Overview This sheet summarizes information related to CSC2555 (*Mathematical Foundations of Algorithmic Fairness*) during Summer 2026. All times listed below are in the **Eastern time zone**.

Webpage **Please consult the course webpage for full and up-to-date details regarding the course.** It will be frequently updated with announcements, schedule of lectures, and assignments. *You are responsible for reading all the announcements on the course website*; please check at least once a week.

https://www.cs.toronto.edu/~nisarg/teaching/2555y26/

Contact

Instructor	Nisarg Shah
Webpage	cs.toronto.edu/~nisarg/
Email	csc2555-2026-05@cs.toronto.edu (Please use this email for all course related emails)
Office	SF 3312 (Please do not drop by unless you have scheduled an in-person meeting)

Discussion Board Piazza will be the preferred forum for asking questions about class material or other topics that are likely to be of general interest to the class. While it may be quicker than scheduling an office hour with an instructor, please do not expect ultra-quick responses.

https://piazza.com/utoronto.ca/summer2026/csc2555

Delivery **Lectures** will be in-person. Please check the course web page for room information. If you occassionally need to attend remotely, you can do so by joining the following zoom room. However, please note that this is only for exceptional circumstances.

https://utoronto.zoom.us/j/87511650816

Assignments submissions will be via MarkUs (link to be circulated soon).

Office hours will be held by appointment in the instructor's office or via Zoom. Send an email to the course email provided above to set up an appointment.

Assignments & Project There will be **two assignments** worth 20% each to be done individually. MarkUs will be used for assignment submissions. For each assignment, you will be asked to submit a SINGLE PDF. You are strongly encouraged to use LaTeX. Scanned PDF of handwritten solutions will be acceptable, but it is your responsibility to ensure that the handwriting is legible!

Additionally, there will be a **course project** to be done in groups of up to 3. More information about the projects will be released during the course.

Due Dates

Assignment 1	Jun 16
Assignment 2	Jul 14

Grading Scheme

- Assignments: 40% (20% each)
- Project: 50%
- Embedded Ethics Module: 5%
- Class participation: 5%

Late Days

- Each student will receive a total of two (4) late days on MarkUs to be used across the assignments.
- You **do not** earn extra late days for minor illnesses, University activities, or other short disruptions; this is precisely what the late days are for. You are responsible for managing your late days.
- If, for some legitimate reason, you absolutely need more late days, you will need to personally request them from the instructor in advance and with proper documentation.

Learning Goals

By the end of this course, you should be able to:

- recognize and understand various mathematical notions of fairness across a range of algorithmic decision-making problems,
- reason about the relations between different fairness notions,
- prove or disprove that an algorithm satisfies a fairness notion,
- design algorithms satisfying given fairness notions or prove the impossibility thereof,
- understand societal implications of imposing (or not imposing) fairness.

Syllabus

Please note that this syllabus is tentative, and can change as the course progresses.

Week	Topic
1	Introduction; Fair allocation I: divisible goods
2	Fair allocation II: indivisible goods
3-4	Proportional representation in voting
5	Fair matching
6	Bias in machine learning
7	Fair classification: individual fairness
8	Fair classification: group fairness
9	Fair clustering
10	Embedded Ethics Module
11	Project presentations
12	Project presentations

No Garbage Policy

- You will receive 20% of the marks for any question or subquestion on any assignment (**except** for any bonus questions or subquestions) if you do not provide a solution (or clearly scratch off any solution you have written).
- If you write a solution that shows a reasonable understanding of the problem and a viable approach, you will generally receive at least 20% marks (up to the instructor's discretion), so you are always encouraged to write a solution if you believe you understood the question well and have a viable approach (even if you do not have the full answer).
- On the other hand, writing nonsense will get you 0% (or close to that), so if you do not understand the question or have no idea how to approach it, you are better off using the no garbage policy.

Textbooks

The primary reference for this course will be the lecture slides, which will be posted to the course webpage. In addition, please refer to the following textbooks.

- *Handbook of Computational Social Choice* edited by Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia.
- *Fairness and Machine Learning* by Solon Barocas, Moritz Hardt and Arvind Narayanan.
- *The Ethical Algorithm* by Aaron Roth and Michael Kearns.

Citation Policy

You are free to discuss the course content with your peers and with large language models (LLMs), and refer to online materials. However, if you obtain critical insights for an assignment question from a peer, a large language model, or an online source, you *must cite them*, by providing their name, chat link, or URL, respectively.