Deep Learning & TensorFlow with Keras

These slides are based on:

Deep Learning for Recommender Systems

Alexandros Karatzoglou (Scientific Director @ Telefonica Research) <u>alexk@tid.es,</u>@alexk_z

> Balázs Hidasi (Head of Research @ Gravity R&D) balazs.hidasi@gravityrd.com, @balazshidasi

> > RecSys'17, 29 August 2017, Como

https://www.slideshare.net/kerveros99/deep-learning-forrecommender-systems-recsys2017-tutorial https://cs224d.stanford.edu/lectures/CS224d-Lecture7.pdf

CS224d: TensorFlow Tutorial

Bharath Ramsundar

Why Deep Learning?



ImageNet challenge <u>error rates</u> (red line = human performance)

Why Deep Learning?





Neural Model



Neuron a.k.a. Unit



Learning



Stochastic Gradient Descent

Generalization of (Stochastic) Gradient Descent

$$E = \frac{1}{2}(f - y)^2$$
$$f = \mathbf{w}^T \mathbf{x}$$

for
$$i = 1, 2, ..., n$$

 $\mathbf{w} := \mathbf{w} - \eta \nabla_f E \mathbf{x}_i$

Stochastic Gradient Descent



Modern Deep Networks

• Ingredients:

 Rectified Linear Activation function a.k.a. ReLu

$$\sigma(x) = max(0, x)$$

$$\sigma(x) = max(\alpha x, x) \quad \alpha < 1$$



Modern Deep Networks

- Mini-batches:
 - Stochastic Gradient Descent (or its evolutions)
 - Compute gradient over many (50 -100) data points (minibatch) and update.
- Shared parameters
- Several layer / network-level templates
 - CNN, RNN, 2vec models, Autoencoders,
 Generative Adversal Networks (GAN),...

- Transfer learning (from one problem to another)

Modern Feedforward Networks

- Ingredients:
- Adagrad a.k.a. adaptive learning rates





Convolutional deep network



[Fig source: http://visiono3.csail.mit.edu/cnn_art/index.html]

Deep-Learning Package Zoo

- Torch
- Caffe
- Theano (Keras, Lasagne)
- CuDNN
- Tensorflow
- Mxnet
- Etc.



theano dmlc **mxnet**





Deep-Learning Package Design Choices

- Model specification: Configuration file (e.g. Caffe, DistBelief, CNTK) versus programmatic generation (e.g. Torch, Theano, Tensorflow)
- For programmatic models, choice of high-level language: Lua (Torch) vs. Python (Theano, Tensorflow) vs others.
- We chose to work with **python** because of rich community and library infrastructure.
- However, the most important is DEAD or ALIVE
 TensorFlow is pretty much alive©

What is TensorFlow?

- TensorFlow is a deep learning library recently(2017?) open-sourced by Google.
- But what does it actually do?
 - TensorFlow provides primitives for defining functions on tensors and automatically computing their derivatives.
 - Much more nowadays
 - <u>https://docs.anaconda.com/anaconda/user-guide/tasks/tensorflow/</u>



But what's a Tensor?

• Formally, tensors are multilinear maps from vector spaces to the real numbers (Vvector space, and V^* dual space)

$$f: \underbrace{V^* \times \cdots V^*}_{p \text{ copies}} \times \underbrace{V \times \cdots V}_{q \text{ copies}} \to \mathbb{R}$$

- A scalar is a tensor ($f : \mathbb{R} \to \mathbb{R}, f(e_1) = c$)
- A vector is a tensor ($f : \mathbb{R}^n \to \mathbb{R}, f(e_i) = v_i$)
- A matrix is a tensor $(f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}, f(e_i, e_j) = A_{ij})$
- Common to have fixed basis, so a tensor can be represented as a multidimensional array of numbers.

Why the "Flow"?

- Tensorflow constructs a graph of the evaluation G
- The nodes in the graph are evaluated upon request (.eval() / session.run()) – TF1.0
- TenorFlow can automatically calculate derivates for each of the nodes and therefor we can use SGD easily to optimize w.r.t. Arbitrary error metric
- Significant changes in TF2.0 interface (but internally it remains a graph all the same)



TensorFlow Session Object (1) # TF1.0

 "A Session object encapsulates the environment in which Tensor objects are evaluated" - <u>TensorFlow Docs</u>

```
In [20]: a = tf.constant(5.0)
In [21]: b = tf.constant(6.0)
In [22]: c = a * b
In [23]: with tf.Session() as sess:
    ....: print(sess.run(c))
    ....:
30.0
30.0
```

Placeholders and Feed Dictionaries (2)



Ex: Linear Regression in TensorFlow (3)



TensorFlow 2.0

- Eager evaluation (no Sessions anymore)
- Model building with TF.Keras (high level abstraction)
- Compile model with error metric and solver
 - Train -> evaluate -> predict
- TF.data.Dataset handling the input data
 - (instead of placeholders)
 - Support for pipelining
- Seemingly much more single-purpose than TF1.0
 - But common things are done easier
 - Heavily researched (outdated tutorials & StOF)

TensorFlow 2.0 Datasets

- Use existing data or create your own e.g. from numpy
- Support for pipelined pre-processing

```
import tesorflow as tf
train, test = tf.keras.datasets.fashion_mnist.load_data()
```

```
images, labels = train
images = images/255
```

```
dataset = tf.data.Dataset.from_tensor_slices((images, labels))
dataset
```

>> <TensorSliceDataset shapes: ((28, 28), ()), types: (tf.float64, tf.uint8)>

Then you want to create a way to iterate through data More details: https://www.tensorflow.org/guide/data

TensorFlow 2.0 Models (TF.Keras)

- Common types of network layers
 - (abstraction on tensors)
- Sequential model
 - Simple pipeline processing (e.g. via dense layers)

```
import tensorflow as tf
from tensorflow import keras
from tensorflow.keras import layers

# Define Sequential model with 3 layers
model = keras.Sequential(
    [
        layers.Dense(2, activation="relu", name="layer1"),
        layers.Dense(3, activation="relu", name="layer2"),
        layers.Dense(4, name="layer3"),
    ]

# Call model on a test input
x = tf.ones((3, 3))
y = model(x)
```

More details: <u>https://www.tensorflow.org/guide/keras/sequential_model</u>

TensorFlow 2.0 Models (TF.Keras)

• Continue with tf2_basics.ipynb

"Large" homework: - see nprg067_04_task.ipynb

- Build a siammese network:
 - Is the pair of data from the same class?
 - One-shot classification task
 - What if a new class emerges?
 - Re-train model (standard ML) or use similarity of embeddings
 - Two "towers" that share parameters and the similarity calculation on top of it
 - Use tf.keras layers + model