

Introduction to AI602: Recent Advances in Deep Learning

Jinwoo Shin

KAIST AI & EE

Course Information

- Goal: Cover a very partial subset of recent advances in deep learning
- Course homepage: http://alinlab.kaist.ac.kr/ai602_2022.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/2947126282?pwd=KzFIVmhkM1h2dCtUbGVwa1puektzdz09>
 - ID: 294 712 6282
 - pwd: G0tVa5
- Office hours: Every Monday, 10:15am-11am, after the class (on demand)

Instructor and TAs

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

- TAs
 - Junsu Kim, junsu.kim@kaist.ac.kr
 - Seong Hyeon Park, seonghyp@kaist.ac.kr
 - Seojin Kim, osikjs@kaist.ac.kr
 - Minkyu Kim, kimmk135@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)
 - Stochastic gradient descent
 - Batch normalization
 - Overfitting, underfitting and regularization
 - Reparameterization tricks
 - Popular generative models (e.g., GAN, VAE)

Schedule

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

Schedule

- Lecture 0: Introduction to AI602
- Lecture 1: Advanced optimizers for neural networks
- Lecture 2: Advanced deep spatial models (e.g., for image recognition)
- Lecture 3: Advanced deep temporal models (e.g., for language modeling)
- Lecture 4: Advanced deep spatial-temporal models (e.g., for video recognition)
- Lecture 5: Advanced deep generative models I: Implicit density models
- Lecture 6: Advanced deep generative models II: Explicit density models
- Lecture 7: Self-supervised representation learning
- Lecture 8: Neural fields
- Lecture 9: Transfer and continual learning
- Lecture 10: Meta learning
- Lecture 11: Robustness
- Lecture 12: Fairness and bias in AI

Assignments: 1 Presentation + 1 Report

- We will provide a list of papers in a Google Sheet by **3/1**.
 - 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 **3/4**, 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 **3/6**

Assignments: 1 Presentation + 1 Report

- Report (free format, e.g., use NeurIPS 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 **6/13** (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