Server configuration

- CPU: Ryzen 2700X (8 cores, 16 threads, 3.7GHz)
- GPU: GTX 1070Ti 8GB x 2
- RAM: 32GB
- OS: Manjaro Linux
- Software:
    - GCC 11.1.0
    - Python 3.9.7
    - OpenSSH 8.8

Connect to the server

1. SSH to ip address `143.89.76.18` using your username and password.
    - Windows: use PuTTY or MobaXterm
    - macOS and Linux: use the command `ssh [username]@143.89.76.18`
2. If you cannot connect to the server directly (e.g., you see a "Host Not Found" error). You need to connect via HKUST VPN.

Use TensorFlow/Keras

Install packages

You can use `pip` to install the packages you want. E.g.,

```
pip install keras
```

Specify the GPU to be used

NOTE: must be executed before importing tensorflow or keras!!!

```
import os
os.environ['CUDA_VISIBLE_DEVICES']= '0'
```

'0' means using the first GPU. **If it doesn't work, then it means the first GPU is already fully occupied. Change '0' to '1' to use the second GPU.**

Limit the memory fraction

By default, TensorFlow will allocate all the memory of a GPU to one single session, which prevents multiple users sharing the same GPU.

You can make TensorFlow to allocate only a fraction of GPU memory to a session, so other users can also use the GPU.

We recommend setting the fraction to 0.4 so that each GPU can be used by 2 users simultaneously. It is not possible to run 2 sessions with fraction 0.5 because some GPU memory is occupied by the desktop environment and maybe the TensorFlow runtime.

From our experiments, as long as your model is not very complicated and your data is not very large (several GBs), the running time is nearly not affected.

```
import tensorflow as tf
config = tf.ConfigProto()
config.gpu_options.per_process_gpu_memory_fraction = 0.4

# if using keras
import keras
sess = tf.Session(config=config)
keras.backend.tensorflow_backend.set_session(sess)
# then use keras as usual

# if using session
sess = tf.Session(config=config)

# if using estimator
runConfig = tf.estimator.RunConfig(session_config=config)
estimator = tf.estimator.Estimator(
    model_fn=...,
    config=runConfig)
```