Predicting Learning Dynamics in Multiple-Choice Decision-Making Tasks Using a Variational Bayes Technique*

Ali Yousefi, Reza Kakooee, Mohammad TH. Beheshti, Darin D. Dougherty, Emad N. Eskandar, Alik S. Widge, Uri T. Eden

Abstract— **Multiple-Choice Decision-Making Tasks are widely used to analyze behavior and infer underlying cognitive states that shape the decision and learning processes. The behavioral signals recorded in these tasks are dynamic and often non-Gaussian – for instance, when learning a multiple choice association task. Previously developed estimation algorithms for latent behavioral variables do not address multiple-choice responses. In this research, we use a state-space modeling framework to predict a cognitive learning state related to multiple choice decisions, which are best described by a multinomial distribution. The proposed algorithm combines a multinomial filter/smoother and a variational Bayes technique to estimate the dynamics of a learning state vector. The algorithm is applied to decision response data recorded from non-human primates (NHPs) performing a Multiple-Choice Decision Task.**

I. Introduction

Understanding mechanisms of learning and specifically decision-making is of great interest in systems and behavioral neuroscience. Decision-making is a dynamic process, which is shaped by a subject's interaction with external stimuli and previous experience. Decision-making involves significant changes in brain dynamics and functionality, and is expressed through behavior. For instance, in a task where monkeys decide which direction a moving target is headed, researchers have found that neural activity in the parietal lobe is closely correlated to the monkey's decision and degree of certainty [1]. To better understand neural mechanisms ruling decisions, researchers are working to record more brain regions and run more complex behavioral tasks. This creates a need to develop new analytical tools to characterize dynamical and statistical aspects of both behavioral and neural signals with a higher level of accuracy and precision.

In this research, we focus on the behavioral aspect of decision making in multiple-choice trial-structured decision-making tasks. We propose a new modeling framework to characterize dynamical and statistical aspects of decision-making, given categorical observed behavioral signals. This modeling framework is an extension of a previously developed state-space modeling framework for binary-decision tasks. Here, the observed decision signal is a categorical variable – 1-of-m possible choices – rather than a binary signal – correct or incorrect. Thus, the proposed model not only characterizes the behavioral signal given the binary correctness result, it also demonstrates how the decision is shaped across multiple choices of a task. Categorical data with time varying properties are common in many research domains; thus, the framework proposed here can be a valuable analytical tool across research areas [2].

State-space models have been successfully used in many fields, including neuroscience, to analyze dynamical signals. In this framework, two sets of equations define the model: the state equation(s) and observation equation(s). For instance, in the decision-making process, the state equation models how a correct decision strategy may be learned over time. The observation equations describe the connection between each observed decision and the learning state, which cannot be observed directly. A general problem in the state-space modeling framework is to estimate the state variable given the observable signals. A true estimate of the state variable requires an accurate model of both the state and observation equations and knowledge of the model's free parameters. A proper solution of this problem requires the simultaneous estimation of the model parameters and the state variable. Smith and Brown [3] proposed an approximate expectation-maximization (EM) algorithm to simultaneously estimate model parameters and state variable/s for range of observed signals including point process, mixed binary and continuous signals. This methodology might be extended to the analysis of other forms of behavioral or neural signals including categorical variables recorded in multiple-choice decision tasks.

Here, we extend the state-space modeling framework to estimate model parameters and state variables from multiple-choice behavior. The state equation is defined by a Markov process (often, an AR(1) process) and the decision probability per category is modeled by a softmax sigmoid. Instead of using an EM algorithm, we here develop a Variational Bayes (VB) method to estimate the posterior distribution of the model parameters. Using VB, we have a fully Bayesian estimate of both parameters and state variables. The methodology developed under VB can be extended to more complex data, including a mixture of categorical and continuous signals or even point processes.

In this paper, we first formulate the state-space model of the multiple-choice decision-task. We then describe the solution of the state and parameter posterior distributions using the VB technique. We then show an application of the approach to estimating learning state progression in multiple-choice decision tasks using NHP data. Finally, we discuss the modeling result and its applications and extension to the analysis of other data types.

*This research was funded [in part] by the Defense Advanced Research Projects Agency (DARPA) under Cooperative Agreement Number W911NF-14-2-0045, issued by the Army Research Office contracting office in support of DARPA'S SUBNETS program. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

- AY, and UTE are with the Department of Mathematics and Statistics of Boston University, Boston, MA
- AY, DDD, ENS, and ASW are with Harvard Medical School and Massachusetts General Hospital, Boston, MA
- RK, and MTHB are with the Department of Electrical and Computer Engineering, Tarbiat Modares University, Tehran, Iran
- AY and RK contributed equally to this research.
- UTE and ASW are the co-senior authors and contributed equally to the manuscript.
- Address for reprint requests and other correspondence: Ali Yousefi, Email: ayousefi@mgh.harvard.edu

II. STATE-SPACE MODEL FOR CATEGORICAL OBSERVATIONS

In this section, we formulate the learning state dynamics in a multiple-choice decision-task using the state-space modeling framework. The state-equation defines the learning evolution in time; here we define it as a 1st order autoregressive process

$$x_{k+1} = a_1 * x_k + a_0 + \sigma_v^2 v \quad v \sim N(0, I) \tag{1}$$

where, x_k corresponds to the learning state at time - or trial – index k. The (a_1, a_0, σ_v^2) are the equation free parameters, which define how the state variable evolves through time. If x_k is scalar valued, it might represent a single learning state, upon which the probability of all possible choices depend. If x_k is vector valued, each dimension might represent a learning state associated with one choice goal or set of choice goals. In that case, a_1 would define an interaction between these learning state components.

The observation signal at each trial is a categorical variable – y_k, which represents the response choice. For a L-way categorical variable, the probability of l^{th} choice is defined by a Softmax sigmoid function:

$$\mathrm{P}_{k,l} = \mathrm{P}(\mathrm{l}^{th} \text{ choice}) = \frac{\exp(b_{1,l} * x_k + b_{0,l})}{\sum_l \exp(b_{1,l} * x_k + b_{0,l})} \quad l \in \{1, \cdots, L\} \tag{2}$$

where, the parameter set $(b_{0,l}, b_{1,l})$ $l = 1, \cdots, L$ represents the relationship between learning state – x_k – and observed decision. Note that in equation (2), in order have an identifiable model, we assume $(b_{1,L}, b_{0,L})$ are fixed - note that, we can pick any pair and set them to a fixed value. The values for $(b_{1,L}, b_{0,L})$ can be optional, but to provide an interpretable result, we set $(b_{1,L}, b_{0,L})$ to (0,0) or a small number. The probability of observing y_k at time k is defined by

$$\mathrm{P}(y_k) = \prod_l \mathrm{P}_{k,l}{}^{y_{k,l}} \qquad \sum_l y_{k,l} = 1 \tag{3}$$

Equations (1) to (3) define the state-space modeling framework for a multi-choice decision task. For example, each direction of a two-alternative direction-discrimination task can be described by a 2-way categorical variable [1]. For the task, it is expected that probability of the correct response will grow as the task proceeds. If the 1st choice represents the correct response and the learning state grows over time, we would expect $b_{1,1}$ to have a positive value. In that case, as the learning state grows, the probability of correct choice would increase. If a_1 is one and a_0 is zero, the learning state undergoes a random walk. A value of a_1 below one might suggest a tendency to forget and a value of a_0 above zero might suggest a positive drift in learning. Finally, σ_v^2 value describes how smooth the learning path will be. Here, the objective is to properly estimate the model parameters – a subset of $(b_{0,l}, b_{1,l})$ $l = 1, \cdots, L-1$ and (a_1, a_0, σ_v^2), which fully describes the observed decisions.

EM approaches have previously been developed for parameter estimation in the state-space modeling framework with different types of observation signals [3]. For our problem, we utilize a VB approach instead of EM to estimate unknown model parameters. VB is an analogue of EM in the fully Bayesian setting. By utilizing VB, we are able to provide prior distributions of model parameters in the estimation process and compute posterior distributions for both the state process and model parameters.

In the following subsections, we describe the steps of the parameter estimation algorithm using the VB technique. We start by defining priors on the model parameters for both the state and observation equations. For the case of a scalar state, we assume the prior distribution on the model parameters is defined by - $\theta = \{a_1, a_0, \sigma_v^2, (b_{0,1}, b_{1,1}), \cdots, (b_{0,L-1}, b_{1,L-1})\}$.

$$\mathrm{P}(\theta) = \mathrm{P}(a_0, a_1) * \mathrm{P}(b_{0,1}, b_{1,1}) * \cdots * \mathrm{P}(b_{0,L-1}, b_{1,L-1}) * \mathrm{P}(\tau_v) \tag{4}$$

where, $\tau_v = 1/\sigma_v^2$ – the inverse of noise variance. We define the prior on the (a_0, a_1) by a bi-variate normal distribution

$$\begin{pmatrix} a_1 \\ a_0 \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_{a_1} \\ \mu_{a_0} \end{pmatrix}, \begin{pmatrix} \sigma_{a_1}^2 & \rho_a * \sigma_{a_1} * \sigma_{a_0} \\ \rho_a * \sigma_{a_1} * \sigma_{a_0} & \sigma_{a_0}^2 \end{pmatrix}\right) \tag{5}$$

where, $(\mu_{a_0}\ \mu_{a_1})'$ is the prior mean of (a_0, a_1) parameters. The $(\sigma_{a_1}, \sigma_{a_0}, \rho_a)$ define elements of prior covariance matrix. The τ_v is a positive variable and its prior is defined by a Gamma distribution

$$P(\tau_v) \sim \mathrm{Gamma}(a_v, b_v) \tag{6}$$

where, (a_v, b_v) are the shape and inverse scale parameters of the prior. We can also define the prior of the initial value of the state-variable - x_0; this can be defined by a normal distribution with (m_0, σ_0^2) parameters. Finally, we define the prior on the parameters of Softmax functions. The prior on these parameters are assumed to be Normal, and have a similar form of equation (5).

For the case of a vector valued state, we will assume that each component evolves independently with its own set of parameters (suggesting that each choice goal has it's own learning state process and observation model and that these do not interact). Note that extending this to interacting learning state components would not be difficult with a good choice of priors.

A. Variational Bayes Solution of the Model

Using the VB technique, we work to derive an analytical approximation to the posterior probability of the unknown variables. In the state-space modeling framework, both the model parameters and state-variable - x_k for all trial indices - are unknown. Thus, we derive the posterior probability over model parameters and state variables given the observed signals [4].

To derive the posterior distribution, we calculate the joint probability distribution given the whole observation sequence of state-variable and model parameters. Let's assume $Y_K = [y_1, \cdots, y_K]$ is the observed data in a task with K trials, and $X_K = [x_1, \cdots, x_K]$ is the corresponding state variable; the joint probability distribution is defined by

$$\mathrm{P}(X_K, Y_K, \theta) = \mathrm{P}(Y_K | X_K, \theta) * \mathrm{P}(X_K | \theta) * \mathrm{P}(\theta)$$
$$= \prod_k \mathrm{P}(y_k | x_k, \theta) * \mathrm{P}(x_k | x_{K-1}, \theta) * \mathrm{P}(x_0 | \theta) * \mathrm{P}(\theta) \tag{7}$$

Under VB, we assume the posterior over the state-variable and parameters can be approximated by

$$\mathrm{P}(X_K, \theta | Y_K) \approx Q(X_K, \theta) = Q_{X_K}(X_K) * Q_\theta(\theta) \tag{8}$$

where, we further assume $Q_\theta(\theta)$ can be factorized by

$$Q_\theta(\theta) = q_{a_1,a_0}(a_1, a_0) * q_{b_{1,1},b_{0,1}}(b_{1,1}, b_{0,1}) * \cdots$$
$$* q_{b_{1,L},b_{0,L}}(b_{1,L-1}, b_{0,L-1}) * q_{\tau_v}(\tau_v) \tag{9}$$

The solution for $Q_{X_K}(X_K)$ and different $q_x(x)$ $x \in \{(a_1, a_0), (b_{1,1}, b_{0,1}), \cdots, (b_{1,L-1}, b_{0,L-1}), \tau_v\}$ is defined by

$$\log Q_{X_K}(X_K) = \mathrm{E}_{\theta | Y_K, X_K}[\log \mathrm{P}(X_K, Y_K, \theta)] \tag{10.a}$$

$$\log q_x(x) = \mathrm{E}_{\theta_{-x},X_K|Y_K}[\log \mathrm{P}(X_K, Y_K, \theta)] \quad (10.b)$$

where, θ_{-x} is the full parameter set with parameter x removed. Here, $\theta|Y_K, X_K$ is the posterior estimate of θ given both Y_K and X_K, and the $\theta_{-x}, X_K|Y_K$ is the posterior estimate over a subset of parameters and X_K.

To find the optimal posteriors over model parameters and state-variable, we iteratively compute the values in equations (10.a) and (10.b). Through this iteration, we increase the likelihood of the observation by updating the posterior over model unknowns. In the following parts, we derive the solution for both of (10.a) and (10.b).

B. State Variable Posterior Estimation

To calculate the posterior over the state-variable, we substitute equation (7) into (10.a) and take its expectation over the model parameters. The parameter distribution is defined by its posterior estimate from the previous iteration. We assume the posteriors over model parameters have the same form as their priors; thus, we require taking the log of the joint probability distribution over a mixture of Normal and Gamma distributions. The right side of equation (10.a) is defined by

$$\mathrm{E}_\theta[\log P(\cdot)] \cong -1/2 * \frac{a_v}{b_v} * \sum_k E_\theta(x_k - a_1 * x_{k-1} - a_0)^2 \cdots$$
$$+\sum_k \sum_l y_{k,l} * (\mu_{b_{1,l}} * x_k + \mu_{b_{0,l}}) \cdots$$
$$-\sum_k \mathrm{E}_\theta \log\left(\sum_l \exp(b_{1,l} * x_k + b_{0,l})\right) + \mathrm{C}_0 \quad (11)$$

where, a_v/b_v is the expected value of τ_v and C_0 refers to all the constant terms not being linked to the state variable. To find the posterior of the state variable, we rewrite this expectation in the form of state- and observation equations. The first two terms of the expectation are linear and quadratic terms of the state-variable – the expectation in the first the term is easy to calculate, and if we approximate the last term using its second-order Taylor approximation; we are able to derive an equivalent state-space representation of the state-variable posterior. The approximate expectation of the last term is

$$\mathrm{E}_\theta \log\left(\sum_l \exp(b_{1,l} * x_k + b_{0,l})\right) \approx \sum_l \underbrace{\mathrm{E}_\theta(b_{1,l} * p_l)}_{f_{k,2,l}} * (x_k - \hat{x}_k) \cdots$$
$$+0.5 * \sum_l \underbrace{\mathrm{E}_\theta(b_{1,l}^2 * p_l * (1 - p_l))}_{f_{k,1,l}} * (x_k - \hat{x}_k)^2 + \mathrm{C}_1 \quad (12)$$

where, C_1 is a constant term and p_l is defined by

$$p_l = \exp(b_{1,l} * \hat{x}_k + b_{0,l}) / \sum_l \exp(b_{1,l} * \hat{x}_k + b_{0,l}) \quad (13)$$

where, to calculate the expectations over θ defined in equation (12) we use a Monte Carlo estimate - note the normal posterior over the parameters. The equivalent state-equation is defined by

$$x_k = \left(\mu_{a_1} + \frac{\sigma_{a_1}^2}{\mu_{a_1}}\right) x_{k-1} + \mu_{a_0} + \frac{\rho_a * \sigma_{a_1} * \sigma_{a_0}}{\mu_{a_1}} + \hat{v} \quad \hat{v} \sim N\left(0, \frac{b_v}{a_v}\left(1 + \frac{\sigma_{a_1}^2}{\mu_{a_1}^2}\right)\right) (14)$$

The observation equations consist of L equations, one per each element of the categorical variable expect the last one. The l^{th} observation equation is defined by

$$y_{k,l} = \frac{f_{k,1,l}}{\mu_{b_{1,l}}} * (x_k - \hat{x}_k) + \frac{f_{k,2,l}}{\mu_{b_{1,l}}} + \varepsilon_{k,l} \qquad \varepsilon_{k,l} \sim N\left(0, \frac{f_{k,1,l}}{\mu_{b_{1,l}}^2}\right) \quad (15)$$

where, we generally pick $\hat{x}_k$ to be $x_{k|k-1}$. The equivalent state-space model defined by equations (14) and (15) is a linear time-varying Gaussian process. We can thus run a Kalman filter and smoother to find the posterior distribution of x_k, which is defined by $x_{k|K}$ [5].

C. Model Parameters Posterior Estimation

To calculate the posterior over the parameters, we take the expectation of the log of the joint probability distribution over the state-variable posterior. The state-variable posterior is approximately Gaussian; thus, the posterior over (a_1, a_0) will be Gaussian given its Gaussian prior – note that the posteriors over θ and X_K are assumed to be independent. Thus, we only require finding update rules for the mean and covariance matrix of the posterior. The update rule is defined by

$$\frac{1}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_1}^2} \propto \frac{1}{(1-\rho_a{}^2) * \sigma_{a_1}^2} + \frac{a_v}{b_v} * \sum_k \mathrm{E}_{X_K|Y_K}(x_{k-1}^2) \quad (16.a)$$

$$\frac{1}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_0}^2} \propto \frac{1}{(1-\rho_a{}^2) * \sigma_{a_0}^2} + \frac{a_v}{b_v} * K \quad (16.b)$$

$$\frac{\hat{\rho}_a}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_0} * \hat{\sigma}_{a_1}} \propto \frac{\rho_a}{(1-\rho_a{}^2) * \sigma_{a_0} * \sigma_{a_1}} - \frac{a_v}{b_v} * \sum_k \mathrm{E}_{X_K|Y_K}(x_{k-1}) \quad (16.c)$$

$$\begin{bmatrix} \frac{1}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_1}^2} & -\frac{\hat{\rho}_a}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_1} * \hat{\sigma}_{a_0}} \\ -\frac{\hat{\rho}_a}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_1} * \hat{\sigma}_{a_0}} & \frac{1}{(1-\hat{\rho}_a{}^2) * \hat{\sigma}_{a_0}^2} \end{bmatrix} * \begin{bmatrix} \hat{\mu}_{a_1} \\ \hat{\mu}_{a_0} \end{bmatrix} =$$
$$\begin{bmatrix} \frac{a_v}{b_v} * \sum_k \mathrm{E}_{X_K|Y_K}(x_k * x_{k-1}) + \frac{\mu_{a_1}}{(1-\rho_a{}^2) * \sigma_{a_1}^2} - \frac{\rho_a \mu_{a_0}}{(1-\rho_a{}^2) * \sigma_{a_0} * \sigma_{a_1}} \\ \frac{a_v}{b_v} * \sum_k \mathrm{E}_{X_K|Y_K}(x_k) + \frac{\mu_{a_0}}{(1-\rho_a{}^2) * \sigma_{a_0}^2} - \frac{\rho_a \mu_{a_1}}{(1-\rho_a{}^2) * \sigma_{a_0} * \sigma_{a_1}} \end{bmatrix} \quad (16.d)$$

where, parameters with a ^ sign represent the posterior estimators. To complete the update rule, we calculate the state-variable mean, variance, and one-step covariance. The first two terms are trivial, and solution of the one-step covariance can found in [6].

The posterior of τ_v will be a Gamma distribution under Gaussian assumption of the state variable posterior. The update rule for distribution parameters are defined by

$$\hat{a}_v = a_v + \frac{K}{2} \quad (17.a)$$

$$\hat{b}_v = b_v + \frac{1}{2} * \sum_k \mathrm{E}_{X_K,(a_1,a_0)|Y_K}(x_k - a_1 * x_{k-1} - a_0)^2 \quad (17.b)$$

where, we assume the posterior on (a_1, a_0) and X_K are independent. To calculate the expectation, we expand the quadratic term and calculate each term expectation over the state-variable and parameters.

To find the posterior of $(b_{0,l}, b_{1,l})$ for different ls, we require calculating expectation of log of softmax function over the state-variable posterior. There is a non-linear function in this expectation; thus, the assumption of a Normal posterior over $(b_{0,l}, b_{1,l})$ is invalid. We approximate the non-linear term with its second-order Taylor expansion; thus, similar to the state-variable, we have an approximate normal posterior. The Taylor expansion of the non-linear term is defined by

$$\mathrm{E}_{X_K|Y_K} \log\left(\sum_l \exp(b_{1,l} * x_k + b_{0,l})\right) = \mathrm{C}_2 +$$
$$\sum_k \sum_l \mathrm{E}_{X_K|Y_K}[x_k * p_{k,l}] * (b_{1,l} - \mu_{b_{1,l}}) + \mathrm{E}_{X_K|Y_K}[p_{k,l}] * (b_{0,l} - \mu_{b_{0,l}}) \cdots$$
$$+0.5 * \mathrm{E}_{X_K|Y_K}[p_{k,l} * (1 - p_{k,l})] * (b_{0,l} - \mu_{b_{0,l}})^2 \cdots$$
$$+0.5 * \mathrm{E}_{X_K|Y_K}[x_k^2 * p_{k,l} * (1 - p_{k,l})] * (b_{1,l} - \mu_{b_{1,l}})^2 \cdots$$
$$+\mathrm{E}_{X_K|Y_K}[x_k * p_{k,l} * (1 - p_{k,l})] * (b_{0,l} - \mu_{b_{0,l}}) * (b_{1,l} - \mu_{b_{1,l}}) \quad (18)$$

where, we can use a Monte Carlo estimate to find expectation terms. Note that $p_{k,l}$ is defined by

$$p_{k,l} = \exp\left(\mu_{b_{1,l}} * x_k + \mu_{b_{0,l}}\right) / \sum_l \exp\left(b_{1,l} * \hat{x}_k + b_{0,l}\right) \quad (19)$$

Using the quadratic approximation defined in equation (18), the update rule for each pair will be similar to the update rule defined for (a_1, a_0). Note that we can include the update rule for x_0 which is easy given the normal posterior assumption for the parameters.

III. Application In Behavioral Signal Analysis

In this section, we apply this approach to model learning during multiple-choice decision-making. The behavioral signal is an NHP's response to a 3-choice decision-task. There are 90 trials in total, and the subject has three choices per trial. The correct response to stimulus A is response A, and so forth. Figure 1.a shows the subject response to each trial; the subject picks response C to stimulus A in the first trial, then picks response B to the same stimulus in the second trial. As the task proceeds, the subject adopts a different strategy. For stimulus A, the subject switches between response A and B, and its response doesn't converge to the correct choice even after many trials. For stimulus B, the subject selects response A and B at the beginning, then switches to B and C responses, and as the task proceeds, correctly responds to each B stimulus. It can be also seen that the subject strategy for stimulus C is completely different from A and B. The data demonstrates how complex learning can be, and why analyzing this as a binary correct/incorrect response would fail to capture all the features of the learning process. For instance, by using the correct/incorrect signal, we would not able to identify the strategy evidenced by the incorrect responses, and determine whether the subject learns the through a sequential process or through random exploration. The goal of this research was to develop an algorithm that makes understanding of complex learning processes possible.

Here, we present the modeling result for stimuli B and C. For each choice, we use equations (1) and (2) with L equal to 3. We fix the (a_1, a_0) with values $(1,0)$ to define a learning process that is a random walk. We pick the initial values for τ_v distribution by $a_v = 20$ and $b_v = 100$; by using these values, we set the state-transition noise to be reasonably high and let the model bring the state-transition noise variance to its proper value given the observation attributes. The remaining parameters of the model – $(\tau_v, b_{0,1}, b_{1,1}, b_{0,2}, b_{1,2}, b_{0,3}, b_{1,3})$ – are updated by running the equivalent state-space model and parameter update rule defined in equations (16) to (19). The initial values for $b_{1,1}$ and $b_{1,2}$ are set to 1.0, and the initial values for $b_{0,1}$ and $b_{0,2}$ are set to 0. The prior covariance for each $b_{*,*}$ pair is set to an identity matrix; this prior gives enough space for the observation parameters to be updated using the learning rule. We also set $b_{1,1} = 0.001$ and $b_{0,1} = 0$; given this value, the learning state growth suggests that the subject has learned the category better.

Figure 1.b and 1.c show the learning state filter and smoother results for both stimuli B and C. The 1.b curve goes continually upward, suggesting that the subject learned stimulus B correctly. The learning of stimulus C was less robust, as shown by the lower plateau. Figures 1.d and 1.e show the mean probability of each response – along with 95% HPD of the probability estimate – when seeing stimuli B and C. The result suggests that the subject learns choice B by ignoring A first and then ignoring C. The result also suggests the subject did not learn stimulus C; its accuracy barely goes above chance and drops at the end of the task.

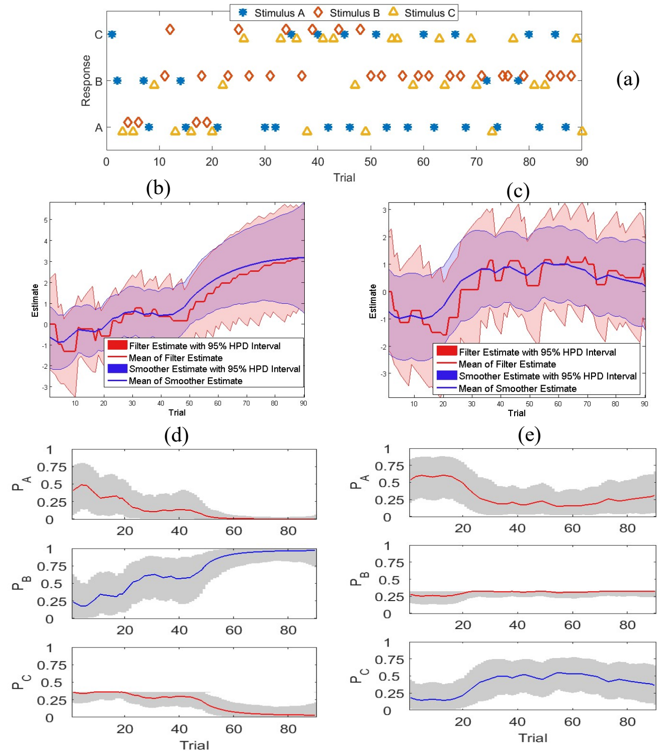

Figure 1 Behavioral data and analysis results. **a.** Decision signal. There are three categories to be learned and there are three choices per trial.l. **b.** Learning state filter and smoother estimate on the B stimulus trials **c.** Learning state filter and smoother estimate on the C stimulus trials **d.** 95% HPD region of different response options on the B stimulus trials **e.** 95% HPD region of different response options on the C stimulus trials

IV. Conclusion

We developed a state-space modeling framework to study multi-choice decision-making. Instead of considering only the correct or incorrect response, we use the categorical decision signal to provide a more elaborate understanding of the learning dynamic in this family of cognitive tasks. The algorithm presented here assumes that the learning states for each stimulus evolve independently; we plan to extend this application to examine dependence between the learning state trajectories in a larger dataset.

References

[1] Wei, Ziqiang, and Xiao-Jing Wang. "Confidence estimation as a stochastic process in a neurodynamical system of decision making." Journal of neurophysiology 114.1 (2015): 99-113.

[2] Agresti, A., & Kateri, M. (2011). *Categorical data analysis* (pp. 206-208). Springer Berlin Heidelberg

[3] Smith, Anne C., and Emery N. Brown. "Estimating a state-space model from point process observations." Neural Computation 15.5 (2003): 965-991.

[4] Beal, Matthew J.; Ghahramani, Zoubin. Variational Bayesian learning of directed graphical models with hidden variables. Bayesian Anal. 1 (2006), no. 4, 793--831

[5] Särkkä, Simo. Bayesian filtering and smoothing. Vol. 3. Cambridge University Press, 2013.

[6] Yousefi, Ali, et al. "Cognitive state prediction using an EM algorithm applied to gamma distributed data." EMBC, 2015 37th Annual International Conference of the IEEE. IEEE, 2015.