APPLIED **DEEP LEARNING**

PROF ALEXIEI DINGLI

UNIVERSITY OF MALTA
L-Università ta' Malta

TECH NEWS

ROBERT MCMILLAN BUSINESS 03.13.13 6:30 AM

GOOGLE HIRES BRAINS THAT HELPED SUPERCHARGE MACHINE LEARNING

Google buying AI startup DeepMind for a reported $400 million

By Kwame Opam on January 26, 2014 08:03 pm Email @kwameopam

TECH NEWS

HOW TO DO IT?

ATARI
TWO MINUTE
PAPERS
GOOGLE DEEPMIND'S DEEP Q-LEARNING &
SUPERHUMAN ATARI GAMEPLAYS
Disclaimer: I was not part of this research project. I am merely providing commentary on this work.

TECH NEWS

NYU "Deep Learning" Professor LeCun Will Head Facebook's New Artificial Intelligence Lab

Posted Dec 9, 2013 by Josh Constine (@joshconstine)

856 SHARES

By teaching a computer to think, Facebook hopes to better understand how its users do too. So today the company announced that one of the world's leading deep learning and machine learning scientists, NYU's Professor Yann LeCun, will lead its new artificial intelligence laboratory.

APPLICATIONS

TECH NEWS


TECH 2/19/2015 @ 1:06PM | 4,873 views

Microsoft's Deep Learning Project Outperforms Humans In Image Recognition


DANIELA HERNANDEZ | BUSINESS | 07.14.14 | 12:00 PM

MICROSOFT CHALLENGES GOOGLE'S ARTIFICIAL BRAIN WITH 'PROJECT ADAM'

TECH NEWS


New Apple Hire Points to Self-Driving iCar
The latest hire from Apple helped create GPUs that were used to power the cameras and radar being used in self-driving car prototypes.
Driverless Transportation
Humanoid Robots
Science & Technology
CAR CONCEPT

NEURAL NETWORKS

- Interconnected set of nodes and edges
- Designed to perform complex tasks
- Deep Learning based upon neural networks

CLASSIFICATION

CLASSIFICATION

NEURAL NETWORKS

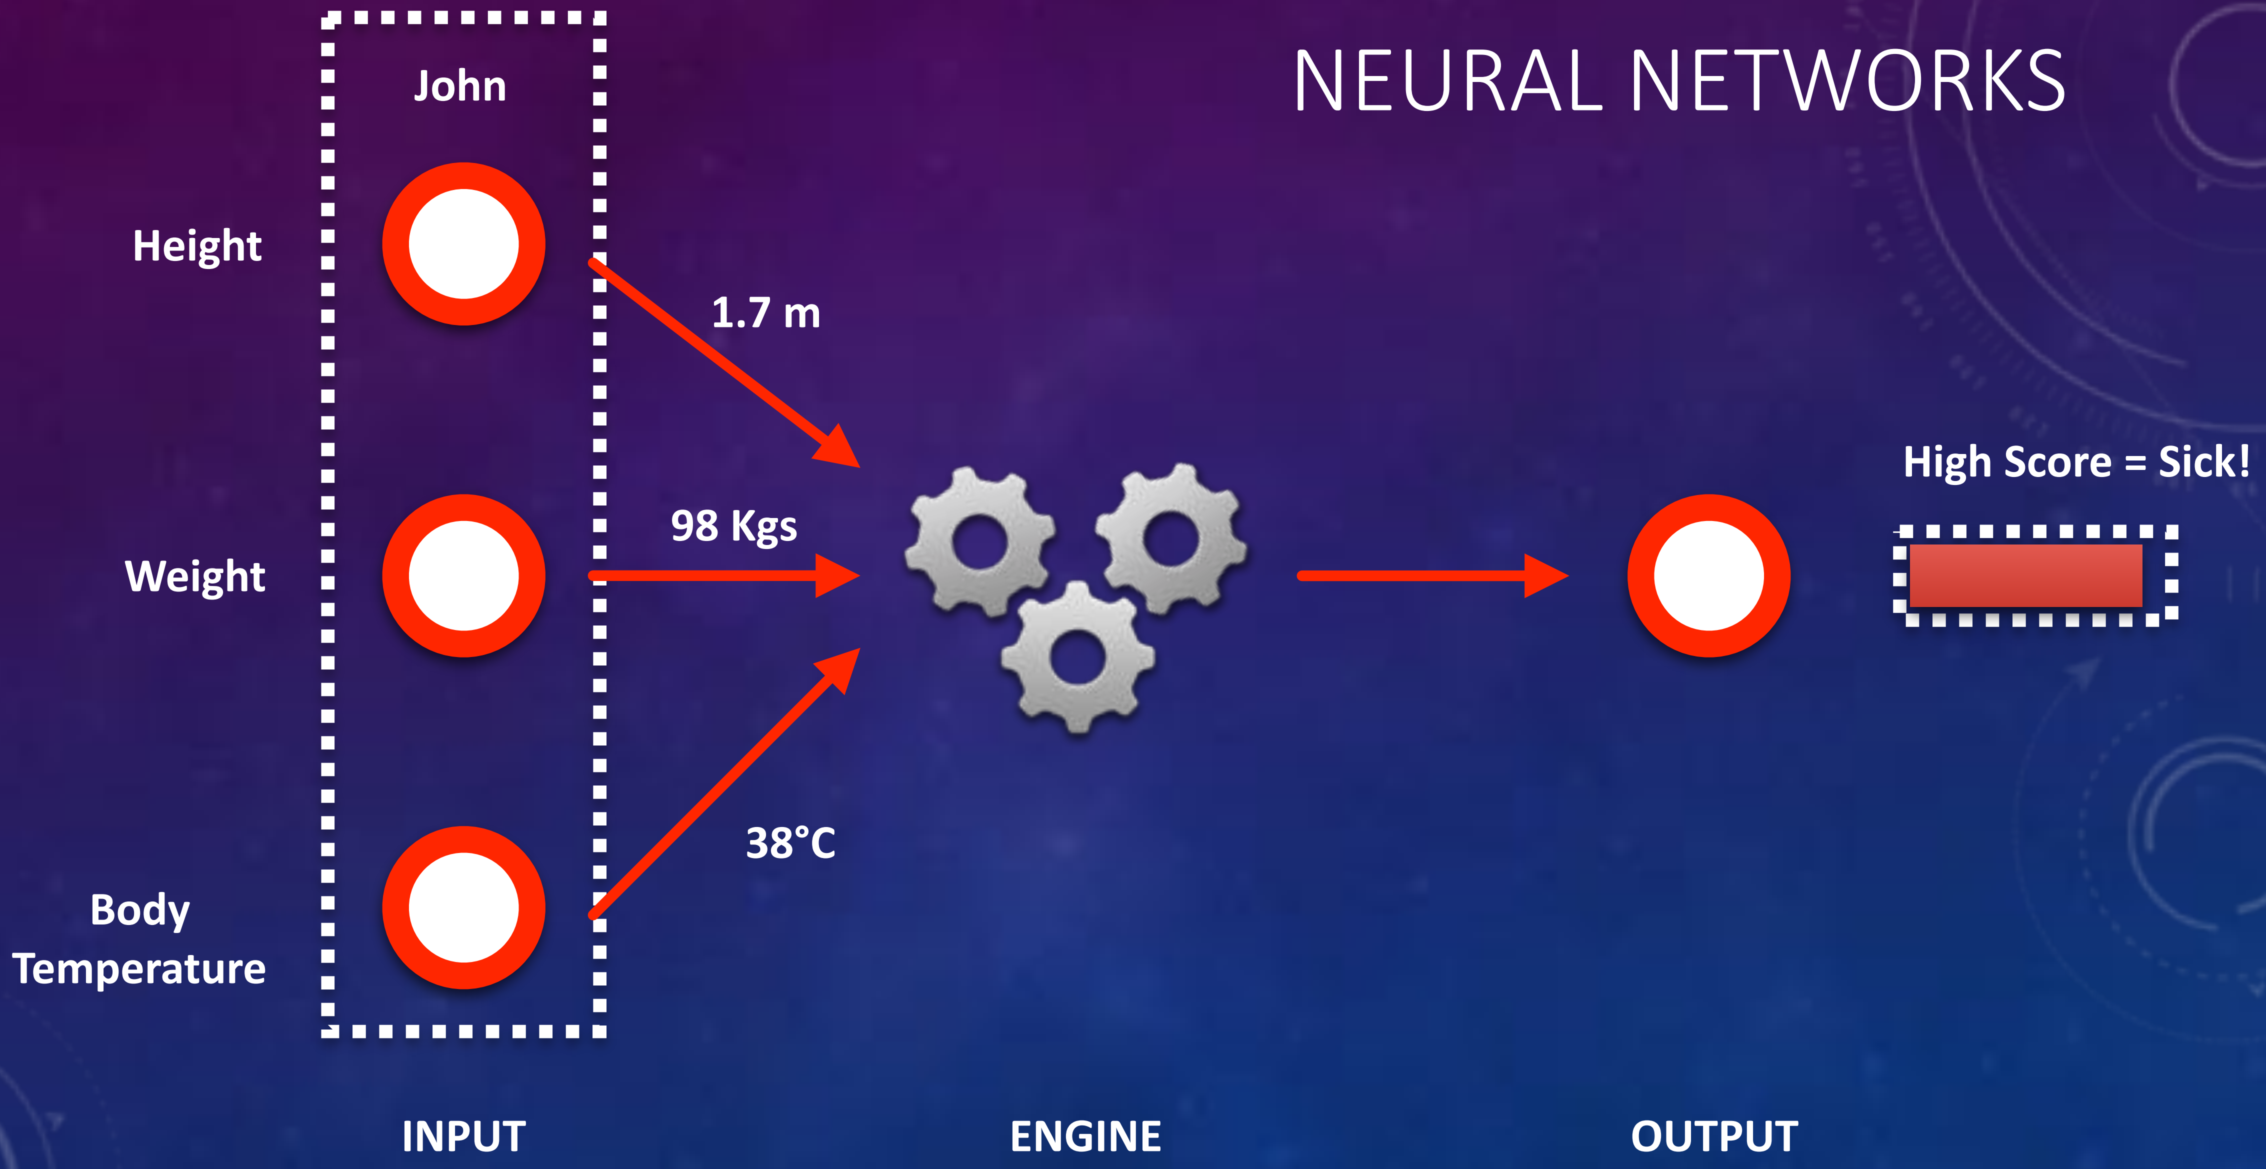

NEURAL NETWORKS

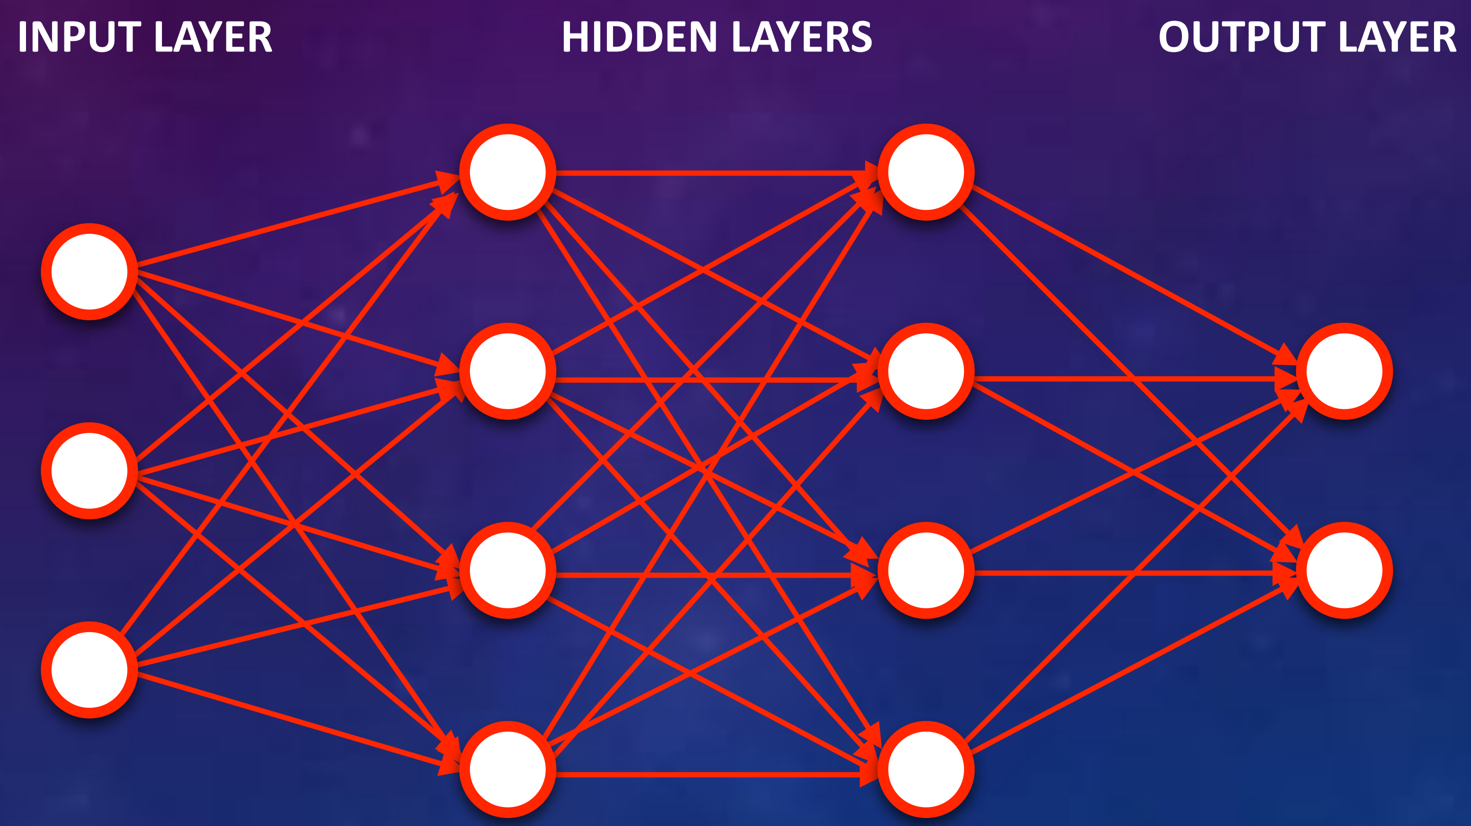

NEURAL NETWORKS EXAMPLE

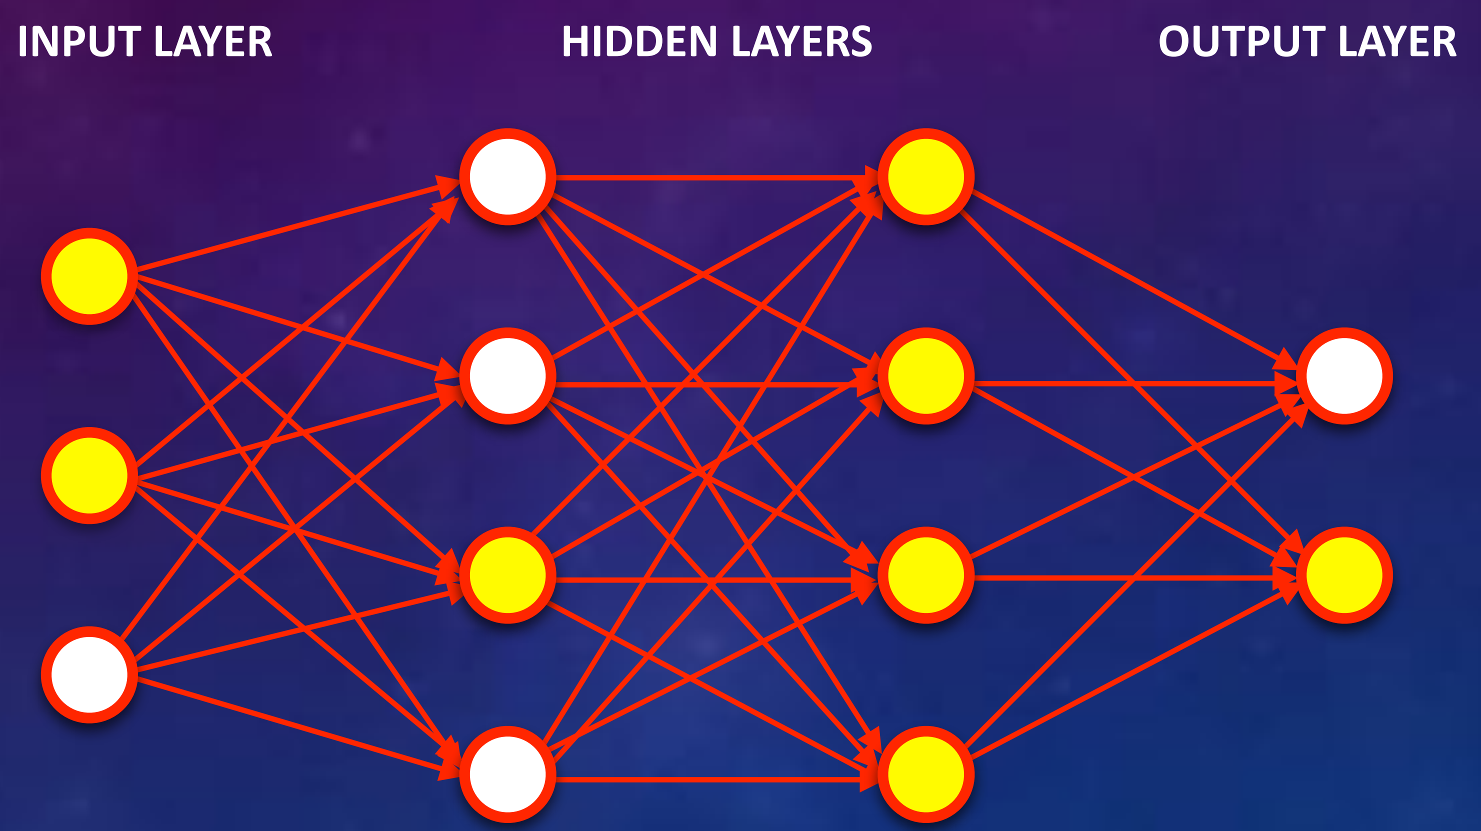

NEURAL NETWORKS EXAMPLE

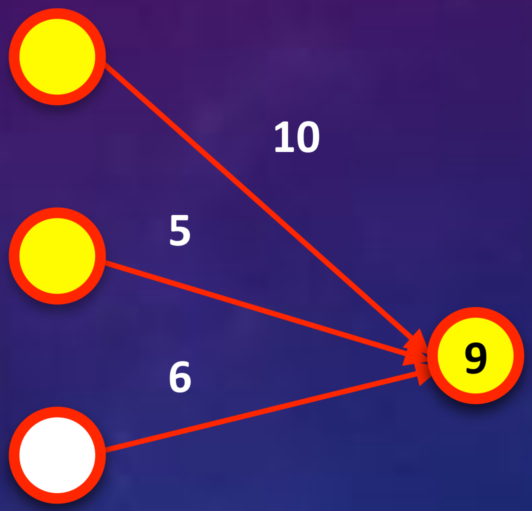

TRAINING

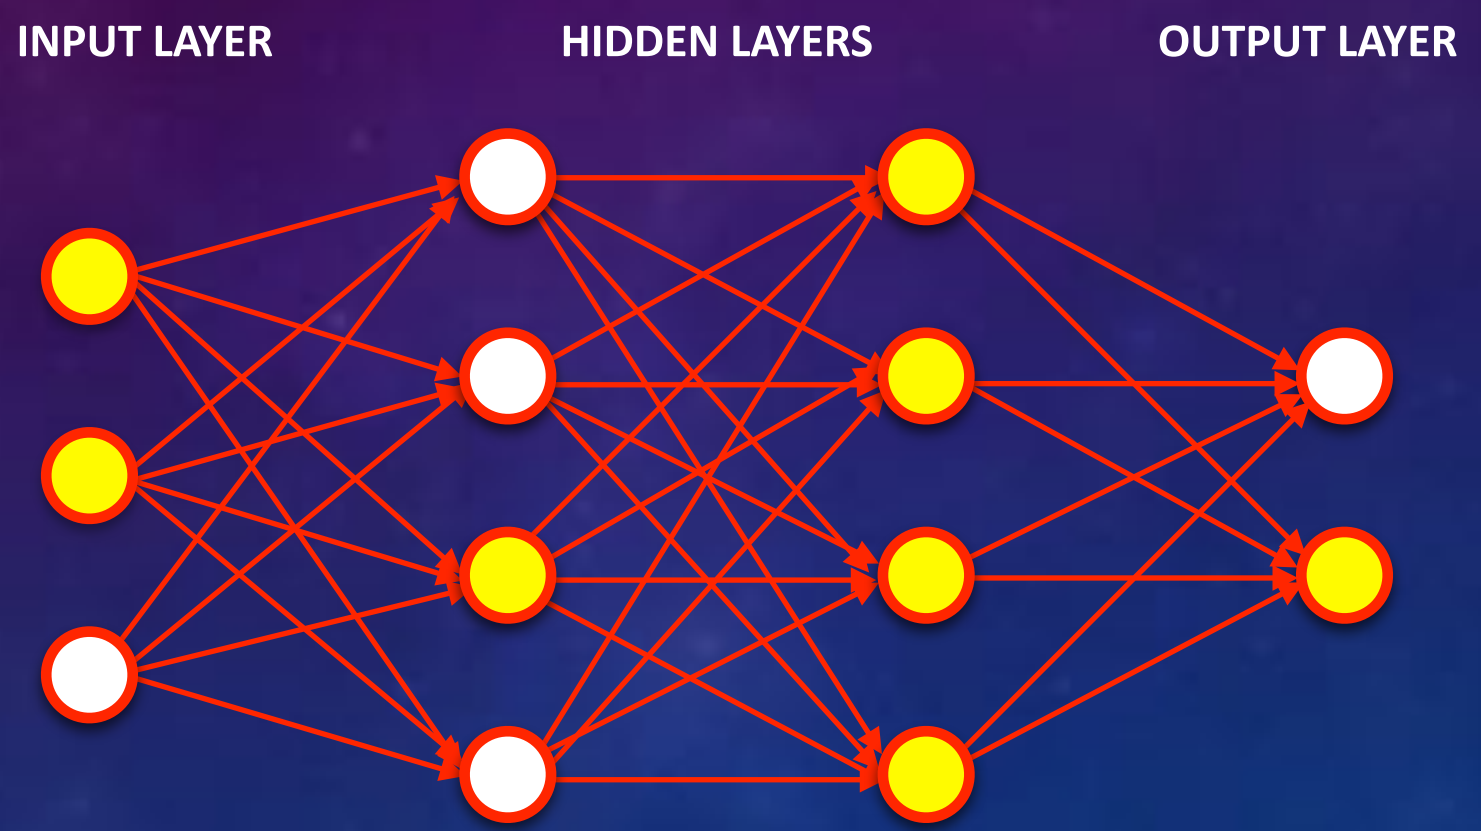

WHY ARE NEURAL NETS HARD?

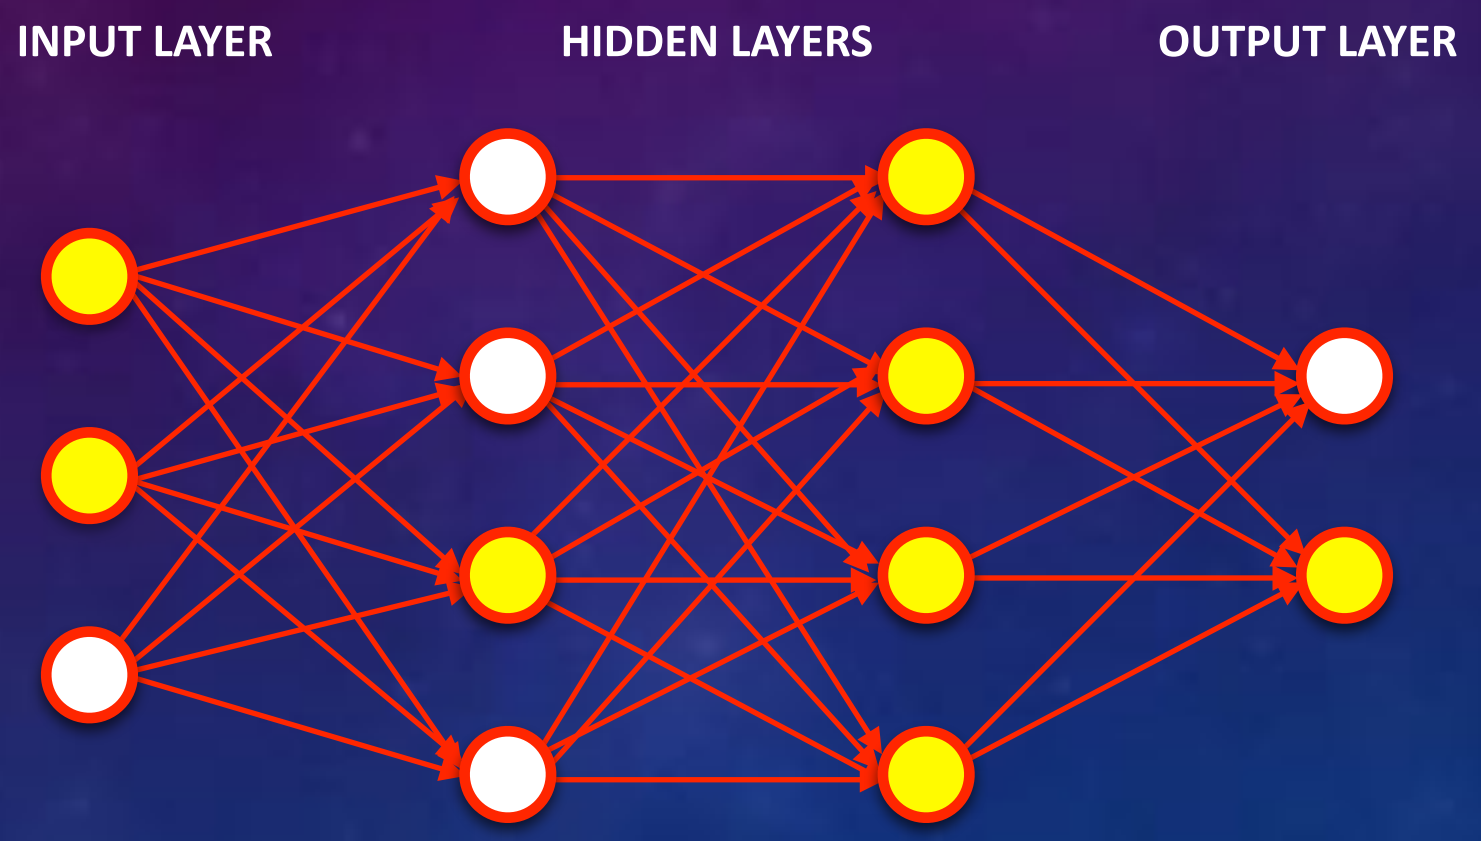

BACK PROPAGATION

Gradients in Hidden layers

Layer 3   0.2500 * 0.33 = 0.0825

Layer 2   0.0825 * 0.25 = 0.0206

Layer 1   0.0206 * 0.20 = 0.004125

… and numbers keep on getting smaller

NEURAL NETWORKS EXAMPLE

- http://playground.tensorflow.org/

PATTERN COMPLEXITY

Simple	Use ML methods like SVM
Moderate	Deep Nets, SVM
Complex	Deep Nets - only practical choice

DEEP NETS BETTER THAN HUMANS

Humans were 5.10% wrong
Microsoft was 4.94% wrong
Google was 4.82% wrong
Baidu was 4.58% wrong

Microsoft, Google Beat Humans at Image Recognition

Deep learning algorithms compete at ImageNet challenge

R. Colin Johnson
2/18/2015 08:15 AM EST
14 comments

PORTLAND, Ore. -- First computers beat the best of us at chess, then poker, and finally Jeopardy. The next hurdle is image recognition -- surely a computer can't do that as well as a human. Check that one off the list, too. Now Microsoft has programmed the first computer to beat the humans at image recognition.

The competition is fierce, with the ImageNet Large Scale Visual Recognition Challenge doing the judging for the 2015 championship on December 17. Between now and then expect to see a stream of papers claiming they have one-upped humans too. For instance, only 5 days after Microsoft announced it had beat the human benchmark of 5.1% errors with a 4.94% error grabbing neural network, Google announced it had one-upped Microsoft by 0.04%.

MIT Technology Review

Login / Register Search

Topics+ The Daily Magazine Business Reports More+

Computing

Baidu's Artificial-Intelligence Supercomputer Beats Google at Image Recognition

Update: On June 1, 2015, Baidu amended its technical paper on its system to admit that it had broken rules governing the ImageNet Challenge that the company had used to claim it had beaten other research teams. The organizers of the challenge reviewed Baidu's conduct and issued a statement saying its results should not be considered directly comparable to results obtained by others.

Chinese search giant Baidu says it has invented a powerful supercomputer that brings new muscle to an artificial-intelligence technique giving software more power to understand speech, images, and written language.

HOW DOES IT WORK?

DEEP ARCHITECTURE

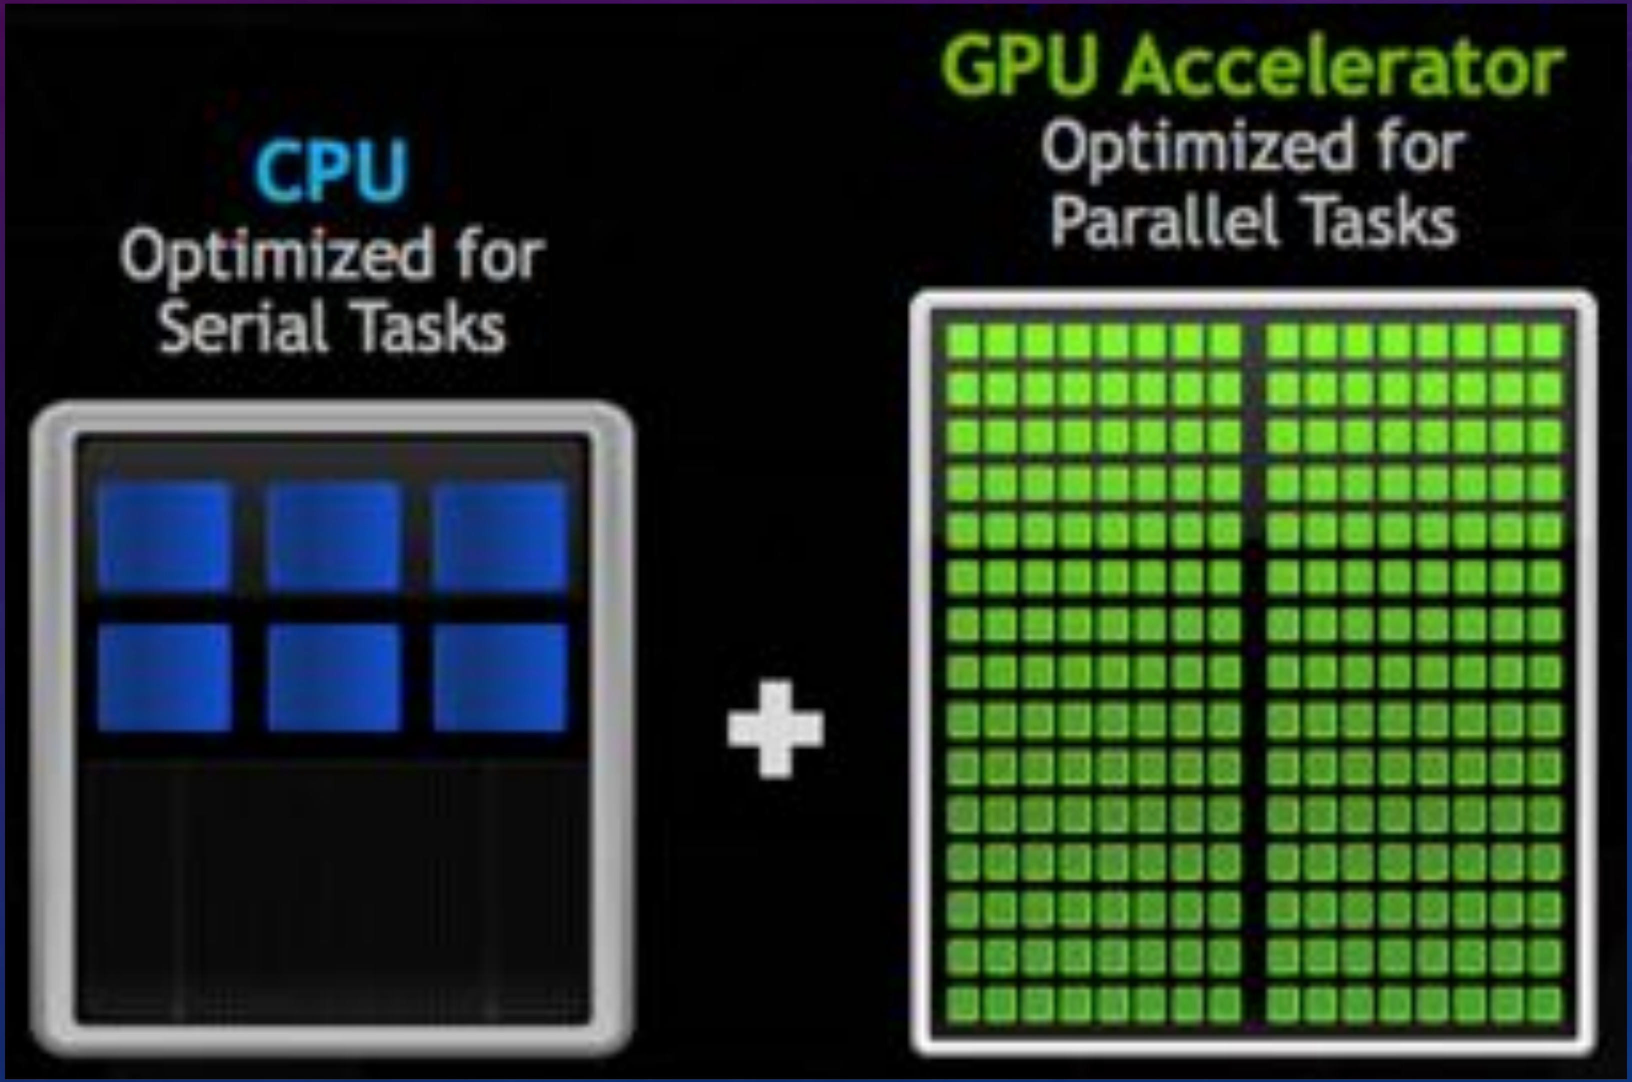

DEEP ARCHITECTURE

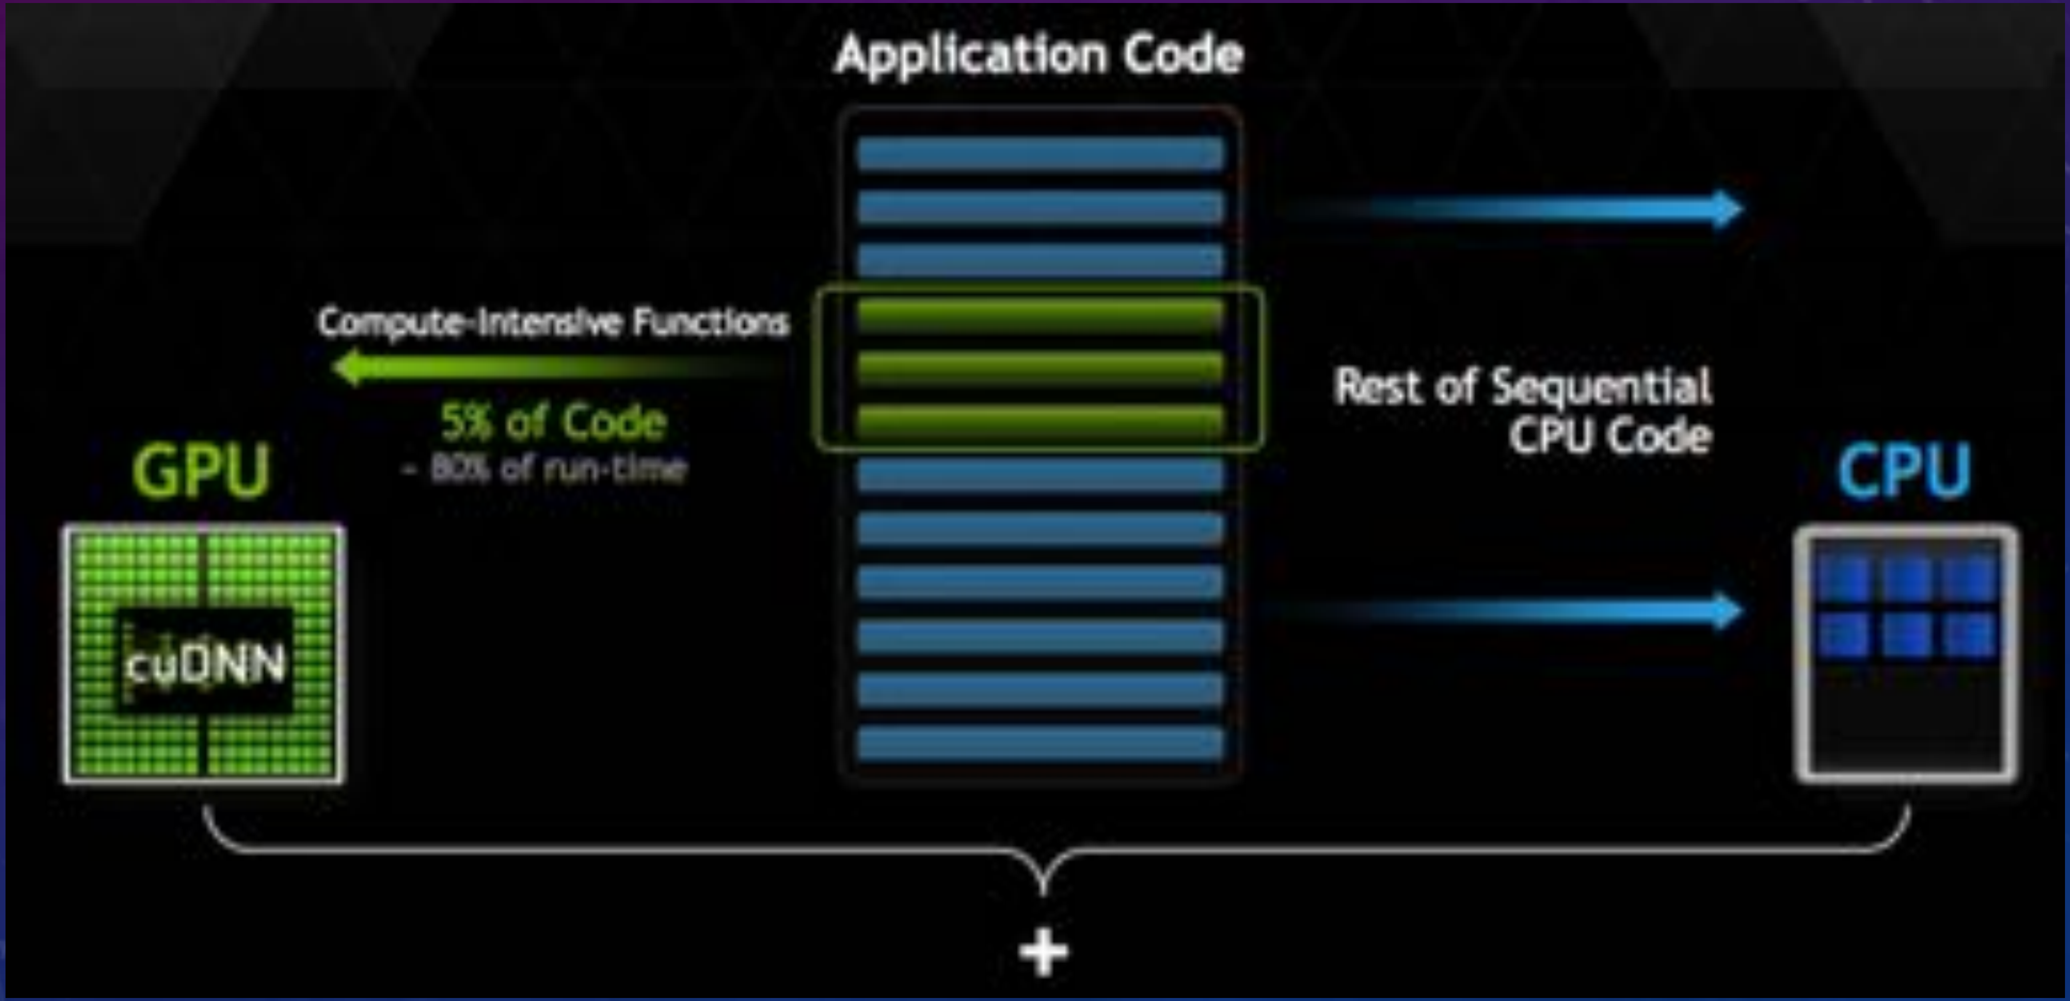

DEEP ARCHITECTURE

Batch Size	Training Time CPU	Training Time GPU	GPU Speed Up
64 images	64 s	7.5 s	8.5X
128 images	124 s	14.5 s	8.5X
256 images	257 s	28.5 s	9.0X

2006 - 2007


Geoffrey Hinton
(Toronto, Google)
Yann LeCun
(NewYork, Facebook)
Yoshua Bengio
(Montreal)

RESTRICTED BOLTZMANN MACHINE (RBM)

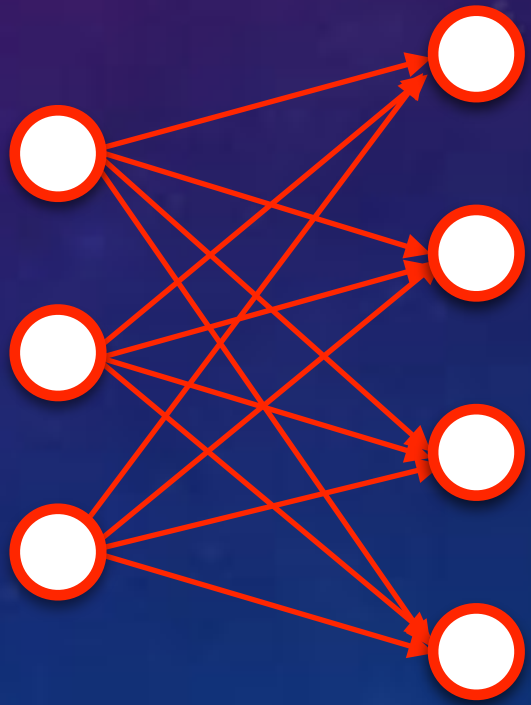

VISIBLE LAYER

HIDDEN LAYERS

DEEP BELIEF NETS (DBNS)

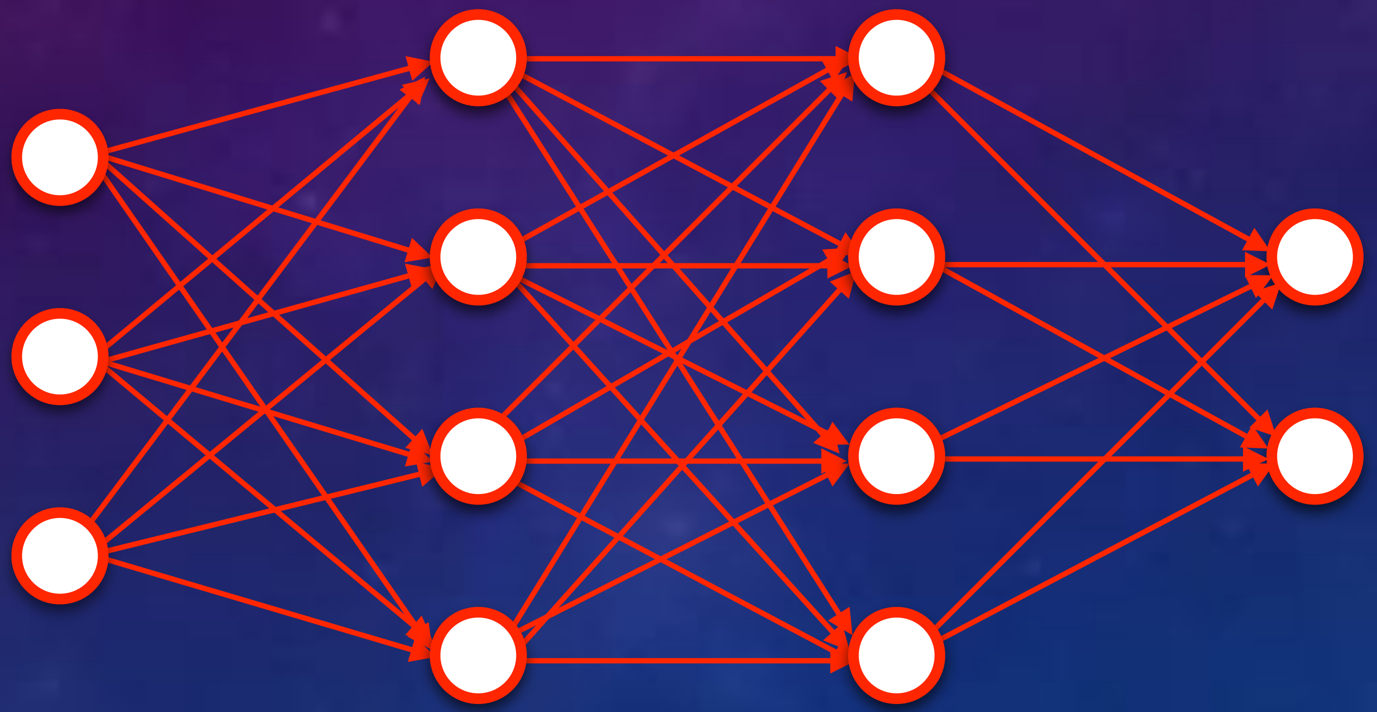

VISIBLE
LAYER

HIDDEN
LAYERS

AUTOENCODERS

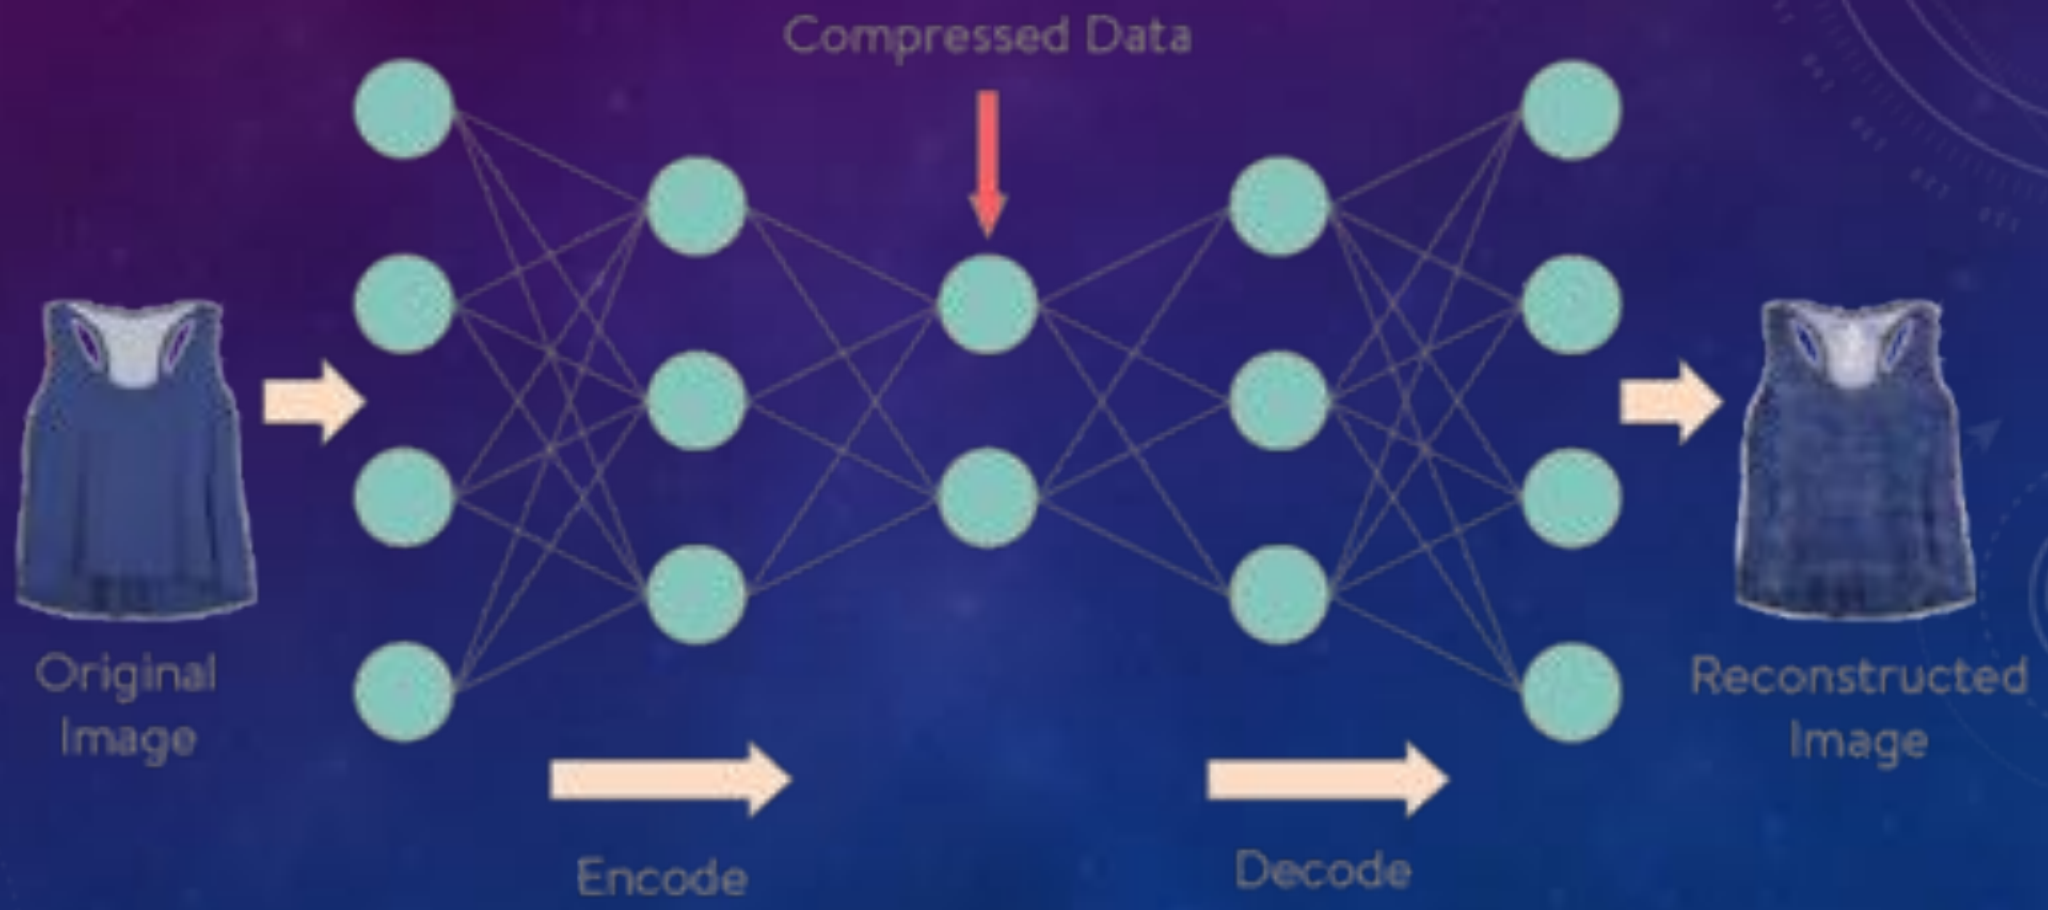

CONVOLUTIONAL NEURAL NETWORK (CNN)

CONVOLUTIONAL NEURAL NETWORK (CNN)

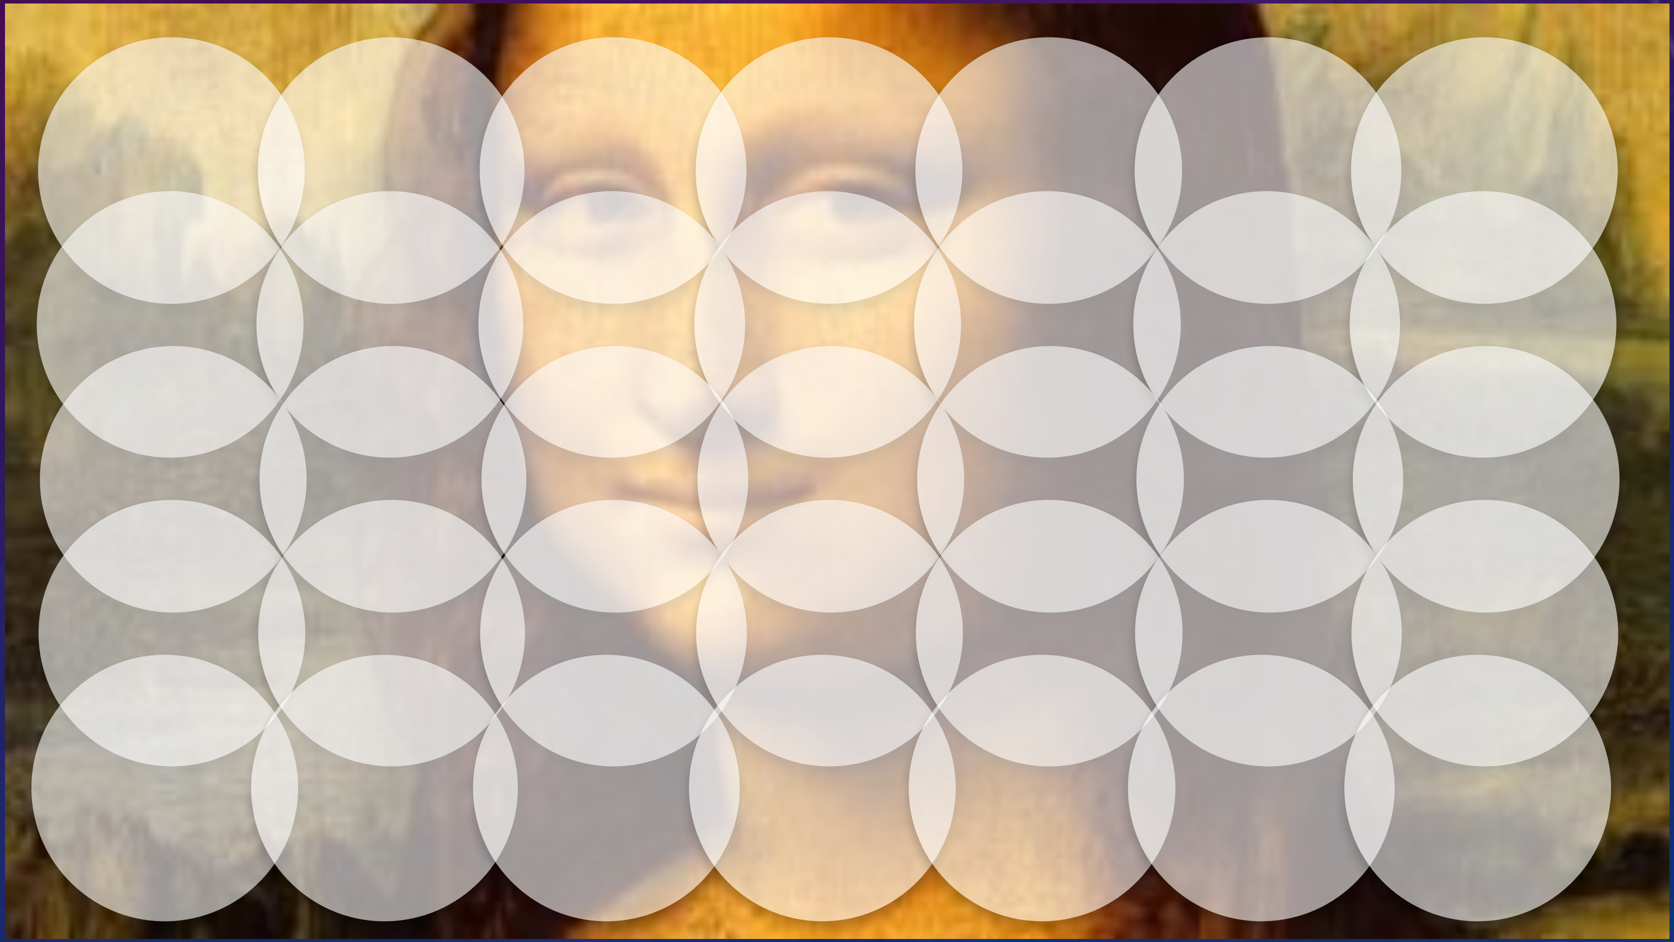

CONVOLUTIONAL NEURAL NETWORK (CNN)

Filter 1

Filter 2

Filter 3

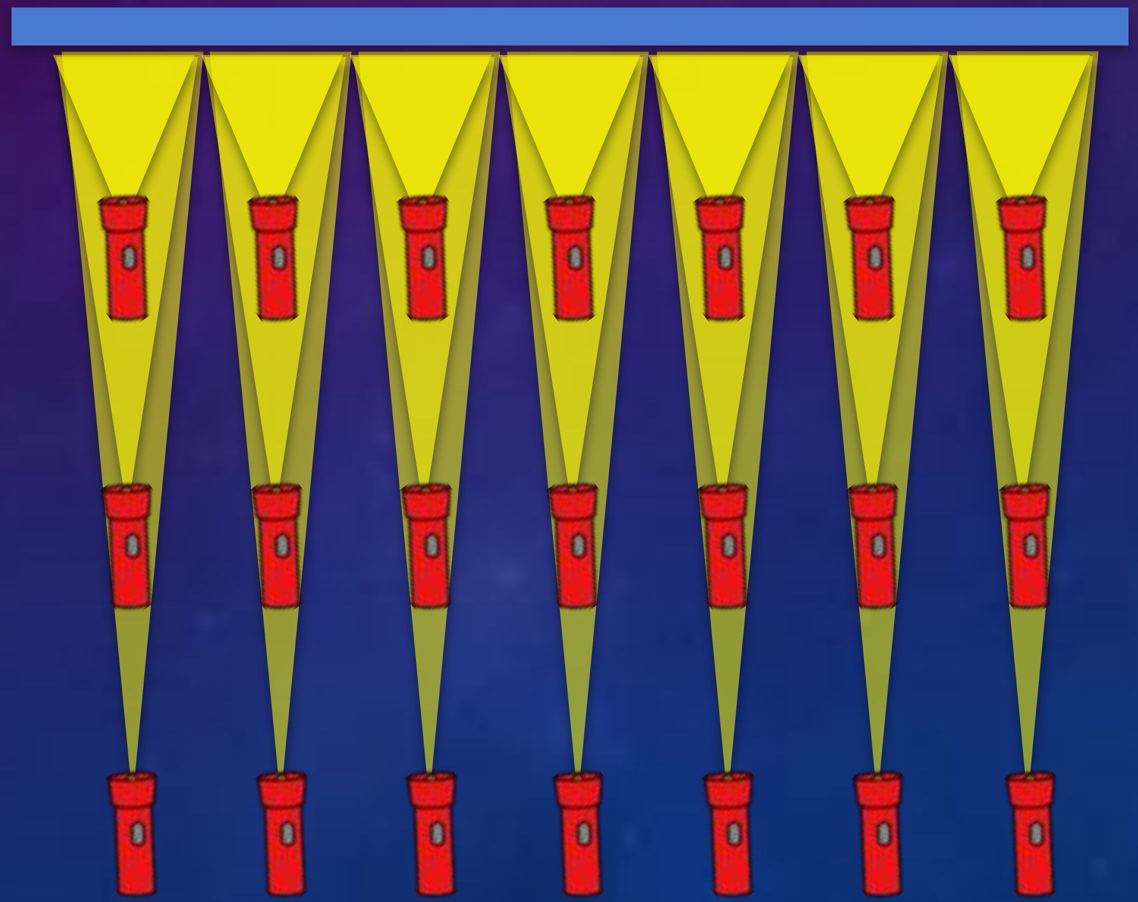

CONVOLUTIONAL NEURAL NETWORK (CNN)

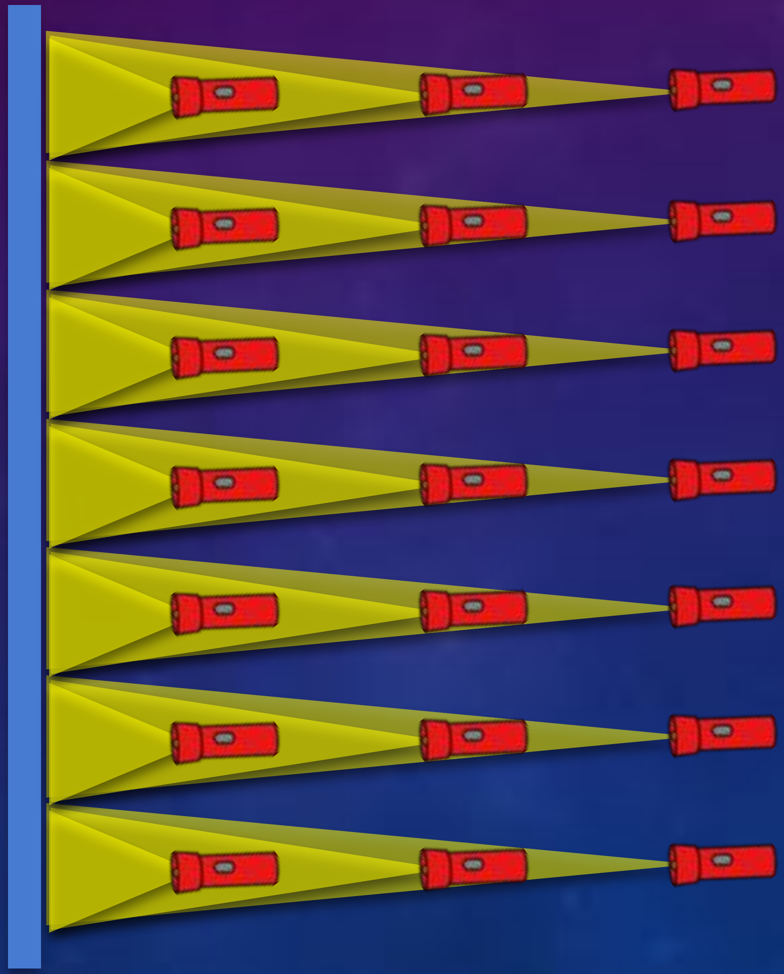

CONVOLUTIONAL NEURAL NETWORK (CNN)

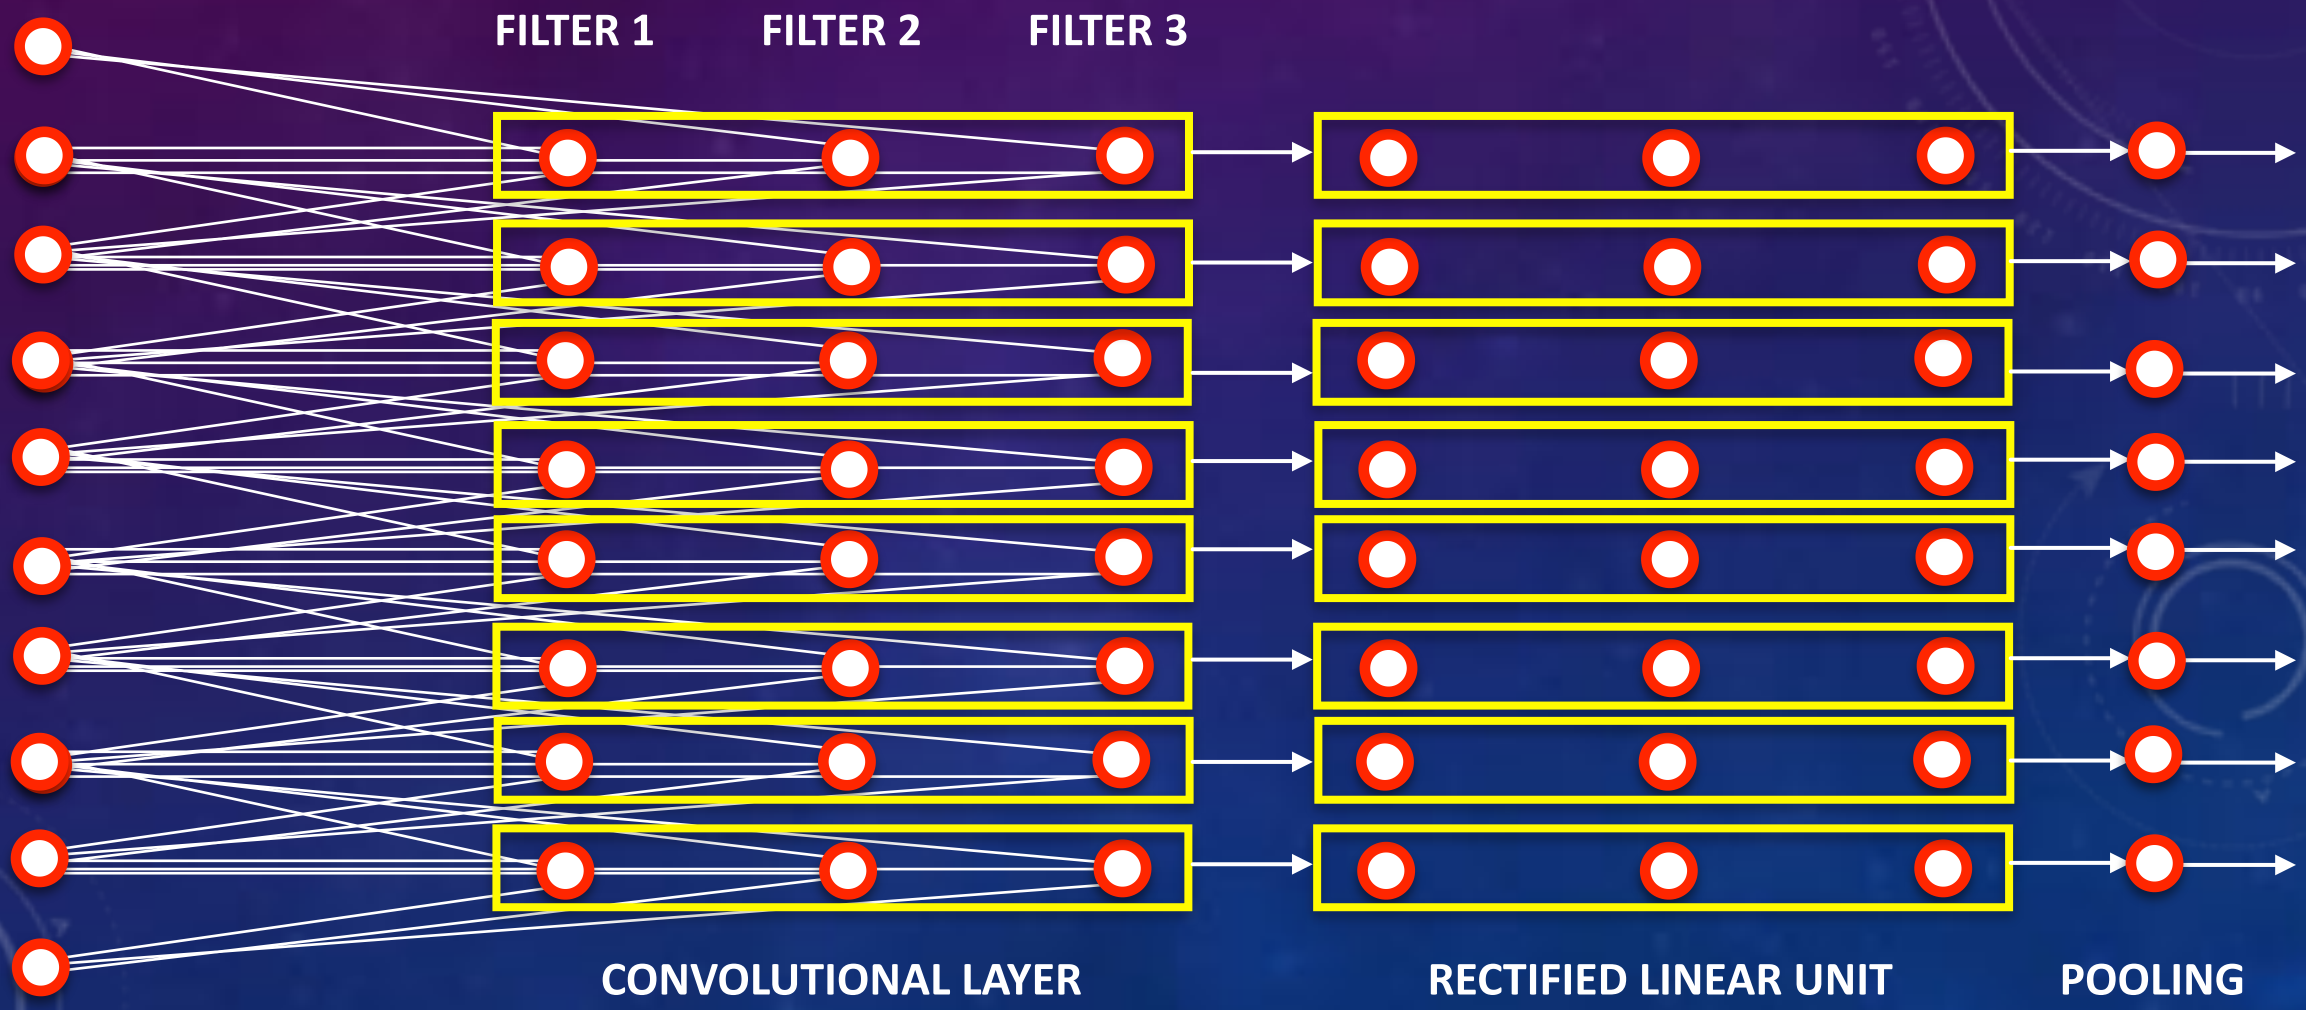

CONVOLUTIONAL NEURAL NETWORK (CNN)

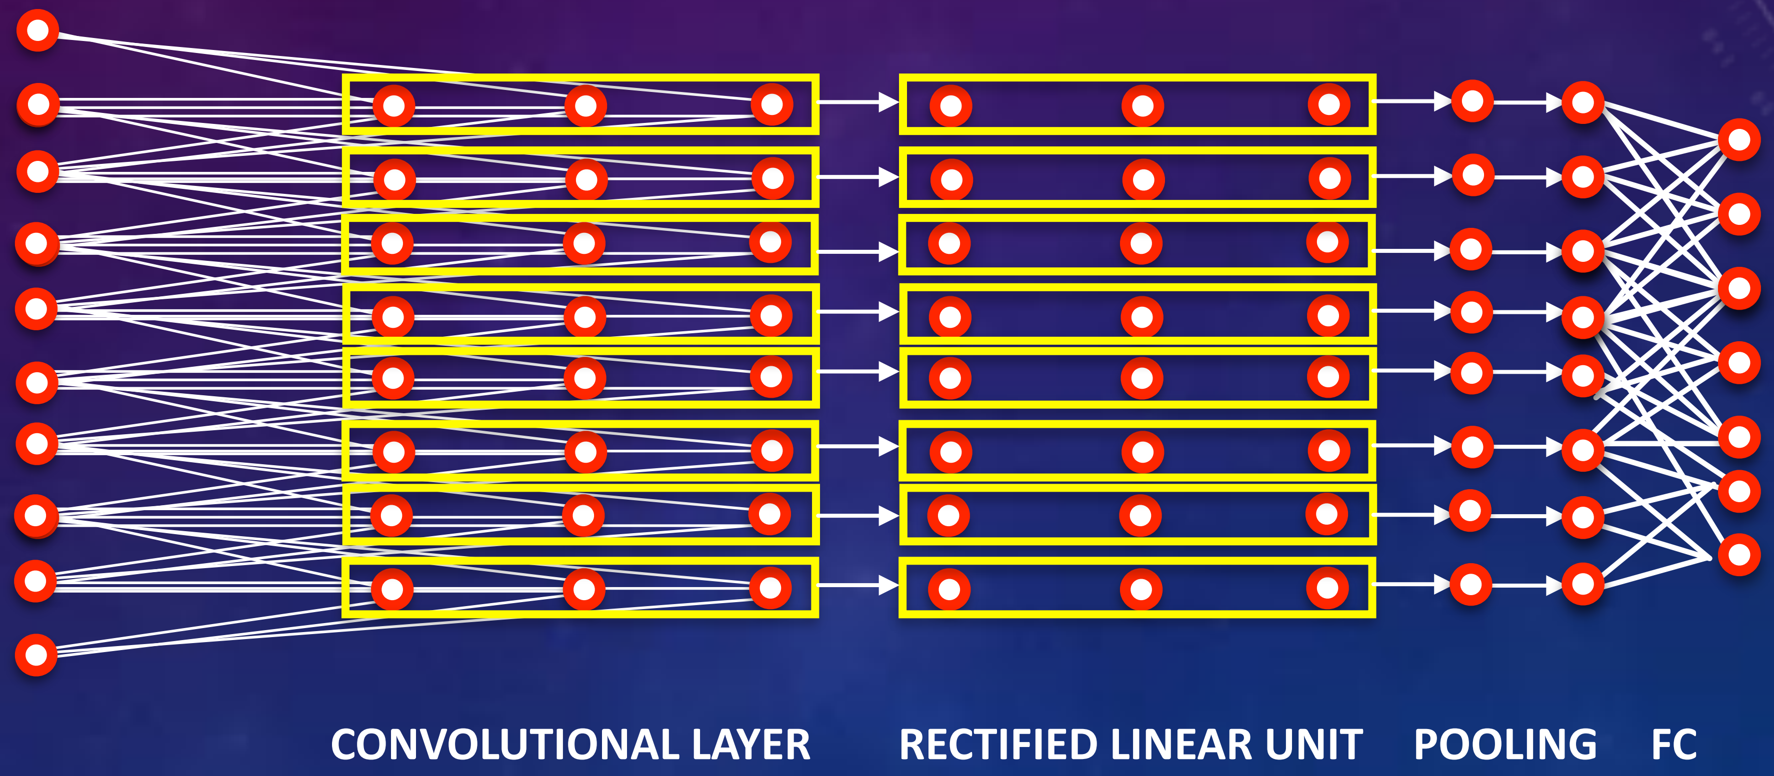

CNN ARCHITECTURE

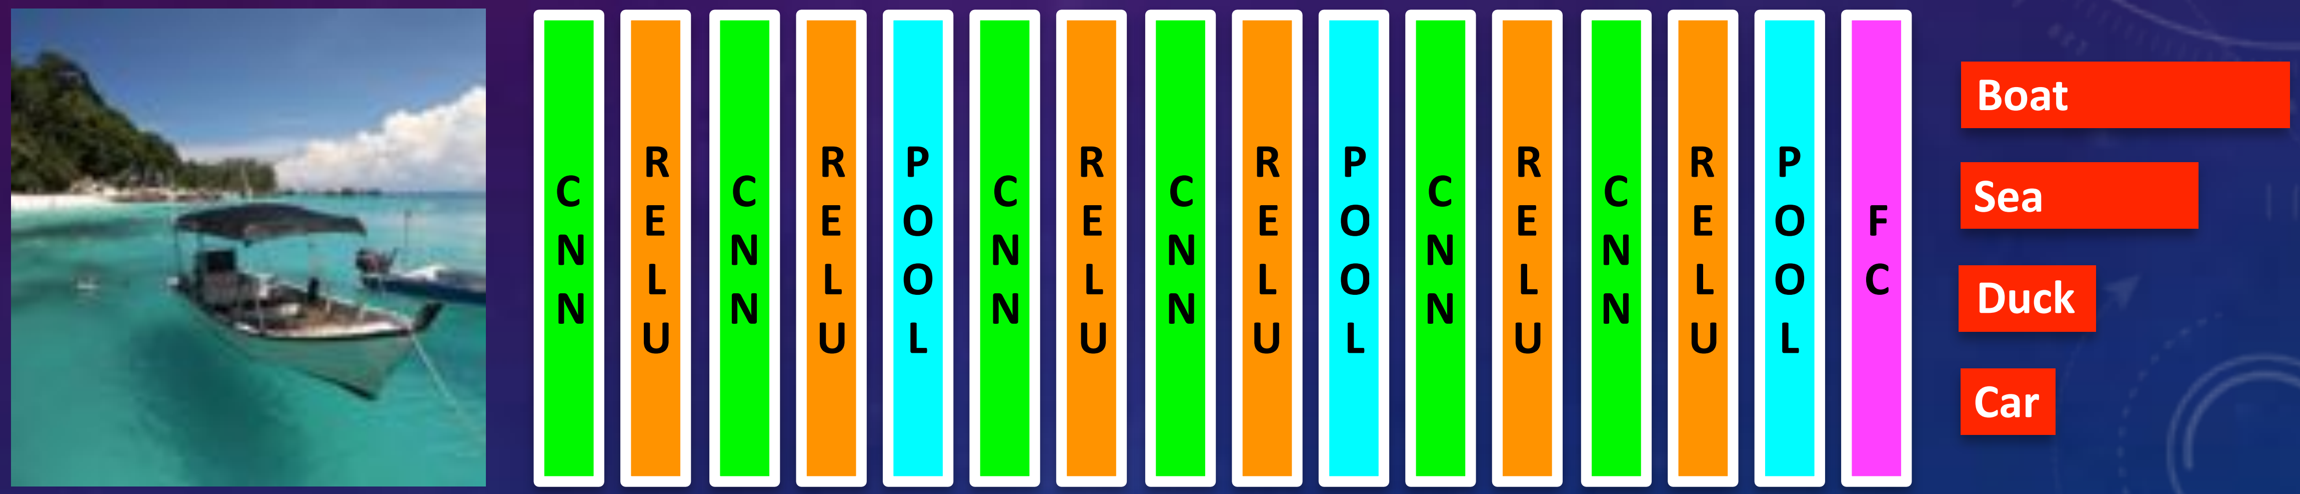

RECURRENT NEURAL NETWORK (RNN)

RNN STRUCTURE

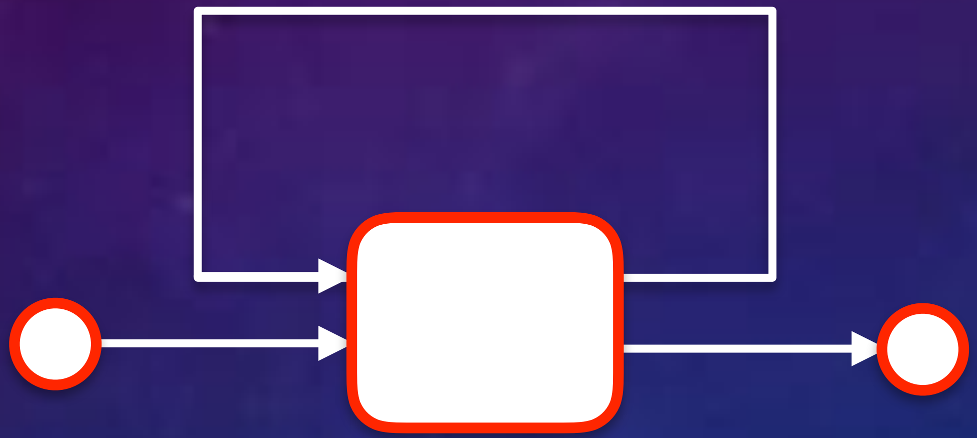

RECURSIVE NEURAL TENSOR NETWORKS (RNTNS)

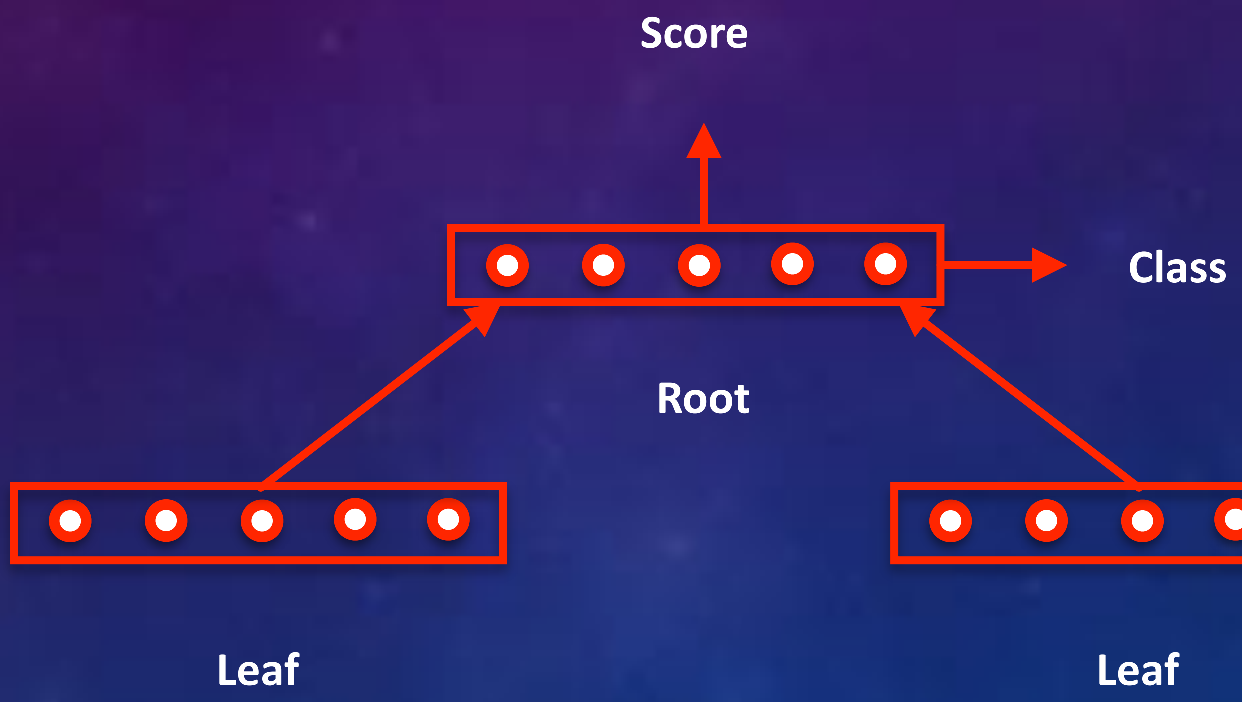

RECURSION

HOW TO CHOOSE A DNN?

- ✔ Unlabelled Data
- ✔ Feature Extraction
- ✔ Unsupervised Learning
- ✔ Pattern Recognition

Restricted Boltzmann Machine (RBM)

Autoencoder

HOW TO CHOOSE A DNN?

- ✔ Labelled Data
- ✔ Textual Processing
- ✔ Sentiment Analysis
- ✔ Parsing
- ✔ Named Entity Recognition

Recursive Neural Tensor Network (RNTN)

Recurrent Net

HOW TO CHOOSE A DNN?

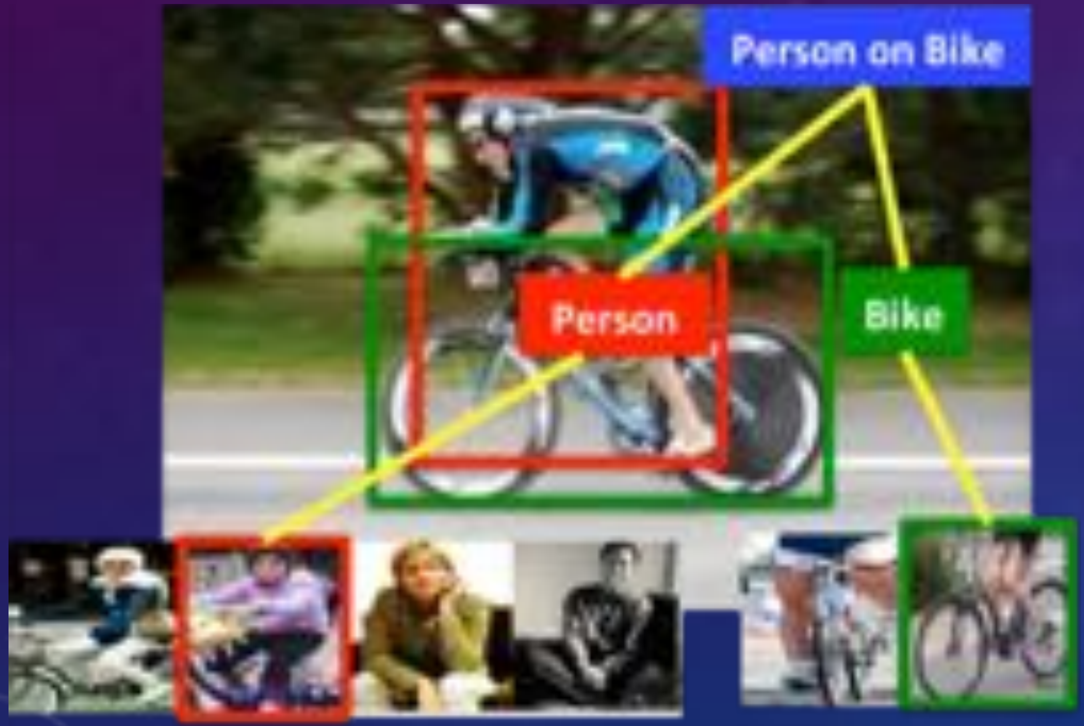

✔ Object Recognition

Recursive Neural Tensor Network (RNTN)

Convolutional Net

HOW TO CHOOSE A DNN?

- ✔ Image Classification

Deep Believe Network (DBN)

Convolutional Net

HOW TO CHOOSE A DNN?

✔ Speech Recognition } **Recurrent Net**

HOW TO CHOOSE A DNN?

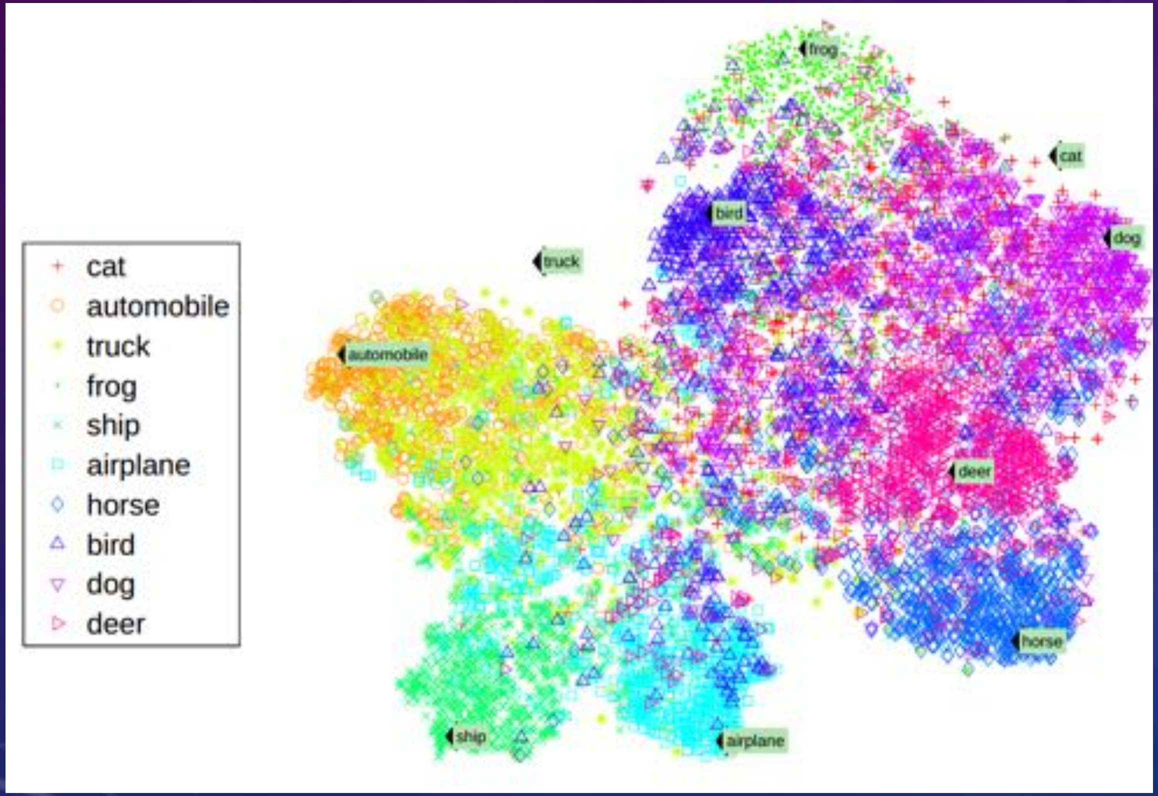

✔ Classification

Multilayer Perceptron (MLP) with Rectified Linear Units (RELU)

Deep Belief Net

DEEP LEARNING PLATFORMS OR LIBRARIES?

H2O.ai

turi

jupyter

DL4J Deep Learning for Java

torch

Caffe

TensorFlow

theano

TENSOR FLOW

"Machine learning is a core, transformative way by which we're rethinking everything we're doing"

Sundar Pichai, Alphabet CEO

WHAT ARE COMPUTATIONAL GRAPHS?

- A directed graph
- Defines computational structures
- Functions chained together to produced a specific output
- We can construct, complex transformation on data using small, well-defined mathematical functions
- Core of TensorFlow Programs

BASIC FUNCTION

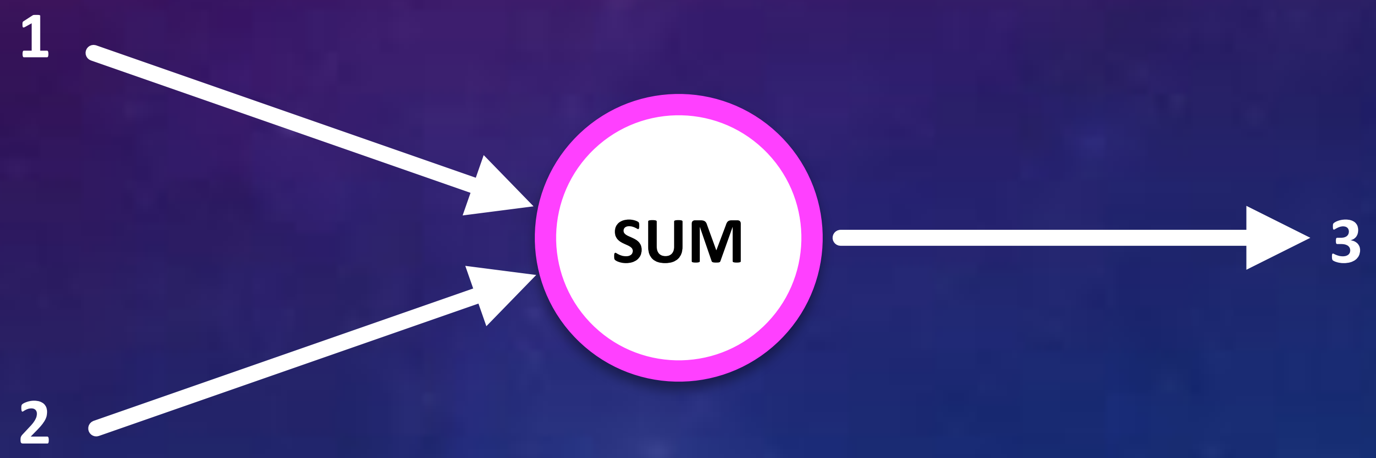

COMPLEX FUNCTION

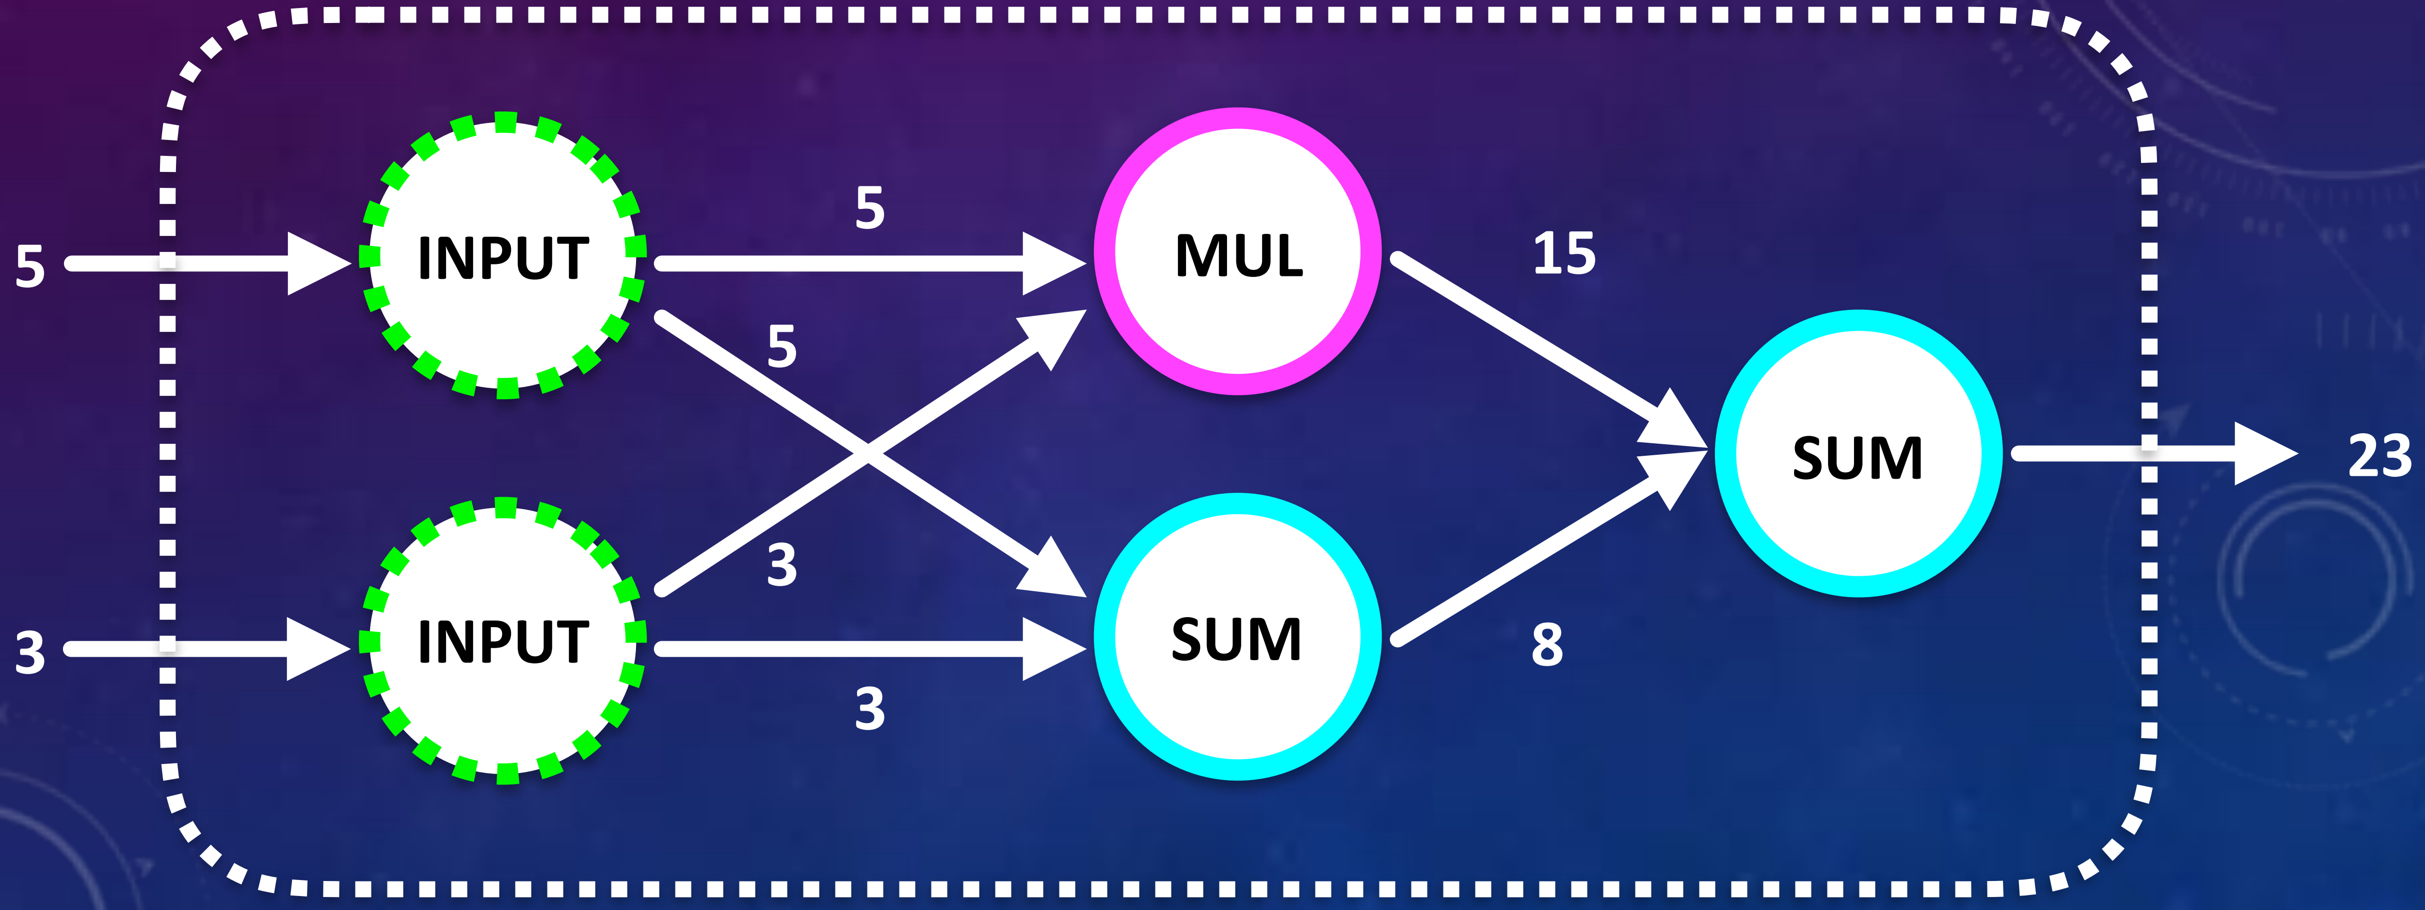

MORE COMPLEX FUNCTION

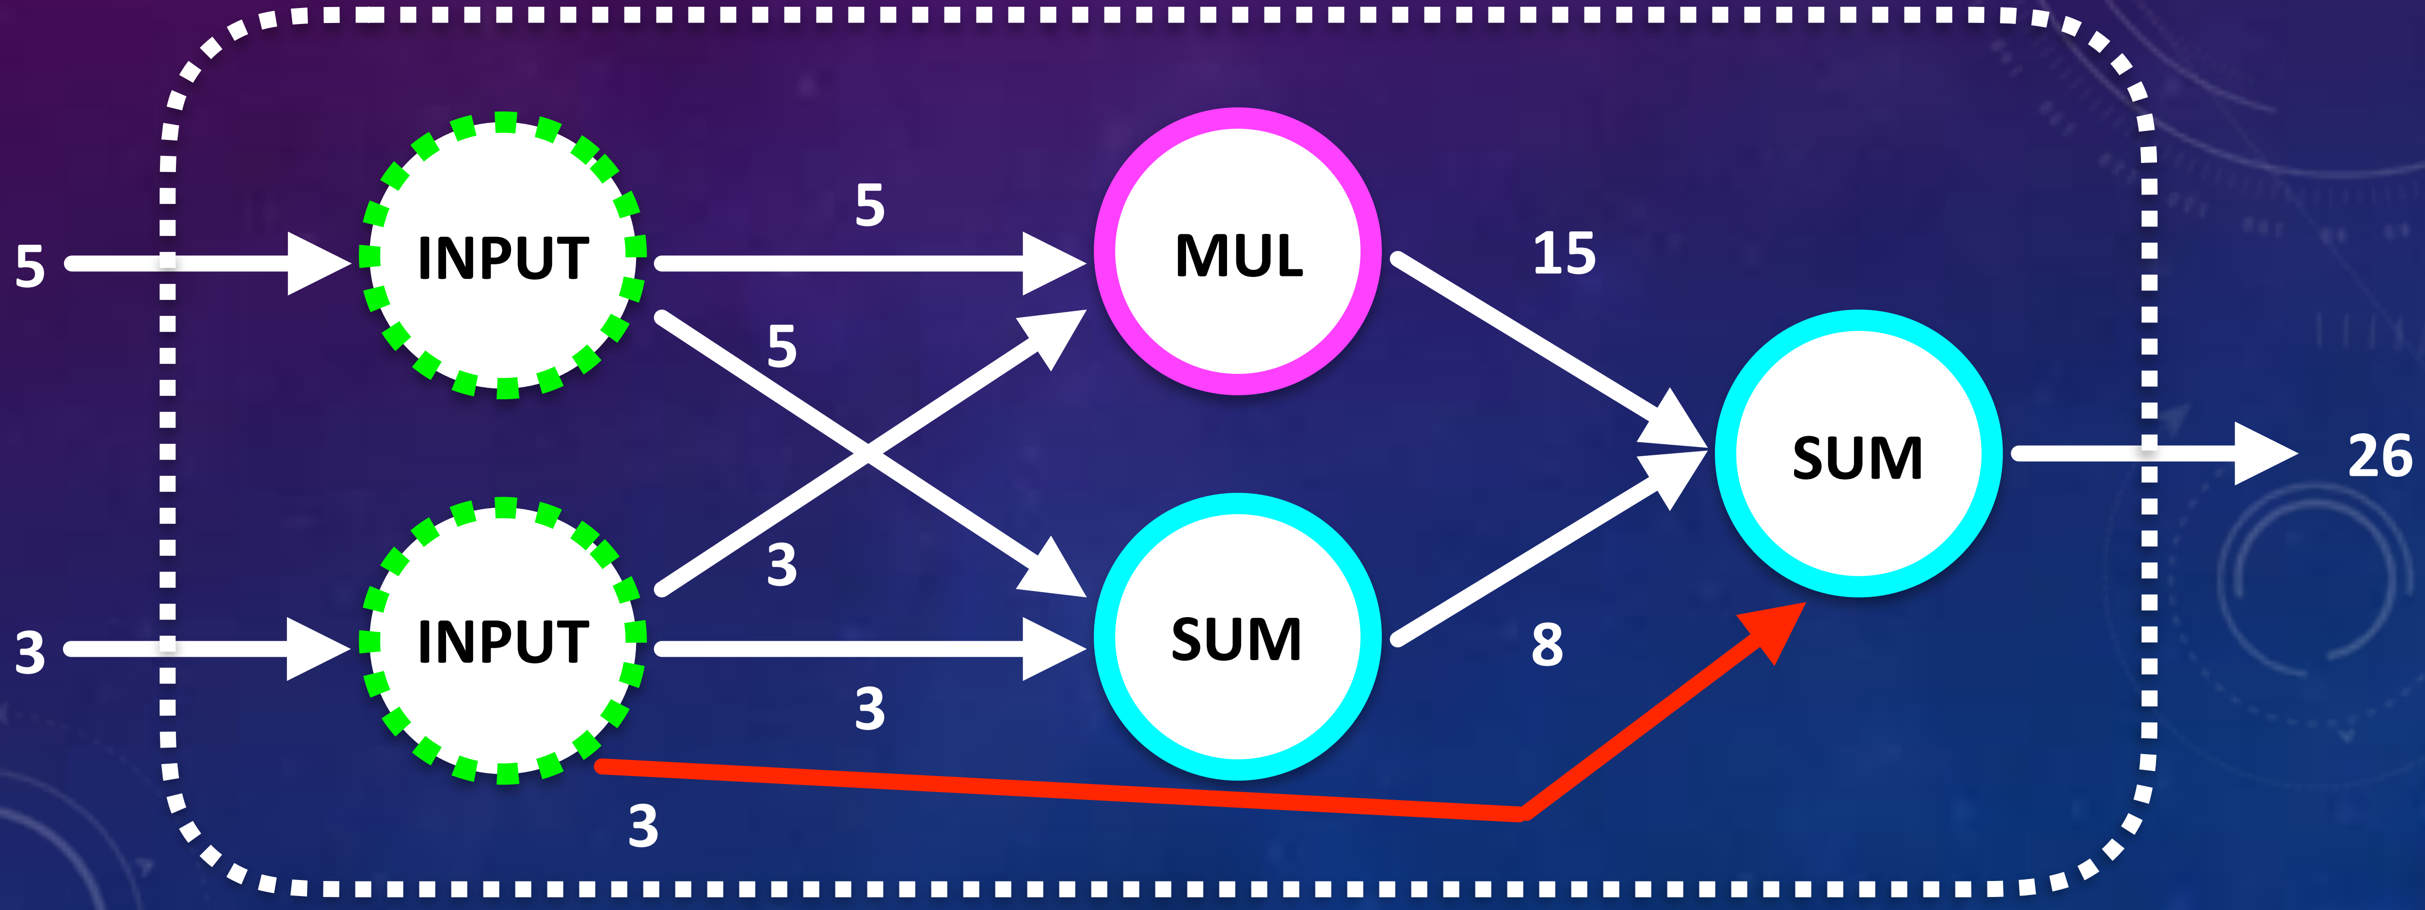

DEPENDENCIES

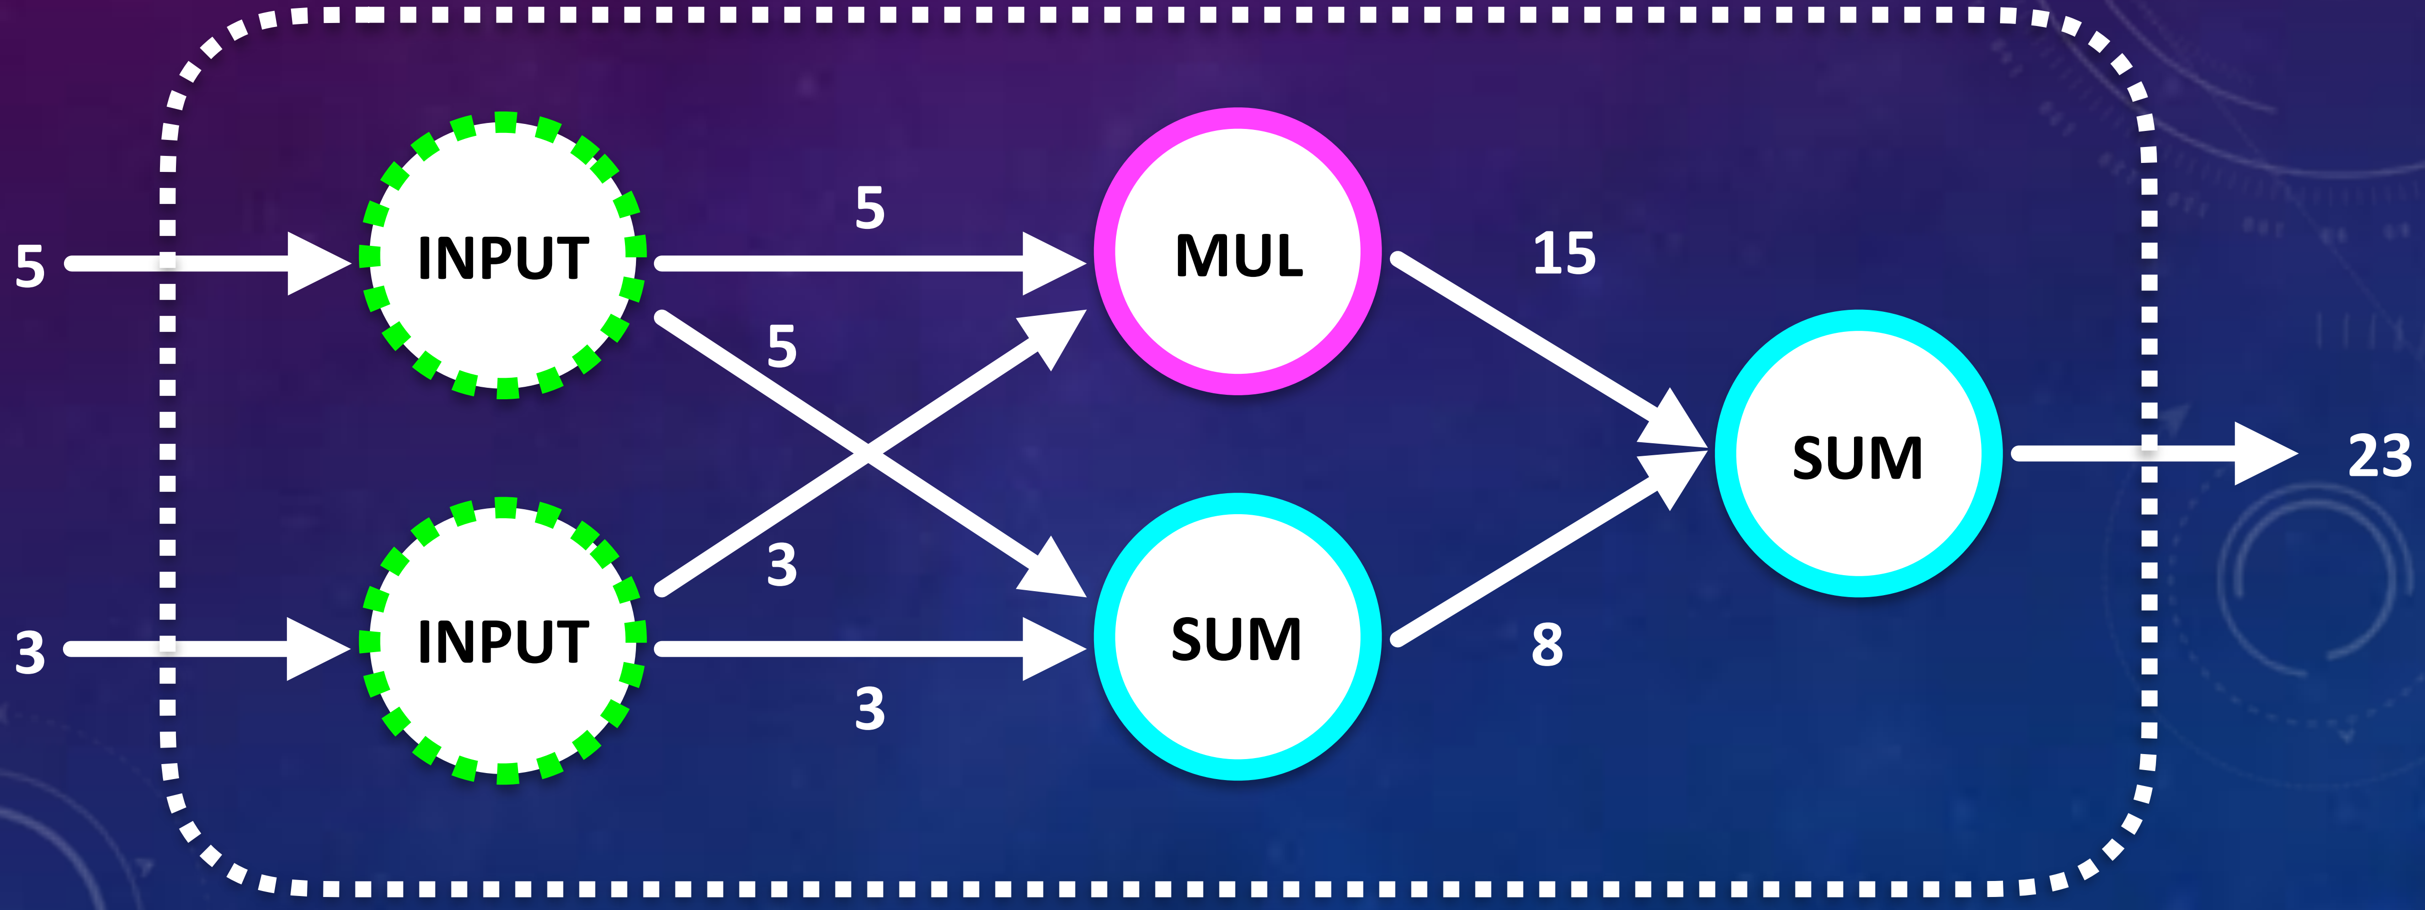

DEPENDENCIES

CIRCULAR DEPENDENCIES

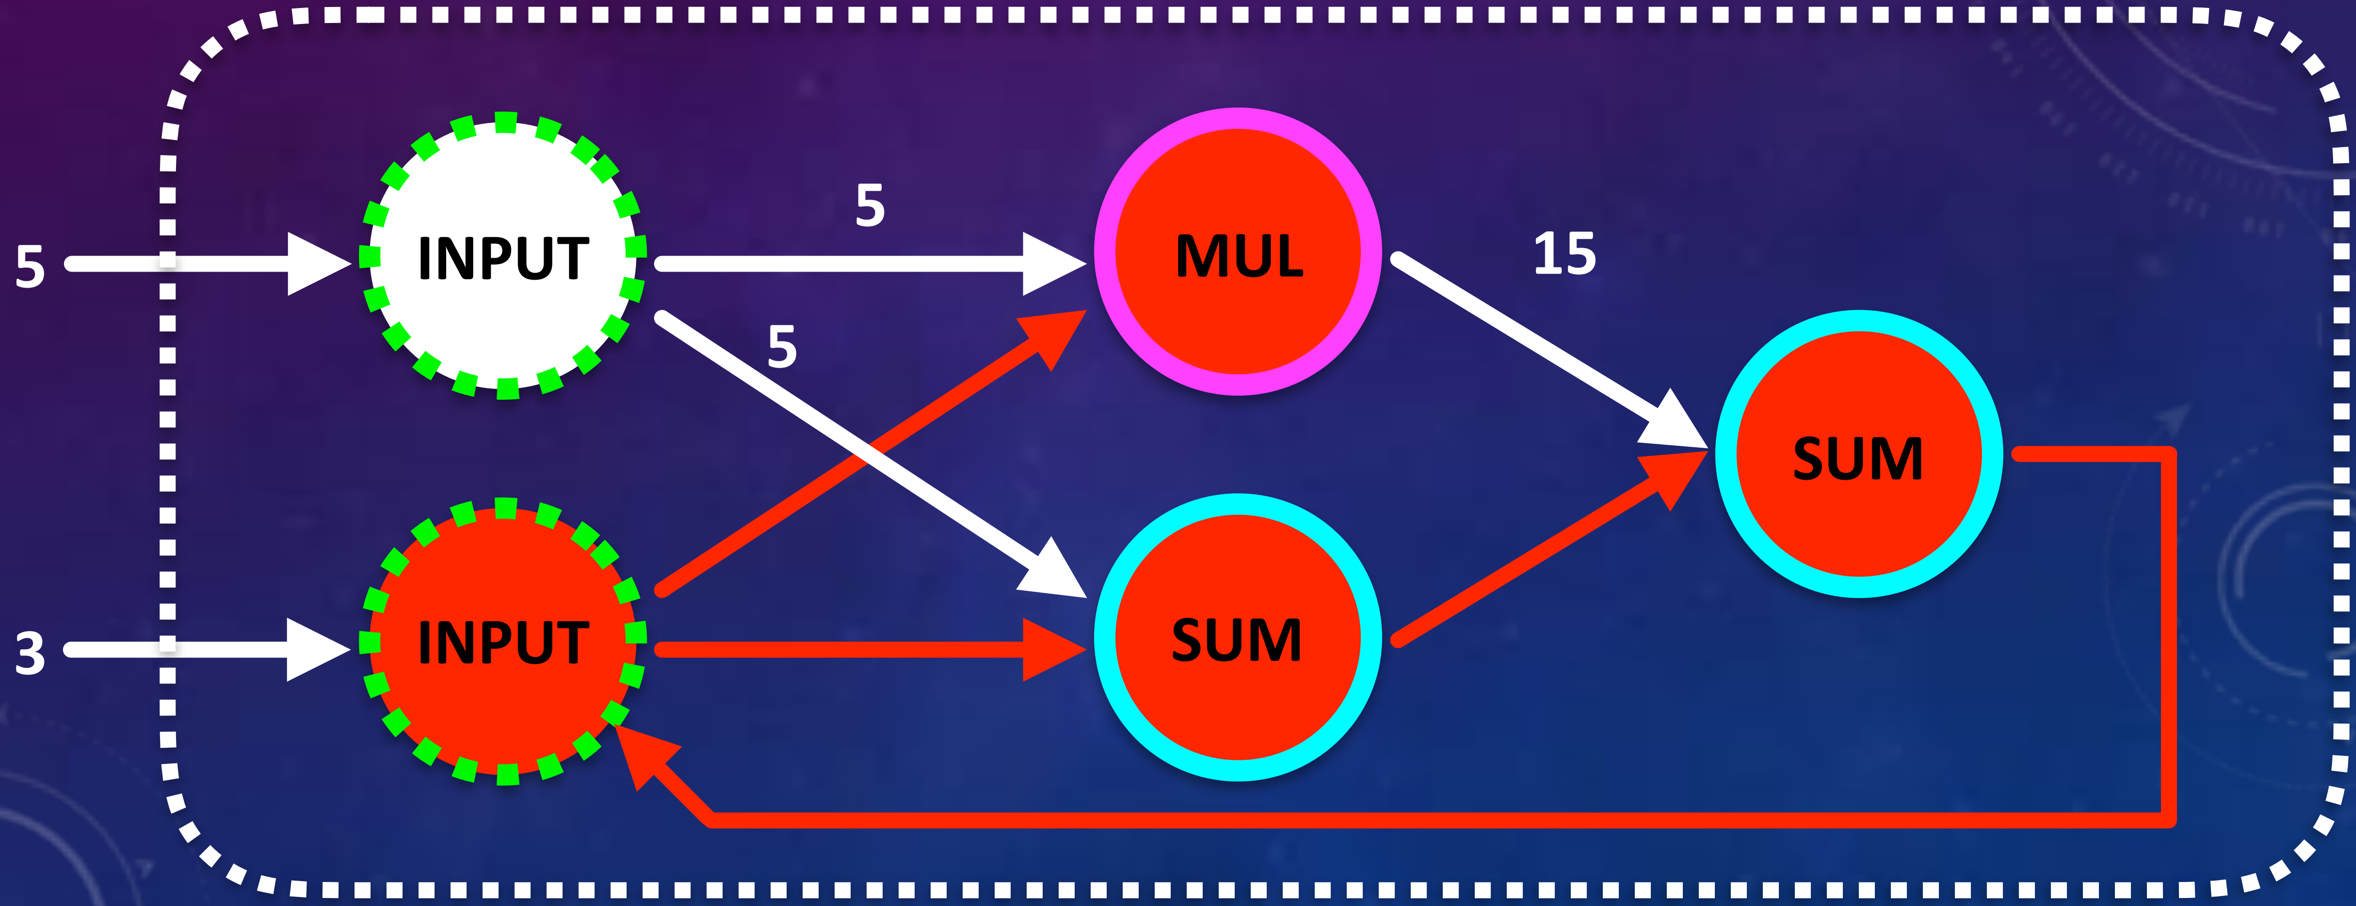

UNROLLING THE GRAPH

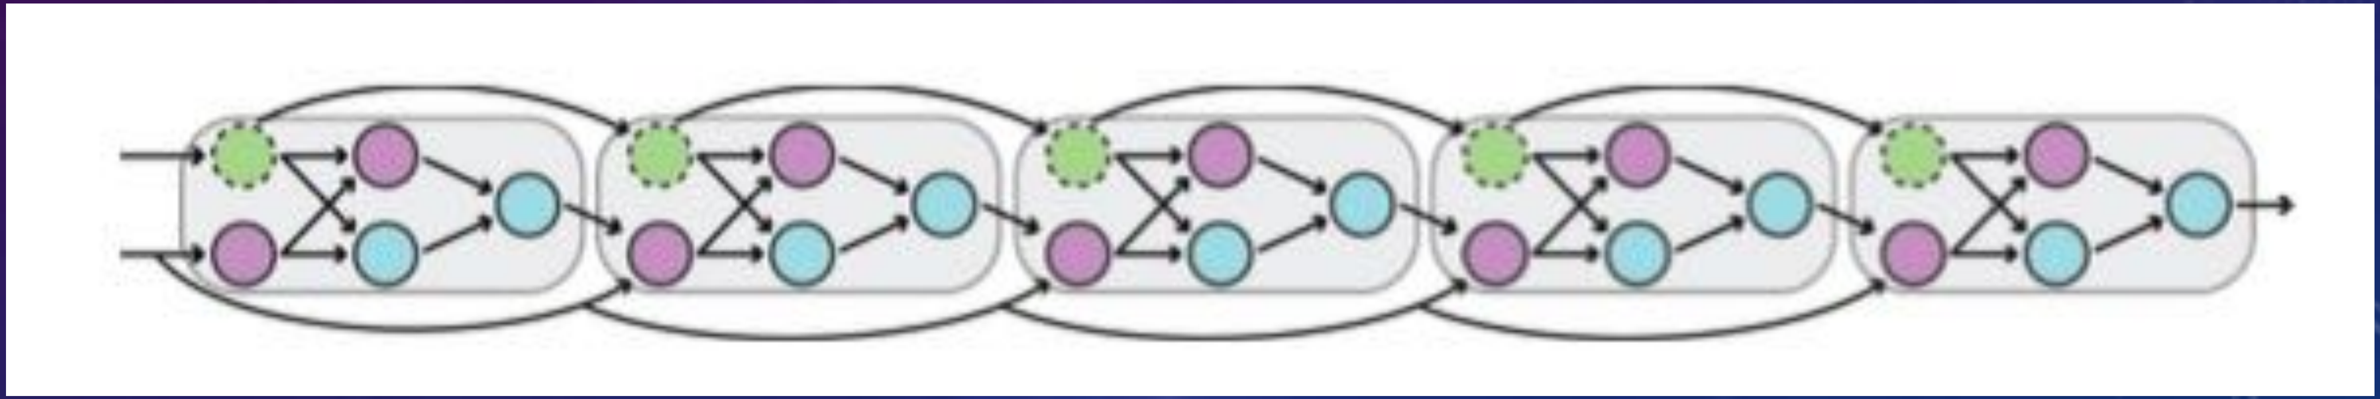

COMPUTATIONAL GRAPHS AND TENSOR FLOW

1. Define Computational Graph
2. Run the Graph (including the data)

MULTIPLICATION FUNCTION

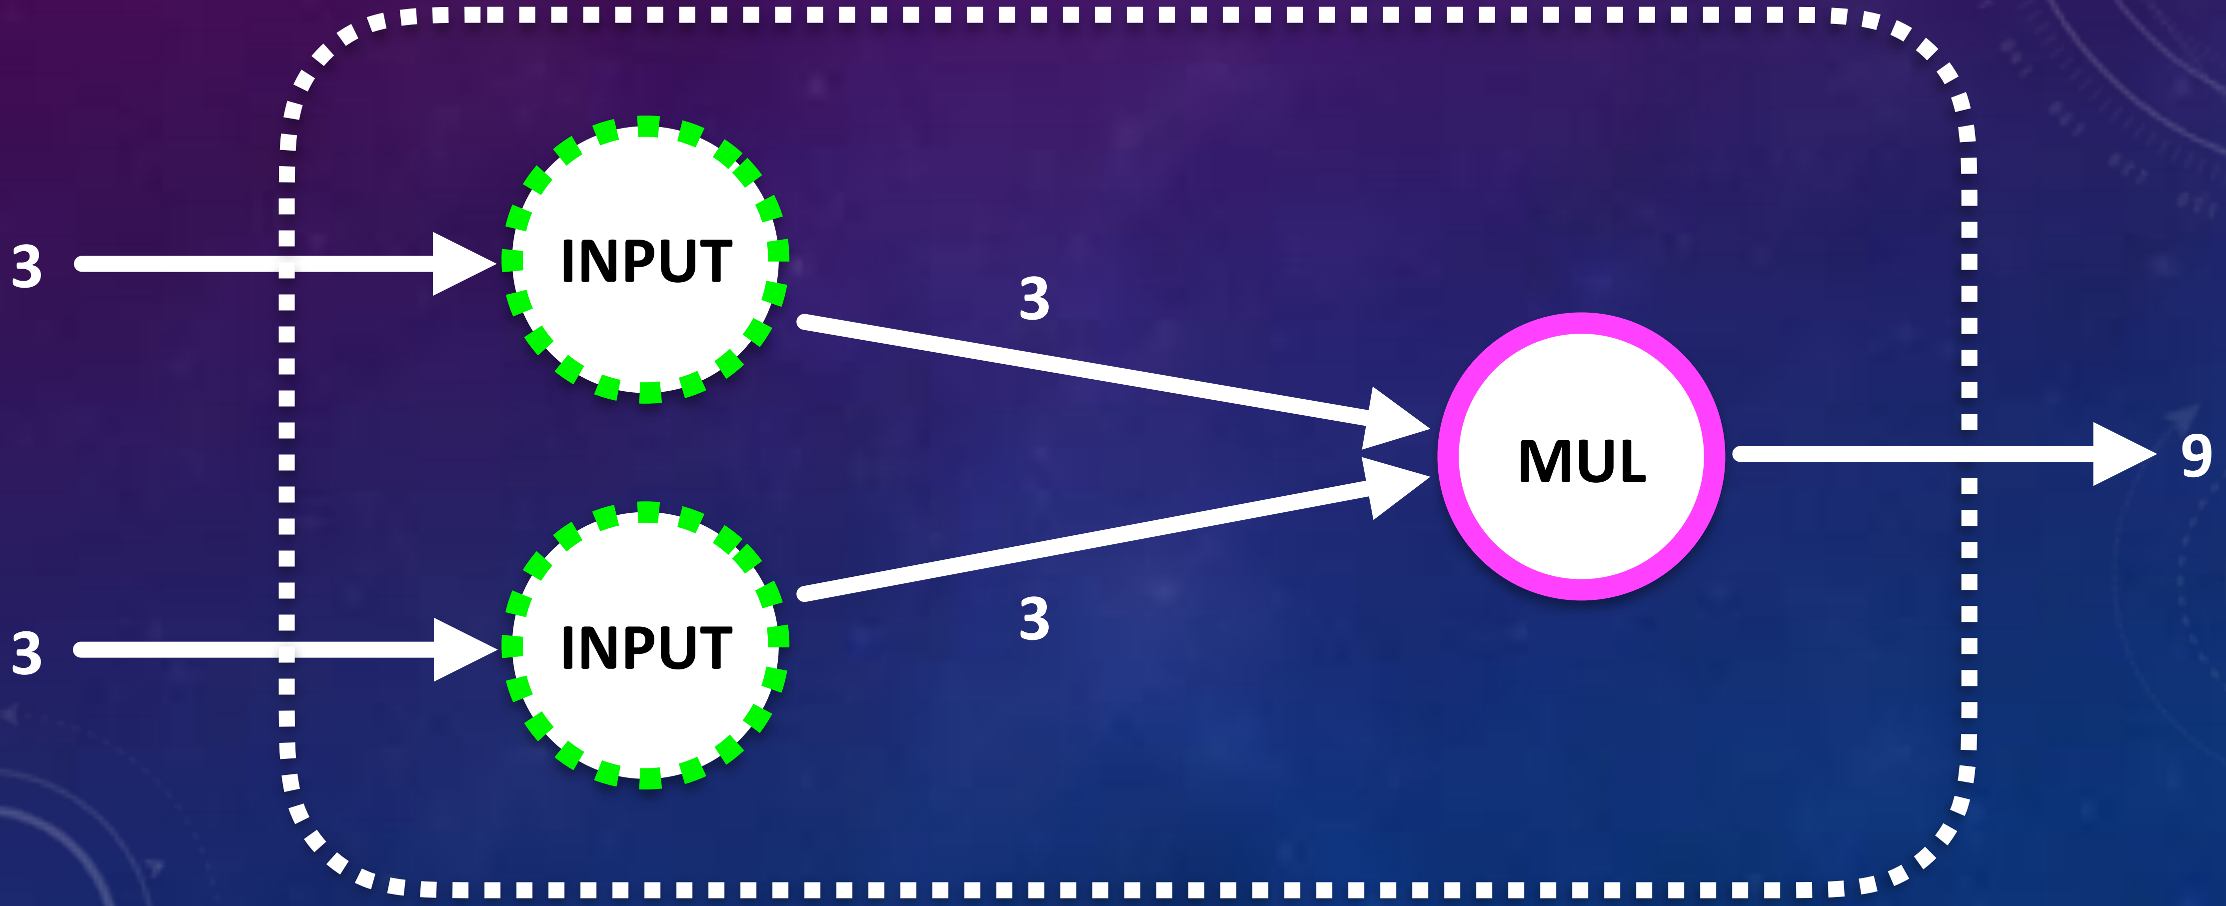

TENSORFLOW - MULTIPLICATION

```
import tensorflow as tf

a = tf.placeholder("float")
b = tf.placeholder("float")

y = tf.mul(a, b)

sess = tf.Session()


print sess.run(y, feed_dict={a: 3, b: 3})
```

TENSORFLOW - LINEAR REGRESSION

```
import numpy as np
import matplotlib.pyplot as plt

num_points = 1000
vectors_set = []

for i in xrange(num_points):
          x1= np.random.normal(0.0, 0.55)
          y1= x1 * 0.1 + 0.3 + np.random.normal(0.0, 0.03)
          vectors_set.append([x1, y1])

x_data = [v[0] for v in vectors_set]
y_data = [v[1] for v in vectors_set]

plt.plot(x_data, y_data, 'ro', label='Original data')
plt.legend()
plt.show()
```

TENSORFLOW LINEAR REGRESSION

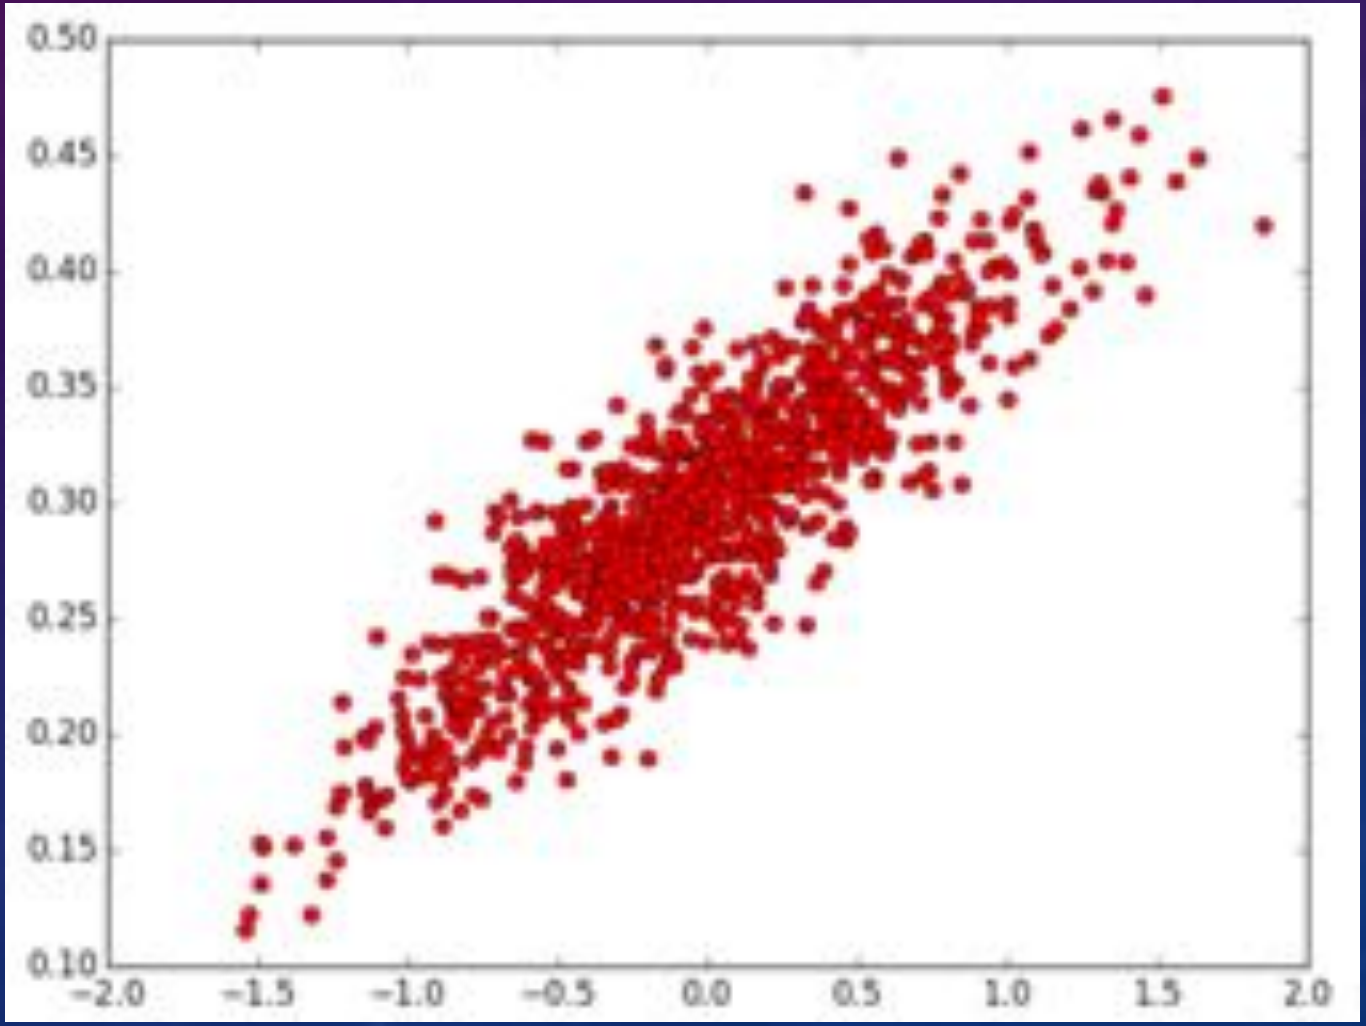

TENSORFLOW LINEAR REGRESSION

```
W = tf.Variable(tf.random_uniform([1], -1.0, 1.0))
b = tf.Variable(tf.zeros([1]))
y = W * x_data + b

loss = tf.reduce_mean(tf.square(y - y_data))

optimizer = tf.train.GradientDescentOptimizer(0.5)
train = optimizer.minimize(loss)


init = tf.initialize_all_variables()
sess = tf.Session()
sess.run(init)

for step in xrange(8):
    sess.run(train)
print step, sess.run(W), sess.run(b)


plt.plot(x_data, y_data, 'ro')
plt.plot(x_data, sess.run(W) * x_data + sess.run(b))
plt.legend()
plt.show()
```

TENSOR FLOW - MNIST DATASET

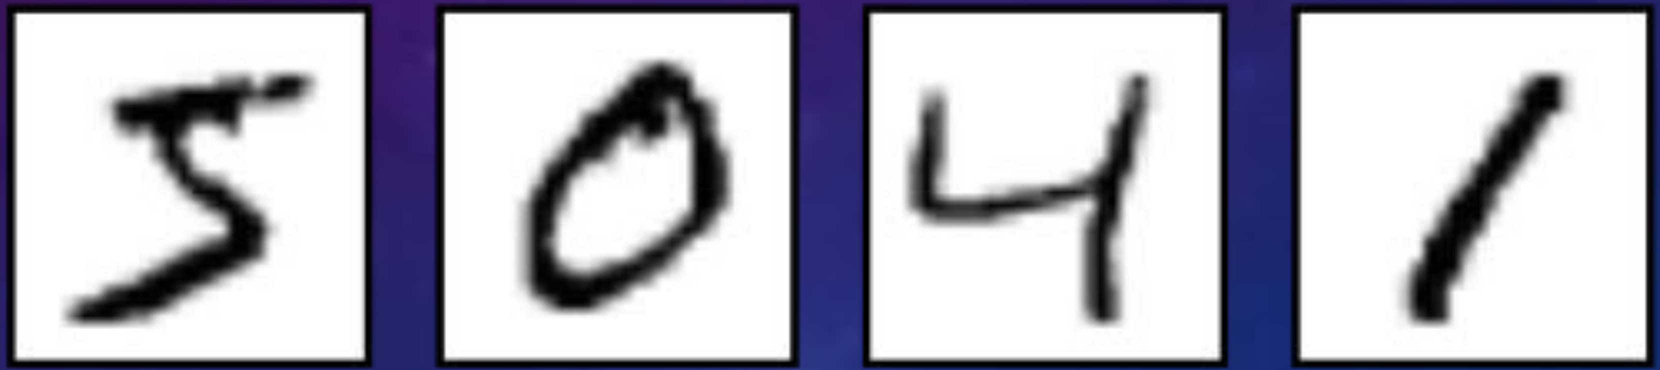

MNIST DATASET USING TENSORFLOW

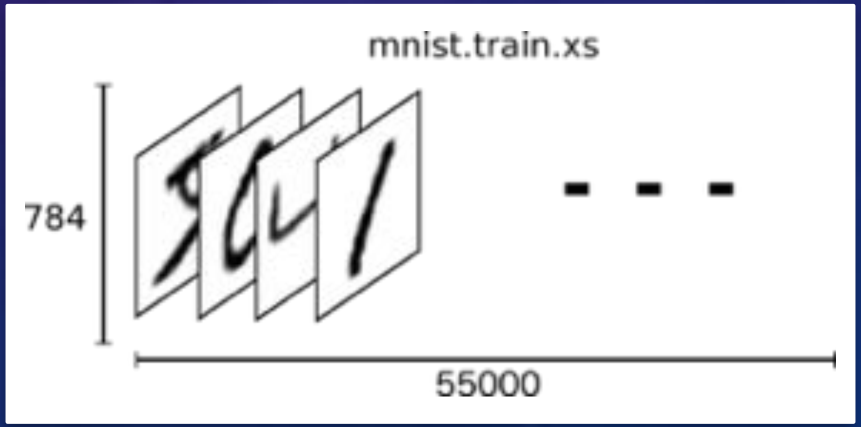

MNIST DATASET USING TENSORFLOW

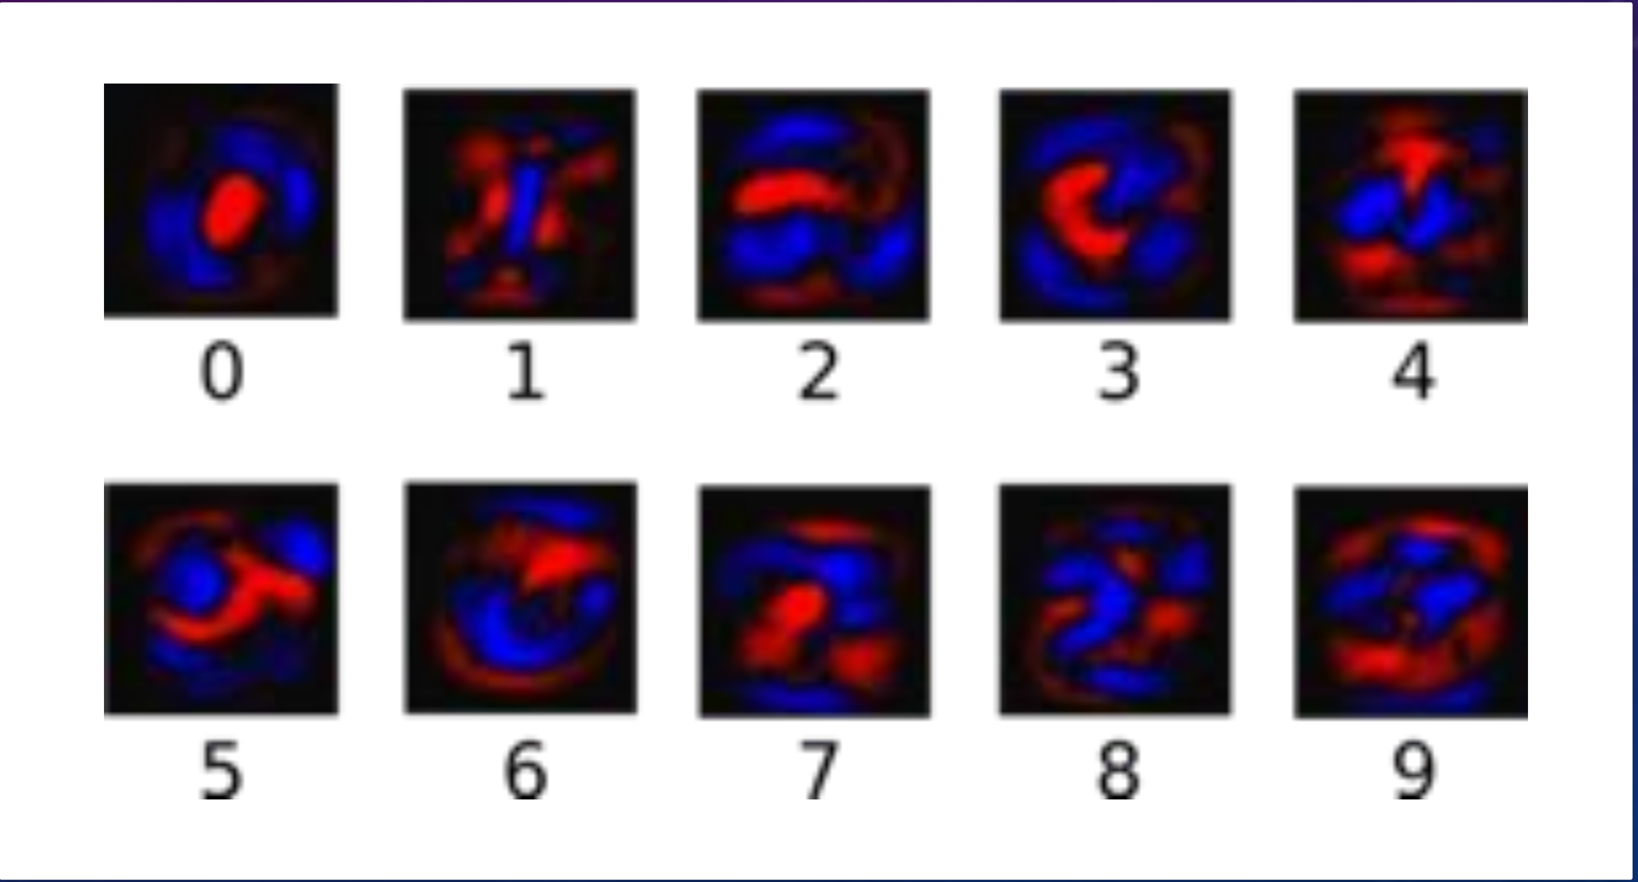

MNIST DATASET USING TENSORFLOW

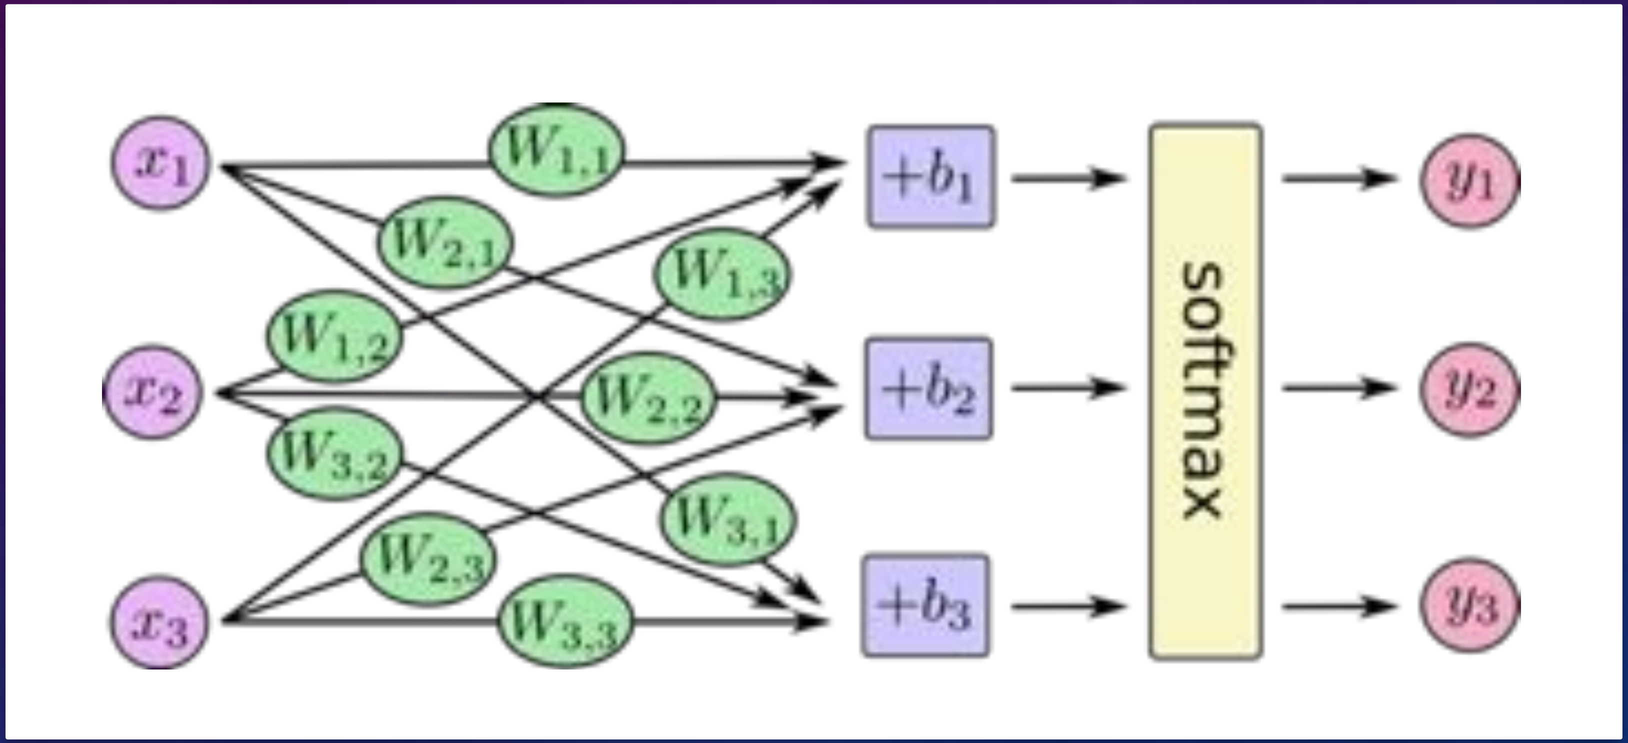

$$y = \text{softmax}(Wx + b)$$

MNIST DATASET USING TENSORFLOW - SETUP

```
import tensorflow as tf

x = tf.placeholder(tf.float32, [None, 784])
W = tf.Variable(tf.zeros([784, 10]))
b = tf.Variable(tf.zeros([10]))

y = tf.nn.softmax(tf.matmul(x, W) + b)

from tensorflow.examples.tutorials.mnist import input_data
mnist = input_data.read_data_sets("MNIST_data", one_hot=True)
```

$$y = \mathrm{softmax}(Wx + b)$$

MNIST DATASET USING TENSORFLOW - TRAINING

```
y_ = tf.placeholder(tf.float32, [None, 10])

cross_entropy = tf.reduce_mean(-tf.reduce_sum(y_ * tf.log(y), reduction_indices=[1]))

train_step = tf.train.GradientDescentOptimizer(0.5).minimize(cross_entropy)


init = tf.initialize_all_variables()
sess = tf.Session()
sess.run(init)
```

$$H_{y'}(y) = -\sum_i y'_i \log(y_i)$$

MNIST DATASET USING TENSORFLOW - TRAINING SETUP

```
for i in range(1000):
  batch_xs, batch_ys = mnist.train.next_batch(100)
  sess.run(train_step, feed_dict={x: batch_xs, y_: batch_ys})
```

MNIST DATASET USING TENSORFLOW - EVALUATION

```
correct_prediction = tf.equal(tf.argmax(y,1), tf.argmax(y_,1))
accuracy = tf.reduce_mean(tf.cast(correct_prediction, tf.float32))
```

WORD EMBEDDINGS

Word embeddings is a function mapping words in some language to a high-dimensional vector

$W: words \rightarrow \mathbb{R}^n$

$W(cat) = (0.3, 0.7, -2.3, \ldots)$

W is initialized randomly and then it learns meaningful vectors to perform some task (E.g. grammatical correctness)

E.g. learn if 5 words are valid

- cat sat on the mat
- cat sat song the mat

- R(W("cat"), W("sat"), W("on"), W("the"), W("mat")) = 1
- R(W("cat"), W("sat"), W("song"), W("the"), W("mat")) = 0

VISUALISE THE WORD EMBEDDINGS

EXAMPLE

- A few people eat well

Becomes

- A couple people eat well

Possible number of 5-grams is enormous

Few and Couple will be located close together thus allowing for generalization between

- Sentences to a
- Class of similar sentences

Handles not just synonyms but other classes like colors like …

- The wall is blue
- The wall is red

WHAT ELSE CAN WE LEARN?

Gender

Man	–	Woman
Uncle	–	Aunt
King	–	Queen

Several other relationships

France	–	Paris
Big	–	Bigger
Einstein	–	Scientist
Berlusconi	–	Silvio
Japan	–	Sushi
Messi	–	Midfielder

THE USE OF WORD REPRESENTATIONS… HAS BECOME A KEY “SECRET SAUCE” FOR THE SUCCESS OF MANY NLP SYSTEMS IN RECENT YEARS, ACROSS TASKS INCLUDING NAMED ENTITY RECOGNITION, PART-OF-SPEECH TAGGING, PARSING, AND SEMANTIC ROLE LABELING.

LUONG ET AL. (2013)

WHAT DOES IT HAS TO DO WITH DEEP LEARNING?

It's a tactic widely utilised as well

- Learn a good representation for task A
- Use it on task B

Called

- Retraining
- Transfer learning
- Multi-task learning

Using this approach, representation can learn from more than one kind of data

LEARN TO MAP MULTIPLE KINDS OF DATA

Another trick used in deep learning is to learn mappings between data in a single representation

E.g. English Words and Mandarin Words

- Words with similar meanings appear close together
- So do words we didn't know about

So using what we've seen so far

- If we know two male matching words in English and Chinese
- We can find the two female equivalent

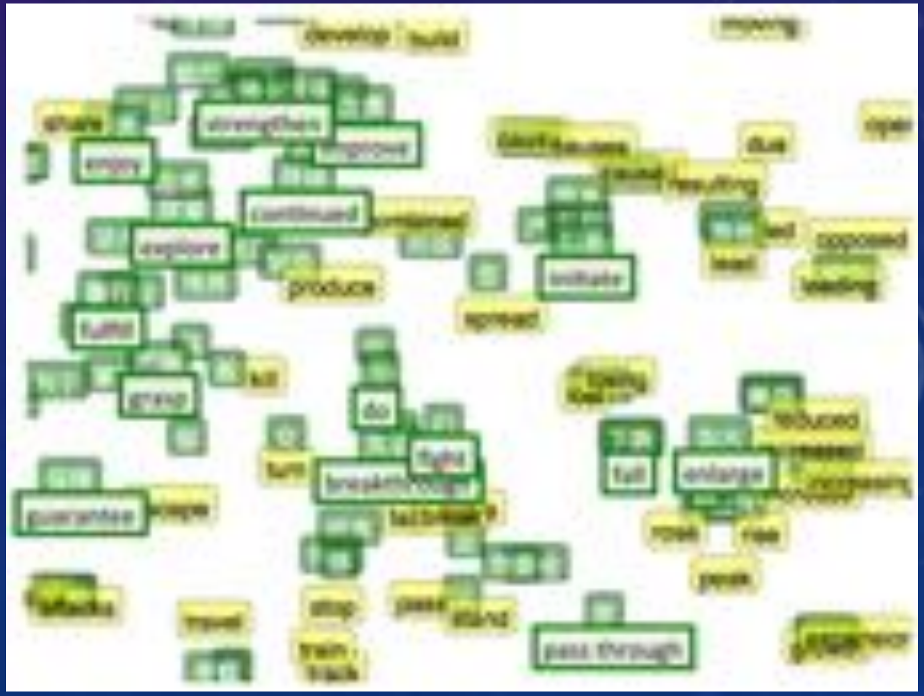

USING DEEP LEARNING TO GO FURTHER

So far we've dealt with words

But what about images?

- Images of dogs, horses and automobiles are mapped next to their word vector

What happens with new classes such as cats?

- They are mapped close to the dog vector
- But a little distant from it

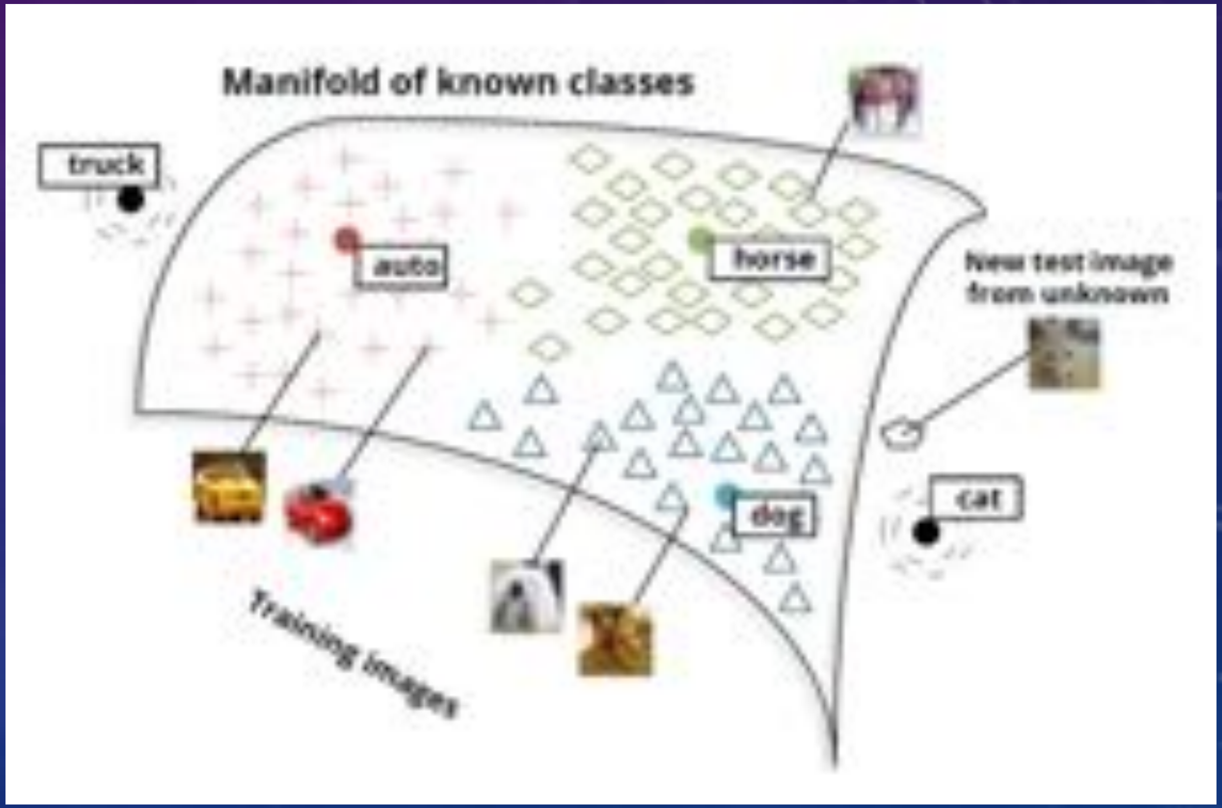


sky
building
tree
STOP
car
road
grass
person

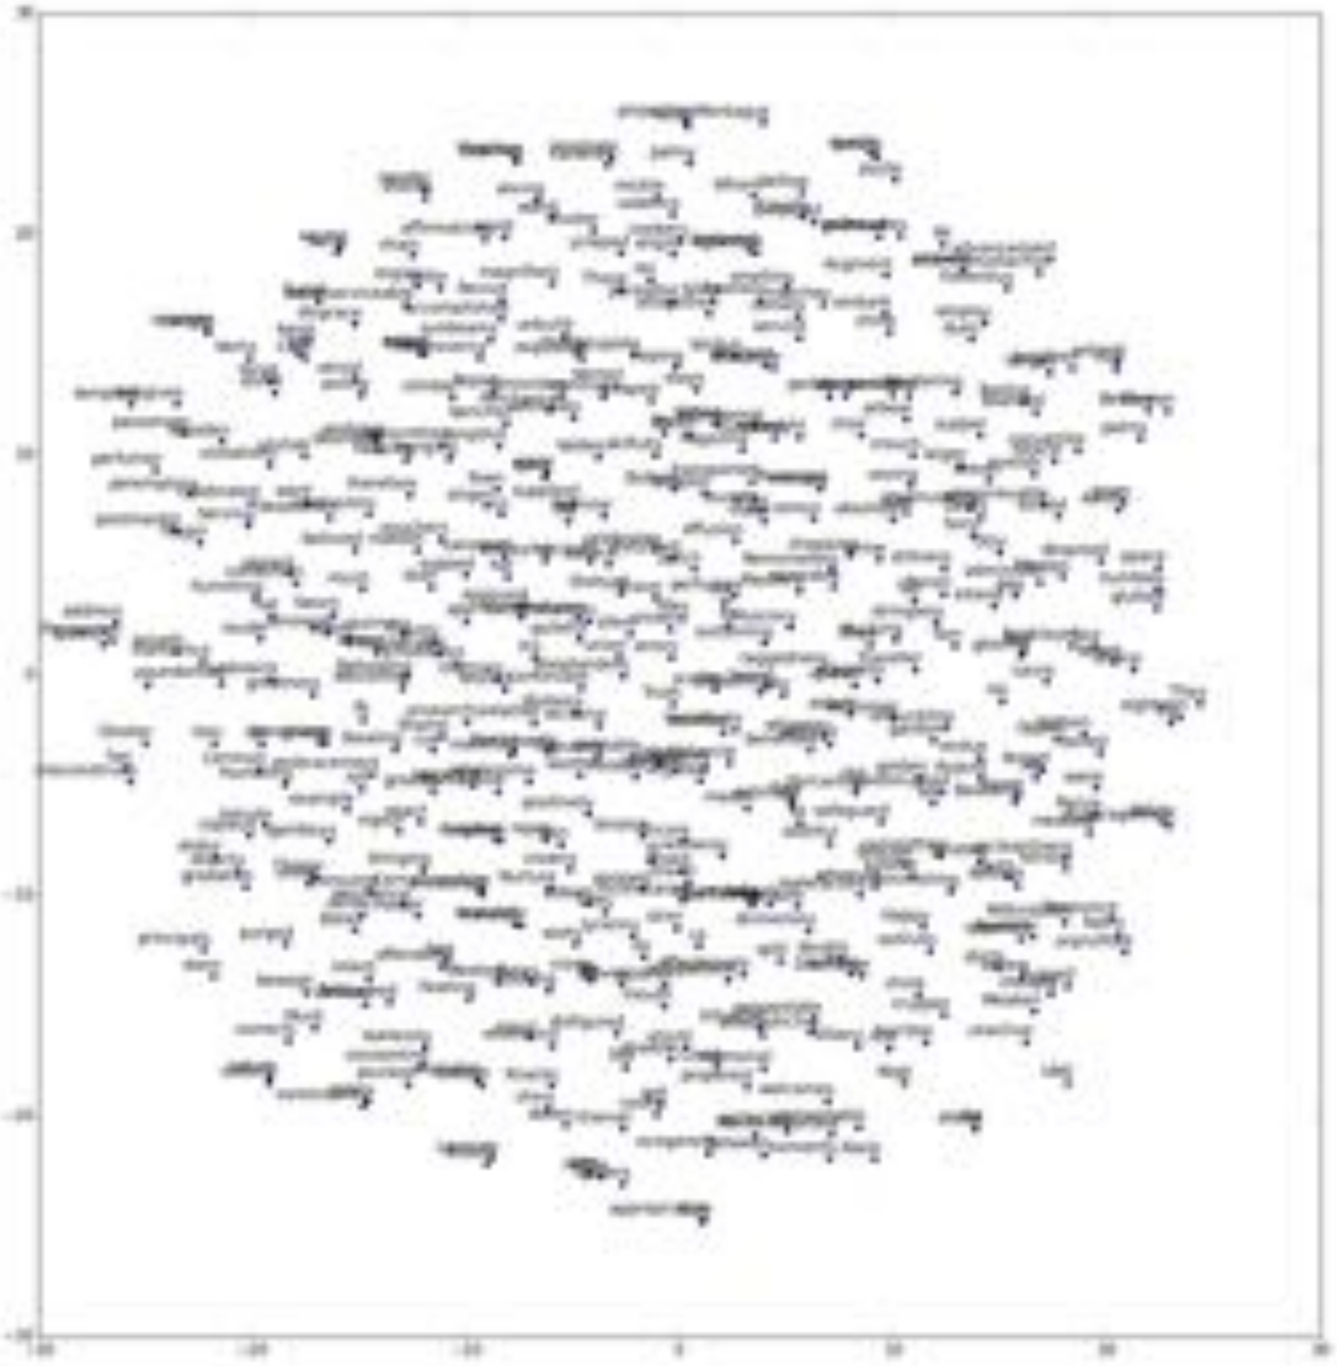

DEMO - AUTO ENCODERS

http://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html

DEMO - CNN MNIST

http://cs.stanford.edu/people/karpathy/convnetjs/demo/mnist.html

DEMO - CNN

http://cs.stanford.edu/people/karpathy/convnetjs/demo/cifar10.html

CONCLUSION

USES

- Predict stock markets
- Predict customer Churn
- Perform sentiment analysis
- Decode ancient texts automatically
- Autonomous Vehicles

QUESTIONS ?

alexiei.dingli@um.edu.mt