



Deep mixtures of unigrams for uncovering topics in textual data

Cinzia Viroli¹ · Laura Anderlucci¹

Received: 1 June 2020 / Accepted: 17 September 2020
© The Author(s) 2021

Abstract

Mixtures of unigrams are one of the simplest and most efficient tools for clustering textual data, as they assume that documents related to the same topic have similar distributions of terms, naturally described by multinomials. When the classification task is particularly challenging, such as when the document-term matrix is high-dimensional and extremely sparse, a more composite representation can provide better insight into the grouping structure. In this work, we developed a deep version of mixtures of unigrams for the unsupervised classification of very short documents with a large number of terms, by allowing for models with further deeper latent layers; the proposal is derived in a Bayesian framework. The behavior of the deep mixtures of unigrams is empirically compared with that of other traditional and state-of-the-art methods, namely k -means with cosine distance, k -means with Euclidean distance on data transformed according to semantic analysis, partition around medoids, mixture of Gaussians on semantic-based transformed data, hierarchical clustering according to Ward's method with cosine dissimilarity, latent Dirichlet allocation, mixtures of unigrams estimated via the EM algorithm, spectral clustering and affinity propagation clustering. The performance is evaluated in terms of both correct classification rate and Adjusted Rand Index. Simulation studies and real data analysis prove that going deep in clustering such data highly improves the classification accuracy.

Keywords Deep learning · Mixture models · Clustering · Text data analysis

1 Introduction

Deep learning methods are receiving an exponentially increasing interest in the last years as powerful tools to learn complex representations of data. They can be basically defined as a multilayer stack of algorithms or modules able to gradually learn a huge number of parameters in an architecture composed by multiple nonlinear transformations (LeCun et al. 2015). Typically, and for historical reasons, a structure for deep learning is identified with advanced neural networks: deep feed forward, recurrent, auto-encoder, convolution neural networks are very effective and used algorithms of deep learning (Schmidhuber 2015). They demonstrated to be particularly successful in supervised classification problems arising in several fields such as image and speech recognition, gene expression data, topic classification. Under the framework of graph-based learning, Peng et al. (2016) proposed an efficient method to produce robust subspace clustering and

subspace learning; deep model-Structured AutoEncoder for subspace clustering was introduced by Peng et al. (2018) to map input data points into nonlinear latent spaces while preserving the local and global subspace structure. Zhou et al. (2018) addressed the data sparsity issue in hashing.

When the aim is uncovering unknown classes in an unsupervised classification perspective, important methods of deep learning have been developed along the lines of mixture modeling, because of their ability to decompose a heterogeneous collection of units into a finite number of sub-groups with homogeneous structures (Fraley and Raftery 2002; McLachlan and Peel 2000). In this direction, van den Oord and Schrauwen (2014) proposed Multilayer Gaussian Mixture Models for modeling natural images; Tang et al. (2012) defined deep mixture of factor analyzers with a greedy layer-wise learning algorithm able to learn each layer at a time. Viroli and McLachlan (2019) developed a general framework for deep Gaussian mixture models that generalizes and encompasses the previous strategies and several flexible model-based clustering methods such as mixtures of mixture models (Li 2005), mixtures of Factor Analyzers (McLachlan et al. 2003), mixtures of factor analyzers with common factor

✉ Cinzia Viroli
cinzia.viroli@unibo.it

¹ Department of Statistical Sciences, University of Bologna, via Belle Arti 41, 40126 Bologna, Italy

loadings (Baek et al. 2010), heteroscedastic factor mixture analysis (Montanari and Viroli 2010) and mixtures of factor mixture analyzers introduced by Viroli (2010). A general ‘take-home-message’ coming from the existing deep clustering strategies is that deep methods vs shallow ones appear to be very efficient and powerful tools especially for complex high-dimensional data; on the contrary, for simple and small data structures, a deep learning strategy cannot improve performance of simpler and conventional methods or, to better say, it is like to use a ‘sledgehammer to crack a nut’ (Viroli and McLachlan 2019).

The motivating problem behind this work derives from ticket data (i.e. content of calls made to the customer service) of an important mobile phone company, collected in Italian language. When a customer calls the assistance service, a ticket is created: the agent classifies it as, e.g. a claim, a request of information for specific services, deals or promotions. Our dataset consists of tickets related to five classes of services, previously classified from independent operators. The aim is to define an efficient clustering strategy to automatically assign the tickets into the same classes without the human judgment of an operator. The data are textual and information are collected in a document-term matrix with raw frequencies at each cell. They have a very complex and a high-dimensional structure, caused by the huge number of tickets and terms used by people that call the company for a specific request and by a relevant degree of sparsity (after a preprocessing step, the tickets have indeed an average length of only 5 words and, thus, the document-term matrix contains zero almost everywhere).

The simplest topic model for clustering document-term data is represented by mixture of unigrams (MoU) (Nigam et al. 2000) MoUs are based on the idea of modeling the word frequencies as multinomial distributions, under the assumption that a document refers to a single topic. Table 5 (last column) shows that the method appears to be the most efficient tool for classifying the complex ticket data, compared to other conventional clustering strategies such as k -means, partition around medoids and hierarchical clustering. The reason is probably related to the fact that, by using proportions, MoU is not affected by the large amount of zeros, differently from the other competitors. We also compared MoU with the latent Dirichlet allocation model (LDA) (Blei et al. 2003), which represents an important and very popular model in textual data analysis, allowing documents to exhibit multiple topics with a different degree of importance. The latent Dirichlet allocation model has demonstrated great success on long texts (Griffiths and Steyvers 2004), and it could be thought of as a generalization of the MoU, because it adds a hierarchical level to it and, hence, much more flexibility. However, when dealing with very short documents like the ticket dataset, it is very rare that a single unit could

refer to more than one topic; in such cases, the LDA model may not improve the clustering performance of MoU.

The aim of this paper is to derive a deep generalization of mixtures of unigrams, in order to better uncover topics or groups in case of complex and high-dimensional data. The proposal will be derived in a Bayesian framework and we will show that it will be particularly efficient for classifying the ticket data with respect to the ‘shallow’ MoU model. We will also show the good performance of the proposed method in a simulation study.

The paper is organized as follows. In the next section, the mixture of unigrams model is described. In Sect. 3, the deep formulation of the model is developed. Section 4 is devoted to the estimation algorithm for fitting the model. Experimental results on simulated and real data are presented in Sects. 5 and 6, respectively. We conclude the paper with some final remarks (Sect. 7).

2 Mixtures of unigrams

Let \mathbf{X} be a document-term matrix of dimension $n \times T$ containing the word frequencies of each document in row and let k be the number of homogeneous groups in which documents could be allocated. Let \mathbf{x}_d be the single document of length T , with $d = 1, \dots, n$.

In MoU, the distribution of each document has a specific distribution function conditional on the values of a discrete latent allocation variable z_d describing the probability of each topic. More precisely,

$$p(\mathbf{x}_d) = \sum_{i=1}^k \pi_i p(\mathbf{x}_d | z_d = i), \quad (1)$$

with $p(z_d = i) = \pi_i$, under the restrictions (i) $\pi_i > 0$ for $i = 1, \dots, k$ and (ii) $\sum_{i=1}^k \pi_i = 1$.

The natural distribution for $p(\mathbf{x}_d | z_d = i)$ is represented by the multinomial distribution with a parameter vector, say ω_i , that is cluster-specific:

$$p(\mathbf{x}_d | z_d = i) = \frac{N_d!}{\prod_{t=1}^T x_{dt}!} \prod_{t=1}^T \omega_{it}^{x_{dt}},$$

with $N_d = \sum_{t=1}^T x_{dt}$ denoting the word-length of the document d . The multinomial distribution assumes that, conditionally to the cluster membership, all the terms can be regarded as independently distributed.

The model is indeed a simple mixture of multinomial distributions that can be easily estimated by the EM algorithm (see Nigam et al. 2000 for further details) under the assumption that a document belongs to a single topic and the number of groups coincides with the number of topics. The

approach has been successfully applied not only to textual data, where it originated from, but to genomic data analysis. In this latter field, a particular improvement of the method consists in relaxing the conditional independence assumption of the variables/terms, by using m -order Markov properties, thus leading to the so-called m -gram models (Tomović et al. 2006). Due to the limited average length of documents in ticket data, co-occurrence information is very rare and this extension would be not effective on these data.

3 Going deep into mixtures of unigrams: a novel approach

We aim at extending MoU by allowing a further layer in the probabilistic generative model, so as to get a nested architecture of nonlinear transformations able to describe the data structure in a very flexible way. At the deepest latent layer, the documents can come from k_2 groups with different probabilities, say $\pi_j^{(2)}$ $j = 1, \dots, k_2$. Conditionally to what happened at this level, at the top observable layer the documents can belong to k_1 groups with conditional probabilities $\pi_{ij}^{(1)}$ $i = 1, \dots, k_1$. For the sake of a simple notation, we refer in the following to a generic document denoted by \mathbf{x} . The distribution of \mathbf{x} conditionally to the two layers is a multinomial distribution with cluster-specific proportions ω :

$$p(\mathbf{x} | z^{(1)} = i, z^{(2)} = j, \omega) = \frac{(\sum_{t=1}^T x_t)!}{\prod_{t=1}^T x_t!} \prod_{t=1}^T \omega_{tij}^{x_t} \tag{2}$$

where $z^{(1)}$ and $z^{(2)}$ are the allocation variables at the top and at the bottom layers, respectively. They are discrete latent variables that follow the distributions:

$$p(z^{(2)} = j) = \pi_j^{(2)}, \tag{3}$$

and

$$p(z^{(1)} = i | z^{(2)} = j) = \pi_{ij}^{(1)}. \tag{4}$$

In mixture of unigrams, the proportions ω were fixed parameters. Here we assume they are realizations of random variables with a Dirichlet distribution. In order to have a flexible and deep model we assume that the parameters of the Dirichlet distribution are a linear function of two sets of parameters that originate in the two subsequent layers. The Dirichlet parameters are $\beta_i + \alpha_j \beta_i = \beta_i(1 + \alpha_j)$, where β_i and α_j are vectors of length T . Since they must be positive and the overall model must be identifiable, we further assume that $\beta_{it} > 0$ and $-1 < \alpha_{jt} < 1$. These choices lead to a nice




Fig. 1 Structure of a deep MoU with components $k_1 = 3$ and $k_2 = 2$

interpretative perspective: at the bottom layer α_j acts as a perturbation on the cluster-specific β_i parameters of the top layer. Therefore, the distribution is:

$$p(\omega | z^{(1)} = i, z^{(2)} = j) = \frac{\Gamma(\sum_{t=1}^T \beta_{it} (1 + \alpha_{jt}))}{\prod_{t=1}^T \Gamma(\beta_{it} (1 + \alpha_{jt}))} \prod_{t=1}^T \omega_{tij}^{\beta_{it}(1+\alpha_{jt})-1}, \tag{5}$$

where Γ denotes the Gamma function.

An example of deep MoU structure is depicted in Fig. 1 for the case $k_1 = 3$ and $k_2 = 2$. Notice that in a model with $k_1 = 3$ and $k_2 = 2$ components we have an overall number M of sub-components equal to 5, but $6 > M$ possible paths for each document. The paths share and combine the parameters of the two levels, thus achieving great flexibility with less parameters to be estimated.

By combining equations (2) and (5), the latent variable ω can be integrated out from the model estimation, thus gaining efficiency without losing flexibility and interpretability. More precisely:

$$\begin{aligned} p(\mathbf{x} | z^{(1)} = i, z^{(2)} = j) &= \int p(\mathbf{x} | z^{(1)} = i, z^{(2)} = j, \omega) \\ &= \int p(\omega | z^{(1)} = i, z^{(2)} = j) d\omega \\ &= \frac{(\sum_{t=1}^T x_t)!}{\prod_{t=1}^T x_t!} \frac{\Gamma(\sum_{t=1}^T \beta_{it} (1 + \alpha_{jt}))}{\prod_{t=1}^T \Gamma(\beta_{it} (1 + \alpha_{jt}))} \\ &\quad \times \int \prod_{t=1}^T \omega_{tij}^{x_t + \beta_{it}(1+\alpha_{jt})-1} d\omega_{ij} \\ &= \frac{(\sum_{t=1}^T x_t)!}{\prod_{t=1}^T x_t!} \frac{\Gamma(\sum_{t=1}^T \beta_{it} (1 + \alpha_{jt}))}{\prod_{t=1}^T \Gamma(\beta_{it} (1 + \alpha_{jt}))} \end{aligned}$$




Fig. 2 DAG specifying the deep MoU model

$$\begin{aligned} & \times \frac{\prod_{t=1}^T x_t + \Gamma(\beta_{it}(1 + \alpha_{jt}))}{\Gamma\left(\sum_{t=1}^T x_t + \beta_{it}(1 + \alpha_{jt})\right)} \\ & = \frac{\sum_{t=1}^T x_t}{\prod_{t=1}^T x_t} \frac{B\left(\sum_{t=1}^T x_t, \sum_{t=1}^T \beta_{it}(1 + \alpha_{jt})\right)}{\prod_{t=1}^T B(x_t, \beta_{it}(1 + \alpha_{jt}))}, \end{aligned} \tag{6}$$

where B denotes the Beta function.

Figure 2 shows the Directed Acyclic Graph (DAG) summarizing the dependence structure of the model.

Finally by combining formulas (6) with (3) and (4), the density of the data is a mixture of mixtures of Multinomial-Dirichlet distributions:

$$\begin{aligned} p(\mathbf{x}) & = \sum_{i=1}^{k_1} \sum_{j=1}^{k_2} \pi_{i|j}^{(1)} \pi_j^{(2)} \\ & \times \frac{\sum_{t=1}^T x_t}{\prod_{t=1}^T x_t} \frac{B\left(\sum_{t=1}^T x_t, \sum_{t=1}^T \beta_{it}(1 + \alpha_{jt})\right)}{\prod_{t=1}^T B(x_t, \beta_{it}(1 + \alpha_{jt}))}. \end{aligned} \tag{7}$$

In other terms, adding a level to the hierarchy resulted in adding a mixing step. A double mixture is deeper, more flexible and it can capture more heterogeneity of the data, than a simple mixture of Multinomial-Dirichlet distributions. Having two layers and two number of groups for each, that are k_1 and k_2 , it is important to define the procedure by which the units are clustered.

3.1 Cluster assignment

Theoretically, under this double mixtures, we could group units into k_1 groups, k_2 groups or $k_1 \times k_2$ groups. However, note that, under the constraint $-1 < \alpha_{jt} < 1$, the role of k_2 components at the deepest layer is confined to add flexibility to the model, while the real cluster-distribution is specified by the β parameters. For this reason, the number of ‘real’ clusters is given by k_1 . Their internal heterogeneity is captured by the k_2 sub-groups that help in adding more flexibility to the model. Therefore, the final allocation of the documents

to the clusters is given by the posterior probability $p(z^{(1)}|\mathbf{x})$ that can be obtained as follows:

$$\begin{aligned} p\left(z^{(1)} = i|\mathbf{x}\right) & = \frac{\sum_{j=1}^{k_2} \pi_{i|j}^{(1)} \pi_j^{(2)} p\left(\mathbf{x}|z^{(1)} = i, z^{(2)} = j\right)}{\sum_{i=1}^{k_1} \sum_{j=1}^{k_2} \pi_{i|j}^{(1)} \pi_j^{(2)} p\left(\mathbf{x}|z^{(1)} = i, z^{(2)} = j\right)}. \end{aligned} \tag{8}$$

The model encompasses the simple MoU, that can be obtained as special case when $k_2 = 1$ and without any prior on ω . When $k_2 = 1$ and a Dirichlet prior is put on ω , a mixture of Dirichlet-Multinomials is defined. In this case, in order to assure identifiability, we assume $\alpha = 0$.

The approach can be generalized to multilayer of latent variables, where at each layer perturbation parameters to the final β are introduced, under the constraint that their values are limited between -1 and 1. However, we will show in the next sections that the structure with just one additional latent layer is generally sufficient to gain large flexibility and very good clustering performance.

4 Model estimation

In this section, we present a Bayesian algorithm for parameter estimation. The prior distribution for the weights of the mixture components is assumed to follow a Dirichlet distribution with hyperparameter δ . We want non-informative priors for the model parameters. Hence, the value of the Dirichlet hyperparameter is $\delta = 1$ in order to have a flat Dirichlet distribution. The prior distributions for each α_j and β_i are given by the Uniform in the interval [-1,1] and in (0, 1000], respectively.

By using the previous model assumptions, the posterior distribution can be expressed as

$$\begin{aligned} p(z^{(1)}, z^{(2)}, \pi^{(1)}, \pi^{(2)}, \alpha, \beta|\mathbf{X}) & \propto \\ & p\left(\mathbf{X}|z^{(1)}, z^{(2)}, \alpha, \beta\right) p\left(z^{(1)}|z^{(2)}, \pi^{(1)}\right) p\left(z^{(2)}|\pi^{(2)}\right) p \\ & (\alpha) p(\beta) p\left(\pi^{(1)}\right) p\left(\pi^{(2)}\right), \end{aligned}$$

where $p(\mathbf{X}|z^{(1)}, z^{(2)}, \alpha, \beta)$ is the likelihood function of the model. By indexing the documents for which it holds that $z_{di}^{(1)} \cdot z_{dj}^{(2)} = 1$, as $d : z_{di}^{(1)} \cdot z_{dj}^{(2)} = 1$, the likelihood function can be expressed as:

$$\begin{aligned} p\left(\mathbf{X}|z^{(1)}, z^{(2)}, \alpha, \beta\right) & = \prod_{i=1}^{k_1} \prod_{j=1}^{k_2} \prod_{d: z_{di}^{(1)} \cdot z_{dj}^{(2)} = 1} \frac{\sum_{t=1}^T x_{dt}}{\prod_{t=1}^T x_{dt}} \psi_d(\alpha_j, \beta_i), \end{aligned} \tag{9}$$

with

$$\psi_d(\alpha_j, \beta_i) = \frac{B\left(\sum_{t=1}^T x_{dt}, \sum_{t=1}^T \beta_{it}(1 + \alpha_{jt})\right)}{\prod_{t=1}^T B(x_{dt}, \beta_{it}(1 + \alpha_{jt}))}$$

In order to sample parameters and latent variables from the posterior distribution, we determine the full conditionals of each unobservable variable given the other ones.

4.1 Full conditionals

The posterior distribution of the parameters and latent allocation variables given the other variables are proportional to known quantities. By using $|\dots$ to denote conditioning on all other variables, they are:

$$p(\pi^{(2)}|\dots) \sim \text{Dirichlet}\left(\sum_{d=1}^n z_{d1}^{(2)} + \delta, \dots, \sum_{d=1}^n z_{dj}^{(2)} + \delta, \dots, \sum_{d=1}^n z_{dk_2}^{(2)} + \delta\right);$$

$$p(\pi_{i|j}^{(1)}|\dots) \sim \text{Dirichlet}\left(\sum_{d=1}^n z_{d1}^{(1)} z_{dj}^{(2)} + \delta, \dots, \sum_{d=1}^n z_{di}^{(1)} z_{dj}^{(2)} + \delta, \dots, \sum_{d=1}^n z_{dk_1}^{(1)} z_{dj}^{(2)} + \delta\right);$$

$$Pr(z_{dj}^{(2)} = 1|\dots) \propto \prod_{i=1}^{k_1} \left(\pi_{i|j}^{(1)} \pi_j^{(2)} \psi_d(\alpha_j, \beta_i)\right)^{z_{di}^{(1)}};$$

$$Pr(z_{di}^{(1)} = 1|\dots) \propto \prod_{j=1}^{k_2} \left(\pi_{i|j}^{(1)} \psi_d(\alpha_j, \beta_i)\right)^{z_{dj}^{(2)}};$$

$$f(\alpha_j|\dots) \propto \prod_{i=1}^{k_1} \prod_{d: z_{di}^{(1)} \cdot z_{dj}^{(2)} = 1} \psi_d(\alpha_j, \beta_i) \mathbb{1}_{[-1 < \alpha_j < 1]};$$

$$f(\beta_i|\dots) \propto \prod_{j=1}^{k_2} \prod_{d: z_{di}^{(1)} \cdot z_{dj}^{(2)} = 1} \psi_d(\alpha_j, \beta_i) \mathbb{1}_{[0 < \beta_i < \infty]}.$$

A Gibbs sampling MCMC algorithm can be thus easily implemented for generating values from the posterior distributions.

Note that in order to get α_j and β_i we need an accept-reject mechanism. We consider as proposal value for β_i , the average value of the parameters generated by $n \times k_2$ Dirichlet-multinomials, given α_j fixed and vice versa.

The computational time of the algorithm depends on the desired number of runs of the MCMC algorithm, the number of nodes k_1 and k_2 and on the length of the vectors α_j and β_i . For a two-class dataset, described in Sect. 6, with $D = 240$, $T = 357$ and $k_2 = 2$ the MCMC algorithm with 5.000

iterations requires about 10 minutes on a processor Intel(R) Core(TM) i7-6500U CPU @ 2.50GHz, 2592 Mhz, 2 cores, under R cran 3.6.1.

We will show the estimation and clustering performance through a simulation study and real applications in the next sections.

5 Simulation study

The performance of the proposed method is evaluated under different aspects in an empirical simulation study. In order to prove the capability of the deep MoU to uncover the clusters in complex data, data were generated with a high level of sparsity. Several simulation studies are presented and discussed in the following.

The first simulation study aimed to check the capability of the deep MoU to cluster well the data, when these are generated according to a deep generative process. More precisely, we set $T = 200$ and $n = 200$, $k_1 = 3$, $k_2 = 2$ and balanced classes. We randomly generated β from a Uniform distribution in $(0,20]$ and α from a Uniform distribution in $[-1,1]$. In order to assure a high level of sparsity, for each document the total number of terms has been generated according to a Poisson distribution with parameter $N_d = 20, \forall d$. Data are then organized in a document-term matrix containing the term frequencies of each pseudo-document.

Panel (a) of Fig. 3 shows the row frequency distributions of the features across the clusters and provides a representation of the group overlapping.

Clustering performance has been measured by using both the Adjusted Rand Index (ARI) and the accuracy rate. The former is a corrected-for-chance version of the Rand Index (Hubert and Arabie 1985), that measures the degree to which two partition of objects agree; Romano et al. (2016) proved that this measure is particularly indicated in presence of large equal sized reference clusters. The accuracy is defined as the complement of the misclassification error rate. Table 1 shows the Adjusted Rand Index and the accuracy obtained by a deep MoU model for different values of k_2 , ranging from 1 to 5. We run an MCMC chain with 5000 iterations, discarding the first 2000 as burn-in. Visual inspection assured that this burn-in was largely sufficient.

The results show how the model with $k_2 > 1$ is really effective in clustering the data and, as desirable, the model with $k_2 = 2$ (i.e. the setting that reflects the generative process of the data) resulted to be the best one in terms of recovering the ‘true’ grouping structure. The gap between $k_2 = 1$ and $k_2 > 1$ is relevant; however, the performance remains elevate for the various $k_2 > 1$, thus indicating that a deep structure can be really effective in clustering such kind of data.

Fig. 3 Heatmap of the conditional distribution of features across the classes. The left panel refers to the first simulation study and the right panel to the second ones

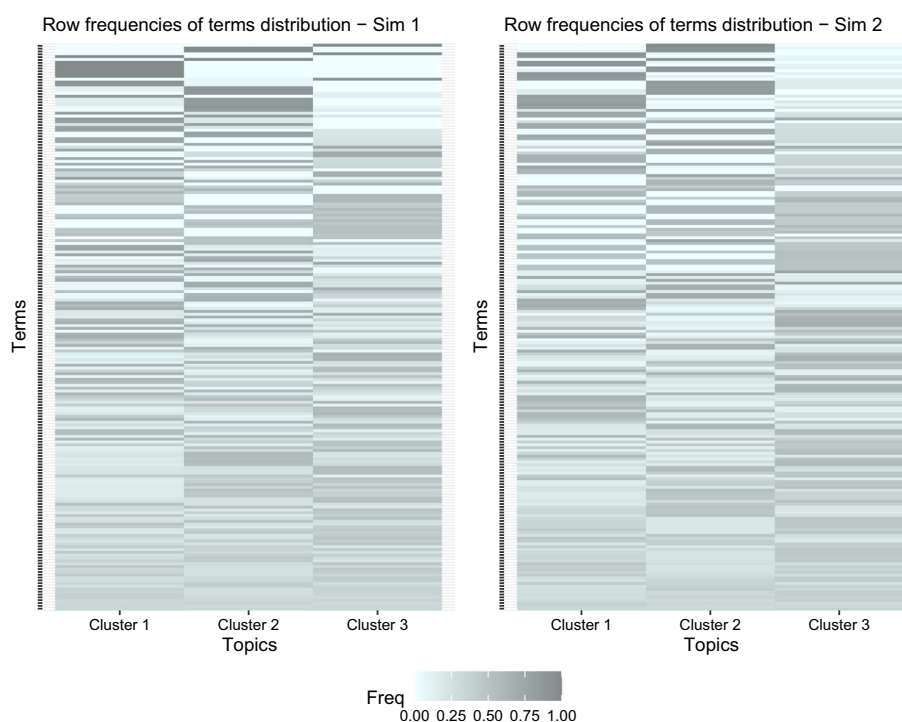


Table 1 Simulated Data, simulation study 1. Adjusted Rand Index (ARI) and Accuracy of deep MoU for $k_1 = 3$ and different values of k_2

k_2	ARI	Accuracy
1	0.468	0.755
2	0.940	0.980
3	0.910	0.970
4	0.934	0.980
5	0.925	0.975

Table 2 Simulated Data, simulation study 2. Adjusted Rand Index (ARI) and Accuracy of deep MoU for $k_1 = 3$ and different values of k_2

k_2	ARI	Accuracy
1	0.824	0.940
2	0.811	0.935
3	0.798	0.930
4	0.811	0.935
5	0.797	0.930

In a second simulation study, we tested the performance of a deep MoU with data that are not originated according to a deep generative process, but simply by $k_1 = 3$ balanced groups, $T = n = 200$ and $N_d = 20, \forall d$. Panel (b) of Fig. 3 displays the row frequency distributions of the generated features. As shown in Table 2, in this situation, the deep model with $k_2 > 1$ does not significantly improve the clustering per-

formance, and the accuracy remains stable as k_2 increases. This suggests that when the data are pretty simple, and are not high-dimensional, a deep algorithm is not more efficient than the conventional MoU.

The third simulation study aimed at measuring the accuracy of the estimated parameters α and β in data with double structure $k_1 = 3$ and $k_2 = 2$, allowing for different combinations of T , n and N , so as to measure the effect of data dimensionality and level of sparsity on the goodness of fit. We considered a total of 8 different scenarios generated according to the combinations of $T = \{100, 200\}$, $n = \{100, 200\}$ and $N = \{10, 20\}$. Table 3 contains the Euclidean distance between the true parameter vectors and the posterior means, normalized over T .

As expectable, for a given T the goodness of fit improves as the number of documents increases. The level of sparsity has a relevant role as well: when N increases the documents are more informative and the parameter estimates become more accurate.

5.1 Application to real data

The effectiveness of the proposed model is demonstrated by using four textual datasets, including the introduced ticket data. We compare the deep MoU with conventional clustering strategies: k -means (Lloyd 1982) with cosine distance (k -means) and with Euclidean distance on data transformed according to semantic analysis (LSA k -means), partition around medoids (PAM) (Kaufman and Rousseeuw 2009),

Table 3 Simulated Data, simulation study 3. Precision of the estimated parameters for different values of n , T and N

n	T	N	$\frac{\ \beta_1 - \hat{\beta}_1\ }{T}$	$\frac{\ \beta_2 - \hat{\beta}_2\ }{T}$	$\frac{\ \beta_3 - \hat{\beta}_3\ }{T}$	$\frac{\ \alpha_1 - \hat{\alpha}_1\ }{T}$	$\frac{\ \alpha_2 - \hat{\alpha}_2\ }{T}$
100	100	10	5.016	5.357	4.737	0.488	0.596
100	200	10	5.360	5.754	5.455	0.511	0.516
100	100	20	3.983	3.917	3.166	0.340	0.329
100	200	20	5.396	4.261	4.220	0.375	0.392
200	100	10	3.999	4.188	5.061	0.317	0.370
200	200	10	4.630	4.540	5.934	0.439	0.441
200	100	20	3.275	5.488	3.205	0.261	0.291
200	200	20	3.912	3.435	3.982	0.311	0.315

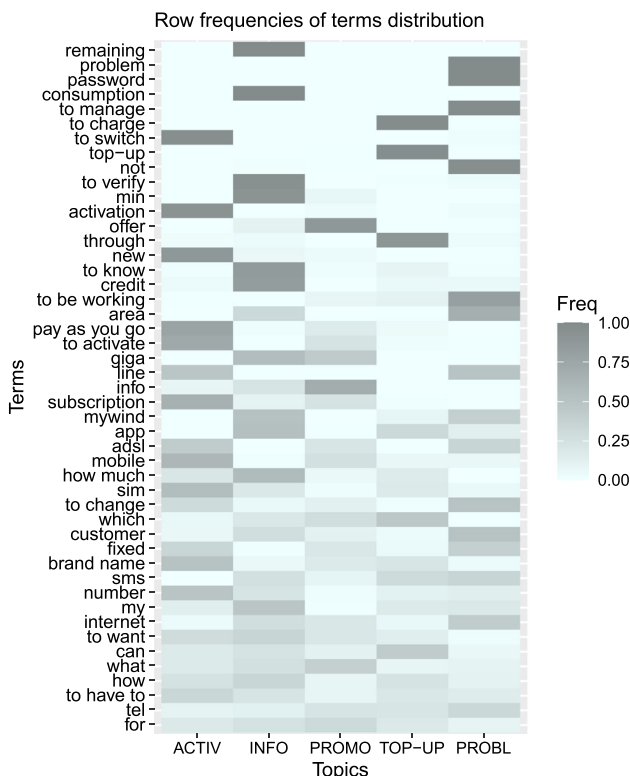


Fig. 4 Heatmap of the conditional distribution of terms across classes. The shades reflect the distribution of row frequencies

mixture of Gaussians on semantic-based transformed data (MoG) (McLachlan and Peel 2000), hierarchical clustering according to Ward’s method (Murtagh and Legendre 2014) with cosine dissimilarity (HC), latent Dirichlet allocation (LDA) (Blei et al. 2003), mixtures of unigrams (MoU) (Nigam et al. 2000) estimated via the EM algorithm (Dempster et al. 1977), spectral clustering (SpeCl) (Ng et al. 2002) and affinity propagation clustering (AffPr) (Frey and Dueck 2007) with normalized linear kernel.

The CNAE-9 dataset contains 1080 documents of free text business descriptions of Brazilian companies categorized into 9 balanced categories (Ciarelli and Oliveira 2009; Cia-

Table 4 Real Data. Number of tickets for each class

Class	Description	Freq.
ACTIV	Activation of SIM, ADSL, new contracts	407
INFO	General information about current balance, consumption, etc.	471
PROMO	Request of information about new offers and promotions	376
TOP-UP	Top-up	435
PROBL	Problems with password, top-up, internet connection, etc.	440

relli et al. 2010) for a total of 856 preprocessed words. Since the classes 4 and 9 are the most overlapped we considered also the reduced CNAE-2 dataset composed by these two groups only which consists of 240 documents and 357 words. This dataset is highly sparse: the 99.22% of the document-term matrix entries are zeros.

The BBC dataset consists of 737 documents from the BBC Sport website corresponding to sport news articles in five topics/areas (athletics, cricket, football, rugby, tennis) from 2004 to 2005 (Greene and Cunningham 2006). After a preprocessing phase aimed at discarding non-relevant words, the total number of features is 1075. This dataset is moderately sparse with a fraction of zeros equal to 92.36%.

The ticket dataset contains $n = 2129$ tickets and $T = 489$ terms obtained after preprocessing: original raw data were processed via stemming, so as to reduce inflected or derived words to their unique word stem, and some terms have been filtered out in order to remove very common non-informative *stopwords* words in the Italian language. The tickets have then been classified by independent operators to $k = 5$ main classes described in Table 4.

The peculiarity and major challenge of this dataset is the limited number of words used, on average, for each ticket. In

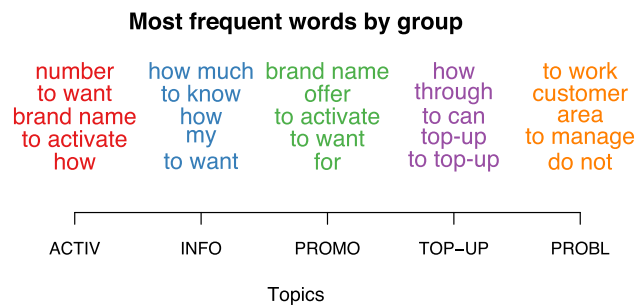


Fig. 5 Most frequent terms within each group

Table 5 Real Data. Accuracy (multiplied by 100) for different methods

Method	CNAE-2	CNAE-9	BBC	Ticket
<i>k</i> -means	55.8	59.8	76.5	51.6
LSA <i>k</i> -means	65.0	54.4	40.6	46.4
PAM	52.5	11.9	28.8	22.2
MoG	50.4	37.6	43.6	43.8
HC	51.2	13.2	25.4	22.5
LDA	53.3	64.6	41.0	31.8
MoU	60.8	69.6	59.0	74.3
SpeCl	72.9	50.3	44.5	53.0
AffPr	59.2	58.9	53.5	51.7

fact, after preprocessing, the tickets have an average length of 5 words. A graphical representation of the data, after the preprocessing step, is shown in Fig. 4. The heatmap shows the row frequency distributions of the most common terms (overall occurrences ≥ 50). As clear from the shades, each class is characterized by a limited set of specific words, while the majority is uniformly distributed across them, thus making the classification problem particularly challenging. In order to have further details about the terms characterizing the tickets within each topic, Fig. 5 displays the most frequent words.

Data are characterized by a large amount of sparsity with the 99.05% of zeros: as a consequence, most conventional clustering strategies fail.

Tables 5 and 6 show, respectively, the accuracy and the ARI of the different methods on the presented datasets. For comparative reasons, the true number of clusters was considered as known for all the methods.

Among the considered methods, spectral clustering, *k*-means and mixtures of unigrams are the most effective method for classifying the datasets. An advantage of MoU, which is inherited by the proposed deep MoU, is that it seems to be not affected by the large number of zeros like the other methods, as it is based on proportions. The latent Dirichlet allocation model, despite its flexibility, is not able to improve the classification on these short documents, because it is based on the assumption of multiple topics for each docu-

Table 6 Real Data. Adjusted Rand Index (ARI) for different methods

Method	CNAE-2	CNAE-9	BBC	Ticket
<i>k</i> -means	0.010	0.418	0.567	0.233
LSA <i>k</i> -means	0.088	0.282	0.038	0.135
PAM	0.002	0.000	0.031	0.000
MoG	0.000	0.092	0.060	0.158
HC	-0.003	-0.005	-0.004	-0.001
LDA	0.000	0.507	0.092	0.000
MoU	0.043	0.562	0.226	0.523
SpeCl	0.207	0.284	0.104	0.224
AffPr	0.030	0.338	0.189	0.219

Table 7 Real Data. Accuracy (multiplied by 100) of deep MoU for different numbers of nodes k_2

k_2	CNAE-2	CNAE-9	BBC	Ticket
1	81.2	50.1	95.5	76.2
2	93.3	73.3	95.8	89.2
3	92.9	77.5	95.9	87.6
4	92.5	77.3	93.9	85.6
5	91.2	76.4	93.9	83.3

Table 8 Real Data. Adjusted Rand Index (ARI) of deep MoU for different numbers of nodes k_2

k_2	CNAE-2	CNAE-9	BBC	Ticket
1	0.388	0.290	0.888	0.616
2	0.750	0.667	0.894	0.757
3	0.736	0.661	0.895	0.721
4	0.721	0.663	0.848	0.681
5	0.679	0.615	0.843	0.651

ment, which is not realistic for short data, as in the case of ticket dataset.

We applied deep MoU with $k_2 = 1, \dots, 5$. For each setting, we run 5000 iterations of the MCMC algorithm, discarding the first 2000 as burn-in. From graphical visualization and diagnostic criteria we observed convergence and stability to the different choices of starting points and hyperparameters δ , so we considered $\delta = 1$.

Tables 7 and 8 contain the clustering results of the deep MoU on the datasets, measured by accuracy and ARI, respectively. The case $k_2 = 1$ corresponds to a Bayesian MoU and classification is improved with respect to the conventional MoU in most all empirical cases. A probable reason for that is the adoption of Dirichlet-multinomials (in Bayesian MoU) instead of classical multinomials (in frequentist MoU), which are more suitable to capture overdispersion with respect to the multinomial framework (Wilson and Koehler 1991). In our analysis, the only exception is represented by the dataset

CNAE-9, which is particularly challenging because it is composed by 9 unbalanced groups. In this case, classical MoU performs slightly better in terms of misclassification rate and ARI. This may depend on a particularly good starting point of the EM algorithm for fitting MoU, which is initialized by the k -means clustering.

When we move from $k_2 = 1$ to a deeper structure with $k_2 = 2$ or $k_2 = 3$ the accuracy improves in all the analyzed experiments. This proves how the introduction of the parameters α_j may be beneficial for classification. When the hidden nodes are greater than 3 ($k_2 = 4$ or $k_2 = 5$) results seem to be little worsened, probably because of the larger number of parameters to be estimated.

6 Final remarks

In this paper, we have proposed a deep learning strategy that extends the mixtures of unigrams model. With respect to other clustering methods, MoUs have desirable features for textual data. Firstly, document-term matrices usually contain the frequencies of the words in each document; for this reason, MoUs represent an intuitive choice, since they are based on multinomial distributions that are the probabilistic distributions for modeling positive frequencies. Moreover, as MoUs naturally model proportions, they are not affected by the large amount of zeros of the datasets like other methods, so they are a proper choice for modeling very short texts and sparse datasets. Furthermore, MoU is based on the idea that documents related to the same topic have similar distributions of terms, which is realistic in practice. Taking a mixture of k multinomials means doing clustering into k topics/groups: there is a unique association between documents and topics.

All these nice characteristic are inherited by the proposed deep MoU model. The proposed deep MoU is particularly effective in clustering with challenging issues (sparsity, overdispersion, short document length and high-dimensionality). Being hierarchical in its nature, the model can be easily estimated by a MCMC algorithm in a Bayesian framework. In our analysis, we chose non-informative priors, because there is little prior information available on the empirical context. The estimation algorithm produces good results in all the simulated and real situations considered here.

The proposed model could be extended in several directions: as discussed in Sect. 3, several hidden layers (instead of just a single one) could be considered. The merging function $\beta_i(1 + \alpha_j)$ has been defined for identifiability reasons under the idea that the number of estimated groups is k_1 and the latent layer is only aimed at perturbing the β_i parameters for capturing some residual heterogeneity inside the groups. Of course, more complex nonlinear functions could be considered, without however losing sight of identifiability. In case of non-extreme sparsity and long documents, the model

could be also extended to allow for deep m -gram models. We leave all these ideas to future research.

Funding Open Access funding provided by Alma Mater Studiorum - Università di Bologna

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

References

- Baek, J., McLachlan, G., Flack, L.: Mixtures of factor analyzers with common factor loadings: applications to the clustering and visualization of high-dimensional data. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(7), 1298–1309 (2010)
- Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022 (2003)
- Ciarelli, P.M., Oliveira, E.: Agglomeration and elimination of terms for dimensionality reduction. In: Proceedings of the 2009 Ninth International Conference on Intelligent Systems Design and Applications, ISDA '09, pp. 547–552. IEEE Computer Society, Washington, DC, USA (2009)
- Ciarelli, P.M., Salles, E.O.T., Oliveira, E.: An evolving system based on probabilistic neural network. In: 2010 Eleventh Brazilian Symposium on Neural Networks, pp. 182–187 (2010)
- Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc. Ser. B Methodol.* **39**(1), 1–38 (1977)
- Fraley, C., Raftery, A.: Model-based clustering, discriminant analysis and density estimation. *J. Am. Stat. Assoc.* **97**, 611–631 (2002)
- Frey, B.J., Dueck, D.: Clustering by passing messages between data points. *Science* **315**(5814), 972–976 (2007)
- Greene, D., Cunningham, P.: Practical solutions to the problem of diagonal dominance in kernel document clustering. In: Proceedings of 23rd International Conference on Machine Learning (ICML'06), pp. 377–384. ACM Press (2006)
- Griffiths, T.L., Steyvers, M.: Finding scientific topics. *Proceed. Natl. Acad. Sci.* **101**(suppl 1), 5228–5235 (2004). <https://doi.org/10.1073/pnas.0307752101>
- Hubert, L., Arabie, P.: Comparing partitions. *J. Classification* **2**(1), 193–218 (1985)
- Kaufman, L., Rousseeuw, P.J.: Finding Groups in Data: An Introduction to Cluster Analysis, vol. 344. Wiley, Hoboken (2009)
- LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
- Li, J.: Clustering based on a multilayer mixture model. *J. Comput. Graph. Stat.* **14**(3), 547–568 (2005). <https://doi.org/10.1198/106186005X59586>
- Lloyd, S.: Least squares quantization in pcm. *IEEE Trans. Inf. Theory* **28**(2), 129–137 (1982)

- McLachlan, G., Peel, D., Bean, R.: Modelling high-dimensional data by mixtures of factor analyzers. *Comput. Stat. Data Anal.* **41**(3), 379–388 (2003)
- McLachlan, G.J., Peel, D.: *Finite Mixture Models*. Wiley, Hoboken (2000)
- Montanari, A., Viroli, C.: Heteroscedastic factor mixture analysis. *Stat. Model.* **10**(4), 441–460 (2010)
- Murtagh, F., Legendre, P.: Ward’s hierarchical agglomerative clustering method: which algorithms implement ward’s criterion? *J. Classification* **31**(3), 274–295 (2014)
- Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: analysis and an algorithm. In: *Advances in Neural Information Processing Systems*, pp. 849–856 (2002)
- Nigam, K., McCallum, A., Thrun, S., Mitchell, T.: Text classification from labeled and unlabeled documents using EM. *Mach. Learn.* **39**, 103–134 (2000)
- van den Oord, A., Schrauwen, B.: Factoring variations in natural images with deep Gaussian mixture models. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 27*, pp. 3518–3526. Curran Associates Inc, New York (2014)
- Peng, X., Feng, J., Xiao, S., Yau, W.Y., Zhou, J.T., Yang, S.: Structured autoencoders for subspace clustering. *IEEE Trans. Image Process.* **27**(10), 5076–5086 (2018)
- Peng, X., Yu, Z., Yi, Z., Tang, H.: Constructing the l2-graph for robust subspace learning and subspace clustering. *IEEE Trans. Cybern.* **47**(4), 1053–1066 (2016)
- Romano, S., Vinh, N.X., Bailey, J., Verspoor, K.: Adjusting for chance clustering comparison measures. *J. Mach. Learn. Res.* **17**(1), 4635–4666 (2016)
- Schmidhuber, J.: Deep learning in neural networks: an overview. *Neural Netw.* **61**, 85–117 (2015)
- Tang, Y., Hinton, G.E., Salakhutdinov, R.: Deep mixtures of factor analyzers. In: J. Langford, J. Pineau (eds.) *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 505–512. ACM, New York (2012)
- Tomović, A., Janičić, P., Kešeljić, V.: n-gram-based classification and unsupervised hierarchical clustering of genome sequences. *Comput. Methods Prog. Biomed.* **81**(2), 137–153 (2006)
- Viroli, C.: Dimensionally reduced model-based clustering through mixtures of factor mixture analyzers. *J. Classification* **27**(3), 363–388 (2010)
- Viroli, C., McLachlan, G.J.: Deep Gaussian mixture models. *Stat. Comput.* **29**(1), 43–51 (2019)
- Wilson, J., Koehler, K.: Hierarchical models for cross-classified overdispersed multinomial data. *J. Business Econ. Stat.* **9**(1), 103–110 (1991)
- Zhou, J.T., Zhao, H., Peng, X., Fang, M., Qin, Z., Goh, R.S.M.: Transfer hashing: from shallow to deep. *IEEE Trans. Neural Netw. Learn. Syst.* **29**(12), 6191–6201 (2018)

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