
deepforge Documentation

Brian Broll

Mar 11, 2019

1	Getting Started	1
2	Quick Start	3
3	Custom Operations	5
4	Overview	9
5	Native Installation	11
6	Dockerized Installation	15
7	Command Line Interface	17
8	Configuration	19
9	Operation Feedback	21

1.1 What is DeepForge?

Deep learning is a very promising, yet complex, area of machine learning. This complexity can both create a barrier to entry for those wanting to get involved in deep learning as well as slow the development of those already comfortable in deep learning.

DeepForge is a development environment for deep learning focused on alleviating these problems. Leveraging principles from Model-Driven Engineering, DeepForge is able to reduce the complexity of using deep learning while providing an opportunity for integrating with other domain specific modeling environments created with [WebGME](#).

1.2 Design Goals

As mentioned above, DeepForge focuses on two main goals:

1. **Improving the efficiency** of experienced data scientists/researchers in deep learning
2. **Lowering the barrier to entry** for newcomers to deep learning

It is important to highlight that although one of the goals is focused on lowering the barrier to entry, DeepForge is intended to be more than simply an educational tool; that is, it is important not to compromise on flexibility and effectiveness as a research/industry tool in order to provide an easier experience for beginners (that's what forks are for!).

1.3 Overview and Features

DeepForge provides a collaborative, distributed development environment for deep learning. The development environment is a hybrid visual and textual programming environment. Higher levels of abstraction, such as creating architectures, use visual environments to capture the overall structure of the task while lower levels of abstraction, such as defining custom training functions, utilize text environments. DeepForge contains both a pipelining language

and editor for defining the high level operations to perform while training or evaluating your models as well as a language for defining neural networks (through installing a DeepForge extension such as [DeepForge-Keras](#)).

1.3.1 Concepts and Terminology

- *Operation* - essentially a function written in torch (such as *SGD*)
- *Pipeline* - directed acyclic graph composed of operations - eg, a training pipeline may retrieve and normalize data, train an architecture and return the trained model
- *Execution* - when a pipeline is run, an “execution” is created and reports the status of each operation as it is run (distributed over a number of worker machines)
- *Artifact* - an artifact represents some data (either user uploaded or created during an execution)
- *Resource* - a domain specific model (provided by a DeepForge extension) to be used by a pipeline such as a neural network architecture

CHAPTER 2

Quick Start

The easiest way to get started quickly with DeepForge is using docker-compose. First, install [docker](#) and [docker-compose](#).

Next, download the docker-compose file for DeepForge:

```
wget https://raw.githubusercontent.com/deepforge-dev/deepforge/master/docker-compose.  
→yml
```

Then start DeepForge using docker-compose:

```
docker-compose up
```

and now DeepForge can be used by opening a browser to <http://localhost:8888>!

For detailed instructions about deployment installations, check out our [deployment installation instructions](#)

Custom Operations

In this document we will outline the basics of custom operations including the operation editor and operation feedback utilities.

3.1 The Basics

Operations are used in pipelines and have named inputs and outputs. When creating a pipeline, if you don't currently find an operation for the given task, you can easily create your own by selecting the *New Operation...* operation from the add operation dialog. This will create a new operation definition and open it in the operation editor. The operation editor has two main parts, the interface editor and the implementation editor.

The interface editor is provided on the left and presents the interface as a diagram showing the input data and output data as objects flowing into or out of the given operation. Selecting the operation node in the operation interface editor will expand the node and allow the user to add or edit attributes for the given operation. These attributes are exposed when using this operation in a pipeline and can be set at design time - that is, these are set when creating the given pipeline. The interface diagram may also contain light blue nodes flowing into the operation. These nodes represent "references" that the operation accepts as input before running. When using the operation, references will appear alongside the attributes but will allow the user to select from a list of all possible targets when clicked.

On the right of the operation editor is the implementation editor. The implementation editor is a code editor specially tailored for programming the implementations of operations in DeepForge. It also is synchronized with the interface editor. A section of the implementation is shown below:

The "Train" operation uses capabilities from the `keras` package to train the neural network. This operation sets all the parameters using values provided to the operation as either attributes or references. In the implementation, attributes are provided as arguments to the constructor making the user defined attributes accessible from within the implementation. References are treated similarly to operation inputs and are also arguments to the constructor. This can be seen with the `model` constructor argument. Finally, operations return their outputs in the `execute` method; in this example, it returns a single output named `model`, that is, the trained neural network.

After defining the interface and implementation, we can now use the "Train" operation in our pipelines! An example is shown below.

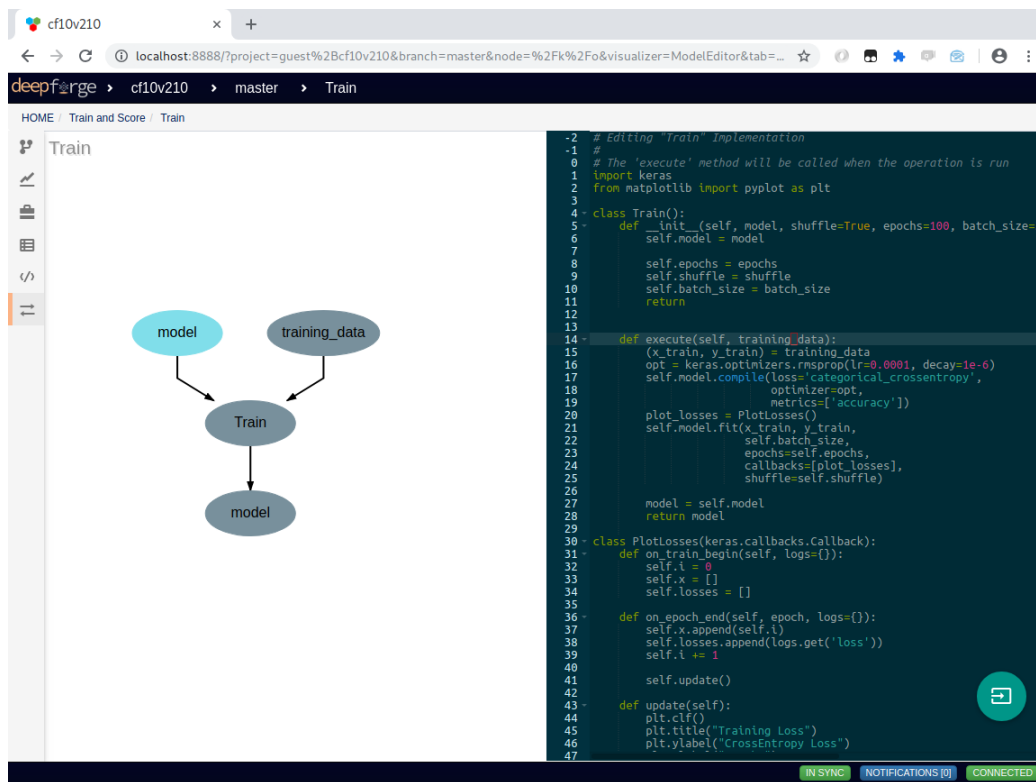


Fig. 1: Editing the “Train” operation from the “CIFAR10” example

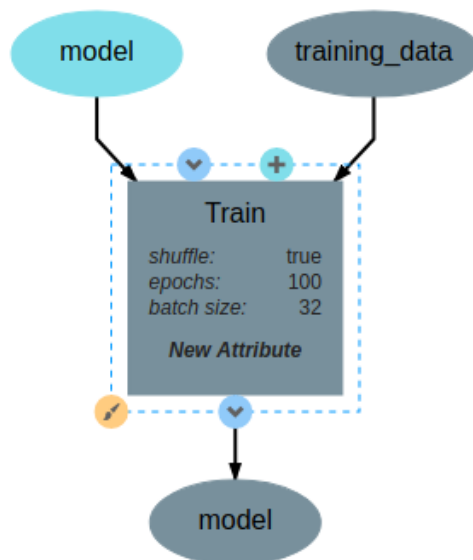


Fig. 2: The train operation accepts training data, a model and attributes for shuffling data, setting the batch size, and the number of epochs.

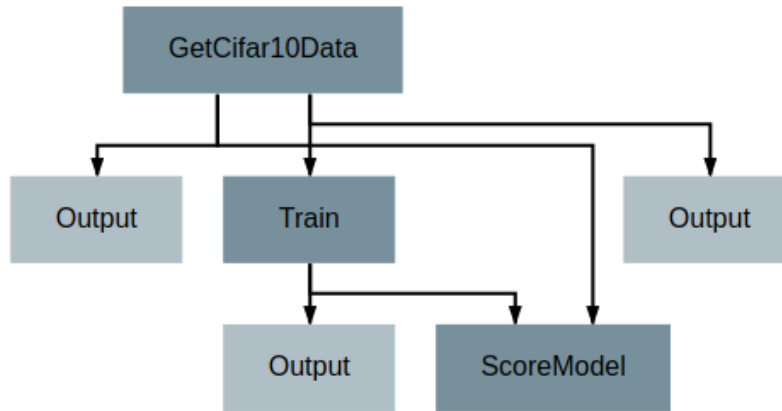


Fig. 3: Using the “Train” operation in a pipeline

3.2 Operation feedback

Operations in DeepForge can generate metadata about its execution. This metadata is generated during the execution and provided back to the user in real-time. An example of this includes providing real-time plotting feedback. When implementing an operation in DeepForge, this metadata can be created using the `matplotlib` plotting capabilities.

Detailed information about the available operation metadata types can be found in the [reference](#).




Fig. 4: An example graph of the loss function while training a neural network

4.1 DeepForge Component Overview

DeepForge is composed of four main elements:

- *Server*: Main component hosting all the project information and is connected to by the clients.
- *Database*: MongoDB database containing DeepForge, job queue for the workers, etc.
- *Worker*: Slave machine performing the actual machine learning computation.
- *Client*: The connected browsers working on DeepForge projects.

Of course, only the *Server*, *Database* (MongoDB) and *Worker* need to be installed. If you are not going to execute any machine learning pipelines, installing the *Worker* can be skipped.

4.2 Component Dependencies

The following dependencies are required for each component:

- *Server* (NodeJS v8.11.3)
- *Database* (MongoDB v3.0.7)
- *Worker*: NodeJS v8.11.3 (used for job management logic) and Python 3. If you are using the deepforge-keras extension, you will also need Keras and [TensorFlow](#) installed.
- *Client*: We recommend using Google Chrome and are not supporting other browsers (for now). In other words, other browsers can be used at your own risk.

4.3 Configuration

After installing DeepForge, it can be helpful to check out [configuring DeepForge](#)

Native Installation

5.1 Dependencies

First, install [NodeJS](#) (LTS) and [MongoDB](#). You may also need to install git if you haven't already.

Next, you can install DeepForge using npm:

```
npm install -g deepforge
```

Now, you can check that it installed correctly:

```
deepforge --version
```

After installing DeepForge, it is recommended to install the [deepforge-keras](#) extension which provides capabilities for modeling neural network architectures:

```
deepforge extensions add deepforge-dev/deepforge-keras
```

DeepForge can now be started with:

```
deepforge start
```

5.1.1 Database

Download and install MongoDB from the [website](#). If you are planning on running MongoDB locally on the same machine as DeepForge, simply start *mongod* and continue to setting up DeepForge.

If you are planning on running MongoDB remotely, set the environment variable “MONGO_URI” to the URI of the Mongo instance that DeepForge will be using:

```
MONGO_URI="mongodb://pathToMyMongo.com:27017/myCollection" deepforge start
```

5.1.2 Server

The DeepForge server is included with the deepforge cli and can be started simply with

```
deepforge start --server
```

By default, DeepForge will start on *http://localhost:8888*. However, the port can be specified with the *--port* option. For example:

```
deepforge start --server --port 3000
```

5.1.3 Worker

The DeepForge worker can be started with

```
deepforge start --worker
```

To connect to a remote deepforge instance, add the url of the DeepForge server:

```
deepforge start --worker http://myaddress.com:1234
```

5.1.4 Updating

DeepForge can be updated with the command line interface rather simply:

```
deepforge update
```

```
deepforge update --server
```

For more update options, check out *deepforge update --help*!

5.2 Manual Installation (Development)

Installing DeepForge for development is essentially cloning the repository and then using *npm* (node package manager) to run the various start, test, etc, commands (including starting the individual components). The deepforge cli can still be used but must be referenced from *./bin/deepforge*. That is, *deepforge start* becomes *./bin/deepforge start* (from the project root).

5.2.1 DeepForge Server

First, clone the repository:

```
git clone https://github.com/dfst/deepforge.git
```

Then install the project dependencies:

```
npm install
```

To run all components locally start with


```
./bin/deepforge start
```

and navigate to <http://localhost:8888> to start using DeepForge!

Alternatively, if jobs are going to be executed on an external worker, run `./bin/deepforge start -s` locally and navigate to <http://localhost:8888>.

5.2.2 DeepForge Worker

If you are using `./bin/deepforge start -s` you will need to set up a DeepForge worker (`./bin/deepforge start` starts a local worker for you!). DeepForge workers are slave machines connected to DeepForge which execute the provided jobs. This allows the jobs to access the GPU, etc, and provides a number of benefits over trying to perform deep learning tasks in the browser.

Once DeepForge is installed on the worker, start it with

```
./bin/deepforge start -w
```

Note: If you are running the worker on a different machine, put the address of the DeepForge server as an argument to the command. For example:

```
./bin/deepforge start -w http://myaddress.com:1234
```

5.2.3 Updating

Updating can be done the same as any other git project; that is, by running `git pull` from the project root. Sometimes, the dependencies need to be updated so it is recommended to run `npm install` following `git pull`.

CHAPTER 6

Dockerized Installation

Each of the components are also available as docker containers. This page outlines the running of each of the main components as docker containers and connecting them as necessary.

6.1 Database

First, you can start the mongo container using:

```
docker run -d -v /abs/path/to/data:/data/db mongo
```

where `/abs/path/to/data` is the path to the mongo data location on the host. If running the database in a container, you will need to get the ip address of the given container:

```
docker inspect <container id> | grep IPAddr
```

The `<container id>` is the value returned from the original `docker run` command.

When running mongo in a docker container, it is important to mount an external volume (using the `-v` flag) to be used for the actual data (otherwise the data will be lost when the container is stopped).

6.2 Server

The DeepForge server can be started with

```
docker run -d -v $HOME/.deepforge/blob:/data/blob \
-p 8888:8888 -e MONGO_URI=mongodb://172.17.0.2:27017/deepforge \
deepforge/server
```

where `172.17.0.2` is the ip address of the mongo container and `$HOME/.deepforge/blob` is the path to use for binary DeepForge data on the host. Of course, if the mongo instance is locating at a different location, `MONGO_URI` can be set to this address as well. Also, the first port (8888) can be replaced with the desired port to expose on the host.

6.3 Worker

As workers may require GPU access, they will need to use the nvidia-docker plugin. Workers can be created using

```
nvidia-docker run -d deepforge/worker http://172.17.0.1:8888
```

where `http://172.17.0.1:8888` is the location of the DeepForge server to which to connect.

Note: The `deepforge/worker` image is packaged with `cuda 7.5`. Depending upon your hardware and `nvidia` version, you may need to build your own docker image or run the worker natively.

Command Line Interface

This document outlines the functionality of the `deepforge` command line interface (provided after installing `deepforge` with `npm install -g deepforge`).

- Installation Configuration
- Starting DeepForge or Components
- Update or Uninstall DeepForge
- Managing Extensions

7.1 Installation Configuration

Installation configuration can be edited using the `deepforge config` command as shown in the following examples:

Printing all the configuration settings:

```
deepforge config
```

Printing the value of a configuration setting:

```
deepforge config worker.dir
```

Setting a configuration option, such as `worker.dir` can be done with:

```
deepforge config worker.dir /some/new/directory
```

For more information about the configuration settings, check out the [configuration](#) page.

7.2 Starting DeepForge Components

DeepForge components, such as the server or the workers, can be started with the `deepforge start` command. By default, this command will start all the necessary components to run including the server, a mongo database (if applicable) and a worker.

The server can be started by itself using

```
deepforge start --server
```

The worker can be started by itself using

```
deepforge start --worker http://154.95.87.1:7543
```

where *http://154.95.87.1:7543* is the url of the deepforge server.

7.3 Update/Uninstall DeepForge

DeepForge can be updated or uninstalled using

```
deepforge update
```

DeepForge can be uninstalled using `deepforge uninstall`

7.4 Managing Extensions

DeepForge extensions can be installed and removed using the `deepforge extensions` subcommand. Extensions can be added, removed and listed as shown below

```
deepforge extensions add https://github.com/example/some-extension
deepforge extensions remove some-extension
deepforge extensions list
```

Configuration of deepforge is done through the *deepforge config* command from the command line interface. To see all config options, simply run *deepforge config* with no additional arguments. This will print a JSON representation of the configuration settings similar to:

```
Current config:
{
  "blob": {
    "dir": "/home/irishninja/.deepforge/blob"
  },
  "worker": {
    "cache": {
      "useBlob": true,
      "dir": "~/.deepforge/worker/cache"
    },
    "dir": "~/.deepforge/worker"
  },
  "mongo": {
    "dir": "~/.deepforge/data"
  }
}
```

Setting an attribute, say *worker.cache.dir*, is done as follows

```
deepforge config worker.cache.dir /tmp
```

8.1 Environment Variables

Most settings have a corresponding environment variable which can be used to override the value set in the cli's configuration. This allows the values to be temporarily set for a single run. For example, starting a worker with a different cache than set in *worker.cache.dir* can be done with:

```
DEEPFORGE_WORKER_CACHE=/tmp deepforge start -w
```

The complete list of the environment variable overrides for the configuration options can be found [here](#).

8.2 Settings

8.2.1 blob.dir

The path to the blob (large file storage containing models, datasets, etc) to be used by the deepforge server.

This can be overridden with the *DEEPFORGE_BLOB_DIR* environment variable.

8.2.2 worker.dir

The path to the directory used for worker executions. The workers will run the executions from this directory.

This can be overridden with the *DEEPFORGE_WORKER_DIR* environment variable.

8.2.3 mongo.dir

The path to use for the *-dbpath* option of mongo if starting mongo using the command line interface. That is, if the *MONGO_URI* is set to a local uri and the cli is starting the deepforge server, the cli will check to verify that an instance of mongo is running locally. If not, it will start it on the given port and use this setting for the *-dbpath* setting of mongod.

8.2.4 worker.cache.dir

The path to the worker cache directory.

This can be overridden with the *DEEPFORGE_WORKER_CACHE* environment variable.

8.2.5 worker.cache.useBlob

When running the worker on the same machine as the server, this allows the worker to use the blob as a cache and simply create symbolic links to the data (eg, training data, models) to prevent having to even perform a copy of the data on the given machine.

This can be overridden with the *DEEPFORGE_WORKER_USE_BLOB* environment variable.

Operations can provide real-time graph feedback. In the future, this will be extended to other forms of feedback such as image feedback.

9.1 Graphs

Real-time graphs can be created using *matplotlib*:

```
import matplotlib.pyplot as plt
plt.plot([1, 4, 9, 16, 25])

plt.draw()
plt.show()
```

The graph can also be configured in regular matplotlib fashion:

```
plt.title('My First Plot')
plt.xlabel('This is the x-axis')
plt.ylabel('This is the y-axis')

plt.draw()
plt.show()
```