



(http://toronto.edu/)

# CSC420: Introduction to Image Understanding

## Fall 2025



Representations of "computer vision" created using GPT4

## Course Description

This class is an introduction to fundamental concepts in image understanding, the subdiscipline of artificial intelligence that tries to make the computers "see". It will survey a variety of interesting vision problems and techniques. Specifically, the course will cover image formation, features, object and scene recognition and learning, multi-view geometry and video processing. The goal of the class will be to grasp a number of computer vision problems and understand basic approaches to tackle them for real-world applications.

## Instructor



David Lindell  
(<https://davidlindell.com/>)

## Teaching Assistants



Javad Rajabi  
(<https://www.cs.toronto.edu/~rajabi/>)



Michael Neumayr  
(<https://github.com/michaelnoi>)



Andrew Guo  
(<https://www.linkedin.com/in/ykguo/>)

## Course Logistics

**Lectures:** Wednesdays 1:00pm-3:00pm (LEC0101; MP 137) and 3:00pm-5:00pm (LEC0201; GB 221). **Lectures will be recorded and made available on Quercus.**

**Tutorials:** Friday 1:00pm-2:00pm (LEC0101; GB 221) and 3:00pm-4:00pm (LEC0201; MP 137).

**Instructor office hours:** Wednesdays 11:10am-12:00pm (BA 7228).

**TA office hours:** Fridays 2:00pm-3:00pm in BA 2270.

**Forum:** Course announcements and general information will be posted on the course forum at Piazza (). Any questions can also be discussed on the course forum, though **do not expect an immediate response** from the TAs. Also, do not expect answers during weekends or at the last minute before assignments are due. Please only email the staff for emergency inquiries that cannot be handled on the forum via a private post: Course Staff Email (<mailto:csc420-2025-09@cs.toronto.edu>).

**MarkUs:** Assignment materials and grades will be posted on MarkUs. (<https://markus.teach.cs.toronto.edu/markus/courses/76>)

**Quercus:** Lecture recordings and tutorial codes will be posted on Quercus. (<https://q.utoronto.ca/>)

**Attendance:** While lectures will be recorded, in-person attendance is expected as the recording will not capture all presented material (e.g., handwritten annotations, chalkboard diagrams, etc.).

**Assignments:** All assignments must be submitted on MarkUs (<https://markus.teach.cs.toronto.edu/markus/courses/76>). You will be automatically added to MarkUs if you're taking the course. Please do not email the staff if you are not yet added at the beginning of the term; it may take a week or two to be added.


Remark requests must be submitted within 1 week of graded assignments being returned. Late remark requests will not be accepted.

**Lateness:** All assignments are due on Thursdays at 12:59 PM (i.e., in the afternoon) as indicated on the schedule. There is a one-hour grace period after the deadline to account for any technical glitches/issues during the submission. If you require an extension you must submit an official request to the course staff using this form (<https://forms.office.com/r/W58tpQkp7g>). **All assignments must be submitted within seven days of the deadline (no exceptions).**


**Textbook:** Readings are assigned from the following texts:

- Computer Vision: Algorithms and Applications 2nd Edition (Szeliski), available online (<http://szeliski.org/Book/>).
- Foundations of Computer Vision (Torralba, Isola, Freeman), available online (<https://visionbook.mit.edu/>).
- 3D Computer Graphics (Gortler); readings are posted on Quercus.
- Multiple View Geometry in Computer Vision (Hartley, Zisserman); readings are posted on Quercus.

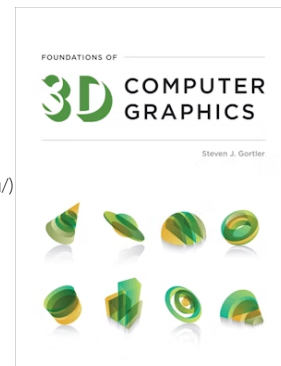
Some topics will have specific additional assigned reading materials as indicated in the course schedule.



(<http://szeliski.org/Book/>)



(<https://visionbook.mit.edu/>)



## Grading

### Assignments (29%)

- **Assignment 1: 8%**
- **Assignment 2: 7%**
- **Assignment 3: 7%**
- **Assignment 4: 7%**

There will be 4 assignments, posted every two weeks, starting with the second week. Assignments will consist of problem sets and programming problems with the goal of deepening your understanding of the material covered in class. *The assigned problem sets will not be graded; the programming problems will be graded.* Note that material for the midterm and final exams will be partly based on questions from the problem sets and programming problems. The programming problems will use Python and Jupyter Notebook (<https://jupyter.org/install>).

**Collaboration Policy:** The *graded portion* of assignments (programming problems) must be implemented individually; however, you may discuss general strategies and discuss the concepts of the course with other students, TAs (e.g., on Piazza), or AI. You must not share code or derivations. the *ungraded portion* of assignments (problem sets) may be completed collaboratively.

**AI Assistants Policy:** You are welcome to use AI assistants in a responsible fashion. You should not ask an AI assistant to complete your assignments for you; this is just as inappropriate as asking another student. However, you may ask an AI assistant questions about lecture material, clarifications about the assignments, tips on how to approach a problem, etc. (i.e., you may discuss the assignments with the AI assistant in the same way you are allowed to discuss with your peers or the TAs). **If you decide to use an AI assistant, you must include a sentence or two describing how you used it at the very top of your submitted assignment.**

### Midterm Exam (30%)

A timed one-hour exam held during the tutorial hour. The exam will be a combination of multiple choice and short-answer/long-answer. There are no make-ups for the midterm. If you miss the midterm, we will add a 30% weight to your final exam (so that the final will be worth 70% of your total grade).

## Final Exam (40%)

A timed three-hour exam consisting of a combination of multiple choice and short-answer/long-answer questions. **The time and date will be announced later in the term.**

## Ethics Module (1%)

Completion of two surveys based on a course ethics module of the course will be worth 0.5% each.

## Accessibility and Accommodations

Students with diverse learning styles and needs are welcome in this course. In particular, if you have a disability or health consideration that may require accommodations, please feel free to approach me and/or the Accessibility Services Office as soon as possible. The Accessibility Services staff are available by appointment to assess specific needs, provide referrals and arrange appropriate accommodations. The sooner you let them and me know your needs, the quicker we can assist you in achieving your learning goals in this course: <https://www.studentlife.utoronto.ca/as>, 416-978-8060, or [accessibility.services@utoronto.ca](mailto:accessibility.services@utoronto.ca)

## Schedule and Syllabus

Week	Date	Description	Material	Readings	Event	Deadline
Week 1	Wed Sep 3	<b>Lecture 1: Introduction &amp; Machine Learning Review</b> Course overview, review of fully connected networks, optimization, backpropagation	[slides (slides/lecture1.pdf)]	FCV 9-14 ( <a href="https://visionbook.mit.edu/part_foundation_learning.html">https://visionbook.mit.edu/part_foundation_learning.html</a> ) What is a neural network? ( <a href="https://www.youtube.com/watch?v=aircAruvnKk">https://www.youtube.com/watch?v=aircAruvnKk</a> ) Gradient descent ( <a href="https://www.youtube.com/watch?v=IHZwWFHWa-w">https://www.youtube.com/watch?v=IHZwWFHWa-w</a> ) Backpropagation ( <a href="https://www.youtube.com/watch?v=Ilg3gGewQ5U">https://www.youtube.com/watch?v=Ilg3gGewQ5U</a> ) Backpropagation calculus ( <a href="https://www.youtube.com/watch?v=tIeHLnjs5U8">https://www.youtube.com/watch?v=tIeHLnjs5U8</a> ) (optional) Notes on backpropagation ( <a href="https://cs231n.github.io/optimization-2/">https://cs231n.github.io/optimization-2/</a> ) (optional) Least squares review ( <a href="https://q.utoronto.ca/courses/396057/files?preview=38575754">https://q.utoronto.ca/courses/396057/files?preview=38575754</a> ) (optional) SVD review ( <a href="https://q.utoronto.ca/courses/396057/files?preview=38575755">https://q.utoronto.ca/courses/396057/files?preview=38575755</a> )	<b>Assignment 1 out on MarkUs</b>	
	Fri Sep 5	<b>Tutorial 1</b> PyTorch tutorial, Jupyter Notebook tutorial, lecture 1 practice problems, A1 Q&A				
Week 2	Wed Sep 10	<b>Lecture 2: Recognition</b> Convolutional neural networks for segmentation, object detection, and classification	[slides (slides/lecture2.pdf)]	FCV 24 ( <a href="https://visionbook.mit.edu/part_neural_architectures.html">https://visionbook.mit.edu/part_neural_architectures.html</a> ) FCV 50 ( <a href="https://visionbook.mit.edu/object_recognition_v3.html">https://visionbook.mit.edu/object_recognition_v3.html</a> ) (optional) Hubel & Wiesel 1979 ( <a href="https://q.utoronto.ca/courses/396057/files?preview=38575741">https://q.utoronto.ca/courses/396057/files?preview=38575741</a> )		
	Fri Sep 12	<b>Tutorial 2</b> PyTorch Tutorial (CNNs), fine-tuning, logging tools, lecture 2 practice problems, A1 Q&A				
Week 3	Wed Sep 17	<b>Lecture 3: Vision Transformers</b> attention, transformers, vision transformers, CLIP, autoregressive text-to-image	[slides (slides/lecture3.pdf)]	FCV 25 ( <a href="https://visionbook.mit.edu/part_neural_architectures.html">https://visionbook.mit.edu/part_neural_architectures.html</a> ) FCV 26 ( <a href="https://visionbook.mit.edu/part_neural_architectures.html">https://visionbook.mit.edu/part_neural_architectures.html</a> ) FCV 51 ( <a href="https://visionbook.mit.edu/object_recognition_v3.html">https://visionbook.mit.edu/object_recognition_v3.html</a> )		
	Thu Sep 18					<b>Ethics Survey 1 due</b>
	Fri Sep 19	<b>Tutorial 3</b> understanding & visualizing attention, lecture 3 practice problems, A1 Q&A				

Week 4	Wed Sep 24	<b>Lecture 4: Generative Models &amp; Embedded Ethics</b> VAEs, GANs, diffusion models, bias in deep learning	[slides (slides/lecture4.pdf)]	FCV 32.1–32.3, 32.8–32.10 ( <a href="https://visionbook.mit.edu/generative_models.html">https://visionbook.mit.edu/generative_models.html</a> ) FCV 33 ( <a href="https://visionbook.mit.edu/generative_modeling_and_rep_learning.html">https://visionbook.mit.edu/generative_modeling_and_rep_learning.html</a> ) FCV 35 ( <a href="https://visionbook.mit.edu/bias_and_shift.html">https://visionbook.mit.edu/bias_and_shift.html</a> ) In Class Exercise: CelebA-HQ Samples ( <a href="slides/celebahq.html">slides/celebahq.html</a> )	<b>Assignment 2 out on MarkUs</b>
	Thu Sep 25				<b>Assignment 1 due 12:59 pm</b>
	Fri Sep 26	<b>Tutorial 4</b> exploring vision latent spaces, lecture 4 practice problems, A2 Q&A			
Week 5	Wed Oct 1	<b>Lecture 5: Digital Photography</b> ray optics, aperture, depth of field, exposure, sensor noise, demosaicking, denoising, deconvolution	[slides (slides/lecture5.pdf)]	FCV 5 ( <a href="https://visionbook.mit.edu/imaging.html">https://visionbook.mit.edu/imaging.html</a> ) Szeliski 2.3	
	Fri Oct 3	<b>Tutorial 5</b> lecture 5 practice problems, A2 Q&A			
Week 6	Wed Oct 8	<b>Lecture 6: Radiometry and Photometric Cues</b> radiometric units, reflectance, photometric stereo, shape from shading, single-image depth	[slides (slides/lecture6.pdf)]	Gortler 21 (Quercus) FCV 43 ( <a href="https://visionbook.mit.edu/3d_learning.html">https://visionbook.mit.edu/3d_learning.html</a> ) (optional) First Principles of Computer Vision (Photometric Stereo I) ( <a href="https://www.youtube.com/watch?v=mw43A8Y6NIg">https://www.youtube.com/watch?v=mw43A8Y6NIg</a> ) (optional) First Principles of Computer Vision (Photometric Stereo II) ( <a href="https://www.youtube.com/watch?v=dNkMLqBUNKI">https://www.youtube.com/watch?v=dNkMLqBUNKI</a> )	
	Fri Oct 10	<b>Tutorial 6</b> lecture 6 practice problems, A2 Q&A			
Week 7	Wed Oct 15	<b>Lecture 7: 3D Vision &amp; Stereo I</b> coordinate systems, calibration, stereo geometry, triangulation, stereo matching	[slides (slides/lecture7.pdf)]	FCV 39 ( <a href="https://visionbook.mit.edu/imaging_geometry.html">https://visionbook.mit.edu/imaging_geometry.html</a> ) FCV 40 ( <a href="https://visionbook.mit.edu/3d_scene_understanding_stereo.html">https://visionbook.mit.edu/3d_scene_understanding_stereo.html</a> ) (optional) Notes on stereo vision ( <a href="https://q.utoronto.ca/courses/396057/files?preview=38575743">https://q.utoronto.ca/courses/396057/files?preview=38575743</a> )	<b>Assignment 3 out on MarkUs</b>
	Thu Oct 16				<b>Assignment 2 due at 12:59 pm</b>
	Fri Oct 17	<b>Tutorial 7</b> lecture 7 practice problems, midterm review, A3 Q&A			
Week 8	Wed Oct 22	<b>Lecture 8: Epipolar Geometry &amp; Stereo II</b> Essential & Fundamental Matrices, 8-point algorithm, active stereo, volumetric stereo	[slides (slides/lecture8.pdf)]	Zisserman & Hartley 9, 11 Szeliski 12.7.2 Szeliski 13.2 In Defense of the 8-point Algorithm ( <a href="https://q.utoronto.ca/courses/396057/files?preview=38575744">https://q.utoronto.ca/courses/396057/files?preview=38575744</a> )	
	Fri Oct 24	<b>Midterm</b>			

Week 9	Wed Oct 29	<b>Reading Week (No Lecture)</b>			
Week 10	Wed Nov 5	<b>Lecture 9: Structure from Motion &amp; Radiance Fields</b> feature matching, bundle adjustment, learned SfM, volume rendering, surface-based representations	[slides (slides/lecture9.pdf)]	FCV 44 ( <a href="https://visionbook.mit.edu/multiview.html">https://visionbook.mit.edu/multiview.html</a> ) FCV 45 ( <a href="https://visionbook.mit.edu/multiview.html">https://visionbook.mit.edu/multiview.html</a> )	
	Fri Nov 7	<b>Tutorial 9</b> lecture 9 practice problems, A3 Q&A			
Week 11	Wed Nov 12	<b>Lecture 10: Optical Flow &amp; Tracking</b> Motion estimation, patch match, 3D motion, brightness constancy, applications	[slides (slides/lecture10.pdf)]	FCV 46-49 ( <a href="https://visionbook.mit.edu/part_understanding_motion.html">https://visionbook.mit.edu/part_understanding_motion.html</a> )	<b>Assignment 4 out on MarkUs</b>
	Thu Nov 13				<b>Assignment 3 due at 12:59pm</b>
Week 12	Mon Nov 17	<b>Tutorial 10 (on zoom)</b> lecture 10 practice problems, A4 Q&A			
	Wed Nov 19	<b>Lecture 11: Modeling Humans</b> parametric human models, pose estimation, mesh recovery, generation	[slides (slides/lecture11.pdf)]	Szeliski 13.6	
	Fri Nov 21	<b>Tutorial 11</b> lecture 11 practice problems, A4 Q&A			
Week 13	Wed Nov 26	<b>Lecture 12: Course Review &amp; Selected Topics Potpourri</b>	[slides (slides/lecture12.pdf)]		
	Thu Nov 27				<b>Assignment 4 due/Ethics Survey 2 due 12:59 pm</b>
	Fri Nov 28	<b>Tutorial 12</b> review and final exam Q&A			
	TBD	<b>Final Exam</b>			

## Miscellanea

### Installing Python

We will only support Python 3.7 and recommend that you install it using Anaconda or Miniconda (see installation instructions here (<https://docs.anaconda.com/anaconda/install/>)).

## Acknowledgements

Some course content is adapted from previous offerings by Sanja Fidler, Ahmed Ashraf, and Babak Taati. Many of the slides are modified from those of Philip Isola, Bill Freeman, Yannis Gkioulekas, Gordon Wetzstein, and Steve Seitz. The webpage is based on the website for CS231N and EE367 at Stanford University.

## Previous offerings

- 2023 (<https://www.cs.toronto.edu/~lindell/teaching/420/2023>)
- 2024 (<https://www.cs.toronto.edu/~lindell/teaching/420/2024>)
- 2025 (<https://www.cs.toronto.edu/~lindell/teaching/420/2025>)