

Preparation of the project environment

For the course 1DT109 Acceleration System with Programmable Logic Components, 2022.
Department IT, Uppsala University

This documentation will walk you through

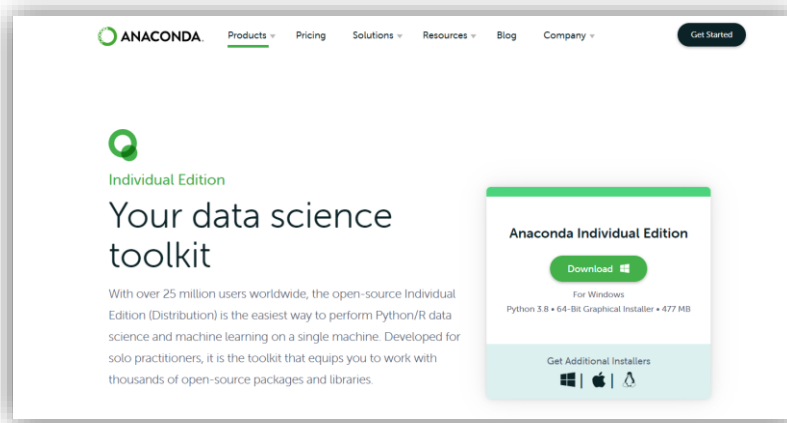
1. Installation of Anaconda, a virtual environment manager
2. Import the premade Anaconda environment file (hdr-nn.yml) into your system. This script will setup everything for you so that you do not need to maneuver yourself.
3. Run the default hdr-nn code and dump trained weights/bias for a given network.
4. Examine the default network structure

(The document is written with Windows 10 Version 21H1. However, if you work with Linux the whole setup should also work. If you indeed encounter problems, please feel free to contact with us at yuan.yao@it.uu.se; shiming.li@it.uu.se.)

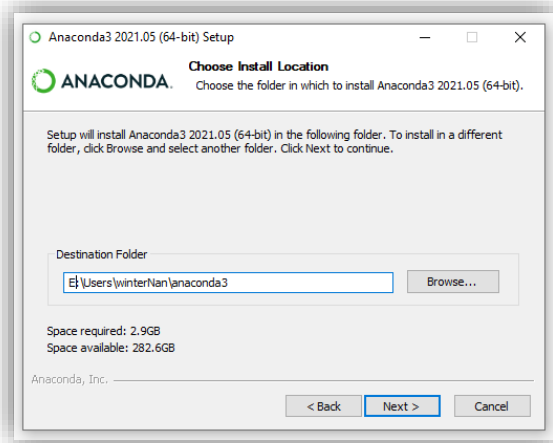
Throughout the documentation, you will find the @TODO labels. Such are concrete ideas/practices that you need to tackle with in your project.

1, Download Anaconda, the virtual environment manager

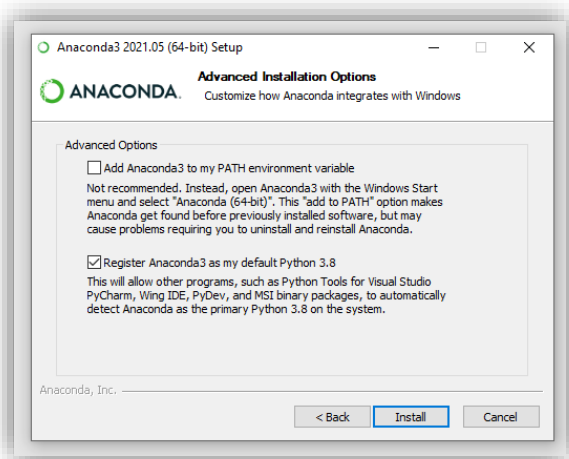
Download from the official website: <https://www.anaconda.com/products/individual>



Install it to your system. You can choose the installation path as you want. All the environment supporting files we will download later will go to this path, too.



Follow the following choices and proceed to Install.




Done!

2, Import the premade Anaconda environment file `hdr-nn.yml`

Download the zip file from studium (`hdr-nn.zip`). You will find a file named `hdr-nn.yml` inside the root. This file lists all the necessary packages/libraries that you need to run `hdr-nn`. You can open the file up using any program that can display ASCII. Take a look at it to get a overview of the packets and their versions that you will install.

Now we are ready to install our first virtual environment (`hdr-nn`) in Anaconda.

Open Windows power shell inside the root of `hdr-nn`.



Type the following command.

```
Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

Try the new cross-platform PowerShell https://aka.ms/pscore6

Loading personal and system profiles took 621ms.
(base) PS C:\Users\winterNan\Desktop\hdr-nn> conda env update -n hdr-nn --file .\hdr-nn.yml
```

Note that there is a “(base)” indicator at the very front of the prompt. This shows that we are now inside the default environment in Anaconda, which is exactly the environment of your system before Anaconda is installed. Inside the (base) environment, all the software/binary/library versions that you’ve installed before Anaconda are used.

```
Windows PowerShell
Try the new cross-platform PowerShell https://aka.ms/pscore6

Loading personal and system profiles took 621ms.
(base) PS C:\Users\winterNan\Desktop\hdr-nn> conda env update -n hdr-nn --file .\hdr-nn.yml
Collecting package metadata (repodata.json): done
Solving environment: done

==> WARNING: A newer version of conda exists. <==
  current version: 4.10.1
  latest version: 4.10.3

Please update conda by running

  $ conda update -n base -c defaults conda

Downloading and Extracting Packages
m2w64-libiconv-1.14          1.5 MB |#####| 100%
python-2.7.18                19.5 MB |#####| 100%
wincertstore-0.2            13 KB |#####| 100%
libpython-2.1               39.2 MB |#####| 100%
m2w64-make-4.1.2351        116 KB |#####| 100%
m2w64-libmangle-glibc-21  21 KB |#####| 100%
m2w64-binutils-2.25        44.3 MB |#####| 100%
m2w64-isl-0.16.1           655 KB |#####| 100%
m2w64-pkg-config-0.2       469 KB |#####| 100%
m2w64-gcc-5.3.0            41.1 MB |#####| 100%
m2w64-gcc-objc-5.3.0       15.1 MB |#####| 100%
m2w64-windows-default     3 KB |#####| 100%
m2w64-crt-git-5.0.0        3.4 MB |#####| 100%
vs2008_runtime-9.00       1017 KB |#####| 100%
ca-certificates-2021       113 KB |#####| 100%
m2w64-gcc-ada-5.3.0        33.5 MB |#####| 100%
m2w64-toolchain-5.3.0      2 KB |#####| 100%
m2w64-mpc-1.0.3            70 KB |#####| 100%
certifi-2020.6.20          155 KB |#####| 100%
sqlite-3.30.1              900 KB |#####| 100%
m2w64-bzip2-1.0.6          200 KB |#####| 100%
m2w64-zlib-1.2.8           197 KB |#####| 100%
m2w64-tools-git-5.0.0      314 KB |#####| 100%
m2w64-gcc-fortran-5.0.0    10.3 MB |#####| 100%
vc-9                        3 KB |#####| 100%
m2w64-headers-git-5.0.0    5.6 MB |#####| 100%
pip-19.3.1                 1.9 MB |#####| 100%
m2w64-winpthread-glibc-45  45 KB |#####| 100%
m2w64-mpfr-3.1.4           293 KB |#####| 100%
setuptools-44.0.0          680 KB |#####| 100%
Preparing transaction: done
Verifying transaction: done
Executing transaction: done
Installing pip dependencies: \ Ran pip subprocess with arguments:
['E:\Users\winterNan\anaconda3\envs\hdr-nn\python.exe', '-m', 'pip', 'install', '-U', '-r', 'C:\Users\winterNan\Desktop\hdr-nn\condaenv.txtxkgm.requirements.txt']
Pip subprocess output:
Collecting numpy==1.16.6
Using cached https://files.pythonhosted.org/packages/14/ef/9f2eeb4ff0c733ad9149f17266e388c308e171fdb8c2415dbb472e2bbc0f/numpy-1.16.6-cp27-cp27m-win_amd64.whl
Collecting scipy==1.2.3
Using cached https://files.pythonhosted.org/packages/91/46/3c982996646325c8c11a33956e0d944d9f1b04e02b147e70e3efee942/scipy-1.2.3-cp27-cp27m-win_amd64.whl
Collecting six==1.16.0
Using cached https://files.pythonhosted.org/packages/d9/5a/e7c31adbe875f2abbb91bd84cf2dc52d792b5a01506781dbcF25c91daf11/six-1.16.0-py2.py3-none-any.whl
Processing c:\users\winterNan\appdata\local\pip\cache\wheels\89\40\74\3a0b7d937890c66c4373120117ebf4ba99f4889b4a0a6a810c\theano-1.0.5-cp27-none-any.whl
Installing collected packages: numpy, scipy, six, theano
Successfully installed numpy-1.16.6 scipy-1.2.3 six-1.16.0 theano-1.0.5

done
#
# To activate this environment, use
#
# $ conda activate hdr-nn
#
# To deactivate an active environment, use
#
```

Import of the virtual environment is now finished. You can check what packages are installed/configured from the standard output dump (your screen). If your system has no python/whatever that is needed installed, now you have them all.

The virtual environment we’ve just imported is called “hdr-nn”, which locates inside the path: \$INSTALL/anaconda3/envs/hdr-nn. \$INSTALL if the path that you choose to where you install Anaconda.

To activate the virtual environment, type the following command:

```
Windows PowerShell
Using cached https://files.pythonhosted.org/packages/14/ef/9f2eeb4ff0c733ad9149f17266e388c308e171fdb8c2415dbb472e2bbc0f/numpy-1.16.6-cp27-cp27m-win_amd64.whl
Collecting scipy==1.2.3
Using cached https://files.pythonhosted.org/packages/91/46/3c982996646325c8c11a33956e0d944d9f1b04e02b147e70ec3efee942/scipy-1.2.3-cp27-cp27m-win_amd64.whl
Collecting six==1.16.0
Using cached https://files.pythonhosted.org/packages/d9/5a/e7c31adbe875f2abbb91bd84cf2dc52d792b5a01506781dbcF25c91daf11/six-1.16.0-py2.py3-none-any.whl
Processing c:\users\winternan\appdata\local\pip\cache\wheels\89\40\74\3a0b7d937890c66c4373120117ebf4ba99f4889b4a0a6a810c\theano-1.0.5-cp27-none-any.whl
Installing collected packages: numpy, scipy, six, theano
Successfully installed numpy-1.16.6 scipy-1.2.3 six-1.16.0 theano-1.0.5

done
#
# To activate this environment, use
#
# $ conda activate hdr-nn
#
# To deactivate an active environment, use
#
# $ conda deactivate
#

(base) PS C:\Users\winterNan\Desktop\hdr-nn> conda activate hdr-nn
(hdr-nn) PS C:\Users\winterNan\Desktop\hdr-nn>
```

If you have successfully active the virtual environment, the env indicator will change from (base) to (hdr-nn). Inside this env, you can access all the libraries with the version number specified in hdr-nn.yml. For example, type “python” and you will see python version 2.7.18:

```
Windows PowerShell
Using cached https://files.pythonhosted.org/packages/14/ef/9f2eeb4ff0c733ad9149f17266e388c308e171fdb8c2415dbb472e2bbc0f/numpy-1.16.6-cp27-cp27m-win_amd64.whl
Collecting scipy==1.2.3
Using cached https://files.pythonhosted.org/packages/91/46/3c982996646325c8c11a33956e0d944d9f1b04e02b147e70ec3efee942/scipy-1.2.3-cp27-cp27m-win_amd64.whl
Collecting six==1.16.0
Using cached https://files.pythonhosted.org/packages/d9/5a/e7c31adbe875f2abbb91bd84cf2dc52d792b5a01506781dbcF25c91daf11/six-1.16.0-py2.py3-none-any.whl
Processing c:\users\winternan\appdata\local\pip\cache\wheels\89\40\74\3a0b7d937890c66c4373120117ebf4ba99f4889b4a0a6a810c\theano-1.0.5-cp27-none-any.whl
Installing collected packages: numpy, scipy, six, theano
Successfully installed numpy-1.16.6 scipy-1.2.3 six-1.16.0 theano-1.0.5

done
#
# To activate this environment, use
#
# $ conda activate hdr-nn
#
# To deactivate an active environment, use
#
# $ conda deactivate
#

(base) PS C:\Users\winterNan\Desktop\hdr-nn> conda activate hdr-nn
(hdr-nn) PS C:\Users\winterNan\Desktop\hdr-nn> python
Python 2.7.18 |Anaconda, Inc.| (default, Apr 23 2020, 17:26:54) [MSC v.1500 64 bit (AMD64)] on win32
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

This is because the python version specified in hdr-nn.yml is indeed 2.7.18. Otherwise, if you type “python” in (base), you will see python version 3.8.8:

```
Windows PowerShell
Using cached https://files.pythonhosted.org/packages/14/ef/9f2eeb4ff0c733ad9149f17266e388c308e171fdb8c2415dbb472e2bbc0f/numpy-1.16.6-cp27-cp27m-win_amd64.whl
Collecting scipy==1.2.3
Using cached https://files.pythonhosted.org/packages/91/46/3c982996646325c8c11a33956e0d944d9f1b04e02b147e70ec3efee942/scipy-1.2.3-cp27-cp27m-win_amd64.whl
Collecting six==1.16.0
Using cached https://files.pythonhosted.org/packages/d9/5a/e7c31adbe875f2abbb91bd84cf2dc52d792b5a01506781dbcF25c91daf11/six-1.16.0-py2.py3-none-any.whl
Processing c:\users\winternan\appdata\local\pip\cache\wheels\89\40\74\3a0b7d937890c66c4373120117ebf4ba99f4889b4a0a6a810c\theano-1.0.5-cp27-none-any.whl
Installing collected packages: numpy, scipy, six, theano
Successfully installed numpy-1.16.6 scipy-1.2.3 six-1.16.0 theano-1.0.5

done
#
# To activate this environment, use
#
# $ conda activate hdr-nn
#
# To deactivate an active environment, use
#
# $ conda deactivate
#

(base) PS C:\Users\winterNan\Desktop\hdr-nn> conda activate hdr-nn
(hdr-nn) PS C:\Users\winterNan\Desktop\hdr-nn> python
Python 2.7.18 |Anaconda, Inc.| (default, Apr 23 2020, 17:26:54) [MSC v.1500 64 bit (AMD64)] on win32
Type "help", "copyright", "credits" or "license" for more information.
>>> ^Z
(hdr-nn) PS C:\Users\winterNan\Desktop\hdr-nn> conda activate base
(base) PS C:\Users\winterNan\Desktop\hdr-nn> python
Python 3.8.8 (default, Apr 13 2021, 15:08:03) [MSC v.1916 64 bit (AMD64)] :: Anaconda, Inc. on win32
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

This is because the default python version in my system is 3.8.8, before Anaconda is installed.

You can also see that you can switch to different virtual envs anytime you want by typing “conda active ENV-NAME”, where ENV-NAME is the name for the target environment.

Now everything is done with import the environment.

Important message.

Notice that all the libraries installed for env `hdr-nn` are kept inside the path `"$INSTALL/anaconda3/envs/hdr-nn"`. That is to say, Anaconda is nothing magical at all but downloads and keeps all the binaries/libraries for each virtual environment in the path `"$INSTALL/anaconda3/envs"` so that you do not need to maintain your library version manually. Note that the more virtual environments you've created, the more disk you need to keep the corresponding files.

3, Run the default `hdr-nn` code and dump trained weights/bias for a given network.

You will find the code for `hdr-nn` inside the folder you've downloaded from `sdudium`. Put it to your favorite working space/path so that you can open it with your best development tool.

In my case, I just put it to my Desktop. 😊


And I will open the project inside VScode. You can open the project up with whatever IDE you want.

Now, switch power shell to the root of `hdr-nn`, that is, `"./Desktop/hdr-nn"`. This is important since you need to run the python file there in order to load the training set correctly.


Important message.

The training images are kept as a compression filed named `"mnist.pkl.gz"` inside `"hdr-nn/data"`. The loader we use to load these images are called `"mnist_loader.py"` which locates inside `"hdr-nn/py2/src"`. `"mnist_loader.py"` will load `"mnist.pkl.gz"` from the relative path `"./data/mnist.pkl.gz"`. Thus, to successfully run the code, your CWD need to be inside `"./Desktop/hdr-nn"`.

There are three files inside the project folder `"./Desktop/hdr-nn/py2/src"`, as shown below (don't mind the `.pyc` files, they are auto-generated from the python interpreter). Let us run the code and then go through each file separately.



To run the code, type the following in the power shell (remember that you need to activate env `"hdr-nn"` beforehand).



Press `"return"` and you will get the following output from the command line. The total run time is `~3.5 min` on my machine featured with Intel i7-9700K with 16GB DDR4. The output looks like this:

```
(hdr-nn) PS C:\Users\winterNan\Desktop\hdr-nn\hdr-nn> python .\py2\src\run.py
Epoch 0: 9062 / 10000
Epoch 1: 9234 / 10000
Epoch 2: 9284 / 10000
Epoch 3: 9306 / 10000
Epoch 4: 9356 / 10000
Epoch 5: 9369 / 10000
Epoch 6: 9375 / 10000
Epoch 7: 9403 / 10000
Epoch 8: 9435 / 10000
Epoch 9: 9439 / 10000
Epoch 10: 9430 / 10000
Epoch 11: 9466 / 10000
Epoch 12: 9442 / 10000
Epoch 13: 9498 / 10000
Epoch 14: 9476 / 10000
Epoch 15: 9491 / 10000
Epoch 16: 9450 / 10000
Epoch 17: 9446 / 10000
Epoch 18: 9472 / 10000
Epoch 19: 9478 / 10000
Epoch 20: 9465 / 10000
Epoch 21: 9483 / 10000
Epoch 22: 9495 / 10000
Epoch 23: 9489 / 10000
Epoch 24: 9468 / 10000
Epoch 25: 9445 / 10000
Epoch 26: 9482 / 10000
Epoch 27: 9498 / 10000
Epoch 28: 9504 / 10000
Epoch 29: 9468 / 10000
```

By default, the network is trained with 30 epochs, with mini-batch size of 10 and learning rate of 3.0. After each epoch of training, the network is tested with another set of test images (10,000 images) and report and accuracy. At initialization, all the weights/biases of the network are randomly picked up. We can see that with the training going on, we indeed get more accuracy, which in-turn goes up and down because of various reasons. (@TODO. Can you tell why is that?)

To dump the trained weights/biases, go to network.py. Uncomment the codes on line 74-76 to output the weights/biases AFTER training to the screen. You can also uncomment the codes on line 40-42 to output the weights/biases BEFORE training to the screen. Compare then to see what are changed and unchanged.

Because the weights on the first layer are pretty of a *huge* number (784), thus, you will observe a lot of dots (...) on the output. To suppress python auto-truncation of outputting large array, uncomment code on line 19-20.

@TODO. In the project, you may want to dump the weights/biases to a file instead of outputting them to the screen.

Now let us examine the three files.

4, Examine the default network structure

4.1, run.py

Let us start with run.py.

```
# -----
# - read the input data:
import mnist_loader
training_data, validation_data, test_data = mnist_loader.load_data_wrapper()
training_data = list(training_data)
# -----
# - run the network with SGD training
import network
```

```
net = network.Network([784, 30, 10])
net.SGD(training_data, 30, 10, 3.0, test_data=test_data)
```

Run.py is short and self-explanatory. On line 4-6, we call the mnist_loader to load three image vectors, that are, training_data to train the network, validation_data to validate our training (which is ignored in our simple network), and test_data which is used to test the network after each mini-batch training. For the scope of the project, you do not need to touch these codes.

In line 11, we import from the file Network.py with the actually network.

In line 12, we instantiate a network called “net”, which serves as our baseline. The “net” is a simple structured network which contains 3 layers: an input layer, a hidden layer, and an output layer.

- The input layer contains 784 neurons, with each neuron getting input from a corresponding pixel from the 28*28 input image.
- The output layer contains 10 neurons, with each neuron output the possibility that the input image is one of the numbers of “0” ~ “9”.
- The hidden layer contains 30 neurons. Each neuron receives all outputs from all neurons in the input layer, and each neuron outputs to all neurons on the output layer. That is, our simple network is a fully connected network without a convolution layer.

You can easily add more hidden layers to the network. For example, the following code will add one additional hidden layer to the network, which contains 60 neurons.

```
net = network.Network([784, 30, 60, 10])
```

All the default values of weights/biases are initialized with random numbers, which can be found on line 35-39 in Network.py.

After the network is instantiated, we will train the network using:

```
net.SGD(training_data, 30, 10, 3.0, test_data=test_data)
```

As the code shows, the training uses stochastic gradient descent (SGD), with training_data (50,000 images) as input picture pool, and test_data (another 10,000 images) as test image pool. Those numbers are specified by the mnist training set (line 68 of mnist_loader.py), rather than parameterized in the python code.

In the above training, we train the network with 30 epochs, with each epoch consists of 10 images. The default learning rate (η) is set to 3.0 empirically. We can apply the same methodology to train our new network which contains one additional hidden layer with 60 neurons.

```
net = network.Network([784, 30, 60, 10])
net.SGD(training_data, 30, 10, 3.0, test_data=test_data)
```

Run run.py and the below are what I got from my system,

```
(hdr-nn) PS C:\Users\winterNan\Desktop\hdr-nn\hdr-nn> python .\py2\src\run.py
Epoch 0: 8872 / 10000
Epoch 1: 9201 / 10000
Epoch 2: 9318 / 10000
Epoch 3: 9399 / 10000
Epoch 4: 9442 / 10000
Epoch 5: 9395 / 10000
Epoch 6: 9397 / 10000
Epoch 7: 9502 / 10000
Epoch 8: 9513 / 10000
Epoch 9: 9531 / 10000
Epoch 10: 9481 / 10000
Epoch 11: 9528 / 10000
Epoch 12: 9504 / 10000
Epoch 13: 9537 / 10000
Epoch 14: 9543 / 10000
Epoch 15: 9557 / 10000
Epoch 16: 9555 / 10000
Epoch 17: 9549 / 10000
Epoch 18: 9548 / 10000
Epoch 19: 9559 / 10000
Epoch 20: 9529 / 10000
Epoch 21: 9551 / 10000
Epoch 22: 9560 / 10000
Epoch 23: 9527 / 10000
Epoch 24: 9534 / 10000
Epoch 25: 9567 / 10000
Epoch 26: 9560 / 10000
Epoch 27: 9578 / 10000
Epoch 28: 9566 / 10000
Epoch 29: 9563 / 10000
```

It seems that the ultimate accuracy is increased from ~94.8% to ~95.6%, which is to our expectation since we have added an additional layer. However, note that the size of the network is also increased dramatically which may render the accuracy harvest hard to achieve for FPGA.

@TODO. You can also experiment with increasing/adding more layers into the network and you will observe that the accuracy will drop instead of increasing!

4.2, network.py

The network.py implements the hdr-nn. It contains the following functions.

- `__init__`, which is the “constructor” of the network initializing all weights/biases. Note how `__init__` creates arrays for weights and biases.
@TODO. What are the dimensions of `self.weights` and `self.biases`?
- `sigmoid(z)`, computes sigmoid function of input vector `z`. Note that input `z` is a vector or Numpy array, Numpy automatically applies the function `sigmoid` elementwise, that is, in vectorized form.
@TODO. How to implement the sigmoid function in FPGA?
- `feedforward(self, a)`, given an input `a` for the network, returns the corresponding output.
@TODO. Figure out the dimensions of `w`, `a`, and `b` through out each loop iteration. Hint. You can use `print x.shape` to figure out the dimension of ndarray `x`.
- `evaluate(self, test_data)`, computes the output vector `x` for input images in `test_data`, where `y` is the reference output read from `test_data`.
@TODO. Figure out how `evaluate` calculates the output of the network.

The other functions such as `SGD`, `update_mini_batch`, etc, are used for training the network, which are not required to be implemented in the FPGA in the project. However, you are strongly recommended to read the code and understand the data flow even if you are not required to work with them in the project. **@TODO.** Understand the algorithm of the remaining functions.

4.3, mnist_loader.py

This is the loader we use to interpret test/training/invalidation images from the mnist data set. You are not required to work with them and please leave them untouched unless you know what to do.