Audit Mechanisms for Privacy Protection in Healthcare Environments

Anupam Datta

Joint work with
Jeremiah Blocki, Nicolas Christin and Arunesh Sinha
Carnegie Mellon University

Position

- Audit mechanisms are essential for privacy protection in healthcare environments
  - Guided by comprehensive study of HIPAA Privacy Rule (WPES'10, CCS'11)

- Principled audit mechanisms based on machine learning and economics can be used to provide operational guidance to organizations on how to conduct audits
  - For "grey" policy concepts: was access for purpose of treatment or curiosity, financial gain etc.?

Learning to Audit

Auditor

Auditing budget: $3000/ cycle
Cost for one inspection: $100
Only 30 inspections per cycle

Access divided into 2 types

Loss from 1 violation (internal, external)

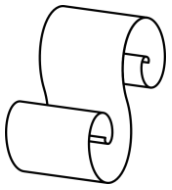

100 accesses

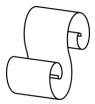

30 accesses

Sandra Bullock

$500, $1000

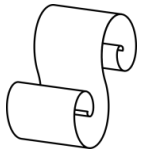

70 accesses

$250, $500

Audit Mechanism Choices

Only 30 inspections

Consider 4 possible allocations of the available 30 inspections

	0	10	20	30
	30	20	10	0
Weights	1.0	1.0	1.0	1.0

Choose allocation probabilistically based on weights

Audit Mechanism Run

No. of Access	Actual Violation
30	2
70	4

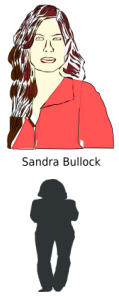

0	10	20	30
30	20	10	0

Observed Loss

Estimated Loss

Int. Caught	Ext. Caught
1	1
2	1

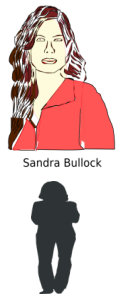

$2000	$1500	$1000	$1000
$750	$1250	$1250	$1500

Updated weights

0.5	0.5	2.0	1.5

Learning from experience: weights updated using observed and estimated loss

Regret Minimizing Audits

- Learns from experience to recommend budget allocation for audit in each audit cycle
- Observed loss used to estimate loss for each action and update probabilities for actions
- Budget allocation is *provably close* to *optimal fixed strategy* in hindsight (e.g., budget allocation)

- Technical approach: New regret minimization algorithm for repeated games of imperfect information

  (Online learning-theoretic technique)

J. Blocki, N. Christin, A. Datta, A. Sinha, Regret Minimizing Audits: A Learning-Theoretic Basis for Privacy Protection, *CSF*, June 2011.

Future Work

- Alternative adversary models
  - Worst-case, rational, well-behaved
- Alternative audit mechanisms
  - Incorporating incentives
- Identifying experts
  - Can experts be learned from logs?
- Experimental evaluation
  - Real hospital logs, user studies