Deep Learning for Computer Vision

113-1/Fall 2024

https://cool.ntu.edu.tw/courses/41702 (NTU COOL)

http://vllab.ee.ntu.edu.tw/dlcv.html (Public website)

Yu-Chiang Frank Wang 王鈺強, Professor

Dept. Electrical Engineering, National Taiwan University

2024/09/03

Deep Learning for Computer Vision

- Time:
  - Tuesday 2, 3, 4
  - ~~09:10-12:20~~ → **09:30-12:20** with 1 short break
- Location:
  - BL-112
- Websites:

  - NTU COOL; primary website: https://cool.ntu.edu.tw/courses/41702

  - Public use: http://vllab.ee.ntu.edu.tw/dlcv.html

Deep Learning for Computer Vision (cont'd)

- **Required** Knowledge & Skills
  - Knowledge of linear algebra, vector calculus, and probability
  - Machine learning & deep learning (this is **NOT** an entry-level course)
  - Programming skills (in Python)

  - *If any of the above skills not available, please come back when ready.*

- This course is offered in **Mandarin** this semester.
  Q&As & discussions can be done in any languages communicable.

- To encourage class attendance, lectures will be recorded
  but **NOT** uploaded to COOL until a week after.

What to Expect from this Course?

- (Deep) Learning-Based Computer Vision
  - ~~Fundamentals of machine learning~~ (Sorry, it's required!)
  - Deep learning technologies for visual classification, synthesis, and beyond (e.g., LLM & foundation models)
  - Different from courses like *Computer Vision, 3D Vision, Applied Deep Learning, etc.*

- Practical Experiences
  - Assignments and projects dealing with real-world visual data
  - Final projects associated with challenges from CVPR, ICCV, etc.

- Lots of work with fast pace, but hopefully helpful with lots of fun!

Disclaimer

- You are responsible for your own computing resources.
  (e.g., GPUs, Colab Pro for $11/month, etc.)
  You **can't** finish this course w/o using GPUs.
- While doing my best to stick to what we announce,
  syllabus, HW policy or details might change over time.
- Yes, up to 20% of students chose to drop the course in previous years.
- And, yes, we did fail students in the past years (~5% each semester).

Required Computing Resources

- If you're using your own machines…(update by HW#1)
  - Ubuntu 20.04.6 LTS
  - NVIDIA TITAN RTX TU102 (24GB)
  - GNU bash, ver. 5.0.17(1)-release

Course Information

- Teaching Team & Office Hours
- Course Policy
- How to enroll in this class if not already in?

About Myself

- **Education**

- **Ph.D./M.S. in Electrical & Computer Engineering** 2002 – 2009
  Carnegie Mellon University, Pittsburgh, USA

- **B.S. in Electrical Engineering** 1997 – 2001
  National Taiwan University, Taipei, Taiwan

- **Experiences**

- **AI研究總監**, NVIDIA Research 2022/08 – present
- **Professor**, Dept. EE, National Taiwan University 2019 – present

- **Principal AI Consultant**, Inventec 2021 – 2022

- **AI Advisory Consultant**, ASUS Intel. Cloud Services (AICS) 2019 – 2022
- **Associate Professor**, National Taiwan University 2017 – 2019
- **Deputy Director**, CITI, Academia Sinica 2015 – 2017

- **Associate/Assistant Research Fellow** 2009 – 2017
  Research Center for IT Innovation (CITI), Academia Sinica

- **Office Hour**
  - After class or by appointment via email (ycwang@ntu.edu.tw)

TAs & Office Hours

- We have 8 TAs this semester.
  - Will post and update the TA info on DLCV website
- If any questions, you are encouraged to
  - Ask at NTU COOL (anonymous is fine)
  - Contact your TA **by email at** ntudlcv@gmail.com
- Q&A @ NTU COOL
  - “Discussions” -> “Q&A for Date” -> Ask your questions @ Reply

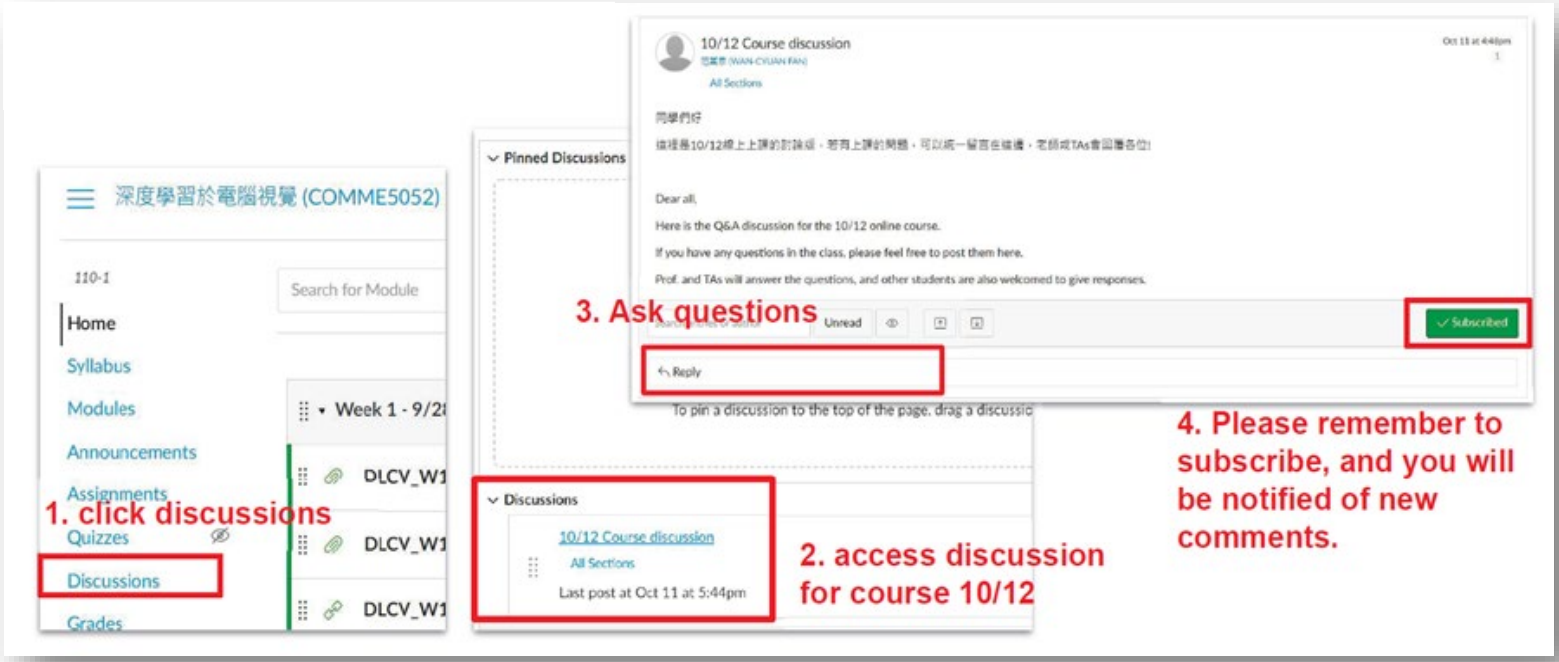

Q&A @ NTU COOL

- “Discussions” -> “Q&A for Date” -> Ask your questions @ Reply

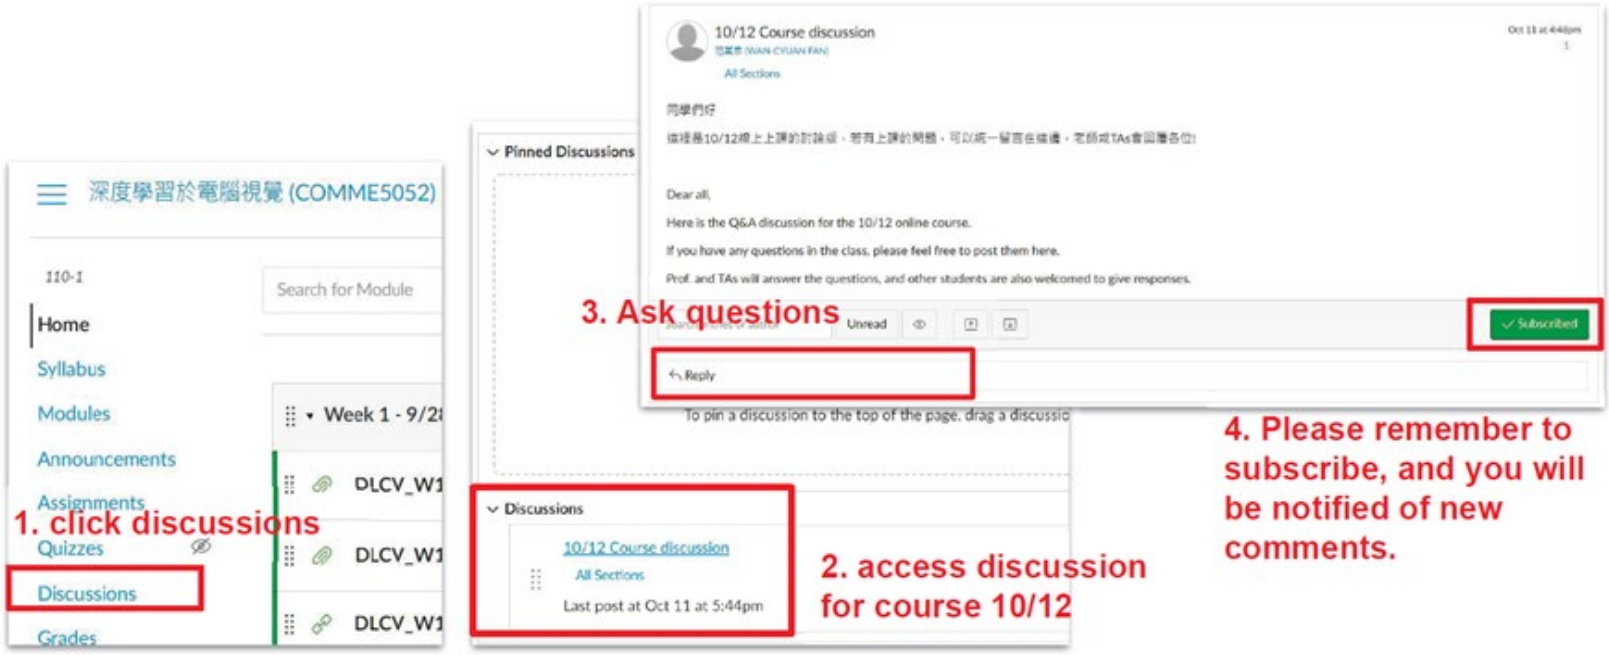

A Very Tight Schedule…

Week	Date	Topic	Course Materials	Remarks
1	09/03	Course Logistics & Registration; Intro to Neural Nets	W1-1	
2	09/10	Convolutional Neural Networks & Image Segmentation		**HW #1 out**
3	09/17	**No class**		Mid-Autumn Festival
4	09/24	Generative Models (I) - AE, VAE & GAN		**HW #1 due**
5	10/01	**Guest Lecture**		ECCV week
6	10/8	Generative Models (II) - Diffusion Model		**HW # 2 out**
7	10/15	Recurrent Neural Networks & Transformer		
8	10/22	Transformer; Vision & Language Models		
9	10/29	Vision & Language Models		**HW #2 due; HW #3 out**
10	11/05	Multi-Modal Learning; Parameter-Efficient Finetuning		
11	11/12	Unlearning, Debiasing, and Interoperability		
12	11/19	3D Vision		**HW #3 due; HW #4 out**
13	11/26	Transfer & Adversarial Learning		**Final Project Announcement**
14	12/03	Federated Learning		**HW #4 due**
15	12/10	**Progress Check for Final Projects**		NeurIPS week
17	12/25 Wed	**Final Project Presentation**		

Textbook (Optional)

- Deep Learning, MIT Press
  - Ian Goodfellow, Yoshua Benjio, and Aaron Courville
  - Free online versions available at www.deeplearningbook.org

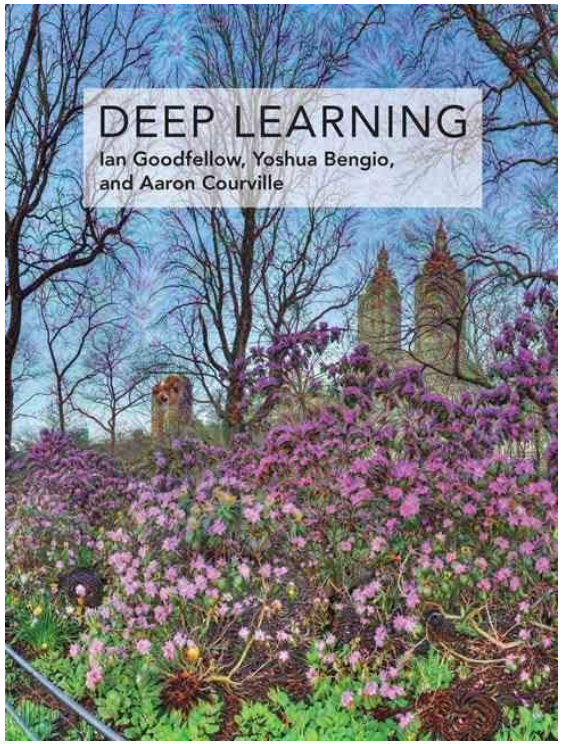

Textbook (Optional)

- Computer Vision: Algorithms and Applications, Springer
  - Richard Szeliski
  - Free online versions available at http://szeliski.org/Book/

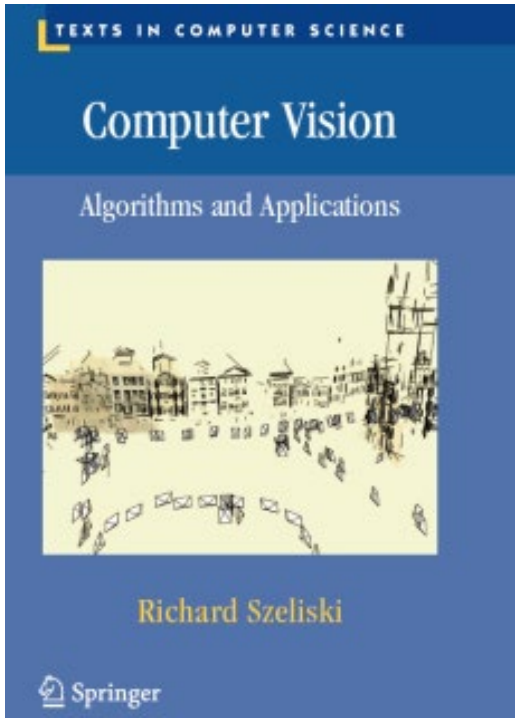

About Grading

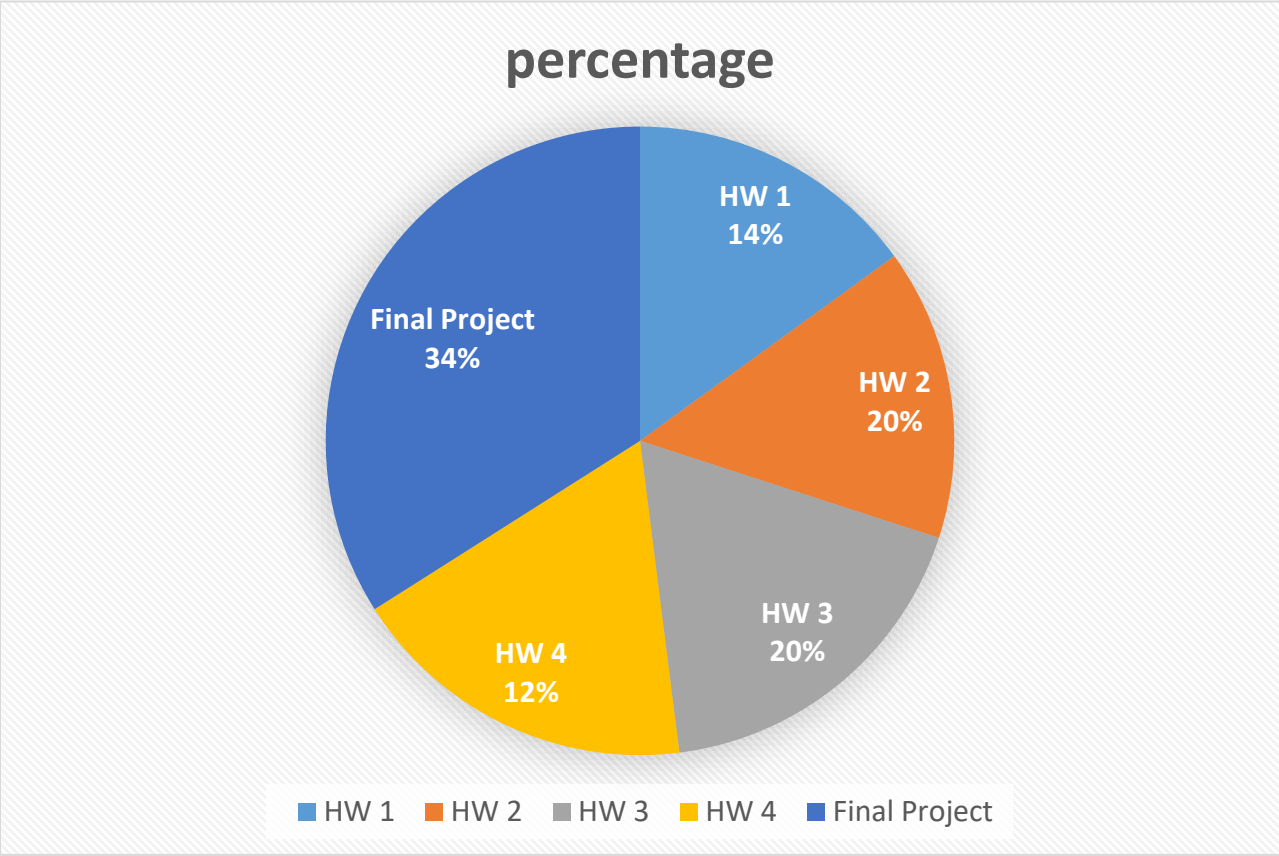

Bonus points possibly available for both HWs & lecture interaction

About Grading (cont'd)

- HW assignments **(66%)**
  - ~~HW 0 (non-DL HW, 0% but required)~~
  - HWs #1~4 (12~20% each with possible bonus points)
  - ~~HW #5 (optional yet with bonus points)~~
- Final project + poster presentation + code, etc. **(34%)**
- Bonus points
  - Course participation (e.g., interaction, Q&A, etc.)
  - Extra challenges in HWs
  - Excellent performance for final project
    (e.g., competitions, publication submissions, etc.)

Final Grade

Letter Grading System	Definition	Grade Points	Conversion Scale
A+	All goals achieved beyond expectation	4.3	90-100
A	All goals achieved	4.0	85-89
A-	All goals achieved, but need some polish	3.7	80-84
B+	Some goals well achieved	3.3	77-79
B	Some goals adequately achieved	3.0	73-76
B- (passing grade for graduate students)	Some goals achieved with minor flaws	2.7	70-72
C+	Minimum goals achieved	2.3	67-69
C	Minimum goals achieved with minor flaws	2.0	63-66
C- (passing grade for undergraduate students)	Minimum goals achieved with major flaws	1.7	60-62
F	Minimum goals not achieved	0	59 and below
X	Not graded due to unexcused absences or other reasons	0	0
W	Withdrawal		
NG	No grade reported		
IP	In progress		
TR	Transfer credit		
EX	Exempted		

http://www.oia.ntu.edu.tw/upload/files/20150616223619.pdf

About Course HWs/Projects

- About HW late policy
  - We offer free late days (up to THREE days)
    in case you have dates, midterms, HW dues for other course, etc.
    (如期中考、專題、社團、約會、找不到人抄)
  - 1 min ~ 23 hr 59 min all count as ONE late day.
  - After HW due day, a penalty of **30%** per day.
  - We'll maximize your final score
    based on HW scores and the late days used.
  - **No** late submission for the final project, obviously.

About Course HWs/Projects

- About Final Project
  - Details to be announced around mid semester.
  - 3~4 people per group (no more than 4)
  - TAs will check your code after presentation (to finalized your scores)
  - Selected topics possibly come with cash prizes.
  - Evaluated by instructor, TAs, and possibly guest judges
  - (Intra/inter-group) peer evaluation will be conducted.
  - Snack/drinks will be provided during final presentation.

Academic Integrity

- Can discuss HW with peers, but DO NOT copy and/or share code
  - Plagiarism is against university policy.
  - Violation in ANY form for HWs & final project would result in F.
  - We gave at least five Fs in previous semesters for the above cases.
- Do not directly use code/results from Internet unless you have permissions.
  - **If not sure, ask!**
  - **If so, do specify in your HW/project.**
- No double dipping!
  - You **CANNOT** use your research work as your final project.
  - However, you are encouraged to extend your previous work.
  - Actually, some students turned their final projects into international publications in the past semesters!

DOs and DONTs for the TAs (& Instructor)

- Do NOT send private messages to the TAs via Facebook, etc.
  - TAs are here to help, but they are not your tutors 24/7.
- TAs will NOT debug for you,
  including coding, environmental, library dependency problems.
- If you cannot make the TA hours, please email & schedule an appointment instead of stopping by the lab directly.
- (Obviously,) TAs do NOT answer questions not related to the course.
- You are encouraged to discuss w/ me if any questions on (but not limited to) DLCV, research, industry, career planning, etc.

How to Sign Up If Not Already In?

- **Capacity**
  - Classroom capacity: 120; currently registered: 100
  - About **20** students can be additionally added.
  - If you do not plan to take the course anymore, please drop it ASAP.
- **Priority**
  - EECS students > students with strong research needs > those are interested in AI/GenAI…
    - Research topic alignment & necessity are key factors
  - If you are still interested in this course and plan to enroll, please fill in the following form Sept 3rd Tue between 9am-1pm: https://forms.gle/M6Grdrm3F2YsBFeFA
  - We will announce the enrollment results via email no later than Sept 6th Fri 6pm. **All decisions are final**.

加簽表單

Any Questions?

Course registration at https://forms.gle/M6Grdrm3F2YsBFeFA
by **Tue Sept. 3rd 1pm**!