DIFFERENTIALLY PRIVATE WASSERSTEIN BARYCENTERS

Anonymous authors
Paper under double-blind review

ABSTRACT

The Wasserstein barycenter is defined as the mean of a set of probability measures under the optimal transport metric, and has numerous applications spanning machine learning, statistics, and computer graphics. In practice these input measures are empirical distributions built from sensitive datasets, motivating a differentially private (DP) treatment. We present, to our knowledge, the first algorithms for computing Wasserstein barycenters under differential privacy. We use the central DP model for empirical distributions. We develop two complementary approaches targeting different regimes. Firstly, we consider *output perturbation*. We show that while this approach is competitive when the number of marginals is large, its privacy guarantee does not improve with the size of the dataset n. As a result, we propose as our primary contribution, *private coreset reduction*, where we construct private measures that are provably close in Wasserstein distance to the originals and run any Wasserstein barycenter on these. We also leverage dimensionality reduction to improve the runtime. Empirically, on synthetic data, MNIST, and large-scale U.S. population datasets, our methods produce high-quality private barycenters with strong accuracy-privacy tradeoffs.

1 INTRODUCTION

In the era of big data and machine learning, users are increasingly concerned about their privacy. Differential privacy (DP) (Dwork et al., 2006) has seen widespread adoption to provide guarantees for user privacy. For example, government bureaus use DP when releasing census data (Abowd, 2018; Hod & Canetti, 2024), and companies such as Apple (Apple, 2017), Microsoft (Ding et al., 2017), and LinkedIn (Rogers et al., 2020) extensively employ DP when releasing data — aiming to protect user data from security threats.

Clustering, summarizing and reducing the size of datasets are fundamental tasks in unsupervised machine learning. Many of these unsupervised learning problems are NP-hard (Megiddo & Supowit, 1984; Altschuler & Boix-Adserà, 2022), leading to the development of polynomial-time approximation algorithms (Charikar et al., 1999; Charikar & Guha, 1999; Jain & Vazirani, 2001; Jain et al., 2003; Charikar & Li, 2012; Cohen-Addad et al., 2022b). A long line of works (Gupta et al., 2010; Balcan et al., 2017; Kaplan & Stemmer, 2018; Jones et al., 2021; Chaturvedi et al., 2020; Ghazi et al., 2020) have further studied clustering under DP, providing polynomial time algorithms with tight approximation bounds.

Defined as the mean of a set of probability measures under the optimal transport metric,[1] the Wasserstein barycenter is a useful notion that contains many of these unsupervised tasks as special cases, with applications to a much more general suite of problems. Specific instances of the Wasserstein barycenter include centroids of probability measures (Zen & Ricci, 2011) and k-means clustering (Canas & Rosasco, 2012). Consequently, Wasserstein barycenters have seen extensive applications in domain adaptation (Montesuma & Mboula, 2021), computer graphics (Pele & Werman, 2009; Solomon et al., 2015), and biology (Nadeem et al., 2020; Heinemann et al., 2022).

[1]Specifically, it is the distribution that minimizes the average Wasserstein distance between itself and each distribution in the set, and generalizes the classical concept of the (Euclidean) mean from datapoints to entire distributions.

Similar to clustering and other unsupervised learning problems mentioned above, privatization of Wasserstein barycenters becomes crucial when working with sensitive data. As one of many possible examples, suppose a company wishes to train and deploy machine learning models for analysis of sensitive data for each of many countries. Prior to actually training these models, to minimize the risk of privacy breaches, a machine learning engineer must design the model architecture and tune hyperparameters on a DP synthetic dataset (see (Ridgeway et al., 2021; McKenna et al., 2021; Xie et al., 2018) for more details on this emerging practice), which by the post-processing property of DP can be used and re-used *ad infinitum* without incurring any additional privacy loss. A private Wasserstein barycenter of a (sub)set of country-level datasets would be a natural candidate for this private synthetic dataset, as (1) it averages across many countries and hence should incur much less privacy cost than maintaining separate private synthetic datasets for each country, and (2) it approximately minimizes the Wasserstein distances to the true distributions, which should maximize the chance that the designed model architecture will work well when applied to each countries' data at deployment time.

More classically, recall the Wasserstein barycenter minimizes the (weighted) average transport cost from the barycenter to each marginal. This can be an interesting optimization problem in its own right, e.g. choosing locations for distribution centers for multiple products each with its own geographic demand distribution, where each center has the same mix of products.

Motivated by these considerations, our work answers the following question in the affirmative:

Do there exist efficient[2] algorithms for computing Wasserstein barycenters under DP?

Contributions To the best of our knowledge, we provide the first algorithms for computing Wasserstein barycenters under the constraints of the central model (Dwork et al., 2006) of DP. We work under the setting where each individual contributes one datapoint to one distribution.

Our main contributions are as follows.

- We provide a lightweight (ϵ, δ)-DP algorithm that works when the number of support points is much smaller than the number of distributions; see Theorem 4.1. This method treats the output barycenter as a vector and uses the Gaussian mechanism to privatize it.
- We provide an efficient ϵ-DP algorithm using a black-box reduction from private Wasserstein distance coresets (Definition 8); see Theorem 5.5. Here, we form a private version of each distribution, and use these to solve the barycenter problem. Privacy is guaranteed by parallel composition, and we use the Johnson-Lindenstrauss to improve runtime.
- We show the efficacy of our more general algorithm when applied to large-scale real-world sensitive data; see Figure 3.

2 Preliminaries

Notation In big-O notation, we use $\tilde{O}$ to hide logarithmic factors and subscripts to hide dependence in those variables. We use $[t]$ to denote the set $\{1, \dots, t\}$. For a function $T : \mathcal{X} \to \mathcal{Y}$ and a measure μ, we use $T_\sharp \mu$ to denote the pushforward measure, e.g. $T_\sharp \mu(B) = \mu(T^{-1}(B))$ for a measurable set $B \subseteq \mathcal{Y}$. For a datapoint x, we use δ_x to denote a Dirac delta at x. We reserve the Greek letter ξ to denote a failure probability. We use $B_x(R) := \{y \mid \|x - y\|_2 \leq R\}$ to denote the closed Euclidean ball of radius R centered at point x. **(OT)** We use Greek letters μ for the dataset and ν for barycenter, which are probability measures. **(DP)** We use the script font $\mathcal{A}$ and $\mathcal{D}$ to denote algorithms and datasets, respectively. We reserve Greek letter ϵ for privacy parameters for DP. **(JL)** We use letters d to represent the dimension of the *ambient* space and d' to represent the dimension of the *projected* space from the JL transform. We use the Greek letter γ for the multiplicative factor in the JL and Π for the projection matrix.

Differential privacy (DP) (Dwork et al., 2006) is a mathematical framework for establishing guarantees on privacy loss of an algorithm, with nice properties such as degradation of privacy loss under composition and robustness to post-processing. We provide a brief introduction and refer to (Dwork & Roth, 2014) for a thorough treatment.

[2]Same (asymptotic) runtime as a non-private algorithm.

Definition 1 (ϵ-DP)**.** *Algorithm $\mathcal{A}$ is said to satisfy ϵ-differential privacy if for all adjacent datasets $\mathcal{D}, \mathcal{D}'$ (datasets differing in at most one element) and all $\mathcal{S} \subseteq \operatorname{range} \mathcal{A}$, it holds*

$$\Pr[\mathcal{A}(\mathcal{D}) \in \mathcal{S}] \leq e^{\epsilon} \Pr[\mathcal{A}(\mathcal{D}') \in \mathcal{S}].$$

Let $(\mathcal{X}, \rho)$ be a metric space and let $\mathcal{P}(\mathcal{X})$ be the set of Borel probability measures on $\mathcal{X}$.

Definition 2 (Wasserstein distance)**.** *For $p \in [1, \infty)$, the p-Wasserstein distance between probability measures $\mu, \nu \in \mathcal{P}(\mathcal{X})$ is defined to be*

$$W_p(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathcal{X} \times \mathcal{X}} \rho(x, y)^p d\pi(x, y) \right)^{1/p},$$

where $\Pi(\mu, \nu) := \{\pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \mid (P_x)_\sharp \pi = \mu, (P_y)_\sharp \pi = \nu\}$ is the set of transport plans, and $P_x(x, y) := x$ and $P_y(x, y) := y$ are the projections onto the first and second coordinates, respectively.

We will be using the Euclidean metric for the cost function, e.g. $\rho(x, y) := \|x - y\|_2$.

In Appendix C, we recall some additional facts on differential privacy, optimal transport, and the Johnson-Lindenstrauss transform.

3 Problem statement

(Agueh & Carlier, 2011) introduced the notion of barycenters on Wasserstein space:

Definition 3 (Wasserstein barycenter)**.** *Given probability distributions $\mu_1, \ldots, \mu_k \in \mathcal{P}(\mathcal{X})$ and weights $\beta_1, \ldots, \beta_k > 0$, the p-Wasserstein barycenter is any distribution ν^* satisfying*

$$\nu^* \in \operatorname*{arg\,min}_{\nu \in \mathcal{P}(\mathcal{X})} \sum_{i=1}^{k} \beta_i W_p^p(\mu_i, \nu). \tag{1}$$

We will be working with *discrete distributions*, where each distribution can be thought of as a subpopulation, and one individual contributes sensitive data to one of the distributions. For instance each of these distributions could represent the data from one country.

Formally, we have k empirical distributions μ_i for $i \in [k]$, each with n point masses, where

$$\mu_i = \frac{1}{n} \sum_{j=1}^{n} \delta_{x_j}, \tag{2}$$

Our goal is to compute a distribution ν consisting of exactly $m \leq n$ point masses and uniform weights that minimizes the objective (1) under the constraints of DP. For any application of DP, a definition of neighboring datasets is required. We use the following slight generalization of the most standard definition.

Definition 4 (Neighboring datasets)**.** *Let the dataset be $\mathcal{D} := \tilde{\mathcal{D}} \times [k]$, where $\tilde{\mathcal{D}} := \{x \mid x \in \cup_{i \in [k]} \operatorname{supp} \mu_i\}$. We say two datasets $\mathcal{D}, \mathcal{D}'$ are neighboring if they differ by exactly one row, e.g. they differ by one point in one distribution.*

Specifically, as is typical in DP, two datasets are neighbors if they differ in one row. This definition of neighboring datasets is motivated by viewing each μ_i is a nonoverlapping subpopulation, i.e. we are essentially assuming that individuals do not appear in multiple of the μ_i. Without loss of generality, we can assume all of the points are distinct. We will also abuse notation and identify $\mathcal{D}$ with $\{\mu_1, \ldots, \mu_k\}$.

In order to limit the influence of any individual, we require an assumption on the support. For simplicity, without loss of generality we assume a support contained in a ball of radius of 1/2.

Assumption 1. *It holds that $\cup_{i \in [k]} \operatorname{supp} \mu_i \subseteq B_0(1/2)$.*

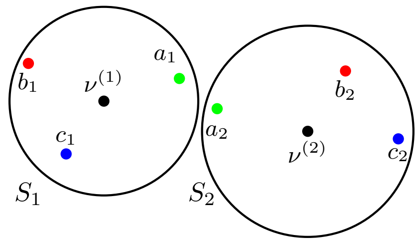

Figure 1: Example of a solution. The input distributions are $\mu_a := \frac{1}{2}\delta_{a_1} + \frac{1}{2}\delta_{a_2}, \mu_b := \frac{1}{2}\delta_{b_1} + \frac{1}{2}\delta_{b_2}, \mu_c := \frac{1}{2}\delta_{c_1} + \frac{1}{2}\delta_{c_2}$ and the candidate barycenter is $\nu := \frac{1}{2}\delta_{\nu^{(1)}} + \frac{1}{2}\delta_{\nu^{(2)}}$. Observe that: $S_1 = \{a_1, b_1, c_1\}, S_2 = \{a_2, b_2, c_2\}, w_1(a_1) = w_1(b_1) = w_1(c_1) = w_2(a_2) = w_2(b_2) = w_2(c_2) = 1$.

This is a standard assumption in private means and medians, e.g. see (Ghazi et al., 2020),[3] and more generally private convex optimization. This assumption is used to simplify the description of the results as it is well known that the additive error of any DP algorithm scales proportionally with respect to the radius of the support of the dataset, e.g. see (Altschuler et al., 2024).

(Izzo et al., 2021) utilizes the Johnson-Lindenstrauss transform to speed up algorithms for Wasserstein barycenters. We start with the following definition, adapted from Definition 2.1 of (Izzo et al., 2021).[4]

Definition 5 (Solution)**.** *Fix a candidate barycenter ν supported on points $\nu^{(1)}, \ldots, \nu^{(m)}$. Define the solution $(\mathbf{S}, \mathbf{w}) := (S_1, \ldots, S_m, w_1, \ldots, w_m)$ as follows. $w_j(x)$ is defined to be the total weight transported from $x \in \cup_{i=1}^k \operatorname{supp} \mu_i$ to point $\nu^{(j)}$ based on the optimal transport plan. Define the set*

$$S_j := \left\{ x \in \mathcal{D} \,\middle|\, w_j(x) > 0 \right\}.$$

We call $(\mathbf{S}, \mathbf{w})$ a solution because the following holds. For each $j \in [m]$, $\nu^{(j)}$ minimizes the objective

$$\sum_{x \in S_j} w_j(x) \|x - \nu^{(j)}\|^p. \tag{3}$$

See Figure 1 for intuition on the definition of this solution. Notice that if we are given the weights w_j, we can easily reconstruct the points $\nu^{(j)}$ using convex optimization. We obtain these weights by solving a corresponding Wasserstein barycenter in the reduced space using any approximation algorithm. Note that by conservation of mass, it holds that $\sum_{j=1}^n w_j(x) = 1$. We constrain the optimization objective in (1) as follows.

Assumption 2. *Assume that the objective (1) has an added constraint that the solution has m equally weighted atoms, where m is specified. Specifically, the solution satisfies $\nu = \frac{1}{m}\sum_{j=1}^m \delta_{\nu^{(j)}}$, where $m \leq n$, and $\beta_i = \frac{1}{k}$.*[5]

For ease of analysis, we assume n is divisible by m, which ensures that the optimal transport plans between μ_i and ν do not split mass. We make two remarks on the uniform weight assumption:

- From a *interpretability* perspective, uniform weights is a reasonable assumption so that each datapoint can be considered as data representing one synthetic person.
- From a *computational* perspective, many papers on Wasserstein barycenters *a priori* solve the problem under the uniform weight assumption, as optimizing weights for barycenters is much more challenging than optimizing for supports, e.g. see the discussion in (Altschuler & Boix-Adsera, 2021).

We will use the following cost function for Wasserstein barycenters.

[3](Ghazi et al., 2020) considers the ball of radius 1, while we consider the ball with diameter 1.

[4]Their definition is slightly different. Our definition is simplified to work better with Assumption 2.

[5]We remark that our second algorithm (Theorem 5.5) does not require any assumption on β_i.

Algorithm 1 WassersteinBarycenterViaOutputPerturbation$^{\epsilon,\delta}$

Require: k discrete distributions $\mu_1, \ldots, \mu_k$ supported on $\mathbb{R}^d$, approximate Wasserstein barycenter algorithm NonPrivateApprox, privacy parameter ϵ, δ

1: $\nu := (\nu^{(1)}, \ldots, \nu^{(m)}) \leftarrow \text{NonPrivateApprox}(\mu_1, \ldots, \mu_k)$

2: $(\tilde{\nu}^{(1)}, \ldots, \tilde{\nu}^{(m)}) \leftarrow (\nu^{(1)}, \ldots, \nu^{(m)}) + \mathcal{N}(0, \sigma^2 \cdot I_{md})$, where

$$\sigma^2 := \frac{2m\ln(1.25/\delta)}{(\epsilon k)^2}.$$

3: **return** $\tilde{\nu}$ with uniform support on $\{\tilde{\nu}^{(j)}\}_{j\in[m]}$

Definition 6 (Cost)**.** *For a solution* $(\mathbf{S}, \mathbf{w})$*, define its cost to be the value of* (1) *when* ν *is reconstructed from* $(\mathbf{S}, \mathbf{w})$*:*

$$\text{cost}(\mathbf{S}) := \min_{\nu} \frac{1}{nk} \sum_{j=1}^{n} \sum_{x\in S_j} w_j(x)\|x - \nu^{(j)}\|^p. \tag{4}$$

Similarly, for a projection Π*, define* $\text{cost}(\Pi_\sharp \mathbf{S})$ *to be the value of* (1) *when we first project each distribution to* $\mathbb{R}^{d'}$ *using* Π*, then compute* $\tilde{\nu}$ *using the original weights* w_j*:*

$$\text{cost}(\Pi_\sharp \mathbf{S}) := \min_{\nu} \frac{1}{nk} \sum_{j=1}^{n} \sum_{x\in S_j} w_j(x)\|\Pi x - \tilde{\nu}^{(j)}\|^p. \tag{5}$$

Above, note that we suppress the dependence on p for the cost. We use the following definition for approximate Wasserstein barycenters.

Definition 7 (Approximate Wasserstein barycenter)**.** *Let* OPT *be the minimum of* (1)*. A* (z,t)*-approximation for the* p*-Wasserstein barycenter is probability measure* ν *such that*

$$\text{cost}(\nu) \le z \cdot \text{OPT}_{(\mu_1,\ldots,\mu_k)} + t,$$

where OPT *is the cost of an optimal barycenter supported on* n *atoms with uniform weights. When it's clear, we suppress the dependence on the input barycenters.*

4 A NATURAL OUTPUT PERTURBATION APPROACH

A natural approach to the problem is to consider output perturbation; however, we show that this will only have good utility if $md \ll k^2$. The issue that inhibits an upper bound that benefits from increasing n is that in a neighboring dataset, the couplings of all n points in the updated distribution could potentially change, so we only obtain an averaging effect due to the k distributions, as opposed to $\frac{nk}{m}$ (number of points that are mapped to each point in the support of the barycenter); see Proposition D.3.

We provide the pseudocode in Algorithm 1. Its guarantees are as follows, and the proof is provided in Appendix E. Note that privacy benefits are derived by increasing the number of marginals k.

Theorem 4.1 (Output perturbation method)**.** *For any* $p \ge 1$*, suppose there exists a (not necessarily private)* (z,t)*-approximation algorithm for the* p*-Wasserstein barycenter problem. Then for every* $\epsilon > 0$*,* $\delta \in (0,1)$*, Algorithm 1 is* (ϵ,δ)*-DP, and yields a* $\left(z, t + O_p\left(\left(\frac{md\ln(1/\delta)}{(\epsilon k)^2}\right)^{p/2}\right)\right)$*-approximation for the* p*-Wasserstein barycenter problem.*

5 IMPROVING WITH n: A PRIVATE CORESET-BASED APPROACH

In this section, we achieve bounds that improve with data size n by first extracting private coresets from each full dataset that are close in Wasserstein distance to the sensitive distributions, yet have desirable privacy guarantees. We use these coresets as input to the approximate Wasserstein barycenter algorithm to obtain the private barycenter via data processing inequality.

5.1 Coreset for Wasserstein distance

We start by introducing the notion of coresets for Wasserstein distance.

Definition 8 (Coreset for Wasserstein distance)**.** *A measure μ' is a (p, z, t)-coreset of μ for the p-Wasserstein distance if for every $\pi \in \mathcal{P}(\mathbb{R}^d)$, we have $W_p(\mu', \pi) \le z \cdot W_p(\mu, \pi) + t$. When p is unambiguous, we drop the p.*

The following proposition is a direct consequence of the triangle inequality.

Proposition 5.1. *If $W_p(\mu, \mu') \le t$, then μ' is a $(p, 1, t)$-coreset of μ for the p-Wasserstein distance problem.*

Now our goal is to find a coreset for the p-Wasserstein distance problem. We use the algorithm from (He et al., 2023). Informally, the algorithm works as follows. First, obtain a hierarchical binary partition over the space of $\log \epsilon n$ levels. Use the (discrete) Laplace mechanism on each cell to compute the number of points in each cell, with noise calibrated to the level. Then, it suffices to choose points in each cell totaling the number of counts independently of the data. The set of all of these points becomes the private data. For a full description of the algorithm, see Algorithm 4 of (He et al., 2023).

We remark that there exists a corresponding with high probability (w.h.p.) algorithm that has the following guarantee.

Theorem 5.2. *For every $\epsilon > 0$ and $\xi \in (0, 1)$, there exists an ϵ-DP algorithm running in time $\tilde{O}(\epsilon d n)$ that with probability $1 - \xi$, outputs an $\left(p, 1, O_p\left(\left(\frac{1}{(\epsilon n)^{1/d}} \cdot \operatorname{poly}\log\left(\frac{1}{\xi}\right)\right)^{1/p}\right)\right)$-approximate coreset of size $O(n \log \epsilon n)$ for the Wasserstein distance.*

Proof. (He et al., 2023) provides an algorithm with guarantees in expectation for the W_1 distance. Due to Lemma C.5, this implies a similar guarantee for W_p distance. To obtain the w.h.p. algorithm, we run $O\left(\log \frac{1}{\xi}\right)$ trials of the algorithm and use the exponential mechanism (McSherry & Talwar, 2007) to choose the best one, e.g. see Appendix D of (Bassily et al., 2014) for an analogous argument. □

Note that we can downsample back down to n points in the private coreset with a small loss in approximation accuracy, e.g. see Figure 7 in the appendix.

Our key technical lemma is the following result to bound the error using Wasserstein distance coresets instead of the true distributions:

Lemma 5.3. *Let $\mu_1, \dots, \mu_k$ be discrete probability measures and suppose $\mu'_1, \dots, \mu'_k$ are $(p, 1, t)$-coresets for each μ_i, respectively. Then,*

$$\mathrm{OPT}_{(\mu'_1,\dots,\mu'_k)} \le \mathrm{OPT}_{(\mu_1,\dots,\mu_k)} + O_p(t^p).$$

Proof (sketch). This follows from Definition 8 and (4). See Appendix E for the full proof. □

5.2 Main result

(Izzo et al., 2021) generalized the breakthrough work of (Makarychev et al., 2019) to show that reducing to $O(\log n)$ dimension suffices to preserve the cost of p-Wasserstein distances for *all* solutions supported on at most n data points. Their main result is the following:

Theorem 5.4. *Let $\mu_1, \dots, \mu_k$ be discrete probability distributions on $\mathbb{R}^d$ such that $|\operatorname{supp} \mu_i| \le \operatorname{poly}(n)$ for all $i \in [k]$. Let $d \ge 1$, $\gamma, \xi \in (0, 1)$. Let $\Pi : \mathbb{R}^d \to \mathbb{R}^{d'}$ be an i.i.d. Gaussian JL map with $d' = O\left(\frac{p^4}{\gamma^2} \log \frac{n}{\gamma\xi}\right)$. Then, with probability $1 - \xi$, it holds that*

$$\mathrm{cost}(\mathbf{S}) \approx_{1+\gamma} \mathrm{cost}(\Pi_\sharp \mathbf{S})$$

for all solutions $(\mathbf{S}, \mathbf{w})$.

Algorithm 2 WassersteinBarycenter$^\epsilon$

Require: k discrete distributions $\mu_1, \ldots, \mu_k$ supported on $\mathbb{R}^d$, projection dimension d', approximate Wasserstein barycenter algorithm NonPrivateApprox, privacy parameter ϵ

1: Sample a JL transform $\Pi \in \mathbb{R}^{d \times d'}$
2: **for** $i \in [k]$ **do**
3: $\quad \mu'_i \leftarrow \text{WassersteinDistanceCoreset}^\epsilon(\mu_i)$ ▷ Algorithm 4 of (He et al., 2023)
4: $\quad \hat{\mu}_i \leftarrow \Pi_\sharp \mu'_i$
5: **end for**
6: $\hat{\nu} \leftarrow \text{NonPrivateApprox}(\hat{\mu}_1, \ldots, \hat{\mu}_k)$ ▷ $\hat{\nu} \in \mathbb{R}^{d'}$
7: $(\mathbf{S}, \mathbf{w}) \leftarrow \text{SolutionWeights}(\hat{\nu}, \hat{\mu}_1, \ldots, \hat{\mu}_k)$
8: $(\nu^{(1)}, \ldots, \nu^{(m)}) \leftarrow \text{SupportPoints}(\mu'_1, \ldots, \mu'_k, \mathbf{S}, \mathbf{w})$ ▷ $(\nu^{(1)}, \ldots, \nu^{(n)}) \in \mathbb{R}^d$
9: **return** ν with uniform support on $\{\nu^{(j)}\}_{j \in [m]}$

Above, for $\gamma \geq 0$, we use $a \approx_{1+\gamma} b$ to denote $\frac{1}{1+\gamma} \leq \frac{a}{b} \leq 1 + \gamma$. We briefly remark that it is also possible to use the fast JL transform using $d' = O\left(\frac{p^6}{\gamma^2} \log \frac{n}{\gamma \xi}\right)$; for details please refer to Appendix B of (Izzo et al., 2021).

Our main result is the following, whose proof we provide in Appendix E.

Theorem 5.5 (Private coreset method)**.** *For any $p \geq 1$, suppose that there exists a (not necessarily private) (z, t)-approximation algorithm that runs in time $2^{O(d)} \cdot \mathrm{poly}(n, k)$ for the p-Wasserstein barycenter problem. Then, for every $\epsilon > 0$ and $\gamma, \xi \in (0, 1)$, there exists a polynomial-time ϵ-DP algorithm that outputs an*

$$\left(z(1+\gamma), O_{p,\gamma,z}\left(\frac{1}{(\epsilon n)^{1/d}} \cdot \mathrm{poly}\log\left(\frac{k}{\xi}\right) + t\right)\right)$$

-approximate p-Wasserstein barycenter, with probability $1 - \xi$.

Remark 1 (Privacy amplification)**.** *Note that we can amplify the privacy using subsampling (Balle et al., 2018), which will improve the dependence on ϵ.*

Remark 2 (Comparison to output perturbation)**.** *While the output perturbation approach of Section 4 focused on gaining privacy through larger k, the private coreset method obtains privacy through reducing individual sensitivity as n increases. By dropping poly-log factors on k, ξ, δ, we see that if $k \lesssim n^{\frac{1}{pd}} \epsilon^{-1+\frac{1}{pd}} \sqrt{md}$, then the private coreset method yields lower error than the output perturbation approach.*

As a result, the private coreset approach should be the better choice in more (but not all) real world settings.

Via dimensionality reduction, we can afford an algorithm that has an *exponential dependence* on the dimension as $d' = O(\log n)$. Unfortunately, many state of the art additive approximation algorithms

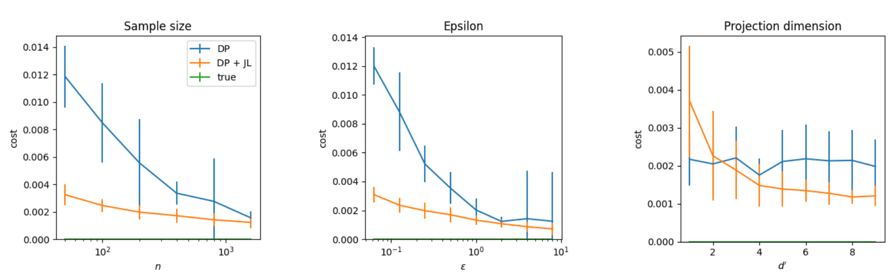

(a) $n = 50, 100, \ldots, 1600$ (b) $\epsilon = 2^{-4}, 2^{-2}, \ldots, 2^3$ (c) $d' = 1, 2, \ldots, 9$

Figure 2: Synthetic experiments testing sample size n, privacy parameter ϵ, and projection dimension d', averaged over 30 runs.

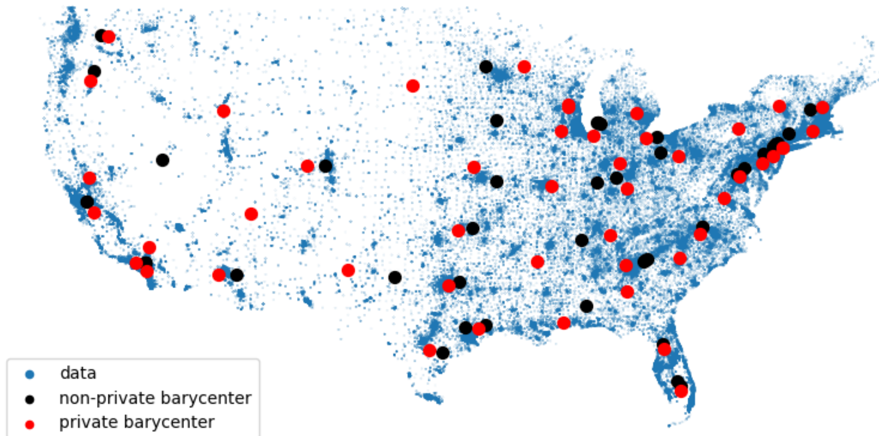

(a) $n = 200000$ and $\epsilon = 1$ for $m = 48$ and $k = 1$. Denoting ν, ν_ϵ as the non-private and private barycenters, respectively, we have $\mathrm{cost}(\nu) = 15.92$, $\mathrm{cost}(\nu_\epsilon) = 21.62$ (squared degrees longitude/latitude), and $W_2(\nu, \nu_\epsilon) = 5.633$ (degrees).

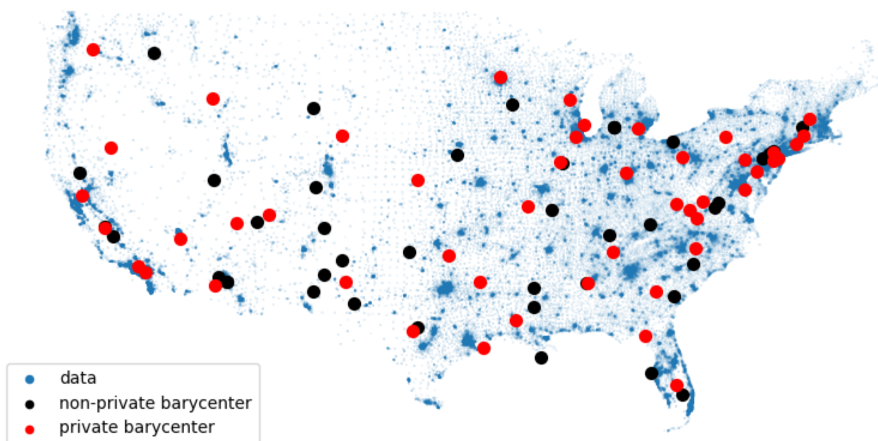

(b) $n = 100000$ and $\epsilon = 1$ for $m = 48$ and $k = 4$ (self-reported White, Asian, Black, Hispanic). Denoting ν, ν_ϵ as the non-private and private barycenters, respectively, we have $\mathrm{cost}(\nu) = 5058.416$, $\mathrm{cost}(\nu_\epsilon) = 6842.462$ (squared degrees) and $W_2(\nu, \nu_\epsilon) = 11.766$ (degrees).

Figure 3: Barycenters on continental US populations.

still do not lend polynomial runtime when combined with dimensionality reduction. For instance, the algorithm of (Altschuler & Boix-Adsera, 2021) runs in time $(nk)^{O(d)}$.

The weights $(\mathbf{S}, \mathbf{w})$ from Definition 5 are computed via optimal transport plans between $\hat{\mu}_i$ and $\hat{\nu}$, e.g. the distributions in low dimension. Due to post-processing, the $\hat{\mu}_i$ are private, so the computation incurs no additional privacy loss. We provide pseudocode in Algorithm 3. To recover the support points, we use empirical risk minimization (Algorithm 4).

6 EXPERIMENTS

We test our method from Section 5 on simple synthetic data, MNIST, and US population data. We provide additional experiments (on MNIST) and additional discussion of all experiment setups in Appendix F. All of our experiments use the Sinkhorn free support barycenter (Flamary et al., 2021) with 50 iterations and 100 inner (Sinkhorn) iterations. In our experiments, we utilize subsampling of the private coresets, which increases cost by a negligible amount, but significantly improves the runtime (see Figure 7 on MNIST in Appendix F).

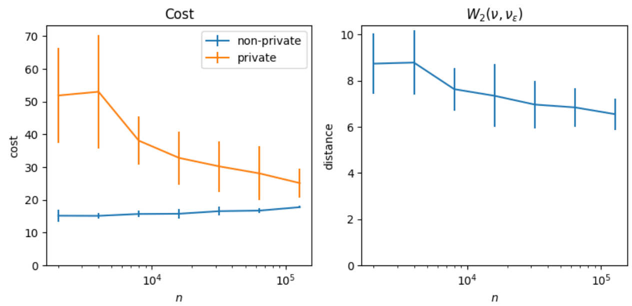

Figure 4: $n = 2000, 4000, \ldots, 128000$ and $\epsilon = 1$ in the same experimental setup as Figure 3a, averaged over 10 trials. On the left, we have cost in squared degrees. On the right, we plot the 2-Wasserstein distance between the private and non-private barycenters (in degrees).

In Figure 2, we consider equally weighted mixtures of 4 Gaussians at $(\pm 0.25, \pm 0.25, 0, \cdots, 0) \in \mathbb{R}^{10}$. We use $m = 8$ and 0.04 for the entropic regularization. For the synthetic experiments, we fix $n = 1000$, $\epsilon = 1$, and $d' = 5$ (when they are not varied). We generally observe that using the JL algorithm provides better utility (for small n or ϵ) under DP because without JL, the algorithm tends to get stuck in local minima, e.g. all points centered at the origin.

In Figure 3, we consider the experiment setup from (Cuturi & Doucet, 2014) ($m = 48$, $k = 1$) and use US population data from the American Community Survey (ACS) 2015.[6] In Figure 3a, we take $k = 1$, where the dataset is the whole US population. We take the (sensitive) data to be multisets of the centers of census tracts (chosen with replacement) of size $n = 200000$. In Figure 3b, we take the sensitive data to be each of $n = 100000$ (uniformly at randomly chosen) points corresponding to the self-reported racial groups White, Asian, Black, Hispanic, where for privacy, we assume the groups are disjoint.

In our private coreset construction, we only sample points that are inside the US border (which is private by data independent post-processing). Our algorithms use $\epsilon = 1$ on the *full* population (or subpopulation), utilizing privacy amplification by subsampling (Balle et al., 2018). Each barycenter computation only take a few minutes to run on CPU; however, the sampling of points inside the US takes a few hours for the largest experiments. In this experiment, we use entropic regularization of 0.001 and do not utilize dimensionality reduction.

In our experiments on the US population, we observe that the cost of privacy is less pronounced when $k \geq 2$. In particular, we find that the costs for non-private and private versions of the $k = 4$ experiment to stay relatively stable across varying values of n. Thus, in Figure 4, we report results for the $k = 1$ setting for the US population experiments.

7 Conclusion

We extended the study of private facility location problems from clustering to Wasserstein barycenters. One limitation of Algorithm 2 is the curse of dimensionality, and future work can study the setting where the data lies near a low dimensional subspace (Weed & Bach, 2019) and alleviate the curse of dimensionality via privatized versions of entropic OT (Mena & Niles-Weed, 2019; Genevay et al., 2019) or Gaussian-smoothed OT (Goldfeld et al., 2020; Goldfeld & Greenewald, 2020; Nietert et al., 2021; Zhang et al., 2021).

In this work, we studied Wasserstein barycenters under the central model of differential privacy, for future work it would also be interesting to obtain results under the local (Kasiviswanathan et al., 2011) and shuffle (Bittau et al., 2017; Erlingsson et al., 2019; Cheu et al., 2019) models. Furthermore, we focused on the setting where one individual contributes a single datapoint out of the k distributions. An interesting direction would be to consider the setting where one person contributes a whole probability measure, as this would allow practitioners to consider *continuous* distributions as input data.

[6] https://www.census.gov/acs/www/data/data-tables-and-tools/data-profiles/2015/

REFERENCES

John M Abowd. The US census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2867–2867, 2018.

Pankaj K. Agarwal, Sharath Raghvendra, Pouyan Shirzadian, and Keegan Yao. Fast and accurate approximations of the optimal transport in semi-discrete and discrete settings. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 4514–4529, 2024. doi: 10.1137/1.9781611977912.159. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611977912.159.

Pankaj K. Agarwal, Sharath Raghvendra, Pouyan Shirzadian, and Keegan Yao. Efficient approximation algorithm for computing Wasserstein barycenter under Euclidean metric. Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2025. doi: 10.1137/1.9781611978322.163. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611978322.163.

Martial Agueh and Guillaume Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.

Noga Alon, Phillip B Gibbons, Yossi Matias, and Mario Szegedy. Tracking join and self-join sizes in limited storage. In *Proceedings of the eighteenth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 10–20, 1999.

Jason Altschuler, Jonathan Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration, 2018. URL https://arxiv.org/abs/1705.09634.

Jason M Altschuler and Enric Boix-Adsera. Wasserstein barycenters can be computed in polynomial time in fixed dimension. *Journal of Machine Learning Research*, 22(44):1–19, 2021. URL http://jmlr.org/papers/v22/20-588.html.

Jason M. Altschuler and Enric Boix-Adserà. Wasserstein barycenters are NP-hard to compute. volume 4, pp. 179–203, 2022. doi: 10.1137/21M1390062. URL https://doi.org/10.1137/21M1390062.

Jason M. Altschuler, Jinho Bok, and Kunal Talwar. On the privacy of noisy stochastic gradient descent for convex optimization. *SIAM Journal on Computing*, 53(4):969–1001, 2024. doi: 10.1137/23M1556538. URL https://doi.org/10.1137/23M1556538.

Apple. Learning with privacy at scale. 2017. URL https://machinelearning.apple.com/research/learning-with-privacy-at-scale.

Maria-Florina Balcan, Travis Dick, Yingyu Liang, Wenlong Mou, and Hongyang Zhang. Differentially private clustering in high-dimensional Euclidean spaces. In *International Conference on Machine Learning*, pp. 322–331. PMLR, 2017.

Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/3b5020bb891119b9f5130f1fea9bd773-Paper.pdf.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th annual symposium on foundations of computer science*, pp. 464–473. IEEE, 2014.

Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative Bregman projections for regularized transportation problems. volume 37, pp. A1111–A1138, 2015. doi: 10.1137/141000439. URL https://doi.org/10.1137/141000439.

Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th symposium on operating systems principles*, pp. 441–459, 2017.

Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The Johnson-Lindenstrauss transform itself preserves differential privacy, 2012. URL https://arxiv.org/abs/1204.2136.

March Boedihardjo, Thomas Strohmer, and Roman Vershynin. Private measures, random walks, and synthetic data. *Probability theory and related fields*, pp. 1–43, 2024.

Stéphane Boucheron, Gábor Lugosi, and Olivier Bousquet. Concentration inequalities. In *Summer school on machine learning*, pp. 208–240. Springer, 2003.

Guillermo Canas and Lorenzo Rosasco. Learning probability measures with respect to optimal transport metrics. *Advances in Neural Information Processing Systems*, 25, 2012.

Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for the facility location and k-median problems. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pp. 378–388. IEEE, 1999.

Moses Charikar and Shi Li. A dependent LP-rounding approach for the k-median problem. In *International Colloquium on Automata, Languages, and Programming*, pp. 194–205. Springer, 2012.

Moses Charikar and Erik Waingarten. The johnson-lindenstrauss lemma for clustering and subspace approximation: From coresets to dimension reduction. pp. 3172–3209. Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, 2025. doi: 10.1137/1.9781611978322.102. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611978322.102.

Moses Charikar, Sudipto Guha, Éva Tardos, and David B Shmoys. A constant-factor approximation algorithm for the k-median problem. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pp. 1–10, 1999.

Anamay Chaturvedi, Huy Nguyen, and Eric Xu. Differentially private k-means clustering via exponential mechanism and max cover. *arXiv preprint arXiv:2009.01220*, 2020.

Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Advances in Cryptology–EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part I 38*, pp. 375–403. Springer, 2019.

Sebastian Claici, Edward Chien, and Justin Solomon. Stochastic Wasserstein barycenters, 2018. URL https://arxiv.org/abs/1802.05757.

Vincent Cohen-Addad, Hossein Esfandiari, Vahab Mirrokni, and Shyam Narayanan. Improved approximations for Euclidean k-means and k-median, via nested quasi-independent sets. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1621–1628, 2022a.

Vincent Cohen-Addad, Anupam Gupta, Lunjia Hu, Hoon Oh, and David Saulpic. An improved local search algorithm for k-median. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1556–1612. SIAM, 2022b.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/af21d0c97db2e27e13572cbf59eb343d-Paper.pdf.

Marco Cuturi and Arnaud Doucet. Fast computation of Wasserstein barycenters. In Eric P. Xing and Tony Jebara (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 685–693, Bejing, China, 22–24 Jun 2014. PMLR. URL https://proceedings.mlr.press/v32/cuturi14.html.

Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(34):211–407, 2014. ISSN 1551-305X. doi: 10.1561/0400000042. URL http://dx.doi.org/10.1561/0400000042.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pp. 265–284. Springer, 2006.

Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2468–2479. SIAM, 2019.

Vitaly Feldman, Audra McMillan, Satchit Sivakumar, and Kunal Talwar. Instance-optimal private density estimation in the wasserstein distance. *Advances in Neural Information Processing Systems*, 37:90061–90131, 2024.

Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78): 1–8, 2021. URL http://jmlr.org/papers/v22/20-451.html.

Aude Genevay, Lénaic Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of sinkhorn divergences. In *The 22nd international conference on artificial intelligence and statistics*, pp. 1574–1583. PMLR, 2019.

Badih Ghazi, Ravi Kumar, and Pasin Manurangsi. Differentially private clustering: Tight approximation ratios. *Advances in Neural Information Processing Systems*, 33:4040–4054, 2020.

Ziv Goldfeld and Kristjan Greenewald. Gaussian-smoothed optimal transport: Metric structure and statistical efficiency. In *International Conference on Artificial Intelligence and Statistics*, pp. 3327–3337. PMLR, 2020.

Ziv Goldfeld, Kristjan Greenewald, Jonathan Niles-Weed, and Yury Polyanskiy. Convergence of smoothed empirical measures with applications to entropy estimation. *IEEE Transactions on Information Theory*, 66(7):4368–4391, 2020.

Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the 2010 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1106–1125, 2010. doi: 10.1137/1.9781611973075.90. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611973075.90.

Yiyun He, Roman Vershynin, and Yizhe Zhu. Algorithmically effective differentially private synthetic data. In *The Thirty Sixth Annual Conference on Learning Theory*, pp. 3941–3968. PMLR, 2023.

Florian Heinemann, Axel Munk, and Yoav Zemel. Randomized wasserstein barycenter computation: Resampling with statistical guarantees. volume 4, pp. 229–259, 2022. doi: 10.1137/20M1385263. URL https://doi.org/10.1137/20M1385263.

Shlomi Hod and Ran Canetti. Differentially private release of Israel's national registry of live births. *arXiv preprint arXiv:2405.00267*, 2024.

Zachary Izzo, Sandeep Silwal, and Samson Zhou. Dimensionality reduction for Wasserstein barycenter, 2021. URL https://arxiv.org/abs/2110.08991.

Kamal Jain and Vijay V Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM (JACM)*, 48(2):274–296, 2001.

Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM (JACM)*, 50(6):795–824, 2003.

William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into Hilbert space. *Contemporary mathematics*, 26:189–206, 1984.

Matthew Jones, Huy L Nguyen, and Thy D Nguyen. Differentially private clustering via maximum coverage. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11555–11563, 2021.

Haim Kaplan and Uri Stemmer. Differentially private k-means with constant multiplicative error. *arXiv preprint arXiv:1804.08001*, 2018.

Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011. doi: 10.1137/090756090. URL https://doi.org/10.1137/090756090.

Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of statistics*, pp. 1302–1338, 2000.

Stuart Lloyd. Least squares quantization in PCM. *IEEE transactions on information theory*, 28(2): 129–137, 1982.

Konstantin Makarychev, Yury Makarychev, and Ilya Razenshteyn. Performance of Johnson-Lindenstrauss transform for k-means and k-medians clustering. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1027–1038, 2019.

Ryan McKenna, Gerome Miklau, and Daniel Sheldon. Winning the nist contest: A scalable and general approach to differentially private synthetic data. *arXiv preprint arXiv:2108.04978*, 2021.

Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pp. 94–103, 2007. doi: 10.1109/FOCS.2007.66.

Nimrod Megiddo and Kenneth J Supowit. On the complexity of some common geometric location problems. *SIAM Journal on Computing*, 13(1):182–196, 1984.

Gonzalo Mena and Jonathan Niles-Weed. Statistical bounds for entropic optimal transport: sample complexity and the central limit theorem. volume 32, 2019.

Eduardo Fernandes Montesuma and Fred Maurice Ngole Mboula. Wasserstein barycenter for multi-source domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16785–16793, 2021.

Shanmugavelayutham Muthukrishnan et al. Data streams: Algorithms and applications. *Foundations and Trends® in Theoretical Computer Science*, 1(2):117–236, 2005.

Saad Nadeem, Travis Hollmann, and Allen Tannenbaum. Multimarginal Wasserstein barycenter for stain normalization and augmentation. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pp. 362–371, Cham, 2020. Springer International Publishing. ISBN 978-3-030-59722-1.

Shyam Narayanan, Sandeep Silwal, Piotr Indyk, and Or Zamir. Randomized dimensionality reduction for facility location and single-linkage clustering. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7948–7957. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/narayanan21b.html.

Sloan Nietert, Ziv Goldfeld, and Kengo Kato. Smooth p-wasserstein distance: structure, empirical approximation, and statistical applications. In *International Conference on Machine Learning*, pp. 8172–8183. PMLR, 2021.

Aleksandar Nikolov. Private query release via the Johnson-Lindenstrauss transform, 2022. URL https://arxiv.org/abs/2208.07410.

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.

Diane Ridgeway, Diane Ridgeway, Mary F Theofanos, Terese W Manley, and Christine Task. *Challenge design and lessons learned from the 2018 differential privacy challenges*. US Department of Commerce, National Institute of Standards and Technology, 2021.

Ryan Rogers, Subbu Subramaniam, Sean Peng, David Durfee, Seunghyun Lee, Santosh Kumar Kancha, Shraddha Sahay, and Parvez Ahammad. Linkedin's audience engagements API: A privacy preserving data analytics system at scale. *arXiv preprint arXiv:2002.05839*, 2020.

Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.

Justin Solomon, Fernando De Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (ToG)*, 34(4):1–11, 2015.

Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620 – 2648, 2019. doi: 10.3150/18-BEJ1065. URL https://doi.org/10.3150/18-BEJ1065.

Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.

Gloria Zen and Elisa Ricci. Earth mover's prototypes: A convex learning approach for discovering activity patterns in dynamic scenes. In *CVPR 2011*, pp. 3225–3232. IEEE, 2011.

Yixing Zhang, Xiuyuan Cheng, and Galen Reeves. Convergence of gaussian-smoothed optimal transport distance with sub-gamma distributions and dependent samples. In *International Conference on Artificial Intelligence and Statistics*, pp. 2422–2430. PMLR, 2021.

A Related work

In the theoretical computer science community, the Wasserstein barycenter falls under the category of facility location problems. This class of problem is concerned with placing points, or "facilities," to minimize some objective given a set of input data. Note that clustering also falls under this category. Clustering (Lloyd, 1982) has seen many non-private approximation algorithms. Over the past few decades, a line of works (Charikar et al., 1999; Charikar & Guha, 1999; Jain & Vazirani, 2001; Jain et al., 2003; Charikar & Li, 2012; Cohen-Addad et al., 2022b) have pushed multiplicative approximation factors to 2.406 and 5.912 for Euclidean k-medians and k-means, respectively (Cohen-Addad et al., 2022a).

(Gupta et al., 2010) initiated the study of facility location algorithms under DP, and provided an inefficient algorithm based on the exponential mechanism (McSherry & Talwar, 2007) that gave constant factor multiplicative approximation. Then a series of works (Balcan et al., 2017; Kaplan & Stemmer, 2018; Jones et al., 2021; Chaturvedi et al., 2020; Ghazi et al., 2020) culminated in polynomial time algorithms for private clustering with the optimal multiplicative approximation ratio and small additive errors.

On the other hand, the Wasserstein barycenter is a much more nascent problem. Initial works provide approximations using methods such as entropic regularization (Cuturi, 2013; Cuturi & Doucet, 2014), iterative Bregman projections (Benamou et al., 2015), or stochastic optimization (Claici et al., 2018); however, these lack worst-case guarantees on the approximations, for instance to the non-entropic setting. Even theoretical guarantees for fast approximations of Wasserstein distances are recent (Altschuler et al., 2018; Agarwal et al., 2024). More recently, some works have provided theoretical guarantees for Wasserstein barycenters in the $p = 2$ setting: (Altschuler & Boix-Adsera, 2021; Agarwal et al., 2025) showed that additive and multiplicative (respectively) approximations for Wasserstein barycenters can be computed in polynomial time for *constant* dimension. Recently (Boedihardjo et al., 2024; He et al., 2023) provided constructions for private measures that are close to input empirical measures over $[0, 1]^d$ in 1-Wasserstein distance and (Feldman et al., 2024) provided instance-optimal constructions for finite metric spaces.

The Johnson-Lindenstrauss (JL) lemma (Johnson & Lindenstrauss, 1984) is a dimensionality reduction method that provides worst-case guarantees on preserving pairwise distances between a collection of points. It has been applied to numerous problems in many areas of computer science, including streaming algorithms (Alon et al., 1999; Muthukrishnan et al., 2005) and DP (Blocki et al., 2012; Nikolov, 2022). The (lower) bound on the dimension required to approximately preserve solutions varies from problem to problem, e.g. see (Narayanan et al., 2021; Charikar & Waingarten, 2025) for a discussion. For facility location problems, (Makarychev et al., 2019) showed that dimension $d' = O(\log k)$ suffices for preserving the cost of solutions to k-means clustering, and (Izzo et al., 2021) showed that dimension $d' = O(\log n)$ suffices for Wasserstein barycenters supported on $\leq n$ points.

B Deferred algorithms

Algorithm 3 SolutionWeights

Require: barycenter ν, k input distributions $\mu_1, \ldots, \mu_k$ supported on $m_1, \ldots, m_k$ points, respectively

1: Obtain $(\mathbf{S}, \mathbf{w})$ based on $(T_1, \ldots, T_k)$ as follows:
2: **for** $i \in [k]$ **do**
3:     $T_i \leftarrow \mathsf{OT}(\mu_i, \nu)$
4:     **for** $(\ell, j) \in [m_i] \times [m]$ **do**
5:         **if** $T_i[\ell, j] > 0$ **then**
6:             $S_j \leftarrow S_j \cup \{\mu_i[\ell]\}$
7:             $w_j(\mu_i[\ell]) \leftarrow T_i[\ell, j]$
8:         **end if**
9:     **end for**
10: **end for**
11: **return** $(\mathbf{S}, \mathbf{w})$

Algorithm 4 SupportPoints

Require: k discrete distributions $\mu_1, \ldots, \mu_k$ supported on $\mathbb{R}^d$, partition $(\mathbf{S}, \mathbf{w})$ as described in Definition 5

1: **for** $S_j \in \mathbf{S}$ **do**
2: $\quad \nu^{(j)} \leftarrow \arg\min \sum_{x \in S_j} w_j(x) \|x - \nu^{(j)}\|^p$
3: **end for**
4: **return** $(\nu^{(1)}, \ldots, \nu^{(m)})$

C Additional preliminaries and lemmata

C.1 Differential privacy

Lemma C.1 (Parallel composition). *Let $\mathcal{A}_1, \ldots, \mathcal{A}_k$ be ϵ-DP algorithms. Suppose $\mathcal{D} = S_1 \cup \cdots \cup S_k$, where $S_i \cap S_j = \emptyset$ for every $i \neq j$. Then $(\mathcal{A}_1(S_1), \ldots, \mathcal{A}_k(S_k))$ is ϵ-DP.*

A nice property of differential privacy is the post-processing property, which informally says that transforming private output does not incur additional privacy loss. Formally, we have the following:

Lemma C.2 (Post-processing). *Let $\mathcal{A}$ be an ϵ-DP algorithm. Then for any (possibly randomized algorithm) g, $g \circ \mathcal{A}(\mathcal{D})$ is ϵ-DP.*

Definition 9 (ℓ_p-sensitivity). *We define the ℓ_p-sensitivity of a function f to be*

$$\Delta_p f := \max_{\mathcal{D}, \mathcal{D}'} \|f(\mathcal{D}) - f(\mathcal{D}')\|_p,$$

where $\mathcal{D}, \mathcal{D}'$ are adjacent datasets.

Lemma C.3 (Gaussian mechanism). *Let f be a function, $\epsilon, \delta \in (0,1)$, and $\sigma^2 \geq (\Delta_2 f)^2 \cdot \frac{2\ln(1.25/\delta)}{\epsilon^2}$. The Gaussian mechanism $f(\mathcal{D}) + \mathcal{N}(0, \sigma^2)$ is (ϵ, δ)-DP.*

C.2 Optimal transport

It can easily be checked that indeed the Wasserstein distance is a metric: in particular, the triangle inequality holds.

Lemma C.4 ((Santambrogio, 2015), Lemma 5.4). *For any $p \geq 1, \mu, \nu, \pi \in \mathcal{P}_p(\mathcal{X})$, we have $W_p(\mu, \pi) \leq W_p(\mu, \nu) + W_p(\nu, \pi)$.*

For bounded spaces, we can bound the p-Wasserstein distance by the 1-Wasserstein distance:

Lemma C.5 ((Santambrogio, 2015)). *Let $\mathcal{X}$ be bounded. Then for any $p \geq 1, \mu, \nu \in \mathcal{P}_p(\mathcal{X})$, we have $W_p(\mu, \nu) \leq \mathrm{diam}(\mathcal{X})^{(p-1)/p} W_1(\mu, \nu)^{1/p}$.*

C.3 JL transform

A JL transform is any (linear) map that satisfies the JL lemma:

Theorem C.6 (Johnson-Lindenstrauss lemma, (Johnson & Lindenstrauss, 1984)). *Given an accuracy parameter $0 < \gamma < 1$, a set of n points X in $\mathbb{R}^d$, and the projection dimension $d' = O(\log n/\gamma^2)$, there exists a linear map $\Pi : \mathbb{R}^d \to \mathbb{R}^{d'}$ such that all pairwise distances are preserved within factor $(1 \pm \gamma)$, i.e., it holds*

$$\frac{1}{1+\gamma}\|x - y\| \leq \|\Pi x - \Pi y\| \leq (1+\gamma)\|x - y\|$$

for every $x, y \in X$, with probability $1 - 1/\mathrm{poly}(n)$.

In other words, the JL transform reduces the dimensionality of the data from d to d' while approximately preserving all pairwise distances (whp). Note that this worst-case guarantee is the main strength of the JL approach, other dimensionality reduction techniques such as principal component analysis typically do not have this guarantee.

C.4 USEFUL LEMMATA

Lemma C.7. *Let $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ and $\mu_\sigma := \mu * \mathcal{N}(0, \sigma^2 I_d)$. Then*

$$W_p(\mu, \mu_\sigma) \lesssim \sigma \left(\sqrt{d} + \sqrt{2p}\right).$$

Proof. By definition of Wasserstein distance, we have

$$W_p(\mu, \mu_\sigma) \le (\mathbb{E}[X - (X + \sigma Z)]^p)^{1/p} = \sigma \left(\mathbb{E}\|Z\|^p\right)^{1/p},$$

where $X \sim \mu$ and $Z \sim \mathcal{N}(0, I_d)$. We have

$$\mathbb{E}\|Z\|^p = \mathbb{E}\left(\|Z\|^2\right)^{p/2} = \mathbb{E}[Y^{p/2}],$$

where $Y := \|Z\|^2 \sim \chi^2_d$ (chi-squared with d degrees of freedom). Using (Laurent & Massart, 2000), it holds

$$\Pr[Y - d \ge 2\sqrt{dt} + 2t] \le \exp(-t)$$
$$\Pr[d - Y \ge 2\sqrt{dt}] \le \exp(-t),$$

so we can apply Theorem 2.3 of (Boucheron et al., 2003), which yields the desired result. □

D COMPARING CLUSTERING AND WASSERSTEIN BARYCENTERS

D.1 CLUSTERING

For completeness, we provide a brief discussion of the construction for private clusterings. We start by formalizing the problem statement for clustering:

Definition 10 ((k,p)-clustering)**.** *Given $k \in \mathbb{N}$ and a dataset $\mathbf{X} = \{x_1, \dots, x_n\}$ in $B_0(1/2)$, we want to find k centers $c_1, \dots, c_k \in \mathbb{R}^d$ that minimizes*

$$\text{cost}_{\mathbf{X}}(c_1, \dots, c_k) := \sum_{i \in [n]} \left(\min_{j \in [k]} \|x_i - c_j\|\right)^p. \tag{6}$$

The optimal cost is denoted as OPT, where we suppress the dependence on $k, p, \mathbf{X}$.

Definition 11 (Approximation for clustering)**.** *A (w,t)-approximation algorithm for (k,p)-clustering outputs $c_1, \dots, c_k$ such that $\text{cost}(c_1, \dots, c_k) \le w \cdot \text{OPT} + t$.*

(Ghazi et al., 2020) showed the following:

Theorem D.1. *For any $p \ge 1$, suppose that there exists a polynomial time (not necessarily private) $(w,0)$-approximation algorithm for the (k,p)-clustering problem. Then, for every $\epsilon > 0$ and $\delta, \gamma, \xi \in (0,1)$, there exists an (ϵ,δ)-DP algorithm that runs in $\left(\frac{k}{\xi}\right)^{O_{p,\gamma}(1)} \cdot \text{poly}(nd)$ time and outputs an*

$$\left(w(1+\gamma), O_{p,\gamma,w}\left(\frac{k\sqrt{d}}{\epsilon} \cdot \text{poly}\log\frac{k}{\delta\gamma} + \frac{(k/\gamma)^{O_{p,\gamma}(1)}}{\epsilon} \cdot \text{poly}\log\frac{n}{\gamma}\right)\right)$$

-approximation for (k,p)-clustering, with probability $1 - \xi$.

The first term in the additive error comes from computing the centers in high dimension. The second term comes from bounding the error in low dimension.

The construction of private (k,p)-clustering in low dimension works as follows. Under DP, we start by finding a centroid set[7] of size $O(k \log n)$ with small multiplicative error. Then, we create a private coreset by "snapping" all the input data to their nearest point in the centroid set and use the (discrete) Laplace mechanism to privatize the counts.

[7] A set that contains a good approximate solution.

D.2 STABILITY

The (k, p)-clustering objective is stable in the following sense: suppose we fix a dataset $\mathbf{X}$ and k centers. If we move one datapoint and update the k centers, a large fraction of the points will remain clustered together. This fact is key to the accuracy of private clustering algorithms.

Due to the stability of the clustering objective, the snapping procedure above incurs small additive error and is easy to reason about. On the other hand, optimal transport plans are highly non-stable:

Proposition D.2. *Let $n \in \mathbb{N}$. Fix a distribution ν supported on n atoms with uniform weights. There exists a distribution over $\mathbb{R}$ such that moving one datapoint by $O\left(\frac{1}{n}\right)$ changes the mapping for every datapoint.*

Proof. We construct μ_n with support over $\left\{\frac{k}{n}\right\}_{k\in[n]}$. Recall that the optimal transport plan in one dimension can be computed from the cumulative density function, e.g. see (Santambrogio, 2015). In the setting of the proposition, this will just be based on the order of the datapoints. Thus moving the particle on $\frac{1}{n}$ to $\frac{2+\epsilon}{n}$ for any $\epsilon > 0$ will yield the desired result. □

Remark 3. *The distribution in Proposition D.2 also shows that in order to use the private coreset construction from k-means clustering for Wasserstein distances requires taking $k = \tilde{\Omega}(n)$ for small snapping error, when we usually should think of $k = \tilde{O}(1)$.*

Proposition D.3. *There exists a distribution over $\mathbb{R}^2$ such that changing one point by $O(1)$ causes all the points in the support of the output barycenter by $\Omega(1)$.*[8]

Proof. See Figure 5. □

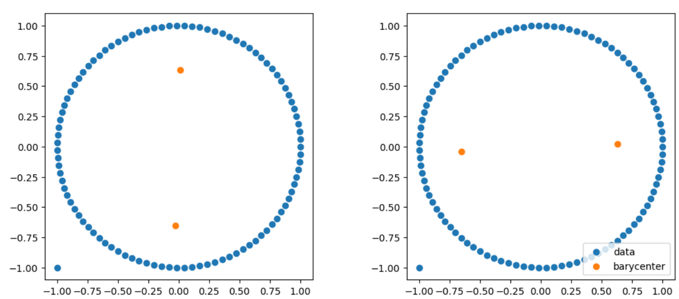

(a) Non-optimal barycenter. (b) Optimal barycenter.

Figure 5: Unperturbed data is uniform over $\mathbb{S}^1$. Here, the averages of any of the two disjoint half-arcs yield an optimal barycenter. However, with a bad initialization, each point in the support of the output distribution can move $\Omega(1)$ as $\Omega(n)$ of the couplings change.

E DEFERRED PROOFS

E.1 PROOF OF THEOREM 4.1

Theorem E.1 (Theorem 4.1, restated)**.** *For any $p \geq 1$, suppose there exists a (not necessarily private) (z, t)-approximation algorithm for the p-Wasserstein barycenter problem. Then for every $\epsilon > 0$, $\delta \in (0, 1)$, Algorithm 1 is (ϵ, δ)-DP, and yields a $\left((z, t + O_p\left(\left(\frac{md\ln(1/\delta)}{(\epsilon k)^2}\right)^{p/2}\right)\right)$-approximation for the p-Wasserstein barycenter problem.*

Proof. We prove the privacy and utility separately.

[8] whereas the expected should be $O\left(\frac{m}{n}\right)$

(Privacy) Consider the ℓ_2 sensitivity of the algorithm, which is a function $\mathcal{X}^{nk} \to \mathbb{R}^{m\times d}$. If we change one datapoint, the couplings of up to n elements could potentially change, namely all of those in the subpopulation. By normalization and Assumption 1, this implies the ℓ_2 sensitivity of $\|\nu^{(j)} - \nu'^{(j)}\|_2$ is $\frac{1}{k}$. Thus, the ℓ_2 sensitivity of the output is

$$\|\nu_1 \circ \cdots \circ \nu_m - \nu_1' \circ \cdots \circ \nu_m'\|_2 = \left(\sum_{j\in[m]} \|\nu_j - \nu_j'\|^2\right)^{1/2} \le \left(m\left(\frac{1}{k}\right)^2\right)^{1/2} = \frac{\sqrt{m}}{k},$$

where $\nu_1\circ\cdots\circ\nu_m \in \mathbb{R}^{md}$ is vector-concatenation. Privacy follows by the guarantees of the Gaussian mechanism (Lemma C.3).

(Utility) We have

$$\begin{aligned}
\mathrm{cost}_{\mu_1,\ldots,\mu_k}(\tilde{\nu}) &= \frac{1}{k}\sum_{i=1}^{k} W_p^p(\mu_i,\tilde{\nu}) \\
&\le \frac{1}{k}\sum_{i=1}^{k}(W_p(\mu_i,\nu) + W_p(\nu,\tilde{\nu}))^p \qquad (7)\\
&\le \frac{1}{k}\sum_{i=1}^{k} W_p^p(\mu_i,\nu) + p2^{p-1}W_p^p(\nu,\tilde{\nu}) \qquad (8)\\
&= \mathrm{cost}_{\mu_1,\ldots,\mu_k}(\nu) + p2^{p-1}W_p^p\left(\nu, \nu * \mathcal{N}\left(0, \frac{2m\ln(1.25/\delta)}{(\epsilon k)^2} I_d\right)\right)\\
&\le \mathrm{cost}_{\mu_1,\ldots,\mu_k}(\nu) + p2^{p-1}\left(\frac{2m\ln(1.25/\delta)}{(\epsilon k)}\right)^{p/2}\left(\sqrt{d}+\sqrt{2p}\right)^p \qquad (9)\\
&\lesssim \mathrm{cost}_{\mu_1,\ldots,\mu_k}(\nu) + O_p\left(\left(\frac{md\ln(1/\delta)}{(\epsilon k)^2}\right)^{p/2}\right),
\end{aligned}$$

where (7) follows from Lemma C.4, (8) follows from Lemma E.2, and (9) follows from Lemma C.7. This concludes the proof. □

E.2 Proof of Lemma 5.3

We will use the following lemma.

Lemma E.2. *For every $p \ge 1$, if $0 \le a, b \le 1$, then it holds*

$$(a+b)^p \le a^p + p(a+b)^{p-1}b \le a^p + p2^{p-1}b.$$

Lemma E.3 (Lemma 5.3, restated)**.** *Let $\mu_1,\ldots,\mu_k$ be discrete probability measures and suppose $\mu_1',\ldots,\mu_k'$ are $(p,1,t)$-coresets for each μ_i, respectively. Then,*

$$\mathrm{OPT}_{(\mu_1',\ldots,\mu_k')} \le \mathrm{OPT}_{(\mu_1,\ldots,\mu_k)} + O_p(t^p).$$

Proof. Consider a candidate barycenter ν. Using (1), we have

$$\begin{aligned}
\mathrm{cost}_{\mu_1',\ldots,\mu_k'}(\nu) &= \frac{1}{k}\sum_{i=1}^{k} W_p^p(\mu_i',\nu)\\
&\le \frac{1}{k}\sum_{i=1}^{k}\left(W_p(\mu_i',\mu_i) + W_p(\mu_i,\nu)\right)^p \qquad (10)\\
&\le \frac{1}{k}\sum_{i=1}^{k} W_p^p(\mu_i,\nu) + \frac{p2^{p-1}}{k}\sum_{i=1}^{k} W_p^p(\mu_i',\mu_i) \qquad (11)\\
&= \mathrm{cost}_{\mu_1,\ldots,\mu_k}(\nu) + \frac{p2^{p-1}}{k}\sum_{i=1}^{k} W_p^p(\mu_i',\mu_i) \qquad (12)
\end{aligned}$$

$$\leq \text{cost}_{\mu_1,\ldots,\mu_k}(\nu) + \frac{p2^{p-1}}{k}\sum_{i=1}^{k} t^p, \tag{13}$$
$$= \text{cost}_{\mu_1,\ldots,\mu_k}(\nu) + O_p(t^p)$$

where (10) follows from the triangle inequality, (11) follows from Lemma E.2, (12) follows from (1), and (13) follows from Definition 8. Note that applying Lemma E.2 uses Assumption 1. The claim follows. □

E.3 Proof of Theorem 5.5

Proof. Let d' be as in Theorem 5.4.

(Runtime) It suffices to bound the runtime of computing the barycenter in low dimensions as it is clear that the pre- and post-processing steps run in polynomial time. With the given d', we have that $2^{O(d')} \cdot \text{poly}(n,k) = \text{poly}(n) \cdot \text{poly}(n,k) = \text{poly}(n,k)$, as desired.

(Privacy) The privacy follows from Lemmas C.1 and C.2 as the input distributions are disjoint.

(Utility) For each μ_i, we invoke Theorem 5.2 with failure probability $\frac{\xi}{2k}$. Then by a union bound, with probability $1-\frac{\xi}{2}$, μ_i' is a $\left(1, O_p\left(\left(\frac{1}{(\epsilon n)^{1/d}} \cdot \text{poly}\log\left(\frac{k}{\xi}\right)\right)^{1/p}\right)\right)$-coreset for μ_i, for each $i \in [k]$. Assuming this event holds, Lemma 5.3 implies

$$\text{OPT}_{(\mu_1',\ldots,\mu_k')} \leq \text{OPT}_{(\mu_1,\ldots,\mu_k)} + O_p\left(\frac{1}{(\epsilon n)^{1/d}} \cdot \text{poly}\log\left(\frac{k}{\xi}\right)\right). \tag{14}$$

Let ν be the output of the algorithm. Now we apply Theorem 5.4, along with the guarantee of the not necessarily private approximation algorithm, which implies with probability $1-\frac{\xi}{2}$,

$$\text{cost}_{(\mu_1',\ldots,\mu_k')}(\nu) \leq z(1+\gamma)\text{OPT}_{(\mu_1',\ldots,\mu_k')} + O_{p,\gamma,z}(t) \tag{15}$$

By a union bound, (14) and (15) both occur with probability $1-\xi$. When this is the case, we deduce

$$\text{cost}_{(\mu_1',\ldots,\mu_k')}(\nu) \leq z(1+\gamma)\text{OPT}_{(\mu_1,\ldots,\mu_k)} + O_{p,\gamma,z}\left(\frac{1}{(\epsilon n)^{1/d}} \cdot \text{poly}\log\left(\frac{k}{\xi}\right) + t\right),$$

which concludes the proof. □

F Experiments

F.1 Additional experiments

We provide an example of the private coreset under uniform noise and scaled-down unifrom noise on the MNIST dataset (as points over $[-0.25, 0.25]^2$ in Figure 6. In particular, for visualization, we treat one image as one distribution.

In the next experiment with MNIST, we follow the setup of (Izzo et al., 2021) and treat each image as a point in $\mathbb{R}^n$. We consider $k = 10$ (each digit is one distribution), $d' = 25$, and $m = 40$. One difference is that we pre-process the data onto $B_{0.5}(0)$. We consider the cost to be the Wasserstein distance between the output and the Wasserstein barycenter over 5000 images from each class. Our experiment parameters are 0.01 for entropic regularization, 50 iterations, and 100 inner (Sinkhorn) iterations. In Figure 7, we show that subsampling on the private coresets to size n has a negligible increase in cost for sufficiently large n. In Figure 8, we observe that the choice of how points are sampled in each cell has a small effect on the cost for sufficiently large n.

In the US population experiment, we use GPT to implement the code for testing whether a point is inside the US population.

F.2 Additional discussion on experiment setup

In all the experiments, we scale the noise down by $(240d)^{1/d}$, e.g. Figure 8. We do not use zero noise as it usually is unstable.

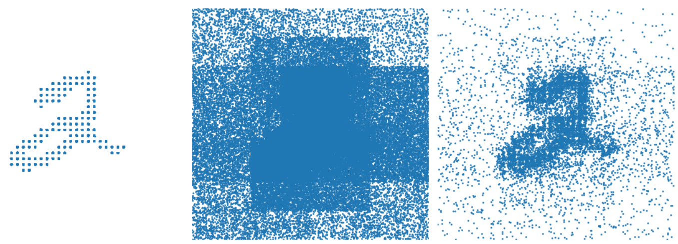

(a) Points chosen uniformly at random from the full cell.

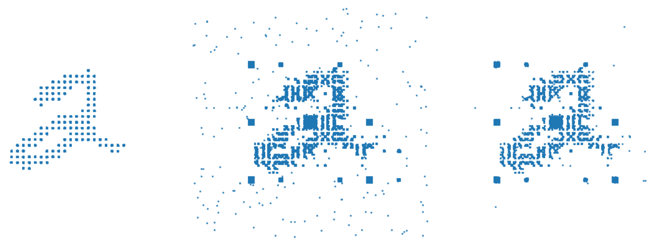

(b) Points chosen uniformly at random from a scaled-down version of the cell.

Figure 6: Example of a MNIST 2 digit with privacy parameter $\epsilon = 1$. On the left, we have the sensitive data, upsampled to $n = 10000$. In the middle, we have the full private coreset (with $O(n \log \epsilon n)$ points). On the right, we choose a subsample the private coreset down to n points.

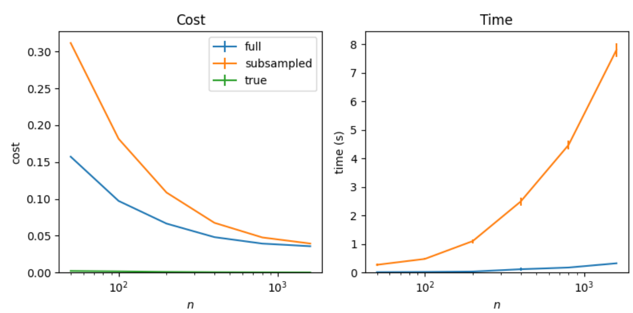

Figure 7: MNIST experiments with $n = 50, 100, \ldots, 1600$ and $\epsilon = 1$, averaged over 10 runs. (left) Cost of solutions. (right) Runtime in seconds.

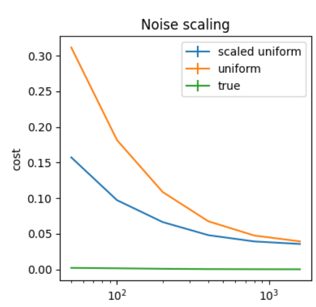

Figure 8: Subsampled experiment in the same setup as Figure 7. Uniform corresponds to the setting as Figure 6a and scaled uniform setup corresponds to the setting as Figure 6b.

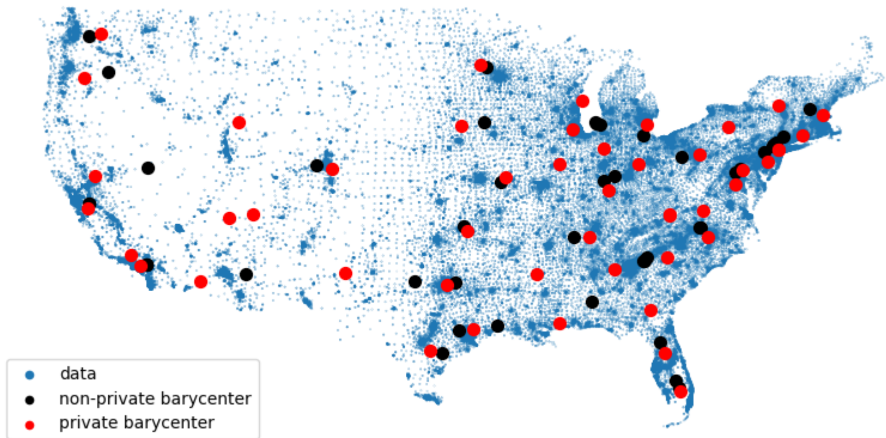

(a) $n = 200000$, $m = 48$, and $k = 1$. Denoting ν, ν_ϵ as the non-private and private barycenters, respectively, we have $\text{cost}(\nu) = 16.536$, $\text{cost}(\nu_\epsilon) = 20.407$ (squared degrees longitude/latitude), and $W_2(\nu, \nu_\epsilon) = 5.375$ (degrees).

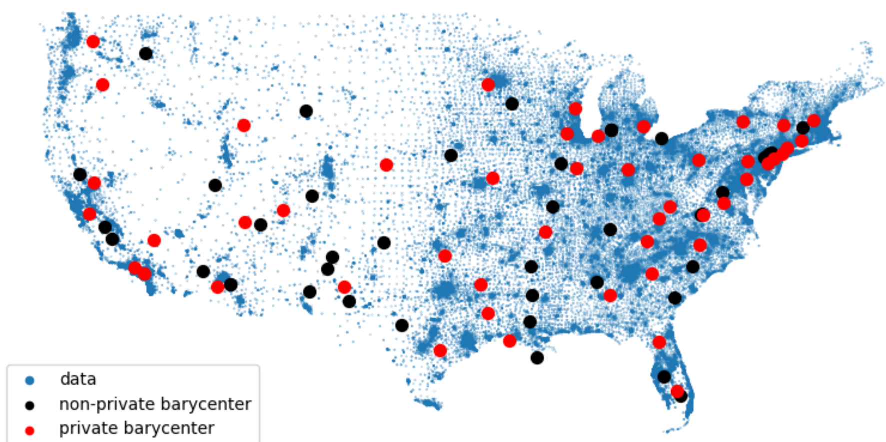

(b) $n = 100000$, $m = 48$, and $k = 4$ (self-reported White, Asian, Black, Hispanic). Denoting ν, ν_ϵ as the non-private and private barycenters, respectively, we have $\text{cost}(\nu) = 4451.297$, $\text{cost}(\nu_\epsilon) = 6343.896$ (squared degrees) and $W_2(\nu, \nu_\epsilon) = 12.76$ (degrees).

Figure 9: Same experimental setup as Figure 3, with $\epsilon = 5$.

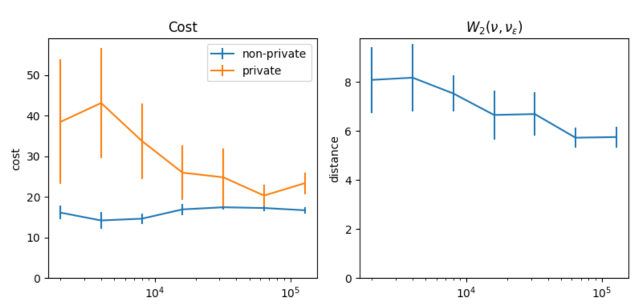

Figure 10: $n = 2000, 4000, \ldots, 128000$ and $\epsilon = 5$, similar to Figure 4, averaged over 10 trials.