00103335: Deep Learning and Reinforcement Learning
Fall 2024

Lectures: Wednesdays 3:10–5:00 pm and odd Mondays 1:00–2:50 pm, 519 Classroom Building 2

Instructor: Wei Lin (weilin@math.pku.edu.cn)
Office hours: Fridays 10:00–11:30 am, 346 Zhihua Building

Teaching Assistant: Yongqi Jin (yongqijin@stu.pku.edu.cn)
Office hours: Tuesdays 7:00-9:00 pm, 102 Building 21

Course Description:
As highly successful and widely applied machine learning methods, deep learning and reinforcement learning are the core techniques underlying the latest major breakthroughs in the field of AI. Building on the general principles and methodology of machine learning and motivated by important practical problems, this course will introduce the basic concepts and methods, mathematical foundations and theory, optimization algorithms, and applications and case studies of deep learning and reinforcement learning. The part on deep learning will cover feedforward neural networks, regularization and optimization for deep learning, convolutional neural networks, recurrent neural networks, and autoencoders and generative models; the part on reinforcement learning will cover multi-armed bandits, Markov decision processes, dynamic programming, Monte Carlo methods, temporal difference learning, and deep reinforcement learning.

Primary Texts:
1. Mohri, M., Rostamizadeh, A. and Talwalkar, A. (2018). *Foundations of Machine Learning* (2nd ed.). MIT Press.
2. Goodfellow, I., Bengio, Y. and Courville, A. (2016). *Deep Learning*. MIT Press.
3. Prince, S. J. D. (2023). *Understanding Deep Learning*. MIT Press.
4. Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.

Supplementary Texts:
1. Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
2. Bach, F. (2024). *Learning Theory from First Principles*. MIT Press.
3. Stevens, E., Antiga, L. and Viehmann, T. (2020). *Deep Learning with PyTorch*. Manning.
4. Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley.
5. Powell, W. B. (2022). *Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions*. Wiley.

Homework:
There will be about five homework assignments due in class. If you missed the class, contact the TA to turn in your homework by the end of the day. No late homework will be accepted.

Exams:
There will be no midterm or final exam, but students are required to form groups of two to four and complete a final project on a topic of their own choice. Proposals of project topics are due by Wednesday, November 27. Oral presentations of projects are scheduled on Wednesdays, December 18 and 25.

Grading:
The course grade breaks down as follows: homework 50% and project 50%.

Website:
Lecture topics and homework assignments will be posted on the course website at http://www.math.pku.edu.cn/teachers/linw/3335f24.html.