A Riemannian Newton trust-region method for fitting Gaussian mixture models

Lena Sembach¹ · Jan Pablo Burgard¹ · Volker Schulz¹

Received: 3 May 2021 / Accepted: 14 November 2021 / Published online: 17 December 2021
 © The Author(s) 2021

Abstract

Gaussian Mixture Models are a powerful tool in Data Science and Statistics that are mainly used for clustering and density approximation. The task of estimating the model parameters is in practice often solved by the expectation maximization (EM) algorithm which has its benefits in its simplicity and low per-iteration costs. However, the EM converges slowly if there is a large share of hidden information or overlapping clusters. Recent advances in Manifold Optimization for Gaussian Mixture Models have gained increasing interest. We introduce an explicit formula for the Riemannian Hessian for Gaussian Mixture Models. On top, we propose a new Riemannian Newton Trust-Region method which outperforms current approaches both in terms of runtime and number of iterations. We apply our method on clustering problems and density approximation tasks. Our method is very powerful for data with a large share of hidden information compared to existing methods.

Keywords Gaussian mixture models · Clustering · Density estimation · Manifold optimization · Trust-region methods

1 Introduction

Gaussian Mixture Models are widely recognized in Data Science and Statistics. The fact that any probability density can be approximated by a Gaussian Mixture Model with a sufficient number of components makes it an attractive tool in statistics. However, this comes with some computational limitations, where some of them are described in Ormoneit and Tresp (1995), Lee and McLachlan (2013) and Coretto (2021). Nevertheless, we here focus on the benefits of Gaussian mixture models. Besides the goal of density approximation, the possibility of modeling latent features by the underlying components make it also a strong tool for soft clustering tasks. Typical applications are to be found in the area of image analysis (Alfö et al. 2008; Dresvyanskiy et al. 2020; Zoran and Weiss 2012), pattern recognition (Wu et al. 2012; Bishop 2006), econometrics (Articus and Burgard 2014; Compiani and Kitamura 2016 and many others).

We state the Gaussian Mixture Model in the following:

Let $K \in \mathbb{N}$ be given. The Gaussian Mixture Model (with K components) is given by the (multivariate) density function p of the form

$$p(x) = \sum_{j=1}^K \alpha_j p_{\mathcal{N}}(x; \mu_j, \Sigma_j), \quad x \in \mathbb{R}^d, \quad (1)$$

with positive mixture components α_j that sum up to 1 and Gaussian density functions $p_{\mathcal{N}}$ with means $\mu_j \in \mathbb{R}^d$ and covariance matrices $\Sigma_j \in \mathbb{R}^{d \times d}$. In order to have a well-defined expression, we impose $\Sigma_j > 0$, i.e. the Σ_j are symmetric positive definite.

Given observations $x_1, \dots, x_m$, the goal of parameter estimation for Gaussian Mixture Models consists in maximizing the log-likelihood. This yields the optimization problem

$$\max_{\substack{\alpha \in \Delta_K \\ \mu_j \in \mathbb{R}^d \\ \Sigma_j > 0}} \sum_{i=1}^m \log \left(\sum_{j=1}^K \alpha_j p_{\mathcal{N}}(x_i; \mu_j, \Sigma_j) \right), \quad (2)$$

where

$$\Delta_K = \left\{ (\alpha_1, \dots, \alpha_K), \alpha_j \in \mathbb{R}^+ \forall j, \sum_{j=1}^K \alpha_j = 1 \right\}$$

✉ Lena Sembach
 sembach@uni-trier.de

Jan Pablo Burgard
 burgardj@uni-trier.de

Volker Schulz
 volker.schulz@uni-trier.de

¹ Universitaet Trier, 54296 Trier, Germany

is the K -dimensional probability simplex and the covariance matrices Σ_j are restricted to the set of positive definite matrices.

In practice, this problem is commonly solved by the Expectation Maximization (EM) algorithm. It is known that the Expectation Maximization algorithm converges fast if the K clusters are well separated. Ma et al. (2000) showed that in such a case, the convergence rate is superlinear. However, Expectation Maximization has its speed limits for highly overlapping clusters, where the latent variables have a non-negligible large probability among more than one cluster. In such a case, the convergence is linear (Xu and Jordan 1996) which might result in slow parameter estimation despite very low per-iteration costs.

From a nonlinear optimization perspective, the problem in (2) can be seen as a constrained nonlinear optimization problem. However, the positive definiteness constraint of the covariance matrices Σ_j is a challenge for applying standard nonlinear optimization algorithms. While this constraint is naturally fulfilled in the EM algorithm, we cannot simply drop it as we might leave the parameter space in alternative methods. Approaches to this problem like introducing a Cholesky decomposition (Salakhutdinov et al. 2003; Xu and Jordan 1996), or using interior point methods via smooth convex inequalities (Vanderbei and Benson 1999) can be applied and one might hope for faster convergence with Newton-type algorithms. Methods like using a Conjugate Gradient Algorithm (or a combination of both EM and CG) led to faster convergence for highly overlapping clusters (Salakhutdinov et al. 2003). However, this induces an additional numeric overhead by imposing the positive-definiteness constraint by a Cholesky decomposition.

In recent approaches Hosseini and Sra (2015, 2020) suggest to exploit the geometric structure of the set of positive definite matrices. As an open set in the set of symmetric matrices, it admits a manifold structure (Bhatia 2007) and thus the concepts of Riemannian optimization can be applied. The concept of Riemannian optimization, i.e. optimizing over parameters that live on a smooth manifold is well studied and has gained increasing interest in the domain of Data Science, for example for tensor completion problems (Heidel and Schulz 2018). However, the idea is quite new for Gaussian Mixture Models and Hosseini and Sra (2015, 2020) showed promising results with a Riemannian LBFGS and Riemannian Stochastic Gradient Descent Algorithm. The results in Hosseini and Sra (2015, 2020) are based on a reformulation of the log-likelihood in (2) that turns out to be very efficient in terms of runtime. By design, the algorithms investigated in Hosseini and Sra (2015, 2020) do not use exact second-order information of the objective function. Driven by the quadratic local convergence of the Riemannian Newton method, we thus might hope for faster algorithms with the availability of the Riemannian Hessian. In the present work, we derive a

formula for the Riemannian Hessian of the reformulated log-likelihood and suggest a Riemannian Newton Trust-Region method for parameter estimation of Gaussian Mixture Models.

An arXiv preprint of this article is available (Sembach et al. 2021).

The paper is organized as follows. In Sect. 2, we introduce the reader to the concepts of Riemannian optimization and the Riemannian setting of the reformulated log-likelihood for Gaussian Mixture Models. In particular, we derive the expression for the Riemannian Hessian in Sect. 2.3 which is a big contribution to richer Riemannian Algorithm methods for Gaussian Mixture Models. In Sect. 3 we present the Riemannian Newton Trust-Region Method and prove the convergence to global optima and superlinear local convergence for our problem. We compare our proposed method against existing algorithms both on artificial and real world data sets in Sect. 4 for the task of clustering and density approximation.

2 Riemannian setting for Gaussian mixture models

We will build the foundations of Riemannian Optimization in the following to specify the characteristics for Gaussian Mixture Models afterwards. In particular, we introduce a formula for the Riemannian Hessian for the reformulated problem which is the basis for second-order optimization algorithms.

2.1 Riemannian optimization

To construct Riemannian Optimization methods, we briefly state the main concepts of Optimization on Manifolds or Riemannian Optimization. A good introduction is Absil et al. (2008) and Boumal (2020), we here follow the notations of Absil et al. (2008). The concepts of Riemannian Optimization are based on concepts from unconstrained euclidean optimization algorithms and are generalized to (possibly non-linear) manifolds.

A manifold is a space that locally resembles Euclidean space, meaning that we can locally map points on manifolds to $\mathbb{R}^n$ via bicontinuous mappings. Here, n denotes the dimension of the manifold. In order to define a generalization of differentials, Riemannian optimization methods require smooth manifolds meaning that the transition mappings are smooth functions. As manifolds are in general not vector spaces, standard optimization algorithms like line-search methods cannot be directly applied as the iterates might leave the admissible set. Instead, one moves along tangent vectors in tangent spaces $T_\theta \mathcal{M}$, local approximations of a point θ on the manifold, i.e. $\theta \in \mathcal{M}$. Tangent spaces are basically first-order approximations of the manifold at

Fig. 1 Retraction-based Riemannian optimization

specific points and the *tangent bundle* $T\mathcal{M}$ is the disjoint union of the tangent spaces $T_\theta\mathcal{M}$. In Riemannian manifolds, each of the tangent spaces $T_\theta\mathcal{M}$ for $\theta \in \mathcal{M}$ is endowed with an inner product $\langle \cdot, \cdot \rangle_\theta$ that varies smoothly with θ . The inner product is essential for Riemannian optimization methods as it admits some notion of length associated with the manifold. The optimization methods also require some local pull-back from the tangent spaces $T_\theta\mathcal{M}$ to the manifold $\mathcal{M}$ which can be interpreted as moving along a specific curve on $\mathcal{M}$ (dotted curve in Fig. 1). This is realized by the concept of *retractions*: Retractions are mappings from the tangent bundle $T\mathcal{M}$ to the manifold $\mathcal{M}$ with rigidity conditions: we move through the zero element 0_θ with velocity $\xi_\theta \in T_\theta\mathcal{M}$, i.e. $DR_\theta(0_\theta)[\xi_\theta] = \xi_\theta$. Furthermore, the retraction of $0_\theta \in T_\theta\mathcal{M}$ at θ is θ itself (see Fig. 1).

Roughly spoken, a step of a Riemannian optimization algorithm works as follows:

- At iterate θ^t , take a new step ξ_{θ^t} on the tangent space $T_{\theta^t}\mathcal{M}$
- Pull back the new step to the manifold by applying the retraction at point θ^t by setting $\theta^{t+1} = R_{\theta^t}(\xi_{\theta^t})$

Here, the crucial part that has an impact on convergence speed is updating the new iterate on the tangent space, just like in the Euclidean case. As Riemannian optimization algorithms are a generalization of Euclidean unconstrained optimization algorithms, we thus introduce a generalization of the gradient and the Hessian.

Riemannian Gradient. In order to characterize Riemannian gradients, we need a notion of differential of functions defined on manifolds.

The *differential* of $f : \mathcal{M} \rightarrow \mathbb{R}$ at θ is the linear operator $Df(\theta) : T\mathcal{M}_\theta \rightarrow \mathbb{R}$ defined by:

$$Df(\theta)[v] = \left. \frac{d}{dt} f(c(t)) \right|_{t=0},$$

where $c : I \rightarrow \mathcal{M}$, $0 \in I \subset \mathbb{R}$ is a smooth curve on $\mathcal{M}$ with $c'(0) = v$.

The Riemannian gradient can be uniquely characterized by the differential of the function f and the inner product associated with the manifold:

The *Riemannian gradient* of a smooth function $f : \mathcal{M} \rightarrow \mathbb{R}$ on a Riemannian manifold is a mapping $\text{grad } f : \mathcal{M} \rightarrow$

$T\mathcal{M}$ such that, for all $\theta \in \mathcal{M}$, $\text{grad } f(\theta)$ is the unique tangent vector in $T_\theta\mathcal{M}$ satisfying

$$\langle \text{grad } f(\theta), \xi_\theta \rangle_\theta = Df(\theta)[\xi_\theta] \quad \forall \xi_\theta \in T_\theta\mathcal{M}.$$

Riemannian Hessian. Just like we defined Riemannian gradients, we can also generalize the Hessian to its Riemannian version. To do this, we need a tool to differentiate along tangent spaces, namely the Riemannian connection (for details see Absil et al. 2008, Section 5.3).

The *Riemannian Hessian* of $f : \mathcal{M} \rightarrow \mathbb{R}$ at θ is the linear operator $\text{Hess } f(\theta) : T_\theta\mathcal{M} \rightarrow T_\theta\mathcal{M}$ defined by

$$\text{Hess } f(\theta)[\xi_\theta] = \nabla_{\xi_\theta} \text{grad } f(\theta),$$

where ∇ is the Riemannian connection with respect to the Riemannian manifold.

2.2 Reformulation of the log-likelihood

In Hosseini and Sra (2015), the authors experimentally showed that applying the concepts of Riemannian optimization to the objective in (2) cannot compete with Expectation Maximization. This can be mainly led back to the fact that the maximization in the M-step of EM, i.e. the maximization of the log-likelihood for a single Gaussian, is a concave problem and thus very easy to solve - it even admits a closed-form solution. However, when considering Riemannian optimization for (2), the maximization of the log-likelihood of a single Gaussian is not geodesically concave (concavity along the shortest curve connecting two points on a manifold). The following reformulation introduced by Hosseini and Sra (2015) removes this geometric mismatch and results in a speed-up for the Riemannian algorithms.

We augment the observations x_i by introducing the observations $y_i = (x_i, 1)^T \in \mathbb{R}^{d+1}$ for $i = 1, \dots, m$ and consider the optimization problem

$$\hat{\mathcal{L}}(\theta) = \sum_{i=1}^m \log \left(\sum_{j=1}^K h^i(\theta_j) \right), \quad (3)$$

where

$$\begin{aligned} h^i(\theta_j) &= \frac{\exp(\eta_j)}{\sum_{k=1}^K \exp(\eta_k)} \frac{\exp\left(\frac{1}{2} \left(1 - y_i^T S_j^{-1} y_i\right)\right)}{\sqrt{(2\pi)^d \det(S_j)}} \\ &= \frac{\exp(\eta_j)}{\sum_{k=1}^K \exp(\eta_k)} q_{\mathcal{N}}(y_i; S_j), \end{aligned} \quad (4)$$

with $q_{\mathcal{N}}(y_i; S_j) = \sqrt{2\pi} \exp(\frac{1}{2}) p_{\mathcal{N}}(y_i; 0, S_j)$ for parameters $\theta_j = (S_j, \eta_j)$, $j = 1, \dots, K - 1$ and $\theta_K = (S_K, 0)$.

This means that instead of considering Gaussians of d -dimensional variables x , we now consider Gaussians of $d + 1$ -dimensional variables y with zero mean and covariance matrices S_j . The reformulation leads to faster Riemannian algorithms, as it has been shown in Hosseini and Sra (2015) that the maximization of a single Gaussian

$$\max_{S > 0} \sum_{i=1}^m \log q_{\mathcal{N}}(y_i; S)$$

is geodesically concave.

Furthermore, this reformulation is faithful, as the original problem (2) and the reformulated problem (3) are equivalent in the followings sense:

Theorem 1 (Hosseini and Sra 2015, Theorem 2.2) *A local maximum of the reformulated GMM log-likelihood $\hat{\mathcal{L}}(\theta)$ with maximizer $\theta^* = (\theta_1^*, \dots, \theta_K^*)$, $\theta_j^* = (S_j^*, \eta_j^*)$ is a local maximum of the original log-likelihood $\mathcal{L}((\alpha_j, \mu_j, \Sigma_j)_j)$ with maximizer $(\alpha_j^*, \mu_j^*, \Sigma_j^*)_j$. Here, $\mathcal{L}$ denotes the objective in the problem (2).*

The relationship of the maximizers is given by

$$S_j^* = \begin{pmatrix} \Sigma_j^* + \mu_j^* \mu_j^{*T} & \mu_j^* \\ \mu_j^{*T} & 1 \end{pmatrix}, \quad (5)$$

$$\eta_j^* = \log \left(\frac{\alpha_j^*}{\alpha_K^*} \right) \quad j = 1, \dots, K-1; \quad \eta_K \equiv 0. \quad (6)$$

This means that instead of solving the original optimization problem (2) we can easily solve the reformulated problem (3) on its according parameter space and transform the optima back by the relationships (5) and (6).

Penalizing the objective. When applying Riemannian optimization algorithms on the reformulated problem (3), covariance singularity is a challenge. Although this is not observed in many cases in practice, it might result in unstable algorithms. This is due to the fact that the objective in (3) is unbounded from above, see ‘‘Appendix A’’ for details. The same problem is extensively studied for the original problem (2) and makes convergence theory hard to investigate. An alternative consists in considering the maximum a posterior log-likelihood for the objective. If conjugate priors are used for the variables μ_j, Σ_j , the optimization problem remains structurally unchanged and results in a bounded objective (see Snoussi and Mohammad-Djafari 2002). Adapted versions of Expectation Maximization have been proposed in the literature and are often applied in practice.

A similar approach has been proposed in Hosseini and Sra (2020) where the reformulated objective (3) is penalized by an additive term that consists of the logarithm of the Wishart

prior, i.e. we penalize each of the K components with

$$\psi(S_j, \Psi) = -\frac{\rho}{2} \log \det(S_j) - \frac{\beta}{2} \text{tr}(\Psi S_j^{-1}), \quad (7)$$

where Ψ is the block matrix

$$\Psi = \begin{pmatrix} \frac{\gamma}{\beta} \Lambda + \kappa \lambda \lambda^T & \kappa \lambda \\ \kappa \lambda^T & \kappa \end{pmatrix}$$

for $\lambda \in \mathbb{R}^d$, $\Lambda \in \mathbb{R}^{d \times d}$ and $\gamma, \beta, \rho, \nu, \kappa \in \mathbb{R}$. If we assume that $\rho = \gamma(d + \nu + 1) + \beta$, the results of Theorem 1 are still valid for the penalized version, see Hosseini and Sra (2020). Besides, the authors introduce an additive term to penalize very tiny clusters by introducing Dirichlet priors for the mixing coefficients $\alpha_j = \frac{\exp(\eta_j)}{\sum_{k=1}^K \exp(\eta_k)}$, i.e.

$$\varphi(\eta, \zeta) = \zeta \left(\sum_{j=1}^K \eta_j - K \log \left(\sum_{k=1}^K \exp(\eta_k) \right) \right). \quad (8)$$

In total, the penalized problem is given by

$$\max_{\theta} \hat{\mathcal{L}}_{\text{pen}}(\theta; \Psi, \zeta) = \sum_{i=1}^m \log \left(\sum_{j=1}^K h^i(\theta_j) \right) + \text{Pen}(\theta), \quad (9)$$

where

$$\text{Pen}(\theta) = \sum_{j=1}^K \psi(S_j, \Psi) + \varphi(\eta, \zeta).$$

The use of such an additive penalizer leads to a bounded objective:

Theorem 2 *The penalized optimization problem in (9) is bounded from above.*

Proof We follow the proof for the original objective (2) from Snoussi and Mohammad-Djafari (2002). The penalized objective reads

$$\hat{\mathcal{L}}_{\text{pen}}(\theta; \Psi, \zeta) = \sum_{i=1}^m \log \left(\sum_{j=1}^K q_{\mathcal{N}}^{\text{AP}}(\theta_j, S_j; \Psi, \zeta) \right),$$

where

$$q_{\mathcal{N}}^{\text{AP}}(\theta_j, S_j; \Psi, \zeta) = h^i(\theta_j) \left(\prod_{k=1}^K \det(S_k)^{-\frac{\rho}{2}} \exp \left(-\frac{1}{2} \text{tr}(S_k^{-1} \Psi) \right) \alpha_k^{\zeta} \right)^{1/m}.$$

We get the upper bound

$$\begin{aligned} q_{\mathcal{N}}^{\text{AP}}(\theta_j, S_j; \Psi, \zeta) &\leq h^i(\theta_j) \prod_{k=1}^K \det(S_k)^{-\frac{\rho}{2}} \exp\left(-\frac{1}{2} \text{tr}(S_k^{-1} \Psi)\right) \alpha_k^\zeta \\ &\leq a \alpha_j \det(S_j)^{-\frac{d}{2}} \prod_{k=1}^K \det(S_k)^{-\frac{\rho}{2}} \exp\left(-\frac{1}{2} \text{tr}(S_k^{-1} \Psi)\right) \\ &= a \alpha_j (\det(S_j))^{-\frac{d+\rho}{2}} \exp\left(-\frac{1}{2} \text{tr}(S_j^{-1} \Psi)\right) \\ &\quad \times \prod_{\substack{k=1 \\ k \neq j}}^K \det(S_k)^{-\frac{\rho}{2}} \exp\left(-\frac{1}{2} \text{tr}(S_k^{-1} \Psi)\right), \quad (10) \end{aligned}$$

where we applied Bernoulli's inequality in the first inequality and used the positive definiteness of S_j in the second inequality. a is a positive constant independent of S_j and S_k .

By applying the relationship $\det(A)^{1/n} \leq \frac{1}{n} \text{tr}(A)$ for $A \in \mathbb{R}^{n \times n}$ by the inequality of arithmetic and geometric means, we get for the right hand side of (10)

$$\begin{aligned} \det(S_k)^{-\frac{b}{2}} \exp\left(-\frac{1}{2} \text{tr}(S_k^{-1} \Psi)\right) &\leq (\det(S_k))^{-\frac{b}{2}} \exp\left(-\frac{d+1}{2} \left(\frac{\det(\Psi)}{\det(S_k)}\right)^{\frac{1}{d+1}}\right) \end{aligned} \quad (11)$$

for a constant $b > 0$. The crucial part on the right side of (11) is when one of the S_k approaches a singular matrix and thus the determinant approaches zero. Then, we reach the boundary of the parameter space. We study this issue in further detail:

Without loss of generality, let $k = 1$ be the component where this occurs. Let S_1^* be a singular semipositive definite matrix of rank $r < d + 1$. Then, there exists a decomposition of the form

$$S_1^* = U^T D U,$$

where $D = \text{diag}(0, \dots, 0, \lambda_{d-r}, \lambda_{d-r+1}, \dots, \lambda_{d+1})$, $\lambda_l > 0$ for $l = d - r, \dots, d + 1$ and U an orthogonal square matrix of size $d + 1$. Now consider the sequence $S_1^{(n)}$ given by

$$S_1^{(n)} = U^T D^{(n)} U, \quad (12)$$

where

$$D^{(n)} = \text{diag}(\lambda_1^{(n)}, \dots, \lambda_{d-r-1}^{(n)}, \lambda_{d-r}, \lambda_{d-r+1}, \dots, \lambda_{d+1})$$

with $(\lambda_l^{(n)})_{l=1, \dots, d-r-1}$ converging to 0 as $n \rightarrow \infty$. Then, the matrix $S_1^{(n)}$ converges to S_1^* . Setting $\lambda^{(n)} = \prod_1^{d-r-1} \lambda_l^{(n)}$ and $\lambda^+ = \prod_{d-r}^{d+1} \lambda_l$, the right side of (11) reads

$$(\lambda^{(n)} \lambda^+)^{-\frac{b}{2}} \exp\left(-\frac{d+1}{2} \left(\frac{\det(\Psi)}{\lambda^{(n)} \lambda^+}\right)^{\frac{1}{d+1}}\right),$$

which converges to 0 as $n \rightarrow \infty$ by the rule of Hôpital. $\square$

With Theorem 2, we are able to study the convergence theory of the reformulated problem (3) in Sect. 3.

2.3 Riemannian characteristics of the reformulated problem

To solve the reformulated problem (3) or the penalized reformulated problem (9), we specify the Riemannian characteristics of the optimization problem. It is an optimization problem over the product manifold

$$\mathcal{M} = (\mathbb{P}^{d+1})^K \times \mathbb{R}^{K-1}, \quad (13)$$

where $\mathbb{P}^{d+1}$ is the set of strictly positive definite matrices of dimension $d + 1$. The set of symmetric matrices is tangent to the set of positive definite matrices as $\mathbb{P}^{d+1}$ is an open subset of it. Thus the tangent space of the manifold (13) is given by

$$T_\theta \mathcal{M} = (\mathbb{S}^{d+1})^K \times \mathbb{R}^{K-1}, \quad (14)$$

where $\mathbb{S}^{d+1}$ is the set of symmetric matrices of dimension $d + 1$. The inner product that is commonly associated with the manifold of positive definite matrices is the *intrinsic inner product*

$$\langle \xi_S, \chi_S \rangle_S = \text{tr}(S^{-1} \xi_S S^{-1} \chi_S), \quad (15)$$

where $S \in \mathbb{P}^{d+1}$ and $\xi_S, \chi_S \in \mathbb{S}^{d+1}$. The inner product defined on the tangent space (14) is the sum over all component-wise inner products and reads

$$\langle \xi_\theta, \chi_\theta \rangle_\theta = \sum_{j=1}^K \text{tr}(S_j^{-1} \xi_{S_j} S_j^{-1} \chi_{S_j}) + \xi_\eta^T \chi_\eta, \quad (16)$$

where

$$\begin{aligned} \theta &= ((S_1, \dots, S_K), \eta) && \in \mathcal{M}, \\ \xi_\theta &= ((\xi_{S_1}, \dots, \xi_{S_K}), \xi_\eta) && \in T_\theta \mathcal{M}, \\ \chi_\theta &= ((\chi_{S_1}, \dots, \chi_{S_K}), \chi_\eta) && \in T_\theta \mathcal{M}. \end{aligned}$$

The retraction we use is the exponential map on the manifold given by

$$R_\theta(\xi) = \left(\begin{array}{c} (S_j \exp(S_j^{-1} \xi S_j))_{j=1, \dots, K} \\ (\eta_j + \xi_{\eta_j})_{j=1, \dots, K-1} \end{array} \right), \quad (17)$$

see Jeuris et al. (2012).

Riemannian Gradient and Hessian. We now specify the Riemannian Gradient and the Riemannian Hessian in order to apply second-order methods on the manifold. The Riemannian Hessian in Theorem 4 is novel for the problem of fitting Gaussian Mixture Models and provides a way of making second-order methods applicable.

Theorem 3 *The Riemannian gradient of the reformulated problem reads $\text{grad } \hat{\mathcal{L}}(\theta) = (\chi_S, \chi_\eta)$, with*

$$\chi_S = \left(\frac{1}{2} \sum_{i=1}^m f_l^i (y_i y_i^T - S_l) \right)_{l=1, \dots, K},$$

$$\chi_\eta = \left(\sum_{i=1}^m f_r^i - \alpha_r \sum_{j=1}^K f_j^i \right)_{r=1, \dots, K-1},$$

where

$$f_l^i = \frac{h^i(\theta_l)}{\sum_{j=1}^K h^i(\theta_j)},$$

$$\alpha_r = \frac{\exp(\eta_r)}{\sum_{k=1}^K \exp(\eta_k)}.$$

The additive terms for the penalizers in (7), (8) are given by

$$\text{grad Pen}(\theta) = \left(\begin{array}{c} (-\frac{1}{2}(\rho S_l - \beta \Psi))_{l=1, \dots, K} \\ \zeta(1 - K\alpha_r)_{r=1, \dots, K-1} \end{array} \right).$$

Proof The Riemannian gradient of a product manifold is the Cartesian product of the individual expressions (Absil et al. 2008). We compute the Riemannian gradients with respect to $S_1, \dots, S_K$ and η .

The gradient with respect to η is the classical Euclidean gradient, hence we get by using the chain rule

$$\begin{aligned} (\text{grad } \hat{\mathcal{L}}(\theta))_{\eta_r} &= (\text{grad}^e \hat{\mathcal{L}}(\theta))_{\eta_r} \\ &= \sum_{i=1}^m \left(\sum_{k=1}^K h^i(\theta_k) \right)^{-1} \sum_{j=1}^K \frac{\partial h^i(\theta_j)}{\partial \eta_l} \end{aligned}$$

$$= \sum_{i=1}^m \sum_{j=1}^K f_j^i \left(\mathbb{1}_{\{j=r\}} - \frac{\exp(\eta_r)}{\sum_{k=1}^K \exp(\eta_k)} \right),$$

for $r = 1, \dots, K-1$, where $\mathbb{1}_{\{j=r\}} = 1$ if $j = r$ and 0, else.

For the derivative of the penalizer with respect to η_r , we get

$$\begin{aligned} (\text{grad Pen}(\theta))_{\eta_r} &= (\text{grad}^e \text{Pen}(\theta))_{\eta_r} \\ &= \zeta \left(1 - K \frac{\exp(\eta_r)}{\sum_{j=1}^K \exp(\eta_j)} \right). \end{aligned}$$

The Riemannian gradient with respect to the matrices $S_1, \dots, S_K$ is the projected Euclidean gradient onto the subspace $T_{S_j} \mathbb{P}^{d+1}$ (with inner product 15), see Absil et al. (2008) and Boumal (2020). The relationship between the Euclidean gradient $\text{grad}^e f$ and the Riemannian gradient $\text{grad } f$ for an arbitrary function $f : \mathbb{P}^n \rightarrow \mathbb{R}$ with respect to the intrinsic inner product defined on the set of positive definite matrices (15) reads

$$\text{grad } f(S) = \frac{1}{2} S \left(\text{grad}^e f(S) + (\text{grad}^e f(S))^T \right) S, \quad (18)$$

see for example Hosseini and Sra (2015) and Jeuris et al. (2012). In a first step, we thus compute the Euclidean gradient with respect to a matrix S_l :

$$(\text{grad}^e \hat{\mathcal{L}}(\theta))_{S_l} = -\frac{1}{2} \sum_{i=1}^m f_l^i (S_l^{-1} y_i y_i^T S_l^{-1} - S_l^{-1}), \quad (19)$$

where we used the Leibniz rule and the partial matrix derivatives

$$\begin{aligned} \frac{\partial (\det(S_l)^{-1/2})}{\partial S_l} &= -\frac{1}{2} (\det(S_l))^{-1/2} S_l^{-1}, \\ \frac{\partial \exp\left(-\frac{1}{2} y_i^T S_l^{-1} y_i\right)}{\partial S_l} &= \frac{1}{2} \exp\left(-\frac{1}{2} y_i^T S_l^{-1} y_i\right) \\ &\quad \times S_l^{-1} y_i y_i^T S_l^{-1}, \end{aligned}$$

which holds by the chain rule and the fact that S_l^{-1} is symmetric.

Using the relationship (18) and using (19) yields the Riemannian gradient with respect to S_l . It is given by

$$(\text{grad } \hat{\mathcal{L}}(\theta))_{S_l} = \frac{1}{2} \sum_{i=1}^m f_l^i (y_i y_i^T - S_l).$$

Analogously, we compute the Euclidean gradient of the matrix penalizer $\psi(S_j, \Phi)$ and use the relationship (18) to get the Riemannian gradient of the matrix penalizer. $\square$

To apply Newton-like algorithms, we derived a formula for the Riemannian Hessian of the reformulated problem. It is stated in Theorem 4:

Theorem 4 Let $\theta \in \mathcal{M}$ and $\xi_\theta \in T_\theta \mathcal{M}$. The Riemannian Hessian is given by

$$\text{Hess}(\hat{\mathcal{L}}(\theta))[\xi_\theta] = (\zeta_S, \zeta_\eta) \in T_\theta \mathcal{M},$$

where

$$\zeta_{S_l} = -\frac{1}{4} \sum_{i=1}^m f_l^i \left[C_l^i - \left(a_l^i - \sum_{j=1}^K f_j^i a_j^i \right) (y_i y_i^T - S_l) \right] \quad (20)$$

$$\begin{aligned} \zeta_{\eta_r} = \frac{1}{2} \sum_{i=1}^m \left[f_r^i \left(a_r^i - \sum_{j=1}^K f_j^i a_j^i \right) \right. \\ \left. - 2\alpha_r \left(\xi_{\eta_r} - \sum_{j=1}^{K-1} \alpha_j \xi_{\eta_j} \right) \right] \end{aligned} \quad (21)$$

for $l = 1, \dots, K$, $r = 1, \dots, K-1$ and

$$a_l^i = y_i^T S_l^{-1} \xi_{S_l} S_l^{-1} y_i - \text{tr}(S_l^{-1} \xi_{S_l}) + 2\xi_{\eta_l},$$

$$f_l^i = \frac{h^i(\theta_l)}{\sum_{j=1}^K h^i(\theta_j)},$$

$$C_l^i = y_i y_i^T S_l^{-1} \xi_{S_l} + \xi_{S_l} S_l^{-1} y_i y_i^T$$

$$\alpha_r = \frac{\exp(\eta_r)}{\sum_{k=1}^K \exp(\eta_k)},$$

$$\xi_{\eta_K} \equiv 0.$$

The Hessian for the additive penalizer reads $\text{Hess}(\text{Pen}(\theta))[\xi_\theta] = (\zeta_S^{\text{pen}}, \zeta_\eta^{\text{pen}})$, where

$$\zeta_{S_l}^{\text{pen}} = \frac{\beta}{4} \left(\Psi S_l^{-1} \xi_{S_l} + \xi_{S_l} S_l^{-1} \Psi \right),$$

$$\zeta_{\eta_l}^{\text{pen}} = K \zeta \alpha_r \left(\xi_{\eta_r} - \sum_{j=1}^{K-1} \alpha_j \xi_{\eta_j} \right).$$

Proof It can be shown that for a product manifold $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ with ∇^1, ∇^2 being the Riemannian connections of $\mathcal{M}_1, \mathcal{M}_2$, respectively, the Riemannian connection ∇ of $\mathcal{M}$ for $X, Y \in \mathcal{M}$ is given by

$$\nabla_Y(X) = \nabla_{Y_1}^1 X_1 + \nabla_{Y_2}^2 X_2,$$

where $X_1, Y_1 \in T\mathcal{M}_1$ and $X_2, Y_2 \in T\mathcal{M}_2$ (Carmo 1992).

Applying this to our problem, we apply the Riemannian connections of the single parts on the Riemannian gradient derived in Theorem 3. It reads

$$\begin{aligned} \text{Hess} \hat{\mathcal{L}}(\theta)[\xi_\theta] &= \nabla_\theta \text{grad} \hat{\mathcal{L}}(\theta) \\ &= \left(\left(\nabla_{\xi_{S_l}}^{(pd)} \text{grad} \hat{\mathcal{L}}(\theta) \right)_l, \left(\nabla_{\xi_{\eta_r}}^e \text{grad} \hat{\mathcal{L}}(\theta) \right)_r \right)^T \end{aligned} \quad (22)$$

for $l = 1, \dots, K$ and $r = 1, \dots, K-1$.

We will now specify the single components of (22). Let

$$\begin{aligned} \text{grad}_{S_l} \hat{\mathcal{L}}(\theta) &= \left(\text{grad} \hat{\mathcal{L}}(\theta) \right)_{S_l}, \\ \text{grad}_{\eta_r} \hat{\mathcal{L}}(\theta) &= \left(\text{grad} \hat{\mathcal{L}}(\theta) \right)_{\eta_r} \end{aligned}$$

denote the Riemannian gradient from Theorem 3 at position S_l, η_r , respectively.

For the latter part in (22), we observe that the Riemannian connection $\nabla_{\xi_{\eta_l}}^e$ for $\xi_{\eta_l} \in \mathbb{R}$ is the classical vector field differentiation (Absil et al. 2008, Section 5.3). We obtain

$$\begin{aligned} \nabla_{\xi_{\eta_l}}^e \text{grad} \hat{\mathcal{L}}(\theta) &= \sum_{j=1}^K D_{S_j}(\text{grad}_{\eta_l} \hat{\mathcal{L}}(\theta))[\xi_{S_j}] \\ &\quad + \sum_{j=1}^{K-1} D_{\eta_r}(\text{grad}_{\eta_j} \hat{\mathcal{L}}(\theta))[\xi_{\eta_j}], \end{aligned} \quad (23)$$

where $D_{S_j}(\cdot)[\xi_{S_j}]$, $D_{\eta_r}(\cdot)[\xi_{\eta_r}]$ denote the classical Fréchet derivatives with respect to S_j, η_j along the directions ξ_{S_j} and ξ_{η_j} , respectively.

For the first part on the right hand side of (23), we have

$$\begin{aligned} &\sum_{j=1}^K D_{S_j}(\text{grad}_{\eta_r} \hat{\mathcal{L}}(\theta))[\xi_{S_j}] \\ &= \frac{1}{2} \sum_{i=1}^m \left[\frac{h^i(\theta_r)}{\sum_{k=1}^K h^i(\theta_k)} \left(y_i^T S_r^{-1} \xi_{S_r} S_r^{-1} y_i - \text{tr}(S_r^{-1} \xi_{S_r}) \right) \right. \\ &\quad \left. - \sum_{j=1}^K \frac{h^i(\theta_j)}{\sum_{k=1}^K h^i(\theta_k)} \left(y_i^T S_j^{-1} \xi_{S_j} S_j^{-1} y_i - \text{tr}(S_j^{-1} \xi_{S_j}) \right) \right] \end{aligned}$$

and for the second part

$$\sum_{j=1}^{K-1} D_{\eta_j}(\text{grad}_{\eta_r} \hat{\mathcal{L}}(\theta))[\xi_{\eta_j}]$$

$$= \sum_{i=1}^m \left[\left(\frac{h^i(\theta_r)}{\sum_{k=1}^K h^i(\theta_k)} - \alpha_r \right) \xi_{\eta_r} + \alpha_r \sum_{j=1}^{K-1} \alpha_j \xi_{\eta_j} - \frac{h^i(\theta_r)}{\sum_{k=1}^K h^i(\theta_j)} \sum_{j=1}^{K-1} \frac{h^i(\theta_j)}{\sum_{k=1}^K h^i(\theta_j)} \xi_{\eta_j} \right]$$

by applying the chain rule, the Leibniz rule and the relationship $\alpha_l = \frac{\exp(\eta_l)}{\sum_{k=1}^K \exp(\eta_k)}$. Plugging the terms into (23), this yields the expression for ζ_{η_r} in (21).

For the Hessian with respect to the matrices S_l , we first need to specify the Riemannian connection with respect to the inner product (15). It is uniquely determined as the solution to the Koszul formula (Absil et al. 2008, Section 5.3), hence we need to find an affine connection that satisfies the formula. For a positive definite matrix S and symmetric matrices ξ_S and $D(\xi_S)[\xi_S]$, this solution is given by (Jeuris et al. 2012; Sra and Hosseini 2015)

$$\nabla_{\nu_S}^{(pd)} \xi_S = D(\xi_S)[\nu_S] - \frac{1}{2}(\nu_S S^{-1} \xi_S + \xi_S S^{-1} \nu_S),$$

where ξ_S , ν_S are vector fields on $\mathcal{M}$ and $D(\xi_S)[\nu_S]$ denotes the classical Fréchet derivative of ξ_S along the direction ν_S . Hence, for the first part in (22), we get

$$\begin{aligned} & \left(\nabla_{\xi_{S_l}}^{(pd)} \text{grad } \hat{\mathcal{L}}(\theta) \right)_l \\ &= \left(\sum_{j=1}^K D_{S_j}(\text{grad}_{S_l} \hat{\mathcal{L}}(\theta))[\xi_{S_j}] + \sum_{j=1}^{K-1} D_{\eta_j}(\text{grad}_{S_l} \hat{\mathcal{L}}(\theta))[\xi_{\eta_j}] - \frac{1}{2} \left(\text{grad}_{S_l} \hat{\mathcal{L}}(\theta) S_l^{-1} \xi_{S_l} + \xi_{S_l} S_l^{-1} \text{grad}_{S_l} \hat{\mathcal{L}}(\theta) \right) \right)_l. \end{aligned} \quad (24)$$

After applying the chain rule and Leibniz rule, we obtain

$$\begin{aligned} & \sum_{j=1}^K D_{S_j}(\text{grad}_{S_l} \hat{\mathcal{L}}(\theta))[\xi_{S_j}] \\ &= -\frac{1}{4} \sum_{i=1}^m f_l^i \left[2\xi_{S_l} - \left((y_i^T S_l^{-1} \xi_{S_l} S_l^{-1} y_i - \text{tr}(S_l^{-1} \xi_{S_l})) \right. \right. \\ & \quad \left. \left. + \sum_{j=1}^K f_j^i (y_i^T S_j^{-1} \xi_{S_j} S_j^{-1} y_i - \text{tr}(S_j^{-1} \xi_{S_j})) \right) \right. \\ & \quad \left. \times (y_i y_i^T - S_l) \right] \end{aligned} \quad (25)$$

and

$$\begin{aligned} & \sum_{j=1}^{K-1} D_{\eta_j}(\text{grad}_{S_l} \hat{\mathcal{L}}(\theta))[\xi_{\eta_j}] \\ &= \frac{1}{2} \sum_{i=1}^m \left(\frac{h^i(\theta_l)}{\sum_{k=1}^K h^i(\theta_j)} \left(\xi_{\eta_l} - \sum_{j=1}^{K-1} \frac{h^i(\theta_j)}{\sum_{k=1}^K h^i(\theta_k)} \xi_{\eta_k} \right) \right) \\ & \quad \times (y_i y_i^T - S_l). \end{aligned} \quad (26)$$

We plug (25), (26) into (24) and use the Riemannian gradient at position S_l for the last term in (24). After some rearrangement of terms, we obtain the expression for ζ_{S_l} in (21).

The computation of $\text{Hess}(\text{Pen}(\theta))[\xi_\theta]$ is analogous by replacing $\hat{\mathcal{L}}$ with $\varphi(\eta, \zeta)$ in (23) and with $\psi(S_l, \Phi)$ in (24).

3 Riemannian Newton trust-region algorithm

Equipped with the Riemannian gradient and the Riemannian Hessian, we are now in the position to apply Newton-type algorithms to our optimization problem. As studying positive-definiteness of the Riemannian Hessian from Theorem 4 is hard, we suggest to introduce some safeguarding strategy for the Newton method by applying a Riemannian Newton Trust-Region method.

3.1 Riemannian Newton trust-region method

The Riemannian Newton Trust-Region Algorithm is the retraction-based generalization of the standard Trust-Region method (Conn et al. 2000)

on manifolds, where the quadratic subproblem uses the Hessian information for an objective function f that we seek to minimize. Theory on the Riemannian Newton Trust-Region method can be found in detail in Absil et al. (2008), we here state the Riemannian Newton Trust-Region method in Algorithm 1. Furthermore, we will study both global and local convergence theory for our penalized problem.

Global convergence. In the following Theorem, we show that the Riemannian Newton Trust-Region Algorithm applied on the reformulated penalized problem converges to a stationary point. To the best of our knowledge, this has not been proved before.

Theorem 5 (Global convergence) *Consider the penalized reformulated objective $\hat{\mathcal{L}}_{\text{pen}}$ from (9). If we apply the Riemannian Newton Trust-Region Algorithm (Algorithm 1) to minimize $f = -\hat{\mathcal{L}}_{\text{pen}}$, it holds*

$$\lim_{t \rightarrow \infty} \text{grad } \hat{\mathcal{L}}_{\text{pen}}(\theta^t) = 0. \quad (27)$$

Algorithm 1: Riemannian Trust-Region algorithm Absil et al. (2008)

Input: objective f with Hessian H , initial iterate $\theta^0 \in \mathcal{M}$, initial TR-radius Δ_0 , maximal TR-radius $\bar{\Delta}$, rejection threshold $\rho' \in [0, 1/4)$, acceptance parameters $0 \leq \omega_1 \leq \omega_2 \leq 1$, $\tau_1 \leq 1/4$, $\tau_2 > 1$, termination criteria

- 1 Set $t = 0$;
- 2 **while** termination criteria not fulfilled **do**
- 3 Obtain s^t by (approximately) solving the TR-subproblem

$$\min_{s \in T_{\theta^t} \mathcal{M}} \hat{m}_{\theta^t}(s) = f(\theta^t) + \langle \text{grad } f(\theta^t), s \rangle_{\theta^t} + \frac{1}{2} \langle H_t[s], s \rangle_{\theta^t}$$
 s.t. $\langle s, s \rangle_{\theta^t} \leq \Delta_t^2$;
- 4 Evaluate $\rho_t = \frac{f(\theta^t) - f(R_{\theta^t}(s^t))}{\hat{m}_{\theta^t}(0_{\theta^t}) - \hat{m}_{\theta^t}(s^t)}$
- 5 **if** $\rho_t < \omega_1$ **then**
- 6 $\Delta_{t+1} = \tau_1 \Delta_t$;
- 7 **else if** $\rho_t > \omega_2$ and $\|s^t\|_{\theta^t} = \Delta_t$ **then**
- 8 $\Delta_{t+1} = \min(\tau_2 \Delta_t, \bar{\Delta})$;
- 9 **else**
- 10 $\Delta_{t+1} = \Delta_t$;
- 11 **if** $\rho_t > \rho'$ **then**
- 12 $\theta^{t+1} = R_{\theta^t}(s^t)$;
- 13 **else**
- 14 $\theta^{t+1} = \theta^t$
- 15 set $t = t + 1$;
- 16 **end**

Proof According to general global convergence results for Riemannian manifolds, convergence to a stationary point is given if the level set $\{\theta : f(\theta) \leq f(\theta^0)\}$ is a compact Riemannian manifold (Absil et al. 2008, Proposition 7.4.5) and the function f is lower bounded. Since the penalized reformulated objective $\hat{\mathcal{L}}_{pen}$ is upper bounded by Theorem 2, the latter is fulfilled. In the following, we will show that the iterates produced by Algorithm 1 stay in a compact set:

Let $\theta^0 = ((S_1^0, \dots, S_K^0), \eta^0) \in \mathcal{M}$ be the starting point of Algorithm 1 applied on the objective $f = -\hat{\mathcal{L}}_{pen}$. We further define the set of successful (unsuccessful) steps $\mathfrak{S}_t$ ($\mathfrak{F}_t$) generated by the algorithm until iteration t by

$$\mathfrak{S}_t = \{l \in \{0, 1, \dots, t\} : \rho_l > \rho'\},$$

$$\mathfrak{F}_t = \{l \in \{0, 1, \dots, t\} : \rho_l \leq \rho'\}.$$

At iterate t , we consider $\theta^t = ((S_1^0, \dots, S_K^0), \eta^0) \in \mathcal{M}$. Let $s^t = ((s_{S_1}^t, \dots, s_{S_K}^t), s_{\eta}^t) \in T_{\theta^t} \mathcal{M}$ be the tangent vector returned by solving the quadratic subproblem in line 3, Algorithm 1 and $R_{S_j^t} = S_j^t \exp((S_j^t)^{-1} s_{S_j}^t)$ be the retraction part of S_j , see (17). Then

$$\begin{aligned} \|\theta^t\| &= \|R_{\theta^{t-1}}(s^t)\| \mathbf{1}_{\{t \in \mathfrak{S}_t\}} + \|\theta^{t-1}\| \mathbf{1}_{\{t \in \mathfrak{F}_t\}} \\ &\leq \sum_{j=1}^K \left(\|R_{S_j^t}(s_{S_j}^t)\| \mathbf{1}_{\{t \in \mathfrak{S}_t\}} + \|S_j^{t-1}\| \mathbf{1}_{\{t \in \mathfrak{F}_t\}} \right) + \|\eta^t\|, \end{aligned} \quad (28)$$

where $\mathbf{1}_{\{t \in \mathfrak{S}_t\}}$ is the indicator function, i.e. $\mathbf{1}_{\{t \in \mathfrak{S}_t\}} = 1$ if $t \in \mathfrak{S}_t$ and 0, else.

Since $\|s^t\|_{\theta^t} \leq \Delta$ for all t , there exists $\bar{\Delta}^1 > 0$, $\bar{\Delta}^2 > 0$ such that $\|s_{S_j}^t\| \leq \bar{\Delta}^1$ for all $j = 1, \dots, K$ and $\|s_{\eta}^t\| \leq \bar{\Delta}^2$. Hence, the second part on the right hand side of (28) yields

$$\begin{aligned} \|\eta^t\| &= \left\| \eta^0 + \sum_{l=1}^{t-1} s_{\eta}^l \mathbf{1}_{\{l \in \mathfrak{S}_l\}} \right\| \\ &\leq \|\eta^0\| + \sum_{l=1}^{t-1} \|s_{\eta}^l\| \\ &\leq \|\eta^0\| + t \bar{\Delta}^2. \end{aligned}$$

For the first part on the right hand side of (28), we take a closer look at $\|R_{S_j^t}(s_{S_j}^t)\|$. For better readability, we omit the index j and simply write $\|R_{S^t}(s_S^t)\|$. Let $\lambda(S^l)$ denote the minimal eigenvalue of S at iteration l . By the inequality of Cauchy-Schwarz and using the retraction (17), we get

$$\begin{aligned} \|R_{S^t}(s_S^t)\| & \leq \|S^{t-1}\| \exp\left(\|(S^{t-1})^{-1}\| \|s_S^t\|\right) \\ & \leq \|S^{t-1}\| \exp\left(\frac{\bar{\Delta}^1}{\|S^{t-1}\|}\right) \\ & \leq \|S^{t-1}\| \exp\left(\frac{\bar{\Delta}^1}{(d+1)\lambda(S^{t-1})}\right) \\ & \leq \|S^{t-2}\| \exp\left(\frac{\bar{\Delta}^1}{(d+1)\lambda(S^{t-1})}\right) \\ & \quad \times \left(\exp\left(\frac{\bar{\Delta}^1}{(d+1)\lambda(S^{t-2})}\right) \mathbf{1}_{\{t-1 \in \mathfrak{S}_t\}} + \mathbf{1}_{\{t \in \mathfrak{F}_t\}} \right). \end{aligned} \quad (29)$$

$$\quad (31)$$

By applying (31) iteratively to (28), we get

$$\|S^t\| \leq \|S^0\| \exp\left(\frac{\bar{\Delta}^1}{d+1} \sum_{l=0}^{t-1} \frac{1}{\lambda(S^l)} \mathbf{1}_{\{l \in \mathfrak{S}_l\}}\right). \quad (32)$$

We now show that the eigenvalues of S_l cannot become infinitesimal small. For this, assume that there exists a subsequence of minimal eigenvalues $\{\lambda(S_{l_i})\}_i$ such that $\lambda(S_{l_i}) \rightarrow 0$ for $i \rightarrow \infty$. Then, according to the proof of Theorem 2, this yields $\hat{\mathcal{L}}_{pen}(\theta^{l_i}) \rightarrow -\infty$ for $i \rightarrow \infty$. This is a contradiction, as the Riemannian Trust-Region method in Algorithm 1 ensures we have a decrease in f (increase in $\hat{\mathcal{L}}_{pen}$) in every iteration.

Thus, the right hand side of (32) cannot become infinitely large and there exists a constant $C > 0$ such that $\|\theta^t\| \leq C$. Hence, we see that our algorithm applied to the (penalized) problem of fitting Gaussian Mixture Models converges to stationary points. $\square$

We have shown that Algorithm 1 applied on the problem of fitting Gaussian Mixture Models converges to a stationary point for all starting values. From general convergence theory for Riemannian Trust-Region algorithms (Absil et al. 2008, Section 7.4.2), under some assumptions, the convergence speed of Algorithm 1 is superlinear. In the following, we show that these assumptions are fulfilled and specify the convergence rate.

Local convergence. Local convergence result on (Riemannian) Trust Region-Methods depend on the solver for the quadratic subproblem. If the quadratic subproblem in Algorithm 1 is (approximately) solved with sufficient decrease, the local convergence close to a maximizer of $\hat{\mathcal{L}}_{pen}$ is superlinear under mild assumptions. The truncated Conjugate Gradient method is a typical choice that returns a sufficient decrease, we suggest to use this matrix-free method for Gaussian Mixture Models. We state the Algorithm in “Appendix B”. The local superlinear convergence is stated in Theorem 6.

Theorem 6 (Local convergence) *Consider Algorithm 1, where the quadratic subproblem is solved by the truncated Conjugate Gradient Method and $f = -\hat{\mathcal{L}}_{pen}$. Let $v \in \mathcal{M}$ be a nondegenerate local minimizer of f , H the Hessian of the reformulated penalized problem from Theorem 4 and $\delta < 1$ the parameter of the termination criterion of tCG (“Appendix B”, line 18 in Algorithm 2). Then there exists $c > 1$ such that, for all sequences $\{\theta^t\}$ generated by Algorithm 1 converging to v , there exists $T > 0$ such that for all $t > T$,*

$$\|\text{grad } \hat{\mathcal{L}}_{pen}(\theta^{t+1})\| \leq c \|\text{grad } \hat{\mathcal{L}}_{pen}(\theta^t)\|^{\delta+1}. \quad (33)$$

Proof Let $v \in \mathcal{M}$ be a nondegenerate local minimizer of f (maximizer of $\hat{\mathcal{L}}_{pen}$). We choose the termination criterion of tCG such that $\delta < 1$. According to Absil et al. (2008, Theorem 7.4.12), it suffices to show that $\text{Hess } \hat{\mathcal{L}}_{pen}(R_\theta(\xi))$ is Lipschitz continuous at 0_θ in a neighborhood of v . From the proof of Theorem 2, we know that close to v , we are bounded away from the boundary of $\mathcal{M}$, i.e. we are bounded away

from points on the manifold with singular S_j , $j = 1, \dots, K$. Thus, there exist neighborhoods $B_{\epsilon_1}, B_{\epsilon_2}$ such that for $\theta \in B_{\epsilon_1}(v) \subset \mathcal{M}$ and $\xi \in B_{\epsilon_2}$, the function $\text{Hess } \hat{\mathcal{L}}_{pen}(R_\theta(\xi))$ is continuously differentiable. From this, the local Lipschitz continuity follows directly with the Heine–Borel Theorem. Thus, all assumptions of Absil et al. (2008, Theorem 7.4.12) are fulfilled and (33) holds. $\square$

3.2 Practical design choices

We apply Algorithm 1 on our cost function (9), where we seek to minimize $f = -\hat{\mathcal{L}}_{pen}$. The quadratic subproblem in line 3 in Algorithm 1 is solved by the truncated Conjugate Gradient method (tCG) with the inner product (16). To speed up convergence of the tCG, we further use a preconditioner: At iteration t , we store the gradients computed in tCG and store an inverse Hessian approximation via the LBFSGS formula. This inverse Hessian approximation is then used for the minimization of the next subproblem $\hat{m}_{\theta^{t+1}}$. The use of such preconditioners has been suggested by Morales and Nocedal (2000) for solving a sequence of slowly varying systems of linear equations and gave a speed-up in convergence for our method. The initial TR radius is set by using the method suggested by Sartenar (1997) that is based on the model trust along the steepest-descent direction. The choice of parameters $\omega_1, \omega_2, \tau_1, \tau_2, \rho'$ in Algorithm 1 are chosen according to the suggestions in Gould et al. (2005) and Conn et al. (2000).

4 Numerical results

We test our method on the penalized objective function (3) on both simulated and real-world data sets. We compare our method against the (penalized) Expectation Maximization Algorithm and the Riemannian LBFSGS method proposed by Hosseini and Sra (2015). For all methods, we used the same initialization by running *k-means++* (Arthur and Vassilvitskii 2007) and stopped all methods when the difference of average log-likelihood for two subsequent iterates falls below $1e-10$ or when the number of iterations exceeds 1500 (clustering) or 3000 (density approximation). For the Riemannian LBFSGS method suggested by Hosseini and Sra (2015, 2020), we used the MixEst package (Hosseini and Mash’al 2015) kindly provided by one of the authors.¹ For the Riemannian Trust region method we mainly followed the code provided by the *pymanopt* Python package provided by Townsend et al. (2016), but adapted it for a faster implementation by computing the matrix inverse S_j^{-1} only once per iteration. We used Python version 3.7. The experiments were conducted

¹ We also tested the Riemannian SGD method, but the runtimes turned out to be very high, so we omit the results here.

on an Intel Xeon CPU X5650 at 2.67 GHz with 24 cores and 20GB RAM.

In Sect. 4.1, we test our method on clustering problems on simulated data and on real-world datasets from the UCI Machine Learning Repository (Dua and Graff 2017). In Sect. 4.2 we consider Gaussian Mixture Models as probability density estimators and show the applicability of our method.

4.1 Clustering

We test our method for clustering tasks on different artificial (Sect. 4.1.1) and real-world (Sect. 4.1.2) data sets.

4.1.1 Simulated data

As the convergence speed of EM depends on the level of separation between the data, we test our methods on data sets with different degrees of separation as proposed in Dasgupta (1999); Hosseini and Sra (2015). The distributions are sampled such that their means satisfy

$$\|\mu_i - \mu_j\|^2 \geq c \max_{i,j} (\text{tr}(\Sigma_i), \text{tr}(\Sigma_j))$$

for $i, j = 1, \dots, K, i \neq j$ and c models the degree of separation. Additionally, a low *eccentricity* (or condition number) of the covariance matrices has an impact on the performance of Expectation Maximization (Dasgupta 1999), for which reason we also consider different values of eccentricity $e = \sqrt{\left(\frac{\lambda_{\max}(\Sigma_j)}{\lambda_{\min}(\Sigma_j)}\right)}$. This is a measure of how much the data scatters.

We test our method on 20 and 40-dimensional data and an equal distribution among the clusters, i.e. we set $\alpha_j = \frac{1}{K}$ for all $j = 1, \dots, K$. Although it is known that unbalanced mixing coefficients α_j also result in slower EM convergence, this effect is less strong than the level of overlap (Naim and Gildea 2012), so we only show simulation results with balanced clusters. Results for unbalanced mixing coefficients and varying number of components K are shown for real-world data in Sect. 4.1.2.

We first take a look at the 20-dimensional data sets, for which we simulated $m = 1000$ data points for each parameter setting. In Table 1, we show the results for very scattered data, that is $e = 1$. We see that, like predicted by literature, the Expectation Maximization converges slowly in such a case. This effect is even stronger with a lower separation constant c . The effect of the eccentricity becomes even more clear when comparing the results of Table 1 with Table 2. Also the Riemannian algorithms converge slower for lower values of eccentricity e and separation levels c . However, they seem to suffer less from hidden information than Expectation Maximization. The proposed Riemannian Newton Trust

Table 1 Simulation results of 20 runs for dimensions $d = 20$, number of components $K = 5$ and eccentricity $e = 1$

		EM	R-NTR	R-LBFGS
$c = 0.2$	Iterations	295	79.4	113.4
	Mean time (s)	3.833	2.716	16.008
	Mean ALL	−42.64	−42.63	−42.63
	MSE weights	0.00139	0.00143	0.008
	MSE means	0.14	0.14	0.13
	MSE cov	2.27	2.5	2.28
$c = 1$	Iterations	262	47.5	102.7
	Mean time (s)	3.65	2.069	14.281
	Mean ALL	−41.2	−41.21	−41.21
	MSE weights	0.00909	0.00985	0.008
	MSE means	0.23	0.22	0.24
	MSE cov	0.67	0.56	0.7
$c = 5$	Iterations	208.8	54.2	92.4
	Mean time (s)	2.747	2.123	12.975
	Mean ALL	−36.98	−36.98	−36.99
	MSE weights	0.00264	0.00282	0.008
	MSE means	0.15	0.17	0.16
	MSE cov	9.81	7.1	10.19

Table 2 Simulation results of 20 runs for dimensions $d = 20$, number of components $K = 5$ and eccentricity $e = 10$

		EM	R-NTR	R-LBFGS
$c = 0.2$	Iterations	66.2	16	33.2
	Mean time (s)	0.873	0.798	4.012
	Mean ALL	−60.06	−60.06	−60.07
	MSE weights	3e−05	3e−05	0.008
	MSE means	0.07	0.07	0.07
	MSE cov	0.31	0.23	0.31
$c = 1$	Iterations	56.6	17.4	30
	Mean time (s)	0.748	0.818	3.624
	Mean ALL	−62.82	−62.82	−62.83
	MSE weights	3e−05	3e−05	0.008
	MSE means	0.09	0.09	0.09
	MSE cov	0.17	0.16	0.17
$c = 5$	Iterations	43.1	14.7	29
	Mean time (s)	0.617	0.748	3.378
	Mean ALL	−61.04	−61.04	−61.05
	MSE weights	4e−05	4e−05	0.008
	MSE means	0.08	0.08	0.08
	MSE cov	0.13	0.14	0.13

Region algorithm (R-NTR) beats the other methods in terms of runtime and number of iterations (see Fig. 2). The Riemannian LBFGS (R-LBFGS) method by Hosseini and Sra (2015) also shows faster convergence than EM, but the gain

Fig. 2 Average penalized log-likelihood reduction for highly overlapping clusters: $d = 20$, $K = 5$, $e = 1$, $c = 0.2$

Fig. 3 Average penalized log-likelihood for highly overlapping clusters: $d = 20$, $K = 5$, $e = 1$, $c = 0.2$

of second-order information available by the Riemannian Hessian is obvious. However, the R-LBFGS results created by the MixEst toolbox show long runtimes compared to the other methods. We see from Fig. 3 that the average penalized log-likelihood is slightly higher for R-LBFGS in some experiments. Still, the objective evaluated at the point satisfying the termination criterion is at a competitive level in all methods (see also Table 1).

When increasing the eccentricity (Table 2), we see that the Riemannian methods still converge faster than EM, but our method is not faster than EM. This is because EM benefits from very low per-iteration costs and the gain in number of iterations is less strong in this case. However, we see that the Riemannian Newton Trust-Region method is not substantially slower. Furthermore, the average log-likelihood values (ALL) are more or less equal in all methods, so we might assume that all methods stopped close to a similar optimum. This is also underlined by comparable mean squared errors

Table 3 Simulation results of 20 runs for dimensions $d = 40$, number of components $K = 5$ and eccentricity $e = 1$ with $m = 1000$ observations

		EM	R-NTR	R-LBFGS
$c = 0.2$	Iterations	57.4	27.9	40.7
	Mean time (s)	1.332	2.296	6.744
	Mean ALL	-84.7935	-84.79	-84.7895
	MSE weights	0.00023	0.00023	0.008
	MSE means	0.104	0.09	0.104
	MSE cov	0.282	0.196	0.281
$c = 1$	Iterations	61	29	48.6
	Mean time (s)	1.391	2.476	8.282
	Mean ALL	-82.3395	-82.3384	-82.3358
	MSE weights	0.0002	0.0002	0.008
	MSE means	0.076	0.084	0.076
	MSE cov	0.139	0.128	0.14
$c = 5$	Iterations	81.8	28.8	49.2
	Mean time (s)	1.812	2.203	8.62
	Mean ALL	-92.4886	-92.4874	-92.4925
	MSE weights	0.00013	0.00013	0.008
	MSE means	0.08	0.095	0.08
	MSE cov	0.116	0.133	0.117

(MSE) to the true parameters from which the input data has been sampled from. In average, Riemannian Newton Trust-Region gives the best results in terms of runtime and number of iterations.

In Table 3, we show results for dimension $d = 40$ and low eccentricity ($e = 1$) and the same simulation protocol as above (in particular, $m = 1000$). We observed that with our method, we only performed very few Newton-like steps and instead exceeded the trust-region within the tCG many times, leading to poorer steps (see also Fig. 5). One possible reason is that the number of parameters increases with d quadratically, that is in $\mathcal{O}(Kd^2)$, while at the same time we did not increase the number of observations $m = 1000$. If we are too far from a local optimum and the clusters are not well initialized due to few observations, the factor f_l^i in the Hessian (Theorem 4) becomes small, leading to large potential conjugate gradients steps (see Algorithm 2). Although this affects the E-step in the Expectation Maximization algorithm as well, the effect seems to be much severe in our method.

To underline this, we show simulation results for a higher number of observations, that is, $m = 10000$, in Table 4 with the same true parameters $\alpha_j, \mu_j, \Sigma_j$ as in Table 3. As expected, the superiority in runtime of our method becomes visible: The R-NTR method beats Expectation maximization with a factor of 4. Just like for the case of a lower dimension $d = 20$, the mean average log-likelihood and the errors are comparable between our method and EM, whereas R-LBFGS shows slightly worse results although it now attains comparable runtimes to our method.

Table 4 Simulation results of 20 runs for dimensions $d = 40$, number of components $K = 5$ and eccentricity $e = 1$ with $m = 10,000$ observations

		EM	R-NTR	R-LBFGS
$c = 0.2$	Iterations	350.4	33.2	69.4
	Mean time (s)	53.621	12.455	15.449
	Mean ALL	-86.5717	-86.5718	-86.5731
	MSE weights	0.00043	0.00043	0.008
	MSE means	0.093	0.086	0.093
	MSE cov	0.207	0.195	0.206
$c = 1$	Iterations	495.6	63.6	107.3
	Mean time (s)	79.955	20.739	23.153
	Mean ALL	-84.3783	-84.3779	-84.3797
	MSE weights	0.00062	0.00065	0.008
	MSE means	0.075	0.064	0.076
	MSE cov	0.075	0.038	0.075
$c = 5$	Iterations	260.4	28.8	54
	Mean time (s)	42.6	10.434	11.692
	Mean ALL	-94.5592	-94.5591	-94.5603
	MSE weights	0.00012	0.00013	0.008
	MSE means	0.071	0.086	0.071
	MSE cov	0.045	0.053	0.045

We thus see that the ratio between number of observations and number of parameters must be large enough in order to benefit from the Hessian information in our method.

4.1.2 Real-world data

We tested our method on some real-world data sets from UCI Machine Learning repository (Dua and Graff 2017) besides the simulated data sets. For this, we normalized the data sets and tested the methods for different values of K .

Combined Cycle Power Plant Data Set (Kaya and Tufekci 2012). In Table 5, we show the results for the combined cycle power plant data set, which is a data set with 4 features. Although the dimension is quite low, we see that we can beat EM both in terms of runtime and number of iterations for almost all K by applying the Riemannian Newton-Trust Region method. This underlines the results shown for artificial data in Sect. 4.1.1. The gain by our method becomes even stronger when we consider a large number of components K where the overlap between clusters is large and we can reach a local optimum with our method in up to 15 times less iterations and a time saving of factor close to 4.

MAGIC Gamma Telescope Data Set (Bock et al. 2004). We also study the behaviour on a data set with higher dimensions and a larger number of observations with the MAGIC Gamma Telescope Data Set, see Table 6. Here, we can also observe a lower number of iterations in the Riemannian Optimization methods. Similar to the combined cycle power plant

Fig. 4 Average penalized log-likelihood reduction for overlapping clusters: $d = 40$, $K = 5$, $e = 1$, $c = 1$

Fig. 5 Average penalized log-likelihood for overlapping clusters: $d = 40$, $K = 5$, $e = 1$, $c = 1$

Table 5 Results of (normalized) combined cycle power plant data set for different number of components

		EM	R-NTR	R-LBFGS
$K = 2$	Time (s)	0.40	0.63	2.38
	Iterations	56	19	34
	ALL	-4.24	-4.24	-4.24
$K = 5$	Time (s)	3.29	2.26	7.50
	Iterations	239	48	70
	ALL	-4.01	-4.01	-4.01
$K = 10$	Time (s)	31.72	4.28	23.40
	Iterations	1097	58	110
	ALL	-3.83	-3.82	-3.83
$K = 15$	Time (s)	28.27	6.79	35.77
	Iterations	677	67	111
	ALL	-3.75	-3.75	-3.75

Number of observations $m = 9568$, dimensions $d = 4$

Table 6 Results of (normalized) magic gamma telescope data set for different number of components

		EM	R-NTR	R-LBFGS
K = 2	Time (s)	1.18	0.57	1.52
	Iterations	30	6	17
	ALL	-7.81	-7.81	-7.81
K = 5	Time (s)	3.96	1.32	5.37
	Iterations	65	9	34
	ALL	-6.53	-6.53	-6.53
K = 10	Time (s)	36.00	6.99	20.78
	Iterations	293	34	77
	ALL	-6.02	-6.02	-6.02
K = 15	Time (s)	56.24	12.61	50.52
	Iterations	354	38	115
	ALL	-5.39	-5.54	-5.51

Number of observations $m = 19020$, dimensions $d = 11$

Table 7 Weighted mean squared errors of (normalized) magic gamma telescope data set for $K = 2$

	EM	R-NTR	R-LBFGS
distance	4.783910	4.783934	4.783922
wMSE weight	0.000034	0.000034	0.000034
wMSE mean	1.883762	1.883765	1.883759
wMSE cov	8.214824	8.214840	8.214828

data set, this effect becomes even stronger for a high number of clusters where the ratio of hidden information is large. Our method shows by far the best runtimes. For this data set, the average log-likelihood values are very close to each other except for $K = 15$ where the ALL is worse for the Riemannian methods. It seems that in this case, the R-NTR and the R-LBFGS methods end in different local maxima than the EM. However, for all of the methods, convergence to global maxima is theoretically not ensured and for all methods, a globalization strategy like a split-and-merge approach (Li and Li 2009) might improve the final ALL values. As the Magic Gamma telescope data set is a classification data set with 2 classes, we further report the classification performance in Table 7. We see that the geodesic distance defined on the manifold and the weighted mean squared errors (wMSE) are comparable between all three methods. In Table 8, we also report the Adjusted Rand Index (Hubert and Arabie 1985) for all methods. Although the clustering performance is very low compared to the true class labels (first row), we see that it is equal among the three methods.

We show results on additional real-world data sets in “Appendix C”.

Table 8 Adjusted rand index for (normalized) magic gamma telescope data set ($K = 2$)

	Truth	EM	R-NTR	R-LBFGS
Truth	1	0.06	0.06	0.06
EM		1	1.00	1.00
R-NTR			1	1.00
R-LBFGS				1

Fig. 6 Probability density of bivariate Beta(0.5,0.5)-Gamma(1, 1) distribution

4.2 Gaussian mixture models as density approximators

Besides the task of clustering (multivariate) data, Gaussian Mixture Models are also well-known to serve as probability density approximators of smooth density functions with enough components (Scott 1992).

In this Subsection, we present the applicability of our method for Gaussian mixture density approximation of a Beta-Gamma distribution and compare against EM and R-LBFGS for the estimation of the Gaussian components.

We consider a bivariate Beta-Gamma distribution with parameters $\alpha_{Beta} = 0.5$, $\beta_{Beta} = 0.5$, $a_{Gamma} = 1$, $\beta_{Gamma} = 1$, where the joint distribution is characterized by a Gaussian copula. The density function surface is visualized in Fig. 6.

We simulated 1000 realizations of the Beta(0.5,0.5)-Gamma(1, 1) distribution and fitted a Gaussian Mixture Model for 100 simulation runs. We considered different numbers of K and compared the (approximated) root mean integrated squared error (RMISE).

Table 9 Simulation results averaged over 100 simulation runs for approximation of a Beta(0.5,0.5)-Gamma(1, 1) distribution by a Gaussian Mixture Models with different values of K

		EM	R-NTR	R-LBFGS
K = 2	Mean RMISE	0.00453	0.00453	0.00453
	SE RMISE	0.00017	0.00017	0.00017
	Iterations	113	27.9	25.9
	Mean time (s)	0.2002	0.1532	0.9877
	Mean ALL	−1.53103	−1.52953	−1.53103
K = 5	Mean RMISE	0.00565	0.00565	0.00567
	SE RMISE	0.00028	0.00024	0.00025
	Iterations	346.5	47	72.9
	Mean time (s)	1.0893	0.5202	6.7941
	Mean ALL	−1.22852	−1.22108	−1.22912
K = 10	Mean RMISE	0.00639	0.00643	0.00643
	SE RMISE	0.00025	0.00027	0.00029
	Iterations	623.1	56.8	112.2
	Mean time (s)	3.4226	1.7108	20.4754
	Mean ALL	−1.06821	−1.02844	−1.06781
K = 15	Mean RMISE	0.00667	0.00669	0.0067
	SE RMISE	0.00026	0.00027	0.00026
	Iterations	791.5	62.9	135.2
	Mean time (s)	6.1649	2.5701	37.6
	Mean ALL	−1.02677	−0.95907	−1.02604
K = 20	Mean RMISE	0.00681	0.00683	0.00685
	SE RMISE	0.00026	0.00025	0.00026
	Iterations	874.8	66.9	144.6
	Mean time (s)	8.7694	3.3773	55.6337
	Mean ALL	−1.02029	−0.9202	−1.02056

Parameter estimation by EM, R-NTR and R-LBFGS

The RMISE and its computational approximation formula is given by (Gross et al. 2015):

$$RMISE(\hat{f}) = \sqrt{\mathbb{E} \left(\int (f(x) - \hat{f}(x))^2 dx \right)}$$

$$\approx \sqrt{\frac{1}{N} \sum_{r=1}^N (f(g_r) - \hat{f}(g_r))^2 \delta_g^2},$$

where f denotes the underlying true density, i.e. Beta(0.5, 0.5)-Gamma(1, 1), $\hat{f}$ the density approximator (GMM), N is the number of equidistant grid points with the associated grid width δ_g .

For our simulation study, we chose 16384 grid points in the box $[0, 5] \times [0, 10]$. We show the results in Table 9, where we fit the parameters of the GMM by our method (R-NTR) and compare against GMM approximations where we fit the parameters with EM and R-LBFGS. We observe that the RMISE is of comparable size for all methods and

Fig. 7 Contours of pointwise root mean squared error (RMSE) for density approximation via GMMs of the Beta (0.5,0.5)-Gamma (1, 1) distribution

even slightly better for our method for $K = 2$, $K = 5$ and $K = 10$. Just as for the clustering results in Sect. 4.1, we have much lower runtimes for R-NTR and a much lower number of total iterations. This is a remarkable improvement especially for a larger number of components. We also observe that in all methods, the mean average log-likelihood (ALL) of the training data sets with 1000 observations attains higher values with an increasing number of components K . This supports the fact that the approximation power of GMMs for arbitrary density function is expected to become higher if we add additional Gaussian components (Scott 1992; Goodfellow et al. 2016). On the other hand, the RMISE (which is not based on the training data) increased in our experiments with larger K 's. This means that we are in a situation of overfitting. The drawback of overfitting is well-known for EM (Andrews 2018) and we also observed this for the R-NTR and the R-LBFGS methods. However, the RMISE are comparable and so none of the methods outperforms another substantially in

terms of overfitting. This can also be seen from Fig. 7 showing the distribution of the pointwise errors for $K = 2$ and $K = 5$. Although the R-LBFGS method shows higher error values on the boundary of the support of the distribution for $K = 5$, the errors show similar distributions among the three methods at a comparable level. We propose methodologies such as cross validation (Murphy 2013) or applying a split-and-merge approach on the optimized parameters (Li and Li 2009) to address the problem of overfitting.

The results show that our method is well-suited for both density estimation tasks and especially for clustering of real-world and simulated data.

5 Conclusion

We proposed a Riemannian Newton Trust-Region method for Gaussian Mixture Models. For this, we derived an explicit formula for the Riemannian Hessian and gave results on convergence theory. Our method is a fast alternative to the well-known Expectation Maximization and existing Riemannian approaches for Gaussian Mixture Models. Especially for highly overlapping components, the numerical results show that our method leads to an enormous speed-up both in terms of runtime and the total number of iterations and medium-sized problems. This makes it a favorable method for density approximation as well as for difficult clustering tasks for data with higher dimensions. Here, especially the availability of the Riemannian Hessian increases the convergence speed compared to Quasi-Newton algorithms. When considering higher-dimensional data, we experimentally observed that our method still works very well and is faster than EM if the number of observations is large enough. We plan to examine this further and take a look into data sets with higher dimensions. Here, it is common to impose a special structure on the covariance matrices to tackle the curse of dimensionality (McLachlan et al. 2019). Adapted versions for Expectation Maximization exist and a transfer of Riemannian method into constrained settings is subject of current research.

Acknowledgements We thank Mrs. Ngan Mach for her support of contributing to the creation of numerical results. We would also like to thank two anonymous reviewers that gave important feedback on the paper.

Funding Open Access funding enabled and organized by Projekt DEAL. This research has been supported by the German Research Foundation (DFG) within the Research Training Group 2126: Algorithmic Optimization, Department of Mathematics, University of Trier, Germany.

Declarations

Conflict of interest The authors declare that there is no conflict of interest.

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

A Unboundedness of the reformulated problem

Theorem 7 *The reformulated objective (3) without penalization term is unbounded from above.*

Proof We show that there exists θ_s such that $\lim_{\theta \rightarrow \theta_s} \hat{\mathcal{L}}(\theta) = \infty$.

We consider the simplified case where $K = 1$ and investigate the log-likelihood of the variable $\theta^{(n)} = S^{(n)}$ since $\alpha_1 = 1$. Similar to the proof of Theorem 2, we introduce a singular positive semidefinite matrix S_s with rank $r < d + 1$ and a sequence of positive definite matrices $S_1^{(n)}$ converging to S_s for $n \rightarrow \infty$. The reformulated objective at $\theta^{(n)}$ reads

$$\hat{\mathcal{L}}(\theta^{(n)}) = \sum_{i=1}^m \left(\log \left(\frac{1}{(2\pi)^{d/2}} \frac{\exp(\frac{1}{2}(1 - y_i^T S_1^{(n)} y_i))}{\det(S_1^{(n)})^{1/2}} \right) \right). \quad (34)$$

With the decomposition $S_1^{(n)} = U^T D^n U$ as in (12), we see that

$$y_i^T S_1^{(n)} y_i = \sum_{l=1}^{d+1} \frac{((U y_i)_l)^2}{\lambda_l} + \sum_{l=r-d}^{d+1} \frac{((U y_i)_l)^2}{\lambda_l}$$

Now assume that one of the y_i is in the kernel of the matrix $\tilde{U}_s$, where

$$\tilde{U}_s = \begin{pmatrix} u_{11} & \dots & u_{1(r-d+1)} \\ \vdots & \ddots & \vdots \\ u_{(r-d+1)1} & \dots & u_{(r-d+1)(r-d+1)} \end{pmatrix}.$$

Then we see that (34) reads

$$\hat{\mathcal{L}}(\theta^{(n)}) = \sum_{i=1}^m \left[\log \left(\frac{1}{(2\pi)^{d/2}} \frac{1}{\lambda^{(n)} \lambda^+} \times \exp \left(\frac{1}{2} \left(1 - \sum_{l=r-d}^{d+1} \frac{((U y_i)_l)^2}{\lambda_l} \right) \right) \right) \right]$$

$$n \rightarrow \infty, \infty,$$

where $\lambda^{(n)} = \prod_{l=1}^{d-r-1} \lambda_l^{(n)}$ and $\lambda^+ = \prod_{l=d-r}^{d+1} \lambda_l$ as in the proof of Theorem 2. $\square$

B Truncated conjugate gradient method

We here state the truncated CG method to solve the quadratic subproblem in the Riemannian Newton Trust-Region Algorithm, line 3 in Algorithm 1.

Algorithm 2: Truncated Conjugate Gradient Method
Absil et al. (2008)

Input: Riemannian gradient $\text{grad } f(\theta^t)$, linear operator H_t , inner product $\langle \cdot, \cdot \rangle_{\theta^t}$, optional preconditioner M_t , TR radius Δ_t , termination parameter δ, κ

- 1 Set $s_0 = 0, r_0 = \text{grad } f(\theta^t), z_0 = M_t r_0, p_0 = -z_0$;
- 2 **for** $n = 0, 1, 2, \dots$ **do**
- 3 **if** $\langle p_n, H_t[p_n] \rangle_{\theta^t} \leq 0$ **then**
- 4 Compute $\tau^{tCG} = \arg \min_{\tau} \hat{m}_{\theta^t}(s_n + \tau p_n)$ s.t.
- 5 $\|s_n + \tau p_n\|_{\theta^t} = \Delta_t$;
- 6 Set $s_{n+1} = s_n + \tau^{tCG} p_n$;
- 7 **return** s_{n+1} ;
- 8 Compute $\alpha_n^{tCG} = \frac{\langle r_n, z_n \rangle_{\theta^t}}{\langle p_n, H_t[p_n] \rangle_{\theta^t}}$;
- 9 **if** $\|s_n + \alpha_n^{tCG} p_n\|_{\theta^t} > \Delta_t$ **then**
- 10 Compute $\tau^{tCG} = \arg \min_{\tau} \hat{m}_{\theta^t}(s_n + \tau p_n)$ s.t.
- 11 $\|s_n + \tau p_n\|_{\theta^t} = \Delta_t$;
- 12 Set $s_{n+1} = s_n + \tau^{tCG} p_n$;
- 13 **return** s_{n+1} ;
- 14 **else**
- 15 $s_{n+1} = s_n + \alpha_n^{tCG} p_n$;
- 16 $r_{n+1} = r_n + \alpha_n^{tCG} H_t[p_n]$;
- 17 $z_{n+1} = M_t r_{n+1}$;
- 18 $\beta_{n+1}^{tCG} = \frac{\langle z_{n+1}, z_{n+1} \rangle_{\theta^t}}{\langle z_n, z_n \rangle_{\theta^t}}$;
- 19 $p_{n+1} = -z_{n+1} + \beta_{n+1}^{tCG} p_n$;
- 20 **if** $\|r_{n+1}\|_{\theta^t} \leq \|r_0\|_{\theta^t} \min(\|r_0\|_{\theta^t}^{\delta}, \kappa)$ **then**
- 21 **return** s_{n+1} ;
- 22 **end**

C Additional numerical results

We present numerical results additional to Sect. 4 both for simulated clustering data and for real-world data.

C.1 Simulated data

See Tables 10, 11, and 12

Table 10 Simulation results of 20 runs for dimensions $d = 20$, number of components $K = 5$ and eccentricity $e = 5$

		EM	R-NTR	R-LBFGS
c = 0.2	Iterations	145.7	38.5	60.4
	Mean time (s)	1.899	1.512	8.165
	Mean ALL	-50.65	-50.66	-50.66
	MSE weights	0.00018	0.00019	0.008
	MSE means	0.08	0.08	0.08
	MSE cov	2.18	1.83	2.21
c = 1	Iterations	258.9	60.4	84.1
	Mean time (s)	3.061	1.951	11.64
	Mean ALL	-57.65	-57.65	-57.65
	MSE weights	0.00037	0.0004	0.008
	MSE means	0.08	0.07	0.07
	MSE cov	1.02	1.39	1.04
c = 5	Iterations	194.6	37.2	65.7
	Mean time (s)	2.367	1.267	8.927
	Mean ALL	-54.86	-54.86	-54.87
	MSE weights	0.00017	0.00018	0.008
	MSE means	0.07	0.07	0.07
	MSE cov	0.15	0.13	0.15

Table 11 Simulation results of 20 runs for dimensions $d = 40$, number of components $K = 5$ and eccentricity $e = 10$ with $N = 1000$ observations

		EM	R-NTR	R-LBFGS
c = 0.2	Iterations	4	16.3	6.6
	Mean time (s)	0.091	1.336	0.622
	Mean ALL	-125.63	-125.63	-125.649
	MSE weights	3e-05	3e-05	0.008
	MSE means	0.091	0.095	0.091
	MSE cov	0.224	0.215	0.224
c = 1	Iterations	4	17.6	8.8
	Mean time (s)	0.097	1.561	0.862
	Mean ALL	-110.112	-110.112	-110.14
	MSE weights	3e-05	3e-05	0.008
	MSE means	0.089	0.084	0.089
	MSE cov	0.461	0.337	0.461
c = 5	Iterations	4	17	7.2
	Mean time (s)	0.094	1.497	0.697
	Mean ALL	-116.276	-116.276	-116.3
	MSE weights	3e-05	3e-05	0.008
	MSE means	0.061	0.066	0.061
	MSE cov	0.351	0.43	0.351

Table 12 Simulation results of 20 runs for dimensions $d = 40$, number of components $K = 5$ and eccentricity $e = 10$ with $N = 10,000$ observations

		EM	R-NTR	R-LBFGS
$c = 0.2$	Iterations	4	2.3	3.3
	Mean time (s)	0.628	1.524	0.74
	Mean ALL	-127.762	-127.762	-127.764
	MSE weights	0	0	0.008
	MSE means	0.089	0.09	0.089
	MSE cov	0.14	0.189	0.14
$c = 1$	Iterations	4	2.8	3.8
	Mean time (s)	0.609	1.676	0.884
	Mean ALL	-112.25	-112.25	-112.252
	MSE weights	0	0	0.008
	MSE means	0.084	0.084	0.084
	MSE cov	0.313	0.29	0.313
$c = 5$	Iterations	4	2.6	3.6
	Mean time (s)	0.63	1.662	0.774
	Mean ALL	-118.337	-118.337	-118.339
	MSE weights	0	0	0.008
	MSE means	0.065	0.071	0.065
	MSE cov	0.228	0.235	0.228

C.2 Real-world datasets

C.2.1 Gas turbine CO and NOx emission data set (Kaya et al. 2019)

See Tables 13, 14, 15, 19, 17 18 and 19.

Table 13 Results of (normalized) Magic Gamma telescope data set for different number of components K . Number of observations $m = 36733$, dimensions $d = 11$

		EM	R-NTR	R-LBFGS
$K = 2$	Time (s)	2.04	2.48	3.22
	Iterations	25	22	28
	ALL	-5.73	-5.73	-5.73
$K = 5$	Time (s)	17.41	9.52	11.48
	Iterations	90	46	58
	ALL	-1.93	-1.99	-1.93
$K = 10$	Time (s)	40.66	29.36	27.31
	Iterations	130	61	75
	ALL	-1.17	-1.16	-1.17
$K = 15$	Time (s)	52.06	63.20	142.77
	Iterations	115	76	233
	ALL	0.04	-0.04	-0.05

Table 14 Results of (normalized) Wine quality data set for different number of components K . Number of observations $m = 6497$, dimensions $d = 11$

		EM	R-NTR	R-LBFGS
$K = 2$	Time (s)	0.54	0.24	1.55
	Iterations	27	8	20
	ALL	-11.02	-11.02	-11.02
$K = 5$	Time (s)	4.75	1.24	5.69
	Iterations	154	34	51
	ALL	-9.74	-9.98	-9.88
$K = 10$	Time (s)	12.61	5.65	22.80
	Iterations	239	83	100
	ALL	-9.23	-9.28	-9.28
$K = 15$	Time (s)	91.16	11.17	51.65
	Iterations	1137	70	147
	ALL	-8.91	-8.88	-8.89

C.2.2 Wine quality data set (Cortez et al. 2009)

The wine quality data set also provides classification labels: we can distinguish between white and red wine or distinguish between 7 quality labels. Besides the clustering performance of the methods (Table 14), we also show the goodness of fit of our method for $K = 2$ and $K = 7$.

Table 15 Weighted mean squared errors of (normalized) Wine data set for $K = 2$ (white wine/ red wine)

	EM	R-NTR	R-LBFGS
Distance	2.757966	2.757982	2.757974
wMSE weight	0.003169	0.003169	0.003169
wMSE mean	0.073776	0.073779	0.073778
wMSE cov	0.562106	0.562113	0.562109

Table 16 Adjusted Rand Index for (normalized) Wine data set for $K = 2$ (white wine/ red wine)

	Ground truth	EM	R-NTR	R-LBFGS
Ground truth	1	0.78	0.78	0.78
EM		1	1.00	1.00
R-NTR			1	1.00
R-LBFGS				1

Table 17 Weighted mean squared errors of (normalized) Wine data set for $K = 7$ (quality label). Weighting by mixing coefficient of the mixing coefficients belonging to the true labels

	EM	R-NTR	R-LBFGS
Distance	89.259333	89.488919	88.894978
wMSE weight	0.019019	0.010065	0.035496
wMSE mean	3.442393	4.214438	3.082073
wMSE cov	12.047944	13.400811	11.327231

Table 18 Weighted mean squared errors of (normalized) Wine data set for $K = 7$ (quality label). Weighting by mixing coefficients of the respective method

	EM	R-NTR	R-LBFGS
Distance	89.259333	89.488919	88.894978
wMSE weight	0.019019	0.010065	0.035496
wMSE mean	6.068149	5.617490	11.996024
wMSE cov	61.808676	37.060684	55.707085

Table 19 Adjusted rand index for (normalized) Wine data set for $K = 7$ (quality label)

	Ground truth	EM	R-NTR	R-LBFGS
Ground truth	1	0.02	0.01	0.02
EM		1	0.73	0.70
R-NTR			1	0.85
R-LBFGS				1

References

- Absil, P.-A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton (2008) ISBN 978-0-691-13298-3
- Alfö, M., Nieddu, L., Vicari, D.: A finite mixture model for image segmentation. *Stat Comput*, 18(2):137–150. <https://doi.org/10.1007/s11222-007-9044-9>
- Andrews, J.L.: Addressing overfitting and underfitting in Gaussian model-based clustering. *Comput Stat Data Anal* **127**, 160–171 (2018)
- Arthur, D., Vassilvitskii, S.: k-means++: the advantages of careful seeding. In: SODA '07: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1027–1035, Philadelphia (2007). Society for Industrial and Applied Mathematics. ISBN 978-0-898716-24-5
- Articus, C., Burgard, J.P.: A finite mixture fay herriot-type model for estimating regional rental prices in Germany. Research Papers in Economics 2014-14, University of Trier, Department of Economics, 2014. <https://ideas.repec.org/p/trr/wpaper/201414.html>
- Bhatia, R.: Positive Definite Matrices. Princeton Series in Applied Mathematics, Princeton University Press, Princeton (2007)
- Bishop, C.M.: Pattern Recognition and Machine Learning (Information Science and Statistics). Springer, Berlin, Heidelberg (2006)
- Bock, R., Chilingarian, A., Gaug, M., Hakl, F., Hengstebeck, T., Jiřina, M., Klaschka, J., Kotř, E., Savický, P., Towers, S., Vaiculis, A., Wittek, W.: Methods for multidimensional event classification: a case study using images from a cherenkov gamma-ray telescope.

- Nucl. Instrum. Methods Phys. Res. Sect. A Acceler. Spectrom. Detect. Assoc. Equip. **516**(2), 511–528 (2004). <https://doi.org/10.1016/j.nima.2003.08.157>
- Boumal, N.: An introduction to optimization on smooth manifolds (2020). <http://www.nicolasboumal.net/book>
- Carmo, M.P.d. (1992) Riemannian Geometry. Mathematics Theory and Applications. Birkhäuser, Boston. ISBN 0-8176-3490-8
- Compiani, G., Kitamura, Y.: Using mixtures in econometric models: a brief review and some new results. *Econom. J.* **19**(3), C95–C127, 10 (2016). <https://doi.org/10.1111/ectj.12068>
- Conn, A.R., Gould, N.I., Toint, P.L.: Trust-region methods. MPS-SIAM series on optimization. SIAM Society for Industrial and Applied Mathematics, Philadelphia (2000)0898714605
- Coretto, P.: Estimation and computations for gaussian mixtures with uniform noise under separation constraints. *Stat. Methods Appl.*, 07 2021. 10.1007/s10260-021-00578-2
- Cortez, P., Cerdeira, A., Almeida, F., Matos, T., Reis, J.: Modeling wine preferences by data mining from physicochemical properties. *Decis. Support Syst.*, 47(4):547–553, (2009). <https://doi.org/10.1016/j.dss.2009.05.016>
- Smart Business Networks: Concepts and Empirical Evidence
- Dasgupta, S.: Learning mixtures of gaussians. In: Proceedings of the 40th Annual Symposium on Foundations of Computer Science, FOCS '99, pp. 634 (1999). IEEE Computer Society. ISBN 0769504094
- Dresvyanskiy, D., Karaseva, T., Makogin, V., Mitrofanov, S., Redenbach, C., Spodarev, E.: Detecting anomalies in fibre systems using 3-dimensional image data. *Stat. Comput.* **30**(4), 817–837 (2020)
- Dua, D., Graff, C.: UCI machine learning repository. <http://archive.ics.uci.edu/ml> (2017)
- Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press. <http://www.deeplearningbook.org> (2016)
- Gould, N., Orban, D., Sartenaer, A., Toint, P.: Sensitivity of trust-region algorithms to their parameters. 4OR Q J Belgian French Italian Oper. Res. Soc. **3**(227–241),(2005). <https://doi.org/10.1007/s10288-005-0065-y>
- Gross, M., Rendtel, U., Schmid, T., Schmon, S., Tzavidis, N.: Estimating the Density of Ethnic Minorities and Aged People in Berlin (2015)
- Heidel, G., Schulz, V.: A Riemannian trust-region method for low-rank tensor completion. *Numer. Linear Algebra Appl.* **25**, 6 (2018)
- Hosseini, R., Mash'al, M.: Mixest: An estimation toolbox for mixture models. *CoRR*, abs/1507.06065. <http://arxiv.org/abs/1507.06065> (2015)
- Hosseini, R., Sra, S.: Matrix manifold optimization for gaussian mixtures. In: Advances in Neural Information Processing Systems, vol. 28. Curran Associates, Inc. (2015)
- Hosseini, R., Sra, S.: An alternative to EM for gaussian mixture models: batch and stochastic Riemannian optimization. *Math. Program.* **181**(1), 187–223 (2020)
- Hubert, L., Arabie, P.: Comparing partitions. *J Classif* **2**(1), 193–218 (1985)
- Jeuris, B., Vandebril, R., Vandereycken, B.: A survey and comparison of contemporary algorithms for computing the matrix geometric mean. *ETNA* **39**, 379–402 (2012)
- Kaya, H., Tufekci, P.: Local and global learning methods for predicting power of a combined gas and steam turbine (2012)
- Kaya, H., Tufekci, P., Uzun, E.: Predicting co and nox emissions from gas turbines: novel data and a benchmark pems. *Turk. J. Electr. Eng. Comput. Sci.* **27**, 4783–4796 (2019)
- Lee, S., McLachlan, G.: On mixtures of skew normal and skew [tex equation: t]-distributions. *Adv. Data Anal. Classif.* **7**, 09 (2013). <https://doi.org/10.1007/s11634-013-0132-8>
- Li, Y., Li, L.: A novel split and merge em algorithm for gaussian mixture model. In: 2009 Fifth International Conference on Natural Compu-

- tation, vol. 6, pp. 479–483 (2009). <https://doi.org/10.1109/ICNC.2009.625>
- Ma, J., Xu, L., Jordan, M.: Asymptotic convergence rate of the em algorithm for gaussian mixtures. *Neural Comput.* **12**, 2881–2907 (2000)
- McLachlan, G.J., Lee, S.X., Rathnayake, S.I.: Finite mixture models. *Ann. Rev. Stat. Appl.* **6**(1), 355–378 (2019). <https://doi.org/10.1146/annurev-statistics-031017-100325>
- Morales, J., Nocedal, J.: Automatic preconditioning by limited memory quasi-newton updating. *SIAM J. Optim.* **10**, 1079–109606 (2000)
- Murphy, K.P.: *Machine Learning: A Probabilistic Perspective*. MIT Press, Cambridge (2013)
- Naim, I., Gildea, D.: Convergence of the em algorithm for gaussian mixtures with unbalanced mixing coefficients. In: *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML'12*, , Madison, pp. 1427–1431 (2012). Omnipress. ISBN 9781450312851
- Ormonet, D., Tresp, V.: Improved gaussian mixture density estimates using bayesian penalty terms and network averaging. In: *Proceedings of the 8th International Conference on Neural Information Processing Systems, NIPS'95*, pp. 542–548. MIT Press, Cambridge (1995)
- Salakhutdinov, R., Roweis, S., Ghahramani, Z.: Optimization with em and expectation-conjugate-gradient. In: *Proceedings of the Twentieth International Conference on International Conference on Machine Learning, ICML'03*, pp. 672–679. AAAI Press (2003). ISBN 1577351894
- Sartenaer, A.: Automatic determination of an initial trust region in nonlinear programming. *SIAM J. Sci. Comput.* **18**(6), 1788–1803 (1997)
- Scott, D.: *Multivariate Density Estimation. Theory, Practice, and Visualization*. Wiley Series in Probability and Mathematical Statistics. Wiley, New York, London (1992)
- Sembach, L., Burgard, J.P., Schulz, V.H.: A Riemannian newton trust-region method for fitting gaussian mixture models (2021)
- Snoussi, H., Mohammad-Djafari, A.: Penalized maximum likelihood for multivariate gaussian mixture. *AIP Conf. Proc.* **617**(1), 36–46 (2002)
- Sra, S., Hosseini, R.: Conic geometric optimization on the manifold of positive definite matrices. *SIAM J. Optim.* **25**, 713–739 (2015)
- Townsend, J., Koep, N., Weichwald, S.: Pymanopt: a python toolbox for optimization on manifolds using automatic differentiation. *J. Mach. Learn. Res.* **17**(137), 1–5 (2016)
- Vanderbei, R.J., Benson, H.Y.: On formulating semidefinite programming problems as smooth convex nonlinear optimization problems. *Comput. Optim. Appl.* **13**(1), 231–252 (1999)
- Wu, B., Mcgrory, C.A., Pettitt, A.N.: A new variational Bayesian algorithm with application to human mobility pattern modeling. *Stat Comput* **22**(1), 185–203 (2012). <https://doi.org/10.1007/s11222-010-9217-9>
- Xu, L., Jordan, M.I.: On convergence properties of the em algorithm for gaussian mixtures. *Neural Comput.* **8**(1), 129–151 (1996)
- Zoran, D., Weiss, Y.: Natural images, Gaussian mixtures and dead leaves. *Adv. Neural Inf. Process. Syst.* **3**, 1736–1744 (2012)

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