Deep Learning Frameworks

The NVIDIA Deep Learning SDK accelerates widely-used deep learning frameworks such as Caffe, CNTK, TensorFlow, Theano and Torch as well as many other deep learning applications. Choose a deep learning framework from the list below, download the supported version of cuDNN and follow the instructions on the framework page to get started.

Caffe is a deep learning framework made with expression, speed, and modularity in mind. Caffe is developed by the Berkeley Vision and Learning Center (BVLC), as well as community contributors and is popular for computer vision.

Caffe supports cuDNN v5 for GPU acceleration.

Supported interfaces: C, C++, Python, MATLAB, Command line interface

Learning Resources

- Deep learning course: Getting Started with the Caffe Framework
- Blog: Deep Learning for Computer Vision with Caffe and cuDNN

Download Caffe | Download cuDNN

The Microsoft Cognitive Toolkit —previously known as CNTK— is a unified deep-learning toolkit from Microsoft Research that makes it easy to train and combine popular model types across multiple GPUs and servers. Microsoft Cognitive Toolkit implements highly efficient CNN and RNN training for speech, image and text data.

Microsoft Cognitive Toolkit supports cuDNN v5.1 for GPU acceleration.

Supported interfaces: Python, C++, C# and Command line interface

Download CNTK | Download cuDNN

TensorFlow is a software library for numerical computation using data flow graphs, developed by Google's Machine Intelligence research organization.

TensorFlow supports cuDNN v5.1 for GPU acceleration.

Supported interfaces: C++, Python

Download TensorFlow | Download cuDNN

Theano is a math expression compiler that efficiently defines, optimizes, and evaluates mathematical expressions involving multi-dimensional arrays.

Theano supports cuDNN v5 for GPU acceleration.

Supported interfaces: Python

Learning resources

- Deep learning course: Getting Started with the Theano Framework

Download Theano | Download cuDNN

Torch is a scientific computing framework that offers wide support for machine learning algorithms.

Torch supports cuDNN v5 for GPU acceleration.

Supported interfaces: C, C++, Lua

Learning resources

- Deep learning course: Getting Started with the Torch Framework
- Blog: Understanding Natural Language with Deep Neural Networks Using Torch

Download Torch | Download cuDNN

MXnet is a deep learning framework designed for both efficiency and flexibility that allows you to mix the flavors of symbolic programming and imperative programming to maximize efficiency and productivity.

MXnet supports cuDNN v5 for GPU acceleration.

Supported Interfaces: Python, R, C++, Julia

Download MXnet | Download cuDNN

Chainer is a deep learning framework that's designed on the principle of define-by-run. Unlike frameworks that use the define-and-run approach, Chainer lets you modify networks during runtime, allowing you to use arbitrary control flow statements.

Chainer supports cuDNN v5.1 for GPU acceleration.

Supported Interfaces: Python

Download Chainer | Download cuDNN

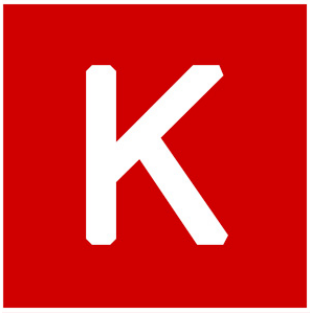

Keras is a minimalist, highly modular neural networks library, written in Python, and capable of running on top of either TensorFlow or Theano. Keras was developed with a focus on enabling fast experimentation.

cuDNN version depends on the version of TensorFlow and Theano installed with Keras.

Supported Interfaces: Python

Download Keras | Download cuDNN

More frameworks

There are several other deep learning frameworks that leverage the Deep Learning SDK, including BidMach, Brainstorm, Kaldi, MatConvNet, MaxDNN, Deeplearning4j, Keras, Lasagne(Theano), Leaf, and more.

If you're a framework developer and would like to see your framework listed here, please get in touch with us at deeplearning@nvidia.com