

From Learning Algorithms to Differentially Private Algorithms

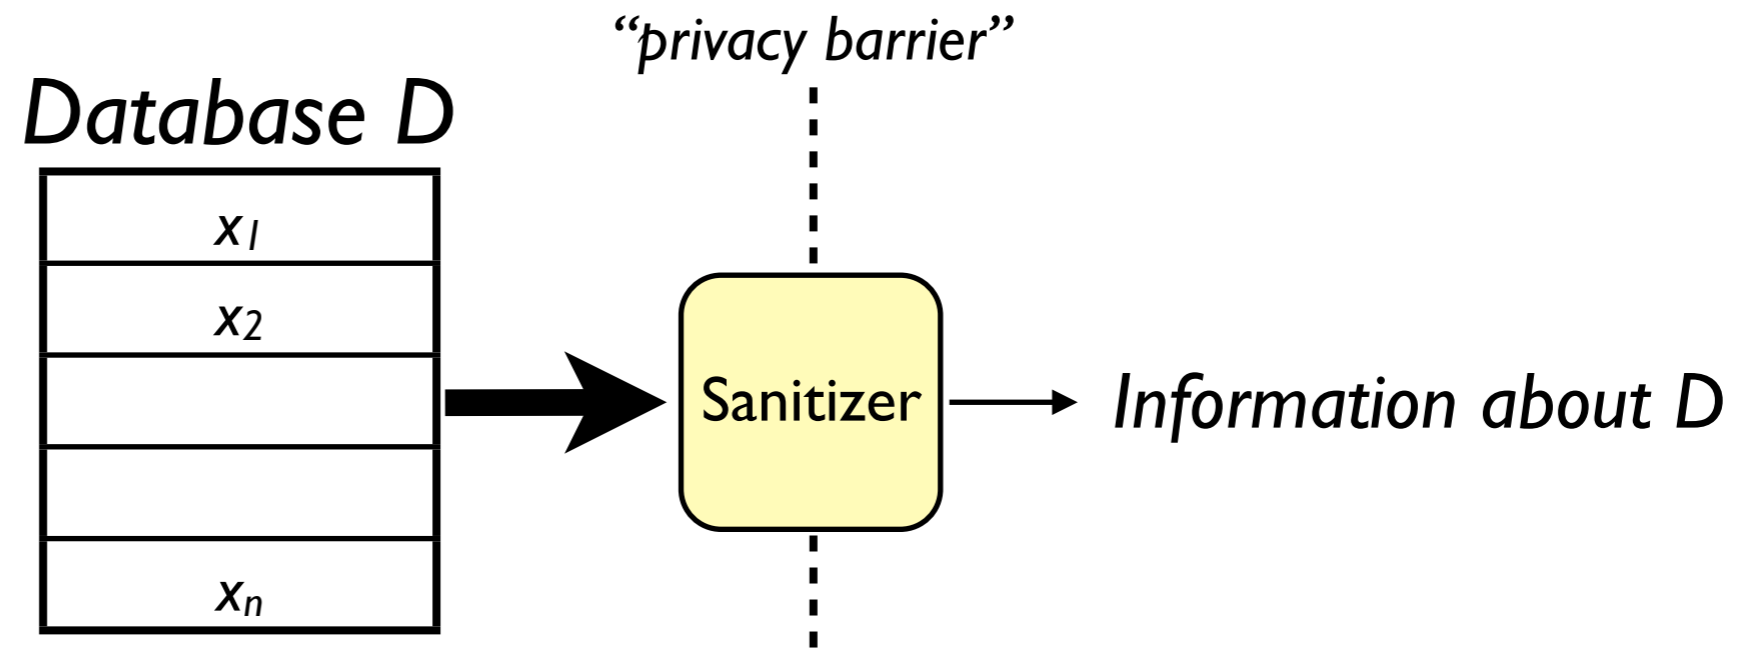
Jonathan Ullman, Harvard University

Big Data and Differential Privacy Workshop
December 12, 2013

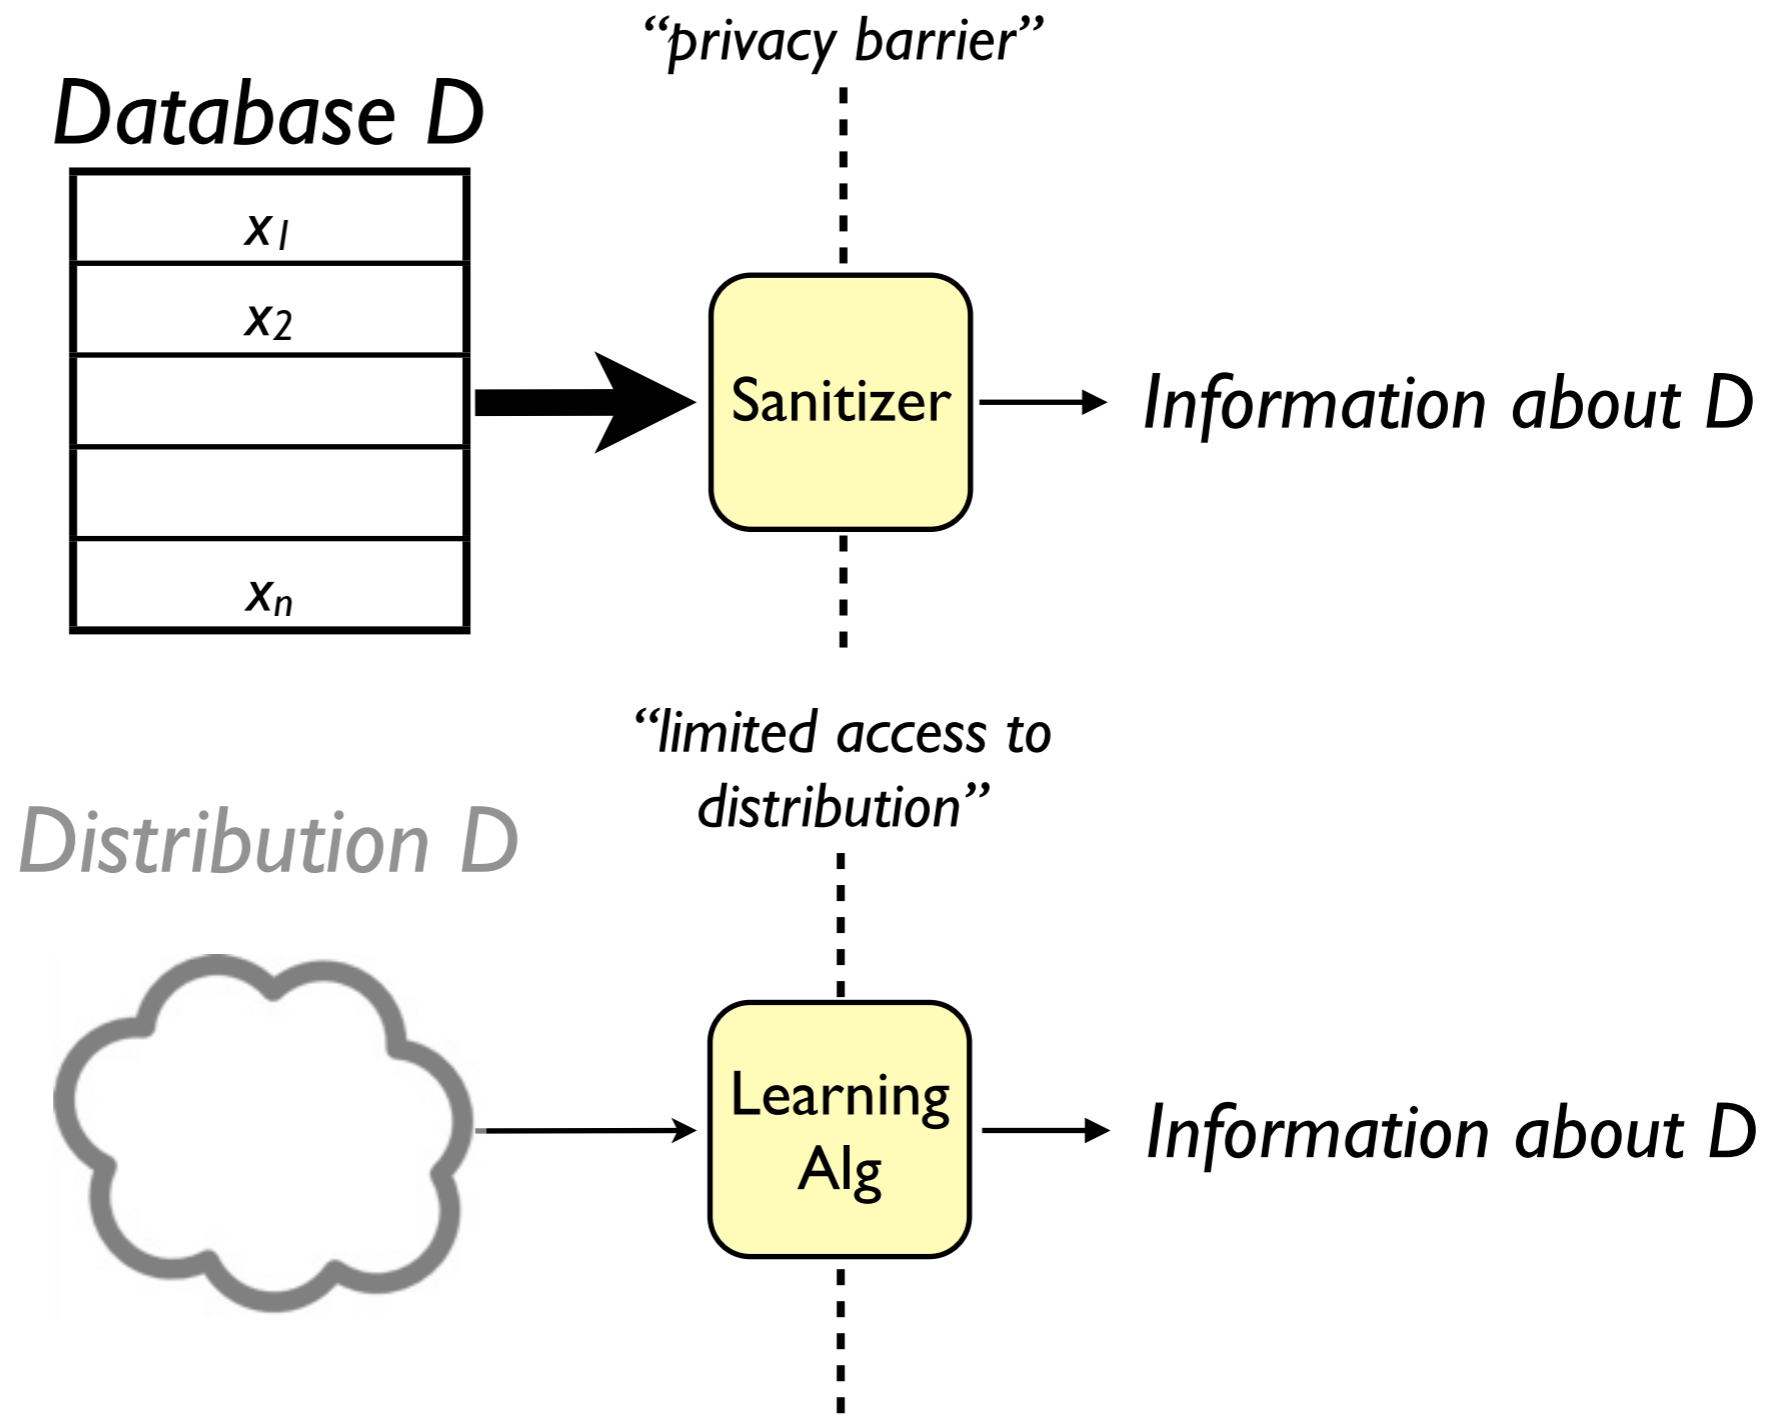
What is this tutorial about?

- Using powerful techniques from learning theory to design differentially private algorithms

Why would we want to do that?



Why would we want to do that?



Why would we want to do that?

- Connections between learning and DP algorithm design first(?) introduced in [BDMN, KLNRS]

Why would we want to do that?

- Clean, qualitatively strong guarantees brought out the potential of differentially private data analysis
- For these strong guarantees, learning-theoretic techniques yield nearly-optimal algorithms

Private Counting Query Release

$$D \in (\{0, 1\}^d)^n$$

Counting query: *What fraction of records satisfy property q ?*

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

d attributes per record

Private Counting Query Release

$$D \in (\{0, 1\}^d)^n$$

Counting query: What fraction of records satisfy property q ? e.g.

$q(x) = \text{GiveYouUp} \vee \text{LetYouDown?}$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

$$q(x_1) = 0$$

$$q(x_2) = 1$$

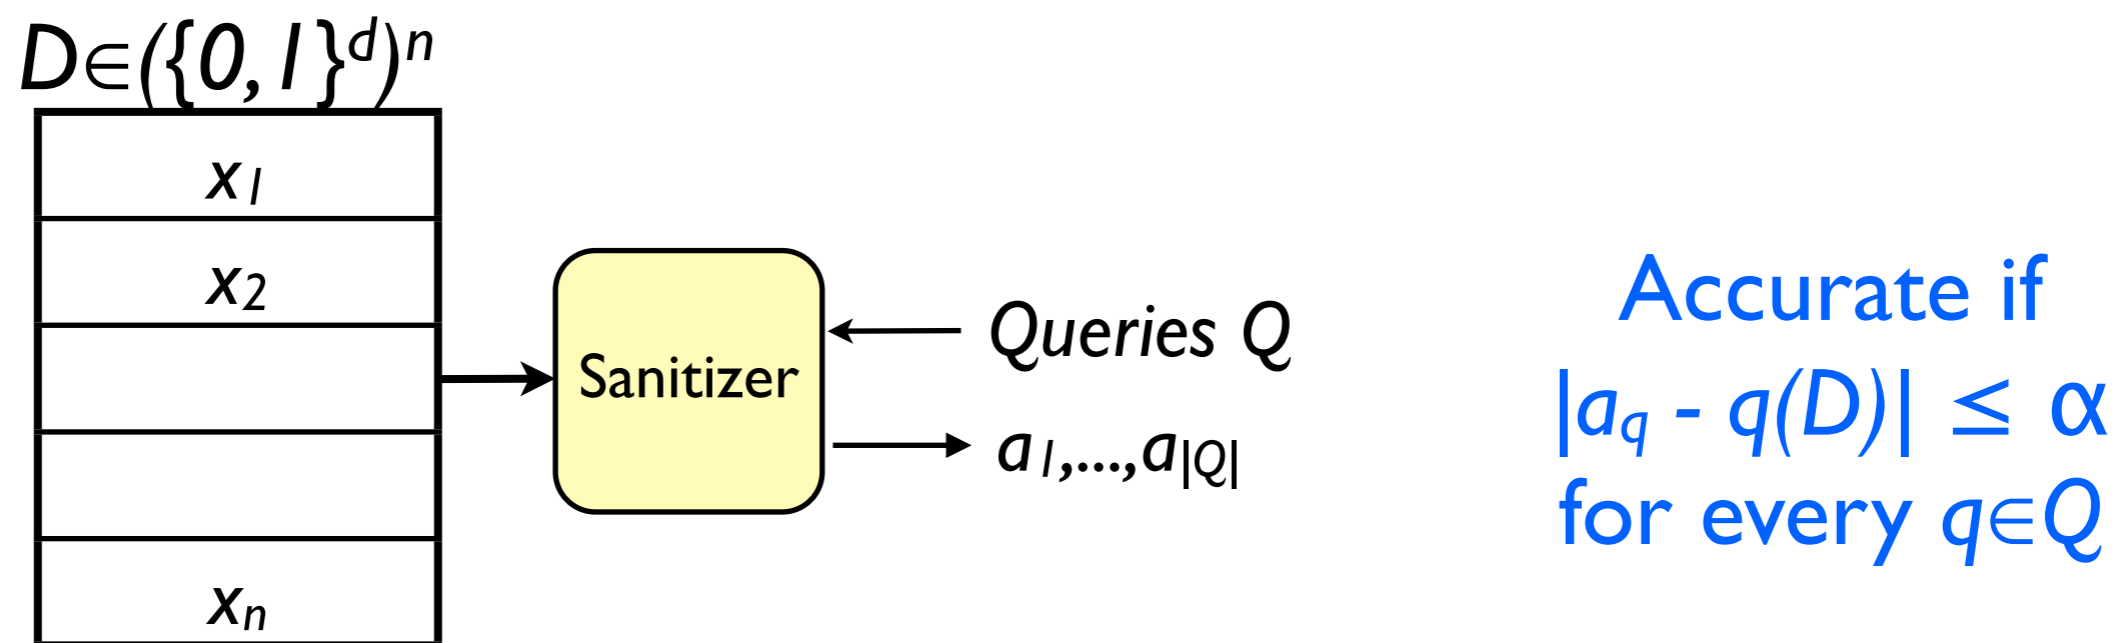
$$q(x_3) = 1$$

$$q(x_4) = 1$$

d attributes per record

$$q(D) = 3/4$$

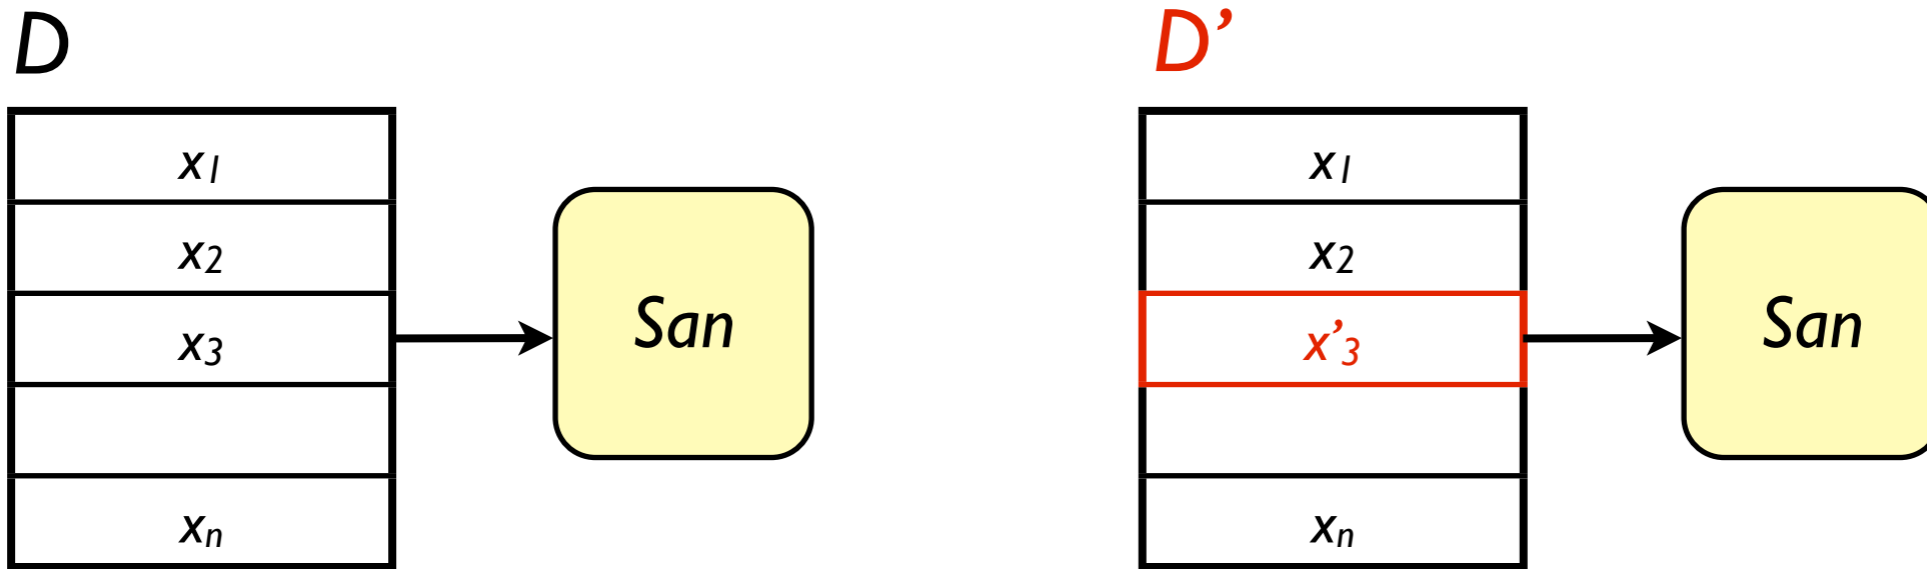
Private Counting Query Release



- Want to design a sanitizer that is simultaneously differentially private and accurate

Differential Privacy

[DN, DN, BDMN, DMNS, D]

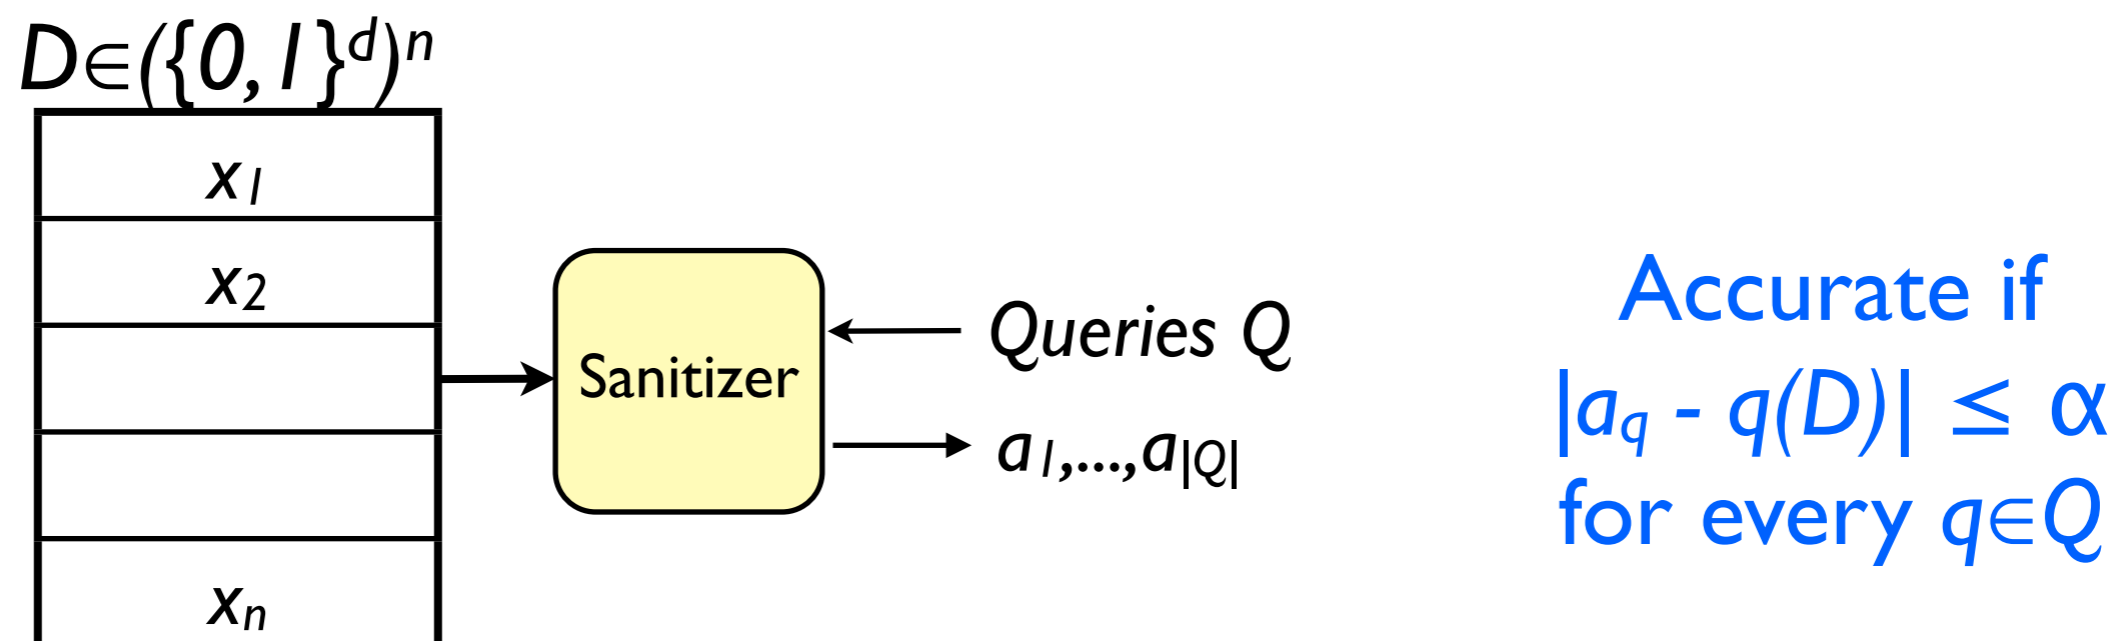


D and D' are neighbors if they differ only on one user's data

Definition: A (randomized) San is (ϵ, δ) -differentially private if for all neighbors D, D' and every $S \subseteq \text{Range}(San)$

$$\Pr[San(D) \in S] \leq e^\epsilon \Pr[San(D') \in S] + \delta$$

Private Counting Query Release



- Want to design a sanitizer that is simultaneously differentially private and accurate
- Want to minimize
 - Amount of data required, n for a given Q, d, α
 - Running time of the sanitizer

Private Query Release: An Abridged History

- Adding independent noise (Laplace mechanism)
requires $n \gtrsim |Q|^{1/2}/\alpha$

Private Query Release: An Abridged History

- Adding independent noise (Laplace mechanism) requires $n \gtrsim |Q|^{1/2}/\alpha$
- [BLR] gave a sanitizer that requires only $n \gtrsim d \log|Q|/\alpha^3$
- Several important improvements by [DNRRV,DRV,RR]

Private Query Release: An Abridged History

- [HR] introduced the private multiplicative weights algorithm, requires only $n \gtrsim d^{1/2} \log|Q|/\alpha^2$

Private Query Release: An Abridged History

- [HR] introduced the private multiplicative weights algorithm, requires only $n \gtrsim d^{1/2} \log|Q|/\alpha^2$
- Put in a general framework, with tight analysis by [GHRU,GRU,HLM]
- Several improvements for special cases of private query release followed [GRU,JT,BR,HR,HRS,TUV,CTUW,...]

Talk Outline

- Differentially private query release
- A blueprint for private query release
 - No-regret algorithms / MW
- Query Release Algorithms
 - Offline MW
 - Online MW
 - Variants
 - Faster algorithms for disjunctions via polynomial approx.

A Blueprint for Query Release

Sanitized (DP) Output

Raw Data


D



A Blueprint for Query Release

Sanitized (DP) Output

D_I



Raw Data


D



A Blueprint for Query Release

Sanitized (DP) Output

D_I



Raw Data

D

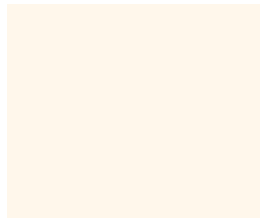


Is D_I
good for Q ?



A Blueprint for Query Release

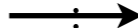

Sanitized (DP) Output

D_I  *(hint_I)*

Raw Data

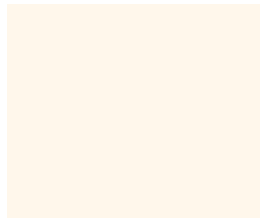
D



Is D_I
good for Q ? 
 Here's *hint_I*

A Blueprint for Query Release

Sanitized (DP) Output



D_I  *(query_I)*

Raw Data

D



Is D_I
good for Q ?

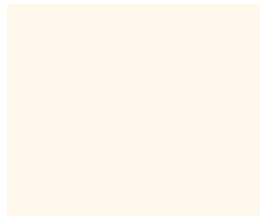



Here's *query_I*

query_I is a query D_I
answers incorrectly

A Blueprint for Query Release

Sanitized (DP) Output

D_1  $(query_1)$



Update Alg: U

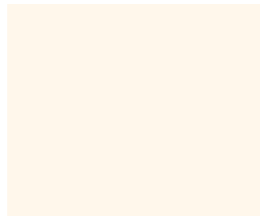
Raw Data


D




A Blueprint for Query Release

Sanitized (DP) Output

D_1  *(query₁)*

D_2 



Update Alg: U

Raw Data

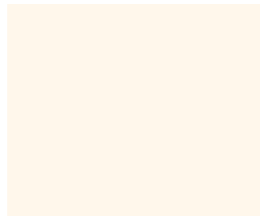
D




A Blueprint for Query Release

Sanitized (DP) Output

Raw Data


D_1  *(query₁)*

D_2  *(query₂)*



Update Alg: U

Is D_2
good for Q ?

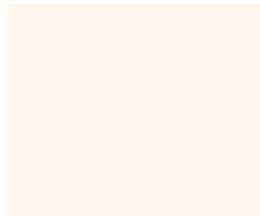


Here's *query₂*

query₂ is a query D_2
answers incorrectly


A Blueprint for Query Release

Sanitized (DP) Output

D_1  *(query₁)*

D_2  *(query₂)*

D_3 



Update Alg: U

Raw Data



D



A Blueprint for Query Release


Sanitized (DP) Output

Raw Data



Update Alg: U

Is D_T
good for Q ?






Yes, approximately!

A Blueprint for Query Release


Sanitized (DP) Output

Raw Data



Update Alg: U

Is D_T
good for Q ?



Yes, approximately!

A Blueprint for Query Release

LET D be the real database

LET D_1 be an “initial guess”

FOR $t = 1, \dots, T$

LET $query_t = \operatorname{argmax}_{q \in Q} q(D_t) - q(D)$

LET $D_{t+1} = \operatorname{Update}(D_t, q_t)$



Why did we do this?

- Decomposed the problem into smaller problems
 - Fortunately, DP has nice composition properties
- We've separated privacy (finding q_t) from the task of learning the database (updating D_t)
 - Means we can choose any update algorithm

A Blueprint for Query Release


Sanitized (DP) Output

Raw Data



Update Alg: U

Is $D_{|Q|}$
good for Q ?



Yes, approximately!

Why did we do this?

- (Hopefully) decomposed the problem into $T \ll |Q|$ smaller problems
 - Fortunately, DP has nice composition properties
- We've separated privacy (finding q_t) from the task of learning the database (updating D_t)
 - Means we can choose any update algorithm

Talk Outline

- Differentially private query release
- A blueprint for private query release
 - [No-regret algorithms / MW](#)
- Query Release Algorithms
 - Offline MW
 - Online MW
 - Variants
 - Faster algorithms for disjunctions via polynomial approx.

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



No-Regret Learning Algorithms

Set of experts X



Distribution over X

D_1 .25 .25 .25 .25

Losses for each expert



No-Regret Learning Algorithms

Set of experts X



Distribution over X

D_t .25 .25 .25 .25

Losses for each expert



$[0, 1]^X$

Loss is $\langle D_t, L_t \rangle$

1 0 1 0 L_t

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



Distribution over X

D_1	.25	.25	.25	.25
D_2	.20	.30	.20	.30

Loss is $\langle D_1, L_1 \rangle$

$[0, 1]^X$

1 0 1 0 L_1

Multiplicative Weights Update [LW]

$D_2 = MWU(D_1, L_1)$:

$$D'_2(x) = (1 - \eta L_1(x)) D_1(x)$$

$$D_2(x) = D'_2(x) / \sum_{x \in X} D'_2(x)$$

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



Distribution over X

D_1	.25	.25	.25	.25
D_2	.20	.30	.20	.30

Loss is $\langle D_1, L_1 \rangle$
 Loss is $\langle D_2, L_2 \rangle$

$[0, 1]^X$

1	0	1	0	L_1
0	0	1	0	L_2

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



Distribution over X

D_1	.25	.25	.25	.25
D_2	.20	.30	.20	.30
D_T	.23	.32	.15	.32

Loss is $\langle D_1, L_1 \rangle$

Loss is $\langle D_2, L_2 \rangle$

⋮

Loss is $\langle D_T, L_T \rangle$

$[0, 1]^X$

1	0	1	0	L_1
0	0	1	0	L_2
0	0	0	1	L_T

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



Distribution over X

					$[0, 1]^X$					
D_1	.25	.25	.25	.25	Loss is $\langle D_1, L_1 \rangle$	1	0	1	0	L_1
D_2	.20	.30	.20	.30	Loss is $\langle D_2, L_2 \rangle$	0	0	1	0	L_2
					\vdots					
D_T	.23	.32	.15	.32	Loss is $\langle D_T, L_T \rangle$	0	0	0	1	L_T

For any distribution D , sequence L_1, \dots, L_T ,

$$\sum_{t=1}^T \langle D_t - D, L_t \rangle \leq \sqrt{T \log |X|}$$

Counting Queries

Counting query: What fraction of records satisfy property q ? e.g.

$q(x) = \text{GiveYouUp} \vee \text{LetYouDown}$

$$D \in (\{0, 1\}^d)^n$$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

$$q(x_1) = 0$$

$$q(x_2) = 1$$

$$q(x_3) = 1$$

$$q(x_4) = 1$$

$$q(D) = 3/4$$

d attributes per record

Counting Queries

Counting query: What fraction of records satisfy property q ? e.g.

$q(x) = \text{GiveYouUp} \vee \text{LetYouDown}$

$$D \in (\{0, 1\}^d)^n$$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

$$q(x_1) = 0$$

$$q(x_2) = 1$$

$$q(x_3) = 1$$

$$q(x_4) = 1$$

$$q(D) = 3/4$$

d attributes per record



D is a distribution on $\{0, 1\}^d$

Counting Queries

Counting query: What fraction of records satisfy property q ? e.g.

$q(x) = \text{GiveYouUp} \vee \text{LetYouDown}$

$$D \in (\{0, 1\}^d)^n$$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

$$q(x_1) = 0$$


$$q(x_2) = 1$$

$$q(x_3) = 1$$


$$q(x_4) = 1$$

$$q(D) = 3/4$$

d attributes per record



q is an indicator vector



D is a distribution on $\{0, 1\}^d$

Linear query: $q(D) = \langle D, q \rangle$

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



Distribution over X

					$[0, 1]^X$					
D_1	.25	.25	.25	.25	Loss is $\langle D_1, L_1 \rangle$	1	0	1	0	L_1
D_2	.20	.30	.20	.30	Loss is $\langle D_2, L_2 \rangle$	0	0	1	0	L_2
					\vdots					
D_T	.23	.32	.15	.32	Loss is $\langle D_T, L_T \rangle$	0	0	0	1	L_T

For any distribution D , sequence L_1, \dots, L_T ,

$$\sum_{t=1}^T \langle D_t - D, L_t \rangle \leq \sqrt{T \log |X|}$$

Multiplicative Weights for Query Release

Set of experts $X = \{0, 1\}^d$



Losses for each expert



Distribution over $X = \{0, 1\}^d$

Truth table of q in $[0, 1]^X$

D_1	.25	.25	.25	.25	Loss is $\langle D_1, q_1 \rangle$	1	0	1	0	q_1
D_2	.20	.30	.20	.30	Loss is $\langle D_2, q_2 \rangle$	0	0	1	0	q_2
					⋮					
D_T	.23	.32	.15	.32	Loss is $\langle D_T, q_T \rangle$	0	0	0	1	q_T

For any database D , sequence q_1, \dots, q_T ,

$$\sum_{t=1}^T \langle D_t - D, q_t \rangle \leq \sqrt{Td}$$

A Blueprint for Query Release

LET D be the real database, viewed as a dist over $\{0, 1\}^d$

LET D_I be the uniform dist on $\{0, 1\}^d$

FOR $t = 1, \dots, T$

LET $q_t = \operatorname{argmax}_{q \in Q} \langle D_t - D, q \rangle$

LET $D_{t+1} = \text{MWU}(D_t, q_t)$

$$D'_{t+1}(x) = (1 - \eta q_t(x)) D_t(x)$$

$$D_{t+1}(x) = \frac{D'_{t+1}(x)}{\sum_{x \in \{0, 1\}^d} D'_{t+1}(x)}$$

Query Release via MW

- Thm: For any database D sequence q_1, \dots, q_T ,

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle$$

Query Release via MW

- Thm: For any database D sequence q_1, \dots, q_T ,

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle$$

- If q_1, \dots, q_T all satisfy $\langle D_t - D, q_t \rangle \geq \alpha$, then we have

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle \geq \alpha T$$

Query Release via MW

- Thm: For any database D sequence q_1, \dots, q_T ,

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle$$

- If q_1, \dots, q_T all satisfy $\langle D_t - D, q_t \rangle \geq \alpha$, then we have

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle \geq \alpha T$$

- If $T \geq d/\alpha^2$, then $\langle D_T - D, q \rangle \leq \alpha$ for all of Q

Query Release via MW

- Thm: For any database D sequence q_1, \dots, q_T ,

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle$$

- If q_1, \dots, q_T all satisfy $\langle D_t - D, q_t \rangle \geq \alpha$, then we have

$$\sqrt{Td} \geq \sum_{t=1}^T \langle D_t - D, q_t \rangle \geq \alpha T$$

- If $T \geq d/\alpha^2$, then $|\langle D_T - D, q \rangle| \leq \alpha$ for all of Q

Q is closed under neg.

A Blueprint for Query Release

LET D be the real database, viewed as a dist over $\{0, 1\}^d$

LET D_1 be the uniform dist on $\{0, 1\}^d$

FOR $t = 1, \dots, T = O(d/\alpha^2)$

LET $q_t = \operatorname{argmax}_{q \in Q} \langle D_t - D, q \rangle$

LET $D_{t+1} = \operatorname{MWU}(D_t, q_t)$

A Blueprint for Query Release


LET D be the real database, viewed as

LET D_1 be the uniform dist on $\{0, 1\}^d$

FOR $t = 1, \dots, T = O(d/\alpha^2)$

LET $q_t = \operatorname{argmax}_{q \in Q} \langle D_t - D, q \rangle$

LET $D_{t+1} = \operatorname{MWU}(D_t, q_t)$



Have to make
this DP

Finding the “Bad” Queries

- How do I find $\operatorname{argmax}_{q \in Q} \langle D_t - D, q \rangle$ privately? Use the exponential mechanism!
- Output q w.p. proportional to $\exp(\epsilon_0 n \langle D_t - D, q \rangle)$

If $n \geq \log|Q|/\alpha\epsilon_0$ then whp EM outputs q_t s.t.
 $\langle D_t - D, q_t \rangle \geq \max_{q \in Q} \langle D_t - D, q \rangle - \alpha/2$

A Blueprint for Query Release

LET D be the real database, viewed as a dist on $\{0, 1\}^d$

LET D_1 be the *uniform distribution* on $\{0, 1\}^d$

FOR $t = 1, \dots, T = O(d/\alpha^2)$

LET $q_t = q$ w.p. proportional to $\exp(\epsilon \langle D_t - D, q \rangle)$

LET $D_{t+1} = \text{MWU}(D_t, q_t)$

A Blueprint for Query Release

Need $n \geq \log|Q|/\alpha\epsilon_0$

LET D be the real database, viewed as a dist on $\{0, 1\}^d$

LET D_1 be the *uniform distribution* on $\{0, 1\}^d$

FOR $t = 1, \dots, T = O(d/\alpha^2)$

LET $q_t = q$ w.p. proportional to $\exp(\epsilon_0 n \langle D_t - D, q \rangle)$

LET $D_{t+1} = \text{MWU}(D_t, q_t)$

A Blueprint for Query Release

Need $n \geq \log|Q|/\alpha\epsilon_0$

LET D be the real database, viewed as a dist on $\{0,1\}^d$

LET D_1 be the *uniform distribution* on $\{0,1\}^d$

FOR $t = 1, \dots, T = O(d/\alpha^2)$

LET $q_t = q$ w.p. proportional to $\exp(\epsilon_0 n \langle D_t - D, q \rangle)$

LET $D_{t+1} = \text{MWU}(D_t, q_t)$

Thm [DRV]: If $\epsilon_0 \leq \epsilon/(8T \log(1/\delta))^{1/2} \approx \epsilon/T^{1/2}$, then running T (adaptively chosen) ϵ_0 -DP algorithms satisfies (ϵ, δ) -DP.

A Blueprint for Query Release

Need
 $n \geq d^{1/2} \log |Q| / \alpha^2 \epsilon$

LET D be the real database, viewed as a dist on $\{0, 1\}^d$

LET D_1 be the *uniform distribution* on $\{0, 1\}^d$

FOR $t = 1, \dots, T = O(d/\alpha^2)$

LET $q_t = q$ w.p. proportional to $\exp(\epsilon \alpha n \langle D_t - D, q \rangle / d^{1/2})$

LET $D_{t+1} = \text{MWU}(D_t, q_t)$

Thm [DRV]: If $\epsilon_0 \approx \epsilon/T^{1/2} \approx \epsilon\alpha/d^{1/2}$, then running T (adaptively chosen) ϵ_0 -DP algorithms satisfies (ϵ, δ) -DP.

Recap

Thm: PMW takes a database $D \in (\{0, 1\}^d)^n$ and a set of counting queries Q , satisfies (ϵ, δ) -DP and, if

$$n \geq d^{1/2} \log |Q| / \alpha^2 \epsilon,$$

it outputs D_T such that for every $q \in Q$,

$$|q(D) - q(D_T)| \leq \alpha$$

Optimality?

- PMW achieves a nearly-optimal data requirement for this level of generality
- **Thm [BUV]:** for every sufficiently large s , there is a family of s queries Q such that any (ϵ, δ) -DP algorithm that is α -accurate for Q requires
$$n \gtrsim d^{1/2} \log |Q| / \alpha^2 \epsilon$$

Recap

Thm: PMW takes a database $D \in (\{0, 1\}^d)^n$ and a set of counting queries Q , satisfies (ϵ, δ) -DP and, if

$$n \geq O(d^{1/2} \log |Q| / \alpha^2 \epsilon),$$

it outputs D_T such that for every $q \in Q$,

$$|q(D) - q(D_T)| \leq \alpha$$

Thm: PMW runs in time $\text{poly}(n, 2^d, |q_1| + \dots + |q_{|Q|}|)$

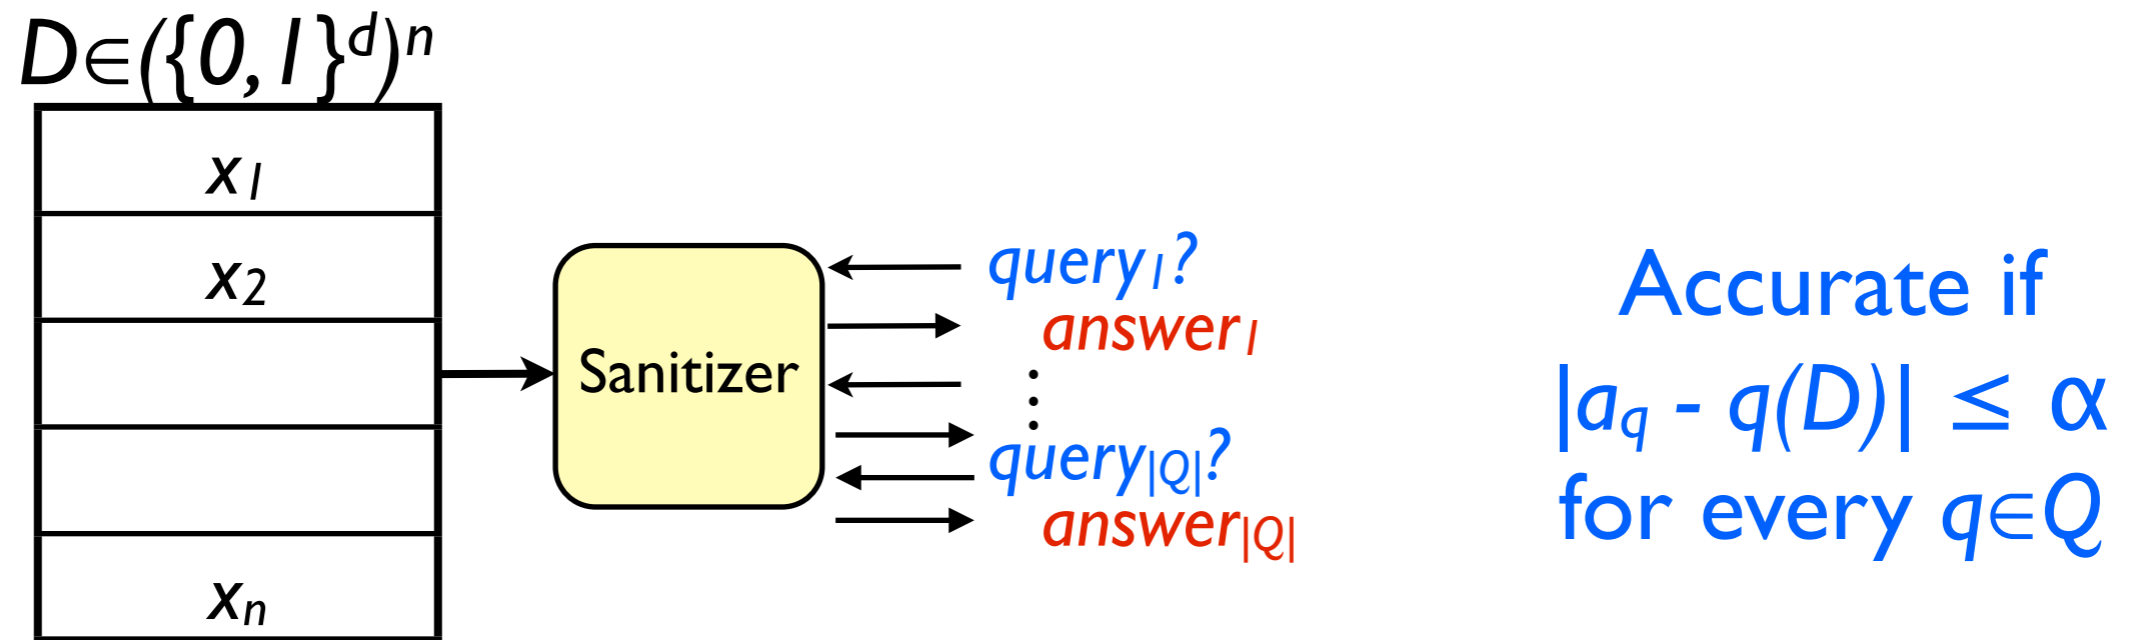
Optimality?

- Private multiplicative weights achieves nearly-optimal running time for this level of generality
- **Thm [U]:** any DP algorithm that takes a database $D \in (\{0, 1\}^d)^n$ and a set of counting queries Q , runs in time $\text{poly}(n, d, |q_1| + \dots + |q_{|Q|}|)$, and accurately answers Q requires $n \geq |Q|^{1/2}$
(assuming secure crypto exists)
- But PMW can be practical! [HLM]

Talk Outline

- Differentially private query release
- A blueprint for private query release
 - No-regret algorithms / MW
- Query Release Algorithms
 - Offline MW
 - **Online MW**
 - Variants
 - Faster algorithms for disjunctions via polynomial approx.

Online Counting Query Release



- Want to design an online sanitizer that is simultaneously differentially private and accurate
- Want to minimize
 - Amount of data required, n as a function of $|Q|, d, \alpha$
 - Running time of the sanitizer per query


A Blueprint for Query Release

Sanitized (DP) Output

Raw Data

Family of queries $Q?$ →

D_1




(*query*₁)

D_2



(*query*₂)

D_3



(*query*₃)

D_4



(*query*₄)

⋮

D_T



(OK)

D



Is D_1
good for $Q?$

Here's *query*₁


*query*₁ is a query D_1
answers incorrectly

A Blueprint for Online Query Release

Sanitized (DP) Output

Query $q_I?$ →

D_I





Raw Data

D




Is D_I
good for $q_I?$




Yes

A Blueprint for Online Query Release

Sanitized (DP) Output

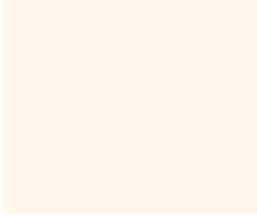


Raw Data



A Blueprint for Online Query Release

Sanitized (DP) Output

Query $q_1?$ \rightarrow D_1  (OK)
 $q_1(D_1)$ \leftarrow

Query $q_2?$ \rightarrow D_1  ($q_2(D)$)
 $q_2(D)$ \leftarrow

Raw Data

D



Is D_1
good for $q_2?$


\rightarrow
 \leftarrow
No, $q_2(D_2)$
should be $q_2(D)$

A Blueprint for Online Query Release

Sanitized (DP) Output

Query $q_1?$ →
 $q_1(D_1)$ ←


D_1



(OK)

Query $q_2?$ →
 $q_2(D)$ ←

D_1



($q_2(D)$)

D_2




Raw Data

D



Is D_1
good for $q_2?$

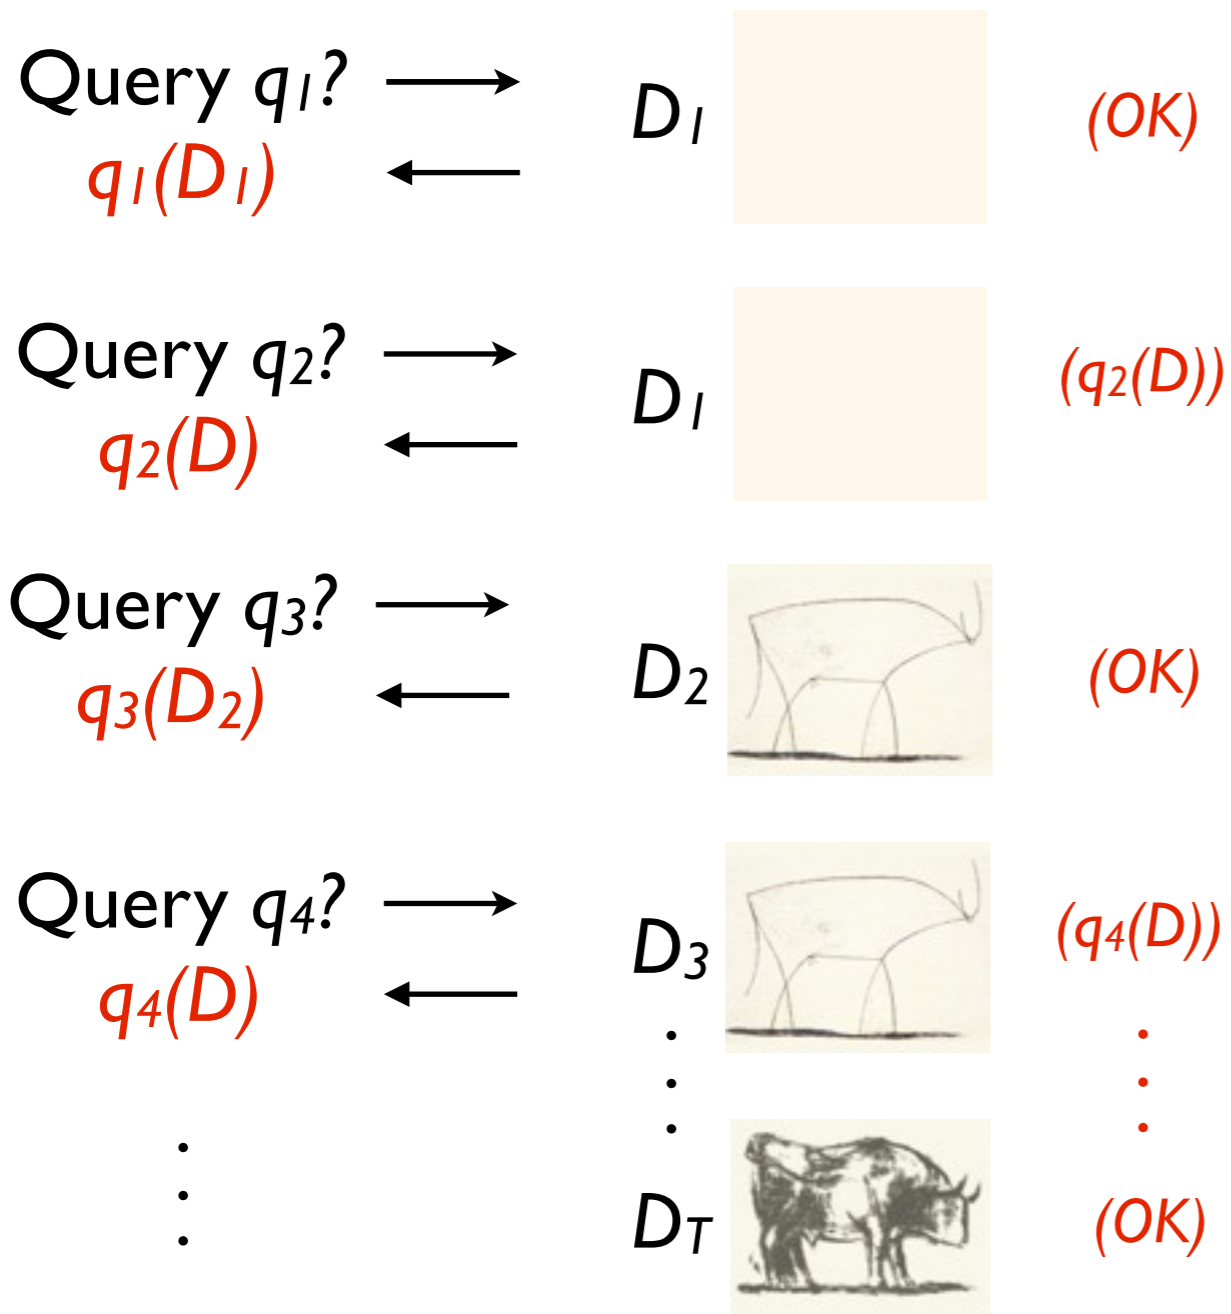


No, $q_2(D_2)$
is too low


A Blueprint for Online Query Release

Sanitized (DP) Output

Raw Data



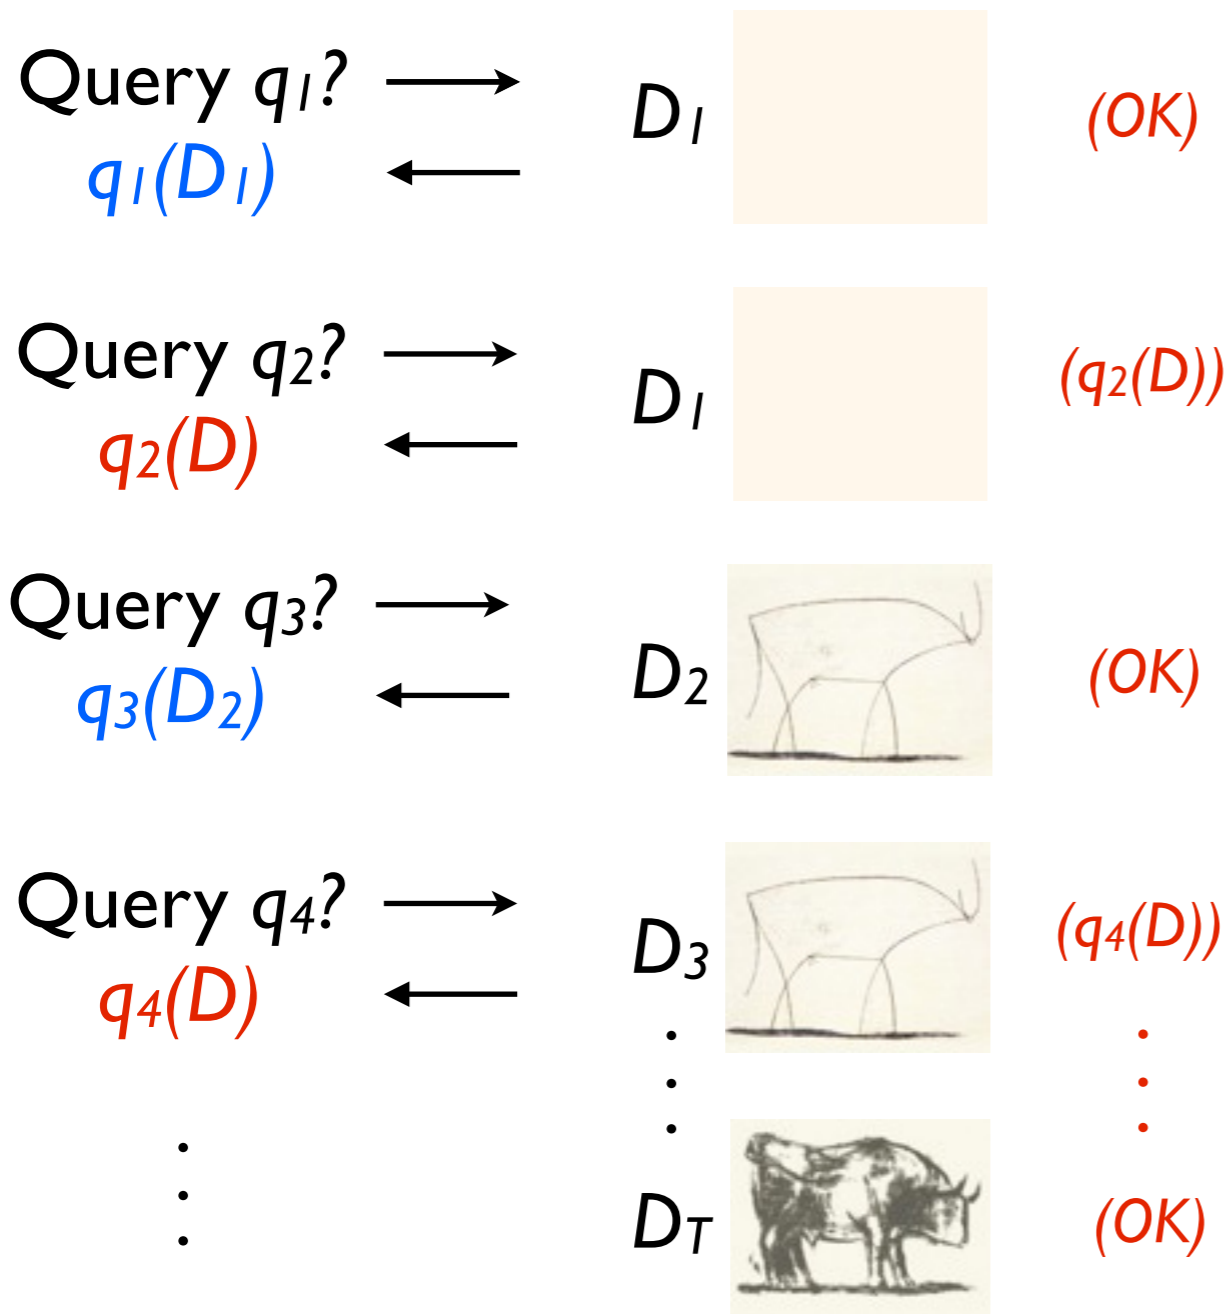
Is D_t
good for $q_k?$



A Blueprint for Online Query Release

Sanitized (DP) Output

Raw Data



Is D_t
good for $q_k?$



A Blueprint for Query Release

LET D be the real database, viewed as a dist over $\{0, 1\}^d$

LET D_t be the uniform dist on $\{0, 1\}^d$

FOR $k = 1, \dots, |Q|$

IF $|\langle D_t - D, q_k \rangle| \leq \alpha$ THEN answer $\langle D_t, q_k \rangle$

ELSE

answer $\langle D, q_k \rangle$, $D_{t+1} = \text{MWU}(D_t, q_k)$

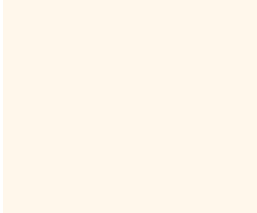
LET $t = t + 1$


$$T \leq d/\alpha^2$$

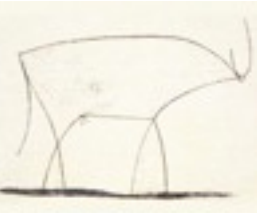
“Threshold” Algorithm


- Suppose we have a stream of queries q_1, \dots, q_k and promise that there is only a single q_i s.t. $q_i(D) \geq \alpha/2$
- Then there is an ϵ_0 -DP algorithm that whp answers every query with accuracy α as long as $n \gtrsim \log(k)/\alpha\epsilon_0$


A Blueprint for Online Query Release

Query $q_1?$ \longrightarrow D_1  (OK)
 $q_1(D_1)$ \longleftarrow

Query $q_2?$ \longrightarrow D_1  ($q_2(D)$)
 $q_2(D)$ \longleftarrow

Query $q_3?$ \longrightarrow D_2  (OK)
 $q_3(D_2)$ \longleftarrow

Query $q_4?$ \longrightarrow D_3  ($q_4(D)$)
 $q_4(D)$ \longleftarrow

\vdots
 \vdots D_T  (OK)
 \vdots

Instance of
threshold algorithm

Instance of
threshold algorithm

Recap

Thm: Online PMW takes a database $D \in (\{0, 1\}^d)^n$ and an online stream of counting queries Q , satisfies (ϵ, δ) -DP and, if $n \geq d^{1/2} \log |Q| / \alpha^2 \epsilon$, is α -accurate for all of Q

Thm: Runs in time $\text{poly}(n, 2^d, |q|)$ for each query q

Talk Outline

- Differentially private query release
- A blueprint for private query release
 - No-regret algorithms / MW
- Query Release Algorithms
 - Offline MW
 - Online MW
 - Variants
 - Faster algorithms for disjunctions via polynomial approx.

Other Applications

- PMW has optimal data requirement and running time in the worst case, but better algorithms are known for special cases
- Modular design makes it easy to construct new algorithms by swapping in different no-regret algorithms

Graph Cuts

[GRU]



- G in $(V \times V)^{|E|}$. Cut query $q_{S,T}(G)$ asks “What fraction of edges cross from S to T ?”
- Counting queries on a database D in $(\{0, 1\}^{2 \log |V|})^{|E|}$
- Can reduce the data requirement for some settings of parameters by replacing MW with an algorithm based on the “cut-decomposition” [FK]



Mirror Descent

[JT]

- Replace MW with algorithms from the mirror descent family
- Reduces the data requirement when the L_p norm of the database and L_q norm of the queries satisfy certain relationships
 - For PMW, we view the database as a distribution over $X = \{0, 1\}^d$ (L_1 norm = 1), we view the query as a vector in $[0, 1]^X$ (L_∞ norm = 1)
- Applications to cut queries, matrix queries

Sparse Queries


[BR]



- Query is sparse if it only accepts $S \ll 2^d$ elements from $\{0, 1\}^d$
- Can design an “implicit” implementation of MW that keeps track of $\sim S$ weights instead of 2^d
 - Improves running time per query from 2^d to $\sim S$
 - Also improves the data requirement slightly

Distance Queries

[HR]



- D in $([0, 1]^d)^n$. Query q_x is a point x in $[0, 1]^d$ and asks “What is the average distance between points in D and x ?”
- Can answer in time $\text{poly}(n, d)$ per query using a specialized no-regret algorithm for distance queries
- Improves data requirement in some cases too

Talk Outline

- Differentially private query release
- A blueprint for private query release
 - No-regret algorithms / MW
- Query Release Algorithms
 - Offline MW
 - Online MW
 - Variants
 - Faster algorithms for disjunctions via polynomial approx.

Private Counting Query Release

Counting query: What fraction of records satisfy property q ? e.g.

$q(x) = \text{GiveYouUp} \vee \text{LetYouDown}$

$$D \in (\{0, 1\}^d)^n$$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

$$q(x_1) = 0$$

$$q(x_2) = 1$$



$$q(x_3) = 1$$

$$q(x_4) = 1$$

$$q(D) = 3/4$$

d attributes per record

$$D \in (\{0, 1\}^d)^n$$



Queries Q

$a_1, \dots, a_{|Q|}$

Accurate if
 $|a_q - q(D)| < \alpha$
 for every $q \in Q$

Private Counting Query Release

Disjunction query: What fraction of records satisfy a given monotone k -way disjunction q_S , $|S| \leq k$?

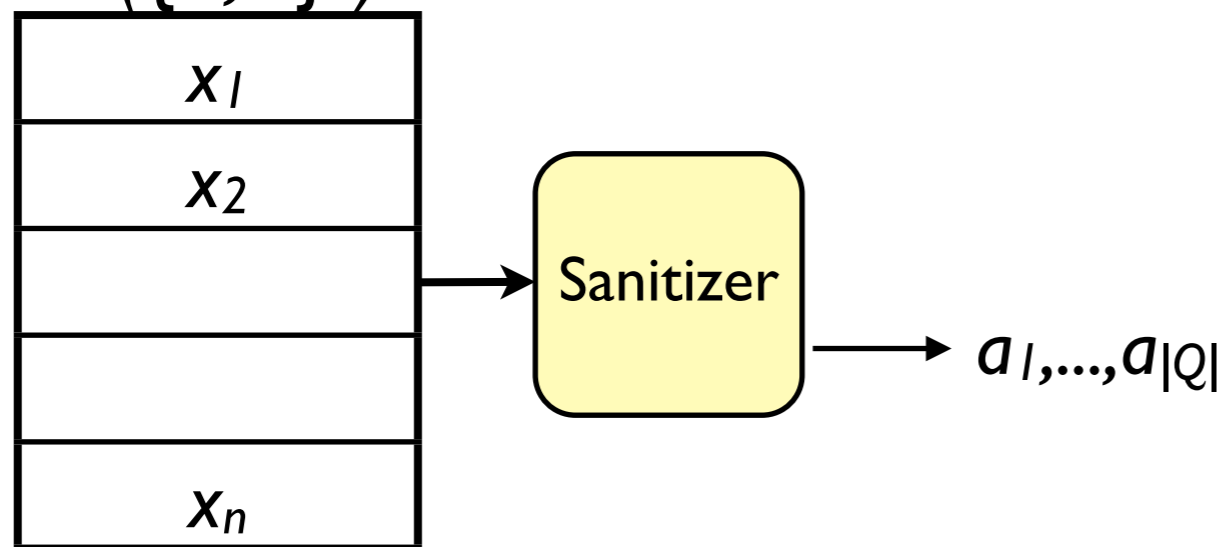
$$q_S(x) = \bigvee_{i \in S} x_i$$

$$D \in (\{0, 1\}^d)^n$$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

d attributes per record

$$D \in (\{0, 1\}^d)^n$$



Accurate if
 $|a_q - q(D)| < .01$
 for every $q \in Q$

Private Counting Query Release

$$D \in (\{0, 1\}^d)^n$$

Disjunction query: What fraction of records satisfy a given monotone k -way disjunction q_S , $|S| \leq k$?

$$q_S(x) = \bigvee_{i \in S} x_i$$

GiveYouUp?	LetYouDown?	RunAround?	DesertYou?
0	0	1	1
1	1	1	1
1	0	0	0
1	1	0	0

d attributes per record

- Useful facts:
 - Number of k -way disj's is d -choose- $k \sim d^k$
 - Equivalent to conjunctions / marginal queries / contingency tables

Algorithms for Disjunctions

Running Time



Minimum DB Size

Algorithms for Disjunctions

Running Time

$$d^k = |Q|$$


• Laplace Mechanism

$$d^{k/2} = |Q|^{1/2}$$

Minimum DB Size

Algorithms for Disjunctions

Running Time



Algorithms for Disjunctions

Running Time

2^d

$d^k = |Q|$

$poly(d, k)$

$k\sqrt{d}$

$d^{k/2} = |Q|^{1/2}$

Minimum DB Size

• PMW

• Laplace Mechanism

• Holy grail

Efficient Reduction to Learning

- The bottleneck in PMW is viewing the database as a distribution over $\{0, 1\}^d$

Efficient Reduction to Learning

- The bottleneck in PMW is viewing the database as a distribution over $\{0,1\}^d$
- Instead, view the database as a map $f_D: Q \rightarrow [0,1]$
 - If Q is “simple”, this map might have a nice structure that leads to more efficient algorithms
 - Doesn't even need to be defined for queries outside Q

Efficient Reduction to Learning

- View the database as a map $f_D: Q \rightarrow [0, 1]$
- **Thm (Approximately)** [HRS]: There is an efficient reduction from answering a family of queries Q to “learning” the family $\{f_D: Q \rightarrow [0, 1]\}_D$
 - Approach was implicit in [GHRU,CKKL]
- Using the learning techniques, without going through the reduction, gives simpler algorithms and stronger guarantees [TUV,CTUW]

Algorithms for Disjunctions

Running Time

2^d

$d^k = |Q|$

$poly(d, k)$

• PMW

• [HRS]

• Laplace Mechanism

• Holy grail

$k\sqrt{d}$


$d^{C\sqrt{k}}$

$d^{k/2} = |Q|^{1/2}$

Minimum DB Size

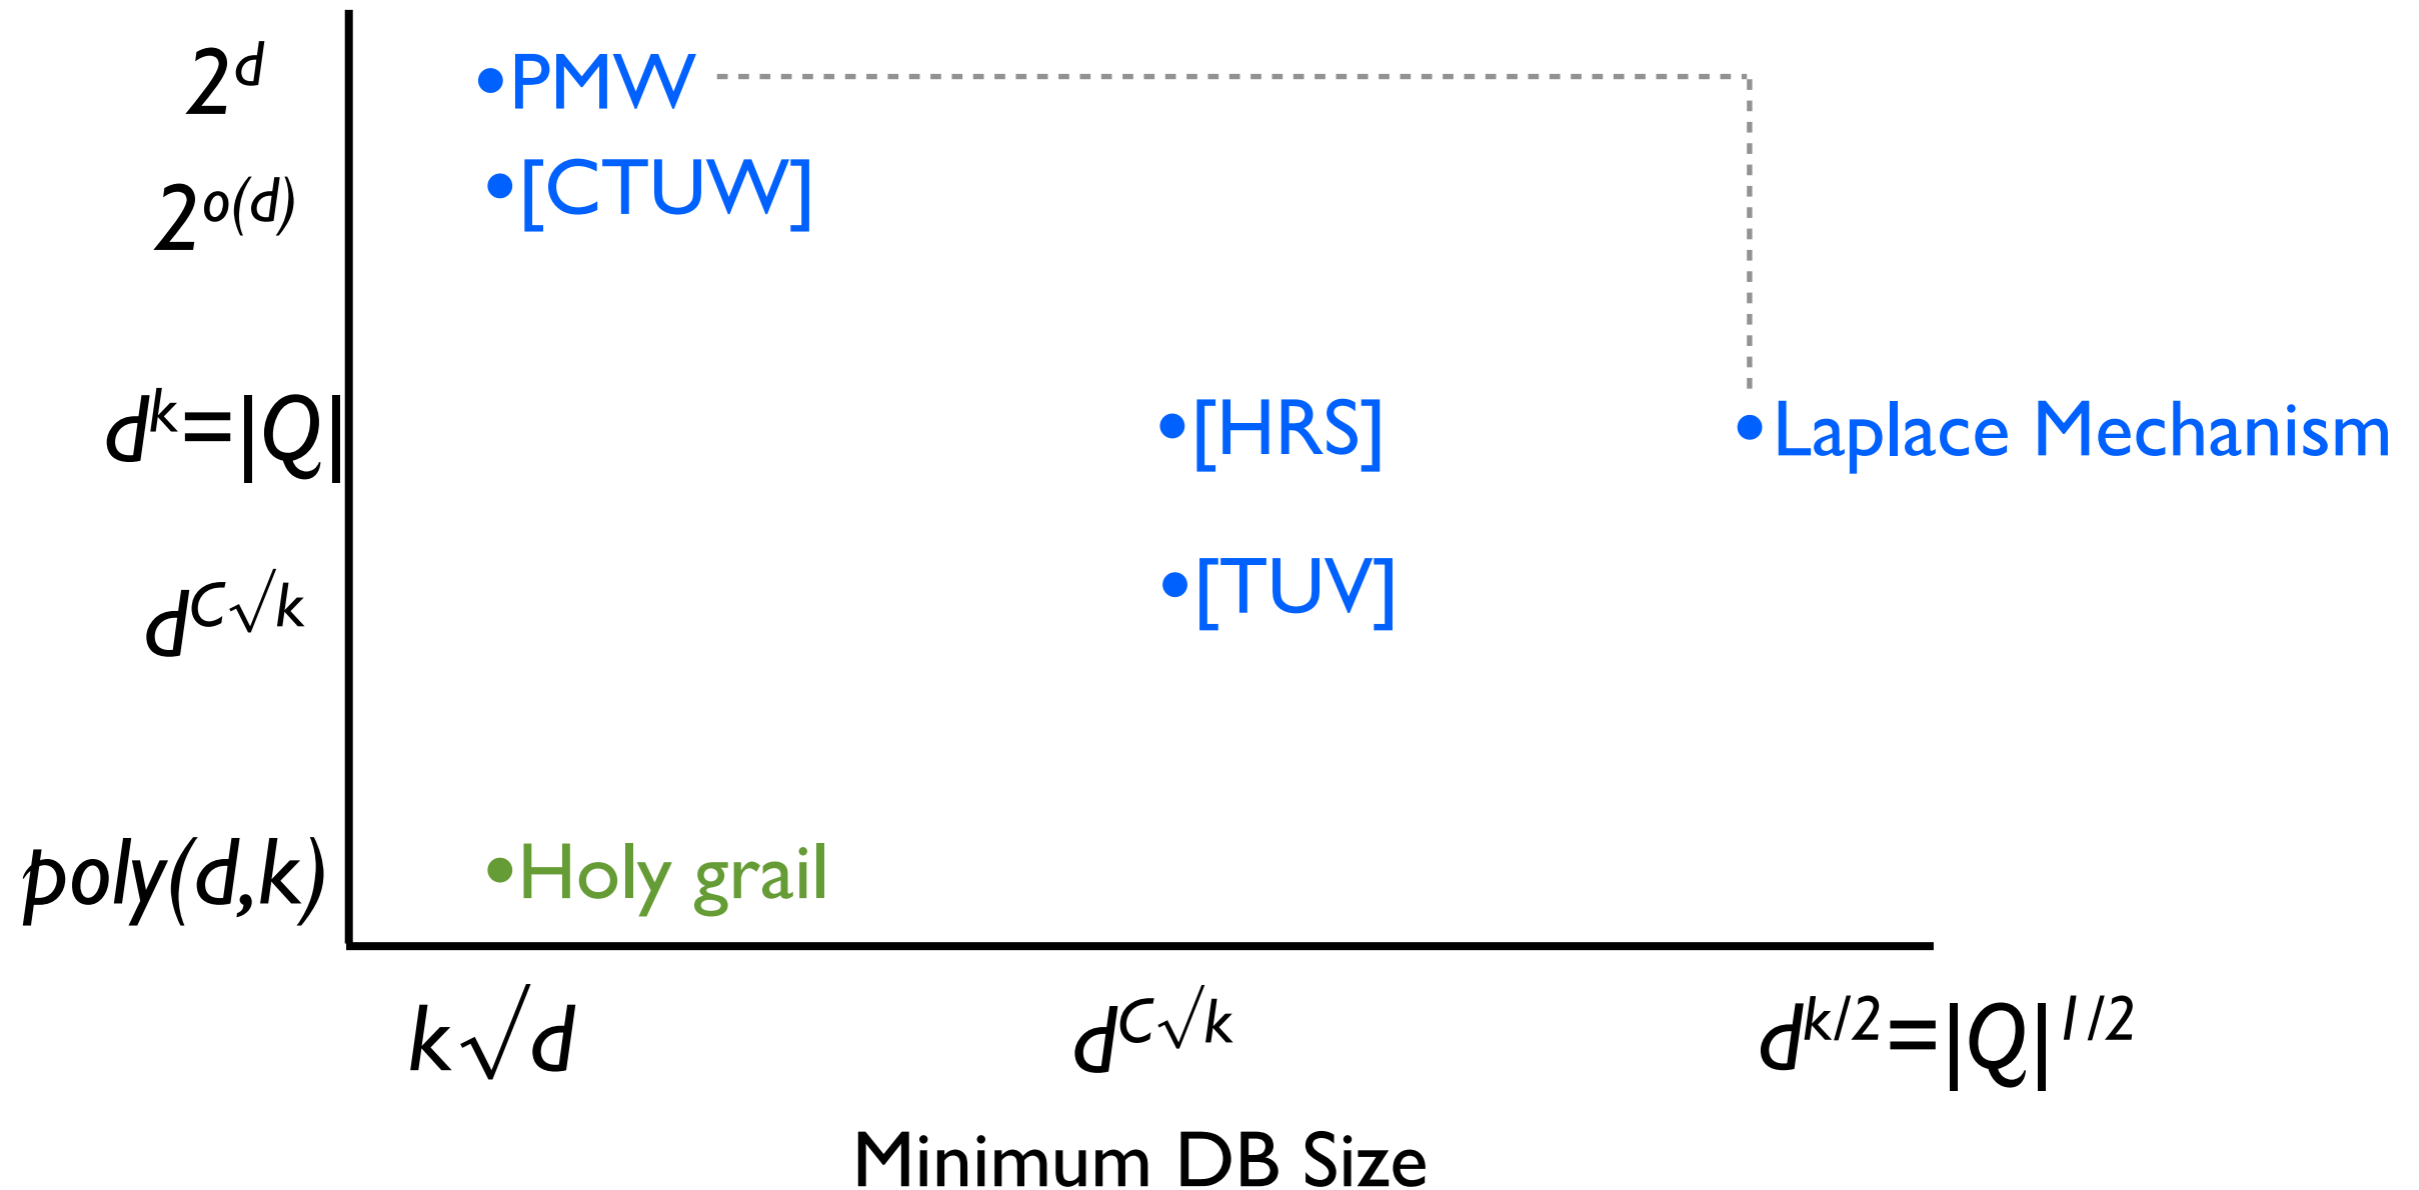
Algorithms for Disjunctions

Running Time




Algorithms for Disjunctions

Running Time



Algorithms for Disjunctions



Low-Weight Bases

- Instead, view the database as a map $f_D: Q \rightarrow [0, 1]$
 - If Q is “simple”, this map might have a nice structure that leads to more efficient algorithms
 - For disjunctions, f_D will be a “low-weight” linear combination of a small number of “basis functions”

Multiplicative Weights

Set of experts $X = \{0, 1\}^d$



Losses for each expert



Distribution over $X = \{0, 1\}^d$

D .25 .25 .25 .25

$q(D) = \langle D, q \rangle$

Truth table of q in $[0, 1]^X$

1 0 1 0 q

$q_x = 1$ iff $q(x) = 1$

Multiplicative Weights

Basis of functions $\{f_x\}$, x in $\{0, 1\}^d$



Losses for each expert



Weight 1 linear comb of fns in $\{f_x\}$

D .25 .25 .25 .25 $q(D) = \langle D, q \rangle$

Truth table of q in $[0, 1]^X$

1 0 1 0 q

Query function on a row:

$$f_x(q) = q(x)$$

Query function on a DB:

$$f_D(q) = (1/n) \sum_i f_{x_i}(q)$$

Losses for an expert x :

$$f_x(q) = q(x)$$

No-Regret Learning Algorithms

Set of experts X



Losses for each expert



Distribution over X

						$[0, 1]^X$				
D_1	.25	.25	.25	.25	Loss is $\langle D_1, L_1 \rangle$	1	0	1	0	L_1
D_2	.20	.30	.20	.30	Loss is $\langle D_2, L_2 \rangle$	0	0	1	0	L_2
					\vdots					
D_T	.23	.32	.15	.32	Loss is $\langle D_T, L_T \rangle$	0	0	0	1	L_T

For any distribution D , sequence L_1, \dots, L_T ,

$$\sum_{t=1}^T \langle D_t - D, L_t \rangle \leq \sqrt{T \log |X|}$$

No-Regret Learning Algorithms

Set of experts $X = F$



Losses for each expert



Weight W linear comb over $X = F$

						$[0, 1]^X$				
D_1	$ $	$ $	$ $	$ $	Loss is $\langle D_1, L_1 \rangle$	$ $	0	$ $	0	L_1
D_2	$.80$	1.20	$.80$	1.20	Loss is $\langle D_2, L_2 \rangle$	0	0	$ $	0	L_2
					\vdots					
D_T	$.92$	1.28	$.60$	1.28	Loss is $\langle D_T, L_T \rangle$	0	0	0	$ $	L_T

For any weight W linear combination D , sequence L_1, \dots, L_T ,

$$\sum_{t=1}^T \langle D_t - D, L_t \rangle \leq W \sqrt{T \log |X|}$$

Multiplicative Weights

Basis of functions $\{f_x\}$, x in $\{0, 1\}^d$



Losses for each expert



Weight 1 linear comb of fns in $\{f_x\}$

D .25 .25 .25 .25 $q(D) = \langle D, q \rangle$

Truth table of q in $[0, 1]^X$

1 0 1 0 q

Query function on a row:

$$f_x(q) = q(x)$$

Query function on a DB:

$$f_D(q) = (1/n) \sum_i f_{x_i}(q)$$

Losses for an expert x :

$$f_x(q) = q(x)$$

The Private MW algorithm treats the database as a weight 1 linear comb. of a set of 2^d functions $f_x: \{\text{All Queries}\} \rightarrow \{0, 1\}$

Multiplicative Weights

Basis of functions $\{f_x\}$, x in $\{0, 1\}^d$



Losses for each expert



Weight 1 linear comb of fns in $\{f_x\}$

D .25 .25 .25 .25 $q(D) = \langle D, q \rangle$

Truth table of q in $[0, 1]^X$

1 0 1 0 q

Query function on a row:

$$f_x(q) = q(x)$$

Query function on a DB:

$$f_D(q) = (1/n) \sum_i f_{x_i}(q)$$

Losses for an expert x :

$$f_x(q) = q(x)$$

Improved algs for disj's treat the database as a weight \mathbf{W} linear comb. of a set of \mathcal{S} functions $f: \{k\text{-way disj's}\} \rightarrow \{0, 1\}$

Low-Weight Bases

- View the database as a map $f_D: Q \rightarrow [0, 1]$
- Let $F = \{f: Q \rightarrow \{0, 1\}\}$ be a set of functions
- **Def:** F is a weight- W approximate basis wrt Q if for every database D , there exists a weight- W linear combination of functions in F , p_D , such that for every $q \in Q$, $|f_D(q) - p_D(q)| \leq .001$

No-Regret Learning Algorithms

Set of experts $X = F$



Losses for each expert



Weight W linear comb over $X = F$

						$[0, 1]^X$				
D_1	$ $	$ $	$ $	$ $	Loss is $\langle D_1, L_1 \rangle$	$ $	0	$ $	0	L_1
D_2	$.80$	1.20	$.80$	1.20	Loss is $\langle D_2, L_2 \rangle$	0	0	$ $	0	L_2
					\vdots					
D_T	$.92$	1.28	$.60$	1.28	Loss is $\langle D_T, L_T \rangle$	0	0	0	$ $	L_T

For any weight W linear combination D , sequence L_1, \dots, L_T ,

$$\sum_{t=1}^T \langle D_t - D, L_t \rangle \leq W \sqrt{T \log |X|}$$

Recap

Thm: PMW takes a database $D \in (\{0, 1\}^d)^n$ and a set of counting queries Q , satisfies (ϵ, δ) -DP and, if

$$n \geq d^{1/2} \log |Q| / \alpha^2 \epsilon,$$

it outputs D_T such that for every $q \in Q$,

$$|q(D) - q(D_T)| \leq \alpha$$

Thm: PMW runs in time $\text{poly}(n, 2^d, |q_1| + \dots + |q_{|Q|}|)$

Recap

Thm [CTUW]: PMW (run with F , a weight- W approximate basis wrt Q) takes a database $D \in (\{0, 1\}^d)^n$, satisfies (ϵ, δ) -DP and, if

$$n \geq Wd^{1/2} \log |Q| / \alpha^2 \epsilon,$$

it outputs D_T such that for every $q \in Q$,

$$|q(D) - q(D_T)| \leq .01$$

Thm: PMW runs in time $\text{poly}(n, |F|, |q_1| + \dots + |q_{|Q|}|)$

Low-Weight Bases

- But where do these low-weight bases come from?
- Polynomial approximations!
 - Extremely prevalent in PAC/agnostic learning. Underlies the most-efficient learning algorithms.
 - First used for disjunctions by [CKKL],[HRS]

Low-Weight Bases

Query on a row:

$$q(x) = x_1 \vee x_2$$

Query on a DB:

$$q(D) = (1/n) \sum_i q(x_i)$$

$$D \in (\{0, 1\}^d)^n$$

$x_1?$	$x_2?$	$x_3?$	$x_4?$
1	1	1	0
1	1	0	0
0	0	1	1
0	0	0	1

Low-Weight Bases

$$D \in (\{0, 1\}^d)^n$$

Query on a row:

$$q_y(x) = x_1 \vee x_2$$

Query on a DB:

$$q_y(D) = (1/n) \sum_i q_y(x_i)$$

$x_1?$	$x_2?$	$x_3?$	$x_4?$
1	1	1	0
1	1	0	0
0	0	1	1
0	0	0	1

Each query described by
a d -bit string $y \in \{0, 1\}^d$

Low-Weight Bases

$$D \in (\{0, 1\}^d)^n$$

$x_1?$	$x_2?$	$x_3?$	$x_4?$
1	1	1	0
1	1	0	0
0	0	1	1
0	0	0	1

Query on a row:

$$q_y(x) = x_1 \vee x_2$$

Query on a DB:

$$q_y(D) = (1/n) \sum_i q_y(x_i)$$

Each query described by
a d -bit string $y \in \{0, 1\}^d$

Query function on a row:

$$f_x(y) = q_y(x)$$

Query function on a DB:

$$f_D(y) = (1/n) \sum_i f_{x_i}(y)$$

Low-Weight Bases

Query on a row:

$$q_y(x) = x_1 \vee x_2$$

Query on a DB:

$$q_y(D) = (1/n) \sum_i q_y(x_i)$$

Each query described by
a d -bit string $y \in \{0, 1\}^d$

Query function on a row:

$$f_x(y) = q_y(x)$$

Query function on a DB:

$$f_D(y) = (1/n) \sum_i f_{x_i}(y)$$

$$D \in (\{0, 1\}^d)^n$$

$x_1?$	$x_2?$	$x_3?$	$x_4?$
1	1	1	0
1	1	0	0
0	0	1	1
0	0	0	1

Approximation: For every x , want $p_x(y)$ s.t.

- p_x has degree T
- p_x has weight W
- for every y corresponding to a k -way disj. $|p_x(y) - f_x(y)| \leq .001$

Low-Weight Bases

Query on a row:

$$q_y(x) = x_1 \vee x_2$$

Query on a DB:

$$q_y(D) = (1/n) \sum_i q_y(x_i)$$

Each query described by
a d -bit string $y \in \{0, 1\}^d$

Query function on a row:

$$f_x(y) = q_y(x)$$

Query function on a DB:

$$f_D(y) = (1/n) \sum_i f_{x_i}(y)$$

$$D \in (\{0, 1\}^d)^n$$

$x_1?$	$x_2?$	$x_3?$	$x_4?$
1	1	1	0
1	1	0	0
0	0	1	1
0	0	0	1

$$f_{(1,1,1,0)}(y_1, \dots, y_d) = y_1 \vee y_2 \vee y_3$$

Disjunction on y
(Coincidentally)

Low-Weight Bases

Query on a row:

$$q_y(x) = x_1 \vee x_2$$

Query on a DB:

$$q_y(D) = (1/n) \sum_i q_y(x_i)$$

Each query described by
a d -bit string $y \in \{0, 1\}^d$

Query function on a row:

$$f_x(y) = q_y(x)$$

Query function on a DB:

$$f_D(y) = (1/n) \sum_i f_{x_i}(y)$$

$$D \in (\{0, 1\}^d)^n$$

$x_1?$	$x_2?$	$x_3?$	$x_4?$
1	1	1	0
1	1	0	0
0	0	1	1
0	0	0	1

$$f(y_1, \dots, y_d) = \text{OR}(y_1, \dots, y_d)$$

Sufficient to approx.
 d -variate OR function
on inputs with at most
 k non-zeros


Recap

- Suppose there is a d -variate polynomial p of deg T and weight W such that for every y in $\{0, 1\}^d$ with at most k non-zeroes $|OR(y) - p(y)| \leq .001$.
- Then there is a weight- W approximate basis wrt k -way disj's of size roughly d -choose- T
 - $F = \{all\ d\text{-variate\ monomials\ of\ degree\ at\ most\ } T\}$

Approximating OR (Low Weight)

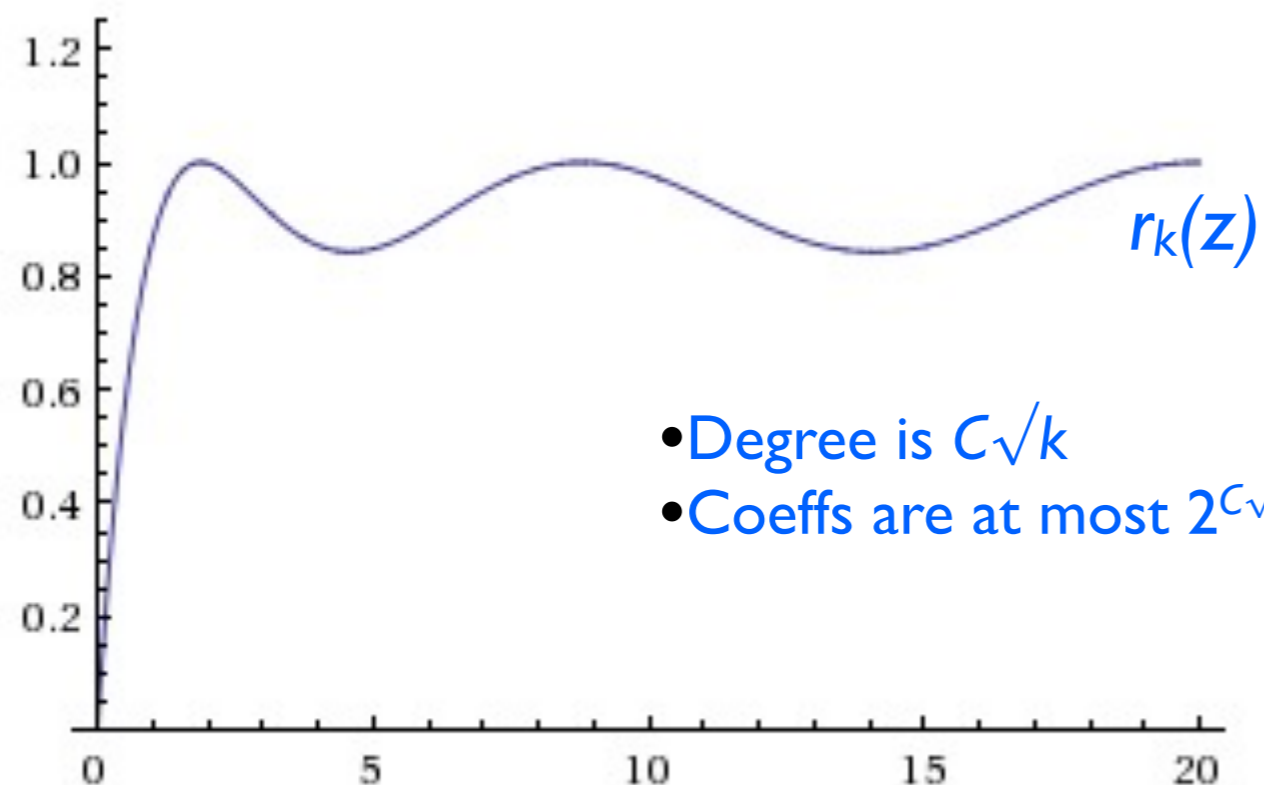
On $[1, k]$, $.999 \leq r(x) \leq 1.001$

- Want to approx $OR(y_1, \dots, y_d)$ on inputs with k non-zeros



Approximating OR (Low Weight)

On $[1, k]$, $.999 \leq r(x) \leq 1.001$




- Want to approx $OR(y_1, \dots, y_d)$ on inputs with k non-zeros

- Set

$$p(y_1, \dots, y_d) = r_k(y_1 + \dots + y_d)$$

Approximating OR (High Weight)

On $[1, k]$, $.999 \leq r(x) \leq 1.001$



- Want to approx $OR(y_1, \dots, y_d)$ on inputs with k non-zeros

- Set

$$p(y_1, \dots, y_d) = r_k(y_1 + \dots + y_d)$$

- If $OR(y_1, \dots, y_d) = 0$, then


$$p(y_1, \dots, y_d) = r_k(0) = 0$$

- If $OR(y_1, \dots, y_d) = 1$, then $1 \leq y_1 + \dots + y_d \leq k$

$$p(y_1, \dots, y_d) = r_k(y_1 + \dots + y_d) \approx 1$$

Approximating OR (High Weight)

On $[1, k]$, $.999 \leq r(x) \leq 1.001$



• Want to approx $OR(y_1, \dots, y_d)$ on inputs with k non-zeros

• Set

$$p(y_1, \dots, y_d) = r_k(y_1 + \dots + y_d)$$

• If $OR(y_1, \dots, y_d) = 0$, then

$$p(y_1, \dots, y_d) = r_k(0) = 0$$


• If $OR(y_1, \dots, y_d) = 1$, then $1 \leq y_1 + \dots + y_d \leq k$

$$p(y_1, \dots, y_d) = r_k(y_1 + \dots + y_d) \approx 1$$

Polynomial has degree $C\sqrt{k}$, weight $d^{C\sqrt{k}}$

Algorithms for Disjunctions

Running Time



Approximating OR (Low Weight)

- Have an approximation with degree $C\sqrt{k}$ and weight $d^{C\sqrt{k}}$


- The “trivial” exact polynomial has degree d and weight 1

OR_d

Approximating OR (Low Weight)

- Have an approximation with degree $C\sqrt{k}$ and weight $d^{C\sqrt{k}}$

- The “trivial” exact polynomial has degree d and weight 1

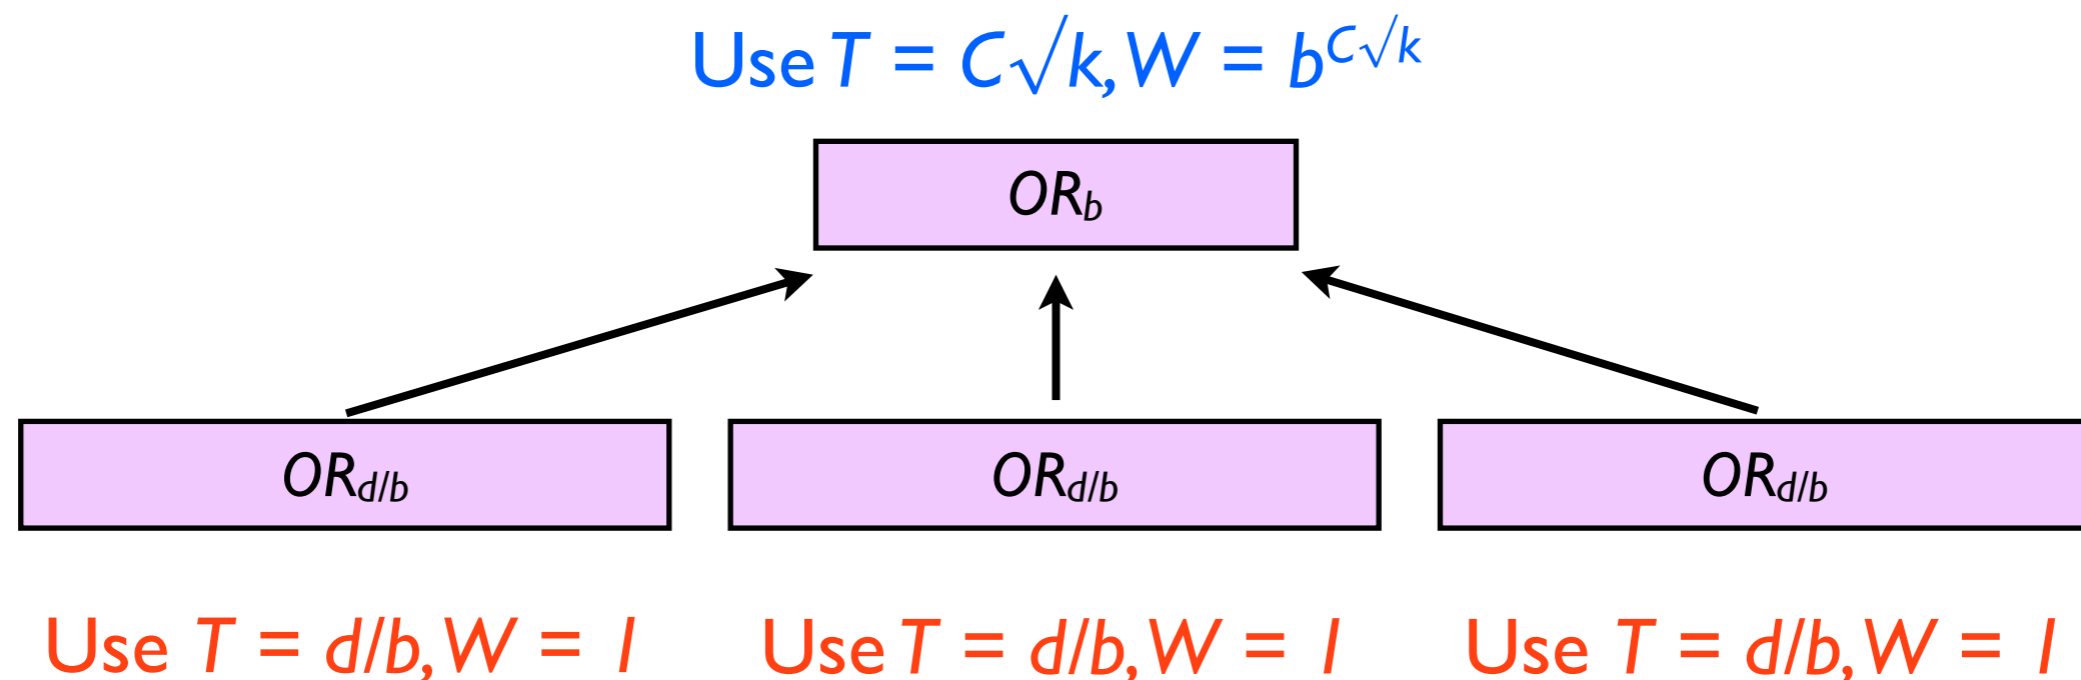


Final polynomial has degree $C(d/b)\sqrt{k}$, weight $b^{C\sqrt{k}}$

Approximating OR (Low Weight)

- Have an approximation with degree $C\sqrt{k}$ and weight $d^{C\sqrt{k}}$

- The “trivial” exact polynomial has degree d and weight 1

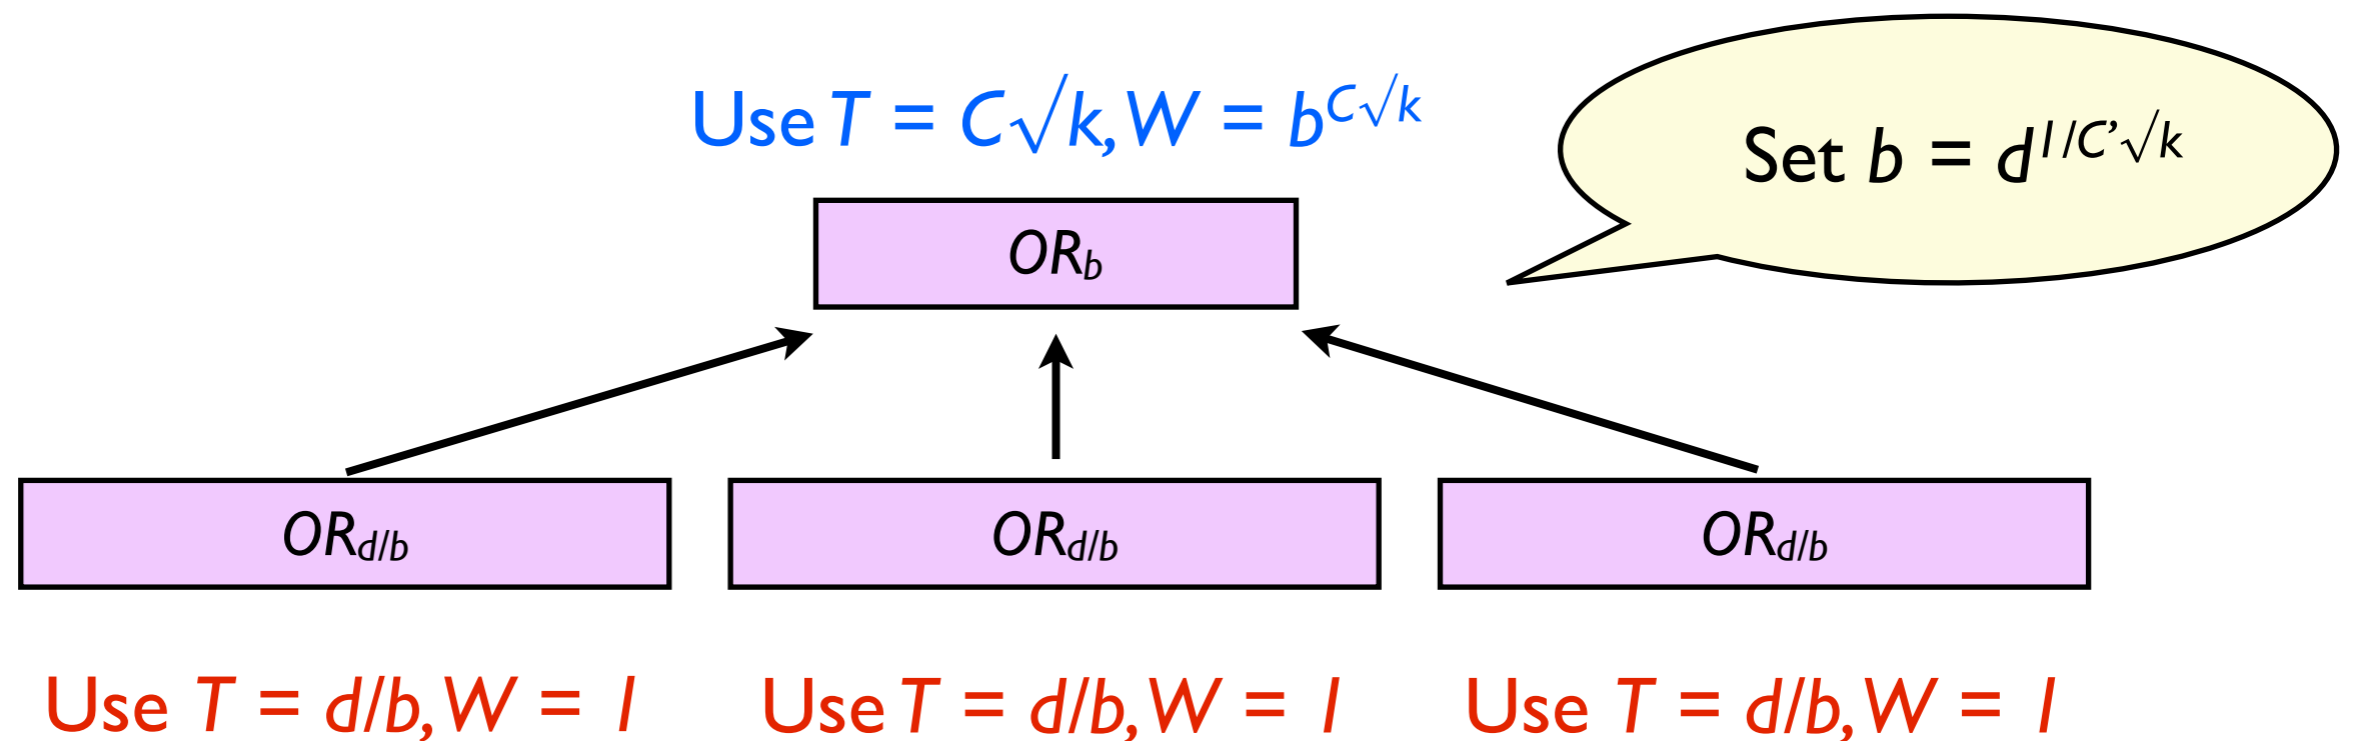


Final polynomial has degree $C(d/b)\sqrt{k}$, weight $b^{C\sqrt{k}}$

Approximating OR (Low Weight)

- Have an approximation with degree $C\sqrt{k}$ and weight $d^{C\sqrt{k}}$


- The “trivial” exact polynomial has degree d and weight 1



Final polynomial has degree $\sim d^{1-1/C\sqrt{k}}$, weight $\sim d^{0.1}$

Algorithms for Disjunctions

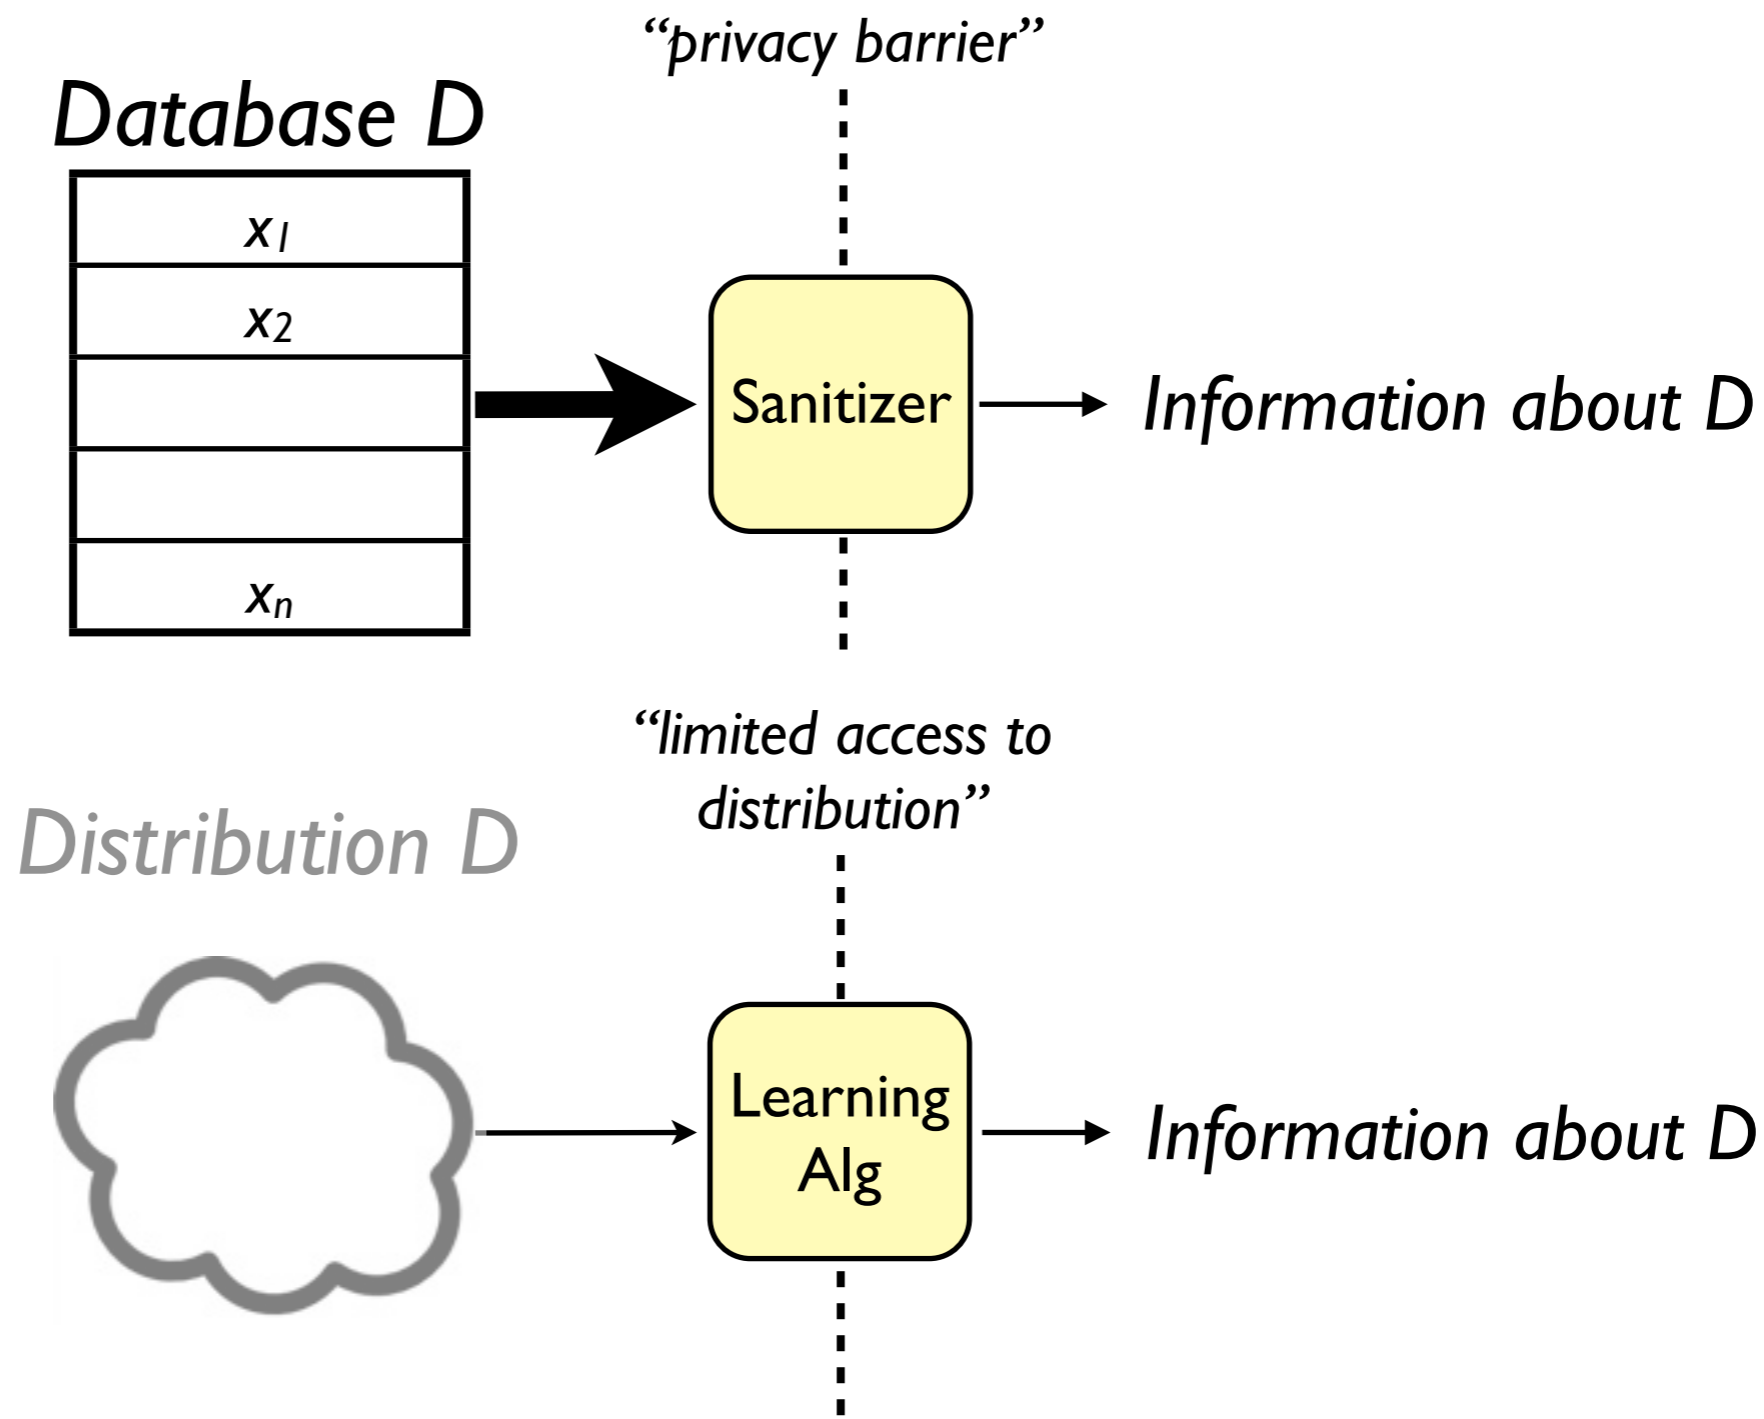
Running Time



Can these results be improved?

- Not using polynomials! [CTUW]
- In the high-weight setting, there is no approximate basis smaller than $d^{C\sqrt{k}}$ [S]
- Open question: What is the smallest weight-poly(d) basis wrt to $\{k\text{-way disj}\}$?


What about using different techniques?



Can these results be improved?

- Sometimes we can improve running time by avoiding learning algorithms altogether.

Algorithms for Disjunctions



Wrap-Up

- There is a flexible, modular framework for deriving differentially private algorithms from learning-theoretic techniques
- For the general private counting query release problem, these techniques (PMW) give optimal accuracy and running time guarantees
- For natural, special cases of query release, learning techniques (often) give best-known algorithms
 - But is this the right approach?

Thanks!