Duality for optimal couplings in free probability

David Jekel (joint work with Wilfrid Gangbo, Kyeongsik Nam, and Dimitri Shlyakhtenko)

University of California, San Diego

Berkeley Free Probabilistic Operator Algebras Seminar June 30, 2021

Acknowledgements

Land Acknowledgement: UCSD stands on the land of the Kumeyaay people.

Funding Acknowledgement: D.J. was supported by the NSF postdoc grant DMS-2002826. W.G. was supported by DMS-1700202 and U.S. Air Force grant FA9550-18-1-0502. D.S. was supported by NSF grant DMS-1762360.

Inspiration acknowledgement: We thank Alice Guionnet, Yoann Dabrowski, Wuchen Li, Ben Hayes, and Srivatsav Kunnawalkam Elayavalli for useful conversations.

Thanks to the conference organizers!

The weak-$*$ topology and the Wasserstein distance

Let $\mathcal{P}(\Omega)$ be the space of Borel measures on a compact metric space Ω (e.g. a compact subset of $\mathbb{R}^m$). Recall that $\mathcal{P}(\Omega)$ is compact in the weak-$*$ topology (as a subset of the dual of $C(\Omega)$). Moreover,

1. The Wasserstein distance metrizes the weak-$*$ topology on $\mathcal{P}(\Omega)$.
2. Finitely supported probability measures on Ω are weak-$*$ dense in $\mathcal{P}(\Omega)$ (hence also dense in the Wasserstein distance).
3. In particular, $(\mathcal{P}(\Omega), d_W)$ is a compact (and hence separable) metric space.

The weak-$*$ topology and the Wasserstein distance

Let $\mathcal{P}(\Omega)$ be the space of Borel measures on a compact metric space Ω (e.g. a compact subset of $\mathbb{R}^m$). Recall that $\mathcal{P}(\Omega)$ is compact in the weak-$*$ topology (as a subset of the dual of $C(\Omega)$). Moreover,

1. The Wasserstein distance metrizes the weak-$*$ topology on $\mathcal{P}(\Omega)$.
2. Finitely supported probability measures on Ω are weak-$*$ dense in $\mathcal{P}(\Omega)$ (hence also dense in the Wasserstein distance).
3. In particular, $(\mathcal{P}(\Omega), d_W)$ is a compact (and hence separable) metric space.

We will show that (1) - (3) fail in the non-commutative setting.

The weak-$*$ topology and the Wasserstein distance

Let $\mathcal{P}(\Omega)$ be the space of Borel measures on a compact metric space Ω (e.g. a compact subset of $\mathbb{R}^m$). Recall that $\mathcal{P}(\Omega)$ is compact in the weak-$*$ topology (as a subset of the dual of $C(\Omega)$). Moreover,

1. The Wasserstein distance metrizes the weak-$*$ topology on $\mathcal{P}(\Omega)$.
2. Finitely supported probability measures on Ω are weak-$*$ dense in $\mathcal{P}(\Omega)$ (hence also dense in the Wasserstein distance).
3. In particular, $(\mathcal{P}(\Omega), d_W)$ is a compact (and hence separable) metric space.

We will show that (1) - (3) fail in the non-commutative setting.

In the sequel, the "weak-$*$ topology" on $\Sigma_{m,R}$ refers to the topology of pointwise convergence as linear functionals on $\mathbb{C}\langle t_1, \dots, t_m \rangle$ and the "Wasserstein topology" refers to the topology induced by the Wasserstein metric.

Non-separability of the NC Wasserstein space

Proposition

For $m > 1$, while $\Sigma_{m,R}$ is compact in the weak-$*$ topology, it is *not separable* with respect to the Wasserstein distance. In particular, the weak-$*$ topology and the Wasserstein topology are different.

Non-separability of the NC Wasserstein space

Proposition

For $m > 1$, while $\Sigma_{m,R}$ is compact in the weak-$*$ topology, it is *not separable* with respect to the Wasserstein distance. In particular, the weak-$*$ topology and the Wasserstein topology are different.

To prove this for all $m > 1$, it suffices to prove it for *some* $m > 1$.

1. If it holds for m, then it also holds for $n \geq m$ because we can map $\Sigma_{m,R} \to \Sigma_{n,R}$ is isometrically by sending the law of X to the law of $(X, 0)$.
2. If (A, τ) is a tracial von Neumann algebra generated by self-adjoints $X_1, \dots, X_m$, then $M_m(A)$ is generated by two elements, $\tilde{X} = \operatorname{diag}(4R + X_1, 8R + X_2, \dots, 4mR + X_m)$ and Y where e^{iY} is a cyclic permutation matrix. One can estimate the Wasserstein distance for the original $X_1, \dots, X_m$ variables in terms of the Wasserstein distance for the $\tilde{X}$, Y variables. Hence, we can obtain the $m = 2$ case.

Non-separability of the NC Wasserstein space

Now we quote the following result:

Theorem (Gromov, Olshanskii, Ozawa)

There exists a group G with property (T) and an uncountable family $(G_\alpha)_{\alpha \in I}$ of non-isomorphic quotients of G.

Non-separability of the NC Wasserstein space

Now we quote the following result:

Theorem (Gromov, Olshanskii, Ozawa)

There exists a group G with property (T) and an uncountable family $(G_\alpha)_{\alpha \in I}$ of non-isomorphic quotients of G.

Pick such a group G. Property (T) means that there exists $C > 0$ and g_1, …, $g_m \in G$ such that for every unitary representation π of G on a Hilbert space H,

$$\|\xi - P_{\text{invariant}}\xi\| \leq C \left(\sum_{j=1}^{m} \|\pi(g_j)\xi - \xi\|^2 \right)^{1/2} \quad \text{for } \xi \in H,$$

where $P_{\text{invariant}}$ is the projection onto the subspace of invariant vectors.

Non-separability of the NC Wasserstein space

Let $q_\alpha : G \to G_\alpha$ be the quotient map. Let $X_\alpha \in L(G_\alpha)_{sa}^{2m}$ be given by the real and imaginary parts of $q_\alpha(g_1)$, ..., $q_\alpha(g_m)$.

Non-separability of the NC Wasserstein space

Let $q_\alpha : G \to G_\alpha$ be the quotient map. Let $X_\alpha \in L(G_\alpha)_{\mathrm{sa}}^{2m}$ be given by the real and imaginary parts of $q_\alpha(g_1), \dots, q_\alpha(g_m)$.

Fix α, β, and let $\epsilon = d_W(\lambda_{X_\alpha}, \lambda_{X_\beta})$. There exists some tracial von Neumann algebra $\mathcal{A}$ and trace-preserving embeddings $\iota_\alpha : L(G_\alpha) \to \mathcal{A}$ and $\iota_\beta : L(G_\beta) \to \mathcal{A}$ such that $\|\iota_\alpha(X_\alpha) - \iota_\beta(X_\beta)\|_{L^2(\mathcal{A})_{\mathrm{sa}}^m} = \epsilon$. Take $H = L^2(\mathcal{A})$ and $\pi(g)h = \iota_\alpha(q_\alpha(g)) h \iota_\beta(q_\beta(g))^{-1}$. Let $\xi = 1$ in H.

$$\left(\sum_{j=1}^m \|\pi(g_j)\xi - \xi\|^2 \right)^{1/2} = \left(\sum_{j=1}^m \|\iota_\alpha(q_\alpha(g_j)) - \iota_\beta(q_\beta(g_j))\|_{L^2(\mathcal{A})}^2 \right)^{1/2}$$
$$= \|\iota_\alpha(X_\alpha) - \iota_\beta(X_\beta)\|_{L^2(\mathcal{A})_{\mathrm{sa}}^m} = \epsilon.$$

Non-separability of the NC Wasserstein space

Hence, there exists an invariant vector η with $\|\eta - \xi\| \leq C\epsilon$. For $g \in G$,

$$\begin{aligned}
|\delta_{q_\alpha(g)=e} - \delta_{q_\beta(g))=e}| &= |\tau_{\mathcal{A}}(\iota_\alpha(q_\alpha(g)) - \iota_\beta(q_\beta(g)))| \\
&\leq \|\iota_\alpha(q_\alpha(g))\xi - \xi\iota_\beta(q_\beta(g))\| \\
&= \|\pi(g)\xi - \xi\| \\
&\leq \|\pi(g)\xi - \pi(g)\eta\| + \|\pi(g)\eta - \eta\| + \|\eta - \xi\| \\
&= 2\|\eta - \xi\| \\
&\leq 2C\epsilon.
\end{aligned}$$

Since G_α and G_β are non-isomorphic, there exists some $g \in G$ such that $q_\alpha(g) = e$ and $q_\beta(g) \neq e$ or vice versa. Hence, $1 \leq 2C\epsilon$, or $d_W(\lambda_{X_\alpha}, \lambda_{X_\beta}) \geq 1/2C$.

Non-separability of the NC Wasserstein space

Hence, there exists an invariant vector η with $\|\eta - \xi\| \leq C\epsilon$. For $g \in G$,

$$\begin{aligned} |\delta_{q_\alpha(g)=e} - \delta_{q_\beta(g))=e}| &= |\tau_{\mathcal{A}}(\iota_\alpha(q_\alpha(g)) - \iota_\beta(q_\beta(g)))| \\ &\leq \|\iota_\alpha(q_\alpha(g))\xi - \xi\iota_\beta(q_\beta(g))\| \\ &= \|\pi(g)\xi - \xi\| \\ &\leq \|\pi(g)\xi - \pi(g)\eta\| + \|\pi(g)\eta - \eta\| + \|\eta - \xi\| \\ &= 2\|\eta - \xi\| \\ &\leq 2C\epsilon. \end{aligned}$$

Since G_α and G_β are non-isomorphic, there exists some $g \in G$ such that $q_\alpha(g) = e$ and $q_\beta(g) \neq e$ or vice versa. Hence, $1 \leq 2C\epsilon$, or $d_W(\lambda_{X_\alpha}, \lambda_{X_\beta}) \geq 1/2C$.

Thus, we have an uncountable family of non-commutative laws which are separated by a distance of $1/2C$, which implies that $\Sigma_{2m,R}$ is not separable.

Connes-embeddability

Next, we will show that (2) in our list of properties of classical probability measures fails: Due the Ji-Natarajan-Vidick-Wright-Yuen's negative answer to the Connes embedding problem, laws of m-tuples in finite-dimensional algebras are not weak-$*$ dense in the space of all non-commutative laws.

Connes-embeddability

Next, we will show that (2) in our list of properties of classical probability measures fails: Due the Ji-Natarajan-Vidick-Wright-Yuen's negative answer to the Connes embedding problem, laws of m-tuples in finite-dimensional algebras are not weak-$*$ dense in the space of all non-commutative laws.

We first recall some background on ultraproducts.

Ultrafilters

The *Stone-Čech compactification of* $\mathbb{N}$ is a compact space $\beta\mathbb{N}$ containing $\mathbb{N}$, such that every map from $\mathbb{N}$ into a compact space Ω extends uniquely to a continuous function $\beta\mathbb{N} \to \Omega$. This property determines the space $\beta\mathbb{N}$ up to canonical homeomorphism. One construction of $\beta\mathbb{N}$ is as a space of ultrafilters; hence, we refer to the points of $\beta\mathbb{N} \setminus \mathbb{N}$ as *free ultrafilters*.

The universal property of $\beta\mathbb{N}$ implies that whenever $\mathcal{U}$ is a free ultrafilter and $(x_n)_{n\in\mathbb{N}}$ is a sequence in a compact space Ω, the limit $\lim_{n\to\mathcal{U}} x_n$ exists in Ω. In particular, any bounded sequence of complex numbers has a limit along $\mathcal{U}$.

Let $\mathcal{A}_n = (A_n, \tau_n)$ be a sequence of tracial von Neumann algebras. Let $\prod_{n \in \mathbb{N}} A_n$ denote the set of sequences $(x_n)_{n \in \mathbb{N}}$ such that $x_n \in A_n$ and $\sup_n \|x_n\| < \infty$. For a free ultrafilter $\mathcal{U}$ on $\mathbb{N}$, let

$$I_{\mathcal{U}} = \{(x_n)_{n \in \mathbb{N}} \in \prod_{n \in \mathbb{N}} A_n : \lim_{n \to \mathcal{U}} \|x_n\|_{L^2(\mathcal{A}_n)} = 0\},$$

and define

$$\tau_{\mathcal{U}} : \prod_{n \to \mathcal{U}} A_n \to \mathbb{C} : \tau_{\mathcal{U}}((x_n)_{n \in \mathbb{N}}) = \lim_{n \to \mathcal{U}} \tau_n(x_n).$$

One can check that $I_{\mathcal{U}}$ is a $*$-ideal in $\prod_{n \in \mathbb{N}} A_n$ and $\tau_{\mathcal{U}}$ vanishes on $I_{\mathcal{U}}$. Hence,

$$A_{\mathcal{U}} := \prod_{n \in \mathbb{N}} A_n / I_{\mathcal{U}}$$

is a $*$-algebra and $\tau_{\mathcal{U}}$ defines a trace on $\mathcal{U}$. In fact, $(A_{\mathcal{U}}, \tau_{\mathcal{U}})$ turns out to be a tracial von Neumann algebra. We denote it by $\prod_{n \to \mathcal{U}} \mathcal{A}_n$, and we call it the *ultraproduct of* $(\mathcal{A}_n)_{n \in \mathbb{N}}$. If $(x_n)_{n \in \mathbb{N}} \in \prod_{n \in \mathbb{N}} A_n$, we denote its equivalence class in $A_{\mathcal{U}}$ by $[x_n]_{n \in \mathbb{N}}$.

Ultraproducts and convergence in law

Lemma (Folklore)

Let $\mathcal{A}_n$ and $\mathcal{A}$ be tracial W^*-algebras, $X_n \in L^\infty(\mathcal{A}_n)_{\mathrm{sa}}^m$ with $\|X_n\|_{L^\infty} \leq R$ and $X \in L^\infty(\mathcal{A})_{\mathrm{sa}}^m$ with $\mathcal{A} = \mathrm{W}^*(X)$. Let $\mathcal{U}$ be a free ultrafilter on $\mathbb{N}$. TFAE:

1. $\lim_{n \to \mathcal{U}} \lambda_{X_n} = \lambda_X$ weak-$*$.
2. There exists a trace-preserving embedding $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ such that $\phi(X) = [X_n]_{n \in \mathbb{N}}$.

Ultraproducts and convergence in law

Lemma (Folklore)

Let $\mathcal{A}_n$ and $\mathcal{A}$ be tracial W^*-algebras, $X_n \in L^\infty(\mathcal{A}_n)_{sa}^m$ with $\|X_n\|_{L^\infty} \leq R$ and $X \in L^\infty(\mathcal{A})_{sa}^m$ with $\mathcal{A} = \mathrm{W}^*(X)$. Let $\mathcal{U}$ be a free ultrafilter on $\mathbb{N}$. TFAE:

1. $\lim_{n\to\mathcal{U}} \lambda_{X_n} = \lambda_X$ weak-$*$.
2. There exists a trace-preserving embedding $\phi : \mathcal{A} \to \prod_{n\to\mathcal{U}} \mathcal{A}_n$ such that $\phi(X) = [X_n]_{n\in\mathbb{N}}$.

(1) $\implies$ (2). Let $Y = [X_n]_{n\in\mathbb{N}}$ in the ultraproduct. For every polynomial p, $\tau_{\mathcal{U}}(p(Y)) = \lim_{n\to\mathcal{U}} \tau_{\mathcal{A}_n}(p(X_n)) = \tau_{\mathcal{A}}(p(X))$, hence $\mathrm{W}^*(Y) \cong \mathrm{W}^*(X) = \mathcal{A}$.

Ultraproducts and convergence in law

Lemma (Folklore)

Let $\mathcal{A}_n$ and $\mathcal{A}$ be tracial W^*-algebras, $X_n \in L^\infty(\mathcal{A}_n)_{sa}^m$ with $\|X_n\|_{L^\infty} \leq R$ and $X \in L^\infty(\mathcal{A})_{sa}^m$ with $\mathcal{A} = \mathrm{W}^*(X)$. Let $\mathcal{U}$ be a free ultrafilter on $\mathbb{N}$. TFAE:

1. $\lim_{n \to \mathcal{U}} \lambda_{X_n} = \lambda_X$ weak-$*$.
2. There exists a trace-preserving embedding $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ such that $\phi(X) = [X_n]_{n \in \mathbb{N}}$.

(1) $\implies$ (2). Let $Y = [X_n]_{n \in \mathbb{N}}$ in the ultraproduct. For every polynomial p, $\tau_{\mathcal{U}}(p(Y)) = \lim_{n \to \mathcal{U}} \tau_{\mathcal{A}_n}(p(X_n)) = \tau_{\mathcal{A}}(p(X))$, hence $\mathrm{W}^*(Y) \cong \mathrm{W}^*(X) = \mathcal{A}$.
(2) $\implies$ (1). For every polynomial p,

$$\lim_{n \to \mathcal{U}} \lambda_{X_n}(p) = \lim_{n \to \mathcal{U}} \tau_{\mathcal{A}_n}(p(X_n)) = \tau_{\mathcal{U}}(p(Y)) = \tau_{\mathcal{A}}(p(X)) = \lambda_X(p).$$

Connes-embeddability

Definition (based on Connes 1976)

A weak-$*$ separable tracial von Neumann algebra $\mathcal{A}$ is *Connes-embeddable* if there exists a trace-preserving embedding from $\mathcal{A}$ into an ultraproduct of finite-dimensional tracial $*$-algebras.

Connes-embeddability

Definition (based on Connes 1976)

A weak-$*$ separable tracial von Neumann algebra $\mathcal{A}$ is *Connes-embeddable* if there exists a trace-preserving embedding from $\mathcal{A}$ into an ultraproduct of finite-dimensional tracial $*$-algebras.

Lemma

Let $\Sigma_{m,R}^{\text{fin}}$ be the set of non-commutative laws of m-tuples from finite-dimensional tracial $*$-algebras. Let $\mathcal{A}$ be a tracial W^*-algebra and $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ such that $\|X\|_{L^\infty} \leq R$ and $\mathcal{A} = \mathrm{W}^*(X)$. Then $\mathcal{A}$ is Connes-embeddable if and only if λ_X is in the weak-$*$ closure of $\Sigma_{m,R}^{\text{fin}}$.

Connes-embeddability

Definition (based on Connes 1976)

A weak-$*$ separable tracial von Neumann algebra $\mathcal{A}$ is *Connes-embeddable* if there exists a trace-preserving embedding from $\mathcal{A}$ into an ultraproduct of finite-dimensional tracial $*$-algebras.

Lemma

Let $\Sigma_{m,R}^{\text{fin}}$ be the set of non-commutative laws of m-tuples from finite-dimensional tracial $*$-algebras. Let $\mathcal{A}$ be a tracial W^*-algebra and $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ such that $\|X\|_{L^\infty} \leq R$ and $\mathcal{A} = \mathrm{W}^*(X)$. Then $\mathcal{A}$ is Connes-embeddable if and only if λ_X is in the weak-$*$ closure of $\Sigma_{m,R}^{\text{fin}}$.

The $\Longleftarrow$ direction is immediate from the previous lemma. For $\Longrightarrow$, we pick an ultraproduct embedding $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$, and use the fact that $\phi(X) = [X_n]_{n \in \mathbb{N}}$ where $\|X_n\|_{L^\infty(\mathcal{A}_n)^m} \leq R$ (using general facts about C^*-algebras / functional calculus).

The weak-$*$ and Wasserstein topologies again

Hence, failure of Connes-embedding implies that the laws coming from finite-dimensional algebras are not weak-$*$ dense.

The weak-$*$ and Wasserstein topologies again

Hence, failure of Connes-embedding implies that the laws coming from finite-dimensional algebras are not weak-$*$ dense.

Biane and Voiculescu noted that the Wasserstein topology is at least as strong as the weak-$*$ topology, so of course they are not Wasserstein dense either.

The weak-$*$ and Wasserstein topologies again

Hence, failure of Connes-embedding implies that the laws coming from finite-dimensional algebras are not weak-$*$ dense.

Biane and Voiculescu noted that the Wasserstein topology is at least as strong as the weak-$*$ topology, so of course they are not Wasserstein dense either.

However, we want to go further and describe a more precise comparison between the weak-$*$ and Wasserstein topologies. To this end, we will describe *Wasserstein* convergence in terms of ultraproducts.

The weak-$*$ and Wasserstein topologies again

Hence, failure of Connes-embedding implies that the laws coming from finite-dimensional algebras are not weak-$*$ dense.

Biane and Voiculescu noted that the Wasserstein topology is at least as strong as the weak-$*$ topology, so of course they are not Wasserstein dense either.

However, we want to go further and describe a more precise comparison between the weak-$*$ and Wasserstein topologies. To this end, we will describe *Wasserstein* convergence in terms of ultraproducts.

We also need to use the notion of factorizable completely positive maps.

Factorizable maps

Definition (Anantharaman-Delaroche)

If $\mathcal{A}$ and $\mathcal{B}$ are tracial W^*-algebras, then a completely positive map $\Phi : \mathcal{A} \to \mathcal{B}$ is said to be *factorizable* if it is realized as a trace-preserving inclusion $\mathcal{A} \to \mathcal{C}$ followed by a trace-preserving conditional expectation $\mathcal{C} \to \mathcal{B}$ (and in this case, we say Φ factorizes through $\mathcal{C}$). We also write $\mathrm{FM}(\mathcal{A}, \mathcal{B})$ for the space of factorizable maps from $\mathcal{A}$ to $\mathcal{B}$.

Facts

1. $\mathrm{FM}(\mathcal{A}, \mathcal{B})$ is convex (prove using direct sums).
2. $\mathrm{FM}(\mathcal{A}, \mathcal{B})$ is compact in the pointwise weak-$*$ topology (prove using ultraproducts).
3. Factorizable maps are closed under composition (prove using amalgamated free products).

Optimal coupling and factorizable maps

Lemma (GJNS)

Let $X \in L^\infty(\mathcal{A})_{sa}^m$ and $Y \in L^\infty(\mathcal{B})_{sa}^m$ be non-commutative tuples. Then

$$C(\lambda_X, \lambda_Y) = \sup_{\Phi \in \mathrm{FM}(\mathcal{A},\mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{sa}^m}$$

Optimal coupling and factorizable maps

Lemma (GJNS)

Let $X \in L^\infty(\mathcal{A})_{sa}^m$ and $Y \in L^\infty(\mathcal{B})_{sa}^m$ be non-commutative tuples. Then

$$C(\lambda_X, \lambda_Y) = \sup_{\Phi \in \mathrm{FM}(\mathcal{A},\mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{sa}^m}$$

$\geq$. Suppose $\Phi \in \mathrm{FM}(\mathcal{A}, \mathcal{B})$ and Φ factorizes using $\iota_1 : \mathcal{A} \to \mathcal{C}$ and $\iota_2 : \mathcal{B} \to \mathcal{C}$. Then $C(\lambda_X, \lambda_Y) \geq \langle \iota_1(X), \iota_2(Y) \rangle_{L^2(\mathcal{C})_{sa}^m} = \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})}$.

Optimal coupling and factorizable maps

Lemma (GJNS)

Let $X \in L^\infty(\mathcal{A})^m_{\mathrm{sa}}$ and $Y \in L^\infty(\mathcal{B})^m_{\mathrm{sa}}$ be non-commutative tuples. Then

$$C(\lambda_X, \lambda_Y) = \sup_{\Phi \in \mathrm{FM}(\mathcal{A},\mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})^m_{\mathrm{sa}}}$$

$\geq$. Suppose $\Phi \in \mathrm{FM}(\mathcal{A}, \mathcal{B})$ and Φ factorizes using $\iota_1 : \mathcal{A} \to \mathcal{C}$ and $\iota_2 : \mathcal{B} \to \mathcal{C}$. Then $C(\lambda_X, \lambda_Y) \geq \langle \iota_1(X), \iota_2(Y) \rangle_{L^2(\mathcal{C})^m_{\mathrm{sa}}} = \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})}$.

$\leq$. Conversely, for any coupling of λ_X and λ_Y on $\mathcal{C}$, there are inclusions $\iota_1 : \mathrm{W}^*(X) \to \mathcal{C}$ and $\iota_2 : \mathrm{W}^*(Y) \to \mathcal{C}$. By composing $\iota_2^* \iota_1$ with the conditional expectation $\mathcal{A} \to \mathrm{W}^*(X)$ and inclusion $\mathrm{W}^*(Y) \to \mathcal{B}$, we get a factorizable map $\mathcal{A} \to \mathcal{B}$.

Corollary

$d_W(\lambda_X, \lambda_Y) = \|X\|^2_{L^2(\mathcal{A})^m} + \|Y\|^2_{L^2(\mathcal{B})^m} - 2 \sup_{\Phi \in \mathrm{FM}(\mathcal{A},\mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})}$.

Ultraproducts and Wasserstein convergence

Definition

Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be a trace-preserving embedding. A sequence of completely positive maps $[\Phi_n]_{n \in \mathbb{N}}$, where $\Phi_n : \mathcal{A} \to \mathcal{A}_n$, is said to *lift* ϕ if $\phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}}$ for all $Z \in L^\infty(\mathcal{A})$.

Ultraproducts and Wasserstein convergence

Definition

Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be a trace-preserving embedding. A sequence of completely positive maps $[\Phi_n]_{n \in \mathbb{N}}$, where $\Phi_n : \mathcal{A} \to \mathcal{A}_n$, is said to *lift* ϕ if $\phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}}$ for all $Z \in L^\infty(\mathcal{A})$.

Lemma (GJNS)

Let $\mathcal{A}_n$ and $\mathcal{A}$ be tracial W^*-algebras, $X_n \in L^\infty(\mathcal{A}_n)_{\mathrm{sa}}^m$ with $\|X_n\|_{L^\infty} \leq R$ and $X \in L^\infty(\mathcal{A})_{\mathrm{sa}}^m$ with $\mathcal{A} = \mathrm{W}^*(X)$. Let $\mathcal{U}$ be a free ultrafilter on $\mathbb{N}$. TFAE:

1. $\lim_{n \to \mathcal{U}} \lambda_{X_n} = \lambda_X$ in the Wasserstein distance.
2. There exists a trace-preserving embedding $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ such that $\phi(X) = [X_n]_{n \in \mathbb{N}}$ and ϕ lifts to a sequence of factorizable completely positive maps.

(1) $\implies$ (2). By the previous lemma, there exist tracial von Neumann algebras $\mathcal{C}_n$ and inclusions $\pi_n : \mathcal{A}_n \to \mathcal{C}_n$ and $\rho : \mathcal{A} \to \mathcal{C}_n$ such that $d_W(\lambda_{X_n}, \lambda_X) = \|\pi_n(X_n) - \rho(X)\|_{L^2(\mathcal{C}_n)}$. These induce maps

$$\pi : \prod_{n \to \mathcal{U}} \mathcal{A}_n \to \prod_{n \to \mathcal{U}} \mathcal{C}_n, \qquad \rho : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{C}_n$$

such that $\pi([X_n]_{n \in \mathbb{N}}) = \rho(X)$.

In particular, $\rho(X)$ is in the image of $\prod_{n \to \mathcal{U}} \mathcal{A}_n$, so ρ codomain-restricts to an embedding $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$, and in fact $\phi = \pi^* \rho$. The maps $\Phi_n = \pi_n^* \rho_n : \mathcal{A} \to \mathcal{A}_n$ are factorizable and lift ϕ (check this by showing that the π^* is the ultraproduct of π_n^*).

Ultraproducts and Wasserstein convergence

(2) $\implies$ (1). Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be an embedding and $\Phi_n : \mathcal{A} \to \mathcal{A}_n$ a lifting of ϕ to a sequence of factorizable maps. Let $\mathcal{M} = \prod_{n \to \mathcal{U}} \mathcal{A}_n$. Then

$$d_W(\lambda_{X_n}, \lambda_X) \leq \|X_n\|_{L^2(\mathcal{A}_n)^m}^2 + \|X\|_{L^2(\mathcal{A})^m}^2 - 2\langle X_n, \Phi_n(X)\rangle_{L^2(\mathcal{A}_n)^m}.$$

Taking the limit as $n \to \mathcal{U}$,

$$\begin{aligned}\lim_{n \to \mathcal{U}} d_W(\lambda_{X_n}, \lambda_X) &\leq \|\phi(X)\|_{L^2(\mathcal{M})_{sa}^m}^2 + \|\phi(X)\|_{L^2(\mathcal{M})_{sa}^m}^2 - 2\langle \phi(X), \phi(X)\rangle_{L^2(\mathcal{M})_{sa}^m} \\ &= 0.\end{aligned}$$

The weak-$*$ and Wasserstein topologies

Comparing the ultraproduct characterizations for weak-$*$ convergence and Wasserstein convergence gives an alternative proof that the Wasserstein topology is at least as strong as the weak-$*$ topology.

The weak-$*$ and Wasserstein topologies

Comparing the ultraproduct characterizations for weak-$*$ convergence and Wasserstein convergence gives an alternative proof that the Wasserstein topology is at least as strong as the weak-$*$ topology.

Our next goal is to characterize the points $\mu \in \Sigma_{m,R}$ at which the two topologies agree using maps into ultraproducts. In particular, we will show that if $\mathrm{W}^*(\mu)$ is Connes-embeddable, then agreement of the two topologies is equivalent to $\mathrm{W}^*(\mu)$ being amenable/hyperfinite.

Agreement of topologies

Definition

Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two topologies on a set Ω and $x \in \Omega$. We say that *$\mathcal{T}_1$ and $\mathcal{T}_2$ agree at x* if every $\mathcal{T}_1$-neighborhood of x contains a $\mathcal{T}_2$-neighborhood of x and vice versa.

Agreement of topologies

Definition

Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two topologies on a set Ω and $x \in \Omega$. We say that *$\mathcal{T}_1$ and $\mathcal{T}_2$ agree at x* if every $\mathcal{T}_1$-neighborhood of x contains a $\mathcal{T}_2$-neighborhood of x and vice versa.

Observation

If $\mathcal{T}_1$ and $\mathcal{T}_2$ are metrizable topologies on Ω, then $\mathcal{T}_1$ and $\mathcal{T}_2$ agree at $x \in \Omega$ if and only if, for every sequence $(x_n)_{n \in \mathbb{N}}$ in Ω, convergence of x_n to x with respect to $\mathcal{T}_1$ is equivalent to convergence of x_n to x with respect to $\mathcal{T}_2$.

Agreement of topologies

Definition

Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two topologies on a set Ω and $x \in \Omega$. We say that *$\mathcal{T}_1$ and $\mathcal{T}_2$ agree at x* if every $\mathcal{T}_1$-neighborhood of x contains a $\mathcal{T}_2$-neighborhood of x and vice versa.

Observation

If $\mathcal{T}_1$ and $\mathcal{T}_2$ are metrizable topologies on Ω, then $\mathcal{T}_1$ and $\mathcal{T}_2$ agree at $x \in \Omega$ if and only if, for every sequence $(x_n)_{n \in \mathbb{N}}$ in Ω, convergence of x_n to x with respect to $\mathcal{T}_1$ is equivalent to convergence of x_n to x with respect to $\mathcal{T}_2$.

Observation

In fact, fix a free ultrafilter $\mathcal{U}$ on $\mathbb{N}$. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be metrizable topologies on Ω. Then $\mathcal{T}_1$ and $\mathcal{T}_2$ agree at $x \in \Omega$, if and only if the following holds: For every sequence $(x_n)_{n \in \mathbb{N}}$, $\lim_{n \to \mathcal{U}} x_n = x$ with respect to $\mathcal{T}_1$ if and only if $\lim_{n \to \mathcal{U}} x_n = x$ with respect to $\mathcal{T}_2$.

Agreement of topologies

The ultraproduct characterizations of weak-$*$ and Wasserstein convergence lead to the following result.

Proposition

Let $\mu \in \Sigma_{m,R}$. Let $\mathcal{A}$ and $X \in L^\infty(\mathcal{A})_{sa}^m$ such that $\lambda_X = \mu$ and $\mathcal{A} = \mathrm{W}^*(X)$. TFAE:

1. The Wasserstein and weak-$*$ topologies agree at μ.
2. Every embedding of $\mathcal{A}$ into an ultraproduct of tracial von Neumann algebras lifts to a sequence of factorizable maps.

If $\mathcal{A}$ satisfies condition (2), we will say that $\mathcal{A}$ is FM-*stable*. This notion is closely related to *tracial stability* for C^*-algebras studied by Hadwin and Shulman, as well as Atkinson and Kunnawalkam Ellayavalli.

Agreement of topologies

Proposition

Let $\mu \in \Sigma_{m,R}$. Let $\mathcal{A}$ and $X \in L^\infty(\mathcal{A})_{sa}^m$ such that $\lambda_X = \mu$ and $\mathcal{A} = \mathrm{W}^*(X)$. Assume that $\mathcal{A}$ is Connes-embeddable (equivalently μ is in the weak-$*$ closure of $\Sigma_{m,R}^{\mathrm{fin}}$). TFAE:

1. $\mathcal{A}$ is approximately finite-dimensional / semi-discrete / amenable / injective (equivalent by Connes 1976).
2. $\mathcal{A}$ is FM-stable.
3. The Wasserstein and weak-$*$ topologies agree at μ.
4. μ is in the Wasserstein closure of $\Sigma_{m,R}^{\mathrm{fin}}$.

Note that (1) $\iff$ (2) can be proved the same was as in Atkinson and Kunnawalkam Ellayavalli's paper (they did it for UCP-liftings).

Agreement of topologies

(1) $\implies$ (2). Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be an embedding. Let $\mathcal{B}_n = \mathcal{A}_n * \mathcal{A} * L^\infty[0,1]$, which is a II_1-factor containing $\mathcal{A}_n$ and $\mathcal{A}$.

Agreement of topologies

(1) $\implies$ (2). Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be an embedding. Let $\mathcal{B}_n = \mathcal{A}_n * \mathcal{A} * L^\infty[0,1]$, which is a II_1-factor containing $\mathcal{A}_n$ and $\mathcal{A}$.

The map ϕ produces an embedding $\tilde{\phi} : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{B}_n$. The embeddings $\psi_n : \mathcal{A} \to \mathcal{B}_n$ produce another embedding $\psi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{B}_n$. Since $\mathcal{A}$ is AFD, any two embeddings into an ultraproduct of II_1 factors are unitarily conjugate, hence $\tilde{\phi}$ and ψ are unitarily conjugate.

Agreement of topologies

(1) $\implies$ (2). Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be an embedding. Let $\mathcal{B}_n = \mathcal{A}_n * \mathcal{A} * L^\infty[0,1]$, which is a II_1-factor containing $\mathcal{A}_n$ and $\mathcal{A}$.

The map ϕ produces an embedding $\tilde{\phi} : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{B}_n$. The embeddings $\psi_n : \mathcal{A} \to \mathcal{B}_n$ produce another embedding $\psi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{B}_n$. Since $\mathcal{A}$ is AFD, any two embeddings into an ultraproduct of II_1 factors are unitarily conjugate, hence $\tilde{\phi}$ and ψ are unitarily conjugate.

By construction, ψ lifts to a sequence of embeddings $\mathcal{A} \to \mathcal{B}_n$, hence so does $\tilde{\phi}$. By taking the conditional expectation back down onto $\mathcal{A}_n$, we get an FM-lifting of ϕ.

Agreement of topologies

(1) $\implies$ (2). Let $\phi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{A}_n$ be an embedding. Let $\mathcal{B}_n = \mathcal{A}_n * \mathcal{A} * L^\infty[0,1]$, which is a II_1-factor containing $\mathcal{A}_n$ and $\mathcal{A}$.

The map ϕ produces an embedding $\tilde{\phi} : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{B}_n$. The embeddings $\psi_n : \mathcal{A} \to \mathcal{B}_n$ produce another embedding $\psi : \mathcal{A} \to \prod_{n \to \mathcal{U}} \mathcal{B}_n$. Since $\mathcal{A}$ is AFD, any two embeddings into an ultraproduct of II_1 factors are unitarily conjugate, hence $\tilde{\phi}$ and ψ are unitarily conjugate.

By construction, ψ lifts to a sequence of embeddings $\mathcal{A} \to \mathcal{B}_n$, hence so does $\tilde{\phi}$. By taking the conditional expectation back down onto $\mathcal{A}_n$, we get an FM-lifting of ϕ.

(2) $\implies$ (3). Done already.

Agreement of topologies

(3) $\implies$ (4) We assumed μ is in the weak-$*$ closure of $\Sigma_{m,R}^{\text{fin}}$, so if the two topologies agree at μ, it is also in the Wasserstein closure.

Agreement of topologies

(3) $\implies$ (4) We assumed μ is in the weak-$*$ closure of $\Sigma^{\text{fin}}_{m,R}$, so if the two topologies agree at μ, it is also in the Wasserstein closure.

(4) $\implies$ (1). Approximating μ by laws μ_n of variables X_n in finite-dimensional tracial $*$-algebras $\mathcal{A}_n$. Let $\Phi_n : \mathcal{A} \to \mathcal{A}_n$ be the associated factorizable maps. Then $\Phi_n^* \Phi_n(X) \to X$ as $n \to \infty$. Using ultraproduct arguments similar to before, we can deduce that $\Phi_n^* \Phi_n \to \mathrm{id}$ pointwise in L^2 on $\mathcal{A}$, which means that $\mathcal{A}$ is semi-discrete by definition.

Remarks

If μ and $\mathcal{A}$ are as above and $\mathcal{A}$ is commutative, then the weak-$*$ and Wasserstein topologies agree at μ. Hence, we recover the classical result that the weak-$*$ topology and the Wasserstein topology agree on $\mathcal{P}([-R, R]^m)$.

Remarks

If μ and $\mathcal{A}$ are as above and $\mathcal{A}$ is commutative, then the weak-$*$ and Wasserstein topologies agree at μ. Hence, we recover the classical result that the weak-$*$ topology and the Wasserstein topology agree on $\mathcal{P}([-R, R]^m)$.

It is unknown whether there exist FM-stable tracial von Neumann algebras that are not Connes-embeddable. Related questions are explored by Isaac Goldbring in the paper "Non-embeddable II_1 factors resembling the hyperfinite II_1 factor."

Optimal couplings and Connes-embeddability

Theorem (Haagerup-Musat 2015)

A completely positive map $M_n(\mathbb{C}) \to M_n(\mathbb{C})$ factorizes through a Connes-embeddable tracial W^*-algebra if and only if it can be approximated by maps that factorize through finite-dimensional algebras. Moreover, the Connes-embedding problem has a positive answer if and only if every factorizable map can be approximated by those that factorize through finite-dimensional algebras.

Optimal couplings and Connes-embeddability

Theorem (Haagerup-Musat 2015)

A completely positive map $M_n(\mathbb{C}) \to M_n(\mathbb{C})$ factorizes through a Connes-embeddable tracial W^*-algebra if and only if it can be approximated by maps that factorize through finite-dimensional algebras. Moreover, the Connes-embedding problem has a positive answer if and only if every factorizable map can be approximated by those that factorize through finite-dimensional algebras.

Ji-Natarajan-Vidick-Wright-Yuen 2020 announced a negative solution to the Connes embedding problem, which would imply the following corollary.

Corollary (GJNS)

For sufficiently large n and m, there exist $X, Y \in M_n(\mathbb{C})_{\mathrm{sa}}^m$ such that every optimal coupling of λ_X and λ_Y uses an non-Connes-embeddable tracial W^*-algebra.

Optimal couplings and Connes-embeddability

Let $\mathrm{FM}_{\mathrm{app}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ be the set of completely positive maps $M_n(\mathbb{C}) \to M_n(\mathbb{C})$ that factorize through a Connes-embeddable von Neumann algebra.

Optimal couplings and Connes-embeddability

Let $\mathrm{FM}_{\mathrm{app}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ be the set of completely positive maps $M_n(\mathbb{C}) \to M_n(\mathbb{C})$ that factorize through a Connes-embeddable von Neumann algebra.

For $X, Y \in M_n(\mathbb{C})_{\mathrm{sa}}^m$, let

$$C_{\mathrm{app}}(X, Y) = \sup_{\Phi \in \mathrm{FM}_{\mathrm{app}}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\mathrm{sa}}^m}.$$

Our goal is to show that $C_{\mathrm{app}}(X, Y)$ is sometimes strictly less than $C(X, Y)$.

Optimal couplings and Connes-embeddability

$\mathrm{FM}_{\mathrm{app}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ and $\mathrm{FM}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ can be viewed as closed convex subsets of the space of real-linear maps $M_n(\mathbb{C})_{\mathrm{sa}} \to M_n(\mathbb{C})_{\mathrm{sa}}$, which is isomorphic to the space of real-linear maps $M_n(\mathbb{C})_{\mathrm{sa}} \otimes_{\mathbb{R}} M_n(\mathbb{C})_{\mathrm{sa}} \to \mathbb{R}$.

If $\Phi \in \mathrm{FM}(M_n(\mathbb{C}), M_n(\mathbb{C})) \setminus \mathrm{FM}_{\mathrm{app}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$, then by the separating hyperplane theorem, there exists $v \in M_n(\mathbb{C})_{\mathrm{sa}} \otimes_{\mathbb{R}} M_n(\mathbb{C})_{\mathrm{sa}}$ such that

$$\Phi(v) > \sup_{\Psi \in \mathrm{FM}_{\mathrm{app}}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \Psi(v).$$

Write v is a sum of simple tensors $v = \sum_{j=1}^m X_j \otimes Y_j$, and let $X = (X_1, \dots, X_m)$ and $Y = (Y_1, \dots, Y_m)$. Then

$$\Phi(v) = \sum_{j=1}^m \langle \Phi(X_j), Y_j \rangle_{L^2(M_n(\mathbb{C}), \mathrm{tr})_{\mathrm{sa}}} = \langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}), \mathrm{tr})_{\mathrm{sa}}^m}.$$

Optimal couplings and Connes-embeddability

Hence,

$$C(X, Y) \geq \langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}),\mathrm{tr})} > C_{\mathrm{app}}(X, Y).$$

This proves the proposition.

In other words, we have shown that given a negative answer to Connes embedding, there exist matrix tuples $X, Y \in M_n(\mathbb{C})_{\mathrm{sa}}^m$ for some n, m such that non-Connes-embeddable tracial von Neumann algebras are needed to witness the value of the Wasserstein distance. In particular, an optimal coupling can only occur in a non-Connes-embeddable von Neumann algebra.

Remark

Since the tensor rank of a vector in $V \otimes W$ is at most $\max(\dim V, \dim W)$, we see that in our argument we can take $m = n^2$.

Remarks/Problems

This result says that even for very simple non-commutative laws μ and ν, the algebra associated to the optimal coupling can be very gnarly. On the other hand, we showed that for X and Y in the optimal coupling $\mathrm{W}^*((1-t)X + tY) = \mathrm{W}^*(X, Y)$ for $t \in (0, 1)$, so a non-Connes-embeddable algebra can be generated by a convex combination of images of matrix tuples.

Remarks/Problems

This result says that even for very simple non-commutative laws μ and ν, the algebra associated to the optimal coupling can be very gnarly. On the other hand, we showed that for X and Y in the optimal coupling $\mathrm{W}^*((1-t)X + tY) = \mathrm{W}^*(X, Y)$ for $t \in (0,1)$, so a non-Connes-embeddable algebra can be generated by a convex combination of images of matrix tuples.

These results inspire many questions for future research:

1. Study optimal couplings for explicit examples of tuples from $M_n(\mathbb{C})$, and try to find the optimal coupling in $M_{nk}(\mathbb{C})$ for $k \in \mathbb{N}$ using a computer.
2. Give an explicit example of matrix tuples such that an optimal coupling requires an infinite-dimensional tracial von Neumann algebra (compare results of Musat and Rørdam on factorizable maps).
3. Explore the optimal coupling problem as a strategy for finding explicit counterexamples to Connes embedding.

References I

1. Wilfrid Gangbo, David Jekel, Kyeongsik Nam, and Dimitri Shlyakhtenko. Duality for optimal couplings in free probability. arXiv:2105.12351
2. Philippe Biane and Dan-Virgil Voiculescu. A free probability analogue of the Wasserstein metric on the trace-state space. Geometric and Functional Analysis, 11:1125–1138, 2001.
3. Narutaka Ozawa. There is no separable universal II_1 factor. Proc. Amer. Math. Soc., 132:487–90, 2004.
4. Alice Guionnet and Dimitri Shlyakhtenko. Free monotone transport. Inventiones Mathematicae, 197(3):613–661, 09 2014.
5. David Jekel, Wuchen Li, and Dimitri Shlyakhtenko. Tracial non-commutative smooth functions and the free Wasserstein manifold. arXiv:2101.06572, 2021.

References II

6. Claire Anantharaman-Delaroche. On ergodic theorems for free group actions on noncommutative spaces. Probab. Theory Rel. Fields, 135:520–546, 2006.
7. Magdalena Musat and Mikael Rørdam. Non-closure of quantum correlation matrices and factorizable channels that require infinite dimensional ancilla (with an appendix by Narutaka Ozawa). Comm. Math. Phys., 375:1761–1776, 2020.
8. Uffe Haagerup and Magdalena Musat. An asymptotic property of factorizable completely positive maps and the Connes embedding problem. Comm. Math. Phys., 338(2):721–752, 2015.
9. Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, Henry Yuen. MIP*=RE. arXiv:2001.04383
10. Alain Connes. Classification of injective factors. Cases II_1, II_∞, III_λ, $\lambda \neq 1$. Ann. of Math. (2), 104(1):73–115, 1976.

References III

11. Scott Atkinson and Srivatsav Kunnawalkam Ellayavalli. On ultraproduct embeddings and amenability for tracial von Neumann algebras. International Mathematics Research Notices, 2021(4):2882–2918, 2021.
12. Isaac Goldbring. Non-embeddable II_1 factors resembling the hyperfinite II_1 factor. arXiv:2101.10467