Solving General Elliptical Mixture Models through an Approximate Wasserstein Manifold

Shengxi Li[1*], Zeyang Yu[1], Min Xiang[1], Danilo Mandic[1]
[1]Imperial College London
South Kensington Campus
London SW7 2AZ, UK
{shengxi.li17, z.yu17, m.xiang13, d.mandic}@imperial.ac.uk

Abstract

We address the estimation problem for general finite mixture models, with a particular focus on the elliptical mixture models (EMMs). Compared to the widely adopted Kullback-Leibler divergence, we show that the Wasserstein distance provides a more desirable optimisation space. We thus provide a stable solution to the EMMs that is both robust to initialisations and reaches a superior optimum by adaptively optimising along a manifold of an approximate Wasserstein distance. To this end, we first provide a unifying account of computable and identifiable EMMs, which serves as a basis to rigorously address the underpinning optimisation problem. Due to a probability constraint, solving this problem is extremely cumbersome and unstable, especially under the Wasserstein distance. To relieve this issue, we introduce an efficient optimisation method on a statistical manifold defined under an approximate Wasserstein distance, which allows for explicit metrics and computable operations, thus significantly stabilising and improving the EMM estimation. We further propose an adaptive method to accelerate the convergence. Experimental results demonstrate the excellent performance of the proposed EMM solver.

Introduction

This paper establishes a general solution to the finite mixture model problem, which has been attracting extensive research effort for decades, due to both its simple representation and potential for universal approximation on arbitrary distributions in $\mathbb{R}^M$. The finite mixture model also provides interpretable and statistical descriptions of data, which makes it a popular choice in a wide range of statistical learning paradigms, such as semi-supervised learning, capsule networks, and various image processing paradigms (e.g., denoising, matching and registration).

The estimation on a finite mixture model boils down to a minimisation problem which considers the mixture of distributions as a parametric model $\rho(\boldsymbol{\theta})$, which is then optimised through a minimisation of a certain discrepancy measure between $\rho(\boldsymbol{\theta})$ and empirical distributions of observed data ρ^*, namely, $\min_{\boldsymbol{\theta}} d\left(\rho(\boldsymbol{\theta}), \rho^*\right)$. This minimisation, although not

explicitly stated, is a constrained problem because $\rho(\boldsymbol{\theta})$ must maintain the property of a probability density throughout, to ensure that $d(\cdot, \cdot)$ is tractable.

Due to this probability constraint, various advanced numerical algorithms (solvers) have been typically restricted by either the requirement of an increasingly flexible $\rho(\boldsymbol{\theta})$ or a powerful $d(\cdot, \cdot)$. Such restrictions, for example, are one of the main rationales for using the expectation-maximisation (EM) algorithm to minimise the KullbackLeibler (KL) divergence in Gaussian mixture model (GMM) problems (Xu and Jordan 1996). On the other hand, gradient-based numerical algorithms typically rest upon additional techniques that only work in particular situations (e.g., gradient reduction (Redner and Walker 1984), positive definite projection (Hosseini and Sra 2015), re-parametrisation (Jordan and Jacobs 1994) and Cholesky decomposition (Naim and Gildea 2012)). Besides the GMM, there exist other solutions that allow for flexible choices of $\rho(\boldsymbol{\theta})$, which still belong to the EM-type methods (e.g., mixtures of t-distributions (Peel and McLachlan 2000), Laplace distributions (Tan and Jiao 2007) and hyperbolic distributions (Browne and McNicholas 2015)). Unfortunately, given other suitable candidates of distributions, those EM-type methods cannot ensure universal convergence (Kent and Tyler 1991; Zhang, Wiesel, and Greco 2013; Sra and Hosseini 2013), which dramatically limits the power of finite mixture models.

Another issue that has been highlighted in the literature is the sensitivity to initialisations when solving GMMs (Xu, Hsu, and Maleki 2016; Jin et al. 2016). One of the main reasons is due to the use of KL divergence, which operates based on a "bin-to-bin" comparison between two density histograms. This means that mixtures which fall into a spurious local minimum cannot be corrected via the points outside. Indeed, with random initialisations for the GMM, Jin has proved that the EM algorithm or any other first-order method which minimises the KL divergence are highly likely to result in arbitrary bad local minima (Jin et al. 2016). This can also be easily verified from the non-smooth optimization space with various local optimum of the KL divergence as illustrated in Fig. 1-(b), even for estimating a simple GMM (Fig. 1-(a)). The gradients on the space are highly

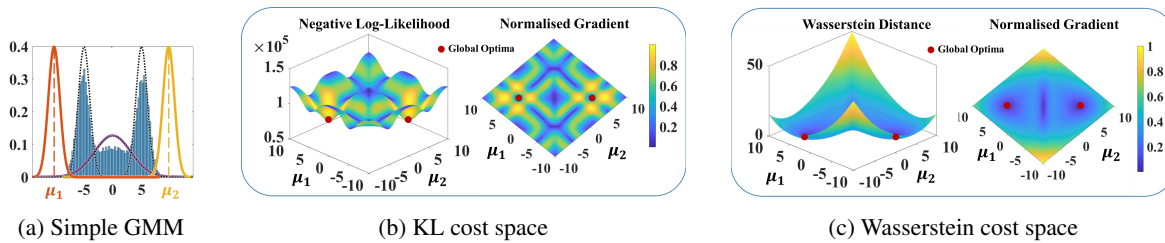(a) Simple GMM   (b) KL cost space   (c) Wasserstein cost space

Figure 1: Illustration of a simple one-dimensional GMM with 3 clusters. (a): The ground-truth of clusters is plotted via the black dotted lines; Parametric models $\rho(\boldsymbol{\theta})$ are denoted as colourful lines; We set ALL the parameters $\boldsymbol{\theta}$ of the GMM to the ground-truth EXCEPT the mean values of only two clusters, i.e., μ_1 and μ_2. (b): The KL divergence and the normalised gradient versus the variance of μ_1 and μ_2. (c): The Wasserstein distance and the normalised gradient versus the variance of μ_1 and μ_2.

concentrated as well, which also leads to ill-posed gradient descent.

On the other hand, by virtue of the reflection of sample space (Kolouri et al. 2017) within the Wasserstein distance[1] (Monge 1781) which employs a "cross-bin" comparison, many practical benefits may be achieved in learning tasks (Arjovsky, Chintala, and Bottou 2017). This property is particularly appealing in mixture model problems, where it ideally provides a comprehensive distance measure over all possible transport plans. The optimization space of the Wasserstein distance is also shown in Fig. 1-(c), of which smoothness is witnessed. Basically, there only exist global optimum in this case. The gradients on this space are also well behavioured, which promises to achieve superior optimisation. Most recently, Kolouri *et al.* (Kolouri, Rohde, and Hoffmann 2018) have adopted the Wasserstein distance for solving GMM problems. However, the aforementioned probability constraint enforces an extremely small stepsize (learning rate) during optimisation. Given the random projections of the sliced Wasserstein distance, this setup leads to extremely slow convergence or even non-convergence (especially in high dimensions as shown in our experiments). At the end of their optimisation, the EM algorithm is still needed to stabilised the algorithm.

Motivations and Contributions: Despite extensive benefits, optimising the GMM under the Wasserstein distance is still not feasible due to the probability constraint. To address this from a different perspective, we resort to the statistical manifold. We should further point out that our work is different from information geometry (Amari 1998), as directly establishing a manifold in the whole density space of mixture models is absolutely intractable and cumbersome in optimisation (Wang et al. 2015). Another problem within the Wasserstein space is that the geodesic between two mixture models may not lie in mixture models of the same type (Chen and Li 2018), which leads to non-convergence.

We propose to resolve this problem by introducing an approximation to the Wasserstein distance, followed by establishing a statistical manifold via the so induced distance, which exhibits the desirable property of being complete within mixture models. The subsequent optimisation along

this manifold intrinsically satisfies the probability constraint and ensures that the solution resides in the same mixture models. More importantly, minimising the induced distance is shown to be truly reducing the discrepancy between two mixture models, unlike most existing solutions which are based on the minimisation of the Euclidean distance from the optimal parameters. This ensures fast and stable convergence in optimisation. By realising that the existing Riemannian adaptive algorithms only make sense in updating vector parameters, we further develop a novel accelerated stochastic gradient descent method for updating the positive definite matrices.

In this way, our proposed framework makes it possible to incorporate a broad family of distributions, $\rho(\boldsymbol{\theta})$, including an important class of multivariate analysis techniques called elliptical distributions (Fang, Kotz, and Ng 1990) and to further investigate the mixture family termed the elliptical mixture model (EMM). We therefore provide computable and identifiable EMMs in a unified way, which demonstrates that EMMs are quite general and flexible and include the GMMs as special cases (Fang, Kotz, and Ng 1990).

Overall, this paper proposes a complete and efficient framework for solving general EMM problems, by establishing a statistical manifold under an approximate Wasserstein distance which promotes stability and efficiency, together with an adaptive stochastic gradient algorithm to further accelerate the optimisation[2]. Compared to the existing literature on mixture problems, the proposed solution achieves consistently superior performance not only in the GMM problems but also for general EMM problems. Our contributions can be summarised as follows:

- A unified framework for dealing with computable and identifiable EMMs, which introduces a rich choice of candidates for flexible finite mixture models.

- Establishment of the statistical manifold through the proposed approximate Wasserstein distance, which provides explicit and complete operations within the manifold.

- An Adaptive accelerated Riemannian gradient descent algorithm on the established manifold, to improve the optimisation and accelerate convergence.

---

[1]Throughout this paper, the term Wasserstein distance refers to the square-Wasserstein distance.

[2]The code of this paper is available at https://github.com/ShengxiLi/wass_emm

Computable and Identifiable EMMs

Elliptical distributions include a wide range of standard distributions, and it therefore comes as no surprise that a unified summary of computable candidates as components in the EMMs is a prerequisite to problem definition and subsequent solutions. A classical summary can be found in Chapter 3 in (Fang, Kotz, and Ng 1990); however, despite progress this framework is not general enough as various elliptical distributions are still missing, and more importantly, it involves complicated representations for each type of elliptical distributions. The existing literature also employs different notations and formulations of particular distributions, which may lead to confusion. To this end, we provide a simple and unified framework for summarising the existing computable elliptical distributions via the stochastic representation, which can then be used to constitute flexible and identifiable EMMs.

Preliminaries on Elliptical Distributions

A random variable, $\boldsymbol{\mathcal{X}} \in \mathbb{R}^m$, is said to exhibit an elliptical distribution if and only if it admits the following stochastic representation (Fang, Kotz, and Ng 1990),

$$\boldsymbol{\mathcal{X}} =^d \boldsymbol{\mu} + \mathcal{R}\boldsymbol{\Lambda}\boldsymbol{\mathcal{S}}, \tag{1}$$

where $\mathcal{R} \in \mathbb{R}^+$ is a non-negative real scalar random variable which models tail properties of the elliptical distribution; $\boldsymbol{\mathcal{S}} \in \mathbb{S}^{(m'-1)}$ is a random vector which is uniformly distributed on a unit spherical surface[3] with the pdf within the class of $2\pi^{-m'/2}\Gamma(m'/2)$; $\boldsymbol{\mu} \in \mathbb{R}^m$ is a mean (location) vector, while $\boldsymbol{\Lambda} \in \mathbb{R}^{m \times m'}$ is a matrix that transforms $\boldsymbol{\mathcal{S}}$ from a sphere to an ellipse, and "$=^d$" designates "the same distribution". For a comprehensive review of elliptical distributions, we refer to (Fang, Kotz, and Ng 1990).

When $m' = m$, that is, for a non-singular scatter matrix $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}\boldsymbol{\Lambda}^T$, the pdf for elliptical distributions does exist and has the following form

$$p(\mathbf{x}) = \underbrace{2\pi^{-\frac{m}{2}}\Gamma(\frac{m}{2})}_{c_m} \det(\boldsymbol{\Sigma})^{-\frac{1}{2}} g(\underbrace{(\mathbf{x} - \boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})}_{t}). \tag{2}$$

In (2), the term c_m serves as a normalisation term and relates solely to m. We denote the Mahalanobis distance $(\mathbf{x} - \boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})$ by t. The density generator, $g(t)$, can be explicitly expressed as $t^{-(m-1)/2}p_{\mathcal{R}}(\sqrt{t})$, where $t > 0$ and $p_{\mathcal{R}}(t)$ denotes the pdf of $\mathcal{R}$. Thus, $\mathcal{R}$, or equivalently[4] $\mathcal{R}^2$, fully characterises $g(\cdot)$, i.e., the type of elliptical distributions. For example, when $\mathcal{R}^2 =^d \chi_m^2$ (χ_m^2 denotes the chi-squared distribution of dimension m), then in (2), $g(t) \propto \exp(-t/2)$, which formulates the multivariate Gaussian distribution. Therefore, the elliptical distribution can be fully characterised by $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ and $\mathcal{R}$. For simplicity, the elliptical distribution in (2) will be denoted by $\boldsymbol{\mathcal{X}} \sim \mathcal{E}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathcal{R})$, where $\mathcal{E}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathcal{R}) = c_m\det(\boldsymbol{\Sigma})^{-1/2}g(t)$ of (2).

---

[3]The term $\mathbb{S}^{m'-1}$ is defined as $\mathbb{S}^{m'-1} := \{\mathbf{x} \in \mathbb{R}^{m'} : \mathbf{x}^T\mathbf{x} = 1\}$.

[4]The term $\mathcal{R}^2$ is frequently used in practice because $\mathcal{R}^2 =^d (\mathbf{x} - \boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})$.

Due to the fact that the $\mathcal{R}^2$ decides the type of elliptical distributions, we here provide a unified summary of elliptical distributions in Table 1; this is achieved through stochastic representations in (1). This makes it possible to avoid complicated formulations, and to instead classify different categories simply through several typical distributions of $\mathcal{R}^2$, which also allows for simple and intuitive sample generations for elliptical distributions. The proof for expressions in this table is provided in Appendix. Uniquely, this further clarifies the commonalities between the members of the elliptical family of distributions. More importantly, constructing an EMM with the candidates in Table 1 can be easily proved to be identifiable based on Theorem 2 in (Holzmann, Munk, and Gneiting 2006). It is thus convenient and safe to establish a well-defined EMM by the candidates summarised in Table 1.

Problem Statement and Notations

For generality, we assume that there are k mixtures in the EMM model, latent variables $\mathcal{Z}_i \in \{0, 1\}$ are binary, and the probability of choosing the i-th mixture is denoted by $p(\mathcal{Z}_i = 1) = \pi_i$, so that $\sum_{i=1}^k \mathcal{Z}_i = 1$ and $\sum_{i=1}^k \pi_i = 1$. If we use the random variable $\boldsymbol{\mathcal{Y}}$ to denote the EMM, then $\boldsymbol{\mathcal{Y}} =^d \sum_{i=1}^k \mathcal{Z}_i\boldsymbol{\mathcal{X}}_i \sim \sum_{i=1}^k \pi_i\mathcal{E}(\mathbf{x}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i, \mathcal{R})$, and the pdf of $\boldsymbol{\mathcal{Y}}$ can be expressed as

$$p(\mathbf{y}) = \sum_{i=1}^k \pi_i c_m \det(\boldsymbol{\Sigma}_i)^{-\frac{1}{2}} g((\mathbf{y} - \boldsymbol{\mu}_i)^T\boldsymbol{\Sigma}_i^{-1}(\mathbf{y} - \boldsymbol{\mu}_i)), \tag{3}$$

where the density generator $g(\cdot)$ can be chosen flexibly from Table 1, for which the identifiability is ensured (Holzmann, Munk, and Gneiting 2006). Thus, without ambiguity in the context, we shall denote the EMM $\boldsymbol{\mathcal{Y}}$ through the pdf of (3), that is, as $\boldsymbol{\mathcal{Y}}_{\boldsymbol{\theta}}$, where $\boldsymbol{\theta} = \{\boldsymbol{\pi}, \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i\}$ ($\boldsymbol{\pi} = [\pi_1, \pi_2, \ldots, \pi_k]^T$ and $i = 1, 2, \ldots, k$).

Statistical Manifold towards EMMs

As the Wasserstein distance possesses a Riemannian structure (Kolouri et al. 2017), it is then natural to treat each EMM as a "point" in the manifold, whereby the metric is defined by the Hessian of the Wasserstein distance. However, as pointed out by Chen (Chen and Li 2018), in the Wasserstein space, the geodesic between two points, i.e., two GMMs, does not necessarily belong to the GMM of the same type. A way to solve this problem is to pull back from the whole density space to the parametric space (Chen and Li 2018). This treatment, nevertheless, is still a toy solution due to two main considerations. The first deficiency is that the metric does not have a closed-form representation, and needs to be numerically obtained for every possible value in advance before the optimisation. This is highly prohibitive, especially for multivariate cases due to the curse of dimensionality. Besides computational intractability, operations such as the exponential mapping and the vector transport in the manifold cannot be well defined as there is typically a second-order differential equation involved to obtain those operations.

Table 1: $\mathcal{R}^2 \leftrightarrow$ Computable elliptical distributions

Types	$\mathcal{R}^2$	$\leftrightarrow$	$c_m \cdot g(t)$	Typical Multivariate Dist.
Kotz Type Ref[a]	$\mathcal{R}^2 =^d \mathcal{G}^{1/s}$, $\mathcal{G} \sim \text{Ga}(\frac{2a+m-2}{2s}, b)$ $a > 1-\frac{m}{2}, b, s > 0$		$= \left(\frac{\Gamma(m/2)sb^{(2a+m-2)/(2s)}}{\Gamma((2a+m-2)/2s)\pi^{m/2}}\right) t^{a-1}\exp(-bt^s)$	Gamma: $s = 1$ Weibull: $a = s$ Generalised Gaussian: $a = 1$ Gaussian: $a = 1,\ s = 1,\ b = \frac{1}{2}$
Scale Mixture of Normals $\mathcal{R}^2 =^d \mathcal{G}\cdot\mathcal{K}$, $\mathcal{G} \sim \text{Ga}(\frac{m}{2}, \frac{1}{2})$, $\mathcal{K}$ has different dist.	Pearson Type VII Ref[a] $\mathcal{K}^{-1} \sim \text{Ga}(s-\frac{m}{2}, 2v)$, $v > 0, s > m/2$		$= \left(\frac{(\pi v)^{-m/2}\Gamma(s)}{\Gamma(s-m/2)}\right)(1+t/v)^{-s}$	T-dist.: $s = \frac{m+v}{2}$ Cauchy: $v = 1, s = \frac{m+1}{2}$
	Hyperbolic Type Ref[b] $\mathcal{K} \sim \text{GIG}(v, a, \lambda)$ $v, a > 0, \lambda \in \mathbb{R}$		$= \left(\frac{(v/a)^{\lambda/2}}{(2\pi)^{m/2}\text{BeK}_\lambda(\sqrt{av})}\right)\frac{\text{BeK}_{(\lambda-m/2)}(\sqrt{av+vt})}{(\sqrt{a/v+t/v})^{m/2-\lambda}}$	Inverse-Gaussian: $\lambda = -1/2$ K-dist. (Ollila et al. 2012): $a \to 0,\ \lambda > 0$ Laplace: $a \to 0,\ \lambda = 1,\ v = 2$
	Other Types Ref[c] $\sqrt{\mathcal{K}} \sim \partial\text{Kov}(\frac{\mathcal{K}}{2})/\partial\mathcal{K}$		$= \exp(-t)/(1+\exp(-t))^2$	Logistic
	$\mathcal{K} \sim \text{S}\alpha\text{S}(\frac{a}{2}), a \in (0,2)$		$\propto \text{S}\alpha\text{S}(a)$	α-stable
Pearson Type II	$\mathcal{R}^2 \sim \text{Beta}(m/2, s), s > 1$		$= \left(\frac{\Gamma(m/2+s)}{\pi^{m/2}\Gamma(s)}\right)(1-t)^{s-1}, t \in [0,1]$	Ref[a]
Notations:	$a, b, s, v, \lambda, \alpha$ are adjustable parameters for different types of dist.; $\mathcal{G}$ and $\mathcal{K}$ are random variables related to $\mathcal{R}^2$; m is the dimension. Gamma dist.: $\text{Ga}(x,y) = y^x t^{x-1}\exp(-yt)/\Gamma(x)$; Inverse Gaussian dist.: $\text{GIG}(x,y,z) = \frac{(x/y)^{z/2}}{2\text{BeK}_z(\sqrt{xy})}t^{z-1}\exp(-\frac{xt^2+y}{2t})$; Kolmogorov-Smirnov dist.: $\text{Kov}(x) = 1 - 2\sum_{n=1}^{\infty}(-1)^{n+1}\exp(-2n^2x^2)$; Beta dist.: $\frac{\Gamma(x+y)}{\Gamma(x)\Gamma(y)}t^{x-1}(1-t)^{y-1}$; $\text{S}\alpha\text{S}(a)$: the symmetric α-stable dist. with index a; $\text{BeK}_x(y)$: the Bessel function of the third kind; $\Gamma(x)$: the Gamma function.			
References:	[a]: (Fang, Kotz, and Ng 1990); [b]: (Browne and McNicholas 2015); [c]: (Andrews and Mallows 1974; Stefanski 1991)			

We thus propose an approximate Wasserstein distance between two EMMs as a means to define a well-behaved manifold for the EMM problems; this is achieved by the property that the Wasserstein distance of two elliptical distributions is completely and explicitly defined. We then provide the Riemannian metric for the EMM problems according to the defined distance.

Approximate Wasserstein Distance between EMMs

We now focus on the distance between two EMMs $\mathcal{Y}_1$ and $\mathcal{Y}_2$, and propose an approximate Wasserstein distance by treating each distribution within an EMM as a "super-point" and defining a transport-like distance between those "super-points". A rigorous definition is given as follows.

Definition 1. *Given two EMMs $\mathcal{Y}_1$ and $\mathcal{Y}_2$, a discrepancy measure is defined as*

$$d_U(\mathcal{Y}_1, \mathcal{Y}_2) = \min_{\gamma(i,j)}\Big(\sum_{i,j}\frac{\gamma(i,j)}{k}d_W^2(\mathcal{X}_{i,1}, \mathcal{X}_{j,2}) + \arccos(\sum_{i,j}\gamma(i,j)\sqrt{\pi_{i,1}\pi_{j,2}})\Big), \quad (4)$$

where $d_W^2(\mathcal{X}_{i,1}, \mathcal{X}_{j,2})$ is the Wasserstein distance between the elliptical distributions $\mathcal{X}_{i,1}$ and $\mathcal{X}_{j,2}$. $\gamma(i,j)$ is binary $\in \{0,1\}$; for each i and j, $\gamma(i,j)$ satisfies $\sum_{i=1}^k \gamma(i,j) = 1$ and $\sum_{j=1}^k \gamma(i,j) = 1$.

Theorem 1. *Given two EMMs $\mathcal{Y}_1$ and $\mathcal{Y}_2$, the discrepancy measure $d_U(\mathcal{Y}_1, \mathcal{Y}_2)$ defines a distance.*

Proof. Please see Appendix. □

We provide our intuitions of defining the distance of (4). The term $\arccos(\sum_{i,j}\gamma(i,j)\sqrt{\pi_{i,1}\pi_{j,2}})$ intrinsically relates

to the probability constraint of $\sum_{i=1}^k \pi_{i,1} = \sum_{j=1}^k \pi_{j,2} = 1$, while $\gamma(i,j)$ operates as a bijection between mixture components in $\mathcal{Y}_1$ and $\mathcal{Y}_2$. $\sum_{i,j}\frac{\gamma(i,j)}{k}d_W^2(\mathcal{X}_{i,1}, \mathcal{X}_{j,2})$ can thus be regarded as a discrete transport between k uniformly distributed "super-points" $\mathcal{X}_{i,1}$ and $\mathcal{X}_{j,2}$, for which their cost is defined as $d_W^2(\mathcal{X}_{i,1}, \mathcal{X}_{j,2})$. We also noticed a computational distance proposed by (Chen, Georgiou, and Tannenbaum 2018). Our metric can be further seen as a restriction (or upper bound) of the Chen's metric because we operate via a one-to-one transport whilst the Chen's metric admits arbitrary transport plans. However, the most advantageous property of our metric is that it defines an explicit manifold over EMMs, which will be introduced shortly.

More importantly, $d_U(\mathcal{Y}_1, \mathcal{Y}_2)$ comprehensively reflects the discrepancy between $\mathcal{Y}_1$ and $\mathcal{Y}_2$ via a summation operation; this means $d_U(\mathcal{Y}_1, \mathcal{Y}_2) = 0$ if and only if both the difference between mixture components $\mathcal{X}_{i,1}$ and $\mathcal{X}_{j,2}$ and the difference between latent variables $\pi_{i,1}$ and $\pi_{j,2}$ equal to 0. The following lemma further proves that for balanced EMMs, $d_U(\mathcal{Y}_1, \mathcal{Y}_2)$ is an upper bound of the Wasserstein distance $d_W^2(\mathcal{Y}_1, \mathcal{Y}_2)$.

Lemma 1. *Given two balanced EMMs $\mathcal{Y}_1$ and $\mathcal{Y}_2$ (i.e., $\pi_{i,1} = \pi_{j,2} = 1/k$ for all i, j), $d_U(\mathcal{Y}_1, \mathcal{Y}_2)$ is an upper bound of the Wasserstein distance:*

$$d_W^2(\mathcal{Y}_1, \mathcal{Y}_2) \leq d_U(\mathcal{Y}_1, \mathcal{Y}_2). \quad (5)$$

The equality holds when $k = 1$.

Proof. Please see Appendix, where the tightness of the upper bound is also analysed. □

Statistical Manifold for EMM Problems

Before introducing the statistical manifold for EMMs, we would like to give credit to several most recent works on

Gaussian distributions and elliptical distributions, which lay a basis of our metric proposed in this section. We would like to mention (Knott and Smith 1984) of the Wasserstein distance of Gaussian measures and (Muzellec and Cuturi 2018) for the elliptical distributions. Most recently, the Riemannian manifold for Gaussian distributions has been established in an explicit form (Takatsu and others 2011; Malagò, Montrucchio, and Pistone 2018). Then, on the basis of the approximated Wasserstein distance, we provide the Riemannian metric of the EMM problems by calculating the Hessian of the proposed distance in Definition 1 as follows.

Lemma 2. *The approximate Wasserstein distance $d_U(\mathcal{Y}_1, \mathcal{Y}_2)$ represents an explicit Riemannian metric in the parametric space, and the corresponding Riemannian manifold is a product manifold of $\mathbb{R}^k \times \prod_{i=1}^{k}(\mathbb{R}^m \times \mathbb{P})$, where $\mathbb{P}$ is the $m \times m$ positive definite manifold:*
1) The manifold for the square root of π_i, i.e., $[\sqrt{\pi_1}, \sqrt{\pi_2}, \cdots, \sqrt{\pi_k}]^T$ is a sphere manifold of $\mathbb{R}^k$.
2) The manifold for $\boldsymbol{\mu}_i$ is the Euclidean space of $\mathbb{R}^m$.
3) The manifold for $\boldsymbol{\Sigma}_i$, i.e., $\mathbb{P}$, is defined by

$$ds^2 = \frac{\mathbb{E}[\mathcal{R}^2]}{m}(\mathbb{L}_{\boldsymbol{\Sigma}_i}[d\boldsymbol{\Sigma}])\boldsymbol{\Sigma}_i(\mathbb{L}_{\boldsymbol{\Sigma}_i}[d\boldsymbol{\Sigma}]). \quad (6)$$

In (6), $\mathbb{L}_{\mathbf{A}}[\mathbf{C}] = \mathbf{B}$ is a Lyapunov operator: $\mathbf{AB} + \mathbf{BA} = \mathbf{C}$, where $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{P}$. More importantly, the sectional curvature is non-negative ($= {}^m/_{\mathbb{E}[\mathcal{R}^2]}k_G$ where $k_G \geq 0$ is the sectional curvature for Gaussian cases (Takatsu and others 2011)). Recall that $\mathcal{R}$ is defined in (1).

Proof. Please see Appendix. □

Remark. *The metric in (6) provides a manifold for positive definite matrices. Compared to the best known manifold belonging to the Hadamard manifold (non-positive sectional curvature) (Sra and Hosseini 2015), the newly developed manifold provides an example of non-negative manifolds. It is actually further stated that an Alexandrov space has the non-negative curvature iff it is a Wasserstein space (Sturm and others 2006). The established metric thus provides a desirable reflection on the curvature when dealing with Wasserstein related EMM problems.*

Adaptively Accelerated Optimisation

By virtue of optimising on a statistical manifold, the probability constraint can be satisfied automatically; this allows us to incorporate various numerical algorithms when solving the EMM problems. More specifically, we first show that the constrained minimisation problem can be transformed to an unconstrained one when restricted on the statistical manifold, which results in the "vanilla" gradient descent on the manifold. We then propose an adaptively accelerated solver.

Vanilla Gradient Descent on Statistical Manifold

Similar to (Martens 2014), for gradient descent methods, the way for gradient descent in our work is given by,

$$\mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}^*} = \underset{\mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}}}{\arg\min} \, d_{SW}(\mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}}, \mathcal{Y}^*) + \frac{d_U(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}})}{c^2},$$

$$\rightarrow \Delta\boldsymbol{\theta}^* = \mathrm{Exp}(-\alpha\nabla_U(\boldsymbol{\theta})) \quad (7)$$

where $\Delta\boldsymbol{\theta}$ is the step size for the next iteration and c^2 denotes a sphere of all realisable distributions, which ensures searching for optimal $\Delta\boldsymbol{\theta}^*$ without being slowed down by the curvature; $d_{SW}(\mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}}, \mathcal{Y}^*)$ denotes the sliced Wasserstein distance (Rabin et al. 2011) between the parametric mixture model $\mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}}$ and the observed samples $\mathcal{Y}^*$. By approximating $d_U(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}})$ with its Hessian term (i.e., the inner product on the tangent space), we naturally obtain an unconstrained Riemannian gradient descent of the second row of (7), where $\nabla_U(\boldsymbol{\theta})$ is the Riemannian gradient on the statistical manifold defined by $d_U(\cdot, \cdot)$; α is the stepsize of each iteration; $\mathrm{Exp}(-\alpha\nabla_U(\boldsymbol{\theta}))$, called exponential mapping, projects the step movement $-\alpha\nabla_U(\boldsymbol{\theta})$ from the tangent space to the statistical manifold along geodesics (Absil, Mahony, and Sepulchre 2009), which ensures the probability constraint. On the other hand, when the proposed distance $d_U(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}})$ is changed to the Euclidean distance, (7) is the trivial gradient descent (Kolouri, Rohde, and Hoffmann 2018). However, as mentioned above, this gradient descent does not satisfy the probability constraint nor does it reflect the probability space curvature, which leads to inefficient and unstable optimisation.

In (7), the sliced Wasserstein distance (Rabin et al. 2011) provides a feasible solution in solving the semi-discrete Wasserstein problem, i.e., $\mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}}$ is continuous and $\mathcal{Y}^*$ is the sum of Dirac masses (Lévy 2015). It uses unit random projections to turn the original problem to a one-dimensional Wasserstein problem via the Radon transform (Rabin et al. 2011), so that a closed-form solution can be obtained[5]. More importantly, as the sliced Wasserstein distance is composed of a set of random projections (Rabin et al. 2011), it is then natural to implement a stochastic gradient descent on the Riemannian manifold for each random projection (denoted as $\mathbf{p} \in \mathbb{S}^{m-1}$). The random projection process also allows for parallel gradient descent. We provide the basic operations of problem in (7) in Table 2, while the details are provided in Appendix.

Adaptively Accelerated Algorithm

On the basis of several accelerated Riemannian stochastic gradient descent methods which adopt the first-order moment information (Zhang and Sra 2016), recently Becigneu and Ganea (Becigneul and Ganea 2019) further proposed a Riemannian adaptive method by employing the second-order moments. Although this adaptive stochastic gradient descent method (Becigneul and Ganea 2019) may be incorporated to our work to further improve the convergence, we argue that it only makes sense for updating vector parameters, i.e., $\boldsymbol{\pi}^h$ and $\boldsymbol{\mu}_i^h$, because in this case the second-order moment can be calculated element-wise by a decomposition into product manifolds (Becigneul and Ganea 2019). This is similar to the *Adam* algorithm in the Euclidean space (Kingma and Ba 2014). When updating the matrix parameter, $\boldsymbol{\Sigma}_i^h$, however, it does not capture the second-order information properly due to the direct accumulation of the

---

[5]It needs to be pointed out that other approximate distances (e.g., entropy relaxed Wasserstein distances) for the semi-discrete problem can also be seamlessly adopted in (7) of our work.

Table 2: Basic operations of the manifold in Lemma 2

For the h-th iteration		
	$\nabla_U(\cdot)$	$\mathrm{Exp}(-\alpha\nabla_U(\cdot))$
$\sqrt{\boldsymbol{\pi}}^h$	$\nabla_E(\sqrt{\boldsymbol{\pi}}^h) - (\sqrt{\boldsymbol{\pi}}^h)^T\nabla_E(\sqrt{\boldsymbol{\pi}}^h)\cdot\sqrt{\boldsymbol{\pi}}^h$	$\cos(\|\alpha\nabla_U(\sqrt{\boldsymbol{\pi}}^h)\|_2)\sqrt{\boldsymbol{\pi}}^h - \frac{\sin(\|\alpha\nabla_U(\sqrt{\boldsymbol{\pi}}^h)\|_2)}{\|\nabla_U(\sqrt{\boldsymbol{\pi}}^h)\|_2}\nabla_U(\sqrt{\boldsymbol{\pi}}^h)$
$\boldsymbol{\mu}_i^h$	$\nabla_E(\boldsymbol{\mu}_i^h)$	$\boldsymbol{\mu}_i^h - \alpha\nabla_U(\boldsymbol{\mu}_i^h)$
$\boldsymbol{\Sigma}_i^h$	$\nabla_E(\boldsymbol{\Sigma}_i^h)\boldsymbol{\Sigma}_i^h + \boldsymbol{\Sigma}_i^h\nabla_E(\boldsymbol{\Sigma}_i^h)$	$(\mathbb{L}_{\boldsymbol{\Sigma}_i^h}[-\alpha\nabla_U(\boldsymbol{\Sigma}_i^h)]+I)\boldsymbol{\Sigma}_i^h(\mathbb{L}_{\boldsymbol{\Sigma}_i^h}[-\alpha\nabla_U(\boldsymbol{\Sigma}_i^h)]+I)$
For the Radon transform $\mathbf{p}\in\mathbb{S}^{m-1}$, specified in the EMM problems:		
$\nabla_E(\sqrt{\boldsymbol{\pi}}^h)$	The i-th dimension of $\nabla_E(\sqrt{\boldsymbol{\pi}}^h)$ is $2\int_{\mathbb{R}}c_m\phi(y)\sqrt{\pi_i}^h(\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p})^{-1/2}g(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}})dy$	
$\nabla_E(\boldsymbol{\mu}_i^h)$	$=\left(-2\int_{\mathbb{R}}c_m\phi(y)\pi_i^h(\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p})^{-3/2}g'(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}})(y-\mathbf{p}^T\boldsymbol{\mu}_i)dy\right)\mathbf{p}$	
$\nabla_E(\boldsymbol{\Sigma}_i^h)$	$=\left(-\int_{\mathbb{R}}c_m\phi(y)\pi_i^h(\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p})^{-3/2}\left(\frac{1}{2}g(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i^h)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}})+g'(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i^h)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}})\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i^h)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}}\right)dy\right)\mathbf{p}\mathbf{p}^T$	
$\phi(y)$ is the Kantorovich potential (Chen and Li 2018). $\nabla_E(\cdot)$ denotes the Euclidean gradient with regard to $\sqrt{\boldsymbol{\pi}}^h$, $\boldsymbol{\mu}_i^h$ and $\boldsymbol{\Sigma}_i^h$.		

Alg. 1: Riemannian adaptively accelerated manifold optimisation

Input: n observed samples $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n$; stepsize $\{\alpha^h\}_{h=1}^H$; hyper-parameters $\{\beta_1^h\}_{h=1}^H$ and β_2
Initialise: 1^{st}-order moment $\{\mathbf{u}_i^0\}_{i=1}^k$; 2^{nd}-order moment $\{\mathbf{v}_i^0\}_{i=1}^k$;
  for: $h=1$ to H **do**
    Random projection: $\mathbf{p}\in\mathbb{S}^{(m-1)}$
    Update $\sqrt{\boldsymbol{\pi}}^{h+1}=\mathrm{Exp}(-\alpha\nabla_U(\sqrt{\boldsymbol{\pi}}^h))$
    for: $i=1$ to k **do**
      Update $\boldsymbol{\mu}_i^{h+1}=\mathrm{Exp}(-\alpha\nabla_U(\boldsymbol{\mu}_i^h))$
      Update $\boldsymbol{\Sigma}_i^h$ by the *Dadam*:
      $\mathbf{u}_i^h = \beta_1^h\varphi_{\boldsymbol{\Sigma}_i^{h-1}\to\boldsymbol{\Sigma}_i^h}(\mathbf{u}_i^{h-1})+(1-\beta_1^h)\nabla_U(\boldsymbol{\Sigma}_i^h)$
      $\mathbf{v}_i^h = \beta_2\mathbf{v}_i^{h-1}+(1-\beta_2)\nabla_E(\boldsymbol{\Sigma}_i^h)\nabla_E(\boldsymbol{\Sigma}_i^h)^T$
      $\mathrm{adp}_i^h = \max\{\mathbf{p}^T\mathbf{v}_i^h\mathbf{p}, \mathrm{adp}_i^{h-1}\}$
      $\boldsymbol{\Sigma}_i^{h+1}=\mathrm{Exp}(-\alpha^h\mathbf{u}_i^h/\sqrt{\mathrm{adp}_i^h})$
    end for
  end for
Return: $\boldsymbol{\pi}^H$, $\{\boldsymbol{\mu}_i^H\}_{i=1}^k$, $\{\boldsymbol{\Sigma}_i^H\}_{i=1}^k$

second-order moments over the whole matrix manifold; this slows down convergence in the optimisation.

The key difficulty for the matrix case is that it is meaningless to accumulate second-order moments in an element-wise manner due to the structure within the matrix. We here view the second-order moments of matrices from another perspective, by realising that a positive definite matrix can be decomposed into a set of eigenvectors and the corresponding eigenvalues. The eigenvectors can be regarded as a set of projection directions and the eigenvalues are scalars that can be connected with the accumulation in those directions. Furthermore, we can see from Table 2 that the Euclidean gradient $\nabla_E(\boldsymbol{\Sigma}_i^h)$ consists of a scalar weight multiplied by a rank-1 matrix characterised by the direction $\mathbf{p}\mathbf{p}^T$. Therefore, instead of the element-wise accumulation, we propose a direction-wise accumulation of second-order moments to adaptively adjust the stepsize when updating the matrix. The details of our algorithm in Algorithm 1, with our directional adaptive accelerated method (*Dadam*) of updating $\boldsymbol{\Sigma}_i^h$ being achieved by $\mathbf{v}_i^h$ and adp_i^h.

Moreover, as there is no explicit parallel transport in the Wasserstein space, we propose a new means of vector transport $\varphi_{\boldsymbol{\Sigma}_i^{h-1}\to\boldsymbol{\Sigma}_i^h}(\mathbf{u}_i^{h-1}) = \mathbb{L}_{\boldsymbol{\Sigma}_i^{h-1}}[\mathbf{u}_i^{h-1}]\boldsymbol{\Sigma}_i^h + \boldsymbol{\Sigma}_i^h\mathbb{L}_{\boldsymbol{\Sigma}_i^{h-1}}[\mathbf{u}_i^{h-1}]$ to accumulate the first-order moments.

As for the convergence analysis, because our goal is to minimise the Wasserstein distance, qualitatively, our manifold, defined by the Hessian of the approximate Wasserstein distance as in Lemma 2, implicitly involves the second-order information and accelerates the convergence. Furthermore, adp_i^h denotes the accumulation in the current projection $\mathbf{p}$, which is a scalar (for computational ease) and does not need any eigen-decomposition operations. This, on the one hand, avoids the high computational complexity of calculating any inverse of matrices, making our method the same computational complexity as that of (Becigneul and Ganea 2019) at each iteration. On the other hand, it also makes our convergence analysis similar to that of (Becigneul and Ganea 2019), upon realising that the sectional curvature is automatically bounded from below; we omit the analysis here.

Experimental Results

We evaluated the effectiveness of our manifold and the proposed *Dadam* on both synthetic data and image data, by employing 4 EMMs, i.e., mixtures of *Gaussian*, *Logistic*, *Cauchy* and *Gamma* ($s=1, a=2, b=0.5$ in Table 1).

Synthetic Data: Each synthetic dataset contains different mixtures of Gaussian distributions, with $10,000$ samples in total. Specifically, we employed three types of synthetic datasets, i.e., $\{m=2, k=3\}$, $\{m=8, k=9\}$ and $\{m=16, k=27\}$. Each type contains 10 randomly generated mixture datasets (the eccentricity ε and the separation c of (Dasgupta 1999) were equal to 10). For every algorithm, we tested on each dataset over 10 random initialisations and recorded the mean and standard deviation. Then, we averaged over all datasets to obtain the performance of each algorithm.

Image Data: We adopted the MNIST (LeCun, Cortes, and Burges 2010) dataset as well as the BSDS500 (Arbelaez et al. 2011) benchmark dataset in our evaluation. For the MNIST dataset, each image in both the training and testing sets was downsampled to 3×3 and 5×5; by vectorising each we obtained two test data with $(n\times m)$ being $(70,000\times 9)$ and $(70,000\times 25)$, respectively. The evaluation on the BSDS500 dataset is slightly different from that on the MNIST dataset, which enables to verify our method on some basic tasks such as the image reconstruction. Specifically, instead of modelling over the whole dataset, each image in the dataset was treated as one test data; this means we had
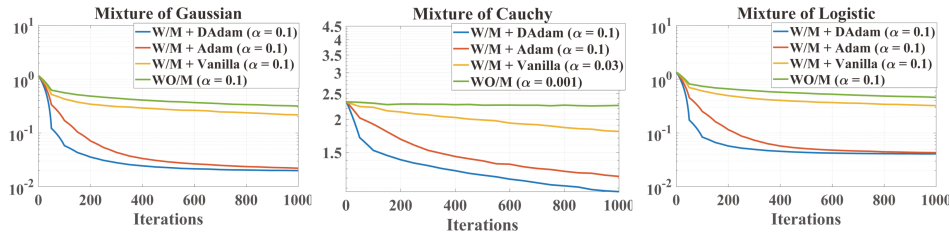
Figure 2: The Wasserstein distance against the number of iterations for the four considered algorithms, averaged over the three types of datasets ($\{m, k\} = \{2, 3\}, \{8, 9\}, \{16, 27\}$). The best learning rate is shown in the legend. *1000* iterations are for illustration purpose. The computational time per iteration is discussed in Table 3.
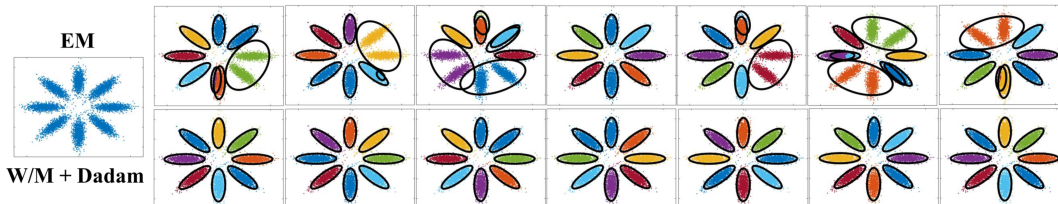

Figure 3: The optimised GMMs via the EM and our (*W/M+Dadam*) methods on a flower-shaped synthetic data ($N = 10,000$) with random initialisations. Each cluster in the flower-shaped data is composed by Gaussian-distributed samples.

Table 3: Comparisons among our, WO/M and EM methods for GMMs by varying m and k.

	$m = 2, k = 3$			$m = 8, k = 9$			$m = 16, k = 27$		
	Wass	NLL	Time per Ite.	Wass	NLL	Time per Ite.	Wass	NLL	Time per Ite.
W/M+Dadam	**0.01 ± 0.00**	**5.10 ± 0.00**	4.35ms	**0.03 ± 0.03**	**19.49 ± 0.03**	9.83ms	**2.39 ± 0.33**	**44.92 ± 1.04**	32.51ms
WO/M	0.03 ± 0.00	5.12 ± 0.00	4.23ms	4.26 ± 3.61	22.27 ± 1.69	9.64ms	5e3 ± 589	Inf	29.77ms
EM	0.65 ± 1.06	5.30 ± 0.35	3.10ms	0.70 ± 0.46	19.83 ± 0.29	13.92ms	2.51 ± 1.27	48.84 ± 1.24	215.00ms

Table 4: Comparison performance of our, WO/M and IRA algorithms over the BSDS500 dataset via five metrics

	Gaussian Mixture			Logistic Mixture			Cauchy Mixture			Gamma Mixture		
	Our	WO/M	IRA	Our	WO/M	IRA	Our	WO/M	IRA	Our	WO/M	IRA
Wass	**35.0**	89.5	53.6	**45.8**	154	68.0	**179**	211	253	**35.5**	510	–
	±4.68	±10.0	±21.9	±12.7	±**10.9**	±24.2	±15.3	±**7.77**	±65.4	±**1.82**	±71.2	–
NLL	**11.8**	12.2	11.8	**10.6**	12.5	10.6	12.1	12.3	**12.0**	**13.3**	19.1	–
	±**0.01**	±0.11	±0.05	±**0.02**	±0.25	±0.06	±**0.01**	±0.01	±0.05	±**0.07**	±0.86	–
PSNR	**18.9**	18.6	18.3	**18.8**	18.2	17.8	19.3	**19.3**	18.4	**21.0**	18.0	–
(dB)	±**0.07**	±0.18	±0.66	±**0.11**	±0.18	±0.67	±**0.05**	±0.06	±0.56	±**0.08**	±0.58	–
SSIM	**0.68**	0.65	0.66	**0.67**	0.62	0.63	**0.70**	0.70	0.69	**0.73**	0.59	–
	±**0.00**	±0.01	±0.03	±**0.01**	±0.01	±0.03	±**0.00**	±0.00	±0.02	±**0.00**	±0.03	–
FailR	**0.20**%	53.1%	2.98%	**0.46**%	75.5%	2.68%	**1.16**%	0.82%	17.16%	**0**%	2.02%	100%

500 test data in the BSDS500 dataset. Furthermore, each test data in both the MNIST and BSDS500 datasets was tested for every algorithm with 10 random initialisations, where the mean and standard deviation were also recorded.

Parameter Settings and Metrics: We found the best learning rates α by searching from $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3\}$. Similar to (Kingma and Ba 2014; Becigneul and Ganea 2019), $\beta_{11} = 0.9$ and $\beta_2 = 0.999$. The maximum number of iterations for testing the synthesis data was $2,000$ and that for the image data was $10,000$. The Wasserstein distance (simplified as *Wass*) and

averaged negative log-likelihood (*NLL*) were mainly employed for comparison. The optimisation fail ratio (*FailR*) was also reported to show the stability of optimisation under various initialisations. For image data, we adopted two well-known metrics, namely, the peak signal-to-noise ratio (*PSNR*) and structural similarity index (*SSIM*) (Wang et al. 2004), to evaluate the quality of reconstructed images via maximising the posterior of optimised EMMs.

Baselines: Each main part of our method was assessed. The vanilla Riemannian gradient descent method on our manifold is denoted by *W/M + Vanilla*. The Riemannian

Adam of (Becigneul and Ganea 2019) with our manifold is *W/M + Adam*. Our proposed *Dadam* is denoted by *W/M + Dadam* or *Our* interchangeably. We also performed comparisons with the trivial gradient descent over the sliced Wasserstein distance without the manifold (denoted by *WO/M*), *which is the basics of (Kolouri, Rohde, and Hoffmann 2018)*. The EM-type methods were also compared, which are denoted as *EM* for GMMs and *IRA* for other EMMs (Kent and Tyler 1991). All the methods were run on the Matlab 2017a under Intel Core(TM) i7-6700 CPU, where the time was recorded.

Assessment over the *Dadam* on Synthesis Data

We first evaluated the effectiveness of our $Dadam$ in adaptively accelerating the convergence. In this part, we only optimised the Σ_i whilst setting the π and μ_i as the ground truth of the synthesis data. Fig. 2 illustrates the convergence speed; observe that the optimisation on our established manifold exhibits much faster convergence. Furthermore, both the adaptive methods of (Becigneul and Ganea 2019) and the $Dadam$ were shown to boost the speed. The proposed $Dadam$ achieved the fastest convergence, which verifies the effectiveness of our *Dadam*.

Overall Comparisons on Synthesis Data

We next compared our algorithm (i.e., *W/M+Dadam*) with the *WO/M* and the *EM* methods under GMM problems. Similar results can be found for other EMMs; the result for GMMs is given in Table 3. It should be pointed out that all parameters (π, μ_i and Σ_i) were optimised in this test and each result was reported via the corresponding best learning rates. From this table, observe that our algorithm not only achieves an extremely stable estimation (the lowest deviations) but also the lowest values of both the Wasserstein distance and the NLL. An illustrative example can be found in Fig. 3, in which our algorithm consistently achieves the optimal clustering but the optimised GMMs via the *EM* are highly unstable. From Table 3, we should also point out that although our algorithm implements a set of manifold operations, the computational burden over that without manifold optimisation (i.e., *WO/M*) is in average around 5%. However, our algorithm enjoys much faster speed compared to the *EM* algorithm per iteration.

Overall Comparisons on the MNIST

We further compared the computational complexity of our algorithm (i.e., *W/M+Dadam*) with the *EM* method, besides the time per iteration reported in Table 3. Fig. 4 plots the decrease of Wasserstein loss with regard to the computational time when optimising GMMs and the trends for other EMMs are the same. We can see clearly from this figure that our algorithm consistently achieves the lowest Wasserstein cost. The computational time of our algorithm is comparable to that of the *EM* algorithm when $m = 3 \times 3 = 9$; our algorithm, however, requires much less time when in higher dimensions such as $m = 5 \times 5 = 25$. We shall also need to point out that the random projections in our algorithm can be implemented parallelly, which can further accelerate the convergence speed of our algorithm.
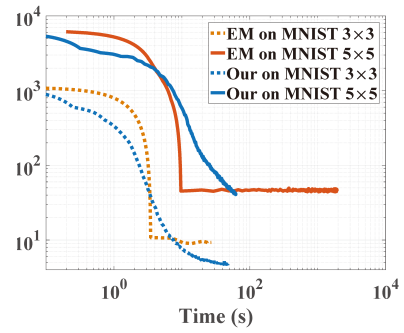

Figure 4: Wasserstein distance against the computational time for the *EM* and *Our*, on the downsampled 3×3 and 5×5 MNIST test data. The best learning rate is chosen as $\alpha = 0.1$. The cluster number is chosen as $k = 3$. The maximal iterations are set to $10,000$ for both methods.

Overall Comparisons on the BSDS500

Finally, we evaluated our algorithm on the BSDS500 dataset, with the results shown in Table 4 shows the results. From this table, we see some improvements of using the Wasserstein distance instead of the KL divergence, by comparing the *WO/M* with the *IRA*. However, due to the probability constraint, the *WO/M* is highly unstable and its fail ratio reaches $> 50\%$ for the Gaussian mixture and Logistic mixture. More importantly, by optimising on our established manifold together with our *Dadam* method (denoted as *Our*), we can consistently achieve best performances over the five metrics. Again, our algorithm also achieve highly stable estimation, taking advantages from the Wasserstein distance, compared to the *IRA* which operates under the KL. Another advantage of our algorithm is its universal convergence, while the *IRA* does not converge for the mixture of *Gamma*, which limits the flexibility of EMMs.

Conclusion

We have introduced a new and complete framework for solving general EMMs. To this end, we have first provided a unified and easy-to-use summary of the candidates for identifiable EMMs, which is achieved via a stochastic representation. Then, an approximate Wasserstein distance for EMMs has been proposed, which unlike the existing Wasserstein distances allows the corresponding metrics to be explicitly calculated. The so established manifold has been shown to consistently improve performances in terms of Wasserstein cost and even the NLL cost, and also significantly stabilise the optimisation on EMMs, making them robust to initialisations. We have further proposed a directional adaptively accelerated algorithm to enhance and stabilise the convergence, the performance of which has been validated through comprehensive experimental results.

Appendix

Proof for expressions in Table 1

Consider the term $p_{\mathcal{R}}(t) = g(t^2) \cdot t^{m-1}$. The pdf of $\mathcal{R}^2$ is then obtained accordingly as,

$$p_{\mathcal{R}^2}(t) = \frac{1}{2} \cdot g(t) \cdot t^{m/2-1}. \tag{8}$$

The term $g(t)$ can be further decomposed as $c \cdot g_c(t)$, where $g_c(t)$ is a nuclear term that only relates to t, and c is a normalisation term of $p_{\mathcal{R}}(t)$. This allows us to prove the results in Table 1 related to computable elliptical probability distribution functions (pdf) as follows.

- **Kotz type distributions:** The nuclear term of density generator is $g_c(t) = t^{a-1}\exp(-bt^s)$ and c can be calculated from the condition $\int p_{\mathcal{R}}(t)dt = \int g(t^2) \cdot t^{m-1}dt = 1$ as,

$$c = (\int_0^\infty t^{2a+m-3}\exp(-bt^{2s})dt)^{-1} = \frac{2sb^{\frac{2a+m-2}{2s}}}{\Gamma(\frac{2a+m-2}{2s})}. \tag{9}$$

We now arrive at the pdf of $\mathcal{R}^2$ for the Kotz type distributions, in the form

$$p_{\mathcal{R}^2}(t) = \frac{sb^{\frac{2a+m-2}{2s}}}{\Gamma(\frac{2a+m-2}{2s})}t^{\frac{m}{2}+a-2}\exp(-bt^s). \tag{10}$$

However, in practice, it is intractable to generate samples of $\mathcal{R}^2$ as in (10). On the other hand, it can be found that $\mathcal{G}^{1/s}$ has the same distribution as $\mathcal{R}^2$ when $\mathcal{G}$ is Gamma distributed. Specifically, when $\mathcal{G} \sim \text{Ga}(\frac{2a+m-2}{2s}, b)$, we obtain the pdf of $\mathcal{G}^{1/s}$ as,

$$p_{\mathcal{G}^{1/s}}(t) = st^{s-1}\frac{b^{\frac{2a+m-2}{2s}}t^{\frac{2a+m-2}{2}-s}\exp(-bt^s)}{\Gamma(\frac{2a+m-2}{2s})}$$
$$= \frac{sb^{\frac{2a+m-2}{2s}}}{\Gamma(\frac{2a+m-2}{2s})}t^{\frac{m}{2}+a-2}\exp(-bt^s). \tag{11}$$

This proves that $\mathcal{R}^2 =^d \mathcal{G}^{1/s}$ for the Kotz type distributions.

- **Scale mixture of normals distributions:** The scale mixture of normal distributions consists of a mixture of zero-mean normal distributions (denoted as $\boldsymbol{\mathcal{X}} = \sqrt{\mathcal{K}}\boldsymbol{\mathcal{N}}$, where $\boldsymbol{\mathcal{N}}$ is zero-mean normal distribution and $\mathcal{K}$ is called the mixing distribution with pdf $p_{\mathcal{K}}(t)$). Correspondingly, we can write the pdf of $\boldsymbol{\mathcal{X}}$ as

$$p_{\boldsymbol{\mathcal{X}}}(\mathbf{x}) = \int_t p_{\boldsymbol{\mathcal{X}}}(\mathbf{x}|t)p_{\mathcal{K}}(t)dt \propto \int_t t^{-\frac{m}{2}}\exp(\frac{\mathbf{x}^T\Sigma^{-1}\mathbf{x}}{t})p_{\mathcal{K}}(t)dt. \tag{12}$$

Furthermore, to represent a scale mixture of normal distributions in the form of stochastic representations, by inspection we see that a normal distribution, $\boldsymbol{\mathcal{N}}$, can be represented by a multiplication of $\sqrt{\chi_m^2}$ and $\boldsymbol{\mathcal{S}}$ (from the stochastic representation), so that the following holds:

$$\boldsymbol{\mathcal{X}} = \sqrt{\mathcal{K}}\boldsymbol{\mathcal{N}} = \sqrt{\mathcal{K}\cdot\chi_m^2}\cdot\boldsymbol{\mathcal{S}} = \sqrt{\mathcal{K}\cdot\mathcal{G}}\boldsymbol{\mathcal{S}}, \tag{13}$$

where $\mathcal{G} \sim \text{Ga}(m/2, 2)$. In this case, $\mathcal{R}^2 =^d \mathcal{K} \cdot \mathcal{G}$ by the stochastic representation. Therefore, although (12)

and (13) are equivalent, (13) provides a unified stochastic representation within the elliptical family. This allows us to discuss different $\mathcal{K}$ for different sub-classes of scale mixture normals in the following.

For the Pearson type VII distributions, where the nuclear density generator is given by $g_c(t) = (1 + t/v)^{-s}$, we can obtain the corresponding normalisation term c as follows,

$$c = (\int_0^\infty t^{m-1}(1 + t^2/v)^{-s}dt)^{-1}$$
$$= \frac{2\Gamma(s)}{(v)^{m/2}\Gamma(s - m/2)\Gamma(m/2)}. \tag{14}$$

Consequently, the pdf of $\mathcal{R}^2$ becomes:

$$p_{\mathcal{R}^2}(t) = \frac{\Gamma(s)}{\Gamma(s - \frac{m}{2})\Gamma(\frac{m}{2})v^{\frac{m}{2}}}(1 + \frac{t}{v})^{-s}t^{\frac{m}{2}-1}. \tag{15}$$

Again, the rather complicated form of (15) makes it impossible to generate samples. It can be found that when $p_{\mathcal{K}}(t) = \frac{(2v)^{s-m/2}}{\Gamma(s-m/2)}t^{m/2-s-1}\exp(-2v/t)$ and $\mathcal{G} \sim \text{Ga}(m/2, 2)$, $\mathcal{R}^2$ in (15) has the same pdf as $\mathcal{K} \cdot \mathcal{G}$. This can be verified as follows,

$$p_{\mathcal{R}^2}(t) = \int_0^\infty p_{\mathcal{G}}(t|\tau)p_{\mathcal{K}}(\tau)d\tau$$
$$= \int_0^\infty \frac{(\frac{2}{\tau})^{\frac{m}{2}}t^{\frac{m}{2}-1}\exp(-\frac{2t}{\tau})}{\Gamma(\frac{m}{2})}\frac{(2v)^{s-\frac{m}{2}}}{\Gamma(s - \frac{m}{2})}\tau^{\frac{m}{2}-s-1}\exp(-\frac{2v}{\tau})d\tau$$
$$= \frac{\Gamma(s)t^{\frac{m}{2}-1}}{\Gamma(s - \frac{m}{2})\Gamma(\frac{m}{2})v^{\frac{m}{2}}(1 + \frac{t}{v})^s}, \tag{16}$$

where $\mathcal{K}^{-1} \sim \text{Ga}(s - \frac{m}{2}, 2v)$ and $\mathcal{R}^2 =^d \mathcal{K} \cdot \mathcal{G}$.

Moreover, for other types within the scale mixture normals, Barndorff *et al.* (Barndorff-Nielsen, Kent, and Sørensen 1982) have proved that the generalised hyperbolic distributions can be formulated in the form of (12) when $\mathcal{K}$ satisfies the inverse-Gaussian distribution. The elliptical logistic distributions are also mixtures of normals, where $\sqrt{\mathcal{K}}$ in (12) relates to the Kolmogorov-Smirnov distribution (Stefanski 1991). Besides, the relationship between the scale mixture normals and the α-stable distribution is given in (Andrews and Mallows 1974), which satisfies our equivalent form of (13). We omit the tedious proofs here.

- **Pearson type II distributions:** The nuclear density generator is $g_c(t) = (1 - t)^{s-1}$ with the constraint $t \in [0, 1]$. We omit the proof here because the calculation of c and the pdf of $\mathcal{R}^2$ are well documented in (Fang, Kotz, and Ng 1990).

This completes the proof of Table 1.

Proof of Theorem 1

Symmetry: It is easy to verify the symmetry of $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)$, i.e., $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2) = d_U(\boldsymbol{\mathcal{Y}}_2, \boldsymbol{\mathcal{Y}}_1)$.
Non-negativity: Due to $\arccos(\cdot) \geq 0$, $d_W^2(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2}) \geq 0$ and $\gamma(i, j) \geq 0$ for all i, j, $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2) \geq 0$. More

importantly, the equality holds if and only if $\boldsymbol{\mathcal{Y}}_1 =^d \boldsymbol{\mathcal{Y}}_2$. In this case, $\min_{\gamma(i,j)} \sum_{i,j} \frac{\gamma(i,j)}{k} d_W^2(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2}) = 0$ and $\min_{\gamma(i,j)} \arccos(\sum_{i,j} \gamma(i,j)\sqrt{\pi_{i,1}\pi_{j,2}}) = \arccos(\sum_i \pi_{i,1}) = \arccos(1) = 0$, which results in $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2) = 0$.

Triangle inequality: Because $\gamma(i,j)$ is binary $\in \{0,1\}$ for each pair $\{i,j\}$, and it satisfies $\sum_{i=1}^k \gamma(i,j) = 1$ and $\sum_{j=1}^k \gamma(i,j) = 1$, the $\gamma(i,j)$ operates as a bijection between the elliptical distributions within the first and second EMMs.

Then, in order to prove the triangle property, we denote the third EMM as $\boldsymbol{\mathcal{Y}}_3 =^d \sum_{h=1}^k z_{h,3} \boldsymbol{\mathcal{X}}_{h,3}$, and investigate the relationship between $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)$ and $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_3) + d_U(\boldsymbol{\mathcal{Y}}_2, \boldsymbol{\mathcal{Y}}_3)$. We use $\gamma^*(\cdot, \cdot)$ to denote the optimal $\gamma(\cdot, \cdot)$ in the defined distance $d_U(\cdot, \cdot)$, and also define the following function

$$f(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2, \gamma(i,j)) =$$
$$\sum_{i,j} \gamma(i,j) d_W^2(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2}) + \arccos(\sum_{i,j} \gamma(i,j)\sqrt{\pi_{i,1}\pi_{j,2}}).$$
(17)

By Definition 1, $\min_{\gamma(i,j)} f(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2, \gamma(i,j)) = f(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2, \gamma^*(i,j)) = d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)$.

More importantly, for two arbitrary $\gamma(i,h)$ and $\gamma(h,j)$, their combination $\gamma(i,h) \cap \gamma(h,j) = \sum_{h=1}^k \gamma(i,h) \cdot \gamma(h,j)$ also formulates a transport plan $\gamma(i,j)$ because $\gamma(i,h) \cap \gamma(h,j)$ is still binary and $\sum_{i=1}^k \gamma(i,h) \cap \gamma(h,j) = \sum_{j=1}^k \gamma(i,h) \cap \gamma(h,j) = 1$.

We therefore now arrive at

$$d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_3) + d_U(\boldsymbol{\mathcal{Y}}_2, \boldsymbol{\mathcal{Y}}_3)$$
$$= f(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_3, \gamma^*(i,h)) + f(\boldsymbol{\mathcal{Y}}_2, \boldsymbol{\mathcal{Y}}_3, \gamma^*(j,h))$$
$$= \frac{1}{k}\left(\sum_{i,h}\gamma^*(i,h)d_W^2(\boldsymbol{\mathcal{X}}_{i,1},\boldsymbol{\mathcal{X}}_{h,3}) + \sum_{h,j}\gamma^*(h,j)d_W^2(\boldsymbol{\mathcal{X}}_{h,3},\boldsymbol{\mathcal{X}}_{j,2})\right)$$
$$+ \arccos(\sum_{i,h}\gamma^*(i,h)\sqrt{\pi_{i,1}\pi_{h,3}}) + \arccos(\sum_{h,j}\gamma^*(h,j)\sqrt{\pi_{h,3}\pi_{j,2}})$$
$$\geq \frac{1}{k}\sum_{i,j}\left(\gamma^*(i,h)\cap\gamma^*(h,j)\right)d_W^2(\boldsymbol{\mathcal{X}}_{i,1},\boldsymbol{\mathcal{X}}_{j,2})$$
$$+ \arccos\left(\sum_{i,j}(\gamma^*(i,h)\cap\gamma^*(h,j))\sqrt{\pi_{i,1}\pi_{j,2}}\right)$$
$$= f(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2, \gamma^*(i,h)\cap\gamma^*(h,j))$$
$$\geq f(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2, \gamma^*(i,j))$$
$$= d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2).$$
(18)

In (18), the first inequality holds due to the fact that both $d_W^2(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})$ and $\arccos(\sum_{i,j}\gamma(i,j)\sqrt{\pi_{i,1}\pi_{j,2}})$ satisfy the triangle property. Moreover, the second inequality is due to the fact that the combined plan $\gamma(i,h)\cap\gamma(h,j)$ is not necessarily the optimal plan between $\boldsymbol{\mathcal{Y}}_1$ and $\boldsymbol{\mathcal{Y}}_2$, as the optimal plan $\gamma^*(i,j)$ achieves the minimum and defines the distance of $d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)$.

This completes the proof of Theorem 1.

Proof of Lemma 1

Recall that $\boldsymbol{\mathcal{Y}} \sim \sum_{i=1}^k \pi_i \mathcal{E}(\mathbf{x}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i, \mathcal{R})$. We can now write the definition of the Wasserstein distance in the form

$$d_W^2(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2) = \inf_{\eta(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)} \int_{m\times m} \eta(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)||\mathbf{x}_1 - \mathbf{x}_2||_2^2 d\mathbf{x}_1 d\mathbf{x}_2,$$
(19)

where $\eta(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)$ denotes the joint distribution between $\boldsymbol{\mathcal{Y}}_1$ and $\boldsymbol{\mathcal{Y}}_2$, and satisfies $\int_m \eta(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)d\mathbf{x}_1 = p_{\boldsymbol{\mathcal{Y}}_2}(\mathbf{x}_2)$ and $\int_m \eta(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)d\mathbf{x}_2 = p_{\boldsymbol{\mathcal{Y}}_1}(\mathbf{x}_1)$.

More importantly,

$$\sum_{i,j}\frac{\gamma^*(i,j)}{k}d_W^2(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2}) =$$
$$\sum_{i,j}\frac{\gamma^*(i,j)}{k}\inf_{\eta(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})}\int_{m\times m}\eta(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})||\mathbf{x}_1 - \mathbf{x}_2||_2^2 d\mathbf{x}_1 d\mathbf{x}_2$$
$$= \int_{m\times m}\sum_{i,j}\frac{\gamma^*(i,j)}{k}\eta^*(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})||\mathbf{x}_1 - \mathbf{x}_2||_2^2 d\mathbf{x}_1 d\mathbf{x}_2,$$
(20)

where $\eta^*(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})$ is the optimal plan that achieves the Wasserstein distance between $\boldsymbol{\mathcal{X}}_{i,1}$ and $\boldsymbol{\mathcal{X}}_{j,2}$. Therefore, by comparing (19) and (20), we can easily observe that $\sum_{i,j}\frac{\gamma^*(i,j)}{k}\eta^*(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})$ consists a subset of joint distribution of $\eta(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2)$, due to $\pi_{i,1} = \pi_{j,2} = 1/k$. In other words, because of the factorisation from $\boldsymbol{\mathcal{Y}}$ to $\boldsymbol{\mathcal{X}}$, the $\sum_{i,j}\frac{\gamma^*(i,j)}{k}\eta^*(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2})$ does not necessarily achieve the optimal transport plan between $\boldsymbol{\mathcal{Y}}_1$ and $\boldsymbol{\mathcal{Y}}_2$.

Moreover, for balanced EMMs, the following holds for arbitrary $\gamma(i,j)$,

$$\arccos(\sum_{i,j}\gamma(i,j)\sqrt{\pi_{i,1}\pi_{j,2}}) = 0.$$
(21)

Thus, we have

$$d_W^2(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2) \leq \sum_{i,j}\frac{\gamma^*(i,j)}{k}d_W^2(\boldsymbol{\mathcal{X}}_{i,1}, \boldsymbol{\mathcal{X}}_{j,2}) = d_U(\boldsymbol{\mathcal{Y}}_1, \boldsymbol{\mathcal{Y}}_2).$$
(22)

This completes the proof of Lemma 1.

Proof of Lemma 2

Due to the fact that $\gamma(i,j)$ in Definition 1 is a bijection between the mixture components in $\boldsymbol{\mathcal{Y}}_1$ and $\boldsymbol{\mathcal{Y}}_2$, for each $\boldsymbol{\pi}_1$, $\boldsymbol{\mu}_{i,1}$ and $\boldsymbol{\Sigma}_{i,1}$ in $\boldsymbol{\mathcal{Y}}_1$, there exist only one corresponding $\boldsymbol{\pi}_2$, $\boldsymbol{\mu}_{j,2}$ and $\boldsymbol{\Sigma}_{j,2}$, respectively, in $\boldsymbol{\mathcal{Y}}_2$. In other words, when casting the problem onto the parameter space, the Hessian of $d_U(\boldsymbol{\mathcal{Y}}_{\boldsymbol{\theta}}, \boldsymbol{\mathcal{Y}}_{\boldsymbol{\theta}+s\Delta\boldsymbol{\theta}})$ can be calculated as follows,

$$\frac{\partial^2 d_U(\boldsymbol{\mathcal{Y}}_{\boldsymbol{\theta}}, \boldsymbol{\mathcal{Y}}_{\boldsymbol{\theta}+s\Delta\boldsymbol{\theta}})}{\partial s^2}|_{s\to 0}$$
$$= \frac{1}{k}\sum_{i=1}^K\frac{\partial^2 d_W^2(\boldsymbol{\mathcal{X}}_{\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i}, \boldsymbol{\mathcal{X}}_{\boldsymbol{\mu}_i+s\Delta\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i+s\Delta\boldsymbol{\Sigma}_i})}{\partial s^2}|_{s\to 0}$$
(23)
$$+ \frac{\partial^2 \arccos(\sqrt{\boldsymbol{\pi}}^T(\sqrt{\boldsymbol{\pi}+s\Delta\boldsymbol{\pi}}))}{\partial s^2}|_{s\to 0},$$

where the square root operation $\sqrt{\cdot}$ is performed element-wise.

In (23), we can see that the manifold defined by $d_U(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}_{\boldsymbol{\theta}+\Delta\boldsymbol{\theta}})$ is a product manifold defined by $d_W^2(\mathcal{X}_{\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i}, \mathcal{X}_{\boldsymbol{\mu}_i+\Delta\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i+\Delta\boldsymbol{\Sigma}_i})$ and $\arccos(\sqrt{\boldsymbol{\pi}}^T(\sqrt{\boldsymbol{\pi}+\Delta\boldsymbol{\pi}}))$.

Furthermore, $\arccos(\sqrt{\boldsymbol{\pi}}^T(\sqrt{\boldsymbol{\pi}+\Delta\boldsymbol{\pi}}))$ defines a sphere manifold (Absil, Mahony, and Sepulchre 2009) (Examples 3.5.1 and 3.6.1), for which basic operations are provided in Examples 4.1.1, 5.4.1 and 8.1.7 of (Absil, Mahony, and Sepulchre 2009). On the other hand, the Wasserstein distance between two elliptical distributions has an explicit representation as follows (Ghaffari and Walker 2018; Muzellec and Cuturi 2018),

$$d_W^2(\mathcal{X}_{\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i}, \mathcal{X}_{\boldsymbol{\mu}_i+\Delta\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i+\Delta\boldsymbol{\Sigma}_i}) = ||\Delta\boldsymbol{\mu}_i||_2^2$$
$$+ \frac{\mathbb{E}[\mathcal{R}^2]}{m}\mathrm{tr}(\boldsymbol{\Sigma}_i + (\boldsymbol{\Sigma}_i+\Delta\boldsymbol{\Sigma}_i) - 2(\boldsymbol{\Sigma}_i^{1/2}(\boldsymbol{\Sigma}_i+\Delta\boldsymbol{\Sigma}_i)\boldsymbol{\Sigma}_i^{1/2})^{1/2}). \tag{24}$$

We can thus conclude that the manifold for $\boldsymbol{\mu}_i$ is the conventional Euclidean space within $\mathbb{R}^m$. Moreover, the manifold of $\boldsymbol{\Sigma}_i$ can also be correspondingly obtained on the basis of (Malagò, Montrucchio, and Pistone 2018).

This completes the proof of Lemma 2.

Calculations within Table 2

The basic operations on a sphere manifold were described in (Absil, Mahony, and Sepulchre 2009) (Examples 3.5.1, 3.6.1, 4.1.1, 5.4.1 and 8.1.7). As the manifold for $\boldsymbol{\mu}_i$ is the conventional Euclidean space, the trivial gradient descent can be employed. Operations for $\boldsymbol{\Sigma}_i$ under the Wasserstein manifold can be found in (Malagò, Montrucchio, and Pistone 2018). Furthermore, we omit the scale weight $\frac{\mathbb{E}[\mathcal{R}^2]}{m}$ in our work as it can be incorporated into the stepsize α during the gradient descent.

To calculate the Euclidean gradients, we first explicitly express our cost function, i.e., the sliced Wasserstein distance $d_{SW}(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}^*)$, as follows,

$$d_{SW}(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}^*) =$$
$$\int_m \inf_{\eta(\mathrm{Ra}(\mathcal{Y}_{\boldsymbol{\theta}},\mathbf{p}),\mathrm{Ra}(\mathcal{Y}^*,\mathbf{p}))} \int_{\mathbb{R}\times\mathbb{R}} \eta_R \cdot (y_1-y_2)^2 dy_1 dy_2 d\mathbf{p}, \tag{25}$$

and

$$\eta_R = \eta(\mathrm{Ra}(\mathcal{Y}_{\boldsymbol{\theta}},\mathbf{p}), \mathrm{Ra}(\mathcal{Y}^*,\mathbf{p})), \tag{26}$$

where y_1 denotes the Radon transform (Rabin et al. 2011) $\mathrm{Ra}(\mathcal{Y}_{\boldsymbol{\theta}},\mathbf{p})$, which projects $\mathcal{Y}_{\boldsymbol{\theta}}$ onto a one-dimensional random variable along the direction $\mathbf{p}$. Similarly, y_2 denotes the Radon transform of $\mathcal{Y}^*$ with the direction $\mathbf{p}$. Recall that $\eta(\cdot,\cdot)$ is a joint distribution. Due to the fact that elliptical distributions belong to the location-scale family, as described by the stochastic representation, the Radon transform of each $\mathcal{X}_{\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i}$ has a simple representation in the form $\mathcal{X}_{\mathbf{p}^T\boldsymbol{\mu}_i,\mathbf{p}^T\boldsymbol{\Sigma}_i\mathbf{p}} \sim \mathcal{E}(x; \mathbf{p}^T\boldsymbol{\mu}_i, \mathbf{p}^T\boldsymbol{\Sigma}_i\mathbf{p}, \mathcal{R})$, where $\mathcal{X}$ is a one-dimensional elliptical distribution.

Fortunately, the one-dimensional Wasserstein distance has the closed-form solution (Kolouri et al. 2017), so that

we can re-write (25) as

$$\min_{\boldsymbol{\pi},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i} d_{SW}(\mathcal{Y}_{\boldsymbol{\theta}}, \mathcal{Y}^*) =$$
$$\min_{\boldsymbol{\pi},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i} \int_m \int_{\mathbb{R}} |y-T(y)|^2 \sum_{i=1}^k \frac{c_m\pi_i}{(\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p})^{\frac{1}{2}}} g\left(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}}\right) dy d\mathbf{p}, \tag{27}$$

where $T(y)$ is the optimal transport plan between $\mathrm{Ra}(\mathcal{Y}_{\boldsymbol{\theta}},\mathbf{p})$ and $\mathrm{Ra}(\mathcal{Y}^*,\mathbf{p})$, which can be explicitly obtained via their cumulative distribution functions (Chen and Li 2018). When optimising (27) using stochastic gradient descent, for each $\mathbf{p}$, we therefore minimise

$$\min_{\boldsymbol{\pi},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i} J(\boldsymbol{\theta}, \mathbf{p}) =$$
$$\min_{\boldsymbol{\pi},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i} \int_{\mathbb{R}} |y - T(y)|^2 \sum_{i=1}^k \frac{c_m\pi_i}{(\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p})^{\frac{1}{2}}} g\left(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}}\right) dy. \tag{28}$$

Then, by calculating the derivatives with regard to $\sqrt{\pi_i}$, $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$, and upon introducing the Kantorovich potential $\phi(y)$ (Chen and Li 2018), we obtain the Euclidean gradients as follows,

$$\nabla_{\sqrt{\pi_i},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i} J(\boldsymbol{\theta}, \mathbf{p}) =$$
$$\int_{\mathbb{R}} \phi(y) \nabla_{\sqrt{\pi_i},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i} \left(\sum_{i=1}^k \frac{c_m\sqrt{\pi_i}^2}{(\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p})^{\frac{1}{2}}} g\left(\frac{(y-\mathbf{p}^T\boldsymbol{\mu}_i)^2}{\mathbf{p}^T\boldsymbol{\Sigma}_i^h\mathbf{p}}\right)\right) dy, \tag{29}$$

where $\nabla_{\sqrt{\pi_i},\boldsymbol{\mu}_i,\boldsymbol{\Sigma}_i}(\cdot)$ denotes the derivatives with regard to $\{\sqrt{\pi_i}, \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i\}$, respectively. The Euclidean gradients in the Table 2 can then be easily calculated.

References

[Absil, Mahony, and Sepulchre 2009] Absil, P.-A.; Mahony, R.; and Sepulchre, R. 2009. *Optimization algorithms on matrix manifolds*. Princeton University Press.

[Amari 1998] Amari, S.-I. 1998. Natural gradient works efficiently in learning. *Neural Computation* 10(2):251–276.

[Andrews and Mallows 1974] Andrews, D. F., and Mallows, C. L. 1974. Scale mixtures of normal distributions. *Journal of the Royal Statistical Society. Series B (Methodological)* 99–102.

[Arbelaez et al. 2011] Arbelaez, P.; Maire, M.; Fowlkes, C.; and Malik, J. 2011. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 33(5):898–916.

[Arjovsky, Chintala, and Bottou 2017] Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein Gan. *arXiv preprint arXiv:1701.07875*.

[Barndorff-Nielsen, Kent, and Sørensen 1982] Barndorff-Nielsen, O.; Kent, J.; and Sørensen, M. 1982. Normal variance-mean mixtures and z distributions. *International Statistical Review/Revue Internationale de Statistique* 145–159.

[Becigneul and Ganea 2019] Becigneul, G., and Ganea, O.-E. 2019. Riemannian adaptive optimization methods. In *Proceedings of the International Conference on Learning Representations*.

[Browne and McNicholas 2015] Browne, R. P., and McNicholas, P. D. 2015. A mixture of generalized hyperbolic distributions. *Canadian Journal of Statistics* 43(2):176–198.

[Chen and Li 2018] Chen, Y., and Li, W. 2018. Natural gradient in Wasserstein statistical manifold. *arXiv preprint arXiv:1805.08380*.

[Chen, Georgiou, and Tannenbaum 2018] Chen, Y.; Georgiou, T. T.; and Tannenbaum, A. 2018. Optimal transport for Gaussian mixture models. *IEEE Access* 7:6269–6278.

[Dasgupta 1999] Dasgupta, S. 1999. Learning mixtures of Gaussians. In *Proceedings of the 40th IEEE Annual Symposium on Foundations of Computer science.*, 634–644. IEEE.

[Fang, Kotz, and Ng 1990] Fang, K. W.; Kotz, S.; and Ng, K. W. 1990. *Symmetric multivariate and related distributions*. London, U.K.: Chapman & Hall.

[Ghaffari and Walker 2018] Ghaffari, N., and Walker, S. 2018. On multivariate optimal transportation. *arXiv preprint arXiv:1801.03516*.

[Holzmann, Munk, and Gneiting 2006] Holzmann, H.; Munk, A.; and Gneiting, T. 2006. Identifiability of finite mixtures of elliptical distributions. *Scandinavian Journal of Statistics* 33(4):753–763.

[Hosseini and Sra 2015] Hosseini, R., and Sra, S. 2015. Matrix manifold optimization for Gaussian mixtures. In *Advances in Neural Information Processing Systems*, 910–918.

[Jin et al. 2016] Jin, C.; Zhang, Y.; Balakrishnan, S.; Wainwright, M. J.; and Jordan, M. I. 2016. Local maxima in the likelihood of Gaussian mixture models: Structural results and algorithmic consequences. In *Advances in Neural Information Processing Systems*, 4116–4124.

[Jordan and Jacobs 1994] Jordan, M. I., and Jacobs, R. A. 1994. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation* 6(2):181–214.

[Kent and Tyler 1991] Kent, J. T., and Tyler, D. E. 1991. Redescending M-estimates of multivariate location and scatter. *The Annals of Statistics* 2102–2119.

[Kingma and Ba 2014] Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[Knott and Smith 1984] Knott, M., and Smith, C. S. 1984. On the optimal mapping of distributions. *Journal of Optimization Theory and Applications* 43(1):39–49.

[Kolouri et al. 2017] Kolouri, S.; Park, S. R.; Thorpe, M.; Slepcev, D.; and Rohde, G. K. 2017. Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine* 34(4):43–59.

[Kolouri, Rohde, and Hoffmann 2018] Kolouri, S.; Rohde, G. K.; and Hoffmann, H. 2018. Sliced Wasserstein distance for learning Gaussian mixture models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3427–3436.

[LeCun, Cortes, and Burges 2010] LeCun, Y.; Cortes, C.; and Burges, C. 2010. Mnist handwritten digit database. *ATT Labs [Online]. Available: http://yann. lecun. com/exdb/mnist* 2.

[Lévy 2015] Lévy, B. 2015. A numerical algorithm for L2 semi-discrete optimal transport in 3d. *ESAIM: Mathematical Modelling and Numerical Analysis* 49(6):1693–1715.

[Malagò, Montrucchio, and Pistone 2018] Malagò, L.; Montrucchio, L.; and Pistone, G. 2018. Wasserstein Riemannian geometry of Gaussian densities. *Information Geometry* 1(2):137–179.

[Martens 2014] Martens, J. 2014. New insights and perspectives on the natural gradient method. *arXiv preprint arXiv:1412.1193*.

[Monge 1781] Monge, G. 1781. *Mémoire sur la théorie des déblais et des remblais*. De l'Imprimerie Royale.

[Muzellec and Cuturi 2018] Muzellec, B., and Cuturi, M. 2018. Generalizing point embeddings using the Wasserstein space of elliptical distributions. In *Advances in Neural Information Processing Systems*, 10237–10248.

[Naim and Gildea 2012] Naim, I., and Gildea, D. 2012. Convergence of the EM algorithm for Gaussian mixtures with unbalanced mixing coefficients. *arXiv preprint arXiv:1206.6427*.

[Ollila et al. 2012] Ollila, E.; Tyler, D. E.; Koivunen, V.; and Poor, H. V. 2012. Complex elliptically symmetric distributions: Survey, new results and applications. *IEEE Transactions on Signal Processing* 60(11):5597–5625.

[Peel and McLachlan 2000] Peel, D., and McLachlan, G. J. 2000. Robust mixture modelling using the t distribution. *Statistics and Computing* 10(4):339–348.

[Rabin et al. 2011] Rabin, J.; Peyré, G.; Delon, J.; and Bernot, M. 2011. Wasserstein barycenter and its application to texture mixing. In *Proceedings of the International Conference on Scale Space and Variational Methods in Computer Vision*, 435–446. Springer.

[Redner and Walker 1984] Redner, R. A., and Walker, H. F. 1984. Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review* 26(2):195–239.

[Sra and Hosseini 2013] Sra, S., and Hosseini, R. 2013. Geometric optimisation on positive definite matrices for elliptically contoured distributions. In *Advances in Neural Information Processing Systems*, 2562–2570.

[Sra and Hosseini 2015] Sra, S., and Hosseini, R. 2015. Conic geometric optimization on the manifold of positive definite matrices. *SIAM Journal on Optimization* 25(1):713–739.

[Stefanski 1991] Stefanski, L. A. 1991. A normal scale mixture representation of the logistic distribution. *Statistics & Probability Letters* 11(1):69–70.

[Sturm and others 2006] Sturm, K.-T., et al. 2006. On the geometry of metric measure spaces. *Acta Mathematica* 196(1):65–131.

[Takatsu and others 2011] Takatsu, A., et al. 2011. Wasserstein geometry of Gaussian measures. *Osaka Journal of Mathematics* 48(4):1005–1026.

[Tan and Jiao 2007] Tan, S., and Jiao, L. 2007. Multivariate statistical models for image denoising in the wavelet domain. *International Journal of Computer Vision* 75(2):209–230.

[Wang et al. 2004] Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 13(4):600–612.

[Wang et al. 2015] Wang, W.; Wang, R.; Huang, Z.; Shan, S.; and Chen, X. 2015. Discriminant analysis on Riemannian manifold of Gaussian distributions for face recognition with image sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2048–2057.

[Xu and Jordan 1996] Xu, L., and Jordan, M. I. 1996. On convergence properties of the EM algorithm for gaussian mixtures. *Neural Computation* 8(1):129–151.

[Xu, Hsu, and Maleki 2016] Xu, J.; Hsu, D. J.; and Maleki, A. 2016. Global analysis of expectation maximization for mixtures of two Gaussians. In *Advances in Neural Information Processing Systems*, 2676–2684.

[Zhang and Sra 2016] Zhang, H., and Sra, S. 2016. First-order methods for geodesically convex optimization. In *Proceedings of the Conference on Learning Theory*, 1617–1638.

[Zhang, Wiesel, and Greco 2013] Zhang, T.; Wiesel, A.; and Greco, M. S. 2013. Multivariate generalized Gaussian distribution: Convexity and graphical models. *IEEE Transactions on Signal Processing* 61(16):4141–4148.