Course Project

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

October 18, 2022

Outline

Announcements

– HW2 is out. Due: Oct 30th

Outline

Timeline & LaTeX template

- Proposal (5%, 1–2 page): Oct 28
- Progress lightning talks (5%, 5 mins): Nov 29
- Progress report (10%, 3–4 pages): Dec 02
- Final report (20%, 7–8 pages): Dec 20

All page counts exclude references

Template for all writeups: ICLR 2023 LaTeX style
https://github.com/ICLR/Master-Template/raw/master/iclr2023.zip
Add `\iclrfinalcopy` to the LaTeX preamble to make your names visible

Groups

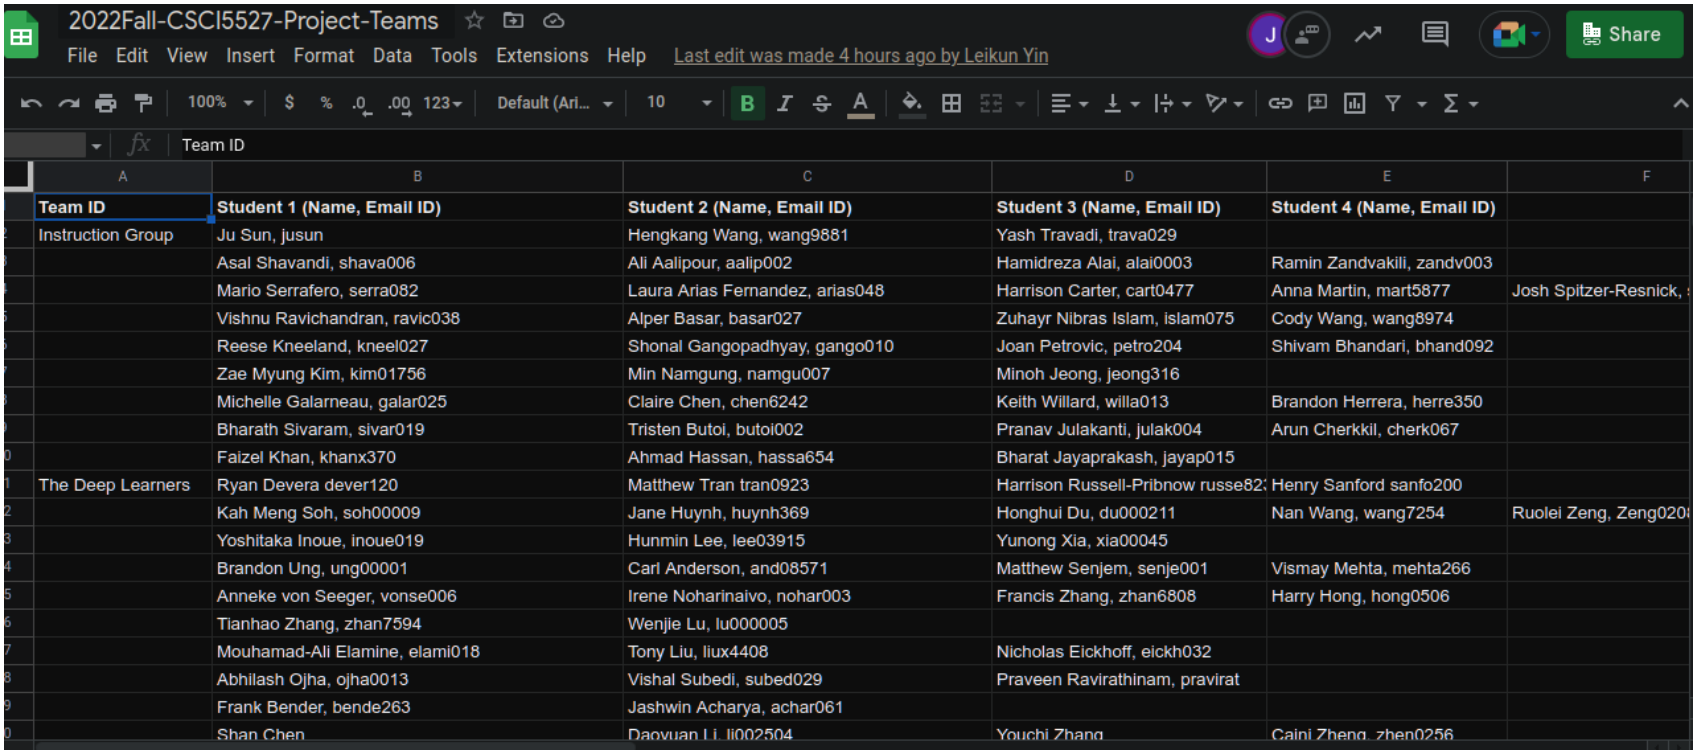

2022Fall-CSCI5527-Project-Teams

File Edit View Insert Format Data Tools Extensions Help Last edit was made 4 hours ago by Leikun Yin

Share

Team ID

	A	B	C	D	E	F
1	**Team ID**	**Student 1 (Name, Email ID)**	**Student 2 (Name, Email ID)**	**Student 3 (Name, Email ID)**	**Student 4 (Name, Email ID)**	
2	Instruction Group	Ju Sun, jusun	Hengkang Wang, wang9881	Yash Travadi, trava029		
3		Asal Shavandi, shava006	Ali Aalipour, aalip002	Hamidreza Alai, alai0003	Ramin Zandvakili, zandv003	
4		Mario Serrafero, serra082	Laura Arias Fernandez, arias048	Harrison Carter, cart0477	Anna Martin, mart5877	Josh Spitzer-Resnick,
5		Vishnu Ravichandran, ravic038	Alper Basar, basar027	Zuhayr Nibras Islam, islam075	Cody Wang, wang8974	
6		Reese Kneeland, kneel027	Shonal Gangopadhyay, gango010	Joan Petrovic, petro204	Shivam Bhandari, bhand092	
7		Zae Myung Kim, kim01756	Min Namgung, namgu007	Minoh Jeong, jeong316		
8		Michelle Galarneau, galar025	Claire Chen, chen6242	Keith Willard, willa013	Brandon Herrera, herre350	
9		Bharath Sivaram, sivar019	Tristen Butoi, butoi002	Pranav Julakanti, julak004	Arun Cherkkil, cherk067	
10		Faizel Khan, khanx370	Ahmad Hassan, hassa654	Bharat Jayaprakash, jayap015		
11	The Deep Learners	Ryan Devera dever120	Matthew Tran tran0923	Harrison Russell-Pribnow russe82	Henry Sanford sanfo200	
12		Kah Meng Soh, soh00009	Jane Huynh, huynh369	Honghui Du, du000211	Nan Wang, wang7254	Ruolei Zeng, Zeng020
13		Yoshitaka Inoue, inoue019	Hunmin Lee, lee03915	Yunong Xia, xia00045		
14		Brandon Ung, ung00001	Carl Anderson, and08571	Matthew Senjem, senje001	Vismay Mehta, mehta266	
15		Anneke von Seeger, vonse006	Irene Noharinaivo, nohar003	Francis Zhang, zhan6808	Harry Hong, hong0506	
16		Tianhao Zhang, zhan7594	Wenjie Lu, lu000005			
17		Mouhamad-Ali Elamine, elami018	Tony Liu, liux4408	Nicholas Eickhoff, eickh032		
18		Abhilash Ojha, ojha0013	Vishal Subedi, subed029	Praveen Ravirathinam, pravirat		
19		Frank Bender, bende263	Jashwin Acharya, achar061			
20		Shan Chen	Daoyuan Li, li002504	Youchi Zhang	Caini Zheng, zhen0256	

– 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
- 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 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

Novel applications

Apply DL to **new** application problems

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

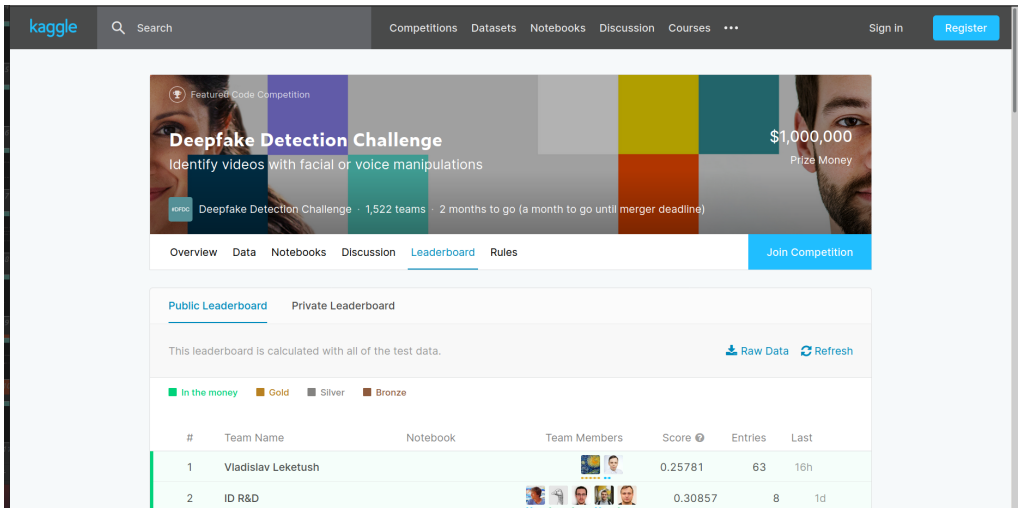

- Think about data availability

  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

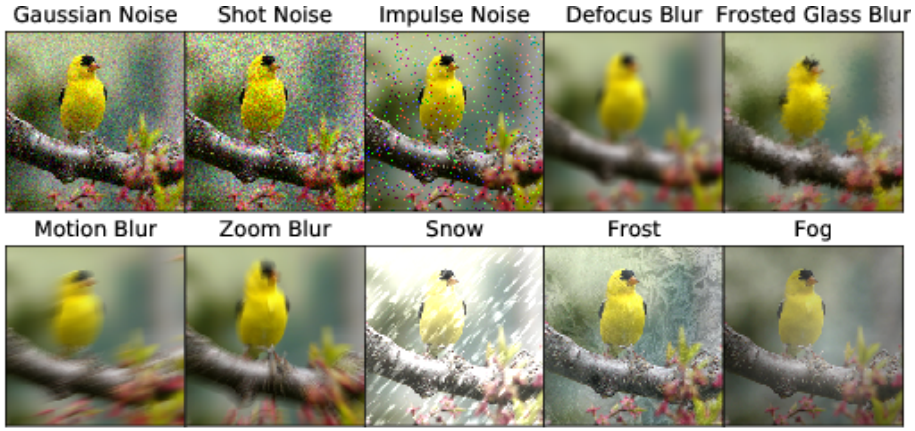

Credit: ImageNet-C https://github.com/hendrycks/robustness

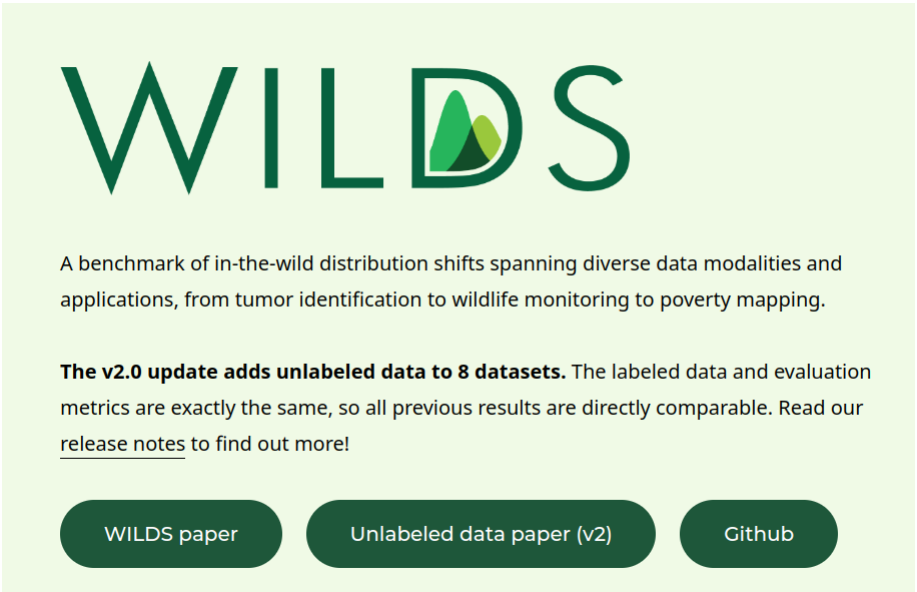

Credit: WILDS https://wilds.stanford.edu/

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

Questions?