

CSE 6243

Advanced Topics in Machine Learning

Bo Dai
School of CSE, Georgia Tech

Course Introduction



Bo Dai

Assistant Professor
School of Computational Science and Engineering

Address: Coda E1342A

Email: bodai@cc.gatech.edu

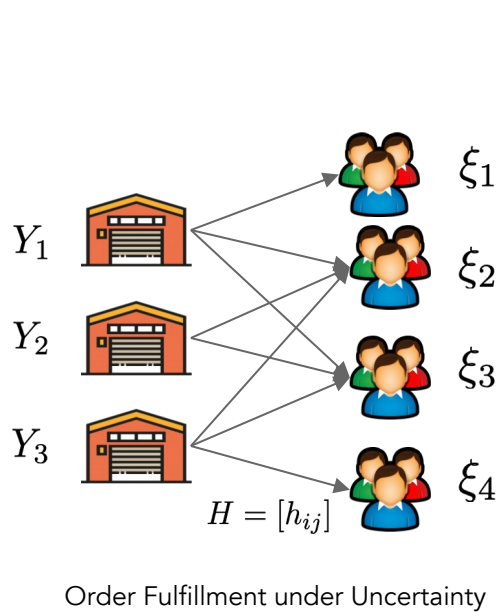
Homepage: <https://bo-dai.github.io>

Teaching:


CSE6243 Advanced Machine Learning (Monday/Wednesday 5:00-06:15 pm)

<https://bo-dai.github.io/CSE6243-fall2024/>


Decision-Making Problem is Everywhere



Searching/Decoding



Robotics




Drug Synthesis

Practical decision-making and planning algorithms with
computational and *statistical* efficiency

Foundation Model for Decision Making

Foundation Models methods to complete

All tasks with Single interface



Sim-to-Real Policy

Sim2Real [IROS 2024]

combating Computation and Sample complexity

in decision-making process, aiming for better Agent

Reinforcement Learning, Representation, Generative Models, Optimization and Sampling

Q: An airport has only 2 planes that fly multiple times a day. Each day, the first plane goes to Greece for three-quarters of its flights, and the remaining flights are split equally between France and Germany. The other plane flies exclusively to Poland, and its trips only amount to half the number of trips the first plane makes throughout each day. How many flights to France does the first plane take in one day?

The first plane makes 44 trips to Greece (3/4 of its flights), and the remaining 1/4 is split equally between France and Germany. So, it makes $1/4 * 4 = 1$ trip to France. The second plane makes 44 trips, which is half the number of trips the first plane makes. So, the first plane makes $44 * 2 = 88$ trips in total. Therefore, the first plane takes 1 trip to France in one day. **#### The answer is 1 trip to France. X**

The second plane makes half the number of trips as the first plane, so the first plane makes $44 * 2 = 88$ trips in one day. The first plane goes to Greece for three-quarters of its flights, which means $88 * 3/4 = 66$ flights to Greece. The remaining flights are split equally between France and Germany, so the first plane makes $(88 - 66) / 2 = 11$ flights to France. **#### The answer is 11 flights to France. O**

Adapted Result

Search Path with Bbox-Adapter

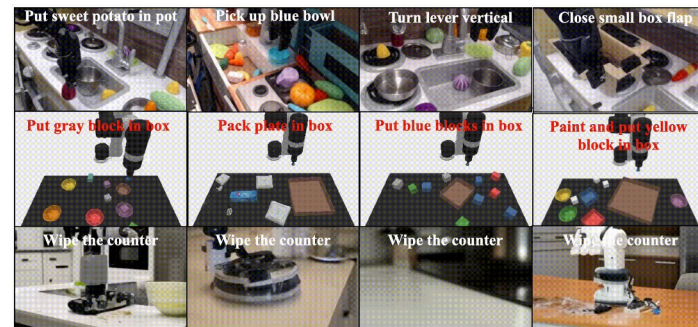
Step 1
Let's first find out how many flights the second plane makes per day.
Let's start by figuring out how many flights the first plane makes in one day.
The second plane makes half the number of trips as the first plane, so the first plane makes $44 * 2 = 88$ trips in one day.

Step 2
There are a total of $44 * 2 = 88$ flights made by the first plane in one day.
The first plane goes to Greece for three-quarters of its flights, which means $88 * 3/4 = 66$ flights to Greece.
If the second plane's 44 trips amount to half the number of trips the first plane makes, the first plane makes $44 * 2 = 88$ trips per day.

Step 3
The remaining flights are split equally between France and Germany, so the first plane makes $(88 - 66) / 2 = 11$ flights to France.
The remaining $88 - 66 = 22$ flights are split equally between France and Germany.
The remaining flights are split equally between France and Germany, so there are $(88 - 66) / 2 = 11$ flights to France.

Step 4
The answer is 11 flights to France. O **#### The answer is 11. O** **#### The answer is 11. O**

LLM Planning & Reasoning [ICML 2023 & 2024]



Uni-Policy [NeurIPS 2023]

Teaching Assistant



Dmitry Shribak

Address: Coda East Wing

Email: shribak@gatech.edu

Logistics

Time: Monday/Wednesday 5:00-06:15 pm

Location: [Molecular Sciences and Engr 1222](#)

Discussion & HW submission: [Ed Discussion & Canvas](#)

Office Hour:

- Instructor: TBD
- TA: TBD

Prerequisite

- Graduate-level Machine Learning
 - Deep neural networks
 - Graphical models
 - Kernel methods...
- Probability and Statistics
 - Random variable, moment generating function
 - Bootstrap, delta method
 - MCMC sampling ...
- Numerical Linear Algebra & Optimization
 - Eigen decomposition, Singular value decomposition
 - Gram-Schmidt process
 - Convex function, duality...

Outline

Providing a unified view for different machine learning methods:

- Module I: Background Knowledge
- Module II: Generative Model
- Module III: Representation Learning
- Module IV: Reinforcement Learning

Outline

Providing a unified view for different machine learning methods:

- Module I: Background Knowledge
- Module II: Generative Model
- Module III: Representation Learning
- Module IV: Reinforcement Learning

WhiteBoard Derivation

Outline

- Module I: Background Knowledge
- Module II: Generative Model
- Module III: Representation Learning
- Module IV: Reinforcement Learning

Guest lectures from prestigious researchers

Reinforcement Learning, and Foundation Models

<https://bo-dai.github.io/CSE6243-fall2024/>

Textbooks

- Boyd & Vandenberghe. [Convex Optimization](#). Cambridge University Press. 2003
- Bishop. [Pattern Recognition and Machine Learning](#). Springer. 2006
- Mohri, Rostamizadeh, & Talwalkar. [Foundations of Machine Learning](#). MIT Press. 2018
- Putman. [Markov Decision Processes: Discrete Stochastic Dynamic Programming](#). John Wiley & Sons, Inc. 1994

The material of the class may go beyond these books

Module I: Background Knowledge

- Convex Optimization
 - Convex function
 - Duality
 - Stochastic gradient descent
- Probabilistic graphical model
 - Directed graphical models (Bayes Nets)
 - Undirected graphical models (Markov Random Fields)
- Sampling
 - Metropolis–Hastings algorithm
 - Gibbs sampling
 - Hamiltonian Monte-Carlo
- Revisit Neural Network

Module II: Generative Model

- Variational auto-encoder
- Autoregressive model
- Generative adversarial net
- Energy-based model
- Diffusion models

Module II: Generative Model AI

- Variational auto-encoder
- Autoregressive model : [ChatGPT](#), [Gemini](#), [Claude](#).....
- Generative adversarial net
- Energy-based model
- Diffusion models : Midjourney, Stability AI, Imagen, Pika....

Module III: Representation Learning

- Representation Learning from EBM view
- Representation Learning from Spectral Decomposition view

Module IV: Reinforcement Learning

- Markov decision process
- Approximate dynamic programming
- Policy gradient
- Imitation learning
- (Offline RL, Exploration)...

Grading

- Participation (20%)
- Scribe Duties (40%)
- Final Project (40%)

Grading

- Participation (20%)
 - In-Class quiz 10%
 - Completing mid-semester evaluation 4%.
 - Machine Learning seminar 6%

This is an in person class, no zoom link, except the guest lectures.

Grading

- Scribe Duties (40%)
 - 2 students as a group
 - 24-26 lectures scribing with [template](#)
 - Submitted in 1 week on Canvas
 - Scribing slots

<https://docs.google.com/spreadsheets/d/11bQt7aMhygQ8gyrSXj8DZb3rrcy61k-KJSSynCQMRT8/edit?usp=sharing>

Grading

- Final Project (40%)
 - 2-4 students as a group
 - **Proposal** : 2 pages excluding references (10%)
 - **Midway Report** : 3 pages excluding references (20%)
 - **Presentation** : oral presentation (20%)
 - **Final Report** : 5 pages excluding references (50%)
 - All write-ups should use the [NeurIPS style](#)

More details: <https://bo-dai.github.io/CSE6243-fall2024/project/>

What is Machine Learning?

Machine learning (ML) is an [umbrella term](#) for solving problems for which development of algorithms by human programmers would be cost-prohibitive, and instead the problems are solved by helping machines 'discover' their 'own' algorithms,^[1] without needing to be explicitly told what to do by any human-developed algorithms.

-- Wikipedia

Machine learning is a branch of artificial intelligence (AI) that focuses on developing computer systems that can learn and adapt without explicit programming. Instead of following rigid rules, these systems learn from data and improve their performance over time.

-- Gemini



Machine learning (ML) is a subset of artificial intelligence (AI) focused on the development of algorithms and statistical models that allow computers to learn from and make predictions or decisions based on data. Instead of being explicitly programmed to perform specific tasks, machine learning systems use data to identify patterns and improve their performance over time.

ChatGPT

Personal Opinion

- Machine Learning is a subfield of AI
- Machine Learning focuses on a special type of algorithm design
 - These algorithms consume data, generates a model for prediction and decision

Top 50 Keywords after Lemmatization




Machine Learning Paradigms

Supervised Learning $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n \in \mathcal{X} \times \mathcal{Y}$ $\text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Y}$

Unsupervised Learning $\mathcal{D} = \{x_i\}_{i=1}^n \in \mathcal{X}$ $\text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Z}$

Reinforcement Learning

Machine Learning Paradigms

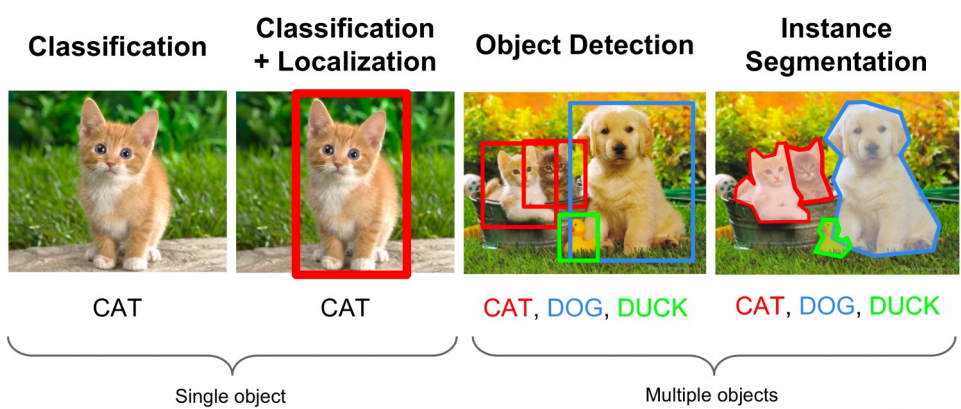



Supervised Learning

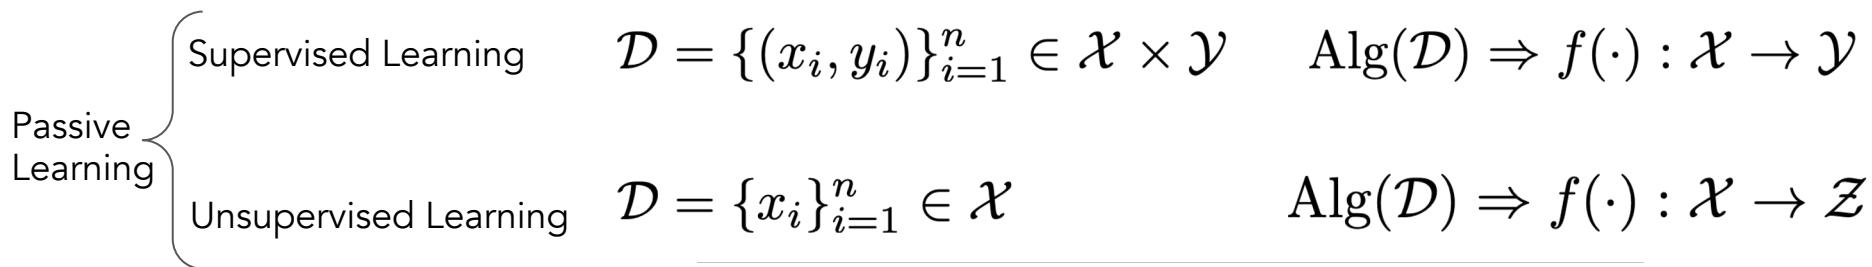
$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n \in \mathcal{X} \times \mathcal{Y} \quad \text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Y}$$

Unsupervised Learning

$$\mathcal{D} = \{x_i\}_{i=1}^n \in \mathcal{X} \quad \text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Z}$$



Machine Learning Paradigms





Reinforcement Learning with Online Interactions

Reinforcement Learning



$$\text{Alg}(\text{Env}) \Rightarrow (\mathcal{D} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^T, \pi(\cdot|s) : \mathcal{S} \rightarrow \Delta(\mathcal{A}))$$

Machine Learning Paradigms



Reinforcement Learning with Online Interactions

Reinforcement Learning



$$\text{Alg}(\text{Env}) \Rightarrow (\mathcal{D} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^T, \pi(\cdot|s) : \mathcal{S} \rightarrow \Delta(\mathcal{A}))$$

Machine Learning Paradigms

Supervised Learning

Unsupervised Learning

Reinforcement Learning



$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n \in \mathcal{X} \times \mathcal{Y} \quad \mathcal{D} = \{x_i\}_{i=1}^n \in \mathcal{X}$$

Semi Supervised Learning

Machine Learning Paradigms

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Offline Reinforcement Learning



$$\mathcal{D} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^T$$

$$\text{Alg}(\mathcal{D}) \Rightarrow \pi(\cdot|s) : \mathcal{S} \rightarrow \Delta(\mathcal{A})$$

Machine Learning Paradigms

Graph Learning

Passive Learning


Supervised Learning

Unsupervised Learning


Reinforcement Learning




Social networks
(Advertisement)




Drug/Material molecules
(Chemistry)



Brain connectivity
(Neuroscience)



Words relationships
(NLP)



Gene Regulatory Network




Recommender systems (Amazon, Netflix)




3D Meshes
(Computer Graphics)




Transportation networks



Knowledge graph
(Causality)



Neutrino detection (High-energy Physics)



Graphs/
Networks

Machine Learning Paradigms

Module I
Basic
Knowledge

Supervised Learning

Unsupervised Learning

Reinforcement Learning

Module II
Generative
Models

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n \in \mathcal{X} \times \mathcal{Y}$$

$$\mathcal{D} = \{x_i\}_{i=1}^n \in \mathcal{X}$$

Reinforcement Learning with Online Interactions



Module III
Differentiable
Programming



$$\text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Y}$$

$$\text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Z}$$



$$\text{Alg}(\text{Env}) \Rightarrow (\mathcal{D} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^T, \pi(\cdot|s) : \mathcal{S} \rightarrow \Delta(\mathcal{A}))$$

Top 50 Keywords after Lemmatization



Machine Learning Paradigms

Module I
Basic
Knowledge

Supervised Learning

Unsupervised Learning

Module IV
Reinforcement Learning

Module II
Generative
Models

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n \in \mathcal{X} \times \mathcal{Y}$$

$$\mathcal{D} = \{x_i\}_{i=1}^n \in \mathcal{X}$$

Module III
Representation
Learning

$$\text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Y}$$

$$\text{Alg}(\mathcal{D}) \Rightarrow f(\cdot) : \mathcal{X} \rightarrow \mathcal{Z}$$

Reinforcement Learning with Online Interactions



$$\text{Alg}(\text{Env}) \Rightarrow (\mathcal{D} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^T, \pi(\cdot|s) : \mathcal{S} \rightarrow \Delta(\mathcal{A}))$$

Tentative Schedule

<https://bo-dai.github.io/CSE6243-fall2024/lectures/>

Q&A