

Course Syllabus for DS 760: Ethics of Data Science

NOTE: This syllabus document contains the basic information of this course. The most current syllabus is available in the full course.

Course Description

This course focuses on the investigation of ethical issues in data science related to privacy, the use of “black box” learning algorithms, censorship on social media, intellectual property rights, piracy, security, confidentiality, and many other issues. Our study of these issues begins broadly, with a look at ethical issues in computer science at large, and then moves to the narrower field of data science. We will consider ethical arguments and positions, the quality and integrity of decisions and inferences based on data, and how important cases and laws have shaped the legality, if not the morality, of data-science-related computing. Case studies will be used to investigate specific issues.

Course Objectives

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

- Identify and analyze ethical issues in data science.
- Identify and avoid some of the complex and unintended negative impacts that can result from leveraging data resources to provide value in strategic decision-making.
- Interpret and apply a professional code of ethics relevant to the data science profession.
- Speak and write in a clear and logically consistent manner about ethical issues in data science.

Course Components

Papers

You will write a total of three papers: two short ones and one longer one.

The two short papers must each be 3 to 5 pages long, double-spaced. You will be given considerable leeway in choosing your topic, but each must address material covered in a chapter from the textbook and/or in any related reading assigned each lesson. For each paper, you must:

1. Summarize the assigned reading material you are addressing (just the relevant parts). Please do not focus on material that extends beyond the readings. Normally the reading material you focus on will involve the author's, or someone else's opinion or position, which you may or may not agree with. Be sure to include the reasoning that person used in arriving at his or her opinion or position.
2. Defend your own position regarding your chosen topic. It might be that you agree with the person whose position you summarized in section 1, or it might be that you disagree. Try to address the REASONS given in section 1.
3. Consider at least one critique of your position from the point of view of someone who would disagree with what you say in section 2. Again, the critique should focus on REASONS given in defense of your position in section 2. The stronger the critique is, the stronger your paper will ultimately be. This section is the most difficult section to write, and the most important.
4. Respond intelligently to the critique. Again, the strongest response will focus on the reasons given in section 3.

The topic for the short paper must be something that has come up in the readings at the point in the course when the paper is due. For instance, Paper 1 is due at the end of Lesson 4. The topic of the paper must therefore be covered in the readings assigned up through Lesson 4. Paper 2 will cover some material in the readings assigned in Lessons 5, 6, or 7.

The third, longer paper requires you to demonstrate a comprehensive understanding of ethical frameworks. Like the shorter papers, it also requires you to exercise critical reasoning skills. This paper should be approximately 10 to 12 pages long, double-spaced. You must select a current topic in data science from the list of ethical issue in data science that your fellow students have presented. You will give a detailed description of the case. Then describe the major ethical issues related to it -- there may be several, and you may choose just one or several. You will then apply two different moral theories/perspectives to the case, perhaps leading to two conflicting conclusions, though not necessarily. You will also do a professional standards analysis in which you apply a relevant professional code of ethics (sometimes called a professional code of conduct). Ideally you will select a code tied to the data science profession, such as one of the ones that are included in your assigned readings for Lesson 8. Finally, you will

draw a conclusion about what is the moral path in this situation and why. The choice to be made must be “non-trivial” in the sense that each option must have ethically nuanced pros and cons. Finally, you will consider an intelligent and detailed objection to the point of view that you defend and you will thoughtfully respond to it. As in the short papers, in considering an objection and then a reply, it is always best to focus on the reasons given for the positions you disagree with.

To assist in the organization of this paper, the details of the format are spelled out for you in the Long Paper assignment.

Grading Policy

Your mastery of course content is assessed using a variety of methods:

Activity	Percentage of Grade
Two Short Papers (3-5 pages)	25%
Peer Review: Critical Analysis of Long Paper (10-12 pages)	20%
Discussions and Peer Review of Long Papers	10%
Response to Ethical Issue in Data Science Presentations	10%
Lesson Comprehension Check Questions	10%
Individual Questions	20%
Ethical Issue in Data Science Presentation	5%
Total	100%

Final grades are assigned using the following scale:

A 90-100%

B 80-89%

C 70-79%

D 60-69%

F At or below 59%