

As you walk in

Welcome!

- 1) Help draw some aliens for our dataset today!
 - Grab a sharpie and some sticky notes
 - Just stick figures, nothing quite this fancy →

Ca



Freepik
Suesse Aliens Bilder - ...



123RF
91,690 Cartoon Alien...



iStock
47,434 Alien Cartoon...



Shutterstock
228,873 Alien Cartoon...



Shutterstock
2,607 Alien Feet Image...



Shutterstock
228,873 Alien Cartoon...


Exercise: Human-defined Feature Space


Try to organize data on a 2-D coordinate plot

to win the following game:

Select three students: A, B, C

1. Student A draws a new alien and hands it to student B
2. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
3. Student C looks at the coordinate and the plot (but not the drawing from A) and **draws a new alien**



An abstract graphic on the left side of the slide, featuring a sphere-like shape composed of a dense grid of intersecting red, green, and blue lines. The lines are curved and follow the contour of the sphere, creating a complex, woven pattern. The sphere is set against a dark gray background.

10-315 Introduction to ML

Instructor: Pat Virtue

Today

Aliens

Course team

What is AI and ML?

ML framework

Example ML models

- Autoencoder (Aliens)

Piazza and website for course info



DALL-E: “Logo of a Scotty dog with a red collar whose brain is made of circuits”

Course Information

Website: <https://www.cs.cmu.edu/~10315>

Communication:

piazza.com



E-mail Joshmin if piazza doesn't work:

joshminr@andrew.cmu.edu

Important posts coming soon!

- Office Hours
- Jupyter notebooks
- Recitation (Fri)
- Pre-reading (due Sun)
- HW0
- HW1

Course Team

Instructor



Pat
Virtue
pvirtue

Best way to contact Pat:

- Post on Piazza (private post as needed)
- Just grab an OH appointment slot

Education Associate



Joshmin Ray
joshminr

Email Joshmin, joshminr@andrew.cmu.edu, for any:

- Exceptions, extensions, etc
- Any course logistics

Course Team

Teaching Assistants



Alex
alextiar




Vincent
vtombari



Deep
dmpatel



Tanay
tbennur



Medha
mpalaval



Myles
msharris



Shreeya
srkhan



Anisha
anishac



Ruthie
rylin



Ethan
ethanwan

Course Team

Students!!



"Poll"

Which of these are artificial?

Which of these are more intelligent?




AI Definition by John McCarthy

What is artificial intelligence

- It is the science and engineering of making intelligent machines, especially intelligent computer programs

What is intelligence

- Intelligence is the computational part of the ability to achieve goals in the world



Attributes of Intelligence

"Poll" Turn to your neighbor

Which of these more intelligent?

- Robot that assemble cars in factory
- Robot that fold clothes



Intelligence and Uncertainty

Uncertainty can have lots of sources, including anything we attribute to random chance

- Hidden information
 - Cards in another player's hand
- Noise
 - Sensor noise
- Way too complicated to model
 - Leaves blowing in the wind
- Infinite number of possible configurations
- More possibilities than any computer can compute in a reasonable time
 - Tic-tac-toe → Checkers → Chess


AI Definition by John McCarthy

What is artificial intelligence

- It is the science and engineering of making intelligent machines, especially intelligent computer programs

What is intelligence

- Intelligence is the computational part of the ability to achieve goals in the world



AI Definition (Pat's Version)

What is intelligence


- Intelligence is the ability to perform well on a task that involves uncertainty

Intelligence is not binary


- How **well** an agent performs and how much **uncertainty** is involved will determine how intelligent we consider the agent to be



Artificial Intelligence vs Machine Learning?



A Brief History of AI




Albert Einstein, Sherlock Holmes, Frankenstein

1800 - 2019 ▼


English (2019) ▼

Case-Insensitive

Smoothing ▼



A Brief History of AI



A Brief History of AI

1940-1950: Early days

- 1943: McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing's "Computing Machinery and Intelligence"

1950—70: Excitement: Look, Ma, no hands!

- 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956: Dartmouth meeting: "Artificial Intelligence" adopted

1970—90: Knowledge-based approaches

- 1969—79: Early development of knowledge-based systems
- 1980—88: Expert systems industry booms
- 1988—93: Expert systems industry busts: "AI Winter"

1990—: Statistical approaches


- Resurgence of probability, focus on uncertainty
- General increase in technical depth
- Agents and learning systems... "AI Spring"?

2012—: Deep learning

- 2012: ImageNet & AlexNet

2023—: Generative AI

- 2023: ChatGPT



ML Applications?

Speech Recognition

1. Learning to recognize spoken words

THEN	NOW
"...the SPHINX system (e.g. Lee 1989) learns speaker-specific strategies for recognizing the primitive sounds (phonemes) and words from the observed speech signal...neural network methods...hidden Markov models..."	

(Mitchell, 1997)

Source: <https://www.stonetemple.com/great-knowledge-box-showdown/#VoiceStudyResults>

8

Robotics

2. Learning to drive an autonomous vehicle

THEN	NOW
"...the ALVINN system (Pomerleau 1989) has used its learned strategies to drive unassisted at 70 miles per hour for 90 miles on public highways among other cars..."	

(Mitchell, 1997)

waymo.com

9

Games / Reasoning

3. Learning to beat the masters at board games


THEN	NOW
"...the world's top computer program for backgammon, TD-GAMMON (Tesauro, 1992, 1995), learned its strategy by playing over one million practice games against itself..."	

(Mitchell, 1997)

11

Computer Vision

4. Learning to recognize images

THEN	NOW
"...The recognizer is a convolution network that can be spatially replicated. From the network output, a hidden Markov model produces word scores. The entire system is globally trained to minimize word-level errors..."  (LeCun et al., 1995)	Revolution of Depth  

Images from <https://blog.openai.com/generative-models/>

12

Learning Theory

5. In what cases and how well can we learn?

Sample Complexity Results

Definition 6.1. The sample complexity of a learning algorithm is the number of examples required to achieve arbitrarily small error (with respect to the optimal hypothesis) with high probability (i.e. close to 1).

Four Cases we care about...

	Realizable	Agnostic
Finite [K]	$N \geq \frac{1}{\epsilon} \left(\frac{1}{\eta} \log \frac{1}{\delta} + \log \frac{1}{\eta} \right)$ Related examples are sufficient so that with probability $1 - \delta$, for all $h \in H$, we have that $R(h) \leq \epsilon$.	$N \geq \frac{1}{\epsilon^2} \left(\frac{1}{\eta} \log \frac{1}{\delta} + \log \frac{1}{\eta} \right)$ Related examples are sufficient so that with probability $1 - \delta$, for all $h \in H$, we have that $R(h) \leq \epsilon$.
Infinite [K]	$N \geq \frac{1}{\epsilon^2} \left(\frac{1}{\eta} \log \frac{1}{\delta} + \log \frac{1}{\eta} \right)$ Related examples are sufficient so that with probability $1 - \delta$, for all $h \in H$, we have that $R(h) \leq \epsilon$.	$N \geq \frac{1}{\epsilon^2} \left(\frac{1}{\eta} \log \frac{1}{\delta} + \log \frac{1}{\eta} \right)$ Related examples are sufficient so that with probability $1 - \delta$, for all $h \in H$, we have that $R(h) \leq \epsilon$.

PK Learning

Can we bound $R(h)$ in terms of $R(h)$?

A: Yes!

PK shows $R(h)$ is bounded by $R(h)$ plus a term that goes to zero as n increases.

Approximately Correct

Def: PK Criterion

$$R(h) - R(h) \leq \epsilon \Rightarrow 1 - \delta$$

Two Types of Error
① Test Error (aka expected risk) (aka Generalization Error) $R(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \ell(h(x), y)$ ② Train Error (aka empirical risk) $\hat{R}(h) = \frac{1}{N} \sum_{i=1}^N \ell(h(x_i), y_i)$

- How many examples do we need to learn?
- How do we quantify our ability to generalize to unseen data?
- Which algorithms are better suited to specific learning settings?

13

Machine Learning and Statistics

Statistics is also about learning from data

Statistics has been around from much longer!

What's the difference?

Until the mid 1990s:

Statistics:

- A branch of mathematics
- Emphasized rigor, correctness, provable properties (“is it correct?”)
- Was not very concerned with scaling

Machine Learning:

- A branch of Computer Science / AI
- Focus on heuristics, making things work in practice (“does it work?”)
- Not much awareness of statistical theory

As you walk in

Welcome!

1) Help draw some aliens for our dataset today!

- See table up front
- Just stick figures, nothing quite this fancy →

Cartoon alien hi-res...



Freepik
Suesse Aliens Bilder - ...



123RF
91,690 Cartoon Alien...



iStock
47,434 Alien Cartoon...



Shutterstock
228,873 Alien Cartoon...



Shutterstock
2,607 Alien Feet Image...



Shutterstock
228,873 Alien Cartoon...


Exercise: Human-defined Feature Space

Try to organize data on a 2-D coordinate plot

to win the following game:

Select three students: A, B, C

1. Student A draws a new alien and hands it to student B
2. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
3. Student C looks at the coordinate and the plot (but not the drawing from A) and **draws a new alien**




Poll 1

What features did you use?


Learning to Organize Data

Neural networks can learn to organization too!

Image $\rightarrow \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \rightarrow$ Image



z_2




z_1

<https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html>


Digit Autoencoder

Demo: Using a learned feature space



Task: Face Generation

<https://thispersondoesnotexist.com/>




Task: Face Generation

<https://thispersondoesnotexist.com/>


Random

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$$



Random

$$\begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$$



Face GAN Slider Demo

<https://github.com/genforce/interfacegan>

<https://colab.research.google.com/github/genforce/interfacegan/blob/master/docs/InterFaceGAN.ipynb>

age:

eyeglasses:

gender:

pose:

smile:

[Show code](#)



Face GAN Slider Demo

<https://github.com/genforce/interfacegan>

<https://colab.research.google.com/github/genforce/interfacegan/blob/master/docs/InterFaceGAN.ipynb>

age: -3

eyeglasses: 2.1

gender: 0

pose: 0

smile: 0

[Show code](#)




The image displays a user interface for a Face GAN slider demo. It features five horizontal sliders for adjusting face attributes: age, eyeglasses, gender, pose, and smile. Each slider has a blue dot indicating the current value, and a numerical value is shown to the right of each slider. Below the sliders is a blue link labeled 'Show code'. At the bottom, there are four generated face images arranged horizontally, showing the results of the sliders: a man with glasses, a woman with pink hair, a girl with dark hair, and a boy with dark hair.

Attribute	Value
age	-3
eyeglasses	2.1
gender	0
pose	0
smile	0

Listen Learner

Chris Harrison, CMU




<https://chrisharrison.net/index.php/Research/ListenLearner>

Dimensionality Reduction with Deep Learning

Hinton, Geoffrey E., and Ruslan R. Salakhutdinov.

"Reducing the dimensionality of data with neural networks."



Science 313.5786 (2006): 504-507.



Digit Autoencoder

<https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html>

Input
28x28 = 784 pixels



Digit Autoencoder

Math for autoencoder model

Linear layer

ReLU layer

Tanh layer