

# Installing and Running Torch

## SYSTEM REQUIREMENTS FOR RUNNING TORCH ON NVIDIA GPUS

The GPU-accelerated version of Torch has the following requirements:

- Ubuntu 14.x (or any 64-bit Linux if you choose to build from source)
- NVIDIA® CUDA® 7.5 or newer (For Pascal GPUs, CUDA 8.0 or newer)
- cuDNN v5.0 or newer

You will also need an NVIDIA GPU supporting compute capability 3.0 or higher. NVIDIA Tesla® P100 and M40 are designed for machine learning workloads. We recommend P100 and M40 for servers and TitanX for PCs.

### How to Download and Install Torch

Torch is built around LuaRocks—a package manager for Lua—and has modular structure. A common collection of Torch modules is distributed under BSD open source license on [GitHub](#). We recommend using a pre-built Torch debian package (Ubuntu 14.x only).

Step 1. Add CUDA and machine learning repositories to apt-get

Get access to machine learning packages from NVIDIA by downloading and installing the cuda-repo-ubuntu1404 and nvidia-machine-learning-repo packages. Run the following commands to get access to the required repositories:

```
1 > CUDA_REPO_PKG=cuda-repo-ubuntu1404_7.5-18_amd64.deb && wget  
http://developer.download.nvidia.com/compute/cuda/repos/ubuntu1404/x86_64/$CUDA_REPO_PKG &&  
sudo dpkg -i $CUDA_REPO_PKG
```

```
2 > ML_REPO_PKG=nvidia-machine-learning-repo_4.0-2_amd64.deb && wget  
http://developer.download.nvidia.com/compute/machine-  
learning/repos/ubuntu1404/x86_64/$ML_REPO_PKG && sudo dpkg -i $ML_REPO_PKG
```

```
3 > sudo apt-get update
```

This provides access to the NVIDIA repositories containing Ubuntu packages for CUDA and ML, like cuda-toolkit-8-0, digits, caffe-nv, torch, and libcudnn5.

Step 2. Install Torch packages via apt-get

Now that you've configured access to NVIDIA ML repositories, install the Torch package and its dependencies:

```
4 > sudo apt-get install libcudnn5 libcudnn5-dev torch7-nv
```

- See more at: <http://www.nvidia.com/object/torch-library-installation.html#sthash.4PJmHekc.dpuf>