CS 360: Machine Learning

Prof. Sara Mathieson

Fall 2019

HAVERFORD
COLLEGE

Admin

- **EVERYONE**: sign in (separate sheets for registered and waitlist) + pick up a handout
- **Waitlist**: come to either lab on Thursday
- **Registered**: must come to your assigned lab section unless you have found someone to switch with you (email me ASAP if that's the case)

Outline for Sept 3

- Welcome + what is Machine Learning (ML)?
- Examples of ML
- Syllabus highlights
- ML terminology and notation
- First algorithm: K-nearest neighbors

Outline for Sept 3

- Welcome + what is Machine Learning (ML)?
- Examples of ML
- Syllabus highlights
- ML terminology and notation
- First algorithm: K-nearest neighbors

Course Staff

- Instructor: Sara Mathieson (can call me Sara or Professor Mathieson)

- TA: Charles Marx (Charlie)
- TA: Pablo Thiel (Pablo)

- Student consultant: Mary Cott (Mary)

Student Introductions: next time (TODO fill out handout and bring it on Thurs!)

Discuss with a Partner

- Introduce yourselves
- Come up with your own definition of "Machine Learning"
- How is Machine Learning different from or similar to:
  - Statistics
  - Data mining
  - Psychology of learning

What is Machine Learning?

- “Machine Learning is the study of methods for programming computers to learn.”

  -Tom Dietterich

- “Machine Learning is about predicting the future based on the past.”

  -Hal Duame III

- “Machine Learning seeks to answer the question: `How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?’ ”

  -Tom Mitchell

What is Machine Learning?

- “Machine Learning is the study of methods for programming computers to learn.”

  -Tom Dietterich

- “Machine Learning is about predicting the future based on the past.”

  -Hal Duame III

- “Machine Learning seeks to answer the question: `How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?’ ”

  -Tom Mitchell

One more definition of ML

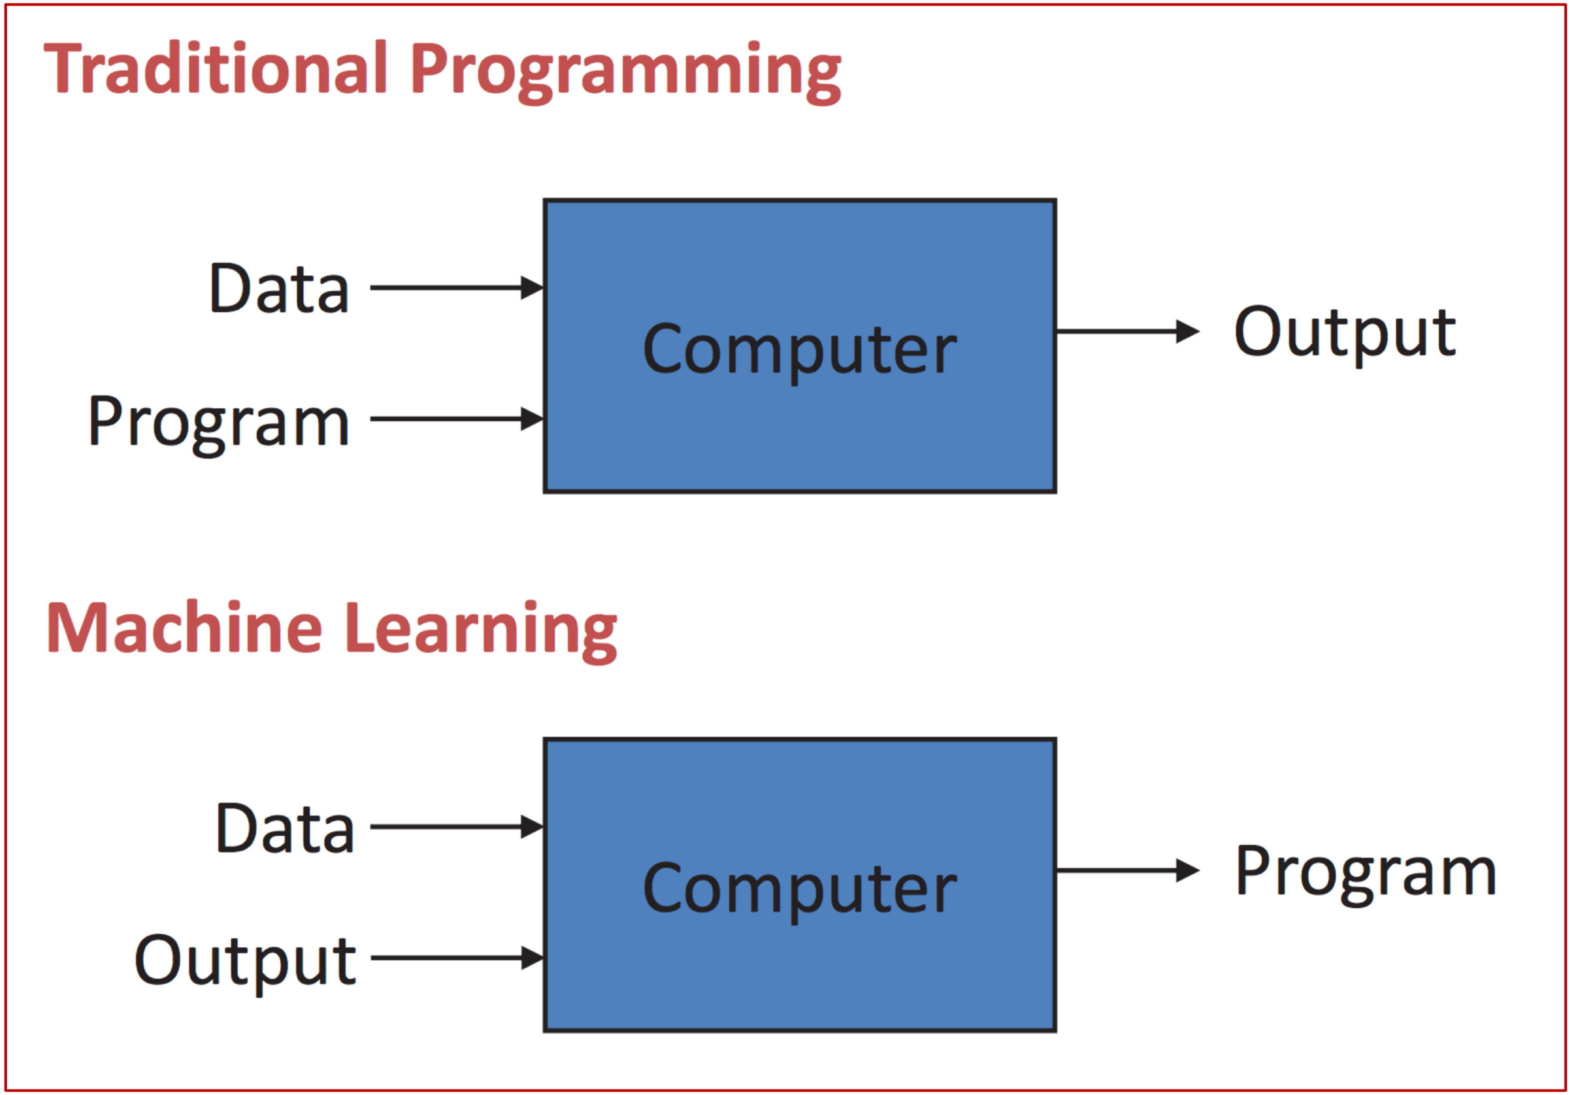

Slide credit: Jessica Wu
Source: Pedro Domingos

ML and related fields

- **Statistics:** understanding phenomenon that generated the data
- **Data Mining:** find patterns in data that are understandable to humans
- **Psychology of learning:** understand the mechanisms behind how humans learn

Adapted from “Machine Learning” by Tom Dietterich

Why would we want ML?

Why would we want ML?

1) No human experts

Example: predicting failure points for new machines

Why would we want ML?

1) No human experts

Example: predicting failure points for new machines

2) Human experts cannot explain expertise

Example: cannot explain exactly what a handwritten “2” looks like

Why would we want ML?

1) No human experts

Example: predicting failure points for new machines

2) Human experts cannot explain expertise

Example: cannot explain exactly what a handwritten “2” looks like

3) Phenomena change rapidly

Example: predicting the stock market

Adapted from “Machine Learning” by Tom Dietterich

Why would we want ML?

1) No human experts

   Example: predicting failure points for new machines

2) Human experts cannot explain expertise

   Example: cannot explain exactly what a handwritten “2” looks like

3) Phenomena change rapidly

   Example: predicting the stock market

4) Customization for each user

   Example: program that adapts to each user’s speech

Adapted from “Machine Learning” by Tom Dietterich

Outline for Sept 3

- Welcome + what is Machine Learning (ML)?
- Examples of ML
- Syllabus highlights
- ML terminology and notation
- First algorithm: K-nearest neighbors

Classic examples of ML

- Email filtering (spam vs. not-spam)

From: cheapsales@buystufffromme.com
To: ang@cs.stanford.edu
Subject: Buy now!

Deal of the week! Buy now!
Rolex w4tchs - $100
Med1cine (any kind) - $50
Also low cost M0rgages
available.

From: Alfred Ng
To: ang@cs.stanford.edu
Subject: Christmas dates?

Hey Andrew,
Was talking to Mom about plans
for Xmas. When do you get off
work. Meet Dec 22?
Alf

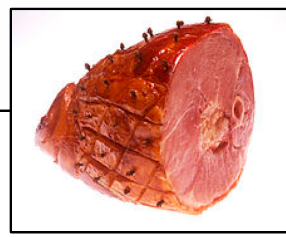

Images: Wikipedia, Matouš Havlena

Classic examples of ML

- Email filtering (spam vs. not-spam)

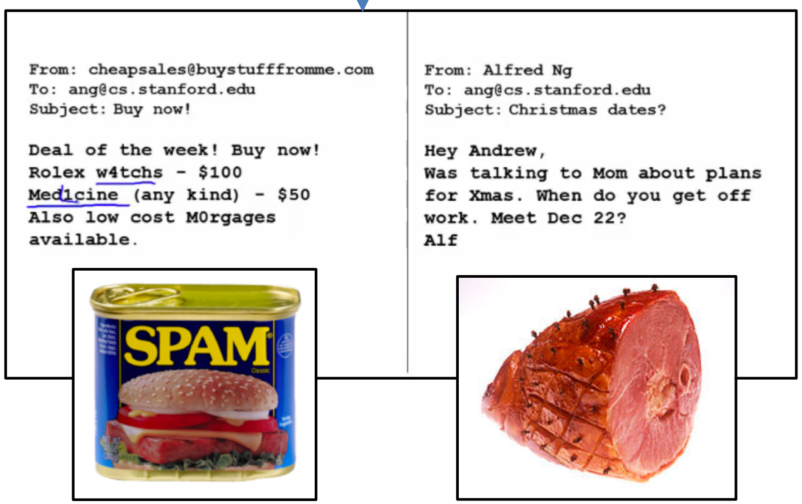

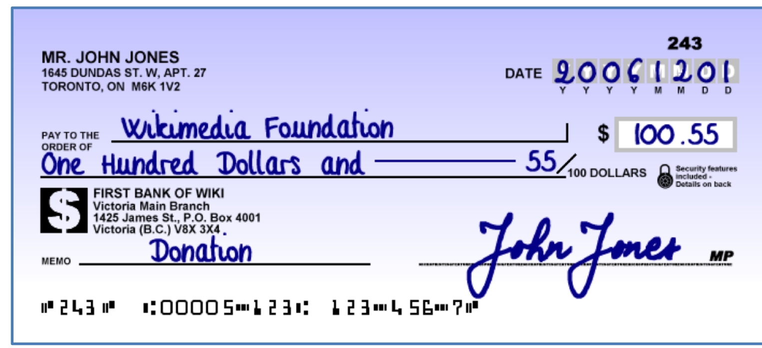

- Handwriting recognition (digits in a check)

Images: Wikipedia, Matouš Havlena

A classic example of a task that requires machine learning: It is very hard to say what makes a 2

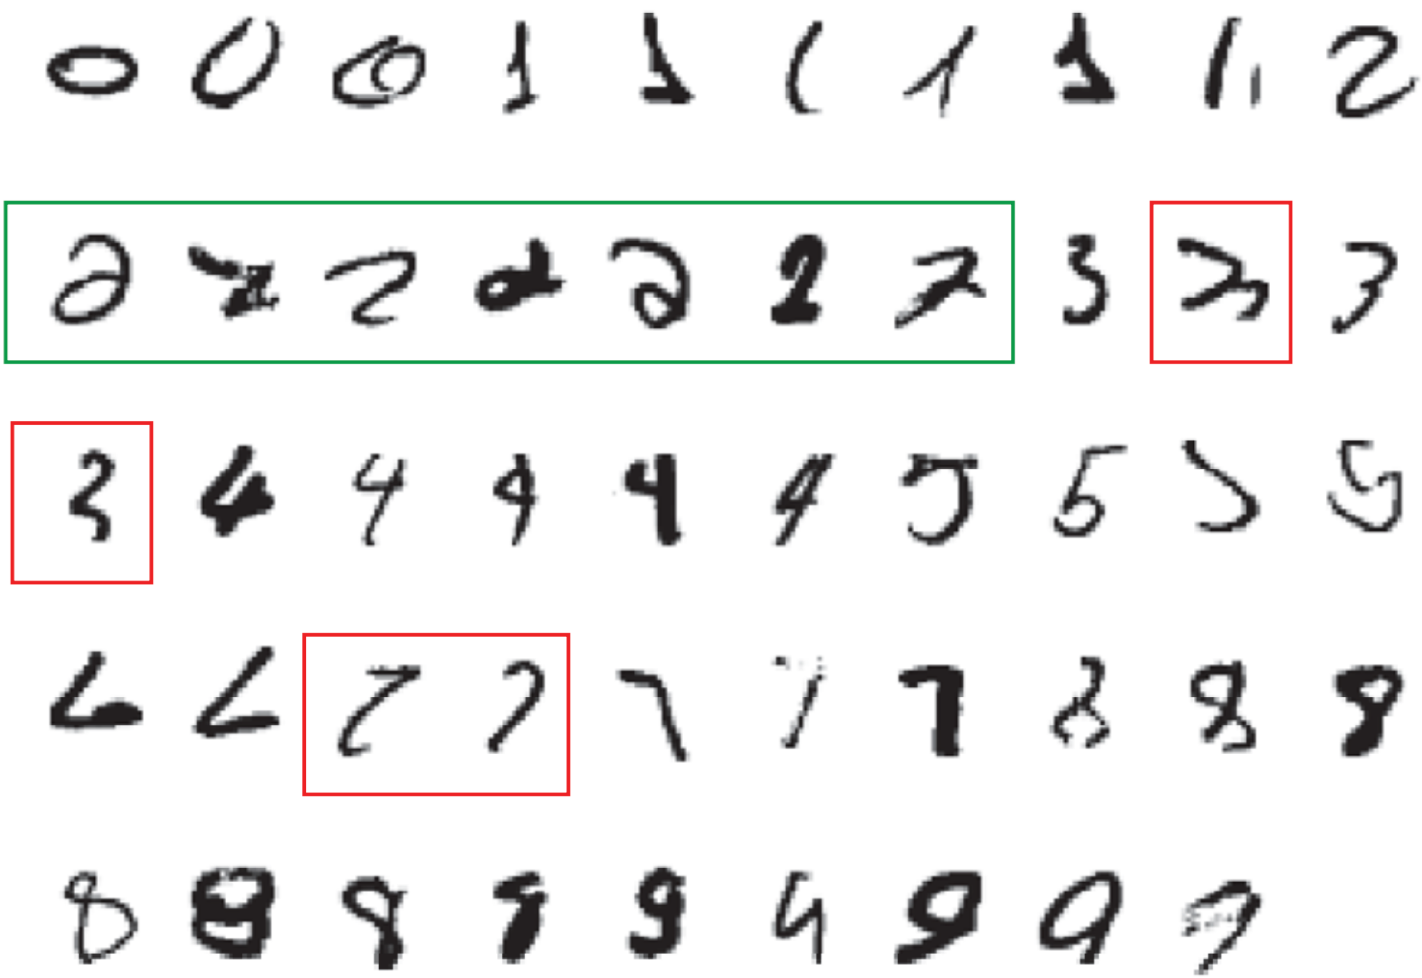

Slide credit: Jessica Wu
Source: Geoffrey Hinton

Classic examples of ML

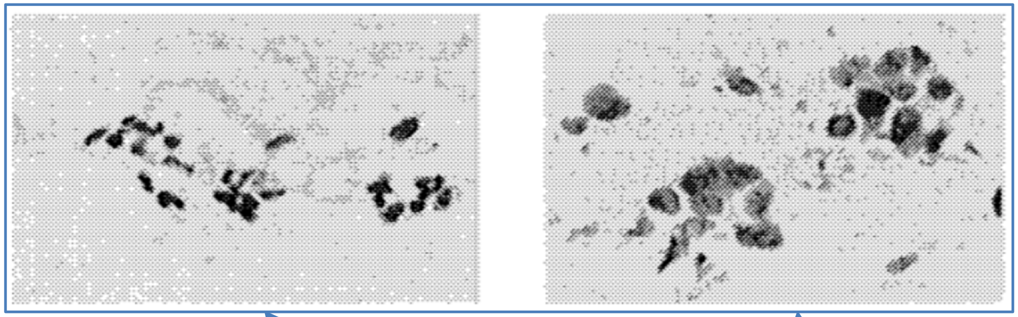

- Tumor detection (benign vs. malignant)

“On Breast Cancer Detection: An Application of Machine Learning Algorithms on the Wisconsin Diagnostic Dataset”

Modern examples of ML

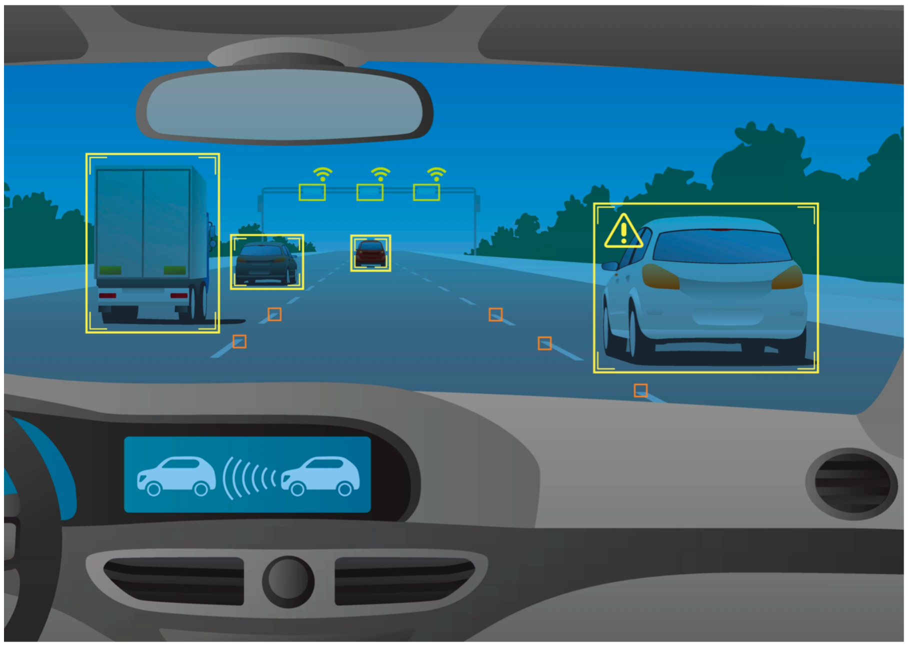

Self-driving cars are in our present and future

AlphaGo: moves humans never thought of

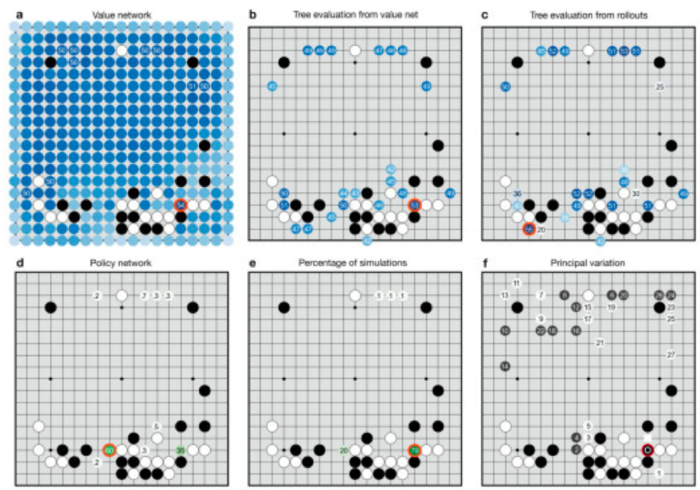

AlphaGo

Images: Scientific American

Modern examples of ML

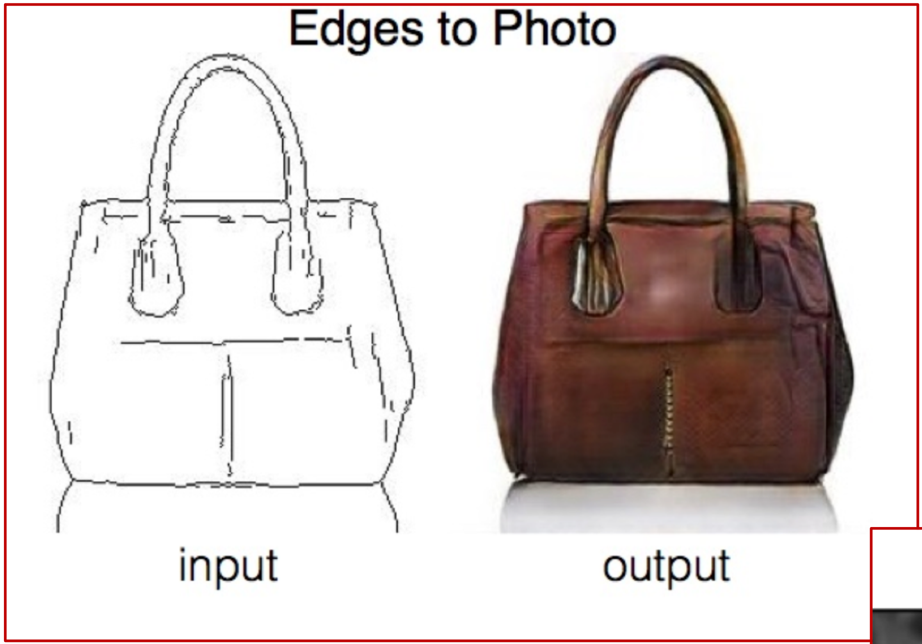

- Algorithms that learn how to create

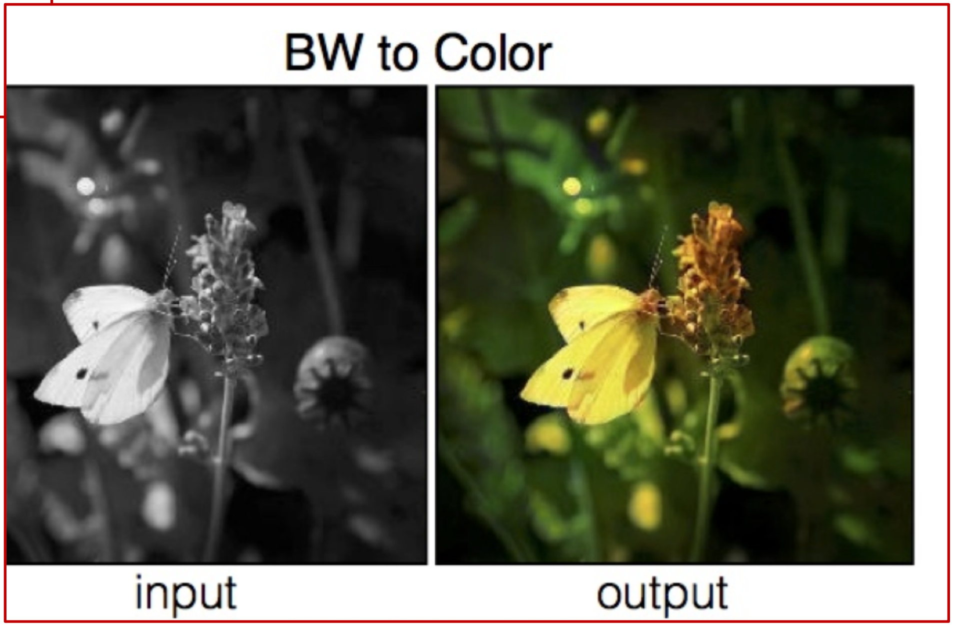

Image-to-Image Translation with Conditional Adversarial Nets (Nov 2016)

Modern examples of ML

"Building the Golden Gate Bridge" (est 1937)

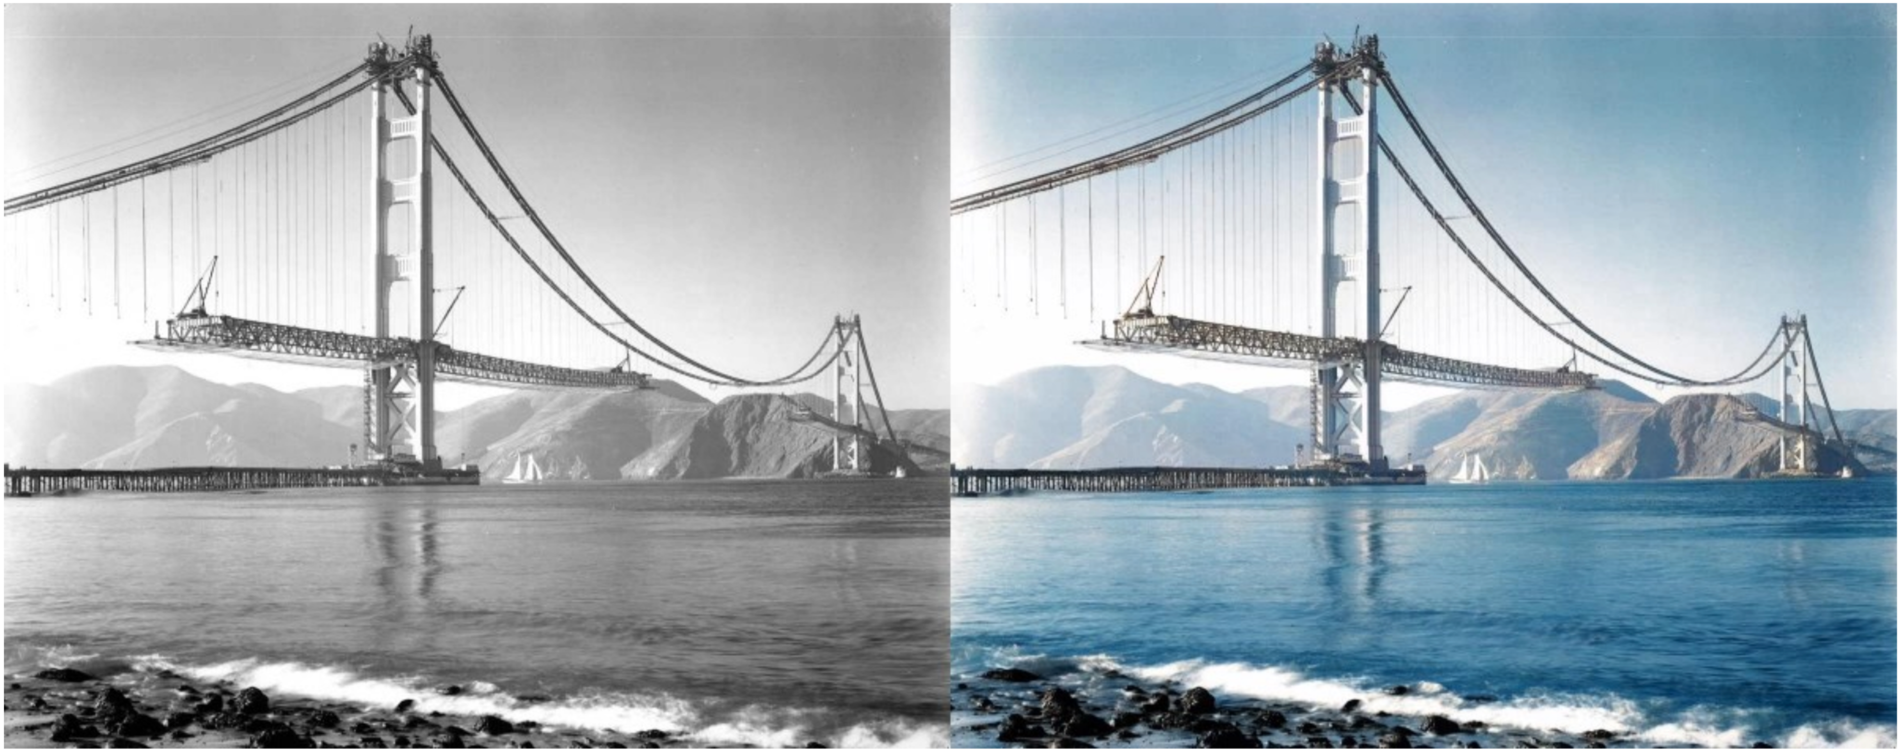

https://github.com/jantic/DeOldify

Face generation over time

Ian Goodfellow
@goodfellow_ian

Follow

4.5 years of GAN progress on face generation. arxiv.org/abs/1406.2661
arxiv.org/abs/1511.06434
arxiv.org/abs/1606.07536
arxiv.org/abs/1710.10196
arxiv.org/abs/1812.04948

4:40 PM - 14 Jan 2019

ML generated painting

Sold for almost half a million dollars

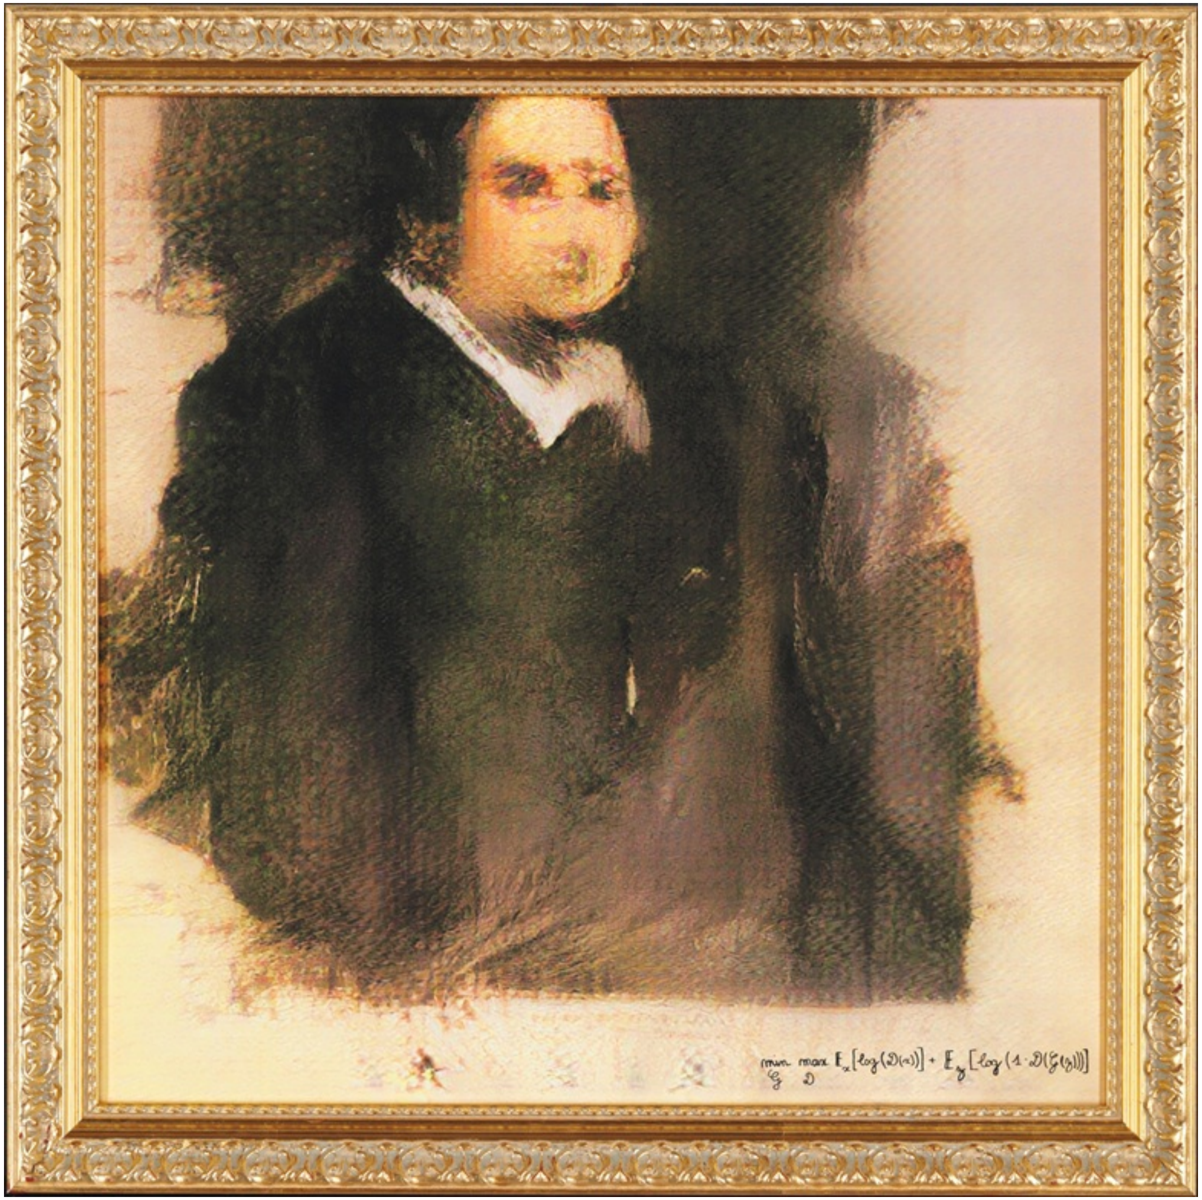

ML in Natural Language Processing

- Text Classification
- Language Modeling
- Speech Recognition
- Caption Generation
- Machine Translation
- Document Summarization
- Question Answering

Machine reading comprehension

Q: What was the theme?
A: "one world, one dream".

Q: What was the length of the race?
A: 137,000 km

Q: Was it larger than previous ones?
A: No

Q: Where did the race begin?
A: Olympia, Greece

Q: Is there anything notable about that place?
A: birthplace of Olympic Games

Q: Where did they go after?
A: Athens

Q: How many days was the race?
A: seven

Q: Did they visit any notable landmarks?
A: Panathinaiko Stadium

Q: And did they climb any mountains?
A:

Target answers: *unknown* or *yes*
Model answer: Everest

https://openai.com/blog/better-language-models/

https://machinelearningmastery.com/applications-of-deep-learning-for-natural-language-processing/

ML in Biology

- Gene location prediction
- Protein structure prediction
- Genes under natural selection
- Protein binding site prediction
- Population size changes

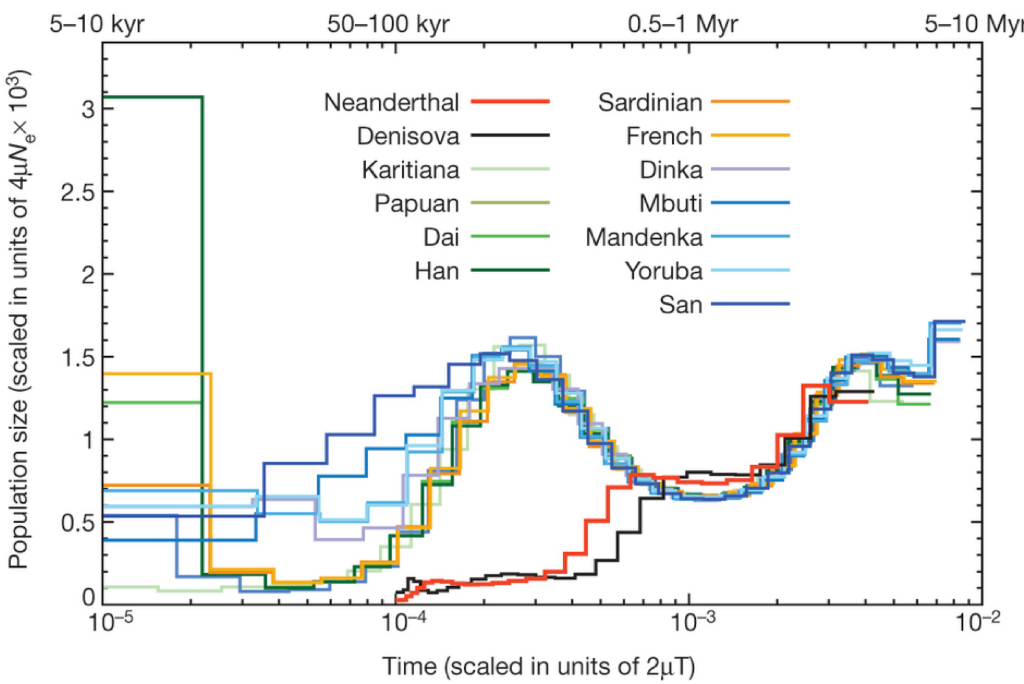

The complete genome sequence of a Neanderthal from the Altai Mountains, Prufer et al (2014)

March 2019
Turing award goes to neural network research

- LeCun, Hinton, Bengio win Turing award for their work on neural networks

https://www.nytimes.com/2019/03/27/technology/turing-award-ai.html

From left, Facebook, via Associated Press; Aaron Vincent Elkaim for The New York Times; Chad Buchanan/Getty Images

ML online

- Every day ML algorithms decide what people see and how their content is seen
  - Search results
  - Targeted ads
  - Newsfeed content
  - Facial recognition
- Example:

We must join the conversation

"When human beings acquired language, we learned not just how to listen but how to speak. When we gained literacy, we learned not just how to read but how to write. And as we move into an increasingly digital reality, we must learn not just how to use programs but how to make them."

-Douglas Rushkoff

Case Study: Google Duplex

- Google Assistant can now make phone calls on behalf of users

https://youtu.be/lXUQ-DdSDoE?t=35

Discussion Questions (small groups)

1) How is this technology an example of ML?

2) How is Google inviting us to imagine its product being used?

3) What other possible uses can we imagine?

4) What are benefits/opportunities of this technology? What are costs/concerns of this technology?

5) What choices are open to us if this is our product?

Questions adapted from Ronni Gura Sadovsky, Department of Philosophy, Harvard University

Outline for Sept 3

- Welcome + what is Machine Learning (ML)?
- Examples of ML
- Syllabus highlights
- ML terminology and notation
- First algorithm: K-nearest neighbors

Learning Goals

- First part of the semester: focus on understanding and implementing algorithms
- Later on: using powerful libraries (i.e. sklearn, TensorFlow, Keras, etc)
- Throughout and during the project: hypothesis development, featurization, algorithm selection, interpretation of results, iteration, conclusions
- Language: Python3, will use numpy/scipy throughout (recommended editor Atom)

Topics (tentative)

- ML terminology and notation
- K-nearest neighbors
- Decision Trees
- Linear regression
- Logistic regression
- Naïve Bayes
- Ensemble methods
- Support vector machines
- Neural networks
- CNNs
- GANs
- Unsupervised learning
- Dimensionality reduction
- Clustering
- Gaussian mixture models
- Special topics in deep learning
- ML and ethics

There will be math!

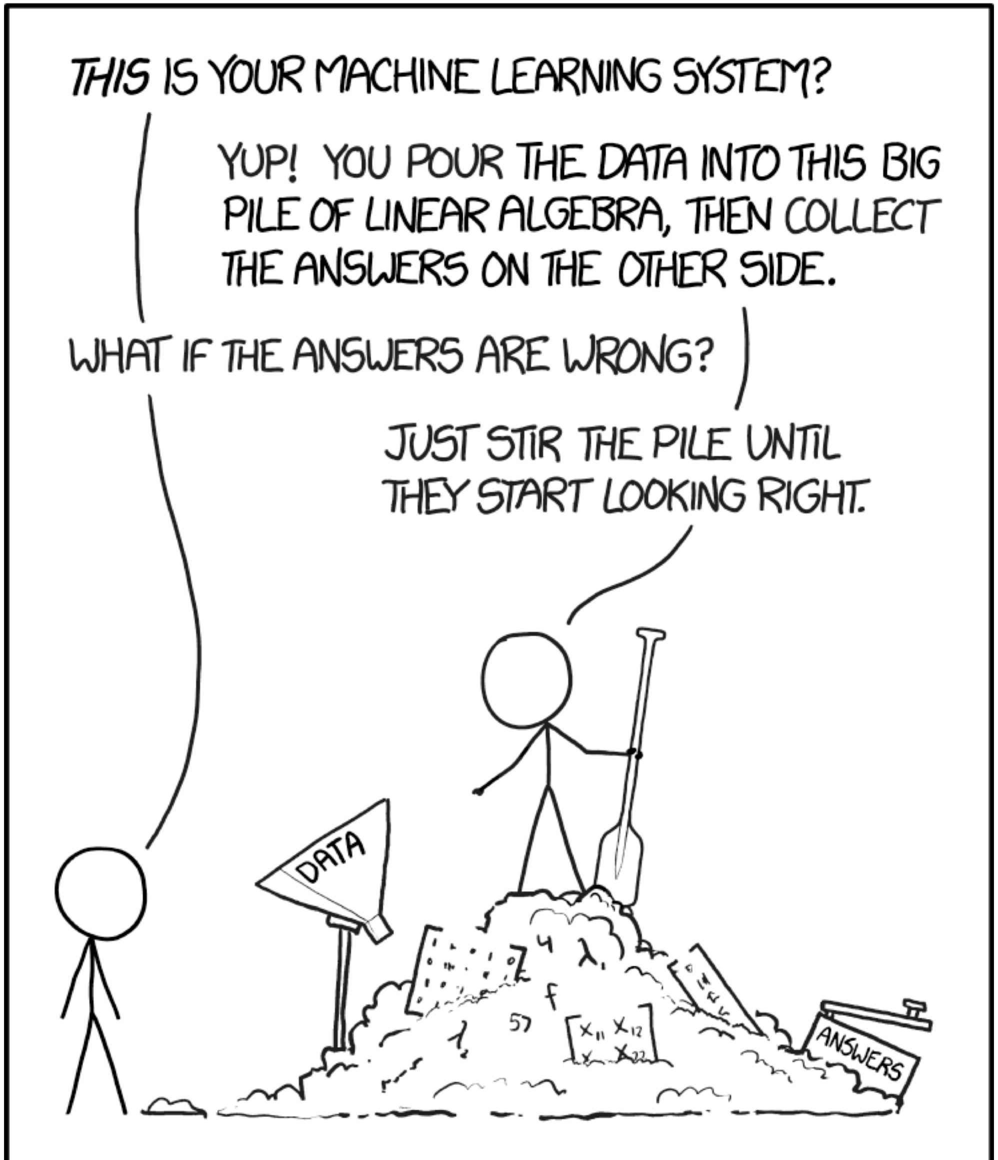

Different Backgrounds

- Prerequisites: Data Structures, Discrete Math, Linear Algebra
- May or may not have statistics or probability
- May or may not have taken artificial intelligence (AI), computer vision, bioinformatics, natural language processing (NLP)

Readings

- Modern machine learning relies heavily on linear algebra, geometry, statistics, and probability
- Our textbook provides a good introduction to terminology, how people think about ML, etc
- Supplemental readings go deeper

A Course in Machine Learning

by Hal Daumé III

Course Components

- Labs (roughly 9-10 total): 35%
- Midterms (2 in-lab + take home, week 5 and week 11): 40% (20% each)
- Final project: 15% (includes an oral presentation and “lab notebook”)
- Participation (includes reading quizzes): 10%

My Expectations

- Come to class (Tu/Th) and lab (Th), **ON TIME**, and actively participate
  - Email me if you will be absent from class
  - If you are sick, do not come to class!

- Complete the weekly reading *before Thurs* (for this week, just read **1.1-1.2**)

- Come to office hours (this week: Tues. 12:30-1:30pm, Fri. 3-5pm)
  - **KINSC L302**

- Post questions on Piazza

WEEK	DAY	ANNOUNCEMENTS	TOPIC & READING	LABS
1	Sep 03		**Introduction to Machine Learning** • Machine learning terminology • Notation • K-nearest neighbors Reading: • **Duame 1.1, 1.2** • **Duame 3.1-3.3** • **Machine Learning** by Tom Dietterich in *Nature Encyclopedia of Cognitive Science* (skim Sections 4-7) • (*optional*) ISL: Chapter 1 • (*optional*) ISL: Sections 2.1, 2.2 (focus on 2.1 and pg. 39-42) • (*optional*) **The Discipline of Machine Learning** by Tom Mitchell	**Lab 1**: K-nearest neighbors
	Sep 05			

Syllabus Notes

(Note: you are responsible for reading the entire syllabus on the course webpage)

1. Notes and slides will be posted *after* class on the course webpage
2. Lab is **mandatory** (attendance will be taken)
3. Labs often have a **pair programming** component (randomly assigned)
4. You will get **2 late days** during the semester
5. Extensions beyond these two days must be arranged with your class dean
6. Email: allow 24 hours for a response (more during weekends)
7. Piazza: should be used for all content/logistics questions

Participation

What counts as participation?

- Asking and answering questions in class (very important!)
  - Raise your hand (because some people are more/less comfortable shouting out answers)
  - Will call on groups, but only after giving you a few minutes to think/discuss
- Actively participating in in-class activities (group work, handouts, etc)
- Working well with your lab partner during lab
  - Switching who is at the keyboard
  - Discussing details instead of just trying to get to the end of the lab
- Asking and answering questions on Piazza
  - Avoid long blocks of code and giving away answers
  - Only non-anonymous posts count toward participation grade
- Attending office hours

Sometimes participation goes too far…

- Try to avoid dominating class discussion, office hours, Piazza, pair-programming, etc

Academic Integrity

Faculty statement on academic integrity

In a community that thrives on relationships between students and faculty that are based on trust and respect, it is crucial that students understand a professor's expectations and what it means to do academic work with integrity. Plagiarism and cheating, even if unintentional, undermine the values of the **Honor Code** and the ability of all students to benefit from the academic freedom and relationships of trust the Code facilitates. Plagiarism is using someone else's work or ideas and presenting them as your own without attribution. Plagiarism can also occur in more subtle forms, such as inadequate paraphrasing, failure to cite another person's idea even if not directly quoted, failure to attribute the synthesis of various sources in a review article to that author, or accidental incorporation of another's words into your own paper as a result of careless note-taking. Cheating is another form of academic dishonesty, and it includes not only copying, but also inappropriate collaboration, exceeding the time allowed, and discussion of the form, content, or degree of difficulty of an exam. Please be conscientious about your work, and check with me if anything is unclear.

Note for this course

Discussing ideas and approaches to problems with others on a general level is fine (in fact, we encourage you to discuss general strategies with each other), but you should never read anyone else's code or let anyone else read your code.

- No code from online
- No code from students who took this course previously

Academic Accommodations

Faculty statement on accommodations

Haverford College is committed to providing equal access to students with a disability. If you have (or think you have) a learning difference or disability – including mental health, medical, or physical impairment - please contact the Office of Access and Disability Services (ADS) at **hc-ads@haverford.edu**. The Coordinator will confidentially discuss the process to establish reasonable accommodations.

Students who have already been approved to receive academic accommodations and want to use their accommodations in this course should share their verification letter with me and also make arrangements to meet with me as soon as possible to discuss their specific accommodations. Please note that accommodations are **not retroactive** and require advance notice to implement.

It is a state law in Pennsylvania that individuals must be given advance notice if they are to be recorded. Therefore, any student who has a disability-related need to audio record this class must first be approved for this accommodation from the Coordinator of Access and Disability Services and then must speak with me. Other class members will need to be aware that this class may be recorded.

https://www.haverford.edu/access-and-disability-services/accommodations/receiving-accommodations

Class Deans

Raquel Esteves-Joyce
Assistant Dean of First Generation/Low Income (FGLI) Student Support & Programming

Stokes 118F
(610) 795-6121
rjoyce@haverford.edu

Make an Appointment
Profile

Amy Feifer
Dean of Career and Professional Advising & Associate Dean of the College

Stokes 300H
(610) 896-1181
afeifer@haverford.edu

Profile

Katrina Glanzer '02
Dean of First Year & Pre-Major Advising, Associate Dean of the College

Stokes 118C
(610) 896-1227
Assigned Students:
Class of 2023 (A-Z)

kglanzer@haverford.edu
Profile

Martha Denney
Dean of the College

Interim Title IX Coordinator

Stokes 103
(610) 896-1232
Assigned Students:
C-F ('20, '21, '22)

mdenney@haverford.edu
Profile
Assistant:
Susan Lewis

Nathan Diehl
Dean of Residential and Community Life & Associate Dean of the College

Stokes 022
(610) 896-1003
Assigned Students:
S ('20, '21, '22)

ndiehl@haverford.edu
Profile

Michael Elias
Dean of Student Engagement & Leadership and Divisional Initiatives, Associate Dean of the College

Stokes 022C
(610) 896-1228
Assigned Students:
T-Z ('20, '21, '22)

melias@haverford.edu
Profile

James Keane
Registrar, Director of Academic Assessment and Operations, Associate Dean of the College

Stokes 108
(610) 896-1023
Assigned Students:
M, Transfer Students ('20, '21, '22)

jkeane@haverford.edu
Profile

Michelle Leao
Director of Student Engagement and Leadership, and Assistant Dean of the College

Stokes 022
(610) 896-1298
Assigned Students:
N-R ('20, '21, '22)

mleao@haverford.edu
Profile

Theresa Tensuan
Assoc. Dean of the College; Dean for Diversity, Access and Community Engagement; Director of the Office of Multicultural Affairs

Stokes 118 A-1
(610) 896-1420
Assigned Students:
A-B ('20, '21, '22)

ttensuan@haverford.edu
Profile
Assistant:
Roxanne Madden

Brian Cuzzolina
Director of the OAR, Associate Dean for Student Academic Success and Persistence

Stokes 118G
(610) 795-6139
Assigned Students:
G-J ('20, '21, '22)

bcuzzolina@haverford.edu

Kelly Wilcox
Dean for Student Health & Learning Resources

Stokes 203D
(610) 795-6140
Assigned Students:
K-L ('20, '21, '22)

kwilcox@haverford.edu

Make an Appointment
Profile

Outline for Sept 3

- Welcome + what is Machine Learning (ML)?
- Examples of ML
- Syllabus highlights
- ML terminology and notation
- First algorithm: K-nearest neighbors

Machine learning terminology

- Alice takes ML for a semester, but the exam is on History of Pottery
- All exam questions are exactly the same as homework questions

Machine learning terminology

- Alice takes ML for a semester, but the exam is on History of Pottery
- All exam questions are exactly the same as homework questions

- Neither is a good judge of Alice's learning!

"Generalization is perhaps the most central concept in machine learning." –Duame

Machine learning terminology

- Alice takes ML for a semester, but the exam is on History of Pottery
- All exam questions are exactly the same as homework questions

- Neither is a good judge of Alice's learning!

"*Generalization* is perhaps the most central concept in machine learning." –Duame

Image: modified from Jessica Wu
Source based on: David Kauchak

Machine learning terminology

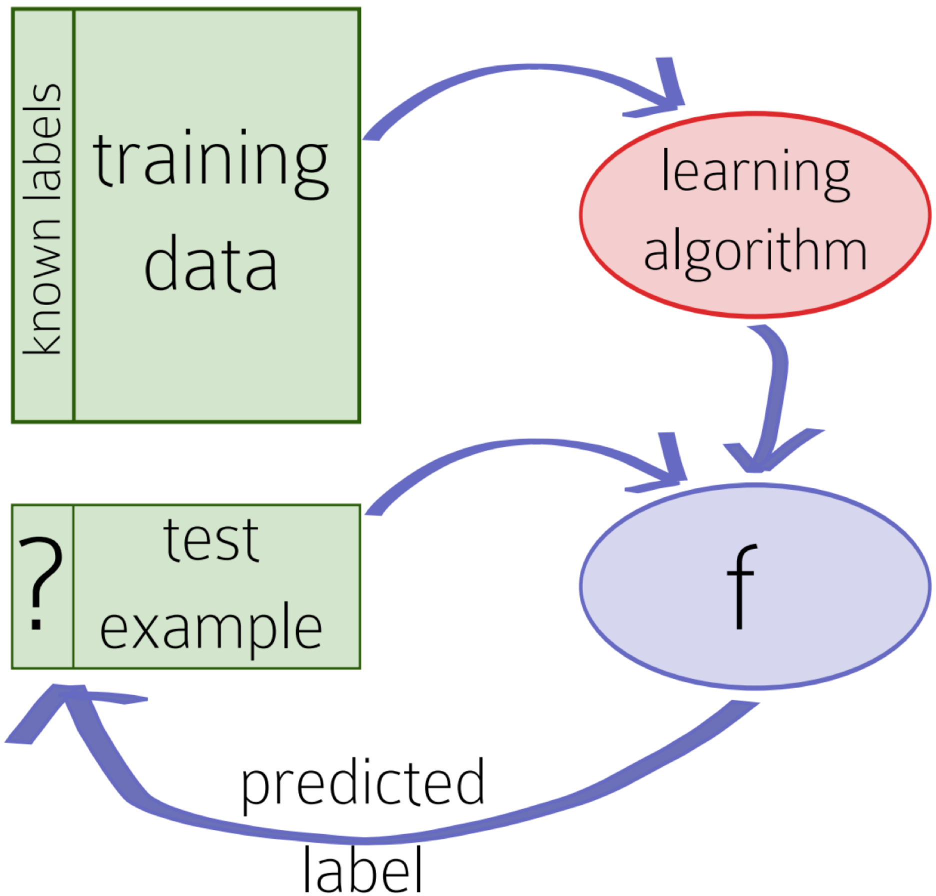

Figure 1.1, Duame

Machine learning terminology

- *Training*: usually involves the program learning from many *examples* (in a supervised setting we know the "answer" or *label* and are using this to learn)

- *Testing*: program predicts output/label for new examples without using their labels

Machine learning terminology

- *Training*: usually involves the program learning from many *examples* (in a supervised setting we know the "answer" or *label* and are using this to learn)

- *Testing*: program predicts output/label for new examples without using their labels

Must never look at the test data!

Machine learning terminology

- *Training*: usually involves the program learning from many *examples* (in a supervised setting we know the "answer" or *label* and are using this to learn)

- *Testing*: program predicts output/label for new examples without using their labels

Must never look at the test data!

Caveat: not all ML problems decompose into training and testing!

Machine learning terminology

- *Supervised learning:* we have information about the output or response variable
  - (can be easier for the computer to learn the function between input and output)

- *Unsupervised learning:* data is unlabeled (no output/class information)
  - Note: there may not be an output to learn

Machine learning terminology

- A common ML task is *regression*
- In this case the output or *response variable* is *continuous*

Example: modeling **house price** as a function of **size, location, year built**, etc

Machine learning terminology

- Another common style of machine learning is *classification*
- Goal: separate examples into two or more *classes* or *categories* (*discrete* setting)
- Example: is a credit card transaction legitimate or fraudulent?

Example: the classic bagel v. dog challenge

How can we *distinguish* between similar objects?

Images: veriy.com

Outline for Sept 3

- Welcome + what is Machine Learning (ML)?
- Examples of ML
- Syllabus highlights
- ML terminology and notation
- First algorithm: K-nearest neighbors

X = input, data, features

$$\underbrace{\begin{bmatrix} \text{———} \; \vec{x}_i \; \text{———} \end{bmatrix}}_{p \,=\, \#\text{ features}} \quad n \times p$$

Y = label / output / class

- regression: $y \in \mathbb{R}$
- binary classification: $y \in \{+1, -1\}$
  or $y \in \{0, 1\}$
- multiclass classification
  $y \in \{1, 2, \dots, C\}$ $C = \#$ classes
- ranking (websearch)

1, 5]

ath.

$\hat{y}$ = prediction for y

accuracy: $\frac{1}{n} \sum_{i=1}^{n} \mathbb{1}(y_i = \hat{y}_i)$

indicator

$$\mathbb{1}(\text{statement}) = \begin{cases} 1 & \text{if statement is true} \\ 0 & \text{if false} \end{cases}$$

$$y = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad \hat{y} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}$$

$n = 5$

accuracy

$$= \frac{1}{5}(1+0+1+1+0)$$

$$= \boxed{\frac{3}{5}}$$

K-nearest neighbors

idea: predict → new example belogs to the same class as the most similar example we've seen before.

bors

example

same

most

mple we've

re.

Euclidean distance

$$d(\vec{x}, \vec{z}) = \sqrt{(x_1-z_1)^2+(x_2-z_2)^2+\cdots(x_p-z_p)^2}$$

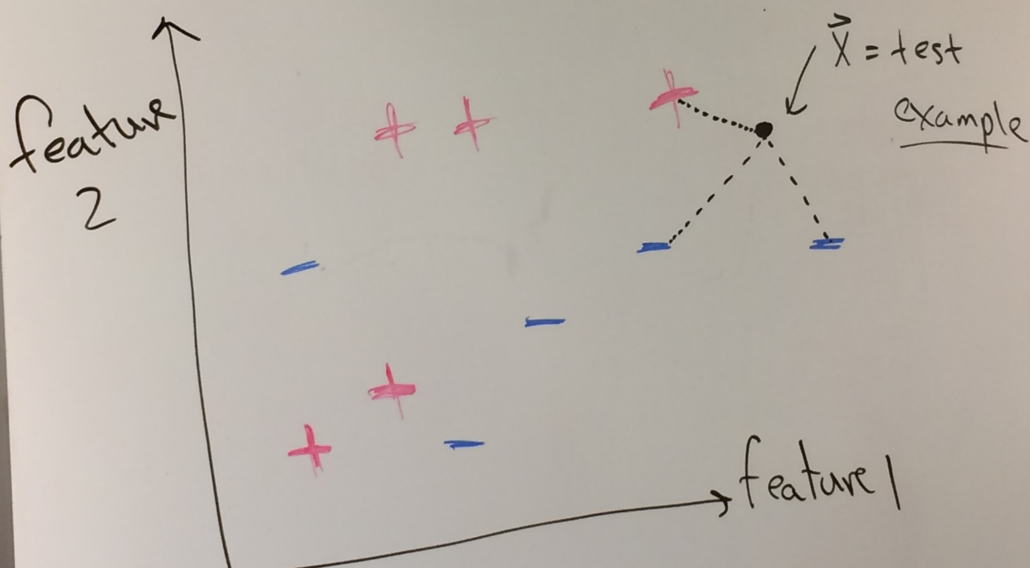

TODO for Thursday

- Reading: Duame 1.1-1.2 (3 pages!)
  - Short reading quiz on these pages on Thurs

- Continue to 3.1-3.3 if you have time

- Fill out Handout 1 and bring back for introductions!