

MANUSCRIPT

Supplemental information

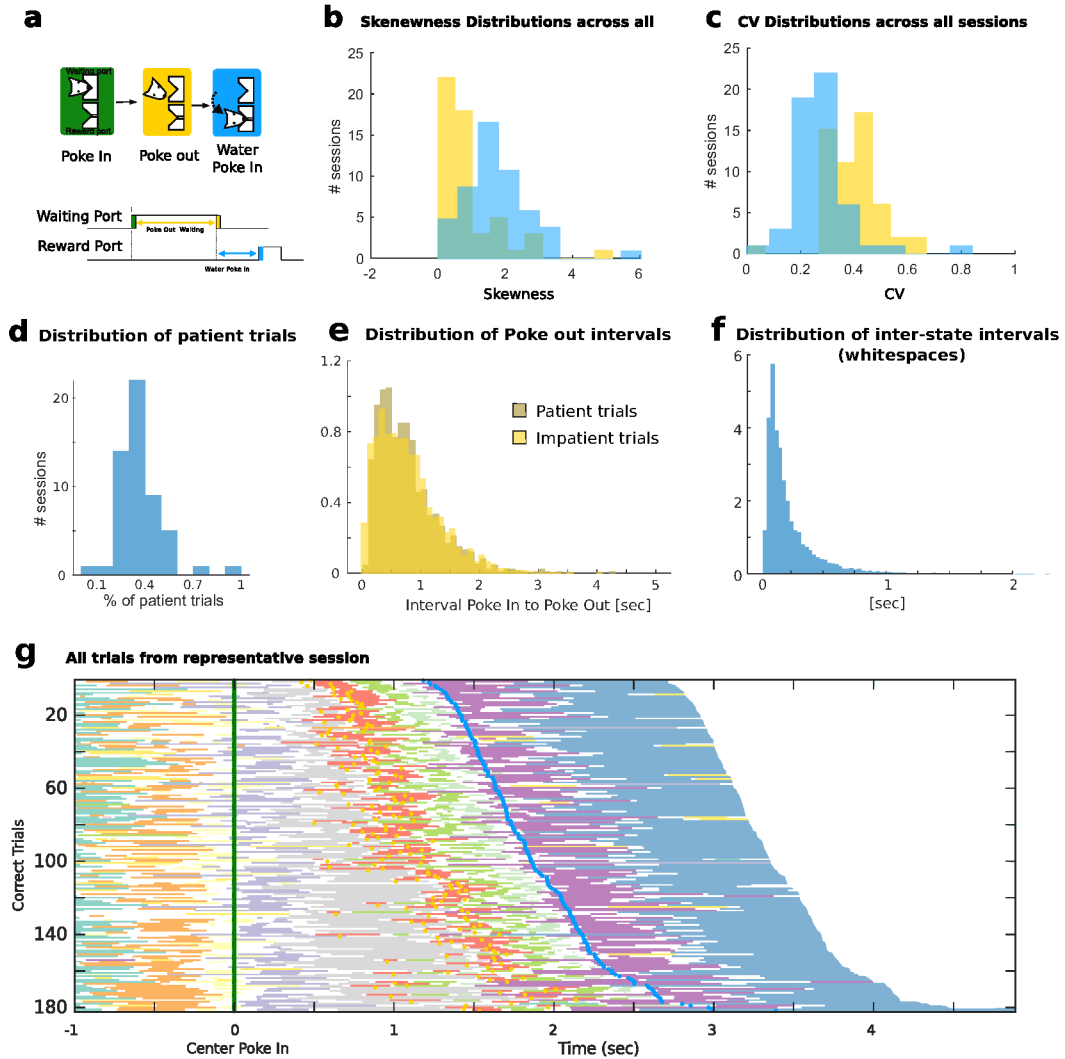


Figure S1: Intervent interval distributions. **a)** Schematic of the interevent time intervals between Poke In and Poke Out (yellow) and between Poke Out and Water Poke In (blue). **b)** Histogram of skewness of interevent interval distributions across all sessions revealed right-skewed distributions. **c)** Histogram of coefficients of variability (CV) of interevent intervals across all sessions revealed large trial-to-trial variability in interevent times. **d)** Percentage of patient trials across all sessions. **e)** Distribution of the interval between Poke in and Poke out for impatient (aligned to tone 1) and patient trials (aligned to the shortest allowed tone 2, 300ms after tone 1). **f)** Distribution of inter-state intervals ("whitespaces"). In between states detected by the HMM analysis there are intervals in between states where the posterior probability of any HMM state is below 80%. These are what we term whitespaces. **g)** Pattern sequences in all trials, ordered by duration (blue dots represent Water Poke In). As opposed to Fig. 1e where trials time courses were stretched, trials in g) represent actual time courses.

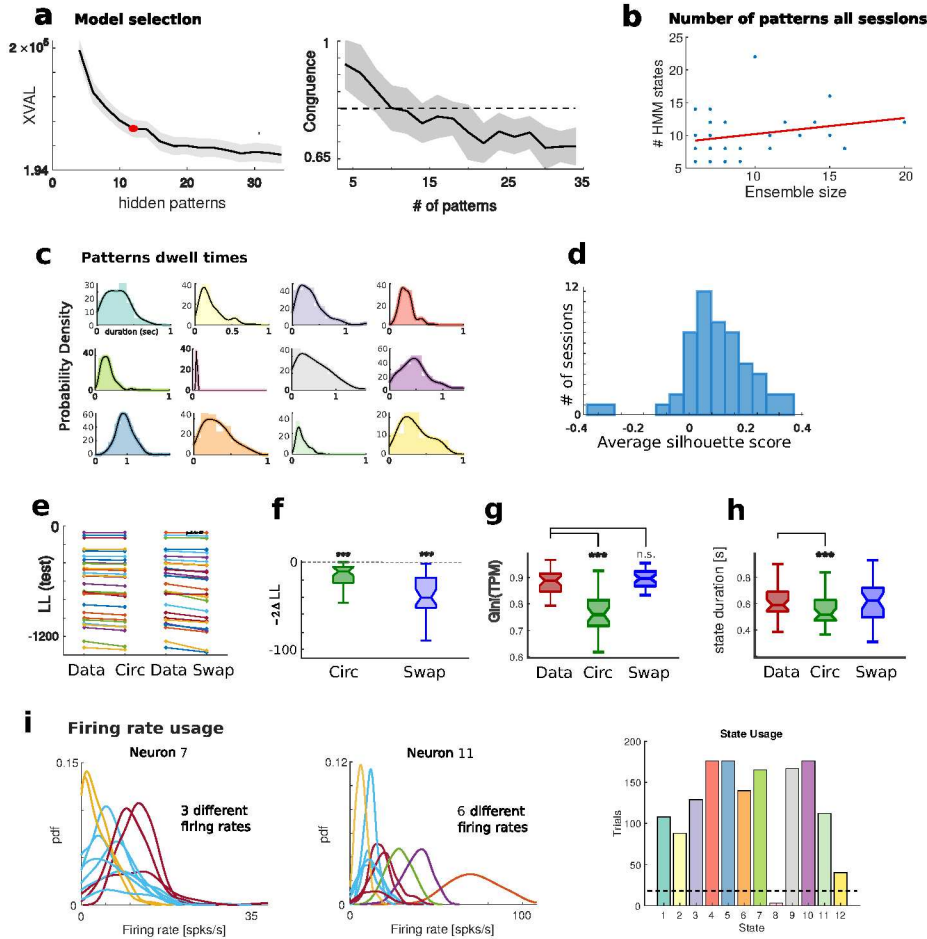


Figure S2: Empirical data fit via hidden Markov models (HMM). **a)** Model selection in a representative session. Left: The number of patterns was selected via an unsupervised 10-fold cross-validation procedure. The number of patterns yielding the largest curvature in the likelihood trend is selected ($XVAL = -2 \times \log$ -likelihood of held-out trials plateaus for increasing number of patterns). Right: For a fixed number of patterns, the similarity between the HMM parameters (cross-validated congruence, see Methods) optimized in different folds drops below 0.8 (dashed line) when the number of patterns grows beyond the value selected via cross-validation (see Williams et al. (2018), Tomasi & Bro (2006)). **b)** The number of patterns across all sessions did not depend on ensemble size (linear regression: $R^2 = 0.07$, $p = 0.12$). **c)** Pattern-specific distribution of dwell times for the representative session in Fig. 1. The average (\pm sd) coefficient of variation is 0.54 ± 0.16 and the average skewness is 0.9926 ± 0.4797 . **d)** Distribution of within- and across-cluster distances between patterns measured with silhouette scores (ranksum test, $p < 2.0 \times 10^{-7}$, see Fig. 3a). **e)** Comparison of log-likelihood of hold-out trials for HMM fit to empirical data compared to circular and swap shuffled datasets shows a significant decrease in both cases (see Fig. 2). **f)** Difference in $-2 \times \log$ -likelihood between HMM fit of empirical dataset and of the 10 surrogate datasets. **g)** Gini coefficient distribution for the pattern transition probability matrices (TPM) across all sessions, compared to shuffled datasets. Empirical TPMs are sparser than TPMs inferred from the circular surrogate dataset. TPMs from the swap surrogate dataset have a similar sparseness to the empirical ones, although the swap surrogate dataset has only two states on average (see Fig. 2, $*** = p < 10^{-6}$). **h)** Average state duration across all sessions, compared to shuffled datasets. **i)** Probability Density Function (pdf) for the firing rate of a single neuron across states. This correspond to the cumulative distribution in Fig. 3c but here individual distribution are colored according to the firing rates that were individuated. Right: number of trials in which each state appears. Dotted-line marks the threshold for considering a firing for that specific state in the analyses of multistability of Fig. 3c.

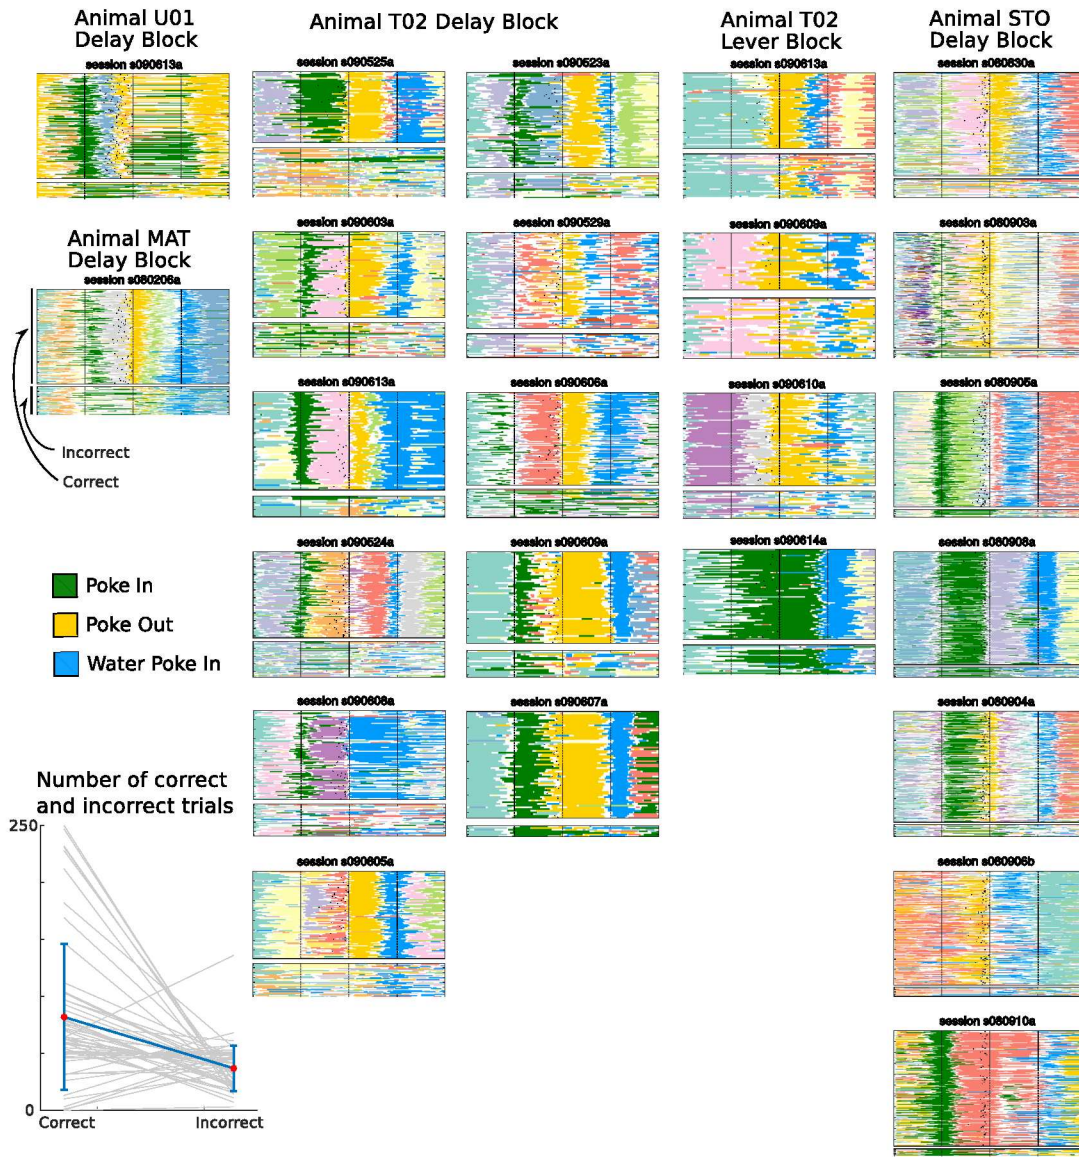


Figure S3: Stretched and event-aligned pattern sequences across sessions. In each session, neural patterns from correct trials (top subplots, same notations as in Fig. 1e) show reliable sequences, as expected from the fixed action sequences to be performed to collect a reward; pattern sequences in incorrect trials are less reliable (bottom subplots, see Fig. 4b), as expected from the inconsistent behavior in those trials. In most sessions, animals only performed the delay task (“Delay block”); in sessions where delay and lever tasks (“Lever block”) were interleaved, block trials from the two tasks were plotted separately (only for animal T02). The patterns tagged to one of the three events analyzed (cf. legend and Fig. 4) are consistently colored across session. The remaining patterns do not follow a consistent color code. Inset in bottom left: Number of correct and incorrect trials in each session. Session s090206a is the representative session in Fig. 1. Only sessions with at least two tagged states are shown.

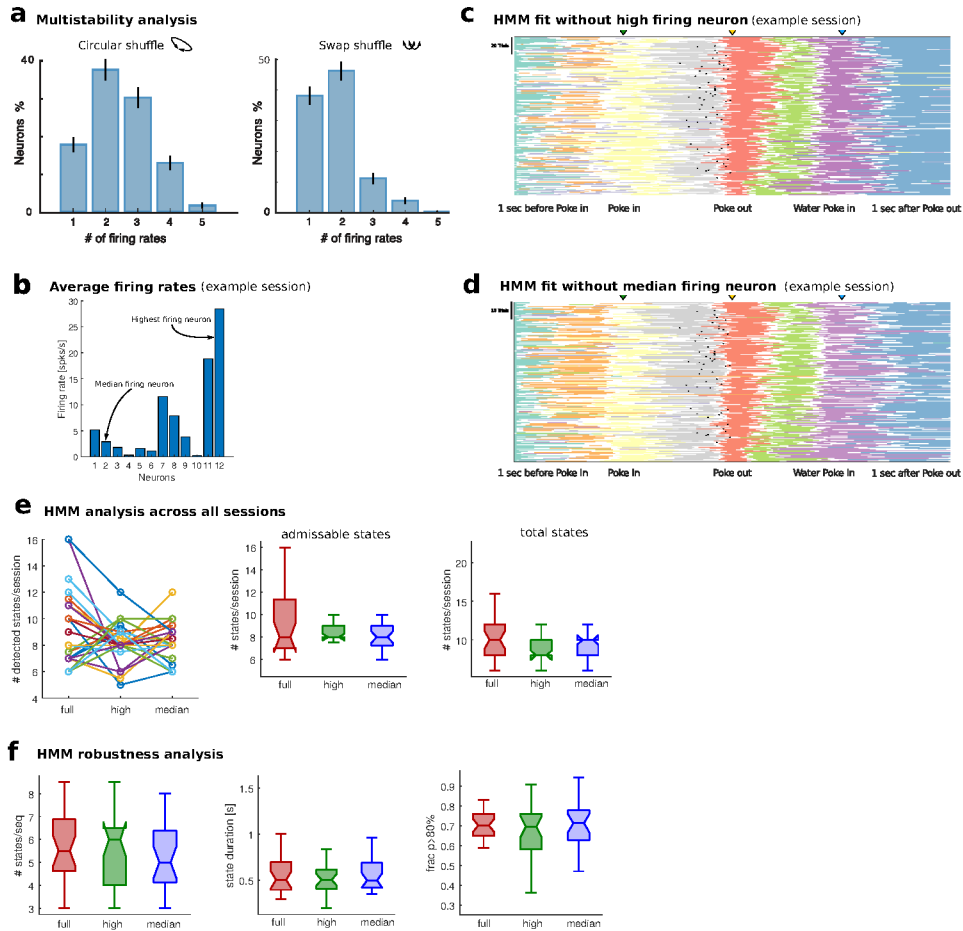


Figure S4: HMM fit to shuffled and subsampled datasets. a) Multistability analysis, same as in Fig. 3c, for shuffled datasets upon applying the HMM analysis. b) We performed a control analysis fitting HMM on two surrogate datasets obtained by removing ensemble neurons with the highest and the median firing rate from each session (same representative session as in Fig. 1). Average firing rate of neurons in example session, same of Fig. 1e. c) HMM sequence of states time-warped to align self-initiated actions across trials for the HMM to the example session without the neuron with highest firing rate. Same as plot Fig. 1e for new HMM fit. d) Same as plot c but with HMM fitted to the neural population activity deprived of the neuron with median firing rate. e) HMM analysis for subsampled neural populations across all sessions. Left: number of states for the HMM fitted to the entire neural population (full) or subsampled of the top firing neuron (high) or median firing neuron (median). Center: Number of states in the three conditions (full, high, median) that used by the HMM fit, pass the conditions of 80% confidence and appear for more than 50 consecutive milliseconds. Right: total number of states in the three conditions. f) We found no significant changes in the number of average number of states per sequence. Left: average number of HMM states per sequence/trial. Center: average state duration. Right: detection confidence defined as the fraction of a trial where states are detected with $p > 80\%$, see Fig. 2.

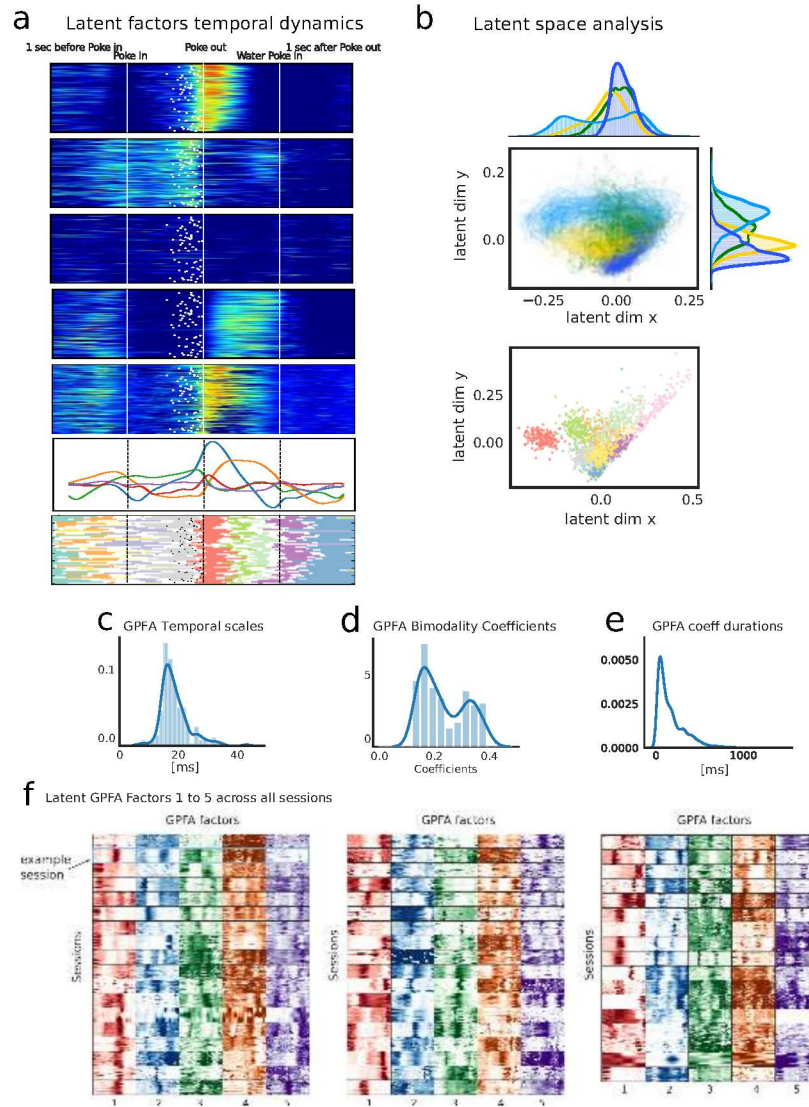


Figure S5: GPFA analysis of M2 neural activity. **a)** Temporal dynamics of the five latent factors for the example session (same as in Fig. 1). The top five panels display the time warped dynamics of each of the five latent factors. The time warping is identical to the one adopted and described in Fig. 1e and each row in each of the five panels corresponds to a trial in the session. The second to last panel shows the average of the five latent factors across all trials in the session. The last panel reproduces Fig. 1e for immediate comparison between the HMM analysis and GPFA fit. **b)** Analyses of the latent space extracted by means of GPFA. Top panel: trajectories of the top two Principal Components of the five latent factors. Temporal intervals in each trial are colored accordingly to the event they lead to (Poke in, Poke out, Water Poke In, Water delivery). On the top and right of the plot the marginal distribution of latent factors are shown projected on PC 1 and 2 are shown. These distributions appear to be multimodal (mixture of distinct distributions) across the four periods individuated, a feature previously linked to the presence of discrete states underlying the latent factor dynamics for a GPFA model Engel et al. (2016). Lower panel: HMM states for each trial (similar to Fig. 3a) projected on the PC space of GPFA latent factors (same as top panel). To visualize the correspondence between the states individuated by the HMM and the GPFA trajectories we projected the HMM states Fig. 3a onto the GPFA latent space. This plot allows to compare the latent spaces extracted by both HMM and GPFA, and the underlying putative influence of discrete states on continuous states and vice versa. **c)** Distribution of time scales of GPFA latent factors across all sessions. **d)** Distribution of bimodality coefficients for GPFA fits to M2 neural activity across all sessions displaying the high number of factors that was strongly bimodal. **e)** Durations of periods in which the first GPFA factor is positive or negative. As the GPFA factor crosses zero in either direction this marks the onset of a new period. The bimodality of latent factors of GPFA fit highlights the separation between periods of sudden change in the temporal evolution of latent factors followed by long approximately constant epochs Engel et al. (2016). GPFA factors, while having fast timescales, appear to switch sign in proximity to HMM state transitions. As a result the duration of temporal intervals where GPFA factors have the same sign are considerably longer than the factor timescale and closer to the timescale of duration of HMM states Fig. 1f. This reveals a clear tension in the GPFA fit to data compared to its underlying assumptions. **f)** GPFA factors across all sessions. For each session (rows) the top 5 time warped GPFA factors are shown. The plots are identical to the ones in Fig. S5a except for the colormap which here is tied to the specific factor. In this synopsis each quadrant corresponds to the full temporal profile across time and trials (similar to Fig. S5a) of a latent factor. This allows the inspection of how many sessions display a characteristic correspondence between a temporal modulation of a unique set of GPFA factors and behavioral events or HMM hidden states. Overall it appears that the onset of specific HMM states, and most often GPFA states, is associated to upcoming actions while in between actions the transition between HMM states appears more rarely corresponding to alternation between GPFA factors.

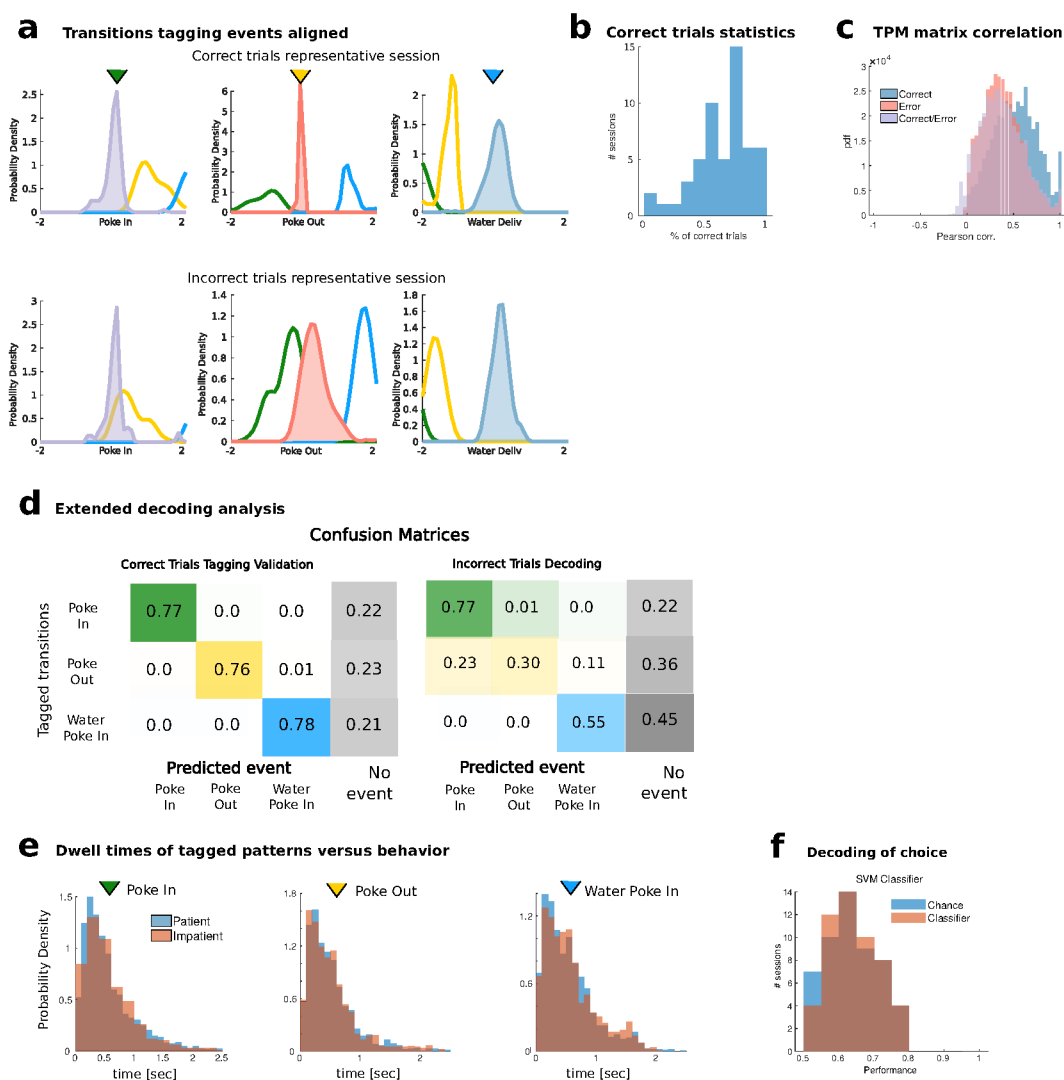


Figure S6: Predicting actions from patterns in incorrect trials. a) After tagging transitions to events using correct trials (top row, left to right: Poke In, Poke Out, Water Poke In; see Fig. 4), transitions were subsequently aligned to the same events but in incorrect trials (bottom row). In each subplot, the distribution of transition times for the corresponding event (e.g. Poke In in the top left: purple-filled histogram shows distribution of “Poke In” transition times aligned to the Poke In events) is compared to events time distribution for other events (yellow and blue curves). b) Percentage of correct trials across all sessions. c) Correlation of TPM matrices between correct and error trials. d) Decoding of actions from tagged state transitions comprising the statistics of trials where the specific event/action was not confused with other events/actions. e) Dwell time distributions of tagged states conditioned on the behavioral condition (patient vs impatient). All distributions were not significantly different. The central plot is similar to Fig. S1e. f) Decoding of patient vs impatient conditions. We computed spikecount vectors binning spikes of neural activity through a non-overlapping moving window of 50ms. In each trial we considered the neural activity from 1sec before Poke In until 1.5sec after Water Poke In. All spikecount vectors were labelled respectively as "patient" or "impatient" if they belonged to a trial where the animal displayed the corresponding behavior. We then computed the performance of an SVM classifier trained to decode the condition from the neural activity across all trial. The distribution of the classifier's performance is compared to chance level (Kolmogorov-Smirnov test p-value=0.87).

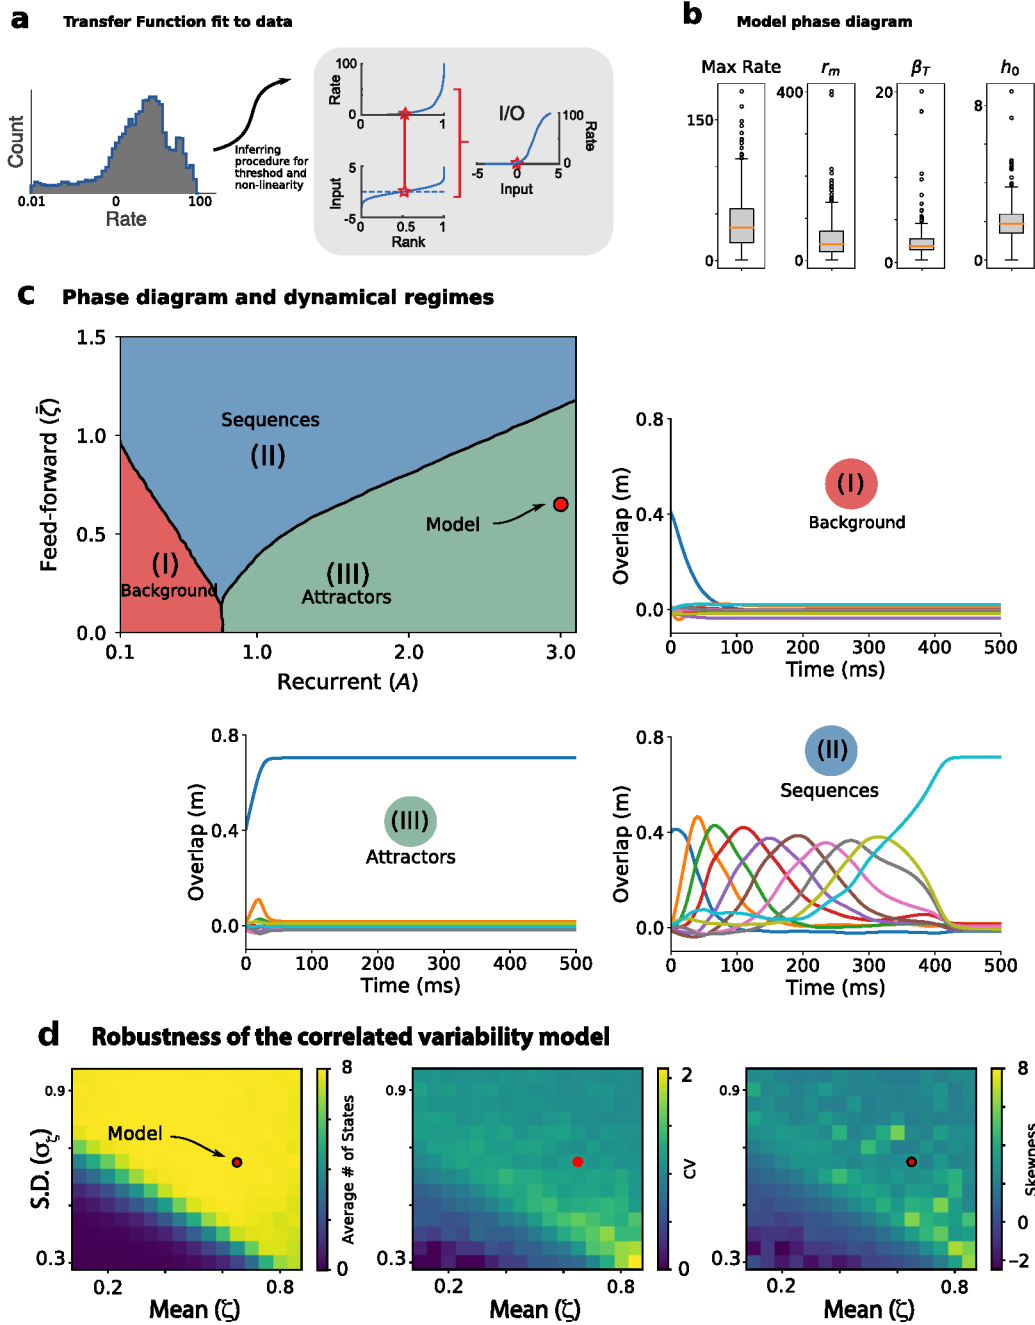


Figure S7: Inferring transfer functions from pattern sequences and model phase structure. *a)* Procedure to infer single-neuron current-to-rate transfer functions from the data. The empirical distribution of firing rates across patterns for a representative neuron (left) was rank-matched to a standardized normal distribution of input currents (top and bottom left panels in grey box), obtaining the current-to-rate function (right). The star in each plot corresponds to the median value. *b)* Each single-cell current-to-rate function was fit to a sigmoidal function, yielding a distribution of fit parameters (328 neurons from 33 sessions; see Methods). *c)* Phase diagram for a recurrent network in the absence of noise (network model with $\sigma_\zeta = 0$; average over 100 network realizations). The number of attractors visited per trial depends on the strength of the average symmetric (x -axis: the parameter A_S represent the strength of J^S) and asymmetric couplings (y -axis: the parameter $\bar{\zeta}$ represents the strength of J^F). The model activity decays to zero for low values of the couplings (phase I, red); it generates stable attractors for large symmetric weights (phase II, green) or sequences of attractors for large asymmetric weights (phase III, blue). In the representative trials, color-coded curves represent the time course of the overlaps between network activity and attractors. The red dot in the phase diagram shows the values of J^S and J^F used in the Results for the correlated variability model (upon adding multiplicative noise for J^F). *d)* Robustness of the correlated variability model. When systematically varying the mean $\bar{\zeta}$ and the standard deviation σ_ζ of the noise $\zeta(t)$ in correlated variability model (corresponding to the value of J^S and J^F from panel *c*), red dot), the network robustly generates long attractor sequences (with an average of $p \approx 8$ attractors in each sequence, left) with large CV (center) and skewness (right) in attractor dwell time distributions over 200 trials.

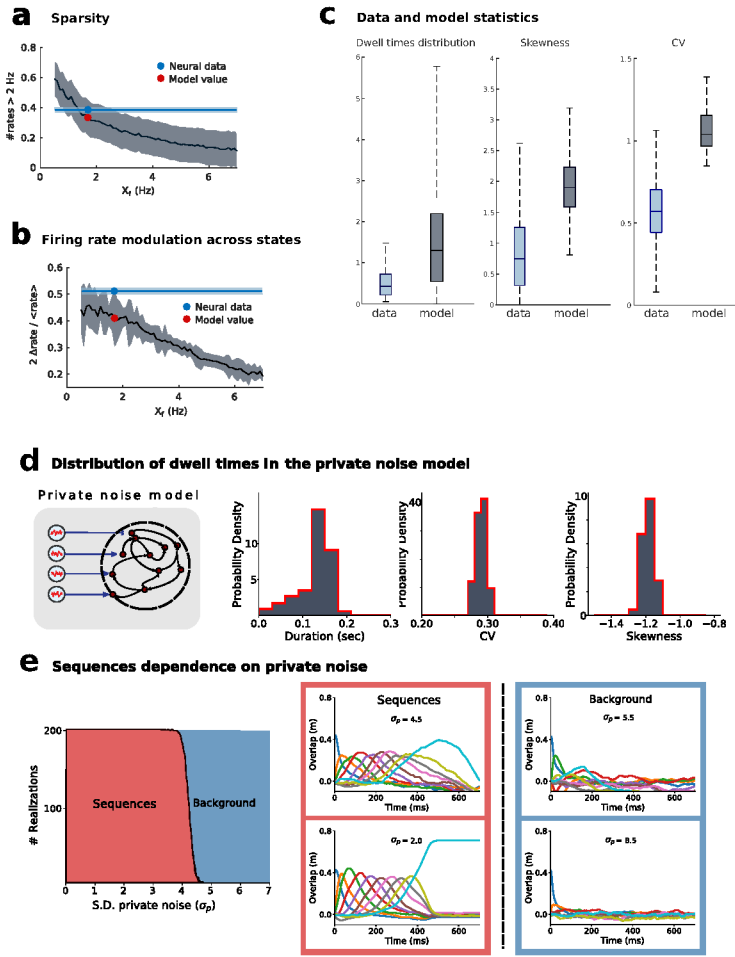


Figure S8: Model analysis and comparison to data. **a)** Sparsity of neural representation as a function of parameter x_f in the model. We computed the percentage of average firing rates for neuron and state that was higher than 2Hz. The parameter x_f regulates the sparseness/denseness of the code in model and the value adopted for simulations of the model in the manuscript $x_f = 1.7$ is compatible with the sparseness of the representation in the data (shaded area represents standard error across sessions). **b)** Variability of firing rates across states. We computed when the absolute difference in firing rates of a neuron between two states i, j , given by $\Delta R_{ij} = |R_i - R_j|$ was larger than half its average firing rate $\langle R \rangle$: $\Delta R_{ij} \geq \langle R \rangle / 2$. This measure is sensitive to the variation of firing rates between states. **c)** Comparison between data and model for distributions of dwell times, skewness and CV. All distributions are significantly different. **d)** In a model with symmetric and asymmetric couplings, adding private noise (see Methods; network parameters as in Fig. 5, except for $A_S = 1$, $\bar{\zeta} = 0.65$, $\sigma_\zeta = 0$, $\sigma_p = 2$) introduces a small amount of trial-to-trial variability in pattern dwell times distributions (left), yielding small non-zero CVs (center) and negative skewness (right) across different networks. The dwell times statistics are qualitatively different from the empirical ones (cf. Fig. 1f). **e)** Left: Beyond a critical value of noise strength σ_p , the private noise model breaks down. For low private noise, networks generate sequential dynamics (red area, networks with sequential dynamics out of 200 network realizations), though with low CV and skewness; for larger values of private noise, activity decays to zero (blue). Right: Representative trials with 4 different values of the noise strength σ_p .

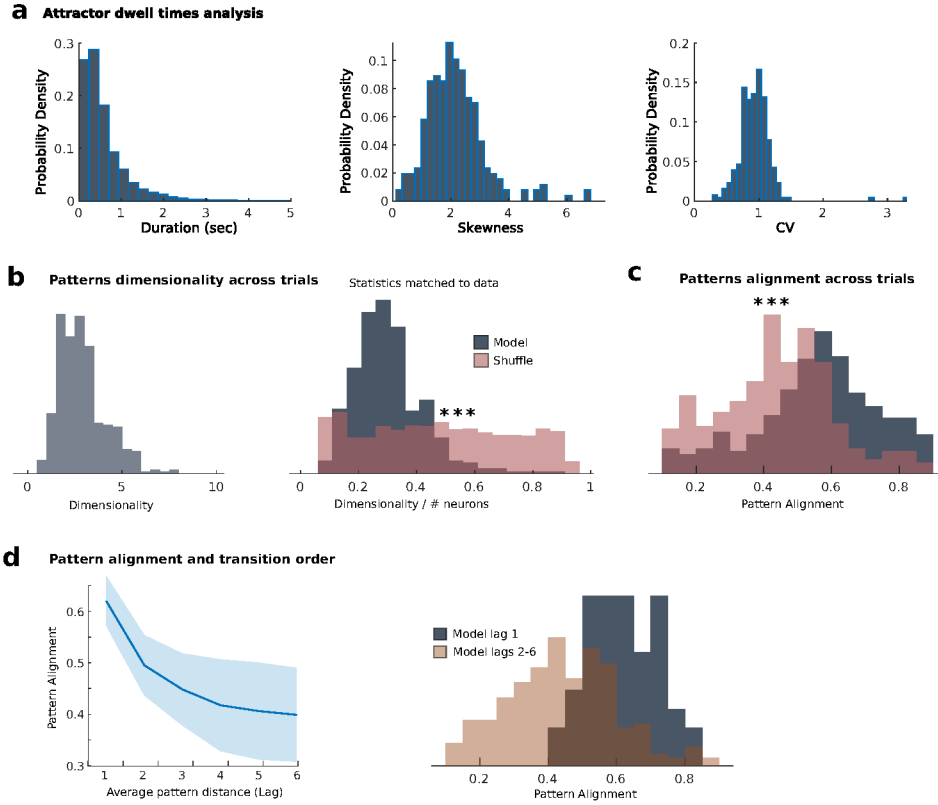


Figure S9: Data analyses on simulated neural activity in the model via HMM fits. **a)** Analysis of dwell times in the model with states inferred by a HMM fit to the model simulated neural traces. This parallels Fig. 5c. **b)** Analysis of dimensionality and pattern alignment in the model with states inferred via HMM fit to the model simulated neural traces. This parallels Fig. 6a. **c)** Same as panel b for the analysis of pattern alignment. This parallels Fig. 6b. **d)** for the analysis of pattern alignment. This parallels Fig. 6c.

Comparison with previous models		
Model	Features	Dynamics
(Abeles 1991, Diesmann et al. 1999)	Synfire chains.	Reliable sequences; pattern dwell times have short duration and low variability. Fine tuned connectivity structure.
(Sompolinsky & Kanter 1986, Kleinfeld 1986, Dehaene et al. 1987, Treves 2005, Murray et al. 2017, Gillett et al. 2020)	Rate networks with both symmetric and asymmetric weights.	Reliable pattern sequences but low trial-to-trial temporal variability.
(Jun & Jin 2007, Liu & Buonomano 2009, Fiete et al. 2010, Pereira & Brunel 2020)	Rate networks trained with unsupervised learning rules.	Reliable pattern sequences but low trial-to-trial temporal variability.
(Litwin-Kumar & Doiron 2012, Deco & Hugues 2012, Mazzucato et al. 2015)	Clustered E-I spiking network generating attractors.	Random sequences with large temporal variability.
(Mongillo et al. 2003, Miller & Katz 2010)	Clustered E-I spiking network generating attractors.	Reliable transitions between attractors with large temporal variability. Architecture is fine tuned to produce a small number of carefully designed metastable attractors.
(Mazzucato et al. 2015, 2016, 2019, Darshan et al. 2017)	Clustered E-I spiking network generating attractors.	Reliable sequences of attractors with large temporal variability. The network is driven by strong time-dependent stimuli designed to pace activity along stimulus-specific sequences.

Table S1: The main features of M2 ensemble activity explained by our model were the reliable identity and order of long-lived neural patterns occurring in a sequence, and the large trial-to-trial variability of pattern dwell times. In our model, both features can be robustly attained when transitions between attractors arise from low-dimensional correlated variability. Previous network models could achieve either sequence reliability or variability in dwell time distributions, but not both; with a notable exception of the model in (Miller & Katz 2010) where both sequence reliability and temporal variability were achieved. Although, the underlying mechanism in this model relies on an architecture that is tuned to produce the small number of carefully designed states. Models generating reliable pattern sequences include synfire chains (Abeles 1991, Diesmann et al. 1999). These models rely on a fine tuned connectivity structure producing pattern dwell times with short duration and low variability. While these dynamics are well suited to explain neural activity observed in songbird HVC (Hahnloser et al. 2002, Fiete et al. 2010) or mammalian hippocampus (Nádasdy et al. 1999), their features are not compatible with the observed M2 ensemble activity. Reliable pattern sequences can otherwise be triggered by specific cues in recurrent networks with asymmetric connectivity structure (Sompolinsky & Kanter 1986, Kleinfeld 1986, Dehaene et al. 1987, Treves 2005, Murray et al. 2017, Gillett et al. 2020), trained with unsupervised learning rules (Jun & Jin 2007, Liu & Buonomano 2009, Fiete et al. 2010, Pereira & Brunel 2020) or in reservoir networks (Rajan et al. 2016). However, pattern dwell times in such models are short, set by single-neuron characteristic time constants, and show little trial-to-trial variability. Pattern dwell times could be increased via synaptic delays (Sompolinsky & Kanter 1986). However, none of these models is capable of generating large trial-to-trial variability in dwell time distributions and are thus incompatible with the observed M2 data. Spiking networks with clustered architecture can give rise to metastable attractors with large trial-to-trial variability in dwell time distributions (Litwin-Kumar & Doiron 2012, Deco & Hugues 2012). However, metastable attractors in these models concatenated in random sequences, incompatible with the highly reliable sequences we observed in M2. Each cluster generates independent fluctuations within activity space, realizing a high-dimensional stochastic process. These fluctuations drive transitions along random directions in activity space, thus unreliable across trials (when concatenating more than two states in a sequence (Mongillo et al. 2003)). To drive a specific transition, independent fluctuations would have to align along a specific direction within the high-dimensional activity space, and the probability of this event occurring vanishes for large network size. One may drive clustered networks with strong time-dependent stimuli to pace activity along stimulus-specific sequences (Miller & Katz 2010, Mazzucato et al. 2015, 2016, 2019, Darshan et al. 2017). However, this would merely shift the problem of reliable sequence generation from the local circuit to an upstream area producing the specific input.