APS360 Artificial Intelligence Fundamentals

Lisa Zhang

Lecture 1; May 6, 2019

Agenda

First Hour:

- Introduction
- Motivation
- Logistics

Second Hour:

- Biological and artificial neurons
- Building an artificial Pigeon

Welcome to APS360!

Introduction

- **Instructor**: Lisa Zhang
- **Email**: lczhang@cs.toronto.edu
  - Please prefix email subject with 'APS360'
- **Office hours**: Thursday 3pm-4pm and by appointment

About your instructor

Before I started teaching, I was...

- a masters student doing research in Machine Learning
- a senior data scientist at an advertising technology company
- a startup founder of a data visualization company
- a software developer intern in various Silicon Valley companies, e.g. Facebook, ContextLogic (Wish)

I studied...

- machine learning at UofT (supervised by Prof. Richard Zemel, Prof. Raquel Urtasun)
- pure math at UWaterloo

Ask me about anything outside of class, or empty office hours!

About you

Introduce yourself to the people sitting in your pod:

- What is your name?
- What is your area of study?
- Why are you here?

Survey: demographics

- **Year of study**:
  - 40% - 2rd years
  - 40% - 3rd years
  - 10% - 4th years
  - 10% - other
- **Area of study**:
  - Engineering Science
  - Electiacal & Computer Engineering
  - Mechanical engineering
  - Others

Why did you take this course?

- AI Minor (50%)
- AI Certificate (30%)
- Practical Skills (65%)
- For Fun (50%)

Previous ML and Neural Networks courses?

- Almost everyone have not

Summer Activities

- Taking another course ~60%
- Working full time ~70%

Interest in Machine Learning

- Application to another field: ~70%
- Becoming a data scientist: ~40%
- Machine Learning Research: ~30%

Programming Proficiency

- Python Experience: ~70%
- Jupyter Notebooks: ~30%
- PyTorch: ~2%

I will assume that you have Python experience. (e.g. loops, conditionals, functions, etc)

What do you know about AI?

What is the difference between:

- Artificial Intelligence,
- Machine Learning, and
- Deep Learning?

AI vs ML vs DL

Artificial Intelligence: Create intelligent machines that work and act like humans.

Machine Learning: Find an algorithm that automatically learns from example data.

Deep Learning: Using deep neural networks to automatically learn from example data.

Relationship

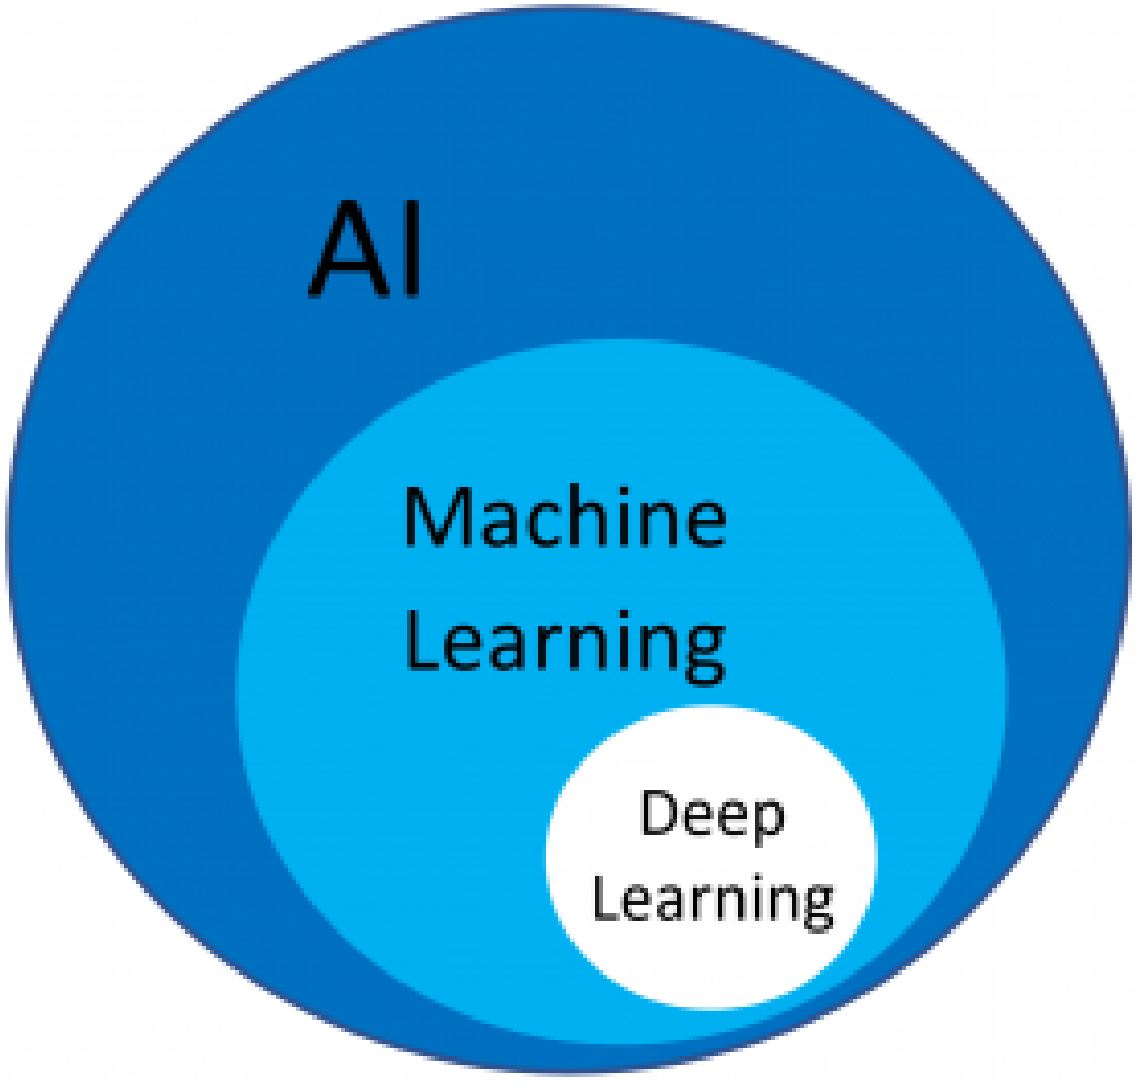

ARTIFICIAL INTELLIGENCE

Artificial Intelligence captures the imagination of the world.

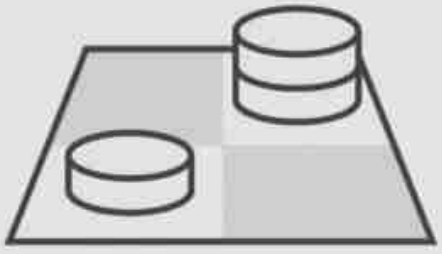

MACHINE LEARNING

Machine learning starts to gain traction.

DEEP LEARNING

Deep learning catapults the industry.

Turing Test Devised
1950

ELIZA
1964 - 1966

Edward Shortliffe writes MYCIN, an Expert or Rule based System, to classify blood disease
1970s

IBM Deep Blue defeats Grand Master Garry Kasparov in chess
1996

ImageNet Feeds Deep Learning
2009

AlphaGo defeats Go champion Lee Sedol
2016

1950s 1960s 1970s 1980s 1990s 2000s 2010s

Why machine learning?

For many problems, it is difficult to program the correct behavior.

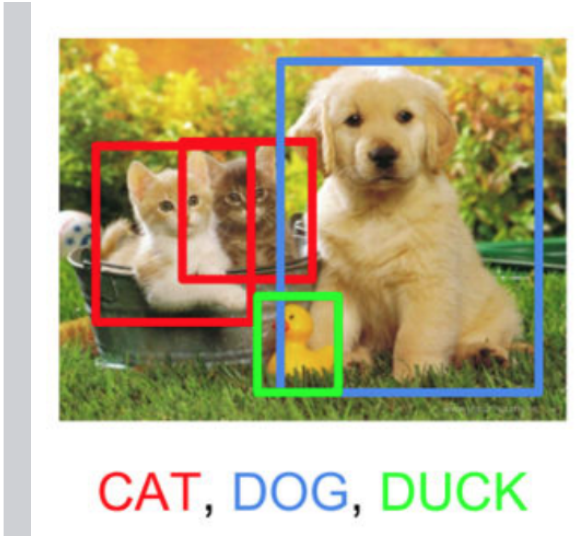

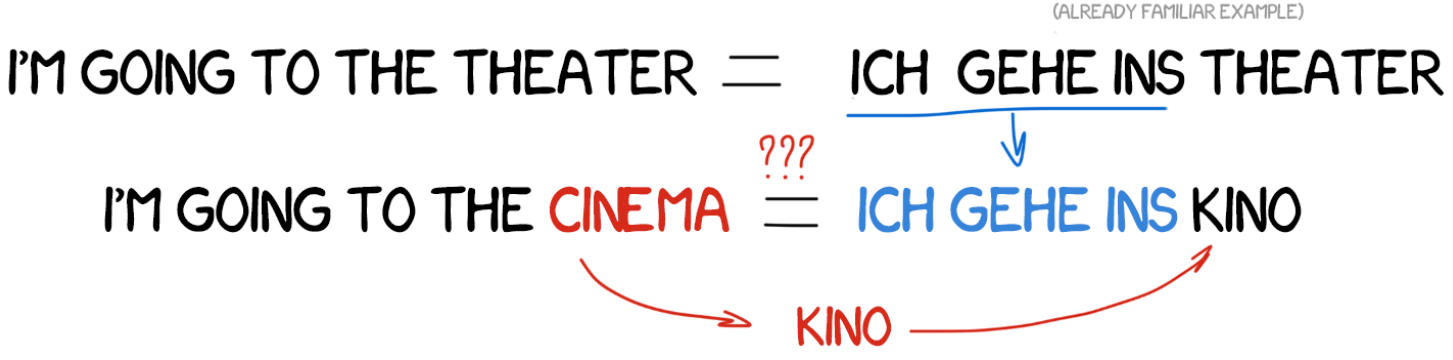

Types of Machine Learning Problems

- Supervised Learning
  - Regression
  - Classification
- Unsupervised Learning
- Reinforcement Learning
- (. . . and more)

Supervised Learning Task

Supervised Learning: learning a function that maps an input to an output based on example input-output pairs

Examples:

- Age prediction given a headshot:
  - Input: headshot image
  - Output: person's age
- Sentiment classification given a tweet:
  - Input: tweet text
  - Output: whether the tweet is happy or sad

Supervised Learning Task

- **Regression**: when the output is a continuous value
  - e.g. age prediction
- **Classification**: when the output is a categorical value
  - e.g. sentiment classification

Unsupervised Learning

Unsupervised Learning: learning the structure of some (unlabelled) data

Example:

- clustering
- generating new images
- style transfer

Reinforcement Learning

Reinforcement Learning: learning what actions to take to optimize long-term reward.

Example:

- playing a video game

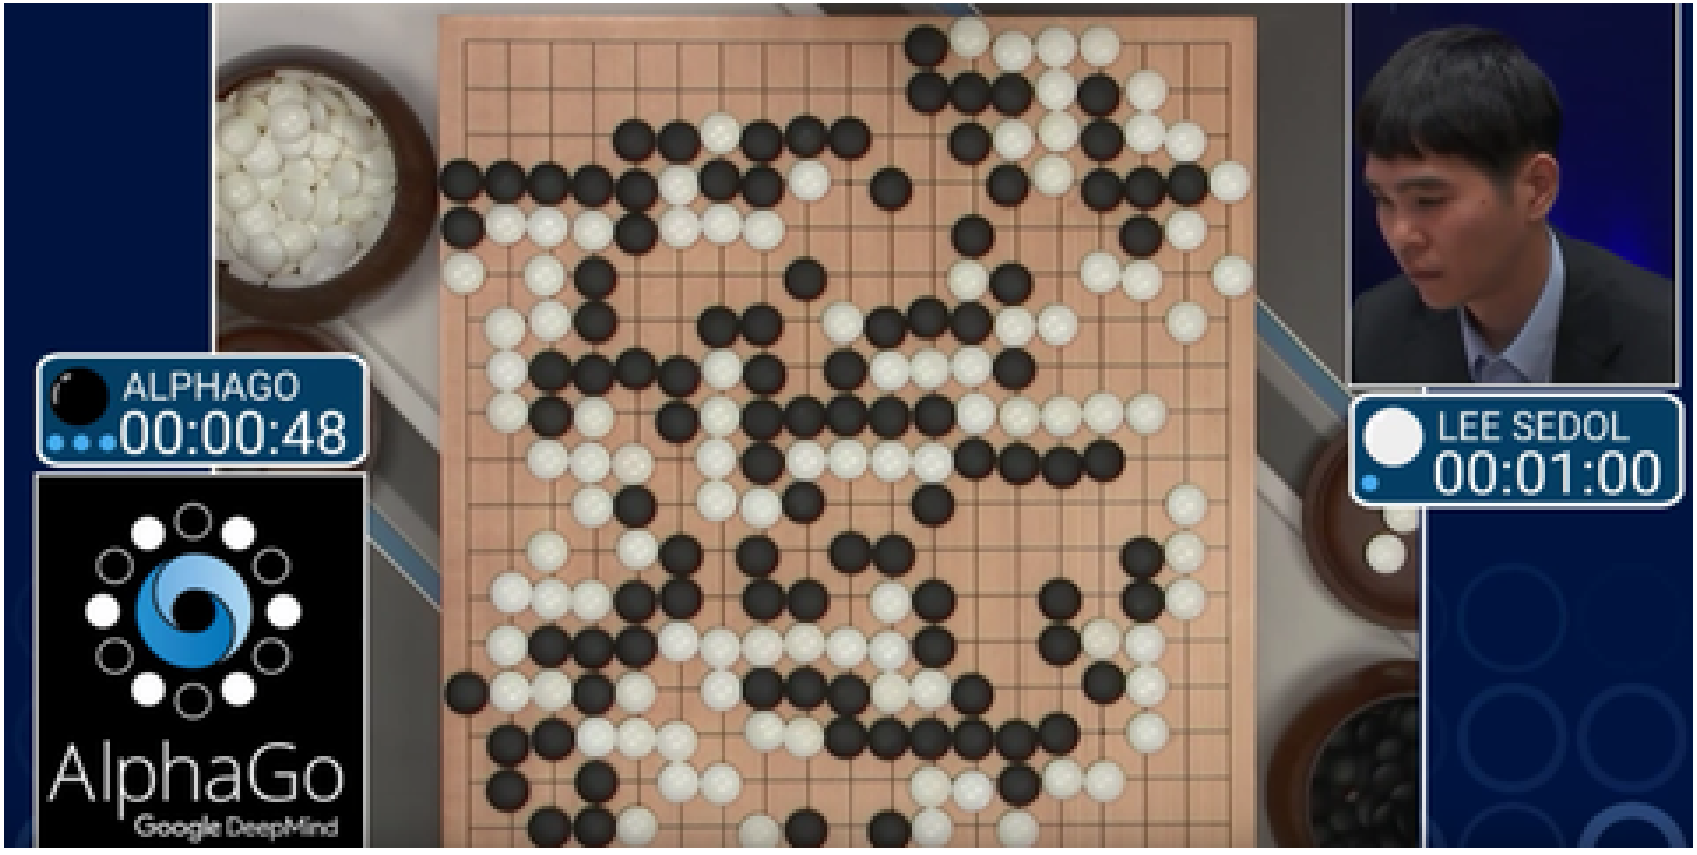

Deep Learning Caveats: Interpretability

Figure 1: from https://xkcd.com/1838/

Deep Learning Caveats: Adversarial Examples

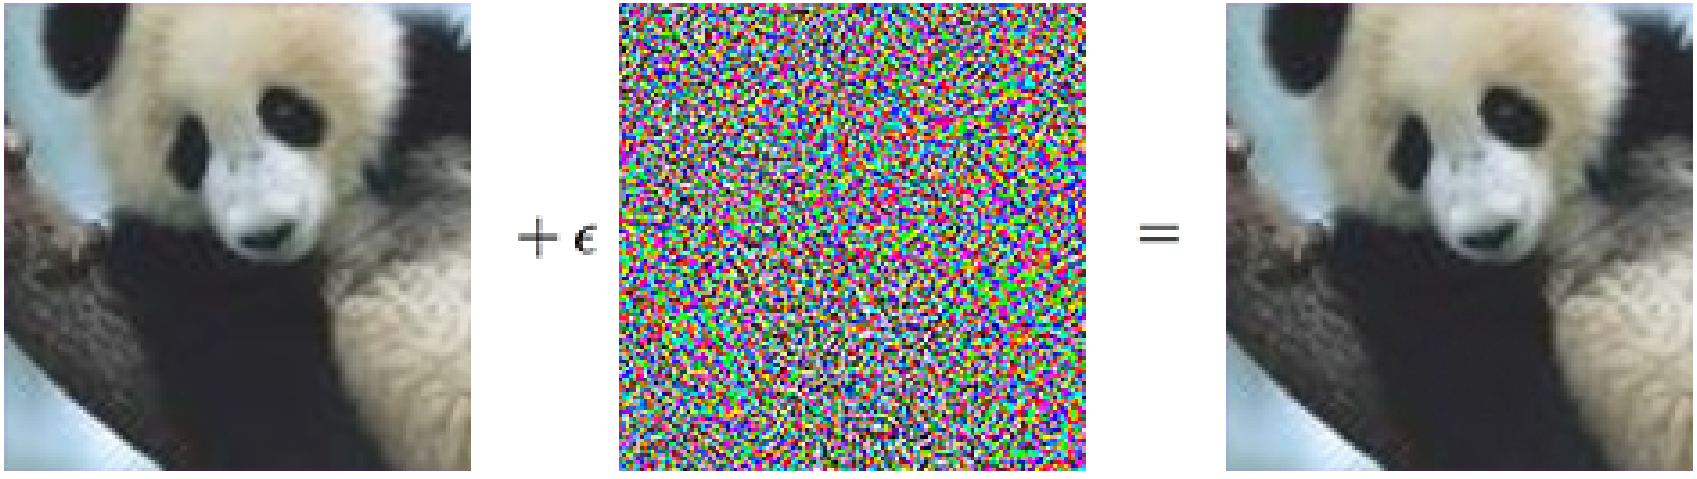

"panda"

57.7% confidence

"gibbon"

99.3% confidence

Deep Learning Caveats: Fairness

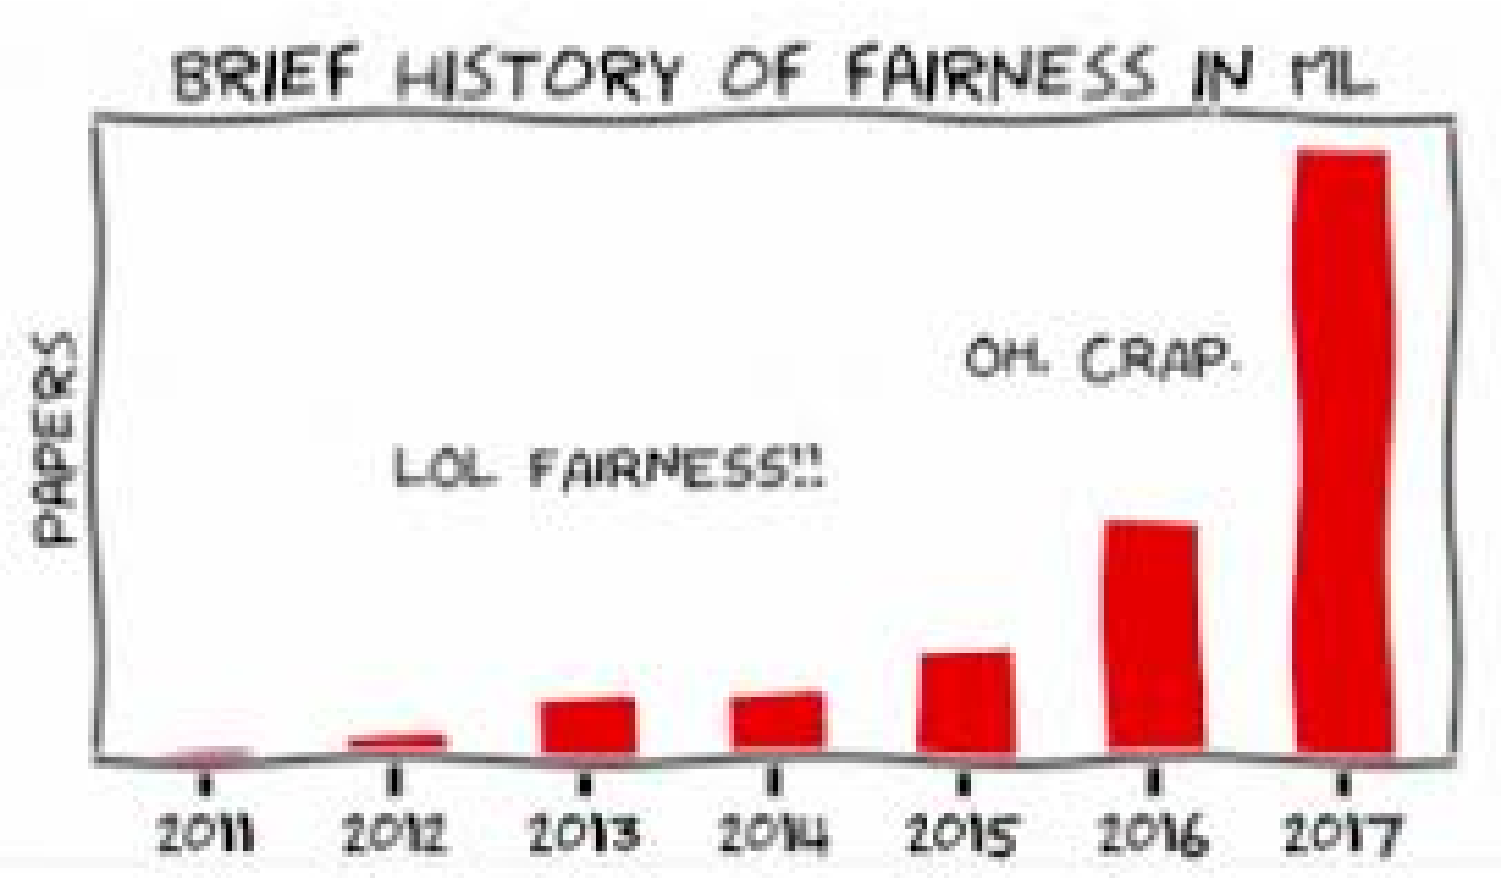

Example

The U.S. military built an AI tool to find suitable combat personnel but had to shut it down because it was discriminating against women

News

Course Coverage

- We will focus exclusively on neural networks and deep learning.
- Mostly supervised learning
- Some unsupervised learning
- A tiny bit of reinforcement learning

Course Philosophy

- Top-down approach
  - Learn by doing
  - Explains the entire system first
  - Details in future courses
- We will introduce very little math
- Focus on implementation and software skills
- Focus on communication skills

Course Website

https://www.cs.toronto.edu/~lczhang/360/

Course Components

- **Lectures**: Monday (2 hr), Thursday (1 hrs)
- **Labs**: Thursday (1 hr), lead by a TA, in the first half of the course
- **Project**: Implementation project in the second half of the course
- **Notes/Readings**: Notes and readings are posted weekly.
- Any material covered in lectures / tutorials / readings is fair game for the midterm, and final term test.

Teaching Assistants (in alphabetical order)

All your TAs are machine learning researchers.

- Andrew Jung
- Huan Ling
- Farzaneh Mahdisoltani
- Jake Snell

Grade Breakdown

- **Labs**: 15%
- **Project**: 30%
- **Midterm**: 20%
- **Final Term Test**: 35%

Labs

- One per week in the first half of the course
- Done **individually** – must be your own work
- You will be given lab time to complete the lab work
- Typically, you can submit your lab by the following Weds 9pm

Late Policy:

- Penalty-free grace period of 1 hours past the deadline
- Penalty of 20% between 1 hour and 24 hours past the deadline
- **No assignments are accepted after the 24 hour deadline**
- The submission time is based on Quercus clock, and **the late policy is strict**

If you have accessibility concerns, speak to me as early as possible.

Software

- Python 3.6
- NumPy
- PyTorch
- Jupyter Notebooks

All lab handouts will be Jupyter notebooks

All assignments are tested on the Google Colab environment.

Course Project

Work in a group of 3 to build a useful machine learning system.

Everyone must contribute to all parts of the project to earn a grade.

- **Project Proposal**: 3%
- **Progress Meeting with TA Mentor**: 3%
- **Progress Report**: 4%
- **Presentation**: 10%
- **Project Repository**: 10%

Midterm

- June 20th, 6pm-8pm
- Length: 110 minutes
- Location: TBD
- No aids permitted

Final Term Test

- August 1th, 6pm-8:30pm
- Length: 150 minutes
- Location: TBD
- No aids permitted

Schedule (Before Reading Week)

Week	Content	Lab
1	Introduction, Basic Neural Networks	PyTorch
2	Neural Network Training	Cats vs Dogs
3	Multi-class Classification	Data Collection
4	Convolutional Neural Networks	Gesture Recognition
5	Deconvolutions and Autoencoders	Autoencoder
6	Language Models and RNNs	Spam Detection
7	Midterm Week	N/A

Schedule (After Reading Week)

Week	Content
8	Text Generation
9	Generative Adversarial Networks
10	Reinforcement Learning
11	Ethics in AI
12	Exam
13	(Project)
14	Presentations

Away

- I am away May 14 - 26
- Your TA Jake will be covering the lectures in my absense

Questions?

Using Pigeons to Detect Cancer

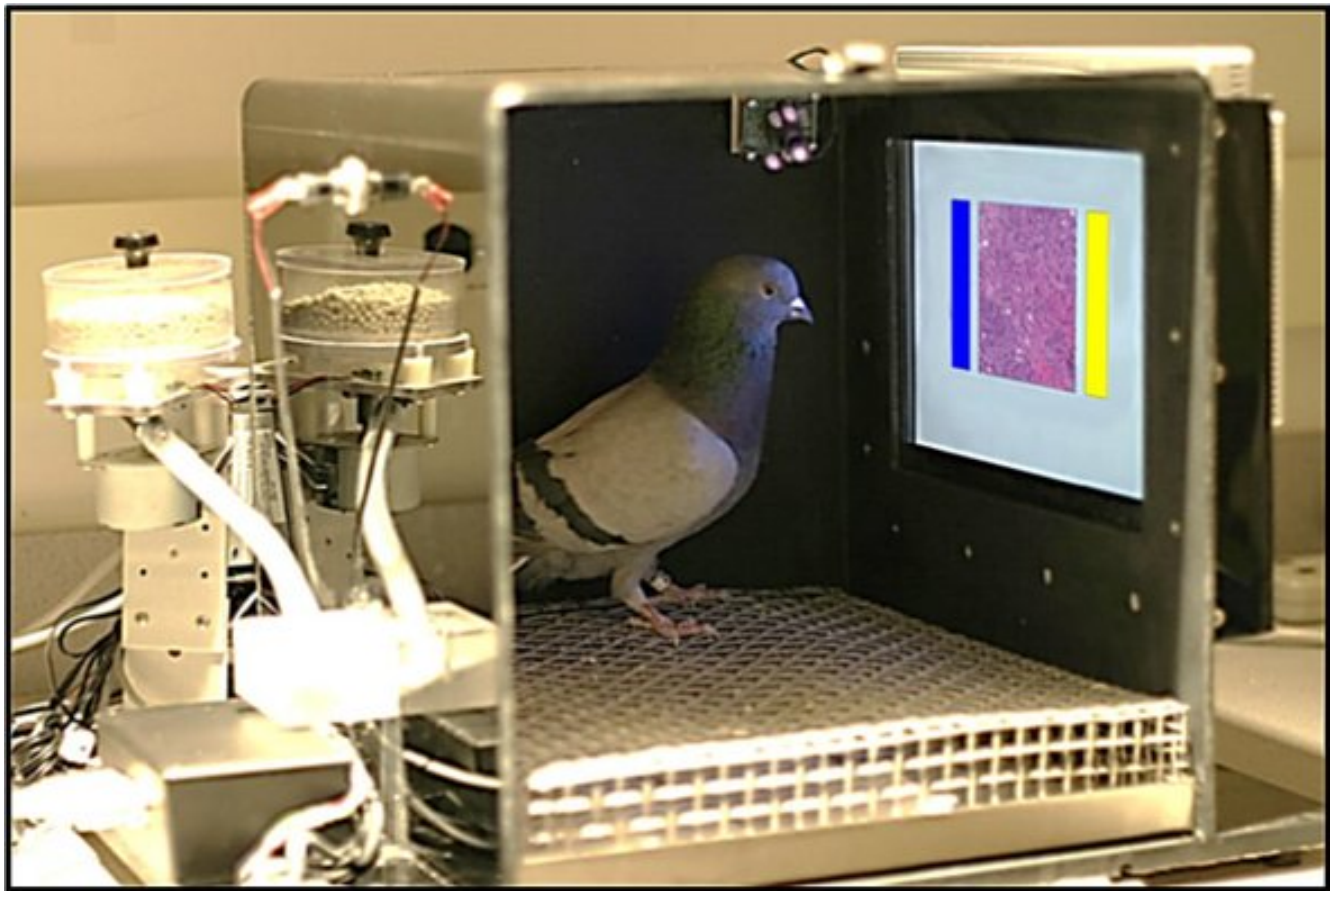

A new study suggests that the common pigeon can reliably distinguish between benign versus malignant tumors and, in doing so, could help researchers develop better cancer screening technologies.

https://www.scientificamerican.com/article/using-pigeons-to-diagnose-cancer/

Training Pigeons

https://www.youtube.com/watch?v=flzGjnJLyS0

1. Show an image of a magnified biopsy to a pigeon
2. Pigeon pecks at one of two answer buttons (cancer or not-cancer)
3. If pigeon picks correctly, reward pigeon with a tasty food pellet

Training an Artificial Neural Network

We need to answer similar questions:

1. How will we reward the pigeon/network?
2. How do we train the pigeon/network quickly and efficiently?
3. How do we know the pigeon/network didn't just memorize the photos?
4. Are there ethical issues in trusting a pigeon/network to detect cancer?

How do pigeons work?

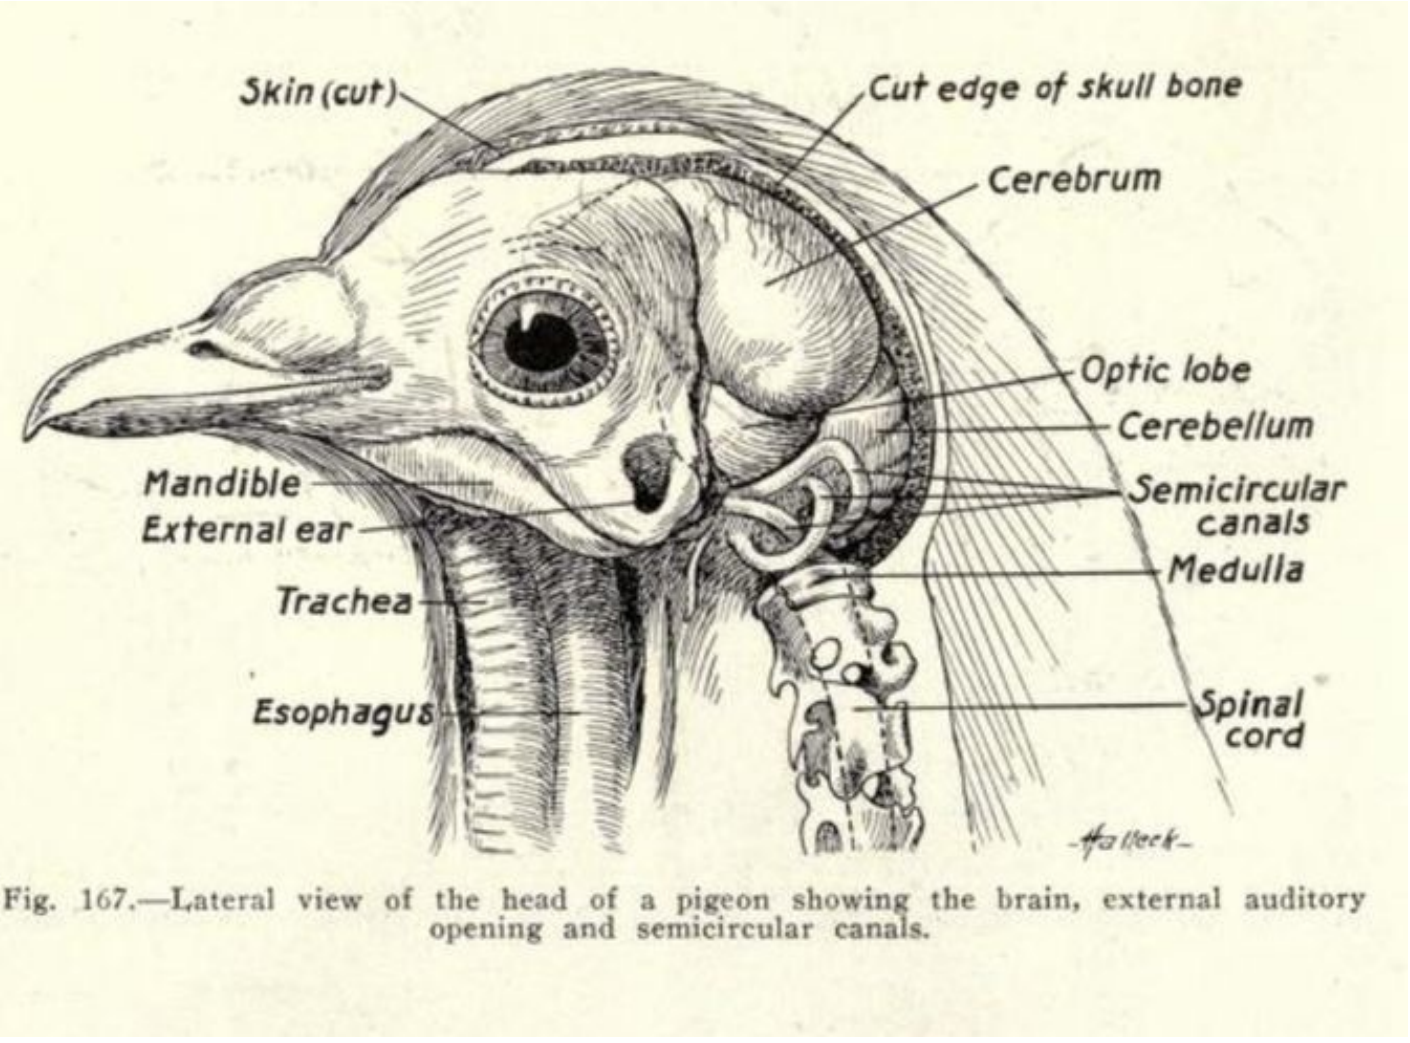

Fig. 167.—Lateral view of the head of a pigeon showing the brain, external auditory opening and semicircular canals.

Neuron

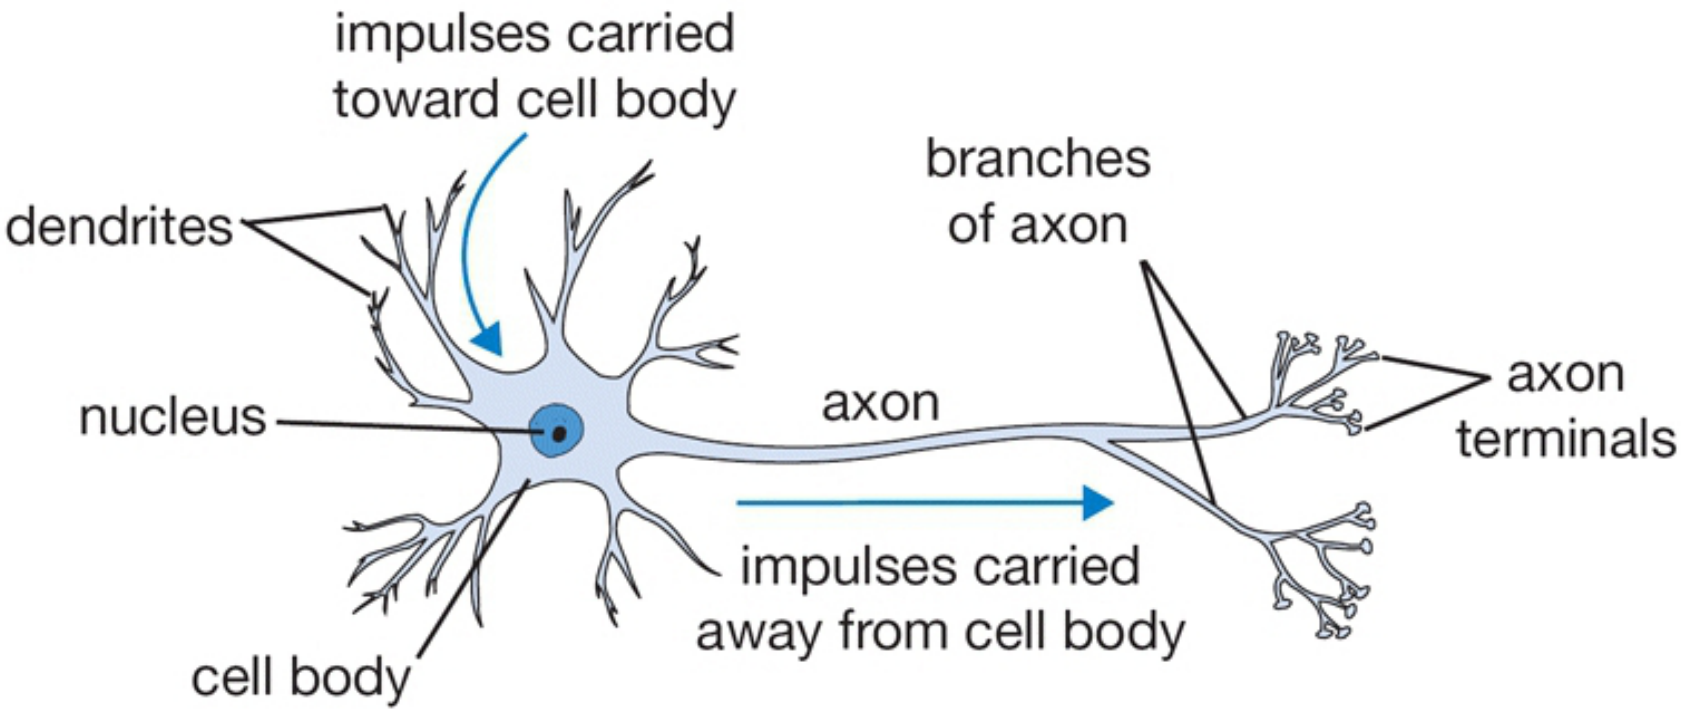
impulses carried
toward cell body
dendrites
branches
of axon
nucleus
axon
axon
terminals
impulses carried
away from cell body
cell body

Neuron (Simplified)

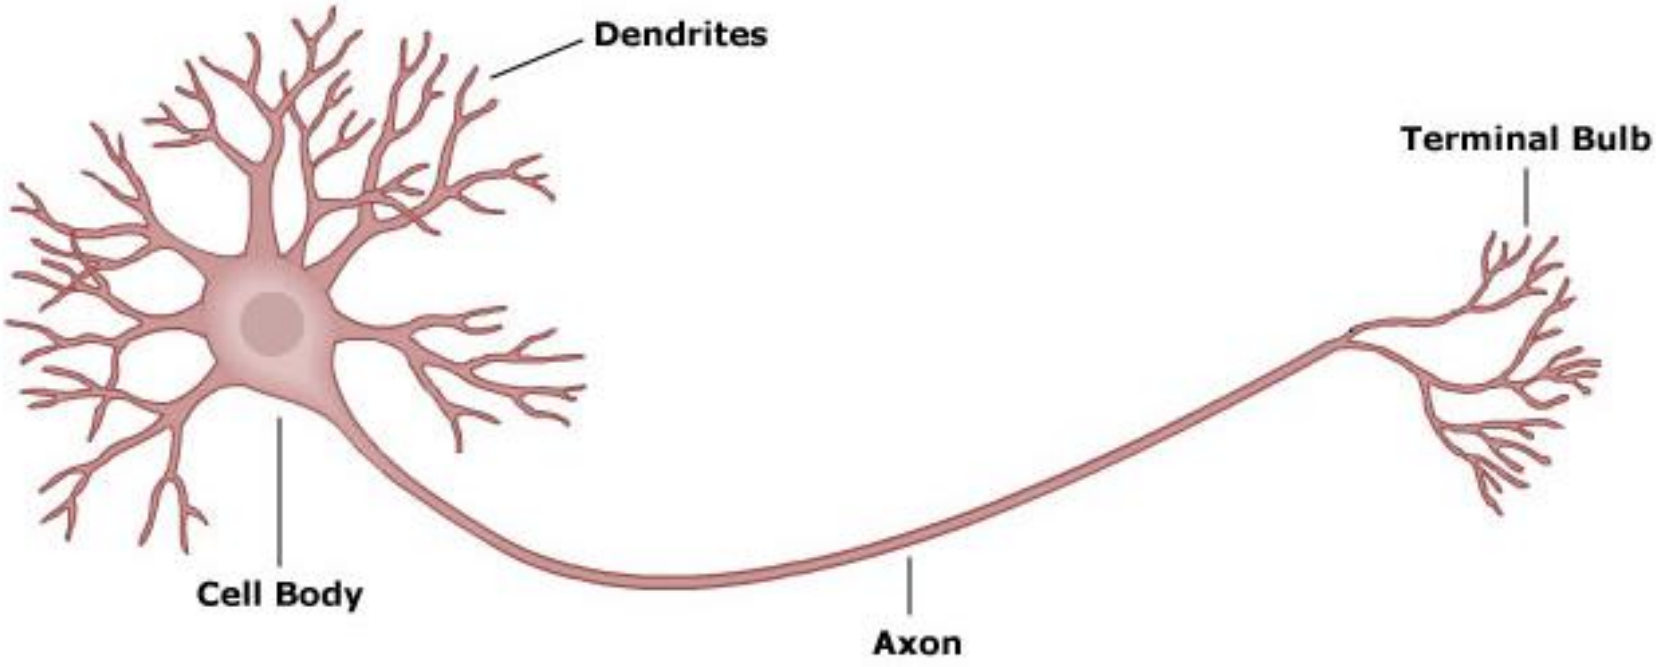

Neuron Anatomy

- The **dendrites**, which are connected to other cells that provides information.
- The **cell body**, which consolidates information from the dendrites.
- The **axon**, which is an extension from the cell body that passes information to other cells.
- The **synapse**, which is the area where the axon of one neuron and the dendrite of another connect.

Synapse

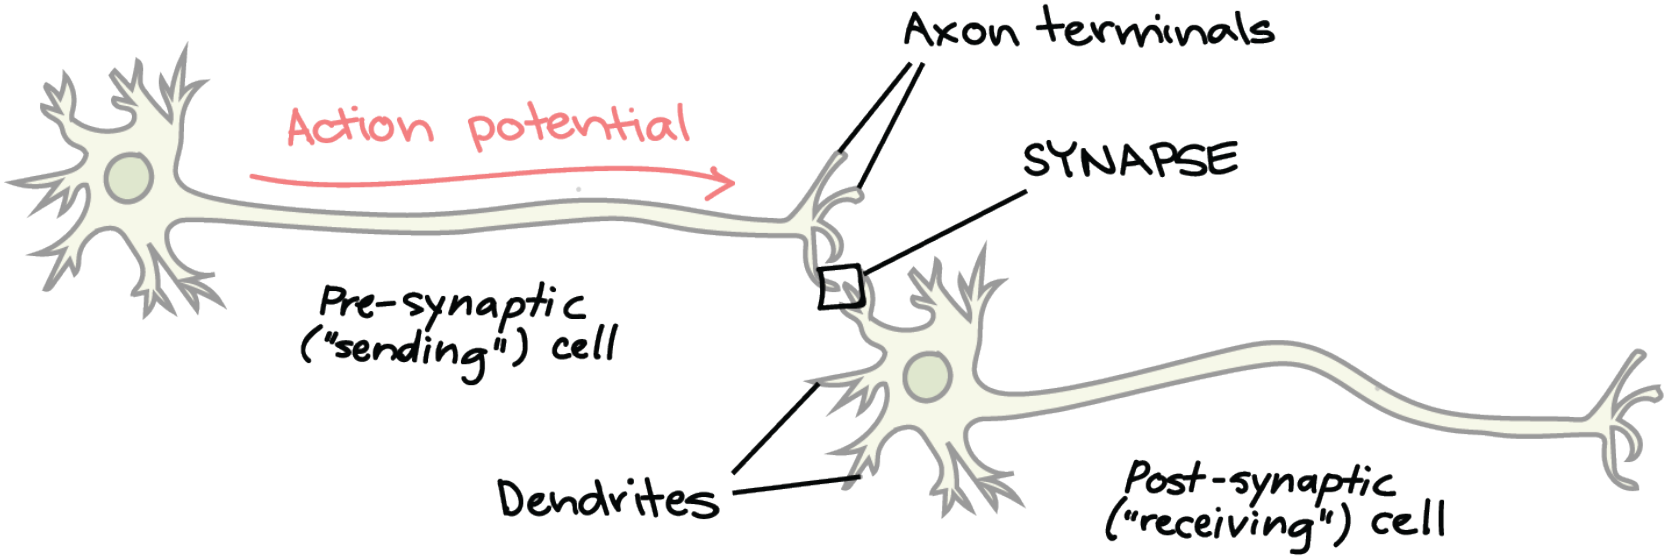
Axon terminals
Action potential
SYNAPSE
Pre-synaptic
("sending") cell
Dendrites
Post-synaptic
("receiving") cell

Synapse

- Small voltage difference between inside and outside of cell
- When a neuron receives "information" in its dendrites, the voltage difference along that part of the cell lowers.
- If the total activity in a neuron's dendrites lowers the voltage difference enough, the entire cell *depolarizes* and the cell **fires**.
- The voltage signal spread along the axon and to the synapse, then to the next cells.

Neural Decoding

What does it mean when a particular neuron fires?

Neuron can fire in response to...

- retinal cells
- certain edges, lines, angles, movements
- hands and faces (in primates)
- specific people like Jennifer Aniston (in humans)

Grandmother Cell

- A neuron that represents a complex but specific concept or object
- Its existence is contested

Distributed encoding

The idea that neuron firing patterns encode information only in a **distributed** fashion

Artificial Pigeon Brain

- Start with an **output** "grandmother cell" that represent the concept that we want to predict.
- Also, start with **input** neurons that activate with each pixel
- Connect input to outputs

Biological Neuron Connectivity

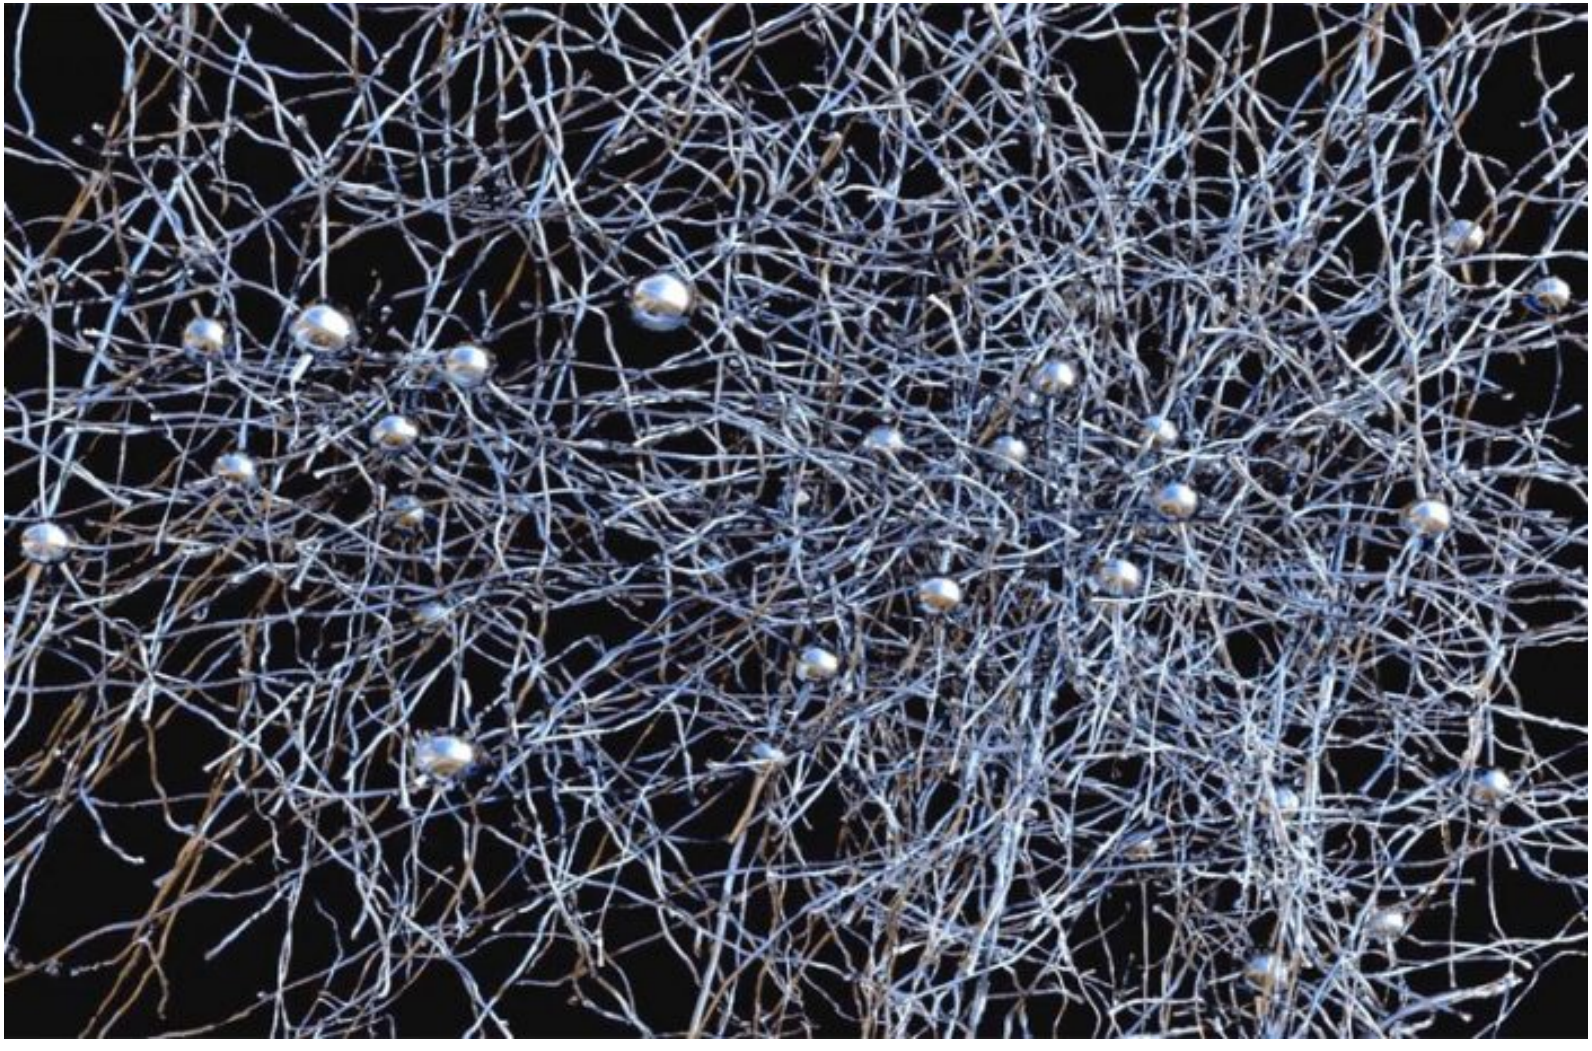

Artificial Neuron Connectivity

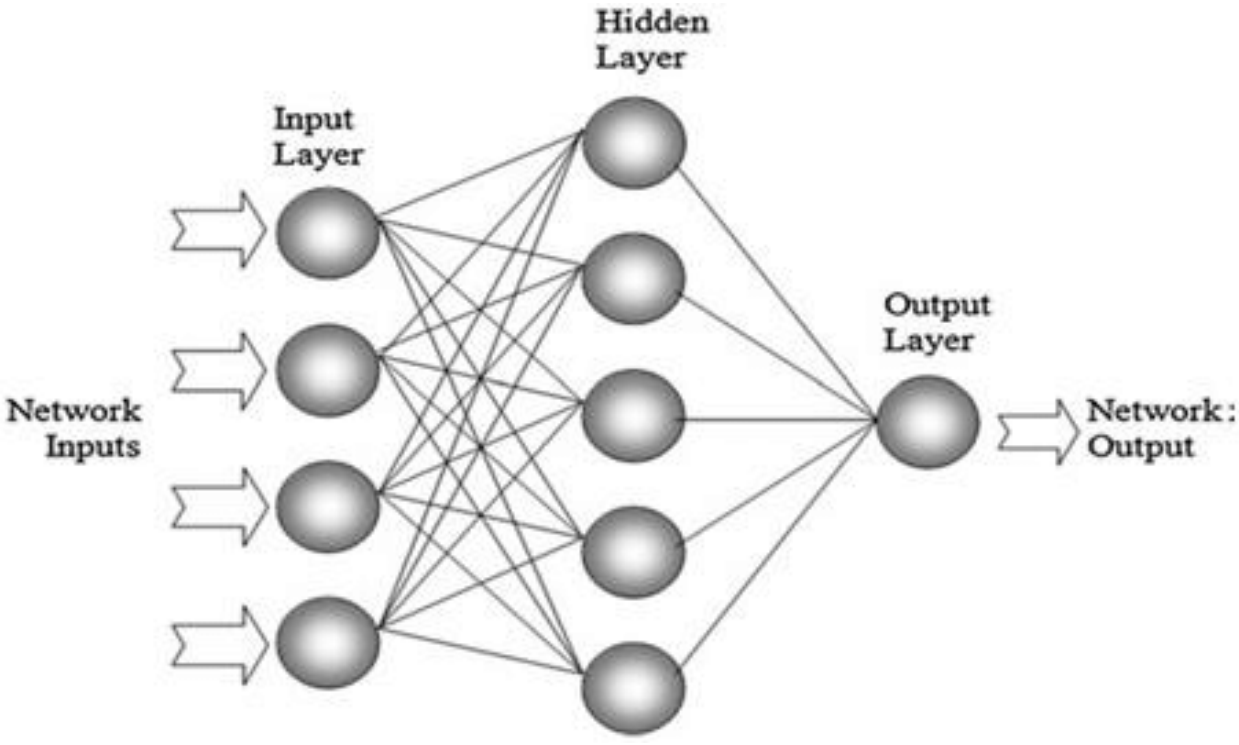

- add a **hidden layer** that don't have specific meaning
- **fully-connected, feed-forward network**

Modelling Individual Neurons

- $x_1, x_2, ... =$ the neurons **activation** of input layer neurons
- $h_1 =$ the neuron activation of a hidden layer neuron
- $y =$ the neuron activation of the output layer neuron

$$h_1 = \sigma(b_1 + \sum_i w_{1,i} x_i)$$

- $w_{1,i} =$ a **weight** summarizing the connectivity of neurons x_i and h_1
- $b_1 =$ a **bias** summarizing the activation requirement of the neuron h_1

Linear Algebra

- **x** - input layer activations
- **h** - hidden layer activations
- y - output unit

$$\mathbf{h} = \sigma(\mathbf{b} + \mathbf{W}\mathbf{x})$$

- W = a matrix of **weights**
- b = a vector of **biases**

But what do the neurons mean?

- Use x_i to encode the input
  - e.g. biopsy image
  - like the neurons that are connected to the receptors in the eye
- Use y to encode the output (of a binary classification problem)
  - e.g. cancer vs. not cancer

Start with a problem (task)

- Input: An 28x28 pixel image
- Output: Whether the digit is a **small** digit (0, 1, or 2)
  - output=1 means that the digit is small
  - output=0 means that the digit is not small

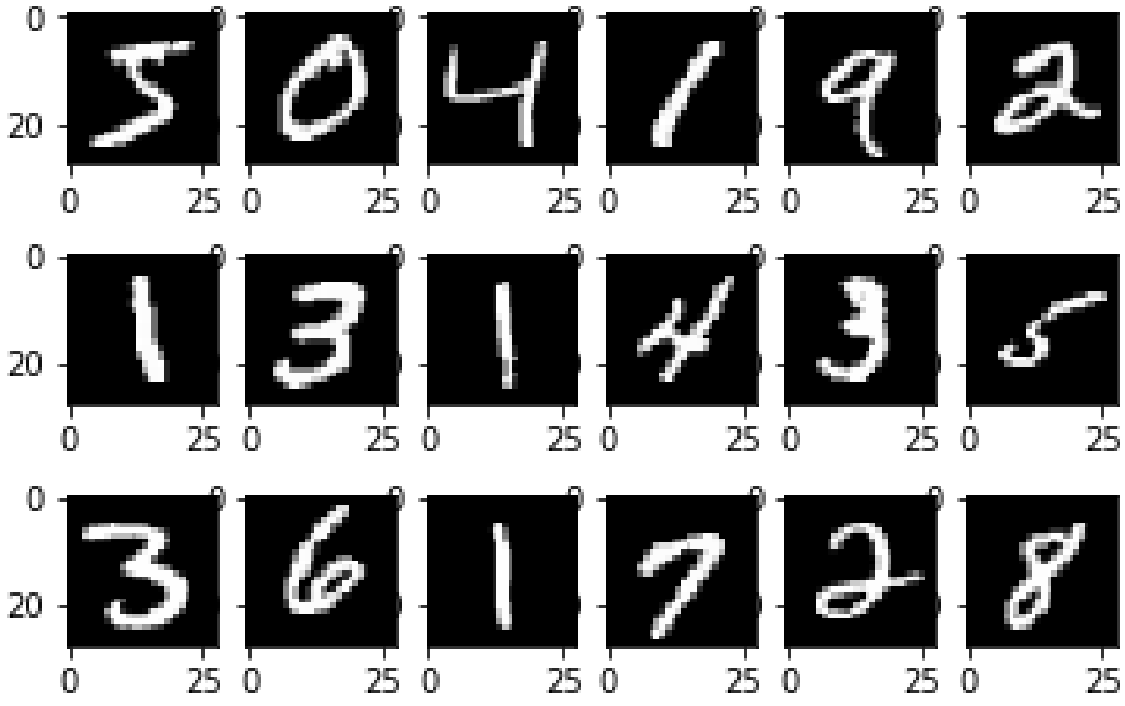

Let's write some code!

Next Class:

- We'll build an artificial pigeon together