Class Logistics and Machine Learning Review

Alexander G. Ororbia II

Introduction to Machine Learning

CSCI-736

1/14/2025

Course Page/Syllabus Up

- Syllabus and policy:
- https://www.cs.rit.edu/~ago/courses/736/index.html
- Prerequisites:
  - CSCI 630: Foundations of Intelligent Systems
  - ***Or*** equivalent background (exposure to ML)

- Introduce yourselves
  - Name, program/department, and research goals/topics of interest

Your Textbook

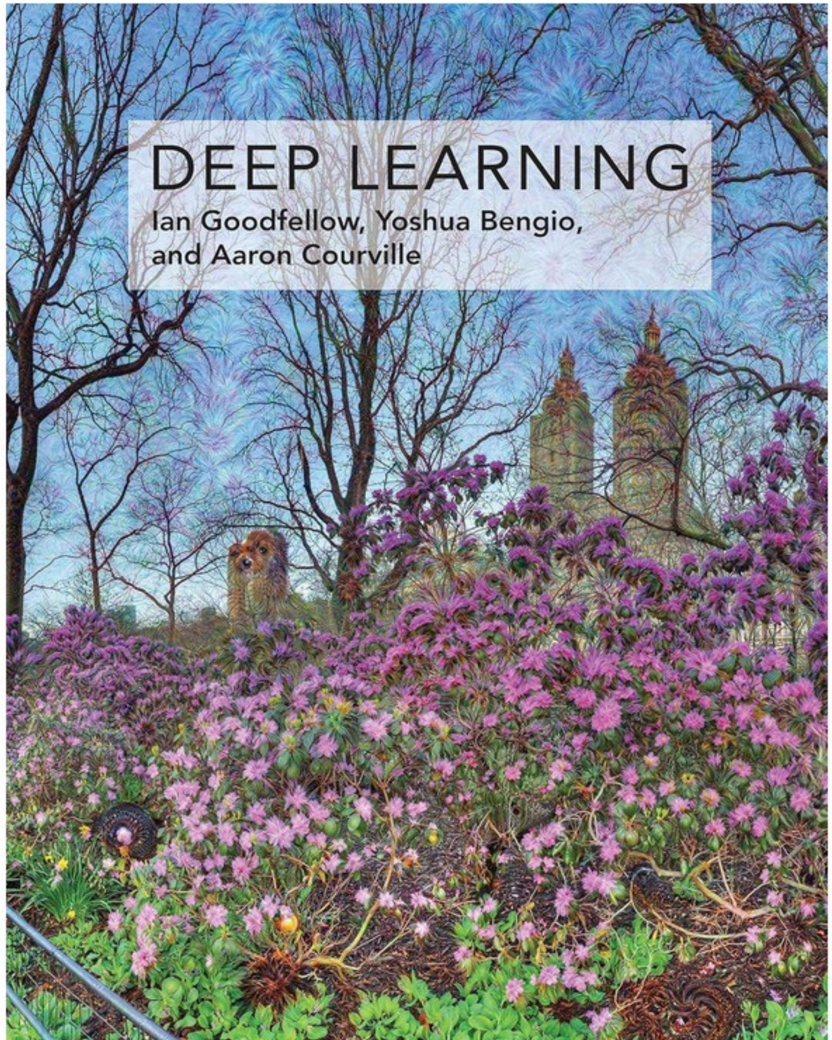

Though it's worth owning and having on your own bookshelf too!

Free online: https://www.deeplearningbook.org/

Objectives for Today

- What is machine learning (ML)?
- What is representation learning?
- Conclusions
- Next time

Some Useful Prerequisites

- Basic algorithmic knowledge
- Some linear algebra (matrices/vectors, operations)
- Multivariate calculus

What is Machine Learning?

- A branch of **artificial intelligence**, concerned with the design and development of algorithms that allow computers to evolve their behavior based on empirical data
  - Automating automation
  - Getting computers to program themselves
  - Writing software is the bottleneck
  - Instead, let the data do the work instead!

- Intelligence requires knowledge, thus it is necessary for computers to acquire knowledge

Traditional Programming

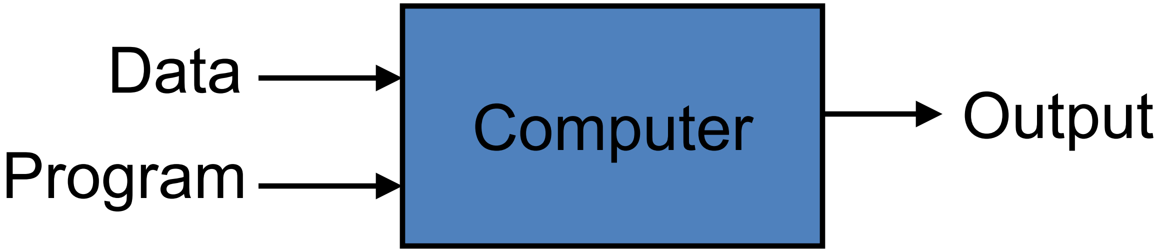

Machine Learning

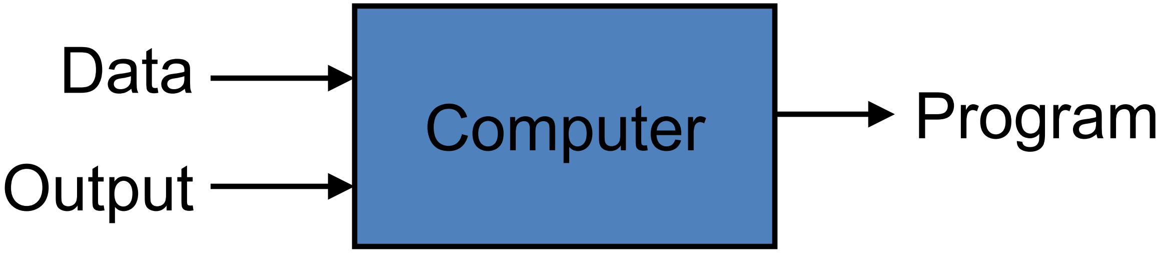

Is it magic??

No, it's more like gardening

- **Seeds** = Algorithms
- **Nutrients** = Data
- **Gardener** = You
- **Plants** = Programs

ML in a Nutshell

- Tens of thousands of ML algorithms
- Hundreds pop up every year

- Every machine learning algorithm has three components or "pillars":
  - **Representation**
  - **Evaluation**
  - **Optimization**

If you took CSCI 635 with me, you probably got sick of these three words!

Representation

- Your data and your model choice
- Some models you might use:
  - Decision trees
  - Sets of rules / Logic programs
  - Instances
  - Graphical models (Bayes/Markov nets)
  - Neural networks
  - Support vector machines
  - Model ensembles
  - Etc., etc.

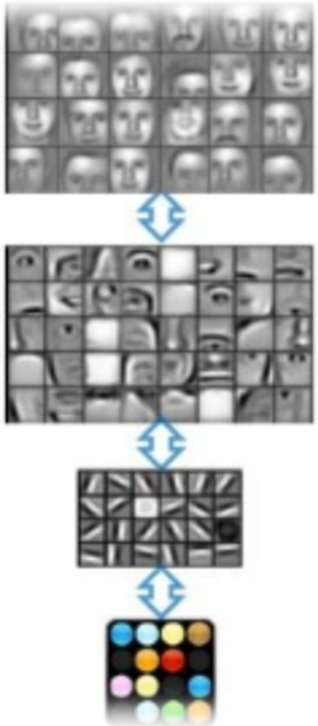

Evaluation

- Your objective function and evaluation metrics
- Objectives
  - Squared error
  - Likelihood
  - Posterior probability
  - Cost / Utility (reward functionals)
  - Margin
  - Entropy, KL divergence
  - Etc.
- Evaluation
  - Accuracy
  - Precision and recall
  - PSNR, SSIM
  - Etc.

Optimization

- How you shape/change free parameters in your model(s)
- Combinatorial optimization
  - E.g.: Greedy search
- Convex optimization
  - E.g.: Gradient descent
- Constrained optimization
  - E.g.: Linear programming

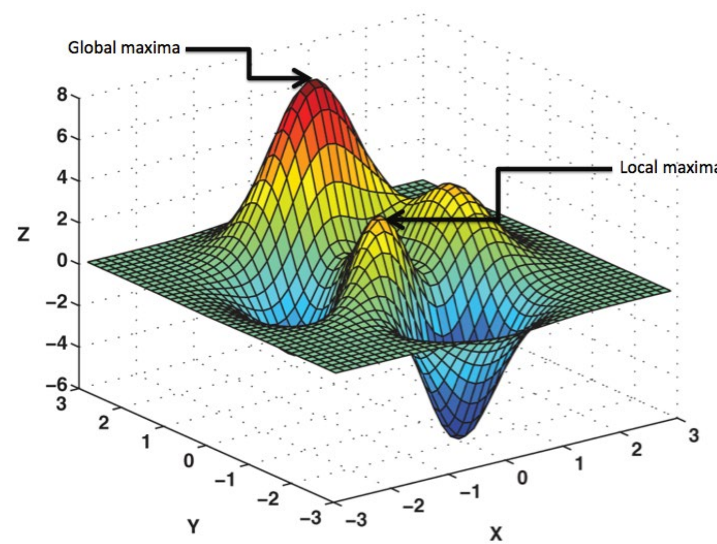

Performance

- There are several factors affecting the performance:
  - **Types of training** provided
  - The form and extent of any initial **background knowledge**
  - The **type of feedback** provided
  - The **learning algorithms** used
- Two important factors: *Modeling* and *Optimization*
- The success of machine learning system also depends on the algorithms used/employed
- Algorithms control the search to find and build knowledge structures
- Learning algorithms should extract useful information from training examples

Types of Learning

- **Supervised (inductive) learning** $\{x_n \in R^d, y_n \in R\}_{n=1}^N$
  - Training data includes desired outputs
  - Prediction / Classification (discrete labels), Regression (real values)
- **Unsupervised learning** $\{x_n \in R^d\}_{n=1}^N$
  - Training data does not include desired outputs
  - Clustering / probability distribution estimation
  - Finding association (in features)
  - Dimension reduction
- **Semi-supervised learning**
  - Training data includes a few desired outputs
- **Reinforcement learning**
  - Rewards from sequence of actions
  - Decision making (robot, chess machine)

Inductive Learning

- **Given** examples of a function *(X, F(X))*
- **Predict** function *F(X)* for new examples *X*
  - Discrete *F(X)*: Classification
  - Continuous *F(X)*: Regression
  - *F(X)* = Probability(*X*): Probability estimation

Questions?