COMP 4230
Fall 2024
Prof. Stucki

Assignment #4: TensorFlow Lab 0

Task 1: [OPTIONAL] Install Tensorflow

- You will want to choose a platform:
  - Your own laptop
  - Lab computer Windows environment or Ubuntu environment
- I recommend using VS Code, since it can be easily configured for different languages and environments. So if you don't have that installed, please download and install it first.
- You will need to have a terminal or command line (or Powershell) open in order to install Tensorflow most easily.
  - Make sure that you know the path to the most recent python installation on your platform (there may be more than one). Then type the following commands (with the path prepended to python if necessary). The dollar sign is the prompt, not part of what you type.
  - `$ python -m pip install --upgrade pip`
  - `$ python -m pip install --upgrade numpy`
  - `$ python -m pip install --user --upgrade tensorflow-cpu`
  - ```
    $ python -c "import tensorflow as tf;print(tf.reduce_sum(
      tf.random.normal([1000, 1000])))"
    ```
  - This last command should result in a Tensor object being displayed (possibly with a compiler flag warning). If so, then you have successfully installed Tensorflow.

Task 2: Explore Colab

- Follow these links to learn more about Google Colab:
  - https://colab.google/
  - https://colab.research.google.com/
  - https://colab.research.google.com/notebooks/basic_features_overview.ipynb

Task 3: Hello, World (sort of…)

- Make sure that you are logged into a Google account on the Colab site.
- Now in Colab, perform the following steps:
  - From the File menu select "New notebook in Drive"
  - From the File menu select "Locate in Drive"
  - Click the three vertical dots at the right end of the row containing Untitled0.ipynb and select "Rename" from the pop-up menu. Name the file "Lab0task3.ipynb".
  - Click the three dots again and select "Open with -> Google Colaboratory"
  - Your screen should look like this:

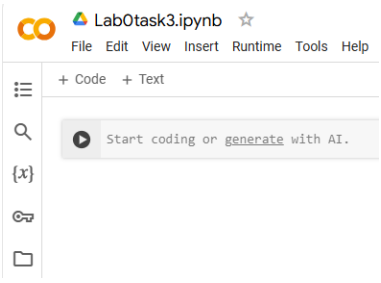
    ```

- The imports that you will need at the top of your source code are as follows:

```
import tensorflow as tf
import numpy as np
from tensorflow import keras
```

- We will start with a simple example similar to the one in Laurence Moroney's first video. We will attempt to train a network to learn the function f(n) = 3n+1. Our neural network will have one layer, a single node, and one input.

- To set up the neural network model you will need code like this:

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
```

- Now we need to compile the model, so we will have to provide the loss an optimizer functions. Like in Moroney's example we will use 'mean squared error' for the loss function and 'stochastic gradient descent' for the optimizer.

```
model.compile(optimizer='sgd', loss='mean_squared_error')
```

- Next, we need some data with which to train the model. Six input/output pairs should be enough for this simple function. We'll use Numpy arrays:

```
xs = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
ys = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

- With that we are ready to initiate training. 500 epochs should do the trick.

```
model.fit(xs, ys, epochs=500)
```

- Finally, we can ask the trained model to predict the value of the function for some input it didn't see during the training, like f(10):

```
print(model.predict(np.array([10.0])))
```

- Save and run your program. What happens? Write a brief description of the output, including the result of your print statement at the end of the program (was it correct?), as well as all the other output generated by Tensorflow (don't copy and paste all of it, just write a brief description). What do you notice as you scroll through all 500 epochs of training log?

Task 3: Big Data, here we come!

- Download and unzip the exercise archive that was posted with this assignment. Open Lab0Ex1.py in VS Code or upload and open Lab0Ex1.ipynb in Colab and review the source. It is the example from Moroney's second video.
- Run this program and briefly describe what happens. How is the training log of the epochs different from the previous one? What does the output from the loop at the end tell you?

Task 4: Science is about exploration

- Open Lab0Ex2.py in VS Code or upload and open Lab0Ex2.ipynb in Colab and review the source. It is the same problem as before but with 1024 hidden nodes rather than 128.
- Run this program and briefly describe what happens. How is the training log of the epochs different from the previous one? What is the effect of the additional nodes?

Task 5: And trial and error

- Go back to Lab0Ex1.py. Modify the number of epochs to 15.
- Run this program. What did that do to the results?
- Now change it to 30 and run it again. What do you observe?

Turn in your answers to the questions in this lab via email.