

Extreme Points of the Local Differential Privacy Polytope

Naoise Holohan^a, Douglas J. Leith^a, Oliver Mason^{b,*}

^a*School of Computer Science and Statistics, Trinity College Dublin, Ireland*

^b*Dept. of Mathematics and Statistics/Hamilton Institute, Maynooth University, Co. Kildare, Ireland & Lero, the Irish Software Research Centre*

Abstract

We study the convex polytope of $n \times n$ stochastic matrices that define locally ϵ -differentially private mechanisms. We first present invariance properties of the polytope and results reducing the number of constraints needed to define it. Our main results concern the extreme points of the polytope. In particular, we completely characterise these for matrices with 1, 2 or n non-zero columns.

Keywords: Data Privacy, Stochastic Matrices, Matrix Polytopes, Differential Privacy.

2010 MSC: 68R01, 68R05, 60C05

1. Introduction

Data privacy has been of interest to researchers in computer science [1], statistics, cryptography [9] and law [7] for decades. The recent emergence of ‘Big Data’, while offering significant potential benefits to business and society, poses very real risks to personal privacy; this naturally has led to increased interest in questions pertaining to data privacy. The concept of *Differential Privacy*, introduced by C. Dwork in 2006 [13], has emerged as a popular theoretical paradigm in privacy research within the computer science community and has been applied to various different types of data and queries [12].

We are interested in the geometry of matrix polytopes arising in the study of differential privacy for categorical or finite-valued datasets. More formally, we consider databases $\mathbf{d} \in D^N$ where the set D is finite and can, without loss of generality, be taken to be $\{1, \dots, n\}$. Each entry in \mathbf{d} , d_i , corresponds to data contributed by an individual; the base set D describes all the values that data entries can take.

The problem we consider is motivated by the construction of differentially private sanitisations, where we are interested in releasing a private, sanitised

*Corresponding author. Tel.: +353 (0)1 7083672; fax: +353 5(0)1 7083913; email: oliver.mason@nuim.ie

version of a database \mathbf{d} . A *sanitisation* is defined by a set of random variables $X_{\mathbf{d}}$ taking values in D^N for every $\mathbf{d} \in D$. Loosely speaking, $X_{\mathbf{d}}$ describes a noisy version of the original \mathbf{d} designed to protect the privacy of individual data contributors.

The differential privacy model specifies two privacy parameters, $\epsilon \geq 0$ and $0 \leq \delta \leq 1$. For any two databases $\mathbf{d}, \mathbf{d}' \in D^N$ that differ in one row only, (ϵ, δ) -differential privacy requires

$$\mathbb{P}(X_{\mathbf{d}} \in A) \leq e^\epsilon \mathbb{P}(X_{\mathbf{d}'} \in A) + \delta, \quad (1)$$

for all $A \subseteq D^N$.

In essence, differential privacy ensures that answers to queries on a database cannot change greatly when one person's information in a database is altered.

The above definition considers global privacy with the mechanism defined on a complete database. Global mechanisms can readily be constructed using locally private mechanism, where subjects perturb/sanitise their own data locally before providing it to a central database upon which queries are answered [11]. The concept of local privacy first appeared over 50 years ago as a way to eliminate bias in surveying [22] and is known in other contexts as *input perturbation* or *randomised response* [22, 15]. A rigorous mathematical framework has been developed which guarantees global differential privacy when local differential privacy methods are applied [17].

We refer to local mechanisms as 1-dimensional mechanisms, as they take a single row of a database as an input, and output another (perturbed/noisy) row. In our context, a 1-dimensional mechanism is specified by giving an appropriate probability mass function p_i for every $i \in D = \{1, \dots, n\}$. More compactly, a 1-dimensional mechanism is defined by a stochastic matrix $A \in \mathbb{R}^{n \times n}$ where a_{ij} denotes the probability of outputting j when the input, or real data, is i . The requirement for local differential privacy is then given by:

$$a_{ik} \leq e^\epsilon a_{jk} + \delta \quad (2)$$

for all i, j, k . These constraints, taken together with the stochastic and non-negativity constraints, define the local differential privacy polytope. We shall consider the simplified case of strict differential privacy (which is what was originally introduced by Dwork) where $\delta = 0$ here.

In practice, we are interested in finding a mechanism (i.e. a matrix in this polytope) which is optimal for some utility function. Understanding the geometry of the polytope guides the design of such mechanisms. For instance, if the utility function happens to be linear, then the optimal mechanism occurs at an extreme point of the polytope. The search for optimal mechanisms in differential privacy has been studied by a number of authors [20, 21, 19]. Local differential privacy has been studied recently in the paper [11], while extremal local differential privacy mechanisms were considered in [18]. Of course, polytopes of stochastic matrices and doubly stochastic matrices have been widely studied in the past [4, 6, 8, 5]. An alternative study on geometrical aspects of differential privacy can be found in [16].

The basic layout of the paper is as follows. In Section 2 we introduce preliminary definitions of polytopes, extreme points, the concept of differential privacy and the polytope with which we will be working. In Section 3 we look at some elementary results for extreme points of this polytope, and in Section 4 we present our main results. In Section 5 we examine a number of special cases for extreme points, and finish with some concluding remarks in Section 6.

2. Notation and Background

To begin, let us introduce the major notation and standard definitions to be used in our results. For a matrix $A \in \mathbb{R}^{n \times n}$ and $1 \leq i \leq n$, we will use $A^{(i)}$ to denote the i th column of A . A^T denotes the usual matrix transpose. We denote by $\mathbf{1}$ the (column) vector of all ones where the dimension will typically be clear from context. We denote by e_i , $1 \leq i \leq n$, the i th standard basis vector of \mathbb{R}^n .

2.1. Polyhedra

In this paper, we adopt the following definitions for polyhedra and polytopes

Definition 1 Let $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}$ be an inner product on a real vector space V , and let $\{c^{(1)}, \dots, c^{(q+l)}\} \subseteq V$ and $b \in \mathbb{R}^{q+l}$ be given. A convex polyhedron $\mathcal{P} \subseteq V$ is defined as:

$$\mathcal{P} = \left\{ v \in V : \begin{array}{l} \langle c^{(i)}, v \rangle = b_i, \quad \forall 1 \leq i \leq q, \\ \langle c^{(q+i)}, v \rangle \leq b_{q+i}, \quad \forall 1 \leq i \leq l. \end{array} \right\}. \quad (3)$$

An inequality constraint is said to be *tight* or *active* on a point v if $\langle c^{(q+i)}, v \rangle = b_{q+i}$.

Definition 2 A convex polytope in a vector space V is the convex hull of a finite collection of points in V .

$$\mathcal{P} = \text{conv}(v_1, \dots, v_k), \quad (4)$$

where $v_i \in V$ for all i .

It is well known that all polytopes are polyhedra, but only bounded polyhedra are polytopes.

An *extreme point* of a polyhedron cannot be written as the convex combination of any other points in the polyhedron.

Definition 3 (Extreme point) Let $\mathcal{P} \subseteq \mathbb{R}^n$ be a convex polyhedron. A point $v \in \mathcal{P}$ is an extreme point of \mathcal{P} if $w, z \in \mathcal{P}$, $\frac{1}{2}(w + z) = v$, implies $w = z = v$.

We denote by $\text{ex}(\mathcal{P})$ the set of all extreme points of a polyhedron \mathcal{P} .

Our primary interest is in characterising the extreme points of the local differential privacy polyhedron. The following theorem from convex geometry shall prove useful in this regard [3].

Theorem 1 Let $\mathcal{P} \subseteq V$ be a polyhedron, and consider a point $v \in \mathcal{P}$. Denote by the set $I_v \subseteq \{1, \dots, l\}$ the indices of the inequality constraints that are tight on v (i.e. $\langle c^{(q+i)}, v \rangle = b_{q+i}$ for all $i \in I_v$ and $\langle c^{(q+i)}, v \rangle < b_{q+i}$ for all $i \in \{1, \dots, l\} \setminus I_v$). Then $v \in \text{ex}(\mathcal{P})$ if and only if

$$\text{span} \left(\left\{ c^{(1)}, \dots, c^{(q)} \right\} \cup \left\{ c^{(q+i)} : i \in I_v \right\} \right) = V.$$

Essentially, this result tells us that v is an extreme point of P if and only if there are n linearly independent constraints tight on v where n is the dimension of V .

2.2. Differential Privacy

As in the Introduction, we take the set $D = \{1, \dots, n\}$ to be the domain of the rows of our database (i.e. each subject contributes a value of D to the database). For local differential privacy to be satisfied, we require:

$$\mathbb{P}(X_i \in I) \leq e^\epsilon \mathbb{P}(X_j \in I) + \delta,$$

for all $i, j \in \{1, \dots, n\}$ and for all $I \subset D$.

For the purpose of this paper, we only consider the case of strict or non-relaxed differential privacy, where $\delta = 0$. In this case, the requirement simplifies to

$$\mathbb{P}(X_i = k) \leq e^\epsilon \mathbb{P}(X_j = k),$$

for all $i, j, k \in \{1, \dots, n\}$.

If we let $A \in \mathbb{R}^{n \times n}$ be given by,

$$a_{ij} = \mathbb{P}(X_i = j)$$

then A defines a valid ϵ -differential privacy mechanism if and only if the following conditions hold:

$$\sum_j a_{ij} = 1, \quad 1 \leq i \leq n, \quad (5a)$$

$$a_{ij} \geq 0, \quad 1 \leq i, j \leq n, \quad (5b)$$

$$a_{ik} \leq e^\epsilon a_{jk}, \quad 1 \leq i, j, k \leq n. \quad (5c)$$

We now define the ϵ -differential privacy polytope, comprised of all matrices satisfying the above constraints.

Definition 4 (Differential Privacy Polytope) Fix $n \in \mathbb{N}$ and $\epsilon \geq 0$. The ϵ -differential privacy polytope, $\mathcal{D} \subset \mathbb{R}^{n \times n}$, is defined as follows:

$$\mathcal{D} = \left\{ A \in \mathbb{R}^{n \times n} : \begin{array}{ll} \sum_j a_{ij} = 1, & \forall 1 \leq i \leq n, \\ a_{ij} \geq 0, & \forall 1 \leq i, j \leq n, \\ a_{ij} \leq e^\epsilon a_{kj}, & \forall 1 \leq i, j, k \leq n. \end{array} \right\}. \quad (6)$$

The non-negativity and stochastic constraints ensure \mathcal{D} is bounded. Therefore it is a polytope.

Note: As the constraint $a_{ij} \leq e^\epsilon a_{kj}$ must hold for all i, j, k , we require $e^{-\epsilon} a_{kj} \leq a_{ij} \leq e^\epsilon a_{kj}$ for each i, j, k . Equivalently, $\max_i a_{ij} \leq e^\epsilon \min_i a_{ij}$ for all j .

Remark: If $\epsilon = 0$, then for A to be in \mathcal{D} , we require that $a_{ij} = a_{kj}$ for all i, j, k .

Using the Hilbert Schmidt inner product

$$\langle X, Y \rangle = \text{tr}(X^T Y),$$

together with the matrices $e_i e_j^T$, $e_i \mathbf{1}^T$, $e_i e_j^T = e^\epsilon e_k e_j^T$, it is not a difficult exercise to represent the constraints defining \mathcal{D} in the form given in Definition 1.

3. Preliminary results

In this section, we present several preliminary results on the structure of the set \mathcal{D} and its extreme points. We first note that the nonnegativity constraint in the definition of \mathcal{D} is redundant in the case where $\epsilon > 0$.

Lemma 1 *Fix $\epsilon > 0$. Let $v \in \mathbb{R}^n$ satisfy $v_i \leq e^\epsilon v_j$ for all i, j . Then $v \geq 0$.*

Proof: Let $v_i < 0$ for some i . Then, for each j , we have:

$$\begin{aligned} & e^{-\epsilon} v_i \leq v_j \leq e^\epsilon v_i \\ \Rightarrow & e^{-\epsilon} v_i \leq e^\epsilon v_i \\ \Rightarrow & e^{-\epsilon} \geq e^\epsilon \\ \Rightarrow & \epsilon \leq 0 \end{aligned}$$

By hypothesis $\epsilon > 0$. Hence, we must have $v_i \geq 0$ for each i . \square

Our next lemma notes that if the differential privacy constraint is tight on two elements in a column, that those two elements must be the minimum and maximum entries of that column.

Lemma 2 *Let $v \in \mathbb{R}^n$ be a vector with $v_i \leq e^\epsilon v_j$ for all $1 \leq i, j \leq n$. Suppose there exists at least one pair i, j where $v_i = e^\epsilon v_j$. Then $\min_k v_k = v_j$ and $\max_k v_k = v_i$.*

Proof: Suppose there exists v_l such that $v_l > v_i$. Then, $v_l > e^\epsilon v_j$, contradicting the differential privacy constraints. Similarly, if $v_l < v_j$, then $e^\epsilon v_l < v_i$. The result follows. \square

Several of our results will relate the extreme points A of \mathcal{D} to the non-zero columns in A . With this in mind, we formally define

$$\gamma(A) = \{i \in \{1, \dots, n\} : A^{(i)} \neq 0\}.$$

So that $\gamma(A)$ consists of the indices of the non-zero columns of A and $1 \leq |\gamma(A)| \leq n$ gives the number of non-zero columns in A .

Our next result concerns the rank of the extreme points of \mathcal{D} ; first we note the simple observation that $\text{rank}(A) \leq |\gamma(A)|$ for all A .

Theorem 2 *Let $A \in \text{ex}(\mathcal{D})$. Then*

$$\text{rank}(A) = |\gamma(A)|.$$

Proof: Suppose $A \in \text{ex}(\mathcal{D})$. As noted before, $\text{rank}(A) \leq |\gamma(A)|$. If A has only one non-zero column, then clearly $\text{rank}(A) = 1 = |\gamma(A)|$.

Let $|\gamma(A)| > 1$ and suppose $\text{rank}(A) < |\gamma(A)|$. Then there exists $\eta \in \mathbb{R}^n$, $\eta \neq 0$ and $\eta_i = 0$ for all $i \notin \gamma(A)$ (i.e. whenever $A^{(i)} = 0$), such that $\sum_i \eta_i A^{(i)} = 0$.

Let $B = A \text{diag}(\eta)$. By construction, $B\mathbf{1} = 0$.

Consider $C = A - \Delta B$, $D = A + \Delta B$, where $0 < \Delta < \frac{1}{\max_i |\eta_i|}$. Then,

1. C and D are stochastic, as A is stochastic and $B\mathbf{1} = 0$;
2. since $a_{ij} \leq e^\epsilon a_{kj}$ and $(1 \pm \Delta \eta_j) > 0$ for all i, j, k , it follows that $c_{ij} \leq e^\epsilon c_{kj}$, $d_{ij} \leq e^\epsilon d_{kj}$; and
3. $C, D \geq 0$.

Hence, C and D are in \mathcal{D} and $C \neq D$ as $B \neq 0$.

However, $\frac{1}{2}(C + D) = A$, so $A \notin \text{ex}(\mathcal{D})$, a contradiction. Therefore, for every $A \in \text{ex}(\mathcal{D})$, $\text{rank}(A) = |\gamma(A)|$. \square

We shall often make implicit use of the following simple corollary to the above result; essentially it states that for an extreme point A with at least 2 non-zero columns, none of these columns can have all their entries equal.

Corollary 1 *Let $A \in \text{ex}(\mathcal{D})$ satisfy $|\gamma(A)| \geq 2$. Then there is no $i \in \gamma(A)$, $k \in \mathbb{R}$ with $A^{(i)} = k\mathbf{1}$.*

Proof: Suppose that there is some $k \in \mathbb{R}$, $i_0 \in \gamma(A)$ such that $A^{(i_0)} = k\mathbf{1}$. Clearly, $k \neq 0$ as $i_0 \in \gamma(A)$ and $k \neq 1$ as $|\gamma(A)| \geq 2$. As A is stochastic,

$$\sum_{i \in \gamma(A)} A^{(i)} = \mathbf{1}$$

which implies that

$$\left(1 - \frac{1}{k}\right)A^{(i_0)} + \sum_{i \in \gamma(A), i \neq i_0} A^{(i)} = 0.$$

This implies that $\text{rank}(A) < |\gamma(A)|$, contradicting Theorem 2. \square

It is clear from the definition that \mathcal{D} is closed under row/column permutations. Our next result notes that this same invariance property also holds for extreme points.

Theorem 3 *Let $A \in \mathcal{D}$ and let $P_1, P_2 \in \{0, 1\}^{n \times n}$ be permutation matrices. Then $P_1 A P_2 \in \mathcal{D}$. Furthermore, $A \in \text{ex}(\mathcal{D})$ if and only if $P_1 A P_2 \in \text{ex}(\mathcal{D})$.*

Proof: From Definition 4, it clearly follows that $P_1AP_2 \in \mathcal{D}$ if $A \in \mathcal{D}$, since permuting a matrix only changes the order of rows and columns, but elements in the same row/column will remain in a common row/column.

Now, let $A \in \text{ex}(\mathcal{D})$ and suppose that $P_1AP_2 \notin \text{ex}(\mathcal{D})$ for some permutation matrices P_1 and P_2 . Then, there exist $B \neq C \in \mathcal{D}$ such that,

$$P_1AP_2 = \frac{1}{2}(B + C).$$

However, since $P^{-1} = P^T$ for any permutation matrix P , we have

$$A = \frac{1}{2}(P_1^TBP_2^T + P_1^TC P_2^T),$$

where $P_1^TBP_2^T \neq P_1^TC P_2^T$ since $B \neq C$. This is a contradiction since $A \in \text{ex}(\mathcal{D})$, hence $P_1AP_2 \in \text{ex}(\mathcal{D})$ also. Thus,

$$A \in \text{ex}(\mathcal{D}) \Rightarrow P_1AP_2 \in \text{ex}(\mathcal{D}),$$

and

$$P_1AP_2 \in \text{ex}(\mathcal{D}) \Rightarrow P_1^TP_1AP_2P_2^T = A \in \text{ex}(\mathcal{D}),$$

hence $A \in \text{ex}(\mathcal{D})$ if and only if $P_1AP_2 \in \text{ex}(\mathcal{D})$ for any permutation matrices P_1 and P_2 . \square

3.1. Tight constraints

We now examine the implications of Theorem 1 for the extreme points of \mathcal{D} . We first note a simple fact concerning the number of linearly independent differential privacy constraints that can be tight on an element of \mathcal{D} .

In the next result, we use \mathcal{C}_j^{dp} to denote the set of all tight differential privacy constraints acting on the j th column of a matrix A . Formally, given A , this consists of all constraints such that $a_{ij} - e^\epsilon a_{kj} = 0$ where $1 \leq i, k \leq n$.

Theorem 4 *Let $A \in \mathcal{D}$ be given. Then, $\dim(\text{span}(\mathcal{C}_j^{dp})) = n$ if and only if $a_{ij} = 0$ for each $i \in \{1, \dots, n\}$.*

Proof: If we make the obvious identification of the j th column of A with a column vector, $A^{(j)}$ in \mathbb{R}^n , then each constraint in \mathcal{C}_j^{dp} can be identified with a vector of the form $(0, \dots, 1, 0, \dots, -e^\epsilon, 0, \dots, 0)^T$ where the 1 occurs in the i th position and e^ϵ occurs in the k th position. If $\dim(\text{span}(\mathcal{C}_j^{dp})) = n$, there are n linearly independent vectors v_1, \dots, v_n such that $v_i^T A^{(j)} = 0$ for $1 \leq i \leq n$ so it follows trivially that $A^{(j)} = 0$.

For the converse, it is enough to note that $A^{(j)} = 0$ implies that every differential privacy constraint acting on the j th column is tight and that there are n linearly independent such constraints. To see this consider the matrix T with: $t_{ii} = 1$ for $1 \leq i \leq n$; $t_{i+1,i} = -e^\epsilon$ for $1 \leq i < n$; $t_{1n} = -e^\epsilon$; $t_{jk} = 0$ otherwise. It can readily be verified that T is non-singular. \square

Remark: A direct consequence of Theorem 4 is that $\dim(\text{span}(\mathcal{C}_j^{dp})) \leq n-1$ for any $j \in \gamma(A)$.

Our later characterisations of the extreme points of \mathcal{D} shall rely on the following concept of *loose entries*.

Definition 5 (Loose entries of a matrix) Given $A \in \mathcal{D}$, define

$$\lambda(A) = \left\{ (i, j) : a_{ij} \notin \left\{ e^\epsilon \min_k a_{kj}, e^{-\epsilon} \max_k a_{kj} \right\} \right\}.$$

For a matrix $A \in \mathcal{D}$, we say the entry a_{ij} is **loose** if $(i, j) \in \lambda(A)$.

It follows from Lemma 2 that for any (i, j) there exists a k such that $a_{ij} = e^{\pm\epsilon} a_{kj}$ if and only if $(i, j) \notin \lambda(A)$.

Example 1 Let $\epsilon = \ln(2)$ and

$$A = \frac{1}{7} \begin{pmatrix} 4 & 1 & 2 \\ 3 & 2 & 2 \\ 2 & 1 & 4 \end{pmatrix}.$$

Then $\lambda(A) = \{(2, 1)\}$, since $3 \notin \{4, 2\}$.

Our next result bounds the number of loose entries of an extreme point in terms of the number of non-zero columns.

Theorem 5 Let $A \in \text{ex}(\mathcal{D})$ with $|\gamma(A)| \geq 2$. Then,

$$|\lambda(A)| \leq n - |\gamma(A)|.$$

Proof: Let $A \in \text{ex}(\mathcal{D})$ and consider the following sets of constraints active on A . We define

$$\mathcal{C}^{dp} = \bigcup_{j \in \gamma(A)} \mathcal{C}_j^{dp}$$

to be the set of tight differential privacy constraints acting on the columns in $\gamma(A)$. Note the following readily verifiable facts:

- (i) for $j \notin \gamma(A)$, every differential privacy constraint acting on column j is tight;
- (ii) the n stochastic constraints are tight;
- (iii) as $|\gamma(A)| \geq 2$, no non-zero column of A is of the form $k\mathbf{1}$ where $k \in \mathbb{R}$.

It follows from (ii) and Theorem 1 that the number of tight, linearly independent differential privacy constraints on A must be $n^2 - n$. Furthermore, Theorem 4 implies that there are n linearly independent differential privacy constraints active on each of the $n - |\gamma(A)|$ zero columns of A . It is not difficult to see that constraints acting on different columns must be linearly independent and hence

there are a total of $(n - |\gamma(A)|)n$ linearly independent tight differentially private constraints arising from the zero columns of A . Putting all of this together, we see that there must be

$$n^2 - n - (n - |\gamma(A)|)n = n|\gamma(A)| - n$$

tight differential privacy constraints acting on the non-zero columns of A . Formally:

$$|\mathcal{C}^{dp}| \geq n|\gamma(A)| - n. \quad (7)$$

From point (iii) above there are no non-zero columns in which all entries are constant; it follows that for each $j \in \gamma(A)$,

$$|\{i : (i, j) \notin \lambda(A)\}| \geq |\mathcal{C}_j^{dp}| + 1.$$

If we let l_j denote the number of loose entries in column j , the previous inequality can be rewritten as

$$|\mathcal{C}_j^{dp}| \leq n - l_j - 1.$$

Combining this with (7) we see that

$$\begin{aligned} n|\gamma(A)| - n &\leq \sum_{j \in \gamma(A)} |\mathcal{C}_j^{dp}| \\ &\leq \sum_{j \in \gamma(A)} n - l_j - 1 \\ &= n|\gamma(A)| - |\lambda(A)| - |\gamma(A)|. \end{aligned}$$

A simple rearrangement now shows that

$$|\lambda(A)| \leq n - |\gamma(A)|$$

as claimed. \square

Note: When $|\gamma(A)| = 1$, $|\lambda(A)| = n$.

To conclude this sub-section, we take a look at the following result for later use, which states that at most one loose entry can appear in any row of an extreme point.

Lemma 3 *Let $A \in \text{ex}(\mathcal{D})$. No row of A has more than one loose entry (i.e. there exist no two distinct pairs $(i_1, j_1), (i_1, j_2) \in \lambda(A)$ with $j_1 \neq j_2$).*

Proof: Let $A \in \text{ex}(\mathcal{D})$, and assume without loss of generality that $(1, 1), (1, 2) \in \lambda(A)$. Let

$$\Delta = \min \left\{ \max_i a_{i1} - a_{11}, a_{11} - \min_i a_{i1}, \max_i a_{i2} - a_{12}, a_{12} - \min_i a_{i2} \right\}.$$

Hence, $A \pm \Delta(E_{11} - E_{12}) \in \mathcal{D}$.

However, $A = \frac{1}{2}((A + \Delta E_{11} - \Delta E_{12}) + (A - \Delta E_{11} + \Delta E_{12}))$, hence, $A \notin \text{ex}(\mathcal{D})$, a contradiction and so the result follows. \square

Finally, for this section we present a number of other results that will add further insight to the behaviour and structure of \mathcal{D} and its extreme points. The next piece of notation will prove useful later.

For $A \in \mathcal{D}$, we define the vector $m' \in \mathbb{R}^n$ where $m'_j = \frac{1}{\min_i a_{ij}}$ for any $j \in \gamma(A)$ and $m'_j = 0$ otherwise. We then denote by \tilde{A} the matrix given by:

$$\tilde{A} = A \operatorname{diag}(m'). \quad (8)$$

Then, for any $A \in \mathcal{D}$, $\tilde{a}_{ij} \in [1, e^\epsilon]$ for any $j \in \gamma(A)$, and $\tilde{a}_{ij} = 0$ otherwise. Hence,

$$\tilde{A} \operatorname{diag} \left(\min_{1 \leq j \leq n} a_{ij} \right) = A.$$

Note: $\gamma(A) = \gamma(\tilde{A})$ and $\lambda(A) = \lambda(\tilde{A})$.

We now show that for any extreme point A , \tilde{A} cannot have a row with equal non-zero values.

Lemma 4 *Let $A \in \operatorname{ex}(\mathcal{D})$ with $|\gamma(A)| > 1$. Then for each row i , there exist two non-zero columns $j, k \in \gamma(A)$ such that $\tilde{a}_{ij} \neq \tilde{a}_{ik}$.*

Proof: We prove this by contradiction. Firstly, suppose there exists a row i such that $\tilde{a}_{ij} = \tilde{a}_{ik}$ for all $j, k \in \gamma(A)$. By Lemma 3, each row cannot have more than one loose element, therefore either $\tilde{a}_{ij} = 1$ or $\tilde{a}_{ij} = e^\epsilon$.

Let $m \in \mathbb{R}^n$ be defined by $m_j = \min_i a_{ij}$. Then $A = \tilde{A} \operatorname{diag}(m)$.

Suppose $\tilde{a}_{ij} = 1$, hence $\sum_{k \in \gamma(A)} m_k = 1$. By Theorem 5, each column j has at least one pair (i, k) such that $a_{ij} = e^\epsilon a_{kj}$, hence there exists a row i^* such that $\tilde{a}_{i^*j} = e^\epsilon$. However, $\tilde{a}_{i^*k} \geq 1$ for every $k \in \gamma(A)$, so $\sum_{k \in \gamma(A)} \tilde{a}_{i^*k} m_k > 1$, contradicting the stochasticity of A .

A similar argument holds for $\tilde{a}_{ij} = e^\epsilon$. The result follows. \square

4. Extreme points for fixed values of $|\gamma(A)|$

In this section, we characterise extreme points with a specified number of non-zero columns. We note that extreme points with one and two non-zero columns are limited to a specific form, while Section 4.3 deals with extreme points with any number of non-zero columns.

4.1. Extreme points with one column non-zero

The first case to consider is that of a single non-zero column in the matrix. Due to the stochastic constraints, there are only n such matrices, and as Theorem 6 below states, each one of these matrices is an extreme point.

Theorem 6 ($|\gamma(A)| = 1$) *Let $E_i \in \mathbb{R}^{n \times n}$ be given by $E_i = \mathbf{1}e_i^T$ for $1 \leq i \leq n$ and define the set $\tilde{\mathcal{D}}'$ as:*

$$\tilde{\mathcal{D}}' = \{E_1, \dots, E_n\}.$$

Then $\tilde{\mathcal{D}}' \subseteq \operatorname{ex}(\mathcal{D})$.

Furthermore, $A \in \operatorname{ex}(\mathcal{D})$, $|\gamma(A)| = 1$ implies that $A \in \tilde{\mathcal{D}}'$.

Proof: Suppose $E_i = \frac{1}{2}(B + C)$ for B, C in \mathcal{D} . As B, C are both nonnegative, it follows immediately that all columns of B and C apart from the i th column are zero. B and C are also both stochastic which immediately implies that $B = C = \mathbf{1}e_i^T$.

Note that if $A \in \mathcal{D}$ with $|\gamma(A)| = 1$, then $A = E_i$ for some i . Hence, if $A \in \text{ex}(\mathcal{D})$ with $|\gamma(A)| = 1$, it follows that $A \in \tilde{\mathcal{D}}'$. \square

The points in $\tilde{\mathcal{D}}'$ are extreme points in all cases, regardless of ϵ . Furthermore, in the trivial case of $\epsilon = 0$, the set $\tilde{\mathcal{D}}'$ are the only extreme points.

Corollary 2 *Let $\epsilon = 0$. Then,*

$$\text{ex}(\mathcal{D}) = \tilde{\mathcal{D}}'.$$

Proof: Let $\epsilon = 0$. Then, for all $A \in \mathcal{D}$, we have $a_{kj} \leq a_{ij} \leq a_{kj}$, hence $a_{ij} = a_{kj}$ for all i, j, k , i.e. entries in the same column are equal. It now follows immediately from Corollary 1 that if A is an extreme point, $|\gamma(A)| = 1$ and hence that $A \in \tilde{\mathcal{D}}'$ as claimed. \square

4.2. Extreme points with two columns non-zero

Next, we consider the case of two non-zero columns. Although Theorem 5 allows for many loose entries to occur in these extreme points, Theorem 7 below states that no loose entries are possible.

Theorem 7 ($|\gamma(A)| = 2$) *Let $A \in \text{ex}(\mathcal{D})$ where $|\gamma(A)| = 2$. Then A has no loose entries.*

Proof: Without loss of generality, assume that $\gamma(A) = \{1, 2\}$. Define $m \in \mathbb{R}^n$ by $m_j = \min_i a_{ij}$ for $1 \leq j \leq n$ and define \tilde{A} so that $A = \tilde{A} \text{diag}(m)$. Then $\tilde{a}_{ij} \in [1, e^\epsilon] \cup \{0\}$ for $1 \leq i, j \leq n$.

By Theorem 5, $|\lambda(A)| \leq n - 2$, so there exist at least two rows with no loose entries. Let row k be one of these rows. Then $\tilde{a}_{k1}, \tilde{a}_{k2} \in \{1, e^\epsilon\}$, but by Lemma 4, $\tilde{a}_{k1} \neq \tilde{a}_{k2}$. We can assume that $\tilde{a}_{k1} = e^\epsilon$ and $\tilde{a}_{k2} = 1$ (otherwise just swap columns 1 and 2). As A is stochastic,

$$m_1 e^\epsilon + m_2 = 1. \tag{9a}$$

By Lemma 3, for all rows j , at least one of $\tilde{a}_{j1}, \tilde{a}_{j2}$ must be in $\{1, e^\epsilon\}$. Moreover, in order to satisfy (9a), $\tilde{a}_{j1} = e^\epsilon$ if and only if $\tilde{a}_{j2} = 1$.

Suppose therefore that there exists a row j where $\tilde{a}_{j1} \in (1, e^\epsilon)$ corresponding to a loose entry in A . It follows from (9a) that $\tilde{a}_{j2} = e^\epsilon$. Hence

$$\begin{aligned} 1 &= m_1 \tilde{a}_{j1} + m_2 e^\epsilon \\ &> m_1 + m_2 e^\epsilon. \end{aligned} \tag{9b}$$

It follows from Corollary 1 that there is some j^* such that $\tilde{a}_{j^*1} = 1$, implying

$$\begin{aligned} 1 &= m_1 + m_2 \tilde{a}_{j^*2} \\ &\leq m_1 + m_2 e^\epsilon, \end{aligned}$$

contradicting (9b). Therefore there are no loose entries in the first column.

Now suppose there exists a row j where $\tilde{a}_{j2} \in (1, e^\epsilon)$. As above, it follows that $\tilde{a}_{j1} = 1$. Hence,

$$\begin{aligned} 1 &= m_1 + m_2 \tilde{a}_{j2} \\ &< m_1 + m_2 e^\epsilon. \end{aligned} \tag{9c}$$

As before, it follows from Corollary 1 that there is some j^* such that $\tilde{a}_{j^*2} = e^\epsilon$, hence,

$$\begin{aligned} 1 &= m_1 \tilde{a}_{j^*1} + m_2 e^\epsilon \\ &\geq m_1 + m_2 e^\epsilon, \end{aligned}$$

contradicting (9c). Therefore there are no loose entries in the second column.

Hence $|\lambda(A)| = 0$. \square

Using this result along with Lemma 4, we can describe the two non-zero columns.

Corollary 3 *Let $A \in \text{ex}(\mathcal{D})$ with $|\gamma(A)| = 2$. Let $\gamma(A) = \{j, k\}$ and \tilde{A} be given by (8). Then, for every $1 \leq i \leq n$, we have*

$$(\tilde{a}_{ij}, \tilde{a}_{ik}) \in \{(1, e^\epsilon), (e^\epsilon, 1)\}. \tag{10}$$

Proof: By Theorem 7, $\tilde{a}_{ij} \in \{1, e^\epsilon\}$ and $\tilde{a}_{ik} \in \{1, e^\epsilon\}$ for each $1 \leq i \leq n$.

By Lemma 4, we must have $\tilde{a}_{ij} \neq \tilde{a}_{ik}$ for each i . So, $\tilde{a}_{ij} = e^\epsilon$ if and only if $\tilde{a}_{ik} = 1$, and $\tilde{a}_{ij} = 1$ if and only if $\tilde{a}_{ik} = e^\epsilon$. \square

The follow example illustrates the consequence of Corollary 3.

Example 2 *Every extreme point $A \in \text{ex}(\mathcal{D})$ with $|\gamma(A)| = 2$ must be of the form shown in (10), and furthermore both non-zero columns of \tilde{A} must contain at least one 1 and one e^ϵ .*

Let $n = 4$ and $A \in \text{ex}(\mathcal{D})$ with $|\gamma(A)| = 2$. One example of such an A is as follows:

$$A = \frac{1}{1 + e^\epsilon} \begin{pmatrix} 1 & 0 & e^\epsilon & 0 \\ 1 & 0 & e^\epsilon & 0 \\ e^\epsilon & 0 & 1 & 0 \\ 1 & 0 & e^\epsilon & 0 \end{pmatrix} \in \text{ex}(\mathcal{D}).$$

4.3. Extreme points with every element constrained

The next definition is necessary before we can state Theorem 8 which is the main result of the paper.

Definition 6 *Let $\tilde{\mathcal{D}} \subset \mathcal{D}$ be defined as follows:*

$$\tilde{\mathcal{D}} = \{A \in \mathcal{D} \mid \text{rank}(A) = |\gamma(A)|, \lambda(A) = \emptyset\}. \tag{11}$$

The set $\tilde{\mathcal{D}}$ contains matrices with between 2 and n non-zero columns, which satisfy the rank condition of Theorem 2 and have no loose entries (i.e. $\tilde{a}_{ij} \in \{0, 1, e^\epsilon\}$ for each i, j). We now show that every one of these matrices is an extreme point of \mathcal{D} .

Theorem 8 *Let $\epsilon > 0$. Then,*

$$\tilde{\mathcal{D}} \subset \text{ex}(\mathcal{D}).$$

Proof: Let $A \in \tilde{\mathcal{D}}$ and let $B, C \in \mathcal{D}$ where $\frac{1}{2}(B + C) = A$. Define $m_j = \min_i a_{ij}$ for each $j \in \{1, \dots, n\}$ (Note that $m_j = 0$ for each $j \notin \gamma(A)$, and $a_{ij} \in \{m_j, e^\epsilon m_j\}$ for each i, j since $\lambda(A) = \emptyset$).

Let $\Delta_j = \frac{1}{2} \max_i |b_{ij} - c_{ij}|$ for each $j \in \gamma(A)$. As B and C are nonnegative, it is not hard to see that:

$$\Delta_j = 0, \quad \forall j \notin \gamma(A). \quad (12a)$$

We shall show that the same conclusion must also hold for $j \in \gamma(A)$. To this end, let $j^* \in \gamma(A)$ be given where $\Delta_{j^*} > 0$. Assume without loss of generality that $b_{i_1 j^*} = a_{i_1 j^*} + \Delta_{j^*}$ for some i_1 (if not, swap B and C).

We claim that $a_{i_1 j^*} \neq m_{j^*}$. Suppose otherwise. Then there exists i_2 where $a_{i_2 j^*} = e^\epsilon m_{j^*}$. However, since $\frac{1}{2}(B + C) = A$, we have $c_{i_1 j^*} = 2a_{i_1 j^*} - b_{i_1 j^*} = a_{i_1 j^*} - \Delta_{j^*}$, and since $C \in \mathcal{D}$, we have

$$\begin{aligned} c_{i_2 j^*} &\leq e^\epsilon c_{i_1 j^*} \\ &= e^\epsilon a_{i_1 j^*} - e^\epsilon \Delta_{j^*} \\ &= a_{i_2 j^*} - e^\epsilon \Delta_{j^*} \end{aligned}$$

By the definition of Δ_{j^*} , we must have $c_{i_2 j^*} \geq a_{i_2 j^*} - \Delta_{j^*}$. Hence it would follow that $\Delta_{j^*} \geq e^\epsilon \Delta_{j^*}$, a contradiction since $\epsilon > 0$. Thus, $a_{i_1 j^*} = e^\epsilon m_{j^*}$ as claimed (i.e. the max change occurs on the max element of the column).

We now know that $b_{i_1 j^*} = e^\epsilon m_{j^*} + \Delta_{j^*}$. Let

$$I_{j^*} = \{i : a_{ij^*} = m_{j^*}\}.$$

Then for every $i \in I_{j^*}$, since $B \in \mathcal{D}$, we get $e^\epsilon m_{j^*} + \Delta_{j^*} = b_{i_1 j^*} \leq e^\epsilon b_{ij^*}$, hence

$$b_{ij^*} \geq m_{j^*} + e^{-\epsilon} \Delta_{j^*}. \quad (12b)$$

Also, for every $i \in I_{j^*}$, since $C \in \mathcal{D}$,

$$\begin{aligned} c_{i_1 j^*} &= e^\epsilon m_{j^*} - \Delta_{j^*} \\ &\leq e^\epsilon c_{ij^*} \\ &= e^\epsilon (2a_{ij^*} - b_{ij^*}) \\ &= 2e^\epsilon m_{j^*} - e^\epsilon b_{ij^*}, \end{aligned}$$

hence $e^\epsilon m_{j^*} - \Delta_{j^*} \leq 2e^\epsilon m_{j^*} - e^\epsilon b_{ij^*}$, or rewriting,

$$b_{ij^*} \leq m_{j^*} + e^{-\epsilon} \Delta_{j^*}. \quad (12c)$$

Hence, from (12b) and (12c),

$$\begin{aligned} b_{ij^*} &= m_{j^*} + e^{-\epsilon} \Delta_{j^*} \\ &= a_{ij^*} + e^{-\epsilon} \Delta_{j^*}, \end{aligned} \quad (12d)$$

for every $i \in I_{j^*}$.

It follows readily that for every $i \in I_{j^*}$, $c_{ij^*} = m_{j^*} - e^{-\epsilon} \Delta_{j^*}$.

We next consider indices $i \notin I_{j^*}$. Choose some $i_2 \in I_{j^*}$. For all $i \notin I_{j^*}$, $a_{ij^*} = e^\epsilon m_{j^*}$, then

$$b_{ij^*} \leq e^\epsilon b_{i_2 j^*} = e^\epsilon m_{j^*} + \Delta_{j^*}, \quad (12e)$$

and

$$\begin{aligned} c_{ij^*} &= 2a_{ij^*} - b_{ij^*} \\ &= 2e^\epsilon m_{j^*} - b_{ij^*} \\ &\leq e^\epsilon c_{i_2 j^*} \\ &= e^\epsilon m_{j^*} - \Delta_{j^*}, \end{aligned}$$

which can be rewritten as

$$b_{ij^*} \geq e^\epsilon m_{j^*} + \Delta_{j^*}. \quad (12f)$$

Hence, from (12e) and (12f),

$$\begin{aligned} b_{ij^*} &= e^\epsilon m_{j^*} + \Delta_{j^*} \\ &= a_{ij^*} + \Delta_{j^*}, \end{aligned} \quad (12g)$$

for all $i \notin I_{j^*}$.

Putting everything together, it follows from (12a), (12d) and (12g),

$$b_{ij} = \begin{cases} 0, & j \notin \gamma(A), \\ m_j + e^{-\epsilon} g_j \Delta_j, & j \in \gamma(A), i \in I_j \\ e^\epsilon m_j + g_j \Delta_j, & j \in \gamma(A), i \notin I_j \end{cases}$$

where $g_j \in \{-1, 1\}$, for all $j \in \gamma(A)$.

Similarly, since $B + C = 2A$,

$$c_{ij} = \begin{cases} 0, & j \notin \gamma(A), \\ m_j - e^{-\epsilon} g_j \Delta_j, & j \in \gamma(A), i \in I_j \\ e^\epsilon m_j - g_j \Delta_j, & j \in \gamma(A), i \notin I_j. \end{cases}$$

Rewriting in terms of \tilde{A} (given by (8)), $b_{ij} = a_{ij} + g_j e^{-\epsilon} \frac{a_{ij}}{m_j} \Delta_j = a_{ij} + g_j e^{-\epsilon} \tilde{a}_{ij} \Delta_j$ and $c_{ij} = a_{ij} - g_j e^{-\epsilon} \tilde{a}_{ij} \Delta_j$ for all i, j .

Hence,

$$\begin{aligned} B &= A + e^{-\epsilon} \tilde{A} \operatorname{diag} (g_j \Delta_j)_{1 \leq j \leq n} \\ C &= A - e^{-\epsilon} \tilde{A} \operatorname{diag} (g_j \Delta_j)_{1 \leq j \leq n}. \end{aligned}$$

Since A, B are stochastic, we require

$$e^{-\epsilon} \tilde{A} \operatorname{diag} (g_j \Delta_j) \mathbf{1} = 0.$$

This equation defines a linear relationship between the columns of \tilde{A} . Moreover, we know that $\Delta_j = 0$ for $j \notin \gamma(A)$. If $\Delta_j \neq 0$ for any $j \in \gamma(A)$, it would imply that the non-zero columns of \tilde{A} and hence those of A are linearly dependent, contradicting the assumption that $\operatorname{rank}(A) = |\gamma(A)|$. It follows that $\Delta_j = 0$ for all j and hence that $B = C = A$. This completes the proof. \square

Furthermore, the set $\tilde{\mathcal{D}}$ contains all extreme points of \mathcal{D} which have no loose entries.

Corollary 4 *Let $A \in \mathcal{D}$ with $\lambda(A) = \emptyset$. Then, $A \in \operatorname{ex}(\mathcal{D})$ if and only if $A \in \tilde{\mathcal{D}}$.*

Proof: “ \Rightarrow ”: Let $A \in \operatorname{ex}(\mathcal{D})$ with $\lambda(A) = \emptyset$. By Theorem 2, $\operatorname{rank}(A) = |\gamma(A)|$, hence $A \in \tilde{\mathcal{D}}$.

“ \Leftarrow ”: $A \in \tilde{\mathcal{D}} \Rightarrow A \in \operatorname{ex}(\mathcal{D})$ by Theorem 8. \square

4.4. Extreme points with all columns non-zero

From an application point of view, it is entirely reasonable to only consider matrices (and the resulting response mechanism) with no zero columns.

Having a zero column in a matrix that defines a response mechanism means that the mechanism never releases a particular (or multiple) values as its output. In many circumstances, this feature will not be required of a mechanism.

Using Theorem 8, we now present the following corollary, which gives a complete characterisation of extreme points without zero columns.

Corollary 5 *Let $A \in \mathcal{D}$, with $|\gamma(A)| = n$. Then, $A \in \operatorname{ex}(\mathcal{D})$ if and only if $A \in \tilde{\mathcal{D}}$*

Equivalently,

$$\{A \in \operatorname{ex}(\mathcal{D}) : |\gamma(A)| = n\} = \{A \in \tilde{\mathcal{D}} : |\gamma(A)| = n\}.$$

Proof: “ \Rightarrow ”: Let $A \in \operatorname{ex}(\mathcal{D})$ have n non-zero columns. Then, $\operatorname{rank}(A) = n$ by Theorem 2 and $\lambda(A) = \emptyset$ by Theorem 5.

“ \Leftarrow ”: Let $A \in \mathcal{D}$ such that $\operatorname{rank}(A) = n$ and $\lambda(A) = \emptyset$. Then $A \in \operatorname{ex}(\mathcal{D})$ by Theorem 8. \square

We now have necessary and sufficient conditions for finding and determining extreme points with n non-zero columns.

5. Discussion

We now take a brief look at a number of useful and interesting consequences of the results given in Sections 3 and 4.

$\text{ex}(\mathcal{D})$ **for small n :** From Theorems 6 and 8, we know that $\tilde{\mathcal{D}}' \cup \tilde{\mathcal{D}} \subseteq \text{ex}(\mathcal{D})$; with the addition of Theorem 7 we can make further observations for small n .

Theorem 9 *Let $n \leq 3$, then*

$$\text{ex}(\mathcal{D}) = \tilde{\mathcal{D}}' \cup \tilde{\mathcal{D}}.$$

Extreme points for $n = 4$: We therefore have a complete characterisation of all extreme points up to $n = 3$. While we lack a formal proof, extensive computer simulations suggest it is also true for $n = 4$ leading to the following conjecture.

Let $n \leq 4$: then

$$\text{ex}(\mathcal{D}) = \tilde{\mathcal{D}}' \cup \tilde{\mathcal{D}}.$$

$\text{ex}(\mathcal{D})$ **for $n \geq 5$:** When $n = 5$, our previous results allow us to characterise all extreme points A for which $|\gamma(A)| = 1, 2, 5$. However, when $|\gamma(A)| = 4$, we can find extreme points with loose entries.

The following point $A \in \mathcal{D}$ can be shown to be an extreme point of \mathcal{D} by using Theorem 1.

$$A = \frac{1}{3 + 2e^\epsilon} \begin{pmatrix} 1 & 1 & 2e^\epsilon & 1 & 0 \\ e^\epsilon & 1 & 2 & e^\epsilon & 0 \\ e^\epsilon & e^\epsilon & 2 & 1 & 0 \\ 1 & e^\epsilon & 2 & e^\epsilon & 0 \\ 1 & 1 & 1 + e^\epsilon & e^\epsilon & 0 \end{pmatrix}.$$

Fitting with Theorem 5, A has only a single loose entry ($\lambda(A) = \{(5, 3)\}$), while we also observe that $\text{rank}(A) = 4$, satisfying Theorem 2.

We therefore have $A \in \text{ex}(\mathcal{D})$, but $A \notin \tilde{\mathcal{D}}' \cup \tilde{\mathcal{D}}$. Hence, $\tilde{\mathcal{D}}' \cup \tilde{\mathcal{D}} \subset \text{ex}(\mathcal{D})$ in general.

6. Conclusion

We have studied the differential privacy polytope of $n \times n$ matrices and described a suite of results characterising its extreme points. In particular, our results describe completely the extreme points of this polytope containing 1, 2 and n non-zero columns. The last fact is of particular practical significance as most implementations of differentially private mechanisms are likely to have no zero columns; this is because a zero column corresponds to a value of the dataset D that is never released by the mechanism. Future work could focus on characterising extreme points with other values of $|\gamma(A)|$; alternative directions for work include considering other convex geometric aspects of the polytope \mathcal{D} such as the structure of its dual set for example.

Acknowledgment

The first named author was supported by the Science Foundation Ireland grant SFI/11/PI/1177.

References

- [1] ADAM, N. R., AND WORTHMANN, J. C. Security-control methods for statistical databases: a comparative study. *ACM Computing Surveys (CSUR)* 21, 4 (1989), 515–556.
- [2] AVIS, D., AND FUKUDA, K. A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete & Computational Geometry* 8, 1 (1992), 295–313.
- [3] BARVINOK, A. *A course in convexity*, vol. 54. American Mathematical Society Providence, 2002.
- [4] BENGTTSSON, I., ERICSSON, Å., KUŚ, M., TADEJ, W., AND ŻYCZKOWSKI, K. Birkhoffs polytope and unistochastic matrices, $n=3$ and $n=4$. *Communications in mathematical physics* 259, 2 (2005), 307–324.
- [5] BRUALDI, R. A. An interesting face of the polytope of doubly stochastic matrices. *Linear and Multilinear Algebra* 17, 1 (1985), 5–18.
- [6] BRUALDI, R. A., AND LIU, B. The polytope of even doubly stochastic matrices. *Journal of Combinatorial Theory, Series A* 57, 2 (1991), 243–253.
- [7] BUDNITZ, M. E. Privacy protection for consumer transactions in electronic commerce: why self-regulation is inadequate. *SCL Rev.* 49 (1997), 847.
- [8] CHAN, C. S., AND ROBBINS, D. P. On the volume of the polytope of doubly stochastic matrices. *Experimental Mathematics* 8, 3 (1999), 291–300.
- [9] DIFFIE, W., AND HELLMAN, M. E. Privacy and authentication: An introduction to cryptography. *Proceedings of the IEEE* 67, 3 (1979), 397–427.
- [10] DOMINGO-FERRER, J., AND MATEO-SANZ, J. M. Practical data-oriented microaggregation for statistical disclosure control. *Knowledge and Data Engineering, IEEE Transactions on* 14, 1 (2002), 189–201.
- [11] DUCHI, J. C., JORDAN, M., WAINWRIGHT, M. J., ET AL. Local privacy and statistical minimax rates. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on* (2013), IEEE, pp. 429–438.
- [12] DWORK, C. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.

- [13] DWORK, C. Differential privacy. In *Encyclopedia of Cryptography and Security*. Springer, 2011, pp. 338–340.
- [14] DYER, M. E. The complexity of vertex enumeration methods. *Mathematics of Operations Research* 8, 3 (1983), 381–402.
- [15] GANTA, S. R., KASIVISWANATHAN, S. P., AND SMITH, A. Composition attacks and auxiliary information in data privacy. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining* (2008), ACM, pp. 265–273.
- [16] HARDT, M., AND TALWAR, K. On the geometry of differential privacy. In *Proceedings of the forty-second ACM symposium on Theory of computing* (2010), ACM, pp. 705–714.
- [17] HOLOHAN, N., LEITH, D. J., AND MASON, O. Differential privacy in metric spaces: Numerical, categorical and functional data under the one roof. *Information Sciences* 305 (2015), 256–268.
- [18] KAIROUZ, P., OH, S., AND VISWANATH, P. Extremal mechanisms for local differential privacy. *CoRR abs/1407.1338* (2014).
- [19] LI, C., HAY, M., RASTOGI, V., MIKLAU, G., AND MCGREGOR, A. Optimizing linear counting queries under differential privacy. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems* (2010), ACM, pp. 123–134.
- [20] MCSHERRY, F., AND TALWAR, K. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on* (2007), IEEE, pp. 94–103.
- [21] NISSIM, K., SMORODINSKY, R., AND TENNENHOLTZ, M. Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd innovations in theoretical computer science conference* (2012), ACM, pp. 203–213.
- [22] WARNER, S. L. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association* 60, 309 (1965), 63–69.