# Course Outline

## Computational Perception COMP 546

## Winter 2024

Tues./Thurs. 2:35 PM – 3:55 PM ENGTR 1080


| | |
|---|---|
| Instructor: | Michael Langer |
| Office: | McConnell Engineering 329 |
| Email: | langer@cim.mcgill.ca |
| Office Hours: | Tues/Thurs (after class) 4–4:30 PM or by appointment |


### Teaching Assistants

The T.A.s are Venissa Carol Quadros and Kenzy Abdel Malek. They will help out mostly with the assignments. Their office hours and contacts will be posted on the assignment PDFs.

# 1 Overview

This course is an introduction to visual perception from a *computational* perspective. Unlike traditional visual perception courses offered in Psychology or Physiology departments which emphasize psychological or neuroscience approaches, this course emphasizes computational aspects of perception. It covers some basic psychology and neuroscience, including some history and some classical experimental findings and theory. These provide the background for the computational models. Indeed as you will learn, the models are often derived directly from what we know from psychology and neuroscience.

Note that in previous years and in the McGill e-Calendar description, the course covered topics in both vision and audition. However, this year *the course will only cover vision only.*

We examine the sensory signals (images) from the visual environment, the information that is contained in them, and how this information is processed. The information includes luminance and color, defocus, binocular disparity, motion, and texture. We examine how properties of the environment (depth, layout, shape, object and scene categories) can be inferred from this image information.

Many students in the course may have no background in vision science. So we will need to spend time on the basics including some classical ideas and theories. Vision science has a long history and one cannot understand where we are today without understanding how theories of vision – especially computational theories – have evolved over time. That said, rather than taking a chronological approach, we will interleave classical and more modern approaches within the various topics of the course.

 

Here is a more detailed breakdown of the course topics. Each of the topics will take 1-2 lectures. We will cover the topics in the order below:

- **image projection**: visual angle and direction
- **focus**: thin lens equation, depth of field, accommodation
- **photoreceptors, color**: spectra, metamers
- **retina**: neural coding, neurotransmitters, spikes, DOG models
- **orientation and V1**: cross-correlation, simple cells, retinotopic maps, Gabor models, complex cells
- **filtering**: contour and texture, multiscale analysis, sparse coding and natural image statistics
- **binocular disparity**: depth, disparity space, accommodation-vergence conflict
- **image motion**: motion constraint equation; normal velocities, population codes
- **egomotion, eye movements**: motion parallax, VOR, smooth pursuit eye movements, saccades
- **attention**: visual search, saliency, change blindness
- **psychophysics**: psychometric curves, thresholds, JNDs, Weber's Law, signal detection theory
- **peripheral vision and texture**: crowding, texture metamers, image similarity metrics, motion textures
- **perceptual organization**: Gestalt grouping, figure-ground, T and L junctions, contour grouping
- **occlusions and transparency**: half-occlusions in stereopsis, accretion-deletion in motion, occlusions and blur, X junctions
- **illumination and reflectance**: intrinsic images, lightness and color constancy
- **3D shape from X, material perception**: shape from shading and texture, slant and tilt, gradients and flows, highlights
- **perspective**: vanishing points, Manhatten frames
- **probabilistic vision**: natural scene statistics, maximum likelihood vs. Bayesian models
- **CNN's** (convolutional neural networks): basic architecture, classification & regression problems
- **object recognition**: history (structural approaches, image-based approaches, neural network approaches), faces
- **scene recognition**: gist, rapid serial processing

 

# 2 Course Materials

The materials will consist of a slides, lecture notes, and exercises which will be made available on mycourses.

## 2.1 Copyright policy

You are not allowed to post any of my course materials on github, coursehero, any other websites. This includes PDFs of lecture slides, lecture notes, exercises, quizzes, assignment questions or anything else that we provide for you.

Stated more formally: "Instructor-generated course materials are protected by law and may not be copied or distributed in any form or in any medium without explicit permission of the instructor(s). Note that infringements of copyright can be subject to follow up by the University under the Code of Student Conduct and Disciplinary Procedures."

# 3 Prerequisites

There are no official prerequisites for the course.

It is assumed students can program in a high level language at a level of least COMP 250. The Assignments will use Python or Matlab. So, if you don't know one of these languages, then now is the time to learn!

The mathematics background needed is:

- basic multivariable calculus e.g. partial derivatives, polar & spherical coordinates
- linear algebra and geometry e.g. vector spaces, matrices, eigenvectors
- probability and statistics – e.g. conditional probabilities, Bayes theorem, mean and variance

The course will introduce some basic psychology and physiology of vision. No prior knowledge of these topics is assumed.

# 4 Evaluation

Your final course grade will consist of the following components.

### Four Assignments ($4 \times 10\% = 40$ %)

As mentioned above, the assignments will be done in Python or Matlab.

- A1 will be posted around January 19. It will cover basic models of retinal image processing, namely models of center-surround cells.
- A2 will be posted around February 2. It will cover orientation selection and binocular disparity.
- A3 will be posted around February 23. It will cover image motion, and psychophysics.
- A4 will be posted in mid-March. It will cover contour grouping, natural scene statistics, and more psychophysics.

You will be given approximately two weeks to do each of the four Assignments. If you do not do an Assignment, then you will receive a grade of 0 for it. Extensions can be given only for unforeseen reasons, such as illness. The instructor reserves the right to ask for documentation.

## Quizzes (0-15 %)

There will be three quizzes, each worth 5 %. If you miss a quiz or do poorly on one of them, you can replace the percentage grade for that quiz with your final exam grade. In this sense, the quizzes are optional.

The quizzes will cover lecture material up to the Study Break. The dates are:

- Quiz 1 – Mon. Jan. 29 – covering lectures 1-5
- Quiz 2 – Mon. Feb. 19 – covering lectures 6-10
- Quiz 3 – Mon. March 11 – covering lectures 11-16

The quiz questions will be a mix of multiple choice, multiselect, and short answer. They will be done online on mycourses.

Since the quizzes are online, they are "open book". But you must do them entirely on your own. No communication is allowed between students about the quiz on the day of the quiz.

You will be given one hour – namely any one hour interval you wish – from 8 AM to 10 PM. (Note you must start by 9 PM.) If you are registered with the SAA office and need more time than that, then you will need to let me know at least one week in advance.

## Final Exam (45-60 %)

The Final Exam will be held during Final Examination Period.

It will be worth either 45%, 50%, 55%, or 60%, depending on whether your percentage grade on the final exam is greater than each of the three quiz percentage grades.

The final exam will be *closed book*. No crib sheet or electronic devices are permitted.

The final exam will contain a mix of short and long answer questions.

## McGill policy on the Evaluation scheme

As stated in Article 3.2.3, of the student assessment policy:
https://www.mcgill.ca/secretariat/files/secretariat/2016-04_student_assessment_policy.pdf
"*In the event of extraordinary circumstances beyond the University's control, the evaluation scheme in a Course is subject to change, provided that there be timely communications to the students regarding the change.*"

## 4.1 Calculation of final course grade

There are many factors that determine your final grade, including how hard you work, how talented you are in this subject, how much time you devote to the course, what your academic background is, what your health or family situation is, etc. However, these factors will not be considered when calculating your final course grade. Rather, your grade will be determined entirely by the above grading scheme.

Your final course grade will be rounded off to the nearest integer. If one has a grade of 84.4 then it rounds to 84 and one gets an A-, whereas if it is 84.6 then it rounds to 85 and one gets an A. If one's grade is 84.5, it will round it up to 85. The same rounding procedure holds for low grades. If one's final course grade is 49.4 then it rounds to 49 which is an F. A very hard line is drawn here, so one does not want to fail then one should stay far away from that line.

Note that graduate students need a 65 to pass.

## 4.2 Regrading

Mistakes can occur when grading assignments or exams. Not surprisingly, requests for re-grading are always in situations in which students feel they received fewer points than deserved, rather than more points than deserved. With that upward tendency in mind, please note that if you wish the instructor or the TAs to re-grade a question on an assignment or exam, we will do so. However, to avoid upward grade ratcheting, we reserve the right to re-grade other questions as well.

## 4.3 Additional Work

You will *not* be given the opportunity to do additional work to upgrade your grade.

## 4.4 Supplemental/Deferred Exam

The Supplemental/Deferred Exam exam will cover the same material as the Final Exam and will replace the Final Exam grade. The same "max" rule for quizzes and final exam will apply.

For information on Supplemental Exams, see
https://www.mcgill.ca/science/student/general/exams/supplemental.

# 5 Communication Policies

## 5.1 Course Announcements

Important information about the course will be announced in class, or on the Ed Discussion board.

## 5.2 Getting help from the instructor (me) or a T.A.

If you have a technical question about the course material or an assignment, *do not email me*. Instead:

- If it is an assignment question, then see a T.A. during their office hours, or post your question on the Ed Discussion board. In general, posting on Ed means that other students can benefit from the answer.

  There is a private setting in case you do not wish other students to see your post. There is also an anonymous setting in case it is ok for other students to see your post but you don't want to identify yourself.

- If it is a question about a quiz or lecture material, then post your question on Ed. (See Discussion Board policies below.) I will do my best to answer your question within 24 hours.

## 5.3 Instructor email policy

I am notified whenever there is a post on Ed, so the only reason to email me rather posting on Ed is if you do not wish the TA's to see your post, for example, if you have an urgent and important personal matter.

You can email one of the TA's for matters that are entirely between you, e.g. following up on a correspondence or arranging for a meeting outside of office hours.

## 5.4 Discussion Board Guidelines

The T.A.s and I will moderate the Discussion Board. We also strongly encouraged you to help each other out by responding to posts. Guidelines for posting on the Discussion Board are as follows.

- Use the search feature to check if your question has been asked before by entering relevant keywords into the search box. Post your question only if you are sure it has not been posted before.

- Choose the appropriate folder (Topic) and subject line.

- If you have multiple questions that are unrelated, then use multiple postings so others can more easily follow the thread.

- Proofread before posting. Ensure that what you have written makes sense before hitting Enter.
- If you would like your posting to be deleted, make a request within the thread.

### 5.5 McGill language policy

*In accord with McGill University's Charter of Students' Rights, students in this course have the right to submit in English or in French any written work that is to be graded.*

https://www.mcgill.ca/study/2017-2018/university_regulations_and_resources/undergraduate/gi_lang_policy

## 6 Academic Integrity

### 6.1 Collaboration on assignments and plagiarism

We encourage you to discuss the assignment problems with each other, and to help each other out with debugging. We also encourage you to use the Discussion Board (Ed Discussion). *However, there are limits to your collaboration.* You can give hints to each other, and the TAs and instructors will give hints sometimes if you ask. However, the hints and discussion should not go so far that you are revealing the solutions to each other, and you must never copy code from each other. Any cases of suspected plagiarism will be reported to the Faculty of Science Disciplinary Office.

### 6.2 McGill policy on academic integrity

*McGill University values academic integrity. Therefore, all students must understand the meaning and consequences of cheating, plagiarism and other academic offenses under the Code of Student Conduct and Disciplinary Procedures. See* ***http://www.mcgill.ca/students/srr/honest/*** *for more information.*