

CPSC 611 / GLBL 6115
Topics in Computer Science and Global Affairs
Fall 2022
Wednesday, 3:30-5:20pm
WLH Room 011

Joan Feigenbaum

Arthur K. Watson Hall, 51 Prospect Street, Room 512

Joan.feigenbaum@yale.edu

Assistant: Judi Paige (judi.paige@yale.edu)

Office hours by appointment. Please contact Ms. Paige to schedule one.

Edward (Ted) Wittenstein

Alwin Hall, 31 Hillhouse Avenue, Room 201

edward.wittenstein@yale.edu

Assistant: Enit Colon (enit.colon@yale.edu)

Office hours in Alwin Hall on Fridays, 2:00-4:30pm and by appointment

Teaching Assistant

Aidan Evans (aidan.evans@yale.edu)

Course Description:

This course focuses on “socio-technical” problems in computing and international relations. These are problems that cannot be solved through technological progress alone but rather require legal, political, or cultural progress as well. Offered jointly by the SEAS Computer Science Department and the Jackson School of Global Affairs, this graduate-level seminar is designed to help bridge the divide across the law, technology, and policy communities at Yale, focusing on four key challenges at the intersection of computer science and global affairs: (1) disinformation; (2) cyberespionage; (3) encryption; and (4) artificial intelligence.

The course is aimed at both STEM graduate students who desire greater exposure to the legal, policy, and ethical dimensions of their research and non-STEM graduate students seeking greater technical fluency. Students engage in interactive discussion, explore socio-technical challenges from diverse perspectives, and collaborate in interdisciplinary teams throughout the semester.

Pandemic Expectations:

The course will meet in-person consistent with university health guidelines. Attendance and participation are required. Every effort will be made to accommodate pandemic-related challenges.

Enrollment:

Enrollment is limited to 18 students. If fewer than 18 graduate students enroll, remaining spots will be open to Yale undergraduates with preference given to CPSC and GLBL majors. Undergraduates seeking course admission should submit a brief statement of interest (max 150 words).

Prerequisites:

There are no prerequisites for this course. Background in the basics of cryptography and computer security (as covered in Yale's CPSC 467), networks (as covered in Yale's CPSC 433), and databases (as covered in Yale's CPSC 437) is helpful but not required.

Course Requirements and Grading:

- 1. *Leading Weekly Discussion (20%)*:** For Weeks #2-10 of the semester, students will lead the seminar discussion, highlighting key questions raised in the assigned readings and facilitating conversation among classmates and professors. Depending on enrollment, discussion leaders may work in pairs. Students may present briefly on their impressions of the readings at the onset of class, but their primary responsibility is to flag questions for group consideration. Students are encouraged meet with Professors Feigenbaum and Wittenstein in advance to review their approach to leading their assigned seminar.
- 2. *Final Project Presentation (20%) and Final Project (40%)*:** Students will work individually or in small groups (**max 3 people**) on a semester-long project, in which they prepare a 15-20 minute class presentation and final project that addresses some aspect of the course's four primary focus areas: (1) disinformation; (2) cyberespionage; (3) encryption; and (4) artificial intelligence. The presentations will occur during Weeks #10-13 of the semester, serving as an opportunity for student and instructor feedback. The final project can take the form of a written paper (max 4000 words) or computer artifact. A project proposal (max 1000 words) is due by **11:00pm on Wednesday, October 5, and the final project is due at 11:00pm on Wednesday, December 14.**
- 3. *Attendance & Participation (20%)*:** Active participation is essential to successful completion of the course, and student attendance is essential. **Students with subpar attendance and/or participation will not receive an "A" in this course.**

Readings:

Wednesday, August 31, Week #1: Course Overview and Introductions

Wednesday, September 7, Week #2: Disinformation Part I: Technical Dimensions of DeepFakes and Influence Operations

- Ferenc Huszár et al., "Algorithmic Amplification of Politics on Twitter" (October 21, 2021).
- Laura Courchesne et al., "Review of Social Science Research on the Effects of Influence Operations," *Empirical Studies of Conflict* (July 17, 2021).
- Matthew Groh et al., "Deepfake Detection by Human Crowds, Machines, and Machine-informed Crowds," *PNAS* (January 2022).
- Yale Cyber Leadership Forum, "Detecting DeepFakes and Coordinated Inauthentic Behavior Online" (March 4, 2021).
- MIT Media Lab: [Take the DeepFake Detection Quiz](#)

Wednesday, September 14, Week #3: Disinformation Part II: Countering Influence Operations: Law and Policy Considerations

- Tarleton Gillespie, *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media* (2018) [**only** chapters 1, 4, 7, and 8 are required reading]
- Dexter Roberts, “China’s Disinformation Strategy: Its Dimensions and Future,” *Atlantic Council* (December 2020).
- Bobby Chesney and Danielle Citron, “Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security,” *California Law Review* (Vol. 107, 2019).
- Ian Bogost and Alexis C. Madrigal, How Facebook Works for Trump, *The Atlantic* (April 17, 2020).
- Katerina Sedova et al., “AI and the Future of Disinformation Campaigns,” *CSET* (December 2021).

Wednesday September 21, Week #4: Cyberespionage Part I: International and American Constitutional Law

- Russell Buchan, *Cyber Espionage and International Law* (2018) [**only** chapters 1 and 6 are required reading: *scans to be provided*].
- *Is the US Government’s use of Section 702 of the PATRIOT Act an example of “good surveillance”?*
 - Office of the Director of National Intelligence, Section 702 Overview
 - Office of the Director of National Intelligence and Office of Civil Liberties, Privacy, and Transparency, Guide to Section 702 Value Examples (October 2017).
 - National Security Agency/Central Security Service, "Section 702" Saves Lives, Protects the Nation and Allies (December 12, 2017).
 - Jim Dempsey, Section 702 Renewal: Opportunities Lost and Gained (January 29, 2018).

Wednesday, September 28, Week #5: Cyberespionage Part II: Russian and Chinese Strategy

- Lennart Maschmeyer, “The Subversive Trilemma: Why Cyber Operations Fall Short of Expectations,” *International Security* 46(2) (2021).
- *Russian Cyberespionage: SolarWinds*
 - Dina Temple-Raston, “A ‘Worst Nightmare’ Cyberattack: The Untold Story of the SolarWinds Hack,” *NPR* (April 16, 2021).
 - The White House, “Imposing Costs for Harmful Foreign Activities by the Russian Government” (April 15, 2021).
 - “Justice Department Announces Court-Authorized Disruption of Botnet Controlled by the Russian Federation’s Main Intelligence Directorate (GRU)” (April 6, 2022).
 - Microsoft, “An Overview of Russia’s Cyberattack Activity in Ukraine” (April 27, 2022).
- *Chinese Intellectual-Property Theft*
 - U.S. Department of Justice, Indictment, U.S. v. Li Xiaoyu & Dong Jiazhi (July 7, 2020).

- The White House, “The United States, Joined by Allies and Partners, Attributes Malicious Cyber Activity and Irresponsible State Behavior to The People’s Republic of China” (July 19, 2021).

Wednesday, October 5, Week #6: Encryption Part I: End-to-End Encryption vs. “Going Dark”

- Berkman Center (Harvard), “Don’t Panic Making Progress on the “Going Dark” Debate” (February 2016).
- Stefan Savage, Lawful Device Access without Mass Surveillance Risk: A Technical Design Discussion, Oct. 2018
- Ronan Farrow, “How Democracies Spy on Their Citizens” *The New Yorker* (April 25, 2022).
- *Optional:*
 - Hal Abelson *et al.*, “Keys under doormats: mandating insecurity by requiring government access to all data and communications” (September 2015).

[Final-project proposals due October 5]

Wednesday, October 12, Week #7: Encryption Part II: Lawful Access and Continued Private Sector Tensions with Law Enforcement

- Alan Z. Rozenshtein, Surveillance Intermediaries, Jan. 2018
- Mayank Varia, A Roadmap for Exceptional Access Research, Dec. 5, 2018
- Riana Pfefferkorn and Jennifer King, “Here’s Why Apple’s New Child Safety Features are So Controversial” (podcast and transcript) (August 10, 2021).
- Jack Nicas, Raymond Zhong, and Daisuke Wakabayashi, Censorship, Surveillance and Profits: A Hard Bargain for Apple in China (June 17, 2021).
- *Optional (technically oriented):*
 - Apple, CSAM Detection: Technical Summary, Aug. 2021
 - Abhishek Bhowmick *et al.*, The Apple PSI System, July 29, 2021

[Fall Break]

Wednesday, October 26, Week #8: Artificial Intelligence Part I: Artificial General Intelligence: Separating Fact from Fiction

- Scott Reed *et al.*, “A Generalist Agent,” *DeepMind* (May 19, 2022).
- Kyle Wiggers, “DeepMind’s New AI Can Perform Over 600 Tasks, from Playing Games To Controlling Robots,” *TechCrunch* (May 13, 2022).
- Gary Marcus, “The New Science of Alt Intelligence: AI Has Lost Its Way. It’s Time To Take A Step Back,” *Substack* (May 14, 2022).
- Tom Simonite, “AI Has a Hallucination Problem That Is Proving Tough to Fix,” *Wired* (March 9, 2019).
- Stuart Russell, “Provably Beneficial Artificial Intelligence,” Yale Jackson School of Global Affairs (May 1, 2022)

Wednesday, November 2, Week #9: Artificial Intelligence Part II: Vulnerabilities, Adversarial Use, and Dual-Use Research

- Fabio Urbina et al., “Dual Use of Artificial-Intelligence-Powered Drug Discovery,” *Nature Machine Intelligence* 4 (189-191) (2022).
- Ilya Shumailov et al., “Sponge Examples: Energy-Latency Attacks on Neural Networks,” *IEEE European Symposium on Security and Privacy* (212-231) (2021).
- Ilya Shumailov et al., “Manipulating SGD with Data Ordering Attacks,” (April 19, 2021).
- Shafi Goldwasser et al., “Planting Undetectable Backdoors in Machine-Learning Models” (April 14, 2022).
- Brian Christian, “The Alignment Problem: Machine Learning and Human Values,” Yale Jackson School of Global Affairs (March 31, 2022).

Wednesday, November 9, Week #10: Artificial Intelligence Part III: Algorithmic Warfare and Lethal Autonomous Weapons

- U.S. Department of Defense, *Responsible Artificial Intelligence Strategy and Implementation Pathway* (June 2022).
- National Security Commission on Artificial Intelligence, *Final Report* (November 2019) [**only** chapters 1-4 are required reading].
- International Committee of the Red Cross, “Autonomous weapons: The ICRC recommends adopting new rules” (August 3, 2021).

Wednesday, November 16, Week #11: Project Presentations Part I

[Thanksgiving Break]

November 30, Week #12: Project Presentations Part II

Wednesday, December 7, Week #13: Project Presentations Part III

[Wednesday, December 14: Projects Due]