Active Learning, Curriculum Learning, & Reinforcement Learning

Danna Gurari
University of Texas at Austin
Spring 2020

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

Review

- Last week:
  - Machine Learning for Unlabeled Data
  - Autoencoders
  - Clustering

- Assignments (Canvas):
  - Project outline with ML system prototype due yesterday
  - Final project video due in two weeks
  - Final project report due in three weeks

- Questions?

Paper Writing: Support

- Writing center: http://uwc.utexas.edu/
  - can schedule four individual 45-minutes consultation per month

- Tutoring:
  - https://utdirect.utexas.edu/apps/ugs/my/tutoring/student/tutoring-agreement/

Plagiarism: Definition

- Material from: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html

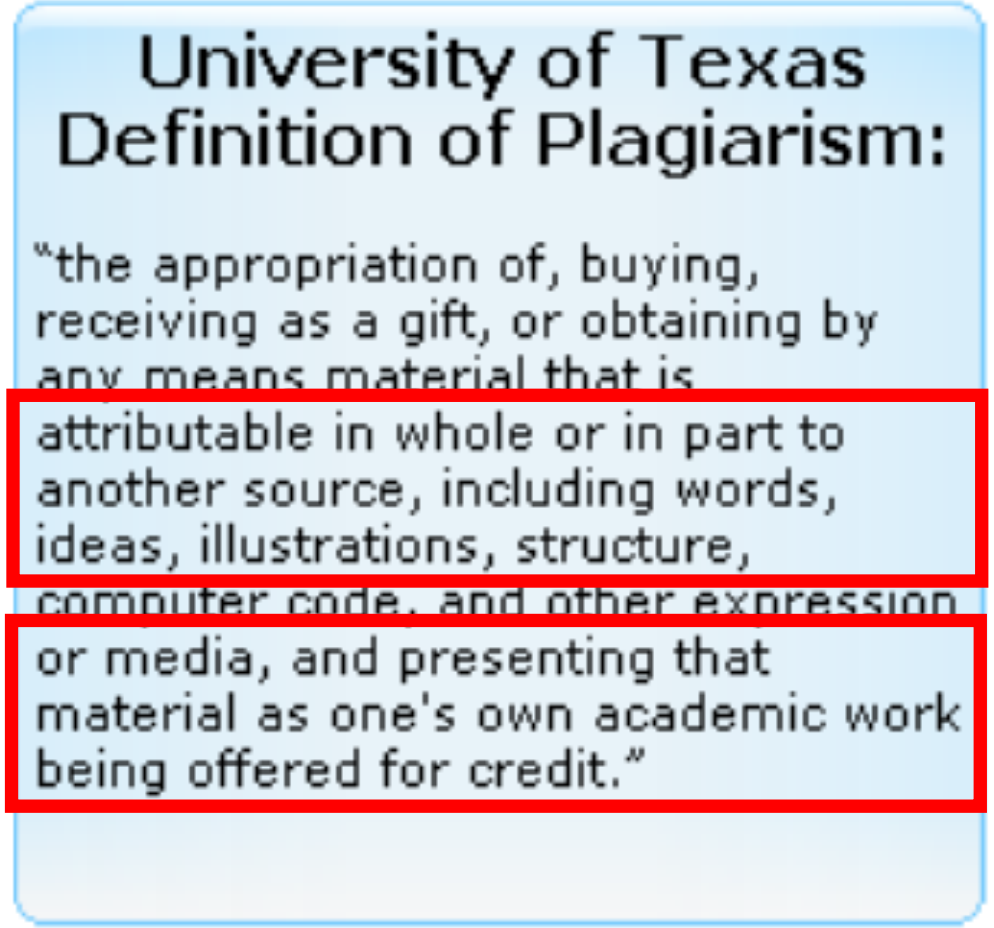

Plagiarism: Definition

- Material from: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html

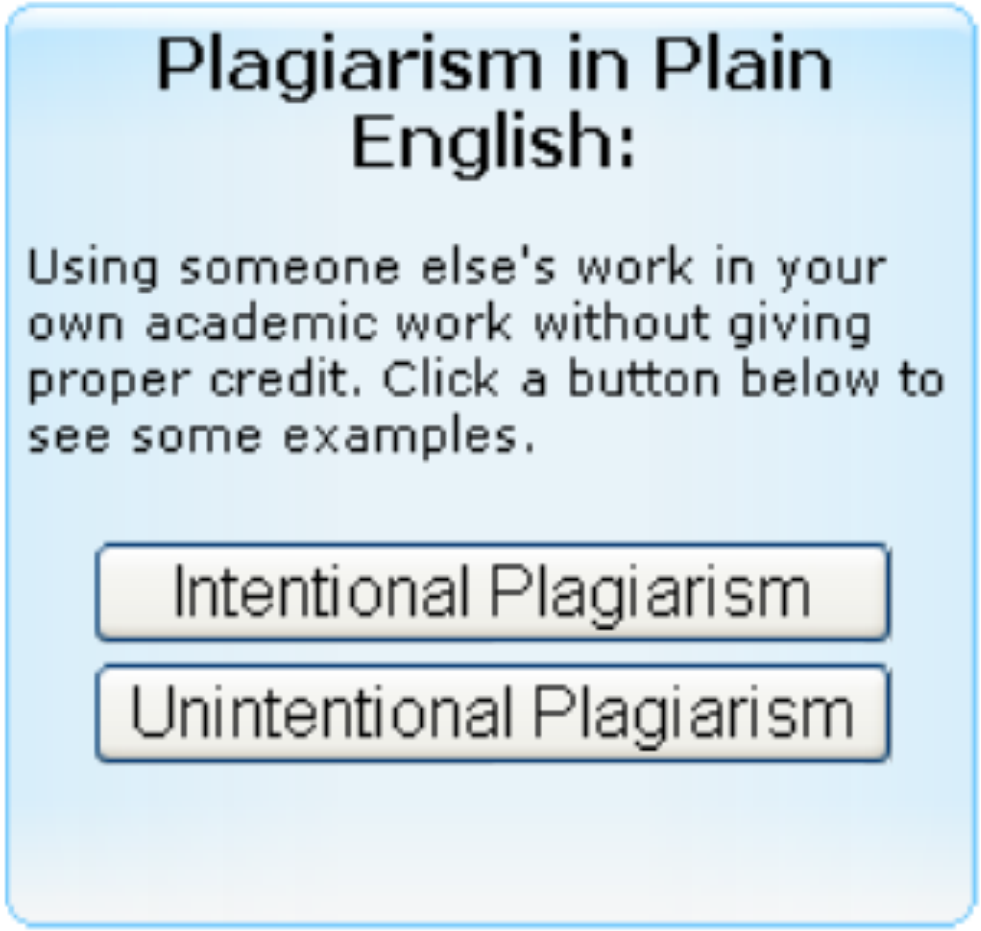

Plagiarism: Play It Safe, Give Credit Generously

- Material from: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html

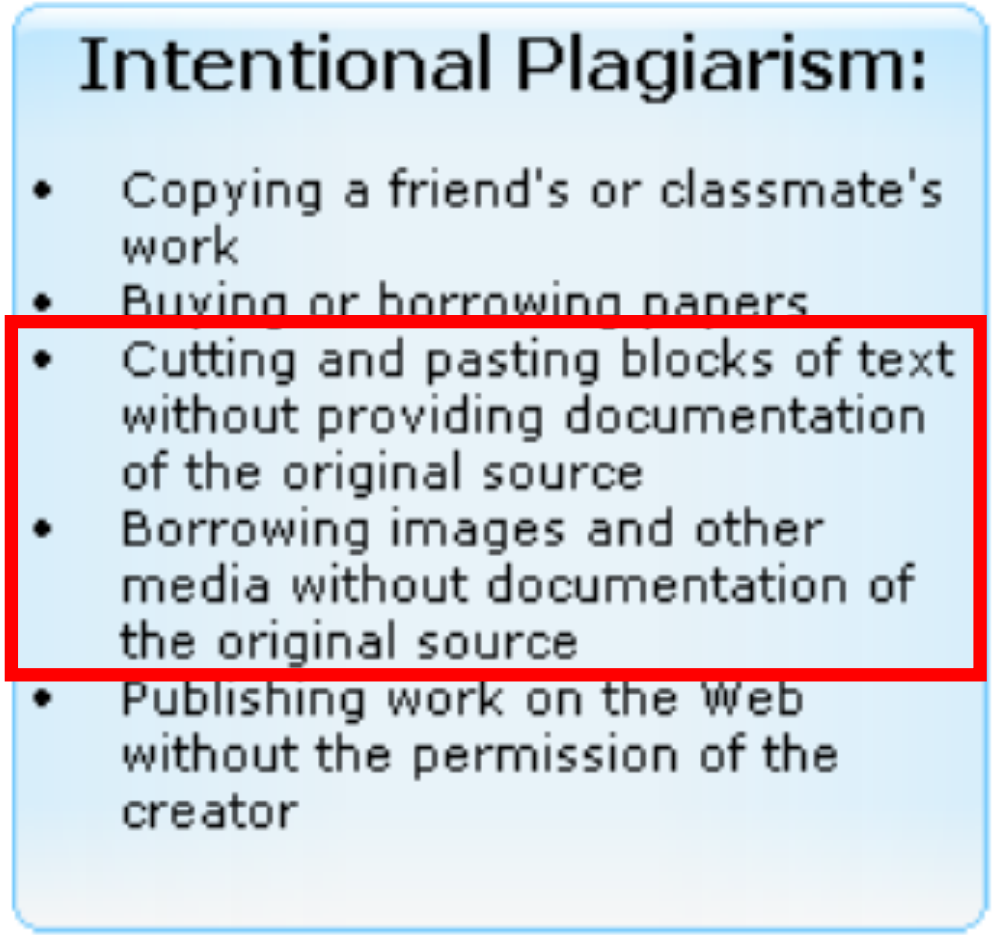

Plagiarism: Play It Safe, Give Credit Generously

- Material from: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html

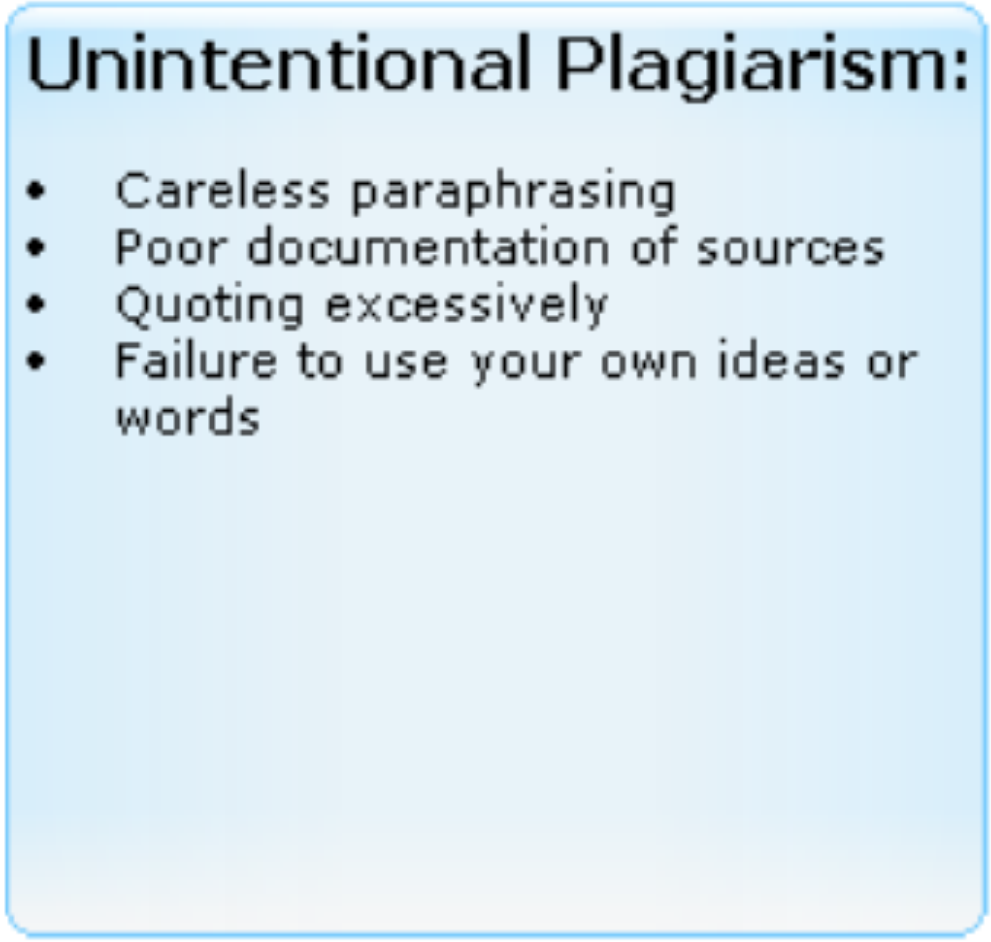

Plagiarism: Play It Safe, Give Credit Generously

- Material from: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html

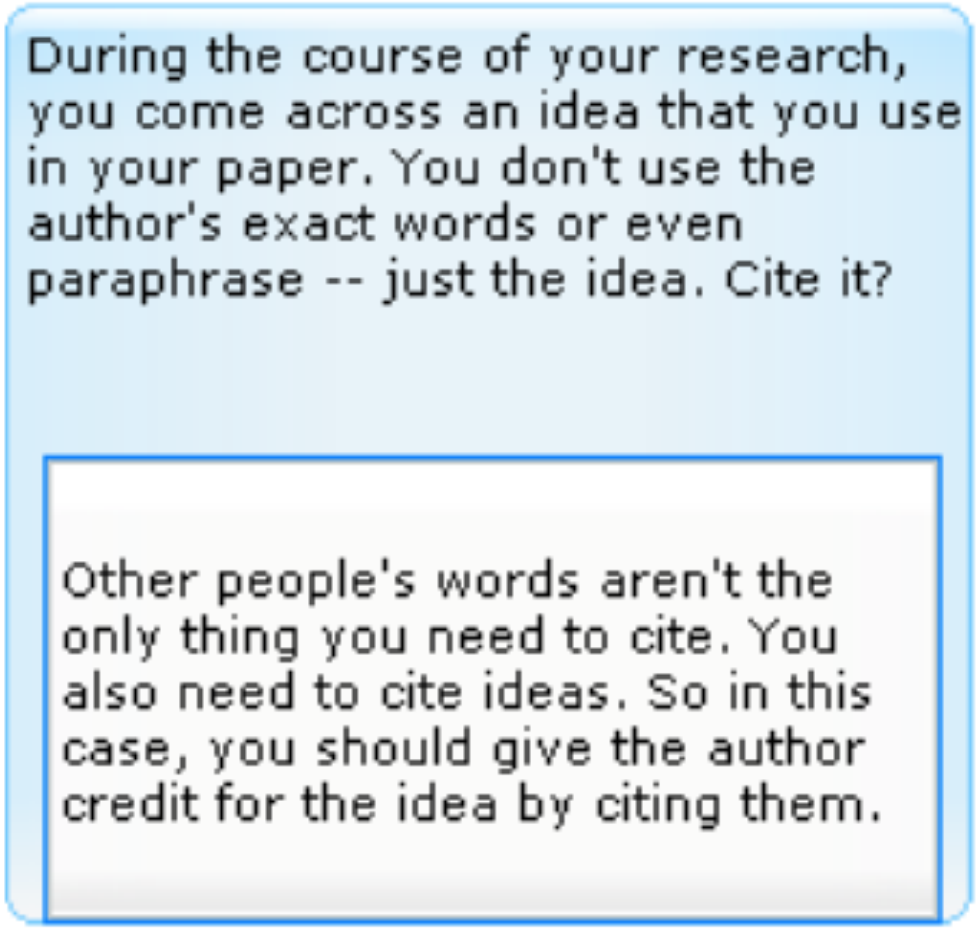

Plagiarism: Play It Safe, Give Credit Generously

- Material from: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html

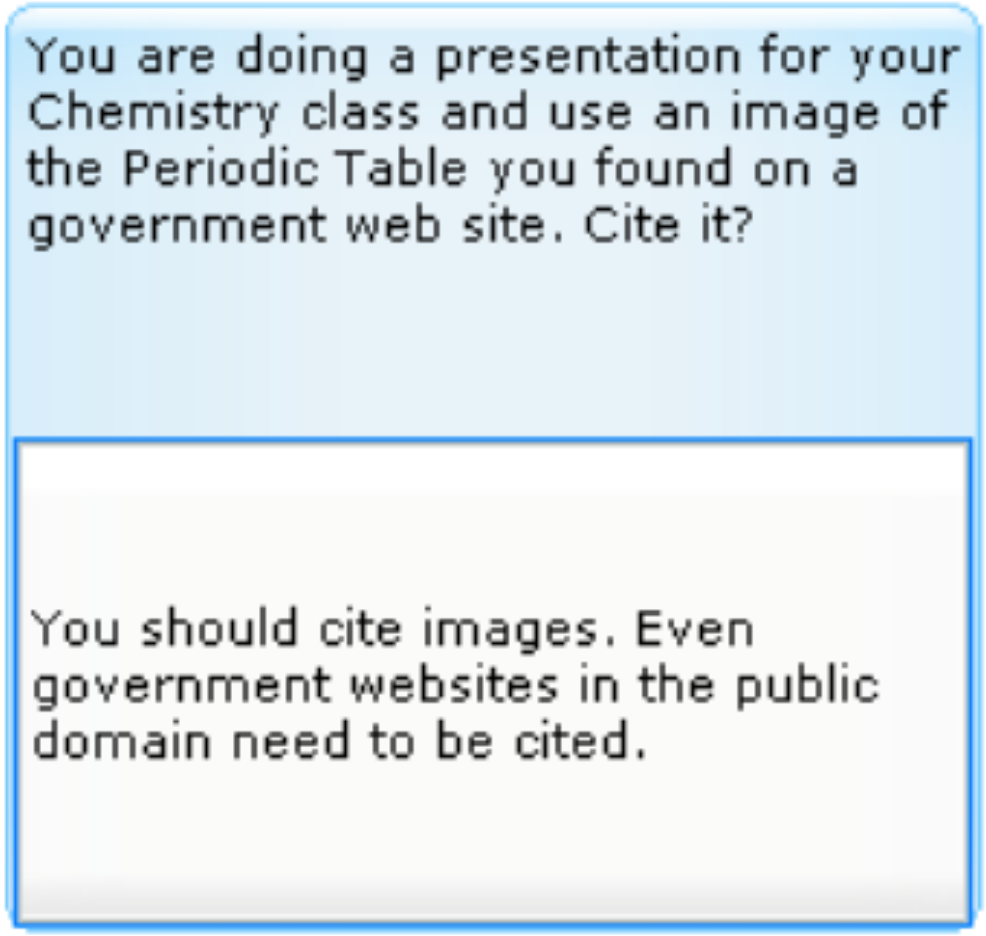

Plagiarism: Play It Safe, Give Credit Generously

- What can happen if you are accused of plagiarism?
  - Redo assignment
  - Receive a failing grade
  - Be suspended
  - Be expelled

- What resources can help you to avoid plagiarism?
  - Review: https://legacy.lib.utexas.edu/services/instruction/avoidplagiarism.html
  - Review: https://legacy.lib.utexas.edu/d7/sites/default/files/services/instruction/AvoidingPlagiarism_guide.pdf
  - Visit writing center: http://uwc.utexas.edu/

- Neither you (I believe) nor I have any desire to talk about plagiarism ☺
- Play it safe and give credit generously!!!

Give Credit Generously

- Idea: add credit page to your presentation for resources used
  - e.g., Microsoft Azure
  - e.g., freely-shared code/libraries
  - e.g., links to all images
  - …

Today’s Topics

- Active Learning
- Curriculum Learning
- Reinforcement Learning
- Guest: Dr. Cheryl Martin from Alegion

Today’s Topics

- Active Learning
- Curriculum Learning
- Reinforcement Learning
- Guest: Dr. Cheryl Martin from Alegion

Idea

Passive Learning

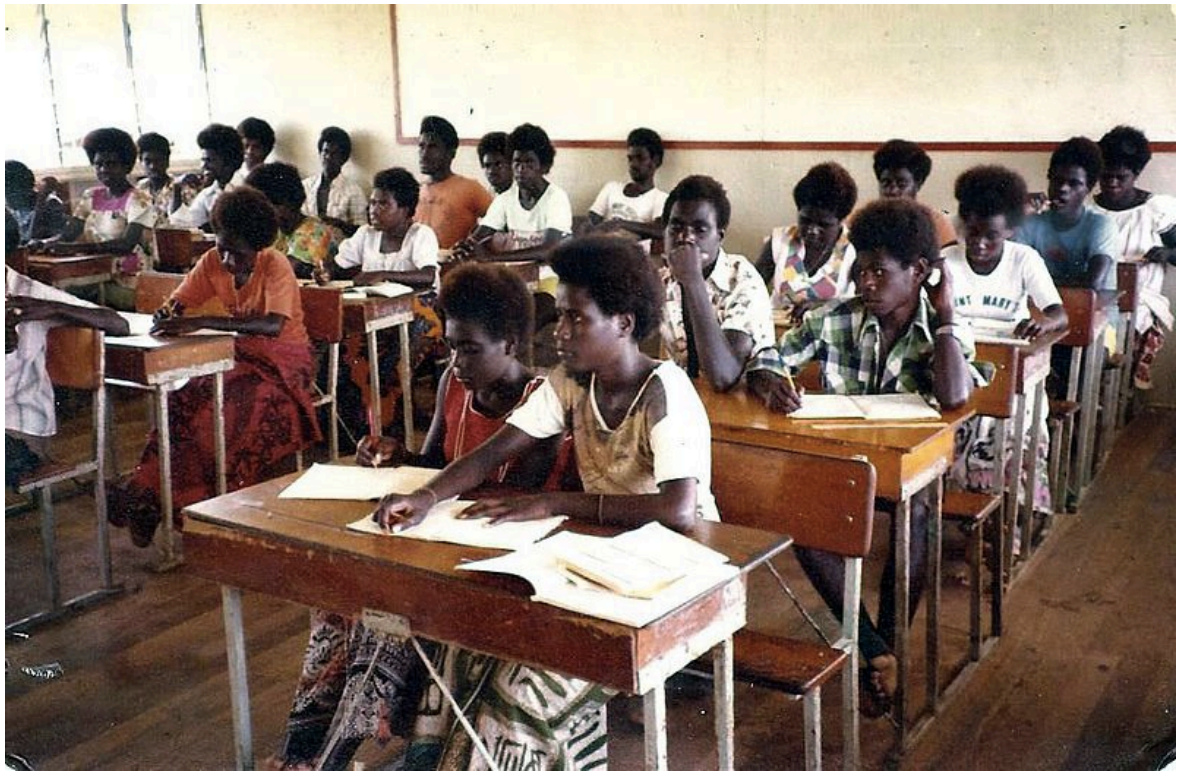

Active Learning

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

Passive Learning: Classical ML Approach

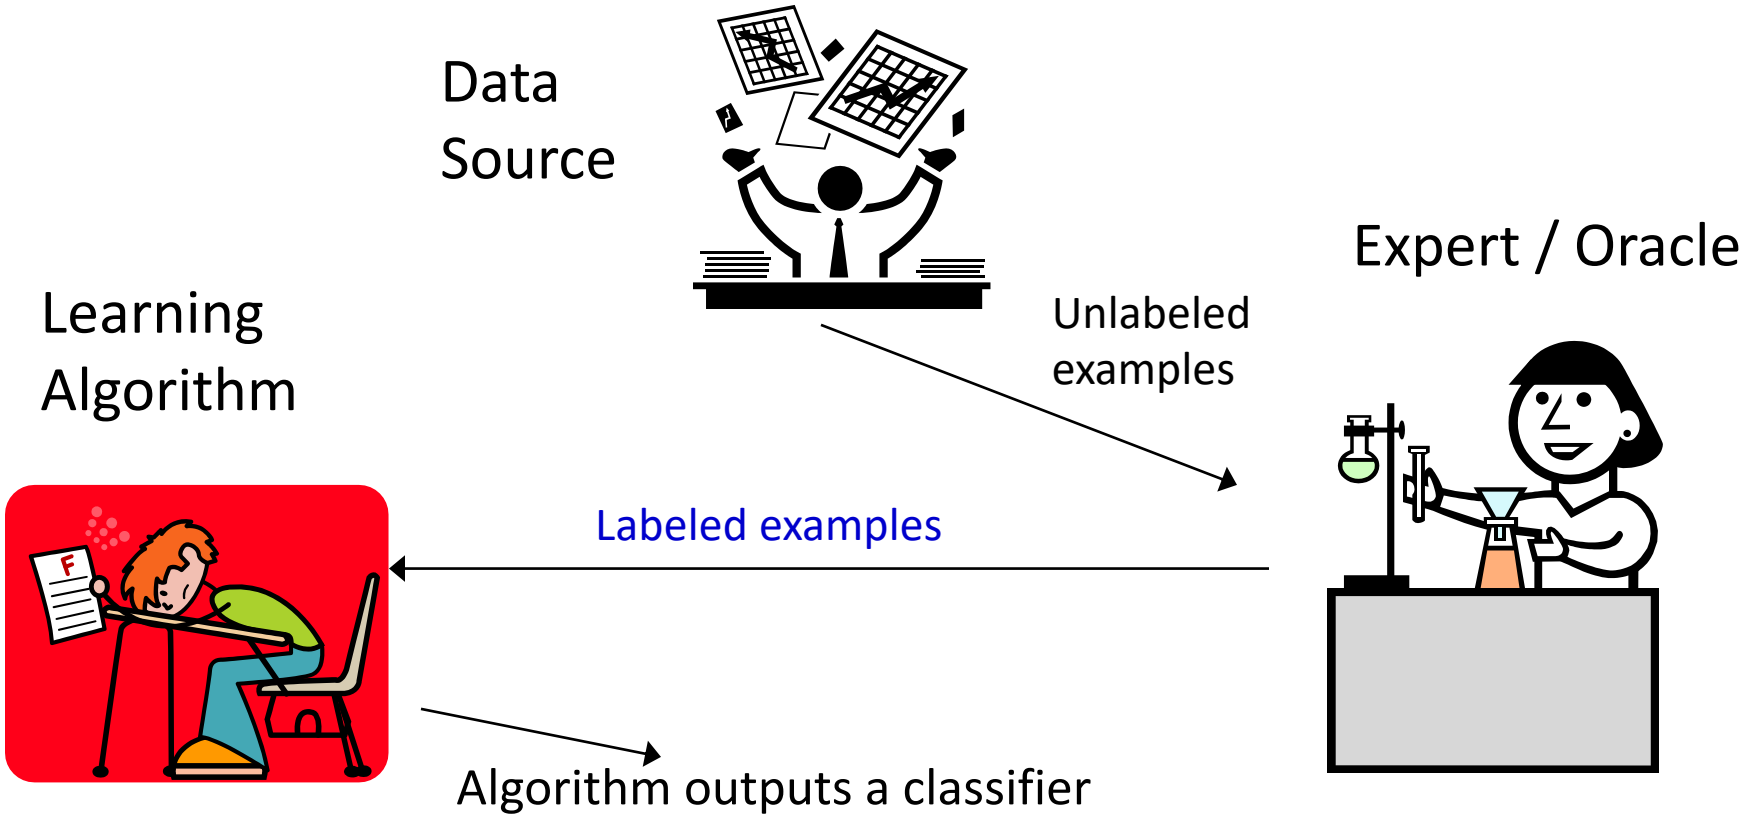

Slide Credit: http://www.cs.cmu.edu/~learning/talks-2007-spring/slides/mll0319.active_learning.ppt

Active Learning

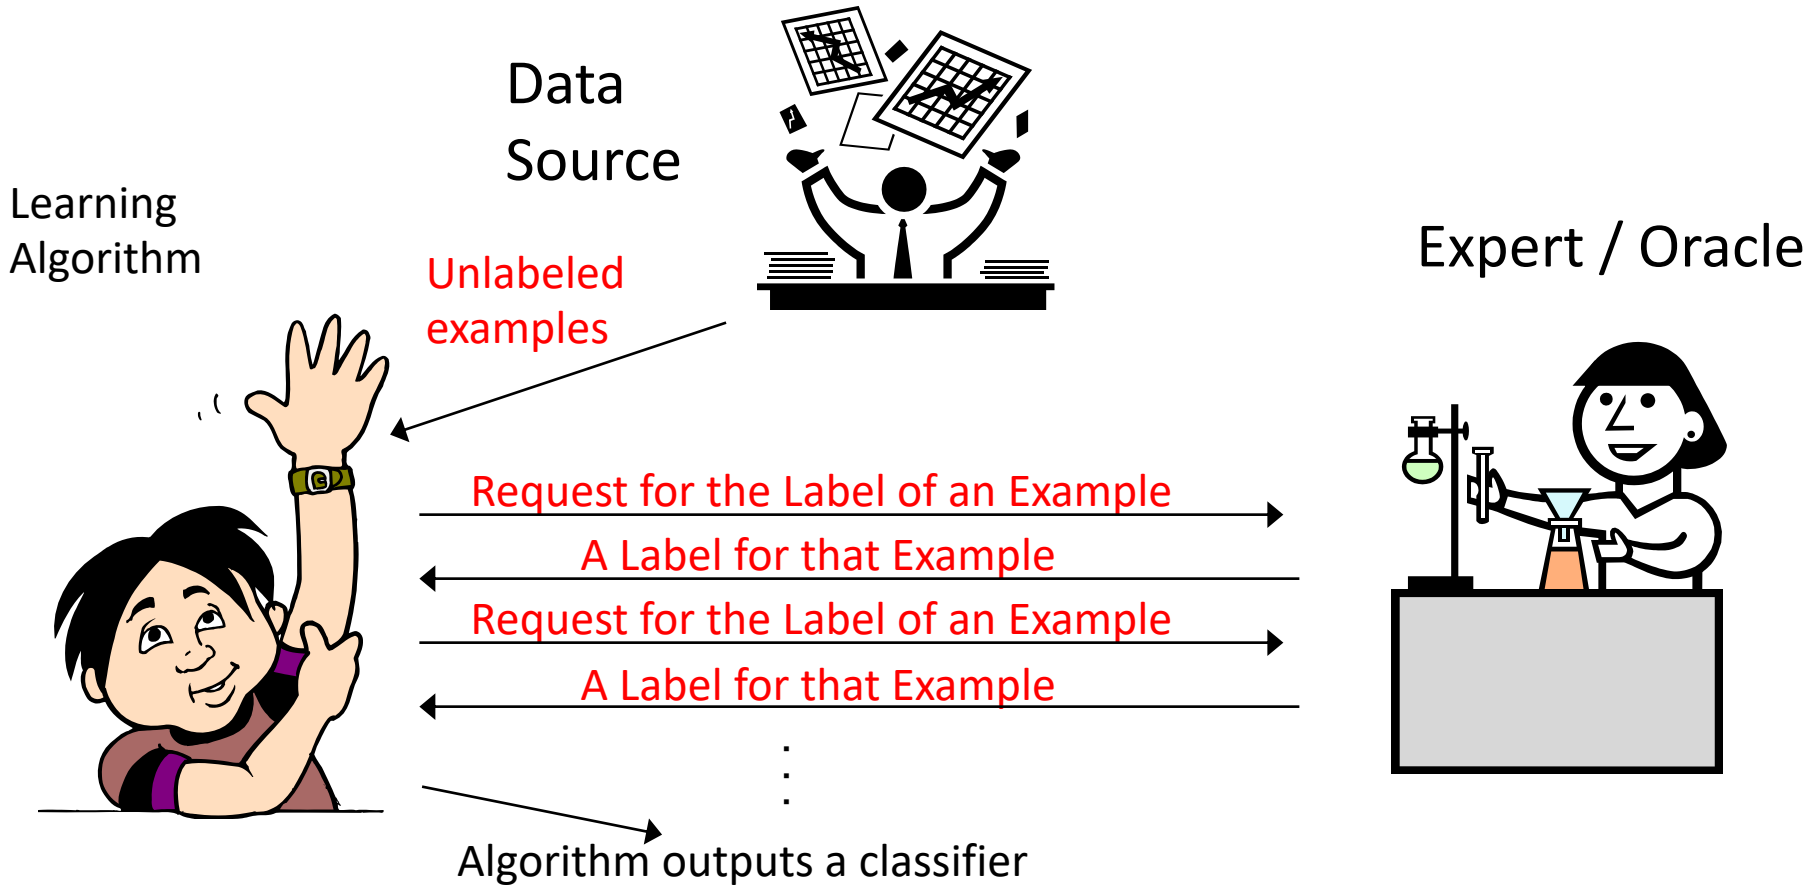

Slide Credit: http://www.cs.cmu.edu/~learning/talks-2007-spring/slides/mll0319.active_learning.ppt

Active Learning

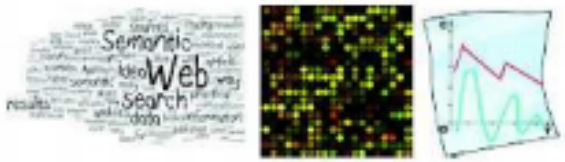

raw unlabeled data

$x_1, x_2, x_3, \ldots$

inspect the unlabeled data

request labels for selected data

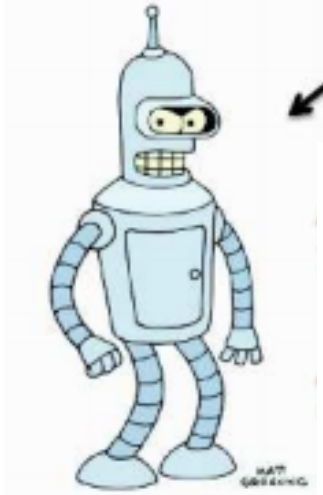

active learner
induces a classifier

expert / oracle
analyzes experiments to determine labels

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

Learning Curves: Active versus Passive Learning

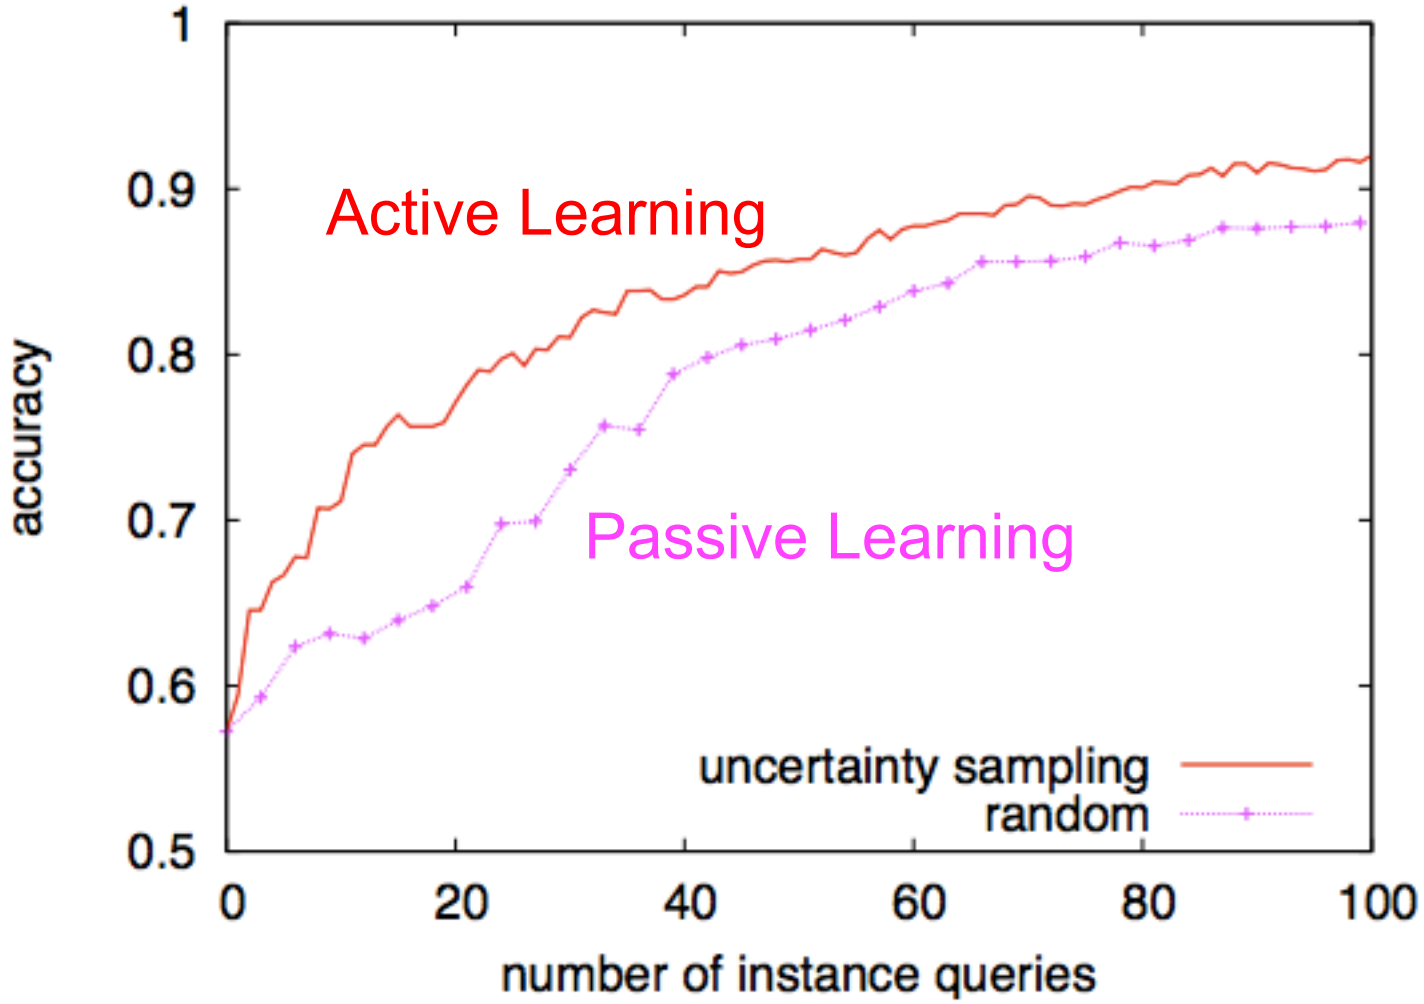

What are benefits of active learning?

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

Learning Curves: Active versus Passive Learning

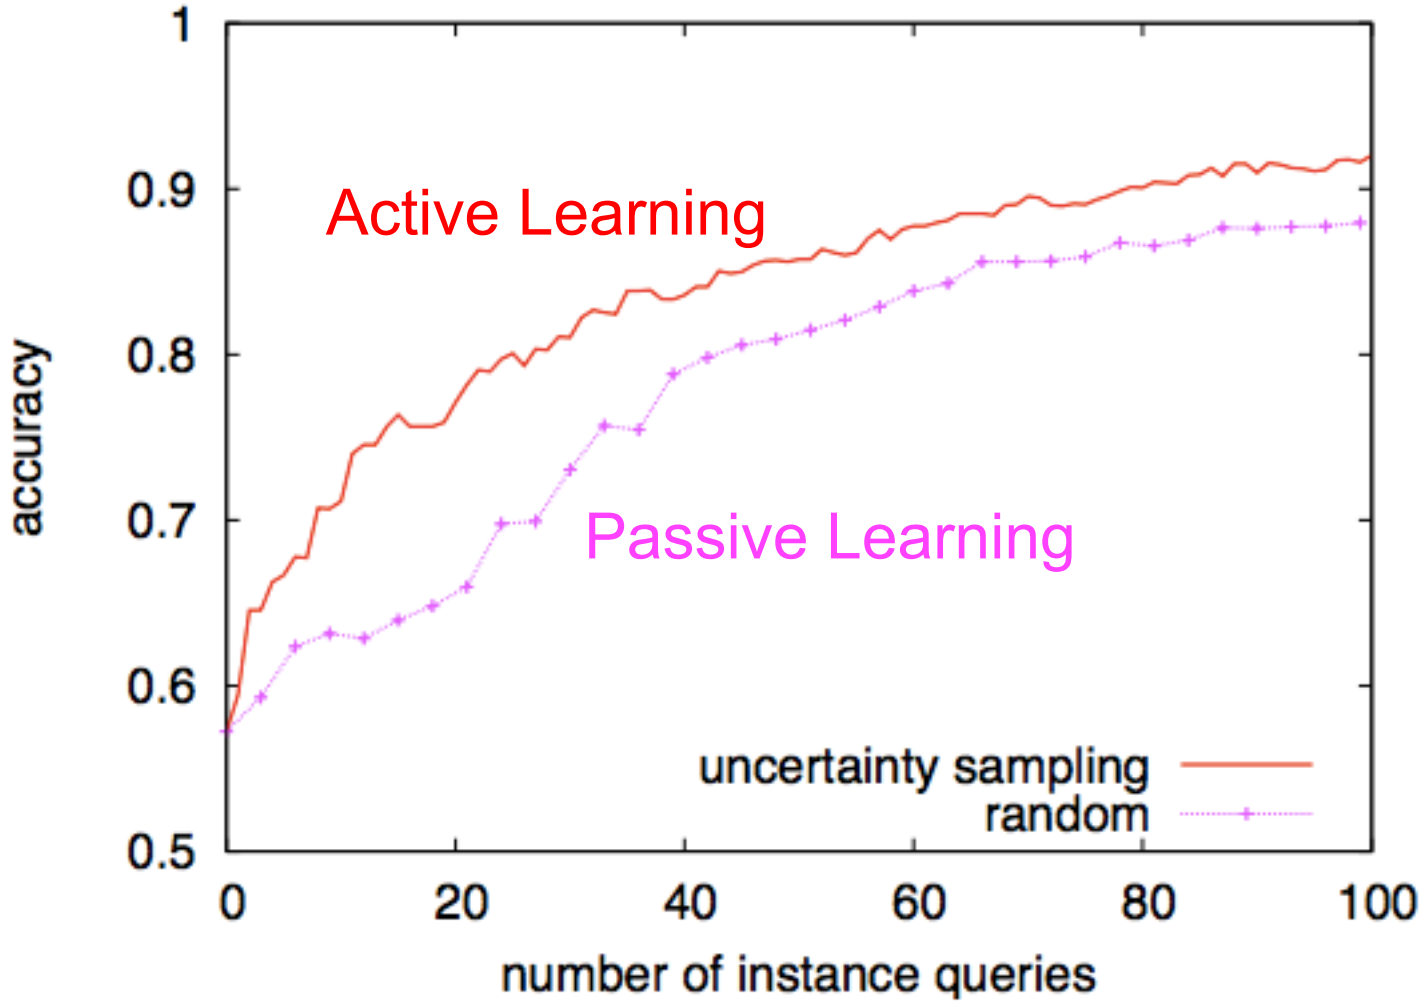

Machines can learn with fewer training instances if they ask questions.

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

Types of Active Learning

1. Stream-Based

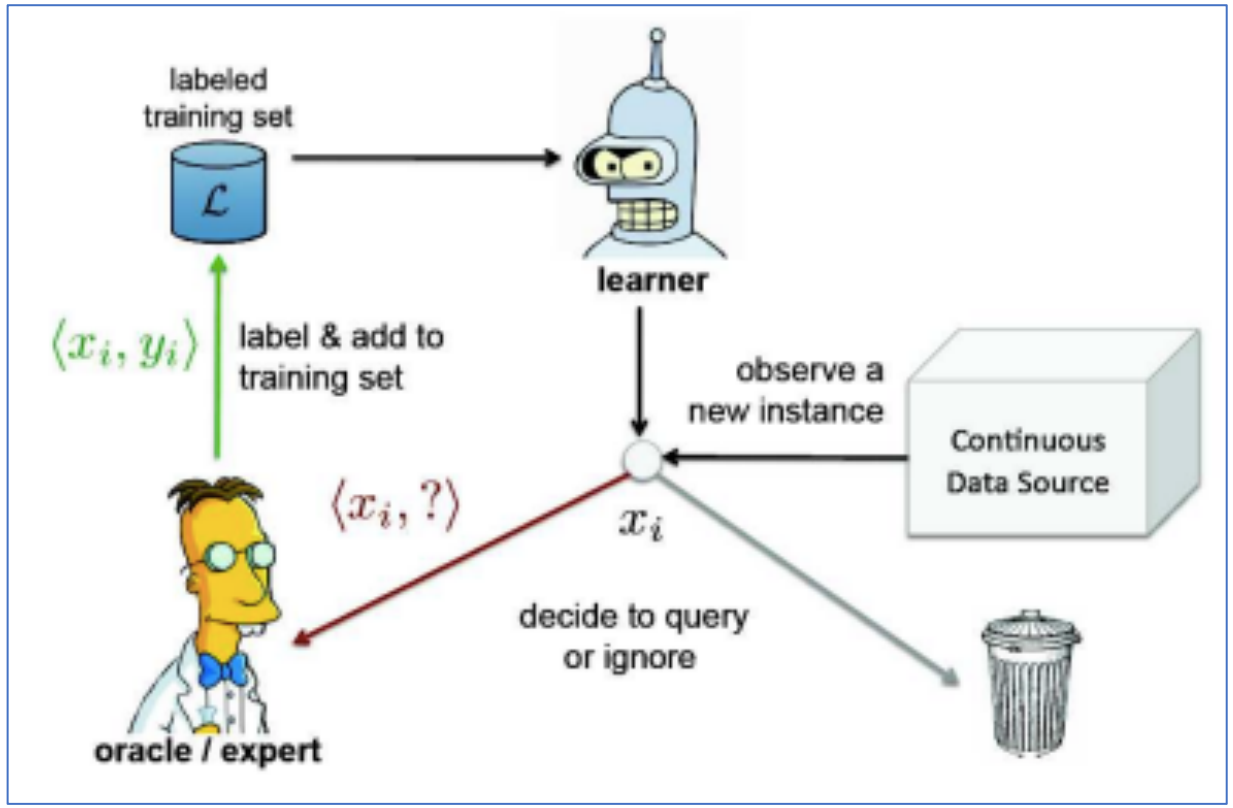

Consider one example at a time

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

Types of Active Learning

2. Pool-Based

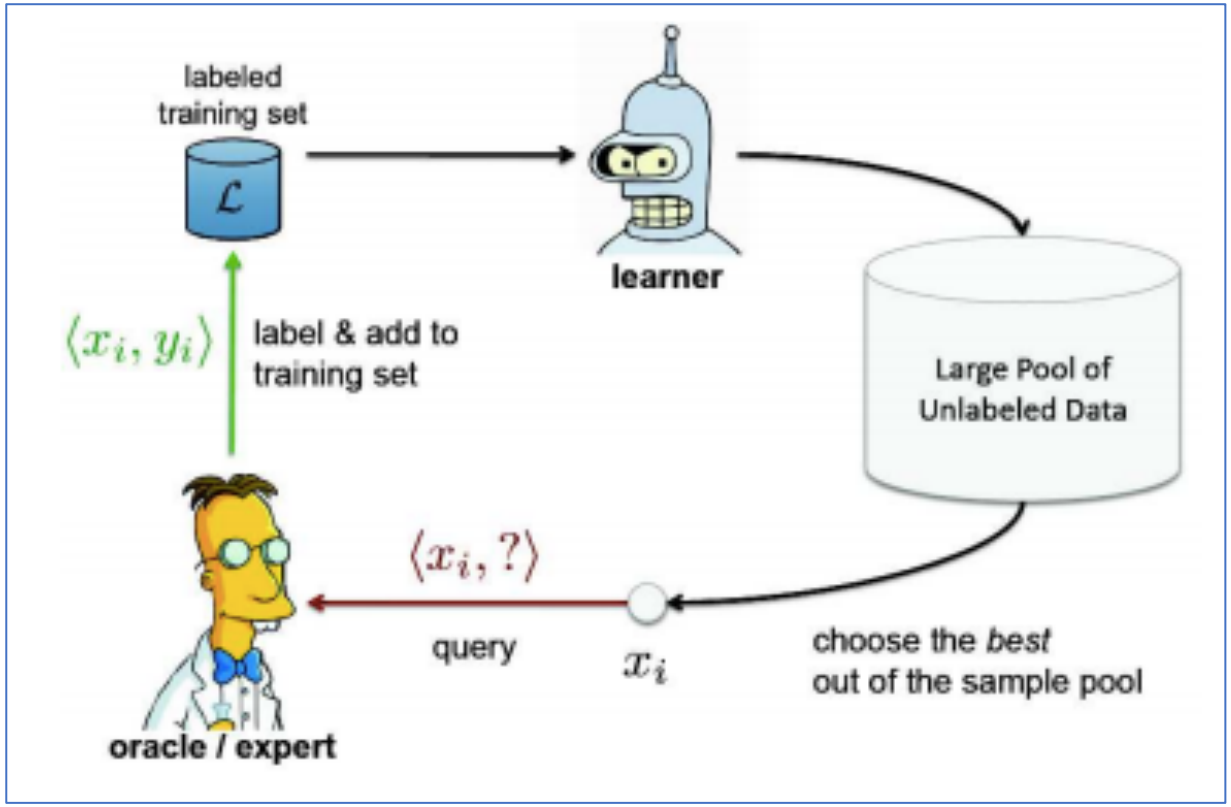

Consider many examples at a time

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

Types of Active Learning

Stream-Based

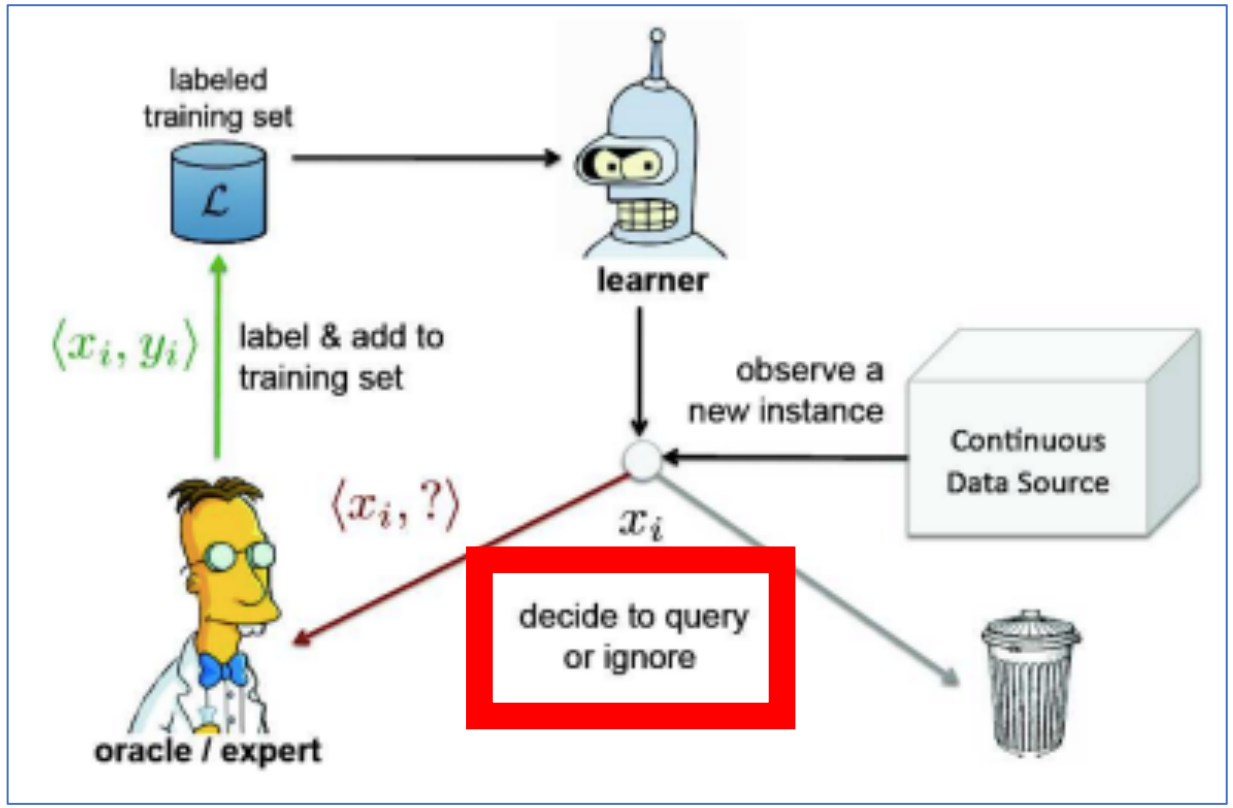

Consider one example at a time

Pool-Based

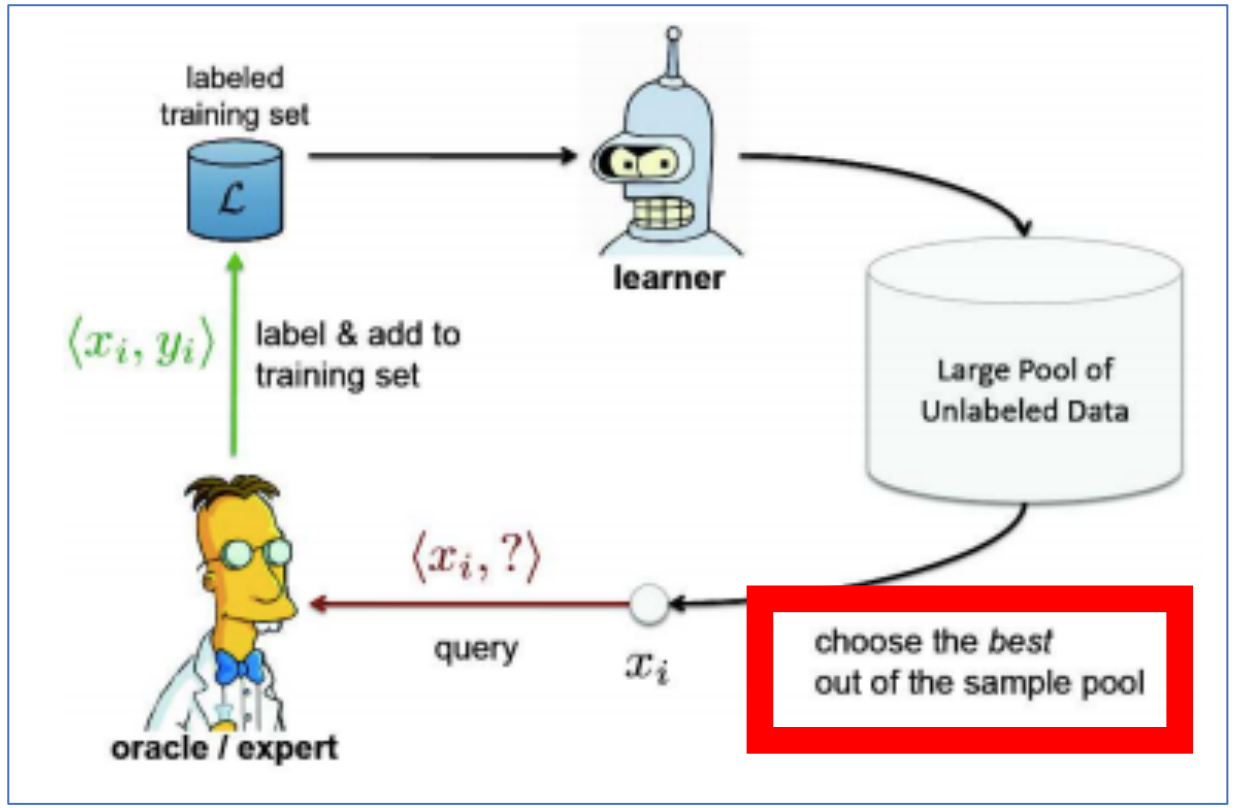

Consider many examples at a time

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

Active Learning Approach

- Active Learning proceeds in rounds
- Each round has a current model (learned using the labeled data seen so far)
- The current model is used to assess informativeness of unlabeled examples
  - .. using one of the query selection strategies
- The most informative example(s) is/are selected
- The labels are obtained (by the labeling oracle)
- The (now) labeled example(s) is/are included in the training data
- The model is re-trained using the new training data

Slide 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

Uncertainty Sampling: e.g., Logistic Classifier

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

True Representation (Assume Labels Are Not Known)

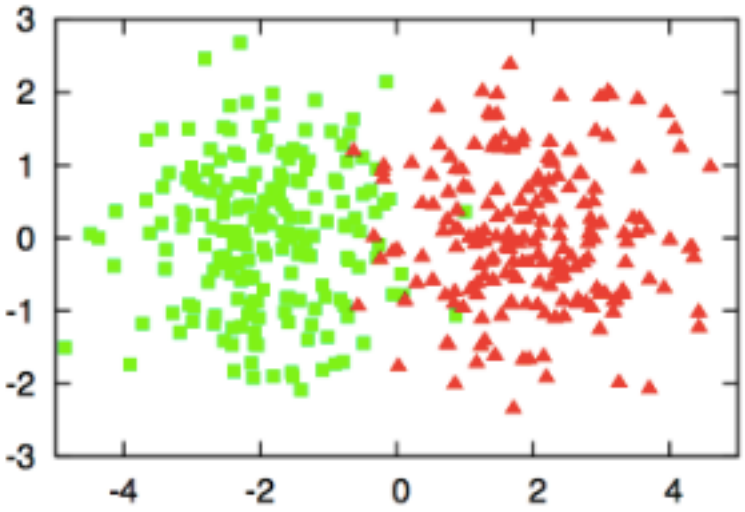

Passive Learner (Random Selection)

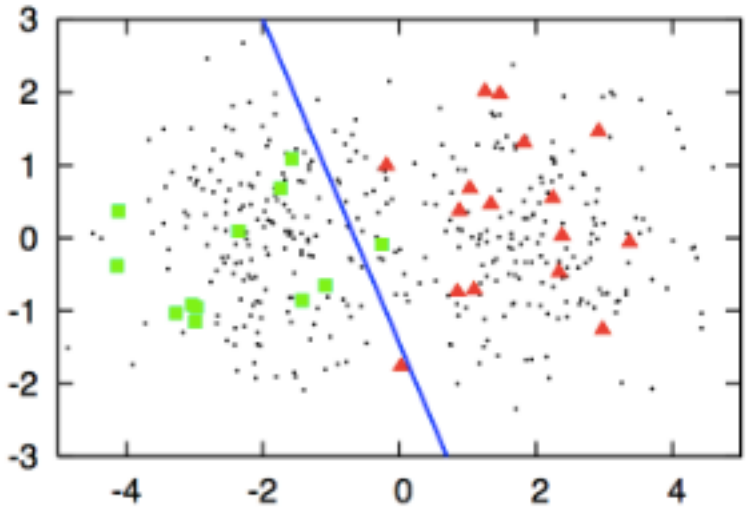

Active Learner (Uncertainty Sampling)

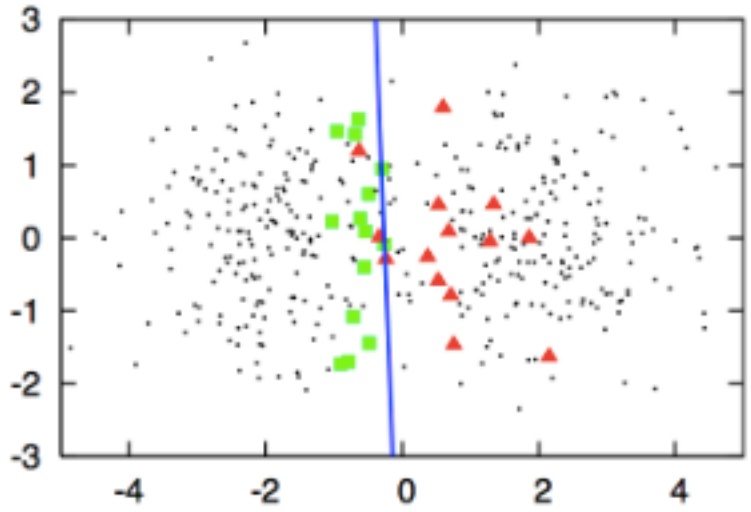

Image Credit: http://burrsettles.com/pub/settles.activelearning.pdf

Uncertainty Sampling: e.g., SVM Classifier

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

e.g., strategy 1: request the label of the example closest to the current separator.

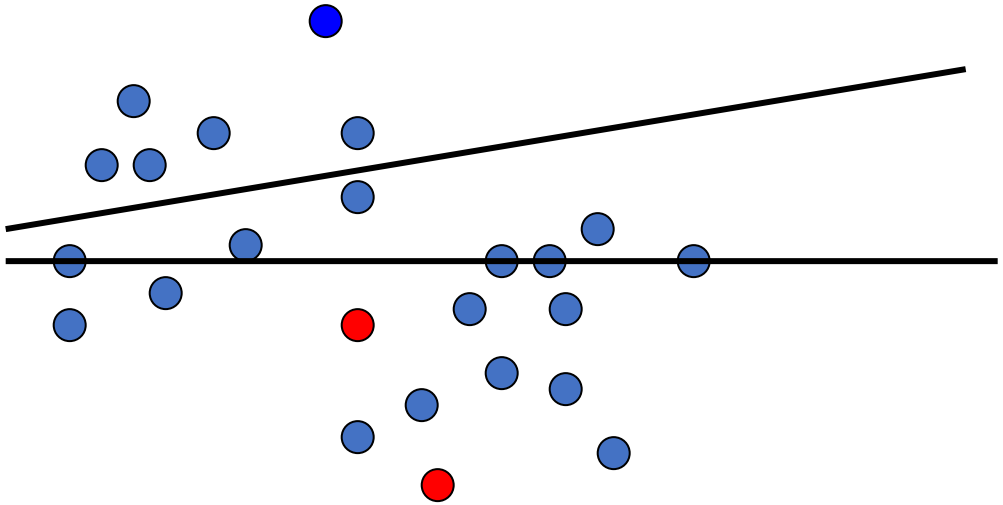

Slide Credit: http://www.cs.cmu.edu/~learning/talks-2007-spring/slides/mll0319.active_learning.ppt

Query By Committee

Query instance(s) different classifiers disagree most about.

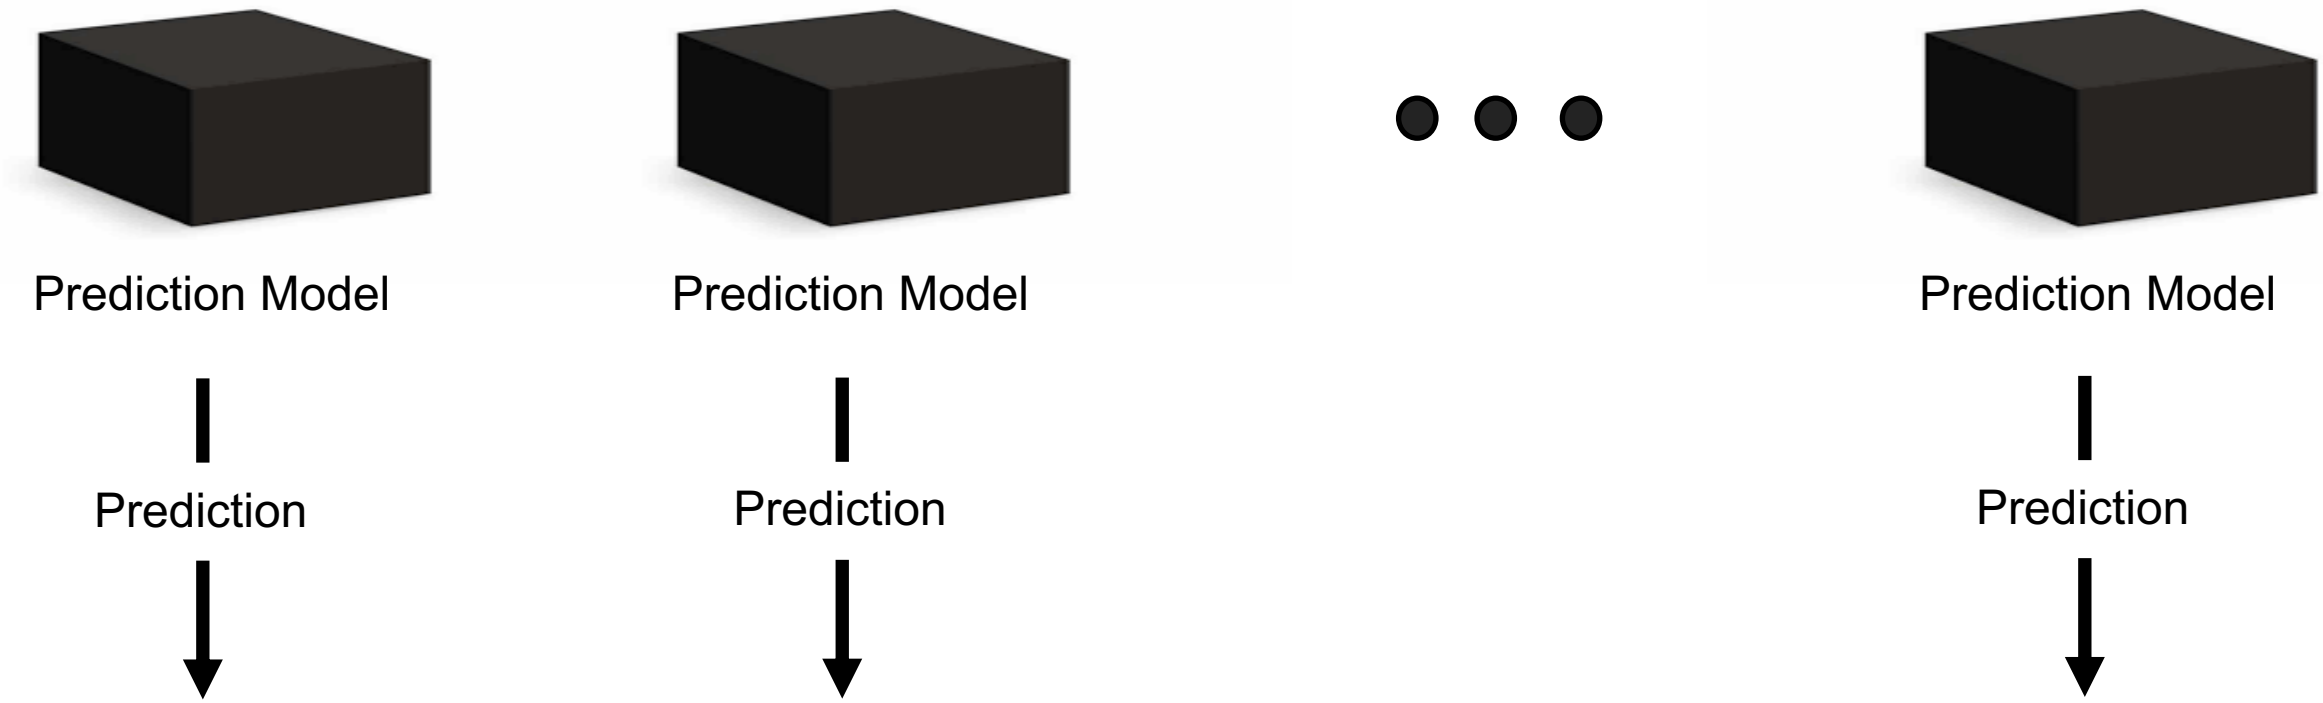

Image Credit: http://burrsettles.com/pub/settles.activelearning.pdf

Group Discussion:

Assume you are hired to build a new face recognition service. How would you design an active learning approach to train an accurate machine learning algorithm while collecting training data efficiently?

Gfycat's facial recognition software can now recognize individual members of K-pop band Twice, but in early tests couldn't distinguish different Asian faces. GFYCAT

https://www.wired.com/story/how-coders-are-fighting-bias-in-facial-recognition-software/

Today’s Topics

- Active Learning
- Curriculum Learning
- Reinforcement Learning
- Guest: Dr. Cheryl Martin from Alegion

Idea

How to teach machines to learn faster?

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

Random Order of Examples

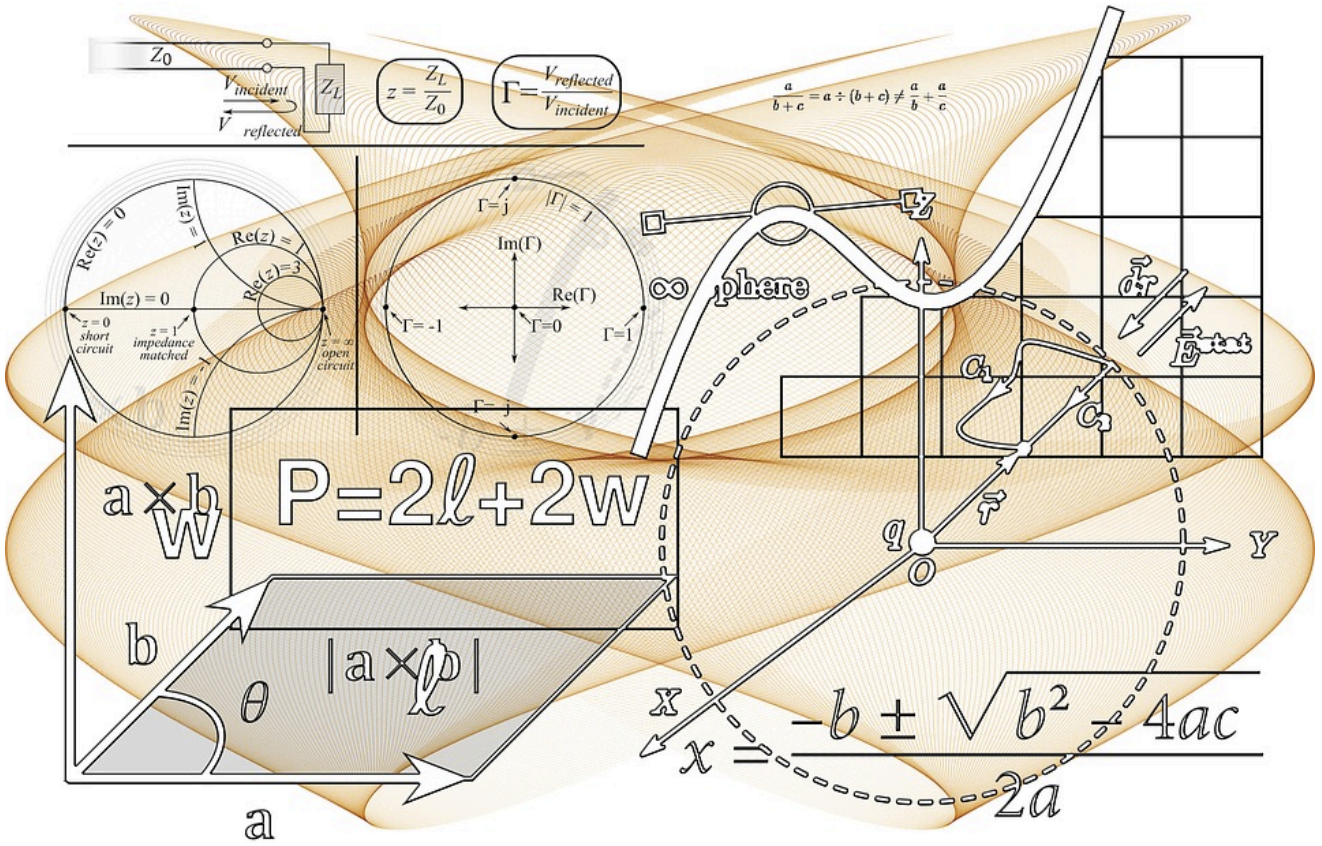

Meaningful Order of Examples

Big Book of Math; Dinah Zike

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

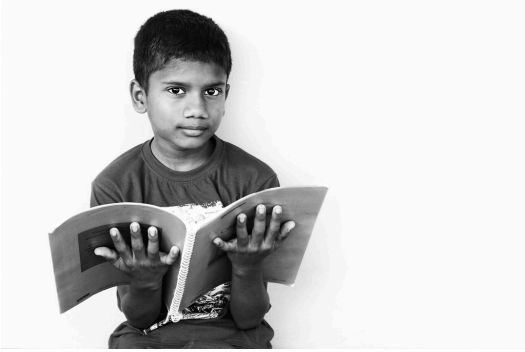

Random Order of Examples

Meaningful Order of Examples

Idea: Teach Machines As We Teach Humans

Curriculum

Train with simpler examples first and progressively harder examples over time.

Jeffrey L. Elman. Learning and development in neural networks: The importance of starting small. Cognition, 1993.

Learning Curves: Shape Variability

Artificial data: classify images into 3 shapes (rectangle, ellipse, triangle)

Input: 32×32 grey-scale image

(Less shape variability)

Easy:
(Basic)

Hard:
(Geom)

Yoshua Bengio et al.; Curriculum Learning; 2009.

Learning Curves: Shape Variability

Artificial data: classify images into 3 shapes (rectangle, ellipse, triangle)

- Training: 3-layer neural network with BasicShapes or GeomShapes (10,000 examples)
- Testing: GeomShapes

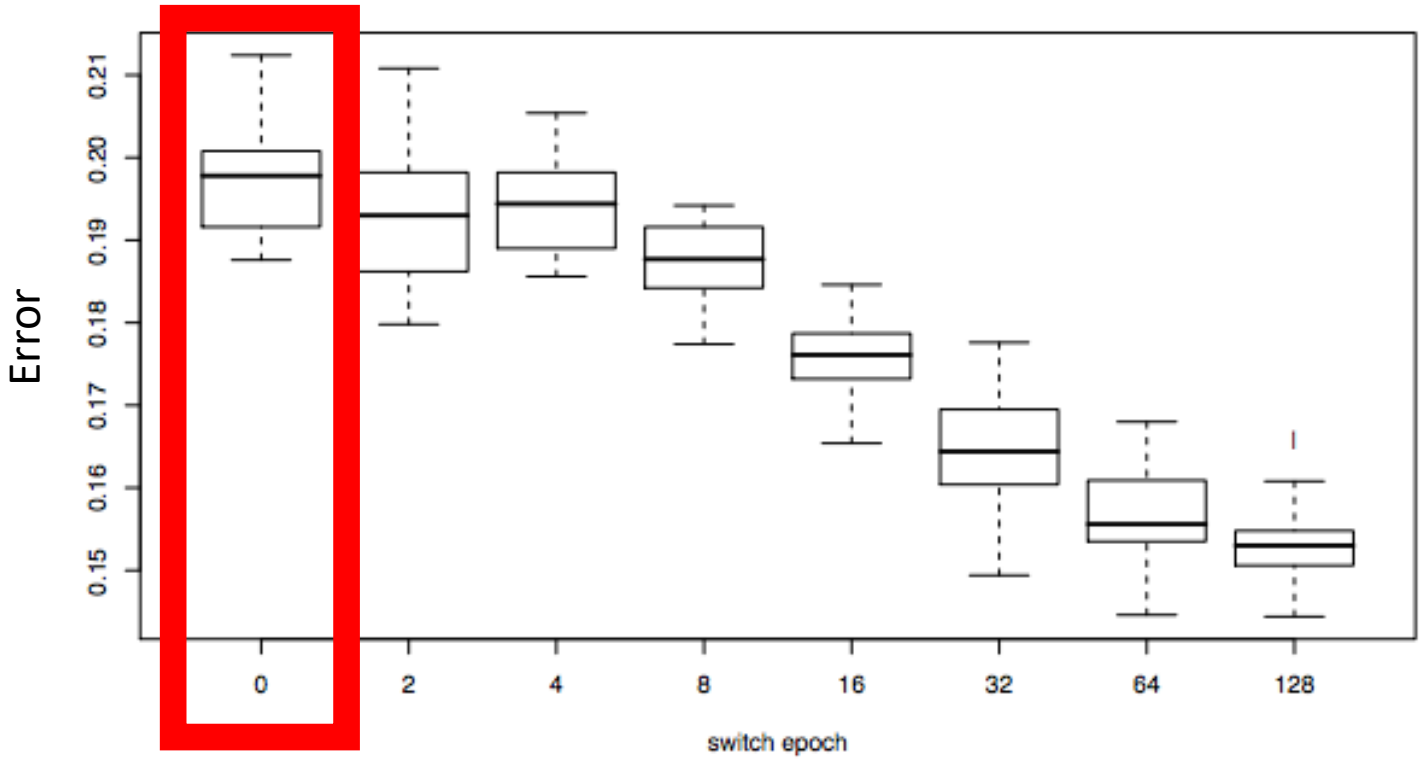

No curriculum

What are benefits of curriculum learning?

How long should the algorithm train with easy examples before switching to difficult examples?

Yoshua Bengio et al.; Curriculum Learning; 2009.

Learning Curves: Word Prediction

Wikipedia: predict next word in a sentence

- Curriculum: grow vocabulary size; 5k most frequent words, then 10k most frequent words, etc
- Target: final vocabulary size is 20, 000 words

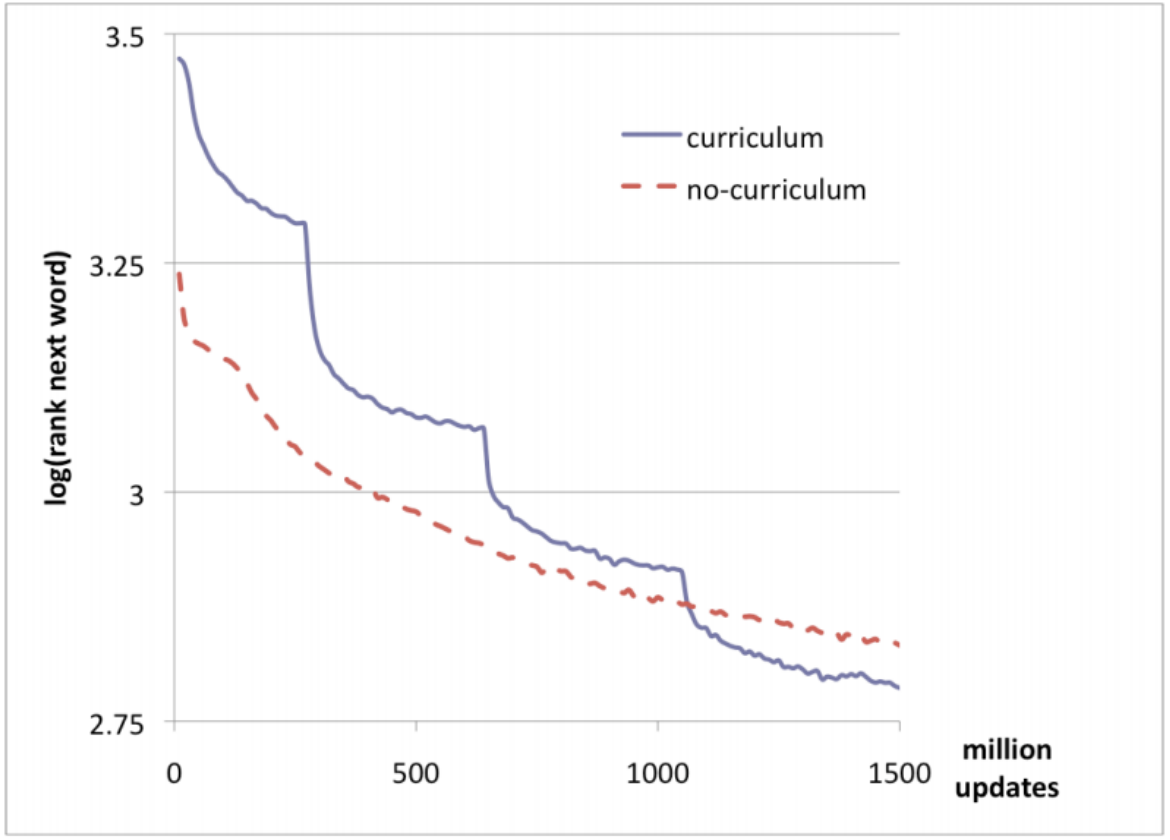

What are benefits of curriculum learning?

How long should the algorithm train with easy examples before switching to difficult examples?

Yoshua Bengio et al.; Curriculum Learning; 2009.

Group Discussion: Curriculum Learning

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

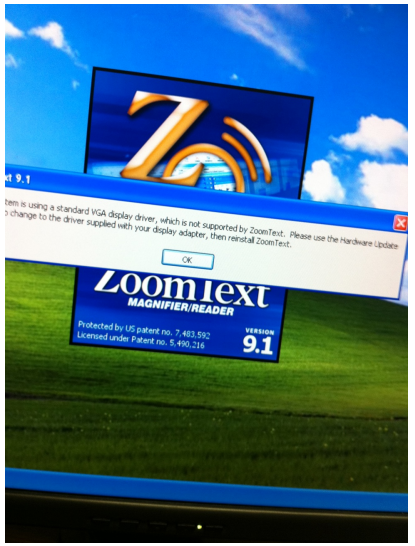

Questions

1. What criteria should be used to order examples?
2. What batches would you use when changing the available data?
3. How often would you make updates?

Today’s Topics

- Active Learning
- Curriculum Learning
- Reinforcement Learning
- Guest: Dr. Cheryl Martin from Alegion

Reinforcement Learning Overview

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

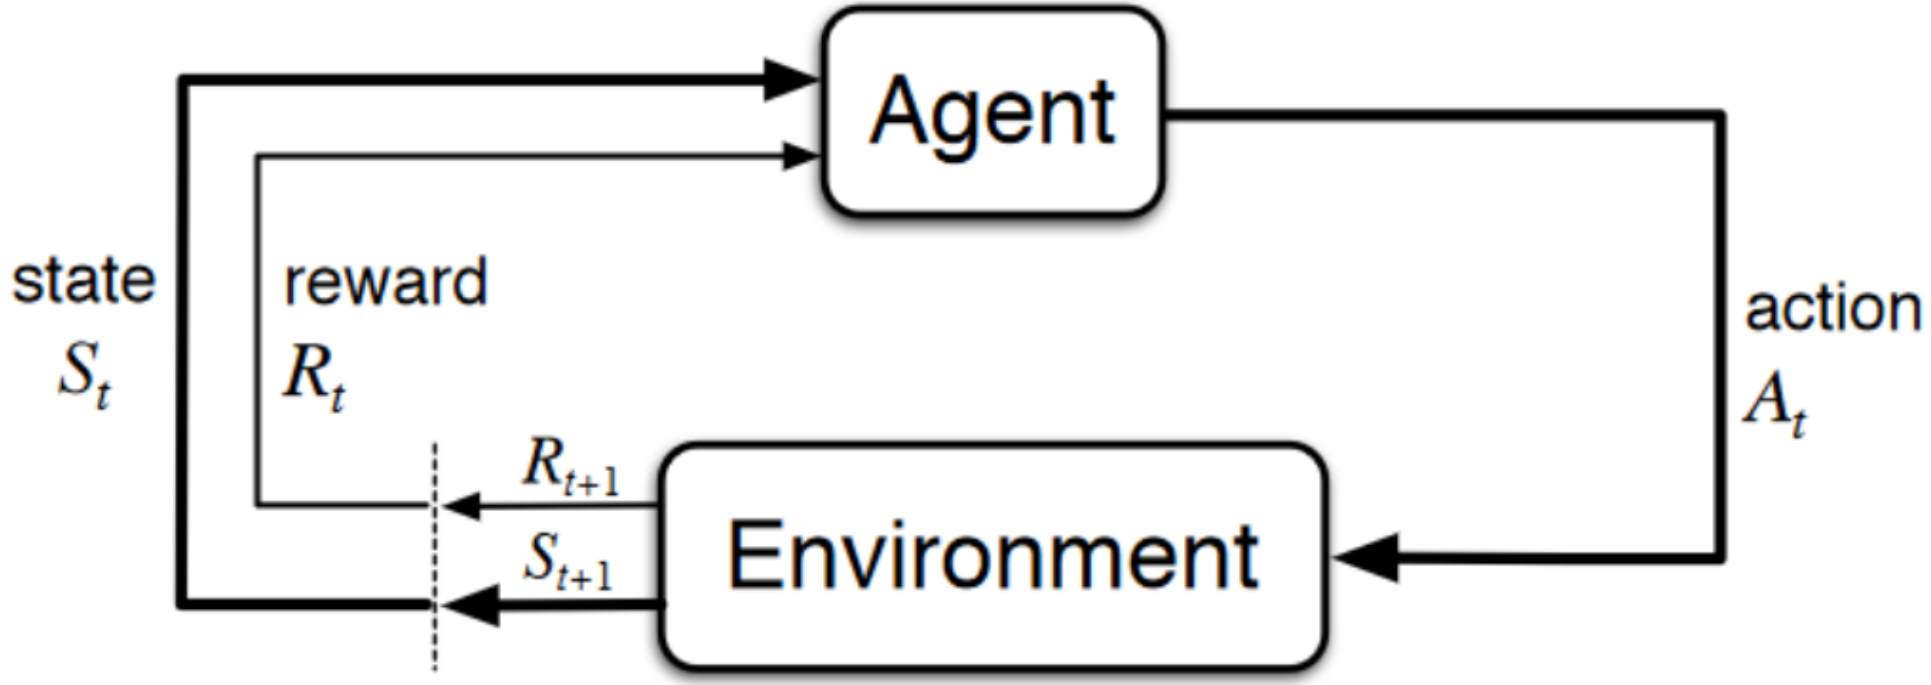

Figure Credit: https://towardsdatascience.com/applications-of-reinforcement-learning-in-real-world-1a94955bcd12

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

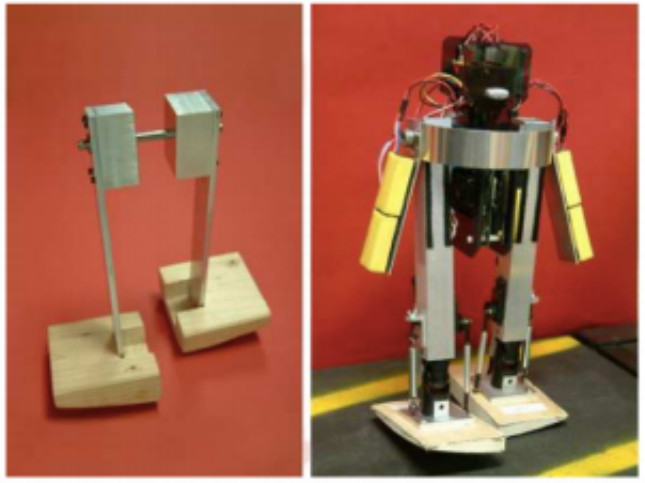

Reinforcement Learning Applications

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

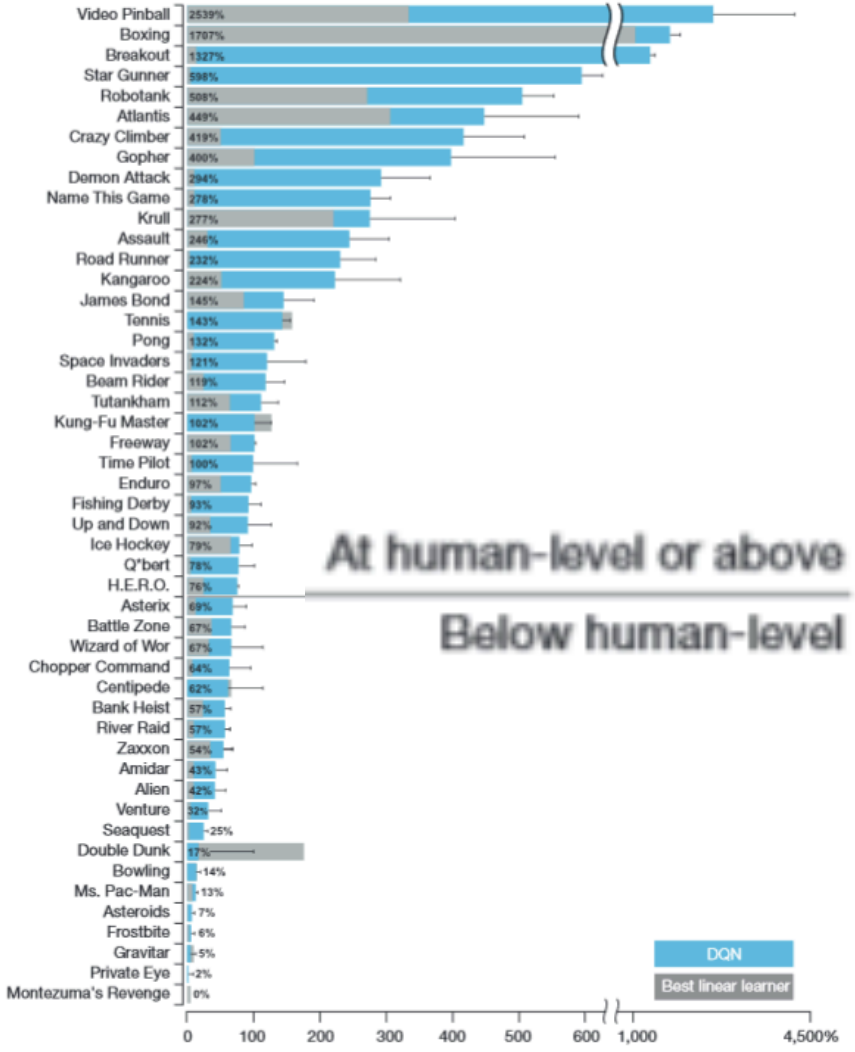

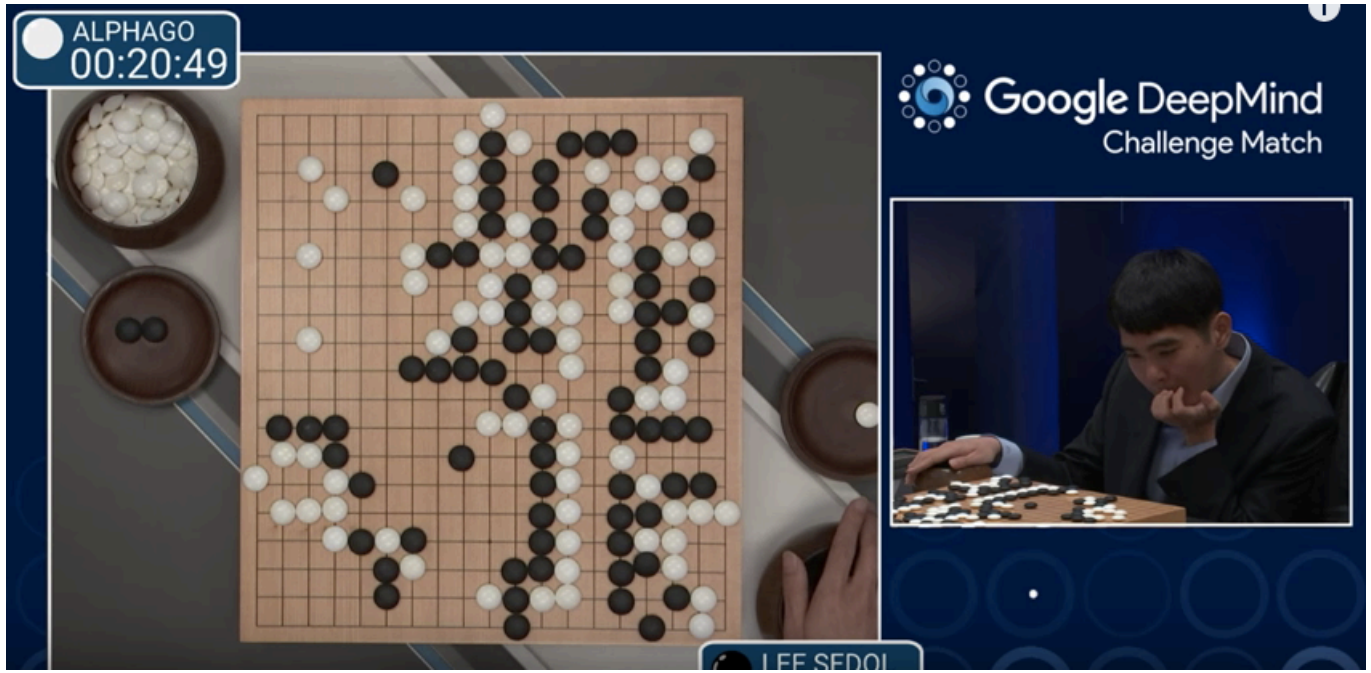

https://www.tastehit.com/blog/google-deepmind-alphago-how-it-works/

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

e.g., Pong Game - Learning Example

Move “up” or “down”

-1 if missed the ball

+1 reward if ball goes past opponent

0 otherwise

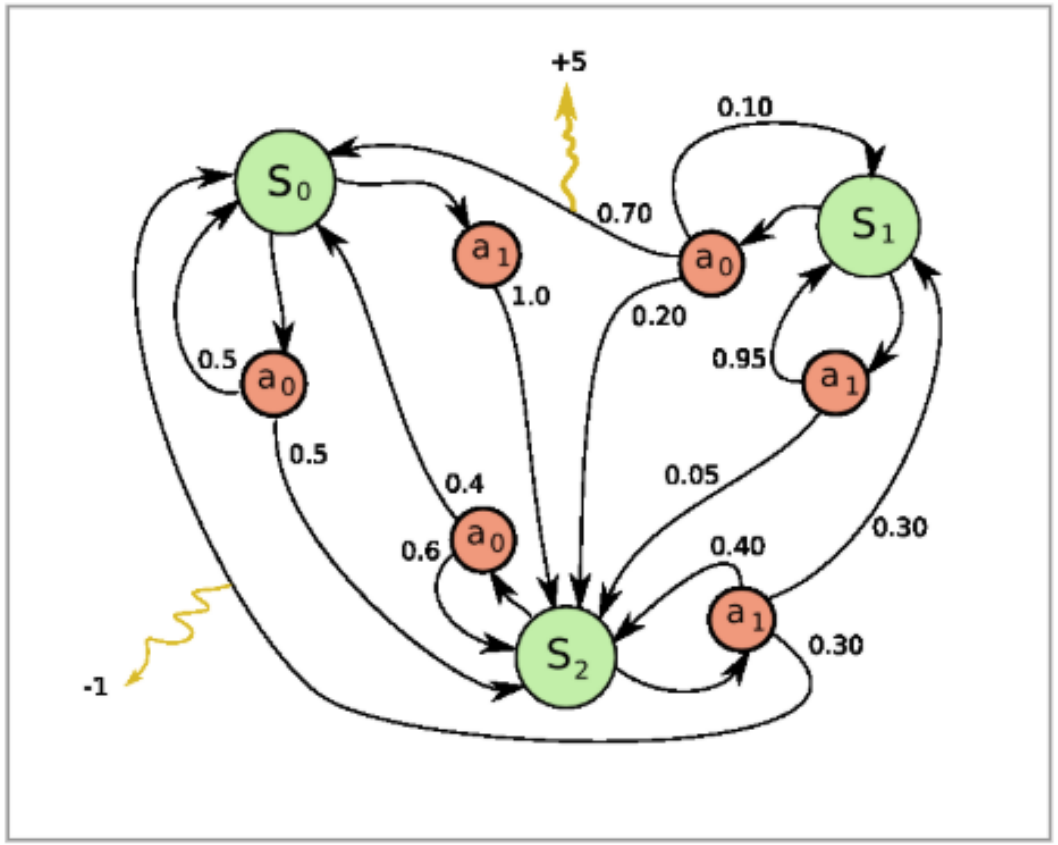

http://karpathy.github.io/2016/05/31/rl/

e.g., Pong Game: Policy Network

Implements our player (or “agent”)

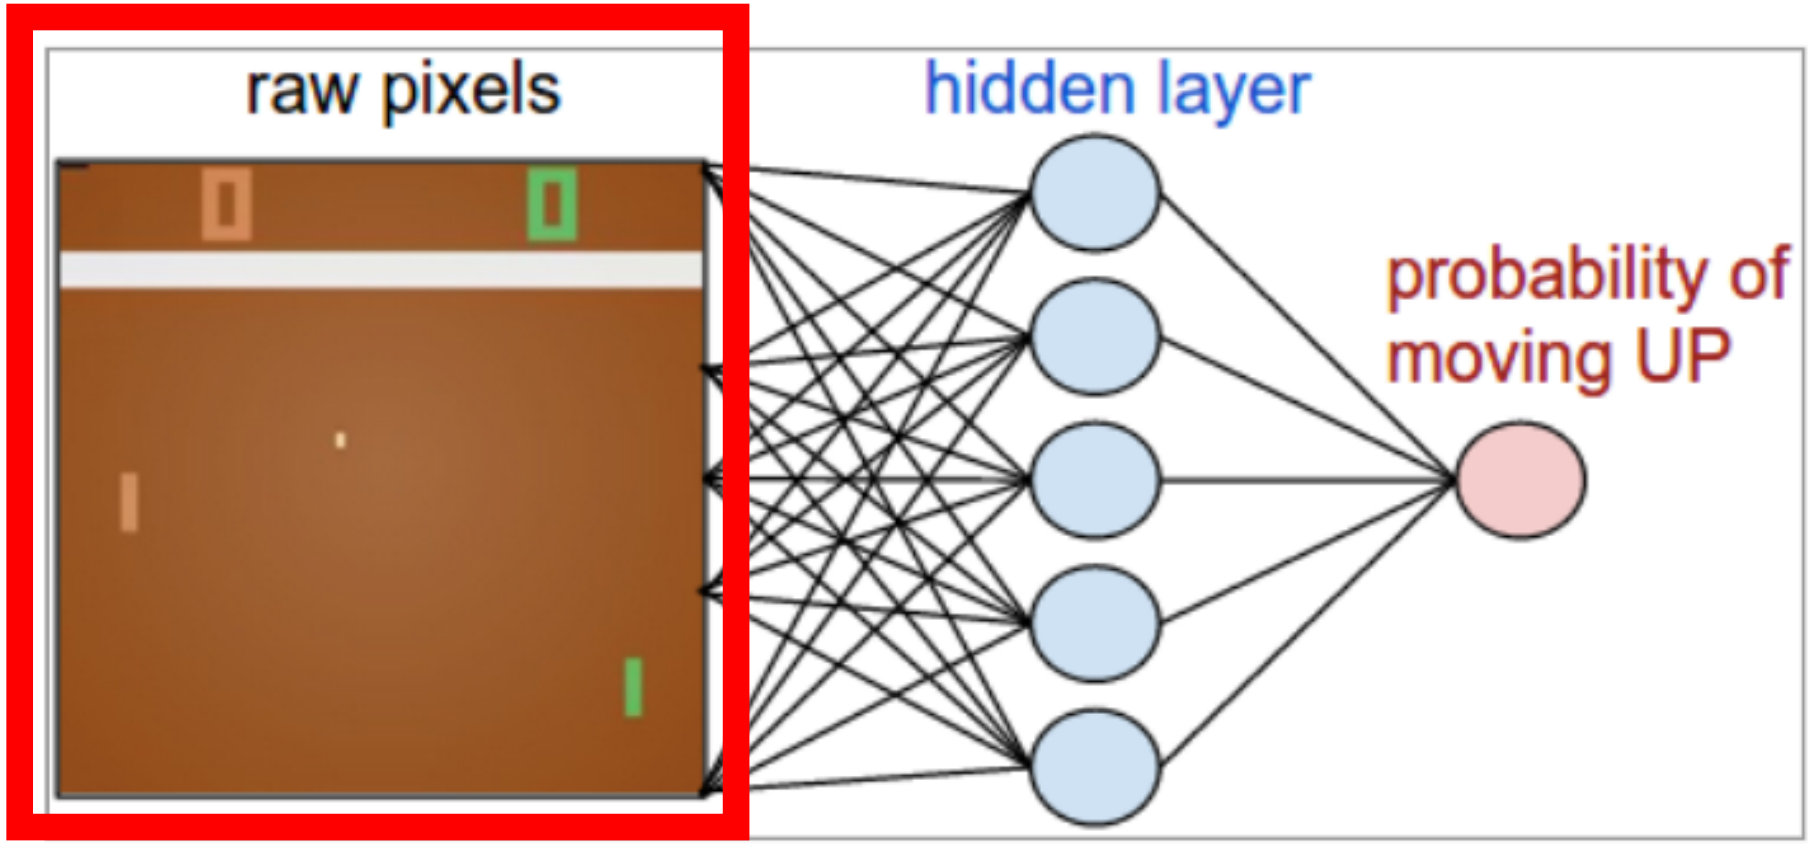

Game State

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

e.g., Pong Game: Training Protocol

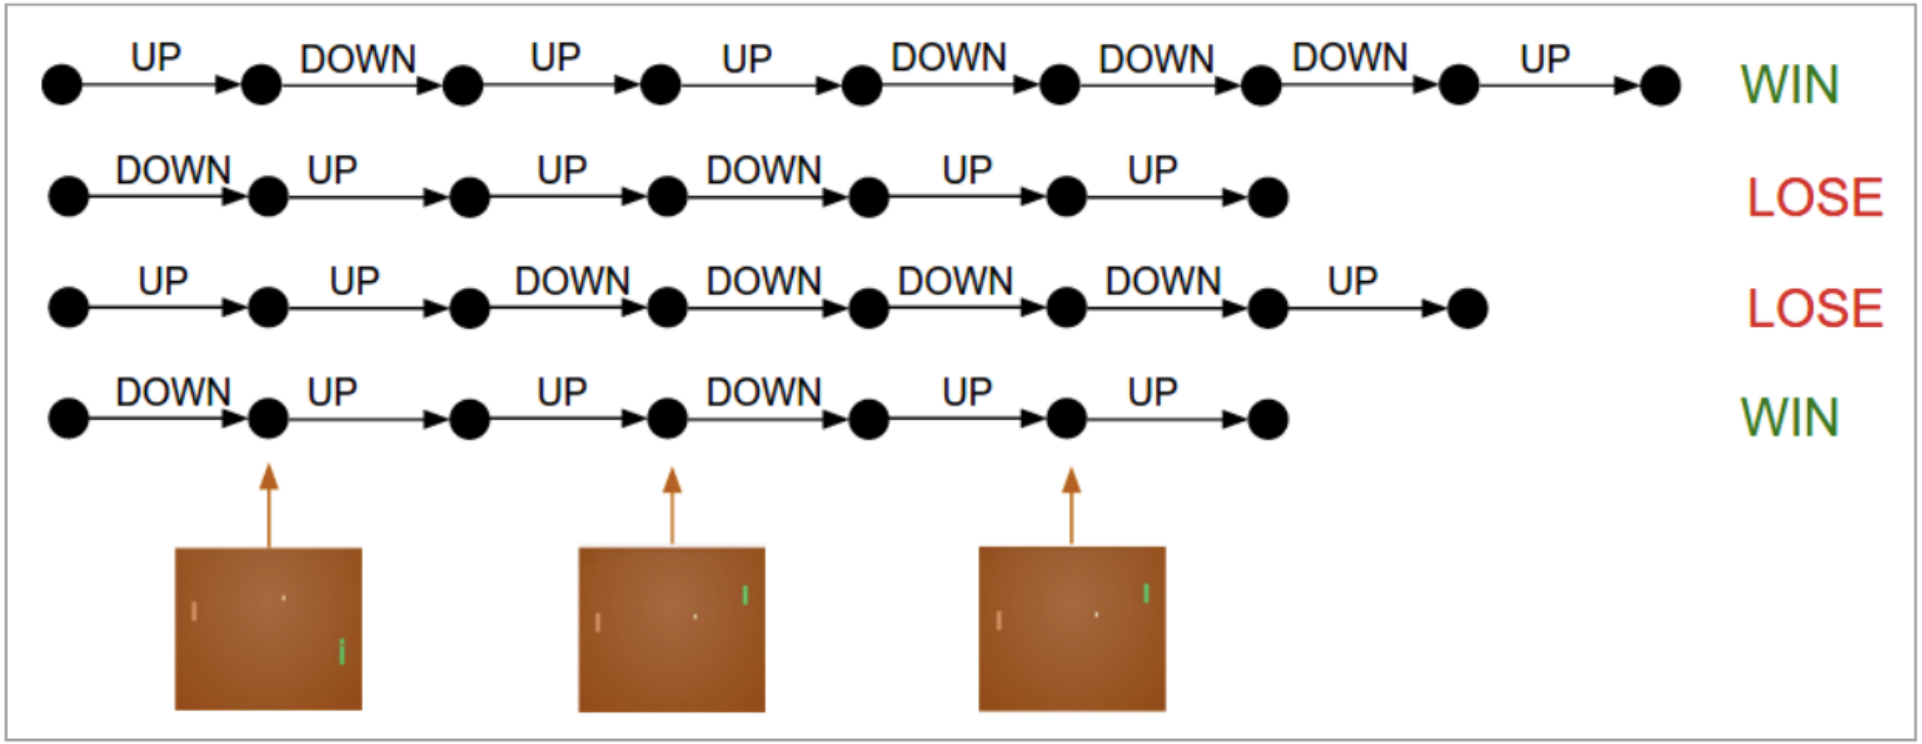

- Play 100 games of Pong; i.e., policy "rollouts" (200 images/game); Suppose: win 12 games, lose 88
- # Winning Decisions = 200*12 = 2400 decisions; positive update (fill in a +1.0 in the gradient for the sampled action, do backprop, and parameter update to encouraging the actions)
- # Losing Decisions: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

Google Form: Guest Speaker & Class Feedback

- Google form:
  - Guest: Dr. Cheryl Martin, Chief Data Scientist at Alegion (https://www.alegion.com/company/leadership); list one question for her for today's visit
- Then, take a short break.
- Class resumes at 4:50pm CST.

Today’s Topics

- Active Learning
- Curriculum Learning
- Reinforcement Learning
- Guest: Dr. Cheryl Martin from Alegion