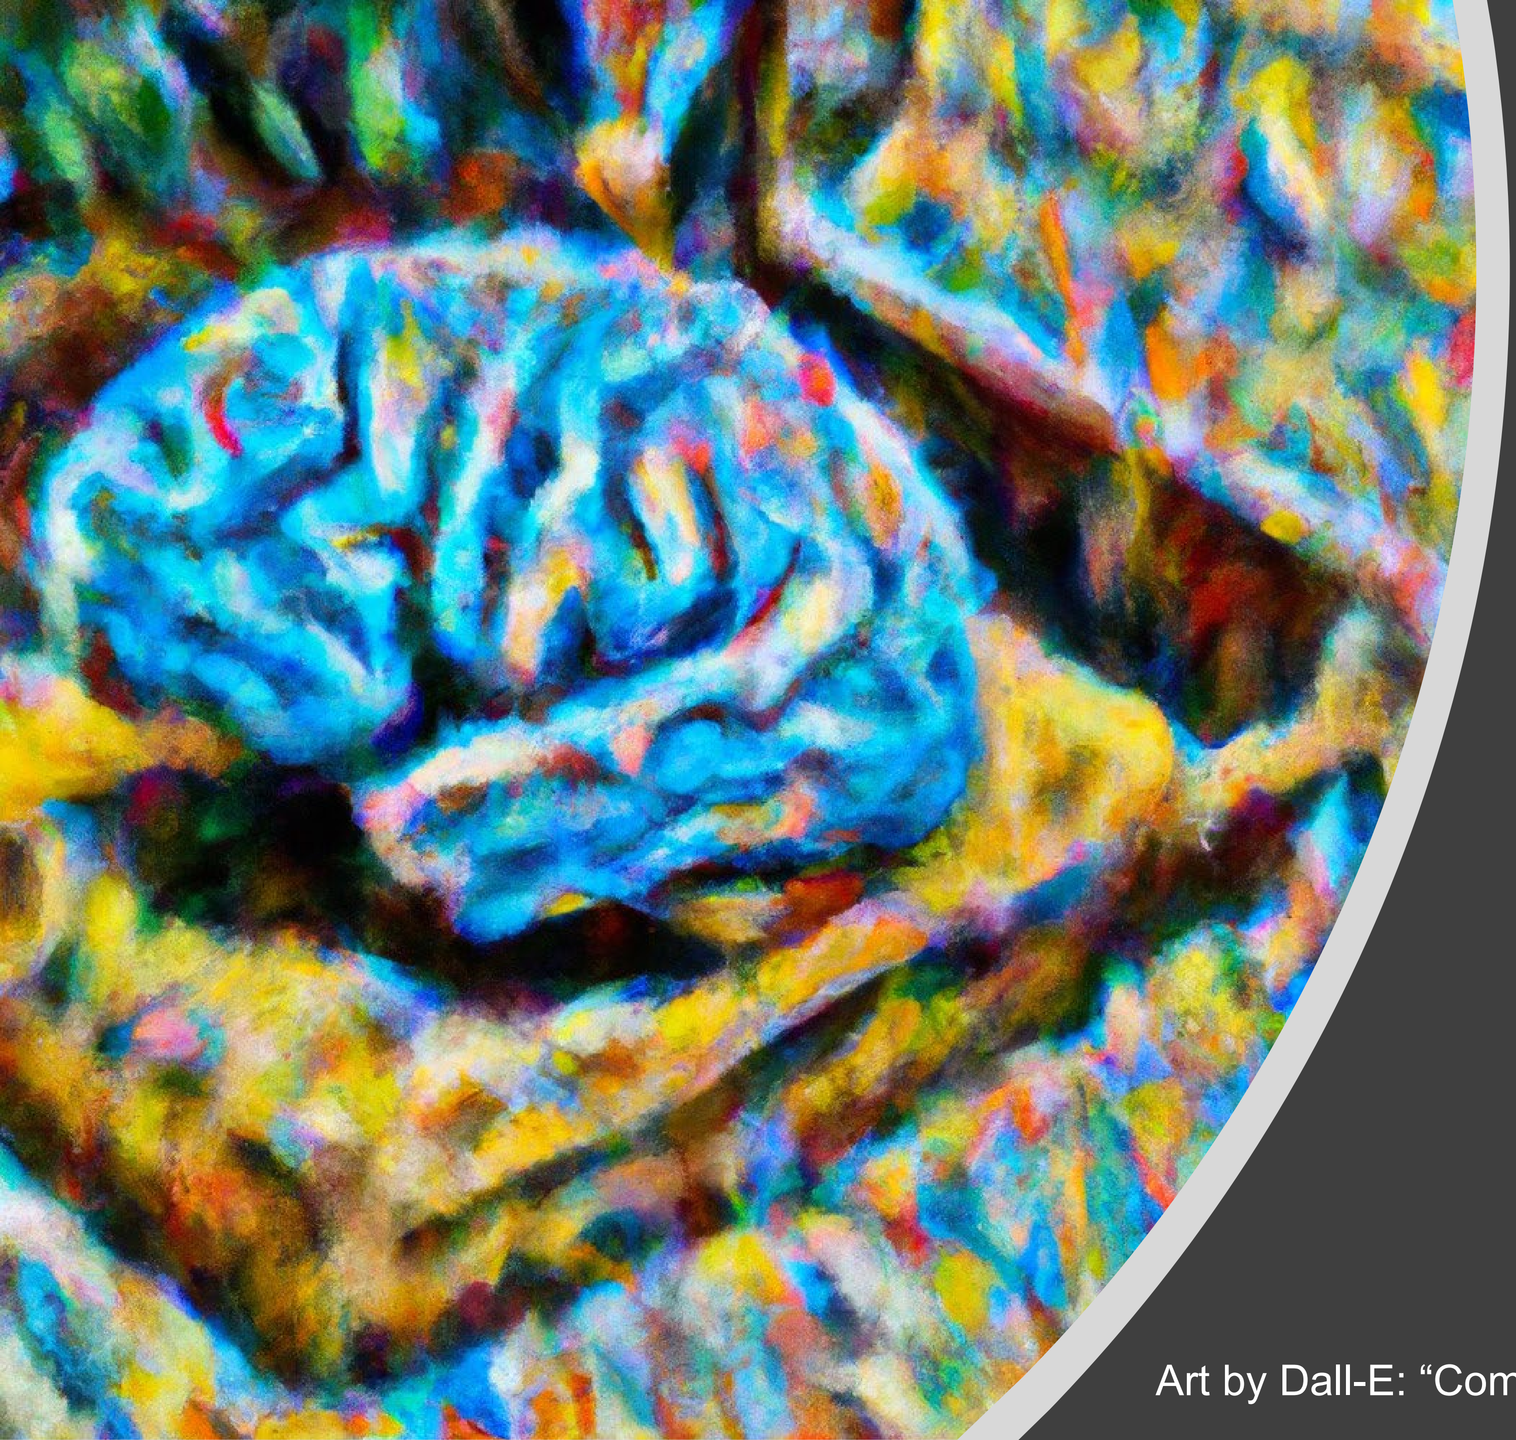

CS441 Applied Machine Learning

Instructor: Derek Hoiem

Art by Dall-E: “Computer brain gathering knowledge, impressionist”

Today's Class

- A little about me
- Intro to Applied Machine Learning
- Course outline and logistics

About me

Raised in “upstate” NY

A little about me

Undergrad at SUNY Buffalo (EE, CE)

PhD in Robotics at Carnegie Mellon

Professor in Computer Science at University of Illinois

Reconstruct: vision for construction

Co-founder and Chief Science Officer of Reconstruct

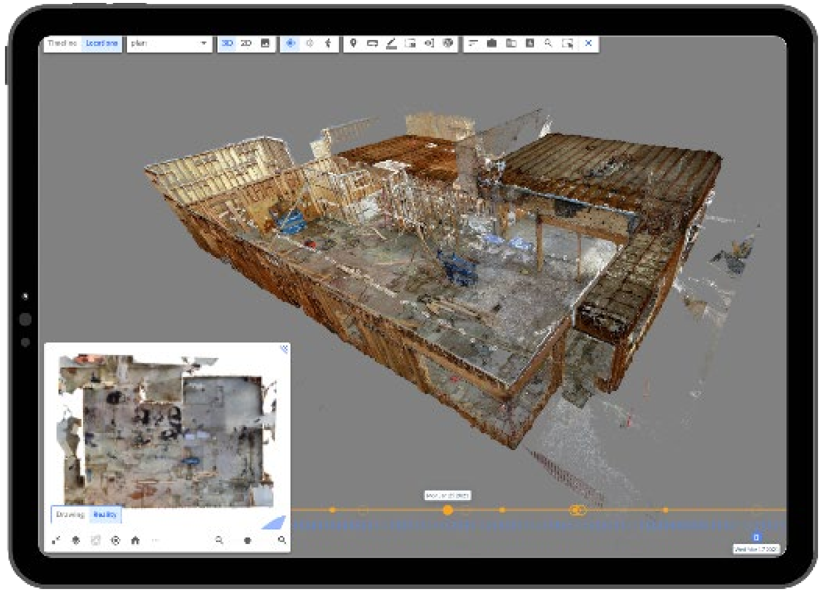

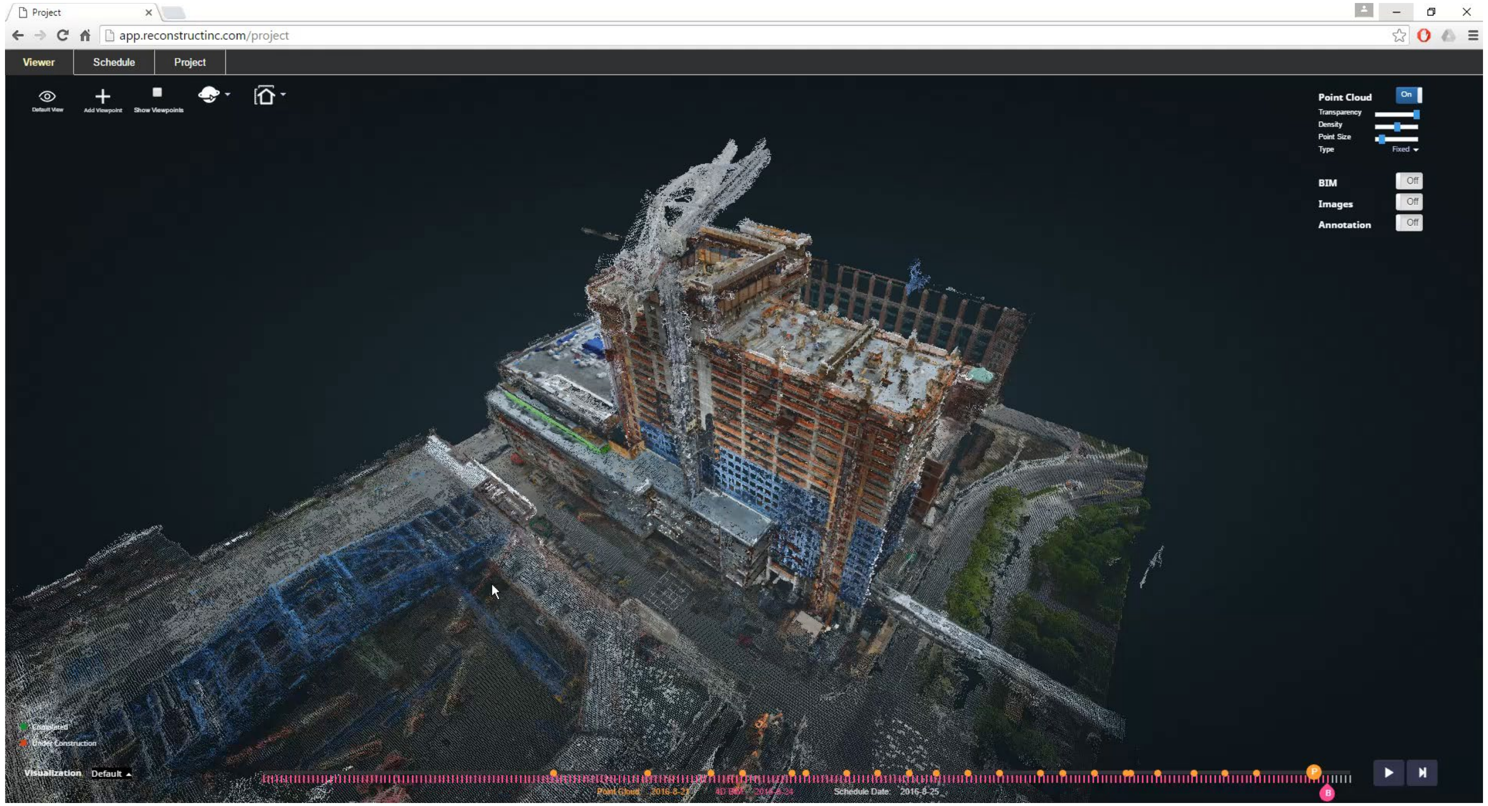

https://vimeo.com/242479887

https://www.reconstructinc.com/

I ❤ ML

My early research: Learning to interpret geometry

Hoiem et al. 2005

MonoPatchNeRF: use deep networks to model 3D scenes

Wu Lee Zou Wang Hoiem. 2025

General Purpose Learners

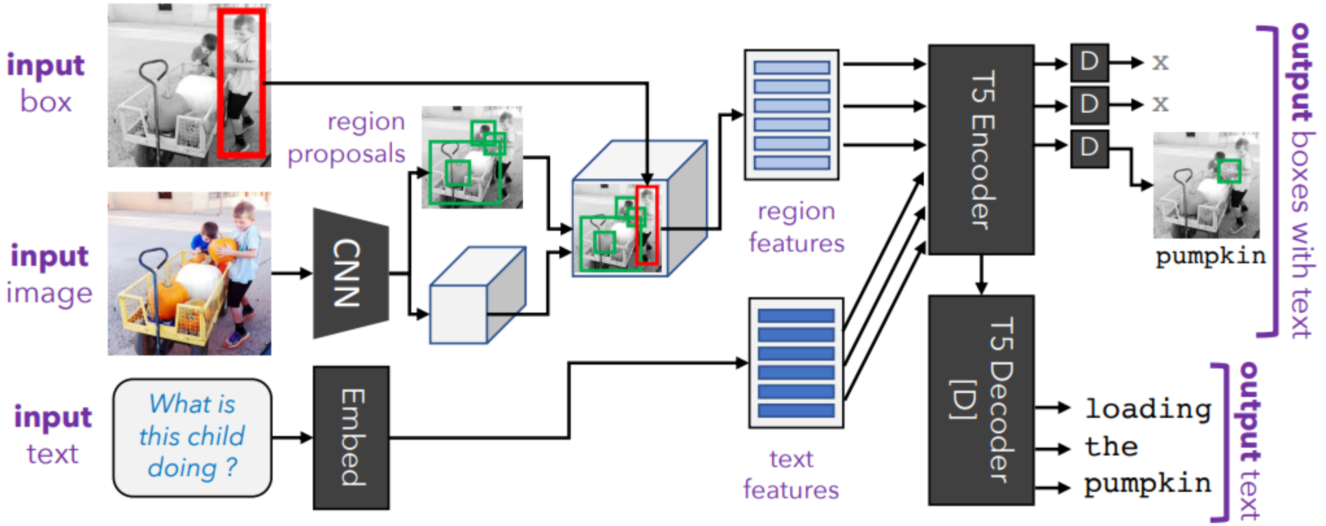

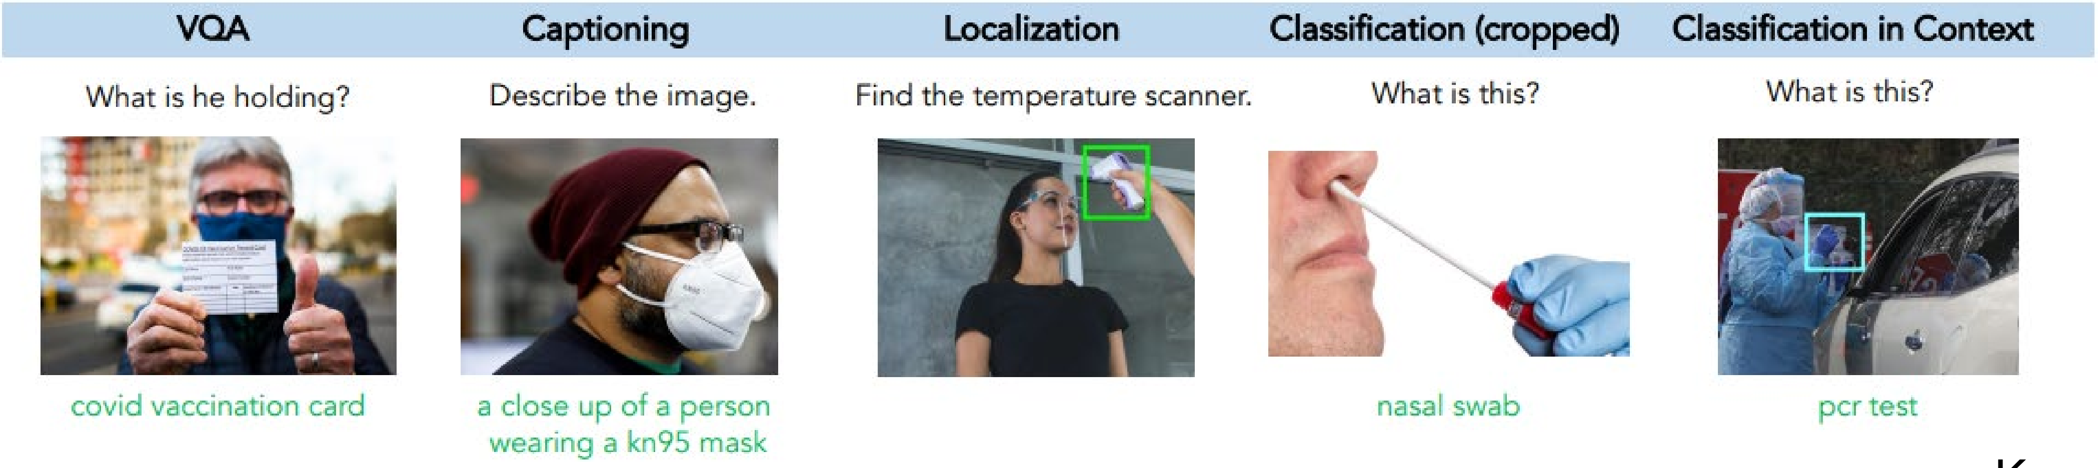

Kamath et al. 2022

Other examples of my research that use machine learning

- Vision
  - Object detection
  - Image classification
  - Photo album organization
  - Image retrieval
  - Describing objects
  - 3D scene modeling
  - 3D object modeling
  - Robot navigation
  - Shadow detection and removal
  - Generating animations

- Vision and Language
  - Visual question answering
  - Phrase grounding
  - Video analysis
  - Large multimodal models

- Audio
  - Sound detection
  - Music identification

Machine learning and the quest for intelligent machines

Rabbi Loew’s Golem

By Mikoláš Aleš https://commons.wikimedia.org/w/index.php?curid=2210897

Vaucanson’s Automata (1738)

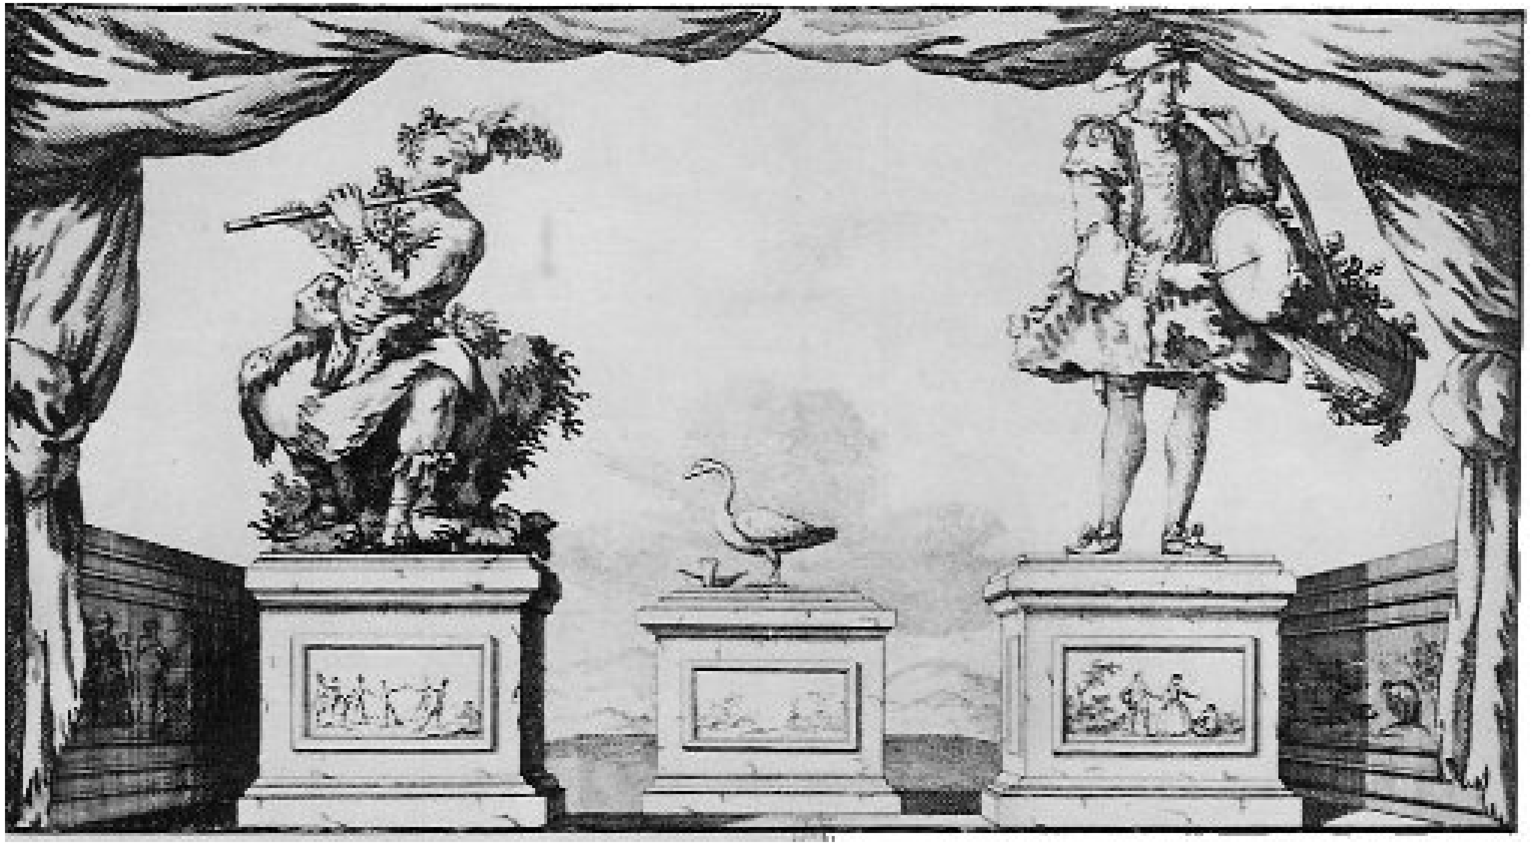

Flute player Defecating Duck Tambourine Player

http://www.francoisjunod.com/automates/eightennth/vaucanson_uk.htm

https://time.graphics/event/1736967

Mechanical Turk (1770)

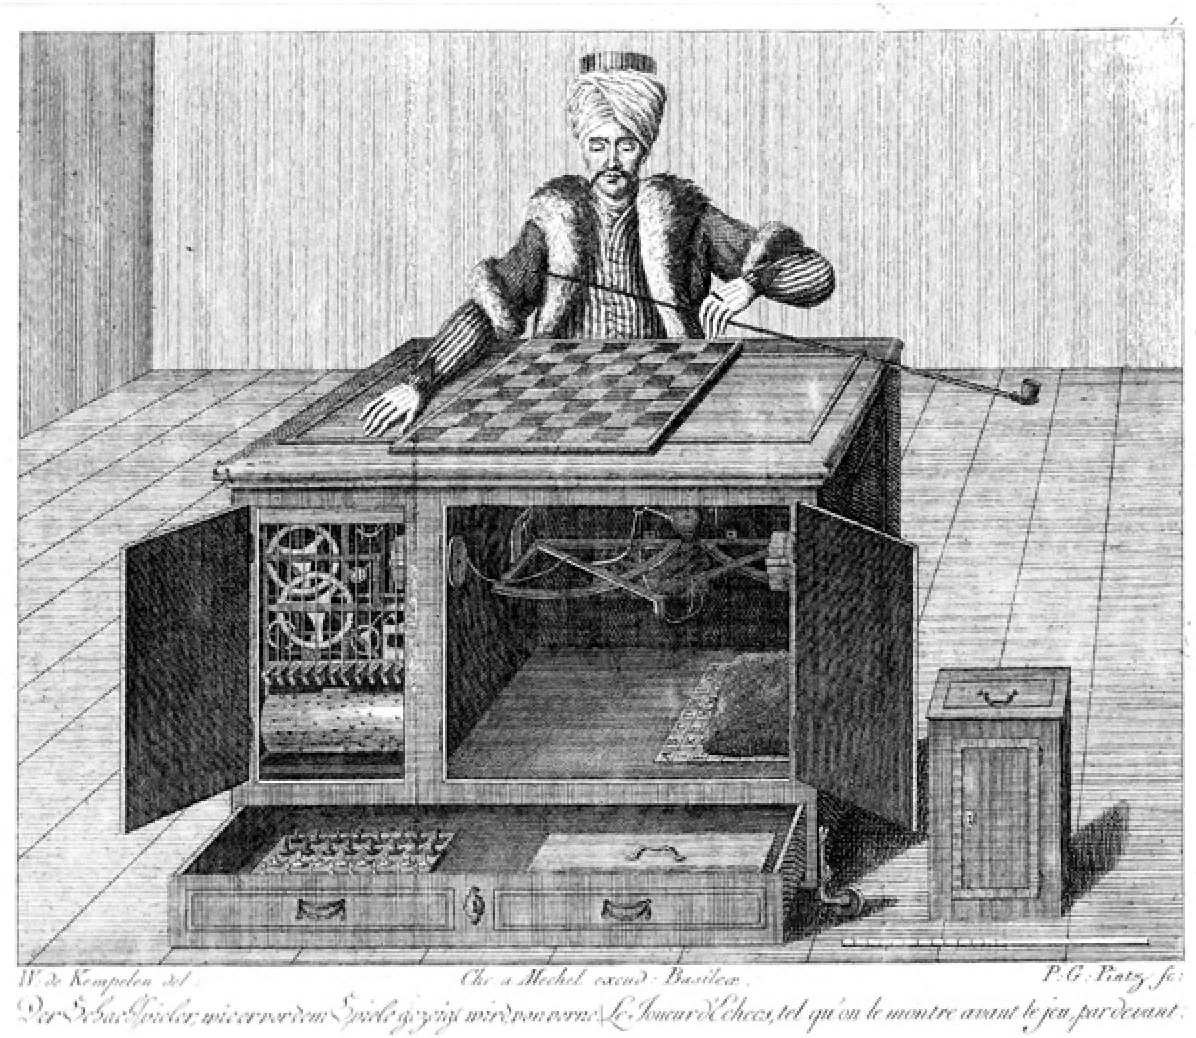

Engraving of Mechanical Turk ([src](src))

Secretly operated by a chess master, the mechanical turk was exhibited as an automatic chess playing machine.

Opponents included Benjamin Franklin and Napoleon Bonaparte.

Babbage's analytical engine (1837)

Not ever completely built

Ada Lovelace described a way to calculate Bernoulli numbers using the machine in 1843 (first computer program)

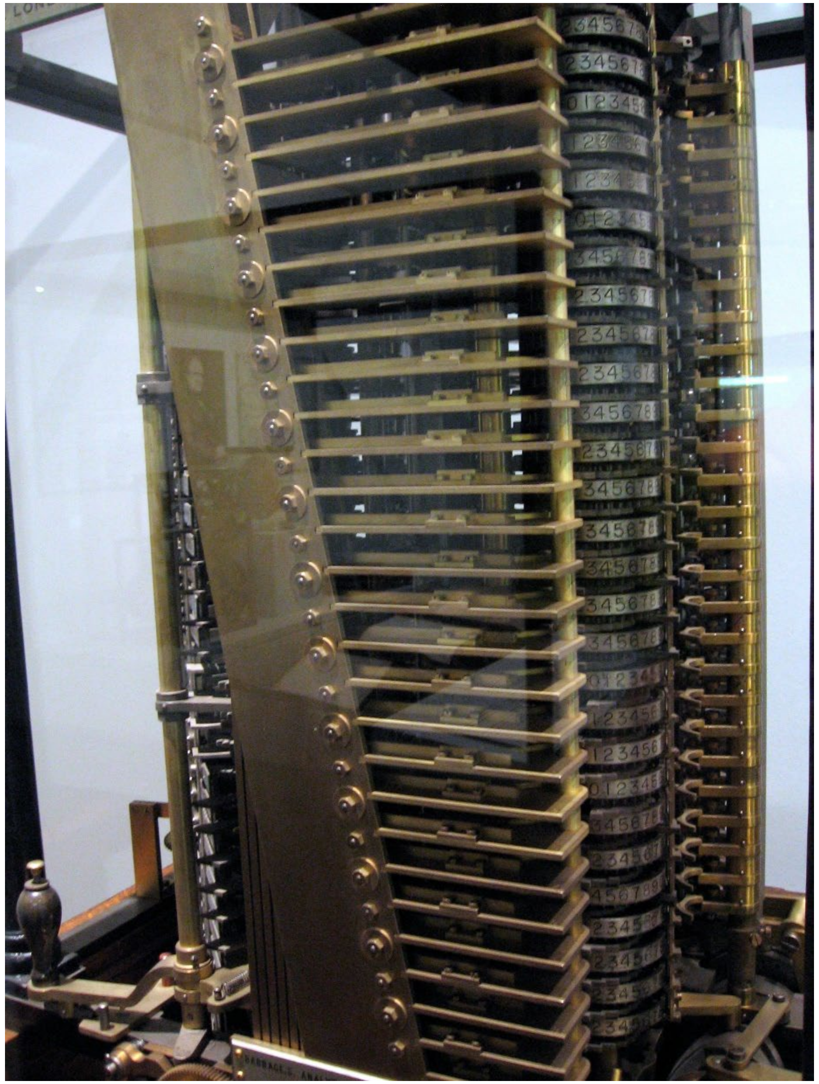

Analytical Engine Mill, built 1910 ([src](src))

Two truths and a lie

1. AGI was predicted to be just around the corner in 1957

2. The first self-driving car was built in 2005

3. Today, AI Chatbots have 1B+ weekly users

Some AI milestones

Early days (1950-1980)

- Alan Turing proposes the Imitation Game (1949)
- The Logic Theorist theorem prover (1955) – Newell and Simon
- The Perceptron - Rosenblatt (1957) (expected to lead to AGI)
- Eliza Chat Bot (1966) – Weizenbaum

A few wins (1980-2010)

- USPS uses OCR for mail processing (1982)
- First commercially successful AI – Expert Systems (1980s)
- Navlab5 drives 2,848 miles across the country averaging 64 mph, 98% autonomously (1995)
- "Deep Blue" defeats Kasparov in Chess (1997)
- Roomba (2002)
- Face detection in digital cameras (2005)
- Google Translate (2006)

Rise of deep learning and generative models (2012-now)

- Convolutional network kicks butt in the ImageNet Challenge (2012)
- AlphaGo defeats Lee Sedol in Go (2016)
- ChatGPT released to public (2022)
- AI ChatBots have 1 billion weekly active users (2025)

Two truths and a lie

1. AGI was predicted to be just around the corner in 1957

2. ~~The first self-driving car was built in 2005~~

3. Today, AI Chatbots have 1B weekly users

Machine Learning

ML spins raw data into gold!

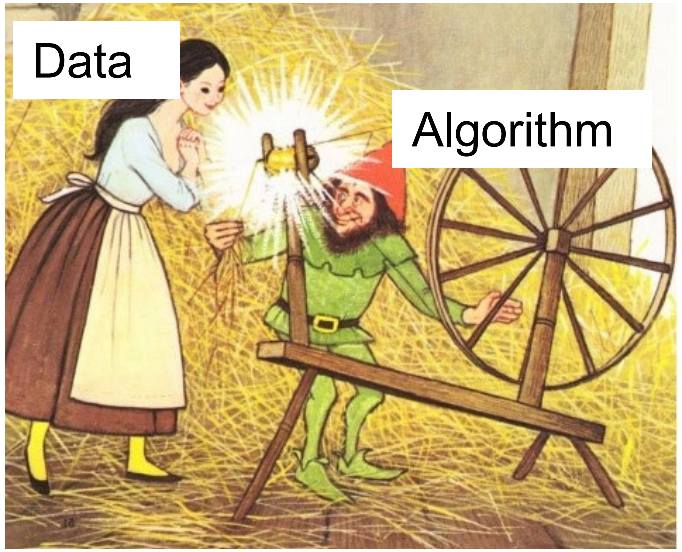

Data
Collect, Annotate, Curate, Format

Train / Evaluate

Design
Model design
Loss, Training Algorithm
Evaluation design

Model

Process of ML

1. Model design
   a. How do we encode information?
   b. How do we predict (decode)?
2. Optimize encoder/decoder models using data

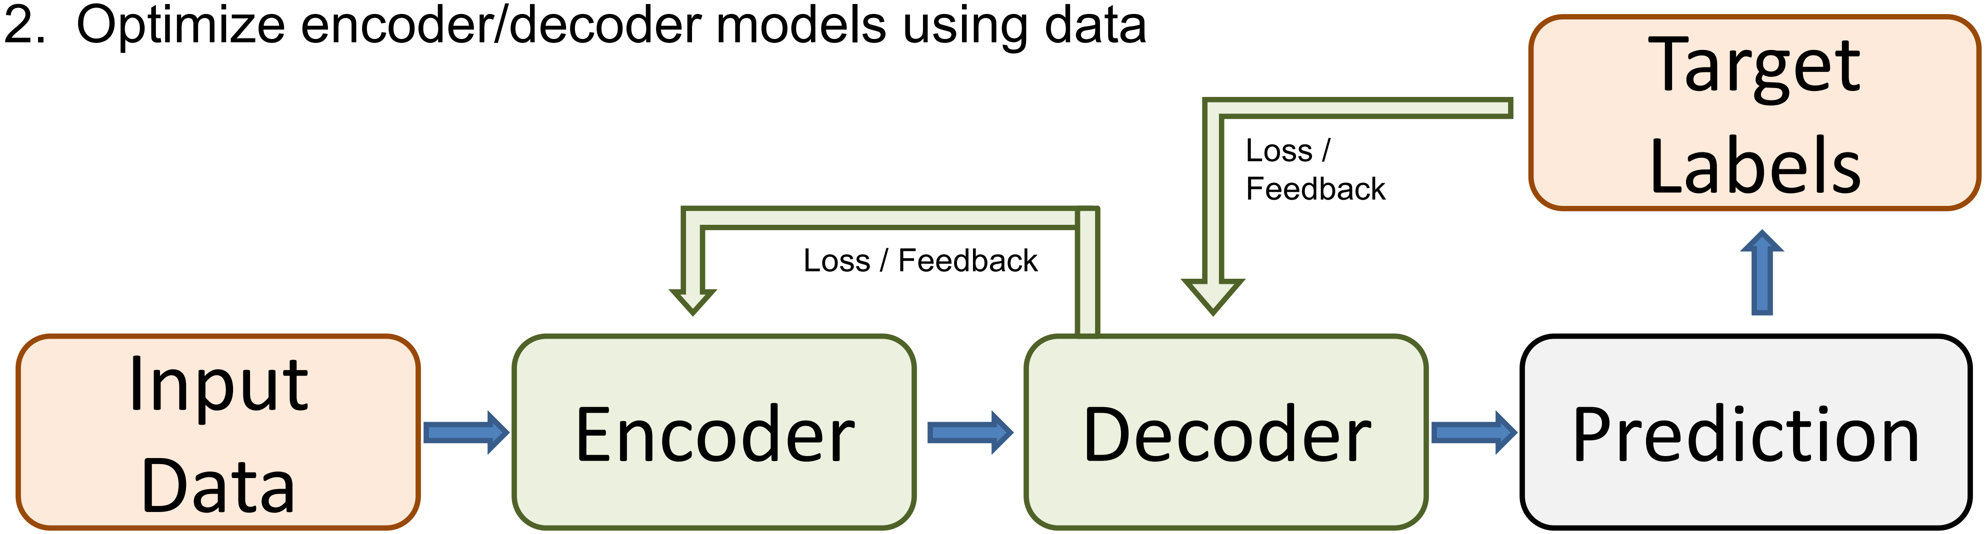

Text
Images
Audio
Discrete/continuous values
Structured/unstructured
Few/many examples
Few/many features
Clean/noisy labels

Manual feature design
Deep networks
Trees
Clustering
Feature selection
Kernels
Density estimation

Linear regressor
SVM
Logistic regressor
Nearest Neighbor
Probabilistic model

Clusters
Low dimensional embedding
Category
Pixel labels
Generated text, image, audio
Positions
Yes/No

Traditional ML: Focus on the Model/Algorithm Design

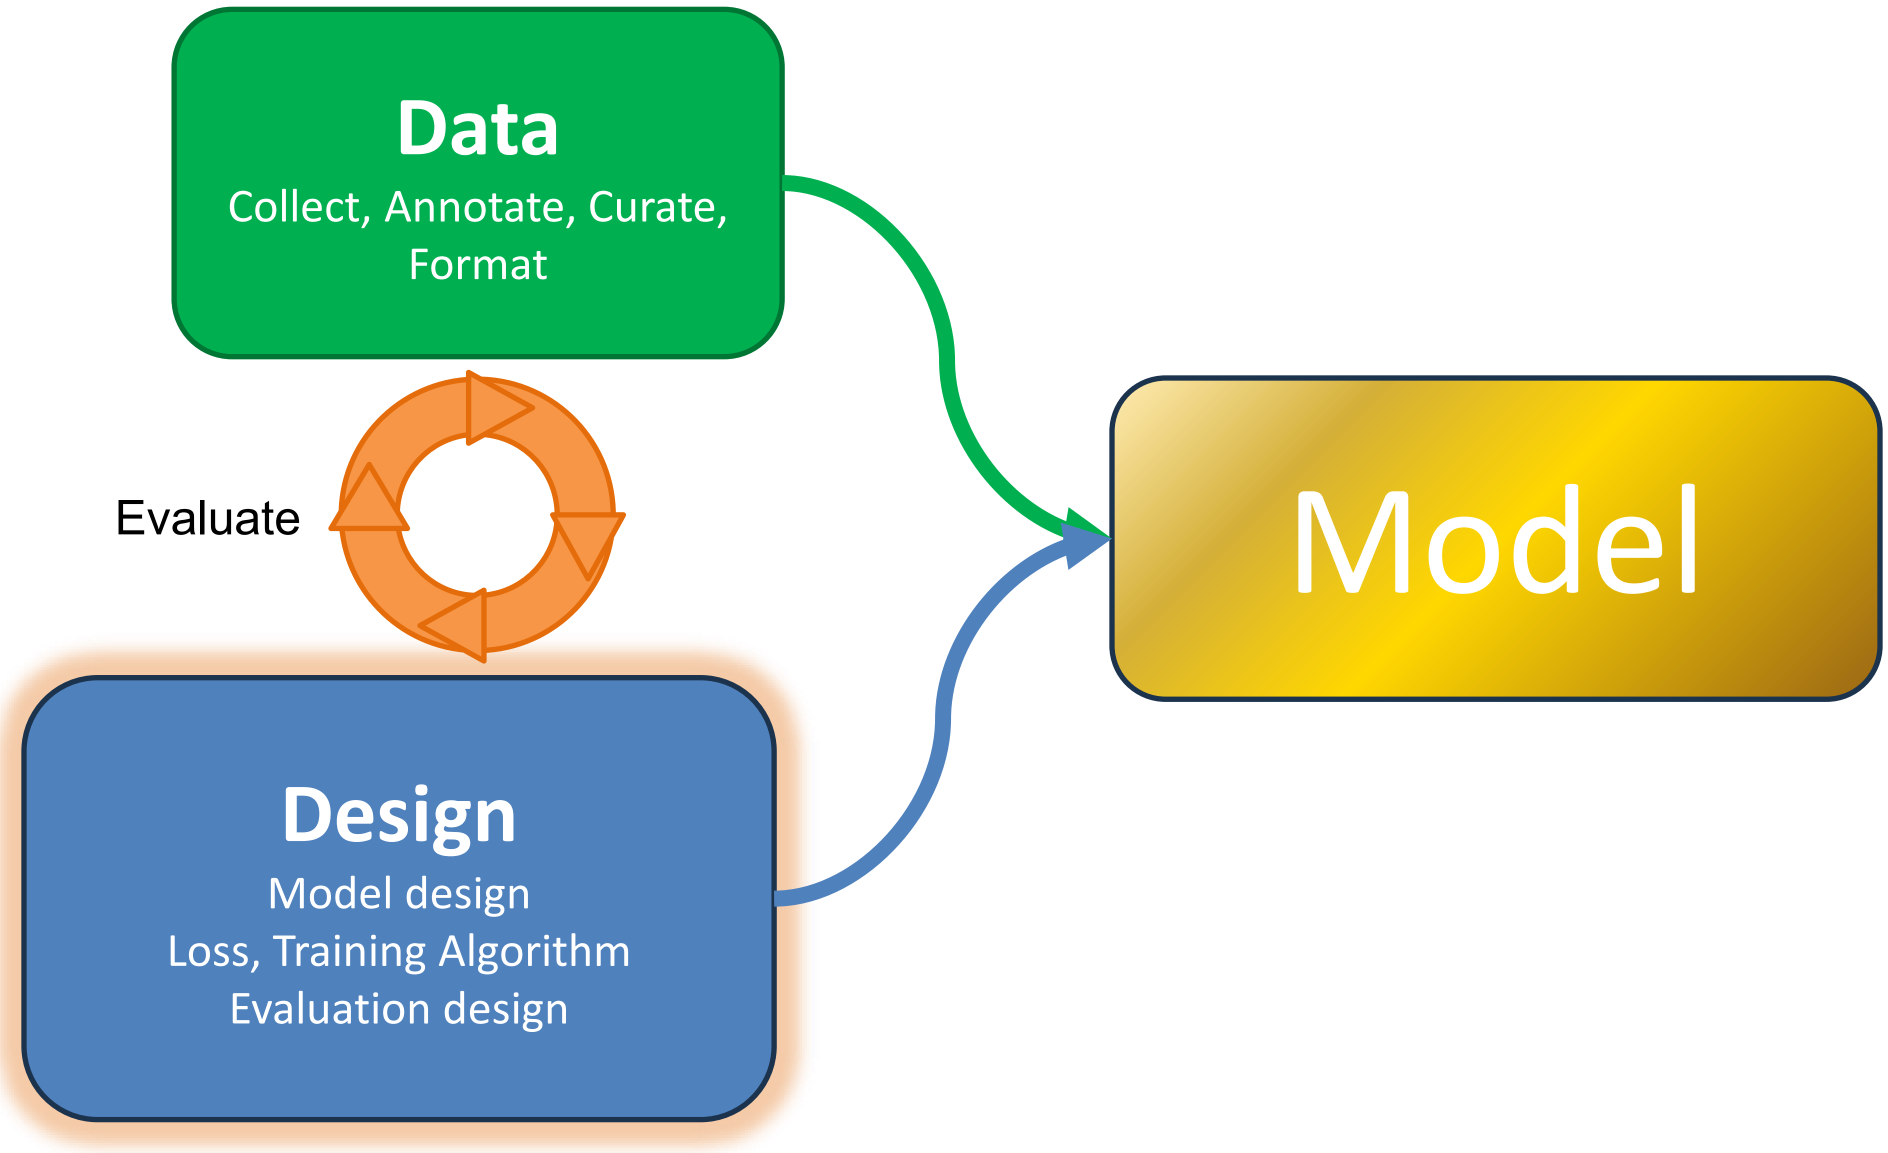

Deep ML (vision/text/audio): Focus on the Data

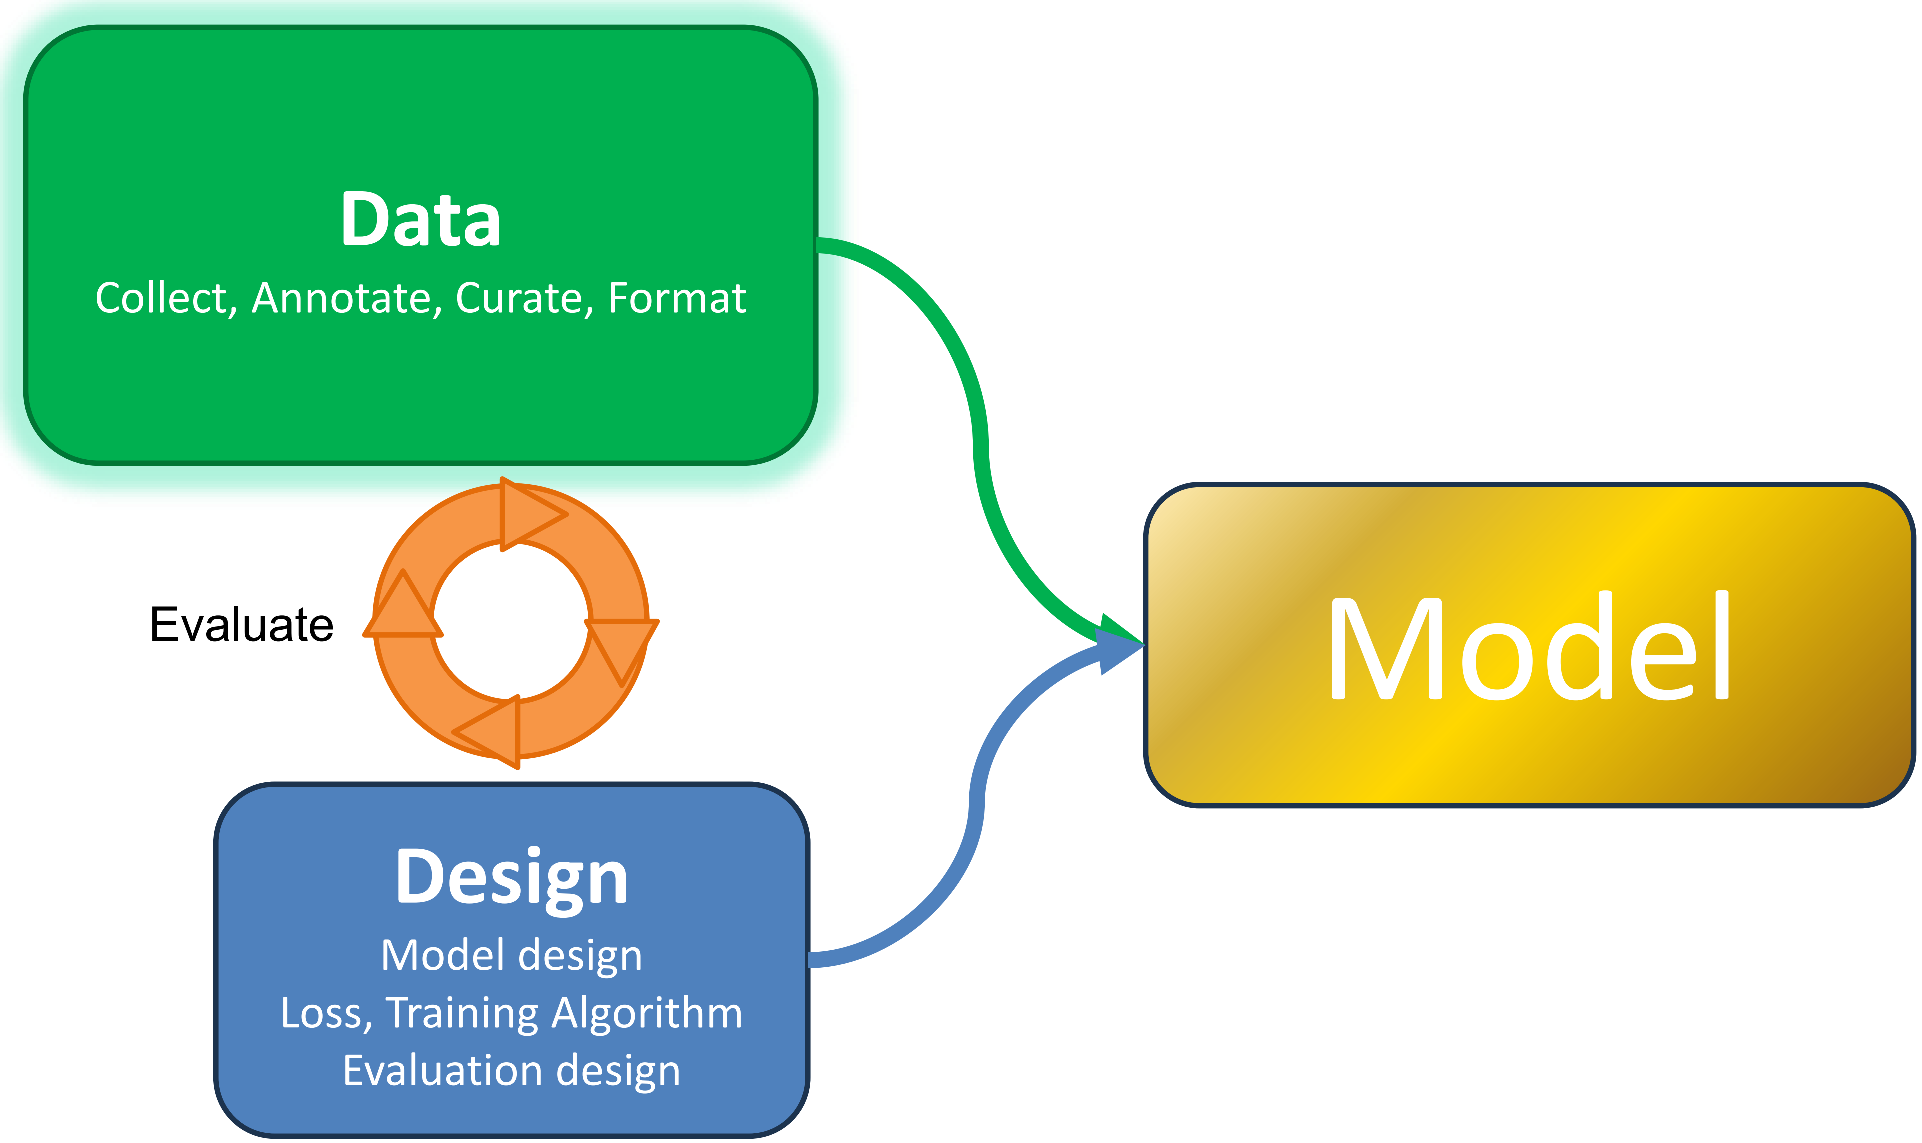

Example: Alexa Speech Recognition

TL;DR The advances in Alexa's speech recognition came from streamlining data acquisition, not algorithms.

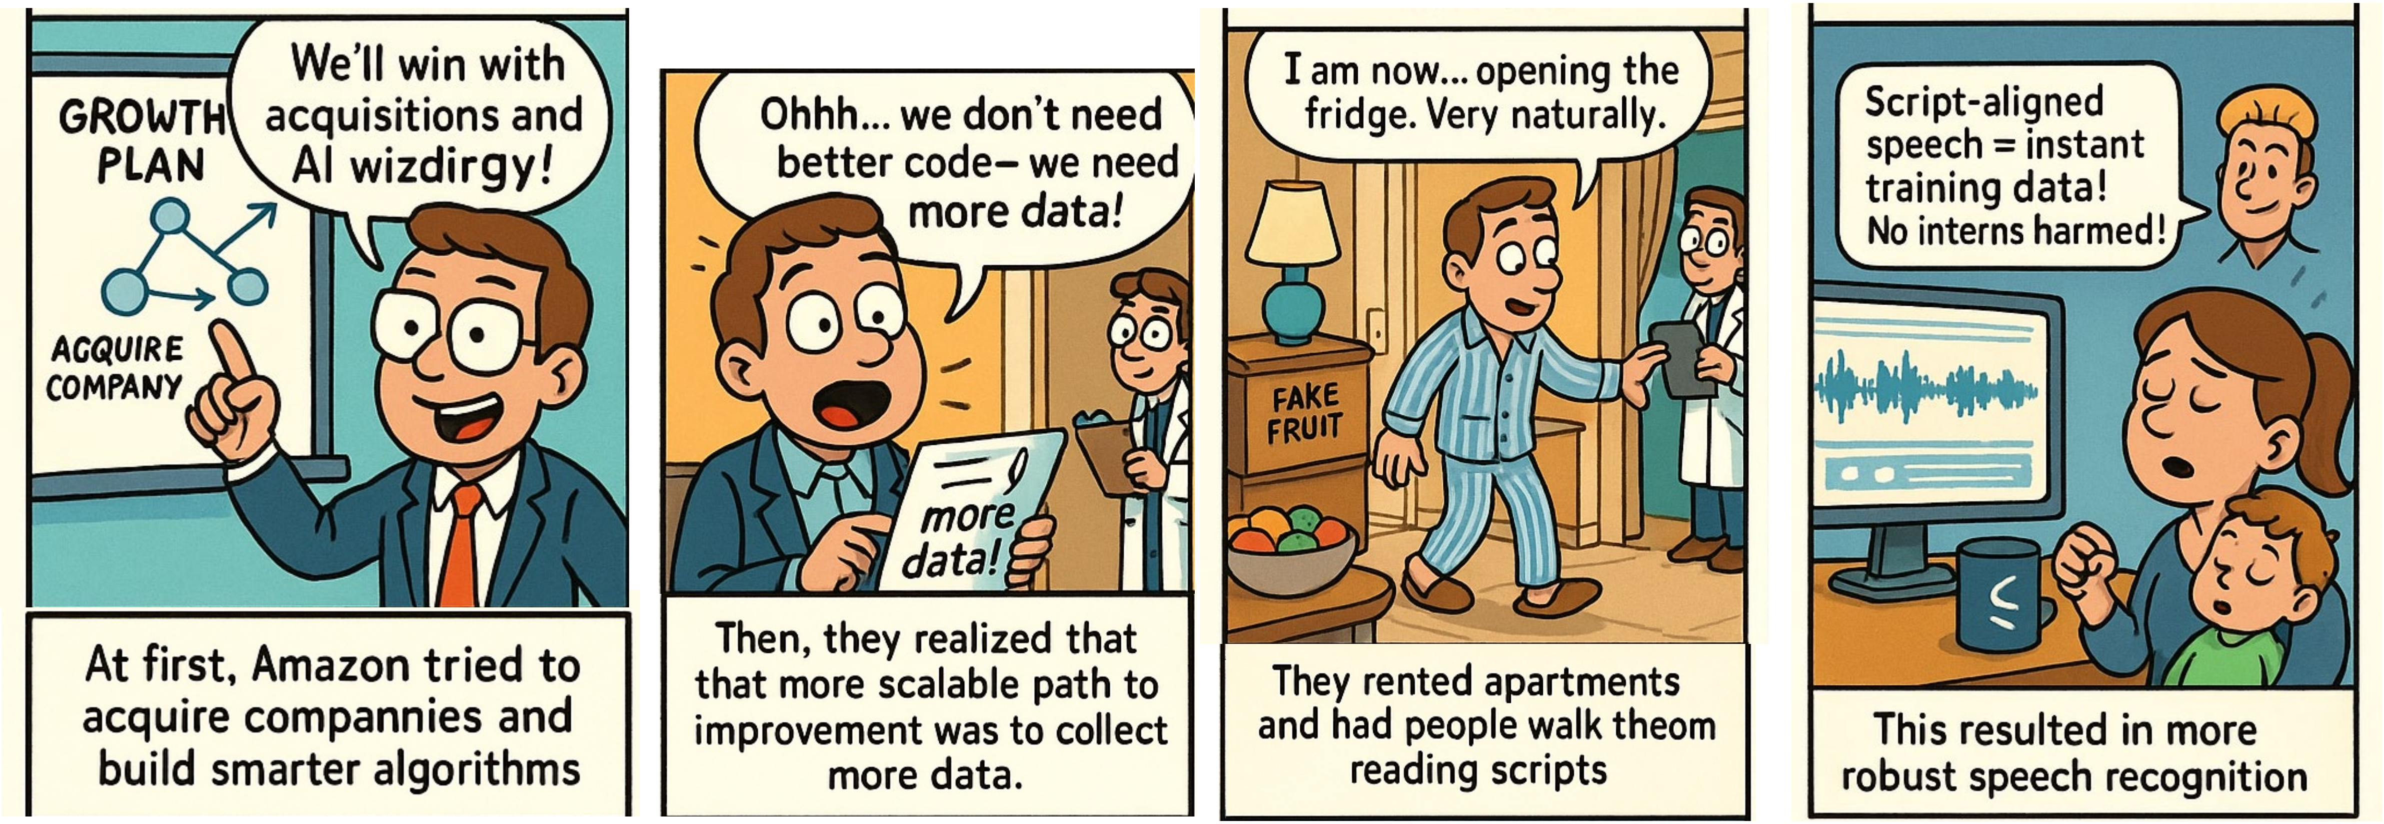

(Comic created by ChatGPT :/)

Why learn ML?

Appreciate your own intelligence

Our identities are largely characterized by what we learn.

Making machines that learn can help understand our own learning.

Learn to use ML to solve problems

- Key concepts and methodologies for learning from data
- Algorithms and their strengths and limitations
- Domain-specific representations
- Ability to select the right tools for the job

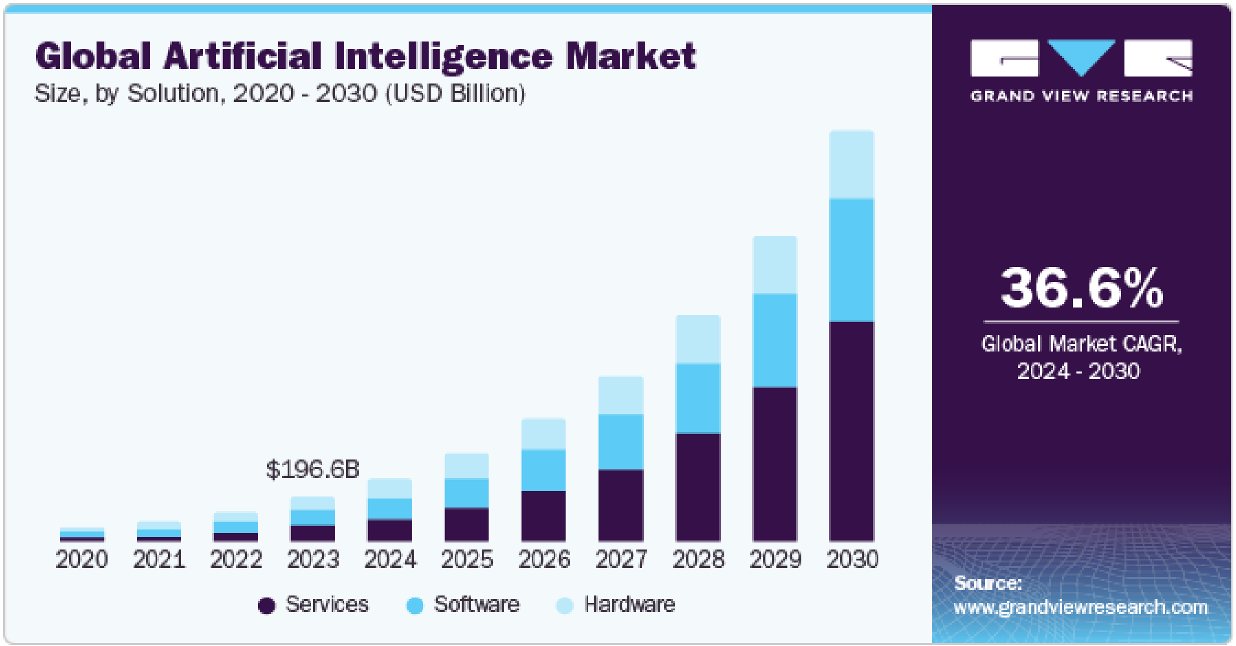

Develop responsibly by understanding societal and individual impact

- Recommending systems
- Surveillance
- Robots
- Smart assistants
- Text generation
- Autonomous cars
- Social media bots

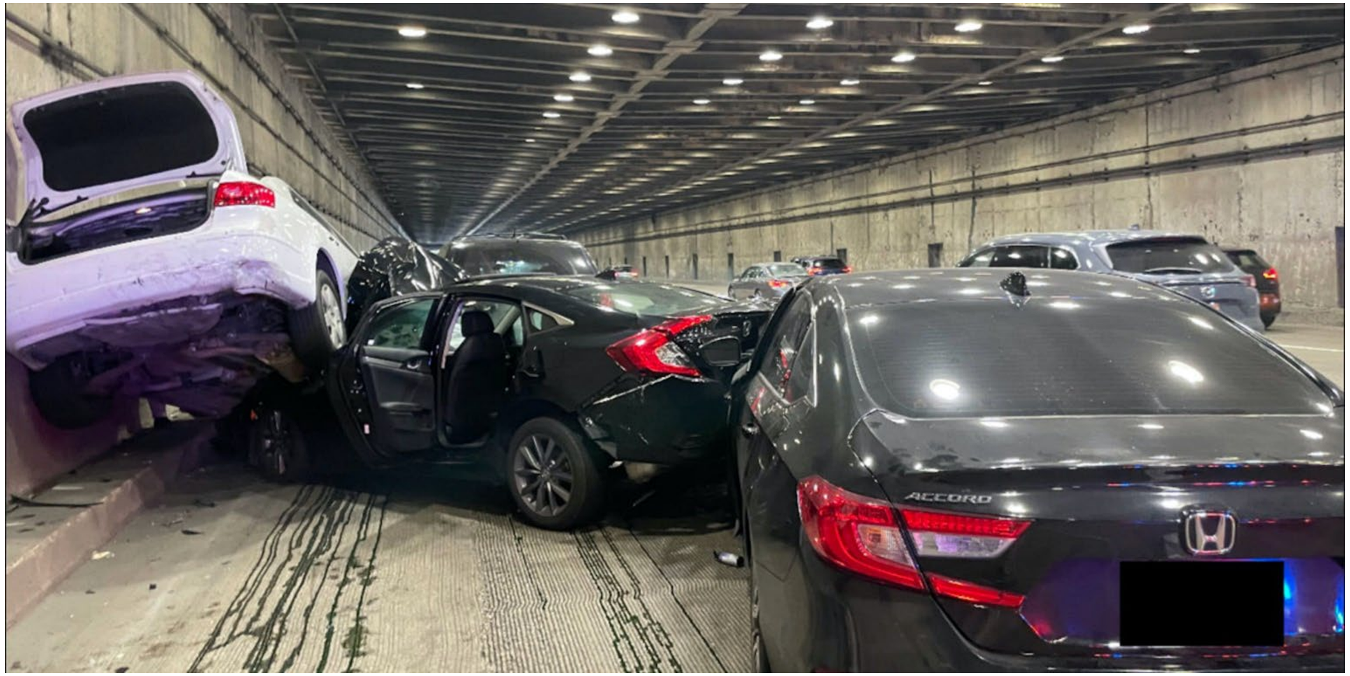

Tesla accident

Course outline

Prof: Derek Hoiem dhoiem@illinois.edu

TAs
Deema Alnuhait deemaa2@illinois.edu
Prachi Garg prachig3@illinois.edu
Nguyen Khoi Hoang nkhoang2@illinois.edu
Savya Khosla savyak2@illinois.edu
Jan Vijay Singh jvsingh2@illinois.edu
Yihang Sun yihangs@illinois.edu
Sethuraman (Sethu) Trichy st34@illinois.edu
Yuqun Wu yuqunwu2@illinois.edu
Yao Xiao yaox11@illinois.edu

Website: https://courses.engr.illinois.edu/cs441/fa2026/

Bookmark this page!!

How the class works

Lectures & Reading

- Attend lecture
- Reading materials and notes are optional
- Recordings available (if they work), slides available

Date	Topic	Link	Reading/Notes
Aug 27 (Tues)	Introduction	ppt ; pdf	Jupyter notebook tutorial vid ipynb cc Numpy tutorial vid cc Linear algebra tutorial vid cc
	Fundamentals of Learning		
Aug 29 (Thurs)	K-NN Classification, Data Representation	ppt ; pdf	AML Ch 1.1-1.2
Sep 3 (Tues)	K-NN Regression, Generalization	ppt ; pdf	AML Ch 1.1-1.2

Exams

- Three 50-minute cumulative exams to test conceptual knowledge
- Exams held in CBTF (closed book/notes)

Application

- 5 HWs
  - Conduct experiments in Jupyter notebooks
  - Report results
  - Answer questions to test your understanding
- Final project
  - Define your own problem
  - Collect data
  - Train and evaluate model(s)

Getting help

- Ask quick questions on CampusWire
  - We aim to answer within 24 hours
- For more in-depth help/understanding, come to office hours
- Schedules for office hours and CW answering will be posted

Topics

- Fundamentals of learning
  - How to build classifiers and regressors based on provided features
  - Working with data, instance-based methods, linear models, probabilistic methods, trees
- Deep representation learning
  - How to learn effective representations
  - Optimization, MLPs, CNNs, transformers, vision, language, foundational models
- Applications
  - Ethics and impact, bias/fairness, building applications, RL, audio and time series

Date	Topic
Aug 25 (Tues)	Introduction
	Fundamentals of Learning
Aug 27 (Thurs)	K-NN Classification, Data Representation
	Probability/Background Review (ask to review in office hours for individual help)
Sep 1 (Tues)	K-NN Regression, Generalization
Sep 3 (Thurs)	Search and Clustering
Sep 8 (Tues)	Dimensionality reduction: PCA, embeddings *No in-person lecture, recording only*
Sep 10 (Thurs)	Linear regression, regularization *No in-person lecture, recording only*
Sep 14 (Mon)	*HW 1 (Instance-based Methods) due*
Sep 15 (Tues)	Linear classifiers: logistic regression, SVM
Sep 17 (Thurs)	Naïve Bayes Classifier
Sep 22 (Tues)	EM and Latent Variables
Sep 24 (Thurs)	Density estimation: MoG, Hists, KDE
Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI
Nov 12 (Thurs)	Bias in AI, Fair ML
Nov 16 (Mon)	*HW 5 (Deep Learning and Applications) due*
Nov 17 (Tues)	Building and Deploying ML *Guest speaker: Chenxi Yu (State Farm)*
Nov 19 (Thurs)	Audio and 1D Signals
Nov 21-Nov 29	Fall Break
Dec 1 (Tues)	Reinforcement Learning
Dec 3 (Thurs)	Review, summary, looking forward
Dec 8 (Tues)	*No Class – Office hours for Final Project in Lecture Hall (CIF 0027) 12:30-1:30*
Dec 3-8	*Exam 3 at CBTF*
Dec 13 (Sun)	*Final Project due (cannot be late)*

Weeks 1-3

Instance-based Methods

How do we represent data --- color, temperature, pixel intensity, name, etc. --- with numbers?

Given a new sample, how do we find similar examples in our dataset and use it for prediction?

How do we represent similar information with fewer numbers and visualize it?

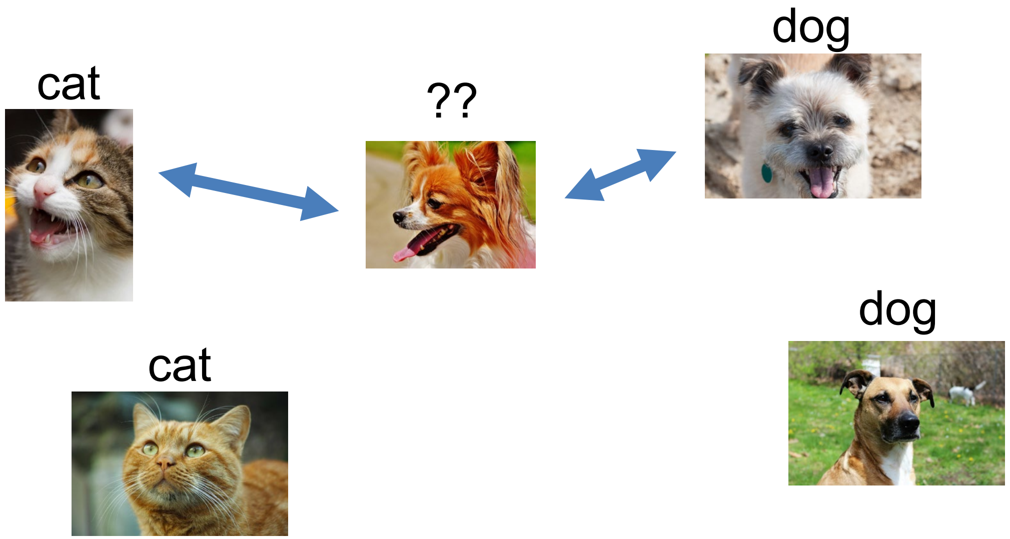

Date	Topic
Aug 25 (Tues)	Introduction
	Fundamentals of Learning
Aug 27 (Thurs)	K-NN Classification, Data Representation
	Probability/Background Review (ask to review in office hours for individual help)
Sep 1 (Tues)	K-NN Regression, Generalization
Sep 3 (Thurs)	Search and Clustering
Sep 8 (Tues)	Dimensionality reduction: PCA, embeddings *No in-person lecture, recording only*
Sep 10 (Thurs)	Linear regression, regularization *No in-person lecture, recording only*
Sep 14 (Mon)	*HW 1 (Instance-based Methods) due*
Sep 15 (Tues)	Linear classifiers: logistic regression, SVM
Sep 17 (Thurs)	Naïve Bayes Classifier
Sep 22 (Tues)	EM and Latent Variables
Sep 24 (Thurs)	Density estimation: MoG, Hists, KDE
Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI

Weeks 3-4

How do we learn a linear model from data, using a weighted combination of features to produce a score or a predicted value?

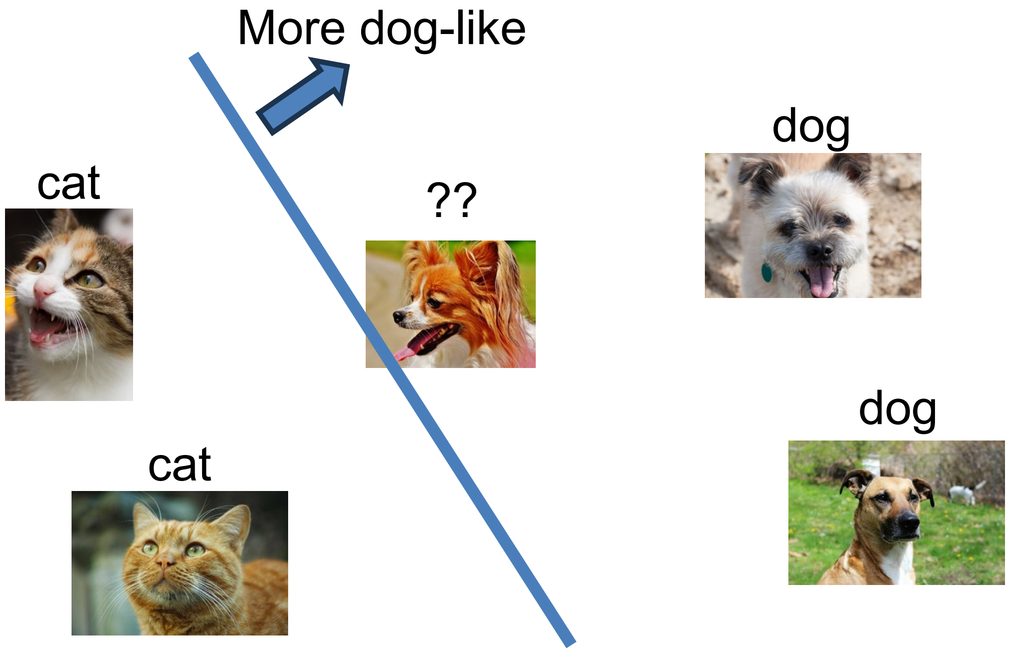

Date	Topic
Aug 25 (Tues)	Introduction
	Fundamentals of Learning
Aug 27 (Thurs)	K-NN Classification, Data Representation
	Probability/Background Review (ask to review in office hours for individual help)
Sep 1 (Tues)	K-NN Regression, Generalization
Sep 3 (Thurs)	Search and Clustering
Sep 8 (Tues)	Dimensionality reduction: PCA, embeddings *No in-person lecture, recording only*
Sep 10 (Thurs)	Linear regression, regularization *No in-person lecture, recording only*
Sep 14 (Mon)	*HW 1 (Instance-based Methods) due*
Sep 15 (Tues)	Linear classifiers: logistic regression, SVM
Sep 17 (Thurs)	Naïve Bayes Classifier
Sep 22 (Tues)	EM and Latent Variables
Sep 24 (Thurs)	Density estimation: MoG, Hists, KDE
Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI

Weeks 4-6

How do we estimate the likelihood of a data sample and use it for prediction, grouping, or to find outliers?

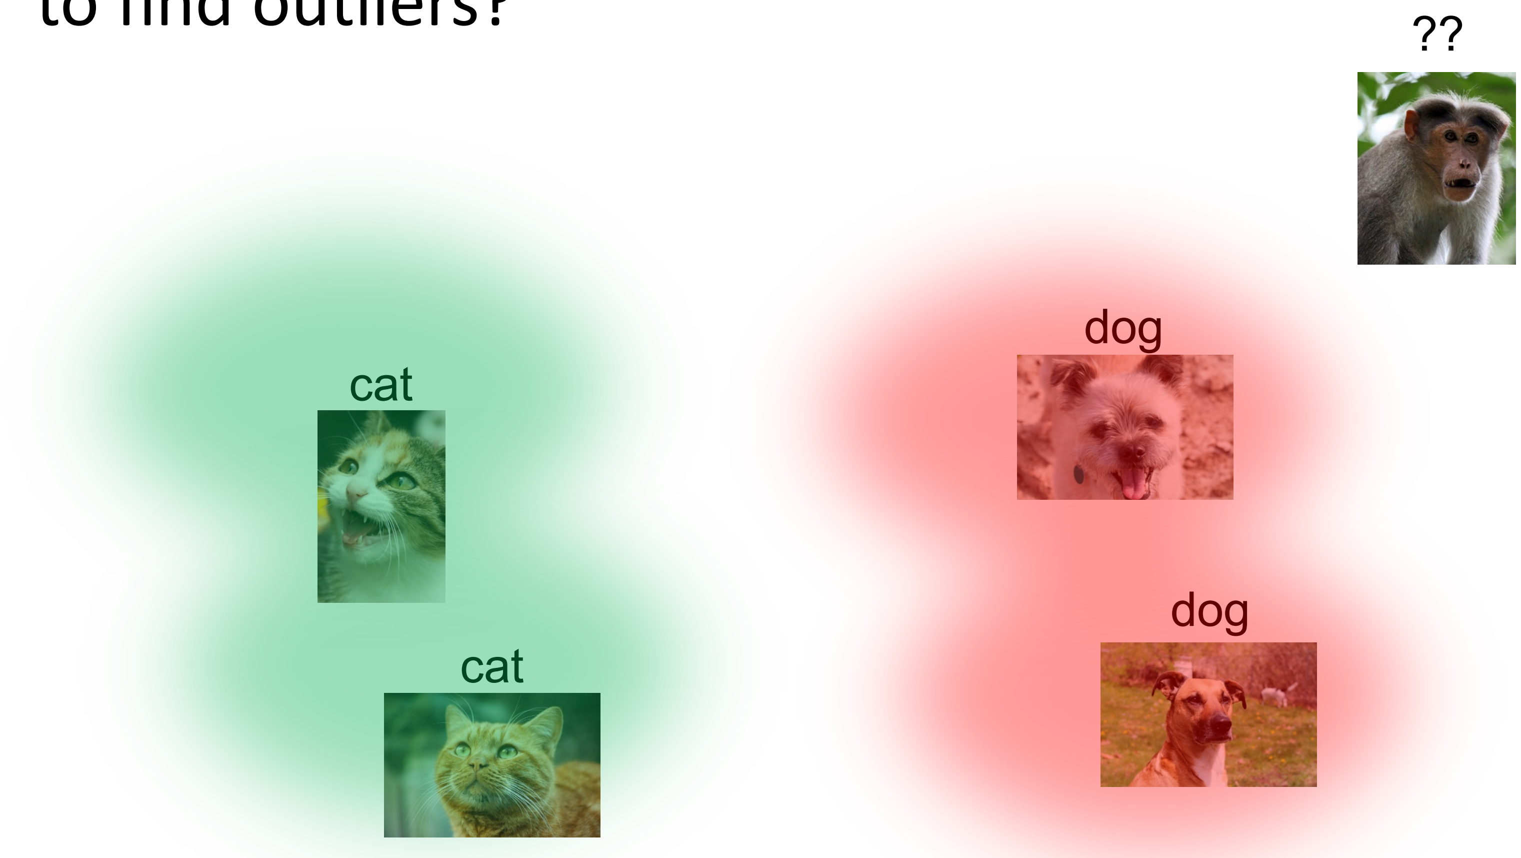

Date	Topic
Aug 25 (Tues)	Introduction
	Fundamentals of Learning
Aug 27 (Thurs)	K-NN Classification, Data Representation
	Probability/Background Review (ask to review in office hours for individual help)
Sep 1 (Tues)	K-NN Regression, Generalization
Sep 3 (Thurs)	Search and Clustering
Sep 8 (Tues)	Dimensionality reduction: PCA, embeddings *No in-person lecture, recording only*
Sep 10 (Thurs)	Linear regression, regularization *No in-person lecture, recording only*
Sep 14 (Mon)	*HW 1 (Instance-based Methods) due*
Sep 15 (Tues)	Linear classifiers: logistic regression, SVM
Sep 17 (Thurs)	Naïve Bayes Classifier
Sep 22 (Tues)	EM and Latent Variables
Sep 24 (Thurs)	Density estimation: MoG, Hists, KDE
Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI

Week 7

How do we learn rules from data and use them to vote for predictions?

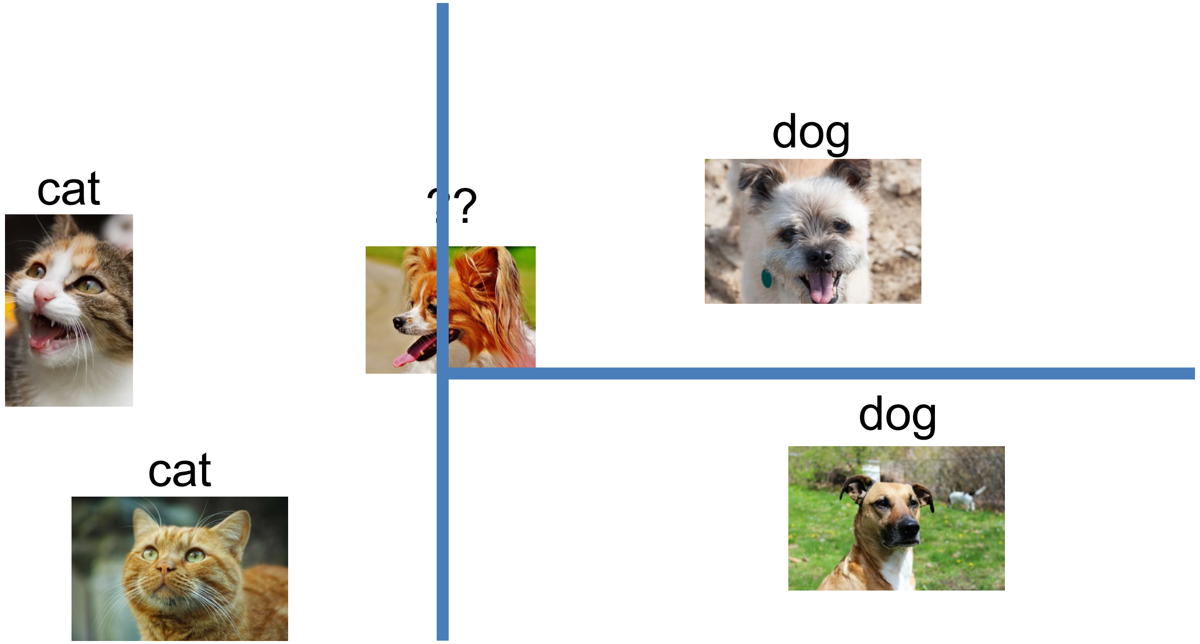

Date	Topic
Aug 25 (Tues)	Introduction
	Fundamentals of Learning
Aug 27 (Thurs)	K-NN Classification, Data Representation
	Probability/Background Review (ask to review in office hours for individual help)
Sep 1 (Tues)	K-NN Regression, Generalization
Sep 3 (Thurs)	Search and Clustering
Sep 8 (Tues)	Dimensionality reduction: PCA, embeddings *No in-person lecture, recording only*
Sep 10 (Thurs)	Linear regression, regularization *No in-person lecture, recording only*
Sep 14 (Mon)	*HW 1 (Instance-based Methods) due*
Sep 15 (Tues)	Linear classifiers: logistic regression, SVM
Sep 17 (Thurs)	Naïve Bayes Classifier
Sep 22 (Tues)	EM and Latent Variables
Sep 24 (Thurs)	Density estimation: MoG, Hists, KDE
Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI

Weeks 8-9

How do we learn new representations of our data using neural networks?

Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI
Nov 12 (Thurs)	Bias in AI, Fair ML
Nov 16 (Mon)	*HW 5 (Deep Learning and Applications) due*
Nov 17 (Tues)	Building and Deploying ML *Guest speaker: Chenxi Yu (State Farm)*
Nov 19 (Thurs)	Audio and 1D Signals
Nov 21-Nov 29	Fall Break
Dec 1 (Tues)	Reinforcement Learning
Dec 3 (Thurs)	Review, summary, looking forward
Dec 8 (Tues)	*No Class – Office hours for Final Project in Lecture Hall (CIF 0027) 12:30-1:30*
Dec 3-8	*Exam 3 at CBTF*
Dec 13 (Sun)	*Final Project due (cannot be late)*

Weeks 8-9

How do we learn new representations of our data using neural networks?

dog

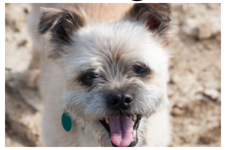

??

cat

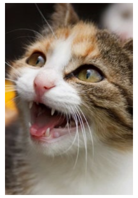

cat

dog

Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI
Nov 12 (Thurs)	Bias in AI, Fair ML
Nov 16 (Mon)	*HW 5 (Deep Learning and Applications) due*
Nov 17 (Tues)	Building and Deploying ML *Guest speaker: Chenxi Yu (State Farm)*
Nov 19 (Thurs)	Audio and 1D Signals
Nov 21-Nov 29	Fall Break
Dec 1 (Tues)	Reinforcement Learning
Dec 3 (Thurs)	Review, summary, looking forward
Dec 8 (Tues)	*No Class – Office hours for Final Project in Lecture Hall (CIF 0027) 12:30-1:30*
Dec 3-8	*Exam 3 at CBTF*
Dec 13 (Sun)	*Final Project due (cannot be late)*

Weeks 9-10

How do we learn to represent images and text?

A small white puppy with a blue tag. The puppy has its mouth open.

Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI
Nov 12 (Thurs)	Bias in AI, Fair ML
Nov 16 (Mon)	*HW 5 (Deep Learning and Applications) due*
Nov 17 (Tues)	Building and Deploying ML *Guest speaker: Chenxi Yu (State Farm)*
Nov 19 (Thurs)	Audio and 1D Signals
Nov 21-Nov 29	Fall Break
Dec 1 (Tues)	Reinforcement Learning
Dec 3 (Thurs)	Review, summary, looking forward
Dec 8 (Tues)	*No Class – Office hours for Final Project in Lecture Hall (CIF 0027) 12:30-1:30*
Dec 3-8	*Exam 3 at CBTF*
Dec 13 (Sun)	*Final Project due (cannot be late)*

Weeks 11-12

Hire (guard dog)

Don't hire

How do we understand the social and personal impact of AI and develop responsibly?

Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI
Nov 12 (Thurs)	Bias in AI, Fair ML
Nov 16 (Mon)	*HW 5 (Deep Learning and Applications) due*
Nov 17 (Tues)	Building and Deploying ML *Guest speaker: Chenxi Yu (State Farm)*
Nov 19 (Thurs)	Audio and 1D Signals
Nov 21-Nov 29	Fall Break
Dec 1 (Tues)	Reinforcement Learning
Dec 3 (Thurs)	Review, summary, looking forward
Dec 8 (Tues)	*No Class – Office hours for Final Project in Lecture Hall (CIF 0027) 12:30-1:30*
Dec 3-8	*Exam 3 at CBTF*
Dec 13 (Sun)	*Final Project due (cannot be late)*

Weeks 13-15

How does ML get deployed in practice?

Other application domains (audio) and methods (RL)

Sep 28 (Mon)	*HW 2 (PCA and Linear Models) due*
Sep 29 (Tues)	Outliers and Robust Estimation
Oct 1-4	*Exam 1 at CBTF* Optional review on Oct 1 in lecture
Oct 6 (Tues)	Decision Trees
Oct 8 (Thurs)	Ensembles and Random Forests
	Deep Learning
Oct 12 (Mon)	*HW 3 (PDFs and Outliers)*
Oct 13 (Tues)	Stochastic Gradient Descent
Oct 15 (Thurs)	MLPs and Backprop
Oct 20 (Tues)	CNNs and Keys to Deep Learning
Oct 22 (Thurs)	Deep Learning Optimization and Computer Vision
Oct 27 (Tues)	Words and Attention
Oct 29 (Thurs)	Transformers in Language and Vision
Nov 2 (Mon)	*HW 4 (Trees and MLPs) due*
Nov 3 (Tues)	Foundation Models: CLIP and GPT
Nov 5-8	*Exam 2 at CBTF (Optional review on Nov 5)*
	Applications
Nov 10 (Tues)	Ethics and Impact of AI
Nov 12 (Thurs)	Bias in AI, Fair ML
Nov 16 (Mon)	*HW 5 (Deep Learning and Applications) due*
Nov 17 (Tues)	Building and Deploying ML *Guest speaker: Chenxi Yu (State Farm)*
Nov 19 (Thurs)	Audio and 1D Signals
Nov 21-Nov 29	Fall Break
Dec 1 (Tues)	Reinforcement Learning
Dec 3 (Thurs)	Review, summary, looking forward
Dec 8 (Tues)	*No Class – Office hours for Final Project in Lecture Hall (CIF 0027) 12:30-1:30*
Dec 3-8	*Exam 3 at CBTF*
Dec 13 (Sun)	*Final Project due (cannot be late)*

Data you’ll use

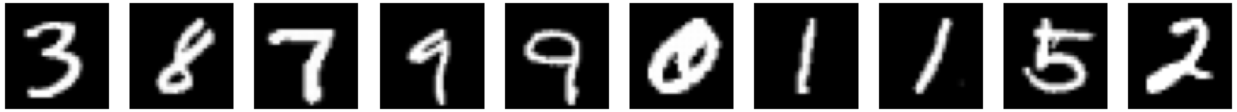

- Handwritten digits
- City Temperatures (predict next day)
- Text messages (spam vs not)
- Salaries
- Penguin measurements (predict species)
- Images of pets and flowers (breed/species classification)
- Your own data (final project)

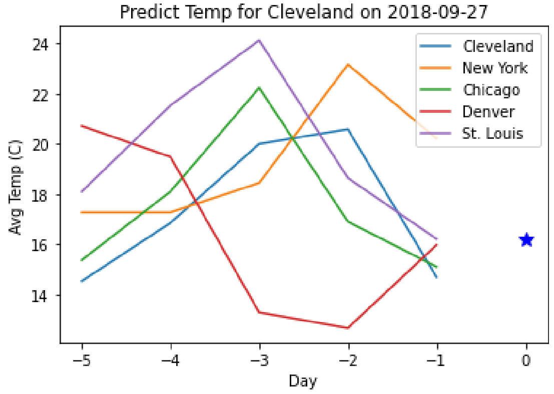

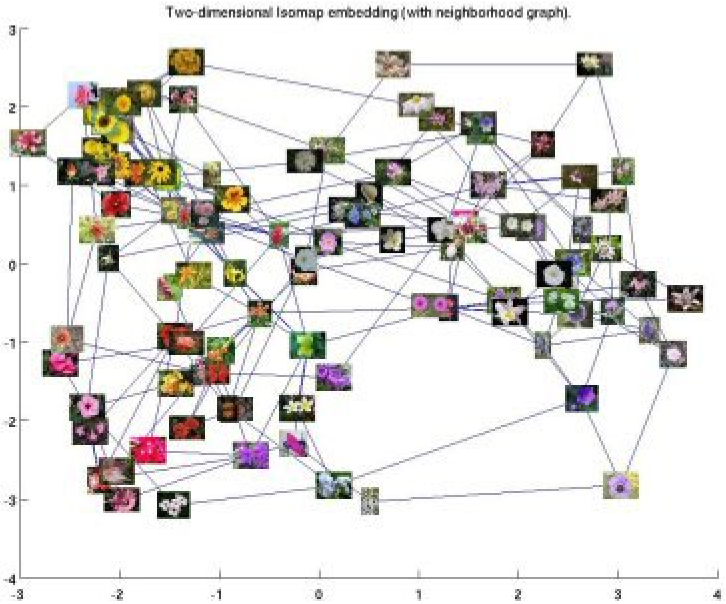

Grading

Experience Points: 400+ points

- 5 homeworks: 100+ points each
- Participation: 2 points per opportunity (up to ~30 opportunities)
- 3 credit: XP target is 400 points
- 4 credit: XP target is 525 points

Late policy (for regular assignments only)

- Up to ten free days total – use them wisely!
- 5 XP penalty per day after that
- Project must be submitted within two weeks of due date to receive any points

Exams: 100-200 pts

- Three exams
- 3 credit: must take third exam; others that increase your exam average count
- 4 credit: exam 3 + higher of first two exams

Final Project: 200 pts

- Planning on 2-3 graded (completion-based) design exercises worth 50-75 pts
- Final project report/video is rest

Final grade calculation

Course_Grade = (XP + exam_score + final_project) / (max(XP, XP_target) + exam_target + 200)

Earn 2 XP now – tell me a little about you

https://tinyurl.com/AML-L1

Sickness

- If you’re well, please come to lectures and office hours. Masks are optional.
- If you’re sick, please stay home. No need to show proof of illness or get permission to miss.

Using ChatGPT / AI models

- OK for assignments and study prep
  - Built into colabs
  - Can use CoPilot, Cursor, Claude Code, or others
  - Use for debugging, ask it to make up questions for you, or ask it to explain why you got something wrong
- Not available in exams

Academic Integrity

These are OK

- Discuss homeworks with classmates (don’t show each other code)
- Use Stack Overflow to learn how to use a Python module
- Use GPT/Co-Pilot/Gemini etc to learn how to use a Python module, or streamline coding
- Get ideas from online (make sure to attribute the source)

Not OK

- Copying or looking at homework-specific code (i.e. so that you claim credit for part of an assignment based on code that you didn’t write)
- Using external resources (code, ideas, data) without acknowledging them

Remember

- Ask if you’re not sure if it’s ok
- You are safe as long as you acknowledge all of your sources of inspiration, code, etc. in your write-up

How is this course different from…

- CS 446 ML
  - This course provides a foundation for ML practice, while 446 provides a foundation for ML research
  - This course has less theory, derivations, and optimization, and more on application, representations, and examples

- Online version of CS 441 AML
  - CS 441 online is being updated
  - This course focuses more on concepts and modern usage of ML, homeworks require more independence

- CS 444 Deep Learning for CV, other domain-oriented courses
  - This course is much broader

What is expected of you?

- You have sufficient background in math and CS, and/or will work hard to catch up in the first couple weeks
- You will attend all lectures (preferred), or watch the recordings
- If you get stuck, you'll try your best to figure it out, but then seek help if you can't
- You are willing to spend ~9-15 hours per week on lectures, reading, review and assignments

Feedback is welcome

- I will occasionally solicit feedback through surveys – please respond
- You can always talk to me after class or send me a message on CampusWire
- My goal is to be a force multiplier on how much you can learn with a given amount of effort

Action Items

- **Bookmark the website**
  https://courses.grainger.illinois.edu/CS441/fa2026/
- **Sign up for campuswire** (if not already signed up)
- **Read the syllabus and schedule**
- Unless you consider yourself highly proficient in Python/numpy and linear algebra, **watch/do the tutorials** linked in the web page

See you Thursday, on “KNN Classification”