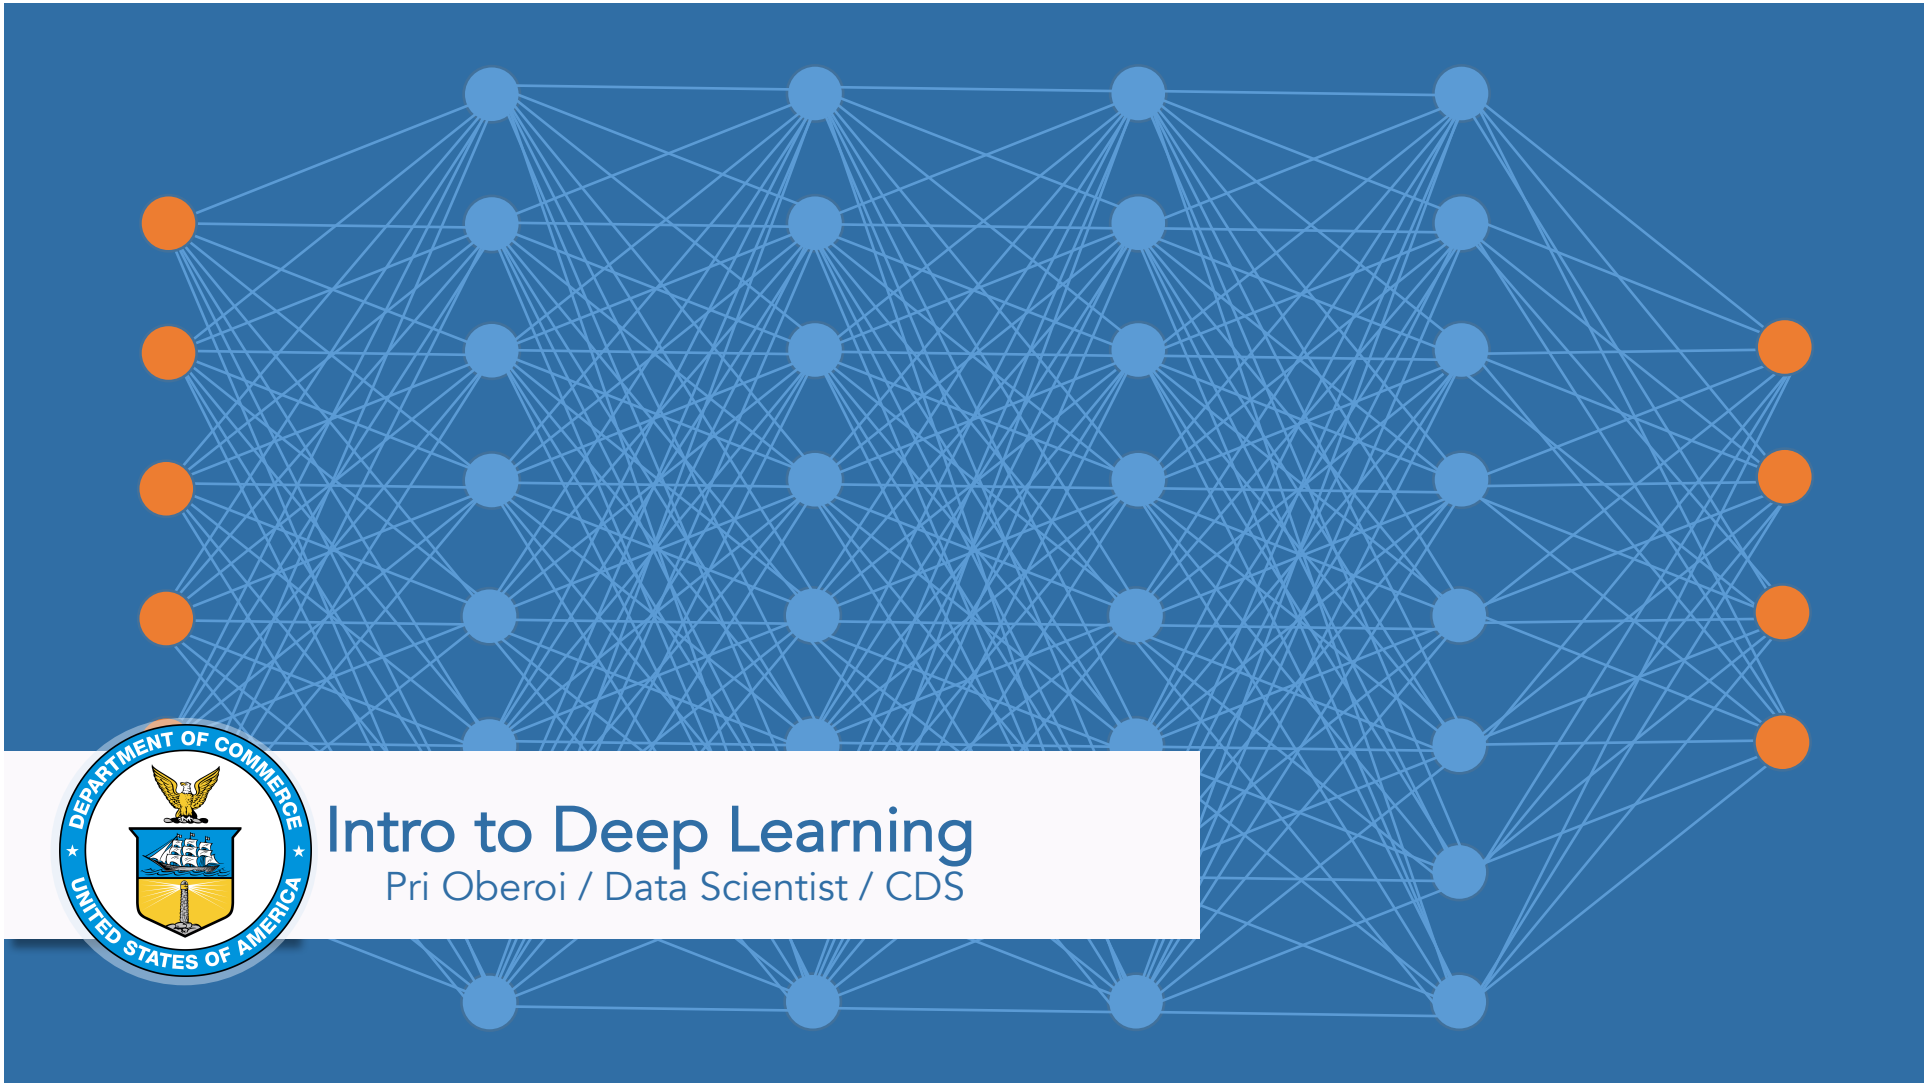

Intro to Deep Learning

Pri Oberoi / Data Scientist / CDS

roadmap

> when is deep learning useful

> what is it

> example

> optimization

> performance

COMMERCE DATA SERVICE

deep learning problems

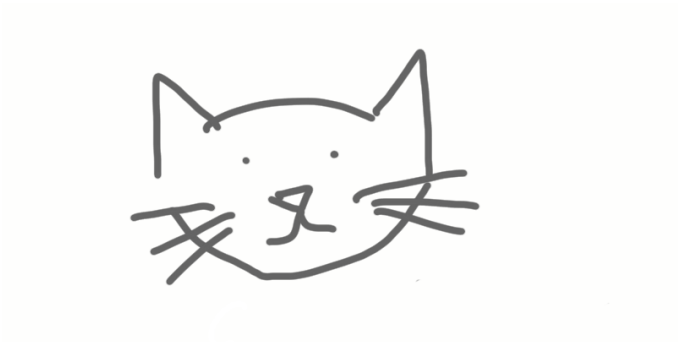

COMMERCE DATA SERVICE

easy for humans

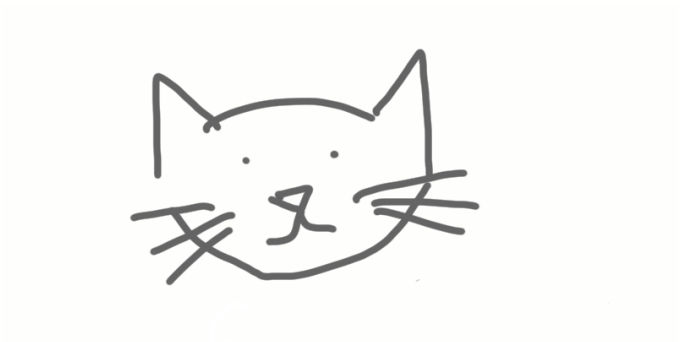

Cat

Not a cat

increasing complexity

Cat

COMMERCE DATA SERVICE

input vector → output vector

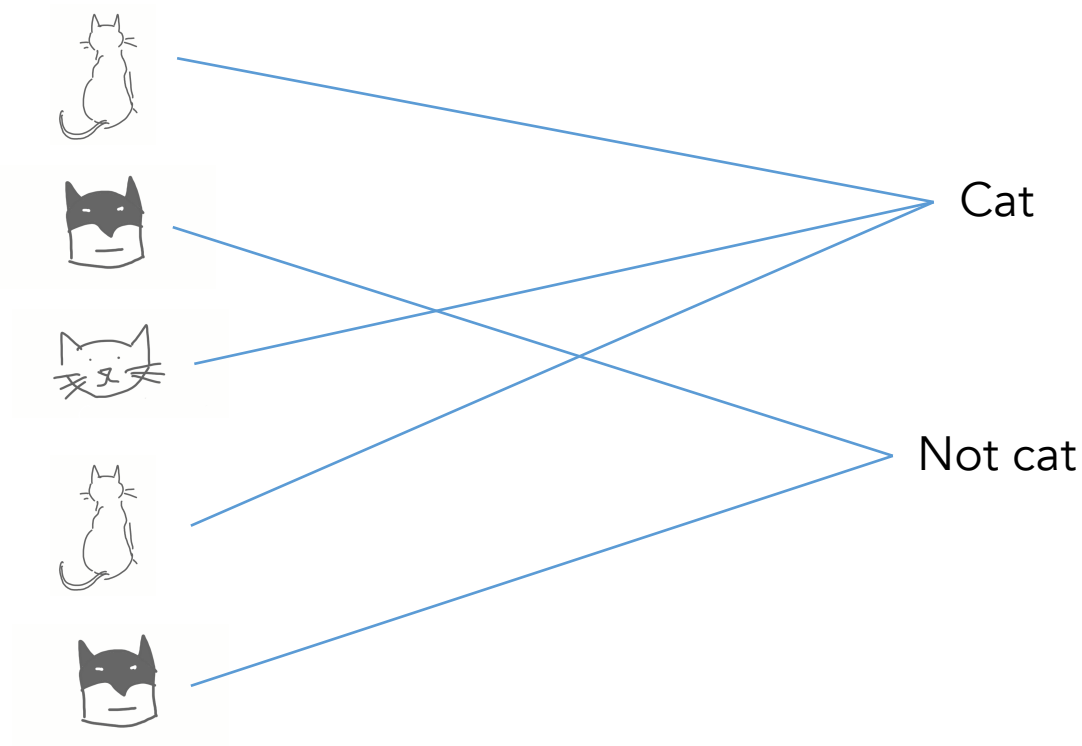

COMMERCE DATA SERVICE

large datasets

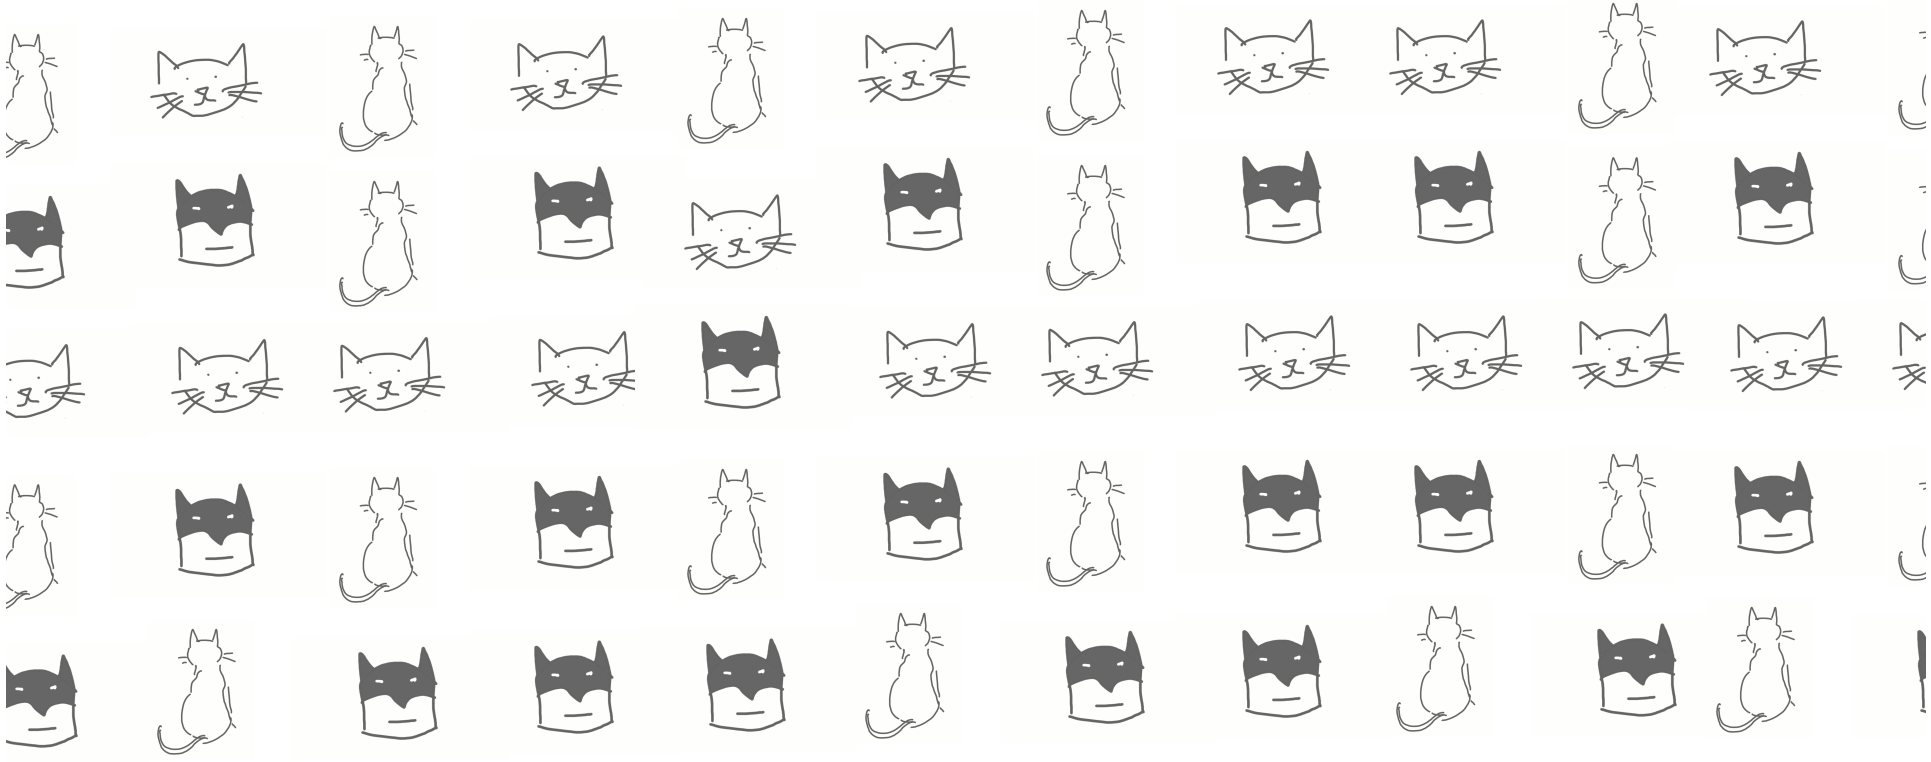

COMMERCE DATA SERVICE

try out a deep learning doodle-classifier

go to **quickdraw.withgoogle.com** on your phone

deep learning or a regular statistical model?

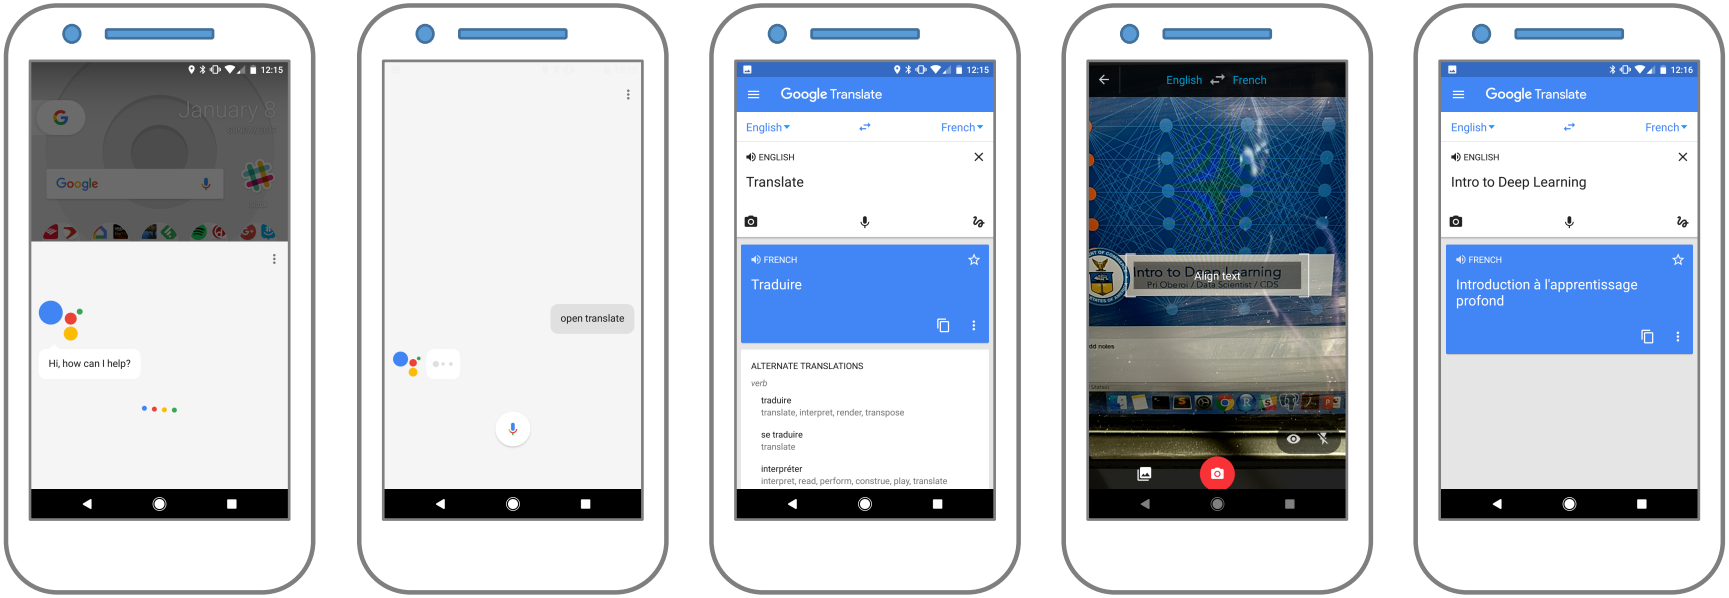

how does it work?

$y = f(x)$

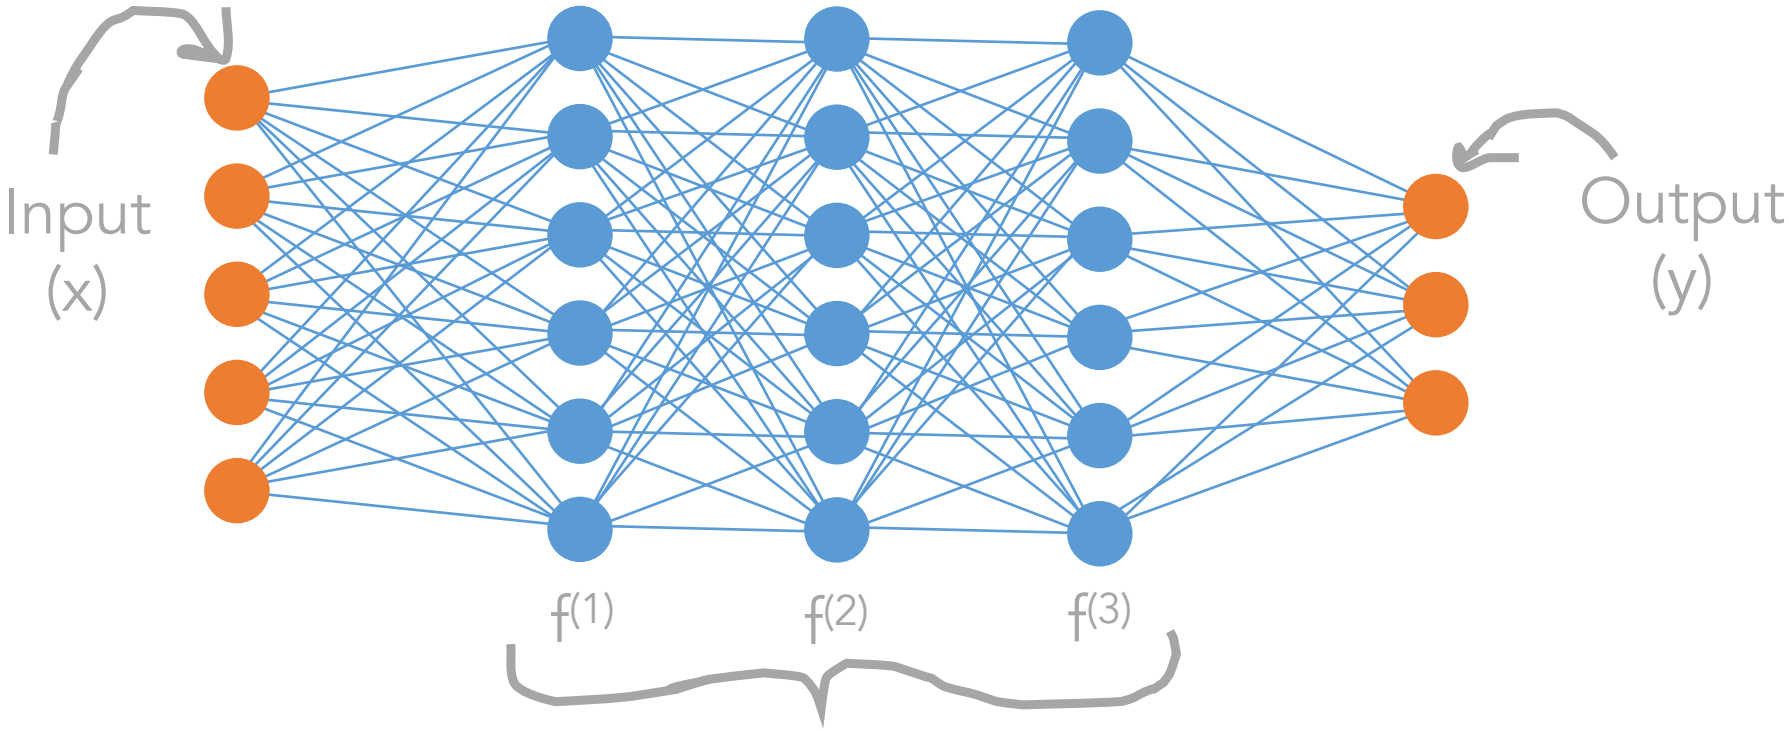

$f(x) = f^{(3)}(f^{(2)}(f^{(1)}(x)))$

COMMERCE DATA SERVICE

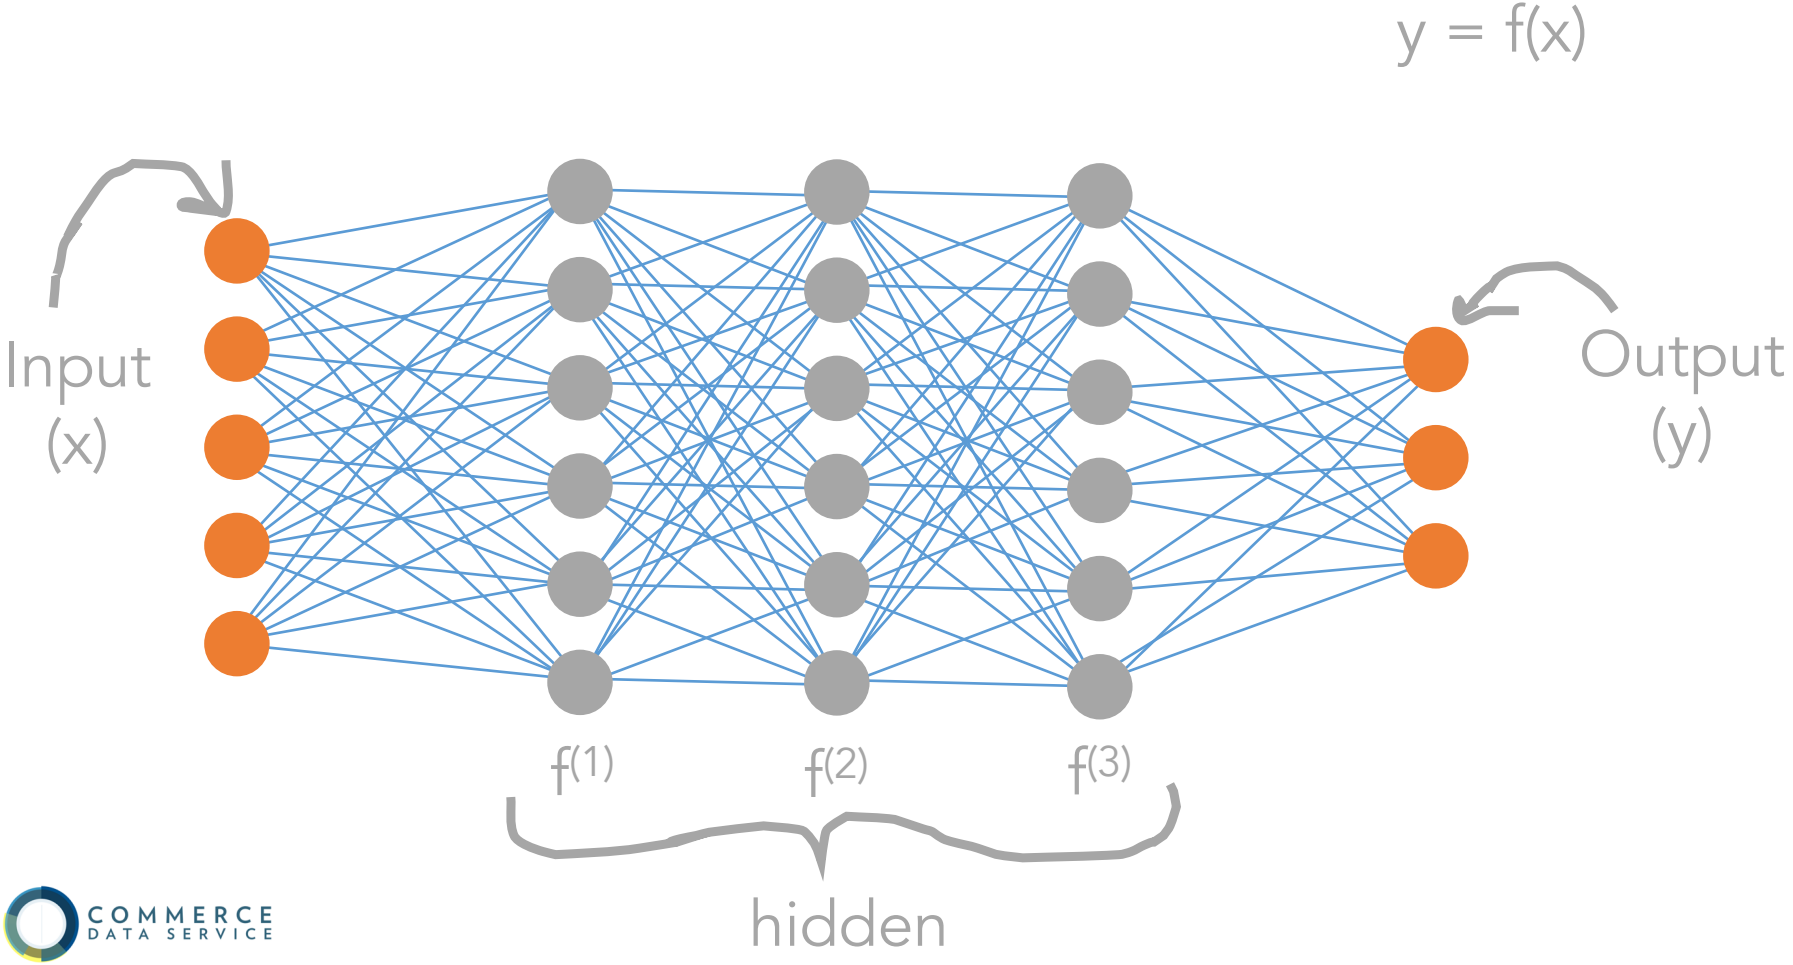

COMMERCE DATA SERVICE

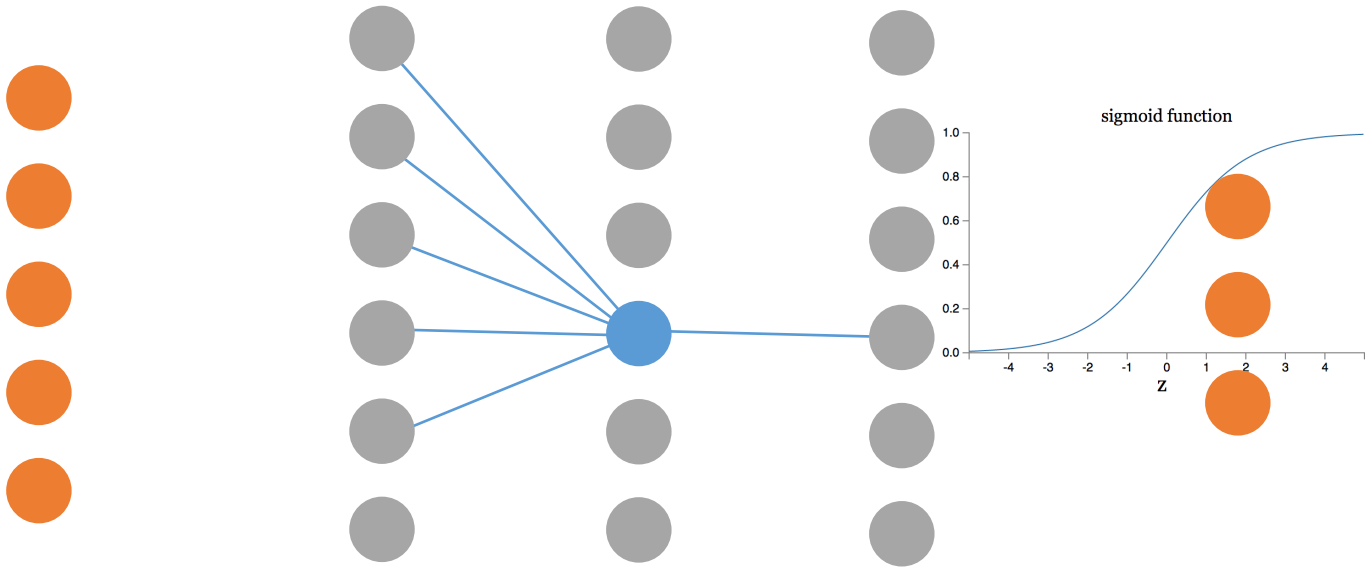

sigmoid neurons

COMMERCE DATA SERVICE

types of neural nets

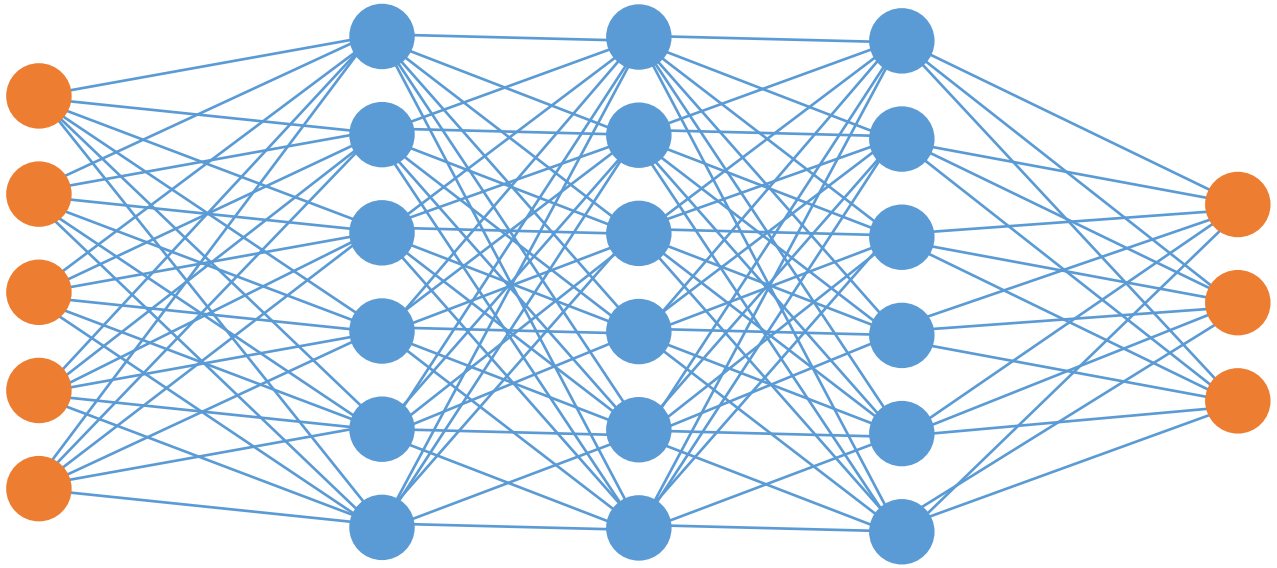

COMMERCE DATA SERVICE

Feedforward neural networks
(multilayer perceptions or MLPs)

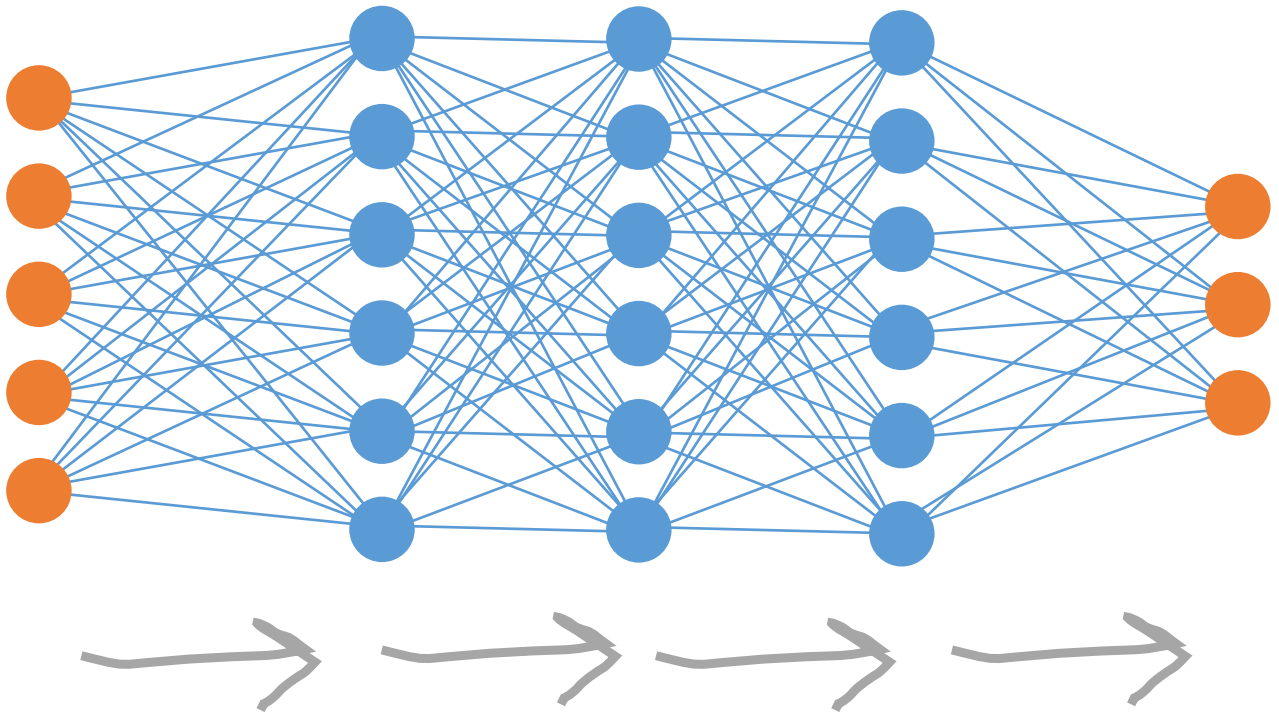

COMMERCE
DATA SERVICE

Recurrent neural nets

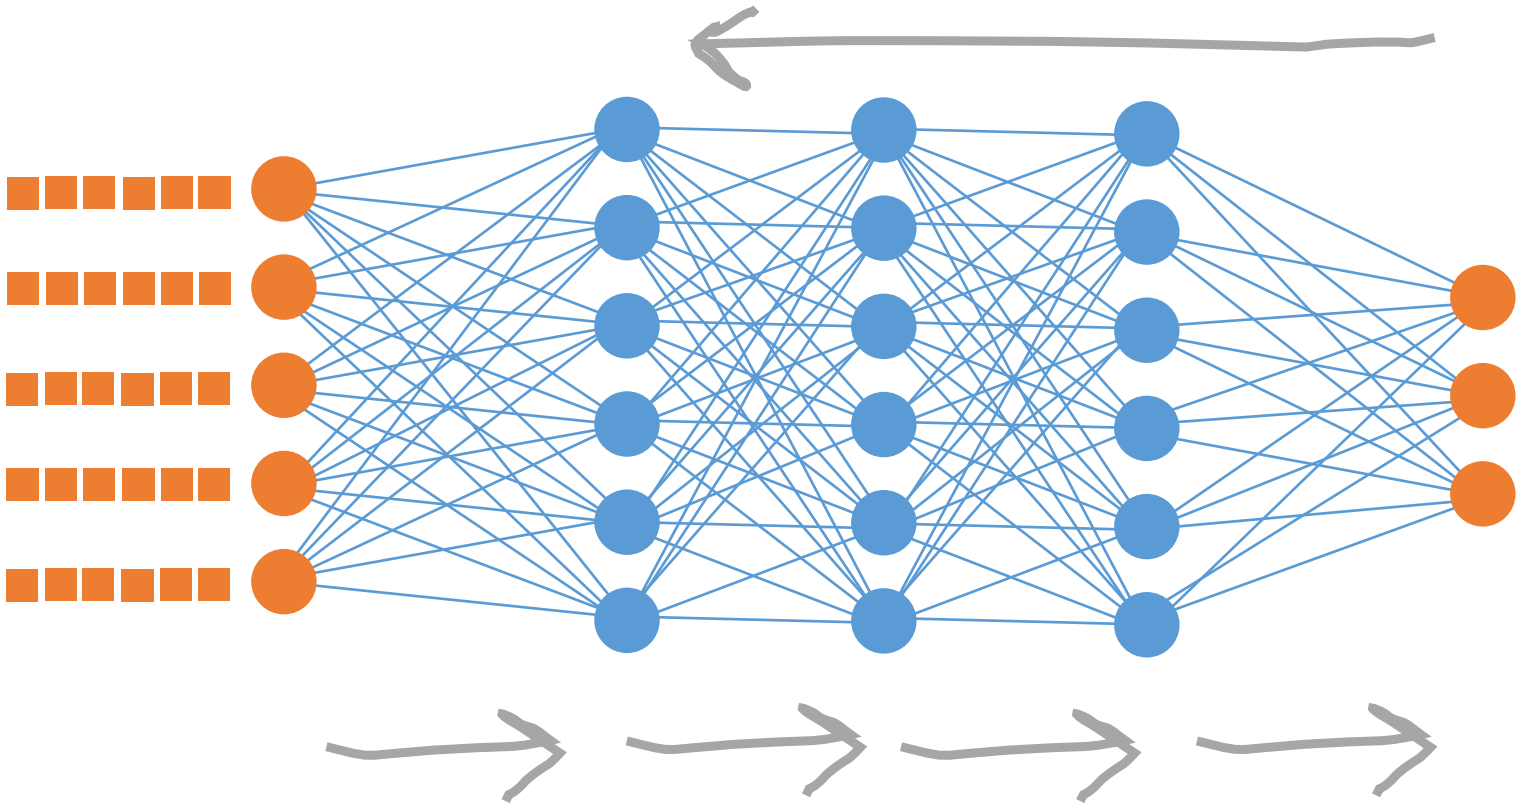

COMMERCE DATA SERVICE

convolutional neural nets

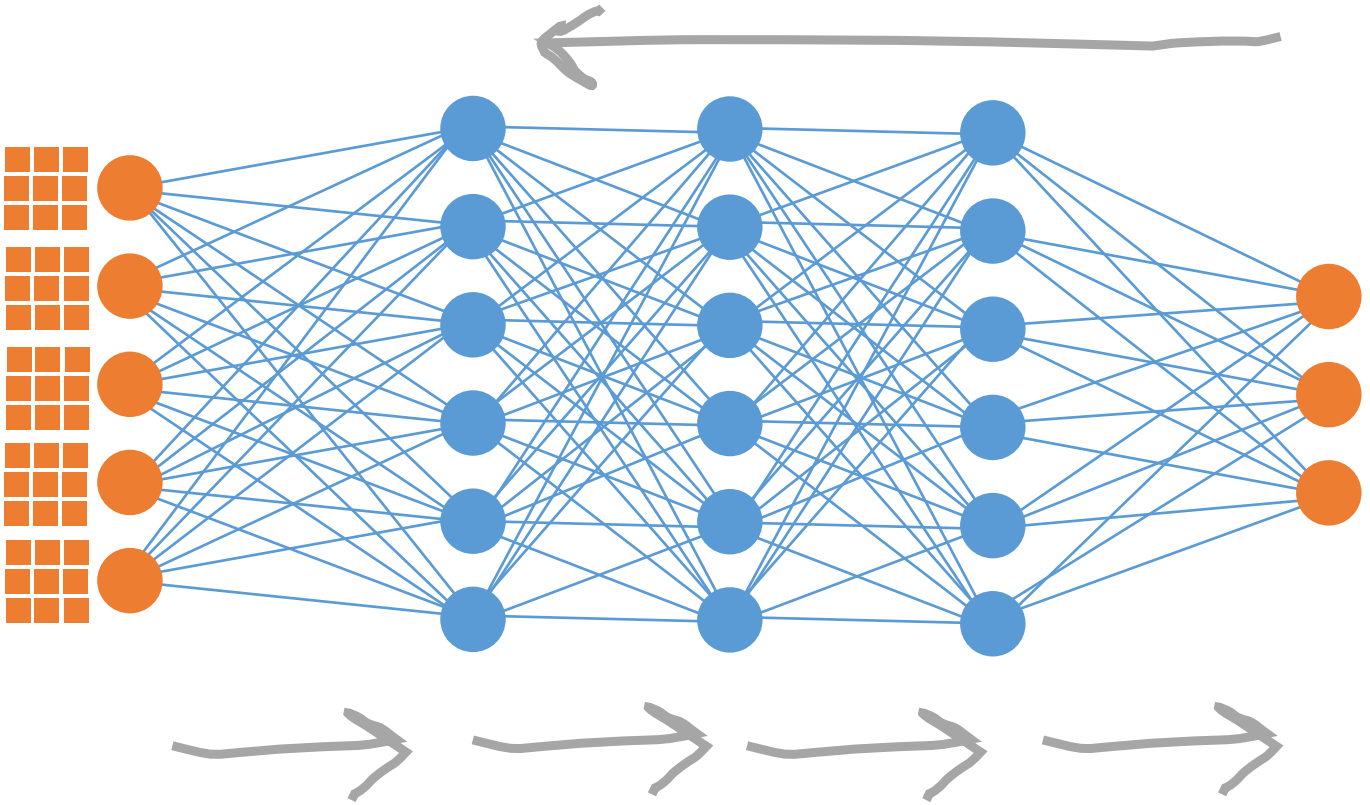

COMMERCE
DATA SERVICE

frameworks

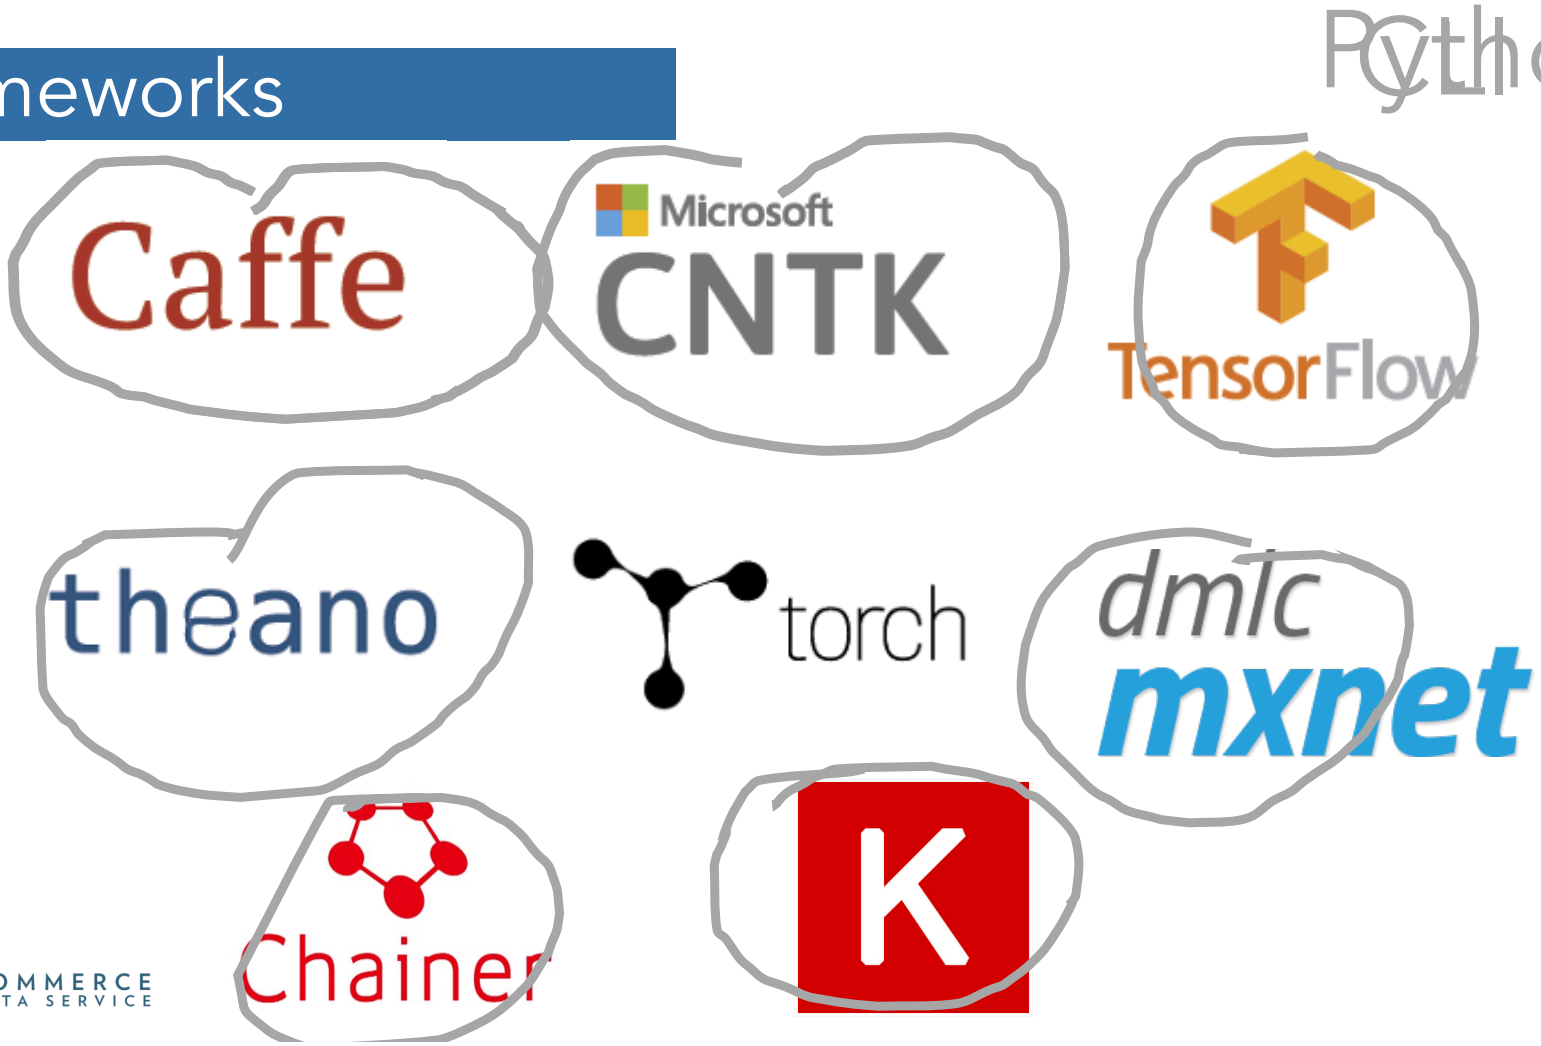

COMMERCE DATA SERVICE

example

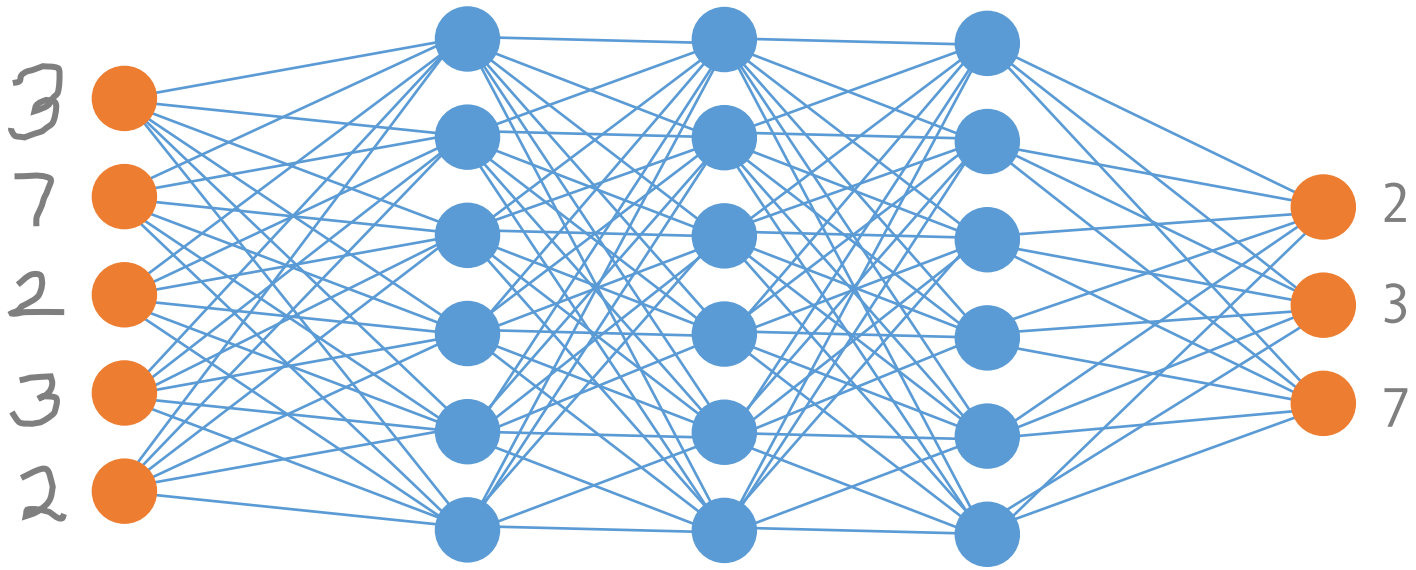

COMMERCE
DATA SERVICE

70,000 labeled handwritten images

55,000 data points of training data

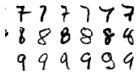

5,000 points of validation data

10,000 points of test data

COMMERCE DATA SERVICE

layers

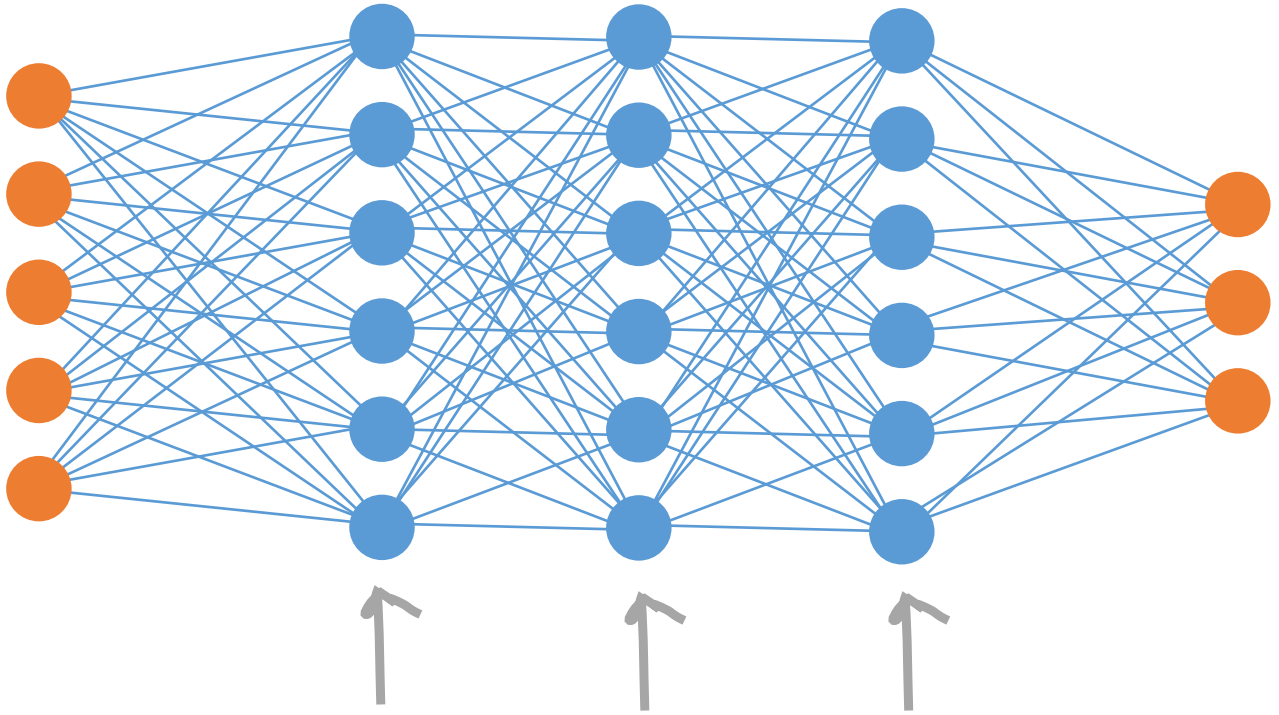

COMMERCE DATA SERVICE

a common CNN

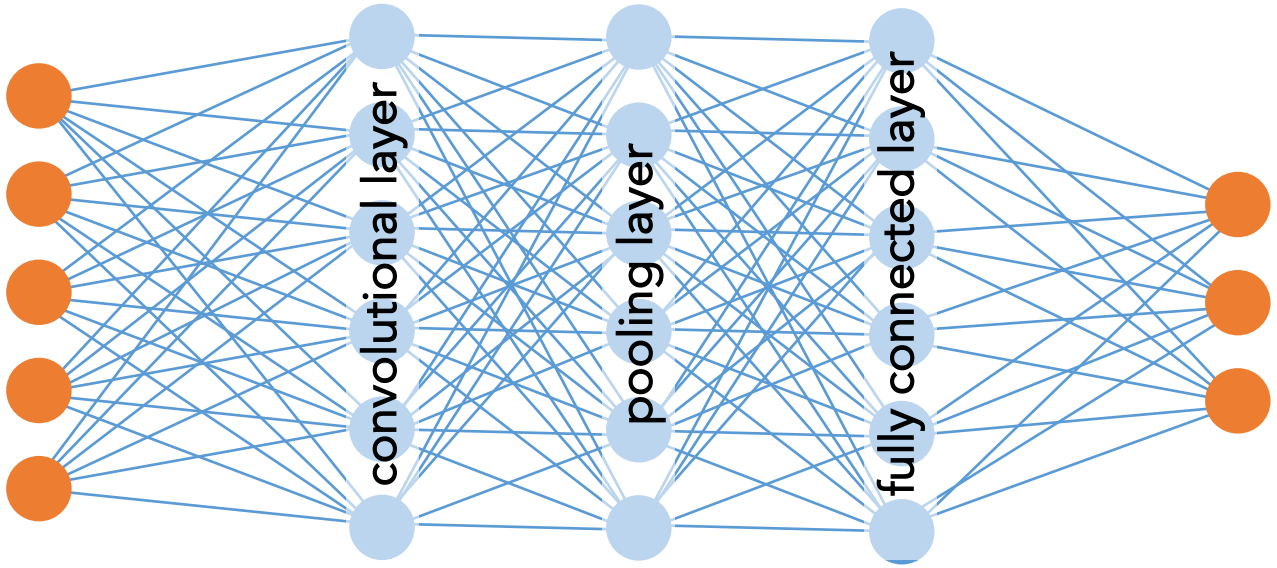

COMMERCE DATA SERVICE

convolutional layer

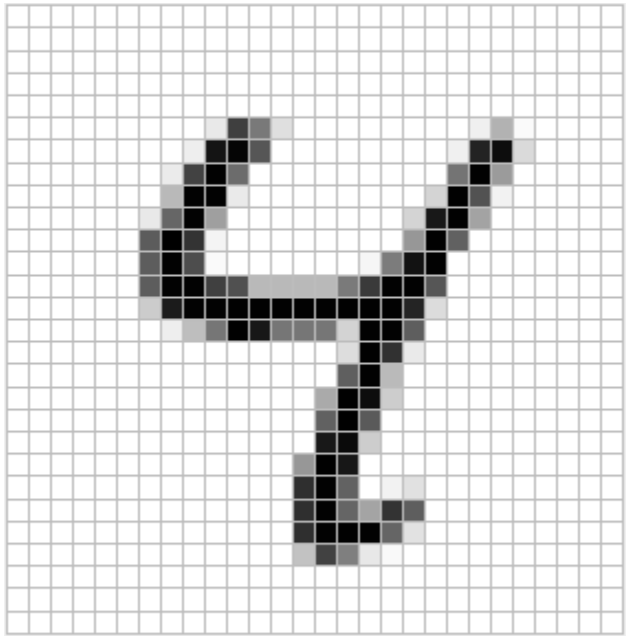

local receptive field
(5x5)

COMMERCE
DATA SERVICE

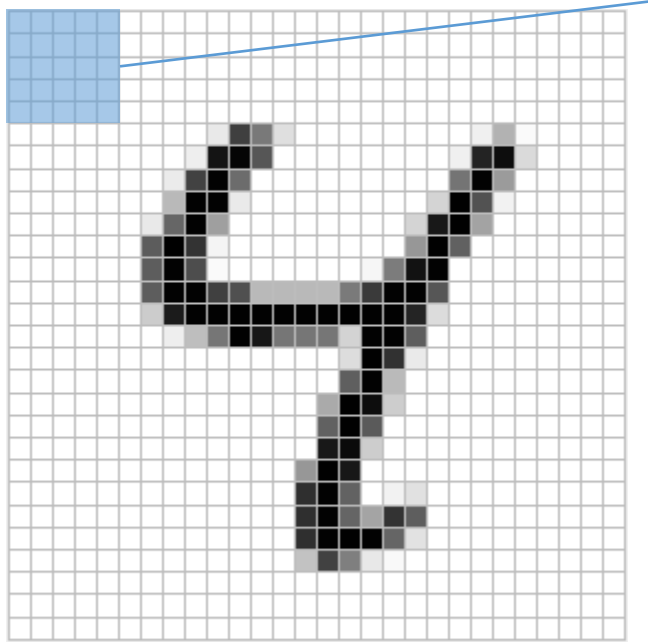

one neuron in the hidden convolutional layer

COMMERCE DATA SERVICE

---> slide one

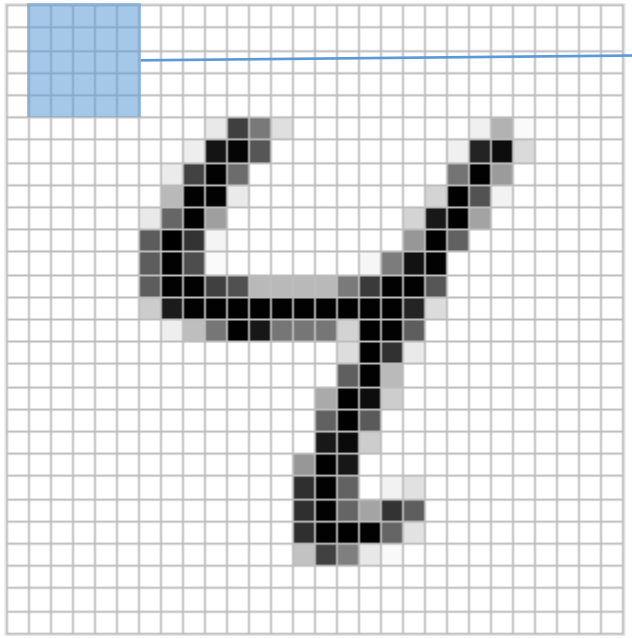

COMMERCE
DATA SERVICE

---> slide one

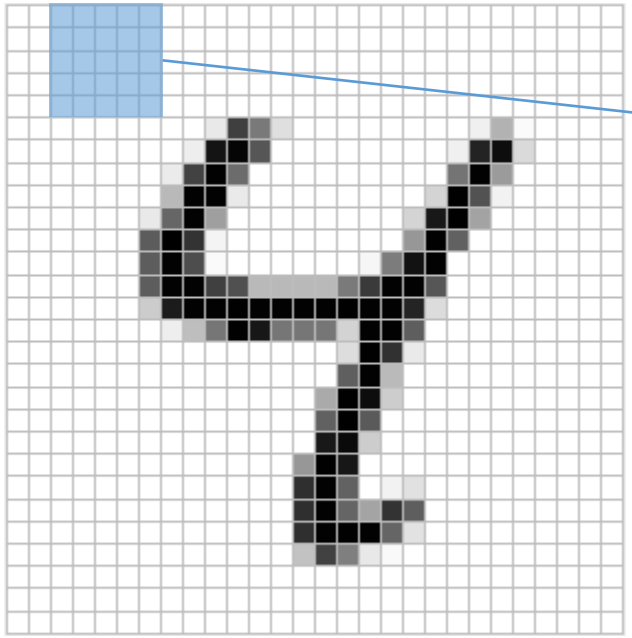

ultimately resulting in 24 neurons

COMMERCE DATA SERVICE

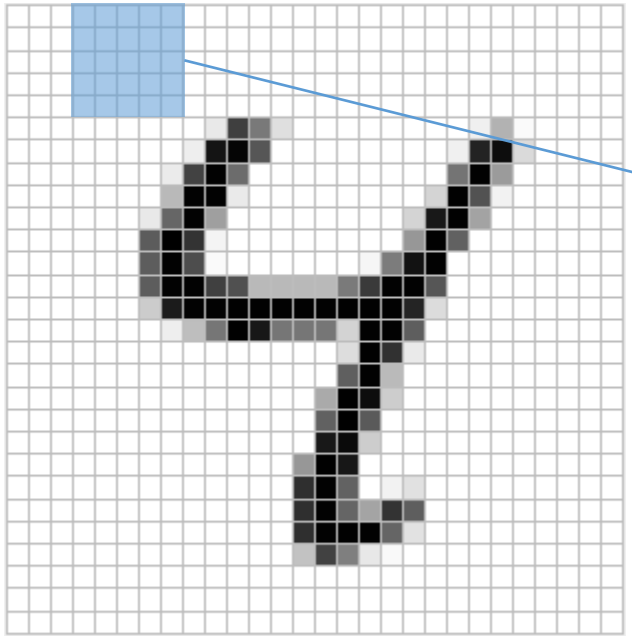

each neuron has a bias and 5 x 5 weights for the local receptive field

COMMERCE DATA SERVICE

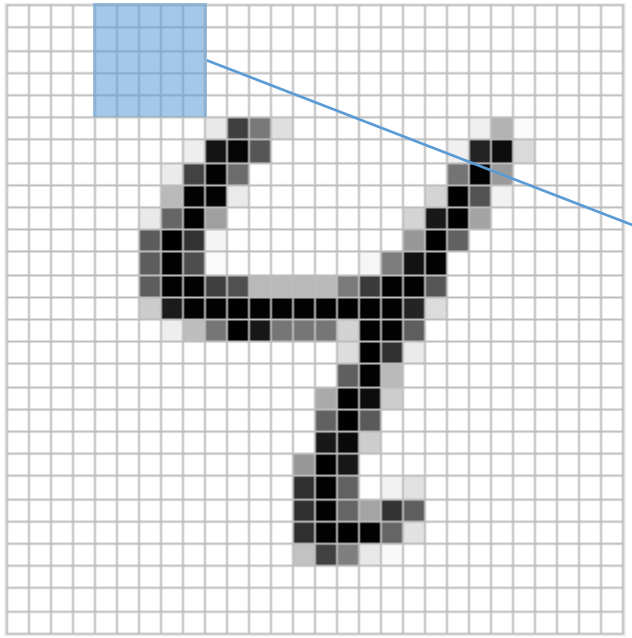

if you look at those weight visually, with larger weights as a darker pixel, you might see something like this:

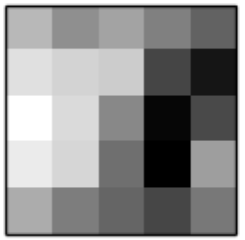

COMMERCE DATA SERVICE

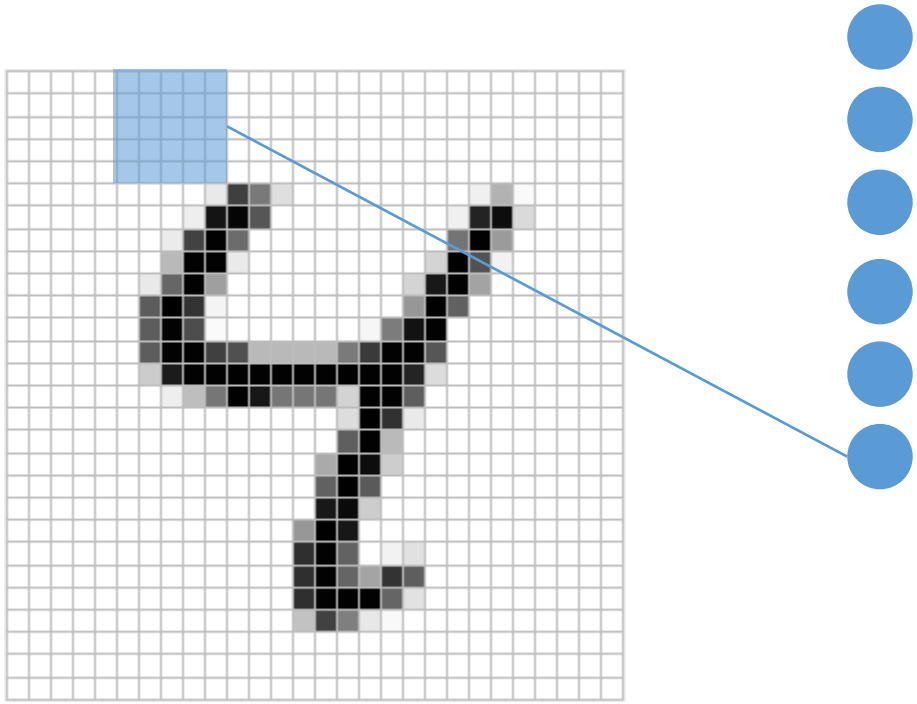

the weights and bias are shared
across all 24 x 24 neurons

COMMERCE
DATA SERVICE

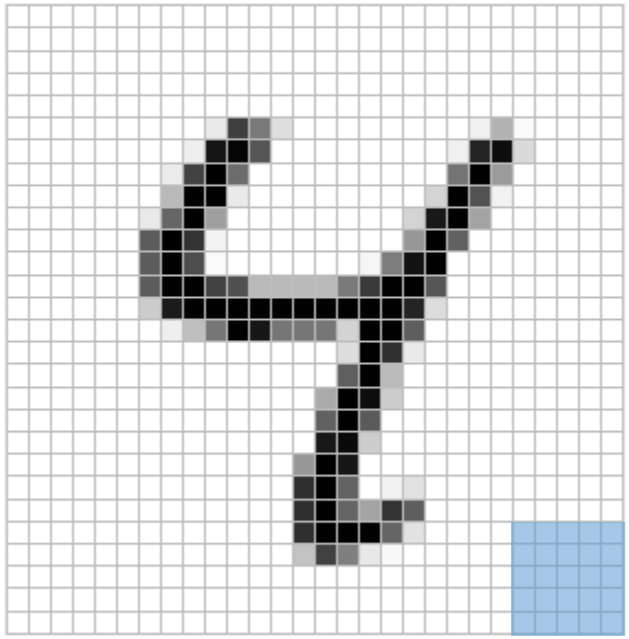

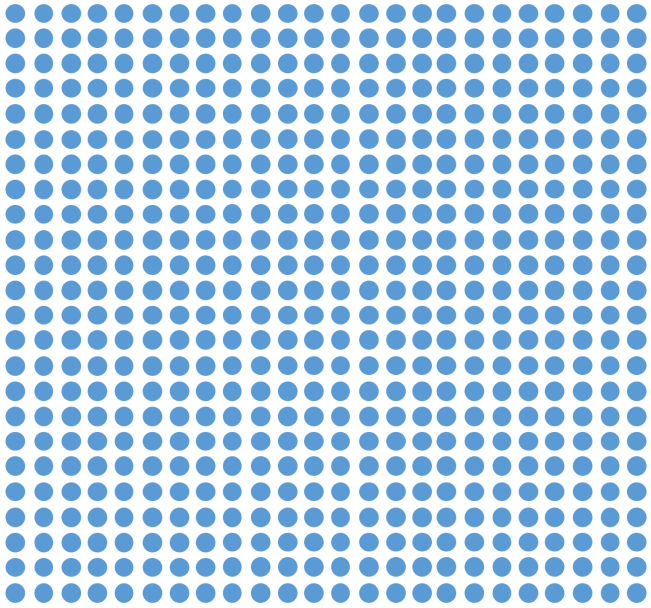

map between input layer and
hidden layer is called a feature map.

COMMERCE
DATA SERVICE

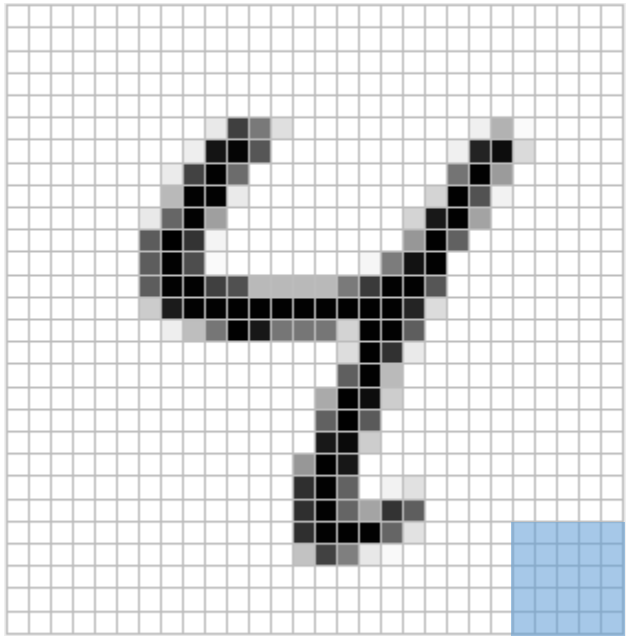

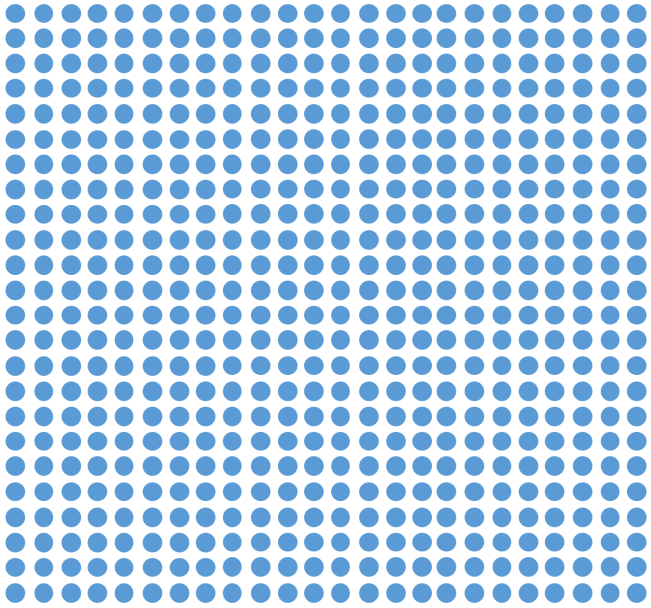

a feature map detects one kind of localized structure

COMMERCE DATA SERVICE

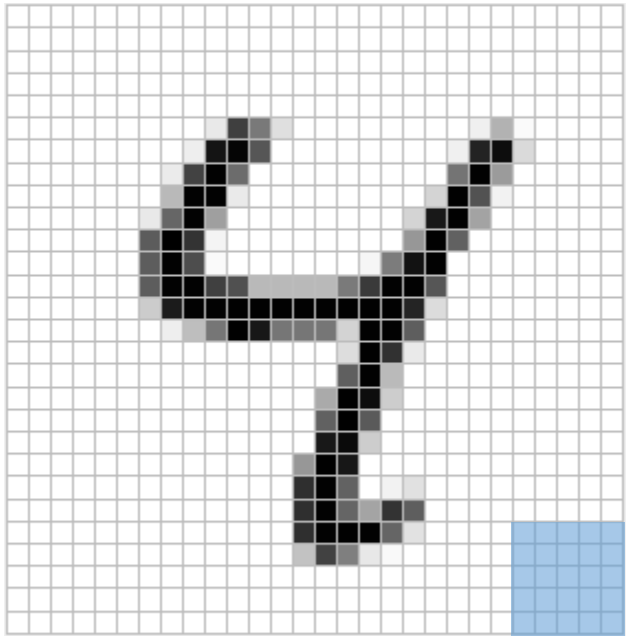

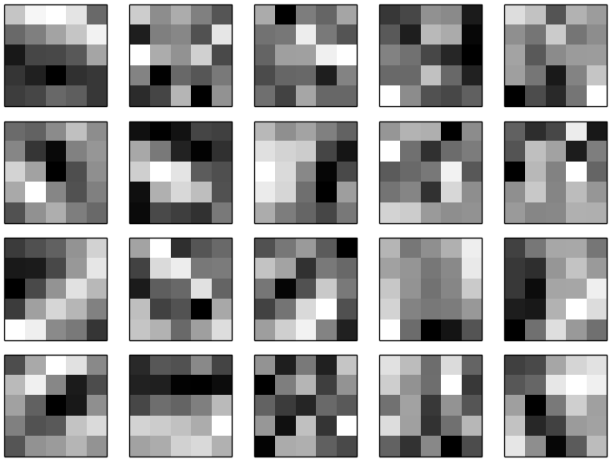

a complete convolutional layer consists of several different feature maps

COMMERCE DATA SERVICE

pooling layer

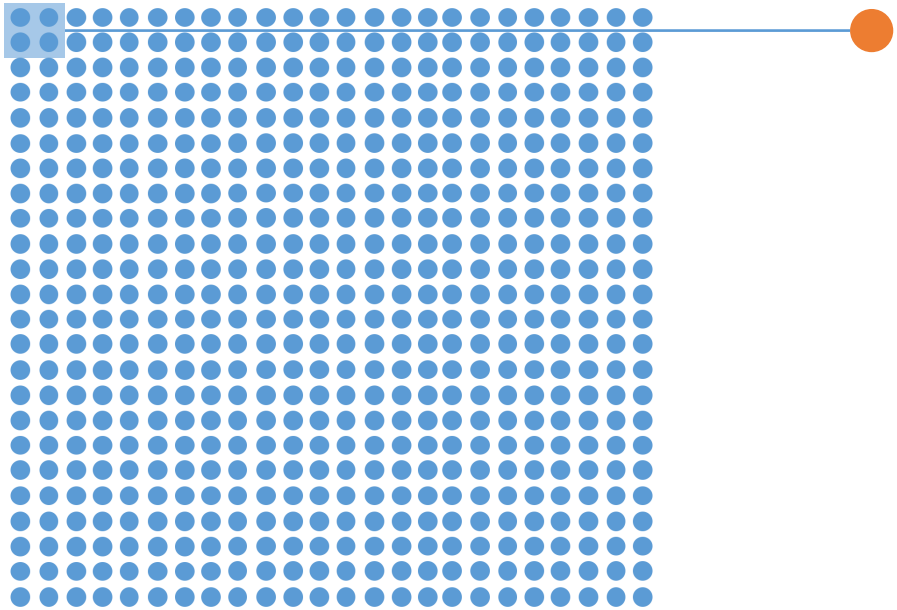

one feature map of the
convolutional layer

COMMERCE
DATA SERVICE

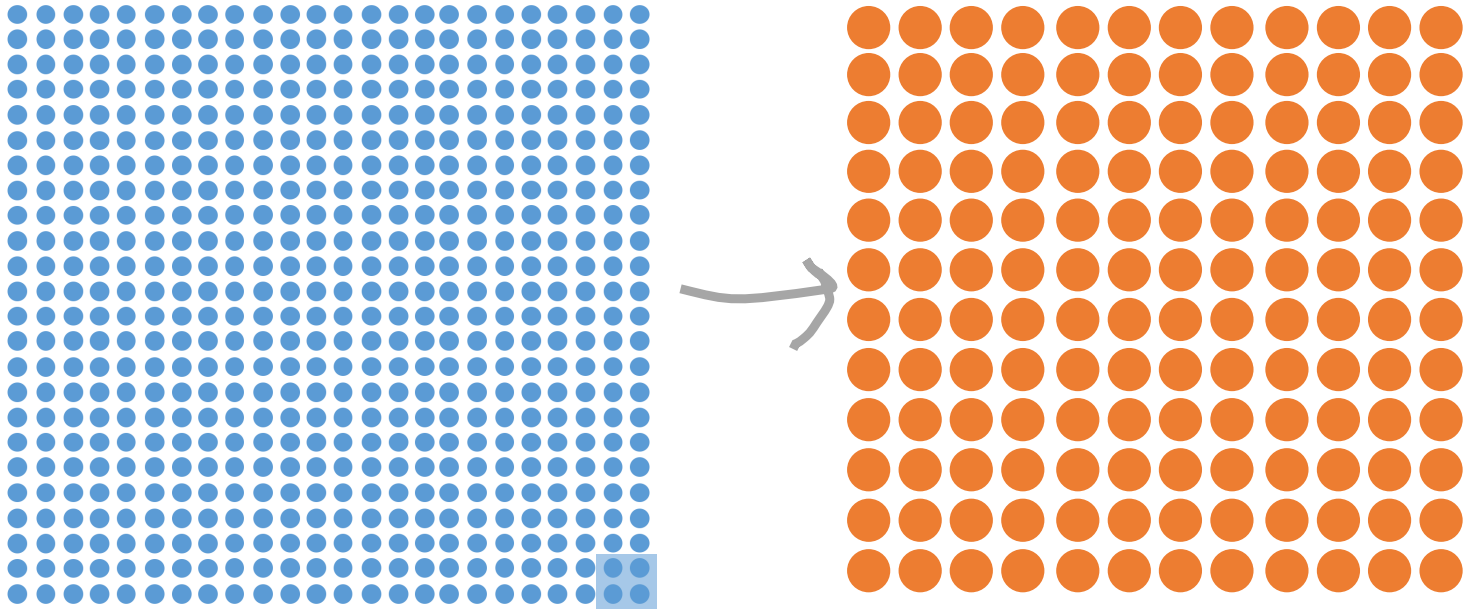

one feature map of the convolutional layer

one max-pooling layer

COMMERCE DATA SERVICE

fully-connected layer

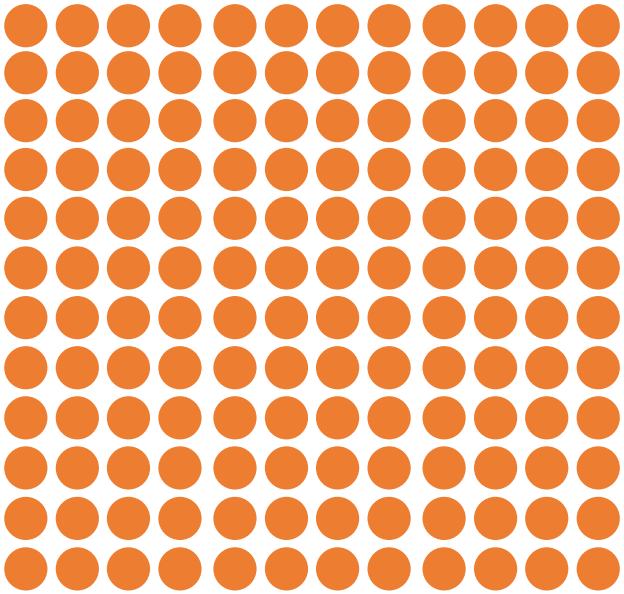

one max-pooling layer

one feature map of the convolutional layer

COMMERCE DATA SERVICE

fully-connected layer

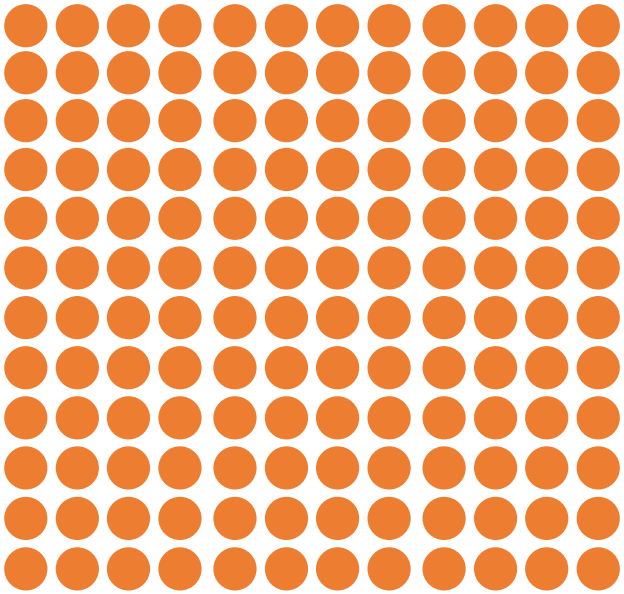

one max-pooling layer

one feature map of the convolutional layer

COMMERCE DATA SERVICE

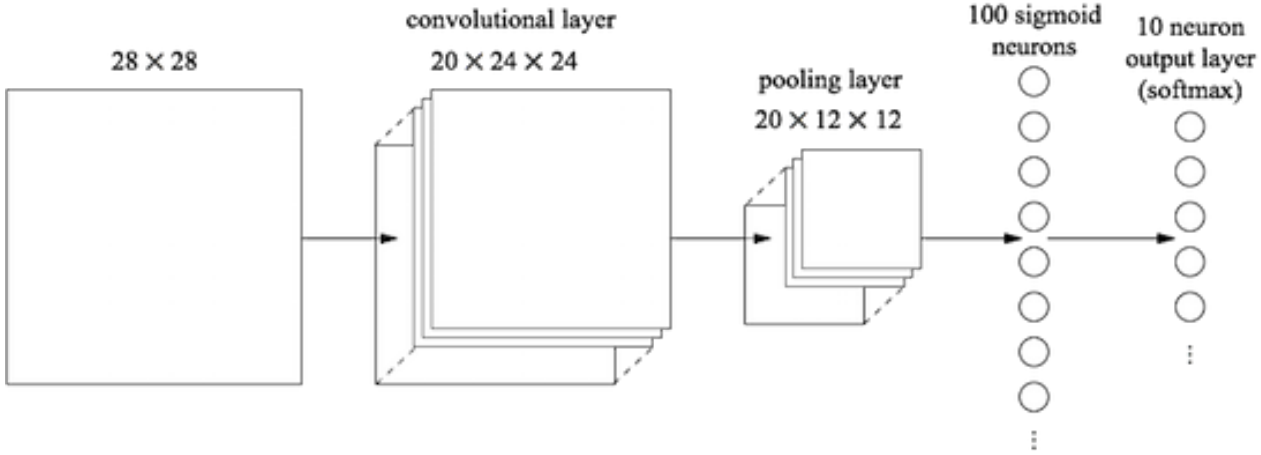

one image of 28 x28 pixels

convolutional layer of 20 feature maps (5 x 5 local receptive field)

pooling layer for each feature map

fully connected layer makes a prediction

COMMERCE DATA SERVICE

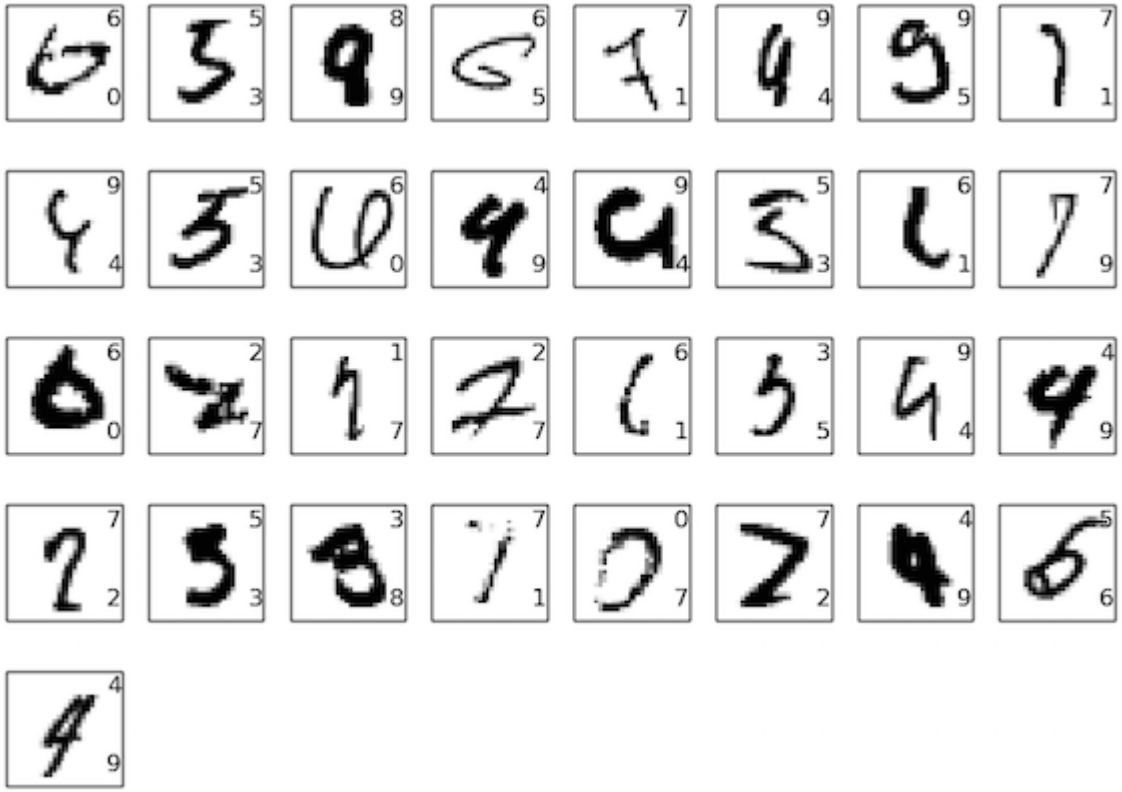

COMMERCE DATA SERVICE

optimization

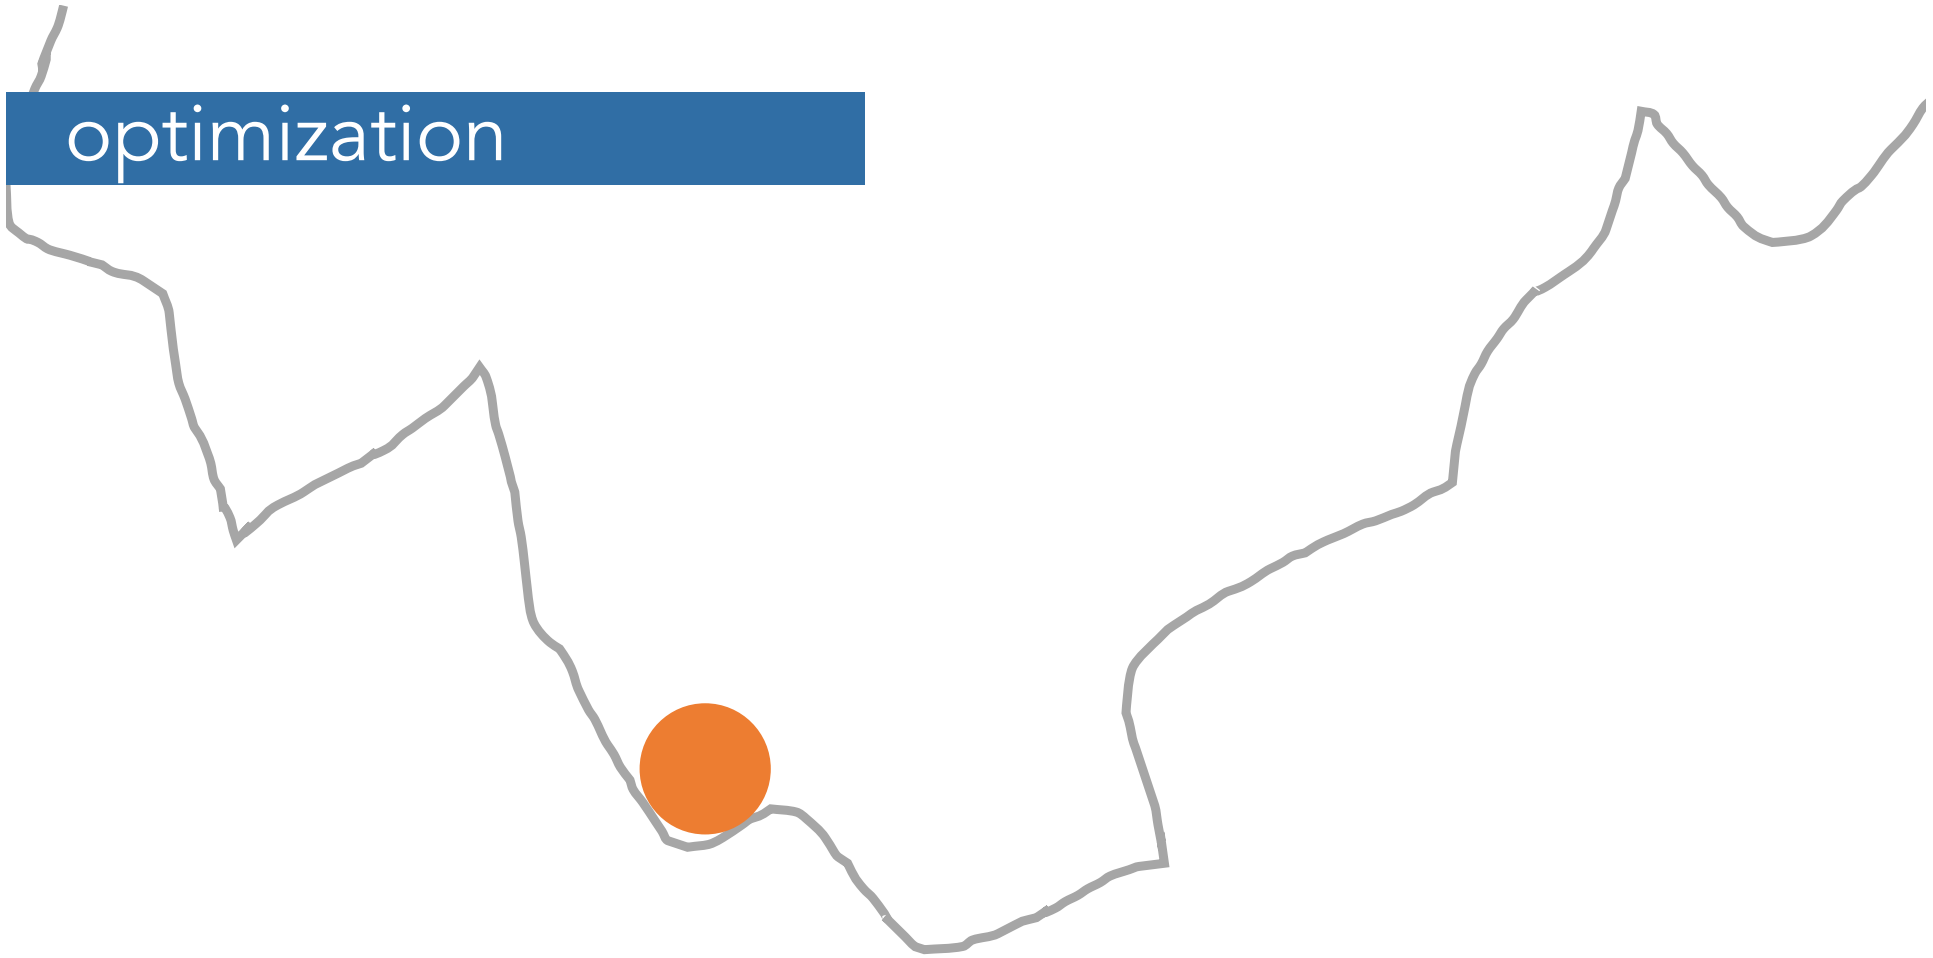

COMMERCE
DATA SERVICE

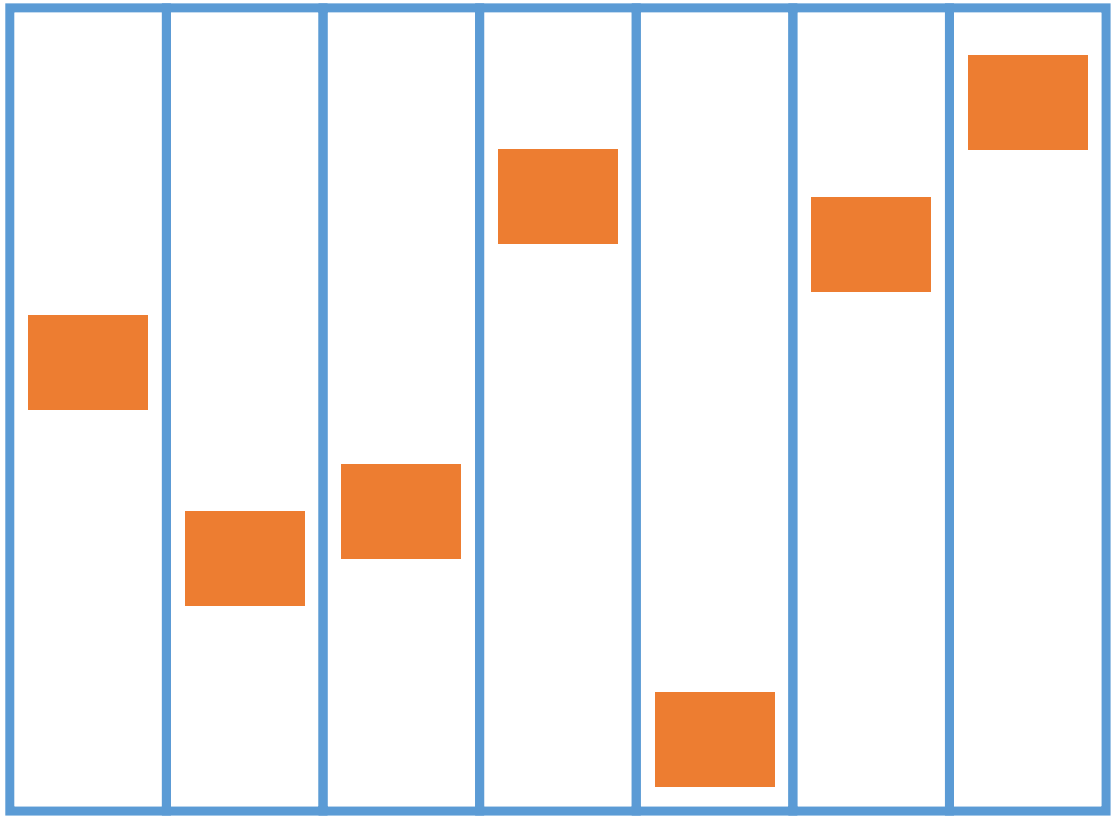

COMMERCE
DATA SERVICE

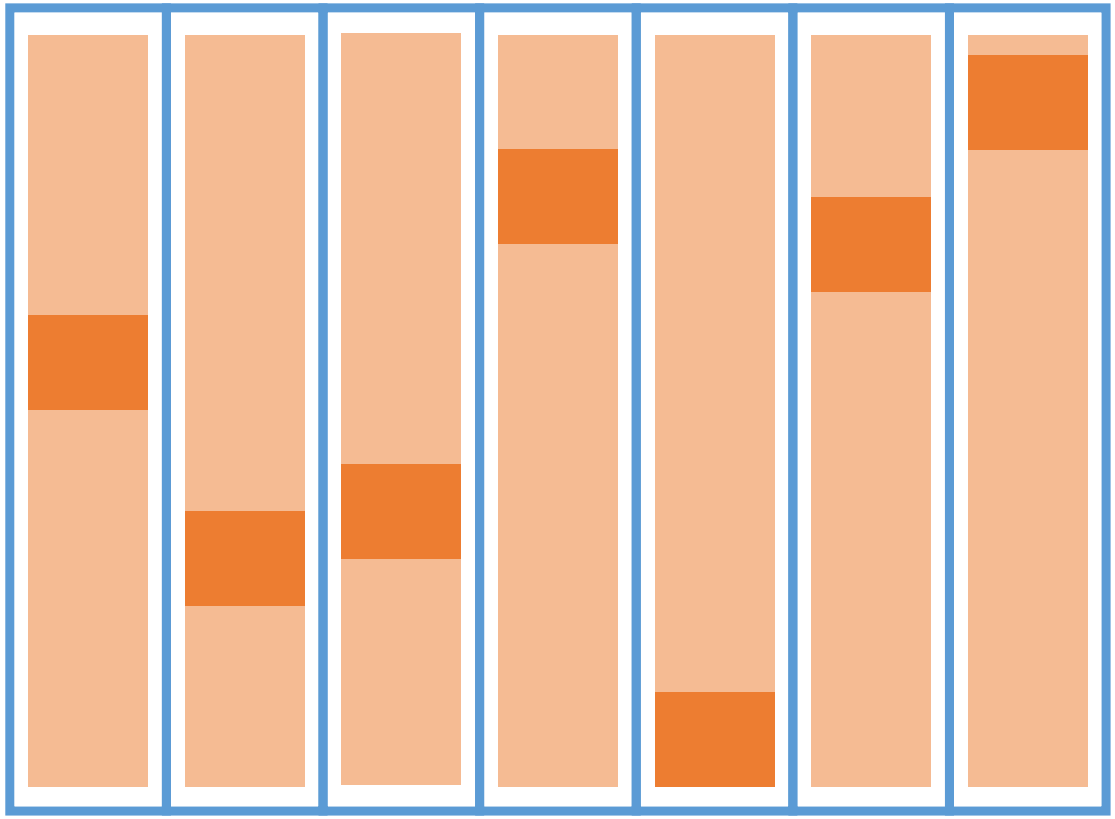

COMMERCE
DATA SERVICE

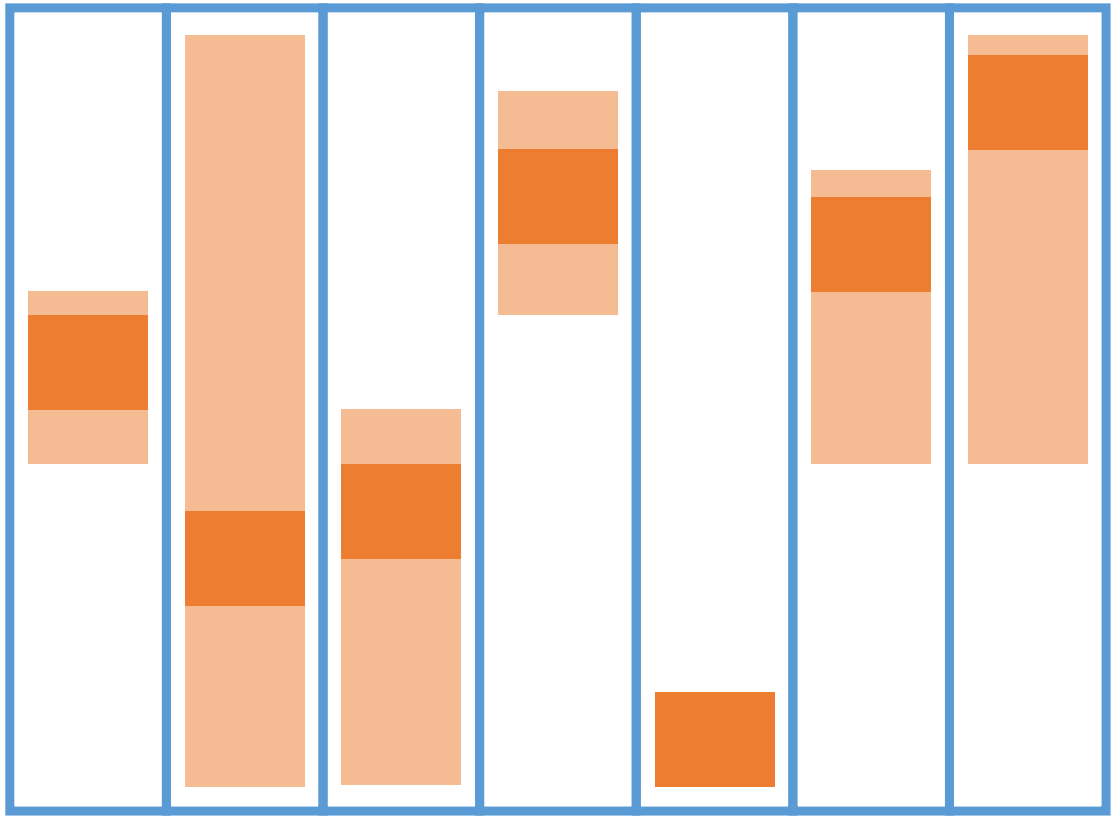

COMMERCE
DATA SERVICE

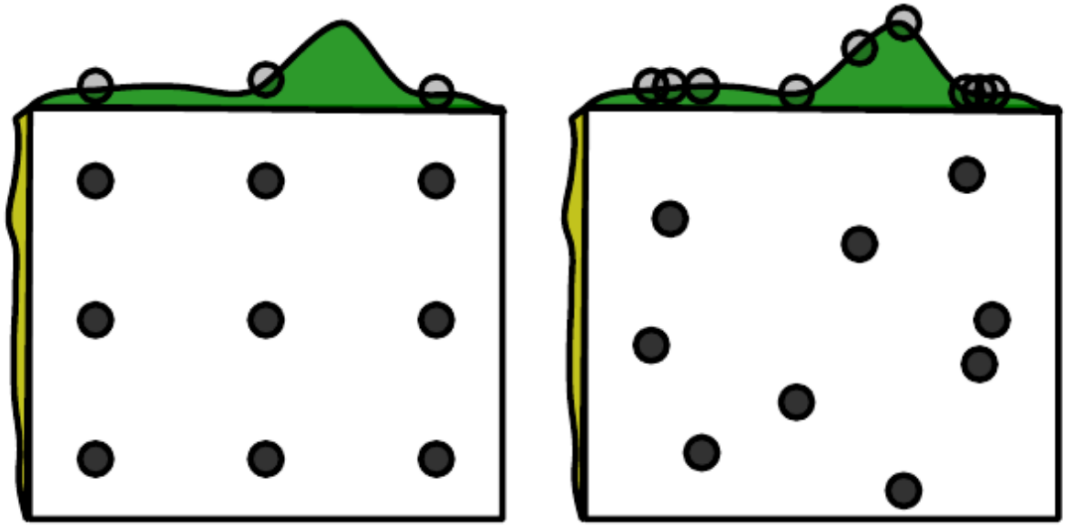

Grid Random

COMMERCE DATA SERVICE

performance

70,000 labeled handwritten images

55,000 data points of training data

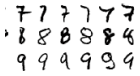

5,000 points of validation data

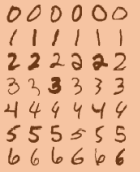

10,000 points of test data

COMMERCE DATA SERVICE

error rate

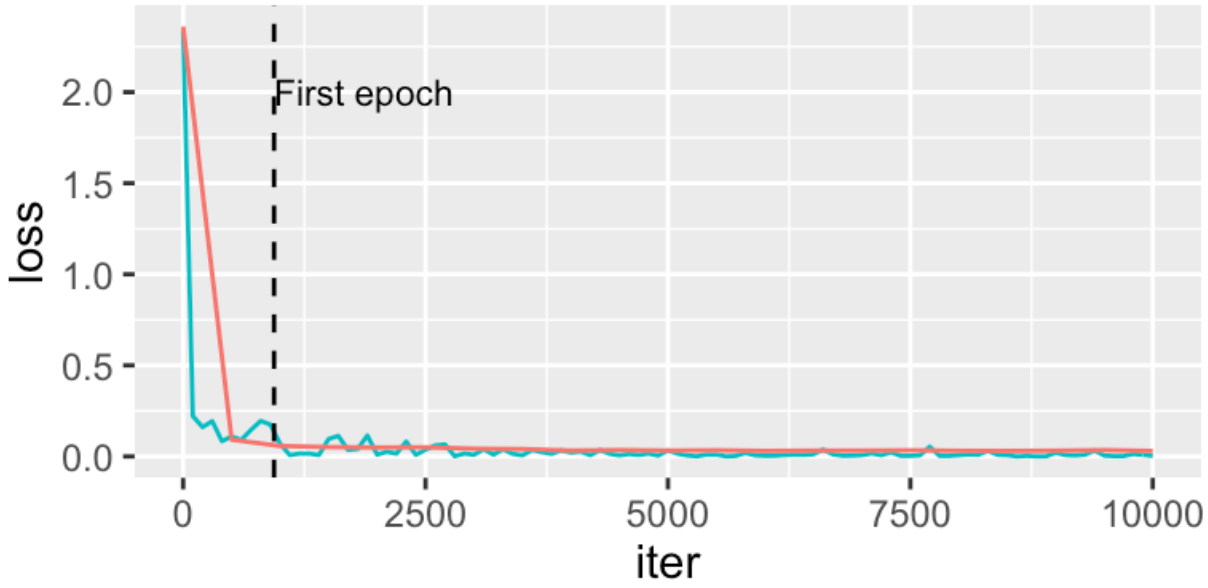

COMMERCE DATA SERVICE

false positive rate or false negative rate

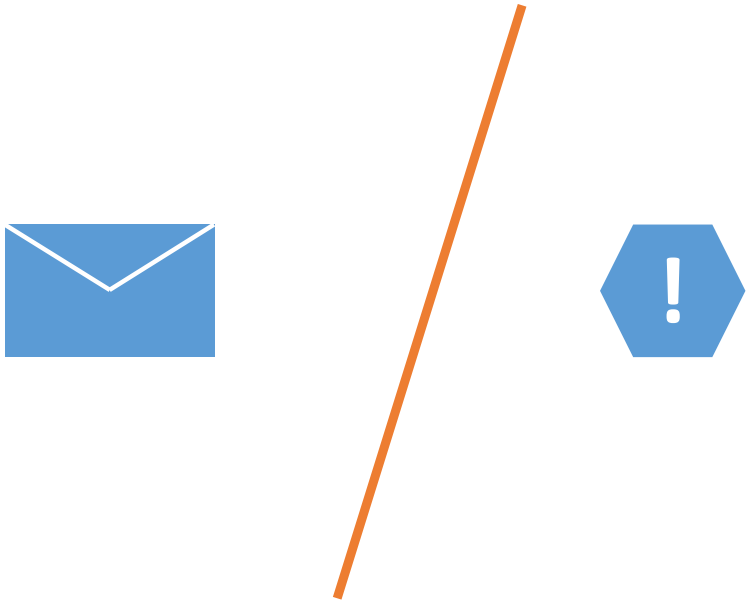

COMMERCE
DATA SERVICE

precision and recall

COMMERCE
DATA SERVICE

tutorials, resources

MNIST TensorFlow tutorial:
https://www.tensorflow.org/tutorials/mnist/pros/
Deep Learning by Ian Goodfellow, Yoshua Bengio, and Aaron Courville:
http://www.deeplearningbook.org/
Machine Learning is Fun! By Adam Geigey:
https://medium.com/@ageitgey/machine-learning-is-fun-part-3-deep-learning-and-convolutional-neural-networks-f40359318721
History of Deep Learning by John Kaufhold:
https://drive.google.com/file/d/0B3aXKp9bt6OXQU5lU0lmOE1ZZjA/edit
Deep Learning talk by Geoffery Hinton:
https://www.youtube.com/watch?v=VhmE_UXDOGs
Intro to Machine Learning with H2O and Python by Erin LeDell:
https://www.youtube.com/watch?v=rAYS6byutlA
Deep Learning Frameworks list by NVIDIA:
https://developer.nvidia.com/deep-learning-frameworks
Neural Networks and Deep Learning by Michael Nielsen:
http://neuralnetworksanddeeplearning.com/

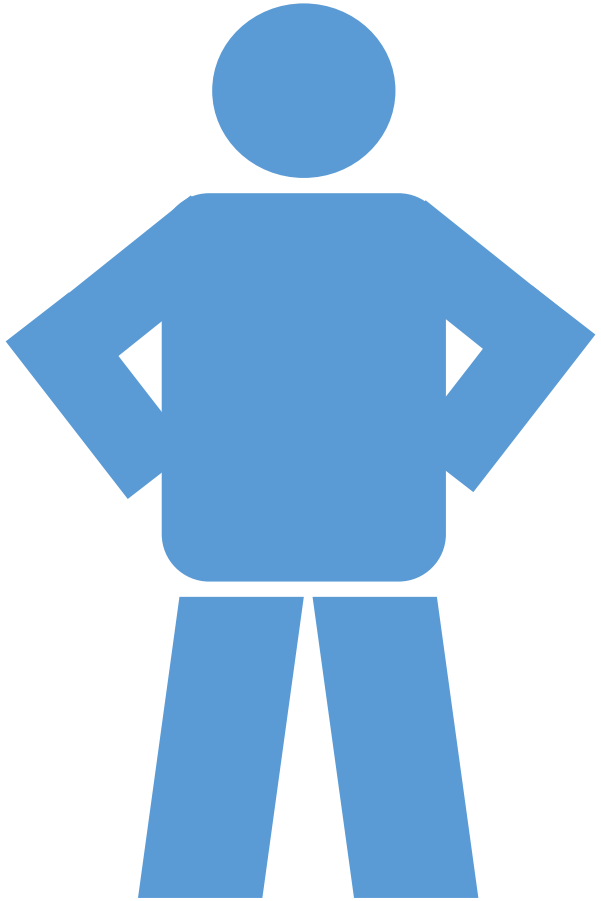

COMMERCE DATA SERVICE

Pri Oberoi, Data Scientist
poberoi@doc.gov