ECE736: 3D Image Processing and Computer Vision

COURSE OUTLINE

Please refer to course website for updated information.

COURSE DESCRIPTION

The goal of computer vision is to use acquired image data to infer something about the world. Central to computer vision are the mathematical models governing image formation and methods for processing and recovering information based on the model and the image data. In this course we concentrate on geometrical and statistical models of visual data. In the first part of this course, we take a geometrical approach to image formation and look at problems such as image blending and stitching and 3D reconstruction. In our discussion of 3D computer vision, we focus on how to make use of the spatial and temporal coherence imposed by camera geometry to reconstruct a 3D scene or track an object using several 2D images. These images are from a moving video camera, stereo camera rig or multiple views from a single camera. We will also look at computational photography where image analysis and processing algorithms are applied to one or more images to create images that go beyond the capabilities of traditional imaging systems. Assuming a statistical model for the visual data, we talk about learning and inference. We cover regression, and classification methods. We will also look at applications of neural networks in classification, regression, and several other computer vision applications.

SCHEDULE and MODE OF DELIVERY

The material for this course will be delivered lectures.

Lecture: Tuesdays and Thursdays 9:30 am – 11:00 am.

INSTRUCTOR

Dr. Shahram Shirani
E-mail: shirani@mcmaster.ca
Office: ITB-225
Phone: 905-525-9140 ext. 27943
Office Hours**:** By appointment – see course website for details

TEACHING ASSISTANTS

No teaching assistant.

COURSE WEBSITE/S

http://avenue.mcmaster.ca

COURSE OBJECTIVES

Learning Objectives:

- Model image formation in single camera and multi-camera setups
- Mathematically understand the relation between the 3D world and its projection in 2D images and learn how to reconstruct a 3D scene model from several 2D images
- Extract features from images and match/track them
- Choose the right regression model for a vision problem
- Choose the right classifier for a vision problem
- Employ the RANSAC algorithm to remove the effects of outliers
- Be able to apply computational photography techniques to solve image processing and computer vision problems
- Be able to devise a neural network-based solution to typical computer vision problems

ASSUMED KNOWLEDGE

Undergraduate level DSP, undergraduate level probability, undergraduate level image processing

COURSE MATERIALS

Textbooks:

- Richard Szeliski, “Computer Vision: Algorithms and Applications, 2nd edition”, 2021. Electronic version of the book is available free at: http://szeliski.org/Book/
- Richard Hartley, Andrew Zisserman, “Multiple View Geometry in Computer Vision” Second Edition, Cambridge University Press, 2004. Electronic version available online.
- Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, “An Introduction to Statistical Learning”, Springer Science, Business Media, New York 2017, available online.
- C. M. Bishop, “Pattern Recognition and Machine Learning”, Springer, 2006. Available online.
- C. C. Aggrawal, “Neural Networks and Deep Learning”, Second Edition, Springer, 2023, available online.

Reference Books:

- Simon J.D. Prince, “Computer vision: models, learning and inference”, Cambridge University Press, 2012.

Electronic version of the book is available free at:
http://www.computervisionmodels.com

Slides, reading material, homework assignments will be posted on Avenue to learn
http://avenue.mcmaster.ca

COURSE OVERVIEW

- Image formation, processing, and 3D reconstruction
  - Pinhole camera model
  - Camera calibration
  - 3D reconstruction
  - Pose estimation
  - Stereo correspondence and epipolar geometry
  - Structure from motion
  - Edge detection
  - Interest points, corners and local image features
  - Feature matching and tracking
  - Model fitting and RANSAC
  - Geometric alignment using deep learning
  - Optical flow
  - Optical flow using deep learning
  - Deformable contours

- Learning and inference in vision
  - Regression models
  - Classification models
  - Kernel methods
  - Support Vector Machine (SVM)
  - Neural networks Basics
  - Convolutional Networks
  - Deep learning (short introduction)
  - Attention and Vision Transformers
  - Generative models?
  - Learnt Image Compression?

- Computational photography
  - Digital camera pipeline
  - Image blending and compositing
  - Image retargeting
  - Texture synthesis and transfer

- Image completion / inpainting
- High dynamic range imaging?
- Image stitching and panorama?

A more detailed timeline is available on the course website.
At certain points in the course, it may make good sense to modify the schedule. The instructor may modify elements of the course and will notify students accordingly (in class, on the course website).

ASSESSMENT

Assessment:

- Homework: 80%
- Project: 20%

Homework: 3-4 mini-project style assignments.
Project: The project can be in the form of a survey about a computer vision related topic, part of a vision related research or developing a vision related application. A one-page project proposal is due by February 25, 2025. The project report is due at the time of presentation. Presentations will be scheduled for April.

Typical project topics:
Some applications of computer vision:
Augmented reality
Virtual reality
SLAM (simultaneous localization and mapping)
Image based tracking
Image based localization
Automotive driver assistance and traffic management
Eye and Head Tracking
Film and Video: Sports analysis
Gesture Recognition
Industrial automation and inspection
People tracking
https://www.cs.ubc.ca/~lowe/vision.html
Theory of computer vision
Deep learning and its application in computer vision

ACADEMIC INTEGRITY

You are expected to exhibit honesty and use ethical behaviour in all aspects of the learning process. Academic credentials you earn are rooted in principles of honesty and academic integrity. It is your responsibility to understand what constitutes academic dishonesty.

Academic dishonesty is to knowingly act or fail to act in a way that results or could result in unearned academic credit or advantage. This behaviour can result in serious consequences, e.g. the grade of zero on an assignment, loss of credit with a notation on the transcript (notation reads: “Grade of F assigned for academic dishonesty”), and/or suspension or expulsion from the university. For information on the various types of academic dishonesty please refer to the Academic Integrity Policy, located at https://secretariat.mcmaster.ca/university-policies-proceduresguidelines/
The following illustrates only three forms of academic dishonesty:
• plagiarism, e.g. the submission of work that is not one’s own or for which other credit has been obtained.
• improper collaboration in group work.
• copying or using unauthorized aids in tests and examinations.

AUTHENTICITY / PLAGIARISM DETECTION

Some courses may use a web-based service (Turnitin.com) to reveal authenticity and ownership of student submitted work. For courses using such software, students will be expected to submit their work electronically either directly to Turnitin.com or via an online learning platform (e.g. A2L, etc.) using plagiarism detection (a service supported by Turnitin.com) so it can be checked for academic dishonesty.
Students who do not wish their work to be submitted through the plagiarism detection software must inform the Instructor before the assignment is due. No penalty will be assigned to a student who does not submit work to the plagiarism detection software. All submitted work is subject to normal verification that standards of academic integrity have been upheld (e.g., on-line search, other software, etc.). For more details about McMaster’s use of Turnitin.com please go to www.mcmaster.ca/academicintegrity.

COURSES WITH AN ON-LINE ELEMENT

Some courses may use on-line elements (e.g. e-mail, Avenue to Learn (A2L), LearnLink, web pages, capa, Moodle, ThinkingCap, etc.). Students should be aware that, when they access the electronic components of a course using these elements, private information such as first and last names, user names for the McMaster e-mail accounts, and program affiliation may become apparent to all other students in the same course. The available information is dependent on the technology used. Continuation in a course that uses on-line elements will be deemed consent to this disclosure. If you have any questions or concerns about such disclosure please discuss this with the course instructor.

Some courses may use online proctoring software for tests and exams. This software may require students to turn on their video camera, present identification, monitor and record their computer activities, and/or lock/restrict their browser or other applications/software during tests or exams. This software may be required to be installed before the test/exam begins.

COPYRIGHT AND RECORDING

Students are advised that lectures, demonstrations, performances, and any other course material provided by an instructor include copyright protected works. The Copyright Act and copyright law protect every original literary, dramatic, musical and artistic work, including lectures by University instructors.
The recording of lectures, tutorials, or other methods of instruction may occur during a course. Recording may be done by either the instructor for the purpose of authorized distribution, or by a student for the purpose of personal study. Students should be aware that their voice and/or image may be recorded by others during the class. Please speak with the instructor if this is a concern for you.

CONDUCT EXPECTATIONS

As a McMaster student, you have the right to experience, and the responsibility to demonstrate, respectful and dignified interactions within all of our living, learning and working communities. These expectations are described in the Code of Student Rights & Responsibilities (the "Code"). All students share the responsibility of maintaining a positive environment for the academic and personal growth of all McMaster community members, whether in person or online. It is essential that students be mindful of their interactions online, as the Code remains in effect in virtual learning environments. The Code applies to any interactions that adversely affect, disrupt, or interfere with reasonable participation in University activities. Student disruptions or behaviours that interfere with university functions on online platforms (e.g. use of Avenue 2 Learn, WebEx or Zoom for delivery), will be taken very seriously and will be investigated. Outcomes may include restriction or removal of the involved students' access to these platforms.

ACADEMIC ACCOMMODATIONS

Students with disabilities who require academic accommodation must contact Student Accessibility Services (SAS) at 905-525-9140 ext. 28652 or sas@mcmaster.ca to make arrangements with a Program Coordinator. For further information, consult McMaster University's Academic Accommodation of Students with Disabilities policy.

ACADEMIC ACCOMMODATION FOR RELIGIOUS, INDIGENOUS OR SPIRITUAL OBSERVANCES (RISO)

Students requiring academic accommodation based on religious, indigenous or spiritual observances should follow the procedures set out in the RISO policy. Students should submit their request to their Faculty Office ***normally within 10 working days*** of the beginning of term in which they anticipate a need for accommodation or to the Registrar's Office prior to their examinations. Students should also contact their instructors as soon as possible to make alternative arrangements for classes, assignments, and tests.

EXTREME CIRCUMSTANCES

The University reserves the right to change the dates and deadlines for any or all courses in extreme circumstances (e.g., severe weather, labour disruptions, etc.). Changes will be communicated through regular McMaster communication channels, such as McMaster Daily News, A2L and/or McMaster email.

RESEARCH ETHICS

The two principles underlying integrity in research in a university setting are these: a researcher must be honest in proposing, seeking support for, conducting, and reporting research; a researcher must respect the rights of others in these activities. Any departure from these principles will diminish the integrity of the research enterprise. This policy applies to all those conducting research at or under the aegis of McMaster University. It is incumbent upon all members of the university community to practice and to promote ethical behaviour. To see the Policy on Research Ethics at McMaster University, please go to http://www.mcmaster.ca/policy/faculty/Conduct/ResearchEthicsPolicy.pdf.

www.eng.mcmaster.ca/ece