Legal Studies 190
Responsible AI, Law, Ethics & Society

Spring 2024
Dates
Graduate Student Instructor: YYY

Talia Schwartz
Taliashwartz@berkeley.edu
Collaborative Co-Teaching: Shlomi Hod
Additional co-lectures TBD
Office hours xxx (or by appt.)

Responsible AI, Law, Ethics & Society

https://learn.responsibly.ai/2024-spring-bu-berkeley/

Course Description and Objectives

The deployment of Artificial Intelligence systems in multiple domains of society raises fundamental challenges and concerns, such as accountability, liability, fairness, transparency and privacy. The dynamic nature of AI systems requires a new set of skills informed by ethics, law, and policy to be applied throughout the life cycle of such systems: design, development and deployment. It also involves ongoing collaboration among data scientists, computer scientists, lawyers and ethicists. Tackling these challenges calls for an interdisciplinary approach: deconstructing these issues by discipline and reconstructing with an integrated mindset, principles and practices between data science, ethics and law. This course aims to do so by bringing together students with diverse disciplinary backgrounds into teams that work together on joint tasks in an intensive series of in-class sessions. These sessions will include lectures, discussions, and group work.

By the end of this class, students will be able to: [1] communicate with professionals from other disciplines, identify gaps in the meaning of terms and perspectives, and develop a shared language; [2] demonstrate understanding of the impact of AI on individuals, groups, society and humanity; [3] proactively spot ethical issues and scan for unintended consequences and potential harms;[4] demonstrate introductory knowledge and skills to oversee and audit AI systems through their life cycle (design, development and deployment); [5] find and use resources to achieve all of the above; [6] apply their knowledge towards first steps in the industry, while shaping their responsibility as professionals.

This class is designed as an experiential learning course: most of class time you'll be spending in working on data science-legal real-world challenges within a small team of law and CS/DS students. Class format is *case-studies* for Law and *iterated and interactive research of data* (e.g., with Jupyter Notebook) for Data Science. These two pedagogies are being used in every class, accessible to all of the students, and integrated together.

Students from computer science or data science backgrounds should have some experience with machine learning.

Evaluation and Grading

Pre-Class Assignments: There are a few assignments to be done and submitted before some of the classes. The students will use the outcomes of these assignments in the class. The submissions are mandatory.
In-Class Assignments: In every class, all teams are required to submit a half-pager memo and a deck of a few slides at the end of each class. Each team will present twice during the course.
Final Project: In their final project, the teams will be asked to develop a new case-study that makes use of data sets and data science techniques to demonstrate a legal and ethical dilemma regarding Responsible AI, Law, Ethics and Society.

Team final project	40%
Team in-class assignments and presentations	20%
Individual contribution (pre-class assignments; presentation; peer-review final project)	40%

Texts

Texts for the course are available in electronic format on bCourses.

Policies

Please refer to Berkeley's Academic Integrity policy (http://sa.berkeley.edu/conduct/integrity). *We take academic integrity and honesty seriously.*

Students requiring accommodation for disability should also make sure that we get the official accommodation notice from DSP ***by the third week of the semester***.

Make sure to check bCourses regularly, since that will be our medium of communication.

Course Readings and Schedule

Date	Theme	Policy Area	Reading materials	Additional reading (optional)
Week 1: 1/22	**AI & Us**	**Social Welfare**	**packet**	Jessica Fjeld, Nele Achten, Hannah Hilligoss, Adam Nagy, and Madhulika Srikumar. 2020. Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. Berkman Klein Center Research Publication 2020-1 (2020)
Week 2: 1/29	**Liability & Robustness**	**Autonomous Vehicles**	**packet**	Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In Proceedings of the conference on fairness, accountability, and transparency. 59–68; Alexandra Chouldechova and Aaron Roth. 2020. A snapshot of the frontiers of fairness in machine learning. Commun. ACM 63, 5 (2020), 82–89; Boeglin, Jack. "The costs of self-driving cars: reconciling freedom and privacy with tort liability in autonomous vehicle regulation." Yale Journal of Law & Technology 17 (1) (spring 2015) 171; JAMES M. ANDERSON ET AL., AUTONOMOUS VEHICLE TECHNOLOGY: A GUIDE FOR POLICYMAKERS 58–95 (RAND Corp. 2014); Hevelke A, Nida-Rümelin J. Responsibility for crashes of autonomous vehicles: an ethical analysis. Sci Eng Ethics. 2015 Jun;21(3):619-30.

Week 3: **2/5**	**Discrimination & Fairness I**	**Labour Market**	**packet**	The Legal and Ethical Implications of Using AI in Hiring. HBR. in: https://hbr.org/2019/04/the-legal-and-ethical-implications-of-using-ai-in-hiring ; Dastin, Jeffrey. "Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women." Reuters, October 9, 2018. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G
Week 4: **2/12**	**Discrimination & Fairness II**	**Healthcare**	**packet**	Blasimme A, Vayena E. The ethics of AI in biomedical research, medicine and public health. In: Dubber MD, et al., editors. The Oxford Handbook of Ethics of AI Oxford. Oxford UP; 2020; Ian Kerr & Jason Millar & Noel Corriveau, "Robots and Artificial Intelligence in Health Care" in Joanna Erdman, Vanessa Gruben, Erin Nelson, eds, *Canadian Health Law and Policy*, 5th ed (Toronto: LexisNexis Canada, 2017) 257.
2/19	**No class - Academic and Administrative Holiday**			
Week 5: **2/26**	**Final project - prep session**			
Week 6: **3/4**	**Privacy**	**Transportation**	**packet**	REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL. Article 6 and Recitals (40) to (49) of the GDPR https://gdpr.algolia.com/gdpr-article-6 ; California Consumer Privacy Act of 2018 (CCPA); Trans-Atlantic Data Privacy Framework

				statement (25 March 2022). See here: https://ec.europa.eu/commission/presscorner/detail/en/ip_22_2087; Collingwood, L. (2017). Privacy implications and liability issues of autonomous vehicles. Information & Communications Technology Law, 26(1), 32–45; Glancy, D. J. (2012). Privacy in autonomous vehicles. Santa Clara Law Review, 52(4), 1171–1239.
Week 7: **3/11** **Berkeley Only**	**Guest lecture**			
Week 8: **3/18**	**Deploying AI applications with foundation models & generative AI**	**ChatGPT**	**packet**	TBD
3/25	**No class - Spring break**			
Week 9: **4/1**	**Transparency & Explainability**	**Finance**	**packet**	Zerilli, J., Knott, A., Maclaurin, J., & Gavaghan, C. (2019). Transparency in Algorithmic and Human Decision-Making: Is There a Double Standard?. Philosophy and Technology, 32 (4), 661-683; Katyal S (2019) Private accountability in the age of artificial intelligence. University of California, Los Angeles Law Review 66: 54–141
Week 10: **4/8**	**Integration: Content Moderation**	**Social Media Platforms**	**packet**	Grimmelmann J (2015) The virtues of moderation. Yale Journal of Law & Technology 17: 42;

				Gorwa, R. What is platform governance?. Information, Communication & Society 22(6): 854–871 (2019)
Week 11: **4/15** **Berkeley Only**	**Guest lecture**			
Week 12: **4/22**	**AI Governance**	**Audit, Governance frameworks**	**packet**	Bommasani, Rishi & Hudson, Drew & Adeli, Ehsan & Altman, Russ & Arora, Simran & Arx, Sydney & Bernstein, Michael & Bohg, Jeannette & Bosselut, Antoine & Brunskill, Emma & Brynjolfsson, Erik & Buch, Shyamal & Card, Dallas & Castellon, Rodrigo & Chatterji, Niladri & Chen, Annie & Creel, Kathleen & Davis, Jared & Demszky, Dora & Liang, Percy. (2021). On the Opportunities and Risks of Foundation Models.
Week 13: **4/29**	**Project Presentations and Course Summary (last meeting)**			
Week 14: **5/6**	**Project Submission Deadline**			