
Closed-form Marginal Likelihood in Gamma-Poisson Matrix Factorization

Louis Filstroff¹ Alberto Lumbreras¹ Cédric Févotte¹

Abstract

We present novel understandings of the Gamma-Poisson (GaP) model, a probabilistic matrix factorization model for count data. We show that GaP can be rewritten free of the score/activation matrix. This gives us new insights about the estimation of the topic/dictionary matrix by maximum marginal likelihood estimation. In particular, this explains the robustness of this estimator to over-specified values of the factorization rank, especially its ability to automatically prune irrelevant dictionary columns, as empirically observed in previous work. The marginalization of the activation matrix leads in turn to a new Monte Carlo Expectation-Maximization algorithm with favorable properties.

1. Introduction

The Gamma-Poisson (GaP) model is a probabilistic matrix factorization model which was introduced in the field of text information retrieval (Canny, 2004; Buntine & Jakulin, 2006). In this field, a corpus of text documents is typically represented by an integer-valued matrix \mathbf{V} of size $F \times N$, where each column \mathbf{v}_n represents a document as a so-called “bag of words”. Given a vocabulary of F words (or in practice semantic stems), the matrix entry v_{fn} is the number of occurrences of word f in the document n . GaP is a generative model described by a dictionary of “topics” or “patterns” \mathbf{W} (a non-negative matrix of size $F \times K$) and a non-negative “activation” or “score” matrix \mathbf{H} (of size $K \times N$), as follows:

$$\mathbf{H} \sim \prod_{k,n} \text{Gamma}(h_{kn} | \alpha_k, \beta_k), \quad (1)$$

$$\mathbf{V} | \mathbf{H} \sim \prod_{f,n} \text{Poisson}(v_{fn} | [\mathbf{W}\mathbf{H}]_{fn}), \quad (2)$$

¹IRIT, Université de Toulouse, CNRS, France. Correspondence to: Louis Filstroff <louis.filstroff@irit.fr>.

where we use the shape and rate parametrization of the Gamma distribution, i.e., $\text{Gamma}(x | \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}$. The dictionary \mathbf{W} is treated as a free deterministic variable.

Though this generative model takes its origins in text information retrieval, it has found applications (with variants) in other areas such as image reconstruction (Cemgil, 2009), collaborative filtering (Ma et al., 2011; Gopalan et al., 2015) or audio signal processing (Virtanen et al., 2008).

Denoting $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_K]^T$, $\boldsymbol{\beta} = [\beta_1, \dots, \beta_K]^T$, and treating the shape parameters α_k as fixed hyperparameters, maximum joint likelihood estimation (MJLE) in GaP amounts to the minimization of

$$C_{\text{JL}}(\mathbf{W}, \mathbf{H}, \boldsymbol{\beta}) \stackrel{\text{def}}{=} -\log p(\mathbf{V}, \mathbf{H} | \mathbf{W}, \boldsymbol{\beta}) \quad (3)$$
$$= D_{\text{KL}}(\mathbf{V} | \mathbf{W}\mathbf{H}) + R_{\boldsymbol{\alpha}}(\mathbf{H}, \boldsymbol{\beta}) + \text{cst} \quad (4)$$

where $D_{\text{KL}}(\cdot | \cdot)$ is the generalized Kullback-Leibler (KL) divergence defined by

$$D_{\text{KL}}(\mathbf{V} | \hat{\mathbf{V}}) = \sum_{f,n} \left(v_{fn} \log \left(\frac{v_{fn}}{\hat{v}_{fn}} \right) - v_{fn} + \hat{v}_{fn} \right) \quad (5)$$

and

$$R_{\boldsymbol{\alpha}}(\mathbf{H}, \boldsymbol{\beta}) = \sum_{k,n} [(1 - \alpha_k) \log(h_{kn}) + \beta_k h_{kn}] - N \sum_k \alpha_k \log \beta_k. \quad (6)$$

Equation (4) shows that MJLE is tantamount to penalized KL non-negative matrix factorization (NMF) (Lee & Seung, 2000) and may be addressed using alternating majorization-minimization (Canny, 2004; Févotte & Idier, 2011; Dikmen & Févotte, 2012).

As explained in Dikmen & Févotte (2012), MJLE can be criticized from a statistical point of view. Indeed, the number of estimated parameters grows with the number of samples N (this is because \mathbf{H} has as many columns as \mathbf{V}). To overcome this issue, they have instead proposed to consider maximum marginal likelihood estimation (MMLE), in which \mathbf{H} is treated as a latent variable over which the

joint likelihood is integrated. In other words, MMLE relies on the minimization of

$$C_{\text{ML}}(\mathbf{W}, \beta) \stackrel{\text{def}}{=} -\log p(\mathbf{V}|\mathbf{W}, \beta) \quad (7)$$

$$= -\log \int_{\mathbf{H}} p(\mathbf{V}|\mathbf{H}, \mathbf{W}, \beta) p(\mathbf{H}|\beta) d\mathbf{H}. \quad (8)$$

We emphasize that MMLE treats the dictionary \mathbf{W} as a free deterministic variable. This is in contrast with fully Bayesian approaches where \mathbf{W} is given a prior, and where estimation revolves around the posterior $p(\mathbf{W}, \mathbf{H}|\mathbf{V})$. For instance, [Buntine & Jakulin \(2006\)](#) place a Dirichlet prior on the columns of \mathbf{W} , while [Cemgil \(2009\)](#) considers independent Gamma priors. [Zhou et al. \(2012\)](#), [Zhou & Carin \(2015\)](#) set a Dirichlet prior on the columns of \mathbf{W} and a Gamma-based non-parametric Bayesian prior on \mathbf{H} , which allows for possible rank estimation.

[Dikmen & Févotte \(2012\)](#) assumed that a closed-form expression of C_{ML} was not available. Besides, they proposed variational and Monte Carlo Expectation-Maximization (MCEM) algorithms based on a complete set formed by \mathbf{H} and a set of other latent components \mathbf{C} that will later be defined. In their experiments, they found MMLE to be robust to over-specified values of K , while MJLE clearly overfit. This intriguing (and advantageous) behavior was left unexplained. In this paper, we provide the following contributions:

- We provide a computable closed-form expression of C_{ML} . The expression is tedious to compute for large F and K , as it involves combinatorial operations, but is workable for reasonably dimensioned problems.
- We show that the proposed closed-form expression reveals a penalization term on the columns of \mathbf{W} that explains the “self-regularization” effect observed in [Dikmen & Févotte \(2012\)](#).
- We show that the marginalization of \mathbf{H} allows to derive a new MCEM algorithm with favorable properties.

The rest of the paper is organized as follows. Section 2 introduces preliminary material (composite form of GaP, useful probability distributions). In Section 3, we propose two new parameterizations of the GaP model in which \mathbf{H} has been marginalized. This yields a closed-form expression of C_{ML} which is discussed in Section 4. Finally, a new MCEM algorithm is introduced in Section 5 and is compared to the MCEM algorithms proposed in [Dikmen & Févotte \(2012\)](#) on synthetic and real data.

2. Preliminaries

2.1. Composite structure of GaP

GaP can be written as a composite model, thanks to the superposition property of the Poisson distribution ([Févotte & Cemgil, 2009](#)):

$$h_{kn} \sim \text{Gamma}(\alpha_k, \beta_k) \quad (9)$$

$$\mathbf{c}_{kn} | h_{kn} \sim \prod_f \text{Poisson}(c_{fkn} | w_{fk} h_{kn}) \quad (10)$$

$$\mathbf{v}_n = \sum_k \mathbf{c}_{kn} \quad (11)$$

The vectors $\mathbf{c}_{kn} = [c_{1kn}, \dots, c_{Fkn}]^T$ of size F and which sum up to \mathbf{v}_n are referred to as *components*. In the remainder, \mathbf{C} will denote the $F \times K \times N$ tensor with coefficients c_{fkn} .

2.2. Negative Binomial and Negative Multinomial distributions

In this section, we introduce two probability distributions that will be used later in the article.

2.2.1. NEGATIVE BINOMIAL DISTRIBUTION

A discrete random variable X is said to have a negative binomial (NB) distribution with parameters $\alpha > 0$ (called the dispersion or shape parameter) and $p \in [0, 1]$ if, for all $c \in \mathbb{N}$, the probability mass function (p.m.f.) of X is given by:

$$\mathbb{P}(X = c) = \frac{\Gamma(\alpha + c)}{\Gamma(\alpha) c!} (1 - p)^\alpha p^c. \quad (12)$$

Its variance is $\frac{\alpha p}{(1-p)^2}$, which is larger than its mean, $\frac{\alpha p}{1-p}$. It is therefore a suitable distribution to model over-dispersed count data. Indeed, it offers more flexibility than the Poisson distribution where the variance and the mean are equal.

The NB distribution can be obtained via a Gamma-Poisson mixture, that is:

$$\mathbb{P}(X = c) = \int_{\mathbb{R}_+} \text{Poisson}(c|\lambda) \text{Gamma}(\lambda|\alpha, \beta) d\lambda \quad (13)$$

$$= \text{NB} \left(\alpha, \frac{1}{\beta + 1} \right). \quad (14)$$

2.2.2. NEGATIVE MULTINOMIAL DISTRIBUTION

The negative multinomial (NM) distribution ([Sibuya et al., 1964](#)) is the multivariate generalization of the NB distribution. It is parametrized by a dispersion parameter $\alpha > 0$ and a vector of event probabilities $\mathbf{p} = [p_1, \dots, p_F]^T$, where $0 \leq p_f \leq 1$ and $\sum_f p_f \leq 1$. Denoting $p_0 =$

$1 - \sum_f p_f$, and for all $(c_1, \dots, c_F) \in \mathbb{N}^F$, the p.m.f. of the NM distribution is given by

$$\mathbb{P}(X_1 = c_1, \dots, X_F = c_F) = \frac{\Gamma(\alpha + \sum_f c_f)}{\Gamma(\alpha) \prod_f c_f!} p_0^\alpha \prod_f p_f^{c_f}, \quad (15)$$

with expectation given by

$$\mathbb{E}(X) = \alpha \left[\frac{p_1}{p_0}, \dots, \frac{p_F}{p_0} \right]^T. \quad (16)$$

In particular, the NM distribution arises in the following Gamma-Poisson mixture, as detailed in the next proposition.

Proposition 1. Let $X = [X_1, \dots, X_F]^T$ be a random vector whose entries are independent Poisson random variables. Assume that each variable X_f is governed by the parameter $w_f \lambda$, where λ is itself a Gamma random variable with parameters (α, β) . Then the joint probability distribution of X is a NM distribution with dispersion parameter α and event probabilities

$$\mathbf{p} = \left[\frac{w_1}{\sum_f w_f + \beta}, \dots, \frac{w_F}{\sum_f w_f + \beta} \right]^T. \quad (17)$$

Proof. $\mathbb{P}(X = [c_1, \dots, c_F]^T)$

$$\begin{aligned} &= \int_{\mathbb{R}_+} \mathbb{P}(X_1 = c_1, \dots, X_F = c_F | \lambda) p(\lambda) d\lambda \\ &= \int_{\mathbb{R}_+} \left(\prod_f \frac{(w_f \lambda)^{c_f} e^{-w_f \lambda}}{c_f!} \right) \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta \lambda} d\lambda \\ &= \left(\prod_f \frac{w_f^{c_f}}{c_f!} \right) \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{\Gamma(\alpha + \sum_f c_f)}{(\sum_f w_f + \beta)^{\alpha + \sum_f c_f}} \\ &= \frac{\Gamma(\alpha + \sum_f c_f)}{\Gamma(\alpha) \prod_f c_f!} \left(\frac{\beta}{\sum_f w_f + \beta} \right)^\alpha \prod_f \left(\frac{w_f}{\sum_f w_f + \beta} \right)^{c_f} \end{aligned}$$

□

The NM distribution can also be obtained with an alternative generative process, as shown in the following proposition.

Proposition 2. Let $Y = [Y_1, \dots, Y_F]^T$ be a random vector following a multinomial distribution with number of trials L and event probabilities $\mathbf{p} = [\frac{w_1}{\sum_f w_f}, \dots, \frac{w_F}{\sum_f w_f}]^T$. Assume that L is a NB random variable with dispersion parameter α and probability $p = \frac{\sum_f w_f}{\sum_f w_f + \beta}$. Then the random vectors Y and X (as defined in Proposition 1) have the same distribution.

Proof. $\mathbb{P}(Y = [c_1, \dots, c_F]^T)$

$$\begin{aligned} &= \mathbb{P}(Y = [c_1, \dots, c_F]^T | L) \times \mathbb{P}(L) \\ &= \frac{L!}{\prod_f c_f!} \prod_f \left(\frac{w_f}{\sum_f w_f} \right)^{c_f} \\ &\quad \times \frac{\Gamma(\alpha + L)}{\Gamma(\alpha) L!} \left(\frac{\beta}{\sum_f w_f + \beta} \right)^\alpha \left(\frac{\sum_f w_f}{\sum_f w_f + \beta} \right)^L \end{aligned}$$

Noting that $L = \sum_f c_f$ completes the proof. □

3. New formulations of GaP

We now show how GaP can be rewritten free of the latent variables \mathbf{H} in two different ways.

3.1. GaP as a composite NM model

Theorem 1. GaP can be rewritten as follows:

$$\mathbf{c}_{kn} \sim \text{NM} \left(\alpha_k, \left[\frac{w_{1k}}{\sum_f w_{fk} + \beta_k}, \dots, \frac{w_{Fk}}{\sum_f w_{fk} + \beta_k} \right]^T \right) \quad (18)$$

$$\mathbf{v}_n = \sum_k \mathbf{c}_{kn} \quad (19)$$

Proof. Combining Equations (9)-(11) with Proposition 1 completes the proof. □

GaP may thus be interpreted as a composite model in which the k^{th} component has a NM distribution with parameters governed by \mathbf{w}_k (the k^{th} column of \mathbf{W}), α_k and β_k . Using straightforward computations, the data expectation can be expressed as

$$\mathbb{E}(\mathbf{v}_n) = \sum_k \mathbb{E}(\mathbf{c}_{kn}) \quad (20)$$

$$= \sum_k \frac{\alpha_k}{\beta_k} \mathbf{w}_k. \quad (21)$$

3.2. GaP as a composite multinomial model

Theorem 2. GaP can be rewritten as follows:

$$L_{kn} \sim \text{NB} \left(\alpha_k, \frac{\sum_f w_{fk}}{\sum_f w_{fk} + \beta_k} \right) \quad (22)$$

$$\mathbf{c}_{kn} \sim \text{Mult} \left(L_{kn}, \left[\frac{w_{1k}}{\sum_f w_{fk}}, \dots, \frac{w_{Fk}}{\sum_f w_{fk}} \right]^T \right) \quad (23)$$

$$\mathbf{v}_n = \sum_k \mathbf{c}_{kn} \quad (24)$$

where ‘‘Mult’’ refers to the multinomial distribution.

Proof. Combining Equations (9)-(11) with Proposition 2 completes the proof. \square

Theorem 2 states that another interpretation of GaP consists in modeling the data as a sum of k independent multinomial distributions, governed individually by \mathbf{w}_k and whose number of trials is random, following a NB distribution governed by \mathbf{w}_k , α_k and β_k .

A special case of the reformulation of GaP offered by Theorem 2 is given by Buntine & Jakulin (2006) using a different reasoning, when it is assumed that $\sum_f w_{fk} = 1$ (a common assumption in the field of text information retrieval, where the columns of \mathbf{W} are interpreted as discrete probability distributions). Theorem 2 provides a more general result as it applies to any non-negative matrix \mathbf{W} .

4. Closed-form marginal likelihood

4.1. Analytical expression

Until now, it was assumed that the marginal likelihood in the GaP model was not available analytically. However, the new parametrization offered by Theorem 1 allows to obtain a computable analytical expression of the marginal likelihood C_{ML} . Denote by \mathcal{C} the set of all ‘‘admissible’’ components, i.e.,

$$\mathcal{C} = \{\mathbf{C} \in \mathbb{N}^{F \times K \times N} \mid \forall(f, n), \sum_k c_{fkn} = v_{fn}\}. \quad (25)$$

By marginalization of \mathbf{C} , we may write

$$p(\mathbf{V}|\mathbf{W}, \boldsymbol{\beta}) = \sum_{\mathbf{C} \in \mathcal{C}} p(\mathbf{C}|\mathbf{W}, \boldsymbol{\beta}) \quad (26)$$

$$= \sum_{\mathbf{C} \in \mathcal{C}} \prod_{k,n} p(\mathbf{c}_{kn}|\mathbf{W}, \boldsymbol{\beta}). \quad (27)$$

Using Equation (18) we obtain

$$\begin{aligned} p(\mathbf{V}|\mathbf{W}, \boldsymbol{\beta}) &= \\ & \sum_{\mathbf{C} \in \mathcal{C}} \prod_{k,n} \left[\frac{\Gamma(\sum_f c_{fkn} + \alpha_k)}{\Gamma(\alpha_k) \prod_f c_{fkn}!} \left(\frac{\beta_k}{\sum_f w_{fk} + \beta_k} \right)^{\alpha_k} \right. \\ & \left. \times \prod_f \left(\frac{w_{fk}}{\sum_f w_{fk} + \beta_k} \right)^{c_{fkn}} \right]. \end{aligned} \quad (28)$$

Introducing the notations

$$\Omega_{\boldsymbol{\alpha}}(\mathbf{C}) = \prod_{k,n} \frac{\Gamma(\sum_f c_{fkn} + \alpha_k)}{\Gamma(\alpha_k) \prod_f c_{fkn}!} \quad (29)$$

and

$$p_{fk} = \frac{w_{fk}}{\sum_f w_{fk} + \beta_k} \quad (30)$$

we may rewrite Eq. (28) as

$$\begin{aligned} p(\mathbf{V}|\mathbf{W}, \boldsymbol{\beta}) &= \left[\prod_k (1 - \sum_f p_{fk})^{N\alpha_k} \right] \\ & \times \sum_{\mathbf{C} \in \mathcal{C}} \left[\Omega_{\boldsymbol{\alpha}}(\mathbf{C}) \prod_{f,k} p_{fk}^{\sum_n c_{fkn}} \right]. \end{aligned} \quad (31)$$

Equation (31) is a computable closed-form expression of the marginal likelihood. It is free of \mathbf{H} and in particular of the integral that appears in Equation (8). However the expression (31) is still semi-explicit because it involves a sum over the set of all admissible components \mathcal{C} . \mathcal{C} is countable set with cardinality $\#\mathcal{C} = \prod_{f,n} \binom{v_{fn}+K-1}{K-1}$. It is straightforward to construct but challenging to compute in large dimension, and for large values of v_{fn} .

The sum over all the matrices in the set \mathcal{C} expresses the convolution of the (discrete) probability distributions of the K components. Unfortunately, the distribution of the sum of independent negative multinomial variables of different event probabilities is not available in closed form.

As already known from Dikmen & Févotte (2012), the value of the marginal likelihood is unchanged when the scales of the columns of \mathbf{W} and the rates $\boldsymbol{\beta}$ are changed accordingly. Let $\boldsymbol{\Lambda}$ be a non-negative diagonal matrix of size K , it can easily be derived from Equation (31) that

$$p(\mathbf{V}|\mathbf{W}\boldsymbol{\Lambda}, \boldsymbol{\beta}\boldsymbol{\Lambda}) = p(\mathbf{V}|\mathbf{W}, \boldsymbol{\beta}). \quad (32)$$

We therefore have a scaling invariance between \mathbf{W} and $\boldsymbol{\beta}$, and as such, we may fix $\boldsymbol{\beta}$ to arbitrary values and leave \mathbf{W} free. Thus, we will treat $\boldsymbol{\beta}$ as a constant in the following and drop it from the arguments of C_{ML} .

4.2. Self-regularization

Dikmen & Févotte (2012) empirically studied the properties of MMLE. In particular, they observed the self-ability of the estimator to regularize the number of columns of \mathbf{W} . For example, one experiment consisted in generating synthetic data according to the GaP model, with a ground-truth number of components K^* . MMLE was run with $K > K^*$ and they noticed that the estimated \mathbf{W} contained $K - K^*$ empty columns. As such, the estimator was able to recover the ground-truth dimensionality. In contrast, MJLE used all K dimensions and overfit the data. They were unable to give a theoretical justification of the observed phenomenon, but provided a first insight thanks to a Laplace approximation of $p(\mathbf{V}|\mathbf{W})$. The closed-form expression (31) offers

a deeper understanding of this phenomenon, as explained next.

Using Equations (31) and (30) and treating β as a constant, the negative log-likelihood can be expressed as

$$\frac{1}{N} C_{\text{ML}}(\mathbf{W}) = -\frac{1}{N} \log \left(\sum_{\mathbf{C} \in \mathcal{C}} \Omega_{\alpha}(\mathbf{C}) \prod_{f,k} p_{f_k}^{\sum_n c_{fkn}} \right) \quad (33)$$

$$+ \sum_k \alpha_k \log(\|\mathbf{w}_k\|_1 + \beta_k) + \text{cst}, \quad (34)$$

where $\text{cst} = -\sum_k \alpha_k \log \beta_k$.

The negative log-likelihood reveals two terms. The first term, Equation (33), captures the interaction between data \mathbf{V} (through \mathbf{C}) and the parameter \mathbf{W} (through the event probabilities $p_{f_k} = w_{f_k}/(\|\mathbf{w}_k\|_1 + \beta_k)$). The second term, Equation (34), only depends on the parameter \mathbf{W} and can be interpreted as a group-regularization term. The non-convex and sharply peaked function $f(\mathbf{x}) = \sum_k \log(x_k + b)$ is known to be sparsity-inducing (Candès et al., 2008). As such, the term (34) will promote sparsity of the norms of the columns of \mathbf{W} . When a norm $\|\mathbf{w}_k\|_1$ is set to zero for some k , the whole column \mathbf{w}_k is set to zero because of the non-negativity constraint. This gives a formal explanation of the ability of MMLE to automatically prune columns of \mathbf{W} , without any explicit sparsity-inducing prior at the modeling stage (recall that \mathbf{W} is a deterministic parameter without a prior).

5. MCEM algorithms form MMLE

5.1. Expectation-Maximization

We now turn to the problem of optimizing Equation (8) by leveraging on the results of Section 4. Despite obtaining a closed-form expression, the direct optimization of the marginal likelihood remains difficult. However, the structure of GaP makes Expectation-Maximization (EM) a natural option (Dempster et al., 1977). Indeed, GaP involves observed variables \mathbf{V} and latent variables \mathbf{C} and \mathbf{H} . As such, we can derive several EM algorithms based on various choices of the complete set. More precisely, we consider three possible choices that each define a different algorithm, as follows.

EM-CH. The complete set is $\{\mathbf{C}, \mathbf{H}\}$ and EM consists in the iterative minimization w.r.t \mathbf{W} of the functional defined by

$$Q_{\text{CH}}(\mathbf{W}|\tilde{\mathbf{W}}) = - \int_{\mathbf{C}, \mathbf{H}} \log p(\mathbf{C}, \mathbf{H}|\mathbf{W}) p(\mathbf{C}, \mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}}) d\mathbf{C} d\mathbf{H}, \quad (35)$$

where $\tilde{\mathbf{W}}$ is the current estimate. Note that \mathbf{V} does not need to be included in the complete set because we have $\mathbf{V} = \sum_k \mathbf{C}_k$. This corresponds to the general formulation of EM in which the relation between the complete set and the data is a many-to-one mapping and slightly differs from the more usual one where the complete set is formed by the union of data and a hidden set (Dempster et al., 1977; Févotte et al., 2011).

EM-H. The complete set is $\{\mathbf{V}, \mathbf{H}\}$ and EM consists in the iterative minimization of

$$Q_{\text{H}}(\mathbf{W}|\tilde{\mathbf{W}}) = - \int_{\mathbf{H}} \log p(\mathbf{V}, \mathbf{H}|\mathbf{W}) p(\mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}}) d\mathbf{H}. \quad (36)$$

EM-C. The complete set is merely $\{\mathbf{C}\}$ and EM consists in the iterative minimization of

$$Q_{\text{C}}(\mathbf{W}|\tilde{\mathbf{W}}) = - \int_{\mathbf{C}} \log p(\mathbf{C}|\mathbf{W}) p(\mathbf{C}|\mathbf{V}, \tilde{\mathbf{W}}) d\mathbf{C}. \quad (37)$$

EM-CH and EM-H have been considered in Dikmen & Févotte (2012). EM-C is a new proposal that exploits the results of Section 4. In all three cases, the posteriors of the latent variables involved – $p(\mathbf{C}, \mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$, $p(\mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$ and $p(\mathbf{C}|\mathbf{V}, \tilde{\mathbf{W}})$ – are untractable and neither are the integrals involved in Equations (35), (36) and (37). To overcome this problem, we resort to Monte Carlo EM (MCEM) (Wei & Tanner, 1990) as described in the next section.

5.2. Monte Carlo E-step

MCEM consists in using a Monte Carlo (MC) approximation of the integrals in Equations (35), (36) and (37) based on samples drawn from the posterior distributions $p(\mathbf{C}, \mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$, $p(\mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$ and $p(\mathbf{C}|\mathbf{V}, \tilde{\mathbf{W}})$. These can be obtained by Gibbs sampling of the joint posterior $p(\mathbf{C}, \mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$, which also returns samples from the marginals $p(\mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$ and $p(\mathbf{C}|\mathbf{V}, \tilde{\mathbf{W}})$ at convergence. The Gibbs sampler can easily be derived because the conditional distributions $p(\mathbf{H}|\mathbf{C}, \mathbf{V}, \tilde{\mathbf{W}}) = p(\mathbf{H}|\mathbf{C}, \tilde{\mathbf{W}})$ and $p(\mathbf{C}|\mathbf{H}, \mathbf{V}, \tilde{\mathbf{W}})$ are available in closed form.

At iteration $j + 1$, Gibbs sampling writes

$$h_{kn}^{(j+1)} \sim \text{Gamma}(\alpha_k + \sum_f c_{fkn}^{(j)}, \beta_k + \sum_f \tilde{w}_{fk}) \quad (38)$$

$$\mathbf{c}_{fn}^{(j+1)} \sim \text{Mult} \left(v_{fn}, \left[\rho_{f1}^{(j+1)}, \dots, \rho_{fK}^{(j+1)} \right]^T \right) \quad (39)$$

where \mathbf{c}_{fn} denotes the vector $[c_{f1n}, \dots, c_{fKn}]^T$ of size K and

$$\rho_{fk}^{(j+1)} = \frac{\tilde{w}_{fk} h_{kn}^{(j+1)}}{\sum_k \tilde{w}_{fk} h_{kn}^{(j+1)}}. \quad (40)$$

Note that $\mathbf{c}_{fn}^{(j+1)}$ only needs to be sampled when $v_{fn} \neq 0$, since $\mathbf{c}_{fn}^{(j+1)} = [0, \dots, 0]^T$ when $v_{fn} = 0$.

5.3. M-step

Given a set of J samples $\{\mathbf{H}^{(j)}, \mathbf{C}^{(j)}\}$ drawn from $p(\mathbf{C}, \mathbf{H}|\mathbf{V}, \tilde{\mathbf{W}})$ (after burn-in), minimization of the MC approximation of Q_{CH} in Equation (35) yields

$$w_{fk}^{\text{MCEM-CH}} = \frac{\sum_{j,n} c_{fkn}^{(j)}}{\sum_{j,n} h_{kn}^{(j)}}, \quad (41)$$

as shown by [Dikmen & Févotte \(2012\)](#). They also show that the following multiplicative update decreases the MC approximation of Q_{H} in Equation (36) at every iteration

$$w_{fk}^{\text{MCEM-H}} = \tilde{w}_{fk} \frac{\sum_{j,n} h_{kn}^{(j)} v_{fn} [\tilde{\mathbf{W}}\mathbf{H}^{(j)}]_{fn}^{-1}}{\sum_{j,n} h_{kn}^{(j)}}. \quad (42)$$

We now derive the new update for EM-C. The MC approximation of Q_{C} in Equation (37) is given by:

$$\hat{Q}_{\text{C}}(\mathbf{W}|\tilde{\mathbf{W}}) \stackrel{\text{def}}{=} -\frac{1}{J} \sum_{j=1}^J \log p(\mathbf{C}^{(j)}|\mathbf{W}). \quad (43)$$

Replacing $p(\mathbf{C}^{(j)}|\mathbf{W})$ by its expression given by Equation (18), we obtain:

$$\begin{aligned} \hat{Q}_{\text{C}}(\mathbf{W}|\tilde{\mathbf{W}}) &= \frac{1}{J} \sum_{j,k,n} \left[\alpha_k \log \left(\sum_f w_{fk} + \beta_k \right) \right. \\ &\quad \left. + \sum_f c_{fkn}^{(j)} \left(\log \left(\sum_f w_{fk} + \beta_k \right) - \log(w_{fk}) \right) + \text{cst} \right]. \end{aligned} \quad (44)$$

The minimization of \hat{Q}_{C} w.r.t. \mathbf{W} leads to K linear systems of equations that we need to solve for each column \mathbf{w}_k :

$$\mathbf{A}_k \mathbf{w}_k = \mathbf{b}_k. \quad (45)$$

The matrix $\mathbf{A}_k \in \mathbb{R}^{F \times F}$ is defined by:

$$a_{fg} = \left(JN\alpha_k + \sum_{j,f,n} c_{fkn}^{(j)} \right) \delta_{fg} - \sum_{j,n} c_{fkn}^{(j)}, \quad (46)$$

where δ_{fg} is the Kronecker symbol, i.e. $\delta_{fg} = 1$ if and only if $f = g$, and zero otherwise. The vector $\mathbf{b}_k \in \mathbb{R}^{F \times 1}$ is defined by:

$$b_{fk} = \beta_k \sum_{j,n} c_{fkn}^{(j)}. \quad (47)$$

The matrix \mathbf{A}_k appears to be the sum of a diagonal matrix with a rank-1 matrix and can be inverted analytically thanks




Figure 1. Illustrative example with $F = 1$ and $K = 2$. A dataset of $N = 300$ (scalar) samples is generated from GaP. The contour plot displays $C_{\text{ML}}(\mathbf{W})$ (thanks to its closed-form expression). We set $\alpha_k = 1$ and $\beta_k = 1$. The red, green and blue lines display the iterates of $\mathbf{W} = [w_1, w_2]$ for the last 80 iterations (out of 500) of MCEM-C, MCEM-H and MCEM-CH, respectively. The iterates generated by MCEM-C satisfy the constraint $w_1 + w_2 = \bar{v}$, as described by Equation (49).

to the Sherman-Morrison formula ([Sherman & Morrison, 1950](#)). This results in the new closed-form update

$$w_{fk}^{\text{MCEM-C}} = \frac{1}{JN} \frac{\beta_k}{\alpha_k} \sum_{j,n} c_{fkn}^{(j)}. \quad (48)$$

In the common case where $\frac{\beta_k}{\alpha_k}$ is equal to a constant γ for all k , we can write:

$$\sum_k w_{fk}^{\text{MCEM-C}} = \sum_k \frac{\beta_k}{\alpha_k} \frac{\sum_{j,n} c_{fkn}^{(j)}}{JN} = \gamma \frac{\sum_{j,n} v_{fn}}{JN} = \gamma \bar{v}_f \quad (49)$$

where $\bar{v}_f = N^{-1} \sum_n v_{fn}$ is the empirical mean of the data for the feature f . Equation (49) implies that the rows of the estimate $\mathbf{W}^{\text{MCEM-C}}$ at every iteration sum to a constant. This behavior is illustrated on Figure 1.

6. Experimental results

We now compare the three MCEM algorithms proposed for MMLE in the GaP model, first using synthetic toy datasets, then real-world data. Python implementations of the three algorithms are available from the first author website.




Figure 2. $C_{ML}(\mathbf{W})$ with respect to CPU time in seconds for the three MCEM algorithms (dataset \mathbf{V}_1).

6.1. Experiments with synthetic data

We generate a dataset of $N = 100$ samples according to the GaP model, with the following parameters:

$$\mathbf{W}_1^* = \begin{bmatrix} 0.638 & 0.075 \\ 0.009 & 0.568 \\ 0.044 & 0.126 \\ 0.309 & 0.231 \end{bmatrix}, \quad \alpha^* = \beta^* = 1. \quad (50)$$

The columns of \mathbf{W}_1^* have been generated from a Dirichlet distribution (with parameters 1). The generated dataset (of size 4×100) is denoted by \mathbf{V}_1 .

We proceed to estimate the dictionary \mathbf{W} using hyperparameters $K = K^* + 1 = 3$, $\alpha_k = \beta_k = 1$ with MCEM-C, MCEM-H and MCEM-CH. The algorithms are run for 500 iterations. 300 Gibbs samples are generated at each iteration, with the first 150 samples being discarded for burn-in (this proves to be enough in practice), leading to $J = 150$. The Gibbs sampler at EM iteration $i + 1$ is initialized with the last sample obtained at EM iteration i (warm restart). The algorithms are initialized from the same deterministic starting point given by

$$w_{fk} = \frac{1}{K} \frac{\beta_k}{\alpha_k} \bar{v}_f, \quad (51)$$

as suggested by Equation (49).

Figure 2 displays the negative log-likelihood $C_{ML}(\mathbf{W})$, computed thanks to the derivations of Section 4, and Figure 3-(a) displays the norm of the three columns of the iterates, both w.r.t. CPU time in seconds. The three algorithms have almost identical computation times, most of the computational burden residing in the Gibbs sampling procedure that is common to the three algorithms. Moreover, the three algorithms converge to the same point, with MCEM-C converging marginally faster than the other two in this case.

Then we proceed to generate a second dataset \mathbf{V}_2 according to the GaP model, with now $\mathbf{W}_2^* = 100 \times \mathbf{W}_1^*$. The




Figure 3. Evolution of the norm of each of the $K = 3$ columns \mathbf{w}_k of the dictionaries w.r.t. CPU time for the three MCEM algorithms. Top : dataset \mathbf{V}_1 , generated according to the GaP model with \mathbf{W}_1^* . Bottom : dataset \mathbf{V}_2 , generated according to the GaP model with \mathbf{W}_2^* .

expectation of \mathbf{V}_2 is thus a hundred times the expectation of \mathbf{V}_1 . We apply the exact same experimental protocol to \mathbf{V}_2 as we did for \mathbf{V}_1 , except that the algorithms are now run for a larger number of 2000 iterations. In this case, because of the combinatorial nature of $\#\mathcal{C}$, it is impossible to compute the likelihood in reasonable time. The norms of the columns of the iterates are displayed on Figure 3-(b). As we can see, MCEM-C clearly outperforms the other two algorithms in this scenario. This behavior has been consistently found when estimating dictionaries from datasets with sufficiently large values. This drastic difference in convergence is unexplained at this stage. However, we conjecture that this is linked to the over-dispersion of the data (i.e. when the variance is greater than the mean), which increases with the scale of \mathbf{W} . This will be studied in future work.

6.2. Experiments with the NIPS dataset

Finally, we consider the NIPS dataset which contains word counts from a collection of articles published at the NIPS conference.¹ The number of articles is $N = 1,500$ and the number of unique terms (appearing at least 10 times after tokenization and removing stop-words) is $F = 12,419$. The matrix \mathbf{V} is quite sparse as 96% of its coefficients are zeros. This saves a large amount of computational effort, because we only need to sample \underline{c}_{fn} for pairs (f, n) such that v_{fn} is non-zero. Moreover, the count values range from 0 to 132.

We applied MCEM-C and MCEM-CH with $K = 10$ and $\alpha_k = \beta_k = 1$. The algorithms are run for 1,000 iterations. 250 Gibbs samples are generated in each iteration with the first half being discarded for burn-in (i.e. $J = 125$). The Gibbs sampler at iteration $i + 1$ is again initialized with warm restart. The algorithms are initialized using Equation (51). MCEM-H results in similar performance than MCEM-CH and is not reported here.

Figure 4 shows that the column norms of the iterates w.r.t. CPU time in seconds. The difference in convergence speed between the two algorithms is again striking. MCEM-C efficiently explores the parameter space in the first iterations and converges dramatically faster than MCEM-CH. The algorithms here converge to different solutions, which confirms the non-convexity of $C_{ML}(\mathbf{W})$. Other runs confirmed that MCEM-C is consistently faster, and also that the two algorithms do not always converge to the same solution.

MCEM-C still takes a few hours to converge. We implemented stochastic variants of the EM algorithm such as SAEM (Kuhn & Lavielle, 2004), however this did not result in significant improvements in our case. Finally, note that large-scale implementations are possible using for example on-line EM (Cappé & Moulines, 2009).

7. Conclusion

In this paper, we have shown how the Gamma-Poisson model can be rewritten free of the latent variables \mathbf{H} . This new parametrization enabled us to come up with a closed-form expression of the marginal likelihood, which revealed a penalization term explaining the “self-regularization” effect described in Dikmen & Févotte (2012). We then proceeded to compare three MCEM algorithms for the task of maximizing the likelihood, and the algorithm taking advantage of the marginalization \mathbf{H} has been empirically proven to have better convergence properties.

In this work, we treated α as a fixed hyperparameter. Fu-

¹<https://archive.ics.uci.edu/ml/datasets/bag+of+words>




Figure 4. Evolution of the norm of each of the $K = 10$ columns of the dictionaries w.r.t. CPU time for MCEM-C and MCEM-CH (NIPS dataset).

ture work will focus on lifting this hypothesis, and on designing algorithms to estimate both \mathbf{W} and α . We will also look into carrying a similar analysis in other probabilistic matrix factorization models, such as the Gamma-Exponential model (Dikmen & Févotte, 2011).

Acknowledgments

We thank Benjamin Guedj, Víctor Elvira and Jérôme Idier for fruitful discussions related to this work. This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under grant agreement No 681839 (project FACTORY).

References

- Buntine, W. and Jakulin, A. Discrete component analysis. In *Subspace, Latent Structure and Feature Selection*, pp. 1–33. Springer, 2006.
- Candès, E. J., Wakin, M. B., and Boyd, S. P. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier analysis and applications*, 14(5):877–905, 2008.
- Canny, J. GaP: a factor model for discrete data. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 122–129, 2004.
- Cappé, O. and Moulines, E. On-line expectation-maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):593–613, 2009.
- Cemgil, A. T. Bayesian inference for nonnegative matrix

- factorisation models. *Computational Intelligence and Neuroscience*, (Article ID 785152), 2009.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 39(1):1–38, 1977.
- Dikmen, O. and Févotte, C. Nonnegative dictionary learning in the exponential noise model for adaptive music signal representation. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2267–2275, 2011.
- Dikmen, O. and Févotte, C. Maximum marginal likelihood estimation for nonnegative dictionary learning in the gamma-poisson model. *IEEE Transactions on Signal Processing*, 60(10):5163–5175, 2012.
- Févotte, C. and Cemgil, A. T. Nonnegative matrix factorizations as probabilistic inference in composite models. In *Proceedings of the European Signal Processing Conference (EUSIPCO)*, pp. 1913–1917, 2009.
- Févotte, C. and Idier, J. Algorithms for nonnegative matrix factorization with the β -divergence. *Neural computation*, 23(9):2421–2456, 2011.
- Févotte, C., Cappé, O., and Cemgil, A. T. Efficient Markov chain Monte Carlo inference in composite models with space alternating data augmentation. In *Proceedings of the IEEE Workshop on Statistical Signal Processing (SSP)*, pp. 221–224, 2011.
- Gopalan, P., Hofman, J. M., and Blei, D. M. Scalable recommendation with hierarchical poisson factorization. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 326–335, 2015.
- Kuhn, E. and Lavielle, M. Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM: Probability and Statistics*, 8:115–131, 2004.
- Lee, D. D. and Seung, H. S. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 556–562, 2000.
- Ma, H., Liu, C., King, I., and Lyu, M. R. Probabilistic factor models for web site recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 265–274, 2011.
- Sherman, J. and Morrison, W. J. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950.
- Sibuya, M., Yoshimura, I., and Shimizu, R. Negative multinomial distribution. *Annals of the Institute of Statistical Mathematics*, 16(1):409–426, 1964.
- Virtanen, T., Cemgil, A. T., and Godsill, S. Bayesian extensions to non-negative matrix factorisation for audio signal modelling. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1825–1828, 2008.
- Wei, G. C. and Tanner, M. A. A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *Journal of the American statistical Association*, 85(411):699–704, 1990.
- Zhou, M. and Carin, L. Negative binomial process count and mixture modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):307–320, 2015.
- Zhou, M., Hannah, L., Dunson, D., and Carin, L. Beta-negative binomial process and Poisson factor analysis. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 1462–1471, 2012.