

# Chapter 7

## Characterizing Complex Human Behaviors and Neural Responses Using Dynamic Models

Sridevi V. Sarma and Pierre Sacré

### 7.1 Introduction

Many experiments conducted in neuroscience entail applications of stimuli and recordings of behavioral responses and neural activity. Traditional approaches to understanding how the brain encodes stimuli often compute correlations between stimuli and neural activity time-locked to behavioral events. For example, when studying motor control, investigators train a participant to move the arm in different directions while activity from premotor and primary motor regions are measured (Carpenter et al. 1999; Schieber 2004). Then, to understand how neurons encode movement direction, firing rates of neurons are modeled as functions of behavior right after the onset of movement (Agarwal et al. 2015). In this example, both direction of movement and neural activity are measured outputs, and behavior is primarily driven by a target cue.

Now, let's consider experiments wherein behavior is not only driven by stimuli provided by the experimentalist, but also by internal factors within the participant that are not easily measurable. A first example is when participants are performing a gambling task, wherein they are betting virtual money and then perhaps get emotional if they are winning or losing (Sacré et al. 2016a,b,c). Although objective measures of emotion have been proposed such as skin conductance response and heart rate variability, these measures are typically delayed or only accurate over several minutes, while emotions can fluctuate at a faster time scale during gambling (Mauss and Robinson 2009). A second example is when participants are performing a Stroop-like task, wherein distractors are present to confuse participants while they attempt to make correct associations between presented stimuli and

---

S.V. Sarma (✉) · P. Sacré  
Johns Hopkins University, Baltimore, MD, USA  
e-mail: [ssarma2@jhu.edu](mailto:ssarma2@jhu.edu); [psacre1@jhu.edu](mailto:psacre1@jhu.edu)

appropriate responses (Shoham et al. 2003; Smith et al. 2015). During this task, participants make errors and their motivational and attentional states vary over the session. Motivation and attention are not directly measurable, yet they may influence behavior in a profound way (Shoham et al. 2003).

In both of the gambling and Stroop-like tasks described above, behavioral responses and neural activity are influenced by external stimuli *and* internal states of participants. Thus, when looking for neural correlates, how behavior changes with stimuli and underlying dynamic state variables must first be characterized. In this chapter, we present a systematic approach based on existing methodologies to (1) estimate internal dynamic states of participants from measured data to explain behavior variability within and across participants, and (2) identify neural substrates of behavior and internal states. The proposed approach is a two-step procedure wherein one first constructs participant-specific state-space models that capture the dynamics of internal states and how they evolve with administered stimuli, and how measured behavior depends on these states and stimuli; and then, one relates stimuli, responses, and states back to neural activity. We discuss the challenges that arise in each step of the process and provide suggestions on how to successfully complete these two steps. We present examples from two data sets involving a gambling and Stroop-like task.

## 7.2 Methods

In this section, we first describe a general dynamic state-space modeling framework and the maximum likelihood procedure used to estimate parameters of participant-specific models of behavior. Then, we discuss how to map model variables and estimated internal states back to neural data using nonparametric statistical tests and point process models (PPMs).

### 7.2.1 Dynamic State-Space Modeling

The first step of the proposed approach is to build a mathematical model with inputs  $\mathbf{u}$  and outputs  $\mathbf{y}$  that explains the variability that we observe in the data. In this context, we can distinguish between two types of models. A model is *static* (or without memory) if the value of the output signal at a particular time depends only on the value of the input signal at the same time. Otherwise, it is *dynamic* (or with memory).

The general *dynamic* state-space model for a discrete-time system can be written as follows

$$\mathbf{x}_{k+1} \sim f_{\theta}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k), \quad (7.1a)$$

$$\mathbf{y}_k \sim h_{\theta}(\mathbf{y}_k \mid \mathbf{x}_k, \mathbf{u}_k), \quad (7.1b)$$

where  $\mathbf{x}_k \in \mathbb{R}^n$  is the  $n$ -dimensional state-vector,  $\mathbf{u}_k \in \mathbb{R}^m$  is the  $m$ -dimensional input-vector, and  $\mathbf{y}_k \in \mathbb{R}^p$  is the  $p$ -dimensional output-vector. The state-transition map  $f_\theta$  and the measurement map  $h_\theta$  are conditional probability distributions of  $\mathbf{x}_{k+1}$  given  $(\mathbf{x}_k, \mathbf{u}_k)$ , and of  $\mathbf{y}_k$  given  $(\mathbf{x}_k, \mathbf{u}_k)$ , respectively. The initial state vector  $\mathbf{x}^0$  is distributed according to  $p_\theta(\mathbf{x}^0)$ . The model parameters are denoted by  $\theta$ . It is often convenient to define its equivalent *static* model, that is,  $\mathbf{y}_k \sim \tilde{h}_\theta(\mathbf{y}_k | \mathbf{u}_k) = h_\theta(\mathbf{y}_k | 0, \mathbf{u}_k)$ , where the state-vector is fixed to zero for all  $k$ .

For our neuroscience applications, the inputs  $\mathbf{u}_k$  represent stimuli (or functions of stimuli) of the task on trial  $k$ , the outputs  $\mathbf{y}_k$  represent measured behavioral responses (or functions of responses) in the task (e.g., reaction time, correct/incorrect answer) on trial  $k$ , and the states  $\mathbf{x}_k$  represent internal states on trial  $k$  that influence behavior (e.g., attentional state, motivation, emotion).

As a first step to model behavioral data, it often is sufficient to begin with a time-invariant state-space model with a linear state equation and a generalized linear output equation, which reduces to

$$\mathbf{x}_{k+1} = \mathbf{A} \mathbf{x}_k + \mathbf{B} \mathbf{u}_k + \mathbf{w}_k, \quad (7.2a)$$

$$\mathbf{y}_k \sim h(\mathbf{y}_k | \mathbf{C} \mathbf{x}_k + \mathbf{D} \mathbf{u}_k, \boldsymbol{\tau}), \quad (7.2b)$$

where  $\mathbf{w}_k \in \mathbb{R}^n$  is the  $n$ -dimensional zero-mean Gaussian noise-vector with unknown covariance matrix  $\boldsymbol{\Sigma}_w$  and  $h$  is a probability distribution from the exponential family, that is conditioned on an affine combination of states and inputs and the dispersion parameter  $\boldsymbol{\tau}$ . The initial state vector  $\mathbf{x}^0$  is assumed to follow a Gaussian distribution with mean  $\bar{\mathbf{x}}^0$  and covariance matrix  $\boldsymbol{\Sigma}^0$ . The model parameters are then  $\theta = \{\mathbf{A}, \mathbf{B}, \boldsymbol{\Sigma}_w, \bar{\mathbf{x}}^0, \boldsymbol{\Sigma}^0, \mathbf{C}, \mathbf{D}, \boldsymbol{\tau}\}$ .

The model estimation problem then boils down to: given  $N$  input-output measurements  $\mathbf{u}_{1:N} = \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$  and  $\mathbf{y}_{1:N} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ , estimate the model parameters  $\theta$  and the state  $\mathbf{x}_{1:N} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ . One approach is to estimate  $\theta$  and  $p_\theta(\mathbf{x}_{1:N} | \mathbf{u}_{1:N}, \mathbf{y}_{1:N})$  from data to maximize the likelihood function (Van Trees 1968; Louis 1991; Moon 1996). The likelihood function is the family of probability distributions considered as a function of  $\theta$ , for fixed  $\mathbf{y}_{1:N}$  and  $\mathbf{u}_{1:N}$ . It is often more convenient to work with its logarithm, which is called the log-likelihood function, and denoted as  $\ell$ :

$$\ell(\theta) = \log p_\theta(\mathbf{y}_{1:N} | \mathbf{u}_{1:N}). \quad (7.3)$$

Now, the problem is to estimate the value of the parameters  $\theta$ . A widely used method, called maximum likelihood estimation, is to estimate  $\theta$  as

$$\hat{\theta}_{\text{ml}} = \arg \max_{\theta \in \Theta} p_\theta(\mathbf{y}_{1:N} | \mathbf{u}_{1:N}) = \arg \max_{\theta \in \Theta} \ell(\theta), \quad (7.4)$$

where  $\theta \in \Theta$  gives the prior information or other constraints on the parameter vector  $\theta$ . In the context of the estimation of a dynamic model, the state is not observed and we can write the likelihood as follows

$$p_\theta(\mathbf{y}_{1:N} | \mathbf{u}_{1:N}) = \int_{\mathcal{X}} p_\theta(\mathbf{x}_{1:N}, \mathbf{y}_{1:N} | \mathbf{u}_{1:N}) d\mathbf{x}. \quad (7.5)$$

One way to solve this problem is to use the expectation-maximization (EM) algorithm. The EM algorithm is an iterative algorithm that is composed of two steps at each iteration: an expectation step and a maximization step.

**E-step** The idea of the E-step is to take the expectation with respect to the unknown underlying states, using the current estimate of the parameters  $\theta^*$  and conditioned upon the observation, that is,

$$Q(\theta \mid \theta^*) = \mathbf{E} \left[ \log \left( \frac{p_\theta(\mathbf{x}_{1:N}, \mathbf{y}_{1:N} \mid \mathbf{u}_{1:N})}{p_{\theta^*}(\mathbf{x}_{1:N}, \mathbf{y}_{1:N} \mid \mathbf{u}_{1:N})} \right) \mid \mathbf{U}_{1:N} = \mathbf{u}_{1:N}, \mathbf{Y}_{1:N} = \mathbf{y}_{1:N}, \theta^* \right], \quad (7.6)$$

$$= \int_{\mathcal{X}} \log \left( \frac{p_\theta(\mathbf{x}_{1:N}, \mathbf{y}_{1:N} \mid \mathbf{u}_{1:N})}{p_{\theta^*}(\mathbf{x}_{1:N}, \mathbf{y}_{1:N} \mid \mathbf{u}_{1:N})} \right) p_{\theta^*}(\mathbf{x}_{1:N} \mid \mathbf{u}_{1:N}, \mathbf{y}_{1:N}) d\mathbf{x}. \quad (7.7)$$

**M-step** The idea of the M-step is to provide a new estimate  $\theta^{**}$  of the parameters, that is,

$$\theta^{**} = \arg \max_{\theta \in \Theta} Q(\theta \mid \theta^*). \quad (7.8)$$

Finally, there are several ways to establish the degree of agreement between the model and observed data. In particular, they are mainly two families of statistics that we can compute: the first family measures the *goodness-of-fit* of the model with the data and the second family measures the *improvement of goodness-of-fit* from a static model to a dynamic model. In both families, we can use different test statistics such as log-likelihood, deviance and Pearson residuals, predictive power, etc. The first family is interested in the absolute value of this statistic; while the second family is interested in the relative difference between the statistics for the dynamic and static models. The statistical significance of these test statistics can be evaluated using a nonparametric permutation test.

## 7.2.2 Neural Correlates Informed by State-Space Model

The second step of the proposed approach is to relate the model variables (inputs, outputs, and estimated states) back to the neural data. Below we describe this second step if one has recorded continuous neural activity or spike train observations.

### 7.2.2.1 Continuous Neural Data

If neural activity measured is continuous (e.g., local field potential, electroencephalogram, electrocorticography), then a common approach to analyzing the data is to move to the spectral domain. In particular, select a time window of interest

(e.g., 500-ms window time-centered to an applied stimulus). Then, compute a spectrogram including frequencies of interest (e.g., 2–150 Hz) for that time window for each trial. This will generate a stack of spectrograms, one for each trial in the recorded session.

Once the stack of spectrograms is computed, identify time-frequency clusters within the stack that distinguish between two conditions of interest (e.g., high bet vs low bet, or moving up vs moving down). In particular, take a model variable of interest (e.g., player's card) and split trials into those when the variable takes on low values and those when the variable takes on high values. Low-value trials may be defined, for example, as the bottom third of the variable distribution over all trials, and high-value trials may be defined as the top third. Then apply a cluster-based nonparametric statistical test to leverage the dependency between adjacent time-frequency windows in order to avoid over-penalizing with multiple comparison corrections (Maris and Oostenveld 2007).

For each window in the spectrogram, create a null distribution by shuffling the condition labels 1000–5000 times between trials within each participant. Within each shuffle, compute a  $t$ -statistic and a  $p$ -value for each window of the newly labeled spectrograms (independent two-sample  $t$ -test with both tails, unequal sample sizes, and unequal variances). Clusters are formed by grouping windows with significant  $p$ -values (e.g.,  $p < 0.05$ ) that are adjacent in either time or frequency. The cluster-level test statistic is then calculated by taking the sum of absolute values of the  $t$ -statistics for each window in the cluster. This prioritizes clusters that have both strong differences and large sizes. A null distribution of cluster statistics is created using the same process but with the 1000–5000 spectrograms obtained from the originally shuffled labels. The observed cluster statistic is then compared against this null distribution of cluster statistics in order to obtain the final  $p$ -value of the test.

Data from all patients can be pooled together but the labels are permuted within each participant only. This process of finding time-frequency cluster correlated to a model or task variable can be repeated for each variable and the estimated state trajectories across participants.

### 7.2.2.2 Spike Train Data

If neural activity is measured as spike trains, then one can use point process models to identify how behavior influences neuronal spiking activity. Several examples of how PPMs are estimated used for different experimental setups are given in Coleman and Sarma (2007), Coleman and Sarma (2010), Santaniello et al. (2010), Santaniello et al. (2012), Sarma et al. (2010), Sarma et al. (2012).

A point process is a series of 0/1 random events that occur in continuous time. For a neural spike train, the 1's are individual spike times and the 0's are the times at which no spikes occur. To define a point process model of neural spiking activity, in this analysis, one can consider an observation interval  $(0, T]$ , and let  $N_k(t)$  be the number of spikes counted in the interval  $(0, t]$  for  $t \in (0, T]$  for a trial  $k$ .

A point process model of a neural spike train can be completely characterized by its conditional intensity function (CIF)  $\lambda_k(t | H_t)$  defined as follows:

$$\lambda_k(t | H_t, \mathbf{u}_k, \mathbf{x}_k) = \lim_{\Delta \rightarrow 0} \frac{\Pr(N_k(t + \Delta) - N_k(t) = 1 | H_t, \mathbf{u}_k, \mathbf{x}_k)}{\Delta}, \quad (7.9)$$

where  $H_t$  denotes the history of spikes and covariates up to time  $t$ . It follows from (7.9) that the probability of a single spike in a small interval  $(t, t + \Delta]$  is approximately

$$\Pr(\text{spike in } (t, t + \Delta] \text{ on trial } k | H_t, \mathbf{u}_k, \mathbf{x}_k) = \lambda_k(t | H_t, \mathbf{u}_k, \mathbf{x}_k) \Delta. \quad (7.10)$$

Details can be found in Cox and Isham (1980), Snyder and Miller (1991).

The CIF generalizes the rate function of a Poisson process to a rate function that is history dependent. Because the CIF completely characterizes a spike train, defining a model for the CIF defines a model for the spike train (Brown et al. 2003).

For neural correlate analyses, use a generalized linear model (GLM) to define CIF models by expressing for each neuron, the log of its CIF in terms of the neurons spike history  $H_t$ , relevant model inputs  $\mathbf{u}_k$ , and the state trajectory  $\mathbf{x}_k$ . The GLM is an extension of the multiple linear regression model, in which the variable being predicted (e.g., in this case spike times) need not be Gaussian (McCullagh and Nelder 1989). GLM also provides an efficient computational scheme for model parameter estimation and a likelihood framework for conducting statistical inferences (McCullagh and Nelder 1989).

One can express the CIF for each neuron at each time step (e.g., millisecond) as a function of task stimuli which can turn on and/or off over time, the state variable value which typically is constant over a trial, and the neuron's spiking history. Instead of estimating the CIF continuously throughout the entire trial, one can estimate it over time windows around key epochs and at discrete time intervals each 1 ms in duration.

In particular, one can express the CIF as follows:

$$\lambda_k(t | H_t, \mathbf{u}_k, \mathbf{x}_k, \boldsymbol{\theta}) = \lambda^S(\mathbf{u}_k | \boldsymbol{\theta}) \lambda^X(\mathbf{x}_k | \boldsymbol{\theta}) \lambda^H(t | H_t, \boldsymbol{\theta}) \quad (7.11)$$

where  $\lambda^S(\mathbf{u}_k | \boldsymbol{\theta})$  describes the effect of the stimulus on the neural response,  $\lambda^X(\mathbf{x}_k | \boldsymbol{\theta})$  describes the effect of the state variable on the neural response, and  $\lambda^H(t | H_t, \boldsymbol{\theta})$  describes the effect of spiking history on the neural response.  $\boldsymbol{\theta}$  is a parameter vector to be estimated from data. The units of  $\lambda^S(t | \boldsymbol{\theta})$  is spikes per second and  $\lambda^H(t | H_t, \boldsymbol{\theta})$  is dimensionless. Finally, one can compute maximum-likelihood (ML) estimates for  $\boldsymbol{\theta}$  and 95% confidence intervals of  $\boldsymbol{\theta}$  for each neuron using `glmfit` in MATLAB.

It is important to establish the degree of agreement between a PPM and observations of the spike train and associated experimental variables is a prerequisite for using the point process analysis to make scientific inferences. One can use Kolmogorov-Smirnov (KS) plots based on the time-rescaling theorem to assess

the model goodness-of-fit. The time-rescaling theorem is a well-known result in probability theory, which states that any point process with an integrable CIF may be transformed into a Poisson process with unit rate (Johnson and Kotz 1970). A KS plot, which plots the empirical cumulative distribution function of the transformed spike times versus the cumulative distribution function of a unit rate exponential, is used to visualize the goodness-of-fit for each model. The model is better if its corresponding KS plot lies near the 45° line. One can compute the 95% confidence bounds for the degree of agreement using the distribution of the KS statistic (Johnson and Kotz 1970). If a model's KS plot was within the 95% confidence bounds, then it can be included in the analyses.

## 7.3 Results

In this section, we present two applications where we applied our approach to reveal new insights on the neural mechanisms involved in a Stroop-like task where spike train observations were made and a gambling task where local field potentials were measured.

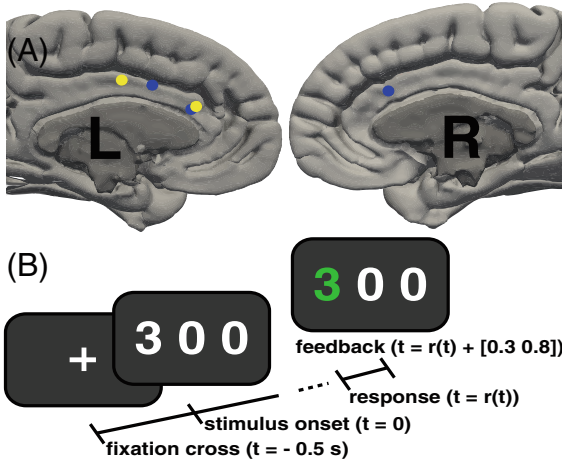
### 7.3.1 *Multi-Source Interference Task*

This example is taken from Sklar et al. (2017). Two participants being treated at the Columbia University Medical Center performed the behavioral task in their hospital rooms using methods previously described in Johnson et al. (2014). Behavioral data were simultaneously acquired on the same time base as the electrophysiology data. Participants performed the multi-source interference task (MSIT) (Shoham et al. 2003). The MSIT is a Stroop-like task in which the participant is presented with three integers ranging from 0 to 3. Two of the three integers presented are the same integer. The goal of the MSIT is to indicate the identity of the different integer on the number pad (e.g., cue: 0 2 0; correct response: button 2; Fig. 7.1b).

Conflict is introduced in this task by changing the position of the target number (e.g., 0 0 1; correct response: button 1; Simon or spatial interference) or by changing the identity of the distracting integers to potential responses (e.g., 1 2 1; correct response: button 2; Eriksen or flanker interference). Additionally, both types of interference can occur (e.g., 3 1 3; correct response: button 1). These four groups of trials were presented randomly, with a uniform frequency distribution.

#### 7.3.1.1 **Dynamic State-Space Modeling**

In this study, we hypothesized that the “cognitive state” of each participant influences behavior and modulates neuronal activity in the dorsal anterior cingulate cortex (dACC). In particular, we hypothesized that when participants require more



**Fig. 7.1** MSIT task and microwire recording locations. (a) Microwire recording locations in the dACC. Colors represent recording locations corresponding to each participant on each hemisphere (L and R). (b) MSIT task diagram showing an example trial structure. In each trial, a fixation cross appears on the screen for 0.5 s prior to the stimulus presentation. The stimulus remains on the screen until the participant indicates her response on the button pad. Feedback is delivered between 0.3 and 0.8 s after the participant indicates her response. Figure reproduced with permission from Sklar et al. (2017)

cognitive control, (1) they are more likely to react to the stimulus slowly and (2) their cingulate neurons are modulated. Since such a cognitive state is not directly measurable, we compute it from measurable data.

Before constructing the state-space model of behavior, we first looked to see whether behavior varied for different stimuli, and for the same stimuli over the session. To examine behavioral variability, we plotted a moving average reaction time for each stimulus type (easy, hard, Simon, Flanker). As shown in Fig. 7.2, the reaction times for each stimulus type change over time, suggesting dynamics in the behavior that may be explained by a latent state variable.

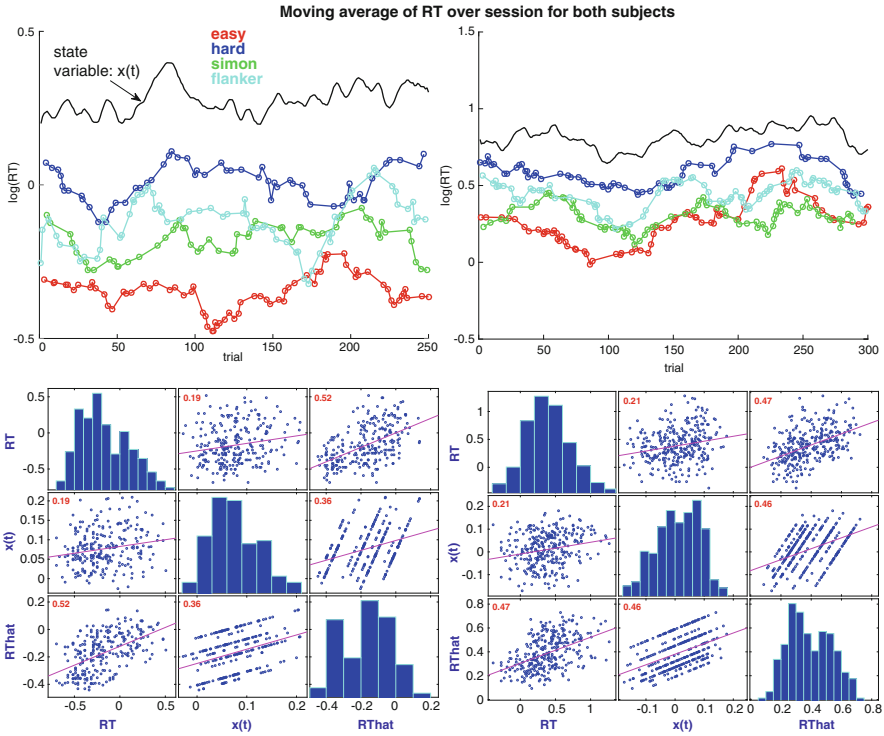
Therefore, we constructed a cognitive state variable  $x_k$  that updates for each trial  $k$  as follows:

$$x_{k+1} = a x_k + \sum_{i=1}^5 b_i u_{i,k} = a x_k + \mathbf{B} \mathbf{u}_k \quad (7.12)$$

where  $\mathbf{u}_k = [u_{1,k}, u_{2,k}, u_{3,k}, u_{4,k}, u_{5,k}]^T$  is an input column vector dependent on the trial conditions:

- $u_{1,k} = 1$  if no interference on trial  $k$ , and 0 otherwise;
- $u_{2,k} = 1$  if both interferences on trial  $k$ , and 0 otherwise;
- $u_{3,k} = 1$  if spatial interference on trial  $k$ , and 0 otherwise;
- $u_{4,k} = 1$  if flanker interference on trial  $k$ , and 0 otherwise;
- $u_{5,k} = 1$  if trial type on trial  $k$  changed from previous trial, and 0 otherwise.





**Fig. 7.2** Relationship between reaction time variability and cognitive state. (Top) Moving average of reaction times for each trial type for participant 1 (left) and participant 2 (right). The estimated cognitive states are overlaid in black. (Bottom) Correlation plots between actual reaction times and  $x_k$  and predicted reaction times for participant 1 (left) and participant 2 (right). Figure reproduced with permission from Sklar et al. (2017)

The parameter  $a$  represents the decaying influence of previous trials on the cognitive state, and  $\mathbf{B} = [b_1, b_2, b_3, b_4, b_5]$  dictate the effects that the trial conditions have on the state  $x_{k+1}$ . The solution to the state-space equation is

$$x_k = a^k x_1 + \sum_{i=1}^5 \sum_{s=1}^{k-1} a^{(k-s-1)} b_i u_{i,s}, \tag{7.13}$$

which can be used to determine the parameters  $\{a, \mathbf{B}\}$  by inserting the solution  $x_k$  as a covariate into a GLM. The output of the GLM is  $y_k$ , defined as the log of the reaction time modeled as

$$y_k = \log(r_k) = x_k + \mathbf{D} \mathbf{u}_k + d_0 + \epsilon_k, \tag{7.14}$$

where  $r_k$  is the reaction time of the trial, the  $\epsilon_k$  are independent zero-mean Gaussian random inputs with variance  $\sigma_\epsilon^2$ , and  $\mathbf{D}$  is a vector of the form of  $\mathbf{B}$ , that represents the direct influence of the current input on the reaction time of a trial.

The state-space model above includes a state that is completely deterministic, and the output is stochastic. If the state is deterministic, then the EM algorithm is unnecessary. It is worth beginning with a deterministic state variable to help identify model structure (what inputs to include into the state-space model) that best explains the observed data.

To estimate the parameters  $a$ ,  $\mathbf{B}$ ,  $\mathbf{D}$ , and  $d_0$  of the state-space model, we gridded the parameter space  $a$  and for each parameter value, (1) we computed each term of the sum of the state trajectory (7.13) for  $b_i = 1$  over the session, then (2) we substituted  $x_k$  by each term of the sum in a GLM and estimated  $\mathbf{B}$  and  $\mathbf{D}$  that maximize the data likelihood function. We then selected  $\theta = \{a, \mathbf{B}, \mathbf{D}, d_0\}$  that produce the maximum of all likelihoods over the entire grid.

Figure 7.2 overlays the estimate state variables (black trajectories) for the two participants. The state trajectories follow the dynamics of mean reaction times over the session for one or more stimuli. For participant 1 (left panel), the estimated  $x_k$  attempts to capture the variability of reaction times over the session for all four task types, but is not able to characterize behavior for all stimuli. In this case, a second state variable may better explain the behavior. On the other hand, the reaction time dynamics for participant 2 (right panel) are very similar across all stimuli suggesting that a scalar state variable is sufficient to explain participant 2's variability in behavior.

The bottom panels in Fig. 7.2 show the correlation plots between actual reaction times and  $r_k$  and predicted reaction times,  $\hat{r}_k$  for participant 1 (left) and participant 2 (right), where  $\log(\hat{r}_k) = x_k + \hat{\mathbf{D}}\mathbf{u}_k + \hat{d}_0$ . The state-space models for both participants suggest that the inclusion of the state helps explain the variability in reaction times over the session that cannot be entirely explained with task stimuli that changes over the session.

### 7.3.1.2 Neural Correlates Informed by Dynamic State-Space Model

Now that behavior is sufficiently explained by the state-space model described above, we search to explain neuronal responses to both task stimuli and cognitive state estimates. We thus formulated a PPM to relate the spiking of each dACC neuron for each participant to factors associated with the neuron's spiking history and the cognitive state variable. We use these model parameters to analyze temporal dynamics in neuronal activity due to the cognitive state variable after the stimulus is shown.

As described in Sect. 7.2.2.2, we use the GLM framework to define the CIFs of our PPMs by expressing, for each neuron, the log of its CIF in terms of the neuron's spike history and relevant covariates (Truccolo et al. 2005). We express the CIF for each neuron as a function of the neuron's spiking history,  $\lambda_k^H$ , in the preceding 240 ms and our derived cognitive state variable,  $\lambda^X$ . Specifically, for trial  $k$  and time bin  $t$ :

$$\lambda_k(t | H_t, \theta) = \lambda^X(x_k | \theta) \lambda^H(t | H_t, \theta), \quad (7.15)$$

such that

$$\lambda^X(x_k | \boldsymbol{\theta}) = \alpha x_k \quad (7.16)$$

and

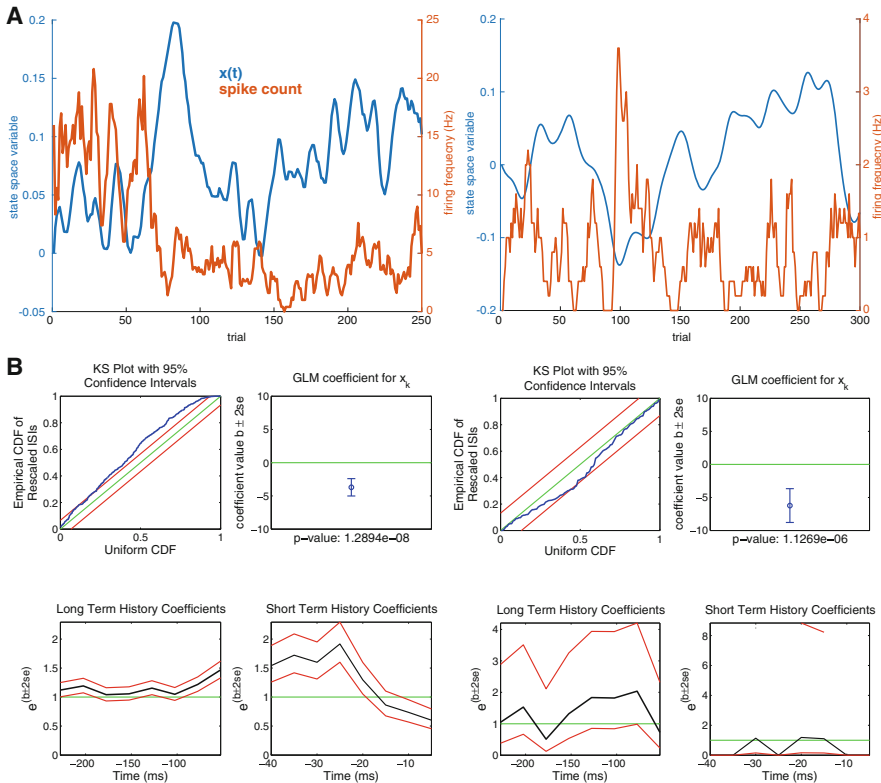
$$\log(\lambda^H(t | H_t, \boldsymbol{\theta})) = \gamma_0 + \sum_{j=1}^8 \gamma_j n_{t-5j:t-5(j-1)} + \sum_{j=1}^8 \beta_j n_{t-40-25j:t-40-25(j-1)}, \quad (7.17)$$

where  $n_{A:B}$  is the number of spikes observed in the time interval  $[A, B)$  during the epoch analyzed. The  $\{\gamma_j\}$  coefficients capture short-term history effects going back to 40 ms in the past in 5-ms bins. The  $\{\beta_j\}$  coefficients capture long-term history effects going back to 240 ms in the past in 25-ms bins, and  $\alpha$  captures the effect of the cognitive state. We computed ML estimates for all coefficients and their associated 95% confidence intervals for each neuron model using `glmfit` in MATLAB.

We examined the activity of 12 units (10 in patient 1 and 2 in patient 2). Figure 7.3a shows the spiking frequency in two units from both participants during the first second after stimulus presentation in each trial. These spike counts are overlaid with the cognitive state variables  $x_k$  for each participant. The neurons' spiking frequencies appear to have a negative correlation with  $x_k$  dropping markedly when  $x_k$  rises.

In participant 1, we found 10 units whose activity was predicted using the cognitive state variable (the GLM fit coefficient for the  $x_k$  covariate was significantly non-zero, with  $p < 0.05$ ). Some units had increased activity as  $x_k$  increased, while some displayed decreased activity. The covariate coefficients for participant 1's PPM are shown in Fig. 7.3b. The dependence on short- and long-term spiking history is displayed in the lower two plots and shows refractoriness in the first 15 ms after a spike, and an increased likelihood to fire in the 25–100 ms interval. The upper right plot displays the PPM coefficient for  $x_k$  with 2 standard deviation error bars. For this unit,  $x_k$  was a strong predictor of the spiking behavior, with spiking probability decreasing for higher  $x_k$ . The upper left plot shows the goodness of fit of the model using a KS plot, with 95% confidence bars. In participant 2, two units' spiking could be significantly predicted by  $x_k$  (see Sklar et al. (2017) for details). The PPM coefficients and goodness of fit for participant 2 for one neuron are also shown in Fig. 7.3b. This unit had significantly longer inter-spike-intervals, so the coefficient values for the short-term history bins have higher uncertainty.

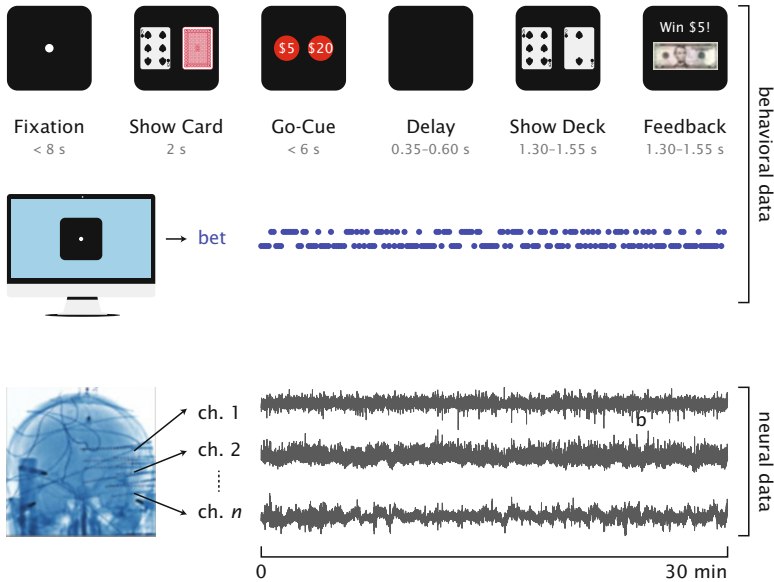
These preliminary results suggest that neurons in the dACC slowly track subjects' overall need for cognitive control, while simultaneously maintaining faster task-related dynamics. A latent cognitive state variable correlates with both reaction times and neuronal activity in two patients. These results provide support for an additional representation of task state or attentional motivation in the dACC.



**Fig. 7.3** Neural correlates of cognitive states. (a) Cognitive states over sessions overlaid with spike counts for one unit for participant 1 (left), and one unit for participant 2 (right). (b) Point process model for same unit as above from participant 1 (the 1st and 2nd columns) for first second after stimulus onset, and same unit as above from participant 2 (the 3rd and 4th columns). Top left: KS plot. Top right: coefficient for  $x_k$ . Bottom left: long-term history coefficients with 95% confidence bounds. Bottom right: short-term history coefficients with 95% confidence bounds. Figure reproduced with permission from Sklar et al. (2017)

### 7.3.2 Gambling Task

This example is taken from Sacré et al. (2016a). Five participants being treated at the Cleveland Clinic Epilepsy Center performed the behavioral task in their hospital rooms using methods previously described in Johnson et al. (2014). The gambling task (Fig. 7.4 top left) is based on a simple game of high card where participants would win virtual money if their card beat the computer's card. Specifically, in the beginning of each trial, the participant controls a cursor via a planar manipulandum to a fixation target. During fixation, participants must center the cursor in less than 8 s. Once centered, the participant is shown his card (only 2, 4, 6, 8, or 10 are in the deck) for a duration of 2 s. The card is randomly chosen with equal distribution.



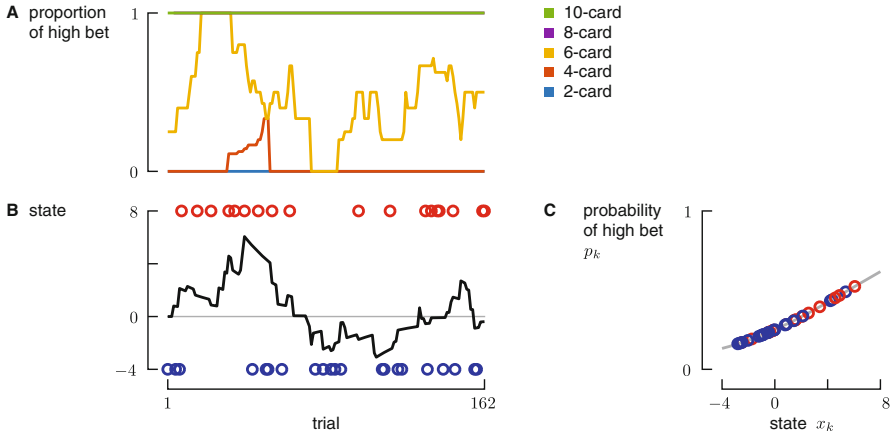
**Fig. 7.4** Timeline of the gambling task. After fixation, subjects were shown their card. Once the bets were shown, subjects selected one of the choices and then were shown the computer’s card following a delay. Feedback was provided afterwards by displaying the amount won or lost

The computer’s card is initially hidden. The screen then shows the two possible choices: a high bet (\$20) or a low bet (\$5). The participant has 6 s to select one with his cursor. Following selection, the computer’s card, which follows the same distribution, is revealed. If the computer’s card is larger than the player’s card, then the participant loses the amount he bets. If the computer’s card is smaller than the player’s card, then the participant wins the amount he bet.

For this task, the expected reward and variance of the reward are functions of the player’s card and bet. For example, on 10-card trials, the expected reward is higher for a high bet than for a low bet and the variance of reward is small for both decisions. On 6-card trials, the expected reward is zero for both betting decisions; but the variance of reward is higher for a high bet than for a low bet.

In this task, bets and reaction times for each trial,  $k$ , were the behavior variables measured. Neural activity was measured with stereotactic EEG depth electrodes. Participants were implanted with 10–14 depth electrodes, each having 10–16 contacts. See Sacré et al. (2016a) for details.

To explore behavioral variability in the data, one can plot behavioral responses to each stimulus type over trials. In our gambling task, we plot the fraction of high bets (smoothed by taking a moving average) on each card-type trials for each patient over his/her session. This is shown in Fig. 7.5a.



**Fig. 7.5** Relationship between betting variability and internal state. **(a)** Moving average of the proportion of high bets over session for different card-type trials (overlapping windows of length  $2w + 1$ , with  $w = 10$ ). **(b)** Estimated state trajectory overlaid with bets on 6-card trials over session for one patient. **(c)** Estimated probability of betting high as a function of estimated state trajectory overlaid with bets (red for a high bet, blue for a low bet) for each 6-card trial

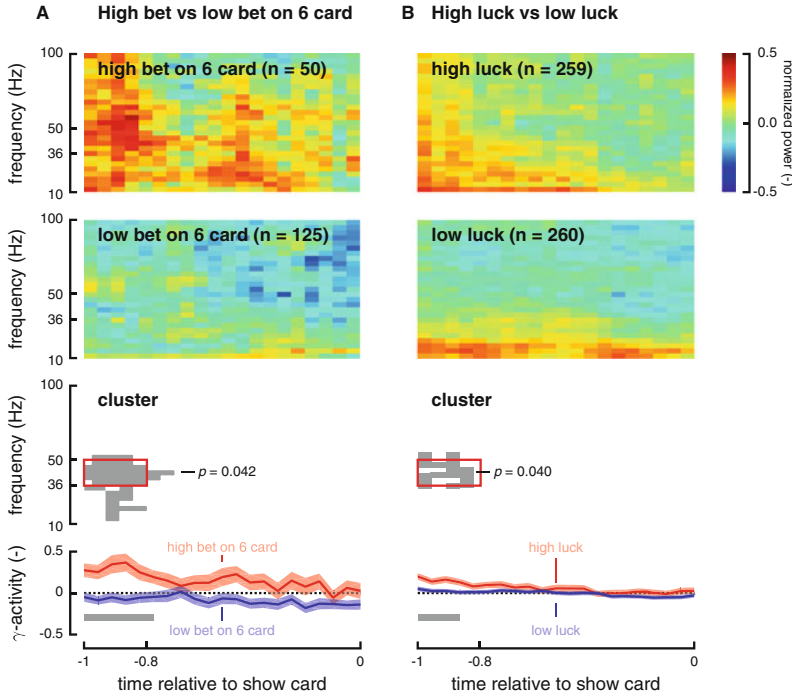
### 7.3.2.1 Dynamic State-Space Modeling

As seen in Fig. 7.5a, most of the within-participant variability is observed on 6-card trials across all participants. We hypothesized that participants bets on 6-card trials were influenced by past outcomes or a latent state variable that accumulated past outcomes. Specifically, we constructed a fading memory state model of cumulative mismatched expectations that we referred to as “luck”  $x_k$  on trial  $k$ . The luck variable is the scalar state variable that updates as follows:

$$x_{k+1} = ax_k + e_k \quad x_0 = 0, \quad (7.18)$$

where  $a$  is a decay factor ( $0 \leq a \leq 1$ ) and  $e_k$  is the mismatched expectations on trial  $k$ , that is, the difference between the actual outcome (loss =  $-1$ , draw =  $0$ , or win =  $1$ ) and expected outcome given the player card  $pc_k$  (computed as  $\frac{1}{5(pc_k - 6)}$ ). Note that  $e_k$  enters the state evolution equation only during trials where expectations are mismatched.

Next, we estimated  $a$  in Eq. (7.18) by varying it between 0 and 1 in 0.01 increments and computed the Pearson’s correlation coefficient between luck and gamma band power in the orbitofrontal cortex (OFC) at the beginning of each trial *before* the player sees his/her card (see Fig. 7.6). Thus, in this case, the OFC gamma power at the beginning of a trial was first found to be correlated to whether or not the player bets high if he/she receives a 6 card on that trial, and then the state-evolution model was constructed through a grid search.



**Fig. 7.6** Oscillatory power before the Show Card. (a) The average spectrograms show differences between high-bet and low-bet conditions on 6-card trials. One significant cluster ( $p = 0.042$ ) resulted from the cluster-based nonparametric statistical test. The cluster contained frequencies between 36 and 50 Hz at a timing between 1000 and 800 ms before the Show Card. This frequency range matches the traditional lower gamma band. Plots of average oscillatory power (36–50 Hz) over time for 6-card trials resulting in high and low bets show the modulation of the power in the gamma band preceding the Show Card. Time bins with significant differences are marked by the grey bar. Error bars represent one standard error of the mean. The number  $n$  denotes the number of trials pooled across patients. (b) The average spectrograms show differences between high-luck and low-luck conditions on all trials. One significant cluster ( $p = 0.040$ ) resulted from the cluster-based nonparametric statistical test. The cluster is located in the similar time-frequency region as the cluster emerging from the high-bet and low-bet conditions on 6-card trials. Figure reproduced with permission from Sacré et al. (2016a)

To see whether  $x_k$  explains the variability of behavior on 6-card trials, one can overlay the state  $x_k$  with bets on all trials across the session as shown in Fig. 7.5b.

To complete the state-space model, the output equation model is then a standard GLM for Bernoulli betting observations:

$$p_k = \frac{1}{1 + e^{-(d_0 + c \cdot x_k)}}, \quad \text{for } k \text{ such that } pc_k = 6. \quad (7.19)$$

Equation (7.19) is a standard GLM when  $x_k$  is known or estimated a priori. If  $x_k$  is not estimated ahead of time, then the EM algorithm can be used. One can also overlay the behavioral data with the output models for  $\hat{p}_k$ . An example of this is shown in Fig. 7.5c.

In this example, since the output equation is constructed using a standard GLM, the fitted model can be evaluated by checking the significance of the parameter  $c$  in front of the state variable. The model showed that the state of “luck” significantly influenced betting decisions ( $c = 0.20$ ,  $p = 0.028$ ). This indicates positive luck biases participants to bet high on 6 cards.

### 7.3.2.2 Neural Correlates Informed by Dynamic State-Space Model

In this example, we obtained continuous local field potential recordings from the OFC and thus analyzed data in the spectral domain. To compute spectrograms, three orthogonal tapers were used with a 300-ms window sliding at 50-ms steps. Frequencies under 10 Hz were dropped because of the Rayleigh criterion and analyzed upwards to 100 Hz. Afterwards, each frequency bin’s power was normalized based on the power across the entire recording session by fitting the log of the power in each frequency bin to a standard normal distribution. The mean and standard deviation used for the normalization were computed from the power between the 5th and 95th percentiles of the data set. This calculation was performed for every electrode’s recording with the final normalized power being averaged across all electrodes in the brain region of interest (OFC in our example). In addition, we removed artifacts by identifying time points in the spectrograms for which the median of the absolute value of the power across all frequencies is larger than 2.5. Finally, in order to remove the effect of 60 Hz power-line noise, we ignored the frequency bins between 56.66 and 63.33 Hz in all analyses.

OFC oscillatory power was compared between the set of trials where subjects end up betting high on a 6 card and the set where they end up betting low. The average normalized spectrograms for both high and low bet trials showed that high bet trials have higher 40–50 Hz oscillatory power about 1000 ms preceding the show card epoch (Fig. 7.6a). To determine statistical significance of this effect, we used a cluster-based nonparametric statistical test described above. Clusters here are defined as a set of adjacent time-frequency windows whose activity is different between trials where the subjects end up betting high versus low.

To examine the correlation between OFC activity and the state variable, we separated the trials between high-luck and low-luck conditions (defined as the bottom third and top third of the values taken by luck variable for all patients) and computed the average normalized spectrograms for both conditions. High-luck trials showed higher OFC oscillatory power than low-luck trials (Fig. 7.6b, first and second panels). Interestingly, the cluster-based nonparametric statistical test identified a significant cluster ( $p = 0.040$ ) in the time-frequency vicinity of the cluster identified when separating trials based on high-bet and low-bet conditions on 6-card trials (Fig. 7.6b, third panel).

These findings suggest OFC may play a pivotal role in processing a subject’s internal (emotional) state during financial decision-making.



## 7.4 Discussion

In this discussion, we highlight four important lessons to use a two-step state-space modeling approach described in this chapter to explore links between behavior and neural activity in humans.

**Lesson 1: Investigating Variability in Behavioral Data** The first lesson is to always start data analyses by exploring the variability in the behavioral data prior to building a model. A good understanding of the variability existing in the behavioral data is the key to a useful model. There are essentially two sources of variability in a participant's behavior. The first source is that the behavior changes as the stimulus changes, which is expected. The second source is that the participant's behavior changes in a "smooth" way over trials during which the same stimulus is applied. This can happen when internal states, such as motivation and attention, vary over trials. If the latter variability is observed in the data, then a state-space modeling framework is appropriate. In the two examples described above, we plot behavior and see both sources of variability and thus move forward with model development. If the second source of variability is not present, then a simple GLM of the behavior may suffice in explaining the first source of behavioral variability, which is how a stimulus impacts behavior.

**Lesson 2: Identifying Model Structure** The second lesson is to identify a model structure that explains the variability that we observed. The design of the measurement map  $h_\theta$  is the easiest part at most of the time: it involves a combination of states  $\mathbf{x}_k$  and inputs  $\mathbf{u}_k$ . The design of the state-transition map  $f_\theta$  is usually more complex. A useful analysis to guide the design of the state-transition map is to investigate the influence of candidate inputs by quantifying the influence of the value of candidate inputs at the previous trial  $k - 1$  on the behavior at trial  $k$ .

**Lesson 3: Estimating State-Space Model Step by Step** The third lesson is to estimate the model parameters and the state for each trial step by step. A good approach is to start by estimating the parameters of the static model, that is, the model where the state is fixed to 0. Then, it is also sometimes useful to estimate the parameters of the dynamic model where we fix the noise in the state evolution to zero. Finally, we can estimate the parameters of the whole dynamic model by using the previous estimates as a first guess for this more complex estimation problem. The decomposition into these different steps helps to interpret the meaning of each parameter and its influence on the state.

**Lesson 4: Dealing with Multiple Comparisons in Neural Data Analysis** The fourth lesson is to deal with multiple comparisons in neural data analysis. Indeed, we are often interested in looking at the neural activity from multiple brain regions (when available) and at different epochs during the task. A standard approach to tackle this multiple comparisons problem is to correct the significance threshold by controlling the false discovery rate (e.g.,  $q = 0.05$ ).

**Acknowledgements** This MIST work was supported by the National Institutes of Health (NIH R01 MH106700, NIH K12 NS080223, NIH S10 OD018211, NIH R01 NS084142) and the Dana Foundation. This gambling work was supported by the National Science Foundation (NSF EFRI-MC3: #1137237). P.S. was supported by the Kavli Neuroscience Discovery Institute at the Johns Hopkins University.

## References

- Agarwal, R., Thakor, N. V., Sarma, S. V., & Massaquoi, S. G. (2015). PMv neuronal firing may be driven by a movement command trajectory within multidimensional Gaussian fields. *Journal of Neuroscience*, *35*(25), 9508–9525.
- Brown, E. N., Barbieri, R., Eden, U. T., & Frank, L. M. (2003). Likelihood methods for neural spike train data analysis. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach*, Chap. 9 (pp. 253–286). London: Chapman & Hall/CRC.
- Carpenter, A. F., Georgopoulos, A. P., & Pellizzer, G. (1999). Motor cortical encoding of serial order in a context-recall task. *Science*, *283*(5408), 1752–1757.
- Coleman, T., & Sarma, S. (2007). Using convex optimization for nonparametric statistical analysis of point processes. In *Proceedings of IEEE International Symposium on Information Theory, 2007* (pp. 1476–1480).
- Coleman, T. P., & Sarma, S. V. (2010). A computationally efficient method for nonparametric modeling of neural spiking activity with point processes. *Neural Computation*, *22*(8), 2002–2030.
- Cox, D. R., & Isham, V. (1980). *Point processes* (Vol. 12). Boca Raton: CRC Press.
- Johnson, M. A., Thompson, S., Gonzalez-Martinez, J., Park, H. J., Bulacio, J., Najm, I., et al. (2014). Performing behavioral tasks in subjects with intracranial electrodes. *Journal of Visualized Experiments* (92), e51947. <https://doi.org/10.3791/51947>. <https://www.jove.com/video/51947/performing-behavioral-tasks-in-subjects-with-intracranial-electrodes>.
- Johnson, N. L., & Kotz, S. (1970). *Continuous univariate distributions*. New York: Houghton Mifflin.
- Louis, L. S. (1991). *Statistical signal processing: Detection, estimation, and time series analysis*. Boston, MA: Addison-Wesley Publishing Company.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190.
- Mauss, I. B., & Robinson, M. D. (2009). Measures of emotion: A review. *Cognition and Emotion*, *23*(2), 209–237.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models* (2nd ed.). London: Chapman & Hall/CRC.
- Moon, T. K. (1996). The expectation-maximization algorithm. *IEEE Signal Processing Magazine*, *13*(6), 47–60.
- Sacré, P., Kerr, M. S. D., Kahn, K., González-Martínez, J., Bulacio, J., Park, H.-J., et al. (2016a). Lucky rhythms in orbitofrontal cortex bias gambling decisions in humans. *Scientific Reports*, *6*, 36206.
- Sacré, P., Kerr, M. S. D., Subramanian, S., Kahn, K., González-Martínez, J., Johnson, M. A., et al. (2016b). Winning versus losing during gambling and its neural correlates. In *Proceedings of 50th Annual Conference on Information Science and Systems, Princeton, NJ* (pp. 557–561).
- Sacré, P., Kerr, M. S. D., Subramanian, S., Kahn, K., González-Martínez, J., Johnson, M. A., et al. (2016c). The precuneus may encode irrationality in human gambling. In *Proceedings of 38th Annual International Conference on IEEE Engineering in Medicine and Biology Society, Orlando, FL* (pp. 3406–3409).

- Santaniello, S., Gale, J. T., Montgomery, E., & Sarma, S. V. (2010). Modeling the motor striatum under deep brain stimulation in normal and MPTP conditions. In *Proceedings of 32nd Annual International Conference on IEEE Engineering in Medicine and Biology Society* (pp. 2065–2068).
- Santaniello, S., Montgomery, E. B., Gale, J. T., & Sarma, S. V. (2012). Non-stationary discharge patterns in motor cortex under subthalamic nucleus deep brain stimulation. *Frontiers in Integrative Neuroscience*, 6, 35.
- Sarma, S. V., Cheng, M. L., Eden, U. T., Williams, Z., Brown, E. N., & Eskandar, E. N. (2012). The effects of cues on neurons in the basal ganglia in Parkinson's disease. *Frontiers in Integrative Neuroscience*, 6, 40.
- Sarma, S. V., Eden, U. T., Cheng, M. L., Williams, Z. M., Hu, R., Eskandar, E. N., et al. (2010). Using point process models to compare neural spiking activity in the subthalamic nucleus of Parkinson's patients and a healthy primate. *IEEE Transactions in Biomedical Engineering*, 57(6), 1297–1305.
- Schieber, M. H. (2004). Motor control: Basic units of cortical output? *Current Biology*, 14(9), R353–R354.
- Shoham, S., Fellows, M. R., & Normann, R. A. (2003). Robust, automatic spike sorting using mixtures of multivariate *t*-distributions. *Journal of Neuroscience Methods*, 127(2), 111–122.
- Sklar, S., Walmer, M., Sacré, P., Schevon, C. A., Srinivasan, S., Banks, G. P., et al. (2017). Neuronal activity in human anterior cingulate cortex modulates with internal cognitive state during multi-source interference task. In *Proceedings of 39th Annual International Conference on IEEE Engineering in Medicine and Biology Society* (pp. 962–965).
- Smith, E. H., Banks, G. P., Mikell, C. B., Cash, S. S., Patel, S. R., Eskandar, E. N., et al. (2015). Frequency-dependent representation of reinforcement-related information in the human medial and lateral prefrontal cortex. *Journal of Neuroscience*, 35(48), 15827–15836.
- Snyder, D. L., & Miller, M. I. (1991). *Random point processes in time and space*. New York: Springer.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., & Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93(2), 1074–1089.
- Van Trees, H. L. (1968). *Detection, Estimation, and Modulation Theory, Part I*. New York, NY: Wiley.