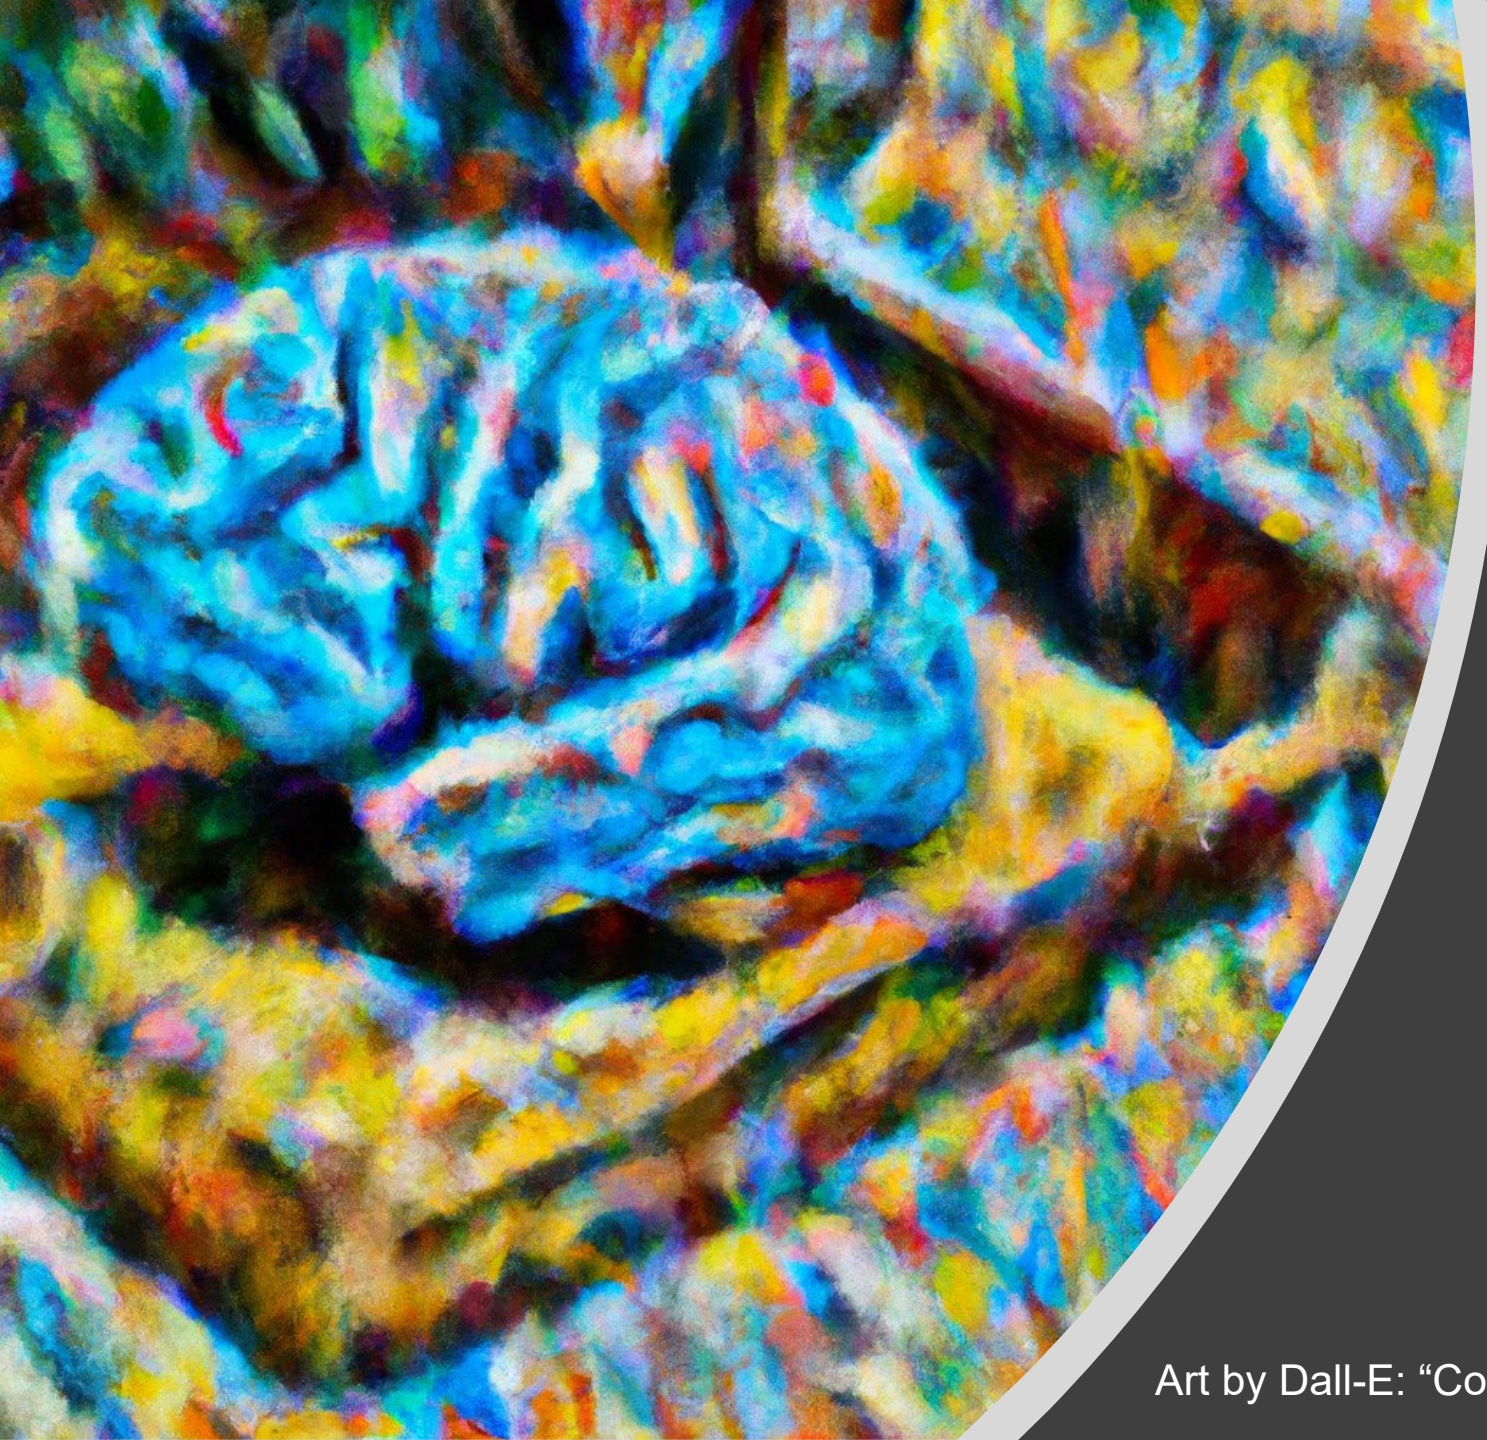

CS441 Applied Machine Learning

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

About me

1998-2002
Undergrad at SUNY Buffalo
B.S., EE and CSE

2002-2007
Grad at Carnegie Mellon
Ph.D. in Robotics

2007-2008
Postdoc at Beckman Institute

2009-
Prof in CS at UIUC

My research

Hoiem et al. 2005

Neural Radiance Fields: use deep networks to model 3D scenes

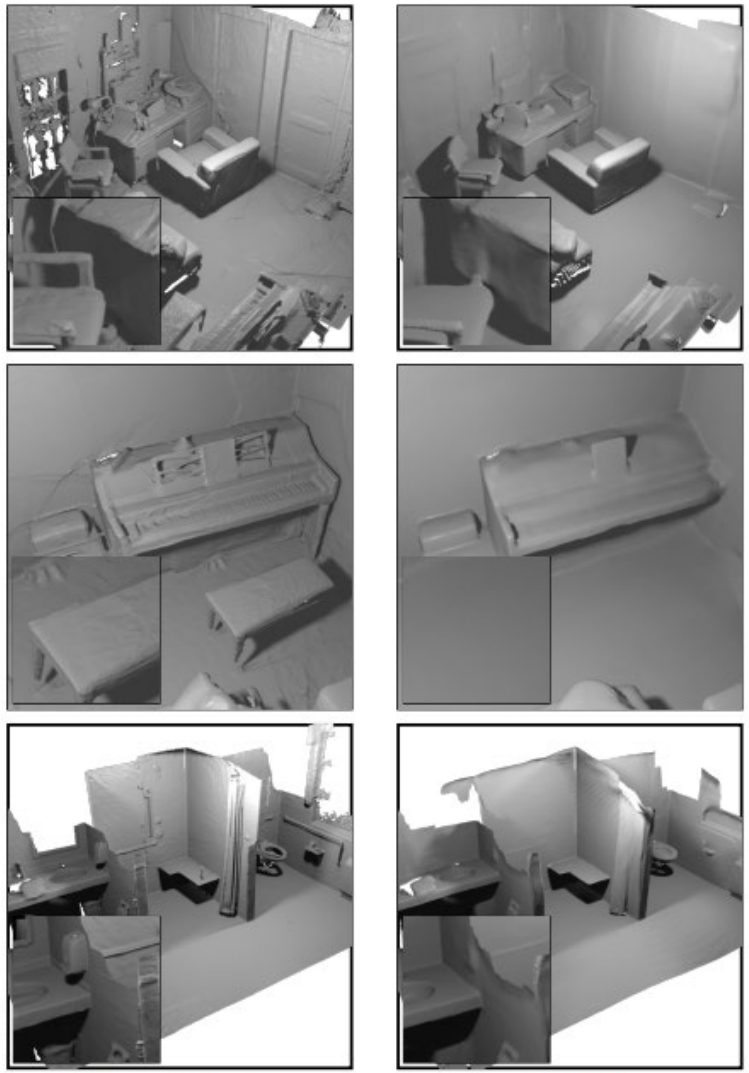

Ground Truth MLP

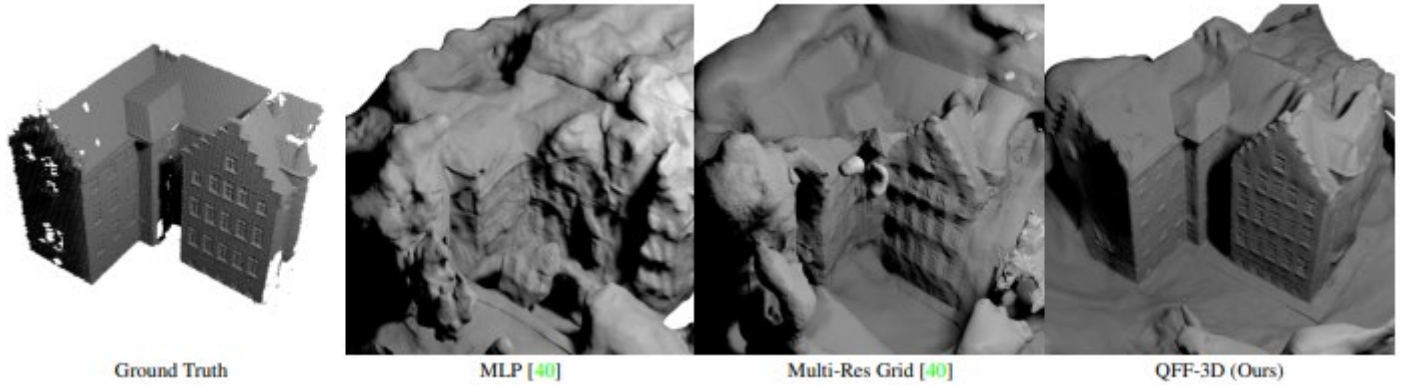

Ground Truth MLP [40] Multi-Res Grid [40] QFF-3D (Ours)

Lee et al. 2022

General Purpose Learners

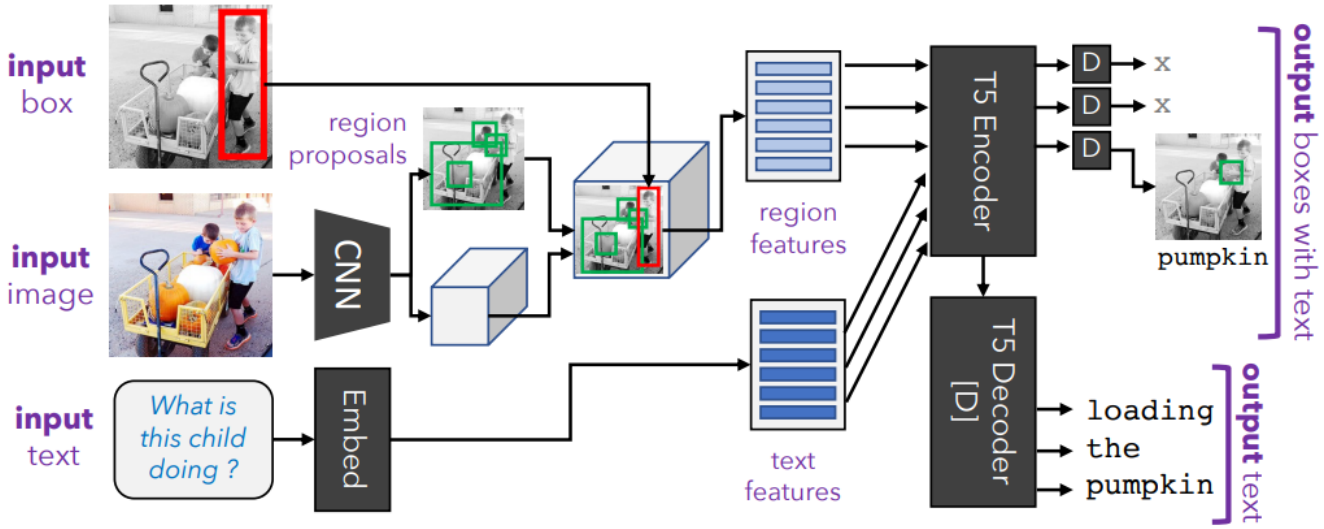

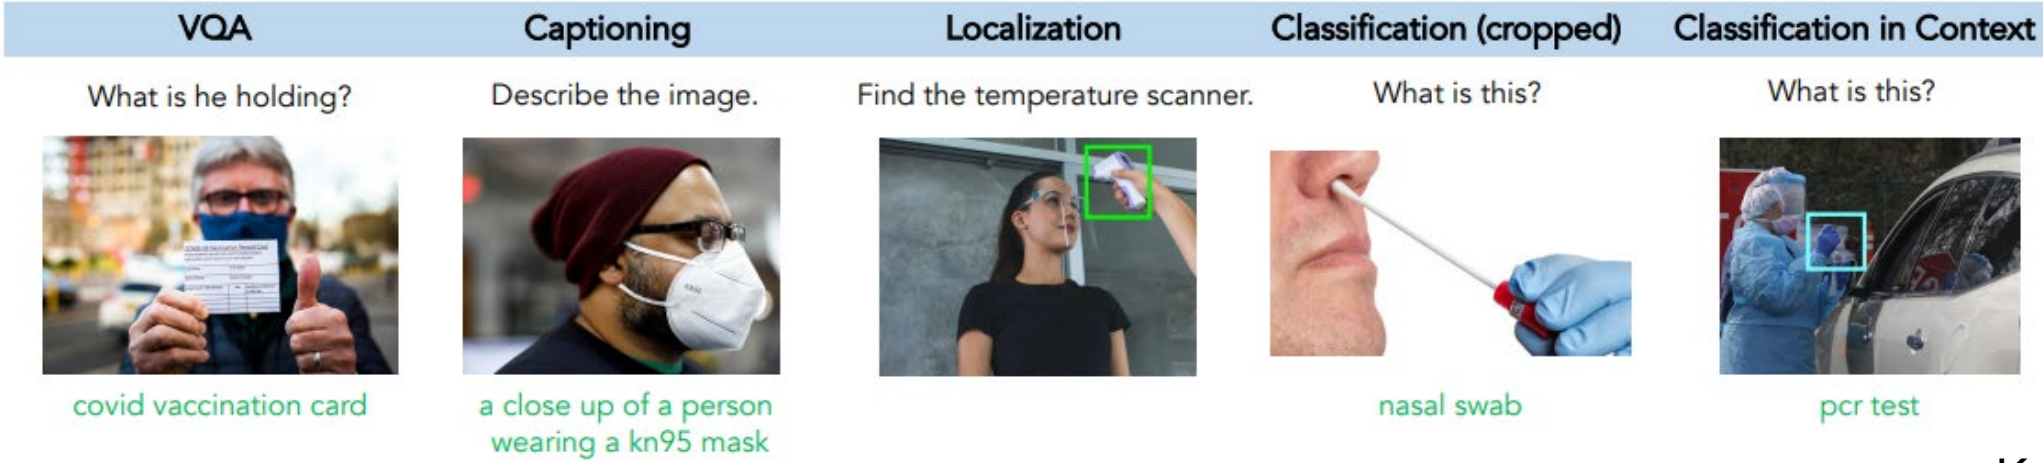

Kamath et al. 2022

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
  - General purpose vision-language

- Audio
  - Sound detection
  - Music identification

Reconstruct: vision for construction

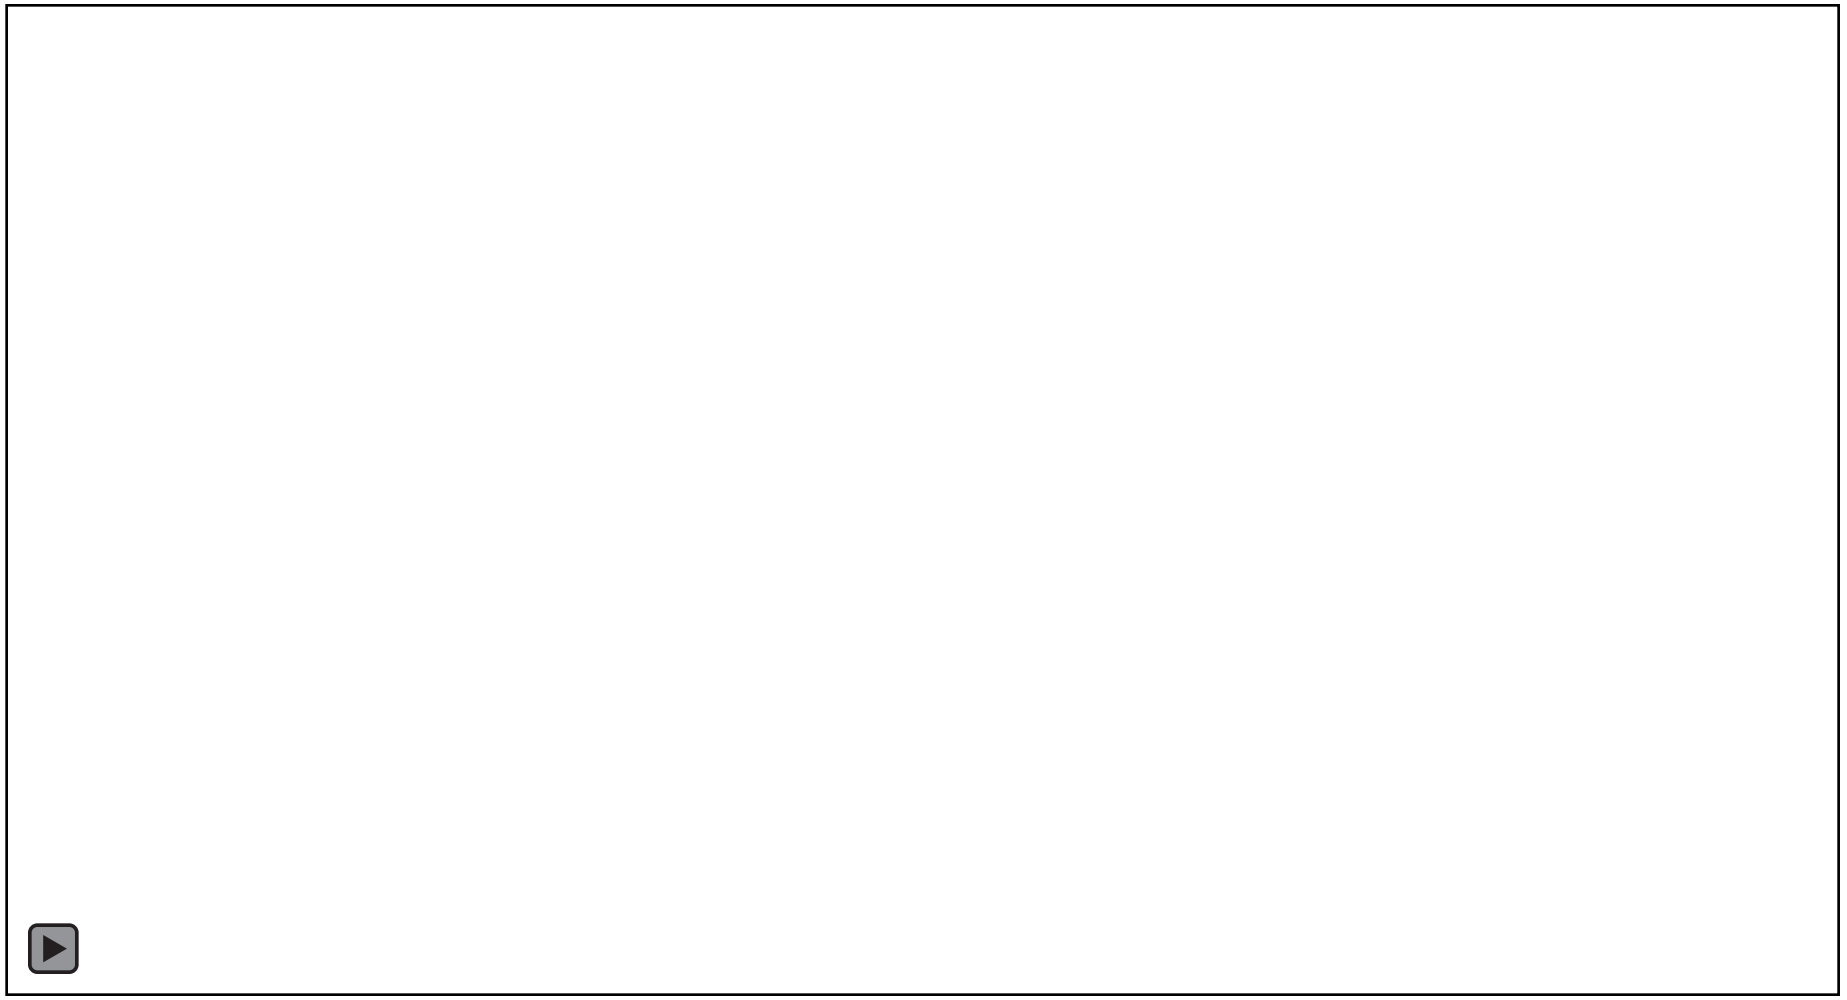

Crunchbase top 50 global startups

https://vimeo.com/242479887

https://www.reconstructinc.com/

What is machine learning?

- Create predictive models or useful insights from raw data
  - Alexa speech recognition
  - Amazon product recommendations
  - Tesla autopilot
  - GPT-3 text generation
  - Image generation
  - Data visualization

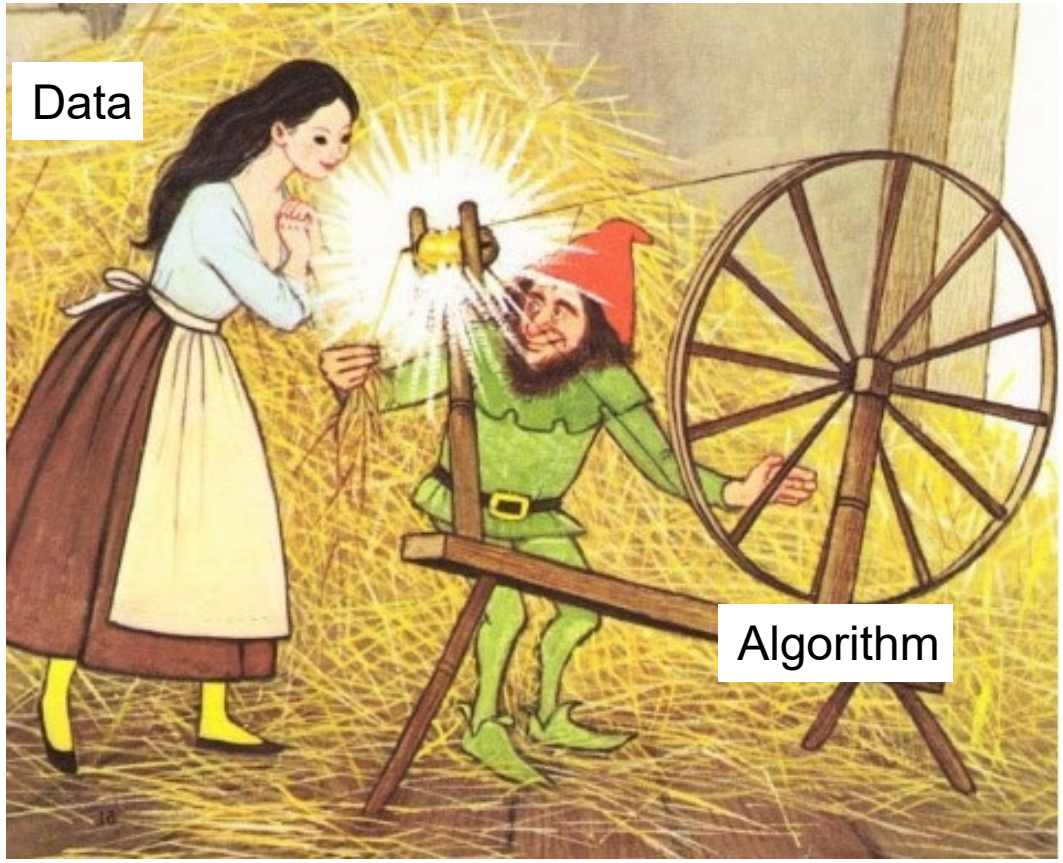

ML spins raw data into gold!

The whole machine learning problem

1. Data preparation
   a. Collect and curate data
   b. Annotate the data (for supervised problems)
   c. Split your data into train, validation, and test sets

2. Algorithm and model development
   a. Design methods to extract features from the data
   b. Design a machine learning model and identify key parameters and loss
   c. Train, select parameters, and evaluate your designs using the validation set

3. Final evaluation using the test set

4. Integrate into your application

Example: voice recognition in Alexa

Our focus, but it's important to understand all of it

Algorithm and model development

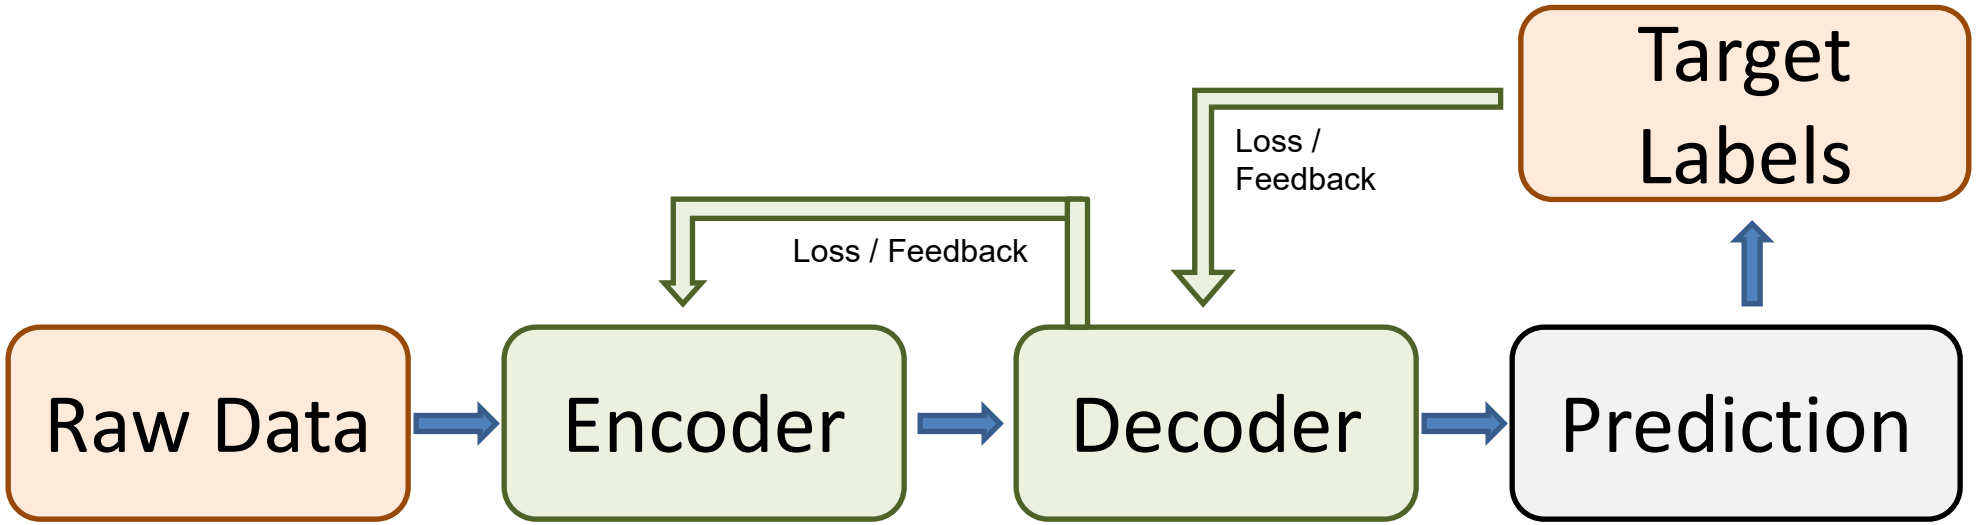

Text
Images
Audio
Discrete/continuous values
Structured/unstructured
Few/many examples
Few/many features
Clean/noisy labels

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

Linear regressor
SVN
Logistic regressor
Nearest Neighbor
Probabilistic model

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

Course objectives

1. Learn how to solve problems with ML
   - 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

The global machine learning market is expected to grow from $21.17 billion in 2022 to $209.91 billion by 2029, at a CAGR of 38.8%. With the field growing at such an exponential rate the number of jobs is growing too and machine learning is one of the most trending career paths of today. - Emeritus

2. Better understanding of real-life application and social implications of machine learning

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

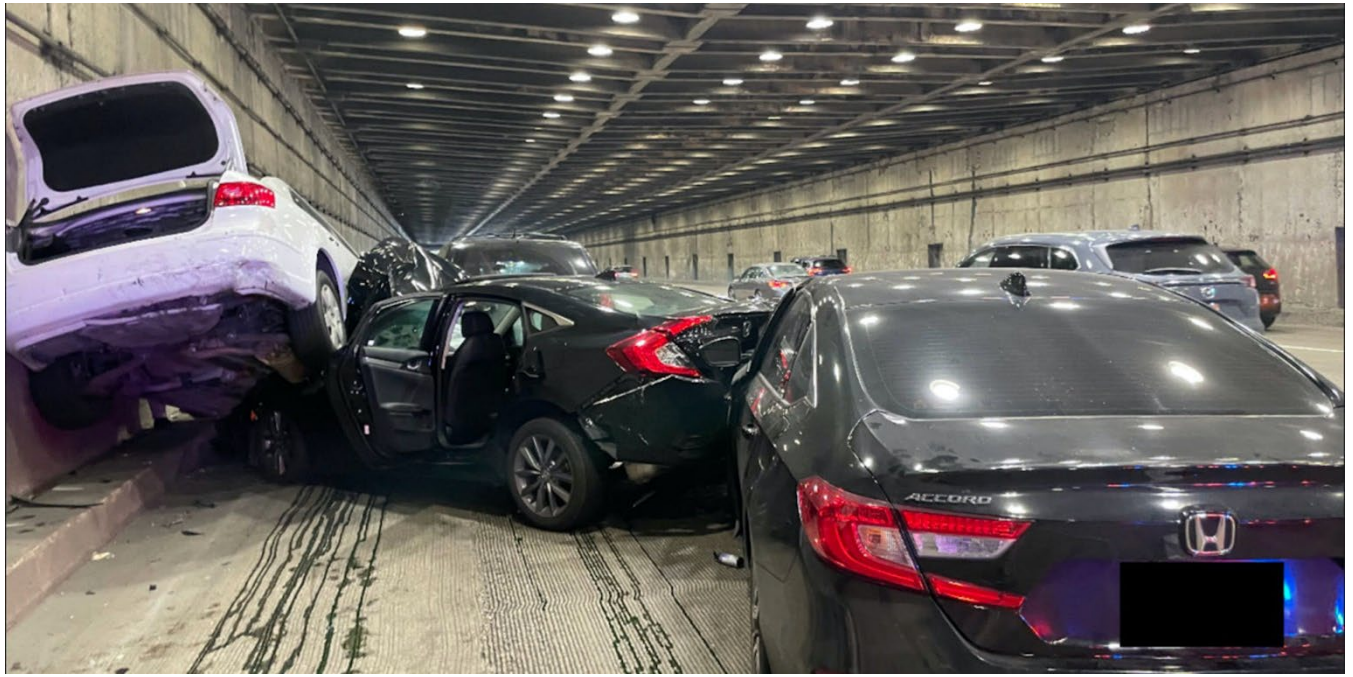

Tesla accident

3. Appreciation for your own constantly learning mind

Course outline

Prof: Derek Hoiem dhoiem@illinois.edu

TAs

- Vatsal Chheda (vchheda2)
- Joshua Levine (joshua45)
- Weijie Lyu (wlyu3)
- Kshitij Phulare (phulare2)
- Yuqun Wu (yuqunwu2)
- Mington Zhang (mz62)
- Wentao Zhang (wentao4)

Topics

- Supervised learning fundamentals
  - KNN, Naïve Bayes, Linear regression, logistic regression, trees, random forests, SVMs, neural networks, deep networks
- Application domains
  - Vision and CNNs, language models, transformers for vision and language, foundation models, task and domain adaptation, audio, ethics and impact, data issues
- Pattern discovery
  - Clustering and retrieval, missing data and EM, density estimation, topic models, outliers, data visualization, CCA

Grades

- Homeworks and final project (80%)
  - 4 homeworks: 100+ points each
  - 1 final project: 100 points (details TBD)
  - 3 credit: graded out of 450 points
  - 4 credit: graded out of 550 points
  - Up to 15 points extra credit

- Exams (20%)
  - Midterm 10%: covers first half
  - Final 10%: covers entire semester

Late policy

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

Covid, masks, sickness

- If you're well, please come to lectures and office hours. Masks are optional, per university policy. You're encouraged to follow CDC guidelines for masking.

- If you're sick, please stay home. No need to show proof of illness or get permission to miss.

- Lectures will be recorded, and exams can be taken from home

Homework details

- Implement and apply machine learning methods in Python notebooks

- Submit Report PDF and Jupyter notebook

Learning resources

Website: https://courses.engr.illinois.edu/cs441/sp2023/

- Syllabus
- Recordings
- CampusWire Discussion
- Canvas Submission
- Assignments
- Schedule
- Lecture slides and readings

Lectures

- In-person, recorded

Office hours

- Will be updated on pinned CampusWire post

Readings/textbook: Forsyth *Applied Machine Learning*

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
- 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

Other comments

Prerequisites

- Probability/stages, linear algebra, calculus
- Experience with Python will help but is not necessary, understanding that it may take more time to complete assignments

- Watch tutorials (see schedule: intro reading) for linear algebra, python/numpy, and jupyter notebooks.

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
  - This course has fewer, larger homeworks, a final project, and exams (vs. many small homeworks and quizzes)
  - This course focuses more on concepts and modern usage of ML
- CS 444 Deep Learning for CV
  - This course is much broader

Should you take this course?

Take this course if …

- You want to learn how to apply machine learning
- You like coding-based homeworks and are OK with math too
- You are willing to spend 10-12 hours per week (maybe even more) on lectures, reading, review, and assignments

Do not take this course if …

- You want more of a theoretical background (take 446 instead)
- You want to focus on one application domain (take vision, NLP, or a special topics course instead)
- You want an “easy A” (it’s not going to be easy)

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

What to do next

- Bookmark the website
- Sign up for campuswire
- 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