Responsible Machine Learning

Lecture 2: ML Basics

CS 4973-05

Fall 2023

Instructor: Avijit Ghosh
ghosh.a@northeastern.edu
Northeastern University, Boston, MA

Northeastern University

Today's lecture credits: Sargur N. Srihari University at Buffalo, State University of New York

Today AI/ML is ubiquitous

- Automate routine labor
  - Search

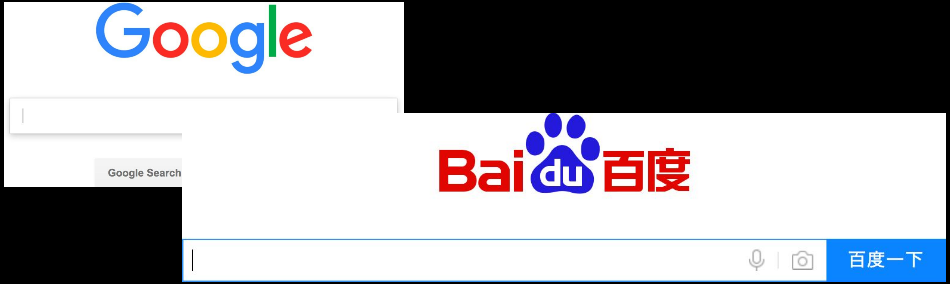

- Understand speech
  - SIRI, Alexa

- Autonomous Vehicles

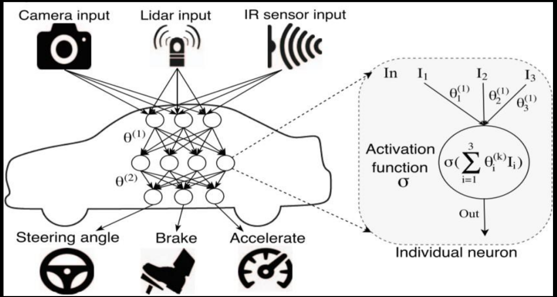

Rule-based programming is hard

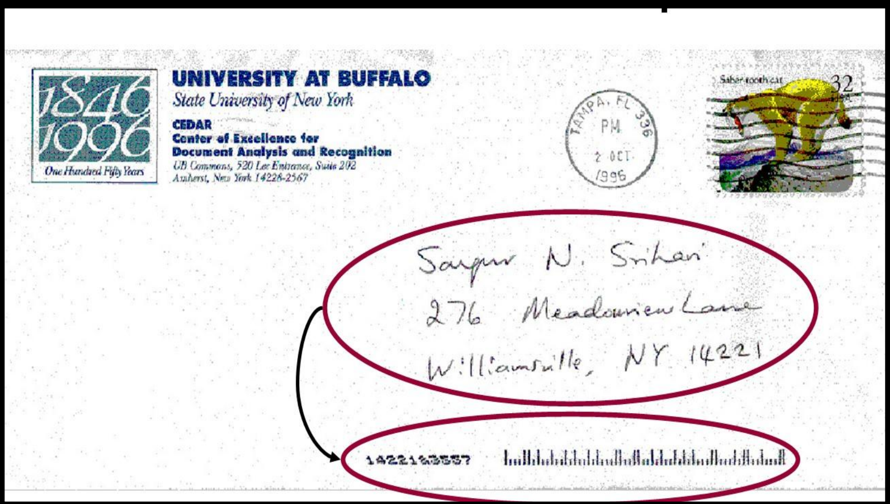

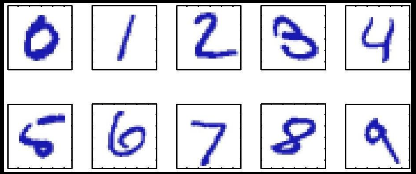

- Many handcrafted rules and exceptions
- Better learning from training set
- Handwriting recognition by a program cannot be done without ML!

AI Paradox

- Hard problems for people are easy for AI

- Easy problems for people are hard for AI

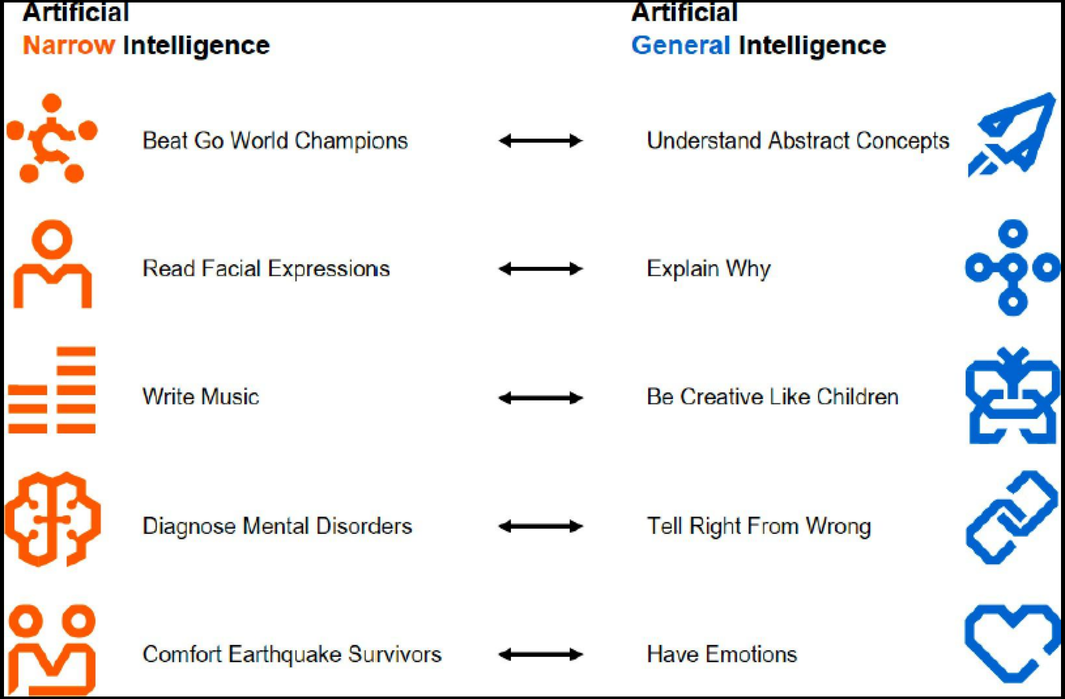

What Tasks Require Intelligence?

- **Reasoning**
  - Puzzles, Judgments
- **Planning**
  - Action sequences
- **Learning**
  - Improve with data
- **Natural language**
- **Integrating skills**
- **Abilities to sense, act**

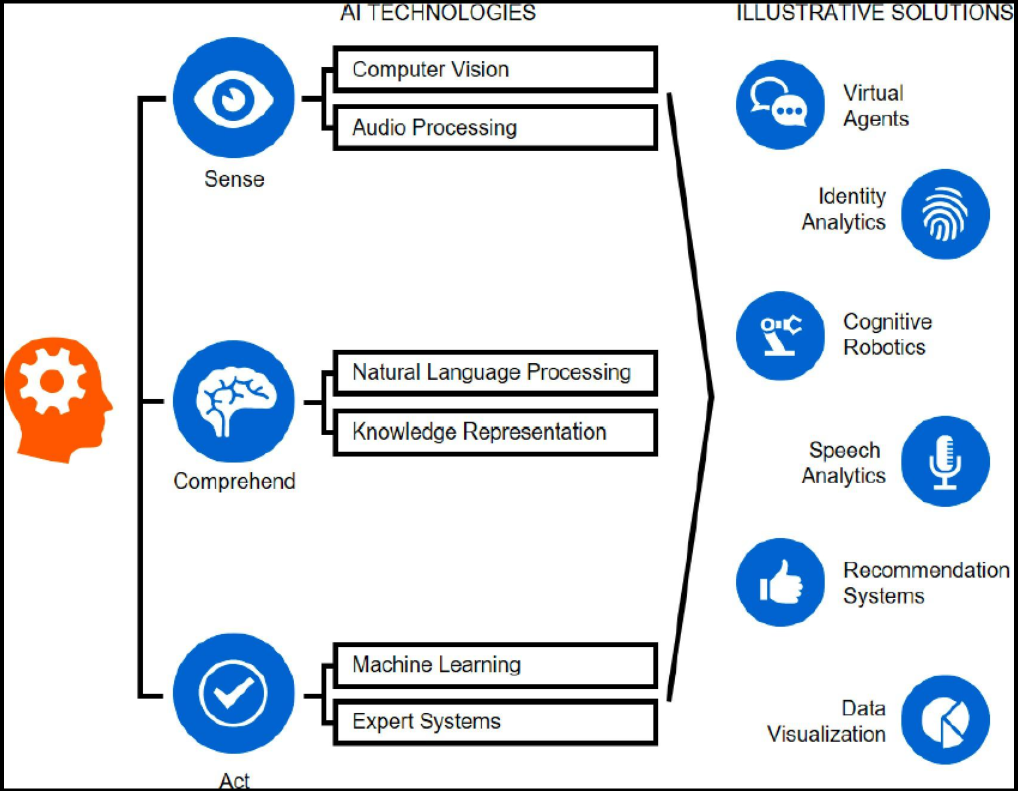

Northeastern University

Everyday life needs knowledge

- Knowledge is intuitive and subjective
  - Key challenge of AI is how to get this informal knowledge into a computer
- Knowledge-based Approach
  - Hard-code knowledge in a formal language
  - Computer can reason about statements in these languages using inference rules
  - Examples:
    - Expert systems for diagnosis (MYCIN, CADUCEUS)
    - Design (VAX)

Knowledge-Based AI

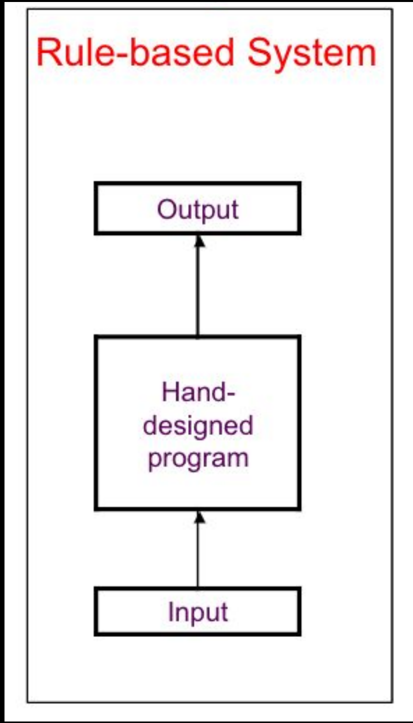

Disadvantage: Unwieldy process

Time of human experts
People struggle to formalize rules with enough complexity to describe the world

The Machine Learning Approach

- Difficulties of hard-coded approach suggests:
  - Allow computers to learn from experience
- First determine what features to use
- Learn to map the features to outputs

Linear classifier

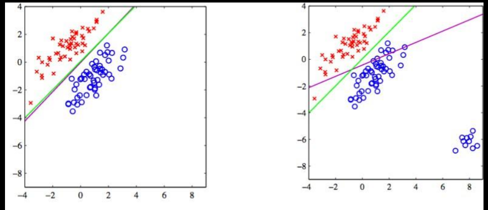

Magenta: Least Squares
Green: Logistic Regression (more robust)

Quadratic classifier

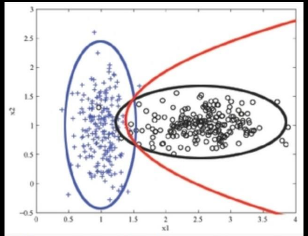

2D binary classification with Naive Bayes. A density contour is drawn for the Gaussian model of each class and the decision boundary is shown in red

Northeastern University

The ML Approach

Data Collection

Models

Model Selection

Prob. Dist. to Model Process

Parameter Estimation

Values/Distributions

↓

Inference

Find responses to queries

Generalization
(Training)

Decision
(Inference
OR
Testing)

ML Problem Types

1. Based on Type of Data
    - Supervised, Unsupervised, Semi-supervised
    - Reinforcement Learning
2. Based on Type of Output
    - Regression, Classification
3. Based on Type of Model
    - Generative, Discriminative

Supervised Learning

Training Data Labelled

- Most widely used methods of ML, e.g.,
  - Spam classification of email
  - Face recognizers over images
  - Medical diagnosis systems

- Inputs ***x*** are vectors or more complex objects
  - documents, DNA sequences or graphs

- Outputs are binary, multiclass(K),
  - Multi-label (more than one class), ranking,
  - Structured:
    - ***y*** is a graph satisfying constraints, e.g., POS tagging
  - Real-valued or mixture of discrete and real-valued

Northeastern University

Unsupervised Learning

Unlabeled data assuming underlying structure

1. Clustering to find partition of data
2. Identify a low-dimensional manifold
   - PCA, Autoencoder

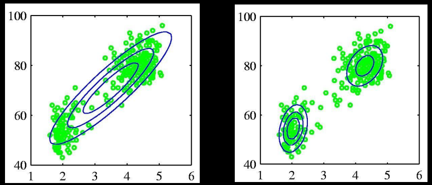

3. Topic modeling
   - Topics are distributions over words
   - Document: a distribution across topics
   - Methods: SVD, Collaborative Filtering

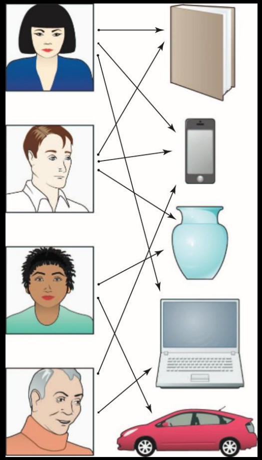

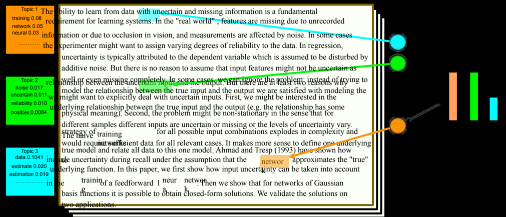

4. Recommendation Systems
   - Data links between users and items
   - Suggest other items to user
   - Solution: SVD, Collaborative Filtering

Reinforcement Learning

- **Training data in between supervised/unsupervised**
  - Indication of whether action is correct or not
  - Reward signal may refer to entire input sequence
  - Mouse is given a reward/punishment for an action
    - Policies: what actions to take in a particular situation
    - Utility estimation: how good is state (used by policy)
- **No supervised output but delayed reward**
- **Credit assignment**
  - what was responsible for outcome
- **Applications:**
  - Game playing, Robot in a maze, Multiple agents, partial observability,...

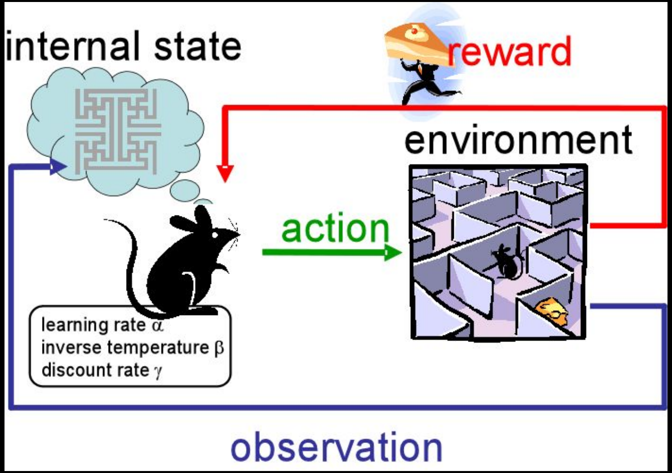

Regression

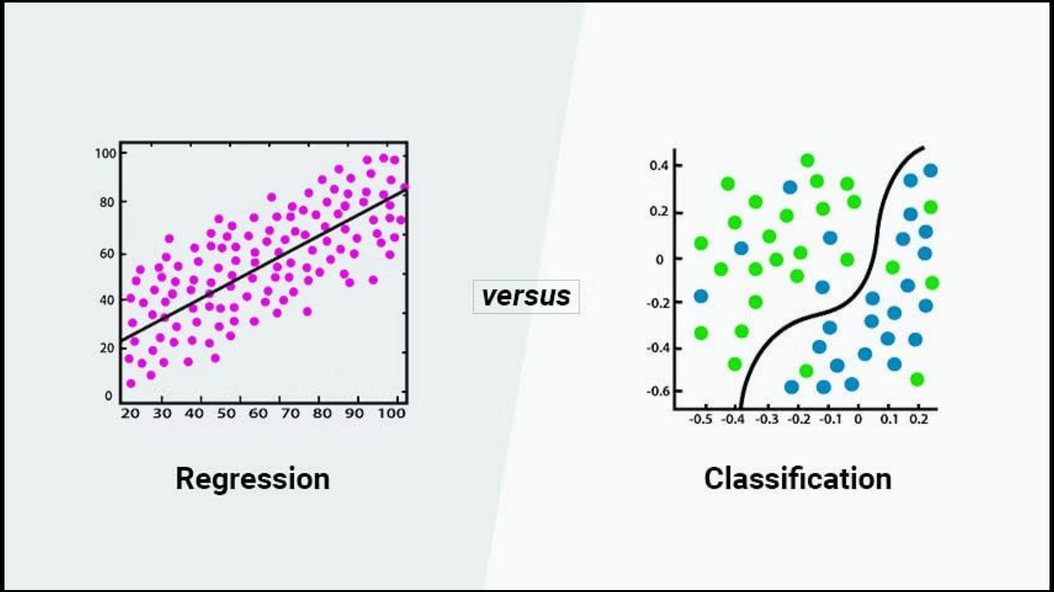

versus
Regression
Classification

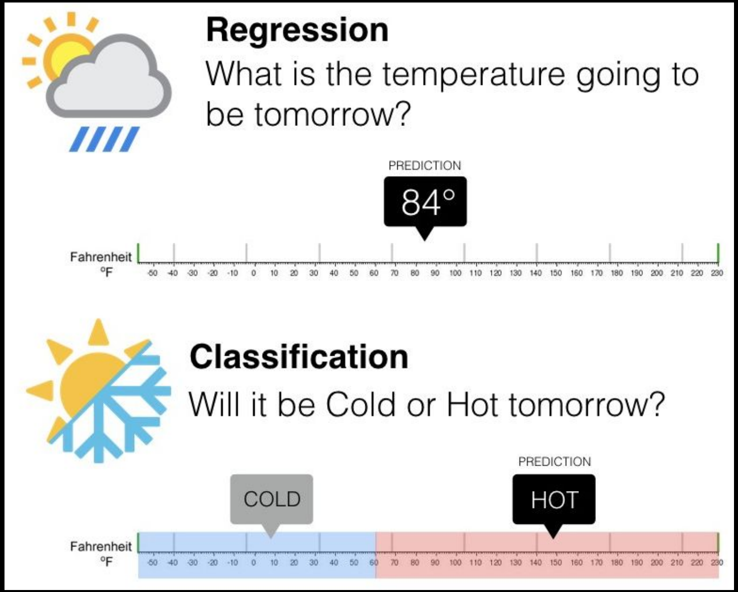

Regression
What is the temperature going to be tomorrow?
PREDICTION
84°
Fahrenheit °F
Classification
Will it be Cold or Hot tomorrow?
PREDICTION
COLD
HOT
Fahrenheit °F

Regression: Learning to Rank

Input (x_i):
(*d* Features of Query-URL pair)

- Log frequency of query in anchor text
- Query word in color on page
- # of images on page
- # of (out) links on page
- PageRank of page
- URL length
- URL contains "~"
- Page length
- Embedding vectors

Output (*y*):
Relevance Value

Target Variable
- Point-wise (0,1,2,3)
- Regression returns continuous value
- Allows fine-grained ranking of URLs

Two Paradigms of AI

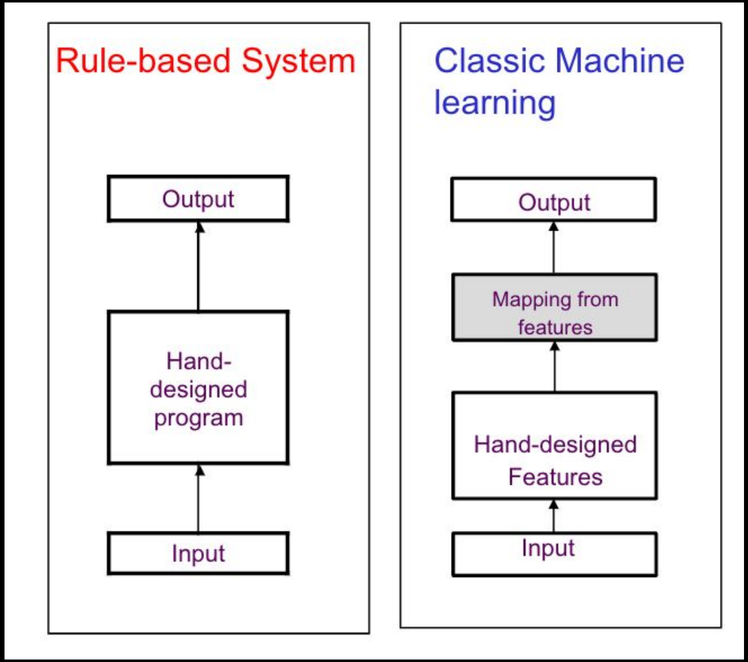

Shaded boxes indicate components that can learn from data

Designing the right set of features

- Simple Machine Learning depends heavily on *representation* of given data
- For detecting a car in photographs
  - Tire shape difficult in terms of pixel values
  - Shadows, Glare, Occlusion etc.

Representation Learning

- Solution: use ML to not only learn mapping from representation to output but representation itself
  - Better results than hand-coded representations
- Allows AI systems to rapidly adapt to new tasks
  - Designing features can take great human effort
    - Can take decades for a community of researchers
- Does not need programmer to have deep knowledge of the problem domain

Feature Learning for Classification

- Function to map pixels to object identity is complicated
- Series of hidden layers extract increasingly abstract features
- Final decision made by a simple classifier

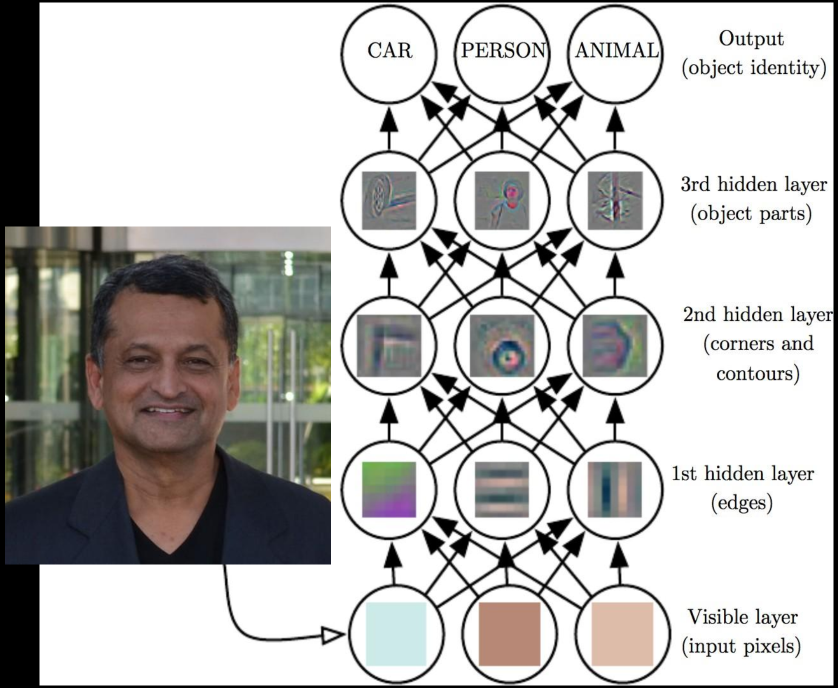

Deep Learning

- Understand the world as hierarchy of concepts
  - How these concepts are built on top of each other is deep, with many layers
  - Weights learnt by gradient descent

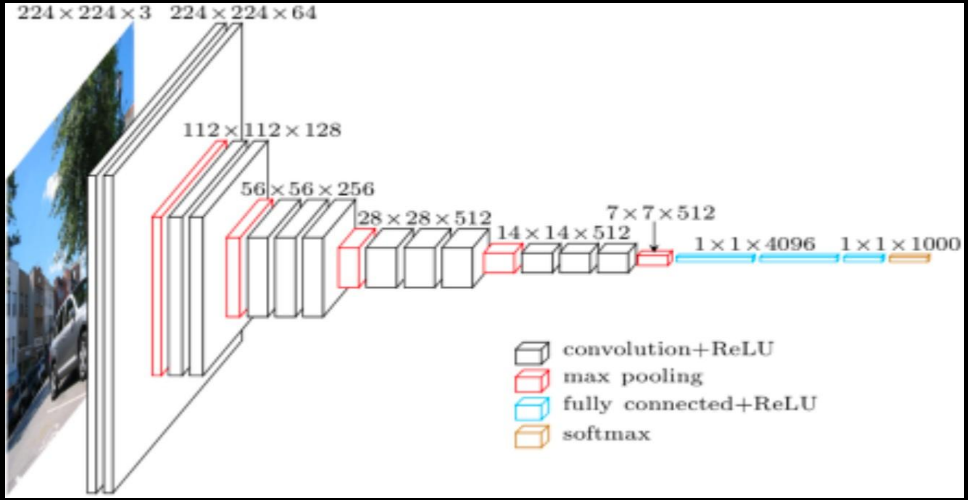

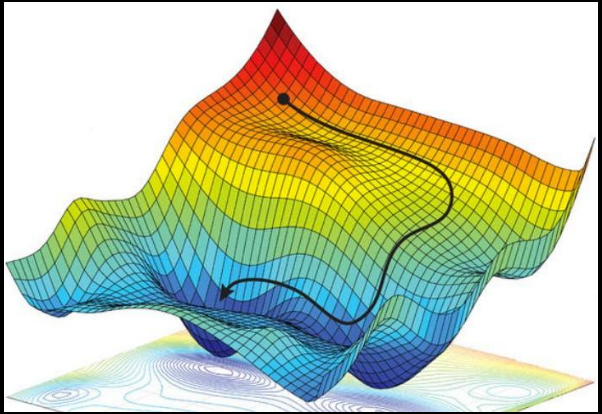

$$x^{t+1} = x^{t} - \varepsilon \nabla f\left(x^{t}\right)$$

Paradigms of AI

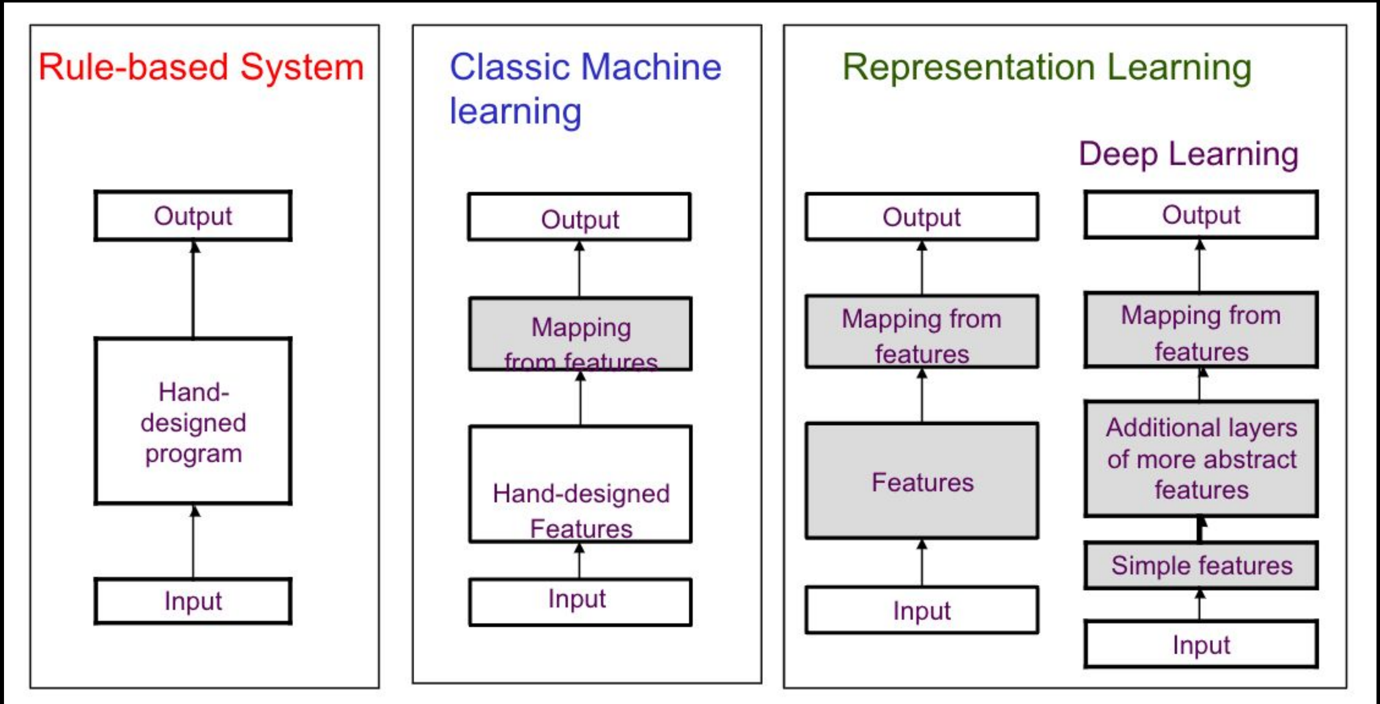

Shaded boxes indicate components that can learn from data

Northeastern University

And Finally: Generative AI

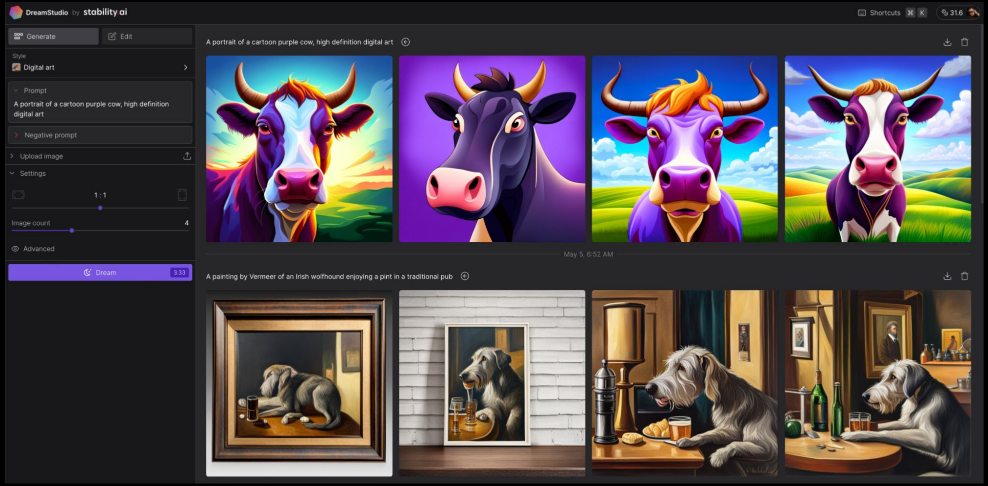

Stable Diffusion (Image Generation)

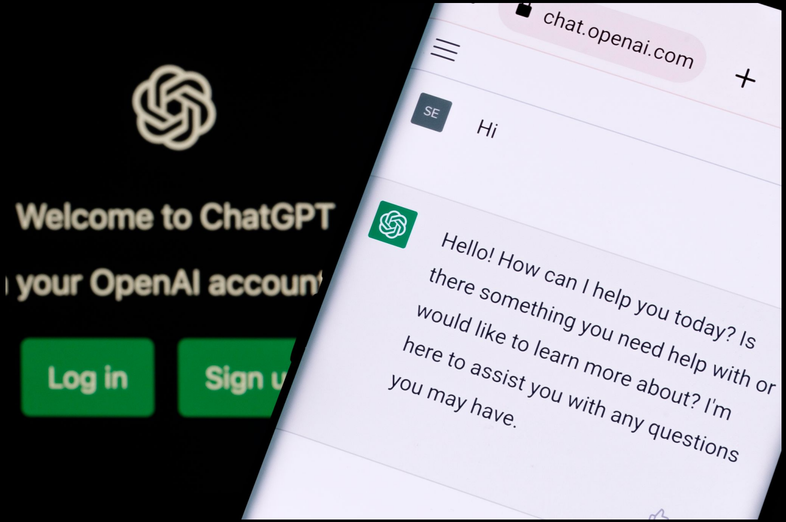

ChatGPT (Text Generation)

Northeastern University

Thank You!

No homework next class

- We will be doing Python+ML setup and some example code in class, led by Jeffrey