

CS331 Computer Vision

Fall 2024

Skidmore College

Syllabus

Instructor:

Dr. Michael Eckmann

Associate Professor

Computer Science Dept.

CIS 230E

email: meckmann@skidmore.edu

Instructor webpage: <https://meckmann.domains.skidmore.edu/>

Course webpage: <https://meckmann.domains.skidmore.edu/2024Fall/cs331/>

Texts

Computer Vision by Richard Szeliski <https://szeliski.org/Book/>

Deep Learning by Ian Goodfellow and Yoshua Bengio and Aaron Courville

<https://www.deeplearningbook.org/>

Appropriate OpenCV, numpy, Keras, and Tensorflow documentation and tutorials

Almost all of the topics presented throughout the semester will have associated readings from the sources just listed. Students will be alerted to which pages or websites to focus on for each topic.

Course Overview

Computer Vision (CV) trains machines to understand images and video. Classifying and detecting/recognizing objects in images are common problems that CV tackles. Machine Learning techniques, specifically Convolutional Neural Networks (CNNs), are often used. Students will learn digital image processing techniques that form the basis of many vision algorithms, as well as the concepts behind CNNs for classification.

Student Learning Objectives

1. Students will learn digital image processing concepts of histograms and convolution operation among others
2. Students will learn the concepts behind how convolutional neural networks are trained and can be used for classification tasks
3. Students will learn to write Python code using software libraries for computer vision and machine learning for vision
4. Students will gain experience in being able to read and present recent work published in the field

Learning Objectives Assessment

The learning objectives will be assessed by five homeworks (one of which is a presentation to the class) and a final project.

Class meetings

Lecture M 1:25 p.m. - 2:20 p.m.

Lecture Tu/Th 2:10 p.m. - 3:30 p.m.

all in CIS286

Assignments

Homework 1: Program involving HSV color space, Histograms, Nearest Neighbors

Homework 2: Linear Regression / Gradient Descent coding

Homework 3: Questions about Cost functions for linear and logistic regression, gradient descent algorithm, learning rates, etc.

Homework 4: Implement a Neural Network based on specifications given

Homework 5: Present an assigned paper to the class (papers will come from recent computer vision conferences / journals)

Final Project: Training CNNs

Grading criteria

Grading of most of the homeworks will be based on correctness.

The presentation will be graded by how accurately and thoroughly the paper is presented, the organization of the talk and to a lesser extent, the quality of the presentation.

The final project will be graded by how well the submission meets the requirements which include training several CNNs with various features and testing them and reporting their classification accuracy and other statistics. A writeup is expected as part of the final project from which I will assess how well the student understood the concepts.

Grading Policy

HW1: 12 %

HW2: 06 %

HW3: 12 %

HW4: 27 %

Presentation: 18 %

Final Project: 25 %

Grade is first computed by the above.

However, your final grade will be influenced either positively or negatively by how I, the instructor, evaluate you on the following:

- a) class participation
- b) good attendance (not just physical)
- c) increase of performance throughout the semester,
- d) quality/effort of your program/homework submissions (aside from just correctness)
- e) lateness of your program/homework submissions

Class participation includes answering questions in class, asking questions in class, visiting me during office hours, or by appointment and asking questions through email. Class participation is based on my assessment of the student's voluntary contribution, plus their responses to questions I ask them. If assignments are habitually submitted late or one is more than a couple of days late, then this can negatively effect your overall course grade.

If you are interested in the Satisfactory/Unsatisfactory (S/U) grade option, you are responsible for being aware of the pertinent deadlines set by the Registrar's Office for choosing this option. Please be advised that the college criterion for "Satisfactory" is the equivalent of a C or better. A grade of "C-" or below will result in the student receiving the grade of "Unsatisfactory" or "U" and, as a consequence, the student will receive no credit for the course.

Academic Integrity

I list here the policies by which the students of this class are expected to abide.

1. Skidmore Honor Code and Academic Integrity.
2. The Ethics of Scholarship.

Failure to abide by these policies results in a negative learning environment and you should expect to be held accountable.

Attendance

Please be on time as class will start promptly at the scheduled time.

Attendance is mandatory.

If you must miss class for a valid reason, please inform me ahead of time, via email preferably.

Topics in the order to be covered

Introduction to Computer Vision course and start Digital Image Processing concepts / techniques that are important to the study of Computer Vision

Convolution operation

Connected Components, Histograms, color spaces

Introduction to OpenCV, a software library for image processing and computer vision and numpy, a software library with mathematical functions and arrays for Python

Supervised Learning

Nearest Neighbors

Machine Learning basics / Linear Regression

Gradient Descent

Linear Regression for higher order polynomials

Logistic Regression for classification

Regularization

Neural Networks

Backpropagation

Computational Graphs

Convolutional Neural Networks

Convolutional layers, Pooling layers

Training tips

Batch Normalization, Xavier Initialization, Stochastic Gradient Descent

Softmax, more Regularization

Build a CNN in code tutorial using Keras / Tensorflow software libraries

Mixup

Learning Rate Schedules, Optimizers (besides SGD)

VGG, GoogLeNet, ResNet

Support Vector Machines

Student Presentations

Approximate weekly breakdown of list of topics with readings

Books:

Szeliski, Richard, "Computer Vision: Algorithms and Applications, 2nd ed.", 2022
available here: <http://szeliski.org/Book/> click the download link within the second sentence on that page and he'll send you a personalized pdf of the text

Goodfellow, et al., "Deep Learning", 2016
available here: <https://www.deeplearningbook.org/>

Software libraries openCV, numpy, keras, tensorflow
Tutorials for these libraries are assigned

Week 1:

Szeliski:

section 1.1 What is computer vision?
pages 18-22 regarding the state of the computer vision field since the year 2000
section 1.5 A note on notation
section 1.6 Additional reading
Digital cameras, sampling, color models
section 2.3 The digital camera
2.3.1 Sampling and aliasing
2.3.2 Color
2.3.3 Compression --- skim it

Week 2:

Szeliski:

In section 3.2 Linear Filtering
read pages 119 (bottom) through 121,
skip page 122
and do not worry too much about the formulas
read page 123 up to middle of page 124 (up to and including 1 st paragraph of section 3.2.1 Separable filtering)
examine figure 3.14 on 125
read section 3.2.2 Examples of linear filtering which starts on bottom of page 125 and ends top of 127

Week 3:

Szeliski:

section 3.3.3 Binary Image Processing, Distance Transforms and Connected Components
–
Pages 138 – top of 142
section 3.1.4 Histogram Equalization which is the end of page 115 to the top of page 119

Week 4:

Online tutorial of openCV

https://docs.opencv.org/master/d6/d00/tutorial_py_root.html

Online tutorial of numpy

https://numpy.org/doc/stable/user/absolute_beginners.html

Week 5:

Chapter 5 Deep Learning in Szeliski

– reading: pages 236 – middle of 243

In Goodfellow text, Chapter 5 Machine Learning Basics
from page 96 – middle of page 120

Week 6:

Readings on Gradient Descent for Linear regression

Readings on Logistic Regression

Week 7:

Readings on Regularization

Week 8:

Readings on Neural Networks, Activation Functions (e.g. ReLU), Forward Propagation, Backward Propagation

Reading on Computational Graphs

Week 9:

Readings on Convolutional Neural Networks (Conv Layers, pooling layers, etc.)

Readings on Training

Week 10:

Readings on Overfitting, Underfitting, Data Preprocessing, Batch Norm, Stochastic Gradient Descent, Weight Initializations

Introduction to Keras and Tensorflow libraries

Week 11:

Reading:

Szeliski section 5.3.6

Training and optimization and Sebastian Ruder's

2017 paper "An overview of gradient descent optimization algorithms" section 4, pages 4-10

Week 12:

Reading in Goodfellow et.al. "Deep Learning":

<https://www.deeplearningbook.org/contents/ml.html>

Section 5.7.2 Support Vector Machines

Pages: 139-140

Week 13:

Presentations

General Health and Safety

The College expects all members of the community to adhere to the College's COVID-19 health and safety guidelines at all times. Please be aware of the COVID-19 guidelines posted on the College's website (<https://www.skidmore.edu/campus-planning/status-levels.php>) and note the differences between Yellow Alert Status (substantial/high transmission levels), and Green Alert Status (low to moderate transmission levels).

If You Are Exhibiting COVID-19 Symptoms

If you think you are exhibiting symptoms of COVID-19, isolate and contact Health Services immediately (health@skidmore.edu, x5550). Please follow Health Services' recommendations regarding testing and when to safely return to class and public spaces. As with any health-related illness, we ask that you contact Health Services as soon as you can.

Proper Mask Wearing

When mask wearing is required based on the College's COVID-19 health and safety guidelines, you must wear a mask covering your mouth and nose fully at all times in the classroom. If your mask is not providing sufficient protection because it has slipped in some way, you are responsible for adjusting the fit. If I ask you to wear your mask properly, you must comply, or you can choose to leave the classroom. Academic Integrity Students are expected to follow the Skidmore College Honor Code and code of conduct to the fullest extent. A recommendation of a maximum penalty will be recommended for all violations of the Honor Code.

Honor Code

I hereby accept membership in the Skidmore College community and, with full realization of the responsibilities inherent in membership, do agree to adhere to honesty and integrity in all relationships, to be considerate of the rights of others, and to abide by the college regulations.

Honor Code Statement for Examinations

While taking this examination, I have not witnessed any wrongdoing, nor have I personally violated any conditions of the Skidmore College honor code.

Accommodating Students with Disabilities and Providing Accessibility

If you are a student with a disability and believe you will need academic accommodation, you must formally request accommodation from Meg Hegener, Coordinator of Student Access Services (mhegener@skidmore.edu). You will also need to provide documentation which verifies the existence of a disability and supports your request. For further information, please call 580-8150 to contact Student Academic Services in Starbuck Center.

Conscientious Religious Observance Policy

If religious observances cause absence from class, campus employment, athletic practice, and/or game days or necessitates accommodations, students should notify their faculty, coaches, or supervisors prior to the date(s) of their absence. New York State policy and Skidmore College policy mandates that students be allowed to make up academic work and/or campus employment requirements without penalty. These accommodations should not reduce the overall expectations of a course nor unduly

burden the student requesting accommodation. Faculty must permit students to take a makeup examination without any penalty if they have to miss an examination due to religious observances. Similarly, faculty must permit students to submit missed assignments by an agreed upon due date, without penalty. Although not required, the College highly recommends that students submit written notification of the pending religious observances at the start of the semester or at least one week before the date. As an option, students may use this form. Distributing the written notification during the first week of classes, campus employment, or the start of the athletic season gives students, faculty, coaches, or supervisors time to prepare for the absence. If a student, supervisor, coach, or faculty member feels the policy is being violated, they should contact the Dean of Faculty Office at 518-580-5705 (Palamountain 416), the Dean of Students Office at 518-580-5760 (Case Center 313), or Human Resources at 518-580-5800 (Barrett Center first floor).

Diversity and Inclusion

Skidmore College is committed to fostering a diverse and inclusive community in which members develop their abilities to live in a complex and interconnected world. Consistent with our educational mission, we recognize ourselves as a community that respects individual identities based on varying sociocultural characteristics such as race, ethnicity, gender identity and expression, sexual orientation, national origin, first language, religious and spiritual tradition, age, ability, socioeconomic status and learning style. We strive to create a socially just world that honors the dignity and worth of each individual, and we seek to build a community centered on mutual respect and openness to ideas—one in which individuals value cultural and intellectual diversity and share the responsibility for creating a welcoming, safe and inclusive environment. We recognize that our community is most inclusive when all members participate to their full capacity in the spirited and sometimes challenging conversations that are at the center of the college's educational mission.

Sexual and Gender-Based Misconduct: Title IX Statement

Skidmore College considers sexual and gender-based misconduct to be one of the most serious violations of the values and standards of the College. Unwelcome sexual contact of any form is a violation of students' personal integrity and their right to a safe environment and therefore violates Skidmore's values. Sexual and gender-based misconduct is also prohibited by federal and state regulations. Skidmore College faculty are committed to supporting our students and upholding gender equity laws as outlined by Title IX. If a student chooses to confide in a member of Skidmore's faculty or staff regarding an issue of sexual or gender-based misconduct, that faculty or staff member is obligated to tell Skidmore's Title IX Coordinator or Title IX Deputy Coordinator. The Title IX Coordinator or Deputy Coordinator will assist the student in connecting with all possible resources for support and options for reporting both on and off campus. Identities and details will be shared only with those who need to know to support the student and to address the situation through the college's processes. If the student wishes to confide in a confidential resource, the Counseling Center Staff, Health Services, and Victim Advocates (anonymous) are all options available. More information can be found at the Sexual and Gender-Based Misconduct website or by contacting the Title IX Coordinator, Joel Aure (jaure@skidmore.edu), 580-5708, or Deputy Coordinator for Student Affairs, Gabriela Melillo (gmelillo@skidmore.edu), 580-5022.