

**CS 523 S2017**  
**Multimedia Systems**  
**(Revised April 18, 2017)**

**FINAL DELIVERABLES - Due May 2, 5pm (see below)**

**Project 3 – Generative Adversarial Networks**

This project may be completed individually or in teams of two or three.

**Part 1 – due 3/28**

**Part 2 – due 4/4**

**Part 3 – due 4/25**

For this final assignment, you will design and implement a novel deep learning project of your choice. The project should incorporate a generative adversarial network (GAN) that synthesizes new multimedia data based on features learned from an input training set (or training sets). The project must be implemented in TensorFlow, with extensive documentation, including well-commented code and a conference-style write-up.

Part 1: Either a) use an existing GAN tutorial/implementation, *other* than the ones we went over in class on 3/7 (such as PPGANs, Stacked GANs, etc) to generate examples from a well-known training set (MNIST, Celebrity Faces, inceptionism, etc). Or, b) port an implementation of a GAN created in another library to TensorFlow. Or, c) implement a GAN algorithm from scratch in TensorFlow based on a published paper. (I will send links to recently published GAN papers before Spring Break.)

For class on 3/28, bring in a working example of code that generates new images. As best you can, be prepared to explain how the code works. Note: for this assignment you do not have necessarily to write your own code, but you must at least be able to use and clearly explain a TF code base in detail.

Part 2: Literature review. Each member of your group will be responsible for summarizing and discussing a different article about deep learning. (The list of articles you can choose from will be provided in a week or two, or you are welcome to suggest your own articles).

A) For class on 4/4, each person in your team will present a short overview of his or her article. You should create a small number of slides that explain the core ideas of the paper, and answer the following questions:

- What are the main innovations of the paper?
- What do the authors claim as their contributions to the literature of deep learning?
- What datasets do they use to test their results?
- Why do you think these results are (or are not) exciting / successful / important?
- How do you think the contributions described in the paper could be used for your own project?

B) Additionally, your group will produce a write-up summarizing each of your presentations (~2 pages for each article). Additionally, it will include a short proposal of your final project, which is due on 4/25 (see below).

Ideally, the papers you choose for Part 2 will be directly related to your proposal for Part 3.

### Part 3: Create a novel deep learning project.

For class on 4/25, you will design and implement a more extensive and polished deep learning project that makes use of a generative model as a core component. You will have decided upon your group and the final project by 4/4.

Deliverables:

4/25: In-class demonstration of your project + slides

5/2: Write-up, code, git repo, representative image, and video documentation due

- Your write-up will be (ideally) of good enough quality that it could be accepted to the CreativeAI workshop at KDD. (See <https://creativeai.mybluemix.net/> for information about this workshop.)
- The write-up will use the ACM SIGCHI Conference template, available here for Word or LaTeX: <https://www.acm.org/publications/proceedings-template>.
- The other deliverables are similar to what was required for Project 2

### **FINAL DELIVERABLES - Due May 2, 5pm**

In addition to the materials submitted by each group for Project 3, by 5pm on 5/2, each individual student will need to provide me with a PDF that has:

- Your name and UIC email
- A link to materials for each of short assignments: Assignment 1 and Assignment 2 (i.e. a PDF, videos, and/or git repo)
- A link to all of the materials for each of the three projects: Project 1, Project 2, and Project 3 (these should each be in a git repo)