



HAL
open science

Accelerated Bregman primal-dual methods applied to optimal transport and Wasserstein Barycenter problems

Antonin Chambolle, Juan Pablo Contreras

► **To cite this version:**

Antonin Chambolle, Juan Pablo Contreras. Accelerated Bregman primal-dual methods applied to optimal transport and Wasserstein Barycenter problems. *SIAM Journal on Mathematics of Data Science*, 2022, 4 (4), pp.1369-1395. 10.1137/22M1481865 . hal-03806188

HAL Id: hal-03806188

<https://hal.science/hal-03806188v1>

Submitted on 7 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ACCELERATED BREGMAN PRIMAL-DUAL METHODS APPLIED TO OPTIMAL TRANSPORT AND WASSERSTEIN BARYCENTER PROBLEMS

ANTONIN CHAMBOLLE AND JUAN PABLO CONTRERAS

ABSTRACT. This paper discusses the efficiency of Hybrid Primal-Dual (HPD) type algorithms to approximate solve discrete Optimal Transport (OT) and Wasserstein Barycenter (WB) problems, with and without entropic regularization. Our first contribution is an analysis showing that these methods yield state-of-the-art convergence rates, both theoretically and practically. Next, we extend the HPD algorithm with linesearch proposed by Malitsky and Pock in 2018 to the setting where the dual space has a Bregman divergence, and the dual function is relatively strongly convex to the Bregman's kernel. This extension yields a new method for OT and WB problems based on smoothing of the objective that also achieves state-of-the-art convergence rates. Finally, we introduce a new Bregman divergence based on a scaled entropy function that makes the algorithm numerically stable and reduces the smoothing, leading to sparse solutions of OT and WB problems. We complement our findings with numerical experiments and comparisons.

Keywords: optimal transport, Wasserstein barycenter, saddle-point, primal-dual method.

AMS subject classifications. 49Q22, 65Y20, 90C05, 90C06, 90C08, 90C47

1. INTRODUCTION

We address the problem of computing the Optimal Transport (OT) between two discrete probability distributions and Wasserstein Barycenter (WB) between several (two or more) discrete probability distributions. Optimal Transport is central in machine learning applications such as classification [27] and unsupervised learning [3, 7]. On the other hand, the Wasserstein Barycenter problem appears naturally in clustering [23] and other problems of imaging [10].

Discrete Optimal Transport and Wasserstein Barycenter are Linear Programming (LP) problems that can be tackled through interior point methods [31], or Network Simplex [11, 20], which are very efficient (*cf* our experiments in Section 5). Nevertheless, the increasing size of the data leads to prohibitive large instances for exact solvers. In addition, some common nonlinear extensions might not be tractable by linear optimization techniques. Thus, finding efficient methods to solve these problems by nonlinear optimization is an exciting task in either optimization or computing theory.

Beyond LP, a classical approach to tackle OT is the well-known Sinkhorn algorithm. This method uses an alternating minimization procedure over the dual variables of an entropic regularized version of the OT problem [16, 40]. Recent papers have analyzed the complexity of Sinkhorn and its variants in terms of the number of arithmetic operations necessary to reach a ε -approximated solution of OT. Most of these rates depend on the error ε , the size of the discrete measures n , and the largest element $\|C\| := \max_{i,j} C_{i,j}$ of the matrix of transportation costs $C \in \mathbb{R}_+^{n \times n}$ (assuming for simplicity, and without loss of generality, $C_{i,j} \geq 0$). In [2] authors proved the rate $\tilde{O}(\frac{n^2\|C\|^3}{\varepsilon^3})$ for Sinkhorn and its greedy variant named Greenkhorn. Then, in [18] the authors improved the analysis of Sinkhorn and proved the rate $\tilde{O}(\frac{n^2\|C\|^2}{\varepsilon^2})$. The same rate was obtained for Greenkhorn in [33]. Using randomized coordinate descent [39], the authors in [32] proposed the Randkhorn algorithm, a randomized version of Sinkhorn, and prove the rate $\tilde{O}(\frac{n^{7/3}\|C\|^{4/3}}{\varepsilon})$.

Recently, there has been a growing interest in adapting accelerated convex optimization methods for the entropic regularized OT. Interestingly, these methods achieve better theoretical convergence rates than Sinkhorn-like methods. To the best of our knowledge, the first of these methods was proposed in [18]. The authors designed and analyzed the APDAGD method, a linesearch version of the Accelerated Gradient Descent (AGD) scheme [37, 46]. Shortly after, Lin et. al. [33] refined the analysis of APDAGD and prove the rate $\tilde{O}(\frac{n^{5/2}\|C\|}{\varepsilon})$. In the same paper, authors proposed the APDAMD method that slightly generalizes the APDAGD and prove the rate $\tilde{O}(\frac{n^2\sqrt{\kappa}\|C\|}{\varepsilon})$, where $\kappa = O(n)$ is a factor associated to the strong convexity of the Bregman prox-function with respect to the uniform norm. The same rate was obtained in [21] using an algorithm based on Accelerated Alternating Minimization (AAM), and also in [22] using Accelerated Primal-Dual Coordinate Descent.

A different approach was introduced by Jambulapati *et. al.* in [24]. Here, the authors proposed a new method based on dual-extrapolation [38] and area convexity [44], and they prove the rate $\tilde{O}(\frac{n^2\|C\|}{\varepsilon})$ (see also [1], [15] and [8] for similar rates). Despite the theoretical optimality of the rate $O(\frac{n^2\|C\|}{\varepsilon})$ (see [8]), methods as AGD or AMM outperform in practice the area-convexity approach as is shown by the numerical experiments in [21] and [17].

The study of the Wasserstein Barycenter problems follows a similar path to OT. For a problem with m marginals, the Iterative Bregman Projection (IBP) method proposed in [6] was proven to achieve the rate $\tilde{O}(\frac{mn^2\|C\|^2}{\varepsilon^2})$ in [26]. Applications of AGD and AAM to WB were also studied in [26] and [21], respectively. Here, authors proved the rate $\tilde{O}(\frac{mn^{5/2}\|C\|}{\varepsilon})$, analogous to the one obtained for OT with the same algorithms. FastIBP algorithm was proposed in [32] and authors proved the rate $\tilde{O}(\frac{mn^{7/3}\|C\|^{4/3}}{\varepsilon^{4/3}})$. Recently, in [17] the idea of area-convexity was extended to WB problems with the near-optimal rate $\tilde{O}(\frac{mn^2\|C\|}{\varepsilon})$. In the same work, authors presented an application of Mirror-Prox [36]) and proved the rate $\tilde{O}(\frac{mn^{5/2}\|C\|}{\varepsilon})$ for the worst-case duality-gap of a saddle-point formulation

of WB. Moreover, they provided numerical evidence where Mirror-Prox outperforms the area-convexity method.

Close to Mirror-Prox is the Hybrid Primal-Dual (HPD). This iterative primal-dual method aims to solve a very general class of saddle-points problems and requires half computations per iteration compared to Mirror-Prox. The method was first analyzed in [12]. Under mild convexity assumptions, the authors proved the rate $O(1/N)$ for the duality-gap after N iterations, and the accelerated rate $O(1/N^2)$ in the presence of strong convexity. Later, it was revisited in [13] with an improvement of the analysis and an extension to general Bregman divergences.

Malitsky and Pock [35] incorporated a linesearch procedure to the HPD in order to locally estimate the parameter of smoothness at each iteration and allow larger step-sizes. The method with linesearch enjoys the same rates $O(1/N)$ and $O(1/N^2)$ as classic HPD, although the original version is restricted to the Euclidean setting. Recently, in [25] it was observed that the analysis of the rate $O(1/N)$ extends to the case in which the dual space is endowed with a Bregman distance while the primal space remains Euclidean. This paper takes one step forward by studying the extension of the rate $O(1/N^2)$. Table 1 brief the results on HPD listed so far.

Paper	primal & dual settings	primal & dual functions	linesearch	rate
[12]	both Euclidean	both convex	✗	$O(1/N)$
[13]	both Bregman	both convex	✗	$O(1/N)$
[13]	both Bregman	convex & strong convex	✗	$O(1/N^2)$
[35]	both Euclidean	both convex	✓	$O(1/N)$
[35]	both Euclidean	convex & strong convex	✓	$O(1/N^2)$
[25]	Euclidean & Bregman	both convex	✓	$O(1/N)$
This paper	Euclidean & Bregman	convex & strong convex	✓	$O(1/N^2)$

TABLE 1. Convergence rates for the duality-gap with HPD

1.1. Contributions. The contribution of this paper is threefold: First, we extend the $O(1/N^2)$ result of [35] to the setting where the dual space has a Bregman distance, and the dual function is strongly convex relative to the kernel of the Bregman distance. To do this, we follow closely the proof of Malitsky and Pock by replacing some of the Hilbertian identities with a well-known 3-points inequality that holds in Bregman settings proposed in [14]. At the end of the proof, we propose a new initialization for the algorithm’s parameters that simplifies the analysis and corrects an issue overlooked in the original proof.

Second, we show how to specialize the HPD into the setting of discrete Optimal Transport and Wasserstein Barycenter problems. To do this, we write these problems in a saddle-point formulation and find tight bounds for the norm of the dual variables. Interestingly,

we find the rates $\tilde{O}(\frac{n^{5/2}\|C\|}{\varepsilon})$ for OT, and $\tilde{O}(\frac{mn^{5/2}\|C\|}{\varepsilon})$ for WB, matching the rates of current state-of-the-art methods in the literature.

Finally, in Section 4.1 we propose a new scaled entropy kernel that can be accomplished with HPD and other methods in the literature, such as AGD. This new entropy function makes the computation numerically stable, as it avoids computing the logarithm of very small numbers. Moreover, the solutions obtained using the new entropy have smaller support than traditional entropic regularized solutions, which may be of interest for specific applications. We end our paper with a remark (Sec. 4.2) on *easy* nonlinear extensions of our algorithm and an experimental section comparing our methods with various other techniques from the literature (Sec. 5).

2. THE HYBRID PRIMAL DUAL METHOD

We consider the saddle-point problem

$$\min_{x \in X} \max_{y \in Y} \mathcal{L}(x, y) := g(x) + \langle Kx, y \rangle - h^*(y), \quad (1)$$

where X, Y are convex sets of general normed spaces $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ (the primal) and $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ (the dual), respectively. Let $K : \mathcal{X} \mapsto \mathcal{Y}^*$ a linear map such that

$$L := \sup_{\|x\|_{\mathcal{X}} \leq 1, \|y\|_{\mathcal{Y}} \leq 1} \langle Kx, y \rangle < \infty,$$

and g, h^* proper, lower semicontinuous, convex functions over X and Y , respectively. Each space is endowed with a Bregman divergence $D(z, \bar{z}) = \xi(z) - \xi(\bar{z}) - \langle \nabla \xi(\bar{z}), z - \bar{z} \rangle$, generated by a kernel or prox-function ξ that is 1-strongly convex with respect to the norm of the space, differentiable in $\text{int}(\text{dom } \xi)$, and $\|\nabla \xi(x)\| \rightarrow \infty$ when $x \rightarrow \partial(\text{dom } \xi)$ (that is, a strongly convex *Legendre function*, as introduced in [42] (Sec. 26), [14], [5]). We denote by $\xi_{\mathcal{X}}$ and $\xi_{\mathcal{Y}}$ the kernel functions on the primal and dual, respectively. Similarly, we denote by $D_{\mathcal{X}}$ and $D_{\mathcal{Y}}$ the Bregman divergence, and we assume that $X \subset \overline{\text{dom } \xi_{\mathcal{X}}}$ and $Y \subset \overline{\text{dom } \xi_{\mathcal{Y}}}$. In addition, we replace without loss of generality g with $g + \delta_X$ (where δ_X is the characteristic function of X , with value 0 in X and $+\infty$ outside) and h^* with $h^* + \delta_Y$, and forget in the remaining of the paper about the convex constraint sets X, Y .

Finally, we define the duality-gap at (x, y) measured at (\tilde{x}, \tilde{y}) as $\mathcal{G}_{\tilde{x}, \tilde{y}}(x, y) := \mathcal{L}(x, \tilde{y}) - \mathcal{L}(\tilde{x}, y)$.

For initial points $x^0, \bar{x}^0 \in \text{dom } \xi_{\mathcal{X}}$, $y_0 \in \text{dom } \xi_{\mathcal{Y}}$, and given non-negative sequences $\{\tau_k\}_k$, $\{\sigma_k\}_k$, $\{\theta_k\}_k$, the main iteration of the HPD method (without linesearch) is given by

$$\begin{aligned} y^{k+1} &= \arg \min_{y \in \mathcal{Y}} h^*(y) - \langle K\bar{x}^k, y \rangle + \frac{1}{\sigma_k} D_{\mathcal{Y}}(y, y^k), \\ x^{k+1} &= \arg \min_{x \in \mathcal{X}} g(x) + \langle Kx, y^k \rangle + \frac{1}{\tau_k} D_{\mathcal{X}}(x, x^k), \\ \bar{x}^{k+1} &= x^{k+1} + \theta_k(x^{k+1} - x^k). \end{aligned}$$

In the simpler case where $\tau_k \equiv \tau_0$ and $\sigma_k \equiv \sigma_0$ are constant, Theorem 1 in [13] establishes that taking $\theta_k \equiv 1$ and $\tau_0\sigma_0L^2 \leq 1$ we have for all $(x, y) \in X \times Y$

$$\mathcal{G}_{x,y}(\hat{x}^N, \hat{y}^N) \leq \frac{1}{T_N} \left(\frac{1}{\tau_0} D_{\mathcal{X}}(x, x^0) + \frac{1}{\sigma_0} D_{\mathcal{Y}}(y, y^0) \right), \quad (2)$$

where $T_N = N/2$, $\hat{x}^N = \frac{1}{N} \sum_{k=1}^N x^k$ and $\hat{y}^N = \frac{1}{N} \sum_{k=1}^N y^k$.

Furthermore, if g is strongly convex relative to the primal prox-function $\xi_{\mathcal{X}}$ (see Definition 1.2 in [34]), then Theorem 4 (together with footnote 1, p. 268) in [13] establishes that taking $\theta_{k+1} = \frac{1}{\sqrt{1+\gamma\tau_k}}$, $\tau_{k+1} = \theta_{k+1}\tau_k$, and $\sigma_{k+1} = \sigma_k/\theta_{k+1}$, equation (2) holds with $T_N \gtrsim \gamma N^2/(L^2\sigma_0)$, where the notation \gtrsim means that the inequality holds up to constant factors and for a sufficiently large N . We remark that a symmetric algorithm and result also holds in the case that h^* is strongly convex relative to $\xi_{\mathcal{Y}}$, which is the case of interest for the next section.

Remark 2.1 (Erratum to [13]). In [13] it is claimed the convergence of the iterates holds in general, yet the proof works only whenever the prox-function is differentiable on the whole domain (as is the case for instance for the scaled entropy of Sec. 4.1) or when the limit points lie in the interior (as in the entropy-regularized problem of Sec. 3.2). This is observed for instance in Theorem 4 in [30], or [29]. The algorithm also extends with easy adaption to strongly convex prox-functions, possibly non-smooth, but with locally bounded subgradients, as pointed out in [28] where a variant of the above Bregman primal-dual descent is introduced. This allows to consider prox-functions with finite gradients on the boundary of their domain. In these cases the convergence proof of [13], which uses the fact that limits of iterates are also fixed points of the algorithm (which is not correct if the prox-function has infinite gradients at such points as the algorithm is then undefined), may be adapted. The authors of [13] apologize for letting this pass through the publication process and warmly thank Chung-Wei Lee from U. South Carolina for pointing out this issue.

2.1. Linesearch. Algorithm 1 presents a primal-dual method with linesearch slightly different from the original proposed by [35]. The results in [35] guarantee the rates $O(1/N)$ and $O(1/N^2)$ for the particular case $D_{\mathcal{X}}(x, \bar{x}) = \frac{1}{2}\|x - \bar{x}\|_2^2$ and $D_{\mathcal{Y}}(y, \bar{y}) = \frac{1}{2}\|y - \bar{y}\|_2^2$,

which is referred to as the Euclidean setting. Recently, [25] proved that the $O(1/N)$ algorithm can be extended to the setting where the Bregman distance $D_{\mathcal{Y}}$ has a non-Euclidean kernel.

Algorithm 1 HPD method with linesearch

Require: $x^1 = x^0 \in \text{dom } \xi_{\mathcal{X}}$, $y^1 \in \text{dom } \xi_{\mathcal{Y}}$, $\beta_0 > 0$, $\tau_0 = \frac{1}{\sqrt{\beta_0 L}}$, $\gamma \geq 0$, $\theta_0 > 0$, $\rho \in (0, 1)$, set $k = 1$.

- 1: **repeat** {mean-loop}
- 2: $\tau_k \leftarrow \tau_{k-1} \sqrt{1 + \theta_{k-1}/\rho}$, $\beta_k \leftarrow \frac{\beta_{k-1}}{1 + \gamma \beta_{k-1} \tau_{k-1}}$
- 3: **repeat** {inner-loop: linesearch}
- 4: $\tau_k \leftarrow \rho \tau_k$, $\theta_k \leftarrow \tau_k / \tau_{k-1}$, $\sigma_k \leftarrow \beta_k \tau_k$
- 5: $\bar{x}^k \leftarrow x^k + \theta_k (x^k - x^{k-1})$
- 6: $y^{k+1} \leftarrow \arg \min_{y \in \mathcal{Y}} h^*(y) - \langle K \bar{x}^k, y \rangle + \frac{D_{\mathcal{Y}}(y, y^k)}{\sigma_k}$
- 7: $x^{k+1} \leftarrow \arg \min_{x \in \mathcal{X}} g(x) + \langle Kx, y^{k+1} \rangle + \frac{\|x - x^k\|_2^2}{2\tau_k}$
- 8: **until** $\frac{1}{2} \|x^{k+1} - \bar{x}^k\|_2^2 + \frac{1}{\beta_k} D_{\mathcal{Y}}(y^{k+1}, y^k) + \tau_k \langle K(x^{k+1} - \bar{x}^k), y^{k+1} - y^k \rangle \geq 0$
- 9: $k \leftarrow k + 1$
- 10: **until** stopping criteria

The following Theorem extends the results of [35] and [25] by observing that if h^* is strongly convex relative to $\xi_{\mathcal{Y}}$, then the accelerated rate $O(1/N^2)$ can be achieved. Its proof is given in Appendix A.1.

Theorem 2.2. *Let $(x^k, \bar{x}^k, y^k, \tau_k, \sigma_k, \theta_k)_k$ the sequences generated by Algorithm 1, and $T_N = \sum_{k=1}^N \tau_k$. For any $(x, y) \in \mathcal{X} \times \text{dom } \xi_{\mathcal{Y}}$ it holds*

$$\mathcal{G}_{x,y}(\hat{x}^N, \hat{y}^N) \leq \frac{1}{T_N} \left(\frac{1}{2} \|x - x^1\|_2^2 + \frac{1}{\beta_1} D_{\mathcal{Y}}(y, y^1) - \tau_1 \theta_1 (\mathcal{L}(\hat{x}^N, y) - \mathcal{L}(x^0, y)) \right), \quad (3)$$

where the ergodic outputs are now $\hat{x}^N = \frac{\tau_1 \theta_1 x^0 + \sum_{k=1}^N \tau_k \bar{x}^k}{\tau_1 \theta_1 + T_N}$ and $\hat{y}^N = \sum_{k=1}^N \frac{\tau_k}{T_N} y^{k+1}$.

The sequence $\beta_k \equiv \beta$ with $\tau_0 \sqrt{\beta} L = 1$ yields $T_N \gtrsim N$. If h^* is γ -strongly convex relative to $\xi_{\mathcal{Y}}$ and $0 < \gamma \leq L/\rho$, then the choice $\beta_0 > 0$, $\tau_0 = 1/(\sqrt{\beta_0} L)$ and $\theta_0 = \gamma \sqrt{\beta_0}/L$ yields $T_N \gtrsim \gamma N^2/L^2$.

If in addition the Lagrangian $\mathcal{L}(x, y)$ is linear with respect to feasible points x , then one can define $\hat{x}^N = \sum_{k=1}^N \frac{\tau_k}{T_N} \bar{x}^k$ and (3) simplifies to

$$\mathcal{G}_{x,y}(\hat{x}^N, \hat{y}^N) \leq \frac{1}{T_N} \left(\frac{1}{2} \|x - x^1\|_2^2 + \frac{1}{\beta_1} D_{\mathcal{Y}}(y, y^1) \right). \quad (4)$$

Remark 2.3 (Difference with respect to [35]). The main differences with respect to the original algorithm are the setup for the strongly convex setting and the linesearch procedure (steps 4 to 8 in Algorithm 1). We choose $\theta_0 = \gamma \sqrt{\beta_0}/L$ instead of $\theta_0 = 1$, and more importantly, we fix the initial value of τ_k before the linesearch as $\tau_k = \tau_{k-1} \sqrt{1 + \theta_{k-1}}$,

while the original algorithm proposes to choose $\tau_k \leq \tau_{k-1} \sqrt{1 + \theta_{k-1}}$. These modifications seem necessary in our analysis (see Appendix A.1) to prove the relation $\tau_k \geq \frac{1}{L\sqrt{\beta_k}}$ for all $k \geq 1$. This property was used without a proof in [35], yet the authors seem to have overlooked that it is unclear in the strongly convex case.

In the linesearch, the stopping criteria for the inner-loop (step 8 in Algorithm 1) was first proposed in [25]. This new condition yields larger steps than the original one, yet it requires to compute x^{k+1} at each iteration of the inner-loop. Hence, it is applicable only when the computation of x^{k+1} is relatively cheap, which will be the case in our implementations. In other cases one can revert to the termination criterion of [35]. Finally, we prove in the Appendix A.2 that the linesearch always terminates, and the overall number of inner-loop iterations of after N main-loop iterations is $O(N/|\ln \rho|)$, therefore does not affect the rate of convergence.

3. APPLICATION TO OPTIMAL TRANSPORT AND WASSERSTEIN BARYCENTER PROBLEMS

We start this section by formally defining the discrete Optimal Transport problem.

Let $C \in \mathbb{R}_+^{n \times n}$ and $\mu, \nu \in \mathbb{R}_+^n$ two discrete probability measures, i.e. $\mathbf{1}_n^\top \mu = \mathbf{1}_n^\top \nu = 1$ where $\mathbf{1}_n = (1, 1, \dots, 1)^\top \in \mathbb{R}^n$. We consider the discrete Optimal Transport problem

$$\min_{X \in \mathbb{R}_+^{n \times n}} \left\{ \langle C, X \rangle : X \mathbf{1}_n = \mu, X^\top \mathbf{1}_n = \nu \right\}. \quad (5)$$

The saddle-point formulation of (5) is

$$\min_{X \in \Delta} \max_{u, v} \langle C, X \rangle + \langle u, \mu - X \mathbf{1}_n \rangle + \langle v, \nu - X^\top \mathbf{1}_n \rangle, \quad (6)$$

where $\Delta = \{X \in \mathbb{R}_+^{n \times n} : \mathbf{1}_n^\top X \mathbf{1}_n = 1\}$ denotes the Simplex of square matrices of size n .

Next, we would like to bound the dual variables as $\|(u, v)\| \leq \lambda$, where $\|\cdot\|$ stands for the norm of the maximum, and λ is a positive constant. Lemma 2 in [24] guarantees that $\lambda = 2\|C\|$ is large enough to bound the dual variables. Here, we improve this estimate by determining the tightness value of the constant λ . See Appendix A.3 for the proof.

Lemma 3.1. *Assume $\min_{i,j} C_{i,j} = 0$. Then problem (6) admits a solution (u^*, v^*) with $\|(u^*, v^*)\| \leq \frac{\|C\|}{2}$. Moreover, there exist a matrix C and marginals μ, ν such that for any dual solution (u^*, v^*) of problem (6) it holds $\|(u^*, v^*)\| \geq \frac{\|C\|}{2}$.*

In equations (1) and (6) the role of primal and dual is interchangeable. As the dual space \mathcal{Y} , we consider the space of square matrices of size n with the norm $\|X\|_1 = \sum_{i,j} |X_{ij}|$. In this setting it is classical to use the entropy kernel given by $\xi_{\mathcal{Y}}(X) = \sum_{ij} X_{ij} \ln X_{ij}$, that

defines the Bregman distance

$$D_{\mathcal{Y}}(X, \bar{X}) = \langle X, \ln X - \ln \bar{X} \rangle.$$

We recall that Pinsker's inequality (see for instance [47]) expresses that the entropy $\xi_{\mathcal{Y}}(X)$ is 1-strongly convex with respect to the norm $\|\cdot\|_1$ as long as X remains in the simplex, and the associated Bregman distance it is nothing but the Kullback-Leibler divergence. For the primal space \mathcal{X} , we consider $\mathbb{R}^n \times \mathbb{R}^n$ with the usual Euclidean norm denoted by $\|\cdot\|_2$ and the kernel $\xi_{\mathcal{X}}(u, v) = \frac{1}{2}\|(u, v)\|_2^2$.

Comparing the structure of (1) and (6), we observe that $g(u, v) = -\langle (u, v), (\mu, \nu) \rangle$, which is linear; $h^*(X) = \langle C, X \rangle + \delta_{\Delta}(X)$, which is the sum of a linear and a characteristic function; and $K(u, v) = u \otimes \mathbf{1}_n + \mathbf{1}_n \otimes v$, where \otimes denotes the tensor product. Moreover, we can compute explicitly

$$L = \sup_{\|X\|_1 \leq 1, \|(u, v)\|_2 \leq 1} \langle X, u \otimes \mathbf{1}_n + \mathbf{1}_n \otimes v \rangle = \sqrt{2}.$$

3.1. Non-regularized optimal transport. We first analyze the non-regularized version of the Optimal Transport problem. Consider problems in steps 6 and 7 of Algorithm 1 which are called the dual and primal proximal problems, respectively. For the saddle-point formulation of OT (6), these problems admits explicit solutions, namely

$$\begin{aligned} X^{k+1} &= \mathcal{N} \left(X^k \exp \left(-\sigma_k (C - \bar{u}^k \otimes \mathbf{1}_n - \mathbf{1}_n \otimes \bar{v}^k) \right) \right), \\ u^{k+1} &= \Pi_{[-\lambda, \lambda]} \left(u^k + \tau_k (\mu - X^{k+1} \mathbf{1}_n) \right), \\ v^{k+1} &= \Pi_{[-\lambda, \lambda]} \left(v^k + \tau_k (\nu - (X^{k+1})^\top \mathbf{1}_n) \right), \end{aligned} \tag{7}$$

where $\mathcal{N}(X) = X/\|X\|_1$ denotes the renormalization of matrix X , and $\Pi_{[-\lambda, \lambda]}$ denotes the component-wise projection on the interval $[-\lambda, \lambda]$. We remark that the computation of (u^{k+1}, v^{k+1}) is practically inexpensive.

As the starting points of Algorithm 1 we take $X^1 = \frac{1}{n^2}(\mathbf{1}_n \otimes \mathbf{1}_n)$, and $u^0 = v^0 = u^1 = v^1 = 0_{\mathbb{R}^n}$. With this, it is easy to check that $D_{\mathcal{Y}}(X, X^1) \leq 2 \ln n$ for all $X \in \Delta$, and $\frac{1}{2}\|(u, v) - (u^0, v^0)\|_2^2 \leq n\lambda^2$ for all pairs (u, v) such that $\|(u, v)\| \leq \lambda$. Note that alternatively we may choose $X^1 = \mu \otimes \nu$ which yields similar bounds.

For slightly improved convergence rates, we rely on the following Lemma which improves the estimate of Lemma 7 in [2] by a factor 2. See Appendix A.4 for its proof.

Lemma 3.2. *Algorithm 2 in [2] (“ROUND”), denoted \mathcal{A} , when receiving a matrix $X \in \Delta$ and marginals μ, ν , returns in $O(n^2)$ arithmetic operations a matrix $\mathcal{A}(X) \in \Delta$ such that $\mathcal{A}(X)\mathbf{1}_n = \mu$, $\mathcal{A}(X)^\top \mathbf{1}_n = \nu$, and*

$$\|X - \mathcal{A}(X)\|_1 \leq \|\mu - X\mathbf{1}_n\|_1 + \|\nu - X^\top \mathbf{1}_n\|_1.$$

Proposition 3.3. *The HPD algorithm with linesearch finds an ε -optimal solution to (5) in $\tilde{O}\left(\frac{n^{5/2}\|C\|}{\varepsilon}\right)$ arithmetic operations.*

Proof. Let \hat{X}^N and (\hat{u}^N, \hat{v}^N) the solutions reported by the algorithm after N iterations. As the functions involved in OT are linear we have $\gamma = 0$.

Denote X^* as the true optimum of (5). Using Cauchy-Schwarz's inequality and Lemma 3.2, we have

$$\begin{aligned} \langle C, \mathcal{A}(\hat{X}^N) - X^* \rangle &= \langle C, \mathcal{A}(\hat{X}^N) - \hat{X}^N + \hat{X}^N - X^* \rangle \\ &\leq \langle C, \hat{X}^N - X^* \rangle + \lambda(\|\mu - \hat{X}^N \mathbf{1}_n\|_1 + \|\nu - (\hat{X}^N)^\top \mathbf{1}_n\|_1). \end{aligned}$$

The last right hand side is exactly the duality-gap at $X = X^*$, $u = \lambda \text{sign}(\mu - \hat{X}^N \mathbf{1}_n)$, and $v = \lambda \text{sign}(\nu - (\hat{X}^N)^\top \mathbf{1}_n)$. Then, using the bound of Theorem 2.2, and optimizing on β, τ_0 with the constraint $\tau_0 \sqrt{\beta} L = 1$ leads to $\tau_0 = \frac{\lambda \sqrt{n}}{L \sqrt{2 \ln n}}$ and $\beta = \frac{2 \ln n}{\lambda^2 n}$. Replacing and recalling that $T_N \gtrsim N$, we obtain the rate $\tilde{O}\left(\frac{\sqrt{n}\|C\|}{\varepsilon}\right)$ for the number of iterations. The result follows by noticing that each iteration requires component-wise operations over a square matrix of size n , then $O(n^2)$ arithmetic operations per iteration. \square

3.2. Entropy regularized Optimal Transport. In this section we consider the entropic regularized version of the saddle-point formulation of OT. To do this, we add an entropy regularization term to (6), so now $h^*(X) = \langle C, X \rangle + \gamma \langle X, \ln X \rangle + \delta_\Delta(X)$, with $\gamma > 0$. Clearly the new function h^* is γ -strongly convex on Δ relative to the entropy kernel and we obtain the entropic regularized saddle-point formulation for OT

$$\min_{X \in \Delta} \max_{u, v} \langle C + \gamma \ln X, X \rangle + \langle u, \mu - X \mathbf{1}_n \rangle + \langle v, \nu - X^\top \mathbf{1}_n \rangle. \quad (8)$$

Note that the same formulation was recently used in [45, Sec. 5.2] to tackle inverse problems on discrete measures.

We emphasize that the new entropy term in equation (8) only affects the solution of the proximal problem associated to the matrix X , which is now given by

$$X^{k+1} = \mathcal{N} \left(\exp \left(\frac{1}{1 + \sigma_k \gamma} \left(\ln X^k - \sigma_k (C - \bar{u}^k \otimes \mathbf{1}_n - \mathbf{1}_n \otimes \bar{v}^k) \right) \right) \right).$$

Proposition 3.4. *Let $\varepsilon > 0$ and $\gamma = \frac{\varepsilon}{4 \ln n}$. The HPD algorithm for γ entropy regularized OT requires $O\left(\frac{n^{5/2}\|C\|}{\varepsilon}\right)$ arithmetic operations to achieve a ε -optimal solution of (5).*

Proof. In order to bound the number of iterations we proceed similarly to the proof of [18, Thm 4]. Let X_γ^* the optimal solution of the entropy regularized problem, then

$$\begin{aligned} \langle C, \mathcal{A}(\hat{X}^N) - X^* \rangle &= \langle C, \mathcal{A}(\hat{X}^N) - \hat{X}^N \rangle + \langle C, \hat{X}^N - X_\gamma^* \rangle + \langle C, X_\gamma^* - X^* \rangle \\ &\leq \langle C, \hat{X}^N - X_\gamma^* \rangle + \gamma(\xi_{\mathcal{Y}}(\hat{X}^N) - \xi_{\mathcal{Y}}(X_\gamma^*)) + \lambda(\|\mu - \hat{X}^N \mathbf{1}_n\|_1 + \|\nu - (\hat{X}^N)^\top \mathbf{1}_n\|_1) \\ &\quad + \langle C, X_\gamma^* - X^* \rangle + \gamma(\xi_{\mathcal{Y}}(X_\gamma^*) - \xi_{\mathcal{Y}}(X^*)) + \gamma(\xi_{\mathcal{Y}}(X^*) - \xi_{\mathcal{Y}}(\hat{X}^N)). \end{aligned}$$

Here, the inequality is due to Lemma 3.2. In the last expression, the first line correspond to the duality-gap at $X = X_\gamma^*$, and $(u, v) = (\lambda \text{sign}(\mu - \hat{X}^N \mathbf{1}_n), \lambda \text{sign}(\nu - (\hat{X}^N)^\top \mathbf{1}_n))$. In the second line, the sum of the first two terms is ≤ 0 because of the optimality of X_γ^* , while the third term is $\leq 2 \ln n$ cause $\xi_{\mathcal{Y}}(X) \in [-2 \ln n, 0]$ for all $X \in \Delta$. We recall that for OT the function $\mathcal{L}(\cdot, y)$ is linear so the bound in equation (4) implies

$$\langle C, \mathcal{A}(\hat{X}^N) - X^* \rangle \leq \frac{1}{T_N} \left(n\lambda^2 + \frac{1}{\beta_1} 2 \ln n \right) + 2\gamma \ln n. \quad (9)$$

Given a tolerance $\varepsilon > 0$, we take $\gamma = \frac{\varepsilon}{4 \ln n}$. On the other hand, we can choose a large β_1 in order to minimize the factor in parenthesis, or we may choose $\beta_1 = O(\frac{2 \ln n}{n\lambda^2})$ in order to equate the two terms in the sum. Note that smaller values of β_1 allow larger initial stepsizes τ_1 as the initialization implies $\tau_1 = 1/\sqrt{2\beta_1}$ (recalling $L = \sqrt{2}$). We discuss the impact of different choices for β_1 in the experimental Section 5. Finally, using Theorem 2.2 we get $T_N \gtrsim \gamma \rho^2 / 32N^2$, so at least $\tilde{O}(\frac{\sqrt{n}\|C\|}{\rho\varepsilon})$ iterations are required to reach the precision ε . \square

Remark 3.5. Interestingly, the entropy regularized algorithm's convergence rate is the same as for the method without regularization. In the experimental section, we observe that the regularized version performs better. On the other hand, the non-regularized algorithm does not depend on the prefixed tolerance ε , so it might be helpful when the tolerance is not defined beforehand or whenever an early stopping criterion is required.

3.3. Application to Wasserstein Barycenter. Given discrete distributions μ_1, \dots, μ_m , cost matrices C_1, \dots, C_l , and weights $w_1, \dots, w_m \geq 0$ with $\sum_i w_i = 1$, the Wasserstein Barycenter problem is to determine a central distribution ν that minimize the weighted OT distance to the given distributions. Similar to OT, we consider the the saddle-point formulation of this problem

$$\min_{\nu \in \Delta, X^l \in \Delta} \max_{u, v} \sum_{l=1}^m w_l \left(\langle C_l, X_l \rangle + \langle u_l, \mu_l - X_l \mathbf{1}_n \rangle + \langle v_l, \nu - X_l^\top \mathbf{1}_n \rangle \right). \quad (10)$$

Here, it may be convenient to eliminate the barycenter (variable ν) by replacing it with one of the constraints, for instance $\nu = X_m^\top \mathbf{1}_n$. By doing this, and defining conveniently v^m such that $\sum_{l=1}^m w_l v_l = 0$, we have

$$\sum_{l=1}^m w_l \langle v_l, \nu - X_l^\top \mathbf{1}_n \rangle = - \sum_{l=1}^m w_l \langle v_l, X_l^\top \mathbf{1}_n \rangle.$$

Analogously to OT, we choose the dual space to be the space of tensors $X = (X_1, \dots, X_m) \in \Delta^m$ with the weighted Bregman distance given by $D_Y(X, \bar{X}) = \sum_l w_l \langle X_l, \ln X_l - \ln \bar{X}_l \rangle$ (which correspond to the weighted entropy kernel $\xi_Y(X) = \sum_l w_l \langle X_l, \ln X_l \rangle$). Similarly, we consider the weighted Euclidean prox $\xi_{\mathcal{X}}(u, v) = \frac{1}{2} \sum_l w_l \| (u_l, v_l) \|_2^2$.

Under the considerations above, we recover explicit solutions for the proximal dual and primal problems similar to the ones of equation (7), namely, for the dual

$$X_l^{k+1} = \mathcal{N} \left(X_l^k \exp \left(-\sigma_k (C_l - \bar{u}_l^k \otimes \mathbf{1}_n - \mathbf{1}_n \otimes \bar{v}_l^k) \right) \right), \quad \text{for all } l = 1, \dots, m,$$

and then for the primal,

$$u_l^{k+1} = \Pi_{[-\lambda, \lambda]} \left(u_l^k + \tau_k (\mu^l - X_l^{k+1} \mathbf{1}_n) \right), \quad \text{for all } l = 1, \dots, m,$$

$$v_l^{k+1} = \Pi_{[-\lambda, \lambda]} \left(v_l^k + \tau_k (X_m^{k+1} - X_l^{k+1})^\top \mathbf{1}_n \right), \quad \text{for all } l = 1, \dots, m-1.$$

Finally, the last variable v_m updates as $v_m^{k+1} = -\frac{1}{w_m} \sum_{l=1}^{m-1} w_l v_l^{k+1}$.

The number of iterations necessary to reach the precision ε can be bounded in the same way as for OT leading to $\tilde{O}(\frac{\sqrt{n}\|C\|}{\varepsilon})$ iterations. We remark that each iteration requires m independent component-wise operations over a square matrix of size n , hence a total of $\tilde{O}(\frac{mn^{5/2}\|C\|}{\varepsilon})$ arithmetic operations, matching the rates of the recent methods in the literature.

Remark 3.6 (Fixed marginal). In the settings described so far we introduce variables (u, v) that correspond to constraints $X\mathbf{1}_n = \mu$ and $X^\top \mathbf{1}_n = \nu$, respectively. We remark that it is possible to eliminate one of these variables by keeping the respective constraints in the optimization problem. For instance, we eliminate the variable u by changing the constraint $X \in \Delta$ for $X \in \Delta_\mu := \{X \in \mathbb{R}_+^{n \times n} : X\mathbf{1}_n = \mu\}$. We refer to this as the “fixed marginal approach”. In this case, the renormalization for the variable X is now given by $[\mathcal{N}_\mu(X)]_{ij} = \mu_i \frac{X_{ij}}{\sum_{j'} X_{i,j'}}$. Note that the reduction on the dimension also has consequences for the computation of the parameter L , which goes down from $\sqrt{2}$ to 1 in OT problems. Although these modifications only improve on the constants, we will see in the experimental section that fixing a marginal leads to smaller values of the duality-gap and improves the overall performance at no additional cost.

4. EXTENSIONS AND VARIANTS

4.1. Scaled entropy kernel. In this section we introduce a new scaled entropy kernel. This new entropy function allows to solve the Optimal Transport problem in a computational stable way by bounding the elements of the matrix X away from 0. Hence, we avoid computing the logarithm of very small numbers, which is an usual drawback of methods

based on entropic regularization. Given $\delta \in (0, 1)$, we define the function:

$$\xi_\delta(X) = \frac{1}{(1-\delta)^2} \xi \left((1-\delta)X + \frac{\delta}{n^2} \mathbf{1}_{n \times n} \right), \quad (11)$$

where $\xi(X) = \sum_{ij} X_{ij} \ln X_{ij}$ is the usual entropy kernel. The natural domain for ξ_δ is $\{X \in \mathbb{R}^{n \times n} : X_{ij} \geq -\frac{\delta}{n^2}, \sum_{ij} X_{ij} = 1\}$, however, we restrict to matrices with non-negative entries by adding to the problem the constraint $X \geq 0$.

The modified entropy inherits the properties of a kernel function. In particular, ξ_δ is 1-strongly convex w.r.t. to the norm 1. Moreover, denoting $X^\delta = (1-\delta)X + \frac{\delta}{n^2} \mathbf{1}_{n \times n}$, we have that the Bregman distance associated to ξ_δ is $D_\delta(X, \bar{X}) = \frac{1}{(1-\delta)^2} D(X^\delta, \bar{X}^\delta)$, where D is the usual Kullback–Leibler divergence generated by ξ . Finally, we remark that if $X \geq 0$ is feasible, i.e. $X \mathbf{1}_n = \mu$ and $X^\top \mathbf{1}_n = \nu$, then $X^\delta \geq \frac{\delta}{n^2}$ and satisfies the marginal constraints $X^\delta \mathbf{1}_n = \mu^\delta$ and $(X^\delta)^\top = \nu^\delta$ for the δ -regularized marginals $\mu^\delta = (1-\delta)\mu + \frac{\delta}{n} \mathbf{1}_n$ and $\nu^\delta = (1-\delta)\nu + \frac{\delta}{n} \mathbf{1}_n$.

At first sight, the idea of the scaled entropy seems nonsense as the usual entropy is precisely introduced to ensure the positivity constraint, while the new scaled entropy makes the computation of the prox difficult again. However, as we show in the next section, for OT and WB, these difficulties are overcome by an iterative subroutine that computes an approximated solution to the prox in a few iterations. Moreover, we show in the numerical experiments that this approach leads to sparse solutions of OT and WB that may be of interest in applications (see for instance [41, 19, 9]).

4.1.1. *Optimal Transport with scaled entropy.* We start by recalling the dual proximal problem (step 6 in Algorithm 1). By incorporating the scaled entropy kernel, and getting rid of the index k , solving this prox problem is equivalent to solve

$$\min_{X^\delta \in \Delta^\delta} \left\langle C - \bar{u} \otimes \mathbf{1}_n - \mathbf{1}_n \otimes \bar{v}, X^\delta \right\rangle + \frac{1}{\sigma} D(X^\delta, \bar{X}^\delta), \quad (12)$$

where \bar{X}^δ is the scaled transformation of the last iterate, and $\Delta^\delta = \{X \in \Delta : X_{ij} \geq \frac{\delta}{n^2}\}$.

These new constraints guarantee that the components of X^δ remain bounded away from 0, nevertheless, there is no longer an explicit solution for problem (12). In the following, we detail a simple and efficient method to compute an accurate solution of (12).

Consider the optimality conditions of (12): there exist multipliers $\varrho \in \mathbb{R}$, $\varrho_{ij} \geq 0$ such that:

$$X_{ij}^\delta = \bar{X}_{ij}^\delta \exp(-\sigma(C_{i,j} - \bar{u}_i - \bar{v}_j) - \varrho + \varrho_{ij}). \quad (13)$$

Furthermore, the complementary conditions implies that if $\varrho_{ij} > 0$ then $X_{ij}^\delta = \frac{\delta}{n^2}$. Denote $s = \exp(\varrho)$, one has therefore

$$X_{ij}^\delta = \max \left\{ \frac{1}{s} \bar{X}_{ij}^\delta \exp(-\sigma(C_{i,j} - \bar{u}_i - \bar{v}_j)), \frac{\delta}{n^2} \right\}.$$

Summing over i and j , and recalling the constraint $\sum_{ij} X_{ij}^\delta = 1$ yields

$$s = \sum_{i,j} \max \left\{ \bar{X}_{ij}^\delta \exp(-\sigma(C_{i,j} - \bar{u}_i - \bar{v}_j)), s \frac{\delta}{n^2} \right\}.$$

Let us define $Z_{ij} := \bar{X}_{ij}^\delta \exp(-\sigma(C_{i,j} - \bar{u}_i - \bar{v}_j))$. Then the problem reduces to computing the value $s > 0$, which is the unique root of the piecewise affine and concave real function

$$F(s) := s - \sum_{ij} \max \left\{ Z_{ij}, s \frac{\delta}{n^2} \right\}. \quad (14)$$

Moreover, one has $F(0) = -\sum_{ij} Z_{ij} < 0$ and $F(n^2 \max_{ij} Z_{ij} / \delta) = (1 - \delta)s \geq 0$, hence s can be approximated using Newton's method starting from $s^0 = 0$. We remark that $F'(s)$ has the simple expression

$$F'(s) = 1 - \frac{\delta}{n^2} \left| \left\{ (i, j) : Z_{ij} \leq s^k \frac{\delta}{n^2} \right\} \right|,$$

where $|\cdot|$ denotes the cardinality of a set. Then, the iteration of Newton's method reads:

$$s^{k+1} = \frac{\sum_{(i,j): Z_{ij} > s^k \delta / n^2} Z_{ij}}{1 - \delta + \frac{\delta}{n^2} |\{(i, j) : Z_{ij} > s^k \frac{\delta}{n^2}\}|}.$$

Alternatively, the value of s can be approximated using the classical Banach-Picard iteration to solve the fixed point problem $s = Ts$ for the operator $Ts = \sum_{ij} \max \{ Z_{ij}, s \frac{\delta}{n^2} \}$. Note that for all $s, s' \geq 0$ we have that $|Ts' - Ts| \leq \delta |s' - s|$, so T is a δ -contraction and we have the estimate $|s^k - s^*| \leq \delta^k |s^* - s^0|$, where s^* is the unique fixed point of T . For instance, with $\delta = 0.1$, we reach the tolerance ϵ after $\left\lceil \log_{10} \frac{|s^*|}{\epsilon} \right\rceil$ iterations, so that in this case the number of significant digits equals the number of iterations. In practice, the Newton method achieves a good precision in even fewer iterations.

Remark 4.1. One can also be tempted to solve the OT problem by alternating maximization on the dual of the scaled entropy-regularized assignment problem, in a sort of “scaled Sinkhorn algorithm”. This works well in practice, but a (good) theoretical bound on the number of iterations needed for solving each maximization is missing.

4.1.2. Accelerated Gradient Descent with scaled entropy. The scaled entropy can also be incorporated into other methods in the literature which rely on entropy smoothing, as is the case for the Accelerated Gradient Descent.

Let us consider the Optimal Transport problem:

$$\min_{X \geq 0} \{ \langle C, X \rangle : X \mathbf{1}_n = \mu, X^\top \mathbf{1}_n = \nu \}.$$

Now, we change the variable X for $X^\delta = (1 - \delta)X + \frac{\delta}{n^2}$. Straightforward computations show that the new problem is equivalent (in the sense that they have the same optimum)

to

$$\min \left\{ \langle C, X^\delta \rangle : X^\delta \geq \frac{\delta}{n^2}, X^\delta \mathbf{1}_n = \mu^\delta, (X^\delta)^\top \mathbf{1}_n = \nu^\delta \right\}.$$

Now we add the entropic regularization term $\gamma \langle X^\delta, \ln X^\delta \rangle$, and, following [18] find the dual:

$$\begin{aligned} \max_{u,v} \varphi(u, v) &:= \langle u, \mu^\delta \rangle + \langle v, \nu^\delta \rangle + V(u, v), \quad \text{with:} \\ V(u, v) &:= \min_{X^\delta \in \Delta^\delta} \langle C - u \otimes \mathbf{1}_n - \mathbf{1}_n \otimes v, X^\delta \rangle + \gamma \langle X^\delta, \ln X^\delta \rangle. \end{aligned} \quad (15)$$

If we denote by $X^\delta(u, v)$ the solution of problem (15), then we have $\nabla_u \varphi(u, v) = \mu^\delta - X^\delta(u, v) \mathbf{1}_n$, and $\nabla_v \varphi(u, v) = \nu^\delta - X^\delta(u, v)^\top \mathbf{1}_n$.

The main difficult is that there is no explicit solution for problem (15). However, following the construction of the past section, we can approximate the solution $X^\delta(u, v)$ by solving the fixed-point equation

$$s = \sum_{ij} \max \left\{ \exp \left(\frac{-C_{ij} + u_i + v_j}{\gamma} \right), s \frac{\delta}{n^2} \right\}.$$

4.2. Nonlinear penalization. Here we observe that an advantage of the saddle-point formulation and the methods studied in this paper is that they can be recast without effort to address nonlinear variants of Optimal Transport and Wasserstein Barycenter problems. Nonlinear extensions cover interesting problems such as unbalanced OT and WB problems, which could not be tackled easily with exact linear optimization-based methods.

We start by defining a simple case of nonlinear Optimal Transport. Let $\psi : \mathbb{R}^n \rightarrow \mathbb{R}_+$ a convex function and consider the following penalized OT problem and its saddle-point formulation:

$$\min_{X \in \Delta_\mu} \langle C, X \rangle + \psi(\nu - X^\top \mathbf{1}_n) = \min_{X \in \Delta_\mu} \max_v \langle C - \mathbf{1}_n \otimes v, X \rangle + \langle v, \nu \rangle - \psi^*(v). \quad (16)$$

The problem above can be solved with the Hybrid Primal Dual method as soon as the proximal problem

$$\arg \min_v \langle v, X^\top \mathbf{1}_n - \nu \rangle + \psi^*(v) + \frac{1}{2\tau} \|v - \bar{v}\|_2^2, \quad (17)$$

can be solved efficiently. Note that the proximal problem associated to X does not depend on the penalization ψ , therefore the solution does not change (see Sec. 3.1).

Formulation (16) can be used to deal with unbalanced Optimal Transport, i.e. an Optimal Transport problem where the two marginals have different total mass (*cf* for instance [40, Sec. 10.2]). In this setting, the classical OT problem (5) does not admit a feasible solution, and one looks for a transportation matrix that satisfies the marginal constraints approximately. (In addition, we may assume that $\mu \cdot \mathbf{1}_n = 1$ so that X still belongs to the

simplex — if not, it belongs to a simplex multiple of the unit simplex and our analysis is easy to adapt).

Several penalization functions can be used to approximate the constraint $\nu - X^\top \mathbf{1}_n = 0$, see again [40]. For instance, defining $\psi = \frac{1}{2\eta} \|\cdot\|_2^2$ leads to quadratic penalization balanced with the parameter $\eta > 0$. Remark that when $\eta \rightarrow 0$, problem (16) becomes equivalent to the classical formulation (5). Moreover, for a quadratic penalization, (17) has an explicit solution given by $v = \frac{\eta}{\eta + \tau} (\bar{v} + \tau(\nu - X^\top \mathbf{1}_n))$. One may also consider a total variation penalization, that is $\psi = \alpha \|\cdot\|_1$. It leads to the explicit update $v = \Pi_{[-\alpha, \alpha]} (\bar{v} + \tau(\nu - X^\top \mathbf{1}_n))$. Note that if $\alpha = \lambda = \|C\|_\infty / 2$, this is already the update of equation (7).

The same ideas can also be extended to the setting of unbalanced Wasserstein Barycenter. In that case, we may consider the saddle-point formulation

$$\min_{\nu \geq 0, X^l \in \Delta_\nu} \max_v \sum_{l=1}^m w_l (\langle C_l - \mathbf{1}_n \otimes v_l, X_l \rangle + \langle v_l, \nu \rangle - \psi_l^*(v_l)), \quad (18)$$

where ψ_l are the penalization functions. Then, we recover explicit updates for all $l = 1, \dots, m$

$$X_l^{k+1} = \mathcal{N}_\mu \left(X_l^k \exp \left(-\sigma_k (C_l - \mathbf{1}_n \otimes \bar{v}_l^k) \right) \right), \quad \nu_l^{k+1} = \nu^k \exp \left(-\sigma_k \sum_{l=1}^m w_l \bar{v}_l \right),$$

and $v_l^{k+1} = \frac{\eta}{\eta + \tau_k} \left(v_l^k + \tau_k (\nu^{k+1} - (X_l^{k+1})^\top \mathbf{1}_n) \right)$ in the case of quadratic penalization, or $v_l^{k+1} = \Pi_{[-\eta, \eta]} \left(v_l^k + \tau_k (\nu^{k+1} - (X_l^{k+1})^\top \mathbf{1}_n) \right)$ for total variation penalization.

An example of this using quadratic penalization is computed in Figure 6 (bottom right). Observe however that a specific complexity analysis should be performed here, depending on the nonlinearity (as for instance we cannot use straight out of the box a result such as lemma 3.1 in this particular setting), which is out of the scope of the present study.

5. NUMERICAL EXPERIMENTS

In this section, we present numerical results over different instances of Optimal Transport and Wasserstein Barycenter problems for the algorithms described in this paper and for methods in the literature. We implement the methods using `Python-torch` and run the algorithms on a MacBook Air M1 Octa Core with 8GB RAM. We thank authors in [21] and [17] for sharing their implementations in `Python-numpy`, which we replicate in `Python-torch` in order to have a fair comparison.

5.1. Instances. For Optimal Transport problems we consider the following instances.

MNIST instance: we randomly choose two images from the MNIST library, normalize and use it as the marginals μ and ν . Note that images in the MNIST library are handwritten numbers of size 28×28 pixels, hence $n = 784$ in this case. For the cost matrix we stick to the Euclidean distance setting.

Gaussian instance: we measure the OT distance between distributions $\mu \sim \mathcal{N}(3, 1) + \mathcal{N}(7, 1)$, and $\nu \sim \mathcal{N}(5, 1)$. To generate a discretization of this problem, we consider a partition of the interval $[0, 10]$ into n equidistant points, then use the Euclidean distance as a metric.

Random instance: given $n > 0$, we draw the components of the vectors $\mu, \nu \in \mathbb{R}^n$, and the matrix $C \in \mathbb{R}^n \times \mathbb{R}^n$ independently from an uniform distribution in $[0, 1]$.

Corner to dense: in these instances, the first marginal correspond to an image of size $n_{\text{pix}} \times n_{\text{pix}}$ where the source is concentrated in the top-left corner (see Figure 1). As a second marginal, we consider an image of the same size with uniform distribution of the mass. We use these instances to measure the performance of exact solvers when transporting the mass from a sparse into a dense marginal. We remark that for those instances the marginals are unbalanced, therefore we normalize it so the total mass sum up to one.

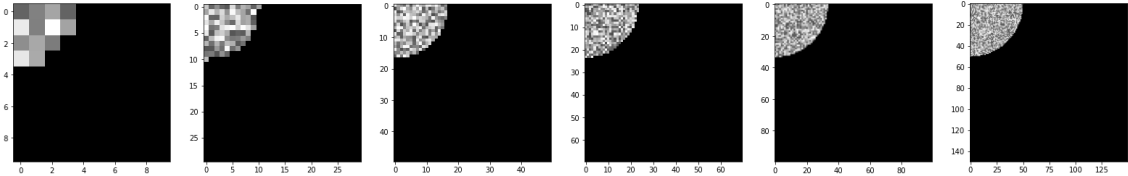


FIGURE 1. Corner images for $n \in \{10, 30, 50, 70, 100, 150\}$

5.2. Exact solutions. In order to compute exact solutions for the instances, we use the Python Optimal Transport (POT) library [20]. In particular, for optimal transport we use the `ot.emd` function that is based on the algorithm in [11], while for Wasserstein Barycenter problems we use the `ot.lp.barycenter` function which is based on Network Simplex. Table 2 shows the performance of the exact solver and our approximated method for different instances and sizes of Optimal Transport problems.

We remark that solvers in the POT library are extremely fast to solve medium size instances (see Table 2). However, for large instances ($n > 10000$), we did not obtain a solution after 30 minutes of computing time. Note that exact solvers like those in POT are not designed to deal with very large problems (see the POT documentation for extra details). Therefore, approximate methods such as those described in this paper can be of practical interest whenever n is large.

Exact solver								
Corner to dense			Gaussian			Random		
n_{pix}	time (s)	ℓ^1 error	n	time (s)	ℓ^1 error	n	time (s)	ℓ^1 error
10	0.0038	1.10e-07	100	0.0011	3.27e-16	100	0.0015	4.18e-16
30	0.5275	1.15e-07	500	0.0238	6.90e-16	500	0.0227	6.97e-16
50	10.5236	2.64e-07	1000	0.1495	1.02e-15	1000	0.1012	7.83e-16
70	92.6429	8.84e-08	5000	3.7442	1.97e-15	5000	4.8426	1.12e-15
100	-	-	10000	148.89	2.39e-15	10000	87.1855	1.60e-15
120	-	-	15000	-	-	20000	-	-
Regularized Hybrid Primal Dual with linesearch								
Corner to dense			Gaussian			Random		
n_{pix}	time (s)	precision	n	time (s)	precision	n	time (s)	precision
10	0.0017	0.0089	100	0.0554	0.0095	100	0.1600	0.0098
30	4.0481	0.0064	500	0.5636	0.0097	500	0.3982	0.0069
50	30.0447	0.0099	1000	1.8527	0.0091	1000	0.7227	0.0059
70	396.647	0.0097	5000	119.69	0.0098	5000	21.9500	0.0046
100	558.39	0.0215*	10000	139.52	0.0175	10000	274.62	0.0038
120	1437.33	0.0230*	15000	502.60	0.0201*	20000	494.80	0.0021*

TABLE 2. Exact solver versus the main algorithm proposed in this paper. The precision measures the distance to the optimal value. For the values with (*), the precision corresponds to the duality gap.

5.3. Results for Optimal Transport instances. In our first experience, we test the effects of the parameter β_1 in the performance of the Hybrid Primal Dual method (HPD), with entropy regularization (γ), linesearch (ls), and the fixed marginals (fm) approach discussed in Remark 3.6. We refer to this method as γ -HPD ls fm. As we discussed in the proof of Proposition 3.4, there is a compromise between the size of the initial stepsize τ_1 (which decrease with β_1 increasing), and the bound (9) (which increase with β_1 increasing). Figure (2) shows the performance for different choices of $\beta_1 = O(\frac{\ln n}{n\lambda^2})$. We observe that the quality of the solution slightly improves if we amplify the quantity $\frac{\ln n}{n\lambda^2}$, meaning that it is of convenience to start with a smaller initial stepsize τ_1 , improving the bound (9). Despite this, the optimal value for β_1 it seems to depend on the characteristic of the problem ($\beta_1 \approx 10^3 \frac{\ln n}{n\lambda^2}$ for MNIST and random instances, while $\beta_1 \approx 10^2 \frac{\ln n}{n\lambda^2}$ for Gaussians instances). A theoretical study concerning a better choice of β_1 than provided by the worse case estimates remains an open question.

Next, we compare the performance of several versions of the HPD method including: non-regularized (HPD), entropy regularized (γ -HPD), and its versions with linesearch (ls) and fixed marginals (fm). Figure 3 shows that, for the three instances, the best performance is attained when we include linesearch and fixed marginals. Moreover, we observe that using fixed marginals also improves the performance of non-regularized versions of the method.

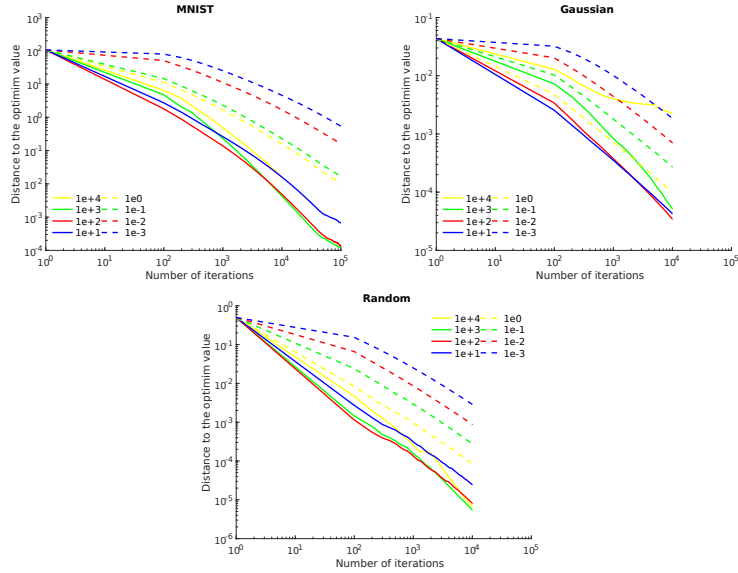


FIGURE 2. log-log distance to the optimum for γ -HPD ls fm with different values $\beta_1 = C \cdot \frac{2 \ln n}{n \lambda^2}$, with $C \in \{10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 10^2, 10^3\}$. For Gaussian and random instances we set $n = 1000$.

In the third experiment, we test the impact of the prefixed parameter $\varepsilon > 0$ in the performance of our algorithm γ -HPD ls fm, as well as the Accelerated Gradient Descent (AGD), both with and without the δ -regularization of Section 4.1 (see Figure 4). We observe that γ -HPD ls fm is very stable when the tolerance approaches zero. We remark that the sensibility to ε highly depends on the implementation strategy to avoid numerical issues. An extensive comparison of different implementation strategies escapes the scope of this paper and we leave it as possible future work. Finally, the δ -regularized methods are also stable and perform slightly better than non-regularized methods. Moreover, in Figure 5, we show the structure of the optimal transport reported by γ -HPD ls and AGD when using the scaled entropy regularization. We remark that these solutions are sparse and have a support similar to the exact optimal solution, contrarily to solutions obtained with entropy regularized methods, which are never sparse.

5.4. Results for Wasserstein Barycenter problems. For Wasserstein Barycenter problems we first use a collection of $m = 5$ images of the number five in the MNIST library. Figure 6 shows the barycenter computed for each method. We also plot the true optimum at the top-left for reference. Then, we compare the performance of the methods when solving a Wasserstein Barycenter problem between 1-dimensional Gaussian measures. For this, we generate $m = 10$ different Gaussian measures with random parameters and $n = 100$ bins in the interval $[-10, 10]$. Figure (7) shows the barycenters obtained with the algorithms of literature and the ones presented in this paper. We also plot the theoretical barycenter computed according to Theorem 2.2 in [43].

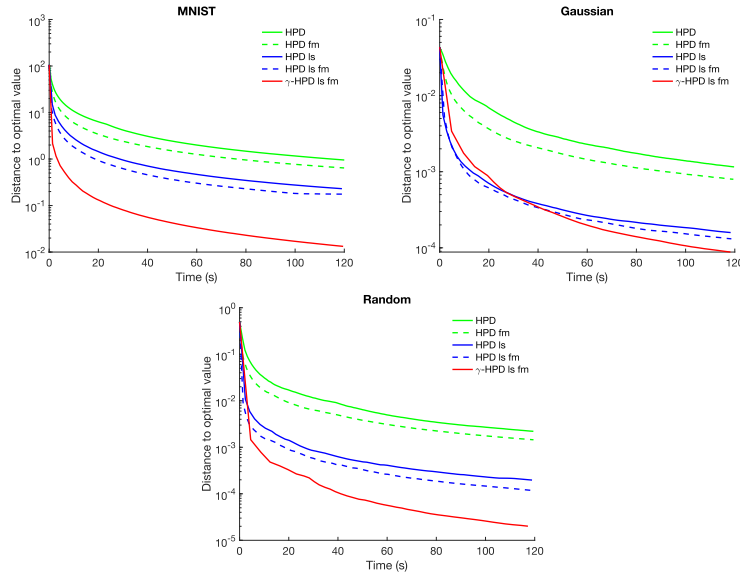


FIGURE 3. log-log distance to the optimum for different versions of the HPD method, including linesearch (ls), fixed-marginals (fm), and entropic regularization (γ).

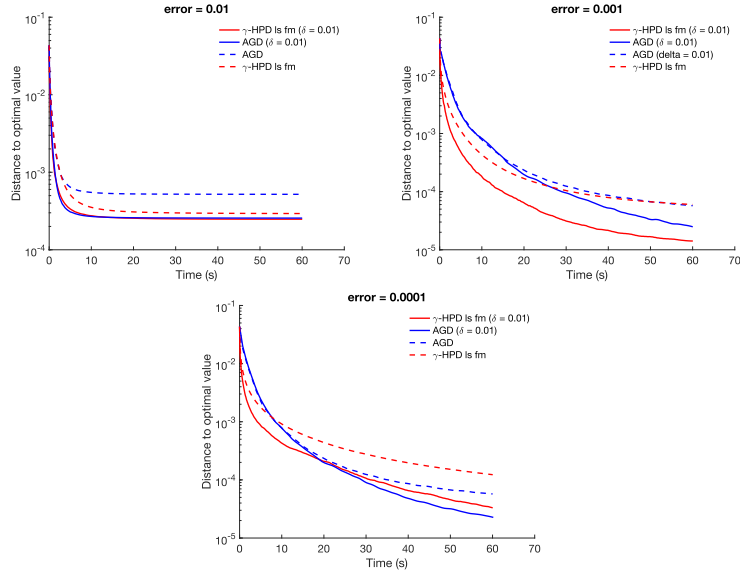


FIGURE 4. Distance to the optimum value for different prefixed tolerances.

The results for WB show that our proposed algorithms are very competitive with the approaches in the literature. In this case, the solutions with δ -regularization are slightly

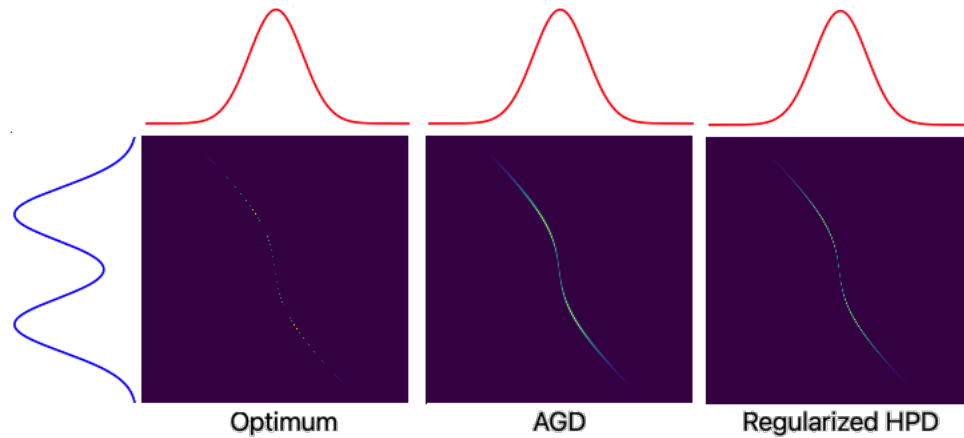


FIGURE 5. Transportation plan between Gaussian distributions reported by γ -HPDs and AGD with *delta*-regularization ($\delta = 0.01$).

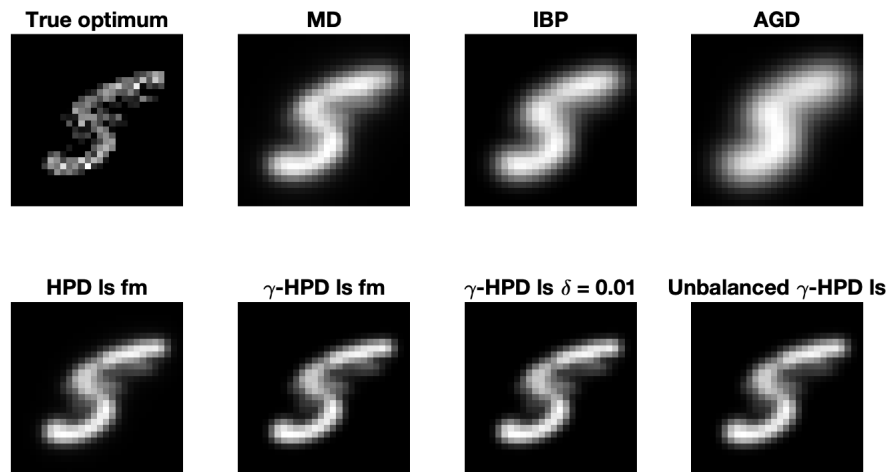


FIGURE 6. Barycenters after 10000 iterations of each method.

sharper yet not as sparse as for OT. An explanation is that the support of the barycenter results from a mixture of the different supports of multiple transportation matrices, and should therefore have a larger support. Finally, we mention that the bottom-right image in Figure 6, which is also sharp, corresponds to the regularized HPD applied directly to the unbalanced MNIST marginals and using a quadratic penalization as described in Section 4.2

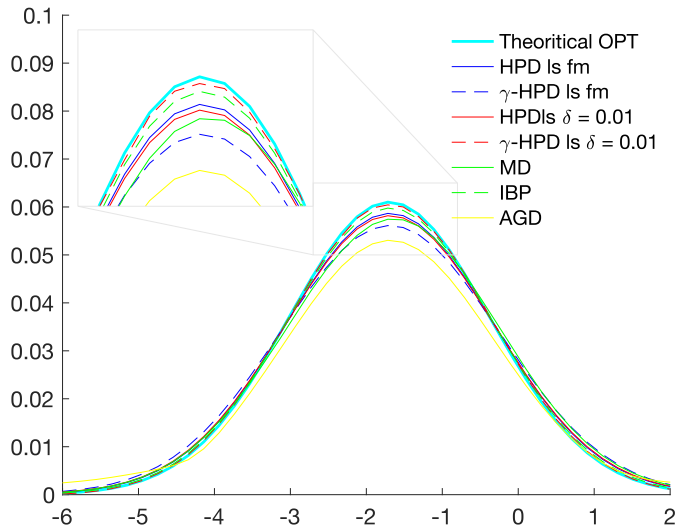


FIGURE 7. Wasserstein Barycenters computed after 10000 iterations of each method.

6. CONCLUSIONS

We have shown how to apply and extend first-order primal-dual methods for saddle-point problems in a non-Euclidean setting to approximate the solutions of Optimal Transport and Wasserstein Barycenter problems. Furthermore, the proposed methods provide a new perspective between the Mirror-Prox and the AGD as they may or may not rely on entropic regularization and use linesearch to improve the numerical performance. We believe that this new perspective is of great interest to understand the dynamics behind the accelerated OT methods and may help tackle open questions, such as the existence of a competitive $O(n^2/\varepsilon)$ algorithm. Afterward, we introduced a modified entropy kernel for the Bregman mirror descent steps that improve the solution's quality at a low computational cost. This idea can also be accomplished with AGD-type methods to obtain sparse transportation plans that are of interest in applications and cannot be obtained with the traditional approach. Numerical experiments show that the presented algorithms are very stable and competitive with the state-of-art-methods in time and precision.

ACKNOWLEDGEMENTS

The work of Juan Pablo Contreras was supported by a doctoral scholarship from ANID-PFCHA/Doctorado Nacional/2019-21190161. The authors warmly thank the reviewers of this manuscript for their comments and their very helpful suggestions.

REFERENCES

- [1] Z. ALLEN-ZHU, Y. LI, R. OLIVEIRA, AND A. WIGDERSON, *Much faster algorithms for matrix scaling*, in 58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017, IEEE Computer Soc., Los Alamitos, CA, 2017, pp. 890–901, <https://doi.org/10.1109/FOCS.2017.87>, <https://doi.org/10.1109/FOCS.2017.87>.
- [2] J. ALTSCHULER, J. NILES-WEED, AND P. RIGOLLET, *Near-linear time approximation algorithms for optimal transport via sinkhorn iteration*, in Advances in Neural Information Processing Systems, vol. 30, 2017, <https://proceedings.neurips.cc/paper/2017/file/491442df5f88c6aa018e86dac21d3606-Paper.pdf>.
- [3] M. ARJOVSKY, S. CHINTALA, AND L. BOTTOU, *Wasserstein generative adversarial networks*, in Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17, JMLR.org, 2017, p. 214–223.
- [4] H. H. BAUSCHKE, J. BOLTE, AND M. TEBoulLE, *A descent lemma beyond Lipschitz gradient continuity: first-order methods revisited and applications*, Math. Oper. Res., 42 (2017), pp. 330–348, <https://doi.org/10.1287/moor.2016.0817>, <https://doi.org/10.1287/moor.2016.0817>.
- [5] H. H. BAUSCHKE, J. M. BORWEIN, AND P. L. COMBETTES, *Essential smoothness, essential strict convexity, and Legendre functions in Banach spaces*, Commun. Contemp. Math., 3 (2001), pp. 615–647, <https://doi.org/10.1142/S0219199701000524>, <https://doi.org/10.1142/S0219199701000524>.
- [6] J.-D. BENAMOU, G. CARLIER, M. CUTURI, L. NENNA, AND G. PEYRÉ, *Iterative Bregman projections for regularized transportation problems*, SIAM J. Sci. Comput., 37 (2015), pp. A1111–A1138, <https://doi.org/10.1137/141000439>, <https://doi.org/10.1137/141000439>.
- [7] J. BIGOT, R. GOUET, T. KLEIN, AND A. LÓPEZ, *Geodesic PCA in the Wasserstein space by convex PCA*, Ann. Inst. Henri Poincaré Probab. Stat., 53 (2017), pp. 1–26, <https://doi.org/10.1214/15-AIHP706>, <https://doi.org/10.1214/15-AIHP706>.
- [8] J. BLANCHET, A. JAMBULAPATI, C. KENT, AND A. SIDFORD, *Towards optimal running times for optimal transport*, arXiv preprint arXiv:1810.07717, (2018), <https://arxiv.org/abs/1810.07717>.
- [9] M. BLONDEL, V. SEGUY, AND A. ROLET, *Smooth and sparse optimal transport*, in Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics, vol. 84 of Proceedings of Machine Learning Research, PMLR, 09–11 Apr 2018, pp. 880–889, <https://proceedings.mlr.press/v84/blondel18a.html>.
- [10] N. BONNEEL, J. RABIN, G. PEYRÉ, AND H. PFISTER, *Sliced and Radon Wasserstein barycenters of measures*, J. Math. Imaging Vision, 51 (2015), pp. 22–45, <https://doi.org/10.1007/s10851-014-0506-3>, <https://doi.org/10.1007/s10851-014-0506-3>.
- [11] N. BONNEEL, M. VAN DE PANNE, S. PARIS, AND W. HEIDRICH, *Displacement interpolation using lagrangian mass transport*, in Proceedings of the 2011 SIGGRAPH Asia conference, 2011, pp. 1–12.

- [12] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, J. Math. Imaging Vision, 40 (2011), pp. 120–145, <https://doi.org/10.1007/s10851-010-0251-1>, <https://doi.org/10.1007/s10851-010-0251-1>.
- [13] A. CHAMBOLLE AND T. POCK, *On the ergodic convergence rates of a first-order primal-dual algorithm*, Math. Program., 159 (2016), pp. 253–287, <https://doi.org/10.1007/s10107-015-0957-3>, <https://doi.org/10.1007/s10107-015-0957-3>.
- [14] G. CHEN AND M. TEBoulLE, *Convergence analysis of a proximal-like minimization algorithm using Bregman functions*, SIAM J. Optim., 3 (1993), pp. 538–543, <https://doi.org/10.1137/0803026>, <https://doi.org/10.1137/0803026>.
- [15] M. B. COHEN, A. MADRY, D. TSIPRAS, AND A. VLADU, *Matrix scaling and balancing via box constrained newton’s method and interior point methods*, in 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), 2017, pp. 902–913, <https://doi.org/10.1109/FOCS.2017.88>.
- [16] M. CUTURI, *Sinkhorn distances: Lightspeed computation of optimal transport*, in Advances in Neural Information Processing Systems, vol. 26, 2013, <https://proceedings.neurips.cc/paper/2013/file/af21d0c97db2e27e13572cbf59eb343d-Paper.pdf>.
- [17] D. DVINSKIKH AND D. TIAPKIN, *Improved complexity bounds in Wasserstein barycenter problem*, in International Conference on Artificial Intelligence and Statistics, PMLR, 2021, pp. 1738–1746.
- [18] P. DVURECHENSKY, A. GASNIKOV, AND A. KROSHNIN, *Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm*, in International conference on machine learning, PMLR, 2018, pp. 1367–1376.
- [19] M. ESSID AND J. SOLOMON, *Quadratically regularized optimal transport on graphs*, SIAM J. Sci. Comput., 40 (2018), pp. A1961–A1986, <https://doi.org/10.1137/17M1132665>, <https://doi.org/10.1137/17M1132665>.
- [20] R. FLAMARY, N. COURTY, A. GRAMFORT, M. Z. ALAYA, A. BOISBUNON, S. CHAMBON, L. CHAPEL, A. CORENFLOS, K. FATRAS, N. FOURNIER, L. GAUTHERON, N. T. GAYRAUD, H. JANATI, A. RAKOTOMAMONJY, I. REDKO, A. ROLET, A. SCHUTZ, V. SEGUY, D. J. SUTHERLAND, R. TAVENARD, A. TONG, AND T. VAYER, *Pot: Python optimal transport*, Journal of Machine Learning Research, 22 (2021), pp. 1–8, <http://jmlr.org/papers/v22/20-451.html>.
- [21] S. GUMINOV, P. DVURECHENSKY, N. TUPITSA, AND A. GASNIKOV, *On a combination of alternating minimization and Nesterov’s momentum*, in International Conference on Machine Learning, PMLR, 2021, pp. 3886–3898.
- [22] W. GUO, N. HO, AND M. JORDAN, *Fast algorithms for computational optimal transport and wasserstein barycenter*, in International Conference on Artificial Intelligence and Statistics, PMLR, 2020, pp. 2088–2097.
- [23] N. HO, X. NGUYEN, M. YUROCHKIN, H. H. BUI, V. HUYNH, AND D. PHUNG, *Multilevel clustering via Wasserstein means*, in International Conference on Machine Learning, PMLR, 2017, pp. 1501–1509.

- [24] A. JAMBULAPATI, A. SIDFORD, AND K. TIAN, *A direct $\tilde{O}(1/\varepsilon)$ iteration parallel algorithm for optimal transport*, in Advances in Neural Information Processing Systems, 2019.
- [25] X. JIANG AND L. VANDENBERGHE, *Bregman primal-dual first-order method and application to sparse semidefinite programming*, Comput. Optim. Appl., 81 (2022), pp. 127–159, <https://doi.org/10.1007/s10589-021-00339-7>, <https://doi.org/10.1007/s10589-021-00339-7>.
- [26] A. KROSHNIN, N. TUPITSA, D. DVINSKIKH, P. DVURECHENSKY, A. GASNIKOV, AND C. URIBE, *On the complexity of approximating Wasserstein barycenters*, in International conference on machine learning, PMLR, 2019, pp. 3530–3540.
- [27] M. KUSNER, Y. SUN, N. KOLKIN, AND K. WEINBERGER, *From word embeddings to document distances*, in International conference on machine learning, PMLR, 2015, pp. 957–966.
- [28] G. LAN AND Y. ZHOU, *An optimal randomized incremental gradient method*, Math. Program., 171 (2018), pp. 167–215, <https://doi.org/10.1007/s10107-017-1173-0>, <https://doi.org/10.1007/s10107-017-1173-0>.
- [29] C. LEE, H. LUO, C. WEI, AND M. ZHANG, *Linear last-iterate convergence for matrix games and stochastic games*, CoRR, abs/2006.09517 (2020), <https://arxiv.org/abs/2006.09517>, <https://arxiv.org/abs/2006.09517>.
- [30] C.-W. LEE, C. KROER, AND H. LUO, *Last-iterate convergence in extensive-form games*. arXiv:2106.14326, 2021, <https://arxiv.org/abs/2106.14326>.
- [31] Y. T. LEE AND A. SIDFORD, *Efficient inverse maintenance and faster algorithms for linear programming*, in 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, IEEE, 2015, pp. 230–249.
- [32] T. LIN, N. HO, X. CHEN, M. CUTURI, AND M. JORDAN, *Fixed-support Wasserstein barycenters: Computational hardness and fast algorithm*, in Advances in Neural Information Processing Systems, vol. 33, Curran Associates, Inc., 2020, pp. 5368–5380, <https://proceedings.neurips.cc/paper/2020/file/3a029f04d76d32e79367c4b325dda4d-Paper.pdf>.
- [33] T. LIN, N. HO, AND M. JORDAN, *On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms*, in Proceedings of the 36th International Conference on Machine Learning, 2019, pp. 3982–3991.
- [34] H. LU, R. M. FREUND, AND Y. NESTEROV, *Relatively smooth convex optimization by first-order methods, and applications*, SIAM J. Optim., 28 (2018), pp. 333–354, <https://doi.org/10.1137/16M1099546>, <https://doi.org/10.1137/16M1099546>.
- [35] Y. MALITSKY AND T. POCK, *A first-order primal-dual algorithm with linesearch*, SIAM J. Optim., 28 (2018), pp. 411–432, <https://doi.org/10.1137/16M1092015>, <https://doi.org/10.1137/16M1092015>.
- [36] A. NEMIROVSKI, *Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems*, SIAM J. Optim., 15 (2004), pp. 229–251, <https://doi.org/10.1137/S1052623403425629>, <https://doi.org/10.1137/S1052623403425629>.
- [37] Y. NESTEROV, *Smooth minimization of non-smooth functions*, Math. Program.,

- 103 (2005), pp. 127–152, <https://doi.org/10.1007/s10107-004-0552-5>, <https://doi.org/10.1007/s10107-004-0552-5>.
- [38] Y. NESTEROV, *Dual extrapolation and its applications to solving variational inequalities and related problems*, Math. Program., 109 (2007), pp. 319–344, <https://doi.org/10.1007/s10107-006-0034-z>, <https://doi.org/10.1007/s10107-006-0034-z>.
- [39] Y. NESTEROV, *Efficiency of coordinate descent methods on huge-scale optimization problems*, SIAM J. Optim., 22 (2012), pp. 341–362, <https://doi.org/10.1137/100802001>, <https://doi.org/10.1137/100802001>.
- [40] G. PEYRÉ AND M. CUTURI, *Computational optimal transport*, (2018), <https://doi.org/10.48550/ARXIV.1803.00567>, <https://arxiv.org/abs/1803.00567>.
- [41] J. RABIN, S. FERRADANS, AND N. PAPADAKIS, *Adaptive color transfer with relaxed optimal transport*, in 2014 IEEE international conference on image processing (ICIP), IEEE, 2014, pp. 4852–4856.
- [42] R. T. ROCKAFELLAR, *Convex analysis*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 1997. Reprint of the 1970 original, Princeton Paperbacks.
- [43] L. RÜSCHENDORF AND L. UCKELMANN, *On the n -coupling problem*, J. Multivariate Anal., 81 (2002), pp. 242–258, <https://doi.org/10.1006/jmva.2001.2005>, <https://doi.org/10.1006/jmva.2001.2005>.
- [44] J. SHERMAN, *Area-convexity, ℓ^∞ regularization, and undirected multicommodity flow*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, 2017, pp. 452–460.
- [45] A. SILVETI-FALLS, C. MOLINARI, AND J. FADILI, *A stochastic bregman primal-dual splitting algorithm for composite optimization*, arXiv preprint arXiv:2112.11928, (2021), <https://arxiv.org/abs/2112.11928>.
- [46] P. TSENG, *On accelerated proximal gradient methods for convex-concave optimization*. (Unpublished), 2008, <https://www.mit.edu/~dimitrib/PTseng/papers/apgm.pdf>.
- [47] A. B. TSYBAKOV, *Introduction to nonparametric estimation*, Springer Series in Statistics, Springer, New York, 2009, <https://doi.org/10.1007/b13794>, <https://doi.org/10.1007/b13794>. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.

APPENDIX A. PROOFS OF THE MAIN RESULTS

A.1. Detailed proof of Theorem 2.2. We mostly focus on the difference with the proofs in [35]. In particular, the termination criterion we consider in Alg. 1 in Step 8 terminates earlier than the one in [35] and it can be shown that it induces a globally bounded (actually, depending on τ_0, β_0, ρ, L in a logarithmic way) extra multiplicative factor in the complexity.

We start by recalling the definition of a relatively strongly convex function.

Definition A.1. [4, 34] Given $\gamma > 0$, the function f is γ -strongly convex relative to ξ if for all $x, y \in \text{int dom } \xi$ it holds

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \gamma D_\xi(y, x).$$

In particular f is γ -strongly convex relative to ξ if and only if $f - \gamma\xi$ is convex over $\text{int dom } \xi$.

Let us recall the following results, the first being a particular case of the second.

Lemma A.2 (Proximal optimality condition). *Let $f : \mathcal{X} \rightarrow \mathbb{R} \cup \{+\infty\}$ a closed, convex function. Then \hat{x} is a solution of the proximal problem $\min_{x \in \mathcal{X}} f(x) + \frac{1}{2}\|x - \bar{x}\|_2^2$, if and only if*

$$\langle \hat{x} - \bar{x}, \hat{x} - x \rangle + f(x) \geq f(\hat{x}), \quad \forall x \in \mathcal{X}.$$

Lemma A.3 (Chen and Teboulle 1993). *Let $f : \mathcal{X} \rightarrow \mathbb{R} \cup \{+\infty\}$ be a closed proper convex function and x^+ an optimal solution to the nonlinear prox problem $\min_{x \in \mathcal{X}} \{f(x) + \tau^{-1}D_\xi(x, \bar{x})\}$. Then, for any $x \in \text{dom } \xi$ it holds*

$$\tau(f(x^+) - f(x)) \leq D_\xi(x, \bar{x}) - D_\xi(x, x^+) - D_\xi(x^+, \bar{x}).$$

Proof of Theorem 2.2. Let $x \in \text{dom } g$ and $y \in \text{dom } \xi_{\mathcal{Y}}$, where we recall that $\xi_{\mathcal{Y}}$ is a Legendre function (see [42]). Remark that in the following analysis (x, y) is not necessarily a saddle-point. We mostly follow [35].

Because of the convexity of g and the proximal optimality condition in Lemma A.2, we have

$$\langle x^k - x^{k-1} + \tau_{k-1}K^*y^k, x - x^k \rangle \geq \tau_{k-1}(g(x^k) - g(x)). \quad (19)$$

Taking $x = x^{k+1}$ in (19) multiplied by $\theta_k = \frac{\tau_k}{\tau_{k-1}}$, plus (19) with $x = x^{k-1}$ multiplied by θ_k^2 , we obtain

$$\langle \bar{x}^k - x^k + \tau_k K^* y^k, x^{k+1} - \bar{x}^k \rangle \geq \tau_k((1 + \theta_k)g(x^k) - g(x^{k+1}) - \theta_k g(x^{k-1})). \quad (20)$$

Using the 3-point inequality of Lemma A.3 with h^* and g gives, respectively

$$D_{\mathcal{Y}}(y, y^k) - D_{\mathcal{Y}}(y^{k+1}, y^k) - (1 + \gamma)D_{\mathcal{Y}}(y, y^{k+1}) \geq \sigma_k(h^*(y^{k+1}) - h^*(y) + \langle K\bar{x}^k, y - y^{k+1} \rangle). \quad (21)$$

$$\|x - x^k\|_2^2 - \|x - x^{k+1}\|_2^2 - \|x^{k+1} - x^k\|_2^2 \geq 2\tau_k(g(x^{k+1}) - g(x) + \langle Kx^{k+1} - Kx, y^{k+1} \rangle). \quad (22)$$

Above we assume h^* is γ -strongly convex relative to $\xi_{\mathcal{Y}}$. If h^* is just convex, then the inequalities holds with $\gamma = 0$.

Summing (20), (21) and (22) and in view of the stopping criteria for the linesearch (step 8 in Algorithm 1), reproducing the same computation as [35] (but here (x, y) is not a saddle-point and we cannot rely on the non-negativity of expressions such as $\mathcal{L}(x^k, y)$):

$$(1 + \theta_k)\mathcal{L}(x^k, y) - \theta_k\mathcal{L}(x^{k-1}, y) - \mathcal{L}(x, y^{k+1}) \leq \frac{1}{2\tau_k}\|x^k - x\|_2^2 - \frac{1}{2\tau_k}\|x - x^{k+1}\|_2^2 + \frac{1}{\sigma_k}D_{\mathcal{Y}}(y, y^k) - \frac{1 + \gamma\sigma_k}{\sigma_k}D_{\mathcal{Y}}(y, y^{k+1}). \quad (23)$$

Let $r_k = (1 + \theta_k)\mathcal{L}(x^k, y) - \theta_k\mathcal{L}(x^{k-1}, y) - \mathcal{L}(x, y^{k+1})$. Equivalently equation (23) can be written as

$$r_k \leq \frac{1}{2\tau_k}\|x^k - x\|_2^2 - \frac{\tau_{k+1}}{\tau_k} \frac{1}{2\tau_{k+1}}\|x - x^{k+1}\|_2^2 + \frac{1}{\sigma_k}D_{\mathcal{Y}}(y, y^k) - \frac{\sigma_{k+1}}{\sigma_k}(1 + \gamma\sigma_k) \frac{1}{\sigma_{k+1}}D_{\mathcal{Y}}(y, y^{k+1}). \quad (24)$$

By construction of the algorithm, $\frac{\tau_{k+1}}{\tau_k} = \frac{\sigma_{k+1}}{\sigma_k}(1 + \gamma\sigma_k)$. Let $R_k := \frac{1}{2\tau_k}\|x - x^k\|_2^2 + \frac{1}{\sigma_k}D_{\mathcal{Y}}(y, y^k)$ then $\tau_k r_k \leq \tau_k R_k - \tau_{k+1} R_{k+1}$. Summing the inequalities for $k = 1, \dots, N$ gives

$$\sum_{k=1}^N \tau_k r_k \leq \tau_1 R_1 - \tau_N R_N \leq \tau_1 R_1 = \frac{1}{2}\|x - x^1\|_2^2 + \frac{1}{\beta_1}D_{\mathcal{Y}}(y, y^1) \quad (25)$$

On the other hand, we have

$$\begin{aligned} \sum_{k=1}^N \tau_k(1 + \theta_k)\mathcal{L}(x^k, y) - \tau_k\theta_k\mathcal{L}(x^{k-1}, y) &= \sum_{k=2}^N (\tau_{k-1}(1 + \theta_{k-1}) - \tau_k\theta_k)\mathcal{L}(x^{k-1}, y) \\ &\quad + \tau_N(1 + \theta_N)\mathcal{L}(x^N, y) - \tau_1\theta_1\mathcal{L}(x^0, y). \end{aligned}$$

Note that $\tau_{k-1}(1 + \theta_{k-1}) - \tau_k\theta_k \geq 0$ and, recalling $T_N = \sum_{k=1}^N \tau_k$:

$$\sum_{k=2}^N (\tau_{k-1}(1 + \theta_{k-1}) - \tau_k\theta_k) + \tau_N(1 + \theta_N) = \tau_1\theta_1 + T_N.$$

Hence, by convexity of $\mathcal{L}(\cdot, y)$, one has:

$$\begin{aligned} &\sum_{k=2}^N (\tau_{k-1}(1 + \theta_{k-1}) - \tau_k\theta_k)\mathcal{L}(x^{k-1}, y) + \tau_N(1 + \theta_N)\mathcal{L}(x^N, y) \\ &\geq (\tau_1\theta_1 + T_N)\mathcal{L}\left(\frac{\sum_{k=2}^N (\tau_{k-1}(1 + \theta_{k-1}) - \tau_k\theta_k)x^{k-1} + \tau_N(1 + \theta_N)x^N}{\tau_1\theta_1 + T_N}, y\right). \end{aligned}$$

Recalling $\bar{x}^k = (1 + \theta_k)x^k - \theta_k x^{k-1}$, one has

$$\sum_{k=2}^N (\tau_{k-1}(1 + \theta_{k-1}) - \tau_k\theta_k)x^{k-1} + \tau_N(1 + \theta_N)x^N = \tau_1\theta_1 x^0 + \sum_{k=1}^N \tau_k \bar{x}^k.$$

Hence we have shown:

$$\begin{aligned} \sum_{k=1}^N \tau_k (1 + \theta_k) \mathcal{L}(x^k, y) - \tau_k \theta_k \mathcal{L}(x^{k-1}, y) \geq \\ (\tau_1 \theta_1 + T_N) \mathcal{L} \left(\frac{\tau_1 \theta_1 x^0 + \sum_{k=1}^N \tau_k \bar{x}^k}{\tau_1 \theta_1 + T_N}, y \right) - \tau_1 \theta_1 \mathcal{L}(x^0, y), \end{aligned} \quad (26)$$

motivating the definition of $\hat{x}^N = \frac{\tau_1 \theta_1 x^0 + \sum_{k=1}^N \tau_k \bar{x}^k}{\tau_1 \theta_1 + T_N}$. Observe that if $\mathcal{L}(\cdot, y)$ is linear in its domain (as is the case for OT), then equation (26) is an equality and the RHS of (26) boils down to $T_N \mathcal{L} \left(\frac{1}{T_N} \sum_{k=1}^N \tau_k \bar{x}^k, y \right)$, motivating a simpler definition

$$\hat{x}^N = \frac{1}{T_N} \sum_{k=1}^N \tau_k \bar{x}^k. \quad (27)$$

Similarly, using the convexity of $-\mathcal{L}(x, \cdot)$ we have

$$\sum_{k=1}^N \tau_k \mathcal{L}(x, y^{k+1}) \leq T_N \mathcal{L} \left(x, \frac{1}{T_N} \sum_{k=1}^N \tau_k y^{k+1} \right), \quad (28)$$

motivating the definition $\hat{y}^N = \frac{1}{T_N} \sum_{k=1}^N \tau_k y^{k+1}$. In brief we have

$$T_N (\mathcal{L}(\hat{x}^N, y) - \mathcal{L}(x, \hat{y}^N)) \leq \frac{1}{2} \|x - x^1\|_2^2 + \frac{1}{\beta_1} D_{\mathcal{Y}}(y, y^1) - \tau_1 \theta_1 (\mathcal{L}(\hat{x}^N, y) - \mathcal{L}(x^0, y)). \quad (29)$$

Again, when $\mathcal{L}(\cdot, y)$ is linear in its domain, we define \hat{x}^N by (27) and obtain the simpler bound:

$$T_N (\mathcal{L}(\hat{x}^N, y) - \mathcal{L}(x, \hat{y}^N)) \leq \frac{1}{2} \|x - x^1\|_2^2 + \frac{1}{\beta_1} D_{\mathcal{Y}}(y, y^1). \quad (30)$$

Complexity estimate: the convex case. If $\gamma = 0$, then $\beta_k \equiv \beta$ is constant because of the update rule in step 2. Moreover, one checks that $\theta_k \leq \bar{\theta} := \frac{1+\sqrt{5}}{2}$ (see [35]), and $\tau_k \geq \frac{\rho}{\sqrt{\beta}L}$, provided that $\tau_0 \geq \frac{\rho}{\sqrt{\beta}L}$, which is the case when we choose $\tau_0 = \frac{1}{\sqrt{\beta}L}$. Hence, $T_N = \sum_{k=1}^N \tau_k \geq \frac{N\rho}{\sqrt{\beta}L} = N\rho\tau_0$. In particular, the bound (3) in Theorem 2.2 reduces to, letting $\sigma_0 = \beta\tau_0 = \sqrt{\beta}/L$:

$$\mathcal{G}_{x,y}(\hat{x}^N, \hat{y}^N) \leq \frac{1}{\rho N} \left(\frac{1}{2\tau_0} \|x - x^1\|_2^2 + \frac{1}{\sigma_0} D_{\mathcal{Y}}(y, y^1) - \tau_1 \theta_1 (\mathcal{L}(\hat{x}^N, y) - \mathcal{L}(x^0, y)) \right). \quad (31)$$

The strongly convex case. Consider now the case $\gamma > 0$. A first point is to show that $\tau_k \geq \frac{\rho}{\sqrt{\beta_k}L}$: this is guaranteed only if the linesearch procedure is active, but since β_k is decreased at each iteration one could fall into a situation where this does not happen.

Let us assume that for some $k \geq 1$, $\tau_k \geq \frac{\rho}{\sqrt{\beta_k L}}$, $\theta_k = \tau_k/\tau_{k-1}$, $\beta_k = \beta_{k-1}/(1 + \gamma\tau_{k-1}\beta_{k-1})$, and consider the next step. Then either the linesearch terminated after more than one iteration and one must have $\tau_{k+1} \geq \frac{\rho}{\sqrt{\beta_{k+1}L}}$, else one has $\tau_{k+1} = \tau_k\sqrt{1 + \theta_k}$. In that case,

$$\begin{aligned} \tau_{k+1} = \tau_k\sqrt{1 + \theta_k} &\geq \frac{\rho\sqrt{1 + \theta_k}}{\sqrt{\beta_k L}} \geq \frac{\rho}{\sqrt{\beta_{k+1}L}} \Leftrightarrow \\ (1 + \theta_k)\frac{\beta_k}{1 + \gamma\tau_k\beta_k} &\geq \beta_k \Leftrightarrow \theta_k = \frac{\tau_k}{\tau_{k-1}} \geq \gamma\tau_k\beta_k \Leftrightarrow 1 \geq \gamma\tau_{k-1}\beta_k = \frac{\gamma\tau_{k-1}\beta_{k-1}}{1 + \gamma\tau_{k-1}\beta_{k-1}} \end{aligned}$$

which is true. Hence, one will have that $\tau_k \geq \rho/(\sqrt{\beta_k L})$ for all $k \geq 1$ provided $\tau_1 \geq \rho/\sqrt{\beta_1 L}$, (θ_1 , β_1 being defined from τ_1 , τ_0 , β_0 as in the algorithm). A choice which guarantees that the induction holds is as follows:

$$\beta_0 > 0, \quad \tau_0 = \frac{1}{\sqrt{\beta_0 L}}, \quad \theta_0 = \frac{\gamma\sqrt{\beta_0}}{L}. \quad (32)$$

Then, one has $\beta_1 = \beta_0/(1 + \gamma\tau_0\beta_0) = \beta_0/(1 + \gamma\sqrt{\beta_0}/L) = \beta_0/(1 + \theta_0)$ so that, at the beginning of the linesearch for the first step, $\tau_1 = \tau_0\sqrt{1 + \theta_0} = \sqrt{(1 + \theta_0)/\beta_0}/L = 1/(\sqrt{\beta_1}L)$ and in particular the first linesearch will terminate immediately. The induction then guarantees that for all $k \geq 2$, $\tau_k \geq \rho/(\sqrt{\beta_k L})$.

Hence, one has for all $k \geq 1$:

$$\beta_k = \frac{\beta_{k-1}}{1 + \gamma\beta_{k-1}\tau_{k-1}} \leq \frac{\beta_{k-1}}{1 + \frac{\gamma\rho}{L}\sqrt{\beta_{k-1}}}.$$

One can write, for any $\alpha \in [0, 1]$:

$$\frac{1}{\beta_k} \geq \frac{1}{\beta_{k-1}} + \frac{\gamma\rho}{L} \frac{1}{\sqrt{\beta_{k-1}}} = \left(\frac{1}{\sqrt{\beta_{k-1}}} + \frac{\alpha\gamma\rho}{2L} \right)^2 + \frac{\gamma\rho}{L} \left(\frac{1 - \alpha}{\sqrt{\beta_{k-1}}} - \frac{\alpha^2\gamma\rho}{4L} \right).$$

Choosing to simplify $\alpha = .5$ one deduces that $\beta_k \leq 1/(1/\sqrt{\beta_l} + \gamma\rho(k-l)/(4L))^2$ provided $\beta_l \leq 64L^2/(\gamma\rho)^2$. Notice, though, that one also has, by induction (assuming $L/(\gamma\rho) \geq 1$):

$$\beta_k \leq \frac{L}{\gamma\rho} \sqrt{\beta_{k-1}} \leq \frac{L^2}{\gamma^2\rho^2} \beta_0^{\frac{1}{2^k}}$$

hence $\beta_l \leq 64L^2/(\gamma\rho)^2$ as soon as $l \geq \log_2 \log_2 \beta_0 - 2$, which is very little whatever the choice of β_0 ($l < 8$ if $\beta_0 \leq 10^{300}$). So for “reasonable” choices of β_0 , and k large enough, one has $\beta_k \leq 16L^2/(\gamma\rho)^2(1/k)^2$, $\tau_k \geq \gamma\rho^2(4L^2)k$, and $T_N \geq \gamma\rho^2/(16L^2)N^2$.

A.2. Termination of the linesearch. In this section we prove that condition in step 8 of Algorithm 1 is always reached, and the number of iterations of the inner-loop does not affect the convergence rate.

If $\tau_k \leq \frac{1}{\sqrt{\beta_k L}}$ we have, recalling that $L = \|K\|_{2 \rightarrow \mathcal{Y}^*} = \|K^*\|_{\mathcal{Y} \rightarrow 2} = \sup_{\|x\|_2 \leq 1, \|y\|_{\mathcal{Y}} \leq 1} \langle x, K^* y \rangle$:

$$\begin{aligned} \tau_k \langle \bar{x}^k - x^{k+1}, K^* y^{k+1} - K^* y^k \rangle &\leq \tau_k \|x^{k+1} - \bar{x}^k\|_2 \|K^* y^{k+1} - K^* y^k\|_2 \\ &\leq \frac{1}{2} \|x^{k+1} - \bar{x}^k\|_2^2 + \frac{\tau_k^2}{2} L^2 \|y^{k+1} - y^k\|_{\mathcal{Y}}^2 \\ &\leq \frac{1}{2} \|x^{k+1} - \bar{x}^k\|_2^2 + \frac{1}{2\beta_k} \|y^{k+1} - y^k\|_{\mathcal{Y}}^2 \\ &\leq \frac{1}{2} \|x^{k+1} - \bar{x}^k\|_2^2 + \frac{1}{\beta_k} D(y^{k+1}, y^k) \end{aligned}$$

Therefore, the condition is reached as soon as $\tau_k \leq \frac{1}{\sqrt{\beta_k L}}$ which eventually will occur since τ_k decrease by a factor $\rho < 1$ at each iteration of the inner-loop. We estimate, following [35], the extra computing time due to the inner loops, up to the k th outer loop.

By construction of the algorithm $\theta_k \leq \sqrt{1 + \theta_{k-1}}$, hence by induction we have that $\theta_k \leq \bar{\theta} := \max \left\{ \frac{1+\sqrt{5}}{2}, \theta_0 \right\}$. Fix iteration k . After i_k iterations of the inner-loop we have $\tau_k(i_k) = \tau_{k-1} \sqrt{1 + \theta_{k-1} \rho^{i_k}}$ and $\theta_k(i_k) = \sqrt{1 + \theta_{k-1} \rho^{i_k}} \leq \bar{\theta} \rho^{i_k}$. Let τ_0 such that the condition in step 8 is true (for instance $\tau_0 = \frac{1}{\sqrt{\beta_0 L}}$), then:

$$\bar{\theta} \rho^{i_k} \geq \theta_k \geq \frac{\tau_0 \rho}{\tau_{k-1}} = \frac{\rho}{\prod_{j=1}^{k-1} \theta_j} \geq \frac{\rho^{k+1}}{\bar{\theta}^k \rho^{\sum_{j=1}^{k-1} i_j}},$$

so that

$$\sum_{j=1}^k i_j \leq (k+1) \left(1 + \frac{\ln \bar{\theta}}{|\ln \rho|} \right). \quad (33)$$

A.3. Proof of Lemma 3.1. We assume as in the statement that $\min_i C_{i,j} = 0$ for all $j = 1, \dots, n$, and $\min_j C_{i,j} = 0$ for all $i = 0, \dots, n$. (Otherwise we remove the minimum on each row and then on each column of C without changing the problem, up to a constant.)

Let (u, v) a dual solution of the saddle point problem. Define $\tilde{u} = u - \mathbf{1}_n \min_i u_i$ and $\tilde{v} = v + \mathbf{1}_n \min_i u_i$. The new dual solution has the same value, $\tilde{u} \geq 0$ and there exist i_0 such that $\tilde{u}_{i_0} = 0$. The dual constraints are $\tilde{u}_i + \tilde{v}_j \leq C_{i,j}$ for all i, j . By optimality of (\tilde{u}, \tilde{v}) we have $\tilde{u}_i = \min_j (C_{i,j} - \tilde{v}_j)$ for each i , and in particular there exist j_0 such that $\tilde{u}_{i_0} + \tilde{v}_{j_0} = \tilde{v}_{j_0} = C_{i_0, j_0}$. Symmetrically, we have $\tilde{v}_j = \min_i (C_{i,j} - \tilde{u}_i) \leq \min_i C_{i,j} = 0$ since $\tilde{u}_i \geq 0$. In particular $\tilde{v}_{j_0} = C_{i_0, j_0} = 0$.

As a consequence, for each i , $\tilde{u}_i \leq C_{i, j_0} - \tilde{v}_{j_0} \leq \|C\|$, and for each j , $\exists i$ with $\tilde{v}_j = C_{i,j} - \tilde{u}_i \geq C_{i,j} - \|C\| \geq -\|C\|$. In brief, we have $0 \leq \tilde{u}_i \leq \|C\|$ and $-\|C\| \leq \tilde{v}_j \leq 0$. Defining $u^* = \tilde{u} - \mathbf{1}_n \frac{\|C\|}{2}$ and $v^* = \tilde{v} + \mathbf{1}_n \frac{\|C\|}{2}$ we have the first result.

For the second part, take $\mu = (1, 0, \dots, 0)^\top$ and $\nu = (0, \dots, 0, 1)^\top$, and $C = \mu \otimes \nu$. The unique optimal transport in this case is $X_{1,n} = 1$ and $X_{ij} = 0$ otherwise. Therefore, for any dual solution (u^*, v^*) we have $C_{1,n} = 1 = u_1^* + v_n^*$, hence $\|(u^*, v^*)\| \geq 1/2 = \frac{\|C\|}{2}$.

A.4. Proof of Lemma 3.2. We show that Algorithm 2, a.k.a. “ROUND”, in [2] satisfies a better error bound than proved in the original paper, at least if its input is already in the matrix unit simplex. We consider $X \in \Delta$: hence one has $\sum_{i,j} X_{i,j} = 1$, and $\sum_i (X\mathbf{1}_n - \mu)_i = 0$, $\sum_j (X^\top \mathbf{1}_n - \nu)_j = 0$.

This algorithm outputs $Y = \mathcal{A}X$ as follows: First, one lets:

$$X'_{i,j} = \min\left\{1, \frac{\mu_i}{(X\mathbf{1}_n)_i}\right\} X_{i,j} = \begin{cases} X_{i,j} & \text{if } \mu_i \geq (X\mathbf{1}_n)_i \\ \frac{\mu_i}{(X\mathbf{1}_n)_i} X_{i,j} & \text{else} \end{cases}$$

and in particular $X' \leq X$ and $X'\mathbf{1}_n \leq \mu$. Then:

$$X''_{i,j} = X'_{i,j} \min\left\{1, \frac{\nu_j}{((X')^\top \mathbf{1}_n)_j}\right\} = \begin{cases} X'_{i,j} & \text{if } \nu_j \geq ((X')^\top \mathbf{1}_n)_j \\ \frac{\nu_j}{((X')^\top \mathbf{1}_n)_j} X'_{i,j} & \text{else} \end{cases}$$

which is such that $X'' \leq X'$ and $X''^\top \mathbf{1}_n \leq \nu$. In particular since $\mu \geq X'\mathbf{1}_n \geq X''\mathbf{1}_n$, $\|\mu - X''\mathbf{1}_n\|_1 = \sum_i (\mu_i - \sum_j X''_{i,j}) = 1 - \sum_{i,j} X''_{i,j} = \sum_j (\nu_j - \sum_i X''_{i,j})$. Eventually the output is given by:

$$Y_{i,j} = X''_{i,j} + \frac{\mu_i - (X''\mathbf{1}_n)_i}{\|\mu - X''\mathbf{1}_n\|_1} (\nu_j - ((X'')^\top \mathbf{1}_n)_j).$$

and we see that $Y\mathbf{1}_n = \mu$, $Y^\top \mathbf{1}_n = \nu$. One has, then:

$$Y - X = Y - X'' - (X' - X'') - (X - X').$$

Let $J^+ = \{j : \nu_j \geq ((X')^\top \mathbf{1}_n)_j\}$ and $J^- = \{1, \dots, N\} \setminus J^+$. Then if $j \in J^+$, $X''_{i,j} = X'_{i,j}$ and

$$\begin{aligned} Y_{i,j} - X_{i,j} &= \frac{\mu_i - (X''\mathbf{1}_n)_i}{\|\mu - X''\mathbf{1}_n\|_1} (\nu_j - ((X'')^\top \mathbf{1}_n)_j) - (X_{i,j} - X'_{i,j}) \\ &= \frac{\mu_i - (X''\mathbf{1}_n)_i}{\|\mu - X''\mathbf{1}_n\|_1} (\nu_j - ((X')^\top \mathbf{1}_n)_j) - (X_{i,j} - X'_{i,j}) \\ &= \frac{\mu_i - (X''\mathbf{1}_n)_i}{\|\mu - X''\mathbf{1}_n\|_1} (\nu_j - (X^\top \mathbf{1}_n)_j) + \frac{\mu_i - (X''\mathbf{1}_n)_i}{\|\mu - X''\mathbf{1}_n\|_1} ((X - X')^\top \mathbf{1}_n)_j - (X_{i,j} - X'_{i,j}) \end{aligned}$$

and

$$\sum_i |Y_{i,j} - X_{i,j}| \leq |\nu_j - (X^\top \mathbf{1}_n)_j| + 2((X - X')^\top \mathbf{1}_n)_j.$$

If $j \in J^-$, then $((X'')^\top \mathbf{1}_n)_j = \nu_j$ hence $Y_{i,j} = X''_{i,j}$, so that:

$$\begin{aligned} Y_{i,j} - X_{i,j} &= -(X'_{i,j} - X''_{i,j}) - (X_{i,j} - X'_{i,j}) \\ &= -\frac{((X')^\top \mathbf{1}_n)_j - \nu_j}{((X')^\top \mathbf{1}_n)_j} X'_{i,j} - (X_{i,j} - X'_{i,j}) \\ &= -\frac{X'_{i,j}}{((X')^\top \mathbf{1}_n)_j} ((X^\top \mathbf{1}_n)_j - \nu_j) - \frac{X'_{i,j}}{((X')^\top \mathbf{1}_n)_j} ((X' - X)^\top \mathbf{1}_n)_j - (X_{i,j} - X'_{i,j}). \end{aligned}$$

and again

$$\sum_i |Y_{i,j} - X_{i,j}| \leq (X^\top \mathbf{1}_n)_j - \nu_j + 2((X - X')^\top \mathbf{1}_n)_j.$$

Hence,

$$\|Y - X\|_1 \leq \|X^\top \mathbf{1}_n - \nu\|_1 + 2\|X - X'\|_1.$$

Eventually, one has:

$$\|X - X'\|_1 = \frac{1}{2} \|X\mathbf{1}_n - \mu\|_1$$

since for all i, j ,

$$X_{i,j} - X'_{i,j} = \frac{((X\mathbf{1}_n)_i - \mu_i)^+}{(X\mathbf{1}_n)_i} X_{i,j}$$

and $\sum_i (X\mathbf{1}_n)_i - \mu_i = 0 \Rightarrow \sum_i ((X\mathbf{1}_n)_i - \mu_i)^+ = \sum_i ((X\mathbf{1}_n)_i - \mu_i)^- = \frac{1}{2} \sum_i |(X\mathbf{1}_n)_i - \mu_i|$.
It follows that

$$\|Y - X\|_1 \leq \|X^\top \mathbf{1}_n - \nu\|_1 + \|X\mathbf{1}_n - \mu\|_1. \quad (34)$$

(A.C.) CEREMADE, CNRS & UNIVERSITÉ PARIS-DAUPHINE, PARIS, FRANCE

Email address: chambolle@ceremade.dauphine.fr

(J.P.C.) UNIVERSIDAD ADOLFO IBÁÑEZ, FACULTAD DE INGENIERÍA Y CIENCIAS, SANTIAGO, CHILE, DEPARTAMENTO DE INGENIERÍA INDUSTRIAL, UNIVERSIDAD CATÓLICA DEL NORTE, ANTOFAGASTA, CHILE

Email address: juan.contrerasff@gmail.com