HHS Public Access

Author manuscript

Conf Proc IEEE Eng Med Biol Soc. Author manuscript; available in PMC 2018 August 31.

Published in final edited form as:

Conf Proc IEEE Eng Med Biol Soc. 2015 ; 2015: 7808–7813. doi:10.1109/EMBC.2015.7320203.

Estimating a Dynamic State to Relate Neural Spiking Activity to Behavioral Signals during Cognitive Tasks*

Xinyi Deng,
Graduate Program in Statistics, Boston University, Boston, MA 02215 USA (xinyi@math.bu.edu).

Rose T. Faghih [Member, IEEE],
Department of Brain and Cognitive Sciences, and the Neuroscience Statistics Research Laboratory at Massachusetts Institute of Technology, Cambridge, MA 02139 USA, and also with Massachusetts General Hospital, Boston, MA 02114 USA (rfaghih@mit.edu).

Riccardo Barbieri [Senior Member, IEEE],
Neuroscience Statistics Research Laboratory, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114 USA, and also with Massachusetts Institute of Technology, Cambridge, MA 02139 USA (barbieri@neurostat.mit.edu).

Angelique C. Paulk,
Department of Neurosurgery, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114 USA (apaulk@mgh.harvard.edu).

Wael F. Asaad,
Departments of Neurosurgery and Neuroscience, Alpert Medical School, Brown University and Rhode Island Hospital, Providence, RI 02912 USA (wael-asaad@brown.edu).

Emery N. Brown [Fellow, IEEE],
Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02115 USA, and also with the Institute for Medical Engineering and Science, and the Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (enb@neurostat.mit.edu).

Darin D. Dougherty,
Department of Psychiatry, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114 USA (ddougherty@partners.org).

Alik S. Widge,
Department of Psychiatry, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114 USA, and also with Picower Institute for Learning and Memory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (awidge@partners.org).

Emad N. Eskandar, and
Nayef Al-Rodhan Laboratories, Department of Neurosurgery, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114 USA (eeskandar@mgh.harvard.edu).

*This research was supported by the Defense Advanced Research Projects Agency (DARPA) under contract number W911NF-14-2-0045 and the National Institute of Neurological Disorders and Stroke [R01 NS073118]. The opinions presented are those of the authors and not of DARPA, NINDS, or our institutions.

Uri T. Eden
Department of Mathematics and Statistics, Boston University, Boston, MA 02215 USA (tzvi@bu.edu).

Abstract

An important question in neuroscience is understanding the relationship between high-dimensional electrophysiological data and complex, dynamic behavioral data. One general strategy to address this problem is to define a low-dimensional representation of essential cognitive features describing this relationship. Here we describe a general state-space method to model and fit a low-dimensional cognitive state process that allows us to relate behavioral outcomes of various tasks to simultaneously recorded neural activity across multiple brain areas. In particular, we apply this model to data recorded in the lateral prefrontal cortex (PFC) and caudate nucleus of non-human primates as they perform learning and adaptation in a rule-switching task. First, we define a model for a cognitive state process related to learning, and estimate the progression of this learning state through the experiments. Next, we formulate a point process generalized linear model to relate the spiking activity of each PFC and caudate neuron to the stimated learning state. Then, we compute the posterior densities of the cognitive state using a recursive Bayesian decoding algorithm. We demonstrate that accurate decoding of a learning state is possible with a simple point process model of population spiking. Our analyses also allow us to compare decoding accuracy across neural populations in the PFC and caudate nucleus.

I. INTRODUCTION

Demonstrating the existence of meaningful relationships between behavior and neural activity is essential to our understanding of the brain and has been a subject of intensive investigation in neuroscience. Behavioral and cognitive neuroscientists aim to understand how the brain uses neural activity to integrate sensory inputs, control movements, facilitate learning and memory, activate and express emotions, etc. Neuroengineers focus on how to decode and stimulate neural activity to assist, supplement or suppress behavior.

Investigations of these questions have led to new multi-faceted experimental design and have generated behavioral data with growing complexity. In addition, technological advances now allow for recording of large quantities of information from the brain at multiple spatial and temporal scales, includes multi-channel electrode arrays, EEG, MEG and fMRI. Access to this type of high-dimensional data, both behavioral and neural, has also presented a challenge for statistical data analysis and modeling: What is an adequate representation of the relation between features of the behavioral task and structures in the neural activity?

Traditionally, studies of neural coding focus on the relation between electrophyiological data and directly observable biological or behavioral signals. For example, place-field models describe spiking activity in hippocampus as a function of an animal’s position in its environment. Recently, there has been increasing interest in models relating neural activity to more general variables that influence multiple aspects of behavior and cognitive function. Understanding the structure of such cognitive variables may be essential to the study of multiple neural disease processes. For example, deficits in cognitive flexibility have been

linked to autism, obsessive-compulsive disorders and schizophrenia. However, cognitive flexibility is only observable through its influence on behavior, and therefore difficult to link to neural activity directly. Other examples of cognitive states that may provide meaningful links between behavior and neural activity include features like attention, affective response tendency, and approach-avoidance level. Finally, such cognitive variables are often dynamic, leading to changing behavioral outcome to stimuli through time. An important statistical challenge is to understand neural representations of these cognitive state processes and to estimate their dynamics through time.

One approach that has been successful in linking neural data to dynamic, unobserved signals is state-space modeling [1]–[8]. For example, state-space modeling has been used to predict the movements of a rat from ensemble place-cell activity [9], [10]. However, such methods require the behavioral signal to be estimated to be low-dimensional and directly observable during a first, encoding step. This would not be possible for the abstract, cognitive state processes, described above.

Here we present one possible solution: a general state-space paradigm to model and fit a low-dimensional cognitive state process that allows us to relate outcomes of various behavioral tasks to simultaneously recorded neural activity across multiple brain areas. The paradigm consists of three steps. First, we estimate the dynamics of a cognitive state variable using previous knowledge of its influence on observed behavioral signals. Second, we construct models that use the estimated state and relevant covariates related to behavior to describe the statistical structure of neural activity. Last, we estimate the dynamic state again, this time using only the neural activity. We illustrate the proposed paradigm with an application to data recorded in the lateral prefrontal cortex (PFC) and caudate nucleus of non-human primates as they perform learning and adaptation in a rule-switching task.

II. COGNITIVE STATE-SPACE DECODING PARADIGM

In this section, we first construct a low-dimensional cognitive state process to relate behavior and neural activity. We then describe a general three-model paradigm to estimate the cognitive state in three steps, summarized in Fig. 1.

A. Model Framework

We model the underlying cognitive function (signal) during a task as a stochastic and dynamic process. The abstract state variable x_k evolves through time according to some dynamics:

$$x_k \mid x_{k-1} \sim f\left(x_k \mid x_{k-1}\right). \quad (1)$$

In most situations, if not all, we do not observe the cognitive state. In other words, x_k is "hidden". Some examples of this unobserved abstract cognitive state are reward motivation level, susceptibility to fear, flexibility of learning, etc.

Instead, we obtain some behavioral data $\vec{z}_k$ related to the task, parameterized by θ_z:

$$\vec{z}_k \mid x_k \sim g\left(\vec{z}_k \mid x_k; \theta_z\right). \quad (2)$$

$\vec{z}_k$ can be a vector with components from distinct distributions, both discrete and continuous. For example, $\vec{z}_k$ can includes a binary component of correct vs. incorrect choices and a continuous component of reaction times.

We also record, in addition to the behavioral data, some electrophysiological activity $\vec{y}_k$, parameterized by θ_y:

$$\vec{y}_k \mid x_k \sim h\left(\vec{y}_k \mid x_k; \theta_y\right). \quad (3)$$

Similarly, $\vec{y}_k$ can also include both discrete neural signals such as single unit spiking activity and continuous neural signals such as local field potentials.

B. Estimating State Dynamics From Behavior

Our paradigm for estimating the cognitive state x_k involves three steps. In the first step, we estimate the state dynamics, $p\left(x_k \mid \vec{z}_{1:k}\right)$, using previous knowledge of the behavioral task structure $g\left(\vec{z}_k \mid x_k; \theta_z\right)$ and some smoothing constraints $f(x_k|x_{k-1})$ on the state process dynamics. Assuming the function g in (2) is known, we use this known structure of g to estimate the state dynamics during behavioral experimental tasks:

$$p(x_k \mid \vec{z}_{1:k}) \propto p(\vec{z}_k \mid x_k) \int p(x_k \mid x_{k-1}) p(x_{k-1} \mid \vec{z}_{1:k-1}) dx_{k-1}. \quad (4)$$

The integral on the right hand side of (4) is the one-step prediction density $p\left(x_k \mid \vec{z}_{1:k-1}\right)$ defined by the Chapman-Kolmogorov equation. Here we have assumed that given the past state value, x_{k-1}, the distribution of the current state does not depend on the past behavior. The integral in (4) typically cannot be solved analytically, but multiple numerical and approximation methods are available to compute its value at each time point. One approach is to apply some kind of filtering algorithm such as Kalman filters, which compute the distribution of the state given parameter estimates $\hat{\theta}$. Extensions of these methods, such as Expectation-Maximization (EM) and sequential Monte Carlo, simultaneously optimize the model parameter estimates $\hat{\theta}$ and the unobserved cognitive state process.

C. Neural Encoding

Second, we characterize the relationship between the cognitive state and features of neural activity. In this encoding step, we use an estimate of x_k from the previous step to identify models for function $h\left(\vec{y}_k \mid x_k; \theta_y\right)$ with unknown structures in (3).

An example of $h\left(\overrightarrow{y}_k \mid x_k;\theta_y\right)$ can be a point process model with condition intensity function $\lambda(x_k)$, when $\overrightarrow{y}_t$ is single unit neural spiking activity:

$$p(\overrightarrow{y}_k \mid x_k) = [\lambda\left(x_k;\theta_y\right)\Delta]^{\overrightarrow{y}_k}\exp[\lambda(x_k;\theta_y)\Delta], \quad (5)$$

where $\lambda(x_k;\theta_y)$ can be estimated by parametric models of generalized linear model form.

Another example of $h\left(\overrightarrow{y}_k \mid x_k;\theta_y\right)$ can be multivariate Gamma models with mean $\mu(x_k;\theta_y)$, when $\overrightarrow{y}_t$ is vector of power estimates in local field potential at specific frequencies.

In any case, we treat the estimated state process from step 1 as known, with some uncertainty, and use model fitting methods, such as maximum likelihood to estimate the unknown model parameters.

D. Neural Decoding

Third, we estimate the dynamic state x_k from a new dataset that includes both neural and behavioral activity. More specifically, in this "decoding" step, we compute the posterior distribution of the state process conditioned on the observed neural activity up until the current time:

$$p\left(x_k \mid \overrightarrow{y}_k, \overrightarrow{z}_k\right) \propto p\left(\overrightarrow{y}_k, \overrightarrow{z}_k \mid x_k\right) \int p\left(x_k \mid x_{k-1}\right)p\left(x_{k-1} \mid \overrightarrow{y}_{1:k-1}, \overrightarrow{z}_{1:k-1}\right)dx_{k-1}. \quad (6)$$

If we assume that conditioning on the state, behavior and neural activity are independent, then

$$p\left(\overrightarrow{y}_k, \overrightarrow{z}_k \mid x_k\right) = p\left(\overrightarrow{y}_k \mid x_k\right)p\left(\overrightarrow{z}_k \mid x_k\right). \quad (7)$$

If we choose to decode x_k during the structured behavioral task, then we use both known g and identified models for h to estimate x_k. If we choose to decode x_k outside of the structured task, we use identified h to estimate x_k.

III. APPLICATION: DECODING LEARNING STATE FROM SPIKING ACTIVITY IN MONKEY PFC AND CAUDATE

A. Experimental Data

We illustrate the application of the proposed paradigm with an example study. The behavioral and neural data are obtained from two monkeys performing a temporally delayed, on-line learning task in which they had to determine by trial-and-error which of four picture

cues or spatial locations was currently rewarded within a learning block. Individual blocks followed either a “spatial” or an “object” rule. In the “spatial” rule, the animal was required to choose the target in the same location on every trial (e.g., always upper right). In the “object” rule, the correct action was to choose a picture that matched a picture cue (e.g., always a blue sailboat). The “spatial” rule is substantially easier and rewards perseverative behavior, while the “object” rule rewards flexibility. Details of the behavioral paradigm, data acquisition, and previous analyses of this experimental data are discussed in [11].

In this specific example, the behavioral data is whether the monkey chose the correct location on each trial. The neural data is spiking activity recorded in the lateral prefrontal cortex (PFC) and caudate nucleus of the monkeys. The cognitive state is whether the subject has learned the rule of the task.

B. Estimate State Dynamics From Behavior

We take advantage of previous development of a dynamic approach to analyzing learning experiments with binary responses [11]–[15]. We use a state-space model of learning in which a Bernoulli probability model describes behavioral task responses and a Gaussian state equation describes the hidden state process.

In other words, $p(\overrightarrow{z}_k \mid x_k)$ in (4) is expressed as the Bernoulli probability mass function:

$$p(\overrightarrow{z}_k \mid x_k) = q_k^{\overrightarrow{z}_k}(1 - q_k)^{1 - \overrightarrow{z}_k}, \quad (8)$$

where q_k is defined by the logistic equation:

$$q_k = \frac{\exp(\mu + x_k)}{1 + \exp(\mu + x_k)}, \quad (9)$$

and μ is determined by the probability of a correct response by chance in the absence of learning or experience. Here x_k defines the learning state of the animal at trial k in the experiment. The unobservable state process $x_k|x_k$ 1 is defined as a random walk:

$$x_k = x_{k-1} + \varepsilon_k, \quad (10)$$

where the ε_k are independent Gaussian random variables with mean 0 and variance σ_ε^2. The one-step prediction density $p(x_k \mid \overrightarrow{z}_{1:k-1})$, or learning curve, is the probability of a correct response as a function of the state process and is calculated using an EM algorithm:

$$f(q \mid \mu, x_{k \mid k}, \sigma^2_{k \mid k}) = \left[\left(2\pi\sigma^2_{k \mid k}\right)^{1/2} q(1-q)\right]^{-1} \exp\left(-\frac{1}{2\sigma^2_{k \mid k}}\left[\log\left[q[(1-q)\exp(\mu)]^{-1}\right] - x_{k \mid k}\right]^2\right). \quad (11)$$

Detailed estimation methods are referred to in [12].

Fig. 2 shows two examples of the learning curves estimated by the EM algorithm in two learning blocks in the rule-switching behavioral task. The correct and incorrect responses are shown, respectively, by black and gray marks above the panels. Neglecting the possibility of behavioral preferences or other biases, the probability of a response occurring by chance is shown as a horizontal line at 0.25. Solid blue lines are the learning curve estimates, and the dotted blue lines are the associated 95% confidence intervals. The lower confidence bounds for the learning trial estimates remained above 0.25 after trial 14 and 63, which are, respectively the learning trials for the two learning block examples shown here.

For simplicity, we further dichotomize the trials within each learning block to be “learned” trials if the lower bound of the learning state estimate remains above 0.25 for the remainder fo the trial block or “not learned” trials if otherwise.

C. Neural Encoding

Because the neural data in our example is spiking activity, we present a point process generalized linear model (GLM) approach [3], [16] for constructing a conditional intensity function that characterizes the spiking activity of PFC and caudate neurons. The conditional intensity function [17], [18] relates spiking probability simultaneously to the temporal features of the behavioral task.

In this case, the conditional intensity model is defined as follows:

$$\log\lambda^c(t) = \sum_{j=1}^{2}\sum_{i=1}^{N} \alpha^c_{i,j} B^c_{i,j}(t). \quad (12)$$

Here $c = 1, \cdots C$ is the index of the neuron. $j = 1,2$ is the binary indicator of the behavioral outcome of the trial, where $j = 1$ and $j = 2$ are “learned” and “not learned” states, respectively. $B^c_{i,j}(t)$ is a basis function for a cardinal spline for neuron c, trial type j. Cardinal splines are locally defined third-order polynomial functions that flexibly approximate arbitrary smooth functions using a small number of basis functions [19]. Here, we use spline functions to capture the firing probability as a function of time relative to the picture cue. N is the number of spling control points used to fit the data. Here we chose $N = 16$ control points. $\theta = \left[\left\{\alpha^c_{i,j}\right\}_{i=1}^{N}\right]$ is a set of unknown parameters which relate the temporal features of the behavioral task to instantaneous spike rate.

It follows from the definition of the conditional intensity function that the probability of a spike from neuron c in a small time interval $[t, t + \Delta)$ is approximately:

$$\Pr(\text{Spike from neuron } c \text{ in} [t, t + \Delta) \mid \theta) \approx \lambda^c(t \mid \theta)\Delta . \quad (13)$$

This spiking intensity function describes a GLM for the spike train data. Such GLMs have a number of nice properties, including convexity of the likelihood surface and asymptotic normality of the parameter estimates, which allow us to compute maximum likelihood estimates for the model parameters in a straightforward manner. We fit these GLMs using the estimated learning state from the behavioral data. We examine the model fits to the data from 500 ms before picture cue to 2500 ms after picture cue.

Fig. 3 shows the model parameters and their uncertainty for the maximum likelihood fit to four example neurons in this spiking data. Each panel shows the spline estimates, in solid lines, and 95% confidence bounds in dashed lines, as a function of time relative to picture cue, represented by the vertical line in cyan. The times of go cue, feedback, and start of inter-trial interval are identified as vertical lines in green, yellow and black, respectively. The estimated intensity and 95% confidence bounds for the learned state and not-learned state are plotted in red and blue, respectively.

Top two and the lower-left panels in Fig. 3 show the model fit for three neurons in the caudate nucleus. For the neuron plotted in the top-left, at around 1000 ms after the picture cue and right before the go cue, the estimated intensity for learned trials in red is significantly higher than the estimated intensity for not-learned trials in blue. For the neuron plotted in the top-right, at around 2000 ms after the picture cue and within a 300 ms lag of the feedback, the estimated intensity for not-learned trials is significantly higher than that of the learned trials. For the neuron plotted in the lower-left, at around 800 ms after the picture cue, the estimated intensity for not-learned trials is significantly higher than that of the learned trials.

The lower-right panel in Fig. 3 shows the model fit for a neuron in the PFC. The 95% confidence bounds for the learned and not-learned trials are always overlapping, which means that the temporal spiking properties during the observation interval are not significantly different for the two learning states.

To assess the goodness-of-fit of the two-state model, we constructed Kolmogorov-Smirnov (KS) plots of time-rescaled inter-spike intervals (ISIs) [20]. The time-rescaling theorem produces a set of rescaled ISIs that are independent with an exponential distribution with mean 1 if the proposed model accurately describes the structure in the observed spiking activity. To construct the KS plot, we plot the empirical cumulative distribution of the rescaled ISIs against the theoretical cumulative distribution of the Exponential(1) distribution. The better the quality of the model fit, the closer the K-S plot should be to a 45 degree line [21]. The left panel in Fig. 4 shows an example K-S plot for the model fit. The model and empirical CDFs demonstrate a good overall fit, with some evidence of misfit in the smaller rescaled ISIs. This may suggest some model misfit related to our assumptions of

a dichotomized cognitive state and a lack of spiking history dependence structure in the point process model. More accurate modeling would likely lead to improvements in the oeverall goodness-of-fit and the resulting decoded state estimates.

D. Neural Decoding

The previous subsection focused on the construction of neural spiking models, which uses relevant covariates related to behavior to describe the statistical structure of neural spiking activity. In this subsection we present a simple recursive Bayesian algorithm to decode the dynamic state from the spiking activity. For each trial k, we compute the posterior distribution of the monkey's learning state given the combined spiking activity of the neural ensemble within the observation interval, $[0, T]$:

$$p\left(\overrightarrow{x}_k \mid \Delta N_{1:T}\right) \propto p\left(\Delta N_{1:T} \mid \overrightarrow{x}_k\right) p\left(\overrightarrow{x}_k\right). \quad (14)$$

Here $p\left(\overrightarrow{x}_k\right)$ is the prior distribution of the state. In this case, we choose a uniform prior for the binary state. The observation model, or likelihood, is given by

$$p(\Delta N_{1:T} \mid \overrightarrow{x}_k) \propto \prod_{t=1}^{T} \prod_{c=1}^{C} \left[\lambda^c(\overrightarrow{x}_k)\Delta\right]^{\Delta N_t} \exp(-\lambda^c(\overrightarrow{x}_k)\Delta). \quad (15)$$

We perform a classification procedure by thresholding the posterior of the state at various cut-off probabilities to determine whether the monkey is in a "learned" or a "not learned" state during a particular trial. The right panel in Fig. 4 shows the receiver-operating characteristic (ROC) curve [22], [23]. It plots the sensitivity of the cut-off, the probability of rejecting the null hypothesis when it is false, versus significance level, the probability of rejecting a null hypothesis when it is true. ROC curves using decoding results from the neural ensemble in PFC and caudate are plotted as a solid blue line and a dashed red line, respectively. The ROC curve based on spiking activity in the caudate is consistently above the ROC curve based on spiking activity in PFC. It shows that neural activity in the caudate provides a better decoding of the learning state than PFC, which corroborates previous findings in the literature that the caudate contributes more closely to learning [24]–[27].

IV. DISCUSSION

The classical two-model state-space paradigm has been successfully applied to relate behavior and neural activity directly in low-dimensional, directly observable data. However, when both the behavioral and neural data become high-dimensional and multi-faceted, this direct approach becomes computationally challenging. Here we proposed a new three-model paradigm to characterize the relationship between behavior and the neural activity. We first introduced a cognitive state process whose dynamics can be estimated from behavior. We then used the state and relevant covariates related to behavior to describe the neural activity. Lastly, we estimated the dynamic state from a combination of behavioral and neural data.s

We illustrated our paradigm with a specific example of two monkeys performing a temporally delayed, on-line learning task. We demonstrated that accurate decoding of the learning state is possible with a simple point process model of population spiking. Our analyses also allowed us to compare decoding accuracy across neural population in the PFC and caudate nucleus.

Immediate extensions to the application of the paradigm shown here are under active development. First, more accurate statistical descriptions of the behavioral data hopefully will lead to a more accurately estimated learning curve. Second, instead of working with the simplified, dichotomized learning state process, the neural encoding and decoding steps can deal directly with a continuous state process. Last, to improve the quality of fit, the point process models used for neural encoding can be expanded to include spiking history.

The essential goal of the proposed paradigm is to demonstrate the existence of meaningful relationships between complex behavior and high-dimensional neural activity. We achieve dimensionality-reduction by using hidden cognitive state processes to represent the relationship. In principle, identifying low-dimensional states that provide meaningful links between behavior and neural data is a major challenge that is specific to particular neural processing tasks. In many cases, neuroscientists may have prior conceptions of cognitive features that could represent such meaningful links. Future work may allow us to develop principled methodologies to identify cognitive states directly from data. Furthermore, by assigning some cognitive meaning to the hidden state, we can design experiments to determine the effect of manipulations of neural activity on cognitive influences of behavior. For the specific example shown here, the cognitive state can be thought of as a measure of learning flexibility, and we can modulate it to facilitate learning in the monkeys.

We envision the proposed paradigm to play a future role in the development of new types of closed-loop experiments, aiming to characterize causal relationships between neural activity and the behavior they encode. The proposed algorithm can allow investigators to identify and manipulate a low-dimensional correlate of cognitive influence in a content-specific way, altering neural activity related to certain cognitive features to modulate behavior. This may be an important step in treating mental diseases such as post-traumatic stress disorder and obsessive-compulsive disorders clinically.

References

[1]. Smith A and Brown EN, “Estimating a State-Space Model from Point Process Observations,” Neural Comput, vol. 15, pp. 965–991, 2003. [PubMed: 12803953]

[2]. Eden UT, Frank LM, Barbieri R, Solo V, and Brown EN, “Dynamic analysis of neural encoding by point process adaptive filtering,” Neural Comput, vol. 16, pp. 971–998, 2004. [PubMed: 15070506]

[3]. Truccolo W, Eden UT, Fellows MR, Dolabelghue JP, and Brown EN, “A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects,” J. Neurophysiol, vol. 94, pp. 1074–1089, 2005. [PubMed: 15872064]

[4]. Srinivasan L, Eden UT, Willsky AS, and Brown EN, “A state-space analysis for reconstruction of goal-directed movement using neural signals,” Neural Comput,vol. 18, pp. 2465–2494, 2006. [PubMed: 16907633]

[5]. Czanner G, Eden UT, Wirth S, Yanike M, Suzuki WA, and Brown EN, "Analysis of between-trial and within-trial neural spiking dynamics," J. Neurophysiol, vol. 99, pp. 2672–2693, 2008. [PubMed: 18216233]

[6]. Kemere C, Santhanam G, Yu BM, Afshar A, Ryu SI, Meng TH, and Shelabely KV, "Detecting neural-state transitions using hidden Markov models for motor cortical prostheses," J. Neurophysiol, vol. 100, pp.2441–52, 2008. [PubMed: 18614757]

[7]. Wu W, Kulkarni J, Hatsopoulos N, and Paninski L, "Neural decoding of goal-directed movements using a linear statespace model with hidden states," IEEE Trans. Neural Sys. and Rehab. Eng, vol. 17, label. 4, 2009.

[8]. Paninski L, Ahmadian Y, Ferreira DG, Koyama S, Rad KR, Vidne M, Vogelstein J, Wu W, "A new look at state-space models for neural data," J. Comput. Neurosci, vol. 29, pp. 107–126, 2010. [PubMed: 19649698]

[9]. Huang Y, Brandon MP, Griffin AL, Hasselmo ME, and Eden UT, "Decoding movement trajectories through a T-maze using point process filters applied to place field data from rat hippocampal region CA1," Neural Comput, vol. 21, pp. 3305–34. [PubMed: 19764871]

[10]. Koyama S, Eden UT, Brown EN,, and Kass RE, "Bayesian decoding of neural spike trains," Ann. Inst. Stat. Math, vol. 62, pp. 37–59, 2010.

[11]. Asaad WF and Eskandar EN, "Encoding of both positive and negative reward prediction errors by neurons of the primate lateral prefrontal cortex and caudate nucleus," J. Neurosci, vol. 31, pp. 17772–87, 2011. [PubMed: 22159094]

[12]. Smith AC, Frank LM, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel AM, Suzuki WA, and Brown EN, "Dynamic analysis of learning in behavioral experiments," J. Neurosci, vol. 24, pp. 447–461, 2004, [PubMed: 14724243]

[13]. Smith AC, Stefani MR, Moghaddam B, and Brown EN, "Analysis and design of behavioral experiments to characterize population learning," J. Neurophysiol, vol. 93, pp. 1776–1792, 2005. [PubMed: 15456798]

[14]. Suzuki WA and Brown EN, "Behavioral and neurophysiological analyses of dynamic learning processes," Behavioral and Cognitive Neuroscience Reviews, vol. 4, pp. 67–95, 2005. [PubMed: 16251726]

[15]. Prerau MJ, Smith AC, Eden UT, Kubota Y, Yanike M, Suzuki W, Graybiel AM, and Brown EN, "Characterizing learning by simultaneous analysis of continuous and binary measures of performance," J. Neurophyisiol, vol. 102, pp. 3060–3072, 2009.

[16]. Brillinger DR. "Nerve cell spike train data analysis: a progression of technique," J. Amer. Statist. Assoc, vol. 87, pp. 260–271, 1992.

[17]. Cox DR and Isham V, Point Processes. Chapman and Hall, London, 1980.

[18]. Daley D and Vere-Jones D, An Introduction to the Theory of Point Processes. Springer-Verlag, New York, 1988.

[19]. Ramsay JO and Silverman BW, Functional Data Analysis. Springer-Verlag, New York, 2010.

[20]. Brown EN, Barbieri R, Ventura V, and Kass RE, "The time-rescaling theorem and its application to neural spike train data analysis," Neural, Comput, vol. 14, pp. 325–46, 2001.

[21]. Johnson A and Kotz S, Distributions in Statistics: Continuous Univariate Distributions, Wiley, New York, 1980.

[22]. Green DM and Swets JA, Signal Detection Theory and Psychophysics. John Wiley & Sons, New York, 1966.

[23]. Zweig MH and Campbell G, "Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine," Clinical Chemistry, vol. 39, pp. 561–577, 1993. [PubMed: 8472349]

[24]. Packard MG and McGaugh JL, "Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning, Neurobiol. Learn Mem, vol. 65, pp. 65–72, 1996. [PubMed: 8673408]

[25]. Wilabelcur G and Eskes G, "Prefrontal cortex and caudate nucleus in conditional associative learning: dissociated effects of selective brain lesions in rats," Behav. Neurosci, vol. 112, pp. 89–101, 1998. [PubMed: 9517818]

[26]. Harulabel M, Kuroda T, Doya K, Toyama K, Kimura M, Samejima K, Imamizu H, and Kawato M. “A neural correlate of reward-based behavioral learning in caudate nucleus: a functional magnetic resonance imaging study of a stochastic decision task,” J. Neurosci, vol. 24, pp. 1660–5, 2004. [PubMed: 14973239]

[27]. Seger CA and Cincotta CM, “The roles of the caudate nucleus in human classification leanring,” J. Neurosci, vol. 25, pp. 2941–51, 2005. [PubMed: 15772354]

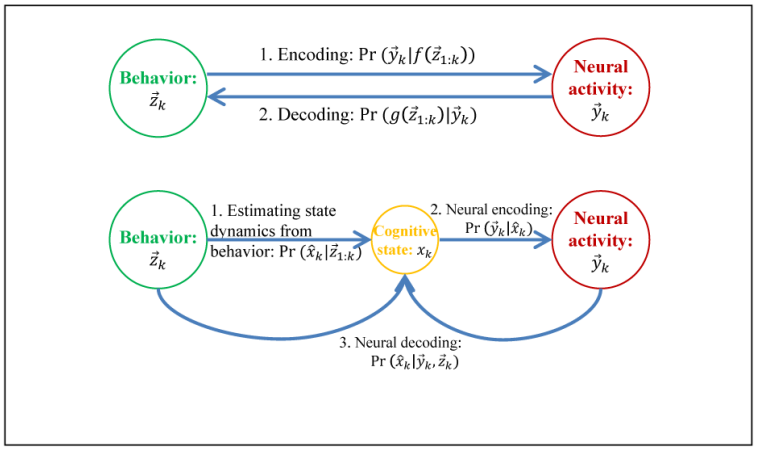

Fig. 1.

Schematic representation of the general three-model cognitive state-space paradigm (*Lower*) in comparison with the classical two-model state-space paradigm (*Upper*).

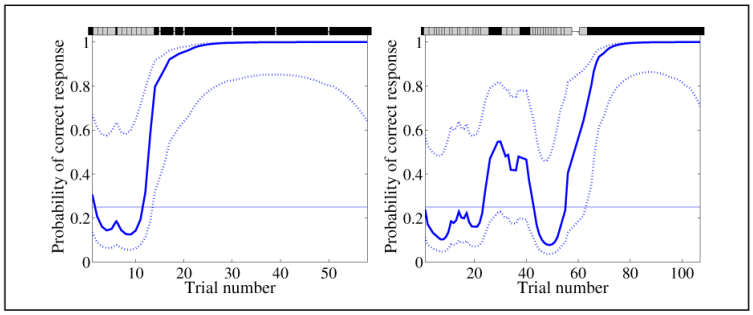

Fig. 2.

Two examples of the EM algorithm applied in the analysis of the state dynamics in a rule-switching task. The correct and incorrect responses are shown, respectively, by black and gray marks above the panels. The probability of a correct response occurring by chance is 0.25 in horizontal line. Solid blue lines are the learning curve estimates, and the dotted blue lines are the associated 95% confidence intervals.

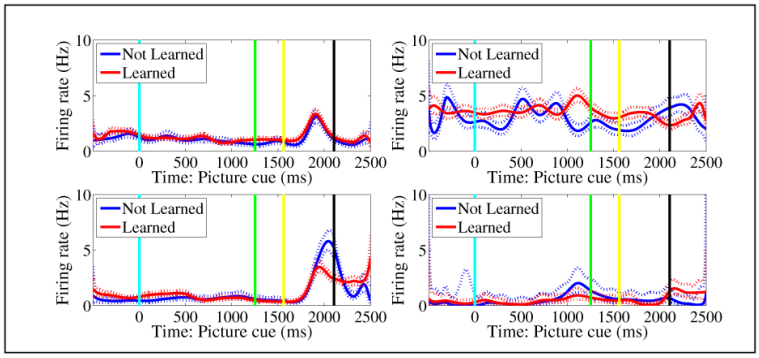

Fig. 3.

Four examples of the estimated conditional intensity with 95% confidence bounds as a function of the temporal features in the task. In each panel, blue and red solid line is the estimated conditional intensity for the “not learned” and “learned” trials, respectively. The vertical lines in cyan, green, yellow and black are the time for picture cue, go cue, feedback, and start of the inter-trial interval, respectively.

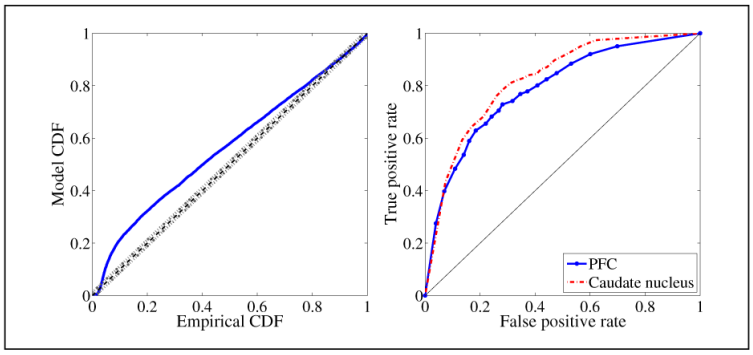

Fig. 4.

Left: An example Kolmogorov-Smirnov (KS) plot of time-rescaled inter-spike intervals for the estimated point process neural encoding model. It plots the empirical cumulative distribution of the rescaled ISIs against the theoretical cumulative distribution of the Exponential(1) distribution. *Right*: An example receiver-operating characteristic (ROC) curve. It plots sensitivity of the cut-off, the probability of rejecting the null hypothesis when it is false, versus significance level, the probability of rejecting a null hypothesis when it is true. ROC curves using decoding results from the neural ensemble in PFC and caudate are plotted in solid blue lines and dashed red line, respectively.