



lrz

FUNDAMENTALS OF DEEP LEARNING MATERIAL

PD. Dr. Juan J. Durillo



DEEP
LEARNING
INSTITUTE

THE GOALS OF THIS COURSE

- Get you up and on your feet quickly
- Build a foundation to tackle a deep learning project right away
- We won't cover the whole field, but we'll get a great head start
- Foundation from which to read articles, follow tutorials, take further classes

AGENDA

Part 1: An Introduction to Deep Learning

Part 2: How a Neural Network Trains

Part 3: Convolutional Neural Networks

Part 4: Data Augmentation and Deployment

Part 5: Pre-trained Models

Part 6: Advanced Architectures

A network diagram on a dark background. The diagram consists of numerous small, glowing nodes connected by thin, light-colored lines. The nodes are primarily white and yellow, with some green nodes. The connections form a complex, interconnected web that is denser in the upper left and upper right areas, with some lines extending towards the center and lower left. The overall effect is that of a digital or social network.

HAVE FUN!

A network diagram consisting of numerous small circular nodes connected by thin, light-colored lines. The nodes are scattered across the frame, with a higher density in the upper right quadrant. Some nodes are white, while others are a bright yellow-green. The lines connecting them form a complex web of connections. The background is a dark, solid color.

HISTORY OF AI

BEGINNING OF ARTIFICIAL INTELLIGENCE



COMPUTERS ARE MADE IN
PART TO COMPLETE HUMAN
TASKS




EARLY ON, GENERALIZED
INTELLIGENCE LOOKED
POSSIBLE



TURNED OUT TO BE HARDER
THAN EXPECTED

EARLY NEURAL NETWORKS



Inspired by biology

Created in the 1950's

Outclassed by Von Neumann Architecture

EXPERT SYSTEMS



Highly complex



Programmed by hundreds of engineers



Rigorous programming of many rules

EXPERT SYSTEMS - LIMITATIONS

What are these three images?



A network graph visualization on a dark background. The graph consists of numerous nodes, some of which are highlighted in a bright yellow color, while others are white. These nodes are interconnected by a dense web of thin, light-colored lines, representing connections or relationships. The overall structure is complex and multi-layered, typical of a neural network or a social network graph.

THE DEEP LEARNING REVOLUTION

DATA

- Networks need a lot of information to learn from
- The digital era and the internet has supplied that data



COMPUTING POWER

Need a way for our artificial “brain” to observe lots of data within a practical amount of time.




THE IMPORTANCE OF THE GPU

A Rendered Image



A Neural Network




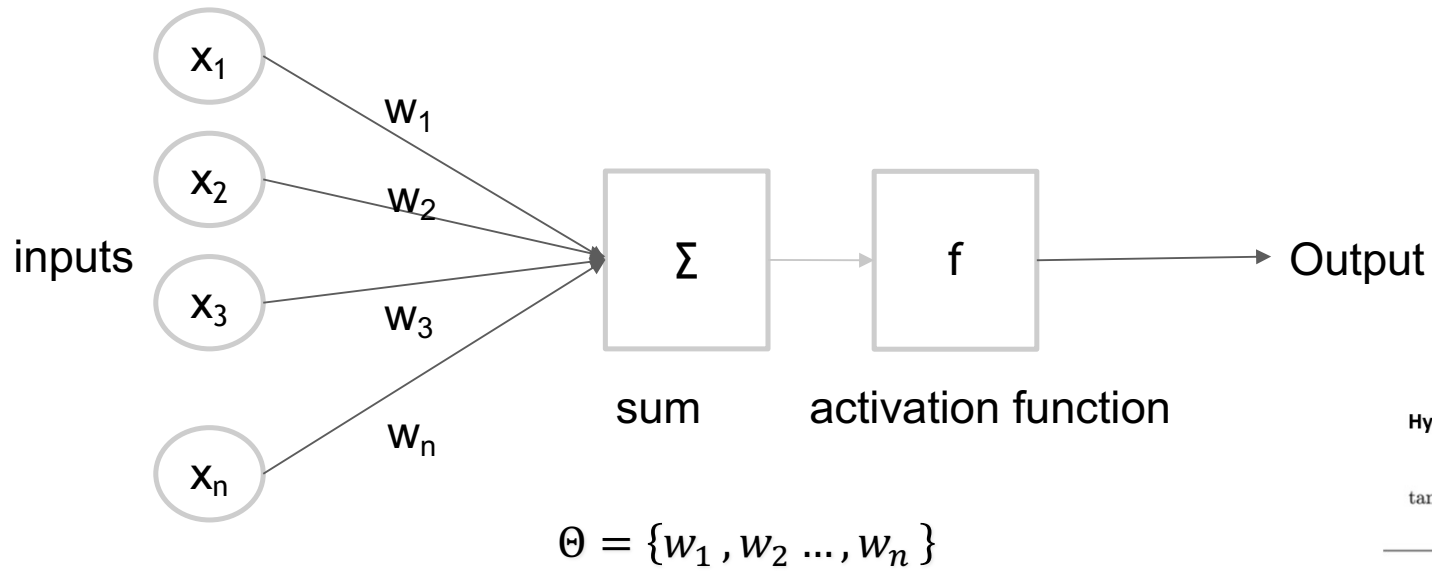


WHAT IS DEEP LEARNING?





A (brief) introduction to ML and DL

PD Dr. Juan J. Durillo

Perceptron - Artificial Neuron




Single artificial neurons work well for linearly separable datasets (indeed output is the activation effect on a linear combination of the input)




most popular activation functions

NEURAL NETWORK



$$\Theta = \{w_{1,1}^1, w_{1,2}^1, w_{1,3}^1, w_{2,1}^1, w_{2,2}^1, w_{2,3}^1, w_{1,1}^2, w_{2,1}^2, w_{2,3}^2\}$$

- Works well even when the data is not linearly separable



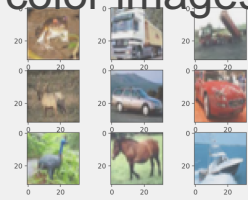
(SUPERVISED) LEARNING

- Data domain Z : $X \times Y$

$X \rightarrow$ domain of the input data

$Y \rightarrow$ set of labels (knowledge)

X : 32 x 32
color images



Y : labels
{truck, car, horse, bird, boat}

Example (CIFAR10 dataset)

- Data Distribution is a probability distribution over a data domain
- Training set z_1, \dots, z_n from Z assumed to be drawn from the Data Distribution D
- Validation set v_1, \dots, v_m from Z also assumed to be drawn from D
- A machine learning model is a function that given a set of parameters θ and z from Z produces a prediction
- The prediction quality is measured by a differentiable non-negative scalar-valued loss function, that we denote $\ell(\theta; z)$

(SUPERVISED) LEARNING

- Given Θ we can define the expected loss as: $L(\Theta) = \mathbb{E}_{z \sim D}[\ell(\Theta; z)]$
- Given D , ℓ , and a model with parameter set Θ , we can define learning as:
“The task of finding parameters Θ that achieve low values of the expected loss, while we are given access to only n training examples”
- The mentioned task before is commonly referred to as *training*
- Empirical average loss given a subset of the training data set $S(z_1, \dots, z_n)$ as:

$$\hat{L}(\Theta) = \frac{1}{n} \sum_{t=1}^n [\ell(\Theta; z_t)]$$

- Usually a proxy function, easier to understand by humans, is used for describing how well the training is performed (e.g., accuracy)

(SUPERVISED) LEARNING

- The dominant algorithms for training neural networks are based on mini-batch stochastic gradient descent (SGD)
- Given an initial point Θ_0 SGD attempt to decrease \hat{L} via the sequence of iterates

$$\Theta_t \leftarrow \Theta_{t-1} - n_t g(\Theta_{t-1}; B_t)$$

$$g(\Theta; B) = \frac{1}{|B|} \sum_{z \in B} \nabla \ell(\Theta; z)$$

B_t : random subset of training examples

n_t : positive scalar (learning rate)


epoch: update the weights after going over all training set

COMPUTER VISION TASKS




predicting the type or class of an object in an image

Image Classification



predicting the type or class on an object in an image and draw a bounding box around
Image Classification + Localization



predicting the location of objects in an image via bounding boxes and the classes of the located objects

Object Detection

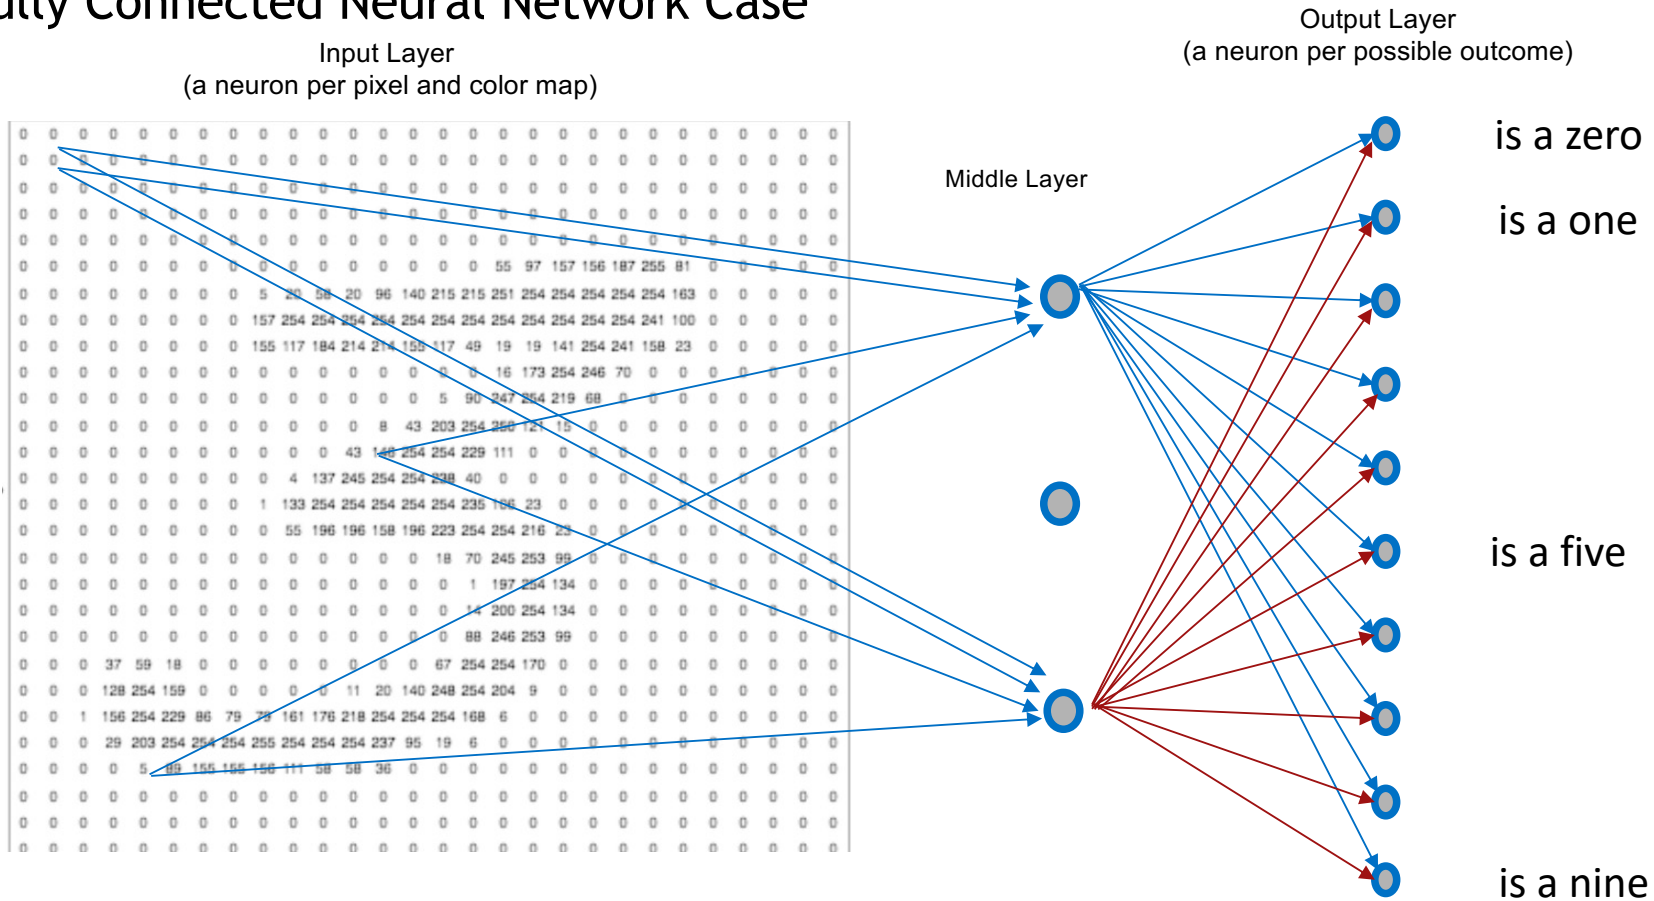


predicting the class to which each pixel in the image belongs to



Image Segmentation

NEURAL NETWORKS FOR IMAGE CLASSIFICATION

Fully Connected Neural Network Case



TRAINING NEURAL NETWORKS



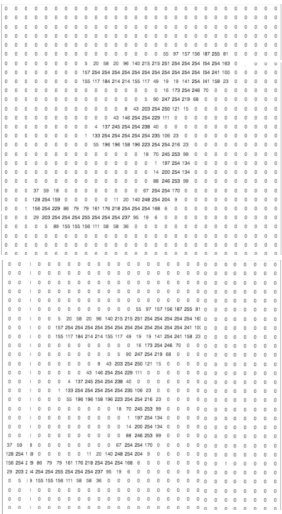
how the surface looks like in reality

Stochastic Gradient Descent



$$\theta_t \leftarrow \theta_{t-1} - n_t g(\theta_{t-1}; B_t)$$

$$g(\theta; B) = \frac{1}{|B|} \sum_{z \in B} \nabla \ell(\theta; z)$$

NEURAL NETWORKS FOR IMAGE CLASSIFICATION





shift to the left



is a zero
is a one
is a five
is a nine

NO MORE FEATURE ENGINEERING



LEARNING FEATURES FROM DATA: CONVOLUTIONS

Input Image

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	0	1	0	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

receptive field

Filter

-1	0	1
-2	1	2
-3	0	3

$$\begin{cases} 1 \times (-1) + 0 \times 0 + 1 \times 1 + \\ 0 \times (-2) + 1 \times 1 + 0 \times 2 + \\ 0 \times (-3) + 0 \times 0 + 1 \times 3 = 4 \end{cases}$$

Convolved Image

		4					

Filter is convolved with all the pixels of the image

1	-1	0	1	0	1	0	1
0	-2	1	2	1	0	1	0
0	-3	0	3	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

How many units the filter moves horizontally or vertically is called **stride** and can be different in both dimensions

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	-1	0	1	0	1	0	1
1	-2	1	2	0	1	0	0
0	-3	0	3	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

The stride defines the size of the convolved image

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	0	1	0	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	-1	0	1
0	0	0	0	0	-2	1	2
0	0	1	0	0	-3	0	3

FILTERS

Input Image:



Can we get only vertical lines out of this picture?

1 0 -1

filter 1

1 0 -1

1 0 -1

1 0 -1

filter 2

1 0 0 0 -1

1 0 0 0 -1

1 0 0 0 -1

1 0 0 0 -1

1 0 0 0 -1

filter 3

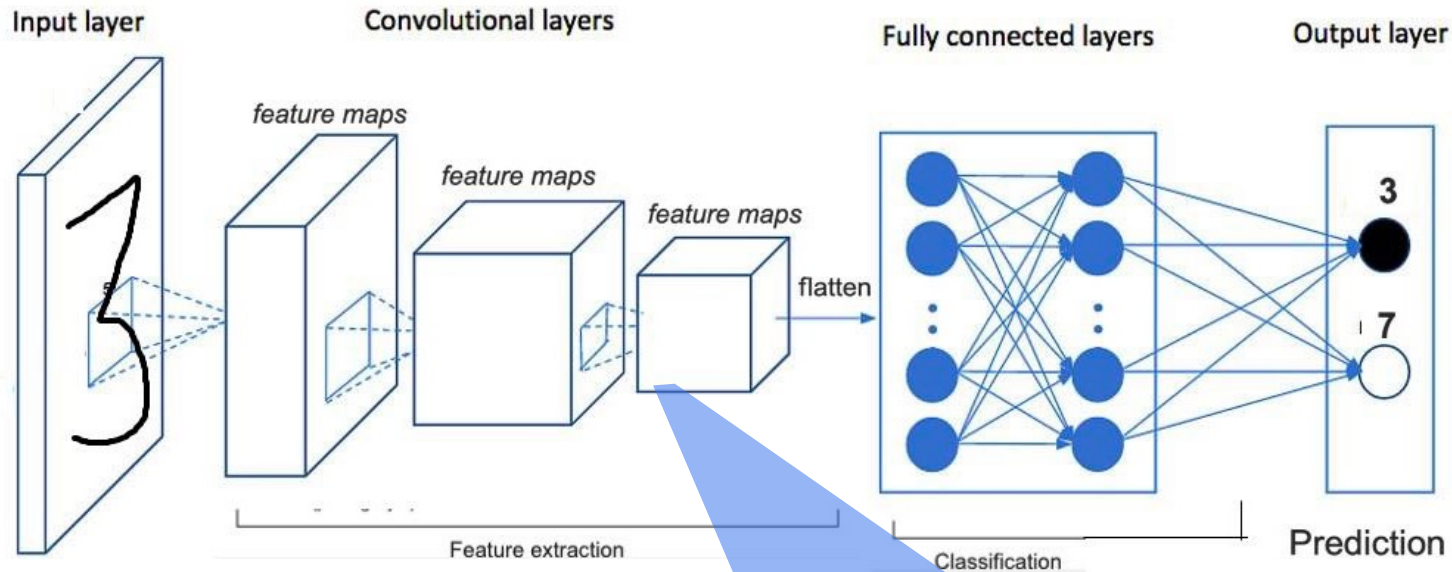
try the code yourself (in octave)!

```
I=imread(<path-to-image>);  
GRAY=rgb2gray(I)  
FILTER=[ 1 0 -1; 1 0 -1; 1 0 -1]; % filter 2  
CONVOLUTED=conv2(GRAY,FILTER);  
Imwrite(CONVOLUTED, <path-to-result>);
```

*The London skyline image is designed by Freepik





CONVOLUTIONAL NEURAL NETWORKS (CNN)



Feature extraction

Classification


Prediction



A pooling layer down sample the feature maps produced by a convolution into smaller number of parameters to reduce the computational complexity.

It is a common practice to add pooling layers after each one or two convolutions layers in the CNN architecture.


CNN ARCHITECTURE: A COMMON PATTERN AND ITS INFLUENCE



The execution time required during a forward pass through a neural network is bounded from below by the number of floating point operations (FLOPs).

This FLOP count depends on the deep neural network architecture and the amount of data.


LENET ARCHITECTURE



Architecture summary :

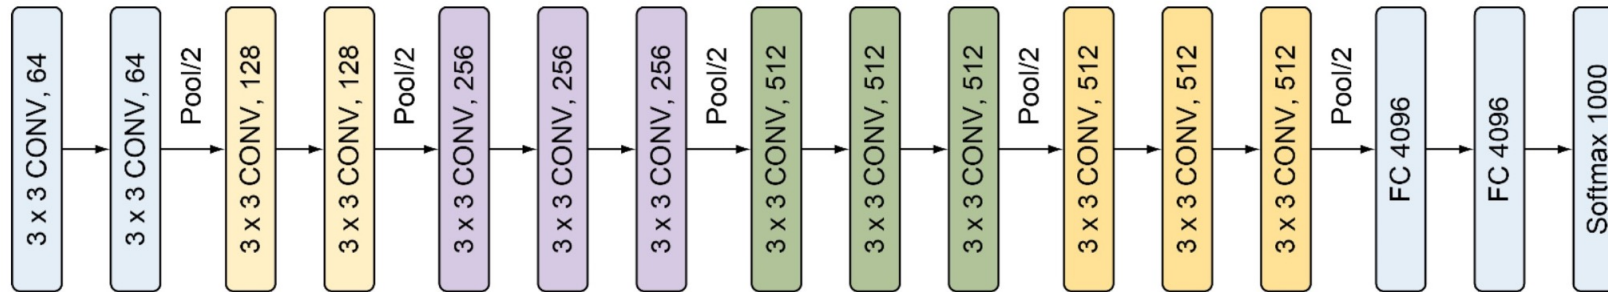
- 3 convolutional layers filters in all the layers equal to 5x5
(layer 1 depth = 6, layer 2 depth = 16, layer 3 depth = 120)
- As activation function the tanh function is used

ALEXNET AND VGG ARCHITECTURES



~60,000,000 parameters


AlexNet



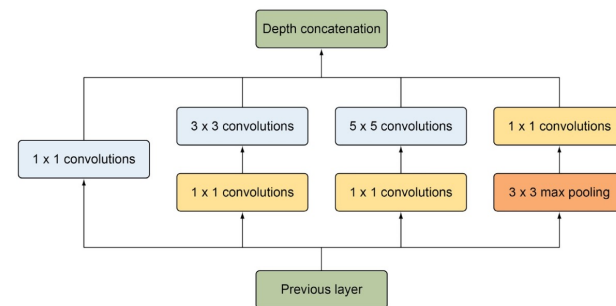
~138,000,000 parameters

VGG16



GOOGLENET



- What is the best kernel size for each layer?
- Concatenating filters instead of stacking them for reducing computational expenses




RESTNET



$\hat{L}(\theta)$

θ_2

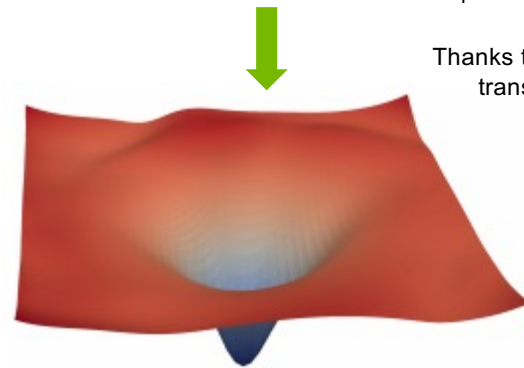


θ_1

Thanks to the shortcut is transformed into

$\hat{L}(\theta)$


θ_2



θ_1


INCREASING COMPLEXITY

7 Exaflops
60 Million Parameters




2015 - Microsoft ResNet
Superhuman Image Recognition

20 Exaflops
300 Million Parameters



2016 - Baidu Deep Speech 2
Superhuman Voice Recognition

100 Exaflops
8700 Million Parameters



2017 - Google Neural Machine Translation
Near Human Language Translation

SUMMARY

Brief introduction to Deep Learning with emphasis in Deep Convolutional Neural Networks

Review of basic concepts: from perceptron to the learning task

Debrief of most important concepts of neural network architectures

A network diagram with nodes and connections on a dark background. The nodes are represented by small circles, some of which are highlighted in a bright yellow-green color. The connections are thin, light-colored lines that form a complex web of relationships between the nodes. The overall aesthetic is modern and technical, typical of a presentation on artificial intelligence or data science.

DEEP LEARNING FLIPS TRADITIONAL PROGRAMMING ON ITS HEAD

TRADITIONAL PROGRAMMING

Building a Classifier

1

Define a set of
rules for
classification

2

Program those
rules into the
computer

3

Feed it examples,
and the program
uses the rules to
classify

MACHINE LEARNING

Building a Classifier

1

Show model the examples with the answer of how to classify

2

Model takes guesses, we tell it if it's right or not

3

Model learns to correctly categorize as it's training. The system learns the rules on its own



THIS IS A FUNDAMENTAL SHIFT

WHEN TO CHOOSE DEEP LEARNING

Classic Programming

If rules are clear
and
straightforward,
often better to just
program it

Deep Learning

If rules are
nuanced, complex,
difficult to discern,
use deep learning

DEEP LEARNING COMPARED TO OTHER AI

Depth and complexity of networks

Up to billions of parameters (and growing)


Many layers in a model

Important for learning complex rules

A network diagram consisting of numerous small circular nodes connected by thin, light-colored lines. The nodes are scattered across the frame, with a higher density in the upper right quadrant. Some nodes are white, while others are a bright green. The lines connecting them form a complex web of connections. The background is a dark, solid color.

HOW DEEP LEARNING IS TRANSFORMING THE WORLD

COMPUTER VISION



ROBOTICS AND
MANUFACTURING



OBJECT
DETECTION



SELF DRIVING
CARS

NATURAL LANGUAGE PROCESSING



REAL TIME
TRANSLATION



VOICE
RECOGNITION



VIRTUAL
ASSISTANTS

RECOMMENDER SYSTEMS



CONTENT
CURATION



TARGETED
ADVERTISING



SHOPPING
RECOMMENDATIONS

REINFORCEMENT LEARNING



ALPHAGO BEATS
WORLD CHAMPION
IN GO



AI BOTS BEAT
PROFESSIONAL
VIDEOGAMERS



STOCK TRADING
ROBOTS

A network diagram consisting of numerous small circular nodes connected by thin, light-colored lines. The nodes are scattered across the frame, with a higher density in the upper right quadrant. Some nodes are white, while others are a bright yellow-green. The lines connecting them form a complex web of connections. The background is a dark, solid color.

OVERVIEW OF THE COURSE

HANDS ON EXERCISES

- Get comfortable with the process of deep learning
- Exposure to different models and datatypes
- Get a jump-start to tackle your own projects



STRUCTURE OF THE COURSE

“Hello World” of Deep Learning



```
graph TD; A["“Hello World” of Deep Learning"] --> B["Train a more complicated model"]; B --> C["New architectures and techniques to improve performance"]; C --> D["Pre-trained models"]; D --> E["Transfer learning"];
```

Train a more complicated model

New architectures and techniques to improve performance

Pre-trained models

Transfer learning

PLATFORM OF THE COURSE



GPU powered cloud server



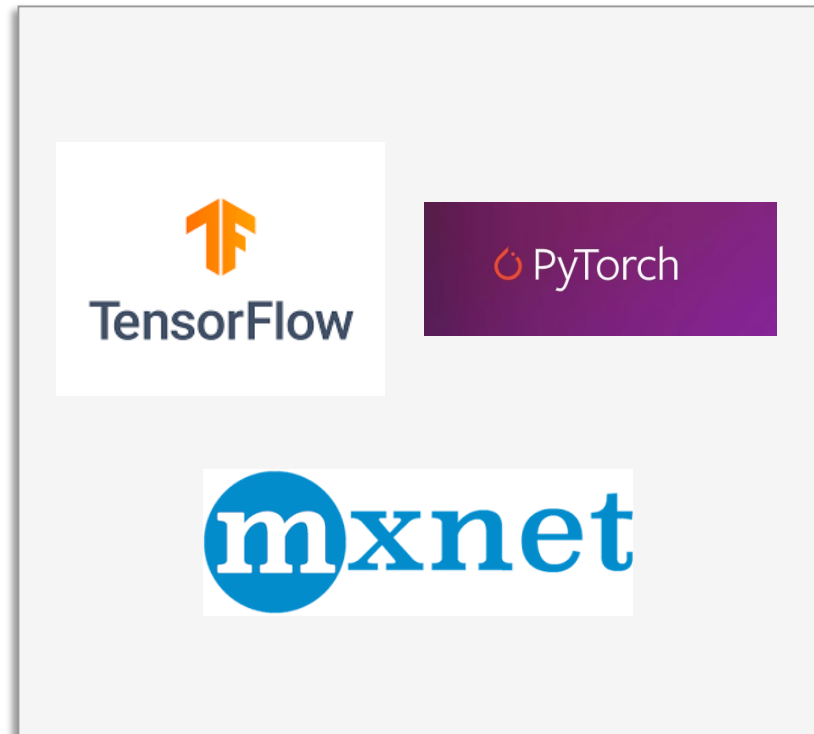
JupyterLab platform



Jupyter notebooks for interactive coding

SOFTWARE OF THE COURSE

- Major deep learning platforms:
 - TensorFlow + Keras (Google)
 - Pytorch (Facebook)
 - MXNet (Apache)
- We'll be using TensorFlow and Keras
- Good idea to gain exposure to others moving forward



A network graph visualization on a dark background. The graph consists of numerous nodes, some of which are highlighted in a bright yellow color, while others are white. These nodes are interconnected by a dense web of thin, light-colored lines, representing connections or relationships between data points. The overall structure is complex and somewhat chaotic, with many lines crisscrossing the space.

**FIRST EXERCISE:
CLASSIFY HANDWRITTEN
DIGITS**

HELLO NEURAL NETWORKS

Train a network to correctly classify handwritten digits


- Historically important and difficult task for computers

Try learning like a Neural Network

- Get exposed to the example, and try to figure out the rules to how it works



LET'S GO!




DEEP
LEARNING
INSTITUTE



DEEP
LEARNING
INSTITUTE

FUNDAMENTALS OF DEEP LEARNING

Part 2: How a Neural Network Trains



Part 1: An Introduction to Deep Learning

Part 2: How a Neural Network Trains

Part 3: Convolutional Neural Networks

Part 4: Data Augmentation and Deployment

Part 5: Pre-trained Models

Part 6: Advanced Architectures


RECAP OF THE EXERCISE

What just happened?

Loaded and visualized our data



Edited our data (reshaped, normalized, to categorical)



Created our model




Compiled our model



Trained the model on our data

DATA PREPARATION

Input as an array



DATA PREPARATION

Targets as categories

0 → [1,0,0,0,0,0,0,0,0,0]


1 → [0,1,0,0,0,0,0,0,0,0]

2 → [0,0,1,0,0,0,0,0,0,0]

3 → [0,0,0,1,0,0,0,0,0,0]

•
•
•

AN UNTRAINED MODEL







A SIMPLER MODEL

A SIMPLER MODEL

$$y = mx + b$$

x	y
1	3
2	5





$$m = ?$$

$$b = ?$$

A SIMPLER MODEL

$$y = mx + b$$

x	y
1	3
2	5





$$m = ?$$

$$b = ?$$

A SIMPLER MODEL

$$y = mx + b$$

x	y	\hat{y}
1	3	4
2	5	3



Start
Random


$$m = -1$$

$$b = 5$$

A SIMPLER MODEL

$$y = mx + b$$

x	y	\hat{y}	err^2
1	3	4	1
2	5	3	4
MSE =			2.5
RMSE =			1.6



$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$$


A SIMPLER MODEL

$$y = mx + b$$

x	y	\hat{y}	err^2
1	3	4	1
2	5	3	4



MSE = 2.5

RMSE = 1.6




```
1 data = [(1, 3), (2, 5)]
2 m = -1
3 b = 5
4
5
6 def get_rmse(data, m, b):
7     """Calculates Mean Square Error"""
8     n = len(data)
9     squared_error = 0
10    for x, y in data:
11        # Find predicted y
12        y_hat = m*x+b
13        # Square difference between
14        # prediction and true value
15        squared_error += (
16            y - y_hat) ** 2
17    # Get average squared difference
18    mse = squared_error / n
19    # Square root for original units
20    return rmse ** .5
21
```


THE LOSS CURVE




THE LOSS CURVE




$m = -1$
 $b = 5$




THE LOSS CURVE




$m = -1$
 $b = 4$



THE LOSS CURVE




$m = 0$
 $b = 4$



THE LOSS CURVE

The Gradient	Which direction loss decreases the most
λ : The learning rate	How far to travel
Epoch	A model update with the full dataset
Batch	A sample of the full dataset
Step	An update to the weight parameters




THE LOSS CURVE

The Gradient	Which direction loss decreases the most
λ : The learning rate	How far to travel
Epoch	A model update with the full dataset
Batch	A sample of the full dataset
Step	An update to the weight parameters



OPTIMIZERS




- Adam
- Adagrad
- RMSprop
- SGD

A network diagram consisting of numerous nodes connected by thin, light-colored lines. The nodes are represented by small circles, some of which are white and others are yellow. The connections form a complex, interconnected web of lines, suggesting a neural network or a data network. The background is a dark, solid color, making the nodes and lines stand out.


FROM NEURON TO NETWORK

BUILDING A NETWORK




- Scales to more inputs

BUILDING A NETWORK



- Scales to more inputs
- Can chain neurons

BUILDING A NETWORK



- Scales to more inputs
- Can chain neurons
- If all regressions are linear, then output will also be a linear regression




ACTIVATION FUNCTIONS

ACTIVATION FUNCTIONS

Linear

$$\hat{y} = wx + b$$


```
1 # Multiply each input
2 # with a weight (w) and
3 # add intercept (b)
4 y_hat = wx+b
```



ReLU

$$\hat{y} = \begin{cases} wx + b & \text{if } wx + b > 0 \\ 0 & \text{otherwise} \end{cases}$$


```
1 # Only return result
2 # if total is positive
3 linear = wx+b
4 y_hat = linear * (linear > 0)
```



Sigmoid


$$\hat{y} = \frac{1}{1 + e^{-(wx+b)}}$$

```
1 # Start with line
2 linear = wx + b
3 # Warp to - inf to 0
4 inf_to_zero = np.exp(-1 * linear)
5 # Squish to -1 to 1
6 y_hat = 1 / (1 + inf_to_zero)
```




ACTIVATION FUNCTIONS


Linear



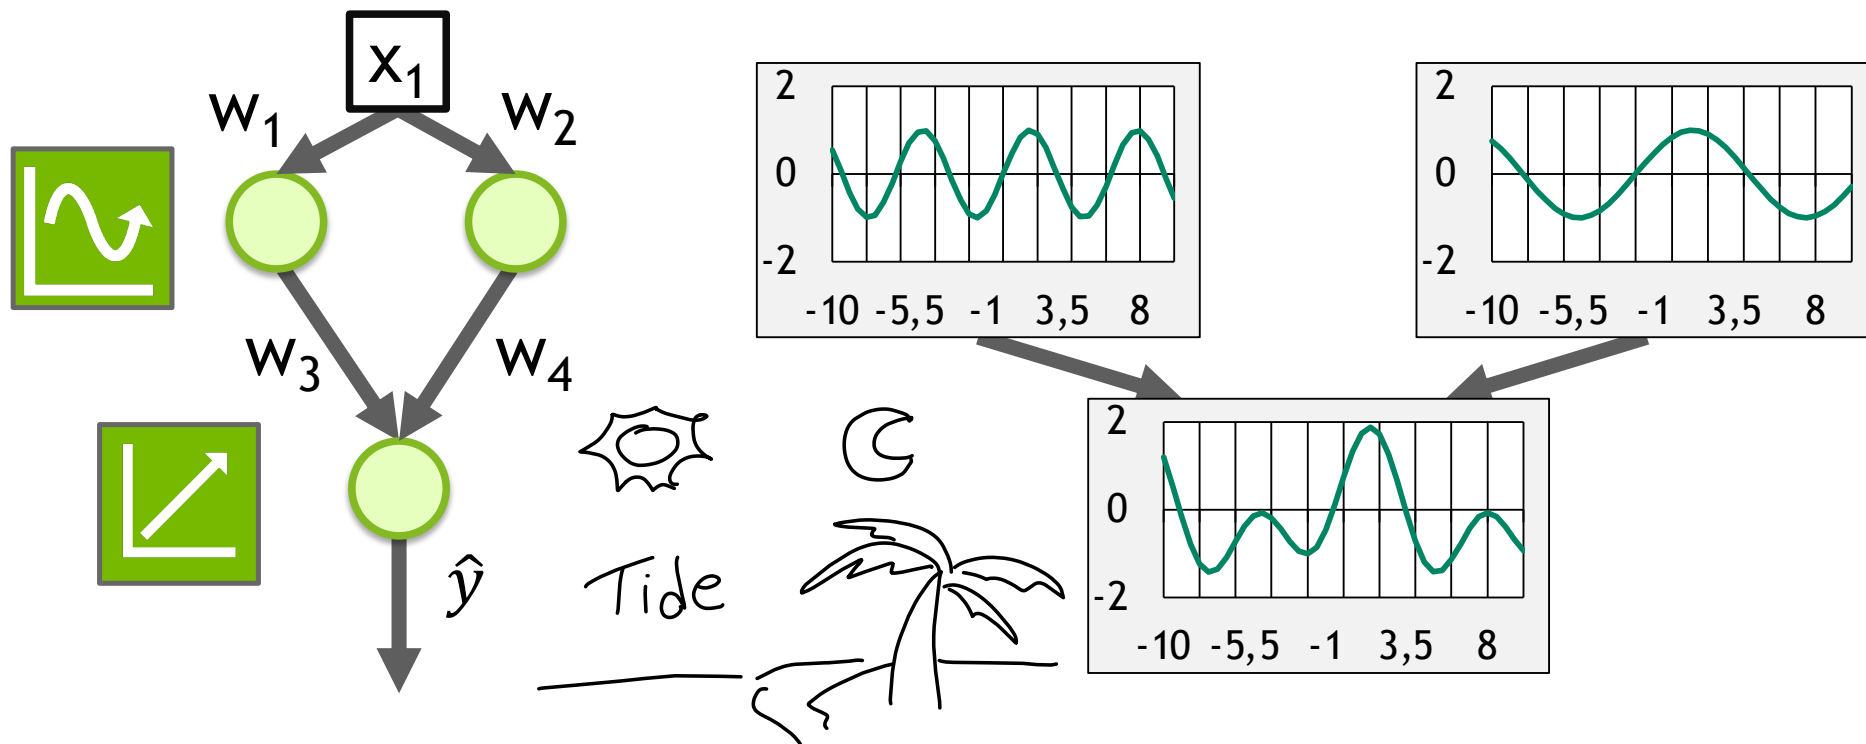
ReLU



Sigmoid



ACTIVATION FUNCTIONS






OVERFITTING



OVERFITTING

Why not have a super large neural network?





OVERFITTING

Which Trendline is Better?



OVERFITTING

Which Trendline is Better?



TRAINING VS VALIDATION DATA

Avoid memorization

Training data


- Core dataset for the model to learn on

Validation data

- New data for model to see if it truly understands (can generalize)

Overfitting


- When model performs well on the training data, but not the validation data (evidence of memorization)
- Ideally the accuracy and loss should be similar between both datasets






FROM REGRESSION TO
CLASSIFICATION


AN MNIST MODEL




AN MNIST MODEL




AN MNIST MODEL





RMSE FOR PROBABILITIES?




RMSE FOR PROBABILITIES?



CROSS ENTROPY




CROSS ENTROPY




$$Loss = -(t(x) \cdot \log(p(x)) + (1 - t(x)) \cdot \log(1 - p(x)))$$

$t(x)$ = target (0 if False, 1 if True)


$p(x)$ = probability prediction of point x



CROSS ENTROPY



```
1 def cross_entropy(y_hat, y_actual):  
2     """Infinite error for misplaced confidence."""  
3     loss = log(y_hat) if y_actual else log(1-y_hat)  
4     return -1*loss
```






BRINGING IT TOGETHER

THE NEXT EXERCISE

The American Sign Language Alphabet






LET'S GO!



DEEP
LEARNING
INSTITUTE

APPENDIX: GRADIENT DESCENT

HELPING THE COMPUTER CHEAT CALCULUS



$$MSE = \frac{1}{n} \sum_{i=1}^n (y - \hat{y})^2 = \frac{1}{n} \sum_{i=1}^n (y - (mx + b))^2$$


$$MSE = \frac{1}{2} ((3 - (m(1) + b))^2 + (5 - (m(2) + b))^2)$$

$$\frac{\partial MSE}{\partial m} = 9m + 5b - 23$$



$$\frac{\partial MSE}{\partial b} = 5m + 3b - 13$$

$$\frac{\partial MSE}{\partial m} = -7$$

$$\frac{\partial MSE}{\partial b} = -3$$



THE LOSS CURVE



THE LOSS CURVE

$$\frac{\partial MSE}{\partial m} = -7$$

$$\frac{\partial MSE}{\partial b} = -3$$




THE LOSS CURVE

$$\frac{\partial MSE}{\partial m} = -7 \quad \frac{\partial MSE}{\partial b} = -3$$

$$m := m - \lambda \frac{\partial MSE}{\partial m}$$

$$b := b - \lambda \frac{\partial MSE}{\partial b}$$




THE LOSS CURVE

$$\frac{\partial MSE}{\partial m} = -7 \quad \frac{\partial MSE}{\partial b} = -3$$

$$m := m - \lambda \frac{\partial MSE}{\partial m} \quad \lambda = .6$$

$$b := b - \lambda \frac{\partial MSE}{\partial b}$$




THE LOSS CURVE

$$\frac{\partial MSE}{\partial m} = -7 \quad \frac{\partial MSE}{\partial b} = -3$$

$$m := m - \lambda \frac{\partial MSE}{\partial m}$$

$$\lambda = .005$$

$$b := b - \lambda \frac{\partial MSE}{\partial b}$$




THE LOSS CURVE

$$\lambda = .1$$

$$m := -1 + 7\lambda = -0.3$$

$$b := 5 + 3\lambda = 4.7$$






DEEP
LEARNING
INSTITUTE



DEEP
LEARNING
INSTITUTE

FUNDAMENTALS OF DEEP LEARNING

Part 3: Convolutional Neural Networks



Part 1: An Introduction to Deep Learning

Part 2: How a Neural Network Trains

Part 3: Convolutional Neural Networks

Part 4: Data Augmentation and Deployment

Part 5: Pre-trained Models

Part 6: Advanced Architectures

RECAP OF THE EXERCISE

Trained a dense neural network model




Training accuracy was high



Validation accuracy was low



Evidence of overfitting

A network diagram consisting of numerous small circular nodes connected by thin, light-colored lines. The nodes are scattered across the frame, with a higher density in the upper right quadrant. Some nodes are white, while others are a bright yellow-green. The lines connecting them form a complex, web-like structure. The background is a dark, solid color.

KERNELS AND CONVOLUTION

KERNELS AND CONVOLUTION







Original Image



KERNELS AND CONVOLUTION



Original Image



KERNELS AND CONVOLUTION

Blur Kernel

.06	.13	.06
.13	.25	.13
.06	.13	.06


*

Original Image

1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0
0	1	1	1	1	0
1	0	1	1	0	1
1	1	0	0	1	1

=

Convolved Image



KERNELS AND CONVOLUTION

Blur Kernel

.06	.13	.06
.13	.25	.13
.06	.13	.06


*

Original Image

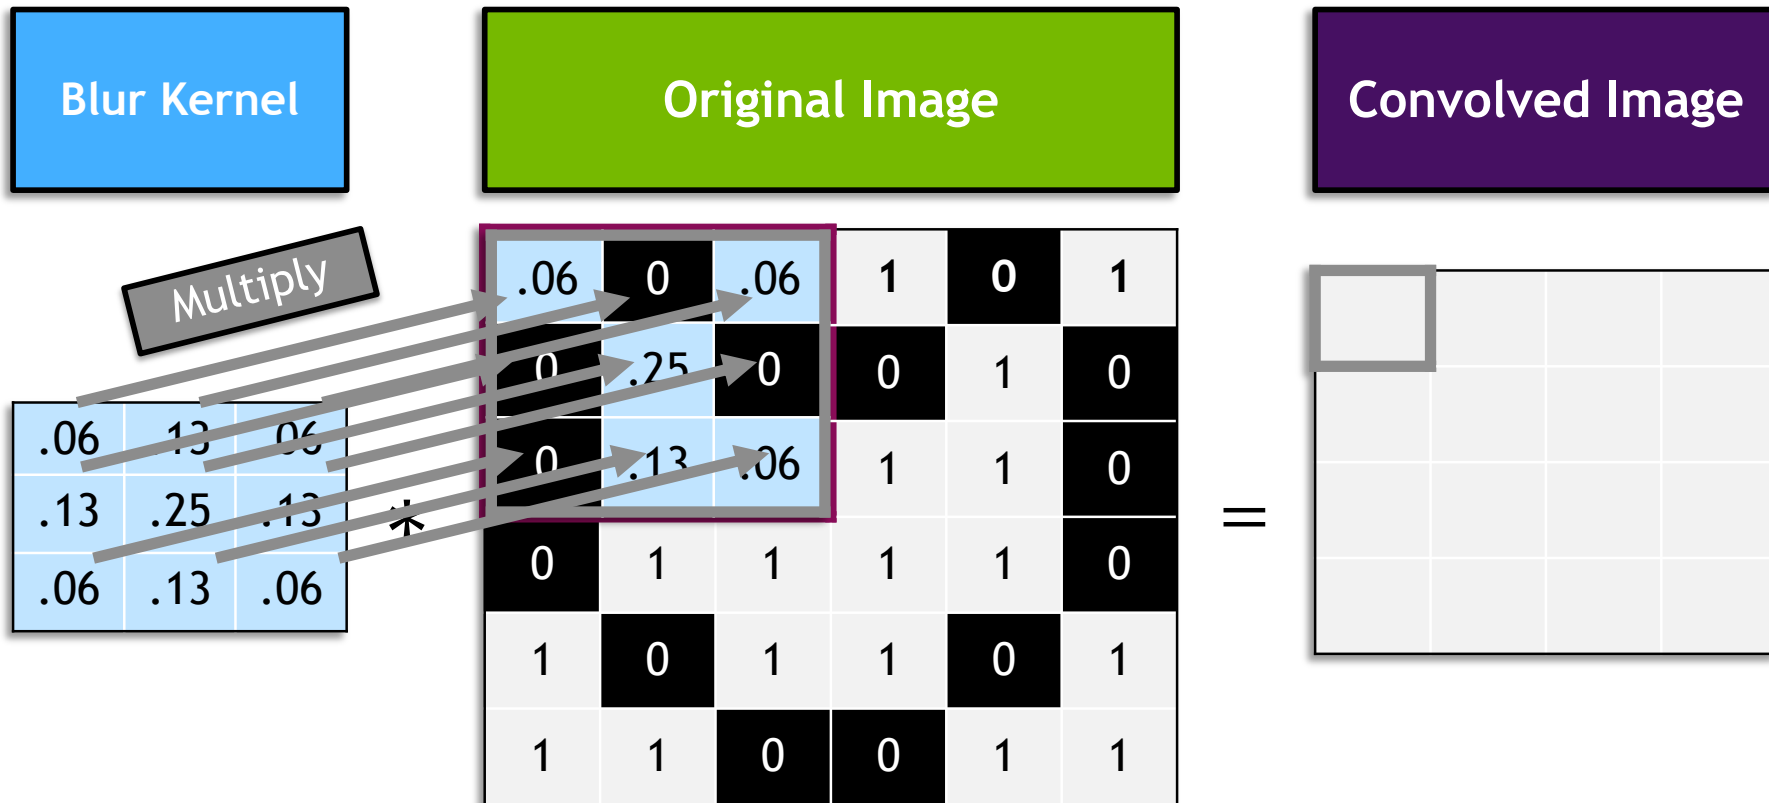
1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0
0	1	1	1	1	0
1	0	1	1	0	1
1	1	0	0	1	1

=

Convolved Image



KERNELS AND CONVOLUTION



KERNELS AND CONVOLUTION

Blur Kernel

.06	.13	.06
.13	.25	.13
.06	.13	.06

*

Original Image

.06	0	.06	1	0	1
0	.25	0	0	1	0
0	.13	.06	1	1	0
0	1	1	1	0	0
1	0	1	1	0	1
1	1	0	0	1	1

Total

=

Convolved Image

.56			

KERNELS AND CONVOLUTION

Blur Kernel

.06	.13	.06
.13	.25	.13
.06	.13	.06

*

Original Image

1	0	.13	.06	0	1
0	.13	0	0	1	0
0	.06	.13	.06	1	0
0	1	1	1	1	0
1	0	1	1	0	1
1	1	0	0	1	1

=

Convolved Image

.56	.57		

KERNELS AND CONVOLUTION

Blur Kernel

.06	.13	.06
.13	.25	.13
.06	.13	.06

*

Original Image

1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0
0	1	1	1	1	0
1	0	1	1	0	1
1	1	0	0	1	1

=


Convolved Image

.56	.57	.57	.56
.7	.82	.82	.7
.69	.95	.95	.69
.64	.69	.69	.64

STRIDE

Stride 1


1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0



.56	.57	.57	.56
-----	-----	-----	-----

Stride 2


1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0



.56	.57
-----	-----

Stride 3

1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0



.56	.56
-----	-----

PADDING

Original Image

1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0
0	1	1	1	1	0
1	0	1	1	0	1
1	1	0	0	1	1

Zero Padding

0	0	0	0	0	0	0	0
0	1	0	1	1	0	1	0
0	0	1	0	0	1	0	0
0	0	1	1	1	1	0	0
0	0	1	1	1	1	0	0
0	1	0	1	1	0	1	0
0	1	1	0	0	1	1	0
0	0	0	0	0	0	0	0

PADDING

Original Image

1	0	1	1	0	1
0	1	0	0	1	0
0	1	1	1	1	0
0	1	1	1	1	0
1	0	1	1	0	1
1	1	0	0	1	1

Mirror Padding

1	1	0	1	1	0	1	1
1	1	0	1	1	0	1	1
0	0	1	0	0	1	0	0
0	0	1	1	1	1	0	0
0	0	1	1	1	1	0	0
1	1	0	1	1	0	1	1
1	1	1	0	0	1	1	1
1	1	1	0	0	1	1	1



KERNELS AND NEURAL NETWORKS

KERNELS AND NEURAL NETWORKS

Kernel


W_1	W_2	W_3
W_4	W_5	W_6
W_7	W_8	W_9

KERNELS AND NEURAL NETWORKS

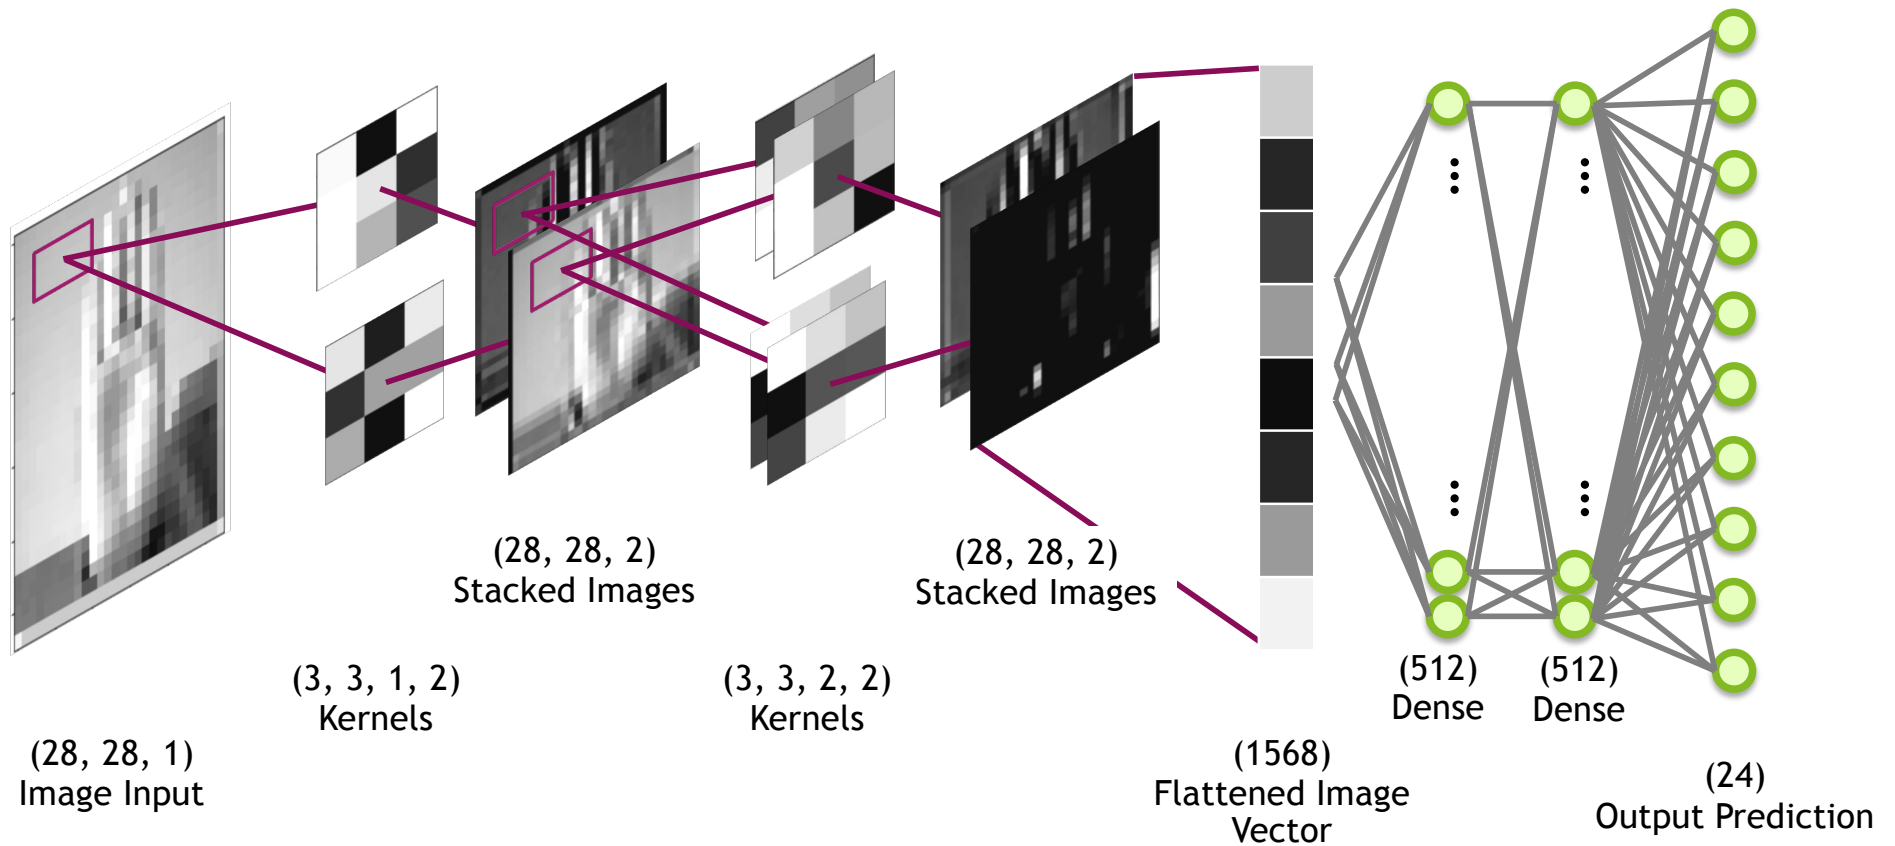
Kernel

w_1	w_2	w_3
w_4	w_5	w_6
w_7	w_8	w_9

Neuron




KERNELS AND NEURAL NETWORKS




FINDING EDGES

Vertical Edges




1	0	-1
2	0	-2
1	0	-1

Original Image



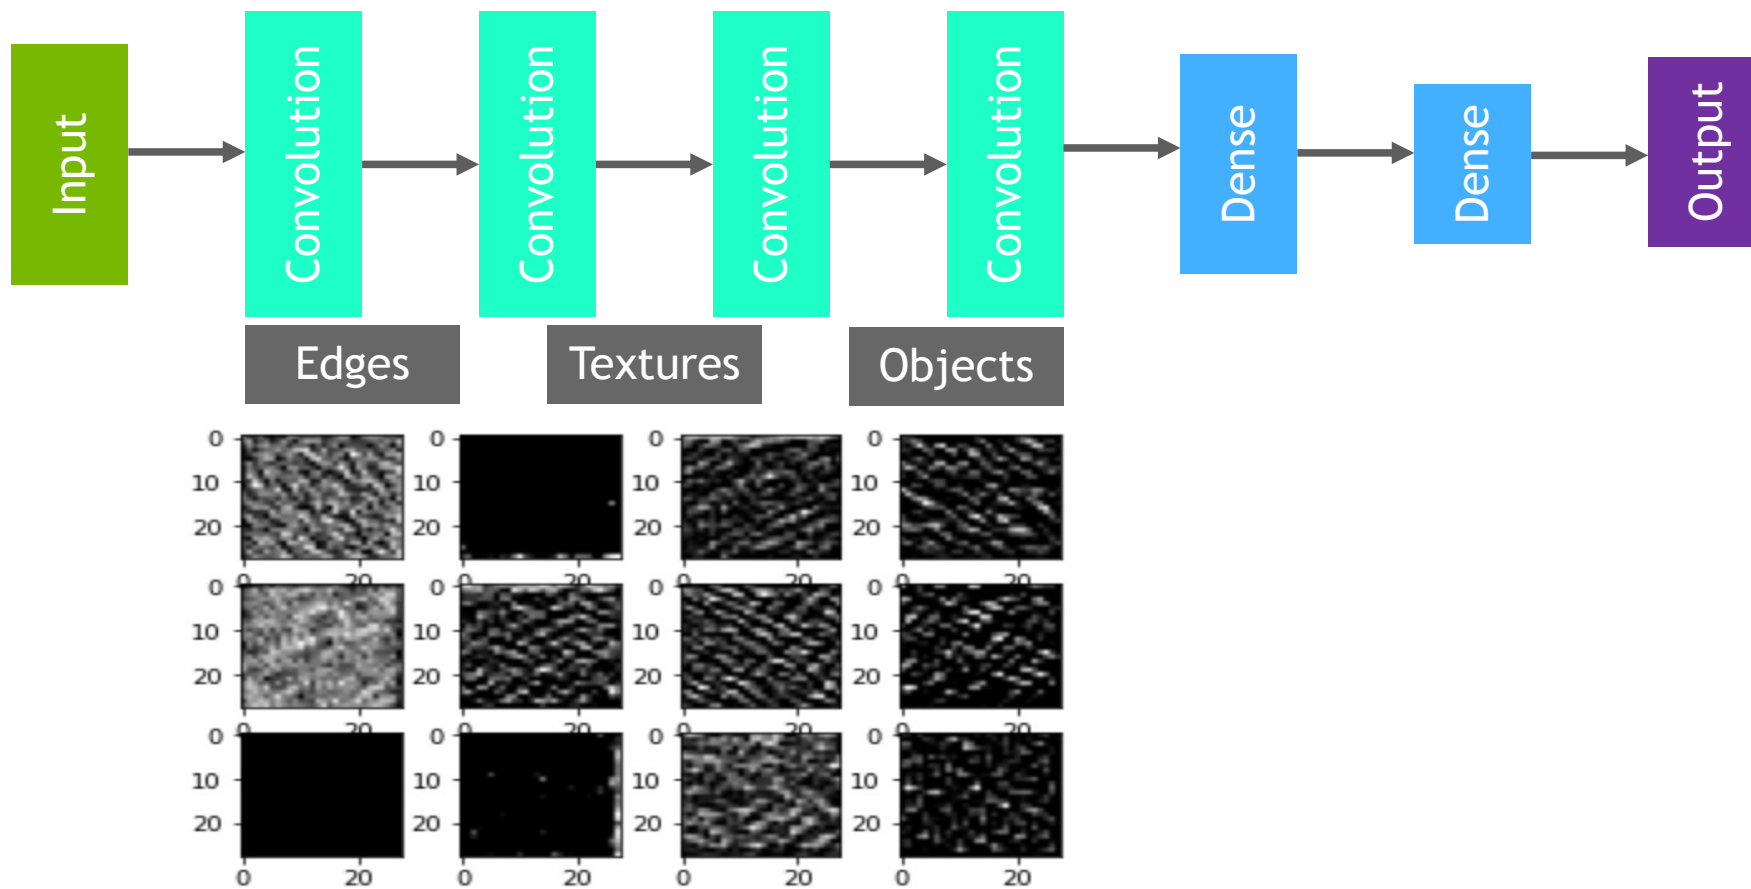
0	0	0
0	1	0
0	0	0

Horizontal Edges



1	2	1
0	0	0
-1	-2	-1

NEURAL NETWORK PERCEPTION



NEURAL NETWORK PERCEPTION






OTHER LAYERS IN THE MODEL


MAX POOLING

110	256	153	67
12	89	88	43
10	15	50	55
23	9	49	23




256	153
23	55


DROPOUT



rate = 0

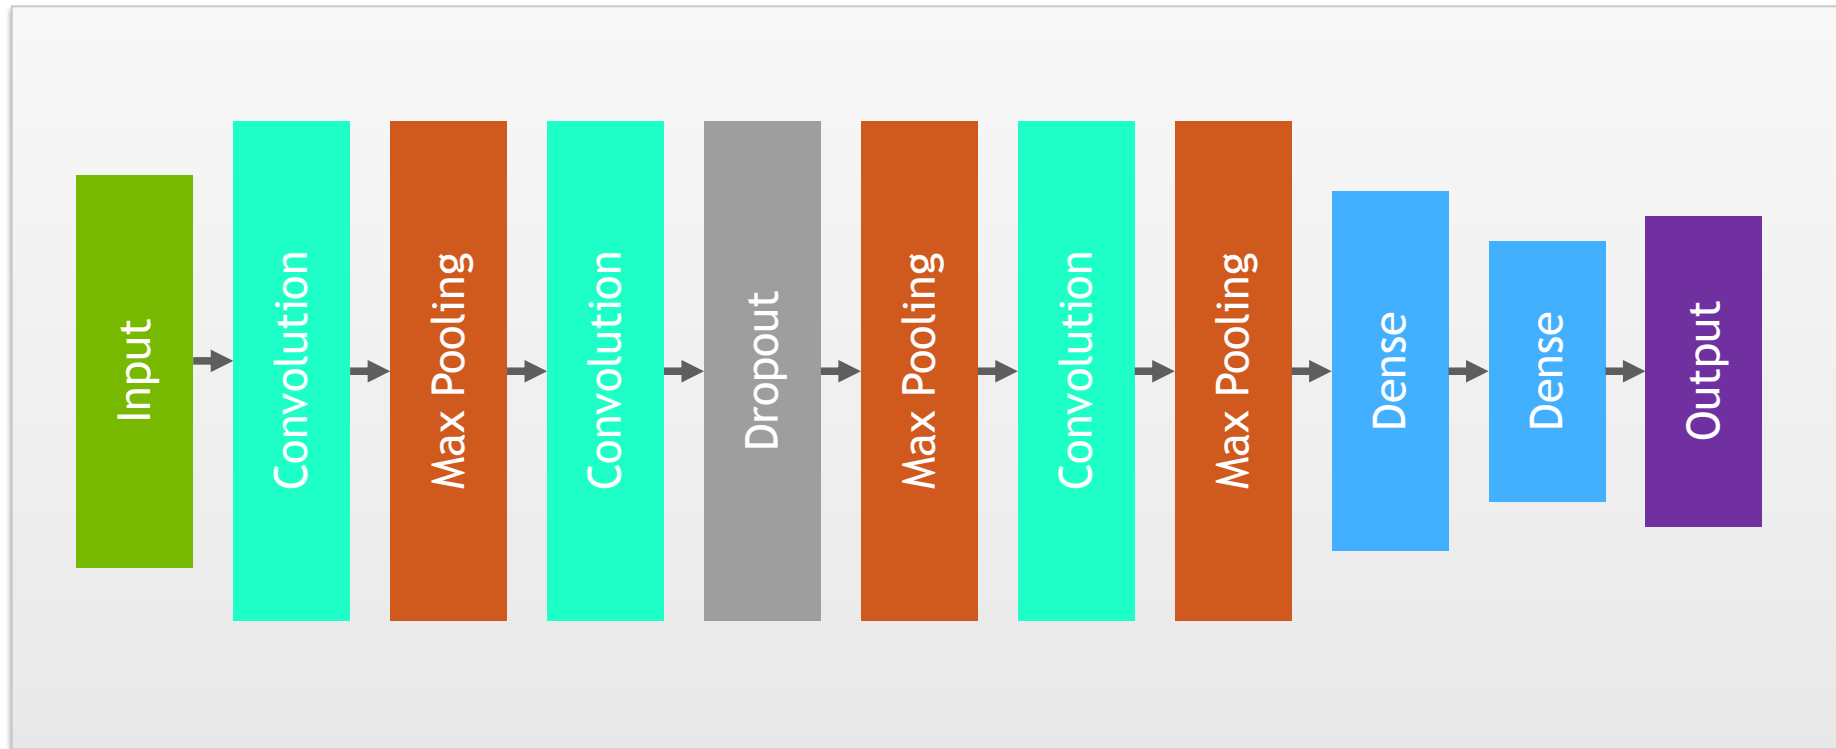


rate = .2




rate = .4

WHOLE ARCHITECTURE





LET'S GO!




DEEP
LEARNING
INSTITUTE



DEEP
LEARNING
INSTITUTE

FUNDAMENTALS OF DEEP LEARNING

Part 4: Data Augmentation and Deployment



Part 1: An Introduction to Deep Learning

Part 2: How a Neural Network Trains

Part 3: Convolutional Neural Networks

Part 4: Data Augmentation and Deployment

Part 5: Pre-trained Models

Part 6: Advanced Architectures

RECAP OF THE EXERCISE

Analysis

- CNN increased validation accuracy
- Still seeing training accuracy higher than validation

Solution

- Clean data provides better examples
- Dataset variety helps the model generalize





DATA AUGMENTATION

DATA AUGMENTATION



IMAGE FLIPPING


Horizontal Flip






Vertical Flip



ROTATION




ZOOMING





WIDTH AND HEIGHT SHIFTING



HOMOGRAPHY



BRIGHTNESS



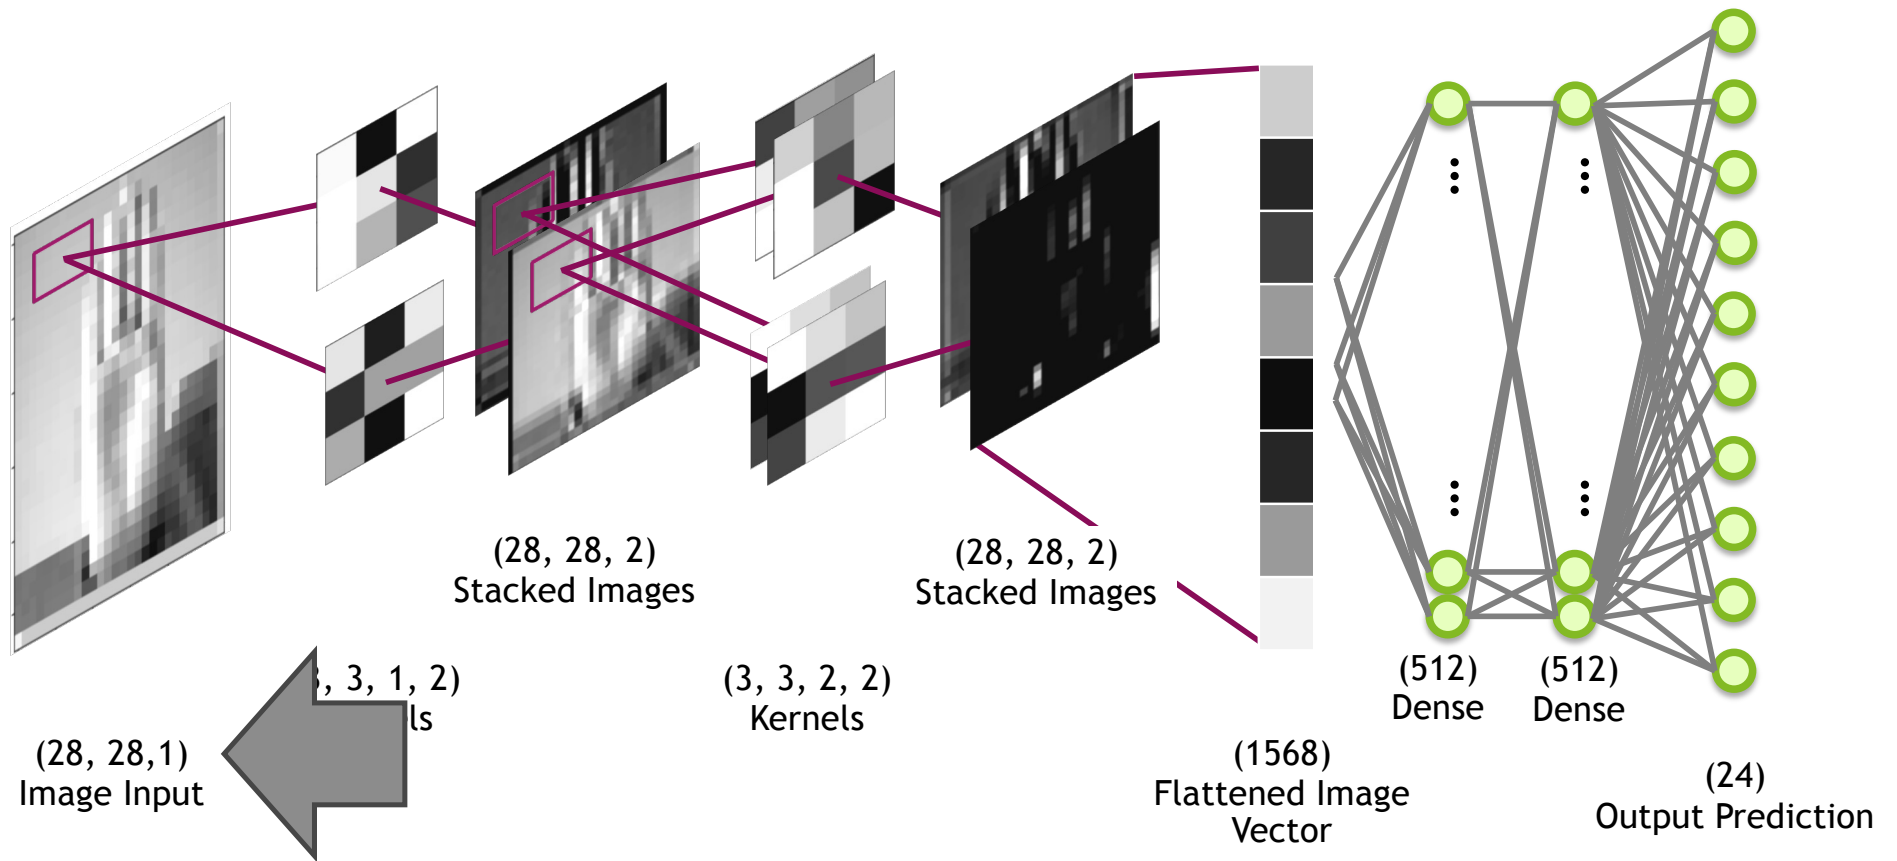
CHANNEL SHIFTING



A network diagram consisting of numerous small circular nodes connected by thin, light-colored lines. The nodes are scattered across the frame, with a higher density in the upper right quadrant. Some nodes are white, while others are a bright yellow-green. The lines connecting them form a complex web of connections. The background is a dark, solid color.

MODEL DEPLOYMENT

MODEL DEPLOYMENT



MODEL DEPLOYMENT

Training
Batch Input



Convolution

Max Pooling

...

MODEL DEPLOYMENT



$(287, 433, 3)$

$(220, 155, 3)$

$(220, 155, 1)$

$(1, 220, 155, 1)$


Resize

Greyscale

“Batch”



LET'S TRY IT OUT!




DEEP
LEARNING
INSTITUTE



DEEP
LEARNING
INSTITUTE

FUNDAMENTALS OF DEEP LEARNING

Part 5: Pre-trained Models



Part 1: An Introduction to Deep Learning

Part 2: How a Neural Network Trains

Part 3: Convolutional Neural Networks

Part 4: Data Augmentation and Deployment

Part 5: Pre-trained Models

Part 6: Advanced Architectures



REVIEW SO FAR

REVIEW SO FAR



- Learning Rate
- Number of Layers
- Neurons per Layer
- Activation Functions
- Dropout
- Data



PRE-TRAINED MODELS

PRE-TRAINED MODELS

TensorFlow Hub

 Keras



PYTORCH
HUB

PRE-TRAINED MODELS

VERY DEEP CONVOLUTIONAL NETWORKS FOR LARGE-SCALE IMAGE RECOGNITION

Karen Simonyan* & Andrew Zisserman⁺

Visual Geometry Group, Department of Engineering Science, University of Oxford
{karen, az}@robots.ox.ac.uk

IM  GENET

THE NEXT CHALLENGE

An Automated Doggy Door





TRANSFER LEARNING

THE CHALLENGE AFTER

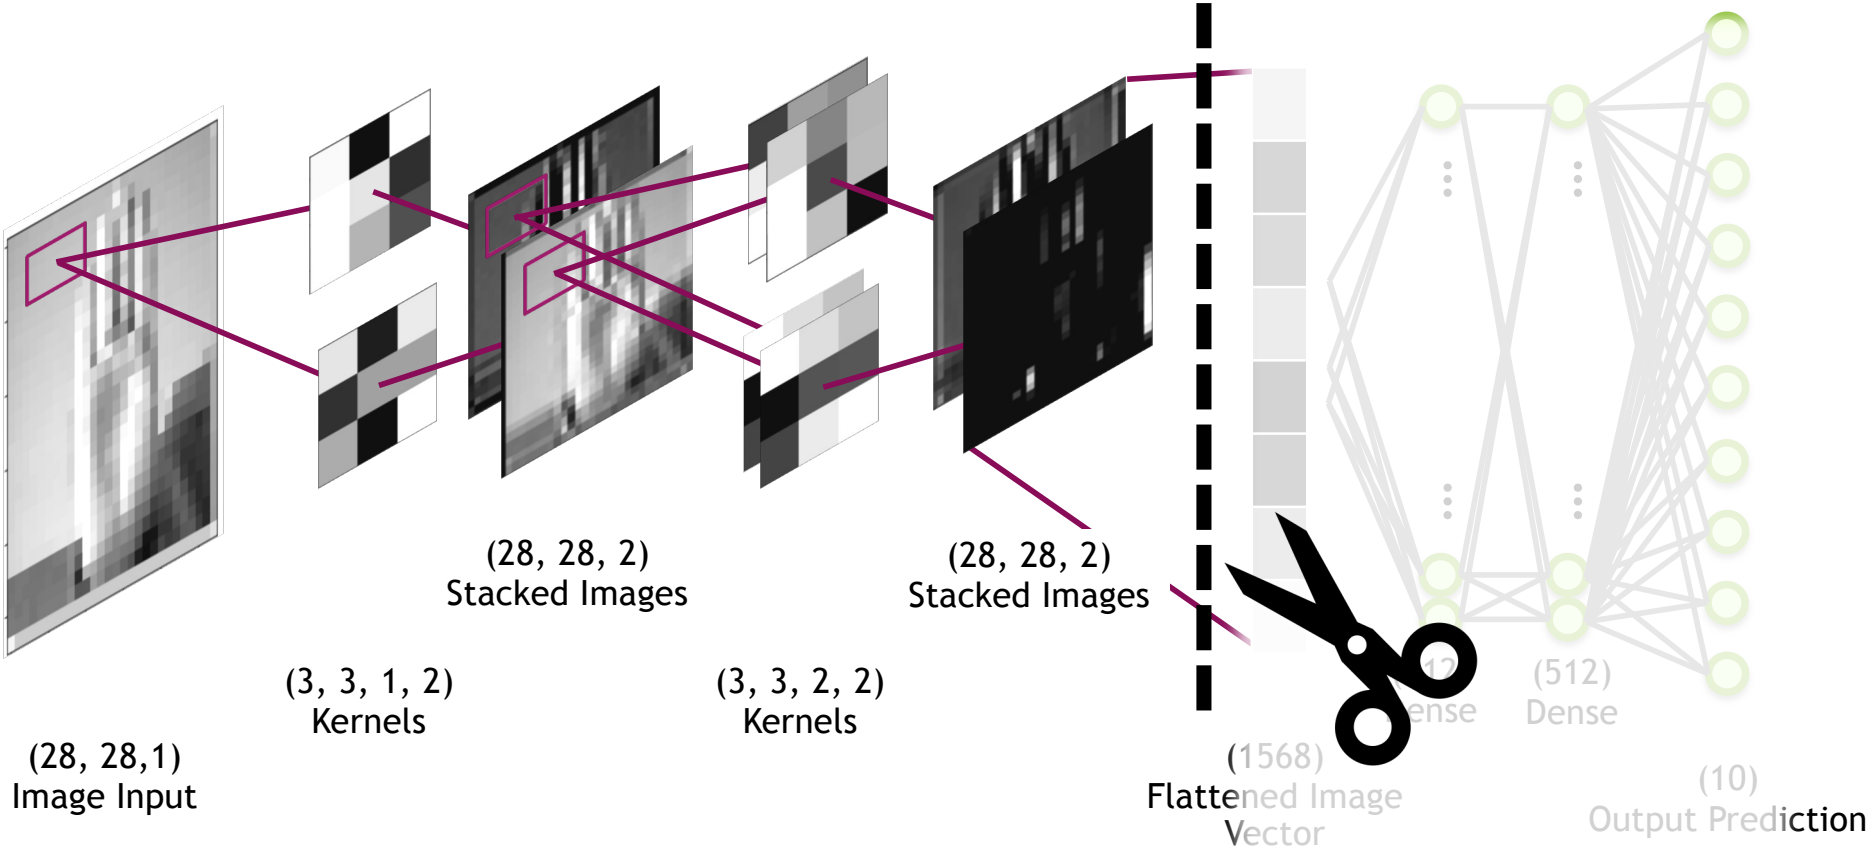
An Automated Presidential Doggy Door



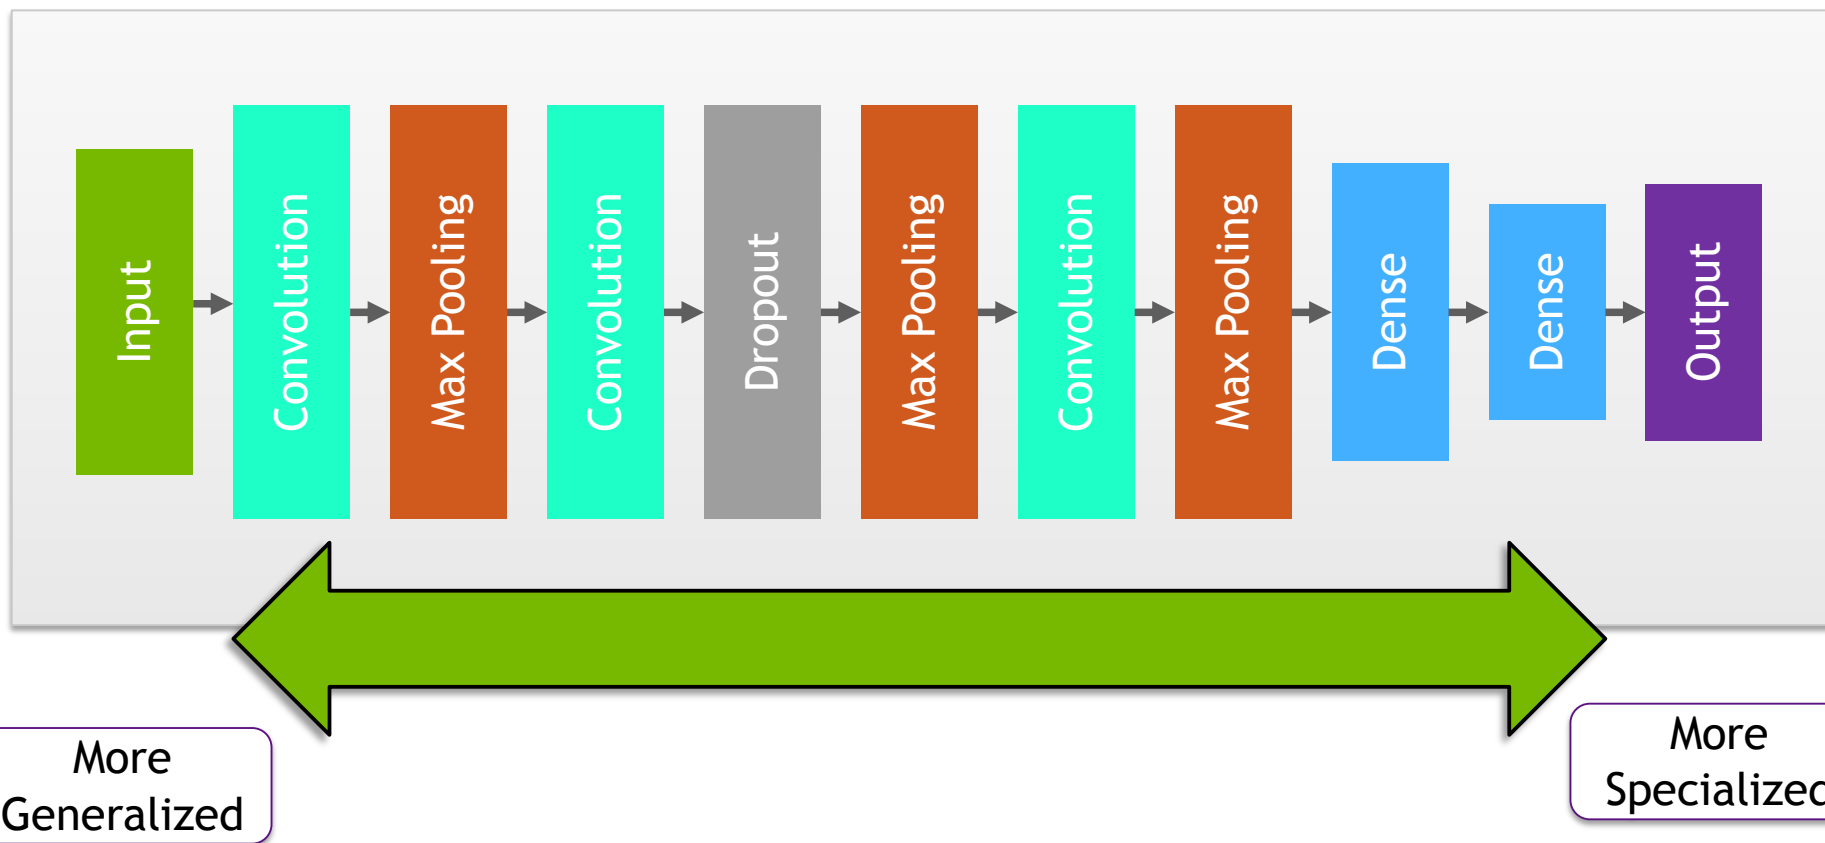
TRANSFER LEARNING



TRANSFER LEARNING



TRANSFER LEARNING





TRANSFER LEARNING

Freezing the Model?




TRANSFER LEARNING





LET'S GET STARTED!




DEEP
LEARNING
INSTITUTE



DEEP
LEARNING
INSTITUTE

FUNDAMENTALS OF DEEP LEARNING

Part 6: Advanced Architectures



Part 1: An Introduction to Deep Learning

Part 2: How a Neural Network Trains

Part 3: Convolutional Neural Networks

Part 4: Data Augmentation and Deployment

Part 5: Pre-trained Models

Part 6: Advanced Architectures



MOVING FORWARD

FIELDS OF AI



Computer Vision

- Optometry



Natural Language Processing

- Linguistics



Reinforcement Learning


- Game Theory
- Psychology



Anomaly Detection

- Security
- Medicine

FIELDS OF AI




Computer Vision

- Optometry



Natural Language Processing

- Linguistics



Reinforcement Learning

- Game Theory
- Psychology



Anomaly Detection


- Security
- Medicine

FIELDS OF AI



Computer Vision

- Optometry



Natural Language Processing

- Linguistics



Reinforcement Learning

- Game Theory
- Psychology



Anomaly Detection

- Security
- Medicine



NATURAL LANGUAGE PROCESSING


FROM WORDS TO NUMBERS

“A dog barked at a cat.”

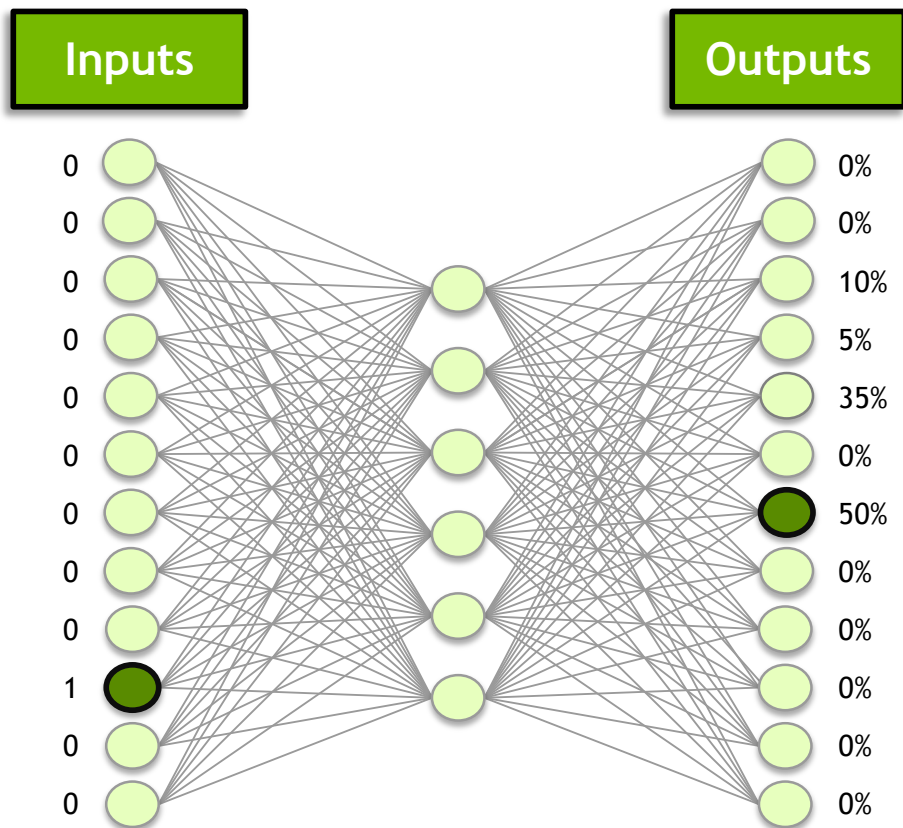
[1, 10, 7, 4, 1, 8]

Dictionary	
1. A	8. Cat
2. An	9. Cats
3. And	10. Dog
4. At	11. Dogs
5. Ate	12. Eat
6. Bark	
7. Barked	

FROM WORDS TO NUMBERS



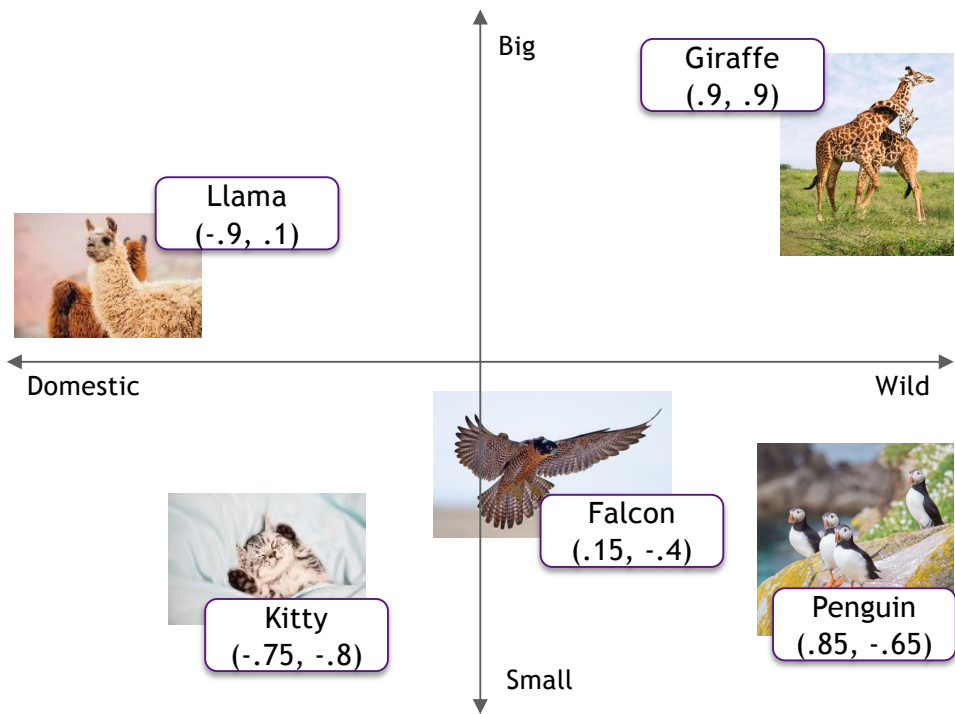
FROM WORDS TO NUMBERS



Dictionary

1.	A	8.	Cat
2.	An	9.	Cats
3.	And	10.	Dog
4.	At	11.	Dogs
5.	Ate	12.	Eat
6.	Bark		
7.	Barked		

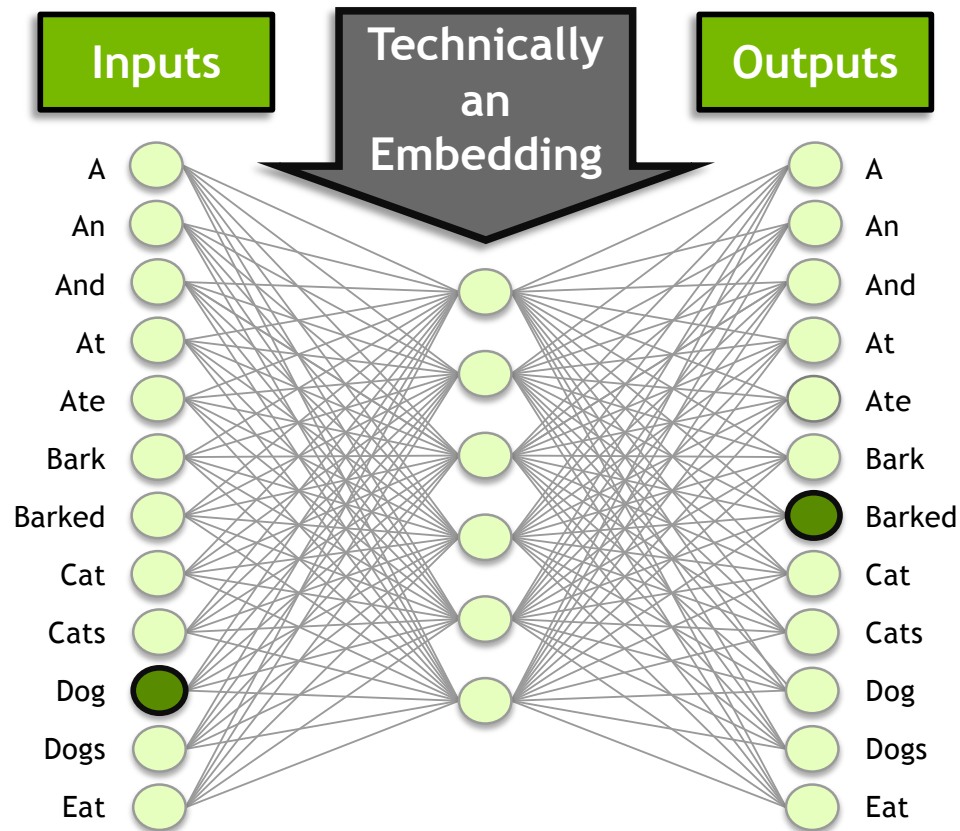
FROM WORDS TO NUMBERS



BIGGER DICTIONARY

1.	A	31.	ATE	61.	CATS
2.	AN	32.	BARK	62.	DOG
3.	AND	33.	BARKED	63.	DOGS
4.	AT	34.	CAT	64.	EAT
5.	ATE	35.	CATS	65.	EATEN
6.	BARK	36.	DOG	66.	A
7.	BARKED	37.	DOGS	67.	AN
8.	CAT	38.	EAT	68.	AND
9.	CATS	39.	EATEN	69.	AT
10.	DOG	40.	A	70.	ATE
11.	DOGS	41.	AN	71.	BARK
12.	EAT	42.	AND	72.	BARKED
13.	EATEN	43.	AT	73.	CAT
14.	A	44.	ATE	74.	CATS
15.	AN	45.	BARK	75.	DOG
16.	AND	46.	BARKED	76.	DOGS
17.	AT	47.	CAT	77.	EAT
18.	ATE	48.	CATS	78.	EATEN
19.	BARK	49.	DOG	79.	...
20.	BARKED	50.	DOGS	80.	...
21.	CAT	51.	EAT	81.	...
22.	CATS	52.	EATEN	82.	...
23.	DOG	53.	A		
24.	DOGS	54.	AN		
25.	EAT	55.	AND		
26.	EATEN	56.	AT		
27.	A	57.	ATE		
28.	AN	58.	BARK		
29.	AND	59.	BARKED		
30.	AT	60.	CAT		

FROM WORDS TO NUMBERS



DICTIONARY

1. A	8. CAT
2. AN	9. CATS
3. AND	10. DOG
4. AT	11. DOGS
5. ATE	12. EAT
6. BARK	
7. BARKED	

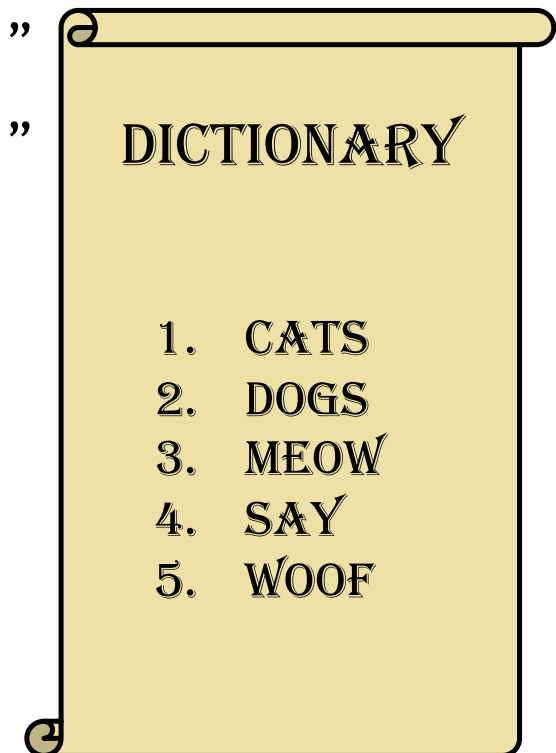
A network diagram consisting of numerous small circular nodes connected by thin, light-colored lines. The nodes are scattered across the frame, with a higher density in the upper right quadrant. Some nodes are white, while others are a bright yellow-green. The lines connecting them form a complex, web-like structure. The background is a dark, solid color.

RECURRENT NEURAL NETWORKS

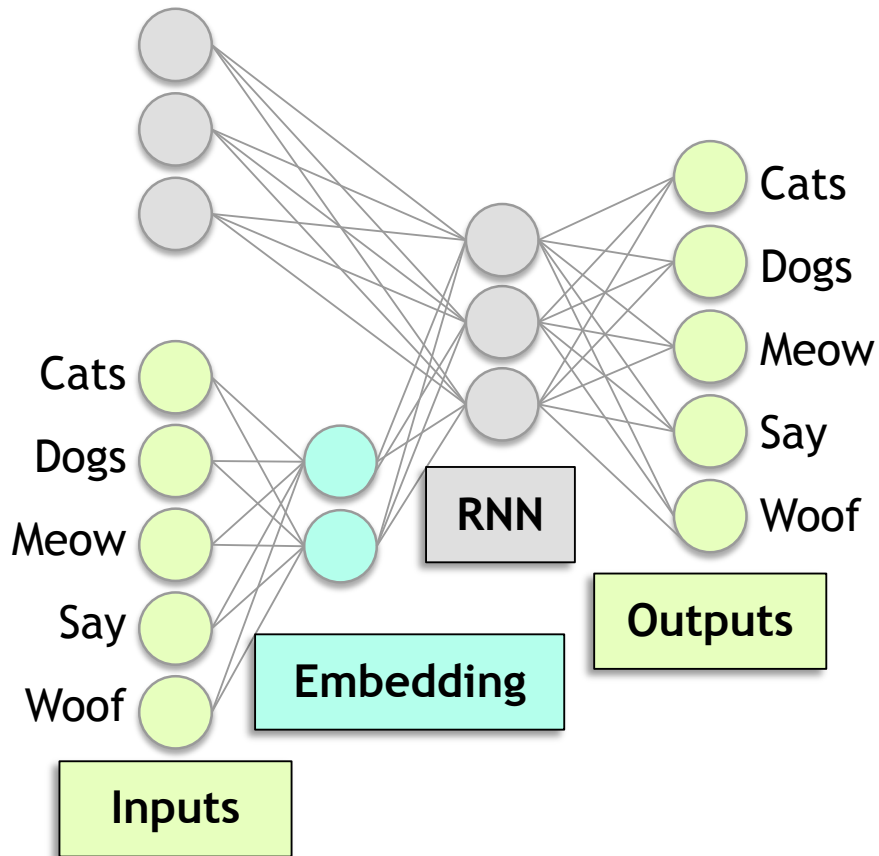
RECURRENT NEURAL NETWORKS

“Cats say ____.”

“Dogs say ____.”



RECURRENT NEURAL NETWORKS




“Cats say ____.”

“Dogs say ____.”

DICTIONARY

1. CATS
2. DOGS
3. MEOW
4. SAY
5. WOOF

RECURRENT NEURAL NETWORKS




“Cats say ____.”

“Dogs say ____.”

DICTIONARY

1. CATS
2. DOGS
3. MEOW
4. SAY
5. WOOF

RECURRENT NEURAL NETWORKS




“Cats say ____.”

“Dogs say ____.”

DICTIONARY

1. CATS
2. DOGS
3. MEOW
4. SAY
5. WOOF

RECURRENT NEURAL NETWORKS




“Cats say ____.”

“Dogs say ____.”

DICTIONARY

1. CATS
2. DOGS
3. MEOW
4. SAY
5. WOOF

RECURRENT NEURAL NETWORKS




“Cats say ____.”

“Dogs say ____.”

DICTIONARY

1. CATS
2. DOGS
3. MEOW
4. SAY
5. WOOF

RECURRENT NEURAL NETWORKS



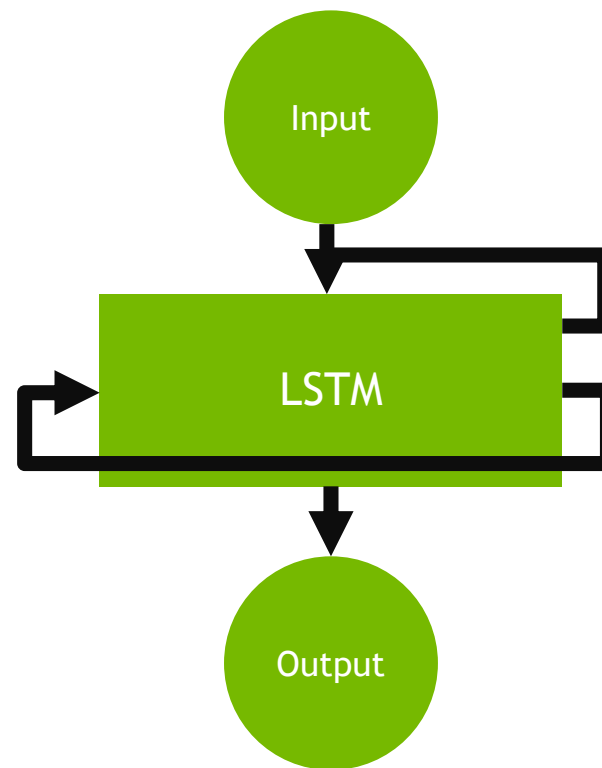

“Cats say ____.”

“Dogs say ____.”

DICTIONARY

1. CATS
2. DOGS
3. MEOW
4. SAY
5. WOOF


RECURRENT NEURAL NETWORKS






OTHER ARCHITECTURES


AUTOENCODERS



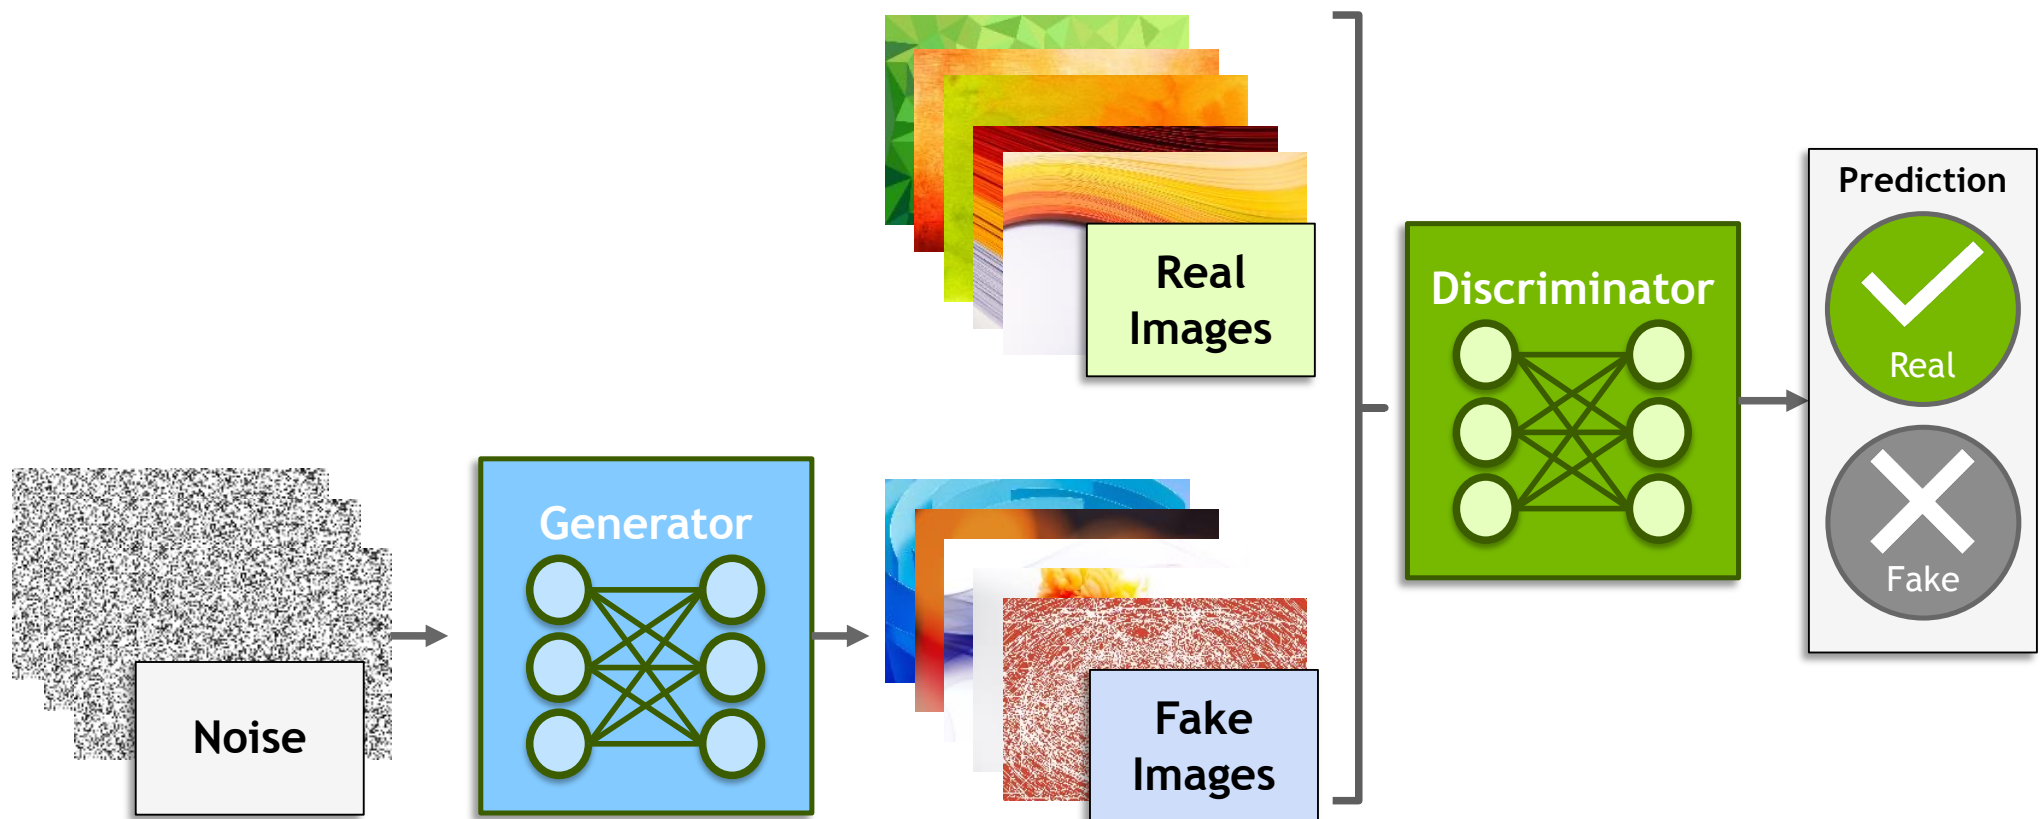
AUTOENCODERS



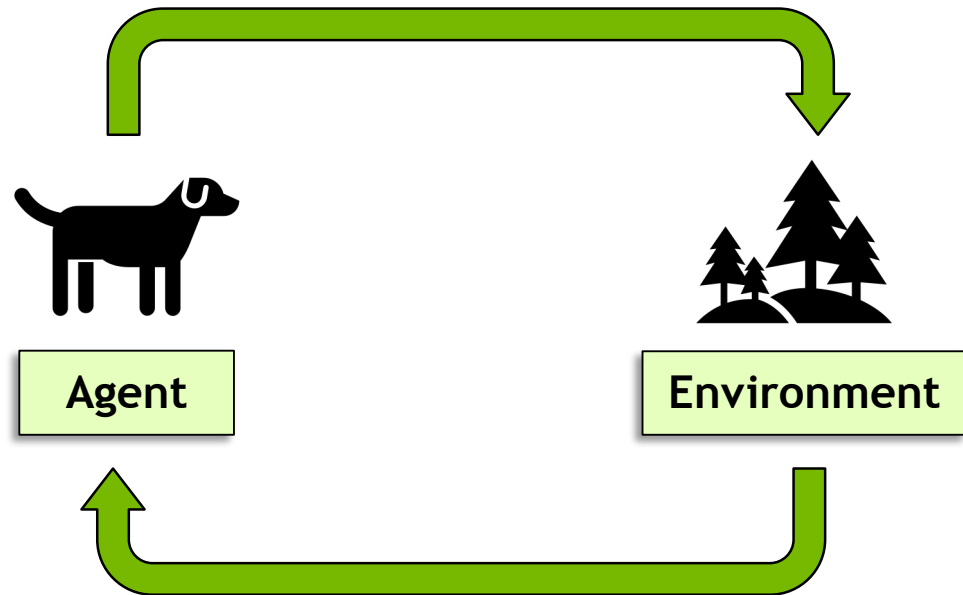
AUTOENCODERS



GENERATIVE ADVERSARIAL NETWORKS (GANS)



REINFORCEMENT LEARNING





NEXT STEPS

ENABLING PORTABILITY WITH NGC CONTAINERS

Extensive

- Diverse range of workloads and industry specific use cases

Optimized

- DL containers updated monthly
- Packed with latest features and superior performance

Secure & Reliable

- Scanned for vulnerabilities and crypto
- Tested on workstations, servers, & cloud instances

Scalable

- Supports multi-GPU & multi-node systems

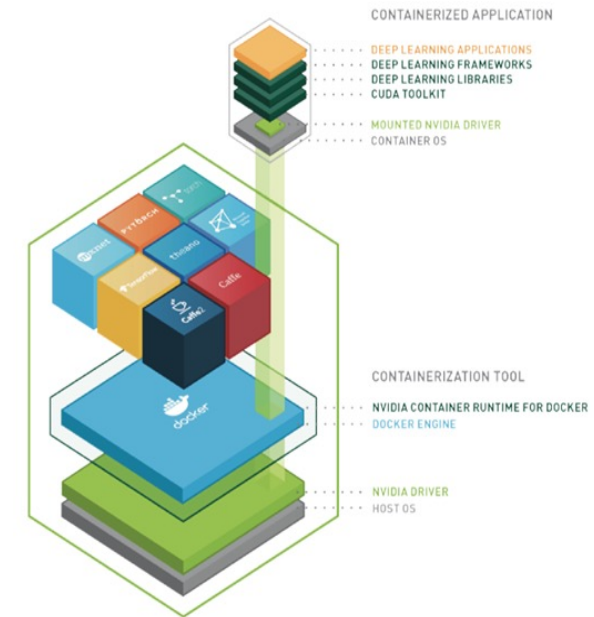
Designed for Enterprise & HPC

- Supports Docker, Singularity & other runtimes

Run Anywhere

- Bare metal, VMs, Kubernetes
- x86, ARM, POWER
- Multi-cloud, on-prem, hybrid, edge

NGC Deep Learning Containers



[Learn more about NGC Containers](#)

NEXT STEPS FOR THIS CLASS

Catalog: Containers / Containers: nvidia:dli-dl-fundamentals

DLI Deep Learning Fundamentals Course -...

Publisher	Built By	Latest Tag	Modified	Size
NVIDIA	NVIDIA	v0.0.1	October 27, 2020	4.19 GB

Multinode Support
No

Multi-Arch Support
No

Description
Base environment used in the NVIDIA Deep Learning Institute (DLI) Course Fundamentals of Deep Learning, along with Next Steps project.

Labels

Computer Vision DLI Jupyter Machine Learning Machine Learning & AI

Pull Command

```
docker pull nvcr.io/nvidia/dli-dl-fundamentals:v0.0.1
```

Step 1 Setup Docker

<https://www.docker.com/>

Step 2 Visit NGC Catalog

<https://ngc.nvidia.com/catalog/containers/nvidia:dli-dl-fundamentals>

Step 3 Pull and Run Container

Visit localhost:8888 to check out a JupyterLab environment with a Next Steps Project




CLOSING THOUGHTS

COPYING ROCKET SCIENCE





LET'S GET STARTED!



DEEP
LEARNING
INSTITUTE