
Theano-MPI Documentation

Release latest

Aug 12, 2017

Contents

1	User Guide	3
2	Evaluation	5

Theano-MPI is a python framework for distributed training of deep learning models built in Theano. It implements data-parallelism in several ways, e.g., Bulk Synchronous Parallel, [Elastic Averaging SGD](#) and [Gossip SGD](#). This project is an extension to [theano_alexnet](#), aiming to scale up the training framework to more than 8 GPUs and across nodes. Please take a look at this [technical report](#) for an overview of implementation details.

Theano-MPI is compatible for training models built in different framework libraries, e.g., [Lasagne](#), [Keras](#), [Blocks](#), as long as its model parameters can be exposed as theano shared variables. Theano-MPI also comes with a light-weight layer library for you to build customized models. Check out the examples of building [Lasagne VGGNet](#), [Wasserstein GAN](#) and [Keras Wide-ResNet](#).

CHAPTER 1

User Guide

The following pages explain how to install Theano-MPI, how to build and train a customized neural network model in Theano-MPI.

CHAPTER 2

Evaluation

This section provides some evaluation results of Theano-MPI.