Introduction to AI602: Recent Advances in Deep Learning

Jinwoo Shin

KAIST AI

Course Information

- Goal: Cover a very partial subset of recent advances in deep learning

- Course homepage: https://alinlab.kaist.ac.kr/ai602_2023.html
  - Slides are made by students in Algorithmic Intelligence Laboratory
  - Reference papers will be uploaded for each class (we have no textbook)

- Zoom link for the class
  - https://kaist.zoom.us/j/87239786758?pwd=eUZZWGN3OWZPL1Z0eXFZNEx0K2dOZz09
- Office hours: Every Monday, 10:15am-11am, after the class (on demand)

Instructor and TAs

- Instructor: Jinwoo Shin
  - Professor, KAIST AI and EE
  - Email: jinwoos@kaist.ac.kr

- TA
  - Kyungmin Lee, kyungmnlee@kaist.ac.kr

Prerequisites

- How much backgrounds do I need?
  - This course is not an introductory course to deep learning
  - I will cover some backgrounds quickly, but not spend too much time
  - For example, I will not teach how to use TensorFlow or PyTorch
  - Very sorry, but if you worry about this, please drop the class

- For example, I assume all students know the following concepts
  - Supervised, unsupervised and reinforcement learning
  - Popular neural architectures (e.g., RNN, CNN, LSTM, GNN, ResNet, Transformer)
  - Stochastic gradient descent
  - Batch normalization
  - Overfitting, underfitting and regularization
  - Reparameterization tricks
  - Popular generative models (e.g., GAN, VAE)

(Tentative) Schedule

- Each Lecture X (X>0) would take a day (or often two or more days)
  - Between lectures, there would be paper presentations by students

- Lecture 0: Introduction to AI602
- Lecture 1: Recent neural architectures for vision I: Discriminative models
- Lecture 2: Recent neural architectures for vision II: Generative models
- Lecture 3: Recent neural architectures for language
- Lecture 4: Foundation models for language
- Lecture 5: Foundation models for vision
- Lecture 6: ML Safety and out-of-distribution detection
- Lecture 7: Neural fields and high-dimensional vision

Assignments: 1 Presentation + 1 Report

- We will provide a list of papers in a Google Sheet by **09/01**.
  - You have to choose a paper
  - The paper is used for your presentation and report
  - You cannot choose a paper chosen by another student (first-come-first-serve)
  - If you do not choose your paper until **09/04**, you will be assigned to a random paper.

- Presentation (free format)
  - Present the paper's contents, e.g., motivation, problem, contribution, method, experiments, etc.
  - Your talk would be around 10-15 minutes, i.e., 10-20 slides.
  - You do not need to include your own experimental results
  - Presentation schedules will be announced on **09/06.**

Assignments: 1 Presentation + 1 Report

- Report (free format, e.g., use NeurIPS or CVPR format)
  - Try to reproduce some results of the paper
  - Try to criticize the weakness of the paper.
  - Try to improve the results of the paper
  - Due is on **12/15** (send your pdf to TA)

- How to criticize the paper?
  - You can criticize the paper upon your reproduced results
  - You can criticize the method fails in a different setup/problem, e.g., if some assumption does not hold
  - You can criticize the method in a way that it is suboptimal, i.e., there is a better method for the same problem

- How to improve the paper?
  - Try to resolve one of criticisms you found by your own idea, with supporting experimental results
  - At least, you can find better hyper-parameters to improve the results

Grading Policy

- Presentation 20% + Report 60% + Attendance 20%
  - You will be graded by the absolute scores, and not by the relative rankings. You will not compete with anybody.
  - You should attend at least 70% of classes (otherwise, 0 credit for attendance).
  - For the attendance criteria, one is considered as "attended" if his/her zoom access log is more than 50 minutes and your video is on (for showing your face). Otherwise, it is considered as "absent".
  - When you enter in Zoom session, please set your Zoom-name as "[student ID] [Name]" (e.g., 20217018 Junsu Kim)