Machine Learning

CPSC 433: Artificial Intelligence
Fall 2024

Jonathan Hudson, Ph.D.
Assistant Professor (Teaching)
Department of Computer Science
University of Calgary

August 8, 2024

UNIVERSITY OF CALGARY

Outlines

- Machine Learning
- Designing Machine Learning

Machine Learning

Machine Learning

- Around concepts around since Optical Character recognition research (1914!)
- Major prominence in 1990s with spam filters

- Machine learning is about **learning patterns from data** in order to make useful **predictions**.

- Good at seeing patterns in images, classifying things with many parts (text), filtering by pattern, chat systems, modeling complex things like, voice recognition, etc.

Relationships

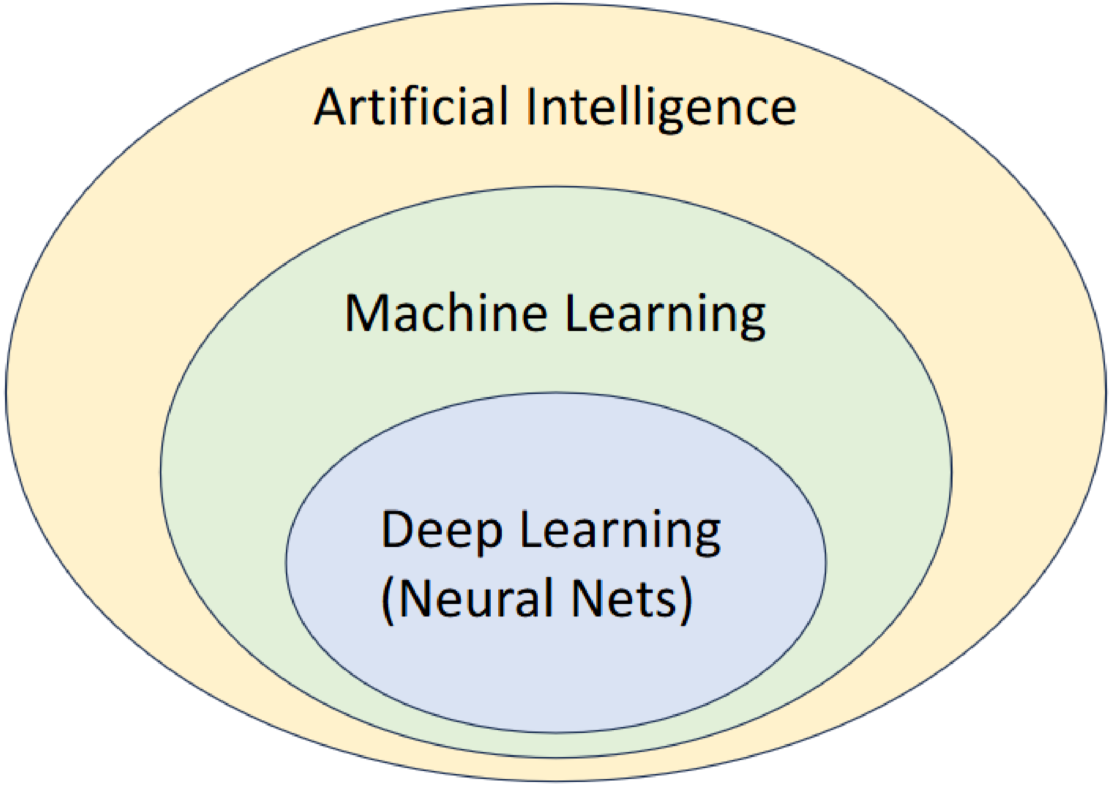

Machine Learning and Humans

- Machine learning is built on the **premise of induction**, i.e. that data we have seen in the past can be used to predict future data.

- Machine learning is **not a magic solution** that solves all problems!

- Even the best ML approaches require a “smart human” to set them up.
  - A human needs to choose
    1. Approach & Assumptions
    2. Model (type of solution you would expect to see)
    3. Data source, encoding
    4. Training algorithm and Evaluation method

Categories

1. **Supervised Learning:** Machine is given guidance
   - Common examples are classifying something under labels, or finding parameters for a model that best matches data
   - e.g. (x, y), where x is an input value and y is the “correct” or “target” output
     - y is also sometimes called a label
2. **Unsupervised Learning:** Machine is left unguided
   - has to decide most everything (often minor guidance like # of things to find)
   - e.g. clustering (determine sub-groups that are ‘similar’)
3. **Reinforcement Learning:** Future actions responses guide what is elevated/added as learned knowledge, and what is reduced or lost

- Of course there are hybrid variants (semi-supervised)

Designing Machine Learning

Design

1. Study the problem you are trying to solve (What?)
2. Choose a model class, hyperparameters (How?)
3. Prepare data (Do.)
4. Run learning algorithm to train the model (Do.)
5. Evaluate trained model (Did it work?)

Design: Model

- A model in machine learning is a hypothesis used for making predictions about data.
- By choosing a model class, you are selecting a space of possible solutions for your learning algorithm to explore.
- e.g. linear regression: $y = mx + b$
  - I have input variable x and output target y (I need learning what m/b match best)

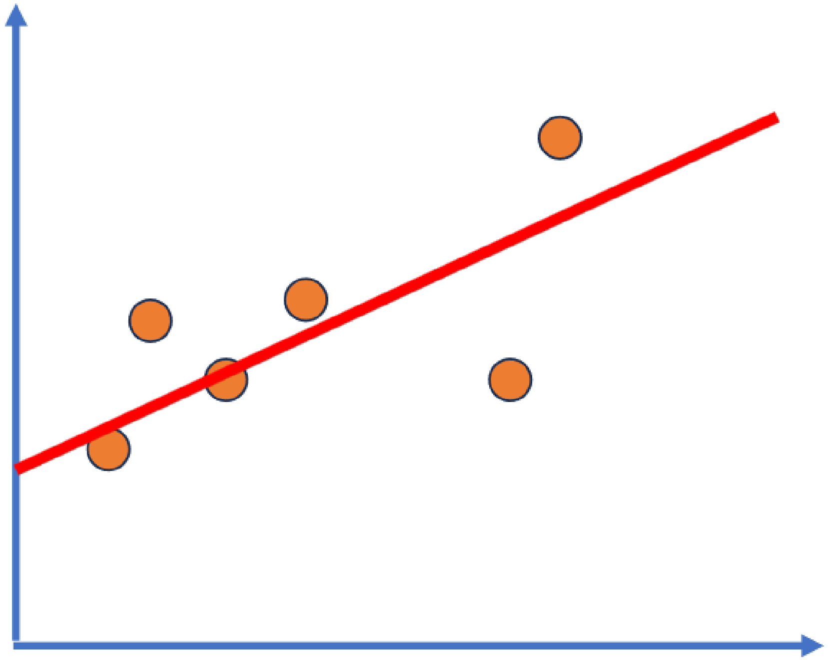

UNIVERSITY OF CALGARY

Loss Function

- How good our model is.
- Measures for each data point x we have, how close the model y was to the actual y
- Ex. Least Squared

$$\min \sum (y_{actual} - y_{model})^2$$

$$\min \sum \left(y_{actual} - (mx_{input} + b)\right)^2$$

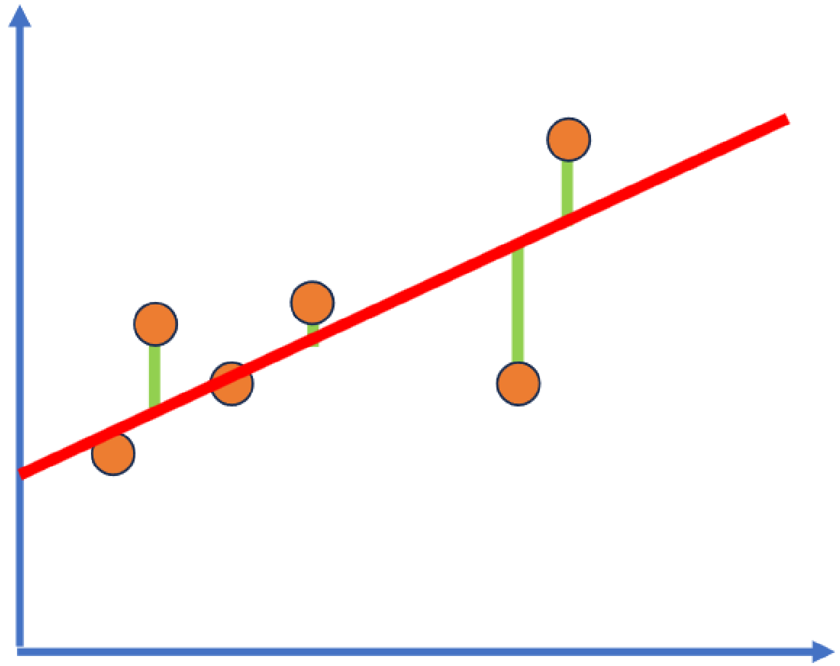

UNIVERSITY OF CALGARY

Challenges

- The four data sets have identical statistics
  - Same mean, and variance
- You might produce same model m/b for all 4 inputs
- But we might disagree that that model will actually predict future y points on the line for new x values as we desire

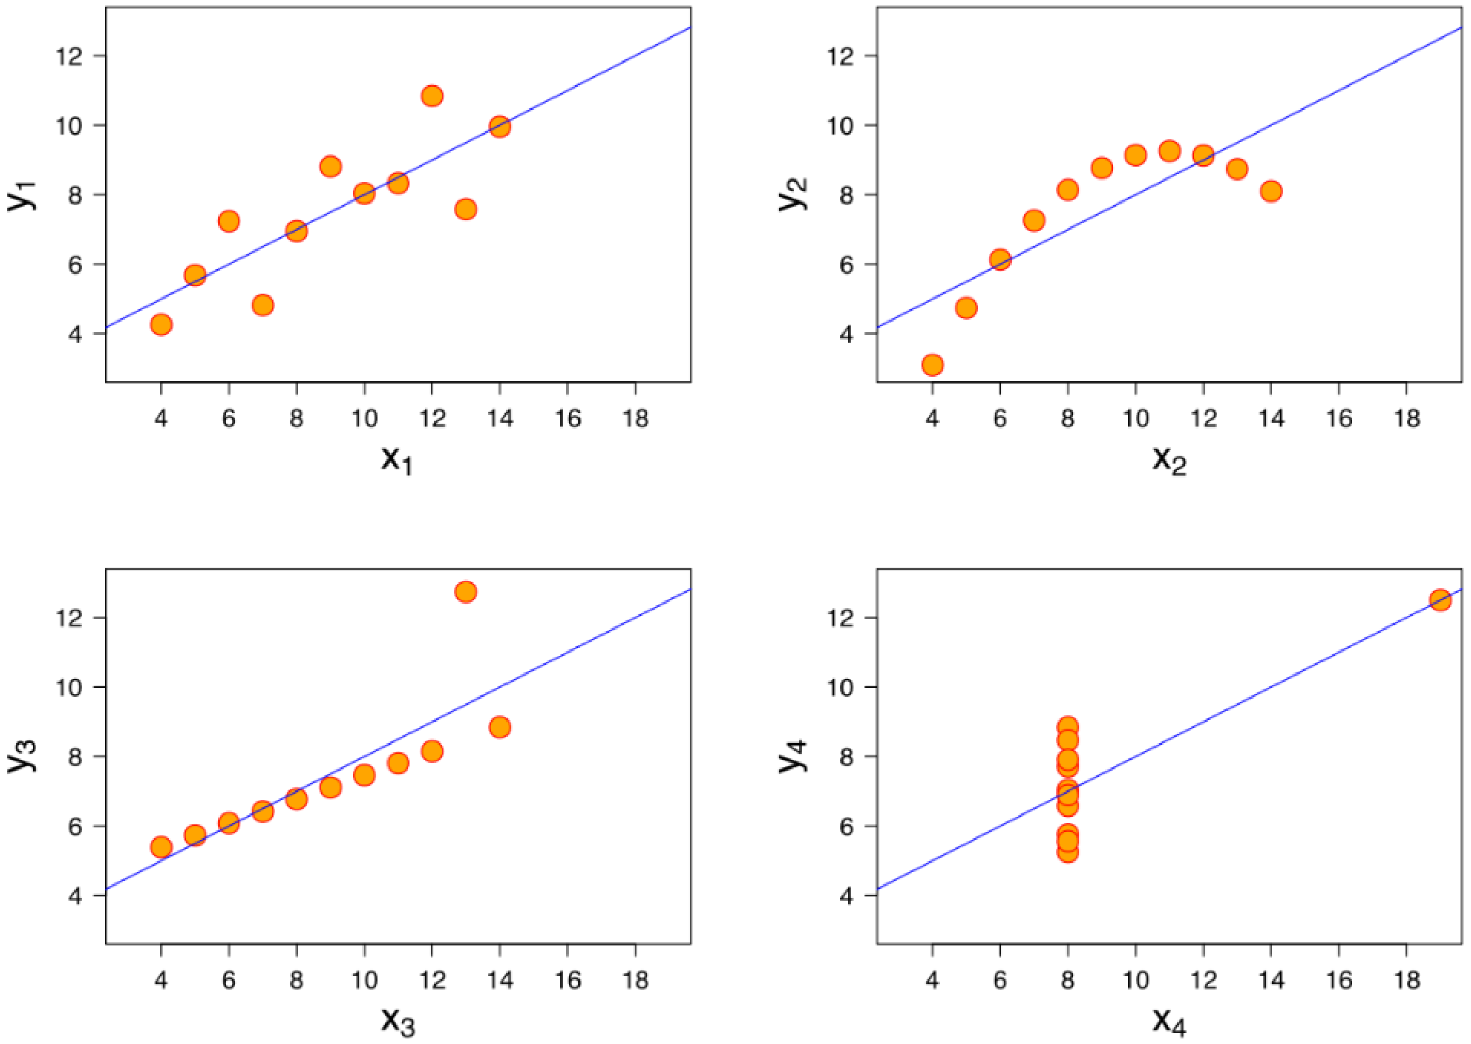

https://en.wikipedia.org/wiki/Anscombe%27s_quartet

Fitting and Regularization

- Under-fitting -> we don't have enough training data so model can't find a pattern and is bad at unknown future input
- Over-fitting -> the model matches data exactly, but is bad with unknown future input
  - Regularization is a technique that allows error in the model while fitting so that it can be more 'wrong' but also more general for future prediction

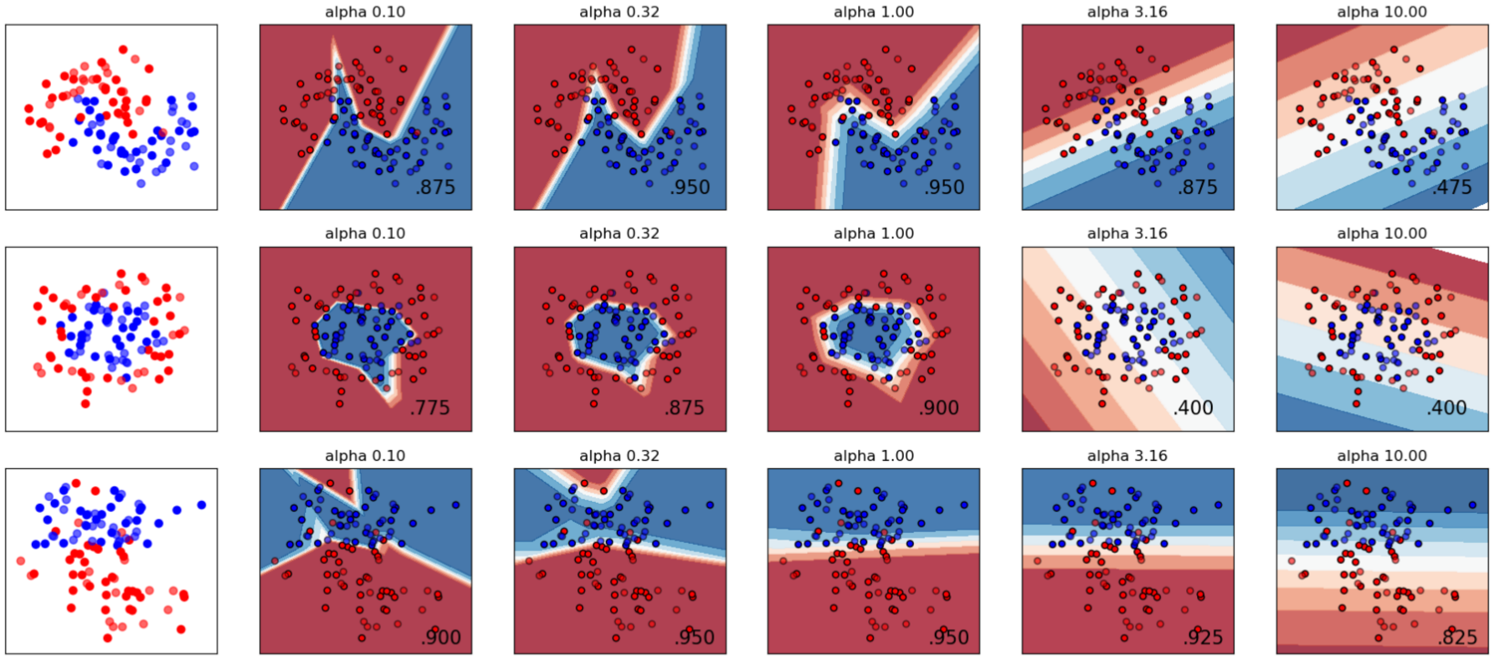

Neural networks often use drop-out for overfitting.
(Only use X% of your brain each time!)

Challenges

- Some data sets don't have natural linear…
- For some like the top-right it could be pre-processed into linear space
  - using polar coordinate system
- Some like bottom-left may have no reasonable re-mapping
- For shapes like the first, we could also make our function more complex by making the model non-linear (Neural networks often chosen here)
  - http://playground.tensorflow.org

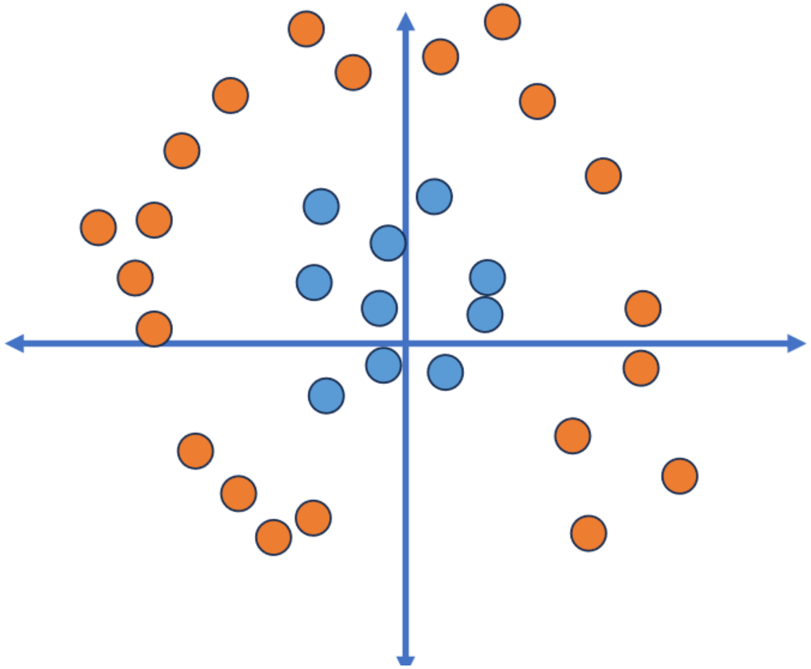

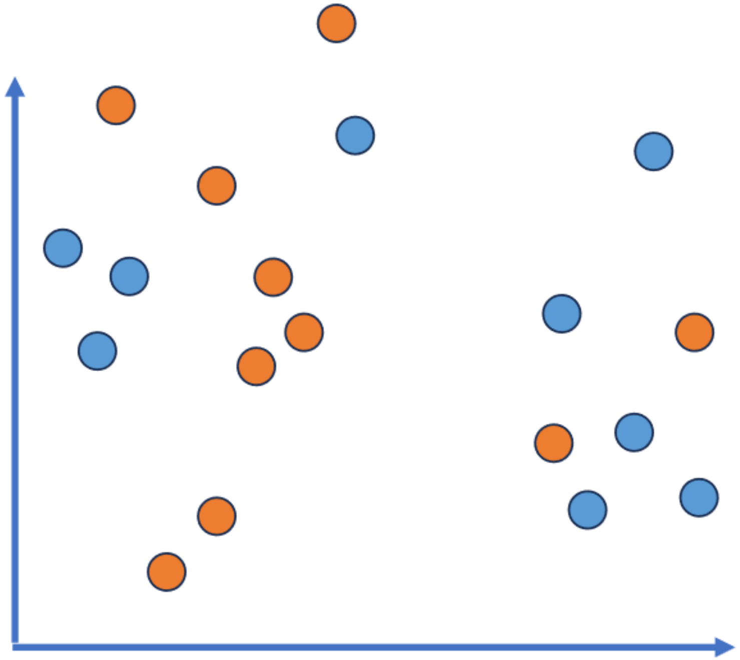

Onward to … neural networks

Jonathan Hudson, Ph.D.
jwhudson@ucalgary.ca
https://cspages.ucalgary.ca/~jwhudson/

UNIVERSITY OF CALGARY