Course Number & Title (A.1)	**CIS399: The Science of Data Ethics, Spring 2019**
Credit Units (A.2)	1
Class/Laboratory Schedule	*Lecture: TueThu 1:30-3PM in 3401 Walnut Active Learning Classroom (3 hrs/week for 14 weeks)*
Instructor (A.3)	Prof Michael Kearns and Prof Ani Nenkova
Text(s)/Required Materials (A.4)	• *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, Cathy O'Neil • *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, Virginia Eubanks • *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World*, Bruce Schneier • *Future Crimes: Inside the Digital Underground and the Battle for Our Connected World*, Marc Goodman • *The Ethical Algorithm* (excerpts), by Michael Kearns and Aaron Roth • Readings from the scientific literature on algorithmic fairness, differential privacy, transparency, interpretability, and related topics
Catalog Description (A.5a)	Hardly a week passes by without a major news article detailing the proliferation of data-driven technology, and the ethical failings of these technologies. Among the topics that garnered most attention in recent years are racial disparities in automated sentencing and parole, gender bias in representations for language technology, and data breaches of stolen customer data, like the Equifax breach and the Cambridge Analytica/Facebook scandal. In this active learning class, we will introduce aspiring data science technologists to the spectrum of ethical concerns, focusing on social norms like fairness, transparency and privacy. We will then introduce technical approaches to a number of these problems, including by hands-on examination of the tradeoffs in fairness and accuracy in predictive technology, introduction to differential privacy and overview of evaluation conventions for predictive technology. Further, we will provide guidelines for examining system training data for bias, representation (of race, gender and other characteristics) and ecological validity. Equipped with this knowledge, students will learn how to conduct informed analysis of the usefulness of predictive systems. They will audit for ethical concerns papers from the contemporary top artificial intelligence venues and the ongoing senior design projects. There will be weekly reading assignments and associated group activities, four technical assignments and a final. At the beginning of class, students will select to read one of the three class books, and after the first month of class, each group will lead a class discussion on the main takeaways of the book they picked.
Prerequisites (A.5b)	*CIS 121 and permission of instructors*
Course Satisfies (A.5c)	[ ] Math [ ] Science [ ] Engineering [ ] Technical Elective [ ] TBS [ ]SSH
Course Web	*Canvas*
Course Outcomes (A.6a)	*Understand the ethical implications of data science technology* *Learn the practical trade-offs needed for fair predictive technology and data privacy* *Be aware of specific ethical concerns about artificial intelligence and data science* *Predictive science technology in medicine, business and law*
Contribution towards Program Outcomes (A.6b)	*G. The ability to communicate effectively* *H. The broad education necessary to understand the impact of engineering solutions in a global, economic, environmental and society context* *I. An ability to recognize the need for and an ability to engage in life-long learning* *J. An understanding of contemporary issues, especially those related to engineering practice*
Contribution towards Professional Component	*Example:* *25% Engineering science* *75% Engineering design*
Topics Covered (A.7)	*Foundations of predictive systems; Accuracy-fairness tradeoffs; System evaluation; Data collection and discovering bias in training data; Data privacy; Auditing predictive systems*
Weekly/Session Schedule (A.7) (Tentative)	Week 1 Jan 17 Class intro and logistics; topic overview, ACM and IEEE code of ethics; assignment 1 release Jan 22 Intro to predictive systems: data, features and approximations Week 2 *Jan 24* Logistic regression and linear classifiers *Jan 29* Issues in fairness and bias

	Week 3 Jan 31 Issues in fairness and bias Feb 4 Algorithmic solutions Week 4 Feb 7 Algorithmic solutions Feb 12 Analyzing bias in training data Week 5 Feb 14 Ecological validity of training data Feb 19 Gender and race variation in system performance Week 6 Feb 21 Individual rights: how inaccurate prediction affect the lives of potential users; case studies Feb 26 Book club discussions Week 7 Feb 28 Book club discussions Spring break Mar 12 Book club discussion and wrap-up Week 8 Mar 14 Privacy and what algorithms can learn from our online traces Mar 19 Mechanisms for collecting user data and privacy laws that can limit that Week 9 Mar 21 Testing of predictive systems; standards for adopting predictive technology Mar 26 Audits of senior design projects Week 10 Mar 28 Audits of senior design projects Apr 2 Discussion with guests from the Penn department of medical ethics and policy Week 11 Apr 4 Analysis of recent news stories Apr 9 Discussion with guests from Penn Law school Week 12 Apr 11 Analysis of recent news stories Apr 16 Discussion with guests from Wharton legal studies and business ethics Week 13 Apr 18 Race and technology: search and vision Apr 23 Gender and technology: speech recognition and language technologies Week 14 Apr 25 Legal context Apr 30 Class wrap up
Grading Details	40% class participation 40% homework 5% final 5% project audits
Program Fee	
Financial Aid	
Prepared By/Date	*Ani Nenkova and Michael Kearns, Sep 18, 2018*