Brain dynamics and spatiotemporal trajectories during threat processing

Joyneel Misra[1] and Luiz Pessoa[1,2]

[1]Departmentof Electrical and Computer Engineering, University of Maryland, College Park, Maryland, United States of America

[2]Department of Psychology and Maryland Neuroimaging Center, University of Maryland, College Park, Maryland, United States of America

Abstract

In the past decades, functional MRI research has investigated mental states and their brain bases in largely static fashion based on evoked responses during blocked and event-related designs. Despite some progress in naturalistic designs, our understanding of threat processing remains largely limited to those obtained with standard paradigms. In the present paper, we applied Switching Linear Dynamical Systems to uncover the dynamics of threat processing during a continuous threat-of-shock paradigm. Importantly, unlike studies in systems neuroscience that frequently assume that systems are decoupled from external inputs, we characterized both endogenous and exogenous contributions to dynamics. First, we demonstrated that the SLDS model learned the regularities of the experimental paradigm, such that states and state transitions estimated from fMRI time series data from 85 ROIs reflected both the proximity of the circles and their direction (approach vs. retreat). After establishing that the model captured key properties of threat-related processing, we characterized the dynamics of the states and their transitions. The results revealed that threat processing can profitably be viewed in terms of dynamic multivariate patterns whose trajectories are a combination of intrinsic and extrinsic factors that jointly determine how the brain temporally evolves during dynamic threat. We propose that viewing threat processing through the lens of dynamical systems offers important avenues to uncover properties of the dynamics of threat that are not unveiled with standard experimental designs and analyses.

1 Introduction

Historically, systems neuroscience has sought to develop experimental paradigms where an effect of interest is varied while as many other variables remain fixed in an attempt to isolate the impact of the former while controlling for the latter. Such time-tested strategy has provided deep insights into brain and behavior. However, with the advent of neurotechnologies allowing the study of animals in more freely-behaving conditions, more naturalistic paradigms have been introduced.[1,2] At the same time, in human neuroimaging, researchers have defended the idea that dynamic and/or naturalistic paradigms are needed to advance our understanding of brain and behavior in a way that is not entirely constrained by rigid paradigms.[3–7] While such ideas resonate with those from ecological psychology and ethological approaches,[8–10] they introduce considerable challenges in the analysis of brain data and associated behaviors. If there are no specific trials or blocks in an experiment, what should constitute the unit of analysis? If the paradigm is more open-ended and lacks specific conditions (e.g., "attend to the visual stimulus" vs. "attend to the audio stimulus"), how should brain and behavior be investigated?

A promising approach is the use of unsupervised methods that consider time series dynamics.[11–24] In particular, Hidden Markov Models (HMMs) can be applied to time series data and partition it into "states" in time. For example, in one study, continuous narratives were presented and state transitions were identified from fMRI signals.[13] Notably, state boundaries corresponded to human-annotated ones better than chance, and the match between model and human boundaries increased for more "high-level" brain regions (e.g., angular gyrus). Another study attempted to determine mental states while participants solved novel mathematical problems.[11] Model states were interpreted as reflecting "encoding", "planning", "solving", and "responding" stages based on fMRI data. Although the authors found evidence for the detected states, the study also highlights some of

the challenges of applying unsupervised methods to complex paradigms. The estimated duration of the final responding state correlated well with the measured time participants spent entering their responses (r=0.53), providing some validation of that state. However, validation of the other mental states remained challenging.[11]

Overall, despite encouraging results,[11–13,15,19,21,25–27] the extent to which state-space models successfully parse fMRI data into meaningful mental/brain states remains poorly understood. Critically, while state-space models like HMM are able to segment time series data, they do not enable discovery of properties of state dynamics. In particular, how do brain signals evolve during a brain state that is maintained for a period of time? In systems neuroscience progress has been made in this context by studying brain signals via dynamical systems theory.[28–32] In particular, it has been suggested that attractor dynamics plays important computational functions.[33,34] However, comparable progress in fMRI research has lagged behind, perhaps because of the slowly evolving nature of the hemodynamic response (but see[35]).

Threat processing is typically studied in terms of discrete events and/or blocks that correspond to experimental conditions (but see[36–39]). For example, a conditioned threat is presented and associated evoked brain responses determined.[40] But in many real-world situations, conditions involving threat unfold temporally in a less-discrete manner.[41] In the present investigation, we investigated brain dynamics during a continuous threat-of-shock paradigm[42,43] (Fig. 1). We sought to characterize the evolution of brain signals in terms of trajectories in multidimensional space, a framework at times called "computing with trajectories".[44] Instead of responses to discrete events possibly observed across multiple brain sites, we determined multivariate and distributed patterns of activity with shared dynamics.

We employed a Switching Linear Dynamical Systems (SLDS) framework that estimates a generative model of the data with both states and transitions like HMMs.[45–48] In addition, each state is described in terms of a linear dynamical system that specifies the trajectory of the brain during the state. Importantly, unlike studies in systems neuroscience that frequently assume that systems do not receive external inputs (see discussion in[34]), we estimated both endogenous and exogenous components of the dynamics. Our approach demonstrated that we can recover key properties of threat-related processing in the brain,[49] while establishing novel properties of threat and safety dynamics. More generally, the approach can be applied to a broad range of continuous and/or naturalistic fMRI paradigms investigated in humans.

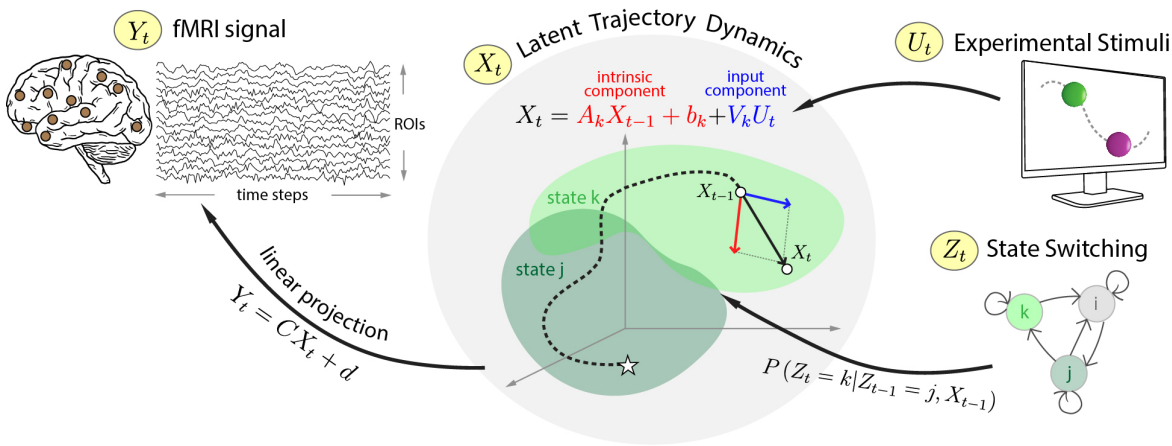

Figure 1: Determining shared dynamics during threat processing. Switching linear dynamical systems (SLDS) were used to model fMRI time series data (Y_t) from a set of brain regions of interest (ROIs). The framework assumes that time series data can be segmented into a set of discrete states. The model represents brain signals in terms of a set of latent variables (X_t). For each state k the temporal evolution of the system is specified via a linear dynamical system with both intrinsic and input-related components. In the diagram, the system starts in state j (white star) and transitions to state k (the colored patches in the middle represent the subspaces associated with the two states). In state k, the system evolves according to the dynamics matrix A_k and input contributions (V_k). Overall, as states switch temporally, so does the corresponding linear dynamical system governing the system's trajectory.

2 Results

2.1 Continuous threat processing switches between states

Participants watched two circles of different colors on the screen, at times moving close to each other, at times moving apart in a smooth but unpredictable fashion[42,43] . Upon circle collision, the participants received mild but unpleasant electrical stimulation together with an aversive sound. To enhance unpredictability and hence anxious apprehension by participants, the movement of the circles included multiple instances of "near misses" in which the circles nearly touched before retreating from each other.

Time series data were analyzed according to the SLDS framework (Fig. 1). Briefly, an SLDS estimates a model of the data using a set of unobserved (also called latent) variables assumed to govern system dynamics. The latent space is typically of considerably smaller dimension relative to the original data space (here 10 latent dimensions and 85 ROI time series, respectively). SLDSs consist of one linear dynamical system model per state, such that the trajectory of the system is governed by a single dynamical systems at at time. Like HMMs, an SLDS segments the data into a fixed number of states, such that the transitions between states follows a Markov process. At a given time, a single state is considered to be "active". Overall, SLDS models provide a piecewise linear approximation of (potentially nonlinear) dynamics. Time series data were obtained from 85 regions involved in threat-related processing, as defined in a prior study.[50]

The results reported below were obtained with 92 participants. A separate group of 30 participants was used to define hyperparameters, including the number of states ($K = 6$, see supplementary methods and Fig. S9)) and the dimensionality of the latent space ($D = 10$, see supplementary methods and Fig. S8). Assessing statistical significance when complex models are applied to multi-participant data requires estimating variability of the estimates across participants. We employed a bootstrap procedure such that an SLDS model was estimated for each bootstrap sample, allowing us to estimate how SLDS parameters varied across the population (see Methods), thus allowing generalization beyond the current sample.

2.1.1 Model states are synchronous with the experimental paradigm

We evaluated if the states estimated by the model reflected experimental stimuli. For example, when the circles were far from each other, participants were "safe", in contrast to when the circles were rather close and could progress to a collision. We encoded the input based on 20 bins defined in terms of circle proximity and direction: A_1-A_{10} for approach, R_1-R_{10} for retreat (see Methods, top inset in Fig. 2). When the circles approached each other, A_1 indicated the farthest distance and A_{10} indicated that they collided; when they retreated from each other, R_{10} indicated that the circles collided and R_1 indicated the farthest distance. We determined the associations between brain states and stimulus categories (i.e., bins) by calculating the probability of a brain state given the stimulus, $P(\text{state} \mid \text{stimulus})$ (see Methods). The probabilities are shown in Fig. 2A (e.g., $P(\textsc{state}2 \mid A_{10}) = 0.63$), where values exceeding those expected by chance are indicated in red outline. A total of five out of six brain states were significantly associated with specific stimuli categories. Examination of Fig. 2A can be used to provide a qualitative description of the estimated brain states. For example, STATE2 was maximally observed when the circles were approaching and then collided, and thus reflects shock delivery (as well as peri-shock periods). As another example, STATE5 was observed when the circles were in close proximity but did not touch. This state was particularly interesting given that the experiment was designed such that multiple "near miss" events would occur during the session, so as to enhance participants' state of apprehension when the circles approached each other.

We found that brain states were observed under specific stimulus conditions. But as the stimulus conditions change, do states transition in a systematic fashion? To evaluate this question, we determined the probability of a state transition given a specific stimulus category (Fig. 2B, e.g., $P([\textsc{state}2 \mapsto \textsc{state}1] \mid A_{10}) = 0.88$; see Methods). Fig. 2B shows state transitions that were significantly associated with a stimulus category. For example, STATE2 (centered on the collision event) transitioned to STATE1 ("post-shock") with very high probability for both stimulus categories A_{10} and R_{10} (i.e., when the circles collided). In Fig. 2B, to facilitate understanding, we ordered the rows by starting with the STATE1 to STATE3 transition (from "post-shock" to "not close") and then chained the transitions in a manner that reflected the movement of the circles. Thus, the second row shows STATE3 $\mapsto$ STATE5 (from "not close" to "near miss"), and so on. For ease of reference, we summarize state descriptions here:

- STATE1: "post-shock". Observed right after the circles collided; only state that followed STATE2 in Fig. 2B.

- STATE2: "shock/peri-shock". Observed during shock and peri-shock inputs.
- STATE3: "not near". Observed when the circles were not near.
- STATE4: "not near". Observed when the circles were not near and similar to STATE3 in terms of inputs; STATE4 followed STATE5 when, after a near-miss collision, the circles continued to move apart from each other.
- STATE5: "near miss". Observed when circles approached, came close to colliding, and then retreated.

Together, brain states and state transitions were systematically related to external stimuli, demonstrating that the model captured important properties of the experimental paradigm in an unsupervised manner.

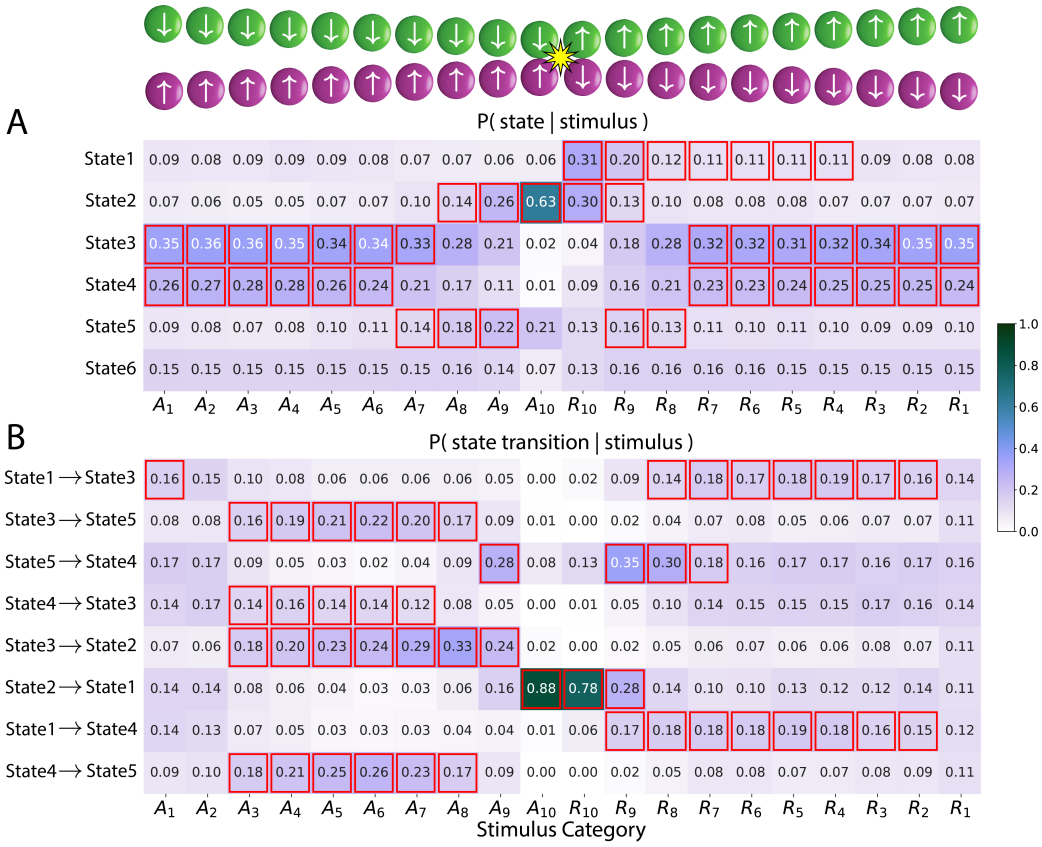

Figure 2: Brain states, state transitions, and input stimuli. Stimuli were categorized into 20 bins based on circle distance and movement direction: A_1-A_{10} for approach, R_1-R_{10} for retreat (1 indicates circles are farthest and 10 indicates circle collision). **(A)** Table entries indicate the probability of being in a state given the input stimulus category. **(B)** Table entries indicate the probability of a state transition given the input stimulus category. In both tables, cells highlighted in red indicate states/transitions significantly associated with the corresponding stimulus category ($p < 0.05$, corrected for multiple comparisons). The color scale indicates probability.

2.1.2 Model states capture threat processing

We reasoned that if states detected by the model reflect known mental states, they should be associated with brain activity with coherent spatial patterns consistent with those observed with standard fMRI designs (such as block and event-related designs). SLDS states and state transitions are defined in a latent space of dimensionality considerably lower than the original brain data space (here 85 brain regions). Thus, to determine brain maps associated with specific states, we performed linear regression with a regressor that was ON during the state in question and OFF otherwise (see Methods). Fig. 3 illustrates some of the states (see Fig. S11 for all states). For STATE5, consistent with the "near miss" interpretation, mutiple brain areas linked to increased

anxious apprehension were positively engaged, including the anterior insula and cingulate cortex.[37,38,42,49–54] STATE4, which tended to follow STATE5, exhibited a deactivation of these regions, consistent with the notion that participants experienced relative safety following a near collision.[49,53,55]

We further investigated the spatial properties of brain states by contrasting states. For example, the contrast of STATE5 vs. STATE4 revealed a pattern of activation very similar to that obtained in the literature for the contrast of higher vs. lower levels of threat. In particular, the contrast identified sectors of ventromedial PFC and posterior cingulate that have been linked to relative safety.[49,55]

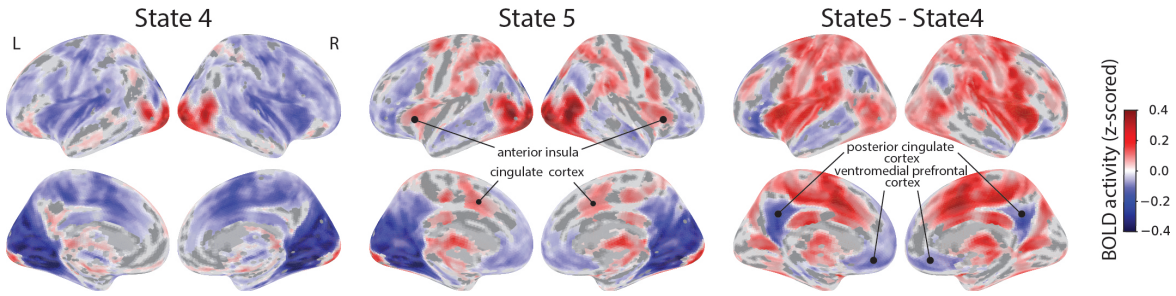

Figure 3: Voxelwise state activity maps and state contrast maps. State activity maps for STATE4 and STATE5 and the contrast of the two states. Maps were corrected for multiple comparisons by thresholding voxels at $p < 0.001$ and at the cluster level at $p < 0.05$.

2.1.3 Temporal evolution of activity

In the threat paradigm, as the circles moved continuously, we expected that brain signals would also vary temporally. If SLDS states captured meaningful information, state transitions should be associated with relevant activity change. To evaluate this possibility, we determined signal evolution during temporal windows centered around state transitions (see Methods). Fig. 4 shows signal evolution in the same order as the state transitions of Fig. 2B in two key regions involved in threat-related processing, the dorsal anterior insula and the ventromedial PFC. For example, the transition from STATE2 to STATE1 reflected transitioning from a shock event to post-shock period. Accordingly, activity in the dorsal anterior insula, which is very sensitive to threat level,[37,38,42,49–54] underwent a drastic decrease. Likewise, when STATE5 transitioned to STATE4, activity in the dorsal anterior insula decreased, consistent with a transition from higher to lower threat and findings about responses in these two regions. In contrast, the transition from STATE3 ("not close") to STATE2 (peri-shock/shock) was associated with a sharp increase in activity strength in the dorsal anterior insula.

SLDS models states and state transitions in a latent space, which makes it possible to investigate activity dynamics (as in Fig. 4) in regions not used for estimating the models, because once a sequence of states is determined one can examine activity in such regions at those times. Furthermore, activity can be determined at different spatial units, including ROI, voxel, or voxel cluster. Here, we probed average signals from two resting-state networks that are engaged during threat-related processing, the salience network which is particularly engaged during higher threat, and the default network which is engaged during conditions of relative safety. As anticipated, signal evolution during the transition STATE5 $\mapsto$ STATE4 was similar to that observed for the dorsal anterior insula and the ventromedial PFC, although signal evolution for the default network deviated to some extent from that of the ventromedial PFC (note that the ROIs and networks discussed here spatially overlapped; they were not selected to be independent).

Overall, the model captured regularities of the experimental paradigm reflected in systematic changes in activity levels across the brain. For example, state transitions STATE3 $\mapsto$ STATE5, STATE4 $\mapsto$ STATE3, and STATE4 $\mapsto$ STATE5 were associated with approaching circles (increased threat) and increased responses in the dorsal anterior insula and salience network following the state transition. Similarly, the transition STATE5 $\mapsto$ STATE4 was associated with retreating circles (decreased threat) and increased responses in the ventromedial PFC and default network following the state transition. Furthermore, responses of the right anterior insula and salience network during STATE2 ("shock/peri-shock") ramped up to the highest activity level with the STATE3 $\mapsto$ STATE2 transition from low to very high threat followed by a sharp decay with the STATE2 $\mapsto$ STATE1 transition that reflected the end of shock delivery.

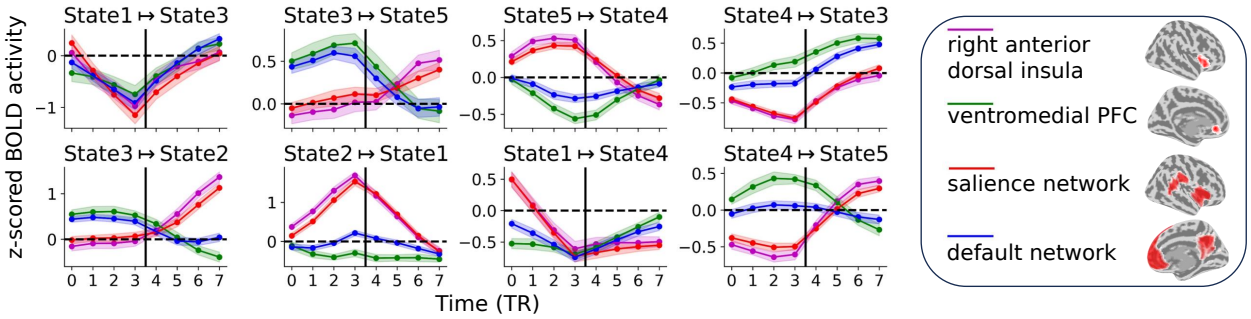

Figure 4: Region and network responses during state transitions. Vertical lines indicate the time of state transition. State transitions are arranged in a chained manner: STATE1 ↦ STATE3 ↦ STATE5 ↦ STATE4 ↦ STATE3 ↦ STATE2 ↦ STATE1 ↦ STATE4 ↦ STATE5 to facilitate continuity in reading responses across state transitions. Error bars correspond to the 95% confidence interval based on the standard error of the mean across participants.

2.2 Threat dynamics

Each state is characterized by a state-specific trajectory in latent space. Such state dynamics are defined by an intrinsic component and an external-input component, which additively define the system's evolution (see Fig. 1 and Methods). In this section, we characterized the dynamics of threat processing based on state trajectories.

2.2.1 Intrinsic state dynamics evolve towards fixed-point attractors

First, we examined *intrinsic* state dynamics, namely the component of the dynamics that is separate from external input contributions. Does intrinsic dynamics converge to an attractor or does it diverge away from a fixed point? It is possible to test the stability of dynamics by evaluating properties of the dynamics matrix (matrix A; see Methods). We found that the intrinsic dynamics of all states studied were attractors (see supplementary Fig. S12).

An attractor state corresponds to a point in latent space. Such vector can be visualized as a map, which we call an "attractor spatial map" by projecting it onto the original brain data space in terms of ROI activity on a brain surface (Fig. 5; see Methods). For example, STATE5 corresponded to "near miss" events and an attractor map with increased responses in the anterior insula/cingulate cortex and decreased responses in the ventromedial PFC and precuneus. The attractor for STATE4 ("not near") exhibited an opposite organization consistent with the context of increased safety/lower threat.

Given that the SLDS model learns a generative model of the fMRI data during the experimental paradigm, state trajectories are based on both endogenous and exogenous contributions. How do these trajectories evolve relative to the state attractor? To investigate this question, we determined the vector field defined by the state's endogenous dynamics, where vectors represent the direction and magnitude of how a system evolves if started at specific points (Fig. 6). Starting at a given point, trajectories evolved for multiple time steps in the overall direction of the attractor but after some time transitioned to other states (the green part of the trajectory indicates the mean lifetime of each state). The plots describe how, in a given state, the effective trajectory was determined by a combination of the endogenous flow and perturbations provided by the input. For example, in STATE2 (shock/peri-shock) the vector field exhibited less curvature than several other states but trajectories were quickly steered away from the attractor because of inputs.

2.2.2 Exogenous contributions to state trajectories

Following the qualitative description above (Fig. 6), we next quantified the contributions of external inputs in steering system trajectories. For each state, we determined the state's centroid, namely, the centroid coordinate of all locations visited by the system during a specific state. Such geometrical location provided a reference position to investigate the input contributions. We considered the state centroid as reference location because we sought to determine how the inputs steered trajectories relative to the "most typical" position of a state.

We considered two scenarios. First, when the system was in state i, the input could drive the trajectory towards state i's centroid or away from it (Fig. 7A). The "effective contribution" was determined by considering

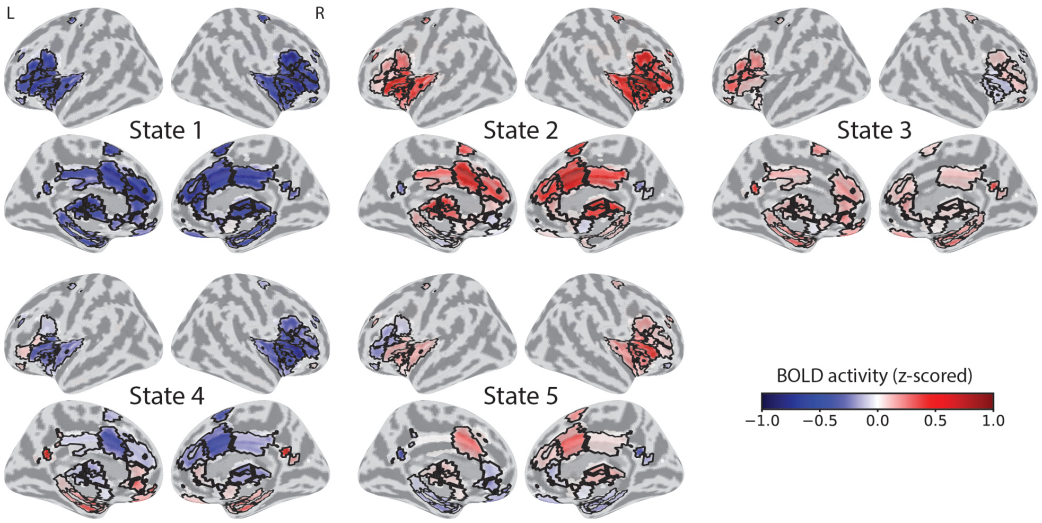

Figure 5: Attractor maps. State attractors projected onto the space of regions of interest (85 ROIs) visualized on a brain surface map. At each ROI, the color scale represents activity strength at the attrator's fixed point. ROI boundaries are marked in black.

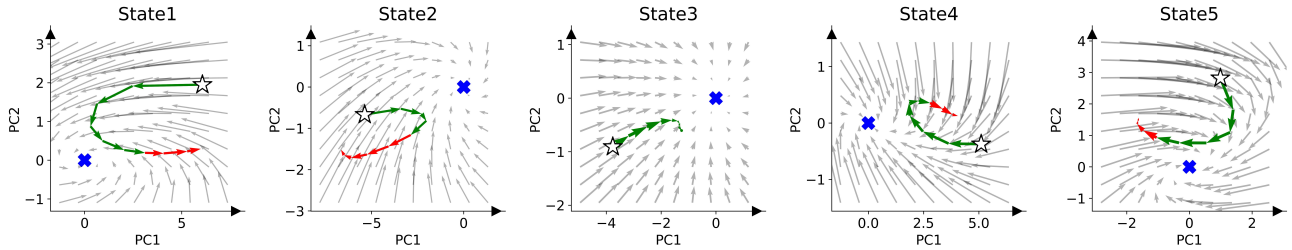

Figure 6: Evolution of state-specific trajectories. Trajectories were determined in the latent space and projected onto a two-dimensional vector field for illustration (coordinate axes are specific to each state). The average trajectory across participants, starting at the white star, is initially shown in green and switches to red when the majority of trajectories (across participants) switch to another state. Whereas the trajectory includes endogenous and exogenous contributions to temporal evolution, the vector field represents the endogenous contribution only. The gray arrows indicate the effect of the state's endogenous dynamics matrix, showing the direction and magnitude of evolution after a single time step. The blue cross indicates the state's fixed point attractor; star indicates the start of the trajectory. PC: principal component.

both direction and magnitude components (see Methods). We investigated the stimulus categories defined previously (A_1—A_{10} and R_{10}—R_1). Consider STATE1 (post-shock), for example (Fig. 7B). When the stimulus was R_{10} or R_9 those inputs tended to push the system in the direction of STATE1's centroid; when the stimulus was R_8, the trajectory was pushed away from the centroid by the input but very weakly. In addition, for STATE2 (shock/peri-shock), we found that stimuli A_8 to R_{10} pushed the system toward the centroid, but stimulis R_9 moved the trajectory away from it. Overall, Fig. 7B displays how inputs steered state trajectories in the direction of the centroid (green cells) or away from it (purple cells).

Next, we investigated how inputs contributed to state transitions by steering trajectories in the direction of a new state's centroid (see Methods). Consider the STATE 2 $\mapsto$ STATE 1 transition (i.e., from shock to post-shock). Input perturbations R_{10} and R_9 provided the strongest push of the trajectory in the direction of STATE1. In addition, when the system was in STATE3 ("not near"), multiple inputs tended to direct the system's evolution in the direction of STATE2 ("shock/peri-shock"). Finally, when the system was in STATE4 ("not close") and the circles were approaching, inputs pushed the system in the direction of STATE5 ("near miss").

Taken together, the results show how external perturbations contributed to the evolution of state trajectories and state transitions.

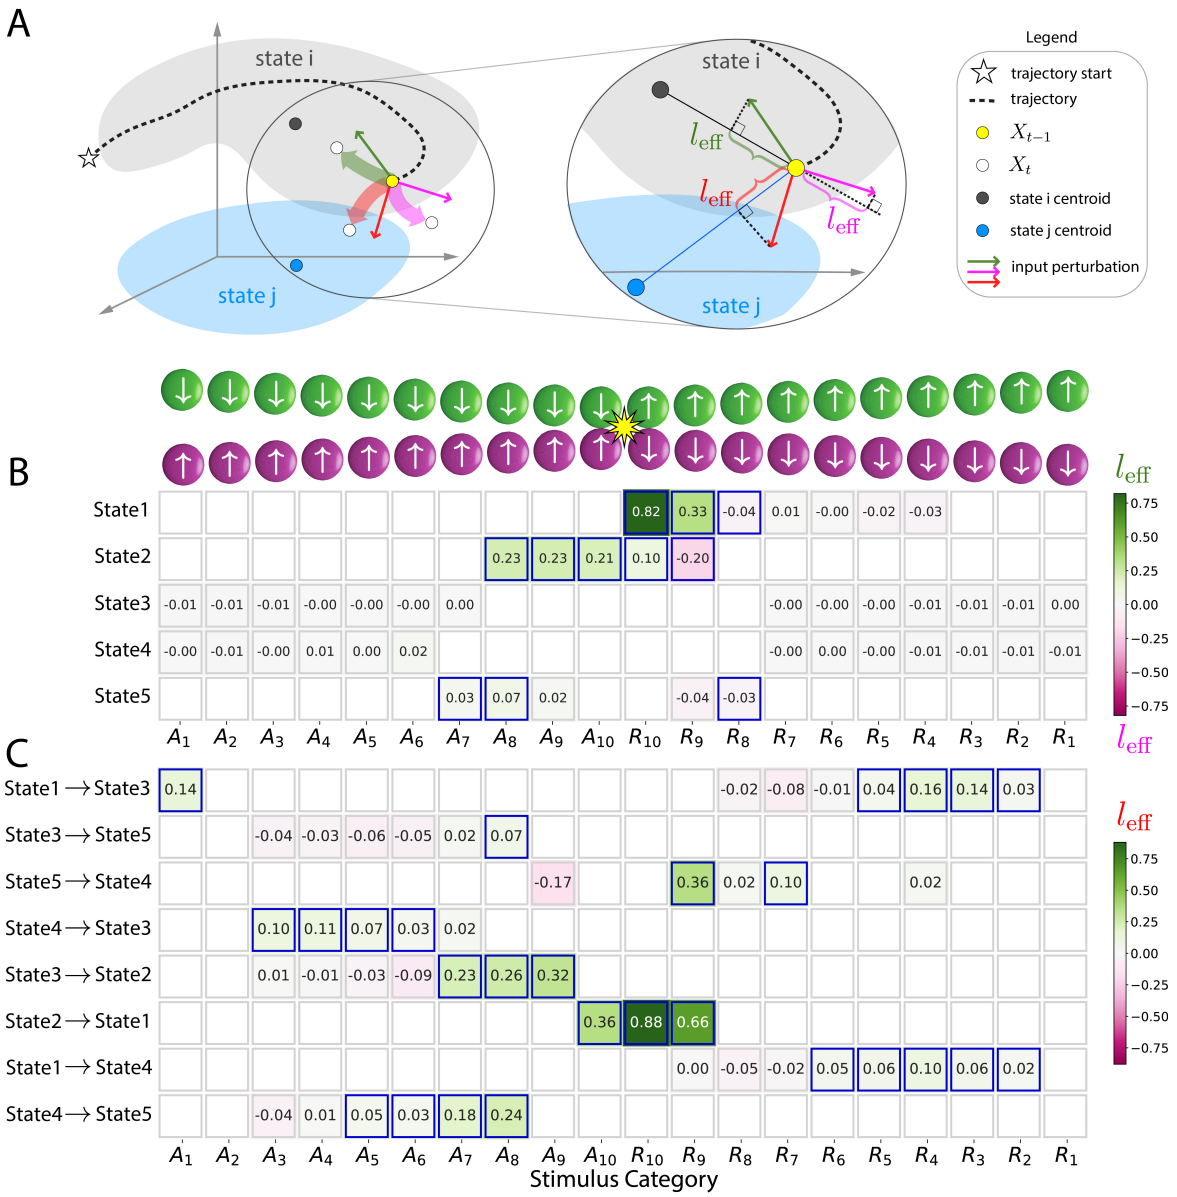

Figure 7: Effect of external inputs on steering state trajectories and driving state transitions. **(A)** Left: Trajectory at $t-1$ in state i. Inputs (colored arrows) could perturb the trajectory (translucent paths of the same color) by steering the evolution towards the state's centroid (green) or away from it (pink). When the input directed it away from the state's centroid, the input could push the system to switch into state j (red). Right: The input effect was measured via the cosine of the angle ϕ between the input vector and the line joining X_{t-1} and state centroids. **(B and C)** Input effects with rows and columns representing states/state-transitions and input categories, respectively. Only states with significant association with inputs and only significant state transitions are shown (see Fig. 2). Cells with significant effects are highlighted in blue ($p < 0.05$, corrected for multiple comparisons).

3 Discussion

In the present paper, we applied Switching Linear Dynamical Systems (SLDS) to uncover the dynamics of threat processing during a continuous threat-of-shock paradigm. First, we demonstrated that the SLDS model learned the regularities of the experimental paradigm, such that states and state transitions estimated from fMRI time series data from 85 ROIs reflected both the proximity of the circles and their direction (approach vs. retreat). After establishing that the model captured key properties of threat-related processing, we characterized the

dynamics of the states and their transitions. The results revealed that threat processing can profitably be viewed in terms of dynamic multivariate patterns whose trajectories are a combination of intrinsic and extrinsic factors that jointly determine how the brain temporally evolves during dynamic threat.

Past work has employed state-space models such as HMMs to estimate states from brain signals in an unsupervised fashion. Previous fMRI studies detected states with blocked design,[56,57] where brain states are expected to align closely to the experimental structure. In continuous experimental paradigms, establishing brain states that meaningfully track experimental variables is considerably more challenging. In fact, in some studies employing HMMs with more unconstrained or naturalistic experiments, states were not explicitly mapped to stimulus events.[21,26,27] A recent interesting study applied HMMs to cell recordings in the rat PFC and found states that correlated with hide-and-seek behaviors, as well as states corresponding to previously undetected behaviors.[58] In general, although HMM models applied to brain data have shown to be a promising strategy to uncover unlabeled states, further work is needed to assess their performance in continuous paradigms in a manner that rigorously establishes the correspondence between states and external stimuli and/or behavior.

Critically, standard HMMs provide limited information about state dynamics, which is essential in understanding temporally unfolding brain processes supporting complex mental states and behaviors, such as the varying levels of threat and safety studied here. To address this gap, we applied the SLDS framework to our continuous threat-of-shock paradigm. There have been a few applications of SLDS models to brain data,[48,59,60] but as far as we know only one fMRI study investigating a blocked working memory task.[25] As one of the challenges with unsupervised state-space methods is that there is no ground truth concerning the states, here we first established that the model estimated states and transitions that captured the structure of the continuous paradigm, namely, circle proximity, shock events, and post-shock period. Of the six states detected by the model, five captured properties of the stimulus in a statistically significant manner. For example, the probability of being in STATE2 given that the circles collided (input: A_{10}) was 0.63, consistent with a state indicative of a "shock/peri-shock" period. State transitions were also associated with specific stimuli. More generally, the states captured the structure of the paradigm, which consisted of periods of increased and decreased threat given the proximity of the circles and the unpleasant shock delivered upon collision.

Further evidence that the states captured important features of threat-related processing was obtained by determining voxelwise brain maps for both individual states and state contrasts. For example, the contrast of STATE5 ("near miss") vs. STATE4 ("not close") revealed a pattern of activation very similar to that obtained in the literature for the contrast of higher vs. lower levels of threat.[37,38,42,49–54] In particular, the contrast identified sectors of ventromedial PFC and posterior cingulate that have been linked to relative safety.[49,53,55] In addition, state transitions reflected notable changes in activity levels, such as the transition of the lower-threat STATE4 to the higher-threat STATE5. Taken together, our analyses demonstrated that the model captured key properties of the continuous paradigm, building on previous studies using HMMs.[13,21,26,56,61]

Our central goal was to discover and characterize the dynamics of continuous threat-related processing. From the perspective developed here brain regions jointly participate in state-specific trajectories. An important property of the SLDS framework is that it estimates state dynamics based on separate and additive endogenous and exogenous contributions. Determining the intrinsic dynamics is particularly informative because it uncovers how a brain state would evolve in the absence of additional inputs. In particular, does the state trajectory tend toward a fixed attractor, or does it evolve in some other more complex manner? Based on the properties of the endogenous state dynamics matrix, we determined that indeed all states behaved as attractors. But what do these attractors represent in terms of brain activity? To visualize the corresponding patterns, we determined "attractor maps" by projecting the model attractor (determined in the latent space) onto the brain space (times series data from ROIs). The attractor maps captured essential properties of the states determined by the model. For example, STATE1 and STATE2 maps exhibited inverse activity patterns, consistent with the fact that they largely reflected the post-shock and the shock/peri-shock periods, respectively. In addition, comparison of STATE5 ("near miss") and STATE4 ("not near") shows how distributed patterns with shared dynamics are involved in mental states of higher and lower anxious apprehension.

The endogenous dynamics can be further understood as providing a flow field that channels the temporal evolution of the system in the absence of inputs. However, in our continuous threat-of-shock paradigm, state trajectories were also determined by the inputs acting at each time point. Thus, visualizing the trajectory in terms of vector fields uncovered both contributions to system behavior. In combination with the information obtained from attractor maps, these results provide novel ways to understand the roles of key brain regions involved in threat-related processing. In addition to the qualitative information provided by vector field trajectories, we quantified the input contributions to states and state transitions. We determined how inputs contributed to

trajectories in a state's subspace and to transitions between states. For example, when the system was in STATE2 (shock/peri shock), the input perturbation R_{10} steered the system in the direction of STATE1 most vigorously, showing that soon after the circles collided, the brain tended to readily transition to a state of decreased activity in the insula and medial PFC. Overall, the analysis revealed how inputs steered a state's trajectory toward or away from the state's centroid, as well as how inputs directed a trajectory's evolution in the direction of a new state.

Work on dynamics of neural circuits in systems neuroscience typically assumes that the target circuit is driven only by endogenous processes.[34] In other words, the external inputs do not vary significantly over time and therefore do not perturb state trajectories. While such a scenario may be adequate in some experimental paradigms, the assumption is problematic in cases where inputs vary continuously. In the present study, we did not assume that the brain was isolated from the environment. Instead, by parsing the dynamics in terms of endogenous and exogenous contributions, we were able to not only study intrinsic state attractors but also to characterize external inputs as perturbations that drove trajectories along meaningful directions, such as pulling or pushing inputs toward or away from state centroids. Additionally, we also identified input stimuli that effectively steered evolution in the direction of other state subspaces. We propose that our approach can be profitably employed in studying neuronal circuit dynamics in systems neuroscience more broadly.

Here, we also discuss some challenges and potential shortcomings of our study. The application of state-space models, including SLDS, to hemodynamic data is challenging because of the low-pass nature of the response. Specifically, the response to a brief event builds up to a peak around 4-6 seconds post onset and decays slowly over the next 4-8 seconds.[62] In the current study, we assumed a constant hemodynamic lag of approximately 5 seconds. Although our paradigm involved continuous movement of the circles, both shock events and transitions from approach to retreat (and vice versa) were discrete in nature. These events call for careful interpretation of the results given the potential for blending of hemodynamic responses. However, our experiment was designed such that most periods of circle approach and retreat lasted for 6-9 seconds so as to minimize consecutive approach/retreat transitions.[43] Thus, the concern about the slow nature of the hemodynamic response was somewhat alleviated in the present context. In general, future studies are needed to develop SLDS approaches that can explicitly consider the properties of the hemodynamic response. Another limitation of the framework developed here is that we assumed, as in all state-model applications, that the underlying systems is well approximated by a series of states. With SLDS models this is less severe an assumption because one can consider the (possibly nonlinear) behavior of a system as modeled by a series of linear dynamical systems that approximate the overall behavior. Nevertheless, it will be valuable to investigate other approaches that address this concern, including considering that multiple states can coexist at a time.[63]

In the present paper, we applied Switching Linear Dynamical Systems to uncover the dynamics of threat processing during a continuous threat-of-shock paradigm. First, we demonstrated that the model learned the regularities of the experimental paradigm, such that states and state transitions estimated from fMRI time series data from 85 ROIs reflected both the proximity of the circles and their direction (approach vs. retreat). After establishing that the model captured key properties of threat-related processing, we characterized the dynamics of the states and their transitions. The results revealed that threat processing can profitably be viewed in terms of dynamic multivariate patterns whose trajectories are a combination of intrinsic and extrinsic factors that jointly determine how the brain temporally evolves during dynamic threat. We propose that viewing threat processing through the lens of dynamical systems offers important avenues to uncover properties of the dynamics of threat that are not uncovered with standard experimental designs and analyses.

Acknowledgements

This research was supported by the National Institute of Mental Health (R01-MH-071589). We thank Xiaoyu Zhou for discussions about the statistical approach developed in the paper.

Declaration of Interests

The authors declare no competing interests.

4 Methods

4.1 Participants

One hundred and twenty-six participants (63 females, ages 18–30 years; average: 20.87, STD: 2.56) with normal or corrected-to-normal vision and no reported neurological or psychiatric disease were recruited from the University of Maryland community. The project was approved by the University of Maryland College Park Institutional Review Board and all participants provided written informed consent before participation. Data from four participants were not used due to technical issues. For the original publication of these data, see.[43]

Out of the 122 participants, we used 30 to perform preliminary testing of the framework and defining model hyperparameters, including the number of states and the dimensionality of the latent space(see following subsections and supplementary). In order to avoid circularity in our analyses, data from this exploratory set were not used further. Thus, all results reported here are based on the held-out set of 92 participants.

4.2 Procedure and Stimuli

The experiment was designed to study the role of proximity on threat-of-shock processing. Participants watched two circles of different colors on the screen, at times moving close to each other, at times moving apart in a smooth but unpredictable fashion. Upon circle collision, the participants received a mild but unpleasant electrical stimulation together with an aversive sound. The movement of the circles was defined such that there were segments of approach and retreat that lasted between 2 and 9 seconds, with most segments lasting more then 6 seconds. To enhance unpredictability and hence anxious apprehension by participants, the movement of the circles included multiple instances of "near misses" in which the circles nearly touched (distance as close as 1.5 circle diameters between circle centers) before retreating away.

The experiment consistent of six 8-minute runs, each having two 3-minute blocks, during which the circles moved. The participants experienced on average 4 circle collisions and 7 near misses per run. All runs used different motion paths, but the visual stimulus was fixed across participants. Finally, visual stimuli were presented using `PsychoPy`(http://www.psychopy.org/) and viewed on a projection screen via a mirror mounted to the scanner's head coil. For further details about the experiment, the reader may refer to.[42,43]

4.3 Regions of Interest

Time series data consisted of fMRI time series from 85 cortical and subcortical ROIs (defined as the average across voxels for each ROI). The ROIs were defined in a prior study by our group and intended to capture brain regions involved in threat-related processing.[50]

4.4 Switching Linear Dynamical Systems

Switching Linear Dynamical Systems (SLDSs) are an example of a latent-variable framework that expresses the dynamics of observed signals using a set of unobserved variables assumed to govern system dynamics. The latent space is typically of considerable smaller dimension relative to the original data space (here 10 and 85, respectively). SLDSs consist of one linear dynamical system model per state, such that the trajectory of the system is governed by a single dynamical system at at each time step. Switching between states is assumed to follow a Markov process. Overall, SLDS models provide a piecewise linear approximation of (potentially nonlinear) dynamics.

Consider activity $Y_t \in \mathbb{R}^N, t = 1, 2, \ldots, T$ (T time steps is the duration of a typical run) from $N = 85$ brain ROIs during an experiment described using M input variables $U_t \in \mathbb{R}^M, t = 1, 2, \ldots, T$. The model represents brain signals Y_t in a low dimensional subspace as a set of D latent variables $X_t \in \mathbb{R}^D, t = 1, 2, \ldots, T$ ($D < N$). Observed BOLD activity activity is related to the latent variables using a linear transformation, such that

$$(Y_t|X_t) \sim \mathcal{N}\left(CX_t + d, \Sigma^v\right), \tag{1}$$

where $C \in \mathbb{R}^{N\times D}$ is the observation matrix, $d \in \mathbb{R}^N$ is a bias term, $\Sigma^v \in \mathbb{R}^{N\times N}$ is a noise covariance matrix and $\mathcal{N}(\mu, \Sigma)$ is a multivariate Gaussian distribution with mean μ and covariance Σ. The expression above is called the observation step.

Latent trajectory dynamics is modeled using a set of K linear dynamical systems that are indexed by a time varying switching variable $Z_t \in \{1, 2, \ldots, K\}$, which we call the state. At each time point t only the linear dynamical system corresponding to the state Z_t drives the temporal evolution of system, such that

$$(X_t|X_{t-1}, Z_t, U_t) \sim \mathcal{N}\left(A_{Z_t} X_{t-1} + V_{Z_t} U_t + b_{Z_t}, \Sigma^h_{Z_t}\right), \tag{2}$$

where $A_{Z_t} \in \mathbb{R}^{D\times D}$ is the linear dynamics matrix, $V_{Z_t} \in \mathbb{R}^{D\times M}$ is the input matrix reflecting the experimental paradigm variables, $b_{Z_t} \in \mathbb{R}^D$ is a bias term and $\Sigma^h_{z_t} \in \mathbb{R}^{D\times D}$ is a noise covariance matrix. The expression above is called the dynamics step.

State transition follows a modification of the Markov property, where the probability of switching to state Z_t from the previous state Z_{t-1} depends on the previous values of the state Z_{t-1} and the latent variables at $t-1$, X_{t-1}:

$$P\left(Z_t = j|Z_{t-1} = i, X_{t-1}\right) \propto \exp\left(\log(Q_{ij}) + r_j^\top X_{t-1}\right), \tag{3}$$

where $r_j \in \mathbb{R}^D$ and $Q \in R^{K\times K}$ is the Markovian transition matrix, whose i^{th} row and j^{th} column element indicates the transition probability from state i to state j. This expression above is called the state switching step. r_j is a weight vector (indexed by the state label, j) that specifies the additional dependency of Z_t on X_{t-1} (on top of the Markov dependence of Z_{t-1} on Z_t), allowing state transitions based on the location of the trajectory in the latent space. SLDS with the above additional dependency is also known as recurrent SLDS[14] and is shown to be more suitable to explain realistic timeseries than SLDS without the recurrent dependencies.[48] Note that, on removing r_j from the above equation the state transition probability is the same as that of a Markov model.

A group-level SLDS model was estimated based on the data of 92 participants by using a previously developed variational Laplace Expectation Maximization algorithm.[14] Specifically, all of the following were estimated: $C, d, \Sigma^v, \{A_k, V_k, b_k, \Sigma^h_k\}_{k=1}^K, Q, r$. In this manner, we obtained model variables X_t and Z_t.

4.5 Modeling experimental stimuli as SLDS inputs

To encode the experimental paradigm, we categorized the stimuli based on the proximity and direction of motion of the moving circles. The instantaneous proximity was calculated as the normalized Euclidean distance between the circles, ranging from 0 (farthest) to 1 (touching). Direction was calculated based on the discrete difference of proximity: $+1$ (approach) or -1 (retreat). Proximity values were quantized into 10 equal-sized bins ($[0.0, 0.1)$, $[0.1, 0.2)$, $\cdots$, $[0.9, 1.0]$), one per binning per direction, resulting in 20 stimulus categories: A_1, A_2, $\cdots$, A_{10} for approach; R_1, R_2, $\cdots$, R_{10} for retreat). The stimulus conditions were thus encoded as a 20-dimensional binary vector U_t at each time point t, with a 1 corresponding to the active category and 0 elsewhere. The number of stimulus categories was arbitrarily defined, but we verified that key results were robust to the number of input bins (see Supplementary Fig. S13 and S14).

For a state k, the SLDS model represents the contribution of input on dynamics using the term $V_k U_t$ (Eq. 2), where U_t encodes the stimulus and V_k maps the input into the latent space. Note V_k is state-specific, allowing distinct representations of the same stimulus category per state.

4.6 Assessing statistical significance using bootstrapping

Data from the held-out set of 92 participants were used to report findings of the study and determine statistical significance, which was determined via bootstrapping.[64] A total of 500 bootstrap samples of the dataset were generated by sampling participants with replacement, and an SLDS model was fit on each bootstrap iteration.

Because state labels are not necessarily preserved across bootstrap samples, we realigned them using a combination of the Hungarian algorithm[65] and k-means clustering. To do so, for each bootstrap sample, state parameters A_k, b_k, Σ^h_k were concatenated to form $K = 6$ "large vectors" θ, one for each state. The resulting vectors were then clustered into K clusters. The K cluster centroids provided representative vectors θ for each state. In this manner, aligning state labels across bootstrap samples amounted to assigning each bootstrap sample's θ vectors to the K cluster centroids. The assignment problem was determined by the Hungarian assignment algorithm, such that the bootstrap sample's state θ vectors were maximally aligned to the K cluster centroids.

To determine the number of bootstrap iterations, we generated samples in steps of 100 until the estimates of the above mentioned θ centroids were stable. Specifically, the elements of each centroid did not change (within 10^{-5} precision) upon adding bootstrap samples. The procedure was stopped at 500 bootstrap samples.

4.7 Calculating the probability of a state given a stimulus

To determine $P(\text{state}|\text{input})$, we defined a state by input table. Each entry corresponded to the proportion of time a state Z_t assumed a specific value z given that the input U_t assumed a specific value u. This counting procedure can be briefly summarize in the (somewhat awkward) expression

$$P(Z_t = z|U_t = u) = \frac{\sum_{t=1}^{T} \mathbb{1}(Z_t = z)\mathbb{1}(U_t = u)}{\sum_{t=1}^{T} \mathbb{1}(U_t = u)},$$

where $\mathbb{1}(\cdot)$ represents an indicator function that takes value 1 when the operand is true and 0 otherwise. Given the hemodynamic delay (approximately 2 secs to initiate a response;[62]), we used a time shift of 3 samples (3.75s) in considering the relationship between SLDS state and stimulus category.

We tested the statistical significance of each entry in the state by input table as follows. To account for sampling variability of the sample we studied, we computed a table for each iteration of the bootstrap, allowing us to estimate the variability of our results across the population (group level). Similarly, we obtained surrogate tables, one for each bootstrap resample, by randomly permuting the order of the states. In this manner, the relative order of the states was random but the other properties were maintained (transition structure in the sequences). By obtaining bootstrap and permuted tables for for each table cell corresponding to a given state and stimulus, we obtained a population-level distribution of table values and a null distribution.

Based on these results, each table entry mean was compared to the mean of the corresponding null distribution via a paired Wilcoxon signed rank test at $p < 0.05$. To correct for multiple comparisons across the number of states and stimulus categories (6×20), Bonferroni correction was applied.

4.8 Calculating the probability of a state transitions given a stimulus

State transitions were linked to the stimuli using an analogous procedure as that to map states to stimuli. We created a table with 8 state transitions of interest (see supplementary methods, Fig. S10) by 20 stimulus categories. Each table entry was calculated by counting the number of times a given state transition was observed given a stimulus category. Formally, this counting procedure can be succinctly as follows: for a row r representing transition state $i \mapsto$ state j and column u, the corresponding table entry measures the probability

$$P(Z_t = i, Z_{t+1} = j|U_t = u) = \frac{\sum_{t=1}^{T} \mathbb{1}(Z_t = i)\mathbb{1}(Z_{t+1} = j)\mathbb{1}(U_t = u)}{\sum_{t=1}^{T} \mathbb{1}(U_t = u)},$$

where $\mathbb{1}(\cdot)$ represents an indicator function. Once again, this is calculated by considering the hemodynamic delay. Finally, statistical inference was performed using a similar approach as described previously for calculating probabilities of states given stimuli.

4.9 State spatial maps

We sought to characterize the spatial organization of brain states: for a given state k, what is the typical pattern of activity throughout the brain. This could be accomplished by using multiple regression with variables for each state: 1 when the state is ON, 0 otherwise. We performed multiple regression in this manner in a voxel-wise manner. Note that this was possible because once the states are estimated (based on 85 ROIs), the regression can be applied at any spatial level desired (ROI or voxel). The regression model can be expressed mathematically as

$$\text{BOLD}_t = \sum_{i=k}^{K} \beta_i \times \mathbb{1}(Z_t = i) + \varepsilon_t,$$

where $\mathbb{1}(Z_t = k)$ is an indicator function denoting whether state k is active at time t, β_k is the corresponding regression coefficient, and ε_t is an error term. For notational simplicity, terms accounting for confounds like scanner drift and head motion have been omitted from the above expression but were included in the regression model.

State maps were evaluated statistically based on the bootstrap approach discussed previously. First, state activity maps were determined for every participant, separately. A group-level map was then obtained by averaging across participants. This procedure was performed for every bootstrap sample, allowing us to estimate sampling variability. Similarly, a set of null state maps was generated by repeating the same procedure but with

randomly permuted state sequences. To identify statistically significant voxels, we performed a paired t-test between the participant-based bootstrapped maps and the null maps. To account for multiple comparisons, we employed a cluster-based approach.[66] Voxels were first thresholded at $p < 0.001$, and clusters with at least $c = 18$ voxels were deemed statistically significant at the cluster level. To determine the minimum cluster size c, we employed null state maps. The cluster size c was such that only 5% of the clusters in the null state maps would be detected, corresponding to a 5% false-positive rate.

For contrast maps, an analogous approach was employed, but at the individual level contrast maps were determined by computing $\beta_i - \beta_j$, where i and j are two states in question.

4.10 Responses during state transitions

Responses at a given ROI were evaluated during temporal windows centered around state transitions. For transitions from state i to state j, we chose windows where the first half of the time points were labelled as state i and the final time points as state j. We chose windows of length 8 time steps such that they were long enough to reveal activity changes but also observed frequently enough to generate a reliable estimate. For a given state transition, participant-level responses were determined by averaging responses across all qualifying time windows. To estimate variability, group- level responses were determined by averaging participant-level responses for each bootstrap sample.

We also determined responses at the "network" level, which were obtained by averaging ROI responses across all ROIs inside the network. We considered the salience (ventral attention) and default networks, as defined in the Schaefer 100-region 7-network parcellation.[67]

4.11 State specific intrinsic dynamics and fixed points

Each state's dynamics in the latent space was modelled using a linear dynamical system, which evolves as follows:

$$X_t = A_k X_{t-1} + b_k + V_k U_t. \tag{4}$$

To characterize the intrinsic dynamics exclusively, assume that the input to the system is zero across time:

$$X_t = A_k X_{t-1} + b_k. \tag{5}$$

Assuming that the dynamics are stable and approach an equilibrium point $\overline{X}_k$ as a function of time (also known as a "fixed point" or attractor), we obtain

$$\overline{X}_k = A_k \overline{X}_k + b_k$$
$$\overline{X}_k = [I - A_k]^{-1} b_k,$$

where I is a $D \times D$ identity matrix. All trajectories in the neighbourhood of the fixed point (also called the basin of attraction) will eventually converge towards the fixed point.

4.12 Testing for stable intrinsic dynamics

The dynamics of a linear dynamical system is determined by the eigenvalues of the dynamics matrix A_k. In particular, the dynamics is stable if all eigenvalues lie inside a unit circle on the complex plane (refer to[68,69] for detailed accounts on stability of linear dynamical systems). Therefore it suffices to test whether norm of the largest eigenvalue is less than unity. Bootstrapping was again used to evaluate the statistical significance of the results. We determined largest eigenvalues for all bootstrap samples, and calculated the p-value as the fraction of the eigenvalues with absolute value greater than 1. States with p-value < 0.05 were considered to follow stable dynamics. Bonferroni correction was applied to correct for multiple comparisons across K states.

4.13 Visualizing state attractors on the brain

Given an attractor state i, its fixed point $\bar{X}_i$ is a vector in the D-dimensional latent space. For the purpose of visualizing state attractors on the brain, we projected them onto the original data space (the time series space of 85 ROIs) using the observation step of the SLDS model (Eq. 1).

$$\bar{X}_i^{\text{ROI}} = C\bar{X}_i + d \qquad\qquad i = 1, 2, \cdots, K,$$

where $\bar{X}_i^{\text{ROI}}$ is state i's attractor in ROI space. State attractors in ROI space were estimated across bootstrap samples, and the average pattern was visualized. The average attractor pattern was thresholded at each ROI ($p < 0.05$) based on a Wilcoxon signed-rank test ($p < 0.05$); that is, by comparing the activity levels to zero. Bonferroni correction was applied to correct for multiple comparisons across K states and 85 ROIs.

4.14 State trajectories and vector fields

For each state, trajectories were determined first in the latent space ($D = 10$). We averaged (across participants) all windows with 10 time steps starting at a specific state k. Note that trajectories correspond to Eq. 4 and consider both endogenous and exogenous terms. At each time step of the average trajectory, a trajectory could continue along a state or switch to a state $j \neq k$. The state label for a time point along the trajectory was assigned to the most typical state at that time point; i.e., the modal state. We colored the trajectory green when the state label was k and red otherwise (state $j \neq k$). To illustrate the trajectory, we plotted it in 2D. To do so, we fit principal component analysis on the 10 points constituting the average trajectory and projected it onto the plane spanned by the top two principal components.

Trajectories were plotted on vector fields that represent how a particle would flow under the endogenous dynamics (i.e., based on A_k). This allowed us to visualized how external inputs contributed to the evolution of the system (note that, by definition, without inputs trajectories would converged to the state's attractor). At each point on the 2D plane, the vectors represent the trajectory evolution for a single time step. Vector field points were sampled uniformly on the 2D plane. The points were projected to the latent space using the top two principal components above, evolved according to endogenous dynamics rule (Eq. 5), and projected back onto the 2D plane again using the top two principal components. Finally, the state's fixed point attractor (a point in the latent space) was also projected onto the 2D plane in the same manner.

4.15 Quantifying contribution of inputs in state dynamics

As described above, state k's trajectory follows the following dynamical rule given in Eq. 4, with the term $V_k U_t$ representing contribution of external inputs as a vector in latent space. To investigate the contribution of external inputs to the latent trajectory, we determined how inputs steered the system's trajectory. For each state, we determined the state's centroid, namely, the vector $X^{c,k}$ that represented the centroid coordinate of all locations visited by the system during state k. Such geometrical location provided a reference position to investigate the input contributions.

Given a state, the external input at a given time can either drive latent trajectories towards the state centroid or away from it. To quantify input effects, we measured the length of the projection of the input vector along the direction connecting the position of the state trajectory and the state centroid:

$$l_{\text{eff}} = \frac{\langle X^{c,k} - X_t, V_k U_t \rangle}{\| X^{c,k} - X_t \|},$$

where $\langle \cdot, \cdot \rangle$ indicates the inner product between two vectors and $\| \cdot \|$ indicates the Euclidean norm. Thus, positive (negative) l_{eff} values indicate that the input pushes state trajectories towards (away from) the state centroid. We focused on inputs as defined throughout the paper (see Methods sec. 4.7). Accordingly, the input effects summarised in a table where the $(i, j)^{\text{th}}$ element represents the effect of input j on trajectories in state i.

Tables were constructed as previously such that each bootstrap iteration generated one instance of the table. The results shown correspond to the average value across bootstrap samples. Statistically, we tested whether the average input effects were significantly different from zero by conducting a Wilcoxon's signed-rank test at each table entry ($p < 0.05$). Bonferroni correction was applied to correct for multiple comparisons across $K = 6$ states and 20 input categories.

4.16 Quantifying contribution of inputs in state transitions

Next, we quantified the contributions of input stimuli to state transitions. Specifically, for particular state transitions $i \mapsto j$, we determined the extent to which input stimuli contributed to such transitions. As above, l_{eff} was used to measure the effect of the external input vector along the direction joining the position of the state trajectory and state j's centroid (in place of state i's centroid). Positive values indicate that the input pushed the latent vector from state i towards state j.

Tables were constructed as previously, such that each bootstrap iteration generated one instance of the table. The results shown correspond to the average value across bootstrap samples. Statistically, we tested whether average input effects were significantly different from zero by conducting a Wilcoxon's signed-rank test at each table entry ($p < 0.05$). Bonferroni correction was applied to correct for multiple comparisons across 8 state transitions of interest and 20 input categories.

5 Supplementary Section

Determining the number of latent dimensions

Data from an exploratory set of participants ($N = 30$) was used to choose the dimensionality of the latent space. To do so, we computed the evidence lower bound (ELBO) as the optimization criterion of the SLDS fits to the data.[48,70] Intuitively, the ELBO is a lower bound on the log-likelihood of the observed data, therefore capturing the goodness of the model fit (i.e., the ability to explain the observed data). For each value $2, \ldots, 16$ an SLDS model was fit on 30 leave-one-out participant samples, and the average ELBO value across leave-one-out samples was determined. ELBO values in the range $[10, 13]$ were very similar according to the Fig. S8, so we chose $D = 10$.

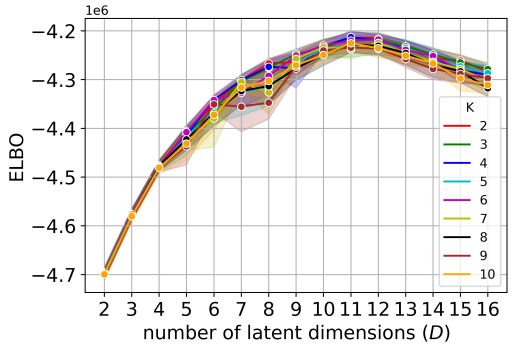

Figure 8: Supplemental figure: Model fit ELBO criterion as a function of the number of latent dimensions determined for multiple values of the number of states (K). Error bars indicate the 95% confidence interval based on standard error across leave-one-out samples.

Determining the number of states

Data from the exploratory dataset also were used to determine the number of states K. We investigated values from 2 to 10, and evaluated the consistency of estimated states across leave-one-participant-out folds. Specifically, we determined state sequences (one state per time point) for different K values. We followed the intuition that state sequences that were most consistent across data samples would be "better".[71] Consistency scores were calculated between all pairs of leave-one-participant-out samples using normalized mutual information. We chose the value $K = 6$ based on Fig. S9.

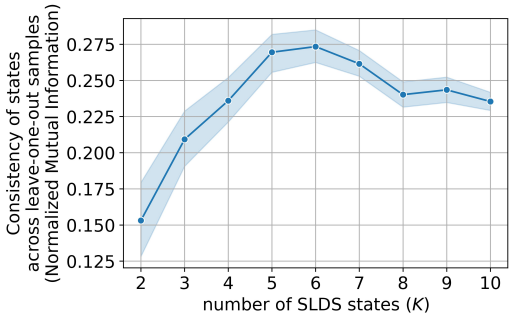

Figure 9: Supplemental figure: Consistency of state sequences across left-out participants as a function of the number of states (K). Error bars indicate the 95% confidence interval based on standard error across leave-one-out samples.

Statistically significant state transitions

We determined state transitions that occurred with frequencies higher than what would be expected by chance. Recall that state transitions were determined based on the state transition probability matrix. Our procedure was as follows. Transition matrices were estimated for each bootstrap sample. To determine a null distribution of transition matrices, for every bootstrap sample we temporally shuffled state sequences while preserving self-transitions (here, state transitions were the main object of study, not self transitions). For each state, randomization of the state sequence was performed by shuffling contiguous blocks attributed to the same state (instead of shuffling at the time point level). Significant state transitions were those for which the mean of the transition probability (based on bootstrap samples) was greater than the mean based on the shuffled versions (paired Wilcoxon signed rank test, $p < 0.05$). To correct for multiple comparisons ($K \times (K-1)$ transitions while ignoring self transitions), we applied Bonferroni correction.

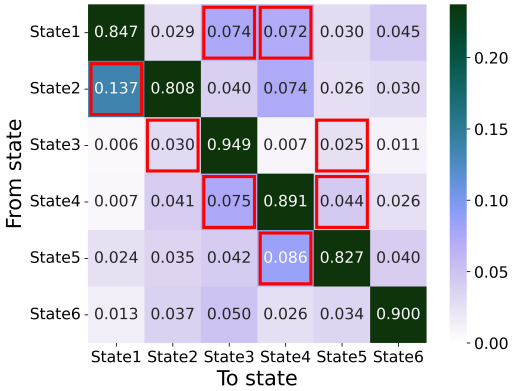

Figure 10: Supplemental figure: Average state transition matrix across bootstrap samples. State transitions highlighted in red have occurred more frequently than expected by chance ($p < 0.05$, corrected for multiple comparisons).

Voxelwise activity maps for all states

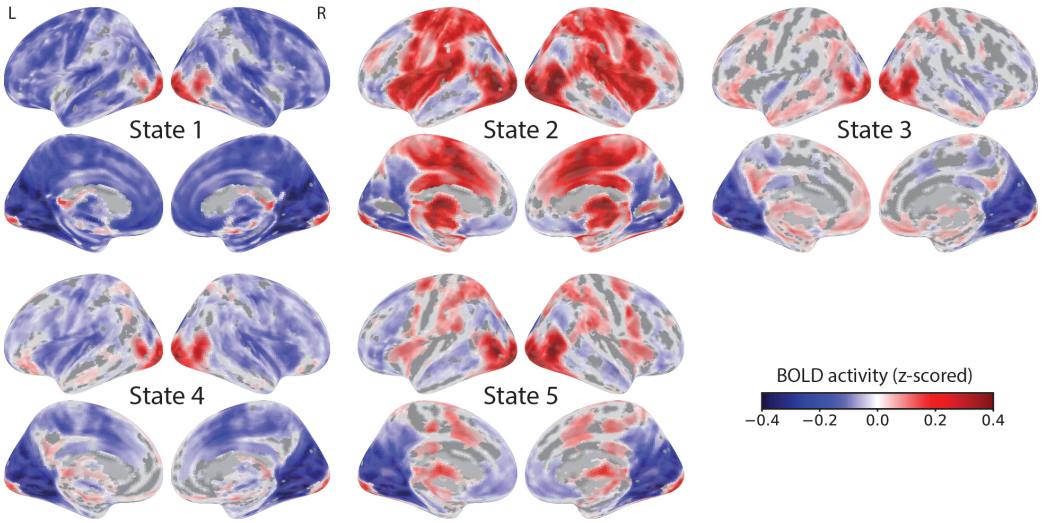

Figure 11: Supplemental figure: Voxelwise state activity maps for all states. Activity maps indicate signal intensity at a voxel when the state was "on". Maps were corrected for multiple comparisons by thresholding voxels at $p < 0.001$ and cluster level at $p < 0.05$.

Stability of state endogenous dynamics

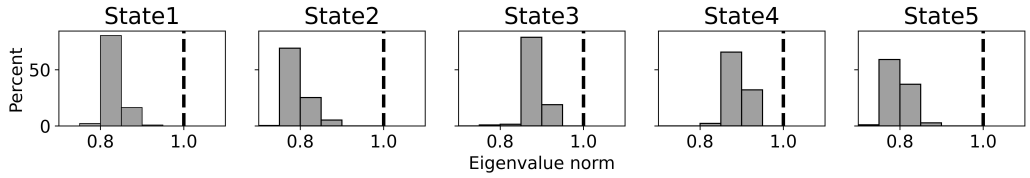

Figure 12: Supplemental figure: Stability of states. Distribution of the norms of the largest eigenvalue of the dynamics matrix (A) across bootstrap samples. A state was considered an attractor if 95% of the norms were smaller than 1.

Effect of the number of input stimulus categories (bins) in the analysis of Figure 2

We evaluated how the choice of input categories (bins) impacted the results shown in Figure 2. Overall, the results generalized well to 10 and 30 bins, instead of the original 20. However, we also note that that a "near miss" state (STATE5) was not well captured with only 10 bins given that it included an input (A_5) that included the circle collision, and it did not include the retreat input that was closest to the collision (R_4).

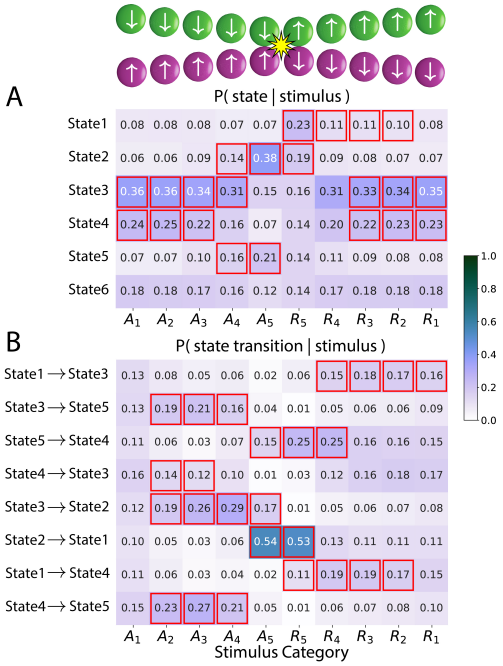

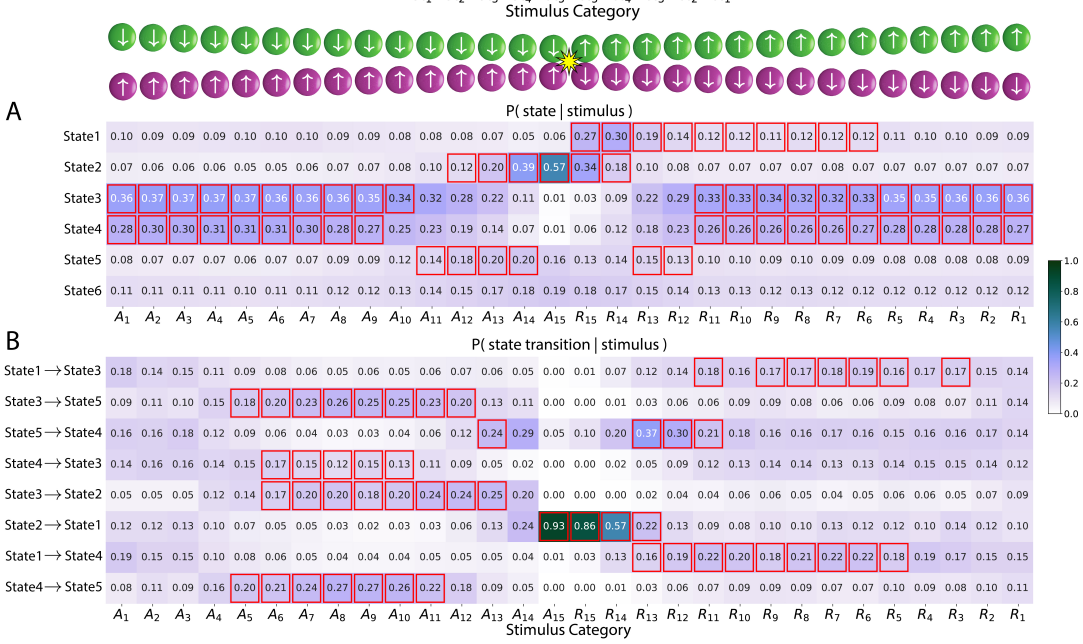

Figure 13: Supplement to Fig. 2: Brain states, state transitions, and input stimuli. **(Top)** Version with 10 stimulus categories. **(Bottom)** Version with 30 stimulus categories. Stimuli were categorized into bins based on circle distance and movement direction indicated with A for approach and R for retreat (1 indicates circles are farthest and 5 or 15 indicates circle collision). **(A)** Table entries indicate the probability of being in a state given the input stimulus category. **(B)** Table entries indicate the probability of a state transition given the input stimulus category. In both tables, cells highlighted in red indicate states/transitions significantly associated with the corresponding stimulus category ($p < 0.05$, corrected for multiple comparisons). The color scale indicates probability.

Effect of the number of stimulus categories (bins) in the analysis of Figure 7

We also evaluated how the number of input stimulus categories (bins) impacted the results of the analysis shown in Figure 7. Broadly, the results generalized well across the number of bins used for inputs. However, some differences were observed. We discuss part (A) first. For example, in STATE2 (shock/peri-shock) in the original case with 20 bins and with 10 bins, inputs at R_9 (20 bins) and R_5 (10 bins) steered the system to STATE1 (post-shock); in contrast, such post-shock influence of the inputs was not detected with 30 bins. In addition, for STATE1 (post-shock), inputs after the shock steered the trajectory away from the state (input R_8 with 20 bins and R_{12} with 30 bins); in contrast, no corresponding tendency was observed with 10 bins. We discuss part (B) next. Input contributions to state transitions were consistent across the binning of inputs with one exception: no input was found to significantly steer the trajectory away from STATE3 in the case of the STATE3 $\mapsto$ STATE5 transition when 10 bins were used (see second row).

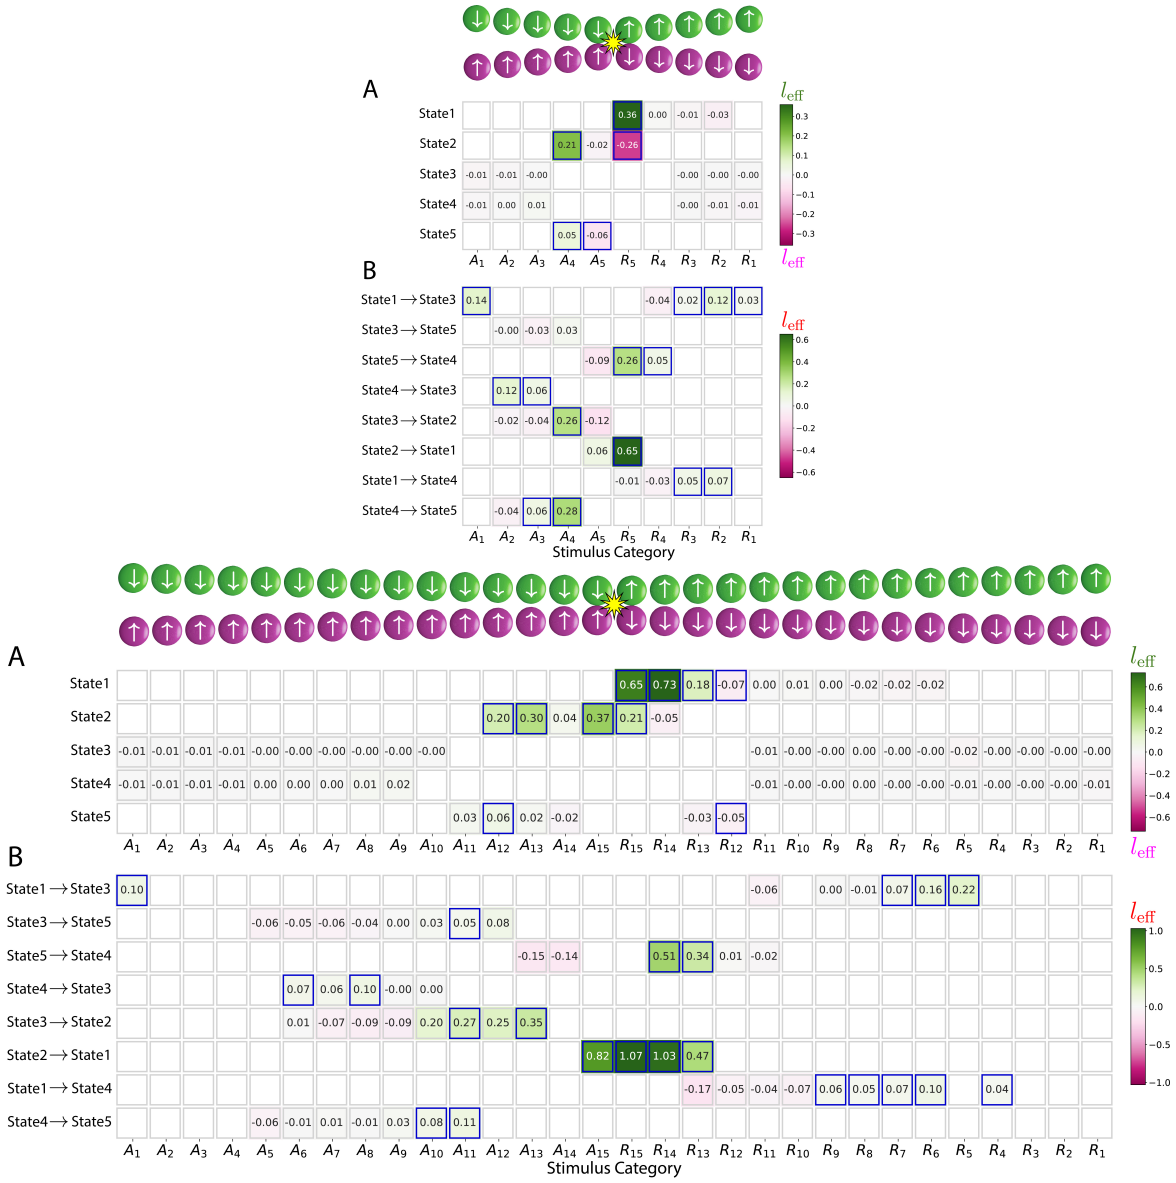

Figure 14: Supplement to Fig. 7: Effect of external inputs on steering state trajectories and driving state transitions. **(Top)** Version with 10 stimulus categories. **(Bottom)** Version with 30 stimulus categories. **(A and B)** Input effects with rows and columns representing states/state-transitions and input categories, respectively. Only states with significant association with inputs and only significant state transitions are shown. Cells with significant effects are highlighted in blue ($p < 0.05$, corrected for multiple comparisons).

References

[1] L. T. Hunt, N. D. Daw, P. Kaanders, M. A. MacIver, U. Mugan, E. Procyk, A. D. Redish, E. Russo, J. Scholl, K. Stachenfeld, C. R. E. Wilson, and N. Kolling. Formalizing planning and information search in naturalistic decision-making. *Nature Neuroscience*, 24(8):1051–1064, August 2021. Publisher: Nature Publishing Group.

[2] Camille Testard, Sébastien Tremblay, Felipe Parodi, Ron W. DiTullio, Arianna Acevedo-Ithier, Kristin L. Gardiner, Konrad Kording, and Michael L. Platt. Neural signatures of natural behaviour in socializing macaques. *Nature*, pages 1–10, March 2024. Publisher: Nature Publishing Group.

[3] Uri Hasson and Christopher J. Honey. Future trends in Neuroimaging: Neural processes as expressed within real-life contexts. *NeuroImage*, 62(2):1272–1278, August 2012.

[4] Saurabh Sonkusare, Michael Breakspear, and Christine Guo. Naturalistic Stimuli in Neuroscience: Critically Acclaimed. *Trends in Cognitive Sciences*, 23(8):699–714, August 2019.

[5] Emily S. Finn, Enrico Glerean, Arman Y. Khojandi, Dylan Nielson, Peter J. Molfese, Daniel A. Handwerker, and Peter A. Bandettini. Idiosynchrony: From shared responses to individual differences during naturalistic neuroimaging. *NeuroImage*, 215:116828, July 2020.

[6] Haemy Lee Masson and Leyla Isik. Functional selectivity for social interaction perception in the human superior temporal sulcus during natural viewing. *NeuroImage*, 245:118741, December 2021.

[7] Clare Grall and Emily S Finn. Leveraging the power of media to drive cognition: a media-informed approach to naturalistic neuroscience. *Social Cognitive and Affective Neuroscience*, 17(6):598–608, June 2022.

[8] Katalin M. Gothard, Clayton P. Mosher, Prisca E. Zimmerman, Philip T. Putnam, Jeremiah K. Morrow, and Andrew J. Fuglevand. New perspectives on the neurophysiology of primate amygdala emerging from the study of naturalistic social behaviors. *WIREs Cognitive Science*, 9(1):e1449, 2018. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcs.1449.

[9] Dean Mobbs, Pete C. Trimmer, Daniel T. Blumstein, and Peter Dayan. Foraging for foundations in decision neuroscience: insights from ethology. *Nature Reviews Neuroscience*, 19(7):419–427, July 2018. Publisher: Nature Publishing Group.

[10] Dominic A. Evans, A. Vanessa Stempel, Ruben Vale, and Tiago Branco. Cognitive Control of Escape Behaviour. *Trends in Cognitive Sciences*, 23(4):334–348, April 2019.

[11] John R. Anderson and Jon M. Fincham. Discovering the Sequential Structure of Thought. *Cognitive Science*, 38(2):322–352, 2014. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cogs.12068.

[12] Diego Vidaurre, Andrew J. Quinn, Adam P. Baker, David Dupret, Alvaro Tejero-Cantero, and Mark W. Woolrich. Spectrally resolved fast transient brain states in electrophysiological data. *NeuroImage*, 126:81–95, February 2016.

[13] Christopher Baldassano, Janice Chen, Asieh Zadbood, Jonathan W. Pillow, Uri Hasson, and Kenneth A. Norman. Discovering Event Structure in Continuous Narrative Perception and Memory. *Neuron*, 95(3):709–721.e5, August 2017.

[14] Scott W. Linderman, Andrew C. Miller, Ryan P. Adams, David M. Blei, Liam Paninski, and Matthew J. Johnson. Recurrent switching linear dynamical systems, October 2016. arXiv:1610.08466 [stat].

[15] Jalil Taghia, Srikanth Ryali, Tianwen Chen, Kaustubh Supekar, Weidong Cai, and Vinod Menon. Bayesian switching factor analysis for estimating time-varying functional connectivity in fMRI. *NeuroImage*, 155:271–290, July 2017.

[16] Chethan Pandarinath, Daniel J. O'Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D. Stavisky, Jonathan C. Kao, Eric M. Trautmann, Matthew T. Kaufman, Stephen I. Ryu, Leigh R. Hochberg, Jaimie M. Henderson, Krishna V. Shenoy, L. F. Abbott, and David Sussillo. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10):805–815, October 2018. Publisher: Nature Publishing Group.

[17] Rory G. Townsend and Pulin Gong. Detection and analysis of spatiotemporal patterns in brain activity. *PLOS Computational Biology*, 14(12):e1006643, December 2018. Publisher: Public Library of Science.

[18] Hiroshi Morioka, Vince Calhoun, and Aapo Hyvärinen. Nonlinear ICA of fMRI reveals primitive temporal structures linked to rest, task, and behavioral traits. *NeuroImage*, 218:116989, September 2020.

[19] Lei Wu, Arvind Caprihan, and Vince Calhoun. Tracking spatial dynamics of functional connectivity during a task. *NeuroImage*, 239:118310, October 2021.

[20] Matthew F. Singh, Anxu Wang, Michael Cole, ShiNung Ching, and Todd S. Braver. Enhancing task fMRI preprocessing via individualized model-based filtering of intrinsic activity dynamics. *NeuroImage*, 247:118836, February 2022.

[21] Hayoung Song, Won Mok Shim, and Monica D Rosenberg. Large-scale neural dynamics in a shared low-dimensional state space reflect cognitive and attentional dynamics. *eLife*, 12:e85487, July 2023. Publisher: eLife Sciences Publications, Ltd.

[22] Yiben Xu, Xian Long, Jianfeng Feng, and Pulin Gong. Interacting spiral wave patterns underlie complex brain dynamics and are related to cognitive processing. *Nature Human Behaviour*, 7(7):1196–1215, July 2023. Publisher: Nature Publishing Group.

[23] Parsa Vahidi, Omid G. Sani, and Maryam M. Shanechi. Modeling and dissociation of intrinsic and input-driven neural population dynamics underlying behavior. *Proceedings of the National Academy of Sciences*, 121(7):e2212887121, February 2024. Publisher: Proceedings of the National Academy of Sciences.

[24] Ruiqi Chen, Matthew Singh, Todd S. Braver, and ShiNung Ching. Dynamical models reveal anatomically reliable attractor landscapes embedded in resting state brain networks, February 2024. Pages: 2024.01.15.575745 Section: New Results.

[25] Jalil Taghia, Weidong Cai, Srikanth Ryali, John Kochalka, Jonathan Nicholas, Tianwen Chen, and Vinod Menon. Uncovering hidden brain state dynamics that regulate performance and decision-making during cognition. *Nature Communications*, 9(1):2505, June 2018. Publisher: Nature Publishing Group.

[26] Johan N. van der Meer, Michael Breakspear, Luke J. Chang, Saurabh Sonkusare, and Luca Cocchi. Movie viewing elicits rich and reliable brain state dynamics. *Nature Communications*, 11(1):5004, October 2020. Publisher: Nature Publishing Group.

[27] Hayoung Song, Bo-yong Park, Hyunjin Park, and Won Mok Shim. Cognitive and Neural State Dynamics of Narrative Comprehension. *Journal of Neuroscience*, 41(43):8972–8990, October 2021. Publisher: Society for Neuroscience Section: Research Articles.

[28] Byron M Yu, Afsheen Afshar, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and Maneesh Sahani. Extracting Dynamical Structure Embedded in Neural Activity. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2005.

[29] Gustavo Deco, Viktor K. Jirsa, Peter A. Robinson, Michael Breakspear, and Karl Friston. The Dynamic Brain: From Spiking Neurons to Neural Masses and Cortical Fields. *PLOS Computational Biology*, 4(8):e1000092, August 2008. Publisher: Public Library of Science.

[30] Krishna V. Shenoy, Maneesh Sahani, and Mark M. Churchland. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience*, 36(Volume 36, 2013):337–359, July 2013. Publisher: Annual Reviews.

[31] Evan D. Remington, Seth W. Egger, Devika Narain, Jing Wang, and Mehrdad Jazayeri. A Dynamical Systems Perspective on Flexible Motor Timing. *Trends in Cognitive Sciences*, 22(10):938–952, October 2018.

[32] Drew N. Robson and Jennifer M. Li. A dynamical systems view of neuroethology: Uncovering stateful computation in natural behaviors. *Current Opinion in Neurobiology*, 73:102517, April 2022.

[33] Paul Miller. Itinerancy between attractor states in neural systems. *Current Opinion in Neurobiology*, 40:14–22, October 2016.

[34] Mikail Khona and Ila R. Fiete. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, pages 1–23, November 2022. Publisher: Nature Publishing Group.

[35] Yohan J. John, Kayle S. Sawyer, Karthik Srinivasan, Eli J. Müller, Brandon R. Munn, and James M. Shine. It's about time: Linking dynamical systems with human neuroimaging to understand the brain. *Network Neuroscience*, 6(4):960–979, October 2022.

[36] Dean Mobbs, Predrag Petrovic, Jennifer L. Marchant, Demis Hassabis, Nikolaus Weiskopf, Ben Seymour, Raymond J. Dolan, and Christopher D. Frith. When Fear Is Near: Threat Imminence Elicits Prefrontal-Periaqueductal Gray Shifts in Humans. *Science*, 317(5841):1079–1083, August 2007. Publisher: American Association for the Advancement of Science.

[37] Dean Mobbs, Jennifer L. Marchant, Demis Hassabis, Ben Seymour, Geoffrey Tan, Marcus Gray, Predrag Petrovic, Raymond J. Dolan, and Christopher D. Frith. From Threat to Fear: The Neural Organization of Defensive Fear Systems in Humans. *Journal of Neuroscience*, 29(39):12236–12243, September 2009. Publisher: Society for Neuroscience Section: Articles.

[38] Dinavahi V. P. S. Murty, Songtao Song, Srinivas Govinda Surampudi, and Luiz Pessoa. Threat and Reward Imminence Processing in the Human Brain. *Journal of Neuroscience*, 43(16):2973–2987, April 2023. Publisher: Society for Neuroscience Section: Research Articles.

[39] Daniel J. Levitas and Thomas W. James. Dynamic threat–reward neural processing under semi-naturalistic ecologically relevant scenarios. *Human Brain Mapping*, 45(4):e26648, 2024. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.26648.

[40] M. A. Fullana, B. J. Harrison, C. Soriano-Mas, B. Vervliet, N. Cardoner, A. Àvila Parcet, and J. Radua. Neural signatures of human fear conditioning: an updated and extended meta-analysis of fMRI studies. *Molecular Psychiatry*, 21(4):500–508, April 2016.

[41] Michael S. Fanselow and Laurie S. Lester. A Functional Behavioristic Approach to Aversively Motivated Behavior:: Predatory Imminence as a Determinant of the Topography of Defensive Behavior. In *Evolution and Learning*. Psychology Press, 1987. Num Pages: 28.

[42] Christian Meyer, Srikanth Padmala, and Luiz Pessoa. Dynamic Threat Processing. *Journal of Cognitive Neuroscience*, 31(4):522–542, April 2019.

[43] Chirag Limbachia, Kelly Morrow, Anastasiia Khibovska, Christian Meyer, Srikanth Padmala, and Luiz Pessoa. Controllability over stressor decreases responses in key threat-related brain areas. *Communications Biology*, 4(1):1–11, January 2021. Publisher: Nature Publishing Group.

[44] Dean V. Buonomano and Wolfgang Maass. State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2):113–125, February 2009. Publisher: Nature Publishing Group.

[45] G. Ackerson and K. Fu. On state estimation in switching environments. *IEEE Transactions on Automatic Control*, 15(1):10–17, February 1970. Conference Name: IEEE Transactions on Automatic Control.

[46] David Barber. Expectation Correction for Smoothed Inference in Switching Linear Dynamical Systems. *Journal of Machine Learning Research*, 7(89):2515–2540, 2006.

[47] Emily Fox, Erik Sudderth, Michael Jordan, and Alan Willsky. Nonparametric Bayesian Learning of Switching Linear Dynamical Systems. In *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008.

[48] Scott Linderman, Annika Nichols, David Blei, Manuel Zimmer, and Liam Paninski. Hierarchical recurrent state space models reveal discrete and continuous dynamics of neural activity in c. elegans. *bioRxiv*, 2019.

[49] Dan W. Grupe and Jack B. Nitschke. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nature Reviews Neuroscience*, 14(7):488–501, July 2013. Publisher: Nature Publishing Group.

[50] Dinavahi V. P. S. Murty, Songtao Song, Kelly Morrow, Jongwan Kim, Kesong Hu, and Luiz Pessoa. Distributed and Multifaceted Effects of Threat and Safety. *Journal of Cognitive Neuroscience*, 34(3):495–516, February 2022.

[51] Dean Mobbs, Rongjun Yu, James B. Rowe, Hannah Eich, Oriel FeldmanHall, and Tim Dalgleish. Neural activity associated with monitoring the oscillating threat value of a tarantula. *Proceedings of the National Academy of Sciences*, 107(47):20582–20586, November 2010. Publisher: Proceedings of the National Academy of Sciences.

[52] Ruben P. Alvarez, Gang Chen, Jerzy Bodurka, Raphael Kaplan, and Christian Grillon. Phasic and sustained fear in humans elicits distinct patterns of brain activity. *NeuroImage*, 55(1):389–400, March 2011.

[53] Leah H. Somerville, Dylan D. Wagner, Gagan S. Wig, Joseph M. Moran, Paul J. Whalen, and William M. Kelley. Interactions Between Transient and Sustained Neural Signals Support the Generation and Regulation of Anxious Emotion. *Cerebral Cortex*, 23(1):49–60, January 2013.

[54] Juyoen Hur, Jason F. Smith, Kathryn A. DeYoung, Allegra S. Anderson, Jinyi Kuang, Hyung Cho Kim, Rachael M. Tillman, Manuel Kuhn, Andrew S. Fox, and Alexander J. Shackman. Anxiety and the Neurobiology of Temporally Uncertain Threat Anticipation. *Journal of Neuroscience*, 40(41):7949–7964, October 2020. Publisher: Society for Neuroscience Section: Research Articles.

[55] Daniela Schiller, Ifat Levy, Yael Niv, Joseph E. LeDoux, and Elizabeth A. Phelps. From fear to safety and back: reversal of fear in the human brain. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 28(45):11517–11525, November 2008.

[56] Diego Vidaurre, Romesh Abeysuriya, Robert Becker, Andrew J. Quinn, Fidel Alfaro-Almagro, Stephen M. Smith, and Mark W. Woolrich. Discovering dynamic brain networks from big data in rest and task. *NeuroImage*, 180:646–656, October 2018.

[57] Ayumu Yamashita, David Rothlein, Aaron Kucyi, Eve M. Valera, and Michael Esterman. Brain state-based detection of attentional fluctuations and their modulation. *NeuroImage*, 236:118072, August 2021.

[58] Bence Bagi, Michael Brecht, and Juan Ignacio Sanguinetti-Scheck. Unsupervised discovery of behaviorally relevant brain states in rats playing hide-and-seek. *Current Biology*, 32(12):2640–2653.e4, June 2022.

[59] Biljana Petreska, Byron M Yu, John P Cunningham, Gopal Santhanam, Stephen Ryu, Krishna V Shenoy, and Maneesh Sahani. Dynamical segmentation of single trials from population neural data. In *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.

[60] Joshua Glaser, Matthew Whiteway, John P Cunningham, Liam Paninski, and Scott Linderman. Recurrent Switching Dynamical Systems Models for Multiple Interacting Neural Populations. In *Advances in Neural Information Processing Systems*, volume 33, pages 14867–14878. Curran Associates, Inc., 2020.

[61] Jamal A. Williams, Elizabeth H. Margulis, Samuel A. Nastase, Janice Chen, Uri Hasson, Kenneth A. Norman, and Christopher Baldassano. High-Order Areas and Auditory Cortex Both Represent the High-Level Event Structure of Music. *Journal of Cognitive Neuroscience*, 34(4):699–714, March 2022.

[62] Geoffrey M. Boynton, Stephen A. Engel, Gary H. Glover, and David J. Heeger. Linear Systems Analysis of Functional Magnetic Resonance Imaging in Human V1. *Journal of Neuroscience*, 16(13):4207–4221, July 1996. Publisher: Society for Neuroscience Section: Articles.

[63] Abigail S. Greene, Corey Horien, Daniel Barson, Dustin Scheinost, and R. Todd Constable. Why is everyone talking about brain state? *Trends in Neurosciences*, 46(7):508–524, July 2023. Publisher: Elsevier.

[64] Bradley Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Chapman and Hall/CRC, New York, May 1994.

[65] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.3800020109.

[66] K. J. Friston, K. J. Worsley, R. S. J. Frackowiak, J. C. Mazziotta, and A. C. Evans. Assessing the significance of focal activations using their spatial extent. *Human Brain Mapping*, 1(3):210–220, 1994. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.460010306.

[67] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and B T Thomas Yeo. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28(9):3095–3114, 07 2017.

[68] David G. Luenberger. *Introduction to dynamic systems: theory, models, and applications.* Wiley, New York, 1979.

[69] Rex Clark Robinson. *An Introduction to Dynamical Systems: Continuous and Discrete.* American Mathematical Soc., 2012. Google-Books-ID: 7VvhilAs3JIC.

[70] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes, December 2022. arXiv:1312.6114 [cs, stat].

[71] Ulrike von Luxburg. Clustering Stability: An Overview. *Foundations and Trends® in Machine Learning*, 2(3):235–274, April 2010. Publisher: Now Publishers, Inc.