Unsupervised Learning, Active Learning, Curriculum Learning, & Reinforcement Learning

Danna Gurari
University of Texas at Austin
Spring 2021

https://www.ischool.utexas.edu/~dannag/Courses/IntroToMachineLearning/CourseContent.html

Review

- Last week:
  - Machine Learning Algorithms that Discriminate
  - FAT (Fair, Accountable, & Transparent) Algorithms
  - Ethics in Machine Learning
  - Guest: Dr. Mehrnoosh Sameki from Microsoft
- Assignments (Canvas):
  - Project video due next week
  - Final project report and code due next week
- Questions?

Today’s Topics

- Machine Learning for Unlabeled Data
- Clustering
- Autoencoders
- Extra: Active Learning, Curriculum Learning, and Reinforcement Learning
- Course Summary

Today’s Topics

- Machine Learning for Unlabeled Data
- Clustering
- Autoencoders
- Extra: Active Learning, Curriculum Learning, and Reinforcement Learning
- Course Summary

How Have Machines Learned So Far in this Class?

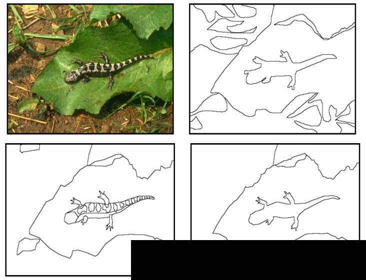

BS

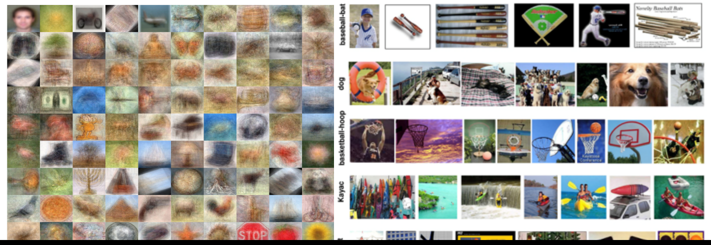

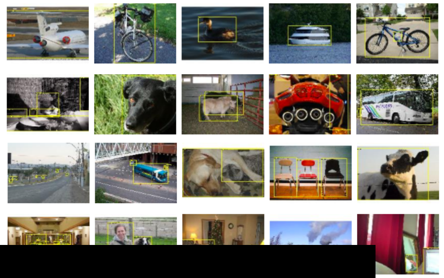

2)

Labe

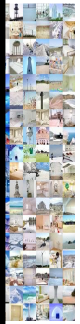

Large labelled datasets

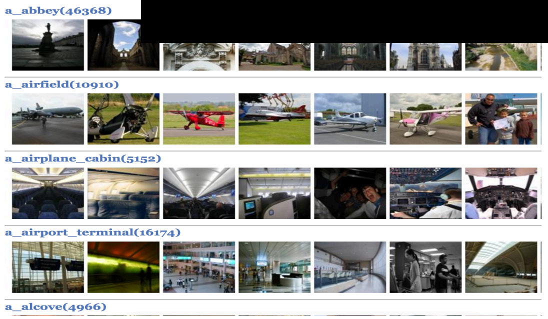

Places (2014)

MS COCO (2014)

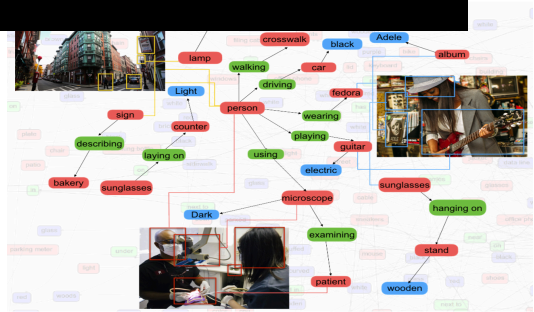

Visual Genome (2016)

Slide Credit: http://vision.cs.utexas.edu/slides/mit-ibm-august2018.pdf

Why Not Rely On Large Labelled Datasets?

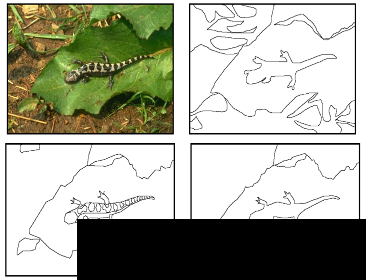

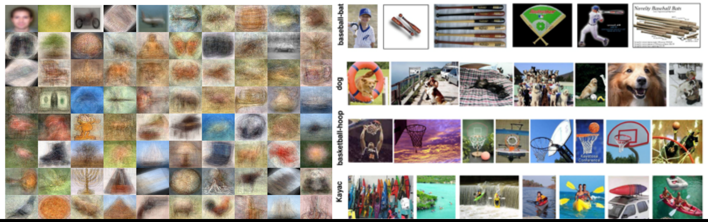

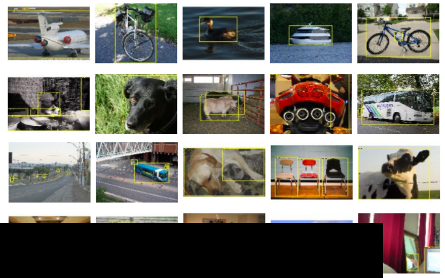

)

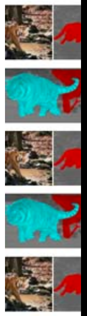

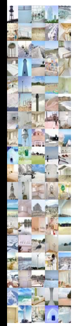

- Expensive
- Relatively Slow
- Disconnect from Human Learning

La

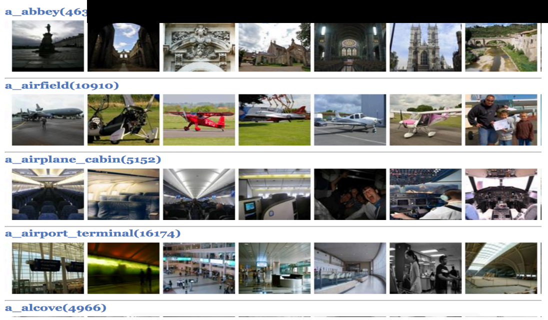

Places (2014)

MS COCO (2014)

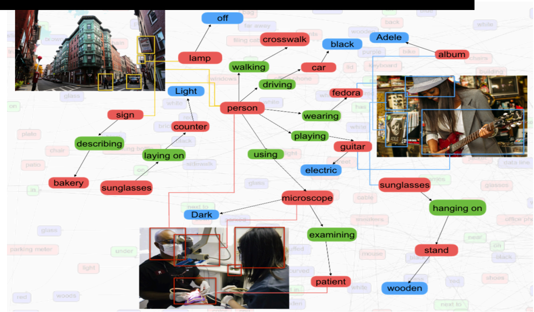

Visual Genome (2016)

Slide Credit: http://vision.cs.utexas.edu/slides/mit-ibm-august2018.pdf

Intuition: How Do Humans Learn?

With Supervision

No Supervision

https://pixabay.com/en/toddler-learning-book-child-423227/

https://www.maxpixel.net/Father-Child-Family-Dad-Baby-Daughter-3046495

Goal: Learn from Experience To Organize Data

https://pixabay.com/en/toddler-learning-book-child-423227/

https://www.maxpixel.net/Father-Child-Family-Dad-Baby-Daughter-3046495

Real-World Applications: Customer Segmentation

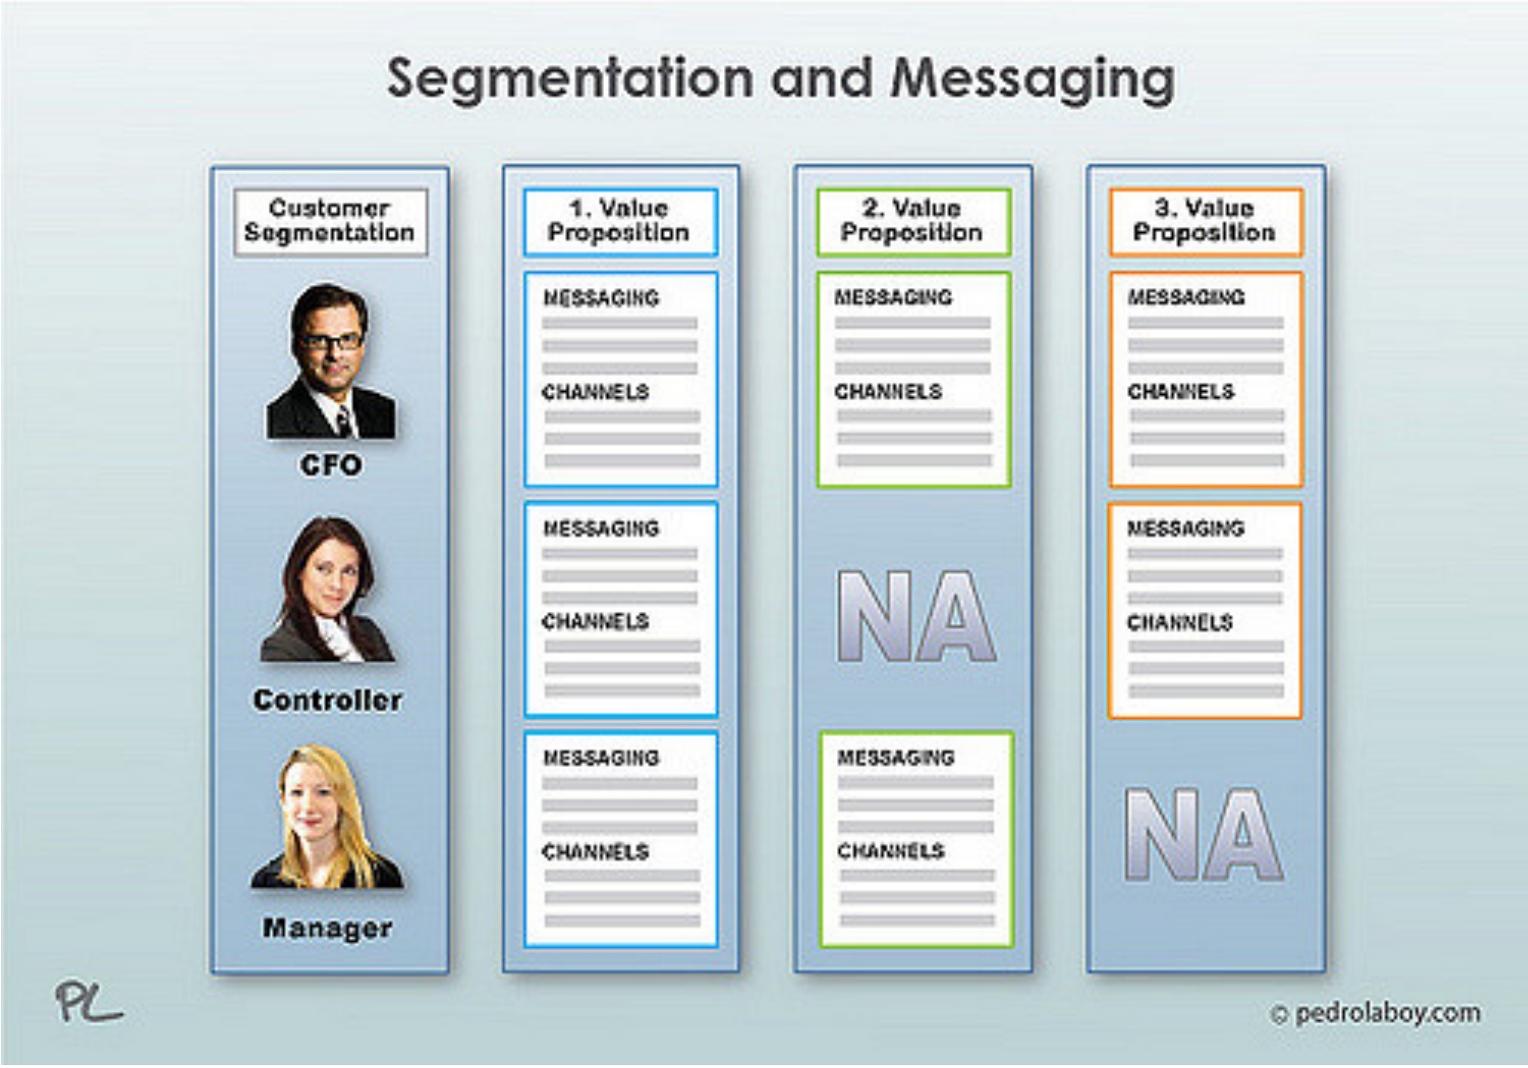

https://www.flickr.com/photos/42565140@N04/3923873188/

Real-World Applications: Recommendations

https://medium.com/@navdeepsingh_2336/scala-machine-learning-projects-recommendation-systems-d41d9eebbb06

Real-World Applications: Social Network Analysis

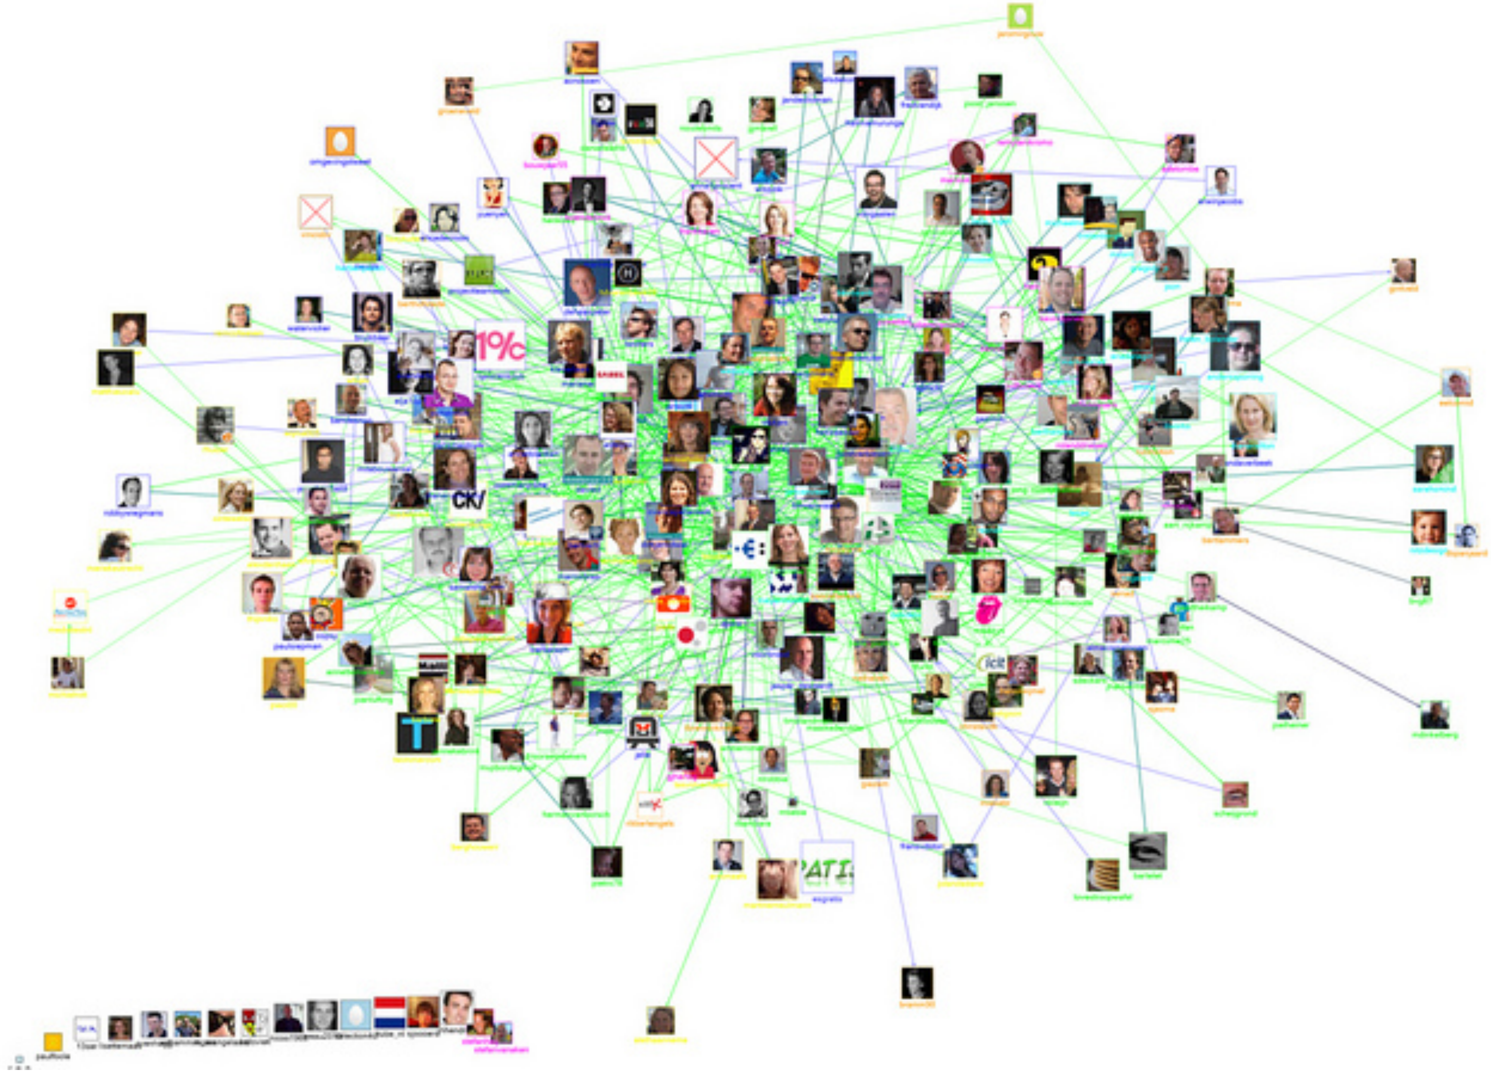

https://www.flickr.com/photos/marc_smith/5529685600

Real-World Applications: Fraud Detection

https://www.lejeune.marines.mil/News/Article/511667/protect-yourself-from-credit-card-fraud/

Today’s Topics

- Machine Learning for Unlabeled Data
- Clustering
- Autoencoders
- Extra: Active Learning, Curriculum Learning, and Reinforcement Learning
- Course Summary

Clustering

A.

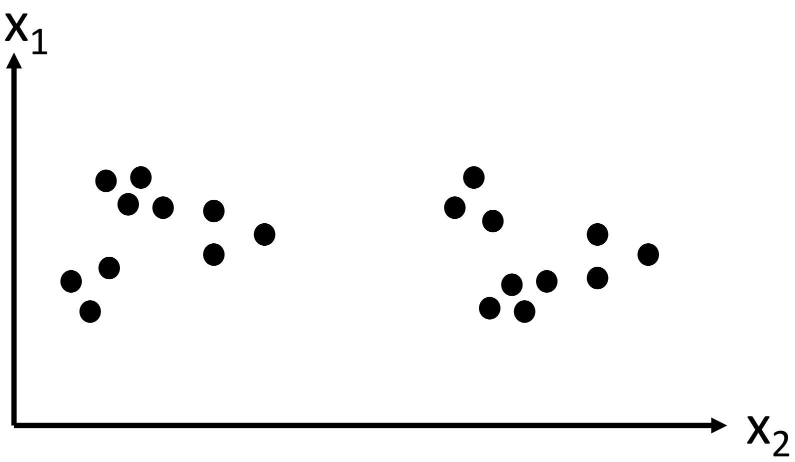

B.

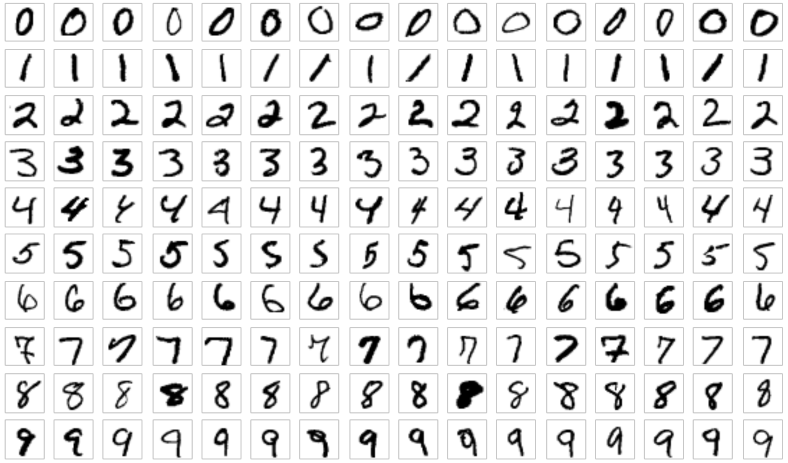

C.

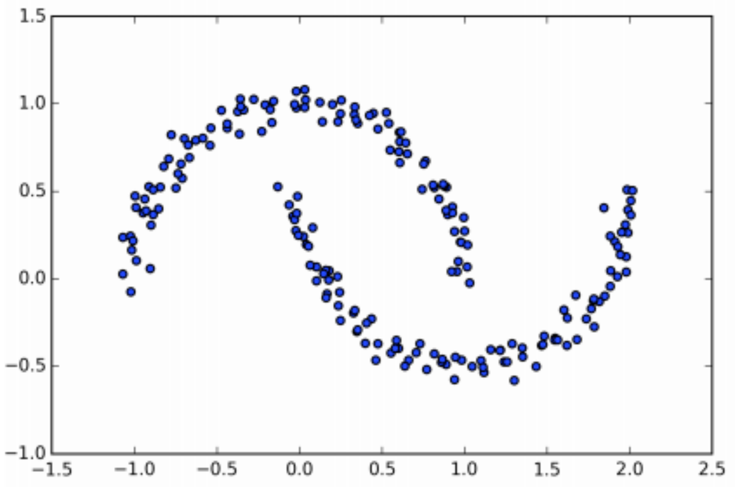

Find groupings such that entities in a group will be similar to each another and different from the entities in other groups.

Raschka and Mirjalili; Python Machine Learning

Clustering: Key Questions

A.

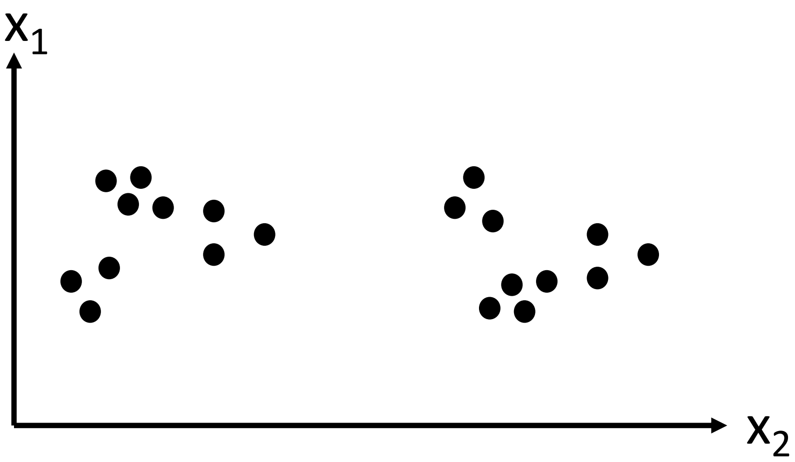

B.

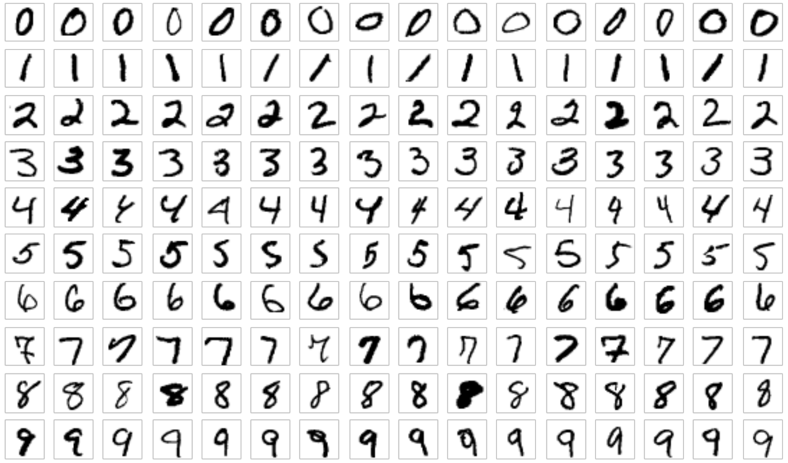

C.

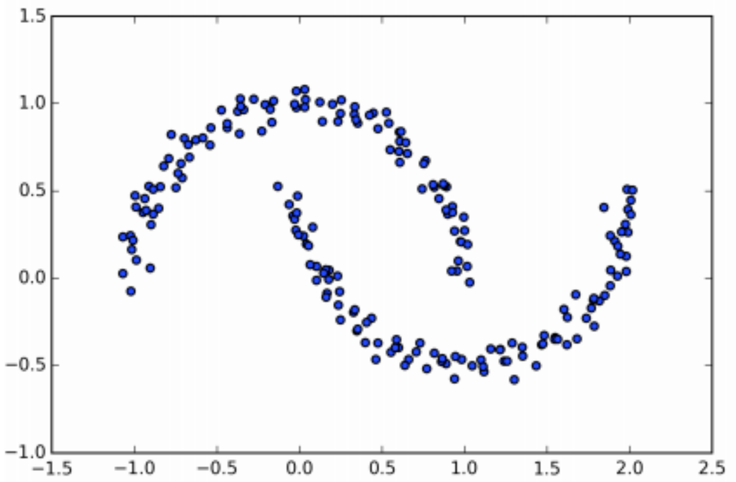

- How many data clusters to create?
- What "algorithm" to use to partition the data?

Raschka and Mirjalili; Python Machine Learning

Breakout Discussion

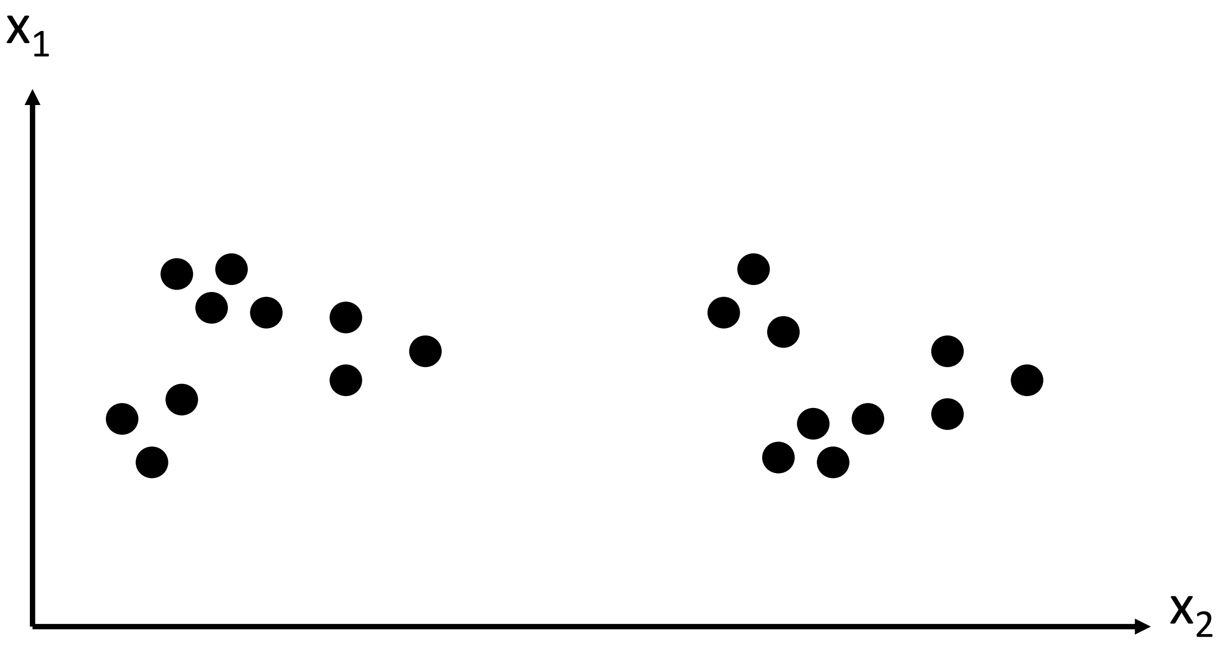

- How many data clusters to create?
- What “algorithm” to use to partition the data?

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

How Many Clusters?

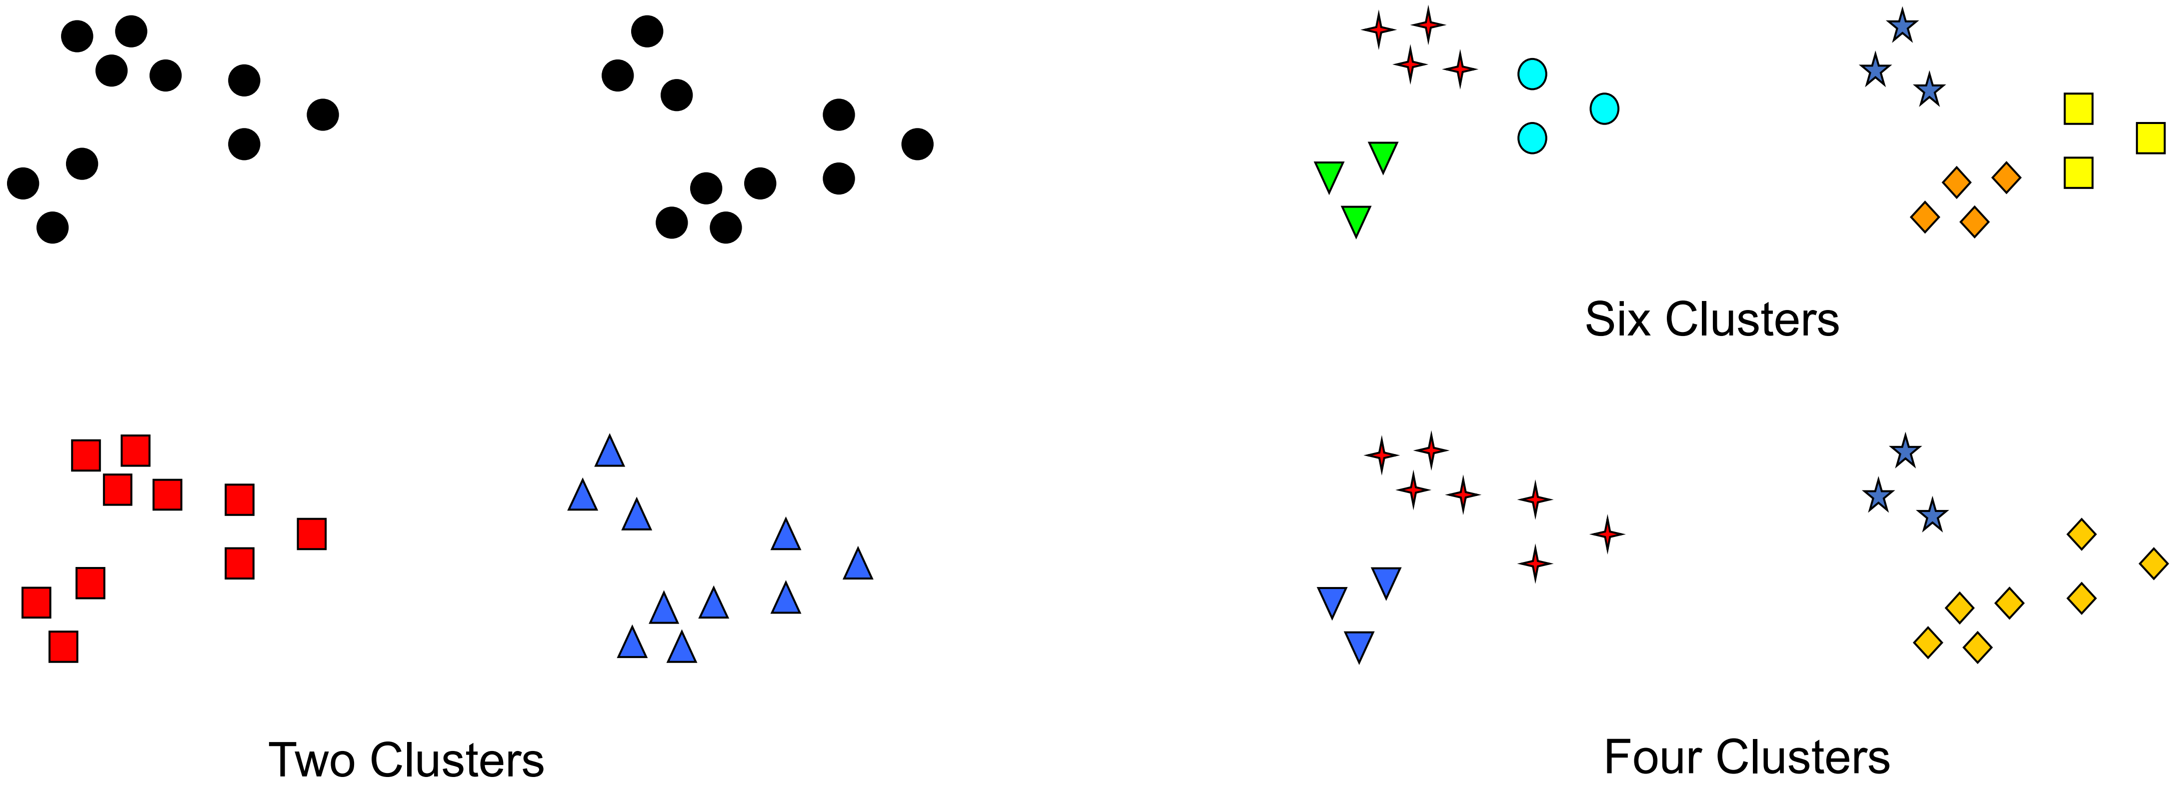

Number of clusters can be ambiguous.

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Types of Clustering

- Partitional Clustering
  - A division of data objects into non-overlapping subsets (clusters) such that each data object is in exactly one subset
- Hierarchical clustering
  - A set of nested clusters organized as a hierarchical tree

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

K-Means Clustering

1: Select K points as the initial centroids.
2: **repeat**
3: Form K clusters by assigning all points to the closest centroid.
4: Recompute the centroid of each cluster.
5: **until** The centroids don't change

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

K-Means Clustering

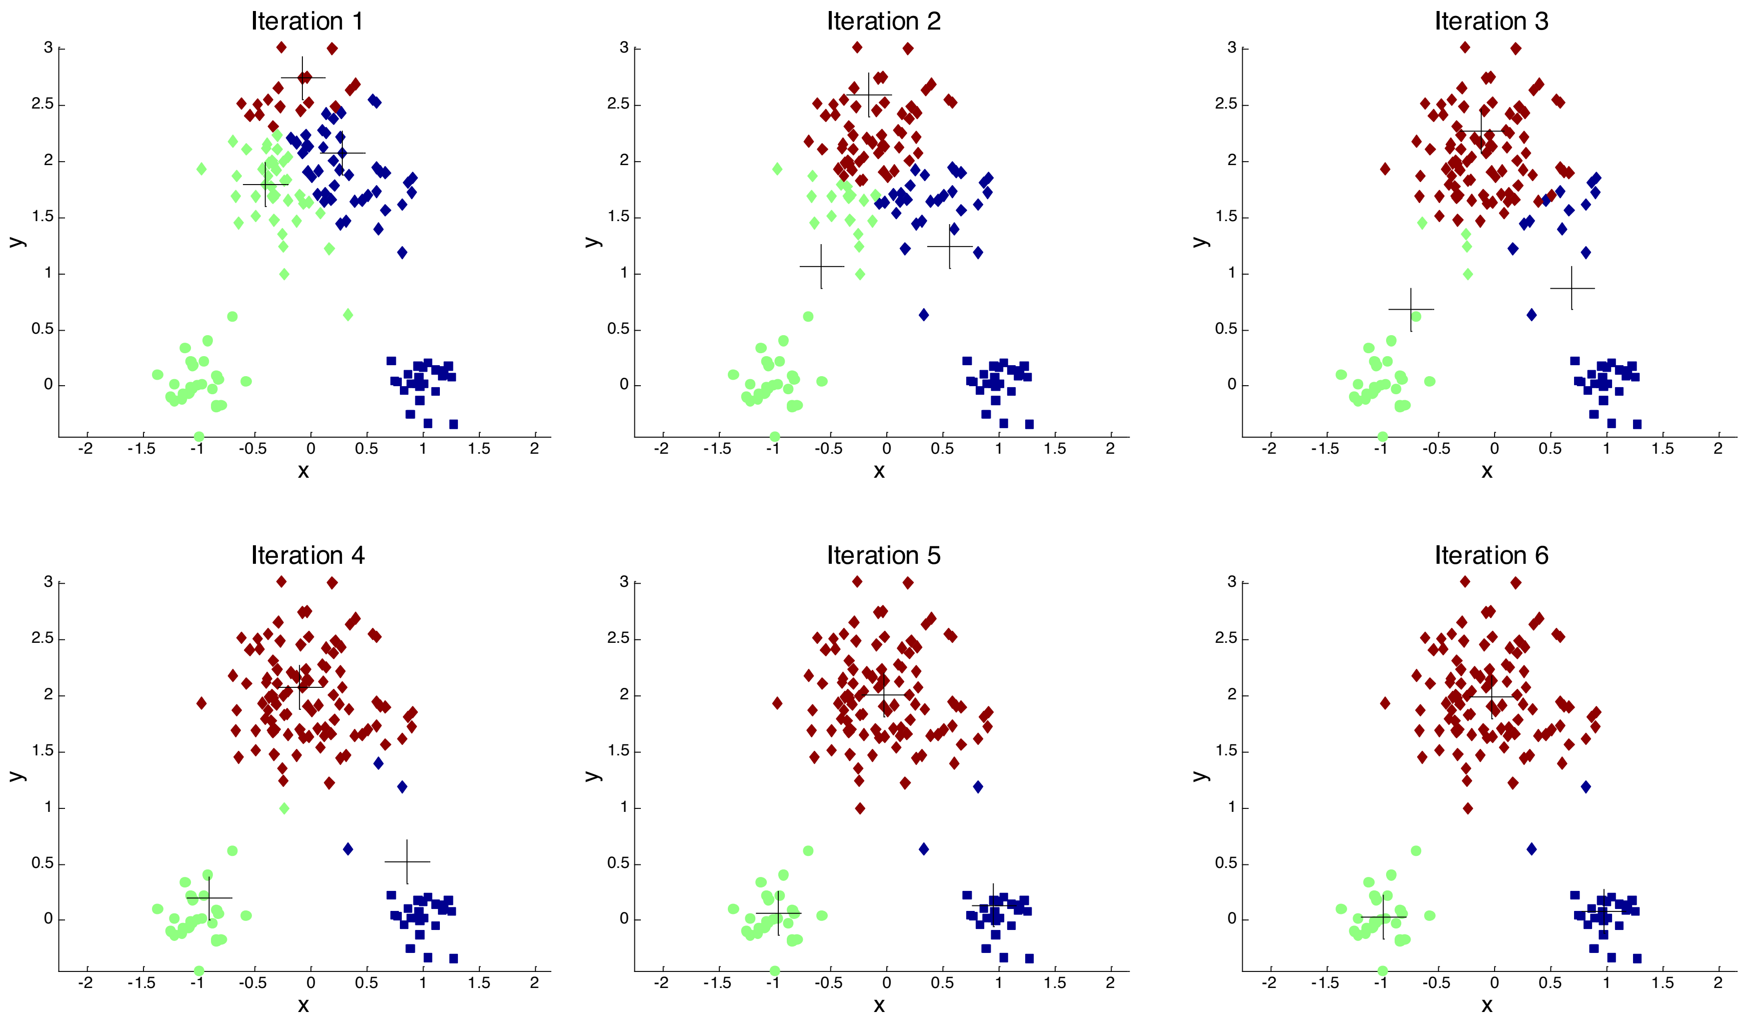

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

K-Means Clustering: Weaknesses?

- Sensitive to initial centroids: different outcomes for same data

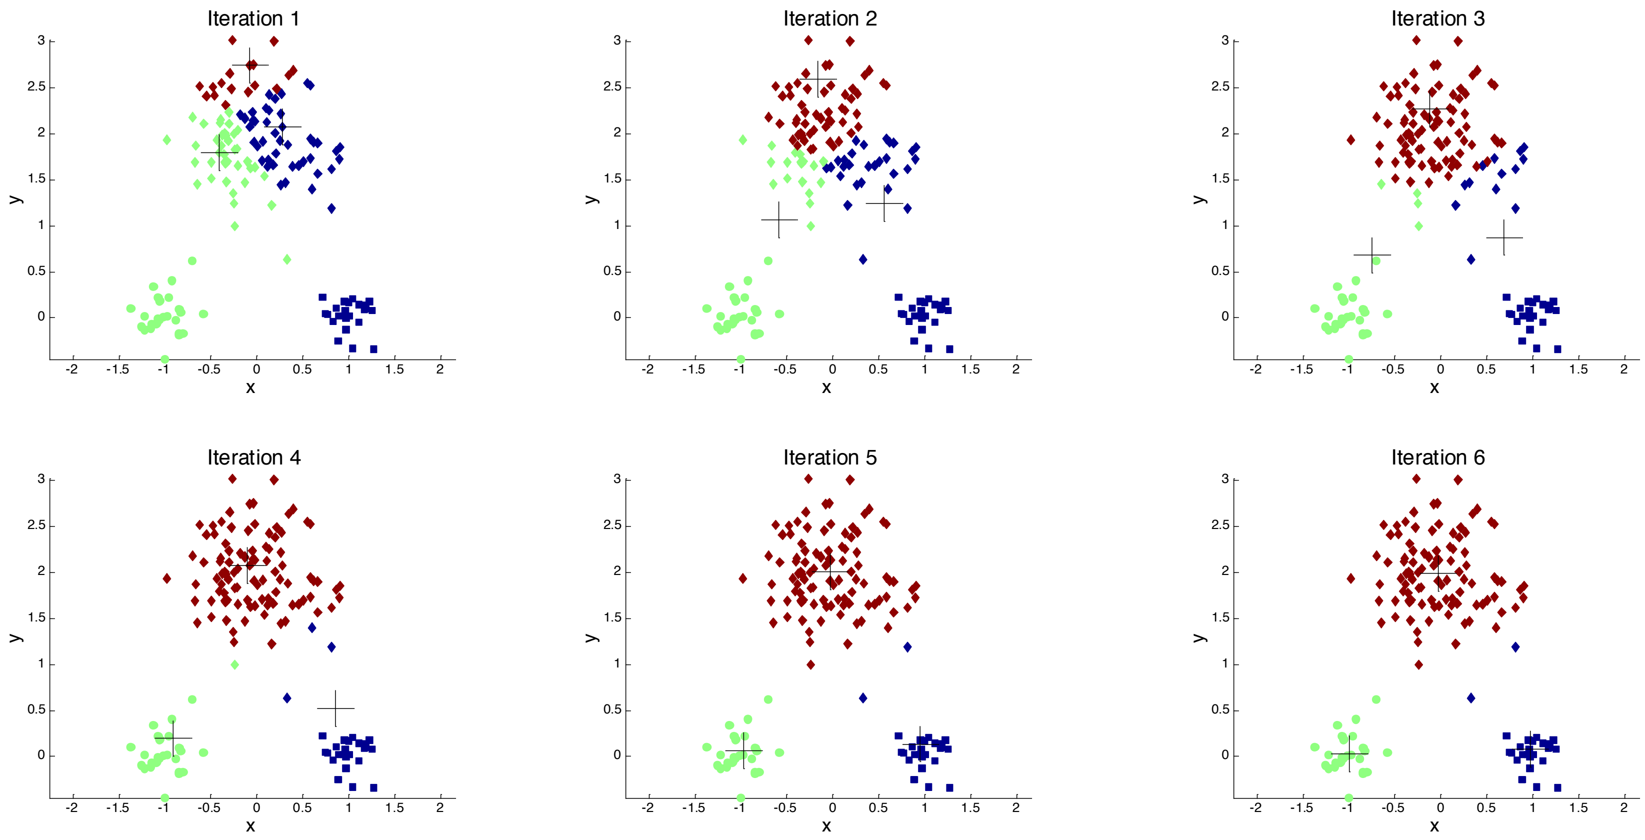

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

K-Means Clustering: Weaknesses?

- Not robust when clusters have different sizes:

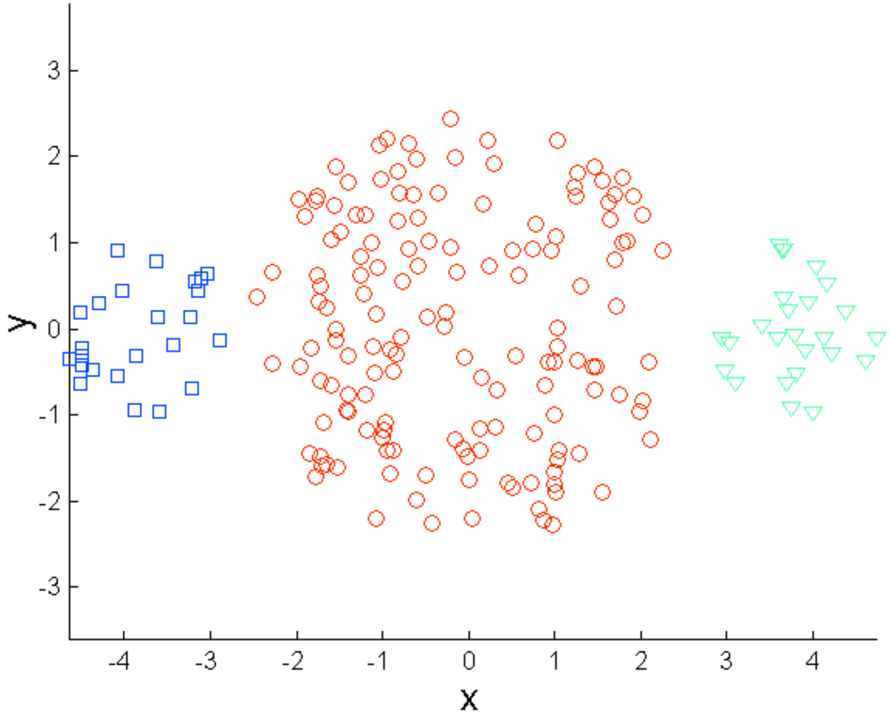

Original Points

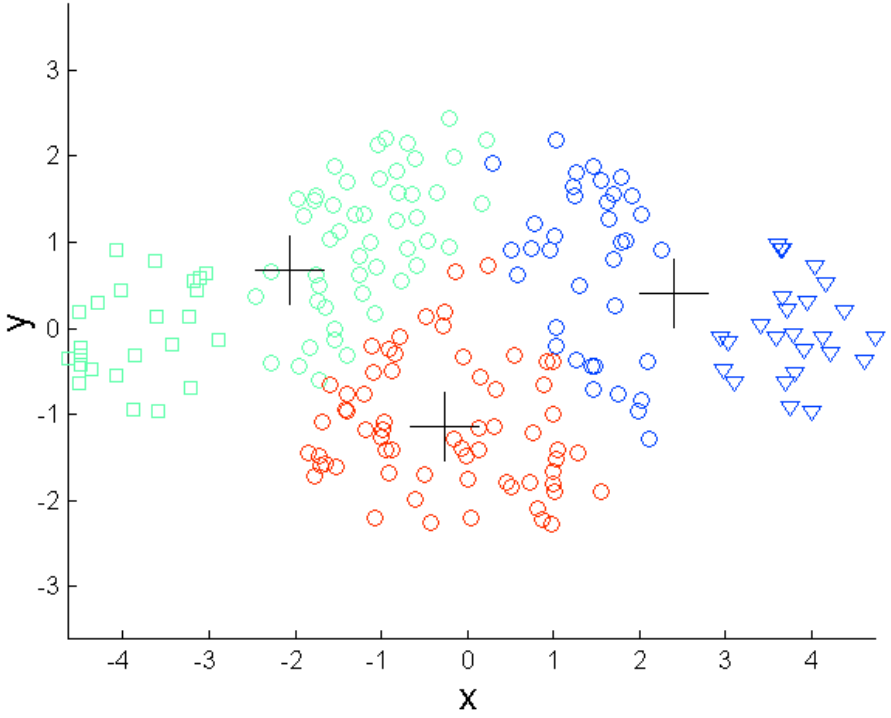

K-means (3 Clusters)

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

K-Means Clustering: Weaknesses?

- Not robust when clusters have different densities:

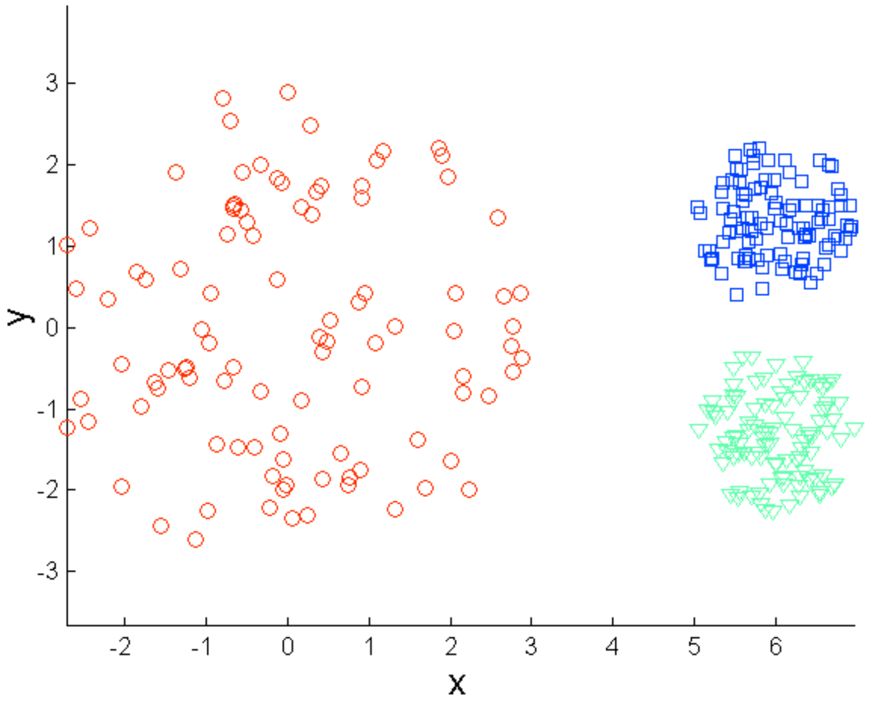

Original Points

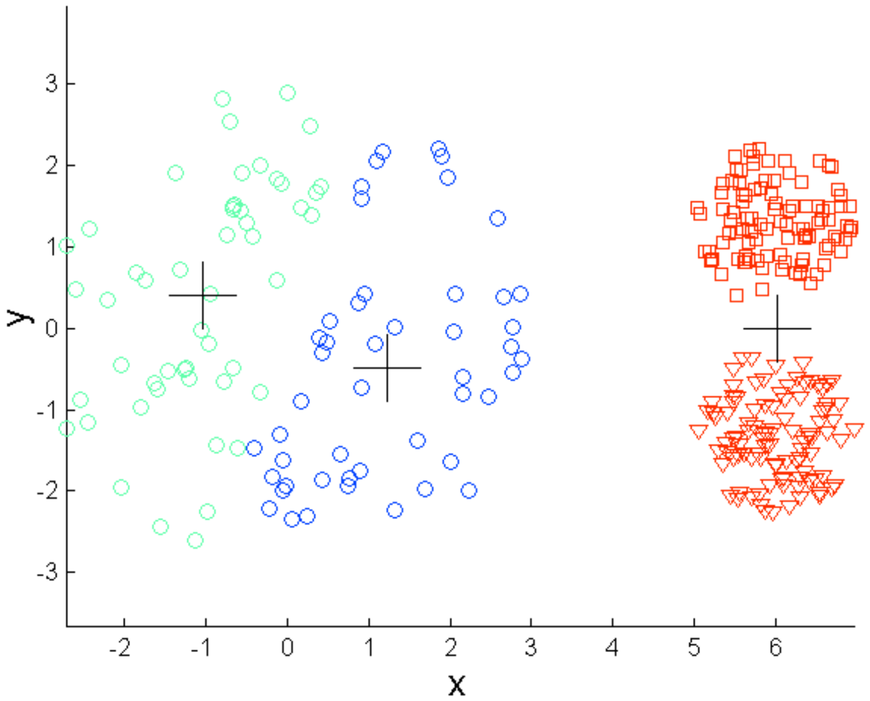

K-means (3 Clusters)

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

K-Means Clustering: Weaknesses?

- Not robust when clusters have different globular shapes:

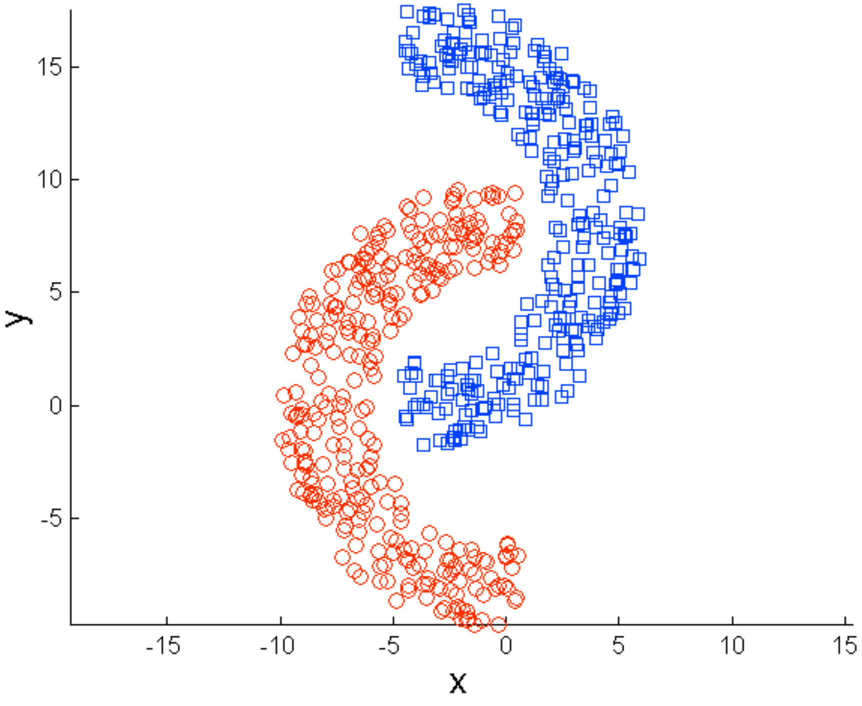

Original Points

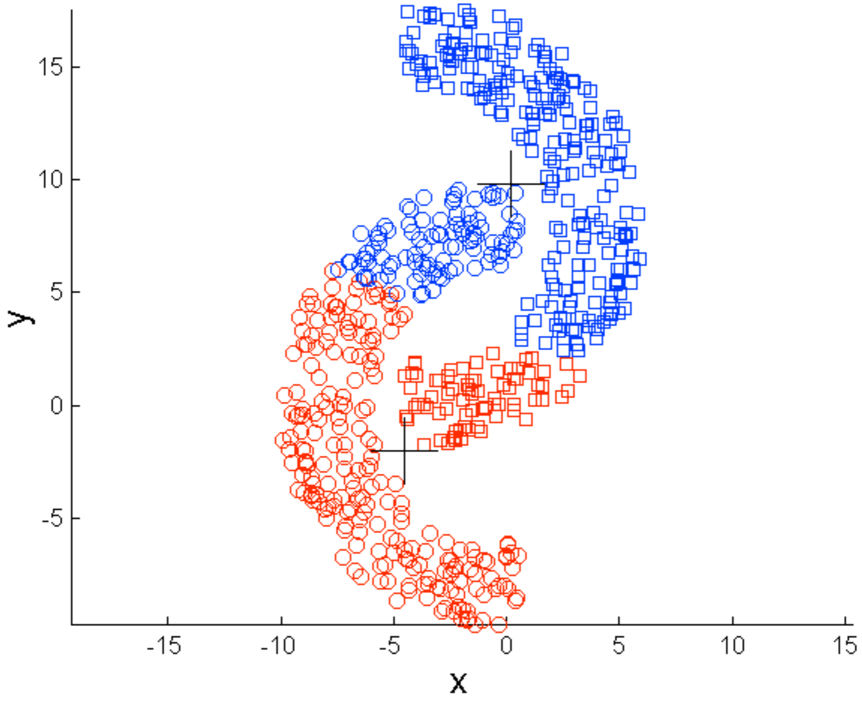

K-means (3 Clusters)

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Hierarchical Clustering

- Set of nested clusters organized in hierarchical tree by merging/splitting
- Dendrogram visualization: shows sequence of merges/splits

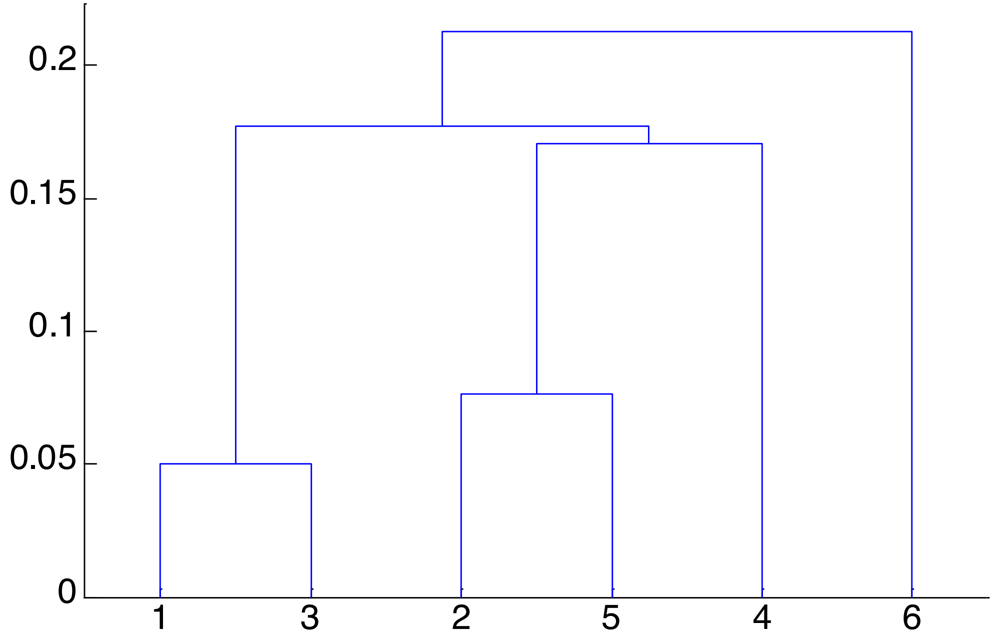

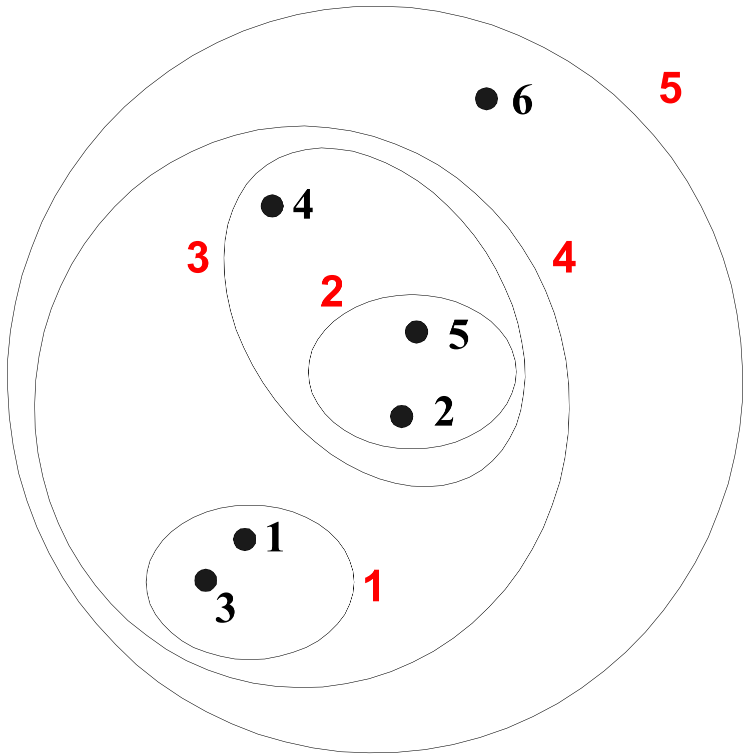

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Hierarchical Clustering: Two Main Approaches

- Agglomerative:
  - Start with points as individual clusters
  - At each step, merge closest pair of clusters until only one cluster (or k clusters) left

- Divisive:
  - Start with one, all-inclusive cluster
  - At each step, split a cluster until each cluster contains an individual point (or there are k clusters)

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Agglomerative Clustering: First Step

- Start with clusters of individual points and a proximity matrix

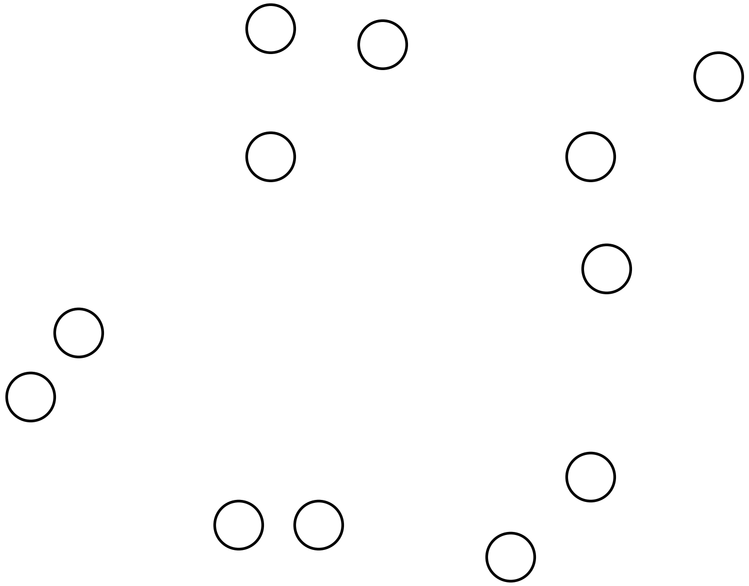

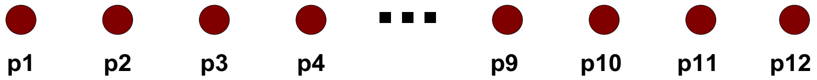

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Agglomerative Clustering: Intermediate Step

- Start with clusters of individual points

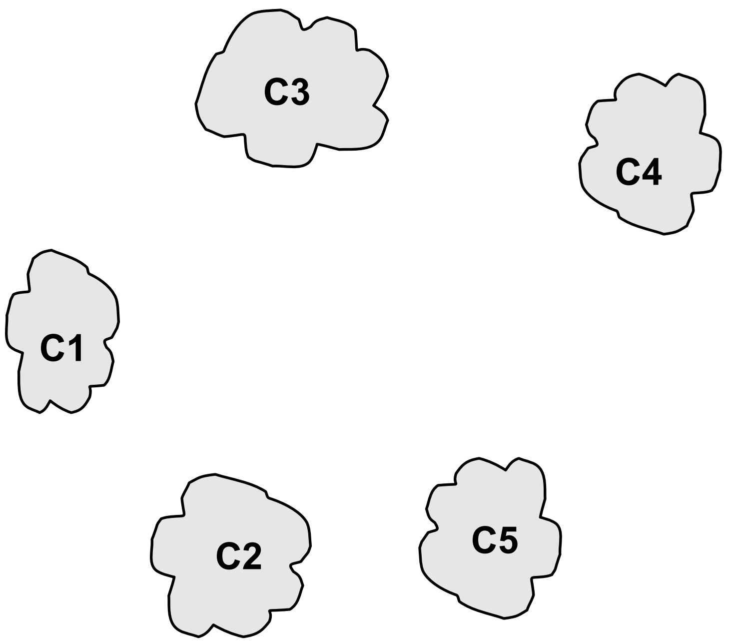

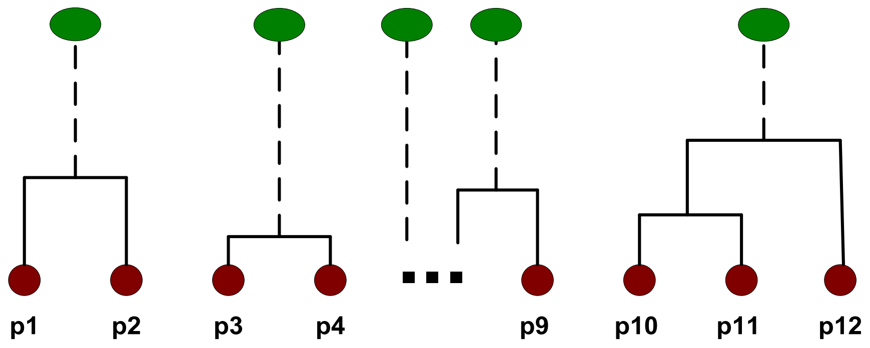

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Agglomerative Clustering: Intermediate Step

- After several merging steps, we have some clusters

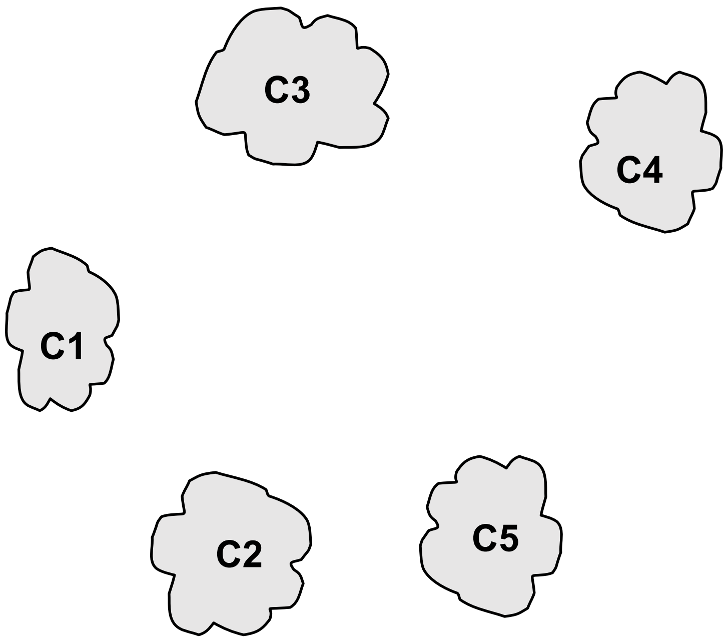

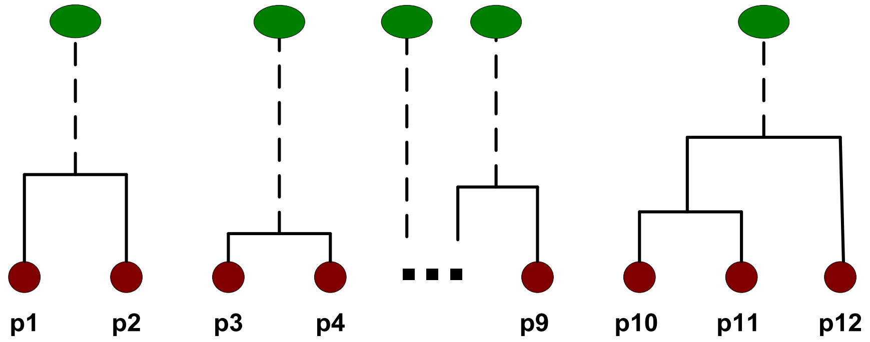

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Agglomerative Clustering: Intermediate Step

- Merge two closest clusters (C2 and C5)

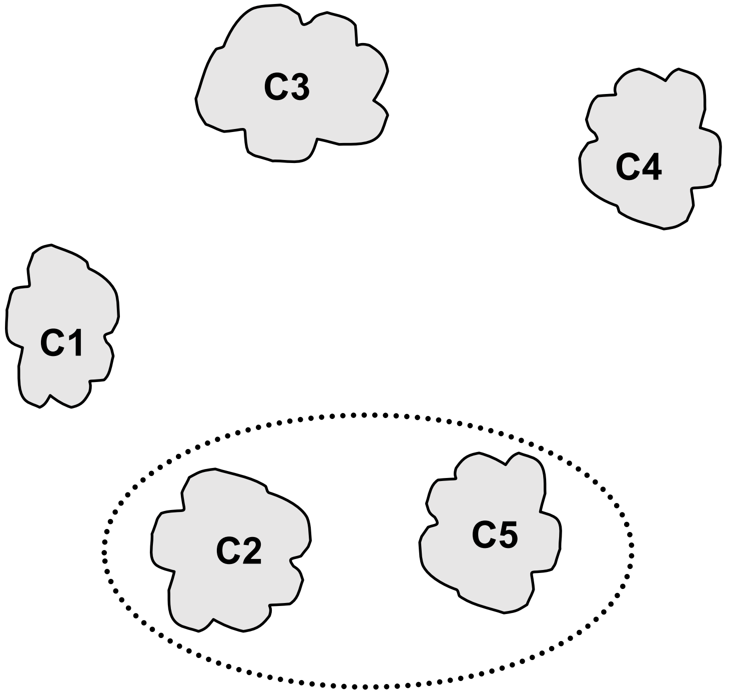

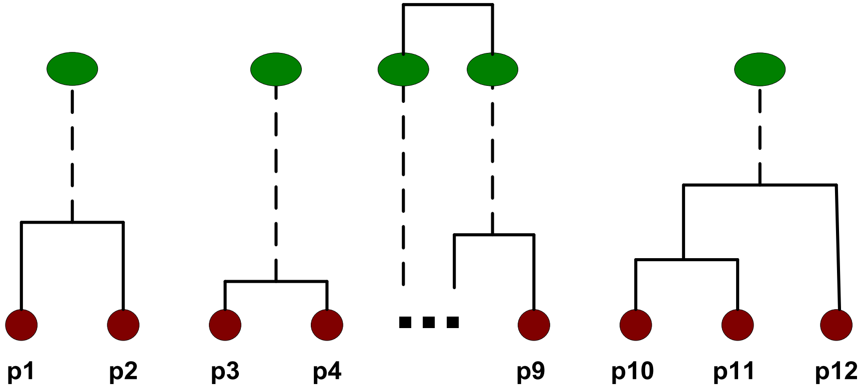

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

How to Measure Inter-Cluster Distance?

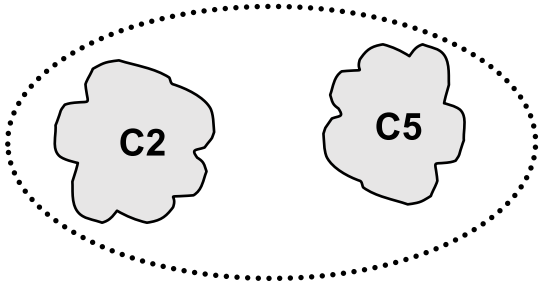

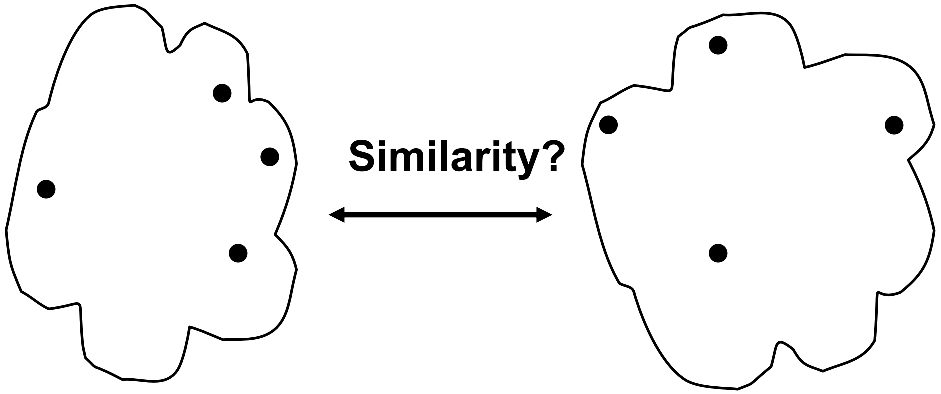

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

How to Measure Inter-Cluster Distance?

- Minimum distance

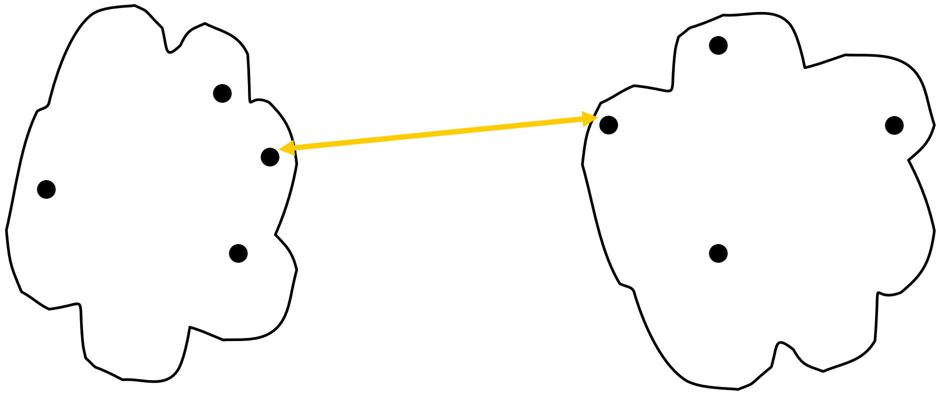

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Minimum Distance: Strengths/Weaknesses?

- Can handle non-elliptical shapes:

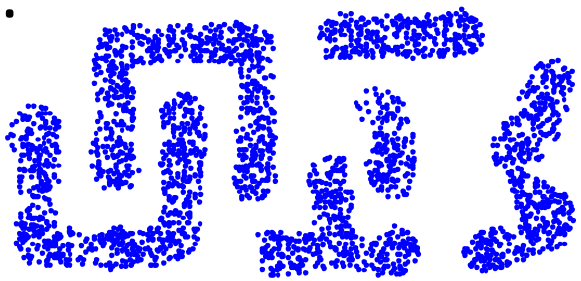

Original Points

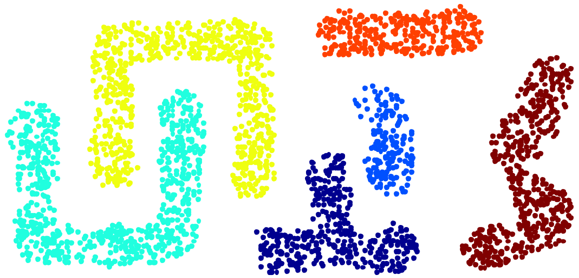

Six Clusters

- Sensitive to noise and outliers:

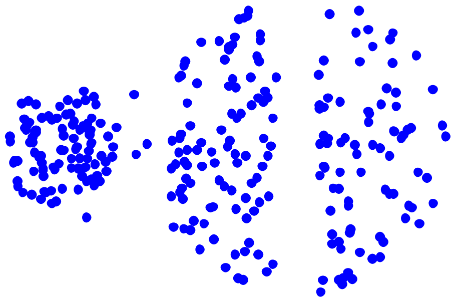

Original Points

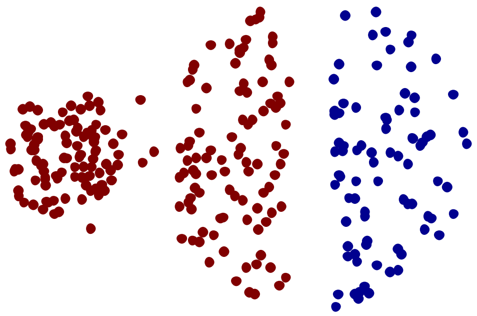

Two Clusters

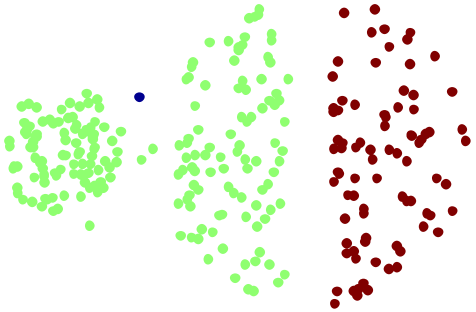

Three Clusters

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

How to Measure Inter-Cluster Distance?

- Minimum distance
- Maximum distance

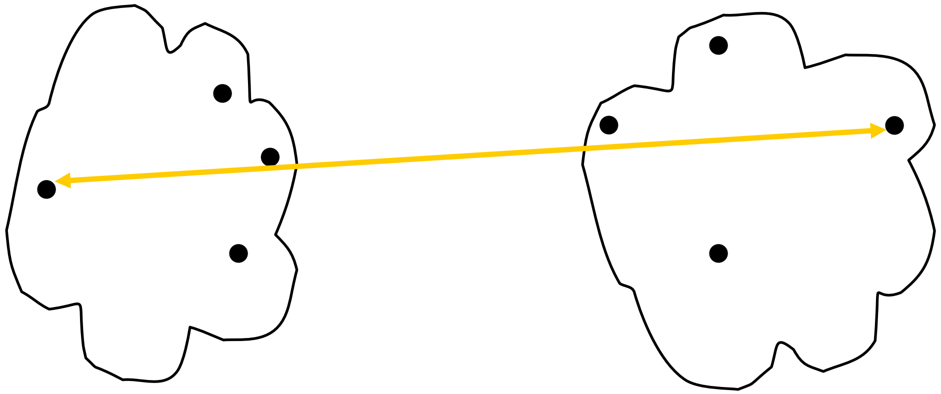

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Maximum Distance: Strengths/Weaknesses?

- Less susceptible to noise and outliers:

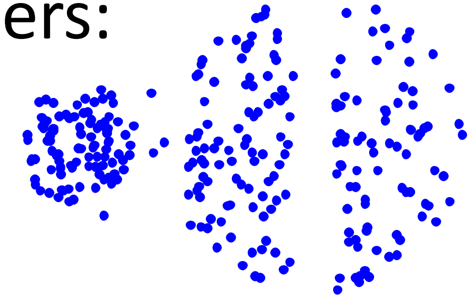

Original Points

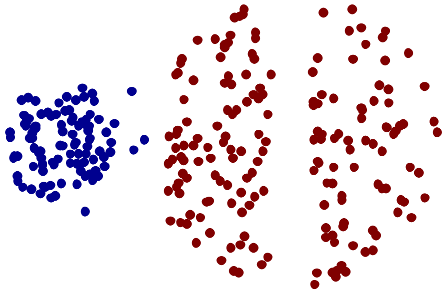

Two Clusters

- Tends to break large clusters:

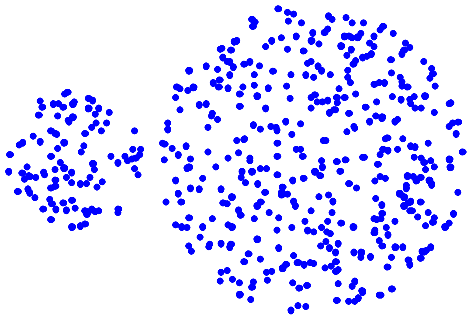

Original Points

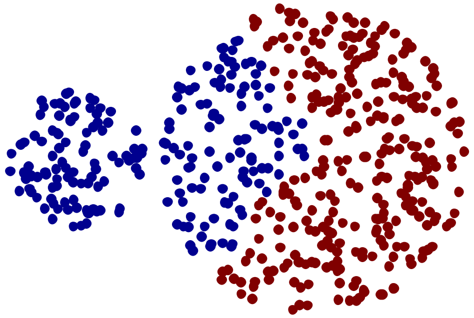

Two Clusters

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

How to Measure Inter-Cluster Distance?

- Minimum distance
- Maximum distance
- Group average

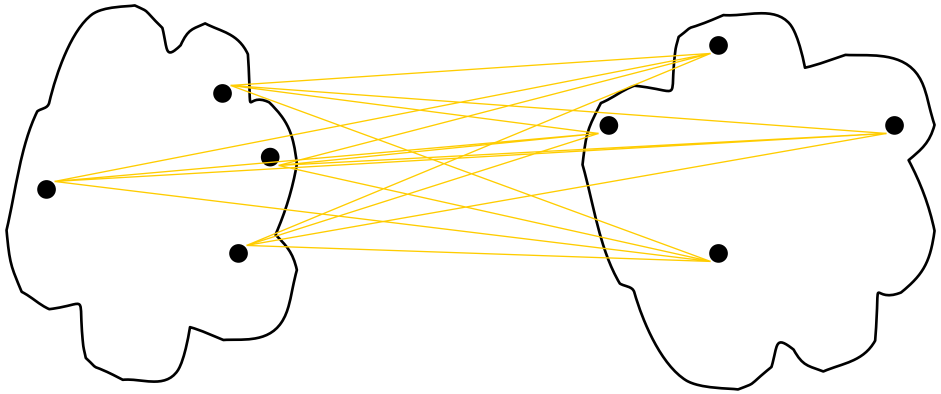

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

How to Measure Inter-Cluster Distance?

- Minimum distance
- Maximum distance
- Group average
- Distance Between Centroids

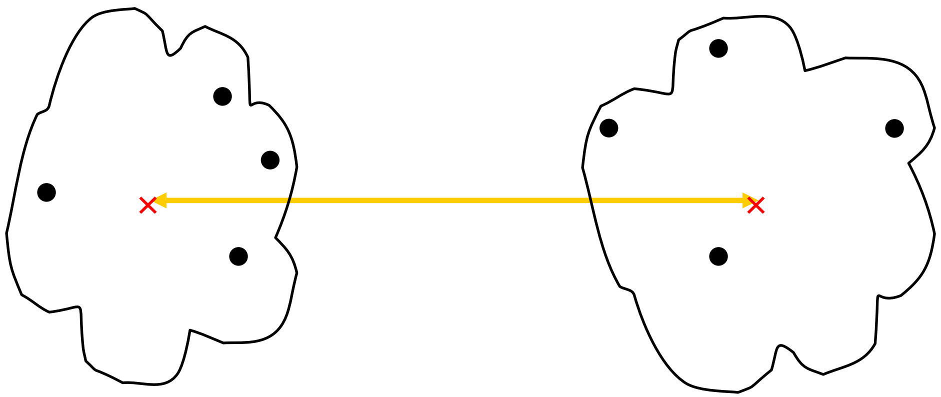

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Hierarchical Clustering: Strengths?

- Any number of clusters can be obtained by 'cutting' the dendrogram at the proper level

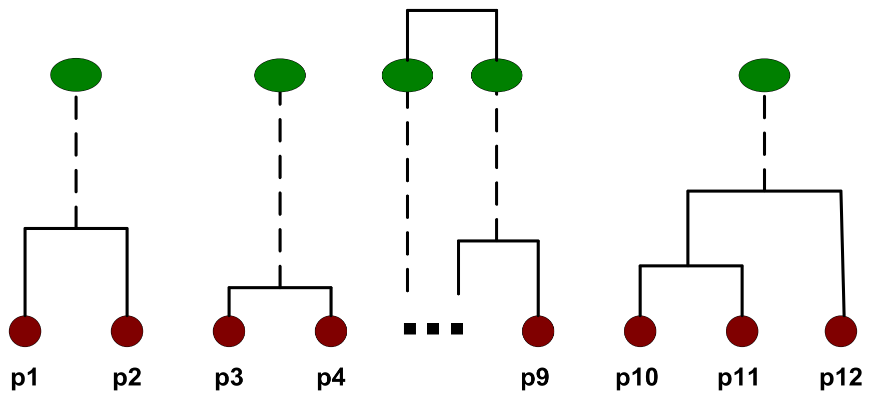

- They may correspond to meaningful taxonomies
  - Example in biological sciences (e.g., animal kingdom, phylogeny reconstruction, …)

Slide adapted from: https://www-users.cs.umn.edu/~kumar001/dmbook/slides/chap7_basic_cluster_analysis.pdf

Today’s Topics

- Machine Learning for Unlabeled Data
- Clustering
- Autoencoders
- Extra: Active Learning, Curriculum Learning, and Reinforcement Learning
- Course Summary

Autoencoder Architecture

- Learn to copy the input to the output

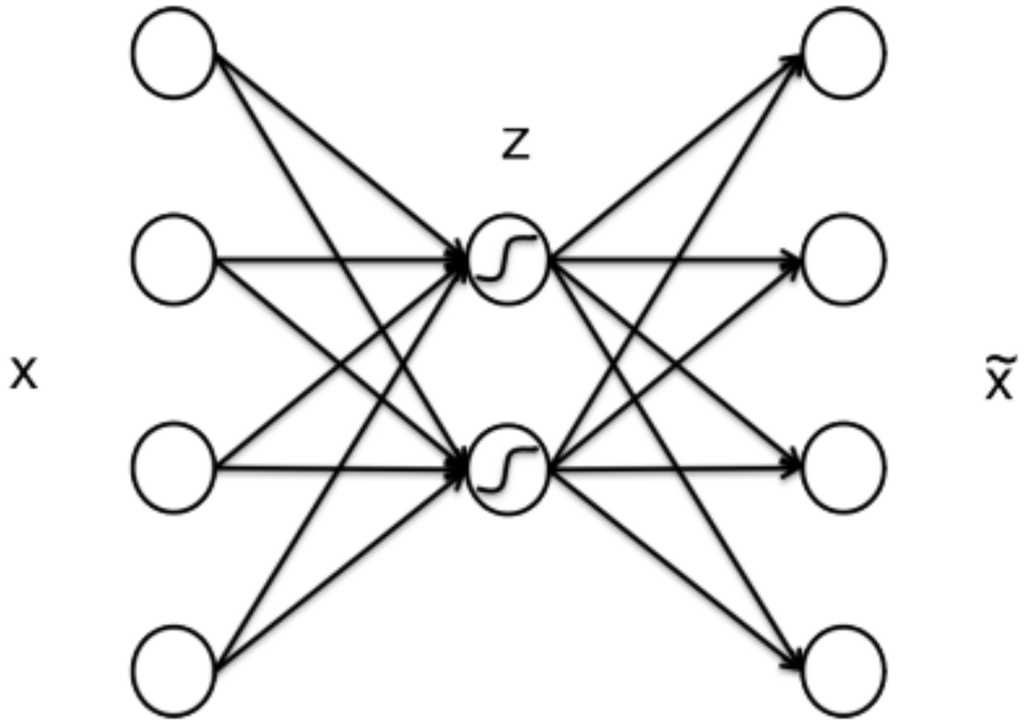

Figure Credit: https://lazyprogrammer.me/a-tutorial-on-autoencoders/

Autoencoder Architecture

- Consists of two parts:
  - **Encoder**: compresses inputs to an internal representation
  - **Decoder**: tries to reconstruct the input from the internal representation

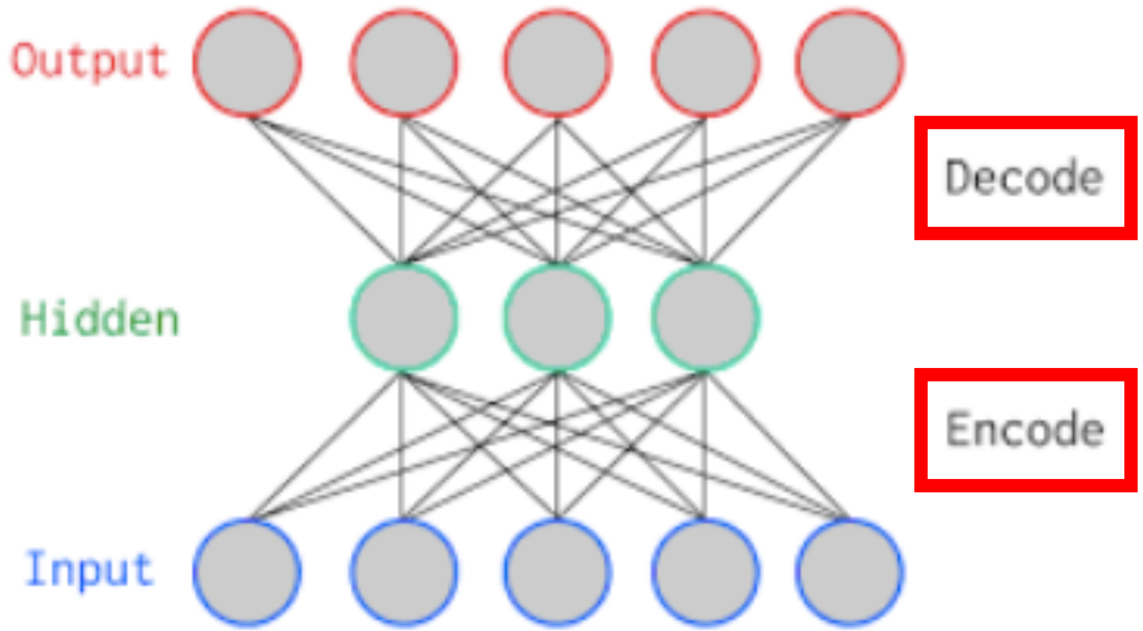

Figure Credit: https://www.datacamp.com/community/tutorials/autoencoder-keras-tutorial

Autoencoder Architecture

- Given this input 620 x 426 image (264,120 pixels):

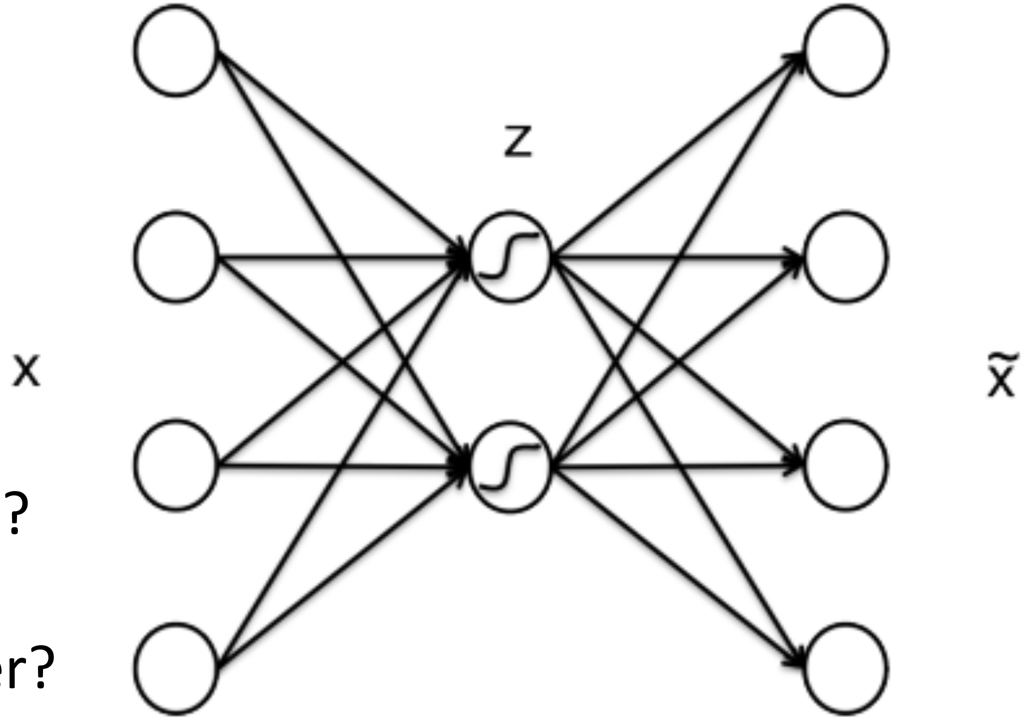

  - What would a perfect autoencoder predict?
    - Itself
  - What number of nodes are in the final layer?
    - 264,120

Figure Credit: https://lazyprogrammer.me/a-tutorial-on-autoencoders/

Autoencoder Training

How do you train a neural network?

Autoencoder Training

Repeat until stopping criterion met:

1. **Forward pass**: propagate training data through network to make prediction
2. **Backward pass**: using predicted output, calculate error gradients backward
3. Update each weight using calculated gradients

Autoencoder

What are useful applications for autoencoders?

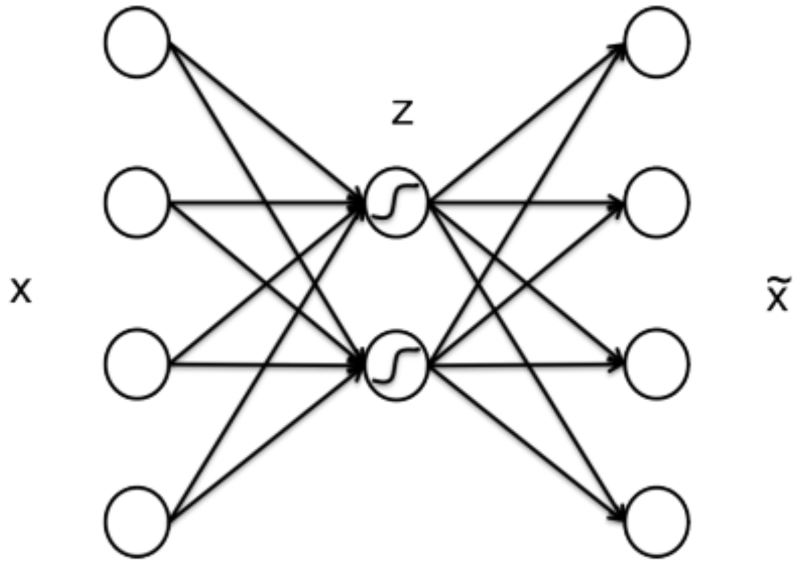

Figure Credit: https://lazyprogrammer.me/a-tutorial-on-autoencoders/

Autoencoders: Dimensionality Reduction

- Intuition: which number sequence is easier to remember?
  - **A:** 30, 27, 22, 11, 6, 8, 7, 2
  - **B:** 30, 15, 46, 23, 70, 35, 106, 53, 160, 80, 40, 20, 10, 5
- **B:** need learn only two rules
  - If even, divide by 2
  - If odd, multiply by 3 and add 1

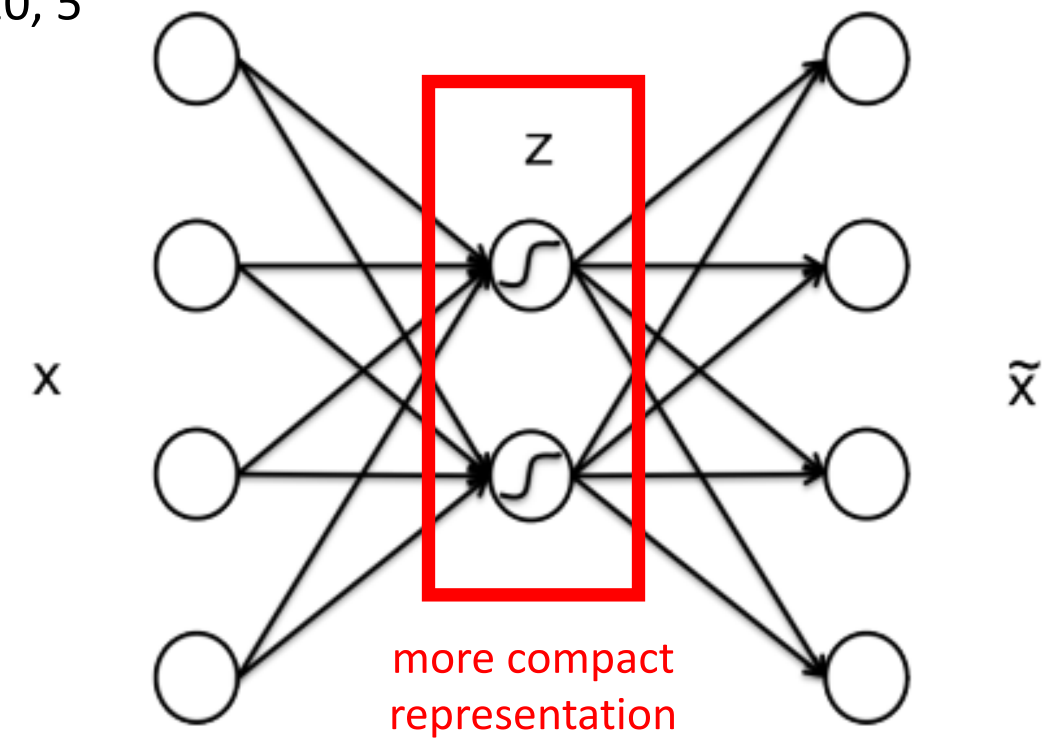

Figure Credit: https://lazyprogrammer.me/a-tutorial-on-autoencoders/

Autoencoders: Feature Extraction

- e.g., training data:
  - 1 image taken from 10 million YouTube videos
  - Each image is in color and 200x200 pixels

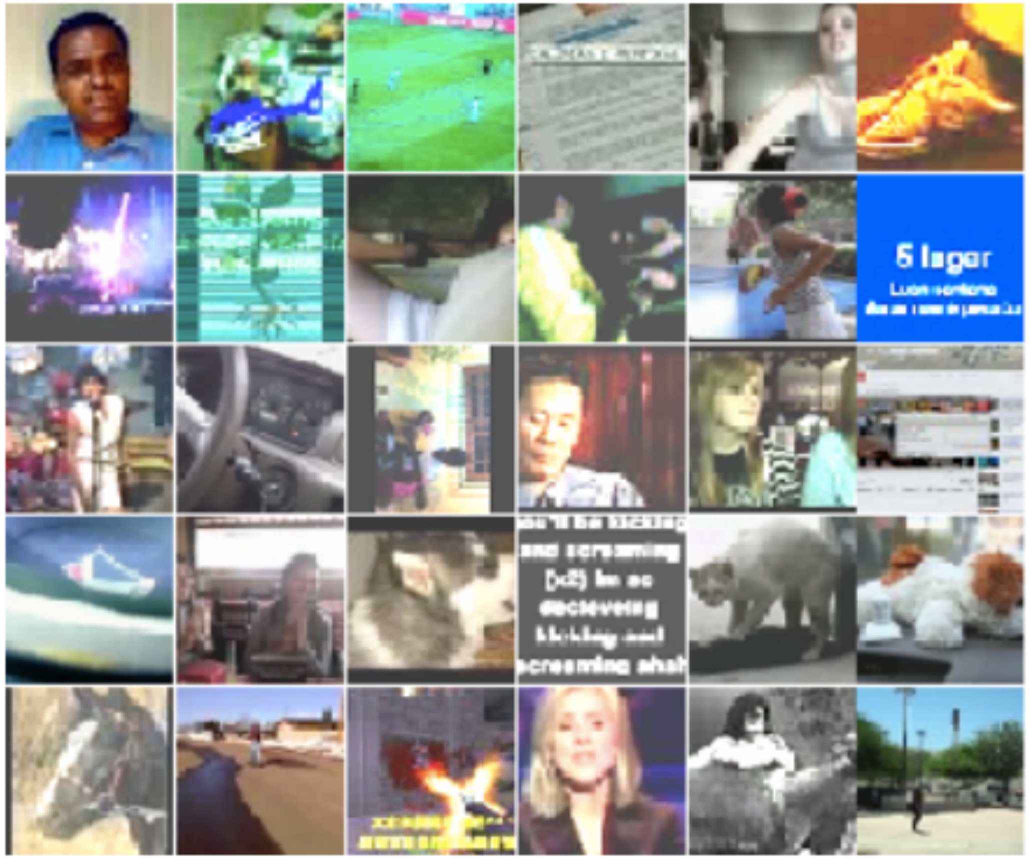

- What features do you think it learned?

Quoc V. Le et al., Building High-level Features Using Large Scale Unsupervised Learning; 2013.

Autoencoders: Feature Extraction

- e.g., features learned include:

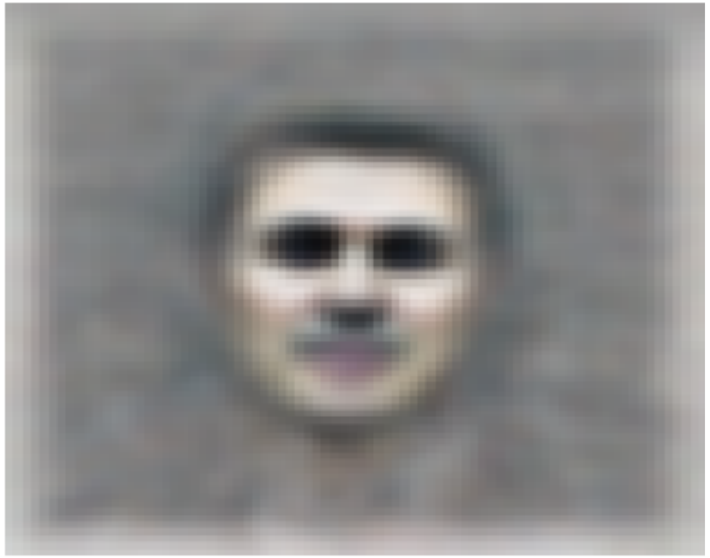

human face

cat face

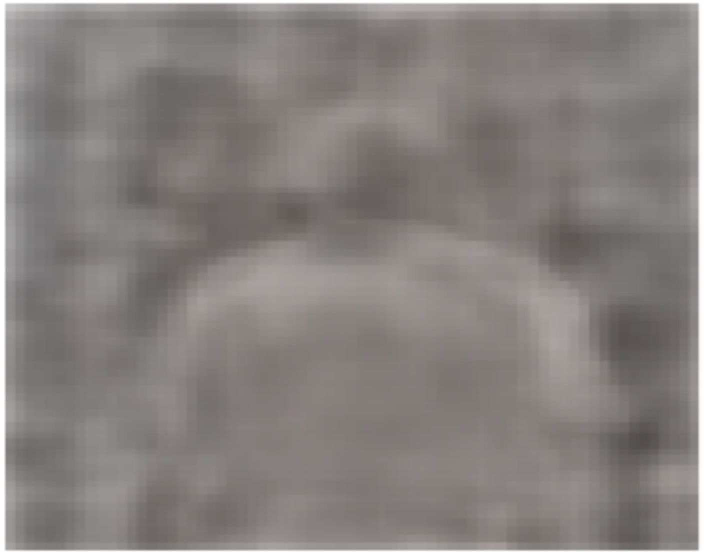

human body

Quoc V. Le et al., Building High-level Features Using Large Scale Unsupervised Learning; 2013.

Autoencoders: Unsupervised Pretraining

- Why use unsupervised pretraining?
  - Little training data is available
  - Too costly and slow to collect labels for exclusive supervised training

- e.g., add layer after highest layer of pretrained autoencoder network (fine-tuning)

Quoc V. Le et al., Building High-level FeaturesUsing Large Scale Unsupervised Learning; 2013.

Autoencoders: Generative Models

Ian Goodfellow et al; Adversarial Autoencoders; 2016

Today’s Topics

- Machine Learning for Unlabeled Data
- Clustering
- Autoencoders
- Extra: Active Learning, Curriculum Learning, and Reinforcement Learning
- Course Summary

Active Learning: Idea

Passive Learning

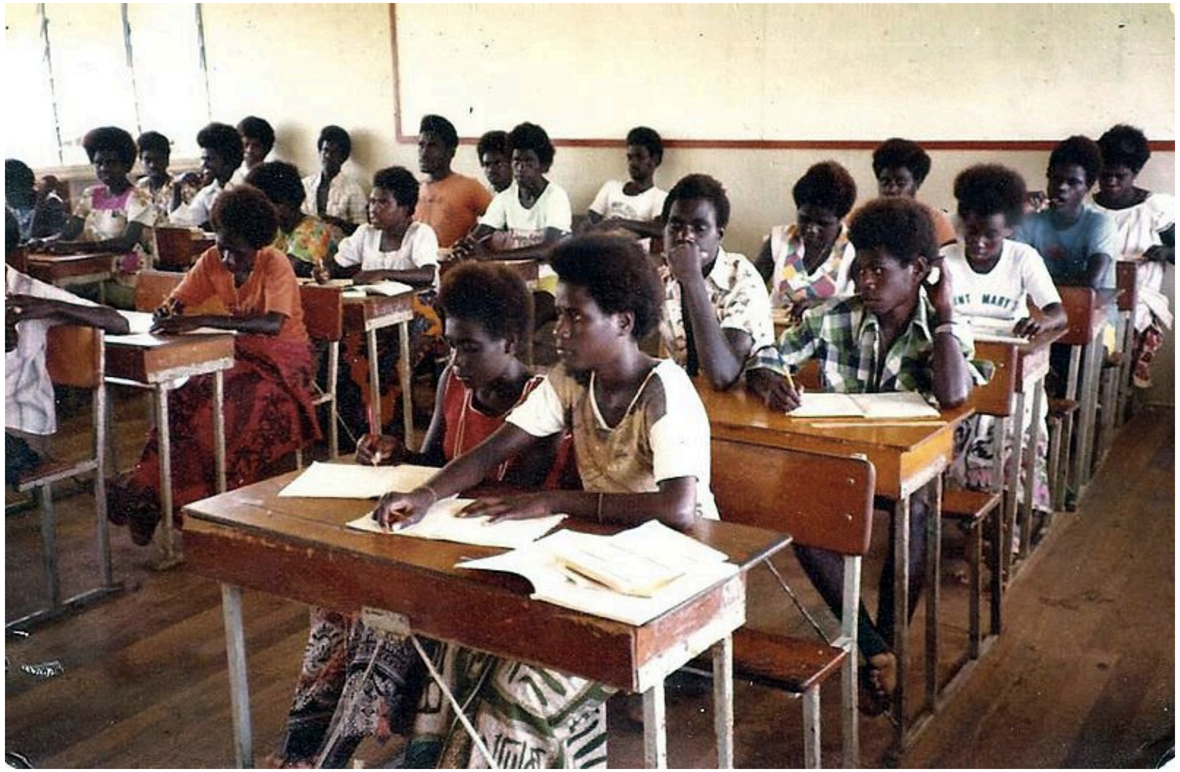

Active Learning

What is the difference between “passive” and “active” learning?

Passive Learning: Classical ML Approach

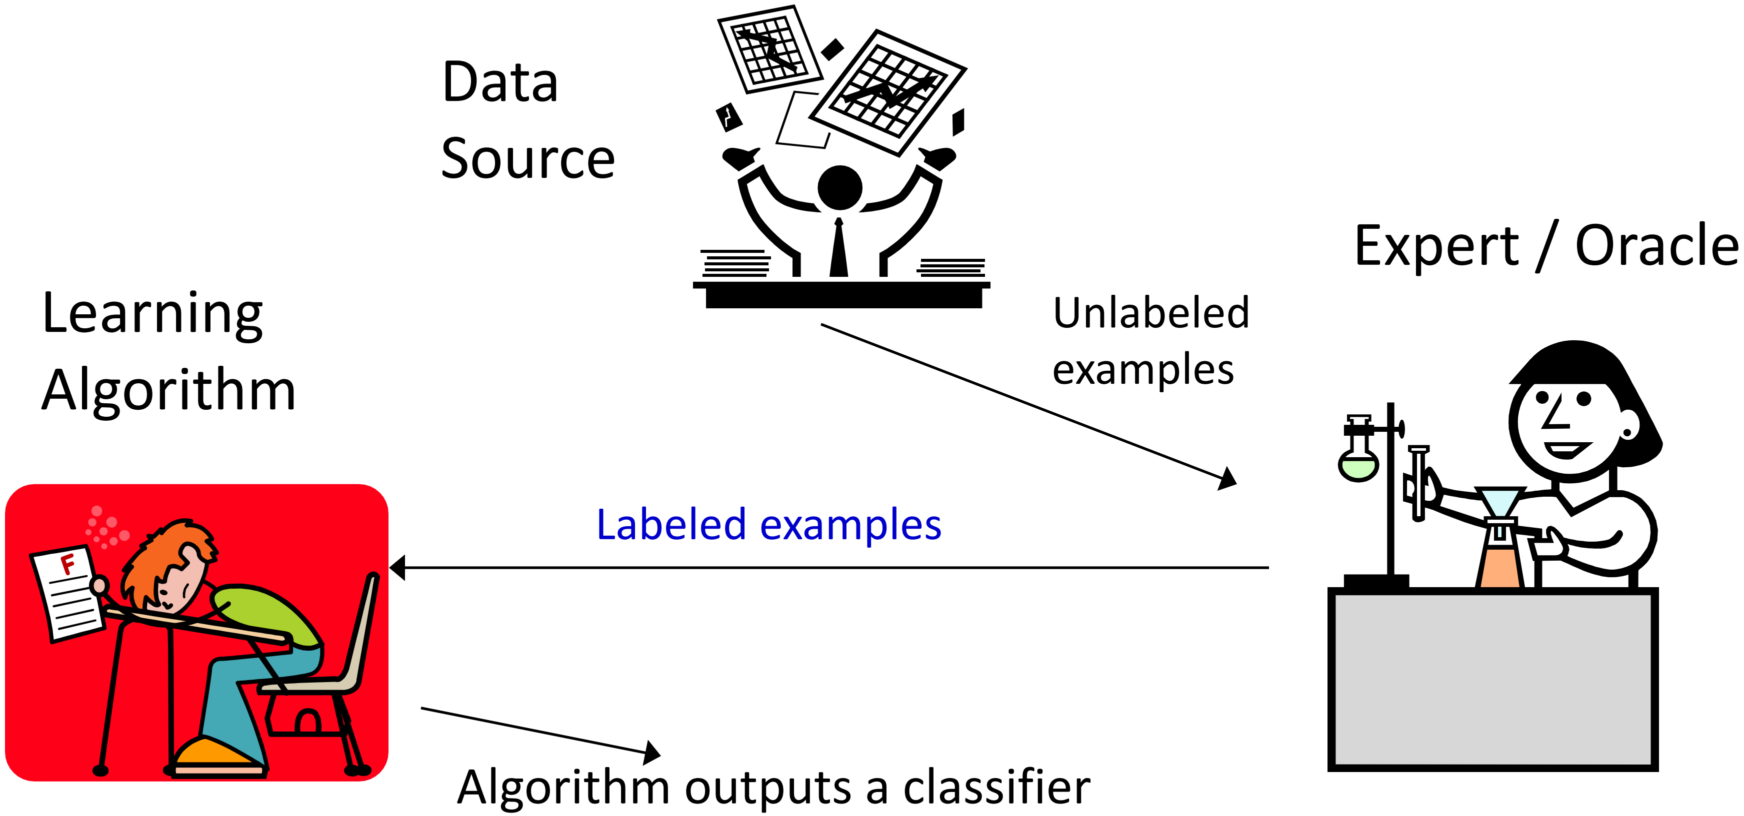

Slide Credit: http://www.cs.cmu.edu/~learning/talks-2007-spring/slides/mll0319.active_learning.ppt

Active Learning

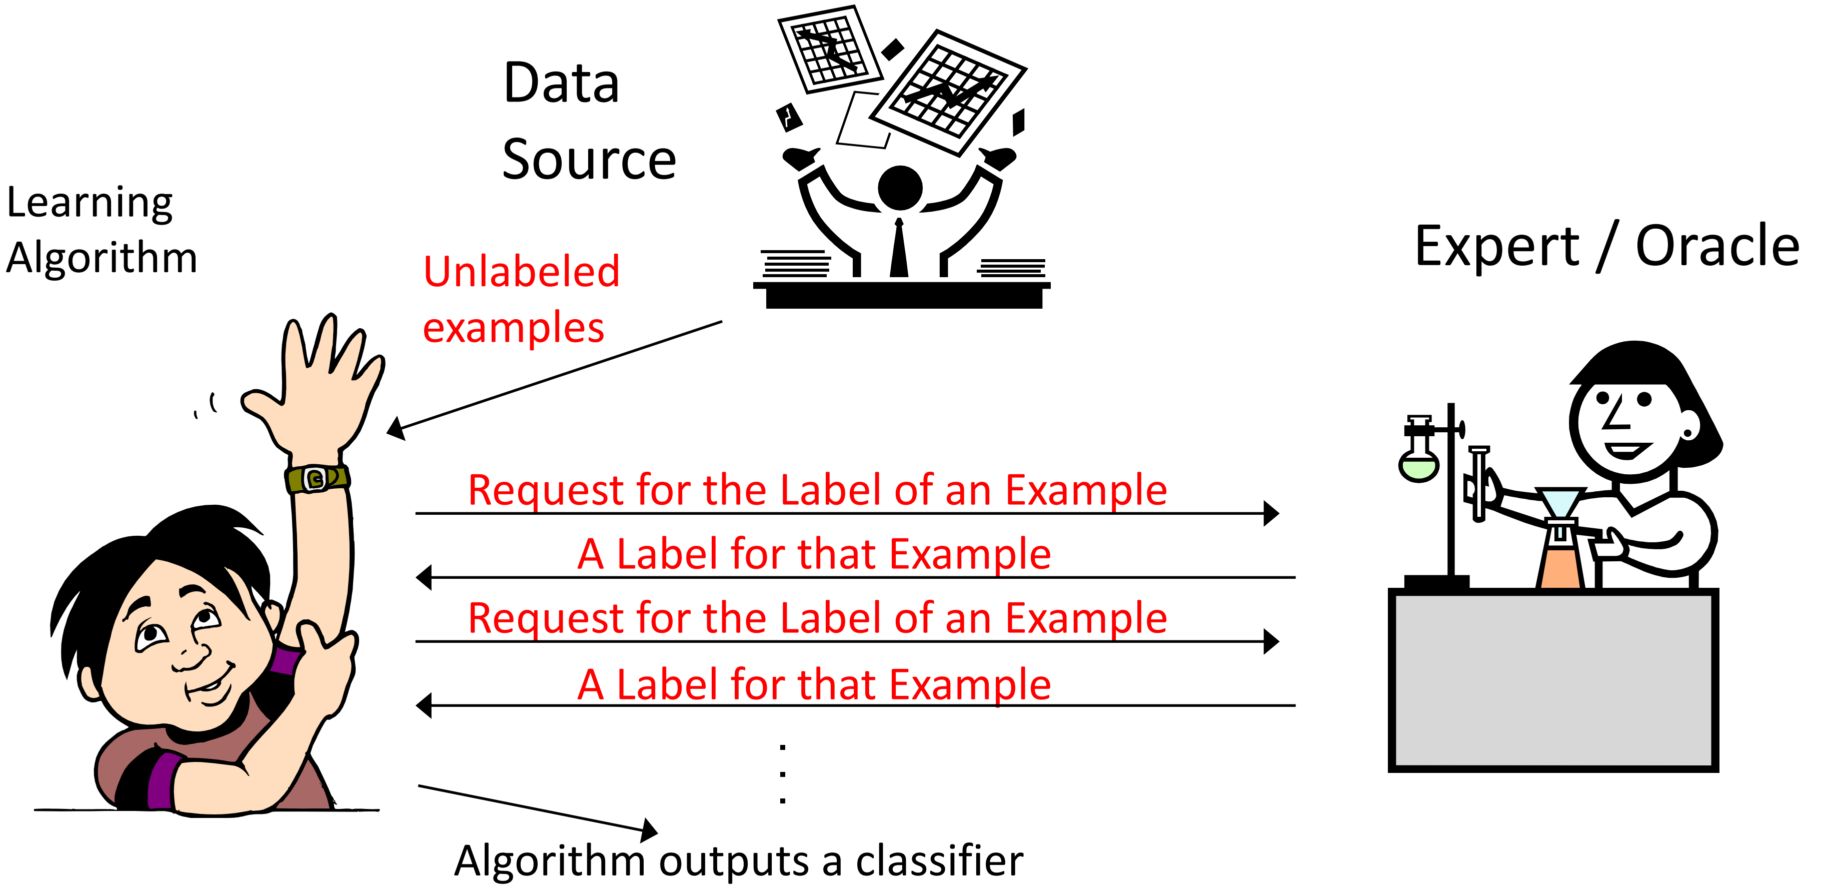

Slide Credit: http://www.cs.cmu.edu/~learning/talks-2007-spring/slides/mll0319.active_learning.ppt

Types of Active Learning

Stream-Based

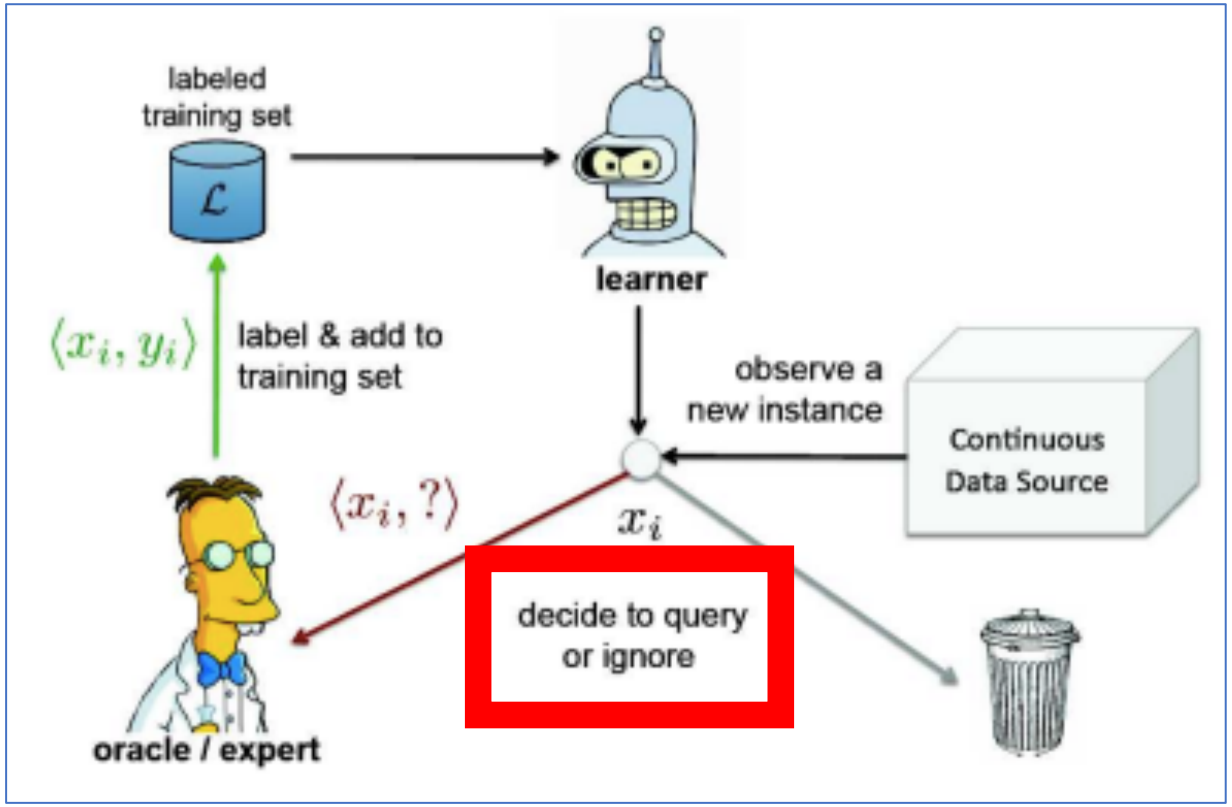

Pool-Based

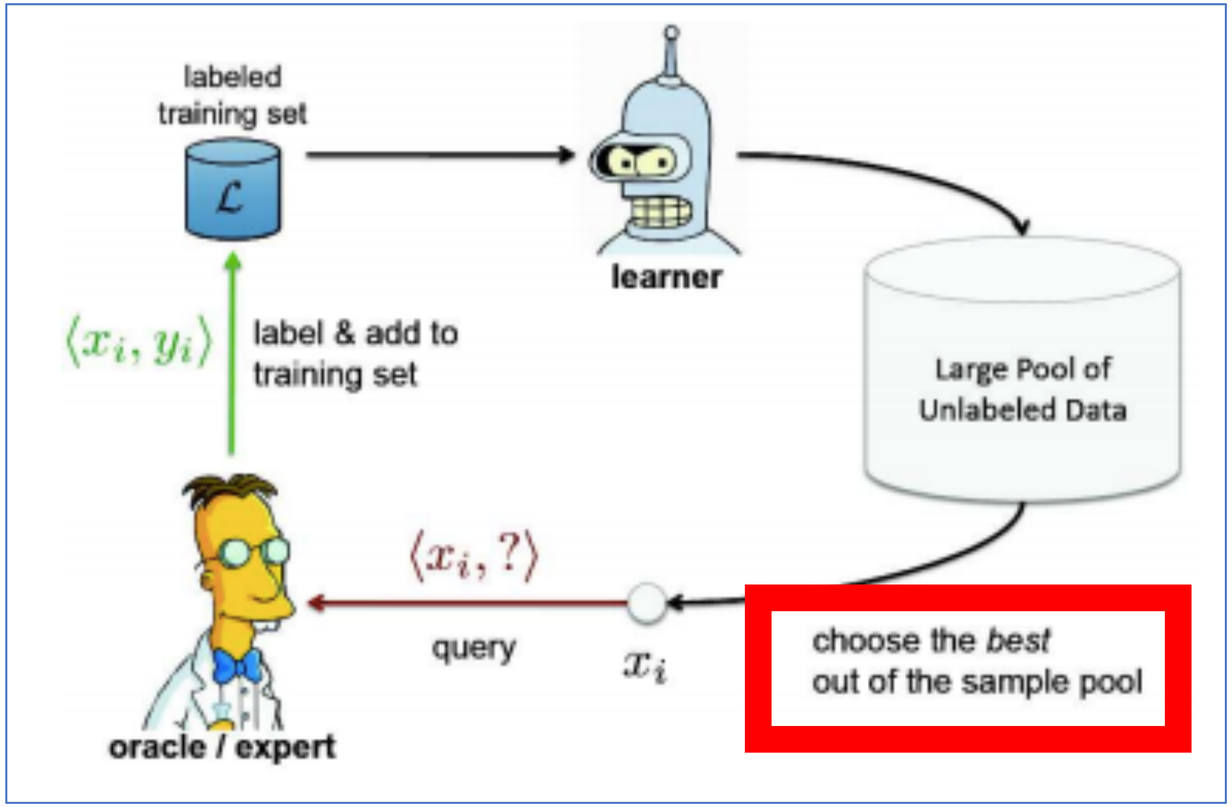

Consider one example at a time

Consider many examples at a time

Image Credit: https://www.cs.utah.edu/~piyush/teaching/10-11-slides.pdf

Active Learning Approach

Approach: query instances based on past queries and their responses (labels)

Problem: how to choose most informative examples to query?

Slide Credit: https://www.cs.utah.edu/~piyush/teaching/10-11-slides.pdf

Active Learning: Uncertainty Sampling

Query instance(s) the classifier is most uncertain about.

e.g., for SVM, request label of example closest to the current separator

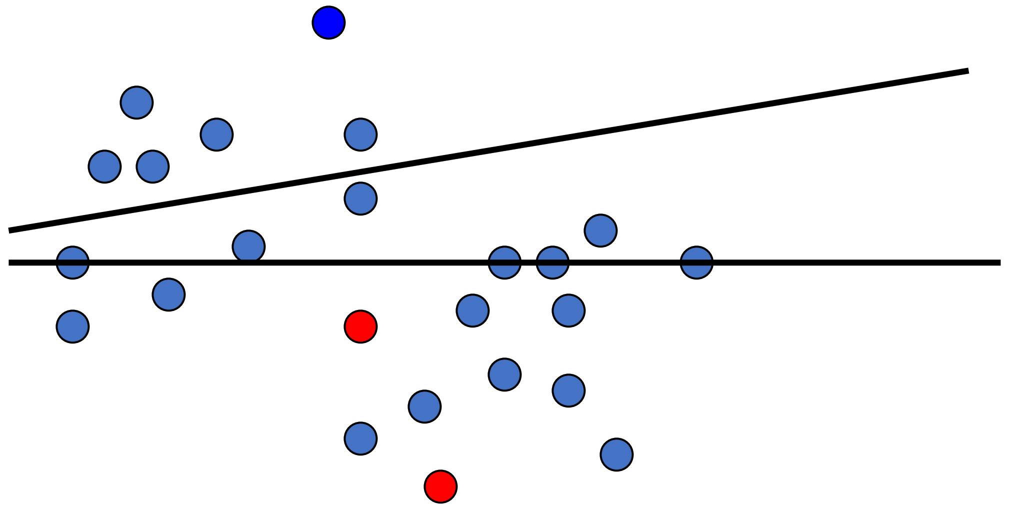

Slide Credit: http://www.cs.cmu.edu/~learning/talks-2007-spring/slides/mll0319.active_learning.ppt

Active Learning: Query By Committee

Query instance(s) different classifiers disagree most about.

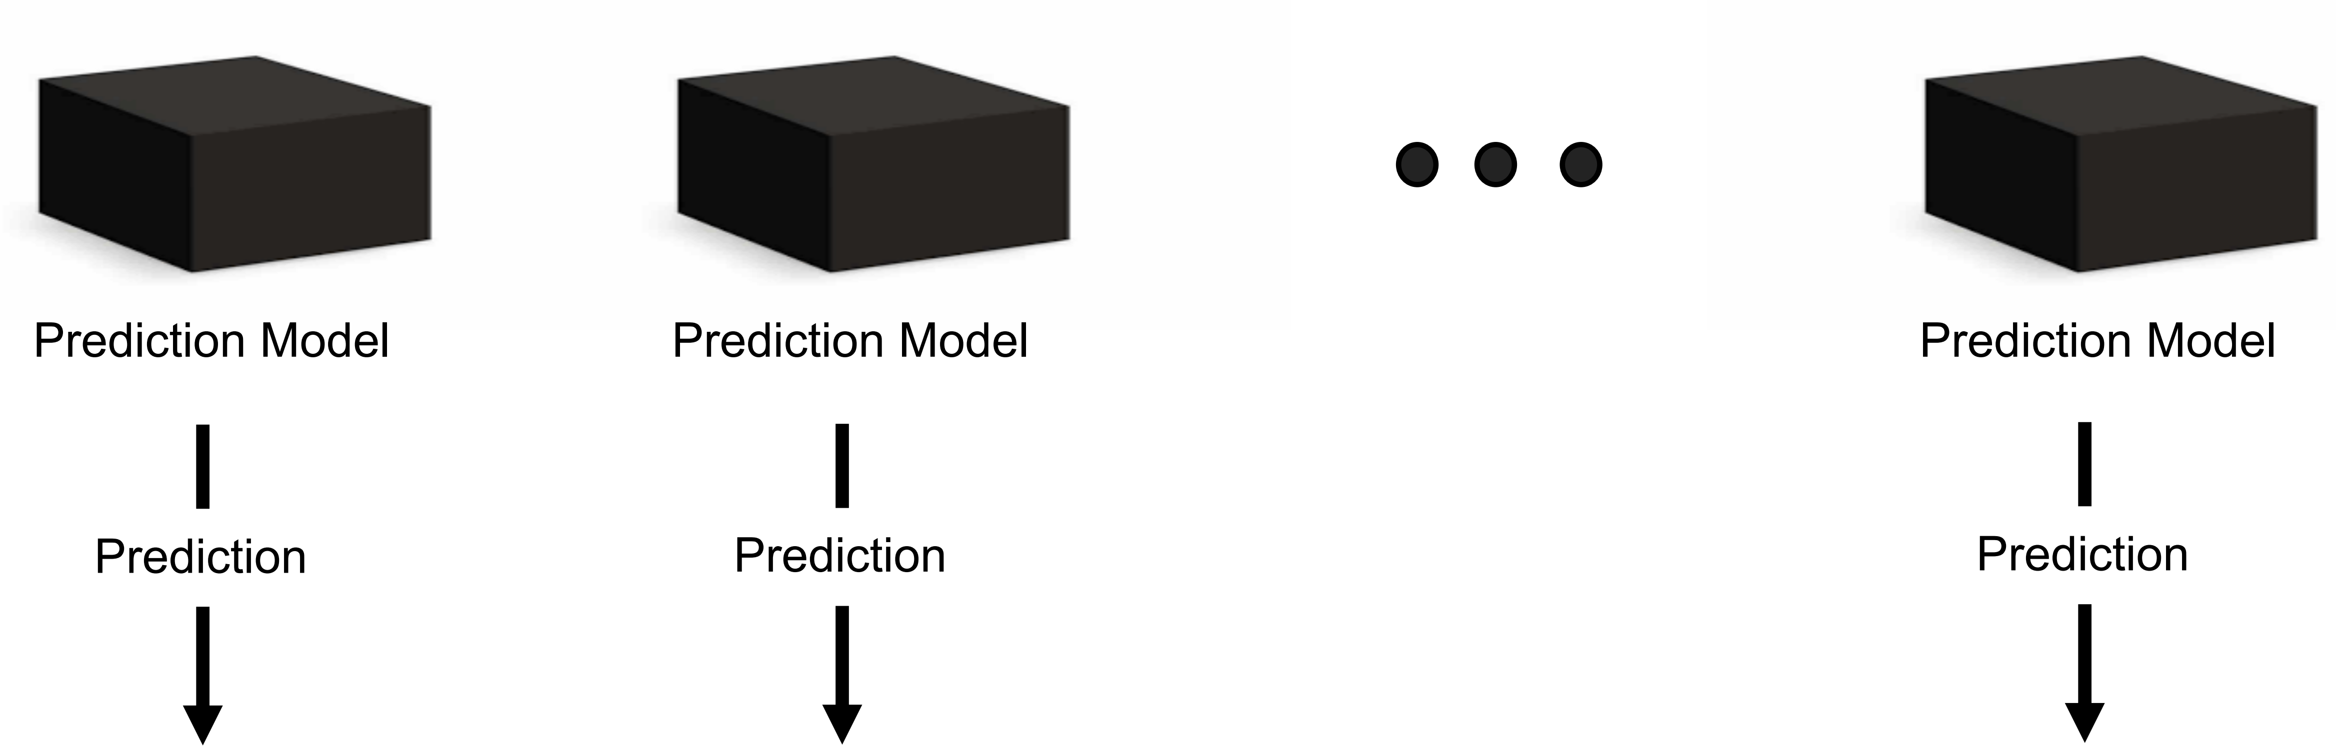

Image Credit: http://burrsettles.com/pub/settles.activelearning.pdf

Curriculum Learning: Idea

How to teach machines to learn faster?

e.g., How to Teach a Child Math?

Random Order of Examples

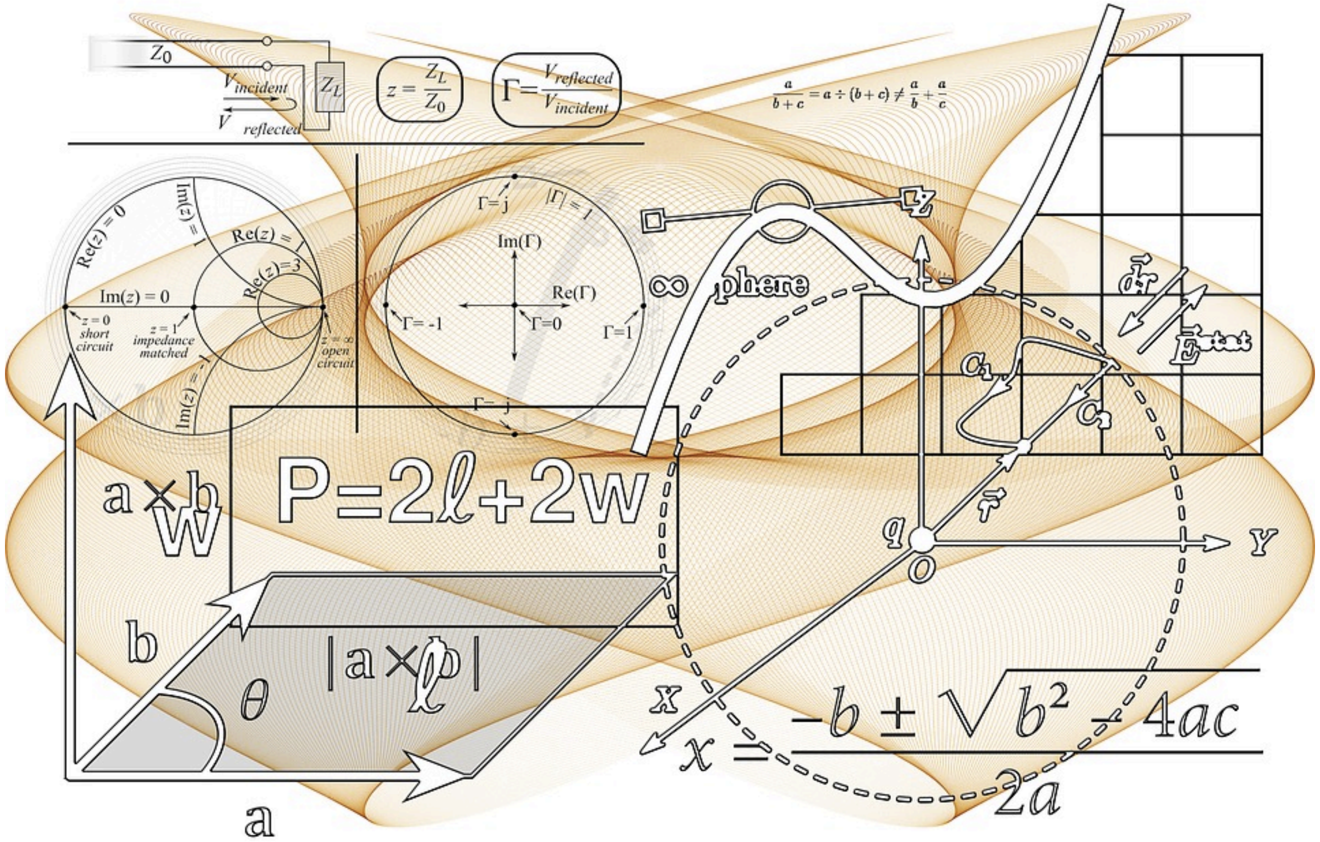

Meaningful Order of Examples

Table of Contents

Big Book of Math; Dinah Zike

e.g., How to Teach a Child To Read?

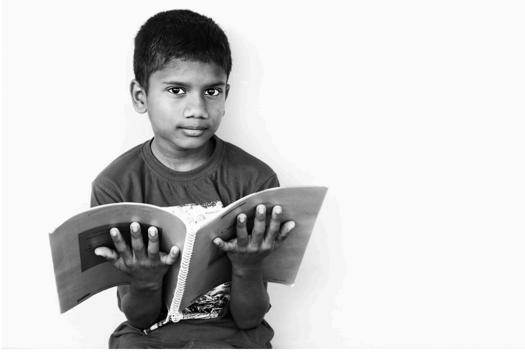

Random Order of Examples

Meaningful Order of Examples

Curriculum Learning

Task: train algorithm to read text in images taken by people who are blind

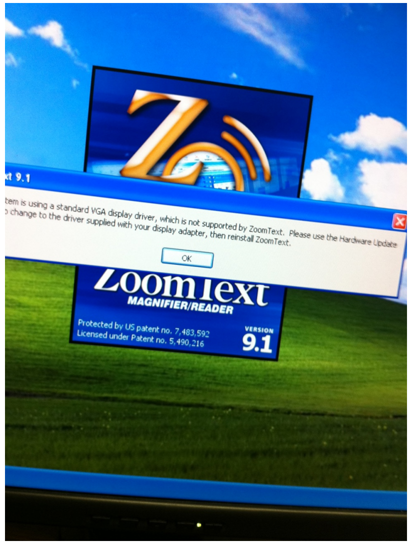

Questions

1. What criteria should be used to order examples?

2. After how long would you make updates to include harder examples?

Reinforcement Learning Overview

Agent takes actions in an environment so as to maximize the total reward.

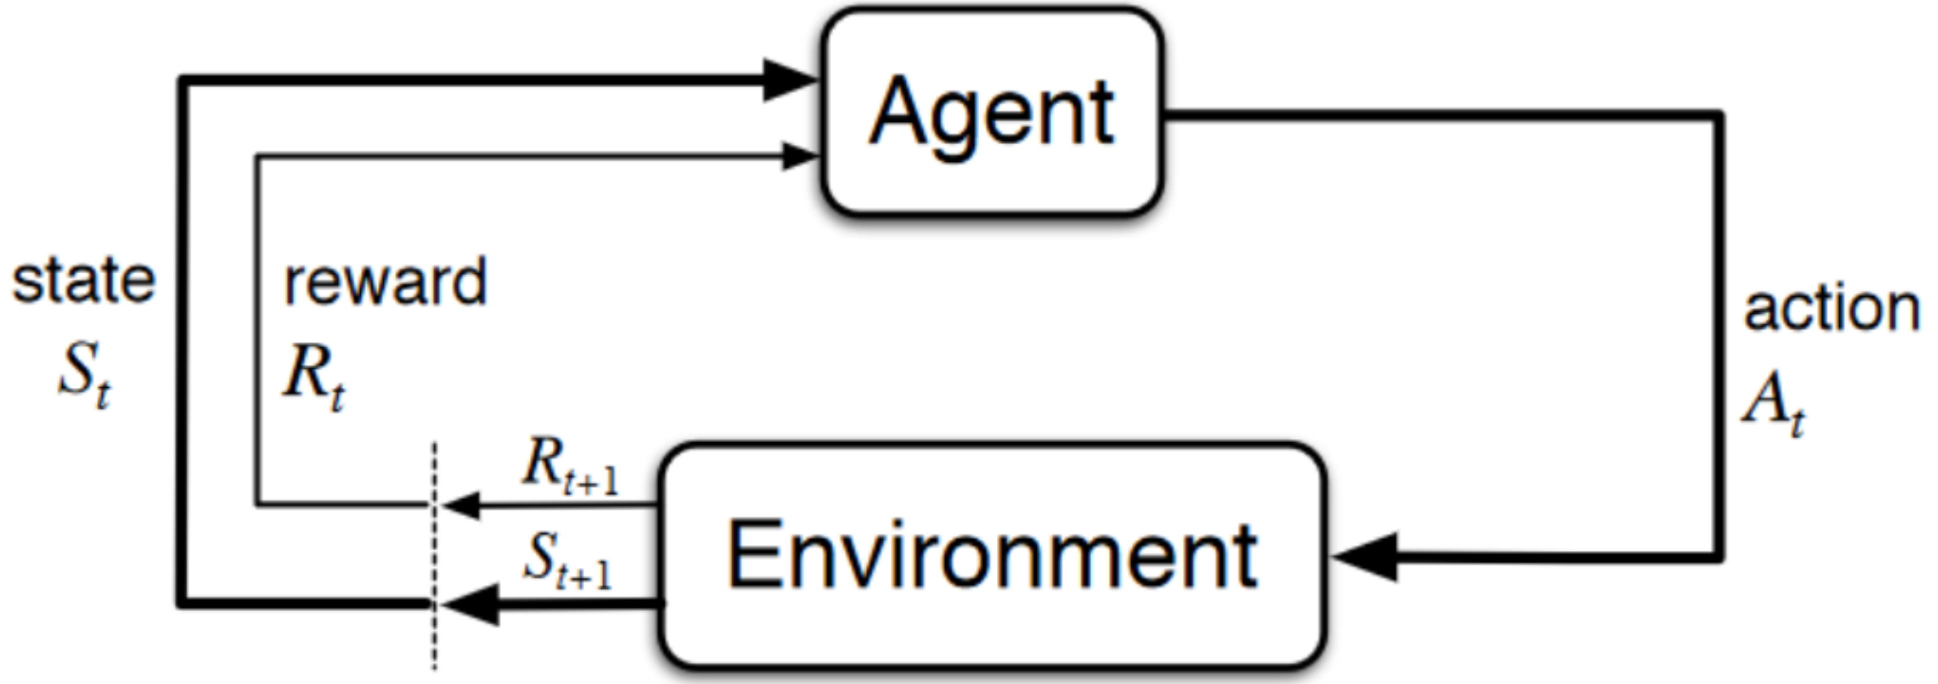

Figure Credit: https://towardsdatascience.com/applications-of-reinforcement-learning-in-real-world-1a94955bcd12

Intuition: Learning to Walk by Trial-and Error

https://en.wikipedia.org/wiki/Crawling_(human)

Reinforcement Learning Applications

Learning to Walk in 20 Minutes

Russ Tedrake
Brain & Cognitive Sciences
Center for Bits and Atoms
Massachusetts Inst. of Technology
Cambridge, MA 02139
russt@csail.mit.edu

Teresa Weirui Zhang
Mechanical Engineering
Department
University of California, Berkeley
Berkeley, CA 94270
resa@berkeley.edu

H. Sebastian Seung
Howard Hughes Medical Institute
Brain & Cognitive Sciences
Massachusetts Inst. of Technology
Cambridge, MA 02139
seung@mit.edu

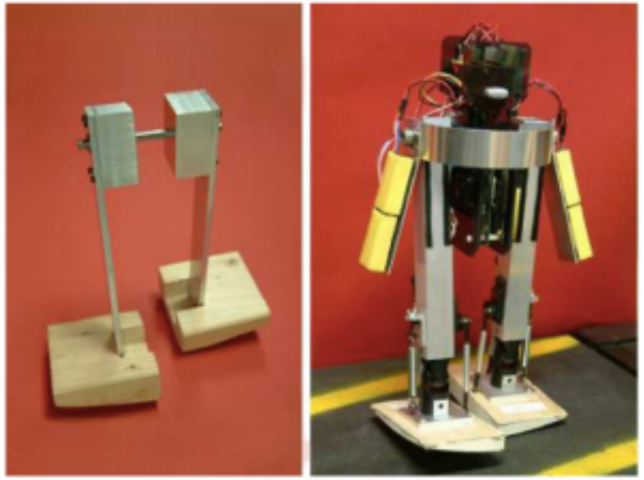

Reinforcement Learning Applications

Autonomous reinforcement learning on raw visual input data in a real world application

Sascha Lange, Martin Riedmiller
Department of Computer Science
Albert-Ludwigs-Universität Freiburg
D-79110, Freiburg, Germany
Email: [slange,riedmiller]@informatik.uni-freiburg.de

Arne Voigtländer
Shoogee GmbH & Co. KG
Krögerweg 16a
D-48155 Münster, Germany
Email: arne@shoogee.com

Fig. 1. The visual slot car racer task. The controller has to autonomously learn to steer the racing car by raw visual input of camera images.

Reinforcement Learning Applications

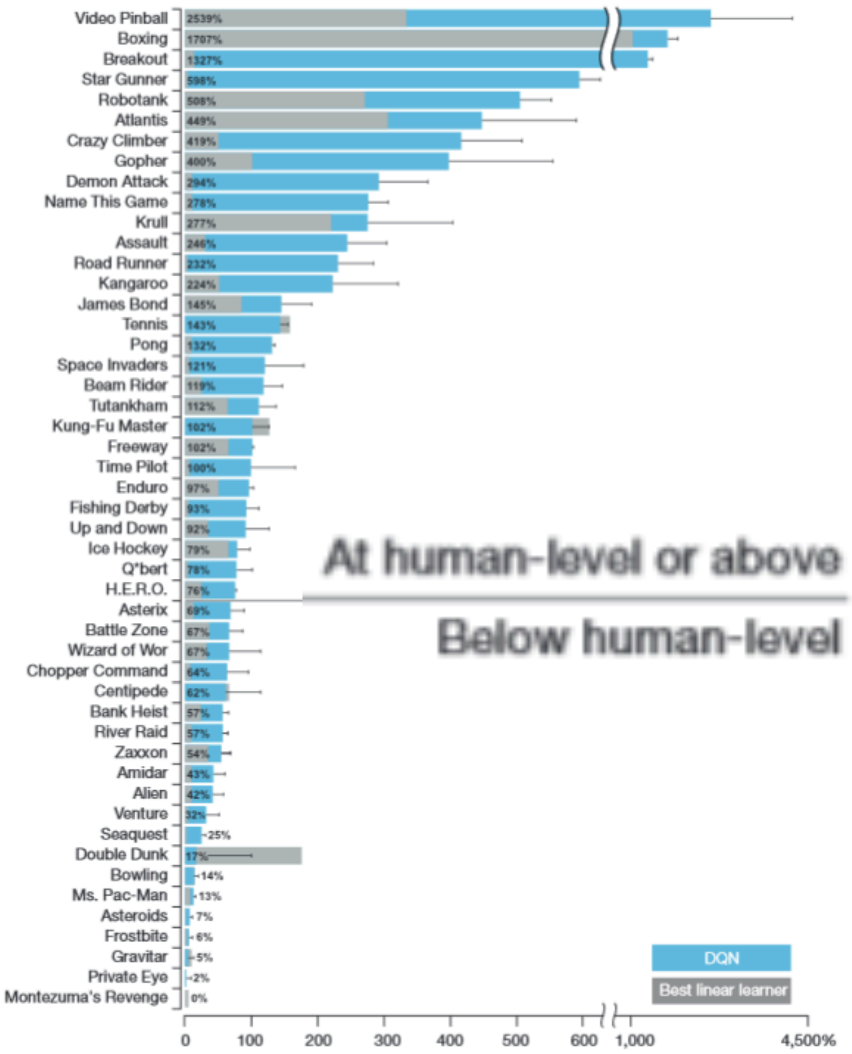

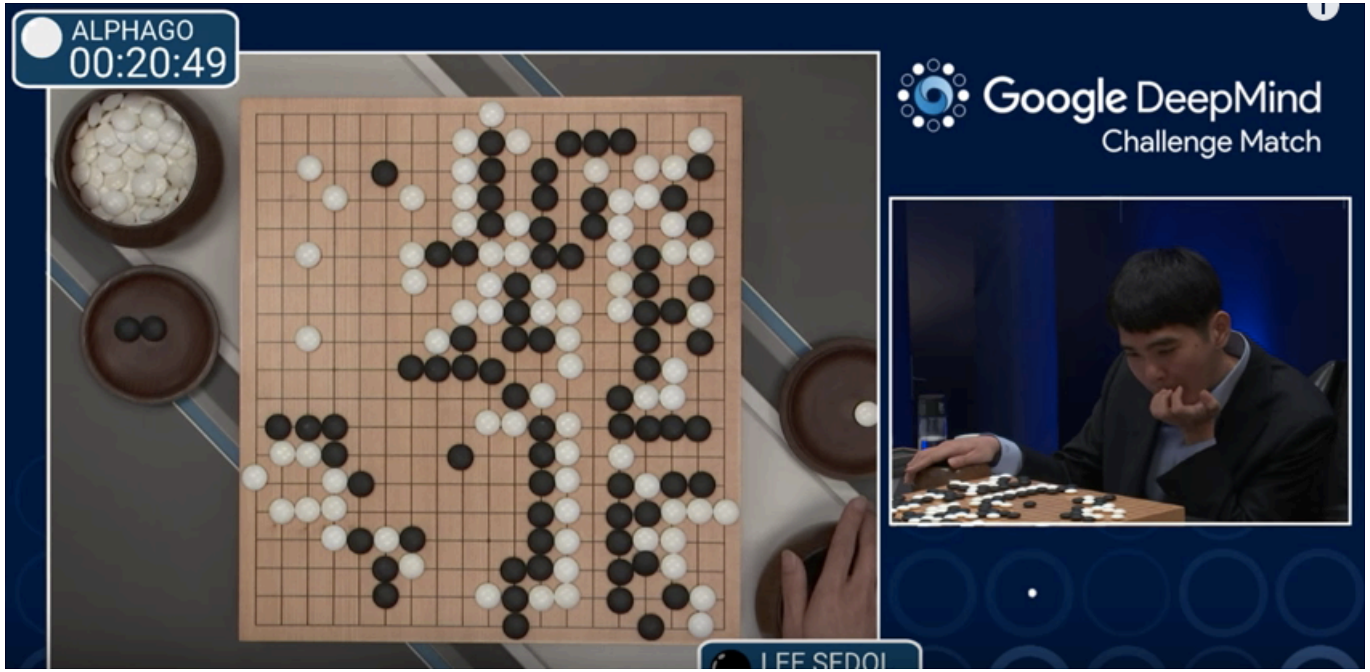

https://www.tastehit.com/blog/google-deepmind-alphago-how-it-works/

https://web.stanford.edu/class/psych209/Readings/MnihEtAlHassibis15NatureControlDeepRL.pdf

e.g., Pong Game - Learning Example

-1 if missed the ball

+1 reward if ball goes past opponent

0 otherwise

Move “up” or “down”

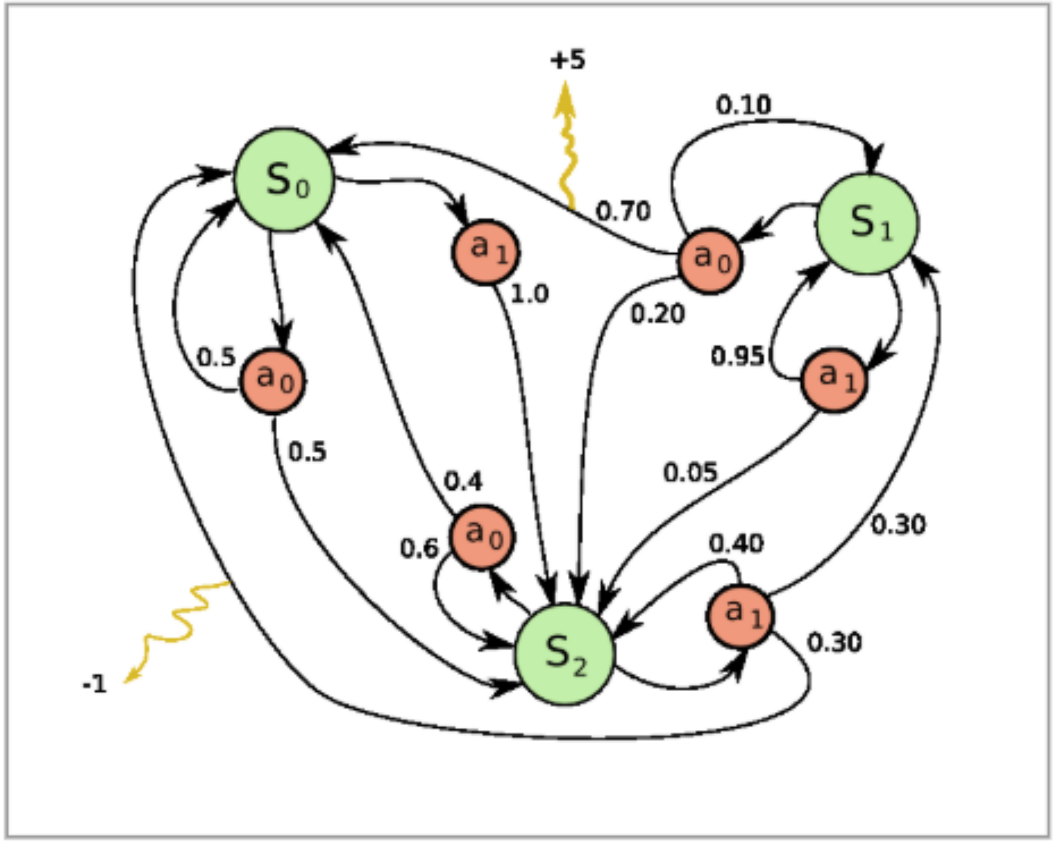

http://karpathy.github.io/2016/05/31/rl/

e.g., Pong Game: Policy Network

Implements our player (or “agent”)

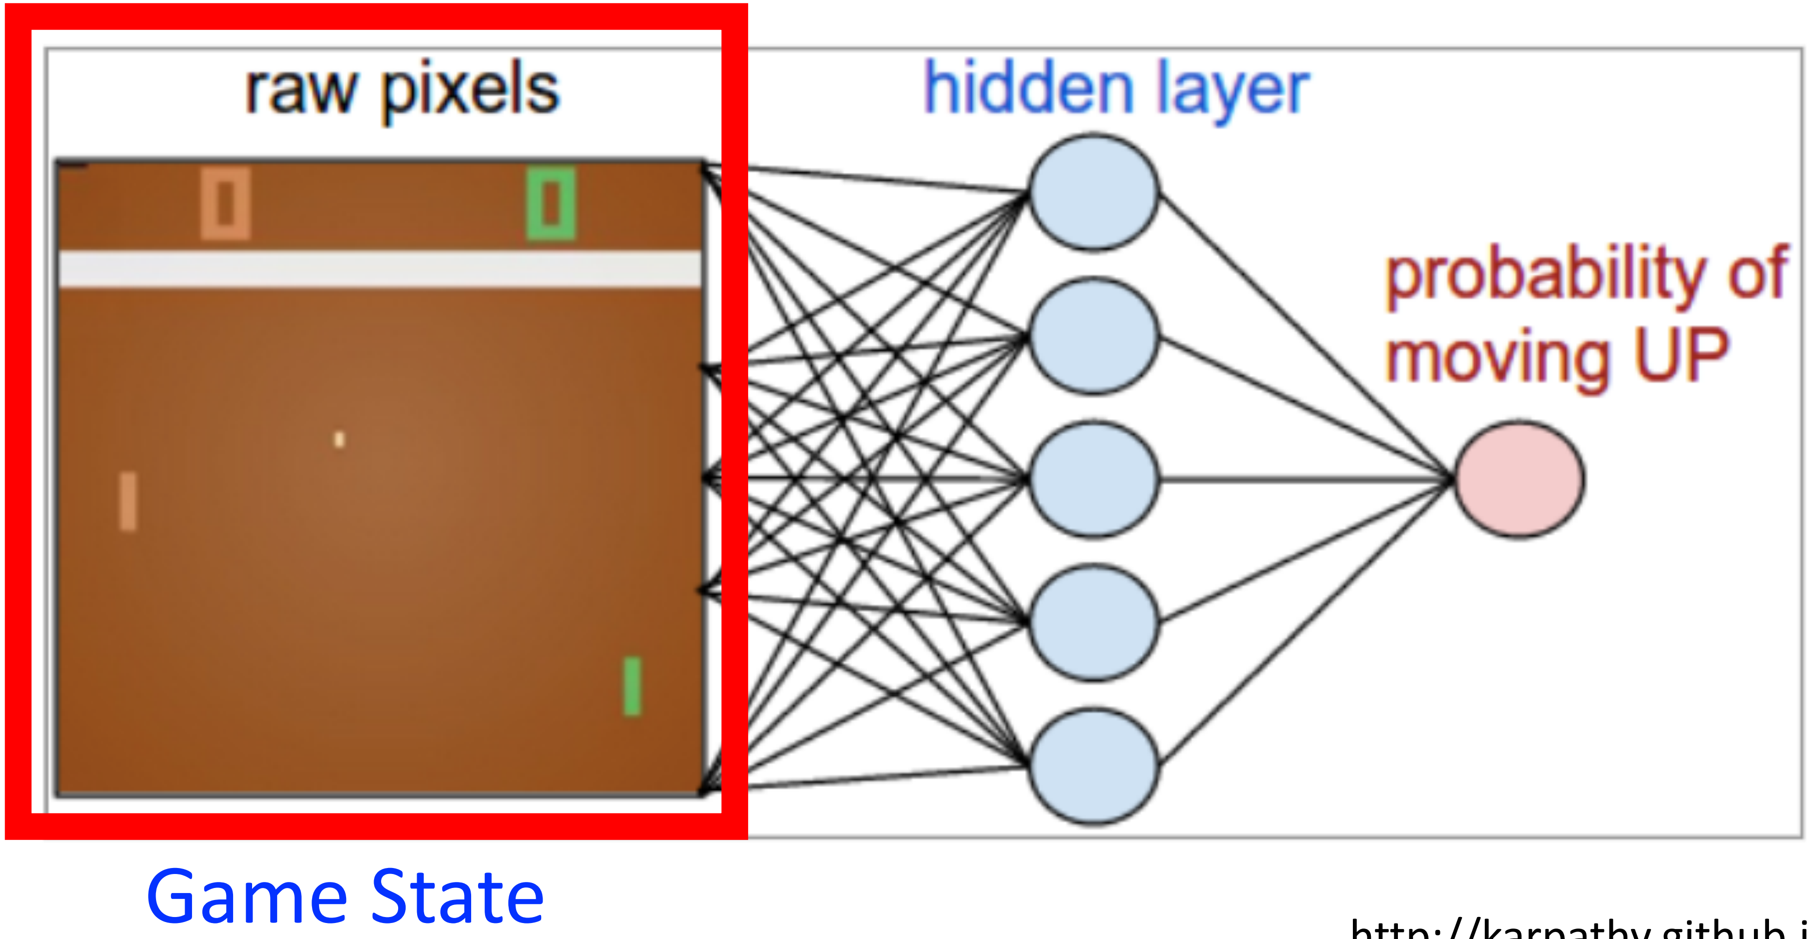

http://karpathy.github.io/2016/05/31/rl/

e.g., Pong Game: Training Protocol

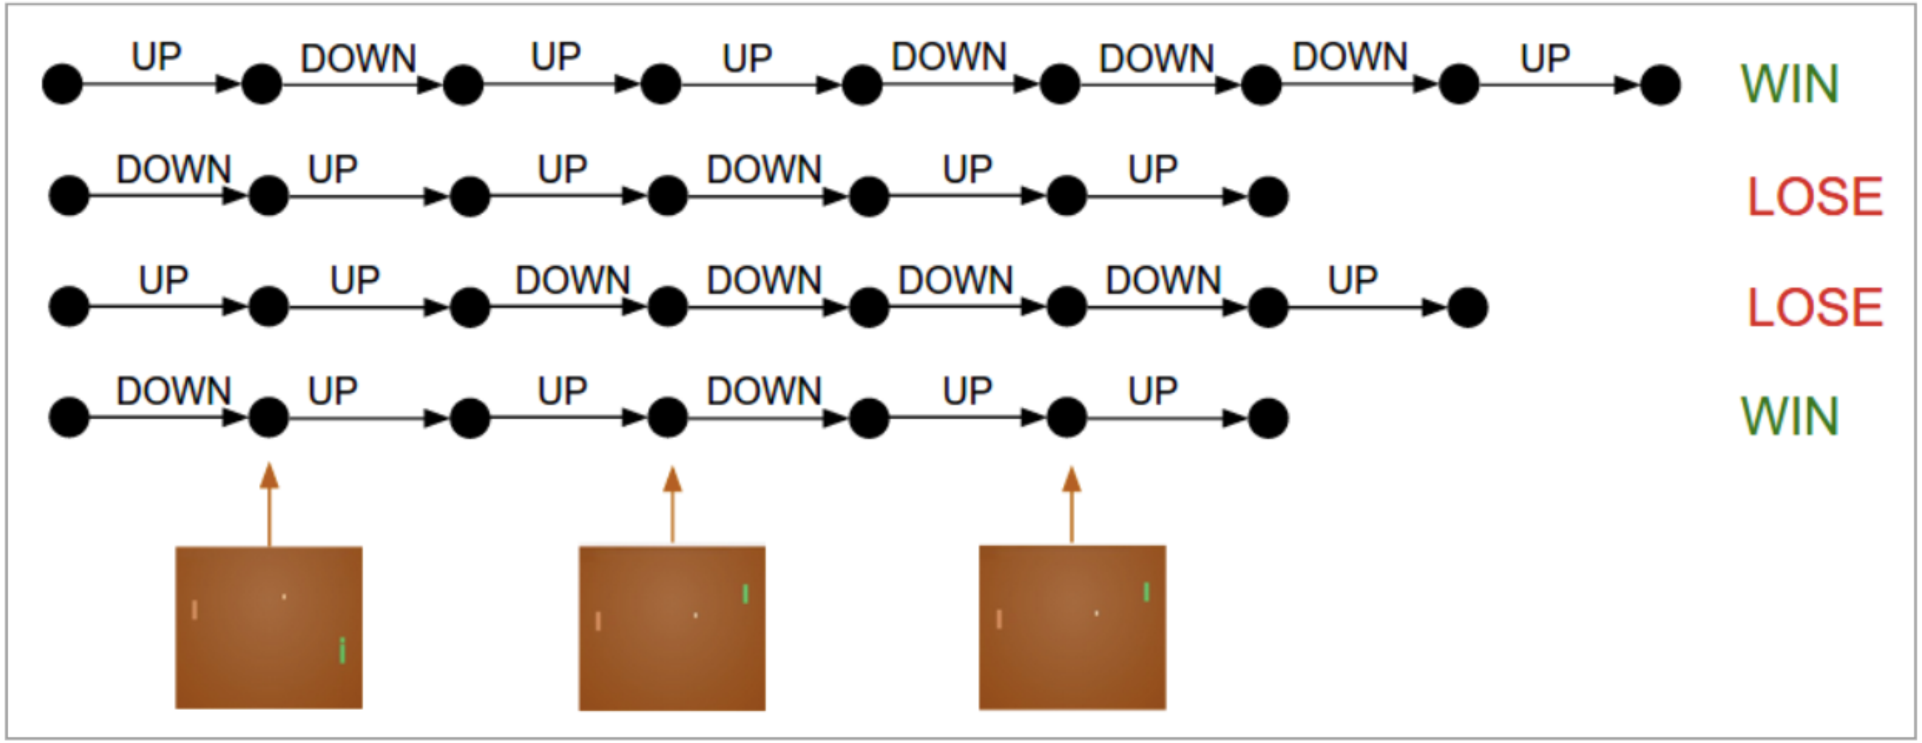

- Play 100 games of Pong; i.e., policy “rollouts” (200 images/game); Suppose: win 12 games, lose 88
- # Winning Decisions = 200*12 = 2400 decisions; positive update (fill in a +1.0 in the gradient for the sampled action, do backprop, and parameter update to encouraging the actions)
- # Losing Decisions: 200*88 = 17600; negative update (as above, but fill in -1.0 in the gradient)

http://karpathy.github.io/2016/05/31/rl/

e.g., Pong Game: Trained for Three Nights

Demo: https://www.youtube.com/watch?time_continue=16&v=YOW8m2YGtRg

e.g., Learning Dexterity

- Demo: https://www.youtube.com/watch?v=jwSbzNHGflM

e.g., Learning to Flip Pancakes

Demo:
https://www.youtube.com/watch?v=W_gxLKSsSIE&list=PL5nBAYUyJTrM48dViibyi68urttMlUv7e

e.g., Learning to Walk

- Demo: https://www.youtube.com/watch?v=gn4nRCC9TwQ

Today’s Topics

- Machine Learning for Unlabeled Data
- Clustering
- Autoencoders
- Extra: Active Learning, Curriculum Learning, and Reinforcement Learning
- Course Summary

Closing Remarks: My "Why" for Teaching You…

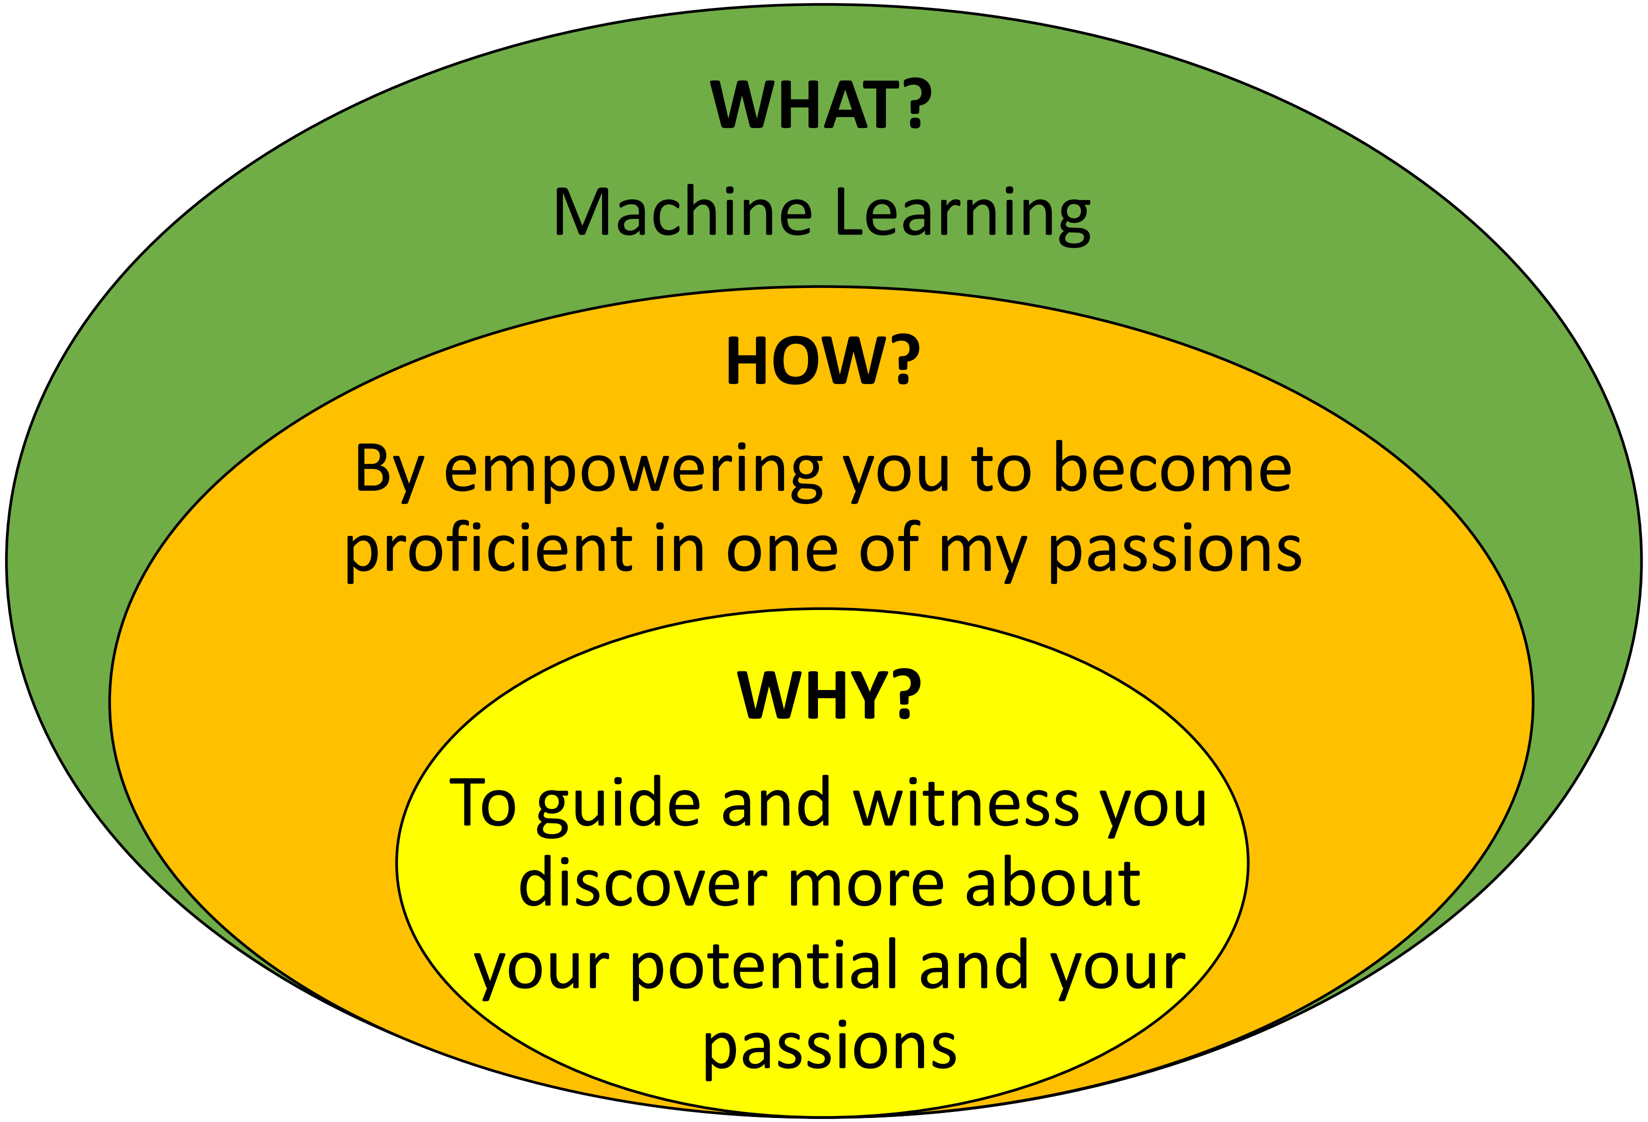

Algorithm Scope for Course: Last 61 Years And More

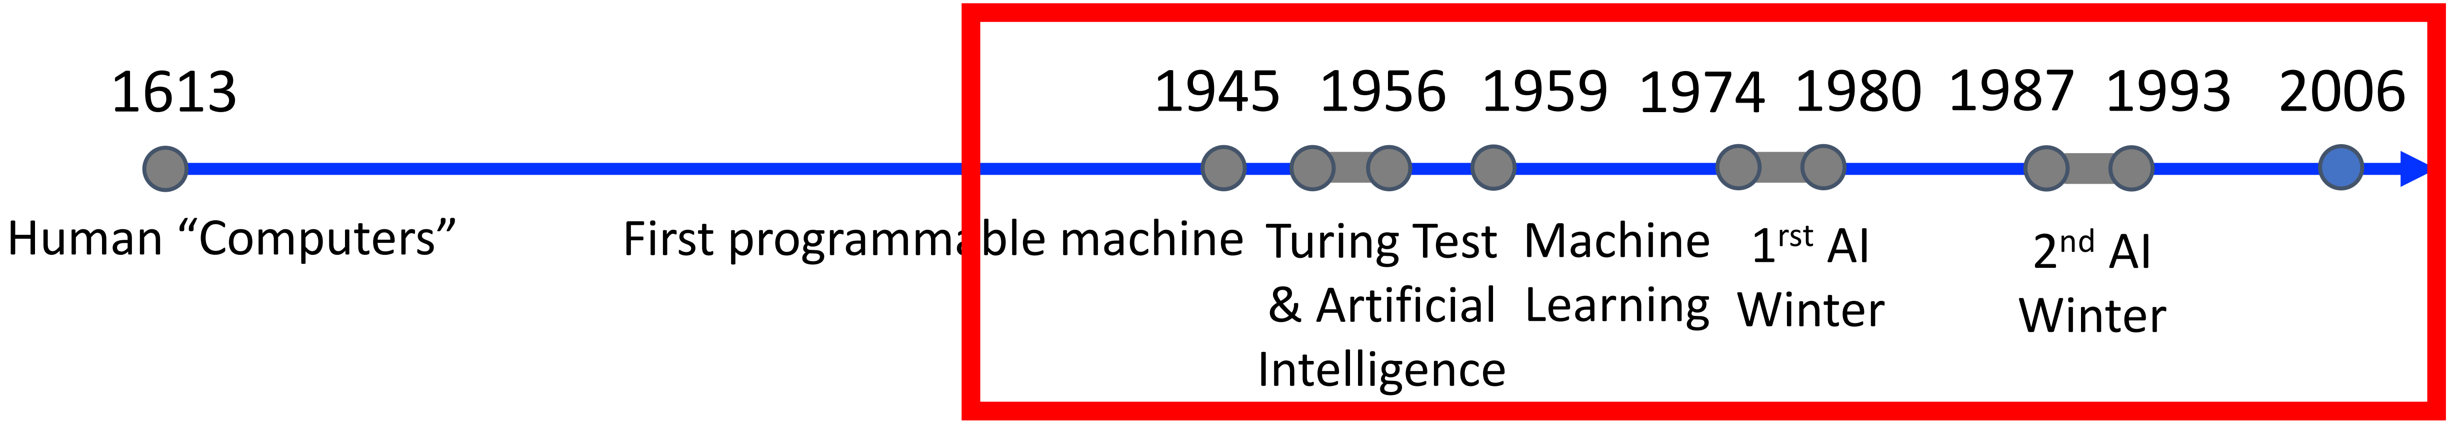

Review: AI, Classical Algorithms, & Deep Learning

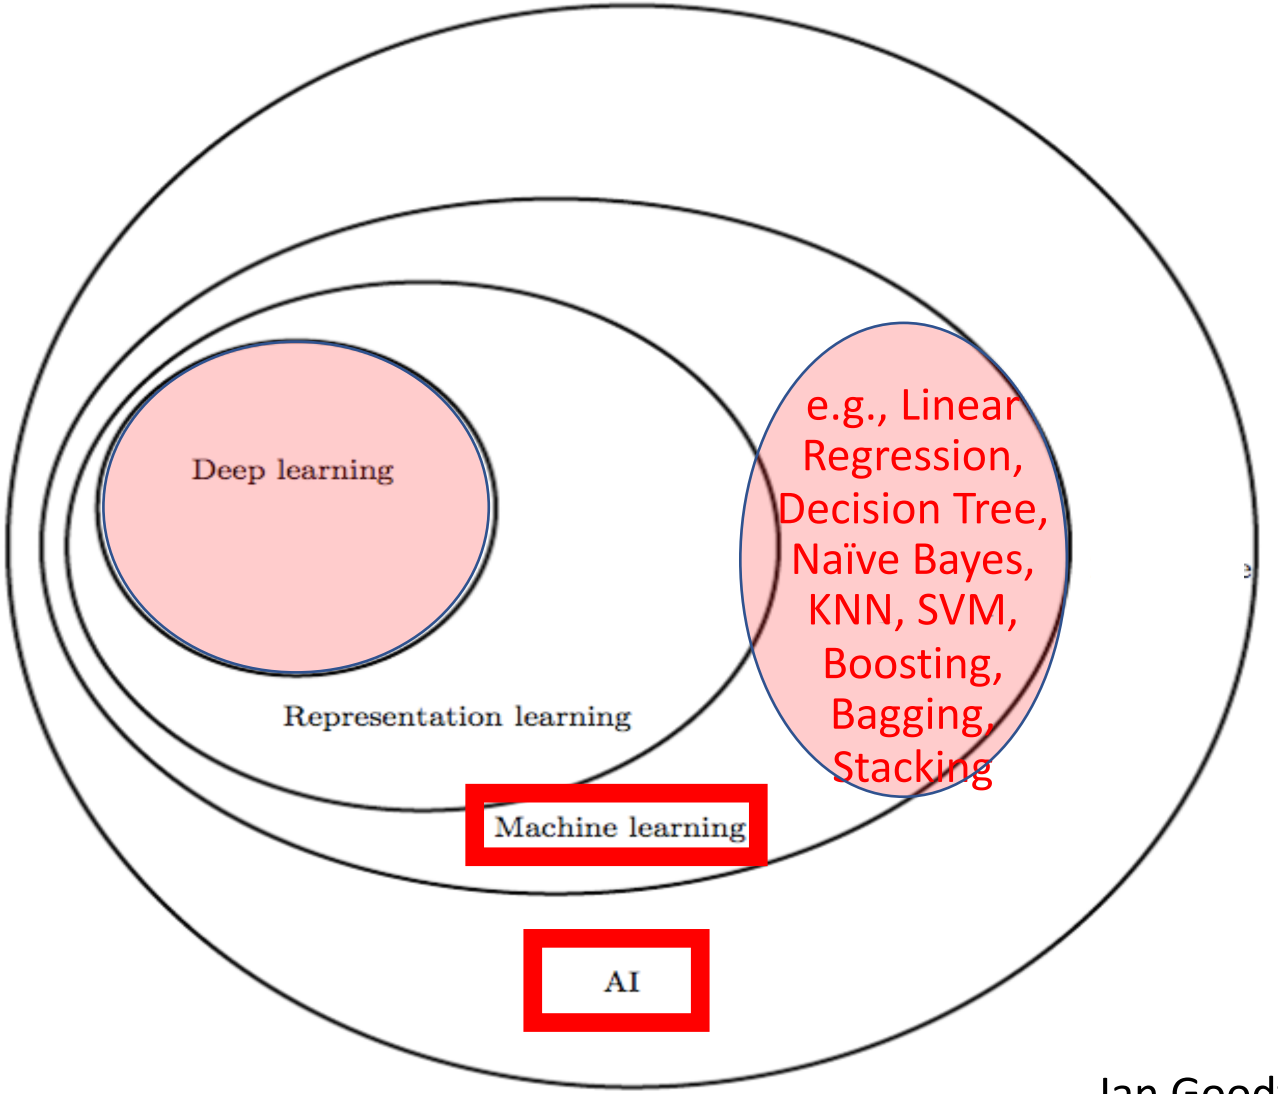

- What is artificial intelligence?
  - 1956: machines that do intelligent things
- What is machine learning?
  - 1959: algorithms that learn on their own
- What machine learning algorithms did we study for the first half of the class?
  - Algorithms other than neural networks
- What machine learning algorithms did we study for the second half of the class?
  - Neural networks (deep learning)

Ian Goodfellow, Yoshua Bengio, and Aaron Courville; Deep Learning, 2016.

Review: Other Topics

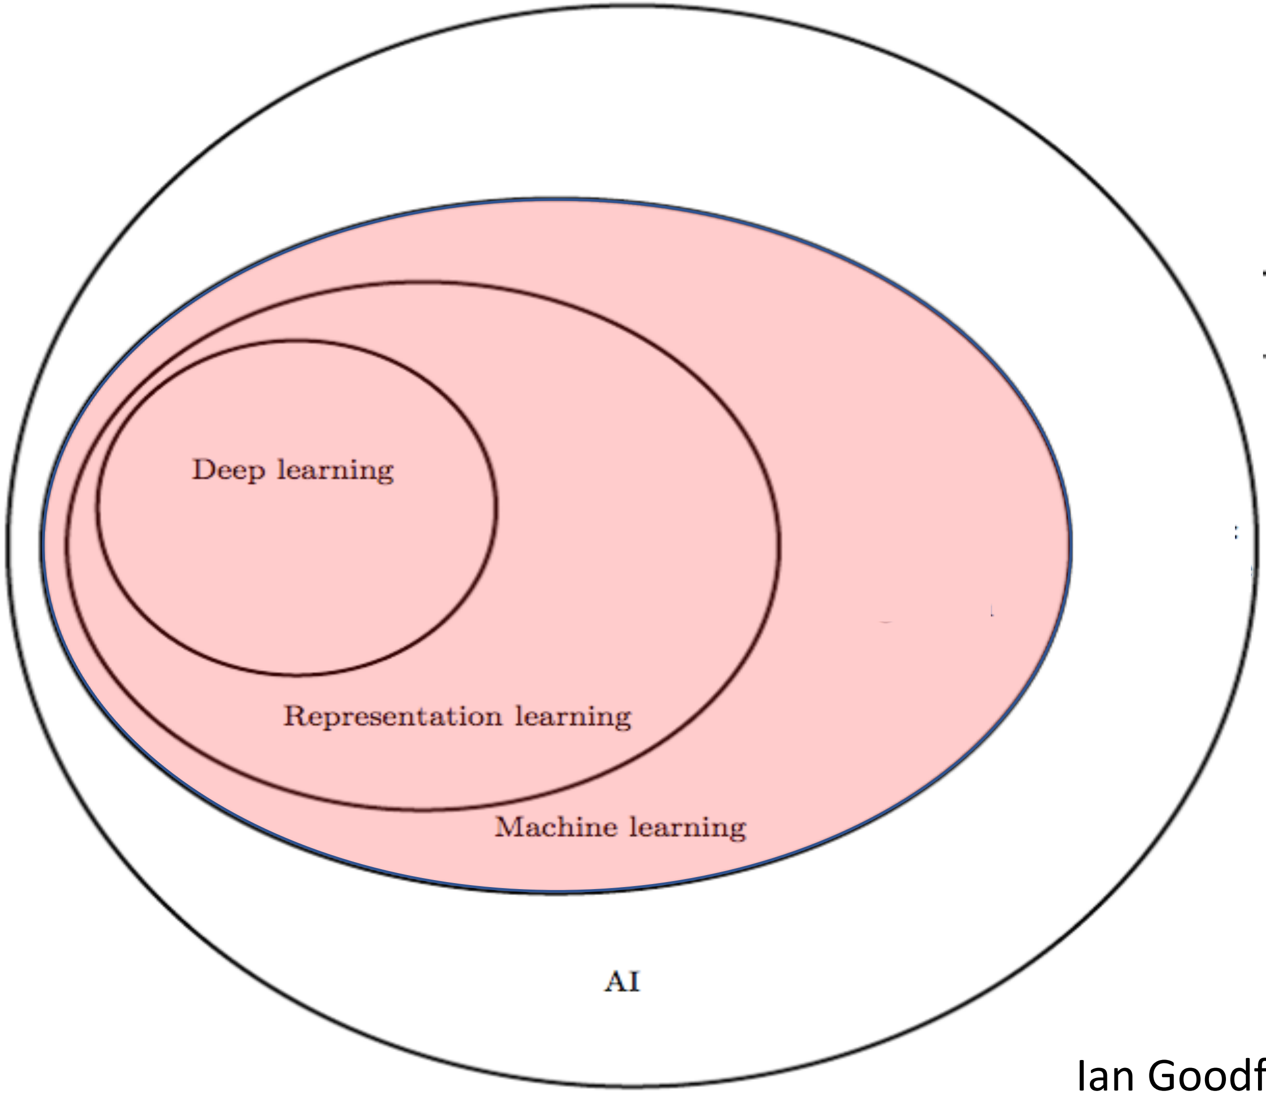

Lecture Topic(s)
Feature Representation, Dimensionality Reduction
Active Learning, Curriculum Learning, Reinforcement Learning
Algorithm Fairness, Accountability, Transparency, and Ethics

Ian Goodfellow, Yoshua Bengio, and Aaron Courville; Deep Learning, 2016.

Review: Towards a More Responsible Future

- Move away from algorithms that discriminate to support diverse populations

Course Review

- Please complete the following:
  - Course review: https://utdirect.utexas.edu/ctl/ecis/
    - Please also consider providing feedback about our TA

Next Steps

Take other machine learning classes at UT Austin from other professors:

http://ml.utexas.edu/index.html

Stay connected on LinkedIn!

Happy Summer Break!

It’s been my pleasure teaching each of you.

Thank you for choosing to take this class.