CSI 436/536
Introduction to Machine Learning

General introduction

Professor Siwei Lyu
Computer Science
University at Albany, State University of New York

What is (machine) learning

- **learning** is the process of **converting experience into** knowledge to perform certain tasks

  

- **machine learning** is to program computers so that they can “learn" from input available to them
  - input to a learning algorithm is **training data**, representing experience
  - the output is expertise, taking the form of another **computer program** that can perform some **task**.

Example: spam filter

- Design a spam filter
  - input: a piece of texts (emails)
  - output: a label (0: no spam, 1: spam email)
- Traditional approach: notice that words like “4U” “free” “credit cards” “amazing” tend to show up in a lot of spam emails, then program rules for classification

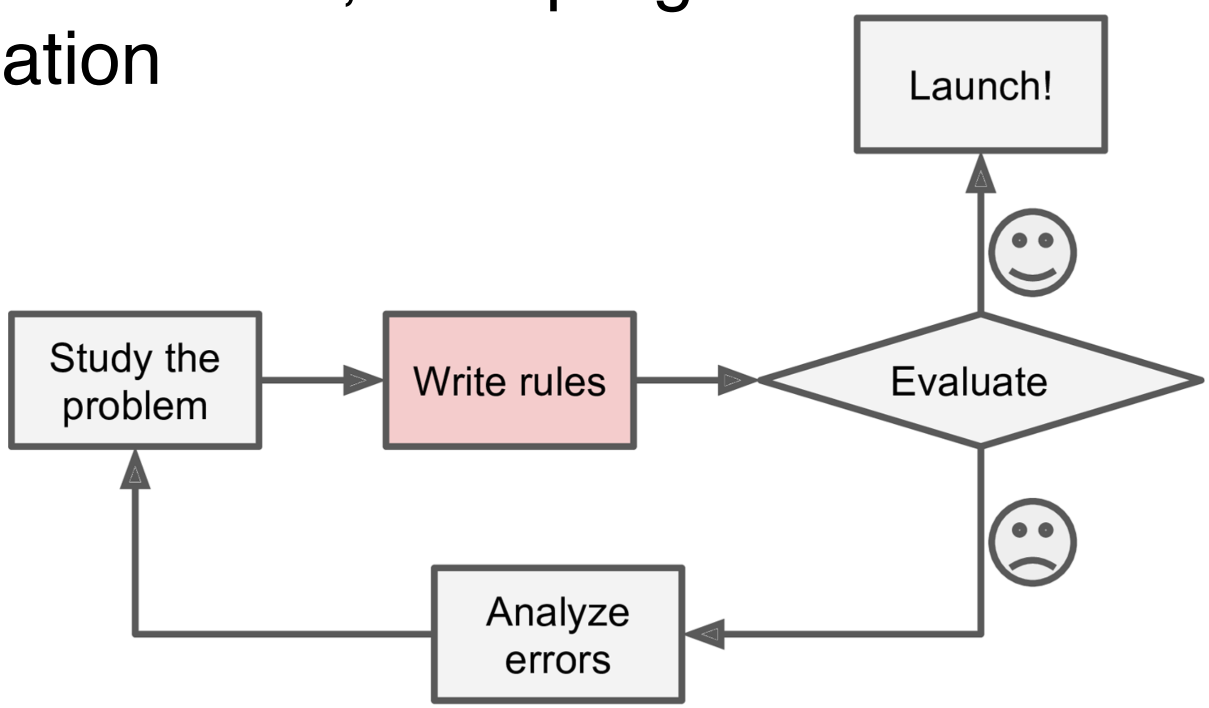

ML approach

- using training data to improve algorithms

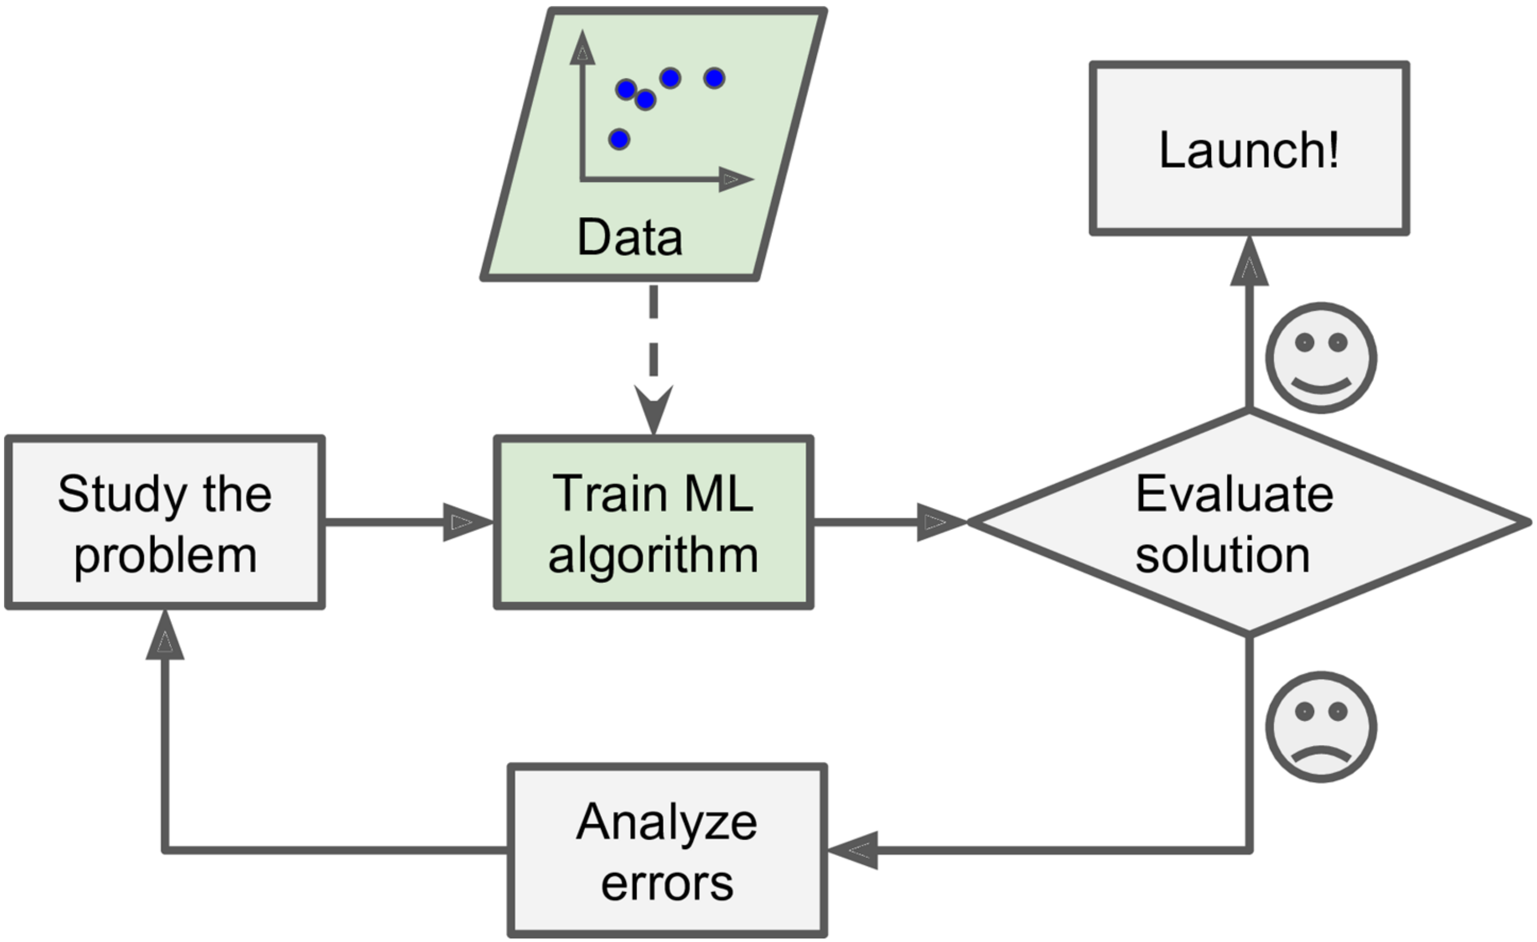

When we need machine learning

- Tasks that are too complex to write program
- Tasks that need program to adapt to input data
- Data are in very high volume

What are the tasks

- Currently most ML tasks are about prediction

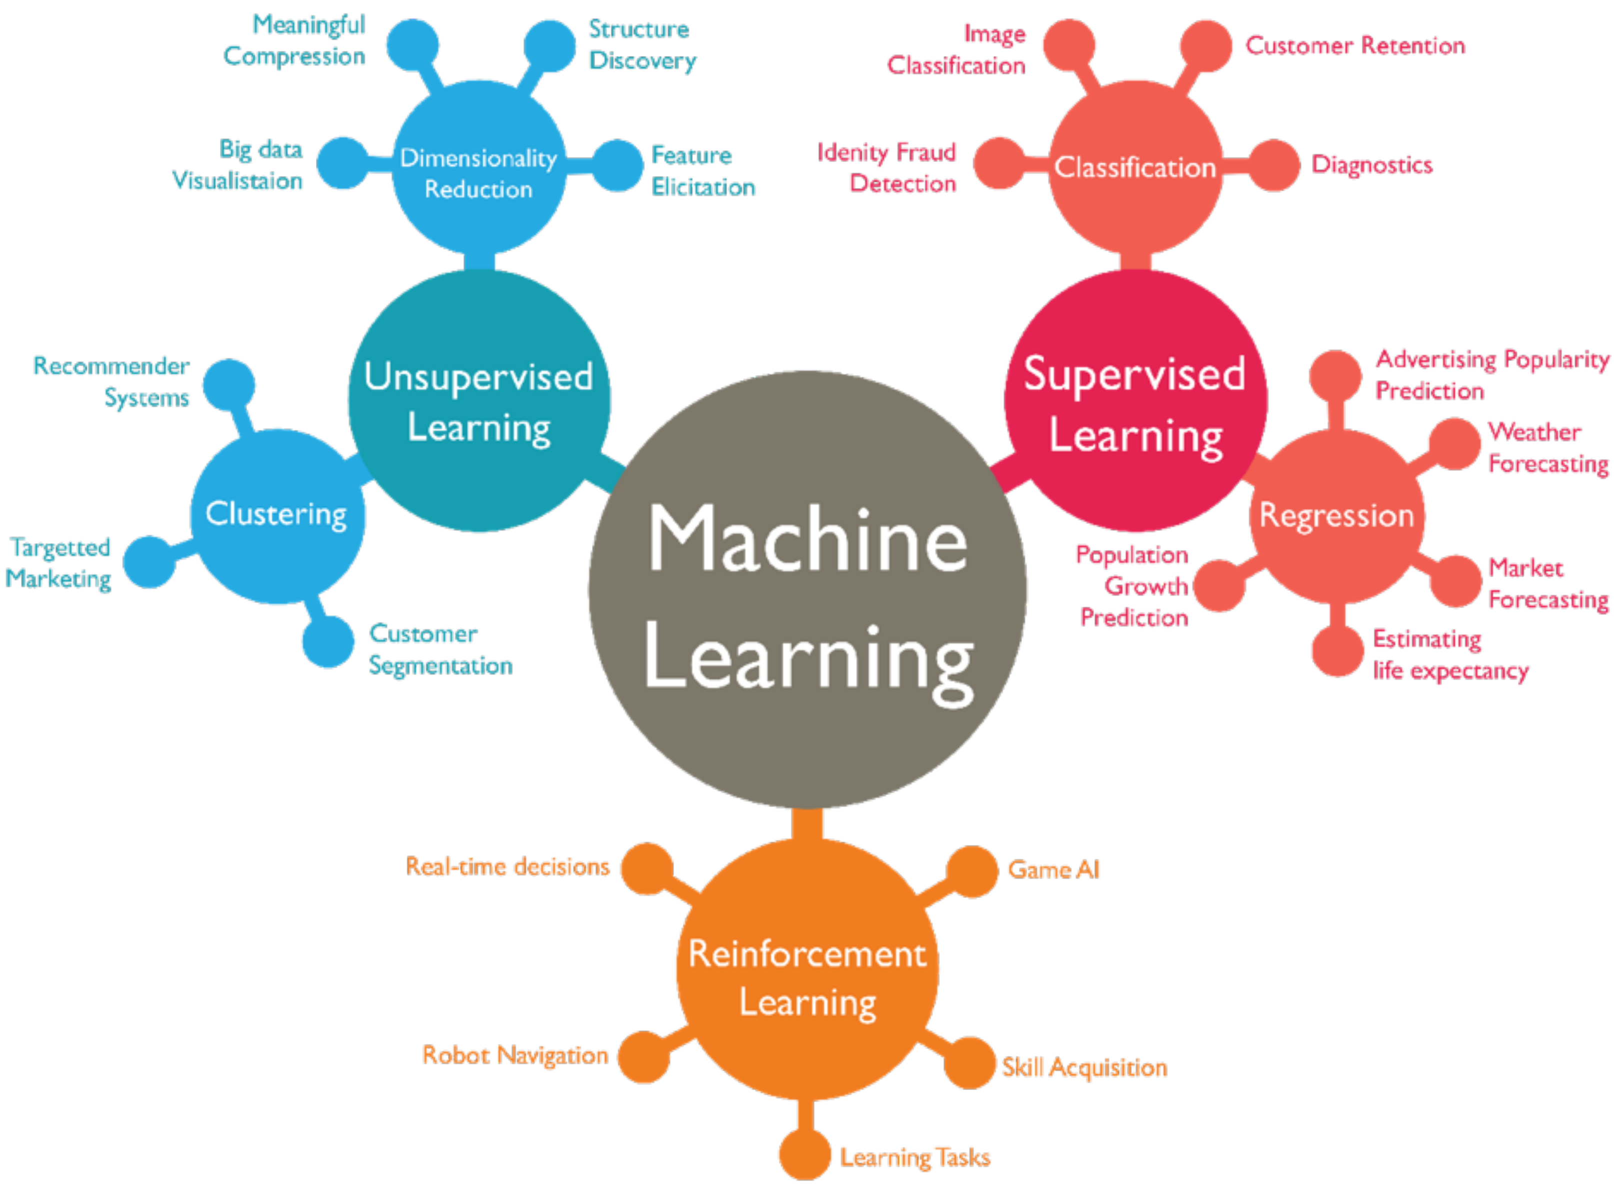

Related fields

 Artificial Intelligence	 Machine Learning	 Deep Learning
Artificial intelligence originated around 1950s.	Machine learning originated around 1960s.	Deep learning originated around 1970s.
AI represents simulated intelligence in machines.	Machine Learning is the practice of getting machines to make decisions without being programmed.	Deep Learning is the process of using Artificial Neural Networks to solve complex problems.
AI is a subset of Data Science.	Machine learning is a subset of AI & Data Science	Deep learning is a subset of Machine learning, AI & Data Science.
Aim is to build machines which are capable of thinking like humans.	Aim is to make machines learn through data so that they can solve problems.	Aim is to build neural networks that automatically discover patterns for feature detection.

Relation with other fields

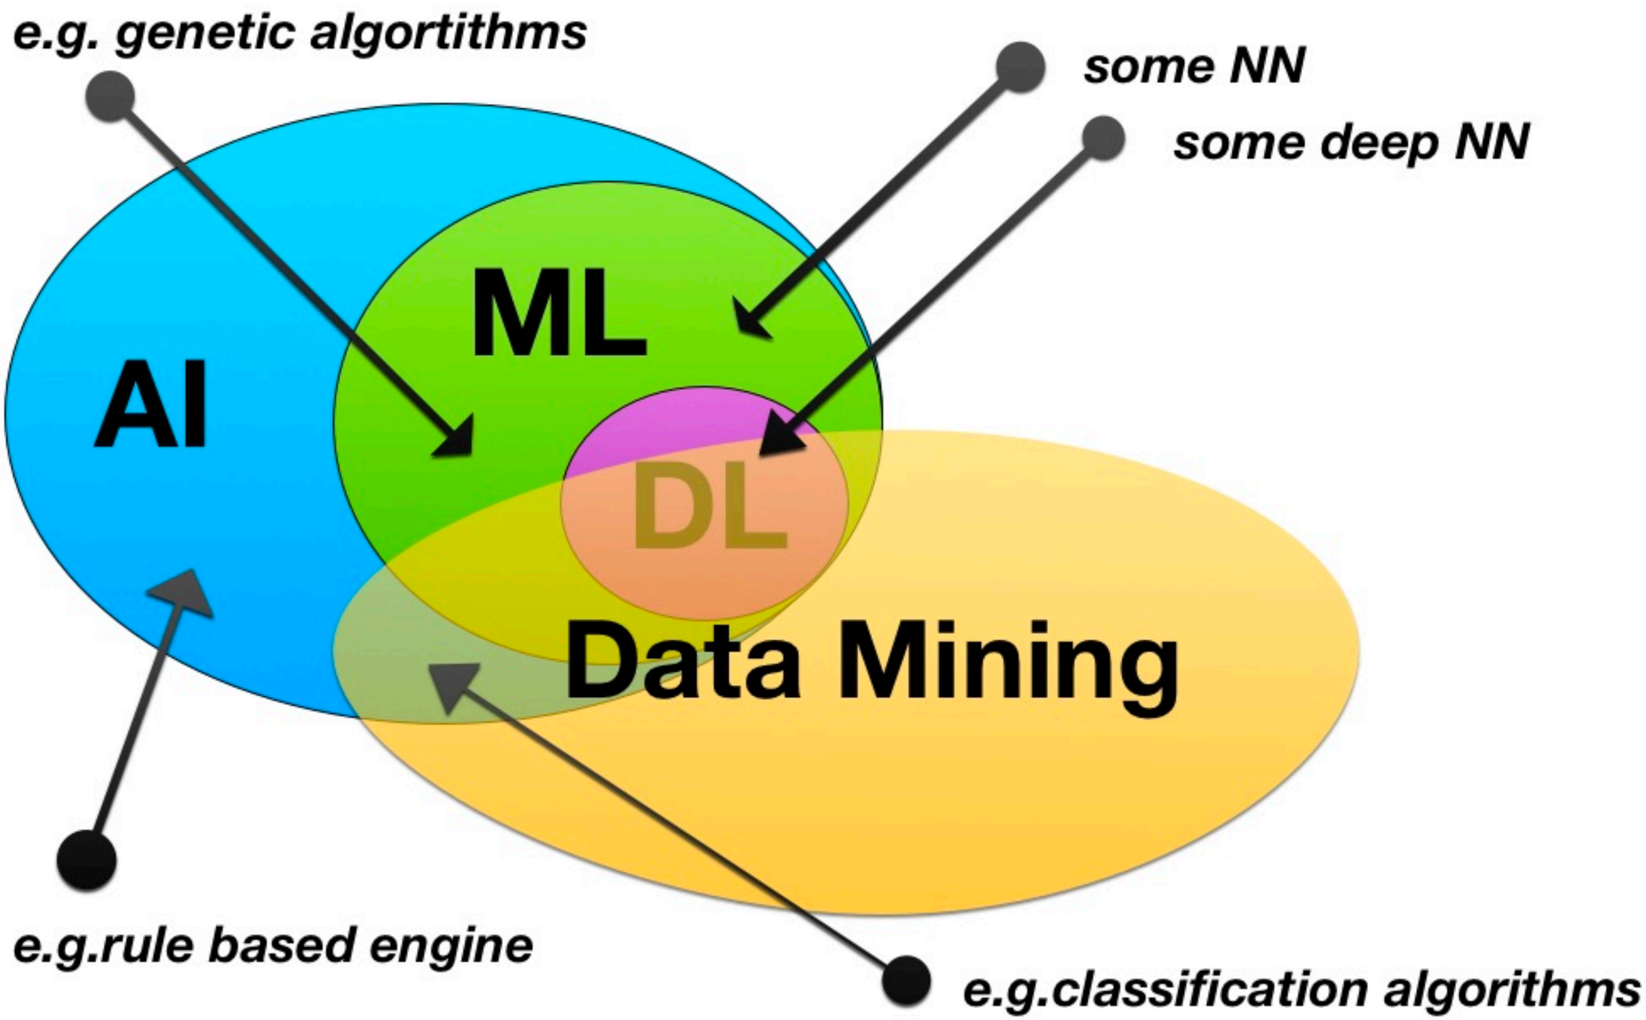

How machine learns (typically)?

- A set of *training data*
- A parametric family of models that relates input & output
- A *training* algorithm to learn the model
  - Usually as minimizing a *learning objective* for the model parameter by *numerical optimization*

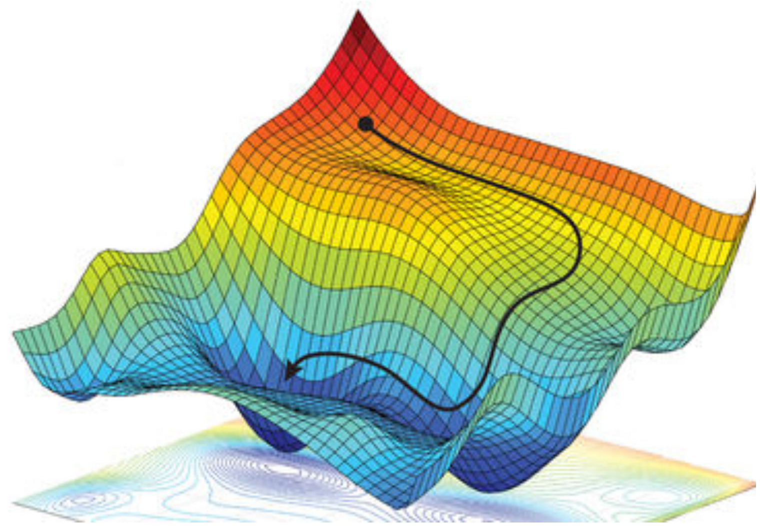

- Parameters of the training algorithm, known as the *meta-parameters*, need to be tuned on an independent *validation* dataset
- A *metric* to evaluate the quality of the learned model on an independent *test* dataset

Keys to successful (machine) learning

- Reflect the experience about learning a new skill (a new language, a new sport, play music instruments, etc)
  - You usually do not start with a clean slate (you already know English before your learn French)
  - You need to practice, practice, practice
  - You may need some sort of feedback on performance
  - As a result, you build knowledge for similar tasks (it helps you to learn Italian)

Keys to successful (machine) learning

- Reflect the experience about learning a new skill (a new language, a new sport, play music instruments, etc)
  - You usually do not start with a clean slate (you already know English before your learn French)
    - **Prior knowledge**
  - You need to practice, practice, practice
    - **Training process with lots of data**
  - You may need some sort of feedback on performance
    - **Task specific objective/metric**
  - As a result, you build knowledge for similar tasks (it helps you to learn Italian)
    - **Generalization and adaptation**

Machine learning topics

- Machine learning theories
  - How learning can occur, what guarantee (worst and best case) we can expect
- **Machine learning algorithms**
  - Design algorithm to solve a learning problem
- Machine learning systems
  - adapt algorithm to code running on specific hardware platform and OS environments
    - CPU, FPGA, GPU, cloud, edge, mobile platform, embedded system, real-time etc
- Machine learning applications
  - Computer vision/NLP/speech/big data, etc

Machine learning algorithms

- **Geometrical and algebraic approaches**
  - LSE, PCA, LDA, SVM, spectral clustering, etc
- **Functional approaches (connectionism)**
  - NN, DNN, SGD
- Statistical and Bayesian approaches
  - EM, MLE, Bayes nets, Markov random fields, MCMC
- Logic (rule-based) approaches
  - Markov logic networks
- Evolution approaches
  - Genetic algorithms
- Meta-learning algorithms

A brief history of ML

- Pre 1980: simple (mostly linear) algorithms, mostly developed in statistics & AI
  - PCA, LDA, MDS, least squares regression perceptron, k-means, EM algorithm, factor analysis
- 1980s: development of nonlinear learning algorithms
  - decision trees, neural networks, Bayes networks
- 1990s: SVMs, probabilistic graphical models, boosting and random forest, sparse coding, NMF, ICA
- 2000 – 2010: nonparametric Bayesian learning
  - Gaussian Process, Dirichlet Process, HDP (Chinese restaurant process)
- 2010 - present: deep learning on very large scale dataset
  - DNN, Deep Auto-encoder, CNN (comeback of NN), RNN, GAN

Why learning is difficult?

- there are many cofounding factors and it is difficult to set up casual relations, learning involves knowing which is important

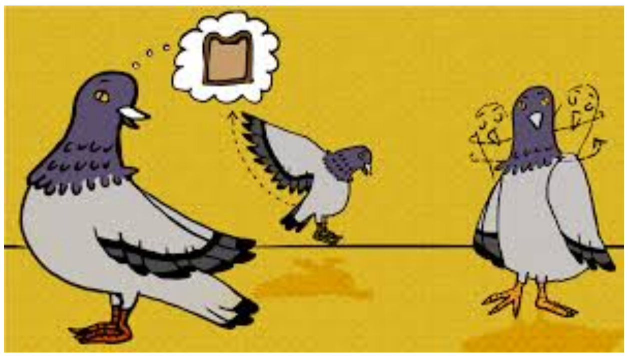

pigeon superstition

rat bait shyness

- learning requires the incorporation of *prior knowledge* that biases the learning mechanism
- Central theme of machine learning is to expressing domain expertise, translating it into a learning bias, and quantifying the effect of such a bias

Overfitting

- Machine learning aims to use finite available data to obtain algorithms that can work for any unseen data
- however, conclusions obtained on finite data set (experiences) are not necessarily true to the unknown data
    - e.g., all swans we have seen are white, and, therefore, all swans are white
    - e.g., the sun rises everyday till today, will it rise tomorrow?

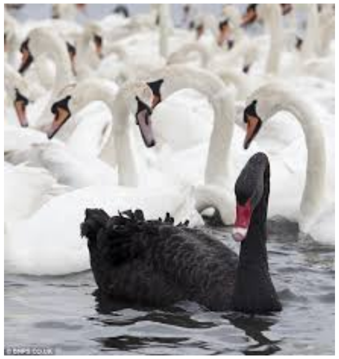

- When an algorithm only works for the finite training data but not the unseen data, we say it "***over-fits***" the training data

Machine learning is not a silver bullet

- ML are powerful tools, but they are still tools
  - Learning is not necessarily understanding
  - Data + model do not always lead to insight
- ML models brings specific issues to the society
  - Explainability and interpretability, and ultimately accountability
  - Security and trustworthiness of ML algorithms
  - Ethics and fairness issues of ML models and datasets
  - Misuse and weaponization of ML models

Current hot topics

- **Deep** learning
- **Big and Fast** learning
- **Real life** learning
- **Human in the loop** machine learning
- **Explainable & Interpretable** machine learning
- **Safe & Secure** machine learning
- **Fairness & Transparency** in machine learning

major conferences/journals

- conferences
  - core conferences: NIPS, ICML, ICLR, UAI, AAAI, IJCAI, IJCNN, AI-STATS, COLT, ECML, ACML
  - computer vision: CVPR, ICCV, ECCV, ACCV, WACV, BMVC
  - robotics: RSS, IROS, ICRA
  - NLP: ACL, COILING
  - Speech: Interspeech, ICASSP
- journals
  - JMLR, ML, IEEE TPAMI, IJCV, PR

Course topics

- Part 1: overview and preliminaries
- Part 2: linear least squares and applications
  - Regression, model selection, online learning, classification, clustering, dimension reduction,
- Part 3: basic machine learning algorithms
  - Dimension reduction (PCA), clustering (k-means, spectral clustering), classification (LDA, SVM), kernel methods
- Part 4: deep neural networks and applications