



Research Repository UCD

Title	Exponential family mixed membership models for soft clustering of multivariate data
Authors(s)	White, Arthur, Murphy, Thomas Brendan
Publication date	2016-12
Publication information	White, Arthur, and Thomas Brendan Murphy. "Exponential Family Mixed Membership Models for Soft Clustering of Multivariate Data." Springer, December 2016. https://doi.org/10.1007/s11634-016-0267-5 .
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/7981
Publisher's statement	The final publication is available at Springer via http://dx.doi.org/10.1007/s11634-016-0267-5 .
Publisher's version (DOI)	10.1007/s11634-016-0267-5

Downloaded 2025-12-04 22:57:32

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Exponential Family Mixed Membership Models for Soft Clustering of Multivariate Data

Arthur White · Thomas Brendan Murphy

Abstract For several years, model-based clustering methods have successfully tackled many of the challenges presented by data-analysts. However, as the scope of data analysis has evolved, some problems may be beyond the standard mixture model framework. One such problem is when observations in a dataset come from overlapping clusters, whereby different clusters will possess similar parameters for multiple variables. In this setting, mixed membership models, a soft clustering approach whereby observations are not restricted to single cluster membership, have proved to be an effective tool. In this paper, a method for fitting mixed membership models to data generated by a member of an exponential family is outlined. The method is applied to count data obtained from an ultra running competition, and compared with a standard mixture model approach.

1 Introduction

The field of model-based clustering (MBC) (Fraley and Raftery, 2002; McLachlan and Peel, 2002) has successfully tackled many of the challenges presented by data-analysts. Within this framework, observations in a dataset are modelled as being drawn from one of several probability distributions. One of the central tenets of MBC, as stated by Fraley and Raftery (2002), is that datapoints may then be classified so that “each component probability distribution corresponds to a cluster.” While more recent developments, such as those by Baudry et al. (2010) have evolved this definition somewhat, fundamentally within this framework a clustering solution is sought whereby observations are partitioned into distinct groups, so that observations which have non-negligible posterior probability of belonging to more than one component are seen as having uncertain group membership, and are perhaps indicative of a poorly fitted model.

However, there are several instances where such a model may prove too restrictive, and it is convenient to introduce a soft clustering approach so that individual observations are modelled by a mixture of components. Examples include: topic

modelling, where documents are often interpreted as covering a combination of topics (Blei et al., 2003; Erosheva et al., 2004); micro cDNA arrays, where overlapping genetic characteristics can be exhibited (Rogers et al., 2005); functional disability surveys, where symptoms may be shared (Erosheva et al., 2007) and elections with preferential voting systems, where voters' political positions can be viewed as some combination of multiple types (Gormley and Murphy, 2009)¹. In each of these examples, the cited authors use mixed membership models to analyse the data. Within this framework, observations may be modelled as possessing multiple attributes from the different component probability distributions which are assumed to form the latent structure of the data. Thus, an observation may possess high posterior membership to two or more components with a high degree of certainty.

The general case of mixed membership models, where quite general component distributions were allowed, has been outlined by Erosheva et al. (2004), however, details of how inference is to be performed are omitted; a variational Bayes approximation is recommended, but not described. Other studies (Blei et al., 2003; Erosheva et al., 2004; Rogers et al., 2005; Gormley and Murphy, 2009) outline a mixed membership approach directly for the problem at hand, and propose to perform inference via either variational Bayes methods (Blei et al., 2003; Erosheva et al., 2007; Rogers et al., 2005) and/or MCMC schemes (Erosheva et al., 2007; Gormley and Murphy, 2009). Airolidi et al. (2006, 2007) discuss mixed membership models with an emphasis on the issue of model selection. See Airolidi et al. (2014) for a detailed overview of the historical development of mixed membership models and the main areas in which they have been applied. In this paper, the mixed membership approach and a variational Bayes method for inference are outlined for the case where component distributions are members of an exponential family.

Examples of the method are applied to count data, where the corresponding component distribution is chosen to be Poisson, are provided. The method is first applied to data obtained from a 24 hour ultra running competition, where the hourly number of laps completed by each competitor has been recorded. A comparison is then made to a mixture model approach consistent with standard MBC practices.

The rest of the paper is detailed as follows: Section 2 outlines the general model specification for a mixed membership framework for members of an exponential family. Parameter estimation and model selection, as well as some model evaluation tools and a brief overview of the mixture framework is then discussed in Section 3. The running data is introduced in Section 4, with mixture and mixed membership models fitted to the data and compared. Possible extensions to the model are then discussed in Section 5.

¹ Note that these examples use different terminology to describe their methods: latent Dirichlet allocation (Blei et al., 2003), latent process decomposition (Rogers et al., 2005) and grade of membership (Erosheva et al., 2007; Gormley and Murphy, 2009). Each of the models allocate individual observations to multiple components in a similar fashion, which we refer in general to as a mixed membership model (Erosheva et al., 2004).

2 Model Specification

We describe the mixed membership framework. Let $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_N)$ denote our dataset, consisting of N observations of M attributes. We assume that some number G of *basis profiles* underwrite the data. We use this term to distinguish from terms such as group or cluster, that are commonly used with respect to mixture models. Rather than treating each observation as belonging to a distinct cluster, observations are considered to be some composition of these profiles.

Weight (or mixed membership) parameters $\boldsymbol{\tau} = (\boldsymbol{\tau}_1, \dots, \boldsymbol{\tau}_N)$ are assigned to observations \mathbf{X} , so that for each $\mathbf{X}_n = (X_{n1}, \dots, X_{nM})$, $\boldsymbol{\tau}_n = (\tau_{n1}, \dots, \tau_{nG})$. Each τ_{ng} can be interpreted as the probability that an observation will have membership to profile g for an attribute m , so each $\tau_{ng} > 0$, $\sum_{g=1}^G \tau_{ng} = 1$. Thus, for a given observation, the *a priori* probability of profile membership is the same for each attribute. Each $\boldsymbol{\tau}_n$ is assumed to follow a Dirichlet distribution, with common hyperparameter $\boldsymbol{\delta} = (\delta_1, \dots, \delta_G)$.

Profile memberships by attribute for \mathbf{X} are denoted by the indicator variable $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_N)$, where $\mathbf{Z}_n = (\mathbf{Z}_{n1}, \dots, \mathbf{Z}_{nM})$. Specifically, profile membership for each X_{nm} is denoted by the indicator variable $\mathbf{Z}_{nm} = (Z_{nm1}, \dots, Z_{nmG})$, where:

$$Z_{nmg} = \begin{cases} 1 & \text{if observation } n \text{ is member of profile } g \text{ for attribute } m; \\ 0 & \text{otherwise.} \end{cases}$$

Each \mathbf{Z}_{nm} is modelled as a multinomial distribution, depending on the probability $\boldsymbol{\tau}_n$.

Lastly, we use $\boldsymbol{\theta}^\top = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$, to denote the distribution of data conditional on profile membership, $\boldsymbol{\theta}_m = (\boldsymbol{\theta}_{1m}, \dots, \boldsymbol{\theta}_{Gm})$. For membership to profile g for attribute m , $\boldsymbol{\theta}_{gm}$ denotes the underlying parameter(s) of a distribution density $p_1(x_{nm} \mid \boldsymbol{\theta}_{gm})$. We restrict $p_1(x_{nm} \mid \boldsymbol{\theta}_{gm})$ to be a member of an exponential family of distributions:

$$p_1(x_{nm} \mid \boldsymbol{\theta}_{gm}) = h(x_{nm})k(\boldsymbol{\theta}_{gm}) \exp \left\{ \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \mathbf{s}(x_{nm}) \right\},$$

where $\mathbf{r}(\boldsymbol{\theta}_{gm})$ is the natural vector of parameters for $\boldsymbol{\theta}_{gm}$, $\mathbf{s}(x_{nm})$ the sufficient statistic for x_{nm} , and $h(x_{nm})$ is a normalising constant. Note that the dimensions of $\boldsymbol{\theta}_{gm}$, $\mathbf{s}(x_{nm})$, and $\mathbf{r}(\boldsymbol{\theta}_{gm})$ depend on the distribution in question.

The generative process for \mathbf{X} is thus assumed to be given by the following steps:

- for each $n \in 1, \dots, N$: $\boldsymbol{\tau}_n \sim \text{Dirichlet}(\boldsymbol{\delta})$
- for each $m \in 1, \dots, M$: $\mathbf{Z}_{nm} \sim \text{Multinomial}(1, \boldsymbol{\tau}_n)$
- $X_{nm} \mid Z_{nmg} = 1 \sim p_1(x_{nm} \mid \boldsymbol{\theta}_{gm})$.

In the special case where profile distributions are $\text{Multinomial}(1, \boldsymbol{\theta}_{gm})$, for all g, m , then at an individual level observations will also follow a multinomial distribution, with parameters that are a convex combination of the profile parameters (Galyardt, 2014). In the more general case, individuals should be interpreted as switching between profiles across attributes.

The complete-data posterior for a mixed membership model takes the form:

$$p(\boldsymbol{\tau}, \boldsymbol{\theta}, \mathbf{Z} | \mathbf{x}, \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) \propto p_2(\mathbf{x} | \mathbf{Z}, \boldsymbol{\theta}) p_3(\mathbf{Z} | \boldsymbol{\tau}) p_4(\boldsymbol{\tau} | \boldsymbol{\delta}) \prod_{g=1}^G \prod_{m=1}^M p_5(\boldsymbol{\theta} | \eta_{gm}, \boldsymbol{\nu}_{gm}), \quad (1)$$

where

$$\begin{aligned} p_2(\mathbf{x} | \mathbf{Z}, \boldsymbol{\theta}) &= \prod_{n=1}^N \prod_{m=1}^M \prod_{g=1}^G p_1(x_{nm} | \boldsymbol{\theta}_{gm})^{Z_{nmg}} \\ p_3(\mathbf{Z} | \boldsymbol{\tau}) &= \prod_{n=1}^N \prod_{g=1}^G \tau_{ng}^{\sum_{m=1}^M Z_{nmg}} \\ p_4(\boldsymbol{\tau} | \boldsymbol{\delta}) &= \prod_{n=1}^N \frac{\Gamma(\sum_{h=1}^G \delta_h)}{\prod_{h=1}^G \Gamma(\delta_h)} \prod_{g=1}^G \tau_{ng}^{\delta_g - 1} \\ p_5(\boldsymbol{\theta}_{gm} | \eta_{gm}, \boldsymbol{\nu}_{gm}) &= h(\eta_{gm}, \boldsymbol{\nu}_{gm}) k(\boldsymbol{\theta}_{gm})^{\eta_{gm}} \exp\{\mathbf{r}(\boldsymbol{\theta}_{gm})^\top \boldsymbol{\nu}_{gm}\}. \end{aligned}$$

where we have assumed conjugate priors for $p_1(\mathbf{x} | \boldsymbol{\theta})$ and $p_3(\mathbf{Z} | \boldsymbol{\tau})$.

Note that the form of the posterior outlined in Equation (1) makes an implicit assumption of the exchangeability of each latent variable \mathbf{Z}_n (see Section 3.1, Blei et al., 2003). That is, the likelihood of the model will be unchanged for any permutation of the variable index $m = 1, \dots, M$. Thus, for any observation \mathbf{X}_n , all of the observed variables (X_{n1}, \dots, X_{nM}) are assumed to be independent, conditional on their respective profile memberships $(\mathbf{Z}_{n1}, \dots, \mathbf{Z}_{nM})$. The use of latent variables in a data augmentation approach can also be motivated by a fundamental representation theorem; see Erosheva et al. (2007, Section 3) for further details.

A graphical depiction of Equation (1) is shown in Figure 1a. For comparison, a mixture model is shown in Figure 1b; this model is formally described in Section 3.4. We repeat notation for the models to highlight similarities in structure. The plate notation in the graph represents the dimensionality of the model parameters. In particular, the different positions of $\boldsymbol{\tau}$ and \mathbf{Z} with respect to this notation illustrate the additional complexity of the mixed membership model.

Note that only the hyperparameter for the prior $p_4(\boldsymbol{\tau} | \boldsymbol{\delta})$ was included in Figure 1a, and that the prior $p_5(\boldsymbol{\theta}_{gm} | \eta_{gm}, \boldsymbol{\nu}_{gm})$ was omitted from the outlined data generative process. This is in keeping with previous studies (Blei et al., 2003; Erosheva et al., 2007; Rogers et al., 2005) where only $\boldsymbol{\tau}$ has been considered a parameter of interest, with $\boldsymbol{\theta}$ treated as a nuisance parameter, with the prior specification for $\boldsymbol{\eta}$ and $\boldsymbol{\nu}$ set as small as possible, so that $p_5(\boldsymbol{\theta}_{gm} | \eta_{gm}, \boldsymbol{\nu}_{gm})$ is as close to a uniform distribution as possible. In either case, calculation of the normalization constant in (1) is intractable (Blei et al., 2003). For completeness, we consider both cases when discussing the inference method for the model, however, when applying the method to data we choose the nuisance parameter method. While we examine the estimated parameters $\hat{\boldsymbol{\theta}}$ in Section 4 in order to interpret the clusters, our primary interest remains the estimation of the underlying mixed membership structure. To perform inference we appeal to variational methods (Beal, 2003; Ormerod and Wand, 2010; Bishop, 2006, Chapter 10).

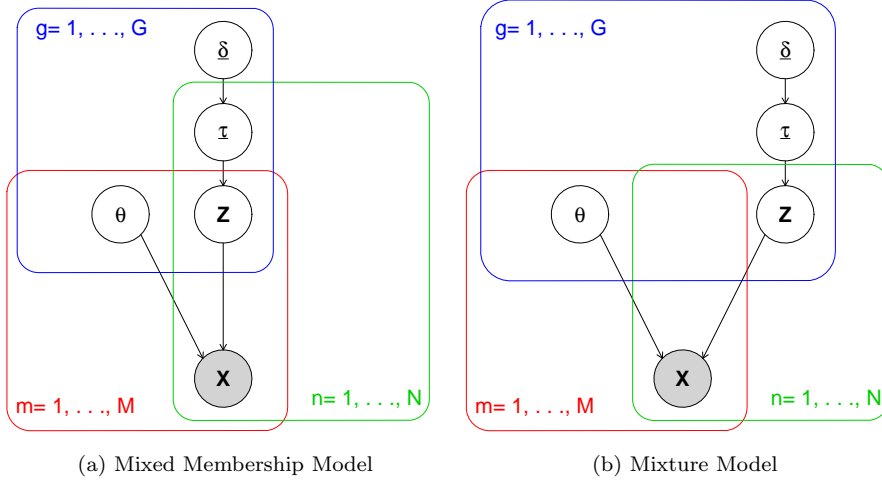


Figure 1: Graphical depiction of the mixed membership formulation (a) in comparison to the standard mixture model (b), for the case where θ is treated as a nuisance parameter. Note in particular the different positions of τ and \mathbf{Z} with respect to the plate notation.

3 Parameter Estimation

In this section parameter estimation for mixed membership exponential family models are outlined. While some of these results are the same as those found in (Blei et al., 2003) the approach as outlined here more closely follows the more general derivation provided in Bishop (2006, Chapter 10). As a running example, we illustrate how these methods are applied to data generated from a Poisson distribution, i.e., the case where

$$p_6(x_{nm} | \theta_{gm}) = \frac{\exp(-\theta_{gm}) \theta_{gm}^{x_{nm}}}{x_{nm}!}. \quad (2)$$

Then $p_6(x_{nm} | \theta_{gm})$ is a member of an exponential family with the following specifications: $h(x_{nm}) = 1/x_{nm}!$, $k(\theta_{gm}) = \exp(-\theta_{gm})$, $s(x_{nm}) = x_{nm}$, and $r(\theta_{gm}) = \log(\theta_{gm})$. A Gamma(α, β) distribution is a conjugate prior for a Poisson distribution:

$$p_7(\theta_{gm} | \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta_{gm}^{\alpha-1} \exp(-\beta \theta_{gm}).$$

Matching notation from the previous section gives $\eta = \beta$, $\nu = \alpha - 1$, and $h(\eta, \nu) = \eta^{\nu+1} / \Gamma(\nu + 1)$. The method applied in Section 4 also uses this distribution.

3.1 Variational Bayes

The posterior (1) is approximated using a variational Bayes method (Blei et al., 2003; Rogers et al., 2005; Erosheva et al., 2007) whereby the posterior $p(\tau, \theta, \mathbf{Z} | \mathbf{x}, \delta, \eta, \nu)$

is replaced by an approximating set of distributions $q(\mathbf{Z}, \tau, \theta | \phi, \gamma, \boldsymbol{\eta}', \boldsymbol{\nu}')$ that factor independently:

$$q(\mathbf{Z}, \tau, \theta | \phi, \gamma, \boldsymbol{\eta}', \boldsymbol{\nu}') = q_1(\tau | \gamma) q_2(\mathbf{Z} | \phi) q_3(\theta | \boldsymbol{\eta}', \boldsymbol{\nu}'), \quad (3)$$

where $\gamma, \phi, \boldsymbol{\eta}'$ and $\boldsymbol{\nu}'$ are free variational parameters of $q_1(\tau)$, $q_2(\mathbf{Z})$ and $q_3(\theta)$ respectively. Note that $\phi, \gamma, \boldsymbol{\eta}'$ and $\boldsymbol{\nu}'$ have the same dimensionality as $\tau, \mathbf{Z}, \boldsymbol{\eta}$ and $\boldsymbol{\nu}$ respectively.

To begin with, we obtain an upper bound to the log posterior $\log p_8(\mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})$ in terms of a posterior q with latent parameters \mathbf{Z} , θ and τ .

$$\log p_8(\mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) = \log \int \int \sum_{\mathbf{Z}} p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) d\tau d\theta \quad (4)$$

$$\begin{aligned} &= \log \int \int \sum_{\mathbf{Z}} \frac{q(\mathbf{Z}, \tau, \theta) p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})}{q(\mathbf{Z}, \tau, \theta)} d\tau d\theta \\ &\geq \int \int \sum_{\mathbf{Z}} q(\mathbf{Z}, \tau, \theta) \log p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) d\tau d\theta \\ &\quad - \int \int \sum_{\mathbf{Z}} q(\mathbf{Z}, \tau, \theta) \log q(\mathbf{Z}, \tau, \theta) d\tau d\theta, \end{aligned} \quad (5)$$

where Eq.(5) is given by Jensen's inequality. It can be shown that the difference between Eq.(5) and Eq.(4) is the Kullback-Liebler divergence $\mathcal{KL}(p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) \| q(\mathbf{Z}, \tau, \theta))$. Thus maximising Eq.(4) amounts to minimising the divergence between the true posterior p and approximate distribution density q .

Introducing the restriction that the approximate distribution density $q(\mathbf{Z}, \tau, \theta | \phi, \gamma, \boldsymbol{\eta}', \boldsymbol{\nu}')$ factors independently, it is then possible to maximise Eq. (5) with respect to $q(\tau)$:

$$\begin{aligned} &\int \int \sum_{\mathbf{Z}} q(\mathbf{Z}, \tau, \theta) \log p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) d\tau d\theta - \int \int \sum_{\mathbf{Z}} q(\mathbf{Z}, \tau, \theta) \log q(\mathbf{Z}, \tau, \theta) d\tau d\theta \\ &= \int q_1(\tau) \int q_3(\theta) \sum_{\mathbf{Z}} q_2(\mathbf{Z}) \log p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu}) d\theta d\tau - \int q_1(\tau) \log q_1(\tau) d\tau + \text{constant} \\ &= \int q_1(\tau) \{ \mathbb{E}_{\mathbf{Z}, \theta} [\log p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})] + \text{constant} - \log q_1(\tau) \} d\tau + \text{constant} \\ &= \int q_1(\tau) \log \left\{ \frac{\exp(\mathbb{E}_{\mathbf{Z}, \theta} [\log p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})] + \text{constant})}{q_1(\tau)} \right\} d\tau + \text{constant} \\ &= -\mathcal{KL}(\exp(\mathbb{E}_{\mathbf{Z}, \theta} [\log p(\tau, \theta, \mathbf{Z}, \mathbf{x} | \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})] + \text{constant}) \| q_1(\tau)) + \text{constant}. \end{aligned}$$

It can thus be shown that maximising Eq. (5) with respect to $q_1(\boldsymbol{\tau})$ is equivalent to setting

$$\begin{aligned}
q_1(\boldsymbol{\tau}) &\propto \exp \left\{ \mathbb{E}_{\mathbf{Z}, \boldsymbol{\theta}} [\log p(\boldsymbol{\tau}, \boldsymbol{\theta}, \mathbf{Z}, \mathbf{x} \mid \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})] \right\} \\
&= \exp \left(\mathbb{E}_{\mathbf{Z}, \boldsymbol{\theta}} [\log p_2(\mathbf{x} \mid \mathbf{Z}, \boldsymbol{\theta}) + \log p_3(\mathbf{Z} \mid \boldsymbol{\tau}) + \log p_4(\boldsymbol{\tau} \mid \boldsymbol{\delta}) + \log p_5(\boldsymbol{\theta} \mid \boldsymbol{\eta}, \boldsymbol{\nu})] \right) \\
&\propto \exp \left(\mathbb{E}_{\mathbf{Z}} [\log p_3(\mathbf{Z} \mid \boldsymbol{\tau}) + \log p_4(\boldsymbol{\tau} \mid \boldsymbol{\delta})] \right) \\
&= \exp \left(\mathbb{E}_{\mathbf{Z}} \left[\sum_{n=1}^N \sum_{m=1}^M \sum_{g=1}^G Z_{nmg} \log \tau_{ng} + \sum_{n=1}^N \sum_{g=1}^G (\delta_g - 1) \log \tau_{ng} \right] \right) \\
&= \prod_{n=1}^N \prod_{g=1}^G \tau_{ng}^{\sum_{m=1}^M \mathbb{E}_{\mathbf{Z}} [Z_{nmg}] + (\delta_g - 1)} \\
&= \prod_{n=1}^N \prod_{g=1}^G \tau_{ng}^{\gamma_{ng} - 1},
\end{aligned}$$

which we recognise as a Dirichlet distribution, and where we have introduced the variational parameter $\boldsymbol{\gamma}$.

Similarly, to maximise Eq. (5) with respect to $q_2(\mathbf{Z})$ set:

$$\begin{aligned}
q_2(\mathbf{Z}) &\propto \exp \left(\mathbb{E}_{\boldsymbol{\tau}, \boldsymbol{\theta}} [\log p(\boldsymbol{\tau}, \boldsymbol{\theta}, \mathbf{Z}, \mathbf{x} \mid \boldsymbol{\delta}, \boldsymbol{\eta}, \boldsymbol{\nu})] \right) \\
&\propto \exp \left(\mathbb{E}_{\boldsymbol{\theta}} [\log p_2(\mathbf{X} \mid \mathbf{Z}, \boldsymbol{\theta})] + \mathbb{E}_{\boldsymbol{\tau}} [\log p_3(\mathbf{Z} \mid \boldsymbol{\tau})] \right) \\
&= \exp \left(\sum_{n=1}^N \sum_{m=1}^M \sum_{g=1}^G Z_{nmg} \mathbb{E}_{\boldsymbol{\theta}} [\log p_1(x_{nm} \mid \boldsymbol{\theta}_{gm})] + \sum_{n=1}^N \sum_{m=1}^M \sum_{g=1}^G Z_{nmg} \mathbb{E}_{\boldsymbol{\tau}} [\log \tau_{ng}] \right) \\
&= \prod_{n=1}^N \prod_{m=1}^M \prod_{g=1}^G \exp \left\{ \mathbb{E}_{\boldsymbol{\theta}} [\log p_1(x_{nm} \mid \boldsymbol{\theta}_{gm})] + \mathbb{E}_{\boldsymbol{\tau}} [\log(\tau_{ng})] \right\}^{Z_{nmg}} \\
&= \prod_{n=1}^N \prod_{m=1}^M \prod_{g=1}^G \phi_{nmg}^{Z_{nmg}}.
\end{aligned}$$

This can be recognised as a multinomial distribution, with the variational parameter $\boldsymbol{\phi}$.

The variational approximation $q_3(\boldsymbol{\theta}_{gm})$ has the form:

$$\begin{aligned}
q_3(\boldsymbol{\theta}_{gm}) &\propto \exp \left\{ \mathbb{E}_{\mathbf{Z}} \left[\sum_{n=1}^N \log p_2(x_{nm} | \boldsymbol{\theta}_{gm}, Z_{nm}) \right] + \log p_5(\boldsymbol{\theta}_{gm} | \eta_{gm}, \boldsymbol{\nu}_{gm}) \right\} \\
&= \exp \left\{ \sum_{n=1}^N \mathbb{E}_{\mathbf{Z}} [Z_{nm} \log p_1(x_{nm} | \boldsymbol{\theta}_{gm})] + \log p_5(\boldsymbol{\theta}_{gm} | \eta_{gm}, \boldsymbol{\nu}_{gm}) \right\} \\
&= \exp \left\{ \sum_{n=1}^N \mathbb{E}_{\mathbf{Z}} [Z_{nm}] \left(\log k(\boldsymbol{\theta}_{gm}) + \log h(x_{nm}) + \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \mathbf{s}(x_{nm}) \right) \right. \\
&\quad \left. + \log h(\eta_{gm}, \boldsymbol{\nu}_{gm}) + \eta_{gm} \log k(\boldsymbol{\theta}_{gm}) + \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \boldsymbol{\nu}_{gm} \right\} \\
&\propto \exp \left\{ \left(\sum_{n=1}^N \mathbb{E}_{\mathbf{Z}} [Z_{nm}] + \eta_{gm} \right) \log k(\boldsymbol{\theta}_{gm}) + \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \left(\sum_{n=1}^N \mathbb{E}_{\mathbf{Z}} [Z_{nm}] \mathbf{s}(x_{nm}) + \boldsymbol{\nu}_{gm} \right) \right\} \\
&= k(\boldsymbol{\theta}_{gm})^{\sum_{n=1}^N \mathbb{E}_{\mathbf{Z}} [Z_{nm}] + \eta_{gm}} \exp \left\{ \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \left(\sum_{n=1}^N \mathbb{E}_{\mathbf{Z}} [Z_{nm}] \mathbf{s}(x_{nm}) + \boldsymbol{\nu}_{gm} \right) \right\} \\
&= k(\boldsymbol{\theta}_{gm})^{\eta'_{gm}} \exp \left\{ \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \boldsymbol{\nu}'_{gm} \right\} \\
&= p_5(\boldsymbol{\theta}_{gm} | \eta'_{gm}, \boldsymbol{\nu}'_{gm}),
\end{aligned}$$

where we have introduced the variational parameters η' and $\boldsymbol{\nu}'$. Thus $q_3(\boldsymbol{\theta}_{gm})$ will be the member of the same exponential family as the prior $p_5(\boldsymbol{\theta}_{gm} | \eta_{gm}, \boldsymbol{\nu}_{gm})$.

Parameter updates in terms of these variational parameters are as follows:

$$\begin{aligned}
\gamma_{ng} &= \sum_{m=1}^M \phi_{nm} + \delta_g; \\
\phi_{nm} &= \exp \left\{ \mathbb{E}_{\boldsymbol{\theta}} \left[\log k(\boldsymbol{\theta}_{gm}) + \log h(x_{nm}) + \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \mathbf{s}(x_{nm}) \right] + \Psi(\gamma_{ng}) - \Psi \left(\sum_{h=1}^G \gamma_{nh} \right) \right\} \\
&\propto \exp \left\{ \mathbb{E}_{\boldsymbol{\theta}} \left[\log k(\boldsymbol{\theta}_{gm}) + \mathbf{r}(\boldsymbol{\theta}_{gm})^\top \mathbf{s}(x_{nm}) \right] + \Psi(\gamma_{ng}) \right\}; \\
\eta'_{gm} &= \sum_{n=1}^N \phi_{nm} + \eta_{gm}; \\
\boldsymbol{\nu}'_{gm} &= \sum_{n=1}^N \phi_{nm} \mathbf{s}(x_{nm}) + \boldsymbol{\nu}_{gm},
\end{aligned}$$

where Ψ denotes the digamma distribution (Abramowitz and Stegun, 1965).

In the case of Poisson/Gamma distributed data, the updates for ϕ , $\boldsymbol{\nu}'$ and η' become:

$$\begin{aligned}
\phi_{nm} &\propto \exp \left(\frac{\nu'_{gm} + 1}{\eta'_{gm}} + (\Psi(\nu'_{gm} + 1) - \log(\eta'_{gm})) X_{nm} + \Psi(\gamma_{ng}) \right); \\
\eta'_{gm} &= \sum_{n=1}^N \phi_{nm} + \beta; \\
\nu'_{gm} &= \sum_{n=1}^N \phi_{nm} x_{nm} + \alpha - 1.
\end{aligned}$$

Nuisance Parameter

When treated as a nuisance parameter, the parameter update for θ can be obtained by direct maximum likelihood estimation of Equation (1). In this case, the log posterior becomes

$$\log p_9(\mathbf{x} \mid \delta, \theta) \geq \int \sum_{\mathbf{Z}} q_4(\mathbf{Z}, \tau) \log p_9(\mathbf{x} \mid \delta, \theta) d\tau - \int \sum_{\mathbf{Z}} q_4(\mathbf{Z}, \tau) \log q_4(\mathbf{Z}, \tau) d\tau.$$

The form of $q_1(\tau)$ and update for γ remain unchanged. While the form of $q_2(\mathbf{Z})$ is the same, the calculation of ϕ differs, however:

$$\begin{aligned} q_2(\mathbf{Z}) &\propto \exp(\log p_2(\mathbf{x} \mid \mathbf{Z}, \theta) + \mathbb{E}_{\tau}[\log p_3(\mathbf{Z} \mid \tau)]) \\ &= p_2(\mathbf{X} \mid \mathbf{Z}, \theta) \times \exp(\mathbb{E}_{\tau}[\log p_3(\mathbf{Z} \mid \tau)]) \\ &= \prod_{n=1}^N \prod_{m=1}^M \prod_{g=1}^G \{p_1(x_{nm} \mid \theta_{gm}) \times \exp(\mathbb{E}_{\tau}[\log \tau_{ng}])\}^{Z_{nmg}}. \end{aligned}$$

Thus the update for ϕ becomes $\phi_{nmg} \propto p_1(x_{nm} \mid \theta_{gm}) \times \exp\{\Psi(\tau_{ng})\}$.

The maximum likelihood estimate $\hat{\theta}$ is achieved by solving

$$\sum_{n=1}^N \nabla \log p_3(x_{nm} \mid \theta_{gm}, Z_{nmg}) \bigg|_{\theta_{gm} = \hat{\theta}_{gm}} = 0.$$

Substituting in the estimate ϕ for \mathbf{Z} , and noting that

$$\begin{aligned} \sum_{n=1}^N \nabla \log p_3(x_{nm} \mid \theta_{gm}, Z_{nmg}) &= \sum_{n=1}^N \phi_{nmg} \nabla \log p_1(x_{nm} \mid \theta_{gm}) \\ &= \sum_{n=1}^N \phi_{nmg} \left(\frac{\nabla k(\theta_{gm})}{k(\theta_{gm})} + \nabla \mathbf{r}(\theta_{gm})^\top \mathbf{s}(x_{nm}) \right), \end{aligned}$$

an estimate of $\hat{\theta}$ can then be obtained by solving:

$$-\nabla \mathbf{r}^{-1}(\hat{\theta}_{gm}) \frac{\nabla k(\hat{\theta}_{gm})}{k(\hat{\theta}_{gm})} = \frac{\sum_{n=1}^N \phi_{nmg} \mathbf{s}(x_{nm})}{\sum_{n=1}^N \phi_{nmg}}.$$

In the case of the Poisson distribution this becomes:

$$\begin{aligned} \phi_{nmg} &\propto \exp(-\theta_{gm}) \theta_{gm}^{x_{nm}} \times \exp(\Psi(\gamma_{ng})) \\ \hat{\theta}_{gm} &= \frac{\sum_{n=1}^N \phi_{nmg} x_{nm}}{\sum_{n=1}^N \phi_{nmg}}. \end{aligned}$$

In addition to estimating the profile memberships \mathbf{Z} and model parameters τ and θ , Erosheva et al. (2007); Airoldi et al. (2006, 2007) propose to estimate the hyperparameter δ using an empirical Bayes method. We omit this step from our analysis, and in the data analysis described in Section 4, we set $\delta_g = 1/G$, for all g .

3.2 Model Selection and Likelihood Estimation

While model assumptions require the number of profiles G to be fixed and known, in reality this is not the case. We therefore run the model over a range of values of $G' = 1, \dots, G^{\max}$, and compare the models post-hoc. While Airoldi et al. (2006) use the variational approximation to Equation (4) as a surrogate for the Bayesian Information Criterion (BIC) (Schwarz, 1978), in our opinion, the fact that the approximation (3) provides only a lower bound to the model posterior (1) makes the use of such a criterion difficult to interpret.

Rogers et al. (2005) propose evaluating the hold-out likelihood of the model, which involves integrating $\boldsymbol{\tau}$ and \mathbf{Z} from the complete-data posterior given in (1). In the case of the Poisson distribution with $\boldsymbol{\theta}$ a nuisance parameter, this becomes:

$$p_9(\mathbf{x}|\boldsymbol{\theta}, \boldsymbol{\delta}, G') \propto \prod_{n=1}^N \left\{ \int_{\boldsymbol{\tau}_n} \prod_{m=1}^M \sum_{g=1}^{G'} \tau_{ng} \frac{\exp\{-\boldsymbol{\theta}_{gm}\} \boldsymbol{\theta}_{gm}^{x_{nm}}}{x_{nm}!} p(\tau_{ng}|\delta_g) d\boldsymbol{\tau}_n \right\}. \quad (6)$$

Equation (6) may be approximated using a Monte Carlo method, by averaging over T draws from the prior $p(\boldsymbol{\tau}|\boldsymbol{\delta})$:

$$p_9(\mathbf{x}|\boldsymbol{\theta}, \boldsymbol{\delta}, G') \approx \prod_{n=1}^N \left\{ \frac{1}{T} \sum_{t=1}^T \prod_{m=1}^M \sum_{g=1}^{G'} \tau_g^{(t)} \frac{\exp\{-\boldsymbol{\theta}_{gm}\} \boldsymbol{\theta}_{gm}^{x_{nm}}}{x_{nm}!} \right\}.$$

3.3 Model Evaluation

While parameter estimates are used to interpret the model fitted in Section 4, we also make use of the following statistics, which further help to summarise the data. For convenience these are briefly described here.

Extent of profile membership (EoM) The extent to which an observation's attributes appear to be generated by multiple profiles can be estimated using a measure such as EoM (Hill, 1973; White et al., 2012), where $\text{EoM}_n = \exp(H(\hat{\boldsymbol{\tau}}_n))$, and H denotes the entropy function, $H(\hat{\boldsymbol{\tau}}_n) = -\sum_{g=1}^G \hat{\tau}_{ng} \log \hat{\tau}_{ng}$. This estimates the number of profiles from which an observation's variables seem to be drawn. Thus considering the EoM over all observations gives an idea of the amount of mixed membership taking place in the data.

Maximum a posteriori ($\hat{\mathbf{Z}}$) We can impose a hard clustering by mapping individuals to their most probable profile memberships for each attribute by setting $\hat{Z}_{nm} = \arg \max_{g=1, \dots, G} \{\mathbb{P}(\text{profile } g|x_{nm})\}$, where $\mathbb{P}(\text{profile } g|x_{nm})$, the probability that the observed value x_{nm} results from profile g , is estimated by $\hat{\phi}_{nmg}$. It can be shown that every mixed membership model can be re-expressed as a finite mixture model with a much larger number of components (Erosheva et al., 2007; Galyardt, 2014). In effect, these components consist of the distinct permutations of profile membership which occur across attributes in the data. One can think of the profile mapping summary statistic $\hat{\mathbf{Z}}$ as an estimate of this quantity.

We use the notation $\{a, b\}$ to indicate the set of individuals whose assigned membership across attributes is some (repeated) permutation of profiles a and

b. In other words, an observation n is an element of $\{a, b\}$, if a and b are the unique elements in $\hat{\mathbf{Z}}_n$. Note that this notation can be used for any number of profiles: for example, $\{1\}$ indicates the individuals who exclusively map to profile 1 across all attributes.

Classification uncertainty (U) Another way to scrutinise classification is to consider the uncertainty associated with an observation's profile assignment for each of their attributes (Bensmail et al., 1997): $U_{nm} = \min_{g=1, \dots, G} \{1 - \mathbb{P}(\text{profile } g | x_{nm})\}$, where the lower the uncertainty, the better the classification.

3.4 Mixture Model Framework

In Section 4 the mixed membership approach is compared to the standard MBC approach. To fit a model using the mixture model framework (Everitt and Hand, 1981), we first assume a fixed number G of groups underly the data. We use this term exclusively for mixture models. Let $\boldsymbol{\tau}^{\text{mix}} = (\tau_1^{\text{mix}}, \dots, \tau_G^{\text{mix}})$ denote the prior probability that an observation belongs to each group. Consequently, the likelihood $p_{\text{mix}}(\mathbf{x} | \boldsymbol{\theta}^{\text{mix}}, \boldsymbol{\tau}^{\text{mix}})$ then takes the form

$$p_{\text{mix}}(\mathbf{x} | \boldsymbol{\theta}^{\text{mix}}) = \prod_{n=1}^N \sum_{g=1}^G \tau_g^{\text{mix}} \prod_{m=1}^M p_6(x_{nm} | \theta_{gm}^{\text{mix}}),$$

where $p_6(x_{nm} | \theta_{gm}^{\text{mix}})$ is defined as in Equation 2. Direct inference of this likelihood is difficult, but can be facilitated with the introduction of missing data $\mathbf{Z}^{\text{mix}} = (\mathbf{Z}_1^{\text{mix}}, \dots, \mathbf{Z}_N^{\text{mix}})$, and $\mathbf{Z}_n^{\text{mix}} = (Z_{n1}^{\text{mix}}, \dots, Z_{nG}^{\text{mix}})$, for each $n = 1, \dots, N$. We define

$$Z_{ng}^{\text{mix}} = \begin{cases} 1 & \text{if observation } n \text{ is member of Group } g; \\ 0 & \text{otherwise.} \end{cases}$$

From a clustering perspective, each Z_n^{mix} can be interpreted as a latent variable indicating cluster membership (Fraley and Raftery, 2002). Note that within the mixture model framework, conditional on group membership, observations are assumed to be drawn independently.

We can use similar summary statistics to evaluate the clustering performance of a mixture model to those described in Section 3.3. In particular, define $\hat{Z}_n^{\text{mix}} = \arg \max_{g=1, \dots, G} \{\mathbb{P}(\text{group } g | \mathbf{x}_n)\}$, and $U_n^{\text{mix}} = \min_{g=1, \dots, G} \{1 - \mathbb{P}(\text{group } g | \mathbf{x}_n)\}$. These map individual observations to groups and assess the uncertainty of this classification respectively. Note that these values assign a single value to each observation (across all attributes), as opposed to the statistics for mixed membership, which potentially assign different values to an observation's attributes.

We omit further details of how inference is performed, except to mention that parameter estimates may be obtained using an EM algorithm (Dempster et al., 1977). To determine the optimal number of clusters in the data, the model was run over a large number of groups, and the BIC was used to identify the optimal number to fit to the data. While the regularity conditions required for the BIC are not met when choosing the number of groups for a mixture model (Biernacki et al., 2000), at a practical level it has proved useful on many occasions (Fraley and Raftery, 2002). To perform inference in a Bayesian setting, conjugate priors can

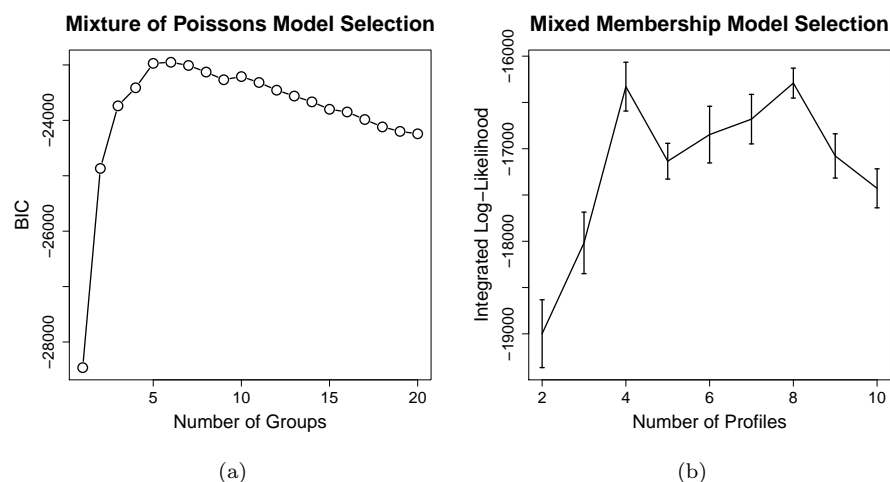


Figure 2: BIC (a) and hold-out likelihood (b) values for mixture and mixed membership models applied to the IAU running data. Within the respective frameworks, 6 component (group) and 4 profile models appear to fit the data optimally.

be chosen in a similar fashion to those already described. The use of priors with different (sensible) choices of hyper-parameters were found to have little effect on the clustering obtained by the application in Section 4.

4 International Association of Ultrarunners 24 Hour World Championships

The International Association of Ultrarunners (IAU) 24 hour World Championships were held in Katowice, Poland on September 8th to 9th, 2013. Two hundred and sixty athletes representing twenty four countries entered the race, which was held on a course consisting of a 1.554 km looped route. An update of the number of laps covered by each athlete was recorded approximately every hour².

Note that the sequential nature of the data means that the exchangeability assumption required by the mixed membership model discussed in Section 2, as well as the conditional independence assumption required by the mixture model, may both be somewhat unrealistic in this setting. Nevertheless, the approaches appear to identify interesting behaviour in the data, and serve to illustrate important differences between the methods. Both mixture and mixed membership models were applied to the dataset, with the BIC and hold out likelihood suggesting that 6-component and 4-profile fits were optimal; this is illustrated in Figure 2.

² A version of this data is available at <http://mathsci.ucd.ie/~brendan/data/24H.xlsx>

4.1 Mixture Model Application

The estimated weight parameters for the 6-component mixture model were $\tau^{\text{mix}} = (0.40 \ 0.33 \ 0.08 \ 0.07 \ 0.06 \ 0.06)$. The estimated values of θ^{mix} are illustrated in Figure 3a. This figure suggests that the two largest groups (Groups 1 and 2) in the dataset ran at a reasonably steady rate over the course of the race, with Group 2's pace declining in a slightly more pronounced manner during the second half of the race. Three of the four remaining smaller groups, Groups 3, 4, and 6, began the race at a similarly high pace to Groups 1 and 2, but were unable to sustain such a rate over the duration of the race. In particular, runners in Groups 3 and 6 failed to complete many laps beyond the 18 and 12 hour marks respectively, while runners clustered in Group 4 maintained a steadier pace throughout the race, and actually improved slightly over the final four hours. Finally, Group 5 consisted of entrants who completed only a very small number of laps over the course of the race, including several runners who completed no laps; this includes race entrants who failed to participate on the day of the race.

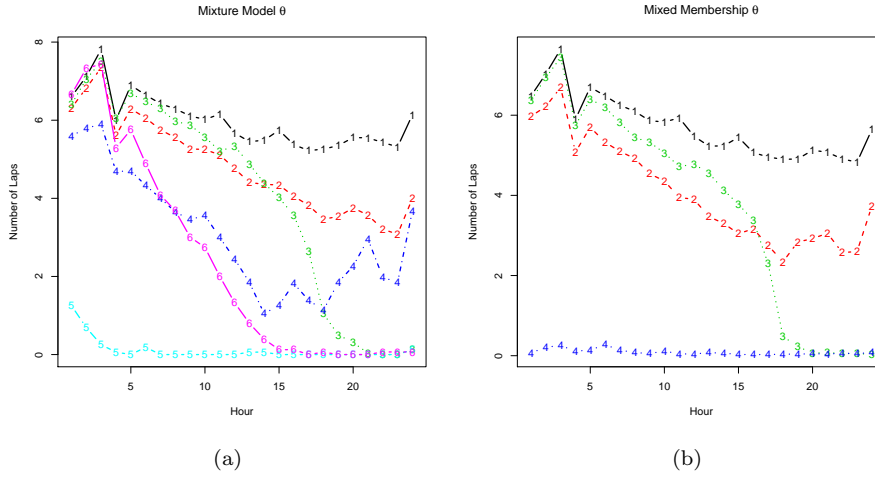


Figure 3: Plots of the expected number of laps completed per hour, conditional on group and profile membership, for (a) mixture model with 6 components (groups) and (b) mixed membership model with 4 profiles respectively.

4.2 Mixed Membership Model Application

The estimated values of θ for the 4-profile mixed membership model are illustrated in Figure 3b. Based on this plot, profile behaviour conveys much of the same information as the mixture model: over the course of the race, the characteristic behaviour of Profile 1 is to perform at a high and steady rate; Profile 2 is at a similarly steady but slower pace; Profile 3 begins brightly but declines sharply

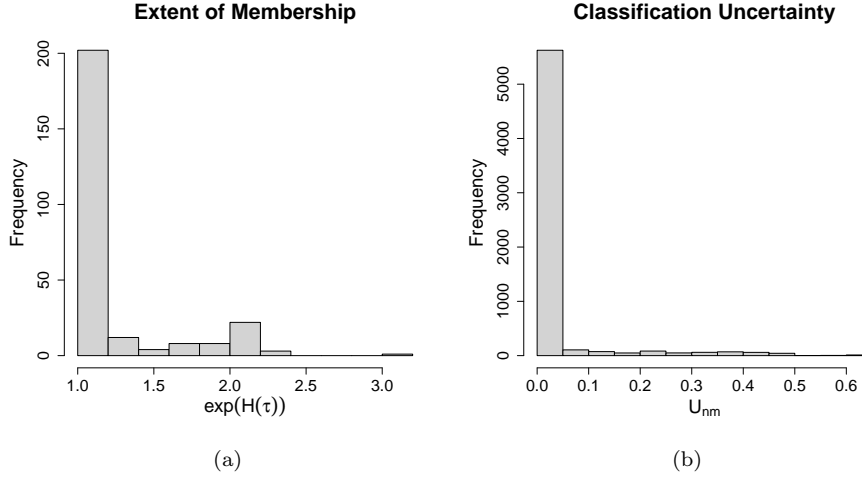


Figure 4: Histograms of the extent of profile membership and classification uncertainty of observations in the running data. Roughly 17% of observations exhibit membership between 2 profiles, with one observation exhibiting membership to 3 profiles. Over 90% of datapoints are mapped to a profile with high certainty.

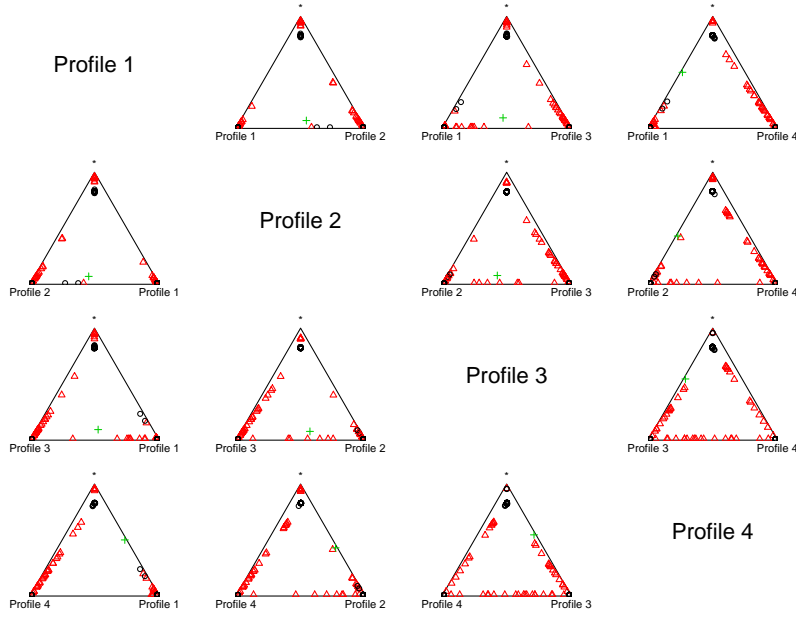
by the final quarter of the race; while Profile 4 can be characterised as exhibiting extremely low-level, non-participatory behaviour. For convenience, we refer to Profiles 1 to 4 by the following names: Fast Pace, Slow Pace, Rapid Decline, and Non-Participation, respectively.

Figure 4a shows that while the majority of observations exhibit membership to only one profile, about 17% of observations exhibit at least some mixed membership, with all but one of these observations displaying membership between two profiles. In the mixed membership setting, about 90% of datapoints are classified with uncertainty less than 5%, substantially higher than the mixture model clustering, in which only 73% of observations were clustered with the same level of certainty. Some datapoints are still classified with high uncertainty by the mixed membership clustering; see Figure 4b.

Table 1: Table detailing which profiles runners map to in mixed membership clustering. Note that the hours of the race which observations map onto different profiles is not provided in this table.

Mapped profile memberships										
{1}	{2}	{3}	{4}	{1,2}	{1,3}	{1, 4}	{1, 2, 3}	{2, 3}	{2,4}	{3, 4}
137	42	16	13	1	9	6	1	7	9	19

A direct inspection of $\hat{\mathbf{Z}}$ shows that a total of 208 of the 260 observations map directly onto one profile, that is, displayed no mixed membership. All except one of the remaining 52 observations display membership across no more than two profiles



(a)

Figure 5: Plots of the marginal simplices representing runners' profile membership. Colour and shape are used to denote the number of profiles datapoints map onto. 80% of observations map to only one profile (black circles), and are closely grouped together. With one exception (the green cross), all remaining datapoints map onto two profiles (red triangles), and are arranged along the edgepoints of the simplices.

at one time. The 3-dimensional simplex is visualised using a ternary plot (van den Boogaart and Tolosana-Delgado, 2008) in Figure 5a. (N.B., recall that we fixed the hyperparameter $\delta_g = 1/4$ in the fitted model.) In cases of mixed membership, the strongest association is between profiles 3 and 4, the Rapid Decline and Non-Participation profiles, as shown in Table 1. The 19 runners exhibiting mixed membership to both these profiles can be characterised as runners starting strongly but whose performance tailed off at various points during the race. While this description is similar to that for the behaviour characterised by the Rapid Decline profile, the behaviour of the two groups is still quite different. Figure 6a shows the percentage of the runners mapped to $\{3, 4\}$ and the 16 runners mapped to $\{3\}$ who completed at least one lap during each hour of the race; the slope of this line for runners belonging exclusively to the Rapid Decline profile is markedly steeper. This indicates that the pace of runners in $\{3, 4\}$ decline over a much wider time frame.

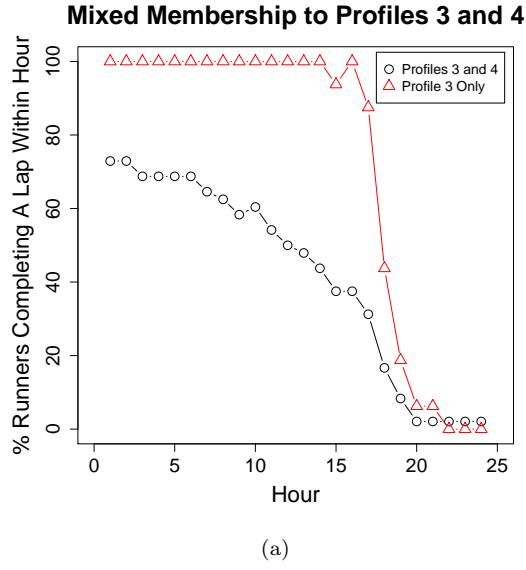


Figure 6: The percentage of runners in $\{3, 4\}$ and $\{3\}$ who completed at least one lap during each hour of the race.

4.3 Comparing the Models

Table 2: Table comparing clusters found using the mixture model framework to the mapped profiles from the mixed membership.

	Mapped profile memberships										
	{1}	{1,3}	{1,4}	{1,2,3}	{2}	{2,1}	{2,3}	{2,4}	{3}	{3,4}	{4}
Group 1	98	0	3	0	0	0	0	0	0	0	0
Group 2	39	8	3	1	33	1	4	0	0	0	0
Group 3	0	1	0	0	0	0	3	0	16	2	0
Group 4	0	0	0	0	9	0	0	8	0	0	0
Group 5	0	0	0	0	0	0	0	1	0	2	13
Group 6	0	0	0	0	0	0	0	0	0	15	0

We now compare the clusters found by the mixed membership and mixture modelling frameworks. Table 2 shows how overlap between the mapped profile memberships from the mixed membership approach compared to the membership of the six groups found using the mixture model framework. Note that 98 of the 101 runners mapped to Group 1 match to $\{1\}$, the Fast Pace profile. The three runners in the group who exhibit mixed membership do so to $\{1,4\}$, the Fast Pace and Non-Participation profiles. The runners mapped to these two profiles all ran at a high pace, but failed to complete any laps (possibly stopping completely for that time) for a single hour at different points in the race. Runners clustered together

in Group 2 by the mixture model approach are mainly split between $\{1\}$ and $\{2\}$, the Fast and Slow Pace profiles in the mixed membership approach. The runners in this group exhibiting mixed membership are similar to those with membership of two profiles in Group 1 in that they run at a high pace but stop, or fail to complete a lap, intermittently, before returning to the previous pace. Group 3 corresponds closely to $\{3\}$, the Rapid Decline profile, while Group 4 matches to either $\{2\}$ or $\{2,4\}$, the Slow Pace and Non-Participation profiles, again indicating that some runners in this group raced only intermittently. Members of Group 5 are mainly clustered to $\{4\}$, the Non-Participation profile, which is perhaps unsurprising. Members of Group 6 are all members of $\{3,4\}$, the Rapid Decline and Non-Participation profiles; this behaviour has been discussed in the previous subsection. This indicates that perhaps Group 6, the smallest group in the fitted mixture model, was a poor fit to the data; rather than being a group of runners whose pace gradually decreased, it consisted of a group of runners completing a large number of laps an hour, with various members of the group withdrawing early at different stages in the race.

4.4 Examples of Mixed Membership

In this section, in order to illustrate the types of mixed membership exhibited by the data, the three race entrants with the highest EoM scores are discussed, in decreasing order. Plots of each runner's lap numbers and profile assignment scores over the course of the race are given in Figure 7.

Runner 158 This was the only race entrant to be mapped to three profiles over the course of the race. Inspecting Figure 7a, it's clear that a high level of uncertainty is associated with this runner's profile membership throughout the race, until the last hour, when their lap time is associated with the Rapid Decline profile with a high level of certainty. Figure 7a shows the runners data, along with the estimated values of θ . From this we can see that for the first half of the race, the runner ran at a good pace, consistent with both the Fast Pace and Rapid Decline profiles. On the 18th hour, this runner experienced a large dip in pace consistent with the Rapid Decline profile, but recovered at hours 21 and 22, again running at a pace more consistent with the Fast and Slow Pace profiles, before eventually fading again for the last two hours.

Runner 237 This runner's performance is characterised as being split between the Rapid Decline and Non-Participation profiles, a type of mixed membership discussed previously. This runner starts well, but does not complete any laps past the 10th hour. Note the high level of uncertainty of profile membership for the last six hours (Figure 7c); this is explained by the fact that the values of θ are very close together for Profiles 3 and 4 for these hours, and that this runner has evenly split profile membership between the two profiles for the hours before that in the race.

Runner 15 This runner's profile membership was split between the Slow Pace and Non-Participation profiles. This runner's race can be characterised as running at a relatively low pace, while stopping for several hours on two occasions before completing a reasonably high number of laps during the final four hours of the race. Despite this runner's erratic behaviour, given an hour m and profile membership g , the number of laps they complete is usually quite close

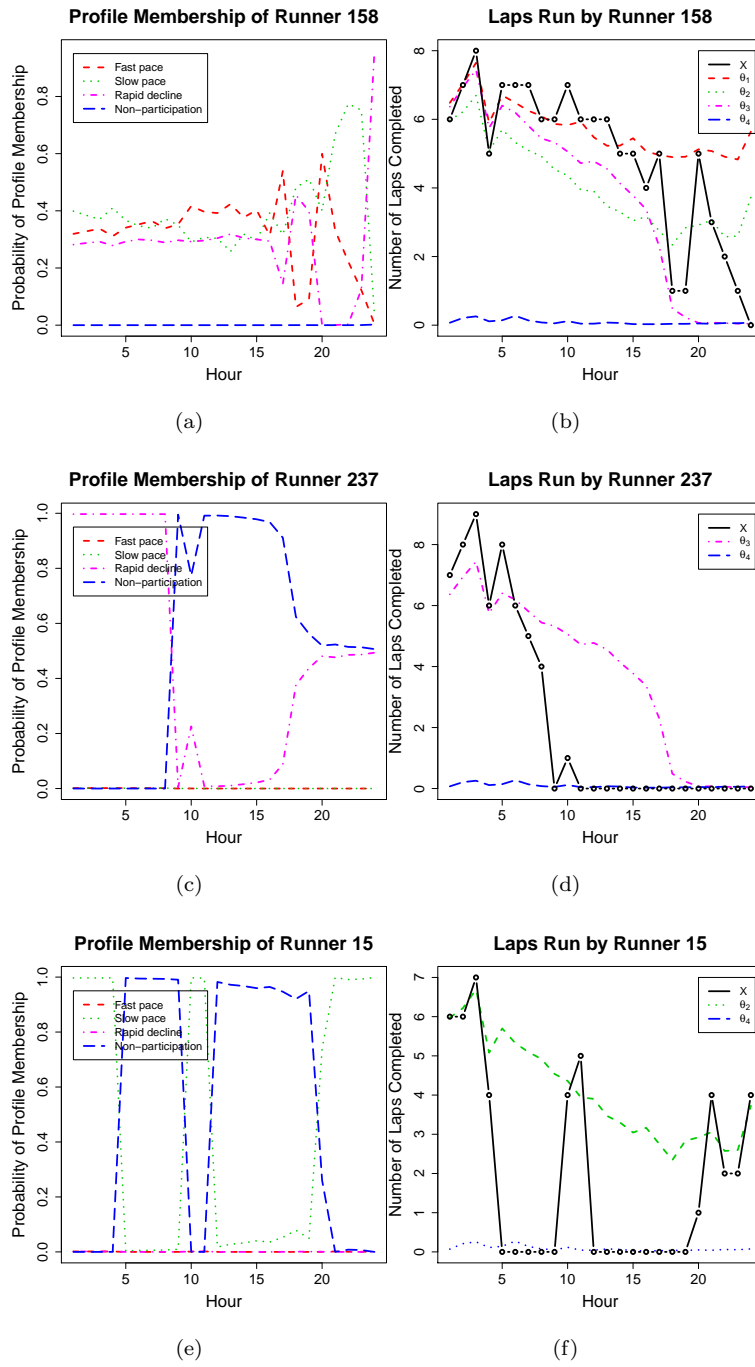


Figure 7: Examples of runners exhibiting mixed membership for the 4-profile fit.

to the value θ_{gm} . In this case, the exchangeability assumption of the mixed membership model is arguably advantageous; a model that incorporated too much dependence between race laps could be over smooth by comparison.

5 Discussion

It is clear that mixed membership methods provide the analyst with tools of greater flexibility than current MBC or standard distance-based clustering methods. While the mixed membership framework is more elaborate than that of the mixture model, our application makes clear the benefits that the method provides, and that its output can be interpreted and understood. While the nature of the running data seems to be better modelled by a mixed membership approach, at least in a qualitative sense, it is difficult to show this quantitatively, and the question of how to compare different types of clustering method in general remains open.

While in theory it is possible to obtain equivalent clusterings of observations using mixed membership and mixture models, we argue that this is unlikely to occur in practice. For example, in the application to the running data, since several observations have unique profile mappings – for example, Runner 15 stops several times – this would suggest an equivalent clustering solution in the mixture model framework would contain many singleton clusters. Typically such clusterings are considered unfavourable. However within the mixed membership framework, the unique aspects of the runner’s behaviour are well explained in this case.

In this paper we have provided a mixed membership formulation for data produced by members of an exponential family with an underlying latent mixed membership structure. It may be of interest to expand this model further to account for mixed-type data, similar to the procedure for mixture models introduced by Vermunt and Magidson (2002). The simplifying assumption of exchangeability made by the model, as discussed in Section 2, may be somewhat unrealistic; for example, in the running data, runners with partial membership to profile 4 tend to be assigned membership later rather than earlier in the race. While in a general sense, as noted by Blei et al. (2003), it may be difficult to justify the epistemological validity of such an assumption, its utility in a clustering framework is clear. In particular, when applied to the running data, the mixed membership approach effectively captures the sporadic nature with which runners stopped throughout the race.

A potential weakness of the model as currently formulated is the use of the Dirichlet distribution to model each observation’s profile membership. The use of this distribution reflects the assumption that the profile membership of an observation’s attributes can be thought of as exchangeable entities, causing any correlation within the data to be ignored. Thus the model may have poor posterior predictive power. While not an explicit aim of this paper, it is a limitation of the current model. One solution is to replace the Dirichlet distribution with a logistic normal distribution (Blei and Lafferty, 2007) although this complicates the inference method. Wang and Blei (2013) have outlined methods for performing inference in a variational Bayes setting when the posterior form is non-conjugate. Additionally, longitudinal mixed membership models have been developed. Manrique-Vallier (2014) explicitly models profile behaviour as a function of time, while Blei

and Lafferty (2006) allow profile behaviour and the *a priori* probability of profile membership to evolve over time using a state space approach.

Acknowledgements This work is supported by Science Foundation Ireland under the Clique Strategic Research Cluster (08/SRC/I1407) and Insight Research Centre grant (SF1/12/RC/2289).

References

- Abramowitz, M., and I. A. Stegun. 1965. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, 1st edn. Dover Publications.
- Airoldi, E. M., D. Blei, E. Erosheva, and S. E. Fienberg. 2014. Introduction to mixed membership models and methods. In *Handbook of mixed membership models*, eds. E. M. Airoldi, D. Blei, E. Erosheva, and S. E. Fienberg. Chapman & Hall/CRC. Chap. 1.
- Airoldi, E. M., S. E. Fienberg, C. Joutard, and T. Love. 2006. Discovering latent patterns with hierarchical Bayesian mixed-membership models., Technical report, Carnegie Mellon University, School of Computer Science, Machine Learning Department.
- Airoldi, E. M., S. E. Fienberg, C. Joutard, and T. Love. 2007. Discovering latent patterns with hierarchical Bayesian mixed-membership models. In *Data mining patterns: New methods and applications*, eds. P. Poncelet, M. Teisseire, and Massegia F. Idea Group Inc. Chap. 11.
- Baudry, J. P., A. E. Raftery, G. Celeux, K. Lo, and R. Gottardo. 2010. Combining mixture components for clustering. *Journal of Computational and Graphical Statistics* 19 (2): 332–353.
- Beal, M. 2003. Variational algorithms for approximate Bayesian inference. PhD diss, University College London.
- Bensmail, H., G. Celeux, A. E. Raftery, and C. Robert. 1997. Inference in model-based cluster analysis. *Statistics and Computing* 7: 1–10.
- Biernacki, C., G. Celeux, and G. Govaert. 2000. Assessing a mixture model for clustering with the integrated completed likelihood. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 22 (7): 719–725. doi:10.1109/34.865189.
- Bishop, C. M. 2006. *Pattern recognition and machine learning*. Secaucus, NJ, USA: Springer.
- Blei, D. M., and J. D. Lafferty. 2006. Dynamic topic models. In *Proceedings of the 23rd International Machine Learning Conference*, eds. W. Cohen and A. Moore.
- Blei, D. M., and J. D. Lafferty. 2007. A correlated topic model of science. *Annals of Applied Statistics* 1 (1): 17–35.
- Blei, D. M., A. Y. Ng, and M. I. Jordan. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research* 3: 993–1022.
- Dempster, A. P., N. M. Laird, and D. B. Rubin. 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39 (1): 1–38. doi:10.2307/2984875. <http://dx.doi.org/10.2307/2984875>.
- Erosheva, E. A., S. E. Fienberg, and C. Joutard. 2007. Describing disability through individual-level mixture models for multivariate binary data. *Annals of Applied Statistics* 1 (2): 502–537.

- Erosheva, E. A., S. E. Fienberg, and J. Lafferty. 2004. Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences of the United States of America* 101: 5220–5227.
- Everitt, B. S., and D. J. Hand. 1981. *Finite mixture distributions*. London: Chapman and Hall.
- Fraley, C., and A. E. Raftery. 2002. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association* 97 (458): 611–631.
- Galyardt, A. 2014. Interpreting mixed membership models: Implications of Erosheva’s representation theorem. In *Handbook of mixed membership models*, eds. E. M. Airoldi, D. Blei, E. Erosheva, and S. E. Fienberg. Chapman & Hall/CRC. Chap. 11.
- Gormley, C., and T. B. Murphy. 2009. A grade of membership model for rank data. *Bayesian Analysis* 4 (2): 265–296.
- Hill, M. O. 1973. Diversity and evenness: A unifying notation and its consequences. *Ecology* 54 (2): 427–432.
- Manrique-Vallier, D. 2014. Longitudinal mixed membership trajectory models for disability survey data. *Ann. Appl. Stat.* 8 (4): 2268–2291.
- McLachlan, G., and D. Peel. 2002. *Finite mixture models*. Wiley.
- Ormerod, J. T., and M. P. Wand. 2010. Explaining variational approximations. *The American Statistician* 64 (2): 140–153.
- Rogers, S., M. Girolami, C. Campbell, and R. Breitling. 2005. The latent process decomposition of cDNA microarray datasets. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 2: 2005.
- Schwarz, G. 1978. Estimating the Dimension of a Model. *Annals of Statistics* 6 (2): 461–464.
- van den Boogaart, K. Gerald, and R. Tolosana-Delgado. 2008. compositions: A unified r package to analyze compositional data. *Computers & Geosciences* 34 (4): 320–338.
- Vermunt, J. K., and J. Magidson. 2002. Latent class cluster analysis. In *Applied latent class analysis*, eds. J. A. Hagenaars and A. McCutcheon, 89–106. Cambridge University Press.
- Wang, C., and D. Blei. 2013. Variational inference in nonconjugate models. *Journal of Machine Learning Research* 14: 1005–1031.
- White, A., J. Chan, C. Hayes, and T. B. Murphy. 2012. Mixed membership models for exploring user roles in online fora. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media (ICWSM 2012)*, eds. N. Ellison, J. G. Shanahan, and Z. Tufekci, 599–602.