The free Wasserstein manifold

David Jekel, Wuchen Li, Dima Shlyakhtenko

February 22, 2021

This work was supported in part by the National Science Foundation.

The talk will focus on the big picture and thus precise definitions will only be given when helpful. Before the rigorous statements, there will be several slides of motivation and introduction aimed at people who are familiar with free probability. Feel free to interrupt with questions about notation.

Motivation

The classical Wasserstein manifold $\mathscr{P}(M)$, whose points are smooth positive probability densities, is an infinite-dimensional Riemannian framework which nicely describes things such as entropy, the heat equation, log-Sobolev inequalities, optimal transport, measure-preserving transformations, etc.

Because many of these notions have analogs in free probability and random matrix theory, we want to define the tracial non-commutative version of the Wasserstein manifold, in which we use non-commutative laws instead of measures.

A big obstacle is that we don't have a direct analog of density in the non-commutative setting. However, there are several indications that the *log-density* is a better behaved notion.

Motivation — random matrices

Given some self-adjoint non-commutative polynomial f (with nice enough behavior at ∞), we can define a function $V^{(N)} : M_N(\mathbb{C})_{sa}^d \to \mathbb{R}$ by $V^{(N)}(x) = \operatorname{tr}_N(f(x))$, where $x = (x_1, \dots, x_d)$ is a d-tuple of self-adjoint $N \times N$ matrices and $\operatorname{tr}_N = (1/N) \operatorname{Tr}$ is the normalized trace on $M_N(\mathbb{C})$. Then we define a probability measure $\mu^{(N)}$ on $M_N(\mathbb{C})_{sa}^d$ by

$$d\mu^{(N)}(x) = \text{constant}(V, N) e^{-N^2 V^{(N)}(x)}\, dx,$$

where dx is Lebesgue measure.

Motivation — random matrices

Letting $X^{(N)}$ be a random matrix tuple chosen according to the measure $\mu^{(N)}$, much past work has found sufficient conditions for $X^{(N)}$ to converge in non-commutative law as $N \to \infty$. More precisely, for certain V, for every non-commutative polynomial p, $\operatorname{tr}_N(p(X^{(N)}))$ converges almost surely to some deterministic limit, which is described by $\tau(p(X))$ for some d-tuple X from a tracial von Neumann algebra $(\mathcal{A}, \tau)$.

Then we might want to say that "$\operatorname{tr}(f)$ is a log-density of the distribution of X."

Motivation — free score function

This idea is closely related to Voiculescu's idea of a free score function (a.k.a. conjugate variable).

The classical score function of a measure with density ρ on $\mathbb{R}^d$ is $-\nabla \log \rho$. If X is a random variable with density ρ and if $\xi = -(\nabla \log \rho)(X)$, then we have the integration-by-parts relation

$$\mathbb{E}[\langle \xi, f(X) \rangle] = \mathbb{E}[\operatorname{Tr}(Df(X))],$$

for all $f \in C_c^\infty(\mathbb{R}^d, \mathbb{R}^d)$, where Df is the Jacobian matrix.

Given a tracial von Neumann algebra $(\mathcal{A}, \tau)$ generated by $X \in \mathcal{A}_{\mathrm{sa}}^d$, we say that $\xi \in \mathcal{A}_{\mathrm{sa}}^d$ is a *free score function* for X if

$$\langle \xi, p(X) \rangle_\tau = \tau \otimes \tau \otimes \operatorname{Tr}(\mathcal{J} p)$$

for every non-commutative polynomial p, where $\mathcal{J} p$ is the matrix of derivatives of p in the sense of Voiculescu's free difference quotient.

Motivation — free score function

In the setting of the random matrix d-tuples $X^{(N)}$ above given by $V^{(N)}(x) = \mathrm{tr}_N(f(x))$, the free score function becomes very concrete. Since $\rho = \mathrm{const} e^{-N^2 V^{(N)}}$, the classical score function is (up to normalization) $\nabla V^{(N)}(x)$, which after some matrix computations works out to $\mathcal{D}^\circ f(x)$, where $\mathcal{D}^\circ f$ is Voiculescu's cyclic gradient.

Classical integration by parts plus some matrix computations tell us that

$$\mathbb{E}[\langle \nabla V^{(N)}(X^{(N)}), p(X^{(N)}) \rangle_{\mathrm{tr}_N}] = \mathbb{E}[\mathrm{tr}_N \otimes \mathrm{tr}_N \otimes \mathrm{Tr}_d[\mathcal{J} p(X^{(N)})]].$$

Thus, if X is the d-tuple of self-adjoint operators from $(\mathcal{A}, \tau)$ describing the large-N limit, then $\xi = \mathcal{D}^\circ p(X)$ is the free score function for X.

The existence of a free score means that "the gradient of the log-density makes sense as an element of L^2." This motivates us to make the log-density the central object of study.

Overview

We're going to define the free Wasserstein manifold $\mathscr{W}(\mathbb{R}^{*d})$ as the space of certain "log-density" functions V, which are a generalization of things like $\operatorname{tr}(f)$.

More precisely, V will some from a space of tracial non-commutative smooth functions that are defined in terms of trace polynomials. (Similar spaces were defined in Dabrowski, Guionnet, and Shlaykhtenko's 2016 preprint on free transport, but we take a different approach to the norms.)

The tangent space of $\mathscr{W}(\mathbb{R}^{*d})$ at V is similarly a space of tracial non-commutative smooth functions W, which are viewed as perturbations of V.

Overview

In order to obtain a Riemannian metric for the Wasserstein manifold, we must associate a non-commutative law μ_V to each V. This is much trickier than in the classical case (where we would just set $d\mu_V(x) = \text{constant}\, e^{-V(x)}\, dx$), but there are two known methods for doing this:

(1) For each V, find a (hopefully unique) law μ that maximizes $\chi(\mu) - \mu(V)$, where χ is Voiculescu's free microstate entropy. This approach is inspired by Voiculescu and is closely related to the random matrix models discussed before.

(2) Set up the free stochastic differential equation $dX_t = dS_t - (1/2)\nabla V(X_t)\, dt$, where S_t is a free Brownian motion (still self-adjoint d-tuple), and (hopefully) recover μ_V as the limiting distribution of X_t as $t \to \infty$. This approach was pioneered by Biane and Speicher (1999) and further developed by Guionnet, Shlyakhtenko, and Dabrowski.

Outline

- Background on non-commutative laws.
- Tracial non-commutative smooth functions.
- Free Wasserstein manifold and diffeomorphism group.
- Riemannian metric.
- Strategy to construct smooth transport.
- Inversion of the Laplacian L_V through heat semigroup and SDE.
- Free Gibbs laws through maximization of $\chi(\mu) - \mu(V)$.
- Geodesics and optimal transport.

Operator algebras and laws

Definition

A *unital* C^**-algebra* is a subalgebra of $B(H)$ (for some Hilbert space H) that is closed under adjoints and limits in operator norm.

Definition

A *tracial* C^**-algebra* is a pair $(\mathcal{A}, \tau)$ where $\mathcal{A}$ is a C^*-algebra and τ is a *faithful trace*, that is,

- $\tau(1) = 1$,
- $\tau(a^* a) \geq 0$ with equality if and only if $a = 0$,
- $\tau(ab) = \tau(ba)$.

Remark

We don't need to go into the definition of von Neumann algebras now. But every tracial C^*-algebra can be completed to a tracial von Neumann algebra.

Operator algebras and laws

Definition

Let $\mathbb{C}\langle x_1, \dots, x_d \rangle$ be the algebra of non-commutative polynomials equipped with the $*$-operation such that $x_j^* = x_j$. (We typically use x to denote formal or generic variables and X to denote a specific tuple of operators.)

Definition

Let $\Sigma_{d,R}$ be the set of linear functionals $\lambda : \mathbb{C}\langle x_1, \dots, x_d \rangle \to \mathbb{C}$ such that

- $\lambda(1) = 1$,
- $\lambda(p * p) \geq 0$,
- $\lambda(pq) = \lambda(qp)$,
- $|\lambda(x_{i_1} \dots x_{i_k})| \leq R^k$,

equipped with the weak-$\star$ topology (as a subset of the dual of $\mathbb{C}\langle x_1, \dots, x_d \rangle$). We call elements of $\Sigma_{d,R}$ *non-commutative laws with exponential bound* R.

Operator algebras and laws

Proposition

If $(\mathcal{A}, \tau)$ is a tracial C^*-algebra and $X = (X_1, \dots, X_d) \in \mathcal{A}_{\text{sa}}^d$, then the map

$$\lambda_X : \mathbb{C}\langle x_1, \dots, x_d \rangle \to \mathbb{C}, p \mapsto \tau(p(X))$$

is a non-commutative law with exponential bound $\|X\|_\infty = \max_j \|X_j\|$. Conversely, every $\lambda \in \Sigma_{d,R}$ can be realized as λ_X for some $(\mathcal{A}, \tau)$ and $X \in \mathcal{A}_{\text{sa}}^d$ with $\|X\|_\infty \leq R$.

The proof is a variant of the GNS construction. The proposition can be interpreted as follows:

1. $\Sigma_{d,R}$ is the space of traces on the C^*-universal free product $C([-R, R])^{*d}$.
2. $\Sigma_{d,R}$ is in bijection with isomorphism classes of triples $(\mathcal{A}, \tau, X)$, where $(\mathcal{A}, \tau)$ is a tracial C^*-algebra and $X \in \mathcal{A}_{\text{sa}}^d$ generates $\mathcal{A}$; here isomorphism means a C^*-isomorphism that preserves the trace and generators.

Trace polynomials

We next describe non-commutative functions that are modeled on trace polynomials in a similar spirit to Dab. Gui. Shl. 2016.

A *trace polynomial* in $(x_1, \dots, x_d)$ is an expression formed through addition, multiplication, and application of a symbol tr, such as

$$f(x_1, x_2, x_3) = -\operatorname{tr}(x_1^2 x_2) \operatorname{tr}(x_3) x_1 + \operatorname{tr}(x_2 x_3) + 5 \operatorname{tr}(x_1 x_2 x_3) \operatorname{tr}(x_1) x_2 x_3^2.$$

These expressions are considered modulo the relations that $\operatorname{tr}(pq) = \operatorname{tr}(qp)$ and $\operatorname{tr}(\operatorname{tr}(p)q) = \operatorname{tr}(p) \operatorname{tr}(q)$.

For any tracial C^*-algebra $(\mathcal{A}, \tau)$ and $X \in \mathcal{A}_{\mathrm{sa}}^d$, we can evaluate a trace polynomial f on X by substituting X_j for the formal symbol x_j and τ for the formal symbol tr. Hence, a trace polynomial f gives rise to a function $f^{\mathcal{A},\tau} : \mathcal{A}_{\mathrm{sa}}^d \to \mathcal{A}$.

Trace polynomials

Trace polynomials have several advantages over non-commutative polynomials.

1. It follows from the work of Procesi (1976) that every function $M_N(\mathbb{C})_{\text{sa}}^d \to M_N(\mathbb{C})$ that is entrywise polynomial and is invariant under unitary conjugation must be given by a trace polynomial.
2. For each trace polynomial f, we can compute the Laplacian of $f^{M_N(\mathbb{C}),\text{tr}_N}$ as a function on $M_N(\mathbb{C})_{\text{sa}}^d$ (equipped with the inner product from tr_N). The Laplacian $(1/N^2)\Delta f^{M_N(\mathbb{C}),\text{tr}_N}$ is a trace polynomial and it converges coefficientwise as $N \to \infty$ to some trace polynomial Lf.

We'll define the non-commutative space $C^k(\mathbb{R}^{*d})$ roughly as functions that such that the first k derivatives can be approximated on operator-norm balls by trace polynomials.

Description of trace C^k functions

The space $C^k_{\mathrm{tr}}(\mathbb{R}^{*d})$ is described as follows:

- Each $f \in C^k_{\mathrm{tr}}(\mathbb{R}^{*d})$ is a collection of functions $f^{\mathcal{A},\tau} : \mathcal{A}_{\mathrm{sa}}^d \to \mathcal{A}$ for tracial C^*-algebras $(\mathcal{A},\tau)$.
- $f^{\mathcal{A},\tau}$ must be a C^k function in the sense of Fréchet differentiation.
- The derivative $\partial^k f^{\mathcal{A},\tau}(\mathrm{X})$ is a multilinear map $\mathcal{A}_{\mathrm{sa}}^d \times \cdots \times \mathcal{A}_{\mathrm{sa}}^d \to \mathcal{A}$.
- Inspired by the non-commutative Hölder's inequality, we define the norm $\|\partial^j f^{\mathcal{A},\tau}(\mathrm{X})\|_{\mathscr{M}^j}$ as the smallest constant such that

$$\|\partial^j f^{\mathcal{A},\tau}(\mathrm{X})[\mathrm{Y}_1,\dots,\mathrm{Y}_k]\|_p \leq \|\partial^j f^{\mathcal{A},\tau}(\mathrm{X})\|_{\mathscr{M}^k} \|\mathrm{Y}_1\|_{p_1} \dots \|\mathrm{Y}_j\|_{p_j}.$$

  where $1/p = 1/p_1 + \cdots + 1/p_j$, and where $j = 0, \dots, k$.
- Then $\|\partial^j f\|_{\mathscr{M}^j,R}$ is the supremum of $\|\partial^j f^{\mathcal{A},\tau}(\mathrm{X})\|_{\mathscr{M}^j}$ over $(\mathcal{A},\tau)$ and $X \in \mathcal{A}_{\mathrm{sa}}^d$ with $\|X\|_\infty \leq R$.
- For $R > 0$ and $j \leq k$, we assume that $\|\partial^j f\|_{\mathscr{M}^j,R}$ is finite and that $\partial^j f$ can be approximated in this norm by trace polynomials of X, Y_1, $\dots$, Y_k that are multilinear in Y_1, $\dots$, Y_k.

Properties of trace C^k functions

There are also spaces $C_{\mathrm{tr}}(\mathbb{R}^{*d}, \mathscr{M}^j(\mathbb{R}^{*d_1}, \dots, \mathbb{R}^{*d_n}))$ of functions where $f^{\mathcal{A},\tau}(X)$ is a multilinear map $\mathcal{A}_{\mathrm{sa}}^{d_1} \times \dots \times \mathcal{A}_{\mathrm{sa}}^{d_n} \to \mathcal{A}$.

The exact definition of these spaces is less important than the properties:

- These spaces are closed under composition, whenever the composition makes sense, and they satisfy the chain rule.
- There is an inverse function theorem: If f is $C_{\mathrm{tr}}^k(\mathbb{R}^{*d})$ self-adjoint d-tuple, and if $\partial f - \mathrm{Id}$ is uniformly bounded by a constant $c < 1$, then f^{-1} is defined and is C_{tr}^k.
- There is a trace map $\mathrm{tr} : C_{\mathrm{tr}}^k(\mathbb{R}^{*d}) \to C_{\mathrm{tr}}^k(\mathbb{R}^{*d})$ given by $\mathrm{tr}(f)^{\mathcal{A},\tau}(X) = \tau(f^{\mathcal{A},\tau}(X))$. The image $\mathrm{tr}(C_{\mathrm{tr}}^k(\mathbb{R}^{*d}))$ consists of those f which are scalar-valued.

Examples of trace C_{tr}^k functions

- Of course, trace polynomials are C_{tr}^∞ functions.
- If $\phi : \mathbb{R} \to \mathbb{R}$ such that $\int_{\mathbb{R}} |2\pi s|^k \widehat{\phi}(s)\, ds < \infty$, then the function $f^{\mathcal{A},\tau}(X) = \phi(X)$ (defined by functional calculus) is in $C_{\mathrm{tr}}^k(\mathbb{R}^{*1})$ and the kth derivative is bounded by $\int_{\mathbb{R}} |2\pi s|^k \widehat{\phi}(s)\, ds$.
- Together with the chain rule, this shows that there is an abundance of $BC_{\mathrm{tr}}^k(\mathbb{R}^{*d})$ functions, that is, functions in $C_{\mathrm{tr}}^k(\mathbb{R}^{*d})$ such that

$$\|\partial^j f\|_{\mathscr{M}^j,u} := \sup_{R>0} \|\partial^j f\|_{\mathscr{M}^j,R} < \infty.$$

- Imposing certain growth conditions at ∞ on a $C_{\mathrm{tr}}(\mathbb{R}^{*d})$ function is not a big restriction. This makes life easier than it would be if we only used trace polynomials.

Differentiation of trace C^k functions

For scalar-valued $g \in C^k_{\mathrm{tr}}(\mathbb{R}^{*d})$, we can define a gradient $\nabla g \in C^k_{\mathrm{tr}}(\mathbb{R}^{*d})$. In the case where $g = \mathrm{tr}(p)$ for some non-commutative polynomial p, then ∇g is the cyclic gradient of p.

The analog of C^k functions from $\mathbb{R}^d$ to $M_d(\mathbb{C})$ is the space $C^k_{\mathrm{tr}}(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d}))$. This is the space that contains the derivative ∂f when $f \in C^{k+1}_{\mathrm{tr}}(\mathbb{R}^{*d})^d_{\mathrm{sa}}$, as well as the Hessian of g when g is a scalar-valued element of $C^{k+1}_{\mathrm{tr}}(\mathbb{R}^{*d})$.

For $F \in C^k_{\mathrm{tr}}(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d}))$, for $X \in \mathcal{A}^d_{\mathrm{sa}}$, the object $F^{\mathcal{A},\tau}(\mathrm{X})$ is a linear transformation $\mathcal{A}^d \to \mathcal{A}^d$. We define $F \# G$ to be the pointwise composition of these linear transformations. It turns out that $C^k_{\mathrm{tr}}(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d}))$ is a $*$-algebra with respect to the $\#$-multiplication.

Differentiation of trace C^k functions

We can define a trace $\mathrm{Tr}_\# : C^k_{\mathrm{tr}}(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d})) \to \mathrm{tr}(C^k_{\mathrm{tr}}(\mathbb{R}^{*d}))$ by

$$(\mathrm{Tr}_\#(F))^{\mathcal{A},\tau}(X) = \langle S, F^{\mathcal{A}*\mathcal{B},\tau*\sigma}(X)[S] \rangle_{\tau*\sigma},$$

where $(\mathcal{B}, \sigma)$ is the tracial C^*-algebra of generated by a free semicircular d-tuple S.

This is the analog of the map $C^k(\mathbb{R}^d, M_d(\mathbb{C})) \to C^k(\mathbb{R}^d)$ defined by pointwise application of the trace Tr_d on $M_d(\mathbb{C})$. This is because the trace of a matrix A can be expressed as $\mathbb{E}\langle Y, AY \rangle$ where Y is a standard Gaussian random vector in $\mathbb{R}^d$, and the analog of the Gaussian in free probability is the semicircular family.

Another motivating example is that if $F \in C^k_{\mathrm{tr}}(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d}))$ is given by $F^{\mathcal{A},\tau}(X)[Y]_i = \sum_j p_{i,j}^{\mathcal{A},\tau}(X) Y_j q_{i,j}^{\mathcal{A},\tau}(X)$ for some matrix $(p_{i,j} \otimes q_{i,j})_{i,j}$ of non-commutative polynomials, then

$$\mathrm{Tr}_\#(F)^{\mathcal{A},\tau}(X) = \sum_i \tau(p_{i,i}(X)) \tau(q_{i,i}(X)).$$

Differentiation of trace C^k functions

The trace $\mathrm{Tr}_\#$ allows us to define the divergence operator

$$\nabla^\dagger : C_{\mathrm{tr}}^{k+1}(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d})) \to \mathrm{tr}(C_{\mathrm{tr}}^k(\mathbb{R}^{*d}))$$

as the trace of the Jacobian, as well as the Laplacian

$$L = \nabla^\dagger \nabla : \mathrm{tr}(C_{\mathrm{tr}}^{k+2}(\mathbb{R}^{*d})) \to \mathrm{tr}(C_{\mathrm{tr}}^k(\mathbb{R}^{*d})).$$

These operators are the limits of the corresponding normalized divergence and Laplacian for functions on $M_N(\mathbb{C})_{\mathrm{sa}}^d$.

Furthermore, the trace $\mathrm{Tr}_\#$ gives rise to a Fuglede-Kadison log-$|$determinant$|$ map

$$\log \Delta_\# : GL(C_{\mathrm{tr}}^k(\mathbb{R}^{*d}, \mathscr{M}(\mathbb{R}^{*d}))) \to \mathrm{tr}(C_{\mathrm{tr}}^k(\mathbb{R}^{*d})).$$

Free Wasserstein manifold and diffeomorphism group

We'll first set up the manifold formally. Afterwards, we'll describe how to extract a non-commutative law μ_V from V and hence define the Riemannian metric.

Definition

The *free Wasserstein manifold* $\mathscr{W}(\mathbb{R}^{*d})$ is the set of $V \in \operatorname{tr}(C_{\mathrm{tr}}^{\infty}(\mathbb{R}^{*d}))$ such that V has "quadratic growth at ∞" in the sense that for some constants $a, a' > 0$ and $b, b' \in \mathbb{R}$, we have

$$a \sum_j \tau(X_j^2) + b \leq V^{\mathcal{A},\tau}(X) \leq a' \sum_j \tau(X_j^2) + b'.$$

Definition

The *free diffeomorphism group* $\mathscr{D}(\mathbb{R}^{*d})$ is the set of $f \in C_{\mathrm{tr}}^{\infty}(\mathbb{R}^{*d})_{\mathrm{sa}}^{d}$ such that f has an inverse function f^{-1} in $C^{\infty}(\mathbb{R}^{*d})_{\mathrm{sa}}^{d}$, and ∂f, ∂f^{-1} are bounded. Note this is a group under composition.

Tangent vectors

A *tangent vector to* $\mathscr{W}(\mathbb{R}^{*d})$ *at* V is an equivalence class of C^1 paths $(-\epsilon, \epsilon) \to \mathscr{W}(\mathbb{R}^{*d}) : t \mapsto V_t$ with $V_0 = V$, where two paths are equivalent if they have the same $\dot{V}_0$. For convenience, we assume that V_t satisfies the quadratic growth bounds with a, a', b, b' independent of t.

A *tangent vector to* $\mathscr{D}(\mathbb{R}^{*d})$ at id is similarly an equivalence class of C^1 paths $t \mapsto f_t$ with $f_0 = \mathrm{id}$, and the equivalence is equality of $\dot{f}_0$. Again, assume that ∂f_t and ∂f_t^{-1} are uniformly bounded.

Here, by "C^1 path", we mean it is continuously differentiable with respect to the Fréchet topology of $C_{\mathrm{tr}}^{\infty}(\mathbb{R}^{*d})$ on the target space (defined by the seminorms of each derivative $\partial^j f$ on each ball of radius R).

The transport action

In the classical case, one studies the action of $\mathrm{Diff}(\mathbb{R}^d)$ on $\mathscr{P}(\mathbb{R}^d)$ by push-forward, which is viewed as an infinite-dimensional Lie group acting on an infinite-dimensional Riemannian manifold.

If $\mu \in \mathscr{P}(\mathbb{R}^d)$ has density e^{-V} and if f is a diffeomorphism, then $f_*\mu$ has density $e^{-(V \circ f^{-1} - \log|\det Df^{-1}|)}$ using the change of variables formula. This motivates the following definition.

Definition

We define the *transport action* $\mathscr{D}(\mathbb{R}^{*d}) \curvearrowright \mathscr{W}(\mathbb{R}^{*d})$ by

$$(f, V) \mapsto f_* V := V \circ f^{-1} - \log \Delta_\#(\partial f^{-1}).$$

One can check this is a well-defined group action.

Differential of the transport action

The key computation behind transport theory is the description of the differential of the transport action. We define

$$\nabla_V^* : C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d})^d \to \mathrm{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d}))$$

by

$$\nabla_V^* f = -\nabla^\dagger f + \partial V \# f = -\operatorname{Tr}_\#(\partial f) + \langle \nabla V, f \rangle_{\mathrm{tr}}.$$

(This is just notation; it is not actually the adjoint.)

Lemma

Let $V \in \mathscr{W}(\mathbb{R}^{*d})$ and let $t \mapsto f_t$ be a tangent vector to $\mathscr{D}(\mathbb{R}^{*d})$ at id. Then

$$\frac{d}{dt}\Big|_{t=0} (f_t)_* V = -\nabla_V^* \dot{f}_0.$$

In other words, $-\nabla_V^*$ is the differential at id of the orbit map $\mathscr{D}(\mathbb{R}^{*d}) \to \mathscr{W}(\mathbb{R}^{*d}) : f \mapsto f_* V$.

$\mathscr{D}(\mathbb{R}^{*d})$ as a Lie group

We saw that the tangent space of $\mathcal{D}(\mathbb{R}^{*d})$ is (a dense subspace of) the space of vector fields $C_{\text{tr}}(\mathbb{R}^{*d})_{\text{sa}}^d$. Conversely:

Lemma

Given a time-dependent vector field $t \mapsto h_t$ (continuous in t) such that ∂h_t is uniformly bounded, there exists a unique path f_t in $\mathscr{D}(\mathbb{R}^{*d})$ such that

$$f_0 = \text{id}, \qquad \dot{f}_t = h_t \circ f_t.$$

The proof is similar to classical ODE theory. If h is independent of t, then we get a one-parameter subgroup of $\mathscr{D}(\mathbb{R}^{*d})$. Combining this with our previous observation:

Lemma

Let $h \in C_{\text{tr}}(\mathbb{R}^{*d})_{\text{sa}}^d$ with ∂h bounded, and let f_t be the corresponding one-parameter subgroup. Then $(f_t)_* V = V$ for all t if and only if $\nabla_V^* h = 0$.

$\mathscr{D}(\mathbb{R}^{*d})$ as a Lie group

By studying the one-parameter subgroups of $\mathscr{D}(\mathbb{R}^{*d})$ as described above, we arrive at the following definition of the Lie bracket, completely analogous to the Lie bracket on vector fields of $\mathbb{R}^d$.

Definition

For two vector fields $h_1,\ h_2 \in C_{\text{tr}}(\mathbb{R}^{*d})^d$, let

$$[h_1, h_2] = \partial h_1 \# h_2 - \partial h_2 \# h_1.$$

This generalizes the definition of Lie brackets for non-commutative polynomials used in Voiculescu's paper "Cyclomorphy."

$\mathscr{D}(\mathbb{R}^{*d})$ as a Lie group

For each $V \in \mathscr{W}(\mathbb{R}^{*d})$, its stabilizer $\{f \in \mathscr{D}(\mathbb{R}^{*d}) : f_* V = V\}$ is a "Lie subgroup," analogous to a classical group of measure-preserving transformations.

By our previous observations, the corresponding Lie subalgebra should be the set of vector fields h with $\nabla_V^* h = 0$. We can verify directly that this is indeed a Lie subalgebra:

Lemma

$\nabla_V^*[h_1, h_2] = \partial(\nabla_V^* h_1) \# h_2 - \partial(\nabla_V^* h_2) \# h_1$, and in particular $\ker(\nabla_V^*)$ is closed under Lie brackets.

Two ingredients for the Riemannian metric

In order to define the Riemannian metric on the tangent space at V, we need two conditions on V. We will worry later about checking when these are true.

Condition 1

There exists a unique non-commutative law μ_V satisfying the *Dyson-Schwinger equation* $\mu_V[\nabla_V^* f] = 0$ for $f \in C_{\mathrm{tr}}(\mathbb{R}^{*d})^d$.

Note that $\nabla_V^* f$ is a *scalar-valued* function approximated by trace polynomials, and $\mu_V[\nabla_V^* f]$ is evaluated as $\nabla_V^* f^{\mathcal{A},\tau}(X)$ for any X with $\lambda_X = \mu_V$.

Condition 2

The operator $-L_V = \nabla_V^* \nabla : \mathrm{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d})) \to \mathrm{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d}))$ has kernel equal to the constant functions, and it has a continuous pseudo-inverse $\Psi_V : \mathrm{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d})) \to \mathrm{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d}))$ with $\mu(\Psi_V f) = 0$ and $-\Psi_V L_V f + \mu_V(f) = f$.

The Riemannian metric

Definition

If V satisfies Conditions 1 and 2, the Riemannian metric on $T_V \mathscr{W}(\mathbb{R}^{*d})$ is given by

$$\langle \dot{V}_1, \dot{V}_2 \rangle_V = \mu_V[\langle \nabla \Psi_V \dot{V}_1, \nabla \Psi_V V_2 \rangle_{\mathrm{tr}}].$$

Remark

This definition relates to the Riemannian metric for measures on $M_N(\mathbb{C})_{\mathrm{sa}}^d$. If $\mu_V^{(N)}$ is the measure with density constant times $e^{-N^2 V^{M_N(\mathbb{C}),\mathrm{tr}_N}}$, then the classical Riemannian metric can be expressed as

$$\int \langle \nabla (L_{V^{(N)}})^{-1} \dot{V}_1, \nabla (L_{V^{(N)}})^{-1} \dot{V}_1 \rangle_{\mathrm{tr}_N} \, d\mu^{(N)} = N^2 \int \dot{V}_1 (L_{V^{(N)}})^{-1} \dot{V}_2 \, d\mu^{(N)}.$$

The expression on the right-hand side seems simpler, but it is dimension-dependent!!

Consequences of Dyson-Schwinger equation

If Conditions 1 and 2 hold for some V, then using the formula for $\nabla_V^*[h_1, h_2]$, one can show that $\ker(\nabla_V^*)$ and $\operatorname{Im}(\nabla)$ are orthogonal with respect to V.

Furthermore, $\nabla \Psi_V \nabla_V^* : C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d})^d \to C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d})^d$ defines a projection onto the space of gradients, allowing us to decompose $C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d})$ into $\operatorname{Im}(\nabla) \oplus \ker(\nabla_V^*)$. The complementary projection is known as the *Leray projection*.

Remark

In the classical setting, the decomposition of vector fields into $\ker(\nabla_V^*)$ and $\operatorname{Im}(\nabla)$ is an infinitesimal version of Brenier's factorization of a diffeomorphism into an optimal transport map and a μ_V-preserving transformation.

Warnings

Although Condition 1 stipulates that μ_V is uniquely determined by V, there are many cases where V is *not* uniquely determined by V. For instance, since μ_V is the law of a d-tuple bounded operators (it is "supported on a operator norm ball"), often modifying V outside an operator norm ball will not change μ_V.

Another way in which degeneracy arises is from the use of trace polynomials. If a particular $(\mathcal{A}, \tau)$ and X are given, and if f is a trace polynomial, then $f^{\mathcal{A},\tau}(X)$ agrees with $p(X)$ for some non-commutative polynomial p. We can easily imagine that many V lead to the same μ for this reason.

Relatedly, the Riemannian metric on the tangent space could have a very large kernel because when we take the inner product in $L^2(\mu_V)$, all the $\mathrm{tr}(p)$ terms are collapsed to constants.

Construction of transport

Closely related to the previous observation about the differential of the transport action, we have:

Lemma

Suppose that $t \mapsto V_t$ is a C^1 path in $\mathscr{W}(\mathbb{R}^{*d})$, for t in some interval containing 0. Let h_t be a vector field with ∂h_t uniformly bounded and $-\nabla^*_{V_t} h_t = \dot{V}_t$. Let f_t be the flow along the vector field h_t. Then $(f_t)_* V_0 = V_t$.

Suppose we are given the path $t \to V_t$ (perhaps interpolating between some given V_0 and V_1) and we want to construct h_t. If each V_t satisfies Conditions 1 and 2, then we can take $h_t = -\Psi_V \nabla \dot{V}_t$. For ∂h_t to be bounded, we require some concrete estimate on Ψ_V. For h_t to depend continuously on t, we need some joint continuous dependence of $\Psi_V f$ on V and f, at least for some family of V's that contains our given path. If these conditions are met, then *some* smooth transport exists.

Construction of transport

The following theorem is similar to previous work such as Guionnet-Shlyakhtenko 2009, Dabrowski-Guionnet-Shlyakhtenko 2016.

Theorem A

Fix C_1, C_2, $C_3 > 0$ with $C_2 < 1$. Consider $V \in \operatorname{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d}))$ such that $\|\nabla V\|_{BC_{\mathrm{tr}}} \leq C_1$ and $\|\partial \nabla V - \mathrm{Id}\|_{BC_{\mathrm{tr}}} \leq C_2$.

- V satisfies Conditions 1 and 2.
- For such V, the map $(V, f) \mapsto \Psi_V f$ is jointly continuous with respect to the Fréchet topology on C_{tr}^∞.
- Let $k \geq 0$. If V is as above and furthermore $\partial^j V$ is bounded by some constant C_j for $j \leq k+2$, then Ψ_V maps BC_{tr}^k into BC_{tr}^k.

The theorem implies that for a path $t \mapsto V_t$, if ∇V_t, $\partial \nabla V_t$, $\partial^2 \nabla V_t$, $\nabla \dot{V}_t$, $\partial \nabla \dot{V}_t$ are uniformly bounded, with $\|\partial \nabla V_t - \mathrm{Id}\|_{BC_{\mathrm{tr}}} \leq C_2 < 1$, then the above construction of transport works.

Construction of transport

From this result, one immediately gets isomorphisms of C^*-algebras associated to the non-commutative laws μ_{V_t}.

Theorem B

For a path $t \mapsto V_t$ satisfying the conditions on the previous slide, there exists a C^1 path $t \mapsto f_t$ of diffeomorphisms with $(f_t)_* V_0 = V_t$. These give rise to isomorphisms between the tracial C^*-algebras (and the von Neumann algebras) associated to the GNS representations of the non-commutative laws μ_{V_t}. In particular, when V is as in the previous theorem, the C^*-algebra of μ_V is isomorphic to the one generated by a free semicircular family.

There is one thing to check to finish the proof: If $f_* V_0 = V_1$, then does $f_* \mu_{V_0} = \mu_{V_1}$? For the potentials V_t as in the previous slide, this can be checked from the free entropy viewpoint, which will be explained later.

Warnings

These results are not true for arbitrary V, even in the one variable case. Indeed, as in Biane-Speicher 1999, consider $V^{\mathcal{A},\tau}(X) = \tau(f(X))$ where $f : \mathbb{R} \to \mathbb{R}$ is a "double well" potential. If the wells are deep enough, then in the large N limit the spectral distribution is supported on a disjoint union of two intervals. Hence, the C^*-algebra is $\cong C[0,1] \oplus C[0,1]$, but the C^*-algebra obtained in the semicircular case is $\cong C[0,1]$.

Actually, Condition 1 fails for a such a potential because other measures satisfying the Dyson-Schwinger equation are obtained by reweighting the two components.

Relatedly, there are non-constant smooth functions such that $L_V \phi$ vanishes in $L^2(\mu_V)$. Namely, we take $\phi(X) = \tau(f(X))$ where $f : \mathbb{R} \to \mathbb{R}$ is constant on each of the two intervals and is smooth. On the other hand, $L_V \phi$ is not zero in $\mathrm{tr}(C_{\mathrm{tr}}^\infty(\mathbb{R}^{*d}))$, but the significance of this is unclear.

Inversion of the Laplacian

Theorem A comes out of two sets of tools:

1. The free entropy approach is used to show existence of a non-commutative law μ satisfying $\mu[\nabla_V^* f] = 0$ for $f \in C_{\mathrm{tr}}(\mathbb{R}^{*d})^d$.
2. The heat semigroup is used to uniqueness of a non-commutative law μ satisfying $\mu[L_V \phi] = 0$ for $\phi \in \mathrm{tr}(C_{\mathrm{tr}}(\mathbb{R}^{*d}))$ as well as constructing Ψ_V.

Let us start with (2). The broad outline is the same as Dab.-Gui.-Shl. 2016, but with different function spaces.

Inversion of the Laplacian

Recalling that $-L_V = \nabla_V^* \nabla$, the heat semigroup is the family of operators e^{tL_V} for $t \geq 0$. The rigorous definition is through free SDE theory. We set

$$[e^{tL_V/2} f]^{\mathcal{A},\tau}(X) = E_{\mathcal{A}}[f^{\mathcal{A}*\mathcal{B},\tau*\sigma}(\mathcal{X}_t(X))],$$

where

$$d\mathcal{X}_t(X) = d\mathcal{S}_t - \frac{1}{2}\nabla V(\mathcal{X}_t(X))\,dt, \qquad \mathcal{X}_0(X) = X,$$

$\mathcal{S}_t$ is a semicircular Brownian motion freely independent of $(\mathcal{A},\tau)$ and hence of the initial condition X, and $(\mathcal{B},\sigma)$ is the tracial C^*-algebra generated by the Brownian motion.

The assumption that $\|\partial \nabla V - \mathrm{Id}\|_{BC_{\mathrm{tr}}} \leq C_2 < 1$ implies that $\partial_X \mathcal{X}_t$ decays like $e^{-t(1-C_2)/2}$ as $t \to \infty$.

Inversion of the Laplacian

This in turn implies $\partial[e^{tL_V} f]$ decays like $e^{-t(1-C_2)}$. To recover the non-commutative law μ_V and the pseudo-inverse Ψ_V, we argue that

$$\mu_V f = \lim_{t\to\infty} e^{tL_V} f$$

and

$$\Psi_V f = \int_0^\infty [e^{tL_V} f - \mu_V(f)]\, dt.$$

These expressions make sense because of the exponential decay.

The smoothness properties as well as the continuous dependence of $\Psi_V f$ on (V, f) are proved by studying the smoothness properties of $\mathcal{X}_t(X)$ as a function of X, with some simpleminded inductive arguments.

Free Gibbs laws — results

A *free Gibbs law for V* is a non-commutative law μ that maximizes $\chi_V(\mu) := \chi(\mu) - \mu(V)$, where χ is the free microstate entropy.

We can show the following:

1. If $V \in \mathscr{W}(\mathbb{R}^{*d})$ with ∂V and $\partial^2 V$ bounded, then a free Gibbs law always exists.
2. Due to the change of variables formula for entropy, any free Gibbs law μ must satisfy the Dyson-Schwinger equation $\mu[\nabla_V^* f] = 0$.
3. Fix $C_1,\ C_2 > 0$. The set of V which have a *unique* free Gibbs law is generic in the set $\mathscr{V}_{C_1,C_2}$ of V with $\|\partial V\|_{BC_{\text{tr}}} \leq C_1$ and $\|\partial^2 V\|_{BC_{\text{tr}}} \leq C_2$, equipped with the subspace topology from $\text{tr}(C_{\text{tr}}(\mathbb{R}^{*d}))$.

Free Gibbs laws — proof with lies

The argument for the existence of free Gibbs laws relies on enlarging the space of laws in order to obtain more compactness. More precisely:

1. We embed the space of non-commutative laws into the dual of a Banach space $\mathcal{C}$ consisting of certain functions with quadratic growth at ∞.
2. Let $\mathcal{E} \subseteq \mathcal{C}^\star$ be the weak-$\star$ closure of the set of NC laws, and let $\mathcal{E}_r$ be the subset with "second moment" bounded by r. Then $\mathcal{E}_r$ is compact by Banach-Alaoglu.
3. χ_V is upper semi-continuous and it goes to $-\infty$ as the as the "second moment" of μ goes to ∞, and thus we get a maximizer using compactness of $\mathcal{E}_r$.
4. Using the change of variables formula for χ, we deduce that any maximizer ν satisfies the Dyson-Schwinger equation (for nice enough test functions).
5. Using the Dyson-Schwinger equation, we show iteratively that moments of ν are finite, and ultimately that $\nu \in \Sigma_{d,R}$ for some R.

Geodesic equations

Definition

The *geodesic equations* on $\mathscr{W}(\mathbb{R}^{*d})$ are the pair of equations

$$\begin{cases} \dot{V}_t = L_{V_t}\phi_t \\ \dot{\phi}_t = -\frac{1}{2}\langle \nabla\phi_t, \nabla\phi_t\rangle_{\mathrm{tr}}. \end{cases}$$

These can be obtained formally as the large N limit of the geodesic equations for measures on $M_N(\mathbb{C})_{\mathrm{sa}}^d$.

Thinking about the classical case, one is led to conjecture that nice enough solutions must have the form

$$V_t = (\mathrm{id} + t\nabla\dot{\phi}_0)_* V_0.$$

It is straightforward to check that when $\partial\nabla\dot{\phi}_0$ is bounded, this formula defines a solution for small enough t. We do not show rigorously that these are the only solutions.

Towards free optimal transport

However, we can show rigorously that these paths minimize length with respect to the L^2-coupling distance when $\partial \nabla \dot{\phi}_0$ is bounded by a constant C and when $t \in (0, 1/C)$. This follows from the more general proposition below.

Definition

For two non-commutative laws μ and ν, we define $d_W(\mu, \nu)$ as the infimum of $\|X - Y\|_{\tau,2}$ over all tracial C^*-algebras $(\mathcal{A}, \tau)$ and $X, Y \in \mathcal{A}_{\mathrm{sa}}^d$ such that $\lambda_X = \mu$ and $\lambda_Y = \nu$.

Proposition

Let $\phi \in \operatorname{tr}(C_{\mathrm{tr}}^2(\mathbb{R}^{*d}))_{\mathrm{sa}}$ such that $\|\partial \nabla \phi - \mathrm{Id}\|_{BC_{\mathrm{tr}}} < 1$. Then for every $(\mathcal{A}, \tau)$ and $X \in \mathcal{A}_{\mathrm{sa}}^d$, we have $d_W(\lambda_X, \lambda_{\nabla \phi(X)}) = \|X - \nabla \phi(X)\|_{\tau,2}$. In other words, X and $\nabla \phi(X)$ are an optimal coupling of their respective laws.

Towards free optimal transport

The proof of the proposition is inspired by the classical Monge-Kantorovich duality.

By the inverse function theorem, $\nabla\phi$ has an inverse function, so define ψ by

$$\psi^{\mathcal{A},\tau}(Z) = \langle Z, ((\nabla\phi)^{-1})^{\mathcal{A},\tau}(Z)\rangle_\tau - (\phi \circ (\nabla\phi)^{-1})^{\mathcal{A},\tau}(Z).$$

for all $(\mathcal{A},\tau)$. Note that $Y = ((\nabla\phi)^{-1})^{\mathcal{A},\tau}(Z)$ maximizes the function

$$\langle Z, Y\rangle_\tau - \phi^{\mathcal{A},\tau}(Y)$$

by calculus and by convexity of ϕ. (So ψ is the Legendre transform of ϕ.) Thus,

$$\phi^{\mathcal{A},\tau}(Y) + \psi^{\mathcal{A},\tau}(Z) \geq \langle Y, Z\rangle_\tau$$

for all $Y, Z \in \mathcal{A}_{\mathrm{sa}}^d$.

Towards free optimal transport

Now fix $(\mathcal{A}, \tau)$ and $X \in \mathcal{A}_{\text{sa}}^d$ as in the proposition. If Y, Z is any coupling of λ_X and $\lambda_{\nabla\phi(X)}$ on some other tracial C^*-algebra $(\mathcal{A}', \tau')$, then

$$\begin{aligned} \langle Y, Z \rangle_{\tau'} &\leq \phi^{\mathcal{A}',\tau'}(Y) + \psi^{\mathcal{A}',\tau'}(Z) \\ &= \phi^{\mathcal{A},\tau}(X) + \psi^{\mathcal{A},\tau}(\nabla\phi^{\mathcal{A},\tau}(X)) \\ &= \langle X, \phi^{\mathcal{A},\tau}(X) \rangle_{\tau}, \end{aligned}$$

where the last inequality follows by the definition of ψ. Thus, $(X, \nabla\phi(X))$ is a coupling that maximizes the inner product between the first and second variable, which is equivalent to minimizing the L^2 distance (since the L^2 norms of Y and Z are uniquely determined by the laws).