DALI

Quick Start Guide

TABLE OF CONTENTS

Chapter 1.
OVERVIEW

Today's deep learning applications include complex, multi-stage pre-processing data pipelines that include compute-intensive steps mainly carried out on the CPU. For instance, steps such as load data from disk, decode, crop, random resize, color and spatial augmentations and format conversions are carried out on the CPUs, limiting the performance and scalability of training and inference tasks. In addition, the deep learning frameworks today have multiple data pre-processing implementations, resulting in challenges such as portability of training and inference workflows and code maintainability.

NVIDIA® Data Loading Library™ (DALI) is a collection of highly optimized building blocks and an execution engine to accelerate input data pre-processing for deep learning applications. DALI provides both performance and flexibility of accelerating different data pipelines, as a single library, that can be easily integrated into different deep learning training and inference applications.

Key highlights of DALI include:

▸ Full data pipeline accelerated from reading disk to getting ready for training/ inference
▸ Flexibility through configurable graphs and custom operators
▸ Support for image classification and segmentation workloads
▸ Ease of integration through direct framework plugins and open source bindings
▸ Portable training workflows with multiple input formats - JPEG, raw formats, LMDB, RecordIO, TFRecord
▸ Extensible for user specific needs through open source license

Chapter 2.
PREREQUISITES

Ensure you meet the following minimum requirements:

▸ NVIDIA CUDA 9.0

▸ Deep learning frameworks
  MXNet
  Version 1.3 beta from the Python package with the following command:

```
pip install mxnet-cu90==1.3.0b20180612
```

PyTorch
Version 0.4

  ▸ If you have Python version 2.7, issue the following commands:

```
pip install
http://download.pytorch.org/whl/cu90/torch-0.4.0-cp27-cp27mu-
linux_x86_64.whl
pip install torchvision
```

  ▸ If you have Python version 3.5, issue the following commands:

```
pip3 install
http://download.pytorch.org/whl/cu90/torch-0.4.0-cp35-cp35m-
linux_x86_64.whl
pip3 install torchvision
```

TensorFlow
Version 1.8 with the following command:

```
pip install tensorflow-gpu
```

Chapter 3.
INSTALLING DALI

DALI can be installed either directly using a pre-built binary or by compiling the sources from GitHub.

3.1. Binary Installation

Install DALI using `pip`.

```
pip install --extra-index-url
https://developer.download.nvidia.com/compute/redist nvidia-dali
```

3.2. GitHub Installation

Ensure you meet the following software requirements before you install DALI from GitHub.

▸ nvJPEG
▸ protobuf version 2.0.0 or above (protobuf version 3.0.0 or above is required for TensorFlow TFRecord file format support
▸ CMake version 3.5 or above
▸ libjpeg-turbo version 1.5.x
▸ OpenCV version 3.1 or above
▸ LMDB version 0.9.x (required for datasets generated by Caffe and Caffe2)

1. Download the DALI source package from GitHub.

```
git clone --recursive https://github.com/NVIDIA/dali
cd dali
```

2. Create the build directory.

```
mkdir build
cd build
```

3. Compile DALI.

   a) If you want to build DALI without LMDB support, issue the following command:

   ```
   cmake ..
   make -j"$(nproc)" install
   pip install dali/python
   ```

   b) If you want to build DALI with LMDB support, issue the following command:

   ```
   cmake -DBUILD_LMDB=ON ..
   make -j"$(nproc)" install
   pip install dali/python
   ```

3.2.1. CMake Build Parameters

Use the following optional CMake build parameters when configuring DALI:

BUILD_PYTHON
Use this parameter to build Python bindings. The default is **ON**.

BUILD_TEST
Use this parameter to include building the test suite. The default is **ON**.

BUILD_BENCHMARK
Use this parameter to include building benchmarks. The default is **ON**.

BUILD_LMDB
Use this parameter to build with support for LMDB. The default is **OFF**.

BUILD_NVTX
Use this parameter to build with NVTX profiling enabled. The default is **OFF**.

BUILD_TENSORFLOW
Use this parameter to build the TensorFlow plugin. The default is **OFF**.

Chapter 4.
EXECUTING RESNET-50 INPUT PIPELINE

After you've installed DALI, you can run a pre-configured, ResNet-50 model accelerated by DALI, on MXNet, PyTorch, and TensorFlow frameworks for image classification training. Each of the following samples offload image loading and augmentation operations onto GPUs.

You can use Python toolchain from the command shell or Jupyter notebook to start the ResNet-50 training session.

The DALI integrated ResNet-50 Python samples are located:

▸ MXNet
▸ PyTorch
▸ TensorFlow

Chapter 5.
UNINSTALLING DALI

Uninstall DALI.

```
pip uninstall -y nvidia-dali
```

Notice

Trademarks

Copyright