



Kantorovich–Rubinstein Distance and Barycenter for Finitely Supported Measures: Foundations and Algorithms

Florian Heinemann¹ · Marcel Klatt¹ · Axel Munk^{1,2,3} 

Accepted: 21 August 2022 / Published online: 7 November 2022
© The Author(s) 2022

Abstract

The purpose of this paper is to provide a systematic discussion of a generalized barycenter based on a variant of unbalanced optimal transport (UOT) that defines a distance between general non-negative, finitely supported measures by allowing for mass creation and destruction modeled by some cost parameter. They are denoted as Kantorovich–Rubinstein (KR) barycenter and distance. In particular, we detail the influence of the cost parameter to structural properties of the KR barycenter and the KR distance. For the latter we highlight a closed form solution on ultra-metric trees. The support of such KR barycenters of finitely supported measures turns out to be finite in general and its structure to be explicitly specified by the support of the input measures. Additionally, we prove the existence of sparse KR barycenters and discuss potential computational approaches. The performance of the KR barycenter is compared to the OT barycenter on a multitude of synthetic datasets. We also consider barycenters based on the recently introduced Gaussian Hellinger–Kantorovich and Wasserstein–Fisher–Rao distances.

The funding was provided by Deutsche Forschungsgemeinschaft (Grant Nos. CRC 1456, Cluster of Excellence MBExC 2067, RTG 2088).

✉ Axel Munk
munk@math.uni-goettingen.de
Florian Heinemann
florian.heinemann@uni-goettingen.de
Marcel Klatt
mklatt@mathematik.uni-goettingen.de

- ¹ Institute for Mathematical Stochastics, University of Göttingen, Goldschmidtstraße 7, 37077 Göttingen, Germany
- ² Max Planck Institute for Multidisciplinary Sciences, Am Faßberg 11, 37077 Göttingen, Germany
- ³ Medical Center Göttingen, Cluster of Excellence 2067 Multiscale Bioimaging - From Molecular Machines to Networks of Excitable Cells, Robert-Koch-Straße 40, 37075 Göttingen, Germany

Keywords Unbalanced optimal transport · Kantorovich–Rubinstein · Wasserstein · Barycenter · Multi-marginal transport · Ultrametric trees

Mathematics Subject Classification 90B80 · 90C08 · 90C46 · 05C05 · 90-08

1 Introduction

Over the past decade, optimal transport (OT) based concepts for data analysis (for a thorough treatment of the mathematical foundations of optimal transport see e.g. [59, 60, 72]) have seen increasing popularity. This is mainly due to the fact that OT based methods respect important features of the data's geometric structure. Furthermore, noteworthy advances have been achieved in various areas, such as optimisation [7, 32, 74], machine learning [24, 56, 75], computer vision [25, 66, 69] and statistical inference [34, 54, 67], among others. This methodological and computational progress recently also paved the way to novel areas of applications including genetics [20, 62] and cell biology [29, 39, 70, 73], to cite but a few. Of particular importance from a data analysis point of view are extensions to compare more than two measures, a prominent proposal being the Fréchet mean [22], in the present context known as *Wasserstein barycenter* [1]. Wasserstein barycenters allow for a notion of average on the space of probability measures, which is well-adapted to the geometry of the data [3, 4]. With recent progress on their computation [9, 12, 19, 28, 36, 41] they establish themselves even further as a promising tool in many fields of data analysis, such as texture mixing [58], distributional clustering [76], histogram regression [10], domain adaptation [51] and unsupervised learning [63], among others.

However, a well known drawback of the Wasserstein distance and its barycenters in various applications is their limitation to measures with equal total mass. In fact, in many real world instances the difference in total mass intensity is of crucial importance. Employing vanilla Wasserstein based tools on general positive measures necessitates the usage of a normalisation procedure to enforce mass equality between the measures. This approach is, by design, oblivious to the mass differences between the original measures and can limit its use in applications. Exemplary, we mention that normalisation destroys stoichiometric features in the analysis of protein interaction and pathways as pointed out in Tameling et al [70]. Overall, this might lead to incorrect conclusions on specific applications. An illustrative example is given in Fig. 1.

1.1 Prior Work

The limitation of OT based concepts dealing only with measures of equal total mass has opened a wealth of approaches to account for more general measures. As an early proposal of this idea, the *partial OT formulation* [11, 21] suggests to fix the total mass of the OT plan in advance, while relaxing the marginal constraints. Comparably

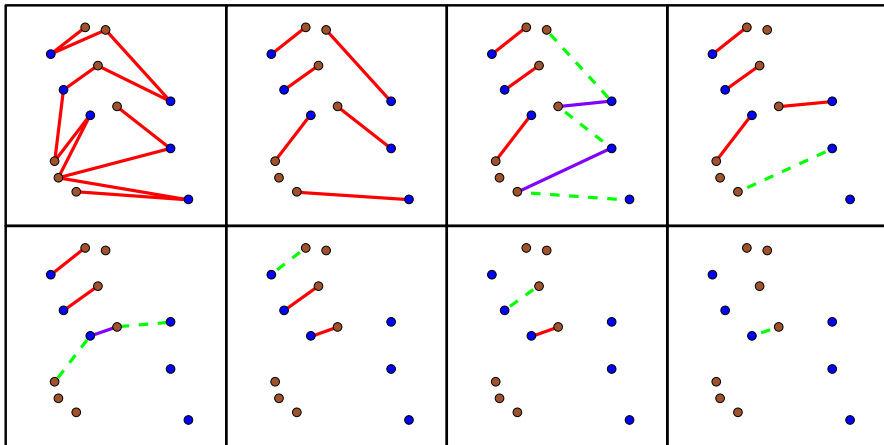


Fig. 1 (Unbalanced) OT between two measures (support in blue and brown, respectively) with weights equal to one at each support point. Top-Left: OT plan (red) between normalised versions of the two measures. Rest: UOT plans (red/purple) between non-normalised measures. From top-left to bottom-right C is decreasing. The edges, which have been removed most recently due to the reduction of C , are shown in green. Edges which have been added to the UOT graph due to the most recent reduction of C are marked in purple (Color figure online)

more recent are *entropy transport formulations*¹. This general framework removes the marginal constraints and instead uses a divergence functional to measure the deviation between the transport marginals and the input measures. The entropy transport framework encompasses the *Hellinger–Kantorovich distance* [15, 45], also known as *Wasserstein–Fisher–Rao distance* [14] and the *Gaussian Hellinger–Kantorovich distance* [45]. Inherent to all of these models is their dependency on parameters whose exact influence on the models’ properties is generally not well understood. An alternative idea is based on extending the well-studied dynamic formulation of OT [5] to measures with different total masses. With a focus on its geodesic properties, this approach has been studied in several works [14, 16, 26].

In this paper, we rely on a simple and intuitive idea based on the seminal work of Kantorovich and Rubinstein [37]. This accounts for mass construction and deletion at a cost modeled by some prespecified parameter (for details see also [33, 35]). It leads to the *Kantorovich–Rubinstein distance (KRD)* which curiously has been revisited several times under different names by various authors. For $p = 1$, it has been referred to as Earth Mover’s Distance [55], and generalized Wasserstein distance [57], while for general $p \geq 1$ common terminology includes Kantorovich distance [31], generalized KRD [61], transport-transform metric [53] and robust optimal transport distance [52].

¹ Critically, this is not to be confused with entropy *regularized* optimal transport, which is a popular computational approach adding an entropy penalty term to the OT problem to allow for efficient, approximate computations [6, 13, 18]

1.2 Contributions

In this work, we define barycenters with respect to the KRD and investigate their fundamental properties from a data analysis point of view. This extends the popular notion of Wasserstein barycenters to unbalanced barycenters (UBCs), i.e., barycenters of measures of different total masses. Similarly, UBCs have been considered explicitly for the Hellinger–Kantorovich distance [17, 23] and for the partial OT distance for absolutely continuous measures [38]. Notably, the well-known approach of matrix scaling algorithms has been shown to provide a general framework to approximate any UBC based on entropy optimal transport [15] of finitely supported measures. Closely related to our approach is the work by Müller et al [53] approximating the KR barycenter in the special case of point patterns.

The KR distance: Let (\mathcal{X}, d) be a finite metric space, where $\mathcal{X} = \{x_1, \dots, x_N\}$ and

$$\mathcal{M}_+(\mathcal{X}) := \left\{ \mu \in \mathbb{R}^{|\mathcal{X}|} \mid \mu(x) \geq 0 \forall x \in \mathcal{X} \right\}$$

is the set of non-negative measures² on \mathcal{X} . For a measure $\mu \in \mathcal{M}_+(\mathcal{X})$ its total mass is defined as $\mathbb{M}(\mu) := \sum_{x \in \mathcal{X}} \mu(x)$ and the subset of non-negative measures with total mass equal to one is the set of probability measures $\mathcal{P}(\mathcal{X})$. If $\pi \in \mathcal{M}_+(\mathcal{X} \times \mathcal{X})$ is a measure on the product space $\mathcal{X} \times \mathcal{X}$ its marginals are defined as $\pi(x, \mathcal{X}) := \sum_{x' \in \mathcal{X}} \pi(x, x')$ and $\pi(\mathcal{X}, x') := \sum_{x \in \mathcal{X}} \pi(x, x')$, respectively. For two measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ we define the set of *non-negative sub-couplings* as

$$\begin{aligned} \Pi_{\leq}(\mu, \nu) := & \{ \pi \in \mathcal{M}_+(\mathcal{X} \times \mathcal{X}) \mid \pi(x, \mathcal{X}) \leq \mu(x), \\ & \pi(\mathcal{X}, x') \leq \nu(x') \forall x, x' \in \mathcal{X} \}. \end{aligned} \tag{1}$$

Similarly, we denote the set of *couplings* between μ and ν as $\Pi_{=}(\mu, \nu)$, where the inequality constraints in (1) are replaced by equalities. For $p \geq 1$ and a parameter $C > 0$, *unbalanced optimal transport* (UOT) between two measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ is defined as

$$\begin{aligned} \text{UOT}_{p,C}(\mu, \nu) := & \min_{\pi \in \Pi_{\leq}(\mu, \nu)} \sum_{x, x' \in \mathcal{X}} d^p(x, x') \pi(x, x') \\ & + C^p \left(\frac{\mathbb{M}(\mu) + \mathbb{M}(\nu)}{2} - \mathbb{M}(\pi) \right). \end{aligned} \tag{2}$$

Notably, $\text{UOT}_{p,C}(\mu, \nu)$ is finite for all measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ with possibly different total masses and a solution of (2) always exists. Here, the parameter C penalizes deviation of mass from the marginals of π with respect to the input measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$. In particular and unlike the (balanced) OT problem

² A non-negative measure on a finite space \mathcal{X} is uniquely characterized by the values it assigns to each singleton $\{x\}$. To ease notation we write $\mu(x)$ instead of $\mu(\{x\})$. The corresponding σ -field is always to be understood as the powerset of \mathcal{X} .

$$\text{OT}_p(\mu, \nu) := \min_{\Pi=(\mu, \nu)} \sum_{x, x' \in \mathcal{X}} d^p(x, x') \pi(x, x')$$

defined only for measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ with equal total mass $\mathbb{M}(\mu) = \mathbb{M}(\nu)$, UOT in (2) relaxes the marginal constraint and allows optimal solutions to have more flexible marginals. Based upon UOT we define the p -th order *Kantorovich-Rubinstein distance* between two measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ as

$$\text{KR}_{p,C}(\mu, \nu) := (\text{UOT}_{p,C}(\mu, \nu))^{1/p}. \tag{3}$$

For any $p \geq 1$, it defines a distance on the space of non-negative measures $\mathcal{M}_+(\mathcal{X})$ and it is an extension of the well-known p -Wasserstein distance $W_p(\mu, \nu) := (\text{OT}_p(\mu, \nu))^{1/p}$ defined only for measures of equal total mass. Indeed, the KRd is shown to interpolate in-between *OT on small scales* and *point-wise comparisons on large scales* (Theorem 2.2) relative to the parameter C . This allows for an intuitive interpretation of the KRd. More precisely, in Lemma 2.1, we detail a clear geometrical connection between the value of C and the structure of the UOT. In particular, this contrasts the closely related partial OT problem [21] mentioned above. Employing Lagrange multipliers one can see that for any choice of C , there exists a fixed mass m of the partial OT problem, such that these two problems are equivalent. However, finding this value of m requires to solve the UOT problem. We stress that the influence of m on the resulting transport is in general hard to determine, while the impact of C is intuitively clear. Thus, this perspective seems better suited to many applications. For the specific case of measures supported on ultrametric trees (Sect. 2.1.1) we prove (Theorem 2.3) an analogue of the well-known closed formula for the p -Wasserstein distance [40]. Additionally, the computation of the KRd is known to be equivalent to solving a related balanced OT problem [33], allowing to apply any state-of-the-art solver with minimal modifications to compute the KRd and plan.

The KR barycenter: The KRd also lends itself to define a notion of a barycenter for a collection of measures as a generalization of the p -Wasserstein barycenter defined for probability measures $\mu_1, \dots, \mu_J \in \mathcal{P}(\mathcal{X})$ as

$$\tilde{\mu} \in \arg \min_{\mu \in \mathcal{P}(\mathcal{Y})} \frac{1}{J} \sum_{i=1}^J W_p^p(\mu, \mu_i). \tag{4}$$

Here, (\mathcal{X}, d) is assumed to be embedded in some ambient space (\mathcal{Y}, d) , e.g., an Euclidean space with $\mathcal{X} \subset \mathcal{Y}$. The distance d on \mathcal{X} is understood to be the distance on \mathcal{Y} restricted to \mathcal{X} . For $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$, any measure

$$\mu^* \in \arg \min_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu) := \frac{1}{J} \sum_{i=1}^J \text{KR}_{p,C}^p(\mu_i, \mu) \tag{5}$$

is said to be a (p, C) -Kantorovich–Rubinstein barycenter or (p, C) -barycenter for short³. We refer to the objective functional $F_{p,C}$ as (unbalanced) (p, C) -Fréchet functional. Notably, (p, C) -barycenters' support is not restricted to the finite space \mathcal{X} which raises fundamental questions on its structural properties. In the following, we establish that there exists a finite set containing the support of any (p, C) -barycenter (Sect. 2.2). Indeed, this set can be explicitly constructed from the support of the individual μ_i 's, but its size grows exponentially in the number of individual measures. However, we prove that there always exists a *sparse* (p, C) -barycenter whose support size is at most linear in the number of measures (Theorem 2.5). We note that these properties are analogs of well-known properties of Wasserstein barycenters [4], that we re-establish for the unbalanced setting.

Comparably, employing more general entropy transport distances, we are not aware of any similar structural description of their barycenters in terms of the input measures and the parameter. Notably, the entropy optimal transport barycenter of dirac measures is not necessarily finitely supported itself (for an example see [23]). In contrast, our explicit structural description of the support of KR barycenters provides an immediate understanding of its properties for a given choice of C . This clear link between C and the (p, C) -barycenter also allows to incorporate previous knowledge of the measures or the ground space into the choice C . The (p, C) -barycenter can be tuned to be more flexible and provide superior performance compared to its p -Wasserstein counterpart by avoiding to normalise each measure. An illustrative example is included in Fig. 2, where the (p, C) -barycenter detects all clusters correctly, while the Wasserstein barycenter does not provide any structural information on the underlying measures. This showcases potentially superior robustness and flexibility of the (p, C) -barycenter compared to the Wasserstein barycenter. We study this comparison in more detail on multiple synthetic data sets in Sect. 4. Here, the computational results⁴ are based on the fact that, due to our structural analysis of the support of the (p, C) -barycenter, it is straightforward to modify any given state-of-the-art solver for the Wasserstein barycenter problem to solve the (p, C) -barycenter problem (Sect. 4.1).

2 Kantorovich–Rubinstein Distance and (p, C) -Barycenter

In this section, we provide some theoretical analysis of the structural properties inherent in the UOT in (2) and as a consequence to the KRD in (3). We also focus on the variational formulation defining the (p, C) -barycenter in (5).

2.1 KR Distance

In this subsection, we focus on structural properties of minimizers for UOT in (2) and their consequences for the KRD. Notably, one can equivalently restate the penalization of total mass in (2) as

³ For the sake of readability, the weights in this definition are fixed to $1/J$, though it is easy to adapt all instances of their occurrence in this work to arbitrary positive weights $\lambda_1, \dots, \lambda_J$, summing to 1.

⁴ An implementation can be found in the R package *WSGeometry* on CRAN.

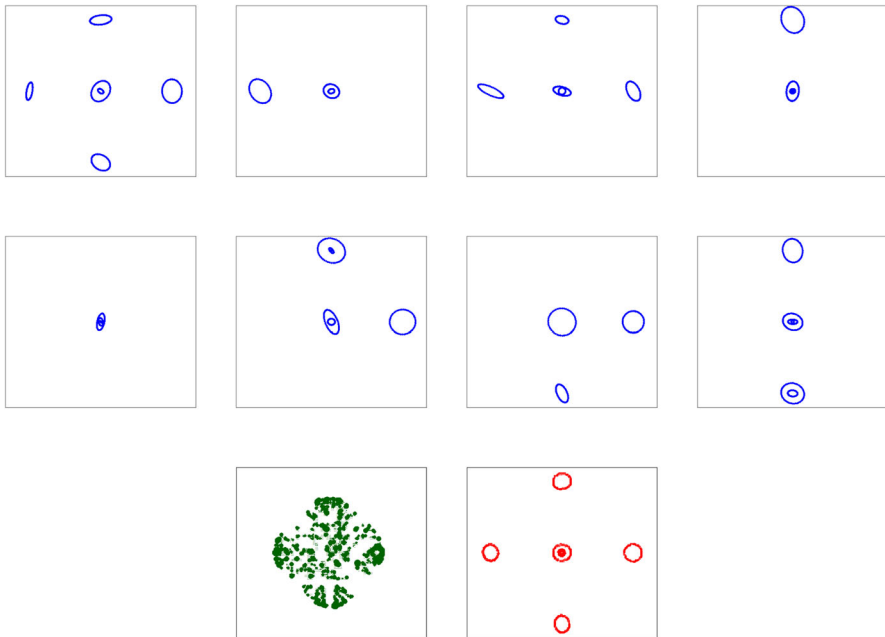


Fig. 2 Upper two rows: An excerpt of eight instances of a dataset of $N = 100$ nested ellipses at up to 5 different clusters in $[0, 1]^2$. The number of ellipses in each cluster follows a Poisson Distribution. For the cluster in the center the intensity is 2 and for the four outer clusters the intensity is 1. Each ellipse is discretized into 50 points with mass 1 at each location. For details on the computational methods refer to Sect. 4. Bottom-Left: The Wasserstein barycenter of the normalized versions of these measures (runtime about 15 hours). Bottom-Right: The $(2, 0.2)$ -barycenter of these measures (runtime about 30 minutes). The $(2, C)$ -barycenter for different values of C can be seen in Fig. 8

$$\begin{aligned}
 & C \left(\frac{\mathbb{M}(\mu) + \mathbb{M}(\nu)}{2} - \mathbb{M}(\pi) \right) \\
 &= \frac{C}{2} \left(\sum_{x \in \mathcal{X}} (\mu(x) - \pi(x, \mathcal{X})) + \sum_{x' \in \mathcal{X}} (\nu(x') - \pi(\mathcal{X}, x')) \right). \tag{6}
 \end{aligned}$$

While in (2) the parameter $C > 0$ controls the deviation of the total mass of π , the alternative representation (6) demonstrates its marginal characterization. Indeed, the parameter C specifies the maximal distance (scale) for which transportation is cheaper than creation or destruction of mass. More precisely, each optimal solution π_C for (2) induces a directed transportation graph $G(\pi_C)$ between the support points of μ (source points) and the support points of ν (sink points). By definition, the graph $G(\pi_C)$ contains a directed edge (x, x') if and only if $\pi_C(x, x') > 0$. For a directed path $P = (x_{i_1}, \dots, x_{i_k})$ in $G(\pi_C)$ its path length is defined as $\mathcal{L}(P) = \sum_{j=1}^{k-1} d^P(x_{i_j}, x_{i_{j+1}})$. The parameter $C > 0$ determines the maximal path length for any path in $G(\pi_C)$ as the following statement demonstrates.

Lemma 2.1 For $p \geq 1$, parameter $C > 0$ and measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ consider the UOT (2) with an optimal solution π_C . The length of any directed path P from the corresponding transport graph $G(\pi_C)$ is bounded by

$$\mathcal{L}(P) \leq C^p.$$

In particular, if $d(x, x') > C$ then for any optimal solution of 2 it holds $\pi_C(x, x') = 0$.

A proof is included in Appendix A.2. Lemma 2.1 shows that the underlying transportation graph has maximal path length C^p which limits the interaction between source and sink points. It will be of crucial importance for closed formulas on ultrametric trees in the following subsection. As an immediate consequence we obtain some important statements on the KRD in (3) along with its metric property.

Theorem 2.2 For any $p \geq 1$ and parameter $C > 0$ the following statements hold:

- (i) The p -th order KRD in (3) defines a metric on the space of non-negative measures $\mathcal{M}_+(\mathcal{X})$.
- (ii) If $C \leq \min_{x \neq x'} d(x, x')$, then it holds that

$$KR_{p,C}^p(\mu, \nu) = \frac{C^p}{2} TV(\mu, \nu),$$

where $TV(\mu, \nu) := 1/2 \sum_{x \in \mathcal{X}} |\mu(x) - \nu(x)|$ is the total variation distance. The same equality holds for all $C > 0$ if $\mu(x) \geq \nu(x)$ for all $x \in \mathcal{X}$ or if $\mu(x) \leq \nu(x)$ for all $x \in \mathcal{X}$.

- (iii) If $C \geq \max_{x, x'} d(x, x')$ and $\mathbb{M}(\mu) = \mathbb{M}(\nu)$, then it holds that

$$KR_{p,C}^p(\mu, \nu) = W_p^p(\mu, \nu).$$

- (iv) If $C_1 \leq C_2$, then it holds

$$KR_{p,C_1}^p(\mu, \nu) \leq KR_{p,C_2}^p(\mu, \nu).$$

We stress that the metric property of the KRD in Theorem 2.2 (i) has already been established in specific instances, e.g., for $p = 1$ [57]. Our proof follows that of Theorem 2 in Müller et al [53] for uniform measures on point patterns with minor modifications.

Theorem 2.2 demonstrates how two measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ are compared with respect to KRD. Depending on the parameter $C > 0$ the optimal value interpolates between p -th order Wasserstein distance on small scales and total variation on larger scales with respect to C . Equivalently, these properties can be shown by considerations of the dual program for UOT in (2) given by

$$UOT_{p,C}(\mu, \nu) = \max_{\substack{f, g: \mathcal{X} \rightarrow \mathbb{R} \\ f \leq C^p/2, g \leq C^p/2}} \sum_{x \in \mathcal{X}} f(x)\mu(x) + \sum_{x' \in \mathcal{X}} g(x')\nu(x') \quad (\text{DUOT}_{p,C})$$

$$\text{s.t. } f(x) + g(x') \leq d^p(x, x'), \forall x, x' \in \mathcal{X},$$

where the equality holds due to *strong duality*. For $p = 1$ this can be further specified to

$$\text{UOT}_{1,C}(\mu, \nu) = \max_{\substack{f: \mathcal{X} \rightarrow \mathbb{R} \\ f \text{ 1-Lipschitz} \\ \|f\|_\infty \leq C/2}} \sum_{x \in \mathcal{X}} f(x)(\mu(x) - \nu(x))$$

which reveals its relation to the *flat metric* [8] as observed in Lellmann et al [44]; Schmitzer and Wirth [65]. As in general $\mathbb{M}(\mu) \neq \mathbb{M}(\nu)$, the bound $f, g \leq C^p/2$ on dual feasible solutions f, g is necessary for the dual to be finite. However, if the measures $\mu, \nu \in \mathcal{M}_+(\mathcal{X})$ have equal total mass $\mathbb{M}(\mu) = \mathbb{M}(\nu)$ and $C \geq \max_{x,x'} d(x, x')$, then the bound on dual feasible solutions is redundant and we obtain the dual of the usual OT problem

$$\begin{aligned} \text{OT}_p(\mu, \nu) &= \max_{f, g: \mathcal{X} \rightarrow \mathbb{R}} \sum_{x \in \mathcal{X}} f(x) \mu(x) + \sum_{x' \in \mathcal{X}} g(x') \nu(x') && \text{(DOT}_p\text{)} \\ \text{s.t. } f(x) + g(x') &\leq d^p(x, x'), \forall x, x' \in \mathcal{X}. \end{aligned}$$

2.1.1 KR Distance on Ultrametric Trees

For OT, the approximations of the underlying distance by a tree metric are common tools for theoretical and practical purposes. The former is usually employed for rates of convergence for the expectation of empirical OT costs [68] while in the latter tree approximations serve to reduce the computational complexity inherent in OT [42]. OT on ultrametric trees is also applied for the analysis of phylogenetic trees [27]. For an efficient computational implementation of UOT on tree metrics we refer to Sato et al [61]. Notably, while OT with tree metric costs has a closed form solution, this fails to hold for its UOT counterpart. An exception is given in terms of ultrametric trees for which not only OT [40] but also UOT admits a closed form solution, which we establish in this subsection.

To this end, consider a tree \mathcal{T} with nodes V , edges E attached with (non-negative) weights $w(e)$ for $e \in E$ and a designated root r . Two nodes $v, w \in V$ are connected by a unique path denoted $\mathcal{P}(v, w)$ either represented by a sequence of nodes or as a sequence of edges. The distance $d_{\mathcal{T}}(v, w)$ is equal to the sum of the weights of those edges contained in $\mathcal{P}(v, w)$. A *leaf* of \mathcal{T} is any node such that its degree (number of edges attached to the node) is equal to one and the set of all leaf nodes is denoted as $L \subset V$. A node v^* is termed *parent* of node v denoted by $\text{par}(v) = v^*$ if both are connected by a single edge but v^* is closer to the root than v . The parent of the root node is set to $\text{par}(r) = r$. For a node v its *children* are the elements of the set $\mathcal{C}(v) = \{w \in V \mid v \in \mathcal{P}(w, r)\}$. Notice that with this definition v is a child of itself (Fig. 3a for an illustration).

A tree \mathcal{T} is termed *ultrametric tree* if all its leaf nodes are at the same distance to the root. Equivalently, there exists a *height function* $h: V \rightarrow \mathbb{R}_+$ that is monotonically decreasing meaning that $h(\text{par}(v)) \geq h(v)$ and such that $h(v) = 0$ for $v \in L$. The




Fig. 3 General Tree Structures: **a** A tree graph \mathcal{T} with root r (orange), internal nodes (black) and leaf nodes L (green). By definition $\text{par}(v_5) = \text{par}(v_6) = v_2$ and the children of v_1 are equal $\mathcal{C}(v_1) = \{v_1, v_3, v_4, v_7, v_8\}$. The distance from each leaf node to the root may vary. **b** An ultrametric tree \mathcal{T} with height function h (red) such that $0 = h_4 < h_3 < h_2 < h_1 < h_0$. Edge weights are defined by the the difference of consecutive height values, e.g. $w(e_1) = h_0 - h_1$. Each leaf node (green) is at the same distance to the root r (orange) (Color figure online)

distance is set to $d_{\mathcal{T}}(v, \text{par}(v)) = |h(v) - h(\text{par}(v))|$ and extended on the full tree (Fig. 3b for an illustration).

Consider an ultrametric tree \mathcal{T} with height function h and measures μ^L, ν^L supported on the leaf nodes $L \subset V$. We prove that the p -th order KR D admits a *closed formula* for such a setting. Intuitively, the parameter C restricts transportation of mass up to a certain threshold allowing to decompose \mathcal{T} into subtrees. Mass transportation is restricted solely within each subtree whereas mass abundance or deficiency is penalized with parameter C for each particular subtree (Fig. 4 for an illustration). We define the set

$$\mathcal{R}(C) := \left\{ v \in V \mid h(v) \leq \frac{C}{2} < h(\text{par}(v)) \right\} \tag{7}$$

with the convention that $\mathcal{R}(C) = \{r\}$ if $C/2 \geq h(r)$ and for a node $v \in V$ set

$$\mu^L(\mathcal{C}(v)) := \sum_{w \in \mathcal{C}(v) \cap L} \mu^L(w).$$

Theorem 2.3 (KR on ultrametric trees) *Consider an ultrametric tree \mathcal{T} with leaf nodes L and height function $h : V \rightarrow \mathbb{R}_+$ inducing the tree metric $d_{\mathcal{T}}$. For any $p \geq 1$ and two measures $\mu^L, \nu^L \in \mathcal{M}_+(L)$ supported on the leaf nodes of \mathcal{T} it holds that*

$$\begin{aligned} & KR_{p,C}^p(\mu^L, \nu^L) \\ &= \sum_{v \in \mathcal{R}(C)} \left(2^{p-1} \sum_{w \in \mathcal{C}(v) \setminus \{v\}} \left((h(\text{par}(w))^p - h(w)^p) \left| \mu^L(\mathcal{C}(w)) - \nu^L(\mathcal{C}(w)) \right| \right) \right) \end{aligned}$$




Fig. 4 Closed formula for the KRD on ultrametric trees: **a** Depending on the regularization $C > 0$ and the underlying height function h the ultrametric tree \mathcal{T} introduced in Fig. 3b is decomposed into two subtrees. Each node in the set $\mathcal{R}(C) = \{v_1, v_2\}$ (orange) serves as a new root and corresponding subtrees $\mathcal{T}(v_1) := \mathcal{C}(v_1)$ and $\mathcal{T}(v_2) := \mathcal{C}(v_2)$ are equal their respective set of children with corresponding edges. **b** The p -th height transformation $\mathcal{T}_p(v_1)$ and $\mathcal{T}_p(v_2)$ of the induced subtrees $\mathcal{T}(v_1)$ and $\mathcal{T}(v_2)$, respectively. Each subtree is extended by a new root (blue) with an edge (lightblue) whose distance is equal the difference of regularization $C^p/2$ and the p -th height transformed value of the former root

$$+ \left(\frac{C^p}{2} - 2^{p-1}h(v)^p \right) \left| \mu^L(\mathcal{C}(v)) - v^L(\mathcal{C}(v)) \right|.$$

The closed formula in Theorem 2.3 decomposes the underlying UOT into two tasks. While summing over subtrees carried out by the outer sum, the inner sum consists of two terms. The first considers OT within each subtree whereas the second accounts for mass deviation on that particular subtree.

The proof of this formula is given in Appendix A.2.1.

2.2 (p, C)-Barycenters

In the finite setting considered in this work a (p, C) -barycenter as defined in (5) always exists, but is not necessarily unique. Moreover, the location and structure of the support of the (p, C) -barycenter are not fixed and hence unknown. For the Wasserstein barycenter there exists a finitely supported, sparse barycenter in this context [4, 43]. We establish analog properties of the (p, C) -barycenter.

Definition 2.4 Let (\mathcal{Y}, d) be a metric space, $p \geq 1$ and $J \in \mathbb{N}$. A Borel barycenter application $T^{J,p}$ associates to any points $(y_1, \dots, y_J) \in \mathcal{Y}^J$ a minimum $y^* \in \mathcal{Y}$ of $\sum_{i=1}^J d^p(y_i, y)$, i.e.,

$$T^{J,p}(y_1, \dots, y_J) \in \arg \min_{y \in \mathcal{Y}} \sum_{i=1}^J d^p(y_i, y).$$

A Borel barycenter application is in general not a function since the minimum does not need to be unique. In particular, $y = T^{J,p}(y_1, \dots, y_J)$ only means that y is one of the minima of the average distance function. As the measures μ_1, \dots, μ_J are defined




Fig. 5 Centroid sets and barycenters: The support points of three ($J = 3$) measures (yellow, brown and blue dots) with unit mass at each position. Top: Different centroid sets $C_{KR}(3, 2, C)$ (red squares) with increasing value of C from left to right. Bottom-Left: The centroid set $C^W(3, 2)$ (dark green squares) corresponding to the 2-Wasserstein barycenter. Bottom-Center: Circles corresponding to C^p (for two different choices of C) balls around the support points. The grey colouring indicates that there is no overlap of at least two circles in this area and thus no $(2, C)$ -barycenter can have mass in this area. Conversely, the green colouring indicates overlap and thus the potential support area of the barycenter. Bottom-Right: The $(2, C)$ -barycenter (red squares) for a specific choice of C and the 2-Wasserstein barycenter (dark green squares) (Color figure online)

on \mathcal{X} we usually restrict the Borel barycenter application to inputs from the space $\mathcal{X} \subset \mathcal{Y}$. We define the *full centroid set* of the measures $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$ as

$$C_{KR}(J, p) = \left\{ y \in \mathcal{Y} \mid \exists L \geq \lceil J/2 \rceil, \exists (i_1, \dots, i_L) \subset \{1, \dots, J\}, \right. \\ \left. x_1, \dots, x_L : x_l \in \text{supp}(\mu_{i_l}) \right. \\ \left. \forall l = 1, \dots, L : y = T^{L,p}(x_1, \dots, x_L) \right\}, \tag{8}$$

and the *restricted centroid set*

$$C_{KR}(J, p, C) = \left\{ y = T^{L,p}(x_1, \dots, x_L) \in C_{KR}(J, p) \mid \forall 1 \leq l \leq L : \right. \\ \left. d^p(x_l, y) \leq C^p; \sum_{i=1}^L d^p(x_i, y) \leq \frac{C^p(2L - J)}{2} \right\}. \tag{9}$$

We stress that for each L -tuple (x_1, \dots, x_L) one fixed representative of $T^{L,p}(x_1, \dots, x_L)$ is chosen for the construction of the centroid set $C_{KR}(J, p, C)$. To streamline the presentation any statement concerning $C_{KR}(J, p, C)$ in the following theorem is to be understood in the sense that there exists a choice of $C_{KR}(J, p, C)$ such that the statement holds true.

Theorem 2.5 Let $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$ be a collection of non-negative measures on the finite discrete space $\mathcal{X} \subset \mathcal{Y}$. For any $C > 0$ it holds that

(i)

$$\inf_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu) = \inf_{\substack{\mu \in \mathcal{M}_+(\mathcal{Y}) \\ \text{supp}(\mu) \subseteq \mathcal{C}_{KR}(J,p,C)}} F_{p,C}(\mu).$$

Moreover, any (p, C) -barycenter μ^* satisfies $\text{supp}(\mu^*) \subseteq \mathcal{C}_{KR}(J, p, C)$ and its total mass is bounded by

$$0 \leq \mathbb{M}(\mu^*) \leq \frac{2}{J} \sum_{i=1}^J \mathbb{M}(\mu_i).$$

(ii) For any (p, C) -barycenter μ^* and any point $y \in \text{supp}(\mu^*)$, there exist UOT plans π_i between μ^* and μ_i for $i = 1, \dots, J$, respectively, such that if $\pi_i(y, x) > 0$, then there exists $L \geq \lceil J/2 \rceil$, $x_l \in \text{supp}(\mu_{i_l})$ for $l = 2, \dots, L$, $(i_2, \dots, i_L) \subset \{1, \dots, J\}$ and $i_l \neq i$ for $l = 2, \dots, L$ with $y = T^{L,p}(x, x_{i_2}, \dots, x_{i_L})$, $\pi_j(y, x_j) > 0$ if $j \in \{i_2, \dots, i_L\}$.

Additionally, if for any $(x_1, \dots, x_L) \in \mathcal{Y}^L$ it holds that

$$T^{L,p}(x_1, \dots, x_L) = T^{L,p}(y_1, x_2, \dots, x_L) \Leftrightarrow x_1 = y_1, \tag{10}$$

then $\pi_i(y, x) \in \{0, \mu^*(y)\}$ for $i = 1, \dots, J$.

(iii) If $M_i := |\text{supp}(\mu_i)|$ for $1 \leq i \leq J$ then there exists a (p, C) -barycenter μ^* such that

$$|\text{supp}(\mu^*)| \leq \min \left\{ |\mathcal{C}_{KR}(J, p, C)|, \sum_{i=1}^J M_i \right\}.$$

(iv) If $C_1 \leq C_2$, then it holds

$$\inf_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C_1}(\mu) \leq \inf_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C_2}(\mu).$$

(v) Furthermore, set $\mathcal{Z} := \bigcup_{i=1}^J \text{supp}(\mu_i) \cup \mathcal{C}_{KR}(J, p)$ and define

$$d'_{\min} := \min_{x \in \mathcal{Z} \setminus \mathcal{C}_{KR}(J,p), y \in \mathcal{C}_{KR}(J,p)} d(x, y).$$

If $C \leq d'_{\min}$, then the (p, C) -barycenter μ^* is given by

$$\mu^* = \sum_{x \in \mathcal{X}} \text{med}(\mu_1(x), \dots, \mu_J(x)) \delta_x.$$

(vi) Let $C > J^{1/p} \text{diam}(\mathcal{Z})$ and let μ_1, \dots, μ_J be ordered such that $\mathbb{M}(\mu_i) \leq \mathbb{M}(\mu_j)$ for $i \leq j$. Suppose that J is odd or there there exists no point $y \in \mathcal{Y}$ contained

in at least $J/2$ different support sets. Then, for any (p, C) -barycenter μ^* it holds that $\mathbb{M}(\mu^*) = \mathbb{M}(\mu_{\lceil J/2 \rceil})$. Else, there exists at least one (p, C) -barycenter with this total mass.

The proof is based on the fact that finding a (p, C) -barycenter can be proven to be equivalent to solving a *multi-marginal optimal transport problem* (Sect. 3.2). Statement (i) provides insights into the structure of the support of any (p, C) -barycenter and its dependency with respect to the magnitude of C . The definition of $\mathcal{C}_{KR}(J, p, C)$ can be understood as a joint restriction on $\sum_{i=1}^L d^p(x_i, y)$ combined with an individual restriction on each $d^p(x_i, y)$ of the original centroid points of $\mathcal{C}_{KR}(J, p)$. The joint restriction ensures that simply deleting any mass at a given centroid point (and thus reducing the total mass of the measure) does not improve the objective value. This is a minimal feasibility assumption on the considered centroid point, as otherwise no measure containing this point can be optimal. The second restriction concerns each point individually. If a point x_i has a distance larger than C^p from a point y , then, by Lemma 2.1, there is no transport between y and x_i . Thus, centroids which have a larger distance to one of the points x_1, \dots, x_L they are constructed from can not be in the support of any (p, C) -barycenter. This also gives rise to some helpful intuition for the support structure of any (p, C) -barycenter. Considering all C^p -neighbourhoods around any of the support points of the μ_i , then a (p, C) -barycenter can only have support in regions where at least balls from $\lceil J/2 \rceil$ different measures intersect. A visual representation of this is given in the center of bottom row of Fig. 5. By definition, the sets $\mathcal{C}_{KR}(J, p, C)$ are equipped with a natural ordering in the sense that $C_1 \leq C_2$ then $\mathcal{C}_{KR}(J, p, C_1) \subseteq \mathcal{C}_{KR}(J, p, C_2)$. Moreover, if C is large enough then $\mathcal{C}_{KR}(J, p, C) = \mathcal{C}_{KR}(J, p)$. We illustrate these sets in the top row Fig. 5. We observe that the cardinality of the restricted centroid set in (9) decreases with decreasing C . In the extremes for large C the restricted centroid sets coincides with the full centroid sets in (8) that is independent of C . For small C , if there is no point which is contained in the support of at least $J/2$ measures, the restricted centroid set is empty. For an illustration we refer to the top row of Fig. 5.

Property (ii) is an analogue to a well-known characterization [4] of the p -Wasserstein barycenter on \mathbb{R}^d with Euclidean distance d_2 , where the transport from the barycenter to the underlying measures is characterized by a transport map. The corresponding statement for the (p, C) -barycenter holds true as well in this context. Indeed, on (\mathbb{R}^d, d_2) condition (10), which can be understood as an injectivity-type assumption on the barycentric application, is satisfied due to the fact that $T^L(x_1, \dots, x_L) = \frac{1}{L} \sum_{i=1}^L x_i$. However, for (\mathbb{R}^d, d_1) this assertion does not hold. Consider $x_1 < x_2 < x_3 < x_4 \in \mathbb{R}$ and measures $\mu_1 = \delta_{x_1} + \delta_{x_2}, \mu_2 = \delta_{x_3} + \delta_{x_4}$, then any measure of the form $\mu^* = 2\delta_y$ for any $y \in [x_2, x_3]$ is a (p, C) -barycenter for $C > 2|x_1 - x_4|$. Thus, there only exist mass-splitting UOT plans between μ^* and μ_1, μ_2 and the transport is not characterized by a transport map. On more general spaces such as a tree \mathcal{T} rooted at r , three leaves x_1, x_2, x_3 and positive edge weights $e_1, \dots, e_3 \in (0, 1)$ the barycenter on \mathcal{T} of any two leafs $x_i \neq x_j$, is the root r . In particular, in this example, or in fact in any tree $\mathcal{T} = (V, E)$ which has a vertex y with degree of at least three⁵ condition (10) fails. The unique $(2, 2)$ -barycenter of two mea-

⁵ The degree of a vertex in a graph is the number of vertices which are adjacent to it.

asures $\mu_1 = \delta_{x_1} + \delta_{x_2}$ and $\mu_2 = \delta_{x_2} + \delta_{x_3}$ is given by $\mu^* = 2\delta_r$. Thus, there are again only mass-splitting UOT plans between μ^* and μ_1 and μ_2 . However, for the unit circle \mathcal{S}^1 equipped with its natural arc-length distance property (10) does hold. Assume $a_0 = T^L(x_1, \dots, x_L) = T^{L,p}(y_1, \dots, x_L)$, $a_1 = T^{L-1,p}(x_2, \dots, x_L)$ and for each $x \in \mathcal{S}_1$ denote $H_r(x)$ and $H_l(x)$ as the halfcircle right and left of x , respectively. It is straightforward to see by contraposition that if it holds $a_1 \in H_r(a_0)$, then this implies $x_1, y_1 \in H_l(a_1)$ and $x_1, y_1 \in H_l(a_0)$. However, it also holds $d(x_1, a_0) = d(y_1, a_0)$, and thus $\langle x_1 - y_1, a_0 \rangle = 0$. In particular, this implies that either $x_1 \in H_l(a_0)$ and $y_1 \in H_r(a_0)$ or vice versa and hence $x_1 = y_1$. The case $a_1 \in H_l(a_0)$ is analog and the case $a_0 = a_1$ clear.

Property (iii) guarantees the existence of *sparse* (p, C) -barycenters. For large C the size $\mathcal{C}_{KR}(J, p, C)$ scales as $\prod_{i=1}^J M_i$, growing essentially exponentially in J . However, here we see that there always exists a (p, C) -barycenter supported on a sparse subset of $\mathcal{C}_{KR}(J, p, C)$ which has cardinality growing only linearly in J . Part (iv) simply extends the monotonicity of the (p, C) -KRD to the (p, C) -Fréchet functional. Statement (v) yields a critical point after which decreasing C does no longer change the resulting (p, C) -barycenter and provides a closed form characterisation of the (p, C) -barycenter in this context. Finally, statement (vi) enables control on the total mass of the (p, C) -barycenter for large values of C . In particular, since the total mass is close to the median of the total masses of the μ_i , we point out that the total mass of the (p, C) -barycenter in this setting is robust against outliers. A small amount of measures with unreasonably high mass has no impact on the total mass of the (p, C) -barycenter.

Naturally, we compare the (p, C) -barycenter to its popular Wasserstein analogue in (4). As proven in Le Gouic and Loubes [43] (and initially for $p = 2$ for \mathbb{R}^d by [4]) the support of any p -Wasserstein barycenter is contained in

$$\mathcal{C}_W(J, p) = \left\{ y \in \mathcal{Y} \mid y = T^{J,p}(x_1, \dots, x_J), x_i \in \text{supp}(\mu_i) \right\}. \tag{11}$$

Compared to the p -Wasserstein barycenter of the probability measures μ_1, \dots, μ_J the restricted centroid set $\mathcal{C}_{KR}(J, p, C)$ allows more flexibility for specific cases and can provide a more reasonable representation of the data. We illustrate this in Fig. 5 (bottom-left/right) where the $(2, C)$ -barycenter clearly represents all clusters while the 2-Wasserstein barycenter fails to capture them. Nevertheless, if C is large enough and all measures have equal total mass both barycenters coincide.

Corollary 2.6 *If $C > 2^{\frac{1}{p}} \text{diam}(\mathcal{Z})$ and $\mathbb{M}(\mu_1) = \mathbb{M}(\mu_2) = \dots = \mathbb{M}(\mu_J)$, then any p -Wasserstein barycenter is also a (p, C) -barycenter and vice versa.*

While this shows that the (p, C) -barycenter is a strict generalisation of the usual p -Wasserstein barycenter as the solutions coincide for large C , for smaller values of C there can be significant differences. One such striking difference between the p -Wasserstein barycenter and the (p, C) -barycenter comes in the form of a localization property. Let $B_1, \dots, B_R \subset \mathcal{Y}$ such that $\text{supp}(\mu_i) \subset \cup_{r=1}^R B_r$ with $\text{diam}(B_r) \leq C$ for all $r = 1, \dots, R$ and $d(B_k, B_l) > 2^{1/p}C$ for all $k \neq l$. Here, the (p, C) -barycenter tends to place mass between the clusters B_1, \dots, B_R . However, a p -Wasserstein barycenter

is obtained by combining R barycenters of the measures restricted to the B_1, \dots, B_R , respectively.

Lemma 2.7 *Let $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$ such that for all $i = 1, \dots, J$ it holds $\text{supp}(\mu_i) \subset \cup_{r=1}^R B_r$ for some $B_1, \dots, B_R \subset \mathcal{Y}$ with $\text{diam}(B_r) \leq C$ for all $r = 1, \dots, R$ and $d(B_k, B_l) > 2^{1/p}C$ for all $k \neq l$. For $r = 1, \dots, R$, let*

$$\mu_r^* \in \arg \min_{\mu \in \mathcal{M}_+(\text{conv}(B_r))} \frac{1}{J} \sum_{i=1}^J KR_{p,C}^p(\mu, \mu_i|_{B_r}),$$

where $\text{conv}(B_r)$ is the convex hull of B_r for $r = 1, \dots, R$. Then, the measure $\sum_{r=1}^R \mu_r^*$ is a (p, C) -barycenter of μ_1, \dots, μ_J .

In particular, Lemma 2.7 implies that the (p, C) -barycenter respects the cluster structure within the supports of the measures if the clustered are sufficiently separated and C is adapted according to the cluster size. Examples of this setting can be seen in Fig. 2 and Fig. 5.

3 A Lift to Optimal Transport, Wasserstein Barycenters and Multi-marginal Optimal Transport

In this section, we provide the necessary tools and framework to establish our results in the previous section. Following the ideas of Guittet [33] we state UOT in (2) as an equivalent balanced OT problem. We extend this idea to the (p, C) -barycenter, showing it to be equivalent to a specific Wasserstein barycenter problem as well as a balanced multi-marginal optimal transport problem.

3.1 A Lift to Optimal Transport

We fix a parameter $C > 0$, introduce an additional dummy point \mathfrak{d} and define the augmented space $\tilde{\mathcal{X}} := \mathcal{X} \cup \{\mathfrak{d}\}$ with metric cost

$$\tilde{d}_C^p(x, x') = \begin{cases} d^p(x, x') \wedge C^p, & x, x' \in \mathcal{X}, \\ \frac{C^p}{2}, & x \in \mathcal{X}, x' = \mathfrak{d}, \\ \frac{C^p}{2}, & x = \mathfrak{d}, x' \in \mathcal{X}, \\ 0, & x = x' = \mathfrak{d}. \end{cases} \tag{12}$$

Notably, $\tilde{d}_C: \tilde{\mathcal{X}} \times \tilde{\mathcal{X}} \rightarrow \mathbb{R}_+$ defines a metric on $\tilde{\mathcal{X}}$ [53, Lemma A1]. Consider the subset $\mathcal{M}_+^B(\mathcal{X}) := \{\mu \in \mathcal{M}_+(\mathcal{X}) \mid \mathbb{M}(\mu) \leq B\} \subset \mathcal{M}_+(\mathcal{X})$ of non-negative measures whose total mass is bounded by B . Setting $\tilde{\mu} := \mu + (B - \mathbb{M}(\mu))\delta_{\mathfrak{d}}$, any measure $\mu \in \mathcal{M}_+^B(\mathcal{X})$ defines an augmented measure $\tilde{\mu}$ on $\tilde{\mathcal{X}}$ such that $\mathbb{M}(\tilde{\mu}) = B$. Hence, for two measures $\mu, \nu \in \mathcal{M}_+^B(\mathcal{X})$ we can define the OT problem on $\tilde{\mathcal{X}}$ between their

augmented measures $\tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\mu}, \tilde{\nu})$. In fact, it holds that

$$\text{UOT}_{p,C}(\mu, \nu) = \text{UOT}_{d^p \wedge C^p, C}(\mu, \nu) = \tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\mu}, \tilde{\nu}),$$

where the first equality follows by Lemma 2.1 as for any optimal solution π_C it holds $\pi_C(x, x') = 0$ if $d^p(x, x') > C^p$ and the second follows by [33, Lemma 3.1]. The same equalities remain valid replacing B by an arbitrarily large constant as summarized by the following lemma.

Lemma 3.1 Consider $\mu, \nu \in \mathcal{M}_+^B(\mathcal{X})$ with extended versions $\tilde{\mu}, \tilde{\nu}$. Then for any $a > 0$ it holds that

$$\tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\mu}, \tilde{\nu}) = \tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\mu} + a\delta_\partial, \tilde{\nu} + a\delta_\partial).$$

Proof For $p = 1$, the result is trivial since by duality $\tilde{\text{OT}}_{\tilde{d}_C}(\tilde{\mu}, \tilde{\nu})$ only depends on the difference of the measures. For $p > 1$ we invoke \tilde{d}_C -cyclical monotonicity [72, Thm. 5.10] of any OT plan π and use the property that $\tilde{d}_C^p(x, \partial) = C^p/2$. This yields that $(\partial, \partial) \in \text{supp}(\pi)$ which leads to the desired conclusion. \square

3.2 A Lift to Wasserstein Barycenters

We can also lift the optimization problem defining a (p, C) -barycenter to an equivalent p -Wasserstein barycenter formulation (4). Augmentation of the underlying measures, however, is not straightforward as the total mass of the (p, C) -barycenter is unknown. A first crude upper bound on its total mass leads to a feasible approach.

Lemma 3.2 Consider $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$ and let $F_{p,C}$ be their associated unbalanced Fréchet functional. Then it holds that

$$\arg \min_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu) = \arg \min_{\substack{\mu \in \mathcal{M}_+(\mathcal{Y}) \\ \mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)}} F_{p,C}(\mu).$$

More precisely, any (p, C) -barycenter μ^* of μ_1, \dots, μ_J satisfies $\mathbb{M}(\mu^*) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)$.

Proof Assume first that there exists a measure $\mu \in \mathcal{M}_+(\mathcal{Y})$ such that $\mu = \nu_1 + \nu_2$ where no transport between ν_2 and any μ_i occurs in the optimal solution of $\text{UOT}_{p,C}(\mu, \mu_i)$ for $1 \leq i \leq J$ and it holds $\mathbb{M}(\nu_2) > 0$. Thus it holds

$$F_{p,C}(\mu) = F_{p,C}(\nu_1 + \nu_2) = F_{p,C}(\nu_1) + (C^p/2)\mathbb{M}(\nu_2) > F_{p,C}(\nu_1)$$

and we improve the objective value of μ by removing ν_2 . Hence, let $\mu \in \mathcal{M}_+(\mathcal{Y})$ be any measure such that $\nu_2 \equiv 0$. Consider π_i the optimal solution for $\text{UOT}_{p,C}(\mu, \mu_i)$ for each $1 \leq i \leq J$. Decompose the measure $\mu = \sum_{i=1}^J \tau_i$, where τ_i is the mass of μ

transported to μ_i according to π_i and which is not yet included in any τ_j for $j < i$. Clearly, $\mathbb{M}(\mu) = \sum_{i=1}^J \mathbb{M}(\tau_i) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)$ and we conclude that

$$\min_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu) = \min_{\substack{\mu \in \mathcal{M}_+(\mathcal{Y}) \\ \mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)}} F_{p,C}(\mu).$$

By our first considerations the claim follows. □

Given the upper bound on the total mass of any (p, C) -barycenter at our disposal we can formulate a lift of the (p, C) -barycenter problem to a related p -Wasserstein barycenter problem. For this, let $\tilde{\mathcal{Y}} := \mathcal{Y} \cup \{\mathfrak{d}\}$ endowed with the metric \tilde{d}_C in (12) (replace \mathcal{X} by \mathcal{Y} and recall that $\mathcal{X} \subset \mathcal{Y}$) and augment the measures μ_1, \dots, μ_J to $\tilde{\mu}_1, \dots, \tilde{\mu}_J$ where $\tilde{\mu}_i = \mu_i + \sum_{j \neq i} \mathbb{M}(\mu_j) \delta_{\mathfrak{d}}$ for $1 \leq i \leq J$. In particular, $\mathbb{M}(\tilde{\mu}_i) = \sum_{j=1}^J \mathbb{M}(\mu_j)$ and we can define the augmented p -Fréchet functional

$$\tilde{F}_{p,C}(\mu) := \frac{1}{J} \sum_{i=1}^J \tilde{\text{OT}}_{\tilde{d}_C}^p(\tilde{\mu}_i, \mu),$$

where by definition $\tilde{F}_{p,C}$ is restricted to measures μ with mass $\mathbb{M}(\mu) = \sum_{i=1}^J \mathbb{M}(\mu_i)$.

Lemma 3.3 *For $1 \leq i \leq J$ consider measures $\mu_i \in \mathcal{M}_+(\mathcal{X})$ and their augmented versions $\tilde{\mu}_i := \mu_i + \sum_{j \neq i} \mathbb{M}(\mu_j) \delta_{\mathfrak{d}}$, respectively. Then it holds that*

$$F_{p,C}(\mu) = \tilde{F}_{p,C} \left(\mu + \left(\sum_{i=1}^J \mathbb{M}(\mu_i) - \mathbb{M}(\mu) \right) \delta_{\mathfrak{d}} \right)$$

for all $\mu \in \mathcal{M}_+(\mathcal{Y})$ such that $\mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)$ and in particular

$$\min_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu) = \min_{\substack{\mu \in \mathcal{M}_+(\tilde{\mathcal{Y}}) \\ \mathbb{M}(\mu) = \sum_{i=1}^J \mathbb{M}(\mu_i)}} \tilde{F}_{p,C}(\mu).$$

The proof of this Lemma is given in Appendix A.1.

Remark 3.4 (Optimal (p, C) -barycenters) Lemma 3.3 states that the optimal objective value for the (p, C) -barycenter is equal the related p -Wasserstein barycenter problem on the augmented space. In particular, the proof also reveals that if $\tilde{\mu}^*$ is a p -Wasserstein barycenter for the augmented measures $\tilde{\mu}_1, \dots, \tilde{\mu}_J$ then $\mu^* := \tilde{\mu}^* - \tilde{\mu}^*(\mathfrak{d}) \delta_{\mathfrak{d}}$ is a (p, C) -barycenter for the measures μ_1, \dots, μ_J . Vice versa, if μ^* is a (p, C) -barycenter for the measures μ_1, \dots, μ_J then $\tilde{\mu}^* := \mu^* + \left(\sum_{i=1}^J \mathbb{M}(\mu_i) - \mathbb{M}(\mu^*) \right) \delta_{\mathfrak{d}}$ is a p -Wasserstein barycenter for the augmented measures $\tilde{\mu}_1, \dots, \tilde{\mu}_J$.

3.3 A Lift to Multi-Marginal Optimal Transport

On the augmented space $\tilde{\mathcal{Y}} := \mathcal{Y} \cup \{\mathfrak{d}\}$ equipped with metric \tilde{d}_C in (12), we define for $p \geq 1$ and $J \in \mathbb{N}$ a Borel barycenter application $\tilde{T}_C^{J,p} : \tilde{\mathcal{Y}}^J \rightarrow \tilde{\mathcal{Y}}$ that takes as input $(y_1, \dots, y_J) \in \tilde{\mathcal{Y}}$ and outputs any minimizer $y \in \tilde{\mathcal{Y}}$ of the function

$$f(y) = \sum_{i=1}^J \tilde{d}_C^p(y_i, y).$$

Of particular interest to us is the barycentric application restricted to inputs from $\tilde{\mathcal{X}}$. However, we collect some of its key properties for general input $(y_1, \dots, y_J) \in \tilde{\mathcal{Y}}^J$. For this, we define the index set

$$\mathcal{B}(y_1, \dots, y_J) := \{i \mid y_i = \mathfrak{d}, 1 \leq i \leq J\}.$$

If clear from the context, then the dependence on y_1, \dots, y_J is suppressed and the set is simply denoted as \mathcal{B} .

Lemma 3.5 *Fix some parameter $C > 0$ and consider the space $\tilde{\mathcal{Y}}$ with metric \tilde{d}_C as defined in (12). For points $(y_1, \dots, y_J) \in \tilde{\mathcal{Y}}^J$ it holds that*

- (i) $\tilde{T}_C^{J,p}(y_1, \dots, y_J) = \mathfrak{d}$ if and only if $\sum_{i \notin \mathcal{B}} \tilde{d}_C^p(y_i, y) \geq (J - 2|\mathcal{B}|)C^p/2$ for any $y \in \tilde{\mathcal{Y}}$. In particular, if strict inequality holds then $\tilde{T}_C^{J,p}(y_1, \dots, y_J) = \mathfrak{d}$ is unique.
- (ii) If $2|\mathcal{B}| \geq J$ then it holds $\tilde{T}_C^{J,p}(y_1, \dots, y_J) = \mathfrak{d}$ with uniqueness if $2|\mathcal{B}| > J$.
- (iii) If $\tilde{T}_C^{J,p}(y_1, \dots, y_J) \neq \mathfrak{d}$ then it holds

$$\tilde{T}_C^{J,p}(y_1, \dots, y_J) = \arg \min_{y \in \mathcal{Y}} \sum_{i \notin \mathcal{B}} \tilde{d}_C^p(y_i, y).$$

- (iv) If $C > 2^{\frac{1}{p}} \text{diam}(\mathcal{Y})$, then for any points $y_1, \dots, y_J \in \mathcal{Y}$ with $|\mathcal{B}| = 0$ it holds that $\tilde{T}_C^{J,p}(y_1, \dots, y_J) = T^{J,p}(y_1, \dots, y_J)$ where the latter one is defined with respect to the usual metric d^p on \mathcal{Y} .

A proof of this result is provided in Appendix A.1. Lemma 3.5 allows to characterize the centroid sets of the augmented measures $\tilde{\mu}_1, \dots, \tilde{\mu}_J$ defined as

$$\tilde{C}_{KR}(J, p, C) := \left\{ y \in \tilde{\mathcal{Y}} \mid y = \tilde{T}_C^{J,p}(x_1, \dots, x_J), x_i \in \text{supp}(\tilde{\mu}_i); \right. \\ \left. d^p(y, x_i) \leq C^p \forall x_i \neq \mathfrak{d} \right\}. \tag{13}$$

Remark 3.6 We point out that computing $\tilde{T}_C^{J,p}$ is in general a difficult optimisation problem. While for squared euclidean distance, computing the barycentric application simply amounts to taking the mean of the x_i , even on the non-augmented space,

there are no closed form solutions available for most choices of distances and values of p . This problem is exacerbated by the truncation of the distance \tilde{d} at C^p (as also pointed out in [53]), since it implies that disregarding a certain subset of points and just computing the barycenter with respect to the remaining x_i might in fact be optimal. However, initially it is not clear which x_i to choose, turning this into a difficult combinatorial problem.

Recall that for any measure μ its support is contained in \mathcal{X} a subset of \mathcal{Y} . The augmented measure $\tilde{\mu}$ is extended by an additional support point at $\{\mathfrak{d}\}$. In particular, while the centroid set is a subset of $\tilde{\mathcal{Y}}$ it only depends on the support of the measures $\tilde{\mu}_i$ contained in $\tilde{\mathcal{X}} := \mathcal{X} \cup \{\mathfrak{d}\}$.

Corollary 3.7 *For the centroid sets of the augmented measures $\tilde{\mu}_i \in \mathcal{M}_+(\tilde{\mathcal{X}})$ with $1 \leq i \leq J$ it holds*

$$\tilde{\mathcal{C}}_{KR}(J, p, C) \subset \mathcal{C}_{KR}(J, p, C) \cup \{\mathfrak{d}\} \subset \mathcal{C}_{KR}(J, p) \cup \{\mathfrak{d}\}.$$

Proof The first inclusion follows by statements (i) and (iii) in Lemma 3.5 and the observation that $|\mathcal{B}| = J - L$. The second by applying $\mathcal{C}_{KR}(J, p, C) \subset \mathcal{C}_{KR}(J, p)$. □

Remark 3.8 One could define $\mathcal{C}_{KR}(J, p, C)$ in terms of \tilde{d}_C instead of d to obtain equality in the first inclusion. Replacing $T^{L,p}$ by $\tilde{T}_C^{L,p}$ in the definition of the centroid set would not alter any of the related proofs and yield slightly sharper control on the support of (p, C) -barycenter. However, as we consider the given definition to be more intuitive, we omit this improvement in the statement of the theorem.

Let $\Pi(\tilde{\mu}_1, \dots, \tilde{\mu}_J)$ be the set of measures on $\tilde{\mathcal{Y}} \times \dots \times \tilde{\mathcal{Y}}$ whose i -th marginal is equal to $\tilde{\mu}_i$ for all $1 \leq i \leq J$. We refer to the elements of this set as *multi-couplings* of $\tilde{\mu}_1, \dots, \tilde{\mu}_J$. For $p \geq 1$ define the *augmented multi-marginal transport problem* as

$$\min_{\pi \in \Pi(\tilde{\mu}_1, \dots, \tilde{\mu}_J)} \int_{\tilde{\mathcal{Y}}^J} c_{p,C}(y_1, \dots, y_J) \pi(dy_1, \dots, dy_J), \tag{14}$$

where

$$c_{p,C}(y_1, \dots, y_J) := \frac{1}{J} \sum_{i=1}^J \tilde{d}_C^p \left(y_i, \tilde{T}_C^{J,p}(y_1, \dots, y_J) \right).$$

The relation between the augmented multi-marginal transport formulation (14) and the (p, C) -barycenter is as follows.

Proposition 3.9 *Let $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$ and $\tilde{\mu}_1, \dots, \tilde{\mu}_J$ be their augmented counterparts. If $\pi \in \Pi(\tilde{\mu}_1, \dots, \tilde{\mu}_J)$ is a solution to the augmented multi-marginal problem (14), then the measure $\mu^* := (\tilde{T}_C^{J,p} \# \pi)|_{\mathcal{Y}} \in \mathcal{M}_+(\mathcal{Y})$ is a (p, C) -barycenter of the measures μ_1, \dots, μ_J , where $\tilde{T}_C^{J,p} \# \pi$ denotes the pushforward of π under $\tilde{T}_C^{J,p}$. Moreover,*

for every (p, C) -barycenter μ^* , there exists a solution π to the augmented multi-marginal transport problem, such that

$$\mu^* + \left(\sum_{i=1}^J \mathbb{M}(\mu_i) - \mathbb{M}(\mu^*) \right) \delta_{\mathfrak{d}} = \tilde{T}_C^{J,p} \# \pi.$$

In particular, it holds that

$$\min_{\pi \in \Pi(\tilde{\mu}_1, \dots, \tilde{\mu}_J)} \int_{\tilde{\mathcal{Y}}^J} c_{p,C}(y_1, \dots, y_J) \pi(dy_1, \dots, dy_J) = \inf_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu).$$

The proof follows straightforwardly along the lines of related statements for the multi-marginal optimal transport problem (43, Theorem 8; 49, Lemma 8 or 54, Proposition 3.1.2). This correspondence between the (p, C) -barycenter problem and a balanced multi-marginal optimal transport serves as one of the key components in the proof of Theorem 2.5.

4 Computational Issues and Numerical Experiments

We present approaches to compute the (p, C) -barycenter problem by solving related OT problems. Based on this, we investigate the performance of the Wasserstein and (p, C) -barycenters on multiple synthetic datasets. For reference, we also report on results for two related concepts of unbalanced barycenters (UBCs), namely the Gaussian-Hellinger–Kantorovich and Wasserstein–Fisher–Rao barycenter.

4.1 Algorithms

Theorem 2.5 and Proposition 3.9 both allow to pose the augmented problem (recall Sect. 3) as a linear program and using Lemma 3.3 one can obtain a solution to the original problem by solving the augmented one. Using any linear program solver this enables the direct computation of an exact solution of this problem. However, the number of variables in this approach scales as the size of $\mathcal{C}_{KR}(J, p, C)$ and hence it turns out to be infeasible already for relatively small instance sizes. To compute (p, C) -barycenters at larger scales we revisit iterative methods to solve the (balanced) Wasserstein barycenter problem and give instructions how to use modifications of them to compute (p, C) -barycenters. In particular, we detail a multi-scale method which solves successive fixed-support (p, C) -barycenter LPs on increasingly refined support sets. This provides a meta-framework to adjust state-of-the-art solvers for the Wasserstein barycenter for (p, C) -barycenter computations.

To construct the augmented problem we add the dummy point \mathfrak{d} to the support of the μ_i 's, while setting its distance to all other locations to be $C^p/2$. Note, that by Lemma 2.1 and Lemma 3.1 the truncation of \tilde{d} at C^p can be omitted if $\mathbb{M}(\tilde{\mu}_i) >$

$3 \max_{i=1, \dots, J} \mathbb{M}(\mu_i)$. If this is not the case, we can enforce it by adding additional mass at δ in all augmented measures without changing the optimal value.

4.1.1 LP-Formulation for the (p, C)-Barycenter

Using property (i) from Theorem 2.5, we can rewrite the augmented (p, C) -barycenter problem as a linear program similarly to the usual p -Wasserstein barycenter problem (4). However, compared to the latter one, we replace the standard centroid set $\mathcal{C}_W(J, p)$ from (11), by the centroid set $\tilde{\mathcal{C}}_{KR}(J, p, C)$ of the augmented measures from (13). This yields

$$\begin{aligned} \min_{\pi^{(1)}, \dots, \pi^{(J)}, a} \quad & \frac{1}{J} \sum_{i=1}^J |\tilde{\mathcal{C}}_{KR}(J, p, C)| \sum_{j=1}^{|\tilde{\mathcal{C}}_{KR}(J, p, C)|} \sum_{k=1}^{M_i} \pi_{jk}^{(i)} c_{jk}^i \\ \text{s.t.} \quad & \sum_{k=1}^{M_i} \pi_{jk}^{(i)} = a_j, \quad \forall i = 1, \dots, J, \forall j = 1, \dots, |\tilde{\mathcal{C}}_{KR}(J, p, C)|, \\ & \sum_{j=1}^{|\tilde{\mathcal{C}}_{KR}(J, p, C)|} \pi_{jk}^{(i)} = b_k^i, \quad \forall i = 1, \dots, J, \forall k = 1, \dots, M_i, \\ & \pi_{jk}^{(i)} \geq 0 \quad \forall i = 1, \dots, J, \forall j = 1, \dots, |\tilde{\mathcal{C}}_{KR}(J, p, C)|, \\ & \quad \quad \quad \forall k = 1, \dots, M_i, \end{aligned}$$

where $M_i = |\tilde{\mathcal{X}}_i|$ is the cardinality of the support of the augmented measure $\tilde{\mu}_i$. Here, c_{jk}^i denotes the distance between the j -th point of $|\tilde{\mathcal{C}}_{KR}(J, p, C)|$ and the k -th point in the support of $\tilde{m}u_i$, while b^i is the vector of masses corresponding to $\tilde{\mu}_i$. For practical purposes it may be advantageous to solve the multi-marginal problem instead of the (p, C) -barycenter problem. This changes the number of variables from $|\tilde{\mathcal{C}}_{KR}(J, p, C)|(1 + \sum_{i=1}^J M_i)$ to $\prod_{i=1}^J M_i$ and the number of constraints from $J|\tilde{\mathcal{C}}_{KR}(J, p, C)| + \sum_{i=1}^J M_i$ to $\sum_{i=1}^J M_i$. Depending on the value of C , and hence the cardinality of $\tilde{\mathcal{C}}_{KR}(J, p, C)$, it is possible to pick the problem with the smaller complexity.

While this formulation is appealing for proving theoretical statements as provided in Theorem 2.5, it quickly becomes computationally infeasible even for small scale problems as the number of variables in the LP grows potentially as $\prod M_i$. However, it still enables exact computations of (p, C) -barycenters for small scale examples, which is currently impossible for general UBCs. Though, while there has been some recent advancement for the 2-Wasserstein barycenter in special cases [2] these LP-based algorithms ultimately do not scale to large instance sizes.

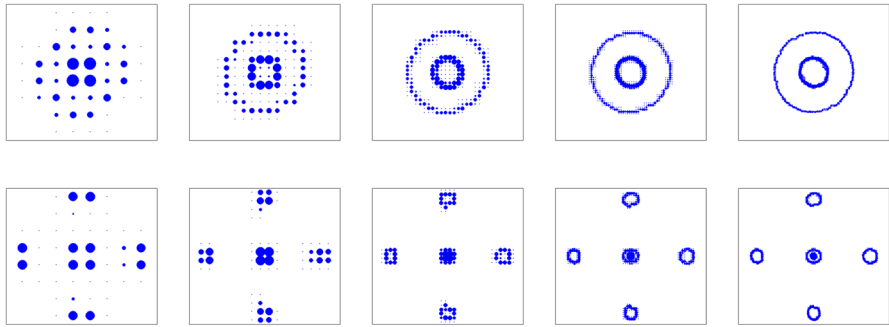


Fig. 6 An illustration of the multi-scale approach on two different datasets. The fixed-support solutions are shown on grids of the sizes 8×8 , 16×16 , 32×32 , 64×64 and 128×128 increasing from left to right. The corresponding run-times on a single core of an Intel Core *i7* 12700K in the first/second row were 2.5/5 seconds, 14/16 seconds, 145/42 seconds, 13/3 minutes and 143/22 minutes. Top: The dataset of nested ellipses from Fig. 7. Bottom: The dataset of ellipses with clustered support structure from Fig. 2

4.2 Iterative Algorithms and the Multi-scale Approach

For the Wasserstein barycenter, iterative methods computing approximate barycenters, with a per iterations complexity only linear in the number of measures, enjoy great popularity. Most well known is the *fixed-support Wasserstein barycenter* [28, 46, 75] approach, aiming to find the best approximation of the barycenter on a pre-specified support set, for which a variety of methods is available. We utilise this fixed-support approach for the augmented (p, C) -barycenter problem by adding the dummy point δ to the given support and constructing the cost as described above. This yields a meta-framework which allows to employ fixed-support Wasserstein barycenter algorithms for fixed-support (p, C) -barycenter computation. One can also modify more general *free support* methods [19, 28, 48], which usually alternate between updating the support set of the barycenter and its weights on this set, to provide approximate (p, C) -barycenters. However, the necessary position updates usually explicitly or implicitly rely on being able to compute the barycentric application $\tilde{T}^{J,p}$ efficiently. Recalling Remark 3.6, this is in general not tractable for the augmented problem, which severely hinders the use of these approaches. Thus, it is tempting to avoid these issues by approximating \mathcal{Y} with a large finite space, i.e., by taking a grid of high-resolution, and solving the fixed support (p, C) -barycenter problem on this set. However, solving the fixed-support problem on this large space requires significant computational effort. We advocate an alternative by adapting the ideas of multi-scale methods for the Wasserstein distance/barycenter [30, 50, 64] to the (p, C) -barycenter setting. The idea of this approach is to start with a coarse version of the problem and then successively solve refined problems, while using the knowledge of the coarse solution to reduce the complexity of the finer ones.

Thus, we initialise the support set of the barycenter as a fixed grid of size $K_1 \times \dots \times K_d$ in \mathbb{R}^d . In the j -th step of the algorithm, after solving the fixed-support problem, we remove the grid points which have zero mass and replace the remaining ones with its 2^d closest points in a refined version of the original grid of size $2^j K_1 \times \dots \times 2^j K_d$. This

can be understood as solving the fixed-support problem on successively finer grids, while incorporating information provided by having already solved a coarser solution of the problem. We terminate the method once a pre-specified resolution has been reached. This allows to obtain fixed-support approximation of the (p, C) -barycenter on fine grids without having to optimise over the full support set.

We point out that this approach, while inspired by multi-scale approaches is more closely related to the formerly mentioned free-support methods. As such it does in general not yield a globally optimal fixed-support (p, C) -barycenter at the finest resolution. Instead it converges to a local minimum of the unbalanced Fréchet functional depending on the resolution of the initial grid. This is a common problem among alternating procedures for the free-support barycenter problem and can be attributed to the fact that the Fréchet functional is non-convex in the support locations of the measures. However, we stress that with this approach we observe reasonable approximations of the (p, C) -barycenter while avoiding the inherent problems of generalising usual position update procedures discussed above. In particular, we do not have to solve the $T_C^{\tilde{J}, p}$ barycenter problem at any point. Additionally, we note that the initial grid size should be chosen at least fine enough that the distance between two adjacent grid points is smaller than C . Otherwise it is possible that support points lying between two grid points, having distance larger C to both, are not accounted for. For a visual illustration of the algorithm we refer to Fig. 6.

4.3 Synthetic Data Simulations

We test the performance of the (p, C) -barycenter as a data analytic tool compared to the usual p -Wasserstein barycenter on a multitude of datasets. We base our computations on the MAAIPM method [28], which allows for high-precision approximations of barycenters up to moderate data sizes. The algorithm has been deployed to solve the fixed-support (p, C) -barycenter problems arising in the multi-scale method detailed above. For all experiments, the initial grid size as been set to 16×16 and the refinement is terminated at a gridsize of 128×128 . Values below 10^{-5} have been considered as zero for the purposes of grid refinement. All experiments have been carried out on a single core of an Intel Core *i7 12700K*. Implementations of our used method and some alternatives can be found as part of the R-package *WSGeometry* (on CRAN).

Mismatched Shapes

This first set of examples mainly serves as starting point to illustrate improved performance of the (p, C) -barycenter compared to the p -Wasserstein barycenter. A prototypical benchmark for the p -Wasserstein barycenter are two nested ellipses as popularized in Cuturi and Doucet [19]. For our example of nested ellipses, we assume that the support of each measure consists of nested ellipses, but the number of ellipses varies between the individual underlying measures. Specifically, we assume that for each μ_i the number of ellipses is uniformly random in $\{1, 2, 3\}$ and that each ellipse is discretised onto M support points with unit mass, respectively. This can be seen

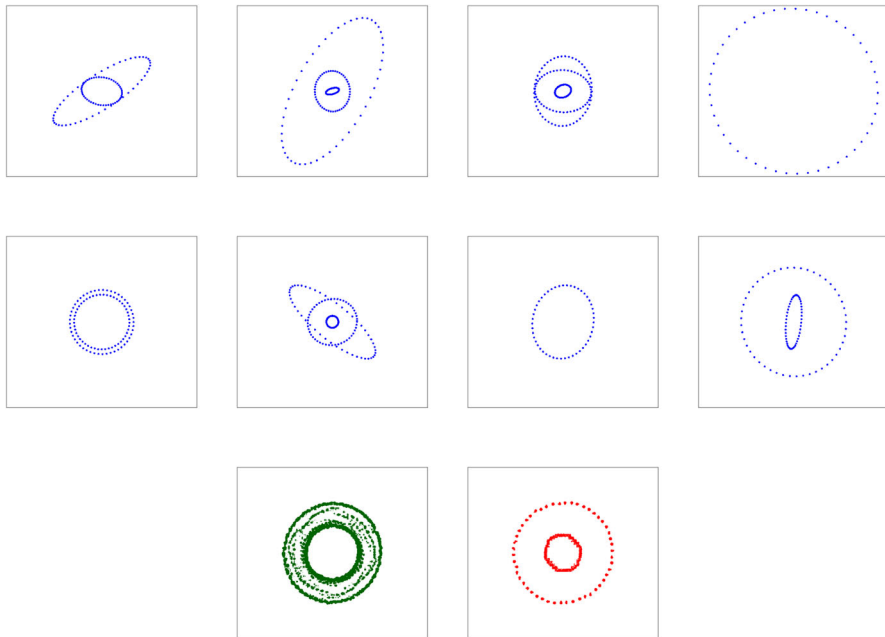


Fig. 7 An excerpt of a dataset of $N = 100$ discretized ellipses. Each measure contains between 1 and 3 ellipses with equal probability. Each ellipse consists of 50 points with mass 1 in $[0, 1]^2$. Left: In darkgreen the 2-Wasserstein barycenter, where all measures are normalized to be probability measures (runtime about 8 hours). Right: In red, the $(2, 1.5)$ -barycenter (runtime about 30 minutes) (Color figure online)

in Fig. 7. We observe that while the p -Wasserstein barycenter recovers the elliptic shape of the underlying measures, it fails to produce distinct ellipses and instead produces something akin to a ring. In contrast, the (p, C) -barycenter yields two distinct ellipses, which coincides with the expected number of ellipses in one of the measures. This aligns well with intuition that the (p, C) -barycenter will simply disregard any additional structures which are not present in a sufficient amount of underlying measures. In contrast, the p -Wasserstein barycenter does not allow for this flexibility which enforces additional support points.

Local Scale Cluster Detection

Recall the setting of Fig. 2. In the following class of examples, we are interested in datasets which possesses a natural cluster structure. Let $B_1, \dots, B_R \subset \mathbb{R}^D$ be convex, disjoint sets and assume that $\text{supp}(\mu_i) \subset \cup_{r=1}^R B_r$ for all $i = 1, \dots, J$. If the diameter of all B_r is bounded from above by C and that the distance between each two B_r, B_s is at least $2^{1/p}C$, then Lemma 2.7 guarantees that the (p, C) -barycenter detects all of the R clusters in which at least $J/2$ measures have positive mass. In particular, by Theorem 2.5 (v) the (p, C) -barycenter will have mass in all of those clusters. Intuitively, this setting is reasonable if, for instance, it is already known that any interactions between support points of different measures are limited to scales

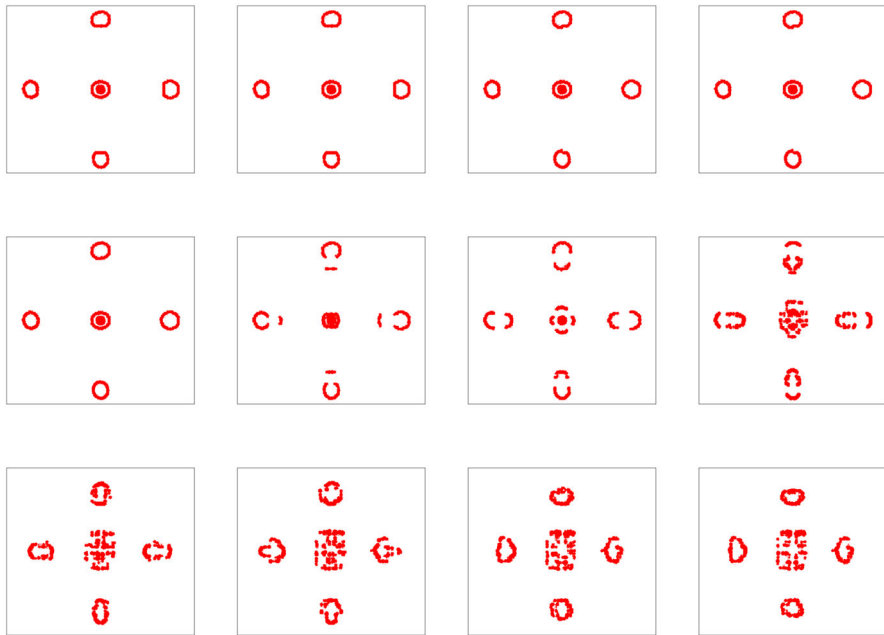


Fig. 8 The $(2, C)$ -barycenters for the measures in Fig. 2 for different values of C . From top-left to bottom-right the values of C /runtime in minutes are equal to $0.1/28$, $0.15/28$, $0.2/26$, $0.25/28$, $0.275/27$, $0.3/33$, $0.35/59$, $0.4/132$, $0.45/156$, $0.5/160$, $0.55/160$, $0.6/155$, respectively

below a certain threshold, which should then be chosen as C . The lower bound on the inter-cluster distance ensures that any pair of two clusters is well-separated, ensuring that it is always possible to distinguish between two different clusters, as they can not be arbitrarily close to each other.

In Fig. 2 the p -Wasserstein barycenter completely fails to capture the geometric data structure. Most of its mass is between the clusters and the outer clusters have nearly no mass. Moreover, the elliptic structure within each cluster is clearly not captured. In contrast, the (p, C) -barycenter not only captures all clusters, it also distinguishes between the difference in intensity (expected number of ellipses) in the clusters, matching the theoretical guarantees of Lemma 2.7. We stress that for this example the choice of C is of particular importance. If we choose C too large, the (p, C) -barycenter will fail to recover the data's support structure (for an illustration of the (p, C) -barycenter in this example over a range of values of C see Fig. 8). Consequently, it is crucial to choose C appropriately. In this example, the barycenter appears to be stable and detect all clusters for $C \in [0.1, 0.275]$. Notably, if the locations of the clusters are already known, this setting also allows for parallel computations of the (p, C) -barycenter, where the problems are solved separately on each cluster and recombined at the end (Lemma 2.7).

Randomly Distorted Measures

In a statistical context it is important to investigate the stability of the (p, C) -barycenter under random distortions. We fix a reference measure μ_0 on \mathbb{R}^d and generate a set of measures by random modifications of μ_0 . We then attempt to recover μ_0 by computing the p -Wasserstein and (p, C) -barycenter of these measures, respectively.

In the following, let $B(p)$ denote a Bernoulli random variable with mean p , $Poi(\lambda)$ a Poisson distribution with mean λ and $U[a, b]$ a uniform distribution on $[a, b]$. We generate μ_1, \dots, μ_J as follows:

For $i = 1, \dots, J$ initialise $\mu_i = \mu_0$, then successively modify μ_i based on the four following steps.

- (i) **Point Deletion:** Fix $p_{del} \in [0, 1]$ and $\lambda_{del} \in \mathbb{R}_+$. We draw a $Ber(p_{del})$ random variable. If it takes the value 1, then we draw $D \sim Poi(\lambda_{del})$ and select $\min(D, |\text{supp}(\mu_0)|)$ points in the support of μ_0 uniformly by drawing without replacement. These points (and their mass) are not contained in μ_i , since they have been deleted.
- (ii) **Point Addition:** We fix parameters $p_{add} \in [0, 1], \lambda_{add} \in \mathbb{R}_+, m_{add} \in \mathbb{R}^2, \sigma_{add} \in \mathbb{R}^{2 \times 2}, u_0, u_1 \in \mathbb{R}$. Draw a $Ber(p_{add})$ random variable. If it takes the value 1, draw a $Poi(\lambda_{add})$ random variable α . Then, generate α random variables following a normal distribution with mean m_{add} and covariance matrix σ_{add} . Add these support points to μ_i , where the weight of each of these points is determined by independent $U[u_0, u_1]$ random variables.
- (iii) **Position Change:** Fix parameters $a_1, a_2, b_1, b_2 \in \mathbb{R}$ with $a_1 \leq b_1$ and $a_2 \leq b_2$. For each x_0 in the support of μ_i , we draw a $U([a_1, b_1] \times [a_2, b_2])$ random variable and shift the position of x_0 by it.
- (iv) **Weight Change:** Fix parameters $l, u \in \mathbb{R}$ with $l \leq u$. For each support point x_0 of μ_0 with weight w_0 , we draw a $U[l, u]$ random variable U and change the weight of x_0 in μ_i to be $w_0 + U$.

An example of this setting can be seen in Fig. 9. Comparing the two barycenters displayed there to the original measure reveals that, while the rough shape of the 2-Wasserstein barycenter is correct, its mass is spread out over a larger area and it has a significantly larger number of support points. Since all measures have been normalised, we have also lost all information on the mass of μ_0 . Contrary to that, the (p, C) -barycenter retrieves the original measures recovering the location and number of the of support points closely. Additionally, it also has a mass which only deviates from the original mass by about 0.23%. If one is only interested in recovering the general shape of the data, both approaches provide comparable performance. However, if the measures total mass and more detailed support structure are of importance the (p, C) -barycenter appears to be preferable.

Total Mass Intensity

While the p -Wasserstein barycenter of J probability measures has mass one, the mass of the (p, C) -barycenter depends on C as well as the geometry of the measures $\mu_1, \dots, \mu_J \in \mathcal{M}_+(\mathcal{X})$. Exact values for the mass of a (p, C) -barycenter without




Fig. 9 An excerpt from a dataset of $N = 100$ noisy nested ellipses supported in $[0, 1]^2$. The parameters are $p_{del} = 1/3, \lambda_{del} = 75, p_{add} = 1/3, \lambda_{add} = 25, m_{add} = (0.5, 0.5)^T, \sigma_{add} = 0.15I_2, u_0 = 0.9, u_1 = 1.1, a_1 = a_2 = -0.025, b_1 = b_2 = 0.025, l = u = 0.1$. Left: The 2-Wasserstein barycenter in dark green (runtime about 4 hours). Center: The original measure μ_0 (black). Right: The $(2, 1.5)$ -barycenter in red (runtime about 20 min) (Color figure online)




Fig. 10 The mass of a (p, C) -barycenter for three sets of measures relative to the median of the total mass intensities of these measures. The green line corresponds to $J = 25$ measures from the same class as considered in Fig. 2. The red line corresponds to the same measures where the four outer clusters have been moved closer to the central one, such that their distance has been halved. The blue line corresponds to $J = 5$ measures with the same cluster structure as in Fig. 2, where the total number of ellipses in all clusters is fixed to be equal to four for all J measures (Color figure online)

detailed computations, are only available in the limiting scenarios where C is extremely small or large relative to the other distances in \mathcal{X} . For the former, we know by Theorem 2.5 (v) that the barycenter has mass zero for disjoint measures and for the latter, Theorem 2.5 (vi) yields that there exists a (p, C) -barycenter with total mass intensity equal to the median of $\mathbb{M}(\mu_1), \dots, \mathbb{M}(\mu_J)$. For intermediate values of C , Theorem 2.5 (i) yields the upper bound by $2J^{-1} \sum_{i=1}^J \mathbb{M}(\mu_i)$. To highlight some possible behaviours of the total mass intensity of (p, C) -barycenter we consider three specific examples in Fig. 10. We note that in all three cases at about $C = 0.6$ the mass

of the barycenters is at the median of their respective μ_1, \dots, μ_J and does no longer change with increasing C . This is significantly smaller than the requirement in Theorem 2.5 (vi), which underlines the fact that while in the worst case, this lower bound is sharp, in many examples the total mass of the (p, C) -barycenter stabilises significantly earlier. Moreover, none of the three curves is monotone. Instead the total mass of the barycenter is increasing up to a certain point, after which it decreases until it reaches the median of the masses. This makes intuitive sense, as the measures are disjoint, thus for small C the barycenter is empty and starts to grow in mass quickly as the points within the clusters can be matched. In particular, the differences in intensity between clusters might lead to a total mass over the median $\mathbb{M}(\mu_1), \dots, \mathbb{M}(\mu_J)$, as by Lemma 2.7 the total mass intensity of the (p, C) -barycenter is $\sum_{r=1}^R \text{med}(\mathbb{M}(\mu_{1|B_r}), \dots, \mathbb{M}(\mu_{J|B_r}))$, where B_1, \dots, B_5 denote the respective cluster locations. For larger C these clusters start to merge and support points between the clusters reduce the total mass. In particular, these points can be seen clearly in the plot. Up until about $C = 0.1$, which is the cluster size, the mass of the barycenters rises sharply, before stabilising until the intercluster distance is reached. This is about 0.3 for the green and blue lines and about 0.15 for the red line (since the measures in this example are generated by halving the intercluster distance from the green one). This behaviour highlights the sensitivity of the mass of the (p, C) -barycenter to the geometry of the measures. It is therefore impossible to infer the total mass of the (p, C) -barycenter from the magnitude of C alone without accounting for the specific measures. However, analysing the structural properties of the support sets of the measures might provide a good indication at what values of C changes in drastic behaviour of the total mass are to be expected.

4.4 Comparison with Related Unbalanced Barycenter Concepts

We compare the (p, C) -barycenter with two alternative UBC approaches.

The Gaussian-Hellinger-Kantorovich Barycenter: This example falls in the general framework of optimal entropy transport problems. Measuring deviation between a feasible solution and the input marginals is carried out via the *Kullback-Leibler divergence* defined for $\mu \ll \nu$ ⁶ as

$$KL(\mu, \nu) = \sum_{x \in X} \mu(x) \log \left(\frac{\mu(x)}{\nu(x)} \right).$$

If $\mu \not\ll \nu$ the value of KL is set to be $+\infty$. For a parameter $\lambda > 0$, the *Gaussian-Hellinger-Kantorovich Distance* [45] is defined as

$$GHK_\lambda(\mu, \nu) = \min_{\pi \in \mathcal{M}_+(\mathcal{X} \times \mathcal{X})} \sum_{x, x' \in \mathcal{X}} d^2(x, x') \pi(x, x') + \lambda KL(\mu, \pi_1) + \lambda KL(\nu, \pi_2),$$

⁶ A measure $\mu \in \mathcal{M}_+(\mathcal{X})$ is said to be *absolutely continuous* (denoted $\mu \ll \nu$) with respect to another measure $\nu \in \mathcal{M}_+(\mathcal{X})$ if $\nu(A) = 0$ implies $\mu(A) = 0$ for any measurable set A .




Fig. 11 Comparison of the three unbalanced barycenters when varying their parameter. All measures are supported on an equidistant 64×64 grid in $[0, 1]^2$. First row: The four underlying measures. Second row: The Gaussian-Hellinger-Kantorovich barycenter for $\lambda = 0.01, 0.15, \dots, 1.83, 1.97$. Third row: The Hellinger-Kantorovich barycenter for $\sigma = 0.01, 0.08, \dots, 0.64, 0.71$. Fourth row: The KR barycenter for $C = 0.01, 0.08, \dots, 0.64, 0.71$

where π_1 and π_2 denote the respective marginals of π . The GHK_λ barycenter is defined as

$$\arg \min_{\mu \in \mathcal{M}_+(\mathcal{Y})} \sum_{i=1}^N GHK_\lambda(\mu_i, \mu).$$

The Hellinger-Kantorovich Barycenter: The Hellinger-Kantorovich distance, also known as Wasserstein-Fisher-Rao distance [14, 45], is closely related to the Gaussian-Hellinger-Kantorovich distance. For fixed parameter $\sigma \in (0, \pi/2]$, referred to as the *cut-locus*, it is defined as

$$\begin{aligned} HK_\sigma(\mu, \nu) = & \min_{\pi \in \mathcal{M}_+(\mathcal{X} \times \mathcal{X})} \sum_{x, x' \in \mathcal{X}} (-\log(\cos_\sigma^2(d(x, y)))) \pi(x, x') \\ & + KL(\mu, \pi_1) + KL(\nu, \pi_2), \end{aligned}$$

where $\cos_\sigma : z \mapsto \cos(\min(z, \sigma))$. For a fixed cut-off locus σ , the HK_σ barycenter is defined as

$$\arg \min_{\mu \in \mathcal{M}_+(\mathcal{Y})} \sum_{i=1}^N HK_\sigma(\mu_i, \mu).$$

Comparing the barycenters: As the resulting barycenters vary significantly in all three cases, depending on the parameters C, λ, σ , we compare their behaviour upon change of parameter.

As a simple example, we consider four measures supported on subsets of a grid on $[0, 1]^2$, displayed in Fig. 11. To ensure fair comparison, we deploy the same method based on the general scaling method [15] to approximate the UBC in all three cases. However, we point out that this implies disregarding the ambient space and instead



Fig. 12 Images displaying the underlying measures used for barycenter computation in Fig. 13. Each row corresponds to a dataset of ten elements of the classical MNIST dataset which have been randomly rescaled and shifted within a 50×50 grid in $[0, 1]^2$. Their total mass intensities have not been normalised

taking the minimum over all positive measures supported on a prespecified grid in $[0, 1]^2$.

For high parameter values all three approaches yield similar results. This is, of course, to be expected, since these distances interpolate between p -Wasserstein distance and total variation/Kullback-Leibler distance and large parameters correspond to a setting being close to the Wasserstein distance. The KR barycenter has mass zero for small choice of C by Theorem 2.5 (iv), since the four measures have disjoint support. After reaching a threshold of $C \approx 0.1$, the mass in the $(2, C)$ -barycenter starts to increase as mass is added in the center of the unit square until at $C \approx 0.3$ the mass of an individual data measure is reached.

For small λ the GHK_λ barycenter has small mass and its support is close to that of a linear mean of the four measures, though the total mass intensity is significantly lower than for the original measures. With increasing λ the mass starts to increase and to smear into the middle of the unit square, until a large square, encompassing all four data supports, is formed. After this point increasing λ causes the square to contract while its mass increases. Finally, we approach a single square at roughly the same size as the squares in the underlying measures for large λ .

The HK_σ barycenter is close to a linear mean of the four measures for small cut-off. Increasing σ initially reduces the mass at each of the square locations. At a threshold of $\sigma \approx 0.34$, we observe a change, where part of the mass is moved vertically or horizontally to the mid points between the squares in a rectangular shape. Until $\sigma \approx 0.43$ all mass is shifted to these "middle-rectangles", at which point a second shift occurs, where the mass from these rectangles starts to move towards a square in the center. At $\sigma \approx 0.6$, all mass has been shifted towards a square in the center and there is no further change in the HK barycenter, when increasing σ .

Additionally, we consider Figs. 12 and 13, where the three unbalanced barycenter models are compared on three exemplary classes based on the MNIST dataset. Here, the original 28×28 images have been rescaled to sizes between 14×14 and 42×42 and embedded in a random subgrid of a 50×50 image. In this setting, there is a notable distinction between the GHK barycenter and the KR and HK barycenters. While for the former, the overall shape is recovered even for small parameter values, the latter two barycenters produce unstructured results for small parameters. The GHK distance is not constructed to have a maximal transport distance comparable to the impact of



(a)



(b)



(c)

Fig. 13 Comparison of the three unbalanced barycenters when varying their parameter. The set of underlying measures for **a** is the first row of Fig. 12. For **b** it is the second for **c** the third. For each class of examples the three different UOT barycenter models are considered in different rows: First row: The Gaussian–Hellinger–Kantorovich barycenter for $\lambda = 0.01, 0.12, 0.23, \dots, 1$. Second row: The Hellinger–Kantorovich barycenter for $\sigma = 0.01, 0.08, \dots, 0.64$. Third row: The KR barycenter for $C = 0.1, 0.2, \dots, 1$

C or σ in the other two cases, which allows to transport across larger distance and recover the correct shape for smaller values of λ . However, the mass of the GHK barycenter is significantly smaller than that of the original measures for small values of λ and only increases to the correct magnitude for larger penalty values. The HK and KR barycenters consist of fragments of the final shape which move towards a joint location for increasing parameters. For large penalties all three models are nearly identical and display the corresponding number correctly. This makes sense, as in this setting the minimisation in any individual term of the (p, C) -Fréchet functional is driven by minimising an OT term. We point out that for the (p, C) -barycenter this

regime is guaranteed to be reached by choosing C larger than the diameter of the space, while for the other two models the suitable parameter choice for this example is ambiguous without actually computing the result for specific values.

Overall, for large parameter values all considered UBCs perform similarly. In small parameter regimes we observe significant differences. This difference in behavior is to be expected as the dependence of the UOT models on their parameters varies significantly. One key advantage of the KR barycenter is that its connection between the choice of C and the properties of the resulting barycenter is immediate and intuitive. While the cut-off locus σ for the HK barycenter fulfils a similar role, imposing control at the maximum scale at which transport does occur, the consequences of changing σ from one value to another are far less immediate due to the involved structure of the cost functional in this setting. Similarly to the KR barycenter, it is worth noticing that the HK barycenter does allow for mass at locations given by centroids of support points of $L < N$ measures. Though, while for the KRD a feature of the underlying measures is only contained in the barycenter if it is present in more than $L = N/2$ measures, the HK barycenter also allows for mass at locations constructed from less support points. Thus, the HK barycenter is prone to being more susceptible to errors due to noise within the data. Compared to the other two choices, the parameter λ of the GHK barycenter does appear to have less interpretation, with the only clear connection being that increasing λ increases the mass of the GHK barycenter. There does also not appear to be any well-founded method how to approach the choice of λ for a given dataset.

Acknowledgements F. Heinemann and M. Klatt gratefully acknowledge support from the DFG Research Training Group 2088 *Discovering structure in complex data: Statistics meets optimization and inverse problems*. A. Munk gratefully acknowledges support from the DFG CRC 1456 *Mathematics of the Experiment A04, C06* and the Cluster of Excellence 2067 MBExC *Multiscale bioimaging—from molecular machines to networks of excitable cells*. We kindly thank two anonymous referees for their helpful comments and suggestions.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Competing interest The authors have not disclosed any competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Proofs

A.1 Proofs of Section 3

Proof of Lemma 3.3 Let $\mu \in \mathcal{M}_+(\mathcal{Y})$ be such that $\mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)$. Then

$$\begin{aligned} F_{p,C}(\mu) &\stackrel{(i)}{=} \frac{1}{J} \sum_{i=1}^J \tilde{\text{OT}}_{d_C^p}(\mu + \mathbb{M}(\mu_i)\delta_{\mathfrak{d}}, \mu_i + \mathbb{M}(\mu)\delta_{\mathfrak{d}}) \\ &\stackrel{(ii)}{=} \frac{1}{J} \sum_{i=1}^J \tilde{\text{OT}}_{d_C^p}\left(\mu + \left(\sum_{i=1}^J \mathbb{M}(\mu_i) - \mathbb{M}(\mu)\right)\delta_{\mathfrak{d}}, \tilde{\mu}_i\right) \\ &= \tilde{F}_{p,C}\left(\mu + \left(\sum_{i=1}^J \mathbb{M}(\mu_i) - \mathbb{M}(\mu)\right)\delta_{\mathfrak{d}}\right), \end{aligned}$$

where (i) follows from the lift to an OT problem (Sect. 3.1) and (ii) follows from Lemma 3.1 by adding mass $\sum_{j \neq i} \mathbb{M}(\mu_j) - \mathbb{M}(\mu)$ at \mathfrak{d} . We then have that

$$\begin{aligned} \min_{\substack{\mu \in \mathcal{M}_+(\mathcal{Y}) \\ \mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)}} F_{p,C}(\mu) &= \min_{\substack{\mu \in \mathcal{M}_+(\mathcal{Y}) \\ \mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)}} \tilde{F}_{p,C}\left(\mu + \left(\sum_{i=1}^J \mathbb{M}(\mu_i) - \mathbb{M}(\mu)\right)\delta_{\mathfrak{d}}\right) \\ &\geq \min_{\substack{\mu \in \mathcal{M}_+(\tilde{\mathcal{Y}}) \\ \mathbb{M}(\mu) = \sum_{i=1}^J \mathbb{M}(\mu_i)}} \tilde{F}_{p,C}(\mu) \end{aligned}$$

and

$$\min_{\substack{\mu \in \mathcal{M}_+(\tilde{\mathcal{Y}}) \\ \mathbb{M}(\mu) = \sum_{i=1}^J \mathbb{M}(\mu_i)}} \tilde{F}_{p,C}(\mu) = \min_{\substack{\mu \in \mathcal{M}_+(\tilde{\mathcal{Y}}) \\ \mathbb{M}(\mu) = \sum_{i=1}^J \mathbb{M}(\mu_i)}} F_{p,C}(\mu|_{\mathcal{Y}}) \geq \min_{\substack{\mu \in \mathcal{M}_+(\mathcal{Y}) \\ \mathbb{M}(\mu) \leq \sum_{i=1}^J \mathbb{M}(\mu_i)}} F_{p,C}(\mu).$$

Combining both inequalities and using Lemma 3.2 then finishes the proof.

Proof of Lemma 3.5 (i) By definition, the objective value for $\tilde{T}_C^{J,p}(y_1, \dots, y_J)$ at \mathfrak{d} is equal to $(J - |\mathcal{B}|)C^p/2$. Thus, $\tilde{T}_C^{J,p}$ outputs \mathfrak{d} if and only if for any $y \in \mathcal{Y}$ it holds

$$\sum_{i=1}^J \tilde{d}_C^p(y_i, y) \geq (J - |\mathcal{B}|)C^p/2$$

which is equivalent to

$$\sum_{i \notin \mathcal{B}} \tilde{d}_C^p(y_i, y) \geq (J - 2|\mathcal{B}|)C^p/2.$$

In particular, if all inequalities are strict ϑ is the unique output for $\tilde{T}_C^{J,p}(y_1, \dots, y_J)$. Statement (ii) is a direct consequence of (i). For statement (iii) we again use that by definition $\tilde{d}_C^p(y, \vartheta) = C^p/2$ for any $y \in \mathcal{Y}$ and hence

$$\min_{y \in \mathcal{Y}} \sum_{i=1}^J \tilde{d}_C^p(y_i, y) = |\mathcal{B}| \frac{C^p}{2} + \min_{y \in \mathcal{Y}} \sum_{i \notin \mathcal{B}} \tilde{d}_C^p(y_i, y).$$

Proving (iv), let $C > 2^{1/p} \text{diam}(\mathcal{Y})$, pick points $y_1, \dots, y_J \in \mathcal{Y}$ and observe that for any $y \in \mathcal{Y}$ it holds that

$$\sum_{i=1}^J \tilde{d}_C^p(y_i, \vartheta) = J \frac{C^p}{2} > J \text{diam}(\mathcal{Y})^p \geq \sum_{i=1}^J d^p(y_i, y).$$

Thus, $\tilde{T}_C^{J,p}(x_1, \dots, x_J) \neq \vartheta$ and since $|\mathcal{B}| = 0$, the claim follows from (iii). □

A.2 Proofs of Section 2

Proof for Lemma 2.1 Suppose that π_C is optimal but its induced graph $G(\pi_C)$ contains a path $P = (x_{i_1}, \dots, x_{i_k})$ such that $\mathcal{L}(P) > C^p$. By definition of $G(\pi_C)$ it holds that $\pi_C(x_{i_j}, x_{i_{j+1}}) > 0$ for all $1 \leq j \leq k - 1$. We define a new transport plan with augmented transport along the path P . For this, define $\epsilon := \min_{1 \leq j \leq k-1} \pi_C(x_{i_j}, x_{i_{j+1}})$ and construct the new plan

$$\tilde{\pi}_C(x, x') = \begin{cases} \pi_C(x, x') - \epsilon, & \text{if } \exists 1 \leq j \leq k - 1, x = x_{i_j}, x' = x_{i_{j+1}} \\ \pi_C(x, x'), & \text{else.} \end{cases}$$

Compared to π_C the transportation cost for $\tilde{\pi}_C$ is reduced by $\epsilon \mathcal{L}(P)$ while the marginal deviation is increased by ϵC^p . In particular, it holds that

$$\begin{aligned} & \sum_{x, x'} d^p(x, x') \tilde{\pi}_C(x, x') + \frac{C^p}{2} \left(\sum_x \mu(x) - \tilde{\pi}_C(x, \mathcal{X}) + \sum_{x'} \nu(x') - \tilde{\pi}_C(\mathcal{X}, x') \right) \\ &= \sum_{x, x'} d^p(x, x') \pi_C(x, x') + \frac{C^p}{2} \left(\sum_x \mu(x) - \pi_C(x, \mathcal{X}) + \sum_{x'} \nu(x') - \pi_C(\mathcal{X}, x') \right) \\ &+ \epsilon (C^p - \mathcal{L}(P)). \end{aligned}$$

As $\epsilon > 0$ and $\mathcal{L}(P) > C^p$ this contradicts the optimality for π_C . Consequently, any path P in the induced graph $G(\pi_C)$ necessarily has path length at most C^p . If $d(x, x') > C$ this implies that $d^p(x, x') > C^p$ and hence by the statement on induced graphs that $\pi_C(x, x') = 0$. □

Proof for Theorem 2.2 We first establish the metric properties (i). It is straightforward to show $\text{KR}_{p,C}(\mu, \nu) = 0$ if and only if $\mu = \nu$ and that $\text{KR}_{p,C}$ is symmetric. For the

triangle inequality let $\mu, \nu, \tau \in \mathcal{M}_+(\mathcal{X})$ and choose $B \geq \max\{\mathbb{M}(\mu), \mathbb{M}(\nu), \mathbb{M}(\tau)\}$. Then by augmenting the measures accordingly (Sect. 3.1) we find that

$$\begin{aligned} \text{KR}_{p,C}(\mu, \nu) &= \left(\tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\mu}, \tilde{\nu})\right)^{1/p} \\ &\leq \left(\tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\mu}, \tilde{\tau})\right)^{1/p} + \left(\tilde{\text{OT}}_{\tilde{d}_C^p}(\tilde{\tau}, \tilde{\nu})\right)^{1/p} = \text{KR}_{p,C}(\mu, \tau) + \text{KR}_{p,C}(\tau, \nu) \end{aligned}$$

where the inequality follows by the triangle inequality for the Wasserstein distance [71, Theorem 7.3]. Statement (ii) follows from Lemma 2.1 by noting that there exists at least one optimal solution π_C equal to zero except on the diagonal for which $\pi_C(x, x) = \mu(x) \wedge \nu(x)$. Plugging into the objective of (2) yields the claim. Additionally, suppose that w.l.o.g. $\mu(x) \geq \nu(x)$ for all $x \in \mathcal{X}$. Then independent to the choice of $C > 0$ and $p \geq 1$ the unique optimal solution is to remain all shared mass at its common place and to delete surplus material which is exactly the solution $\pi_C(x, x) = \mu(x) \wedge \nu(x)$ described before. Statement (iii) follows by noting that for $C \geq \max_{x,x'} d(x, x')$ the dual formulation in (DUOT_{p,C}) and in (DOT_p) coincide.

Finally, for statement (iv) we note that by construction it holds $\tilde{d}_{C_1}^p(x, y) \leq \tilde{d}_{C_2}^p(x, y)$ for all $x, y \in \tilde{\mathcal{Y}}$. Hence, for any coupling π of the augmented measures $\tilde{\mu}, \tilde{\nu}$ it holds

$$\sum_{x,x' \in \tilde{\mathcal{Y}}} \tilde{d}_{C_1}^p(x, x')\pi(x, x') \leq \sum_{x,x' \in \tilde{\mathcal{Y}}} \tilde{d}_{C_2}^p(x, x')\pi(x, x').$$

Taking the minimum over all couplings of $\tilde{\mu}$ and $\tilde{\nu}$ on both sides completes the proof. □

A.2.1 Proof for Theorem 2.3

Using the lift to the OT problem, we can now start to prove the closed formula on ultrametric trees. For this, consider an ultrametric tree \mathcal{T} with height function $h : V \rightarrow \mathbb{R}_+$ and define its p -height transformed tree denoted $\mathcal{T}_p := \mathcal{T}$ as the same tree but with height function $h_p(v) = 2^{p-1}h(v)^p$. An illustration is given in Fig. 4. Notice that by monotonicity \mathcal{T}_p is again an ultrametric tree.

Lemma A.1 *Let \mathcal{T} be an ultrametric tree with height function $h : V \rightarrow \mathbb{R}_+$ and consider its p -height transformed tree \mathcal{T}_p . Then it holds that*

$$d_{\mathcal{T}}^p(v, w) = d_{\mathcal{T}_p}(v, w)$$

for all leaf nodes $v, w \in L \subset V$.

Proof Let $v, w \in L$ be two leaf nodes in the ultrametric tree \mathcal{T} with height function h and let a be their common ancestor⁷. Since paths between any two vertices are unique

⁷ If $v, w \in L$ are leaf nodes their common ancestor is defined as the node included in the path from v to w closest to the root.

and all leaf nodes have the same distance to the root, it holds that

$$\begin{aligned} d_{\mathcal{T}}(v, a) - d_{\mathcal{T}}(w, a) &= d_{\mathcal{T}}(v, a) + d_{\mathcal{T}}(a, r) - d_{\mathcal{T}}(a, r) - d_{\mathcal{T}}(w, a) \\ &= d_{\mathcal{T}}(v, r) - d_{\mathcal{T}}(w, r) = 0. \end{aligned}$$

Hence,

$$(d_{\mathcal{T}}(v, w))^p = (d_{\mathcal{T}}(v, a) + d_{\mathcal{T}}(w, a))^p = 2^p(h(a) - h(v))^p = 2^p h(a)^p,$$

where we use that $h(v) = 0$. Repeating the argument for the ultrametric tree \mathcal{T}_p we conclude that $d_{\mathcal{T}_p}(v, w) = 2d_{\mathcal{T}_p}(v, a) = 2^p h(a)^p$.

Equipped with this result we are now able to prove the closed formula from Theorem 2.3.

Proof for Theorem 2.3 Let $KR_{p,C}^p(\mu, \nu) = UOT_{p,C}(\mu, \nu)$ refer to UOT w.r.t. the distance on \mathcal{T} , which only depends on the distance between individual leaf nodes. Considering the p -th height transformed tree \mathcal{T}_p and applying Lemma A.1 we conclude that

$$\begin{aligned} KR_{p,C}^p(\mu^L, \nu^L) &= \min_{\pi \in \mathcal{M}_+(L \times L)} \sum_{v, v' \in L} d_{\mathcal{T}_p}(v, v') \pi(v, v') \\ &\quad + \frac{C^p}{2} \left(\sum_{v \in L} (\mu^L(v) - \pi(v, L)) + \sum_{v' \in L} (\nu^L(v') - \pi(L, v')) \right) \\ \text{s.t. } \pi(v, L) &\leq \mu^L(v), \quad \forall v \in L, \\ \pi(L, v') &\leq \nu^L(v'), \quad \forall v' \in L. \end{aligned}$$

The linear optimization problem can be decomposed on several subtrees. For this recall that by Lemma 2.1 (i) there exists an optimal solution such that mass transportation is only considered on metric scales between two leaf nodes $v, v' \in L$ such that $d_{\mathcal{T}_p}(v, v') \leq C^p$. If v_0 is the common ancestor of v, v' then by the ultrametric tree properties of \mathcal{T} (see also the proof of Lemma A.1) the inequality $d_{\mathcal{T}_p}(v, v') \leq C^p$ is equivalent to the height function $h(v_0) \leq \frac{C}{2}$. Consider the set $\mathcal{R}(C)$ in (7) and for each $v \in \mathcal{R}(C)$ define subtrees $\mathcal{C}(v)$ consisting of the children of v and the subset of corresponding edges. By construction if $v_i, v_j \in \mathcal{R}(C)$ with $v_i \neq v_j$ then the subtrees are disjoint $\mathcal{C}(v_i) \cap \mathcal{C}(v_j) = \emptyset$ (Fig. 4a for an illustration). In particular, the linear optimization problem $KR_{p,C}^p(\mu^L, \nu^L)$ is decomposed on each individual subtree $\mathcal{C}(v)$ for each $v \in \mathcal{R}(C)$. The distance on individual subtrees is set to be the p -th height transformed tree distance $d_{\mathcal{T}_p}$ which exactly captures the pairwise p -th power distance between leaf nodes belonging to the same subtree (Lemma A.1). For an element $v \in \mathcal{R}(C)$ consider its subtree $\mathcal{C}(v)$ with distance $d_{\mathcal{T}_p}$. By definition the maximal distance between its leaf nodes is bounded by $C^p/2$. We augment the subtree $\mathcal{C}(v)$ with a dummy node \tilde{v} and introduce an edge $e = (v, \tilde{v})$ with edge weight $\frac{C^p}{2} - 2^{p-1}h(v)^p$ (Fig. 4b for an illustration). Denote the augmented tree by $\tilde{\mathcal{C}}(v)$. Considering the measures μ^L, ν^L restricted to $\mathcal{C}(v)$ we augment μ^L adding mass $(\mu^L(\mathcal{C}(v)) - \nu^L(\mathcal{C}(v)))_+$

at \tilde{v} and vice versa augment v^L adding mass $(v^L(\mathcal{C}(v)) - \mu^L(\mathcal{C}(v)))_+$ at \tilde{v} . This construction defines an equivalent OT problem on $\tilde{\mathcal{C}}(v)$ [33]. Hence, applying the closed formula for OT on general metric trees [20, p.575] yields

$$2^{p-1} \sum_{w \in \mathcal{C}(v) \setminus \{v\}} \left((h(\text{par}(w))^p - h(w)^p) \left| \mu^L(\mathcal{C}(w)) - v^L(\mathcal{C}(w)) \right| \right) + \left(\frac{C^p}{2} - h(v) \right) \left| \mu^L(\mathcal{C}(v)) - v^L(\mathcal{C}(v)) \right|.$$

Summing over all subtrees indexed by the set $\mathcal{R}(C)$ finishes the proof. □

Proofs for the Barycenter

Proof of Theorem 2.5 (ii) Let μ be a (p, C) -barycenter and $\tilde{\mu}$ its augmented counterpart. Then, by Proposition 3.9 there exists an optimal multi-coupling π , such that it holds $\mu = \tilde{\mu}|_{\mathcal{Y}} = (\tilde{T}_C^{J,p} \# \pi)|_{\mathcal{Y}}$. Hence, for each $y \in \text{supp}(\tilde{\mu})$ there exists $K_y \geq 1$ and K_y J -tupels $(x_{y,1}^1, \dots, x_{y,1}^J), \dots, (x_{y,K_y}^1, \dots, x_{y,K_y}^J)$ such that for $k = 1, \dots, K_y$ it holds

$$y = \tilde{T}_C^{J,p}(x_{y,k}^1, \dots, x_{y,k}^J)$$

and $\mu(y) = \sum_{k=1}^{K_y} a_k^y$, where $a_k^y = \pi(x_{y,k}^1, \dots, x_{y,k}^J)$. For $i = 1, \dots, J$ define $\tilde{\pi}_i \in \mathcal{M}_+(\mathcal{Y} \times \mathcal{Y})$ by $\tilde{\pi}_i(y, x_{y,k}^i) = a_k^y$ for all $y \in \text{supp}(\tilde{\mu}), i = 1, \dots, J$ and $k = 1, \dots, K_y$. Set $\tilde{\pi}_i$ to be zero everywhere else for $i = 1, \dots, J$. By construction, $\tilde{\pi}_i$ defines an OT plan between $\tilde{\mu}$ and $\tilde{\mu}_i$ for $i = 1, \dots, J$. It holds

$$\begin{aligned} \frac{1}{J} \sum_{i=1}^J \tilde{O}T_{\tilde{d}_C^p}(\tilde{\mu}, \tilde{\mu}_i) &= \sum_{x \in \text{supp}(\pi)} c_{p,C}(x) \pi(x) \\ &= \sum_{x \in \text{supp}(\pi)} \frac{1}{J} \sum_{i=1}^J \tilde{d}_C^p(x_i, \tilde{T}_C^{J,p}(x)) \pi(x) \\ &= \frac{1}{J} \sum_{y \in \text{supp}(\tilde{\mu})} \sum_{i=1}^J \sum_{k=1}^{K_y} \tilde{d}_C^p(x_{y,k}^i, y) a_k^y \\ &= \frac{1}{J} \sum_{i=1}^J \sum_{y \in \text{supp}(\tilde{\mu})} \sum_{k=1}^{K_y} \tilde{d}_C^p(x_{y,k}^i, y) \tilde{\pi}_i(y, x_{y,k}^i), \end{aligned}$$

where the first equality follows from Proposition 3.9 and the third and fourth by construction. Since $\tilde{\pi}_i$ is an OT plan between $\tilde{\mu}$ and $\tilde{\mu}_i$ it holds for all $i = 1, \dots, J$

that

$$\sum_{y \in \text{supp}(\tilde{\mu})} \sum_{k=1}^{K_y} \tilde{d}_C^p(x_{y,k}^i, y) \tilde{\pi}_i(y, x_{y,k}^i) \geq \tilde{O}T_{p,C}^p(\tilde{\mu}, \tilde{\mu}_i).$$

Thus, it follows together with the previous equations that

$$\tilde{O}T_{p,C}^p(\tilde{\mu}, \tilde{\mu}_i) = \sum_{y \in \text{supp}(\tilde{\mu})} \sum_{k=1}^{K_y} \tilde{d}_C^p(x_{y,k}^i, y) \tilde{\pi}_i(y, x_{y,k}^i),$$

i.e. $\tilde{\pi}_i$ is optimal. Lemma 3.5 now yields the first part of the statement.

For the second part assume that for any $(x_1, \dots, x_L) \in \mathcal{Y}^L$ it holds that $T^{L,p}(x_1, \dots, x_L) = T^{L,p}(y_1, x_2, \dots, x_L)$ is equivalent to $x_1 = y_1$. Let $y \in \text{supp}(\tilde{\mu})$ and consider OT plans $\tilde{\pi}^1, \dots, \tilde{\pi}^J$ between $\tilde{\mu}$ and $\tilde{\mu}_i$, respectively. For $i = 1, \dots, J$ consider x^i such that $\tilde{\pi}_i(y, x^i) = a_i > 0$. Assume that it holds $y \neq \tilde{T}_C^{J,p}(x_1, \dots, x_J)$. Denote the minimum of the a_i as $a_0 = \min_{i=1, \dots, J} a_i$. By construction, it follows that

$$\tilde{F}_{p,C}(\tilde{\mu} - a_0 \delta_y + a_0 \delta_{\tilde{T}_C^{J,p}(x_1, \dots, x_J)}) < F_{p,C}(\tilde{\mu}),$$

which is a contradiction to $\tilde{\mu}$ being a barycenter of $\tilde{\mu}_1, \dots, \tilde{\mu}_J$. Thus, it holds $y = \tilde{T}_C^{J,p}(x_1, \dots, x_J)$. Now, assume w.l.o.g. there exists $x_1, z_1 \in \mathcal{Y}$, such that it holds $\pi^1(y, x_1) > 0$ and $\pi^1(y, z_1) > 0$. However, by the previous argument this implies

$$\tilde{T}_C^{J,p}(x_1, \dots, x_J) = y = \tilde{T}_C^{J,p}(z_1, \dots, x_J).$$

By assumption this is equivalent to $x_1 = z_1$, thus it holds for all $x, y \in \mathcal{Y}$ and $i = 1, \dots, J$ that $\pi^i(y, x) \in \{0, \mu(y)\}$.

(i) By Proposition 3.9 the objective value of the balanced multi-marginal and (p, C) -barycenter problem coincide and a (p, C) -barycenter is obtained as the push-forward of an optimal balanced multi-coupling under the map $\tilde{T}_C^{J,p}$ restricted to \mathcal{Y} . By construction and Corollary 3.7 any such measure is supported in $\mathcal{C}_{KR}(J, p, C)$. Thus, there always exists a (p, C) -barycenter whose support is restricted to $\mathcal{C}_{KR}(J, p, C)$ and the minimum over Y and $\mathcal{C}_{KR}(J, p, C)$ coincide.

The second part is similar and we let $\tilde{\mu}$ be any p -Wasserstein barycenter. Then by Proposition 3.9, there exists a multi-coupling of $\tilde{\mu}_1, \dots, \tilde{\mu}_J$, such that $\tilde{\mu} = \tilde{T}_C^{J,p} \# \tilde{\pi}$. Since any such push-forward measure can only have support in $\mathcal{C}_{KR}(J, p, C) \cup \{\emptyset\}$, it holds for $\mu = \tilde{\mu}|_{\mathcal{Y}}$ that $\text{supp}(\mu) \subset \mathcal{C}_{KR}(J, p, C)$. It remains to show the upper bound on the total mass. By the equivalence to the multi-marginal problem and by Lemma 3.5

(ii) any (p, C) -barycenter μ cannot have mass on a point which is constructed from a set of points (x_1, \dots, x_J) for which $2|\mathcal{B}(x_1, \dots, x_J)| \geq J$. Additionally, by part (ii) we know that there exists UOT plans, such that the mass of each (p, C) -barycenter support point is fully transported to points it is constructed from. Let (a_1, \dots, a_K) be

the weight vector of the support points of the (p, C) -barycenter, then it holds that

$$\sum_{i=1}^J \mathbb{M}(\mu_i) - \lceil J/2 \rceil \sum_{k=1}^K a_k \geq 0,$$

since by the previous argument and Lemma 3.5, any (p, C) -barycenter support point x_k reduces the maximum available mass by at least $\lceil J/2 \rceil a_k$ and by Lemma 3.2, the total mass of the (p, C) -barycenter is bounded by the sum of the total masses of the μ_i . Therefore it holds that

$$\mathbb{M}(\mu) = \sum_{k=1}^K a_k \leq \lceil J/2 \rceil^{-1} \sum_{i=1}^J \mathbb{M}(\mu_i) \leq \frac{2}{J} \sum_{i=1}^J \mathbb{M}(\mu_i).$$

(iii) The multi-marginal problem between $\tilde{\mu}_1, \dots, \tilde{\mu}_J$ is a balanced problem, thus we can pose this as a linear program with a total of $\prod_{i=1}^J M_i$ variables and $\sum_{i=1}^N M_i + J$ constraints. As all measures have the same total mass, we can drop one arbitrary marginal constraint for each measure besides the first. Thus, the rank of the constraint matrix in the corresponding constraint is bounded by $\sum_{i=1}^N M_i + 1$. Hence, each basic feasible solution of the linear program has at most $\sum_{i=1}^N M_i + 1$ non-zero entries (see [47] for details). Let π be such a solution. By Proposition 3.9 the measure $\tilde{\mu} = \tilde{T}_C^{J,p} \# \pi$ is a p -Wasserstein barycenter and by construction it has at most $\sum_{i=1}^N M_i + 1$ support points. Due to the upper bound on the total mass of the (p, C) -barycenter in property (i), we can guarantee that there is non-zero mass at \mathfrak{d} for $J > 2$, hence in this case, restricting the measure to \mathcal{Y} reduces the support size by one. For $J = 2$, we note that the multi-marginal problem is just the augmented UOT problem. By construction we either have a point x in the support of one of the two measures, such that there is transport between x and \mathfrak{d} or both measures have equal mass at \mathfrak{d} and it is optimal to leave this mass in place. In the first case, we have mass at $\tilde{T}_C^{J,p}(x, \mathfrak{d}) = \mathfrak{d}$, thus the support size can be reduced by one and in the second the problem is equivalent to the OT problem and thus the barycenter has at most $M_1 + M_2 - 1$ support points. Finally, by property (i) the support of any (p, C) -barycenter is contained in $\mathcal{C}_{KR}(J, p, C)$, thus the cardinality of this set also provides a trivial upper bound on the support size of any (p, C) -barycenter. Taking the minimum over both quantities, we conclude

$$|\text{supp}(\mu)| \leq \min \left\{ |\mathcal{C}_{KR}(J, p, C)|, \sum_{i=1}^J M_i \right\}.$$

(iv) For any $\mu \in \mathcal{M}_+(\mathcal{Y})$, it holds

$$F_{p,C_1}^p(\mu) = \frac{1}{J} \sum_{i=1}^J \text{KR}_{p,C_1}^p(\mu, \mu_i) \leq \frac{1}{J} \sum_{i=1}^J \text{KR}_{p,C_2}^p(\mu, \mu_i) = F_{p,C_2}^p(\mu),$$

where the inequality follows from Theorem 2.2 (iv). Taking the infimum over all measures in $\mathcal{M}_+(\mathcal{Y})$ on both sides completes the proof.

(v) Let $C \leq d'_{\min}$, then by Theorem 2.2 (ii) it holds

$$\begin{aligned} \arg \min_{\mu \in \mathcal{M}_+(\mathcal{Y})} F_{p,C}(\mu) &= \arg \min_{\mu \in \mathcal{M}_+(\mathcal{Y})} \frac{C^p}{2J} \sum_{i=1}^J TV(\mu, \mu_i) \\ &= \arg \min_{a \in \mathbb{R}_+^K} \frac{C^p}{2J} \sum_{i=1}^J \sum_{k=1}^K |a_k - a_k^i| \\ &= \arg \min_{a \in \mathbb{R}_+^K} \frac{C^p}{2J} \sum_{k=1}^K \sum_{i=1}^J |a_k - a_k^i|, \end{aligned}$$

where the change in the arg min in the second line follows from the fact that the total variation can only increase if we place mass outside of the support of the measures. Thus it suffices to consider measures supported on the union of the supports. Now, we note that the K summands are independent to each other, thus we can minimise them separately. Hence, for the k -th entry of a it holds that

$$a_k \in \arg \min_{a \in \mathbb{R}_+} \sum_{i=1}^N |a_k - a_k^i| = \text{med}(a_k^1, \dots, a_k^J)$$

which yields the claim.

(vi) Let $\mathbb{M}_i = \mathbb{M}(\mu_i)$ for $i = 1, \dots, J$ and set $\mathbb{M}_0 = 0$. Assume that J is odd. Let μ be a (p, C) -barycenter of μ_1, \dots, μ_J with $\mu(\mathcal{Y}) \in [\mathbb{M}_{k-1}, \mathbb{M}_k]$. In particular, μ fulfills the non-mass-splitting property in (ii). Let $a \in (0, \mathbb{M}_k - \mathbb{M}(\mu)]$ and $\tilde{\mu}$ the augmented measure for μ . By construction, we can find support points $x_k, \dots, x_J \neq \vartheta$ of the augmented measures $\tilde{\mu}_k, \dots, \tilde{\mu}_J$ from which w.l.o.g. mass a is transported to ϑ in μ . If one of the points has mass smaller a , we can just replace a with the minimum of the masses of the points and repeat the argument until we have considered a total mass of a . Set $x_0 = \tilde{T}_C^{J,p}(\vartheta, \dots, \vartheta, x_k, \dots, x_J)$ and notice that if $x_0 = \vartheta$, we do not change the objective function in the augmented problem (Lemma 3.1) by adding this point which means w.l.o.g. $x_0 \neq \vartheta$. In this case, we have

$$x_0 = \arg \min_{x \in \mathcal{Y}} \sum_{i=k}^J \tilde{d}_C^p(x_i, x).$$

Now, the objective cost of not having mass a at x_0 is $aC^p(J - k)/2$, while the cost of adding $a\delta_{x_0}$ to μ is equal to $a(kC^p/2 + \sum_{i=k}^J \tilde{d}_C^p(x_i, x_0))$. Hence, adding the point improves the value of the Fréchet functional, if

$$\sum_{i=k}^J \tilde{d}_C^p(x_i, x_0) \leq C^p(J - 2k)/2.$$

For $2k > J$, the right hand side will always be negative, so we can not improve. Thus, we assume $2k < J$. By assumption it holds $C \geq J^{\frac{1}{p}} \text{diam}(\mathcal{Z})$. Hence,

$$\begin{aligned} \frac{C^p}{2} &\geq \frac{J}{2} \text{diam}(\mathcal{Z})^p \\ \Leftrightarrow \frac{C^p}{2 \text{diam}(\mathcal{Z})^p} &\geq \frac{J}{2} \geq \frac{J-k}{J-2k} \\ \Leftrightarrow C^p(J-2k)/2 &\geq \text{diam}(\mathcal{Z})^p(J-k) \geq \sum_{i=k}^J \tilde{d}_C^p(x_i, x_0). \end{aligned}$$

Therefore, for $2k < J$ the objective value of μ can always be improved by increasing its mass by a , as long as $k < \lceil J/2 \rceil$. Thus, since μ is a barycenter it holds $\mathbb{M}(\mu) \geq \mathbb{M}\mu_{\lceil J/2 \rceil}$.

An analog, converse argument yields that if $k > J/2$, we can always improve the objective value of μ , since removing and then re-adding any mass to μ increases the objective value by the previous argument. Hence, it holds $\mathbb{M}(\mu) = \mathbb{M}\mu_{\lceil J/2 \rceil}$.

Now, assume J is even. For $2k \neq J$ nothing in the previous argument changes. However, for $2k = J$ (note that this can only hold now that J is even), the right hand side is zero, however, if all the x_i for $i = k, \dots, J$, are identical to x_0 (in particular, there exists a point contained in the support of at least half of the measures), then the left hand side will also be zero. In this case, the presence of this point does not change the objective value and there are (p, C) -barycenters of different total masses. However, we can still always choose to not place mass in such cases, to obtain a (p, C) -barycenter of the desired total mass.

Proof for Lemma 2.7 It suffices to show that there is no centroid point, which is constructed from points from two or more different sets B_r . Assume there is a point $y_0 \in \mathcal{C}_{KR}(J, p, C)$, such that y_0 is constructed, among others, from $x_1 \in B_r$ and $x_2 \in B_s$ for $r \neq s$. We distinguish two cases. Assume $y_0 \in B_r$, then it holds $d^p(x_1, y_0) > 2^{p-1}C^p \geq C^p$ and y_0 would not be in the restricted centroid set. The analogue argument holds for $y_0 \in B_s$. Now, assume y_0 is neither in B_r nor B_s . Since $d(B_r, B_s) > 2^{1/p}C$, it holds either $d^p(B_r, y_0) > C^p$ or $d^p(B_s, y_0) > C^p$. Thus, we obtain another contradiction to $y_0 \in \mathcal{C}_{KR}(J, p, C)$. Hence, $\mathcal{C}_{KR}(J, p, C)$ only contains centroids constructed from points within one B_r and by convexity of the B_r , any centroid point constructed from points within B_r is again in B_r . Theorem 2.5 (ii) yields that there will always be an optimal solution which only transports within each B_r , thus the R problems are in fact independent and we can separate them without changing the objective value. □

References

1. Agueh, M., Carlier, G.: Barycenters in the Wasserstein space. *SIAM J. Math. Anal.* **43**(2), 904–924 (2011)
2. Altschuler, J.M., Boix-Adsera, E.: Wasserstein barycenters can be computed in polynomial time in fixed dimension. *J. Mach. Learn. Res.* **22**, 44–1 (2021)

3. Álvarez-Esteban, P.C., Del Barrio, E., Cuesta-Albertos, J., et al.: A fixed-point approach to barycenters in Wasserstein space. *J. Math. Anal. Appl.* **441**(2), 744–762 (2016)
4. Anderes, E., Borgwardt, S., Miller, J.: Discrete Wasserstein barycenters: optimal transport for discrete data. *Math. Methods Oper. Res.* **84**(2), 389–409 (2016)
5. Benamou, J.D., Brenier, Y.: A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik* **84**(3), 375–393 (2000)
6. Benamou, J.D., Carlier, G., Cuturi, M., et al.: Iterative Bregman projections for regularized transportation problems. *SIAM J. Sci. Comput.* **37**(2), A1111–A1138 (2015)
7. Bertsimas, D., Tsitsiklis, J.N.: Introduction to Linear Optimization, vol. 6. Athena Scientific, Belmont (1997)
8. Bogachev, V.I.: Measure Theory, vol. 1. Springer, Berlin (2007)
9. Bonneel, N., Rabin, J., Peyré, G., et al.: Sliced and Radon Wasserstein barycenters of measures. *J. Math. Imaging Vis.* **51**(1), 22–45 (2015)
10. Bonneel, N., Peyré, G., Cuturi, M.: Wasserstein barycentric coordinates: histogram regression using optimal transport. *ACM Trans. Gr.* **35**(4), 71–1 (2016)
11. Caffarelli, L.A., McCann, R.J.: Free boundaries in optimal transport and Monge-Ampère obstacle problems. *Ann. Math.* pp. 673–730 (2010)
12. Carlier, G., Oberman, A., Oudet, E.: Numerical methods for matching for teams and Wasserstein barycenters. *ESAIM: Math. Model. Numer. Anal.* **49**(6), 1621–1642 (2015)
13. Carlier, G., Duval, V., Peyré, G., et al.: Convergence of entropic schemes for optimal transport and gradient flows. *SIAM J. Math. Anal.* **49**(2), 1385–1418 (2017)
14. Chizat, L., Peyré, G., Schmitzer, B., et al.: An interpolating distance between optimal transport and Fisher-Rao metrics. *Found. Comput. Math.* **18**(1), 1–44 (2018)
15. Chizat, L., Peyré, G., Schmitzer, B., et al.: Scaling algorithms for unbalanced optimal transport problems. *Math. Comput.* **87**(314), 2563–2609 (2018)
16. Chizat, L., Peyré, G., Schmitzer, B., et al.: Unbalanced optimal transport: dynamic and Kantorovich formulations. *J. Funct. Anal.* **274**(11), 3090–3123 (2018)
17. Chung, N.P., Phung, M.N.: Barycenters in the Hellinger–Kantorovich space. *Appl. Math. Optim.*, pp. 1–30 (2020)
18. Cuturi, M.: Sinkhorn distances: lightspeed computation of optimal transport. *Adv. Neural Inf. Process. Syst.* **26**, 2292–2300 (2013)
19. Cuturi, M., Doucet, A.: Fast computation of Wasserstein barycenters. In: International Conference on Machine Learning, PMLR, pp. 685–693 (2014)
20. Evans, S.N., Matsen, F.A.: The phylogenetic Kantorovich-Rubinstein metric for environmental sequence samples. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **74**(3), 569–592 (2012)
21. Figalli, A.: The optimal partial transport problem. *Arch. Ration. Mech. Anal.* **195**(2), 533–560 (2010)
22. Fréchet, M.: Les éléments aléatoires de nature quelconque dans un espace distancié. In: *Annales de l’institut Henri Poincaré*, pp. 215–310 (1948)
23. Friesecke, G., Matthes, D., Schmitzer, B.: Barycenters for the Hellinger-Kantorovich distance over \mathbb{R}^d . *SIAM J. Math. Anal.* **53**(1), 62–110 (2021)
24. Frogner, C., Zhang, C., Mobahi, H., et al.: Learning with a Wasserstein loss. *Adv. Neural Inf. Process. Syst.* **28**, 2053–2061 (2015)
25. Gangbo, W., McCann, R.J.: Shape recognition via Wasserstein distance. *Q. Appl. Math.*, pp. 705–737 (2000)
26. Gangbo, W., Li, W., Osher, S., et al.: Unnormalized optimal transport. *J. Comput. Phys.* **399**(108), 940 (2019)
27. Gavryushkin, A., Drummond, A.J.: The space of ultrametric phylogenetic trees. *J. Theor. Biol.* **403**, 197–208 (2016)
28. Ge, D., Wang, H., Xiong, Z. et al.: Interior-point methods strike back: Solving the Wasserstein barycenter problem. In: *Advances in Neural Information Processing Systems*, vol. 32 (2019)
29. Gellert, M., Hossain, M.F., Berens, F.J.F., et al.: Substrate specificity of thioredoxins and glutaredoxins—towards a functional classification. *Heliyon* **5**(12), e02,943 (2019)
30. Gerber, S., Maggioni, M.: Multiscale strategies for computing optimal transport. *J. Mach. Learn. Res.* **18**, 1–32 (2017)
31. Gramfort, A., Peyré, G., Cuturi, M.: Fast optimal transport averaging of neuroimaging data. In: International Conference on Information Processing in Medical Imaging, pp. 261–272. Springer, Berlin (2015)

32. Grötschel, M., Lovász, L., Schrijver, A.: Geometric Algorithms and Combinatorial Optimization, vol. 2. Springer, Berlin (2012)
33. Guittet, K.: Extended Kantorovich norms: a tool for optimization. Tech. rep., Technical Report 4402, INRIA (2002)
34. Hallin, M., Mordant, G., Segers, J.: Multivariate goodness-of-fit tests based on Wasserstein distance. *Electron. J. Stat.* **15**(1), 1328–1371 (2021)
35. Hanin, L.G.: Kantorovich-Rubinstein norm and its application in the theory of Lipschitz spaces. *Proc. Am. Math. Soc.* **115**(2), 345–352 (1992)
36. Heinemann, F., Munk, A., Zemel, Y.: Randomized Wasserstein barycenter computation: resampling with statistical guarantees. *SIAM J. Math. Data Sci.* **4**(1), 229–259 (2022)
37. Kantorovich, L.V., Rubinstein, S.: On a space of totally additive functions. *Vestnik St Petersburg. Univ.: Math.* **13**(7), 52–59 (1958)
38. Kitagawa, J., Pass, B.: The multi-marginal optimal partial transport problem. In: *Forum of Mathematics, Sigma*, Cambridge University Press (2015)
39. Klatt, M., Taming, C., Munk, A.: Empirical regularized optimal transport: statistical theory and applications. *SIAM J. Math. Data Sci.* **2**(2), 419–443 (2020)
40. Kloeckner, B.R.: A geometric study of Wasserstein spaces: ultrametrics. *Mathematika* **61**(1), 162–178 (2015)
41. Kroshnin, A., Tupitsa, N., Dvinskikh, D., et al.: On the complexity of approximating Wasserstein barycenters. In: *International Conference on Machine Learning*, PMLR, pp. 3530–3540 (2019)
42. Le, T., Yamada, M., Fukumizu, K., et al.: Tree-sliced variants of Wasserstein distances. In: *Advances in Neural Information Processing Systems*, pp. 12304–12315 (2019)
43. Le Gouic, T., Loubes, J.M.: Existence and consistency of Wasserstein barycenters. *Probab. Theory Relat. Fields* **168**(3), 901–917 (2017)
44. Lellmann, J., Lorenz, D.A., Schonlieb, C., et al.: Imaging with Kantorovich-Rubinstein discrepancy. *SIAM J. Imaging Sci.* **7**(4), 2833–2859 (2014)
45. Liero, M., Mielke, A., Savaré, G.: Optimal entropy-transport problems and a new Hellinger-Kantorovich distance between positive measures. *Inventiones Mathematicae* **211**(3), 969–1117 (2018)
46. Lin, T., Ho, N., Chen, X., et al.: Fixed-support Wasserstein barycenters: computational hardness and fast algorithm. *Adv. Neural Inf. Process. Syst.* **33** (2020)
47. Luenberger, D.G., Ye, Y., et al.: *Linear and Nonlinear Programming*, vol. 2. Springer, Berlin (1984)
48. Luis, G., Salzo, S., Pontil, M., et al.: Sinkhorn barycenters with free support via frank-wolfe algorithm. *Adv. Neural Inf. Process. Syst.* **32** (2019)
49. Masarotto, V., Panaretos, V.M., Zemel, Y.: Procrustes metrics on covariance operators and optimal transportation of Gaussian processes. *Sankhya A* **81**(1), 172–213 (2019)
50. Mérigot, Q.: A multiscale approach to optimal transport. In: *Computer Graphics Forum*, Wiley Online Library, pp. 1583–1592 (2011)
51. Montesuma, E.F., Mboula, F.M.N.: Wasserstein barycenter for multi-source domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16785–16793 (2021)
52. Mukherjee, D., Guha, A., Solomon, J.M., et al.: Outlier-robust optimal transport. In: *International Conference on Machine Learning*, PMLR, pp. 7850–7860 (2021)
53. Müller, R., Schuhmacher, D., Mateu, J.: Metrics and barycenters for point pattern data. In: *Statistics and Computing*, pp. 1–20 (2020)
54. Panaretos, V.M., Zemel, Y.: *An Invitation to Statistics in Wasserstein Space*. Springer, Berlin (2020)
55. Pele, O., Werman, M.: A linear time histogram metric for improved SIFT matching. In: *European Conference on Computer Vision*, pp. 495–508. Springer (2008)
56. Peyré, G., Cuturi, M., et al.: Computational optimal transport: with applications to data science. *Found. Trends® Mach. Learn.* **11**(5–6), 355–607 (2019)
57. Piccoli, B., Rossi, F.: Generalized Wasserstein distance and its application to transport equations with source. *Arch. Ration. Mech. Anal.* **211**(1), 335–358 (2014)
58. Rabin, J., Peyré, G., Delon, J., et al.: 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)
59. Rachev, S.T., Rüschendorf, L.: *Mass Transportation Problems: Volume I: Theory*, vol. 1. Springer, Berlin (1998)
60. Santambrogio, F.: *Optimal Transport for Applied Mathematicians*, vol. 55. Springer, Berlin (2015)

61. Sato, R., Yamada, M., Kashima, H.: Fast unbalanced optimal transport on a tree. *Adv. Neural Inf. Process. Syst.* **33**, 19039–19051 (2020)
62. Schiebinger, G., Shu, J., Tabaka, M., et al.: Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell* **176**(4), 928–943 (2019)
63. Schmitz, M.A., Heitz, M., Bonneel, N., et al.: Wasserstein dictionary learning: optimal transport-based unsupervised nonlinear dictionary learning. *SIAM J. Imaging Sci.* **11**(1), 643–678 (2018)
64. Schmitzer, B.: Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM J. Sci. Comput.* **41**(3), A1443–A1481 (2019)
65. Schmitzer, B., Wirth, B.: A framework for Wasserstein-1-type metrics. *J. Convex Anal.* **26**(2), 353–396 (2019)
66. Solomon, J., De Goes, F., Peyré, G., et al.: Convolutional Wasserstein distances: efficient optimal transportation on geometric domains. *ACM Trans. Gr.* **34**(4), 1–11 (2015)
67. Sommerfeld, M., Munk, A.: Inference for empirical Wasserstein distances on finite spaces. *J. R. Stat. Soc. Ser. B* **80**(1), 219–238 (2018)
68. Sommerfeld, M., Schrieber, J., Zemel, Y., et al.: Optimal transport: Fast probabilistic approximation with exact solvers. *J. Mach. Learn. Res.* **20**(105), 1–23 (2019)
69. Su, Z., Zeng, W., Wang, Y., et al.: Shape classification using Wasserstein distance for brain morphometry analysis. In: *International Conference on Information Processing in Medical Imaging*, pp. 411–423. Springer (2015)
70. Tameling, C., Stoldt, S., Stephan, T., et al.: Colocalization for super-resolution microscopy via optimal transport. *Nat. Comput. Sci.* **1**(3), 199–211 (2021)
71. Villani, C.: *Topics in Optimal Transportation*, vol. 58, American Mathematical Society (2003)
72. Villani, C.: *Optimal Transport: Old and New*, vol. 338. Springer, Berlin (2008)
73. Wang, S., Yuan, M.: Revisiting colocalization via optimal transport. *Nat. Comput. Sci.* **1**(3), 177–178 (2021)
74. Wolsey, L.A., Nemhauser, G.L.: *Integer and Combinatorial Optimization*, vol. 55. Wiley, New York (1999)
75. Xie, Y., Wang, X., Wang, R., et al.: A fast proximal point method for computing exact Wasserstein distance. In: *Uncertainty in Artificial Intelligence*, PMLR, pp. 433–453 (2020)
76. Ye, J., Wu, P., Wang, J.Z., et al.: Fast discrete distribution clustering using Wasserstein barycenter with sparse support. *IEEE Trans. Signal Process.* **65**(9), 2317–2332 (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.