Convergence of a variable metric scheme with application to optimal transport

Adrien Vacher

ORSAY – 20th November, 2025

Joint work with François-Xavier Vialard – accepted at ICML 2023

Outline

Introduction

The (quadratic) optimal transport problem

Given two probability measures μ, ν on $\mathbb{R}^{\mathrm{d}}$, the squared Wasserstein distance is defined as

$$\mathrm{W}_2^2(\mu, \nu) = \inf_{\pi \geq 0} \int \|\mathrm{x} - \mathrm{y}\|^2 \mathrm{d}\pi(\mathrm{x}, \mathrm{y}) + \iota(\pi_1 = \mu) + \iota(\pi_2 = \nu)\,,$$

which, under several mild conditions (Santambrogio, 2015), also admits the following dual formulation

$$\mathrm{W}_2^2(\mu, \nu) = \sup_{\varphi(\mathrm{x}) + \psi(\mathrm{y}) \leq \|\mathrm{x} - \mathrm{y}\|^2} \int \varphi(\mathrm{x}) \mathrm{d}\mu(\mathrm{x}) + \int \psi(\mathrm{y}) \mathrm{d}\nu(\mathrm{y})\,.$$

Motivations

Theorem (Brenier, 90s)

If μ has compact, convex support and has a C^1 density bounded from above and below, then φ is C^2 and verifies

$$\mathrm{T}_{\#}(\mu) = \nu\,.$$

where $\mathrm{T}(\mathrm{x}) = \mathrm{x} - \nabla\varphi(\mathrm{x})$ is the optimal transport map.

- For ML applications, one seeks to estimate T by $\hat{\mathrm{T}}$ from empirical measures $\hat{\mu} = \frac{1}{\mathrm{n}}\sum_{\mathrm{i}=1}^{\mathrm{n}} \delta_{\mathrm{x_i}}$, $\hat{\nu} = \frac{1}{\mathrm{n}}\sum_{\mathrm{i}=1}^{\mathrm{n}} \delta_{\mathrm{y_i}}$ where the $(\mathrm{x_i})$'s (resp. the $(\mathrm{y_i})$'s) are iid samples from μ (resp. ν).
- Ideally, $\hat{\mathrm{T}}$ should be "not too expensive" to compute and get closer to the original OT map T as n grows.

Approach via entropic formulation

A popular approach is to regularize the primal with the addition of an entropic term $+\varepsilon\mathrm{KL}(\pi|\mu\otimes\nu)$ (Léonard, 2014). One can then recover approximate potentials $(\hat{\varphi}_\varepsilon, \hat{\psi}_\varepsilon)$ by solving the following dual problem

$$\sup_{\varphi,\psi} \int \varphi(\mathrm{x})\mathrm{d}\hat{\mu}(\mathrm{x}) + \int \psi(\mathrm{y})\mathrm{d}\hat{\nu}(\mathrm{y}) - \varepsilon \int \mathrm{e}^{\frac{\varphi(\mathrm{x})+\psi(\mathrm{y})-\|\mathrm{x}-\mathrm{y}\|^2}{\varepsilon}} \mathrm{d}\hat{\mu}(\mathrm{x})\mathrm{d}\hat{\nu}(\mathrm{y})\,.$$

This problem admits a finite reparametrization and can be approximately solved by the well-known Sinkhorn algorithm (Cuturi, 2013).

Drawbacks of entropic approach

1. Slow of convergence of Sinkhorn as regularization ε decreases: after k iterations, the error scales as $\mathsf{min}((1 - e^{-1/\varepsilon})^{2k}, 1/(k\varepsilon))$ (Peyré and Cuturi, 2019; Léger, 2021).
2. Approximation of the original potential degrades exponentially with the dimension: $\|\nabla\varphi - \nabla\hat{\varphi}_\varepsilon\| \sim \varepsilon^{-d/2}/\sqrt{n}$ (Pooladian and Niles-Weed, 2021).

These two drawbacks find the same origin: as ε goes to 0 we get closer to the original linear dual formulation. Hence

1. The problem lacks strong convexity, leading to slow rates.
2. We end up by simply discretizing the cost constraint $\varphi(x) + \psi(y) \leq \|x - y\|^2$ on the sample points, which is a way too loose relaxation (Vacher et al., 2021).

Semi-dual approach

We pre-optimize the dual w.r.t the potential ψ so that the OT map is given by the gradient of:

$$\operatorname{arg\,min}_{\mathrm{f}} \mathrm{J}(\mathrm{f}) := \int \mathrm{f}(\mathrm{x}) \mathrm{d}\mu(\mathrm{x}) + \int \mathrm{f}^*(\mathrm{y}) \mathrm{d}\nu(\mathrm{y}),$$

where f^* given by: $\mathrm{f}^*(\mathrm{y}) = \sup_{\mathrm{x}} \mathrm{x}^\top \mathrm{y} - \mathrm{f}(\mathrm{x})$.

- The new objective gains convexity: if f has M-Lipschitz gradient $\mathrm{J}(\mathrm{f}) - \mathrm{J}(\mathrm{f}_0) \geq \frac{1}{2\mathrm{M}} \|\nabla \mathrm{f} - \mathrm{T}\|^2_{\mathrm{L}^2(\mu)}$ where $\mathrm{T} = \nabla \mathrm{f}_0$ (Hütter and Rigollet, 2021).
- The cost constraint is removed: more appealing statistical rates (Divol et al., 2025).

$\Longrightarrow$ we aim at developing an algorithm to solve the semi-dual formulation when $\mu, \nu = \hat{\mu}, \hat{\nu}$.

A variable metric scheme

Euclidean gradient descent

The simplest way to minimize a Euclidean function $\theta : \mathbb{R}^{\mathrm{d}} \to \mathbb{R}$ is to apply gradient descent

$$\mathrm{x}_{\mathrm{k}+1} = \mathrm{x}_{\mathrm{k}} - \tau \nabla \theta(\mathrm{x}_{\mathrm{k}}) \,.$$

However, the gradient is not an intrinsic quantity, it is implicitly dependent of the metric:

$$\mathrm{x}_{\mathrm{k}+1} = \operatorname*{arg\,min}_{\mathrm{x}} \mathrm{d}\theta_{\mathrm{x}_{\mathrm{k}}}(\mathrm{x}) + \frac{1}{2\tau} \|\mathrm{x} - \mathrm{x}_{\mathrm{k}}\|^2 \,,$$

with $\mathrm{d}\theta_{\mathrm{y}}$ the differential of θ at y. Can be go beyond a euclidean metric?

Mirror descent: relative smoothness/strong convexity (Lu et al., 2018)

Let h be a differentiable convex function. Instead of using the euclidean metric, one can use the Bregman divergence associated to h and solve

$$x_{k+1} = \operatorname*{arg\,min}_x d\theta_{x_k}(x) + \frac{1}{2\tau}\Delta_h(x, x_k)\,,$$

with $\Delta_h(x, x_k) = h(x_x) - h(x) - dh_{x_k}(x - x_k)$. If θ is convex with minimum θ^* and:

- if there exists $\beta \geq 0$ s.t. $\Delta_\theta(x, y) \leq \beta\Delta_h(x, y)$ and $1/\tau = \beta$, then $\theta(x_k) - \theta^* \leq O(\beta/k)$ (relative smoothness).
- if there further exists $\alpha > 0$ s.t. $\Delta_\theta(x, y) \geq \alpha\Delta_h(x, y)$, then $\theta(x_k) - \theta^* \leq O(\frac{\alpha}{(1+\alpha/(\beta-\alpha))^k - 1})$ (relative strong convexity).

Taking $h(x) = \|x\|^2/2$ recovers standard euclidean g.d. results.

Variable relative smoothness and strong convexity of the semi-dual

We wish to apply a mirror descent type algorithm for the semi-dual: which divergence should we choose?

Proposition (Stability of the semi-dual)

If (f, g) are γ-strongly convex potentials then

$$\Delta_{\mathrm{J}}(\mathrm{f}, \mathrm{g}) \leq \frac{1}{2\gamma} \|\nabla \mathrm{f} - \nabla \mathrm{g}\|^2_{\mathrm{L}^2((\nabla \mathrm{g}^*)_\#(\nu))} .$$

If (f, g) have M-Lispchitz gradients and $\nabla \mathrm{g}^*$ is well-defined over the support of ν then

$$\Delta_{\mathrm{J}}(\mathrm{f}, \mathrm{g}) \geq \frac{1}{2\mathrm{M}} \|\nabla \mathrm{f} - \nabla \mathrm{g}\|^2_{\mathrm{L}^2((\nabla \mathrm{g}^*)_\#(\nu))} .$$

Remarks

- These results generalize previous stability results recovered for $g = f_0$ the optimal transport potential.
- Unfortunately, we cannot directly apply the previous results on mirror descent since $\|\nabla f - \nabla g\|^2_{(\nabla g^*)_\#(\nu)}$ cannot be taken as the Bregman divergence of a fixed function h: such a function h would solve for all (f, g)

$$\lim_{\lambda \to \infty} \frac{h(\lambda f)}{\lambda^2} = \|\nabla f\|^2_{L^2((\nabla g^*)_\#(\nu))} .$$

The left hand side does not depend on g while the right hand side does.

A variable metric scheme

- We thus study the convergence of the following type of schemes

$$\mathrm{x}_{\mathrm{k}+1} = \operatorname*{arg\,min}_{\mathrm{x}\in\mathrm{C}} \mathrm{d}\theta_{\mathrm{x}_\mathrm{k}}(\mathrm{x}) + \frac{1}{2\tau}\|\mathrm{x}-\mathrm{x}_\mathrm{k}\|^2_{\mathrm{x}_\mathrm{k}},$$

  with C a convex set and $\|\cdot\|_\mathrm{x}$ a pseudo-metric depending of the point x. We need the constraint C since we need the potentials over which we optimize to remain strongly convex/and or smooth in order to get stability.
- This variable metric scheme was previously studied but it either assumes that the metric gets finer across each iterations $\|\cdot\|_{\mathrm{x}_\mathrm{k}} \leq \|\cdot\|_{\mathrm{x}_{\mathrm{k}+1}}$ either assume that the metrics are all equivalent (Combettes and Vũ, 2014) $\implies$ this is typically not the case for the semi-dual.

Convergence result: convexity and relative smoothness

Theorem

Let E be a Banach space, let F be a real-valued convex function with Gateaux derivative dF satisfying for all $(\mathrm{x},\mathrm{y}) \in \mathrm{E}$, $\Delta_{\mathrm{F}}(\mathrm{x},\mathrm{y}) \leq \frac{\beta}{2}\mathrm{A}^{\mathrm{y}}(\mathrm{x}-\mathrm{y})$ where for all $\mathrm{y} \in \mathrm{E}$, $\mathrm{A}^{\mathrm{y}}(\cdot)$ is a 2-homogeneous form over E depending on y and where β is a strictly positive constant and let $\mathrm{C} \subset \mathrm{E}$ be a closed convex subset of E. Assuming that $\sup_{(\mathrm{x},\mathrm{y})\in \mathrm{C}^2} \mathrm{A}^{\mathrm{y}}(\mathrm{x}-\mathrm{y}) \leq \mathrm{K}$, that a minimizer $\bar{\mathrm{x}} \in \mathrm{C}$ exists and that the iterates $\mathrm{x}_0 \in \mathrm{C}$, $(\mathrm{x}_{\mathrm{k}})$ generated as

$$\mathrm{x}_{\mathrm{k}+1} \in \arg\min_{\mathrm{x}\in \mathrm{C}} \mathrm{dF}(\mathrm{x}_{\mathrm{k}})(\mathrm{x}-\mathrm{x}_{\mathrm{k}}) + \frac{\beta}{2}\mathrm{A}^{\mathrm{x}_{\mathrm{k}}}(\mathrm{x}-\mathrm{x}_{\mathrm{k}}), \tag{1}$$

exist, we have $\mathrm{F}(\mathrm{x}_{\mathrm{k}}) - \mathrm{F}(\bar{\mathrm{x}}) \leq \frac{2\beta \mathrm{K}}{\mathrm{k}+1}$.

Convergence result: relative convexity and relative smoothness

Theorem

Let E be a Banach space, let F be a real-valued convex function with Gateaux derivative dF, and let $\mathrm{C} \subset \mathrm{E}$ be a closed convex subset of E. If there exists $\alpha, \beta > 0$ and $\mathrm{A}^{\mathrm{y}}(\cdot)$ a 2-homogeneous form such that for all $(\mathrm{x}, \mathrm{y}) \in \mathrm{E}$, $\frac{\alpha}{2}\mathrm{A}^{\mathrm{y}}(\mathrm{x}-\mathrm{y}) \leq \Delta_{\mathrm{F}}(\mathrm{x}, \mathrm{y}) \leq \frac{\beta}{2}\mathrm{A}^{\mathrm{y}}(\mathrm{x}-\mathrm{y})$. If a minimizer $\bar{\mathrm{x}} \in \mathrm{C}$ and the iterates $\mathrm{x}_0 \in \mathrm{C}$, (x_k) generated by (1) exist, we have

$$\mathrm{F}(\mathrm{x}_\mathrm{k}) - \mathrm{F}(\bar{\mathrm{x}}) \leq \left(1 - \frac{\alpha}{\beta}\right)^{\mathrm{k}} [\mathrm{F}(\mathrm{x}_0) - \mathrm{F}(\bar{\mathrm{x}})].$$

Application to the minimization of the semi-dual

An exponentially convergent scheme

Theorem (Balanced case, smooth and strongly convex)

Let C be a convex set of γ-strongly convex, M-smooth functions. The minimum $\min_{f \in C} J(f)$ is attained at $\bar{f} \in C$. Furthermore, the iterates

$$f_{k+1} = \arg\min_{f \in C} dJ(f_k)(f - f_k) + \frac{1}{2\gamma} \|\nabla f - \nabla f_k\|^2_{L^2((\nabla f_k^*)_\#(\nu))},$$

are well defined and verify

$$J(f_{k+1}) - J(\bar{f}) \leq (1 - \frac{\gamma}{M})^k (J(f_1) - J(\bar{f})).$$

Remarks

- By standard regularity of quadratic OT results, the ground truth optimal transport potential verifies the smoothness and strong convexity constraints when μ, ν have compact, convex support (Caffarelli, 2000). Hence in some cases, these extra constraints are benign.
- These results generalize to unbalanced quadratic OT.
- In Jacobs and Léger (2020), a similar type of scheme is considered yet instead the authors assume that μ, ν are supported on a fixed grid Ω and thus consider the fixed metric $\|\nabla \mathrm{f} - \nabla \mathrm{g}\|^2_{\mathrm{L}^2(\Omega)}$. Yet the potentials are not constrained.

Practical implementation?

Recalling that $dJ(f)(g) = \int g(x)(d\mu(x) - d(\nabla f^*)_\#(\nu)(x))$, the iterations we need to solve for empirical measures $\hat{\mu}, \hat{\nu}$ are given by

$$\begin{aligned} f_{k+1} = \operatorname*{arg\,min}_{f \in C} \frac{1}{n} \sum_{i=1}^{n} f(x_i) - f((\nabla f_k^*)(y_i)) \\ + \frac{1}{2n\gamma} \sum_{i=1}^{n} \|\nabla f((\nabla f_k^*)(y_i)) - y_i\|^2 \,. \end{aligned}$$

Even though we observe that f and ∇f are evaluated on a finite set of points, the set C remains infinite dimensional and the problem above may not admit a finite reparametrization.

Finite reparametrization: the case of smooth and strongly convex functions.

Proposition (Taylor et al. (2017))

For $C := C_{\gamma,M}$, the set of M-smooth, γ-strongly convex functions, the previous problem can be reformulated as

$$\inf_{\substack{u\in\mathbb{R}^{2n}\\ g\in\mathbb{R}^{(2n)\times d}}} \quad \sum_{i=1}^{n} u_i - \sum_{i=n+1}^{2n} u_i + \sum_{i=n}^{n+1} \|g_i - z_i\|^2/(2\gamma)\,,$$

$$u_i \geq u_j + g_j^\top(z_i - z_j) + \frac{\|(g_i - g_j)/\sqrt{M} - \sqrt{\lambda}(z_i - z_j)\|^2}{2(1-\lambda/M)}$$

where $z_i = x_i$ if $\in \{1, \cdots, n\}$ and $z_i = y_i$ if $i \in \{n+1, \cdots, 2n\}$. Furthermore, the cost to solve this problem with an Interior Point Method (Nesterov and Nemirovskii, 1994) with precision τ requires $O(n^3 d^3 \log(\log(\tau)))$ operations.

Recovering iterates

Proposition

Denoting (u, g) a solution of the previous problem, then point-wise, the corresponding potential $f(x)$ as well as its gradient $\nabla f(x)$ can be recovered as the solution (v, p) of the following program

$$\min_{v,g} v$$

$$v \geq u_i + g_i(x - z_i) + \frac{\|(g_i - p)/\sqrt{M} - \sqrt{\lambda}(z_i - x)\|^2}{2(1 - \lambda/M)}.$$

In particular, one can compute $\nabla f^*(y)$ point-wise via standard gradient descent. Note that since explicit regularity bounds are known for the recovered potential f, we can achieve optimal rates.

A word about the statistical/computational trade-off

- Even though the number of iterations is $\mathrm{O}(1)$, the overall complexity is dominated by the cost per iteration and scales as $\tilde{\mathrm{O}}(\mathrm{n}^3)$ which may seem disappointing. However, it yields an approximation guarantee of the original OT map of order $\mathrm{n}^{-2/\mathrm{d}}$ hence the overall approximation cost of the original OT map scales as $\tau^{-\min(7,3\mathrm{d}/2+1)}$
- For Sinkhorn, once the regularization parameter is optimally chosen, we obtain $\tau^{-2(\mathrm{d}+1)-7}$ (Vacher and Vialard, 2023).

For $M = +\infty$, the problem we seek to solve is now

$$\inf_{f \in C_{\lambda,\infty}} J(f)\,.$$

This problem was recently studied in Gallouët et al. (2024).
Strikingly, it admits a weak optimal transport reformulation.
Especially we may obtain a better complexity than $\tilde{O}(n^3)$.

Numerical application

We apply this method the compute the OT map between μ, a discrete uniform probability measure across 50-regularly spaced points on $[-1, 1]$ and $\nu = (\nabla \mathrm{f}_0)_\#(\mu)$ with f_0 convex so that an OT map T is given by $\mathrm{T} = \nabla \mathrm{f}_0$. We choose for this experiment, $\mathrm{f}_0 = |\mathrm{x}| + 0.25\mathrm{x}^2/2$.

Frank-Wolfe as a baseline

In the introduction, we recalled that gradient descent was implicitly dependent of a metric. One way to avoid choosing such a metric is simply to compute the iterates as

$$\tilde{\mathrm{f}} = \arg\min_{\mathrm{f} \in \mathrm{C}_{\lambda,\mathrm{M}}} \mathrm{dF}(\mathrm{f}_\mathrm{k})(\mathrm{f}) \,,$$

and take $\mathrm{f}_{\mathrm{k}+1} = \mathrm{f}_\mathrm{k} + \frac{2}{\mathrm{k}+1}(\tilde{\mathrm{f}} - \mathrm{f}_\mathrm{k})$ which is exactly the Frank-Wolfe algorithm for the semi-dual.

Visual comparison

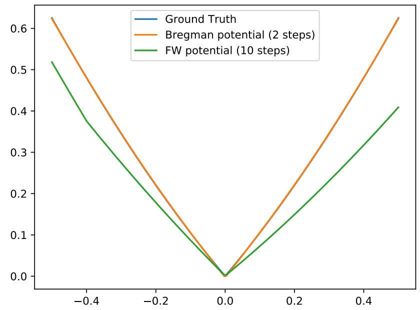

Figure 1: Potential generated by Frank-Wolfe after 10 iterations (in green) vs generated by our algorithm after 2 iterations (in orange) vs. ground truth (in blue).

Towards O(n) optimal transport?

The choice $C_{\lambda,M}$ is way too expressive $\implies$ not so good statistical rates + computationally quite expensive. How about a parametric model? For instance take

$$C_{z,\varepsilon} := \{f(x) = \sum_{i=1}^{p} w_i \sqrt{\|x - z_i\|^2 + \varepsilon} | w_i \geq 0\},$$

where the centroids z_i and the regularization ε are fixed. In this case the iterates cost $O(p^2dn)$ and if the ground truth OT potential belongs to $C_{z,\varepsilon}$, we can recover an approximation rate in $O(1/n)$.

Conclusion

The choice of C?

How to chose the set C such that:

1. The iterations admit a finite reparametrization and are cheap to compute (typically O(n)).
2. C is not "too complex" in order to recover appealing statistical rates.
3. C can approximate reasonably well a large class of OT potentials $\implies$ need additional structure on OT, probably very hard problem.

A geometric question

Mirror descent is roughly Riemannian gradient descent with hessian metric when the step size goes to zero. What is the space $(C_{\lambda,M}, g_f(p,p) = \|\nabla p\|_{L^2((\nabla f^*)_\#(\nu))})$?

References

Caffarelli, L. A. (2000). Monotonicity properties of optimal transportation and the fkg and related inequalities. Communications in Mathematical Physics.

Combettes, P. L. and Vũ, B. C. (2014). Variable metric forward–backward splitting with applications to monotone inclusions in duality. Optimization.

Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In NeurIPS.

Divol, V., Niles-Weed, J., and Pooladian, A.-A. (2025). Optimal transport map estimation in general function spaces. The Annals of Statistics.

Gallouët, T., Natale, A., and Todeschi, G. (2024). From geodesic extrapolation to a variational bdf2 scheme for wasserstein gradient flows. Mathematics of Computation.

Hütter, J.-C. and Rigollet, P. (2021). Minimax estimation of smooth optimal transport maps. The Annals of Statistics.

Jacobs, M. and Léger, F. (2020). A fast approach to optimal transport: the back-and-forth method. Numerische Mathematik.

Léger, F. (2021). A gradient descent perspective on sinkhorn. Applied Mathematics and Optimization.

Lu, H., Freund, R. M., and Nesterov, Y. (2018). Relatively smooth convex optimization by first-order methods, and applications. SIAM Journal on Optimization.

Léonard, C. (2014). A survey of the schrodinger problem and some of its connections with optimal transport. Discrete and Continuous Dynamical Systems.

Nesterov, Y. E. and Nemirovskii, A. S. (1994). Interior point methods in convex programming: theory and applications. Society for Industrial and Applied Mathematics, Philadelphia.

Peyré, G. and Cuturi, M. (2019). Computational optimal transport: With applications to data science. Foundations and Trends in Machine Learning.

Pooladian, A.-A. and Niles-Weed, J. (2021). Entropic estimation of optimal transport maps. In NeurIPS, OTML workshop.

Santambrogio, F. (2015). Optimal transport for applied mathematicians. Birkäuser, NY.

Taylor, A. B., Hendrickx, J. M., and Glineur, F. (2017). Exact worst-case performance of first-order methods for composite convex optimization. SIAM Journal on Optimization.

Vacher, A., Muzellec, B., Rudi, A., Bach, F., and Vialard, F.-X. (2021). A dimension-free computational upper-bound for smooth optimal transport estimation. In COLT.

Vacher, A. and Vialard, F.-X. (2023). Semi-dual unbalanced quadratic optimal transport: fast statistical rates and convergent algorithm. In ICML.