Text analysis using Apache MXNet

Qiang Kou (KK) qkou@umail.iu.edu

Indiana University

About me

- Qiang Kou, KK
- · PhD student in Informatics from Indiana University
- R/C++ developer
- Summer intern in AWS from next Monday
- https://github.com/thirdwing/r_fin_2017

MXNet - Deep Learning Framework of Choice at AWS

By Werner Vogels on 22 November 2016 09:00 AM | Permalink | Comments (12)





Contact Info

Werner Vogels CTO - Amazon.com

http://www.allthingsdistributed.com/2016/11/mxnet-default-framework-deep-learning-aws.html

INCUBATOR



MXNet Project Incubation Status

This page tracks the project status, incubator-wise. For more general project status, look on the project website.

Description

MXNet is an open-source deep learning framework that allows you to define, train, and deploy deep neural networks on a wide array of devices, from cloud infrastructure to mobile devices. It is highly scalable, allowing for fast model training, and supports a flexible programming model and multiple languages.

News

2017-01-23 Project enters incubation.

http://incubator.apache.org/projects/mxnet.html

- Open-source under Apache License 2.0
- Scalability
- Portability
- Multiple programming languages: C++, JavaScript, Python, R, Matlab, Julia, Scala, Go and Perl



MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems

Why do we need GPU?

Why do we need GPU?



Why do we need GPU?

- Inception-BN network
- + 224 \times 224 figures, batch size 64
- GTX 1080: 1.739 seconds per batch
- Intel(R) Core(TM) i7-5960X CPU @ 3.00GHz: 354.949 seconds per batch¹

¹OpenBLAS used. Much better performance using MKL and MKLdnn.

R and Windows support



data <- mx.symbol.Variable("data")</pre>

fc1 <- mx.symbol.FullyConnected(data,name="fc1",num_hidden=128)
act1 <- mx.symbol.Activation(fc1,name="relu1",act_type="relu")
fc2 <- mx.symbol.FullyConnected(act1,name="fc2",num_hidden=64)
act2 <- mx.symbol.Activation(fc2,name="relu2",act_type="relu")
fc3 <- mx.symbol.FullyConnected(act2,name="fc3",num_hidden=10)
softmax <- mx.symbol.SoftmaxOutput(fc3,name="sm")</pre>

graph.viz(softmax)



Text classification

Convolutional Neural Network



Convolutional Kernel Networks, arXiv:1406.3332 [cs.CV]

Character-level Convolutional Networks

Character-level Convolutional Networks

Character-level Convolutional Networks for Text Classification, arXiv:1509.01626 [cs.LG]

Amazon review data

- Alphabet size: 69
- VDCNN model
- Examples
 - **Books**: "Great story. Bit slow to start but picks up pace. Can't wait for the next instalment."
 - Clothing, Shoes & Jewelry: "Good for gym. Love that you hands can breathe, but they don't give as much traction for gloves"
- 43.23 samples/sec using CPU
- 638.36 samples/sec using GPU
- 100000 training samples and 20000 validation samples

Amazon review data

Epoch

Recurrent neural network

Feedforward neural network

A Critical Review of Recurrent Neural Networks for Sequence Learning, arXiv:1506.00019 [cs.LG]

Recurrent neural network

A Critical Review of Recurrent Neural Networks for Sequence Learning, arXiv:1506.00019 [cs.LG]

Recurrent neural network

http://karpathy.github.io/2015/05/21/rnn-effectiveness/

Long short-term memory (LSTM)

LSTM in MXNet

```
model <- mx.lstm(</pre>
 X.train, X.val, ctx = mx.gpu(),
 num.round = num.round,
  update.period = update.period,
 num.lstm.layer = num.lstm.layer,
  seq.len = seq.len,
 num.hidden = num.hidden,
  num.embed = num.embed,
  num.label = vocab,
  batch.size = batch.size,
  input.size = vocab,
  initializer = mx.init.uniform(0.1),
  learning.rate = learning.rate,
  wd = wd, clip_gradient = clip_gradient
```

LSTM in MXNet

The input:

First Citizen: I say unto you, what he hath done famously, he did it to that end: though soft-conscienced men can be content to say it was for his country he did it to please his mother and to be partly proud; which he is, even till the altitude of his virtue. Second Citizen: What he cannot help in his nature, you account a vice in him. You must in no way say he is covetous.

.

LSTM in MXNet

After "learning" for half an hour on my laptop, a 3-layer LSTM network said:

at like you, Eim, and her bear to told not we a fits and my tongue That clothess your quardency let our ends, That those where I knot thou,' are you have fortune of your banner to say Old which friends thou have prout to your bed

MOSTENSIO:

And me salt, your sowest amated all alord, Suscatly the more, legd and than the proting the till'd of mercy, Thou knee strafks now Genetty, let and far of good man?

Summary

- · Character-level CNN and RNN models
- Model training using R with GPU support
- Multiple programming languages
- https://github.com/thirdwing/r_fin_2017

MXNet resources

mxnet.io

Blogs from Microsoft

Tag: MXNet

Deep Learning on the New Ubuntu-Based Data Science Virtual Machine for Linux

April 18, 2017 by Cortana Intelligence and ML Blog Team // 0 Comments

Authored by Paul Shealy, Senior Software Engineer, and Gopi Kumar, Principal Program Manager, at Microsoft. Deep learning has received significant attention recently for its ability to create machine learning models with very high accuracy. It's especially popular in image and speech recognition tasks, where the availability of massive datasets with rich information make it feasible... Read more

Tags: Al, Caffe, Caffe2, CNTK, Cognitive Toolkit, Cortana Intelligence, Data Science, Data Science VM, Deep

MXNet repo

		GitHub, Inc.	Ċ		
C Features Busine	ss Explore Pricing		This repository Se	earch	Sign in or Sign up
dmlc / mxnet			• Watch 915	★ Star 9,6	317 ¥ Fork 3,614
↔ Code ① Issues 1,6	106 11 Pull requests 65	III Projects 4	III Graphs		
for Python, R, Julia, Scala deep-learning deep-neural	a, Go, Javascript and mor	re http://mxnet.io uted-systems machine-learning	artificial-intelligence		onodalor,
3 5,292 commits	2 11 branches	S≥ 21 releases	🞎 340 contribu	itors	화 Apache-2.0
Branch: master • New pu	II request			Find file	Clone or download -
😛 sifmelcara committed w	ith piiswrong Fix RMSProp up	odate rule (#6235)		Latest comm	it 38f7c55 18 hours ago
💼 .github	Create ISSUE_TEMPLATE.md				5 months ago
R-package	[R] fix message error on Jupyter notebook (#5854)				29 days ago
	Cleanup compilations & tests (#5309)				
amalgamation	Cleanup compilatio	ns & tests (#5309)			2 months ago

Acknowledgement

Thank you!

Qiang Kou (KK) qkou@umail.iu.edu