

CS37300

Data Mining & Machine Learning

Ethics Issues in Data Mining
Prof. Steve Hanneke and Chris Clifton
3 April 2023



Ethics Issues for Data Mining & ML

What's the Problem?

- Privacy
 - Training data
 - Allowed uses
- Fairness
 - Inequitable outcomes
 - Variance in accuracy
- Data inaccuracy
- Explainability
- Redress
 - What if someone disputes results?

Privacy: What's the Harm?

- My company collected the data, isn't it ours?
 - Or scraped it from the web
- It's just for training
 - We're producing models, not releasing data
 - Or not even releasing models, just using them
- You agreed to let us use your data

4

What is Privacy?

- “The right to be let alone” - *Warren & Brandeis, 4 Harvard L.R. 193 (Dec. 15, 1890)*
 - My information protected so it doesn't adversely affect me in the future
- Control over data
 - My information used only in ways I approve
- Issues:
 - Disclosure / sharing
 - Approved use
 - Recourse

5

Data Privacy: The Goal

- Protect the Individual
 - “Everyone has the right to the protection of personal data concerning him or her. Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data which has been collected concerning him or her, and the right to have it rectified.” – Charter of Fundamental Rights of the European Union
- Challenges: What do we mean by
 - “concerning” an individual
 - Protection
 - Consent
 - Access / rectified



6

“Obvious” answers

- Concerning an individual
 - Has your name/address/other identifying information
- Protection
 - Only used/accessed in expected, intended, authorized ways
- Consent
 - You know and agree to what is done with the data
- Access/Rectify
 - You can see the data and correct errors

7

Consent

- When you apply for a (job, grad school, ...), do you consent to that data being used with an ML model to decide if you should be accepted?
 - Amazon tried it:
<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- What about having your data used as training data to make decisions about others?
 - *Ungraded assignment (post-midterm): Read the terms of service and privacy policy of Facebook or some other social media you use, and think about this question.*

8

“Obvious” answers

- Concerning an individual
 - Has your name/address/other identifying information
- Protection
 - Only used/accessed in expected, intended, authorized ways
- Consent
 - You know and agree to what is done with the data
- Access/Rectify
 - You can see the data and correct errors

9

Concerning an Individual: IC 24-4.9-2-10

Sec. 10. "Personal information" means:

- (1) a Social Security number that is not encrypted or redacted; or
- (2) an individual's first and last names, or first initial and last name, and one (1) or more of the following data elements that are not encrypted or redacted:
 - (A) A driver's license number.
 - (B) A state identification card number.
 - (C) A credit card number.
 - (D) A financial account number or debit card number in combination with a security code, password, or access code that would permit access to the person's account.

10


The AOL Awakening

- In Aug 2006, AOL released its customers web searches for research studies
- 20 Million unique queries of 650K unique users
- <user-id> AOL fired its CTO over this issue;
- NY Times Two researchers were forced out
individual from the queries
 - Queries included “60 single men” “landscapers in Lilburn, Ga”
 - Many more queries contained enough information to uniquely identify the person
- *And it keeps going (Netflix, NYC Taxi, ...)*

11

Re-identifying “anonymous” data (Sweeney '01)

- 37 US states mandate collection of information
- Dr. Sweeney purchased the voter registration list for Cambridge Massachusetts
 - 54,805 people
- 69% unique on postal code and birth date
- 87% US-wide with all three



- Solution: k-anonymity
 - Any combination of values appears at least k times
- Developed systems that guarantee k-anonymity
 - Minimize distortion of results

12

Redaction: [IC 24-4.9-2-11](#)

(a) Data are redacted for purposes of this article if the data have been altered or truncated so that not more than the last four (4) digits of:

- (1) a driver's license number;
- (2) a state identification number; or
- (3) an account number;

is accessible as part of personal information.

(b) For purposes of this article, personal information is "redacted" if the personal information has been altered or truncated so that not more than five (5) digits of a Social Security number are accessible as part of personal information.

13

Anonymity: The Goal

- Prevent Disclosure of Personal Information
 - GDPR: ‘personal data’ means any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly
 - Qatar Law 13 of 2016: Personal Data: Data belonging to an Individual with specified or reasonably specifiable identity whether through such Personal Data or through combining the same with any other data
 - *But still use the data where appropriate!*
- Problem: It can’t be done!
 - “Perfect” privacy requires zero utility (e.g., the data must be encrypted.)
 - As soon as we can use the data (e.g., decrypt), it is at risk

15

Why Perfect Privacy is Impossible (Dwork, McSherry, Nissim, and Smith ‘06)

- Background Knowledge
 - Adversary may already know a lot
 - Whatever we provide (even de-identified or anonymized data) may add to that knowledge
- It may just take that “last bit of knowledge” to give the adversary the ability to violate privacy
 - *We can formally prove 1 bit may be too much*

16

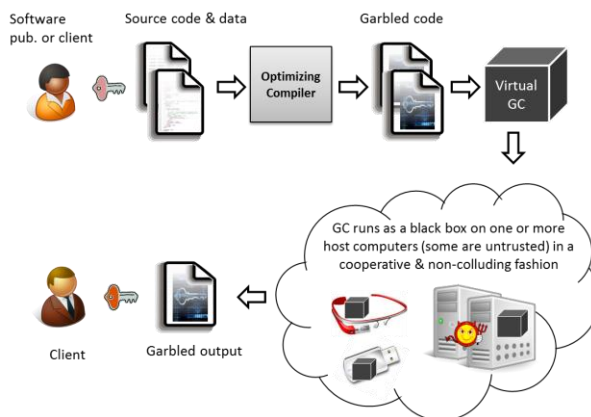
What We Can Do

- Encryption
 - Reduce risk to minimal levels when data not in use
- Anonymization
 - Produce usable data that is hard to link to individuals
- Noise addition
 - Usable data where any link to individuals (or information we surmise about individuals) is guaranteed to be uncertain/suspect

What We Can Do: Encryption

- Goal: Reduce risk to minimal levels when data not in use
- Encrypted Computation
 - Process the data while it is encrypted
 - Decrypt final output: Generalized, non-individual results
- Basic tools
 - Homomorphic Encryption, Commutative Encryption, Order Preserving Encryption
- Research Prototypes can accomplish many data processing and analysis tasks using these tools
 - Garbled Computing: Compute without revealing either the data or the program

- Garbled Computing.



What We Can Do: Anonymization

- Ensure protected/sensitive data not directly identifiable
 - Remove links between protected data and identifiers
- Generalize “quasi-identifiers”: Information that when combined with external data enables re-identification
 - Birth dates, addresses, workplace, etc.
 - E.g., instead of birth date, only give year
- Anonymized data still useful for data analysis
 - Goal is general knowledge, not learning specifics about individuals
- Example: “Anatomized” database from “Private Data in the Cloud” project


Patient	ID
Roan	1
Lisa	2
Roan	3
Elyse	4
Carl	5
Roan	6
Lisa	7
Roan	8

ID	Manufacturer	Drug Name
	Raphe Healthcare	Retinoic Acid
	Raphe Healthcare	Retinoic Acid
	Raphe Healthcare	Retinoic Acid
	Envie De Neuf	Mild Exfoliation
	Emedoutlet	Nexium
	Gep-Tek	Abiraterone
	Jai Radhe	Adapalene
	Hangzhou Btech	Cytarabine

20


What We Can Do: Noise Addition

- Idea: Impact of noise on what we learn from the data larger than impact of any individual’s data
- Formally: For $S \subseteq \text{Range}(f)$, an ϵ -differentially private mechanism M satisfies $\frac{\Pr[M_f(D_1) \in S]}{\Pr[M_f(D_2) \in S]} \leq e^\epsilon$ where D_1 and D_2 differ on at most one element
- U.S. Census Bureau is starting to use Differential Privacy



21

What We Need: Legal Incentives

- “Notice and Consent” framework discourages application of technological advances
 - We can’t guarantee your privacy, so please allow us to use your data in unsafe ways
 - U.S.: [Enforcement action against Snapchat](#) for promising to protect privacy and not doing a good enough job 
 - Companies get away with not even trying, as long as they tell you so
- Can legal frameworks acknowledge that privacy is at risk?
 - Require efforts to manage, not eliminate, that risk