Linear Bandits beyond Inner Product Spaces

the case of Bandit Optimal Transport

Lorenzo Croissant

CREST, INRIA FairPlay

ICML 2026

THE KANTOROVICH OPTIMAL TRANSPORT PROBLEM

- Parameters: $(\mu, \nu) \in \mathscr{P}(\mathcal{M}_\mu) \times \mathscr{P}(\mathcal{M}_\nu)$, $c^* : \mathcal{M}_\mu \times \mathcal{M}_\nu \to \mathbb{R}$.
- Constraint set:

$$\Pi(\mu, \nu) := \{\pi \in \mathscr{P}(\mathcal{M}_\mu \times \mathcal{M}_\nu) : \pi(\cdot, \mathcal{M}_\nu) = \mu, \pi(\mathcal{M}_\mu, \cdot) = \nu\}$$

The Kantorovich optimal transport problem

- Parameters: $(\mu, \nu) \in \mathscr{P}(\mathcal{M}_\mu) \times \mathscr{P}(\mathcal{M}_\nu), c^* : \mathcal{M}_\mu \times \mathcal{M}_\nu \to \mathbb{R}$.
- Constraint set:

$$\Pi(\mu, \nu) := \{\pi \in \mathscr{P}(\mathcal{M}_\mu \times \mathcal{M}_\nu) : \pi(\cdot, \mathcal{M}_\nu) = \mu, \pi(\mathcal{M}_\mu, \cdot) = \nu\}$$

- The *Kantorovich* problem is

$$\text{Kant.}(\mu, \nu, c^*) := \inf_{\pi \in \Pi(\mu,\nu)} \int c^*(x, y) \mathrm{d}\pi(x, y)$$

- Applications: economics, ML, etc.

The Bandit OT problem

1. Marginals μ and ν are given but c^* is unknown.
2. One must play plans $\pi_t \in \Pi(\mu, \nu)$ sequentially over $t \in \mathbb{N}$.
3. Each turn, a bandit feedback C_t is generated

$$C_t := \int c^*(x, y) \mathrm{d}\pi_t(x, y) + \xi_t \,.$$

4. Performance of an algorithm $\boldsymbol{\pi} := (\pi_t)_{t \in \mathbb{N}}$ is evaluated online with regret

$$\mathscr{R}_T(\boldsymbol{\pi}) := \sum_{t=1}^{T} C_t - \text{Kant.}(\mu, \nu, c^*) \quad \text{for} \quad T \in \mathbb{N}$$

Linear Bandits are solved right?

$$\mathscr{R}_T(\boldsymbol{\pi}) := \sum_{t=1}^{T} c_t - \text{Kant.}(\mu, \nu, c^*) \quad \text{for} \quad T \in \mathbb{N}$$

- Optimism in the Face of Uncertainty (OFU).
  - Ridge regression + confidence sets + optimism
  - Abbasi-yadkori et al. (2011).
  - Requires Kant. to be a bilinear functional on a *Hilbert* space.

What about BOT?

The Kantorovich problem

$$\text{Kant.}(\mu, \nu, c^*) := \inf_{\pi \in \Pi(\mu,\nu)} \int c^*(x,y) \mathrm{d}\pi(x,y) = \inf_{\pi \in \Pi(\mu,\nu)} \langle c^* | \pi \rangle$$

is an infinite-dimensional linear program but

1. The actions and cost functions do not live in the same space.
2. Actions are measure-valued and measures are not a Hilbert space.

What about BOT?

The Kantorovich problem

$$\text{Kant.}(\mu, \nu, c^*) := \inf_{\pi \in \Pi(\mu,\nu)} \int c^*(x,y)\mathrm{d}\pi(x,y) = \inf_{\pi \in \Pi(\mu,\nu)} \langle c^* | \pi \rangle$$

is an infinite-dimensional linear program but

1. The actions and cost functions do not live in the same space.
2. Actions are measure-valued and measures are not a Hilbert space.

When are sub-linear rates attainable in BOT?

Solving BOT: the key idea

1. Rewrite the transport payoff in the Fourier domain so BOT *looks* like a linear bandit on transformed costs and plans.
2. Add a small entropic penalty to optimism so that plans admit densities and live in the Hilbert space where OFUL-style confidence sets apply.
3. At round t: estimate Fc^*, build a confidence set, and choose the optimistic entropic OT plan.

Main regret guarantee

Assume $c^* \in L^2(\mathbb{R}^d; \varrho)$ is Lipschitz and the noise is sub-Gaussian. With $\varepsilon_t = \alpha t^{-\alpha}$, the optimistic entropic algorithm satisfies, with probability at least $1-\delta$,

$$\mathscr{R}_T \leq \sigma\sqrt{2T \log\left(\frac{2}{\delta}\right)} + 2C\beta_T(\delta)\sqrt{T \log\det\left(\mathrm{Id} + \frac{1}{2\lambda C} M_T (\mathrm{D}\Lambda)^{-1} M_T^*\right)}$$
$$+ \frac{\kappa\alpha}{1-\alpha}\left(T^{1-\alpha}\log(T) + \frac{\alpha}{2^\alpha}\log(6)\right).$$

- First line: the usual OFUL-style exploration term.
- Second line: approximation error of Kantorovich OT by entropic OT.
- The remaining question is how the learning term $\beta_T(\delta)$ scales with regularity.

Learning + conclusion

- The learning term is controlled by how fast the Fourier coefficients of c^* decay in a basis of $L^2(\mathbb{R}^d; \varrho)$.
- If the truncation error behaves like $\zeta(n) \propto 1 - n^{-q}$, choosing a model size $n_t \asymp t^{1/(q+2)}$ gives

$$\mathscr{R}_T = \mathcal{O}\left(T^{\frac{q+4}{2q+4}}\right)$$

  up to logarithmic factors.
- As $q \to \infty$, we recover the parametric $\mathcal{O}(\sqrt{T})$ rate; as $q \downarrow 0$, the problem becomes unlearnable.
- Take-home message: BOT admits sublinear regret despite going beyond the standard Hilbert setting, and the exact rate is dictated by the regularity of c^*.

References

[plain,noframenumbering]

References

Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. Improved Algorithms for Linear Stochastic Bandits. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 214–223. PMLR, July 2017. ISSN: 2640-3498.

Sinho Chewi, Jonathan Niles-Weed, and Philippe Rigollet. Statistical optimal transport, July 2024. arXiv:2407.18163 [math, stat].

Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

Nicolas Fournier and Arnaud Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3-4): 707–738, August 2015. ISSN 0178-8051, 1432-2064. doi: 10.1007/s00440-014-0583-7.

Alfred Galichon. The unreasonable effectiveness of optimal transport in economics, July 2021. arXiv:2107.04700 [econ].

Wenxuan Guo, YoonHaeng Hur, Tengyuan Liang, and Christopher Thomas Ryan. Online Learning to Transport via the Minimal Selection Principle. In *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178, pages 4085–4109. PMLR, 2022.

Jan-Christian Hütter and Philippe Rigollet. Minimax estimation of smooth optimal transport maps. *The Annals of Statistics*, 49(2):1166–1194, April 2021. ISSN 0090-5364, 2168-8966. doi: 10.1214/20-AOS1997.

Leonid Vitaliyevich Kantorovich. On the Translocation of Masses. *Journal of Mathematical Sciences*, 133(4):1381–1382, 2006. Originally published in Dokl. Akad. Nauk SSSR, 37, No. 7-8, 227–229 (1942); translation by A.N. Sobolevskiĭ.

Vladik Kreinovich, Woraphon Yamaka, and Supanika Leurcharusmee, editors. *Applications of Optimal Transport to Economics and Related Topics*, volume 556 of *Studies in Systems, Decision and Control*. Springer Nature Switzerland, Cham, 2024. ISBN 978-3-031-67769-4 978-3-031-67770-0. doi: 10.1007/978-3-031-67770-0.

Tudor Manole, Sivaraman Balakrishnan, Jonathan Niles-Weed, and Larry Wasserman. Plugin estimation of smooth optimal transport maps. *The Annals of Statistics*, 52(3), June 2024. ISSN 0090-5364. doi: 10.1214/24-AOS2379.

Philippe Rigollet and Austin J. Stromme. On the sample complexity of entropic optimal transport, June 2022. arXiv:2206.13472 [math].

Long Tran-Thanh and Jia Yuan Yu. Functional Bandits, May 2014. arXiv:1405.2432 [stat].

Cédric Villani. Optimal Transport, old and new. volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-540-71049-3 978-3-540-71050-9. doi: 10.1007/978-3-540-71050-9.

Yifei Wang, Tavor Baharav, Yanjun Han, Jiantao Jiao, and David Tse. Beyond the Best: Distribution Functional Estimation in Infinite-Armed Bandits. *Advances in Neural Information Processing Systems*, 35:9262–9273, December 2022.

Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A): 2620–2648, November 2019. ISSN 1350-7265. doi: 10.3150/18-BEJ1065. Publisher: Bernoulli Society for Mathematical Statistics and Probability.

Yinchu Zhu and Ilya O. Ryzhov. Semidiscrete optimal transport with unknown costs, November 2023. arXiv:2310.00786 [econ].