

Law and Algorithms – Spring 2022 Syllabus

A joint class between the School of Law and the faculty of Computing and Data Sciences
CDS 657 & 457 / JD 673

This cross-cutting and interdisciplinary course, taught jointly between the School of Law, the faculty of Computing and Data Sciences, and Computer Science investigates the role that algorithms and automated decision-making systems play in law and society. The course connects legal and computational concepts of *transparency*, *fairness*, *bias*, *trust*, and *privacy*, through a series of case studies that present recent applications of technology to legal and regulatory situations and explore the challenges in regulating algorithms.

Legal concepts explored will include evidence and expert witnesses, anti-discrimination law concepts of disparate impact and disparate treatment, sectoral information privacy regimes, and public access and transparency laws. Computational concepts explored will include artificial intelligence and machine learning, secure multi-party computation, differential privacy, and zero-knowledge proofs.

Grades will be based on a series of assignments that correspond with each case study, to be completed collaboratively in mixed teams of law and computing/data science students.

1. Instructor Information

[Gabriel Kaptchuk](#)

Research Assistant Professor of CS
kaptchuk@bu.edu

[Ran Canetti](#)

Professor of Computer Science
canetti@bu.edu

[Andy Sellars](#)

Clinical Associate Professor of Law
sellars@bu.edu

Consistent with our ongoing COVID protocols designed to minimize situations where disease transmission can be high, the instructors are not on campus as frequently as they would be in pre-pandemic years. That said, we still want to be as available to students as possible. Please use email or Microsoft Teams to schedule a time to meet. Note that meetings will likely be a mix of in-person or over Zoom.

2. Course Websites

For course readings: <https://cs-people.bu.edu/kaptchuk/teaching/ds457/sp22-classpage.html>

For class discussions and announcements: see our Microsoft Teams page.

We have opted to use Microsoft Teams as our main base of operation as it is available to all students at the university. CS & CDS students don't have access to Law's default platform (Blackboard) and Law students don't have access to CS & CDS' default platform (Piazza). You will be invited to join the Microsoft Teams channel through your @bu.edu email accounts. More information about BU's use of Teams can be found [here](#), and we recommend that you [download the desktop version](#) of the app to make your access easier. (You may, but are not expected, to download a smartphone version of the app at the same link.)

Please refer to these resources for the most recent assignments and reading material. We will try to avoid alterations to class material with less than a week's notice. If there ever are last-minute changes, we will let you know.

3. Course Information

Meetings: Thursdays, 4:20—6:20pm, Jan. 20 to April 21 (except March 10)

Location: BU Law Tower, Room 204

Credit Hours: For CDS 657: 4 Credits
For JD 673: 3 Credits

To account for the difference in credits, CDS/CS students will have one additional assignment, wherein they will develop a short guide that explains the computational aspects of a socio-technological system addressed in the course, written for non-CS audiences. More information on that assignment will be distributed separately.

For the law students, per ABA guidelines you should anticipate a workload of roughly 42.5 hours per credit for the semester, which includes both in-class and out-of-class time. For elaboration, you may consult BU Law's Credit Hour Policy.

4. Course Objectives

The goal of this class is to help both law and computer/data sciences students to understand the importance of the other's field to their home discipline, and how law and algorithms work in concert to regulate human behavior. We specifically expect that students will:

- Learn and appreciate the complicated relationship between law and algorithmic systems, and how the two act as interrelated regulators with different systems of adjudication and affordances for human input.
- Understand the fundamental systems of law as they relate to algorithmic regulation including basics of the common law system, as well as the rules and policies that inform the legal domains addressed in our case studies, including evidence, administrative law, legislation, criminal procedure, intellectual property, anti-discrimination law, election law, and information privacy.
- Understand the fundamental systems of computing and data sciences as they relate to law and policy questions, including computational thinking, probabilities, optimization, cryptography, artificial intelligence and machine learning, zero-knowledge proofs, secure multi-party computation, differential privacy, and risk-limiting audits.
- Examine how both law and computer/data science reinforce and counter broader powers within social systems, including how both can perpetuate or mitigate bias and discrimination in criminal, civil, and administrative systems.
- Learn how to communicate concepts from their home discipline to those working in either law or computer/data science, and how to collaborate across disciplines to achieve mutual goals and policy outcomes.

5. Prerequisites

There are no course prerequisites and no formal prior knowledge of law or computer science required to participate in this course.

For law students, we ask that you come to the class with an open mind for computer science and mathematical thinking and vocabulary, and a willingness to explore the way in which algorithms practically operate in computational systems.

For CS/CDS students: we think you will get the most out of this course if you have a good grasp of computer systems, algorithms and their analysis, AI basics, and computer security. We also ask that you come to the class with an open mind for understanding legal thinking and language, as well as the social aspects of information systems.

While this is primarily a graduate level course, advanced undergraduate students in CS and CDS may enroll after receiving permission from the instructors.

6. Course Materials

There are no required textbooks for this course, and all class material will be free.

Readings will be made available through the course websites. The specific readings will be released over the course of the semester, so please refer to the website for the latest information. And because the material will change, do not read more than a week ahead without checking with us first.

For the caselaw readings we'll have this semester we are using the [OpenCasebook platform](#), which allows you to both read our excerpt of the case and click through to see what we're omitting. Unless otherwise indicated, students will be expected only to read/view the article, case excerpt, or blog post indicated, and not any other content on the site.

There may also be optional readings associated with each class day. Optional readings are, indeed, optional. We've selected them because we think they may be interesting or engaging, but you are not required to read them.

The quality of an interdisciplinary class like this really rises and falls on whether the students have done the reading, and we really appreciate the CDS students changing up their usual method of class preparation by doing a good degree of reading before each class. (We appreciate that from the law students too, of course, but it is more generally expected in legal education.) We ask that all students come to class having carefully read what is assigned and prepared to discuss the readings in class.

7. Classroom Attendance and Expectations

The heart of this class is to provide a forum in which the disciplines of law and computer/data science can learn from each other. That learning is best fostered by active and engaged student participation. To that end, *we ask that students attend each of the thirteen class sessions and actively participate in every class.* (More on class participation below.)

That said, we are aware that unavoidable conflicts do come up, especially in these times. If one does arise, please contact one of the instructors in advance of the class so we can discuss it. We do not expect any student to miss more than two days of class, barring highly unusual circumstances.

Boston University is now back to fully on-campus instruction, but public health concerns remain top of mind, especially with the Omicron variant continuing to spread. Please follow all university protocols for COVID testing and symptom attestation. Specifically:

- Do not attend class if you show even light COVID symptoms (fever, cough or other respiratory issues, nausea, etc.) or if you have any reason to believe that you are contagious with COVID or have been exposed. Please err on the side of absence here; we want to do all we can to protect each other.
- Follow all COVID protocols provided by Boston University and the Commonwealth of Massachusetts. Wear a mask at all times, and practice social distancing as much as is possible.

As you surely know by now, the public health situation can change abruptly with COVID. Please be ready for changes to the classroom format and presentation, including a move to virtual classes if required by the University.

8. Assignments and Grading

There is no exam for this course. Your performance in the above objectives will be evaluated through active participation in weekly classes, as well as in assignments that engage with our four case studies. Your grade is specifically based on the following:

8.1. Participation (25% of Grade)

You will be expected to have read the assigned readings each week and participate actively in class discussion with substantive contributions. You satisfy this requirement by making *at least one substantive contribution every week*, in one of two ways:

- during the class session, or
- on the Microsoft Teams page in the section for the day's class.

If you choose to participate by making a substantive contribution before class on Microsoft teams, please be sure to post your comment far enough in advance that other students will have time to react to what you say. This is a good way to contribute if you are forced to miss a class or if you prefer written contributions to oral discussion.

8.2. Written Projects (75% of Grade, Split Evenly Across Five Assignments)

Over the course of the semester, students will complete five short projects in mixed Law/CDS teams of three to four students (with each team including at least one Law and one CDS student). Because one of the goals of this course is to develop your skills at collaborating and communicating across disciplines, you should seek out as many different teammates as possible. You will not be allowed to complete more than one project with the exact same team.

As you will see, these projects will focus on one of the five primary topics that we will explore in this course through our case studies — transparency, fairness, bias, trust, and privacy. The team will be asked to prepare a paper that addresses a prompt related to the current module and suggest a legal, technological, and/or mixed legal/technological response to the problem presented. We will expect the project to engage with the relevant written material for the case study, conduct external research as is appropriate for the assignment, and present a response in a way that thoughtfully engages with existing literature and solutions, including any possible consequences or shortfalls in their response.

Further details on each project will follow. Subject to modification based on the pace of the course, the deadlines for each project will be:

- A. Transparency – due ***before class on February 10***
- B. Fairness – due ***before class on February 24***
- C. Bias – due ***before class on March 17***
- D. Trust – due ***before class on March 31***
- E. Privacy – due ***before class on April 21***

Students interested in further developing their projects into a more substantial (and potentially publishable) work are welcome to discuss their goals with one of the instructors. We have had successful public papers more out of prior versions of this class. For law students, though, please note that we do not expect any of these written projects to be enough to satisfy the Law School's upper-level writing requirement, though we can discuss how you can meet that requirement through an alternative assignment.

9. Accommodations

Boston University is committed to equal access for all students. If you require any ability accommodations in this class, please let the Law Registrar (lawreg@bu.edu) know early in the semester so that appropriate accommodations can be made. You must provide the Law Registrar with a letter of needed accommodations prepared by Disability & Access Services. Contact information for that office is as follows: (617) 353-3658 V/TTY or access@bu.edu. All discussions and written materials will be kept confidential.

An overview of the class-by-class topics follows.

Course Topics

Please note that this is a general overview of the topics we'll have in class this year. The substance is likely to change, so *please refer to the course websites for all topics and readings.*

Introduction

- *Class 1 – Intro to Law, Intro to Algorithms (Jan. 20):* We begin the class with a primer on law and legal thinking, and a primer on computer science and computational theory. For your new domain, this will serve as an exposure to the key concepts and methods within the domain. For your home domain, this will be a chance to think anew about the fundamentals of your discipline.

Transparency – Probabilistic Genotyping Algorithms

- *Class 2 – The Development and Legal Protection of Software (Jan. 27):* to tee up our discussion of a particular form of criminal forensic algorithm, we start with how software is built, how flaws can come into software construction, and how our default protections in intellectual property law (and especially trade secrets law) can put our desire to create an economic incentive for software development at loggerheads with our desire to understand how software is operating.
- *Class 3 – Putting the TrueAllele Algorithm on Trial (Feb. 3):* with our knowledge on how software is built and protected, we turn a particular probabilistic forensic algorithm called “TrueAllele,” which used in criminal cases to identify the presence of a person’s DNA in situations where traditional (and more generally accepted) methods DNA identification are not possible. We look at how the rules of evidence and criminal procedure, including the role of scientific expert witnesses, have been applied to interrogate the reliability of TrueAllele.

Fairness – Criminal Risk Assessment Algorithms

- *Class 4 – The COMPAS Algorithm and the Optimization Paradox (Feb. 10):* One of the most famous collisions between law and algorithm concerned an algorithm called “COMPAS,” which has been used in multiple states to assist courts in evaluating the likelihood that a criminal defendant will reoffend. A major exposé in *Pro Publica* revealed substantial issues in the algorithm, and how differently it treated applicants of different racial and ethnic identities. We’ll review the COMPAS story and how the Supreme Court of Wisconsin approached challenges to the use of the algorithm. We’ll also review a curious statistical issue that emerged from the COMPAS saga around the inherent tradeoffs that must be made between an algorithm’s predictive value and its error rates.
- *Class 5 – Is There a “Right” Way to Use Algorithms in Criminal Sentencing? (Feb. 17):* The errors of COMPAS were clear, but it is equally clear that there will be continued calls to use such algorithms in criminal sentencing to safeguard against human errors and unfairness. So, should we? What should that look like? On what data should those algorithms base their risk assessments? What safeguards, limits, and standards should apply? Or, should we abandon them entirely, and would that leave us in a more fair space?

Bias – AI-Enabled Housing and Credit Tools

- *Class 6 – Artificial Intelligence and Anti-Discrimination Laws (Feb. 24):* We'll begin exploring the interplay of law with artificial intelligence and machine learning, through the lens of federal anti-discrimination laws. We'll contrast a 2019 attempt by the Department of Housing and Urban Development to shield those who use automated decision-making systems in mortgage approval and housing credit situations with the Federal Trade Commission's 2021 decision to crack down on use of biased algorithms more broadly, regardless of intent.
- *Class 7 – Can Algorithms Mitigate Bias? (March 3):* Humans are biased, and so some have called for greater use of computational systems and objective data to remove the bias of humans in areas like housing and employment. But, of course, humans create those algorithms, and few if any forms of data are free of deep social impacts and meanings. So are we at an impasse, or could one solve bias problems with algorithms?

No class on March 10 – Spring Break

Trust – Election Security

- *Class 8 – Vote by Paper, Vote by Mail, Vote by Smartphone (March 17):* In a democracy, so much of our law and policy comes down to votes—from electing officials, to approving ballot initiatives, to issuing jury verdicts. Today we examine the security, privacy, and accountability properties we expect out of a voting system, and how they map onto our systems of voting, both current and (maybe) future.
- *Class 9 – How Do We Trust the Vote? (March 24):* Maybe you heard this one: there was a national presidential election in 2020, and by the beginning of 2021 the lies that were spread about how the election was conducted resulted in an unprecedented attack on the United States Capitol, a second impeachment on the then-President, and grave concerns about what may happen in the 2022 and 2024 national elections. We look at how transparency laws and computational concepts—including “risk limiting audits”—combine to give all of us confidence on the outcomes of our votes. We also look at how more novel forms of secure computation—including “zero knowledge proofs”—could be employed in future elections to boost (or hinder) trust.

Privacy – Privacy-Preserving Computational Techniques

- *Class 10 – “Privacy,” “Security,” and “Encryption” (March 31):* To begin our classes on privacy, we return to how law and computer/data science treat different similar concepts differently. We look in particular about the differences between the privacy of information, the security of information and the computation tool of encryption. We look in particular how these different concepts can be weaponized against each other, including how encryption can be used as a tool of perfect surveillance and how laws on privacy can entrench some forms of surveillance in the guise of stopping others.
- *Class 11 – Differential Privacy and The Census (April 7):* For the 24th time in our country's history, we have just completed a national census, generating many millions of pieces of information about the United States population. Out of a desire to get as truthful a set of information as possible, we place very severe restrictions on the Census Bureau releasing that information. And

yet, this information is incredibly valuable and used for a variety of research and policy applications. So, our Census Bureau and similar statistical organizations around the world have turned to a particular computational technique—“differential privacy”—to share information from that data without fully revealing its contents. We dig into this technique, how the Census went about adopting it, and the legal challenge brought by the State of Alabama for its use in the 2020 census.

- *Class 12 – Conducting Analysis Over Secret Data (April 14)*: Many data privacy statutes—including the sectoral data privacy regimes in the United States, like HIPAA, FERPA, and the Gramm Leach Bliley Act, place heavy restrictions on a data custodian’s ability to disclose data to others, and yet the data held by these custodians can be tremendously useful for a variety of social questions. Today we’ll explore how you can use one advanced computational technique—“secure multi-party computation”—to have both privacy and insights into data, and whether this should cause us to rethink how we implement data privacy laws.

Synthesis

- *Class 13 – Law and Algorithms (April 22)*: We’ll close our class by discussing how we can intelligently address issues at the intersection of law and algorithms, and what lessons we can take from these case studies to other policy debates and problems.