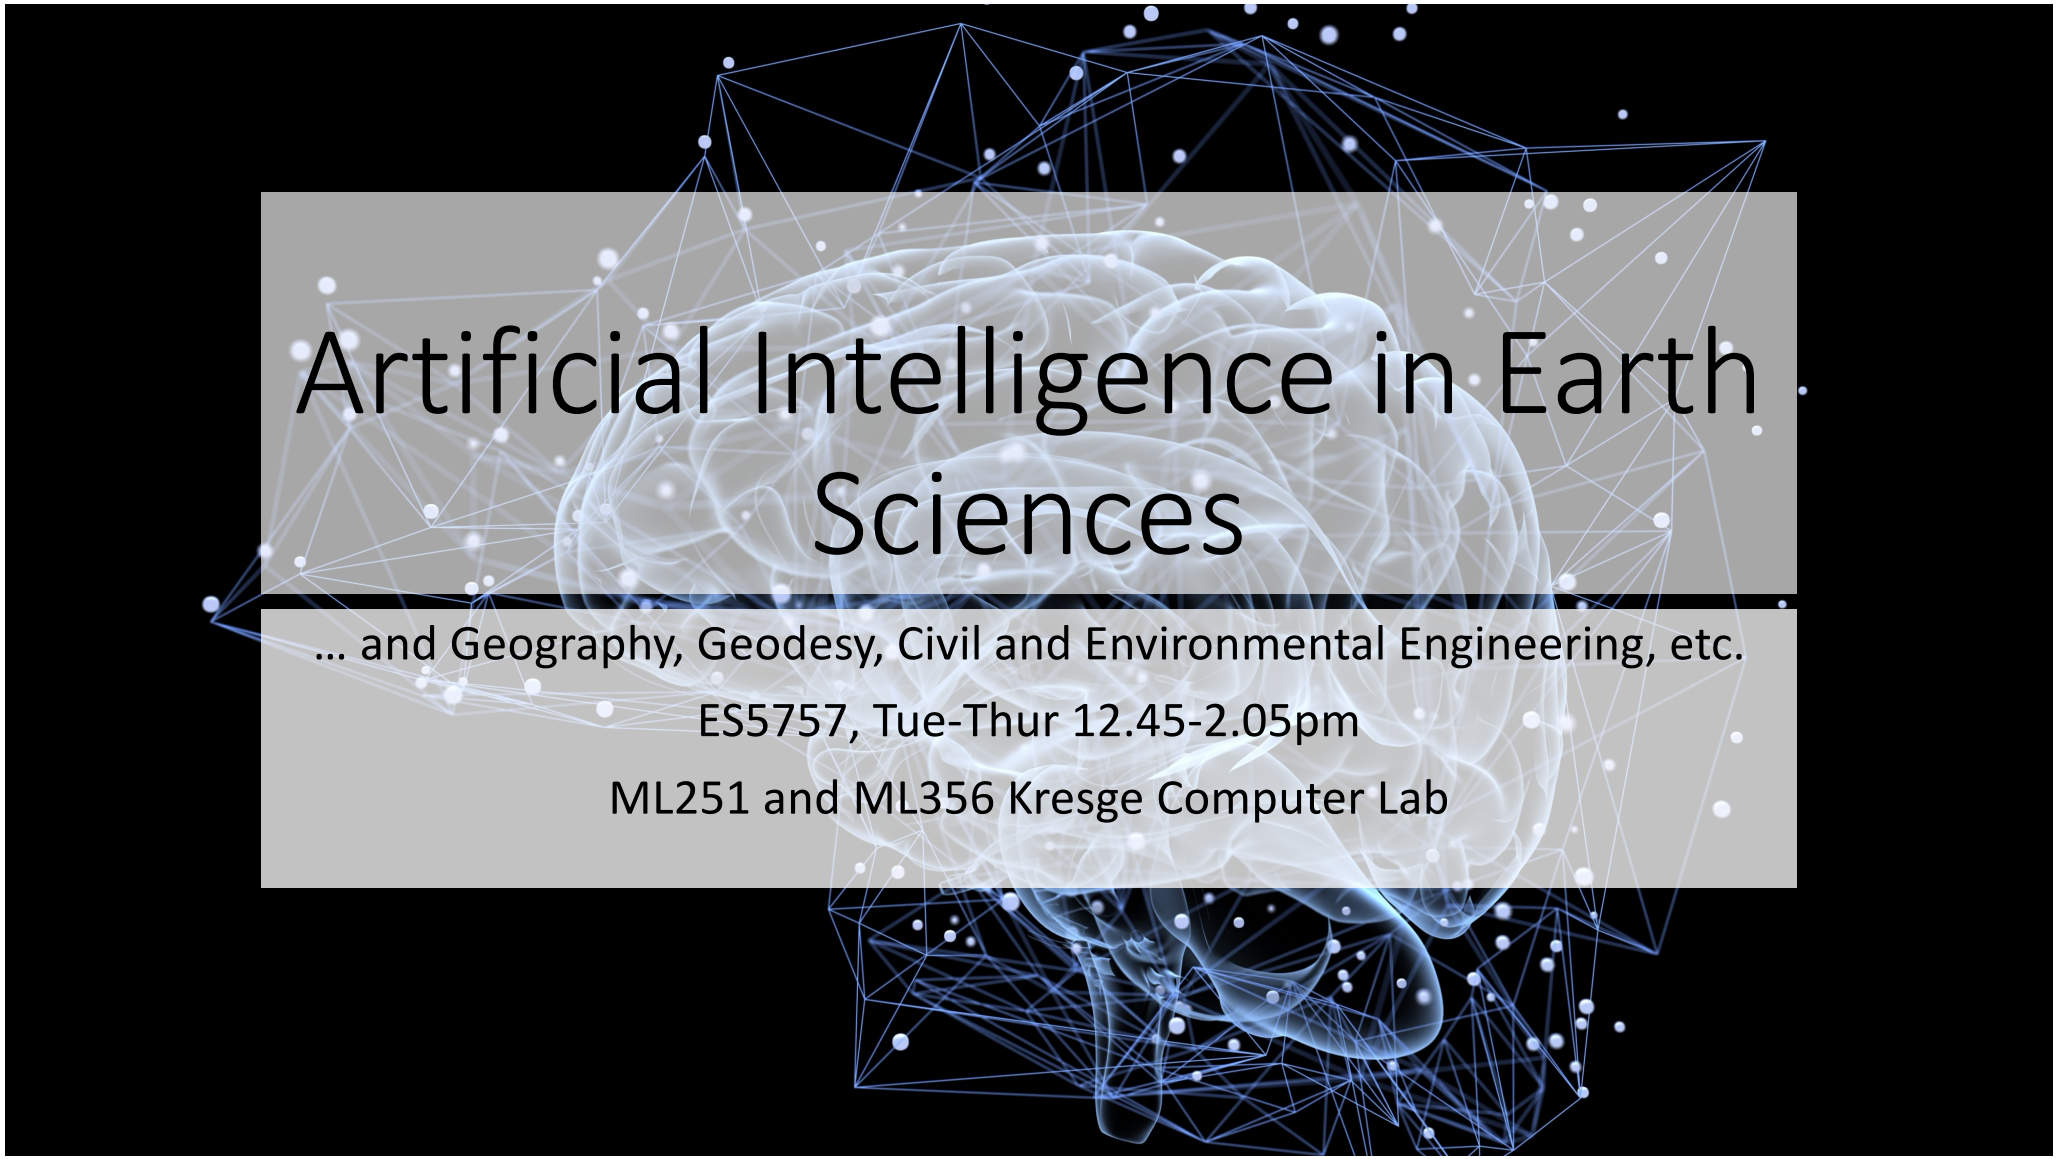

Artificial Intelligence in Earth Sciences

… and Geography, Geodesy, Civil and Environmental Engineering, etc.

ES5757, Tue-Thur 12.45-2.05pm

ML251 and ML356 Kresge Computer Lab

TA: Haniyeh Tajer
PhD student Earth Sciences and Astronomy
Studies exoplanets
McPherson building, room 4042
Office hours: Thursday 11am-12pm (email before?)

Why should you care about AI?

Some numbers:

- 300-350B$ spending by tech giants in 2025 alone! 1.8TB by 2030
- 78% of companies use AI, 100% jump in 2025, 56% higher wages
- All sectors

- Short term: accelerate your research

Goals

- Learn how to program in Python
- Use interactive Jupyter notebooks
- Learn how powerful machine learning algorithms work (no black box)
- Apply machine learning to Science & Engineering problems

Course Structure

- Tuesdays: lectures on how ML/DL algorithms work: ML251
- Remainder: hands-on implementation and applications: ML356, Kresge computer lab
  - Weekly labs
  - Final independent ML project
- Some independent work expected between in-person sessions

Syllabus

- Syllabus, lecture & lab notes shared on Carmen: https://osu.instructure.com/courses/189631
- Week-to-week schedule may be adapted based on collective progress
- Brief overview:

Week 1 (this week)

- Set up google account and Colab(oratory) environment
- Create your first Jupyter notebook
- Work through tutorials of Python basics

Subsequent Weeks, Topics

- Machine Learning terminology
  - Artificial Intelligence vs Machine Learning vs Deep Learning
  - Cost/loss functions, features, gradient descent methods, etc.
- Linear numerical regression in one variable
- Linear numerical regression in multiple (many!) variables
- Non-linear numerical regression

Subsequent Weeks, Topics

- Logistical regression = classification
- Implement logistical regression of one, then multiple, variables

- Other ML methods: Random Forests, Decision Trees, etc.

Subsequent Weeks, Topics

- Neural Networks
  - Perceptrons
  - Artificial Neural Networks (ANN)
  - (Fully) Convolutional Neural Networks

- Applications: Remote Sensing with optical satellite imagery

Course Materials

- Powerpoints of lectures
- Jupyter notebooks with more detailed and interactive lecture notes
- Jupyter notebooks for labs, applying ML/DL algorithms

- Final assignment: solve a science problem with ML/DL from scratch

Evaluation

- Finished jupyter notebooks of weekly labs uploaded to Carmen
- Final independent project
- Final grade based on participation (10%), completion of lab assignments (50%) and final project (40%).

Use of LLM AI

- Encouraged, but be smart about it

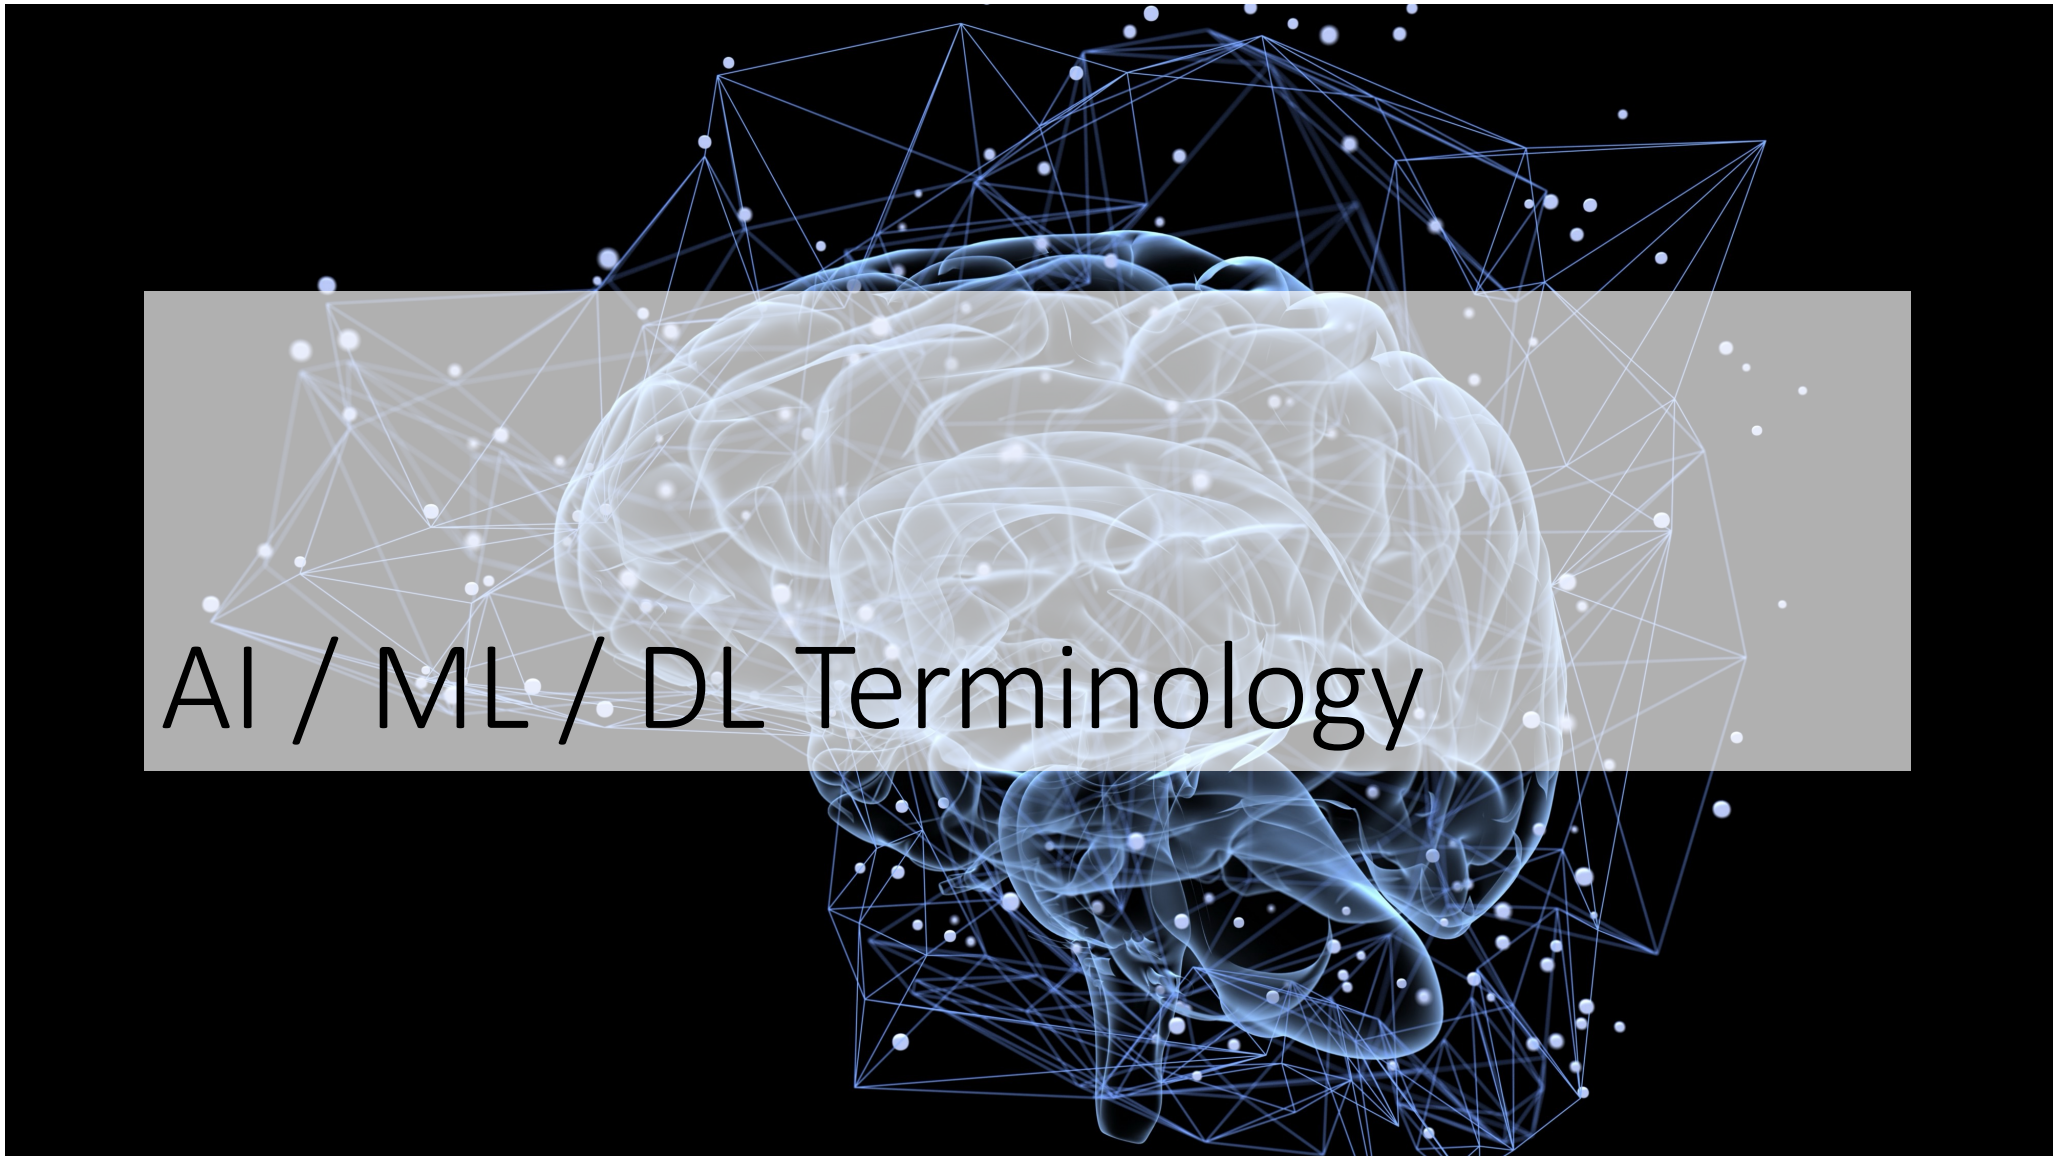

AI / ML / DL Terminology

A bit of terminology

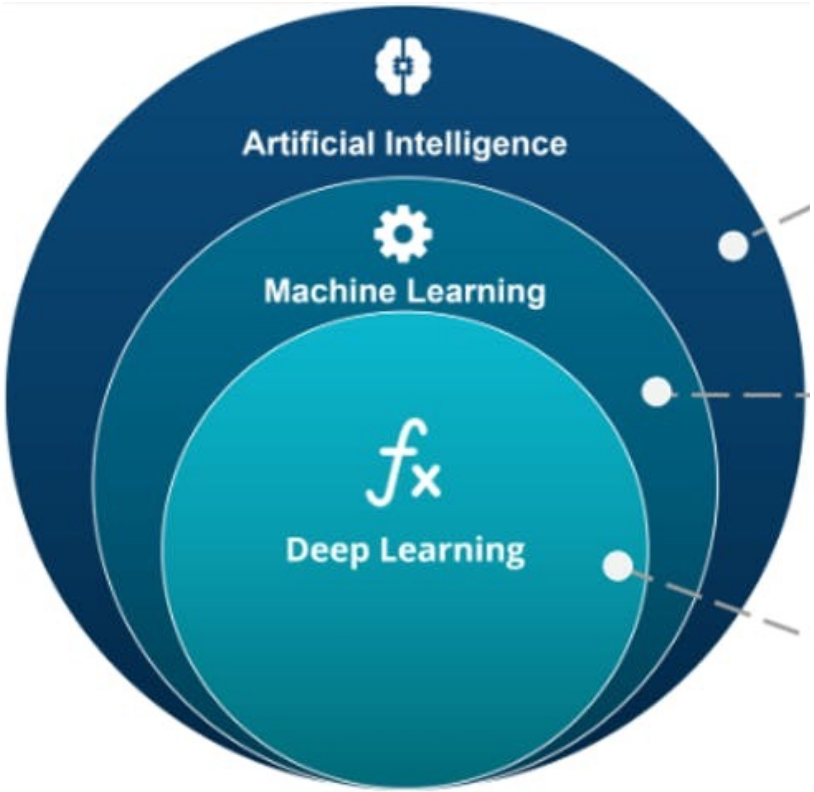

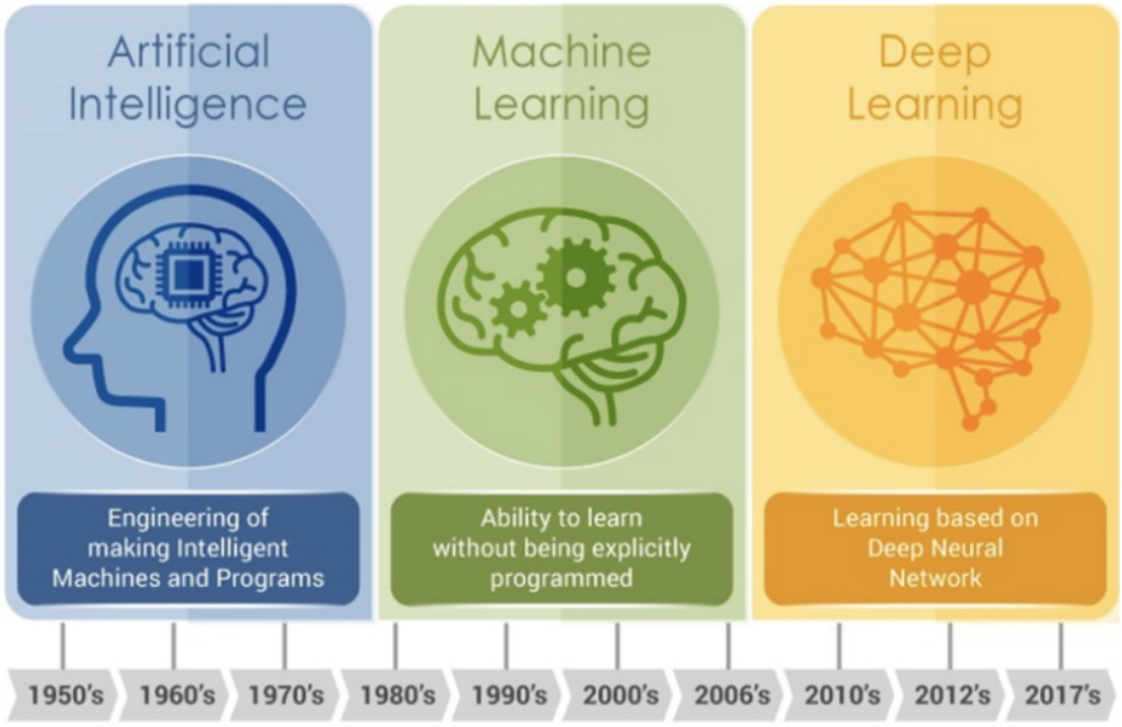

Evolution of AI — Source: https://www.embedded-vision.com/

Definitions of Artificial Intelligence

Oxford:

"*an area of study concerned with making computers copy intelligent human behavior.*"

Webster:

"the capability of a machine to imitate intelligent human behavior"

Britannica:

"Research in AI has focused chiefly on the following components of intelligence: learning, reasoning, problem solving, perception, and using language."

Definitions of Machine Learning

Oxford:

"*the use and development of computer systems that are able to learn and adapt **without following explicit instructions**, by using algorithms and statistical models to analyze and draw inferences from patterns in data.*"

Webster:

"the process by which a computer is able to improve its own performance (as in analyzing image files) by continuously incorporating new data into an existing statistical model"

Definitions of Deep Learning

Oxford:

"*A type of machine learning considered to be in some way more dynamic or complete than others; esp. machine learning based on **artificial neural networks** in which multiple layers of processing are used to extract progressively more features from data.*"

Webster:

*"a form of machine learning in which the computer network rapidly teaches itself to understand a concept without human intervention by performing a large number of **iterative calculations on an extremely large dataset**"*

Definitions of Deep Learning

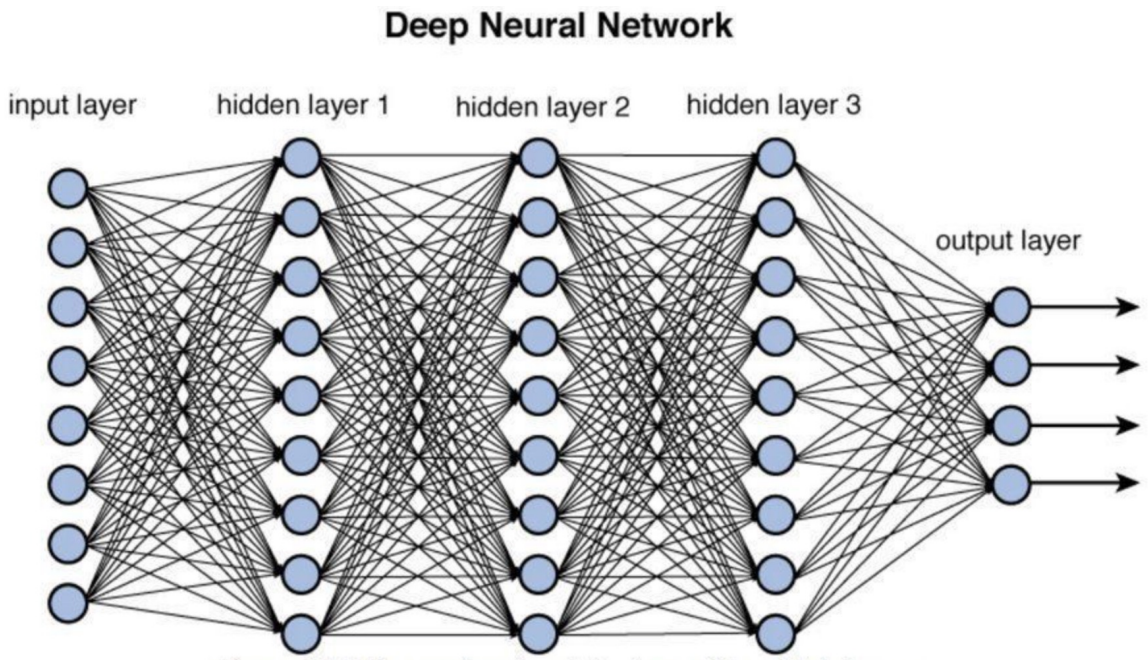

Figure 12.2 Deep network architecture with multiple layers.

Type of ML that uses **neural networks** with many layers

Supervised learning

Numerical regression

- Provide training data of **features** + **target** =
- Independent vs dependent variables

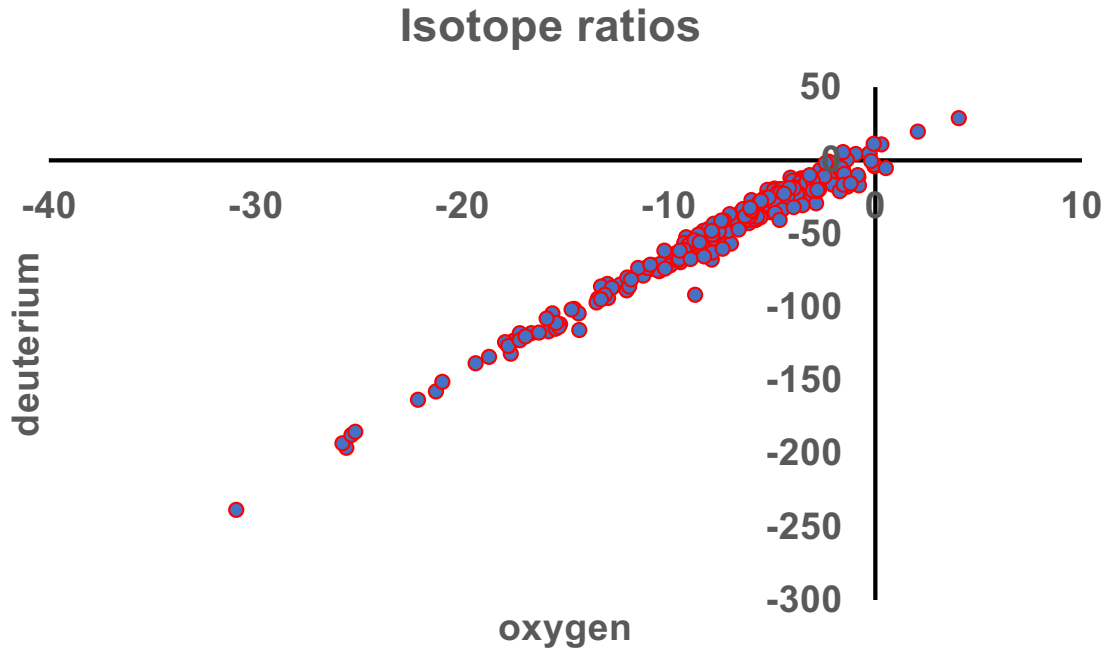

Data	
δ18O [‰]	δ2H [‰]
-9.54	-67.92
-15.23	-112.41
-9.12	-52.92
-12.93	-84.58
-13.26	-86.33
-16.61	-118.56
-11.78	-80.79
-10.71	-72.27
-2.46	-11.03
-8.28	-48.42
-4.09	-12.13
-4.85	-21.23
-11.94	-84.09
-7.74	-49.37
-9.24	-62.31
-12.30	-85.06
-6.02	-35.14
-15.80	-117.16
-7.89	-68.13
-9.34	-67.75
-11.02	-73.50
-13.40	-94.67
-7.21	-53.15
-10.99	-75.04
-8.71	-92.00
-3.88	-23.25
-0.81	-11.85
-5.14	-32.15
-3.08	-24.50
-1.33	-18.22
-5.03	-36.05
-2.39	-3.97
-7.02	-47.58
-8.45	-59.48

Numerical regression

- Use Machine Learning to fit model (**hypothesis**) to **training data**,
- Minimizing error = **loss**, or **cost function** = e.g., RMSE or R^2

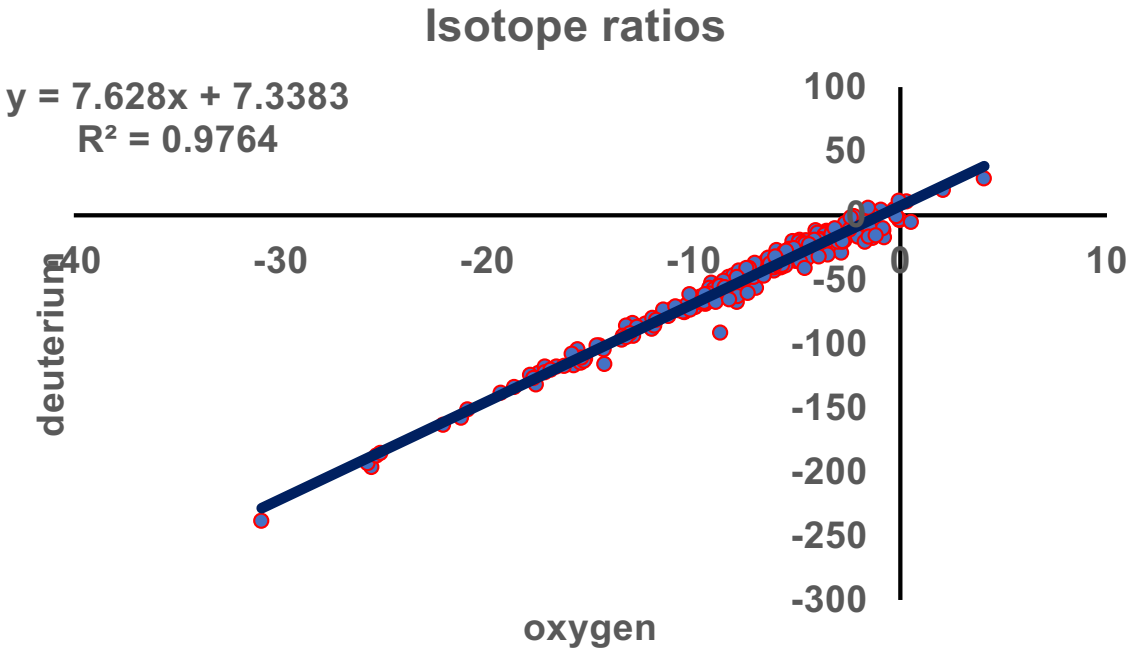

Numerical regression

- Use *independent* **validation data** to test how well model will generalize to make future predictions

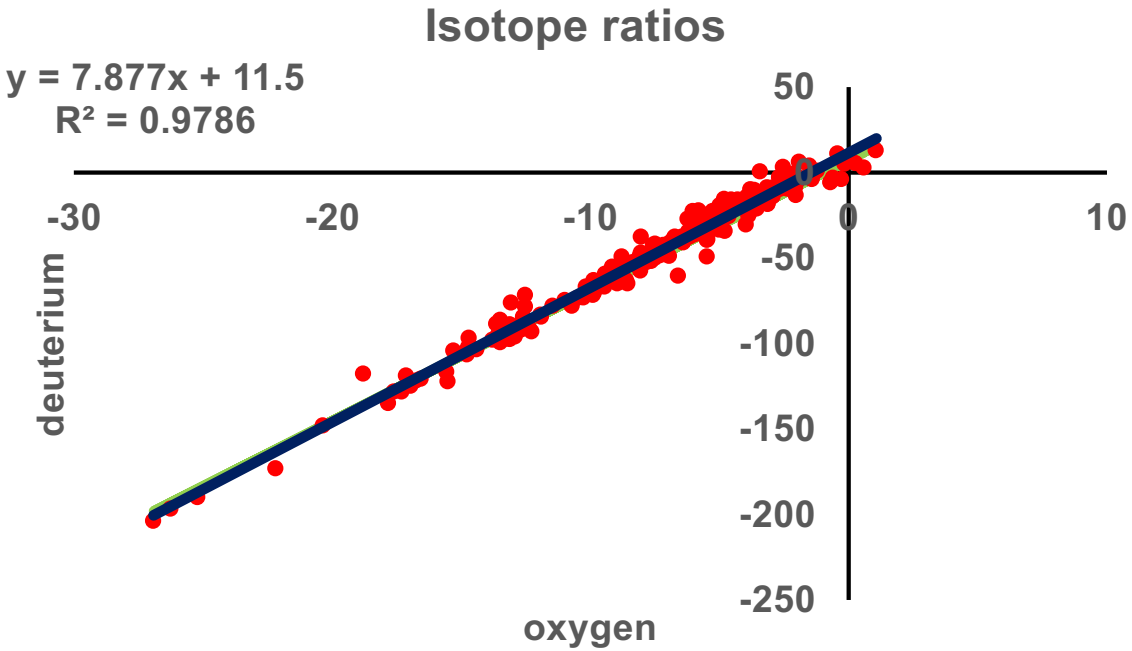

Numerical regression more generally

- Multiple / many features, e.g., precipitation, temperature, pressure etc.
- Higher-order polynomial models or other functions (exponential etc.)
- Neural networks can find extremely complex non-linear functions of one more many input features.

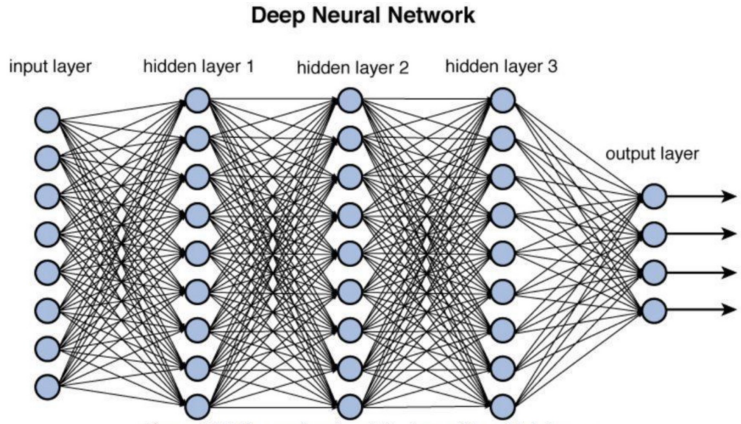

Figure 12.2 Deep network architecture with multiple layers.

Logistical regression

- Instead of continuous predictions, discrete predictions into 2 or more categories: **classification**

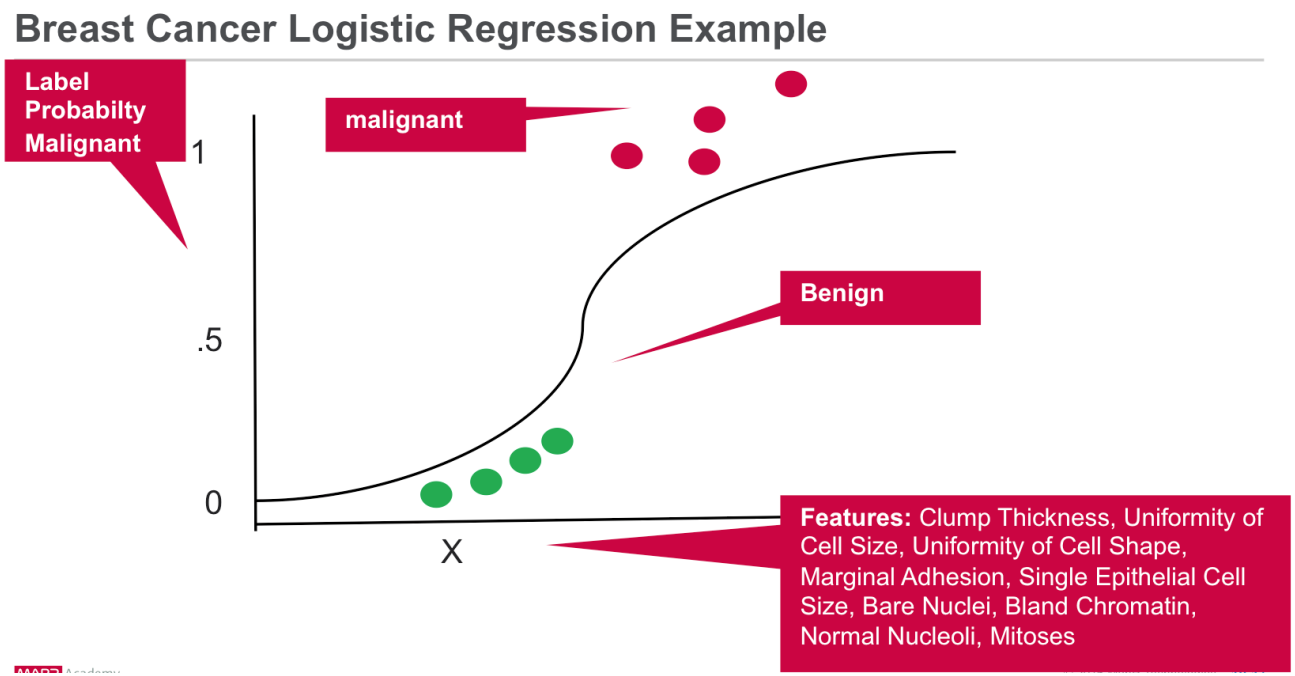

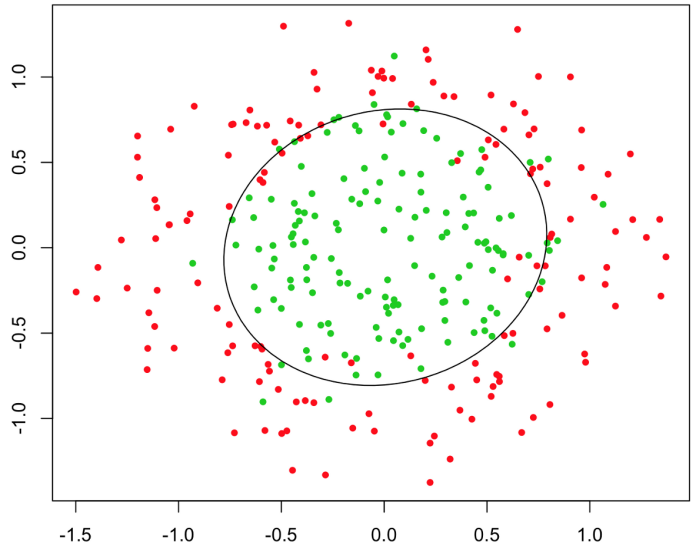

NLP example: Spam filter

Land-Use-Land-Cover (LUCL) Classification

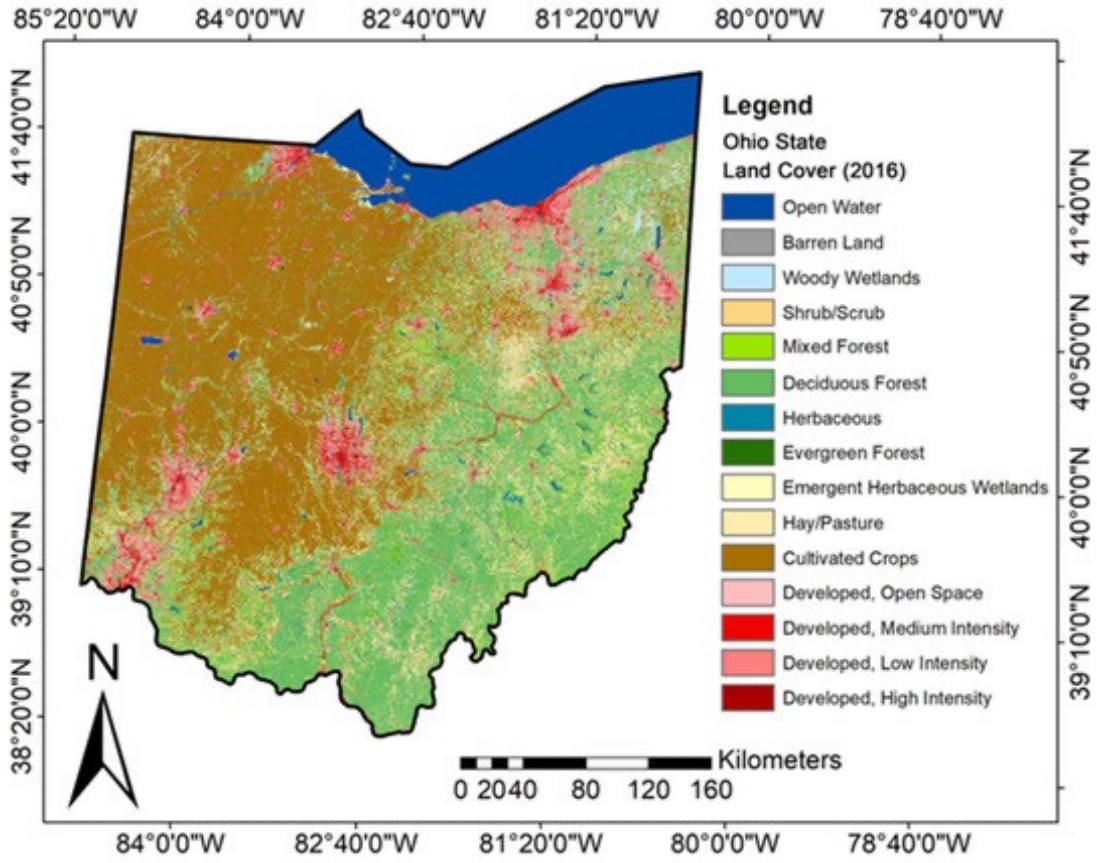

Unsupervised Learning

- Feed computer data, often massive amounts, and let it find patterns.
- Different statistical algorithms, such as Principal Component Analyses, clustering algorithms (e.g., k-means)

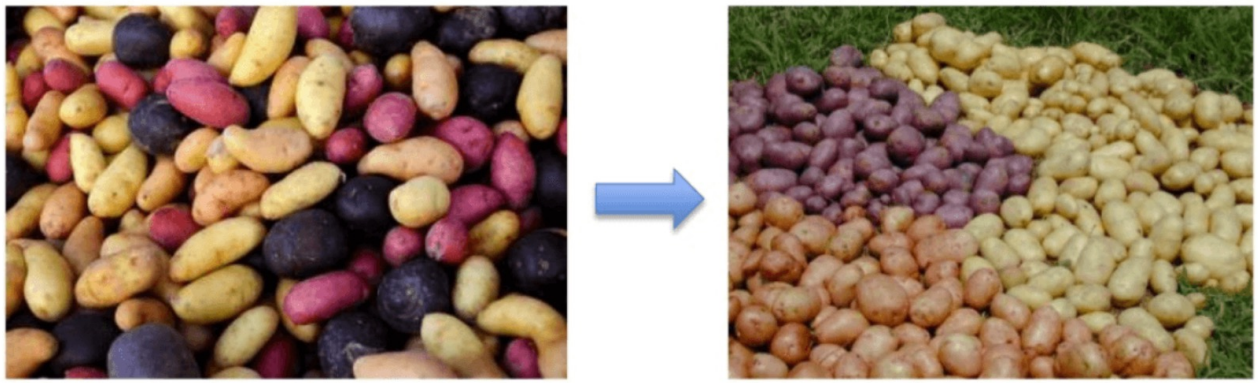

sample

Cluster/group

Unsupervised Learning

- Find different mineral facies in images of rocks

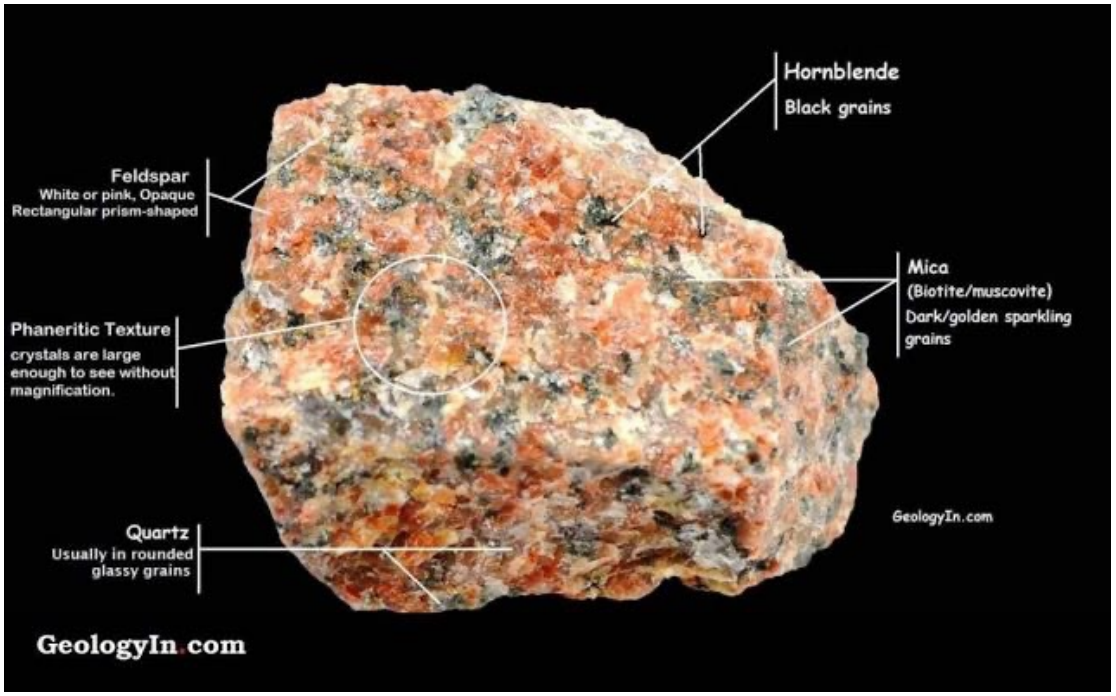

Unsupervised Learning

- Identify different tree categories

Supervised Learning

- Identify different specific tree species, learn from labeled training data

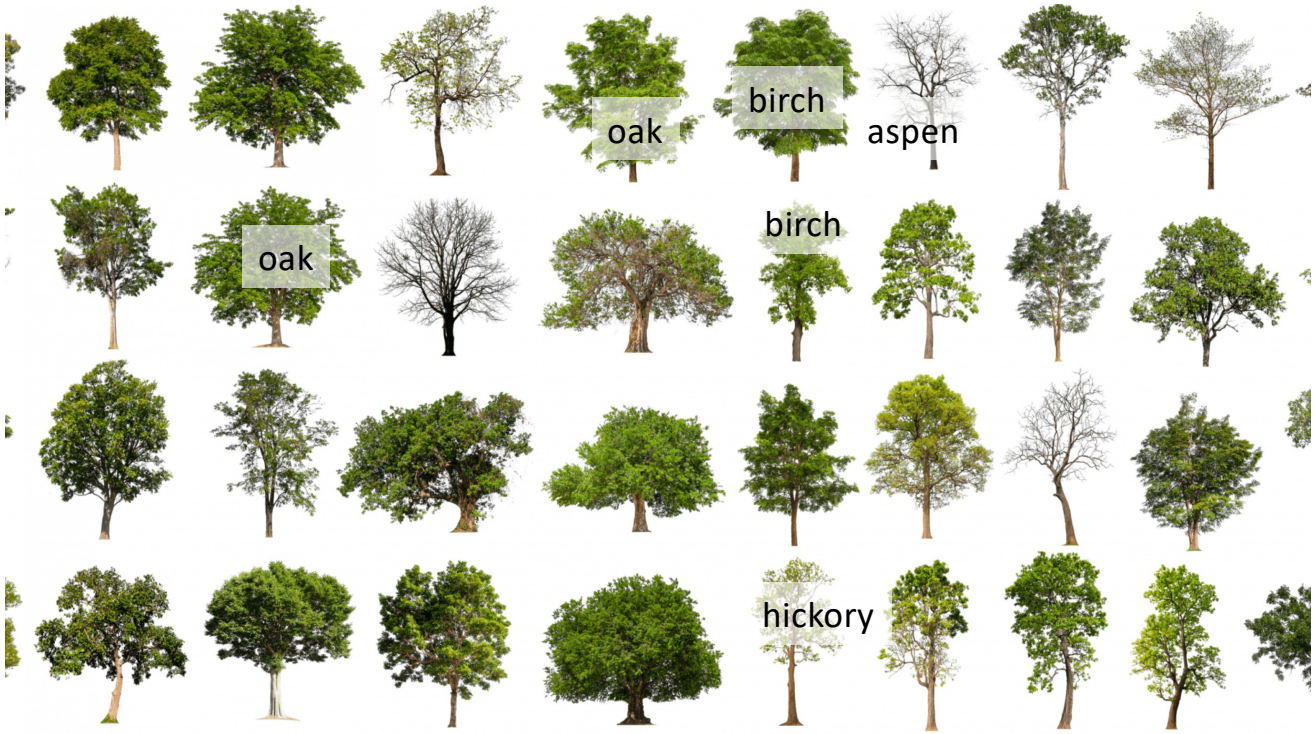

- Make predictions for new input data
- Compute error and optimize
- Provide more labeled training data
- iterate

Supervised learning in computer vision

Classification

CAT

Classification + Localization

CAT

Single object

Object Detection

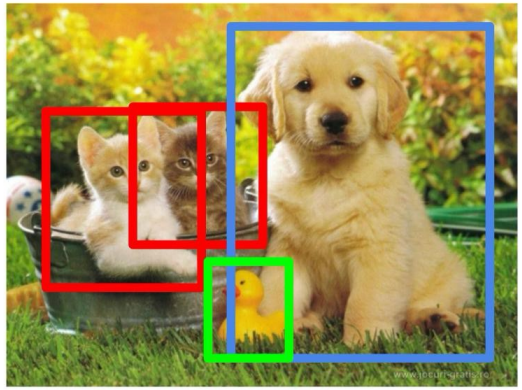

CAT, DOG, DUCK

Instance Segmentation

CAT, DOG, DUCK

Multiple objects

How machine learning works

- **Decision process**: based on labeled or unlabeled input data, make numerical or classification prediction
- **Error/loss/cost function**: evaluate accuracy of predictions
- **Model optimization**: reduce errors between predictions and known values by autonomously adjusting fitting weights

Underfitting and overfitting

- **Underfitting**: model/hypothesis is not complex enough to fit data.
  - Example: fitting linear function to quadratic data
- **Overfitting:** model is too complex for data.
  - Example: fitting quadratic, or higher polynomial, function to linear data
- Test: how well does model perform on new data, especially extrapolating

- To reduce overfitting, use simple model and/or increase training data
- '**Feature engineering**' and **regularization**

What's next?

- Today/Thursday: Lab time to familiarize with Jupyter and Python
- Follow tutorial instructions on Canvas.
- Complete "Getting to Know You" Quiz on Carmen

- Next week: First ML algorithms for Linear Regression