

Efficient and Accurate Optimal Transport with Mirror Descent and Conjugate Gradients

Anonymous authors

Paper under double-blind review

Abstract

We propose *Mirror Descent Optimal Transport* (MDOT), a novel method for solving discrete optimal transport (OT) problems with high precision, by unifying temperature annealing in entropic-regularized OT (EOT) with mirror descent techniques. In this framework, temperature annealing produces a sequence of EOT dual problems, whose solution gradually gets closer to the solution of the original OT problem. We solve each problem efficiently using a GPU-parallel nonlinear conjugate gradients algorithm (PNCG) that outperforms traditional Sinkhorn iterations under weak regularization. Moreover, our investigation also reveals that the theoretical convergence rate of Sinkhorn iterations can exceed existing non-asymptotic bounds when its stopping criterion is tuned in a manner analogous to MDOT.

Our comprehensive ablation studies of MDOT-PNCG affirm its robustness across a wide range of algorithmic parameters. Benchmarking on 24 problem sets of size $n = 4096$ in a GPU environment demonstrate that our method attains high-precision, feasible solutions significantly faster than a representative set of existing OT solvers—including accelerated gradient methods and advanced Sinkhorn variants—in both wall-clock time and number of operations. Empirical convergence rates range between $O(n^2\varepsilon^{-1/4})$ and $O(n^2\varepsilon^{-1})$, where ε is the optimality gap. For problem sizes up to $n = 16384$, the empirical runtime scales as $\tilde{O}(n^2)$ for moderate precision and as $\tilde{O}(n^{5/2})$ at worst for high precision. These findings establish MDOT-PNCG as a compelling alternative to current OT solvers, particularly in challenging weak-regularization regimes.

1 INTRODUCTION

When a statistical distance is required for an event space equipped with a metric, optimal transport (OT) distances, such as the Wasserstein metric, provide an intuitive means to account for the inherent structure of the metric space. Consequently, fast, scalable, and accurate computation of OT distances is a major problem encountered in various scientific fields. Example application areas include point cloud registration (Shen et al., 2021), color transfer (Pitie et al., 2005; Ferradans et al., 2014; Rabin et al., 2014), shape matching (Feydy et al., 2017), texture mixing (Ferradans et al., 2013; Bonneel et al., 2015) and meshing (Digne et al., 2014) in computer vision and graphics, quantum mechanics (Léonard, 2012), astronomy (Frisch et al., 2002; Levy et al., 2021) and quantum chemistry (Bokanowski & Grébert, 1996) in physics, and generative modeling (Gulrajani et al., 2017; Genevay et al., 2018), reinforcement learning (Ferns et al., 2004; Dadashi et al., 2021), and neural architecture search (Kandasamy et al., 2018) in machine learning. Exact solvers for the discrete OT problem encounter significant computational hurdles in high dimensions, with theoretical complexity $\tilde{O}(n^{5/2})$ and practical complexity $\tilde{O}(n^3)$ (Lee & Sidford, 2014; Pele & Werman, 2009).

Entropic regularization, as pioneered by Cuturi (2013), has mitigated challenges in scalability by regularizing the classical problem, thereby allowing approximate solutions in $\tilde{O}(n^2)$ time via the Sinkhorn-Knopp (SK) matrix scaling algorithm. This advancement, together with GPU parallelization, has yielded substantial speed improvements, making it several orders of magnitude faster than conventional CPU-based solvers (e.g., linear programming) in high dimensions (Peyré et al., 2019). However, these methods necessitate a delicate balance between regularization strength and convergence speed, a trade-off that can compromise the precision of the solution. Despite significant progress in recent years, many state-of-the-art solvers still struggle to strike a better trade-off than aggressively tuned Sinkhorn iterations in practice (Dvurechensky

et al., 2018; Jambulapati et al., 2019; Lin et al., 2019). Although they offer superior theoretical guarantees, their practical performance is often less compelling, particularly in terms of speed and scalability. Existing algorithms either suffer from high computational complexity or do not take advantage of modern hardware capabilities, such as GPU parallelization (Tang et al., 2024). To understand and combat these challenges, we make the following contributions:

1. We empirically show that in a GPU environment the decades-old Sinkhorn-Knopp algorithm for OT can still outperform many theoretically grounded recent OT algorithms in practice, especially when tuned with a seemingly unconventional stopping criterion formula proposed here (Fig. 4).
2. We introduce mirror descent optimal transport (MDOT), a method which generalizes temperature annealing in entropic OT (EOT) (Schmitzer, 2019; Feydy, 2020), and connects temperature annealing to mirror descent (Alg. 1).
3. We introduce an instantiation of MDOT that empirically improves speed and robustness to temperature (regularization strength) decay rate compared to ε -scaling of Schmitzer (2019) (Fig. 1).
4. We show that MDOT can compute high precision, feasible solutions and its performance can be boosted by adopting a specialized GPU-parallel conjugate gradients (CG) algorithm developed here (Alg. 2); this method is highly competitive in practice, as we show empirically (Figs. 4, 5, 8-17).

The remainder of this paper is organized as follows. In the next section, we introduce our notation and the necessary background, followed by related work in Sec. 3. In Sec. 4.1-4.2, we introduce the MDOT framework and establish its connection to temperature annealing strategies, and make some practical recommendations. In Sec. 4.3, we introduce the non-linear CG algorithm to be used within MDOT as an alternative to SK. In Sec. 5, we benchmark various algorithms on upsampled MNIST ($n = 4096$) under L_1 and squared L_2 costs, and a color transfer problem *in terms of wall-clock time*, and further study the operation count dependence of the proposed algorithm on problem size n . Lastly, we present concluding remarks in Sec. 6.

2 Background

Here, we present our notation, the basics of EOT, and the necessary background on mirror descent and CG.

Notation and Definitions. We consider discrete OT, where the event space is finite with n particles and $\Delta_n \subset \mathbb{R}_{\geq 0}^n$ is the $(n-1)$ -simplex. The row sum of an $n \times n$ matrix P is $\mathbf{r}(P) := P\mathbf{1}$ and the column sum is $\mathbf{c}(P) := P^\top \mathbf{1}$. Given marginals $\mathbf{r}, \mathbf{c} \in \Delta_n$, the transportation polytope is written as $\mathcal{U}(\mathbf{r}, \mathbf{c}) = \{P \in \mathbb{R}_{\geq 0}^{n \times n} \mid \mathbf{r}(P) = \mathbf{r}, \mathbf{c}(P) = \mathbf{c}\}$. Division, exp and log over vectors or matrices are element-wise. Vectors in \mathbb{R}^n are column vectors, and (\mathbf{x}, \mathbf{y}) denotes the concatenation of \mathbf{x} and \mathbf{y} . Vector and Frobenius inner products alike are given by $\langle \cdot, \cdot \rangle$. An $n \times n$ diagonal matrix with $\mathbf{x} \in \mathbb{R}^n$ along the diagonal is written as $\mathbf{D}(\mathbf{x})$, and the vector formed by the diagonal entries of a matrix Q is $\mathbf{diag}(Q)$. LogSumExp reductions over the rows and columns of $X \in \mathbb{R}^{n \times n}$ are given by $\text{LSE}_r(X) := \log(\exp\{X\}\mathbf{1})$ and $\text{LSE}_c(X) := \log(\exp\{X^\top\}\mathbf{1})$. The Shannon entropy of $\mathbf{r} \in \Delta_n$ is denoted $H(\mathbf{r}) = -\langle \mathbf{r}, \log \mathbf{r} \rangle$ with the convention that $0 \cdot \log 0 = 0$. Under the same convention, the KL divergence $D_{\text{KL}}(\mathbf{r} \mid \mathbf{r}') = \langle \mathbf{r}, \log(\mathbf{r}/\mathbf{r}') \rangle + \langle \mathbf{r}' - \mathbf{r}, \mathbf{1} \rangle$ for $\mathbf{r}, \mathbf{r}' \in \mathbb{R}_{\geq 0}^n$ given \mathbf{r} absolutely continuous with respect to \mathbf{r}' .

2.1 Optimal Transport

Given a cost matrix $C \in [0, 1]^{n \times n}$, where C_{ij} is the transportation cost between the i^{th} and j^{th} particles, we study the EOT problem given by:

$$\underset{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})}{\text{minimize}} \quad \langle P, C \rangle - \frac{1}{\gamma} H(P), \quad (1)$$

where $\gamma > 0$. Here, the regularization weight γ^{-1} is called *temperature*. The Lagrangian of (1) is strictly convex in P , which renders the solution $P^*(\gamma)$ unique. $P^*(\gamma)$ converges to a solution of the unregularized OT problem as $\gamma \rightarrow \infty$ and admits the following form (Cuturi, 2013):

$$P_{ij}(\mathbf{u}, \mathbf{v}; \gamma) = \exp\{u_i + v_j - \gamma C_{ij}\}, \quad (2)$$

where $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$. An optimal pair (\mathbf{u}, \mathbf{v}) minimizes the following convex dual problem (Lin et al., 2019):

$$\underset{\mathbf{u}, \mathbf{v} \in \mathbb{R}^n}{\text{minimize}} \quad g(\mathbf{u}, \mathbf{v}; \gamma) = \mathbf{1}^\top P(\mathbf{u}, \mathbf{v}; \gamma) \mathbf{1} - \langle \mathbf{u}, \mathbf{r} \rangle - \langle \mathbf{v}, \mathbf{c} \rangle, \quad (3)$$

where $\nabla_{\mathbf{u}} g = \mathbf{r}(P) - \mathbf{r}$ and $\nabla_{\mathbf{v}} g = \mathbf{c}(P) - \mathbf{c}$. Given some initial (\mathbf{u}, \mathbf{v}) , solving (3) amounts to a KL projection of $P(\mathbf{u}, \mathbf{v}; \gamma)$ onto $\mathcal{U}(\mathbf{r}, \mathbf{c})$. That is, $P(\mathbf{u}^*, \mathbf{v}^*; \gamma) = \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} D_{\text{KL}}(P | P(\mathbf{u}, \mathbf{v}; \gamma))$; see derivation in Appx. A.1.

The Sinkhorn-Knopp (SK) algorithm (see Alg. 4 in Appx. A) can be used to solve (3) and carry out an approximate projection of $P(\mathbf{u}, \mathbf{v}; \gamma)$ onto $\mathcal{U}(\mathbf{r}, \mathbf{c})$. Sinkhorn updates guarantee convergence to dual-optimal variables as the number of iterations $k \rightarrow \infty$ (Sinkhorn & Knopp, 1967; Sinkhorn, 1967; Franklin & Lorenz, 1989; Knight, 2008). Dvurechensky et al. (2018) showed that SK can be used to compute a solution $P \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ satisfying $\langle P - P^*, C \rangle \leq \varepsilon$ with complexity $\tilde{O}(n^2/\varepsilon^2)$, where P^* is an optimal solution of the unregularized OT problem. In particular, one first minimizes the dual objective until the L_1 -norm of its gradient is below a prescribed threshold, then applies the rounding algorithm of Altschuler et al. (2017) on the infeasible plan given by (2) to obtain $P \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ with an upper bound on the primal cost increase.

2.2 Mirror Descent

Rather than the entropic regularized problem (1), we will approximately solve a sequence of *relative-entropy* regularized problems as in mirror descent, where each problem has the form (Nemirovski & Yudin, 1983):

$$P^{(t+1)} = \arg \min_{P \in \mathcal{F} \cap \mathcal{D}} \{ \langle \nabla_P f(P^{(t)}), P \rangle + \frac{1}{\Delta^{(t)}} D_h(P | P^{(t)}) \}. \quad (4)$$

Here $f : \mathcal{D} \rightarrow \mathbb{R}$ is the objective, \mathcal{F} is a given feasible set, $\Delta^{(t)} > 0$ is the step size and the Bregman divergence

$$D_h(Q | P) = h(Q) - h(P) - \langle \nabla h(P), Q - P \rangle \quad (5)$$

given a strictly convex and differentiable function $h : \mathcal{D} \rightarrow \mathbb{R}$ called the *mirror map*. Equivalent to (4) is:

$$\hat{P}^{(t+1)} = \nabla h^{-1} \left(\nabla h(P^{(t)}) - \Delta^{(t)} \nabla f(P^{(t)}) \right) \quad (6)$$

$$P^{(t+1)} = \arg \min_{P \in \mathcal{F} \cap \mathcal{D}} D_h(P | \hat{P}^{(t+1)}). \quad (7)$$

Here, (6) takes a gradient step in the dual space and maps the new point back onto the primal space via ∇h^{-1} , while (7) defines a *Bregman projection* of $\hat{P}^{(t+1)}$ onto the feasible set \mathcal{F} in the primal space.

Throughout, we only use the negative entropy mirror map $h(P) = -H(P)$ in domain $\mathcal{D} = \mathbb{R}_{\geq 0}^{n \times n}$, which yields $D_h(Q | P) = D_{\text{KL}}(Q | P)$. Di Marino & Gerolin (2020) provide a treatment of more general divergence-regularized optimal transport problems (using mirror maps besides negative Shannon entropy), albeit running only a single step of (4) with $P^{(0)} \in \mathbb{R}_{\geq 0}^{n \times n}$ set to the independence coupling $\mathbf{r} \mathbf{c}^\top$, $\mathcal{F} = \mathcal{U}(\mathbf{r}, \mathbf{c})$ and objective $f(P) = \langle P, C \rangle$; we consider multiple steps.

3 Related Work

Acceleration of approximate OT solvers has been a focus of machine learning research since the seminal work of Cuturi (2013). For instance, Altschuler et al. (2017) proposed the Greenkhorn algorithm, which greedily selects individual rows or columns to scale at a given step and requires fewer row/column updates than SK to converge, but performs poorly due to low GPU utilization unless n is extremely large. Dvurechensky et al. (2018) proposed an adaptive primal-dual accelerated gradient descent (APDAGD) algorithm. Lin et al. (2019) later proposed adaptive primal-dual accelerated *mirror* descent (APDAMD) with theoretical guarantees. Lin et al. (2019) showed APDAMD to outperform APDAGD *in terms of number of iterations*, but not SK. Further, these tests only covered a high relative error regime ($>50\%$); we investigate a broader scope down to 10^{-9} error in Section 5. Modest gains over SK in terms of number of iterations in the same regime were later obtained by Lin et al. (2022) via an accelerated alternating minimization (AAM) algorithm similar to that of Guminov et al. (2021). Notably, APDAMD applies mirror descent to the dual (3) of the EOT problem, while we apply it to the primal of the unregularized OT problem, i.e., problem (1) as $\gamma \rightarrow \infty$.

Application of mirror descent to the primal of the OT problem has also been considered. [Yang & Toh \(2022\)](#) discuss an algorithm similar to MDOT, although their approach differs from ours in a number of ways. They require a rounding procedure onto $\mathcal{U}(\mathbf{r}, \mathbf{c})$ after *each* mirror descent step, and verify complicated stopping criteria for all Bregman projections; we did not find these to be necessary experimentally or mathematically. Recently, [Ballu & Berthet \(2023\)](#) introduced Mirror Sinkhorn (MSK), which also takes gradient steps in the dual space as in (6), but instead of approximately projecting onto the feasible set $\mathcal{U}(\mathbf{r}, \mathbf{c})$ as in (7) (as we do here), they alternately project onto $\mathcal{U}(\cdot, \mathbf{c})$ and $\mathcal{U}(\mathbf{r}, \cdot)$ via Sinkhorn updates, satisfying only half of the equality constraints at a time. Our experiments in Sec. 5 suggest that this approach is efficient only in the low precision regime. Furthermore, MSK requires maintaining a running average of the transport plan at each iteration, precluding a straightforward $O(n)$ memory implementation. One could, in principle, store all past dual variables and recompute each plan to form the average at each step, but this would require $O(nT)$ memory and $O(n^2T)$ recomputation time at iteration T , which is typically impractical. In contrast, all algorithms presented here admit $O(n)$ memory implementations, assuming individual cost matrix entries can be computed on-the-fly in $O(1)$ time. [Xie et al. \(2020\)](#) previously proposed an algorithm (IPOT) similar to MSK with a fixed, even number of Sinkhorn updates (usually 2) following temperature updates; we omitted additional empirical comparison to IPOT given its similarity with MSK. Alg. 3.5 of [Feydy \(2020\)](#) is also similar to these algorithms in spirit and is discussed thoroughly in Sec. 5.2. As discussed in detail in Sec. 4.1, well-known ε -scaling strategies are also closely related ([Kosowsky & Yuille, 1994](#); [Schmitzer, 2019](#)).

An alternative line of acceleration research focuses on multi-scale strategies, which employ clustering or grid-based methods to solve a series of coarse-to-fine OT problems and are sometimes combined with ε -scaling ([Schmitzer, 2016](#); [2019](#); [Feydy, 2020](#)). These are known to provide performance gains when the marginals are defined over well-clustered particles or in low-dimensional event spaces ([Peyré et al., 2019](#)). Lastly, in a similar spirit to our use of non-linear CG here, curvature-aware convex optimization techniques such as L-BFGS have also been considered for OT, e.g., [Mérigot \(2011\)](#); [Blondel et al. \(2018\)](#); however, scalability, precision and better performance than SK on GPUs has not been demonstrated simultaneously to our knowledge. [Tang et al. \(2024\)](#) recently adopted Newton’s method with Hessian sparsification to efficiently use second order information, but their key sparsification strategy is maximally utilized only on CPUs.

4 A Mirror Descent Framework for Optimal Transport

4.1 Temperature Annealing as Mirror Descent

The OT objective $\langle P, C \rangle$ has a constant gradient $\nabla_P \langle P, C \rangle = C$. Given step sizes $\Delta_\gamma^{(t)} > 0$ at time $t \geq 0$, mirror descent iterates with the negative entropy mirror map $h(P) = -H(P)$ are thus given by (cf. 4):

$$P^{(t+1)} = \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} \left\{ \langle P, C \rangle + \frac{1}{\Delta_\gamma^{(t)}} D_{\text{KL}}(P | P^{(t)}) \right\}. \quad (8)$$

A useful convention here is to take $P^{(0)} = \mathbf{r}\mathbf{c}^\top$, the maximum entropy transport plan in $\mathcal{U}(\mathbf{r}, \mathbf{c})$. However, as the second item of the next proposition, which describes the properties of $P^{(t)}$ computed via (8), suggests, this initialization is not strictly necessary and any rank-1 matrix with positive entries may be selected.

Proposition 4.1. *Suppose $P^{(0)} \in \mathbb{R}_{>0}^{n \times n}$ is rank-1 and $P^{(t)}$ are computed via (8) for $t \geq 1$. Let $\gamma^{(0)} = 0$ and $\gamma^{(t+1)} = \gamma^{(t)} + \Delta_\gamma^{(t)}$, which together imply $\gamma^{(t+1)} = \sum_{t'=0}^t \Delta_\gamma^{(t')}$. Given $P^* \in \mathcal{U}(\mathbf{r}, \mathbf{c})$, a minimizer of $\langle P, C \rangle$, and $H_{\min}(\mathbf{r}, \mathbf{c}) := \min(H(\mathbf{r}), H(\mathbf{c}))$, the following are true:*

1. $P^{(t)} = P^*(\gamma^{(t)})$, where $P^*(\gamma^{(t)})$ is the unique solution of (1) for $\gamma = \gamma^{(t)}$.
2. Let $\tilde{P}^{(t)} \in \mathbb{R}_{>0}^{n \times n}$ be any matrix of the form $\tilde{P}_{ij}^{(t)} = \exp\{u_i + v_j - \gamma^{(t)} C_{ij}\}$ for some $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$. In addition to (8), we also have $P^{(t+1)} = \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} \left\{ \langle P, C \rangle + \frac{1}{\Delta_\gamma^{(t)}} D_{\text{KL}}(P | \tilde{P}^{(t)}) \right\}$.
3. $\langle P^{(t)} - P^*, C \rangle \leq H_{\min}(\mathbf{r}, \mathbf{c}) / \gamma^{(t)}$.
4. $\langle P^{(t)} - P^{(t+1)}, C \rangle = \frac{1}{\Delta_\gamma^{(t)}} \left(D_{\text{KL}}(P^{(t)} | P^{(t+1)}) + D_{\text{KL}}(P^{(t+1)} | P^{(t)}) \right)$ for all $t \geq 0$.

The equivalence of $P^{(t)}$ to the solution of (1) can be seen as follows. By setting the partial derivatives (with respect to P) of the Lagrangian of (8) to zero, we find that $P^{(t+1)}$ must have the form

$$P_{ij}^{(t+1)} = P_{ij}^{(t)} \exp\{u_i + v_j - \Delta_\gamma^{(t)} C_{ij}\}. \quad (9)$$

These relations allow us to unroll the recursion to get $P^{(t+1)}$ as follows (using the fact that $P^{(0)}$ is rank-1):

$$P_{ij}^{(t+1)} = \exp\{u_i^*(\gamma^{(t)}) + u_i + v_j^*(\gamma^{(t)}) + v_j - \gamma^{(t+1)} C_{ij}\}. \quad (10)$$

Thus, P in (9-10) has the same form as (2). That is, the mirror descent procedure in (8) mapped to the dual space can be understood as a temperature annealing strategy (called ε -scaling by Schmitzer (2019)), in which dual-optimal \mathbf{u}, \mathbf{v} are found at each value of the decaying temperature $1/\gamma^{(t)}$. Notice here that due to strong duality, optimal $\mathbf{u}^*(\gamma^{(t)}), \mathbf{v}^*(\gamma^{(t)})$ are minimizers of the EOT dual objective $g(\mathbf{u}, \mathbf{v}; \gamma^{(t)})$ in (3). The sequence of dual problems correspond to a sequence of Bregman projections like (7) in the primal space.

The second item in Prop. 4.1 shows that in fact, intermediate dual problems need not be solved exactly to optimality for this procedure to ultimately arrive at $P^*(\gamma)$. In other words, we can use approximations for $\mathbf{u}^*(\gamma^{(t)})$ and $\mathbf{v}^*(\gamma^{(t)})$ in (10) and solve for the update vectors \mathbf{u} and \mathbf{v} to approximate $P^{(t+1)}$. Notice that an initialization $(\mathbf{u}^{(0)}, \mathbf{v}^{(0)})$ here corresponds to $\tilde{P}_{ij}^{(0)} = \exp\{u_i^{(0)} + v_j^{(0)}\}$ with $\gamma^{(0)} = 0$, which is why any rank-1 $P^{(0)} \in \mathbb{R}_{>0}^{n \times n}$ suffices. These observations will form the basis of MDOT (Alg. 1).

The third item in Prop. 4.1 bounds the excess cost of the primal objective at a given step $t \geq 1$ in terms of the entropies of the marginals \mathbf{r}, \mathbf{c} . It should be compared to the more standard upper bound $\gamma^{-1} \log n$ used in prior work (Altschuler et al., 2017; Dvurechensky et al., 2018; Lin et al., 2019).¹ Note $H_{\min}(\mathbf{r}, \mathbf{c}) \leq \log n$ for $\mathbf{r}, \mathbf{c} \in \Delta_n$ and therefore this is a tighter bound. The final item in Prop. 4.1 shows the one-step improvement in the linear objective *with equality*; this is in contrast to the more standard analysis of mirror descent procedures where the improvement is bounded above with an inequality. See Appx. A.2 for proofs.

4.2 A mirror descent method for optimal transport: MDOT

Here we introduce the generic MDOT framework shown in Alg. 1 which will be exemplified throughout this section. Note that a routine in L9 is defined using “...” as a placeholder for additional parameters that may be required by specific implementations. MDOT accepts as input problem parameters $C, \mathbf{r}, \mathbf{c}$, as well as user-provided positive scalars γ_i, γ_f, p, q . Here, γ_i will be the initial step size $\Delta_\gamma^{(0)} = \gamma^{(1)}$ of the (inexact) mirror descent procedure mimicking (8), while γ_f will be the final (inverse) temperature. As seen in L8, near the end of each outer loop iteration, MDOT decays the temperature under a schedule $\gamma^{(t+1)} = q\gamma^{(t)}$ until $\gamma^{(t)} = \gamma_f$ similarly to Schmitzer (2019); Feydy (2020). Observe that due to L1 and L3, if called with $\gamma_i \geq \gamma_f$, MDOT terminates after a single iteration of the while loop.

Algorithm 1 MDOT($C, \mathbf{r}, \mathbf{c}, \gamma_i, \gamma_f, p \geq 1, q > 1$)

```

1:  $t \leftarrow 1$ , done  $\leftarrow$  False,  $\gamma \leftarrow \gamma_i \wedge \gamma_f$ 
2: while not done do
3:   done  $\leftarrow \gamma == \gamma_f$ 
4:    $\varepsilon_d \leftarrow H_{\min}(\mathbf{r}, \mathbf{c})/\gamma^p$ 
5:    $(\tilde{\mathbf{r}}, \tilde{\mathbf{c}}) \leftarrow (1 - \frac{\varepsilon_d}{4}) \cdot (\mathbf{r}, \mathbf{c}) + \frac{\varepsilon_d}{4n} \cdot \mathbf{1}_{2n}$ 
6:   if  $t == 1$  then  $\mathbf{z}^{(t)} \leftarrow (\log \tilde{\mathbf{r}}, \log \tilde{\mathbf{c}})$ 
7:    $\mathbf{z}^{(t)} \leftarrow \text{BregmanProject}(\mathbf{z}^{(t)}, \gamma, C, \tilde{\mathbf{r}}, \tilde{\mathbf{c}}, \varepsilon_d/2)$ 
8:    $\gamma \leftarrow q\gamma \wedge \gamma_f$ 
9:    $\mathbf{z}^{(t+1)} \leftarrow \text{WarmStart}(\mathbf{z}^{(t)}; \dots)$ 
10:   $t \leftarrow t + 1$ 
11: end while
12:  $(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{z}^{(t-1)}$ ,  $P \leftarrow \exp\{\mathbf{u}\mathbf{1}_n^\top + \mathbf{1}_n\mathbf{v}^\top - \gamma_f C\}$ 
13: Output  $P \leftarrow \text{Round}(P, \mathbf{r}, \mathbf{c})$ 

```

At each value of the temperature, in L4 we first select a tolerance $\varepsilon_d = H_{\min}(\mathbf{r}, \mathbf{c})/\gamma^p$ for dual objective gradient norm $\|\nabla g\|_1$, where $p \geq 1$. This is in contrast to the more standard usage $\varepsilon_d \propto (\log n)/\gamma$ in prior work. In Appx. C, we show the advantage of using $H_{\min}(\mathbf{r}, \mathbf{c})$ over $\log n$ empirically (for $p = 1$), demonstrating substantial speedups of order $\log n/H_{\min}(\mathbf{r}, \mathbf{c})$ for Sinkhorn iteration. The case of $p > 1$ is discussed in detail in Sec. 4.2.2. In L5 of Alg. 1, we “smooth” the target marginals \mathbf{r} and \mathbf{c} by mixing in uniform distributions, similarly to prior work by Dvurechensky et al. (2018) and Lin et al. (2019). This step helps provide convergence guarantees for certain choices of minimization algorithms for problem (3) used in L7. Since the mixing weight used in L5 is proportional to γ^{-p} , i.e., it gradually decreases with the temperature, this scheme smoothes marginals more aggressively in earlier iterations of MDOT. In Appx. D, we empirically show the performance benefits of this variable smoothing.

¹This $\log n$ term appears in the time complexity of various algorithms, but is often hidden in \tilde{O} -notation.

In L6, the dual variables are initialized as $\mathbf{u} = \log \tilde{\mathbf{r}}$ and $\mathbf{v} = \log \tilde{\mathbf{c}}$, which corresponds to initializing $\tilde{P}^{(0)} = \tilde{\mathbf{r}}\tilde{\mathbf{c}}^\top \in \mathcal{U}(\tilde{\mathbf{r}}, \tilde{\mathbf{c}})$, i.e., the maximum entropy plan in $\mathcal{U}(\tilde{\mathbf{r}}, \tilde{\mathbf{c}})$, which is rank-1 as required by Prop. 4.1. Next in L7, given the initial guess for the dual variables $\mathbf{z} = (\mathbf{u}, \mathbf{v})$ the objective in (3) is minimized under the smoothed marginals $\tilde{\mathbf{r}}, \tilde{\mathbf{c}}$ until $\|\mathbf{r}(P) - \tilde{\mathbf{r}}\|_1 + \|\mathbf{c}(P) - \tilde{\mathbf{c}}\|_1 \leq \varepsilon_d/2$. This step is called BregmanProject as it carries out (inexact) Bregman projection of the plan $P(\mathbf{u}^{(t)}, \mathbf{v}^{(t)}; \gamma)$ onto $\mathcal{U}(\tilde{\mathbf{r}}, \tilde{\mathbf{c}})$ as in (7). The algorithm to be used for minimizing the dual objective is left unspecified here for generality; for example, Sinkhorn iteration can be used. In 4.3, we introduce a new algorithm (PNCG) for this purpose. L8 sets $\gamma^{(t+1)} = q\gamma^{(t)}$, which corresponds to a default mirror descent step size $\Delta_\gamma = (q-1)\gamma$. In L9 we use the current near-optimal dual variable $\mathbf{z}^{(t)}$ (and perhaps additional $\mathbf{z}^{(t-k)}$'s for $k \geq 1$) to produce a warm-starting guess for the optimal dual variables at the next temperature value. Finally in L13, after an approximation P of $P^*(\gamma_f)$ has been obtained, P is rounded onto $\mathcal{U}(\mathbf{r}, \mathbf{c})$ via Alg. 2 of Altschuler et al. (2017); see Alg. 3 in Appx. A.

Remark 4.2. If $\gamma_f = 5H_{\min}(\mathbf{r}, \mathbf{c})/2\varepsilon$, the output $P \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ of Alg. 1 satisfies $\langle P - P^*, C \rangle \leq \varepsilon + \tilde{O}(\varepsilon^2)$.

Note that the above guarantee can be obtained considering only the very last iteration of MDOT. That is, the specific assignments in L4-5 in prior iterations, choices p, q and implementations of L7 and L9 play a significant role on the overall runtime of the algorithm (as we shall see), but not the guaranteed precision of the final solution returned as per Remark 4.2 above.

In the remainder of this section, we will:

1. Develop a warm-starting routine and compare it to a prior approach in temperature annealing.
2. Investigate the implications of various choices for the p and q parameters via ablation experiments.
3. Show the potential advantage of choosing $p > 1$ theoretically (Prop. 4.4).
4. Investigate a new BregmanProject routine based on non-linear CG for problem (3) in L7 of MDOT.

4.2.1 Warm-starting Bregman Projections

Assume that at each prior temperature value, we obtained the dual-optimal $\mathbf{u}^*(\gamma^{(t)})$, $\mathbf{v}^*(\gamma^{(t)})$ without error. How should \mathbf{u}, \mathbf{v} be initialized for $\gamma^{(t+1)}$? A simple, memory-efficient approach is to consider a Taylor expansion around recent γ to predict $\mathbf{u}^*(\gamma^{(t+1)})$, $\mathbf{v}^*(\gamma^{(t+1)})$. Letting $\mathbf{z} = (\mathbf{u}, \mathbf{v})$ to reduce clutter:

$$\mathbf{z}^*(\gamma^{(t+1)}) = \mathbf{z}^*(\gamma^{(t)}) + \frac{\partial \mathbf{z}^*}{\partial \gamma}(\gamma^{(t)})(\gamma^{(t+1)} - \gamma^{(t)}) + \dots \quad (11)$$

As we cannot compute $\partial \mathbf{z}^*/\partial \gamma$ analytically, we use a numerical approximation (backward finite differencing):

$$\frac{\partial \mathbf{z}^*}{\partial \gamma}(\gamma^{(t)}) \approx \frac{\mathbf{z}^*(\gamma^{(t)}) - \mathbf{z}^*(\gamma^{(t-1)})}{\gamma^{(t)} - \gamma^{(t-1)}}. \quad (12)$$

Keeping the first two terms in (11) and rearranging:

$$\mathbf{z}^*(\gamma^{(t+1)}) \approx \mathbf{z}^*(\gamma^{(t)}) + \frac{\Delta_\gamma^{(t)}}{\Delta_\gamma^{(t-1)}} \left(\mathbf{z}^*(\gamma^{(t)}) - \mathbf{z}^*(\gamma^{(t-1)}) \right). \quad (13)$$

In contrast, the ε -scaling approach of Schmitzer (2019) and Feydy (2020) maintains reparametrized dual variables $\tilde{\mathbf{z}} := \mathbf{z}/\gamma$ as the temperature decays. Rewriting (2) in terms of $\tilde{\mathbf{z}}$ reveals that ε -scaling amounts to predicting $\mathbf{z}^*(\gamma^{(t+1)}) \approx (\gamma^{(t+1)}/\gamma^{(t)})\mathbf{z}^*(\gamma^{(t)})$, i.e., simply scaling the dual variables instead of modelling the trajectory of \mathbf{z}^* with a Taylor approximation, which we argue is a better approach.

Note that in practice, optimal \mathbf{z}^* are unknown; only approximate minimizers are available. Since the Taylor approximation and backward finite differencing in (12-13) already incur some amount of error, care should be taken not to over-optimize (3) in an effort to better approximate \mathbf{z}^* in intermediate steps t .

In Fig. 1, we present an empirical study with varying step sizes $\Delta_\gamma = (q-1)\gamma$ by ablating q , where the advantage of (13) over the ε -scaling warm-start of Schmitzer (2019) is demonstrated. On the left, MDOT warm-start initializes each dual problem closer to the solution than the ε -scaling approach. The quality of initial guesses increase markedly with decreasing temperature (left); at high temperatures dual problems are initialized very close to the solution with the gradient norm just a small multiple of the target ε_d . In contrast, the ε -scaling warm-start stays relatively fixed. For the same decay rate q , this translates to about $10\times$ gains in convergence speed or precision (mid-right). The performance gap widens for slow temperature decay (lower q), as MDOT benefits from reduced Taylor approximation errors given smaller step sizes Δ_γ .

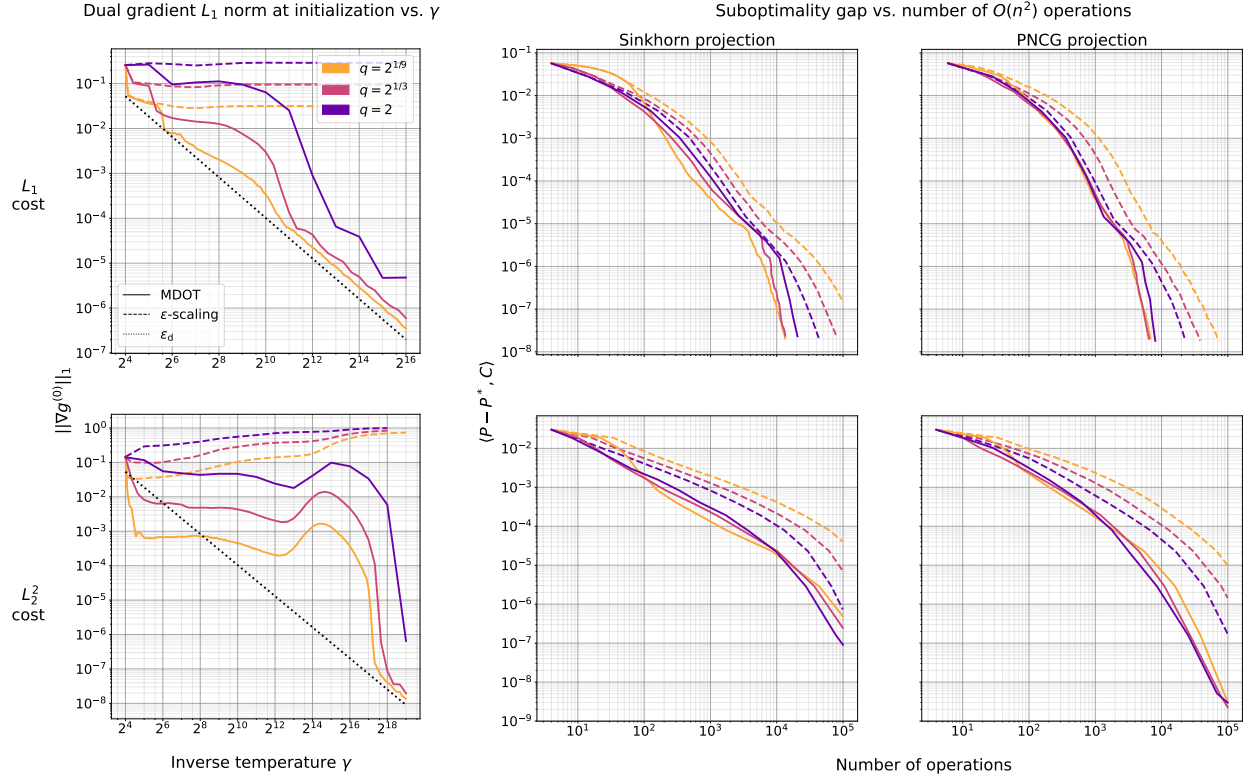


Figure 1: Comparison of the MDOT warm-start proposed in Sec. 4.2.1 to ε -scaling. Curves show the median over 36 upsampled-MNIST problems ($n=4096$) under L_1 (**top**) and L_2^2 (**bottom**) distance costs (see Sec. 5 for details). In all experiments, $p = 1.5$ and $\gamma_i = 2^4$. For the L_1 cost, $\gamma_f = 2^{16}$ and for the L_2^2 cost, $\gamma_f = 2^{19}$.

4.2.2 Bregman Projection Stopping Criteria

Our assignment $\varepsilon_d \propto \gamma^{-p}$ for some $p \geq 1$ in L4 of MDOT departs from the conventional wisdom of choosing $\varepsilon_d \propto \gamma^{-1}$ (Altschuler et al., 2017; Dvurechensky et al., 2018; Lin et al., 2019). Here, we aim to provide justification for this departure. Consider first the fixed-temperature problem (3) for simplicity. building on the results of Cominetti & Martín (1994), Weed (2018) showed in his Prop. 4 and Thm. 5 that there is both a uniform bound $\langle P^*(\gamma) - P^*, C \rangle \leq \log n / \gamma$ (slow rate), and a fast asymptotic rate $O(\exp(-\gamma K))$ which takes over for large enough γ , where the constant $K > 0$ is problem-dependent. Taking these as a starting point, the following remark generalizes the third statement of Prop. 4.1.

Remark 4.3. For any constant $p \in [1, \infty)$ and OT problem given by $(\mathbf{r}, \mathbf{c}, C)$, there exists a $\gamma_0 > 0$ such that for any $\gamma \geq \gamma_0$, we have $\langle P^*(\gamma) - P^*, C \rangle \leq H_{\min}(\mathbf{r}, \mathbf{c}) / \gamma^p$.

That is, below some temperature γ_0^{-1} , a stronger bound $H_{\min}(\mathbf{r}, \mathbf{c}) \gamma^{-p}$ for some $p > 1$ replaces the uniform bound $H_{\min}(\mathbf{r}, \mathbf{c}) \gamma^{-1}$. Thus, the SK algorithm (see Alg. 4 in the Appx.) can be tuned (via the p parameter in Alg. 1) to enjoy a rate substantially better than $O(n^2 \log n / \varepsilon^2)$ given by Dvurechensky et al. (2018).

Proposition 4.4. *Sinkhorn iteration, as instantiated by calling Alg. 1 (L7) with $p \in [1, \infty)$ and a sufficiently large $\gamma_i = \gamma_f = \sqrt[p]{5H_{\min}(\mathbf{r}, \mathbf{c})/2\varepsilon}$, returns a plan $P \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ satisfying $\langle P - P^*, C \rangle \leq \varepsilon + \tilde{O}(\varepsilon^2)$ in at most*

$$O\left(n^2 H_{\min}(\mathbf{r}, \mathbf{c})^{1/p} / \varepsilon^{\frac{p+1}{p}}\right) \text{ arithmetic operations.} \quad (14)$$

This result is consistent with the empirical findings of Jambulapati et al. (2019), who noted “The [tuned] Sinkhorn algorithm converged at rates much faster than the predicted ε^{-2} rate on all experiments, outperforming all other methods, which we believe merits further investigation.” We believe Prop. 4.4 sheds some light on this phenomenon, and further present an ablation of p in Fig. 2 for SK and MDOT algorithms.

For the SK algorithm, Fig. 2 (left) verifies the insight derived from Prop. 4.4. The choice $p = 1$ is better at low precision, but the trend gradually shifts in favor of higher p with (sufficiently) higher γ_f . That is,

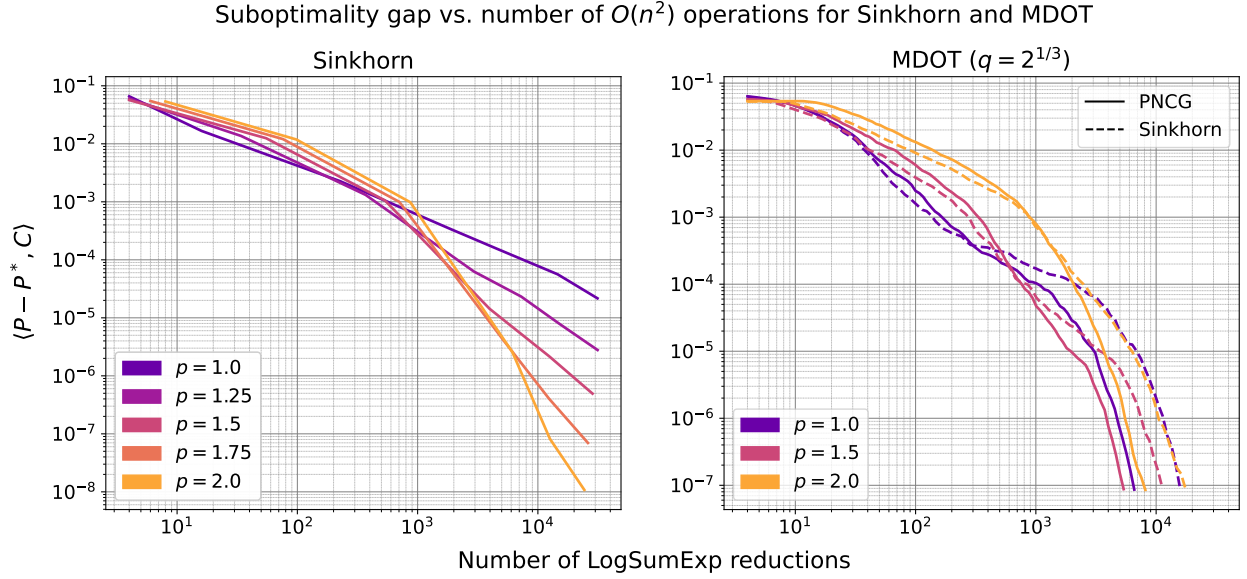


Figure 2: Ablation of stopping criterion parameter p (Sec. 4.2.2) for the SK algorithm (**left**) and MDOT with Sinkhorn and PNCG as Bregman projectors (**right**). The SK algorithm (**left**) is called by running MDOT (Alg. 1) with $\gamma_i = \gamma_f$, where higher precision is achieved by increasing γ_f . Results show the median over 36 random problems from the upsampled MNIST dataset ($n = 4096$) with the L_1 cost.

for sufficiently low temperature γ^{-1} , it is advantageous to reduce the gradient norm error tolerance, from H_{\min}/γ to H_{\min}/γ^p for $p > 1$. In contrast, MDOT is more robust to the p parameter in the high precision regime (right). Moreover, the use of PNCG projections (Alg. 2) for Bregman projections in MDOT (L7 of Alg. 1) provides a speedup of $2 - 3\times$ over Sinkhorn projections. PNCG is introduced and discussed next.

4.3 Preconditioned Non-linear Conjugate Gradients for Bregman Projections

SK converges more slowly at low temperatures (Kosowsky & Yuille, 1994).² For a faster alternative, we develop Alg. 2 based on non-linear CG (NCG) methods (Fletcher & Reeves, 1964; Nocedal & Wright, 2006), which we now briefly review. Given an objective g , NCG takes descent directions $\mathbf{p}^{(0)} = -\nabla g(\mathbf{z}^{(0)})$ and $\mathbf{p}^{(k)} \leftarrow -\nabla g^{(k)} + \beta^{(k)} \mathbf{p}^{(k-1)}$, and iterates $\mathbf{z}^{(k+1)} \leftarrow \mathbf{z}^{(k)} + \alpha^{(k)} \mathbf{p}^{(k)}$, where $\alpha^{(k)}$ is the step size. Optimal $\alpha^{(k)}$ has a closed-form solution for quadratics, but for general non-linear objectives, line search is necessary to find suitable step sizes $\alpha^{(k)}$. Various formulas for computing $\beta^{(k)}$ exist; for quadratic objectives, they are equivalent and guarantee convergence in at most n' iterations, where $n' \leq n$ is the number of distinct eigenvalues of $\nabla^2 g$. Further, the objective decreases faster if eigenvalues are tightly clustered (Stiefel, 1958; Kaniel, 1966; Nocedal & Wright, 2006). For example, the Hestenes-Stiefel formula sets (Nocedal & Wright, 2006):

$$\beta^{(k)} = \frac{\langle \nabla g^{(k)} - \nabla g^{(k-1)}, \nabla g^{(k)} \rangle}{\langle \nabla g^{(k)} - \nabla g^{(k-1)}, \mathbf{p}^{(k-1)} \rangle}. \quad (15)$$

A practical way to further improve the convergence rate of CG methods is via *preconditioning*. By making a change of variables $\mathbf{z} = M^{-1/2} \hat{\mathbf{z}}$ given some symmetric positive-definite matrix M , one reduces the

²In Appx. A we provide a proposition that applies Thm. 4 of Franklin & Lorenz (1989) on matrix scaling to the MD setting, which also characterizes this behavior, albeit asymptotically.

condition number of the problem or tightens the clustering of eigenvalues for improved convergence (ideally, $M^{-1} \approx \nabla^2 g^{-1}$). We refer the reader to [Hager & Zhang \(2006b\)](#) for further details on CG methods.

For the EOT problem, recall the 1st and 2nd order derivatives of the dual objective g in (3) at $\mathbf{z} = (\mathbf{u}, \mathbf{v})$:

$$\nabla g = (\mathbf{r}(P) - \mathbf{r}, \mathbf{c}(P) - \mathbf{c}), \quad \nabla^2 g = \begin{pmatrix} \mathbf{D}(\mathbf{r}(P)) & P \\ P^\top & \mathbf{D}(\mathbf{c}(P)) \end{pmatrix}. \quad (16)$$

A typical choice of a preconditioner M , known to be effective for diagonally-dominant matrices ([Golub & Van Loan, 2013](#)), is the diagonal approximation of the Hessian, which yields the following descent direction:

$$\tilde{\mathbf{s}} = -\mathbf{D}(\text{diag}(\nabla^2 g))^{-1} \nabla g = \left(\frac{\mathbf{r} - \mathbf{r}(P)}{\mathbf{r}(P)}, \frac{\mathbf{c} - \mathbf{c}(P)}{\mathbf{c}(P)} \right) = \left(\frac{\mathbf{r}}{\mathbf{r}(P)}, \frac{\mathbf{c}}{\mathbf{c}(P)} \right) - \mathbf{1}_{2n}. \quad (17)$$

Observe, however, that if at any point in the optimization $\mathbf{r}(P)$ or $\mathbf{c}(P)$ has infinitesimal entries, numerical instabilities may occur when evaluating $\tilde{\mathbf{s}}$. We propose using the *Sinkhorn direction*, \mathbf{s} , in place of $\tilde{\mathbf{s}}$:

$$\mathbf{s} = \left(\log \frac{\mathbf{r}}{\mathbf{r}(P)}, \log \frac{\mathbf{c}}{\mathbf{c}(P)} \right). \quad (18)$$

The Sinkhorn direction can be understood as the result of an alternative diagonal preconditioner, namely $M = \mathbf{D}(-\nabla g / \mathbf{s})$, since $\mathbf{s} = -M^{-1} \nabla g$. Furthermore, for any sub-optimal (\mathbf{u}, \mathbf{v}) , we have

$$-\langle \mathbf{s}, \nabla g \rangle = D_{\text{KL}}(\mathbf{r}(P) | \mathbf{r}) + D_{\text{KL}}(\mathbf{r} | \mathbf{r}(P)) + D_{\text{KL}}(\mathbf{c}(P) | \mathbf{c}) + D_{\text{KL}}(\mathbf{c} | \mathbf{c}(P)) > 0,$$

and therefore \mathbf{s} is also a descent direction. Empirically we find that this Sinkhorn preconditioner results in improved numerical stability. Finally, note that near the solution (for $\mathbf{r} \approx \mathbf{r}(P)$ and $\mathbf{c} \approx \mathbf{c}(P)$) we have

$$\mathbf{s} = \left(\log \frac{\mathbf{r}}{\mathbf{r}(P)}, \log \frac{\mathbf{c}}{\mathbf{c}(P)} \right) \approx \left(\frac{\mathbf{r}}{\mathbf{r}(P)}, \frac{\mathbf{c}}{\mathbf{c}(P)} \right) - \mathbf{1}_{2n} = \tilde{\mathbf{s}}, \quad (19)$$

where we have used $\log x \approx x - 1$ for $x \approx 1$. Therefore, near the solution, the Sinkhorn direction \mathbf{s} approaches the direction $\tilde{\mathbf{s}}$ obtained using the common preconditioner from the diagonal of the Hessian.

Plugging the preconditioner $M = \mathbf{D}(-\nabla g / \mathbf{s})$ into the preconditioned Hestenes-Stiefel formula ([Al-Baali & Fletcher, 1996](#)), we take $\beta^{(k)}$ in L7 of Alg. 2:

$$\beta^{(k)} = \frac{\langle \nabla g^{(k)} - \nabla g^{(k-1)}, -\mathbf{s}^{(k)} \rangle}{\langle \nabla g^{(k)} - \nabla g^{(k-1)}, \mathbf{p}^{(k-1)} \rangle}, \quad (20)$$

where $\mathbf{s}^{(k)}$ is the Sinkhorn direction as in (18), $\beta^{(1)} = 0$, $\mathbf{p}^{(0)} = \mathbf{0}_{2n}$. Observe that $-\mathbf{s}^{(k)}$ above simply replaces a $\nabla g^{(k)}$ term in the numerator of (15).

We defer details of the line search in L11 of Alg. 2 to Appx. B, but note that by design, the proposed line search only carries out the same form of LogSumExp reductions as the log-domain stabilized SK algorithm (Alg. 4 in the Appx A), so that its output is reused when evaluating the Sinkhorn direction \mathbf{s} in (18) at the next iteration (see L11 of Alg. 2). This also allows for a fair comparison of the two algorithms' performance. Indeed, Fig. 3 plots the runtime of the two algorithms in terms of LogSumExp evaluations; PNCG outshines SK empirically, especially at lower temperatures (higher γ). Further, to see whether the added numerical stability of the newly proposed Sinkhorn preconditioner comes at a performance trade-off, we implement an alternative stabilization scheme for the diagonal Hessian preconditioner. In particular, for this alternative we use \mathbf{s} only if the vector $(\mathbf{r}/\mathbf{r}(P), \mathbf{c}/\mathbf{c}(P))$ has any entries outside the range $[0.01, 100]$ and otherwise assign $\tilde{\mathbf{s}}$ given by (17) in L6 of Alg. 2. The results shown in Fig. 3 suggest that, on the contrary, the Sinkhorn preconditioner also provides a performance benefit over the diagonal Hessian in addition to numerical stability.

5 EXPERIMENTS

In this section, we first detail the MNIST experimental setup in Figs. 1-3. Then, we describe an additional color transfer task we use for benchmarking. Next, performance evaluations in terms of precision vs. wall-clock time are discussed given the results over 4 sets of problems shown in Fig. 4. In Appx. F, we add 20 more problem sets from the DOTmark benchmark of [Schrieber et al. \(2017\)](#) showing similar results both in terms of wall-clock time and operation counts. Lastly, the dependence on problem size n is investigated in Fig. 5. All experiments were performed on an NVIDIA GeForce RTX 2080 Ti GPU with 64-bit precision.

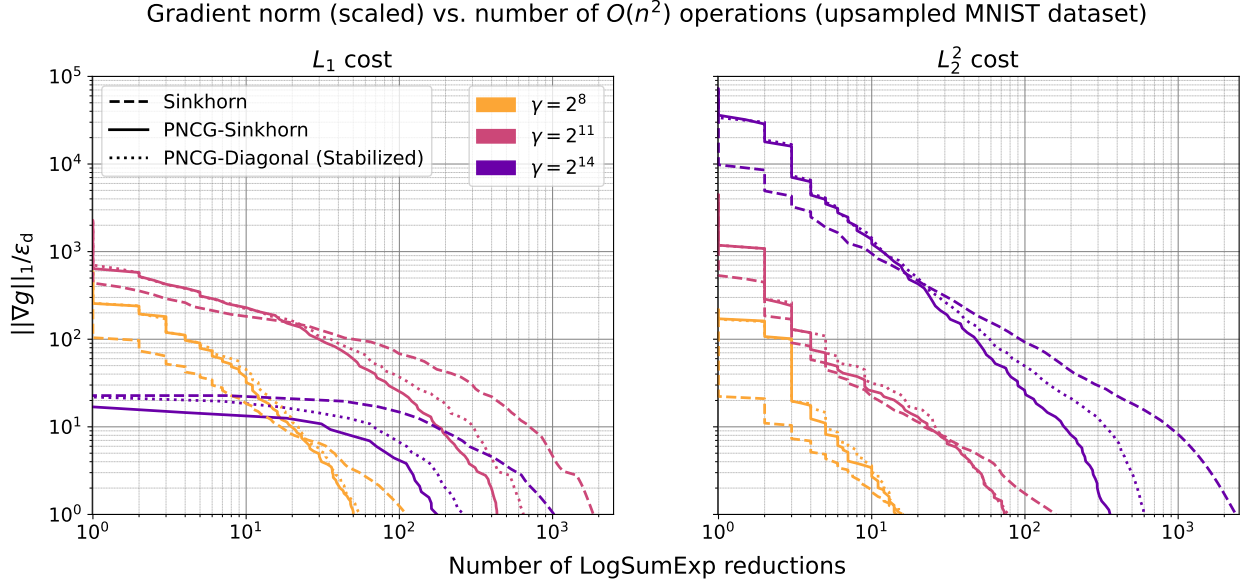


Figure 3: Comparison of BregmanProject algorithms (used in L7 of Alg. 1) over the upsampled MNIST dataset with L_1 (left) and L_2^2 (right) costs. Algorithms evaluated are the SK algorithm, the newly proposed PNCG given in Alg. 2, and a variant (see text). In all 36 problems, $n = 4096$, $p = 1.5$ and $q = 2$. Each curve shows the convergence behavior, at one specific temperature $1/\gamma$, in terms of the median number of LogSumExp reductions (x-axis) until gradient norm (y-axis) reaches below target dual gradient norm $\varepsilon_d(\gamma)$.

5.1 Experimental Setup

Upsampled MNIST. In line with prior work (Cuturi, 2013; Altschuler et al., 2017; Lin et al., 2022; Tang et al., 2024), we first consider the MNIST dataset, where each pixel represents an event and each image a probability distribution. Unlike prior work, we form higher dimensional problems by upsampling the original 28×28 images to be 64×64 (with bilinear interpolation) so that $n = 4096$. Cost matrices C are constructed by measuring the L_1 or squared L_2 distances between pixel locations on a 2D grid, and dividing all entries by the maximum distance value so that all entries of C lie in $[0, 1]$. The probability of each pixel is proportional to its intensity value; marginals \mathbf{r}, \mathbf{c} are obtained by flattening the pixel intensity matrices and subsequent L_1 normalization. To select m random problems, we sample $2m$ images from the dataset without replacement, and compute the OT distances between the first and second halves of the samples. Our selection of $n = 4096$ is driven by the objective of conducting a large number of tests per configuration to ensure statistically significant results, rather than by any inherent limitations of the algorithm. In fact, our MDOT code supports the use of on-the-fly CUDA kernels to evaluate entries of the cost matrix on the go using the PyKeOps package (Charlier et al., 2021). In this case, MDOT leaves an $O(n)$ memory footprint (with both Sinkhorn and PNCG projections) rather than $O(n^2)$; it has been verified to scale to much larger problems ($n \approx 100,000$).

Color Transfer. For the color transfer problem, each image is viewed as a point cloud in RGB space (pixel locations carry no importance). Cost matrices C are constructed by measuring the L_1 or L_2^2 distances between pixels in RGB space and dividing all entries by the maximum distance. Marginals \mathbf{r}, \mathbf{c} are taken to be uniform over Δ_n . With the help of GPT-4, we prompt DALL-E 2 to generate 20 vibrant and colorful images with intricate details or patterns. To match the dimensionality of the upsampled MNIST problem set, we downsample the original 1024×1024 images to 64×64 so that $n = 4096$. Once again, cost matrix entries are normalized to lie in $[0, 1]$.

5.2 Wall-clock Time Comparisons With Prior Work

In Fig. 4, we present wall-clock time benchmarking of MDOT (with both Sinkhorn and PNCG projections) against existing algorithms on the upsampled MNIST and color transfer problems. All benchmark methods were implemented in PyTorch and run on the GPU. For MDOT, we use $q=2^{1/3}$, $p = 1.5$ and $\gamma_i=2^4$ in all experiments. For the closely related Mirror Sinkhorn (MSK) algorithm of Ballu & Berthet (2023) the variable step size schedule prescribed by their Thm. 3.3 is used in our implementation. For Alg. 3.5 of Feydy (2020),

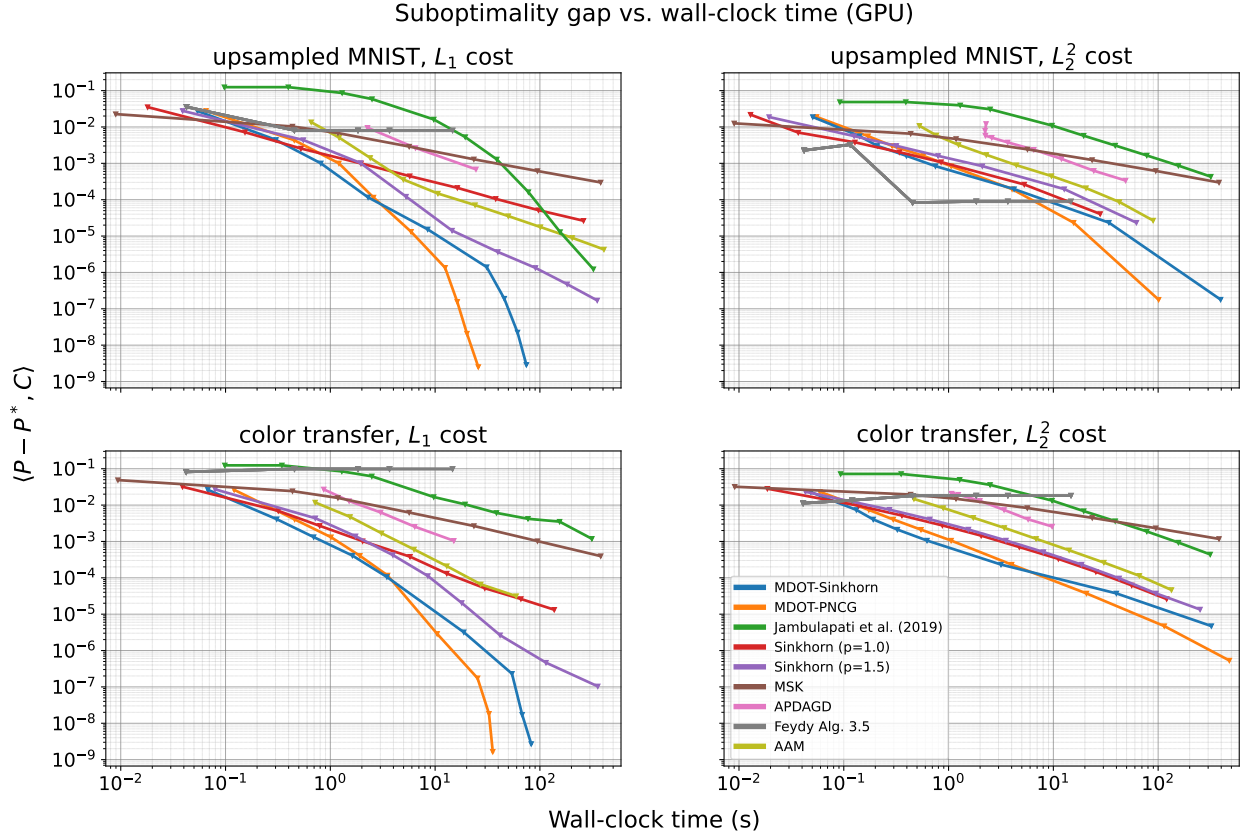


Figure 4: Wall-clock time vs. error benchmarking over the upsampled MNIST (top) and color transfer (bottom) problems using L_1 (left) and L_2^2 (right) distances as cost functions. Each marker shows the median time to converge (over 18 random problems) for each algorithm at a given hyperparameter setting, which controls the precision level, and the error $\langle P - P^*, C \rangle$ after rounding the output of the algorithm onto $\mathcal{U}(\mathbf{r}, \mathbf{c})$ – with the exception of Alg. 3.5 of Feydy (2020); see text.

we decay temperature at a rate $q = 0.7^{-1}$, which interpolates their *fast* ($q=0.5^{-1}$) and *safe* ($q=0.9^{-1}$) settings. For AAM (Guminov et al., 2021), Mirror Prox Sherman Optimized (Jambulapati et al., 2019) and APDAGD (Dvurechensky et al., 2018), each implementation closely follows an open-source NumPy implementation. Our PyTorch implementation was verified to produce identical results to the publicly available NumPy code. We additionally attempted comparison with APDAMD (Lin et al., 2019) and PDASMD (Luo et al., 2023), but observed extremely long convergence times for $n = 4096$ and omitted the results. For further details on the implementation of benchmark methods, we refer the reader to Appx. E.

While MDOT optimizes (3) to satisfy a convergence criterion following each temperature decrease, Alg. 3.5 of Feydy (2020) performs a *single* (symmetrized) Sinkhorn update instead, i.e., it does not minimize the sequence of dual objectives sufficiently despite taking increasingly large gradient steps in the dual space (cf. 6-7). This causes an accumulation of projection errors and results in the algorithm hitting a precision wall. Their *debiasing* option for estimating the OT distance via *Sinkhorn divergences* (introduced by Ramdas et al. (2017)) fares slightly better and is used here to comprise a stronger baseline, albeit this approach does not find a member of $\mathcal{U}(\mathbf{r}, \mathbf{c})$, which may be a strict requirement in some applications. MSK also runs a single row/column scaling update after a temperature decrease, but takes increasingly smaller steps and maintains a running average of transport plans to ensure convergence. It performs well at low precision, but shrinking step sizes slow it down, so that it exhibits $O(n^2 \varepsilon^{-2})$ convergence behavior. Sinkhorn iteration (log-domain stabilized, see Alg. 4 in the Appx. A) benefits substantially from setting $p = 1.5$ rather than $p = 1$ at sufficiently low temperatures for L_1 costs (see also Sec. 4.2.2). APDAGD underperforms SK with $p = 1$ and AAM performs similarly to it. Mirror Prox Sherman Optimized of Jambulapati et al. (2019) overtakes SK ($p = 1.0$) in one case only (top-left) in the high precision range. Meanwhile, MDOT-Sinkhorn enjoys faster convergence than the more competitive SK ($p = 1.5$) owing to warm-started temperature annealing,

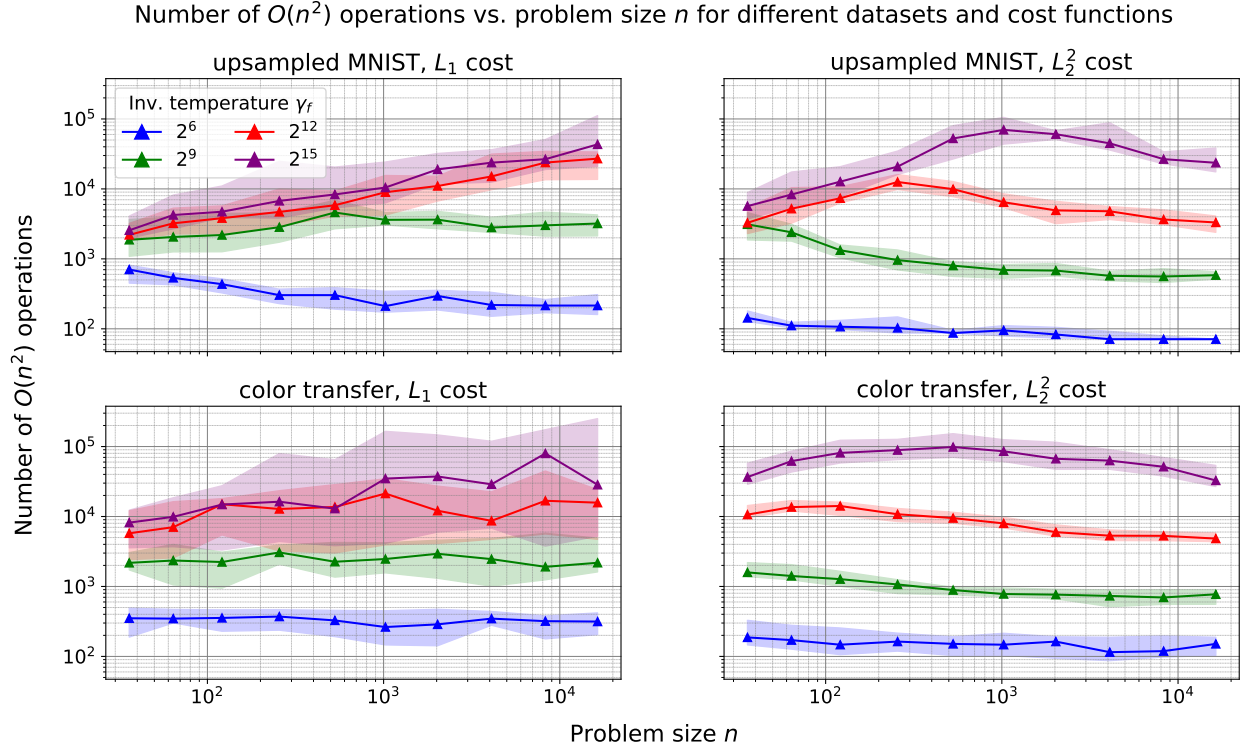


Figure 5: Problem size dependence of MDOT-PNCG convergence over the upsampled MNIST (top) and color transfer (bottom) problems using L_1 (left) and L_2^2 (right) distance cost functions with $\gamma_f \in \{2^6, 2^9, 2^{12}, 2^{15}\}$. Each marker displays the median over 20 problems and shaded areas show 75% confidence intervals.

especially in the high precision range (near the solution). MDOT-PNCG is the quickest to converge in all cases. The performance gap with its close second, MDOT-Sinkhorn, grows with higher precision.

5.3 Empirical Dependence on Problem Size of MDOT-PNCG

Our last set of experiments investigates the practical dependence of MDOT-PNCG on the problem size n . Over the same 4 problem sets as Fig. 4, we change n from 36 to 16,384 by up- or down-sampling images. The n values are selected to be approximately equally spaced on a logarithmic scale. In Fig. 5, we plot the behavior of MDOT-PNCG for a range of final temperature values $\gamma_f \in \{2^6, 2^9, 2^{12}, 2^{15}\}$. At medium precision (green and blue), we find that the algorithm behaves no worse than $O(n^2)$ in practice as implied by the flatness of the curves. As seen visually, at higher precision (roughly 5-decimals) with $\gamma_f \in \{2^{12}, 2^{15}\}$, the proposed GPU-parallel algorithm behaves roughly as $O(n^{5/2})$ at worst and even better for some of the problems in practice. These should be compared to the $\tilde{O}(n^{5/2})$ theoretical and $\tilde{O}(n^3)$ practical complexity of CPU-based exact solvers (Pele & Werman, 2009; Lee & Sidford, 2014).

6 CONCLUSION

In this work, we first presented a general procedure, MDOT, for computing OT distances with high precision and described its relation to a well-known temperature annealing strategy (ε -scaling). MDOT employs a novel warm-starting of the sequence of EOT dual problems encountered in temperature annealing, which was empirically shown to be highly effective compared to existing approaches. In addition, a specialized non-linear CG algorithm was developed as an alternative to Sinkhorn iteration and was shown to be more effective at low temperatures (under weak regularization). Over 24 different problem sets, the combined MDOT-PNCG algorithm outperforms aggressively tuned Sinkhorn iteration and many other recent baselines in terms of convergence of the primal suboptimality gap measured in wall-clock time. The algorithm was also shown to behave well with respect to the problem size. Interesting directions for future research include the theoretical convergence behavior of PNCG and bounds on the gradient norm of our warm-started initialization, as well as the development of faster Bregman projection algorithms and warm-starting methods. Adaptive temperature decay and dual problem stopping criteria are also of interest.

References

- Mehiddin Al-Baali and Robert Fletcher. On the order of convergence of preconditioned nonlinear conjugate gradient methods. *SIAM Journal on Scientific Computing*, 17(3):658–665, 1996.
- Jason Altschuler, Jonathan Niles-Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. *Advances in neural information processing systems*, 30, 2017.
- Marin Ballu and Quentin Berthet. Mirror Sinkhorn: Fast online optimization on transport polytopes. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 1595–1613. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/ballu23a.html>.
- Mathieu Blondel, Vivien Seguy, and Antoine Rolet. Smooth and sparse optimal transport. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 880–889. PMLR, 09–11 Apr 2018. URL <https://proceedings.mlr.press/v84/blondel18a.html>.
- Olivier Bokanowski and Benoit Grébert. Deformations of density functions in molecular quantum chemistry. *Journal of Mathematical Physics*, 37(4):1553–1573, 1996.
- Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45, 2015.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Benjamin Charlier, Jean Feydy, Joan Alexis Glaunes, François-David Collin, and Ghislain Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *The Journal of Machine Learning Research*, 22(1):3457–3462, 2021.
- Roberto Cominetti and J San Martín. Asymptotic analysis of the exponential penalty trajectory in linear programming. *Mathematical Programming*, 67:169–187, 1994.
- Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- Robert Dadashi, Leonard Hussenot, Matthieu Geist, and Olivier Pietquin. Primal Wasserstein imitation learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=TtYSU29zgR>.
- Simone Di Marino and Augusto Gerolin. Optimal transport losses and sinkhorn algorithm with general convex regularization. *arXiv preprint arXiv:2007.00976*, 2020.
- Julie Digne, David Cohen-Steiner, Pierre Alliez, Fernando De Goes, and Mathieu Desbrun. Feature-preserving surface reconstruction and simplification from defect-laden point sets. *Journal of mathematical imaging and vision*, 48:369–382, 2014.
- Pavel Dvurechensky, Alexander Gasnikov, and Alexey Kroshnin. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm. In *International conference on machine learning*, pp. 1367–1376. PMLR, 2018.
- Stephan Eckstein and Marcel Nutz. Quantitative stability of regularized optimal transport and convergence of sinkhorn’s algorithm. *SIAM Journal on Mathematical Analysis*, 54(6):5922–5948, 2022. URL <https://doi.org/10.1137/21M145505X>.
- Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite Markov decision processes. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI ’04, pp. 162–169, Arlington, Virginia, USA, 2004. AUAI Press.

- Sira Ferradans, Gui-Song Xia, Gabriel Peyré, and Jean-François Aujol. Static and dynamic texture mixing using optimal transport. In *Scale Space and Variational Methods in Computer Vision*, pp. 137–148, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- Sira Ferradans, Nicolas Papadakis, Gabriel Peyré, and Jean-François Aujol. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3):1853–1882, 2014. URL <https://doi.org/10.1137/130929886>.
- Jean Feydy. *Geometric data analysis, beyond convolutions*. PhD thesis, Université Paris-Saclay, 2020.
- Jean Feydy, Benjamin Charlier, François-Xavier Vialard, and Gabriel Peyré. Optimal transport for diffeomorphic registration. In *Medical Image Computing and Computer Assisted Intervention- MICCAI: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20*, pp. 291–299. Springer, 2017.
- Reeves Fletcher and Colin M Reeves. Function minimization by conjugate gradients. *The computer journal*, 7(2):149–154, 1964.
- Joel Franklin and Jens Lorenz. On the scaling of multidimensional matrices. *Linear Algebra and its applications*, 114:717–735, 1989.
- Uriel Frisch, Sabino Matarrese, Roya Mohayaee, and Andrei Sobolevski. A reconstruction of the initial conditions of the universe by optimal mass transportation. *Nature*, 417(6886):260–262, 2002.
- Aude Genevay, Gabriel Peyré, and Marco Cuturi. Learning generative models with Sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617. PMLR, 2018.
- Gene H Golub and Charles F Van Loan. *Matrix computations*. JHU press, 2013.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of Wasserstein GANs. *Advances in neural information processing systems*, 30, 2017.
- Sergey Guminov, Pavel Dvurechensky, Nazarii Tupitsa, and Alexander Gasnikov. On a combination of alternating minimization and Nesterov’s momentum. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3886–3898. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/guminov21a.html>.
- William W Hager and Hongchao Zhang. CG_DESCENT, a conjugate gradient method with guaranteed descent. *ACM Transactions on Mathematical Software (TOMS)*, 32(1):113–137, 2006a.
- William W Hager and Hongchao Zhang. A survey of nonlinear conjugate gradient methods. *Pacific journal of Optimization*, 2(1):35–58, 2006b.
- Arun Jambulapati, Aaron Sidford, and Kevin Tian. A direct $\tilde{O}(1/\epsilon)$ iteration parallel algorithm for optimal transport. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/024d2d699e6c1a82c9ba986386f4d824-Paper.pdf.
- Kirthevasan Kandasamy, Willie Neiswanger, Jeff Schneider, Barnabas Poczos, and Eric P. Xing. Neural architecture search with Bayesian optimisation and optimal transport. *Advances in neural information processing systems*, 31, 2018.
- Shmuel Kaniel. Estimates for some computational techniques in linear algebra. *Mathematics of Computation*, 20(95):369–378, 1966.
- Philip A Knight. The Sinkhorn–Knopp algorithm: convergence and applications. *SIAM Journal on Matrix Analysis and Applications*, 30(1):261–275, 2008.
- Jeffrey J Kosowsky and Alan L Yuille. The invisible hand algorithm: Solving the assignment problem with statistical physics. *Neural networks*, 7(3):477–490, 1994.

- Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in \tilde{O} (vrank) iterations and faster algorithms for maximum flow. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pp. 424–433. IEEE, 2014.
- Christian Léonard. From the Schrödinger problem to the Monge–Kantorovich problem. *Journal of Functional Analysis*, 262(4):1879–1920, 2012.
- Bruno Levy, Roya Mohayaee, and Sebastian von Hausegger. A fast semidiscrete optimal transport algorithm for a unique reconstruction of the early Universe. *Monthly Notices of the Royal Astronomical Society*, 506(1):1165–1185, 06 2021. URL <https://doi.org/10.1093/mnras/stab1676>.
- Tianyi Lin, Nhat Ho, and Michael I. Jordan. On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3982–3991. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/lin19a.html>.
- Tianyi Lin, Nhat Ho, and Michael I. Jordan. On the efficiency of entropic regularized algorithms for optimal transport. *Journal of Machine Learning Research*, 23(137):1–42, 2022. URL <http://jmlr.org/papers/v23/20-277.html>.
- Yiling Luo, Yiling Xie, and Xiaoming Huo. Improved rate of first order algorithms for entropic optimal transport. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 2723–2750. PMLR, 25–27 Apr 2023. URL <https://proceedings.mlr.press/v206/luo23a.html>.
- Quentin Mérigot. A multiscale approach to optimal transport. *Computer Graphics Forum*, 30(5):1583–1592, 2011. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2011.02032.x>.
- Arkadi Nemirovski and Dmitry Yudin. Problem complexity and method efficiency in optimization. *John Wiley & Sons*, 1983.
- Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, 2e edition, 2006.
- Ofir Pele and Michael Werman. Fast and robust Earth Mover’s distances. In *2009 IEEE 12th international conference on computer vision*, pp. 460–467. IEEE, 2009.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- François Pitie, Anil Kokaram, and Rozenn Dahyot. N-dimensional probability density function transfer and its application to color transfer. In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, volume 2, pp. 1434–1439 Vol. 2, 2005.
- Julien Rabin, Sira Ferradans, and Nicolas Papadakis. Adaptive color transfer with relaxed optimal transport. In *2014 IEEE international conference on image processing (ICIP)*, pp. 4852–4856. IEEE, 2014.
- Aaditya Ramdas, Nicolás García Trillos, and Marco Cuturi. On Wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017.
- Bernhard Schmitzer. A sparse multiscale algorithm for dense optimal transport. *Journal of Mathematical Imaging and Vision*, 56:238–259, 2016.
- Bernhard Schmitzer. Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM Journal on Scientific Computing*, 41(3):A1443–A1481, 2019. URL <https://doi.org/10.1137/16M1106018>.
- Jörn Schrieber, Dominic Schuhmacher, and Carsten Gottschlich. Dotmark – a benchmark for discrete optimal transport. *IEEE Access*, 5:271–282, 2017.

- Zhengyang Shen, Jean Feydy, Peirong Liu, Ariel H Curiale, Ruben San Jose Estepar, Raul San Jose Estepar, and Marc Niethammer. Accurate point cloud registration with robust optimal transport. In *Advances in Neural Information Processing Systems*, volume 34, pp. 5373–5389. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/2b0f658cbffd284984fb11d90254081f-Paper.pdf.
- Richard Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly*, 74(4):402–405, 1967.
- Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.
- Eduard L Stiefel. Kernel polynomials in linear algebra and their numerical applications. *Nat. Bur. Standards Appl. Math. Ser.*, 49:1–22, 1958.
- Xun Tang, Michael Shavlovsky, Holakou Rahmanian, Elisa Tardini, Kiran Koshy Thekumparampil, Tesi Xiao, and Lexing Ying. Accelerating Sinkhorn algorithm with sparse Newton iterations. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Kuj5gVp5GQ>.
- Jonathan Weed. An explicit analysis of the entropic penalty in linear programming. In *Conference On Learning Theory*, pp. 1841–1855. PMLR, 2018.
- Philip Wolfe. Convergence conditions for ascent methods. *SIAM review*, 11(2):226–235, 1969.
- Philip Wolfe. Convergence conditions for ascent methods. ii: Some corrections. *SIAM review*, 13(2):185–188, 1971.
- Yujia Xie, Xiangfeng Wang, Ruijia Wang, and Hongyuan Zha. A fast proximal point method for computing exact Wasserstein distance. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pp. 433–453. PMLR, 22–25 Jul 2020. URL <https://proceedings.mlr.press/v115/xie20b.html>.
- Lei Yang and Kim-Chuan Toh. Bregman proximal point algorithm revisited: A new inexact version and its inertial variant. *SIAM Journal on Optimization*, 32(3):1523–1554, 2022.
- G Zoutendijk. Nonlinear programming: a numerical survey. *SIAM Journal on Control*, 4(1):194–210, 1966.

A Proofs

A.1 Derivation of dual objective minimization as Bregman projection

For the interested reader, we provide here a short sketch of the relationship

$$P(\mathbf{u}^*, \mathbf{v}^*; \gamma) = \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} D_{\text{KL}}(P|P(\mathbf{u}', \mathbf{v}'; \gamma)), \quad (21)$$

given any initial $\mathbf{u}', \mathbf{v}' \in \mathbb{R}^n$ and $\mathbf{u}^*, \mathbf{v}^* \in \arg \min_{\mathbf{u}, \mathbf{v} \in \mathbb{R}^n} g(\mathbf{u}, \mathbf{v}; \gamma)$. First, given the definition in Sec. 2 of D_{KL} for un-normalized $P \in \mathbb{R}_{>0}^{n \times n}$, observe that

$$\begin{aligned} D_{\text{KL}}(P|P(\mathbf{u}', \mathbf{v}'; \gamma)) &= \langle P, \log P - \log P(\mathbf{u}', \mathbf{v}'; \gamma) \rangle + \langle P(\mathbf{u}', \mathbf{v}'; \gamma) - P, \mathbf{1} \rangle \\ &= \langle P, \log P - \mathbf{u}' \mathbf{1}^\top - \mathbf{1} \mathbf{v}'^\top + \gamma C \rangle + \langle P(\mathbf{u}', \mathbf{v}'; \gamma) - P, \mathbf{1} \rangle \\ &= \langle P, -\mathbf{1} + \gamma C + \log P \rangle - \langle \mathbf{u}', \mathbf{r}(P) \rangle - \langle \mathbf{v}', \mathbf{c}(P) \rangle + \text{const}. \end{aligned}$$

We shall apply the method of Lagrange multipliers. Given the equality constraints $\mathbf{r}(P) = \mathbf{r}$ and $\mathbf{c}(P) = \mathbf{c}$, we write the Lagrangian:

$$\mathcal{L}(P, \mathbf{u}, \mathbf{v}) = \langle P, -\mathbf{1} + \gamma C + \log P \rangle - \langle \mathbf{u}', \mathbf{r}(P) \rangle - \langle \mathbf{v}', \mathbf{c}(P) \rangle - \langle \mathbf{u}, \mathbf{r}(P) - \mathbf{r} \rangle - \langle \mathbf{v}, \mathbf{c}(P) - \mathbf{c} \rangle.$$

Taking the partial derivative with respect to P_{ij} yields

$$\frac{\partial \mathcal{L}}{\partial P_{ij}} = \gamma C_{ij} + \log P_{ij} - u'_i - u_i - v'_j - v_j.$$

Setting the partial to zero, we obtain:

$$P_{ij}^* = \exp(u_i + u'_i + v_j + v'_j - \gamma C_{ij}).$$

Plugging the above into the Lagrangian reveals that $-\mathcal{L}(P^*, \mathbf{u}, \mathbf{v}) = g(\mathbf{u}, \mathbf{v})$, i.e., minimizing the KL divergence $D_{\text{KL}}(P|P(\mathbf{u}', \mathbf{v}'; \gamma))$ over the feasible set $\mathcal{U}(\mathbf{r}, \mathbf{c})$ can be cast as minimizing the dual (3).

A.2 Proof of Proposition 4.1

First, we write the following helper lemma.

Lemma A.1 (A mirror descent bound for linear objectives). *Given a linear objective function $f(P) = \langle P, C \rangle$, an initial point $P^{(0)} \in \mathcal{F}$, an optimal point P^* and any $T > 0$, a sequence $[P^{(t)}]_{t \in \mathbb{N}}$ obtained via (4) satisfies:*

$$f(P^{(T)}) - f(P^*) \leq \frac{D_h(P^*|P^{(0)})}{\sum_{t=0}^{T-1} \Delta^{(t)}}. \quad (22)$$

Proof. Recall the definition of $\hat{P}^{(t+1)}$ from mirror descent iterates in (6):

$$\hat{P}^{(t+1)} = \nabla h^{-1} \left(\nabla h(P^{(t)}) - \Delta^{(t)} \nabla f(P^{(t)}) \right)$$

For any $P \in \mathcal{D}$,

$$\begin{aligned} f(P^{(t+1)}) - f(P) &= \langle \nabla f(P^{(t)}), P^{(t+1)} - P \rangle && \text{(since } f \text{ is linear)} \\ &= \frac{1}{\Delta^{(t)}} \langle \nabla h(P^{(t)}) - \nabla h(\hat{P}^{(t+1)}), P^{(t+1)} - P \rangle && \text{(due to (6))} \\ &\leq \frac{1}{\Delta^{(t)}} \langle \nabla h(P^{(t)}) - \nabla h(P^{(t+1)}), P^{(t+1)} - P \rangle && \text{(by Lemma 4.1 in Bubeck (2015))} \\ &= \frac{1}{\Delta^{(t)}} \left(D_h(P|P^{(t)}) - D_h(P|P^{(t+1)}) - D_h(P^{(t+1)}|P^{(t)}) \right) && \text{(by Eq. 4.1 in Bubeck (2015))} \\ &\leq \frac{1}{\Delta^{(t)}} \left(D_h(P|P^{(t)}) - D_h(P|P^{(t+1)}) \right), && \text{(since } D_h \geq 0) \end{aligned}$$

which implies

$$\Delta^{(t)}(f(P^{(t+1)}) - f(P)) \leq D_h(P|P^{(t)}) - D_h(P|P^{(t+1)}).$$

The above inequality proves monotonic improvement in each step t once we take $P = P^{(t)}$. Letting $P = P^*$, taking a telescopic sum and dividing both sides by $\sum_{s=0}^{T-1} \Delta^{(s)}$ we arrive at:

$$\begin{aligned} \frac{\sum_{t=0}^{T-1} \Delta^{(t)}(f(P^{(t+1)}) - f(P^*))}{\sum_{s=0}^{T-1} \Delta^{(s)}} &\leq \frac{D_h(P^*|P^{(0)}) - D_h(P^*|P^{(T)})}{\sum_{s=0}^{T-1} \Delta^{(s)}} \\ &\leq \frac{D_h(P^*|P^{(0)})}{\sum_{s=0}^{T-1} \Delta^{(s)}}, \end{aligned}$$

which implies (22) since improvement is monotonic and the first term on the LHS is a convex combination of objective values. \blacksquare

Proposition 4.1. Suppose $P^{(0)} \in \mathbb{R}_{>0}^{n \times n}$ is rank-1 and $P^{(t)}$ are computed via (8) for $t \geq 1$. Let $\gamma^{(0)} = 0$ and $\gamma^{(t+1)} = \gamma^{(t)} + \Delta_\gamma^{(t)}$, which together imply $\gamma^{(t+1)} = \sum_{t'=0}^t \Delta_\gamma^{(t')}$. Given $P^* \in \mathcal{U}(\mathbf{r}, \mathbf{c})$, a minimizer of $\langle P, C \rangle$, and $H_{\min}(\mathbf{r}, \mathbf{c}) := \min(H(\mathbf{r}), H(\mathbf{c}))$, the following are true:

1. $P^{(t)} = P^*(\gamma^{(t)})$, where $P^*(\gamma^{(t)})$ is the unique solution of (1) for $\gamma = \gamma^{(t)}$.
2. Let $\tilde{P}^{(t)} \in \mathbb{R}_{>0}^{n \times n}$ be any matrix of the form $\tilde{P}_{ij}^{(t)} = \exp\{u_i + v_j - \gamma^{(t)} C_{ij}\}$ for some $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$. In addition to (8), we also have $P^{(t+1)} = \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} \left\{ \langle P, C \rangle + \frac{1}{\Delta_\gamma^{(t)}} D_{\text{KL}}(P|\tilde{P}^{(t)}) \right\}$.
3. $\langle P^{(t)} - P^*, C \rangle \leq H_{\min}(\mathbf{r}, \mathbf{c})/\gamma^{(t)}$.
4. $\langle P^{(t)} - P^{(t+1)}, C \rangle = \frac{1}{\Delta_\gamma^{(t)}} \left(D_{\text{KL}}(P^{(t)}|P^{(t+1)}) + D_{\text{KL}}(P^{(t+1)}|P^{(t)}) \right)$ for all $t \geq 0$.

Proof. We prove each of the four statements in order.

Proof of the 1st statement. Here, we provide the derivations for (9-10). Consider a mirror descent problem given by (8):

$$P^{(t+1)} = \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} \left\{ \langle P, C \rangle + \frac{1}{\Delta_\gamma^{(t)}} D_{\text{KL}}(P|P^{(t)}) \right\}.$$

To derive the dual problem, we write the (scaled) Lagrangian with dual variables $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ corresponding to equality constraints $\mathbf{r}(P) = \mathbf{r}$ and $\mathbf{c}(P) = \mathbf{c}$:

$$\begin{aligned} \mathcal{L}(P, \mathbf{u}, \mathbf{v}) &= \Delta_\gamma^{(t)} \langle P, C \rangle + D_h(P|P^{(t)}) + \langle \mathbf{u}, \mathbf{r} - \mathbf{r}(P) \rangle + \langle \mathbf{v}, \mathbf{c} - \mathbf{c}(P) \rangle \\ &= \Delta_\gamma^{(t)} \langle P, C \rangle + \langle P, \log P \rangle - \langle P, \log P^{(t)} \rangle - \langle P - P^{(t)}, \mathbf{1} \rangle + \langle \mathbf{u}, \mathbf{r} - \mathbf{r}(P) \rangle + \langle \mathbf{v}, \mathbf{c} - \mathbf{c}(P) \rangle. \end{aligned}$$

Taking the first derivative with respect to the ij^{th} entry P_{ij} of P :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial P_{ij}} &= \Delta_\gamma^{(t)} C_{ij} + 1 + \log P_{ij} - \log P_{ij}^{(t)} - 1 - u_i - v_j \\ &= \Delta_\gamma^{(t)} C_{ij} + \log P_{ij} - \log P_{ij}^{(t)} - u_i - v_j \end{aligned}$$

Setting the partial to 0:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial P_{ij}} = 0 &\iff \log P_{ij} = \log P_{ij}^{(t)} + u_i + v_j - \Delta_\gamma^{(t)} C_{ij} \\ &\iff P_{ij} = P_{ij}^{(t)} \exp\{u_i + v_j - \Delta_\gamma^{(t)} C_{ij}\}, \end{aligned}$$

which is equivalent to (9). Under this reparametrization, we have:

$$\begin{aligned} P &= P^{(t)} \odot \exp\{\mathbf{u}\mathbf{1}^\top + \mathbf{1}\mathbf{v}^\top - \Delta_\gamma^{(t)}C\} \\ \log P &= \log P^{(t)} + \mathbf{u}\mathbf{1}^\top + \mathbf{1}\mathbf{v}^\top - \Delta_\gamma^{(t)}C, \end{aligned}$$

where \odot denotes the Hadamard product.

Now, plugging the above into the Lagrangian and maximizing with respect to the dual variables (and dropping constant terms), we obtain the dual problem (cf. 3):

$$\underset{\mathbf{u}, \mathbf{v} \in \mathbb{R}^n}{\text{minimize}} \quad \sum_{ij} P(\mathbf{u}, \mathbf{v})_{ij} - \langle \mathbf{u}, \mathbf{r} \rangle - \langle \mathbf{v}, \mathbf{c} \rangle. \quad (23)$$

Let $\hat{\mathbf{u}}^*(\gamma^{(t)})$ denote the dual-optimal variables at step t and $\mathbf{u}^*(\gamma^{(t)}) = \sum_{t'=0}^t \hat{\mathbf{u}}^*(\gamma^{(t')})$ (similarly for \mathbf{v}). Unrolling the recursion backwards in time:

$$P^{(t)} = P^{(0)} \odot \exp\left\{\mathbf{u}^*(\gamma^{(t)})\mathbf{1}^\top + \mathbf{1}\mathbf{v}^*(\gamma^{(t)})^\top - \gamma^{(t)}C\right\},$$

where we used the definition $\gamma^{(t)} = \sum_{t'=0}^t \Delta_\gamma^{(t')}$. Since $P^{(0)} \in \mathbb{R}_{>0}^n$ is rank-1 by construction, we can write it in the form $P^{(0)} = \exp\{\mathbf{u}_0\mathbf{1}^\top + \mathbf{1}\mathbf{v}_0^\top\}$ given some vectors $\mathbf{u}_0, \mathbf{v}_0 \in \mathbb{R}^n$. That is,

$$P^{(t)} = \exp\left\{(\mathbf{u}^*(\gamma^{(t)}) + \mathbf{u}_0)\mathbf{1}^\top + \mathbf{1}(\mathbf{v}^*(\gamma^{(t)}) + \mathbf{v}_0)^\top - \gamma^{(t)}C\right\}.$$

Then, we observe that problem (23) is identical to the dual EOT problem (3) up to a simple reparametrization of the dual variables. The result follows since strong duality holds (Slater's condition is satisfied given the feasible plan $\mathbf{r}\mathbf{c}^\top$) and the objective is strictly convex, i.e., the solution $P^*(\gamma)$ of the primal is unique.

Proof of the 2nd statement. Observe that at no point in the above proof of the 1st statement did we use the fact that $P^{(t)} \in \mathcal{U}(\mathbf{r}, \mathbf{c})$, nor did we use the optimality of $\hat{\mathbf{u}}^*(\gamma^{(t)}), \hat{\mathbf{v}}^*(\gamma^{(t)}) \in \mathbb{R}^n$. As long as each member of the sequence can be written in the form $\tilde{P}^{(t)} = \tilde{P}^{(t-1)} \odot \exp\{\hat{\mathbf{u}}(\gamma^{(t)})\mathbf{1}^\top + \mathbf{1}(\hat{\mathbf{v}}(\gamma^{(t)}))^\top - \gamma^{(t)}C\}$ for some dual variables at $\hat{\mathbf{u}}, \hat{\mathbf{v}} \in \mathbb{R}^n$, the recursion can be unrolled in the same way. That is, the dual problem remains identical in structure (as does its solution), but only the initial values of the dual variables change.

Proof of the 3rd statement. By Lemma A.1, for $P^{(0)} \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ we have

$$\langle P^{(t)} - P^*, C \rangle \leq \frac{D_h(P^*|P^{(0)})}{\sum_{t'=0}^{t-1} \Delta_\gamma^{(t')}}.$$

Given $\gamma = \gamma^{(t)} = \sum_{t'=0}^{t-1} \Delta_\gamma^{(t')}$, it remains to show that $D_h(P^*|P^{(0)}) \leq H_{\min}(\mathbf{r}, \mathbf{c})$.

Recall that for the negative entropy $h(\mathbf{x}) = \sum_i x_i \log x_i$, we have $D_h(\mathbf{x}|\mathbf{y}) = D_{\text{KL}}(\mathbf{x}|\mathbf{y})$. Suppose we take $P^{(0)} = \mathbf{r}\mathbf{c}^\top$:

$$\begin{aligned} D_{\text{KL}}(P^*|P^{(0)}) &= \sum_{ij} P_{ij}^* (\log P_{ij}^* - \log r_i c_j) \\ &= \sum_{ij} P_{ij}^* (\log P_{ij}^* - \log r_i - \log c_j) \\ &= -H(P^*) - \sum_i \log r_i \sum_j P_{ij}^* - \sum_j \log c_j \sum_i P_{ij}^* \\ &= -H(P^*) - \sum_i r_i \log r_i - \sum_j c_j \log c_j \quad (\text{since } P^* \in \mathcal{U}(\mathbf{r}, \mathbf{c})) \\ &= H(\mathbf{r}) + H(\mathbf{c}) - H(P^*) \\ &= \max(H(\mathbf{r}), H(\mathbf{c})) + \min(H(\mathbf{r}), H(\mathbf{c})) - H(P^*) \\ &\leq \min(H(\mathbf{r}), H(\mathbf{c})). \end{aligned}$$

The last inequality holds since $H(P) \geq H(\mathbf{r})$ and $H(P) \geq H(\mathbf{c})$ for any $P \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ (Cover, 1999), which together imply $H(P) \geq \max(H(\mathbf{r}), H(\mathbf{c}))$.

Given the 2nd statement, $P^{(t)}$ remains the same if $P^{(0)}$ was any non-negative rank-1 matrix, which can be written as $\exp\{\mathbf{u}_0 \mathbf{1}^\top + \mathbf{1} \mathbf{v}_0^\top\}$. Hence, the result holds for all such initializations, and not just $P^{(0)} = \mathbf{r} \mathbf{c}^\top$.

Proof of the 4th statement. First, note that given $h(P) = \sum_{ij} P_{ij} \log P_{ij}$, we have $\nabla h(P)_{ij} = 1 + \log P_{ij}$ and $\nabla h^{-1}(Q)_{ij} = \exp(Q_{ij} - 1)$. Then, given the definition of $\hat{P}^{(t+1)}$ from mirror descent iterates in (6):

$$\begin{aligned} \hat{P}^{(t+1)} &= \nabla h^{-1} \left(\nabla h(P^{(t)}) - \Delta_\gamma^{(t)} \nabla f(P^{(t)}) \right) \\ &= \exp(\log P^{(t)} - \Delta_\gamma^{(t)} C) \\ &= \exp \left(\mathbf{u}^*(\gamma^{(t)}) \mathbf{1}^\top + \mathbf{1} \mathbf{v}^*(\gamma^{(t)})^\top - (\gamma^{(t)} + \Delta_\gamma^{(t)}) C \right). \quad (\text{given } \mathbf{u}^*(\gamma^{(t)}), \mathbf{v}^*(\gamma^{(t)}) \in \arg \min g(\mathbf{u}, \mathbf{v}; \gamma^{(t)})) \\ &= \exp \left(\mathbf{u}^*(\gamma^{(t)}) \mathbf{1}^\top + \mathbf{1} \mathbf{v}^*(\gamma^{(t)})^\top - \gamma^{(t+1)} C \right). \end{aligned} \quad (24)$$

In the third equality, we used the known closed-form expression (2) to expand $P^{(t)}$.

In the special case that the feasible set $\mathcal{F} = \mathcal{U}(\mathbf{r}, \mathbf{c})$,

$$\begin{aligned} &\langle P^{(t)}, C \rangle - \langle P^{(t+1)}, C \rangle \\ &= \langle \nabla f(P^{(t)}), P^{(t)} - P^{(t+1)} \rangle \quad (\text{since } f = \langle P, C \rangle \text{ is linear}) \\ &= \frac{1}{\Delta_\gamma^{(t)}} \langle \nabla h(P^{(t)}) - \nabla h(\hat{P}^{t+1}), P^{(t)} - P^{(t+1)} \rangle \quad (\text{due to (6)}) \\ &= \frac{1}{\Delta_\gamma^{(t)}} \langle \nabla h(P^{(t)}) - \nabla h(P^{(t+1)}), P^{(t)} - P^{(t+1)} \rangle \quad (\text{see below}) \\ &= \frac{1}{\Delta_\gamma^{(t)}} \left(D_h(P^{(t)} | P^{(t+1)}) + D_h(P^{(t+1)} | P^{(t)}) \right) \quad (\text{by definition of the Bregman divergence as in (5)}) \\ &= \frac{1}{\Delta_\gamma^{(t)}} \left(D_{\text{KL}}(P^{(t)} | P^{(t+1)}) + D_{\text{KL}}(P^{(t+1)} | P^{(t)}) \right). \end{aligned}$$

To see why the third equality holds, observe that $P_{ij}^{t+1} = \hat{P}_{ij}^{t+1} \exp\{\hat{u}_i^* + \hat{v}_j^*\}$ for some optimal update vectors $\hat{\mathbf{u}}^*, \hat{\mathbf{v}}^* \in \mathbb{R}^n$ given the closed-forms (2) and (24). Then, for any $P, P' \in \mathcal{U}(\mathbf{r}, \mathbf{c})$,

$$\begin{aligned} &\langle \nabla h(P^{(t+1)}), P - P' \rangle \\ &= \sum_{ij} (1 + \log \hat{P}_{ij}^{t+1} + \hat{u}_i^* + \hat{v}_j^*) (P_{ij} - P'_{ij}) \\ &= \langle \nabla h(\hat{P}^{t+1}), P - P' \rangle + \sum_i \hat{u}_i^* \sum_j (P_{ij} - P'_{ij}) + \sum_j \hat{v}_j^* \sum_i (P_{ij} - P'_{ij}) \\ &= \langle \nabla h(\hat{P}^{t+1}), P - P' \rangle + \langle \hat{\mathbf{u}}^*, \mathbf{r} - \mathbf{r} \rangle + \langle \hat{\mathbf{v}}^*, \mathbf{c} - \mathbf{c} \rangle \quad (\text{since } P, P' \in \mathcal{U}(\mathbf{r}, \mathbf{c}) \text{ by construction}) \\ &= \langle \nabla h(\hat{P}^{t+1}), P - P' \rangle. \end{aligned} \quad \blacksquare$$

Algorithm 3 Round($P, \mathbf{r}, \mathbf{c}$) (Altschuler et al., 2017)

- 1: $X \leftarrow \mathbf{D}(\mathbf{x})$ with $\mathbf{x} = \mathbf{r}/\mathbf{r}(P) \wedge 1$
 - 2: $F \leftarrow XP$
 - 3: $Y \leftarrow \mathbf{D}(\mathbf{y})$ with $\mathbf{y} = \mathbf{c}/\mathbf{c}(F) \wedge 1$
 - 4: $F' \leftarrow FY$
 - 5: $\text{err}_r \leftarrow \mathbf{r} - \mathbf{r}(F'), \text{err}_c \leftarrow \mathbf{c} - \mathbf{c}(F')$
 - 6: Output $G \leftarrow F' + \text{err}_r \text{err}_c^\top / \|\text{err}_r\|_1$
-

Algorithm 4 Sinkhorn($\mathbf{z}, \gamma, C, \mathbf{r}, \mathbf{c}, \varepsilon_d$)

- 1: $(\mathbf{u}, \mathbf{v}) \leftarrow \mathbf{z}$
 - 2: $\log \mathbf{r}(P) \leftarrow \mathbf{u} + \text{LSE}_r(\mathbf{1}_n \mathbf{v}^\top - \gamma C)$
 - 3: **while** $\|\nabla g\|_1 = \|\mathbf{r} - \mathbf{r}(P)\|_1 > \varepsilon_d$ **do**
 - 4: $\mathbf{u} \leftarrow \mathbf{u} + \log \mathbf{r} - \log \mathbf{r}(P)$
 - 5: $\mathbf{v} \leftarrow \log \mathbf{c} - \text{LSE}_c(\mathbf{u} \mathbf{1}_n^\top - \gamma C)$
 - 6: $\log \mathbf{r}(P) \leftarrow \mathbf{u} + \text{LSE}_r(\mathbf{1}_n \mathbf{v}^\top - \gamma C)$
 - 7: **end while**
 - 8: Output $\mathbf{z} \leftarrow (\mathbf{u}, \mathbf{v})$
-

A.3 Proof of Proposition 4.4

In the remainder of this section, the L_1 norm $\|P\|_1$ of a matrix denotes the L_1 norm of the vectorized form of the matrix, and not the L_1 matrix norm.

First we state the following lemma, which is a simple combination of Lemmas 6 and 8 by [Weed \(2018\)](#).

Lemma A.2 (Entropy increase from mixing ([Weed, 2018](#))). *Let $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3 \in \Delta_n$ and $\mathbf{r}_2 = (1 - \varepsilon)\mathbf{r}_1 + \varepsilon\mathbf{r}_3$, where $\varepsilon \in (0, 1]$. We have,*

$$H(\mathbf{r}_2) \leq (1 - \varepsilon)H(\mathbf{r}_1) + \varepsilon H(\mathbf{r}_3) + \varepsilon(1 - \log \varepsilon) < H(\mathbf{r}_1) + \varepsilon(1 + \log \frac{n}{\varepsilon}). \quad (25)$$

Next, we provide a simple proof for Remark 4.3

Remark 4.3. For any constant $p \in [1, \infty)$ and OT problem given by $(\mathbf{r}, \mathbf{c}, C)$, there exists a $\gamma_0 > 0$ such that for any $\gamma \geq \gamma_0$, we have $\langle P^*(\gamma) - P^*, C \rangle \leq H_{\min}(\mathbf{r}, \mathbf{c})/\gamma^p$.

Proof. Recall from Thm. 5 of [Weed \(2018\)](#) that the quantity $\langle P^*(\gamma) - P^*, C \rangle$ decays at an exponential rate with increasing γ for sufficiently large γ . Since the exponential function $\exp\{-\gamma K\}$ decays more quickly than γ^{-p} for any constant $K > 0$ and finite p , we conclude that there exists some constant $\gamma_0 > 0$ such that

$$\langle P^*(\gamma) - P^*, C \rangle \leq H_{\min}(\mathbf{r}, \mathbf{c})/\gamma^p \quad (26)$$

for all optimal transport problems given by $\mathbf{r}, \mathbf{c}, C$ provided that $\gamma \geq \gamma_0$. ■

Proposition 4.4. *Sinkhorn iteration, as instantiated by calling Alg. 1 (L7) with $p \in [1, \infty)$ and a sufficiently large $\gamma_i = \gamma_f = \sqrt[p]{5H_{\min}(\mathbf{r}, \mathbf{c})/2\varepsilon}$, returns a plan $P \in \mathcal{U}(\mathbf{r}, \mathbf{c})$ satisfying $\langle P - P^*, C \rangle \leq \varepsilon + \tilde{O}(\varepsilon^2)$ in at most*

$$O\left(n^2 H_{\min}(\mathbf{r}, \mathbf{c})^{1/p} / \varepsilon^{\frac{p+1}{p}}\right) \text{ arithmetic operations.} \quad (14)$$

Proof. Let $B \in \mathcal{U}(\mathbf{r}', \mathbf{c}')$ be the transport plan $P(\mathbf{u}, \mathbf{v}) = \exp\{\mathbf{u}\mathbf{1}^\top + \mathbf{1}\mathbf{v}^\top - \gamma_f C\}$ after the termination of the main loop (before rounding in L13) of Alg. 1, which takes place after a single outer loop iteration in this setting, since $\gamma_i = \gamma_f$ by construction. Since B is the output of Sinkhorn iteration (Alg. 4), it lies on the simplex, as do its row and column marginals (specifically, we have $\mathbf{c}' = \mathbf{c}(B) = \tilde{\mathbf{c}} \in \Delta_n$ from Alg. 4). Furthermore, B is the unique optimizer of the EOT problem over $\mathcal{U}(\mathbf{r}', \mathbf{c}')$ due to Prop. 4.1 and the fact that it has the form $B_{ij} = \exp\{\mathbf{u}_i + \mathbf{v}_j - \gamma_f C_{ij}\}$:

$$B = \arg \min_{P \in \mathcal{U}(\mathbf{r}', \mathbf{c}')} \langle P, C \rangle - \frac{1}{\gamma_f} H(P). \quad (27)$$

Sinkhorn iteration used in L7 returns a solution \mathbf{u}, \mathbf{v} such that

$$\begin{aligned} & \|\tilde{\mathbf{r}} - \mathbf{r}(B)\|_1 + \|\tilde{\mathbf{c}} - \mathbf{c}(B)\|_1 \leq \varepsilon'/2 \\ \implies \|\nabla g\|_1 &= \|\mathbf{r} - \mathbf{r}(B)\|_1 + \|\mathbf{c} - \mathbf{c}(B)\|_1 \\ &\leq \|\mathbf{r} - \tilde{\mathbf{r}}\|_1 + \|\tilde{\mathbf{r}} - \mathbf{r}(B)\|_1 + \|\mathbf{c} - \tilde{\mathbf{c}}\|_1 + \|\tilde{\mathbf{c}} - \mathbf{c}(B)\|_1 \quad (\text{triangle inequality}) \\ &\leq \|\mathbf{r} - \tilde{\mathbf{r}}\|_1 + \|\mathbf{c} - \tilde{\mathbf{c}}\|_1 + \varepsilon'/2. \end{aligned}$$

Then, given mixing weights $\varepsilon'/4$ in L5 of Alg. 1:

$$\|\nabla g\|_1 \leq \varepsilon'. \quad (28)$$

Now, we make the following definitions:

- $\hat{B} = \text{Round}(B, \mathbf{r}, \mathbf{c})$, the rounding of B onto $\mathcal{U}(\mathbf{r}, \mathbf{c})$ via Alg. 2 of [Altschuler et al. \(2017\)](#), returned by our Alg. 1,
- $B^* \in \arg \min_{P \in \mathcal{U}(\mathbf{r}', \mathbf{c}')} \langle P, C \rangle$, an optimal plan in the feasible set $\mathcal{U}(\mathbf{r}', \mathbf{c}')$,

- $P^* \in \arg \min_{P \in \mathcal{U}(\mathbf{r}, \mathbf{c})} \langle P, C \rangle$, an optimal plan in the feasible set $\mathcal{U}(\mathbf{r}, \mathbf{c})$.

We have that,

$$\begin{aligned}
\langle \hat{B} - P^*, C \rangle &= \langle \hat{B} - B, C \rangle + \langle B - B^*, C \rangle + \langle B^* - P^*, C \rangle \\
&= \langle \hat{B} - B, C - \frac{1}{2} \mathbf{1}_{n \times n} \rangle + \langle B - B^*, C \rangle + \langle B^* - P^*, C \rangle \quad (\text{since } \hat{B}, B \in \Delta_{n \times n}) \\
&\leq \frac{1}{2} \left\| \hat{B} - B \right\|_1 + \langle B - B^*, C \rangle + \langle B^* - P^*, C \rangle \quad (\text{Hölder's ineq., given } C_{ij} \in [0, 1] \ \forall i, j \in [n]) \\
&\leq \|\nabla g\|_1 + \langle B - B^*, C \rangle + \langle B^* - P^*, C \rangle \quad (\text{by Lemma 7 of Altschuler et al. (2017)}) \\
&\leq \|\nabla g\|_1 + \frac{H_{\min}(\mathbf{r}', \mathbf{c}')}{\gamma_f^p} + \langle B^* - P^*, C \rangle \quad (\text{given (26-27), assuming } \gamma_f \text{ sufficiently large}) \\
&\leq \varepsilon' + \frac{H_{\min}(\mathbf{r}', \mathbf{c}')}{\gamma_f^p} + \langle \tilde{B} - P^*, C \rangle, \tag{29}
\end{aligned}$$

where \tilde{B} is any transport plan in $\mathcal{U}(\mathbf{r}', \mathbf{c}')$. We take \tilde{B} to be the “shadow” of P^* in the sense of Definition 3.1 of Eckstein & Nutz (2022), under the discrete metric. In other words, letting

$$\tilde{B} = \arg \min_{P \in \mathcal{U}(\mathbf{r}', \mathbf{c}')} \|P - P^*\|_1,$$

and noting that the 1-Wasserstein distance under the discrete metric is equal to the total variation (TV) distance, the first equation in Lemma 3.2 of Eckstein & Nutz (2022) yields the equality (*) below:

$$\frac{1}{2} \left\| \tilde{B} - P^* \right\|_1 = \text{TV}(B, P^*) \stackrel{(*)}{=} \text{TV}(\mathbf{r}, \mathbf{r}') + \text{TV}(\mathbf{c}, \mathbf{c}') = \frac{1}{2} \|\nabla g\|_1. \tag{30}$$

Then, continuing from (29),

$$\begin{aligned}
\langle \hat{B} - P^*, C \rangle &\leq \varepsilon' + \frac{H_{\min}(\mathbf{r}', \mathbf{c}')}{\gamma_f^p} + \langle \tilde{B} - P^*, C \rangle \\
&= \varepsilon' + \frac{H_{\min}(\mathbf{r}', \mathbf{c}')}{\gamma_f^p} + \langle \tilde{B} - P^*, C - \frac{1}{2} \mathbf{1}_{n \times n} \rangle \quad (\text{since } \tilde{B}, P^* \in \Delta_{n \times n}) \\
&= \varepsilon' + \frac{H_{\min}(\mathbf{r}', \mathbf{c}')}{\gamma_f^p} + \left\| \tilde{B} - P^* \right\|_1 \left\| C - \frac{1}{2} \mathbf{1}_{n \times n} \right\|_\infty \\
&\leq \frac{3}{2} \varepsilon' + \frac{H_{\min}(\mathbf{r}', \mathbf{c}')}{\gamma_f^p} \quad (\text{given (28-30)}) \\
&\leq \frac{3}{2} \varepsilon' + \frac{H_{\min}(\mathbf{r}, \mathbf{c})}{\gamma_f^p} + \frac{\varepsilon'}{\gamma_f^p} (1 + \log(n/\varepsilon')) \quad (\text{by Lemma A.2}) \\
&= \frac{5H_{\min}(\mathbf{r}, \mathbf{c})}{2\gamma_f^p} + \tilde{O}(\gamma_f^{-2p}) \quad (\text{since } \varepsilon' = \frac{H_{\min}(\mathbf{r}, \mathbf{c})}{\gamma_f^p} \text{ in L4 of Alg. 1}) \\
&= \varepsilon + \tilde{O}(\varepsilon^2). \quad (\text{since } \gamma_f = \left(5H_{\min}(\mathbf{r}, \mathbf{c})/2\varepsilon \right)^{1/p} \text{ by construction})
\end{aligned}$$

The computational complexity of the algorithm follows simply from the same line of reasoning as Thm. 1 and Thm. 2 of Dvurechensky et al. (2018). In particular, they show that Sinkhorn iteration converges in $O(R/\varepsilon')$ steps, where $R = O(\gamma_f) = O(H_{\min}(\mathbf{r}, \mathbf{c})^{1/p} \varepsilon^{-1/p})$ in our case. The complexity result $O(n^2 H_{\min}(\mathbf{r}, \mathbf{c})^{1/p} / \varepsilon^{\frac{p+1}{p}})$ follows since $\varepsilon' = O(H_{\min}(\mathbf{r}, \mathbf{c}) \gamma_f^{-p}) = O(\varepsilon^{-1})$, and each Sinkhorn step costs $O(n^2)$. ■

B An Efficient Line Search Algorithm

B.1 Background: Line Search

Given a descent direction $\mathbf{p}^{(k)} \in \mathbb{R}^n$, i.e., a direction that satisfies $\langle \mathbf{p}^{(k)}, \nabla f(\mathbf{x}^{(k)}) \rangle \leq 0$, line search algorithms aim to find an appropriate step size α , where $\mathbf{x}^{(k+1)} \leftarrow \mathbf{x}^{(k)} + \alpha \mathbf{p}^{(k)}$. Perhaps the most well-known of desirable

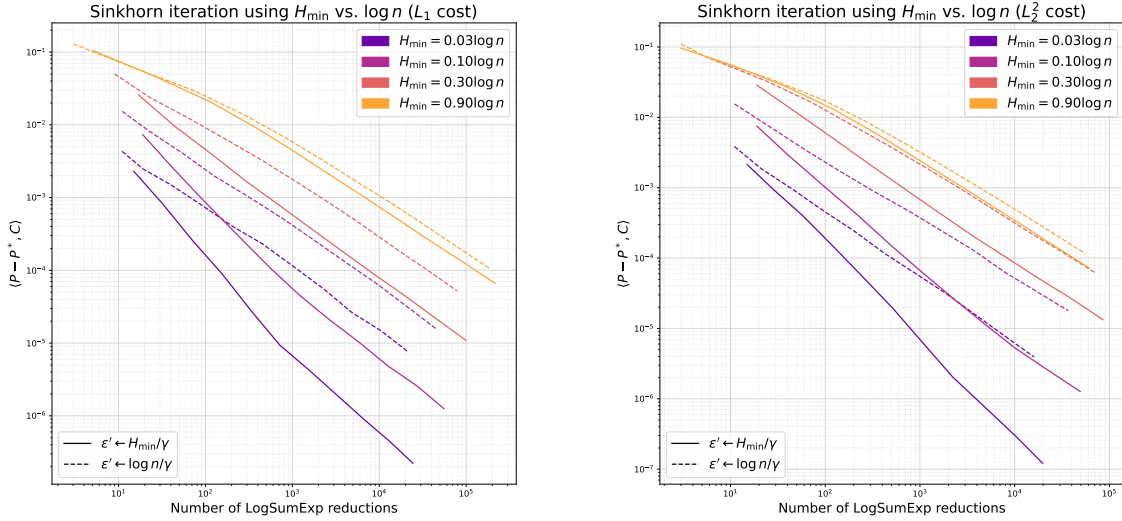


Figure 6: We control the problem parameter $H_{\min}(\mathbf{r}, \mathbf{c})$ over synthetically sampled OT problems $(\mathbf{r}, \mathbf{c}, C)$ and show that using the entropy-aware stopping criterion can yield substantial performance gains, with the gap between the two approaches growing proportionally to the gap between $H_{\min}(\mathbf{r}, \mathbf{c})$ and $\log n$. The experiments carried out use $\gamma \in \{2^4, 2^5, \dots, 2^{14}\}$ to control precision. We sample 18 problems for each γ value and plot the median along both axes. The results are consistent between L_1 (left) and L_2^2 (right) costs.

properties that a step size α should satisfy at any given optimization step are the Wolfe conditions (Wolfe, 1969; 1971). Given $\phi(\alpha) := f(\mathbf{x}^{(k)} + \alpha \mathbf{p}^{(k)})$:

$$\frac{\phi(\alpha) - \phi(0)}{\alpha} \leq c_1 \phi'(0) \quad (31a)$$

$$\phi'(\alpha) \geq c_2 \phi'(0). \quad (31b)$$

where $0 < c_1 < c_2 < 1$ and (31a) and (31b) are known as the *sufficient decrease* and *curvature* conditions respectively (Nocedal & Wright, 2006). It is well-known that given step sizes satisfying the Wolfe conditions and descent directions $\mathbf{p}^{(k)}$ that are *not* nearly orthogonal to the steepest descent directions $-\nabla f(\mathbf{x}^{(k)})$, line search methods ensure convergence of the gradient norms to zero (Zoutendijk, 1966; Wolfe, 1969; 1971). Instead of satisfying (31), some algorithms or theoretical analyses consider *exact* line search, where $\alpha^* \in \arg \min_{\alpha \in \mathbb{R}} \phi(\alpha)$, which has a unique closed-form solution for quadratic objectives with a positive definite Hessian. However, a rule of thumb for general non-linear objectives is to not spend too much time finding α^* (Nocedal & Wright, 2006). Hager & Zhang (2006a) proposed *approximate* Wolfe conditions, derived by replacing the $\phi(\alpha)$ and $\phi(0)$ terms in (31a) with $q(\alpha)$ and $q(0)$, where q is a quadratic interpolant of ϕ such that $q(0) = \phi(0)$, $q'(0) = \phi'(0)$ and $q'(\alpha) = \phi'(\alpha)$:

$$(2c_1 - 1)\phi'(0) \geq \phi'(\alpha) \geq c_2\phi'(0). \quad (32)$$

A key advantage of replacing (31) by (32) stems from the fact that one only needs to evaluate ϕ' rather than both ϕ and ϕ' to check whether the conditions are satisfied, thereby halving the amount of computation necessary per iteration in cases where their evaluation has similar computational cost.

Bisection is a line search strategy with convergence guarantees when the objective is convex. One simply maintains a bracket $[\alpha_{lo}, \alpha_{hi}]$, where $\phi'(\alpha_{lo}) < 0$ and $\phi'(\alpha_{hi}) > 0$, and recursively considers their average and updates either endpoint of the bracket given the sign of $\phi'((\alpha_{hi} + \alpha_{lo})/2)$.

B.2 PNCG Line Search

To perform line search in PNCG (Alg. 2), we adopt a hybrid strategy combining bisection and the secant method to find α_k satisfying approximate Wolfe conditions (32). Given α_{lo}, α_{hi} , the secant method computes the minimizer of a quadratic interpolant \hat{q} that satisfies $\hat{q}'(\alpha_{lo}) = \phi'(\alpha_{lo})$ and $\hat{q}'(\alpha_{hi}) = \phi'(\alpha_{hi})$ as follows:

$$\alpha_{\text{sec}} = \frac{\alpha_{lo}\phi'(\alpha_{hi}) - \alpha_{hi}\phi'(\alpha_{lo})}{\phi'(\alpha_{hi}) - \phi'(\alpha_{lo})}. \quad (33)$$

Thanks to the convexity of the objective g , by ensuring $\phi'(\alpha_{\text{lo}}) < 0$ and $\phi'(\alpha_{\text{hi}}) > 0$ with simple algorithmic checks, we can guarantee that $\alpha_{\text{lo}} < \alpha_{\text{sec}} < \alpha_{\text{hi}}$. Thus, the updated bracket is guaranteed to be smaller once we replace either of α_{lo} or α_{hi} by α_{sec} for the next bracket given the sign of $\phi'(\alpha_{\text{sec}})$. If ϕ behaves like a quadratic inside the bracket, the secant method converges very quickly, but convergence can be arbitrarily slow otherwise. For this reason, we simply average the bisection estimate and α_{sec} for a less aggressive but more reliable line search that still converges quickly, i.e., $\alpha_{\text{hybrid}} = 0.5\alpha_{\text{sec}} + 0.5(\alpha_{\text{hi}} + \alpha_{\text{lo}})/2$.

Evaluation of ϕ' has computational complexity $O(n^2)$ as does a single step of Sinkhorn’s algorithm (given by the two LogSumExp reductions seen in L5-6 of Alg. 4):

$$\phi'(\alpha) = \langle \mathbf{p}_u, \mathbf{r}(P_\alpha) - \mathbf{r} \rangle + \langle \mathbf{p}_v, \mathbf{c}(P_\alpha) - \mathbf{c} \rangle, \quad (34)$$

where $(\mathbf{p}_u, \mathbf{p}_v)$ is the descent direction. Since evaluating ϕ' requires the computation of $\mathbf{r}(P_\alpha)$ and $\mathbf{c}(P_\alpha)$ for the new matrix $P_\alpha := \exp\{(\mathbf{u} + \alpha\mathbf{p}_u)\mathbf{1}^\top + \mathbf{1}(\mathbf{v} + \alpha\mathbf{p}_v)^\top - \gamma C\}$, the last step of the line search readily carries out the LogSumExp reductions necessary for computing the Sinkhorn direction in the next step of PNCG (see L11 of Alg. 2). Observe also that at the next PNCG iteration, $\phi'(0)$ can also be computed in $O(n)$ time rather than $O(n^2)$ since $\mathbf{r}(P_0), \mathbf{c}(P_0)$ are already known from the last line search step of the previous PNCG iteration. With these important implementation details in place, we find that the average number of ϕ' evaluations necessary to find an α that satisfies (32) is typically between 1.5 – 2.5 for the PNCG algorithm. While the approach outlined here is easy to implement (including as a batch process) and works well in practice, better line search methods may further benefit Alg. 2.

C Entropy-aware Stopping Criteria on the Dual Objective Gradient Norm

Here, we show the effect of choosing $H_{\min}(\mathbf{r}, \mathbf{c})$ over the weaker bound $\log n$ in L4 of Alg. 1, where the stopping criterion ε' is selected. To control the problem setting $H_{\min}(\mathbf{r}, \mathbf{c})$, we construct synthetic problems by randomly sampling \mathbf{r} from the simplex via a Dirichlet distribution constructed to meet a target entropy level $H(\mathbf{r})$ as a fraction of the maximum possible entropy $\log n$. The column marginal \mathbf{c} is simply taken to be the uniform distribution $\mathbf{1}_n/n$, so that $H_{\min}(\mathbf{r}, \mathbf{c}) = H(\mathbf{r})$. Cost matrices are constructed by sampling n points $\mathbf{x} \in \mathbb{R}^3$ from a multivariate normal distribution and assigning $C_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|_r^r$ for $r \in \{1, 2\}$, before entrywise division by $\max_{ij} C_{ij}$ to ensure $C_{ij} \in [0, 1]$.

Fig. 6 illustrates the effect of this choice by ranging $H_{\min}(\mathbf{r}, \mathbf{c})/\log n \in \{0.03, 0.1, 0.3, 0.9\}$. Towards the RHS of the plots, we observe an improvement in precision roughly proportional to $\log n/H_{\min}(\mathbf{r}, \mathbf{c})$ for the same number of operations, which agree with our complexity result $O(n^2 H_{\min}(\mathbf{r}, \mathbf{c})/\varepsilon^{-2})$ for $p = 1$ in (26) vs. the $O(n^2 \log n/\varepsilon^{-2})$ result by Dvurechensky et al. (2018).

D Variable vs. Fixed Smoothing of the Marginals in MDOT

As discussed in Sec. 4.2, MDOT smoothes the marginals \mathbf{r}, \mathbf{c} (by mixing in the uniform distribution) with a weighting factor that tracks the temperature. Since MDOT anneals the temperature, this means that the smoothing weight is higher in earlier iterations of MDOT. In particular, the mixture weight gradually decays from $H_{\min}(\mathbf{r}, \mathbf{c})/4\gamma_i^p$ to $H_{\min}(\mathbf{r}, \mathbf{c})/4\gamma_f^p$ given input parameter $p \geq 1$. Here, we study the effect of this design choice as it influences the convergence of two Bregman projection algorithms used in L7 of Alg. 1: Sinkhorn iteration and the newly proposed PNCG algorithm (Alg. 2). The approach is benchmarked against a baseline that fixes the smoothing weight at $H_{\min}(\mathbf{r}, \mathbf{c})/4\gamma_f^p$ all throughout instead. For these experiments, we fix $p = 1.5$ following our experimental setup in Sec. 5. Fig. 7 shows that while MDOT-Sinkhorn is largely unaffected by this design choice, MDOT-PNCG enjoys a notable speedup from variable smoothing. We thus conclude that the approach provides a performance benefit.

E Details of Baseline Algorithm Implementations

Here, we provide details and sources on the implementation of various algorithms shown in Fig. 4. Our implementations of other algorithms will be open-sourced for transparency.

Mirror Prox Sherman Optimized (Jambulapati et al., 2019). For this algorithm, the source code is originated in the NumPy code at [this repository](#). The owner of the repository notes that this NumPy implementation is based on a Julia implementation by the original authors, which was provided in a private

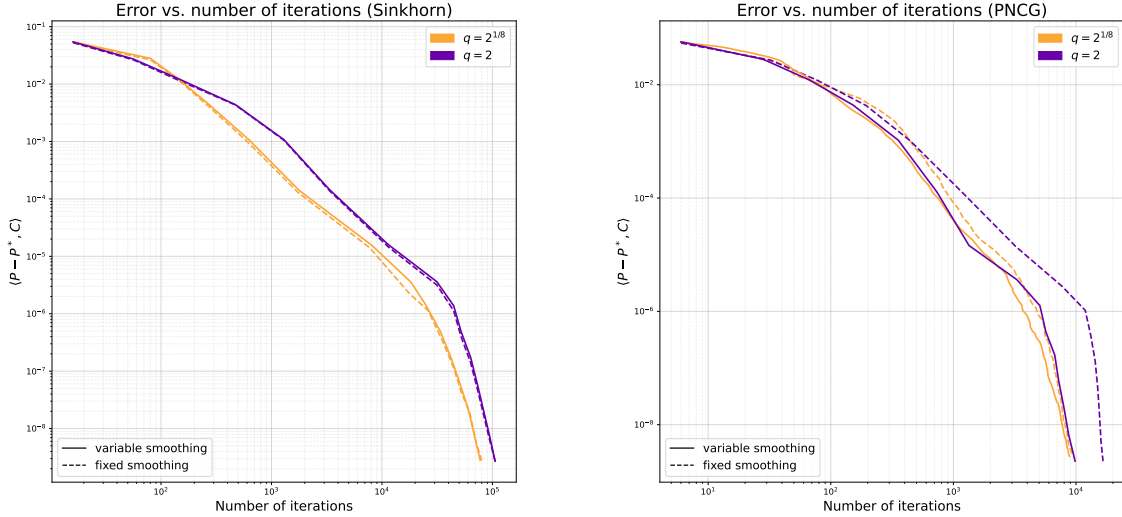


Figure 7: The variable smoothing scheme has almost no impact on the convergence behavior of MDOT-Sinkhorn (**left**) for both settings $q = 2$ (rapid temperature decay) and $q = 2^{1/8}$ (slow temperature decay). MDOT-PNCG (**right**) enjoys a speedup of nearly $2\times$ under rapid decay and a more modest speedup under slow decay. All curves show the median over 36 sample problems from the MNIST dataset (L_1 cost).

exchange. The code used in this paper is a PyTorch adaptation of the NumPy code and has been verified to produce identical output as the NumPy version over multiple problems. The algorithm was called with *entropy factor* parameter set to the default 2.75 in all experiments. The number of iterations for the algorithm was varied from 2 to 2^{15} to achieve different levels of precision.

APDAGD (Dvurechensky et al., 2018). For APDAGD, a similar strategy was used, except with [this code repository](#). A PyTorch version of the original NumPy code was written and verified to produce identical output. For different levels of precision, the ε parameter of the algorithm was varied from 2^{-1} to 2^{-6} . For smaller ε , non-convergence was observed.

AAM (Guminov et al., 2021). The implementation is based on NumPy code by the original authors at [this repository](#). A PyTorch version was verified to produce identical output for GPU execution. The ε parameter was varied from 2^{-1} to 2^{-10} . For smaller ε , numerical errors were encountered.

Feydy, Alg. 3.5 (Feydy, 2020). The implementation is based on the algorithm as presented in the original work. For different levels of precision, the number of total iterations was varied from 2 to 2^{12} . Beyond the upper bound, numerical errors were observed. As it produced better estimates than the alternative, the algorithm was called with *debiasing* turned on; hence, the error $\langle P - P^*, C \rangle$ was instead measured in absolute value as $|\langle P - P^*, C \rangle|$ for this algorithm only. Scaling ratio was set to an intermediate 0.7, which is between the listed 0.5 (fast) and 0.9 (safe) settings.

Sinkhorn (Cuturi, 2013). A log-domain stabilized implementation was used. For different precision levels, γ was varied from 2^5 to 2^{14} for L_1 distance cost and to 2^{15} for L_2^2 distance cost. Stopping criteria were given by our formula in L4 of Alg. 1, and the results obtained by calling Alg. 1 with $\gamma_i = \gamma_f$, so that the algorithm terminates after a single Bregman projection via SK iteration.

Mirror Sinkhorn (MSK) (Ballu & Berthet, 2023). The implementation is based on the algorithm presented in the original paper. For different levels of precision, the number of total iterations was varied from 2^5 to 2^{16} .

F Additional Benchmarking on DOTmark

Figs. 8-17 add further benchmarking on 10 more datasets from the DOTmark benchmark of Schrieber et al. (2017), which include various kinds of randomly generated images, classical test images and real data from microscopy. Each dataset contains 45 unique pairs of marginals (\mathbf{r}, \mathbf{c}) obtained from pixel values. The cost matrix is constructed from distances in 2D pixel locations; we evaluate on both L_1 and L_2 distance costs for $n = 4096$ following our setup in Sec. 5.

Besides clock time, we additionally plot here the total number of $O(n^2)$ -costing operations for each algorithm, e.g., matrix-vector products, row/column sums of matrices, vector outer products, element-wise operations on matrices. We count primitive operations for consistency across algorithms; counting a higher-level function call such as the number of gradient evaluations would be unfair due to inherent differences in the design of various algorithms. For instance, some require costly line search between gradient evaluations. These plots show that operation counts of the baseline algorithms follow similar trends to wall-clock time and no algorithm is unfairly advantaged via low-level optimizations.

For each of 20 problem sets (10 image datasets \times 2 cost functions), 20 out of 45 problems are sampled without replacement. The wall-clock time plots for the respective cost functions (L_1 and L_2^2) follow similar trends as Fig. 4. In addition to the median, we also include 75% confidence intervals along both axes, which show that MDOT is generally robust.

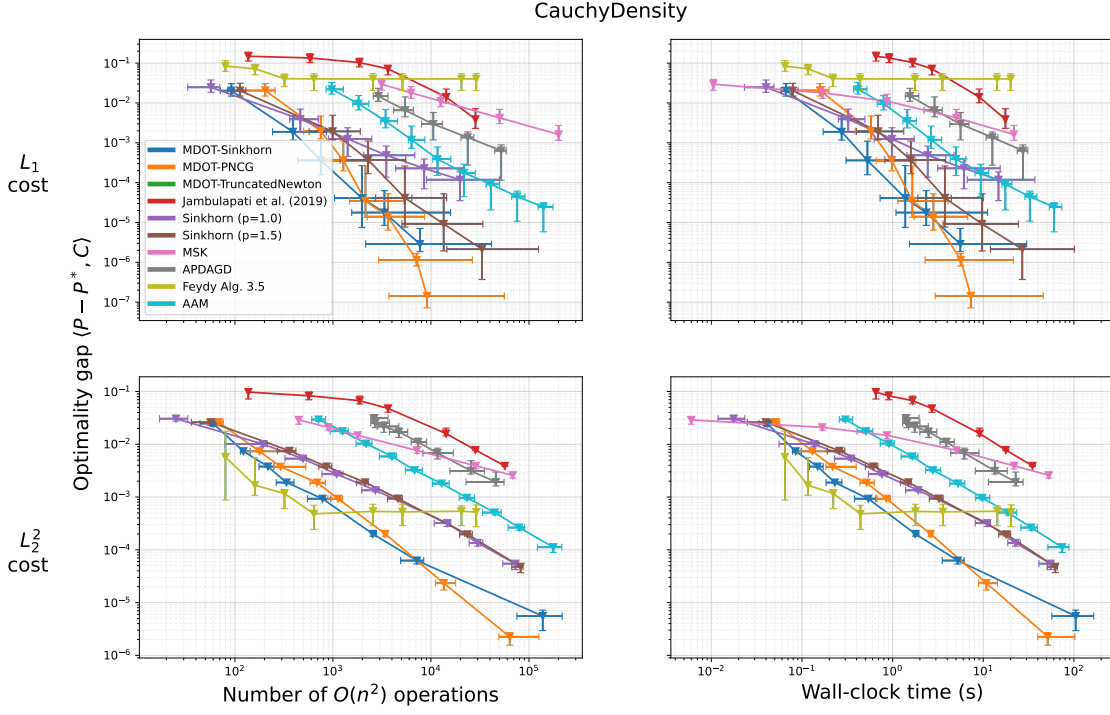


Figure 8: CauchyDensity problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

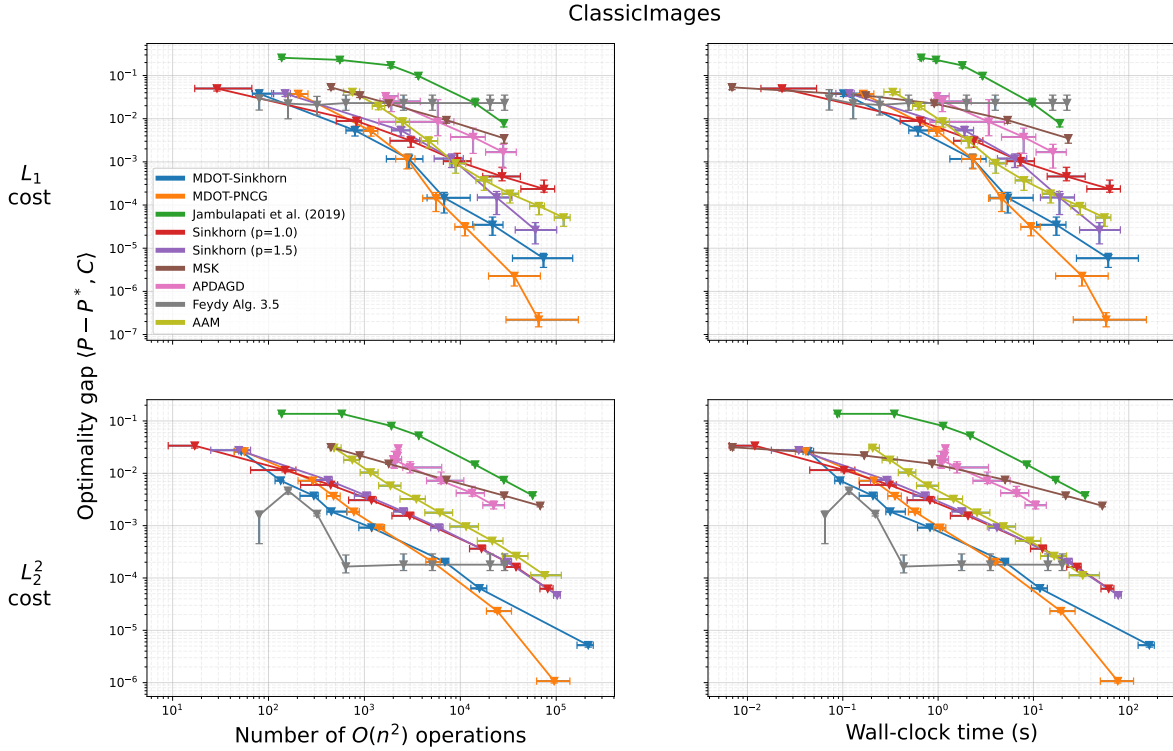


Figure 9: ClassicImage problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

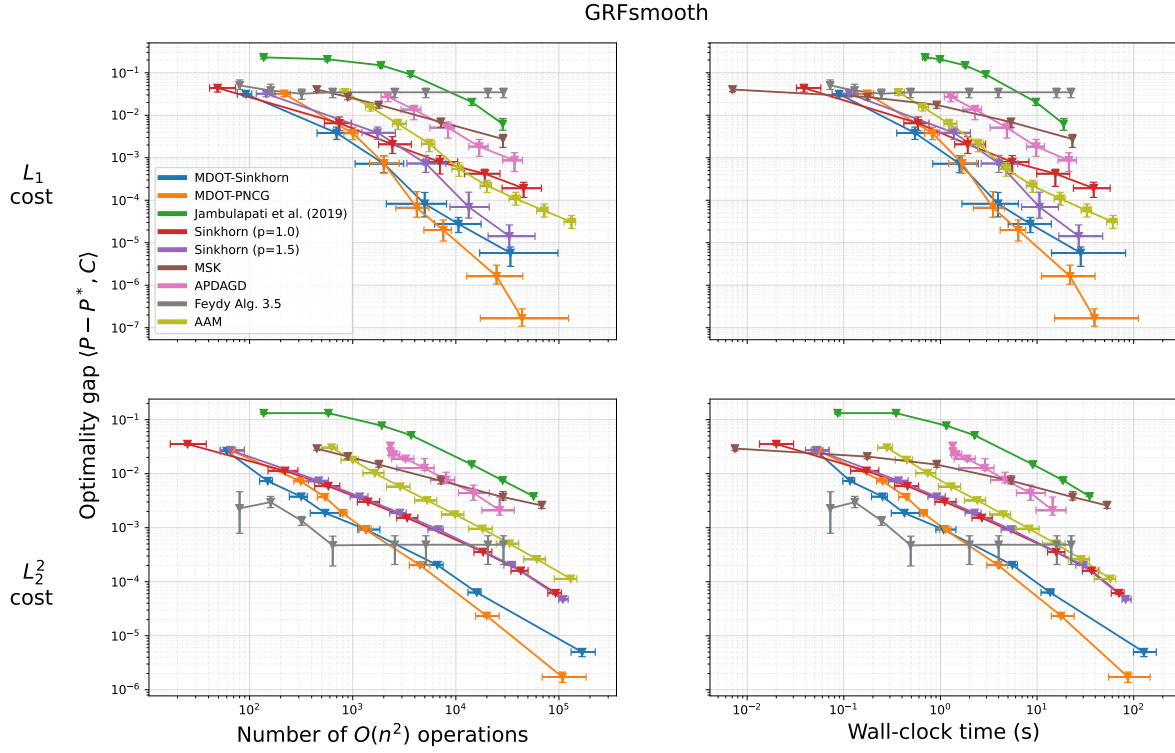


Figure 10: GRFsmooth problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

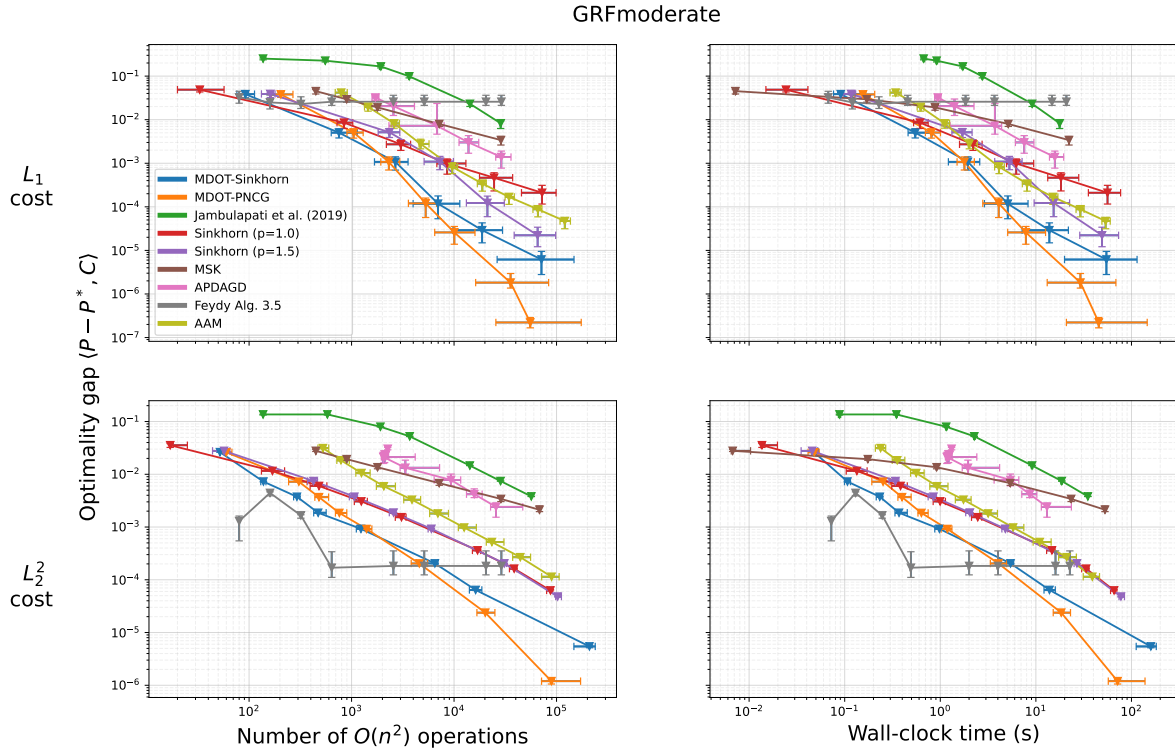


Figure 11: GRFmoderate problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

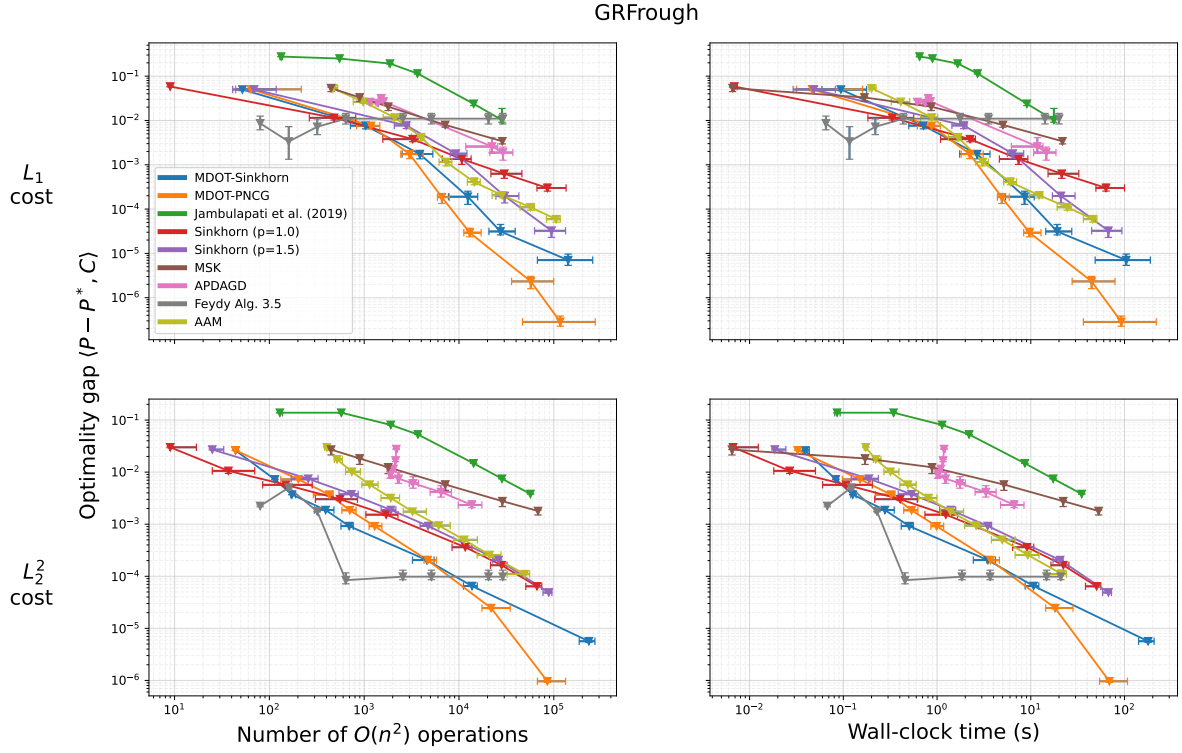


Figure 12: GRFrough problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

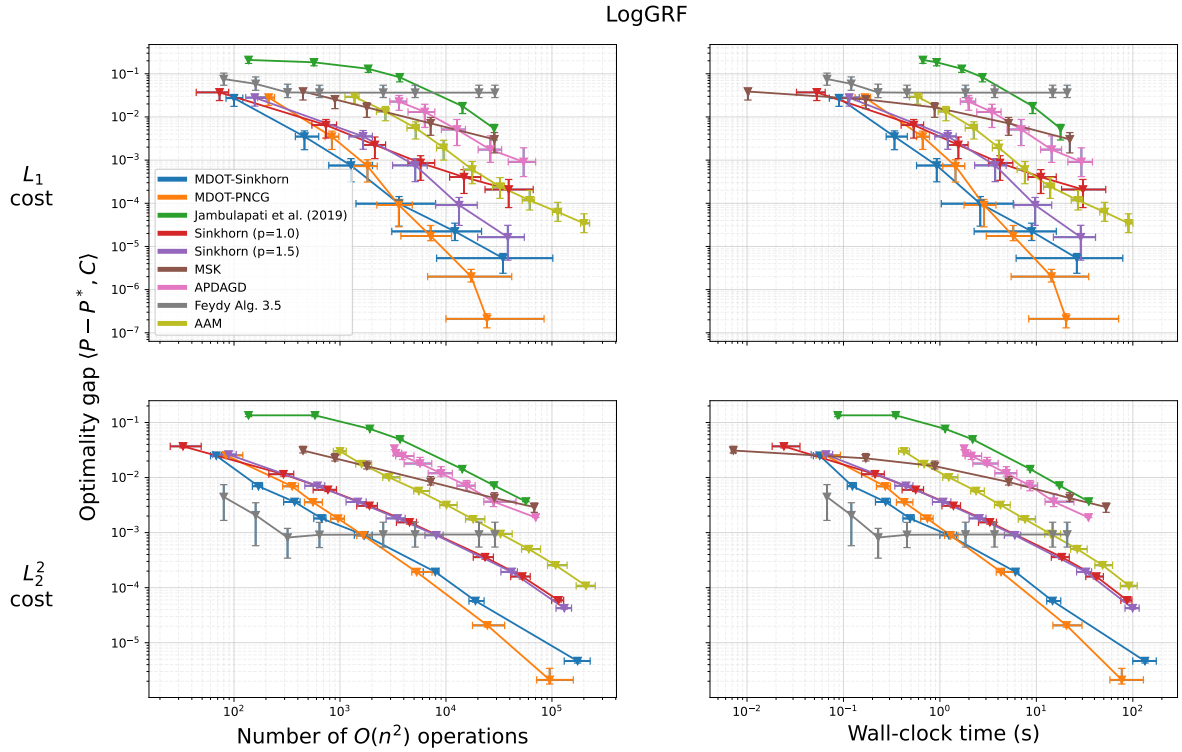


Figure 13: LogGRF problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

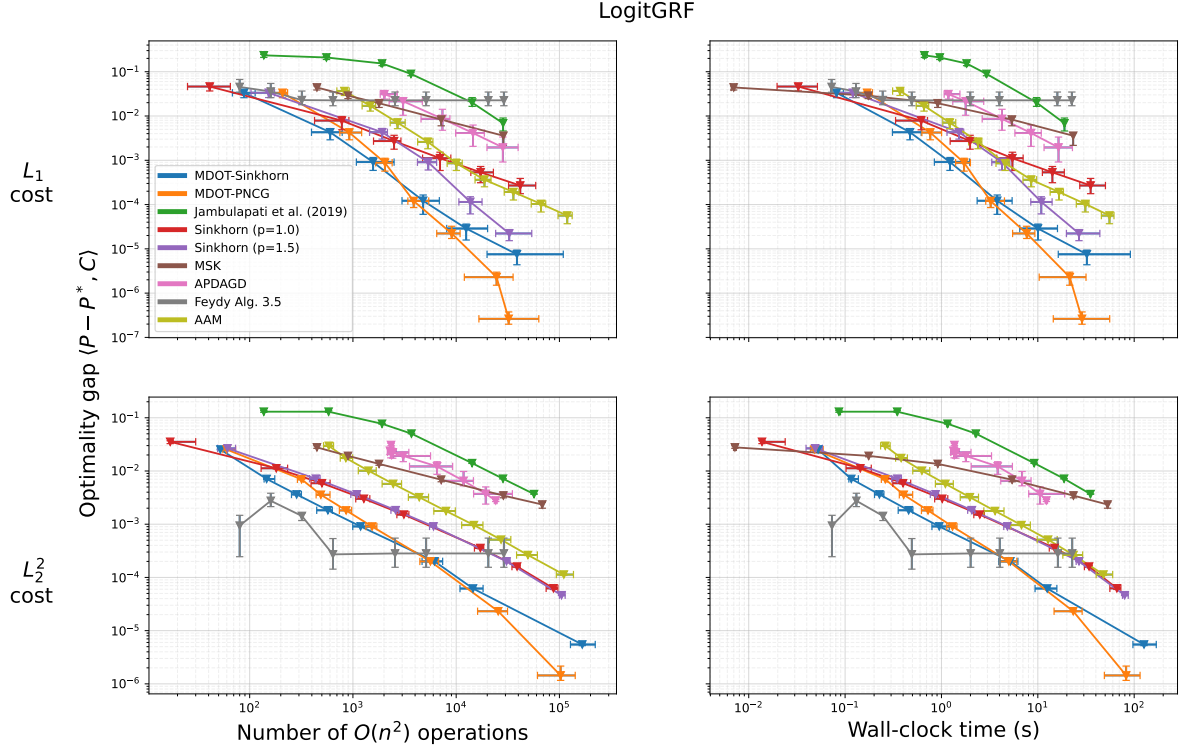


Figure 14: LogitGRF problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

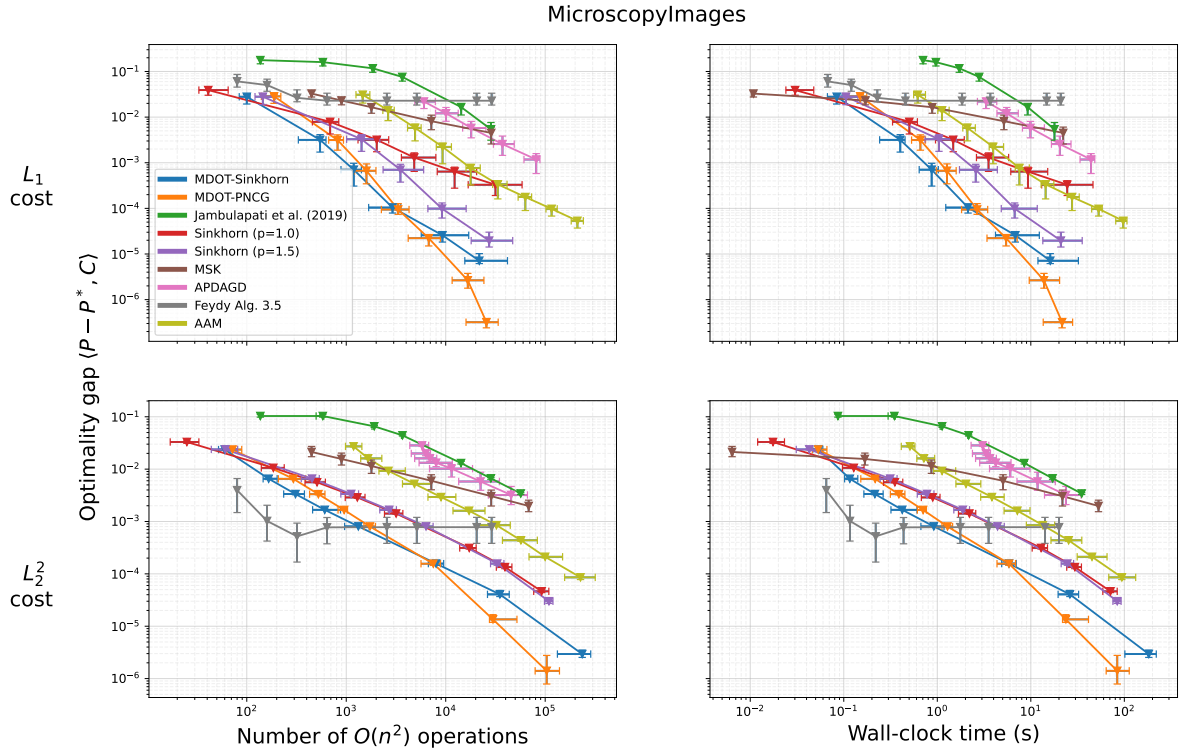


Figure 15: MicroscopyImage problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

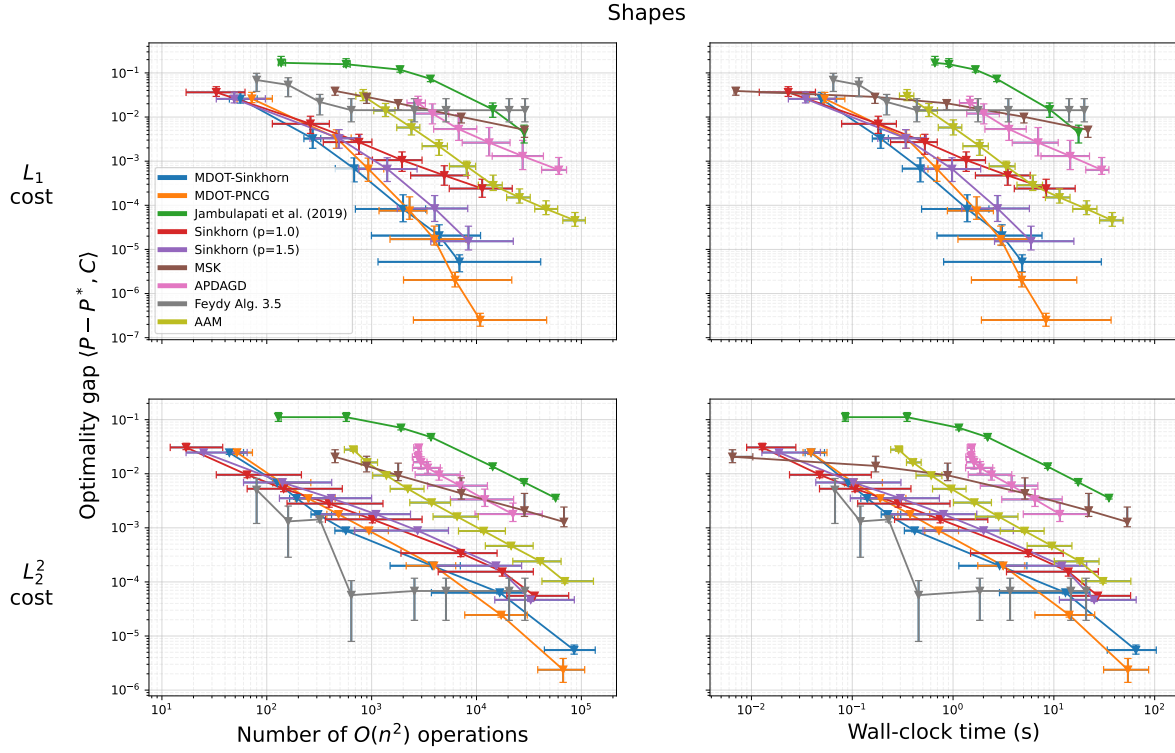


Figure 16: **Shape** problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).

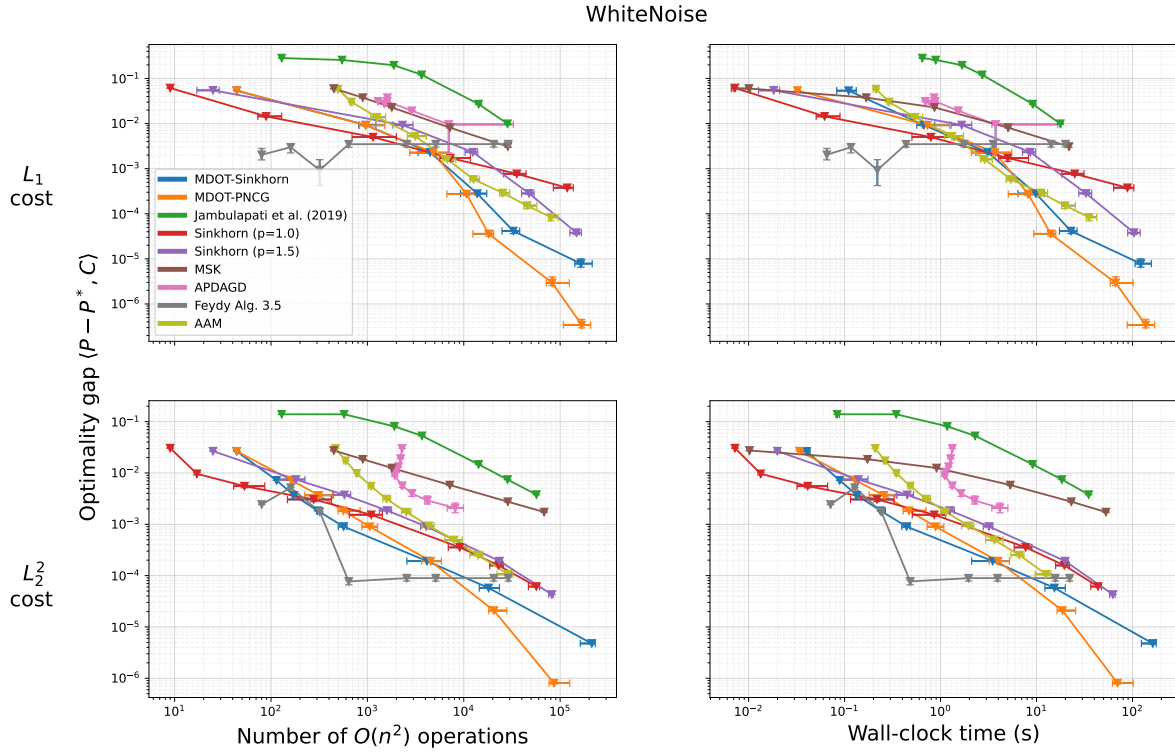


Figure 17: **WhiteNoise** problem with L_1 (top) and L_2^2 (bottom) costs, showing excess cost (error) vs. number of $O(n^2)$ operations (left) and wall-clock time (right).