Law and Algorithms – Spring 2024 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 graduate course, taught jointly between the School of Law and the Faculty of Computing & Data Sciences, investigates the role that algorithms and automated decision-making systems play in law and society. The course connects legal and computational concepts of transparency, trustworthiness, privacy, secrecy, bias, discrimination, and fairness through a series of case studies that present recent applications of technology to legal and regulatory situations and explore the challenges in regulating algorithms and using algorithms in legal systems.

Legal concepts explored will include evidence and expert witnesses, anti-discrimination law concepts of disparate impact and disparate treatment, regulation of civic data gathering activities like the census, sectoral information privacy regimes, and public access and transparency laws. Computational concepts explored will include basics of modeling and automation, software design and development, artificial intelligence and machine learning, and privacy-enhancing technologies.

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

Andy Sellars
Clinical Associate Professor of Law
sellars@bu.edu
BU Law Tower, Room 1204D

Mayank Varia
Associate Professor of Computing & Data Sciences
varia@bu.edu
CDS Building, Room 1343

Please use email or Microsoft Teams to schedule a time to meet with us, either virtually or in person.

2. Course Websites

For course readings: https://github.com/afsellars/Law-and-Algorithms/
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, as we know many CS & CDS students don't use Blackboard (Law's default platform), and most Law students don't use Piazza (CS & CDS' default platform). 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.

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, 2:10—4:10pm, January 18 to April 18 (except March 14)
Location: BU Law Tower, Room 203
Credit Hours: For CDS 457 & 657: 4 Credits
For JD 673: 3 Credits

To account for the difference in credits, CDS/CS students will have an additional assignment separate from this syllabus. More information will be distributed separately to those students.

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 the rules and policies that inform the legal domains addressed in our case studies. These include evidence, administrative law, legislation, criminal procedure, intellectual property, anti-discrimination law, and information privacy.
- Understand the fundamental systems of computing and data sciences as they relate to law and policy questions, including computational thinking and modeling, probabilities, optimization, artificial intelligence and machine learning, cryptography, differential privacy, and secure multi-party computation.
- Examine how both law and computer/data science reinforce and counter broader powers within social systems, including how both can perpetuate or mitigate bias, discrimination, and other exclusionary harms in criminal, civil, and administrative systems.
- Consider the limitations of both law and computing systems, including how to spot errors, pitfalls, overlooked values, and other shortcomings in both software and legal systems, and how one might address those shortcomings.
- 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. Fostering an Inclusive Learning Environment

We intend this class to be a fully inclusive learning environment. We want you to feel safe and comfortable in bringing your best self to class, expressing viewpoints, and participating in seminar discussions. If there is anything we are doing that is preventing you from bringing your best self, please let the instructors know at any time, and together we will address it.

We expect all participants to show respect to one another. If you ever have a concern about our behavior or language, or that of another student, please feel free to approach us in person, by email, over Teams, or with an anonymous note.

Class rosters are provided to us with each student's legal name. We will gladly honor any request to address you by an alternative name and/or pronoun. We will ask for this in our first class, but if we miss this information for any reason please advise us of your preferences as soon as possible.

6. Prerequisites

There are no course prerequisites and no prior cross-disciplinary experience is required to participate in this course. In short, we don't expect or require the law students to have any background in computing or data sciences, nor do we expect the CS, CDS, or other students to have any prior background in the law.

What do we expect, though, is that each student comes into the class willing to work hard to understand the other side's discipline and how it relates to your home discipline. We discuss more about good cross-disciplinary collaboration in Section 10 below.

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 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. Finally, we ask that you come to the class willing to complete assignments that may be different from other CS/CDS courses, including completing significant amounts of reading and persuasive essay writing.

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

7. 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 website. 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 the instructors first.

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 the reading before each class and reflecting on the content therein. (We appreciate that from the law students too, of course, but this is the default mode 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.

8. 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.)

We are aware that unavoidable conflicts do come up. If one does arise, please contact one of the instructors in advance of the class so we can discuss it. We expect that no student will miss more than two days of class, barring highly unusual circumstances.

We follow BU's policy on absences for religious observance.

We can record missed classes by request or as needed under a classroom accommodation (see below). In the class participation section below, we discuss how you can still participate asynchronously on days that you miss class.

9. 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:

9.1. Participation (20% 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, ideally at least a day before the relevant class. Adding a comment on Teams is a great way to contribute if you will miss a class or if you prefer written contributions to oral discussion.

9.2. Written Projects (80% of Grade, Split Evenly Across Four Assignments)

Over the course of the semester, students will complete four short projects in mixed Law/CDS teams of three to five students. The instructors will randomly assign students to teams for each project.

As you will see, these projects will focus on one of each of the four major topics for our course — modeling and automation; embodying algorithms in software; layering in secrecy; and large data models. The team will be asked to address a prompt related to the current module and provide both legal and computational analysis on the problem presented. We will expect the project to engage with the relevant written material for the module, 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. Modeling and automation – due ***before class on February 8***
B. Embodying algorithms in software – due ***before class on February 29***
C. Layering in secrecy – due ***before class on March 28***
D. Large data models – due ***before class on April 18***

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 arise from prior versions of this class.

For law students, 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.

We will provide submission instructions with each assignment.

10. Collaboration Across Disciplines and Assignment Expectations

We have designed the assignments in this class to center around joint group projects and aim to have a similar proportion of law and CS/CDS students in each group. The projects are intentionally open-ended and require lively collaboration and brainstorming within groups to do well.

That said, we also recognize that the open-ended and interdisciplinary character of the projects can lead to some uncertainty and concerns about the roles of students in group projects, and concerns about unequal contribution to projects. We also recognize that the norms in law and in CS/CDS can be a bit different when it comes to collaboration. And despite interdisciplinary collaboration being a standard part of professional life, we know that your prior academic experience provides little frame of reference on what's expected of you.

To that end, we have found it helpful in this course to set some explicit expectations around how we expect groups to collaborate. We expect *all students* to:

- Have read the material from the class days relevant to the assignment in full, as well as the assignment prompt.
- Prepare thoughts as to how you would respond to *all* the prompts in the assignment, and to do so ahead of any discussion meeting that your team may schedule. You should enter your first team meeting with some developed opinions on all aspects of the assignment.
- Engage thoughtfully with your peers across disciplines and collaborate as a team to address the prompts of the assignment.
- Contribute in roughly equal portion to the drafting of the final assignment.
    - *Note:* there can be a temptation in drafting to break these assignments into sections, divide the sections among the students, and have each student write an equal amount in each section. We advise against this, and found this approach to produce lower quality work (and grades) in prior iterations of this course. A thoughtful response to an assignment may require considerably more attention on some of these prompts over others depending on their relative importance. It will often be better to have collaborative writing across all sections.

We expect the CS and CDS students to pay special attention to the computational properties of assignments, to take the lead on any supplemental research needed in the areas of computing and data sciences, and to provide their expert thoughts regarding the computational issues present in the assignment.

We expect the law students to pay special attention to legal and regulatory aspects of assignments, to take the lead on any supplemental research needed in the law, and to provide their expert thoughts regarding the legal and regulatory issues present in the assignment.

We do not want any student:

- Drafting the entire assignment on behalf of the group or serving as the “final editor” of the work.
- Confining their opinions only to some of the portions of the assignment.
- Missing internal team deadlines or delaying the team’s timely completion of assignments.

And finally, a few level-setting notes on the assignments themselves:

1. We expect the documents to have generous citations, but we do not need sentence-by-sentence citation support like you might see in formal legal writing. Instead, reference the literature as you debate or engage with it, or as you are relying on it to make a substantive factual point.
2. Similarly, we have intentionally not asked for a particular citation format. The most important thing for us is that we can see and understand what literature you are using to make your points, and (when the work is paginated) where in the work we can find the specific substantive support.
3. These are consensus documents. No submission will perfectly capture the views of any one individual on the team, nor should they. We expect some disagreement along the way and some effort made to debate and reach a consensus view.
4. If there is a disagreement on a key question and you are unable to reach a consensus after debate, you may indicate as such in the assignment by presenting the conflicting views and their relative strengths/weaknesses.

5. While we hope there won't be, if there is an interpersonal issue on the team and you are unable to resolve it, please let us know as soon as possible. There is more we can do to remedy that situation if we catch it well in advance of an assignment deadline.
6. If you'd like a little extra coaching with your writing don't forget that BU has its Educational Resource Center that provides writing and other academic skills assistance.

11. Use of Large Language Models

It would be hard for us to teach a class like this without spending some time discussing so-called "Generative AI" and large language models. We hope that by the end of the course you will have an appreciation of the affordances and limitations of such models, and are equipped to critically examine their deployment. As you may know, the BU CDS faculty were among the first in the country to adopt a "Generative AI Assistance Policy," and BU Law has been carefully studying their adoption and use for well over a year.

For this class, we see pedagogical value in considering their use for aspects of the work, but also see a number of technical and policy concerns that strongly counsel against their use. In keeping with the overall theme of the course, our preference is that any use of any large language models or other tools that synthesize text in response to short prompts be done in an intentional, transparent, and reflective way. To that end:

- Generative AI tools (e.g., ChatGPT, Bard, LLaMa, etc.) may only be used with the prior approval of the course faculty based on a clear statement of the scope of the intended use. We expect to see proposed uses that supplement, not replace, other techniques, and a proposal that engages with the risks as well as the benefits of the use.

- Any approved use of Generative AI must be documented throughout, including the purpose of the use; the tools, queries, or other techniques used; and a reflection on the results generated.

- The student or team assumes full responsibility for reliance on any of the outputs of a Generative AI tool.

We expect that you will find that, once you fully consider their use, the benefits of relying on a Generative AI tool are most often substantially outweighed by other forms of writing. Indeed, the art of crafting an argument and carefully selecting how to communicate technical and legal concepts across disciplines is one of the key skills we hope you get out of this class, and would not want to see you deprived of that learning opportunity. But that said, we are open to hearing any proposals for use of large language models, provided it is sought in advance of the use.

12. Accommodations

Boston University is committed to equal access for students with disabilities. If you have a specific disability and require accommodations in this class, please submit the BU Disability & Access Services (DAS) online Intake Form and contact BU DAS to make an appointment by calling (617) 353-3658 before the start or early in the semester, so that appropriate accommodations can be made in a timely manner. BU DAS will provide the Law Registrar with a letter of approved accommodations. Faculty are only informed of accommodations that may affect the operation of the classroom. All discussions with,

and written materials provided to, BU DAS will be kept confidential. You can email DAS as access@bu.edu.

13. 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 website for all topics and readings.*

Module 1: Modeling and Automation

We begin the year with the building blocks of all algorithmic systems. Whether we're talking about the algorithms that occupy all of the space in our present debates (like large language models) or the tools that are "algorithms" but aren't often thought of as such (like the U.S. Sentencing Commission's guidelines), these systems start with a similar premise. In order to be useful they must capture the world in some sort of mathematically-expressed system (modeling) that is then capable of repeating with limited human input (automation). And it is right here at this first step that we begin to see how policy problems can creep in, including issues with how models are constructed, what data is used, how those systems are embedded in law, and who is centered in the design of those systems.

- *Class 1: Intro to Modeling and Automation (Jan. 18)*. We begin with an introduction to the social construction of data and algorithms, including how seemingly-objective pieces of information can impart specific policies and world views, how the field of data science and its regulation is shaped by who is allowed to participate in those projects, and why regulation of informational systems can be so complicated.

- *Class 2: Automation Bias vs. Non-Automation Bias (Jan. 25)*. Technology law and policy scholars have developed two, seemingly-contradictory views around how our decision-making processes can be skewed by the introduction of automated systems. One line of argument suggests that humans overly defer to these systems, and believe them to be more accurate, objective, or fair than indeed they usually are. A second line of argument suggests that humans reject these systems, and overly distrust information when they learn that it comes from an automated system. We'll look at both phenomena and attempt to reconcile the two.

- *Class 3: Fairness in Automated Systems (Feb. 1)*. A commonly-stated goal of algorithmic systems is that they should be "fair." But what does that mean? How will we know when we have it? Can it only be detected by its absence? We'll explore these questions through the lens of the 2016 exposé in *Pro Publica* about the COMPAS criminal risk assessment algorithm—the news story that in many ways launched contemporary algorithmic accountability discussion.

The assignment for Module 1 is due before class on February 8.

Module 2: Embodying Algorithms in Software

As the COMPAS story reveals, there are issues that may arise at the design stage of a system, and related-but-distinct issues that can arise when an algorithmic system becomes encoded in software. Our next module considers what we need to know about both the algorithm and its software embodiment in order to assess the system's trustworthiness. Our focus for this module is the use of software as evidence in criminal trials, specifically the use of *probabilistic genotyping software*—that is, software

that estimates the statistical likelihood that a DNA sample (usually one that is too degraded to go through conventional DNA tests) is a match for a particular person. Legal concepts in this module will include evidence in criminal proceedings, due process, and the role of expert witnesses. Computational concepts will include software and software stack development, software verification and validation, and sources of error in programming.

- *Class 4: The Development and Legal Protection of Software (Feb. 8).* On the computational side, we will discuss how software is built and how flaws can come into software construction and various stages of the process. On the legal side, we'll explore how our default protections in intellectual property law (and especially trade secrets law) can limit our ability to access information on how software is operating.

- *Class 5: Putting Software on Trial (Feb. 15).* With our knowledge on how software is built and protected, we turn to a particular probabilistic forensic algorithm called TrueAllele. 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, and how those procedures confront the trade secrets laws we discussed in the prior week.

- *Class 6: Creating an Ecosystem of Trustworthy Software (Feb. 22)*. What could the law do in order to better ensure trustworthy software? If we are dissatisfied with our present method of dueling experts and the limitations of the system of evidence, what would we offer instead? Class today will largely focus on a debate between four different proposals for constructing more trustworthy software for probabilistic genotyping software.

The Module 2 assignment will be due before class on February 29.

Module 3: Layering in Secrecy

Our focus for much of the last module was what sort of transparency is best suited for ensuring trustworthiness of computational systems, but what if there are policy reasons to keep the inputs private? Full transparency may be fitting for experts analyzing the performance of a forensic technology like TrueAllele, but there are many systems where full transparency is unacceptable. This module looks at secrecy and privacy as related-but-distinct concepts, including how law and computer science have evolved differently in defining and measuring privacy, the difficulty of administering and examining privacy-protective systems, and how privacy can be used as a pretext for abuse and other harms.

- *Class 7: Anonymization, Identification, and Formal Notions of Privacy (Feb. 29).* We celebrate Leap Day with a leap into privacy law. Much of privacy law in the United States assumes that some information is "personally identifiable" and some is instead "anonymized," and draws different rules for those two categories. But a wealth of computational literature over the past two decades has shown that distinction is not clean. This has led to the adoption of a rigorous mathematical concept of privacy, "differential privacy," which seeks to carefully measure exactly how well a particular form of data processing prevents others from learning information about specific inputs. We will explore how differential privacy works, its privacy guarantees, and how it harmonizes with information privacy laws.

- *Class 8: Bringing Formal Privacy into Public Administration (March 7).* The advent of privacy-preserving computational techniques and formal notions of privacy allow for profound

new ways in which insights can be learned from data without revealing specific inputs. But these systems are complicated, even for experienced data scientists, and public adoption of them remains very limited. We look at why that is, focused on the contentious adoption of differential privacy by the U.S. Census Bureau for the 2020 decennial census.

No class March 14 – Spring Break

- *Class 9: The Weaponization of Privacy (March 21)*. At first glance, privacy-preserving computational techniques sound as though they give us all of the positives (e.g., the ability to process data and do useful things with it) and none of the negatives (e.g., the privacy harms that can flow from disclosure of information). A closer inspection reveals a more complicated story. Indeed, these techniques can retrench or augment the power of already-powerful institutions at the expense of others, and privacy laws and tools can be used to evade regulatory oversight. In this class we consider the more complicated social tradeoffs with privacy-enhancing technologies.

The Module 3 assignment will be due before class on March 28.

Module 4: Large Data Models

Our final module takes on the computational technology that's attracting the most discussion at our present moment: machine learning models, including large language models. Unlike most of the algorithms discussed above, which are engineered to take in a specific set of inputs and produce a specific set of outputs, machine learning models take a set of parameters and then "train" them based on large amounts of preexisting data to then produce their outputs. But the *scale* at which they do this makes them especially inscrutable, and the hype that surrounds them presents a rather obvious and substantial danger of automation bias. Not to mention, the scale of data access and computing needed for the most advanced of these models presents even further risk of entrenchment of existing economic, political, and social powers. Our next three classes dive into these models, looking particularly at how classification machine learning models can collide with anti-discrimination laws in housing, employment, credit-worthiness, and other determinations.

- *Class 10: Data Access as a Measurement of Power (March 28).* We begin the module with a review of how machine learning models work, and how they differ from some of the algorithms that we have examined previously. As you will see, many of the advantages of these models come from access to large amounts of data—often measured in the millions and billions, or more. We'll examine how data access can entrench existing power, and the growing legal challenges to the use of information from or about oneself in these models.

- *Class 11: Harms in Machine Learning (April 4).* We'll explore the interplay of machine learning and federal anti-discrimination laws. We'll unpack the various ways in which bias can enter such systems, and how the general regulatory structure of anti-discrimination law—and its concepts like "disparate treatment," "disparate impact"—can hold these systems to account.

- *Class 12: Correcting Harms in Machine Learning (April 11).* Even if you know how to detect bias in an algorithmic system, the act of correcting that bias ends up being harder than one might think. Anti-discrimination law's tendency to see any differentiation on the basis of a protected class as potentially actionable makes it very difficult to know when corrective measures to

address disparate treatment may be viewed as unlawful disparate impact. And efforts to computationally design systems to avoid disparate impact may result in unlawful quotas or miss the various ways discrimination extends beyond the so-called “four fifths rule” from disparate impact caselaw. We’ll explore these tensions with an in-class simulation of trying to de-bias an employment algorithm.

The Module 4 assignment will be due before class on April 18.

Synthesis

Finally, 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.

- *Class 13: Law and Algorithms (April 18).* We’ll examine some of the big-picture proposals for algorithmic regulatory reform, and discuss which proposals we think best capture all of the competing values and concerns that we’ve identified over the semester.