

# Deep Learning for Computer Vision: **Tufts** Spring 2017

Spring 2017, TR 7:30 to 8:45pm, Halligan Hall 111B.



Multilabel Convolutional Neural Network (CNN) Classification results from the COCO-Attributes Dataset

## Course Description

### [Course Catalog Entry](#)

This course provides a practical foundation for deep learning, with a special emphasis on those methods used in computer vision. The first part of the class will introduce students to simple neural networks, convolutional neural networks, and some elements of recurrent neural networks, such as long short-term memory networks (LSTMs). Students will implement, train, and test their own end-to-end models for image classification. This course will cover the algorithmic fundamentals of deep learning (backpropagation, optimization, etc.) as well as engineering heuristics for getting networks to perform well.

The second part of the course investigates current research topics in deep learning and object detection. These topics will be pursued through independent reading, class discussion and presentations, and a final project that relates to current research problems in computer vision.

The goal of this course is to give students the background and skills necessary to perform research in deep learning and computer vision. Students should understand the strengths and weaknesses of current approaches to research problems and identify interesting open questions and future research directions. Students will hopefully improve their critical reading and communication skills, as well.

COMP 150-02 and the accompanying lectures, notes, and assignments are adapted from the Stanford CS class [CS231n](#) the UMassAmherst class [COMPSCI 697L](#). Many thanks to Li Fei-Fei, Andrej Karpathy, Justin Johnson and Erik Learned-Miller for graciously letting us use their course materials!

# Course Requirements

## Coding Assignments

There will be 4 coding assignments that lead students through a practical introduction to neural nets. Assignments will build on each other. In early assignments students will implement the basic components of deep nets. Later assignments will require implementing and testing increasingly complex networks. Code will be in python, and students will use [Jupyter notebook](#) to debug, test, and demo their work. If you are unfamiliar with Jupyter notebooks, [try them out!](#)

## Reading and Summaries

Students will be expected to read 10-11 papers from the deep learning literature. These papers will be presented in class by other students, and all student audience members are expected to be prepared for discussion. For each assigned paper, students must write a two or three sentence summary and identify at least one question or topic of interest for class discussion. Interesting topics for discussion could relate to strengths and weaknesses of the paper, possible future directions, connections to other research, uncertainty about the conclusions of the experiments, etc. Reading summaries must be posted to the [class blog](#) by **11:59pm the day before each class**. Feel free to reply to other comments on the blog and help each other understanding confusing aspects of the papers. The blog discussion will be the starting point for the class discussion. If you are presenting you don't need to post a summary to the blog.

## Presentations

Depending on enrollment, students will lead the discussion of one or two papers during the semester. Ideally, students would implement some aspect of the presented material and perform experiments that help understand the algorithms. Presentations and all supplemental material should be ready one week before the presentation date so that students can meet with the instructor, go over the presentation, and possibly iterate before the in-class discussion. For the presentations it is fine to use slides and code from outside sources (for example, the paper authors) but be sure to give credit.

## Semester projects

Students are expected to complete a state-of-the-art research project on topics relevant to the course. Students are encouraged to work in groups of 3-4. Students will propose a research topic part way through the semester. After a project topic is finalized, students will meet occasionally with the instructor to discuss progress. Students will present their progress on their semester project twice during the course and the course will end with final project presentations. Students will also produce a write-up of their project. Projects will be published on the this web page. The ideal project is something with a clear enough direction to be completed in a couple of months, and enough novelty such that it could be published in a peer-reviewed venue with some refinement and extension.

## Prerequisites

Strong mathematical skills (linear algebra, calculus, probability and statistics) and previous imaging (graphics, vision, or computational photography) courses are needed. It is strongly recommended that students have taken one of the following courses (or equivalent courses at other institutions):

- [COMP 135, Introduction to Machine Learning](#)
- [COMP 136, Statistical Pattern Recognition](#)

## Textbook

We will not rely on a textbook, although the free, online textbook "[Deep Learning](#)" by Ian Goodfellow and Yoshua Bengio and Aaron Courville is a helpful resource. The [Deep Learning Tutorial](#) by the Stanford Deep Learning group may also come in handy.

The assignments and lectures in class are largely based on the the Stanford CS class [CS231n](#). The course notes by Li Fei-Fei, Andrej Karpathy, and Justin Johnson will be linked where appropriate.

Additional notes and extensions on the CS231n course are generously contributed by Erik Learned-Miller and Hang Su from their course [compsci697l](#).

## Grading

Your final grade will be made up from

- 10% Homework [#1](#)
- 10% Homework [#2](#)
- 10% Homework [#3](#)
- 10% Homework [#4](#)
- 10% Reading summaries posted to [class blog](#)
- 10% [Paper presentation\(s\)](#), including partial system implementation or testing
- 40% Semester [Project](#)

## Course Questions

Tufts students can post questions about assignments or projects to [Piazza](#). Questions about discussion papers can also go on the [class blog](#).

## Staff



Instructor: **Genevieve  
Patterson**

---



TA: **Rishit Sheth**

---

## Office Hours:

Genevieve Patterson, ??? Rishit Sheth, ???

Course questions? Email [gen@cs.tufts.edu](mailto:gen@cs.tufts.edu)  
Tufts Computer Science