



Hosted by Diane Mueller
Director, Community Development, Red Hat

@pythondj

#TNT

TensorFlow on OpenShift

Guest Speakers:

Gokila Dorai, Florida State University
Eleanor Mehlenbacher, Red Hat Consulting

Bringing together thinkers, creators, and coders
to discuss some of the most interesting ideas and innovations
in technology today

TensorFlow and Openshift

Gokila Dorai (FSU) with

Eleanor Mehlenbacher & Ryne Fagin
(Red Hat Consulting)

Google's Tensorflow

Google's TensorFlow is one of the most popular tools for deep learning. With TensorFlow, you can build your own machine learning systems and use them to power your own applications.

Today, we will learn how to use TensorFlow to build and deploy machine learning models.

What is Tensorflow?

TensorFlow is one of the most popular deep learning frameworks available. It's used for everything from cutting-edge machine learning research to building new features for the hottest start-ups in Silicon Valley.

If you are new to machine learning, you may be wondering what TensorFlow actually is. TensorFlow is a software framework for building and deploying machine learning models.

Machine learning is a different approach than traditional programming. With traditional programming, we write the program that tells the computer exactly what to do to complete the task.

What is Machine Learning?

With machine learning, we don't explicitly tell the computer how to do something. Instead we show a training data and the machine learning algorithm uses the training data that come up with its own rules to complete the task.

Deep neural networks built with TensorFlow are used in many different areas like

- Image recognition, where you recognize what objects appear in a picture.

- Speech recognition where you turn speech into text.

- Image style transfer where you can make a photograph look like it was painted in the style of a famous artist

- Language translation where you translate from one human language to another.

Why do we prefer TensorFlow?

Developed at Google, TensorFlow was designed to be a common platform for building machine learning applications internally.

In late 2015, they decided to share it with the world and release it as an open source application.

TensorFlow is licensed under the Apache open source software license. That means you can use it, modify it, and even redistribute modified version of it. You don't have to pay any licensing fees to use TensorFlow. TensorFlow gives you the basic building blocks that you need to design, train, and deploy machine learning models. While it's flexible enough to be used for several different types of machine learning types of algorithms, it's typically used to build deep neural networks.

Data in Tensorflow

The name TensorFlow comes from the design of the system. TensorFlow is designed to work with large data sets made up of many different individual attributes.

Any data that you want to process with TensorFlow has to be stored in the multi-dimensional array.

These multi-dimensional arrays are also called tensors. To run operations on the data set, you construct a computational graph similar to a flow chart that determines how data flows from one operation to the next.

3	3
1	2
2	4

2	1
3	3
2	4



Add



Square



25	9
9	36
16	64

Openshift

First, we will cover how to get TensorFlow installed and running on your system. Then, we will explore how to create and train a machine learning model in TensorFlow.

We will be using OpenShift which is an open-source container application platform from Redhat.

Next, we will use TensorFlow's visualization tools to analyze and improve our machine learning model. Finally we will deploy the model to the cloud where we can call it from any program. Let's get started!



Hosted by Diane Mueller
Director, Community Development, Red Hat

@pythondj

#TNT

Learn more at

<https://commons.openshift.org>

Bringing together thinkers, creators, and coders
to discuss some of the most interesting ideas and innovations
in technology today