



**HAL**  
open science

# A note on the relations between mixture models, maximum-likelihood and entropic optimal transport

Titouan Vayer, Etienne Lasalle

► **To cite this version:**

Titouan Vayer, Etienne Lasalle. A note on the relations between mixture models, maximum-likelihood and entropic optimal transport. 2025. hal-04901607v2

**HAL Id: hal-04901607**

**<https://hal.science/hal-04901607v2>**

Preprint submitted on 22 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# A note on the relations between mixture models, maximum-likelihood and entropic optimal transport

Titouan Vayer

Inria, ENS de Lyon, CNRS, UCBL, LIP, UMR 5668.

`titouan.vayer@inria.fr`

Etienne Lasalle

ENS de Lyon, CNRS, UCBL, Inria, LIP, UMR 5668.

`etienne.lasalle@ens-lyon.fr`

January 22, 2025

## Abstract

This note aims to demonstrate that performing maximum-likelihood estimation for a mixture model is equivalent to minimizing over the parameters an optimal transport problem with entropic regularization. The objective is pedagogical: we seek to present this already known result in a concise and hopefully simple manner. We give an illustration with Gaussian mixture models by showing that the standard EM algorithm is a specific block-coordinate descent on an optimal transport loss.

**Notations.** Any vector  $\mathbf{x} \in \mathbb{R}^K$  is treated as column matrix. The discrete probability simplex with  $K$  bins is noted as  $\Delta_K = \{\mathbf{a} \in \mathbb{R}_+^K : \sum_{j=1}^K a_j = 1\}$ . The vector of  $K$  ones is denoted as  $\mathbf{1}_K$ .  $\delta_x$  is the dirac mass supported at  $x$ . For simplicity in this note, we will deliberately remain vague in certain places regarding “edge cases” like  $0 \log 0$ . To be entirely rigorous, we would need to reason on the supports of the matrices and define objects that can take  $\infty$  as a value.

## 1 Introduction and preliminaries on optimal transport

The relations between maximum-likelihood and optimal transport (OT) have already been discussed in multiple works (Rigollet and Weed, 2018; Mena et al., 2020; Diebold et al., 2024). The purpose of this brief note is to provide the key tools used to establish these connections. The primary aim is pedagogical: we will focus on the (discrete) mixtures case, adopting a “computational OT” perspective. Hopefully, readers will find this exercise insightful. Our analysis will largely rely on the approach described in Rigollet and Weed (2018), though adapted to a different formalism and applied to a slightly different problem (mixture estimation rather than Gaussian deconvolution).

To fix the notations, we first briefly recall the fundamentals of (entropic) discrete optimal transport (OT) and readers seeking more details can refer to [Peyré et al. \(2019\)](#). Let  $\mathbf{C} \in \mathbb{R}^{n \times K}$  be a cost matrix, representing, for example, the distances between points from two distributions. Let  $\mathbf{a} \in \Delta_n, \mathbf{b} \in \Delta_K$  be two probability vectors, encoding for instance the quantities to be transported and supplied, respectively. The goal of OT is to determine a way to move these quantities while respecting supply constraints, such that the total transportation cost is minimized (defined by  $\mathbf{C}$ ). This is formalized through the set of couplings, or transport plans, with marginals  $\mathbf{a}, \mathbf{b}$ , which is defined by

$$U(\mathbf{a}, \mathbf{b}) \triangleq \{\mathbf{P} \in \mathbb{R}_+^{n \times K} : \mathbf{P}\mathbf{1}_K = \mathbf{a}, \mathbf{P}^\top \mathbf{1}_n = \mathbf{b}\}. \quad (1.1)$$

When  $\mathbf{P} \in U(\mathbf{a}, \mathbf{b})$ ,  $P_{ij} \in [0, 1]$  represents the of probability mass transported from the  $i$ -th point to the  $j$ -th. When we only care about transporting the mass of the input measure without constraint on the supply we can consider *semi-relaxed* transport plans as

$$U_K(\mathbf{a}) \triangleq \{\mathbf{P} \in \mathbb{R}_+^{n \times K} : \mathbf{P}\mathbf{1}_K = \mathbf{a}\}. \quad (1.2)$$

Given a transport plan  $\mathbf{P}$ , the total cost of displacement is given by  $\langle \mathbf{C}, \mathbf{P} \rangle$  and the goal of standard OT is to find the transport plan that minimizes this cost.

Entropic regularization, introduced in [Cuturi \(2013\)](#), was proposed to accelerate the computation of the optimal transport plan. It relies on the Kullback-Leibler divergence between two matrices  $\mathbf{P}, \mathbf{Q} \in \mathbb{R}_+^{n \times K}$  which is defined as

$$\text{KL}(\mathbf{P}|\mathbf{Q}) \triangleq \sum_{i=1}^n \sum_{j=1}^K P_{ij} \log \frac{P_{ij}}{Q_{ij}}. \quad (1.3)$$

The entropic-regularized optimal transport problem (EOT) is expressed as

$$\text{OT}_\varepsilon(\mathbf{a}, \mathbf{b}, \mathbf{C}) \triangleq \min_{\mathbf{P} \in U(\mathbf{a}, \mathbf{b})} \langle \mathbf{C}, \mathbf{P} \rangle + \varepsilon \text{KL}(\mathbf{P}|\mathbf{ab}^\top), \quad (1.4)$$

In (1.4), the goal is to find the transport plan that minimizes a trade-off between the transport cost and a measure of “distance” to the uniform coupling  $\mathbf{ab}^\top$ , which distributes every source point to every target point uniformly. When  $\mathbf{P}$  is a semi-relaxed transport plan the problem (1.4) will be called a semi-relaxed entropic OT problem.

## 2 Maximum-likelihood for mixture models is minimization of EOT

In this note, we consider a mixture model as described below.

**Definition 2.1** ((Discrete) mixture model). *The generative process of a discrete mixture model consists in*

- $Y \sim P_Y$  where  $P_Y = \sum_{j=1}^K \pi_j^* \delta_j$  with  $\boldsymbol{\pi}^* \in \Delta_K$  represents the discrete distribution on  $K$  labels/classes. In other words,  $P_Y$  is the distribution of the “latent variables”.

- $X|y = j \sim P_{X|Y}(\cdot|j, \boldsymbol{\theta}^*)$ , where  $\boldsymbol{\theta}^* \in \Theta$  and  $P_{X|Y}$  is the parametrized distribution of the data given the label.

We note  $P_{X,Y}$  the corresponding joint distribution. A simple example of this generative process is the Gaussian mixture model where the parametrized distribution has density  $P_{X|Y}(\mathbf{x}|j, \boldsymbol{\theta}^*) \propto \exp(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_j^*)^\top \boldsymbol{\Sigma}^{*-1}(\mathbf{x} - \boldsymbol{\mu}_j^*))$  where  $\boldsymbol{\mu}_j^* \in \mathbb{R}^d$  is the true mean associated to the  $j$ -th class and  $\boldsymbol{\Sigma}^* \succ 0$  the true covariance (assumed to be identical for each class). In this case  $\boldsymbol{\theta}^* = (\boldsymbol{\mu}_1^*, \dots, \boldsymbol{\mu}_K^*, \boldsymbol{\Sigma}^*)$ .

Now suppose that we observe some samples  $\mathbf{x}_1, \dots, \mathbf{x}_n \sim P_X$  i.i.d. where  $P_X$  is the distribution of the data, according to the generative model above. The goal of maximum-likelihood estimation is to infer the parameters  $\boldsymbol{\pi}^*, \boldsymbol{\theta}^*$  from these observations. By independence, the negative log-likelihood for a parameter  $\boldsymbol{\beta} = (\boldsymbol{\theta}, \boldsymbol{\pi})$  writes

$$\begin{aligned} \mathcal{L}(\boldsymbol{\beta}) &= - \sum_{i=1}^n \log P_X(\mathbf{x}_i|\boldsymbol{\beta}) \\ &\stackrel{*}{=} - \sum_{i=1}^n \log \left( \sum_{j=1}^K P_{X,Y}(\mathbf{x}_i, j|\boldsymbol{\beta}) \right) \stackrel{**}{=} - \sum_{i=1}^n \log \left( \sum_{j=1}^K \pi_j P_{X|Y}(\mathbf{x}_i|j, \boldsymbol{\theta}) \right) \quad (2.1) \\ &= - \sum_{i=1}^n \log \left( \sum_{j=1}^K \pi_j \exp(\log(P_{X|Y}(\mathbf{x}_i|j, \boldsymbol{\theta}))) \right). \end{aligned}$$

In  $*$  we used the law of total probability and in  $**$  we used the Bayes' formula. In the last line we use a (at first glance) stupid reparametrization.

We will prove three facts: first, the negative log-likelihood (2.1) can be rewritten as a certain semi-relaxed entropic OT problem; second, there is an entropic OT problem that is an upper bound for the negative log-likelihood; and third, minimizing  $\mathcal{L}(\boldsymbol{\beta})$  with respect to  $\boldsymbol{\pi}$  results in equality with this upper bound.

The key result to make the connections between OT and log-likelihood is to rewrite the “logsumexp” term as a minimization problem over the probability simplex. This is next in the following lemma, which is sometimes referred to as the Gibbs variational principle or the dual formulation of the KL divergence.

**Lemma 2.2.** *Let  $\pi_1, \dots, \pi_K$  be positive real numbers and  $h_1, \dots, h_K \in \mathbb{R}$ . Then*

$$\log \left( \sum_{j=1}^K \pi_j \exp(h_j) \right) = \max_{\mathbf{p} \in \Delta_K} \sum_{j=1}^K h_j p_j - \sum_{j=1}^K p_j \log \left( \frac{p_j}{\pi_j} \right).$$

The optimal solution is given by  $\forall k \in \llbracket K \rrbracket$ ,  $p_k = \frac{\pi_k \exp(h_k)}{\sum_{j=1}^K \pi_j \exp(h_j)}$ .

For now, we postpone the proof of this result (see Section 5) but we can use it to rewrite the negative log-likelihood. Combining (2.1) and Lemma 2.2 with  $\log(P_{X|Y}(\mathbf{x}_i|j, \boldsymbol{\beta}))$  in the role of  $h_j$  we obtain

$$\begin{aligned} \mathcal{L}(\boldsymbol{\beta}) &= - \sum_{i=1}^n \left( \max_{\mathbf{p} \in \Delta_K} \sum_{j=1}^K \log(P_{X|Y}(\mathbf{x}_i|j, \boldsymbol{\theta})) p_j - \sum_{j=1}^K p_j \log \left( \frac{p_j}{\pi_j} \right) \right) \\ &= - \max_{\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(n)} \in \Delta_K} \sum_{i=1}^n \sum_{j=1}^K \log(P_{X|Y}(\mathbf{x}_i|j, \boldsymbol{\theta})) p_j^{(i)} - \sum_{i=1}^n \sum_{j=1}^K p_j^{(i)} \log \left( \frac{p_j^{(i)}}{\pi_j} \right). \end{aligned}$$

Equivalently,

$$\frac{1}{n}\mathcal{L}(\boldsymbol{\beta}) = \min_{\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(n)} \in \Delta_K} \sum_{i,j} C_{ij}(\boldsymbol{\theta}) \frac{p_j^{(i)}}{n} + \sum_{i,j} \frac{p_j^{(i)}}{n} \log\left(\frac{p_j^{(i)}/n}{\pi_j/n}\right),$$

where we introduced  $C_{ij}(\boldsymbol{\theta}) \triangleq -\log(P_{X|Y}(\mathbf{x}_i|j, \boldsymbol{\theta}))$ . Now suppose that we have solutions  $\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(n)}$  of the minimization problem above and that we consider  $\mathbf{P} = \frac{1}{n}(\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(n)})^\top \in \mathbb{R}_+^{n \times K}$ . Then obviously  $\mathbf{P}\mathbf{1}_K = \frac{1}{n}\mathbf{1}_n$ . Conversely, any matrix  $\mathbf{P} \in \mathbb{R}_+^{n \times K}$  with  $\mathbf{P}\mathbf{1}_K = \frac{1}{n}\mathbf{1}_n$  can be written as  $\mathbf{P} = \frac{1}{n}(\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(n)})^\top$  for some probability vectors  $\mathbf{p}^{(1)}, \dots, \mathbf{p}^{(n)} \in \Delta_K$  (which are simply the rows of  $\mathbf{P}$ ). In other words, this proves that

$$\begin{aligned} \frac{1}{n}\mathcal{L}(\boldsymbol{\beta}) &= \min_{\substack{\mathbf{P} \in \mathbb{R}_+^{n \times K} \\ \mathbf{P}\mathbf{1}_K = \frac{1}{n}\mathbf{1}_n}} \sum_{i,j} C_{ij}(\boldsymbol{\theta}) P_{ij} + \sum_{i,j} P_{ij} \log\left(\frac{P_{ij}}{\pi_j/n}\right) \\ &= \min_{\mathbf{P} \in U_K(\frac{1}{n})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}(\mathbf{P} | \frac{1}{n}\boldsymbol{\pi}^\top). \end{aligned} \quad (2.2)$$

We almost have the desired entropic OT problem, albeit with a ‘‘semi-relaxed’’ constraint instead of a standard coupling constraint: if  $\mathbf{P}$  were in  $U(\frac{1}{n}, \boldsymbol{\pi})$  we would be done.

To obtain an entropic OT problem we only need to rewrite a bit the quantity above: we show that minimizing the negative log-likelihood with respect to  $\boldsymbol{\pi}$  leads to a coupling constraint rather than a semi-relaxed one. First, for any admissible  $\mathbf{P} \in U_K(\frac{1}{n})$  of the minimization problem (2.2), we have (see Lemma 5.2)

$$\text{KL}(\mathbf{P} | \frac{1}{n}\boldsymbol{\pi}^\top) = \text{KL}\left(\mathbf{P} | \frac{1}{n}(\mathbf{P}^\top \mathbf{1}_n)^\top\right) + \text{KL}(\mathbf{P}^\top \mathbf{1}_n | \boldsymbol{\pi}). \quad (2.3)$$

This implies that

$$\frac{1}{n}\mathcal{L}(\boldsymbol{\beta}) = \min_{\mathbf{P} \in U_K(\frac{1}{n})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}\left(\mathbf{P} | \frac{1}{n}(\mathbf{P}^\top \mathbf{1}_n)^\top\right) + \text{KL}(\mathbf{P}^\top \mathbf{1}_n | \boldsymbol{\pi}). \quad (2.4)$$

Now consider  $\mathbf{P}^*$  the solution of the entropic OT problem  $\min_{\mathbf{P} \in U(\frac{1}{n}, \boldsymbol{\pi})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}(\mathbf{P} | \frac{1}{n}\boldsymbol{\pi}^\top)$ . By the constraints,  $\mathbf{P}^* \in U_K(\frac{1}{n}\mathbf{1}_n)$  and  $\text{KL}(\mathbf{P}^{*\top} \mathbf{1}_n | \boldsymbol{\pi}) = 0$ . Thus, by suboptimality in (2.4)  $\frac{1}{n}\mathcal{L}(\boldsymbol{\beta}) \leq \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P}^* \rangle + \text{KL}\left(\mathbf{P}^* | \frac{1}{n}(\mathbf{P}^{*\top} \mathbf{1}_n)^\top\right)$ . Hence we first obtain an upper-bound on the negative log-likelihood:

$$\frac{1}{n}\mathcal{L}(\boldsymbol{\beta}) \leq \text{OT}_{\varepsilon=1}\left(\frac{1}{n}, \boldsymbol{\pi}, \mathbf{C}(\boldsymbol{\theta})\right). \quad (2.5)$$

To obtain an equality we will minimize with respect to  $\boldsymbol{\pi}$ . Using that  $\min_{\mathbf{b} \in \Delta_K} \text{KL}(\mathbf{a} | \mathbf{b}) = \text{KL}(\mathbf{a} | \mathbf{a}) = 0$ , if we minimize the RHS in (2.4) over  $\boldsymbol{\pi}$  we get,

$$\begin{aligned} &\min_{\boldsymbol{\pi} \in \Delta_K} \min_{\mathbf{P} \in U_K(\frac{1}{n})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}\left(\mathbf{P} | \frac{1}{n}(\mathbf{P}^\top \mathbf{1}_n)^\top\right) + \text{KL}(\mathbf{P}^\top \mathbf{1}_n | \boldsymbol{\pi}) \\ &= \min_{\mathbf{P} \in U_K(\frac{1}{n})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}\left(\mathbf{P} | \frac{1}{n}(\mathbf{P}^\top \mathbf{1}_n)^\top\right) \stackrel{*}{=} \min_{\boldsymbol{\pi} \in \Delta_K} \min_{\mathbf{P} \in U(\frac{1}{n}, \boldsymbol{\pi})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}(\mathbf{P} | \frac{1}{n}\boldsymbol{\pi}^\top) \\ &= \min_{\boldsymbol{\pi} \in \Delta_K} \text{OT}_{\varepsilon=1}\left(\frac{1}{n}, \boldsymbol{\pi}, \mathbf{C}(\boldsymbol{\theta})\right). \end{aligned}$$

For the equality in  $\star$  we used that the second marginal in the optimization problem is redundant in the RHS. Precisely, if  $\mathbf{P}_0$  is a solution of the LHS then  $\mathbf{b}_0 = \mathbf{P}_0^\top \mathbf{1}_n \in \Delta_K$  thus  $\mathbf{P}_0 \in U(\frac{\mathbf{1}_n}{n}, \mathbf{b}_0)$  and  $\text{KL}(\mathbf{P}_0 | \frac{\mathbf{1}_n}{n} (\mathbf{P}_0^\top \mathbf{1}_n)^\top) = \text{KL}(\mathbf{P}_0 | \frac{\mathbf{1}_n}{n} \mathbf{b}_0^\top)$  which implies that  $\text{LHS} \geq \text{RHS}$ . Conversely, if  $(\boldsymbol{\pi}_1, \mathbf{P}_1)$  is a solution of the RHS problem then  $\mathbf{P}_1 \in U_K(\frac{\mathbf{1}_n}{n})$  and  $\boldsymbol{\pi}_1 = \mathbf{P}_1^\top \mathbf{1}_n$ , hence  $\text{LHS} \leq \text{RHS}$ .

In particular for any  $\boldsymbol{\theta} \in \Theta$ ,  $\min_{\boldsymbol{\pi}} \frac{1}{n} \mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\theta}) = \min_{\mathbf{P}} \text{OT}_{\varepsilon=1}(\frac{\mathbf{1}_n}{n}, \boldsymbol{\pi}, \mathbf{C}(\boldsymbol{\theta}))$ . This gives the final result written below.

**Proposition 2.3** (MLE for mixture models is minimization of an EOT problem). *Consider a mixture model as in Definition 2.1 and  $\mathcal{L}$  the negative log-likelihood on  $n$  i.i.d. samples. First, we have the identity*

$$\forall (\boldsymbol{\pi}, \boldsymbol{\theta}) \in \Delta_K \times \Theta, \quad \frac{1}{n} \mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\theta}) = \min_{\mathbf{P} \in U_K(\frac{\mathbf{1}_n}{n})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}(\mathbf{P} | \frac{\mathbf{1}_n}{n} \boldsymbol{\pi}^\top), \quad (2.6)$$

where the cost matrix is  $\mathbf{C}(\boldsymbol{\theta}) \triangleq (-\log P_{X|Y}(\mathbf{x}_i | j, \boldsymbol{\theta}))_{ij}$ . Second, the EOT problem is an upper-bound of the negative log-likelihood:

$$\forall (\boldsymbol{\pi}, \boldsymbol{\theta}) \in \Delta_K \times \Theta, \quad \frac{1}{n} \mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\theta}) \leq \text{OT}_{\varepsilon=1}(\frac{\mathbf{1}_n}{n}, \boldsymbol{\pi}, \mathbf{C}(\boldsymbol{\theta})). \quad (2.7)$$

Third, minimizing the negative log-likelihood with respect to the parameters is equivalent to minimizing over the parameters an EOT problem. Precisely,

$$\min_{(\boldsymbol{\pi}, \boldsymbol{\theta}) \in \Delta_K \times \Theta} \frac{1}{n} \mathcal{L}(\boldsymbol{\pi}, \boldsymbol{\theta}) = \min_{(\boldsymbol{\pi}, \boldsymbol{\theta}) \in \Delta_K \times \Theta} \text{OT}_{\varepsilon=1}(\frac{\mathbf{1}_n}{n}, \boldsymbol{\pi}, \mathbf{C}(\boldsymbol{\theta})). \quad (2.8)$$

### 3 Illustration with Gaussian Mixture Models

To illustrate the previous results we will show that the updates of the Expectation–Maximization algorithm (EM) for the Gaussian mixture model (GMM) can be interpreted as a block-coordinate descent (BCD) on the EOT loss (see e.g. [Murphy 2012](#), Section 11.4 for a description of the EM algorithm). According to Proposition 2.3 minimizing the negative log-likelihood is equivalent to solving

$$\min_{(\boldsymbol{\pi}, \boldsymbol{\theta}) \in \Delta_K \times \Theta} \min_{\mathbf{P} \in U_K(\frac{\mathbf{1}_n}{n})} \langle \mathbf{C}(\boldsymbol{\theta}), \mathbf{P} \rangle + \text{KL}(\mathbf{P} | \frac{\mathbf{1}_n}{n} \boldsymbol{\pi}^\top). \quad (3.1)$$

The BCD strategy consists in alternating between minimizing (3.1) in  $\boldsymbol{\pi}, \boldsymbol{\theta}, \mathbf{P}$  while keeping the other variables fixed.

The update of  $\mathbf{P}$  with  $\boldsymbol{\pi}, \boldsymbol{\theta}$  fixed consists in solving a semi-relaxed entropic OT problem. As described e.g. in [Flamary et al. \(2016\)](#) (or also in the proof of Proposition 2.3) the problem decouples with respect to the rows of  $\mathbf{P}$  and the solution is given by

$$\forall (i, j) \in \llbracket n \rrbracket \times \llbracket K \rrbracket, \quad P_{ij} = \frac{1}{n} \frac{\pi_j \exp(-C_{ij}(\boldsymbol{\theta}))}{\sum_{k=1}^K \pi_k \exp(-C_{ik}(\boldsymbol{\theta}))} = \frac{1}{n} \frac{\pi_j P_{X|Y}(\mathbf{x} | j, \boldsymbol{\theta})}{\sum_{k=1}^K \pi_k P_{X|Y}(\mathbf{x} | k, \boldsymbol{\theta})}. \quad (3.2)$$

This step actually corresponds to finding the conditional distribution in the ‘‘E step’’ of the EM algorithm. The update for  $\boldsymbol{\pi}$ , with  $\mathbf{P}$  held fixed, can be found using (2.3) and is simply

$$\boldsymbol{\pi} = \mathbf{P}^\top \mathbf{1}_n \text{ i.e. } \forall j \in \llbracket K \rrbracket, \quad \pi_j = \sum_{i=1}^n P_{ij}. \quad (3.3)$$

Finally we derive the update of  $\boldsymbol{\theta}$  in the GMM case. We consider

$$P_{X|Y}(\mathbf{x}|j, \boldsymbol{\theta}) = (2\pi)^{-d/2} \det \boldsymbol{\Sigma}^{-1/2} \exp \left( -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_j)^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}_j) \right),$$

and the goal is to update  $\boldsymbol{\theta} = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K, \boldsymbol{\Sigma})$  with  $\boldsymbol{\mu}_j \in \mathbb{R}^d$  and  $\boldsymbol{\Sigma} \succ 0$ . With other variables fixed this boils down to solving

$$\min_{(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K, \boldsymbol{\Sigma})} \frac{1}{2} \sum_{ij} (\mathbf{x}_i - \boldsymbol{\mu}_j)^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}_i - \boldsymbol{\mu}_j) P_{ij} + \frac{n}{2} \log \det \boldsymbol{\Sigma}.$$

Setting the gradient of this loss to zero, one can show that the update of  $\boldsymbol{\Sigma}$  (with  $(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K)$  fixed) reads

$$\boldsymbol{\Sigma} = \frac{1}{n} \sum_{ij} P_{ij} (\mathbf{x}_i - \boldsymbol{\mu}_j)(\mathbf{x}_i - \boldsymbol{\mu}_j)^\top, \quad (3.4)$$

and the update of the means (with  $\boldsymbol{\Sigma}$  fixed) are

$$\forall j \in \llbracket K \rrbracket, \boldsymbol{\mu}_j = \frac{1}{\sum_i P_{ij}} \sum_i P_{ij} \mathbf{x}_i. \quad (3.5)$$

These updates exactly corresponds to the updates of the EM algorithm apply to a GMM: first update  $\mathbf{P}$  according to (3.2) which is the ‘‘E step’’, then the proportion of the classes with (3.3) and the means and covariance with (3.4) and (3.5), which is the ‘‘M step’’.

## 4 Discussions

To finish we make a few comments. First, the proof described above can be easily generalized to infinite mixtures, with appropriate assumptions on  $P_Y$ , see e.g. [Rigollet and Weed \(2018, Definition 2\)](#). With these assumptions we would obtain that minimizing the negative log-likelihood is equivalent to solving a problem of the form  $\inf_{(P_Y, \boldsymbol{\theta})} \text{OT}_{\varepsilon=1}(P_Y, \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}; \boldsymbol{\theta})$  where the cost of the OT is  $c(\mathbf{x}, y; \boldsymbol{\theta}) = -\log(P_{X|Y}(\mathbf{x}|y, \boldsymbol{\theta}))$ .

Also, as discussed in [Mena et al. \(2020\)](#), the same relations between negative log-likelihood and entropic OT can be obtained for more general generative models where  $P_Y$  are  $P_X$  are coupled via a joint distribution  $Q_{X,Y}^{\boldsymbol{\theta}^*}$  (such as  $dQ_{X,Y}^{\boldsymbol{\theta}^*}(\mathbf{x}, y) = \exp(-g_{\boldsymbol{\theta}^*}(\mathbf{x}, y))dP_X(\mathbf{x})dP_Y(y)$ ). This setting encompasses the GMM case and the principle of the proof remains similar to the one described here.

## 5 Postponed proofs

**Lemma 5.1.** *Let  $\pi_1, \dots, \pi_K$  be positive real numbers and  $h_1, \dots, h_K \in \mathbb{R}$ . Then*

$$\log \left( \sum_{j=1}^K \pi_j \exp(h_j) \right) = \max_{\mathbf{p} \in \Delta_K} \sum_{j=1}^K h_j p_j - \sum_{j=1}^K p_j \log \left( \frac{p_j}{\pi_j} \right).$$

*The optimal solution is given by  $\forall k \in \llbracket K \rrbracket, p_k = \frac{\pi_k \exp(h_k)}{\sum_{j=1}^K \pi_j \exp(h_j)}$ .*

*Proof.* The optimization problem above is a maximization of a strictly concave function. Consider the Lagrangian  $L(\mathbf{p}, \lambda) = \sum_j h_j p_j - \sum_j p_j \log(\frac{p_j}{\pi_j}) + \lambda(\sum_j p_j - 1)$ . Then  $\partial_{p_k} L(\mathbf{p}, \lambda) = h_k - \log(p_k/\pi_k) - 1 + \lambda$  thus  $\partial_{p_k} L(\mathbf{p}, \lambda) = 0 \iff p_k = \exp(\lambda - 1)\pi_k \exp(h_k)$ . By primal constraints  $\sum_j p_j = 1 \implies \exp(\lambda - 1) = \frac{1}{\sum_j \exp(h_j)\pi_j}$ . This gives the desired optimal solution. Also  $\forall k \in \llbracket K \rrbracket, \log(p_k/\pi_k) = h_k - \log(\sum_{j=1}^K \pi_j \exp(h_j))$  hence  $\sum_j p_j \log(p_j/\pi_j) = \sum_j h_j p_j - \log(\sum_{j=1}^K \pi_j \exp(h_j)) \sum_j p_j$ . Using that  $\sum_j p_j = 1$  gives the result.  $\square$

We also used the following result about the KL divergence.

**Lemma 5.2.** *Let  $\mathbf{P} \in \mathbb{R}_+^{n \times K}$  be a matrix and  $\mathbf{a} \in \mathbb{R}_+^n, \mathbf{b} \in \mathbb{R}_+^K$  then*

$$\text{KL}(\mathbf{P}|\mathbf{a}\mathbf{b}^\top) = \text{KL}(\mathbf{P} | (\mathbf{P}\mathbf{1}_K)(\mathbf{P}^\top \mathbf{1}_n)^\top) + \text{KL}(\mathbf{P}\mathbf{1}_K | \mathbf{a}) + \text{KL}(\mathbf{P}^\top \mathbf{1}_n | \mathbf{b}). \quad (5.1)$$

*Proof.* By definition

$$\begin{aligned} \text{KL}(\mathbf{P}|\mathbf{a}\mathbf{b}^\top) &= \sum_{ij} P_{ij} \log\left(\frac{P_{ij}}{a_i b_j}\right) = \sum_{ij} P_{ij} \log\left(\frac{P_{ij}}{(\mathbf{P}\mathbf{1}_K)_i (\mathbf{P}^\top \mathbf{1}_n)_j} \frac{(\mathbf{P}\mathbf{1}_K)_i (\mathbf{P}^\top \mathbf{1}_n)_j}{a_i b_j}\right) \\ &= \text{KL}(\mathbf{P} | (\mathbf{P}\mathbf{1}_K)(\mathbf{P}^\top \mathbf{1}_n)^\top) + \sum_{ij} P_{ij} \log\left(\frac{(\mathbf{P}\mathbf{1}_K)_i}{a_i}\right) + \sum_{ij} P_{ij} \log\left(\frac{(\mathbf{P}^\top \mathbf{1}_n)_j}{b_j}\right) \\ &= \text{KL}(\mathbf{P} | (\mathbf{P}\mathbf{1}_K)(\mathbf{P}^\top \mathbf{1}_n)^\top) + \sum_i (\mathbf{P}\mathbf{1}_K)_i \log\left(\frac{(\mathbf{P}\mathbf{1}_K)_i}{a_i}\right) + \sum_j (\mathbf{P}^\top \mathbf{1}_n)_j \log\left(\frac{(\mathbf{P}^\top \mathbf{1}_n)_j}{b_j}\right). \end{aligned}$$

$\square$

## References

- Philippe Rigollet and Jonathan Weed. Entropic optimal transport is maximum-likelihood deconvolution. *Comptes Rendus. Mathématique*, 356(11-12):1228–1235, 2018.
- Gonzalo Mena, Amin Nejatbakhsh, Erdem Varol, and Jonathan Niles-Weed. Sinkhorn em: an expectation-maximization algorithm based on entropic optimal transport. *arXiv preprint arXiv:2006.16548*, 2020.
- Jean-Frédéric Diebold, Nicolas Papadakis, Arnaud Dessein, and Charles-Alban Deledalle. A unified framework for hard and soft clustering with regularized optimal transport, 2024. URL <https://arxiv.org/abs/1711.04366>.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Neural Information Processing Systems (NeurIPS)*, 26, 2013.
- Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- Remi Flamary, Cédric Févotte, N. Courty, and Valentin Emyia. Optimal spectral transportation with application to music transcription. In *Neural Information Processing Systems (NeurIPS)*, 2016.