Private Release and Learning of Thresholds

Mark Bun (3rd year Ph.D., supported by NDSEG Fellowship)

Joint work with Kobbi Nissim, Uri Stemmer, and Salil Vadhan

Privacy Tools for Sharing Research Data

A National Science Foundation Secure and Trustworthy Cyberspace Project

MOTIVATING QUESTION

How many **data samples** do we need to achieve both **differential privacy** and **statistical accuracy**?

i.e. How big a study do we need to conduct to answer our questions and preserve privacy?

PRIVATE QUERY RELEASE

Counting queries: What fraction of rows in a database satisfy property q?

e.g. $q(x) = \text{Age}(x) \geq 42$?

DarkSide?	Age	Home	Weight	
0	896	Dagobah	17	$q(x_1)=1$
0	19	Alderaan	49	$q(x_2)=0$
0	19	Tatooine	77	$q(x_3)=0$
1	42	Tatooine	136	$q(x_4)=1$

$q(D)=1/2$

Goal: Privately answer *all* $q \in Q$ to within 0.05 error

PRIVATE (PROPER) LEARNING

Examples drawn from a distribution and labeled by an unknown predicate $q \in Q$

e.g. $q(x) = \text{Age}(x) \geq 42$?

Age	$q(x)$
896	1
19	0
19	0
42	1

Goal: Output $q' \in Q$ that classifies new examples with 95% accuracy

THRESHOLD FUNCTIONS

$$THRESH^y(x) = \begin{cases} 1 & \text{if } x \geq y \\ 0 & \text{otherwise} \end{cases}$$

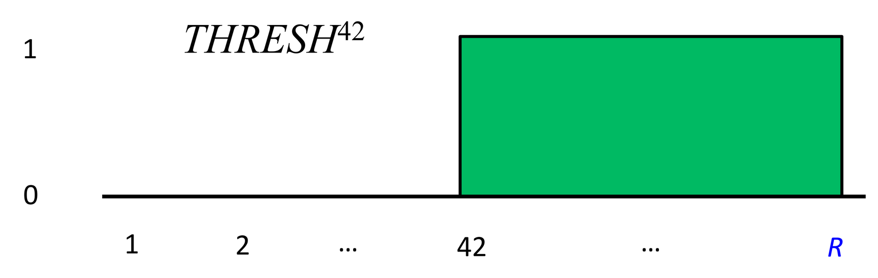

For any (even infinite) domain, VC($THRESH$) = 1
=> $THRESH$ can be learned (non-privately) with O(1) samples

QUESTION FOR THIS WORK

Can we *privately* release/learn thresholds over infinite domains?

If not, does the sample complexity depend on R?

Answer: Thresholds require $n > \log^* R$

PRIOR WORK & OUR RESULTS

	Q arbitrary	**Q** = $POINT_R$	**Q** = $THRESH_R$
No privacy	$n = \Theta(\mathrm{VC}(Q))$ [e.g. Vap98]	$n = \Theta(1)$ [Vap98]	$n = \Theta(1)$ [Vap98]
(ε,δ)-diff. priv. Query release	$\Theta(\log\lvert Q\rvert \cdot (\log R)^{1/2})$ [HR10, BUV14]	$\Theta(1)$ [BNS13]	$n < 8^{\log^* R}$ [BNS13] **OUR WORK:** $\log^* R < n < 2^{\log^* R}$
(ε,δ)-diff. priv. Proper learning	$O(\log\lvert Q\rvert)$ [KLNRS08] …but no general lower bounds	$\Theta(1)$ [BNS13]	$n < 8^{\log^* R}$ [BNS13] **OUR WORK:** $\log^* R < n < 2^{\log^* R}$

DIFFERENTIAL PRIVACY

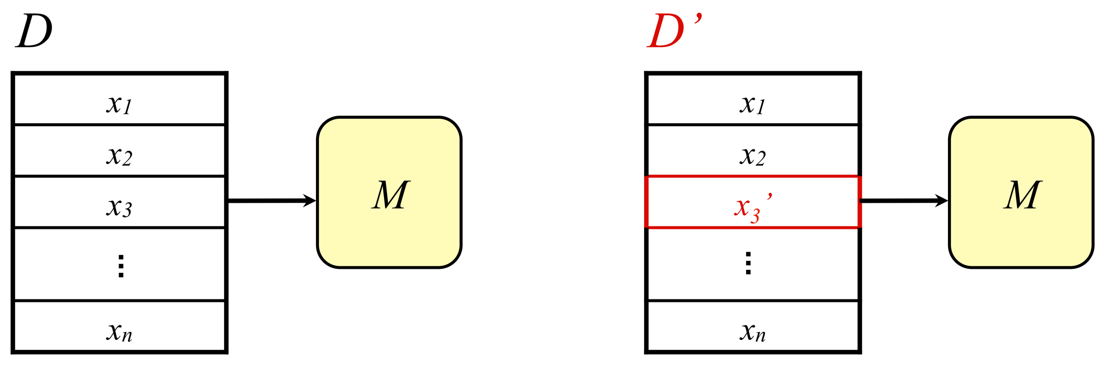

D and D' are neighbors if they differ only on one user's data

An algorithm M is (ε,δ)-differentially private if for all neighbors D, D' and every $S \subseteq \mathrm{Range}(M)$,

$$\Pr[M(D) \in S] \leq e^{\varepsilon} \Pr[M(D') \in S] + \delta$$

Think of $\varepsilon = \Theta(1)$ and $\delta = o(1/n)$

OUR TECHNIQUES

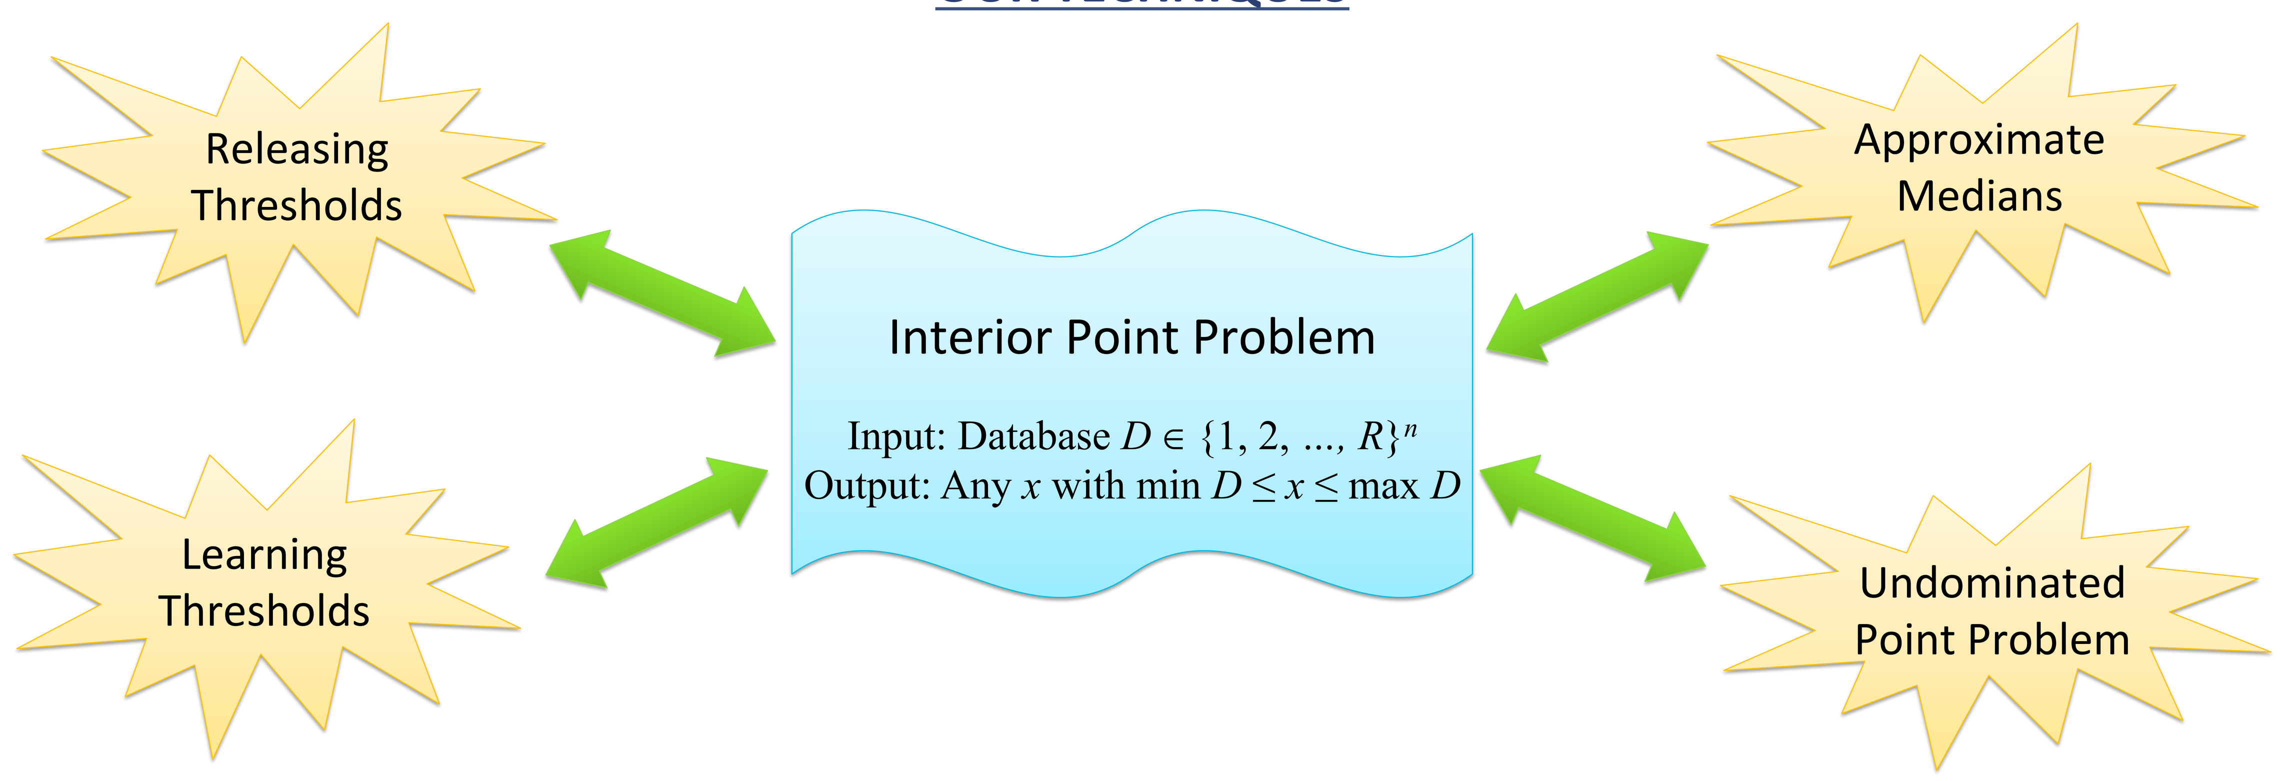

INTERIOR POINT LOWER BOUND

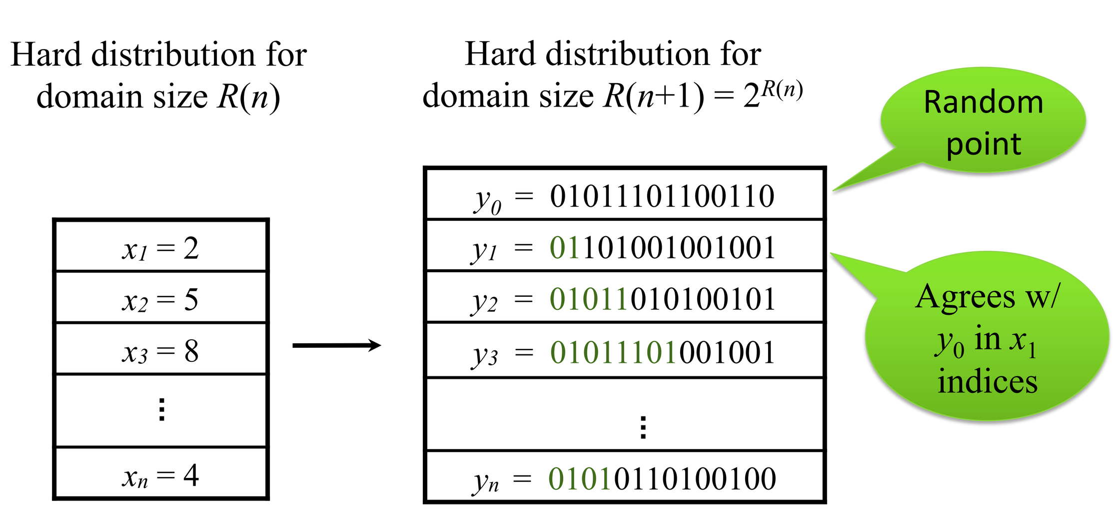

INTERIOR POINT UPPER BOUND

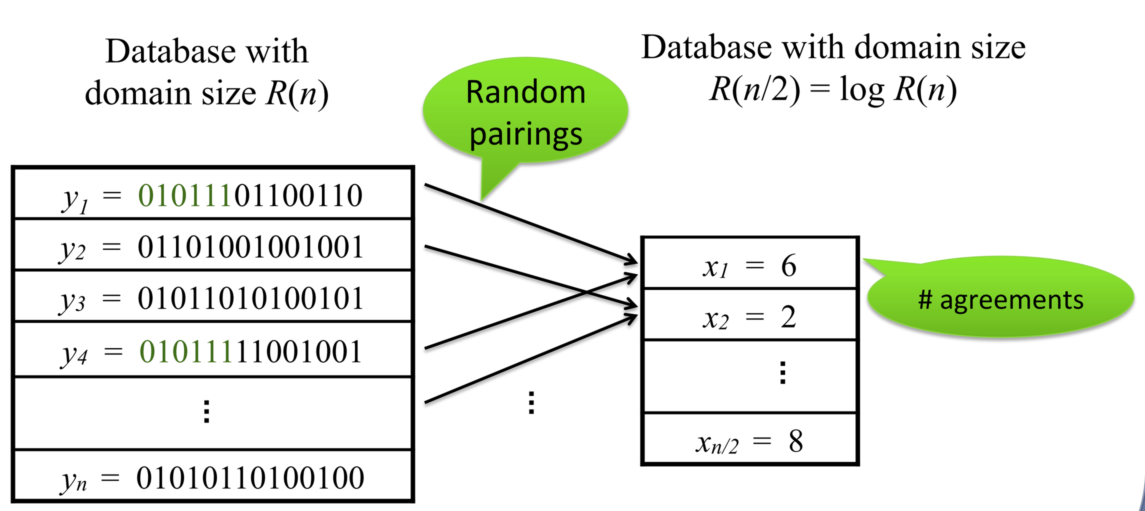

CONCLUSIONS

- Releasing/learning thresholds requires sample complexity growing with R
- Separates private release/learning from non-private cases, even for VC(Q) = 1
- Open questions: Can we characterize the difference between private/non-private sample complexity? Extend results to improper learning?

REFERENCES

[BUV14] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *STOC*, 2014.

[BNS13] Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: pure vs. approximate differential privacy. In *RANDOM*, 2013.

[BNSV14] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Differentially private release and learning of thresholds. *Manuscript*, 2014.

[HR10] Moritz Hardt and Guy Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS*, 2010.

[KLNRS08] Shiva Kasiviswanathan, Homin Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? In *FOCS* 2008.

[Vap98] Vladimir Vapnik. Statistical Learning Theory. 1998.

CONTACT

mbun@seas.harvard.edu

Harvard School of Engineering and Applied Sciences
Maxwell Dworkin 138
33 Oxford St.
Cambridge, MA 02138

CRCS Center for Research on Computation and Society at Harvard School of Engineering and Applied Sciences

BERKMAN CENTER FOR INTERNET & SOCIETY AT HARVARD UNIVERSITY

The Institute for Quantitative Social Science at Harvard University

DATA PRIVACY LAB