Computer Science and Law: Opportunities and Research Directions

Joan Feigenbaum
http://www.cs.yale.edu/homes/jf/

Columbia Univ., Nov. 4, 2020

What is “Computer Science and Law”?

- Formulating and solving problems that are **simultaneously computational problems and legal problems**

- Examples:
  - ❑ Security, privacy, encryption, and surveillance
  - ❑ Data sharing in a decentralized online world
  - ❑ Freedom of expression online (or the lack thereof)
  - ❑ Cyber espionage, cyber war, and cyber diplomacy
  - ❑ Cyber crime, cyber law enforcement, and digital forensics
  - ❑ Online market structure, platform monopolies, and antitrust law
  - ❑ Online government services
  - ❑ Digital intellectual property
  - ❑ Automation of legal reasoning and legal services
  - ❑ Fairness, accountability, transparency, *etc.* (FAT*) in machine learning and data mining

What is "Computer Science and Law"?

- Formulating and solving problems that are **simultaneously computational problems and legal problems**

- Examples:
  - ❑ **Security, privacy, encryption, and surveillance**
  - ❑ **Data sharing in a decentralized online world**
  - ❑ **Freedom of expression online (or the lack thereof)**
  - ❑ Cyber espionage, cyber war, and cyber diplomacy
  - ❑ Cyber crime, cyber law enforcement, and digital forensics
  - ❑ Online market structure, platform monopolies, and antitrust law
  - ❑ Online government services
  - ❑ Digital intellectual property
  - ❑ Automation of legal reasoning and legal services
  - ❑ Fairness, accountability, transparency, *etc.* (FAT*) in machine learning and data mining

Talk Outline

- Three examples
  - ❑ Security, privacy, encryption, and surveillance
  - ❑ Data sharing in a decentralized online world
  - ❑ Freedom of expression online (or the lack thereof)

- What do we mean by “And” and why it matters

- Pointers to two other CS-and-Law events

Talk Outline

- Three examples
  - ❑ **Security, privacy, encryption, and surveillance**
  - ❑ Data sharing in a decentralized online world
  - ❑ Freedom of expression online (or the lack thereof)

- What do we mean by "And" and why it matters

- Pointers to two other CS-and-Law events

Ubiquitous Encryption vs. Lawful Surveillance

- 1990's "Crypto War"
  - ❑ US Gov't: Need "key escrow" to deregulate Cold-War era, strong crypto.
  - ❑ (Most) Technologists and civil-liberties advocates: Key escrow is hard to implement securely and would boost foreign competitors of US technology companies.
  - ❑ Opponents of key escrow won this war.

- 2010's: Tech industry reacts to the Snowden revelations.
  - ❑ Broader and deeper use of E2E, default encryption.
  - ❑ Law enforcement (LE) claims that it is "going dark." It calls upon vendors to enable **LE access to locked devices – with a duly authorized warrant but without users' passcodes.**
  - ❑ Vendors object, saying that LEA would hurt customers' security and privacy.

- (Perfect) example: FBI vs. Apple

Brief Summary of the Case (1)

- Terrorists Syed Rizwan Farook and Tashfeen Malik shot up the San Bernandino, CA health-dept building where they worked, killing 14 and injuring 22.
- The FBI took possession of an iPhone that the health dept had issued to Farook. The phone was locked, Farook was dead, and exhaustive search of the passcode space would not work.
- The FBI asked Apple to unlock the phone.

James Comey (2014)

“Those charged with protecting our people aren’t always able to access the evidence we need to prosecute crime and prevent terrorism even with lawful authority. We have the legal authority to intercept and access communications and information pursuant to court order, but we often lack the technical ability to do so.”

Tim Cook (2016)

“The government is asking Apple to hack our own users and undermine decades of security advancements that protect our customers — including tens of millions of American citizens — from sophisticated hackers and cybercriminals. … We can find no precedent for an American company being forced to expose its customers to a greater risk of attack.”

Brief Summary of the Case (2)

- Apple said that it could not unlock an iPhone running iOS 9 without the user's passcode.
- FBI: Motion to compel Apple to develop software that would unlock ***this*** phone.
- Apple: Motion to dismiss
- The legal question remains unresolved: The FBI discovered that it could use a commercially available "gray-hat" hacking tool to unlock the phone, and it withdrew its motion to compel.

Competing Rights and Fears

- We have a right to privacy.
- Our governments routinely violate that right.
- The only effective grass-roots response to mass surveillance is mass encryption.

- Monopoly tech platforms have too much power.
- As “surveillance intermediaries,” they amass private data and decide whether law enforcement can access them.
- Citizens and their elected government should make those decisions.

Three Approaches to LE Access

- **Secure-systems community: Try to build it.**
  - ❑ Ozzie, Savage, and others: Protocols to unlock a device given an authenticated warrant and a manufacturer’s key
  - ❑ Wright and Varia: Protocols that use reduced-entropy key spaces discoverable in “crypto crumple zones”

Three Approaches to LE Access

- **Secure-systems community: Try to build it.**
  - ❑ Ozzie, Savage, and others: Protocols to unlock a device given an authenticated warrant and a manufacturer's key
  - ❑ Wright and Varia: Protocols that use reduced-entropy key spaces discoverable in "crypto crumple zones"
- **Crypto-research community: Fight it tooth and nail.**
  - ❑ Claim 1: It can't be done securely and cost-effectively.
  - ❑ Claim 2: It isn't needed: The sought-after information is often available in plaintext form (*e.g.*, in a social-media account); if all else fails, vulnerability-based, gray-hat hacking will always work.

Three Approaches to LE Access

- **Secure-systems community: Try to build it.**
  - ❑ Ozzie, Savage, and others: Protocols to unlock a device given an authenticated warrant and a manufacturer's key
  - ❑ Wright and Varia: Protocols that use reduced-entropy key spaces discoverable in "crypto crumple zones"
- **Crypto-research community: Fight it tooth and nail.**
  - ❑ Claim 1: It can't be done securely and cost-effectively.
  - ❑ Claim 2: It isn't needed: The sought-after information is often available in plaintext form (*e.g.*, in a social-media account); if all else fails, vulnerability-based, gray-hat hacking will always work.
- **JF: Use the crypto toolkit, and don't be hypocritical.**
  - ❑ Claim 1: Prove it!
  - ❑ Claim 2: Since when do we praise ubiquitous plaintext and buggy software?

Talk Outline

- Three examples
  - ❑ Security, privacy, encryption, and surveillance
  - ❑ **Data sharing in a decentralized online world**
  - ❑ Freedom of expression online (or the lack thereof)

- What do we mean by “And” and why it matters

- Pointers to two other CS-and-Law events

Privacy-Preserving, Interorganizational Data Sharing

(Idan-F., 2020 ACM Workshop on Privacy in the Electronic Society)

- Data owner: Organization that creates and owns data records
- Data subject: Each record contains data about an individual – the subject.
- Data client: Organization that uses the data records
- Data user: Employee of the client who uses records for assigned tasks

Example: Credit Reports

- Data owner: Credit-reporting agency **(CRA)**
- Data subject: Each record is a **credit report about an individual** – the subject.
- Data client: **Insurance company**
- Data user: **Insurance agent** who crafts pre-approved insurance-policy offers

- The **Fair Credit Reporting Act (1970)** empowers CRAs to collect and manage credit-related information about consumers. Companies (e.g., banks, insurance companies, and credit-card issuers) are allowed to access this information, but to do so they must have a **permissible purpose**. Currently, the three largest CRAs in the US are Experian, TransUnion, and Equifax.

Why is this a Hard Data-Sharing Problem?

- <u>Data owner</u>: Organization that creates and owns data records
- <u>Data subject</u>: Each record contains data about an individual – the subject.
- <u>Data client</u>: Organization that uses the data records
- <u>Data user</u>: Employee of the client who uses records for assigned tasks

- The parties have different and sometimes competing privacy and security concerns.
- Each organization uses its own “vocabulary,” *i.e.*, a set of proprietary metadata attributes that it does not want to share with other organizations.

Example: Credit Reports

- Data owner: Credit-reporting agency **(CRA)**
- Data subject: Each record is a **credit report about an individual** – the subject.
- Data client: **Insurance company**
- Data user: **Insurance agent** who crafts pre-approved insurance-policy offers

- CRA must comply with the “permissible purpose” requirement, *e.g.*, not give credit reports of people who don’t own cars to a user who is selling car insurance.
- Data subject does not want his credit report revealed to anyone except a user who at least *might* want to offer something that he’s interested in.
- Client wants credit reports that are actually relevant to its business purpose and needs to limit each user’s access to credit reports that are needed for her assigned tasks.

Additional Parties

- Intermediary organizations: Organizations that **enrich shared data** with additional information that is needed for clients' tasks, *e.g.*:
  - ❑ Internal Revenue Service
  - ❑ Department of Motor Vehicles
  - ❑ Banks, credit-card issuers, and other financial companies
- Security mediators (aka "proxies"): Each intermediary and client delegates to a proxy server the task of **translating data** into its vocabulary.
  - Proxies are semi-trusted (*i.e.*, "honest but curious").
  - Semi-trusted mediators were proposed by Boneh, Ding, Tsudik, and Wong (2001) and have since been used in many cryptographic protocols, including many with one or more corporate or government parties.

Additional Requirements

- The owner should be able to **encrypt the records once, store them on a cloud-service provider (CSP)** and not have to be online to authorize or serve clients' data requests.
- Data-access **policies** are specified by multiple parties (including the client). They may use attributes that are not directly comparable with the ones used by the owner to specify the data.
- **Privacy of payload data, attribute vocabularies, and auxiliary information** is guaranteed.
- Proxies can **update ciphertexts' attributes dynamically** according to up-to-date auxiliary information. Dynamic updates do not require re-encryption (or any other action) by the data owner.

Solution (1): Start with Attribute-Based Encryption (specifically, Key-Policy ABE)

- Each record is encrypted by an **encryptor** under a set of attributes.
- **Users'** private keys reflect decryption policies (aka access policies) defined by sets of attributes. Keys are issued offline by **trusted authorities** (TAs).
- Invented by Sahai and Waters (2005).
- Provides the kind of **flexible**, **fine-grained access** that we need, but …

- Assumes that encryptor and TAs use the same vocabulary
- Assumes that the payload can always be updated by the encryptor

Solution (1): Start with Attribute-Based Encryption (specifically, Key-Policy ABE)

- Each record is encrypted by an **encryptor** under a set of attributes.
- **Users'** private keys reflect decryption policies (aka access policies) defined by sets of attributes. Keys are issued offline by **trusted authorities** (TAs).
- Invented by Sahai and Waters (2005).
- Provides the kind of **flexible**, **fine-grained access** that we need, but …

X Assumes that encryptor and TAs use the same vocabulary

X Assumes that the payload can always be updated by the encryptor

Solution (2): Attribute-Based Encryption With Oblivious Attribute Translation (OTABE)

- Proxies **obliviously translate ciphertexts** on behalf of clients and intermediaries in a manner that guarantees **privacy of payloads, vocabularies, and auxiliary information**.
- Technical ingredients include:
  - ❑ "Large-universe" KP-ABE scheme of **Rouselakis and Waters (2013)**
  - ❑ **Secret sharing** of R and W's **binder term** – one share in the attribute components of the user and the translators.
- OTABE also
  - ❑ **Prevents unauthorized sharing** of keys by users
  - ❑ **Supports efficient revocation** of expired keys

Example: Credit Reports

C

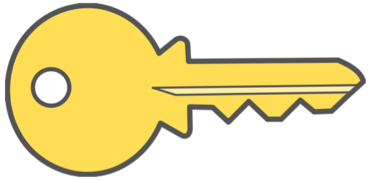

CREDIT_SCORE>600 ∧ #ACCIDENTS<2 ∧
IS-CREDIT-RATIO-LESS-THAN-AVERAGE=TRUE

O

Attributes:
CREDIT-UTILIZATION-RATIO=0.35,
CREDIT-SCORE=650,
DLN=xyz ,
DATE=11/9/2020

Payload:
Credit report

C

CREDIT_SCORE>600 ∧ #ACCIDENTS<2 ∧
IS-CREDIT-RATIO-LESS-THAN-AVERAGE=TRUE

I

Auxiliary information
Driving records

DLN -> #ACCIDENTS
DLN(xyz)=0

I

Auxiliary information
AVERAGE_CREDIT_UTILIZATION_RATIO
=0.4

O

Attributes:
CREDIT-UTILIZATION-RATIO=0.35,
CREDIT-SCORE=650,
DLN=xyz ,
DATE=11/9/2020

Payload:
Credit report

Relevance to Computer Science and Law

- OTABE is not useful for all interorganizational data-sharing governed by law.
  - ❑ Not for scenarios in which the data subject is available to give explicit consent (such as those governed by GDPR)
  - ❑ Not for scenarios with clear, efficiently decidable, and universal rules that govern which users can access which portions of the data

- It's designed for scenarios in which there are **legally mandated, general principles that govern access to sensitive data**, but **instantiating those principles in the form of efficiently decidable rules requires data-specific and dynamically changing knowledge.**
  - ❑ Fair Credit Reporting Act
  - ❑ Electronic Communications Privacy Act
  - ❑ Constitutional rights

Talk Outline

- Three examples
  - ❑ Security, privacy, encryption, and surveillance
  - ❑ Data sharing in a decentralized online world
  - ❑ **Freedom of expression online (or the lack thereof)**

- What do we mean by “And” and why it matters

- Pointers to two other CS-and-Law events

Freedom of Expression vs. Freedom from Harm

Section 230 of the Communications Decency Act (1996):

(1) No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.

(2) No provider or user of an interactive computer service shall be held liable on account of

(a) any action voluntarily taken in good faith to restrict access to or availability of material that the provider or user considers to be obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected; or

(b) any action taken to enable or make available to information content providers or others the technical means to restrict access to material described in paragraph (1).

Treat Social Media Platforms as Publishers?
(*i.e.*, throw out (1) in Section 230?)

Three approaches to "content moderation" at scale*:

- Editorial review: Modeled on traditional broadcasting and publishing, but might be automated to some extent
- Community flagging: Having transferred content creation from employees to users, do the same for content moderation.
- Automated detection: Use technology to lessen the severity of some problems inherent in "community" flagging.

*Tarleton Gillespie, **Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media**, Yale University Press, 2018

First Two Approaches Don’t “Scale” Adequately

- “Given the scale that Twitter is at, a one-in-a-million chance happens 500 times a day. … Say 99.999 percent of tweets pose no risk to anyone. There’s no threat involved. . . . After you take out that 99.999 percent, that tiny percentage of tweets remaining works out to roughly 150,000 per month.” (Del Harvey, vice president of Trust and Safety, Twitter, 2014)

- Facebook’s task is two orders of magnitude larger: It receives millions of community flags to review each day.

Mandate a 1st Amendment for Social-Media Platforms?

(*i.e.*, throw out (2) in Section 230?)

Jack Balkin at the inaugural ACM Symposium on CS and Law:

- **This is exactly what we should not do!** The 1st Amendment, like the rest of the Bill of Rights, restricts *government* power. Corporations aren't analogous to governments, *and we don't want them to be*.

- We want a broad range of social-media platforms, with varying content-moderation policies and a robust culture of professional responsibility. Together, they would constitute a *digital public square*.

“CS and Law” Research on Content Moderation

<u>Combine (the best of) human and automated moderation:</u>

- Explore and formalize “content-moderation policies.”
  - ❑ Exactly what are “community flaggers” supposed to flag?
  - ❑ Is there an (approximately) equivalent policy in which *much but not all* of this flagging can be automated?

- Use distributed algorithms, *e.g.*, collaborative-filtering schemes, to aggregate and act on large numbers of (human-generated) flags?

Talk Outline

- Three examples
  - ❑ Security, privacy, encryption, and surveillance
  - ❑ Data sharing in a decentralized online world
  - ❑ Freedom of expression online (or the lack thereof)

- **What do we mean by "And" and why it matters**

- **Pointers to two other CS-and-Law events**

Interdisciplinarity: “And = ∩” Will Be Hard!

Computer Science	Law
Papers are about results.	**Papers are about persuasion.**
Values breakthroughs	Values intellectual continuity
Outdated technology usually dies.	Outdated laws often live forever.
Powerful, new technologies will have unforeseen uses.	**Regulation can be based on “what are you going to use it for”?**
Often tries to ignore “political reality”	**Sometimes tries to ignore technical reality**

Examples in which “And = ∩” Has Been Hard

- US copyright law
  - Fundamentally at odds with digital distribution
  - This has been obvious for about 25 years.
  - Copyright industries fight change tooth and nail (of course).
  - ☹ **Legal scholars have suggested tweaks, not clean-slate redesign.**
- Ubiquitous encryption vs. lawful surveillance
  - ☹ **LE community “demands” technically sound access to locked devices: “Nerd harder.”**
  - ☹ **Crypto-research community says “this is just the key-escrow fight, which we have already won.”**

ACM INAUGURAL SYMPOSIUM ON
COMPUTER SCIENCE AND LAW

October 28, 2019, New York Law School, Tribeca, NYC

http://computersciencelaw.org/ (click on “Proceedings and Videos”)

How to Regulate
(and Not Regulate)
Social Media

What Cryptography
Can Bring to Law

When Code
Isn’t Law

DIMACS Workshop on Co-Development of Computer Science and Law (Nov. 10-12, 2020)

- From the Call for Participation: “The workshop will consider how to create technical definitions and solutions in concert with the creation of legal language so that the two fields can work together to solve (and proactively prevent) problems.”
- Sample of program:
  - ❑ Fireside chat by James Grimmelmann (Cornell) and Solon Barocas (Microsoft) about “CPU, Esq: Should Law Be More Like Software?”
  - ❑ Lecture by Jonathan Zittrain (Harvard) about “Intellectual Debt: What's Wrong When Machine Learning Gets It Right”
  - ❑ CS and Law Career Opportunities Panel (moderated by JF)
- Online workshop. There is no registration fee, but **participants must register by Nov. 6**: http://dimacs.rutgers.edu/ws-materials/cslaw