

Omitted Variable Bias in GLMs of Neural Spiking Activity

Ian H. Stevenson

ian.stevenson@uconn.edu

*Department of Psychological Sciences, Department of Biomedical Engineering,
and CT Institute for Brain and Cognitive Sciences, University of Connecticut,
Storrs, CT 06269, U.S.A.*

Generalized linear models (GLMs) have a wide range of applications in systems neuroscience describing the encoding of stimulus and behavioral variables, as well as the dynamics of single neurons. However, in any given experiment, many variables that have an impact on neural activity are not observed or not modeled. Here we demonstrate, in both theory and practice, how these omitted variables can result in biased parameter estimates for the effects that are included. In three case studies, we estimate tuning functions for common experiments in motor cortex, hippocampus, and visual cortex. We find that including traditionally omitted variables changes estimates of the original parameters and that modulation originally attributed to one variable is reduced after new variables are included. In GLMs describing single-neuron dynamics, we then demonstrate how postspike history effects can also be biased by omitted variables. Here we find that omitted variable bias can lead to mistaken conclusions about the stability of single-neuron firing. Omitted variable bias can appear in any model with confounders—where omitted variables modulate neural activity and the effects of the omitted variables covary with the included effects. Understanding how and to what extent omitted variable bias affects parameter estimates is likely to be important for interpreting the parameters and predictions of many neural encoding models.

1 Introduction ---

Regression models have been widely used in systems neuroscience to explain how external stimulus and task variables, as well as internal state variables, may relate to observed neural activity (Brown, Barbieri, Eden, & Frank, 2003; Kass, Ventura, & Brown, 2005). However, in many cases, the full set of variables that explain the activity of the observed neurons is not observed or is not even known. It is important to recognize that in these cases, omitted variables can cause the parameter estimates for the effects that are included in a regression model to be biased. That is, parameter estimates for the modeled effects would be different if other, omitted

variables were to be included in the model (Box, 1966). In experiments from behaving animals (Niell & Stryker, 2010; Reimer et al., 2014), but also in more controlled sensory tasks (Arandia-Romero, Tanabe, Drugowitsch, Kohn, & Moreno-Bote, 2016; Kelly, Smith, Kass, & Lee, 2010), there is growing evidence that neural activity is affected by many more variables than are typically considered relevant (Kandler, Mao, McNaughton, & Bonin, 2017; Stringer et al., 2018). At the same time, although it has long been a concern in statistics (Pearl, 2009) and has received some attention in other fields (Clarke, 2005; Greenland, 1989), omitted variable bias, as a general problem, appears underappreciated in systems neuroscience. Here we demonstrate why systematically considering omitted variable bias may be important in neural data analysis and examine how omitted variable bias can affect one popular framework for describing neural spiking activity: the generalized linear model (GLM) with Poisson observations.

In general, regression methods aim to estimate variations in a response variable as a function of other variables or covariates. When the goal of modeling is to maximize prediction accuracy, such as with brain-machine interfaces, interpreting the model parameters may not be a high priority. However, in many other cases, parameter estimates are, at least to some extent, interpreted and analyzed. For instance, tuning curves or receptive fields may be measured and compared under different stimulus or task conditions or before and after a manipulation. In fully controlled experiments, where the covariates are assigned at random, estimated coefficients can often be interpreted as estimates of causal effects (Gelman & Hill, 2007). However, for many cases in neuroscience, it may be difficult or impossible to completely control or randomize all the relevant variables.

In modeling neural activity, omitted variable bias can appear in any situation where neurons are modulated by omitted variables and the omitted variables (often called confounders) are not independent from the variables included in the model—the ones whose effects we are trying to estimate. Minimizing the influence of confounding variables is a major part of most experimental design (Rust & Movshon, 2005), and the statistical effects of confounding variables are well understood (Wasserman, 2004). However, when the goal of modeling is description or explanation (Shmueli, 2010), the effects of these omitted variables are frequently neglected. To give a concrete example, imagine an idealized neuron in primary motor cortex (M1) whose firing, unlike typical M1 neurons (Georgopoulos, Kalaska, Caminiti, & Massey, 1982), is not at all modulated by reach direction but instead is modulated by reach speed (see Figure 1). In a typical experimental setting, an animal's reach directions are randomized, but reach speed cannot be randomized or tightly controlled. If the average speed differs across reach directions, such a hypothetical neuron will appear to be tuned to reach direction, despite not being directly affected by direction. First, fitting a typical tuning curve for reach direction, we would infer that such a neuron has a clear preferred direction and nonzero modulation depth. If we then fit a

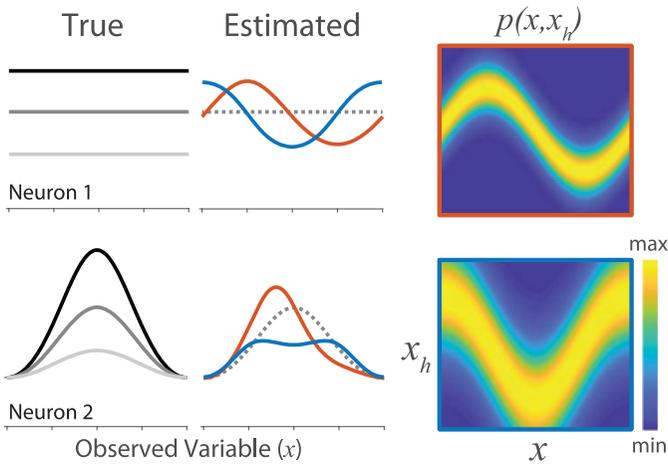


Figure 1: When relevant variables are omitted from the model, estimates of the included effects can be biased. Consider two hypothetical neurons tuned to an observed variable x and an omitted variable x_h . Neuron 1 is not tuned to the observed variable x , but its rate is modulated by the omitted variable with true tuning curves denoted by the gray curves (top left). If x and x_h covary, the apparent tuning of this neuron to x when the tuning curve is estimated using x alone will then be biased (red and blue curves, top middle). This neuron will appear to be tuned to x despite not actually being tuned to this variable. In addition, to this type of illusory tuning, there can also be more subtle biases. Here, neuron 2 is tuned to x and x_h (gray curves, bottom left). However, depending on how x and x_h covary, the preferred stimulus or the modulation can be misestimated. Here the true tuning to x is shown at three different fixed values of x_h (three gray curves, left panels). The estimated tuning when x_h is omitted is shown at the center with red and blue curves corresponding to the estimates under two different joint distributions (matching borders, right). The dashed line denotes the effect of x when x_h is fixed.

second model that included both reach direction and speed, we would infer that the neuron is modulated by speed alone, and it would be apparent that the original preferred direction and modulation depth estimates were biased due to the omitted variable.

In adding more variables, previous studies have largely focused on the fact that including previously omitted variables improves model accuracy or the fact that neural activity is often influenced by a host of task variables. In M1, for instance, including speed improves model accuracy (Moran & Schwartz, 1999), but the presence of many correlated motor variables (e.g., kinematics, end-point forces, muscle activity) makes it difficult to interpret how neurons represent movement overall (Humphrey, Schmidt, & Thompson, 1970; Omrani, Kaufman, Hatsopoulos, & Cheney, 2017). In this letter,

instead of focusing on the advantages or complexities of models with many variables, we focus on the well-known, but underdiscussed, fact that the parameters describing the original effects change as additional variables are included. The hypothetical M1 neuron above points to a more general question about regression models of neural activity. What happens when we cannot or do not include variables that are relevant to the process that we are modeling?

Here we first evaluate the statistical problem of omitted variable bias in the canonical generalized linear model with Poisson observations. Then, as a case study, we examine how speed affects estimates of direction tuning of neurons in primary motor cortex, as well as two other case studies where the spike counts are modeled as a function of external variables: orientation tuning in primary visual cortex (V1) and place tuning in the hippocampus (HC). In each of these case studies, we find that commonly omitted variables (speed in M1, population activity in V1, and speed and heading in HC) can bias the estimated effects of commonly included variables (reach direction in M1, stimulus orientation or direction in V1, and place in HC). Across all three case studies, including the omitted variables reduces the estimated modulation due to typical tuning effects. We also illustrate how omitted variable bias can affect generalized linear models of spike dynamics where a postspike history filter aims to describe refractoriness and bursting (Truccolo, Eden, Fellows, Donoghue, & Brown, 2005). The goal of these models is typically to differentiate aspects of spike dynamics that are due to the neurons' own properties (e.g., membrane time constant, resting potential, after-hyperpolarization currents) from those due to input to the neuron from other sources (Brillinger & Segundo, 1979; Paninski, 2004). In this setting, the input to the neuron is typically not directly observed but is approximated by stimulus or behavioral covariates, local field potential, or the activity of other neurons. Here we show that omitting the input can lead to large biases in postspike history filters and that including omitted variables describing the input can change the interpretation and stability of the estimated history effects.

GLMs have been used in many settings to disentangle the effects of multiple, possibly correlated, stimulus or task variables (Fernandes, Stevenson, Phillips, Segraves, & Kording, 2014; Park, Meister, Huk, & Pillow, 2014; Runyan, Piasini, Panzeri, & Harvey, 2017) and also to model neural mechanisms such as postspike dynamics, interactions between neurons, and coupling to local fields (Harris, Csicsvari, Hirase, Dragoi, & Buzsáki, 2003; Pillow et al., 2008; Truccolo et al., 2005). It is often argued that GLMs are advantageous because they have unique maximum likelihood estimates and can be more robust to nonspherical covariate distributions than other methods, such as spike-triggered averaging (Paninski, 2004; Pillow, 2007). Although these advantages are important, GLMs are not immune to bias. Here we show how the possibility of omitted variable bias, in particular, should encourage researchers to be cautious in their interpretation of model

parameters, even in cases where a GLM achieves high predictive accuracy (Shmueli, 2010).

2 Results

Here we introduce the problem of omitted variable bias and examine differences between omitted variable bias in linear models and the canonical Poisson GLM. We then consider three tuning curve estimation problems: estimating direction tuning in primary motor cortex, place tuning in hippocampus, and orientation tuning in primary visual cortex, and show how omitted variables in each of these three cases can alter parameter estimates. Finally, we consider a GLM that aims to describe the dynamics of postspike history and show how omitted inputs can bias the estimated history effects and qualitatively change model stability.

2.1 Omitted Variable Bias in Linear Regression and Canonical Poisson GLMs. When relevant variables are not included in a regression model, the estimated effects for the variables that are included can be biased (Box, 1966). Omitted variable biases can cause the parameters describing the effects of the original variables to be over- or underestimated, and model fits can change qualitatively when omitted variables are included (see Figure 1).

To understand the problem of omitted variable bias, it will be helpful to briefly review the well-known case of multiple linear regression, where the bias can be described analytically (Box, 1966). In the linear setting, consider the generative model,

$$y = X\beta + X_h\beta_h + \epsilon,$$

where observations y are a linear combination of observed X and omitted X_h variables plus normally distributed independent and identically distributed (i.i.d.) noise $\epsilon \sim N(0, \sigma)$. For simplicity, we ignore the intercept term, but in the analysis that follows, it may also be considered as part of X . If we then fit the (misspecified) model without X_h using maximum likelihood estimation (equivalent to the ordinary least squares solution, in this case), the estimated parameters will be

$$\begin{aligned}\hat{\beta} &= (X^T X)^{-1} X^T y \\ &= \beta + (X^T X)^{-1} X^T X_h \beta_h + \xi,\end{aligned}$$

where ξ denotes the effect of the noise and the bias $(X^T X)^{-1} X^T X_h \beta_h$ will generally be nonzero. There will be no bias only in the cases where the omitted variables do not affect the observations ($\beta_h = 0$) or when the omitted variables and observed variables are not collinear ($X^T X_h = 0$). Note that $(X^T X)^{-1} X^T X_h$ is the matrix of regression coefficients for the omitted

variables using the observed variables as predictors. For linear regression, the omitted variable bias thus depends on both the extent to which the omitted variables affect the observations β_h and the extent to which the omitted variables can be (linearly) predicted from the observed variables.

Although there is a closed-form solution for the omitted variable bias for linear regression, the generalized linear setting is not as tractable (Clogg, Petkova, & Shihadeh, 1992; Drake & McQuarrie, 1995; Gail, Wieand, & Piantadosi, 1984). We will consider the case of a canonical Poisson GLM, in particular, where

$$\begin{aligned}\lambda &= \exp(X\beta + X_h\beta_h), \\ y &\sim \text{Poisson}(\lambda),\end{aligned}$$

In the more general case, GLMs have

$$E[y] = g^{-1}(X\beta + X_h\beta_h),$$

where $g^{-1}(\cdot)$ is the inverse link function and y is distributed following an exponential family distribution (McCullagh & Nelder, 1989). For a canonical GLM the log likelihood takes the form

$$\mathcal{L}(\beta, \beta_h) \propto \sum_i y_i(X_i\beta + X_{h,i}\beta_h) - G(X_i\beta + X_{h,i}\beta_h) + \text{const}(\beta, \beta_h),$$

where the nonlinear function $G(\cdot)$ depends on both the link function and the noise model. For canonical GLMs, this log likelihood is concave, and the maximum likelihood estimate $\hat{\beta}$ satisfies $\frac{\partial \mathcal{L}}{\partial \beta} |_{\hat{\beta}} = 0$. The exact form of $G(\cdot)$ will depend on the model, but for linear regression, $G(x)$ is proportional to $\frac{x^2}{2}$, and for canonical (log-link) Poisson regression, $G(\cdot) = \exp(\cdot)$.

Now, with omitted variables, instead of maximizing the correct log likelihood, we maximize instead

$$\mathcal{L}_o(\beta) \propto \sum_i y_i(X_i\beta) - G(X_i\beta) + \text{const}(\beta).$$

For the omitted variable bias in $\hat{\beta}$ to be 0, we need both $\frac{\partial \mathcal{L}}{\partial \beta} |_{\hat{\beta}} = 0$ and $\frac{\partial \mathcal{L}_o}{\partial \beta} |_{\hat{\beta}} = 0$ at the same value of β . Although neither MLE has a closed-form solution, this condition implies that if there is no bias due to the omitted variables,

$$X^T G'(X\beta + X_h\beta_h) = X^T G'(X\beta),$$

where $G'(\cdot)$ is the derivative of $G(\cdot)$. For linear regression, this equality reduces to the OLS form derived above, and for canonical Poisson regression we have

$$X^T \exp(X\beta + X_h\beta_h) = X^T \exp(X\beta).$$

This equality is satisfied when observations are not modulated by the omitted variables $\beta_h = 0$ or, more generally, when the effect of the omitted variables $\delta\lambda = \exp(X\beta + X_h\beta_h) - \exp(X\beta)$ is orthogonal to the included variables X . Note that with linear regression, $X^T X_h = 0$ implies that the estimates will not be biased, but here this is not the case unless $X^T \delta\lambda = 0$ as well. Due to the structure of the canonical Poisson GLM, omitted variable bias can thus occur even in a properly randomized, controlled experiment (Gail et al., 1984).

It is important to note that the maximum likelihood estimates themselves are consistent. That is, the estimators converge (in probability) to their true values when the generative model is correct. The bias here is a result of the model being misspecified. This misspecification affects the location of the maximum and the shape of the likelihood. Optimization methods such as Newton's method typically contain omitted variable bias in each parameter update. For canonical Poisson regression, for instance, the updates take the form

$$\hat{\beta}^{k+1} \leftarrow \hat{\beta}^k + (X^T W X)^{-1} X^T (y - \lambda)$$

at iteration k , where the weight matrix W is diagonal with entries $W_{ii} = \lambda_i$ and $(X^T W X)^{-1}$ is the Fisher scoring matrix (inverse Hessian of the log likelihood) at the current estimate $\hat{\beta}^k$. Since the misspecified model will use $\lambda = \exp(X\beta)$ instead of the $\exp(X\beta + X_h\beta_h)$ both the weight matrix and the gradient $X^T (y - \lambda)$ will be biased at each step of the optimization (except when $X^T \delta\lambda|_{\beta^k} = 0$). Traditional standard errors for the MLE will also typically be influenced by omitted variables, since $\widehat{se}_i = \sqrt{((X^T W X)^{-1})_{ii}}$. Moreover, as previous studies have shown, omitted variables can lead to misestimation of the variability in $E[y]$ and dispersion $\text{var}(y)$ (Czanner et al., 2008; Goris, Movshon, & Simoncelli, 2014). If the omitted variables affect the observations, they will generally increase the variability of $E[y]$. Then, unless the omitted variables are perfectly predicted by the included variables, the explained variance $\text{var}(E[y])$ of the misspecified model will be lower than that of the full model. This may in turn lead to overestimates of dispersion, since $\text{var}(y) = E[\text{var}(y)] + \text{var}(E[y])$.

2.2 Omitted Variable Bias in Tuning Curve Estimation. When fitting tuning curve models to spike count data, omitted variable bias can cause preferred stimuli and modulation depths to be misestimated and can even lead to completely illusory tuning (see Figure 1). To illustrate how omitted variable bias affects GLMs of neural spiking, not just in theory but in practice, we consider three case studies where we fit typical tuning curve models that omit potentially relevant variables along with augmented models that

include these additional variables. We first consider modeling spike counts across trials and on relatively slow (over 100 ms) timescales. Here we assess (1) the tuning of neurons in motor cortex to reach direction, with speed as a potential omitted variable; (2) the tuning of neurons in hippocampus to position, with both speed and head direction as potential omitted variables; and (3) the tuning of neurons in visual cortex to the direction of motion of a sine-wave grating, with population activity as a potential omitted variable. In each of these cases studies, we show how the omitted variables are not independent of the commonly included variables and how neural responses are modulated by the omitted variables. Together, these two properties, can lead to omitted variable biases.

In our first case study, we model data recorded from primary motor cortex (M1) of a macaque monkey performing a center-out, planar reaching task. In this task, speed differs systematically across reach directions (see Figure 2A), with average speed differing by as much as $35 \pm 3\%$ (for the targets at 45 and 225 degrees relative to right; see Figure 2B). To model neural responses, we first fit a traditional tuning curve model (Amirikian & Georgopoulos, 2000; Georgopoulos et al., 1982), where the predicted responses depend only on target direction. Here we use a circular, cubic B-spline basis (five equally spaced knots) to allow for deviations from sinusoidal firing, but in most cases, the responses of the $n = 81$ neurons in this experiment are well described by cosine-like tuning curves with clear modulation for reach direction. We then fit a second model that includes effects from movement speed. Here we use covariates based on Moran and Schwartz (1999), including a linear speed effect, as well as cosine-tuned direction-speed interactions (see section 4). This model captures the responses of individual neurons, where spike counts can increase (see Figure 2C, top) or decrease (see Figure 2C, middle) as a function of speed, and, in some cases, speed and direction appear to interact (see Figure 2C, bottom). Together, the fact that direction and speed are not independent, along with the fact that neural responses appear to be modulated by speed, could lead to biased parameter estimates for the model where speed is omitted.

Comparing the models with and without omitted variables, we find that, averaged across the population, there are only minimal shifts in the preferred direction (3 ± 2 degrees) when speed is included in the traditional tuning curve model, and there do not appear to be large, systematic shifts in the population distribution of PDs (Kuiper's test, $p > 0.1$). At the same time, there is substantial variability between neurons in the size of the PD shift (circular SD 32 ± 5 degrees). Across the population, modulation depth (measured using the standard deviation of the tuning curve) decreases slightly on average ($3 \pm 2\%$), and the size of the modulation change also varies substantially between individual neurons (SD of changes $18 \pm 3\%$). An example neuron in Figure 1C (bottom), for instance, has a modulation decrease of $9 \pm 5\%$, and the preferred direction changes 4 ± 9 degrees when speed is included in the model (standard error from bootstrapping). Overall,

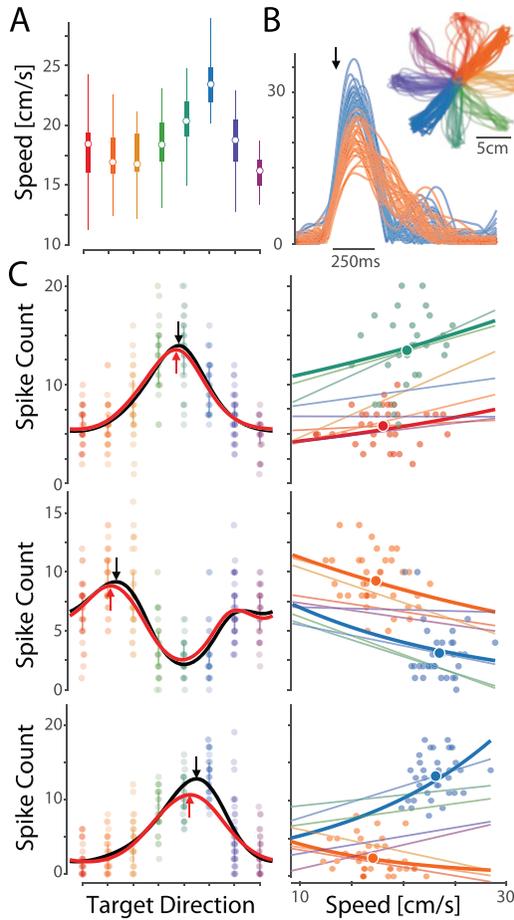


Figure 2: Speed as an omitted variable in M1 tuning for reach direction. (A) The distribution of reach speeds differs by target direction in a center-out task. Circles denote median, and boxes denote interquartile range. (B) Speed profiles for the two targets showing the largest speed differences. Individual traces denote individual trials aligned to the half-max (black arrow). The inset shows the position of each trial with colors denoting reach direction. (C) The responses of three M1 neurons show typical tuning for reach direction. The tuning curve estimated using direction covariates alone (black) changes when speed covariates are included (red). Red curves denote the direction effect within the full model and are generated by assuming speed is constant (equal to the mean speed across all trials). Right panels illustrate the speed dependence for the preferred direction and its opposite. Dark lines denote the estimated effect of speed under the full model. Data points show single trial data, along with the mean speed and rate for each direction (big data point). Light lines show linear trends (OLS) using only the trials from each specific target.

approximately 10% of neurons have statistically significant changes in PD, and approximately 14% have significant changes in modulation (bootstrap tests $\alpha = 0.05$, not corrected for multiple comparisons). For some individual neurons, at least, the parameters of the model without speed thus have clear omitted variable bias. However, since individual neurons have diverse speed dependencies, in this case, the average biases across the population are minimal.

When speed is included in the model, model accuracy (pseudo- R^2) does increase slightly ($p = 0.01$, one-sided paired t -test across neurons). The average cross-validated (jackknife) pseudo- R^2 for the original model is 0.23 ± 0.01 and for the model with speed, $0.24 \pm .01$ (see Figure 5). However, it seems likely that in other experimental contexts, the effects of omitting speed could be more pronounced. By requiring the animal to make reaches to the same targets at different speeds, previous studies have more clearly demonstrated that responses in M1 are modulated by speed (Churchland, Santhanam, & Shenoy, 2006). Here we demonstrate how this type of modulation can lead to omitted variable biases in the estimated parameters of typical tuning curve models without speed.

In our second case study, we examine the activity of neurons in the dorsal hippocampus of a rat foraging in an open field. Here we consider to what extent the practice of omitting speed and head direction from a place field model biases estimates of a neuron's position tuning. As in the first case study, omitted variable bias can occur if neural activity is modulated by omitted variables and the omitted variables covary with the included variables. In the case of the hippocampus, neural activity is known to be modulated by both movement speed and head direction (McNaughton, Barnes, & O'Keefe, 1983), in addition to an animal's position (O'Keefe & Dostrovsky, 1971). Additionally, behavioral variables can be highly nonuniform across the open field (Walsh & Cummins, 1976)—for instance, near and far from the walls. Together the fact that the omitted variables may covary with position and the fact that neurons appear to be modulated by the omitted variables suggest that there may be omitted variable bias.

Here, in one recording during an open field foraging task, we find that the average speed (see Figure 2A) and heading (see Figure 2B) differ extensively as a function of position. Within a given neuron's place field, the distributions of speed and heading may be very different from their marginal distributions. Across the population of $n = 68$ place cells (selected from 117 simultaneously recorded neurons; see section 4), average in-field speed was between 80% and 135% of the average overall speed (5.5 cm/s), and the animal's heading can be either more or less variable in-field (circular SD 57-80 deg) compared to overall (75 degrees).

As previous studies have shown, we also find that neural responses are modulated by speed and head direction. Responses due to place, speed, and heading are shown for one example neuron in Figure 3. This neuron shows a stereotypical place-dependent response (see Figure 3B), but

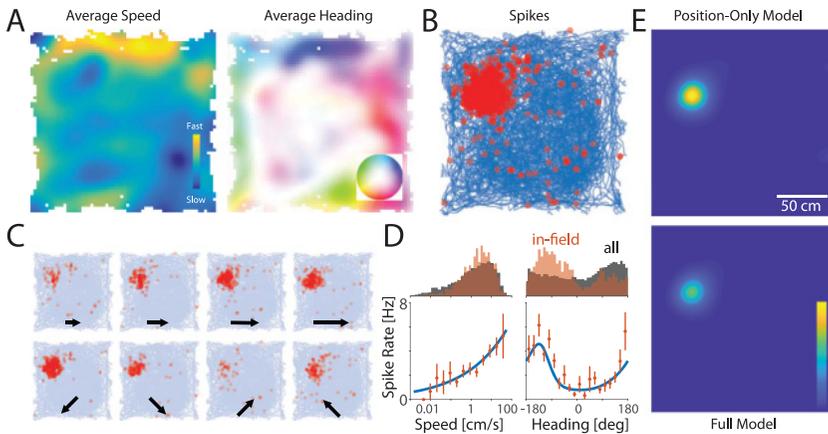


Figure 3: Speed and heading as omitted variables in hippocampal place cells. (A) Average speed and heading as a function of position for a rat foraging in an open field. (B) An example place cell tends to spike (red dots) when the animal is at a specific position in space. (C) The activity of this neuron is modulated by the animal's speed (top row) and heading (bottom row). Speed is split into quartiles, and subplots include all headings. The heading is split into quadrants, and subplots include all speeds. (D) The distributions of speed and heading within the place field differ from the overall distributions, and the neuron is tuned to these variables. The blue curve shows model fit. (E) After modeling the effect of speed and heading within the place field, the location of the place field does not change, but the apparent modulation due to position is reduced.

splitting the observations by speed (see Figure 3C, top) or heading (see Figure 3C, bottom) by quartiles/quadrants reveals that there is also tuning to these variables. The neuron appears to increase its firing with increasing speed and responds most strongly when the rat is facing the left. These dependencies are well fit by the full model where the firing rate depends not just on position, but also on the (log-transformed) speed and the heading (see Figure 3D, bottom). For the example place cell shown here, the location of the place field does not change substantially when the omitted variables are included (see Figure 3E). However, the modulation (SD of the rate map) decreases by 27%. That is, 27% of the apparent modulation due to position when it is modeled alone can be explained by speed and heading effects.

Across the population of place cells, there were no clear, systematic differences in the place field locations, but the modulation (SD of the rate map $\lambda(x)$) decreases by $9 \pm 1\%$ on average when speed and heading are included. Individual neurons showed substantial variability in their modulation differences (population SD $10 \pm 1\%$). As in M1, including the omitted variables increased spike prediction accuracy: the average cross-validated

(10-fold) pseudo- R^2 was $.29 \pm .02$ for the original model and $.31 \pm .02$ for the model, including speed and heading activity. This difference seems small since there is large variability in pseudo- R^2 values across the population, but the average increase in pseudo- R^2 was $11 \pm 3\%$ (see Figure 5). Given that neurons appear to be modulated by speed and heading, it is unsurprising that including these variables improves model fit. However, as before, it is important to note that this modulation can lead to biases in the place field estimates for the model with only position.

In our third case study, we examine the activity of neurons in a more controlled sensory experiment. Here we use data recorded from primary visual cortex (V1) of an anaesthetized monkey viewing oriented sine-wave gratings moving in 1 of 12 different directions (see section 4). In this experiment, variability in the animal's behavior is purposefully minimized, and instead of considering the effect of omitting a behavioral variable, we consider the effect of omitting a variable relating to the animal's internal state: the total population activity. Several studies have shown that population activity alters neural responses in V1 (Arandia-Romero et al., 2016; Arieli, Sterkin, Grinvald, & Aertsen, 1996; Kelly et al., 2010; Okun et al., 2015). If the distribution of population activity also varies with stimulus direction, there is the potential for omitted variable bias.

Here we assess neural activity from $n = 90$ simultaneously recorded neurons across many (2400) repeated trials with 12 different movement directions. We find high trial-to-trial variability in the population rate (see Figure 4A), and the average firing across all neurons differs across stimulus directions—up to about 50%. For this recording, the most extreme differences were between the 180 degree stimulus, where the average rate across the population was 3.4 ± 0.1 Hz, and the 60 degree stimulus, where the average rate was 6.3 ± 0.1 Hz (see Figure 4B). By adding the (log-transformed) population rate as a covariate to a more typical model of direction tuning, we find that population activity may lead to omitted variable bias in models of direction tuning alone.

As in the case studies above, there do not appear to be any consistent or systematic effects on the preferred stimulus direction at the population level (Kuiper's test, $p = 0.1$). However, the modulation depth (measured using SD of the tuning curve) decreases substantially $15 \pm 2\%$ when population rate is included in the model, and there is again high variability across neurons (SD $20 \pm 2\%$). In this case, model accuracy increases substantially when the omitted variable is included. The cross-validated (10-fold) pseudo- R^2 is $.26 \pm .02$ for the original model and $.43 \pm .02$ for the model including population activity, with an average increase of $164 \pm 31\%$ (see Figure 5).

Unlike in M1, where the effect of speed was highly diverse for different neurons, in this case study, the effect of the population rate is largely consistent. Higher population rates are associated with higher firing rates, and for most neurons, the effect of the population rate is stronger in the preferred direction(s), consistent with a multiplicative effect. Note that, we do

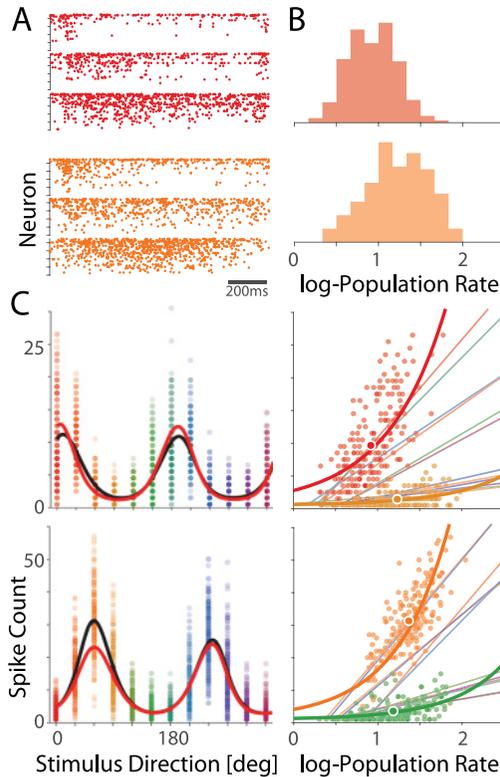


Figure 4: Population rate as an omitted variable in primary visual cortex. (A) Correlated trial-to-trial variability. Population rasters for three trials of the same drifting grating stimulus (0 degree, red; 30 degrees, orange). Neurons are sorted by overall firing rate. (B) Histograms of the population rate across trials. As a population, the neurons respond at higher rates to 30 degree stimuli, but there is high trial-to-trial variability. (C) The responses of 2 V1 neurons show typical tuning for direction of motion. The tuning curve estimated using direction covariates alone (black) changes when the population rate covariates are included (red). Right panels illustrate the dependence for the preferred direction and an orthogonal direction. Dark lines denote the estimated effect of speed under the full model. Data points show single trial data, along with the mean count and rate (big data point). Light lines show linear trends (OLS) using only the trials from each specific stimulus.

not include the neuron whose rate we are modeling in the calculation of the population rate. However, using the population rate as an omitted variable requires some interpretation. The population rate will certainly be affected by the tuning of the relatively small sample of neurons that we observe.

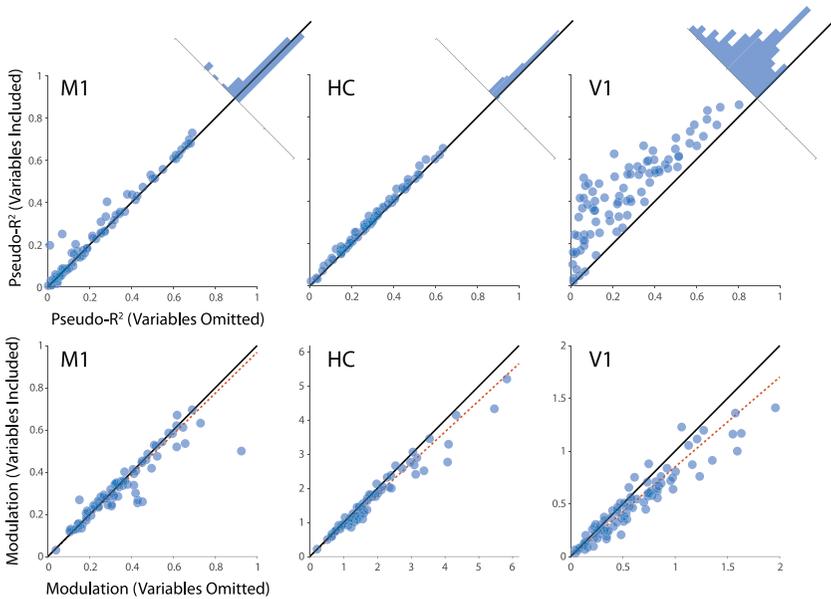


Figure 5: For each of the case studies, on average, the model accuracy increases when omitted variables are included (top) and the modulation due to the original variables decreases (bottom). Scatter plots indicate cross-validated pseudo- R^2 values for each neuron under the two models. Modulation denotes the standard deviation of the tuning to the original variable(s) under each model. Here, modulation values are normalized by the average rate of each neuron. Black lines denote equality. Red dashed lines denote linear fit with 0 intercept.

If we have a disproportionate number of neurons tuned to a specific preferred direction, the population rate in those directions will be higher. This suggests that in a different recording, the covariation between the stimulus and the population rate could very likely be different. However, it appears that the omitted variable biases in this case are mostly driven by noise correlations, where neural activity is correlated on single trials even within the same stimulus condition, rather than stimulus correlations, where neural activity is correlated due to similar tuning. When we shuffle the data within each stimulus condition (removing noise correlations) the average change in the modulation depth is $-1 \pm 2\%$ (SD $18 \pm 3\%$), and the effect of the omitted variable becomes negligible.

2.3 Omitted Variable Bias in the Estimation of Postspike History Effects. In addition to modeling spike counts over trials or on relatively slow (more than 100 ms) behavioral timescales, GLMs are also often used

to describe detailed, single-trial spike dynamics on fast (less than 10 ms) timescales. One common covariate used in these types of models is a post-spike history effect where the probability of spiking at a given time depends on the recent history of spiking. Modeling these effects allows us to describe refractoriness, bursting (Paninski, 2004; Truccolo et al., 2005), and a whole host of other dynamics (Weber & Pillow, 2017). Conceptually, the goal of these models is to disentangle the sources of rate variation based only on observations of a neuron's spiking, with history effects ideally reflecting intrinsic biophysics. However, since the full synaptic input is typically not known with extracellular spike recordings, there is potential for omitted variable biases.

To illustrate the potential pitfalls of omitting the input to a neuron, consider using the GLM to capture single-neuron dynamics in the complete absence of external covariates

$$\lambda(t) = \exp(\mu + \alpha h(t)),$$

where the rate λ is determined by a baseline parameter μ along with a filtered version of the neuron's past spiking with $h_i(t) = \sum_{\tau>0} f_i(\tau)n(t - \tau)$. This is a perfectly acceptable model of intrinsic dynamics, but for most spike data that we observe, this isolated neuron model may not provide a realistic description of a neuron receiving thousands of time-varying synaptic inputs. If we fit this model to data where the input to the neuron did vary over time,

$$\lambda(t) = \exp(\mu + \alpha h(t) + \beta_h x_h(t)),$$

then the history filter in the first model will attempt to capture variation in spiking due to the time-varying input, in addition to any intrinsic dynamics. For example, when x_h is periodic, the estimated history filters of the original model will attempt to capture this periodic structure (see Figures 6A and 6B). Just as in the tuning curve examples above, the fact that history effects covary with the input and the fact that the input modulates the neuron's firing leads to omitted variable bias. When the input is omitted from the model, the biased history effects simply provide the best (maximum likelihood) explanation of the observed spiking (see Figure 6C).

These examples with strong, periodic input are not necessarily biologically realistic, but they make it apparent how the postspike history can be biased by omitted input variables. In vivo, neurons instead appear to be in a high-conductance state, where membrane potential fluctuations have approximately $1/f$ power spectra (Destexhe, Rudolph, Fellous, & Sejnowski, 2001; Destexhe, Rudolph, & Paré, 2003). When these naturalistic input statistics are used to drive the GLM, omitted variable bias can occur as well. Here we simulate a GLM receiving $1/f^\alpha$ noise input with $\alpha = 0$ (white

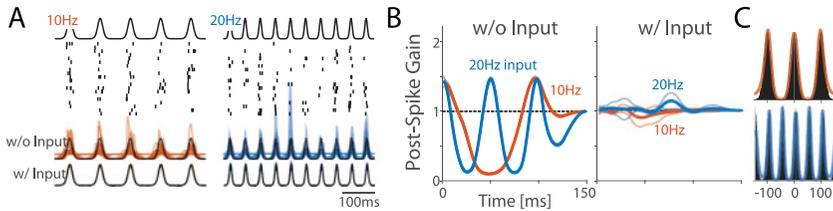


Figure 6: Estimated postspike history filters can be heavily biased when the input is not included in the model. (A) We simulate from an inhomogeneous Poisson model with sinusoidal input (no postspike history effects). The input and spike responses from 20 trials are shown. Although there are no history effects in the generative model, a GLM with history effects that is missing the correct input covariate will use the history terms to capture the structure in the autocorrelation (C). Traces denote the estimated rate for the 20 trials shown above. When the history term is included in the model, but the input is not, the GLM can still reconstruct peristimulus time histogram responses using the postspike history alone. (B) Postspike filters for the models in panel A with 95% confidence bands. Note that when input is included in the model, the filters correctly reconstruct the true (lack of) filter and there is higher uncertainty around the regions where the ISI distribution does not constrain the model.

noise) 1 and 2 (see Figure 7). For white noise input, the MLE accurately recovers the simulated postspike history filter when the input is omitted from the model, but when $\alpha = 1$ or 2 the estimates become increasingly biased (see Figures 7A and 7C). With the full model, where the input is included as a covariate, the history is recovered accurately no matter what the input statistics are. Just as in the periodic case, however, these different input statistics alter the autocorrelation, and when the input is omitted from the model, the maximum likelihood history filter simply aims to capture these patterns.

In GLMs for single-neuron dynamics, one effect of omitted variable bias is that it may lead us to misinterpret how stable a neuron's dynamics are. Even if the true history filter only reduces the neuron's firing rate following a spike (as in Figure 7C), the estimated filter can be biased upward when the input is omitted. If we were to simulate the activity of this neuron based on the biased filter, the bias could cause the neuron's rate to diverge if the rate becomes high enough. To assess the stability of the estimated postspike history effects quantitatively, here we make use of a quasi-renewal approximation analysis introduced in Gerhard, Deger, and Truccolo (2017). Given a history filter, this approach finds an approximate transfer function describing the neuron's future firing rate (output) given its recent (input) firing rate (see section 4). For all estimated models, the transfer function has a stable fixed point near the neuron's baseline firing rate. When the true input is omitted and $\alpha > 0$, the estimated history filters also have an unstable

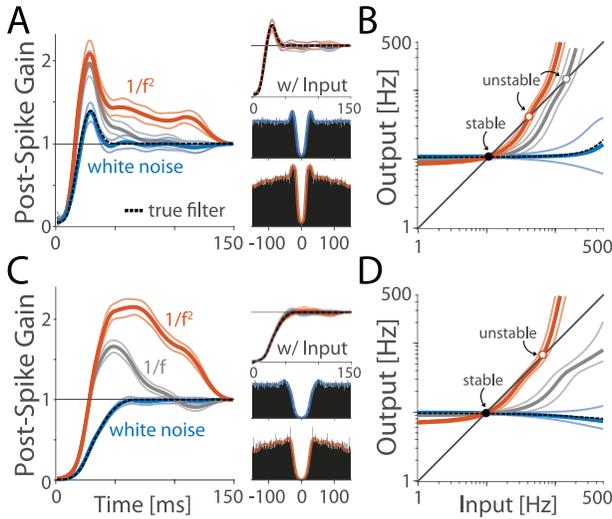


Figure 7: Postspike filters can show omitted variable bias even in a more realistic scenario. Here we simulate from a GLM with a refractory postspike filter and drive the neurons with $1/f^\alpha$ noise. Excepting the case of white noise ($\alpha = 0$), the postspike filters estimated for the GLM without input are heavily biased (A). (C) Even when the effect of the true postspike filter is to strictly decrease the firing rate, the estimated filters can increase the firing rate. (B, D) Approximate transfer functions from a quasi-renewal approximation. When the true filter is stable, the estimated filters can result in fragile dynamics.

fixed point where the neuron's firing rate will diverge if the rate exceeds this point (Gerhard et al., 2017). Here we find that omitted variable bias leads to apparent fragility (see Figures 7B and 7D). The stable region shrinks as α increases, and even when the true dynamics are strictly stable (as in Figures 7C and 7D), omitted variable bias can lead us to mistakenly conclude that the neuron has fragile dynamics.

With most extracellular spike recordings, the synaptic input that the neuron receives is unknown. However, there may also be omitted variable bias when history effects are estimated from real data. In this case, the input to a neuron can be approximated by stimulus or behavioral variables, local field potentials, or the activity of simultaneously recorded neurons (Gerhard et al., 2013; Harris et al., 2003; Kelly et al., 2010; Pillow et al., 2008; Truccolo et al., 2005; Truccolo, Hochberg, & Donoghue, 2010; Volgushev, Ilin, & Stevenson, 2015). Just as in the simulations above, including or omitting these variables can alter the estimated history effects, even though they are not as directly related to spiking as the synaptic input itself. Here we consider total population spiking activity as a proxy for synaptic

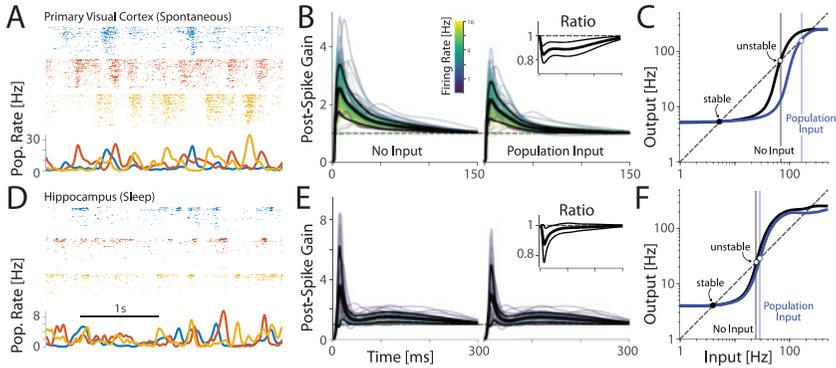


Figure 8: Postspike filters estimated from real data decrease when population activity is included as a covariate. Segments of spontaneous activity are shown for V1 (A) and during sleep for hippocampus (D). Neurons are sorted by firing rate. (B, E) Estimated postspike filters. Black lines denote the average filter (thick) and standard deviation (thin). For clarity, only filters for neurons with firing rates more than 1 Hz are shown. (C, F) Average quasi-renewal transfer functions for the same set of neurons. All neurons appear to have fragile dynamics with one stable fixed point near the neuron’s average firing rate and an unstable fixed point, beyond which the neuron’s firing rate diverges. Including population covariates increases the region of stability.

input and consider how including population activity alters the history filters when compared to a model of history alone.

We examine two data sets: spontaneous activity from primary visual cortex of an anaesthetized monkey with $n = 62$ simultaneously recorded neurons and activity from dorsal hippocampus of a sleeping rat with $n = 39$ simultaneously recorded neurons. To model population covariates, we sum the spiking of all neurons, except the one whose spiking we aim to predict, and low-pass filter the signal (see section 4). Similar to previous results (Okun et al., 2015), we find that, since neurons often have correlated fluctuations in their spiking (see Figures 8A and 8D), the population rate is a good predictor for single-neuron activity. Moreover, when we add population covariates to a GLM with postspike history effects, the history filter changes.

In the V1 data set, the postspike gain decreases by $7.8 \pm 0.5\%$ on average when population covariates are included and $14.9 \pm 0.8\%$ when considering only the first 50 ms after a spike (see Figure 8B). The effects of adding population covariates are less pronounced in the hippocampal data set. The postspike gain decreases by $2.5 \pm 0.3\%$ on average and $9.5 \pm 1.2\%$ when considering only the first 50 ms after the spike (see Figure 8E). Based on the quasi-renewal approximation, all neurons in both the V1 and hippocampal

data sets have fragile transfer functions where there is a stable fixed point (near the neuron's average firing rate) and an unstable fixed point where the neuron's rate diverges if the input becomes too strong. For V1, the average upper limit of the stable region is 80 ± 3 Hz for the models with history only and 143 ± 7 Hz for the models with population covariates (see Figure 8C). In the hippocampal data, the average upper limit of the stable region is 38 ± 6 Hz for the models with history only and 75 ± 13 Hz for the models with population covariates (see Figure 8F). Each neuron is thus apparently more stable after the population covariates are included.

As in the case studies using tuning curves, adding covariates also improves spike prediction accuracy. In the V1 data set, the average log-likelihood ratio relative to a homogeneous Poisson model is 2.2 ± 0.3 bits/s for the history model and 3.3 ± 0.3 bits/s for the model with population covariates. In hippocampus, the log-likelihood ratio is 0.9 ± 0.3 bits/s for the history model and 2.0 ± 0.5 bits/s for the model with population covariates. The larger effects in V1 are likely explained by the fact that the population rate is predictive for many more neurons here than for the hippocampal data. In the hippocampus, only 26% of the neurons have an increase of over 0.5 bits/s when the population covariates are included, compared to 85% of neurons in V1. Altogether, these results demonstrate how omitted variable bias could affect estimates of postspike history filters *in vivo*. In both data sets, we find that when population covariates are included in the GLM, spike prediction accuracy increases, postspike gain decreases, and apparent stability increases.

3 Discussion

When the goal of modeling is causal inference or understanding of biological mechanisms, the potential for biases due to omitted variables is often clear. The statistical effects of confounders (Wasserman, 2004), as well as the limits that they place on neuroscientific understanding, are widely appreciated (Jonas & Kording, 2017; Krakauer, Ghazanfar, Gomez-Marín, MacIver, & Poeppel, 2017; Yoshihara & Yoshihara, 2018). However, when the goal of modeling is to create an abstract explanation or summary of observed neural activity, the fact that omitted variables can bias these explanations is not always widely acknowledged. Here we have illustrated the potential for omitted variable bias in two types of commonly used GLMs for neural spiking activity: tuning curve models using spike counts across trials and models that capture single-neuron dynamics with a postspike history filter. In each model, adding a previously omitted variable improved spike prediction accuracy, as expected. However, what we emphasize here is that when omitted variables were included, the estimates of the original parameters changed. For three case studies using tuning curves, we found that by adding a traditionally omitted variable, tuning curves showed less modulation due to the originally included variables. In models of single-neuron

dynamics, adding omitted variables led to decreased postspike gain and greater apparent stability. Importantly, omitted variables can arise in GLMs in any situation where an omitted variable affects neural activity and the effect of the omitted variable is not independent of the included variables.

The case studies here are not unique; many studies have described how adding variables to a tuning curve or single-neuron model can improve prediction accuracy. In M1, in addition to movement speed, joint angles, muscle activity, end-point force, and many other variables also appear to modulate neural responses (Fetz, 1992; Kalaska, 2009). In addition to speed and head direction in the hippocampus, theta-band local field potential, sharp-wave ripples, and environmental features, such as borders, appear to modulate neural activity (Hartley, Lever, Burgess, & O'Keefe, 2014). And in V1, there is growing evidence that population activity (Lin, Okun, Carandini, & Harris, 2015) and nonvisual information (Ghazanfar & Schroeder, 2006) modulate neural responses. In each of these systems, neural responses are affected by many factors. Responding to many task variables may even be functional, allowing downstream neurons to more effectively discriminate inputs (Fusi, Miller, & Rigotti, 2016). In any case, it seems clear that our models do not yet capture the full complexity of neural responses (Carandini et al., 2005). By omitting relevant variables, current models are likely to be not just less accurate but also biased.

Parameter bias may be problematic in and of itself. However, omitted variable bias may also have an important effect on generalization performance. As Box (1966) noted, in a new context, the effect of the omitted variables and the relationship between the omitted and included variables may be different. Since the parameters of the included variables are biased, this change can reduce generalization accuracy. This phenomena may explain, to some extent, why tuning models fit in one condition often do not generalize to others (Graf, Kohn, Jazayeri, & Movshon, 2011; Oby, Ethier, & Miller, 2013). For models of single-neuron dynamics, omitted variable bias can also have a negative effect on the accuracy of simulations. Previous work has shown that simulating a GLM with postspike filters estimated from data often results in unstable, diverging simulations. Although several methods for stabilizing these simulations have recently been developed (Gerhard et al., 2017; Hocker & Park, 2017), one, perhaps primary, reason for this instability may be that the postspike filters are biased due to omitted synaptic input. Since estimated postspike filters may reflect not just intrinsic neuron properties but also the statistics of the input, interpreting and comparing postspike filters may be difficult. Different history parameters may be different due to intrinsic biophysics (Tripathy, Padmanabhan, Gerkin, & Urban, 2013) or due to differing input, and resolving this ambiguity will likely involve more accurately accounting for the input itself (Kim & Shinomoto, 2012).

The possibility of omitted variable bias does not mean that estimated parameters, predictions, and simulations from simplified model are useless,

but it may mean that we need to be cautious in interpreting these models and their outputs. When reporting the results of regression, in addition to avoiding describing associations with causal language, it may be generally useful to discuss known and potential confounds. Previous studies have already identified several specific cases of omitted variable bias where careful interpretation is necessary. For instance, omitted common input can bias estimates of interactions between neurons (Brody, 1999), and omitted history effects can bias receptive field estimates (Pillow & Simoncelli, 2003). In estimating peristimulus time histograms, omitting variables that account for trial-to-trial variation may cause biases (Czanner et al., 2008) or issues with identifiability (Amarasingham, Geman, & Harrison, 2015). Similarly, biases due to spike sorting errors (Ventura, 2009) could be framed as a result of omitting variables related to missing or excess spikes. Since we typically do not model or observe all the variables that affect neural activity, omitted variable problems are likely to be pervasive in systems neuroscience far beyond these specific cases.

Although we have focused on GLMs here, omitted variable bias can affect any model, and other types of model misspecification can also result in biased parameter estimates. Adding input nonlinearities (Ahrens, Paninski, & Sahani, 2008; David, Mesgarani, Fritz, & Shamma, 2009), interaction effects (McFarland, Cui, & Butts, 2013), or higher-order terms to the GLM (Berger, Song, Chan, & Marmarelis, 2010; Park, Archer, Priebe, & Pillow, 2013) may fix certain types of model misspecification, but any model that omits relevant variables is still likely to suffer from the same problems. This includes both machine learning methods that may provide better prediction accuracy than GLMs (Benjamin et al., 2017) and single-neuron models aiming to describe greater biophysical detail (Herz, Gollisch, Machens, & Jaeger, 2006). Unlike over-fitting or nonconvergence (Zhao & Iyengar, 2010), omitted variable bias will generally not be resolved by including additional data or by adding regularization. Moreover, adding one omitted variable, as we have done with the case studies here, is no guarantee that there are not other relevant variables being omitted.

One approach that could potentially reduce omitted variable bias is latent variable modeling, where the effects of unknown covariates are explicitly included (constrained by simplifying assumptions). Recent work has introduced latent variables for neural activity with linear dynamics (Kulkarni & Paninski, 2007; Paninski et al., 2010; Smith & Brown, 2003), switching dynamics (Putzky, Franzen, Bassetto, & Macke, 2014), rectification (White-way & Butts, 2017), and oscillations (Arai & Kass, 2017). These models appear to outperform GLMs on population data in retina (Vidne et al., 2012), visual (Archer, Koster, Pillow, & Macke, 2014), and motor cortices (Chase, Schwartz, & Kass, 2010; Macke et al., 2011). Inferring latent variables requires making (sometimes strong) assumptions about the nature of the variables and may require observations from multiple neurons or across multiple trials, but by approximating some of the effects of relevant omitted

variables, latent variables may reduce omitted variable bias. However, generally determining when relevant variables are omitted from a model and what those variables are is not a trivial problem.

There is a well-known aphorism from George E. P. Box: “All models are wrong, but some are useful.” The lengthier version of this quip is, “All models are approximations. Essentially, all models are wrong, but some are useful. However, the approximate nature of the model must always be borne in mind” (Box & Draper, 1987). GLMs are certainly useful descriptions of neural activity. They are computationally tractable, can disentangle the relative influence of multiple covariates, and often provide the core components for Bayesian decoders. Here we emphasize, however, one ubiquitous circumstance in systems neuroscience where the “approximate nature” of the models should be “borne in mind.” Namely, omitted variables can bias estimates of the included effects.

4 Methods

4.1 Neural Data. All data analyzed here was previously recorded and shared by other researchers through the Collaborative Research in Computation Neuroscience (CRCNS) Data Sharing Initiative (crcns.org).

Data from primary motor cortex are from CRCNS data set DREAM-Stevenson_2011 (Walker & Kording, 2013). These data were recorded using a 100 electrode Utah array (Blackrock Microsystems, 400 mm spacing, 1.5 mm length) chronically implanted in the arm area of primary motor cortex of an adult macaque monkey. The monkey made center-out reaches in a 20×20 cm workspace while seated in a primate chair, grasping a two-link manipulandum in the horizontal plane (arm roughly in a sagittal plane). Each trial for the center-out task began with a hold period at a center target (0.3–0.5 s). After a go cue, subjects had 1.25 s to reach one of eight peripheral targets and then hold this outer target for at least 0.2 s to 0.5 s. Each success was rewarded with juice, and feedback (1–2 cm diam) about arm position was displayed on-screen as a circular cursor. Spike sorting was performed by manual cluster cutting using an offline sorter (Plexon) with waveforms classified as single- or multi-unit based on action potential shape and minimum ISIs greater than 1.6 ms (yielding $n = 81$ single units). Here we model tuning curves using spike counts between 150 ms before to 350 ms after the speed reached its half-max. Average movement speed for each trial was calculated from 0 to 250 ms after the speed reached its half-max (290 trials in total). (For details of the surgery, recording, and spike sorting, see Stevenson et al., 2011.)

Hippocampal data are from CRCNS HC-3 (Mizuseki, Sirota, Pastalkova, Diba, & Buzsáki, 2013). Here we use recording sessions ec16.19.272 and ec014.215, where a Long-Evans rat was sleeping and foraging in a 180×180 cm maze, respectively. For both recordings, 12 shank silicon probes (with 8 recording sites each, 20 μm separation) were implanted in CA1

(8 shanks) and EC3-5 (4 shanks) (based on histology). Spikes were sorted automatically using KlustaKwick and then manually adjusted (Klusters), yielding 85 units for the sleep data and 117 for the open field data. For the sleep data, where we model postspike history effects, the spike trains were binned at 1 ms and the recording length was 27 min. Here we model all neurons with firing rates above 0.5 Hz ($n = 39$). For the open field data, where we model place tuning, spike trains were binned at 250 ms, and the recording length was 93 min. Place cells ($n = 68$) were selected based on having an overall firing rate below 5 Hz (to rule out interneurons), a peak firing rate above 2 Hz, and a contiguous set of pixels (after smoothing with an isotropic gaussian $\sigma = 8$ cm) of at least 200 cm² where the firing rate was above 10% of the peak rate. For details of the surgery, recording, and spike sorting see (Diba & Buzsáki, 2008; Mizuseki, Sirota, Pastalkova, & Buzsáki, 2009).

Data from primary visual cortex are from CRCNS data set PVC-11 (Kohn & Smith, 2016). Here we use spontaneous activity, during a gray screen, (monkey 1) and responses to drifting sine-wave gratings (monkey 1) both from an anaesthetized (Sufentanil: 4–18 microg/kg/hr) adult monkey (*Macaca fascicularis*). Recordings were made in parafoveal V1 (RFs within 5 degrees of the fovea) using a 96-channel multielectrode array (Blackrock Microsystems), 400 mm electrode spacing, and 1 mm depth. After automatic spike sorting and manual cluster adjustment, 87 and 106 units were recorded during spontaneous activity and grating presentation, respectively. Only neurons with waveform SNR >2 and firing rates >1 Hz were analyzed, $n = 62$ for spontaneous and $n = 90$ for grating data. For the spontaneous activity, we bin spike counts at 1 ms, and the recording length was 20 min. For the drifting grating data, we analyzed spike counts from 200 ms to 1.2 s after stimulus onset on each trial: 12 directions, 2400 trials total. Gratings had a spatial frequency of 1.3 cyc/deg, temporal frequency of 6.25 Hz, and size of 8 to 10 degrees (to cover receptive fields of all recorded neurons) and were presented for 1.25 s with a 1.5 s intertrial interval between stimuli. (For surgical, stimulus, and preprocessing details see (Kelly et al., 2010, and Smith & Kohn, 2008.)

4.2 Tuning Curve Models. For the M1 data we use a circular, cubic B-spline basis with five equally spaced knots,

$$\lambda(\theta) = \exp\left(\mu + \sum_i \beta_i g_i(\theta)\right),$$

where $g(\cdot)$ are the splines that depend on the reach direction θ , weighted by parameters β , and the parameter μ defines a baseline firing rate. To include the effect of speed, we then add three covariates,

$$\lambda(\theta, s) = \exp\left(\mu + \sum_i \beta_i g_i(\theta) + \alpha_1 s + \alpha_2 s \cos(\theta) + \alpha_3 s \sin(\theta)\right),$$

where s indicates the speed, and the parameters α allow for a multiplicative speed effect, as well as possible cosine-tuned speed \times direction interactions as in Moran and Schwartz (1999).

For place fields in hippocampus, we use isotropic gaussian radial basis functions $f(\cdot)$ equally spaced (30 cm) on a 6×6 square lattice with a standard deviation of 30 cm:

$$\lambda(x) = \exp\left(\mu + \sum_i \beta_i f_i(x)\right).$$

We find that the effect of speed is well modeled using the log-transformed speed s , and to model head direction-dependence, we use circular, cubic B-splines $g(\cdot)$ with six equally spaced knots:

$$\lambda(x, s, \theta) = \exp\left(\mu + \sum_i \beta_i f_i(x) + \gamma \log s + \sum_j \alpha_j g_j(\theta)\right).$$

For the V1 data, we again use a circular, cubic B-spline basis for the direction of the sine-wave grating (seven equally spaced knots):

$$\lambda(\theta) = \exp\left(\mu + \sum_i \beta_i g_i(\theta)\right).$$

We find that the effect of population activity is well modeled using the total log-transformed firing rate of all neurons except the one being modeled,

$$\lambda(\theta, z) = \exp\left(\mu + \sum_i \beta_i g_i(\theta) + \alpha z\right),$$

where $z = \sum_{i \neq j} \log(n_i + 1)$. In all models, to avoid overfitting, especially for low firing rate neurons, we add a small L2 penalty to the log likelihood with a fixed hyperparameter of 10^{-4} .

4.3 Postspike History Simulations and Population Rate Models. In addition to capturing tuning curves, many studies have used GLMs to describe the dynamics of single spike trains (Brillinger, 1988; Harris et al., 2003; Okatan, Wilson, & Brown, 2005; Paninski, 2004; Truccolo et al., 2005; Weber & Pillow, 2017). Here, to account for postspike history effects, we use a GLM

taking the form

$$\lambda(t) = \exp(\mu + \alpha h(t)),$$

$$n(t) \sim \text{Poisson}(\lambda(t)\Delta t),$$

where $h(t)$ denotes the vector of spike history covariates representing the recent history of spiking and μ determines a baseline firing rate. Here we assume $h_i(t) = \sum_{\tau>0} f_i(\tau)n(t - \tau)$, and we use neuron-specific, cubic B-spline bases $f(\cdot)$ whose knots are determined by the quantiles of each neuron's ISI distribution. Specifically, we choose knots spaced between 10 and 400 ms (HC) or 2 and 200 ms (V1), where the spacing follows equal percentile regions of the ISI distribution in that same range. This gives six basis functions, and coefficients α to capture the spike history. To enforce refractoriness, we fix the coefficient of the fastest basis (which peaks at 0 and ends at 10 ms) to be -5 , leaving five coefficients to be estimated.

The population rate model simply adds covariates where, for each neuron i :

$$\lambda_i(t) = \exp\left(\mu_i + \alpha_i h_i(t) + \beta_i g\left(\sum_j n_{j \neq i}(t)\right)\right),$$

$$n_i(t) \sim \text{Poisson}(\lambda_i(t)\Delta t).$$

Here we use a set of acausal gaussian filters for $g(\cdot)$ with standard deviations 20, 50, and 100 ms. Note that spikes from the neuron being modeled are excluded from the population covariates.

4.4 Stability Analysis. Here we make use of a stability analysis proposed in Gerhard et al. (2017). Briefly, we use a quasi-renewal approximation of the conditional intensity by considering the effect of the most recent spike, at time t' , and averaging over possible spike histories preceding this spike,

$$\lambda_0(t, t') = \exp(\mu + H(t - t')) \langle \exp(H * S) \rangle_{S(t < t')},$$

where $H(t - t') = \alpha f(t - t')$ and S represents the history of spiking. By assuming that S is generated from a homogeneous Poisson process with firing rate A_0 , the second term can be approximated by

$$\langle \exp(H * S) \rangle_{S(t < t')} \approx \exp\left(A_0 \int_{t-t'}^{\infty} (e^{H(u)} - 1) du\right).$$

Given this approximation, we can then estimate the interspike interval distribution as we would for a true renewal process and the steady-state distribution of interspike intervals is given by

$$P(\tau) = \exp\left(-\int_0^\tau \lambda_0(u)du\right) \lambda_0(\tau),$$

and the predicted steady-state firing rate is $f(A_0) = 1/E_{P(\tau)}[\tau]$.

To assess stability, we can then examine how the predicted steady-state firing rate depends on the assumed rate of the homogeneous Poisson process A_0 . In particular, when $f(A_0) = A_0$ the quasi-renewal model has a fixed point. To allow for external input, we incorporate the average effect of the covariates X into the conditional intensity approximation:

$$\lambda_0(t, t') = \exp(\mu + h(t - t') + \langle X\beta \rangle)(\exp(h * \mathcal{S}))_{\mathcal{S}(t < t')},$$

$$\lambda_0(t, t') = \exp(\mu + h(t - t'))(\exp(X\beta))_t(\exp(h * \mathcal{S}))_{\mathcal{S}(t < t')}.$$

Note that in general, adding inputs X will only change the stability of the model to the extent that these covariates change the estimate of h .

References

- Ahrens, M. B., Paninski, L., & Sahani, M. (2008). Inferring input nonlinearities in neural encoding models. *Network: Computation in Neural Systems*, 19(1), 35–67.
- Amarasingham, A., Geman, S., & Harrison, M. T. (2015). Ambiguity and nonidentifiability in the statistical analysis of neural codes. *Proceedings of the National Academy of Sciences of the United States of America*, 112(20), 6455–6460. <http://doi.org/10.1073/pnas.1506400112>
- Amirikian, B., & Georgopoulos, A. P. (2000). Directional tuning profiles of motor cortical cells. *Neuroscience Research*, 36(1), 73–79.
- Arai, K., & Kass, R. E. (2017). Inferring oscillatory modulation in neural spike trains. *PLoS Computational Biology*, 13(10), e1005596. <http://doi.org/10.1371/journal.pcbi.1005596>
- Arandia-Romero, I., Tanabe, S., Drugowitsch, J., Kohn, A., & Moreno-Bote, R. (2016). Multiplicative and additive modulation of neuronal tuning with population activity affects encoded information. *Neuron*, 89(6), 1305–1316. <http://doi.org/10.1016/J.NEURON.2016.01.044>
- Archer, E. W., Koster, U., Pillow, J. W., & Macke, J. H. (2014). Low-dimensional models of neural population activity in sensory cortical circuits. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 27 (pp. 343–351). Red Hook, NY: Curran.
- Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science*, 273(5283), 1868–1871. <http://doi.org/10.1126/science.273.5283.1868>

- Benjamin, A. S., Fernandes, H. L., Tomlinson, T., Ramkumar, P., VerSteeg, C., Miller, L., & Kording, K. P. (2017). *Modern machine learning far outperforms GLMs at predicting spikes*. BioRxiv:111450. <http://doi.org/10.1101/111450>
- Berger, T. W., Song, D., Chan, R. H. M., & Marmarelis, V. Z. (2010). The neurobiological basis of cognition: Identification by multi-input, multioutput nonlinear dynamic modeling: A method is proposed for measuring and modeling human long-term memory formation by mathematical analysis and computer simulation of nerve-cell. *Proceedings of the IEEE*, 98(3), 356–374. <http://doi.org/10.1109/JPROC.2009.2038804>
- Box, G. E. P. (1966). Use and abuse of regression. *Technometrics*, 8(4), 625. <http://doi.org/10.2307/1266635>
- Box, G. E., & Draper, N. R. (1987). *Empirical model-building and response surfaces*. New York: Wiley.
- Brillinger, D. R. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, 59(3), 189–200.
- Brillinger, D. R., & Segundo, J. P. (1979). Empirical examination of the threshold model of neuron firing. *Biological Cybernetics*, 35(4), 213–220. <http://doi.org/10.1007/BF00344204>
- Brody, C. D. (1999). Disambiguating different covariation types. *Neural Computation*, 11(7), 1527–1535.
- Brown, E., Barbieri, R., Eden, U., & Frank, L. (2003). Likelihood methods for neural data analysis. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach* (pp. 253–286). London: Chapman and Hall.
- Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., . . . Rust, N. C. (2005). Do we know what the early visual system does? *Journal of Neuroscience*, 25(46), 10577.
- Chase, S. M., Schwartz, A. B., & Kass, R. E. (2010). Latent inputs improve estimates of neural encoding in motor cortex. *Journal of Neuroscience*, 30(41), 13873–13882.
- Churchland, M. M., Santhanam, G., & Shenoy, K. V. (2006). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of Neurophysiology*, 96(6), 3130.
- Clarke, K. A. (2005). The phantom menace: Omitted variable bias in econometric research. *Conflict Management and Peace Science*, 22(4), 341–352. <http://doi.org/10.1080/07388940500339183>
- Clogg, C. C., Petkova, E., & Shihadeh, E. S. (1992). Statistical methods for analyzing collapsibility in regression models. *Journal of Educational Statistics*, 17(1), 51. <http://doi.org/10.2307/1165079>
- Czanner, G., Eden, U. T., Wirth, S., Yanike, M., Suzuki, W. A., & Brown, E. N. (2008). Analysis of between-trial and within-trial neural spiking dynamics. *Journal of Neurophysiology*, 99(5), 2672–2693. <http://doi.org/10.1152/jn.00343.2007>
- David, S. V., Mesgarani, N., Fritz, J. B., & Shamma, S. A. (2009). Rapid synaptic depression explains nonlinear modulation of spectro-temporal tuning in primary auditory cortex by natural stimuli. *Journal of Neuroscience*, 29(11), 3374.
- Destexhe, A., Rudolph, M., Fellous, J. M., & Sejnowski, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1), 13–24. [http://doi.org/10.1016/S0306-4522\(01\)00344-X](http://doi.org/10.1016/S0306-4522(01)00344-X)

- Destexhe, A., Rudolph, M., & Paré, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nature Reviews. Neuroscience*, 4(9), 739–751. <http://doi.org/10.1038/nrn1289>
- Diba, K., & Buzsáki, G. (2008). Hippocampal network dynamics constrain the time lag between pyramidal cells across modified environments. *Journal of Neuroscience*, 28(50), 13448–13456. <http://doi.org/10.1523/JNEUROSCI.3824-08.2008>
- Drake, C., & McQuarrie, A. (1995). A note on the bias due to omitted confounders. *Biometrika*, 82(3), 633–638. <http://doi.org/10.1093/biomet/82.3.633>
- Fernandes, H. L., Stevenson, I. H., Phillips, A. N., Segraves, M. A., & Kording, K. P. (2014). Saliency and saccade encoding in the frontal eye field during natural scene search. *Cerebral Cortex*, 24(12), 3232–3245. <http://doi.org/10.1093/cercor/bht179>
- Fetz, E. E. (1992). Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences*, 15(4), 679–690. <http://doi.org/10.1017/S0140525X00072599>
- Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, 37, 66–74. <http://doi.org/10.1016/J.CONB.2016.01.010>
- Gail, M. H., Wieand, S., & Piantadosi, S. (1984). Biased estimates of treatment effect in randomized experiments with nonlinear regressions and omitted covariates. *Biometrika*, 71(3), 431. <http://doi.org/10.2307/2336553>
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge: Cambridge University Press.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 2(11), 1527–1537.
- Gerhard, F., Deger, M., & Truccolo, W. (2017). On the stability and dynamics of stochastic spiking neuron models: Nonlinear Hawkes process and point process GLMs. *PLoS Computational Biology*, 13(2). e1005390. <http://doi.org/10.1371/journal.pcbi.1005390>
- Gerhard, F., Kispersky, T., Gutierrez, G. J., Marder, E., Kramer, M., & Eden, U. (2013). Successful reconstruction of a physiological circuit with known connectivity from spiking activity alone. *PLoS Computational Biology*, 9(7), e1003138. <http://doi.org/10.1371/journal.pcbi.1003138>
- Ghazanfar, A. A., & Schroeder, C. E. (2006). Is neocortex essentially multisensory? *Trends in Cognitive Sciences*, 10(6), 278–285. <http://doi.org/10.1016/J.TICS.2006.04.008>
- Goris, R. L. T., Movshon, J. A., & Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nature Neuroscience*, 17(6), 858–65. <http://doi.org/10.1038/nn.3711>
- Graf, A. B. A., Kohn, A., Jazayeri, M., & Movshon, J. A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. *Nature Neuroscience*, 14(2), 239–245. <http://doi.org/10.1038/nn.2733>
- Greenland, S. (1989). Modeling and variable selection in epidemiologic analysis. *American Journal of Public Health*, 79(3), 340–349. <http://doi.org/10.2105/AJPH.79.3.340>
- Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., & Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424(6948), 552–556.

- Hartley, T., Lever, C., Burgess, N., & O'Keefe, J. (2014). Space in the brain: How the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 369(1635), 20120510. <http://doi.org/10.1098/rstb.2012.0510>
- Herz, A. V. M., Gollisch, T., Machens, C. K., & Jaeger, D. (2006). Modeling single-neuron dynamics and computations: A balance of detail and abstraction. *Science*, 314(5796), 80–85. <http://doi.org/10.1126/science.1127240>
- Hocker, D., & Park, I. M. (2017). Multistep inference for generalized linear spiking models curbs runaway excitation. In *Proceedings of the 2017 8th International IEEE/EMBS Conference on Neural Engineering* (pp. 613–616). Piscataway, NJ: IEEE. <http://doi.org/10.1109/NER.2017.8008426>
- Humphrey, D. R., Schmidt, E. M., & Thompson, W. D. (1970). Predicting measures of motor performance from multiple cortical spike trains. *Science*, 170(3959), 758–762. <http://doi.org/10.1126/SCIENCE.170.3959.758>
- Jonas, E., & Kording, K. P. (2017). Could a neuroscientist understand a micro-processor? *PLoS Computational Biology*, 13(1), e1005268. <http://doi.org/10.1371/journal.pcbi.1005268>
- Kalaska, J. F. (2009). From intention to action: Motor cortex and the control of reaching movements. In D. Sternad (Ed.), *Progress in Motor Control: A Multidisciplinary Perspective* (pp. 139–178). New York: Springer.
- Kandler, S., Mao, D., McNaughton, B. L., & Bonin, V. (2017). *Encoding of tactile context in the mouse visual cortex*. BioRxiv:199364. <http://doi.org/10.1101/199364>
- Kass, R. E., Ventura, V., & Brown, E. N. (2005). Statistical Issues in the analysis of neuronal data. *Journal of Neurophysiology*, 94(1), 8–25.
- Kelly, R. C., Smith, M. A., Kass, R. E., & Lee, T. S. (2010). Local field potentials indicate network state and account for neuronal response variability. *Journal of Computational Neuroscience*, 29(3), 567–579. <http://doi.org/10.1007/s10827-009-0208-9>
- Kim, H., & Shinomoto, S. (2012). Estimating nonstationary input signals from a single neuronal spike train. *Physical Review E*, 86(5), 051903. <http://doi.org/10.1103/PhysRevE.86.051903>
- Kohn, A., & Smith, M. A. (2016). *Utah array extracellular recordings of spontaneous and visually evoked activity from anesthetized macaque primary visual cortex (V1)*. CRCNS.org. <http://doi.org/10.6080/K0NC5Z4X>
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience needs behavior: Correcting a reductionist bias. *Neuron*, 93(3), 480–490. <http://doi.org/10.1016/J.NEURON.2016.12.041>
- Kulkarni, J. E., & Paninski, L. (2007). Common-input models for multiple neural spike-train data. *Network: Computation in Neural Systems*, 18(4), 375–407.
- Lin, I.-C., Okun, M., Carandini, M., & Harris, K. D. (2015). The nature of shared cortical variability. *Neuron*, 87(3), 644–656. <http://doi.org/10.1016/J.NEURON.2015.06.035>
- Macke, J., Büsing, L., Cunningham, J., Yu, B., Shenoy, K., & Sahani, M. (2011). Empirical models of spiking in neural populations. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 24 (pp. 1350–1358). Red Hook, NY: Curran.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models* (2nd ed.). Boca Raton, FL: CRC Press.

- McFarland, J. M., Cui, Y., & Butts, D. A. (2013). Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS Computational Biology*, 9(7), e1003143. <http://doi.org/10.1371/journal.pcbi.1003143>
- McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52(1), 41–49.
- Mizuseki, K., Sirota, A., Pastalkova, E., & Buzsáki, G. (2009). Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron*, 64(2), 267–280.
- Mizuseki, K., Sirota, A., Pastalkova, E., Diba, K., & Buzsáki, G. (2013). *Multiple single unit recordings from different rat hippocampal and entorhinal regions while the animals were performing multiple behavioral tasks*. CRCNS.org. <http://doi.org/10.6080/K09G5JRZ>
- Moran, D. W., & Schwartz, a B. (1999). Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82(5), 2676–2692. <http://doi.org/citeulike-article-id:560567>
- Niell, C. M., & Stryker, M. P. (2010). Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65(4), 472–479. <http://doi.org/10.1016/J.NEURON.2010.01.033>
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34, 171–175.
- Oby, E. R., Ethier, C., & Miller, L. E. (2013). Movement representation in the primary motor cortex and its contribution to generalizable EMG predictions. *Journal of Neurophysiology*, 109(3), 666–678. <http://doi.org/10.1152/jn.00331.2012>
- Okatan, M., Wilson, M. A., & Brown, E. N. (2005). Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Computation*, 17(9), 1927–1961.
- Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barthó, P., . . . Harris, K. D. (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, 521(7553), 511–515. <http://doi.org/10.1038/nature14273>
- Omrani, M., Kaufman, M. T., Hatsopoulos, N. G., & Cheney, P. D. (2017). Perspectives on classical controversies about the motor cortex. *Journal of Neurophysiology*, 118(3), 1828–1848. <http://doi.org/10.1152/jn.00795.2016>
- Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15, 243–262.
- Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rahnama Rad, K., Vidne, M., . . . Wu, W. (2010). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29(1), 107–126. <http://doi.org/10.1007/s10827-009-0179-x>
- Park, I. M., Archer, E. W., Priebe, N., & Pillow, J. W. (2013). *Spectral methods for neural characterization using generalized quadratic models*. IN C. J. C. Burges, L. Bottou, M. Welling, Z. Ohahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 26. Red Hook, NY: Curran.
- Park, I. M., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience*, 17(10), 1395–1403. <http://doi.org/10.1038/nn.3800>

- Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys*, 3, 96–146. <http://doi.org/10.1214/09-SS057>
- Pillow, J. (2007). Likelihood-based approaches to modeling the neural code. In K. Doya, S. I. Kenji, A. Pouget, and R. P. N. Rao (Eds.), *Bayesian brain: Probabilistic approaches to neural coding* (pp. 53–70). Cambridge, MA: MIT Press.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., & Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207), 995–999.
- Pillow, J. W., & Simoncelli, E. P. (2003). Biases in white noise analysis due to non-Poisson spike generation. *Neurocomputing*, 52–54, 109–115. [http://doi.org/10.1016/S0925-2312\(02\)00822-6](http://doi.org/10.1016/S0925-2312(02)00822-6)
- Putzky, P., Franzen, F., Bassetto, G., & Macke, J. H. (2014). A Bayesian model for identifying hierarchically organised states in neural population activity. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 27. Red Hook, NY: Curran.
- Reimer, J., Froudarakis, E., Cadwell, C. R., Yatsenko, D., Denfield, G. H., & Tolias, A. S. (2014). Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron*, 84(2), 355–362. <http://doi.org/10.1016/j.NEURON.2014.09.033>
- Runyan, C. A., Piasini, E., Panzeri, S., & Harvey, C. D. (2017). Distinct timescales of population coding across cortex. *Nature*, 548(7665), 92–96. <http://doi.org/10.1038/nature23020>
- Rust, N. C., & Movshon, J. A. (2005). In praise of artifice. *Nature Neuroscience*, 8(12), 1647–1650. <http://doi.org/10.1038/nn1606>
- Shmueli, G. (2010). To explain or to predict? *Statistical Science*, 25, 289–310. <http://doi.org/10.2307/41058949>
- Smith, A. C., & Brown, E. N. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15(5), 965–991.
- Smith, M. A., & Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience*, 28(48), 12591–12603.
- Stevenson, I. H., Cherian, A., London, B. M., Sachs, N. A., Lindberg, E., Reimer, J., . . . Kording, K. P. (2011). Statistical assessment of the stability of neural movement representations. *Journal of Neurophysiology*, 106(2), 764–774. <http://doi.org/10.1152/jn.00626.2010>
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2018). *Spontaneous behaviors drive multidimensional, brain-wide population activity*. [BioRxiv:306019](https://doi.org/10.1101/306019). <http://doi.org/10.1101/306019>
- Tripathy, S. J., Padmanabhan, K., Gerkin, R. C., & Urban, N. N. (2013). Intermediate intrinsic diversity enhances neural population coding. *Proceedings of the National Academy of Sciences of the United States of America*, 110(20), 8248–8253. <http://doi.org/10.1073/pnas.1221214110>
- 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.
- Truccolo, W., Hochberg, L. R., & Donoghue, J. P. (2010). Collective dynamics in human and monkey sensorimotor cortex: Predicting single neuron spikes. *Nature Neuroscience*, 13(1), 105–111.

- Ventura, V. (2009). Traditional waveform based spike sorting yields biased rate code estimates. *Proceedings of the National Academy of Sciences*, *106*(17), 6921.
- Vidne, M., Ahmadian, Y., Shlens, J., Pillow, J. W., Kulkarni, J., Litke, A. M., . . . Paninski, L. (2012). Modeling the impact of common noise inputs on the network activity of retinal ganglion cells. *Journal of Computational Neuroscience*, *33*(1), 97–121. <http://doi.org/10.1007/s10827-011-0376-2>
- Volgushev, M., Ilin, V., & Stevenson, I. H. (2015). Identifying and tracking simulated synaptic inputs from neuronal firing: Insights from in vitro experiments. *PLoS Computational Biology*, *11*(3), e1004167. <http://doi.org/10.1371/journal.pcbi.1004167>
- Walker, B., & Kording, K. (2013). The database for reaching experiments and models. *PLoS One*, *8*(11), e78747. <http://doi.org/10.1371/journal.pone.0078747>
- Walsh, R. N., & Cummins, R. A. (1976). The open-field test: A critical review. *Psychological Bulletin*, *83*(3), 482–504. <http://doi.org/10.1037/0033-2909.83.3.482>
- Wasserman, L. (2004). *All of statistics*. New York: Springer. <http://doi.org/10.1007/978-0-387-21736-9>
- Weber, A. I., & Pillow, J. W. (2017). Capturing the dynamical repertoire of single neurons with generalized linear models. *Neural Computation*, *29*(12), 3260–3289. http://doi.org/10.1162/neco_a_01021
- Whiteway, M. R., & Butts, D. A. (2017). Revealing unobserved factors underlying cortical activity with a rectified latent variable model applied to neural population recordings. *Journal of Neurophysiology*, *117*(3), 919–936. <http://doi.org/10.1152/jn.00698.2016>
- Yoshihara, M., & Yoshihara, M. (2018). “Necessary and sufficient” in biology is not necessarily necessary: Confusions and erroneous conclusions resulting from misapplied logic in the field of biology, especially neuroscience. *Journal of Neurogenetics*, *32*(2), 53–64. <http://doi.org/10.1080/01677063.2018.1468443>
- Zhao, M., & Iyengar, S. (2010). Nonconvergence in logistic and Poisson models for neural spiking. *Neural Computation*, *22*(5), 1231–1244.

Received April 18, 2018; accepted August 6, 2018.