

Functional connectivity models for decoding of spatial representations from hippocampal CA1 recordings

L. Posani¹, S. Cocco¹, K. Jezek², and R. Monasson³

¹Laboratoire de Physique Statistique, Ecole Normale Supérieure and CNRS UMR 8550, PSL Research, Paris Sorbonne UPMC, 24 rue Lhomond, 75005 Paris, France

²Laboratory of Experimental Neurophysiology, Biomedical Center, Charles University in Prague, Faculty of Medicine in Pilsen, alej Svobody 1655/76, 32300 Pilsen, Czech Republic

³Laboratoire de Physique Théorique, Ecole Normale Supérieure and CNRS UMR 8549, PSL Research, Paris Sorbonne UPMC, 24 rue Lhomond, 75005 Paris, France

Abstract

Hippocampus can store spatial representations, or maps, which are recalled each time a subject is placed in the corresponding environment. We consider the problem of decoding the recalled maps as a function of time from multi-cellular recordings in CA1. Maps corresponding to different environments in CA1 mostly differ by changes in firing rates rather than firing fields, and are harder to identify than in CA3, in which maps are essentially orthogonal. We introduce a functional-connectivity-based decoder, which accounts for the pairwise correlations between the spiking activities of neurons in each map and does not require any positional information, i.e. any knowledge about place fields. We first show, on recordings of hippocampal activity in constant environmental conditions, that our decoder outperforms existing decoding methods in CA1. Our decoder is then applied to data from teleportation experiments, in which instantaneous switches between environmental conditions trigger the recall of the corresponding maps. We test the sensitivity of our approach on the transition dynamics between the respective memory states (maps). We find that the rate of spontaneous state shifts (flickering) after a teleportation event is increased not only within the first few seconds as already reported, but the network also shows a higher instability level on much longer (> 1 min) intervals.

1 Introduction

Over the recent decades, multi-cell recording techniques have provided insights into the nature of brain representations and their internal dynamics. While many works have focused on the input-output transfer functions in primary sensory systems (visual, olfactory, etc.), understanding functions corresponding to complex representations in higher cortical circuits is very hard as they are often based on mixed selectivities (Rigotti et al., 2013). In relatively rare cases, such as in the entorhino-hippocampal system, a highly processed neural activity can be reliably correlated with behavior. The so-called ‘place cells’ in the CA1 and CA3 of hippocampus exhibit sharp spatially tuned and environment specific activity (O’Keefe and Dostrovsky, 1971), see Fig. 1. Collective activity of the place-cell population coding for the environment defines its neural representation, or *map*. Simultaneous recording of multiple place-cell activity thus allows one to identify a general memory state of the network (specific map), as well as to decode the accurate position of the rodent in the corresponding environment (Zhang et al., 1998).

Recently, Jezek et al. (2011) have studied the dynamics of transient change between the spatial maps encoding two different environments in CA3 at high temporal resolution (ca 120 ms time windows). The two environments differed by light cues that could be switched instantaneously ('teleportation procedure'), while the animal hippocampal neural activity was recorded to monitor the course of activation of the proper spatial map. An unstable state generally emerged for some seconds after the light switch, as both maps started to flicker back and forth. This phenomenon, called flickering, was identified through measure of the similarity between the place-cell population activity and its averaged patterns across both environments, recorded earlier in respective reference sessions. Typically, a given 120 ms time window activity of the test data strongly correlated with the average reference activity in one map, and had essentially no correlation with the reference activity in the other map.

Success of such comparison-based decoding methods reflects the strong orthogonality of spatial maps in CA3: across two environments, activity of place cells broadly differ in their mean frequencies and receptive field locations, see Fig. 1(b). Hence, simple map decoders, essentially assuming that cells fire independently of each other, are sufficient to reliably identify the representation expressed by the animal. In contrast, remapping between environments (especially of similar geometry) is less orthogonal in CA1 as it shows higher number of cells firing at corresponding places across rooms, see Fig. 1(a). The population activity vector often correlates well with both concurrent reference templates, which hinders the use of comparison-based methods for map decoding.

Here we address the challenging goal of map decoding in CA1 by introducing a probabilistic graphical model for the neural activity configurations in each map. Graphical models, in which a functional connectivity network accounts for the pairwise correlation structure between neuronal firing events in the recorded population (MacKay, 2003), have been applied to various areas so far (Stevenson et al., 2008), e.g. to estimate the information conveyed by (Tkacik et al., 2010) or the activity of (Pillow et al., 2008; Schneidman et al., 2006) retinal ganglion cells in the presence of visual stimuli, to detect learning-related changes in functional connectivity in the prefrontal cortex (Tavoni et al., 2015, 2016).

We apply our graphical-model decoder to new recordings of the hippocampal activity in CA1, performed within the teleportation setup of Jezek et al. (2011). Our decoder shows very good performances in terms of precision and statistical properties in CA1. It allows us, in particular, to identify transitions between spatial representations in CA1 in a statistically robust way. Remarkably, we find that the frequency of these flickering events is increased even minutes after a teleportation switch.

It is important to stress that, in contradistinction with previously used map decoders, ours does not use any position information. It can therefore be applied to decode and study the dynamics of general brain states with unknown input correlates, the only working hypothesis being that we dispose of reference sessions to build statistical models of the corresponding internal states.

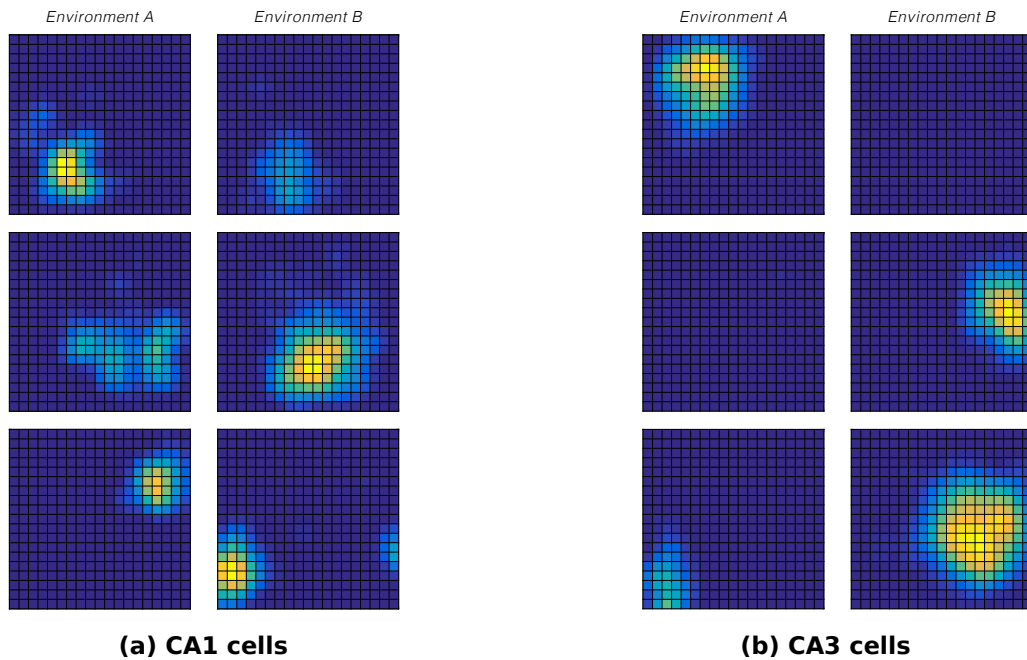


Figure 1: **Hippocampal representations are less orthogonal in CA1 (a) than in CA3 (b).** Each panel shows six firing fields from CA1 (a) and CA3 (b) corresponding to three place cells (rows) in the recorded neuronal population, computed from 10 min recordings of the activity during free exploration of environments A and B (same shapes). Whereas CA3 coding is highly sparse and representations are largely orthogonal, CA1 population shows higher amount of cells active in corresponding locations across the two rooms, with peak rates (color scale) changing from one environment to the other. The non-orthogonality of environment representations in CA1 makes identification of the represented map from neural activity difficult compared to the situation in CA3. CA3 data were taken from Jezek et al. (2011).

2 Materials and Methods

2.1 Experimental methods

Electrode preparation and surgery. Single unit neuronal activity was recorded in adult Long Evans male rats in hippocampal subfields CA1. Rats were implanted with a “hyperdrive” allowing for an independent positioning of 16 tetrodes organized into an ellipsoid bundle. Tetrodes were twisted from 17 μm insulated platinum - iridium wire (90% and 10%, respectively, California Fine Wire Company). Impedance of electrode tips was adjusted by platinum plating to 120 – 250 kOhm (at 1 kHz). Anesthesia was introduced by placing the rat into a plexiglas chamber with seal top filled with isoflurane vapour. Then the animal was shaved and placed into the stereotaxic frame and continued the isoflurane delivery with a face mask. Breathing, heart action and reflexes were monitored continuously. Hyperdrive was then implanted above the right dorsal hippocampus at coordinates AP 3.8 mm and ML 3.2 mm relative to bregma. Stainless steel screws and dental acrylic were used to stabilize the implant on the skull. Two of the screws served as the hyperdrive ground.

Tetrode position. The tetrodes were slowly approached towards CA1 within 2-3 weeks after the surgery while the rat was resting in a comfortable pot on a pedestal. To maintain stable recordings, electrodes were not moved at all before and during the experiment on a given day. The recording reference electrode was positioned in corpus callosum. Additional reference for EEG was placed in stratum lacunosum moleculare.

Recording procedures. Neural activity was recorded while the rat was behaving in an apparatus described by Jezek et al. (2011). Signal was recorded differentially against the reference tetrode. Hyperdrive was connected to a multichannel, impedance matching, unity gain headstage and its output conducted through a 82-channel commutator to a Neuralynx digital 64 channel data acquisition system. Signal was band-pass filtered at 600 Hz–6 kHz. Unit waveforms above individually set thresholds (45-70 μV) were time-stamped and digitized at 32 kHz. Position of the light emitting diodes on the headstage was tracked at 50 Hz to assess the animal’s position. For the purpose of this study only data from intervals when the rat’s movement speed exceeded 5 cm/sec were used. Broadband EEG from each tetrode was recorded continuously at 2000 Hz.

Spike sorting and cell classification. Spikes were sorted manually using 3D graphical cluster-cutting software (SpikeSort, Neuralynx) The feature space consisted of three-dimensional projections of multidimensional waveform amplitudes and energies. Autocorrelation and crosscorrelation functions were used as additional separation tools. Putative pyramidal cells were distinguished from putative interneurons by average rate, spike width and occasional complex spikes.

Histology. After the experiment was finished, the rat was overdosed with Equithesin and was perfused intracardially with saline followed by 4 % formaldehyde. Brain coronal sections (30 μm) were stained with cresyl violet. Traces of all 14 tetrode locations were

identified. Each tip location was considered as the place in the section before the tissue damage became negligible. Only recordings from tetrodes with their tips in CA1 were used in this study.

Behavioral procedure. Animals were first pre-trained according to the procedure described in Jezek et al. (2011). Briefly speaking, the apparatus consisted of two identical black plastic boxes (60×60 cm, 50 cm in height). The two environments differed only by sets of light cues, one placed on the upper rim of the box, the second was positioned under the semi-transparent floor with an additional cue on one wall, respectively. There were no other visual cues present as the experiment was otherwise carried in darkness provided by surrounding light-proof curtains. The training consisted of four phases. Initially, the two boxes were connected with an alley so the rat could freely explore both of them within three 20 min. sessions for 3 days. In the second phase, after the first 20 min. session, the alley was removed and the animal was placed into box *A* or *B*, respectively, in a quasi-random manner so that it received two 10 min. sessions in each of them, respectively. The next day the rat received two 10 min. sessions in each environment as the day before. Then we removed the double maze and replaced it with a single box equipped with both sets of lights that was presented at the original locations with just one cue set switched on at the given session. The rat was given another two 10 min. sessions in each environment that day. Finally, the next day, after two sessions in the original locations, the box was presented in a central location. Again, the animal was presented another two 10 min sessions in each environment, respectively, in a quasi-random order. In all stages, the running sessions were separated by a 20 min. break in the resting pot. On the test day, both environments were presented in two “reference” recording sessions (10 min each). After a 20 minutes break, the test session begun. The animal was inserted to the box with one set of lights on, and the lights were switched between the both sets after couple of minutes of recording.

2.2 Data structure

Cross validation of environment decoding methods. For the validation of environment decoding methods (Section 3.1) a total amount of four recordings were used. Two of these sessions, one in the environment *A* and one in the environment *B*, were used to infer activity models and reference statistics. The other two (again one in environment *A* and one in environment *B*) have then been used to test inferred models in the task of inferring which environment is internally-represented by the rodent. After the discretization in constant time windows described in section 2.3, only non-silent time bins (containing at least one firing event) were considered for the cross validation analysis.

Teleportation sessions. In the post-teleportation analysis shown in section 3.2 we used recordings from two different experiments. The light switch is performed once, and the activity is recorded for some minutes before and after the teleportation. Data comes from two recording sessions, each of which includes a reference session for both

environments and two teleportation sessions, one from A to B and one from B to A .

2.3 Map decoding methods

We consider two classes of decoders: *Rate-map based decoders*, which expressly use the knowledge of place fields and the rat trajectory as an input, and *Activity-only decoders* that do not rely on any information about the correspondence between position and neural firing. Throughout this section neural activities are binned with time resolution Δt ; we define the number of spikes of neuron i in time bin t , $n_{i,t}$, and the binary activity, $s_{i,t} = \min(n_{i,t}, 1)$. Little information is lost when considering s instead of n as long as Δt is smaller than the typical inter-spike interval of the cells.

Activity-only decoders

Bayesian approach to map decoding. We introduce probabilistic models for the distribution of activities $\{s_i\}$ in a time bin, $P(\{s_i\}, \Theta)$. Those models are parametrized by a set of variables, Θ , which are fit to maximize the likelihood of the data in reference sessions. Two sets of parameters $\Theta^{(m)}$ are fitted, one for each reference session $m = A, B$. We then define the difference in log-likelihood, for each time bin t :

$$\mathcal{E}(\{s_{i,t}\}) = \log \frac{P(\{s_{i,t}\}|\Theta^{(A)})}{P(\{s_{i,t}\}|\Theta^{(B)})}. \quad (1)$$

The sign of the quantity may be used to decode the map at time t . Significance levels, based on the percentiles of the distribution of \mathcal{E} can be imposed, see Results, Section 3.2.

Independent-cell model. We here assume that the neural activities s_i are independent from cell to cell. The probability distribution P is parametrized by a set $\Theta = \{h_i\}$ of N ‘inputs’ h_i :

$$P(\{s_i\}|\Theta) = \prod_i \frac{e^{h_i s_{i,t}}}{(1 + e^{h_i})} \quad (2)$$

h_i is fit such that the average value of s_i with P matches $\mu_i^{(m)}$, the mean value of $s_{i,t}$ across the time bins t in reference session relative to map m , yielding $h_i^{(m)} = \log \frac{\mu_i^{(m)}}{(1-\mu_i^{(m)})}$.

Graphical Ising model. We now take into account pairwise correlations between neural activities s_i in a time bin. We introduce couplings J_{ij} to express the conditional probability that cell i is active given the activity of cell j . The probability distribution P is now parametrized by the set $\Theta = \{h_i, J_{ij}\}$ of N inputs h_i and $\frac{1}{2}N(N-1)$ couplings J_{ij} :

$$P(\{s_i\}|\Theta) = \frac{1}{\mathcal{Z}[\{h_i, J_{ij}\}]} \exp \left(\sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j \right) \quad (3)$$

where \mathcal{Z} is the normalization constant. Parameters $h_i^{(m)}$ and $J_{ij}^{(m)}$ are computed such that the average values of s_i and $s_i s_j$ with P match the mean values of, respectively, $s_{i,t}$ and $s_{i,t} s_{j,t}$ across the time bins t in reference session relative to map m . This hard computational problem can be approximately solved with the Adaptive Cluster Expansion (ACE) algorithm (Barton et al., 2016; Cocco and Monasson, 2011), which provides estimates of the parameters $\{h_i^{(m)}, J_{ij}^{(m)}\}$ and $\mathcal{Z}^{(m)}$ in Eq. (3).

Rate-map based decoders

Computation of rate maps. The squared box is partitioned into a 20×20 grid of 3×3 cm² bins, and the rat position during the two reference sessions is discretized with respect to this grid. The coordinates (x_t, y_t) associated to time bin t correspond to the first spatial bin visited by the rat in the time interval $[t - \Delta t; t]$. We define the average firing rate $r_i^{(m)}(x, y)$ as the total number of spikes emitted by neuron i in the reference session m when the rat is at position (x, y) , divided by the total time $T^{(m)}(x, y)$ spent by the animal in this spatial bin. These rate maps are then smoothed to fill missing bins through discrete cosine transform (Garcia, 2010).

Pearson decoder. The observed firing pattern $\{n_{i,t}\}$ is compared to the average firing rates $\{r_i^{(m)}(x_t, y_t)\}$ in the position occupied by the animal at the time the neural activity is recorded (Jezek et al., 2011). This comparison is made through the Pearson correlation

$$\mathcal{C}^{(m)}(\{n_{i,t}\}) = \frac{\langle n r^{(m)}(x_t, y_t) \rangle_t - \langle n \rangle_t \langle r^{(m)}(x_t, y_t) \rangle_t}{\sqrt{(\langle n^2 \rangle_t - \langle n \rangle_t^2) (\langle r^{(m)}(x_t, y_t)^2 \rangle_t - \langle r^{(m)}(x_t, y_t) \rangle_t^2)}} \quad (4)$$

where the notation $\langle f \rangle_t := \frac{1}{N} \sum_{i=1}^N f_{i,t}$ denotes the average of the quantity $f_{i,t}$ over the N neurons i in time bin t . The decoding of the map is done according to the sign of

$$\mathcal{E}(\{n_{i,t}\}) = \mathcal{C}^{m=A}(\{n_{i,t}\}) - \mathcal{C}^{m=B}(\{n_{i,t}\})$$

Dot-product decoder. The second method used in Jezek et al. (2011) compares directly the activity to the firing rates at the rat position. The decoding of the map m is done according to the sign of

$$\mathcal{E}(\{n_{i,t}\}) = \langle n r^{(m=A)}(x_t, y_t) \rangle_t - \langle n r^{(m=B)}(x_t, y_t) \rangle_t$$

Bayesian Poisson rate model. This model assumes that each neuron fires independently following a Poisson statistics, with a position-dependent firing rate $r_i^{(m)}(x, y)$ in map m . The likelihood of the number of spikes $\{n_i\}$ emitted by the cells in a time bin when the rat is at position (x, y) reads

$$P^{(m)}(\{n_i\} | (x, y)) = \prod_i \frac{(r_i^{(m)}(x, y) \Delta t)^{n_i}}{n_i!} e^{-r_i^{(m)}(x, y) \Delta t} \quad (5)$$

The prior probability over positions is

$$P^{(m)}(x, y) = T^{(m)}(x, y) / \sum_{x', y'} T^{(m)}(x', y')$$

decoder output	A	B	A	B
cue	A	A	B	B
denomination	True Positive	False Negative	False Positive	True Negative

Table 1: Denomination used for the four possible events, depending on the output of the decoder and on the environment-defining cue.

and we assume that both maps m are *a priori* equally likely. We obtain the likelihood of the activity conditioned to the map m by marginalizing over positions

$$P(\{n_i\}|m) = \sum_{x,y} P(\{n_i\}|(x,y)) \times P^{(m)}(x,y)$$

We then define the difference in log-likelihoods

$$\mathcal{E}(\{n_i\}) = \log P(\{n_i\}|m = A) - \log P(\{n_i\}|m = B)$$

2.4 Performance measure of a binary decoder

To quantitatively assess decoding performance of map-decoding methods we refer to binary classifier theory (Bradley, 1997; Hanley and McNeil, 1982; Metz, 1978; Pencina et al., 2008).

Receiver Operating Characteristic (ROC) diagrams. A standard framework to assess the performance of binary decoders is the so-called ROC diagram (Hanley and McNeil, 1982). For each time bin t the decoder outputs either map A or map B . To match the vocables used in the ROC framework we will arbitrarily say that the output is *positive* if the map is decoded to be A , and *negative* if the map is predicted to be B . If the output of the decoder matches the environment defined by the light cues at the same time t , the prediction is said to be *True*, otherwise it is said to be *False*. For instance, a time bin such that the decoder predicts A , in agreement with the cues, corresponds to a True Positive event. The 2×2 possible events are shown in Table 1. Two important quantities are: the True Positive Rate (TPR, also called Recall), that is, the number of true positive predictions divided by the total number of positive events, and the False Positive Rate (FPR), that is, the number of false positive predictions, divided by the total number of negative events. In other words, the TPR measures the fraction of time bins with A -cues that are correctly decoded as A , while the FPR is the fraction of time bins with B -cues that are incorrectly predicted to be A .

Our binary decoders are all based on thresholding the *estimator* variable \mathcal{E} . Within the Bayesian framework, for instance, we compute \mathcal{E} as the difference between the logarithms of the posterior probabilities of A and B , and output Positive if the difference is larger than $\theta = 0$, Negative otherwise. The value of the significance threshold θ can be arbitrarily changed, with the consequence of modifying the TPR and FPR values. A ROC curve shows the parametric plot of TPR vs. FPR as the threshold varies, and describes a curve in the unit square, see Results, Section 3.1. The two extreme points

of the ROC curves have coordinates (0,0), and (1,1); (0,0) is obtained for a very large significance threshold θ , the decoder never outputs Positive and both TPR and FPR vanish; (1,1) is obtained when the significance threshold is very low, the decoder always outputs Positive and both TPR and FPR are equal to unity. Very good decoders are such that the TPR is close to unity, while maintaining a very low value for the FPR. A random-guessing decoder would give equal values for the TPR and FPR, and the ROC curve would coincide with the diagonal of the unit square.

A complementary measure of decoding performances is the Precision versus Recall (or TPR) curve, obtained by scanning the values of the significance threshold θ , see Results, Section 3.1. The Precision is defined as the number of true positive events, divided by the total number of positive predictions. When lowering the significance threshold the Precision decreases from 1 to 0, while the Recall increases from 0 to 1.

Area Under the Curve (AUC). A quantitative measure of the decoding performances is the Area Under the (ROC) Curve (Hanley and McNeil, 1982). According to this measure, the ideal decoder has $AUC = 1$, while random guessing would give $AUC = 0.5$. Note that this measure is invariant with respect to the arbitrary choice of assigning *positive* value to environment *A*: if we assign positive to *B* and negative to *A* instead of the previous choice, ROC curves will undergo a symmetry transformation with respect to the top-left/bottom-right diagonal, resulting in an identical area under the curve. This is granted by the fact that positive and negative values are mutually exclusive and complementarily cover the whole data set: for each Θ value the fraction of False Positive (*B* decoded as *A*) equals one minus the fraction of True Negative (*B* decoded as *B*) events.

2.5 Continuity prior for map decoding

A continuity prior can be included in map inference in order to reduce noise in the decoding and highlight clusters of contiguous transited time bins. To do so, we consider the output $\{\mathcal{E}_t\}$ of the map decoder (see Section 2.3); for Bayesian decoders \mathcal{E}_t is the difference between the log-likelihoods of the two maps $m_t = +1$ and -1 in time bin t . We then introduce a prior, controlled by a strength parameter K , which favors persistence between decoded maps in nearby time bins. Informally speaking, K is the cost (in log-likelihood) we are willing to pay for flipping the map index in time bin t predicted by the sign of \mathcal{E}_t to its opposite value, if it then matches the map indices of the neighboring time bins, $t - 1$ or $t + 1$. The prior may thus be effective in changing the map prediction m_t if the differences between \mathcal{E}_{t-1} , \mathcal{E}_t , \mathcal{E}_{t+1} , ... are of the order of K (in absolute value). Two situations are encountered: (1) for some decoders, *e.g.* Pearson, \mathcal{E}_t takes value in $[-2; 2]$, and the variations of \mathcal{E} over successive time bins is bounded; (2) for other decoders, *e.g.* Independent-Cell, Poisson and Ising, the difference between \mathcal{E}_t and \mathcal{E}_{t+1} can take arbitrarily large values and show wide fluctuations as t varies across the recording. In the latter case, a uniform prior K is unadequate in large portions of the recording. To circumvent this difficulty we introduce a scale factor $\beta < 1$, and multiply all outcomes \mathcal{E}_t by this factor. As a result we get a smoother time course of \mathcal{E}_t over the time index t , on which a uniform prior can now be applied.

The joint probability of the time sequence of map predictions $\{m_t\}$ reads

$$P(m_1, m_2, \dots, m_T) = \frac{1}{\mathcal{Z}} \exp \left(\frac{\beta}{2} \sum_{t=1}^T \mathcal{E}_t m_t + K \sum_{t=1}^{T-1} m_t m_{t+1} \right) \quad (6)$$

where \mathcal{Z} is a normalization coefficient. To decode the map in time bin t we compute the marginal probability P_t over m_t from the joint distribution P . Exploiting the analogy with the one-dimensional model of statistical physics, this computation can be done with the transfer matrix method, also called dynamic programming in a time scaling linearly with the total number of time bins. Then the outcome of our combined decoder+prior is

$$\mathcal{E}_t^{\text{decoder+prior}} = \frac{1}{\beta} (\log P_t(m_t = +1) - \log P_t(m_t = -1)) \quad (7)$$

The presences of the $\frac{1}{2}$ and $\frac{1}{\beta}$ factors in, respectively, Eq. (6) and Eq. (7) ensure that, for $K = 0$, $\mathcal{E}_t^{\text{decoder+prior}}$ and \mathcal{E}_t coincide. In practice we choose $\beta = \frac{1}{|\mathcal{E}_0|}$, where $\mathcal{E}_0 := \max_t \{|\mathcal{E}_t|\}$.

Induced correlation as a function of K . The transfer matrix technique allows us to compute also the correlation between the maps decoded τ bins apart, defined as

$$C(\tau) := \frac{1}{T - \tau} \sum_{t=1}^{T-\tau} (\langle m_t m_{t+\tau} \rangle - \langle m_t \rangle \langle m_{t+\tau} \rangle) \quad (8)$$

where the angular-bracket notation denotes the average over the probability distribution in Eq. (6). $C(\tau)$ decays exponentially with τ , over a characteristic ‘time’ monotonically growing with K in Eq. (6), see Results, Section 3.1.

3 Results

3.1 Cross validation of map-decoding methods

Decoding methods and number of parameters used

We start by presenting map-decoding methods and their performances. For each environment, A and B , we have two recorded sessions with constant light cues: the first one, called reference session, is used to infer the decoder parameters. The second one, called test session, is used for cross-validation, *i.e.* to assess the performances of the decoder. We compare the performances of five different decoders, described in Methods, Section 2.3. Our decoders mainly differ by the fact that they may use or not knowledge of the rat positions and of the spatial rate maps (place fields). They are also based on simple comparison methods or on more sophisticated probabilistic frameworks.

Rate-map based decoders require the computation of the rate maps during the reference session. Knowledge of the position $\vec{x}(t)$ of the rats and of the neural firing rates $r_i(t)$ as a function of time t allows one to build the rate maps, that is, the average firing rate of each cell i as a function of the rat position \vec{x} , $r_i^{(E)}(\vec{x})$ for environment $E = A, B$. The similarities between those reference population activities and the activity measured during the test sessions may then be used as a simple estimator of the map retrieved by the rodent. We consider two such comparison-based approaches, called *Dot Product* and *Pearson* (Jezek et al., 2011). A more sophisticated decoder, called *Poisson*, consists in assuming that each place cell i fires with a Poisson process, with average rate $r_i^{(E)}(\vec{x})$ when the rodent is at position \vec{x} , and in estimating the likelihood of the test spiking activity with this multiple Poisson process and for maps $E = A$ and B . The posterior distribution for the (binary) map variable E can then be computed, and we decode the map as the one with larger posterior probability. Poisson is based on a more solid probabilistic framework than Dot Product and Pearson, while making use of the same rate maps estimated from the reference sessions.

Activity-only decoders do not need any information about rat position and place fields. Those models provide approximate expressions for the probability distribution of population neural activity over short time bins, *i.e.* of binary (silent or active neuron in the time bin) strings of length N (the number of recorded neurons). The *independent-cell* model is the simplest maximum-entropy model (Jaynes, 1957); it reproduces the N average activities of the neuron only. The second model, called *Ising* in statistical physics, is a graphical model that, in addition, reproduces the pairwise correlations between the neural activities in a time bin (Cocco and Monasson, 2011; Jaynes, 1957; Schneidman et al., 2006). The Ising model requires the inference of pairwise effective couplings between every two cells, which we have performed with the Adaptive Cluster Expansion method (Barton et al., 2016; Cocco and Monasson, 2011). Similarly to Poisson, the independent-cell and Ising models provide estimates of the likelihood of the population

activity in a time bin, and can be used to compute the posterior distribution for the map variable, E , and to decode the retrieved map through maximization over E .

As a consequence the numbers of parameters to be learned from the reference sessions vary a lot with the decoders. For N recorded neurons (38 in one of the data sets studied here, see Materials) and a discretization of the environment into S ($=20 \times 20$ in the present analysis) spatial bins, the numbers of parameters to be extracted from the reference sessions are, respectively $N = 38$ for the independent-cell decoder, $\frac{1}{2}N(N+1) = 741$ for the Ising decoder, and $N \times S = 15,200$ for the Poisson, Pearson, and Dot Product decoders.

Inferred Ising couplings are fingerprints of environment representation in CA1

As a result of rate remapping taking place in CA1 (Fig. 1) the populations of active cells in the two environments are similar. This property can be seen from the comparison of the inputs $\{h_i\}$ in the Ising models inferred in the reference sessions of the two environments, see Fig. 2. The input h_i to place cell i take similar values across the environments; its value is indicative of the average firing rate of the cell (Methods, Section 2.3).

Distinction between the neural representations of the environments in CA1 can, however, be drawn from the correlational structure of firing events in the place-cell population. Place cells with overlapping firing fields in one environment are indeed more likely to be simultaneously active during the animal's exploration, and their activity is thus correlated. Due to remapping the amplitudes of these correlations are specific to each environment. The inferred Ising couplings $\{J_{ij}\}$, which capture the direct correlation between cells i, j not mediated by other recorded cells (Methods, Section 2.3), are different from one environment to the other, as shown in Fig. 2.

The set of effective couplings $\{J_{ij}\}$ is therefore a *fingerprint* of the environment (Okatan et al., 2005), which we can exploit to distinguish between maps, *i.e.* to decode the neural representation. Note that these effective, functional couplings are not directly related to the physiological synaptic interactions, which are not accessible from the data.

Inferring effective couplings allows us to *score* any configuration of the population activity, that is, to quantitatively assess its similarity with typical activities in each environment, as shown below. This score is, in practice, given by the Ising likelihood, see Methods, Eq. (3), and heavily relies on the inferred couplings.

Comparison of performances of map-decoding methods

We present a systematic study of the performances of map-decoding methods in CA1 within the framework of binary-decoder theory, see Methods, Section 2.4 for a detailed description. Results are reported in Fig. 3.

We plot in Fig. 3 (a) the Receiver Operating Characteristic (ROC) curve for the Ising and Pearson decoders. Briefly speaking, ROC curve shows the value of the True Positive Rate (fractions of time bins in reference session for environment A for which the decoder

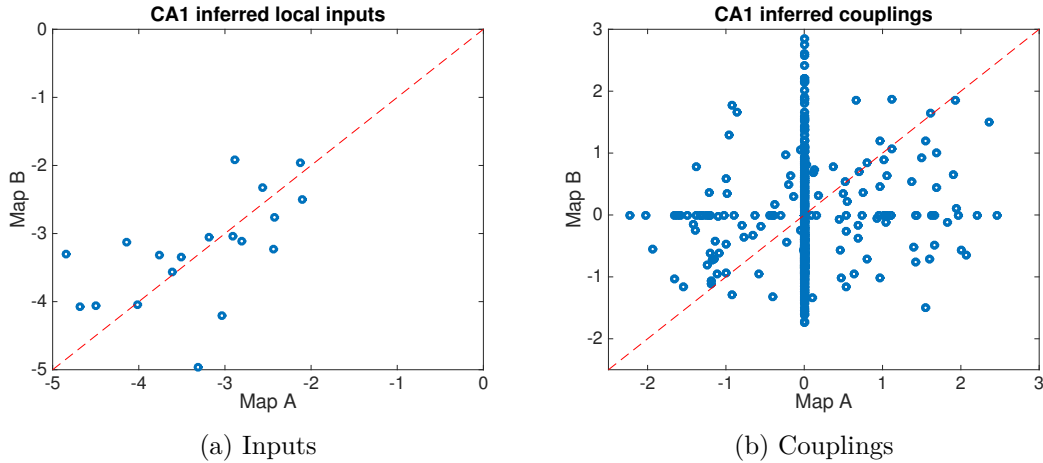


Figure 2: **Comparison of inferred Ising parameters across the two maps.** Top: Inputs h_i of the Ising models inferred from reference sessions. Only values greater than -5 , corresponding to a firing rate of c.a. 0.05 Hz in the independent-cell model, are shown. Bottom: Couplings J_{ij} of the Ising models inferred from reference sessions. Note the presence of many zero couplings in both environments. Analysis performed with discretization time bin $\Delta t = 120$ ms.

rightly decodes map A) as a function of the False Positive Rate (fractions of time bins in reference session for environment B for which the decoder erroneously recognizes map A). A random decoder would have equal values for TPR and FPR, and lies on the diagonal line of the unit square in Fig. 3 (a). A perfect decoder would always recognize map A in environment A and never in environment B , and would thus correspond to TPR = 1, FPR = 0. Varying the threshold for significance of the decoder changes both the values of TPR and FPR, with the resulting ROC shown in Fig. 3 (a). We observe that the Ising decoder shows much better performances than the Pearson decoder. An alternative representation of the decoder performances is given by the Precision-Recall curve, shown in Fig. 3 (b), see Methods, Section 2.4 for definition.

A measure of the accuracy of the decoder is given by the integral of the ROC curve, called Area Under the Curve (AUC), which ranges from 0.5 for a random decoder to 1 for a perfect decoder. To compare the five decoders we plot in Fig. 3 (c) their AUC values as a function of the elementary time bin Δt , ranging from 10 ms to 1s. The Ising model, which takes into account the correlational structure of the population activity, has higher decoding precision and retrieval capacity than other decoders in CA1 recordings (Fig. 3 (a,b)). As a consequence, in terms of AUC (Fig. 3 (c)), *Ising* is generally the most performant model, followed by *Poisson* and lastly by equally-performant *Pearson* and *independent-cell* decoders. *Dot Product* method is the best performant on very short time scales (< 20 ms), but its performance increases very slowly with the time bin width, and as a consequence it has the worst performance for $\Delta t > 100$ ms.

This behavior has an explanation in terms of sensitivity of the different models to the *average number of active neurons per time bin*. Bayesian models, whose predictions do not depend on the specific position of the rat at each time, rely on information conveyed

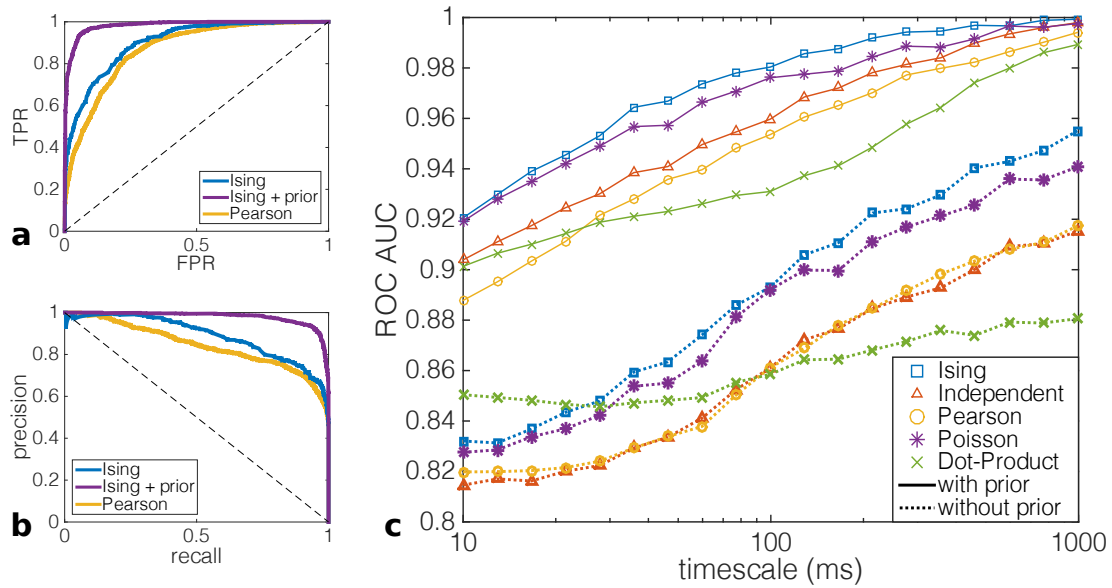


Figure 3: Performances of spatial representation decoders. ROC (a) and Precision-Recall (b) curves computed at fixed time scale $\Delta t = 120$ ms for a combination of two test sessions in environments *A* and *B*, recorded in CA1. Maps *A* and *B* correspond, respectively, to positive and negative predictions, see Table 1. The True Positive Rate, also called Recall, is the number of true positive predictions divided by the total number of positive events. The False Positive rate is the number of false positive predictions, divided by the total number of negative events. Precision is defined as the fraction of identified positive events that are true positives. (c) performances of *Ising*, *Independent-cell*, *Poisson*, *Pearson*, and *Dot Product* decoders (with and without the addition of a continuity prior) as a function of the discretization time scale, applied to CA1 neural recordings. Full and dashed curves correspond to predictions, respectively, without and with continuity prior; in the latter case the correlation C in Eq. (8) decays over $t_0 = 2$ time bins (Methods, Section 2.5 and Fig. 4 (a)).

by activity alone. As a consequence, when the number of simultaneously active neurons for each time bin is very small, Bayesian models may be less accurate than decoders that take into account spatial information, like Dot Product.

As a general feature we observe that the performances of all decoders improve on larger discretization time scales. Statistical fluctuations in the activity are indeed smaller for large time bins, resulting in smaller errors in the decoding task. A similar reduction of statistical fluctuations can be obtained through the introduction of a continuity prior, which prevents switching back and forth between spatial maps.

Map decoding with continuity prior

Map decoding can be combined with a continuity prior that enhances persistence in the decoded maps over consecutive time bins, see Methods, Section 2.5. The motivation for the continuity prior is two-fold. First, in situations where the latency between a delivery of external stimulus and the network state change is the main parameter to be measured (e.g. after pharmacology treatment, etc.), one needs to search for a single time point of the state transition. This can be achieved by imposing a strong continuity prior, allowing for the presence of a single transition between maps along the whole recording session.

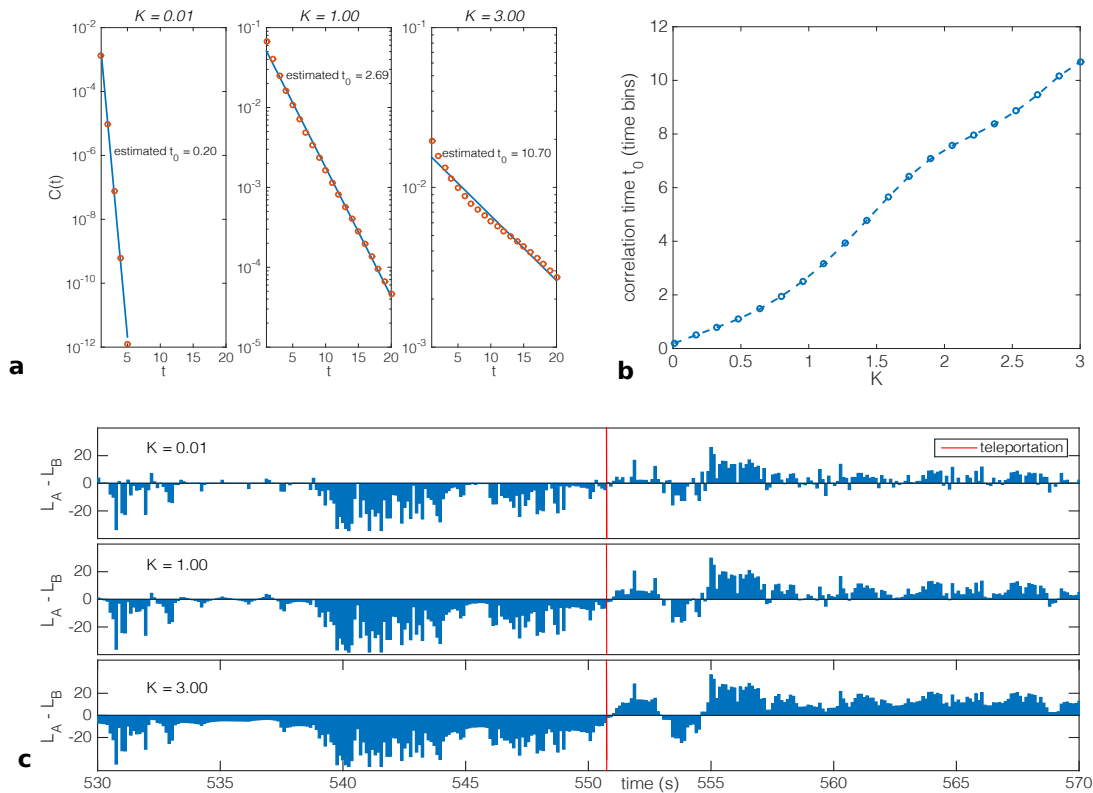


Figure 4: **Continuity prior for map decoding.** (a) Correlation (Methods, Eq. (8)) between maps decoded in two time bins as a function of their separation t (measured in units of time bins), for three values of the prior strength K . Correlations are well fitted by exponential decaying functions, over a characteristic number of bins t_0 . (b) Value of t_0 as a function of the prior strength K . (c) Application of the prior on CA1 teleportation session for different values of prior strength parameter K . Ising decoder, with a discretization time bin $\Delta t = 120$ ms.

Secondly, with moderate continuity prior, dynamical events (such as state transitions) can be detected with more precision, at the price of discarding events that happen on a time-scale shorter than a given temporal resolution. When the prior is applied, the resulting correlation between decoded maps in two time-bins that are t bins apart decays

exponentially with t , see Fig. 4 (a) and Methods, Eq. (8). A persistence ‘time’ t_0 can be computed through an exponential fit of the correlation: t_0 is the characteristic number of bins over which decoded maps are persistent. Its value can be chosen at our convenience by tuning the prior strength parameter K , see Fig. 4 (b) and Methods, Section 2.5. Hence, we can choose a temporal resolution t_0 and exploit the noise-cancelling property of the continuity prior over larger time scales.

Unless otherwise specified we set in the following the characteristic persistence time to the small value $t_0 = 2$ time bins. As shown in Fig. 3, the use of this weak continuity prior sensibly improves decoding performances with the Ising method. The AUC increases by about 10%, see Fig. 3 (c). For direct comparison, if one instead increases the time-bin resolution Δt by a factor 2, the increase in AUC is much lower (Fig. 3 (c)): for instance, Ising AUC is equal to 0.90 for $\Delta t = 120$ ms and to 0.92 for $\Delta t = 240$ ms, while it reaches 0.98 for $\Delta t = 120$ ms with a continuity prior such that $t_0 = 2$ time bins. This results shows the statistical efficiency of the continuity prior.

3.2 Transitions between maps in "teleportation" experiment

The Ising decoder presented above can be used to identify transitions between different activity states of dynamic systems. Brain hippocampal memory circuitry is a dynamic system expressing distinct states of activity - neural representations of surrounding space - with attractor properties (Colgin et al., 2010; Jezek et al., 2011; Wills et al., 2005). We recorded new CA1 data in a ‘teleportation’ setup first introduced for CA3 recordings in Jezek et al. (2011), in which the appearance of recording box is abruptly changed by switching between two familiar light cue settings (A and B , respectively) while the laboratory rat continuously explores it (Methods, Section 2.2). This procedure was shown to induce a rapid exchange of corresponding hippocampal representations in CA3, including periods of instability with spontaneous fast flickering between them. Transitions between the maps were identified based on activity models of representations A and B , respectively, inferred from reference recordings in both environments under stable conditions preceding the ‘teleportation test session’.

Teleportation procedure induces long-term network instability

We use our Ising decoder to assess the effect of teleportation on subsequent network state expression in CA1. Representative evolution of the log-likelihood (with the Ising model inferred from the reference sessions) of the identified network state is shown before and after two teleportation events in Fig. 5(a). The criterion for accepting given bin as corresponding to representation of environment A or B , respectively, was set to match 1% error derived from stable reference sessions, see Fig. 5 (b). This ensures that a time bin is identified as A only if there is 99% confidence that its signal value does not come from state B (and vice-versa).

To characterize the kinetics of network state development, we identify the amount of time bins expressing a neural representation that was incongruent (non-corresponding) with the present environment, *i.e.* coding for the environment presented before the

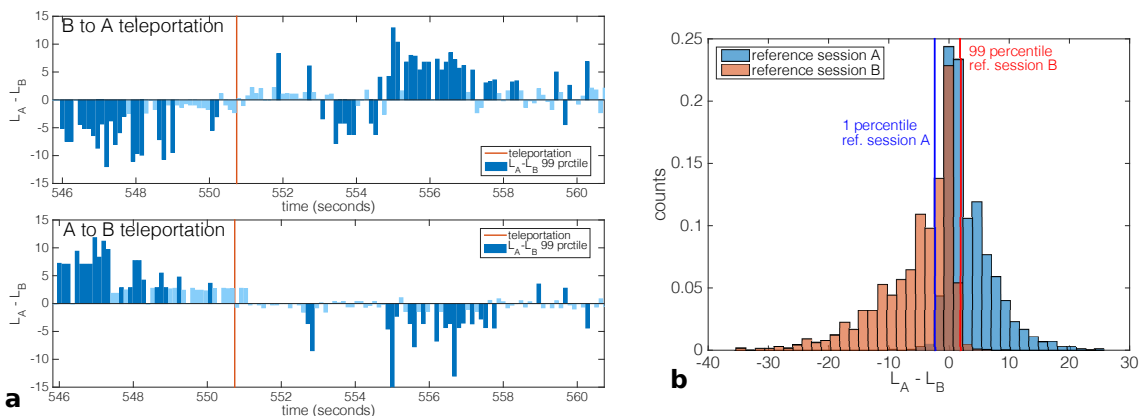


Figure 5: Log-likelihoods of neural activities around teleportation events. Difference $L_A - L_B$ of the log-likelihoods computed with the Ising decoder applied to the neural activity recorded in a teleportation session (a), with light-cue switch from environment B to A (top) and from A to B (bottom). The light switch is marked with a red line, predictions higher than 99 percentile value of reference sessions are colored in dark blue, weaker prediction are colored in light blue. Panel (b) shows the distributions of log-likelihoods in reference sessions. A percentile value θ in $[0, 100]$ (normally in the interval $[90, 100]$) is defined. We consider a test time bin as significantly decoded as A only if the log-likelihood of the activity configuration in the time bin is higher than the θ percentile value of reference session B , and as B only if its value is lower than the $100 - \theta$ percentile value of reference session A . The underlying reasoning is to decode a test time bin as A only if it is very unlikely that it comes from reference population B , and vice-versa.

teleportation. We estimate the short term effect within interval of first 10 seconds, and an eventual long term effect beyond the first 60 seconds from teleportation. The rates of incongruent bins are shown in Fig. 6. The amount of non-corresponding events were larger during the short-term period (less than 10 sec) after the light switch than those identified in the stable conditions before the teleportation test (0.017 ± 0.004 S.E.M. stable; 0.101 ± 0.029 S.E.M. after switch, $F = 11.24$, $p < 0.015$). These results show that the short-term instability found in CA3 in Jezek et al. (2011) extends to CA1. In addition an interesting finding comes from the analysis of long-term (> 60 sec) effects shown in Fig. 6. Compared to stable environment conditions, the rate of flickering is found to be significantly larger (0.047 ± 0.007 S.E.M. long-term post switch, $F = 15.21$, $p < 0.008$), and this increase persists over several minutes after the switch, see Fig. 6. Notice that the apparent difference between the long-term and short-term enhancements in network instabilities is not statistically significant ($F = 4.56$, $p = 0.076 > 0.05$).

Identification of transitions with strong continuity prior

The network state-decoding procedure with implemented continuity prior can be used to detect internal state-shifts under predefined criteria. For instance, when the prior strength is brought to extreme values the decoding procedure discards the fast instability-




Figure 6: Teleportation enhances network instability over both short and long term periods. Percentage of temporal bins expressing the environment-incongruent coding computed in rest conditions (pre), during the first 10 seconds after a light switch (short post), and in long-term period after the teleportation event (long post, more than 60 seconds after light switch). Only bins expressing likelihood values higher than 99 percentile of reference sessions have been taken into account; Similar results are obtained with 90 and 95 percentiles. Values were averaged over four sessions, with recording durations before (pre) / after (post) teleportation equal to, respectively, 10/8, 9/9, 10/8, and 9/7 minutes. Analysis performed with Ising environment decoder with discretization time bin $\Delta t = 120$ ms.

driven dynamics and, instead, returns a single state transition time point that reflects the evolution of log-likelihood values across the continuum of temporal bins. Taking as an illustration the CA1 teleportation session in Fig. 7, we see that the response of network activity state to the teleportation event is identified with high accuracy. This is a valuable tool to measure the most probable moment of network remapping even under widely fluctuating dynamics.

4 Discussion

Functional connectivity-based models for map decoding: Ising and other models. The Ising decoder introduced in this article yields the highest performance on all time scales Δt in CA1 (Fig. 3). While we have here considered the activity vectors as discretized in regular-spaced time windows of duration Δt it would be easy to extend our analysis to process activity in elementary windows in correspondence to Theta cycles. Such an analysis would be interesting to assess the plausibility of transition scenarios put forward by theoretical studies (Monasson and Rosay, 2015) and deepen our understanding of the role of Theta oscillations for the dynamics of transitions (Jezek et al., 2011).

In this regard repeating the present study with probabilistic models capable of capturing some aspects of the activation dynamics in recorded spiking sequences, such as Generalized-Linear Models (Truccolo et al., 2005) could be potentially interesting. Con-




Figure 7: **Network state transitions identified by implementation of continuity prior.** Ising decoder and Viterbi algorithm with strong continuity constraint applied to neural activity in a CA1 teleportation session with enlarged examples. Light switches are marked with red lines. Analysis performed with time bin $\Delta t = 120$ ms.

trary to their Ising model counterparts effective couplings in the GLM approach are not necessarily symmetric, and may reflect specific ordering in neuron activations. However, some basic assumptions underlying GLM, such as the Poissonian nature of firing events are questionable for hippocampal place cell activity (Fenton and Muller, 1998). Another potentially interesting alternative is provided by reverse engineering of networks of Integrate-and-Fire neurons (Koyama and Paninski, 2010; Makarov et al., 2005; Monasson and Cocco, 2011), which were already applied to recordings, *e.g.* of retinal data with tens of neurons.

Graphical models for brain state identification in the absence of input correlates. Methods for decoding spatial representations considered in this work can be divided in two classes, depending on whether they make use of positional information or not. Remarkably, the latter methods do not show worse performances than the former approaches. In CA3, efficient decoding does not require the use of sophisticated probabilistic models: due to the quasi-orthogonality of maps, the simple independent-cell decoder, which compares the activity at any time to the average activities in environments A and B irrespectively of the rat position shows very good performances (Posani et al., 2016). In CA1, the similarity in the firing fields across environments constrained us to consider a graphical model, the pairwise Ising model, which not only captures the average activity of the place cells but also their pairwise correlations. The higher performance of the Ising model, combined with the lower number of parameters involved in the inference process compared to firing field-based decoding methods, suggests that the correlational structure of neural firing activities conveys essential information about the internal representation of memorized environments.

A substantial advantage of this approach is that it can be effectively applied to other brain regions with much weaker correlation between the local activity and its inputs,

e.g. the prefrontal cortex, or without any known input-output relation. The use of graphical models does not require any knowledge about the network inputs, because the activity-states are merely derived from a (high-dimensional) fit of the correlation structure of the spiking data. The core idea, first put forward in the context of retinal data modeling (Schneidman et al., 2006), is that the model obtained after inference of the functional network is an approximate (albeit quantitatively accurate, compared to principal-component based approaches (Lin and Gervasoni, 2008)) description of the distribution of activities characterizing specifically one brain state. Provided that we have at our disposal different data sets for well-identified states (here, the reference sessions) we may later use the inferred model to decode the activity at any time. This approach has recently been applied to identify transient activation of memory-related cell assemblies in the rat prefrontal cortex (Tavoni et al., 2015, 2016), and we may expect, owe to its generality and its applicability to very fast time scales (down to ~ 10 msec), further applications in future.

Instabilities in hippocampal space representations. In the CA3 area of hippocampus, patterns of place cell activity across different environments behave as uncorrelated network states with attractor properties (Wills et al., 2005). Transitions between those hippocampal activity states were recently studied based on recordings taken during a free exploration in two environments in an experimental paradigm shown to induce rapid switches (Jezek et al., 2011). In the present paper we used multiunit recordings from hippocampal area CA1. Both CA3 and CA1 are parts of the entorhino-hippocampal loop, an essential circuit for spatial memory and navigation in mammalian brain. Despite being directly connected in series (CA3 signalling into CA1), they very much differ in their architecture - while CA3 is organized as a recurrent network with attractor properties, CA1 has a feed forward structure - and in their connections with other brain areas involved into space representation (Knierim, 2006).

Use of the Ising model allowed us to robustly decode the memory state expressed in the CA1 network, with temporal resolution high enough to reflect natural time patterning of activity provided by local theta oscillation (ca. 6-11 Hz). We could track the network state kinetics following the sensory input switch. In agreement with previous report in CA3 (Jezek et al., 2011), we detected a high degree of flickering in CA1 following the switch. Importantly, we found so far unseen flickering increase in long-term scale after the teleportation event (more than 60 seconds) in CA1 compared with the stable environment conditions, see Fig. 6. A similar long-term instability can also be found in CA3 data (Posani et al., 2016). This effect is rather surprising as the network usually achieves relative stability within a couple of seconds after the cue switch. Whereas the short-term instability can be explained by reverberatory activity in CA3 and by competing visual and self-motion inputs, the long-term flickering will probably reflect inner, sensory-independent processing with so far unknown origin.

Acknowledgements We are indebted to S. Rosay, who contributed to the early stage of the data analysis. We are grateful to J. Tübiana and A. Treves for useful discussions and suggestions. This study benefited from partial fundings from the CNRS-InphyNiTi INFERNEUR project and from GACR 15-20008S, PRVOUK P-36 and NPU I LO1503 of the Czech Republic.

References

- John Barton, Eleonora De Leonardis, Alice Coucke, and Simona Cocco. Ace: adaptive cluster expansion for maximum entropy graphical model inference. *Bioinformatics*, 2016.
- Andrew P Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.
- Simona Cocco and Rémi Monasson. Adaptive cluster expansion for inferring boltzmann machines with noisy data. *Physical review letters*, 106(9):090601, 2011.
- Laura L Colgin, Stefan Leutgeb, Karel Jezek, Jill K Leutgeb, Edvard I Moser, Bruce L McNaughton, and May-Britt Moser. Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *Journal of neurophysiology*, 104(1):35–50, 2010.
- A.A. Fenton and R.U. Muller. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proc. Natl. Acad. Sci. USA*, 95:3182–3187, 1998.
- Damien Garcia. Robust smoothing of gridded data in one and higher dimensions with missing values. *Computational statistics & data analysis*, 54(4):1167–1178, 2010.
- James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- Karel Jezek, Espen J Henriksen, Alessandro Treves, Edvard I Moser, and May-Britt Moser. Theta-paced flickering between place-cell maps in the hippocampus. *Nature*, 478(7368):246–249, 2011.
- James J Knierim. Neural representations of location outside the hippocampus. *Learning & Memory*, 13(4):405–415, 2006.
- S. Koyama and L. Paninski. Efficient computation of the maximum a posteriori path and parameter estimation in integrate-and-fire and more general state-space models. *J. Comp. Neurosci.*, 29:89, 2010.

- S.C. Lin and D. Gervasoni. Defining global brain states using multielectrode field potential recordings. In Nicolelis MAL, editor, *Methods for Neural Ensemble Recordings*. CRC Press/Taylor and Francis, 2008.
- David JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- V. A. Makarov, F. Panetsos, and O. de Feo. A method for determining neural connectivity and inferring the underlying network dynamics using extracellular spike recordings. *J. Neurosci. Methods*, 244:165, 2005.
- Charles E Metz. Basic principles of roc analysis. In *Seminars in nuclear medicine*, volume 8, pages 283–298. Elsevier, 1978.
- R. Monasson and S. Cocco. Fast inference of interactions in assemblies of stochastic integrate-and-fire neurons from spike recordings. *J. Comp. Neurosci.*, 31:199–227, 2011.
- Rémi Monasson and Sophie Rosay. Transitions between spatial attractors in place-cell models. *Physical Review Letters*, 115:09810, 2015.
- Murat Okatan, Matthew A Wilson, and Emery N Brown. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural computation*, 17(9):1927–1961, 2005.
- John O’Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 34(1):171–175, 1971.
- Michael J Pencina, Ralph B D’Agostino, Ralph B D’Agostino, and Ramachandran S Vasan. Evaluating the added predictive ability of a new marker: from area under the roc curve to reclassification and beyond. *Statistics in medicine*, 27(2):157, 2008.
- Jonathan W Pillow, Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M Litke, EJ Chichilnisky, and Eero P Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999, 2008.
- Lorenzo Posani, Simona Cocco, Karel Jezek, and Rémi Monasson. Persistence of positional information in unstable spatial representations in CA3. *preprint*, 2016.
- Mattia Rigotti, Omri Barak, Melissa R Warden, Xiao-Jing Wang, Nathaniel D Daw, Earl K Miller, and Stefano Fusi. The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–590, 2013.
- Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

- Ian H Stevenson, James M Rebeco, Lee E Miller, and Konrad P Körding. Inferring functional connections between neurons. *Current opinion in neurobiology*, 18(6):582–588, 2008.
- Gaia Tavoni, Ulisse Ferrari, Francesco Paolo Battaglia, Simona Cocco, and Rémi Monasson. Inferred model of the prefrontal cortex activity unveils cell assemblies and memory replay. *bioRxiv*, page 028316, 2015.
- Gaia Tavoni, Simona Cocco, and Rémi Monasson. Neural assemblies revealed by inferred connectivity-based models of prefrontal cortex recordings. *Journal of Computational Neuroscience*, 41:269–293, 2016.
- Gasper Tkacik, Jason S Prentice, Vijay Balasubramanian, and Elad Schneidman. Optimal population coding by noisy spiking neurons. *Proceedings of the National Academy of Sciences*, 107(32):14419–14424, 2010.
- W. Truccolo, U.T. Eden, M.R. Fellows, J.P. Donoghue, and E.P. Brown. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.*, 93:1071–89, 2005.
- Tom J Wills, Colin Lever, Francesca Cacucci, Neil Burgess, and John O’Keefe. Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876, 2005.
- Kechen Zhang, Iris Ginzburg, Bruce L McNaughton, and Terrence J Sejnowski. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of neurophysiology*, 79(2):1017–1044, 1998.