRLan  
4C3rytUAPW9mujysdRcuCW_z9GVrcdrE9-XxKhCjKn2hSRgDhnpwvY9J-WMzAlRpx1frpPVaZgHKZaiegdoG3rQqyFj4e42Am7_xtnVIk [suivant]  
--2019-03-28 15:02:32-- https://developer.download.nvidia.com/compute/cuda/10.0/secure/Prod/local_installers/cuda_10.0.130_410.48_linux.run?vtXso4  
w0X1dpkIpfSi4C3rytUAPW9mujysdRcuCW_z9GVrcdrE9-XxKhCjKn2hSRgDhnpwvY9J-WMzAlRpx1frpPVaZgHKZaiegdoG3rQqyFj4e42Am7_xtnVIk  
Résolution de developer.download.nvidia.com (developer.download.nvidia.com)... 192.229.221.58, 2606:2800:233:ef6:15dd:1ece:1d50:1e1  
Connexion vers developer.download.nvidia.com (developer.download.nvidia.com)|192.229.221.58|:443...connecté.  
requête HTTP transmise, en attente de la réponse...200 OK  
Longueur: 2020126691 (1,9G) [application/octet-stream]  
Sauvegarde en : «cuda_10.0.130_410.48_linux»  
  
100%[=====]  
  
2019-03-28 15:02:47 (143 MB/s) - «cuda_10.0.130_410.48_linux» sauvegardé [2020126691/2020126691]  
  
/pbs/throng/creatis/users/tbaudier/cuda100(0)>sh cuda_10.0.130_410.48_linux  
□
```

# Requirements - Cuda

- Download and install Cuda 10.0 - without interactive job

```
/pbs/throng/creatis/users/tbaudier/cuda100(0)>sh cuda_10.0.130_410.48_linux
Logging to /scratch/users/t/tbaudier/cuda_install_27619.log
Using less to view the EULA.
Do you accept the previously read EULA?
accept/decline/quit: accept

Install NVIDIA Accelerated Graphics Driver for Linux-x86_64 410.48?
(y)es/(n)o/(q)uit: no

Install the CUDA 10.0 Toolkit?
(y)es/(n)o/(q)uit: yes

Enter Toolkit Location
[ default is /usr/local/cuda-10.0 ]: /pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0

Do you want to install a symbolic link at /usr/local/cuda?
(y)es/(n)o/(q)uit: no

Install the CUDA 10.0 Samples?
(y)es/(n)o/(q)uit: yes

Enter CUDA Samples Location
[ default is /pbs/home/t/tbaudier ]: /pbs/throng/creatis/users/tbaudier/cuda100/

Installing the CUDA Toolkit in /pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0 ...
□
```

- Set environment variables

```
/pbs/home/t/tbaudier(0)>cat ~/.profile | grep "Cuda 10.0" -A 3
#Cuda 10.0 for pytorch & tensorflow
export PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/bin:${PATH}
export LD_LIBRARY_PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/lib64:${LD_LIBRARY_PATH}

/pbs/home/t/tbaudier(0)>□
```

# Installation

- Start interactive GPU jobs

```
/pbs/home/t/tbaudier(0)>qlogin -l GPU=1 -l GPUtype=K80 -q mc_gpu_interactive -pe multicores_gpu 4
```

- Set environment variables & start virtual environment

```
[tbaudier@ccwgige010 ~]$ . python_env.sh 3.6.7
You are using version 3.6.7 of python.
[tbaudier@ccwgige010 ~]$ export PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/bin:${PATH}
[tbaudier@ccwgige010 ~]$ export LD_LIBRARY_PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/lib64:${LD_LIBRARY_PATH}
[tbaudier@ccwgige010 ~]$ source /pbs/throng/creatis/users/tbaudier/venvAI/bin/activate
(venvAI) [tbaudier@ccwgige010 ~]$
(venvAI) [tbaudier@ccwgige010 ~]$ deactivate
[tbaudier@ccwgige010 ~]$ █
```

- Install PyTorch

```
(venvAI) [tbaudier@ccwgige010 ~]$ pip install https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl torchvision
Collecting torch==1.0.1.post2 from https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl
  Downloading https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl (636.8MB)
    100% |██████████████████████| 636.8MB 15kB/s
Could not install packages due to an EnvironmentError: [Errno 28] Aucun espace disponible sur le périphérique
```






# Installation

- Set environment variables & start virtual environment

```
/pbs/home/t/tbaudier(0)>. python_env.sh 3.6.7
You are using version 3.6.7 of python.
/pbs/home/t/tbaudier(0)>export PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/bin:${PATH}
/pbs/home/t/tbaudier(0)>export LD_LIBRARY_PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/lib64:${LD_LIBRARY_PATH}
/pbs/home/t/tbaudier(0)>source /pbs/throng/creatis/users/tbaudier/venvAI/bin/activate
(venvAI) /pbs/home/t/tbaudier(0)>
```

- Install PyTorch

 **PyTorch** <https://pytorch.org/get-started/locally/>

PyTorch Build	Stable (1.0)		Preview (Nightly)	
Your OS	Linux		Mac	Windows
Package	Conda	Pip		LibTorch Source
Language	Python 2.7	Python 3.5	Python 3.6	Python 3.7 C++
CUDA	8.0	9.0	10.0 None	

Run this Command:

```
pip3 install https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl
pip3 install torchvision
```

# Installation

- Set environment variables & start virtual environment

```

/pbs/home/t/tbaudier(0)>. python_env.sh 3.6.7
You are using version 3.6.7 of python.
/pbs/home/t/tbaudier(0)>export PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/bin:${PATH}
/pbs/home/t/tbaudier(0)>export LD_LIBRARY_PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/lib64:${LD_LIBRARY_PATH}
/pbs/home/t/tbaudier(0)>source /pbs/throng/creatis/users/tbaudier/venvAI/bin/activate
(venvAI) /pbs/home/t/tbaudier(0)>

```

- Install PyTorch

```

(venvAI) /pbs/home/t/tbaudier(0)>pip install https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl torchvision
Collecting torch==1.0.1.post2 from https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl
  Using cached https://download.pytorch.org/whl/cu100/torch-1.0.1.post2-cp36-cp36m-linux_x86_64.whl
Collecting torchvision
  Downloading https://files.pythonhosted.org/packages/fb/01/03fd7e503c16b3dc262483e5555ad40974ab5da8b9879e164b56c1f4ef6f/torchvision-0.2.2.post1-cp36-cp36m-linux_x86_64.whl
    100% |████████████████████████████████████████| 71kB 1.7MB/s
Collecting pillow>=4.1.1 (from torchvision)
  Downloading https://files.pythonhosted.org/packages/85/5e/e91792f198bbc5a0d7d3055ad552bc4062942d27eaf75c3e2783cf64eae5/Pillow-5.4.1-cp36-cp36m-linux_x86_64.whl
    100% |████████████████████████████████████████| 2.0MB 1.8MB/s
Collecting six (from torchvision)
  Downloading https://files.pythonhosted.org/packages/73/fb/00a976f728d0d1fecfe898238ce23f502a721c0ac0ecfedb80e0d88c64e9/six-1.12.0-py2.py3-noarch-anycpu.whl
Collecting numpy (from torchvision)
  Downloading https://files.pythonhosted.org/packages/35/d5/4f8410ac303e690144f0a0603c4b8fd3b986feb2749c435f7cddb288f17e/numpy-1.16.2-cp36-cp36m-linux_x86_64.whl
    100% |████████████████████████████████████████| 17.3MB 904kB/s
Installing collected packages: torch, pillow, six, numpy, torchvision
Successfully installed numpy-1.16.2 pillow-5.4.1 six-1.12.0 torch-1.0.1.post2 torchvision-0.2.2.post3
(venvAI) /pbs/home/t/tbaudier(0)>

```





# Installation

- Check installation

```
(venvAI) /pbs/home/t/tbaudier(1)>pip freeze
absl-py==0.7.1
astor==0.7.1
gast==0.2.2
grpcio==1.19.0
h5py==2.9.0
Keras-Applications==1.0.7
Keras-Preprocessing==1.0.9
Markdown==3.1
mock==2.0.0
numpy==1.16.2
pbr==5.1.3
Pillow==5.4.1
protobuf==3.7.1
six==1.12.0
tensorboard==1.13.1
tensorflow==1.13.1
tensorflow-estimator==1.13.0
tensorflow-gpu==1.13.1
termcolor==1.1.0
torch==1.0.1.post2
torchvision==0.2.2.post3
Werkzeug==0.15.1
(venvAI) /pbs/home/t/tbaudier(0)>
```

# Installation

- Check installation – in an interactive GPU job

```
/pbs/home/t/tbaudier(0)>qlogin -l GPU=1 -l GPUtype=K80 -q mc_gpu_interactive -pe multicores_gpu 4
```

```
[tbaudier@ccwgi010 ~]$ . python_env.sh 3.6.7
```

```
You are using version 3.6.7 of python.
```

```
[tbaudier@ccwgi010 ~]$ export PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/bin:${PATH}
```

```
[tbaudier@ccwgi010 ~]$ export LD_LIBRARY_PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/lib64:${LD_LIBRARY_PATH}
```

```
[tbaudier@ccwgi010 ~]$ source /pbs/throng/creatis/users/tbaudier/venvAI/bin/activate
```

```
(venvAI) [tbaudier@ccwgi010 ~]$
```

- Check PyTorch:

```
(venvAI) [tbaudier@ccwgi010 ~]$ python
```

```
Python 3.6.7 (default, Dec 5 2018, 09:50:29)
```

```
[GCC 4.8.5 20150623 (Red Hat 4.8.5-36)] on linux
```

```
Type "help", "copyright", "credits" or "license" for more information.
```

```
>>> import torch
```

```
>>> torch.cuda.is_available()
```

```
True
```

```
>>> torch.cuda.get_device_name(0)
```

```
'Tesla K80'
```

```
>>> print(torch.rand(3,3).cuda())
```

```
tensor([[0.5956, 0.5658, 0.5196],  
        [0.7455, 0.1029, 0.7484],  
        [0.0840, 0.4392, 0.6628]], device='cuda:0')
```

```
>>> □
```

# Installation

- Check TensorFlow:

```
(venvAI) [tbaudier@ccwige010 ~]$ python
Python 3.6.7 (default, Dec 5 2018, 09:50:29)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-36)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import tensorflow as tf
>>> tf.Session(config=tf.ConfigProto(log_device_placement=True))
2019-03-28 16:53:39.787518: I tensorflow/core/platform/cpu_feature_guard.cc:141] Your CPU supports instructions that this TensorFlow binary was not compiled to use: AVX2 FMA
2019-03-28 16:53:39.883653: I tensorflow/compiler/xla/service/service.cc:150] XLA service 0x43c9460 executing computations on platform CUDA. Devices:
2019-03-28 16:53:39.883699: I tensorflow/compiler/xla/service/service.cc:158]   StreamExecutor device (0): Tesla K80, Compute Capability 3.7
2019-03-28 16:53:39.885903: I tensorflow/core/platform/profile_utils/cpu_utils.cc:94] CPU Frequency: 2599805000 Hz
2019-03-28 16:53:39.886049: I tensorflow/compiler/xla/service/service.cc:150] XLA service 0x44308c0 executing computations on platform Host. Devices:
2019-03-28 16:53:39.886068: I tensorflow/compiler/xla/service/service.cc:158]   StreamExecutor device (0): <undefined>, <undefined>
2019-03-28 16:53:39.886779: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1433] Found device 0 with properties:
name: Tesla K80 major: 3 minor: 7 memoryClockRate(GHz): 0.8235
pciBusID: 0000:05:00.0
totalMemory: 11.17GiB freeMemory: 11.11GiB
2019-03-28 16:53:39.886807: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1512] Adding visible gpu devices: 0
2019-03-28 16:53:39.888456: I tensorflow/core/common_runtime/gpu/gpu_device.cc:984] Device interconnect StreamExecutor with strength 1 edge matrix:
2019-03-28 16:53:39.888471: I tensorflow/core/common_runtime/gpu/gpu_device.cc:990]   0
2019-03-28 16:53:39.888478: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1003] 0: N
2019-03-28 16:53:39.889097: I tensorflow/core/common_runtime/gpu/gpu_device.cc:1115] Created TensorFlow device (/job:localhost/replica:0/task:0/device:GPU:0 with 10805 MB memory)
: Tesla K80, pci bus id: 0000:05:00.0, compute capability: 3.7)
Device mapping:
/job:localhost/replica:0/task:0/device:XLA_GPU:0 -> device: XLA_GPU device
/job:localhost/replica:0/task:0/device:XLA_CPU:0 -> device: XLA_CPU device
/job:localhost/replica:0/task:0/device:GPU:0 -> device: 0, name: Tesla K80, pci bus id: 0000:05:00.0, compute capability: 3.7
2019-03-28 16:53:39.889718: I tensorflow/core/common_runtime/direct_session.cc:317] Device mapping:
/job:localhost/replica:0/task:0/device:XLA_GPU:0 -> device: XLA_GPU device
/job:localhost/replica:0/task:0/device:XLA_CPU:0 -> device: XLA_CPU device
/job:localhost/replica:0/task:0/device:GPU:0 -> device: 0, name: Tesla K80, pci bus id: 0000:05:00.0, compute capability: 3.7

<tensorflow.python.client.session.Session object at 0x2b41c68fdb00>
>>> □
```

# Use

- Create ~/testAI.job:

```
# Correct path in sps
cd /sps/creatis/tbaudier/testAI/

# Set correct python
. python_env.sh 3.6.7
export LC_ALL=fr_FR.utf-8
export LANG=fr_FR.utf-8

# Set Cuda 10.0 environment variables
export PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/bin:${PATH}
export LD_LIBRARY_PATH=/pbs/throng/creatis/users/tbaudier/cuda100/cuda-10.0/lib64:${LD_LIBRARY_PATH}

# Start python virtual environment
source /pbs/throng/creatis/users/tbaudier/venvAI/bin/activate

which python

# Main part
python ./testAI.py

# Close virtual environment
deactivate
```

# Use

- Create `/sps/creatis/tbaudier/testAI/testAI.py`:

```
import torch
import tensorflow as tf

#Pytorch test
print(torch.cuda.is_available())
print(torch.cuda.get_device_name(0))
print(torch.rand(3,3).cuda())

#Tensorflow test
a = tf.constant([1.0, 2.0, 3.0, 4.0, 5.0, 6.0], shape=[2, 3], name='a')
b = tf.constant([1.0, 2.0, 3.0, 4.0, 5.0, 6.0], shape=[3, 2], name='b')
c = tf.matmul(a, b)
sess = tf.Session(config=tf.ConfigProto(log_device_placement=True))
print(sess.run(c))
```

# Use

- Run GPU job:

```
/pbs/home/t/tbaudier(0)>qsub -l h_fsize=1G,s_cpu=1:00:00,h_rss=1G,GPU=1,GPUtype=V100 -pe multicores_gpu 4 testAI.job
Your job 2215873 ("testAI.job") has been submitted
```

- Output log:

```
/pbs/home/t/tbaudier(0)>cat testAI.job.o2215873
Cleared Accounted PIDs for GPU 00000000:AF:00.0.
All done.
*****
*                               *
*      Grid Engine Batch System  *
*                               *
*      IN2P3 Computing Centre, Villeurbanne FR  *
*                               *
*****
* User:          tbaudier          *
* Group:         creatis          *
* Jobname:       testAI.job       *
* JobID:         2215873          *
* Queue:         mc_gpu_medium    *
* Worker:        ccwgpg012.in2p3.fr *
* Operating system: Linux 3.10.0-957.5.1.el7.x86_64 *
* Project:       P_creatis        *
*****
* Submitted on:   jeu. mars 28 17:13:49 2019  *
* Started on:    jeu. mars 28 17:14:02 2019  *
*****

/pbs/throng/creatis/users/tbaudier/venvAI/bin/python
True
Tesla V100-PCI-E-32GB
tensor([[0.3183, 0.6610, 0.1380],
        [0.6139, 0.6171, 0.4197],
        [0.4921, 0.9900, 0.3816]], device='cuda:0')
[[22. 28.]
 [49. 64.]]
Device mapping:
/job:localhost/replica:0/task:0/device:XLA_GPU:0 -> device: XLA_GPU device
/job:localhost/replica:0/task:0/device:XLA_CPU:0 -> device: XLA_CPU device
/job:localhost/replica:0/task:0/device:GPU:0 -> device: 0, name: Tesla V100-PCI-E-32GB, pci bus id: 0000:af:00.0, compute capability: 7.0
MatMul: (MatMul): /job:localhost/replica:0/task:0/device:GPU:0
a: (Const): /job:localhost/replica:0/task:0/device:GPU:0
b: (Const): /job:localhost/replica:0/task:0/device:GPU:0
```

```
*****
* GPU requested: 1  *
*****
* TeslaV100-PCI-E-32GB: GPU-a1275a52  *
*   time:          10317ms  *
*   gpu_utilization: 1%  *
*   mem_utilization: 0%  *
*   max_memory_usage: 0MiB  *
# lines inserted: 1

*****
* Ended on:        jeu. mars 28 17:14:25 2019  *
* Exit status:     0  *
* Consumed  *
*   cpu (HS06):    00:00:00  *
*   cpu scaling factor: 17.770000  *
*   cpu time:      0 / 14400  *
*   efficiency:    00 %  *
*   io:            0.00000GB  *
*   vmem:          N/A  *
*   maxvmem:      N/A  *
*   maxrss:       N/A  *
*****
```



# Application

- Monte-Carlo simulation in medical physics
  - For dose computation in radiation therapy
  - For nuclear imaging (SPECT, PET)
- Investigate the use of Neural Network in Variance Reduction Technique (VRT) to speed up simulation
- Current project: learn Phase-Space file (PHSP) with GAN
  - PHSP =  $1e8$  particles,  $E, x, y, z, dx, dy, dz$ ,  $\sim 3$  GB
  - GAN: Generative Adversarial Network
  - Output is  $\sim 3$  MB
  - PyTorch
  - 1 job 1 GPU, 1h-5h according parameter numbers



Marginal distribution  
Courtesy of David SARRUT

