



Wasserstein Archetypal Analysis

Katy Craig¹ · Braxton Osting² · Dong Wang^{3,4} · Yiming Xu⁵

Accepted: 6 August 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

Archetypal analysis is an unsupervised machine learning method that summarizes data using a convex polytope. In its original formulation, for fixed k , the method finds a convex polytope with k vertices, called archetype points, such that the polytope is contained in the convex hull of the data and the mean squared Euclidean distance between the data and the polytope is minimal. In the present work, we consider an alternative formulation of archetypal analysis based on the Wasserstein metric, which we call Wasserstein archetypal analysis (WAA). In one dimension, there exists a unique solution of WAA and, in two dimensions, we prove the existence of a solution, as long as the data distribution is absolutely continuous with respect to the Lebesgue measure. We discuss obstacles to extending our result to higher dimensions and general data distributions. We then introduce an appropriate regularization of the problem, via a Rényi entropy, which allows us to obtain the existence of solutions of the regularized problem for general data distributions, in arbitrary dimensions. We prove a consistency result for the regularized problem, ensuring that if the data are iid samples from a probability measure, then as the number of samples is increased, a subsequence of the archetype points converges to the archetype points for the limiting data distribution, almost surely. Finally, we develop and implement a gradient-based computational

✉ Katy Craig
kcraig@math.ucsb.edu

Braxton Osting
osting@math.utah.edu

Dong Wang
wangdong@cuhk.edu.cn

Yiming Xu
yiming.xu@uky.edu

¹ Department of Mathematics, University of California, Santa Barbara, USA

² Department of Mathematics, University of Utah, Salt Lake City, Utah, USA

³ School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China

⁴ Shenzhen International Center for Industrial and Applied Mathematics, Shenzhen Research Institute of Big Data, Guangdong, China

⁵ Department of Mathematics, University of Kentucky, Lexington, KY, USA

approach for the two-dimensional problem, based on the semi-discrete formulation of the Wasserstein metric. Detailed numerical experiments are provided to support our theoretical findings.

Keywords Archetypal analysis · Optimal transport · Wasserstein metric · Unsupervised learning · Multivariate data summarization

Mathematics Subject Classification 62H12 · 62G07 · 65K10 · 49Q22

1 Introduction

Given a probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$, *archetypal analysis* (AA) aims to find the convex polytope $\Omega \subseteq \mathbb{R}^d$ with k vertices that best approximates μ . As originally introduced by Culter and Breiman in 1994 [9], given data $X = \{x_i\}_{i=1}^N \subseteq \mathbb{R}^d$ and $k \geq d + 1$, AA finds k vertices, $A = \{a_\ell\}_{\ell=1}^k \subseteq \mathbb{R}^d$, that belong to the convex hull of the data, for which the convex hull $\text{co}(A)$ explains the most variation of the dataset. In particular, AA can be framed in terms of the following constrained optimization problem

$$\min_{A=\{a_\ell\}_{\ell=1}^k \subseteq \mathbb{R}^d} \left\{ \frac{1}{N} \sum_{i=1}^N d^2(x_i, \text{co}(A)) : A \subseteq \text{co}(X) \right\}, \quad (1.1)$$

where $d^2(\cdot, \cdot)$ denotes the squared Euclidean distance. The archetypes, $A = \{a_\ell\}_{\ell=1}^k \subseteq \mathbb{R}^d$, may be interpreted as exemplars of extreme patterns of the dataset, a mixture of which explain the general characteristics of the associated distribution; see [6, 21, 26] for applications of AA in astronomy, biology, and many others.

AA is closely related to other unsupervised learning methods, such as k -means, principal component analysis, and nonnegative matrix factorization [21]. Under bounded support assumptions, AA is well-posed [9, 22] and the consistency and convergence of AA were recently established in [22], laying the foundation for AA to apply to large-scale datasets [17]. In practice, however, it is often more appropriate to assume that the distribution generating the data has finite moments but unbounded support, in which case AA is not well-posed [22]. Also, due to the definition of the loss, AA is unstable with respect to outliers [12]. In particular, one can construct datasets for which a single additional data point leads to an arbitrarily large change of the archetypal polytope. Likewise, AA is insensitive to perturbations of the data distribution on the interior of the archetypal polytope.

To address both issues, the present work considers a different formulation of the AA problem, based on the Wasserstein metric.

1.1 Main Results

Let $\mathcal{P}_2(\mathbb{R}^d)$ denote the space of Borel probability measures on $\mathbb{R}^d$ with finite second moment, i.e. $M_2(\mu) := \int_{\mathbb{R}^d} |x|^2 d\mu(x) < +\infty$. Given $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and a number of vertices $k \geq d + 1$, we seek to find the (nondegenerate) convex k -gon $\Omega \subseteq \mathbb{R}^d$ that

is *closest* to μ , in the sense that the uniform probability measure on Ω is as close as possible to μ in the 2-Wasserstein metric:

$$\min_{\Omega \in S_k} W_2(\mu, 1_{\Omega}/|\Omega|), \quad S_k = \left\{ \Omega \subset \mathbb{R}^d : \Omega \text{ is a convex } k\text{-gon with nonempty interior} \right\}, \quad (\text{WAA})$$

where

$$1_{\Omega}(x) = \begin{cases} 1 & \text{if } x \in \Omega, \\ 0 & \text{otherwise} \end{cases}.$$

Here, we use the term *k-gon* to mean a bounded polytope with k vertices.

We make mild abuse of notation in the above problem formulation and throughout this manuscript: if a measure $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ is absolutely continuous with respect to d -dimensional Lebesgue measure $\mathcal{L}^d$, we will write $\nu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$ and use the same symbol to denote both the measure and its density, $d\nu(x) = \nu(x) d\mathcal{L}^d(x)$. In this way, we will use $1_{\Omega}/|\Omega|$ to denote the uniform probability measure on Ω .

Informally, the 2-Wasserstein metric measures the distance between probability measures in terms of the amount of *effort* it takes to rearrange one to look like the other. More precisely, given $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, the 2-Wasserstein metric is defined by

$$W_2^2(\mu, \nu) = \inf_{X \sim \mu, Y \sim \nu} \mathbb{E}[|X - Y|^2], \quad (1.2)$$

where the expectation is taken with respect to a rich enough underlying probability space and the infimum is taken over all couplings (X, Y) with marginals μ and ν . If $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$, then there is a unique optimal coupling and, furthermore, the coupling is *deterministic*: there exists a measurable function T , which is unique μ -a.e., so that

$$Y = T(X), \quad (1.3)$$

see work by Gigli for sharp conditions guaranteeing the existence of deterministic couplings [16]. For further background on optimal transport and the 2-Wasserstein metric, we refer the reader to one of the many excellent textbooks on the subject [1, 2, 13, 23, 27, 28].

A notable difference between classical AA, in Eqs. (1.1), and (WAA) is that, in classical AA, the support of the admissible convex polytope Ω is constrained to be in the support of the empirical measure $\mu = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$. On the other hand, in (WAA), proximity to the measure μ is enforced by the fact that the uniform probability measure on Ω has minimal 2-Wasserstein distance from μ . This leads to two notable differences between AA and (WAA). In classical AA, given an empirical measure μ and its archetypal polytope Ω , *any other empirical measure* ν that agrees with μ on Ω^c will have the exact same archetypal polytope Ω , regardless of how μ and ν differ inside Ω . On the other hand, (WAA) is sensitive to changes of the measure μ anywhere in $\mathbb{R}^d$, in proportion to the size of that change in the W_2 metric. A second key difference between AA and (WAA) is that there exist empirical measures $\mu = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$ for

which one can construct perturbations of the measure by the addition of a single Dirac mass that lead to an arbitrarily large changes in the archetypal polytope. In contrast, by definition, (WAA) is robust against outliers that correspond to small perturbations in W_2 .

Alternative generalizations of the classical AA problem, based on probabilistic interpretations, can be found in [24], where data are assumed to be generated from a parametric model and the corresponding archetypes are found using the maximum likelihood estimation, resulting in a similar formulation as (1.1). For nonparametric approaches, aside from the Euclidean and 2-Wasserstein metrics, other discrepancy measures such as the Kullback–Leibler (KL) divergence may also be used. However, a KL divergence would treat outliers of a given mass in the same way regardless of their spatial closeness, whereas a Wasserstein approach places a larger focus on outliers that are farther from the bulk of the dataset. More generally, one could also consider p -Wasserstein metrics, where larger choices of p would further penalize the distance from the dataset. In this way, the choice of metric or statistical divergence is problem-dependent and encodes modeling assumptions about the structure of the dataset.

Our result is similar in spirit to work by Digne et. al. [11], which considered the problem of surface reconstruction. As in the present case, Digne et. al. seek a shape for which the uniform distribution on that shape is as close as possible to the data distribution in the 2-Wasserstein metric. However, the goal of this previous work was to consider a data distribution concentrated on a lower dimensional surface and find a general simplicial complex that best approximates this surface.

When $d = 1$, there is a closed-form solution for the 2-Wasserstein metric, and it is straightforward to directly compute the unique minimizer of (WAA), as previously done in related work by Cuesta-Albertos, Matrán, and Rodríguez-Rodríguez [8]; see below for a discussion of this and other related previous works. Our first main result is that when $d = 2$, a solution of (WAA) exists, provided that μ is absolutely continuous. We summarize the results in the one and two-dimensional setting in the following theorem.

Theorem 1.1 (Existence of minimizer in one and two dimensions) *If $d = 1$ and $k \geq 2$, then for all $\mu \in \mathcal{P}_2(\mathbb{R})$, there exists a unique solution of (WAA), and this solution may be expressed explicitly as $\Omega = [a, b]$ for*

$$a = 4C_0 - 6C_1, \quad b = 6C_1 - 2C_0, \quad C_0 = \int_{\mathbb{R}} x \, d\mu(x), \quad C_1 = \int_{\mathbb{R}} xF(x) \, d\mu(x), \quad (1.4)$$

where F is the cumulative distribution function (CDF) of μ . If $d = 2$ and $k \geq 3$, then for all $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^2)$, there exists a solution of (WAA).

In the two-dimensional case, our proof of Theorem 1.1 relies on an explicit characterization of the closure of $\{1_{\Omega}/|\Omega| : \Omega \in S_k\}$ in the narrow topology; see Proposition 3.1. The key difficulty in proving a minimizer exists arises when considering the possibility that the minimizing sequence converges to a limit point with lower dimensional support. Note, in particular, that limit points with lower dimensional support are in




Fig. 1 In this figure, we illustrate how limit points of $\{1_{\Omega}/|\Omega| : \Omega \in S_k\}$ in the narrow topology are not, in general, given by uniform measures on their support. For example, the top row of this figure illustrates the limit of a sequence of triangles Ω_{δ} , where the base of the triangle remains fixed, while the height of the triangle δ converges to zero. The second row of this figure shows the limit of the corresponding probability measures. We see that the limiting measure ν_0 is supported on the one-dimensional interval prescribed by the base of the triangle, but its density with respect to $\mathcal{L}^1$ is not constant, but piecewise linear, concentrating more mass in the interior of the interval, where the third vertex converged. For reference, in the bottom right hand panel, we plot the distribution of ν_0 in the plane in black, with dashed points representing infinity, and we plot the density ν_0 with respect to one-dimensional Lebesgue measure in blue. In particular, we see that $\lim_{\delta \rightarrow 0} 1_{\Omega_{\delta}}/|\Omega_{\delta}| = \nu_0 \neq 1_{\Omega_0}/|\Omega_0|$

general *not* uniformly distributed on their support; see Fig. 1. For this reason, not all limit points have a valid interpretation in terms of *archetypal analysis*, in which the k vertices or *archetypes* of their support completely describe the measure via their convex combinations. For this reason, we must rule out the possibility that a limit point of this type achieves the infimum of $W_2(\mu, 1_{\Omega}/|\Omega|)$ over $\Omega \in S_k$. We succeed in doing this when $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^2)$ by adapting a perturbation argument of Cuesta-Albertos, et. al. [7] to the case of convex polygons. In particular, we construct perturbations around degenerate limit points that are both feasible for our constraint set and improve the value of the objective function. It remains an open question how to extend this technique to μ which is not absolutely continuous with respect to $\mathcal{L}^2$; see Remark 3.2. Furthermore, our approach to proving the value of the objective function strongly leverages the structure of the 2-Wasserstein metric and a different approach would be needed to extend the result to p -Wasserstein metrics for $p \neq 2$.

While this perturbative approach allows us to overcome the difficulty of degenerate limit points in two dimensions, our approach fails in dimensions $d \geq 3$, since in the higher dimensional setting, edges of polygons can cross in the limit, creating artificial vertices in the perturbations we consider, so that they no longer belong to the constraint set; see Remark 3.3 and Fig. 2. For this reason, it remains an open problem whether minimizers of (WAA) exist for dimensions $d \geq 3$.

Finally, in terms of uniqueness, we observe that, in dimensions higher than one, the minimizer is not unique: for example, if μ is the uniform distribution on the unit ball, any rotation of an optimal polytope would also be optimal. Understanding whether uniqueness holds up to such invariances, remains an interesting open question.

As described above, our approach to proving the existence of (WAA) in two dimensions proceeds by adapting a perturbation argument of Cuesta et al. [7], which considered a related problem: given $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$, find the *convex set* Ω that minimizes $W_2(\mu, 1_\Omega/|\Omega|)$. This built on previous work by the same authors, which also examined optimal W_2 approximation of measures $\mu \in \mathcal{P}(\mathbb{R})$ in one dimension by empirical measures, uniform measures on intervals, and ellipsoids [8]. The motivation of this work was to describe the *shape* and *flatness* of a measure μ . Subsequent work by Belili and Heinich [4] considered W_2 approximation of a measure μ over more general types of measures, including uniform measures on convex sets and uniform measures on sets of the form $\Omega_1 \setminus \Omega_2$ where Ω_1, Ω_2 are convex. Furthermore, Belili and Heinich succeeded in proving the existence of a closest approximation under more general hypotheses on μ : instead of requiring $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$, they require that the affine hull of the support of μ is d -dimensional. However, an essential hypothesis in the work of Belili and Heinich is that the class of approximating measures satisfies certain symmetry properties [4, condition (3)]. For example, if ν_0 were the limiting measure shown in Fig. 1, there would have to exist a triangle Ω so that the first marginal of $1_\Omega/|\Omega|$ is ν_0 and its second marginal is symmetric about the origin. Unfortunately, no such triangle exists. For this reason, while our study of the existence of solutions to (WAA) is closely related to the aforementioned works, the fact that (WAA) constrains Ω to be a k -gon requires the development of new techniques.

Motivated by the fact that the key challenge in proving the existence of minimizers to (WAA) arises when the support of the minimizing sequence collapses to a lower dimensional set, we introduce a regularized version of the problem that prevents this degeneracy. For $m \geq 1$, the m -Rényi entropy is given by

$$\mathcal{U}_m(\nu) = \begin{cases} \frac{1}{m-1} \int \nu(x)^m dx & \text{if } \nu \ll \mathcal{L}^d \text{ and } d\nu(x) = \nu(x) dx, \\ +\infty & \text{otherwise.} \end{cases}$$

In particular, if $\nu = 1_\Omega/|\Omega|$ for $\Omega \in \mathcal{S}_k$, we have

$$\mathcal{U}_m(1_\Omega/|\Omega|) = \begin{cases} \frac{1}{(m-1)|\Omega|^{m-1}} & \text{if } m > 1, \\ -\log(|\Omega|) & \text{if } m = 1. \end{cases} \quad (1.5)$$

For fixed $\varepsilon > 0$, $k \geq d + 1$, $m \geq 1$, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we consider the regularized problem:

$$\min_{\Omega \in \mathcal{S}_k} W_2(\mu, 1_\Omega/|\Omega|) + \varepsilon \mathcal{U}_m(1_\Omega/|\Omega|). \quad (\text{WAA}_\varepsilon)$$

Due to the fact that the regularization prevents the minimizer from collapsing to a lower dimensional set, we are able to obtain the existence of solutions in all dimensions, for general $\mu \in \mathcal{P}_2(\mathbb{R}^d)$.

Theorem 1.2 (Existence of minimizer for (WAA_ε) , in all dimensions) *Fix $\varepsilon > 0$, $k \geq d + 1$, $m \geq 1$, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a solution of (WAA_ε) .*

The proof of Theorem 1.2 follows via Prokhorov's theorem and the lower semi-continuity of the Wasserstein metric and Rényi entropy in the narrow topology. More generally, the regularizing role of the Rényi entropy could be played by any nonnegative, decreasing function of the volume $|\Omega|$ that diverges to $+\infty$ as $|\Omega| \rightarrow 0$ and is lower semicontinuous with respect to narrow convergence of the probability measures $1_\Omega/|\Omega|$. Uniqueness again fails in dimensions higher than two, due to the rotational invariance of the Rényi entropy.

Remark 1.3 (Limit as $\varepsilon \rightarrow 0$) If $d = 1$ or 2 and $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$, similar arguments as in the proof of Theorem 1.1 can be used to show that, for any $\varepsilon_n \rightarrow 0$ as $n \rightarrow +\infty$, the solutions to (WAA_ε) with $\varepsilon = \varepsilon_n$, form a minimizing sequence to (WAA) and converge (up to a subsequence) to a solution of (WAA) . This shows the consistency of the regularized problem with the original problem, as the regularization is removed.

Our final results consider a practical application of the Wasserstein archetypal analysis to data. In particular, we seek to understand how solutions of the regularized problem (WAA_ε) behave when a measure $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is approximated by a sequence of empirical measure $\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, $X_i \stackrel{\text{iid}}{\sim} \mu$. We show that almost surely and up to a subsequence, optimizers Ω_n for the empirical measure μ_n converge as $n \rightarrow \infty$ to an optimizer Ω for μ . We also provide a convergence rate in terms of the value of the objective function, based on a sharp estimate in [15].

Theorem 1.4 (Consistency and convergence for (WAA_ε)) *Fix $\varepsilon > 0$, $k \geq d + 1$, $m \geq 1$, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Suppose*

$$\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}, \quad X_i \stackrel{\text{iid}}{\sim} \mu.$$

For each $n \in \mathbb{N}$, let $\{\Omega_n\}_{n \in \mathbb{N}}$ be a minimizer to (WAA_ε) for μ_n . Then the following hold:

(i) *Almost surely, there exists $\Omega \in S_k$ so that, up to a subsequence,*

$$1_{\Omega_n}/|\Omega_n| \rightarrow 1_\Omega/|\Omega| \text{ narrowly,} \quad (1.6)$$

and Ω solves (WAA_ε) for μ ;

(ii) If $M_q(\mu) = \int_{\mathbb{R}^d} |x|^q d\mu(x) < \infty$ for some $q > 4$, then there exists a constant $C(q, d) > 0$ such that

$$\begin{aligned} & \mathbb{E} \left[W_2 \left(\mu, \frac{1}{|\Omega_n|} 1_{\Omega_n} \right) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) - \left(\min_{\Omega \in S_k} W_2 \left(\mu, \frac{1}{|\Omega|} 1_{\Omega} \right) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|) \right) \right] \\ & \leq C(q, d) M_q^{1/q}(\mu) \cdot \begin{cases} n^{-\frac{1}{4}} & d < 4 \\ n^{-\frac{1}{4}} \sqrt{\log n} & d = 4 \\ n^{-\frac{1}{d}} & d > 4 \end{cases} \end{aligned}$$

In addition to ensuring consistency of WAA under empirical approximation, the previous theorem is also relevant from the perspective of numerical methods. Our final main result is the development of a numerical method for solving (WAA) and (WAA_ε), based on approximating a given measure $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ by a sequence of finitely supported measures $\mu_n = \sum_{i=1}^n \delta_{x_i} m_i$ and approximating the solution of (WAA) or (WAA_ε) for μ_n . Such an approximation greatly simplifies the problem numerically, since the Wasserstein distance $W_2(\mu_n, 1_{\Omega}/|\Omega|)$ can now be computed using the *semidiscrete* method introduced by Mérigot [20]. Based on this perspective, we introduce an alternating gradient-based method to approximate the optimizer; see Sect. 5 and Algorithm 1.

In Sect. 6, we consider several numerical experiments for approximating solutions of (WAA) and (WAA_ε) in two dimensions. First, we consider the case when μ is the uniform distribution on a disk (Sect. 6.1) and a normal distribution (Sect. 6.2). In this case, the solution for each k is a regular k -gon. Next, we study the sensitivity to the parameter ε in (WAA_ε) (Sect. 6.3) and provide an example that demonstrates the non-convexity of the energy landscape in (WAA) (Sect. 6.4). Finally, we consider an example of the early spread of the COVID-19 virus in the U.S. (Sect. 6.5). We conclude in Sect. 7 with a discussion of our contributions and directions for future work.

2 Preliminaries

Given a closed, convex set $S \subseteq \mathbb{R}^d$, let $\pi_S : \mathbb{R}^d \rightarrow S$ denote the projection on the set. Given any (Borel) measurable function $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\mu \in \mathcal{P}(\mathbb{R}^d)$, the *push-forward of μ through T* , denoted $T\#\mu$, is the probability measure defined by

$$\int_{\mathbb{R}^d} f(T(x)) d\mu(x) = \int_{\mathbb{R}^d} f(x) d(T\#\mu)(x), \quad \forall f \text{ bounded and (Borel) measurable.}$$

We now recall several facts about the Wasserstein metric and uniform measures on convex polygons. First, recall that in one dimension, we have an explicit formula for the Wasserstein metric in terms of cumulative distribution functions (CDFs) [27, Theorem 2.18]. In particular, if $\mu, \nu \in \mathcal{P}_2(\mathbb{R})$ have CDFs F and G , then

$$W_2(\mu, \nu) = \left(\int_0^1 |F^{-1}(t) - G^{-1}(t)|^2 dt \right)^{1/2}, \quad (2.1)$$

where the *generalized inverse* of F is given by $F^{-1}(t) = \inf\{x \in \mathbb{R} : F(x) > t\}$.

Next, we recall relevant notions of convergence, beginning with narrow convergence.

Definition 2.1 (Narrow convergence of probability measures in $\mathbb{R}^d$) We say that $\{\mu_n\}_{n=1}^\infty \subseteq \mathcal{P}(\mathbb{R}^d)$ converges narrowly to $\mu \in \mathcal{P}(\mathbb{R}^d)$ if

$$\int_{\mathbb{R}^d} f(x) \, d\mu_n(x) \rightarrow \int_{\mathbb{R}^d} f(x) \, d\mu(x), \quad \forall f \in C_b(\mathbb{R}^d), \quad (2.2)$$

where $C_b(\mathbb{R}^d)$ denotes the set of bounded continuous functions on $\mathbb{R}^d$.

In the probability literature, narrow convergence is also called weak convergence or convergence in distribution. Narrow convergence is slightly weaker than convergence in W_2 , and they are equivalent when the second moments also converge.

Lemma 2.2 ([2, Remark. 7.1.11]) *Given $\mu_n, \mu \in \mathcal{P}_2(\mathbb{R}^d)$, the following statements are equivalent:*

$$\lim_{n \rightarrow +\infty} W_2(\mu_n, \mu) = 0 \iff \mu_n \rightarrow \mu \text{ narrowly and } \lim_{n \rightarrow +\infty} M_2(\mu_n) = M_2(\mu). \quad (2.3)$$

Recall that, for any measures $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, if we consider their translations to have mean zero, $\mu' = (\text{Id} - \int x \, d\mu(x))\#\mu$ and $\nu' = (\text{Id} - \int x \, d\nu(x))\#\nu$, then

$$W_2^2(\mu, \nu) = W_2^2(\mu', \nu') + \left| \int x \, d\mu(x) - \int x \, d\nu(x) \right|^2. \quad (2.4)$$

See, for example, [7, Proposition 2.5].

Remark 2.3 (mean of minimizers of WAA and (WAA_ε)) An immediate consequence of Eq. (2.4) is that any minimizer Ω of either (WAA) or the regularized problem (WAA_ε) must have the same mean as μ , that is,

$$\frac{1}{|\Omega|} \int_{\Omega} x \, dx = \int_{\mathbb{R}^d} x \, d\mu(x).$$

Next, we recall an elementary lemma for the CDFs of real-valued random variables. For the reader's convenience, we include a proof.

Lemma 2.4 ([7, Equation 2] and [27, Theorem 2.18]) *Let X be a real-valued random variable with CDF $F(x)$. Assume $\mathbb{E}[|X|] < +\infty$ and the law of X is not a Dirac mass. Then*

$$\mathbb{E}[XF(X)] > \frac{1}{2}\mathbb{E}[X]. \quad (2.5)$$

Furthermore, if $X \sim \mu$ for $\mu \in \mathcal{P}_{2,ac}(\mathbb{R})$, then for any $a < b$, $(b-a)(F(x) + a)$ is the optimal transport map from μ to the probability measure $1_{[a,b]}/(b-a)$.

Proof First, we show (2.5). To prove this, note

$$\mathbb{E}[XF(X)] = \int_{\mathbb{R}} xF(x) dF(x) = \frac{1}{2} \int_{\mathbb{R}} x d(F(x)^2) = \frac{1}{2} \mathbb{E}[\tilde{X}],$$

where $\tilde{X}$ has CDF $F(x)^2$. Since $F(x) \in [0, 1]$, $\tilde{X}$ stochastically dominates X , i.e. $\mathbb{P}(\tilde{X} \geq t) = 1 - F(t)^2 \geq 1 - F(t) = \mathbb{P}(X \geq t)$. As a consequence of stochastic dominance [25, Theorem 1.A.8], $\mathbb{E}[\tilde{X}] \geq \mathbb{E}[X]$, with equality holding if and only if $F(x)^2 = F(x)$ for all $x \in \mathbb{R}$. Thus, equality can only hold if $F(x)$ is either equal to zero or one for all $x \in \mathbb{R}$, which implies the law of X is a Dirac mass. Consequently, we conclude $\mathbb{E}[XF(X)] = \mathbb{E}[\tilde{X}]/2 > \mathbb{E}[X]/2$.

Now, suppose $\mu \in \mathcal{P}_{2,ac}(\mathbb{R})$. [27, Remark 2.19] ensures that $(b-a)(F(x) + a)$ is the optimal transport map from the law of X to $1_{[a,b]}/(b-a)$. $\square$

Our next lemma concerns optimal transport maps from a measure $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$ to a measure $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ with $\text{supp } \nu \subseteq \pi_1(\mathbb{R}^d) = \{(x_1, \dots, x_d) \in \mathbb{R}^d : x_i = 0 \ \forall i = 2, \dots, d\}$.

Lemma 2.5 *Given $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$ and $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ with $\text{supp } \nu \subseteq \pi_1(\mathbb{R}^d)$, if T is the optimal transport map from μ to ν , then $T = (T_1, 0, \dots, 0)$, where $T_1: \mathbb{R} \rightarrow \mathbb{R}$ is the optimal transport map from $\pi_1\#\mu$ to $\pi_1\#\nu$.*

Proof Since $\text{supp } \nu \subseteq \pi_1(\mathbb{R}^d)$, for any $Y = (Y_1, \dots, Y_d)$ with $Y \sim \nu$, we have $Y_i = 0$ almost everywhere, for $i = 2, \dots, d$. Thus, by definition of the Wasserstein metric in Eq. (1.2),

$$\begin{aligned} W_2^2(\mu, \nu) &= \inf_{X \sim \mu, Y \sim \nu} \mathbb{E}[(X_1 - Y_1)^2] + \sum_{i=2}^d \mathbb{E}[(X_i - Y_i)^2] \\ &= \inf_{X \sim \mu, Y \sim \nu} \mathbb{E}[(X_1 - Y_1)^2] + \sum_{i=2}^d \mathbb{E}[X_i^2] \\ &= M_2(\pi_{2,\dots,d}\#\mu) + \inf_{X_1 \sim (\pi_1\#\mu), Y_1 \sim (\pi_1\#\nu)} \mathbb{E}[(X_1 - Y_1)^2], \end{aligned}$$

where the second term in the last expression corresponds with $W_2^2(\pi_1\#\mu, \pi_1\#\nu)$. Since $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$, $\pi_1\#\mu \in \mathcal{P}_{2,ac}(\mathbb{R})$, so there exists an optimal transport map T_1 from $\pi_1\#\mu$ to $\pi_1\#\nu$; see Eq. (1.3). Furthermore, the above computation shows that, $\tilde{T} = (T_1, 0, \dots, 0)$ is optimal from μ to ν . Uniqueness of optimal transport maps for $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$ then gives the result. $\square$

Another useful fact is that, for uniform distributions on convex k -gons in $\mathbb{R}^d$, uniformly bounded second moments imply uniformly bounded support, as well as convergence, up to a subsequence. This is proved in a slightly different setting in [7, Lemmas 3.2–3.3] and [4, Proposition 1]. We recall the result in the present setting for the reader's convenience, with the proof deferred to Appendix 8.1.

Lemma 2.6 *Fix $d \geq 1$ and $k \geq d + 1$. Suppose $\{\Omega_n\}_{n=1}^{+\infty} \subseteq S_k$ satisfies $\sup_n M_2(1_{\Omega_n}/|\Omega_n|) < +\infty$.*

- (i) There exists $R > 0$ so that $\Omega_n \subseteq B_R(0) := \{x \in \mathbb{R}^d : |x| \leq R\}$ for all $n \in \mathbb{N}$.
- (ii) Up to a subsequence, the vertices of Ω_n converge, and if we let Ω be their convex hull, then there exists $v \in \mathcal{P}_2(\mathbb{R}^d)$ so that $1_{\Omega_n}/|\Omega_n| \rightarrow v$ narrowly, $1_{\Omega_n}(x) \rightarrow 1_{\Omega}(x)$ pointwise $\mathcal{L}^d$ -almost everywhere, and $\text{supp } v = \Omega$.

Our next lemma identifies the density of a narrowly convergent sequence in $\mathcal{P}_{2,ac}(\mathbb{R}^d)$, with uniformly bounded support and density.

Lemma 2.7 Consider $v_n \in \mathcal{P}_{2,ac}(\mathbb{R}^d)$, where $dv_n(x) = v_n(x) d\mathcal{L}^d(x)$, and suppose

- (i) v_n has uniformly bounded support,
- (ii) $\sup_n \|v_n\|_{L^\infty(\mathbb{R}^d)} < +\infty$,
- (iii) $v_n \rightarrow v \in \mathcal{P}_2(\mathbb{R}^d)$ narrowly.

Then, if there exists $f \in L^\infty(\mathbb{R}^d)$ so that $v_n(x) \rightarrow f(x)$ pointwise, we must have $dv = f(x) d\mathcal{L}^d(x)$.

Proof Fix $\varphi \in C_b(\mathbb{R}^d)$. Since $v_n(x)$ is uniformly bounded, with uniformly bounded support, by the dominated convergence theorem.

$$\lim_{n \rightarrow +\infty} \int \varphi(x) v_n(x) d\mathcal{L}^d(x) = \int \varphi(x) f(x) d\mathcal{L}^d(x). \quad (2.6)$$

Furthermore, narrow convergence of v_n to v ensures that the left-hand side coincides with $\int \varphi(x) dv(x)$. Thus, the Riesz–Markov–Kakutani representation theorem implies $dv(x) = f(x) dx$. $\square$

We close with the following technical lemma, focused on the two-dimensional case, which characterizes the projection of measures of the form $1_\Omega/|\Omega|$ onto a one-dimensional affine space.

Lemma 2.8 Let $d = 2$ and $k \geq d + 1$. For every $\Omega \in S_k$ and affine space $V \subseteq \mathbb{R}^d$ with $\dim(V) = 1$, the density of the marginal distribution of $1_\Omega/|\Omega|$ along V , denoted $\pi_V \# 1_\Omega/|\Omega|$, with respect to one dimensional Lebesgue measure on V , is a piecewise linear function on Ω , with at most k vertices, that is concave on its support.

Proof When $k = 3$, the result can be checked by brute force. For $k > 3$, we observe that any measure of the form $1_\Omega/|\Omega|$ for $\Omega \in S_k$ can be written as a conical combination of uniform distributions on triangles,

$$1_\Omega/|\Omega| = \sum_{i=1}^l 1_{T_i}/|\Omega|.$$

Thus, by linearity of the push forward,

$$\pi_V \# 1_\Omega/|\Omega| = \sum_{i=1}^l (\pi_V \# 1_{T_i})/|\Omega|.$$

By the $k = 3$ case, we know that $\pi_V \# 1_\Omega/|\Omega|$ is piecewise linear. The concavity part follows from Brunn’s concavity principle [3, Theorem 1.2.1]. $\square$

3 Existence of Minimizers

To prove Theorem 1.1 on the existence of a minimizer for (WAA), we begin with the following lemma, characterizing the closure of the constraint set in the narrow topology. We consider three cases, depending on the dimension of the affine hull of the support of the measure.

Proposition 3.1 *Let $d \geq 1$ and $k \geq d + 1$, and suppose $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ is the narrow limit of $1_{\Omega_n}/|\Omega_n|$, for $\{\Omega_n\}_{n=1}^{+\infty} \subseteq S_k$ with $\sup_n M_2(1_{\Omega_n}/|\Omega_n|) < +\infty$. Let $\Omega = \text{supp } \nu$.*

- (i) *If $\dim(\text{aff}(\Omega)) = 0$, then ν is a Dirac mass at $\text{supp}(\nu)$, i.e., $\nu = \delta_{\text{supp}(\nu)}$.*
- (ii) *If $d = 2$ and $\dim(\text{aff}(\Omega)) = 1$, then the projection of ν onto $\text{aff}(\Omega)$ is absolutely continuous with respect to the Lebesgue measure on $\text{aff}(\Omega)$ and has a piecewise linear density with at most k vertices, which is concave on its support.*
- (iii) *If $\dim(\text{aff}(\Omega)) = d$, then $\nu = 1_{\Omega}/|\Omega|$ for $\Omega \in S_k$.*

Proof of Proposition 3.1 First, recall from Lemma 2.6, that Ω is a convex k -gon, Ω_n and Ω are uniformly bounded in $\mathbb{R}^d$, and 1_{Ω_n} converges pointwise a.e. to 1_{Ω} .

We now consider part (i). This result is immediate, since the limit of any narrowly convergent sequence of probability measures must be a probability measure, and the only probability measure concentrated on a point is a Dirac mass at that point.

Next, we show part (ii), under the assumption that $d = 2$. Since Ω is a bounded, convex k -gon with $\dim(\text{aff}(\Omega)) = 1$, Ω must be a bounded line segment. For simplicity of notation, let $V = \text{aff}(\Omega)$, and let $\pi_V \# (1_{\Omega_n}/|\Omega_n|)$ denote the marginal distribution of $1_{\Omega_n}/|\Omega_n|$ along V . By Lemma 2.8, the density of $\pi_V \# (1_{\Omega_n}/|\Omega_n|)$ with respect to one-dimensional Lebesgue measure on V , which we denote by f_n , is a piecewise linear function with at most k vertices, which is concave on its support. Without loss of generality, suppose V coincides with the x_1 -axis, so f_n is a function of x_1 .

We now show that f_n is uniformly bounded for all n . To do this, we begin by showing that, up to a subsequence, $\inf_n \mathcal{L}^1(\pi_V(\Omega_n)) > 0$. By Lemma 2.6, the vertices $\{v_i^n\}_{i=1}^k$ of Ω_n converge to $\{v_i\}_{i=1}^k$, where Ω is the convex hull of these limit points. By definition of V , $v^i = (v_1^i, 0)$ and $\pi_V(v^i) = v_1^i$ for all $i = 1, \dots, k$. Furthermore, note that $\pi_V(\Omega_n)$ is a sequence of uniformly bounded intervals on the x_1 -axis, with nonempty interior with respect to $\mathbb{R}$. Thus, by Lemma 2.6, up to a subsequence, the endpoints of the intervals converge, and we may let $\Omega' \subseteq \mathbb{R}$ denote the convex hull of their limits. By uniqueness of limits and the continuity of π_V , we have $\Omega' = \pi_V(\Omega)$, so Lemma 2.6 ensures $1_{\pi_V(\Omega_n)}(x_1) \rightarrow 1_{\pi_V(\Omega)}(x_1)$ pointwise $\mathcal{L}^1$ almost everywhere. Thus, by the dominated convergence theorem,

$$\lim_{n \rightarrow +\infty} \mathcal{L}^1(\pi_V(\Omega_n)) = \lim_{n \rightarrow +\infty} \int 1_{\pi_V(\Omega_n)} d\mathcal{L}^1 = \int 1_{\pi_V(\Omega)}(x_1) dx_1 = \mathcal{L}^1(\pi_V(\Omega)).$$

Since $\dim(\text{aff}(\Omega)) = 1$, $\mathcal{L}^1(\pi_V(\Omega)) > 0$. Thus, up to another subsequence, we may assume that $\inf_n \mathcal{L}^1(\pi_V(\Omega_n)) > 0$.

Since f_n is a concave piecewise linear density, for any $x_1 \in \mathbb{R}$, the area of the triangle with base $\mathcal{L}^1(\pi_V(\Omega_n))$ and height $f_n(x_1)$ is always smaller than $\int f_n d\mathcal{L}^1$.

Therefore, for all $n \in \mathbb{N}$ and $x_1 \in \pi_V(\Omega_n)$,

$$\begin{aligned} \frac{1}{2} \left(\inf_m \mathcal{L}^1(\pi_V(\Omega_m)) \right) f_n(x_1) &\leq \frac{1}{2} \mathcal{L}^1(\pi_V(\Omega_n)) f_n(x_1) \\ &\leq \int f_n d\mathcal{L}^1 = 1 \Rightarrow f_n(x_1) \leq 2 \left(\inf_m \mathcal{L}^1(\pi_V(\Omega_m)) \right)^{-1}. \end{aligned}$$

This shows that the densities f_n are uniformly bounded in n .

We now seek to apply Lemma 2.7. By the Continuous Mapping Theorem,

$$f_n d\mathcal{L}^1 = \pi_V \# (1_{\Omega_n} / |\Omega_n|) \rightarrow \pi_V \# \nu \text{ narrowly.}$$

Given that $f_n(x)$ is piecewise linear with at most k vertices, we may let $x_{n,1} \leq \dots \leq x_{n,k}$ denote the vertices in increasing order. Since the support of f_n is bounded, up to a subsequence, $x_{n,i}$ converges to some x_i for all $i = 1, \dots, k$, and since f_n is uniformly bounded, $f_n(x_{n,i})$ converges to some y_i as for all $i = 1, \dots, k$. Let f be the piecewise linear functions that interpolates between $\{(x_i, y_i)\}_{i=1}^k$. Then $f_n \rightarrow f$ pointwise so, by Lemma 2.7, $d(\pi_V \# \nu) = f d\mathcal{L}^1$. Finally, since f_n are concave on their support, so is f . This completes the proof of part (ii).

It remains to show part (iii), which holds for general $d \geq 1$. Since 1_{Ω_n} converges pointwise $\mathcal{L}^d$ a.e. to 1_Ω and both have uniformly bounded support, the dominated convergence theorem ensures

$$\lim_{n \rightarrow +\infty} |\Omega_n| = \lim_{n \rightarrow +\infty} \int 1_{\Omega_n} d\mathcal{L}^d = \int 1_\Omega d\mathcal{L}^d = |\Omega|.$$

Therefore, up to a subsequence, we may assume $\inf_n |\Omega_n| > 0$. Applying the dominated convergence theorem again, we conclude that for any Borel measurable set $A \subseteq \mathbb{R}^d$,

$$\lim_{n \rightarrow +\infty} \frac{1}{|\Omega_n|} \int_{\Omega_n} 1_A d\mathcal{L}^d = \lim_{n \rightarrow +\infty} \frac{1}{|\Omega_n|} \int 1_A 1_{\Omega_n} d\mathcal{L}^d = \frac{\int 1_A 1_\Omega d\mathcal{L}^d}{|\Omega|}.$$

This shows that $1_{\Omega_n} / |\Omega_n| \rightarrow 1_\Omega / |\Omega|$ strongly as measures, hence narrowly. By uniqueness of limits, we conclude that $\nu = 1_\Omega / |\Omega|$. Furthermore, since $\dim(\text{aff}(\Omega)) = d$, $\Omega \subseteq \mathbb{R}^d$ must have non-empty interior, so $\Omega \in S_k$. $\square$

We now turn to the proof of Theorem 1.1.

Proof of Theorem 1.1 First, suppose $d = 1$. Let F be the CDF of μ , and let G be the CDF of the probability measure $\frac{1}{|b-a|} 1_{[a,b]}$. Then $G^{-1}(t) := a + (b-a)t$, so, by Eq. (2.1),

$$\begin{aligned} W_2^2 \left(\mu, \frac{1}{|b-a|} 1_{[a,b]} \right) &= \int_0^1 (F^{-1}(t) - G^{-1}(t))^2 dt \\ &= \int_0^1 (F^{-1}(t) - [a + (b-a)t])^2 dt \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{6} \begin{pmatrix} a \\ b \end{pmatrix}^t \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} - 2 \begin{pmatrix} C_0 - C_1 \\ C_1 \end{pmatrix}^t \begin{pmatrix} a \\ b \end{pmatrix} + C_2 \\
&=: \varphi(a, b),
\end{aligned}$$

where

$$\begin{aligned}
C_0 &= \int_0^1 F^{-1}(t) \, dt = \int_{\mathbb{R}^d} x \, d\mu(x), \\
C_1 &= \int_0^1 t F^{-1}(t) \, dt = \int_{\mathbb{R}^d} x F(x) \, d\mu(x), \\
C_2 &= \int_0^1 \left(F^{-1}(t) \right)^2 \, dt = \int_{\mathbb{R}^d} x^2 \, d\mu(x).
\end{aligned}$$

The result then follows from the fact that $\varphi(a, b)$ is a strictly convex quadratic function, with a unique minimizer given in (1.4).

Now, we turn to the case $d = 2$. By Eq. (2.4), we may assume without loss of generality that μ has mean zero and restrict the minimization problem to the set $\{\Omega \in S_k : \int_{\Omega} x \, d\mathcal{L}^d(x) = 0\}$. Consider a minimizing sequence $\Omega_n \in S_k$ with $\int_{\Omega_n} x \, d\mathcal{L}^d(x) = 0$ so that

$$\lim_{n \rightarrow +\infty} W_2(\mu, 1_{\Omega_n}/|\Omega_n|) = \inf_{\Omega \in S_k} W_2(\mu, 1_{\Omega}/|\Omega|). \quad (3.1)$$

By the triangle inequality,

$$\begin{aligned}
\sup_n M_2(1_{\Omega_n}/|\Omega_n|) &= \sup_n W_2^2(1_{\Omega_n}/|\Omega_n|, \delta_0) \\
&\leq \sup_n 2 \left(W_2^2(\mu, \delta_0) + W_2^2(\mu, 1_{\Omega_n}/|\Omega_n|) \right) < +\infty.
\end{aligned}$$

Thus, Lemma 2.6 ensures that, up to a subsequence, there exists $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and a closed, convex k -gon Ω so that $1_{\Omega_n}/|\Omega_n| \rightarrow \nu$ narrowly and $\text{supp } \nu = \Omega$. Since $1_{\Omega_n}/|\Omega_n|$ has mean zero for all $n \in \mathbb{N}$ and uniformly bounded support, ν also has mean zero. By the lower semicontinuity of the Wasserstein metric with respect to narrow convergence,

$$W_2(\mu, \nu) \leq \lim_{n \rightarrow +\infty} W_2(\mu, 1_{\Omega_n}/|\Omega_n|) = \inf_{\Omega \in S_k} W_2(\mu, 1_{\Omega}/|\Omega|). \quad (3.2)$$

Thus, it suffices to show that $\nu = 1_{\Omega}/|\Omega|$ for $\Omega \in S_k$ in order to conclude that a minimizer exists. By Proposition 3.1, if $\dim(\text{aff}(\Omega)) = 2$, the result holds. Thus, it remains to exclude the possibility that $\dim(\text{aff}(\Omega)) \leq 1$. We accomplish this by showing that, if either $\dim(\text{aff}(\Omega)) = 0$ or $\dim(\text{aff}(\Omega)) = 1$, then there exists $\tilde{\Omega} \in S_k$ with $W_2(\mu, 1_{\tilde{\Omega}}/|\tilde{\Omega}|) < W_2(\mu, \nu)$, contradicting Eq. (3.2).

Let $X = (X_1, X_2)$ be a random variable with distribution μ , and let $F(x_1)$ denote the CDF of X_1 . Let $X_2|X_1$ denote the conditional distribution of X_2 given X_1 , with

CDF $F_{X_1}(x_2)$. By Lemma 2.4, F is the optimal transport map from $\pi_1\#\mu$ to $1_{[0,1]}$, $F - 1$ is the optimal transport map from $\pi_1\#\mu$ to $1_{[-1,0]}$, and, almost everywhere, F_{X_1} is the optimal transport map from the law of $X_2|X_1$ to $1_{[0,1]}$. Suppose that

$$\begin{aligned} g: \mathbb{R} \rightarrow [0, +\infty) \text{ is compactly supported, } \int g = 1, \text{ and, on its support,} \\ g \text{ is concave and piecewise linear, with at most } k \text{ vertices.} \end{aligned} \quad (3.3)$$

In particular, we have

$$g(x_1) > 0 \text{ for } x_1 \in \text{int}(\text{supp } g). \quad (3.4)$$

For $\alpha \in \mathbb{R} \setminus \{0\}$, $\beta > 0$, consider the following family of convex k -gons:

$$\Omega(\alpha, \beta) := \{(x_1, x_2) : x_2 \in [0, \alpha g(x_1/\beta)]\} \in S_k.$$

Note that $|\Omega(\alpha, \beta)| = \alpha\beta$. Let T_1 be the optimal transport map from $\pi_1\#\mu$ to $\pi_1\#1_{\Omega(1,1)}$. Define the random variables

$$\begin{aligned} Z_1 &= T_1(X_1), \\ Z_2 &= g(T_1(X_1))(F_{X_1}(X_2) - C_\alpha), \quad \text{for } C_\alpha = \begin{cases} 0 & \text{if } \alpha > 0, \\ 1 & \text{if } \alpha < 0, \end{cases} \\ Y &= (\beta Z_1, |\alpha| Z_2). \end{aligned}$$

By construction, $Y \sim 1_{\Omega(\alpha, \beta)}/|\Omega(\alpha, \beta)|$. Since $\pi_1\#1_{\Omega(1,1)} \in \mathcal{P}_{2,ac}(\mathbb{R})$ and $T_1(X_1) \in \text{int}(\pi_1\#1_{\Omega(1,1)}) = \text{int}(\text{supp } g)$ almost everywhere, Eq. (3.4) and Lemma 2.4 ensure

$$\begin{aligned} |\alpha| \mathbb{E}[Z_2 X_2] &= |\alpha| \mathbb{E}[g(T_1(X_1))(F_{X_1}(X_2) - C_\alpha) X_2] \\ &= |\alpha| \mathbb{E}[g(T_1(X_1)) (\mathbb{E}[F_{X_1}(X_2) X_2 | X_1] - C_\alpha \mathbb{E}[X_2 | X_1])], \\ &> |\alpha| \mathbb{E} \left[g(T_1(X_1)) \left(\frac{1}{2} \mathbb{E}[X_2 | X_1] - C_\alpha \mathbb{E}[X_2 | X_1] \right) \right] \\ &= \begin{cases} \frac{|\alpha|}{2} \mathbb{E}[g(T_1(X_1)) (\mathbb{E}[X_2 | X_1])] & \text{if } \alpha > 0, \\ -\frac{|\alpha|}{2} \mathbb{E}[g(T_1(X_1)) (\mathbb{E}[X_2 | X_1])] & \text{if } \alpha < 0. \end{cases} \end{aligned} \quad (3.5)$$

We now apply this construction to rule out the possibility $\dim(\text{aff}(\Omega)) \leq 1$. First, suppose $\dim(\text{aff}(\Omega)) = 0$. By Proposition 3.1 and the fact that ν has mean zero, $\nu = \delta_0$, so $W_2^2(\mu, \nu) = \mathbb{E}[X_1^2 + X_2^2]$. Define

$$g(x_1) = \begin{cases} 1 + x_1 & \text{if } x_1 \in [-1, 0], \\ 1 - x_1 & \text{if } x_1 \in [0, 1]. \end{cases}$$

Then g satisfies (3.3). Furthermore,

$$\begin{aligned} W_2^2(\mu, \nu) - W_2^2(\mu, 1_{\Omega_{\alpha, \beta}}/|\Omega(\alpha, \beta)|) \\ \geq \mathbb{E}[X_1^2 + X_2^2] - \mathbb{E}[(X_1 - \beta Z_1)^2 + (X_2 - |\alpha|Z_2)^2] \\ = 2\beta\mathbb{E}[X_1 Z_1] - \beta^2\mathbb{E}[Z_1^2] + 2|\alpha|\mathbb{E}[Z_2 X_2] - \alpha^2\mathbb{E}[Z_2^2]. \end{aligned} \quad (3.6)$$

We now apply inequality (3.5) to bound this strictly from below. If $\mathbb{E}[g(T_1(X_1))(\mathbb{E}[X_2|X_1])] \geq 0$, then $\mathbb{E}[Z_2 X_2] > 0$ for all $\alpha > 0$. Thus, we may choose $\alpha > 0$ and $\beta > 0$ sufficiently small so that the right-hand side of (3.6) is strictly positive. On the other hand, if $\mathbb{E}[g(T_1(X_1))(\mathbb{E}[X_2|X_1])] \leq 0$, then $\mathbb{E}[Z_2 X_2] > 0$ for all $\alpha < 0$. Thus, we may choose $\alpha < 0$ sufficiently close to zero and $\beta > 0$ sufficiently small so that the right-hand side of (3.6) is strictly positive. In particular, in either case, there exist $\alpha \in \mathbb{R} \setminus \{0\}$ and $\beta > 0$ so that

$$W_2^2(\mu, \nu) > W_2^2(\mu, 1_{\Omega(\alpha, \beta)} / |\Omega(\alpha, \beta)|),$$

which contradicts (3.2). This shows that $\dim(\text{aff}(\Omega)) = 0$ is impossible.

It remains to exclude the possibility that $\dim(\text{aff}(\Omega)) = 1$. We proceed by contradiction, assuming $\dim(\text{aff}(\Omega)) = 1$. Without loss of generality, we may rotate μ and ν so that $\text{aff}(\Omega)$ coincides with the x_1 -axis. By Proposition 3.1, $d(\pi_1 \# \nu) = g \, d\mathcal{L}^1$, for g satisfying (3.3) above. Note that, in this case, $\pi_1 \# (1_{\Omega(1, 1)} / |\Omega(1, 1)|) = \pi_1 \# \nu$. By Lemma 2.5, the optimal coupling between μ and ν is of the form $(X, T(X))$, where $T(x_1, x_2) = (T_1(x_1), 0)$ and T_1 is the optimal transport map from $\pi_1 \# \mu$ to $\pi_1 \# \nu$. Therefore,

$$\begin{aligned} W_2^2(\mu, \nu) - W_2^2(\mu, 1_{\Omega(\alpha, 1)} / |\Omega(\alpha, 1)|) \\ \geq \mathbb{E}[(X_1 - T_1(X_1))^2 + X_2^2] - \mathbb{E}[(X_1 - Z_1)^2 + (X_2 - |\alpha|Z_2)^2] \\ = \mathbb{E}[(X_1 - T_1(X_1))^2 + X_2^2] - \mathbb{E}[(X_1 - T_1(X_1))^2 + (X_2 - |\alpha|Z_2)^2] \\ = 2|\alpha|\mathbb{E}[X_2 Z_2] - \alpha^2\mathbb{E}[Z_2^2]. \end{aligned} \quad (3.7)$$

As before, we apply inequality (3.5) to bound this from below. If $\mathbb{E}[g(T_1(X_1))(\mathbb{E}[X_2|X_1])] \geq 0$, then $\mathbb{E}[X_2 Z_2] > 0$ for all $\alpha > 0$. Thus, we may choose $\alpha > 0$ sufficiently small so that the right-hand side of (3.6) is strictly positive. On the other hand, if $\mathbb{E}[g(T_1(X_1))(\mathbb{E}[X_2|X_1])] \leq 0$, then $\mathbb{E}[X_2 Z_2] > 0$ for all $\alpha < 0$. Thus, we may choose $\alpha < 0$ sufficiently close to zero so that the right-hand side of (3.6) is strictly positive. In particular, in either case, there exists $\alpha \in \mathbb{R} \setminus \{0\}$

$$W_2^2(\mu, \nu) > W_2^2(\mu, 1_{\Omega(\alpha, \beta)} / |\Omega(\alpha, \beta)|),$$

which contradicts (3.2). This shows that $\dim(\text{aff}(\Omega)) = 1$ is impossible. $\square$

Remark 3.2 (regularity of μ when $d = 2$) While in the one-dimensional case, our proof of the existence of minimizers to (WAA) holds for all $\mu \in \mathcal{P}_2(\mathbb{R})$, in the two-dimensional case, we require μ to be absolutely continuous with respect to $\mathcal{L}^2$. This is

used in our application of Lemma 2.4. In particular, we must have the following: after an arbitrary rotation of the coordinate plane, if $X = (X_1, X_2)$ is a random variable with law μ , then, on a set of positive measure, the conditional distribution of X_2 given X_1 , $X_2|X_1$, is not a single Dirac mass. Note that this fails to be true when μ is an empirical measure. The assumption that $\mu \in \mathcal{P}_{2,ac}(\mathbb{R}^2)$ also implies that $\pi_1\#\mu$ and the conditional distribution $d\mu_{x_1}(x_2)$ are absolutely continuous with respect to one dimensional Lebesgue measure, hence that optimal transport maps exist from these measures to any other measure in $\mathcal{P}_2(\mathbb{R}^d)$.

It remains an open question how to extend our result to more general $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, particularly μ with lower dimensional support, such as an empirical measure. While related work due to Belili and Heinich [4] on approximating measures by general convex sets Ω succeeded in weakening the condition on μ to merely require $\dim(\text{aff}(\text{supp } \mu)) = d$, their approach strongly uses symmetry arguments, which fail in our setting.

Remark 3.3 (existence of minimizers for $d \geq 3$) There are two key gaps preventing obtaining the existence of minimizers to (WAA) when $d \geq 3$ by a similar approach. The first is an analog of Proposition 3.1, characterizing how limits of minimizing sequences could degenerate. This becomes more difficult in higher dimensions, as a degeneracy can occur downwards by more than one dimension: for example, a three-dimensional polytope can collapse to a line segment.

The second, more significant, gap is that even given an appropriate analog of Proposition 3.1, the perturbation argument used in the proof of Theorem 1.1 would fail. For example, consider the following sequence $\Omega_n \in S_k$ for $d = 3$ with $k = 4$, illustrated in Fig. 2,

$$\Omega_n = \text{conv} \left(\left\{ (-1, 0, 0), (0, 1, 0), (0, -1, 0), \left(\cos \left(\frac{\pi}{2n} \right), 0, \sin \left(\frac{\pi}{2n} \right) \right) \right\} \right), \quad n \in \mathbb{N}.$$

While the polygons Ω_n converge to the square $\Omega = \text{conv}\{(-1, 0, 0), (0, 1, 0), (0, -1, 0), (1, 0, 0)\}$ in the xy -plane, the measures $1_{\Omega_n}/|\Omega_n|$ narrowly converge to the measure $\nu = f \, d\mathcal{L}^2$, where f is the piecewise linear density, supported on Ω , that interpolates between $(-1, 0, 0)$, $(0, 1, 0)$, $(0, -1, 0)$, $(1, 0, 0)$, and $(0, 0, 3/2)$. In other words, f is a tent-like function with a peak in the interior. In this case, one cannot construct a competitive element in S_k , $k = 4$, following the same path as in the proof of Theorem 1.1. In particular, perturbing ν by scaling the z -direction according to f , since would result in an element in S_5 , instead of S_4 .

While there do exist elements in S_4 whose marginal distributions in the xy -plane coincide with ν , e.g. the polytopes with vertices $(-1, 0, 0)$, $(0, 1, 0)$, $(0, -1, 0)$, $(1, 0, h)$ for any $h \neq 0$, this example illustrates how a different approach from that used in Theorem 1.1 would be needed to deal with the more intricate structure of the narrow closure of the constraint set found in higher dimensions.

In spite of the difficulty of obtaining the existence of minimizers to (WAA) when $d \geq 3$, by introducing an arbitrarily small regularization term, we are able to obtain existence in all dimensions. We begin with a lemma, providing compactness of the constraint set for the regularized problem (WAA $_{\varepsilon}$).




Fig. 2 Visualization of Ω_n as $n \rightarrow +\infty$

Lemma 3.4 Consider $\mu_n \in \mathcal{P}_2(\mathbb{R}^d)$ with $\sup_n M_2(\mu_n) < +\infty$. Fix $\varepsilon > 0$, $k \geq d + 1$, and $m \geq 1$. Suppose $\Omega_n \in S_k$ satisfy

$$\sup_n W_2(\mu_n, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) < +\infty.$$

Then there exists $\Omega \in S_k$ so that, up to a subsequence, $1_{\Omega_n}/|\Omega_n| \rightarrow 1_{\Omega}/|\Omega|$ narrowly.

Proof When $m > 1$, $\mathcal{U}_m(1_{\Omega}/|\Omega|) \geq 0$, and when $m = 1$, a Carleman-type estimate [5, Lemma 4.1] shows that, for all $\delta > 0$, $\mathcal{U}_1(1_{\Omega}/|\Omega|) \geq -(2\pi/\delta)^{d/2} - \delta M_2(1_{\Omega}/|\Omega|)$. Thus, along the minimizing sequence, taking $\delta = 1/(4\varepsilon)$ and applying the triangle inequality gives

$$\begin{aligned} \sup_n M_2(1_{\Omega_n}/|\Omega_n|) &= \sup_n 2W_2^2(1_{\Omega_n}/|\Omega_n|, \delta_0) - M_2(1_{\Omega_n}/|\Omega_n|) \\ &\leq \sup_n 4 \left(W_2^2(\mu_n, \delta_0) + W_2^2(\mu_n, 1_{\Omega_n}/|\Omega_n|) \right) - M_2(1_{\Omega_n}/|\Omega_n|) \\ &\leq \sup_n 4M_2(\mu_n) + 4W_2^2(\mu_n, 1_{\Omega_n}/|\Omega_n|) + 4\varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) \\ &\quad + 4\varepsilon(8\varepsilon\pi)^{d/2} < +\infty. \end{aligned}$$

Thus, Lemma 2.6 ensures that, up to a subsequence, there exists $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and a closed, convex k -gon Ω so that $1_{\Omega_n}/|\Omega_n| \rightarrow \nu$ narrowly, $1_{\Omega_n}(x) \rightarrow 1_{\Omega}(x)$ pointwise almost everywhere, and $\text{supp } \nu = \Omega$. Furthermore, there exists $R > 0$ so that $\Omega_n \subseteq B_R(0)$ for all $n \in \mathbb{N}$. Since $\varepsilon > 0$, $\mathcal{U}^m(1_{\Omega_n}/|\Omega_n|)$ is bounded above uniformly along a minimizing sequence, so by Eq. (1.5), $|\Omega_n|$ is bounded below. Hence, by the dominated convergence theorem,

$$|\Omega| = \int_{\mathbb{R}^d} 1_{\Omega} = \lim_{n \rightarrow +\infty} \int_{\mathbb{R}^d} 1_{\Omega_n} = \lim_{n \rightarrow +\infty} |\Omega_n| > 0. \quad (3.8)$$

Therefore $\dim(\text{aff}(\Omega)) = d$, so by Proposition 3.1, we have $\nu = 1_{\Omega}/|\Omega|$. $\square$

Now, we turn to the proof of Theorem 1.2, which ensures the existence of solutions of the regularized problem $(\text{WAA}_{\varepsilon})$, for all $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $d \geq 1$.

Proof of Theorem 1.2 Our proof begins similarly to the proof of Theorem 1.1. By Eq. (2.4) and the fact that $\mathcal{U}_m$ is invariant under translations, we may assume without loss of generality that μ has mean zero and restrict the minimization problem to the set $\{\Omega \in S_k: \int_{\Omega} x \, d\mathcal{L}^d(x) = 0\}$. Consider a minimizing sequence $\Omega_n \in S_k$ with $\int_{\Omega_n} x \, d\mathcal{L}^d(x) = 0$ so that

$$\lim_{n \rightarrow +\infty} W_2(\mu, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) = \inf_{\Omega \in S_k} W_2(\mu, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|). \quad (3.9)$$

By Lemma 3.4, with $\mu_n \equiv \mu$, there exists $\Omega \in S_k$ so that, up to a subsequence $1_{\Omega_n}/|\Omega_n| \rightarrow 1_{\Omega}/|\Omega|$ narrowly. By lower semicontinuity of the Wasserstein metric and $\mathcal{U}_m$ in the narrow topology, we have

$$\begin{aligned} \inf_{\Omega \in S_k} W_2(\mu, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|) &= \liminf_{n \rightarrow +\infty} W_2(\mu, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) \\ &\geq W_2(\mu, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|). \end{aligned} \quad (3.10)$$

Therefore, Ω is a minimizer of (WAA_ε). $\square$

4 Consistency

We now turn to the proof of consistency for the regularized Wasserstein archetypal analysis problem, as the measure μ is approximated by a sequence of empirical measures.

Proof of Theorem 1.4 First, note that by definition of Ω_n as the solution of (WAA_ε), for any $\tilde{\Omega} \in S_k$, we have

$$W_2(\mu_n, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) \leq W_2(\mu_n, 1_{\tilde{\Omega}}/|\tilde{\Omega}|) + \varepsilon \mathcal{U}_m(1_{\tilde{\Omega}}/|\tilde{\Omega}|). \quad (4.1)$$

We now turn to part (i). Since $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, the law of large numbers ensures $M_2(\mu_n) \rightarrow M_2(\mu)$ and $\mu_n \rightarrow \mu$ narrowly, almost surely. Thus, Eq. (2.3) ensures $W_2(\mu_n, \mu) \rightarrow 0$ almost surely. Thus, Lemma 3.4 ensures that, almost surely, there is a subsequence such that $1_{\Omega_n}/|\Omega_n| \rightarrow 1_{\Omega}/|\Omega|$ narrowly, where $\Omega \in S_k$.

To see Ω is optimal, note that, by inequality (4.1) and lower semicontinuity of the Wasserstein metric and $\mathcal{U}_m$ with respect to narrow convergence, almost surely, for any $\tilde{\Omega} \in S_k$,

$$\begin{aligned} W_2(\mu, 1_{\tilde{\Omega}}/|\tilde{\Omega}|) + \varepsilon \mathcal{U}_m(1_{\tilde{\Omega}}/|\tilde{\Omega}|) &= \liminf_{n \rightarrow +\infty} W_2(\mu_n, 1_{\tilde{\Omega}}/|\tilde{\Omega}|) + \varepsilon \mathcal{U}_m(1_{\tilde{\Omega}}/|\tilde{\Omega}|) \\ &\geq \liminf_{n \rightarrow +\infty} W_2(\mu_n, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) \\ &\geq W_2(\mu, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|), \end{aligned}$$

so Ω solves (WAA_ε).

Finally, to get the convergence rate in part (ii), note that for all $n \in \mathbb{N}$,

$$\begin{aligned} & W_2(\mu, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) \\ & \leq W_2(\mu_n, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) + W_2(\mu_n, \mu) \\ & \leq W_2(\mu_n, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|) + W_2(\mu_n, \mu) \\ & \leq W_2(\mu, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|) + 2W_2(\mu_n, \mu). \end{aligned}$$

Taking expectations on both sides and rearranging,

$$\begin{aligned} & \mathbb{E}[W_2(\mu, 1_{\Omega_n}/|\Omega_n|) + \varepsilon \mathcal{U}_m(1_{\Omega_n}/|\Omega_n|) - (W_2(\mu, 1_{\Omega}/|\Omega|) + \varepsilon \mathcal{U}_m(1_{\Omega}/|\Omega|))] \\ & \leq 2\mathbb{E}[W_2(\mu_n, \mu)] \leq 2(\mathbb{E}[W_2^2(\mu_n, \mu)])^{1/2}, \end{aligned}$$

where the last term can be further bounded using [15, Theorem 1], completing the proof. $\square$

5 Computational Approach in Two Dimensions

Here, we develop a computational approach for solving (WAA) and (WAA $_{\varepsilon}$) in dimension $d = 2$, when μ is a sum of Dirac masses, based on the semidiscrete approach introduced by Mérigot [20]; see also [23, Section 6.4.2]. This approach is based on the dual formulation of the 2-Wasserstein metric, which we now recall; see also [2, Theorem 6.1.1, Theorem 6.1.4].

5.1 Semidiscrete Formulation of (WAA) and (WAA $_{\varepsilon}$)

Suppose that $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$, where X, Y are compact subsets of $\mathbb{R}^d$.¹ Then

$$W_2(\mu, \nu) = \max_{\varphi \in C(X)} \Phi_{\nu}(\varphi)$$

for

$$\Phi_{\nu}(\varphi) := \int_X \varphi(x) d\mu(x) + \int_Y \varphi^*(y) d\nu(y), \quad (5.1)$$

$$\varphi^*(y) := \inf_{x \in X} |x - y|^2 - \varphi(x). \quad (5.2)$$

In the special case that μ is a sum of Dirac masses, $\mu = \sum_{i=1}^n m_i \delta_{x_i}$, we may suppose $X = \{x_i\}_{i=1}^n$ and identify $C(X) \simeq \mathbb{R}^n$, abbreviate $\varphi_i := \varphi(x_i)$, for any $\varphi \in C(X) \simeq \mathbb{R}^n$. Given any such $\varphi \in \mathbb{R}^n$, we may define the weighted Voronoi

¹ The compactness assumptions on X and Y can be removed, but then the maximum in the dual problem becomes a supremum. Also, X, Y compact is sufficient for our purposes, since our discrete measure μ and uniform distribution on a polygon ν will always be contained in compact sets.

tessellation corresponding to the points $\{x_i\}_{i=1}^n$ by

$$\text{Vor}_\mu^\varphi(x_i) = \{y \in \mathbb{R}^d : |x_i - y|^2 - \varphi_i \leq |x_j - y|^2 - \varphi_j \text{ for all } j = 1, \dots, n\}.$$

Using this partition of the domain $\mathbb{R}^d$, we may rewrite Φ_v as

$$\begin{aligned} \Phi_v(\varphi) &= \sum_{i=1}^n \left(\varphi_i m_i + \int_{\text{Vor}_\mu^\varphi(x_i)} \varphi^*(y) \, dv(y) \right), \\ &= \sum_{i=1}^n \left(\varphi_i m_i + \int_{\text{Vor}_\mu^\varphi(x_i)} |x_i - y|^2 - \varphi_i \, dv(y) \right). \end{aligned}$$

In this way, we may equivalently reformulate (WAA) and (WAA $_\varepsilon$) by writing, for $\varepsilon \geq 0$,

$$\begin{aligned} &\min_{\Omega \in S_k} W_2^2(\mu, 1_\Omega/|\Omega|) + \varepsilon \mathcal{U}_m(1_\Omega/|\Omega|) \\ &= \min_{\Omega \in S_k} \max_{\varphi \in \mathbb{R}^n} \Phi_{1_\Omega/|\Omega|}(\varphi) + \varepsilon \frac{1}{(m-1)|\Omega|^{m-1}} \\ &= \min_{\Omega \in S_k} \max_{\varphi \in \mathbb{R}^n} \sum_{i=1}^n \left(\varphi_i m_i + \frac{1}{|\Omega|} \int_{\text{Vor}_\mu^\varphi(x_i) \cap \Omega} (|y - x_i|^2 - \varphi_i) dy + \varepsilon \frac{1}{(m-1)|\Omega|^{m-1}} \right), \end{aligned} \quad (5.3)$$

where in the second equality, we apply Eq. (1.5) for $\mathcal{U}_m(1_\Omega/|\Omega|)$, with the convention that $|\Omega|^0/0 = -\log(|\Omega|)$.

5.2 Optimization Approach

We now discuss our approach for finding a polygon $\Omega \in S_k$ and dual vector $\varphi \in \mathbb{R}^n$ that solves Eq. (5.3), via an alternating gradient descent/ascent method in Ω and φ .

First, we consider the gradient ascent step in φ . Since φ only appears in the first two terms of our objective functional, it suffices to compute the variation of $\Phi_{1_\Omega/|\Omega|}(\varphi)$ with respect to φ , as in Mérigot's original work [20]. For the reader's convenience, we recall the form of this variation, following the presentation of Santambrogio [23, Sect. 6.4.2], with the derivation deferred to Appendix 8.2.

Proposition 5.1 *Suppose $\mu = \sum_{i=1}^n \delta_{x_i} m_i$ and $v \ll \mathcal{L}^d$. Then, for all $l = 1, \dots, n$,*

$$\frac{\partial \Phi_v}{\partial \varphi_l} = m_l - v(\text{Vor}_\mu^\varphi(x_l)). \quad (5.4)$$

Now, we turn to the gradient descent step in Ω , which we will perform by perturbing the vertices of Ω . In particular, in the two-dimensional case, we may assume that Ω is a convex k -gon with vertices $\{a_\ell\}_{\ell=1}^k$, ordered in the clockwise direction. We denote

the edge connecting vertices a_ℓ and $a_{\ell+1}$ by $E_\ell = [a_\ell, a_{\ell+1}]$, where we use circular indexing, and we denote the outward normal vector to E_ℓ by n_ℓ . The following theorem will allow us to compute the gradient of the objective function in (5.3) with respect to the vertex a_ℓ , $\ell = 1, \dots, k$.

Proposition 5.2 *Let $g \in L^1(\mathbb{R}^2) \cap W^{1,1}(\mathbb{R}^2)$ and consider the function $J(\Omega) = \int_\Omega g(x) dx$, where $\Omega = \text{conv}(\{a_\ell\}_{\ell=1}^k)$ is a convex k -gon. Using the above notation, the gradient of J with respect to vertex a_ℓ , $\ell = 1, \dots, k$, is given by*

$$\nabla_{a_\ell} J = L_\ell n_{\ell-1} + R_\ell n_\ell, \quad (5.5)$$

where

$$L_\ell = \int_{E_{\ell-1}} g(x) F_{\ell-1,\ell}(x) dx \quad \text{and} \quad R_\ell = \int_{E_\ell} g(x) F_{\ell+1,\ell}(x) dx$$

are constants computed on the intervals to the left and right of vertex a_ℓ and $F_{j,k}: \mathbb{R}^2 \rightarrow \mathbb{R}$ is the affine function defined by

$$F_{j,k}(x) = \frac{\langle x - a_j, a_k - a_j \rangle}{\|a_k - a_j\|^2}. \quad (5.6)$$

The proof of Proposition 5.2 is deferred to Appendix 8.3. We may now compute the gradient descent step in (5.3) by rewriting the second two terms in the objective function as

$$\begin{aligned} & \sum_{i=1}^n \frac{1}{|\Omega|} \int_{\text{Vor}_\mu^\varphi(x_i) \cap \Omega} (|y - x_i|^2 - \varphi_i) dy + \varepsilon \frac{1}{(m-1)|\Omega|^{m-1}} \\ &= \frac{1}{|\Omega|} \int_\Omega f_\varphi(y) dy + \varepsilon \frac{1}{(m-1)|\Omega|^{m-1}}, \end{aligned} \quad (5.7)$$

for

$$f_\varphi(y) = \sum_{i=1}^n 1_{\text{Vor}_\mu^\varphi(x_i)}(y) (|y - x_i|^2 - \varphi_i),$$

and taking the gradient of (5.7) with respect to each vertex by applying Proposition 5.2, the formula for $|\Omega|$ in terms of its vertices, and the quotient rule.

In summary, our method aims to optimize the saddle point problem (5.3) by alternating gradient ascent steps according to (5.4) and gradient descent steps of (5.7), for each vertex of Ω ; see Algorithm 1. We emphasize that the output of our algorithm is the extremal points $\{a_\ell\}_{\ell=1}^k$. The convex hull of these points determines the archetypal polygon. While our algorithm performs well in practice (see Sect. 6) we leave convergence analysis to future work. In particular, as illustrated in Fig. 9, we generically expect the energy landscape to have multiple stationary points so that further hypotheses would be needed to ensure our method converges to a global optimizer.

Algorithm 1: An alternating method to approximately solve (WAA) and (WAA_ε).

Input: A distribution $\mu = \sum_{i=1}^n \delta_{x_i} m_i$, an initial dual variable φ^0 , a number of vertices $k \geq 3$, an initial polygon $\Omega^0 = \text{conv}(\{a_\ell^0\}_{\ell=1}^k)$, regularization parameters $\varepsilon \geq 0$ and $m \geq 1$, step sizes $\tau_1, \tau_2 > 0$, and tolerances $\delta_1, \delta_2 > 0$.

Output: A polygon Ω^n that is a candidate approximate solution for (WAA) or (WAA_ε).

```

1 for  $s \in \mathbb{N}$  do
2   1. Gradient ascent in  $\varphi$ . Set  $\bar{\varphi}^{s,1} = \varphi^s$ 
3   for  $r \in \mathbb{N}$  do
4     Find the power diagram at  $\bar{\varphi}^{s,r}$ , compute  $|\text{Vor}_\mu^\varphi(x_i) \cap \Omega|/|\Omega|$  and compute
        
$$\bar{\varphi}_i^{s,r+1} = \bar{\varphi}_i^{s,r} + \tau_1(m_i - |\text{Vor}_\mu^\varphi(x_i) \cap \Omega|/|\Omega|)$$

5     If  $\|\bar{\varphi}^{s,r+1} - \bar{\varphi}^{s,r}\|_1 \leq \delta_1$ , STOP
6   end for
7   Set  $\varphi^{s+1} = \bar{\varphi}^{s,r}$ 
8   2. Gradient descent in  $\Omega$ . Update the vertices  $\{a_\ell^s\}_{\ell=1}^k$  of  $\Omega$  by
        
$$a_\ell^{s+1} = a_\ell^s - \tau_2 \nabla_{a_\ell} \left( \frac{1}{|\Omega|} \int_\Omega f_{\varphi^{s+1}}(y) \, dy + \varepsilon \frac{1}{(m-1)|\Omega|^{m-1}} \right).$$

9   If  $\|\alpha^{s+1} - \alpha^s\|_\infty < \delta_2$ , STOP
9 end for

```

6 Numerical Experiments

In this section, we illustrate the performance of the algorithm via several computational experiments. In all numerical experiments, we set $\tau_1 = 0.5$, $\tau_2 = 0.1$, $\delta_1 = 10^{-3}$, and $\delta_2 = 10^{-5}$. In our numerical studies, we observe that the algorithm is not very sensitive to the choice of step sizes τ_1 and τ_2 . We choose the initial dual variable φ^0 to be a zero vector with dimension n in all experiments.

6.1 Example 1: Uniform Distribution in a Disk

We first consider the case where the measure μ is the uniform probability measure on the unit disk. The measure μ is approximated as follows. We first generate 10,000 data points, where each data point $x_i = (x_i^1, x_i^2)$ is randomly generated by

$$\begin{cases} x^1 = \sqrt{r} \cos(2\pi\theta), \\ x^2 = \sqrt{r} \sin(2\pi\theta), \end{cases}$$

with r and θ drawn from uniform distributions on $[0, 1]$. We generate a uniform 15×15 discretization of the square $[x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$ where $x_{\min}/\max$, $y_{\min}/\max$ are the smallest or largest value in x , y , respectively. We then set the center of the i -th cell as x_i and the ratio of random points located in this cell as m_i to have the approximate





Fig. 3 Snapshots of the evolution of the triangle computed by Algorithm 1 until convergence. See Subsect. 6.1

Fig. 4 Figure 3 of [22]: A minimizer for 30,033 random data points. The red curve is the boundary of the convex hull of the archetype points and the magenta curve is the boundary of the convex hull of random data points. See Subsect. 6.1



measure $\mu_n = \sum_{i=1}^n m_i \delta_{x_i}$ by dropping those cells where $m_i = 0$. In this experiment, $n = 172$. We seek to solve (WAA), with no regularization ($\varepsilon = 0$), and $k = 3$.

Figure 3 illustrates the evolution of the triangle computed using Algorithm 1. We observe that the triangle evolves gradually to a regular triangle centered at the origin. Note that all three vertices lie outside of the disk. This is in contrast to the classical archetypal analysis problem (1.1), where the solution is an inscribed triangle, as depicted in Fig. 4. For more detail see [22, Proposition 3.1].

In Fig. 5, we display the results of (WAA) for $k = 3, 4, \dots, 10$. In all experiments, the optimal solutions appear to be regular polygons. It is known that the solution for the classical archetypal analysis is a regular k -gon [22, Proposition 3.1] and we conjecture that this holds true for the solution to (WAA) as well. As k increases, the k -gon becomes closer to a disk. In Fig. 5, we also plot the change of the radius of the polygons as k increases, and the asymptotic behavior is consistent with the fact that the polygon converges to a disk as $k \rightarrow \infty$.




Fig. 5 Solutions obtained from Algorithm 1 for different numbers of vertices, $k = 3, 4, \dots, 10$, where μ is the uniform probability measure on the unit disk. **Bottom:** The change of radius of the optimal k -gon as the number of vertices increases. See Subject. 6.1

6.2 Example 2: Normal Distribution

In this experiment, we study the behavior of the solution of (WAA) when μ is a normal distribution $\mathcal{N}(0, I)$, where I is the identity matrix. The approximate measure $\mu_n = \sum_{i=1}^n m_i \delta_{x_i}$ is generated in a similar way as those described in Sect. 6.1, except that the empirical measure is directly generated by a 2-dimensional normal distribution with $n = 92$. In Fig. 6, we plot the solution for the cases when $k = 3, 4$. The initial choice of points is given by k randomly generated vertices on the circle. We observe that the solution is in the interior of the convex hull of data points randomly generated from the normal distribution. For the case $k = 4$, we observe that the solution is a square with a side length ≈ 3.3855 , which approximately coincides with the length of the optimal interval (≈ 3.3862) for the analogous 1-dimensional problem; see Eq. (1.4).

6.3 Example 3: Sensitivity to ε in (WAA $_{\varepsilon}$)

In this experiment, we consider the solution of the problem (WAA $_{\varepsilon}$) with

$$\mu = \mathcal{N}(0, \Sigma), \quad \Sigma = \begin{bmatrix} 5 & 0 \\ 0 & 1 \end{bmatrix}$$




Fig. 6 Solutions for $k = 3$ and $k = 4$, when μ is a normal distribution $\mathcal{N}(0, I)$





Fig. 7 Solutions for $k = 4$ with target distribution $\mathcal{N}\left(0, \begin{bmatrix} 5 & 0 \\ 0 & 1 \end{bmatrix}\right)$ and different choices of $\varepsilon \in \{0, 5, 10, 15, 20, 30, 40, 50\}$

Fig. 8 Change of ratio between the longest side and shortest side of the obtained quadrilateral with ε



approximating μ as in Example 2 with $n = 90$. As before, we compute the approximate optimizer via Algorithm 1. In Fig. 7, we study how the solution changes as the value of ε increases from 0 to 50. In all cases, we set $m = 2$ and initialize the points at the $\varepsilon = 0$ archetypes. We observe that as ε increases, the quadrilateral becomes larger and closer to a square, that is, the ratio between the long side and the short side decreases, as shown in Fig. 8. This is consistent with the motivation on introducing the regularization term, as it rewards larger areas.

6.4 Example 4: Non-convexity of the Energy Landscape

We now consider an example that provides numerical evidence suggesting (WAA) is a non-convex optimization problem. Let $\mu = 1_T/|T| \in \mathcal{P}(\mathbb{R}^2)$, where $T \subseteq \mathbb{R}^2$ is the triangle with base length two, height one, and center of mass zero, where the base and height are chosen to be parallel with the coordinate axes: see black triangles on the left-hand side of Fig. 9. In this case, (WAA), with $k = 3$, has an obvious unique global minimizer: $\Omega = T$. We seek to compute the values of $W_2(\mu, 1_\Omega/|\Omega|)$ for many different triangles Ω to investigate the nonconvexity of the energy landscape.

In order to compute the 2-Wasserstein distance, we begin by approximating μ and $1_\Omega/|\Omega|$ by $n = 3792$ Dirac masses, arranged on a uniform grid. We then apply the `emd` function from the Python Optimal Transport library [14] to compute the 2-Wasserstein distance between the approximations. Our approximation by Dirac masses introduces numerical error on the order of ~ 0.01 in the distance computation.

On the left-hand side of Fig. 9, we compute $W_2(\mu, 1_\Omega/|\Omega|)$ as the triangle Ω varies according to two parameters: p_1 controls the height (different values of p_1 are shown in each row) and p_2 controls the width of the base (different values of p_2 are shown in each column). The triangles Ω are constrained to always have the center of mass 0 since the optimal choice of Ω will have the center of mass coinciding with μ . Each cell in the grid shows the varying triangle Ω in blue, the target triangle T in black, and the value of the 2-Wasserstein distance between them. We see that the minimum distance is obtained when $\Omega = T$, shown in the second row, third column. Due to the error of our numerical approximation of the 2-Wasserstein distance, the distance between them is shown as 0.01.

The right-hand side of Fig. 9 depicts the value of $W_2(\mu, 1_\Omega/|\Omega|)$, as the triangles Ω vary in the same manner, visualized as a contour plot. We observe the global minima when $\Omega = T$ at the top of the figure. The bottom of the figure shows a potential local minimum, though due to the accuracy limits of our numerical computation, it could also be a flat area of the energy landscape. Either way, it is clear from the contour plot that the energy landscape is nonconvex. For this reason, it is an important direction for future work to understand under what conditions our gradient-based minimization algorithm is guaranteed to converge to the global minimum, as well as to develop non-convex minimization methods for the general setting.

6.5 Example 5: Early Spread of the Covid-19 Virus in the U.S

In this section, we apply (WAA) to explore the early-stage evolution of the Covid-19 virus in the U.S. The dataset used for analysis is freely available at <https://covidtracking.com/data/api>. In total, there are 51 data points (50 states + D.C.), each corresponding to a time series of the average positivity rates (PRs), where

$$\text{PR} = \frac{\text{total \# of positive cases by the day}}{\text{total \# of tests by the day}},$$




Fig. 9 We numerically investigate the energy landscape of (WAA), when $k = 3$ and $\mu = 1_T/|T|$, when $T \subseteq \mathbb{R}^2$ is given by the black triangle shown in each cell. On the left, the figure shows how $W_2(\mu, 1_{\Omega}/|\Omega|)$ changes as the form of the triangles Ω is varied according to parameters p_1 (height) and p_2 (width). The blue triangle in each cell shows the shape of Ω for a given choice of (p_1, p_2) , and the number in each cell shows the value of $W_2(\mu, 1_{\Omega}/|\Omega|)$. On the right, the figure shows contour plots for the value of $W_2(\mu, 1_{\Omega}/|\Omega|)$ as Ω varies, depicting a nonconvex energy landscape

and the average is computed using the centered 5-day moving average scheme. The time range is chosen between April 20 and September 20, 2020, the relatively early stage of the pandemic. Visualization of the dataset can be found in Fig. 10.

In this example, the dimension of the time series data is 154, which is much higher than the number of data points. For convenience, we first use PCA to reduce the dimension of the data before applying WAA, with the explained variances by the first five principal components (PCs) plotted in Fig. 10. In this example, the first two PCs combined account for about 99% variation of the dataset. As a result, we use the first two PCs to obtain a reduced representation for the dataset. A quantitative analysis of such an approximation procedure in the context of AA can be found in [17].

In Fig. 11, we apply WAA to the dataset, choosing $k = 3$, using the empirical elbow rule. The initial points for Algorithm 1 are given by three arbitrarily chosen points, the convex hull of which contains the entire dataset. We also compare WAA to both classical AA and k -means, applied to the same dataset. It can be seen that both WAA and AA yield similar results, with the former demonstrating a more robust performance to the outliers. In both cases, the archetypes correspond to exemplars of the different evolutionary patterns. The upper right archetype is close to the northeastern states like New York (NY) and New Jersey (NJ), which is where the first outbreak in the U.S. took place. The bottom archetype is close to many southern states, which corresponds to the second outbreak. The upper left archetype is surrounded by states that have a low population density and experienced a relatively slow positive rate curve at the




Fig. 10 Visualization of average PRs of the 50 states plus D.C. from April 20 to September 20, 2022 (left) and the variances explained by the first 5 PCs (right)




Fig. 11 Three archetypes (red triangles) given by WAA (left) and the original AA (right) applied to the reduced dataset. Green squares correspond to the k -means centers

beginning stage of the pandemic. In contrast, the k -means centers are more difficult to interpret.

7 Discussion

In this paper, we considered a Wasserstein metric-based Archetypal Analysis, where a probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ is approximated by a distribution that is uniformly supported on a polytope with a fixed number of vertices. We established that, in one dimension, there is a unique minimizer of (WAA), and in two dimensions, we showed that a minimizer exists, under the additional assumption that μ is absolutely continuous with respect to Lebesgue measure (Theorem 1.1). By adding an appropriate regularization in terms a Rényi entropy, we were able to prove that a solution of the regularized problem (WAA $_{\varepsilon}$) exists for all $\varepsilon > 0$, $d \geq 1$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ (Theorem 1.2). Finally, we proved a consistency and convergence result for the (WAA $_{\varepsilon}$)

problem (Theorem 1.4). In Sect. 5, we introduced a computational approach for the (WAA) and (WAA_ε) problems using the semi-discrete formulation of the Wasserstein metric; see Algorithm 1. We concluded by implementing the algorithm and conducting several numerical experiments in Sect. 6 to support our analytical findings.

There are many interesting future directions for this work.

- Extend the analysis for $\varepsilon = 0$ to higher dimensions (see remark 3.3) and μ with lower dimensional support (see remark 3.2).
- Consider formulations of the archetypal analysis problem for more general optimal transport metrics, e.g. p -Wasserstein, $p \neq 2$.
- Analyze the uniqueness of solutions up to global invariances.
- For μ the uniform distribution on the unit disk (Sect. 6.1) and a standard normal distribution (Sect. 6.2), establish that the optimal solution is given by a regular k -gon.
- Develop computational methods for higher dimensional problems. Here one could use the back and forth method [19] or the entropic approximations of the 2-Wasserstein distance [10].
- Further study the differences between the original archetypal analysis problem and WAA.
- Analysis of a variant of WAA in which the archetypes are projected onto the convex hull of the support of the measure, either during the iterations of Algorithm 1 or as a post-processing step.
- Analysis of sufficient conditions on the data distribution and the initialization of our numerical method to ensure convergence to an optimizer.

8 Appendices

8.1 Proof of Lemma 2.6

First, we prove part (i). By Hölder's inequality and the fact that $\mu_n := 1_{\Omega_n}/|\Omega_n|$ are probability measures,

$$\int x \, d\mu_n(x) \leq \left(\int |x|^2 \, d\mu_n(x) \right)^{1/2} = M_2(\mu_n)^{1/2},$$

so that the means of μ_n are uniformly bounded. Assume, for the sake of contradiction, that item (i) fails. Then there must exist a line segment $L_n \subseteq \Omega_n$ so that, up to a subsequence, $|L_n| = \mathcal{L}^1(L_n) \rightarrow +\infty$. Then, using the fact from [7, Lemma 3.2] that there exists $\gamma = \gamma(d) > 0$ so that $M_2(\pi_{L_n} \# \mu_n) \geq \gamma |L_n|^2$, we obtain

$$\int |x|^2 \, d\mu_n(x) \geq \int |\pi_L(x)|^2 \, d\mu_n(x) = M_2(\pi_L \# \mu_n) \geq \gamma |L_n|^2 \rightarrow +\infty,$$

which is a contradiction.

Now, since (i) holds, by Prokhorov's theorem, the sequence is tight. Thus, up to a subsequence, the sequence converges in the narrow topology to some $\nu \in \mathcal{P}(\mathbb{R}^d)$. By

lower-semicontinuity of M_2 with respect to narrow convergence, we have $M_2(\nu) < +\infty$, so $\nu \in \mathcal{P}_2(\mathbb{R}^d)$. By Heine-Borel, up to a subsequence, each of the vertices $\{v_n^i\}_{i=1}^k$ of Ω_n converges to some $\{v^i\}_{i=1}^k$. Let Ω be the closed, convex hull of these limit points.

In order to prove convergence of the indicator functions $\mathcal{L}^d$ -a.e., we first prove the following:

Claim *Let $\{\Omega_n\}_{n=1}^\infty$ be a sequence of convex k -gons with nonempty interior so that their vertices converge, and let Ω be the convex hull of the limit points (which may have empty interior). Then if x belongs to the interior of Ω , there exists $\delta > 0$ so that $B_\delta(x) \subseteq \Omega_n$ for n sufficiently large.*

Proof of Claim Let $\{w_n^j\}_{j=1}^l$ denote the vertices of Ω . By assumption, for all $j = 1, \dots, l$, there exists $w_n^j \in \Omega_n$ so that $\lim_{n \rightarrow +\infty} w_n^j = w^j$. Thus, for all $\epsilon > 0$, there exists N so that $n > N$ ensures $|w^j - w_n^j| < \epsilon$ for all $j = 1, \dots, l$. Since Ω_n is convex, this implies there exists $\delta > 0$ so that $B_\delta(x) \subseteq \Omega_n$ for all $n > N$.

We now apply this claim to prove convergence of the indicator functions $\mathcal{L}^d$ -a.e. In particular, if $x \in \Omega^\circ$, then the previous claim shows $x \in \Omega_n^\circ$ for n sufficiently large. Likewise, if $x \in \Omega^c$, then convergence of the vertices ensures that $x \in \Omega_n^c$ for n sufficiently large. Combining these gives

$$1_\Omega(x) = 1 = \lim_{n \rightarrow +\infty} 1_{\Omega_n}(x) \quad \text{if } x \in \Omega^\circ, \quad (8.1)$$

$$1_\Omega(x) = 0 = \lim_{n \rightarrow +\infty} 1_{\Omega_n}(x) \quad \text{if } x \in \Omega^c. \quad (8.2)$$

Since $\mathcal{L}^d(\partial\Omega) = 0$, this shows that $1_{\Omega_n} \rightarrow 1_\Omega$ $\mathcal{L}^d$ -a.e.

We now show $\text{supp } \nu = \Omega$. First, we show $\text{supp } \nu \subseteq \Omega$. Suppose $x \in \Omega^c$. Since Ω^c is open, there exists an open ball so that $x \in B \subset \subset \Omega^c$. Convergence of the vertices ensures $B \subseteq \Omega_n^c$ for n sufficiently large. Thus, by the fact that $\mu_n \rightarrow \nu$ narrowly,

$$0 = \liminf_{n \rightarrow +\infty} \mu_n(B) \geq \nu(B).$$

This shows $x \notin \text{supp } \nu$.

Now, we show $\Omega \subseteq \text{supp } \nu$. Let π be a projection onto the affine hull of Ω , $\text{aff}(\Omega)$. If $\dim(\text{aff}(\Omega)) = 0$, then Ω is a single point, and the previous implication that $\text{supp } \nu \subseteq \Omega$, combined with the fact that $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ must have nontrivial support, ensures $\text{supp } \nu = \Omega$. Now, assume $\dim(\text{aff}(\Omega)) > 0$. Fix $x \in \Omega$ and let $C \subseteq \mathbb{R}^d$ be a closed ball containing x . It suffices to show that $\nu(C) > 0$. Note that the convex k -gons $\pi(\Omega_n)$ have nonempty interior with respect to the topology on $\pi(\mathbb{R}^d)$ and $\pi(\Omega)$ is the convex hull of the limits of their vertices. Furthermore, since $\pi(\Omega)$ is convex, there exists $y \in \pi(\Omega)^\circ$ (interior taken with respect to the topology on $\pi(\mathbb{R}^d)$) and $\epsilon > 0$ so that $B_\epsilon(y) \subseteq \pi(C) \cap \pi(\Omega)$, where $B_\epsilon(y)$ is an open ball in the affine subspace $\pi(\mathbb{R}^d)$.

By the above claim, there exists $\delta > 0$ so that $B_\delta(y) \subseteq \pi(\Omega_n) \cap \pi(C)$ for n sufficiently large. Then, by [4, Lemma 5], there exists $k > 0$ so that

$$k \leq \frac{|\Omega_n \cap \pi^{-1}(B_\delta(y))|}{|\Omega_n|} = \mu_n(\pi^{-1}(B_\delta(y))) = \pi\#\mu_n(B_\delta(y)) \leq \pi\#\mu_n(\overline{B_\delta(y)}). \quad (8.3)$$

Since $\overline{B_\delta(y)}$ is closed and the continuous mapping theorem ensures $\pi\#\mu_n \rightarrow \pi\#v$ narrowly, taking the limsup of (8.3) as $n \rightarrow +\infty$ gives

$$k \leq \pi\#v(\overline{B_\delta(y)}) = v(\pi^{-1}(\overline{B_\delta(y)})) \leq v(C),$$

where the last inequality follows since $\overline{B_\delta(y)} \subseteq \pi(C)$, so $\pi^{-1}(\overline{B_\delta(y)}) \cap \text{supp } v \subseteq C$. This gives the result. $\square$

8.2 Proof of Proposition 5.1

By definition of Φ_v and φ^* in Eqs. (5.1)–(5.2), we may rewrite this as

$$\frac{d}{d\varphi_l} \Phi_v(\varphi) = \frac{d}{d\varphi_l} \sum_{i=1}^n \varphi_i m_i + \int \left(\inf_j |x_j - y|^2 - \varphi_j \right) dv(y) \quad (8.4)$$

$$= m_l + \frac{d}{d\varphi_l} \int \left(\inf_j |x_j - y|^2 - \varphi_j \right) dv(y). \quad (8.5)$$

To compute the remaining derivative, we follow Santambrogio [23, p243]. We proceed by decomposing the domain of integration into the following three regions:

$$S_1 := \{y : |x_l - y|^2 - \varphi_l < |x_j - y|^2 - \varphi_j \text{ for all } j \neq l\}$$

$$S_2 := \{y : |x_l - y|^2 - \varphi_l > |x_j - y|^2 - \varphi_j \text{ for some } j \neq l\}$$

$$S_3 := \{y : |x_l - y|^2 - \varphi_l = |x_j - y|^2 - \varphi_j = \inf_{j'} |x_{j'} - y|^2 - \varphi_{j'} \text{ for some } j \neq l\}.$$

These regions may be interpreted as the region for which l is the unique index that attains the infimum in $\inf_j |x_j - y|^2 - \varphi_j$, the region for which l does not attain the infimum, and the region for which l attains the infimum, but not uniquely. Then,

$$\begin{aligned} \frac{d}{d\varphi_l} \int_{S_1} \left(\inf_j |x_j - y|^2 - \varphi_j \right) dv(y) &= \frac{d}{d\varphi_l} \int_{S_1} (|x_l - y|^2 - \varphi_l) dv(y) \\ &= -v(S_1) = -v(\text{Vor}_\mu^\varphi(x_l)) \\ \frac{d}{d\varphi_l} \int_{S_2} \left(\inf_j |x_j - y|^2 - \varphi_j \right) dv(y) &= \frac{d}{d\varphi_l} \int_{S_2} (|x_j - y|^2 - \varphi_j) dv(y) = 0 \\ \frac{d}{d\varphi_l} \int_{S_3} \left(\inf_j |x_j - y|^2 - \varphi_j \right) dv(y) &= \frac{d}{d\varphi_l} 0 = 0 \end{aligned}$$

where the last equality follows since S_3 has zero Lebesgue measure and $v \ll \mathcal{L}^d$. Combining these estimates with Eq. (8.4) above gives (5.4).

8.3 Proof of Proposition 5.2

If we perturb the ℓ -th vertex $a_\ell \mapsto a_\ell + \varepsilon u$ for $u \in \mathbb{R}^2$ and $\varepsilon \in \mathbb{R}$, this perturbs the two adjacent edges $E_{\ell-1}$ and E_ℓ . Let $F_{j,k}: \mathbb{R}^2 \rightarrow \mathbb{R}$ be the affine function defined in (5.6) which satisfies $F_{j,k}(a_j) = 0$ and $F_{j,k}(a_k) = 1$. Let $v \in W^{1,\infty}(\mathbb{R}^2)$ satisfy

$$v(x) = \begin{cases} F_{\ell-1,\ell}(x)u, & x \in E_{\ell-1} \\ F_{\ell+1,\ell}(x)u, & x \in E_\ell \\ 0 & x \in E_m, m \neq \ell-1, \ell. \end{cases}.$$

Denote the perturbed polygon by $\Omega_\varepsilon = (I + \varepsilon v)(\Omega)$. Noting that $\Omega = \Omega_0$ has Lipschitz boundary, we apply the Leibniz integral formula (see, e.g., [18, Thm.5.2.2]) to obtain

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} J(\Omega_\varepsilon) = \int_{\partial\Omega} g(x) v(x) \cdot n(x) \, dx,$$

where $n = n(x)$ denotes the outward normal vector at $x \in \partial\Omega$. Observing that

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} J(\Omega_\varepsilon) = \int_{\partial\Omega} g(x) v(x) \cdot n(x) \, dx = (L_\ell n_{\ell-1} + R_\ell n_\ell) \cdot u$$

for arbitrary $u \in \mathbb{R}^2$, we obtain the desired result.

Acknowledgements K. Craig would like to thank Jun Kitagawa and Dejan Slepčev for helpful discussions on semi-discrete optimal transport and the challenges of WAA in higher dimensions.

Funding K. Craig acknowledges partial support from NSF DMS 1811012, NSF DMS 2145900, and a Hellman Faculty Fellowship. K. Craig also acknowledges the support from the Simons Center for Theory of Computing, at which part of this work was completed. B. Osting acknowledges partial support from NSF DMS 17-52202. D. Wang is partially supported by National Natural Science Foundation of China (Grant No. 12101524), Guangdong Basic and Applied Basic Research Foundation (Grant No. 2023A1515012199), Shenzhen Science and Technology Innovation Program (Grant No. JCYJ20220530143803007, RCYX20221008092843046), Guangdong Provincial Key Laboratory of Mathematical Foundations for Artificial Intelligence (2023B1212010001), and Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone Project (No.HZQSW-S-KCCYB-2024016).

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

1. Ambrosio, L., Brué, E., Semola, D.: Lectures on optimal transport. Springer, Cham (2021)
2. Ambrosio, L., Gigli, N., Savaré, G.: Gradient flows in metric spaces and in the space of probability measures. In: Lectures in Mathematics ETH Zürich, 2nd edn. Birkhäuser Verlag, Basel (2008)
3. Artstein-Avidan, S., Giannopoulos, A., Milman, V.D.: Asymptotic geometric analysis, Part I, vol. 261. American Mathematical Society, Providence (2021)
4. Belili, N., Heinrich, H.: Approximation of distributions. Stat. Prob. Lett. **76**(3), 298–303 (2006)
5. Carrillo, J.A., Patacchini, F.S., Sternberg, P., Wolansky, G.: Convergence of a particle method for diffusive gradient flows in one dimension. SIAM J. Math. Anal. **48**(6), 3708–3741 (2016)

6. Chan, B.H.P., Mitchell, D.A., Cram, L.E.: Archetypal analysis of galaxy spectra. *Mon. Not. R. Astron. Soc.* **338**(3), 790–795 (2003)
7. Cuesta-Albertos, J., Matrán, C., Rodríguez-Rodríguez, J.: Approximation to probabilities through uniform laws on convex sets. *J. Theor. Probab.* **16**(2), 363–376 (2003)
8. Cuesta-Albertos, J.A., Bea, C.M., Rodríguez, J.M.R.: Shape of a distribution through the l2-wasserstein distance. In: *Distributions with Given Marginals and Statistical Modelling*, pp. 51–61. Springer, Cham (2002)
9. Cutler, A., Breiman, L.: Archetypal analysis. *Technometrics* **36**(4), 338 (1994)
10. Cuturi, M.: Sinkhorn distances: lightspeed computation of optimal transport. *Adv. Neural Inform. Process. Syst.* **26**, 2292 (2013)
11. Digne, J., Cohen-Steiner, D., Alliez, P., De Goes, F., Desbrun, M.: Feature-preserving surface reconstruction and simplification from defect-laden point sets. *J. Math. Imaging Vision* **48**, 369–382 (2014)
12. Eugster, M.J., Leisch, F.: Weighted and robust archetypal analysis. *Comput. Stat. Data Anal.* **55**(3), 1215–1225 (2011)
13. Figalli, A., Glaudo, F.: *An Invitation to Optimal Transport, Wasserstein Distances, and Gradient Flows*. EMS Textbooks in Mathematics, Berlin (2021)
14. Flamary, R., Courty, N., Gramfort, A., Alaya, M.Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N.T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D.J., Tavenard, R., Tong, A., Vayer, T.: Pot: python optimal transport. *J. Mach. Learn. Res.* **22**(78), 1–8 (2021)
15. Fournier, N., Guillin, A.: On the rate of convergence in Wasserstein distance of the empirical measure. *Probab. Theory Relat. Fields* **162**(3), 707–738 (2015)
16. Gigli, N.: On the inverse implication of Brenier-McCann theorems and the structure of $(P_2(M), W_2)$. *Methods Appl. Anal.* **18**(2), 127–158 (2011)
17. Han, R., Osting, B., Wang, D., Xu, Y.: Probabilistic methods for approximate archetypal analysis. *Inform. Inference J. IMA* **12**, 466 (2022)
18. Henrot, A., Pierre, M.: *Shape Variation and Optimization: A Geometrical Analysis*. EMS Press, Berlin (2018)
19. Jacobs, M., Léger, F.: A fast approach to optimal transport: the back-and-forth method. *Numer. Math.* **146**(3), 513–544 (2020)
20. Mérigot, Q.: A multiscale approach to optimal transport. *Comput. Graph. Forum* **30**(5), 1583–1592 (2011)
21. Mørup, M., Hansen, L.K.: Archetypal analysis for machine learning and data mining. *Neurocomputing* **80**, 54–63 (2012)
22. Osting, B., Wang, D., Xu, Y., Zosso, D.: Consistency of archetypal analysis. *SIAM J. Math. Data Sci.* **3**(1), 1–30 (2021)
23. Santambrogio, F.: *Optimal transport for applied mathematicians*. Birkhäuser (NY) **55**(58–63), 94 (2015)
24. Seth, S., Eugster, M.J.: Probabilistic archetypal analysis. *Mach. Learn.* **102**(1), 85–113 (2016)
25. Shaked, M., Shanthikumar, J.G.: *Stochastic orders*. Springer, Cham (2022)
26. Shoval, O., Sheftel, H., Shinar, G., Hart, Y., Ramote, O., Mayo, A., Dekel, E., Kavanagh, K., Alon, U.: Evolutionary trade-offs, pareto optimality, and the geometry of phenotype space. *Science* **336**(6085), 1157–1160 (2012)
27. Villani, C.: *Topics in optimal transportation*. In: *Graduate Studies in Mathematics*, vol. 58. American Mathematical Society, Providence (2003)
28. Villani, C.: *Optimal transport*, volume 338 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, (2009). Old and new

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.