Scoring Sequences of Hippocampal Activity using Hidden Markov Models

Etienne Ackermann[1*] and Caleb Kemere[1,2]

Abstract— We propose a novel sequence score to determine to what extent neural activity is consistent with trajectories through latent ensemble states—virtual place fields—in an associated environment. In particular, we show how hidden Markov models (HMMs) can be used to model and analyze sequences of neural activity, and how the resulting joint probability of an observation sequence and an underlying sequence of states naturally lead to the development of a two component sequence score in which the sequential and contextual information are decoupled. We also show how this score can discriminate between true and shuffled sequences of hippocampal neural activity.

I. INTRODUCTION

The activity of ensembles of neurons within the hippocampus is thought to enable the memory formation, storage, and recall and even potentially decision making. Temporally-ordered, *sequential* activity of these neurons is thought to enable associations across time and episodic memories that span longer periods. As hippocampal neural activity is communicated to nearly the entire neocortex, variability in these sequences may potentially have a significant impact. Despite numerous phenomenological expositions, our understanding of the variability of hippocampal sequences is lacking, due in large part to a lack of quantitative metrics of sequential variability. Two forms of variability exist, namely variability in the co-firing of subsets of cells (which we refer to as contextual variability) as well as variability in the sequential ordering of neural activity. Here we use hidden Markov models (HMMs) to characterize sequences of neural activity, and we use the HMM framework to derive a two component sequence score that allows us to determine to what extent a candidate sequence is sequentially consistent. In particular, this score quantifies both contextual as well as sequential (trajectory) information of putative sequences.

In rodents, hippocampal "place cells" are known to encode an animal's location in its environment as it explores [1]. Hence, populations of these neurons fire in sequences corresponding to the spatiotemporal trajectories the animals traverse. HMMs are well suited to model this sort of sequential activity—even more so for behavioral correlates of memory that are non-spatial (e.g., odors). Of particular interest are hippocampal replay events in which neurons recapitulate their spatially ordered sequences during periods of quiescence or sleep. However, in this paper we limit our attention to sequences of neural activity for which behavioral correlates (the animal's position as a function of time) are known, so that we can evaluate the efficacy of our novel sequence score with ground truth data.

A recent study used a HMM to study population activity in the hippocampus during spatial navigation [2], where they showed that nonparametric Bayesian extensions to the HMM were useful to dynamically infer the number of hidden states in the model. However, they did not explicitly describe ways to score candidate sequences, and did not discuss how this approach can apply to replay.

In studies of replay, the "quality" of a candidate sequence has been evaluated by comparison to a firing rate model constructed using average neural responses over multiple behaviorally-correlated ensemble firing instances. Specific metrics have included a combinatorial approach where it was determined how many cells in a sequence could fire out of place but still be considered a "valid" sequence (see e.g. [3]) and the quality-of-fit of a decoded position [4], [5]. These approaches critically rely on the availability of non-neural behavioral data to properly generate models. What if we don't have access to this data?

Here we describe how we can use the HMM framework to use only ensemble activity to effectively learn consistent underlying hidden states which we demonstrate correspond with behavioral correlates when they are available. We also demonstrate how our novel sequence score can be used to determine to what extent the neural activity is consistent with trajectories through the latent ensemble states or virtual place fields in an associated environment.

II. HIDDEN MARKOV MODELS OF NEURAL ACTIVITY

HMMs are statistical models where the systems being modeled are Markov chains (or more generally Markov processes) with unobserved or hidden states, and they have been widely used for sequential pattern recognition and processing in fields ranging from speech recognition to bioinformatics (see [6] for an excellent tutorial introduction).

In this section we loosely follow the approach and notation presented by [2] to demonstrate how HMMs can be used to model sequential neural activity, and we also present our two-component sequence score, which we use to characterize sequences in terms of their sequential consistency, as well as their contextual activity.

A. Model specification

Let $\mathbf{y}_t$ denote the *observation* at time t, where $\mathbf{y}_t \in \mathbb{Z}^C$ is a vector of spike counts for C hippocampal place cells. It is assumed that the observations are sampled at discrete, equally-spaced time intervals, so that t can be an integer-valued index, with some associated Δt.

We further assume that the hidden state space is *discrete*. That is, $S_t \in \{1, \ldots, m\}$ can take on one of m possible states, with each possible state loosely corresponding to a particular location in the environment (a "virtual place cell").

To define a probability distribution over sequences of observations, we then need to specify a probability distribution over the initial state $P(S_1)$, with $\boldsymbol{\pi}_i \equiv \Pr(S_1 = i)$, the $m \times m$ state transition probability matrix, $\mathbf{P}$, with $\mathbf{P}_{ij}$ defining $P(S_t = j | S_{t-1} = i)$, and the output or emissions model defining $P(\mathbf{y}_t | S_t)$.

We will further assume that our model is time-invariant: the state transition probability matrix and the output model does not change over time. The state sequences are also assumed Markovian, so that $P(S_t | S_{t-1}, \ldots S_1) = P(S_t | S_{t-1})$—that is, the state sequence forms a first order Markov chain—and we also assume that observations are conditionally independent given the underlying states.

For the output model we assume Poisson firing statistics for each spike train, so that the emission probability for the ith state is modeled by a spatially varying (state-dependent) multivariate Poisson process:

$$
\begin{aligned}
P(\mathbf{y}_t | S_t = i; \theta) &= \prod_{c=1}^{C} P(y_{c,t} | S_t = i; \theta) \\
&= \prod_{c=1}^{C} \prod_{j=1}^{m} P(y_{c,t} | S_t = j; \theta)^{S_{t,i}} \\
&= \prod_{c=1}^{C} \prod_{j=1}^{m} \left(\frac{\exp(-\lambda_{jc}) \lambda_{jc}^{y_{c,t}}}{y_{c,t}!} \right)^{S_{t,i}}
\end{aligned}
$$

where $\theta = \{\boldsymbol{\pi}, \mathbf{P}, \boldsymbol{\Lambda}\}$ are the model parameters, $\boldsymbol{\Lambda} \in \mathbb{R}^{m \times C}$ are the tuning curve parameters (a spike firing rate λ for every possible state $j \in \{1 \ldots m\}$ for each place cell $c \in \{1, \ldots C\}$), and $S_{t,i} = 1$ iff $S_t = i$, and 0 otherwise.

Given a training set $\mathcal{D} = \{\mathbf{y}_{1:T_1}^{(1)}, \ldots, \mathbf{y}_{1:T_N}^{(N)}\}$, containing N sequences of observations, and since the training sequences are assumed to have been drawn independently, the complete data likelihood takes the form

$$
P(\mathcal{D}, \mathbf{S} | \theta) = \prod_{n=1}^{N} P\left(\mathbf{y}_{1:T_n}^{(n)} | \theta, S_{1:T_n}^{(n)} \right) P\left(S_{1:T_n}^{(n)} \right). \quad (1)
$$

The model parameters $\theta = \{\boldsymbol{\pi}, \mathbf{P}, \boldsymbol{\Lambda}\}$ can then be estimated using standard methods such as expectation maximization, variational Bayes, or Monte Carlo methods (the solutions presented in [6] use expectation maximization, for example).

Making the dependence on θ implicit, the joint distribution of a sequence of states and the resulting observations can then be factored as follows:

$$
P(S_{1:T}, \mathbf{y}_{1:T}) = P(S_1) P(\mathbf{y}_1 | S_1) \prod_{t=2}^{T} P(S_t | S_{t-1}) P(\mathbf{y}_t | S_t), \quad (2)
$$

which is equivalent to

$$
P(S_{1:T}, \mathbf{y}_{1:T}) = \underbrace{\left[P(S_1) \prod_{t=2}^{T} P(S_t | S_{t-1}) \right]}_{\text{sequential (seq)}} \underbrace{\left[\prod_{t=1}^{T} P(\mathbf{y}_t | S_t) \right]}_{\text{contextual (ctx)}}. \quad (3)
$$

We base our novel sequence score on (3).

B. Sequence scoring

Scoring [replay] sequences allow us to determine if they are consistent with sequences observed during active behavior. More specifically, we are interested in evaluating (i) how well a particular sequence of observations fit in an underlying context, and (ii) how well that sequence captures the sequential nature of an associated (experienced or imaginary) trajectory through the environment, and (iii) we want to be able to compare different sequences (possibly of different lengths) to each other using these criteria.

1) Existing sequence and replay scores: Probably the most widely used scores for replay sequences include the "replay score" by [4], and the weighted correlation, defined as Pearson's product-moment correlation weighted by the decoded posterior probability [7]. A decoding-free approach was also presented by [8]. More recently, [5] summarized six previously used metrics, namely the (i) weighted correlation, (ii) maximum jump distance, (iii) slope of the best linear fitted trajectory along the entire track, (iv) aforementioned replay score, (v) sharpness of decoded probability, and (vi) position occupancy. They found that the maximum jump distance along with correlation were the strongest differentiating factors. The jump distance captures the continuity of movement, while the correlation captures the population neural activity (including both sequential and contextual information).

However, the maximum jump distance is a very coarse metric, and the weighted correlation mixes a lot of sequential information into the contextual component. The abovementioned scores also generally assume linear, constant velocity motion, and some require several empirical filtering steps, so that a more elegant and more general solution is needed. Our novel two component score is an attempt to improve on these existing scores.

2) Novel sequence score: We consider the *contextual* and *sequential* factors of the joint distribution of a (known) sequence of states and its associated sequence of observations, as given in (3), and we modify them as follows:

$$
q_{\text{ctx}} = \frac{1}{T} \sum_{t=1}^{T} \max_{S} \left\{ \log P(\mathbf{y}_t | S) \right\} \quad \text{and} \quad (4)
$$

$$
q_{\text{seq}} = \log |\mathcal{S}| + \frac{1}{T} \left(\log P(S_1) + \sum_{t=2}^{T} \log P(S_t | S_{t-1}) \right) \quad (5)
$$

where $|\mathcal{S}|$ denotes the cardinality of the set of states visited: $\mathcal{S} = \{S_t\}_{t=1}^{T}$. That is, $|\mathcal{S}|$ is simply the number of unique states visited in the sequence, but since the true state sequence might be unknown, we will use the Viterbi sequence to determine the most likely sequence of states instead. Here we also divide the log probabilities by the sequence length, T, to normalize the scores for sequences of different lengths.

Note that in evaluating the contextual component (which tells us how well the neural activity matches the tuning curves associated with the virtual place fields) of a sequence, we evaluate the probabilities using the symbol-by-symbol memoryless maximum a posteriori (MAP) state estimates at each time point: we do this so that the contextual score is free from any model-imposed trajectory structure, and only depends on the unordered, spatially independent population activity in any given context.

III. RESULTS

To demonstrate the efficacy of our novel sequence score in discriminating true sequences from random data, we used data from the online repository CRCNS.org. In particular, we used the hippocampal data set hc-3, with animal gor1 which ran on two linear tracks per day[1] (see [9] for details). Here, results are shown for the second session of day one, but similar results are obtained for every session[2].

A. Model selection and validation

Completely specifying the model requires us to fix (i) the number of states, (ii) the time bin size, and (iii) the velocity threshold for classifying the animal's behavioral state as either active or inactive. To further facilitate model selection, the data was split into train, validation, and test sets. Only the train data was used to train the HMM while the validation set was used to determine a suitable number of states, and the test set was used for reporting final decoding accuracy.

A time bin of 125 ms was chosen because (i) it is fast enough to capture the behavioral dynamics of the animal, and (ii) it captures a full theta cycle (≈ 8 Hz in rodents). Similarly, a running velocity of 8 units/second was chosen so that the running behavior corresponded primarily to traversals on the track[3].

As shown in Fig. 1 there is very little improvement in the data log likelihood (of the validation set) after about $m = 25$ states, which we then used for all subsequent analyses.

To verify that the learned model states correspond to spatial behavioral correlates, we determined their spatial selectivity by associating decoded states with the animal's true position (which was never shown to the model!). In this way we can visualize where each state is active as shown in Fig. 2. Indeed, the learned model states show remarkable spatial selectivity, reminiscent of real place cell fields.

[1] The linear tracks are actually the same track, just in a different position/orientation.

[2] The second session was chosen for this day, since the position data for the first session are incomplete.

[3] It is important to note that even though we only present results for these parameter values, we have repeated the analysis on a wide range of parameters, all of which yielded comparable results. In other words, this analysis is relatively insensitive to the actual parameter choices.
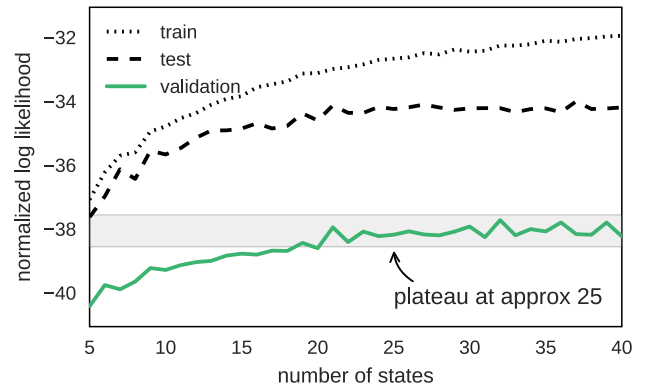

Fig. 1. Model selection; model log likelihood reaches a plateau (on both the validation and test sets) at around $m = 25$ states.
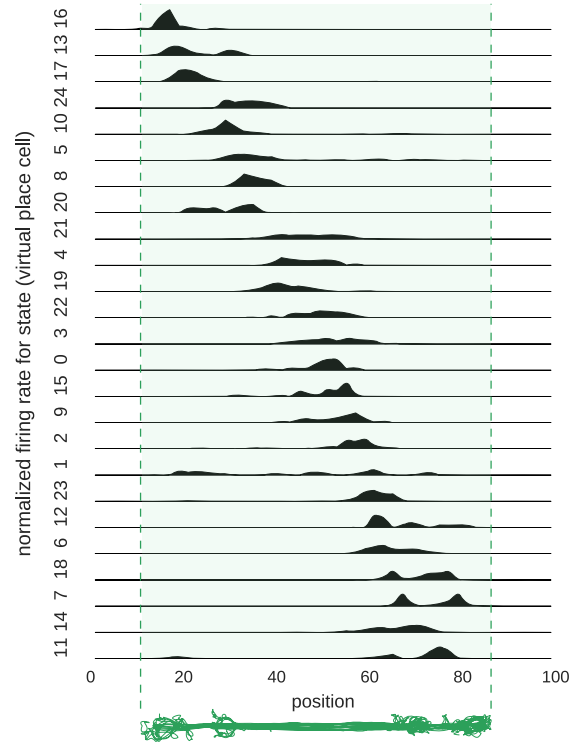

Fig. 2. Place fields corresponding to latent ensemble states, ordered by peak location on the track. Only x-dimension considered. Bottom shows actual (x, y) coordinates of animal.

As a final validation step, we used the HMM and the place fields of Fig. 2 to decode (symbol-by-symbol MAP) trajectories from the test set (which were never seen by the model). An example of a decoded sequence is shown in Fig. 3, and the entire test set had an RMSE of 5.60 units.

B. Sequence score results

Using this HMM, we computed the contextual and sequential scores for the remaining subsequences during active behavior. The scores are shown in Fig. 4 where the sequences are clearly separated from the trajectory-shuffled data. The trajectory shuffle only shuffles the binned spike count data so
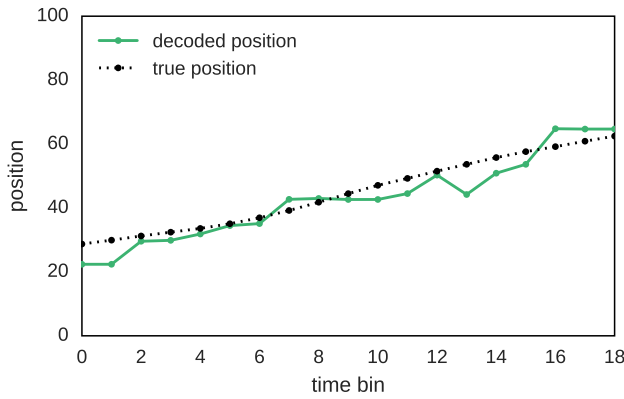
Fig. 3. Example of a decoded sequence from the test set; RMSE = 4.31.

that the time bins are out of sequence. This type of shuffle should make it more difficult to discriminate between real and shuffled sequences than other commonly used shuffling strategies including the unit identity shuffle (see e.g. [4]), since unlike most shuffling strategies, the trajectory shuffle leaves the correlations between place cells unaltered.
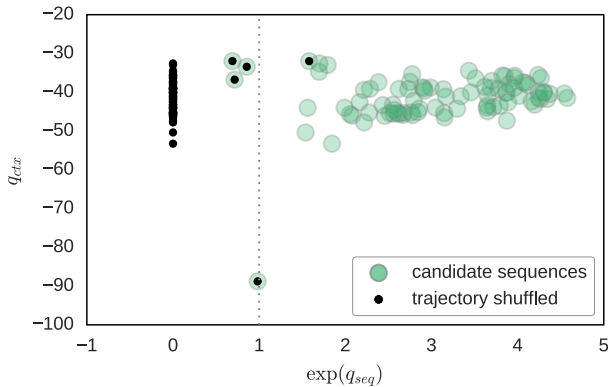

Fig. 4. Sequence scores for previously unseen sequences and trajectory-shuffled versions of those sequences. Note that the contextual component is invariant to trajectory-shuffling, and that the sequential component clearly discriminates between true and shuffled sequences.

Aside from the few samples close to the margin (which are there mostly due to being short sequences), the sequential component alone effectively discriminates true sequences from shuffled data.

IV. DISCUSSION

The results are shown for sequences during active behavior to facilitate a comparison to ground truth data. However, the idea is to use this sequence score for replay analysis, where the imagined position is unobservable. The ability to train these HMMs in the absence of any behavioral correlates really sets them apart from the conventional template matching and Bayesian decoding approaches for which having behavioral data is a prerequisite.

It is encouraging that the factorization of the joint distribution of (3) naturally leads to the sequential and con-

textual factors, and that [5] found the maximum jump distance (= sequential information) and weighted correlation (= contextual information) the most useful in discriminating between real sequences and random data. However, our score is a refinement of these ideas where the sequential component summarizes the sequential nature of the entire sequence, and not just the single worst time bin, as is the case for the maximum jump distance, and our contextual component mixes less of the sequential information than the weighted correlation.

Indeed, our two component score (i) is easy to interpret (without the need to first compare them to scores obtained from shuffle distributions), (ii) accommodates sequences of different lengths, (iii) is appropriate for complex environments without modification (including nonlinear movement, both in terms of speed and trajectory), and (iv) directly penalizes sequences of neural activity for which an animal is almost stationary.

V. CONCLUSION

HMMs are ideally suited for sequential analysis, and even more so in the case of replay, where the underlying states are unobservable. We have shown how HMMs can be used to model and analyze neural activity, and we have shown how it leads to our novel two component sequence score, which is interpretable even in the absence of scores of shuffled data, and which can easily discriminate between real sequences and shuffled data. We will use this score to perform replay detection and analysis, and we believe that our score is an improvement on (refinement of) the two metrics used by [5].

ACKNOWLEDGMENT

REFERENCES

[1] J. O'Keefe and J. Dostrovsky, "The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat," *Brain research*, vol. 34, no. 1, pp. 171–175, 1971.
[2] S. W. Linderman, M. J. Johnson, M. A. Wilson, and Z. Chen, "A bayesian nonparametric approach for uncovering rat hippocampal population codes during spatial navigation," *Journal of neuroscience methods*, vol. 263, pp. 36–47, 2016.
[3] D. Ji and M. A. Wilson, "Coordinated memory replay in the visual cortex and hippocampus during sleep." *Nature neuroscience*, vol. 10, no. 1, pp. 100–7, jan 2007.
[4] T. J. Davidson, F. Kloosterman, and M. A. Wilson, "Hippocampal replay of extended experience," *Neuron*, vol. 63, no. 4, pp. 497–507, 2009.
[5] D. Silva, T. Feng, and D. J. Foster, "Trajectory events across hippocampal place cells require previous experience," *Nature Neuroscience*, vol. 18, no. 12, pp. 1–10, 2015.
[6] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, 1989.
[7] X. Wu and D. J. Foster, "Hippocampal replay captures the unique topological structure of a novel environment," *The Journal of Neuroscience*, vol. 34, no. 19, pp. 6459–6469, 2014.
[8] A. S. Gupta, M. A. van der Meer, D. S. Touretzky, and A. D. Redish, "Hippocampal replay is not a simple function of experience," *Neuron*, vol. 65, no. 5, pp. 695–705, 2010.
[9] K. Mizuseki, K. Diba, E. Pastalkova, J. Teeters, A. Sirota, and G. Buzsaki, "Neurosharing: large-scale data sets (spike, LFP) recorded from the hippocampal-entorhinal system in behaving rats," *F1000Res*, vol. 3, no. 0, p. 98, 2014.