Stat
The ISI's Journal for the Rapid
Dissemination of Statistics Research
(wileyonlinelibrary.com) DOI: 10.100X/sta.0000

Nonparametric Bayes to Infer Playing Strategies Adopted in a Population of Mobile Gamers

Seppo Virtanen[a]*, Mattias Rost[b], Matthew Higgs[b], Alistair Morrison[b], Matthew Chalmers[b], Mark Girolami[a]

Received 13 January 2015; Accepted 00 Month 2015

Analysis of trace logging data collections of interactions of a heterogenous and diverse population of consumers of digital software with mobile devices provides unprecedented possibilities for understanding how software is actually used and for finding recurring patterns of software usage over the population that are exhibited to greater or lesser degree in each individual software user. In this work, we consider an elementary mobile game played by a population of mobile gamers and collect pieces of game sessions over an extended period of time resulting in a collection of users' trace logs for multiple sessions. We develop a simple, yet flexible, nonparametric Bayes approach to infer playing strategies adopted in the population from the logged traces of game interactions. We demonstrate our approach finds interpretable strategies and provides good predictive performance compared to alternative modelling assumptions using a nonparametric Bayes framework. Copyright © 2012 John Wiley & Sons, Ltd.

Keywords: Bayesian methods; nonparametric methods; state-space models; statistical inference; statistical modelling

1. Introduction

The ubiquitous ability to monitor and measure many aspects of daily life is producing vast repositories of heterogenous data taking on many forms from tagged images to the time series describing, for example, an individuals blood-pressure and heart rate on a second by second basis. The need to analyse such data collections in a sound statistical manner that does not introduce bias, or highlights many confirmatory hypotheses which are just artefacts of the amount of data available, presents a number of methodological research challenges.

[a]**Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK**
[b]**School of Computing Science, University of Glasgow, Glasgow, G12 8RZ, UK**
***Email: s.virtanen@warwick.ac.uk**

In the particular case we consider in this paper our interest focuses on a population of consumers of online content in the form of mobile games. The longer term goal of this research is to consider statistical methods which will provide a means of extracting shared patterns of behaviour within a population that are exhibited to greater or lesser degree in each individual software user.

In this proof of concept study we consider a simple mobile game application played by a small population of users. The question posed is whether it is feasible to infer, in this case, playing strategies exhibited in the population from the logged traces of game interactions.

One of the important issues is to consider the evaluation of evidence supporting the possible number of emergent strategies being employed by the population and from this infer the plausible strategies themselves in an interpretable manner. In this setting nonparametric Bayesian methods lend themselves well to these two stages of inference and as such we consider them in some more detail.

Nonparametric Bayesian methods have recently been shown to provide useful modelling frameworks for a range of applications in varying domains, producing good predictive performance due to increased flexibility and robustness. In this work we develop a model based on the particular game being studied to assess the feasibility of inferring population traits in this setting. This form of hierarchical modelling allows to combine information from a heterogeneous and diverse population of users and accordingly make inferences.

In this work, we represent a logged software trace as a discrete sequence of states (actions) from a specific vocabulary and capture transition probabilities between consecutive states using stationary first-order discrete-state Markov models. Finite mixtures of Markov models and mixed membership Markov models have been proposed to account for inter-user variability previously by Andrei et al. (2014); Cadez et al. (2003); Girolami & Kabán (2004) and Higgs et al. (2013). Here, mixture components correspond to transition probability matrices. The mixture model assumes trace logs for each user are drawn from a single component ensuring that components capture meaningful structure and alleviating interpretation of the model. In the mixed membership model, the trace logs for each user is modelled with a mixture model, where the mixture proportions are user-specific, and the mixture components are shared across users. The latter model provides increased flexibility over mixture modelling but may comprise interpretability, since one is forced to look at user-specific combinations of components. While finite models require the number of components to be known, the nonparametric Bayesian formulations via Dirichlet and hierarchical Dirichlet processes, which are natural prior candidates for mixture and mixed membership modelling, respectively, determine posterior distribution over the number of components necessary to explain the data favouring parsimonious models.

We propose a novel approach combining the benefits of the two modelling approaches described above employing application-specific knowledge regarding the multi-session structure for each user. That is, we assume a user follows the same component during a single game session, but mixture assignments may vary from session to session following user-specific mixture proportions. Our hierarchical model formulation via HDP is flexible; subpopulations of users may choose highly specific combinations of components. This flexibility is required for coping with noisy logging data from a heterogenous population. We derive a variational Bayes algorithm approximating the analytically intractable posterior distribution of the model to allow subsequent inference and analytic study of what the data inform us about the population of gamers.

In Section 2, we first describe briefly the mobile game studied in this work and the corresponding data structure of trace logs. Second, we present our nonparametric Bayesian model for the data. Section 2.1 describes our variational Bayes inference algorithm. Section 3 presents alternative modelling assumptions. In Section 4, we first compare these different assumptions quantitatively in terms of predictive ability as well as computational performance and then present key findings of our model.

2. Model

We observe multiple users' trace logging data from a *Hungry Yoshi* game (McMillan et al., 2010). In brief, the game is about viewing creatures called yoshis and feeding them fruits they request. That is, each yoshi requires certain fruits that can be picked to a basket from plantations. Users can cultivate particular fruits on plantations by picking and planting seeds. We label each feed as good or bad depending on yoshi's request. The trace logs for each user (sequences of user actions) are structured into distinct (game) sessions of varying length indicated by starting and closing the application.

One session is denoted as $\mathbf{w}_j^{(m)}$ for the mth user and jth session, respectively. The number of sessions for the mth user is denoted as J_m. Each session consists of an ordered sequence of states,

$$\mathbf{w}_j^{(m)} = [w_{j,1}^{(m)} w_{j,2}^{(m)} \dots w_{j,N_j^{(m)}}^{(m)}],$$

where $w_{j,n}^{(m)}$ denotes the state at a discrete time stamp $n \in \{1, \dots, N_j^{(m)}\}$. Each state $w_{j,n}^{(m)}$ may correspond to one of S different values, encoded as integer values mapping to a vocabulary. The data collection for M users is referred to as $\mathcal{D}$.

We assume a generative model for the $\mathbf{w}_j^{(m)}$ is

$$p(\mathbf{w}_j^{(m)} | \mathbf{T}^{(C_j^{(m)})}) = p(w_{j,1}^{(m)}) \prod_{n=2}^{N_j^{(m)}} p(w_{j,n}^{(m)} | w_{j,n-1}^{(m)}). \tag{1}$$

In the following, we assume the initial state $w_{j,1}^{(m)}$ is given. A latent $S \times S$ state transition matrix $\mathbf{T}^{(C_j^{(m)})}$ underlies sequence generation,

$$p(w_{j,n}^{(m)} = i | w_{j,n-1}^{(m)} = l) = T_{l,i}^{(C_j^{(m)})}.$$

We assume a set of different transition matrices (termed as strategies) explain the data $\mathcal{D}$. The mth user selects a particular strategy for the jth session, indicated by an indicator variable $C_j^{(m)}$, following his preference (a distribution over the strategies). We present a nonparametric hierarchical prior distribution for the model (1). The prior is based on Dirichlet processes (DP; Ferguson, 1973) and their hierarchical extension (HDP; Teh et al., 2006). The generative description as well as constructive formulations for our model use both stick-breaking (Sethuraman, 1994) and normalised gamma processes (Ferguson, 1973; Ishwaran & Zarepour, 2002).

We assume *a priori* an infinite number of different strategies drawn from a base distribution. The strategies $\mathbf{T}^{(k)}$, where $k = 1, \dots, \infty$, are associated with probabilities p_k that are proportional to an overall popularity of the strategy among M users. We assume the base distribution is a Dirichlet distribution,

$$\mathbf{T}_{s,:}^{(k)} \sim \text{Dirichlet}(\gamma_s \mathbf{1})$$

for $s = 1, \dots, S$, and we define

$$p_k = V_k \prod_{l=1}^{k-1} (1 - V_l),$$

where

$$V_k \sim \text{Beta}(1, \alpha),$$

following the stick breaking process. The V_k for $k = 1, \ldots, \infty$ are weight parameters and α is a concentration parameter.

The mth user is assigned a user-specific probability distribution over the strategies, denoted as $\boldsymbol{\theta}^{(m)}$, reflecting his/her preference for a certain combination of strategies. We construct the $\boldsymbol{\theta}^{(m)}$ using the normalised gamma process. First, we draw an auxiliary variable

$$Z_k^{(m)} \sim \text{Gamma}(\beta p_k, 1),$$

where $\mathbf{p}$ are the strategy probabilities and β is another concentration parameter. Then, the user-specific strategy distribution (preference) is obtained by normalising the $\mathbf{Z}^{(m)}$,

$$\theta_k^{(m)} = \frac{Z_k^{(m)}}{\sum_{l=k'}^{\infty} Z_{k'}^{(m)}}.$$

Given the user-specific strategy distribution for the mth user, we draw indicator variables, which indicate the strategy chosen for the jth session,

$$C_j^{(m)} \sim \text{Discrete}(\boldsymbol{\theta}^{(m)}).$$

Finally, given the initial state $w_{j,1}^{(m)}$ and the $C_j^{(m)}$, where $j = 1, \ldots, J_m$, the consecutive states $w_{j,n}^{(m)}$, where $n = 2, \ldots, N_j^{(m)}$, are generated sequentially from the Markov model (1). This process is repeated for M users.

To summarise, a joint distribution of the model over the unobserved variables, denoted as Θ, and the data $\mathcal{D}$ (see also Figure 1 for a graphical illustration of the model) is

$$p(\mathcal{D}, \Theta) = \prod_{m=1}^{M} \prod_{j=1}^{J_m} \prod_{s=1}^{S} \prod_{k=1}^{\infty} p(\mathbf{w}_j^{(m)} | \mathbf{T}^{(C_j^{(m)})}) p(C_j^{(m)} | \mathbf{Z}^{(m)}) p(Z_k^{(m)} | \beta, \mathbf{V}_{1:k}) p(V_k | \alpha) p(\alpha) p(\beta) p(\mathbf{T}_{s,:}^{(k)} | \gamma_s),$$

where the distributions are

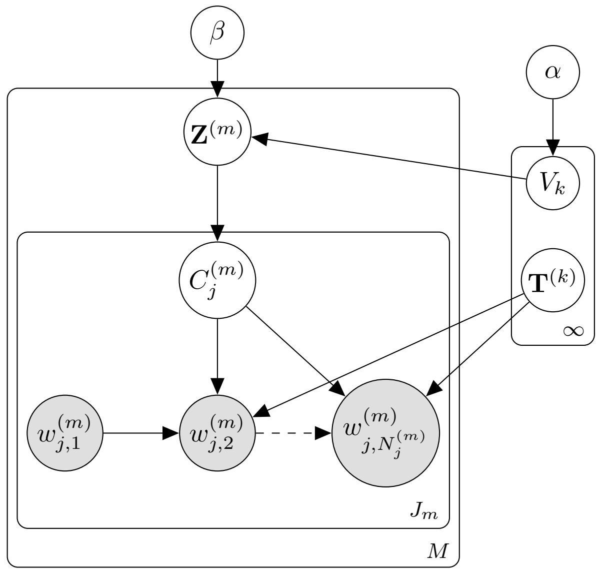

Figure 1. Plate diagram of the model. Unshaded nodes correspond to unobserved variables, whereas shaded nodes correspond to observed variables. Hyperpriors for the concentration parameters α and β are omitted from the visualisation for clarity. Plates indicate replication over strategies, users and sessions. The hidden variables may be divided into user-specific variables and variables common to all users. That is, the unnormalised strategy proportions $\mathbf{Z}^{(m)}$ and the strategy indicators $C_j^{(m)}$ are defined for each user, whereas the set of strategies $\mathbf{T}^{(k)}$ is common to all users.

$$
\begin{aligned}
w_{j,n}^{(m)} | (w_{j,n-1}^{(m)} = l) &\sim \text{Discrete}(\mathbf{T}_{l,:}^{(C_j^{(m)})}), & n &= 2, \ldots, N_j^{(m)} \\
C_j^{(m)} &\sim \text{Discrete}(\boldsymbol{\theta}^{(m)}), & j &= 1, \ldots, J_m, \\
\theta_k^{(m)} &= \frac{Z_k^{(m)}}{\sum_{k'=1}^{\infty} Z_{k'}^{(m)}}, & m &= 1, \ldots, M, \\
Z_k^{(m)} &\sim \text{Gamma}(\beta p_k, 1), & k &= 1, \ldots . \infty, \\
\mathbf{T}_{s,:}^{(k)} &\sim \text{Dirichlet}(\gamma_s \mathbf{1}), & s &= 1, \ldots, S, \\
p_k &= V_k \prod_{j=1}^{k-1} (1 - V_j), \\
V_k &\sim \text{Beta}(1, \alpha).
\end{aligned}
$$

To complete the generative formulation, we assign a Gamma$(1, c)$ prior for the concentration parameters α and β.

2.1. Variational Bayesian inference

The posterior distribution of the model is analytically intractable to compute. To work around this problem both sampling and variational Bayesian (VB) algorithms have been presented for mixed membership modelling via HDP (Bryant & Sudderth, 2012; Hoffman et al., 2013; Liang et al., 2007; Teh et al., 2006, 2007; Wang et al., 2011;

Wang & Blei, 2012). In this work, we use VB, since our focus of inference is not only in density estimation but also in interpreting the model, in particular, the posterior distribution of the strategies $\mathbf{T}^{(k)}$. For VB we get a closed form approximation, whereas for sampling based approaches simple averaging of posterior samples is meaningless. More specifically, we search for a truncated (that is, finite) approximation to the posterior distribution of the model building on the work by Blei & Jordan (2006); Liang et al. (2007) and Paisley et al. (2012).

Variational algorithms search for a factorised distribution over the unobserved variables Θ (here, Θ contains the unshaded nodes appearing in Figure 1) $q(\Theta)$ that minimises Kullback-Leibler divergence between the $q(\Theta)$ and the true posterior distribution. Alternatively, the algorithm maximises a lower bound of the marginal likelihood of the model,

$$\ln p(\mathcal{D}) \geq \mathcal{L} = \mathbb{E}[\ln p(\mathcal{D}, \Theta)] - \mathbb{E}[p(\Theta) \ln p(\Theta)],$$

where expectations, abbreviated as $\mathbb{E}[\cdot]$, are taken with respect to the $q(\Theta)$.

We assume the truncated factorised posterior approximation is

$$q(\Theta) = \prod_{k=1}^{K} \prod_{s=1}^{S} \prod_{m=1}^{M} \prod_{j=1}^{J_m} q(\mathbf{T}_{s,:}^{(k)}) q(Z_k^{(m)}) q(C_j^{(m)}) q(\alpha) q(\beta) q(V_k).$$

The variational algorithm proceeds by updating the parameters of each $q(\cdot)$ until convergence. Truncation of the approximation is obtained by setting $V_K = 1$ for a truncation level K. We use point (delta) distributions for α, β and V_k to simplify computations. In practice, use of point distributions equals finding point estimates that maximise the lower bound of the marginal likelihood $\mathcal{L}$. Nontrivial distributions are used for the strategies $q(\mathbf{T}_{s,:}^{k})$, the strategy indicators $q(C_j^{(m)})$ and the unnormalised strategy proportions $q(Z_k^{(m)})$. Note that the strategy probabilities (preferences) p_k and the normalised strategy proportions $\boldsymbol{\theta}^{(m)}$ for $k = 1, \ldots, K$ and $m = 1, \ldots, M$ are not included in the factorisation since they are definitions used to ease model description.

We select the following distributions

$$\begin{aligned} q(C_j^{(m)}) &= \text{Discrete}(C_j^{(m)} | \boldsymbol{\phi}_j^{(m)}), \\ q(Z_k^{(m)}) &= \text{Gamma}(Z_k^{(m)} | a_k^{(m)}, b_k^{(m)}), \\ q(\mathbf{T}_{s,:}^{(k)}) &= \text{Dirichlet}(\mathbf{T}_{s,:}^{(k)} | \boldsymbol{\eta}_s^{(k)}). \end{aligned}$$

The parameters of the $q(\mathbf{T}_{s,:}^{k})$, $q(C_j^{(m)})$ and $q(Z_k^{(m)})$ are

$$\begin{aligned}
\ln \phi_{j,k}^{(m)} &\propto \mathbb{E}[\ln Z_k^{(m)}] + \sum_{i,l} \mathbb{E}[\ln T_{i,l}^{(k)}] \#[\mathbf{w}_j^{(m)} : i \to l], \\
a_k^{(m)} &= \beta p_k + \sum_{j=1}^{J_m} \phi_{j,k}^{(m)}, \\
b_k^{(m)} &= 1 + \frac{J_m}{\sum_{k=1}^{T} \mathbb{E}[Z_k^{(m)}]}, \\
\eta_{i,l}^{(k)} &= \gamma_i + \sum_{m,j} \phi_{j,k}^{(m)} \#[\mathbf{w}_j^{(m)} : i \to l],
\end{aligned}$$

where

$$\begin{aligned}
\mathbb{E}[\ln T_{i,l}^{(k)}] &= \psi(\eta_{i,l}^{(k)}) - \psi(\sum_{l=1}^{S} \eta_{i,l}^{(k)}), \\
\mathbb{E}[\ln Z_k^{(m)}] &= \psi(a_k^{(m)}) - \ln b_k^{(m)}, \\
\mathbb{E}[Z_k^{(m)}] &= \frac{a_k^{(m)}}{b_k^{(m)}},
\end{aligned}$$

and $\#[\mathbf{w}_j^{(m)} : i \to l]$ counts how many times $\mathbf{w}_{j,n-1}^{(m)} = i$ precedes $\mathbf{w}_{j,n}^{(m)} = l$ in $\mathbf{w}_j^{(m)}$ and $\psi(\cdot)$ denotes the digamma function. We note that the inference algorithm scales well with respect to the dimensionality of the state space S, since the algorithm requires computing sums of variables over the state space, which may become an issue for extremely large state spaces. For later convenience, we note that

$$\begin{aligned}
\mathbb{E}[T_{i,l}^{(k)}] &\propto \exp(\psi(\eta_{i,l}^{(k)})), \\
\mathbb{E}[\theta_k^{(m)}] &\propto \mathbb{E}[Z_k^{(m)}].
\end{aligned}$$

We update V_k, where $k = 1, \ldots, K-1$, by maximising the $\mathcal{L}$ using gradient based optimisation. Retaining all the variables that depend on V_k, the corresponding cost function is

$$\mathcal{L}_{\mathrm{p}} = \beta \sum_k \sum_m \mathbb{E}[\ln Z_k^{(m)}] p_k - M \ln \Gamma(\beta p_k) + \ln p(V_k|\alpha),$$

where we represent V_k via p_k and $\Gamma(\cdot)$ denotes the gamma function. In a similar spirit, we obtain the cost function for the concentration parameter β by replacing the last term with $\ln p(\beta)$. Since closed form solutions for V_k and β may not be computed analytically, we use general purpose (box-constrained) L-BFGS optimisation algorithm (Byrd et al., 1995). The required gradients are straightforward to compute. The cost function with respect to the concentration parameter α is $\mathcal{L}_\alpha = \ln p(\mathbf{V}, \alpha)$ and the α can be updated in closed form.

3. Alternative model assumptions

The simplicial mixtures of Markov chains model proposed by Girolami & Kabán (2004) is related to our model and can be applied to model trace logging data. Following our notation, the model assumes strategy indicator variables are *state-specific*. That is, we draw an indicator variable $C_{j,n}^{(m)}$, where $n = 1, \ldots, N_j^{(m)}$, for each state. This assumption may produce spurious combinations of strategies for a single session complicating interpretations. One is forced to look at specific combinations that may vary from user to user. Girolami & Kabán (2004) assumed a fixed number of strategies corresponding to a finite model formulation, drawing the user-specific strategy distributions from a Dirichlet distribution. For inference they derived a variational EM-algorithm, approximately marginalising strategy proportions and using point estimates for the $\mathbf{T}^{(k)}$. However, no consensus exists for determining the number of strategies of the finite model.

As a side result, we present a novel Bayesian nonparametric formulation and inference algorithm for SMMC sidestepping the selection of the number of strategies. The nonparametric SMMC model is written as

$$
\begin{aligned}
w_{j,n}^{(m)} | (w_{j,n-1}^{(m)} = j) &\sim \text{Discrete}(\mathbf{T}_{j,:}^{(C_{j,n}^{(m)})}), \\
C_{j,n}^{(m)} &\sim \text{Discrete}(\boldsymbol{\theta}^{(m)}).
\end{aligned}
$$

The distributions for the remaining variables is similar to our model.

A simple mixture model formulation assumes users follow the same strategy for all sessions. The model is essentially similar to a finite mixture of stationary first-order Markov models proposed by Cadez et al. (2003). This model is also related to our model but ignores inter-session variability. Instead, the model assumes for the mth user a single strategy indicator variable.

We present a nonparametric formulation for this mixture model, referred to as Dirichlet process mixture model (DPMM), that is written as

$$
\begin{aligned}
w_{j,n}^{(m)} | (w_{j,n-1}^{(m)} = j) &\sim \text{Discrete}(\mathbf{T}_{j,:}^{(C^{(m)})}), \\
C^{(m)} &\sim \text{Discrete}(\mathbf{p}),
\end{aligned}
$$

where $\mathbf{p}$ denotes the strategy preferences drawn from the stick-breaking process. Our model simplifies to a DP mixture model when assuming a single indicator variable $C^{(m)}$, where $m = 1, \ldots, M$, since the user-specific strategy preferences become redundant and may be removed. For this simple model we follow the variational algorithm presented by Blei & Jordan (2006).

Finally, we also describe a baseline model, global Markov chain (GMC), that assumes all the users follow the same strategy. Due to simplicity it is unable to capture user-specific aspects. The model is

$$
w_{j,n}^{(m)} | (w_{j,n-1}^{(m)} = j) \sim \text{Discrete}(\mathbf{T}_{j,:}).
$$

When a Dirichlet prior is adopted for the (rows of) $\mathbf{T}$, the model equals the hierarchical Dirichlet language model introduced by MacKay & Peto (1995). Instead of modelling discrete software trace logging data, the model was originally proposed to model text (sequences of words or individual letters). In the context of language modelling, Wallach (2006) presented a finite bigram topic model that is similar to the model previously presented by Girolami & Kabán (2004).

4. Results

We start by describing the experimental setting for assessing different modelling assumptions via predictive ability. We compare our model (Section 2) against SMMC and DPMM using nonparametric formulations and VB for inference (see Section 3). For all the models we initialise the parameters of the factorised approximation randomly and set the truncation level to $K = 50$. The results are shown for $c = 1$, although we note that our model showed low statistical difference for a wide range of values between 10^{-3} and 1 corresponding to commonly applied values in the previous work. We assign a uniform Dirichlet distribution ($\gamma_s = 1$, where $s = 1, \dots, S$) for the strategies. We also include GMC in the comparison as a baseline model.

We have data from $M = 745$ users, $S = 11$ and a median value for session length over users is 56. We randomly sample one fifth of the data (users) as a test data set $\widetilde{\mathcal{D}}$ not used for learning the strategies. This process is repeated 50 times. We partition trace logs (sessions) for the test users into two halves with proportions 4/5 and 1/5, respectively. The first half, $\mathbf{w}_j^{(m)}$, where $m \in \widetilde{\mathcal{D}}$, is used to compute the user-specific variables (for example, $q(Z_k^{(m)})$ and $q(C_j^{(m)})$, for our model) given the global variables, whereas the other half, $\tilde{\mathbf{w}}_j^{(m)}$, is used to compute the marginal probability,

$$p(\tilde{\mathbf{w}}_j^{(m)}|\mathbf{w}_j^{(m)}, \mathcal{D}).$$

We compare models by computing predictive perplexities (lower is better),

$$\mathcal{P} = \exp\left\{ - \frac{\sum_{(m,j)\in\widetilde{\mathcal{D}}} \ln p(\tilde{\mathbf{w}}_j^{(m)}|\mathbf{w}_j^{(m)}, \mathcal{D})}{\sum_{(m,j)\in\widetilde{\mathcal{D}}} \sum_{i,l} \#[\tilde{\mathbf{w}}_j^{(m)} : i \to l]} \right\}.$$

To approximate the analytically intractable marginal probability, we select a maximum *a posteriori* strategy and use $\mathbb{E}[\mathbf{T}^{(\arg\max \boldsymbol{\phi}_j^{(m)})}]$ for computing a quantity similar to that in Equation 1. We adopt this strategy also for DPMM, whereas for SMMC we follow the approach adopted by Girolami & Kabán (2004), although, using the expectations of the corresponding variational distributions.

Table 1. Mean predictive perplexities as well as standard deviations computed over 50 different partitions of the data set $\mathcal{D}$ to train and test data sets.

Model	$\mathcal{P}$
Our model	3.16 ± 0.26
SMMC	3.19 ± 0.42
DPMM	3.28 ± 0.47
GMC	3.55 ± 0.61

Table 1 shows the predictive perplexities. We see that the baseline model using a single strategy is not flexible enough to capture intriguing complexity of the data and results in poor predictions. The DP mixture model improves performance supporting a prior assumption of a heterogenous population. Our model and SMMC further improve significantly over the DPMM suggesting that users tend to apply combinations of strategies (paired one-sided Wilcoxon test; $p < 0.003$). We see that our model is as (or slightly more) powerful than the less constrained and more expressive SMMC model. Although the predictive perplexities are useful for comparing predictive ability of the models, they may not give a good measure for Bayesian model selection. Accordingly, we carried out a cross-model selection between our model and SMMC. We used the corresponding lower bounds of model evidence to compute approximate Bayes factors. The results show decisive support for our model over SMMC (mean difference for the lower bound is 10^4). Considering the number of variables in the models the large difference is intuitive; SMMC assigns a separate strategy indicator variable for each transition, whereas our model assigns a separate indicator for each session.

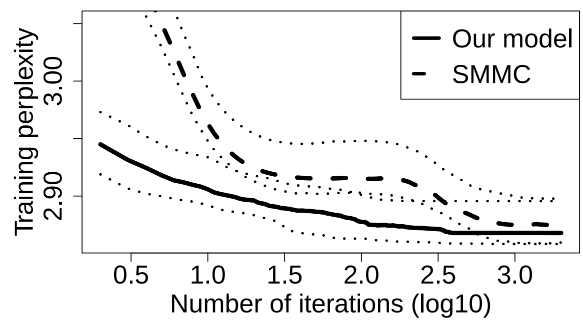

Figure 2. Training perplexity versus iteration count. The curves represent 50% quantiles, dotted lines around them correspond to 25% and 75% quantiles, respectively, computed over the 50 runs.

We found that our model equipped with the crucial assumption, that each user follows same strategy during one session, does not only lead to good predictive performance but also simplifies inference leading to faster and more stable posterior computations than SMMC. Figure 2 shows perplexities computed on training data for each iteration of the VB inference algorithms. We see that our model converges within few hundred iterations, but SMMC suffers from a plateau reaching similar performance to our model after roughly a thousand iterations. We note that DPMM has a similar learning curve to our model, however, converging to a solution with worse perplexities.

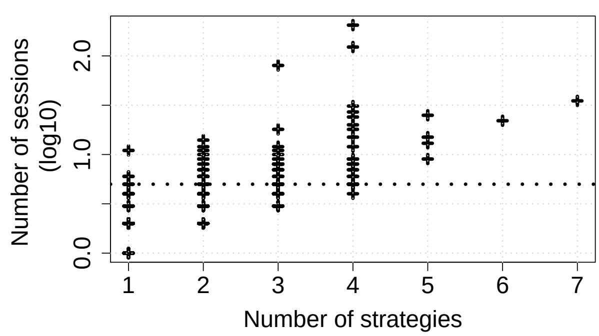

Figure 3. The number of strategies against the number of sessions for each user. To ease inspection of the figure, we added a dotted horizontal line denoting five sessions.

After quantitative model comparison and before interpreting the strategies, we inspect the strategy assignments $\arg\max \boldsymbol{\phi}_j^{(m)}$, which may be used to compute the number of strategies over the population and combinations of strategies for each user. Our model inferred in total 23 strategies out of which four strategies were adopted by one

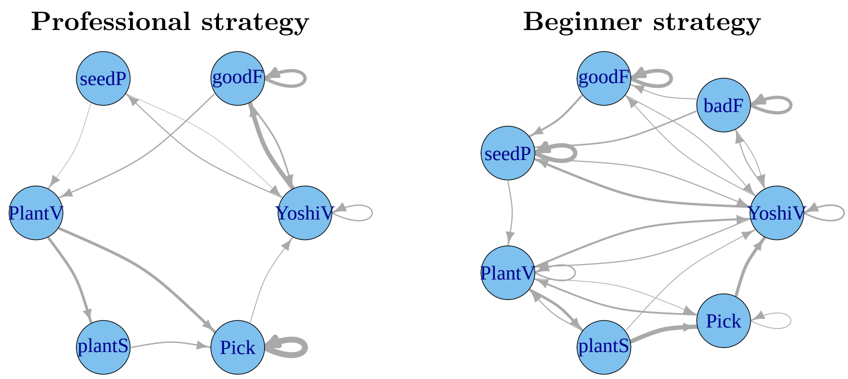

Figure 4. Visualisation of the strategies $\mathbb{E}[\mathbf{T}^{(k)}]$ inferred from the proposed model. See text for a more detailed interpretation of the strategies. Table 2 explains the state abbreviations.

user. Figure 3 shows the number of strategies against the number of sessions for each user. From the figure, we see that some users, who posses a small number of sessions, follow a single strategy. (This is is evident for users with only one session). However, more experienced users, indicated by a larger number of sessions, tend to apply a combination of two to five strategies. Despite the small number of strategies for each user, the unique combinations (preferences) of strategies are highly dispersed. The users disperse altogether to 102 different (small) subpopulations. The number of users per subpopulation varies between one and 25. This finding reflects inherent variability of the trace logging data for the heterogenous and uncontrolled population.

Figure 4 highlights the main findings of our model visualising two strategies $\mathbb{E}[\mathbf{T}^{(k)}]$. Table 2 explains abbreviations for the states used in the figure. In the visualisation, nodes represent states and directed edges indicate a transition between states. Loops indicate repeating the same state. Width of the edges are proportional to transition probabilities. We threshold the probabilities at a moderate value of 0.05. Consequently, we omit from the visualisation states that are either impossible to visit or have no state(s) to transit to. This simplification is well justified by our assumption of user following the same strategy during one session but does not apply to SMMC due to potential switches of strategies between consecutive states hindering interpretation. Based on our knowledge of the game, we label the strategies according to our interpretation, detailed in the following. We discovered neither SMMC or DPMM were able to find as intuitive and interpretable strategies as our model. The strategies inferred by SMMC represent parts of playing strategies reflecting the switching property of the model. On the other hand, the strategies of DPMM fail to capture interpretable and meaningful patterns.

Table 2. Abbreviations and their explanations for the state symbols used in the visualisation of the strategies (see Figure 4).

Abbreviation	Explanation
YoshiV	view a yoshi
goodF	give good (correct) feed
badF	give bad (wrong) feed
seedP	pick fruit seeds
PlantV	view a plantation
plantS	plant seeds
Pick	pick a fruit

Figure 4 (**left**) presents a professional and sensible strategy that may be adopted by more experienced users (experts). After viewing a yoshi (**YoshiV**) there are two states to transit to. The first transition means feeding the yoshi correctly (**goodF**). Note the strong self link (loop) meaning that the user tends to give multiple correct fruits in a sequence. The second transition means picking fruit seeds (**seedP**). After feeding a yoshi with correct fruits or picking seeds, the user may either view (another) yoshi or view a plantation (**plantV**). In the latter case, after viewing the plantation, the user picks (multiple) fruits in a basket from the plantation or plants the seeds (**plantS**) before picking the fruits. Then, the cycle repeats by viewing yoshi again.

Figure 4 (**right**) presents an explorative strategy that may be adopted by less experienced users (beginners). Here, a user is likely to transit back and forth between two states, learning which fruits yoshis request and which fruits are available on the plantations. After viewing a yoshi the user may give wrong fruits (**badF**) in addition to correct fruits. It is interesting to note the transition from bad feed to good feed. We anticipate this behaviour corresponds to user learning to give correct fruits. After feeding the yoshi, the user proceeds by picking seeds, viewing a plantation for planting them and picking fruits, in general. However, we note that there is no common efficient single path between these state transitions as in the expert strategy. Instead, the user prefers to either view yoshi potentially checking which fruits were needed or view different plantations to see which fruits they provide. In addition, there is a lower probability for picking sequentially multiple fruits in a basket before viewing a yoshi. This is not efficient. Altogether, these findings and differences to the expert strategy indicate this strategy is investigative.

Based on the interpretations, we checked whether the proposed model supports our labelings. We expect that the number of sessions for a user indicates an overall grade for his/her playing expertise. Accordingly, we inspected the strategy assignments for the two strategies as well as the number of game sessions for the users. We found the first (professional) strategy is frequently adopted by users, who have played a large number of sessions, whereas the second (beginner) strategy is commonly adopted by users, who have played a small number of sessions. Indeed the model supports our expectation. The more experienced users have adopted a reasonable and efficient strategy, whereas the less experienced users, who may not be familiar with the game, follow an explorative strategy, learning the idea of the game and getting familiar with different yoshis and their diets as well as range of different fruits and plantations providing them.

5. Conclusions

In this work, we explored nonparametric Bayes approaches to modelling several trace logs of a heterogenous and diverse population of mobile gamers for inferring how users interact with the game. We presented a novel model formulation targeting a particular mobile application and compared the model against competing alternative models in the nonparametric Bayes framework. We demonstrated that our model retains not only predictive ability to a same degree as a more flexible model formulation of mixed membership modelling but also simplified interpretability of a standard mixture model. We inspected the model and found interpretable playing strategies; one is adopted by more experienced users capturing the essence of the game and the other provides insight into how beginners play the game. These findings support our proof of concept is feasible and calls for future modelling developments taking into account dynamics and hierarchies of strategies. Considering the evolution of users' expertise over time may improve predictive ability and provide a more detailed understanding of the strategies. While we have focused in this paper to first order Markov chains, the model formulation and inference algorithm require minor changes to account for arbitrary model order.

Acknowledgement

This research is supported by EPSRC Programme Grant *A Population Approach to Ubicom System Design* (EP/J007617/1).

References

Andrei, O, Calder, M, Higgs, M & Girolami, M (2014), 'Probabilistic model checking of DTMC models of user activity patterns,' in *QEST*.

Blei, D & Jordan, MI (2006), 'Variational inference for Dirichlet process mixtures,' *Bayesian Analysis*, **1**(1), pp. 121–143.

Bryant, M & Sudderth, EB (2012), 'Truly nonparametric online variational inference for hierarchical Dirichlet processes,' in *Advances in Neural Information Processing Systems*.

Byrd, RH, Lu, P, Nocedal, J & Zhu, C (1995), 'A limited memory algorithm for bound constrained optimization,' *SIAM Journal on Scientific Computing*, **16**(5), pp. 1190–1208.

Cadez, I, Heckerman, D, Meek, C, Smyth, P & White, S (2003), 'Model-based clustering and visualization of navigation patterns on a web site,' *Data Mining and Knowledge Discovery*, **7**(4), pp. 399–424.

Ferguson, TS (1973), 'A Bayesian analysis of some nonparametric problems,' *The Annals of Statistics*, pp. 209–230.

Girolami, M & Kabán, A (2004), 'Simplicial mixtures of markov chains: Distributed modelling of dynamic user profiles,' in *Advances in Neural Information Processing Systems*.

Higgs, M, Morrison, A, Girolami, M & Chalmers, M (2013), 'Analysing user behaviour through dynamic population models,' in *CHI Extended Abstracts on Human Factors in Computing Systems*.

Hoffman, MD, Blei, DM, Wang, C & Paisley, J (2013), 'Stochastic variational inference,' *Journal of Machine Learning Research*, **14**(1), pp. 1303–1347.

Ishwaran, H & Zarepour, M (2002), 'Exact and approximate sum representations for the Dirichlet process,' *Canadian Journal of Statistics*, **30**(2), pp. 269–283.

Liang, P, Petrov, S, Jordan, MI & Klein, D (2007), 'The infinite PCFG using hierarchical Dirichlet processes,' in *EMNLP*.

MacKay, D & Peto, L (1995), 'A hierarchical Dirichlet language model,' *Natural Language Engineering*, **1**(03), pp. 289–308.

McMillan, D, Morrison, A, Brown, O, Hall, M & Chalmers, M (2010), 'Further into the wild: Running worldwide trials of mobile systems,' in *Pervasive Computing*.

Paisley, J, Wang, C & Blei, D (2012), 'The discrete infinite logistic normal distribution,' *Bayesian Analysis*, **7**(4), pp. 997–1034.

Sethuraman, J (1994), 'A constructive definition of Dirichlet priors,' *Statistica Sinica*, **4**, pp. 639–650.

Teh, YW, Jordan, MI, Beal, MJ & Blei, D (2006), 'Hierarchical Dirichlet processes,' *Journal of the American Statistical Association*, **101**(476), pp. 1566–1581.

Teh, YW, Kurihara, K & Welling, M (2007), 'Collapsed variational inference for HDP,' in *Advances in Neural Information Processing Systems*.

Wallach, HM (2006), 'Topic modeling: Beyond bag-of-words,' in *International Conference on Machine learning*.

Wang, C & Blei, DM (2012), 'Truncation-free online variational inference for Bayesian nonparametric models,' in *Advances in Neural Information Processing Systems*.

Wang, C, Paisley, JW & Blei, DM (2011), 'Online variational inference for the hierarchical Dirichlet process,' in *International Conference on Artificial Intelligence and Statistics*.