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_2021.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://zoom.us/j/2947126282?pwd=KzFlVmhkM1h2dCtUbGVwa1puektzdz09
  - ID: 294 712 6282
  - pwd: G0tVa5

- Office hours: Every Monday, 12pm-2pm, N1-914

Instructor and TAs

- Instructor: Jinwoo Shin (N1-914)
  - Associate professor, KAIST AI and EE
  - Email: jinwoos@kaist.ac.kr

- TAs (N1-917)
  - Hyuntak Cha hyuntak.cha@kaist.ac.kr
  - Junsu Kim junsu.kim@kaist.ac.kr
  - Seunghyun Lee seunghyun.lee@kaist.ac.kr
  - Sumin Shin sym807@kaist.ac.kr
  - Jihoon Tack jihoontack@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)

  - 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: Paper presentation for Lecture 1
  - Lecture 5: Advanced deep generative models I (e.g., for image generation)
  - Lecture 6: Paper presentation for Lecture 2
  - Lecture 7: Advanced deep generative models II (e.g., for density estimation)
  - Lecture 8: Paper presentation for Lecture 3
  - Lecture 9: Novelty and uncertainty estimation
  - Lecture 10: Paper presentation for Lecture 5
  - Lecture 11: Unsupervised and self-supervised representation learning
  - Lecture 12: Paper presentation for Lecture 7
  - Lecture 13: Transfer and continual learning
  - Lecture 14: Paper presentation for Lecture 9
  - Lecture 15: Few-shot learning
  - Lecture 16: Paper presentation for Lecture 11
  - Lecture 17: Network_compression
  - Lecture 18: Paper presentation for Lecture 13
  - Lecture 19: Adversarial robustness
  - Lecture 20: Paper presentation for Lecture 15
  - Lecture 21: Interpretable and explainable learning
  - Lecture 22: Paper presentation for Lecture 17
  - Lecture 23: Paper presentation for Lecture 19
  - Lecture 24: Paper presentation for Lecture 21

Assignments

- Assignment: 1 presentation + 1 report

- We will no longer receive emails regarding presentation topic assignment.
- Instead, we will move on to Google Sheets.
  - If you have not chosen your topic yet, please do so (due **3/12**)
  - **Update**: you may now freely choose **any topic** by **3/12**; Otherwise, you will be assigned to a random topic.
  - Once you have chosen the topic, you must choose the paper you want to present on (due one week before your presentation)
    - You must choose a paper published in 2019 - 2021
    - You cannot choose a paper chosen by another student (first-come-first-serve)

- Presentation schedules will be announced on **3/14**

- FAQ
  - Q. Is it enough to present on the paper's content only, or do we need to include our own experiment results?
  - A. You only have to present on the paper's content only.