AIDC

INTEL AI DEVCON 2018

AIDC
INTEL AI DEVCON 2018

Seeing AI: Currency Classifier

Sandeep.Alur @Microsoft.com
@saalur

intel AI

Microsoft

Seeing AI

Turning the visual world into an audible experience.

Microsoft

Well! You may know this…

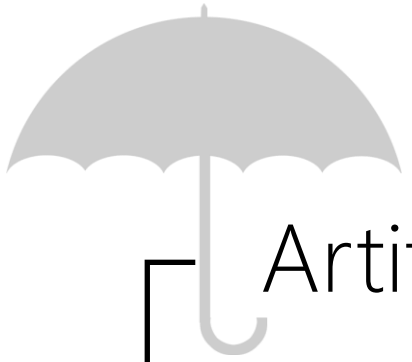

Artificial Intelligence

Machine Learning

Deep Learning

Deep Learning

Artificial Neural Network

Biological Neural Network

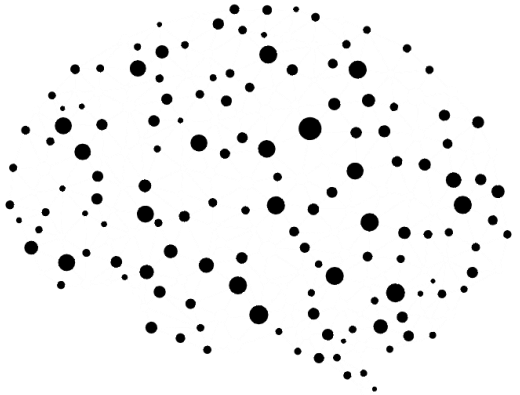

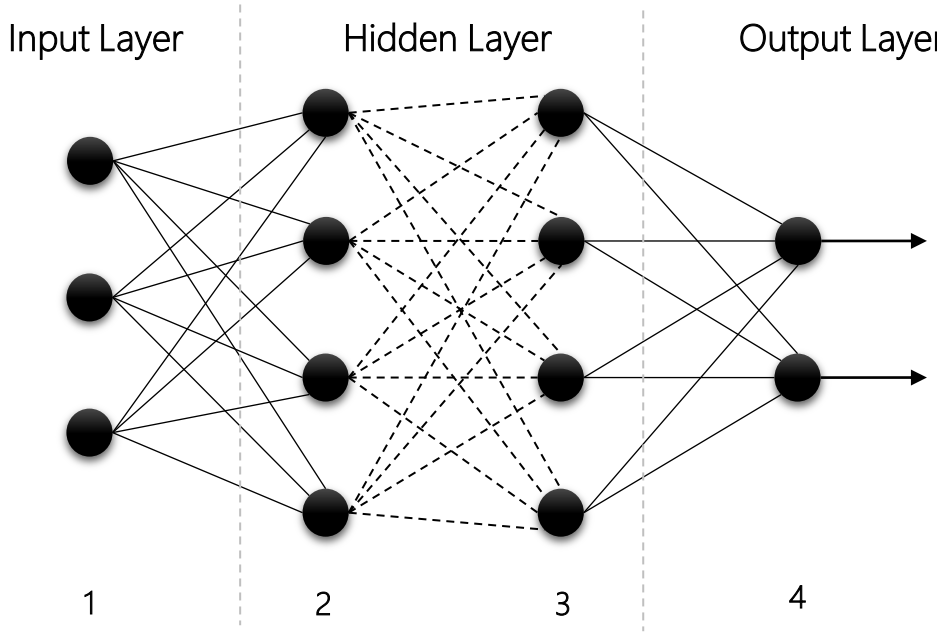

Microsoft

A FEW ALGORITHMS

- Convolution Neural Network (CNN)
- Recurrent Neural Network (RNN)
- Long Short Term Memory (LSTM)
- Generative Adversarial Network (GAN)
- Autoencoder
- ⋮

Data -> CORPUS

SPEECH	IMAGES	VIDEO	TEXT
LDC Linguistic Data Consortium	IMAGENET	YouTube 8M	Wikipedia

Microsoft

Well! You may know this...a refresher

Supervised Learning

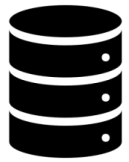

Unsupervised Learning

Reinforcement Learning

• • •

Machine Learning Model

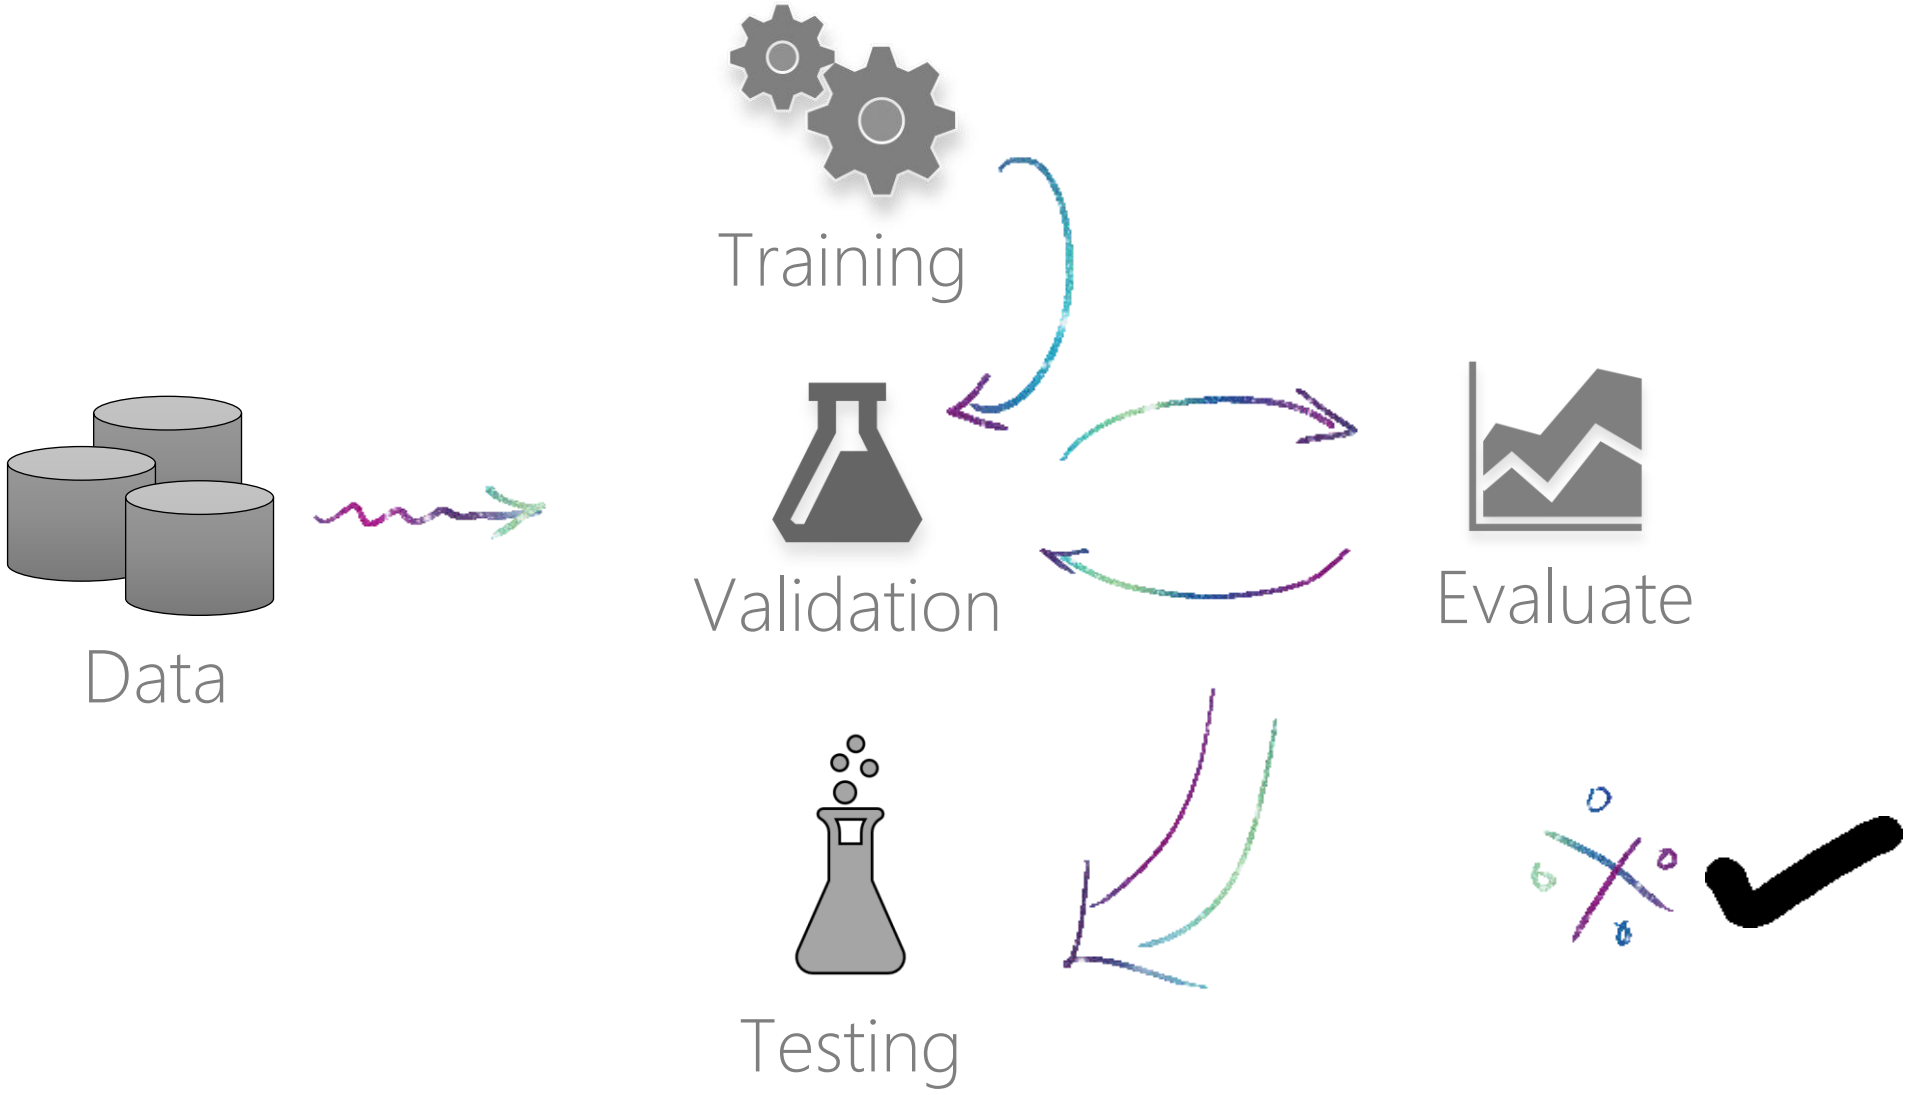

Popular Deep Learning Frameworks

Foundation: A Neural Network Algorithm

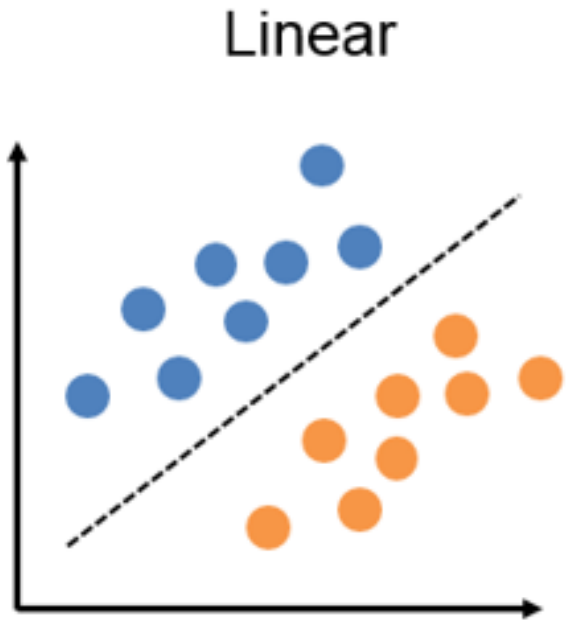

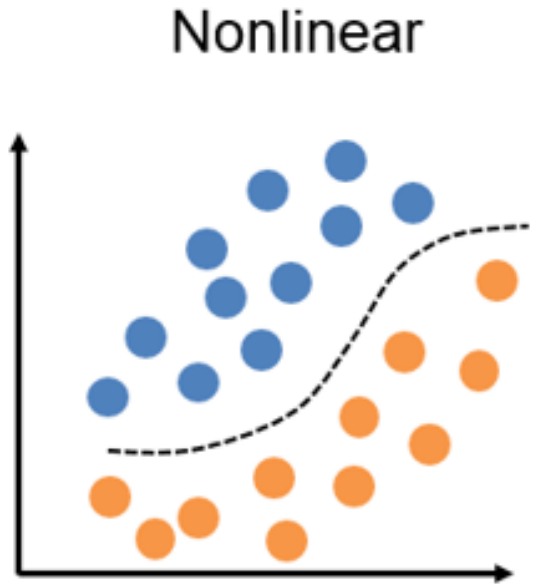

Single Perceptron

Multilayer Perceptron

Neural Network

Feed Forward

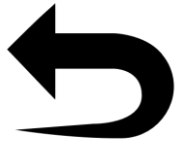

Back Propagation

The Beginning…

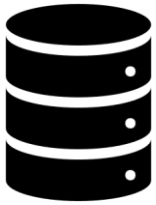

Lots of Training Data

Significant Computational Resources

Coffee Mode

Model Development and Training

Data Collection…A Task

- Image Re-sizing and Standardization
- Resolution & Features
- Variety (Angles, Background)

Data Augmentation

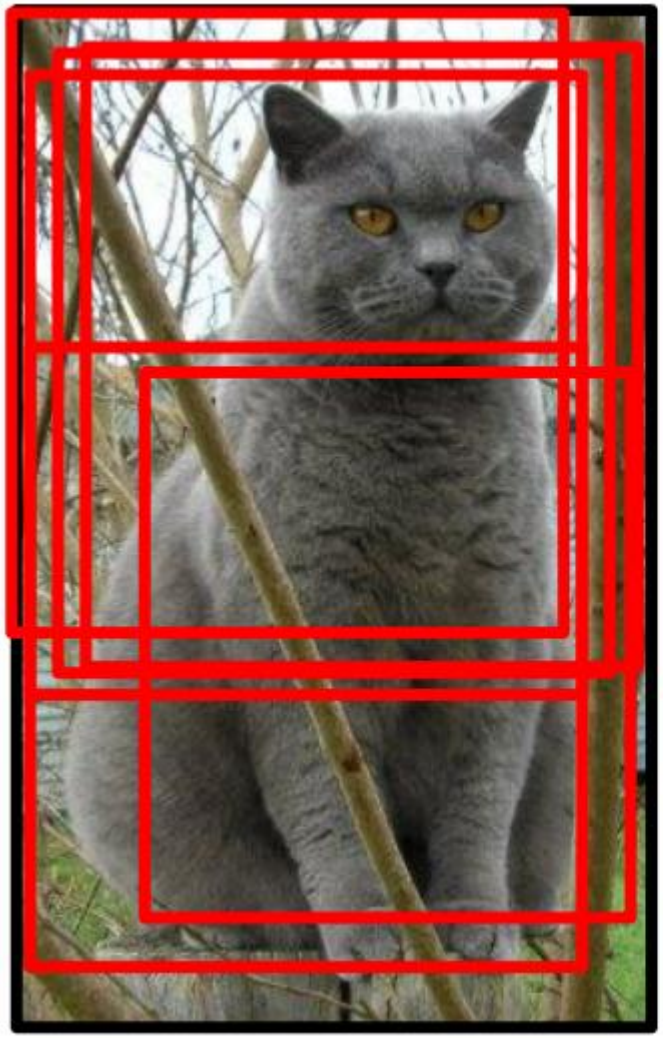

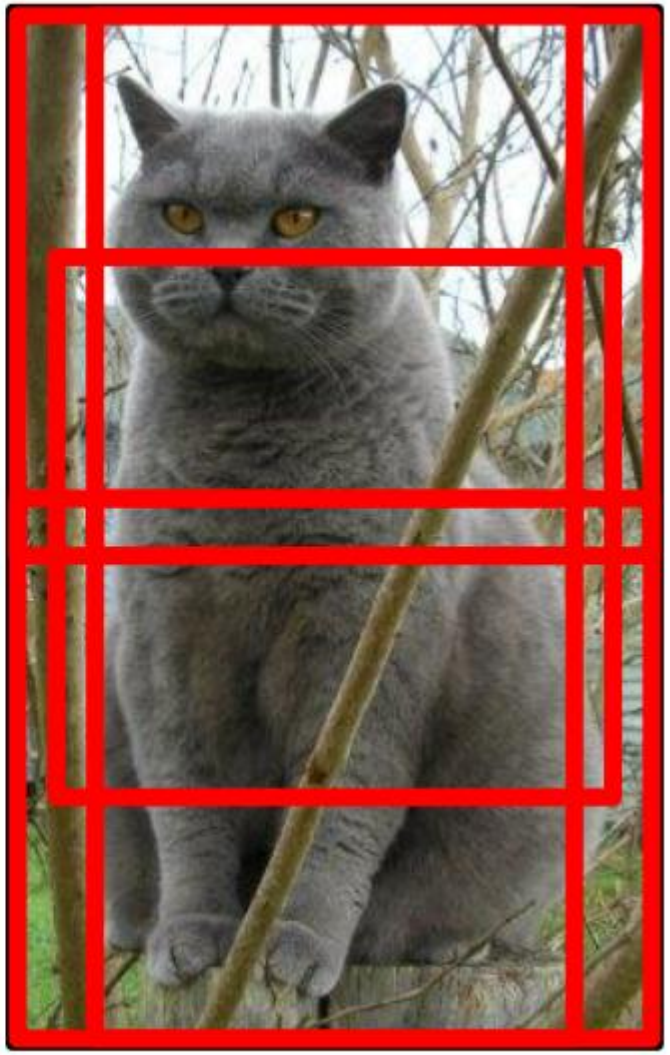

Image Source: https://deeplearningsandbox.com/

Transfer Learning

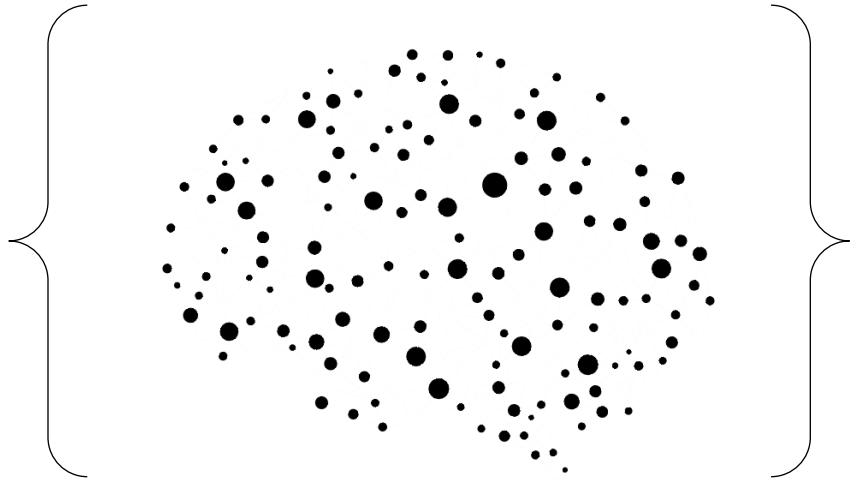

Pre-Trained Models

Reload the weights
&
Freeze Pre-trained Layers

Aggregate Model

Designing a DNN

- Art
- Science
- Trial & Error

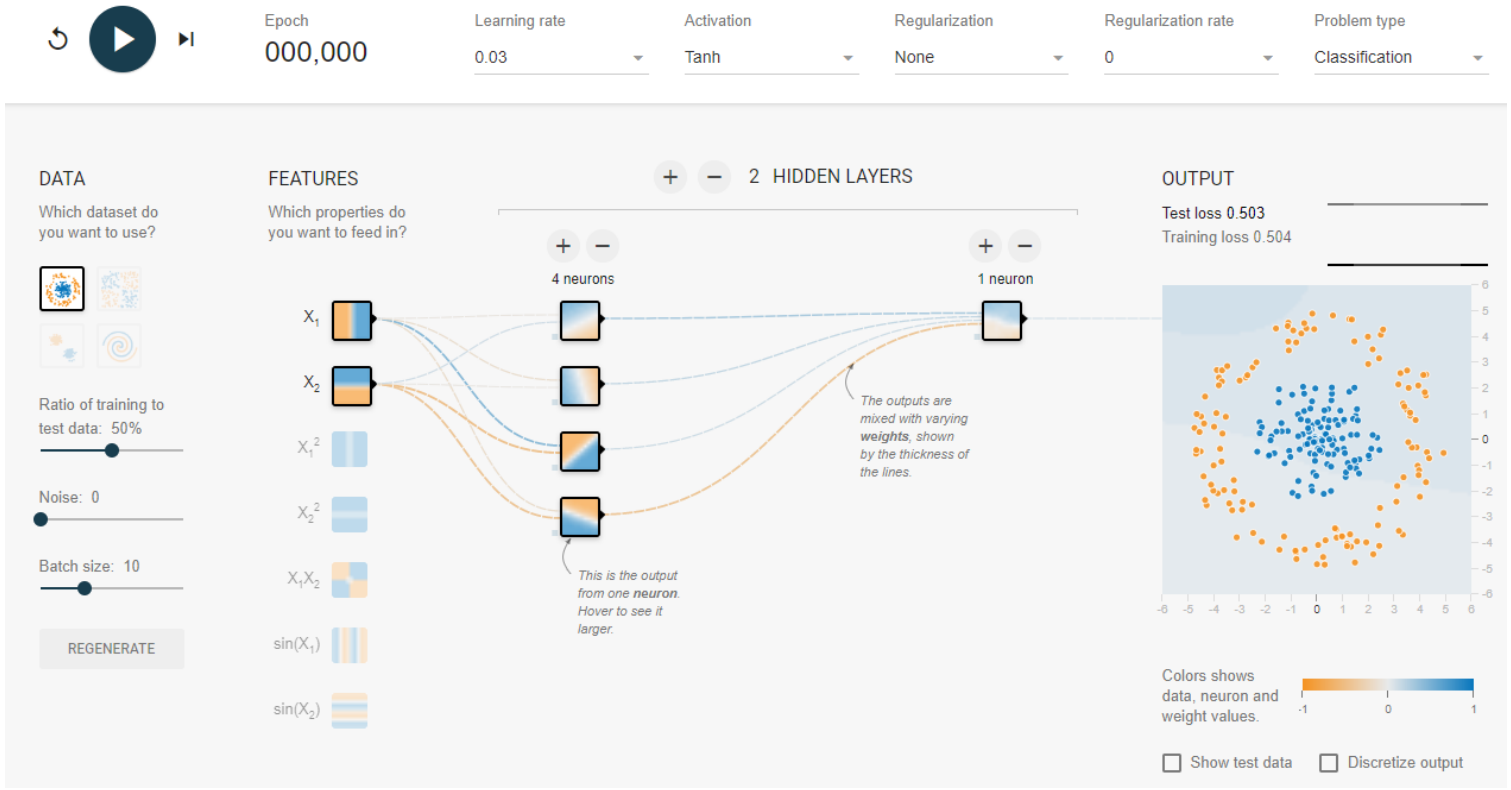

Resource: http://playground.tensorflow.org

A few knobs to adjust

Data

Features (X)

Learning Rate

Problem Type

Hidden Layers

# of Neurons

Activation

Epoch

Iterations

Batch Size

Pool Size

Test Loss

Training Loss

Output (Y)

Convolutional Neural Network

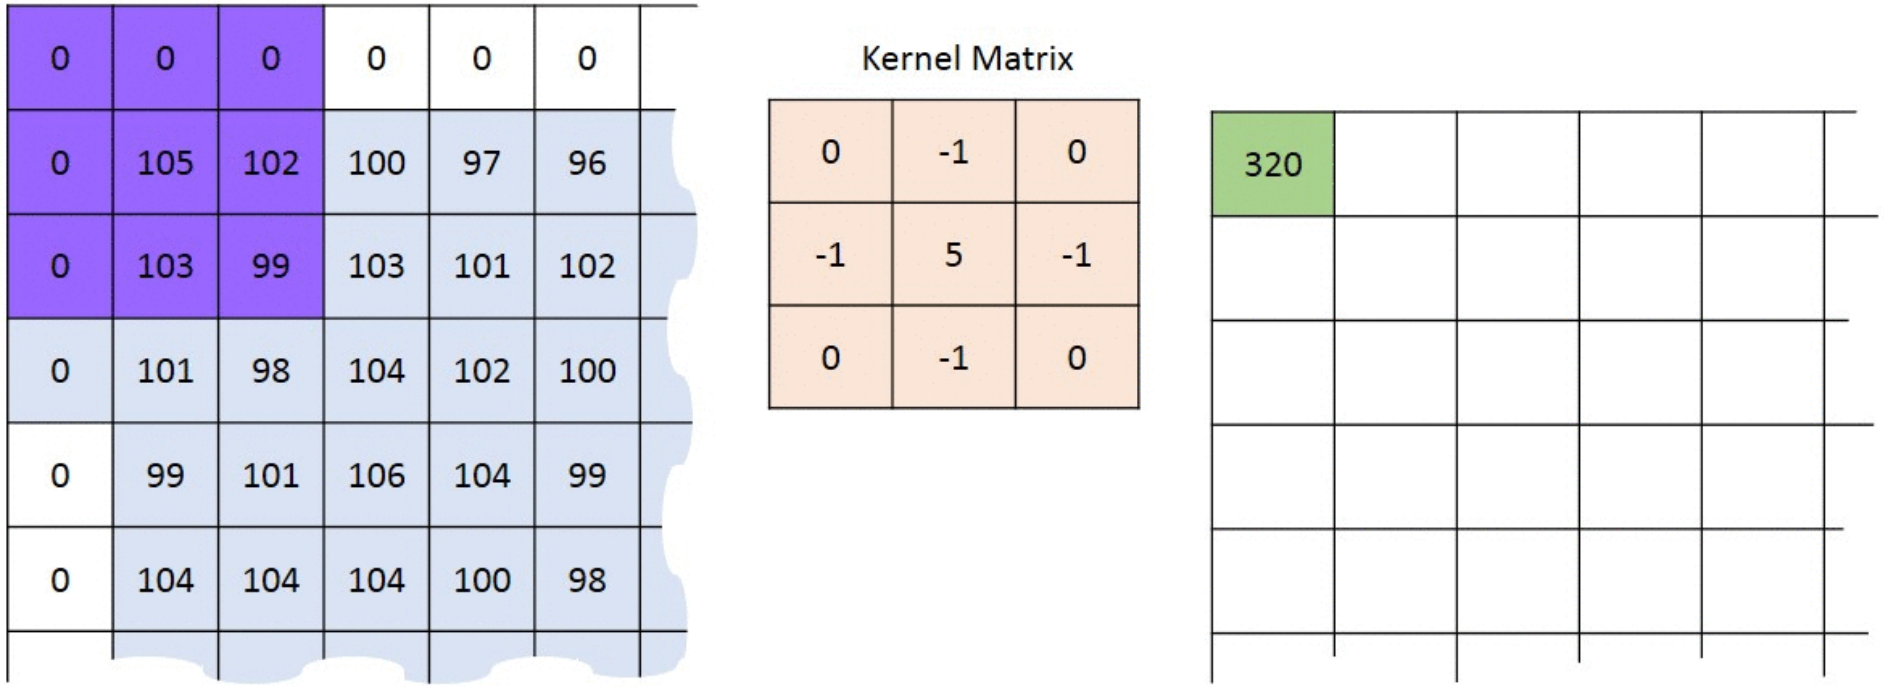

$$0*0+0*-1+0*0$$
$$+0*-1+105*5+102*-1$$
$$+0*0+103*-1+99*0 = 320$$

Convolution with horizontal and vertical strides = 1

Image Source: http://machinelearninguru.com/computer_vision/basics/convolution/convolution_layer.html

End to End – A Summary

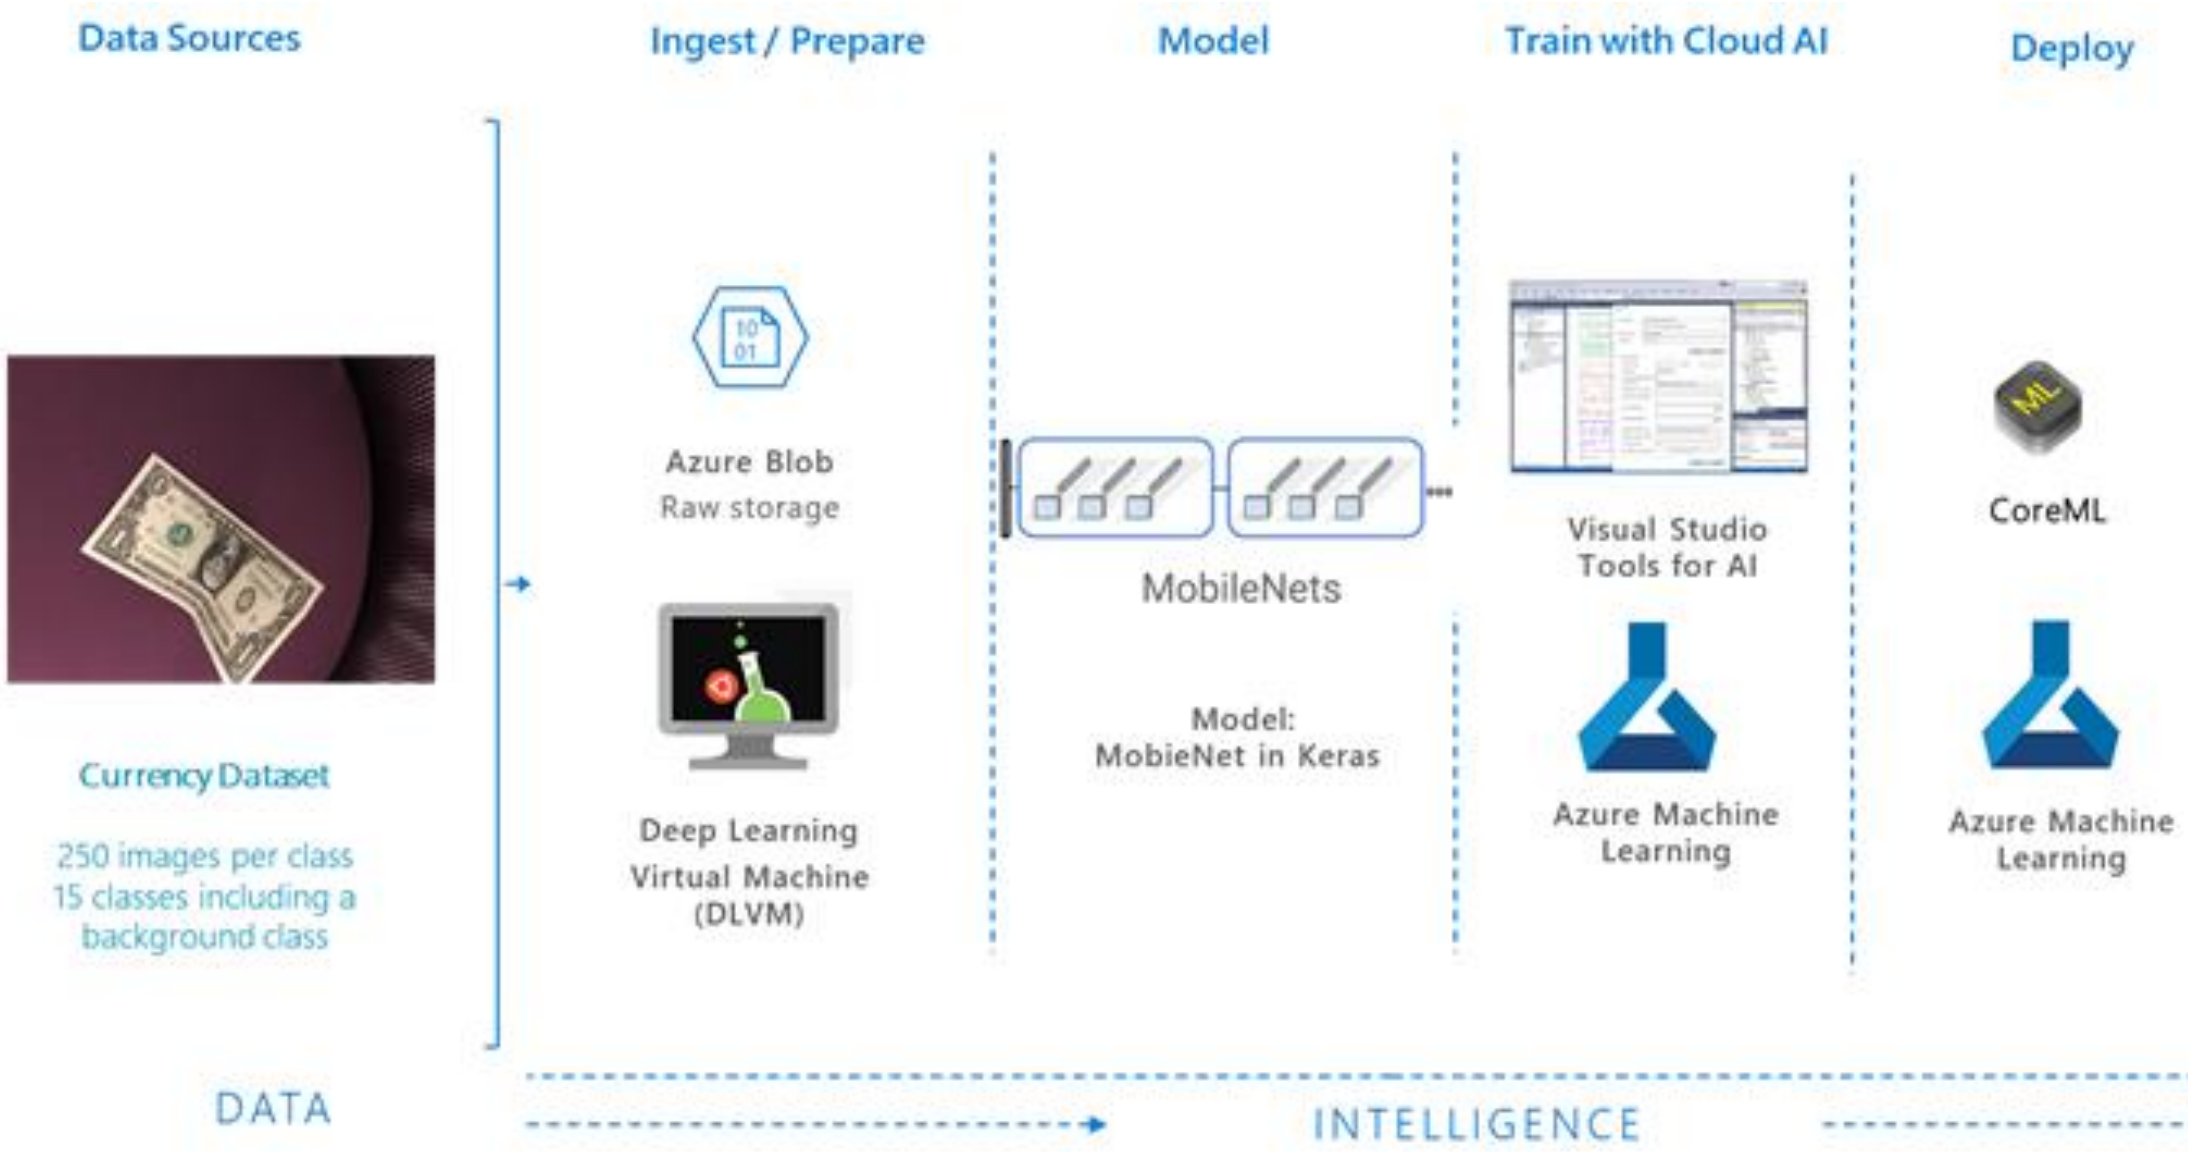

Azure Machine Learning

Azure Machine Learning Studio

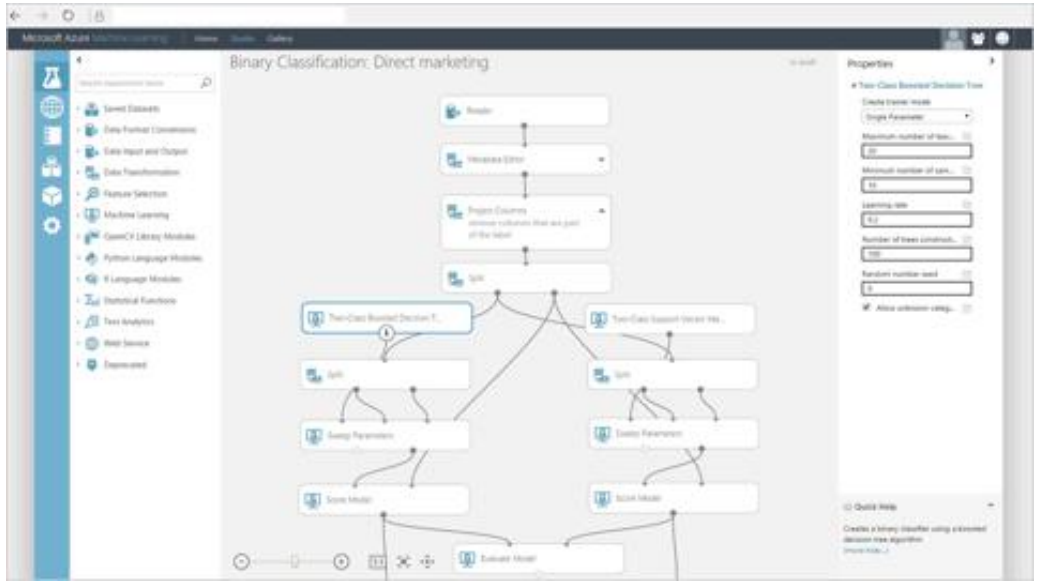

Deep Learning Virtual Machine

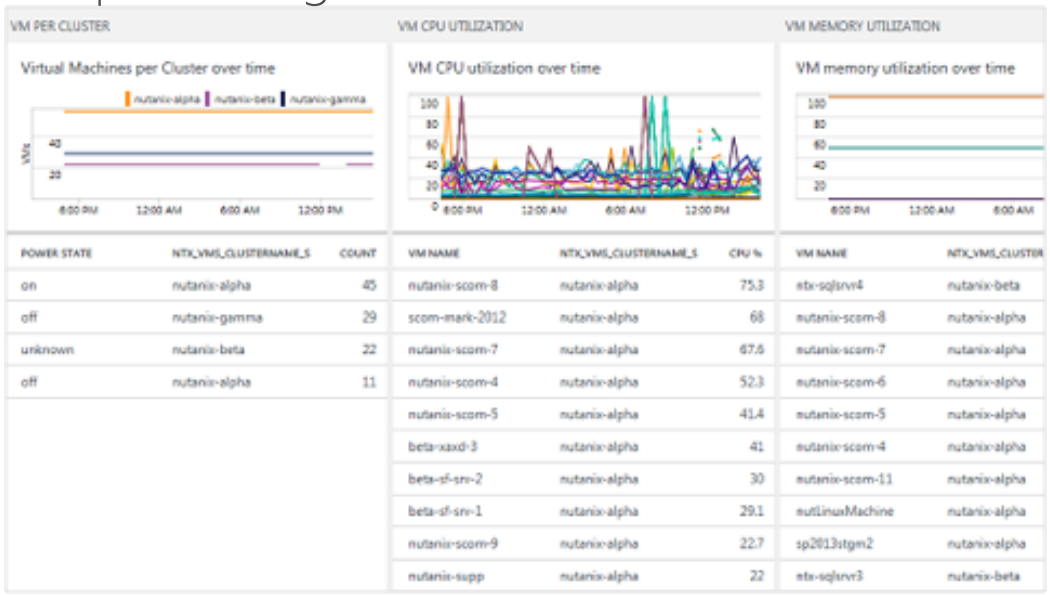

Azure Machine Learning services

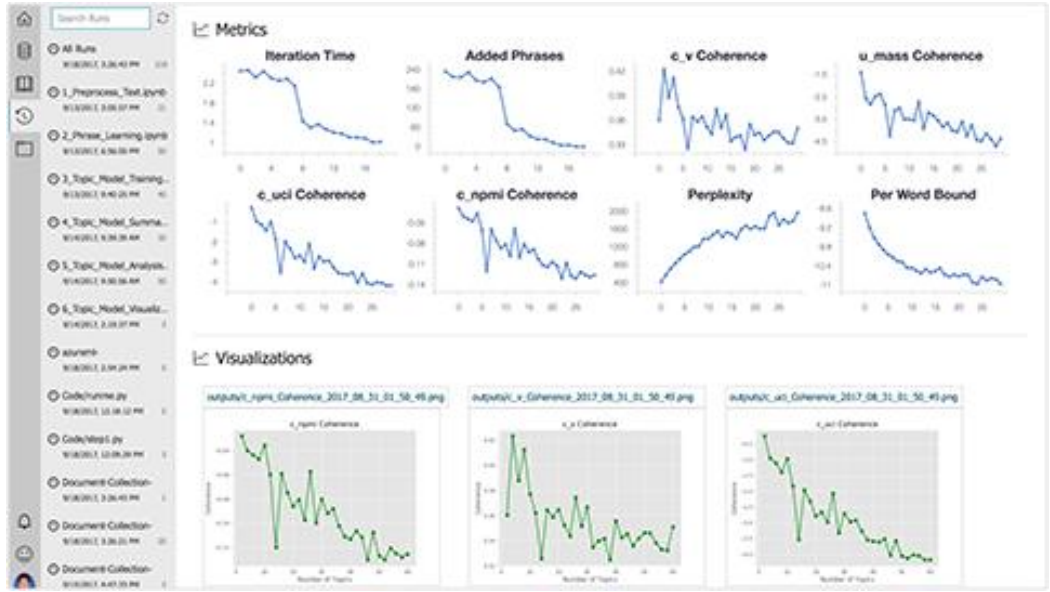

Microsoft

THANK YOU

Sandeep.Alur @Microsoft.com

 @saalur

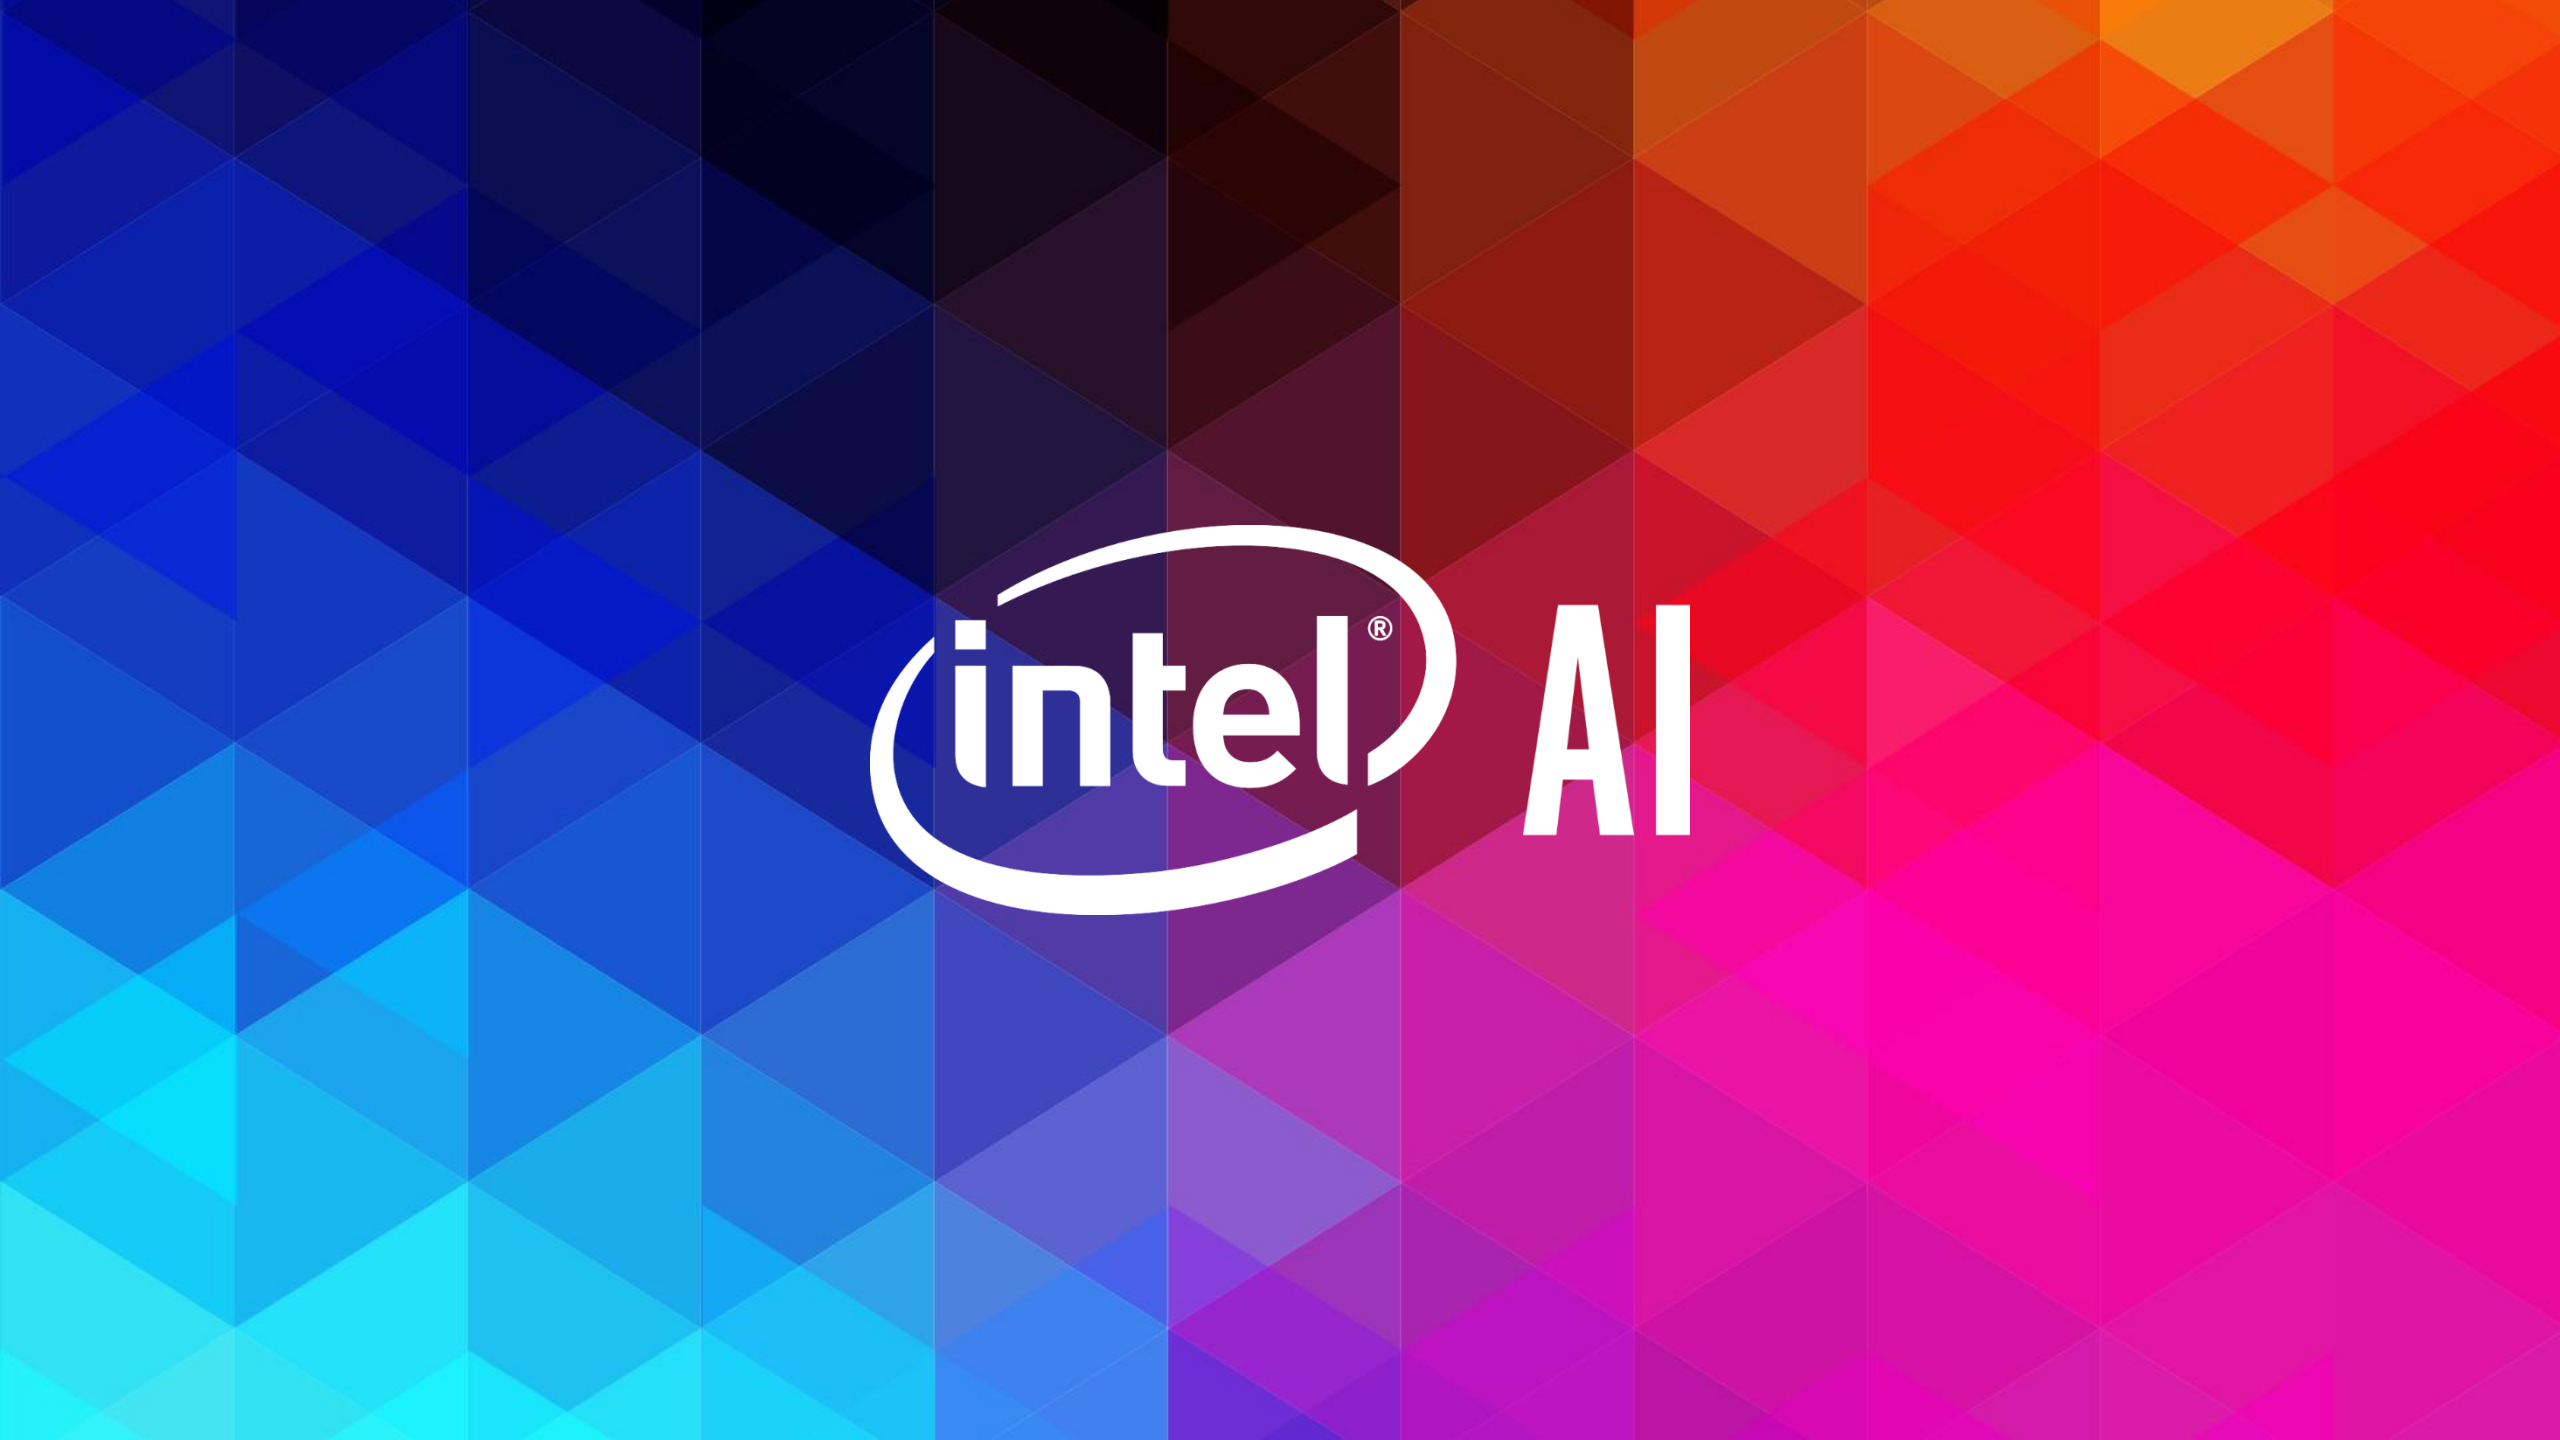
intel®
AI