

Ethics Module for 15-110 Principles of Computing

Instructor's Guide

Mike Tasota, Ramesh Nagarajah
16-735 (Spring 2021)

Table of Contents

Table of Contents	2
Cover Sheet	3
Selected Reading and Additional Links	4
Class Outline: Unit Reviews	6
Class Outline: Ethics Lecture	9
Assessment Material	11

Cover Sheet

This instructor's guide includes guidance on an ethics module crafted for a lecture-centric, large class size, broad topic, introductory course in computing. This course is not specific to CS students and is implemented to assist students who are instead more likely to work closely alongside CS professionals or in a management role of CS.

The ethics module design was based on the limited time available per class to focus on additional content, with time cut out for one dedicated full lecture. The lecture-heavy design of the course and limited extra time made implementing active learning a challenge. These opportunities came in the form of discussion questions included in each of the unit review ethics slides, with an opportunity to invite students to continue the discussion outside of class, possibly using discussion boards or Piazza. An active learning activity is also included at the end of the ethics lecture.

The guide includes learning objectives for the module, a suggested pre-class reading (plus additional readings), a breakdown of the unit review ethics slides and ethics lecture with a timeline and tips, an active learning activity description, and assessment questions (with answer key and grading rubrics) for a final exam.

Learning Objectives: At the end of the entire module, students should be able to...

- Critically reflect on individual ethical concerns in relation to Computer Science
- Highlight examples of ethics-related instances in the real world
- Understand the significance of ethics from the perspective of the students' own academic disciplines and desired careers
- Engage in discussions of ethical questions with peers

Zipped Archive Contents: This document should be included within a zipped archive containing the following documents:

- **Instructor Guide.pdf** - PDF document (this document)
- **Presentation.pdf** - PDF document; helps facilitate the use of this ethics module by other courses by setting the stage and describing the design from 15-110. Other courses should consider more applicable unit review topics for their own needs, if necessary.
- **Ethics Lecture** - Directory containing the ethics lecture slides, in both PDF and powerpoint formats
- **Unit Reviews** - Directory containing ethics reflection slides to be amended to the unit reviews for units 1 through 4, in both PDF and powerpoint formats

Selected Reading and Additional Links

The instructor is provided with supplemental readings sparsely into each of the four unit reviews as hyperlinks, a final page in each unit review containing links for additional reading, as well as supplemental academic readings provided in a folder included with all other final deliverables.

Class Pre-Reading: The link for the suggested required pre-class reading for the ethics lecture is below.

Thomsen, S. (2020, June 09). When algorithms go wrong: How relying on automated tech is a world of pain for governments. Retrieved April 29, 2021, from <https://www.startupdaily.net/2020/06/when-algorithms-go-wrong-how-relying-on-automated-tech-is-a-world-of-pain-for-governments/>

Additionally, the complete set of supplemental links are below. These include all hyperlinks within the amended unit review slides as well as the ethics lecture slides.

Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. Retrieved April 30, 2021, from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

Dizikes, P. (n.d.). Study: On twitter, false news travels faster than true stories. Retrieved April 30, 2021, from <https://news.mit.edu/2018/study-twitter-false-news-travels-faster-true-stories-0308>

Dreyfuss, E. (n.d.). Facebook is changing news feed (again) to stop fake news. Retrieved April 30, 2021, from <https://www.wired.com/story/facebook-click-gap-news-feed-changes/>

Feinstein, B., & Werbach, K. (2021, April 14). Don't fear cryptocurrencies. manage them. Retrieved April 30, 2021, from <https://www.nytimes.com/2021/04/14/opinion/coinbase-ipo-cryptocurrencies.html>

Harford, T. (2012, August 10). High-frequency trading and the \$440m Mistake. Retrieved April 30, 2021, from <https://www.bbc.com/news/magazine-19214294>

A history of ransomware attacks: The biggest and worst ransomware attacks of all time. (2020, December 01). Retrieved April 30, 2021, from <https://digitalguardian.com/blog/history-ransomware-attacks-biggest-and-worst-ransomware-attacks-all-time>

Leaked: Cambridge Analytica's blueprint for Trump victory. (2018, March 23). Retrieved April 30, 2021, from

<https://www.theguardian.com/uk-news/2018/mar/23/leaked-cambridge-analytica-blueprint-for-trump-victory>

McCausland, P. (2019, November 11). Self-driving uber car that hit and killed woman did not recognize that pedestrians jaywalk. Retrieved April 30, 2021, from <https://www.nbcnews.com/tech/tech-news/self-driving-uber-car-hit-killed-woman-did-not-recognize-n1079281>

Metz, C. (2021, March 15). Who is making sure the A.I. machines aren't racist? Retrieved April 30, 2021, from <https://www.nytimes.com/2021/03/15/technology/artificial-intelligence-google-bias.html>

Revealed: 50 million Facebook profiles harvested for Cambridge analytica in major data breach. (2018, March 17). Retrieved April 30, 2021, from <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>

Roose, K. (2019, August 04). 'Shut the Site Down,' says the creator of 8chan, a megaphone for Gunmen. Retrieved April 30, 2021, from <https://www.nytimes.com/2019/08/04/technology/8chan-shooting-manifesto.html>

Roose, K., Isaac, M., & Frenkel, S. (2020, November 24). Facebook struggles to balance civility and growth. Retrieved April 30, 2021, from <https://www.nytimes.com/2020/11/24/technology/facebook-election-misinformation.html>

Smith, C. (2021, February 23). A.I. here, there, everywhere. Retrieved April 30, 2021, from <https://www.nytimes.com/2021/02/23/technology/ai-innovation-privacy-seniors-education.html>

Team, I. (2019, June 05). Forbes insights: Managing the ethics of algorithms. Retrieved April 30, 2021, from <https://www.forbes.com/sites/insights-intelai/2019/03/27/managing-the-ethics-of-algorithms/?sh=6494c0e73481>

Class Outline: Unit Reviews

Below are class outlines for the slides to be amended to the class unit reviews. They are designed to be no longer than 10 minutes for each unit review leading up to the ethics lecture.

The topics for the ethics reflection in each unit review was based on lecture topics presented in that unit. The mapping of relevant topics are as follows:

Unit Topics	Related Ethics Reflection
Unit 1: Intro to Programming & Algorithms	AI Explainability
Unit 2: Data Structures	Facebook-Cambridge Analytica data scandal
Unit 3: Internet, Encryption	Encryption and the impact of anonymity online
Unit 4: Concepts in AI & ML	MIT Study on Misinformation & Social Media

Notice the outline for the first review is the longest of the set because it also provides an introductory slide.

Unit 1 Review (10 minutes)	Description/Important Notes
Ethics Overview (3 min)	This section is intended to introduce students to ethics and its relevance/application to computer science. This is to ensure that students are thinking of ethics early in the semester rather than waiting until the full lecture at the end of the course.
Ethics Reflection: AI Explainability (2 min)	Instead of focusing on the rudimentary coding topics that are covered in this early unit, we explain how once these seemingly simple skills are used to develop complex technology, like AI, comprehension and explainability become increasingly essential. Additionally, AI and explainability remains a crucial ethical question for our society as AI becomes more widespread in areas of significant impact. Invite students curiosity.
Discussion (5 min)	If time permits, allow students to break into small groups and discuss these questions in order to support active learning. Additionally,

<ul style="list-style-type: none"> As a society, should we preference safety and caution over rapid innovation? How would this look in terms of AI development and implementation? Can you imagine any scenarios where a “black box AI” would cause significant social or legal problems? 	<p>once the class reforms, invite some students to share their thoughts and comment on those of others. Following the class, open a discussion board or Piazza thread and invite students to engage with each other on thoughts they were not able to share in class.</p>
---	---

Unit 2 Review (7 minutes)	Description/Important Notes
Ethics Reflection: data structures & access (2 min)	<p>Highlights how the FB-CA data scandal relates to data structures: the CA app would collect the FB data of its users, plus all of their friends.</p> <p>Connects with concepts from Graphs lecture (“All Neighbors of a Node == All FB Friends of a FB User”)</p>
Discussion (5 min) <ul style="list-style-type: none"> Should technology services be responsible for disclosing to their users what data is being collected and/or how data is used? 	Other possible topics of discussion: app platform policies, targeted advertising and microtargeting , 2016 US Elections

Unit 3 Review (9 minutes)	Description/Important Notes
Ethics Reflection: encryption (4 min) <ul style="list-style-type: none"> Tor/VPN Anonymous forums Cryptocurrencies, Ransomware 	<p>Highlights a few examples of encryption innovations and how they have impacted society.</p> <p>Reference to the 2021 documentary is included as a resource to learn more about how the imageboard 8chan relates to the online communities and controversies mentioned (Pizzagate, Gamergate, QAnon)</p>
Discussion (5 min) <ul style="list-style-type: none"> What are some other ways in which online anonymity can create a positive or negative impact on society? 	Other possible topics of discussion: censorship by governments, expectation of privacy in US

<ul style="list-style-type: none"> Should users be allowed a reasonable expectation of privacy online? What is the responsibility of online companies to disclose when this is not the case? 	
---	--

Unit 4 Review (8 minutes)	Description/Important Notes
Ethics Reflection: Misinformation and MIT Study (3 min)	This section focuses on the spread of misinformation as a consequence of heuristic and algorithm design. It quickly pivots to social media as a real world example where ethical questions are being asked and discussed regularly.
Discussion (5 min) <ul style="list-style-type: none"> The algorithms on social media sites demonstrably preference misinformation. Should companies like Facebook and Twitter be legally responsible for the spreading of misinformation that occurs on their sites? 	This is a hotbed topic in conversations about regulation, accountability, and general misinformation in our society (“fake news”). It likely will encourage a lot of discussion from students. Consider blocking out as much time as possible to talk about this. Following the class, open a discussion board or Piazza thread and invite students to engage with each other on thoughts they were not able to share in class.

Class Outline: Ethics Lecture

Below is the class outline for the main ethics module lecture. This spans the entire 50-minute lecture, including a 10-minute activity to facilitate active learning as well as a 4-minute wrapping up at the end of the lecture.

Section	Description/Important Notes
<ul style="list-style-type: none"> ● Title Slide, Learning Goals (2 min) 	<p>Introduces students to the lecture and its purpose</p>
<ul style="list-style-type: none"> ● Ethics in Computer Science (3 min) 	<p>A brief description of the importance of ethics in computer science, highlighting responsibility.</p>
<ul style="list-style-type: none"> ● Data (10 min) <ul style="list-style-type: none"> ○ Collecting Data ○ Data Economy <ul style="list-style-type: none"> ■ Questions to consider <ul style="list-style-type: none"> ■ <i>Who are the stakeholders in the data economy?</i> ■ <i>Is this a fair representation of those who participate in it?</i> ■ <i>Should the people whose data is being used have ownership and a say over its usage?</i> 	<p>A section covering data collection and the data economy. A set of questions is offered for students to reflect on, consider inviting students to share thoughts.</p> <p>Slide 6 suggests that students check out the data their web browser is sharing via the provided link: https://webkay.robinlinus.com/, consider inviting students to share their surprises/concerns, if time allows.</p>
<ul style="list-style-type: none"> ● Bias (10 min) <ul style="list-style-type: none"> ○ Data Bias ○ Facial Recognition ○ Other AI Algorithms 	<p>This section focuses on bias in machine learning and AI resulting from data and how that can bear real world impacts - citing issues with facial recognition and algorithms such as those used for hiring.</p>
<ul style="list-style-type: none"> ● Misinformation (3 min) 	<p>This brief section makes a comment on misinformation resulting from algorithms and data collection on social media.</p> <p>This slide originally also referenced the MIT study on fake news but was removed since it is now covered in the unit 4 ethics reflection.</p>
<ul style="list-style-type: none"> ● Responsibility & AI (3 min) 	<p>This slide returns to discussion about ethical AI design, emphasizing the earlier definition of ethics offered about ““What should we</p>

	value broadly as a society? And what should we do to realize those values?”. This section prompts the active learning activity and references content from the first unit review on explainability in AI.
<ul style="list-style-type: none"> AI Algorithm Design Activity (10 min) 	The design activity centers around AI adoption and the considerations that students would weigh now that they have learned about ethics. This is intended to be an active learning component of the lecture. See a detailed description and guide later in the lecture.
<ul style="list-style-type: none"> Learning Goals, Wrapping up (4 min) 	To close the lecture, the instructor circles back to the initially stated learning goals for the lecture to emphasize the purpose of the lecture and what students should have gotten out of it.

Notes on Active Learning Activity (slide 16 of Ethics Lecture)

<p>Students will work in small groups from the prompt stated in the slides:</p> <p><i>Suppose that you are part of a team that is deciding whether to adopt an AI hiring algorithm for your company to choose potential job candidates.</i></p> <p><i>What questions might you have for the designers?</i></p> <p><i>What concerns might you have about the data used by the algorithm?</i></p> <p><i>How will you ensure the adoption of the technology is not causing harm or reinforcing existing structural inequalities?</i></p>
<p>Here are some questions that the instructor might consider to hear from the students:</p> <ul style="list-style-type: none"> - What data will be used by the algorithm? - How will the data be used? - How might you deal with bias in the data? - Questions related to ethics and management for AI designers

Assessment Material

We were encouraged by the course professor to be considerate of the current course workload and, rather than crafting additional assignments, to work in opportunities for her to consider new assignments and draft ethics questions for the final exam. We were asked to avoid any extra credit or additional homework.

Questions were crafted in support of our learning goals for the students experiencing the ethics module, with a hope that the content we offer and assess will emphasize the need to consider ethics in computer science, and that they feel better prepared to engage with ethics in their sphere of work and when collaborating in the CS space. We put together five multiple choice questions, which are able to be auto-graded through Canvas, as well as two short answer questions with provided rubrics for TAs to efficiently grade. Multiple choice questions were based on the ethics content covered in the ethics review slides and the ethics lecture. Short answer questions focus on content covered in the main ethics lecture.

The assessment material is included below.

Multiple Choice/Answer Questions

- 1) What was the main idea behind the Facebook-Cambridge Analytica Scandal?
 - a. The account credentials (including passwords) of 50 million Facebook users were compromised
 - b.** An app developed by Cambridge Analytica collected the data of its users along with all their Facebook friends' data, including friends who did not even use the app
 - c. Hackers gained access to Facebook users' accounts and threatened to release their personal data unless a ransom was paid
 - d. Facebook filed a lawsuit against Cambridge Analytica for breach of terms of the Facebook platform policy

- 2) The main technology that ransomware implements to lock victims out of their data is _____.
 - a.** Asymmetric Encryption
 - b. Virtual Private Networks (VPNs)
 - c. Man-in-the-middle Attack
 - d. DDOS Attack

- 3) The bias against those with darker skin discovered in machine learning algorithms for facial recognition software was mostly the result of:
 - a. Prior existing biases in law enforcement
 - b.** Two popular training sets that were overwhelmingly composed of lighter-skinned subjects.
 - c. User misimplementation rather than from the training data
 - d. A failure to recognize the limits of the software itself.

- 4) An MIT study found that true stories on Twitter:
- Spread at the same rate as false stories.
 - Spread 6x slower than false stories.
 - Bear less weight on opinions of users than false ones
 - Are more likely to be promoted by the platform's algorithm
- 5) Explainability in AI refers to:
- The ability for AI to communicate with other AI and cooperate
 - The solutions produced by an AI process of getting there being understandable by humans
 - The ability for an AI to explain, in human terms, how it reaches its decisions
 - Being able to explain to non-technologists how AI works

Short Answer Questions

- 1) Briefly describe one way that bias can emerge in a machine learning algorithm.

Rubric: Grant full credit if the student, in some form, mentions any of the following:

- Unproportionate representation of subpopulations
- A failure to match reality
- A reflection of the biases of the designers
- A poor sampling method for the derived data
- Training data is unrepresentative of the entire population or set of scenarios

- 2) In class, we mentioned that data has become the economy of the Internet. Name one of the many ways that companies profit from data collected on the internet.

Rubric: Grant full credit if the student, in some form, mentions any of the following:

- Advertising
- Sell data to other companies
- Aggregate data