San José State University
Department of Computer Science
CS 256, Selected Topics in Artificial Intelligence, Section 2, Fall, 2019

Practical Computer Vision using Convolutional Neural Networks

Course and Contact Information

Instructor:	Mashhour Solh
Office Location:	Virtual (Zoom or Hangouts)
Phone	408-924-7171
Email:	mashhour.solh@sjsu.edu
Office Hours:	By email appointment
Class Days/Time:	MW 7:30 AM – 8:45 AM
Classroom:	Online (Zoom)
Prerequisites:	Preferred CS122 and CS156, or proficiency with Python, high-level familiarity with C++, Calculus and Linear Algebra

Course Format

This course will be a hands-on project-based class. The lecture time will be focused on lectures and recurring project presentations. Class projects are to be conducted in groups of 5 and each group will be provided AWS credits to complete their project. In-class participation and discussions are expected from every student and the instructor will randomly pick a member of each group to provide updates on the group's progress.

Canvas Course Site

Course materials such as syllabus, lecture slides, project discussions and exams can be found on the Canvas Leaning Management System course website at http://sjsu.instructure.com. You are responsible for regularly checking with Canvas to learn of any updates.

Course Description

Basics of computer vision and image classification: neural networks, linear classification, regression, backpropagation and loss functions. Walkthrough convolutional neural networks for computer vision: architectures, training, hyper-parameters, and applications. Introduction to deep learning software platforms: pyTorch, Tensorflow, and MXNet.

Course Goals

Upon successful completion of this course, students will be able to:

1. Explain the fundamentals of convolutional neural networks (CNN's).

2. Explain the difference between various deep learning architectures and basic computer vision topics such as pose estimation, object and human detections, face recognition, and depth estimation.
3. Read and understand state-of-the-art works of literature on deep learning. Build and run codes to validate it using one of the three main software platforms (MXNet, pyTorch or Tensorflow).
4. Work with large public datasets on computer vision.

Required Texts/Readings

Textbook

There is not textbook required for the course. The following e-books are sufficient reference for students: **" *Deep Learning*"** (https://www.deeplearningbook.org/) and **" Dive into Deep Learning**" (http://d2l.ai/)

Other Readings

Students will be asked to perform literature surveys on selected topics. Here are two resources that are useful to do the paper search: arxiv.org (https://arxiv.org/) and papers with code (https://paperswithcode.com/).

Other technology requirements / equipment / material

For the project the users will have to choose a software platform of the following: TensorFlow: (https://www.tensorflow.org/learn), pyTorch (https://pytorch.org/tutorials/) or MXNet (https://mxnet.incubator.apache.org/versions/master/tutorials/)

Pretrained models can also be found on ModelZoo: https://modelzoo.co/

Course Requirements and Assignments

The learning in the lectures will supplement your hands-on work on the project. The final assessment will be the project and final examination. The class project will be conducted in groups and is broken into four milestones:

- *Milestone 1*: Each group will be assigned a topic for reading and the group will research and read state-of-the-art papers on the assigned topic and produce a report and presentation.
- *Milestone 2*: Checkpoint A where each group will have to present their idea of project and to run a chosen state-of-the-art network on assigned topic and report the results in form of a report and presentation.
- *Milestone 3*: Checkpoint B where each group will have to present initial results of the project in form of a report and presentation.
- *Milestone 4*: Project shall be completed and final report, presentations and github page are expected.

Final Examination or Evaluation

The final examination will be based on the lectures provided in class as well as testing your knowledge of the topic assigned to your group. The location and time are TBA.

Grading Information

Determination of Grades

The final grade in the course will be calculated based on the following percentages:

- Attendance: 5%
- Project: 75% - broken as follows:
  - Milestone 1 20%
  - Milestone 2 10%
  - Milestone 3 10%
  - Milestone 4 35%
- Final Examination: 20%

Late Work or Rescheduling

Presentations, project reports, and the final exam will receive a zero if not performed on time. If you will be unable to make your group's scheduled presentation, let the instructors know at least 2 weeks in advance.

Grade Scale

The letter grade will be determined based on the following scale:
A+ = 98% - 100% **A** = 93% - 97% **A-** = 90% - 92% **B+** = 87% - 89% **B** = 83% - 86%
B- = 80% - 82% **C+** = 77% - 79% **C** = 73% - 76% **C-** = 70% - 72% **D** =60%-69%
F = below 60

Classroom Protocol

Regular attendance is an integral part of the learning process. Please arrive to class on time and make sure your cell phones are silent during the lecture. Your cameras shall be turned on and microphones muted unless you have a question to ask.

University Policies

Per University Policy S16-9, university-wide policy information relevant to all courses, such as academic integrity, accommodations, etc. will be available on Office of Graduate and Undergraduate Programs' Syllabus Information web page at http://www.sjsu.edu/gup/syllabusinfo/ . Make sure to review these policies and resources.

Class Schedule

This schedule is subject to change. Fair notice will be provided: Updates will be given in class and the schedule available on the course website will be updated.

Week	Date	Topic	
1	8-19	Course logistics	
2	8-24	Introduction	
2	8-26	Linear Classification	
3	8-31	Optimization	
3	9-2	Backpropagation	
4	9-7	Convolutional Neural Networks	
4	9-9	*Labor Day*	Milestone1 report is due on 9-13
5	9-14	Introduction to MXNet (AWS)	
5	9-16	Literature Survey Presentations I	
6	9-21	Literature Survey Presentations II	
6	9-23	Hardware and Software	
7	9-28	Training a Network I	
7	9-30	Training a Network II	
8	10-5	Architectures: AlexNet, VGG, GoogleNet	
8	10-7	Architectures: ResNet, MobileNet, HRNet	
9	10-12	Recurrent Networks	Milestone2 report is due on 10-14
9	10-14	Guest Lecture: Use-case Study	
10	10-19	Checkpoint A presentations I	
10	10-21	Checkpoint A presentations II	
11	10-27	Generative Models	
11	10-29	Reinforcement Learning	
12	11-2	Detection: R-CNN, Fast and Faster R-CNN	Milestone3 report is due on 11-4
12	11-4	Detection: YOLO and SSD	
13	11-9	Checkpoint B presentations I	
13	11-11	Checkpoint B presentations II	
14	11-16	*Veteran's Day*	
14	11-18	Segmentation	
15	11-23	Pose Estimation	
15	11-25	*Thanksgiving NI – No Class*	Final report is due on 12-1
16	11-30	*Review*	
16	12-2	Final Project presentations I	
17	12-7	Final Project presentations II	
Final Exam	12-11	07:15am -09:30am	