

# Course Outline

Computational Perception

COMP 546

Winter 2019

Tues./Thurs. 1:05 PM – 2:25 PM

ENGTR 1100

Instructor: Michael Langer  
Office: McConnell Engineering 329  
Email: langer@cim.mcgill.ca  
Office Hours: TBA on mycourses Announcements

## Teaching Assistants (T.A.)

T.A. office hours and contacts will be posted on mycourses Announcements.

## Overview

This course examines fundamental computational problems in visual and auditory perception. Unlike traditional perception courses offered in Psychology or Physiology departments which emphasize neural mechanisms, this course emphasizes computational aspects of perception. The course consists of two main topics, namely vision and audition. For both of these perceptual modalities, we begin by examining the signals from the environment, namely visual and auditory images, and the information that is contained in these images. For vision, we consider color, shading, binocular disparity, motion, and focus. For audition, we consider information carried by impact vs. non-impact sounds, echos, as well as binaural timing and intensity differences.

For both vision and audition, we then examine how images are processed by the sensory system, using concepts and tools from linear system theory. For vision, we discuss retinal and cortical processing. For audition, we discuss how sound waves are decomposed into frequency bands by the ear and encoded by the auditory cortex. We then examine how properties of the environment can be inferred from the information that is extracted from images. For vision, we consider how depth is estimated and for audition, we consider how depth and direction are estimated. We will also briefly discuss problems of object and scene recognition.

The course will consist of 25 lectures, each 80 minutes.

Lecture slides, notes, and exercises will be given as PDFs on mycourses.

# Prerequisites

There are no official prerequisites for the course. It is assumed students can program in a high level language at least that level of COMP 250 and that you are very comfortable with basic mathematics needed for an undergrad degree in computer science, in particular:

- multivariable calculus (MATH 222 or equivalent)
- linear algebra (MATH 223 or equivalent) - e.g. vector spaces, complex numbers
- probability (normal/Gaussian distributions and definitions such as mean and variance, joint and conditional probabilities).
- waves and optics (CEGEP level or PHYS 101/102).

The course will cover some basic psychology and physiology of vision and audition. It will also cover basic tools of linear system theory (convolution, Fourier transforms). No prior knowledge of these topics is assumed.

# Evaluation

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 have available because of other commitments, what your academic background is, what your health situation or family situation is, etc. *However, we do not consider these factors when we calculate your final course grade. Rather, we calculate your final grade according to the following breakdown:*

- Assignments + Research Paper presentation (50 %)
- Exams (50 %)

Details on each of the components are given below.

Your final course grade will be rounded off to the nearest integer. If your grade is 84.4 then it rounds to 84 and you get an A-, whereas if it is 84.6 then it rounds to 85 and you get an A. If your grade is 84.5, it will round it up to 85. The same round off procedure holds for low grades. If your final course grade is 49.4 then it rounds to 49 which is an F. I draw a very hard line on this, so if you don't want to fail then you should stay far away from that line.

## Assignments + Research Paper Presentation (50 % total)

There will be three assignments and a research paper presentation. These four components will be weighted equally (12.5 % each) and will be done in the following order:

- A1 will be posted around January 16; it will cover basic retinal image processing
- A2 will be posted in early February; it will cover early visual processing: orientation selection, binocular disparity, and motion;
- The research paper will be presented in a poster session that will be scheduled in March, after the Study Break. Details to be announced later.
- A3 will be posted by April 1 and will cover linear systems theory (Fourier transforms) and sound analysis using spectrograms;

You will be given approximately two weeks to do each of the three Assignments. If you do not do an Assignment, then you will receive a grade of 0 for it. There are no exceptions. Extensions can be given only for unforeseen reasons, such as illness. The instructor reserves the right to ask for documentation.

The research paper component of the course is new. *Logistics and further specification will be given by Monday Jan. 21, when I know better how many of students will remain in the course. (As of January 1, 2019 there were 80 students registered.) If the number of students remains more than say 40, then mostly likely I will require you to work in pairs.*

You will read a research paper in the area of computational perception, present this paper in a poster session, and write up a summary of the paper. *A list of research papers to choose from will be provided for you.*

We will most likely have to hold two poster sessions, to avoid scheduling conflicts. The poster sessions will be held in the late afternoon or early evening.

The poster itself should contain the title and bibliographic information, but otherwise it should only contain figures from the paper. It should not contain any text, other than the figure caption or figure annotations. The reason for this restriction is that the poster is only a visual aid to help you explain to the other students what the paper is about.

Before your poster presentation, you must submit a one page summary of what you plan to say in your poster (single space, 12 point font). Consider this to be a transcript of exactly what you plan to say when explaining the poster. *In the case that we ask you to work in pairs, you will be required to submit a single summary (mycourses logistics to be determined) and by default you will receive the same grade.*

After the poster session, you will submit a self-evaluation. What went right? What went wrong? What, if anything, would you have done differently? In the case that you worked in pairs, you will also be asked to submit a statement about the work done by each of you.

## Midterm Exam and Final Exam (total 50 %)

The midterm exam will take place in class on the Tuesday March 12 which is the first Tuesday after the Study Break. There will be no makeup exam.

The Final Exam will be held during Final Examination Period.

Your grade for the exam component is a total of 50%, which will either be 15% for midterm and 35 % for final, or 0 % for the midterm and 50 % for final, whichever is higher. If you do not write the midterm exam, your grade will be 0 and your final exam will automatically be worth 50 % of your course grade.

Both exams will be *closed book*. No crib sheet or electronic devices are permitted.

Both exams will contain a mix of multiple choice questions and short answer questions. For the multiple choice, there will be four choices on each question. If you answer a question correctly, you get 1/1. If you do not answer a question, you get 0/1. If you answer a question but your answer is incorrect, you will receive -0.2, that is, you will be penalized by 0.2 points. This penalty is meant to discourage you from guessing for questions in which you have no idea.

## Miscellaneous Policies

### 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 exam or assignment, we will do so. However, to avoid upward grade ratcheting, *we reserve the right to re-grade other questions as well.*

### Additional Work

If you receive a grade of D, F or J, you will *not* be given the opportunity to complete additional work to upgrade your grade.

### Collaboration on assignments

We strongly 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 mycourses Discussion Boards.

*However, there are limits to this collaboration.* You can give hints (and the TAs and instructors will give hints sometimes if you ask). However, your 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 higher authorities.

## Supplemental/Deferred Exam

The S/D exam will be held in August. It will cover the same material as the Final Exam and will replace the Final Exam grade. The same “max” rule for midterm and final will apply.

For information on Supplemental Exams, see

<https://www.mcgill.ca/science/student/general/exams/supplemental>.

## 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](https://www.mcgill.ca/study/2017-2018/university_regulations_and_resources/undergraduate/gi_lang_policy)

## 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*