

CPSC 310: Syllabus for Spring 2025 (Version 5)

Joan Feigenbaum
www.cs.yale.edu/homes/jf

Artur Pericles Lima Monteiro
jackson.yale.edu/person/artur-pericles-lima-monteiro/

This document contains the schedule for lectures, homework assignments, and exams and a tentative set of reading assignments for the Spring 2025 offering of CPSC 310.

Twenty-first century societies are faced with both threats and opportunities that combine sophisticated computation with politics and international relations in critical ways. Examples include online disinformation campaigns, the tension between online privacy and lawful surveillance, AI regulation and transparency, net neutrality, and digital copyright. CPSC 310 examines some of the political challenges wrought by massive increases in the power of information and communication technologies and the potential for citizens and governments to harness those technologies to solve problems. It is co-taught by Joan Feigenbaum (SEAS, CPSC Department) and Artur Pericles Lima Monteiro (Jackson School & Yale Law School).

The prerequisite for CPSC 310 is CPSC 223 or the equivalent. There will be two lectures and one section per week. Students will be required to complete four written homework assignments, each of which may include both programming exercises and short essays; each HW will account for 12% of the course grade. There will be two in-class exams, each of which will account for 16% of the course grade, but no final exam during exam week. Approximately 20% of a student's course grade will be determined by his or her in-person attendance in both lectures and sections and active participation in sections.

Throughout the semester, the policy for late HWs submitted without Dean's extensions is that, for seven days after the due date, a 5% per day penalty is imposed. After seven days, the HW is no longer accepted, and the student receives a grade of zero for the assignment. There will be **no exceptions** to this late-HW rule; in order to be able to submit a HW assignment after the due date and not be penalized, you **must** get a Dean's extension.

In previous years, CPSC 310 and PLSC 369 / CPSC 210 / EP&E 365 were taught in an overlapping fashion with common lectures and separate sections. PLSC 369 / CPSC 210 / EP&E 365 will not be taught this year. If you have already taken it for credit, you may not take CPSC 310 for credit this year or in the future.

Schedule and reading assignments (subject to revision)

January 14, 2025: Lecture 1 (Feigenbaum and Monteiro): Administrative Matters and “Teaser”

January 16, 2025: Lecture 2 (Feigenbaum): Internet Background

Reading assignment:

David D. Clark, [“The Design Philosophy of the DARPA Internet Protocols.”](#) *ACM SIGCOMM Computer Communication Review*, Vol. 18, No. 4, pp. 106–114, August 1988.

J. H. Salter, D. P. Reed, and D. D. Clark, [“End-To-End Arguments in System Design.”](#) *ACM Transactions on Computer Systems*, Vol. 2, No. 4, pp. 277–288, November 1984.

January 21, 2025: Lecture 3 (Feigenbaum): Internet and Cryptography&Security Background

Reading assignment:

Marjory S. Blumenthal and David D. Clark, [“Rethinking the Design of the Internet: The End-to-End Arguments vs. The Brave New World.”](#) *ACM Transactions on Internet Technology*, Vol. 1, No. 1, pp. 70–109, August 2001.

David Wagner *et al.*, Chapters [1 \(Security Principles\)](#), [5 \(Introduction to Cryptography\)](#), [6 \(Symmetric-Key Cryptography\)](#), and [7 \(Cryptographic Hashes\)](#) of [Computer Security](#).

Optional: A more detailed version of the material in the “Security Principles” chapter above can be found in Matt Bishop, [Chapter 1 \(Introduction\)](#) of **Computer Security: Art and Science**, Pearson Education, Inc., 2019.

January 23, 2025: Lecture 4 (Feigenbaum & Monteiro): Cryptography&Security Background

Reading assignment:

Whitfield Diffie and Martin E. Hellman, [“New Directions in Cryptography.”](#) *IEEE Transaction on Information Theory*, Vol. IT-22, No. 6, pp. 644–654, November 1976. (Sections I, II, III, and VII are required; the rest are optional.)

Arvind Narayanan, [“What Happened to the Crypto Dream?. Part 1.”](#) *IEEE Security & Privacy*, Vol. 11, No. 2, pp. 75–76, March/April, 2013.

Daniel J. Solove, [“The Nothing to Hide Argument,”](#) Chapter 2 of [Nothing to Hide: The False Tradeoff between Privacy and Security](#), Yale University Press, 2011.

Topic 1: Privacy

January 28, 2025: HW1 posted

January 28, 2025, Lecture 5 (Monteiro): Privacy Overview

Reading assignment: Neil M. Richards, [“What Privacy Is,”](#) Pages 17 –34 of [Why Privacy Matters](#), Oxford University Press, Published Online November 2021; Published in print March 2022.

January 30, 2025, Lecture 6 (Monteiro): Data Privacy Online

Reading assignment: Solon Barocas and Helen Nissenbaum. "[Big Data's End Run around Anonymity and Consent](#)," Chapter 2 in **Privacy, Big Data, and the Public Good: Frameworks for Engagement**, Cambridge: Cambridge University Press, 2014.

February 4, 2025, Lecture 7 (Feigenbaum): Encryption vs. Lawful Surveillance, part 1

Reading assignment:

Susan Hennessey and Benjamin Wittes, "[Apple is selling you a phone, not civil liberties](#)," February 18, 2016.

Steven M. Bellovin, *et al.*, "[Op-ed: Ray Ozzie's crypto proposal—a dose of technical reality](#)," *Ars Technica*, May 7, 2018.

Joan Feigenbaum, "[Encryption and Surveillance: Why the law-enforcement access question will not just go away](#)," *Communications of the ACM*, Vol. 62, No. 5, pp. 27–29, May 2019.

February 6, 2025, Lecture 8 (Monteiro): Encryption vs. Lawful Surveillance, part 2

Reading assignment:

Hal Abelson *et al.*, "[Keys under doormats: mandating insecurity by requiring government access to all data and communications](#)," *Journal of Cybersecurity*, Vol. 1, No. 1, pp. 69–79, 2015.

Hal Abelson *et al.*, "[Bugs in our pockets: the risks of client-side scanning](#)," *Journal of Cybersecurity*, Vol. 10, Issue 1, pp. 1–18, 2024.

February 7, 2025: HW1 due

February 11, 2025: HW2 posted

February 11, 2025, Lecture 9 (Monteiro & Feigenbaum): Privacy, Surveillance, & Intermediaries

Viewing and reading assignment:

Michael Veale, Youtube video: "[Data is dead](#)," February 2024.

Alan Z. Rozenstein, "[Surveillance Intermediaries](#)," *Stanford Law Review*, Vol. 70, pp. 99–189, January 2018. (Intro and Sections I, II, and IV are required; the rest is optional.)

Joseph Menn, "[U.K. orders Apple to let it spy on users' encrypted Accounts](#)," *Washington Post*, February 7, 2025.

February 13, 2025, Lecture 10 (Feigenbaum): Anonymity

Reading assignment:

Gary T. Marx, "[Identity and Anonymity](#)," in J. Caplan and J. Torpey, **Documenting Individual Identity**, Princeton University Press, 2001.

David Wagner *et al.*, “[Anonymity and Tor](#),” Chapter 39 of [Computer Security](#).

Optional: A more detailed and precise explanation of the material in Chapter 39 above can be found in the second and third sections (entitled “Onion Routing and Tor” and “Attacks on Onion Routing”) of [this article](#).

Topic 2: Platforms (esp. social media)

February 18, 2025, Lecture 11 (Monteiro): Introduction to Platform Governance

Reading assignment:

Jack Balkin, “[How to Regulate \(and Not Regulate\) Social Media](#),” *Journal of Free Speech Law*, Vol. 1, No. 1, pp. 71–96, 2021.

Kate Klonick, “[The New Governors: The People, Rules, and Processes Governing Online Speech](#),” *Harvard Law Review*, Vol. 131, pp. 1598–1670, April 2018. (Pages 1625–1657 are required; the rest is optional.)

Tarleton Gillespie, **Custodians of the Internet: Platform Content Moderation, and the Hidden Decisions that Shape Social Media**, Yale University Press, June 2018. PDF version available [here](#). (Chapters 1, 2, and 4 are required; the rest optional.)

February 20, 2025, Lecture 12 (Feigenbaum): Misinformation and Polarization

Reading assignment:

Christopher A. Bail *et al.*, “[Exposure to opposing views on social media can increase political polarization](#),” *Proceedings of the National Academy of Science*, Vol. 115, No. 37, pp. 9216–9221, September 2018.

Killian L. McLoughlin *et al.*, “[Misinformation exploits outrage to spread online](#),” *Science*, Vol. 386, Issue 6725, pp. 991–996, November 2024.

Rachel Kleinfeld, [Polarization, Democracy, and Political Violence in the United States: What the Research Says](#), Carnegie Endowment for International Peace, September 2023. (Summary, Introduction (pp. 1–7), “Media Bubbles” (pp. 25–27), and Conclusion (pp. 40–46) are required; the rest is optional.)

February 21, 2025: HW2 due

February 25, 2025: Lecture 13 (Monteiro): Recommender Systems

Reading assignment:

Robert Gorwa, Reuben Binns, and Christian Katzenbach. “[Algorithmic Content Moderation: Technical and Political Challenges in the Automation of Platform Governance](#),” *Big Data & Society*, Vol. 7, Issue 1, pp. 1–15, January–June 2020.

Arvind Narayanan, [Understanding Social Media Recommendation Algorithms](#), Algorithmic Amplification and Society, New York: Knight First Amendment Institute at Columbia University, March 9, 2023.

Tom Cunningham *et al.*, [“What We Know about Using Non-Engagement Signals in Content Ranking.”](#) arXiv.org/pdf/2402.06831.

February 27, 2025: Exam 1

March 4, 2025, Lecture 14 (Feigenbaum): Accountability: Technical vs. Political

Reading assignment:

Deven R. Desai and Joshua A. Kroll, [“Trust But Verify: A Guide to Algorithms and the Law,”](#) *Harvard Journal of Law and Technology*, Vol. 31, No. 1, Winter 2018. (Read Chapters I, II, IV, and V. Skim Chapter VI.)

Aleksandra Kuczerawy and Jef Ausloos, [“From Notice-And-Takedown To Notice-And-Delist: Implementing Google Spain.”](#) *Colorado Technology Law Journal*, Vol. 14, No. 2, pp. 219–258, 2016. (Pages 220-246 are required; the rest is optional.)

March 6, 2025: Last class before midterm drop date. Grade predictions posted.

March 6, 2025, Lecture 15 (Feigenbaum): Balance of Power Between Users & Platforms, part 1

Reading assignment:

Ayelet Gordon-Tapiero, Alexandra Wood, and Katrina Ligett, [“The Case for Establishing a Collective Perspective to Address the Harms of Platform Personalization,”](#) in *Proceedings of the 2022 ACM Symposium on Computer Science and Law*.

Caleb Malchik and Joan Feigenbaum, [Toward User Control over Information Access: A Sociotechnical Approach](#), in *Proceedings of the 2022 New Security Paradigms Workshop*. (Read Abstract, Sec. 1, and Sec. 3 in full. Skim Sec. 2, paying attention to Subsec. 2.4.)

March 7, 2025: Spring break begins

March 25, 2025: First class after Spring break

March 25, 2025, Lecture 16 (Monteiro): Balance of Power Between Users & Platforms, part 2

Reading and listening assignment:

Francis Fukuyama, [“Making the Internet Safe for Democracy.”](#) *Journal of Democracy* Vol. 32, No. 2, pp. 37–44, 2021.

Mike Masnick, [“Protocols, Not Platforms: A Technological Approach to Free Speech.”](#) Free Speech Futures. New York: Knight First Amendment Institute at Columbia University, August 21, 2019.

Vergecast (podcast) with Bluesky CEO, Jay Graber, <https://www.theverge.com/2024/3/25/24108872/bluesky-ceo-graber-federation-social-media-decoder-interview>

Topic 3: AI Governance

March 27, 2025, Lecture 17 (Monteiro): What is AI, and what is AI governance?

Reading assignment:

Michael Veale, Kira Matus, and Robert Gorwa, "[AI and Global Governance: Modalities, Rationales, Tensions](#)," *Annual Review of Law and Social Science*, Vol. 19, No. 1, pp. 255–275, 2023.

Frank Pasquale, *The Second Wave of Algorithmic Accountability*, 2019, <https://lpeproject.org/blog/the-second-wave-of-algorithmic-accountability/>

March 29, 2025: HW3 posted

April 1, 2025, Lecture 18 (Feigenbaum): Introduction to Machine Learning

Reading and viewing assignment:

Solon Barocas and Andrew D. Selbst, "[Big Data's Disparate Impact](#)," *California Law Review*, Vol. 104, pp. 671–732, 2016.

Arvind Narayanan, "[21 Fairness Definitions and Their Politics](#)" (YouTube video).

April 3, 2025: Lecture 19 (Monteiro): AI Accountability, part 1: Transparency & Interpretability

Reading assignment:

Mike Ananny and Kate Crawford, "[Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability](#)," *New Media & Society*, Vol. 20, No. 3, pp. 973–89, 2018.

Tim G. J. Rudner and Helen Toner, [Key Concepts in AI Safety: Interpretability in Machine Learning](#), 2021.

April 8, 2025: HW3 due

April 8, 2025, Lecture 20 (Guest lecturer: A. Feder Cooper): Introduction to Generative AI

Reading assignment: A. Feder Cooper *et al.*, "[Machine Unlearning Doesn't Do What You Think: Lessons for Generative AI Policy, Research, and Practice](#)," <https://arxiv.org/abs/2412.06966>.

Optional: Katherine Lee, A. Feder Cooper, and James Grimmelmann, "[Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain](#)." Unpublished draft: Do not redistribute.

April 10, 2025, Lecture 21 (Monteiro): AI Accountability, part 2: Ethics, Trust & Risk Management

Viewing, reading, and listening assignment:

Philippe Hacker, "[The Finalised EU AI Act: Implications for Businesses, Engineers and Entrepreneurs](#)" (Youtube video).

Luke Munn, "[The Uselessness of AI Ethics](#)," *AI and Ethics*, Vol. 3, pp. 869–877, 2023.

Lawfare (podcast), Margot Kaminski, Regulating AI Risks, April 20, 2023,
<https://shows.acast.com/lawfare/episodes/margot-kaminski-on-regulating-ai-risks>

April 12, 2025: HW4 posted

April 15, 2025, Lecture 22 (Guest lecturer: Anat Lior): AI Insurance

Reading assignment:

Anat Lior, "[AI Entities as AI Agents: Artificial Intelligence Liability and the AI Respondeant Superior Analogy](#)," *Mitchell Hamline Law Review*, Vol. 46, Issue 5, Article 2, 2020. (Sections I, III & IV are required; the rest is optional.)

Anat Lior, "[Insuring AI: The Role of Insurance in Artificial Intelligence Regulation](#)," *Harvard Journal of Law and Technology*, Vol. 35, No. 1, Spring 2022. (Chapters I, II, III, V.B are required; the rest is optional.)

April 17, 2025, Lecture 23 (Feigenbaum): Antitrust

Reading assignment: Lina Khan, "[Amazon's Antitrust Paradox](#)," *The Yale Law Journal*, Vol. 126, No. 3, pp. 710–805. (Introduction; Sections 1, 2, 3, and 6; and Conclusion are required; the rest is optional.)

April 22, 2025: HW4 due

April 22, 2025, Lecture 24 (Feigenbaum & Monteiro): What's been happening, and what's next?

Reading assignment: Pamela Samuelson, "[Generative AI Meets Copyright](#)," *Science*, Vol. 381, Issue 6654, July 2023, pp. 158–161.

April 24, 2025: Exam 2