CS 360: Machine Learning

Sara Mathieson, Sorelle Friedler

Spring 2024

HAVERFORD
COLLEGE

Please sit somewhere you did not sit in CS260!

Admin

- **EVERYONE**: (registered and waitlist)
  - Sign in
  - Pick up a handout
  - Pick up a notecard
  - Pick up a construction paper sheet

- We still have an imbalance in the labs – if you can come to **Lab B** please let me know!

Outline for Jan 23

- Welcome and syllabus highlights
- Classification review
- Evaluation metrics review
- Naïve Bayes review
- Thursday: gradient descent review

Outline for Jan 23

- Welcome and syllabus highlights

- Classification review

- Evaluation metrics review

- Naïve Bayes review

- Thursday: gradient descent review

Course Staff

- **Instructor**: Sara Mathieson
- **Lab instructor**: Sorelle Friedler
- **TAs**:
  - Trinity Kleckner
  - Grace Proebsting
  - Ben Menko
- **Peer tutors**:
  - Wahub Ahmed
  - Seun Eisape

Notecard and Name card

- **Notecard**:
  - Preferred first name
  - Pronouns (optional)
  - Which semester you took CS260
  - One topic you're hoping we'll cover in CS360
  - Anything else that would be helpful for us to know

- **Name card** ("tent")
  - Preferred first name
  - Pronouns (optional)
  - (sharpies going around!)

Discuss with a Partner

- Introduce yourselves
  - Be ready to introduce your partner to the class!
- Come up with your own definition of "Machine Learning"
- Share the topic you're hoping to cover in CS360

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

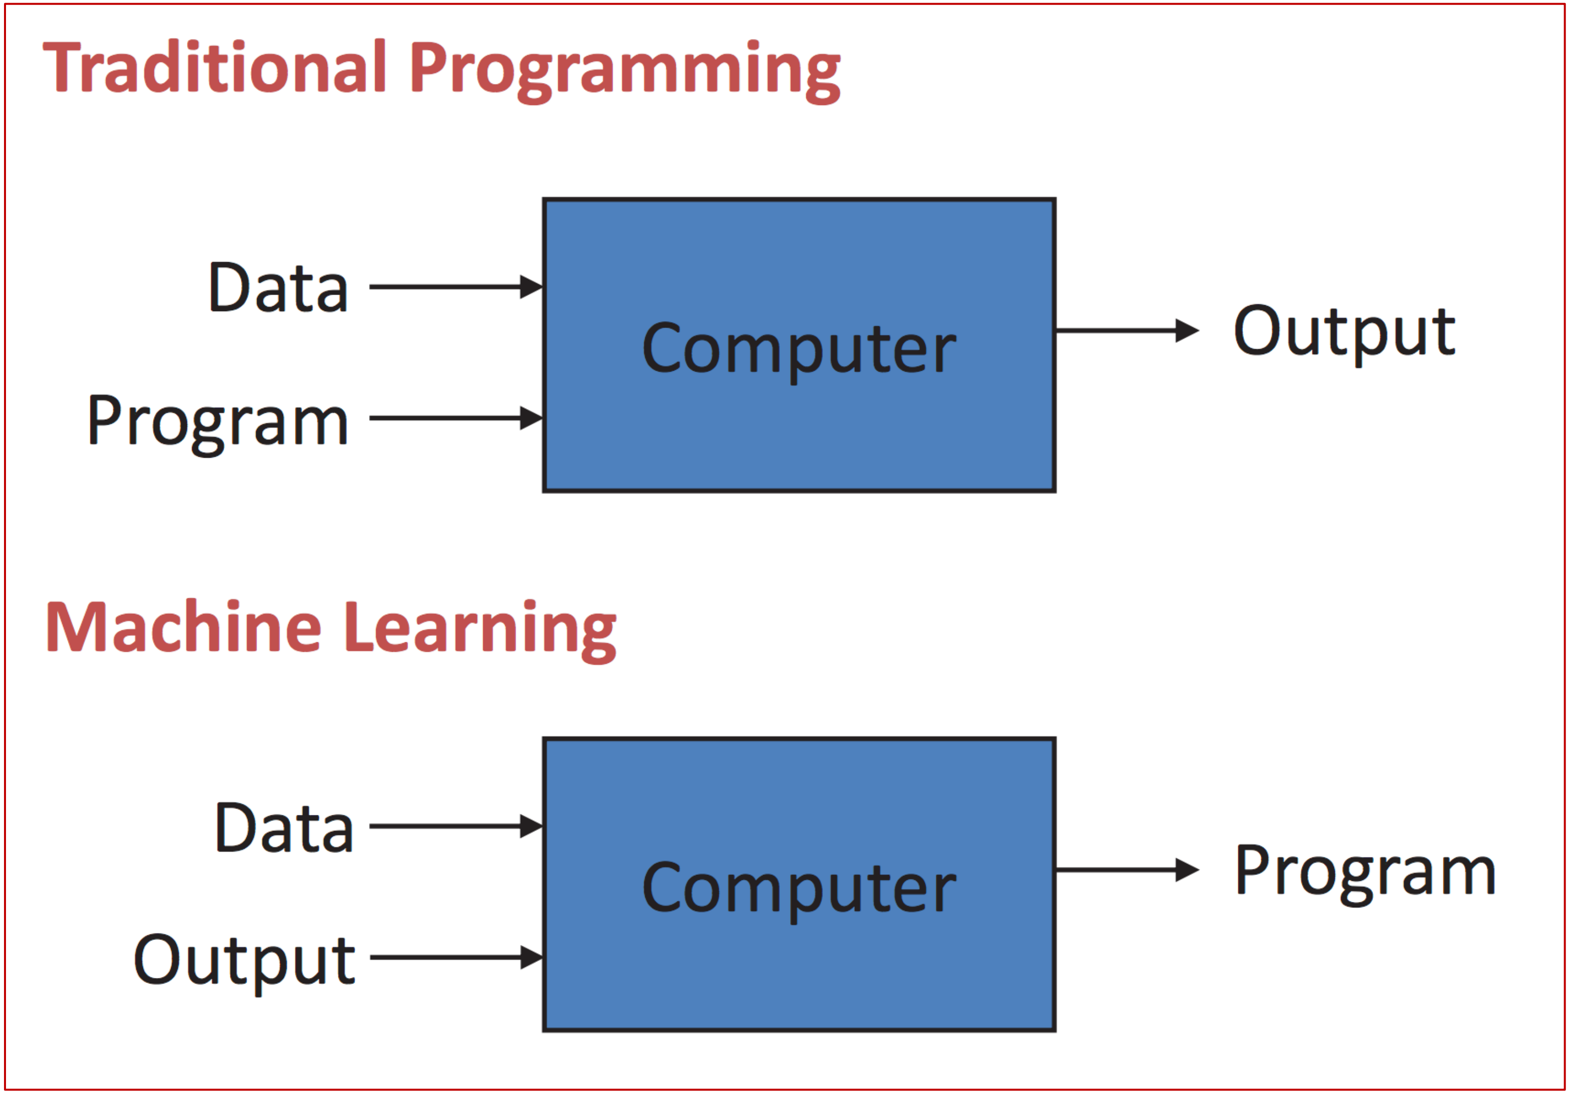

Slide credit: Jessica Wu
Source: Pedro Domingos

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?

Adapted from “Machine Learning” by Tom Dietterich

Why would we want ML?

1) No human experts

Example: predicting failure points for new machines

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

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

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

Learning Goals

- Understanding and implementing machine learning algorithms

- Using libraries (i.e. sklearn, TensorFlow, Keras, etc) in a principled machine learning workflow

- Throughout and during the project: hypothesis development, featurization, algorithm selection, interpretation of results, iteration, conclusions

- Language: Python3, will use numpy/scipy throughout (recommended editor VScode)

Topics (tentative)

- CS260 review
- Machine learning pipeline
- Evaluation metrics (AUC, cross validation)
- Model cards
- K-nearest neighbors and KD trees
- Decision trees
- Ensemble methods
- Advanced regression
- Fairness and explainability
- Support vector machines
- Neural networks
- CNNs
- Transformers
- GANs
- Unsupervised learning
- Dimensionality reduction
- Clustering
- Gaussian mixture models

Textbook

- Required textbook:

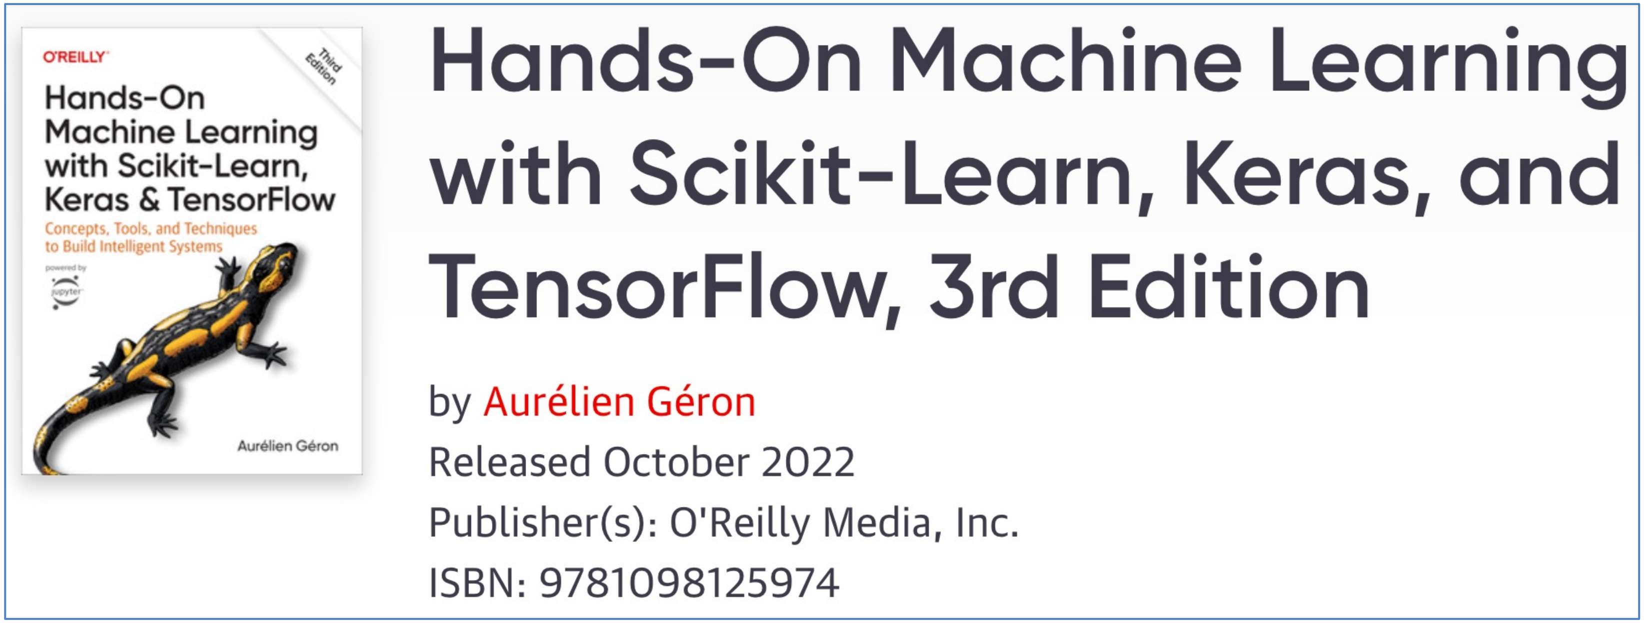
O'REILLY
Third Edition
Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow
Concepts, Tools, and Techniques to Build Intelligent Systems
powered by jupyter
Aurélien Géron

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 3rd Edition

by Aurélien Géron

Released October 2022

Publisher(s): O'Reilly Media, Inc.

ISBN: 9781098125974

- Supplemental readings provided as needed

Course Components

- Labs (roughly 8 total): 35%
- Midterms (2 in class): 40% (20% each)
- Final project: 15% (includes an oral presentation and writeup)
- Participation (including attendance): 10%

My Expectations

- Come to class (Tu/Th) and lab (Tu), **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*
- Come to office hours
- Post questions on Piazza

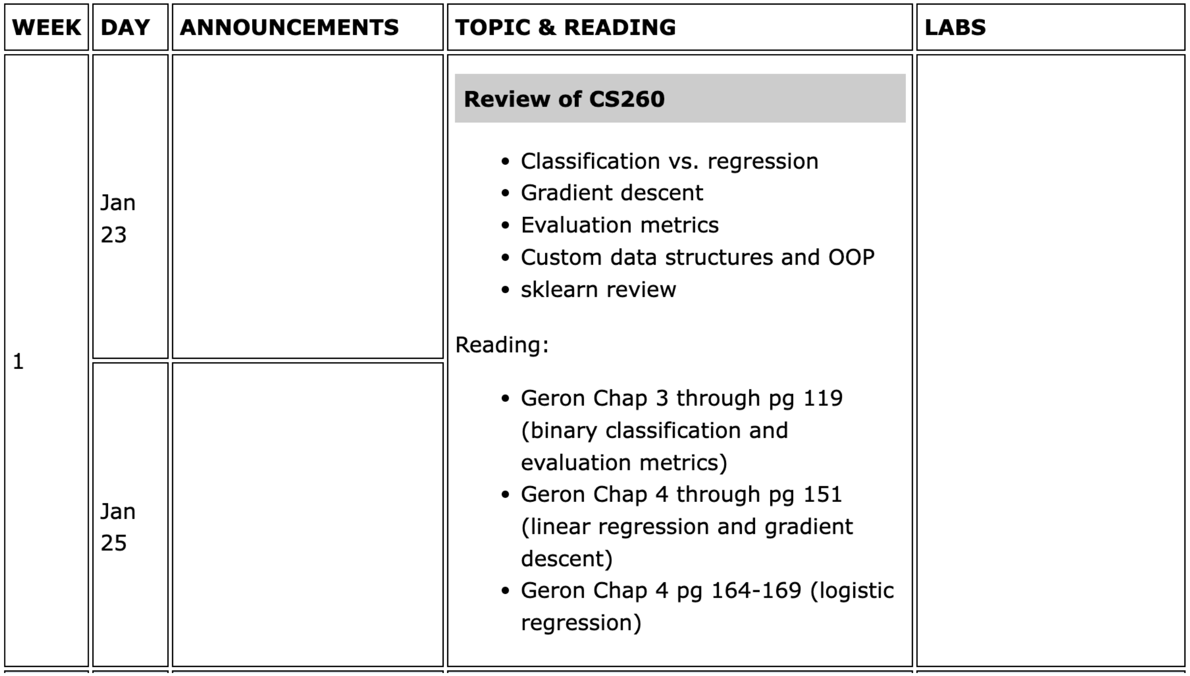

WEEK	DAY	ANNOUNCEMENTS	TOPIC & READING	LABS
1	Jan 23		**Review of CS260** • Classification vs. regression • Gradient descent • Evaluation metrics • Custom data structures and OOP • sklearn review Reading: • Geron Chap 3 through pg 119 (binary classification and evaluation metrics) • Geron Chap 4 through pg 151 (linear regression and gradient descent) • Geron Chap 4 pg 164-169 (logistic regression)	
	Jan 25			

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

Academic Integrity

Faculty statement on academic integrity

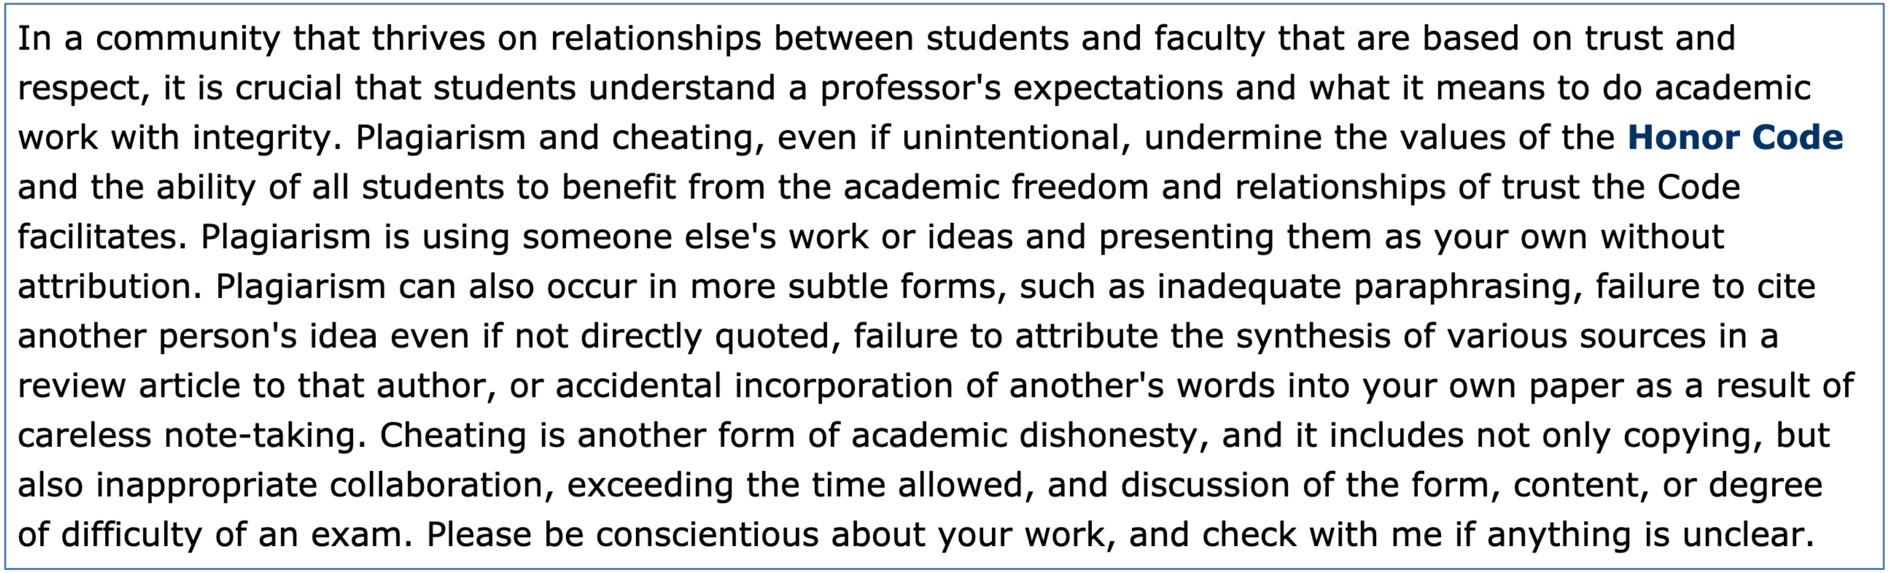

In a community that thrives on relationships between students and faculty that are based on trust and respect, it is crucial that students understand a professor's expectations and what it means to do academic work with integrity. Plagiarism and cheating, even if unintentional, undermine the values of 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

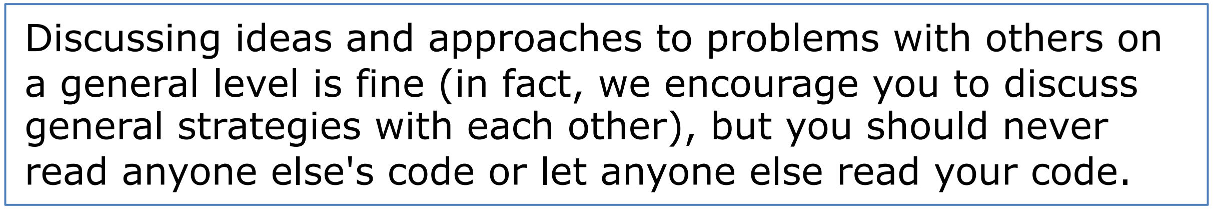

Discussing ideas and approaches to problems with others on a general level is fine (in fact, we encourage you to discuss general strategies with each other), but you should never read anyone else's code or let anyone else read your code.

- No code from online
- No code generation programs (until the final project)
- No code from students who took this course previously

Academic Accommodations

Faculty statement on accommodations

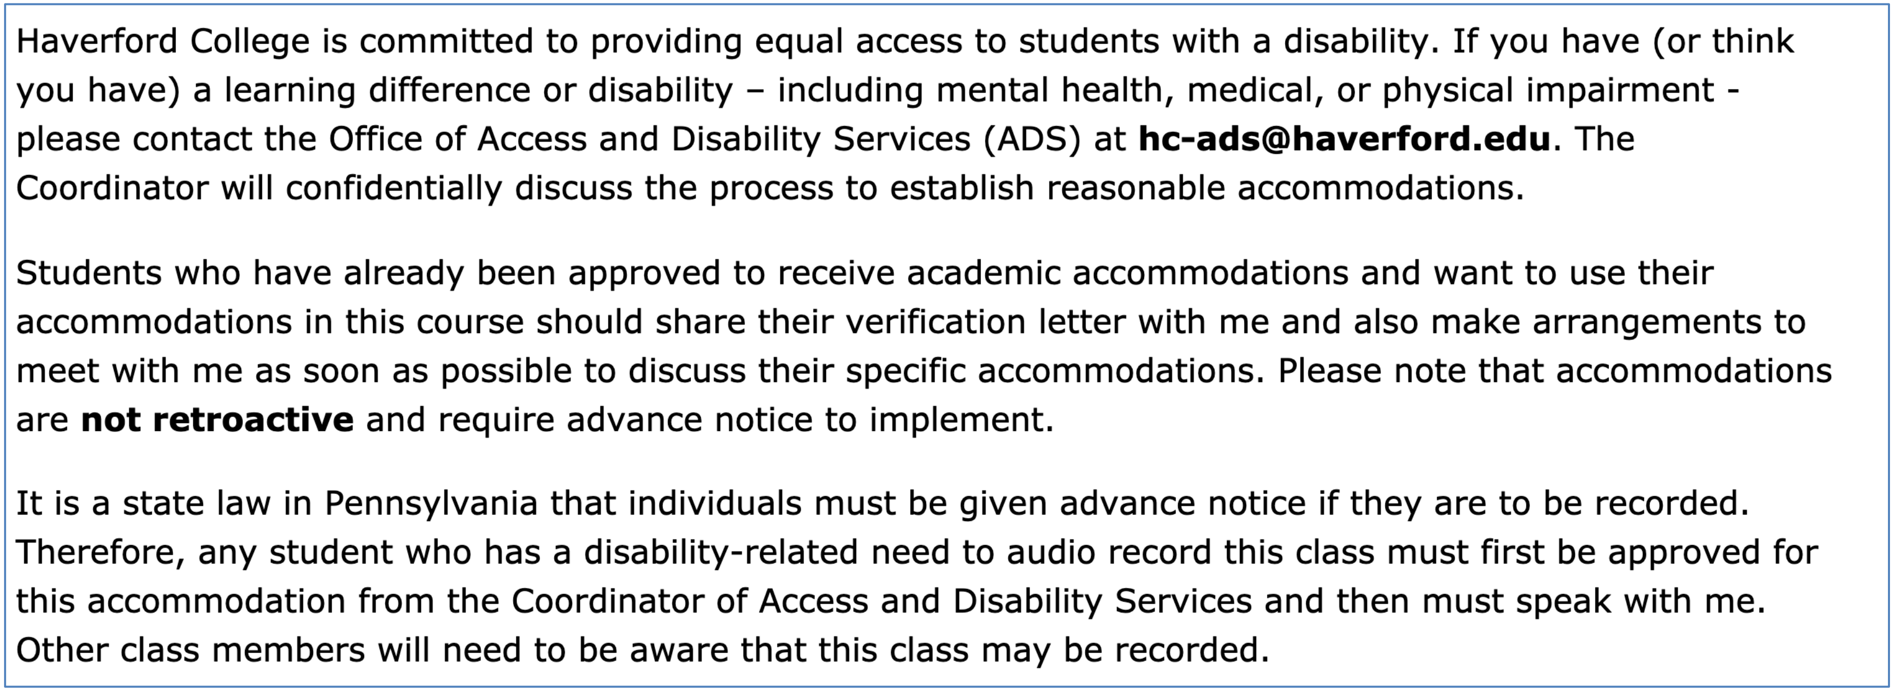

Haverford College is committed to providing equal access to students with a disability. If you have (or think you have) a learning difference or disability – including mental health, medical, or physical impairment - please contact the Office of Access and Disability Services (ADS) at **hc-ads@haverford.edu**.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

Outline for Jan 23

- Welcome and syllabus highlights
- Classification review
- Evaluation metrics review
- Naïve Bayes review
- Thursday: gradient descent review

Traditional Approach

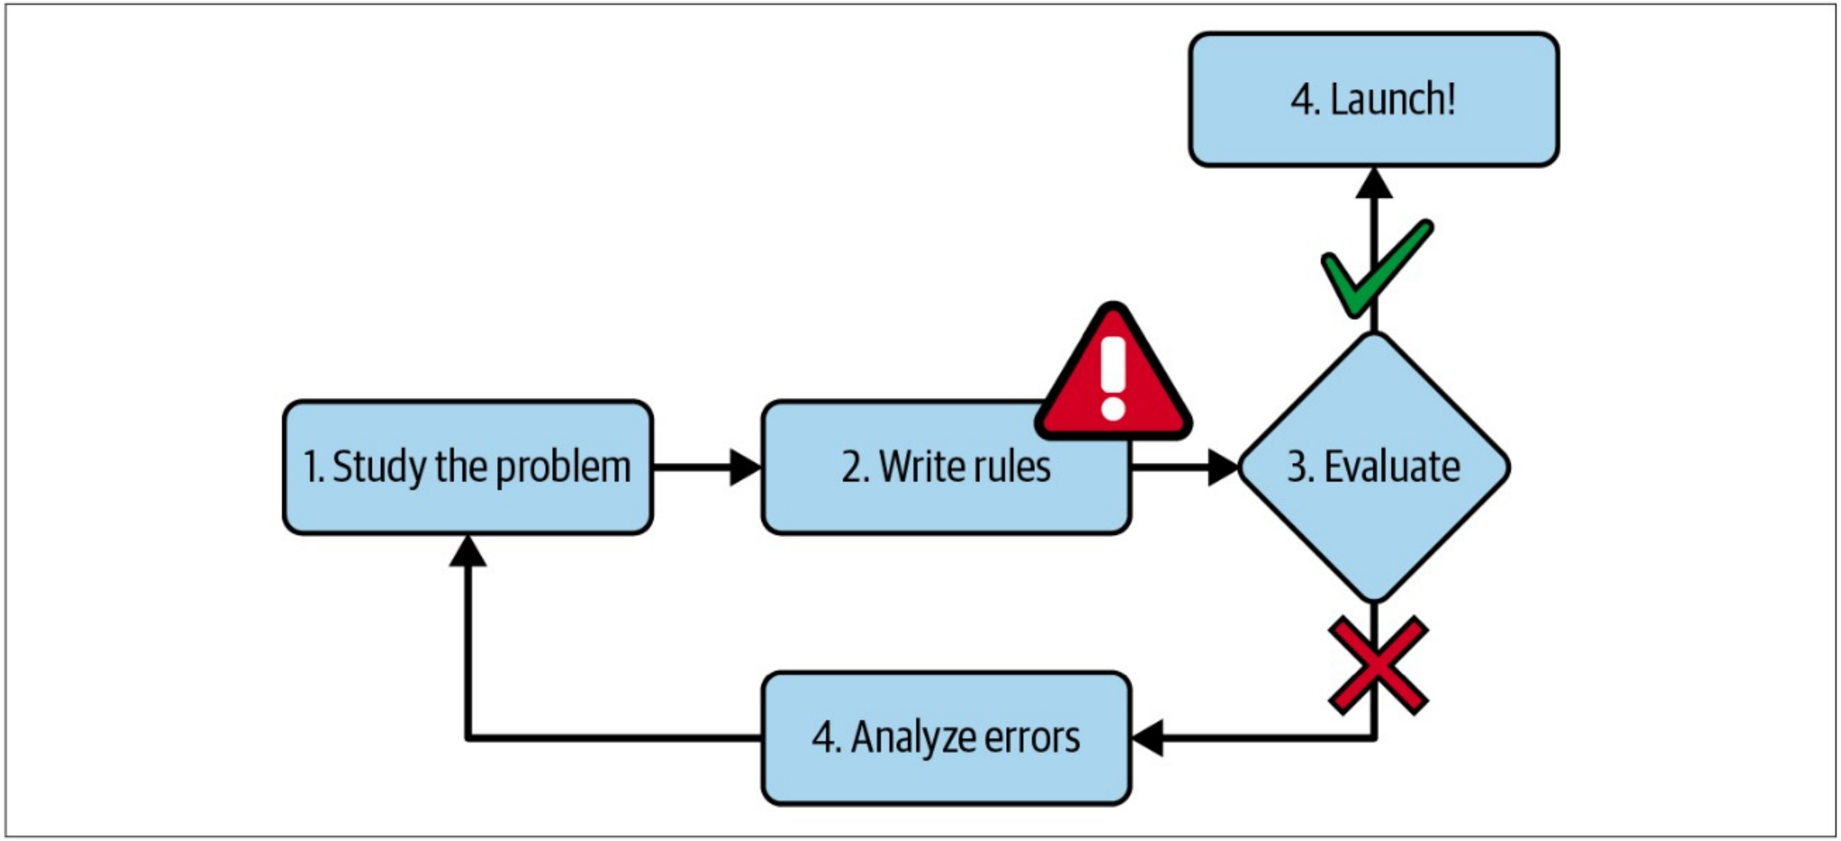

Figure 1-1. The traditional approach

Geron (textbook)

Machine Learning approach

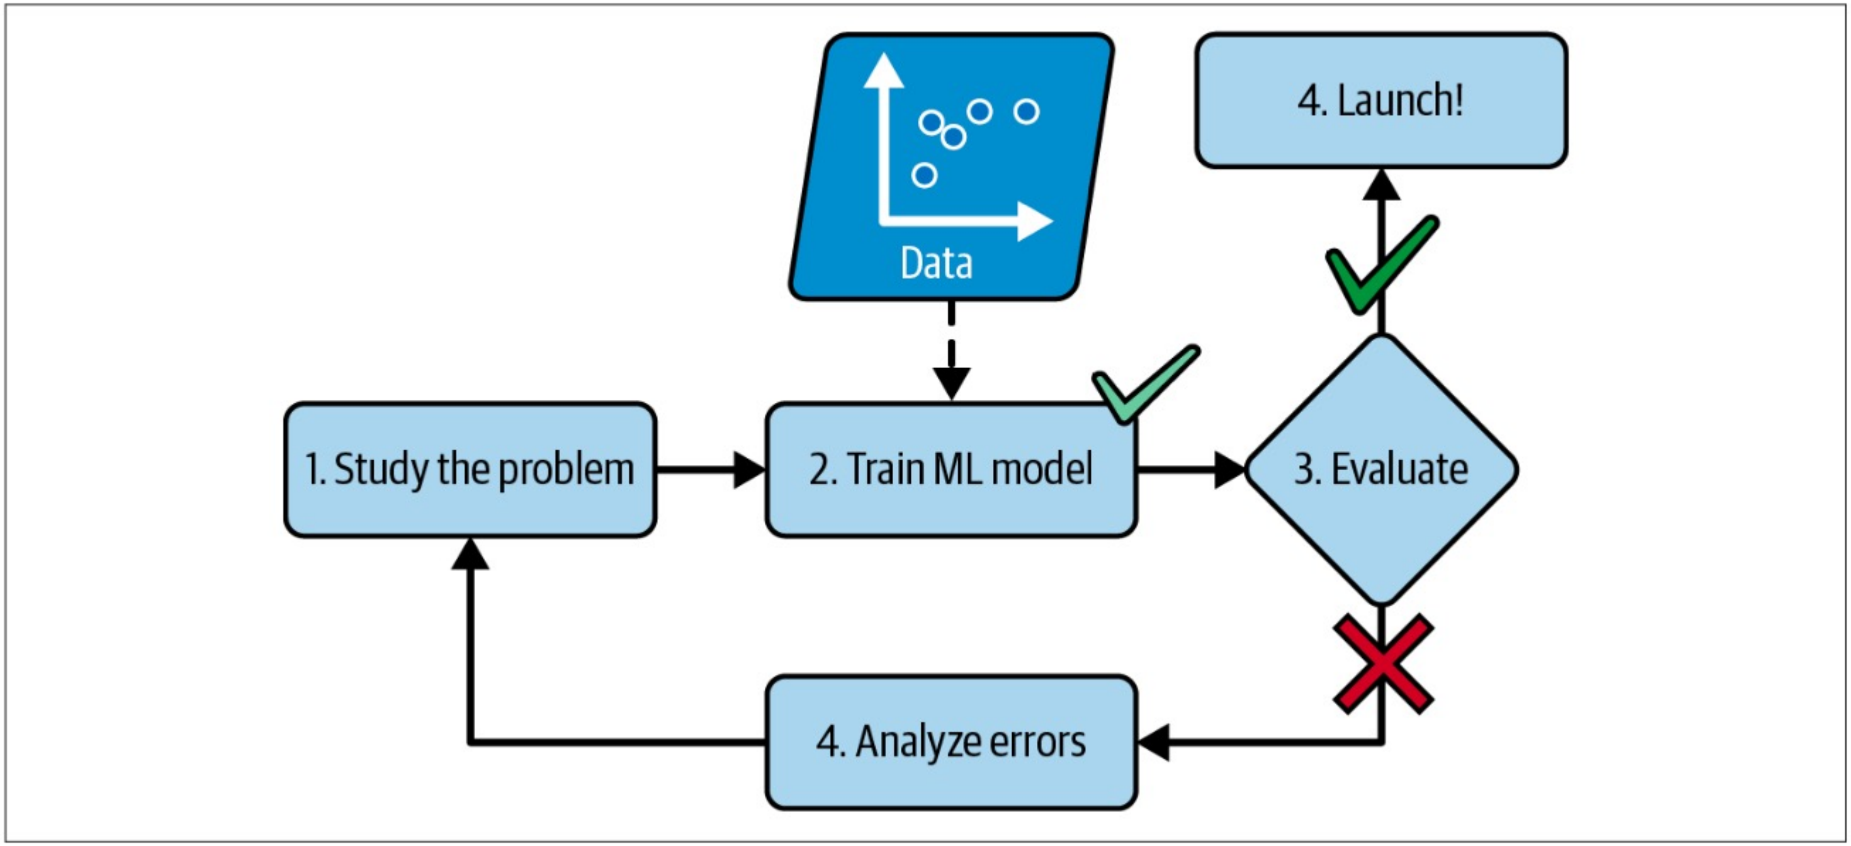

Figure 1-2. The machine learning approach

Geron (textbook)

Adaptive machine learning approach

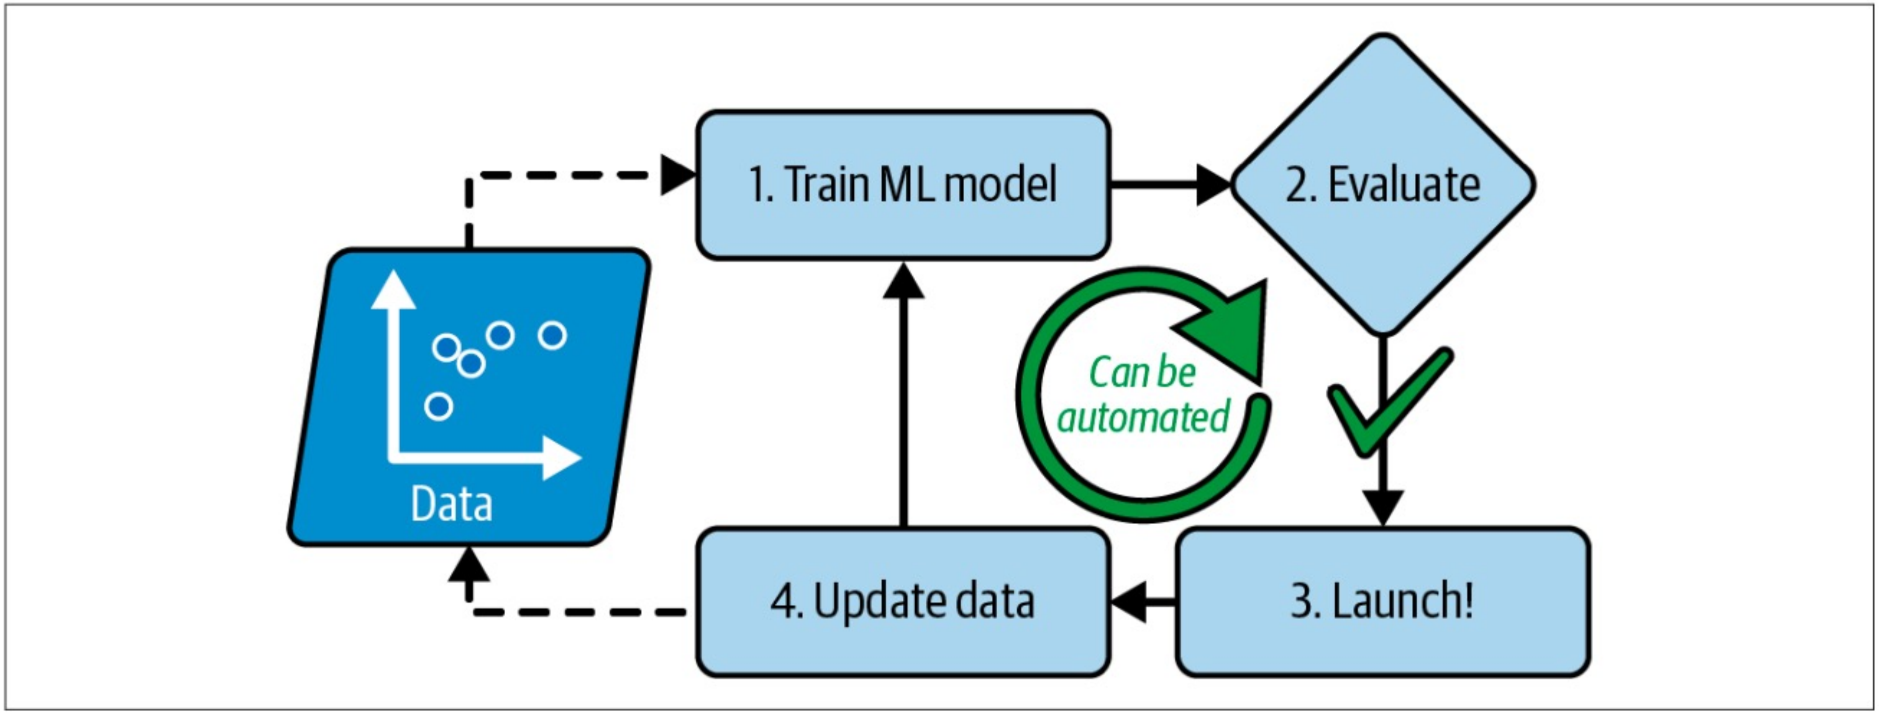

Figure 1-3. Automatically adapting to change

Geron (textbook)

Best case scenario: we learn something!

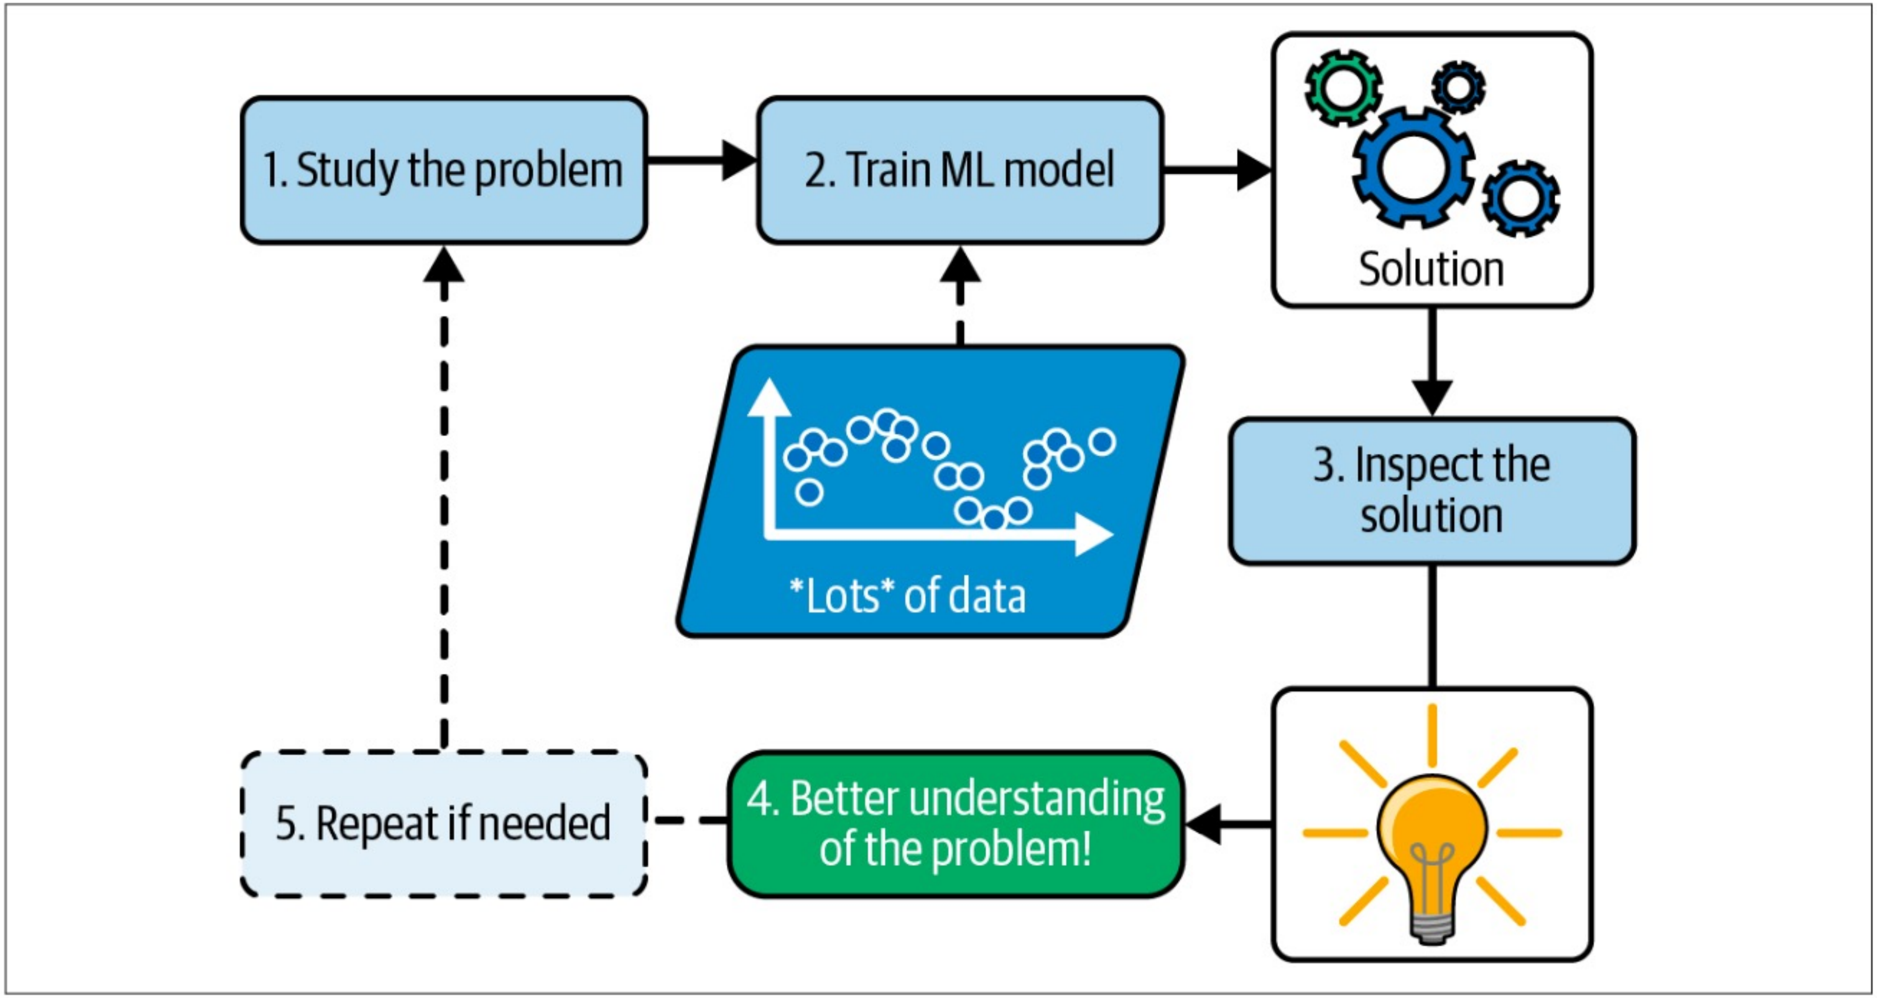

Figure 1-4. Machine learning can help humans learn

Geron (textbook)

Another view of the ML pipeline

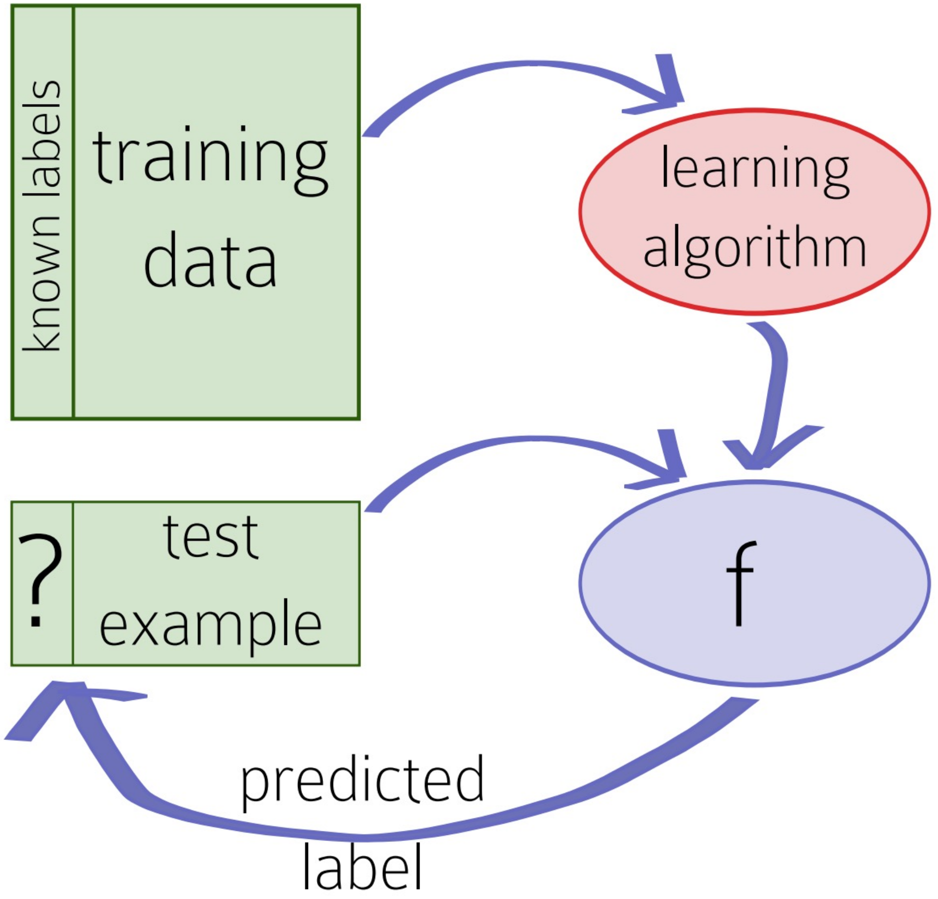

Figure 1.1, Duame

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

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

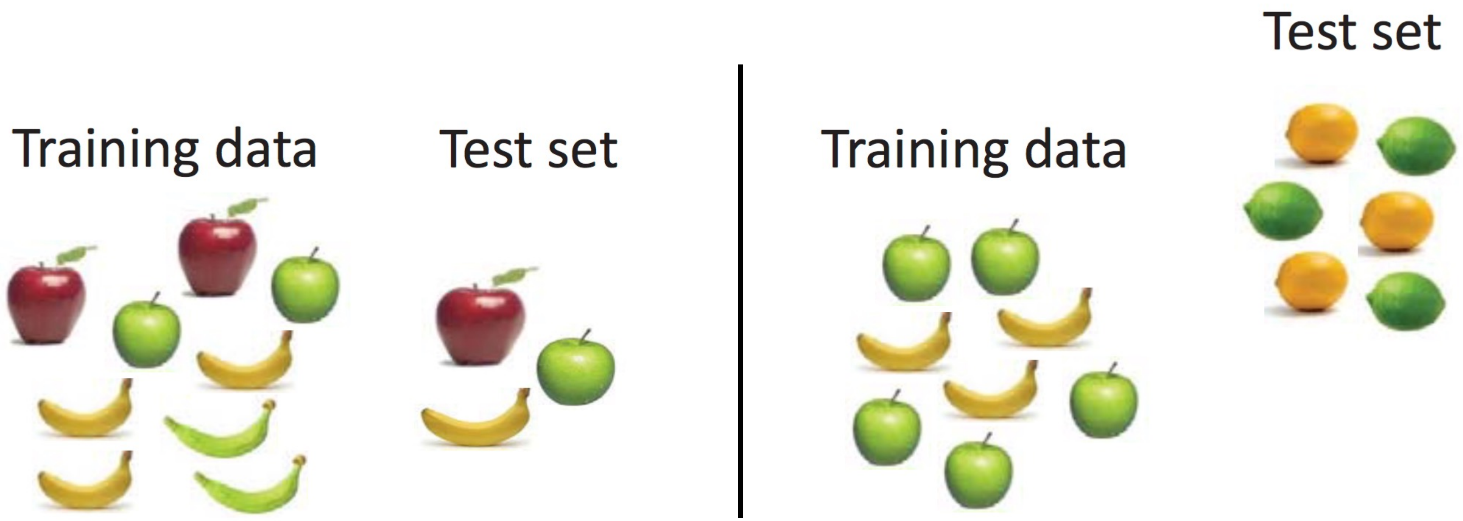

Training data

Test set

Training data

Test set

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

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*

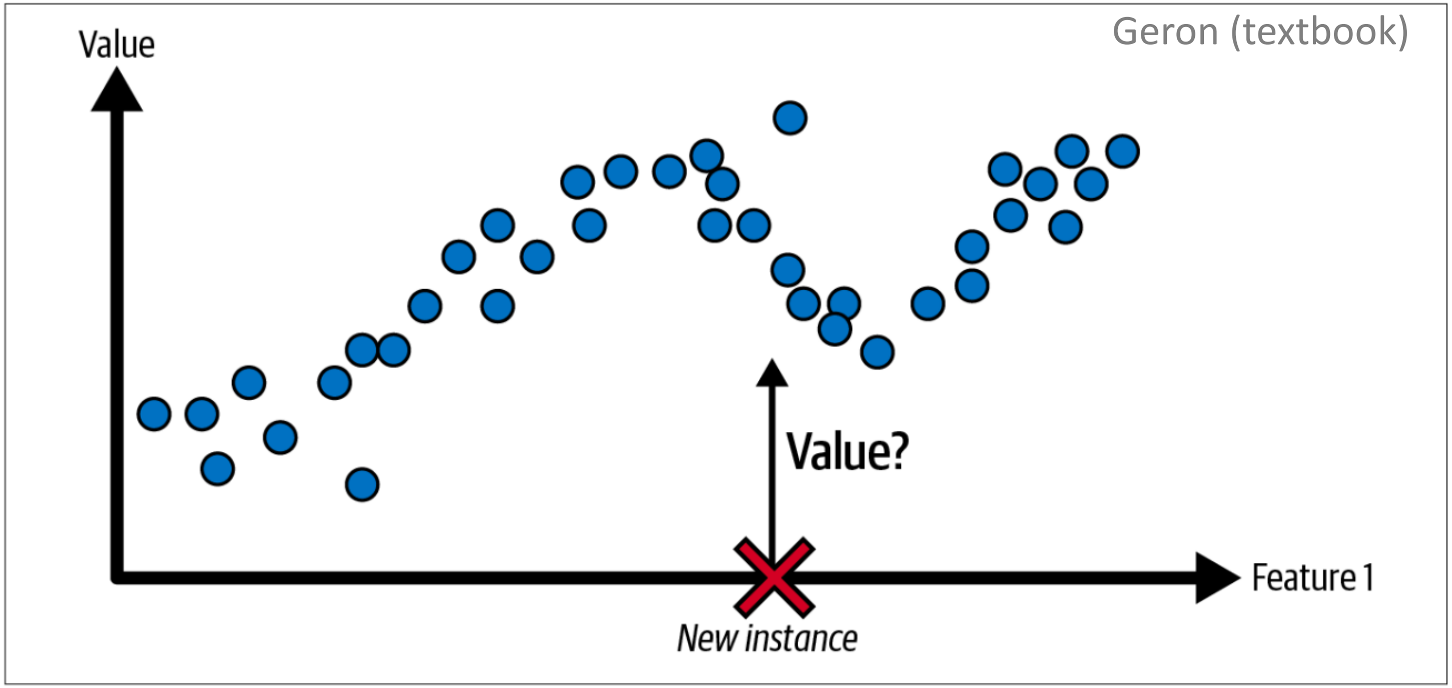

Geron (textbook)

Figure 1-6. A regression problem: predict a value, given an input feature (there are usually multiple input features, and sometimes multiple output values)

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)

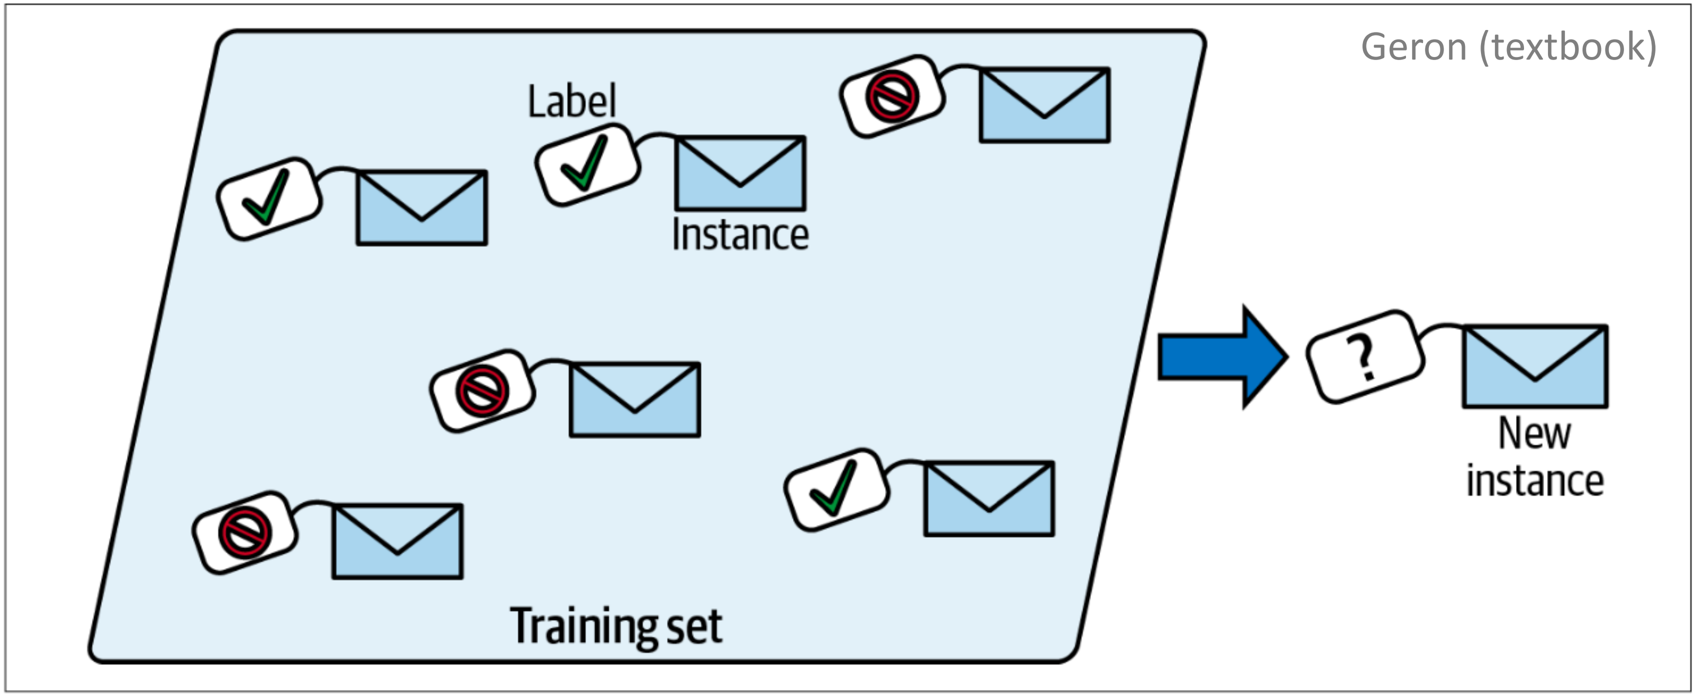

Geron (textbook)

Figure 1-5. A labeled training set for spam classification (an example of supervised learning)

Example: the classic bagel v. dog challenge

How can we *distinguish* between similar objects?

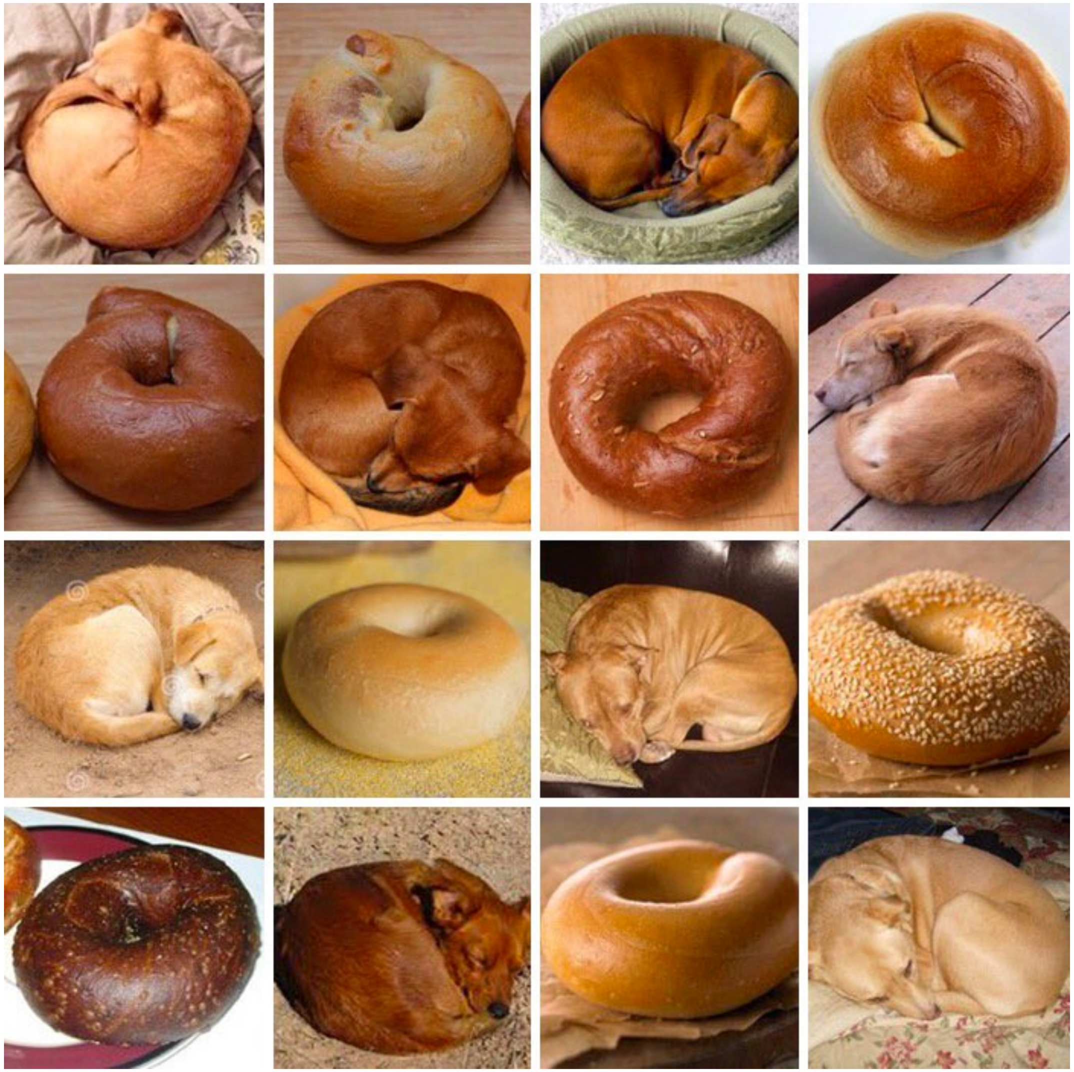

Images: veriy.com

Outline for Jan 23

- Welcome and syllabus highlights
- Classification review
- Evaluation metrics review
- Naïve Bayes review
- Thursday: gradient descent review

Binary Classification: asymmetric examples

- Computing the probability an email message is **spam**, given the **words** of the email

- Another example: what is the probability of **Trisomy 21** (Down Syndrome), given the amount of sequencing of each chromosome?

- Credit card fraud given timing and type of transactions

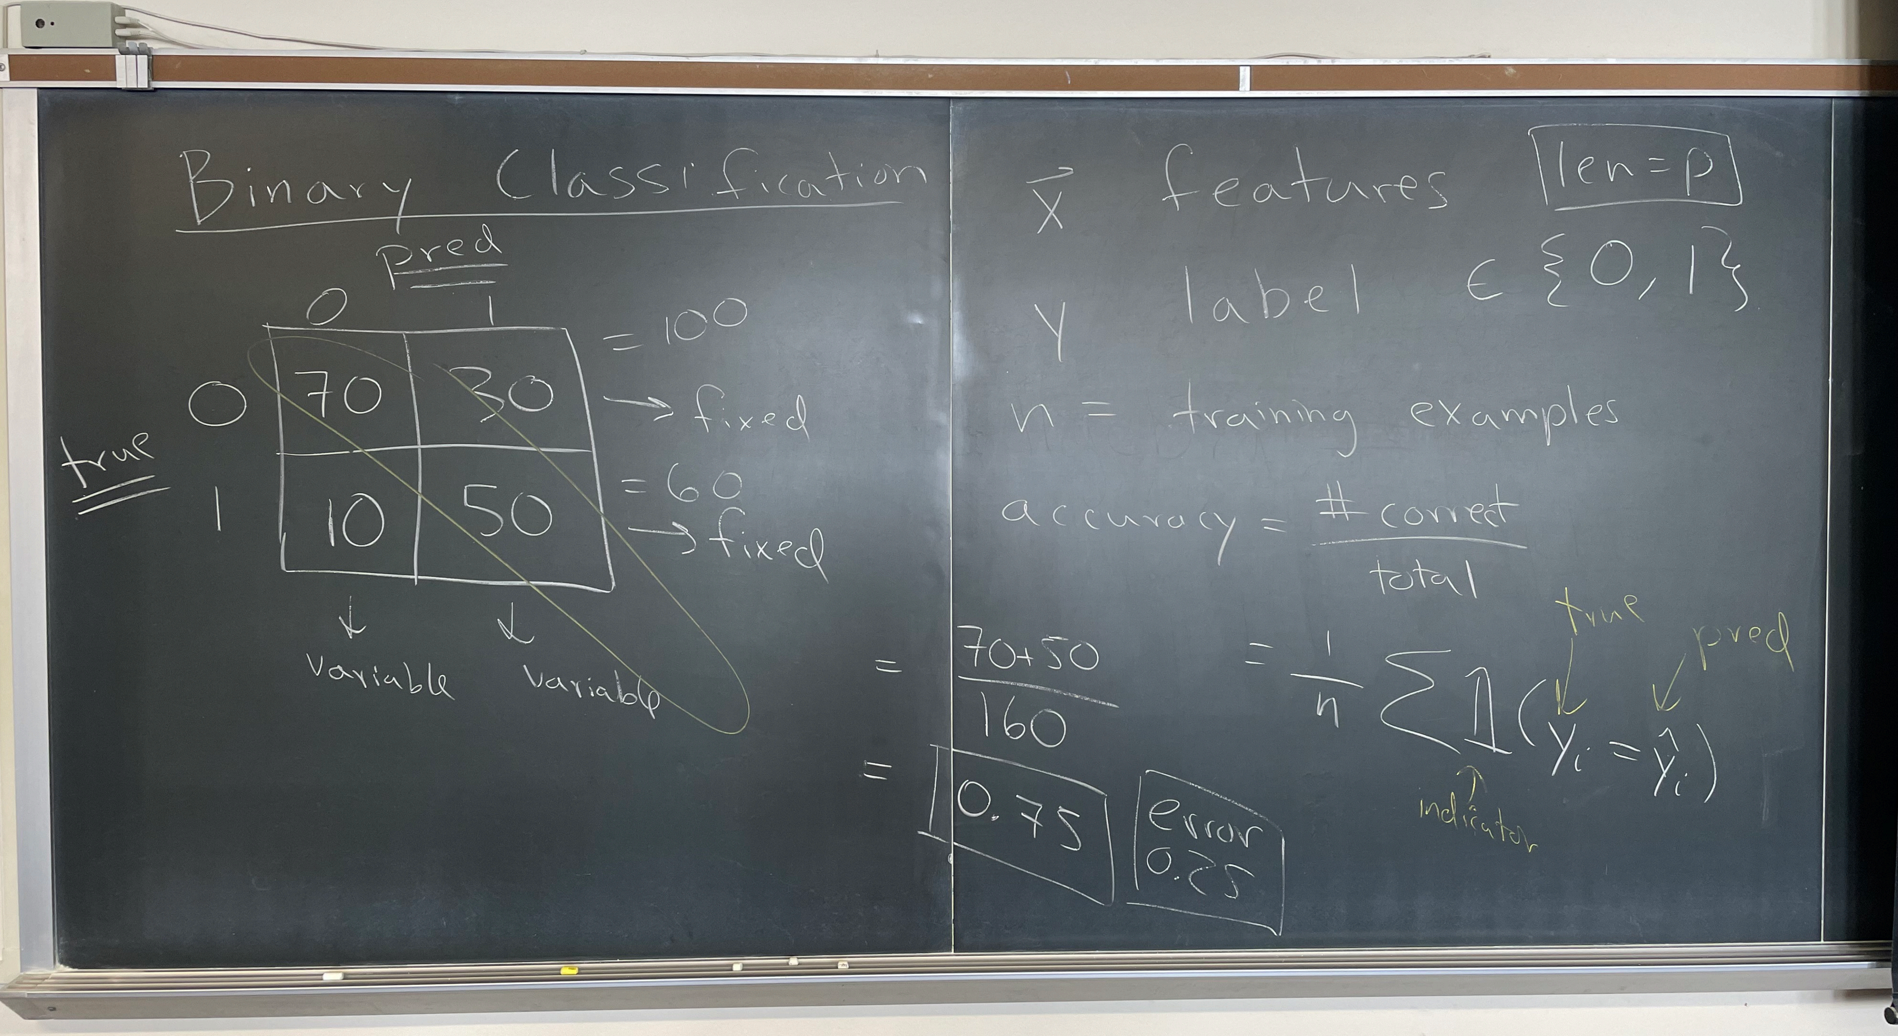

Binary Classification

pred

		0	1	
true	0	70	30	= 100 → fixed
	1	10	50	= 60 → fixed
		↓ Variable	↓ Variable	

$$= \frac{70+50}{160} = \boxed{0.75} \quad \boxed{\begin{array}{l}\text{error}\\ 0.25\end{array}}$$

$\vec{x}$ features $\boxed{\text{len} = p}$

y label $\in \{0, 1\}$

n = training examples

$$\text{accuracy} = \frac{\#\text{ correct}}{\text{total}} = \frac{1}{n} \sum \underset{\text{indicator}}{\mathbb{1}} \left(\overset{\text{true}}{y_i} = \overset{\text{pred}}{\hat{y}_i} \right)$$

Normalize

$$\begin{array}{c|cc|l} & 0 & 1 & \\ \hline 0 & 0.7 & 0.3 & \rightarrow \text{sum to } 1 \\ 1 & 0.17 & 0.83 & \rightarrow \text{sum to } 1 \end{array}$$

Confusion Matrices

		Predicted class	
		Negative	Positive
True class	Negative	True negative (TN)	False positive (FP)
	Positive	False negative (FN)	True positive (TP)

Modified from Jessica Wu

Confusion Matrices

		Predicted class		
		Negative	Positive	
True class	Negative	True negative (TN)	False positive (FP) “false alarm”	N (total number of true negatives)
	Positive	False negative (FN) “miss”	True positive (TP)	P (total number of true positives)
		N* (what we said was negative)	P* (what we said was positive “flagged”)	

Modified from Jessica Wu

Confusion Matrices

		Predicted class		
		Negative	Positive	
True class	Negative	True negative (TN) ✓	False positive (FP) “false alarm” ✗	N
	Positive	False negative (FN) “miss” ✗	True positive (TP) ✓	P
		N*	P*	

Modified from Jessica Wu

Confusion Matrices

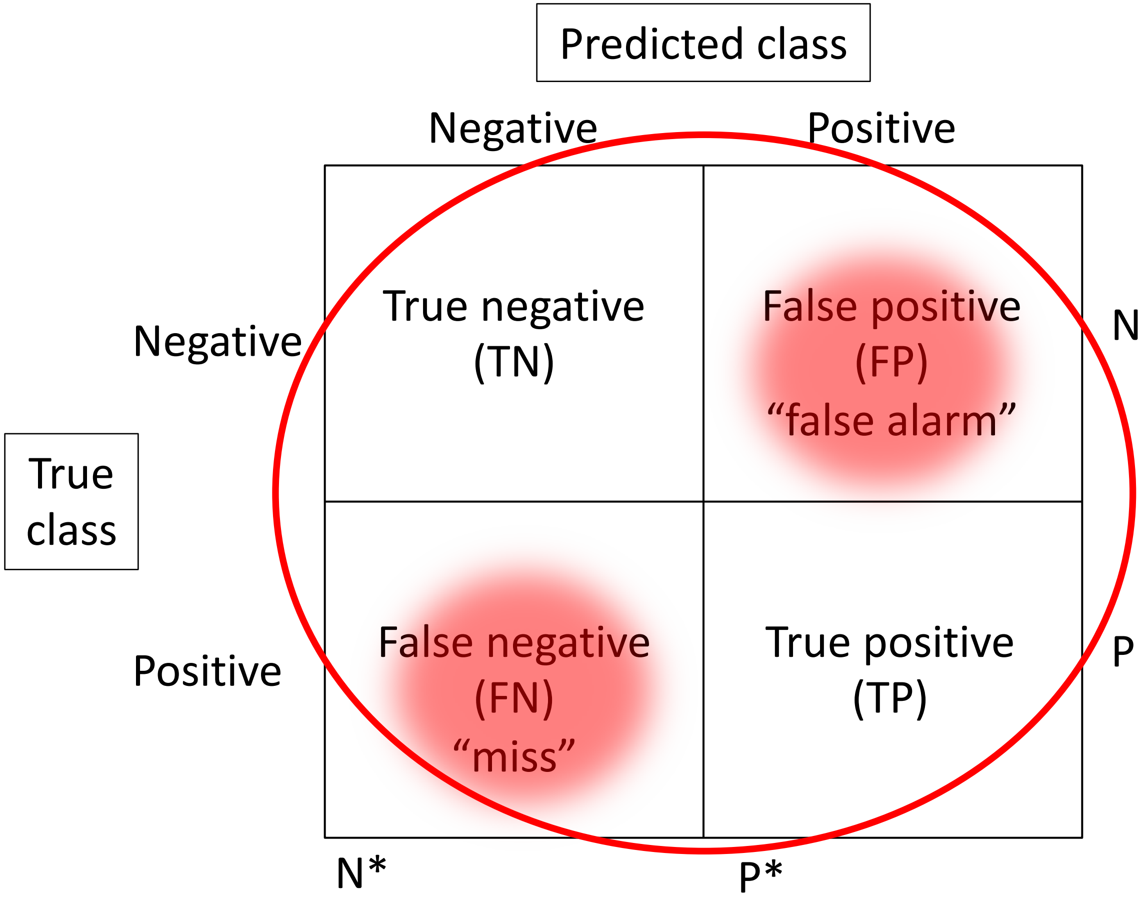

Predicted class

True class

		Negative	Positive	
True class	Negative	True negative (TN)	False positive (FP) “false alarm”	N
	Positive	False negative (FN) “miss”	True positive (TP)	P
		N*	P*	

Error:

$$(FN+FP)/(TN+FP+FN+TP)$$

$$= (FN+FP)/(N+P)$$

Modified from Jessica Wu

Confusion Matrices

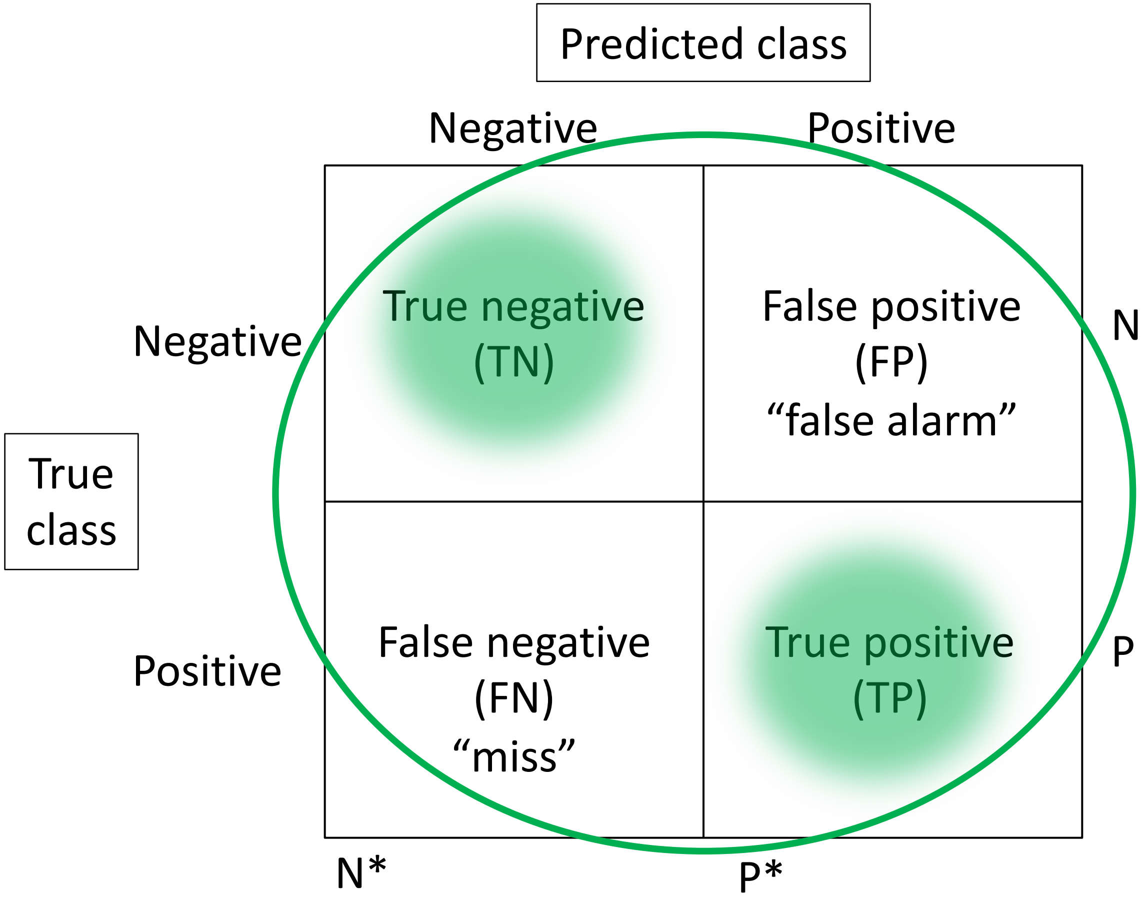

		Predicted class		
		Negative	Positive	
True class	Negative	True negative (TN)	False positive (FP) “false alarm”	N
	Positive	False negative (FN) “miss”	True positive (TP)	P
		N*	P*	

Accuracy = 1-Error:

(TN+TP)/(TN+FP+FN+TP)

= (TN+TP)/(N+P)

Modified from Jessica Wu

Confusion Matrices

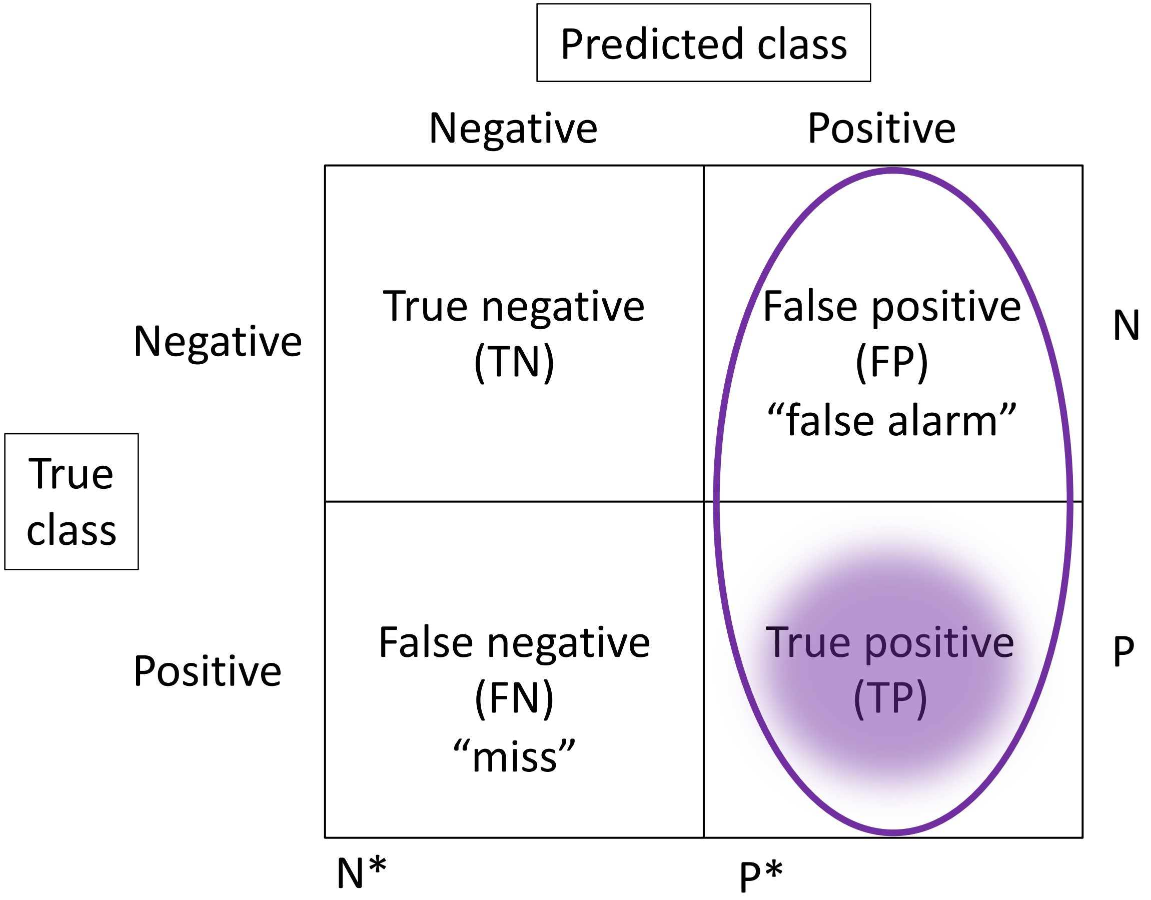

		Predicted class		
		Negative	Positive	
True class	Negative	True negative (TN)	False positive (FP) “false alarm”	N
	Positive	False negative (FN) “miss”	True positive (TP)	P
		N*	P*	

Precision:

TP/(FP+TP) = TP/P*

Modified from Jessica Wu

Confusion Matrices

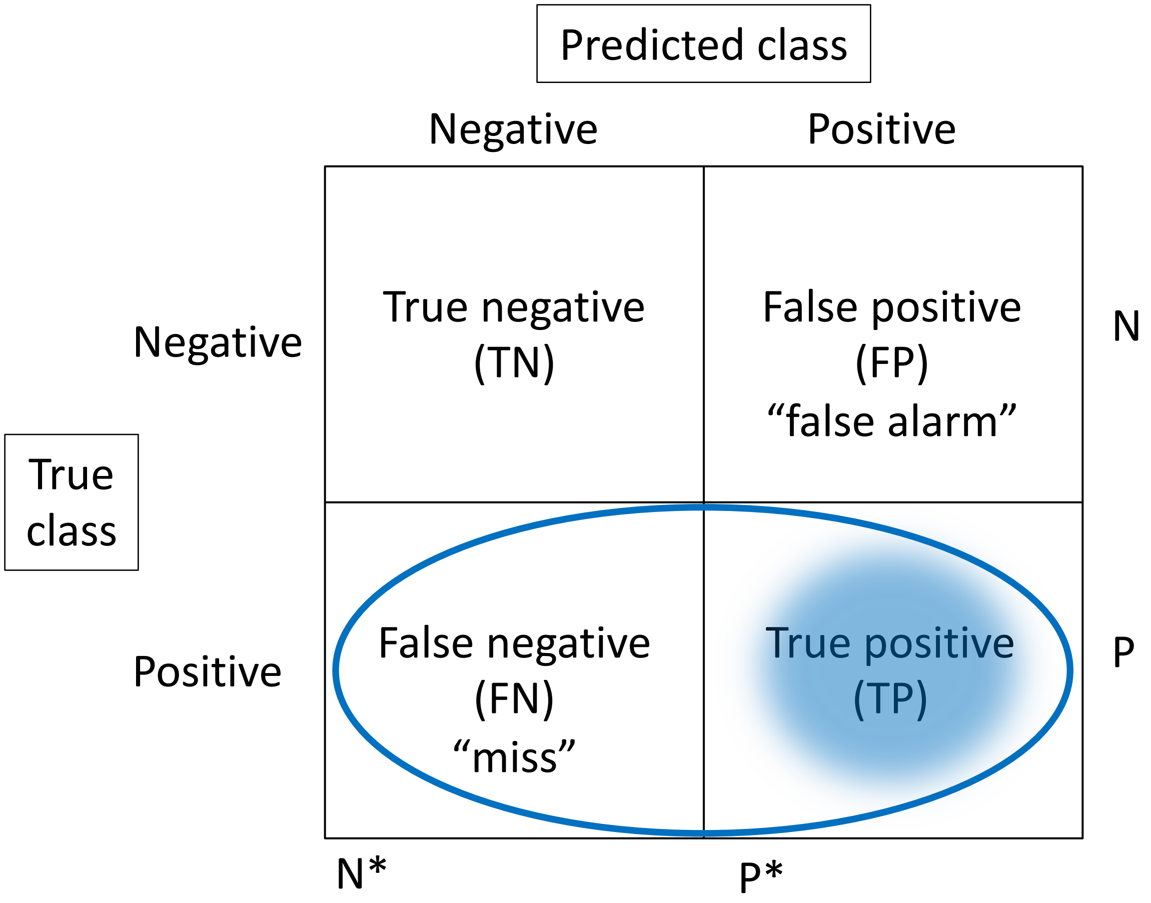

		Predicted class		
		Negative	Positive	
True class	Negative	True negative (TN)	False positive (FP) “false alarm”	N
	Positive	False negative (FN) “miss”	True positive (TP)	P
		N*	P*	

Recall
(True Positive Rate):

TP/(FN+TP) = TP/P

Modified from Jessica Wu

Confusion Matrices

		Predicted class		
		Negative	Positive	
True class	Negative	True negative (TN)	False positive (FP) “false alarm”	N
	Positive	False negative (FN) “miss”	True positive (TP)	P
		N*	P*	

False Positive Rate:

FP/(TN+FP) = FP/N

Modified from Jessica Wu

Handout 1, Q1

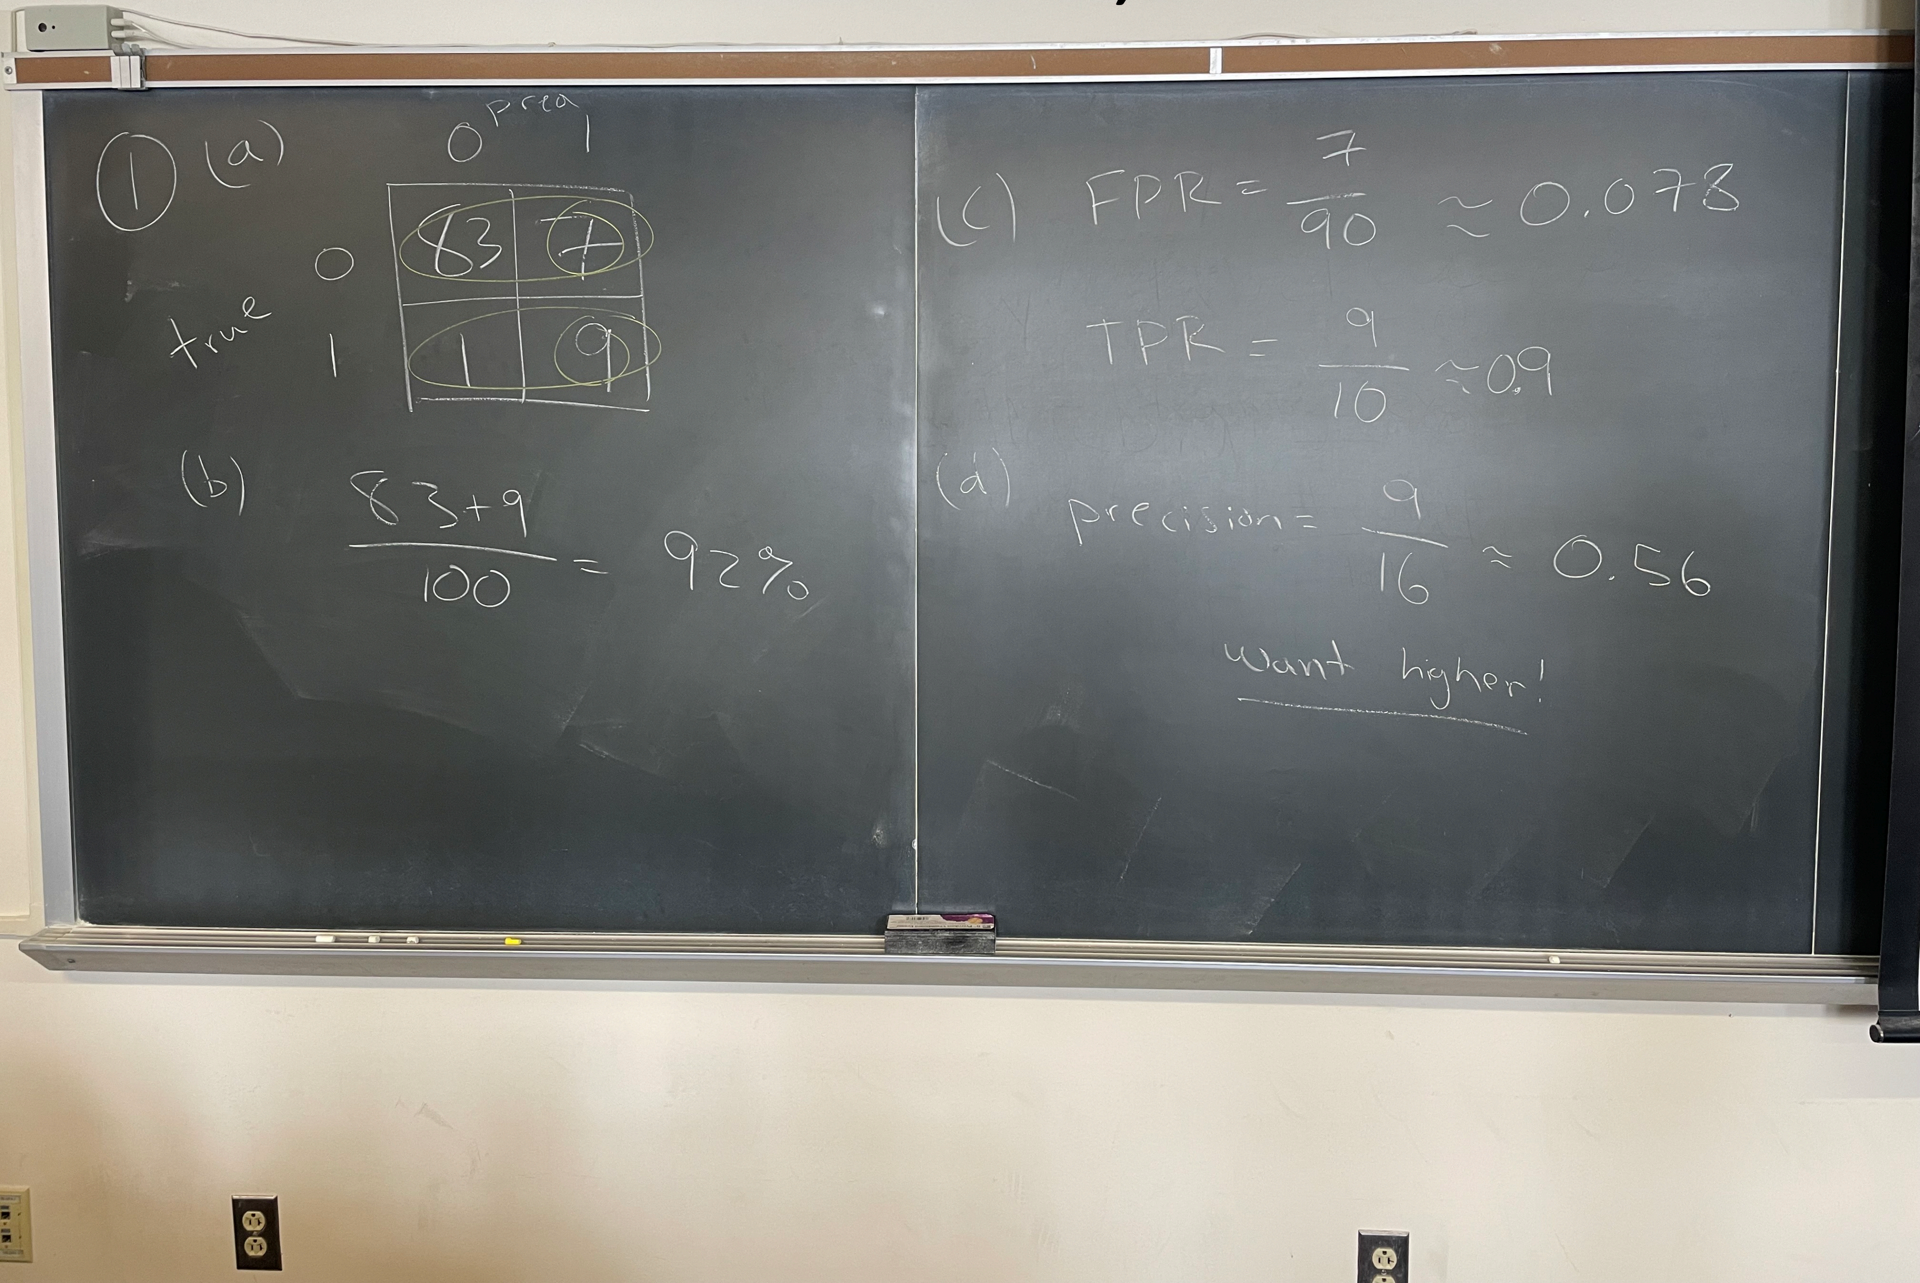

1 (a)

$$\begin{array}{cc|cc} & & \text{pred} & \\ & & 0 & 1 \\ \hline \text{true} & 0 & 83 & 7 \\ & 1 & 1 & 9 \end{array}$$

(b) $\frac{83+9}{100} = 92\%$

(c) $\text{FPR} = \frac{7}{90} \approx 0.078$

$\text{TPR} = \frac{9}{10} \approx 0.9$

(d) $\text{precision} = \frac{9}{16} \approx 0.56$

want higher!

Outline for Jan 23

- Welcome and syllabus highlights
- Classification review
- Evaluation metrics review
- Naïve Bayes review
- Thursday: gradient descent review

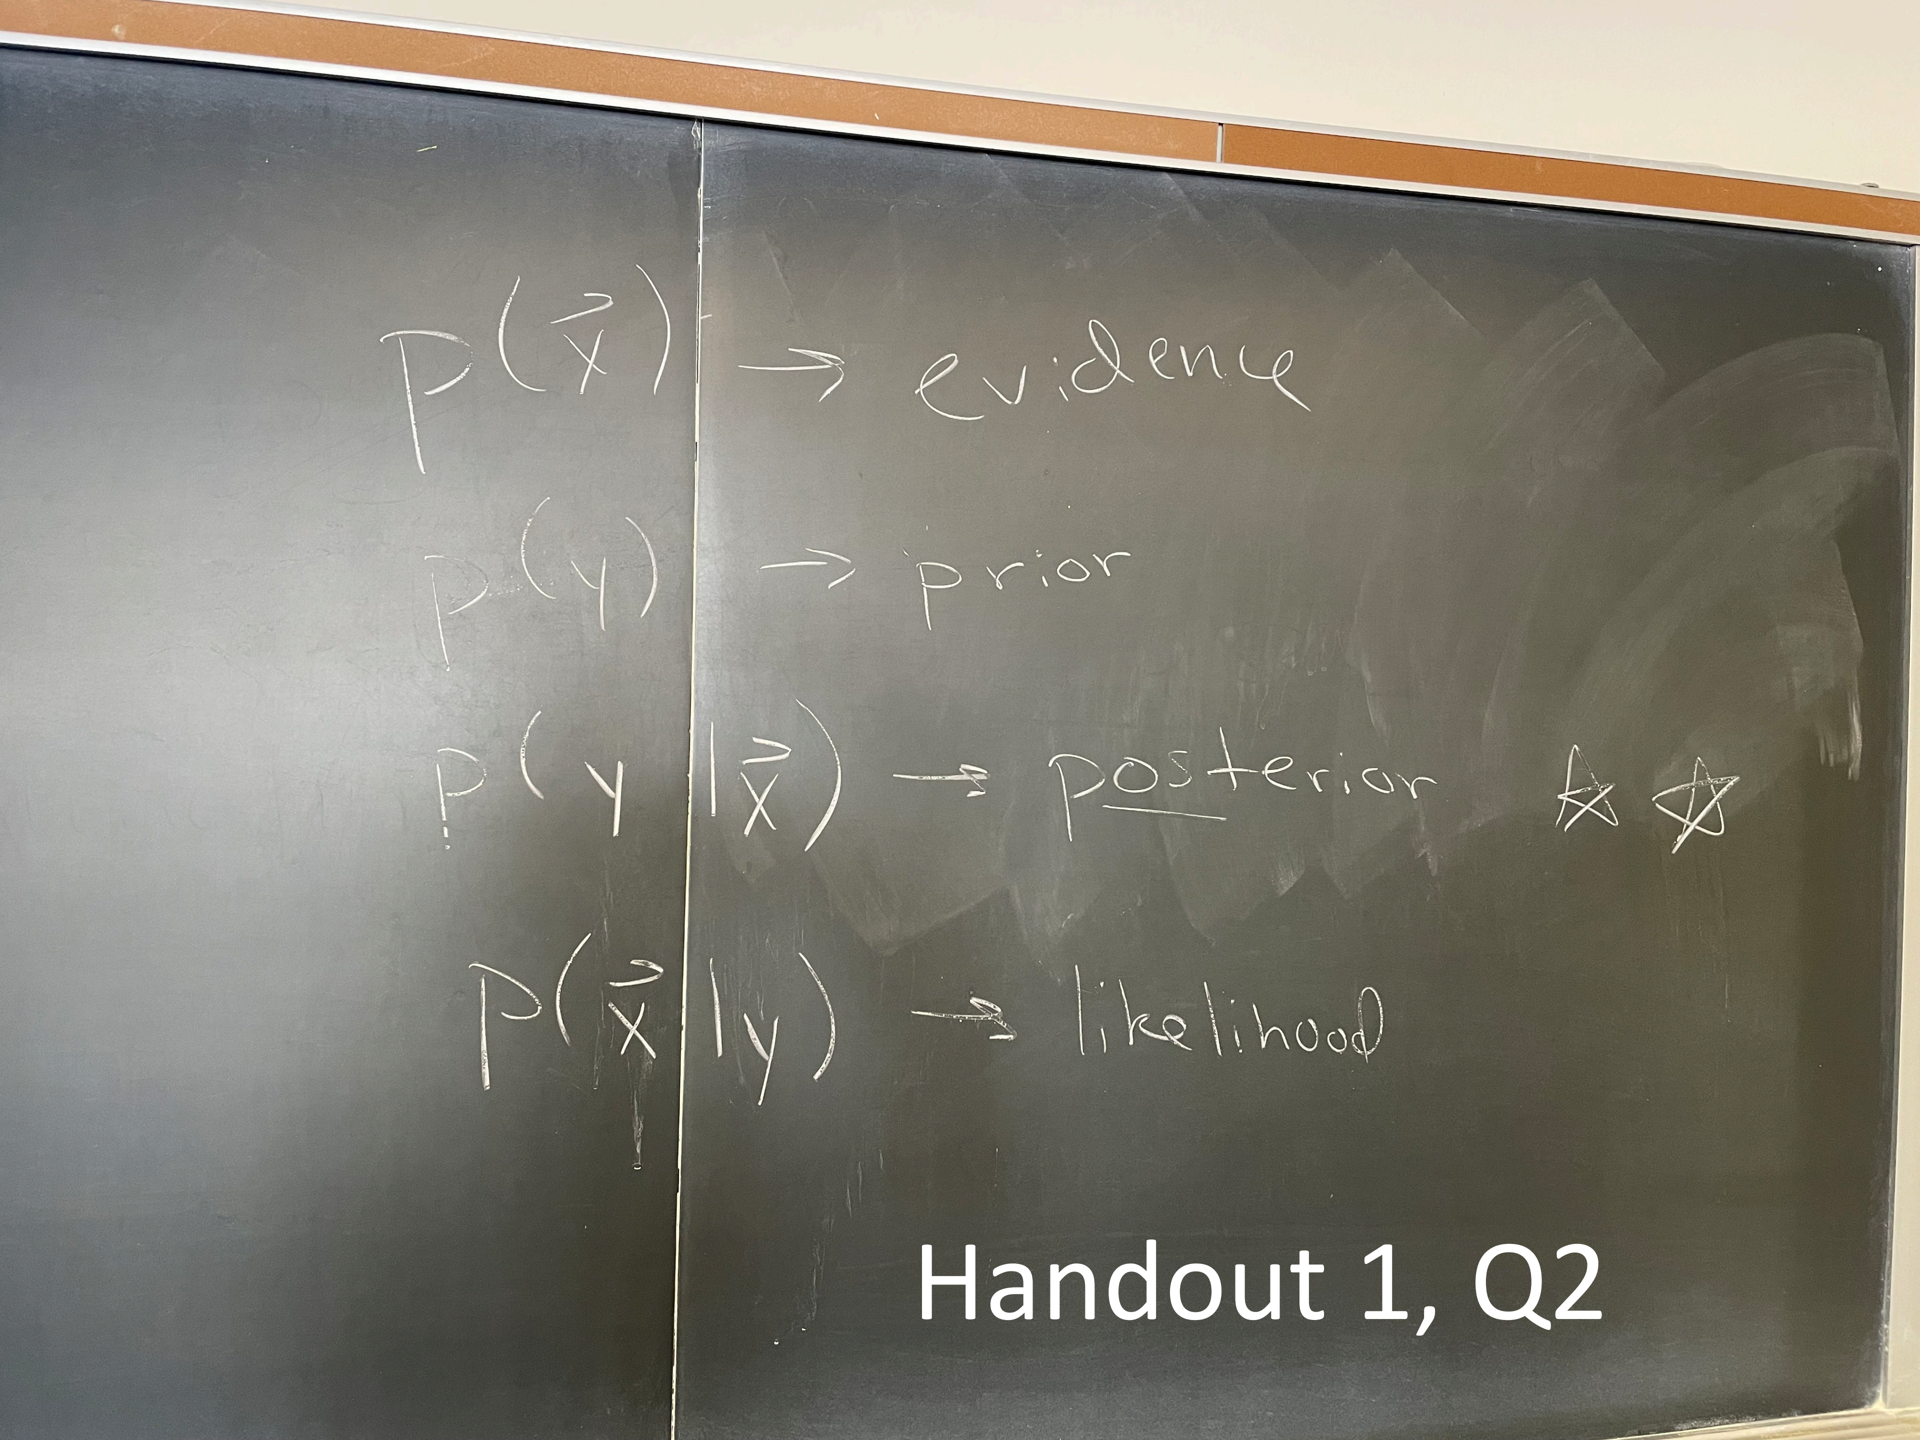

$p(\vec{x})$ → evidence

$p(y)$ → prior

$p(y|\vec{x})$ → posterior ☆ ☆

$p(\vec{x}|y)$ → likelihood

Handout 1, Q2

Components of a Bayesian Model

- Identify the evidence, prior, posterior, and likelihood in the equation below

$$p(y = k|\boldsymbol{x}) = \frac{p(y = k)p(\boldsymbol{x}|y = k)}{p(\boldsymbol{x})}$$

Components of a Bayesian Model

- Identify the evidence, prior, posterior, and likelihood in the equation below

$$p(y=k|\boldsymbol{x}) = \frac{p(y=k)p(\boldsymbol{x}|y=k)}{p(\boldsymbol{x})}$$

- **Evidence**: this is the data (features) we actually observe, which we think will help us predict the outcome we're interested in

Components of a Bayesian Model

- Identify the evidence, prior, posterior, and likelihood in the equation below

$$p(y = k|\boldsymbol{x}) = \frac{p(y = k)p(\boldsymbol{x}|y = k)}{p(\boldsymbol{x})}$$

- **Prior**: without seeing any evidence (data), what is our prior believe about each outcome (intuition: what is the outcome in the population as a whole?)

Components of a Bayesian Model

- Identify the evidence, prior, posterior, and likelihood in the equation below

$$p(y = k|\boldsymbol{x}) = \frac{p(y = k)p(\boldsymbol{x}|y = k)}{p(\boldsymbol{x})}$$

- **Posterior**: this is the quantity we are actually interested in. **Given** the evidence, what is the probability of the outcome?

Components of a Bayesian Model

- Identify the evidence, prior, posterior, and likelihood in the equation below

$$p(y=k|\boldsymbol{x}) = \frac{p(y=k)p(\boldsymbol{x}|y=k)}{p(\boldsymbol{x})}$$

- **Likelihood**: given an outcome, what is the probability of observing this set of features?

Reading for Week 1

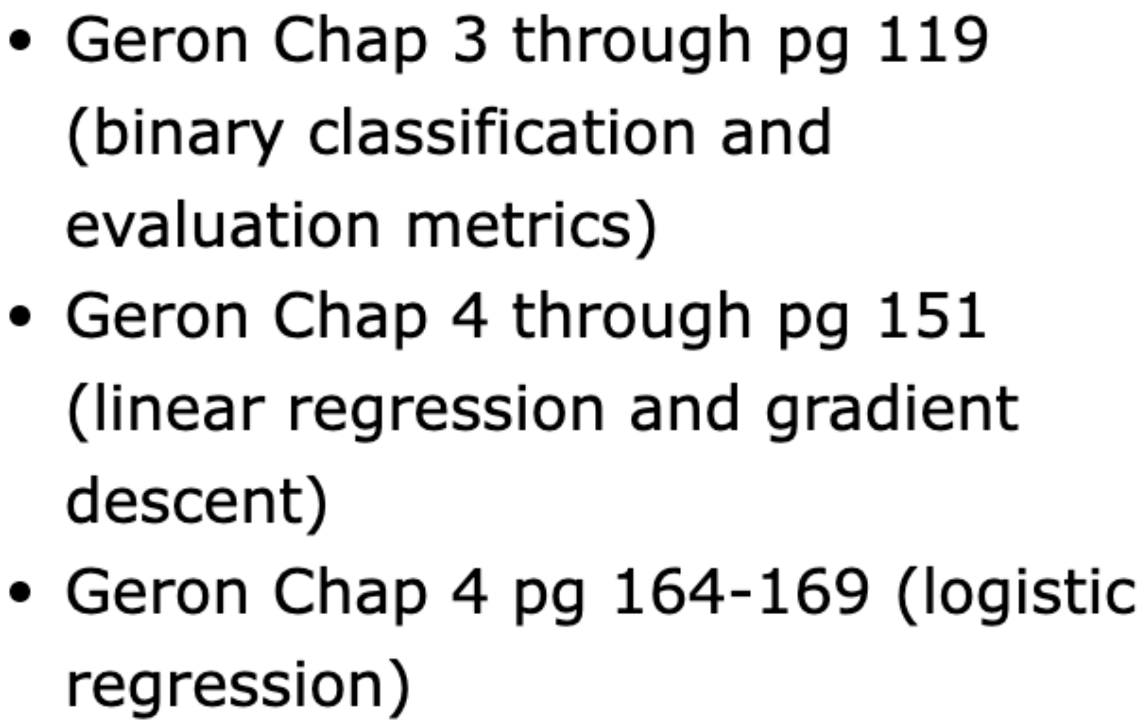

- Geron Chap 3 through pg 119 (binary classification and evaluation metrics)
- Geron Chap 4 through pg 151 (linear regression and gradient descent)
- Geron Chap 4 pg 164-169 (logistic regression)