



Analysis of a generalised expectation–maximisation algorithm for Gaussian mixture models: a control systems perspective

Sarthak Chatterjee, Orlando Romero & Sérgio Pequito

To cite this article: Sarthak Chatterjee, Orlando Romero & Sérgio Pequito (2022) Analysis of a generalised expectation–maximisation algorithm for Gaussian mixture models: a control systems perspective, International Journal of Control, 95:10, 2734-2742, DOI: [10.1080/00207179.2021.1931964](https://doi.org/10.1080/00207179.2021.1931964)

To link to this article: <https://doi.org/10.1080/00207179.2021.1931964>



Published online: 28 May 2021.



Submit your article to this journal [↗](#)



Article views: 106



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)



Analysis of a generalised expectation–maximisation algorithm for Gaussian mixture models: a control systems perspective

Sarthak Chatterjee ^a, Orlando Romero ^b and Sérgio Pequito ^b

^aDepartment of Electrical, Computer, and Systems Engineering Rensselaer Polytechnic Institute, Troy, NY, USA; ^bDepartment of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA

ABSTRACT

The expectation–maximisation (EM) algorithm is one of the most popular methods used to solve the problem of parametric distribution-based clustering in unsupervised learning. In this paper, we propose to analyse a generalised EM (GEM) algorithm in the context of Gaussian mixture models, where the maximisation step in the EM is replaced by an increasing step. We show that this GEM algorithm can be understood as a linear time-invariant (LTI) system with a feedback nonlinearity. Therefore, we explore some of its convergence properties by leveraging tools from robust control theory. Lastly, we explain how the proposed GEM can be designed and present a pedagogical example to understand the advantages of the proposed approach.

ARTICLE HISTORY

Received 25 March 2020
Accepted 12 May 2021

KEYWORDS

Statistical data analysis; linear multivariable systems; output regulation; Robust control applications

1. Introduction

A fundamental problem in unsupervised learning is the problem of *clustering*, where the task in question is to group certain objects of interest into subgroups called *clusters*, such that all objects in a particular cluster share features (in some predefined sense) with each other, but not with objects in other clusters (Bishop, 2006; Tan et al., 2005).

The expectation–maximisation (EM) algorithm (Dempster et al., 1977) is one of the most commonly used methods in parametric distribution-based clustering analysis (Nowak, 2003). Notably, Gaussian mixture models (GMMs) (i.e. a finite convex combination of multivariate Gaussian distributions) have found several applications in real-world problems (Bishop, 2006; Tan et al., 2005). In this setting, clustering consists of estimating the parameters in a GMM that maximise its likelihood function (iteratively maximised through the EM algorithm), followed by assigning to each data point the ‘cluster’ corresponding to its most likely multivariate Gaussian distribution in the GMM.

The convergence of the EM algorithm is well-studied in the literature (Wu, 1983), particularly in the context of determining the parameters of GMMs (Xu & Jordan, 1996). Nonetheless, it is worth analysing the EM algorithm as a dynamical system and possibly gain insights that enable us to design more efficient variations of the EM algorithm. Therefore, in Romero et al. (2019), the authors proposed to change the perspectives on local optimisers and convergence of the EM algorithm by assessing, respectively, the equilibria and asymptotic stability (in the sense of Lyapunov) of a nonlinear dynamical system that represents the standard EM algorithm, through explicit use of discrete-time Lyapunov stability theory.

In this paper, we build upon the recent work in Romero et al. (2019) and propose to analyse a generalised EM (GEM) algorithm (Dempster et al., 1977; Neal & Hinton, 1998) in the context of Gaussian mixture models, where the maximisation step in the EM is replaced by an increasing step. GEM algorithms have also been used in applications such as computer vision (Fessler & Hero, 1995) and noise estimation in communication channels (Krisjansson et al., 2001), and, in general, the study of the EM algorithm and its myriad variants constitute an active area of research (Moon, 1996; Roche, 2011). The main contributions of this work are as follows. First, we show that this GEM algorithm can be understood as a linear time-invariant (LTI) system with a feedback nonlinearity. Secondly, we explore some of its convergence properties by leveraging tools from robust control theory. Lastly, we explain how the proposed GEM can be designed, and present a pedagogical example to understand the advantages of the proposed approach.

2. Problem statement

Let $\theta \in \Theta \subseteq \mathbb{R}^p$ be some vector of unknown (but deterministic) parameters characterising a distribution of interest, which we seek to infer from a collected dataset $y \in \mathbb{R}^d$ (from now on assumed fixed) and a statistical model composed by a family of joint probability density or mass functions (possibly mixed) $p_\theta(x, y)$ indexed by $\theta \in \Theta$, where

$$x \in \mathcal{X} := \{x \in \mathbb{R}^n : p_\theta(x, y) > 0\} \quad (1)$$

is some latent (hidden) random vector.

The EM algorithm seeks to find a local maximiser of the *incomplete* likelihood function $\mathcal{L} : \Theta \rightarrow \mathbb{R}$ given by

$$\mathcal{L}(\theta) := p_\theta(y) = \begin{cases} \int_{\mathcal{X}} p_\theta(x, y) dx, & \text{if } x \text{ is continuous,} \\ \sum_{x \in \mathcal{X}} p_\theta(x, y), & \text{if } x \text{ is discrete.} \end{cases} \quad (2)$$

The mapping $\theta \mapsto p_\theta(x, y)$ is, naturally, referred to as the *complete* likelihood function. To optimise $\mathcal{L}(\theta)$, the EM algorithm alternates at each iteration k between two steps. First, in the *expectation* step (E-step), we compute $Q(\theta, \theta^{(k)})$, defined through

$$Q(\theta, \theta') := \mathbb{E}_{p_{\theta'(\cdot|y)}}[\log p_\theta(\cdot, y)] \quad (3a)$$

$$= \int_{\mathcal{X}} p_{\theta'}(x|y) \log p_\theta(x, y) dx, \quad (3b)$$

so that $Q(\cdot, \theta^{(k)})$ denotes the expected value of the complete log-likelihood function with respect to $\theta = \theta^{(k)}$. Second, in the *maximisation* step (M-step), we maximise $Q(\cdot, \theta^{(k)})$ and update the current estimate as that maximiser.

Before formally stating the EM algorithm, let us make some mild simplifying assumptions that will avoid pathological behaviour on the Q -function, $Q : \Theta \times \Theta \rightarrow \mathbb{R}$.

Assumption 2.1: \mathcal{X} does not depend on $\theta \in \Theta$ and has positive Lebesgue measure.

Assumption 2.2: \mathcal{L} is twice continuously differentiable in Θ .

Notice that, from Assumption 2.1, the conditional distribution $p_{\theta'}(x|y) = p_{\theta'}(x, y)/p_{\theta'}(y)$ is well defined in \mathcal{X} , since $p_{\theta'}(y) > 0$ for every $\theta' \in \Theta$. Finally, we make the following simplifying assumption, which makes the M-step well defined.

Assumption 2.3: $Q(\cdot, \theta')$ has a unique global maximiser in Θ .

With all these ingredients and assumptions, we summarise the EM algorithm in Algorithm 1.

However, it is to be kept in mind that when we implement the EM algorithm, for most parametric distributions, we do not obtain a closed-form expression for the M-step. As a consequence, to determine a solution (i.e. an approximation) in the M-step, we need to rely on numerical optimisation schemes. For

Algorithm 1 Expectation- -maximisation (EM).

Input: $y \in \mathbb{R}^d$, $p_\theta(x, y)$, $\theta^{(0)} \in \Theta$.

Output: $\hat{\theta}$.

- 1: **for** $k = 0, 1, 2, \dots$ (until some stopping criterion) **do**
 - 2: **E-step:** Compute $Q(\theta, \theta^{(k)})$
 - 3: **M-step:** Determine $\theta^{(k+1)} = \operatorname{argmax}_{\theta \in \Theta} Q(\theta, \theta^{(k)})$
 - 4: **end for**
 - 5: **return** $\hat{\theta} =$ last computed iteration in $\{\theta^{(k)}\}$.
-

instance, we can consider first-order optimisation algorithms (e.g. gradient ascent), i.e.

$$\theta^{(k+1)} = \theta^{(k)} + \eta \left. \frac{\partial Q(\theta, \theta^{(k)})}{\partial \theta} \right|_{\theta = \theta^{(k)}}, \quad (4)$$

for some $\eta > 0$. Notice that this could constitute a problem by itself since first-order algorithms are known to have slow convergence rates that get aggravated by the increase in the dimension of the search space. Furthermore, any variant of Algorithm 1 that does not explicitly maximise $Q(\cdot, \theta^{(k)})$ at the M-step, but instead is such that $Q(\theta^{(k+1)}, \theta^{(k)}) > Q(\theta^{(k)}, \theta^{(k)})$ is referred to as a *generalised* EM (GEM) algorithm.

As previously mentioned, a particularly important class of models are the Gaussian mixture models (GMMs). In these models, each component of the mixture is given by

$$p_{\theta_i}(y) = \frac{\alpha_i}{\sqrt{\det(2\pi \Sigma_i)}} e^{-\frac{1}{2}(y-\mu_i)^\top \Sigma_i^{-1}(y-\mu_i)}, \quad (5)$$

with $i = 1, 2, \dots, K$, $y, \mu_i \in \mathbb{R}^d$, $\Sigma_i \in \mathbb{R}^{d \times d}$ is positive definite, and $\alpha_i \in [0, 1]$ such that $\sum_{i=1}^K \alpha_i = 1$. The vector of unknown parameters θ lumps together the scalar parameters within $\alpha_i, \mu_i, \Sigma_i$ for $i \in \{1, \dots, K\}$, as follows:

$$\theta = [\alpha^\top, \mu^\top, \operatorname{vec}[\Sigma]^\top]^\top, \quad (6)$$

where $\alpha^\top = [\alpha_1, \dots, \alpha_K]^\top$, $\mu^\top = [\mu_1, \dots, \mu_K]^\top$, and $\operatorname{vec}[\Sigma] = [\operatorname{vec}[\Sigma_1]^\top, \dots, \operatorname{vec}[\Sigma_K]^\top]^\top$, with $\operatorname{vec}(M)$ denoting the vector obtained by stacking the column vectors of M .

In this setting, an alternative is to replace the M-step by (4), and we obtain a GEM that is able to recover similar (asymptotic) convergence rates available in the literature (Balakrishnan et al., 2017). Nonetheless, (asymptotic) convergence rates can be misleading as they do not reflect the practical number of iterations required to converge. Furthermore, as it is clear in the GMM, there are some additional constraints that are implicit and are not necessarily satisfied by (4) (i.e. $\alpha_1 + \dots + \alpha_K = 1$ and $\Sigma_i > 0$ for $i = 1, \dots, K$).

That said, we need to further understand the transient and the local behaviour of the GEM algorithm, for which dynamical systems theory provides us with the proper framework. Subsequently, in this paper, we propose to step away from the dynamics without an explicit control (e.g. the M-step in Algorithm 1), towards one where we can consider an additive control, and therefore, study its properties.

In summary, we seek to address the following questions.

Problem 2.1: (1) Is it possible to replace the M-step in Algorithm 1 by a parameter update step given by

$$\theta^{(k+1)} = \theta^{(k)} + u^{(k)}, \quad (7)$$

where we can design a feedback control law $u^{(k)} = \phi(\theta^{(k)})$ to obtain a GEM algorithm?

(2) What insights (particularly with respect to design) can such control laws provide us with?

3. Main results

In this section, we provide the main result of the paper. Specifically, we show how we can leverage tools from control systems

theory to analyse a GEM algorithm as an LTI system connected in feedback with a nonlinearity. Furthermore, we also show how to derive the convergence rate for such an algorithm using tools from robust control. Lastly, we briefly describe how we can look into certain aspects of designing new GEM-like algorithms.

3.1 GEM algorithms as LTI systems with a feedback nonlinearity

We first show how we can leverage tools from dynamical systems and control theory to cast a GEM algorithm into the framework of an LTI system with an interconnected feedback nonlinearity. We begin with the following lemma that provides us with expressions for the closed-form solution of the problem of estimating the parameters of a GMM using a generalised EM algorithm.

Lemma 3.1 (Dempster et al. (1977)): *Given K possible mixtures in the GMM, and independently and identically distributed (i.i.d.) samples $\{x^{(t)}\}_{t=1}^N$, we can estimate the parameter vector θ by maximising the log-likelihood $\mathcal{L}(\theta)$, that, in the context of a GMM has a closed-form solution given as follows:*

$$\alpha_j^{(k+1)} = \frac{1}{N} \sum_{t=1}^N h_j^{(k)}(t), \quad (8)$$

$$\mu_j^{(k+1)} = \frac{1}{\sum_{t=1}^N h_j^{(k)}(t)} \sum_{t=1}^N h_j^{(k)}(t) x^{(t)}, \quad (9)$$

and

$$\Sigma_j^{(k+1)} = \frac{1}{\sum_{t=1}^N h_j^{(k)}(t)} \sum_{t=1}^N h_j^{(k)}(t) z_j^{(t),(k+1)} (z_j^{(t),(k+1)})^\top, \quad (10)$$

with

$$z_j^{(t),(k+1)} = x^{(t)} - \mu_j^{(k+1)}, \quad (11)$$

where the posterior probabilities $h_j^{(k)}$ are given by

$$h_j^{(k)}(t) = \frac{\alpha_j^{(k)} p(x^{(t)} | \mu_j^{(k)}, \Sigma_j^{(k)})}{\sum_{i=1}^K \alpha_i^{(k)} p(x^{(t)} | \mu_i^{(k)}, \Sigma_i^{(k)})}. \quad (12)$$

With the above closed-form solution, if we, instead, consider a ‘shifted-update’ of the covariance as

$$\Sigma_j^{(k+1)} = \frac{1}{\sum_{t=1}^N h_j^{(k)}(t)} \sum_{t=1}^N h_j^{(k)}(t) z_j^{(t),(k)} (z_j^{(t),(k)})^\top \quad (13)$$

(i.e. the update of $\Sigma_j^{(k+1)}$ is done with respect to $\mu_j^{(k)}$ instead of $\mu_j^{(k+1)}$), we can summarise in the following lemma the relationships between the updates of the parameters of the GMM that we aim to estimate, i.e. the mixing weights, the means, and the covariance matrices.

Lemma 3.2: *For the shifted updates of the covariance matrices considered in (13), the following relations hold:*

$$\alpha^{(k+1)} - \alpha^{(k)} = P_{\alpha^{(k)}} \frac{\partial \mathcal{L}}{\partial \alpha} \Big|_{\alpha=\alpha^{(k)}}, \quad (14)$$

$$\mu_j^{(k+1)} - \mu_j^{(k)} = P_{\mu_j^{(k)}} \frac{\partial \mathcal{L}}{\partial \mu_j} \Big|_{\mu_j=\mu_j^{(k)}}, \quad (15)$$

and

$$\text{vec}[\Sigma_j^{(k+1)}] - \text{vec}[\Sigma_j^{(k)}] = P_{\Sigma_j^{(k)}} \frac{\partial \mathcal{L}}{\partial \text{vec}[\Sigma_j]} \Big|_{\Sigma_j=\Sigma_j^{(k)}}, \quad (16)$$

with

$$P_{\alpha^{(k)}} = \frac{1}{N} (\text{diag}[\alpha_1^{(k)}, \dots, \alpha_K^{(k)}] - \alpha^{(k)} \alpha^{(k)\top}), \quad (17)$$

$$P_{\mu_j^{(k)}} = \frac{1}{\sum_{t=1}^N h_j^{(k)}(t)} \Sigma_j^{(k)}, \quad (18)$$

$$P_{\Sigma_j^{(k)}} = \frac{2}{\sum_{t=1}^N h_j^{(k)}(t)} (\Sigma_j^{(k)} \otimes \Sigma_j^{(k)}), \quad (19)$$

where $j \in \{1, \dots, K\}$ denotes the indices of the mixture components, k denotes the iteration number, and \otimes denotes the Kronecker product.

Therefore, by combining Equations (14)–(16) of Lemma 3.2, we can briefly write the evolution of the parameters as

$$\theta^{(k+1)} = \theta^{(k)} + P(\theta^{(k)}) \nabla \mathcal{L}(\theta^{(k)}), \quad (20)$$

where

$$P(\theta) = \text{diag}[P_{\alpha}, P_{\mu_1}, \dots, P_{\mu_K}, P_{\Sigma_1}, \dots, P_{\Sigma_K}]. \quad (21)$$

In this case, the term $\phi(\theta^{(k)}) = P(\theta^{(k)}) \nabla \mathcal{L}(\theta^{(k)})$ could be understood as a nonlinearity driving the system. Nonetheless, it is not guaranteed that some essential implicit constraints hold, i.e.

$$\theta^{(k+1)} \in \Theta = \left\{ \theta : \sum_{j=1}^K \alpha_j = 1, \Sigma_j = \Sigma_j^\top > 0 \right\}. \quad (22)$$

Therefore, towards incorporating these inter-dependencies, we can consider the ‘shifted’ subspace

$$\Theta_s = \left\{ \theta' : \sum_{j=1}^K \alpha_j = 0 \right\}, \quad (23)$$

where $\theta' = \theta - \theta_0 \in \Theta$ for a shift θ_0 . Furthermore, let the coordinates of θ' under the basis $\{e_1, \dots, e_m\}$ be denoted by θ_c , where the e_i -s are canonical (orthonormal) basis vectors and m is the dimension of Θ . Then, $\theta - \theta_0 = E\theta_c$, or equivalently, $\theta = E\theta_c + \theta_0$, where $E = [e_1, \dots, e_m]$.

Now, notice that $E^\top \theta = E^\top E \theta_c + E^\top \theta_0$ (or, equivalently, $\theta_c = E^\top \theta - E^\top \theta_0$), by multiplying both sides by E^\top , and noticing that $E^\top E = I$. Thus, $\theta' = \theta - \theta_0 = E^\top E (\theta - \theta_0) = E E^\top \theta' \in \Theta$, by observing that Θ is an open convex set since we only consider local differential properties of the log-likelihood, and consequently, the constraint on positive definiteness of Σ_j holds.

Therefore,

$$\begin{aligned} \theta^{(k+1)} &= E E^\top (\theta^{(k)} + P(\theta^{(k)}) \nabla \mathcal{L}(\theta^{(k)})) \\ &= \theta^{(k)} + E E^\top P(\theta^{(k)}) \nabla \mathcal{L}(\theta^{(k)}) \end{aligned} \quad (24)$$

Algorithm 2 Projection-based GEM (PB-GEM).**Input:** $y \in \mathbb{R}^d$, $p_\theta(y)$, $\theta^{(0)} \in \Theta$.**Output:** $\hat{\theta}$.

- 1: **for** $k = 0, 1, 2, \dots$ (until some stopping criterion) **do**
- 2: $\theta^{(k+1)} = \theta^{(k)} + EE^\top P(\theta^{(k)}) \frac{\partial \mathcal{L}}{\partial \theta} \Big|_{\theta=\theta^{(k)}}$
- 3: **end for**
- 4: **return** $\hat{\theta} =$ last computed iteration in $\{\theta^{(k)}\}$.

belongs to Θ , which constitutes the parameter update of a GEM algorithm that we shall refer to as *projection-based GEM* (PB-GEM) – see Algorithm 2.

Consequently, the term $\phi(\theta^{(k)}) = EE^\top P(\theta^{(k)}) \frac{\partial \mathcal{L}}{\partial \theta} \Big|_{\theta=\theta^{(k)}}$ can be understood as a nonlinearity driving a linear time-invariant (LTI) system. As such, we can consider $\phi(\theta) = \nabla f(\theta)$ to be (locally) Lipschitz and for which there is a (locally) strongly convex function f . Before stating the theorem that shows the rate of convergence for the PB-GEM algorithm, we introduce the following preliminary definitions and results.

Definition 3.3 (Q-convergence (Jay, 2001)): Given a sequence $\{\theta^{(k)}\} \rightarrow \theta^*$ with $\theta^{(k)} \neq \theta^*$ for $k = 0, 1, 2, \dots$, the *order of convergence* β is a nonnegative number satisfying

$$\limsup_{k \rightarrow \infty} \frac{\|\theta^{(k+1)} - \theta^*\|}{\|\theta^{(k)} - \theta^*\|^\beta} = \rho < \infty, \quad (25)$$

with ρ being the *rate* of convergence.

Definition 3.4 (Sector Integral Quadratic Constraint (IQC) for the gradient map): For a *strongly convex* function f with strong convexity parameter μ_f , having *Lipschitz continuous* gradients with Lipschitz constant L , the gradient map ∇f satisfies the sector IQC defined by

$$\begin{aligned} & \begin{bmatrix} \theta - \theta^* \\ \nabla f(\theta) - \nabla f(\theta^*) \end{bmatrix}^\top \begin{bmatrix} -2\mu_f L I & (L + \mu_f) I \\ (L + \mu_f) I & -2I \end{bmatrix} \\ & \times \begin{bmatrix} \theta - \theta^* \\ \nabla f(\theta) - \nabla f(\theta^*) \end{bmatrix} \geq 0 \end{aligned} \quad (26)$$

for all θ, θ^* .

Lemma 3.5 (A modified version of Theorem 4 in Lessard et al. (2016)): Consider a first-order linear optimisation scheme represented as the dynamical system

$$\xi[k+1] = A\xi[k] + Bu[k], \quad (27a)$$

$$\theta[k] = C\xi[k] + Du[k], \quad (27b)$$

$$u[k] = \phi(\theta[k]), \quad (27c)$$

with nonlinearity $\phi(\theta) = \nabla f(\theta)$. If ∇f satisfies the sector IQC defined by (26), then the linear matrix inequality (LMI)

$$\begin{aligned} & \begin{bmatrix} A^\top \\ B^\top \end{bmatrix} R \begin{bmatrix} A^\top \\ B^\top \end{bmatrix}^\top - \begin{bmatrix} \rho^2 R & 0 \\ 0 & 0 \end{bmatrix} + \lambda \begin{bmatrix} C & D \\ 0 & I \end{bmatrix}^\top \\ & \times \begin{bmatrix} -2\mu_f L I & (L + \mu_f) I \\ (L + \mu_f) I & -2I \end{bmatrix} \begin{bmatrix} C & D \\ 0 & I \end{bmatrix} \leq 0 \end{aligned} \quad (28)$$

is feasible for some $R > 0$, $\lambda \geq 0$. Specifically, $\{\xi[k]\} \rightarrow \xi^*$ with respect to a suitable norm $\|\cdot\|$, with a convergence rate of ρ , where ξ^* is a fixed point of (27a) satisfying $\xi^* = A\xi^*$.

With the above ingredients, we are ready to state our main result concerning the convergence rate of the PB-GEM algorithm, which builds upon tools from robust control theory.

Theorem 3.6: Consider a function $f(\theta)$ that is μ_f -strongly convex, has an L -Lipschitz gradient, and satisfies $\nabla f(\theta) = EE^\top P(\theta) \nabla \mathcal{L}(\theta)$. Then, $\theta^{(k+1)} = \theta^{(k)} + u^{(k)}$, with $u^{(k)} = \nabla f(\theta^{(k)})$ is a GEM algorithm (i.e. $\{\theta^{(k)}\} \rightarrow \theta^*$, where θ^* is the maximum-likelihood estimate) with convergence rate ρ bounded by

$$\rho \leq \max\{|1 - \mu_f|, |1 - L|\}. \quad (29)$$

Proof: That the Projection-Based GEM algorithm presented in Algorithm 2 indeed constitutes a generalised EM can be shown using an argument similar to one presented in Salakhutdinov et al. (2003). In particular, if Assumption 2.3 is satisfied for models of the exponential family (a special case being the GMMs considered in this paper), the PB-GEM algorithm evolves in a way such that we have $\mathcal{L}(\theta^{(k+1)}) > \mathcal{L}(\theta^{(k)})$ for all $k \in \mathbb{Z}_+$, provided $\nabla \mathcal{L}(\theta^{(0)}) \neq 0$. Secondly, we notice that the iterative scheme can be represented as the LTI system in (27a) with a feedback nonlinearity given by $\phi(\cdot) = \nabla f(\cdot)$.

Due to the regularity of f and the fact that the PB-GEM algorithm can be represented as the dynamical system (27a), with $A = B = C = I$ and $D = 0$, we can invoke the results of Definition 3.4 and Lemma 3.5 to recover bounds on the convergence rate of the PB-GEM algorithm using the LMI in (28). Remarkably, due to the general block-diagonal structure of optimisation algorithms like gradient ascent, we can then use a ‘lossless dimensionality reduction argument’ and reduce the case of the feasibility of the above LMI to analyse the corresponding semidefinite program for the single-dimensional case without loss of generality (Lessard et al., 2016).

This ascertains the local convergence for the maximum of the function f as long as the following LMI holds:

$$\begin{bmatrix} (1 - \rho^2)R & R \\ R & R \end{bmatrix} + \lambda \begin{bmatrix} -2\mu_f L & L + \mu_f \\ L + \mu_f & -2 \end{bmatrix} \leq 0, \quad (30)$$

for some scalar $R > 0$, and $\lambda \geq 0$, where $\rho \in (0, 1)$ denotes the convergence rate. Since R is a scalar, we can consider $R = 1$ without loss of generality. This gives us the LMI

$$\begin{bmatrix} 1 - \rho^2 - 2\mu_f L \lambda & 1 + \lambda(L + \mu_f) \\ 1 + \lambda(L + \mu_f) & 1 - 2\lambda \end{bmatrix} \leq 0. \quad (31)$$

As a consequence, to ensure the negative semidefiniteness of the above matrix, both $1 - 2\lambda$ (which is present in the bottom right block) and the Schur complement of the bottom right block need to be negative semidefinite. Thus, we have

$$\lambda \geq \frac{1}{2}, \quad (32a)$$

$$0 \geq 1 - \rho^2 - 2\mu_f L \lambda - \frac{(\lambda(L + \mu_f) + 1)^2}{1 - 2\lambda}. \quad (32b)$$

Combining these two, we have

$$\rho^2 \geq 1 - 2\mu_f L \lambda - \frac{(\lambda(L + \mu_f) + 1)^2}{1 - 2\lambda}, \quad (33)$$

which yields $\rho \leq \max\{|1 - \mu_f|, |1 - L|\}$. ■

Additionally, the transformation matrix $P(\cdot)$ also provides us with valuable insights regarding the rate of convergence of the PB-GEM algorithm. Indeed, differentiating the equation

$$F^{\text{PB-GEM}}(\theta) = \theta + EE^\top P(\theta) \nabla \mathcal{L}(\theta), \quad (34)$$

we have

$$\frac{\partial F^{\text{PB-GEM}}}{\partial \theta}(\theta) = I + EE^\top \frac{\partial P}{\partial \theta}(\theta) \overline{\nabla} \mathcal{L}(\theta) + EE^\top P(\theta) S(\theta), \quad (35)$$

where $\frac{\partial P}{\partial \theta} = \begin{bmatrix} \frac{\partial P}{\partial \theta^1} & \dots & \frac{\partial P}{\partial \theta^m} \end{bmatrix}$ with $\theta = [\theta^1, \dots, \theta^m]$,

$$\overline{\nabla} \mathcal{L}(\theta) = \begin{bmatrix} \nabla \mathcal{L}(\theta) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \nabla \mathcal{L}(\theta) \end{bmatrix}, \quad (36)$$

and $S(\theta) = \nabla^2 \mathcal{L}(\theta)$ denotes the Hessian matrix of $\mathcal{L}(\cdot)$.

Therefore, near a stationary point $\theta = \theta^*$ of $\mathcal{L}(\theta)$ over which $P(\theta)$ is bounded, we have

$$P(\theta) \approx (EE^\top)^{-1} \left(\frac{\partial F^{\text{PB-GEM}}}{\partial \theta}(\theta) - I \right) [S(\theta)]^{-1}. \quad (37)$$

As a consequence, it follows that, under the aforementioned conditions, the projection-based GEM algorithm exhibits superlinear convergence when $\nabla \mathcal{L}(\theta)$ approaches zero. In particular, the nature of convergence is dictated by the eigenvalues of the matrix $\frac{\partial F^{\text{PB-GEM}}}{\partial \theta}(\theta)$. If the eigenvalues are near zero, then the transformation matrix scales the EM update step by approximately the scaled negative inverse Hessian, and the EM algorithm behaves like Newton's method. On the other hand, if the eigenvalues are near unity (in absolute value), then the PB-GEM algorithm exhibits first-order convergence.

3.2 Towards the design of GEM algorithms

We can, therefore, propose to design a GEM algorithm by changing the control law. Nonetheless, we have to be careful with the updates on the different parameters as, implicitly, they possess constraints on the updates. Specifically, we require the α -s to sum up to unity, and the Σ -s to be symmetric positive definite.

Subsequently, in what follows, we focus only on the change of the mean by considering the following weighted function $f_W(\theta)$ such that $f_W(\theta)$ satisfies

$$\nabla f_W(\theta) = EE^\top DP(\theta) \nabla \mathcal{L}(\theta), \quad (38)$$

where $D = \text{diag}(I_K, W, I_{d^2K})$, and with $W \in \mathbb{R}^{dK \times dK}$ being a weight matrix that mixes the different means. In particular, we can consider $W = \text{diag}(\beta_1 I_d, \dots, \beta_K I_d)$ where $\beta_i > 0$ denotes a scaling of the mean similar to a learning rate but applied only

Algorithm 3 Weighted projection-based GEM (W-PB-GEM).

Input: $y \in \mathbb{R}^d, p_\theta(y), \theta^{(0)} \in \Theta, W \in \mathbb{R}^{dK \times dK}$.

Output: $\hat{\theta}$.

- 1: **for** $k = 0, 1, 2, \dots$ (until some stopping criterion) **do**
- 2: $\theta^{(k+1)} = \theta^{(k)} + EE^\top DP(\theta^{(k)}) \frac{\partial \mathcal{L}}{\partial \theta} \Big|_{\theta=\theta^{(k)}}$
- 3: **end for**
- 4: **return** $\hat{\theta} = \text{last computed iteration in } \{\theta^{(k)}\}$.

on the component rates of the means of the mixture model. Note that we can extend this design step only on the means because the means are the only parameters of the GMMs under consideration that do not have implicit constraints associated with them. This allows us to introduce the following parameter update step:

$$\theta^{(k+1)} = \theta^{(k)} + EE^\top DP(\theta^{(k)}) \nabla \mathcal{L}(\theta^{(k)}), \quad (39)$$

for an algorithm which we will refer to as the *weighted projection-based GEM (W-PB-GEM) algorithm* – see Algorithm 3.

As a result, we have the following corollary on the convergence rate of the W-PB-GEM algorithm.

Corollary 3.7: *Suppose that there exists a function $f_W(\theta)$ that is μ_f -strongly convex, has an L -Lipschitz gradient, and satisfies $\nabla f_W(\theta) = EE^\top DP(\theta) \nabla \mathcal{L}(\theta)$, where $D = \text{diag}(I_K, W, I_{d^2K})$, and with $W \in \mathbb{R}^{dK \times dK}$ being the matrix of weights that determine the component-wise mixture of the means of the GMM whose parameters are to be estimated. Then, $\theta^{(k+1)} = \theta^{(k)} + u^{(k)}$, with $u^{(k)} = \nabla f_W(\theta^{(k)})$ is a GEM algorithm (i.e. $\{\theta^{(k)}\} \rightarrow \theta^*$, where θ^* is the maximum-likelihood estimate) with convergence rate ρ bounded by*

$$\rho \leq \max\{|1 - \mu_f|, |1 - L|\}. \quad (40)$$

Remark 3.1: The convergence rates obtained for the W-PB-GEM algorithm are the same as those obtained for the PB-GEM algorithm. It is to be noted, however, that the update equations associated with the α -s and the Σ -s cannot be arbitrarily changed because of the explicit constraints associated with them.

Remark 3.2: It is worth mentioning here that the convergence rates as obtained in (29) and (40) are merely upper bounds, and, unfortunately, do not shed any light on the transient behaviour of the PB-GEM or W-PB-GEM algorithm – see the inset of Figures 2 and 4.

4. Pedagogical examples

In this section, we seek to demonstrate a pedagogical example that shows the efficacy of the methods extended in this paper in identifying the parameters of unknown GMMs. To do this, we first sample 1000 arbitrary points from a mixture of two Gaussians with the following parameters:

$$\mu_1^* = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \text{and} \quad \mu_2^* = \begin{bmatrix} -1 \\ -1 \end{bmatrix},$$

$$\Sigma_1^* = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \Sigma_2^* = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

and

$$\alpha^* = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}.$$

Further, we initialised the algorithms with the following parameters:

$$\mu_1 = \begin{bmatrix} -3 \\ 3 \end{bmatrix}, \quad \mu_2 = \begin{bmatrix} 3 \\ -3 \end{bmatrix},$$

such that they lie on the line which is orthogonal to the direction defined by μ_1^* and μ_2^* . Additionally, Σ_1 and Σ_2 are initialised to be arbitrary positive definite matrices and α is arbitrarily initialised such that $\alpha_1 + \alpha_2 = 1$.

4.1 The PB-GEM algorithm

We first consider a pedagogical example to demonstrate the performance of the proposed PB-GEM algorithm on estimating the parameters of the synthetic Gaussian mixture model specified above. The results of running the PB-GEM algorithm




Figure 1. 1000 randomly sampled points from a synthetic two-component Gaussian mixture model and the contour plots of the initial and final estimated probability distribution functions using the PB-GEM algorithm. The circles show the initial positions, and the left and right traces show the evolutions of μ_1 and μ_2 , respectively.

to determine the parameters of the above mixture model are shown in Figures 1 and 2. We see that the proposed PB-GEM algorithm is able to successfully determine the parameters of the synthetic GMM from which the points have been sampled. Convergence is obtained in 316 iterations (i.e. to attain a relative change in log-likelihood smaller than 10^{-10}).

4.2 The W-PB-GEM algorithm

Next, we test the performance of the W-PB-GEM algorithm. The matrix of weights W that determine the mixture of proportions during the updates of the means is given by

$$W = \begin{bmatrix} 0.996I_2 & 0 \\ 0 & 0.996I_2 \end{bmatrix}.$$

The results of running the W-PB-GEM algorithm with the same initialisations for $\mu_1, \mu_2, \Sigma_1, \Sigma_2$, and α are documented in Figures 3 and 4. Convergence is obtained in 279 iterations with the same stopping criterion used in the previous section.

4.3 Multi-class classification

In what follows, we also present an illustrative example where




Figure 5. 1000 randomly sampled points from a synthetic four-component Gaussian mixture model and the contour plots of the initial and final estimated probability distribution functions using the PB-GEM algorithm. The circles show the initial positions, and the traces show the evolutions of $\mu_i, i = 1, 2, 3, 4$.

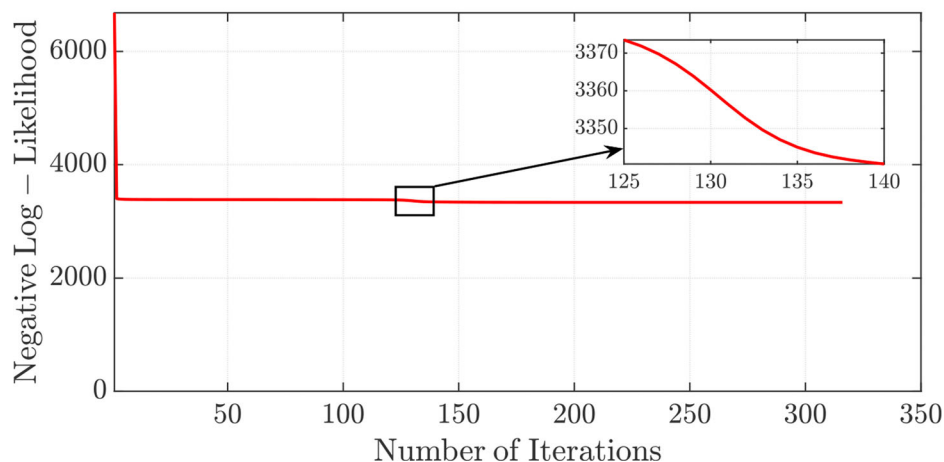


Figure 2. Plot of negative log-likelihood versus the number of iterations for the PB-GEM algorithm. The transient behaviour between iterations 125 and 140 is shown in the inset.




Figure 3. 1000 randomly sampled points from a synthetic two-component Gaussian mixture model and the contour plots of the initial and final estimated probability distribution functions using the W-PB-GEM algorithm. The circles show the initial positions, and the left and right traces show the evolutions of μ_1 and μ_2 , respectively.




Figure 4. Plot of negative log-likelihood versus the number of iterations for the W-PB-GEM algorithm. The transient behaviour between iterations 125 and 140 is shown in the inset.




Figure 7. Means and standard deviations of the negative log-likelihood function for 30 instances of the PB-GEM and W-PB-GEM algorithm for the two-class classification example running with the same initial conditions for a particular instance.

both the PB-GEM and the W-PB-GEM algorithms are used in order to identify the parameters of a GMM with more than two Gaussians. We sample 1000 arbitrary points from a mixture of four Gaussians with the following parameters:

$$\mu_1^* = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \mu_2^* = \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \quad \mu_3^* = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \text{and} \\ \mu_4^* = \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \quad \Sigma_1^* = \Sigma_2^* = \Sigma_3^* = \Sigma_4^* = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

and

$$\alpha^* = \begin{bmatrix} 0.25 \\ 0.25 \\ 0.25 \\ 0.25 \end{bmatrix}.$$

Further, we initialised both the algorithms with the following parameters:

$$\mu_1 = \begin{bmatrix} -3 \\ 3 \end{bmatrix}, \quad \mu_2 = \begin{bmatrix} 3 \\ -3 \end{bmatrix}, \quad \mu_3 = \begin{bmatrix} -3 \\ -3 \end{bmatrix}, \quad \text{and} \\ \mu_4 = \begin{bmatrix} 3 \\ 3 \end{bmatrix}.$$

Additionally, $\Sigma_1, \Sigma_2, \Sigma_3,$ and Σ_4 are initialised to be arbitrary positive definite matrices and α is arbitrarily initialised such that $\sum_{i=1}^4 \alpha_i = 1$. The matrix of weights W that determine the mixture of proportions during the updates of the means for the W-PB-GEM algorithm is once again given by

$$W = \begin{bmatrix} 0.996I_2 & 0 \\ 0 & 0.996I_2 \end{bmatrix}.$$

The results of running the PB-GEM and W-PB-GEM algorithms for this problem are shown in Figures 5 and 6, respectively. The results are similar to the case on two-class classification. Convergence (i.e. attaining a relative change in log-likelihood smaller than 10^{-10}) is obtained in 1822 iterations for the PB-GEM algorithm and in 472 iterations for the W-PB-GEM algorithm.

4.4 Discussion of results

The reason why the initial conditions on the means are selected such that they lie on a line orthogonal to the means characterising the synthetic GMM considered above is to intentionally




Figure 6. 1000 randomly sampled points from a synthetic four-component Gaussian mixture model and the contour plots of the initial and final estimated probability distribution functions using the W-PB-GEM algorithm. The circles show the initial positions, and the traces show the evolutions of $\mu_i, i = 1, 2, 3, 4$.

make the convergence of the PB-GEM and W-PB-GEM algorithms more difficult. We also illustrate in Figure 7 a comparative study of the PB-GEM and W-PB-GEM algorithms for the two-class example (the matrix W was selected identical to the one in Section 4.2) by plotting the mean and standard deviation of the negative log-likelihood function for 30 instances of both the algorithms when they are initialised with the same set of parameters for a particular instance. In general, such worst-case initialisation conditions are useful in order to gain insights into the transient behaviours of such algorithms.

It is also instructive here to note that for the problem of identifying the parameters of a high-dimensional GMM, the number of iterations to convergence would grow exponentially. In such a case, it would be extremely important to have convergence to the actual parameters in as few iterations as possible, since each iteration would involve a pass over the entire dataset, and when the dataset is large, having a lower number of iterations to convergence would amount to a reduction in the amount of time taken for the estimation of the parameters.

We also reiterate that the convergence rates presented in (29) and (40) are merely upper bounds. As demonstrated in the insets of Figures 2 and 4, these have no relationship with the transient behaviours of the PB-GEM and W-PB-GEM algorithms. Although, in practice, we can improve the convergence rates of these algorithms by designing new and more efficient varieties (as detailed in the construction of the W-PB-GEM algorithm), the upper bound of the obtained convergence rates does not change.

5. Conclusions and future work

In this paper, we analysed a GEM algorithm to estimate the parameters of GMMs from a dynamical systems perspective. In particular, we showed that this algorithm can be understood as an LTI system connected in feedback with a nonlinearity. The convergence properties of the proposed algorithm are studied by utilising tools from robust control theory. We also explored the simple design of this class of GEM algorithms and provided evidence using pedagogical examples that it might be possible to improve the transient and the practical convergence of these algorithms despite the fact that they exhibit the same asymptotic convergence rates. Future work will consist of using tools from adaptive systems theory to accelerate practical convergence properties for GEM algorithms. Additionally, fundamental connections exist between the EM algorithm and proximal point methods (Chrétien & Hero, 2000, 2008; Figueiredo, 2008) and future work will focus on analysing proximal interpretations of the EM algorithm using tools from robust control theory (Fazlyab et al., 2018; Lessard et al., 2016).

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Sarthak Chatterjee  <http://orcid.org/0000-0002-4774-9436>

Orlando Romero  <http://orcid.org/0000-0001-7545-761X>

Sérgio Pequito  <http://orcid.org/0000-0002-5143-1543>

References

- Balakrishnan, S., Wainwright, M. J., & Yu, B. (2017). Statistical guarantees for the EM algorithm: From population to sample-based analysis. *The Annals of Statistics*, 45(1), 77–120. <https://doi.org/10.1214/16-AOS1435>
- Bishop, C. (2006). *Pattern recognition and machine learning*. Springer Verlag.
- Chrétien, S., & Hero, A. O. (2000). Kullback proximal algorithms for maximum-likelihood estimation. *IEEE Transactions on Information Theory*, 46(5), 1800–1810. <https://doi.org/10.1109/18.857792>
- Chrétien, S., & Hero, A. O. (2008). On EM algorithms and their proximal generalizations. *ESAIM: Probability and Statistics*, 12, 308–326. <https://doi.org/10.1051/ps:2007041>
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B (Methodological)*, 39(1), 1–22. <https://doi.org/10.1111/rssb.1977.39.issue-1>
- Fazlyab, M., Ribeiro, A., Morari, M., & Preciado, V. M. (2018). Analysis of optimization algorithms via integral quadratic constraints: Nonstrongly convex problems. *SIAM Journal on Optimization*, 28(3), 2654–2689. <https://doi.org/10.1137/17M1136845>
- Fessler, J. A., & Hero, A. O. (1995). Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms. *IEEE Transactions on Image Processing*, 4(10), 1417–1429. <https://doi.org/10.1109/83.465106>
- Figueiredo, M. A. T. (2008). Lecture notes on the EM algorithm. http://www.lx.it.pt/mtf/Figueiredo_EM_Algorithm.pdf
- Jay, L. O. (2001). A note on Q-order of convergence. *BIT Numerical Mathematics*, 41(2), 422–429. <https://doi.org/10.1023/A:1021902825707>
- Krisjansson, T., Frey, B., Deng, L., & Acero, A. (2001). Joint estimation of noise and channel distortion in a generalized EM framework. In *Proceedings of the IEEE Workshop on Automatic Speech Recognition and Understanding* (pp. 155–158). IEEE.
- Lessard, L., Recht, B., & Packard, A. (2016, January 1). Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1), 57–95. <https://doi.org/10.1137/15M10-09597>
- Moon, T. K. (1996). The expectation-maximization algorithm. *IEEE Signal Processing Magazine*, 13(6), 47–60. <https://doi.org/10.1109/79.543975>
- Neal, R. M., & Hinton, G. E. (1998). A view of the EM algorithm that justifies incremental, sparse, and other variants. In M. I. Jordan (Ed.), *Learning in graphical models* (pp. 355–368). MIT Press.
- Nowak, R. D. (2003, August). Distributed EM algorithms for density estimation and clustering in sensor networks. *IEEE Transactions on Signal Processing*, 51(8), 2245–2253. <https://doi.org/10.1109/TSP.2003.814623>
- Roche, A. (2011). EM algorithm and variants: an informal tutorial. *arXiv preprint arXiv:1105.1476*.
- Romero, O., Chatterjee, S., & Pequito, S. (2019, July). Convergence of the expectation-maximization algorithm through discrete-time Lyapunov stability theory. In *Proceedings of the 2019 American Control Conference* (pp. 163–168). IEEE.
- Salakhutdinov, R., Roweis, S., & Ghahramani, Z. (2003). Optimization with EM and expectation-conjugate-gradient. In T. Fawcett & N. Mishra (Eds.), *Proceedings of the 20th International Conference on Machine Learning* (pp. 672–679). AAAI Press.
- Tan, P. N., Steinbach, M., & Kumar, V. (2005). *Introduction to data mining* (1st ed.). Addison-Wesley Longman Publishing Co., Inc.
- Wu, C. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11(1), 95–103. <https://doi.org/10.1214/aos/1176346060>
- Xu, L., & Jordan, M. I. (1996, January). On convergence properties of the EM algorithm for Gaussian mixtures. *Neural Computation*, 8(1), 129–151. <https://doi.org/10.1162/neco.1996.8.1.129>