UNIVERSITY OF MACAU
FACULTY OF SCIENCE AND TECHNOLOGY
DEPARTMENT OF COMPUTER AND INFORMATION SCIENCE
CISC7414 Advanced Topics in Computer Science
Syllabus
2st Semester 2023/2024
Part A – Course Outline

Elective course in Computer Science

Course Description:
3 credits, 15 weeks. This is an advanced course in Computer Vision, focusing on Visual Recognition in the real open world. Visual Recognition aims at high-level understanding of visual data (e.g., an RGB image) and is indispensable in numerous applications such as autonomous driving, medical / biological image analysis, agriculture, entertainment, and forensics. Understanding complex visual data is challenging, therefore state-of-the-art methods adopt data-driven techniques (e.g., by machine learning and deep learning) with massive visual data. This course will cover multiple fundamental problems of visual recognition: object detection, image classification, and semantic segmentation. It will introduce both classic and data-driven approaches with a focus on practical training in developing such methods through paper reading, and several lightweight course projects.

Course type:
Theoretical

Prerequisites:
- Proficiency in Python programming.
- Linear algebra (e.g., matrix computations, eigenvalues and eigenvectors, singular value decomposition).
- Basic machine learning (e.g., supervised learning, linear regression, logistic regression, support vector machines).
- Basic deep learning (convolutional neural networks, backpropagation, logistic regression).

Preferably access to GPUs (e.g., Google Colab).

Recommended background
- **Proficiency in Python**. All class assignments will be in Python. For those who aren't as familiar with Python, please follow this tutorial. In the tutorial, go over all topics under "Python Tutorial", "Python NumPy". It will also be beneficial to go over all topics under "Machine Learning".
- **Familiar with Jupyter Notebook**. We will use Jupyter Notebook (https://jupyter.org) to implement algorithms, demonstrate experiment qualitative and quantitative results, and write up assignment report. Converting a Jupyter Notebook file to pdf / html / python script is trivial (https://nbconvert.readthedocs.io/en/latest/usage.html).
- **Familiar with PyTorch**. Some coding assignments involving deep learning will use PyTorch, which is an excellent python-based toolbox for machine learning and deep learning. For those who haven't used it before, please refer to its official tutorial. In the tutorial, go over "Introduction" topics, "Learning PyTorch", and "Image and Video".

- **College Level Linear Algebra**. You should be comfortable taking derivatives and understanding matrix vector operations and notation. Go over the "Essence of linear algebra" playlists by "3Blue1Brown" at https://www.youtube.com/c/3blue1brown/playlists
- **Basic Probability and Statistics.** You should be familiar with basics of probabilities, Gaussian distributions, mean, standard deviation, etc. Go over the "Probabilities of probabilities" playlist by "3Blue1Brown" at https://www.youtube.com/c/3blue1brown/playlists
- **Foundations of Machine Learning, Deep Learning, Computer Vision.**

Textbook(s) and other required material:
There are no required textbooks. We will focus on papers (see links in the Course Schedule). The following textbooks may also be useful as references:
- Szeliski. *Computer Vision: Algorithms and Applications, 2nd ed. draft* (free online)
- Goodfellow, Bengio, Courville. *Deep Learning*. (free online)

Course objectives:
The objective of this course is to teach fundamental problems of visual recognition with focus on data-driven approaches, practical implementations, and experimentation. Upon completion of the course, students will:
1. Have a good understanding of fundamental visual recognition problems.
2. Be aware of assumptions, strengths, and weaknesses of popular approaches.
3. Be able to develop data-driven methods to solve real-world problems.

The learning outcomes will be assessed based on course projects.

Topics covered:
- Image classification, object detection, semantic segmentation, deep learning, convolutional neural networks
- long-tailed recognition, open-set recognition, out-of-distribution detection, uncertainty, calibration
- multi-modal learning, multi-modal fusion,
- few-shot learning, continual learning,

Class/laboratory schedule:

Timetabled work in hours per week			**No of teaching weeks**	**Total hours**	**Total credits**	**No/Duration of exam papers**
Lecture	**Tutorial**	**Practice**				
3	**0**	**0**	**14**	**42**	**3**	**n/a**

Student study effort required:

Class contact:	
Lecture	42 hours
Tutorial	0 hours
Other study Effort	
Self-study	20 hours

Course project	48 hours
Total student study effort	110 hours

Student assessment:
Final assessment will be determined on the basis of
* Three course projects: 60%=3*20%
* Final Project: 40%

Course Outline:

Weeks	Topics	Course Work	Homework / project
1	Introduction	python tutorial pytorch tutorial	
2	Image Recognition	Features: [SIFT] [HoG] Methods: [ScSPM] [LeCun et al 1989] [AlexNet] [ResNet]	Project 1 out: install PyTorch and Jupyter; use CIFAR10 and MNIST to implement recognition methods based on PCA, KNN, SVM, etc.; study accuracy
3	Recognition in the real world: the long tail	Long-tail: [Survey] [StrongBaseline]	
4	Recognition in the real world: the open set	Open-set / anomaly: [MaxSoftmaxProb] [StrongBaseline] [OutlierExposure] [OpenGAN]	Project 1 due Project 2 out: implement MLP and CNN for recognition on CIFAR10 and MNIST; use each other as open-set to study uncertainty (e.g., max of softmax score, and entropy)
5	Segmentation	Segmentation: [Shi & Malik 2000] [MeanShift] [gPb] Semantic segmentation: [U-Net] [FCN] [DeepLab]	
6	Detection (part 1): Object Proposals	[SelectiveSearch] [FasterRCNN] [PixelGrouping]	Project 2 due Project 3 out: Make CIFAR10 long-tail, retrain CNN to report per-class accuracy; improve accuracy using long-tailed techniques (e.g., weight balancing).
7	Detection (part 2): Top-Down Methods	[DPM] [FPN] [MaskRCNN]	

8	Recognition with More Sensors and Modalities	lidar: [CenterPoint] [MVP] [DeepFusion] audio-visual [Owens & Efros 2018]	
9	Visual Learning with Less labels	Transfer learning: [Razavian et al. 2014] Few-shot: [Snell et al 2017] Semi-supervised: [self-training] [FixMatch] Weakly-supervised: [Detic]	Project 3 due Project 4 out: propose a research project worthy of 1.5 months to work on; do literature review; design experiments; prepare datasets; design methods; benchmark methods; write project report
10	Visual Learning with Unlabeled Data	[SimCLR] [MoCo] [MaskFeat]	Project proposal and literature review due as a part of project report
11	Visual Recognition and Language	[CLIP] [DALL·E 2]	
12	Motion and Optical Flow	[Lucas & Kanade 1981] [Sun et al. 2010] [PWC-Net] [Sun et al. 2021]	
13	Visual Learning from Videos	[spatial-time correspondence] [siam-MAE]	
14	Presentation of course project		Final project report due

Student Disabilities Support Service:
The University of Macau is committed to providing an equal opportunity in education to persons with disabilities. If you are a student with a physical, visual, hearing, speech, learning or psychological impairment(s) which substantially limit your learning and/or activities of daily living, you are encouraged to communicate with your instructors about your impairment(s) and the accommodations you need in your studies. You are also encouraged to contact the Student Disability Support Service of the Student Counselling and Development Section (SCD), which provides appropriate resources and accommodations to allow each student with a disability to have an equal opportunity in education, university life activities and services at the University of Macau. To learn more about the service, please contact SCD at scd.disability@umac.mo, or 8397 4901 or visit the following website:
http://www.umac.mo/sao/scd/sds/aboutus/en/scd_mission.php.

Coordinator:
Shu Kong, Assistant Professor of CIS

Persons who prepared this description:
Shu Kong, December 26, 2023

Part B General Course Information and Policies

2nd Semester 2023/2024
Instructor: Prof. Shu Kong Office: E11-4025
Office Hour: Tuesday and Wednesday at 15:00-16:00, or by appointment
Email: skong@um.edu.mo

Time/Venue

Lecture	Thursday 7pm-10pm	E11-1015

Grading Distribution:

Percent. Grade	Final Grade	Percent. Grade	Final Grade	Percent. Grade	Final Grade
100 – 93	A	77 – 73	B-	57 – 53	D+
92 – 88	A-	72 – 68	C+	52 – 50	D
87 – 83	B+	67 – 63	C	below 50	F
82 – 78	B	62 – 58	C-		

Homework Policy:
The completion and correction of homework are helpful for learning. As a result,

- There are four course projects.
- Projects are due two weeks after assignment unless otherwise noted.
- There is no late penalty but is a reward if submitting >24 hours earlier by the deadline.
- Possible revision of grades may be discussed with the grader within one week after deadline.
- The course grade will be based on the average of the project grades.

Other Important Notes:

- Check course web pages for announcements, and lectures.
- Cheating is absolutely prohibited by the university.