
Parrots Documentation

Release 0.0.0

Parrots Team

March 11, 2015

1	Contents	3
1.1	Overview of Parrots	3

This document provides all information related to the development of *Parrots*, a next generation of deep learning platform designed specifically for Big data and distributed computing.

1.1 Overview of Parrots

Parrots is an open-source deep learning platform designed specifically to provide a series of features (see below) that make deep learning easy and performant in real world.

- High extensibility
- High performance
- High productivity
- Support of distributed computing
- Support of procedural control
- Support of data and model concealment

1.1.1 Why Parrots

The success of *deep learning* in a variety of application domains, such as speech recognition, image analysis, and natural language understanding, has made it one of the most popular techniques in machine learning. Whereas computations involved in deep learning are *conceptually* simple (*e.g.* gradient descent), implementing such computations efficiently, especially in the context with large data sets, is usually a nontrivial procedure that requires a considerable amount of engineering efforts.

In recent years, software libraries for deep learning, such as [Caffe](#) and [Torch](#), have been developed to provide user-friendly ways for training and applying deep networks. These libraries, however, encounter substantial difficulties when met with *large datasets*, *new model designs* (*e.g.* models involving new layers or recurrent networks), and *distributed computing environments*. This motivated us to launch this project and develop a new framework, with these challenges taken into account from the very beginning.

1.1.2 Main Features

Below, we explain the main features of the *Parrots* platform.

High Extensibility

High Performance

High Productivity

Distributed Computing

Procedural Control

Data and Model Concealment