链滴

矩池云上安装 NVCaffe 教程

作者：matpool

原文链接：https://ld246.com/article/1628134152734

来源网站：链滴

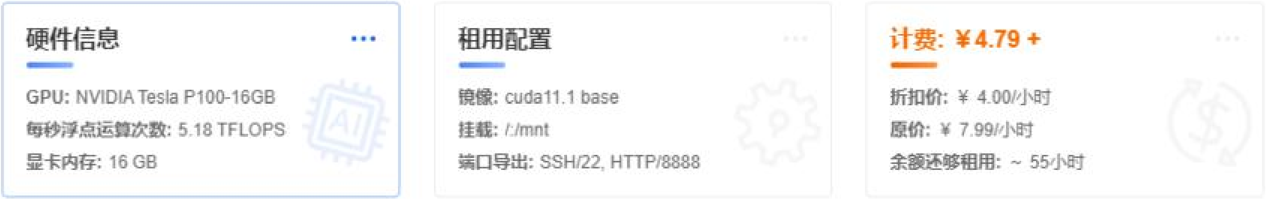

硬件信息

GPU: NVIDIA Tesla P100-16GB
每秒浮点运算次数: 5.18 TFLOPS
显卡内存: 16 GB

租用配置

镜像: cuda11.1 base
挂载: /:/mnt
端口导出: SSH/22, HTTP/8888

计费: ¥4.79 +

折扣价: ¥ 4.00/小时
原价: ¥ 7.99/小时
余额还够租用: ~ 55小时

使用的是P100，cuda11.1base镜像

创建虚拟环境

```
conda create -n py36 python=3.6
conda deactivate
conda activate py36
```

安装依赖包

```
apt update
apt-get install libopencv-dev libopenblas-dev libopenblas-base libhdf5-dev protobuf-compile libgoogle-glog-dev libgflags-dev libprotobuf-dev libboost-dev libleveldb-dev liblmdb-dev libturbojpeg0-dev  libboost-filesystem-dev libboost-system-dev libboost-thread-dev libboost-rgex-dev libsnappy-dev
```

下载NVIDIA caffe

```
cd /home/

# 官方链接wget https://github.com/NVIDIA/caffe/archive/refs/tags/v0.17.4.tar.gz
我这里用了镜像来下载
wget https://download.fastgit.org/NVIDIA/caffe/archive/refs/tags/v0.17.4.tar.gz

tar -xvf v0.17.4.tar.gz
cd caffe-0.17.4

for req in $(cat python/requirements.txt); do pip install $req; done
pip install --upgrade google-api-python-client

cp Makefile.config.example Makefile.config
```

修改Makefile.config

直接复制进去，保存即可。

```
## Refer to http://caffe.berkeleyvision.org/installation.html
# Contributions simplifying and improving our build system are welcome!

# cuDNN acceleration switch (uncomment to build with cuDNN).
# cuDNN version 6 or higher is required.
USE_CUDNN := 1

# NCCL acceleration switch (uncomment to build with NCCL)
# See https://github.com/NVIDIA/nccl
USE_NCCL := 1

# Builds tests with 16 bit float support in addition to 32 and 64 bit.
# TEST_FP16 := 1

# uncomment to disable IO dependencies and corresponding data layers
# USE_OPENCV := 0
# USE_LEVELDB := 0
# USE_LMDB := 0

# Uncomment and set accordingly if you're using OpenCV 3/4
OPENCV_VERSION := 3

# To customize your choice of compiler, uncomment and set the following.
# N.B. the default for Linux is g++ and the default for OSX is clang++
# CUSTOM_CXX := g++

# CUDA directory contains bin/ and lib/ directories that we need.
CUDA_DIR := /usr/local/cuda
# On Ubuntu 14.04, if cuda tools are installed via
# "sudo apt-get install nvidia-cuda-toolkit" then use this instead:
# CUDA_DIR := /usr

# CUDA architecture setting: going with all of them.
CUDA_ARCH :=    -gencode arch=compute_60,code=sm_60 \
           -gencode arch=compute_61,code=sm_61 \
           -gencode arch=compute_70,code=sm_70 \
           -gencode arch=compute_75,code=sm_75 \
           -gencode arch=compute_75,code=compute_75

# BLAS choice:
# atlas for ATLAS
# mkl for MKL
# open for OpenBlas - default, see https://github.com/xianyi/OpenBLAS
BLAS := open
# Custom (MKL/ATLAS/OpenBLAS) include and lib directories.
BLAS_INCLUDE := /opt/OpenBLAS/include/
BLAS_LIB := /opt/OpenBLAS/lib/

# Homebrew puts openblas in a directory that is not on the standard search path
# BLAS_INCLUDE := $(shell brew --prefix openblas)/include
# BLAS_LIB := $(shell brew --prefix openblas)/lib

# This is required only if you will compile the matlab interface.
# MATLAB directory should contain the mex binary in /bin.
# MATLAB_DIR := /usr/local
# MATLAB_DIR := /Applications/MATLAB_R2012b.app

# NOTE: this is required only if you will compile the python interface.
# We need to be able to find Python.h and numpy/arrayobject.h.
#PYTHON_INCLUDE := /usr/include/python2.7 \
#          /usr/lib/python2.7/dist-packages/numpy/core/include
# Anaconda Python distribution is quite popular. Include path:
```

```
# Verify anaconda location, sometimes it's in root.
# ANACONDA_HOME := $(HOME)/anaconda
# PYTHON_INCLUDE := $(ANACONDA_HOME)/include \
		# $(ANACONDA_HOME)/include/python2.7 \
		# $(ANACONDA_HOME)/lib/python2.7/site-packages/numpy/core/include \

# Uncomment to use Python 3 (default is Python 2)
PYTHON_LIBRARIES := boost_python3 python3.6m
PYTHON_INCLUDE := /root/miniconda3/envs/py36/include/python3.6m \
       /root/miniconda3/envs/py36/lib/python3.6/site-packages/numpy/core/include

# We need to be able to find libpythonX.X.so or .dylib.
PYTHON_LIB := /root/miniconda3/envs/py36/lib
# PYTHON_LIB := $(ANACONDA_HOME)/lib

# Homebrew installs numpy in a non standard path (keg only)
# PYTHON_INCLUDE += $(dir $(shell python -c 'import numpy.core; print(numpy.core.__file__)
))/include
# PYTHON_LIB += $(shell brew --prefix numpy)/lib

# Uncomment to support layers written in Python (will link against Python libs)
WITH_PYTHON_LAYER := 1

# Whatever else you find you need goes here.
INCLUDE_DIRS := $(PYTHON_INCLUDE) /usr/local/include /usr/include/hdf5/serial
LIBRARY_DIRS := $(PYTHON_LIB) /usr/local/lib /usr/lib /usr/lib/x86_64-linux-gnu/hdf5/serial

# If Homebrew is installed at a non standard location (for example your home directory) and
ou use it for general dependencies
# INCLUDE_DIRS += $(shell brew --prefix)/include
# LIBRARY_DIRS += $(shell brew --prefix)/lib

# Uncomment to use `pkg-config` to specify OpenCV library paths.
# (Usually not necessary -- OpenCV libraries are normally installed in one of the above $LIBR
RY_DIRS.)
# USE_PKG_CONFIG := 1

BUILD_DIR := build
DISTRIBUTE_DIR := distribute

# Uncomment for debugging. Does not work on OSX due to https://github.com/BVLC/caffe/i
sues/171
# DEBUG := 1

# The ID of the GPU that 'make runtest' will use to run unit tests.
TEST_GPUID := 0

# enable pretty build (comment to see full commands)
Q ?= @

# shared object suffix name to differentiate branches
LIBRARY_NAME_SUFFIX := -nv
```

想自己找到上面修改的路径，可以使用下面的命令查找

python -c "from distutils.sysconfig import get_python_inc; print(get_python_inc())"
python -c "import distutils.sysconfig as sysconfig; print(sysconfig.get_config_var('LIBDIR'))"

find /root/miniconda3/envs/py36/lib/ -name numpy

设置环境变量

export PYTHONPATH=/home/caffe-0.17.4/python/:$PYTHONPATH
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/root/miniconda3/envs/py36/lib

开始编译

make clean
make all -j12
make pycaffe -j12

使用python环境测试

python

import caffe
caffe.set_mode_gpu()
caffe.__version__

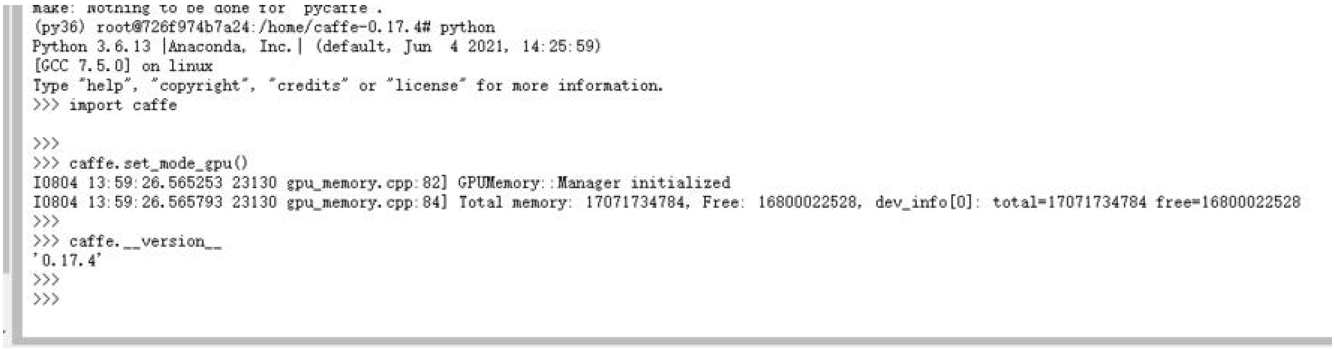

```
make: Nothing to be done for `pycaffe'.
(py36) root@726f974b7a24:/home/caffe-0.17.4# python
Python 3.6.13 |Anaconda, Inc.| (default, Jun  4 2021, 14:25:59)
[GCC 7.5.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import caffe

>>>
>>> caffe.set_mode_gpu()
I0804 13:59:26.565253 23130 gpu_memory.cpp:82] GPUMemory::Manager initialized
I0804 13:59:26.565793 23130 gpu_memory.cpp:84] Total memory: 17071734784, Free: 16800022528, dev_info[0]: total=17071734784 free=16800022528
>>>
>>> caffe.__version__
'0.17.4'
>>>
>>>
```

使用官方examples测试

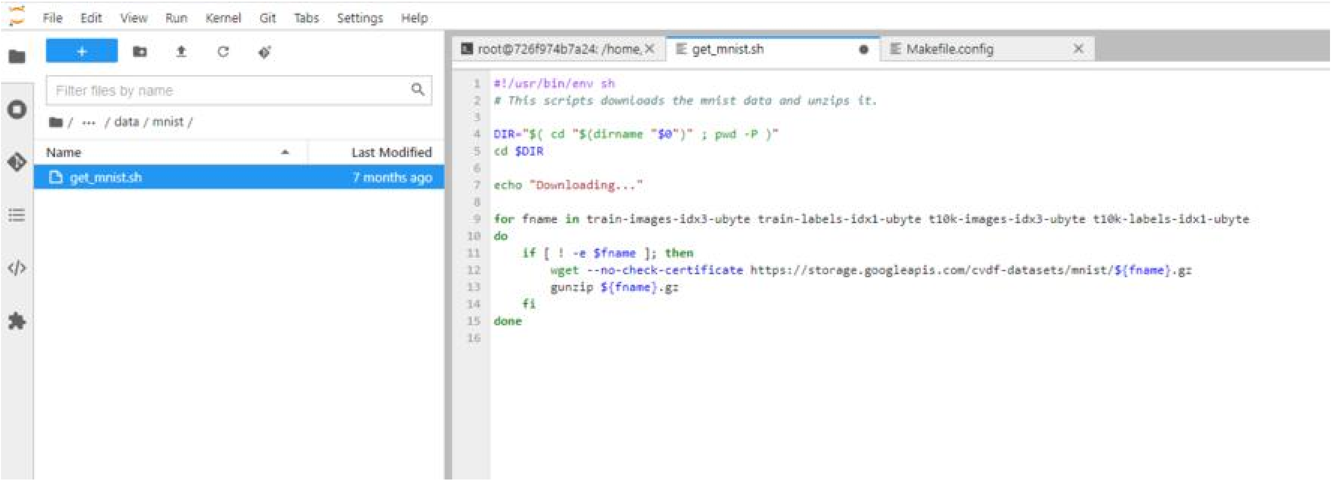

```
#!/usr/bin/env sh
# This scripts downloads the mnist data and unzips it.

DIR="$( cd "$(dirname "$0")" ; pwd -P )"
cd $DIR

echo "Downloading..."

for fname in train-images-idx3-ubyte train-labels-idx1-ubyte t10k-images-idx3-ubyte t10k-labels-idx1-ubyte
do
    if [ ! -e $fname ]; then
        wget --no-check-certificate https://storage.googleapis.com/cvdf-datasets/mnist/${fname}.gz
        gunzip ${fname}.gz
    fi
done

```

#!/usr/bin/env sh

```
# This scripts downloads the mnist data and unzips it.

DIR="$( cd "$(dirname "$0")" ; pwd -P )"
cd $DIR

echo "Downloading..."

for fname in train-images-idx3-ubyte train-labels-idx1-ubyte t10k-images-idx3-ubyte t10k-laels-idx1-ubyte
do
    if [ ! -e $fname ]; then
        wget --no-check-certificate https://storage.googleapis.com/cvdf-datasets/mnist/${fname.gz
        gunzip ${fname}.gz
    fi
done
```

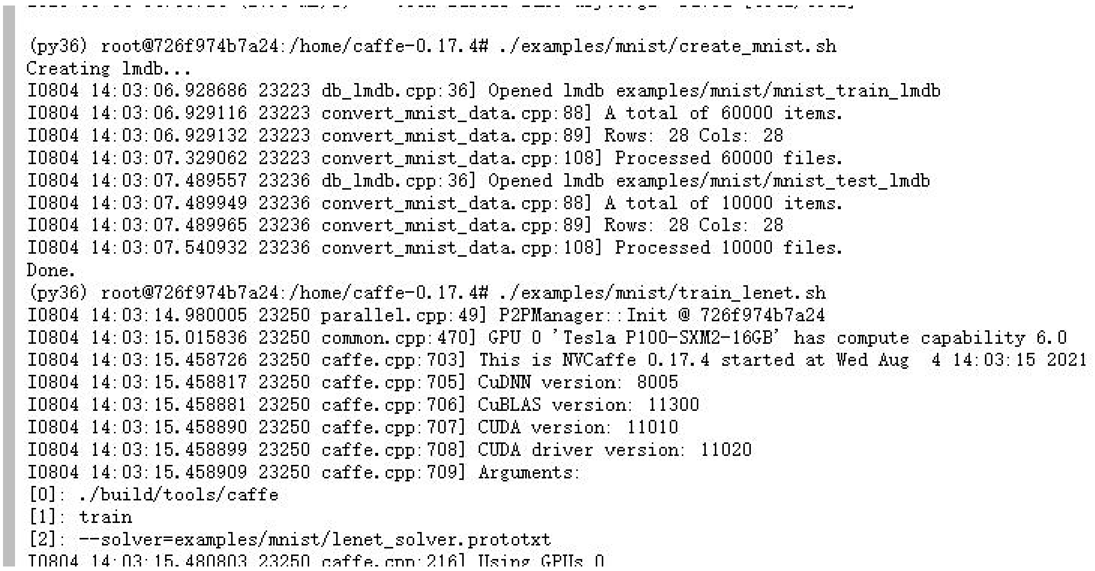

```
(py36) root@726f974b7a24:/home/caffe-0.17.4# ./examples/mnist/create_mnist.sh
Creating lmdb...
I0804 14:03:06.928686 23223 db_lmdb.cpp:36] Opened lmdb examples/mnist/mnist_train_lmdb
I0804 14:03:06.929116 23223 convert_mnist_data.cpp:88] A total of 60000 items.
I0804 14:03:06.929132 23223 convert_mnist_data.cpp:89] Rows: 28 Cols: 28
I0804 14:03:07.329062 23223 convert_mnist_data.cpp:108] Processed 60000 files.
I0804 14:03:07.489557 23236 db_lmdb.cpp:36] Opened lmdb examples/mnist/mnist_test_lmdb
I0804 14:03:07.489949 23236 convert_mnist_data.cpp:88] A total of 10000 items.
I0804 14:03:07.489965 23236 convert_mnist_data.cpp:89] Rows: 28 Cols: 28
I0804 14:03:07.540932 23236 convert_mnist_data.cpp:108] Processed 10000 files.
Done.
(py36) root@726f974b7a24:/home/caffe-0.17.4# ./examples/mnist/train_lenet.sh
I0804 14:03:14.980005 23250 parallel.cpp:49] P2PManager::Init @ 726f974b7a24
I0804 14:03:15.015836 23250 common.cpp:470] GPU 0 'Tesla P100-SXM2-16GB' has compute capability 6.0
I0804 14:03:15.458726 23250 caffe.cpp:703] This is NVCaffe 0.17.4 started at Wed Aug  4 14:03:15 2021
I0804 14:03:15.458817 23250 caffe.cpp:705] CuDNN version: 8005
I0804 14:03:15.458881 23250 caffe.cpp:706] CuBLAS version: 11300
I0804 14:03:15.458890 23250 caffe.cpp:707] CUDA version: 11010
I0804 14:03:15.458899 23250 caffe.cpp:708] CUDA driver version: 11020
I0804 14:03:15.458909 23250 caffe.cpp:709] Arguments:
[0]: ./build/tools/caffe
[1]: train
[2]: --solver=examples/mnist/lenet_solver.prototxt
I0804 14:03:15.480803 23250 caffe.cpp:216] Using GPUs 0
```

```
./data/mnist/get_mnist.sh
./examples/mnist/create_mnist.sh
./examples/mnist/train_lenet.sh
```

查看显存使用率

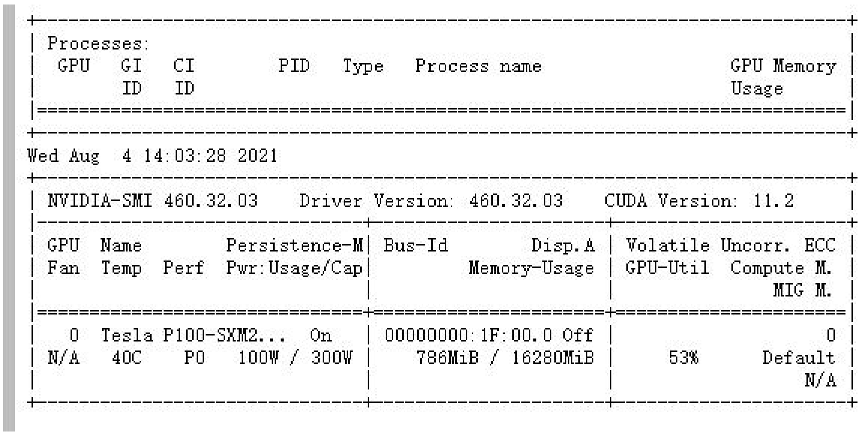

```
+-----------------------------------------------------------------------------+
| Processes:                                                                  |
|  GPU   GI   CI        PID   Type   Process name                  GPU Memory |
|        ID   ID                                                   Usage      |
|=============================================================================|
+-----------------------------------------------------------------------------+
Wed Aug  4 14:03:28 2021
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 460.32.03    Driver Version: 460.32.03    CUDA Version: 11.2     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  Tesla P100-SXM2...  On   | 00000000:1F:00.0 Off |                    0 |
| N/A   40C    P0   100W / 300W |    786MiB / 16280MiB |     53%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
```

```
nvidia-smi -l 5
```

参考文章

https://stackoverflow.com/questions/36183486/importerror-no-module-named-google

https://stackoverflow.com/questions/28190534/windows-scipy-install-no-lapack-blas-resourcs-found/29860484#29860484

https://github.com/xianyi/OpenBLAS/issues/1114

https://pypi.org/project/scipy/0.17.0/

https://github.com/NVIDIA/caffe/releases/tag/v0.17.4