Trust and Technology: Growing Gaps and Urgent Opportunities

MIT Future of Data Initiative
January 21, 2021

Daniel J. Weitzner
3Com Founders Principal Research Scientist, MIT CSAIL
Founding Director, MIT Internet Policy Research Initiative

Director, CSAIL Computing and Society Community of Research

Internet Policy Research Initiative

MIT CSAIL

Four Decades of Computing Leadership: 1982 - 2018

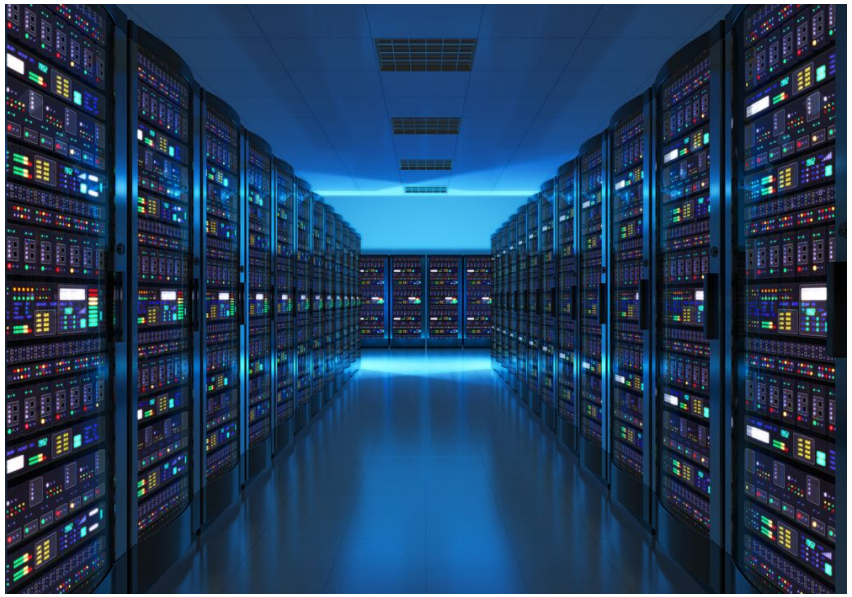

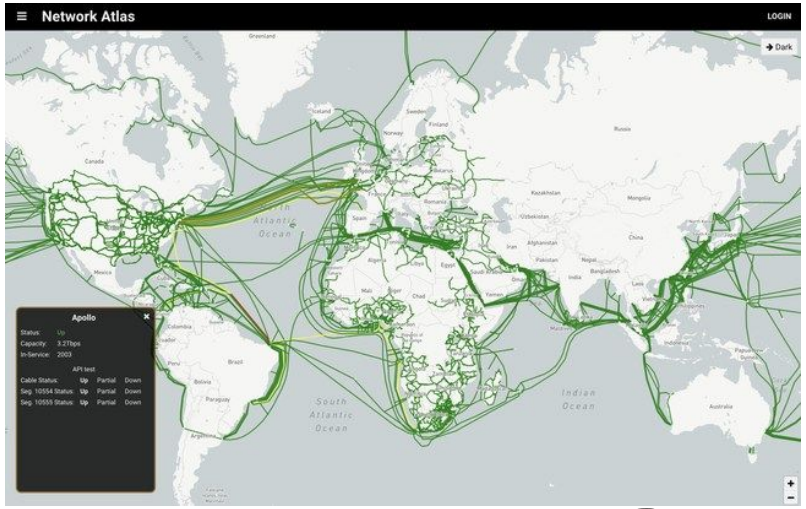

Internet Policy Research Initiative

MIT CSAIL

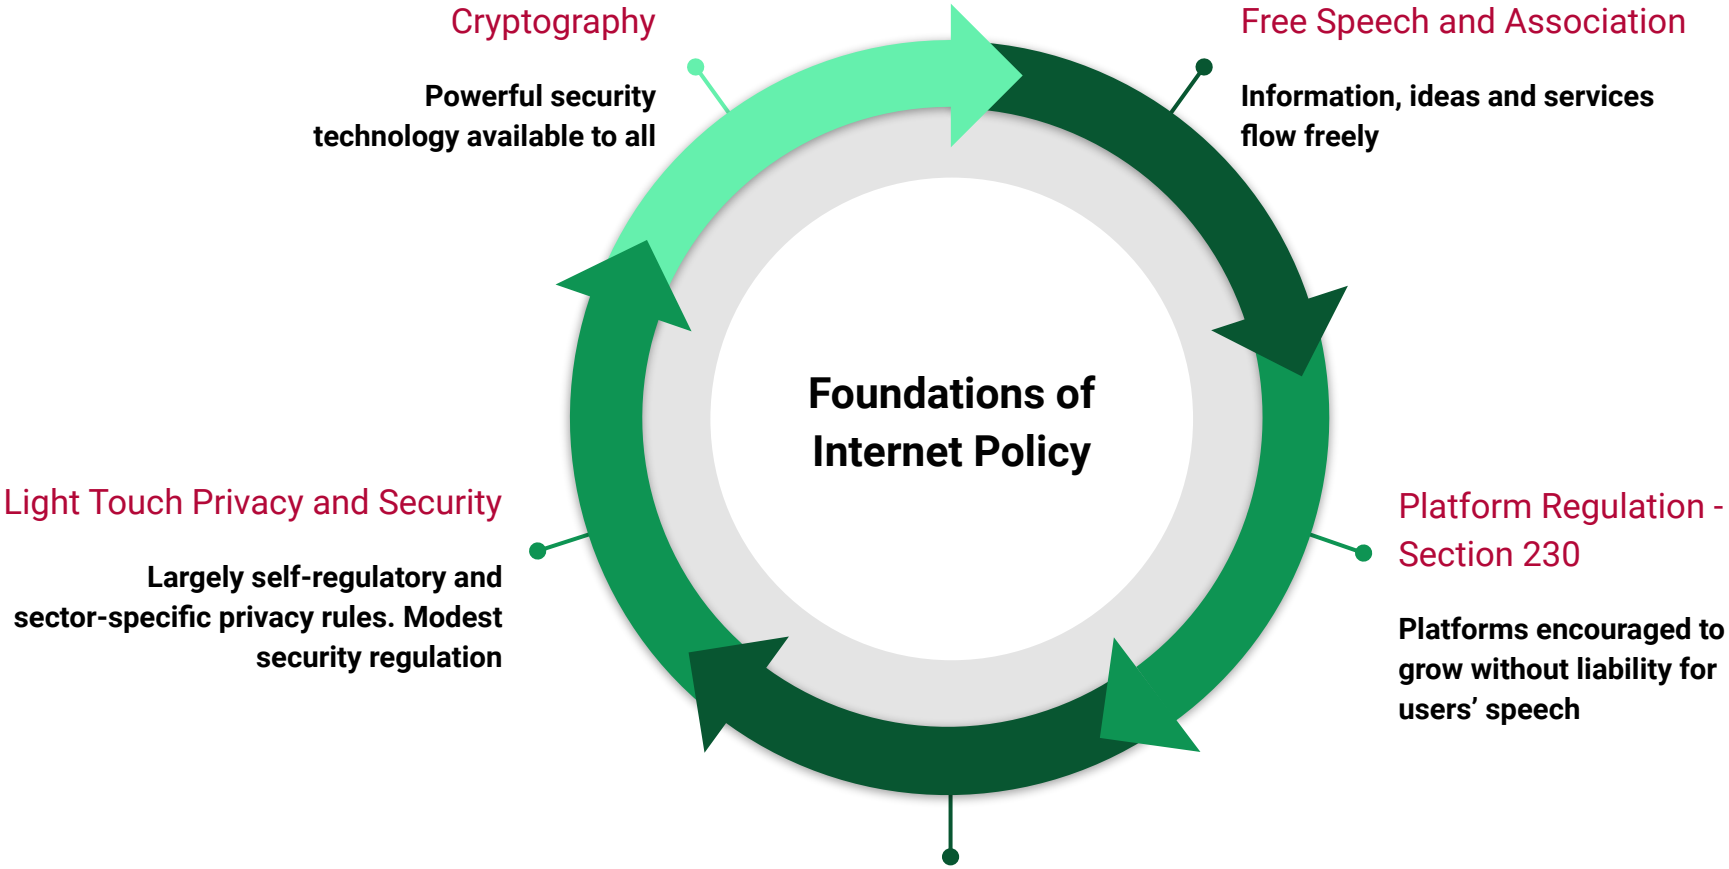

Internet Policy Research Initiative

MIT CSAIL

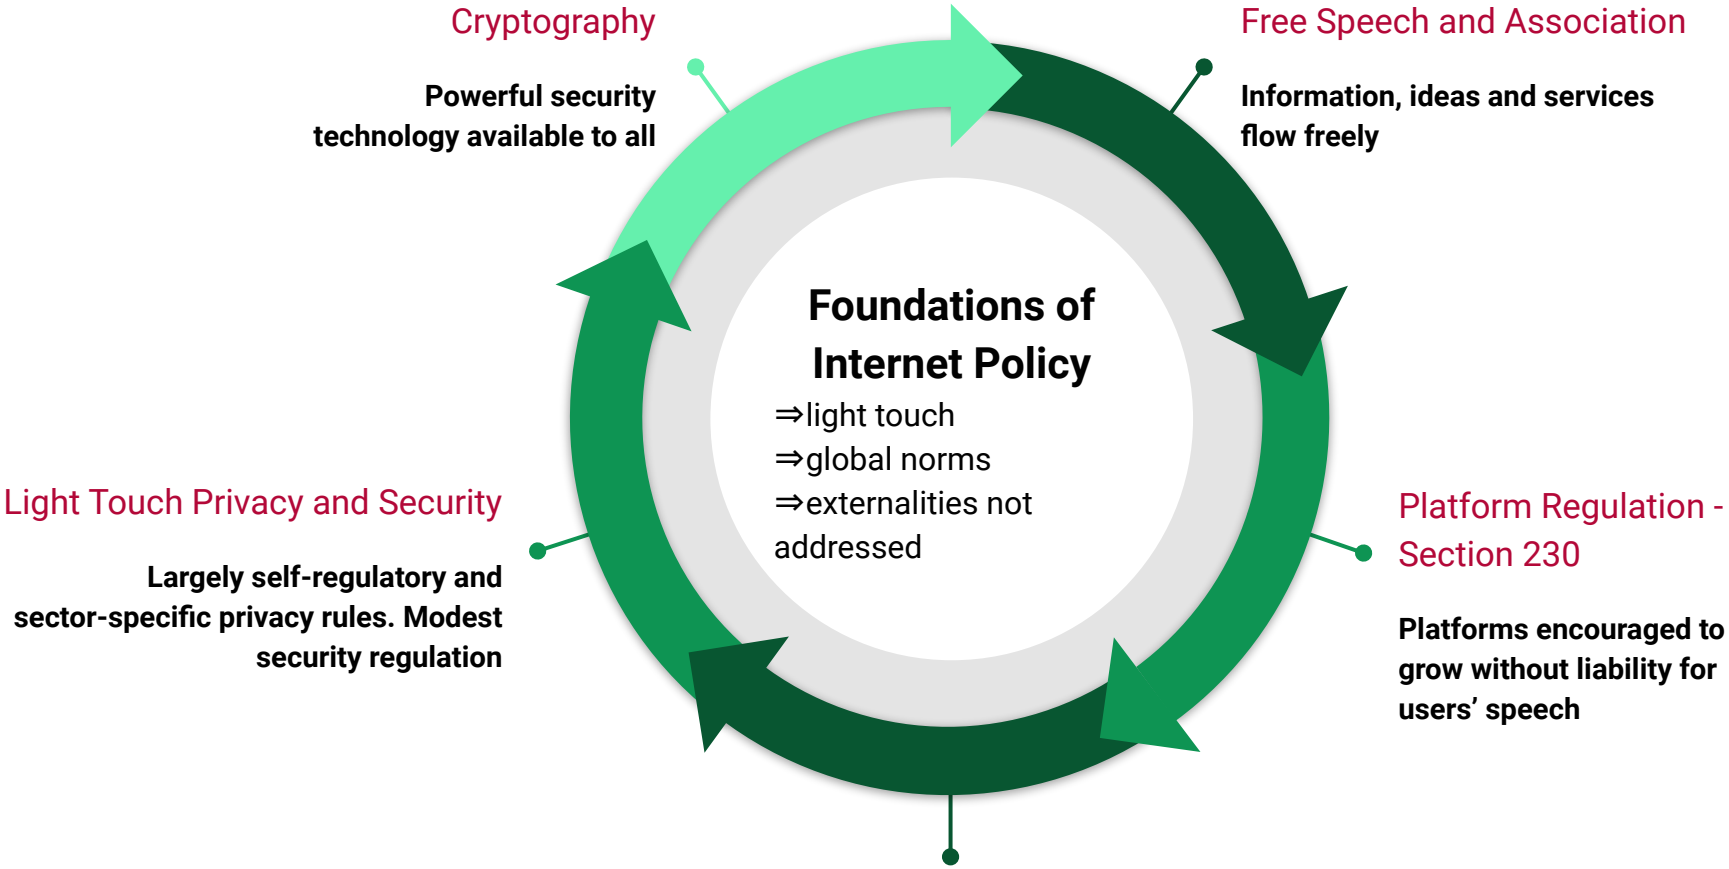
Cryptography
Powerful security technology available to all
Free Speech and Association
Information, ideas and services flow freely
Foundations of Internet Policy
⇒light touch
⇒global norms
⇒externalities not addressed
Light Touch Privacy and Security
Largely self-regulatory and sector-specific privacy rules. Modest security regulation
Platform Regulation - Section 230
Platforms encouraged to grow without liability for users' speech
Global Interoperability Standards
Global, industry-driven technical standards tie the Internet together

$6T+/3B Person-enabling Internet Policy framework

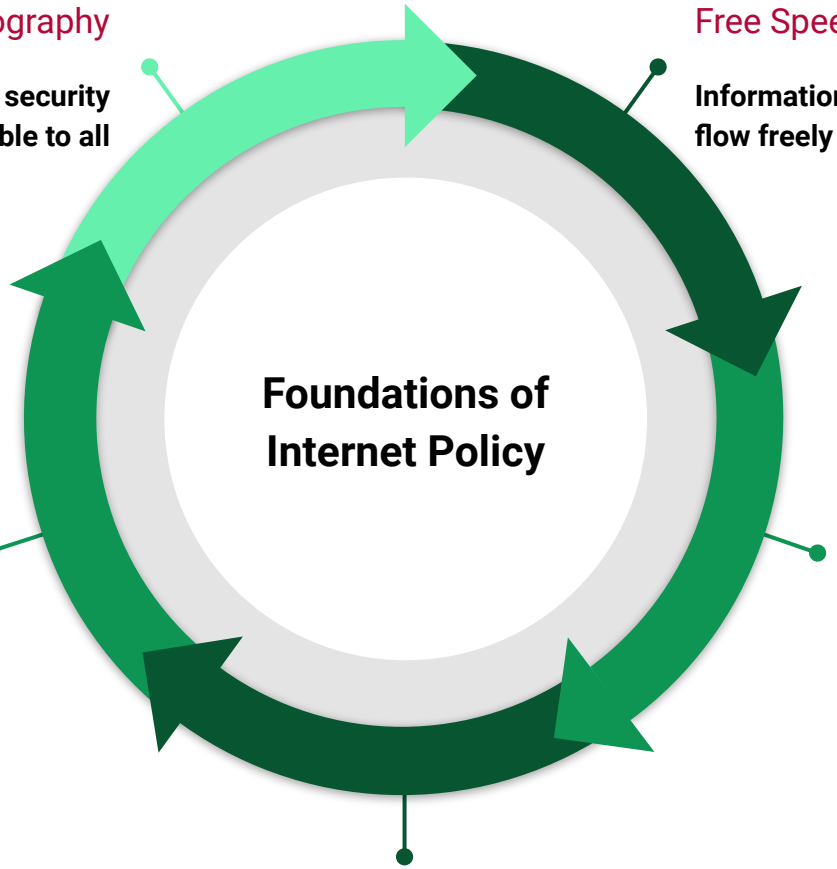

$2.2T

Cryptography

Powerful security technology available to all

Free Speech and Association

Information, ideas and services flow freely

$781B

Light Touch Privacy and Security

Largely self-regulatory and sector-specific privacy rules. Modest security regulation

Platform Regulation - Section 230

Platforms encouraged to grow without liability for users' speech

$1.2T

$1.6T

Global Interoperability Standards

Global, industry-driven technical standards tie the Internet together

Internet Policy Research Initiative

MIT CSAIL

Gaps Emerge, Opportunities too.

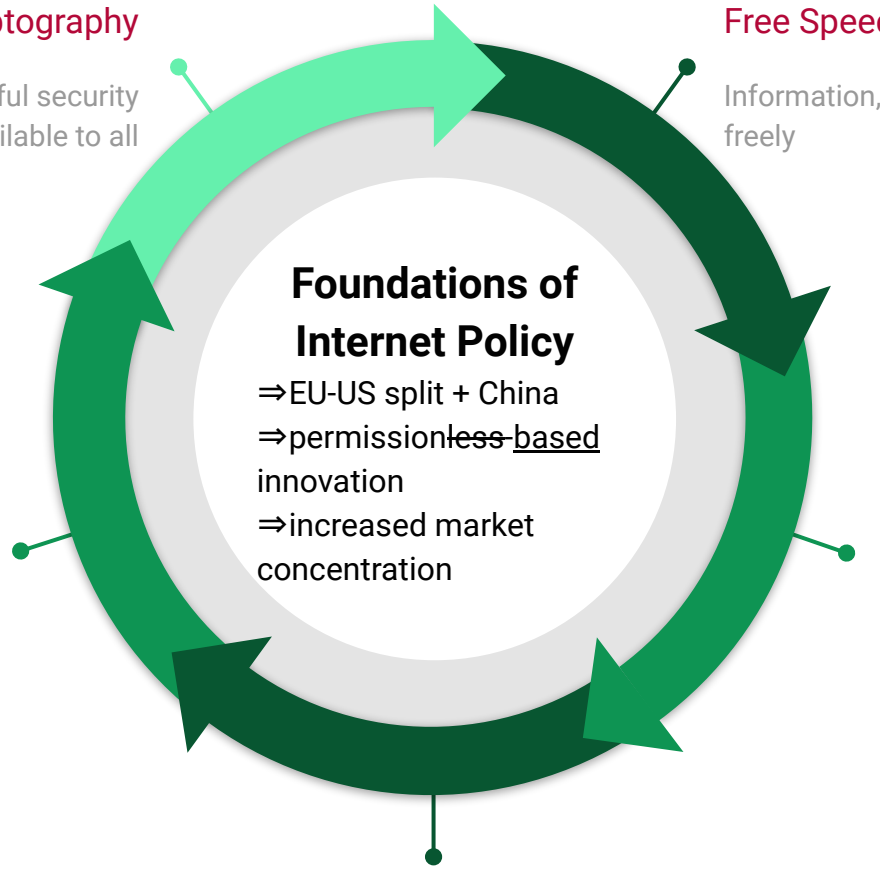

Cryptography

Powerful security technology available to all

Free Speech and Association

Information, ideas and services flow freely

Light Touch Privacy and Security

Gap: GDPR is new global norm
Risk: slower data innovation, bleed from privacy to AI regulation
Opportunity: Data Governance, AI Reg, New security metrics

Platform Regulation - Section 230

Platforms encouraged to grow without liability for users' speech

Global Interoperability Standards

Global, industry-driven technical standards tie the Internet together

Internet Policy Research Initiative

MIT CSAIL

Gaps Emerge, Opportunities too.

Cryptography

Gap: pressure to reduce E2E crypto to help law enforcement
Risk: market-wide security vulns
Opportunity: new secure computing techniques

Free Speech and Association

Information, ideas and services flow freely

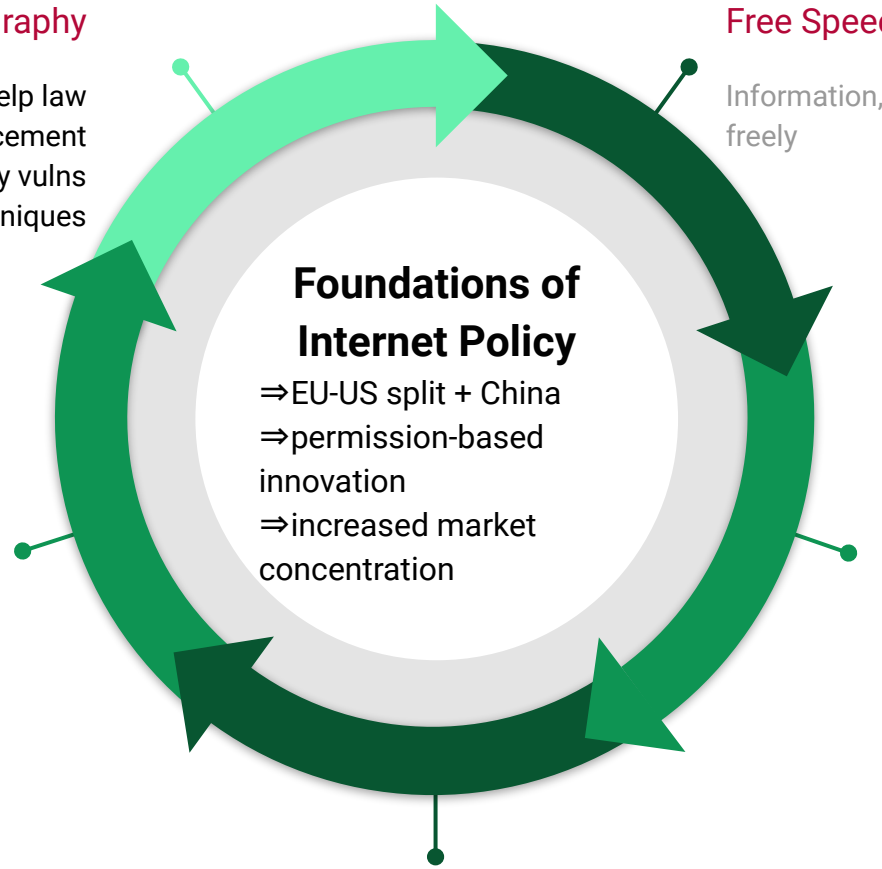

Light Touch Privacy and Security

Gap: GDPR is new global norm
Risk: slower data innovation, bleed from privacy to AI regulation
Opportunity: Data Governance, High-trust AI Rules, New security metrics

Platform Regulation - Section 230

Platforms encouraged to grow without liability for users' speech

Global Interoperability Standards

Global, industry-driven technical standards tie the Internet together

Gaps Emerge, Opportunities too.

Cryptography

Gap: pressure to reduce E2E crypto to help law enforcement
Risk: market-wide security vulns
Opportunity: new secure computing techniques

Free Speech and Association

Information, ideas and services flow freely

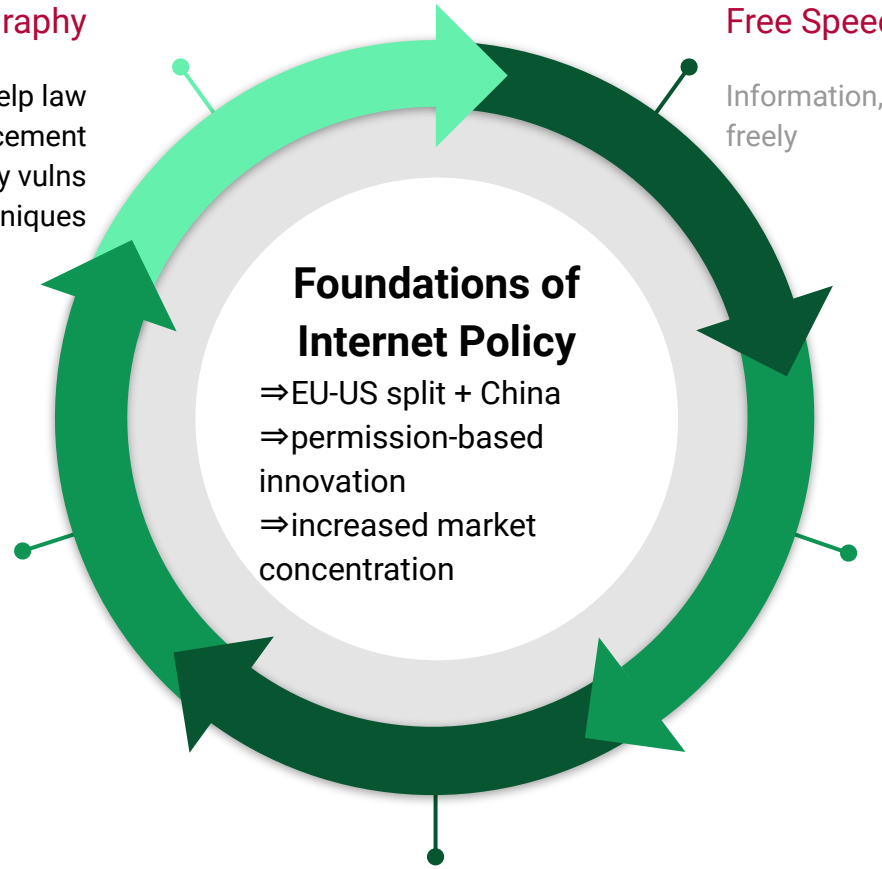

Light Touch Privacy and Security

Gap: GDPR is new global norm
Risk: slower data innovation, bleed from privacy to AI regulation
Opportunity: Data Governance, High-trust AI Rules, New security metrics

Platform Regulation - Section 230

Gap: Pressure to address disinformation, curb crime
Risk: increased compliance burden -> market concentration and reduced services innovation
Opportunity: new content moderation tech and norms

Global Interoperability Standards

Global, industry-driven technical standards tie the Internet together

Gaps Emerge, Opportunities too.

Cryptography

Gap: pressure to reduce E2E crypto to help law enforcement
Risk: market-wide security vulns
Opportunity: new secure computing techniques

Free Speech and Association

Gap: IoT and sensor networks put pressure on freedom of association.
Risk: Chilling effect or tech rejectionism
Opportunity: Policy and best practices for sensor networks and IoT collection sensitive data.

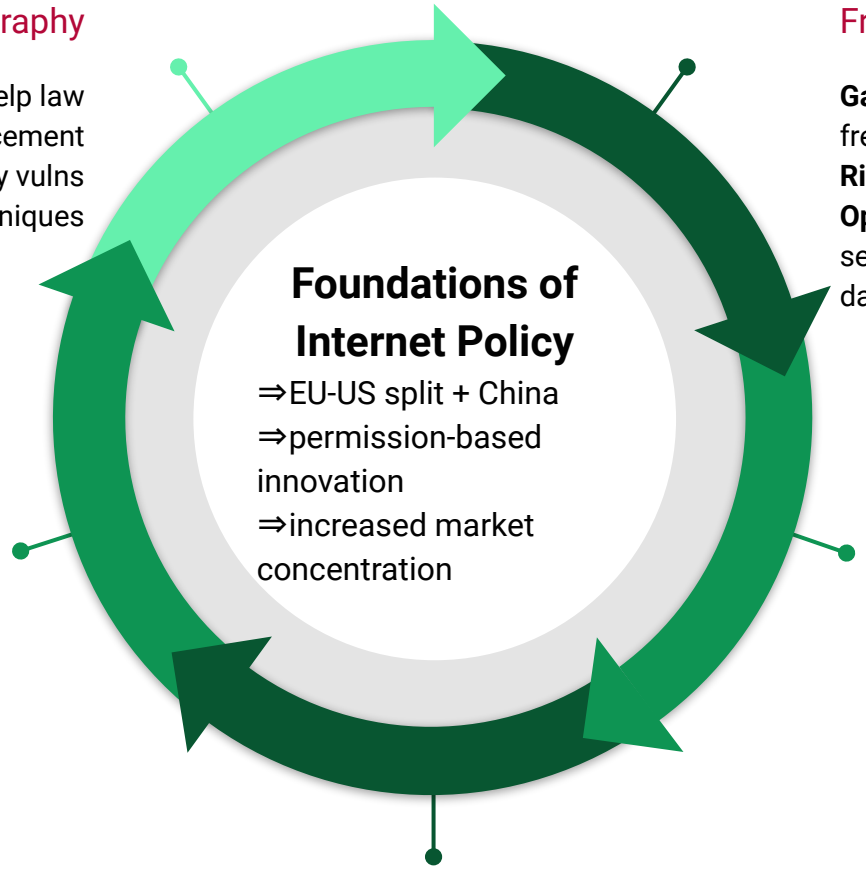

Light Touch Privacy and Security

Gap: GDPR is new global norm
Risk: slower data innovation, bleed from privacy to AI regulation
Opportunity: Data Governance, High-trust AI Rules, New security metrics

Platform Regulation - Section 230

Gap: Pressure to address disinformation, curb crime
Risk: increased compliance burden -> market concentration and reduced services innovation
Opportunity: new content moderation tech and norms

Global Interoperability Standards

Gap: National-security (US-EU-UK)/mercantilist (China) policies
Risk: balkanization -> lost Metcalfe's law value
Opportunity: empirically-based security/reliability and new trade policies

Frontiers of Computing no longer purely technical

M.I.T. Plans College for Artificial Intelligence, Backed by $1 Billion

The Massachusetts Institute of Technology is taking a particularly ambitious step in preparing students to develop, and consider the implications of, artificial intelligence. It is creating a new college, backed by a planned investment of $1 billion. Cody O'Loughlin for The New York Times

By Steve Lohr

Oct. 15, 2018

"The MIT Schwarzman College of Computing will seek to be not only a center of advances in computing, but also **a place for teaching and research on relevant policy and ethics to better ensure that the groundbreaking technologies of the future are responsibly implemented in support of the greater good.**"
-MIT President Rafael Reif

Future of Data: Research Agenda

SCRAM

Secure Cyber Risk Aggregation and Measurement

de Castro, L., Lo, A. W., Reynolds, T., Susan, F., Vaikuntanathan, V., Weitzner, D., & Zhang, N. (2020). SCRAM: A Platform for Securely Measuring Cyber Risk . Harvard Data Science Review. https://doi.org/10.1162/99608f92.b4bb506a

Problem: Cyber attacks happen all the time, but we collectively learn very little from them because firms are reluctant to disclose how they were attacked and the magnitude of their losses.

Gap: Lack of cyber risk pricing models:

- Impair CISO investment and prioritization decisions
- Limit quality and efficiency of cyber insurance
- Leave policymakers and regulators making uninformed choices in setting security standards
- Sow public distrust

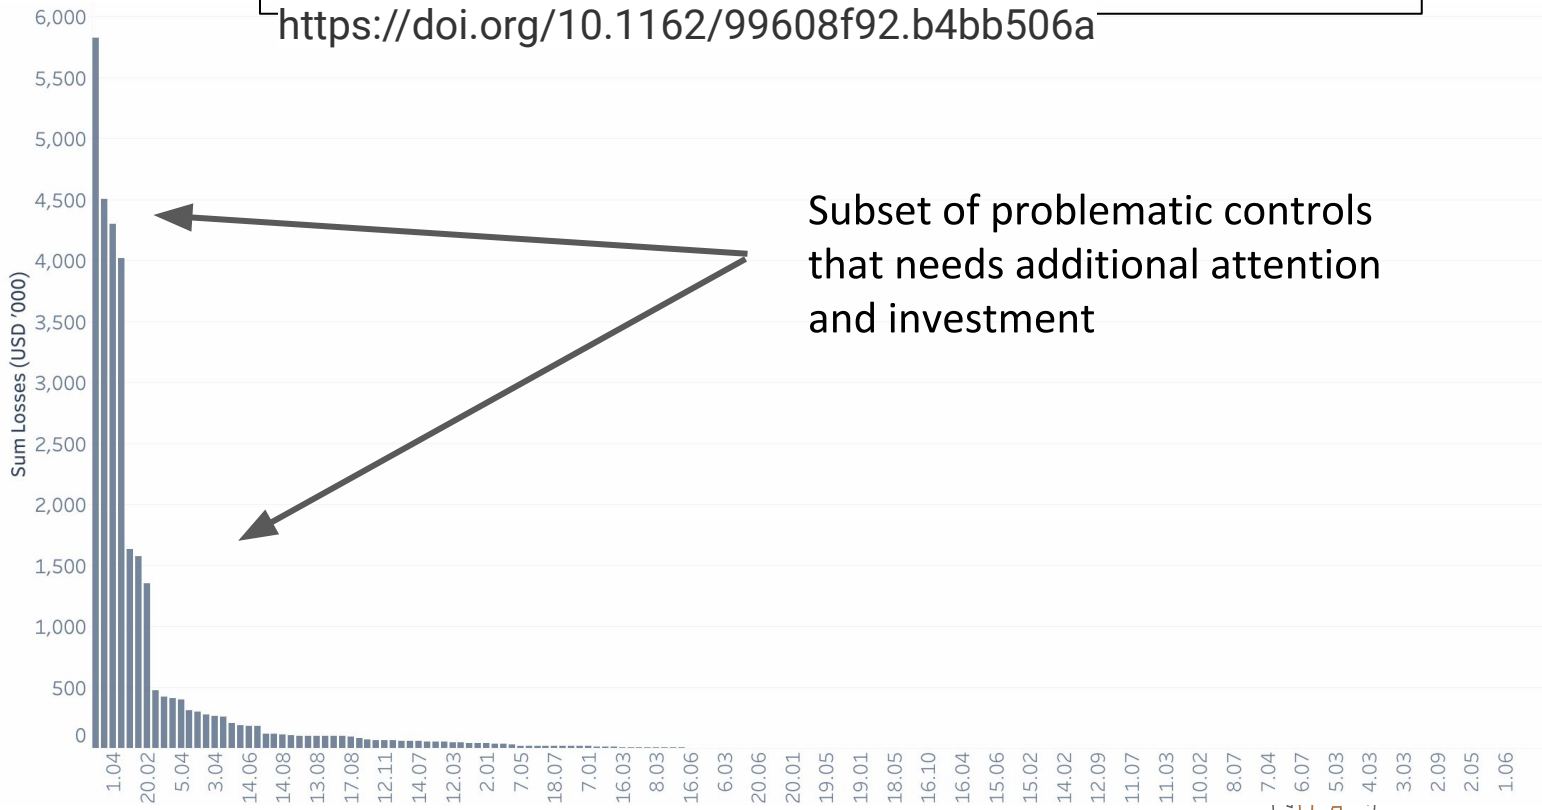

Subset of problematic controls that needs additional attention and investment

https://scram.mit.edu

Internet Policy Research Initiative

MIT CSAIL

We built a solution

Using our new cryptographic platform (multi-party computation), firms can securely and privately contribute sensitive data for calculating aggregate frequency and loss data without disclosure to anyone - including MIT!

Homomorphic encryption → Elegant way of computing on encrypted data

SCRAM
Secure Cyber Risk Aggregation and Measurement

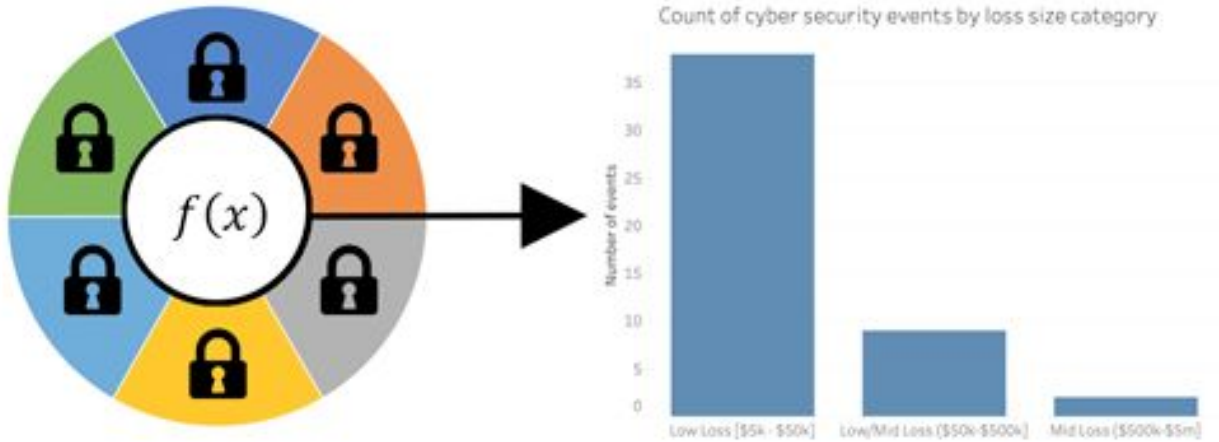

scram.mit.edu

Early results: Losses by category

“Implemented” does not mean protected

Correlation of control adoption with losses

10.3%

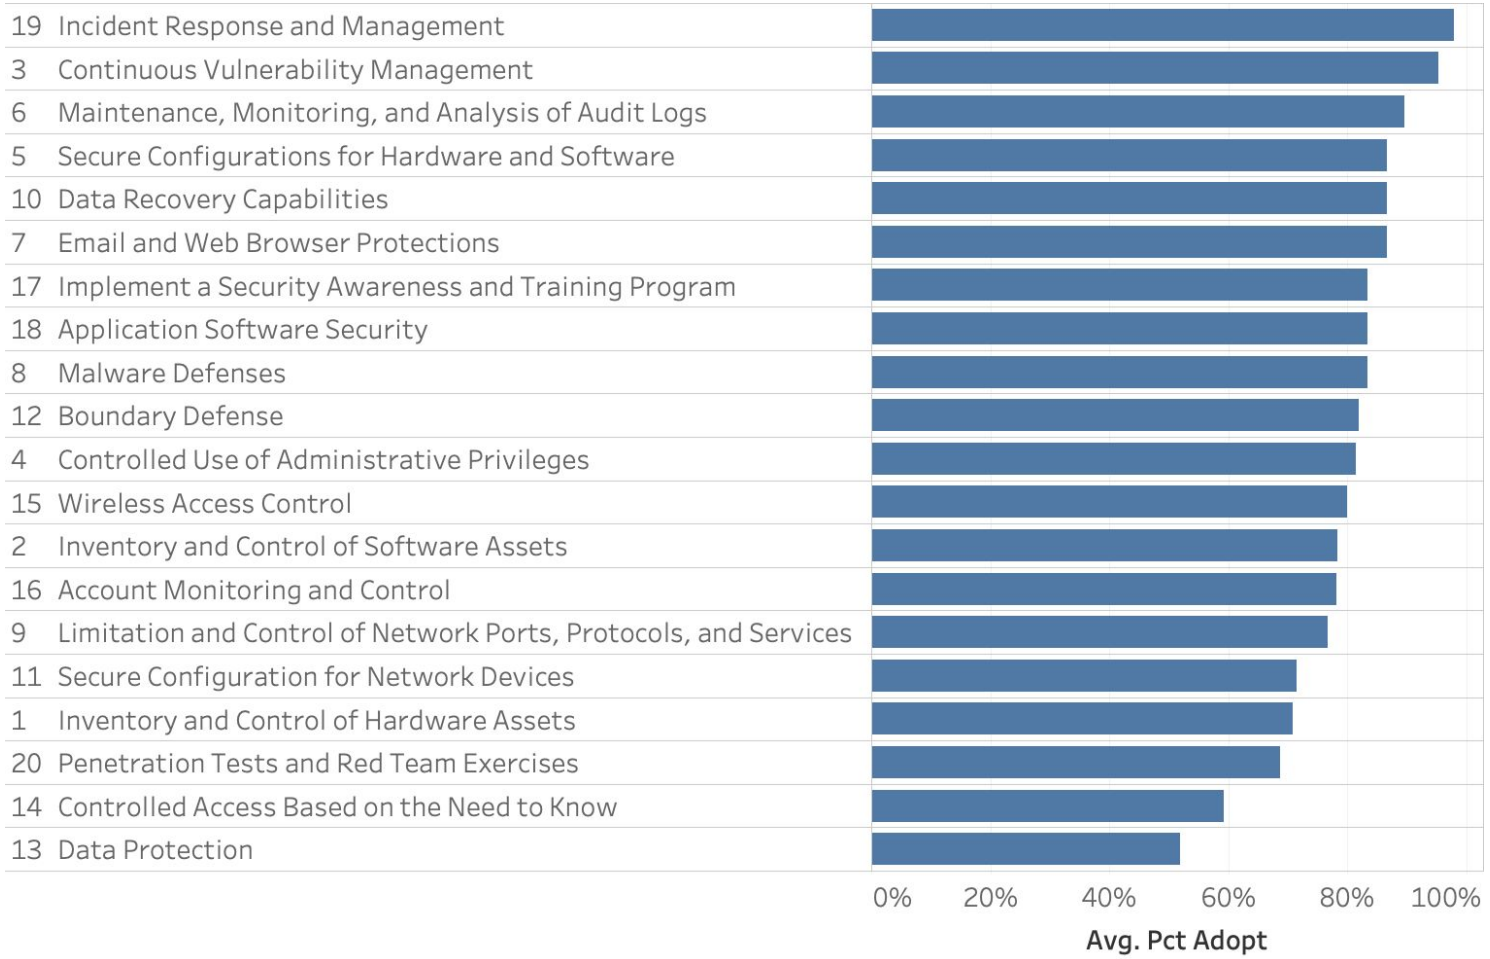

Internet Policy Research Initiative

MIT CSAIL

Tools: GDPR-aware database architecture - unsolved problems

GDPR	Consumer Privacy Bill of Rights	Cal. Consumer Privacy Act	Apple proposal
Lawful basis - Consent, etc.	Right individual control		
legitimate interest,etc	Right to respect for context		
Right to be informed	Right to transparency	Right to kn	Right to know
Right of access	Right to a		Right to access
Right to rectification	Right to accu		
Right to erasure		Rig a deleted	
Right to restrict processing		No discrimination for exercise	Right to minimization
Right to portability	Machine- s		
Right to object	Right to c		
Right to avoid automated decisionmaking & explanation			
Data Breach Notification	Security & Breach Notification		rity
Accountability	Accountability		
Fines < 4% annual revenue	Fines	$7500/incident, 30 day cure	

Purpose limitation enforcement

Effective Notice - HCI/UX

Hard Delete

Graph privacy

ML explanation

Policy-Aware Event Logging

Purpose-Aware Database Architecture

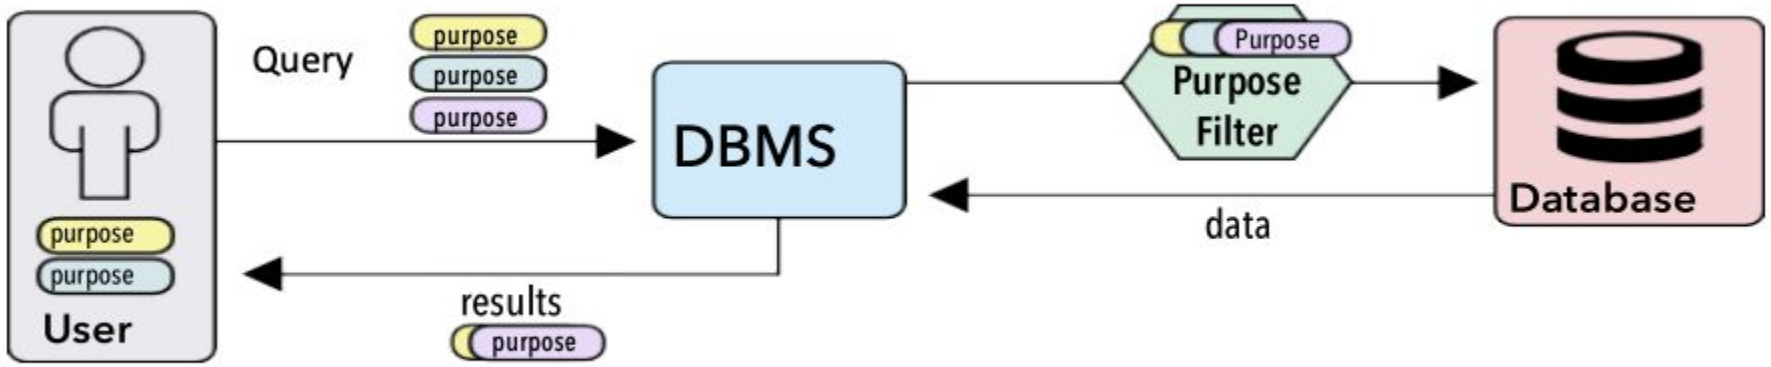

Stonebraker, M., Brodie, M., Kraska, T., Servan-Schreiber, S., Weitzner, D. J., SchengenDB: A Data Protection Database., VLBD Workshop Poly'19

Voatz

Internet Policy Research Initiative

MIT CSAIL

Bringing Technical Rigor to Policy Debates

The New York Times

Voting on Your Phone: New Elections App Ignites Security Debate

New Hampshire residents voting in the presidential primary on Tuesday. A smartphone app could let some absentee voters this year cast ballots from home.
Alyssa Schukar for The New York Times

By Matthew Rosenberg

Feb. 13, 2020

In the new paper, the M.I.T. researchers, Michael A. Specter, James Koppel and Daniel J. Weitzner, go beyond speculation and detail how they found serious security issues by reverse-engineering Voatz's app and recreating what they could of the company's server from publicly available information.

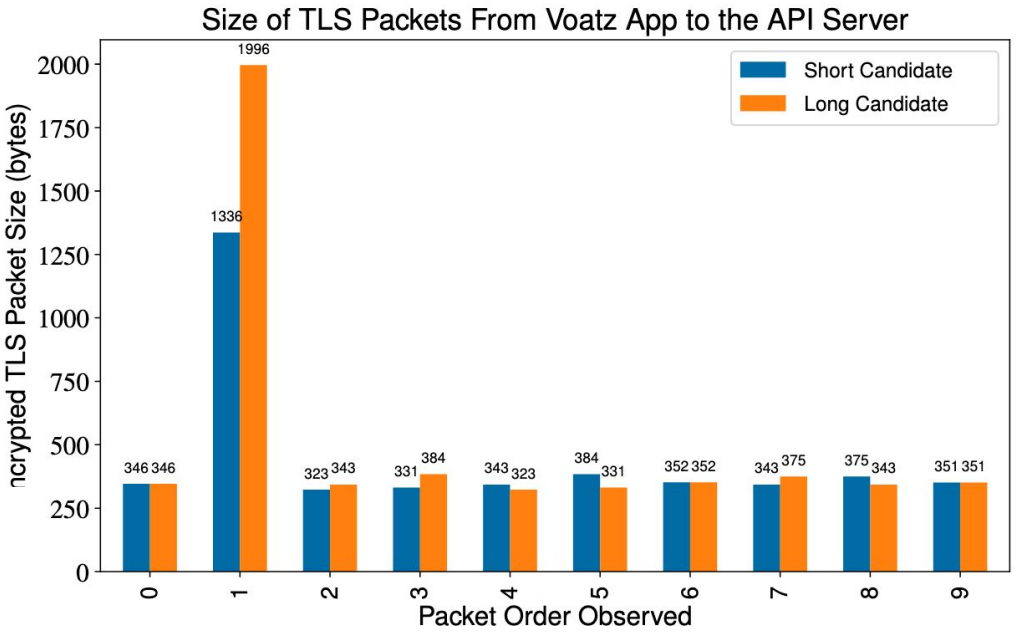

Specter, Michael A., James Koppel, and Daniel Weitzner. "The Ballot is Busted Before the Blockchain: A Security Analysis of Voatz, the First Internet Voting Application Used in US Federal Elections." In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pp. 1535-1553. 2020.

Internet Policy Research Initiative

MIT CSAIL

Findings on mobile voting security

- 5 high-severity vulnerabilities & a serious privacy issue
- Many basic implementation failures, e.g.:
  - Mandated use of weak passwords
  - Anti-tamper/AV solution was easily circumventable
  - Sends a photo of user's ID, and location to a third party *without alerting the user*
- API Server has complete control
  - No proofs of inclusion (where's the Blockchain?)
  - Weak receipt validation, not E2E-V

Adversary	**Attacker Capability**				
	Suppress Ballot	Learn Secret Vote	Alter Ballot	Learn User's Identity	Learn User IP
Passive Network (§5.3)		✓			✓
Active Network (§5.3)	✓	✓			✓
3rd-Party ID Svc. (§5.4)	✓			✓	✓
Root On-Device (§5.1)	✓	✓	✓	✓	✓
Voatz API Server (§5.2)	✓	✓	✓	✓	✓

Internet Policy Research Initiative

MIT CSAIL

Future of Data: Engagement and Education

Tools: Engagement

"All the News That's Fit to Print"

The New York Times

VOL. CLXIV . . . No. 56,921 NEW YORK, WEDNESDAY, JULY 8, 2015

Security Experts Oppose Government Access to Encrypted Communication

By NICOLE PERLROTH JULY 7, 2015

SAN FRANCISCO — An elite group of security technologists has concluded that the American and British governments cannot demand special access to encrypted communications without putting the world's most confidential data and critical infrastructure in danger.

The New York Times

A.I. Policy Is Tricky. From Around the World, They Came to Hash It Out.

Nicolas Miailhe, a co-founder of the Future Society, asking a question during a gathering of global policymakers last week at the Massachusetts Institute of Technology.
Kayana Szymczak for The New York Times

By Steve Lohr

Jan. 20, 2019

MIT News

ON CAMPUS AND AROUND THE WORLD

Browse

AI, the law, and our future

MIT "Policy Congress" examines the complex terrain of ar regulation.

Peter Dizikes | MIT News Office
January 18, 2019

Press Inquiries

Scientists and policymakers converged at MIT on Tuesday to discuss one of the hardest problems in artificial intelligence: How to govern it.

The first MIT AI Policy Congress featured seven panel discussions sprawling across a variety of AI applications, and 25 speakers — including two former White House chiefs of staff, former cabinet secretaries, homeland security and defense policy chiefs, industry and civil society leaders, and leading researchers.

Internet Policy Research Initiative

MIT CSAIL

Computing & Society CoR Student Research

ML : Diagnosing bias in predictive modeling

Creating and democratizing conversational AI

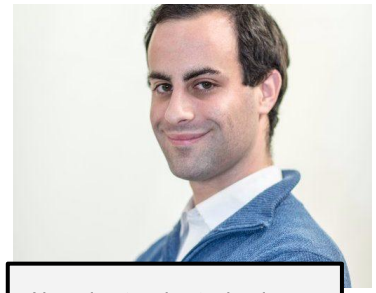

Neural networks, technology policy, and the Lottery Ticket Hypothesis

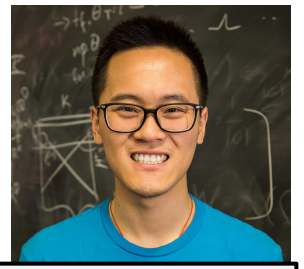

App Inventor voice-based conversational programming system

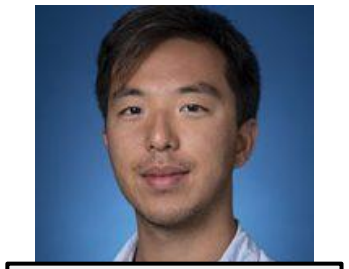

ML Bias in Predictive Policing

Differential Privacy for Machine Learning

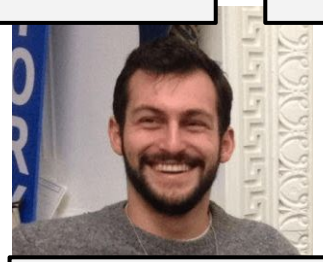

How organizations develop structures for technology policy decision making

Efficient two-party computation from plaintext-ciphertext homomorphic operations

Personal audio classifier for App Inventor

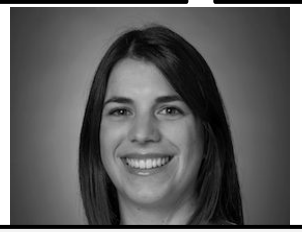

The impact of vulnerabilities and benefits of security in the Internet's core routing protocols

Cybersecurity policy, trust, and privacy online

International security, the political economy of technology, and Chinese foreign policy

Denial of Service vulnerabilities, network protocols, Linux performance bottlenecks, and Future Internet Architectures

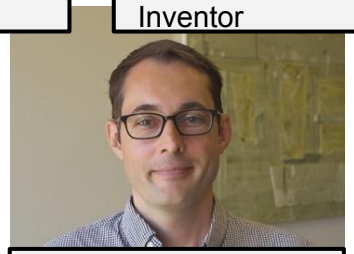

Evaluating the quasi-voluntary public-private partnership approach to critical infrastructure cybersecurity

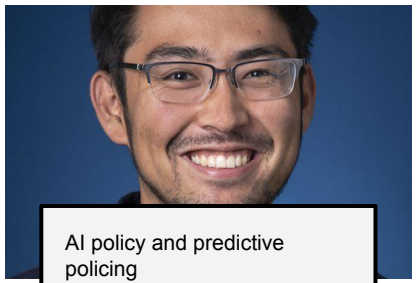

AI policy and predictive policing

Systems security and applied cryptography, particularly in areas relevant to public policy

Internet Policy Research Initiative

MIT CSAIL

https://internetpolicy.mit.edu/

Internet Policy Research Initiative

MIT CSAIL

Techniques for Data Privacy: "Old" and New

Srini Devadas

MIT CSAIL

Motivation

- Being able to compute on private data is essential
- Many questions:
  - What is the privacy guarantee? Does it conform to applicable laws governing use and sharing of data?
  - What is the computational cost of ”private” computation?
  - What is the utility or accuracy cost of “private” computation?
- Not surprisingly, there are many approaches with differing tradeoffs

Approach 1: "Anonymize" the Data

Name	Sex	Blood	…	HIV?
Chen	F	B	…	Y
Jones	M	A	…	N
Smith	M	O	…	N
Ross	M	O	…	Y
Lu	F	A	…	N
Shah	M	B	…	Y

Problems?

Reidentification via Linkage

Name	Sex	Blood	…	HIV?
Chen	F	B	…	Y
Jones	M	A	…	N
Smith	M	O	…	N
Ross	M	O	…	Y
Lu	F	A	…	N
Shah	M	B	…	Y

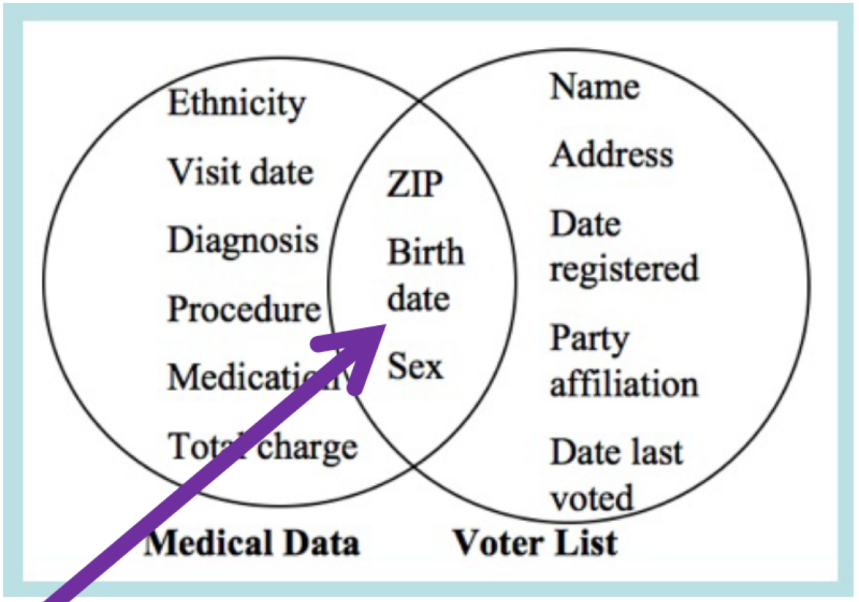

[Sweeney `97]

Uniquely identify > 60% of the US population [Sweeney `00, Golle `06]

All it takes is a knowledge of a small number of attributes to identify/name the person!

Approach 2: Encrypt the Data

Name	Sex	Blood	⋯	HIV?
Chen	F	B	⋯	Y
Jones	M	A	⋯	N
Smith	M	O	⋯	N
Ross	M	O	⋯	Y
Lu	F	A	⋯	N
Shah	M	B	⋯	Y

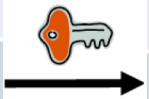

Name	Sex	Blood	⋯	HIV?
100101	001001	110101	⋯	110111
101010	111010	111111	⋯	001001
001010	100100	011001	⋯	110101
001110	010010	110101	⋯	100001
110101	000000	111001	⋯	010010
111110	110010	000101	⋯	110101

Challenges: How to search over data or compute statistics? How efficient/general is this?

Ongoing work on Fully Homomorphic Encryption and Secure Multiparty Computation

Approach 3: Use Secure Hardware

- Secure processors, such as Intel SGX, provide attested execution inside enclaves
  - Encrypted data from user is decrypted inside the enclave and processed

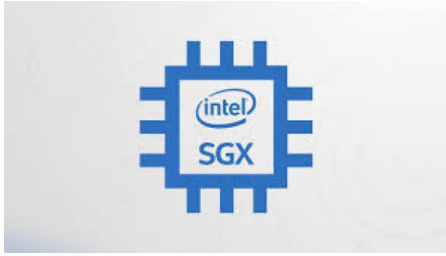

Challenges: Side Channel attacks! Spectre, Meltdown, Foreshadow, …

Ongoing work on RISC-V based secure enclaves

Approach 4: Mediate Access

Name	Sex	Blood	⋯	HIV?
Chen	F	B	⋯	Y
Jones	M	A	⋯	N
Smith	M	O	⋯	N
Ross	M	O	⋯	Y
Lu	F	A	⋯	N
Shah	M	B	⋯	Y

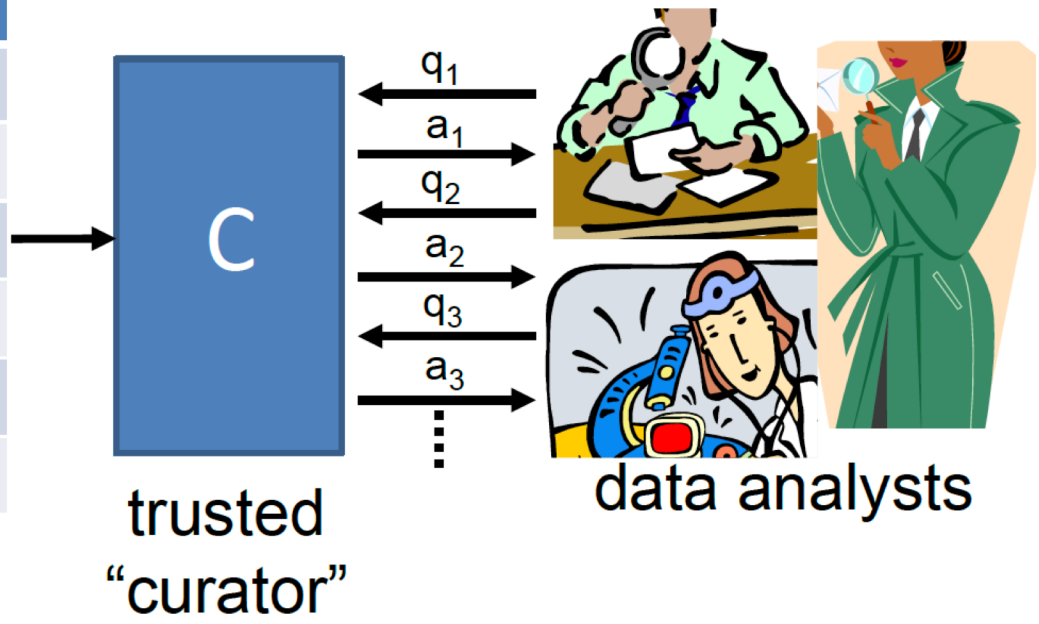

Problems: Curator sees all the data. What queries are allowed? How much do they leak?

Differential privacy

Sex	Blood	…	HIV?
F	B	…	Y
M	A	…	N
M	O	…	N
M	O	…	Y
F	A	…	N
M	B	…	Y

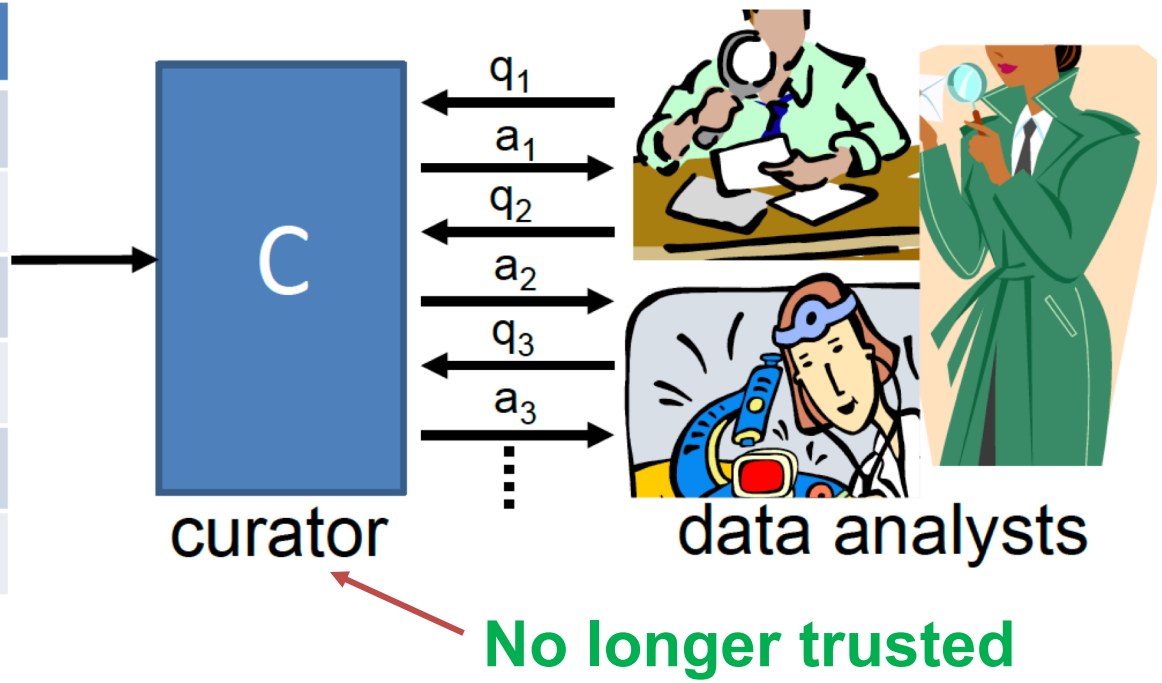

- Requirement: effect of each individual should be "hidden"

Differential privacy

Sex	Blood	…	HIV?
F	B	…	Y
M	A	…	N
M	O	…	N
M	O	…	Y
F	A	…	N
M	B	…	Y

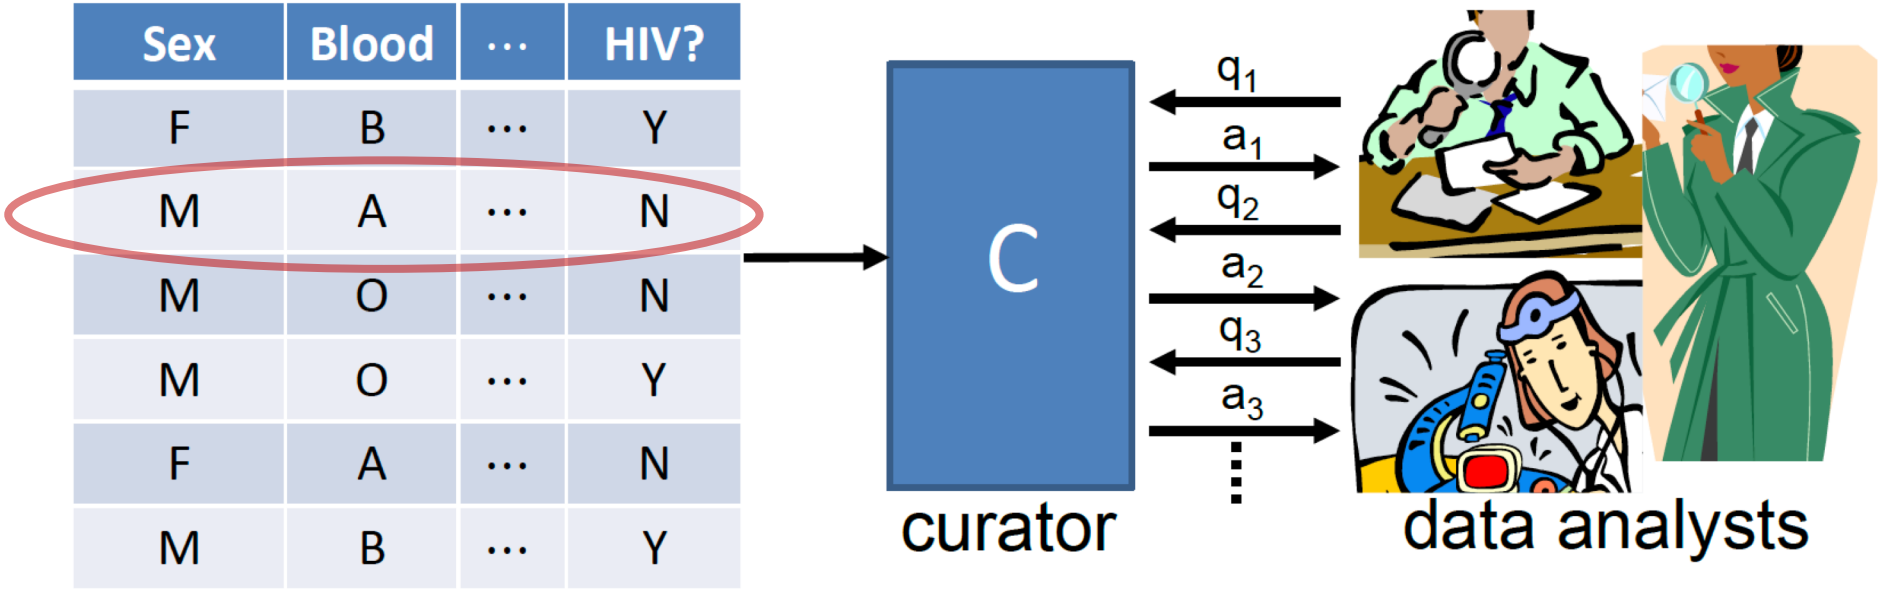

- Requirement: Adversary should not be able to tell if any one person's data were changed arbitrarily

Differential privacy

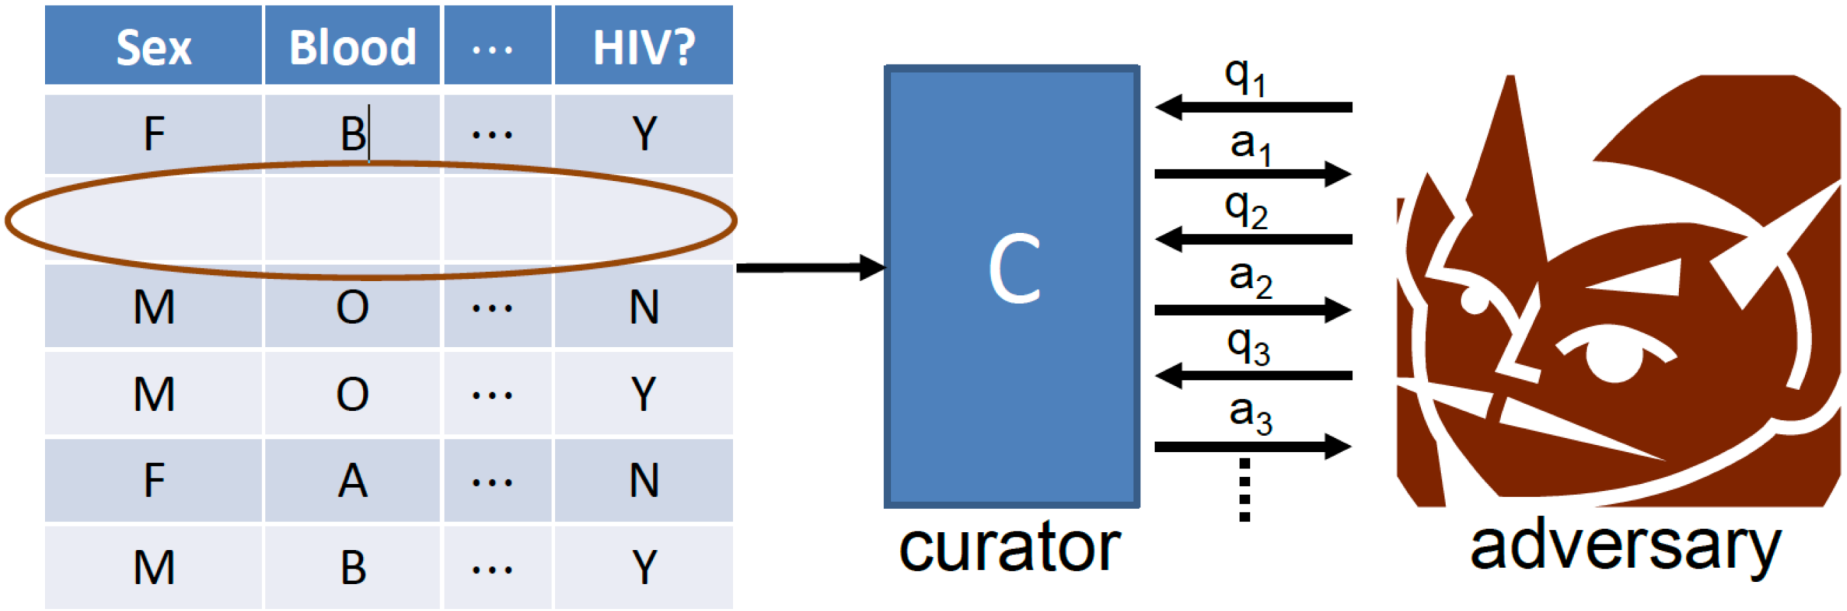

- Requirement: Adversary should not be able to tell if any one person's data were changed arbitrarily

Differential privacy

Sex	Blood	…	HIV?
F	B	…	Y
F	A		Y
M	O	…	N
M	O	…	Y
F	A	…	N
M	B	…	Y

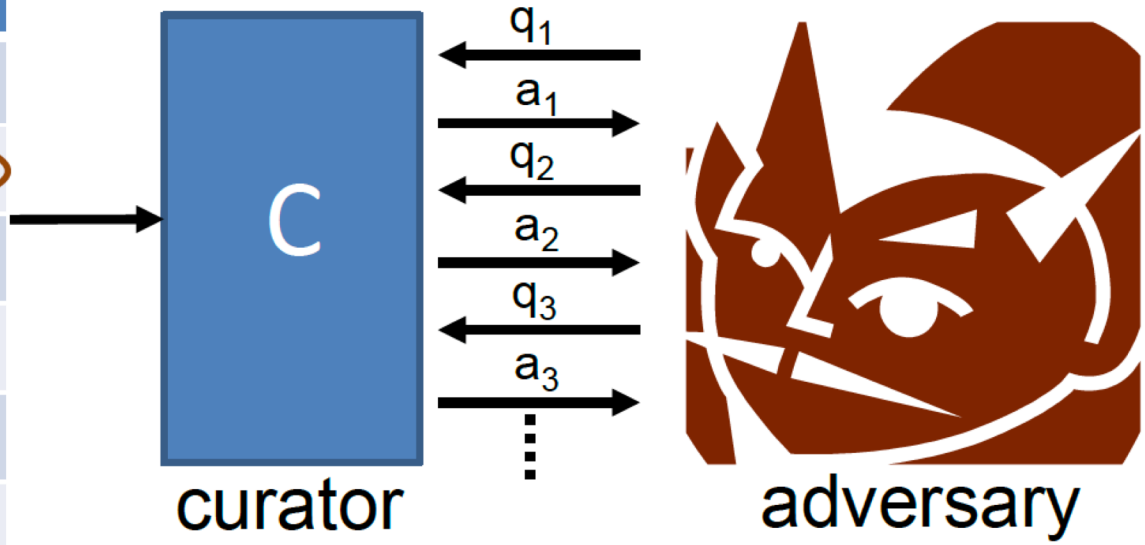

- Requirement: Adversary should not be able to tell if any one person's data were changed arbitrarily

Simple approach: random noise

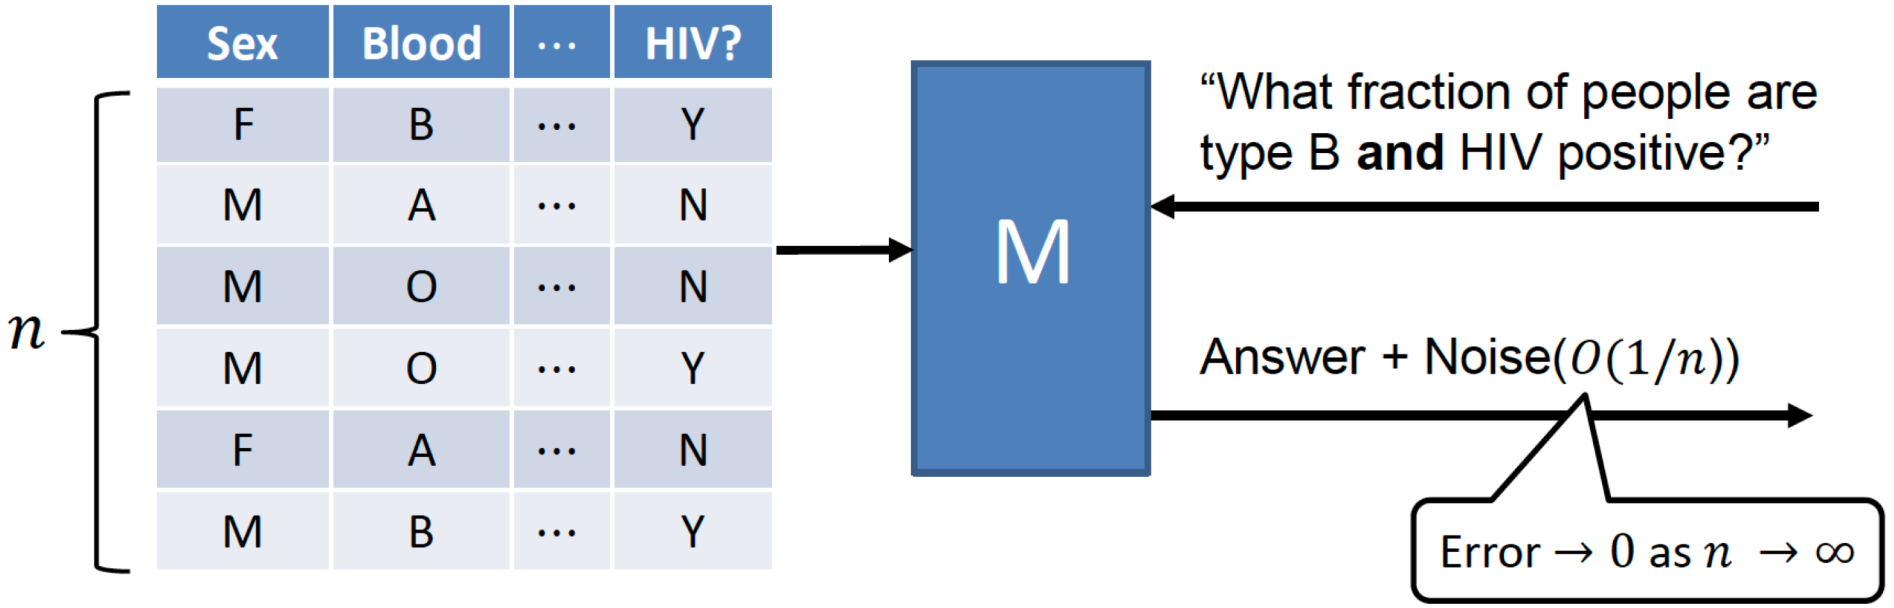

- Very little noise needed to hide each person as $n \to \infty$
- This is just for one query

Differential Privacy Deployed

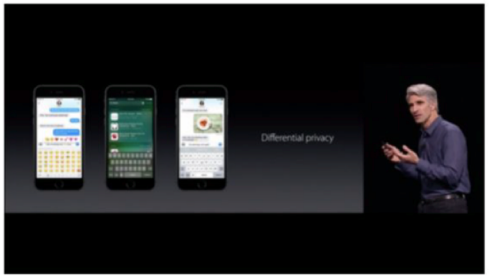

Apple

Google Security Blog

The latest news and insights from Google on security and safety on the Internet

Learning statistics with privacy, aided by the flip of a coin

October 30, 2014

Cross-posted on the Research Blog and the Chromium Blog

At Google, we are constantly trying to improve the techniques we use to protect our users' security and privacy. One such project, RAPPOR (Randomized Aggregatable Privacy-Preserving Ordinal Response), provides a new state-of-the-art, privacy-

Google

United States™ Census Bureau

Census Bureau

Uber

mostly focused on count and average statistics

Challenges for DP in Practice

- Accuracy for “small data” (small n)
- Modeling and managing privacy loss over time
- Analysts are used to working with raw data, not querying (slightly) noisy data
- Matching guarantees with privacy law and regulation

- Significant work at CSAIL addressing these challenges

A New Data Transformation Approach

- Dataset X belongs to Owner
- Owner transforms X using a private transform T to T(X) and shares with untrusted entity U
- U learns model M based on T(X) and returns M to Owner
- To predict using M, owner queries $M(T(x_{new}))$
- Key questions:
  - How much does T(X) leak about X?
  - How accurate is M relative to model based on X?

Data Augmentation

- Mixup is a popular data augmentation strategy
- Given two samples X and Y, choose random r, $0 < r < 1$ and generate a new sample

$$r\,X + (1 - r)\,Y$$

+ Can increase the number of training samples
+ Data augmentation can hide private Y with secret r, even if one sample X is public

Questions and Observations

- How to choose the transform T so we can show a privacy property?
- What is the resultant utility of learning/computing on the transformed data?

+ One-time transformation of data is computationally efficient
+ Positive preliminary results for Support Vector Machines, fully connected neural networks for transform with bounded leakage

Summary

- Being able to legally and efficiently compute on private data is essential
- This is a rich area of research with new technologies constantly being developed
- We look forward to working with you to address research challenges!