CaffeOnSpark: Deep Learning on Spark Cluster

Andy Feng, Jun Shi and Mridul Jain
Yahoo! Inc.

SPARK SUMMIT 2016
DATA SCIENCE AND ENGINEERING AT SCALE
JUNE 6-8, 2016 SAN FRANCISCO

Agenda

- Why Deep Learning on Spark?
- CaffeOnSpark
  - Architecture
  - API: Scala + Python
- Demo
  - CaffeOnSpark on Python Notebook

Deep Learning

Handwritten digits (MNIST)

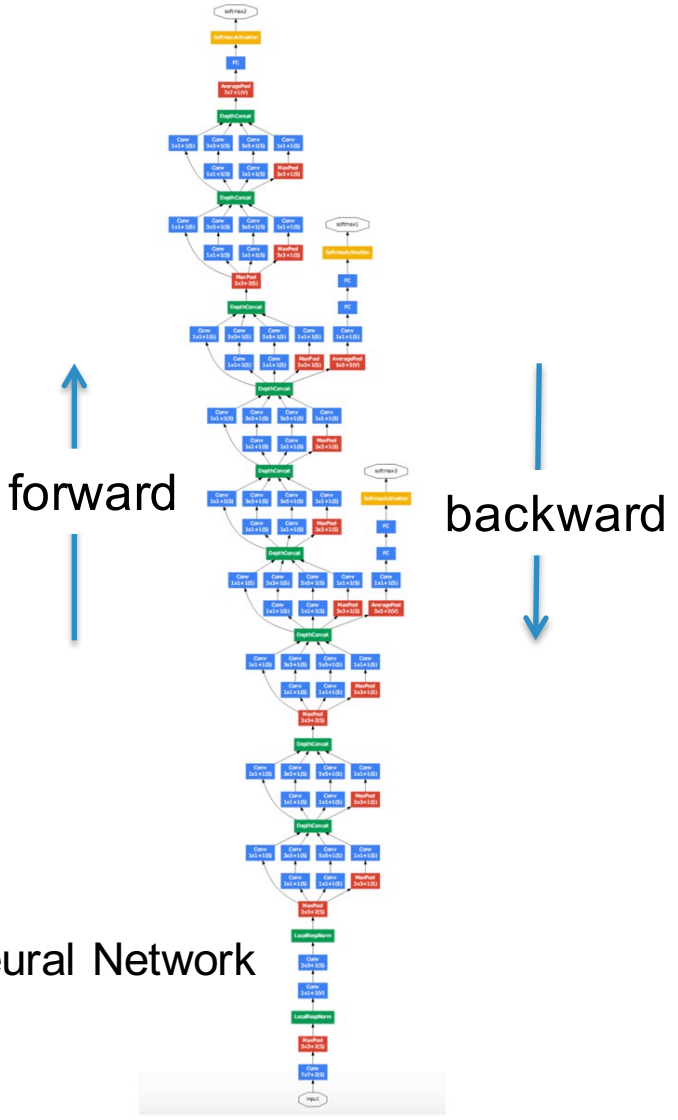

Deep Neural Network

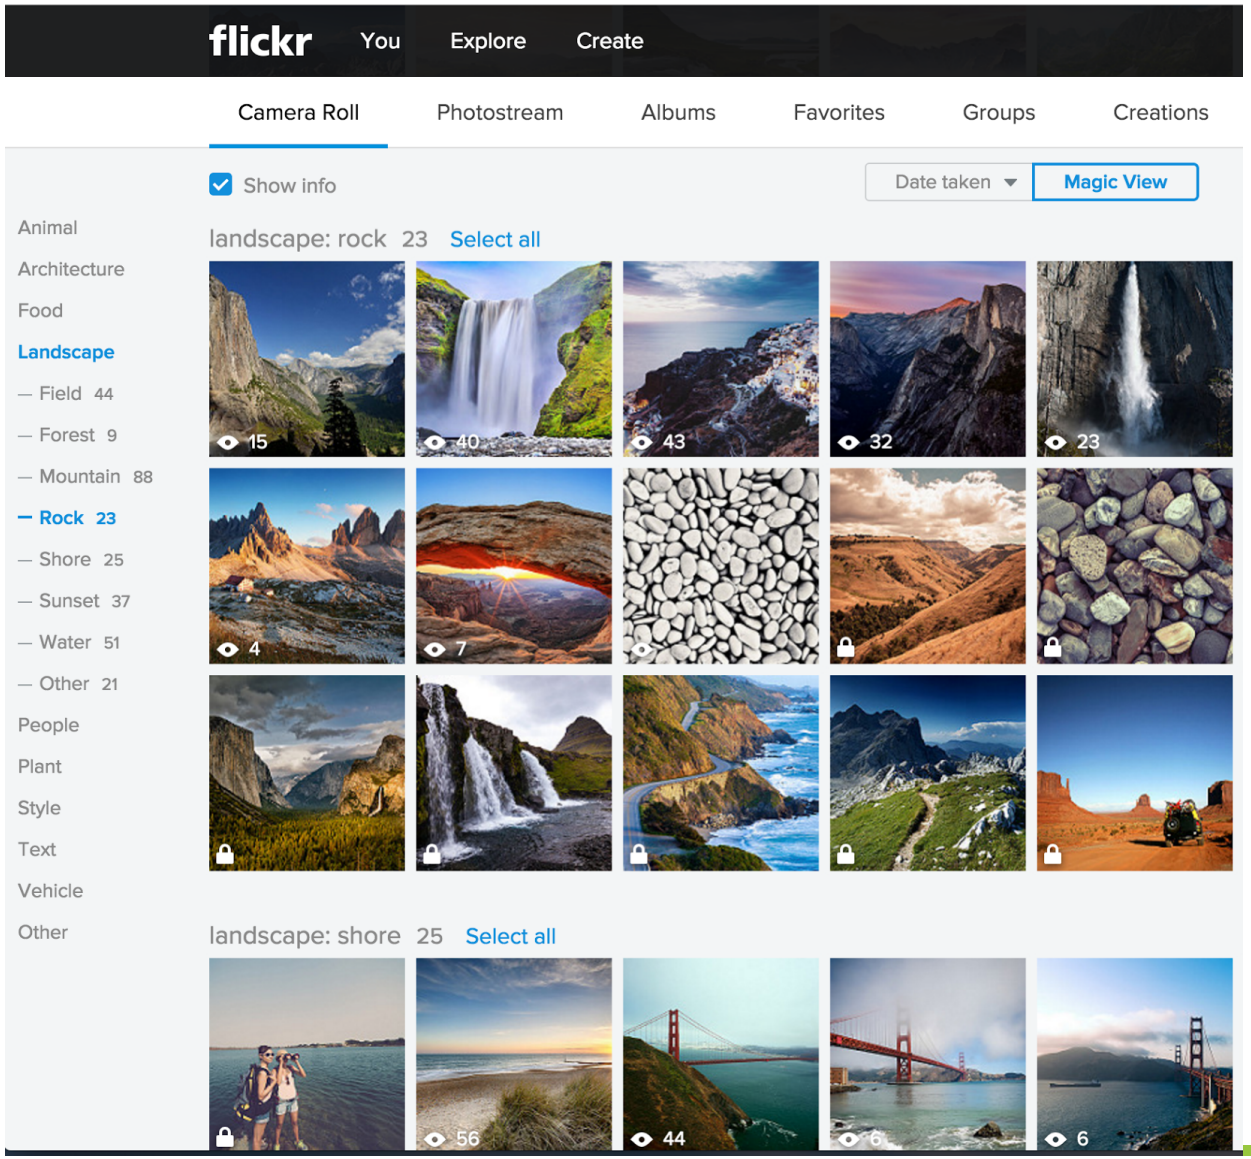

Flickr Magic View:

https://flickr.com/cameraroll

- Photos organized according to 70 categories
- Empowered by deep learning & machine learning

Flickr DL/ML Pipeline

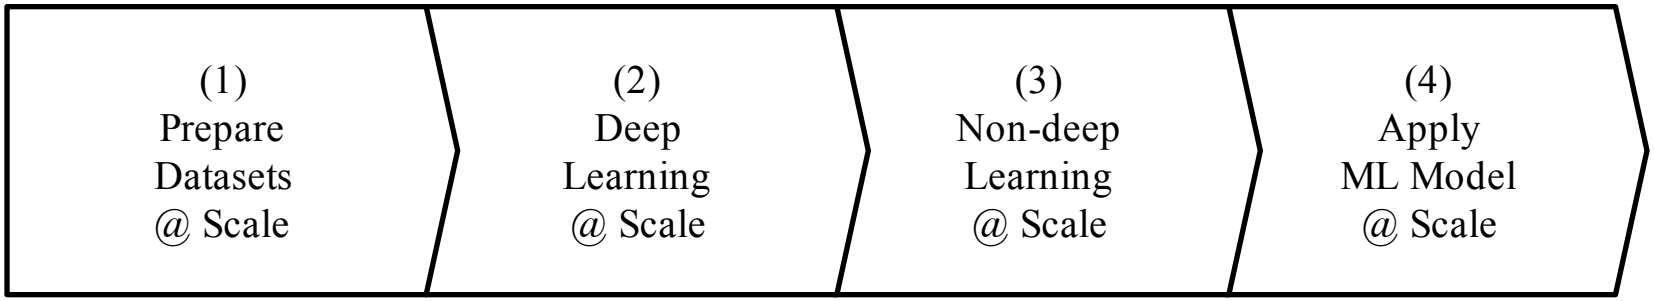

hadoop Spark STORM

* http://bit.ly/1KIDfof by Pierre Garrigues, Deep Learning Summit 2015

Deep Learning vs. Spark

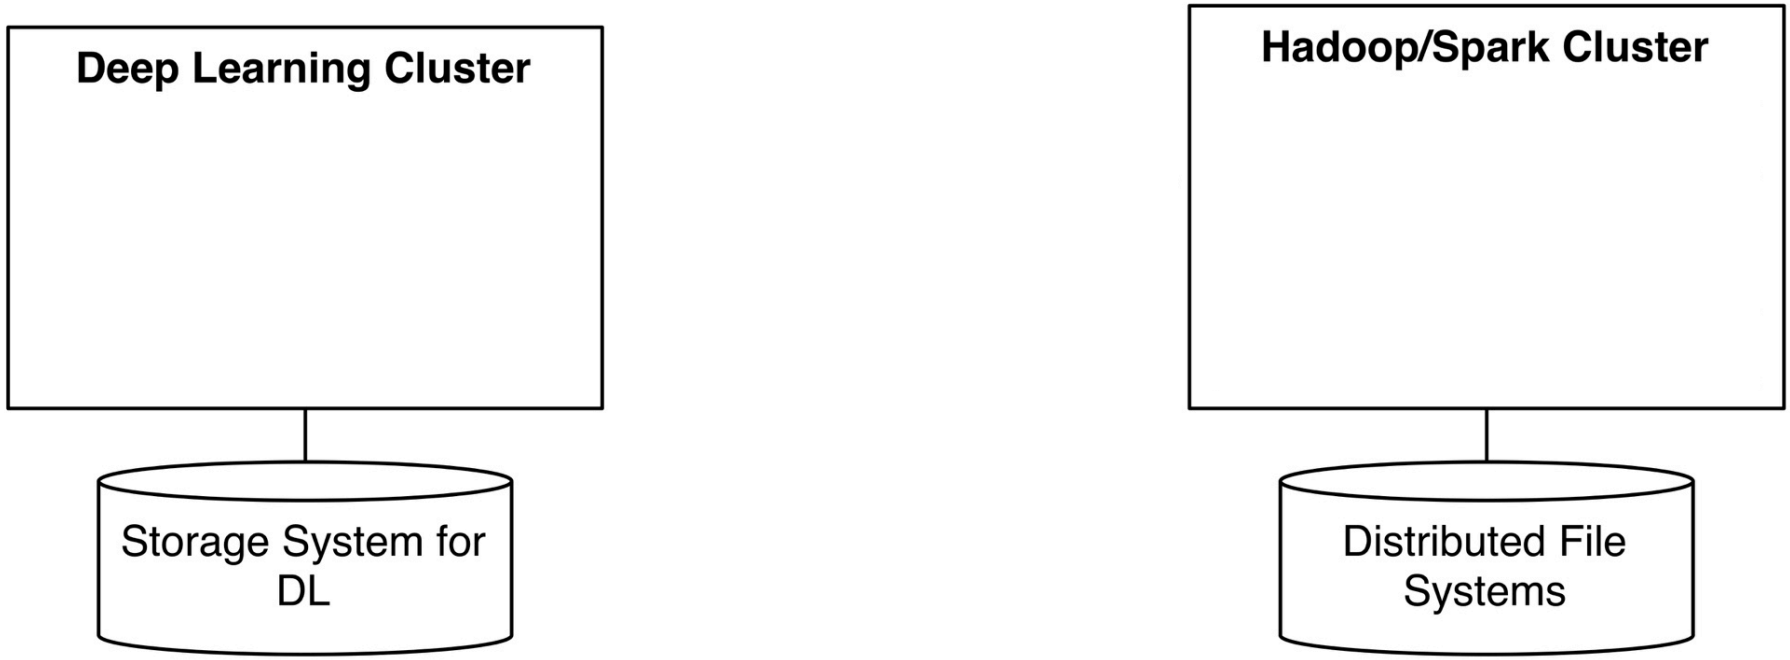

Deep Learning Frameworks

- Theano
- Torch
- Caffe
  - Popular choice for vision community
  - Widely used in Yahoo
- TensorFlow
- …

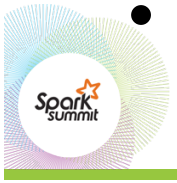

Deep Learning on Spark

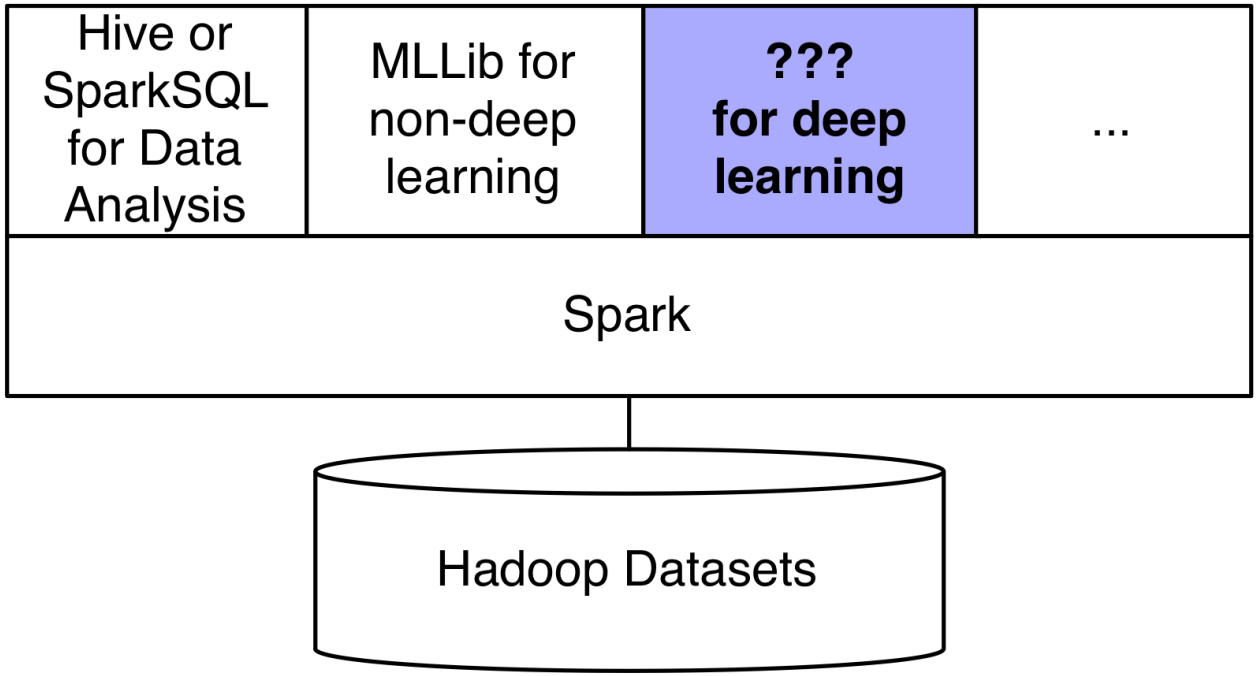

Related Work: SparkNet & DL4J

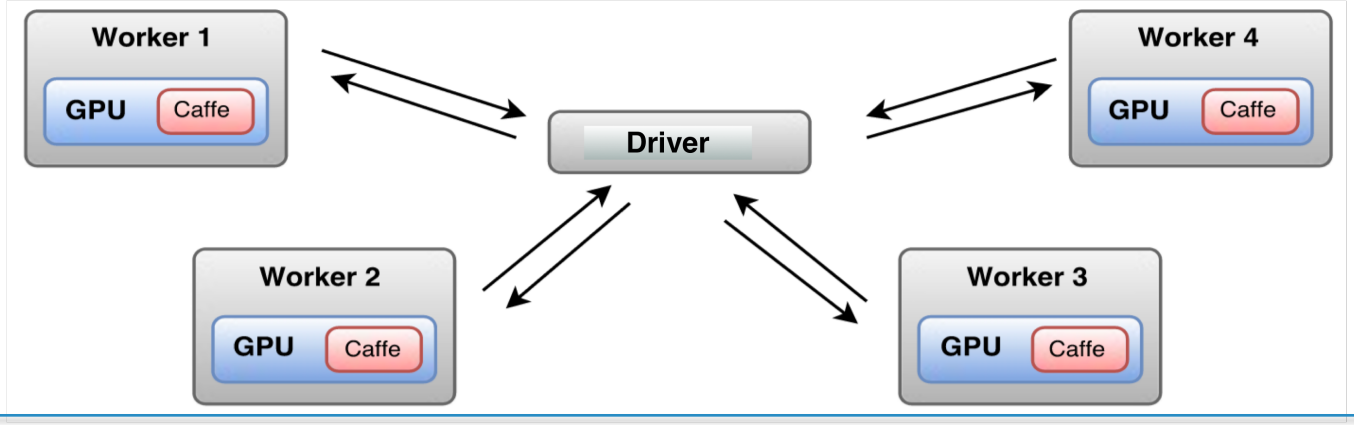

REPEAT

1) [driver] sc.broadcast(model) to executors
2) [executor] apply DL training against a mini-batch of dataset to update models locally
3) [driver] aggregate(models) to produce a new model

CaffeOnSpark Open Sourced

github.com/yahoo/CaffeOnSpark

- Apache 2.0 license
- Distributed deep learning
  - GPU or CPU
  - Ethernet or InfiniBand
- Easily deployed on public cloud or private cloud

CaffeOnSpark: Scalable Architecture

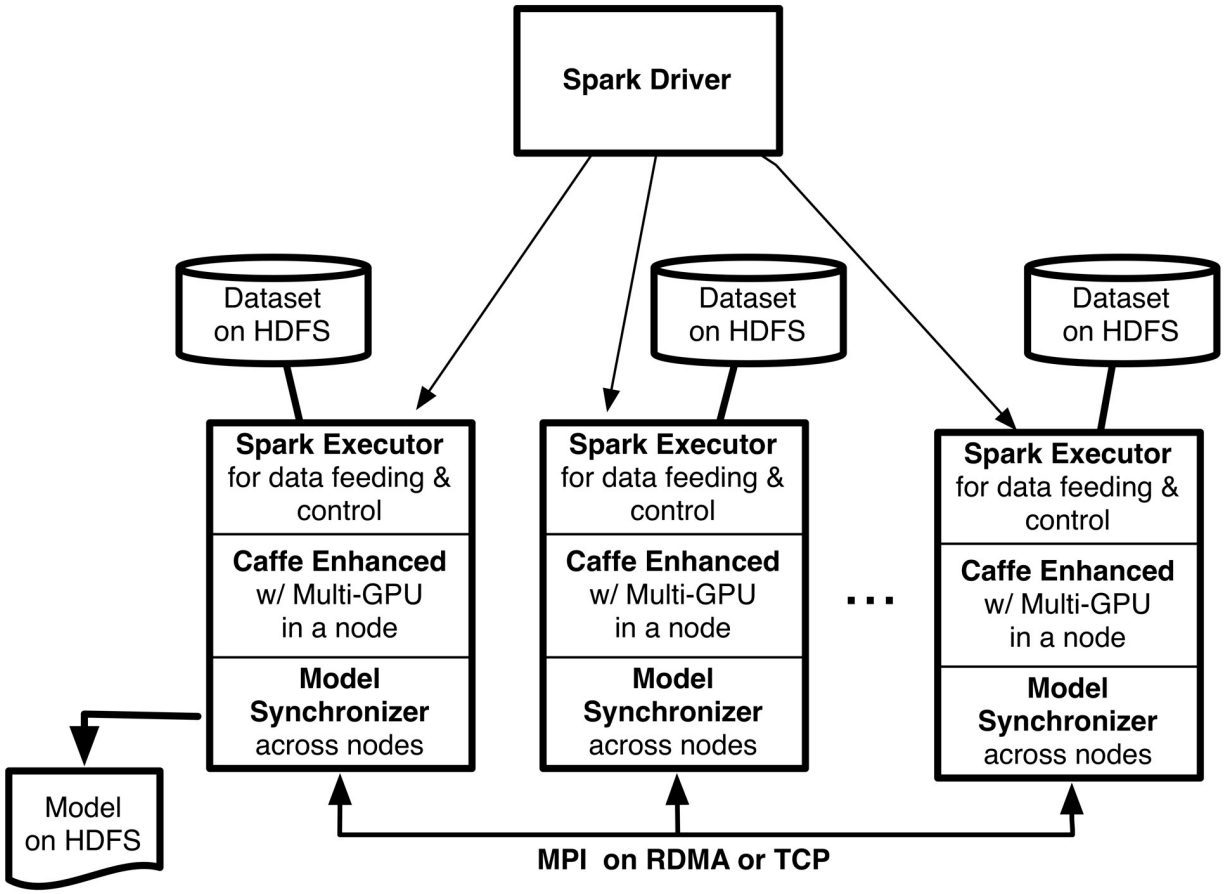

CaffeOnSpark: Deployment Options

- Single node
  - Spark-submit –master local
- Multiple nodes w/ ethernet connection
  - Spark-submit –master URL –connection ethernet
  - Ex. EC2
- Multiple nodes w/ Infiniband connection
  - Spark-submit –master URL –connection infiniband
  - Ex., Yahoo Hadoop cluster

Deep Learning: 19x Speedup (est.)

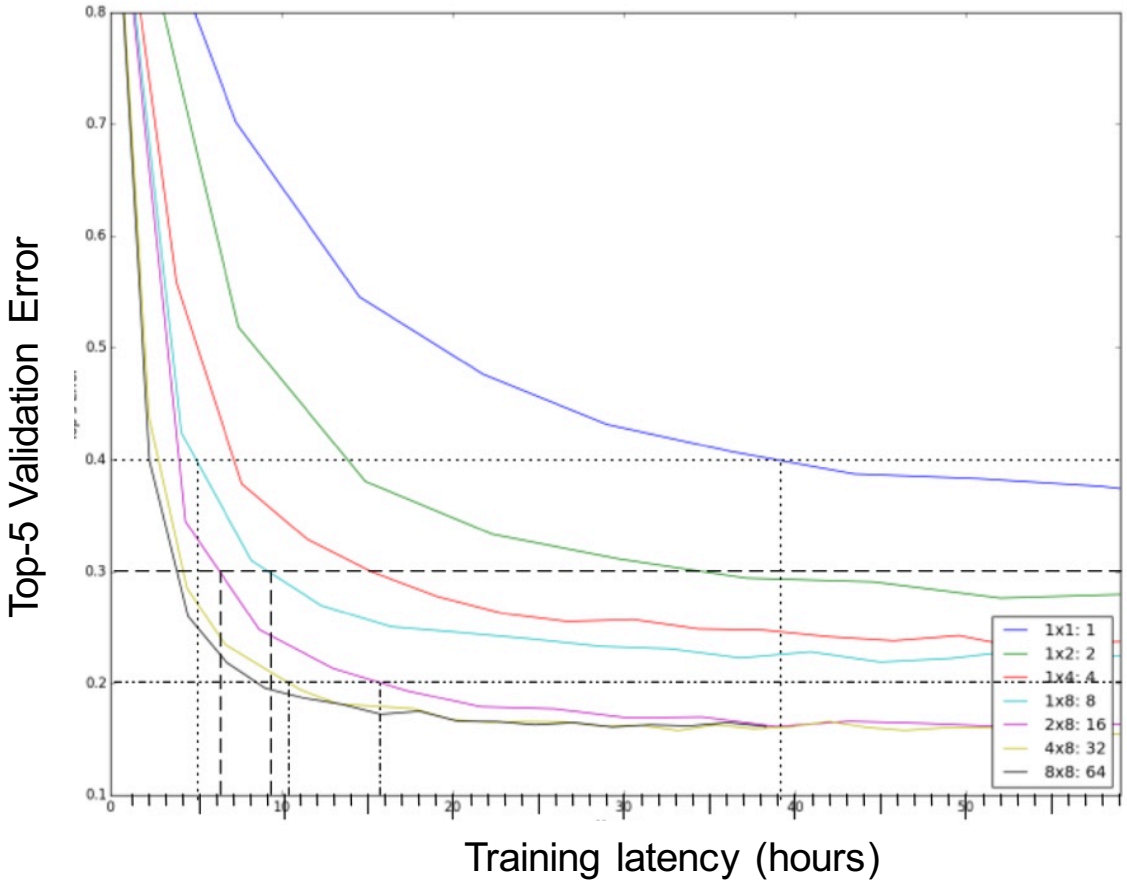

CaffeOnSpark: DL Made Easy

Spark CLI

- spark-submit
  --num-executors *#_Processes*
  --class com.yahoo.ml.CaffeOnSpark
  caffe-on-spark.jar
  -devices *#_gpus_per_proc*
  -conf *solver_config_file*
  -model model_file
  -train | *-test* | *-feature*

Caffe Configuration

```
layer {
  name: "data"
  type: "MemoryData"
  source_class="com.yahoo.ml.caffe.LMDB"
  memory_data_param {
    source: "hdfs:///mnist/trainingdata/"
    batch_size: 64;
    channels: 1;
    height: 28;
    width: 28;
  }
  …
}
```

CaffeOnSpark: One Program (Scala)

http://bit.ly/21ZY1c2

```
cos = new CaffeOnSpark(ctx)  conf = new Config(ctx, args).init()

// (1) training DL model
dl_train_source = DataSource.getSource(conf, true)  cos.train(dl_train_source)
// (2) extract features via DL
lr_raw_source = DataSource.getSource(conf, false)  ext_df = cos.features(lr_raw_s
// (3) apply ML
lr_input=ext_df.withColumn("L", cos.floats2doubleUDF(ext_df(conf.label)))
.withColumn("F", cos.floats2doublesUDF(ext_df(conf.features(0))))  lr = new
LogisticRegression().setLabelCol("L").setFeaturesCol("F")  lr_model = lr.fit(lr_inpu
```

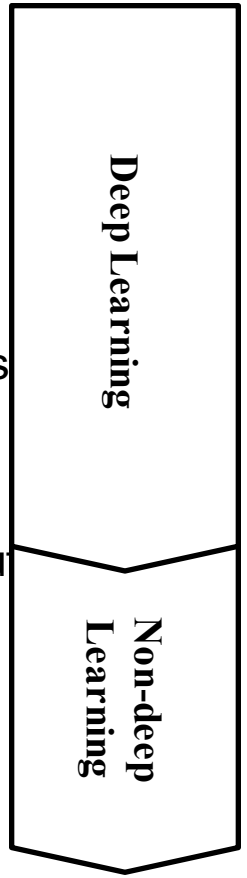

CaffeOnSpark: One Notebook (Python)

http://bit.ly/1REZ0cN

Feature Extraction

In [50]:
```
args['features']='accuracy,ip1,ip2'
args['label']='label'
cfg=Config(sc,args)
```

In [52]:
```
dl_feature_source = DataSource(sc).getSource(cfg,False)
```

In [54]:
```
f=cos.features(dl_feature_source)
```

In [55]:
```
f.show(5)
```

```
+--------+--------+--------------------+--------------------+-----+
|SampleID|accuracy|                 ip1|                 ip2|label|
+--------+--------+--------------------+--------------------+-----+
|00000000|   [1.0]|[-0.0, 3.109636, ...|[-0.6478175, -1.4...|[7.0]|
|00000001|   [1.0]|[1.3683326, -0.0,...|[2.0906663, 1.048...|[2.0]|
|00000002|   [1.0]|[1.5641443, -0.0,...|[-0.773368, 10.61...|[1.0]|
|00000003|   [1.0]|[-0.0, 1.9505613,...|[16.46351, -6.917...|[0.0]|
|00000004|   [1.0]|[0.5979191, 0.075...|[-0.48371825, -2....|[4.0]|
+--------+--------+--------------------+--------------------+-----+
```

In [45]:
```
dl_train_source = DataSource(sc).getSource(cfg,True)
```

In [46]:
```
cos.train(dl_train_source)
```

CaffeOnSpark: UI & Logs

Spark 1.5.1 Jobs Stages Storage Environment Executors com.yahoo.ml.dl.CaffeOnSpark application UI

Spark Jobs (?)

Total Uptime: 3.2 h
Scheduling Mode: FIFO
Active Jobs: 1
Completed Jobs: 47

▸ Event Timeline

Active Jobs (1)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
47	reduce at CaffeOnSpark.scala:203	2016/01/26 05:49:00	3.1 min	0/2	690/1892

Completed Jobs (47)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total	Tasks (for all stages): Succeeded/Total
46	reduce at CaffeOnSpark.scala:203	2016/01/26 05:44:51	4.1 min	1/1 (1 skipped)	945/945 (947 skipped)
45	reduce at CaffeOnSpark.scala:203	2016/01/26 05:40:42	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
44	reduce at CaffeOnSpark.scala:203	2016/01/26 05:36:31	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
43	reduce at CaffeOnSpark.scala:203	2016/01/26 05:31:48	4.7 min	1/1 (1 skipped)	945/945 (947 skipped)
42	reduce at CaffeOnSpark.scala:203	2016/01/26 05:27:33	4.3 min	1/1 (1 skipped)	945/945 (947 skipped)
41	reduce at CaffeOnSpark.scala:203	2016/01/26 05:23:24	4.1 min	1/1 (1 skipped)	945/945 (947 skipped)
40	reduce at CaffeOnSpark.scala:203	2016/01/26 05:19:15	4.1 min	1/1 (1 skipped)	945/945 (947 skipped)
39	reduce at CaffeOnSpark.scala:203	2016/01/26 05:15:06	4.1 min	1/1 (1 skipped)	945/945 (947 skipped)
38	reduce at CaffeOnSpark.scala:203	2016/01/26 05:10:48	4.3 min	1/1 (1 skipped)	945/945 (947 skipped)
37	reduce at CaffeOnSpark.scala:203	2016/01/26 05:06:38	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
36	reduce at CaffeOnSpark.scala:203	2016/01/26 05:02:27	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
35	reduce at CaffeOnSpark.scala:203	2016/01/26 04:58:17	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
34	reduce at CaffeOnSpark.scala:203	2016/01/26 04:54:07	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
33	reduce at CaffeOnSpark.scala:203	2016/01/26 04:49:50	4.3 min	1/1 (1 skipped)	945/945 (947 skipped)
32	reduce at CaffeOnSpark.scala:203	2016/01/26 04:45:40	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
31	reduce at CaffeOnSpark.scala:203	2016/01/26 04:41:22	4.3 min	1/1 (1 skipped)	945/945 (947 skipped)
30	reduce at CaffeOnSpark.scala:203	2016/01/26 04:37:12	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
29	reduce at CaffeOnSpark.scala:203	2016/01/26 04:33:00	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
28	reduce at CaffeOnSpark.scala:203	2016/01/26 04:28:40	4.3 min	1/1 (1 skipped)	945/945 (947 skipped)
27	reduce at CaffeOnSpark.scala:203	2016/01/26 04:24:29	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
26	reduce at CaffeOnSpark.scala:203	2016/01/26 04:20:19	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
25	reduce at CaffeOnSpark.scala:203	2016/01/26 04:16:10	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)
24	reduce at CaffeOnSpark.scala:203	2016/01/26 04:11:59	4.2 min	1/1 (1 skipped)	945/945 (947 skipped)

Firefox File Edit View History Bookmarks Tools Window Help

Logs for container_e01_144364...

gpbl191n07.blue.ygrid.yahoo.com:8042/node/containerlogs/container_e01_1443645088453_152667

http://gpbl191n07.blue.ygrid.yahoo.com:8042/node/containerlogs
/container_e01_1443645088453_1526676_01_000002/afeng/stderr/?start=0

```
  }
  layer {
    name: "accuracy"
    type: "Accuracy"
    bottom: "ip2"
    bottom: "label"
    top: "accuracy"
    include {
      phase: TEST
    }
  }
  layer {
    name: "loss"
    type: "SoftmaxWithLoss"
    bottom: "ip2"
    bottom: "label"
    top: "loss"
  }
I1025 21:56:58.067378 23673 layer_factory.hpp:75] Creating layer data
I1025 21:56:58.067386 23673 net.cpp:98] Creating Layer data
I1025 21:56:58.067391 23673 net.cpp:408] data -> data
I1025 21:56:58.067399 23673 net.cpp:408] data -> label
I1025 21:56:58.067405 23673 net.cpp:130] Setting up data
I1025 21:56:58.068318 23673 net.cpp:139] Top shape: 100 1 28 28 (78400)
I1025 21:56:58.068330 23673 net.cpp:139] Top shape: 100 (100)
I1025 21:56:58.068334 23673 layer_factory.hpp:75] Creating layer label_data_1_split
I1025 21:56:58.068346 23673 net.cpp:98] Creating Layer label_data_1_split
I1025 21:56:58.068351 23673 net.cpp:452] label_data_1_split <- label
I1025 21:56:58.068356 23673 net.cpp:408] label_data_1_split -> label_data_1_split_0
I1025 21:56:58.068363 23673 net.cpp:408] label_data_1_split -> label_data_1_split_1
I1025 21:56:58.068370 23673 net.cpp:130] Setting up label_data_1_split
I1025 21:56:58.068377 23673 net.cpp:139] Top shape: 100 1 1 1 (100)
I1025 21:56:58.068382 23673 net.cpp:139] Top shape: 100 1 1 1 (100)
I1025 21:56:58.068384 23673 layer_factory.hpp:75] Creating layer conv1
I1025 21:56:58.068390 23673 net.cpp:98] Creating Layer conv1
I1025 21:56:58.068394 23673 net.cpp:452] conv1 <- data
I1025 21:56:58.068402 23673 net.cpp:408] conv1 -> conv1
I1025 21:56:58.068408 23673 net.cpp:130] Setting up conv1
I1025 21:56:58.068493 23673 net.cpp:139] Top shape: 100 20 24 24 (1152000)
I1025 21:56:58.068502 23673 layer_factory.hpp:75] Creating layer pool1
I1025 21:56:58.068507 23673 net.cpp:98] Creating Layer pool1
I1025 21:56:58.068511 23673 net.cpp:452] pool1 <- conv1
I1025 21:56:58.068516 23673 net.cpp:408] pool1 -> pool1
I1025 21:56:58.068521 23673 net.cpp:130] Setting up pool1
I1025 21:56:58.068528 23673 net.cpp:139] Top shape: 100 20 12 12 (288000)
I1025 21:56:58.068531 23673 layer_factory.hpp:75] Creating layer conv2
I1025 21:56:58.068538 23673 net.cpp:98] Creating Layer conv2
I1025 21:56:58.068542 23673 net.cpp:452] conv2 <- pool1
I1025 21:56:58.068547 23673 net.cpp:408] conv2 -> conv2
I1025 21:56:58.068552 23673 net.cpp:130] Setting up conv2
I1025 21:56:58.068768 23673 net.cpp:139] Top shape: 100 50 8 8 (320000)
I1025 21:56:58.068778 23673 layer_factory.hpp:75] Creating layer pool2
I1025 21:56:58.068783 23673 net.cpp:98] Creating Layer pool2
I1025 21:56:58.068788 23673 net.cpp:452] pool2 <- conv2
I1025 21:56:58.068792 23673 net.cpp:408] pool2 -> pool2
I1025 21:56:58.068797 23673 net.cpp:130] Setting up pool2
I1025 21:56:58.068804 23673 net.cpp:139] Top shape: 100 50 4 4 (80000)
I1025 21:56:58.068810 23673 layer_factory.hpp:75] Creating layer ip1
I1025 21:56:58.068815 23673 net.cpp:98] Creating Layer ip1
I1025 21:56:58.068826 23673 net.cpp:452] ip1 <- pool2
```

Demo: CaffeOnSpark on EC2

- https://github.com/yahoo/CaffeOnSpark/wiki
  - Get started on EC2
  - Python for CaffeOnSpark

Summary

- CaffeOnSpark open sourced
  - https://github.com/yahoo/CaffeOnSpark

  - Empower Flickr and other Yahoo services
  - Scalable DL made easy

THANK YOU.

bigdata@yahoo-inc.com

SPARK SUMMIT 2016
DATA SCIENCE AND ENGINEERING AT SCALE
JUNE 6-8, 2016 SAN FRANCISCO