Course Project

Ju Sun
Computer Science & Engineering
University of Minnesota, Twin Cities

October 18, 2023

Timeline & LaTeX template

- Teaming up: Oct 29
  https://docs.google.com/spreadsheets/d/1dKLKW7dailnLtcrTu9Cyn1lZeute97QvuLYJM5yV6oM/edit?usp=sharing
- Proposal (5%, 1–2 pages): Nov 05
- Progress lightning talks (5%, 5 mins): Nov 22 (or Nov 21)?
- Progress report (5%, 3–4 pages): Dec 03
- Final report (25%, 7–8 pages): Dec 20

All page counts exclude references

Template for all writeups: ICLR 2024 LaTeX style
https://github.com/ICLR/Master-Template/raw/master/iclr2024.zip
Add `\iclrfinalcopy` to the LaTeX preamble to make your names visible

Groups

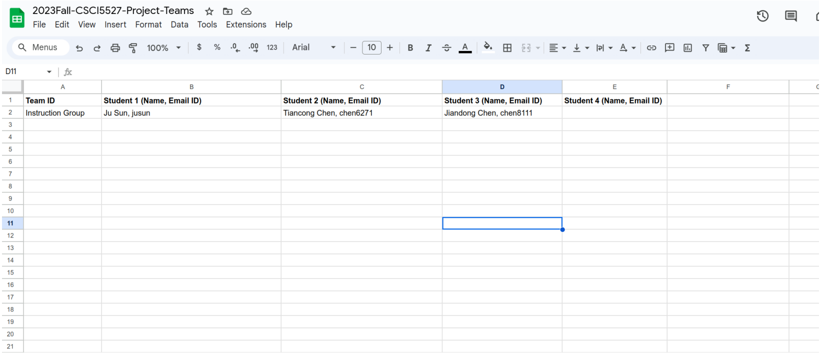
2023Fall-CSCI5527-Project-Teams

Team ID	Student 1 (Name, Email ID)	Student 2 (Name, Email ID)	Student 3 (Name, Email ID)	Student 4 (Name, Email ID)
Instruction Group	Ju Sun, jusun	Tiancong Chen, chen6271	Jiandong Chen, chen8111	

– All submissions as a group (in Canvas as group assignment); the group gets the same score

Computing resources

- Prototyping
  * Google Colab https://colab.research.google.com/
  * Local installation of Jupyter Notebook https://jupyter.org/
  * MSI notebook notebooks.msi.umn.edu (https://www.msi.umn.edu/support/faq/how-do-i-get-started-jupyter-notebooks)
- Large-scale jobs: submit them to MSI GPU queues
  * MSI quick start https://www.msi.umn.edu/quick-start-guides
  * Slurm scheduler tutorial https://www.msi.umn.edu/slurm

Proposal

Five necessary components

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

Outline

Overview

Roughly by ascending level of difficulty

- Literature survey/review (least favorable given the good summarization capabilities of AI tools nowadays)
- Novel applications
- Novel methods
- Novel theories

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

Literature survey/review

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

https://paperswithcode.com/rc2022

Random topics

- DL for noneuclidean data (e.g., graph NN, manifold NN)
- transformer models for sequential data
- generative models (e.g., GAN, VAE, normalization flow, diffusion models)
- 2nd order methods for deep learning
- constrained optimization for deep learning
- differential programming
- universal approximation theorems
- DL for 3D reconstruction
- DL for video understanding and analysis
- DL for solving PDEs
- DL for material discovery
- DL for inverse problems
- RL for games
- RL for robotics
- DL for medical imaging
- DL for (astro)physics
- DL for chemistry
- 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
- large vision-language models

Novel applications

Apply DL to **new** application problems

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

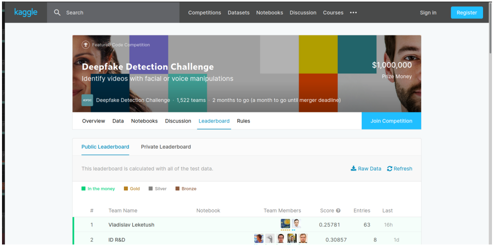

- Think about data availability

  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
- MLRC

https://paperswithcode.com/rc2020

Novel methods

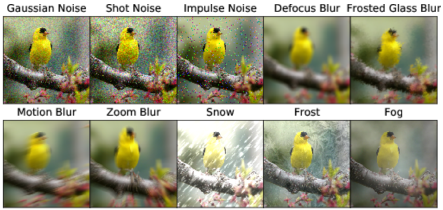

Credit: ImageNet-C https://github.com/hendrycks/robustness

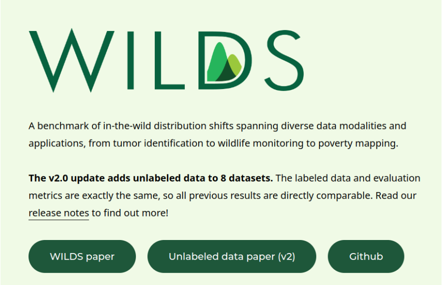

Credit: WILDS https://wilds.stanford.edu/

Equally interesting to fool/fail the state-of-the-art, e.g., exploring robustness of DL, finding common limitations of state-of-the-art

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
- https://sunju.org/teach/TMML-Fall-2021/

Questions?