Gibbs Manifolds

Bernd Sturmfels
MPI Leipzig

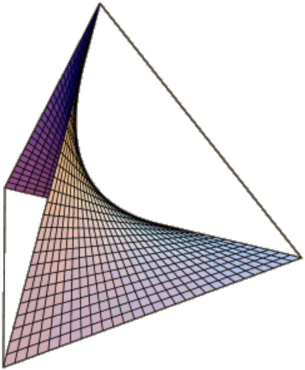

Paper with Dmitrii Pavlov and Simon Telen

Back-to-the-Roots Seminar

KU Leuven, September 11, 2023

The World is Toric

Toric varieties provide the geometric foundations for many successes in the mathematical sciences. In statistics they appear as discrete exponential families. In optimization, they furnish nonnegativity certificates and they govern entropic regularization of linear programming.

The World is Toric

Toric varieties provide the geometric foundations for many successes in the mathematical sciences. In statistics they appear as discrete exponential families. In optimization, they furnish nonnegativity certificates and they govern entropic regularization of linear programming. Notable sightings in phylogenetics, stochastic analysis, Gaussian inference and chemical reaction networks led to the slogan **The World is Toric**.

Section 8.3 in

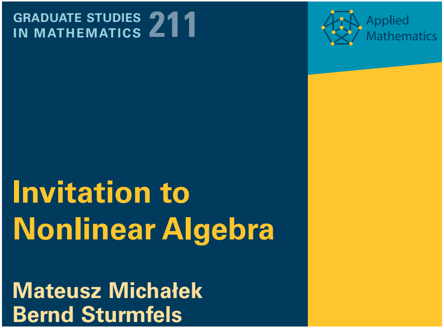

Being Positive

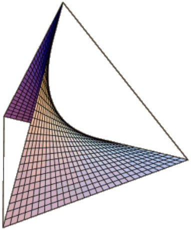

Example (Two binary random variables)

$$P = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} = \begin{bmatrix} st & s(1-t) \\ (1-s)t & (1-s)(1-t) \end{bmatrix}$$

Among all nonnegative matrices P with fixed row and column sums, a unique matrix satisfies $\det(P) = p_{00}p_{11} - p_{01}p_{10} = 0$.

LP → Optimal Transport

Being Positive

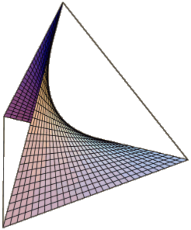

Example (Two binary random variables)

$$P = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} = \begin{bmatrix} st & s(1-t) \\ (1-s)t & (1-s)(1-t) \end{bmatrix}$$

Among all nonnegative matrices P with fixed row and column sums, a unique matrix satisfies $\det(P) = p_{00}p_{11} - p_{01}p_{10} = 0$.

LP → Optimal Transport

In applications, the key player is the positive part of the toric variety. That manifold is identified with a convex polytope by the moment map. The fibers of the underlying linear map are polytopes of complementary dimension, and each fiber intersects the toric variety in a unique point. This is the unique maximizer of the entropy over the fiber. In statistical physics and computer science, this is known as the **Gibbs distribution**.

Abstract

Gibbs manifolds are images of linear spaces of symmetric matrices under the exponential map. Arising in applications like optimization, statistics and quantum physics, they extend the *ubiquitous role of toric geometry*.

$$\left\{ \begin{bmatrix} a_0+b_0 & 0 & 0 & 0 \\ 0 & a_0+b_1 & 0 & 0 \\ 0 & 0 & a_1+b_0 & 0 \\ 0 & 0 & 0 & a_1+b_1 \end{bmatrix} \right\}$$

$$p_{00} = \exp(a_0+b_0),\; p_{01} = \exp(a_0+b_1),$$
$$p_{10} = \exp(a_1+b_0),\; p_{11} = \exp(a_1+b_1).$$

$$p_{00}p_{11} - p_{01}p_{10} \;=\; 0.$$

The Gibbs variety is the zero locus of all polynomials that vanish on the Gibbs manifold. We compute these polynomials and show that the Gibbs variety is low-dimensional. Our theory is applied to a wide range of scenarios, including matrix pencils and quantum optimal transport.

Exponentials and Logarithms

The space $\mathbb{S}^n$ of symmetric $n \times n$-matrices has dimension $\binom{n+1}{2}$.
The cone of positive semidefinite (PSD) matrices is denoted $\mathbb{S}^n_+$.

The PSD cone $\mathbb{S}^n_+$ is self-dual under the inner product

$$\langle X, Y \rangle \;:=\; \mathrm{trace}(XY) \quad \text{for} \quad X, Y \in \mathbb{S}^n.$$

Exponentials and Logarithms

The space $\mathbb{S}^n$ of symmetric $n \times n$-matrices has dimension $\binom{n+1}{2}$.
The cone of positive semidefinite (PSD) matrices is denoted $\mathbb{S}^n_+$.

The PSD cone $\mathbb{S}^n_+$ is self-dual under the inner product

$$\langle X, Y \rangle := \mathrm{trace}(XY) \quad \text{for} \quad X, Y \in \mathbb{S}^n.$$

The exponential function maps symmetric matrices to positive definite symmetric matrices:

$$\exp \; : \mathbb{S}^n \to \mathrm{int}(\mathbb{S}^n_+)\,,\; X \mapsto \sum_{i=0}^{\infty} \frac{1}{i!} X^i.$$

This map is invertible, with inverse given by the logarithm:

$$\log \; : \mathrm{int}(\mathbb{S}^n_+) \to \mathbb{S}^n\,,\; Y \mapsto \sum_{j=1}^{\infty} \frac{(-1)^{j-1}}{j} (Y - \mathrm{id}_n)^j.$$

Players

Fix $A_0, A_1, A_2, \ldots, A_d \in \mathbb{S}^n$, where the last d are linearly independent. They determine a d-dim'l affine space of symmetric matrices (ASSM):

$$\mathcal{L} \; = \; A_0 + \mathrm{span}_{\mathbb{R}}(A_1, A_2, \ldots, A_d) \;\; \subset \;\; \mathbb{S}^n \simeq \mathbb{R}^{\binom{n+1}{2}}.$$

If $A_0 = 0$, then $\mathcal{L}$ is a *linear space of symmetric matrices* (LSSM).

We are interested in the image of $\mathcal{L}$ under the exponential map.

Players

Fix $A_0, A_1, A_2, \ldots, A_d \in \mathbb{S}^n$, where the last d are linearly independent. They determine a d-dim'l affine space of symmetric matrices (ASSM):

$$\mathcal{L} \;=\; A_0 + \mathrm{span}_{\mathbb{R}}(A_1, A_2, \ldots, A_d) \;\subset\; \mathbb{S}^n \simeq \mathbb{R}^{\binom{n+1}{2}}.$$

If $A_0 = 0$, then $\mathcal{L}$ is a *linear space of symmetric matrices* (LSSM).

We are interested in the image of $\mathcal{L}$ under the exponential map.

Definition

The *Gibbs manifold* is $\mathrm{GM}(\mathcal{L}) := \exp(\mathcal{L}) \;\subset\; \mathbb{S}^n_+$.

This is diffeomorphic to $\mathcal{L} = \mathbb{R}^d$ via the logarithm map.

Consider all polynomials that vanish on $\mathrm{GM}(\mathcal{L})$.

What is their common zero set?

Players

Fix $A_0, A_1, A_2, \ldots, A_d \in \mathbb{S}^n$, where the last d are linearly independent. They determine a d-dim'l affine space of symmetric matrices (ASSM):

$$\mathcal{L} \; = \; A_0 + \operatorname{span}_{\mathbb{R}}(A_1, A_2, \ldots, A_d) \;\; \subset \;\; \mathbb{S}^n \simeq \mathbb{R}^{\binom{n+1}{2}}.$$

If $A_0 = 0$, then $\mathcal{L}$ is a *linear space of symmetric matrices* (LSSM).

We are interested in the image of $\mathcal{L}$ under the exponential map.

Definition

The *Gibbs manifold* is $\mathrm{GM}(\mathcal{L}) := \exp(\mathcal{L}) \;\subset\; \mathbb{S}^n_+$.

This is diffeomorphic to $\mathcal{L} = \mathbb{R}^d$ via the logarithm map.

Consider all polynomials that vanish on $\mathrm{GM}(\mathcal{L})$.

What is their common zero set?

Definition

The *Gibbs variety* $\mathrm{GV}(\mathcal{L})$ is the Zariski closure of $\mathrm{GM}(\mathcal{L})$ in $\mathbb{C}^{\binom{n+1}{2}}$.

Quiz: What do we get if $\mathcal{L}$ consists of diagonal matrices?

Players

Fix $A_0, A_1, A_2, \ldots, A_d \in \mathbb{S}^n$, where the last d are linearly independent. They determine a d-dim'l affine space of symmetric matrices (ASSM):

$$\mathcal{L} \;=\; A_0 + \operatorname{span}_{\mathbb{R}}(A_1, A_2, \ldots, A_d) \;\subset\; \mathbb{S}^n \simeq \mathbb{R}^{\binom{n+1}{2}}.$$

If $A_0 = 0$, then $\mathcal{L}$ is a *linear space of symmetric matrices* (LSSM).

We are interested in the image of $\mathcal{L}$ under the exponential map.

Definition

The *Gibbs manifold* is $\mathrm{GM}(\mathcal{L}) := \exp(\mathcal{L}) \;\subset\; \mathbb{S}^n_+$.

This is diffeomorphic to $\mathcal{L} = \mathbb{R}^d$ via the logarithm map.

Consider all polynomials that vanish on $\mathrm{GM}(\mathcal{L})$.

What is their common zero set?

Definition

The *Gibbs variety* $\mathrm{GV}(\mathcal{L})$ is the Zariski closure of $\mathrm{GM}(\mathcal{L})$ in $\mathbb{C}^{\binom{n+1}{2}}$.

Quiz: What do we get if $\mathcal{L}$ consists of diagonal matrices?

Answer: **Toric geometry**

Example

Fix $n = 3$ and consider the LSSM

$$\mathcal{L} = \left\{ \begin{bmatrix} y_1 + y_2 + y_3 & y_1 & y_2 \\ y_1 & y_1 + y_2 + y_3 & y_3 \\ y_2 & y_3 & y_1 + y_2 + y_3 \end{bmatrix} : \; y_1, y_2, y_3 \in \mathbb{R} \right\}$$

The Gibbs manifold $\mathrm{GM}(\mathcal{L}) \subset \mathrm{int}(\mathbb{S}^3_+)$ has dimension 3.

Its points are the matrices $X \in \mathbb{S}^3$ whose logarithm has constant diagonal, with entries equal to the sum of the off-diagonal entries.

Example

Fix $n = 3$ and consider the LSSM

$$\mathcal{L} = \left\{ \begin{bmatrix} y_1+y_2+y_3 & y_1 & y_2 \\ y_1 & y_1+y_2+y_3 & y_3 \\ y_2 & y_3 & y_1+y_2+y_3 \end{bmatrix} : \ y_1, y_2, y_3 \in \mathbb{R} \right\}$$

The Gibbs manifold $\mathrm{GM}(\mathcal{L}) \subset \mathrm{int}(\mathbb{S}^3_+)$ has dimension 3.

Its points are the matrices $X \in \mathbb{S}^3$ whose logarithm has constant diagonal, with entries equal to the sum of the off-diagonal entries.

The Gibbs variety $\mathrm{GV}(\mathcal{L})$ has dimension 5.
It is the cubic hypersurface

$$\left\{ X \in \mathbb{S}^3 : \begin{array}{c} (x_{11}-x_{22})(x_{11}-x_{33})(x_{22}-x_{33}) \;\; = \\ x_{33}(x_{13}^2-x_{23}^2) + x_{22}(x_{23}^2-x_{12}^2) + x_{11}(x_{12}^2-x_{13}^2) \end{array} \right\}$$

Q: How to find such polynomials?

Example

Fix $n = 3$ and consider the LSSM

$$\mathcal{L} \;=\; \left\{ \begin{bmatrix} y_1+y_2+y_3 & y_1 & y_2 \\ y_1 & y_1+y_2+y_3 & y_3 \\ y_2 & y_3 & y_1+y_2+y_3 \end{bmatrix} :\; y_1, y_2, y_3 \in \mathbb{R} \right\}$$

The Gibbs manifold $\mathrm{GM}(\mathcal{L}) \subset \mathrm{int}(\mathbb{S}^3_+)$ has dimension 3.

Its points are the matrices $X \in \mathbb{S}^3$ whose logarithm has constant diagonal, with entries equal to the sum of the off-diagonal entries.

The Gibbs variety $\mathrm{GV}(\mathcal{L})$ has dimension 5.
It is the cubic hypersurface

$$\left\{ X \in \mathbb{S}^3 : \begin{array}{c} (x_{11}-x_{22})(x_{11}-x_{33})(x_{22}-x_{33}) \;\;= \\ x_{33}(x_{13}^2-x_{23}^2)+x_{22}(x_{23}^2-x_{12}^2)+x_{11}(x_{12}^2-x_{13}^2) \end{array} \right\}$$

Q: How to find such polynomials?

A1: Numerically   A2: Symbolically

Maximizing Entropy

The quotient map from $\mathbb{S}^3 \simeq \mathbb{R}^6$ onto $\mathbb{S}^3/\mathcal{L}^\perp \simeq \mathbb{R}^3$ takes matrices $X = [x_{ij}]$ to their inner products with a basis of $\mathcal{L}$:

$$\pi : \mathbb{S}^3_+ \to \mathbb{R}^3 : X \mapsto \left(\mathrm{trace}(X)+2x_{12}, \mathrm{trace}(X)+2x_{13}, \mathrm{trace}(X)+2x_{23}\right)$$

Each fiber $\pi^{-1}(b)$ is a 3-dimensional spectrahedron:

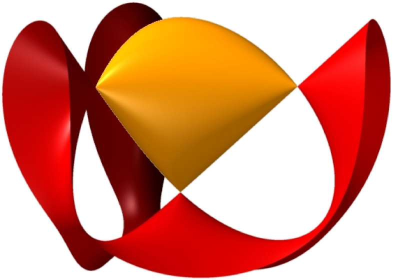

Maximizing Entropy

The quotient map from $\mathbb{S}^3 \simeq \mathbb{R}^6$ onto $\mathbb{S}^3/\mathcal{L}^\perp \simeq \mathbb{R}^3$ takes matrices $X = [x_{ij}]$ to their inner products with a basis of $\mathcal{L}$:

$$\pi : \mathbb{S}^3_+ \to \mathbb{R}^3 : X \mapsto \big(\mathrm{trace}(X)+2x_{12}, \mathrm{trace}(X)+2x_{13}, \mathrm{trace}(X)+2x_{23}\big)$$

Each fiber $\pi^{-1}(b)$ is a 3-dimensional spectrahedron:

The *von Neumann entropy* $h(X) = \mathrm{trace}(X - X \cdot \log(X))$ is maximized at the Gibbs point $\pi^{-1}(b) \cap \mathrm{GM}(\mathcal{L})$. The Gibbs manifold $\mathrm{GM}(\mathcal{L})$ is the set of all Gibbs points, in the various spectrahedra $\pi^{-1}(b)$, for all $b \in \mathbb{R}^3$.

Dimension

Theorem

Let $\mathcal{L} \subset \mathbb{S}^n$ be an ASSM of dimension d. The dimension of the Gibbs variety $\mathrm{GV}(\mathcal{L})$ is at most $n + d$. If $A_0 = 0$, i.e. $\mathcal{L}$ is an LSSM, then $\dim \mathrm{GV}(\mathcal{L})$ is at most $n + d - 1$.

Dimension

Theorem

Let $\mathcal{L} \subset \mathbb{S}^n$ be an ASSM of dimension d. The dimension of the Gibbs variety $\mathrm{GV}(\mathcal{L})$ is at most $n+d$. If $A_0 = 0$, i.e. $\mathcal{L}$ is an LSSM, then $\dim \mathrm{GV}(\mathcal{L})$ is at most $n+d-1$.

Theorem

Let $\mathcal{L}$ be an affine space of pairwise commuting symmetric matrices. The Gibbs variety $\mathrm{GV}(\mathcal{L})$ is a toric variety whose dimension is determined by the arithmetic of the eigenvalues.

Example

Let $n = 3$, $d = 1$, and fix the LSSM $\mathcal{L} = \mathbb{R}\begin{bmatrix} 4 & 1 & 1 \\ 1 & 3 & 1 \\ 1 & 1 & 3 \end{bmatrix}$. The Gibbs manifold $\mathrm{GM}(\mathcal{L})$ is a curve. Its Zariski closure $\mathrm{GV}(\mathcal{L})$ is the surface

$$\{X \in \mathbb{S}^3 : x_{11} - x_{23} - x_{33} = x_{12} - x_{13} = x_{22} - x_{33} = x_{23}^4 - 4x_{23}^3x_{33} \\ +6x_{23}^2x_{33}^2 - 4x_{23}x_{33}^3 + x_{33}^4 + 2x_{13}^2 - x_{23}^2 - 2x_{23}x_{33} - x_{33}^2 = 0\}.$$

Symbolic Algorithm

Algorithm 1 Implicitization of the Gibbs variety of an LSSM $\mathcal{L}$, defined over $\mathbb{Q}$

Input: Linearly independent matrices $A_1, \ldots, A_d \in \mathbb{S}^n$ with rational entries
Output: Polynomials that define $\mathrm{GV}(\mathcal{L})$, where $\mathcal{L} = \mathrm{span}_{\mathbb{R}}(A_1, \ldots, A_d)$

Compute the characteristic polynomial $P_{\mathcal{L}}(\lambda; y) = c_0(y) + c_1(y)\lambda + \cdots + c_n(y)\lambda^n$
Require: $P_{\mathcal{L}}(\lambda; y)$ has n distinct roots in $\overline{\mathbb{R}(y)}$
$E_1' \leftarrow \{$the n polynomials $(-1)^i \sigma_{n-i}(\lambda) - c_i(y)$ in (8)$\}$
$E_1 \leftarrow \{$generators of any associated prime over $\mathbb{Q}$ of $\langle E_1' \rangle\}$
$E_2 \leftarrow \{$the entries of $\phi(y, \lambda, z) - X\}$, with $X = (x_{ij})$ a symmetric matrix of variables
$E_2, D \leftarrow$ clear denominators in E_2 and record the least common denominator D
if the roots $\lambda_1, \ldots, \lambda_n$ of $P_{\mathcal{L}}(\lambda; y)$ are $\mathbb{Q}$-linearly dependent **then**
    $E_3 \leftarrow \{z^\alpha - z^\beta \,:\, \sum \alpha_i \lambda_i = \sum \beta_j \lambda_j,\ \alpha, \beta \in \mathbb{Z}^n_{\geqslant 0}\}$
else
    $E_3 \leftarrow \emptyset$
$I \leftarrow$ the ideal generated by E_1, E_2, E_3 in the polynomial ring $\mathbb{R}[y, \lambda, z, X]$
$I \leftarrow I : D^\infty$
$J \leftarrow$ elimination ideal obtained by eliminating y, λ, z from I
return a set of generators of J

Theorem

Let $\mathcal{L} \subset \mathbb{S}^n$ be an LSSM with distinct eigenvalues.
The Gibbs variety $\mathrm{GV}(\mathcal{L})$ is irreducible and unirational.
The ideal J found by Algorithm 1 is its prime ideal.

Pencils

A pencil of quadrics is an LSSM of dimension $d = 2$. They are classified by *Segre symbols*. We compute their Gibbs varieties.

Example

The pencil $\mathcal{L}$ for Segre symbol $\sigma = [(3),(1)]$ is spanned by

$$\begin{bmatrix} 0 & 0 & \alpha_1 & 0 \\ 0 & \alpha_1 & 1 & 0 \\ \alpha_1 & 1 & 0 & 0 \\ 0 & 0 & 0 & \alpha_2 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{for} \quad \alpha_1, \alpha_2 \in \mathbb{R} \text{ distinct.}$$

Here, $\dim \mathrm{GV}(\mathcal{L}) = 5$, our upper bound. Algorithm 1 finds the ideal

$$J = \left\langle x_{14}, x_{24}, x_{34},\ x_{13} - x_{22} + x_{33},\ x_{12}^2 - x_{11}x_{22} - x_{12}x_{23} + x_{11}x_{33} + x_{22}x_{33} - x_{33}^2 \right\rangle$$

Pencils

A pencil of quadrics is an LSSM of dimension $d = 2$. They are classified by *Segre symbols*. We compute their Gibbs varieties.

Example

The pencil $\mathcal{L}$ for Segre symbol $\sigma = [(3),(1)]$ is spanned by

$$\begin{bmatrix} 0 & 0 & \alpha_1 & 0 \\ 0 & \alpha_1 & 1 & 0 \\ \alpha_1 & 1 & 0 & 0 \\ 0 & 0 & 0 & \alpha_2 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{for} \quad \alpha_1, \alpha_2 \in \mathbb{R} \text{ distinct.}$$

Here, $\dim \mathrm{GV}(\mathcal{L}) = 5$, our upper bound. Algorithm 1 finds the ideal

$$J = \langle x_{14}, x_{24}, x_{34},\, x_{13}-x_{22}+x_{33},\, x_{12}^2-x_{11}x_{22}-x_{12}x_{23}+x_{11}x_{33}+x_{22}x_{33}-x_{33}^2 \rangle$$

If $\alpha_1 = \alpha_2$, then $\sigma = [(3,1)]$. Now dim $= 4$, given by the additional cubic

$$x_{11}x_{22}x_{33} + 2x_{12}x_{13}x_{23} - x_{13}^2x_{22} - x_{11}x_{23}^2 - x_{12}^2x_{33} - x_{44} \quad \in \quad J.$$

We have several general results on Gibbs varieties of pencils.

Convex Optimization

Fix an LSSM $\mathcal{L} = \mathrm{span}_{\mathbb{R}}(A_1, \ldots, A_d)$ and its linear map

$$\pi \,:\, \mathbb{S}^n_+ \to \mathbb{R}^d,\; X \,\mapsto\, \big(\langle A_1, X\rangle, \langle A_2, X\rangle, \ldots, \langle A_d, X\rangle\big).$$

Remark: The image of the PSD cone is the *spectrahedral shadow* $\pi(\mathbb{S}^n_+)$.

Semidefinite programming (SDP)

$$\text{Minimize} \quad \langle C, X\rangle \quad \text{subject to} \quad X \in \mathbb{S}^n_+ \;\text{ and }\; \pi(X) = b.$$

The feasible region $\pi^{-1}(b)$ is a *spectrahedron*.

Remark: These spectrahedra are compact if and only if $\mathcal{L} \cap \mathrm{int}(\mathbb{S}^n_+) \neq \emptyset$.

Entropic regularization of SDP

$$\text{Minimize} \quad \langle C, X\rangle - \epsilon \cdot h(X) \quad \text{subject to} \quad X \in \mathbb{S}^n_+ \;\text{ and }\; \pi(X) = b$$

Remark: $h(X) = \mathrm{trace}\big(X - X \cdot \log(X)\big)$ is the von Neumann entropy.

Convex Optimization: Example

Consider the linear map $\pi : \mathbb{S}^3_+ \to \mathbb{R}^3$ given by

$$\mathcal{L} = \left\{ \begin{bmatrix} y_1 + y_2 + y_3 & y_1 & y_2 \\ y_1 & y_1 + y_2 + y_3 & y_3 \\ y_2 & y_3 & y_1 + y_2 + y_3 \end{bmatrix} : y_1, y_2, y_3 \in \mathbb{R} \right\}$$

$\pi(\mathbb{S}^n_+)$ is the convex cone over the region defined by the red sextic

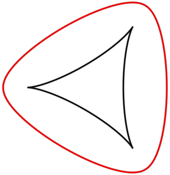

Black curve separates two types of spectrahedra seen as fibers

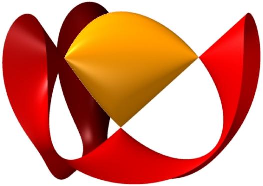

Samosa

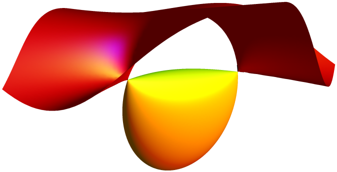

Teardrop

Entropic Regularization

We incorporate ϵ and the cost matrix C in the ASSM

$$\mathcal{L}_\epsilon \;:=\; \mathcal{L} - \frac{1}{\epsilon} C \quad \text{for any } \epsilon > 0.$$

For $\epsilon = \infty$, this is the LSSM, i.e. $\mathcal{L}_\infty = \mathcal{L}$.

Theorem

For $b \in \pi(\mathbb{S}^n_+)$, the intersection of $\pi^{-1}(b)$ with the Gibbs manifold $\mathrm{GM}(\mathcal{L}_\epsilon)$ consists of a single point X^*_ϵ. This point is the optimal solution to the regularized SDP. For $\epsilon = \infty$, it is the maximizer of von Neumann entropy over the spectrahedron $\pi^{-1}(b)$.

Entropic Regularization

We incorporate ϵ and the cost matrix C in the ASSM

$$\mathcal{L}_\epsilon \ := \ \mathcal{L} - \frac{1}{\epsilon} C \quad \text{for any } \epsilon > 0.$$

For $\epsilon = \infty$, this is the LSSM, i.e. $\mathcal{L}_\infty = \mathcal{L}$.

Theorem

For $b \in \pi(\mathbb{S}^n_+)$, the intersection of $\pi^{-1}(b)$ with the Gibbs manifold $\mathrm{GM}(\mathcal{L}_\epsilon)$ consists of a single point X^*_ϵ. This point is the optimal solution to the regularized SDP. For $\epsilon = \infty$, it is the maximizer of von Neumann entropy over the spectrahedron $\pi^{-1}(b)$.

Remark: The limit $\lim_{\epsilon \to 0} X^*_\epsilon$ exists and it is an optimal solution to the SDP. It is unique for generic C. This limit process is entropic regularization of SDP.

Quiz: What do we get if $\mathcal{L}$ consists of diagonal matrices?

Entropic Regularization

We incorporate ϵ and the cost matrix C in the ASSM

$$\mathcal{L}_\epsilon \;:=\; \mathcal{L} - \frac{1}{\epsilon} C \quad \text{for any } \epsilon > 0.$$

For $\epsilon = \infty$, this is the LSSM, i.e. $\mathcal{L}_\infty = \mathcal{L}$.

Theorem

For $b \in \pi(\mathbb{S}^n_+)$, the intersection of $\pi^{-1}(b)$ with the Gibbs manifold $\mathrm{GM}(\mathcal{L}_\epsilon)$ consists of a single point X^*_ϵ. This point is the optimal solution to the regularized SDP. For $\epsilon = \infty$, it is the maximizer of von Neumann entropy over the spectrahedron $\pi^{-1}(b)$.

Remark: The limit $\lim_{\epsilon \to 0} X^*_\epsilon$ exists and it is an optimal solution to the SDP. It is unique for generic C. This limit process is entropic regularization of SDP.

Quiz: What do we get if $\mathcal{L}$ consists of diagonal matrices?

Answer: **Linear Programming** and **Toric Geometry**

Quantum Optimal Transport

Fix the space $\mathbb{S}^{d_1 d_2}$ of symmetric matrices X of size $d_1 d_2 \times d_1 d_2$. Write $X = (x_{ijkl})$, where $(i,j),(k,l) \in [d_1] \times [d_2]$. The *marginalization map* is

$$\pi \, : \, \mathbb{S}_{+}^{d_1 d_2} \to \mathbb{S}^{d_1} \times \mathbb{S}^{d_2}, \;\; X \mapsto (Y, Z).$$

The $d_1 \times d_1$ matrix $Y = (y_{ik})$ and the $d_2 \times d_2$ matrix $Z = (z_{jl})$ are the *partial traces of X*. They satisfy $y_{ik} = \sum_{j=1}^{d_2} x_{ijkj}$ and $z_{jl} = \sum_{i=1}^{d_1} x_{ijil}$.

Quantum Optimal Transport

Fix the space $\mathbb{S}^{d_1 d_2}$ of symmetric matrices X of size $d_1 d_2 \times d_1 d_2$. Write $X = (x_{ijkl})$, where $(i,j),(k,l) \in [d_1] \times [d_2]$. The *marginalization map* is

$$\pi \,:\, \mathbb{S}_+^{d_1 d_2} \to \mathbb{S}^{d_1} \times \mathbb{S}^{d_2}, \;\; X \mapsto (Y, Z).$$

The $d_1 \times d_1$ matrix $Y = (y_{ik})$ and the $d_2 \times d_2$ matrix $Z = (z_{jl})$ are the *partial traces of X*. They satisfy $y_{ik} = \sum_{j=1}^{d_2} x_{ijkj}$ and $z_{jl} = \sum_{i=1}^{d_1} x_{ijil}$.

Quantum optimal transport (QOT) is the task of minimizing a linear function $X \mapsto \langle C, X \rangle$ over the spectrahedra $\pi^{-1}(Y, Z)$. For this SDP, the Gibbs manifold equals the Gibbs variety inside the PSD cone $\mathbb{S}_+^{d_1 d_2}$.

Theorem

The Gibbs manifold for QOT is a semialgebraic subset of $\mathbb{S}_+^{d_1 d_2}$, namely

$$\mathrm{GM}(\mathcal{L}) \;=\; \left\{ Y \otimes Z \;:\; Y \in \mathbb{S}_+^{d_1} \text{ and } Z \in \mathbb{S}_+^{d_2} \right\}.$$

Gibbs variety $\mathrm{GV}(\mathcal{L})$ is the cone over the Segre variety $\mathbb{P}^{\binom{d_1+1}{2}-1} \times \mathbb{P}^{\binom{d_2+1}{2}-1}$.

QOT Example

Fix $d_1 = d_2 = 2$. The map π takes any 4×4 PSD matrix

$$X = \begin{bmatrix} x_{1111} & x_{1112} & x_{1121} & x_{1122} \\ x_{1112} & x_{1212} & x_{1221} & x_{1222} \\ x_{1121} & x_{1221} & x_{2121} & x_{2122} \\ x_{1122} & x_{1222} & x_{2122} & x_{2222} \end{bmatrix}$$

to its partial traces

$$Y = \begin{bmatrix} x_{1111} + x_{1212} & x_{1121} + x_{1222} \\ x_{1121} + x_{1222} & x_{2121} + x_{2222} \end{bmatrix} \text{ and } Z = \begin{bmatrix} x_{1111} + x_{2121} & x_{1112} + x_{2122} \\ x_{1112} + x_{2122} & x_{1212} + x_{2222} \end{bmatrix}.$$

Here $\pi(\mathbb{S}^4_+)$ is the cone over the product of two disks. The fibers $\pi^{-1}(Y, Z)$ are the 5-dimensional transportation spectrahedra.

Quiz: What if we restrict to diagonal matrices?

QOT Example

Fix $d_1 = d_2 = 2$. The map π takes any 4×4 PSD matrix

$$X = \begin{bmatrix} x_{1111} & x_{1112} & x_{1121} & x_{1122} \\ x_{1112} & x_{1212} & x_{1221} & x_{1222} \\ x_{1121} & x_{1221} & x_{2121} & x_{2122} \\ x_{1122} & x_{1222} & x_{2122} & x_{2222} \end{bmatrix}$$

to its partial traces

$$Y = \begin{bmatrix} x_{1111} + x_{1212} & x_{1121} + x_{1222} \\ x_{1121} + x_{1222} & x_{2121} + x_{2222} \end{bmatrix} \text{ and } Z = \begin{bmatrix} x_{1111} + x_{2121} & x_{1112} + x_{2122} \\ x_{1112} + x_{2122} & x_{1212} + x_{2222} \end{bmatrix}.$$

Here $\pi(\mathbb{S}^4_+)$ is the cone over the product of two disks. The fibers $\pi^{-1}(Y, Z)$ are the 5-dimensional transportation spectrahedra.

Quiz: What if we restrict to diagonal matrices?

Answer: We get the familiar toric picture

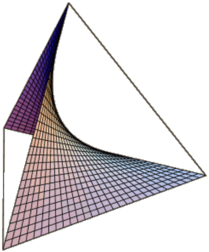

Conclusion

Gibbs manifolds are images of linear spaces of symmetric matrices under the exponential map. Arising in applications like optimization, statistics and quantum physics, they extend the *ubiquitous role of toric geometry*.

$$\left\{ \begin{bmatrix} a_0+b_0 & 0 & 0 & 0 \\ 0 & a_0+b_1 & 0 & 0 \\ 0 & 0 & a_1+b_0 & 0 \\ 0 & 0 & 0 & a_1+b_1 \end{bmatrix} \right\}$$

$$p_{00} = \exp(a_0+b_0),\ p_{01} = \exp(a_0+b_1),$$
$$p_{10} = \exp(a_1+b_0),\ p_{11} = \exp(a_1+b_1).$$

$$p_{00}p_{11} - p_{01}p_{10} \ = \ 0.$$

The Gibbs variety is the zero locus of all polynomials that vanish on the Gibbs manifold. We compute these polynomials and show that the Gibbs variety is low-dimensional. Our theory is applied to a wide range of scenarios, including matrix pencils and quantum optimal transport.

References

D. Pavlov, BSt and S. Telen: *Gibbs manifolds*, Information Geometry (2023)

D. Pavlov: *Logarithmically sparse symmetric matrices*, arXiv:2301.10042

E. Duarte, D. Pavlov and M. Wiesmann: *Algebraic geometry of quantum graphical models*, arXiv:2308.11538

References

D. Pavlov, BSt and S. Telen: *Gibbs manifolds*, Information Geometry (2023)

D. Pavlov: *Logarithmically sparse symmetric matrices*, arXiv:2301.10042

E. Duarte, D. Pavlov and M. Wiesmann: *Algebraic geometry of quantum graphical models*, arXiv:2308.11538

BSt, S. Telen, F.-X. Vialard and M. von Renesse: *Toric geometry of entropic regularization*, J. Symbolic Computation (2024)

S. Cole, M. Eckstein, S. Friedland and K. Zyczkowski: *Quantum Monge–Kantorovich problem and transport distance between density matrices*, Physical Review Letters (2022)

References

D. Pavlov, BSt and S. Telen: *Gibbs manifolds*, Information Geometry (2023)

D. Pavlov: *Logarithmically sparse symmetric matrices*, arXiv:2301.10042

E. Duarte, D. Pavlov and M. Wiesmann: *Algebraic geometry of quantum graphical models*, arXiv:2308.11538

BSt, S. Telen, F.-X. Vialard and M. von Renesse: *Toric geometry of entropic regularization*, J. Symbolic Computation (2024)

S. Cole, M. Eckstein, S. Friedland and K. Zyczkowski: *Quantum Monge–Kantorovich problem and transport distance between density matrices*, Physical Review Letters (2022)

BSt and C. Uhler: *Multivariate Gaussians, semidefinite matrix completion, and convex algebraic geometry*, Annals of the Institute of Statistical Math. (2010)

Y. Jiang and BSt: *Bad projections of the PSD cone*, Collectanea Mathematica (2021)

C. Fevola, Y. Mandelshtam and BSt: *Pencils of quadrics: old and new*, Le Matematiche (2021)

Algebraic Statistics

is the setting of the
papers in blue

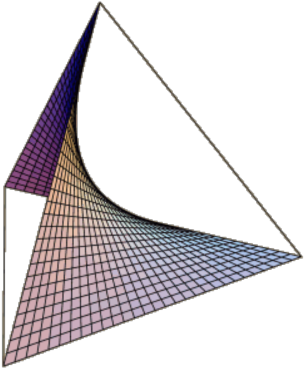

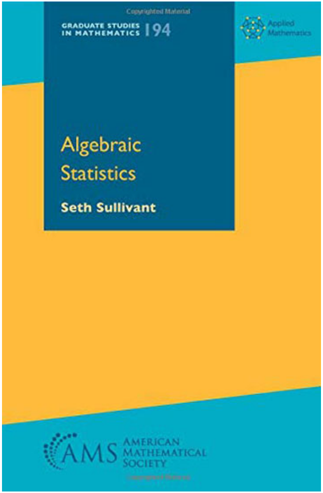

And, yes, let's chat about finding those roots