



**HAL**  
open science

## Markov Decision Process for MOOC Users Behavioral Inference

Firas Jarboui, Célya Gruson-Daniel, Alain Durmus, Vincent Rocchisani, Sophie-Helene Goulet Ebongue, Anneliese Depoux, Wilfried Kirschenmann, Vianney Perchet, Sophie-Helene Goulet Ebongue

► **To cite this version:**

Firas Jarboui, Célya Gruson-Daniel, Alain Durmus, Vincent Rocchisani, Sophie-Helene Goulet Ebongue, et al.. Markov Decision Process for MOOC Users Behavioral Inference. EMOOCs 2019: Digital Education: At the MOOC Crossroads Where the Interests of Academia and Business Converge, pp.70-80, 2019, 10.1007/978-3-030-19875-6\_9 . hal-03089780

**HAL Id: hal-03089780**

**<https://hal.science/hal-03089780>**

Submitted on 16 Jan 2024

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

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

# Markov Decision Process for MOOC users behavioral inference\*

Firas JARBOUI<sup>1,2</sup>, Célya GRUSON-DANIEL<sup>3,5</sup>, Pierre Chanial<sup>1</sup>, Alain DURMUS<sup>2</sup>, Vincent ROCCHISANI<sup>1</sup>, Sophie-Helene GOULET EBONGUE<sup>3</sup>, Anneliese DEPOUX<sup>3,4</sup>, Wilfried KIRSCHENMANN<sup>1</sup>, and Vianney PERCHET<sup>2</sup>

<sup>1</sup> ANEO, Boulogne Billancourt, France

<sup>2</sup> CMLA, École normale supérieur Paris Saclay, Université Paris Saclay, France

<sup>3</sup> Centre Virchow-Villermé for Public Health Paris-Berlin, Université Sorbonne Paris-Cité, France,

<sup>4</sup> GRIPIC - EA 1498, , Sorbonne Université, France,

<sup>5</sup> DRISS (Digital Research in Science & Society)

**Abstract.** Studies on massive open online courses (MOOCs) users discuss the existence of typical profiles and their impact on the learning process of the students. However defining the typical behaviors as well as classifying the users accordingly is a difficult task. In this paper we suggest two methods to model MOOC users behaviour given their log data. We mold their behavior into a Markov Decision Process framework. We associate the user's intentions with the MDP reward and argue that this allows us to classify them.

**Keywords:** User Behaviour Studies · Learning Analytics · Markov Decision Process · Inverse Reinforcement Learning.

## 1 Introduction

Finding an efficient way to identify behavioural patterns of MOOC users community is a recurring issue in e-learning. However, as detailed in the review of Romero and Ventura [1] on educational data science research, the way this problem was studied was either testing correlations given conjectures or trying to identify communities of look-a-likes.

The main approaches study aggregates of data generated by users in order to identify their respective behaviors with respect to some typology of the students. For instance, Ramesh A. et al. [2] distinguish learner behavior according to their engagement into *active*, *passive* and *disengaged* learners. They predict their behavior based on a probabilistic soft logic model taking into account users features relating to their engagement on the Mooc. Corin L. et al. [3] describes the behaviour of users through flow diagrams of the state transitions through comparing different sets of behaviours with graphical models. Cheng Y. and

---

\* Supported by ANR-15-IDFN-0003-04.

Gautam B. [4] predicted users drop-out through temporal granularities in features suspected to influence the drop-out of the users. Finally, Christopher B. et al. [5] cover the prediction of the users achievement through their levels of activity in the course. Unfortunately, given a different typology of users profile, we can not easily transpose these approaches because the features used to characterize the learner are selected with respect to the definition of the classes. For example, performance related typology can be matched to users' quizzes success rate, the drop-out rate oriented classes can be tracked through the times series of connection history. However the task becomes more difficult when the definition of classes can not be reduced to a simple quantifiable measurement, for example the user intentions, goals and motivations. Chase G. and Cheng Z. [6] aimed for a generalized method by modelling the user behaviour as a two layer hidden Markov model. They used log data to construct their characterization of the users. They cluster the transition probabilities to define the hidden states, and they compare the transitions between these hidden states for the high and low performing students. However, if we only observe the transitions probabilities, identifying an interpretation of the associated behavior is a complex task.

We aim to define general models to study any kind of user behavior without loosing the interpretability of results. We consider two different main models. In the first one, we assume that each user can adopt his own different policy. Each user has a reward function over the MOOC and tries to optimize it. We can cluster these rewards into a finite number of classes that represent the behaviors explaining the observations. In the second one, we assume that there are a limited number of rewards the users can optimize. Each reward translates into a typical behavior and the users are switching between them along the MOOC.

## 2 Mathematical Preliminary

### 2.1 Markov Decision Process

Consider a Markov Decision Process (MDP)  $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \nu, \mathcal{X}\}$  where  $\mathcal{S} = \{1, 2, \dots, N_s\}$  is a finite state space,  $\mathcal{A} = \{1, 2, \dots, N_a\}$  is a finite action space,  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the probability distribution of state transition such as  $\mathcal{P}(s, a, s') = P(s_{t+1} = s' | s_t = s, a_t = a)$  is a probability of going to state  $s'$  under action  $a$  from state  $s$ ,  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  is the reward function,  $\nu \in ]0, 1[$  is a discount factor and  $\mathcal{X}$  is an initial distribution over the states i.e.  $P(s_0 = s) = \mathcal{X}(s)$ . Any transition matrix compatible with the MDP on  $\mathcal{S} \times \mathcal{A}$  is referred to as a policy, and we denote  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  such policy. An agent following the policy  $\pi$  would take at time  $t$  the action  $a_t = a$  at state  $s_t = s$  with probability  $P(a_t = a | s_t = s) = \pi(s, a)$ . The value of the policy  $\pi$  is  $V^\pi = \sum_s \mathcal{X}(s) V^\pi(s)$  where  $V^\pi(s) = \mathbf{E}^\pi[\sum_t \nu^t \mathcal{R}(s_t, a_t) | s_0 = s]$  is the state value function. Similarly we denote the state-action value function by  $Q^\pi(s, a)$  defined by  $Q^\pi(s, a) = \mathbf{E}^\pi[\sum_t \nu^t \mathcal{R}(s_t, a_t) | s_0 = s, a_0 = a]$ . We denote by  $\pi^*$  the optimal policy maximizing the expected discounted reward given any starting state, and

by  $V^*$  and  $Q^*$  the corresponding values defined as [7]:

$$\begin{aligned}\pi^* &= \operatorname{argmax}_{\pi} V^{\pi}(s) = \operatorname{argmax}_{\pi} Q^{\pi}(s, a) \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A} \\ V^*(s) &= V^{\pi^*}(s) \quad \forall s \in \mathcal{S} \\ Q^*(s, a) &= Q^{\pi^*}(s, a) \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A}.\end{aligned}$$

Given an MDP  $\mathcal{M}$ , the optimal policy  $\pi^*$  can be computed using well-known methods such as Value iteration or Policy iteration [7]. We use a Value iteration approach and we refer to this procedure as  $Q^* = \text{MDP}(\mathcal{M})$

**Reward parametrized MDP:** In this paper we consider the case of MDPs with linearly parametrized rewards  $\mathcal{M}_{\theta} = \{\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}_{\theta}, \phi, \nu, \mathcal{X}\}$  i.e. such that:

$$\forall s \in \mathcal{S} \quad \forall a \in \mathcal{A} \quad \mathcal{R}_{\theta}(s, a) = \theta \cdot \phi(s, a)^T,$$

where  $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbf{R}^N$  is a feature map from the state-action space to a real valued  $N$  dimensional space and  $\theta \in \mathbf{R}^N$  is an  $N$  dimensional real weights.

For a given  $\theta$ , we will denote the optimal policy by  $\pi_{\theta}^*$ , and the corresponding optimal value functions by  $V_{\theta}^*$  and  $Q_{\theta}^*$ .

## 2.2 Inverse Reinforcement Learning

Let be  $\mathcal{M}_{\theta}$  a reward parametrized MDP whose parameter  $\theta$  is unknown. We denote by  $\mathcal{W} = \{\mathcal{S}, \mathcal{A}, \mathcal{P}, \phi, \nu, \mathcal{X}\}$ , and  $\mathcal{D}_M = \{(y_{0:T_i}^i)_{i=1}^M\}$  the behavioral data where  $y_{0:T_i}^i = \{y_t^i = (s_t^i, a_t^i)_{i=1}^{T_i}\}$  are the  $i^{th}$  individual following an unknown policy  $\pi_{\theta}^*$ . The goal of the Inverse Reinforcement Learning (IRL) problem is to identify parameters  $\hat{\theta}$  such that  $\pi_{\hat{\theta}}^*$  are as likely as  $\pi_{\theta}^*$  to generate the observations  $\mathcal{D}_M$ . The IRL problem is ill-posed [8] as there exists infinitely many reward parameters that yield  $\pi_{\theta}^*$  as an optimal policy. For example with  $\theta = 0$ , any policy is optimal for any IRL problem.

To circumvent this issue, many approaches have been proposed to define preferences over the reward space. These approaches can be broadly divided in two settings: Optimization IRL and Bayesian IRL. Optimization oriented approaches define objective function that encode such preferences [8,9,10]. Bayesian approaches formulate the reward preferences in the form of a prior distribution over the rewards and define behavior compatibility as a likelihood function [11,12,13]. We will follow the latter setting [13]. The model we consider assumes that that agents are not following an optimal policy  $\pi_{\theta}^*$  but rather an aproximal one. more precisely, we assume that:

$$\begin{aligned}\tilde{\pi}_{\theta}^{\eta}(s, a) &= \frac{\exp(\eta Q_{\theta}^*(s, a))}{\sum_{a_i} \exp(\eta Q_{\theta}^*(s, a_i))} \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A} \\ P(a_t = a | s_t = s, \mathcal{M}_{\theta}) &= \tilde{\pi}_{\theta}^{\eta}(s, a) \quad \forall s \in \mathcal{S} \quad \forall a \in \mathcal{A} \\ P(s_{t+1} = s' | s_t = a, a_t = a, \mathcal{M}_{\theta}) &= \mathcal{P}(s, a, s') \quad \forall s, s' \in \mathcal{S} \quad \forall a \in \mathcal{A}\end{aligned}$$

Therefore, under this model, the likelihood is given by:

$$P(\mathcal{D}_M | \mathcal{M}_{\theta}) = \prod_{i=1}^M \prod_{t=1}^{T_i} \tilde{\pi}_{\theta}^{\eta}(s, a) \times \prod_{i=1}^M \prod_{t=1}^{T_i} P(s_{t+1} | s_t, a_t, \mathcal{M}_{\theta})$$

We have assumed that  $(a_t, s_t)_{t \geq 0}$  is still Markovian, with transition probabilities given by  $\mathcal{P}$  which does not depend on  $\theta$ . Thus,  $\prod_{i=1}^M \prod_{t=1}^{T_i} P(s_{t+1}|s_t, a_t, \mathcal{M}_\theta)$  can be treated as a multiplicative constant with respect to  $\theta$ . We define:

$$\mathcal{L}(\theta; (\mathcal{D}_M, \mathcal{W})) = \prod_{i=1}^M \prod_{t=1}^{T_i} \frac{\exp(\eta Q_\theta^*(s_t^i, a_t^i))}{\sum_a \exp(\eta Q_\theta^*(s_t^i, a))},$$

where  $\eta$  can be interpreted as a confidence parameter. The bigger it gets, the closer are the policies  $\tilde{\pi}_\theta^\eta$  and  $\pi_\theta^*$ , as  $\lim_{\eta \rightarrow \infty} \tilde{\pi}_\theta^\eta(s, a) = \pi_\theta^*$ . The posterior distribution is given by Bayes Theorem, where we choose  $\theta \rightarrow P(\theta)$  to be a uniform distribution over a subset of the parameter space.

$$P(\theta|\mathcal{D}_M) \propto \mathcal{L}(\theta; (\mathcal{D}_M, \mathcal{W}))P(\theta)$$

We use approximate samples from the distribution  $P(\theta|\mathcal{D}_M)$  to compute the a posteriori mean or median which are optimal under the square or linear loss function respectively [13,15]. Iterating Algorithm 1 generates the samples.

---

**Algorithm 1** SampleTheta
 

---

- 1: **procedure**  $\theta = \text{SampleTheta}(\theta^0, \mathcal{W}, P, \sigma, \eta, \mathcal{D}_M)$
  - 2:   sample  $\epsilon \sim \mathcal{N}(0, 1)$  and set  $\tilde{\theta} = \theta^0 + \sigma\epsilon$
  - 3:    $\tilde{Q}_\theta^* = \text{MDP}(\mathcal{M}_{\tilde{\theta}})$
  - 4:   Set  $\theta = \tilde{\theta}$  with probability:
  - 5:          $\min(1, \frac{\mathcal{L}(\mathcal{D}_M, \mathcal{M}_{\tilde{\theta}}P(\tilde{\theta}))}{\mathcal{L}(\mathcal{D}_M, \mathcal{M}_\theta P(\theta))})$
  - 6:   else set  $\theta = \theta^0$
  - 7:   Return  $\theta$
- 

### 2.3 Switched Markov Decision Process sMDP

Inspired from switched Linear Dynamical Systems [14], switched Markov Decision Process allow us to simplify complex phenomena into transitions among a set of simpler models. For example, non-linear behaviors such as an individual's movement in a crowd, can be viewed as an array of linear behavior among which the person is temporally switching [11].

Let  $(\mathcal{M}_{\theta_i})_{i=1}^L$  be a set of MDP models with corresponding parameters  $(\theta_i)_{i=1}^L$  and policies  $(\tilde{\pi}_{\theta_i}^\eta)_{i=1}^L$ . Switching between these models is governed by a discrete Markov process with transitions  $\zeta$ . We denote  $z_t$  as the latent mode of the system at time  $t$ , thus, it is sampled according to  $\zeta_{z_{t-1}, \cdot}$ . We also denote by  $y_t = (s_t, a_t)$  the observations which obey to a Markov decision process model.

$$z_t | \{\zeta_i\}_{i=1}^L, z_{t-1} \sim \zeta_{z_{t-1}, \cdot}$$

$$y_t | y_{t-1}, z_t, \{\theta_i\}_{i=1}^L \sim \pi_{\theta_{z_t}}^\eta(y_t) \mathcal{P}(s_{t-1}, a_{t-1}, s_t)$$

The hidden modes and the MDP associated with each one of them provide a Hidden Markov Model (HMM) structure leading to repeating simple behaviors. A common approach to solve HMM model is the forward backward method

developed by Andrew Viterbi[17]. We call Viterbi the procedure that evaluates the latent modes  $\hat{z}_{1:T}$  and the transition probabilities  $\mathcal{F}_{ij} = P(z_t = i, z_{t+1} = j)$  and we denote it by  $[z_{1:T}, \mathcal{F}] = \text{Viterbi}(y_{1:T}, \zeta_{kk'}, \pi_{\theta_i}^\eta)$ . We denote by  $\mathcal{M}^S = (\mathcal{W}, L, (\theta_i)_{i=1}^L, \zeta, \mathcal{X}_m)$  the sMDP model where  $\mathcal{X}_m$  is the initial mode distribution. As an application we will tackle the case where the parameters  $((\theta_i)_{i=1}^L, \zeta)$  are unknown and will be learned from the data.

## 2.4 Label Propagation

Let  $(x_i, y_i)_{i=1}^l$  be a set of labeled data, and  $(x_i, y_i)_{i=l+1}^{l+u}$  be a set of unlabeled data, i.e.  $(y_i)_{i=l+1}^{l+u}$  are unobserved. Where  $x_i \in \mathbf{R}^D$  and  $y_i \in \mathcal{C}$  for  $i = 1, \dots, l+u$  and where  $\mathcal{C}$  is a finite set. We define  $X = \{x_1, x_2, \dots, x_{l+u}\}$ ,  $Y_L = \{y_1, \dots, y_l\}$  and  $Y_U = \{y_{l+1}, \dots, y_{l+u}\}$ . The problem is to estimate  $Y_U$  given  $X$  and  $Y_L$ . we also denote by  $Y$  the matrix of label probability where  $Y_{ic} = P(y_i = c)$ . We want to find a matrix  $Y$  that satisfies the following:

$$\begin{cases} Y_i. = \frac{(TY)_i.}{\|(TY)_i.\|_2} & \forall i > l+1 \\ Y_i. = \delta(y_i) & \forall i \leq l \end{cases}$$

Where  $T$  is the matrix of label transition probabilities through the set  $X$ .  $T_{ij}$ , the probability that  $x_i$  will inherit the label of  $x_j$ , is proportional to the distances between the two points. Algorithm 2 [16] solves for  $Y$ .

$$T_{ij} = P(y_j \rightarrow y_i) = \frac{w_{ij}}{\sum_{k=1}^{l+u} w_{kj}}$$

and  $w_{ij} = \exp(-\frac{d_{ij}^2}{\sigma^2}) = \exp(-\frac{|x_i - x_j|_2^2}{\sigma^2})$

---

### Algorithm 2 Label Propagation

---

- 1: **procedure**  $Y = \text{LabelProp}(X, Y_L, \sigma)$
  - 2:     initialize  $Y$  with:  $Y_{ic} = \delta(c = y_i)$  if  $i \leq l$  and  $Y_{ic} = \frac{1}{C}$  if  $i > l$
  - 3:     repeat until convergence of  $Y$ :
  - 4:          $Y = TY$
  - 5:         row-normalize  $Y$
  - 6:      $Y_{ic} = \delta(c = y_i)$  for  $i \leq l$
  - 7:     return  $Y$
- 

## 3 Behavior inference with IRL

We now suggest two ways to develop a classification for MOOC user behaviors. In the first one, **Static Behavior Clustering (SBC)**, we consider that each user follows a policy that optimizes his own reward function and that generates their log data. We use the reward parameter associated to each user as features and we propagate labels that experts define on a restricted set of users. In the second one, **Dynamic Behavior Clustering (DBC)**, we consider that there is a small number of behaviors a user can adopt. We model their behaviors with a sMDP, the log data is then simplified into a set of typical behavior

successions. We denote in the sequel by  $\mathcal{W}$  the MDP associated to the MOOC and by  $\mathcal{D}_M = \{(y_{0:T_i}^i)_{i=1}^M\}$  the collected observation from  $M$  users where  $y_t^i = (s_t^i, a_t^i)$  indicate user  $i$  being at state  $s_t^i$  at time  $t$  and taking action  $a_t^i$ . We denote by  $\mathcal{D}_M(i) = \{y_{0:T_i}^i\}$  the data associated to the  $i^{\text{th}}$  user.

### 3.1 Construction of the MOOC MDP

Given a MOOC, we first define some associated MDP parameters  $\mathcal{W}$ . The construction is straightforward: we define  $\mathcal{S}$  as the different pages a user can access along with a resting state (associated to logging out of the website),  $\mathcal{A}$  is associated to the different actions available to the user such as playing a video, clicking on a given link or answering a quiz.  $\mathcal{X}$  and  $\mathcal{P}$  are computed empirically given the data  $\mathcal{D}_M$ . The feature function  $\phi$  however gives some flexibility to our approach. If we do not have much knowledge about the behaviors we are trying to track, we can define  $\phi(s, a) = 1_{\mathcal{S} \times \mathcal{A}}(s, a)$  as the indicator function of each state-action combination. Unfortunately, this will become unhandy for higher dimensions as  $\theta \in \mathbf{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ . An expert can however define a set of features to which the set of state-actions can be mapped. The discount parameter  $\nu$  reflects the ability of the agents of long term planning. It should be learned along with other parameters as it might not be the same for each user. However, we will consider a shared parameter that we fix at 0.9 for the sake of simplicity.

### 3.2 Static Behavior Clustering

We assume the existence of a set of behavior classes  $\mathcal{C} = \{1, \dots, N\}$ . Let  $(c_i)_{i=1}^M$  be the classes associated to each user in the data. With the help of a human expert, using highly restrictive conditions, we identify  $l$  users classes. Without loss of generality, we assume that  $C_L = \{c_1, \dots, c_l\}$  is the set of known classes. Let  $\mathbf{P}^c$  be the class probability matrix where  $\mathbf{P}_{ij}^c = P(c_i = j)$ . We consider also that each user behaves according to the MDP  $\mathcal{M}_{\theta_i} = \{\mathcal{W}, \mathcal{R}_{\theta_i}\}$  such that:

$$\forall t \leq T_i \quad a_t^i \sim \pi_{\theta_i}^\eta(s_t^i, \cdot) \text{ and } s_{t+1}^i \sim \mathcal{P}(s_t^i, a_t^i, \cdot)$$

For user  $i$ , we propose to infer such parameters given  $\mathcal{D}_M(i)$ , which will allow us to infer  $(\theta_i)_{i=1}^M$ . The objective is then to identify  $\mathbf{P}^c$  given  $\Theta$  and  $C_L$ . In Algorithm 3 we suggest a method to solve this problem.

---

#### Algorithm 3 Static Behavior Clustering

---

- 1: **procedure**  $\mathbf{P}^c = \text{SBC}(\mathcal{D}_M, \mathcal{W}, C_L, P_\theta, \eta, \sigma_{MDP}, \sigma_{LP})$
  - 2:   initialize  $\Theta = \{\theta_i\}$  to void values
  - 3:   for  $i=1, \dots, M$ :
  - 4:      $\theta_i = \text{BIRL}(\mathcal{W}, P_\theta, \sigma, \eta, N_{\max}, \mathcal{D}_M(i))$
  - 5:    $\mathbf{P}^c = \text{LabelProp}(\Theta, C_L, \sigma_{LP})$
  - 6:   return  $\mathbf{P}^c$
-

### 3.3 Dynamic Behavior Clustering

We suggest here a simpler version of the model of the sticky Hierarchical Dirichlet Process for sMDP [11]. To simplify the problem we assume that the number of clusters  $L$  is given. We suppose that the transition distribution from the  $i^{th}$  mode are generated according to a Dirichlet distribution  $\text{Dir}(\alpha_L + \delta_i)$  where  $\alpha_L = \alpha \cdot 1_L$ ,  $\delta_i$  is the  $i^{th}$  vector of the canonical base of  $\mathbf{R}^L$ , and  $\alpha \in \mathbf{R}$ . This avoids jumping between modes as the probability of remaining in the same mode is higher than switching to a new one. The reward parameter associated to each mode is sampled from  $U$  the uniform distribution over some subset of the parameter space. The observations obey to an sMDP model defined with  $\mathcal{M}^S = (\mathcal{W}, L, \{\theta_i, \zeta_i\}_{i=1}^L, \mathcal{X}_m)$ . The full generative model is given below:

$$\begin{aligned} \zeta(i, \cdot) | \alpha &\sim \text{Dir}(\alpha_L + \delta_i) & \forall i = 1, \dots, L \\ \theta_i &\sim U & \forall i = 1, \dots, L \\ z_t^i | \{\zeta(i, \cdot)\}_{j=1}^L &\sim \zeta(z_{t-1}, \cdot) & \forall i = 1, \dots, M \text{ and } t = 1, \dots, T_i \\ y_t^i | y_{t-1}^i, z_t^i, \{\theta_j\}_{j=1}^L &\sim \pi_{\theta_{z_t^i}}^\eta(s_t^i, a_t^i) \mathcal{P}(s_{t-1}^i, a_{t-1}^i, s_t^i) & \forall i = 1, \dots, M \text{ and } t = 1, \dots, T_i \end{aligned}$$

The intuition behind this model is that each user can adopt at each time step one of the behavior modes (i.e. he behaves accordingly to one of the MDPs). To alleviate notations we denote in the following  $\Pi = \{\pi_{\theta_i}^\eta\}_{i=1}^L$ , and  $\mathcal{D}^k = \{y_t^i \mid \forall i \leq L \forall t \leq T_i \mid z_t^i = k\}$  the observations associated to the  $k^{th}$  behavior mode. We suggest a MCMC approach to solve this inference problem where each step looks as developed in Algorithm 4. We start by inferring the latent modes according to the previous parameters values. We define the new sample of the HMM parameters using the normalization of the frequencies probabilities  $\mathcal{F}$  in step 6. In step 7, we split the data set according to the  $z_t^i$  into  $\mathcal{D}^i$  and use Algorithm 1 to sample the new MDP parameters  $\theta_i$  and their policies.

---

#### Algorithm 4 Dynamic Behavior Clustering

---

- 1: **procedure**  $\Theta^n, \zeta^n, \Pi^n = \text{DBC}(\mathcal{D}_M, \mathcal{W}, P_\theta, \eta, \sigma, \Theta^{n-1}, \zeta^{n-1}, \Pi^{n-1}, \alpha)$
  - 2:     Set  $\mathcal{F} = [f_{ij}]_{i \in [1, L], j \in [1, L]}$  to zeros
  - 3:     for  $i=1, \dots, M$ :
  - 4:          $[z_{1:T_i}^i, \mathcal{F}^i] = \text{Viterbi}(y_{1:T_i}^i, \zeta^{n-1}, \Pi^{n-1})$
  - 5:          $\mathcal{F} = \mathcal{F} + \mathcal{F}^i$
  - 6:      $[\zeta^n] = \text{SampleHMMPParam}(\mathcal{D}_M, \mathcal{F}, \zeta^{n-1}, \alpha, \sigma)$
  - 7:      $[\Theta^n, \Pi^n] = \text{SampleMDPPParam}(\mathcal{D}_M, \Theta^{n-1}, z_{1:T_i}^i, \eta, \sigma)$
  - 8:     return  $[\Theta^n, \zeta^n, \Pi^n]$
- 

## 4 Experiment

The experiments were conducted on MOOCs published in the framework of the research project #MOOCLive under the leadership of the Centre Virchow-Villermé for Public Health. The project aimed to substantially improve the efficiency of the MOOCs through a deeper understanding of the participants and their behaviors. In the modelization of the MOOC's MDP, we used an indicator feature function over the state-action space  $\mathcal{S} \times \mathcal{A}$ .



#### 4.1 Static Behavior Clustering

As mentioned before, our objective is to introduce a procedure that can be applied irrespectively of the experts classification. We experimented with multiple behavior classes  $\mathcal{C}$ . Each time, we first defined  $\mathcal{C}$  and then identified the subset  $C_L$ . For instance defining the *collaborative behavior* as users who finish all quizzes and courses, get highest scores, and participate on the forum, the *targeting behavior* which corresponds to a super student on only one chapter of the MOOC. Such perfect representation are rare, and only few users satisfy the required criterion. Afterward, we randomly select a testing set and ask experts either they agree or disagree with the classification. The experiment can have one of two possible outcomes, either we add the test set to the labeled set  $C_L$  and run the algorithm again for better accuracy, or, we find that the expert becomes aware of some limitations and improves his behavior classes  $\mathcal{C}$  definitions [18]. This iterative process can be repeated as much as needed until the expert is satisfied with the outcome. In our case we converged to the following classification:

- **Participant (P)** Does all the chapters, and answer all the quizzes;
- **Collaborative (C)** Does at least 70% of all the chapters and quizzes but highly active on the forum;
- **Targeting (T)** Targets a chapter, solves the relevant quizzes;
- **Auditor (A)** Reads 70% the chapters but answers 30% of the quizzes;
- **Clicker (Cl)** Does not stay on the same page longer than 5 seconds;
- **Big Starter (BS)** Has a participant behavior up to the first 3 chapters;
- **Late Quitter (LQ)** Has a participant behavior up to the last 3 chapters.

#### 4.2 Dynamic Behavior Clustering

The results of **SBC** motivated the development of **DBC**. We observed that the classes that were satisfactory for the analysis requirement, are actually temporally characterized by a smaller number of simple behavior.

- **Exploration** where the user is randomly skipping through the MOOC;
- **Learning** where the user pays attention to the content of pages;
- **Certification** where the user is interested in the certification and tries to fulfill the courses requirements.

We considered a three dimensional feature space where the weight in each dimension reflects the probability of following the associated behavior. As expected, the users behaviors could be explained by the succession of simple behavior, we converged to 3 modes each of them optimizes one of the behaviors  $\theta_1, \theta_2, \theta_3 \in \{[1, 0, 0], [0, 1, 0], [0, 0, 1]\}$ .

In Figure 1 we observe the expected temporal evolution of modes ( $z_{t=0:T_i}^i$ ) for three users. Some behaviors such as the *clicker* behavior, can be observed as an agent with the unique goal of exploration as shown in the case of user #1. Other behaviors correspond to other patterns such as the *late quitter*, whose behavior is

quite similar to the *participant/collaborative* behavior. The user start exploring a little bit, oscillates between *exploring* and *learning* before adopting mostly the *certifying* behavior. We end up sometimes with an exploration phase before leaving the courses entirely. The difference between the LQ and (C/P) behaviors is the length of the sequences. For instance we observe the fact that LQ tend to explore more by the end.



Fig. 1. Three #MOOCLive users' latent behavior modes

## 5 Discussion

From a practical point of view, both **SBC** and **DBC** were satisfactory as the results satisfied the experts who lead the experimentation process with us. We were able to improve our understanding of the users learning behaviors without requiring additional informations when treating different sets  $\mathcal{C}$  of user classes. The results of the **SBC** are easily interpretable as the outputted behaviors are defined with the help of an expert. However, in the case of **DBC**, the task is more difficult. It mainly depends of the considered feature map and our ability to identify behaviors when observing the outputted parameters. In our cases, we did not struggle as we were anticipating such results.

A big drawback of the **SBC** is the assumption that users behave according to a unique policy throughout the course. To circle around this, the behavior classes had to be specified enough to capture the nuances between the users. Even though **DBC** resolves partially this issue by allowing the users to jump among typical behaviors, a temporal explanation of the mode switching is far from being comprehensively satisfactory. In fact, the users are more likely to switch from an exploration behavior to a learning one because they visited a set of different pages (or states) rather than because they spent a certain amount

of time exploring the state space. The reward is likely to be non-Markovian as it depends of the trajectory a user follows and not just the last state he visits. Indeed, for a user trying to learn the content of a MOOC, a chapter is more rewarding when visited for the first time. An interesting direction for future work would be to tackle such challenging problem.

## References

1. Cristobal Romero and Sebastian Ventura. Educational data science in massive open online courses. Wiley interdisciplinary Reviews-Data Mining and knowledge discovery, 2017.
2. Arti Ramesh, Dan Goldwasser, Bert Huang, Hal Daume III and Lise Getoor. Modeling learner engagement in MOOCs using probabilistic soft logic. NIPS Workshop on Data Driven Education, p. 62, 2013.
3. Linda Corrin, Paula de Barba, Carleton Corin and Gregor Kennedy. Visualizing patterns of student engagement and performance in moocs. 2014.
4. Cheng Ye and Gautam Biswas. Early prediction of student dropout and performance in moocs using higher granularity temporal information. Journal of learning analytics, 2014.
5. Craig Thompson Christopher Brooks and Stephanie Teasley. Towards a general method for building predictive models of learner success using educational time series data. Workshops of the International Conference on Learning Analytics and Knowledge, 2014.
6. Chase Geigle and Cheng Xiang Zhai. Modeling mooc student behavior with two-layer hidden markov models. Learning at Scale, 2017.
7. Richard Sutton and Andrew Barto. Reinforcement Learning: An Introduction. The MIT Press, 2005.
8. Ng Andrew and Stuart J. Russell. Algorithms for Inverse Reinforcement Learning. ICML '00 Proceedings of the Seventeenth International Conference on Machine Learning, 2000.
9. Markus Wulfmeier, Peter Ondruska and Ingmar Posner. Maximum Entropy Deep Inverse Reinforcement Learning, 2016.
10. Sergey Levine, Zoran Popovic and Vladlen Koltun. Nonlinear Inverse Reinforcement Learning with Gaussian Processes. Advances in Neural Information Processing Systems 24 (NIPS), 2011.
11. Amit Surana and Kunal Srivastava. Bayesian Nonparametric Inverse Reinforcement Learning for Switched Markov Decision Processes. In ICMLA 2014 13th International Conference on Machine Learning and Applications, 2014.
12. Monica Babe, Vroman Vukosi Marivate, Kaushik Subramanian and Michael Littman. Apprenticeship learning about multiple intentions. Proceedings of the International Conference on International Conference on Machine Learning, 2011.
13. Ramachandran, Deepak and Eyal Amir. Bayesian Inverse Reinforcement Learning. IJCAI'07 Proceedings of the 20th international joint conference on Artificial intelligence, 2007.
14. Emily Fox, Erik Sudderth, Michael Jordan and Alan Willsky. Nonparametric bayesian learning of switching linear dynamical systems. Neural Information Processing Systems, 2009.
15. Bernard Michini and Jonathan How. Improving the efficiency of Bayesian inverse reinforcement learning. ICRA, 2012.

16. Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical Report, Carnegie Mellon University, 2002
17. Viterbi, Andrew. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm, 1967.
18. Firas Jarboui, Vincent Rocchisani, and Wilfried Kirchenmann. Users behavioural inference with markovian decision process and active learning. In IAL@PKDD/ECML, 2017.