

Course Project

Ju Sun

Computer Science & Engineering

University of Minnesota, Twin Cities

February 6, 2020

Logistics

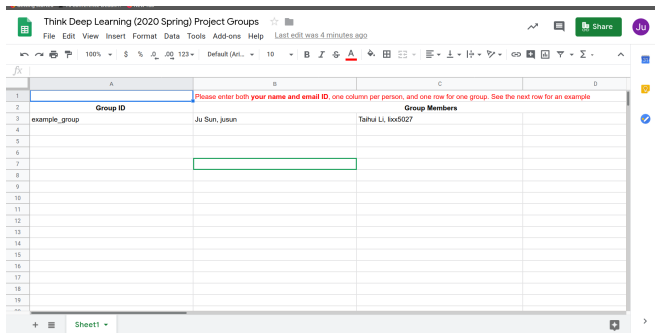
Project ideas

- Proposal (5%, 1 page): Feb 16th
- Progress presentation (5%, 2-3 mins): Mar 26th
- Progress report (5%, 2 pages): Mar 28th
- Final report (25%, 6 – 8 pages): May 12th
- Poster presentation?
- **Publisable results \implies A!**

Template for all writeups: NeurIPS 2019 style

<https://neurips.cc/Conferences/2019/PaperInformation/StyleFiles>

Groups



- Each group: 2 or 3 students; 1 permitted but discouraged
- All submissions as a group (in Canvas as group assignment); the group gets the same score

Proposal

- What problem?
- Why interesting?
- Previous work
- Your goal
- Plan and milestones

Logistics

Project ideas

Roughly by ascending level of difficulty

- Literature survey/review
- Novel applications
- Novel methods
- Novel theories

Excerpt from a research project is fine, but you should describe your own contributions

A coherent account of recent papers in a focused topic

- Description and comparison of main ideas, or
- Implementation and comparison of performance, or
- Both of the above

should complement the topics we cover in the course


Random topics

- DL for noneuclidean data (e.g., graph NN, manifold NN)
- transformer models for sequential data
- generative models (e.g., GAN, VAE)
- 2nd order methods for deep learning
- differential programming
- universal approximation theorems
- DL for 3D reconstruction
- DL for video understanding and analysis
- DL for solving PDEs
- DL for games
- RL for robotics
- adversarial attacks; robustness of DL
- privacy, fairness in DL
- visualization for DNN
- network quantization and compression
- hardware/software platforms for DL
- automated ML; architecture search
- optimization/generalization theory of DL

Novel applications

Apply DL to **new** application problems

- A good place to start: Kaggle <https://www.kaggle.com/>



- Think about data availability

Google dataset search

<https://datasetsearch.research.google.com/>

- Think about GPUs

Where to find inspirations

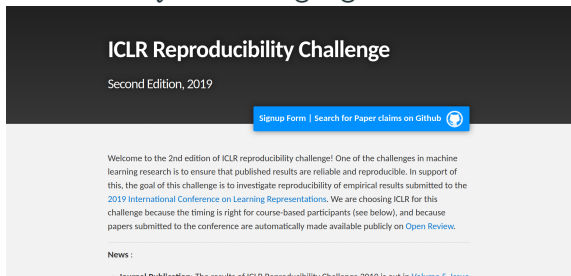
- arXiv machine learning
<https://arxiv.org/list/cs.LG/recent>
- Recent conference papers
 - ML: NeurIPS, ICML, ICLR, etc
 - CV: ICCV, ECCV, CVPR, etc
 - NLP: ACL, EMNLP, etc
 - Robotics: ICRA, etc
 - Graphics: SIGGRAPH, etc
- Talk to researchers!

Novel methods

Create new **NN models or training algorithms** to improve the state-of-the-art


Where to start:

- Kaggle (again)!
- arXiv machine learning and recent conference papers
- ICLR reproducibility challenge: https://reproducibility-challenge.github.io/iclr_2019/



Novel methods

Equally interesting to fool/fail the state-of-the-art, i.e., exploring robustness of DL



Credit: ImageNet-C

Novel theories

Nothing is more practical than a good theory. – V. Vapnik

- universal approximation theorems
- nonconvex optimization
- generalization

Where to start:

- Analyses of Deep Learning (Stanford, fall 2019)
<https://stats385.github.io/>
- Theories of Deep Learning (Stanford, fall 2017)
https://stats385.github.io/stats385_2017.github.io/
- Toward theoretical understanding of deep learning (ICML 2018 Tutorial)
<https://unsupervised.cs.princeton.edu/deeplearningtutorial.html>

Questions?