

---

# Projection Robust Wasserstein Barycenters

---

Minhui Huang<sup>1</sup> Shiqian Ma<sup>2</sup> Lifeng Lai<sup>1</sup>

## Abstract

Collecting and aggregating information from several probability measures or histograms is a fundamental task in machine learning. One of the popular solution methods for this task is to compute the barycenter of the probability measures under the Wasserstein metric. However, approximating the Wasserstein barycenter is numerically challenging because of the curse of dimensionality. This paper proposes the projection robust Wasserstein barycenter (PRWB) that has the potential to mitigate the curse of dimensionality, and a relaxed PRWB (RPRWB) model that is computationally more tractable. By combining the iterative Bregman projection algorithm and Riemannian optimization, we propose two algorithms for computing the RPRWB, which is a max-min problem over the Stiefel manifold. The complexity of arithmetic operations of the proposed algorithms for obtaining an  $\epsilon$ -stationary solution is analyzed. We incorporate the RPRWB into a discrete distribution clustering algorithm, and the numerical results on real text datasets confirm that our RPRWB model helps improve the clustering performance significantly.

## 1. Introduction

The Wasserstein barycenter (WB) problem is attracting a lot of interest recently due to its wide applications in statistics and machine learning, including but not limited to image processing (Rabin et al., 2011), multi-level clustering (Ho et al., 2017), and text mining (Ye et al., 2017a;b). The WB serves as a geodesic interpolation between two or more distributions. It aggregates the underlying geometric structures of the input distributions under the Wasserstein metric. Therefore, the WB model provides deep insight when collecting information from probability distributions.

---

<sup>1</sup>Department of Electrical and Computer Engineering, University of California, Davis, CA, USA <sup>2</sup>Department of Mathematics, University of California, Davis, CA, USA. Correspondence to: Shiqian Ma <sqma@ucdavis.edu>.

However, computing the WB for a set of probability distributions is notoriously hard. The hardness comes from two aspects: the representation of the measure support and the potential curse of dimensionality. In many applications, the underlying distributions are unknown and we only have sampled data from these distributions. We wish to estimate the WB using the sampled data only. Therefore, the task reduces to compute WB from sampled discrete measures on a fixed number of support points. However, solving the free-support discrete WB is still very difficult (Borgwardt & Patterson, 2019). In this paper, we mainly consider the fixed-support WB problem. On the other hand, computing fixed-support WB can be challenging if the dimension of the problem is high. Recent theoretical developments have revealed that the sample complexity of approximating Wasserstein distances grows exponentially in dimension (Dudley, 1969; Weed & Bach, 2019). For the WB problem, (Altschuler & Boix-Adsera, 2021) has proved that computing WB is NP-hard since its runtime scales exponentially in the dimension. However, the sample complexity of original WB is still not well understood. Since the WB model minimizes a sum of Wasserstein distances, we conjecture that the WB problem would also have the issue of curse of dimensionality. To overcome this difficulty, we adopt a technique used in computing the Wasserstein distance (Paty & Cuturi, 2019) to the WB problem. The resulting projection robust WB (PRWB) model is an inf-sup-inf problem, which is computationally intractable. We further propose a relaxation of PRWB that is computationally more tractable. The idea of the new technique is to project the sampled data to a common low dimensional subspace and compute the WB of the projected data as an approximation to the original WB. The resulting problem is a max-min problem with Stiefel manifold constraint, and we propose two algorithms that can find an  $\epsilon$ -stationary point of it efficiently.

**Related work:** Most existing works for fixed-support WB focus on designing efficient algorithms. (Cuturi & Doucet, 2014) proposed to add an entropy regularizer and solve its dual problem that is smooth. This idea was further studied by (Benamou et al., 2015) under the name of iterative Bregman projection (IBP) algorithm. The convergence behavior of IBP was studied in (Kroshnin et al., 2019). There exist some other algorithms for computing fixed-support WB, including the dual accelerated gradient descent method

(Kroshnin et al., 2019; Lin et al., 2020b), the stochastic gradient descent method (Claici et al., 2018), the Bregman ADMM method (Ye et al., 2017b) and the interior-point method (Ge et al., 2019). On the other hand, people have proposed some efficient ways to mitigate the curse of dimensionality of the optimal transport (OT) problem (Niles-Weed & Rigollet, 2019). Specifically, (Bonneel et al., 2015) proposed the sliced Wasserstein distance and applied it to the WB problem. The sliced OT projects the sampled data to a random line and reduces the problem to a one-dimensional OT, which can be solved very efficiently by sorting. This idea motivated the work of (Paty & Cuturi, 2019; Niles-Weed & Rigollet, 2019) that suggest projecting the data to a low dimensional subspace. This leads to the projection robust Wasserstein (PRW) distance, and algorithms for computing it include (Lin et al., 2020a; Huang et al., 2020).

**Contributions:** Our main contributions are below.

(i) We propose a projection robust Wasserstein barycenter (PRWB) model. The PRWB model has the potential to mitigate the curse of dimensionality by projecting the probability measures onto a low dimensional subspace. Since PRWB is still numerically challenging to solve, we further propose a relaxation of PRWB (RPRWB) that is more tractable. Our numerical results indicate that RPRWB is more robust to noise compared with the WB.

(ii) We propose two algorithms: Riemannian block coordinate descent (RBCD) and Riemannian gradient ascent with IBP algorithm (RGA-IBP), to compute the RPRWB. The RGA-IBP incorporates the IBP algorithm to a Riemannian gradient ascent algorithm, and the RBCD is based on a reformulation of the max-min problem that is suitable for BCD type algorithms. The complexities of arithmetic operations of both algorithms for obtaining an  $\epsilon$ -stationary point are analyzed.

(iii) We conduct extensive numerical experiments to show the robustness and the practicality of the RPRWB model. We adopt RPRWB to the discrete distribution (D2) clustering algorithm, which we call the projected D2 clustering. We test this new algorithm on the real text datasets, and the numerical results show that the projected D2 clustering achieves better performance than the D2 clustering.

## 2. Optimal Transport and Wasserstein Barycenter

In this section, we review some background in optimal transport and Wasserstein barycenter.

Denote  $\mathcal{P}(\mathbb{R}^d)$  as the set of Borel probability measures in  $\mathbb{R}^d$  and  $\mathcal{P}_2(\mathbb{R}^d)$  as a subset of  $\mathcal{P}(\mathbb{R}^d)$  whose elements have finite second moment. The 2-Wasserstein distance between

probability measures  $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$  is defined as

$$\mathcal{W}(\mu, \nu) := \left( \inf_{\pi \in \Pi(\mu, \nu)} \int \|x - y\|^2 d\pi(x, y) \right)^{1/2}, \quad (1)$$

where  $\Pi(\mu, \nu)$  is the set of all joint distributions with marginals  $\mu$  and  $\nu$ . We denote  $\Delta^q = \{u \in \mathbb{R}_+^q \mid u^\top \mathbf{1}_q = 1\}$  as the probability simplex in  $\mathbb{R}^q$ . The WB of  $m$  probability measures  $\boldsymbol{\mu} := \{\mu^l\}_{l \in [m]}$  is the solution of the following problem:

$$\inf_{\boldsymbol{\omega} \in \mathcal{P}_2(\mathbb{R}^d)} \mathcal{WB}(\boldsymbol{\mu}, \boldsymbol{\omega}) := \sum_{l=1}^m \omega^l \mathcal{W}^2(\mu^l, \nu), \quad (2)$$

where  $\boldsymbol{\omega} \in \Delta^m$  is a given weighting vector and  $[m] := \{1, \dots, m\}$ . We use  $\text{Proj}_E$  to denote the orthogonal projector onto  $E$  for any  $E \in \mathcal{G}_k$ , where the Grassmannian  $\mathcal{G}_k := \{E \in \mathbb{R}^d \mid \dim(E) = k\}$  is the set of all  $k$ -dimensional subspaces of  $\mathbb{R}^d$ . For Wasserstein distance, (Paty & Cuturi, 2019) proposed the projection robust Wasserstein distance as follows:

$$\mathcal{P}_k(\mu, \nu) := \sup_{E \in \mathcal{G}_k} \mathcal{W}(\text{Proj}_E \mu, \text{Proj}_E \nu). \quad (3)$$

That is, the probability measures  $\mu$  and  $\nu$  are projected onto the  $k$ -dimensional subspace  $E$ , and the Wasserstein distance between the projected measures is computed as an approximation to the original Wasserstein distance. Moreover, to measure the worst case approximation, the subspace  $E$  that maximizes this Wasserstein distance is sought. The study in (Niles-Weed & Rigollet, 2019) shows that the projection robust Wasserstein distance is able to improve the sample complexity from  $O(n^{-1/d})$  for Wasserstein distance to  $O(n^{-1/k})$ , where  $n$  denotes the number of sampled data. This is a significant improvement since usually  $k \ll d$  for high dimensional OT. Therefore, the projection robust Wasserstein distance can mitigate the curse of dimensionality.

## 3. Projection Robust Wasserstein Barycenter

Our projection robust Wasserstein barycenter is motivated by the success of the projection robust Wasserstein distance and the sliced Wasserstein barycenter proposed in (Bonneel et al., 2015). By replacing the Wasserstein distance in (2) with the PRW distance (3), the fixed-support PRWB is defined as the solution of the following problem:

$$\inf_{\boldsymbol{\omega} \in \mathcal{P}_2(\mathbb{R}^d)} \sum_{l=1}^m \omega^l \mathcal{P}_k^2(\mu^l, \nu). \quad (4)$$

Plugging (3) into (4), we have

$$\begin{aligned} & \inf_{\nu \in \mathcal{P}_2(\mathbb{R}^d)} \sum_{l=1}^m \omega^l \sup_{E_\ell \in \mathcal{G}_k} \mathcal{W}^2(\text{Proj}_{E_\ell} \mu^l, \text{Proj}_{E_\ell} \nu) \\ &= \inf_{\nu \in \mathcal{P}_2(\mathbb{R}^d)} \sum_{l=1}^m \omega^l \sup_{U_\ell \in \text{St}(d, k)} \inf_{\pi^l \in \Pi(\mu^l, \nu)} \\ & \quad \int \|U_\ell^\top(x^l - y)\|^2 d\pi^l(x^l, y). \end{aligned} \quad (5)$$

According to (Paty & Cuturi, 2019)[Proposition 1], PRW is a well defined distance over  $\mathcal{P}_2(\mathbb{R}^d)$  and can be formulated as a sup-inf problem. Moreover, the support of the barycenter is fixed and our target barycenter  $\nu$  lies on a probability simplex. Our PRWB formulation (5) is a inf-sup-inf problem over  $m$  Stiefel manifolds. Solving (5) directly is extremely difficult, because of the complex inf-sup-inf structure and also the existence of  $m$  Stiefel manifolds constraints. Therefore, we propose the following relaxation to RPWB (5) that is more computationally tractable:

$$\begin{aligned} & \sup_{E \in \mathcal{G}_k} \inf_{\nu \in \mathcal{P}_2(\mathbb{R}^d)} \sum_{l=1}^m \omega^l \mathcal{W}^2(\text{Proj}_E \mu^l, \text{Proj}_E \nu) \\ &= \sup_{U \in \text{St}(d, k)} \inf_{\nu \in \mathcal{P}_2(\mathbb{R}^d)} \sum_{l=1}^m \omega^l \inf_{\pi^l \in \Pi(\mu^l, \nu)} \\ & \quad \int \|U^\top(x^l - y)\|^2 d\pi^l(x^l, y). \end{aligned} \quad (6)$$

More specifically, we first use a common projector  $\text{Proj}_E(\cdot)$  for all PRW distances, and then we switch the order of sup and the first inf. The relaxed model (6) searches for a common low-dimensional subspace, the union of all subspaces of  $m$  PRW distances, that maximizes the barycenter objective. Roughly speaking, we solve an easier problem in a low-dimensional subspace to approximate the original WB problem. We call (6) the Relaxed PRWB (RPRWB) and focus on solving this relaxed version in the rest of the paper. We first study some properties of RPRWB. The following proposition shows the existence of the optimal subspace  $E^*$ .

**Proposition 1** *Given a probability measure set  $\mu$ , the support of the barycenter  $\nu$ , the weight vector  $\omega$ , and  $k \in [d]$ , there exists an optimal  $E^*$  for the problem (6).*

Notice that the target barycenter  $\nu$  lies on a probability simplex. This combined with Proposition 1 indicates that the fixed-support RPRWB problem can be written as a max-min problem. Using  $U \in \text{St}(d, k)$  to denote an orthonormal basis of  $E$ , the RPRWB can be formulated as

$$\max_{U \in \text{St}(d, k)} \min_{\pi^l \in \Pi(\mu^l, \nu)} \sum_{l=1}^m \omega^l \int \|U^\top(x^l - y)\|^2 d\pi^l(x^l, y), \quad (7)$$

where  $\text{St}(d, k)$  denotes the Stiefel manifold,  $x^l$  is the support of  $\mu^l$  and  $y$  is the support of  $\nu$ .

**Remark 2** *We remark here that analyzing the sample complexity PRWB is highly nontrivial and the analysis in (Niles-Weed & Rigollet, 2019) for PRW does not apply here. In fact, we are not aware of any results for the sample complexity of the empirical discrete WB problem. There are only some computational hardness results (Altschuler & Boix-Adsera, 2021) showing WB is NP-hard because of the ‘‘curse of dimensionality’’. Since the WB problem minimizes the sum of a set of Wasserstein distances, we conjecture that the ‘‘curse of dimensionality’’ should be inherited by WB. Deriving the sample complexity of WB and PRWB is an important future topic.*

In this paper, we consider solving WB for a set of discrete distributions. Specifically, we denote  $X^l = [x_1^l; \dots; x_n^l] \in \mathbb{R}^{d \times n}$  as the support of each  $\mu^l$  and write  $\mu^l = \sum_{i=1}^n p_i^l \delta_{x_i^l}$ , where  $p^l \in \Delta^n$  and  $\delta_x$  denotes the Dirac function at  $x$ . The support of the barycenter  $\nu$  is given and denoted as  $Y = \{y_1, \dots, y_n\} \in \mathbb{R}^{d \times n}$ . Therefore, the barycenter can be written as  $\nu = \sum_{i=1}^n q_j \delta_{y_j}$  with  $q \in \Delta^n$ . Denote  $\pi = \{\pi^l\}_{l \in [m]}$ . Throughout this paper, we denote  $\mathcal{M} = \text{St}(d, k)$ . Computing the fixed-support RPRWB is equivalent to solving

$$\begin{aligned} & \max_{U \in \mathcal{M}} \min_{q \in \Delta^n} \sum_{l=1}^m \omega^l \mathcal{W}^2(\text{Proj}_E \mu^l, \text{Proj}_E \nu) \\ &= \max_{U \in \mathcal{M}} \min_{\pi \in \Pi(\mathbf{p})} f(\pi, U), \end{aligned} \quad (8)$$

where  $f(\pi, U) := \sum_{l=1}^m \omega^l \sum_{i,j=1}^n \pi_{i,j}^l \|U^\top(x_i^l - y_j)\|^2$ , and  $\Pi(\mathbf{p}) = \{\pi \mid \pi^l \in \mathbb{R}_+^{n \times n}, \pi^l \mathbf{1} = p^l, (\pi^l)^\top \mathbf{1} = (\pi^{l+1})^\top \mathbf{1}, l \in [m]\}$ .

## 4. The Riemannian Gradient Ascent and Riemannian BCD Algorithms

In this section, we propose two algorithms for solving (8): RGA-IBP and RBCD. We can show that both algorithms find an  $\epsilon$ -stationary point of (8) defined as follows.

**Definition 3** *We call  $(\hat{\pi}, \hat{U}) \in \Pi(\mathbf{p}) \times \mathcal{M}$  an  $\epsilon$ -stationary point of the fixed-support RPRWB problem (8), if the following two inequalities hold:*

$$\|\text{grad}_U f(\hat{\pi}, \hat{U})\|_F \leq \epsilon, \quad (9)$$

$$f(\hat{\pi}, \hat{U}) - f(\pi^*(\hat{U}), \hat{U}) \leq \epsilon, \quad (10)$$

where  $\text{grad}_U f(\hat{\pi}, \hat{U})$  is the Riemannian gradient w.r.t.  $U$ ,  $\pi^*(\hat{U})$  is the optimal solution of the inner minimization problem of (8) when fixing  $U$  as  $\hat{U}$ . The corresponding  $\epsilon$ -approximate barycenter  $q \in \Delta^n$  can be computed as  $q = (\hat{\pi}^\top)^\top \mathbf{1}, \forall l \in [m]$ .

Before we present the algorithms, we define some useful notation first.

**Algorithm 1** The RGA-IBP Algorithm

---

```

1: Input:  $\{\mu^l = (X^l, p^l)\}_{l \in [m]}$ ,  $\{Y\}$ , accuracy tolerance
    $\epsilon > 0$ . Set parameters  $\tau$ ,  $\eta$ , and  $\rho$  as in (24).
2: Initialization:  $U_0 \in \text{St}(d, k)$ .
3: for  $t = 0, 1, 2, \dots$ , do
4:    $\pi_{t+1} = \text{IBPsolver}(\mu, Y, U_t, \eta, \epsilon)$ ;
5:    $\xi_{t+1} = \text{Proj}_{\text{T}_{U_t}\mathcal{M}}(2V_{\pi_{t+1}}U_t)$ ;
6:    $U_{t+1} = \text{Retr}_{U_t}(\tau\xi_{t+1})$ ;
7:   if  $\|\xi_{t+1}\|_F \leq \epsilon$  then
8:     break;
9:   end if
10: end for
11: Output:  $\hat{U} = U_t$ ,  $\hat{\pi} = \pi_{t+1}$ .
    
```

---

**Definition 4** (*Cost and Correlation Matrices*) Given the support vectors  $\{X^l\}_{l \in [m]}$  and  $Y$ , the cost matrices, denoted as  $\{C^l\}_{l \in [m]}$ , are defined as  $C_{i,j}^l = \|x_i^l - y_j\|^2, \forall l \in [m]$ . The correlation matrix, denoted as  $V_{\pi}$ , is defined as  $V_{\pi} = \sum_{l=1}^m \omega^l \sum_{i,j=1}^n \pi_{i,j}^l (x_i^l - y_j)(x_i^l - y_j)^\top \in \mathbb{R}^{d \times d}$ .

#### 4.1. The Riemannian Gradient Ascent with IBP Iterations

The RGA-IBP algorithm is a natural extension of the RGAS algorithm (Riemannian gradient ascent with Sinkhorn's iteration) that was proposed by (Lin et al., 2020a) for computing the projection robust Wasserstein distance. Here we extend it to solve the RPRWB problem (8). The RGA-IBP algorithm solves the following problem, which is obtained by adding an entropy regularization to (8).

$$\begin{aligned} \max_{U \in \mathcal{M}} \min_{\pi \in \Pi(\mathbf{p})} f_{\eta}(\pi, U) := \\ \sum_{l=1}^m \omega^l \sum_{i,j=1}^n \pi_{i,j}^l \|U^\top (x_i^l - y_j)\|^2 - \eta H(\pi^l), \end{aligned} \quad (11)$$

where  $H(\pi) := -\sum_{i,j=1}^n (\pi_{i,j} \log \pi_{i,j} - \pi_{i,j})$  is the entropy regularizer, and  $\eta > 0$  is a weighting parameter. By further defining

$$f_{\eta}(U) := \min_{\pi \in \Pi(\mathbf{p})} f_{\eta}(\pi, U), \quad (12)$$

we know that (11) is equivalent to the following Riemannian optimization problem with smooth objective  $f_{\eta}(U)$ :

$$\max_{U \in \mathcal{M}} f_{\eta}(U). \quad (13)$$

Problem (13) can be naturally solved by a Riemannian gradient ascent algorithm whose  $t$ -th iteration is:

$$U_{t+1} := \text{Retr}_{U_t}(\tau \text{grad} f_{\eta}(U_t))$$

where  $\text{Retr}$  denotes the retraction operation,  $\text{grad} f_{\eta}$  denotes the Riemannian gradient of  $f_{\eta}$ , and  $\tau > 0$  is a step size.

Moreover, it is easy to verify that

$$\text{grad} f_{\eta}(U) = \text{Proj}_{\text{T}_U \mathcal{M}}(2V_{\pi_{\eta}^*(U)}U), \quad (14)$$

where  $\text{T}_U \mathcal{M}$  denotes the tangent space of  $\mathcal{M}$  at  $U$ , and  $\pi_{\eta}^*(U)$  is the optimal solution of (12) that can be found by the IBP algorithm (see details in Algorithm 4).

The RGA-IBP algorithm is detailed in Algorithm 1, where the IBP solver solves (12) up to an accuracy  $\epsilon$  (see Algorithm 4 in the supplementary material).

#### 4.2. The Riemannian Block Coordinate Descent Algorithm

Notice that the RGA-IBP requires to solve an optimization problem (12) in each iteration using an iterative solver. This can be quite expensive in practice. In this section, we propose the RBCD algorithm that can alleviate this computational burden. The RBCD algorithm presented here can be regarded as an extension of the algorithm recently proposed in (Huang et al., 2020) for computing the projection robust Wasserstein distance.

First, note that the optimization problem in (12) is convex and we have the following result about its dual.

**Lemma 5** *The dual problem of (12) is equivalent to the following problem:*

$$\max_{\mathbf{u}, \mathbf{v} \in \mathbb{R}^{m \times n}, \sum_{l=1}^m \omega^l v^l = 0} - \sum_{l=1}^m \omega^l \left\{ \sum_{i,j=1}^n \pi_{i,j}^l - \langle u^l, p^l \rangle \right\}, \quad (15)$$

where  $\pi_{i,j}^l = [\pi(u^l, v^l, U)]_{i,j}$  is given by:

$$\pi(u^l, v^l, U)_{i,j} = \exp \left( -\frac{\|U^\top (x_i^l - y_j)\|^2}{\eta} + u_i^l + v_j^l \right). \quad (16)$$

As a result, we know that (11) is equivalent to:

$$\begin{aligned} \min_{\substack{\mathbf{u}, \mathbf{v} \in \mathbb{R}^{m \times n}, U \in \mathcal{M}, \\ \sum_{l=1}^m \omega^l v^l = 0}} g(\mathbf{u}, \mathbf{v}, U) := \\ \sum_{l=1}^m \omega^l \left\{ \sum_{i,j=1}^n \pi_{i,j}^l - \langle u^l, p^l \rangle \right\}. \end{aligned} \quad (17)$$

Note that (17) has three block variables and it is suitable for block coordinate descent method. Our RBCD for solving (17) updates the iterates as follows:

$$\mathbf{u}_{t+1} := \underset{\mathbf{u}}{\text{argmin}} g(\mathbf{u}, \mathbf{v}_t, U_t) \quad (18)$$

$$\mathbf{v}_{t+1} := \underset{\mathbf{v}: \sum_{l=1}^m \omega^l v^l = 0}{\text{argmin}} g(\mathbf{u}_{t+1}, \mathbf{v}, U_t) \quad (19)$$

$$U_{t+1} := \text{Retr}_{U_t}(-\tau \text{grad}_U g(\mathbf{u}_{t+1}, \mathbf{v}_{t+1}, U_t)). \quad (20)$$

**Algorithm 2** The RBCD Algorithm

- 1: **Input:**  $\{\mu^l = (X^l, p^l)\}_{l \in [m]}, \{Y\}$ , accuracy tolerance  $\epsilon > 0$ . Set parameters  $\tau, \eta$  and  $\rho$  as in (27).
- 2: **Initialization:**  $U_0 \in \text{St}(d, k)$ ,  $\mathbf{u}_0, \mathbf{v}_0 \in \mathbb{R}^{m \times n}$ ,
- 3: **for**  $t = 0, 1, 2, \dots$ , **do**
- 4:   Compute  $\mathbf{u}_{t+1}, \mathbf{v}_{t+1}$  by (21)-(22);
- 5:   Compute  $U_{t+1}$  by (20);
- 6:   **if**  $\sum_{l=1}^m \omega^l \|q_t^l - \bar{q}_t\|_1 \leq \frac{\underline{\omega}^{3/2} \epsilon}{(12\bar{c})}$ , and  $\eta \|\text{grad}_{U^l} g(\mathbf{u}_{t+1}, \mathbf{v}_{t+1}, U_t)\|_F \leq \epsilon/3$  **then**
- 7:     **break;**
- 8:   **end if**
- 9: **end for**
- 10: **Output:**  $q = \bar{q}_t = \sum_{l=1}^m \omega^l q_t^l$ ,  $\hat{U} = U_t$ , and  $\hat{\pi}^l = \text{Round}(\pi^l(u_{t+1}^l, v_{t+1}^l, U_t), p^l, q), \forall l \in [m]$ .

According to (Benamou et al., 2015)[Propositions 1, 2],  $\mathbf{u}$  and  $\mathbf{v}$  steps in (18) and (19) are the same as the steps in the IBP algorithm, and they admit closed-form solutions below:

$$u_{t+1}^l = u_t^l + \log \frac{p^l}{\pi(u_t^l, v_t^l, U_t) \mathbf{1}}, \quad \forall l \in [m], \quad (21)$$

$$v_{t+1}^l = v_t^l + \log \frac{q_{t+1}^l}{q_t^l}, \quad \forall l \in [m], \quad (22)$$

where we denote  $q_t^l = (\pi(u_{t+1}^l, v_t^l, U_t))^\top \mathbf{1}$  and  $q_{t+1} = \exp(\sum_{l=1}^m \omega^l \log q_t^l)$ . Notice that (21)-(22) renormalize the sum of rows and columns of each  $\pi^l$  to be  $p^l$  and  $q_{t+1}$ , which yields  $\langle q_t^l, \mathbf{1} \rangle = 1, \forall l \in [m]$ . Moreover, the update (20) requires to compute  $\text{grad}_{U^l} g$ , and from (16) and (17) we know that

$$\text{grad}_{U^l} g(\mathbf{u}, \mathbf{v}, U) = \text{Proj}_{\mathcal{T}_{U^l} \mathcal{M}} \left( -\frac{2}{\eta} V_{\pi(\mathbf{u}, \mathbf{v}, U)} U \right). \quad (23)$$

By combining (18)-(23), we can summarize the details of the RBCD in Algorithm 2, in which we have adopted the following notation for the simplicity of presentation:

$$\bar{c} := \max_l \|C^l\|_\infty, \quad \underline{\omega} = \min_l \omega^l.$$

Note that in Algorithm 2 we adopted a rounding procedure for the output. This is because that  $\pi$  computed according to (16) does not necessarily lie in the constraint set  $\Pi(\mathbf{p})$ . The rounding procedure proposed in (Altschuler et al., 2017) and outlined in Algorithm 3 can help round the solution to set  $\Pi(\mathbf{p})$ . Note that this rounding procedure is also adopted in the IBP algorithm and thus in the RGA-IBP algorithm.

## 5. Convergence Analysis

In this section, we give the complexities of both the iteration number and the arithmetic operations for both RGA-IBP and RBCD for obtaining an  $\epsilon$ -stationary point of (8) as defined in

**Algorithm 3**  $\text{Round}(\pi, p, q)$ 

- 1: **Input:**  $\pi \in \mathbb{R}^{n \times n}, p \in \mathbb{R}^n, q \in \mathbb{R}^n$ .
- 2:  $X = \text{Diag}(x)$  with  $x_i = \frac{p_i}{[\pi \mathbf{1}]_i} \wedge 1$
- 3:  $\pi' = X \pi$
- 4:  $Y = \text{Diag}(y)$  with  $y_j = \frac{q_j}{[(\pi')^\top \mathbf{1}]_j} \wedge 1$
- 5:  $\pi'' = \pi' Y$
- 6:  $\text{err}_p = p - \pi'' \mathbf{1}, \text{err}_q = q - (\pi'')^\top \mathbf{1}$
- 7: **Output:**  $\pi'' + \text{err}_p \text{err}_q^\top / \|\text{err}_p\|_1$ .

**Definition 3.** The proofs are provided in the supplementary materials.

The next theorem and corollary are for RGA-IBP algorithm.

**Theorem 6** Choose parameters

$$\tau = \frac{1}{8L_2 \bar{c} + 2\rho L_1^2}, \quad \eta = \frac{\epsilon}{4 \log(n) + 2}, \quad \rho = 2\bar{c} + \frac{4\bar{c}^2}{\eta}. \quad (24)$$

The Algorithm 1 returns an  $\epsilon$ -stationary point defined in Definition 3 in

$$T = O(\log(n)^2 L_1^2 \bar{c}^4 k \epsilon^{-4}) \quad (25)$$

iterations.

**Corollary 7** The per iteration arithmetic operations complexity of Algorithm 1 is  $O(mn^2 dk + mn^2 \bar{c}^6 \log(n)^2 \epsilon^{-6})$ . Therefore, the total arithmetic operations complexity of Algorithm 1 is

$$O(mn^2 dk + mn^2 \bar{c}^6 \log(n)^2 \epsilon^{-6}) \cdot O(\log(n)^2 L_1^2 \bar{c}^4 k \epsilon^{-4}) \\ = \log(n)^2 L_1^2 \bar{c}^4 k O(mn^2 dk \epsilon^{-4} + mn^2 \bar{c}^6 \log(n)^2 \epsilon^{-10}). \quad (26)$$

The next theorem and corollary are for RBCD algorithm.

**Theorem 8** Choose parameters

$$\tau = \frac{1}{4L_2 \bar{c} / \eta + \rho L_1^2}, \quad \eta = \frac{\epsilon}{4 \log(n) + 2}, \quad \rho = \frac{2\bar{c}}{\eta} + \frac{4\sqrt{k} \bar{c}^2}{\eta^2}. \quad (27)$$

The Algorithm 2 returns an  $\epsilon$ -stationary point defined in Definition 3 in

$$T = O\left(L_1^2 \bar{c}^5 \log(n) \sqrt{k} \underline{\omega}^{-3} \epsilon^{-3}\right) \quad (28)$$

iterations.

**Corollary 9** The per iteration arithmetic operations complexity of Algorithm 2 is  $O(mn^2 d)$ . Therefore, the total arithmetic operations complexity of Algorithm 2 is

$$O\left(mn^2 \log(n) d L_1^2 \bar{c}^5 k^{\frac{3}{2}} \underline{\omega}^{-3} \epsilon^{-3}\right). \quad (29)$$

**Remark 10** Comparing (26) with (29), we see that RBCD has a better complexity dependence on  $\epsilon$  and  $\bar{c}$ . However, RBCD has an extra term  $\omega \leq \frac{1}{m}$ . Therefore, RGA-IBP has a better theoretical complexity when  $m$  is large.

## 6. Numerical Experiments

In this section, we conduct numerical experiments on both synthetic datasets and real datasets to evaluate the proposed RPRWB model (8)<sup>1</sup>. For the synthetic dataset, we consider solving RPRWB for a set of Gaussian distributions, which has closed-form solutions (Agueh & Carlier, 2011). We compare the convergence rate of WB and RPRWB to the ground truth for the sampled discrete distributions, as well as the robustness against noise for the RPRWB model. For real datasets, we incorporate the RPRWB model to the discrete distribution (D2) clustering algorithm (Ye et al., 2017b) and test it on text datasets. All experiments are conducted on a Linux server with a 32-core Intel Xeon CPU (E5-2667, v4, 3.20GHz per core).

### 6.1. Synthetic Dataset

**Multi-variable Gaussian Distributions:** It is well-known that the Wasserstein barycenter of a set of multi-variable Gaussian distributions  $\{\mu^l\}_{l \in [m]}$  with  $\mu^l = \mathcal{N}(a^l, \Sigma^l)$ , where  $a^l$  is the mean and  $\Sigma^l$  is the covariance matrix, has a closed-form formula (Agueh & Carlier, 2011). Specifically, we have the following theorem.

**Theorem 11 (Agueh & Carlier, 2011)** Let  $\mu^1, \dots, \mu^m$  be Gaussian distributions with respective means  $a^1, \dots, a^m$  and covariance matrices  $\Sigma^1, \dots, \Sigma^m$ . The barycenter of  $\mu^1, \dots, \mu^m$  with weights  $\omega^1, \dots, \omega^m$  is the Gaussian distribution with mean  $\bar{a} = \sum_{l=1}^m \omega^l a^l$  and covariance matrix  $\Sigma$  defined as the only positive definite matrix satisfying the equation

$$S = \sum_{l=1}^m \omega^l \left( S^{1/2} \Sigma^l S^{1/2} \right)^{1/2}. \quad (30)$$

In this subsection, we compute the WB and RPRWB of a given set of zero-mean multi-variable Gaussian distributions  $\{\mu^l\}_{l \in [m]}$ ,  $\mu^l = \mathcal{N}(0, \Sigma^l)$ . We set  $\omega^l = \frac{1}{m}, \forall l \in [m]$  in all experiments.

**The dependence of RPRWB on  $k$ .** We first explore the dependence of the objective function value on  $k$ . For each  $\mu^l$ , we sample an empirical measure  $\mu_n^l$ . Specifically, we sample  $n$  points according to the Gaussian distribution  $\mu^l = \mathcal{N}(0, \Sigma^l)$  to form the support matrix  $X^l \in \mathbb{R}^{d \times n}$  and set  $p_i^l = \frac{1}{n}, i \in [n]$ . We set each of the covariance matrices

<sup>1</sup>Code available at <https://github.com/mhhuang95/PRWB>.

$\Sigma^l$  to be a SPD matrix with rank  $k^*$ . Therefore,  $X^l$  lies in a  $k^*$ -dimensional subspace and the barycenter of  $\{\mu_n^l\}_{l \in [m]}$  should be in a  $(m \times k^*)$ -dimensional subspace. The support of the barycenter  $Y \in \mathbb{R}^{d \times n}$  is obtained by applying  $k$ -means clustering on  $X = [X^1; \dots; X^m] \in \mathbb{R}^{d \times mn}$ . We set parameters as  $d = 100, m = 3, n = 10$ . We further set the step size  $\tau = 0.0005$  for both RBCD and RGA-IBP algorithms and  $\eta = 0.5 \cdot \text{mid}(\{C^l\}_{l \in [m]})$ , where  $\text{mid}(\{C^l\}_{l \in [m]})$  is the median of the entries of  $\{C^l\}_{l \in [m]}$ .

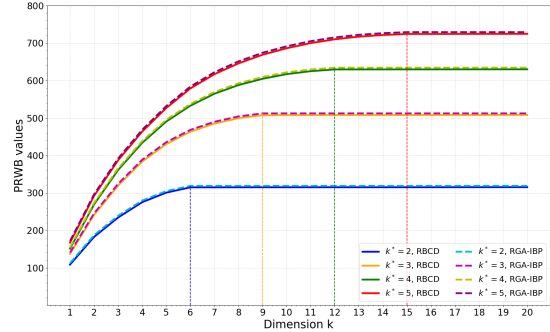


Figure 1. RPRWB function value versus projection dimension  $k$ . We run Algorithms 2 and 1 on different  $k$  and averaging over 100 runs,

We run both RBCD and RGA-IBP for solving (8) with different  $k^*$  and  $k$ , and report the results in Figure 1. From Figure 1 we see that the RPRWB values computed by the two algorithms are almost the same. We also notice that the RPRWB value increases when  $k < m \times k^*$  and remains as a constant when  $k \geq m \times k^*$ , which verifies the fact that the barycenter of  $\{\mu_n^l\}_{l \in [m]}$  lies in a  $(m \times k^*)$ -dimensional subspace.

**Robustness Against Noise.** We further conduct experiments on comparing the robustness of WB and RPRWB against noise. Specifically, we add Gaussian noise  $\sigma \mathcal{N}(0, I)$ , where  $\sigma$  is the noise level, to the discrete support  $\{X^l\}_{l \in [m]}$ . We compare the relative error of the objective function value for WB and RPRWB under different noise level  $\sigma$ . The relative error for WB and RPRWB is defined as

$$\text{Relative Error} = \frac{OBJ(\{\mu_n^l\}_\sigma) - OBJ(\{\mu_n^l\}_0)}{OBJ(\{\mu_n^l\}_0)},$$

where  $\{\mu_n^l\}_\sigma$  denotes the distributions after adding noise  $\sigma \mathcal{N}(0, I)$  and  $OBJ$  denotes the objective function of WB (the discrete version of (2)) or RPRWB (8). We set parameters as  $d = 100, m = 3, n = 10, \sigma \in [0.01, 0.1, 1, 2, 4, 7, 10]$ . We choose the step size  $\tau = 0.001$  when  $\sigma < 7$  and  $\tau = 0.0005$  otherwise for both RBCD

and RGA-IBP algorithms and  $\eta = 0.5 \cdot \text{mid}(\{C^l\}_{l \in [m]})$ . The results are shown in Figure 2, which shows that the proposed RPRWB model is more robust to noise compared to the WB.

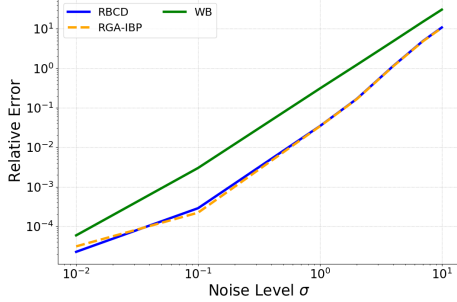


Figure 2. Relative error of the WB and RPRWB function value on different noise level  $\sigma$ . The results are averaged on 100 runs.

**Convergence rate to the ground truth.** We further consider approximating the Wasserstein barycenter for a set of continuous distributions by sampling data. Note that (Niles-Weed & Rigollet, 2019) proved that for a so-called spiked transport model, the mean projection robust Wasserstein distance between the sampled empirical distributions is  $O(n^{-1/k})$ , which improves the corresponding complexity of  $O(n^{-1/d})$ . We conjecture that similar results hold for WB and RPRWB and give some numerical evidence in this section. We set  $d = 10, m = 2, k = 2$ . The covariance matrices  $\Sigma^1, \Sigma^2 \in \mathbb{R}^{d \times d}$  are diagonal matrices with  $\Sigma^1(1, 1) = 10.1, \Sigma^2(2, 2) = 10.1$  and the rest of diagonal elements are all 0.1. In this case, a 2-dimensional subspace catches most of the information about the barycenter. We then sample  $n$  points as the support for each of  $\mu^l$ . To have a better estimation, the probability for  $x_i^l$  is computed according to the Gaussian PDF:

$$\text{Prob}(x_i^l) = \frac{1}{(2\pi)^{d/2} \det(\Sigma^l)^{1/2}} e^{-\frac{1}{2}(x_i^l)^T (\Sigma^l)^{-1} (x_i^l)}.$$

We sampled the support of the barycenter  $Y$  according to a uniform distribution over  $[-2, 2]^d$ . The barycenter Mean Estimation Error is defined as

$$\text{MEE} = |\text{OBJ}(\{\mu^l\}) - \text{OBJ}(\{\mu_n^l\})|,$$

where the ground truth objective function  $\text{OBJ}(\{\mu^l\})$  is calculated by solving (30) and  $\text{OBJ}(\{\mu_n^l\})$  is the sampled barycenter objective function value of the WB or RPRWB model. We set the step size  $\tau = 0.05$  for both RBCD and RGA-IBP algorithms and  $\eta = 0.5 \cdot \text{mid}(\{C^l\}_{l \in [m]})$  and select  $n \in \{20, 50, 100, 250, 500, 1000\}$ . The results are shown in Figure 3, which shows that the proposed RPRWB model converges to the ground truth much faster than the WB.

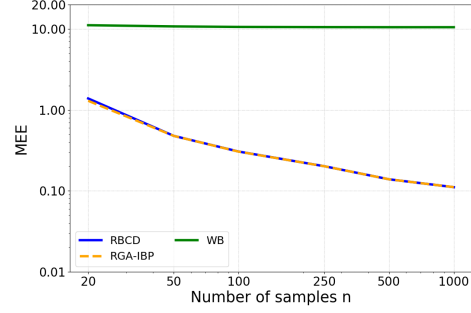


Figure 3. Mean Estimation Error (MEE) on different  $n$ . The results are averaged on 500 runs.

**Computational time comparison.** We compare the mean computational time of the WB solved by the IBP algorithm (Benamou et al., 2015) and the proposed RPRWB solved by RBCD and RGA-IBP. We set  $d = 100, m = 3, k = 2$  and select  $n \in \{20, 50, 100, 250, 500, 1000\}$ . We generate the support matrices  $X^l \in \mathbb{R}^{d \times n}$  from  $\mu^l = \mathcal{N}(0, \Sigma^l)$  by empirical sampling. The support of the barycenter  $Y \in \mathbb{R}^{d \times n}$  is obtained by k-means clustering. We further set the step size  $\tau = 0.01$  for both RBCD and RGA-IBP algorithms and  $\eta = 0.5 \cdot \text{mid}(\{C^l\}_{l \in [m]})$ . We stop the RBCD algorithm when  $\eta \|\text{grad}_U g(\mathbf{u}_{t+1}, \mathbf{v}_{t+1}, U_t)\|_F \leq \epsilon, \frac{1}{m} \sum_{l=1}^m \|q_t^l - \bar{q}_t\|_1 \leq \epsilon$ , and the RGA-IBP algorithm when  $\|\text{grad}_U f(\xi_{t+1})\|_F \leq \epsilon$ , and we set  $\epsilon = 10^{-4}$ . The results are shown in Figure 4, which shows that RBCD always runs faster than RGA-IBP. Note that the IBP for solving WB runs much faster than the other two algorithms, and this is because the latter two solve a more difficult problem.

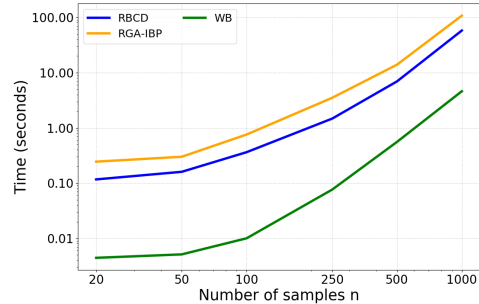


Figure 4. Computational time of the WB model solved by the IBP algorithm and the RPRWB model solved by RBCD and RGA-IBP algorithms on different  $n$ . The results are averaged on 100 runs.

## 6.2. Real Dataset: text data

We consider the discrete distributions (D2) clustering model proposed in (Ye et al., 2017b), which requires to solve the

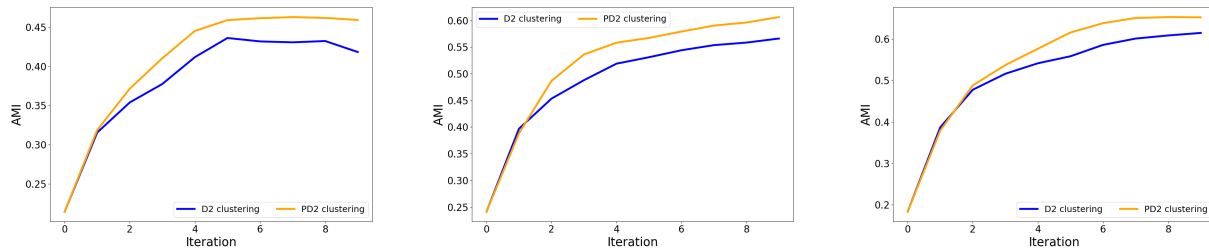


Figure 5. AMI scores for each iteration. **Left:** the “Reuters Subset” dataset, **Middle:** the “BBCsport Abstract” dataset, **Right:** the “BBCnews Abstract” dataset. The results are averaged on 5 runs.

free-support discrete Wasserstein barycenter model:

$$\begin{aligned} & \min_{\nu} \frac{1}{m} \sum_{l=1}^m \mathcal{W}(\mu^l, \nu) \\ & = \min_{\pi \in \Pi(\mathbf{p}), Y \in \mathbb{R}^{d \times n}} \frac{1}{m} \sum_{l=1}^m \sum_{i,j=1}^n \pi_{i,j}^l \|x_i^l - y_j\|^2, \end{aligned} \quad (31)$$

Note that there are two block variables:  $\pi$  and  $Y$ . (Ye et al., 2017b) proposed to solve (31) using an alternating minimization algorithm. That is, one alternately minimizes the objective function (31) with respect to one variable and with the other one fixed. This procedure is repeated until no progress can be made. When  $Y$  is fixed, problem (31) becomes a fixed-support WB problem. When  $\Pi = [\pi^1; \dots; \pi^m]$  is fixed, we have the following closed-form solution for  $Y$ :

$$y_i = \frac{1}{mq_i} \sum_{l=1}^m \sum_{j=1}^n \pi_{i,j}^l x_j^l, \quad (32)$$

which can be written more compactly as  $Y = \frac{1}{m} X \Pi^\top \text{diag}(1/q)$ .

Since we have numerically demonstrated that RPRWB might be a better model than WB, we propose to replace the WB problem in D2 clustering by RPRWB. We call the resulting algorithm projection robust D2 clustering (PD2 clustering). More details of the D2 and PD2 clusterings can be found in the supplementary material. We compare the performance of D2 and PD2 clusterings on three text datasets listed in Table 1. The “Reuters Subset” is a 5-class subset of the “Reuters” dataset<sup>2</sup>. The “BBCnews Abstract” and “BBCsport Abstract”<sup>3</sup> (Greene & Cunningham, 2006) are truncated versions of 2,225 and 737 posts. Each document retains only the title and the first sentence of the original post.

**Preprocessing.** We follow the idea of treating each document as a bag of word-vectors. For all three datasets in

Table 1. Text Datasets.  $N$  is the number of data,  $n$  is the number of samples, and  $K$  is the number of clusters.

Dataset	$N$	$d$	$n$	$K$
Reuters Subset	1209	300	16	5
BBCnews Abstract	2225	300	16	5
BBCsport Abstract	737	300	16	5

Table 1, we use the pre-trained word-vector dataset GloVe (Pennington et al., 2014) to transform a list of words to a measure over  $\mathbb{R}^{300}$ . The weight of each word is the normalized frequency modified by the TF-IDF scheme. We use the GloVe 300d (word vectors  $\in \mathbb{R}^{300}$ ) that was trained on 6 billion tokens and contains a 400,000 lower case vocabulary. Before we transform words into vectors, we lower the capital letters, remove all punctuations and stop words and lemmatize each document. Finally, we restrict the number of support points to  $n$  by recursively merging the closest words. Specifically, when the number of different words in  $\mu^l$  is larger than  $n$ , we solve the following discrete optimization problem:

$$\min_{i,j} p_i^l p_j^l \|x_i^l - x_j^l\|^2 / (p_i^l + p_j^l), \quad (33)$$

and merge  $p_i^l, p_j^l$  as  $\bar{p}^l = p_i^l + p_j^l$ ,  $\bar{x} = (p_i^l x_i^l - p_j^l x_j^l) / \bar{p}^l$ .

**Parameter setting and initialization.** In each iteration of PD2 clustering, we run the RBCD algorithm with the step size  $\tau = 0.05$ , the regularization parameter  $\eta = 1$ . We choose  $k = 2$  for the “BBCsport Abstract” dataset and  $k = 3$  for the “Reuters Subset” and the “BBCnews Abstract” datasets. The initial  $K$  barycenters are chosen randomly from documents with more than  $n$  different words and recursively merged so the number of support points remains  $n$ .

**The Adjusted Mutual Information.** To measure the performance of the clustering results, we use the Adjusted Mutual Information (AMI) (Vinh et al., 2010). Denote  $P_V(i) = |V_i|/N$  as the probability of clus-

<sup>2</sup><https://www.nltk.org/book/ch02.html>

<sup>3</sup><http://mlg.ucd.ie/datasets/bbc.html>



ter  $i$  in the partition  $V$ . The entropy  $H(V)$  is defined as  $H(V) = -\sum_{i=1}^{S_V} P_V(i) \log P_V(i)$ , where  $S_V$  is the number of clusters in  $V$ . The mutual information between the two partitions  $V_1, V_2$  is defined as  $MI(V_1, V_2) = \sum_{i,j=1}^{S_{V_1}, S_{V_2}} P_{V_1, V_2}(i, j) \log \frac{P_{V_1, V_2}(i, j)}{P_{V_1}(i)P_{V_2}(j)}$ , where  $P_{V_1, V_2}(i, j) = |V_{1,i} \cap V_{2,j}|/N$ . The AMI score between two partitions  $V_1, V_2$  is computed by

$$AMI(V_1, V_2) = \frac{MI(V_1, V_2) - \mathbb{E}MI(V_1, V_2)}{(H(V_1) + H(V_2))/2 - \mathbb{E}MI(V_1, V_2)}.$$

The AMI score lies in the interval  $[0, 1]$ , and it remains unchanged when we permute the cluster labels. In our experiments, we present the AMI scores between the ground truth labels and the predicted labels.

**Clustering results.** We run D2 and PD2 on the two real datasets in Table 1. The final AMI score and the average number of iterations for different datasets are given in Tables 2 and Table 3 respectively. We apply k-means clustering on the raw TF-IDF vectors as a baseline. Each result is averaged over five runs with different initialization. We stop the D2 and PD2 algorithms when the labels for each cluster are stable. Comparing the AMI scores in Table 2, we see that the proposed RPRWB model improves the performance of text clustering. One possible reason is that for many real high dimensional datasets, a low dimensional subspace catches most of the information. Notice that the D2 clustering AMI scores reported here are smaller than those in (Ye et al., 2017a). This is because the clustering performance highly depends on the barycenter initialization, and we are reporting the average AMIs with different initialization while (Ye et al., 2017a) reported the best AMI they obtained. We further see that the average number of iterations of the PD2 algorithm is smaller. Moreover, we plot the AMI scores for the first ten iterations of the D2 and PD2 clustering algorithm in Figure 5. We see that the PD2 clustering algorithm gives better AMI scores than the D2 clustering algorithm, which shows the advantage of the proposed RPRWB model.

Table 2. AMI scores for clustering results.

Dataset	k-means	D2	PD2
Reuters Subset	0.4627	0.4200	<b>0.4713</b>
BBCnews Abstract	0.3877	0.6095	<b>0.6557</b>
BBCsport Abstract	0.4276	0.6510	<b>0.6892</b>

## 7. Conclusion

In this paper, we have proposed a novel WB model called the projection robust Wasserstein barycenter, which has the potential to mitigate the curse of dimensionality for the WB

Table 3. Average number of clustering iteration.

Dataset	D2	PD2
Reuters Subset	24.2	23.2
BBCnews Abstract	23.8	22.4
BBCsport Abstract	29.8	14.4

problem. To resolve the computational issue of the PRWB, we have proposed a relaxed PRWB model: RPRWB. We have proposed two algorithms, the RBCD algorithm and the RGA-IBP algorithm for solving the fixed-support RPRWB problem. We have analyzed the iteration complexity and complexity of arithmetic operations for both algorithms. Numerical results on synthetic datasets have demonstrated the robustness and the better sample complexity of the proposed RPRWB model comparing with the WB model. Moreover, we have incorporated the RPRWB model to the D2 clustering algorithm, and proposed the projection robust D2 clustering algorithm. Numerical results on real text datasets show that the PD2 clustering improves the performance of the D2 clustering. Future directions include deriving sample complexity for WB, PRWB and RPRWB.

## Acknowledgements

We would like to thank the AC and all anonymous reviewers for very constructive suggestions that improved the quality of the paper. We are also grateful to Tianyi Lin for discussions on this topic. This research was partially supported by NSF HDR TRIPODS grant CCF-1934568, NSF grants CCF-1717943, CNS-1824553, CCF-1908258, DMS-1953210, ECCS-2000415 and CCF-2007797, and UC Davis CeDAR (Center for Data Science and Artificial Intelligence Research) Innovative Data Science Seed Funding Program.

## References

- Absil, P.-A., Mahony, R., and Sepulchre, R. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- Agueh, M. and Carlier, G. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2): 904–924, 2011.
- Aldaz, J. A monotonicity property of variances. *Statistics & Probability Letters*, 83(5):1416–1419, 2013.
- Altschuler, J., Niles-Weed, J., and Rigollet, P. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in neural information processing systems*, pp. 1964–1974, 2017.
- Altschuler, J. M. and Boix-Adsera, E. Wasserstein

- barycenters are NP-hard to compute. *arXiv preprint arXiv:2101.01100*, 2021.
- Benamou, J.-D., Carlier, G., Cuturi, M., Nenna, L., and Peyré, G. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- Bonneel, N., Rabin, J., Peyré, G., and Pfister, H. Sliced and radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, 2015.
- Borgwardt, S. and Patterson, S. On the computational complexity of finding a sparse Wasserstein barycenter. *arXiv preprint arXiv:1910.07568*, 2019.
- Boumal, N., Absil, P.-A., and Cartis, C. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1):1–33, 2019.
- Claici, S., Chien, E., and Solomon, J. Stochastic Wasserstein barycenters. In *International Conference on Machine Learning*, pp. 999–1008, 2018.
- Cuturi, M. and Doucet, A. Fast computation of Wasserstein barycenters. 2014.
- Dudley, R. M. The speed of mean Glivenko-Cantelli convergence. *The Annals of Mathematical Statistics*, 40(1): 40–50, 1969.
- Ge, D., Wang, H., Xiong, Z., and Ye, Y. Interior-point methods strike back: Solving the Wasserstein barycenter problem. In *Advances in Neural Information Processing Systems*, pp. 6894–6905, 2019.
- Greene, D. and Cunningham, P. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *Proceedings of the 23rd international conference on Machine learning*, pp. 377–384, 2006.
- Ho, N., Nguyen, X., Yurochkin, M., Bui, H. H., Huynh, V., and Phung, D. Multilevel clustering via Wasserstein means. In *International Conference on Machine Learning*, pp. 1501–1509, 2017.
- Huang, M., Ma, S., and Lai, L. A Riemannian block coordinate descent method for computing the projection robust Wasserstein distance. *arXiv preprint arXiv:2012.05199*, 2020.
- Kroshnin, A., Tupitsa, N., Dvinskikh, D., Dvurechensky, P., Gasnikov, A., and Uribe, C. On the complexity of approximating Wasserstein barycenters. In *International conference on machine learning*, pp. 3530–3540. PMLR, 2019.
- Lin, T., Fan, C., Ho, N., Cuturi, M., and Jordan, M. Projection robust Wasserstein distance and Riemannian optimization. In *NeurIPS*, volume 33, 2020a.
- Lin, T., Ho, N., Chen, X., Cuturi, M., and Jordan, M. I. Fixed-support Wasserstein barycenters: Computational hardness and fast algorithm. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Niles-Weed, J. and Rigollet, P. Estimation of Wasserstein distances in the spiked transport model. *arXiv preprint arXiv:1909.07513*, 2019.
- Paty, F.-P. and Cuturi, M. Subspace robust Wasserstein distances. In *International Conference on Machine Learning*, pp. 5072–5081, 2019.
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Rabin, J., Peyré, G., Delon, J., and Bernot, M. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 435–446. Springer, 2011.
- Vinh, N. X., Epps, J., and Bailey, J. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *The Journal of Machine Learning Research*, 11:2837–2854, 2010.
- Weed, J. and Bach, F. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- Ye, J., Li, Y., Wu, Z., Wang, J. Z., Li, W., and Li, J. Determining gains acquired from word embedding quantitatively using discrete distribution clustering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1847–1856, 2017a.
- Ye, J., Wu, P., Wang, J. Z., and Li, J. Fast discrete distribution clustering using Wasserstein barycenter with sparse support. *IEEE Transactions on Signal Processing*, 65(9): 2317–2332, 2017b.