Title:

Slow ramping emerges from spontaneous fluctuations in spiking neural networks

Authors: Jake Gavenas[1,2], Ueli Rutishauser[5,6,7,8], Aaron Schurger[1,3,9,10], Uri Maoz[1,2,3,4,8,11]

[1]Institute for Interdisciplinary Brain and Behavioral Sciences, Chapman University, Orange, CA, USA
[2]Schmid College of Science and Technology, Chapman University, Orange, CA, USA
[3]Crean College of Health and Behavioral Sciences, Chapman University, Orange, CA, USA
[4]Fowler School of Engineering, Chapman University, Orange, CA, USA
[5]Department of Neurosurgery, Cedars-Sinai Medical Center, Los Angeles, CA, USA
[6]Department of Neurology, Cedars-Sinai Medical Center, Los Angeles, CA, USA
[7]Center for Neural Science and Medicine, Department of Biomedical Sciences, Cedars-Sinai Medical Center, Los Angeles, CA, USA
[8] Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA, USA
[9]INSERM U992, Cognitive Neuroimaging Unit, NeuroSpin Center, Gif sur Yvette 91191, France
[10]Commissariat à l'Energie Atomique, Direction des Sciences du Vivant, I2BM, NeuroSpin Center, Gif sur Yvette 91191, France
[11]Anderson School of Management, University of California, Los Angeles, CA, USA

Keywords:
self-initiated action, spontaneous movement, volition, readiness potential, slow-ramping signal, medial frontal cortex, slow synapses, spiking neural network, clustered connectivity, spontaneous neural activity

Highlights:

1. We reveal a mechanism for slow-ramping signals before spontaneous voluntary movements.
2. Slow synapses stabilize spontaneous fluctuations in spiking neural network.
3. We validate model predictions in human frontal cortical single neuron recordings
4. The model recreates the readiness potential in an EEG proxy signal.
5. Neurons that ramp together had correlated activity before ramping onset.

Abstract

The capacity to initiate actions endogenously is critical for goal-directed behavior. Spontaneous voluntary actions are typically preceded by slow-ramping medial frontal cortex activity that begins around two seconds before movement, which may reflect spontaneous fluctuations that influence action timing. However, the mechanisms by which these slow ramping signals emerge from single-neuron and network dynamics remain poorly understood. Here, we developed a spiking neural network model that produces spontaneous slow ramping activity in single neurons and population activity with onsets ~2 seconds before threshold crossings. A key prediction of our model is that neurons that ramp together have correlated firing patterns before ramping onset. We confirmed this model-derived hypothesis in a dataset of human single neuron recordings from medial frontal cortex. Our results suggest that slow ramping signals reflect bounded spontaneous fluctuations that emerge from quasi-winner-take-all dynamics in clustered networks that are temporally stabilized by slow-acting synapses.

Introduction

Humans and other animals can initiate actions spontaneously, devoid of immediate external triggers. Such spontaneous voluntary actions are critical to purposeful goal-directed behavior,

but the neural mechanisms underlying this kind of action are poorly understood. Studies investigating the neural precursors of spontaneous voluntary actions have found that movement onset is preceded by neural signals that slowly ramp up or down more than 2 seconds before movement onset. This phenomenon is most evident in human medial frontal cortex and has been captured by scalp electroencephalography (EEG), where it is termed the readiness potential or "RP"[1–3], as well as by fMRI[4–6], and in single-neuron firing rates[7], where it is termed the readiness discharge. Other mammals also exhibit slow-ramping signals before spontaneous actions in analogous brain regions[8–10]; Similarly, slow ramping has also been reported in invertebrates[11]. This therefore suggests that slow ramping is an evolutionarily conserved mechanism related to self-initiated action. While widely reported, the origin and significance of slow ramping signals remain debated.

Slow-ramping signals were originally interpreted as a reflection of preparatory processes that begin at the onset of ramping[2,12]. However, more recently, it has been proposed that slow ramping signals, like the RP, may instead be the result of autocorrelated fluctuations in neural activity that trigger movement upon crossing a threshold[13–17] or bias the precise time of movement onset[18]. In contrast to classic interpretations, these models (hereafter collectively referred to as "stochastic fluctuation models" or SFMs) ascribe no special meaning to the onset of the slow-ramping signals. Rather, SFMs posit that ramping signals appear only because of back-averaging autocorrelated signals aligned to movement onset. SFMs therefore offer drastically different neural and cognitive implications for action initiation than classic interpretations[19]. However, a major challenge for SFMs is that the time constant of experimentally-observed autocorrelations in neural spike trains are much shorter than the observed rate of ramping by about an order of magnitude (hundreds of milliseconds for autocorrelation[20,21] versus several seconds for ramping in some neurons[7]). SFMs predict that more highly autocorrelated fluctuations will approach the threshold more gradually and thus take longer to ramp up, while less autocorrelated (and less stable) processes are more likely to "jump up" and cross the threshold more abruptly[13]. Remarkably, at the single-neuron level, slow ramping on the time scale of seconds has been observed in human medial frontal cortex[7]. Neurons in this brain area have autocorrelation time constants that are much shorter than this time of ramping[20]. Thus, the emergence of slow ramping, if tied to autocorrelated fluctuations, likely stems from some currently unknown network-level mechanism.

Here, we investigate the single-neuron and network-level mechanisms that lead to slow-ramping signals via a combination of spiking-neural network simulations and new analysis of human single-neuron experimental data[7]. These data were recorded[7] from human medial frontal cortex during a spontaneous voluntary movement task similar to the classic Libet task[2]. Here, we show that spiking neural networks can be configured to produce spontaneous fluctuations with properties similar to those seen in our experimental data. Our model is a spiking neural-network with 400 units with clustered connectivity. Such networks can exhibit spontaneous fluctuations[22,23], but prior studies found that neurons in these networks exhibit "abrupt switching" activity—i.e. abrupt switches between slow and fast firing—which is unlike the gradual changes seen during slow ramping. We hypothesized that we could stabilize activity by inserting slow synapses, which are critical for producing temporally stable activity [24,25] and increased autocorrelation in simulated spiking neurons [26]. We show that the presence of slow synapses temporally stabilized spontaneous fluctuations and facilitated the emergence of slow

ramping. The properties of ramping resembled the ramping signals we saw in our experimental data both at the single-neuron and population level. Furthermore, our model makes novel predictions related to correlation structure before ramping onsets. We validated the model-derived hypotheses in the aforementioned experimental data. We conclude that slow ramping signals before self-initiated action likely emerge from spontaneous fluctuations aligned to threshold-crossings, and that clustered connectivity and slow synapses are the key mechanism that facilitates those fluctuations. These results are compatible with SFMs of spontaneous voluntary action generation.

Results

Synaptic dynamics temporally stabilize spontaneous fluctuations in simulated networks.

We simulated networks consisting of N=400 spiking leaky integrate-and-fire neurons; 80% of the neurons were excitatory, i.e., they provided only positive inputs to other neurons, while the remainder were inhibitory. Excitatory neurons were grouped into 4 clusters of 80 neurons each, such that within-cluster connections were more likely than between-cluster connections. The extent of clustering is quantified by the parameter R_{EE} (clustering degree), which is equal to the ratio of the probability of a within-cluster connection to the probability of a between-cluster connection (within-cluster connections were also stronger compared to between-cluster connections as in[22]; see Connectivity Matrix Construction in Methods). Inhibitory neurons were not clustered, and their connectivity probability with excitatory neurons did not depend on R_{EE} [27].

R_{EE} describes a continuum between homogenous networks without clustering ($R_{EE} = 1$) and winner-take-all networks with strong clustering ($R_{EE} >> 1$) (Figure 1A). Homogenous networks with balanced excitation and inhibition exhibit asynchronous spiking activity without slow spontaneous fluctuations in firing rates[28]. Conversely, in winner-take-all networks, neurons in a single cluster dominate and suppress neurons in other clusters via lateral inhibition[29]. In between these two extremes, clustered networks (R_{EE} roughly between 2-5, depending on network size, number of clusters, and other factors) exhibit slow spontaneous fluctuations due to a combination of lateral inhibition and lateral excitation between clusters. Note that in our networks, spontaneous fluctuations were entirely due to internal interactions because external input to the network was fixed (besides noise).

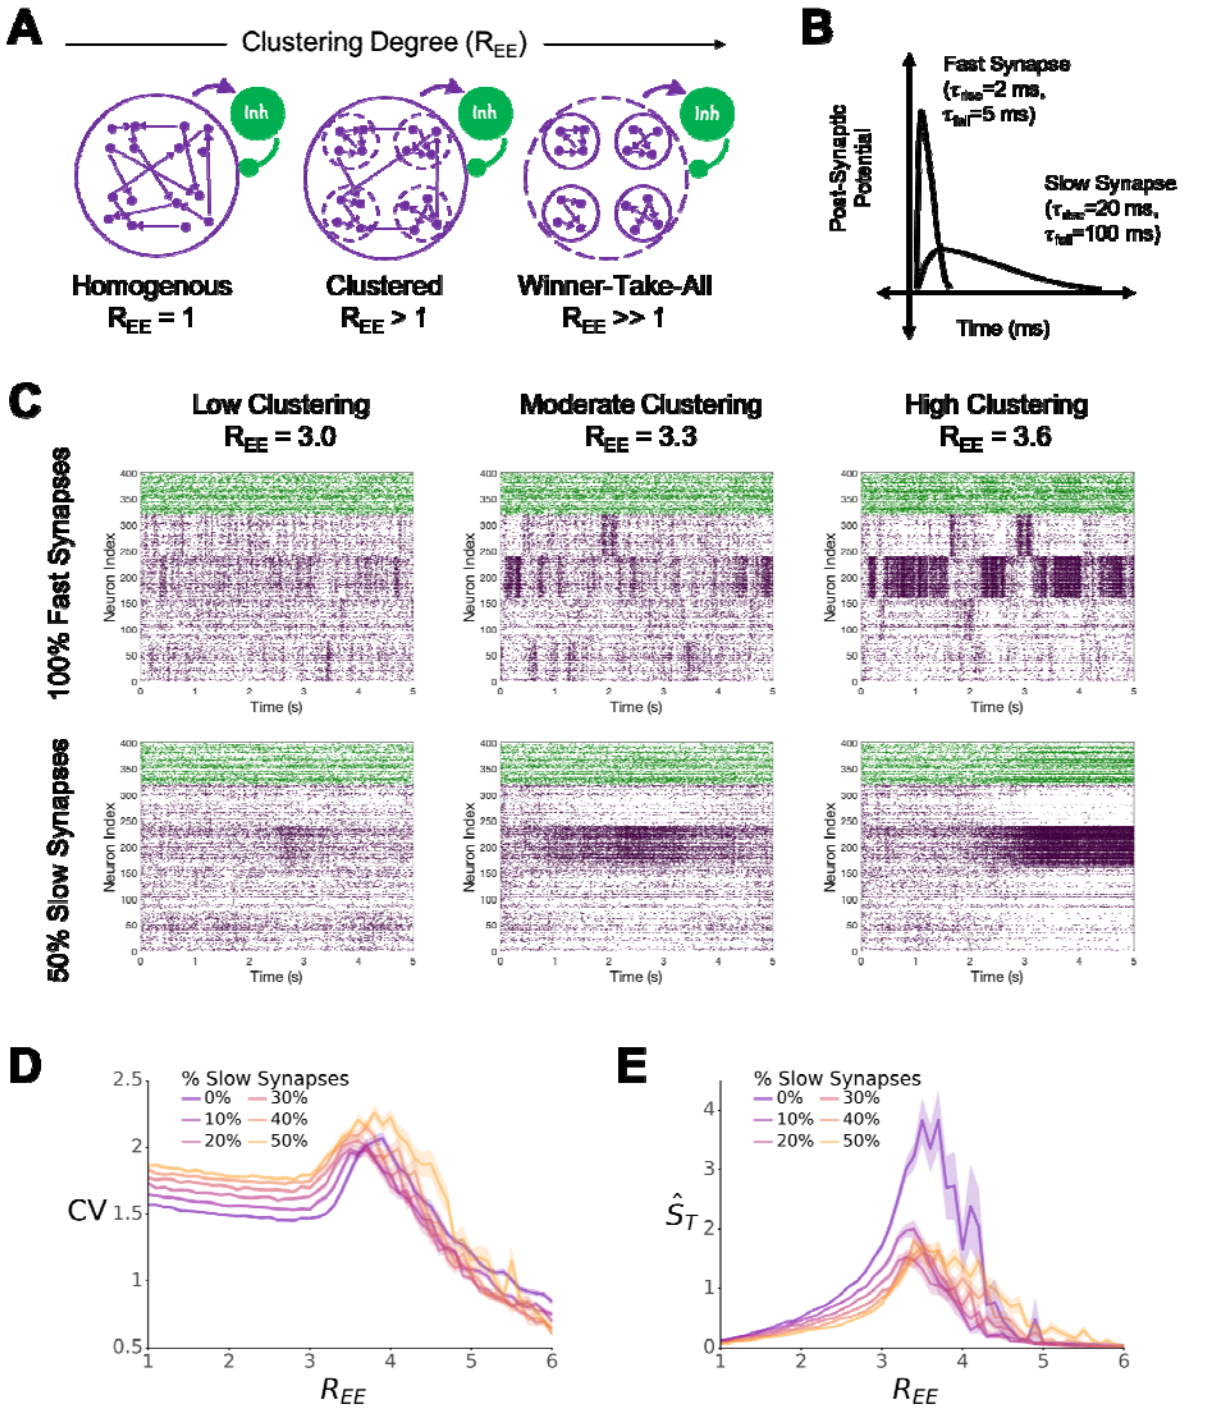

Figure 1: Effects of clustering & synaptic dynamics on spontaneous fluctuations in simulated networks. A: As the clustering degree (R_{EE}) increases, connectivity transitions from homogenous/balanced, through clustered connectivity, to winner-take-all. **B:** Schematic of fast vs. slow synapses. **C:** Example 5-second simulations for networks with 100% fast synapses (top row) and 50% slow synapses (bottom row) at various clustering degrees (same underlying network structure). Purple are excitatory neurons (4 clusters of 80 neurons), green are inhibitory neurons. Clustering leads to spontaneous fluctuations, but fluctuation characteristics depend on synaptic dynamics. Networks with 100% fast synapses switch between high- and low-activity states, whereas networks with 50% slow

synapses show smoother, more temporally stable fluctuations. **D:** Firing rate coefficient of variation (CV) across time for different R_{EE} and synaptic ratios (Mean ± standard error across n=50 simulations each). The CV characterizes fluctuations in firing rates over time. Slow synapses result in higher CV for low to moderate R_{EE} values. CV sharply increases in the fluctuation regime ($3 < R_{EE} < 4$) and then decreases as networks transition to a winner-take-all regime. **E:** $\hat{S}_T$ across time for different R_{EE} and synaptic ratios (Mean ± standard error across n=50 simulations each). $\hat{S}_T$ characterizes spiking variability over time and detects slow-switching behavior. Slow synapses result in lower $\hat{S}_T$ for low to moderate R_{EE} values, with a pronounced effect in the fluctuation regime ($3 < R_{EE} < 4$), suggesting networks with slow synapses fluctuate without slow-switching behavior. Like CV, $\hat{S}_T$ sharply increases in the fluctuation regime and then decreases as networks transition to winner-take-all behavior.

Previous work had only simulated clustered networks with fast synaptic dynamics, in which the entire postsynaptic potential is delivered within ~5 ms [22,23,30]. We were interested in the role of slow synapses in modulating spontaneous fluctuations in such networks. We therefore simulated networks with both fast and slow synapses. We used double-exponential functions to model synapses[22,31], with fast synapses having a rising and falling time-constant of 2 and 5 ms respectively, and slow synapses having a rising and falling time-constant of 20 and 100 ms respectively[31] (Figure 1B). Because different types of synapses engage either fast or slow dynamics (AMPA and GABA-A for fast, NMDA and GABA-B for slow), we modeled mixtures of synapses as a weighted sum of these two types of synapses, parameterized as the ratio of fast to slow synaptic signaling (i.e., the percent of post-synaptic potential (PSP) delivered via fast versus slow synapses; note that this applied to both excitatory and inhibitory connections; see LIF Network Simulations in Methods).

Recreating prior results for clustered networks with 100% fast synapses, we found that increasing R_{EE} was associated with a transition from asynchronous activity to more winner-take-all dynamics, with networks in an intermediate range exhibiting slow fluctuations (Figure 1C top row). Notably, fluctuations manifested as bistable switching activity, with entire clusters switching between states of low and high activity. These fluctuations exhibited quasi-winner-take-all dynamics in that activations in one cluster suppress firing in all the other clusters. Notably, switches between these states were rapid with no apparent slow ramping (Fig. 1C, top row; see Fig. 2 for quantification).

In contrast, networks with slow synapses exhibited slow ramping: with 50% slow synapses, fluctuations in networks with intermediate R_{EE} values were temporally stable, with rises and falls that occur over a longer timespan when compared to networks with 100% fast synapses (Figure 1C bottom row; see below for quantification). Fluctuations in 50% slow synapse networks still exhibited quasi-winner-take-all dynamics in that activations in one cluster suppress other clusters, but without the all-or-none nature of fluctuations in 100% fast synapse networks. Note that fluctuations are less clear in the inhibitory population (Figure 1C green raster), and henceforth we exclusively analyze properties of the excitatory neurons.

To quantify the slow fluctuations we observed, we calculated two measures of spiking variability. We first calculated the coefficient of variation (CV) for individual neurons' firing rates across time. The CV equals the standard deviation of firing rates across time (average firing rates in 100 ms non-overlapping bins; see Simulating & Analyzing Fluctuations in Methods), normalized by dividing by the trial-average firing rate. High CV values are expected for

networks that exhibit slow fluctuations in firing rates, as those fluctuations entail higher firing rate standard deviations across time for the same mean firing rate. We simulated 5 seconds of activity in 50 networks for R_{EE} values between 1 and 6 with synaptic ratios ranging from 0% slow (100% fast) to 50% slow. We calculated the average CV across excitatory neurons after discarding the first 500 milliseconds. CV sharply increased after a critical R_{EE} value (around 3-3.5) for all networks, suggesting the emergence of slow fluctuations past that point (Figure 1D). CV then decreased for higher R_{EE} values, indicating the transition to a winner-take-all regime. Notably, a greater percentage of slow synapses consistently resulted in a higher CV, suggesting that slow synapses increase firing rate variability even without clustering.

We next calculated $\hat{S}_T$, a measure of spiking variability[23] that characterizes switching behavior in clustered networks, specifically like those observed in fast-synapse networks (Fig. 1C top row). $\hat{S}_T$ equals the standard deviation of cluster-averaged firing rates (100 ms non-overlapping bins) across time, averaged across clusters, normalized by subtracting the same measure applied to random groupings of neurons rather than by cluster identity. $\hat{S}_T$ therefore quantifies variability across time and slow-switching activity. High values of $\hat{S}_T$ indicate slow fluctuations specifically of the slow-switching variety (e.g. Figure 1C top middle-right). $\hat{S}_T$ increased steadily as R_{EE} increased, reaching maximal values for R_{EE} between 3 and 4 (Figure 1E), and then decreasing. Notably, a greater percentage of slow synapses resulted in lower $\hat{S}_T$ values. These results indicate that networks with slow synapses exhibit slow fluctuations, but that these fluctuations manifest as gradual fluctuations rather than slow-switching activity. See Supplementary Figure 1 for quantification of heterogeneity across networks. Taken together, clustering and slow synapses jointly facilitate the emergence of spontaneous yet temporally stable fluctuations. We next turned to quantifying the speed of ramping in this network architecture.

Slow-ramping at the single-neuron level

We next examined the relation between spontaneous fluctuations and slow ramping signals at the single-neuron level. We simulated activity in 20 randomly connected networks with intermediate clustering degrees that led to spontaneous fluctuations (Min R_{EE} = 2.8; Max R_{EE} = 4; manually tuned for each network so that dynamics were in the fluctuation regime and not the winner-take-all regime; see Network Selection in Methods). We simulated 100 trials of 10 seconds each for each network when synaptic dynamics were either 100% fast or 50% slow. We then aligned data to times when activity in a specific cluster increased their activity above a threshold (termed threshold-crossing, defined as 50% of maximum normalized firing rate across the whole trial; see Threshold Alignment in Methods). Neurons that significantly increased or decreased their firing rates in the 400 milliseconds before threshold-crossing relative to a baseline (-3 to -2 seconds; as in Fried et al. [7]) were categorized as increasing or decreasing, respectively (rank-sum T-test, the cutoff was $p<0.01$, as in Fried et al.[7]). The ratio of neurons categorized as increasing to decreasing was 50:50 in 100% fast synapse networks and 39:61 in 50% slow synapse networks.

Networks with 50% slow synapses demonstrated slow ramping with increasing and decreasing of firing rates before threshold-crossing, whereas networks with fast synapses showed only a steep jump shortly before threshold-crossing (Figure 2A). Notably, our networks did not have a specified ‘ramping onset,’ nor were there any changes to the network’s input after trial onset (which was many seconds before the spontaneous ramping onset) beyond those attributable to noise. Therefore, the observed ramping must have occurred due to spontaneous fluctuations. To

further explore the relationship between slow synapses and ramping behavior, we conducted simulations (same as above) for various ratios of fast-to-slow synapses. Those simulations revealed that slow ramping emerged gradually as we increased the proportion of slow synaptic signaling (Figure 2B). Similarly, we find that the larger the proportion of slow synapses, the longer the autocorrelation of neuronal activity (see Supplementary Figure 2A). Together, these simulations support the notion that slow synapses facilitate temporally stable autocorrelated activity.

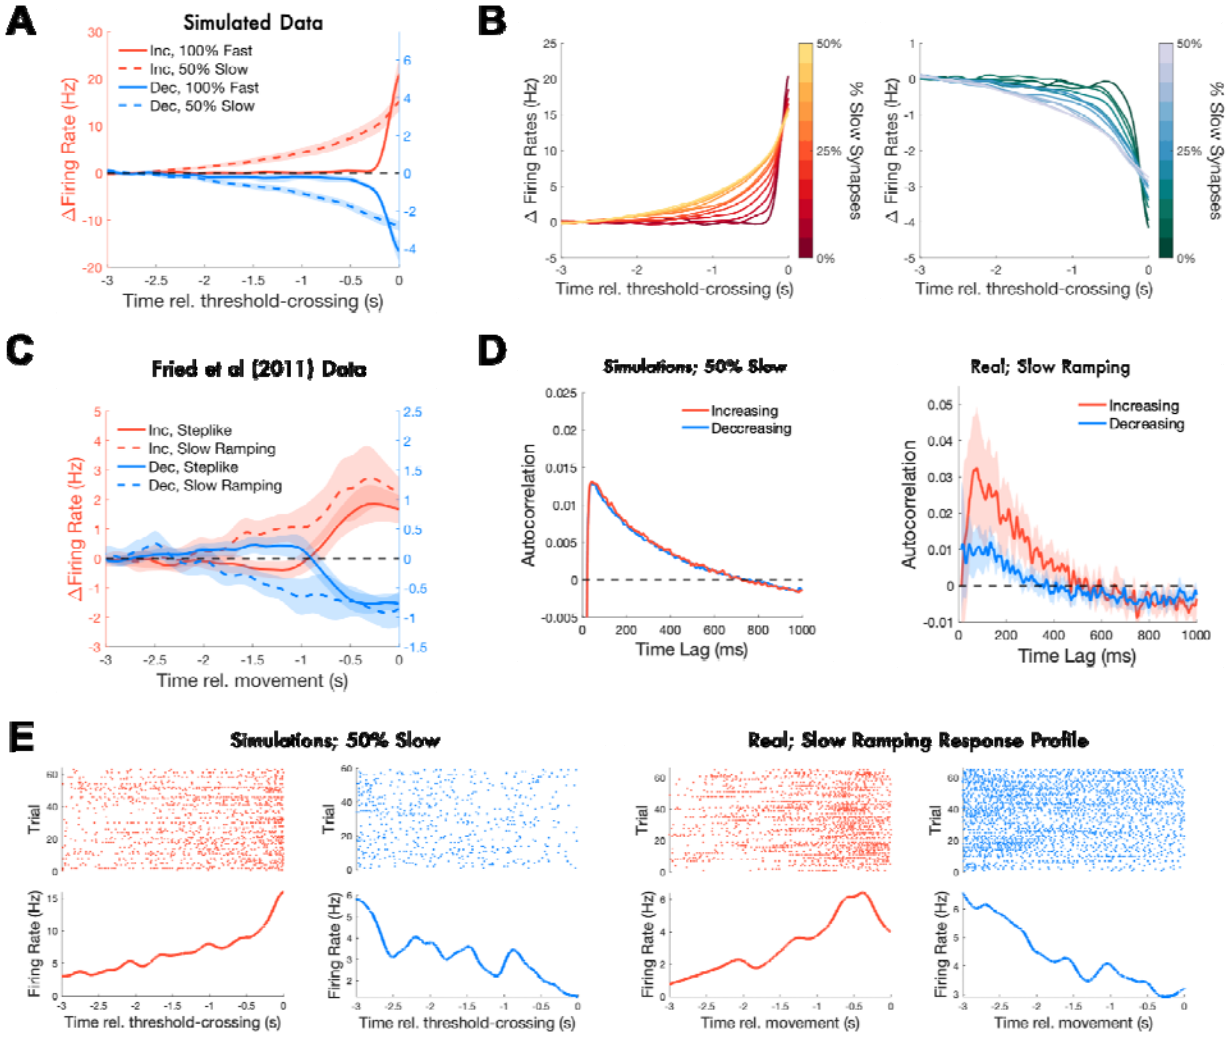

Figure 2: Slow ramping at the single-neuron level in simulated networks and recorded human frontal neurons. A: Firing rates (Δ relative to baseline; mean ± 95% confidence intervals across networks) of neurons that increased (inc.; red) or decreased (dec.; blue) activity before threshold-crossings in networks with 100% fast (solid lines) or 50% slow synapses (dashed lines). **B:** Slow ramping of firing rates gradually emerges with the presence of slow synapses, increasing and decreasing on the left and right, respectively. **C:** Slow ramping before movement in human frontal neurons (Fried et al.[7] dataset; mean ± 95% confidence intervals across neurons). Some neurons show a steplike response profile (solid lines), and others show slow ramping (dashed lines). **D:** Slow ramping emerges despite low autocorrelation at the single-neuron level in both simulated (left; mean) and real neurons that exhibit slow ramping (right; mean ± 95% confidence interval). **E:** Example rasters and trial-averaged firing rates from

simulated (left) and real neurons (right). Notably, slow ramping is not apparent in most single-trial responses.

Comparison of single-neuron level slow ramping in simulated and experimental data

We next assessed whether the degree of slow ramping observed in the model is comparable to that typically seen in the motor system preceding the onset of volitional actions. One of the key brain areas within which such slow ramping has been observed is in medial frontal cortex (MFC) of humans, including the pre-supplementary motor area, supplementary motor area proper, and anterior cingulate cortex[3,32,33]. We examined whether the properties of empirically measured ramping are comparable to those seen in our model by re-analyzing human MFC single-neuron data recorded by Fried and colleagues[7] that show slow ramping and other firing rate changes before spontaneous voluntary actions.

We analyzed n=512 MFC neurons. Of these, 153 (30%) neurons significantly changed their response relative to movement onset (Using the same selection criteria we used for our model data; note that the original analysis aligned data to the time of declared intention rather than movement onset). Of these, 51 had increasing, and 102 had decreasing firing rates before movement onset relative to baseline. We grouped the variety of MFC neuron response profiles into slow ramping and steplike groups, which exhibited gradual and sudden changes of firing rates, respectively. We distinguished between response profiles by fitting a sigmoid to trial-averaged firing rates and extracting an inverse-gain parameter. Neurons with an inverse-gain >0.4 were categorized as slow ramping, and otherwise as steplike (see Neuron Categorization and Response Profiles in Methods). 41 MFC neurons were slow-ramping (27% of all ramping MFC neurons; 15 increasing and 26 decreasing) and 112 were steplike (73% of all ramping MFC neurons; 36 increasing and 76 decreasing). Notably, slow ramping neurons showed progressive changes to their firing rates beginning around 2 seconds before movement (see Supplementary Figure 2C for distribution of ramping onset times), while steplike neurons showed a sharper firing rate change around 1 second before movement (Figure 2C).

Comparing the response profiles of simulated and recorded MFC neurons reveals that neurons in the slow ramping group had response profiles that were similar to those in simulated networks with slow-synapses (cf Figure 2A vs Figure 2C). There were also similarities between the step-like pre-threshold responses in 100% fast synapse networks and the MFC neurons with step-like response profiles. Here, our focus will be on the slow-ramping responses in MFC (Aflalo and colleagues[34] analyzed step-like responses in parietal cortex).

The autocorrelation of slow ramping neurons decayed to zero after ~600 milliseconds (Figure 2D; estimated during the ramping period of -3 seconds to movement or threshold crossing). In contrast, slow ramping occurred over periods of time 2-4 times longer than the time for which neurons had above-zero autocorrelation. An investigation of single-trial data reveals how this network-level phenomena emerges. Although slow ramping in trial-average response profiles is apparent for both simulated and real neurons, a smooth increase in firing is not apparent on many individual trials (examples in Figure 2E; more examples in Supplementary Figure 3; this fact is also noted by Schurger et al.[19]). Rather, movement is triggered after an entire cluster's average activity crosses a threshold. Therefore, individual neurons can and will show variable activity on individual trials. Still, trial-averaged response profiles will show smooth ramps because their

activity is, across many trials, entrained to a cluster which either increases or decreases net firing before threshold-crossings.

Slow ramping at the population-level and comparison to experimental data

We next investigated slow ramping at the population level of all recorded/simulated neurons. We trained a linear discriminant analysis classifier (LDA) to differentiate between 'pre-ramping' and 'baseline' (-3.0 to -2.6 seconds) at a given point of time. Classification accuracy for 50% slow synapse networks reached 75% 641±7 milliseconds (average across networks) before threshold-crossing, while classification accuracy for fast synapse networks only increased above 75% on average 31±30 milliseconds before threshold-crossing (Figure 3A shows results from 100 simulations on a single network, aligned to threshold-crossings as in Figure 2A; shown are mean ± standard error across 5 cross-validations; further examples in Supplementary Figure 3A). Classification accuracy in the 50% slow synapse network resembled the gradual rise in accuracy obtained by an analogous analysis of the Fried et al. (2011) dataset (Figure 3A right), where decoding accuracy reached 75% at 1.2 seconds before movement onset. We conducted a similar analysis on the recorded MFC neurons, pooling across participants using a pseudopopulation for decoding[35]. Decoding accuracy reached ~90% at movement onset, presumably because many real neurons show a sudden 'step like' change in their firing rates at around 0.5s before movement onset (see Figure 2C). Our model does not contain such neurons, which are likely related to motor preparation. These results suggest that our simulations capture the early, gradually increasing accuracy in decoding activity from baseline up to 0.5s before movement onset (Fig. 3A, left vs. right plot).

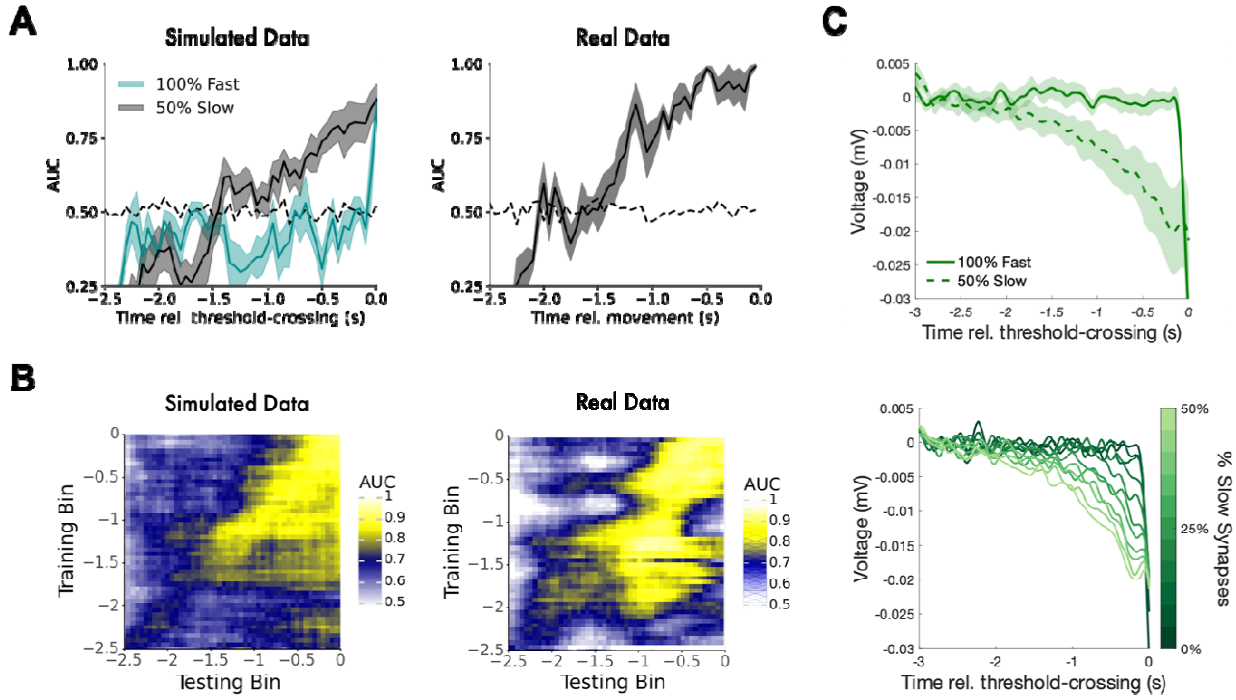

Figure 3: Slow ramping at the population level; decoding analysis & EEG proxy. A: Classification accuracy versus baseline (linear discriminant analysis; spike count in 400 ms window with shown times referring to the leading edge of that window, versus -3 to -2.6 sec baseline) in simulated (left) and real data (right, shaded region is 95% confidence interval across cross-validations; dashed line is chance level

across 100 shuffles of the data). Accuracy slowly ramps before threshold-crossing / movement, suggesting slow ramping at the population level. **B:** Performance of LDA classifiers trained at one time bin and tested on another time bin (relative to threshold-crossing or movement). Left: simulated data. Right: real data. Accuracy gradually increased leading up to the time of threshold-crossing or movement and was higher when the training bin was earlier than the testing bin compared to vice-versa (above versus below the main diagonal). **C:** Top: EEG proxy (PSPs onto excitatory populations; insert) shows slow ramping in 50% slow but not 100% fast networks (confidence intervals over networks). Bottom: slow ramping in EEG proxy signal gradually emerges with an increased percentage of slow synaptic signaling.

To investigate whether gradually increasing accuracy reflected a monotonic drift away from baseline firing rates, we investigated the temporal generalization of decoders, which can be used to ascertain the temporal evolution of neural activity[36]. When classifying activity against a baseline, decoders trained on activity that is on average close to baseline in neural state space (i.e. close in time to the baseline) will find a separatrix that is near baseline, which will perform well when applied to activity further away in neural state-space. However, decoders trained on activity that is far from baseline in neural state space (i.e. later in time) will find a separatrix that is far from baseline activity, which will perform poorly when applied to activity closer to baseline in neural state-space (i.e. earlier in time), Thus, we should observe an asymmetry in across-time generalization accuracy such that classifiers perform better when generalizing from earlier to later timepoints compared to the reverse.

To test this theoretical prediction, we trained an LDA classifier to differentiate between ramping activity at a given time point relative to threshold crossing and activity during a baseline period (and tested that classifier's accuracy at different time points (see Decoding Analyses in Methods). For simulated data, decoding accuracy gradually increased leading up to threshold-crossing (Figure 3B left; more examples in Supplementary Figure 4B). Notably, as expected for ramping signals, across-time generalization accuracy was higher when training on earlier and testing on later timepoints compared to when training on later and testing on earlier timepoints (Fig. 3B). This can be seen by comparing the average values in the upper vs. lower triangle of the temporal generalization matrix (one-tailed rank-sum t-test $p < 0.001$). Similarly, for the MFC neurons, decoding accuracy showed a similar increase and asymmetry (rank-sum t-test comparing upper triangle to lower triangle decoding accuracies, $p < 0.001$, Figure 3B right; see Supplementary Figure 4C for explicit demonstration of asymmetry). To verify that this asymmetry reflects ramping, we generated spike trains that slowly ramped up and down to create a pseudopopulation of neurons that we know a priori exhibit linear ramping (n=20, 50 trials, see Simulated Ramping in Methods). We conducted the same decoding and temporal generalization procedure on those spike trains and found the same asymmetry that decoders generalized to future time points better than to past time (one-tailed rank-sum t-test $p < 0.001$; Fig. S4D).

We also investigated ramping at the population level by calculating an EEG proxy from our network to compare to the readiness potential, a negative deflection in trial-averaged scalp EEG that slowly builds up over the last ~1 second prior to the time of movement [3,19]. Notably, we focus only on the early (negative) rise of the readiness potential, which is thought to originate in bilateral SMA, relative to the later parts closer to movement onset, which originate in lateral motor cortex[3]. The EEG signal is thought to reflect the summation of post-synaptic potentials onto neurons[37]. So, as a proxy signal, we calculated the total signed post-synaptic input (both

inhibitory and excitatory) delivered to each excitatory neuron in the network, then averaged across neurons, and investigated changes before threshold-crossings (see Threshold Alignment in Methods). Networks with 50% slow synapses exhibited slow negative ramping in the EEG proxy. In contrast, networks with 100% fast synapses showed no changes until close to threshold-crossing (Figure 3C top). Repeating this analysis for multiple ratios of fast-to-slow synaptic signaling revealed that slow ramping in the EEG proxy signal gradually emerged with a greater percentage of slow synapses (Fig 3C bottom). Results were similar for an EEG proxy that uses PSPs onto all simulated neurons rather than just excitatory neurons (Supplementary Figure 4E), accounting for suggestions that human cortex contains larger proportions of inhibitory neurons than other species[38].

Pre-ramping correlation structure relates to slow ramping activities

Our model, like other SFMs, predicts that slow ramping is a consequence of averaging autocorrelated fluctuations aligned to times of threshold-crossings. Those fluctuations do not begin at ramping onset, but occur throughout the entire trial (see Figure 1C bottom). Therefore, a novel prediction of our model is that two neurons that show ramping behavior in the same direction (both increasing or decreasing) will have more correlated activity *before ramping onset* in comparison to pairs that show ramping behavior in different directions (one increasing and one decreasing). This is because ramping behavior relates to cluster membership—two neurons that both increase or decrease their firing rates are likely part of the same cluster, whereas two neurons that show opposite ramping directions likely belong to different clusters (Figure 4A left). In contrast to SFMs, traditional conceptual models of the RP and other slow ramping signals suggest ramping is initiated by an external input after some neural or cognitive event[2] (Implicit Decision Model). Such models would not make any particular prediction about correlations before ramping onset, because ramping onset would be triggered by activity in, presumably, a different brain region (Figure 4A right).

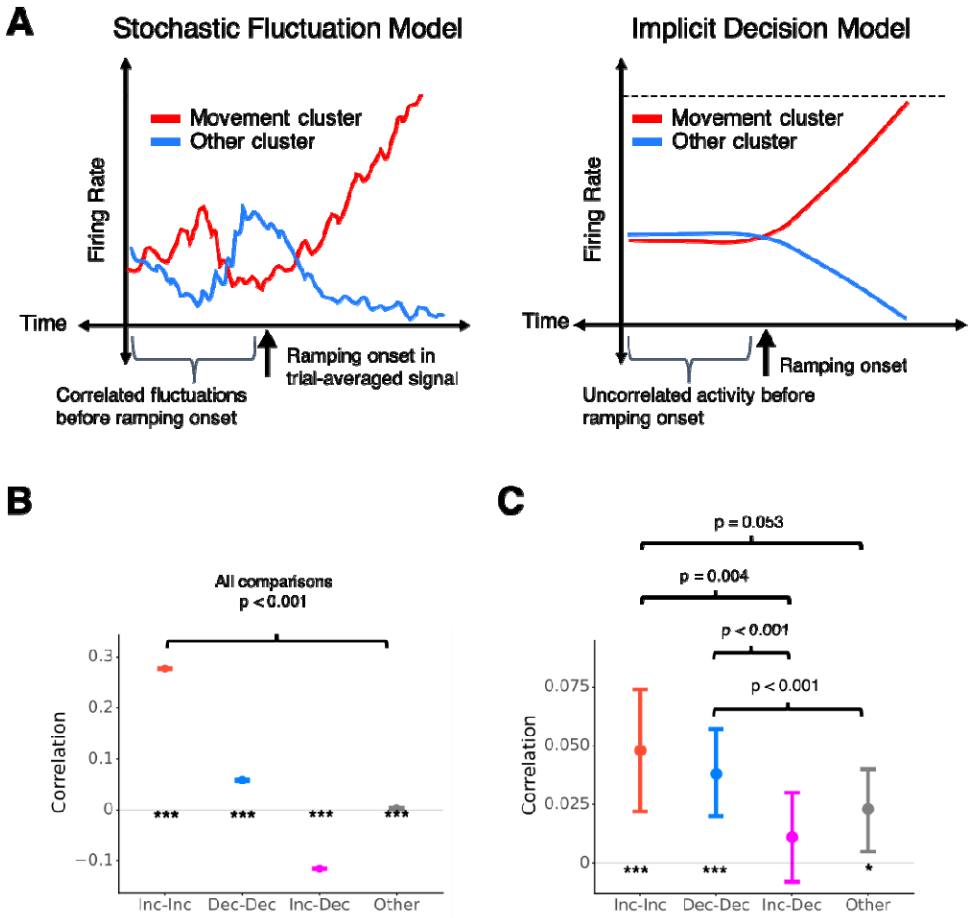

Figure 4: Model schematic and correlation structure relates to ramping direction. **A:** Schematic of two models and their diverging predictions about pre-ramping correlations. Left: fluctuation-based model, where ramping reflects autocorrelated fluctuations aligned to threshold crossing. Because the fluctuations occur throughout a trial, neurons with the same ramping behavior should demonstrate correlated fluctuations even before ramping onset. Right: input-based model, where ramping is determined by some input at ramping onset. Because the input occurs only at ramping onset, correlations before ramping onset are not predicted. **B:** In 50% slow synapse networks, pairs of simulated neurons that exhibited the same ramping behavior (increasing or decreasing) were significantly correlated even before ramping onset. Furthermore, neurons that showed opposite ramping behaviors were anticorrelated at baseline. **C:** Pairs of real neurons that showed the same ramping behavior (increasing or decreasing) were significantly correlated before ramping onset (-3.5 to -2.5 s relative to movement; steplike response were categorized as "other"; means & 95% confidence intervals, linear mixed effects). Pairs of neurons that exhibited the same type of ramping (both increasing or decreasing) were more correlated at baseline compared to pairs that exhibited opposing types of ramping (increasing vs. decreasing), although pairs showing opposite ramping behaviors were not anticorrelated, contrary to simulations.

As predicted, in 50% slow synapse networks, activity in pairs of increasing neurons (i.e., both belonging to the specific cluster associated with movement) was positively correlated during a pre-ramping baseline period [-3,-2.5] (Figure 4B, $\rho = 0.277$, $p < 0.001$; see Analysis of Pairwise Correlations in Methods). In addition, activity in pairs of decreasing neurons (i.e., both belonging to other clusters not associated with movement) was also positively correlated, though to a lesser extent than pairs of increasing neurons ($\rho = 0.059$, $p < 0.001$). In contrast, activity in pairs

involving one increasing and one decreasing neuron (i.e., very likely belonging to different clusters) was negatively correlated ($\rho = -0.115$, $p < 0.001$). Finally, some neurons neither increased nor decreased their activity before threshold-crossing, likely due to the random nature of network connectivity leading to them receiving heterogeneous inputs. Activity in pairs involving such neurons had very low but significantly positive correlations before ramping onset ($\rho = 0.003$, $p < 0.001$).

To differentiate between the two models, we next examined pairwise correlations between firing rates in the pre-ramping baseline period [-3.5 s,-2.5 s] for simultaneously recorded MFC neurons (i.e. within the same session; see Analysis of Pairwise Correlations in Methods; Figure 4C). Pairs of neurons that later exhibited the same type of ramping (i.e., both increasing or both decreasing) were significantly positively correlated before ramping onset (increasing-increasing N=108: $\rho = 0.048$, $p < 0.001$; decreasing-decreasing N=865: $\rho = 0.038$, $p < 0.001$), whereas pairs with opposing types of ramping were not significantly correlated at baseline (N=532 pairs; $\rho = 0.011$, $p = 0.285$; Fig. 4C). Further, same-type pairs were more correlated than opposite-type pairs (increasing-increasing vs. increasing-decreasing: $p = 0.004$; decreasing-decreasing vs. increasing-decreasing: $p < 0.001$) and were also more correlated than pairs involving neurons that did not increase or decrease or had a steplike response profile (N = 18,384 pairs; increasing-increasing vs. other: $p = 0.053$, only trending; decreasing-decreasing vs. other: $p < 0.001$). It is worth noting that correlations in the experimental data were smaller in magnitude overall compared to our model predictions, and opposite-type and other pairs had more positive correlations than the model (e.g. opposite: $\rho_{simulated} = -0.115$, $\rho_{data} = 0.011$). Together, this analysis therefore qualitatively confirms the prediction made by a SFM (Figure 4A left), indicating that the slow ramping signals likely emerge from correlated fluctuations that are always present (compare Figure 4B and C).

Discussion

We simulated spiking neural networks with clustered connectivity and a mixture of fast and slow synapses. We found that slow synapses temporally stabilized neuronal activity and, in the aggregate, facilitated the emergence of slow ramping before threshold-crossings (Fig. 1). In networks with 50% slow synapses, simulated neurons slowly ramped their firing rates, with 61% of slow ramping neurons increasing their firing rates. The accuracy of classifiers trained to distinguish activity from a baseline period slowly ramped up before threshold-crossings. Further, classifiers generalized to later time points better than they generalized to earlier time points (Fig. 3B left), indicating ramping at the population-level[36].

Furthermore, all of these features of the model matched analogous analyses that we carried out on intracranially recorded human MFC neurons[7]. In addition, an aggregate signal derived from our model, which could be interpreted as a proxy for EEG, slowly ramped negatively, akin to the readiness potential[1] (RP) in terms of timing, shape, and ramping direction. Importantly, all simulated slow-ramping signals emerged spontaneously, without necessitating an external input at ramping onset to drive changes in activity. Furthermore, individual neurons slowly ramped their firing rates—some as early as 2+ seconds before threshold-crossing or movement—despite autocorrelation diminishing to zero within 600 milliseconds for experimentally measured and simulated neurons. Finally, our model (and other Stochastic Fluctuation Models, SFMs)

predicted that neurons that show the same ramping behavior (increase or decrease) after ramping onset would exhibit elevated pairwise correlations in their activity already before ramping onset. This prediction was borne out in the MFC neurons, which showed a similar correlation structure to our model even before ramping onset, suggesting that ramping reflects correlated fluctuations aligned to threshold-crossings. That prediction contrasts with traditional conceptual models of the RP that suggest that ramping onset reflects a specific event (e.g. Implicit Decision Models[2]). If so, there should not be a relation between pre-ramping activity and later ramping behavior (Fig. 4). Hence, following the above, spontaneous fluctuations in the networks we simulated are sufficient to explain many aspects of neural activity before self-initiated action.

Our work expands on research into self-initiated actions by offering what is, to our knowledge, the first model of slow ramping in spiking neural networks. Previous work had investigated spontaneous neural activity[22,23,26,28,30] or autocorrelation in spiking neural networks[25,26,39], but did not explore the connection to slow-ramping signals. Our model provides an implementation at the level of spiking neurons of previous more abstract spontaneous fluctuation models (SFMs) of self-initiated action, demonstrating how the fluctuations that are central to those models emerge from specific neuronal dynamics[13–17]. The fact that our model can parsimoniously explain many features of MFC activity before self-initiated actions suggests that slow fluctuations in neural activity may underlie slow ramping signals and play a key role in the genesis of self-initiated action. Our model also goes beyond prior abstract computational models such as stochastic accumulation[14] by describing a process that does not terminate at threshold crossing but rather continues in time. In doing so, our model offers possibilities for investigating strings of actions rather than single actions that occur in isolation.

Critically, our results have major implications for the use of slow-ramping signals for predicting movement intention, for instance in brain-computer interfaces. Because activity can stochastically increase or decrease, movement timing isn't determined until threshold-crossing. Therefore, slow-ramping signals are not predictive of movement onset time and should therefore be of limited utility in an online real-time prediction system. Indeed, predictions based on whole-brain EEG activity could predict movement timing in real time only ~600 milliseconds before movements[40] (see also[41,42]), whereas the RP and other slow ramping signals emerge 2 seconds or more before movement onset. Furthermore, the EEG proxy we derived closely matched the RP. Our proxy is based on post-synaptic potentials, which exhibits a negative deflection before threshold-crossings due to an overall increase of inhibitory signaling and decrease of excitatory signaling as part of the quasi-winner-take-all process inherent to the networks we investigated. Taken together, we suggest that the apparent onset of the RP and other slow ramping signals several seconds before movement is the result of aligning spontaneous fluctuations to threshold-crossings rather than the many cognitive interpretations that have been offered since the RP's discovery.

Although our model suggests a mechanism for how slow-ramping signals might emerge, many questions remain about the neural mechanisms underlying self-initiated action. A key open question is regarding the mechanism that implements the threshold operation. Prior models suggested that threshold-crossing leads to movement within about 200 milliseconds[14]. However, many non-ramping neurons in MFC and other brain areas abruptly change their firing rates 0.5-1s before movement (Figure 2C), around which time activity also abruptly changes in parietal

cortex[34]. Such activations may reflect the spontaneous ignition of global activity patterns following threshold-crossings[17,43,44]. Additionally, while slow ramping activity before self-initiated action is most prominent in MFC, such signals have been observed in other regions including medial temporal cortex[7], sensory areas[6], and even subcortical areas such as locus coeruleus[45] (indexed by pupil dilations before movement). Indeed, recordings in rats find ramping activity in subcortical structures such as the dopaminergic substantia nigra[46–48]. Recent large-scale recordings of neural activity in mice further suggest that spontaneous, self-initiated movements involve widely distributed rather than localized activity[49,50], as suggested by Schurger and Uithol[51] . One possibility is that these fluctuations are related to fluctuations in arousal[52], which may explain correlations between the RP and respiration[53] and pupil dilations before freely timed movement in humans[45]. Taken together, fluctuations influencing movement timing may reflect a distributed, brain-wide process rather than one localized only in MFC.

Because slow-ramping signals before self-initiated action may reflect stochastic population-level neural fluctuations, they also offer the opportunity to validate models that describe how temporally stable fluctuations emerge from neural dynamics. Brain activity fluctuates over several seconds or tens of seconds [54], and the temporal properties of such fluctuations correlate with behavior [55] change with age[56], and are disrupted in disease[57]. It is therefore of interest to understand how these fluctuations emerge more generally. Prior work suggested that such fluctuations emerge when neural systems are poised at the edge of a phase transition[22,23,30,58–60] (i.e., at criticality). Indeed, networks exhibit spontaneous fluctuations for a critical range of clustering of connectivity, below or above which they exhibit asynchronous or winner-take-all dynamics, respectively[22,23]. However, our results suggest that other factors that are relevant to neural processing (the type of synaptic signaling in our case) can also alter the spatiotemporal characteristics of fluctuations. Future work should therefore consider how different mechanisms relevant to fluctuations might interact.

More broadly, slow ramping signals also emerge before other types of spontaneous behaviors, including generation of creative ideas[61], spontaneous abstract decisions[62,63], eureka moments in problem solving[64], free recall of memories[65,66], and switches between bistable percepts during binocular rivalry[67]. Notably, slow ramping signals before creative ideas have been linked to the temporal properties of spontaneous resting-state fluctuations[61], similar to the relationship between autocorrelation and slow ramping[13] that we investigated. Others have noticed these similarities and proposed that similar mechanisms underlie different types of spontaneous behaviors[17,68]. The neural mechanism we describe here may apply to spontaneous cognitive and perceptual processes more generally, offering an intriguing direction for future empirical studies.

Acknowledgements
We thank Gabriel Kreiman for granting us access to the neural recordings we analyzed as part of this study and for feedback and discussion. J.G., A.S., and U.M. were supported by the John Templeton Foundation and the Fetzer Institute (Consciousness and Free Will: A Joint Neuroscientific-Philosophical Investigation (John Templeton Foundation #61283; Fetzer Institute, Fetzer Memorial Trust #4189)). U.R. has been supported by the US National Institutes of Health (U01NS117839) and the NSF (BCS-2219800).

Code and Data Availability

Code used for simulating data & analyzing simulated data will be made publicly available on OSF upon acceptance.

References

1. Kornhuber, H. H. & Deecke, L. [CHANGES IN THE BRAIN POTENTIAL IN VOLUNTARY MOVEMENTS AND PASSIVE MOVEMENTS IN MAN: READINESS POTENTIAL AND REAFFERENT POTENTIALS]. *Pflugers Arch Gesamte Physiol Menschen Tiere* **284**, 1–17 (1965).
2. Libet, B., Gleason, C. A., Wright, E. W. & Pearl, D. K. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain* **106 (Pt 3)**, 623–642 (1983).
3. Shibasaki, H. & Hallett, M. What is the Bereitschaftspotential? *Clin Neurophysiol* **117**, 2341–2356 (2006).
4. Cunnington, R., Windischberger, C., Deecke, L. & Moser, E. The Preparation and Execution of Self-Initiated and Externally-Triggered Movement: A Study of Event-Related fMRI. *NeuroImage* **15**, 373–385 (2002).
5. Cunnington, R., Windischberger, C. & Moser, E. Premovement activity of the pre-supplementary motor area and the readiness for action: Studies of time-resolved event-related functional MRI. *Human Movement Science* **24**, 644–656 (2005).
6. Sakata, H. *et al.* Slow Accumulations of Neural Activities in Multiple Cortical Regions Precede Self-Initiation of Movement: An Event-Related fMRI Study. *eNeuro* **4**, ENEURO.0183-17.2017 (2017).
7. Fried, I., Mukamel, R. & Kreiman, G. Internally Generated Preactivation of Single Neurons in Human Medial Frontal Cortex Predicts Volition. *Neuron* **69**, 548–562 (2011).
8. Bauer, H. H. & Rebert, C. S. Preliminary study on subcortical slow potentials related to the readiness potential in monkey. *Int J Psychophysiol* **9**, 269–278 (1990).
9. Murakami, M., Vicente, M. I., Costa, G. M. & Mainen, Z. F. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat Neurosci* **17**, 1574–1582 (2014).
10. Mitelut, C. *et al.* Mesoscale cortex-wide neural dynamics predict self-initiated actions in mice several seconds prior to movement. *eLife* **11**, e76506 (2022).
11. Kagaya, K. & Takahata, M. Readiness Discharge for Spontaneous Initiation of Walking in Crayfish. *J. Neurosci.* **30**, 1348–1362 (2010).
12. Libet, B. Unconscious cerebral initiative and the role of conscious will in voluntary action. *THE BEHAVIORAL AND BRAIN SCIENCES* 38 (1985).
13. Schurger, A. Specific Relationship between the Shape of the Readiness Potential, Subjective Decision Time, and Waiting Time Predicted by an Accumulator Model with Temporally Autocorrelated Input Noise. *eNeuro* ENEURO.0302-17.2018 (2018) doi:10.1523/ENEURO.0302-17.2018.
14. Schurger, A., Sitt, J. D. & Dehaene, S. An accumulator model for spontaneous neural activity prior to self-initiated movement. *PNAS* **109**, E2904–E2913 (2012).
15. Maoz, U., Yaffe, G., Koch, C. & Mudrik, L. Neural precursors of decisions that matter—an ERP study of deliberate and arbitrary choice. *eLife* **8**, e39787 (2019).
16. Khalighinejad, N., Schurger, A., Desantis, A., Zmigrod, L. & Haggard, P. Precursor processes of human self-initiated action. *NeuroImage* **165**, 35–47 (2018).
17. Moutard, C., Dehaene, S. & Malach, R. Spontaneous Fluctuations and Non-linear Ignitions: Two Dynamic Faces of Cortical Recurrent Loops. *Neuron* **88**, 194–206 (2015).

18. Schmidt, S., Jo, H.-G., Wittmann, M. & Hinterberger, T. ‘Catching the waves’ – slow cortical potentials as moderator of voluntary action. *Neuroscience & Biobehavioral Reviews* **68**, 639–650 (2016).
19. Schurger, A., Hu, P. ‘Ben’, Pak, J. & Roskies, A. L. What Is the Readiness Potential? *Trends in Cognitive Sciences* **25**, 558–570 (2021).
20. Cavanagh, S. E., Wallis, J. D., Kennerley, S. W. & Hunt, L. T. Autocorrelation structure at rest predicts value correlates of single neurons during reward-guided choice. *eLife* **5**, e18937 (2016).
21. Murray, J. D. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience* **17**, 1661–1663 (2014).
22. Litwin-Kumar, A. & Doiron, B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience* **15**, 1498–1505 (2012).
23. Schaub, M. T., Billeh, Y. N., Anastassiou, C. A., Koch, C. & Barahona, M. Emergence of Slow-Switching Assemblies in Structured Neuronal Networks. *PLOS Computational Biology* **11**, e1004196 (2015).
24. Wong, K.-F. & Wang, X.-J. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
25. Li, Y., Kim, R. & Sejnowski, T. J. Learning the Synaptic and Intrinsic Membrane Dynamics Underlying Working Memory in Spiking Neural Network Models. *Neural Computation* **33**, 3264–3287 (2021).
26. Klinshov, V. V., Kovalchuk, A. V., Franović, I., Perc, M. & Svetec, M. Rate chaos and memory lifetime in spiking neural networks. *Chaos, Solitons & Fractals* **158**, 112011 (2022).
27. Fino, E. & Yuste, R. Dense inhibitory connectivity in neocortex. *Neuron* **69**, 1188–1203 (2011).
28. van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
29. Oster, M., Douglas, R. & Liu, S.-C. Computation with spikes in a winner-take-all network. *Neural Comput* **21**, 2437–2465 (2009).
30. Cramer, B. *et al.* Autocorrelations from emergent bistability in homeostatic spiking neural networks on neuromorphic hardware. Preprint at https://doi.org/10.48550/arXiv.2208.08329 (2022).
31. Gerstner, W. & Kistler, W. M. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. (Cambridge University Press, 2002).
32. Soon, C. S., Brass, M., Heinze, H.-J. & Haynes, J.-D. Unconscious determinants of free decisions in the human brain. *Nature Neuroscience* **11**, 543–545 (2008).
33. Zapparoli, L. *et al.* Dissecting the neurofunctional bases of intentional action. *PNAS* **115**, 7440–7445 (2018).
34. Aflalo, T. *et al.* Implicit mechanisms of intention. *Current Biology* **32**, 2051-2060.e6 (2022).
35. Fu, Z. *et al.* The geometry of domain-general performance monitoring in the human medial frontal cortex. *Science* **376**, eabm9922 (2022).
36. King, J.-R. & Dehaene, S. Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn Sci* **18**, 203–210 (2014).
37. Nunez, P. L. & Srinivasan, R. *Electric Fields of the Brain: The Neurophysics of EEG, 2nd Edition*. (Oxford University Press, 2005).

38. Vanderhaeghen, P. & Polleux, F. Developmental mechanisms underlying the evolution of human cortical circuits. *Nat Rev Neurosci* **24**, 213–232 (2023).
39. Kim, R. & Sejnowski, T. J. Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks. *Nature Neuroscience* **24**, 129–139 (2021).
40. Bai, O. *et al.* Prediction of human voluntary movement before it occurs. *Clinical Neurophysiology* **122**, 364–372 (2011).
41. Lew, E., Chavarriaga, R., Zhang, H., Seeck, M. & Millán, J. del R. Self-paced movement intention detection from human brain signals: Invasive and non-invasive EEG. in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* 3280–3283 (2012). doi:10.1109/EMBC.2012.6346665.
42. Lew, E. Y. L., Chavarriaga, R., Silvoni, S. & Millán, J. del R. Single trial prediction of self-paced reaching directions from EEG signals. *Frontiers in Neuroscience* **8**, (2014).
43. Baars, B., Franklin, S. & Ramsøy, T. Global Workspace Dynamics: Cortical "Binding and Propagation" Enables Conscious Contents. *Frontiers in Psychology* **4**, (2013).
44. Baars, B. J., Geld, N. & Kozma, R. Global Workspace Theory (GWT) and Prefrontal Cortex: Recent Developments. *Frontiers in Psychology* **12**, (2021).
45. Richer, F. & Beatty, J. Pupillary Dilations in Movement Preparation and Execution. *Psychophysiology* **22**, 204–207 (1985).
46. Hamilos, A. E. *et al.* Slowly evolving dopaminergic activity modulates the moment-to-moment probability of reward-related self-timed movements. *eLife* **10**, e62583 (2021).
47. Lee, I. H. & Assad, J. A. Putaminal Activity for Simple Reactions or Self-Timed Movements. *Journal of Neurophysiology* **89**, 2528–2537 (2003).
48. Schultz, W. & Romo, R. Role of primate basal ganglia and frontal cortex in the internal generation of movements. I. Preparatory activity in the anterior striatum. *Exp Brain Res* **91**, 363–384 (1992).
49. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, eaav7893 (2019).
50. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci* **22**, 1677–1686 (2019).
51. Schurger, A. & Uithol, S. Nowhere and Everywhere: The Causal Origin of Voluntary Action. *Rev.Phil.Psych.* **6**, 761–778 (2015).
52. Raut, R. V. *et al.* Global waves synchronize the brain's functional systems with fluctuating arousal. *Science Advances* **7**, eabf2709 (2021).
53. Park, H.-D. *et al.* Breathing is coupled with voluntary action and the cortical readiness potential. *Nature Communications* **11**, 1–8 (2020).
54. Huang, Z., Liu, X., Mashour, G. A. & Hudetz, A. G. Timescales of Intrinsic BOLD Signal Dynamics and Functional Connectivity in Pharmacologic and Neuropathologic States of Unconsciousness. *J. Neurosci.* **38**, 2304–2317 (2018).
55. Palva, J. M. *et al.* Neuronal long-range temporal correlations and avalanche dynamics are correlated with behavioral scaling laws. *Proceedings of the National Academy of Sciences* **110**, 3585–3590 (2013).
56. Hill, A. T., Clark, G. M., Bigelow, F. J., Lum, J. A. G. & Enticott, P. G. Periodic and aperiodic neural activity displays age-dependent changes across early-to-middle childhood. *Dev Cogn Neurosci* **54**, 101076 (2022).

57. Wiest, C. *et al.* The aperiodic exponent of subthalamic field potentials reflects excitation/inhibition balance in Parkinsonism. *eLife* **12**, e82467 (2023).
58. Chaudhuri, R., He, B. J. & Wang, X.-J. Random Recurrent Networks Near Criticality Capture the Broadband Power Distribution of Human ECoG Dynamics. *Cereb Cortex* **28**, 3610–3622 (2018).
59. Poil, S.-S., Hardstone, R., Mansvelder, H. D. & Linkenkaer-Hansen, K. Critical-state dynamics of avalanches and oscillations jointly emerge from balanced excitation/inhibition in neuronal networks. *J. Neurosci.* **32**, 9817–9823 (2012).
60. di Santo, S., Villegas, P., Burioni, R. & Muñoz, M. A. Landau–Ginzburg theory of cortex dynamics: Scale-free avalanches emerge at the edge of synchronization. *Proceedings of the National Academy of Sciences* **115**, E1356–E1365 (2018).
61. Broday-Dvir, R. & Malach, R. Resting-State Fluctuations Underlie Free and Creative Verbal Behaviors in the Human Brain. *Cerebral Cortex* **31**, 213–232 (2021).
62. Soon, C. S., He, A. H., Bode, S. & Haynes, J.-D. Predicting free choices for abstract intentions. *PNAS* **110**, 6217–6222 (2013).
63. Alexander, P. *et al.* Readiness potentials driven by non-motoric processes. *Consciousness and Cognition* **39**, 38–47 (2016).
64. Jung-Beeman, M. *et al.* Neural Activity When People Solve Verbal Problems with Insight. *PLOS Biology* **2**, e97 (2004).
65. Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R. & Fried, I. Internally generated reactivation of single neurons in human hippocampus during free recall. *Science* **322**, 96–101 (2008).
66. Norman, Y., Yeagle, E. M., Harel, M., Mehta, A. D. & Malach, R. Neuronal baseline shifts underlying boundary setting during free recall. *Nat Commun* **8**, 1301 (2017).
67. Gelbard-Sagiv, H., Mudrik, L., Hill, M. R., Koch, C. & Fried, I. Human single neuron activity precedes emergence of conscious perception. *Nat Commun* **9**, 2057 (2018).
68. Fried, I. Neurons as will and representation. *Nat Rev Neurosci* **23**, 104–114 (2022).

Methods

LIF Network Simulations

We simulated networks of leaky integrate-and-fire (LIF) neurons where the membrane potential V_m was governed (and updated at a resolution of 1 ms) by the equations:

$$\tau_m \frac{dV_m}{dt} = (E_L - V_m) + I_{app} + RI + \zeta$$

Here τ_m is the synaptic time-constant (10 ms for our simulations). E_L is the reversal potential (-70 mV). I_{app} is a constant bias current set at the rheobase (300 pA). $\zeta \sim N(0, 0.5\ mV)$ is the noise term, which changes every 10 ms. In other words, each neuron receives constant noise for 10 ms, which then changes. The membrane potential, V_m, had a threshold of -40 mV, and upon crossing it was reset to -65 mV and held constant for a refractory period of 2 ms. In every simulation $V_{m\ initial} \sim U(-70, -50\ mV)$ for each neuron.

RI is the synaptic current vector. For each neuron it is a sum (weighted according to the connectivity matrix W; details below) of the other neurons' synaptic output. Each such synaptic output is itself a weighted sum of fast and slow synaptic currents. RI follows the equation:

$$RI = W * \left(0.9 \ * \ K_{syn} * \ r_{slow} + \left(1 - K_{syn}\right) * r_{fast}\right)$$

Here, $0 \leq K_{syn} \leq 1$ is the ratio of fast to slow synapses. Hence, if $K_{syn} = 0$ or $K_{syn} = 1$ the network's activity would entirely be governed by fast or slow synapses, respectively. And $0 < \ K_{syn} < 1$, the network's activity is governed by a weighted sum of fast and slow synaptic currents. We multiplied the slow synaptic currents by 0.9 to afford better comparisons between networks with similar clustering values (see below). Synaptic currents, r, follow a double-exponential time course including a rising and falling time-constant which differs between fast and slow synapses. Thus r follows the equations:

$$\frac{dr}{dt} = -\frac{1}{\tau_{fall}} r + s$$

$$\frac{ds}{dt} = \ -\frac{1}{\tau_{rise}} s + \frac{1}{\tau_{rise}\tau_{fall}} * Spikes(t)$$

Here τ_{rise} and τ_{fall} are the time-constants for the synaptic currents to rise to their peak and then decay, respectively. Fast synapses had $\tau_{rise} = 2\ ms$ and $\tau_{fall} = 5ms$ and slow synapses had $\tau_{rise} = 20\ ms$ and $\tau_{fall} = 100\ ms$ [31]. $Spikes(t)$ is a binary (0 or 1) spiking vector reflecting, for each time-step in the simulation, whether each neuron had spiked on the previous timestep. These spiking values were normalized by the rising and falling time constants, which ensured that the same total current was delivered for fast and slow synapses (see Figure 1B).

Connectivity Matrix Construction

Much of the dynamics of neural networks is determined by how neurons are interconnected, encoded in the connectivity matrix W. Previous research has found that *balanced* networks with roughly equal excitatory and inhibitory signaling exhibit stable asynchronous dynamics[28]. Balanced networks are usually implemented with more excitatory than inhibitory neurons (usually 75-80% excitatory), though with relatively weaker synaptic signaling in those excitatory neurons (to compensate for their higher proportion). To introduce slow changes in firing rates, we grouped the excitatory neurons into "clusters" both topographically and synaptically. Topographically, we raised the probability of connection for inter-cluster excitatory neurons in comparison with intra-cluster excitatory neurons. Synaptically, we increased the weights representing connections between excitatory neurons in the same cluster[22] (although similar slow fluctuations in firing rates can be achieved by increasing intra-cluster connectivity strength alone, as well as through other methods[23]). Notably, however, the slow fluctuations in firing rates that emerged due to this clustering were wildly different depending on the synaptic dynamics (see LIF Network Simulations section above).

We wanted to maintain sparse connectivity in our network, hence we used a 20% connection probability between any two excitatory neurons (P_{EE}), taking into account the different number of intra- and inter-cluster connections. In a network with more than two clusters, there will be a larger percentage of inter-cluster connection opportunities than intra-cluster ones. So, calculating

the precise values of the intra-cluster connection probability (P_{EE}^{In}) and inter-cluster connection probability (P_{EE}^{Out}) is non-trivial. To do so, we first note that the overall probability of connection (P_{EE}) is a sum of intra- and inter-cluster connections divided by the total number of possible connections between excitatory neurons:

$$P_{EE} = \frac{E(\#\ Intra - cluster\ connections) + E(\#\ Inter - cluster\ connections)}{{N_E}^2} = 0.2$$

Next, we find the expected number of intra-cluster connections:

$$E(\#\ Intra - cluster\ connections) = P_{EE}^{In} * N_{Clusters} * C^2 = P_{EE}^{In} \frac{{N_E}^2}{N_{Clusters}}$$

Where $N_{Clusters}$ is the number of clusters, C is the number of neurons in each cluster (which equals $\frac{N_E}{N_{Clusters}}$), and N_E is the total number of excitatory neurons. The expected number of inter-cluster connections similarly follows:

$$E(\#Inter - cluster\ connections) = P_{EE}^{Out} * N_{Clusters}(N_{Clusters} - 1) * C^2$$
$$= P_{EE}^{Out} \frac{{N_E}^2 (N_{Clusters} - 1)}{N_{Clusters}}$$

Then, letting $A = \frac{{N_E}^2}{N_{Clusters}}$ and $B = \frac{{N_E}^2 (N_{Clusters} - 1)}{N_{Clusters}}$ for simplicity, we obtain:

$$P_{EE} = \frac{P_{EE}^{In} * A + P_{EE}^{Out} * B}{{N_E}^2}$$

Finally, letting the clustering ratio $R_{EE} = \frac{P_{EE}^{In}}{P_{EE}^{Out}}$ we can then solve explicitly for P_{EE}^{In} and P_{EE}^{Out}:

$$P_{EE}^{In} = \frac{{N_E}^2 P_{EE} R_{EE}}{B + AR_{EE}}$$
$$P_{EE}^{Out} = \frac{{N_E}^2 P_{EE}}{B + AR_{EE}}$$

Solving for these probabilities explicitly thereby allows for a relatively easy construction of connectivity matrices corresponding to clustered networks, while maintaining a constant probability of connection overall (P_{EE}=0.2, or 20% between excitatory neurons). Furthermore, these equations demonstrate that the precise values of P_{EE}^{In} and P_{EE}^{Out} depend not only on the clustering coefficient (R_{EE}), but also the number of excitatory neurons (N_E), and the number of clusters ($N_{Clusters}$) (or, equivalently, on the number of clusters and size of each cluster).

Simulating & Analyzing Fluctuations

We simulated 5 seconds of activity in 50 unique network architectures for a range of R_{EE} values (1 to 6 at increments of 0.1) and a range of synapse ratios (0% to 50% slow synapses at increments of 10%). After discarding the first 500 milliseconds of each simulation, we calculated the firing rates of each excitatory neuron i ∈ {1, …, N}, where N is 320 (80% of 400), over time using non-overlapping bins of 100 milliseconds. From these we calculated the coefficient of variation of firing rates across time (standard deviation divided by mean). For each neuron, with T time bins and firing rates at each bin we get a sequence Fi = [fi(t1), fi(t2), …, fi(tT)], over which we calculated the standard deviation σi and mean μi. Then, the network-level coefficient of variation is as follows:

$$CV_{Net} = \frac{1}{N}\sum_{i=1}^{N}\frac{\sigma_i}{\mu_i}$$

We calculated a single CV_{Net} value for each network, and plot grand means and standard error across networks in Figure 1D. For Figure S1 we plot individual network-average values for the first ten of fifty networks to improve visibility.

In order to quantify spiking variability that was due to slow-switching or bistable activity, we also calculated the $\hat{S}T$ value used by Schaub and colleagues[23] to investigate slow-switching behavior. This measure quantifies spiking variability over time by taking the average firing rate in each cluster $i \in \{1, \ldots, C\}$ in 100 ms non-overlapping bins. Then, for each cluster with firing rates in T time bins we get a sequence of firing rates $F_i = [f_i(t_1), f_i(t_2), \ldots, f_i(t_T)]$ where t_k represents the kth time bin, for which we calculate the standard deviation $\sigma_T(i)$, and then sum to obtain an initial measure of spiking variability.

$$S_T = \frac{1}{C}\sum_{i=1}^{C}\sigma_T(i)$$

We then normalize this value by subtracting a bootstrapped value S_{Boot} calculated on shuffled data (shuffling each neuron's cluster membership), averaged over 20 random shuffles.

$$\hat{S}_T = S_T - S_{Boot}$$

Like for CV, we calculated a single $\hat{S}_T$ value for each network, and plot grand means and standard error across networks in Figure 1E. For Figure S1 we once again plot individual network-average values for the first ten of fifty networks.

Network Selection for Threshold-Aligned Analyses

Prior work had shown that clustered networks exhibit fluctuations for a critical range of clustering coefficients ($2 < R_{EE} < 5$ typically, although this depends on network specifics; Litwin-Kumar & Doiron, 2012; Schaub et al., 2015). Because we generated network structures pseudo-randomly (setting the seed for random generation using MATLAB's rng command), the critical range was slightly different for each network. We identified an R_{EE} value suitable for simulation in each network by visually inspecting activity at a range of values. We selected R_{EE} values for each network such that the network exhibited fluctuations at the given value for both 100% fast and 50% slow synapses, but exhibited winner-take-all behavior in at least one case if we increased the R_{EE} by 0.3. Prior work had linked the R_{EE} value demarking the transition to the winner-take-all regime to the leading eigenvalue of the connectivity matrix exceeding a value of 1 [23]. However, our simulations did not exhibit such a relation, likely because that prior work used non-dimensionalized equations for simulating our activity whereas we used the dimensionalized versions of those equations.

Threshold Alignment

We hypothesized that spontaneous fluctuations trigger movements upon crossing a threshold, leading to slow ramping in trial-averaged activity. To investigate this hypothesis, we simulated

10-second trials 100 times for each of the 20 network architectures and then aligned to threshold-crossings (defined below). We identified the most-active cluster across all simulated trials and aligned to threshold crossings in that cluster (reflecting the notion that threshold crossings in only some subpopulation of SMA would lead to a movement). If no threshold crossing was detected in a trial, the trial was omitted from further analyses.

Threshold crossings were defined as the time at which the most-active cluster crossed 50% of its maximum normalized firing rate across the entire simulation period. (obtained by averaging firing rates across neurons in each cluster, subtracting the mean firing rate of the first 500 milliseconds, and then dividing by the maximum firing rate on each trial). That choice was motivated by the fact that networks show heterogeneity in their fluctuations both in timing and magnitude, and that increasing neurons reached about 50% of their maximum normalized firing rate right before movement (Figure 4 in Fried et al., 2011). Only trials where crossing occurred more than 3 seconds into the trial were retained for analysis of neural data (as is commonly done in self-initiated action paradigms[7]), as we were looking at changes across this period.

Neurons were categorized as increasing or decreasing by comparing their firing rates in the 400 ms before threshold crossings to their firing rates during a baseline period (-3 to -2 seconds, rank-sum test with a cut-off of p=0.01, as in[7]). Firing rates were smoothed with a gaussian kernel (width 400 ms), then averaged across response types (increasing and decreasing) to get a single trace for each network. Those traces were then averaged across networks for Figure 2A; the shaded regions reflect the confidence intervals across networks.

Our EEG proxy was obtained by averaging the post-synaptic potentials (PSPs) onto excitatory neurons[37]. We retained the PSPs onto each neuron (*RI* in the above equations). We then sub-selected the PSPs to retain only those going onto excitatory neurons, averaging across neurons to obtain a single EEG proxy per trial. We then averaged across trials to get a single EEG proxy trace per network. Those traces were again averaged across networks for Figure 2B, and the shaded regions reflect the confidence intervals across networks.

Autocorrelation at the single-neuron and cluster level was obtained using the autocorr function in MATLAB on the firing rates (smoothed using a moving average filter with a window of 10 milliseconds to avoid smoothing artifacts; then averaged across neurons in each cluster for cluster-average firing rates) for each -3 to 0.5-second trial aligned to threshold-crossings. Those traces were then averaged across trials to get a single autocorrelation trace for each neuron or cluster. Then, those traces were averaged across neurons or clusters to get a single trace per network and then averaged across networks to obtain grand averages for Figure 2C.

Neuron Categorization & Response Profiles

We categorized real and simulated neurons as increasing or decreasing following the process described by Fried et al. [7]. We compared firing rates in the 400 milliseconds before threshold crossing or movement to firing rates during a baseline period (-3 to -2 seconds) using a non-parametric rank-sum t-test. Neurons that significantly increased (decreased) their activity relative to baseline were categorized as increasing (decreasing; cut-off was p=0.01). Notably, this process omits neurons with more complex response profiles (e.g. Figure 3 in Fried et al., 2011[7]), but these are out of the scope of this study.

MFC neurons had a mixture of response profiles. In particular, some neurons slowly ramped their firing rates up or down before movement, and others exhibited a steep jump up or down before movement. Simulated neurons tended to exhibit the same type of response (slow ramp or steplike), but responses in real data were mixed and therefore required a method to distinguish them. To do so, followed a process similar to the supplementary analysis conducted by Fried et al.[7] (supplementary analysis/Figure 2). We fit a sigmoid function to the trial-averaged smoothed firing rates (400 ms Gaussian) using nonlinear least-squares via the Levenberg-Marquardt algorithm in MATLAB. We fit a different sigmoid depending on if the neuron was increasing or decreasing:

$$\frac{F_W \cdot (-1)^{\,d}}{1 + exp(-(x - t_0)/\alpha)} + F_B.$$

Here F_B is the firing rate at baseline, F_W is the firing rate in the 400 ms before movement. t_0 is the (fitted) sigmoid inflection point. α is the (fitted) inverse-gain parameter (i.e. higher values of α correspond to a smoother ramp). And d is 0 for increasing neurons and 1 for decreasing ones. We selected an inverse-gain of 0.4 as the cut-off for steplike versus slow ramp response profiles based on visualizing averaged firing rates from the Fried et al. [7] (2011) dataset (see Figure 2C). Thus, neurons with an α parameter greater than 0.4 were categorized as slow-ramping, and neurons with an α parameter less than 0.4 were categorized as steplike.

Decoding Analyses

We also assessed whether above-chance decoding accuracy in the seconds before movement[7,32] might emerge from aligning fluctuations like those in our networks to threshold-crossings. For simulated data, at each time point, we compared spike counts in the time bin [t-0.4, t] s to spike counts in the baseline of (-3,-2.6) s for all neurons. We trained a linear discriminant analysis classifier on 80% of the data and recorded its performance on a reserved 20% test set (i.e., 5-fold cross-validation). Figure 3A shows the average accuracy and standard error across 5 cross-validation splits for individual networks (the same splits were used for $K_{syn} = 0$ and $K_{syn} = 0.5$), as well as performance on 100 random splits of the data to assess chance-level performance. For real data, we constructed a pseudo population of all recorded MFC neurons. We aggregated across subjects with at least 50 trials and conducted the same analysis as above, with time relative to movement onset rather than threshold-crossing.

To assess the temporal generalization of decoding (Figure 3B), we completed a similar analysis as above. However, rather than splitting the data at a single time point into a training and testing set, we trained the classifier on the data from a given time point and then tested the trained classifier on a different time point[36]. Therefore, the main diagonal (train and test timepoint are the same) corresponds to the classifiers' training accuracy; the off-diagonal results refer to the classifiers' test accuracy.

Simulating Ramping Spike-Trains

We simulated spike trains where we knew that the ground-truth was a gradual ramp in order to see if the asymmetry we observed in our temporal generalization analysis (where decoders trained to differentiate activity from a baseline generalize better into the future than into the past) is due to ramping. We simulated 50 spike trains for a pseudopopulation of 20 neurons. Half of the neurons increased their firing rates from 5 Hz to 8 Hz, and the other half decreased their firing rates from 8 Hz to 5 Hz. We simulated 3 seconds of activity (labeled as -3 to 0 to match

other analyses) and generated spike trains in 50 millisecond bins, where the instantaneous firing rate was obtained by linearly interpolating the starting and ending firing rates. Then, we conducted the same temporal generalization analysis as above using a baseline of [-3,-2.6]. We found a striking asymmetry in the AUCs such that decoders generalized better into the future than into the past (one-tailed rank-sum t-test on upper vs lower triangle of AUC matrix, $p < 0.001$; Fig S4D).

Analysis of Pairwise Correlations at Baseline

Our model suggests that early ramping reflects fluctuations in clustered networks aligned to threshold-crossings. Crucially, those fluctuations emerge naturally due to the clustered connectivity, and also because neurons that ramp up tend to be in the same cluster, and those that ramp down tend to belong to clusters other than the one ramping up. However, in our model, connectivity is static, as would be expected in biological networks at the time scales in which we simulate here. Neurons in the same cluster therefore show correlated activity across the entire trial, not just during ramping. Hence, our model predicts that neurons that ramp in up or down together before movement or threshold crossings should have correlated activity even before ramping onset.

We thus investigated pairwise correlations at baseline in real neurons in a baseline period of (-3.5,-2.5) (a baseline of (-3,-2.5) was used for simulated data because we did not save data earlier than -3 seconds). We first smoothed (Gaussian kernel, 400 ms wide), then soft-normalized (divided by the square root of the maximum firing rate[35]) firing rates, and finally down-sampled to a sample every 50 milliseconds to avoid inflating correlations due to autocorrelations in the data. After that, we correlated firing rates between every pair of neurons that were recorded in the same session, keeping track of whether the pair was both increasing, both decreasing, one increasing and one decreasing, or involved a neuron that didn't change its firing rate (for these analyses we categorized neurons with a steplike response as "other" in order to focus solely on slow-ramping behavior). At that point, we used linear mixed-effects models (with the network or subject as a random intercept) to investigate whether correlation structure at baseline depended on later ramping behavior. We found that increasing and decreasing pairs show significantly correlated activity at baseline. In contrast, increasing-decreasing pairs do not; they also show significantly lower correlations than increasing-increasing pairs and decreasing-decreasing pairs. Notably, running the same analysis without separating slow ramping and steplike responses led to all shown statistical tests in Figure 4C becoming highly significant (both increasing vs increasing-decreasing $p = 0.001$; both increasing vs other $p = 0.013$; both decreasing vs increasing-decreasing $p < 0.001$; both decreasing vs other $p < 0.001$), as well as the difference between increasing-decreasing and other pairs becoming significant ($p = 0.040$). Additionally, omitting pairwise comparisons involving two neurons on the same channel (accounting for potential issues due to spike sorting) largely did not change the results for decreasing-decreasing pairs, which had many examples. However, only few slow ramping increasing-increasing pairs remained after this correction, reducing statistical power and leading to insignificant differences between increasing-increasing vs increasing-decreasing pairs ($p = 0.439$), and increasing-increasing vs other pairs ($p = 0.777$). However, the estimated correlation was similar to decreasing neurons after that correction, so we suggest that the lack of statistical power is the main cause of these differences.

Supplementary Material

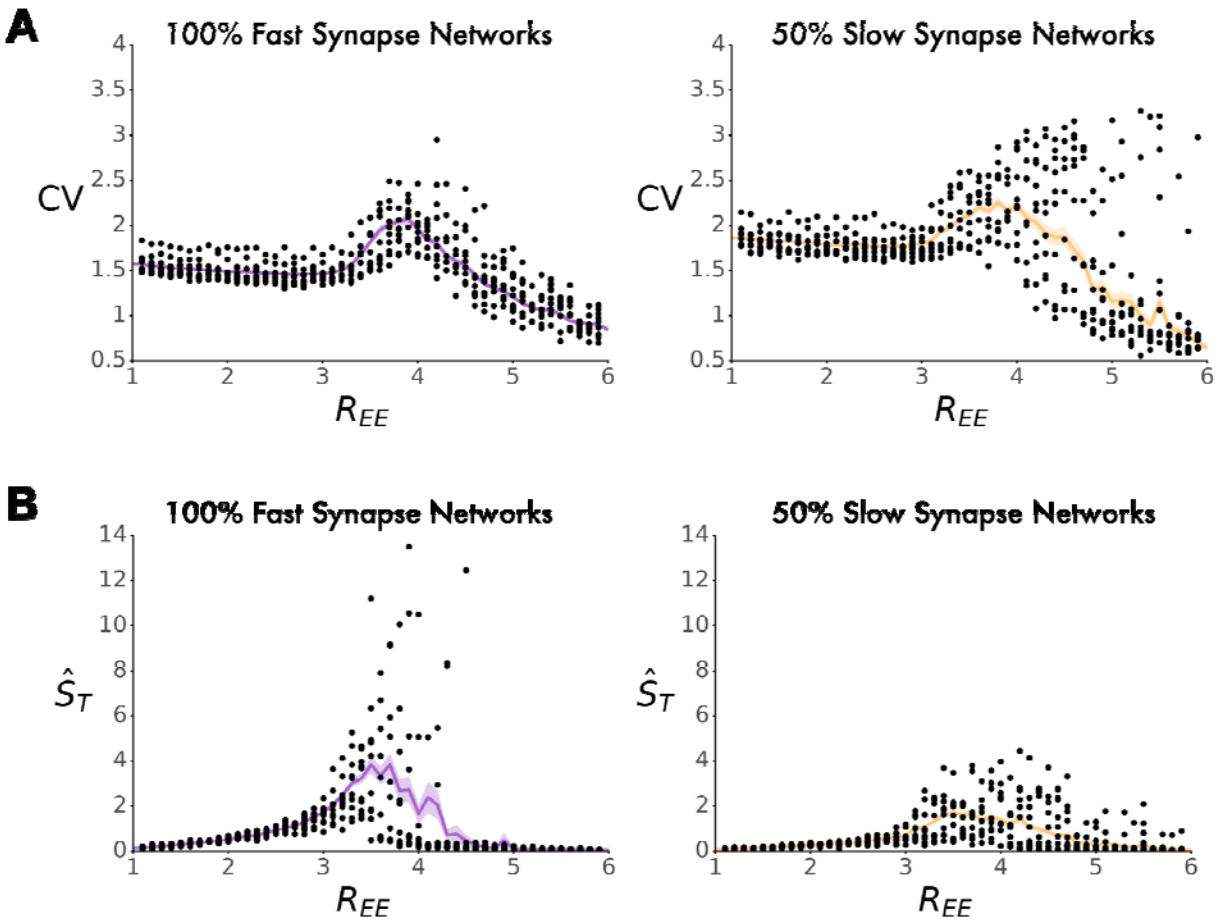

Supplementary Figure 1: Demonstration of the heterogeneity of fluctuations on different underlying network structures. Like Figure 1D-E, we calculated CV (top) and $\hat{S}_T$ (bottom) for a range of R_{EE} values. Plotted are means & standard errors across networks for 100% fast (left) and 50% slow networks. Black dots are average values for individual network architectures. Critically, the fluctuation regime (R_{EE} between 3 and 4) is not the same for every network—a wide range of CV and $\hat{S}_T$ values occur in this range. It is thus likely that the specific network architecture affects the R_{EE} values at which fluctuations occur, potentially due to the presence of particular connectivity motifs (e.g. cycles, hubs, etc.).

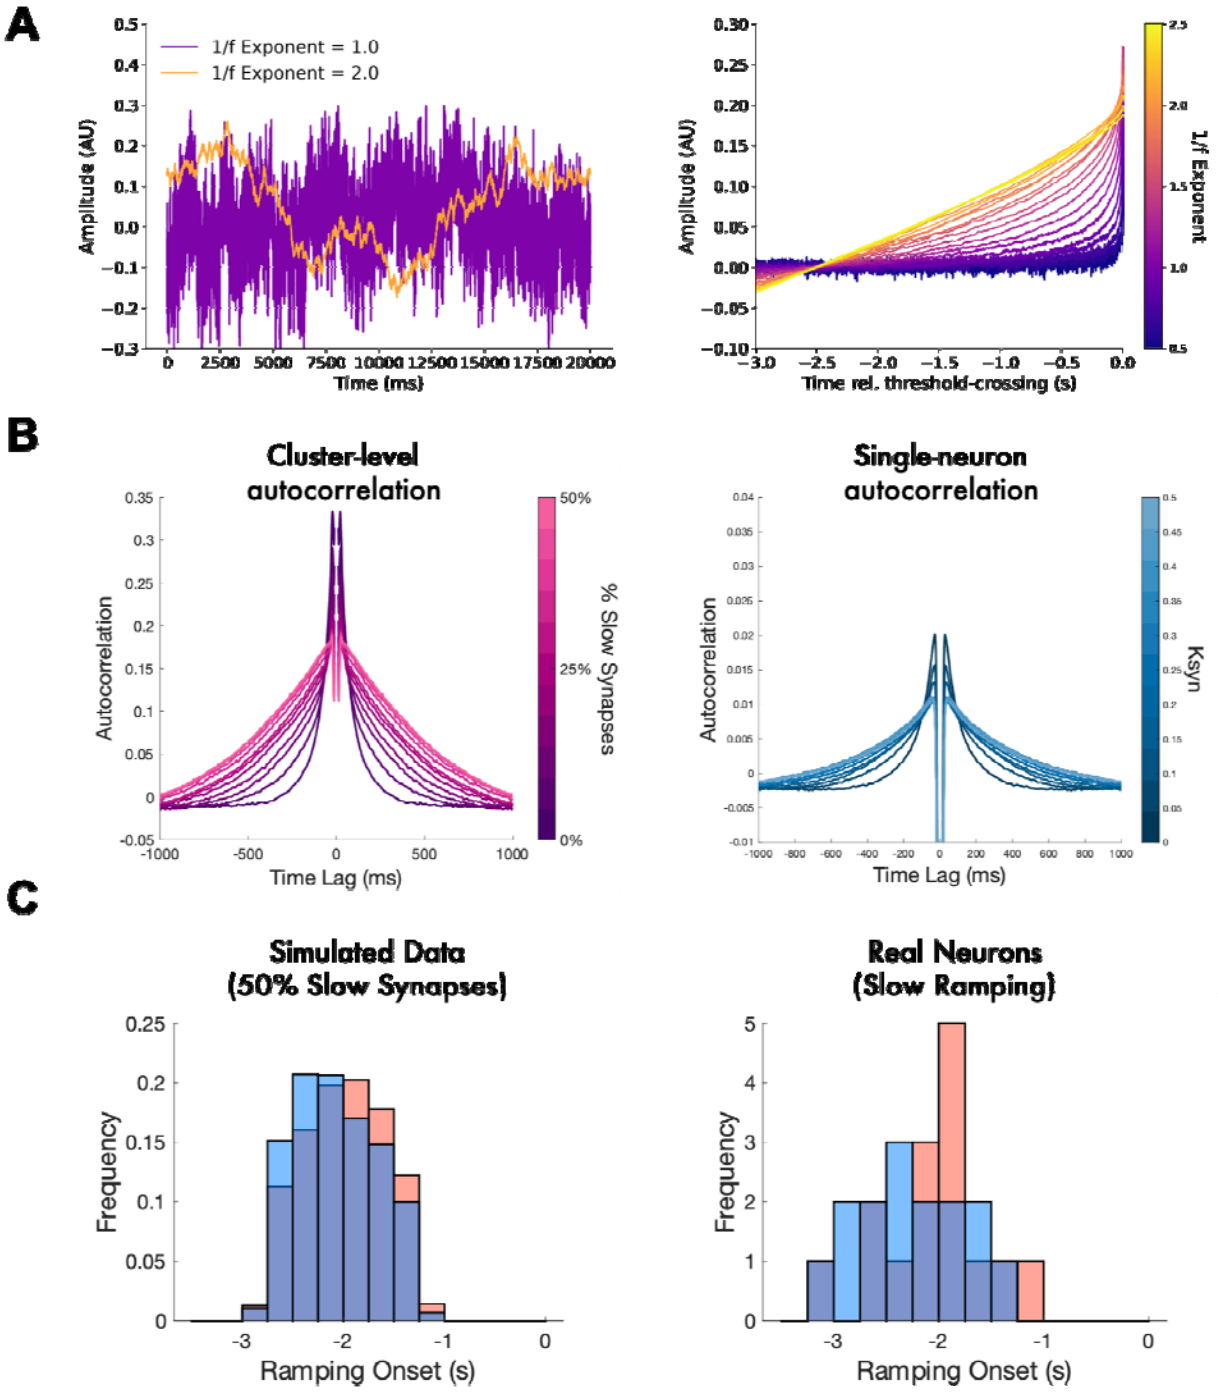

Supplementary Figure 2: Relation between autocorrelation and ramping onset & characterization of ramping onset in simulated & real data. **A:** Relation between autocorrelation and slow ramping in pink noise. We simulated pink noise with 1/f exponents ranging from 0.5 to 2.5 using the *colorednoise* Python package. We simulated 10,000 trials of 20 seconds of noise and aligned to threshold-crossings (threshold of 0.3, baselined using the period [-3, -2] as in our simulations). Left: Two example simulated trials with lower and higher autocorrelation (1/f exponent of 1 and 2, respectively). Right: When aligned to threshold crossings, high-autocorrelated noise shows slow ramping in the preceding signal, whereas low-autocorrelated

noise shows little to no ramping. As autocorrelation (1/f exponent) increases, slow ramping behavior emerges and, to some degree, has earlier onsets. **B:** Autocorrelation at the cluster-level (left, autocorrelation calculated for each cluster-average firing rate and then averaged across clusters) and single-neuron level (right, autocorrelation calculated for each neuron's firing rate and then averaged across neurons). Autocorrelation increases with increased proportion of slow synapses. Note from the Y-Axis that cluster-level autocorrelation is much larger than the single-neuron level. **C:** Characterization of slow ramping onset in simulated (left) and real neurons (right) (red = increasing neurons, blue = decreasing neurons). We used the MATLAB function findchangepts with the linear statistical method to conduct a breakpoint analysis in the timeframe of -3.5 to -1 s to determine the onset of ramping. Determining the onset of a change in slope in a noisy signal is a difficult problem, and methods using means are often biased, hence our usage of the linear statistical method, which takes into account both mean and slope to find a breakpoint. In simulated and real data, a range of ramping onsets are apparent, with an average ramping onset of about -2 seconds.

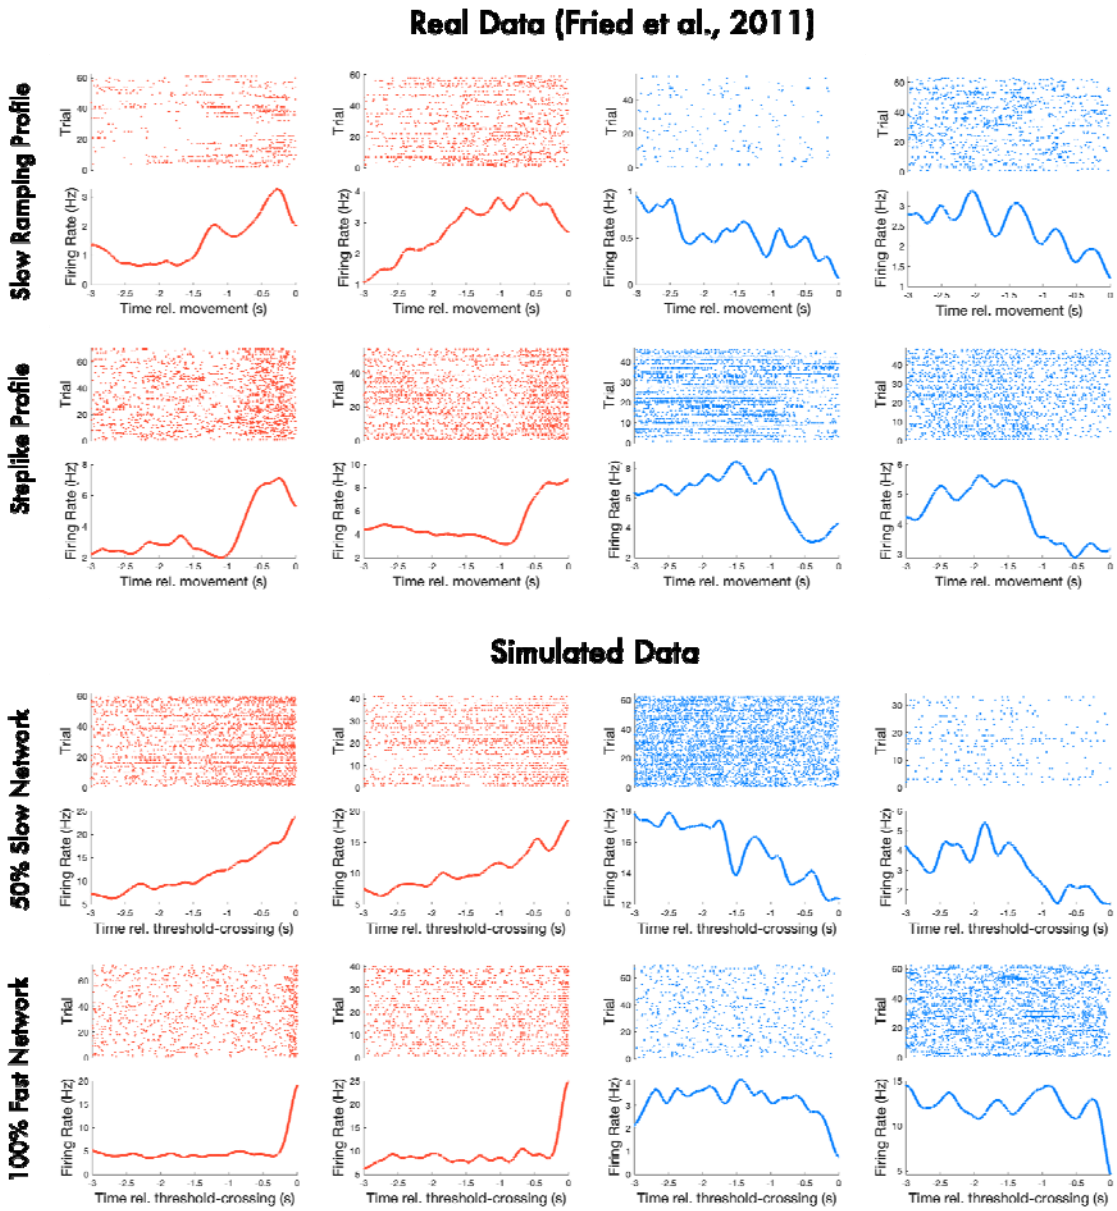

Supplementary Figure 3: Further example rasters & trial-average firing rates for real and simulated neurons (red is increasing, blue is decreasing).

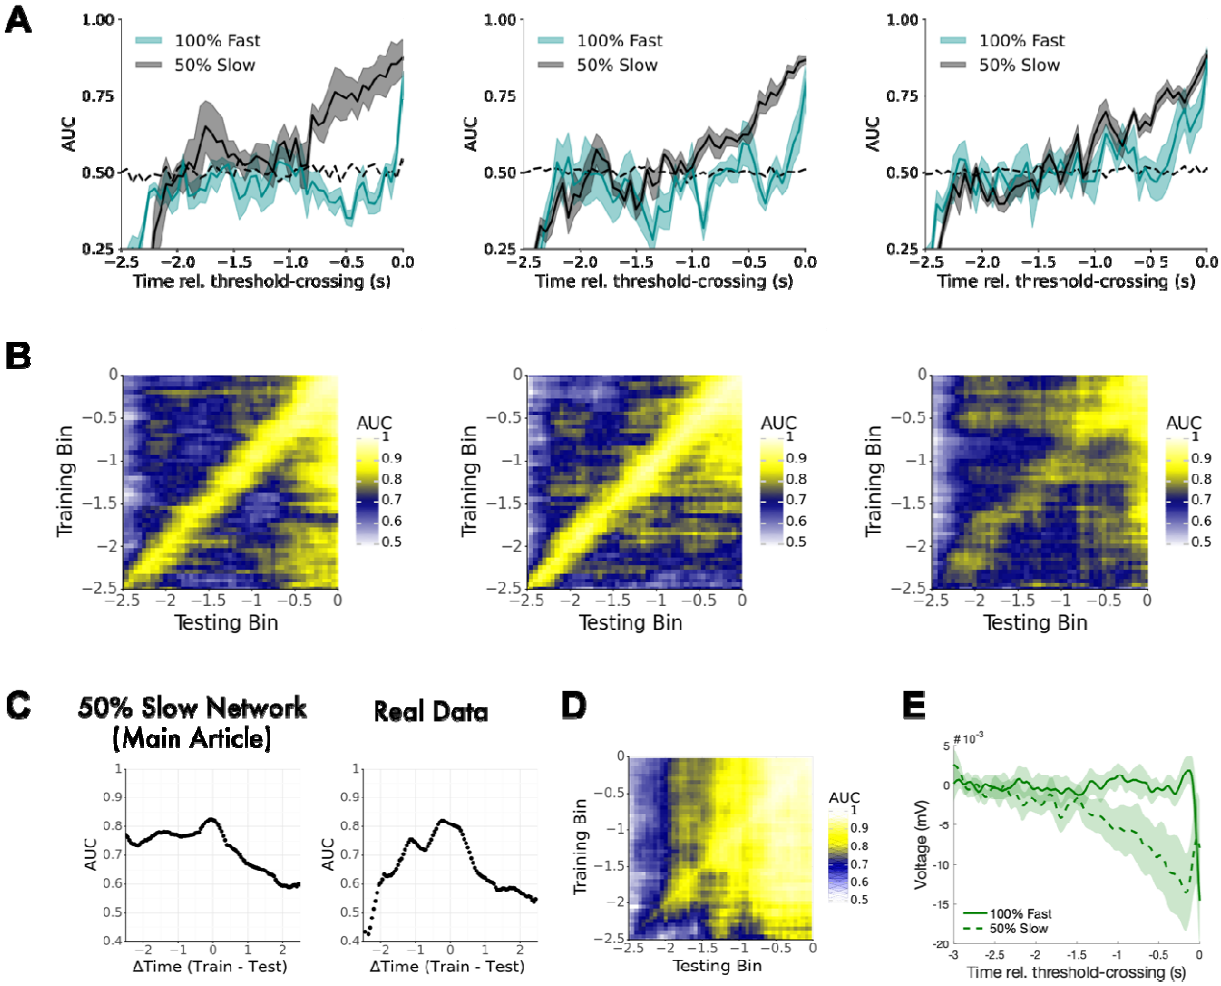

Supplementary Figure 4: More examples of slow ramping decoding accuracy in simulated networks. **A:** As in Figure 3A, decoding AUC over time for 3 other example networks with 100% fast (blue) and 50% slow (black) synapses. **B:** As in Figure 3B, temporal generalization of decoders trained at one time bin and tested on another time bin. **C:** AUC averaged across train-test bin pairs that are the same distance from diagonal. AUC was higher when the training bin was earlier than the test bin (i.e. for negative x-axis values) for both simulated (left & middle) and real data (right). That asymmetry is consistent with a monotonic drift in population firing rates away from a baseline. **D:** Results of temporal generalization analysis conducted on poisson spike-trains that ramp. In a case where ramping is a priori known to be present, we see an asymmetry such that decoders that classify activity at a given time from baseline ([-3,-2.6] in this case) generalize better to the future than they do to the past, indicated by higher AUC in the lower-triangle compared to the upper-triangle of this matrix. **E:** Pre-threshold changes in total PSP delivered to all network neurons (EEG proxy). Lines & shaded region is grand-average mean & 95% confidence intervals across 20 simulated networks. Magnitude is smaller than the EEG proxy used in main text (PSPs onto excitatory neurons only), but still shows an early negative deflection reminiscent of the EEG readiness potential before self-initiated actions.