A probabilistic framework for decoding behavior from in vivo calcium imaging data

Guillaume Etter[1], Frederic Manseau[1], Sylvain Williams[1,*]

[1]McGill University & Douglas Mental Health University Institute, Montreal, QC, Canada

* Correspondance:
Sylvain Williams
sylvain.williams@mcgill.ca

Keywords: calcium imaging, decoding, Bayesian inference, hippocampus, spatial coding

Abstract

Understanding the role of neuronal activity in cognition and behavior is a key question in neuroscience. Previously, *in vivo* studies have typically inferred behavior from electrophysiological data using probabilistic approaches including bayesian decoding. While providing useful information on the role of neuronal subcircuits, electrophysiological approaches are often limited in the maximum number of recorded neurons as well as their ability to reliably identify neurons over time. This can be particularly problematic when trying to decode behaviors that rely on large neuronal assemblies or rely on temporal mechanisms, such as a learning task over the course of several days. Calcium imaging of genetically encoded calcium indicators has overcome these two issues. Unfortunately, because calcium transients only indirectly reflect spiking activity and calcium imaging is often performed at lower sampling frequencies, this approach suffers from uncertainty in exact spike timing and thus activity frequency, making rate-based decoding approaches used in electrophysiological recordings difficult to apply to calcium imaging data. Here we describe a simple probabilistic framework that can be used to robustly infer behavior from calcium imaging recordings. Our method discriminates periods of activity and periods of inactivity to compute conditional and unconditional probabilities. Neuronal activity can then be described in terms of joint probability, specificity, and confidence. We next devise a simple method to decode behavior from calcium activity and propose different metric to quantify decoding accuracy. Finally, we show that neuronal activity can be predicted from behavior, and that the accuracy of such reconstructions can guide the understanding of relationships that may exist between behavioral states and neuronal activity.

A probabilistic framework for decoding behavior from in vivo calcium imaging data

Introduction

Early *in vivo* studies have established relationships between external variables and neuronal activity, including (but not restricted to) auditory information in the auditory cortex (Katsuki et al., 1956), visual stimuli in the visual cortex (Hubel and Wiesel, 1962), and spatial information in the hippocampus (O'Keefe and Dostrovsky, 1971). Based on the widely influential information theory (Shannon, 1948), it has previously been proposed that neurons can act as 'communication channels' between physiological variables (input) and spike trains (output) (Richmond and Optican, 1990; Richmond et al., 1990; Skaggs et al., 1993). In addition to providing metrics to quantify the extent to which external variables can be encoded in neurons, these studies laid the first foundations in establishing computational tools to predict animal behavior merely using neuronal activity. This process, termed decoding, is essential in understanding the role of neuronal activity in behavior, and the success rate of predictions can be used as a metric of knowledge of a given system. Among techniques that have been described in this context, Bayesian decoding in particular has been relatively popular and widely used (Zhang et al.; Brown et al., 1998; Gerwinn, 2009; Quian Quiroga and Panzeri, 2009; Koyama et al., 2010).

While the literature on *in vivo* neuronal physiology has been largely dominated by electrophysiological studies, calcium imaging methods have recently gained popularity. Originally performed at the single cell level with the aid of calcium sensors (Grynkiewicz et al., 1985; Persechini et al., 1997), calcium imaging can now be performed *in vivo,* in large neuronal assemblies, and over very long periods of time (Ziv et al., 2013; Sheintuch et al., 2017; Gonzalez et al., 2019). These major improvements coincided with the development of genetically encoded calcium indicators (GECI), including the popular GCaMP (Nakai et al., 2001; Tian et al., 2009; Ohkura et al., 2012; Chen et al., 2013). In the recent years, calcium imaging methods have seen the development of various computational tools that solve the problem of signal extraction from raw calcium imaging video recordings. In particular, several groups have proposed open-source software codes to perform fast, recursive motion correction (Pnevmatikakis and Giovannucci, 2017), offline (Pnevmatikakis et al., 2016; Zhou et al., 2018) and online (Giovannucci et al., 2017) extraction of neuronal spatial footprints and their associated calcium activity, temporal registration of neurons across days (Sheintuch et al., 2017), and complete pipelines have been made available (Giovannucci et al., 2018). The aforementioned open source codes have significantly facilitated the analysis of calcium imaging datasets. Most often, one of the objectives when using such tool is to understand the neural basis of behaviors. Unfortunately, there are only few open source analysis toolbox that can relate calcium imaging data to behavior to this day (Tegtmeier et al., 2018; www.miniscope.org). While these useful analytical tools allow the exploration of relationships between calcium signals and behavior, they are mostly restricted to visualization and correlation. Ultimately, a thorough understanding of a given system would translate into perfect predictions of neuronal activity given the behavioral state, and vice versa.

While calcium imaging does not allow the determination of exact spike timing, some methods have been proposed to better approximate spiking activity from calcium imaging data by deconvolving calcium transients (Deneux et al., 2016; Pachitariu et al., 2018; Rahmati et al., 2018). Consequently, one strategy that can be employed is to first estimate deconvolution parameters from ground truth data (e.g. *in vitro* unit recording in brain slices combined with calcium imaging) to then apply them to recordings performed *in vivo*. However, one major caveat with this approach is that physiological responses can differ greatly between *in vivo* and *in vitro* conditions (Belle et al., 2018) leading to erroneous parameter estimation. Another obstacle to using deconvolved signals and estimated spikes to decode calcium activity is that the very nature of calcium imaging does not allow to estimate

A probabilistic framework for decoding behavior from in vivo calcium imaging data

exact spike timing. While unit recordings are typically done at sampling rates exceeding 10 KHz, 1-photon microendoscopes used in freely moving animals usually sample images at 30 Hz (30 frames per second) or less, and spike trains will generally be associated with large calcium transients of varying size and duration. Consequently, one could for example successfully estimate that a neuron fires 10 action potentials based on the observation of a single calcium transient, however the exact timing of each spike would remain unknown, and could happen anywhere within a ~33 ms window (for calcium imaging performed at 30 Hz).

Importantly, another issue encountered when performing calcium imaging with GCaMP is photobleaching, which leads to a progressive loss of signal due to the destruction of fluorescent proteins that report calcium influx. Unlike electrophysiological unit recordings that can be performed for several hours, calcium imaging is thus typically done during shorter windows of time. While it is possible to follow GCaMP-positive cell assemblies over months (Ziv et al., 2013; Sheintuch et al., 2017), each recording session has to be limited in duration to avoid photobleaching. This results in low sampling that can be problematic when trying to associate neuronal activity with a certain behavior: some behavioral states can be over- or underrepresented and concurrently, calcium activity can be too sparse to establish tuning curves of neuronal activity and make conclusions.

Here we propose simple analytical methods to relate calcium activity to behavior by (1) extracting periods of activity in calcium imaging data without approximating spike timing and subjecting actual data to null hypothesis testing in order to solve the problem of low sampling, (2) performing Bayesian decoding, and (3) reconstructing neuronal activity from behavior to assess the quality of neuronal coding.

Results

1. Establishment of probabilistic neural tuning curves

To demonstrate the usefulness of our method, we performed calcium imaging in a well characterized system: CA1 pyramidal cells of the dorsal hippocampus (fig. 1a). These neurons are known to display spatial tuning and are often referred to as place cells (O'Keefe and Dostrovsky, 1971). We trained a mouse to run back and forth on a 100 cm long linear track by providing sucrose water rewards at each end of the track and scheduling homecage access to water every day (fig. 1b). We recorded ~400 neurons in these conditions (fig. 1c). After extracting neuronal spatial footprints (fig. 1d), we visualized corresponding calcium activity along with the position and locomotor speed of the animal (fig. 1e). Previous studies have shown that immobility periods are associated with replay of experience (Foster and Wilson, 2006; Diba and Buzsáki, 2007; Davidson et al., 2009). In order to focus on the spatial tuning curves of CA1 neurons, we therefore excluded periods of immobility ($< 5cm.s^{-1}$) that could potentially contain periods of neuronal activity that reflect internal, rather than external variables.

1.1. Conditional probabilities

In order to compute probabilities that will be used in latter analyzes of tuning curves, we sought to discriminate periods of activity versus inactivity (fig. 1f). To this end, we devised a simple binarizing method where raw calcium signals are first filtered (to remove high frequency fluctuations that could erroneously be detected as transient rise periods), and we considered periods of activity as following the two criteria: (1) the signal amplitude of a normalized trace has to be above 2 standard-deviations, and (2) the first order derivative has to be positive (thus corresponding to a transient rise period).

Following the binarization of raw calcium traces, we propose a probabilistic framework to describe how the activity of a neuron encodes a certain behavior or state (fig. 1g). To this end, we can first compute the probability of a neuron to be active P(A) using the following formula:

$$(1) \qquad P(A) \; = \; \frac{time\ active}{total\ time}$$

P(A) only informs on the activity rate of a neuron over the course of a recording session. We can also compute the probability of spending time in a given behavioral state:

$$(2) \qquad P(S) \; = \; \frac{time\ in\ state}{total\ time}$$

P(S) can be informative on whether the distribution of behavioral states is homogeneous or inhomogeneous, which can potentially lead to biases in further analyzes. Both P(A) and P(S) are unconditional probabilities that will be used later when performing Bayesian inference.

1.2. Joint probabilities

We can then compute the joint probability that a given cell is active while the animal is in a given state:

$$(3) \qquad P(A \cap S) \; = \; \frac{time\ active\ while\ in\ state}{time\ in\ state}$$

$P(A \cap S)$, or tuning curve, is more informative as it allows a direct quantification of firing probabilities under specific states, and is easy to interpret. For instance, a joint probability value of 0.8 means that a cell is active 80% of the time when the animal is in a given behavioral state. It is noteworthy that the complement of $P(A \cap S)$ can easily be derived:

$$(4) \qquad P(A^c \cap S) \; = \; 1 - P(A \cap S)$$

$P(A^c \cap S)$ can be particularly helpful to describe cells that are active in most states, but inactive in very specific states.

In our example, we isolated running periods when the mouse was running towards the right hand side of the linear track (fig. 1h), and divided the track in 3 cm bins. Each bin thus represents a given state, and while visualizing the activity of neuron #4 for each run, it is apparent that this cell displays some spatial tuning (fig. 1i). We thus computed $P(A \cap S)$ for that cell and found a peak joint probability value of 0.78 at ~64.5 cm from the left hand side of the track (fig. 1j).

1.3. Probability density function (PDF)

It is noteworthy that while $P(A \cap S)$ and $P(A^c \cap S)$ are very informative on the relationship between cell activity and behavior, they do not necessarily give much information about state specificity. In our example, a cell could very well be active 78% in every behavioral state While this is quite a high value, it would be rather hard to make conclusion on whether it would encode specific external variables or not. Therefore, another step can be taken to normalize these values by simply dividing each joint probability by the sum of joint probabilities for every state (for N number of states). Because the sum of all values of a

normalized distribution is 1, this type of representation is also termed probability density function or PDF:

$$(5) \qquad PDF(i) = \frac{P(A \cap S_x)}{\sum\limits_{i=1}^{N} P(A \cap S_i)}$$

Where i is a given state (e.g. a particular location on the linear track), PDF informs on the specificity of neuronal activity against behavioral states, regardless of the activity rate of the neuron. It is also possible to further summarize the state-specificity of neuronal coding with a single value. The Kullback-Leibler divergence (KLD) can be used to compare our PDF to a uniform distribution U:

$$(6) \qquad D_{KL} = \sum\limits_{i=1}^{N} PDF(i) \, log_2 \frac{PDF(i)}{U(i)}$$

Because a uniform distribution is equivalent to a PDF with the same value $\frac{1}{N}$, equation (5) can be simplified to:

$$(7) \qquad D_{KL} = \sum\limits_{i=1}^{N} PDF(i) \, log_2(PDF(i) \times N)$$

Since we use a log base 2, equation (6) and (7) can be expressed in bits of information. A helpful analogy would be to consider the number of coin flips required to accurately predict the position of the mouse on the linear track given that the cell is active. In the event where a neuron is equally active during every state, the PDF would be flat (i.e. it would take the same value at every behavioral state) and the KLD would be equal to zero. In our example, cell #4 displays 2.21 bits of information (fig. 1k).

1.4. Confidence testing

One current issue with calcium imaging is photobleaching, which prevents long recordings and thus restricts the sampling of both neuronal activity and behavioral data. Experimenters can thus be frequently faced with one of four cases: first, sampling of both behavior and neuronal activity are sufficient, and there is no apparent relationship between these two variables (fig. 1l, case 1). Secondly, sampling is sufficient and there is a very clear relationship between behavior and neuronal activity (fig. 1l, case 2). Thirdly, sampling is too low to observe a clear phenomenon (not enough coverage of behavioral states, sparse neuronal activity; fig. 1l, case 3). Lastly, behavioral sampling is sufficient, but neuronal activity is sparse and while there is an apparent relationship between behavior and neuronal activity, there is no absolute confidence that this could indeed be the case (fig. 1l, case 4). One solution we propose to this problem is to confront the actual data to a null hypothesis that there is no relationship between behavior and neuronal activity. To this end, we generated a distribution of tuning curves that are computed from the actual calcium trace, but that has been shifted in time so that any potential relationship to behavior is lost. The shuffling method we propose is to take a random time point in the data, and swap the two segments (AB becomes BA; fig. 1m). We recommend this shuffling method because it preserves the temporal nature of calcium transients and leads to more conservative results, as opposed to a complete randomization of every data point which often gives rise to very non-physiological data. The choice of a null hypothesis should however be determined carefully depending on

A probabilistic framework for decoding behavior from in vivo calcium imaging data

the nature of the question asked. In our example, shuffling the calcium trace temporally gave
rise to much lower joint probabilities, with lower spatial specificity (fig. 1n). We performed n
= 1000 random permutations and computed the mean and SEM, and compared it to our actual
tuning curve (fig. 1o). To convert these values to a normalized score that could be used to
describe the confidence level of our results, we can z score our actual data using parameters of
the shuffled distribution:

$$(8) \qquad z = \frac{a - \bar{s}}{(\sigma_{\bar{s}})}$$

where a is the actual tuning curve, $\bar{s}$ the average tuning curve of the surrogate data, and $\sigma_{\bar{s}}$
the standard-deviation of the surrogate data. This z-scored tuning curve (fig. 1p) then
expresses the degree of divergence from the surrogate data. Typically, values above 2 (or
bellow -2) are considered significantly different from a random distribution, which represents
our null hypothesis. We can also visualize the distance between our peak joint probability and
the shuffled distribution to assess the difference between our data and the probability
distribution of values under the null hypothesis (fig. 1q). We compared these three different
metrics (joint probability, PDF, and z-score values) for a variety of cells (supplementary fig.
1a-d) and confirmed that cell that did not classify as place cells displayed peak z-scores under
2. Importantly, peak joint probability and z-scores were positively correlated (Pearson
correlation, $R^2 = 0.81$, P < 0.000, n = 417 neurons; supplementary fig. 1e).

2. Bayesian decoding of behavior from calcium imaging data
2.1. Computing posterior probabilities

Extracting tuning curves for each individual neuron can shed light about their activity
pattern but does not fully explain a particular behavior. Importantly, the temporal coordination
of large neuronal assemblies is likely to provide more information about the specificity of
generated behaviors. In our example, we would like to understand the relationship between
location (without discriminating left/right trajectories at first) and the activity patterns of a
large (~400) cell assembly. To this end, we develop a simple algorithm based on classical
Bayesian inference, that typically takes the general form:

$$(9) \qquad P(H|E) = \frac{P(E|H) \times P(H)}{P(E)}$$

255 where $P(H|E)$ represents the posterior probability of H given E, $P(E|H)$ is the probability of E
256 given H, $P(H)$ is the prior probability and is an estimate of the probability of the hypothesis,
257 and $P(E)$ is the probability of the evidence. In our case, we would like to estimate the
258 probability that a mouse is in a given behavioral state (here, a particular location on the linear
259 track), given the state of each cell (active or inactive). Our equation then takes the form:

$$(10) \qquad P(S|A) = \frac{P(A|S) \times P(S)}{P(A)}$$

where $P(S|A)$ is the probability of the mouse to be in a given state S (in our example, a
specific location) given that a cell is active, $P(A|S)$ is the probability that the cell is active
while in a given state S, $P(S)$ is the prior probability that the mouse is in a given state, and
$P(A)$ the probability that a cell is active (this term will also be referenced as the 'bias' later on).
Both $P(A)$ and $P(S)$ are unconditional probabilities. In our linear track example, $P(A|S)$
corresponds to the joint probability $P(A \cap S)$ that a cell fires in a given location (fig. 2a).

P(S) can be measured directly (in our case, it is the general likelihood of finding the mouse in
any given location; fig. 2a, teal line in the bottom panel) or kept uniform. In the latter case, we
make no prior assumption about the potential location of the mouse on the linear track and
attribute equal probability for each location (fig. 2a, orange line in the bottom panel). Note
that the inactivity of a cell can also be used to determine the state of the mouse. Because of
the dichotomy of cell states (active or inactive), the corresponding equation can be derived:

$$(11) \quad P(S|A^c) = \frac{[1 - P(A|S)] \times P(S)}{1 - P(A)}$$

where *P(S|A^c)* represents the probability of the mouse being in a given state *S* given that a cell
is inactive. To predict the mouse state using multiple neurons, it is more efficient to take the
product (rather than the sum) of *a posteriori* probabilities, because null values translate into
an absolute certainty that the mouse cannot be in a given state considering the activity of a
given neuron. It is therefore crucial that tuning curves are properly established, otherwise
decoding could become impossible due to null factorization. We can then rewrite our equation
to include tuning curves from multiple neurons:

$$(12) \quad y = \prod_{i=1}^{N} \frac{P(A|S) \times P(S)}{P(A)}$$

where *y* is a vector of *a posteriori* behavioral states and *N* corresponds to the number of
neurons used. In our example, we can measure the activity of every cell at a given time point
(fig. 2b), derive the associated tuning curves (fig. 2c, top panel), and corresponding posterior
location probability (fig. 2c, bottom panel). Importantly, while equation (12) is fundamentally
correct, the repeated product of small values (such as probability values that are comprised
between 0 and 1) will lead to numerical underflow when computed on most softwares
available currently. Although, this is not a problem when decoding activity from a small
number of cells, numerical underflow will prevent decoding activity from large sets of cell
assemblies. One solution to this problem is to perform computations on a log scale.
Additionally, using *exp(x)-1* and *log(1+x)* allows very small values to converge toward *x*
instead of 0. Our equation can then take the form:

$$(13) \quad y = exp\left[\sum_{i=1}^{N} log\left(1 + \frac{P(A|S) \times P(S)}{P(A)}\right) - 1\right]$$

302 Finally, while *y* is a vector of states (corresponding to specific locations on the linear track in
303 our example), the mouse can only be in one location at a given time point. We can thus
304 convert y to a probability density function for each time point, and take the maximum *a*
305 *posteriori* (MAP) as the decoded state:

$$(14) \quad \hat{y} = arg\ max\ exp\left[\sum_{i=1}^{N} log\left(1 + \frac{P(A|S) \times P(S)}{P(A)}\right) - 1\right]$$

309 where $\hat{y}$ is the decoded (i.e. most likely) state.

311 **2.2. Assessing decoding accuracy**

A probabilistic framework for decoding behavior from in vivo calcium imaging data

In our example, we can compute the posterior probabilities for each timestep based on neuronal activity, and compare the actual versus decoded location on the linear track (fig. 2d). To visualize which states are associated with better/worse decoding error, we can compute a confusion matrix, which expresses the portion of time points where the actual state was successfully decoded (fig. 2e). This representation is also useful to identify which states are more difficult to decode. While confusion matrices are useful, they are not practical when it comes to summarizing decoding accuracy for large datasets and perform statistical analyzes. We thus propose two metrics: (1) decoding agreement, and (2) decoding error. We define decoding agreement as the portion of time where the exact state of the mouse was successfully decoded:

$$(15) \quad decoding\ agreement\ =\ \frac{time\ points\ successfully\ decoded}{Total\ time}$$

Therefore, decoding agreement is a value between 0 and 1. For instance, a value of 0.5 means that 50 % of time points have been successfully decoded. This approach is quite conservative: when the decoded state is only one bin away from the actual behavioral state, it would lead to a null decoding agreement while still being close to reality. Therefore, another metric commonly used in decoding analyzes is decoding error, which is the distance between the decoded behavior and the actual behavior. Note that in our case, the distance is explicitly euclidean and can be expressed in cm. For one-dimensional data, equation (16) can be used to compute decoding error:

$$(16) \quad decoding\ error\ =\ |decoded\ state\ -\ actual\ state|$$

Decoding error may or may not be useful depending on the variables under study. For instance, in the case of auditory stimuli, the distance between tone frequencies might not necessarily be as meaningful as an actual spatial distance, as it is the case in our example. Moreover, its computation will be different for two-dimensional space, or head orientation, to list a few of the variables commonly studied. Importantly, to assess decoding accuracy, it is recommended not to test the decoder on epochs that were used to train the Bayesian decoder. Some epochs, such as periods of immobility in our case, can be excluded for both training and testing altogether. We propose here to train and test Bayesian decoders on non-overlapping sets of random epochs, repeat the process several times, and compute the average decoding agreement and decoding error (fig. 2g). Using this approach, we found in our conditions that using a uniform prior led to higher decoding agreement (0.37 ± 0.002, n = 30 trials; data expressed as mean ± SEM) compared to using observed prior probabilities (0.07 ± 0.004, n = 30 independent trials), or replacing the biasing term P(A) by 1 (condition which we term 'unbiased' here; 0.07 ± 0.001, n = 30 independent trials; 1ANOVA, $F_{(2, 87)} = 4521$, P < 0.0001; fig. 2h). Similarly, decoding error was lower using a uniform prior (8.12 ± 0.08 cm, n = 30 independent trials) compared to using an observed prior (13.18 ± 0.15 cm, n = 30 independent trials) or in unbiased conditions (49.34 ± 0.08 cm, n = 30 trials; $F_{(2, 87)} = 44710$, P < 0.0001; fig. 2i).

2.3. Adding temporal constraints

Although decoding can be performed for each individual time point, this temporal independence can lead to spurious results (see decoded trajectory in fig. 2d, pink line in the bottom panel). Rationaly, if the mouse is at one end of the linear track, it is extremely unlikely to be found at the opposite end on the next frame. There are several ways to solve this

problem and improve state estimation. A first method could to build a transition matrix (such as one that would be used in a HMM), and attribute null values to impossible transitions (such as going from one end of the linear track to the other), as well as uniform probabilities to adjacent states. One could then replace the observed or uniform prior P(S) by the appropriate realistic transition values at each individual timestep. Another method is temporal smoothing. In that case, we would decode the mouse state using the following equation:

$$(17) \quad \hat{y} = arg\ max\ exp\ [\sum_{i=0}^{t} log(1 + x_t) - 1]$$

where t is a given timestep, and x_t the a posteriori distribution of states at timestep t. Note that fundamentally, the exp-sum-log is again only used to avoid numerical underflow. The advantage of this method is that it does not require to determine transition probabilities experimentally. For instance, if the state under study is location, the mouse cannot travel from one end to the other over the course of one timestep. On the other hand, if the state represents head direction (or any circular variable), the two ends of the state vector (0 and 360º) represent similar states, so such transition would be very likely. Because this adjacence would be represented in the individual neurons tuning curves, applying temporal filtering using the equation (17) would lead to a better reconstruction regardless of the behavior under study. On the other hand, if transition matrices are established *a priori*, such property would have to be assumed or determined experimentally. In our conditions, temporal filtering can greatly improve reconstruction and remove erratic 'jumps' that can sometimes occur during decoding (fig. 3a).

2.4. Parameter optimization

It is possible to optimise every parameter to find the best conditions to decode neuronal activity. For instance, we performed decoding on 30 sets of random and non-overlapping epochs using several temporal filtering values ranging from 0 (no filtering) to 2 s, and found that better reconstructions could be achieved using a 0.5 s filtering window, leading to smaller decoding errors (4.73 ± 0.04 cm, n = 30 independent trials per filtering window; 1ANOVA, $F_{(8, 261)}$ = 1002, P < 0.0001; fig. 3b). Interestingly, the bigger the temporal filtering window, the lower the decoding agreement (Pearson correlation, R^2 = 0.96, P < 0.0001, n = 30 independent trials per filtering window; fig. 3c). As expected, the more cells used during the reconstruction, the lower the decoding error (Pearson correlation, R^2 = 0.75, P < 0.0012, n = 30 independent trials per cell number; fig. 3d). Likewise, a higher number of cells was associated with higher decoding agreement (Pearson correlation, R^2 = 0.87, P < 0.0001, n = 30 independent trials per cell number; fig. 3e). We also tested the influence of the training/testing length ratio on reconstruction accuracy and found that good reconstruction can be achieved by using testing epochs that represent beyond 30 % of the total experiment length. Mean decoding error decreased as the training set portion increased (Pearson correlation, R^2 = 0.64, P = 0.01, n = 30 independent trials per training set portion tested; fig. 3f), while mean decoding agreement increased (Pearson correlation, R^2 = 0.79, P = 0.0013, n = 30 independent trials per training set portion tested; fig. 3g). We next assessed the robustness of tuning curves to random noise. To this end, we computed tuning curves as described previously, then replaced a portion (between 0 and 1, with 0.1 incremental steps) of the tuning curves data with random probability values (fig. 3h). Addition of noise was correlated with decreased decoding agreement (Pearson correlation, R^2 = 0.80, P = 0.0014, n = 30 independent trials per noise amount; fig. 3i), and increased decoding error (Pearson

correlation, $R^2 = 0.77$, $P = 0.0004$, n = 30 independent trials per noise amount; fig. 3j). Finally, we tested the impact of smoothing tuning curves on decoding accuracy. Gaussian smoothing is often performed in the context of place cell studies, presumably to improve the visualization of assumed place fields (O'Keefe and Burgess, 1996; Hetherington and Shapiro, 1997). In our conditions, we found that gaussian smoothing of tuning curves was systematically associated with decreased decoding agreement (Pearson correlation, $R^2 = 0.92$, $P = 0.0025$; n = 30 independent trials per gaussian sigma value; fig. 3l), together with increasing decoding error (Pearson correlation, $R^2 = 0.97$, $P = 0.0003$, n = 30 independent trials per gaussian sigma value; fig. 3m).

2.5. Optimal method to binarize neuronal activity

In our conditions, we used a simple binarizing algorithm that transformed rising periods of calcium transients into periods of activity. We compared this method to a simple z-score threshold where all activity above 2 standard-deviations is considered active, and deconvolved signal, where all values above zero are considered as periods of activity. To quantify the accuracy of these methods, we performed *in vitro* electrophysiological recordings in the cell attached configuration, in conjunction with 1-photon calcium imaging (supplementary fig. 2a). We extracted calcium transients from the recorded cell (supplementary fig. 2b) so as to contrast these signals with ground truth spiking activity (supplementary fig. 2c). Interestingly, calcium transients appeared much longer in these conditions, and our binarizing method only matched the later portion of transients rising periods (supplementary fig. 2d). On the other hand, using a deconvolved trace to estimate neuronal activity resulted in a higher number of action potentials successfully detected as active periods (0.94 ± 0.032) compared to our binarizing algorithm (0.49 ± 0.067) or a simple z-score threshold (0.65 ± 0.075; 1ANOVA, $F_{(2, 108)} = 13.71$, $P < 0.0001$, n = 37 detection windows; supplementary fig. 2e). Furthermore, both the portion of true negatives (epochs binarized as inactive, that indeed contained no action potential) and the portion of true positives (epochs binarized as active, that indeed contained at least one action potential) were comparable between methods (supplementary fig. 2f & g respectively).

Interestingly, these *in vitro* results did not compare to our *in vivo* conditions. When computing tuning curves for the neuron presented in fig. 1, using a simple threshold resulted in a larger place field, while binarizing data from a deconvolved trace resulted in two peaks (supplementary fig. 3a). While there are no ground truth data to conclude which method is best to compute tuning curves, decoding analyzes can shed a light on this question, because animal behavior can be used a ground truth data (the higher the decoding accuracy, the closer to ground truth). We thus trained a Bayesian decoder using tuning curves computed from binarized activity derived using a simple z-score threshold, a deconvolved trace, or using our binarizing method. We found that using both our binarizing method (4.74 ± 0.0039 cm) or a deconvolved trace (4.81 ± 0.048 cm) led to lower decoding errors compared to using a simple threshold (5.18 ± 0.051 cm, $F_{(2, 87)} = 26.22$, $P < 0.0001$, n = 30 independent trials for each binarizing method.)

2.5. Decoding two-dimensional behavioral variables

The decoding method presented above is scalable to a large variety of behaviors. However, sometimes it can be useful to represent behaviors in more than one dimension. This is for instance the case with spatial location in larger environments. We will now show that the principles presented above can easily be translated to more dimensions. To this end, we recorded neuronal activity using calcium imaging in a mouse exploring an open-field for the first time. Calcium transients are then extracted and binarized, along with the x and y mouse

position (fig. 4a). It is possible to plot periods of activity of one cell in 2D space, and color code this activity to visualize the stability of such activity in time/space (fig. 4b). Relative occupancy (fig. 4c) and joint probabilities can be computed for each state (3 x 3 cm spatial bin) the same way as presented above (fig. 4d). To assess the confidence of such result, it is also possible to shuffle actual data and compute joint probability maps (fig. 4e). From these shuffled joint probability maps, we can derive the standard deviation (fig. 4f), average joint probability (fig. 4g), and compute a z-scored map (fig. 4h) using equation (8). A significant place field can be considered as the area exceeding 2 standard-deviations. The specificity of such activity can also be assessed by normalizing the joint probability map into a PDF using equation (5), and KLD can inform on the information carried by the neuron under study (fig. 4i). Importantly, it is noteworthy that the standard-deviation of the shuffled distribution is negatively correlated to the relative occupancy (Pearson correlation, $R^2 = 0.47$, $P < 0.0001$; fig. 4j). This suggests that for states with high joint probabilities, the z-score will be higher if the state displays high occupancy, and lower if the state displays low occupancy. We also assessed the effect of temporal filtering on the quality of the reconstructions and found that in our conditions, a 1.6 s filtering window yielded best results (1ANOVA, $F_{(39, 1160)} = 72.31$, $P < 0.0001$, n = 30 independent trials per temporal filter window size; fig. 4k). As for one-dimensional data, gaussian filtering of tuning maps (2D tuning curves) consistently increased the decoding error (Pearson correlation, $R^2 = 0.99$, $P < 0.0001$, n = 30 independent trials per gaussian sigma value; fig. 4l).

3. Reconstruction of neuronal activity from behavior and model refinement

The ultimate goal of decoding neuronal activity is to improve our understanding of the relationship that may exist between neurons and behavior. In other terms, in addition to predicting behavior from neuronal activity, one should be able to predict neuronal activity from behavior.

3.1 Bayesian inference of neuronal activity given behavioral state

This can easily be achieved after determining tuning curves and using an equation analogous to equation (10):

$$(18) \quad P(A|S) = \frac{P(S|A) \times P(A)}{P(S)}$$

This time we want to determine $P(A|S)$, the probability that a neuron is active given that the animal is in state S. Because the date used to generate tuning curves is binary activity (active or inactive), we can reconstruct binary neuronal activity by considering $P(A|S) > 0.5$ as representing activity for a given neuron. 0.5 is of course a subjective threshold that will lead to some errors and depends on the way tuning curves are determined. If there is too much uncertainty in the activity rate of a neuron in a given behavioral state, subsequent reconstructions will suffer from that uncertainty.

3.2 Model refinement by inclusion of both location and orientation variables

In our linear track example, it has been known for a long time that place cell activity on a linear track tend to be unidirectional: some cell will fire in one given location, but only when being traversed in one direction (McNaughton et al., 1983; Markus et al., 1995). If the peak joint probability of a neuron that displays a prominent place field is only 0.5, it could be due to the fact that this cell fires only 50 % of the time, when the animal is running in one direction, but not in the other. We will now demonstrate that it is possible to predict neuronal

A probabilistic framework for decoding behavior from in vivo calcium imaging data

activity from estimated tuning curves, and that refining our behavioral states by including directional information can increase our success rate. To this end, we extracted tuning curves from neurons being active on the linear track using either location only (fig. 5a), or by considering both location and direction (fig. 5b). Note that using the later method, the peak joint probabilities greatly increase. When comparing the peak joint probabilities obtained from the same cells but using either technique, it seems that most cells joint probability increase when including direction information (Pearson correlation, $r = 0.8781$, $R^2 = 0.77$, $P < 0.0001$; fig. 5c) although, this is not the case for a minority of cells, indicating that some cells still display preferential activity for absolute location, regardless of direction. We then reconstructed neuronal activity using the same protocol used in behavior reconstruction (non-overlapping training/testing sets of data representing 50 % of the total experiment time). To compare the quality of reconstructions, we quantified the portion of activity that was correctly reconstructed (fig. 5e). Because inactivity represents most of the neuronal activity on a frame-by-frame basis, focusing only on neuronal activity using this method is more conservative. Importantly, we found that including the mouse direction in addition to its location improved reconstruction drastically: reconstruction success rate went from 0.1494 ± 0.0009 (n = 30 independent trials) when only using the mouse location when establishing tuning curves, to 0.4476 ± 0.0012 (n = 30 independent trials) when using both location and direction (two-tailed unpaired t-test, $t_{58} = 189.8$, $P < 0.0001$).

Discussion

We show here that representing neuronal activity extracted from calcium imaging data by a binary state (active vs inactive) is sufficient to approximate the state of a neuronal assembly. While such binarization was previously proposed as an intermediate step to perform decoding (Ziv et al., 2013), here we generalize this principle and propose several additional metrics to describe the strength of neuronal tuning to behavioral variables. In particular, several methods can be used to binarize calcium activity, but because the rise time of calcium transients contains the vast majority of action potentials, binarizing methods should aim at labeling specifically these epochs as periods of activity. Importantly, optimizing methods and parameters using in vitro conditions cannot necessarily be translated to data acquired in vivo because calcium transients differ fundamentally across conditions, even if most variables are the same (mouse strain/age, viral construct and dilution).

Information on neuronal coding can be extracted using simple methods and minimal data processing. Importantly, three metrics can be used to describe neurons: the joint probability of being active in a given behavioral state, the specificity of such activity (that can be further summarized by a single value using the Kullback-Leibler divergence), and z-scored tuning curve (derived from shuffled surrogates). Using such probabilistic approach allows to perform Bayesian decoding using joint probability values along with unconditional probabilities computed from neuronal activity and animal behavior. We found that z-scored tuning curves are significantly correlated with joint probabilities, while in parallel, behavioral states with low occupancy display more variability in the surrogate data, which indicates that higher z-scores can be achieved when the sampling of behavioral states is higher. On the other hand, the associated PDF and KLD only inform on specificity, not on the neuronal coding reliability. Thus, these two latter metrics alone cannot accurately predict the strength of neuronal coding. Nevertheless, these methods are useful to describe neuronal activity in function of behavior, and models can be derived to predict behavior from neuronal activity and vice versa. The method we describe here can be applied to any type of behavior, and while we present examples of one- and two-dimensional data, the number of dimensions being studied is not a limiting factor.

In our conditions, we found that minimal a priori (uniform distribution of states likelihood) yielded better results. Adding temporal constraints could decrease decoding error but not decoding agreement. Consequently, these filters have to be optimized on a case by case basis depending on the goal of the study. Interestingly, smoothing probability distributions had negative effects in our conditions, most likely due to the asymmetric nature of place fields when unfiltered. Such post-processing methods thus have to be used with caution, and while they can improve the visualization of particular phenomena such as place fields, they can result in spurious interpretations.

Finally, we propose here a new method to predict and reconstruct neuronal activity from previous beliefs of behavior-neuronal activity relationships. This method is useful in refining the behavioral components that can predict neuronal activity. As such, the quality of models that can be drawn from observations largely depends on the very nature and accuracy of these observations. In particular, increasing the amount of information concerning a certain behavior can result in a refinement of the underlying model of neuronal activity. Perfect predictions of neuronal activity on the simple basis of behavior is a difficult endeavor however, because such activity is not only determined by external variables (behavior) but also internal variables (animal state, and pre-synaptic activity that is mostly unobserved). Along those lines, whether one assumes dependence or independence of neuronal activity (may it be on the basis of neuroanatomical or physiological data) could also improve the quality of predictions. Rather than proposing a linear process, the methods presented here should be taken as a continuous process that will help neuroscientist to progressively refine their theories based on increasingly complex sampled behavior, to ultimately improve predictions of both behavior and neuronal activity, as well as producing and communicating metrics that could be scaled across different systems.

A probabilistic framework for decoding behavior from in vivo calcium imaging data

Materials and methods
Surgical procedures

All procedures were approved by the McGill University Animal Care Committee and the Canadian Council on Animal Care. For the linear track and open field data, one adult mouse (~ 2 months) was anesthetized with isoflurane (5% induction, 0.5-2% maintenance) and placed in a stereotaxic frame (Stoelting). The skull was completely cleared of all connective tissue, and a ~500 μm hole was drilled. We then injected the AAV5.CamKII.GCaMP6f.WPRE.SV40 virus (Addgene # 100834; 200 nL at 1 nl.s^{-1}) in hippocampal CA1 using the following coordinates: anteroposterior (AP) -1.86 mm from bregma, mediolateral (ML) 1.5 mm, dorsoventral (DV) 1.5 mm. 2 weeks following the injection, the mouse was anesthetized with isoflurane and the skull was cleared. A ~2 mm diameter hole was perforated in the skull above the injection site. An anchor screw was placed on the posterior plate above the cerebellum. The dura was removed, and the portion of the cortex above the injection site was aspirated using a vacuum pump, until the corpus callosum was visible. These fiber bundles were then gently aspirated without applying pressure on the underlying hippocampus, and a 1.8 mm diameter gradient index (GRIN; Edmund Optics) lens was lower at the following coordinates: AP -1.86 mm from bregma, ML 1.5 mm, DV 1.2 mm. The GRIN lens was permanently attached to the skull using C&B-Metabond (Patterson dental), and Kwik-Sil (World Precision Instruments) silicone adhesive was placed on the GRIN to protect it. 4 weeks later, the silicone cap was removed and CA1 was imaged using a miniscope mounted with an aluminium base plate while the mouse was under light anesthesia (~0.5 % isoflurane) to allow the visualization of cell activity. When a satisfying field of view was found (large neuronal assembly, visible landmarks), the base plate was cemented above the GRIN lens, the miniscope was removed, and a plastic cap was pace on the base plate to protect the GRIN lens.

Behavior and miniscope recordings

After baseplating, the mouse was gently handled for ~5 min per day for 7 days. The mouse was then water-scheduled (2 h access per day), and place on a 1 m long linear track for 15 min. 10% sucrose in water rewards were placed at each end of the linear track, and the mouse had to consume one reward before getting the next one delivered. Miniscope recordings were performed at 30 Hz for 15 min every day, and decoding was performed on the last training day (day 7). The following week, the mouse was allowed to freely explore for 15 min a 45 x 45 cm dark gray open-field that contained visual cues, and miniscope recordings were performed at 30 Hz for the entire duration of the exploration (15 min).

Miniscope and behavior video acquisition

Miniscopes were manufactured using open source plans available on www.miniscope.org and as described previously (Ghosh et al., 2011; Cai et al., 2016; Aharoni and Hoogland, 2019). Imaging data was acquired using a CMOS imaging sensor (Aptina, MT9V032) and multiplexed through a lightweight coaxial cable. Data was acquired using a data acquisition (DAQ) box connected via a USB host controller (Cypress, CYUSB3013). Data was recorded using a custom written acquisition software relying on Open Computer Vision (OpenCV) librairies. Video streams were recorded at ~30 frames per second (30 Hz) and saved as uncompressed .avi files. Animal behavior was recorded using a webcam, and the DAQ

A probabilistic framework for decoding behavior from in vivo calcium imaging data

624 software simultaneously recorded timestamps for both the miniscope and behavior camera in
625 order to perform subsequent alignment.
626

Calcium imaging analysis

Calcium imaging video were analyzed using the MiniscopeAnalysis pipeline
(https://github.com/etterguillaume/MiniscopeAnalysis). In particular, we first applied rigid
motion correction using NoRMCorre (Pnevmatikakis and Giovannucci, 2017). 1000 frame
videos were then concatenated into a large video file after a 2 fold spatial downsampling.
Spatial components as well as calcium traces were then extracted using CNMFe (Zhou et al.,
2018) using the following parameters: gSig = 3 pixels (width of gaussian kernel), gSiz = 15
pixels (approximate neuron diameter), background_model = 'ring', spatial_algorithm = 'hals',
min_corr = 0.8 (minimum pixel correlation threshold), min_PNR = 8 (minimum
peak-to-noise ratio threshold). When applicable, calcium traces were deconvolved with
OASIS (Friedrich et al., 2017), using an autoregressive model with order p = 1 and using the
'constrained' method.

In vitro **patch-clamp electrophysiology**

One adult mouse (~ 2 months) was stereotaxically injected with a GCaMP6f construct
(AAV5.CamKII.GCaMP6f.WPRE.SV40 virus, Addgene # 100834; 0.4 µL at 0.06 µl/min) in
hippocampal CA1. 2 weeks later, it was deeply anesthetized using
ketamine/xylazine/acepromazine mix (100, 16, 3 mg/kg, respectively, intraperitoneal
injection), and intracardially perfused with cold N- methyl-d-glutamine (NMDG) recovery
solution (4° C), oxygenated with carbogen (5% CO2 / 95% O2). The NMDG solution
contained the following (in mM): 93 NMDG, 93 HCl, 2.5 KCl, 1.2 NaH2PO4, 30 NaHCO3,
20 HEPES, 25 glucose, 5 sodium ascorbate, 2 thiourea, 3 sodium pyruvate, pH adjusted to 7.4
with HCl before adding 10 MgSO$_4$ and 0.5 CaCl$_2$. Following NMDG perfusion, brains were
quickly removed and immersed for an additional 1 minute in cold NMDG recovery solution.
Coronal slices (300 µm) were cut using a vibratome (Leica-VT1000S), then collected in a
32°C NMDG recovery solution for 12 minutes. Slices were transferred to room temperature
and oxygenated artificial cerebrospinal fluid (aCSF) containing the following (in mM): 124
NaCl, 24 NaHCO3, 2.5 KCl, 1.2 NaH2PO4, 2 MgSO4, 5 HEPES, 2 CaCl$_2$ and 12.5 glucose
(pH 7.4). Patch pipettes (3–5 MΩ) were filled with internal solution, containing the following
(in mM): 140 K gluconate, 2 MgCl2, 10 HEPES, 0.2 EGTA, 2 NaCl, 2 mM Na2-ATP and 0.3
mM Na2-GTP, pH adjusted to 7.3 with KOH, 290 mOsm. Slices were transferred to a
submerged recording chamber filled with aCSF (2-3 ml /min flow rate, 30 °C), continuously
oxygenated with carbogen. All reagents were purchased from Sigma-Aldrich, unless stated
otherwise. Extracellular cell-attached patch-clamp recordings were used for monitoring
spontaneous cell firing activity from hippocampal pyramidal neurons expressing GcAMP6f
(identified under EGFP-fluorescence). The recording pipette was held at a potential of -70
mV. Imaging of GcAMP6f-expressing pyramidal cells was performed with a 60x Olympus
water immersion objective (LUMPLFLN60X/W, NA 1.0) and acquired at 10 Hz using
Olympus cellSens software. Electrophysiological signals were amplified, using a Multiclamp
700B patch-clamp amplifier (Axon Instruments), sampled at 20 kHz, and filtered at 10 kHz.

Statistics

GraphPad Prism version 6.00 (GraphPad Software, La Jolla, California USA). All data are presented as mean ± standard error of the mean (SEM) and statistic test details are described in the corresponding results. All t-tests are two-tailed. Normality distribution of each group was assessed using Shapiro-Wilk normality test and parametric tests were used only when distributions were found normal (non-parametric tests are described where applicable). 1ANOVA: one-way ANOVA; 2ANOVA: two-way ANOVA; RMANOVA: repeated measure ANOVA. $p < 0.05$ was considered statistically significant. *, $p < 0.05$; **, $p < 0.01$; ***, $p < 0.001$, ****, $p < 0.0001$.

Code and data availability

All the code and data presented here can be downloaded at the following address: https://github.com/etterguillaume/CaImDecoding

Acknowledgments

We thank Ke Cui for support maintaining the colony and perfusing mice, Daniel Aharoni for support using the miniaturized fluorescence microscopes (www.miniscope.org), Bruno Rivard for manufacturing the miniscopes, and Suzanne van der Veldt for helpful comments on the manuscript. This work is supported by Brain Canada, Fonds de la recherche en Santé du Québec (FRSQ), the Canadian Institutes of Health Research (CIHR), the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Alzheimer Society of Canada.

Contributions

GE and SW designed the study. GE performed surgeries, behavioral experiments, and analyzes. FM performed patch clamp *in vitro* electrophysiology. GE and SW wrote the manuscript.

Competing interests

The authors declare no competing interests.

A probabilistic framework for decoding behavior from in vivo calcium imaging data

References

Aharoni, D., and Hoogland, T. M. (2019). Circuit investigations with open-source miniaturized microscopes: past, present and future. *Front. Cell. Neurosci.* 13. doi:10.3389/fncel.2019.00141.

Belle, A. M., Enright, H. A., Sales, A. P., Kulp, K., Osburn, J., Kuhn, E. A., et al. (2018). Evaluation of in vitro neuronal platforms as surrogates for in vivo whole brain systems. *Sci. Rep.* 8. doi:10.1038/s41598-018-28950-5.

Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., and Wilson, M. a (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J. Neurosci.* 18, 7411–25.

Cai, D. J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., et al. (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* 534, 115–118. doi:10.1038/nature17955.

Chen, T., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., et al. (2013). Ultra-sensitive fluorescent proteins for imaging neuronal activity. 499, 295–300.

Davidson, T. J., Kloosterman, F., and Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron* 63, 497–507. doi:10.1016/j.neuron.2009.07.027.

Deneux, T., Kaszas, A., Szalay, G., Katona, G., Lakner, T., Grinvald, A., et al. (2016). Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations in vivo. *Nat. Commun.* 7. doi:10.1038/ncomms12190.

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* 10, 1241–1242. doi:10.1038/nn1961.

Foster, D. J., and Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440, 680–683. doi:10.1038/nature04587.

Friedrich, J., Zhou, P., and Paninski, L. (2017). Fast online deconvolution of calcium imaging data. *PLOS Comput. Biol.* 13, e1005423. doi:10.1371/journal.pcbi.1005423.

Gerwinn, S. (2009). Bayesian population decoding of spiking neurons. *Front. Comput. Neurosci.* 3, 1–14. doi:10.3389/neuro.10.021.2009.

Ghosh, K. K., Burns, L. D., Cocker, E. D., Nimmerjahn, A., Ziv, Y., Gamal, A. E., et al. (2011). Miniaturized integration of a fluorescence microscope. *Nat. Methods* 8, 871–878. doi:10.1038/nmeth.1694.

Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Koay, S. A., Taxidis, J., et al. (2018). CaImAn: An open source tool for scalable Calcium Imaging data Analysis. *bioRxiv*, 339564. doi:10.1101/339564.

Giovannucci, A., Friedrich, J., Kaufman, M., Churchland, A., Chklovskii, D., Paninski, L., et al. (2017). "OnACID: Online Analysis of Calcium Imaging Data in Real Time," in *Advances in Neural Information Processing Systems 30*, eds. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, et al. (Curran Associates, Inc.), 2381–2391. Available at: http://papers.nips.cc/paper/6832-onacid-online-analysis-of-calcium-imaging-data-in-real-time.pdf [Accessed July 9, 2019].

Gonzalez, W. G., Zhang, H., Harutyunyan, A., and Lois, C. (2019). Persistence of neuronal representations through time and damage in the hippocampus. *Science* 365, 821–825. doi:10.1126/science.aav9199.

Grynkiewicz, G., Poenie, M., and Tsien, R. Y. (1985). A new generation of Ca2+ indicators with greatly improved fluorescence properties. *J. Biol. Chem.* 260, 3440–3450.

Hetherington, P. A., and Shapiro, M. L. (1997). Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. *Behav. Neurosci.* 111, 20–34. doi:10.1037/0735-7044.111.1.20.

Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* 160, 106–154.

Katsuki, Y., Uchiyama, H., Watanabe, T., and Maruyama, N. (1956). Electric responses of auditory neurons in cat to sound stimulation. II. *Proc. Jpn. Acad.* 32, 504–506. doi:10.2183/pjab1945.32.504.

Koyama, S., Eden, U. T., Brown, E. N., and Kass, R. E. (2010). Bayesian decoding of neural spike trains. *Ann. Inst. Stat. Math.* 62, 37–59. doi:10.1007/s10463-009-0249-x.

Markus, E. J., Qin, Y. L., Leonard, B., Skaggs, W. E., McNaughton, B. L., and Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* 15, 7079–7094. doi:10.1523/JNEUROSCI.15-11-07079.1995.

McNaughton, B. L., Barnes, C. A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp. Brain Res.* 52, 41–49. doi:10.1007/BF00237147.

Nakai, J., Ohkura, M., and Imoto, K. (2001). A high signal-to-noise Ca(2+) probe composed of a single green fluorescent protein. *Nat. Biotechnol.* 19, 137–141. doi:10.1038/84397.

Ohkura, M., Sasaki, T., Sadakari, J., Gengyo-Ando, K., Kagawa-Nagamura, Y., Kobayashi, C., et al. (2012). Genetically encoded green fluorescent Ca2+ indicators with improved detectability for neuronal Ca2+ signals. *PLOS ONE* 7, e51286. doi:10.1371/journal.pone.0051286.

O'Keefe, J., and Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature* 381, 425–428. doi:10.1038/381425a0.

O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175.

Pachitariu, M., Stringer, C., and Harris, K. D. (2018). Robustness of spike deconvolution for neuronal calcium imaging. *J. Neurosci.* 38, 7976–7985. doi:10.1523/JNEUROSCI.3339-17.2018.

Persechini, A., Lynch, J. A., and Romoser, V. A. (1997). Novel fluorescent indicator proteins for monitoring free intracellular Ca2+. *Cell Calcium* 22, 209–216. doi:10.1016/S0143-4160(97)90014-2.

Pnevmatikakis, E. A., and Giovannucci, A. (2017). NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *J. Neurosci. Methods* 291, 83–94. doi:10.1016/j.jneumeth.2017.07.031.

Pnevmatikakis, E. A., Soudry, D., Gao, Y., Machado, T. A., Merel, J., Pfau, D., et al. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* 89, 299. doi:10.1016/j.neuron.2015.11.037.

Quian Quiroga, R., and Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nat. Rev. Neurosci.* 10, 173–185. doi:10.1038/nrn2578.

Rahmati, V., Kirmse, K., Holthoff, K., and Kiebel, S. J. (2018). Ultra-fast accurate reconstruction of spiking activity from calcium imaging data. *J. Neurophysiol.* 119, 1863–1878. doi:10.1152/jn.00934.2017.

Richmond, B. J., and Optican, L. M. (1990). Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J. Neurophysiol.* 64, 370–380. doi:10.1152/jn.1990.64.2.370.

Richmond, B. J., Optican, L. M., and Spitzer, H. (1990). Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus-response relations. *J. Neurophysiol.* 64, 351–369. doi:10.1152/jn.1990.64.2.351.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.* 27,

379–423. doi:10.1002/j.1538-7305.1948.tb01338.x.

Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., et al. (2017). Tracking the same neurons across multiple days in ca2+ imaging data. *Cell Rep.* 21, 1102–1115. doi:10.1016/j.celrep.2017.10.013.

Skaggs, W. E., McNaughton, B. L., and Gothard, K. M. (1993). An information-theoretic approach to deciphering the hippocampal code. in *Advances in neural information processing systems*, 1030–1037.

Tegtmeier, J., Brosch, M., Janitzky, K., Heinze, H.-J., Ohl, F. W., and Lippert, M. T. (2018). CAVE: an open-source tool for combined analysis of head-mounted calcium imaging and behavior in MATLAB. *Front. Neurosci.* 12. doi:10.3389/fnins.2018.00958.

Tian, L., Hires, S. A., Mao, T., Huber, D., Chiappe, M. E., Chalasani, S. H., et al. (2009). Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat. Methods* 6, 875–881. doi:10.1038/nmeth.1398.

Zhang, K., Ginzburg, I., McNaughton, B. L., and Sejnowski, T. J. Interpreting neuronal population activity by reconstruction: A unified framework with application to hippocampal place cells. *J. Neurophysiol.*, 57.

Zhou, P., Resendez, S. L., Rodriguez-Romaguera, J., Jimenez, J. C., Neufeld, S. Q., Giovannucci, A., et al. (2018). Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *eLife* 7, e28728. doi:10.7554/eLife.28728.

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., et al. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* 16, 264–266. doi:10.1038/nn.3329.

A probabilistic framework for decoding behavior from in vivo calcium imaging data

        Figure legends
        Figure 1. Rationale for extracting spatial coding characteristics of CA1
 pyramidale cells. a, diagram of GRIN lens implanted over CA1 pyramidal cells of the dorsal
 hippocampus. b, calcium imaging was performed as a mouse was running in alternating
 directions on a linear track. c, maximum projection of the corresponding field of view. d,
 corresponding extracted spatial footprints using CNMFe. e, example traces from a subset of
 extracted cells aligned with position on a linear track and locomotor speed. Running epochs
 are indicated with green stripes. f, example raw transient (top) from one cell and
 corresponding filtered, z-scored, first-derivative, and binarized signals. g, rationale used to
 extract unconditional and joint probabilities from binarized data. h, mouse location on the
 linear track with corresponding raw calcium activity and derived binary trace (blue). Only
 runs to the right are considered here. i, (top) mouse trajectory on the linear track (gray) with
 corresponding locations where one cell's binarized activity was detected (blue dots), and
 (bottom) location of binarized activity on the linear track for each run (n = 16 runs). j, joint
 probability of cell #4 to be active in each given linear track location. k, derived probability
 density function (ocher) compared to uniformity (dotted line), and (top) corresponding
 Kullback-Leibler divergence. l, example cases of poor variable coding (case 1), superior
 variable coding (case 2), poor variable coding with sparse information (case 3), and superior
 variable coding with sparse information (case 4). m, actual (a) calcium trace, corresponding
 permutations (s_n), and corresponding location (green). n, joint probability computed from one
 example permuted trace. o, actual joint probability distribution (blue) and corresponding
 average shuffled distribution from n = 1000 surrogates (the thickness of the line represents the
 SEM). p, z-scored distribution of cell activity on the linear track, computed from actual and
 shuffled data. q, histogram distribution of peak joint probability corresponding to each
 permutation (magenta) compared to actual peak probability (blue).

        Figure 2. Bayesian decoding of behavior from calcium imaging recording. a,
 spatial tuning curves for each individual CA1 neuron (data sorted from location of peak joint
 probability), and corresponding probability of being active (right-hand side), and probability
 of being in a given state (= location; bottom). b, raster plot of binarized cell activity and
 corresponding position on the linear track (bottom). c, tuning curves of cells corresponding to
 their state at frame 11392 (in b) and subsequent posterior probability of being in a given
 location on the linear track (bottom). Location was estimated using maximum a posteriori
 (MAP). d, posterior probabilities for each frame estimated from ongoing binarized calcium
 activity, and corresponding actual (green) and decoded (pink) location estimated with MAP. e,
 confusion matrix of actual vs decoded position. f, method used to compute euclidean
 decoding error (top) and decoding agreement (bottom). g, paradigm used to train and test the
 decoder on different epochs of the dataset. h, effect of prior and bias (cell probability of being
 active) on decoding agreement. i, same for decoding error. Color codes in a,c,d,e: dark blue =
 low probability, yellow = high probability.

        Figure 3. Decoding parameter estimation. a, example posterior probabilities when
 using a 0.5 s temporal filtering window (top), and corresponding decoded location estimated
 from MAP (bottom). b & c, effect of temporal filtering window size on decoding error and
 agreement, respectively. d & e, effect of the number of cell used on decoding error and
 agreement, respectively. f & g, effect of training set portion on decoding error and agreement,
 respectively. h, effect of random noise on spatial tuning curves. i & j, corresponding decoding
 agreement and error, respectively. k, effect of gaussian smoothing on spatial tuning curves. l
 & m, corresponding decoding agreement and error, respectively. Color codes in a & h,: dark

blue = low probability, yellow = high probability.

Figure 4. Decoding two-dimensional behaviors. a, *x,y* coordinates of mouse location in an open field (bottom) and corresponding raw calcium trace of one example cell and binarized activity (top). b, top view of mouse trajectory (beige trace) with overlaid location corresponding to neuronal activity (green, early activity; magenta, late activity). c, relative occupancy in the open field. d, joint probability tuning map of one neuron. e, example tuning maps computed from shuffled calciu traces. f, standard-deviation of the shuffled distribution. g, average joint probability tuning map of the shuffled distribution. h, corresponding z-scored tuning map, computed from the actual tuning map, and standard-deviation as well as mean from the shuffled distribution. d, probability density function derived from the joint probability tuning map. The number on top indicates the corresponding spatial information. j, scatter plot comparing the standard-deviation of the shuffled distribution, and the mouse open field occupancy. k, effect of temporal filtering on decoding error in the open field. The red arrow indicates the the temporal filtering window size yielding the lower decoding error. l, effect of gaussian smoothing of tuning maps on decoding error in the open field.

Figure 5. Reconstructing neuronal activity and refining tuning curves
a, tuning curves of every neuron sorted by peak activity in the linear track. b, same, but after discriminating left and right trajectories. c, relationship between peak joint probabilities of tuning curves computed either method (location only versus location and direction). d, actual location of the mouse in the linear track (top), and corresponding actual and reconstructed neuronal activity using location only (middle), as well as actual and reconstructed neuronal activity using both location and direction (bottom). e, method used to assess reconstruction success rate. f, mean reconstruction success using either location data (purple) or location and direction concurrently (green). Color codes in a & b: dark blue = low probability, yellow = high probability.

a1

in vivo

GRIN lens

CA1

b

100 cm

c

cell 4

d

cell 4

e

Cell ID

Position (cm)

100
50
0

Speed (cm.s⁻¹)

30
20
10
0

10 s

f

Raw

10 % ΔF/F

Filtered

10 % ΔF/F

Z

SD

10
6
2
-2

f'(x)

ΔF/Δt

1.5
0
-1.5

Binarized

1 s

g

$P(a)$ $P(s)$

$P(a \cap s)$ **Joint probability**: 'Likelihood of being active in state *s*'

$P(a^c \cap s)$ **Unconditional probability**: 'Likelihood of being inactive in state *s*'

$P(a \cap s^c)$ **Unconditional probability**: 'Likelihood of being active outside of state *s*'

h

Runs to the right

Position (cm)

100
50
0

Bin.

Raw

100 s

i

Lap number

Binarized

0   25   50   75   100
Position on track (cm)

j

Joint probability

1.0
0.8
0.6
0.4
0.2
0.0

0   25   50   75   100
Position on track (cm)

k

2.2105 bits

Probability

0.3
0.2
0.1

0   25   50   75   100
Position on track (cm)

l

CASE 1

CASE 2

CASE 3

CASE 4

m

a

s₁

s₂

s₃

s₄

s₅

...

Position (cm)

100
50
0

n

Joint probability

0.5
0.4
0.3
0.2
0.1

0   25   50   75   100

o

Actual
Shuffled (*n* = 1000)

Joint probability

1.0
0.8
0.6
0.4
0.2
0.0
-0.2

0   25   50   75   100
Position on track (cm)

p

Z score

60
40
20
0
-20

2 sd

0   25   50   75   100
Position on track (cm)

q

shuffled   actual

Count

0   0.2   0.4   0.6   0.8   1
Peak joint probability

a

Direction

Cell ID
417
1

Joint probability
0.5
0.25
0

Position (cm)
0   50   100

b

Direction

Cell ID
417
1

Joint probability
1
0.8
0.6
0.4
0.2
0

Position (cm)
0   50   100   50   0

c

R² = 0.77
****

Peak joint probability (location and direction)
1.0
0.5
0.0

Peak joint probability (location only)
0.0   0.5   1.0

d

Left trajectory   Right trajectory

Actual location (cm)
100
50
0

Cell index sorted by location only

Actual neuronal activity

Estimated activity using location only

Cell index sorted by location & direction

Actual neuronal activity

Estimated activity using location only

30 s

e

Frame:	1	2	3	4	5	
	1	0	0	0	0	
	1	1	1	0	0	
	0	1	1	0	0	Actual
	0	0	1	1	0	
	0	0	0	1	0	
	0	0	0	0	1	

Cell ID 6 ... 1

	1	0	0	0	0	
	1	0	0	0	0	
	0	1	1	0	0	Reconstructed
	0	0	1	1	0	
	0	0	0	0	0	
	0	0	0	0	1	

Cell ID 6 ... 1

Success rate: 1   0.5   0.67   1   1   Average --> 0.83

f

Mean reconstruction success
0.5
0.4
0.3
0.2
0.1
0.0

****

Location only   Location + direction