TITLE PAGE

Title: Population temporal structure supplements the rate code during sensorimotor transformations

Abbreviated title: Population temporal dynamics in gaze control

Authors: Uday K. Jagadisan[1,3] and Neeraj J. Gandhi[1,2,3]

Author affiliations:
[1]Department of Bioengineering
[2]Department of Neuroscience
[3]Center for the Neural Basis of Cognition
University of Pittsburgh, Pittsburgh, PA 15213

Corresponding author: Uday K. Jagadisan
153 Eye and Ear Institute
203 Lothrop St
Pittsburgh, PA 15213
USA
Tel: +1-315-4099934
Correspondence and requests for more information should be addressed to U.K.J (kj.udayakiran@gmail.com) or N.J.G (neg8@pitt.edu).

Author Contributions: U.K.J and N.J.G designed the study. U.K.J. performed the experiments, analyzed the data, and performed model simulations. U.K.J and N.J.G wrote the manuscript.

Number of pages: 44

Number of figures: 6 + 9 extended figures

Number of words: Abstract - 156, Main text - 4154, Supplementary sections - 2, Methods - 4520

Conflict of interest: The authors declare no competing financial interests.

Acknowledgements: We thank Drs. A. Batista, C. Olson, M. Smith, and B. Yu for scientific discussions and critical feedback on previous versions of the manuscript and J. McFerron for programming assistance. The study was funded by the following NIH R01 grants: EY022854 and EY024831 awarded to N.J.G.

Short Summary

Sensorimotor transformations are mediated by premotor brain networks where individual neurons represent sensory, cognitive, and movement-related information. Such multiplexing poses a conundrum – how does a decoder know precisely when to initiate a movement if its inputs are active at times when a movement is not desired (e.g., in response to sensory stimulation)? Here, we propose a novel hypothesis: movement is triggered not only by an increase in firing rate, but critically by a reliable temporal pattern in the population response. Laminar recordings in the superior colliculus (SC), a midbrain hub of orienting control, and pseudo-population analyses in SC and cortical frontal eye fields (FEF) corroborated this hypothesis. We also used spatiotemporally patterned microstimulation to causally verify the importance of temporal structure. A spiking neuron model with dendritic integration was able to decode temporal structure. These findings offer an alternative perspective on movement generation and highlight the importance of short-term population history in neuronal communication and behaviour.

Long Summary

Sensorimotor transformations are mediated by premotor brain networks where individual neurons represent sensory, cognitive, and movement-related information. Such multiplexing poses a conundrum – how does a decoder know precisely when to initiate a movement if its inputs are active at times when a movement is not desired (e.g., in response to sensory stimulation)? Here, we propose a novel hypothesis: movement is triggered not only by an increase in firing rate, but critically by a reliable temporal pattern in the population response. Laminar recordings in the superior colliculus (SC), a midbrain region that plays an essential role in orienting eye movements, indicate that the temporal structure across neurons is a factor governing movement initiation. Specifically, using a measure that captures the fidelity of the population code - here called temporal stability - we show that the temporal structure fluctuates during the visual response but becomes increasingly stable during the movement command, even when the mean population activity is similar between the two epochs. Analyses of pseudo-populations in SC and cortical frontal eye fields (FEF) corroborated this model. We also used spatiotemporally patterned microstimulation to causally test the contribution of population temporal stability to movement initiation and found that stable stimulation patterns were more likely to evoke a movement, even when other features of the patterns such as mean pulse rates and population state subspaces were matched. Finally, a spiking neuron model was able to discriminate between stable and unstable input patterns, providing a putative biophysical mechanism for decoding temporal structure. These findings offer an alternative perspective on the relationship between movement preparation and generation by situating the correlates of movement initiation in the temporal features of activity in shared neural substrates. They also suggest a need to look beyond the instantaneous rate code at the single neuron or population level and consider the effects of short-term population history on neuronal communication and behaviour.

In order to successfully interact with the environment, the brain must funnel down the sensory inputs it receives to specific movements at specific times. Such sensory-to-motor transformations are critically mediated by premotor brain networks where evolving activities in individual neurons represent sensory, cognitive, and movement-related information[1-3]. For example, in brain regions involved in the control of gaze, including the SC, so-called visuomovement neurons burst a volley of spikes both in response to a visual stimulus and for generating a gaze shift or saccade to the location of the stimulus[1]. The dual nature of visuomovement neurons is best illustrated by examining their activity in the delayed response paradigm (left panel in Figure 1A), which requires the subject to withhold a saccade to a stimulus in the visual periphery until after the disappearance of a central fixation cue. We recorded population activity from the SC using laminar microelectrode arrays (right panel in Figure 1A) in two monkeys (*Macaca mulatta*). Figures 1B and 1C show responses on individual channels and population activity averaged across channels, respectively, aligned on target (left) and saccade (right) onsets for three example trials in one session. The population exhibits a high frequency visual burst following target onset and a subsequent premotor burst prior to a saccade. The peak magnitude of the visual burst is lower than that of the premotor burst on some trials (light traces in Figure 1C), but it is not uncommon for the peak visual response to match or exceed the premotor activity (medium and dark traces in Figure 1C), especially when accounting for the 10-20 ms efferent delay from neural initiation to movement onset[4] (vertical dashed line in Figure 1C). Yet, on these trials, the visual burst does not trigger a movement, an observation that casts doubt on thresholding[5,6] as a singular mechanism. Given that such neurons project directly to the brainstem saccade burst generator that initiates and guides saccadic gaze shifts[7,8], we asked how downstream structures are able to differentiate between the two bursts.

Quantification of population temporal structure

We reasoned that if the mean rate of population activity is insufficient to discriminate between visual and premotor bursts, the answer likely resides in the spatiotemporal structure of activity during the two bursts. Indeed, a growing body of work has revealed that precise coordination in the timing of input spikes is more efficient at driving downstream cortical and subcortical neurons[9-11]. Specifically, we hypothesized that a critical criterion for a decoder to generate a movement should be a high level of certainty in the instructed movement and its metrics. Since the exact decoding scheme is unknown, we hypothesized that certainty is likely provided by consistency in the population pattern over the course of a burst, under any generalized weighted pooling scheme. To quantify this consistency, we first constructed a population firing rate trajectory as a function of time. Next, we normalized the population vector at each time point, constraining it to a unit hypersphere in state space. This step factors out global changes in firing across the population and focuses on the relative activity pattern. We then computed the dot product between two of these unit vectors separated in time (parametrized by τ) - we call this measure the temporal stability of the population code. This procedure is schematized in Figure 1D (for details, see Methods).

Evolution of population temporal structure during sensorimotor tasks

Figure 1E shows the evolution of temporal stability averaged across all sessions (n = 14, mean +/- s.e.m., coloured traces, also see EDF 1A). The stability of the population pattern decreased relative to baseline

during the visual burst and increased during the premotor burst. Moreover, this property was preserved when considering only the subset of trials where the peak visual response matched or exceeded the premotor activity 15 ms before saccade onset (gray trace in Figure 1E), and was significantly different from the trend in the null condition when the saccade was directed to a stimulus in the opposite hemifield (black trace in Figure 1E). The anti-phase relationship of stability between the visual and premotor bursts was also present across a range of separation times between the population vectors (EDF 1B). We also computed temporal stability by realigning activity on the peak population visual response, to discount any effect of visual response onset latency differences across the population, and found similar (if not stronger) effects (Figure 1F). Finally, scrambling the population code by shuffling the activations of individual neurons at each time point lowered the stability profile for the entire trial significantly (EDF 1C). Thus, population temporal stability seems to impose a constraint that prevents movement initiation at an undesirable time (visual epoch), allowing the animal to successfully perform the task at hand. Once cued to execute a gaze shift, the activity in the same population rises in a stable manner allowing saccade initiation.

Next, we explored the robustness of the temporal stability hypothesis. We tested whether the properties of stability observed in other populations of neurons or conditions were consistent with its predicted role in movement initiation. For this part of the study we used data obtained from single-unit recordings in SC and the frontal eye fields (FEF), the latter of which plays a major role in the cortical control of saccade initiation[5,12]. Figure 2A shows the average normalized population activity of 57 SC and 22 FEF neurons, affirming the visuomovement pattern of activity discussed above. We then used the trial-averaged activity of individual neurons to construct a pseudo-population, resulting in an "expected neural trajectory" on a given trial (Figure 2B), and used this as the input to the temporal stability computation (for more, see Methods and EDF 2). Since trial-averaged responses are much smoother than single trial population responses, the stability profiles observed here were relatively smoother and closer to unity overall compared to the profiles in the previous section, especially during the delay period (Figure 2C). Notably, the key observations were in line with those observed with simultaneous population recordings – the dramatic reduction in the stability during the visual burst, and stable activity during the premotor burst – in both SC and FEF, and were consistent across time separations (EDF 3A) and subjects (EDF 3B).

We then tested whether the temporal stability framework was obeyed by SC neurons beyond those that are involved in the generation of large saccades. Neurons in the rostral SC are tonically active during fixation and reduce their activity during larger movements[13] (Figure 2D). Importantly, they also project to the saccade burst generator[14]. Assuming population stability modulates the input drive to downstream structures, we hypothesized that the temporal structure in rostral SC neurons must decrease during large saccades but remain elevated during fixation, even when a visual burst occurs in other parts of SC and FEF. In other words, we expected the evolution of temporal stability across rostral SC neurons to be the inverse of what occurs in caudal SC - Figure 2E confirms this prediction. In addition to the antiphase relationship with caudal SC during large saccades, rostral SC is also known to play a causal role in the generation of microsaccades[15]. Therefore, we considered whether the rostral SC population exhibits stable temporal structure during the burst that generates microsaccades. This was indeed the case (Figure 2F). Thus, the population activity of neurons in rostral SC, including its temporal structure, supplements the pattern in other parts of SC in suppressing and initiating movements. Finally,

158 pooling neurons in SC and FEF to form a combined population (since they all project to saccade-
159 generating structures) did not impact the main result (EDF 3C).

160

161 **Causal discrimination of population temporal structure**

162 Thus far, we have shown that population temporal structure is a candidate measure that can be used to
163 discriminate between visual and premotor bursts, using only correlation with the presence or absence of
164 a movement. To test whether this measure is actively used by the brain in a causal manner, we
165 performed multi-site patterned microstimulation in SC. We first identified individual contacts of the
166 laminar probe that evoked low-latency saccades with suprathreshold stimulation, to ensure their
167 position within the intermediate layers of SC. We then designed temporally stable and unstable
168 stimulation patterns restricted to these contacts, limiting stimulation parameters for individual sites to
169 the sub-threshold regime and verifying that the overall stimulation was near-threshold (for details, see
170 Methods). Stable patterns were created with a linearly decreasing inter-pulse interval (IPI) sequence for
171 one contact (to simulate a burst) and scaling the IPIs for other contacts by a uniformly spaced factor. For
172 each stable pattern, we also created a *paired* unstable pattern by jittering the pulse times within a
173 window and shuffling spikes between contacts, preserving both the pulse count for each site and the
174 average pulse rate across the "population". An example pair of stable-unstable pulse trains is shown in
175 Figure 3A (top row). The bottom row shows the pulse rates determined from these trains, illustrating
176 the scaled rates for the stable pattern and the fluctuating rates for the unstable pattern for individual
177 sites, despite the comparable population rates on individual trials (thick traces in the bottom row of
178 Figure 3A; see EDF 4B,C for more examples). Note that the assignment of high and low rates to different
179 contacts was randomized across trials, resulting in similar trial-averaged pulse rates for any given
180 channel, for either type of pattern (Figure 3B). Figure 3C shows the temporal stability for stable and
181 unstable patterns for all trials in the example session, confirming the impression provided by the pulse
182 trains and rates, i.e., the scaled patterns are highly stable compared to the relative instability of the
183 jittered patterns.

184 We delivered these stimulation patterns during the "gap period" in a gap saccade task (EDF 4A). Figure
185 3D shows a scatter plot of the stimulation-aligned saccade latencies observed using stable versus
186 unstable stimulation patterns for all pairs from one session. For the majority of stable-unstable trial
187 pairs, a saccade was evoked with the stable but not the unstable pattern (points in the blue shaded
188 window spanning the stimulation duration in Figure 3D). The relative likelihood of evoking a movement
189 with the stable stimulation pattern only (as described in the inset in Figure 3D) is shown as a function of
190 session number in Figures 3E (blue points). For most sessions, the observed relative likelihood values
191 were significantly higher (i.e., biased towards the stable pattern; $p < 0.01$, permutation tests) than the
192 null distribution (gray points) obtained by shuffling trials with randomly assigned stable-unstable
193 identities, reinforcing the observation that the stable pattern was more likely to evoke a movement
194 compared to a *state-matched* unstable pattern. We also analyzed the movement vectors evoked by
195 stable and unstable patterns – both sets of saccades, when they occurred, were similar to each other in
196 both amplitude and direction (EDF 5; $p > 0.01$ for most sessions, Wilcoxon signed-rank tests) and to the
197 "fixed vector saccades" evoked by constant frequency multi-channel stimulation, in amplitude (EDF 5; p
198 > 0.01 for most sessions, Wilcoxon rank-sum tests; for details, see Methods).

199

Comparison with other population-based models of movement generation

Having demonstrated the role of population temporal structure in saccade initiation using both correlative and causal approaches, we sought to disambiguate the stability framework from extant models of movement initiation. We earlier argued for the implausibility of threshold-based gating based on the fact that the population activity during the visual burst can sometimes exceed premotor activity (Figure 1C). Could other population activity-based mechanisms, such as the optimal or potent subspace model that seems to govern movement generation in the skeletomotor system[16-18], play a role here? To verify this, we first used factor analysis (FA) to visualize the low-dimensional neural states in the visual and premotor bursts[19]. An example 3-dimensional FA projection of the two bursts is shown in Figure 4A – the two sets of states are clearly separable, likely an effect of the subtle yet distinct trends in visual versus premotor activity levels along the dorso-ventral extent of the SC (e.g., Figure 1B). Indeed, a linear discriminant analysis (LDA) classifier was able to easily discriminate between the visual and premotor states, confirming the result of visualization (purple points in Figure 4B, also see EDF 6). The properties of neural activity in SC therefore seem to be consistent with a static state space code as well.

Next, we used data from the patterned microstimulation experiments to disambiguate between the optimal subspace and temporal stability frameworks. Using a representative low-dimensional FA projection to visualize the stimulation patterns was infeasible because all the electrode sites contributed equally to the patterns, due to the random assignment of pulse rates across sites on different trials. Indeed, it was evident that the stable and unstable patterns were indiscriminable in state space as shown in an example 3-dimensional projection (Figure 4C; the eigenspectrum in the inset shows gradually increasing cumulative variance with FA dimension). To confirm this, we trained an LDA classifier to discriminate between stable and unstable patterns based on population pulse states alone, and never observed above-chance classification (Figure 4D, also see EDF 6). This result provides a key piece of evidence in support of the notion that the brain in fact uses temporal stability information, since the evoked behavior reflected the difference between stable and unstable patterns while a linear readout of population states did not. However, it is still possible that certain sites or dimensions are more potent in evoking movements compared to others, a property which may not be revealed when classifying patterns based on stability alone.

To explicitly test this possibility, we trained another LDA classifier to discriminate between trials in which stimulation evoked a movement and those where no movements were evoked. In order to estimate the effect of population states independent of the contribution of temporal stability, we divided the pairs of stable-unstable trials into two subsets – the stability subset (SS), which was made up of trial pairs where only one of the stable-unstable pair evoked a movement, and the neutral subset (NS), made up of trial pairs where either both or neither of the stable-unstable pair evoked a movement (Figure 4E). When trained on the NS trials alone, the linear decoder was able to successfully discriminate trials in which a movement was evoked (orange points in Figure 4F), indicating that the population pulse states contained significant information about movement initiation likelihood. Since the stable-unstable pairs were matched in terms of whether a movement was evoked or not for this subset, this readout of the population pattern was independent of its temporal structure. In contrast, a linear decoder trained on the SS trials was unable to determine whether a movement was evoked on a given trial with above-chance accuracy (brown points in Figure 4F). Crucially, when the population pulse states were supplemented with stability information (added as a native dimension in the input to the decoder), classifier performance increased to reflect the trend in relative likelihood estimates in Figure 3E (green

points in Figure 4F, also see EDF 6). This analysis thus reveals two independent contributors to movement initiation in SC – a static "laminar" code, presumably related to the optimal or potent subspace, that may be a function of the distribution of preferred stimulation sites, and a dynamic "temporal" code, where stability of the population pattern controls movement initiation even under matched state space conditions.

We also tested how the temporal stability hypothesis was related to modified threshold-based mechanisms that rely on pooling the activity of a correlated ensemble of neurons[20]. In this framework, correlations in the accumulation rates of neurons influence movement initiation (reaction times) under a given pooling scheme, with higher correlations leading to shorter reaction times. Spike count correlations between SC neurons in the visual and premotor epochs (for details, see Methods) were inconsistent with this notion – the correlations during the visual epoch were slightly, but significantly, higher compared to the premotor epoch (EDF 7A, mean difference between visual and premotor epoch correlations = 0.0167 > 0, p = 2.5E-6, one-tailed t-test). We also computed pulse count correlations for the stable and unstable stimulation patterns and found no difference (EDF 7B, median difference in correlations between conditions = 4.11E-4, not significantly different from 0, p = 0.279, Wilcoxon signed-rank test), despite the significant difference in their effects on behavior. Thus, the temporal stability framework is, to a large extent, independent of extant models of movement initiation.

Biophysical models for decoding population temporal structure

Finally, we sought to identify a mechanism by which downstream neurons could discriminate between stable and unstable population codes. We modelled the decoder as a spiking neuron that receives population inputs through its network of dendrites (Figure 5A). The decoder can be thought to represent neurons in the pons that receive and integrate inputs from the superior colliculus and burst for saccades[14]. To mimic the potent inhibitory gating (and disinhibition during saccades) provided by the omnipause neurons (OPNs) on the burst neurons, we also included a spiking disinhibitor unit with reciprocal inhibitory connections with the decoder. The disinhibitor also received both excitatory and inhibitory inputs from the same population as the decoder[21], creating a balance between excitation and inhibition that must be overcome in order to produce a saccade[22]. How might the population pattern be converted into a signal that initiates a movement only when the inputs are stable? Conceptually, the decoder should have a mechanism to keep track of the short-term history of population activity, use this history to evaluate temporal structure, and respond selectively when the activity pattern is deemed stable over the time scale of integration. We incorporated these heuristic requirements by using state-dependent modulation of input-evoked excitatory post-synaptic potentials in the decoder, i.e., inputs arriving when the local post-synaptic potential was depolarized caused a higher subsequent change in the potential[23,24] (Figure 5B).

We simulated the model with the stable and unstable patterns from the microstimulation experiments as inputs, since they readily offered mean-matched input sequences differing only in temporal structure. We characterized the efficacy of integration by the decoder as the time of first spike in the burst, a proxy for movement initiation latency. Although there were several pairs of trials for which both the stable and unstable inputs caused the decoder to spike or neither did (Figure 5C), there were a number of pairs for which only the stable input pattern was decoded (arrow in Figure 5C). To facilitate visualization and comparison with experimental data (Figure 3D), latency values greater than the

stimulation duration were randomly assigned when the decoder wasn't recruited. Figure 5D shows the membrane potential output of the disinhibitor (cyan and magenta traces) and decoder (blue and red traces) units for a pair of stable (top row) and unstable (bottom row) input patterns from this subset of trial pairs (see EDF 8 for examples from the other subsets). For these pairs, the disinhibitor maintained a tonic firing rate throughout the trial when the inputs were unstable, and sharply reduced its activity at some point during stable stimulation (Figure 5E - magenta and cyan traces, respectively). This latter disinhibition on stable trials was in anti-phase relationship with the bursting exhibited by the decoder, which was absent on unstable trials (Figure 5E - blue and red traces, respectively). The firing rate profiles of the decoder and disinhibitor units share a striking resemblance to medium-lead burst neurons and OPNs, respectively, in the pons[25]. Thus, a relatively simple, yet biophysically realistic, module seems capable of discriminating between population inputs based solely on their temporal characteristics, offering a putative mechanism by which downstream networks could use temporal information to make decisions about movement generation. We also found that a firing rate-based accumulator model with short-term synaptic plasticity produces comparable results (EDF 9), demonstrating the flexibility of various biophysical mechanisms in their ability to decode temporal structure and hence, the model-independence of this result.

Perspectives on the role of temporal structure

Neurons in premotor structures are constantly bombarded with information from thousands of presynaptic neurons that are active during sensorimotor processing. It is unclear how activity relevant for movement initiation is discriminated from activity related to other processing. Extant models of movement generation that rely on firing rate, including rise-to-threshold[5], inhibitory gating[25], and dynamical switches at the population level[16,17], leave certain explanatory gaps unfilled. A canonical model of movement initiation, especially in the oculomotor system, is threshold-based gating[5] (Figure 6A-B). Current knowledge points to a role by the OPNs in defining the threshold and controlling saccade initiation[5,26]. However, thresholds vary across behavioural paradigms[6,27], raising the question of how the threshold is set in a particular condition. Furthermore, evidence that the threshold changes during the course of a trial purely based on OPN activity is limited[26]. Critically, the existence of trials in which the population activity during the visual response exceeds premotor activity strongly reduces the likelihood of thresholding operating as a singular mechanism. An extension of simple thresholding is a mechanism based on the pooled activity of neurons with varying degrees of correlation in the population[20]. However, we show that population correlation statistics are not necessarily related to temporal structure or saccade initiation (EDF 7). A related explanation is that a movement is initiated only when movement-related neurons are active, but most neurons in sensorimotor structures likely span a continuum between having visuomovement activity to pure movement activity[1,28].

Other models, primarily in the skeletomotor system, posit that muscles are recruited and a movement is initiated when neural activity traverses certain optimal regions of the population state space[16] and is inhibited otherwise, e.g., during movement preparation[17,18] (Figure 6C). While the optimal subspace and nullspace hypotheses certainly seem to be consistent with recorded neural activity in SC (Figure 4A,B), they cannot readily account for the observation that neck[29] and upper limb[30] muscles are recruited time-locked to the visual target, as proxied through electromyography. We instead show that temporal stability plays an independent role in determining movement initiation, even when the putative potent

subspace is matched (Figure 4E, green symbols). We reason that for a decoder downstream, it is important to ensure the stability - or consistency over time - of the input code while processing it in order to influence the motor output. Population activity that creates a high firing rate drive but is inconsistent should be prevented from triggering a movement. This requirement is critical especially in the case of ballistic movements such as saccades, where the ability to reverse the decision once the movement has been initiated is limited. Thus, we present a novel mechanism, wherein both high firing rate and consistent temporal structure are necessary conditions for movement initiation, allowing for population trajectories that traverse away (increasing overall rate) but within hypersectors that fan out (stable temporal pattern) from the state space origin (Figure 6D). Note that this model does not preclude population activity from traversing a localized "optimal" region of state space, possibly due to hardwired network constraints. Moreover, the proposed mechanism is also consistent with classical population vector decoding schemes, since any weighted population vector readout during the sensory and movement epochs will produce decoded movement vectors that are largely similar. This allows for the same physical movement to be planned and executed by a given population of neurons in the labelled line sense, a possibility that is precluded by weighted neural readout mechanisms of movement preparation and generation, such as the potent/nullspace models.

Our findings are closely related to the premotor theory of attention[31] and offer a way to reconcile the attention- intention debate. They could also account for the mirror-like activity recorded during both action observation and execution in neurons known to project directly to motoneurons in the spinal cord[32]. In both cases, it is unclear how the same neuronal population represents two distinct signals that serve different functional roles. The results presented here suggest that this multiplexing ability may be provided by the distinct temporal structures of population activity patterns. Indeed, intracellular recordings have demonstrated that visual stimulation drives cortical networks into an asynchronous state[33], which may be a critical requirement for sensory processing. In addition, the differential effects of stable and unstable population stimulation patterns are reminiscent of the desynchronizing effects of patterned deep brain stimulation recently developed for the treatment of Parkinson's disease[34]. In this so-called "coordinated reset" approach, sequential, non-overlapping spatiotemporal sequences are more effective at resetting the affected neuronal population from the pathological synchronized state to an asynchronous state, enabling better movement control[35].

Previous studies have looked at the role of precise coordination in the timing of incoming spikes in the transmission of information and the efficacy of driving the recipient neuron[9,10]. However, input firing rates may vary greatly across the population, limiting the ability to compare to spike times. Our study proposes a mean-field equivalent to the spike-based temporal or correlation code[11,36] by looking at the temporal structure of population firing rates, thus tying together the notion of rate, temporal, and population codes. We suggest that temporal structure of population activity is critical to understanding movement generation as well as, more broadly, neuronal communication and its relationship to behaviour.

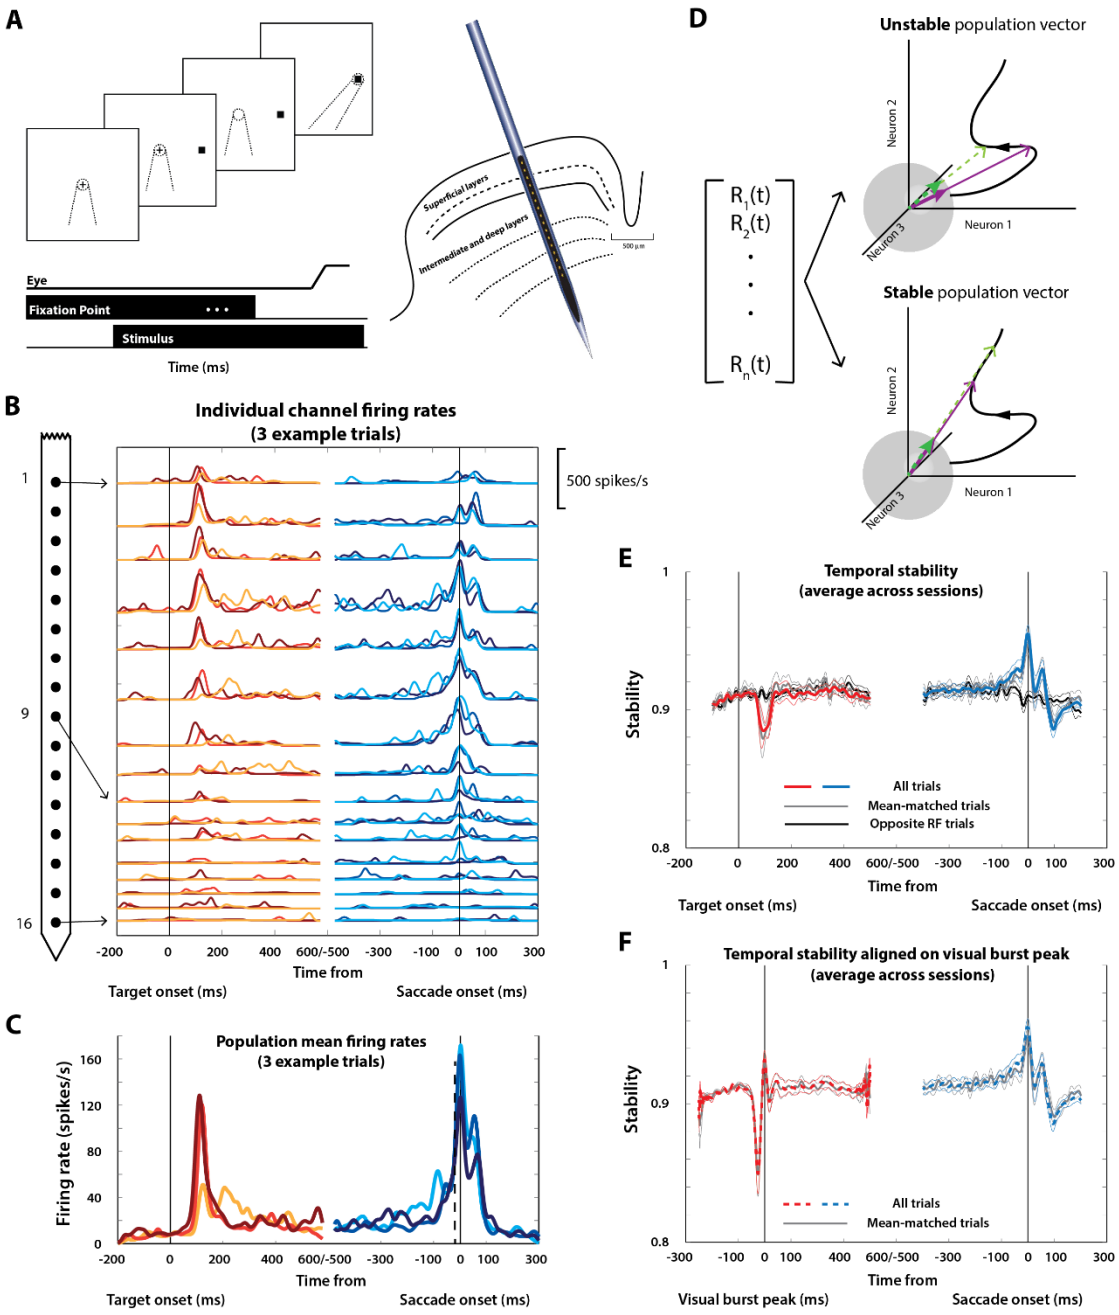

Figure 1. **Temporal structure of population activity during sensorimotor multiplexing. A.** Left - Sequence of events in the delayed saccade task. The top row shows a typical display sequence and the bottom rows show the timeline. The fixation target is shown as a plus symbol for illustration purposes. Dotted lines depict line of sight of the animal. Right – Linear electrode arrays with 16-24 contacts were used to penetrate the SC normal to the surface. Thus, the recordings were along a dorso-ventral "column" of the SC. **B.** Example recording session showing activity from 16 channels aligned on target (left column) and saccade (right column) onsets for 3 different trials (traces with different colour saturations). **C.** Population activity averaged across the 16 channels for the 3 trials shown in B. Note the considerable variation in the amplitude of the visual burst. **D.** Schematic depicting computation of temporal stability. The population vector (left column) was used to construct a single-trial neural

trajectory, which was then normalized at each instance to represent a unit length population vector traversing along a hypersphere (right column). The top panel in the right column highlights an unstable part of a schematic trajectory and the bottom panel emphasizes a stable part. Note that in both cases, the two population vectors (thin arrows in green and purple) are separated by a similar length of time. In the unstable case, the normalized vectors (thick arrows) move around on the surface of the hypersphere (green and purple vectors are spatially separate), whereas in the stable case, they stay pointed roughly in the same direction (the two vectors overlap in space). **E.** Temporal stability is computed as a dot product between two normalized population vectors separated by $2\tau = 20\ ms$ aligned on target (left) and saccade (right) onsets. The stability of population activity drastically decreases during the visual burst (red traces, mean +/- s.e.m.; $p < 0.01$ for Wilcoxon signed-rank test with respect to baseline) and increases before and during movement onset (blue traces; $p < 0.01$ for Wilcoxon signed-rank test with respect to baseline). The gray trace is the stability computed only for those trials where the peak visual burst matched or exceeded activity during the premotor burst at saccade initiation (after accounting for an efferent delay of 15 ms). The black trace shows population stability on trials in which the target and saccade were directed away from the response field (RF) of the neurons. **F.** Temporal stability aligned to peak visual burst. The peak time was computed on the average population response of each trial, and the peak-aligned trace was averaged across trials and sessions. Note the sharper dip in the stability of the visual response when aligned to the peak. Like in E, the gray trace shows peak-aligned stability for trials in which the peak visual activity matched or exceeded premotor activity at saccade initiation.

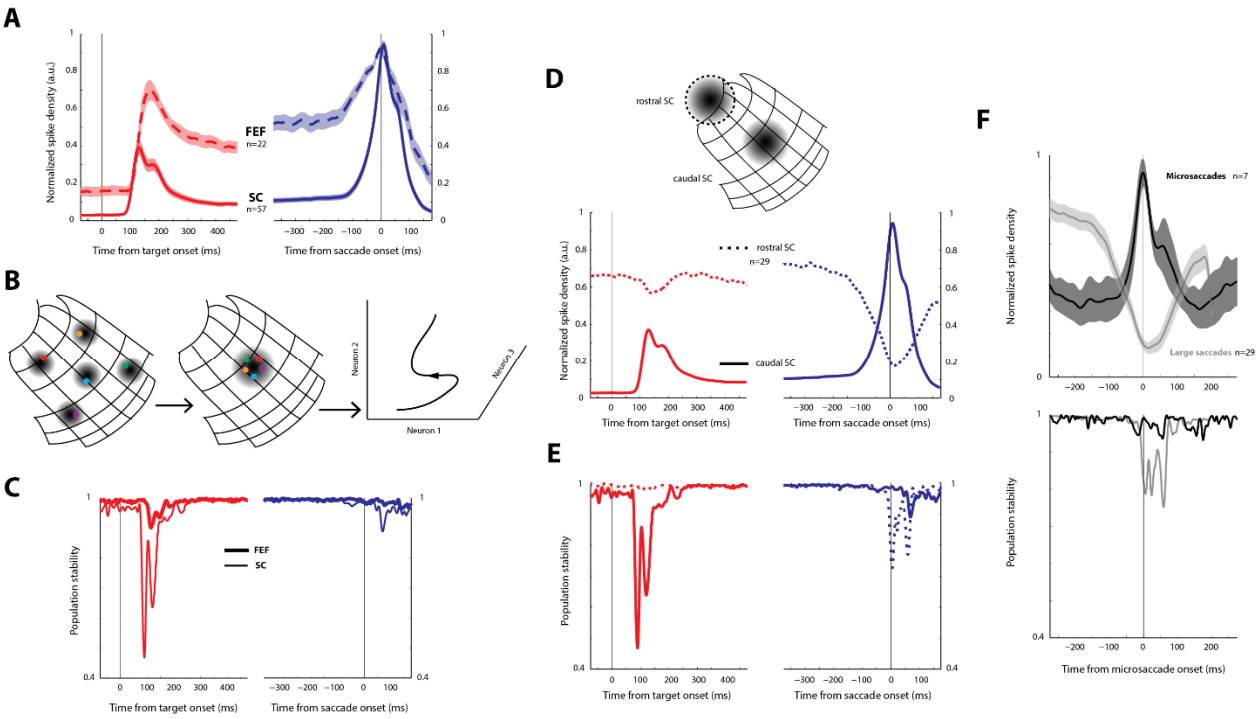

Figure 2. **Temporal stability in sub-populations of SC and FEF. A.** Normalized mean population activity of visuomovement neurons recorded in SC (solid) and FEF (dashed) during the delayed saccade task. The left and right panels show data aligned on target and saccade onsets, respectively. **B.** Inferring population dynamics from single-unit recordings. We combined recordings made in different active populations on the SC map (left) into one pseudo-population (middle) that is active for any saccade, enabling the construction of an expected population trajectory (right) in neural state space (see *Estimating Population Dynamics* in Methods and EDF2 for more details). **C.** Temporal stability for the SC (thin traces) and FEF (thick traces) pseudo-populations. Population stability shows a drastic reduction at the time of the visual burst but is stable before and during the onset of the movement, consistent with the simultaneously recorded population in the previous figure. **D.** Top row - Schematic of the SC depicting active population regions in the rostral SC (dotted boundary) and caudal SC (plain). Bottom row – Mean normalized activity in rostral (dotted) and caudal (solid) SC neurons during the delayed saccade task. **E.** Temporal stability of the rostral SC population (dotted) with the caudal population shown for comparison. **F.** Population activity of rostral SC neurons is stable during microsaccades. Top – Normalized population activity of microsaccade-related neurons in rostral SC (black). Note the strong burst compared to the suppression in fixation-related neurons in the same region for large saccades (gray). Bottom – Temporal stability of the rostral SC population for microsaccades shown in comparison to large saccades. Neurons which burst for microsaccades remain stable during these fixational eye movements, consistent with the hypothesis that both an increase in firing rate and high stability is required for movement generation.

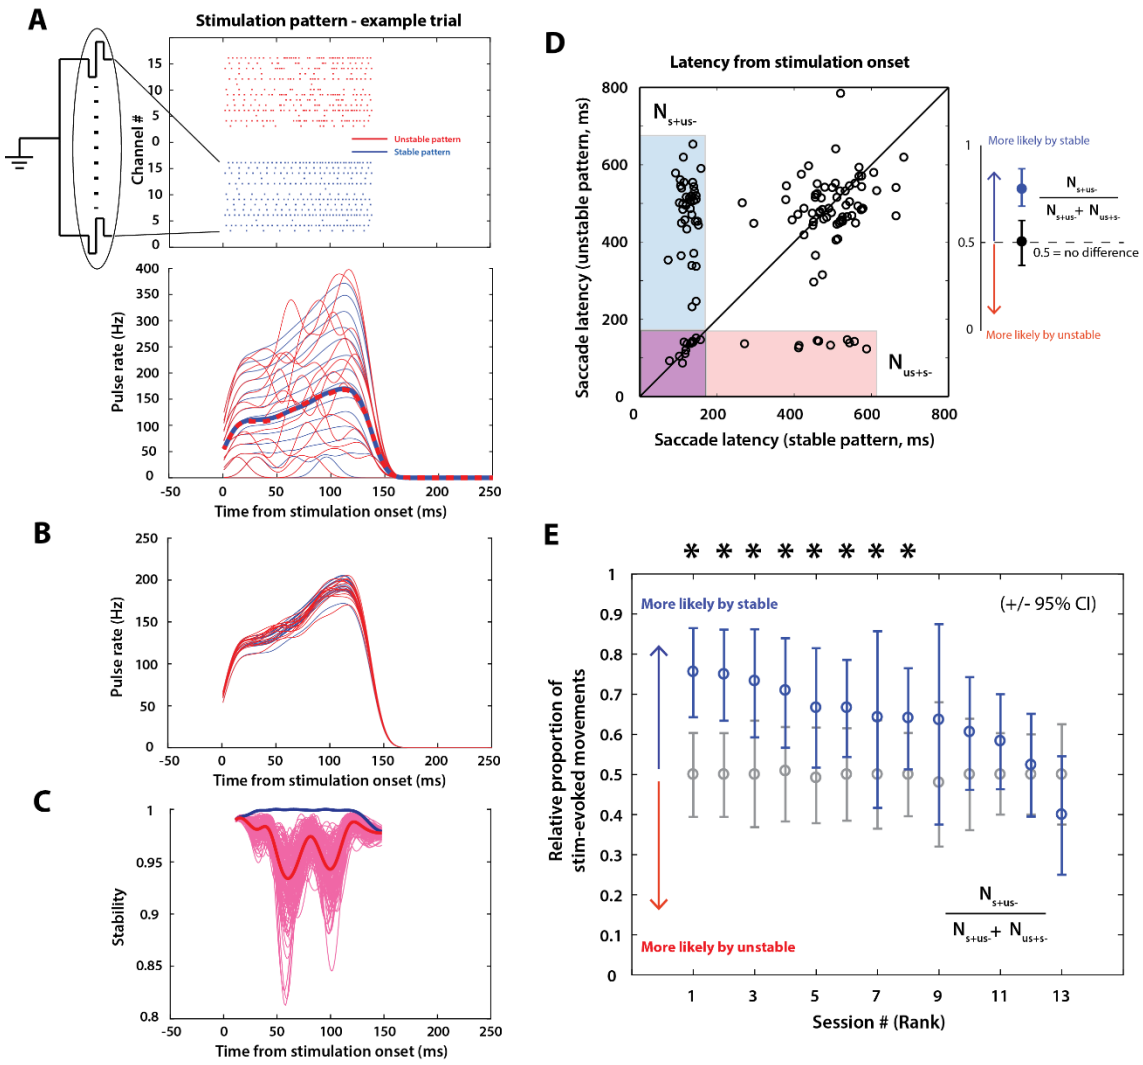

Figure 3. **Patterned microstimulation supports temporal stability as a model of movement initiation. A.** Example pair of stable-unstable stimulation patterns. Top row – Pulse trains for 14 of 16 channels that were stimulation-viable with suprathreshold parameters for this session. Red and blue trains indicate an example pair of unstable and stable patterns, respectively. Bottom row – Pulse rates for the 14 channels in the stable-unstable pulse patterns, with the population rates averaged across channels overlaid as thick traces. Note that the rates in the unstable pattern are highly fluctuating despite the matched population average. **B.** Average pulse rates for each channel across all trials for the stable (blue) and unstable (red) patterns used in this session. The clustering of trial-averaged rates in a narrow band is the result of randomization of pulse rate assignment to individual channels on different trials. **C.** Temporal stability of the stable and unstable patterns (blue and pink, respectively) for all trials in the example session. **D.** Scatter plot of saccade latencies (relative to stimulation onset) for the example session. Each point reflects the outcome of stimulation with stable and unstable pulse patterns constituting each pair. The points in shaded regions are stimulation-evoked for at least one condition. N_{s+us-}, points in the light blue shaded box, reflects the number of trials in which a stable stimulation pattern evoked a saccade but its unstable pair did not. N_{us+s-}, points in the light red shaded box, denotes the number of trials in which an unstable stimulation pattern evoked a saccade but its stable pair did not. Points in the purple box show subset of trials in which both stable and unstable pairs evoked a saccade. Points in the unshaded (white) region refer to trials in which neither stable or

unstable pair evoked a saccade. For these pairs, the trial was assigned the latency of the saccade (relative to stimulation onset) directed to a target presented after the gap period. The inset shows the calculation of relative likelihood of the stable pattern evoking a stimulation-evoked movement for a given session. **E.** Relative proportions of the stable pattern evoking a movement for each session (blue points, error bars represent 95% CI of the bootstrapped distribution). The sessions are sorted in descending order of proportion for viewing clarity. Note that this sorting order is used in all subsequent figures depicting individual sessions (in Figure 4). Gray points and error bars are computed from a surrogate dataset in which stable/unstable trial identities are completely shuffled (mean +/- 95% CI). Asterisks above a particular session denote a significant effect ($p < 0.01$ on the permutation test).

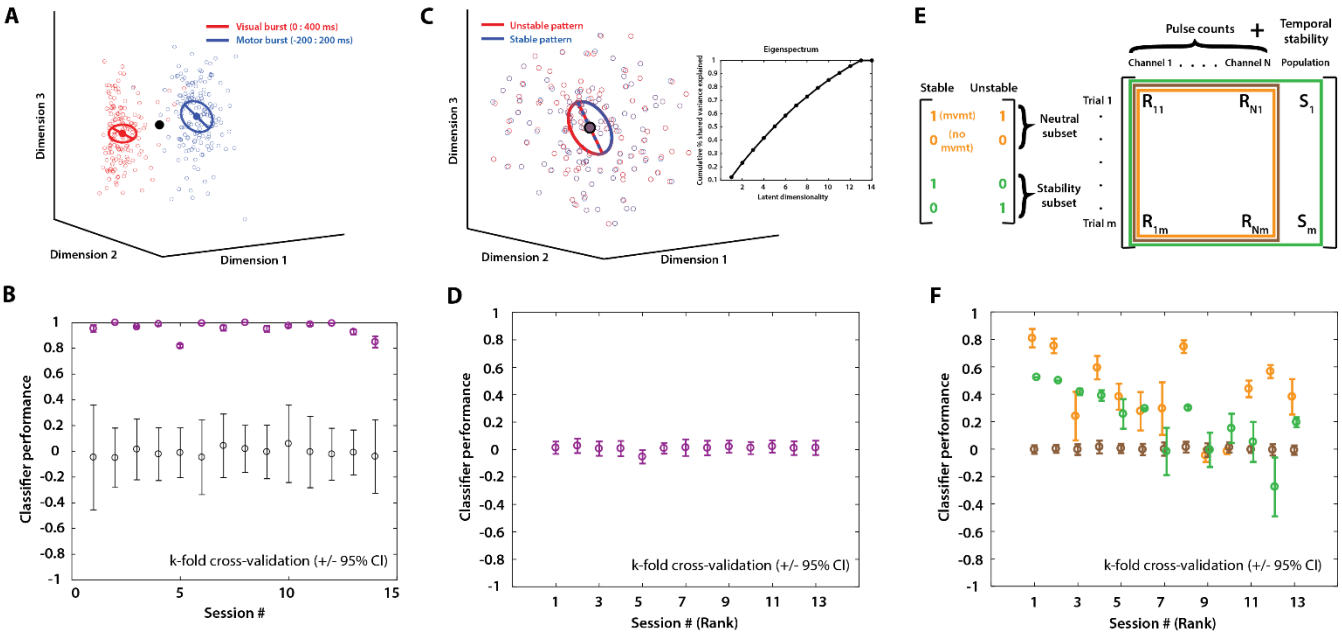

Figure 4. Linear discriminability of population activity and microstimulation patterns. A. Three-dimensional factor analysis (FA) projection of population activity states during the visual and premotor bursts (red and blue points, respectively; visual burst: 0 to 400 ms after target onset, premotor burst: -200 to 200 ms around saccade onset). The dimensions were chosen arbitrarily to illustrate the separation of the two states. The solid lines and ellipses show the axes of maximum variance and covariance ellipses for each cluster, respectively. **B.** Performance of a linear discriminant analysis (LDA) classifier on discriminating between visual and premotor states estimated using the Matthews correlation coefficient (MCC, chance is 0). The purple points show the mean +/- 95% CI computed from k-fold cross-validation on the actual data. The black points show the same but for a shuffled dataset in which the visual and premotor labels were randomly permuted, indicating chance performance distribution for each dataset. For all sessions, classification performance was significantly above chance ($p < 0.01$, one-tailed t-test). These sessions were ordered arbitrarily do not correspond to those associated with patterned microstimulation experiments. **C.** Three-dimensional FA projection of population pulse states for stable and unstable microstimulation patterns (blue and red points, respectively). The projection dimensions were chosen arbitrarily, but no view showed any reasonable separation between the two clusters. The eigenspectrum in the inset illustrates the difficulty of choosing a good projection, since all native dimensions contributed equally to the variance in the data, due to the randomized assignment of pulse rates to channels across trials. **D.** LDA classification performance on discriminating stable and unstable stimulation patterns based on population pulse states alone. For all sessions, classification performance was not significantly different from chance ($p > 0.05$, two-tailed t-test). **E.** Schematic of the approach to partition each dataset for classification of stimulation-evoked movement occurrence. Left column - Trial pairs in which both or neither the stable and unstable pattern evoked a movement were grouped into the neutral subset, and pairs in which only one of the stable or unstable pattern evoked a movement were grouped into the stability subset. Right column - For both subsets, the classifier was first trained on pulse counts alone (part of the matrix highlighted by the orange and brown boundaries – match colours in next panel). Separately, for the stability subset, the classifier was also trained on an additional dimension of temporal stability values (full matrix highlighted by the green boundary). **F.** LDA classification performance on discriminating trials in which stimulation evoked a movement from those in which it did not. The orange points correspond to the neutral subset, and the green and brown points correspond to the stability subset. The brown points are with the classifier trained solely on the population pulse states. The green points are with the addition of temporal stability in the input to the classifier. For most sessions, addition of temporal information increased

475 classification performance ($p < 0.01$, one-tailed t-test of a difference between the cross-validated performance
476 distributions with and without addition of temporal stability). In all relevant panels, sessions are sorted by the
477 order determined in Figure 3E based on relative likelihood of evoking a movement.

478

479

480

481

482

483

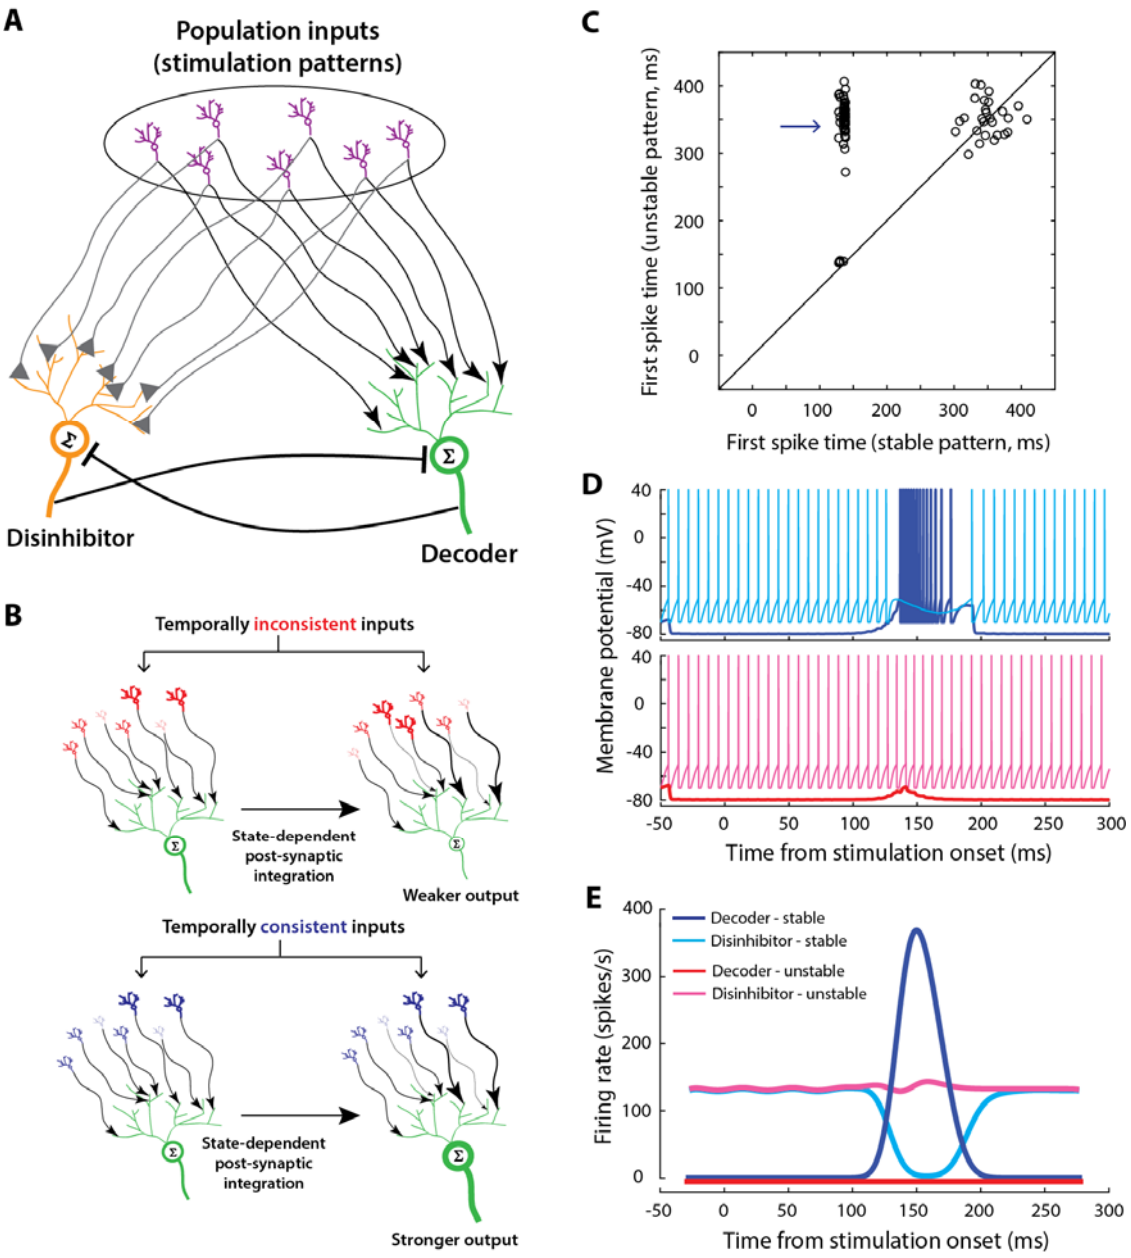

Figure 5. Spiking neuron model with dendritic processes can discriminate population temporal structure. A. Schematic of the model architecture. The decoder (green neuron), representing high frequency burst neurons in the brainstem, receives population inputs (output of purple neurons, or, in this case, microstimulation pulses) through its dendritic network (excitatory inputs, represented by the arrowhead terminals). The disinhibitor (orange neuron), representing pontine omnipause neurons, also receives inputs from the same population (both excitatory and inhibitory[21], represented by the triangular terminals), in addition to a constant current that produces tonic firing activity (not shown). The decoder and disinhibitor are both modelled as leaky integrate-and-fire spiking neurons and mutually inhibit each other (flat terminals). **B.** Schematic of the model heuristic. Each input neuron's activation level is represented by its size and thickness. The effect of an input spike at the local dendritic site (i.e., excitatory post-synaptic conductance changes) is represented by the thickness of the arrow. Two time points are shown for illustration (left column – arbitrary initial point, right column – subsequent time point). In the top row,

497 unstable inputs lead to a scenario where the post-synaptic conductances are no longer aligned to the strong inputs
498 that initially created them, whereas in the bottom, stable inputs lead to matched strong input and post-synaptic
499 conductances, resulting in stronger accumulation and firing. **C.** Scatter plot of first spike latencies in the decoder
500 (putatively representing saccade initiation) for stable versus unstable inputs from matched pairs. The input
501 duration was 150 ms and thus only latency values <150 ms correspond to actual first spikes produced by the
502 decoder. To facilitate visualization, latency values were assigned randomly for trials in which the decoder did not
503 generate any spikes, sampled from a distribution with an arbitrarily selected mean of 350 ms. Note the occurrence
504 of a number of trial pairs for which only the stable input causes the decoder emit spikes (blue arrow). **D.** Simulated
505 membrane potential of the decoder (blue and red traces) and disinhibitor (cyan and magenta traces) for an
506 example matched trial pair with stable (top row) and unstable (bottom row) input patterns. **E.** Average activity of
507 the disinhibitor (cyan and magenta traces, stable and unstable inputs, respectively) and decoder (blue and red
508 traces, stable and unstable inputs, respectively) for all trial pairs in which only the stable input produced a
509 spike/movement (i.e., trial pairs indicated by the arrow in panel D).

510

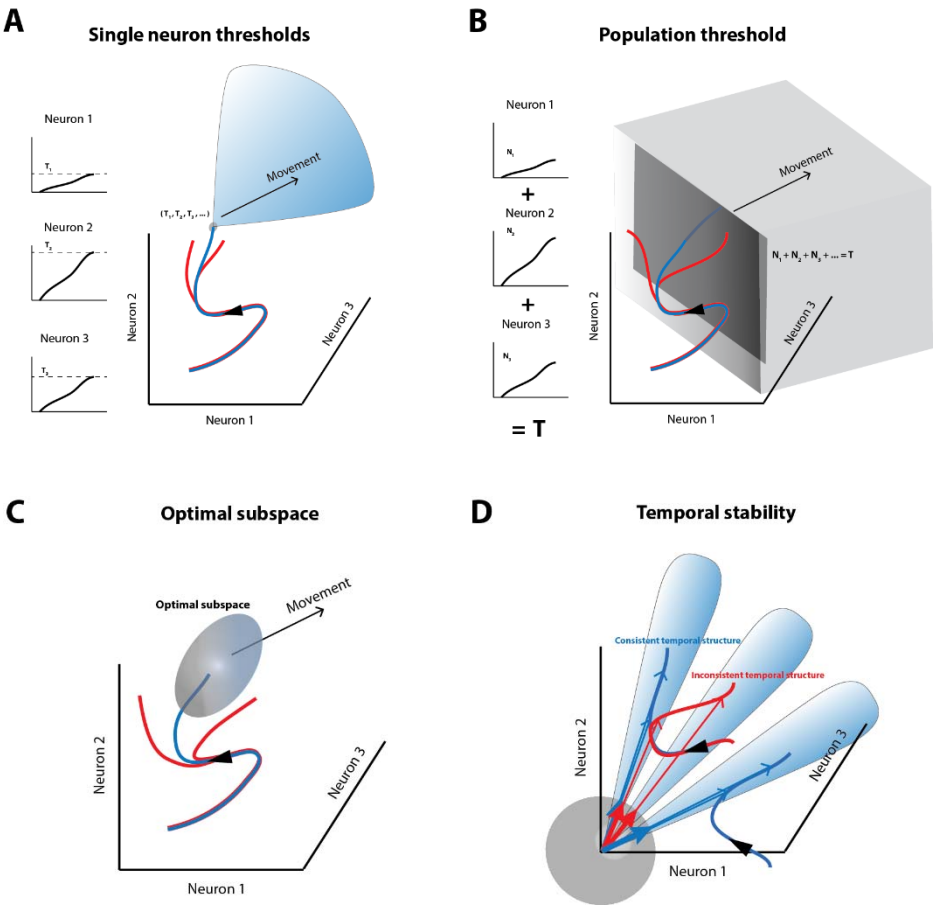

511

512

513 **Figure 6. Summary of models of movement preparation in a state space framework. A** and **B.** Multi-dimensional
514 representations of the threshold hypotheses. Panel A depicts single neuron thresholds, i.e., the activity of each
515 neuron must rise to a fixed threshold at the same time in order to initiate a movement. Fixed thresholds for each
516 neuron are equivalent to a point (or a small region, represented by the gray sphere) in neural state space. Activity
517 profiles that meet this criterion (e.g., blue trajectory) and beyond (blue region) lead to movement generation while
518 those that don't meet this criterion (red traces) do not result in a movement. Alternatively, activity may need to
519 cross a fixed threshold at the population level (sum of neurons' activities = constant, i.e., an n-1 dimensional
520 hyperplane, dark gray plane), depicted in panel B. **C.** The optimal subspace hypothesis offers more leeway by
521 allowing neuronal activity to reach a relatively larger region of population state space (blue trajectory and gray
522 ellipsoid) over a period of time in order to signal the initiation of a movement. Activity trajectories that evolve
523 outside this "optimal subspace" (red traces) do not lead to a movement. **D.** The temporal stability hypothesis
524 suggests that a burst of neural activity that is consistent over time in state space, i.e., an activity trajectory that
525 points in the same direction (blue trajectories and vectors) is likely to be interpreted as a movement command by
526 a decoder, while high-frequency activity that is inconsistent (fluctuating directions, red trajectories and arrows) is
527 not. It does not matter where in state space the activity happens to evolve – a different subpopulation of neurons
528 could be active, but as long as they are pointing in the same direction (i.e., evolving within one of the blue sectors),
529 it will lead to a movement. The gray sphere around the origin represents a unit hypersphere for visualization of
530 back-projected unit vectors.

531

532

Online Methods

General and surgical procedures

All experimental and surgical procedures were approved by the Institutional Animal Care and Use Committee at the University of Pittsburgh and were in compliance with the US Public Health Service policy on the humane care and use of laboratory animals. We used three adult rhesus monkeys (*Macaca mulatta*, 2 male, ages 8 and 6, and 1 female, age 10) for our experiments. Both SC and FEF were recorded in monkeys BB and BL whereas only SC was recorded in monkey WM. Under isoflourane anaesthesia, recording chambers were secured to the skull over craniotomies that allowed access to the SC and FEF. In addition, posts for head restraint and scleral search coils to track gaze shifts were implanted. Post-recovery, the animal was trained to perform standard eye movement tasks for a liquid reward.

Visual stimuli and behaviour

Visual stimuli were displayed either by back-projection onto a hemispherical dome or on a LED-backlit flat screen monitor. Stimuli were white squares on a dark grey background, 4x4 pixels in size and subtended approximately 0.5° of visual angle. Eye position was recorded using the scleral search coil technique or using an EyeLink 1000 eye tracker, both sampled at 1 kHz. Stimulus presentation and the animal's behaviour were under real-time control with a Labview-based controller interface[37]. All monkeys were trained to perform standard oculomotor tasks. In the delayed saccade task (Figure 1A), the monkey was required to initiate the trial by acquiring fixation on a central fixation target. Next, a target appeared in the periphery but the fixation point remained illuminated for a variable 500-1200 ms, and the animal was required to delay saccade onset until the fixation point was extinguished (GO cue). The gap task (EDF 4A) was used for the patterned microstimulation experiments. In this task, initial fixation on a central target was followed by a gap period (200-300 ms) during which the fixation point disappeared, while the animal was required to maintain fixation at the now vacant location. On stimulation trials, microstimulation pulses were delivered 100 ms into the gap period and window constraints were relaxed to allow for the stimulation-evoked movements. This was followed by the appearance of a saccade target in the periphery which was also the GO cue for the animal to make a target-directed saccade. All animals performed these tasks with >95% accuracy. Incorrectly performed trials were removed from further analyses. The tasks were occasionally interleaved with visual search paradigms used for a different study.

Experimental sessions

We used data from several different types of sessions in this study. Each session was either a laminar recording session (n = 16 sessions), a single-unit recording session (n = 108 sessions), or a patterned microstimulation session (n = 13 sessions). The laminar and single-unit recording sessions employed mostly the delayed saccade task with a few microstimulation trials to verify electrode presence in SC (or FEF) and estimate the location on the topographic map. The patterned microstimulation sessions were restricted to the gap task, with stimulation trials and non-stimulation trials interleaved.

572

573 *Laminar recordings and single-unit electrophysiology*

574 During each laminar recording session, a linear microelectrode array (LMA, AlphaOmega, Inc.) or a
575 Plexon V-probe (Plexon, Inc.) was inserted into the SC chamber using a hydraulic microdrive. Neural
576 activity was amplified, digitized and recorded using the Grapevine Neural Interface Processor (Ripple,
577 Inc.) and visualized using the associated Trellis interface. Neural activity was band-pass filtered between
578 500 Hz and 5 kHz to record spiking activity and between 0.1-250 Hz to record local field potentials for
579 another study. Approach towards the SC surface was identified by luminance-based visual modulation of
580 activity in the lowermost channels, after which the electrode was driven down a further 2-3 mm until
581 known hallmarks of SC activity were observed in a majority of the channels. The presence of several
582 contacts in the intermediate layers was further confirmed by the ability to evoke movements with
583 single-channel microstimulation at these contacts. For single units, a tungsten microelectrode was
584 lowered into the FEF or SC. Neural activity was amplified and band-pass filtered between 200 Hz and 5
585 kHz and fed to a digital oscilloscope for visualization and spike discrimination. A window discriminator
586 was used to threshold and trigger spikes online, and the corresponding spike times were recorded. Both
587 SC and FEF were confirmed by the presence of visual and movement-related activity as well as the
588 ability to evoke fixed vector saccadic eye movements at low stimulation currents (20-40 µA, 400 Hz, 100
589 ms). Before beginning data collection for a given neuron or laminar recording site, their response field
590 was roughly estimated. In most data collection sessions with either electrode, the saccade target was
591 placed either in the neurons' response field or at the diametrically opposite location in a randomly
592 interleaved manner. In addition, stimulation-evoked saccades were recorded to identify the response
593 field centers (or "hotspots") for the cells recorded during that session. For recordings in rostral SC,
594 stimuli were presented at one of four locations at an eccentricity sufficient to induce a reduction in
595 activity during the large amplitude saccade.

596

597 *Data analysis and pre-processing*

598 Data were analyzed using a combination of in-house software and Matlab. Eye position signals were
599 smoothed with a phase-neutral filter and differentiated to obtain velocity traces. Saccades were
600 detected using standard velocity criteria. The animal was considered to be maintaining fixation if the
601 gaze remained within a 2-3° window around the fixation target. We also detected any microsaccades
602 that occurred during the delay period in each trial by using a velocity criterion based on the noise in the
603 velocity signal for that trial. Only one of the two monkeys (WM) in whom we recorded neural activity in
604 the rostral SC made sufficient number of microsaccades to permit further analysis.

605 Raw spike density waveforms were computed for each neuron (or multi-unit activity cluster) and each
606 trial by convolving the spike trains with a Gaussian kernel (width = 4ms; in some instances, we used 10
607 ms for display purposes only). For the laminar recordings, we analyzed channel activity on single trials
608 independently. Because direct kernel-based estimation of firing rates from binary spike trains on single
609 trials can be noisy, we first computed the inverse of the inter-spike intervals as a measure of the firing
610 rate prior to convolution with the smoothing kernel. All thresholded units were considered for further
611 analyses, regardless of their characteristics (so number of neurons matched number of channels for
612 these datasets). For single-electrode recordings, the spike densities were averaged across condition-

matched trials (same target location) following alignment with target or saccade onset. Neurons were classified as visuomovement neurons if the spike density was significantly elevated above baseline during the visual epoch (50-200 ms following target onset) and during the premotor epoch (50 ms before and after saccade onset). This resulted in 57 caudal SC neurons and 22 FEF neurons. In addition, rostral SC neurons were defined as fixation-related if the activity during the premotor epoch of large saccades was significantly reduced below baseline (29 neurons). A subset of these neurons also elevated their discharge around the onset of microsaccades (7 neurons). To minimize the effect of noise in the spike density waveforms due to insufficient number of trials in our analysis, we used only neurons which had at least 10 trials for a given condition. This was not a factor in most cases (we typically had 50-100 trials).

Estimating population dynamics from laminar and single-unit recordings

We define $\mathbf{R}(t)$ as a population activity vector:

$$\mathbf{R}(t) = \begin{bmatrix} R_1(t) \\ R_2(t) \\ \dots \\ R_n(t) \end{bmatrix}$$

where $\mathbf{R}(t)$ represents the instantaneous activity at time t as a point in an n-dimensional space, n is the number of neurons, and $R_i(t)$ is the spike density function of the i^{th} neuron. The curve connecting successive points over time is the neural trajectory that describes the evolution of population activity. For laminar recording data, a neural trajectory was determined for each trial. Analysis of single electrode data relied on pseudo-population analysis, for which each neuron's firing rate waveform $R_i(t)$ is the average across many matched trials (identical stimulus/response conditions). Thus, the neural trajectory is the *expected trajectory* of population activity. Since these neurons have a fairly broad RF, many neurons contribute to the active population for any given visual stimulus or saccade[38]. Our neurons were sampled roughly (but not exactly) around the hotspot of the active population for a given session. Therefore, the pooled data from individual sessions can be thought of as an approximation of the population mound active for an arbitrary location/RF in the visual field on any given trial. Many recent studies have reconstructed such "pseudo-populations" from sequentially recorded neurons and found comparable properties from simultaneous and serial recordings[17,39,40]. Indeed, this is also expected of our pseudo-population under the assumption of isotropy – that each neuron's contribution to its respective local active population is similar regardless of the locus of the population, and consistency between the results we observed in the laminar dataset and pseudo-population confirms this. To better demonstrate this, we estimated the location of a given neuron in the active pseudo-population as follows (we use SC for illustrative purposes because of its convenient topography). We used the point image of the target location on the SC map as a representation of the center of the active population for that session, and used the stimulation-evoked saccade vector to identify the location of that neuron on the SC. We referenced the point image of each target location to a single location to create an active pseudo-population and translated the neuron locations relative to this population center (EDF 2). All mathematical equations for transforming between visual and SC tissue coordinates have been defined previously[41]. Inclusion of stimuli/saccades in the anti-preferred RF of the neurons allowed us to also estimate a pseudo-population of neurons in the ipsilateral SC. To complete the

representation of activity across the SC topography, we also recorded from and included in our analyses neurons in the rostral portion of SC, which are active during fixation[13], burst during microsaccades[15], and are suppressed during large saccades[13].

Temporal stability analysis

To assess temporal stability, we first normalized the population trajectory $\mathbf{R}(t)$ by its Euclidean norm ($\|\mathbf{R}(t)\|$), equivalently its magnitude, at each time point to yield $\widehat{\mathbf{R}}(t)$:

$$\widehat{\mathbf{R}}(t) = \frac{\mathbf{R}(t)}{\|\mathbf{R}(t)\|}$$

The normalized trajectory can be visualized as a unity length population vector that points in an n-dimensional direction at each instant in time. That is, while $\mathbf{R}(t)$ is free to traverse the n-dimensional activity space, $\widehat{\mathbf{R}}(t)$ is constrained to the surface of an n-dimensional hypersphere (Figure 1D). Temporal stability or consistency of the evolving population was then quantified by the dot product of two time-shifted unity length vectors:

$$S(t) = \widehat{\mathbf{R}}(t-\tau) \cdot \widehat{\mathbf{R}}(t+\tau)$$

The stability metric ($S(t)$) tracks the running similarity of the normalized trajectory separated in time by 2τ ms. Crucially, the normalization constrains $S(t)$ between 0 and 1. Thus, if $S(t) \to 1$, the population activity is considered stable since the relative contribution of each neuron is consistent. If $S(t) \ll 1$, the population activity is deemed unstable because the relative contribution is variable. If the neural trajectory is not normalized, the dot product quantifies similarity across the vectors' magnitude and direction. It roughly mimics the quadratic of the firing rate (exactly so for $\tau = 0$). When the vector direction remains constant, the dot product yields no additional information than that already present in the firing rate. In contrast, the normalization scales the neural trajectory so it always has unity magnitude. It neither alters the relative contributions of the neurons nor compromises the vector direction. The dot product therefore performs an unbiased evaluation of stability based only on vector direction and is an estimate of the fidelity of the population code modulo a multiplicative gain factor. Intuitively, the stability measure is analogous to the correlation between the neurons' activities at two different time points. We chose the dot product, however, because of its interpretability as a measure of pattern similarity in n-dimensional activity space.

We assessed the significance of the stability profiles in two ways. First, we compared the average stability profiles across sessions during the visual response to the premotor epoch (indicated by the s.e.m. bounds in Figure 1E). Next, for each trial, we randomly shuffled the activities of various neurons (channels for laminar recordings) at each time point. This shuffle retains the average firing rate at each instant but removes any temporal correlation in the firing rate across neurons. We performed multiple such shuffles and re-computed temporal stability for each instance, followed by across-trial averaging, baseline correction, and across-session averaging to obtain the distribution of stability profiles expected to occur by chance if the neurons' activities were uncoordinated (gray trace, EDF 1C).

To mitigate the effect of potentially variable visual response onset latencies (on different trials) on the trial- and session-averaged mean temporal stability profiles in the visual epoch, we also performed the

688 averaging after realigning the data to the peak visual response on individual trials (Figure 1F). In order to
689 do this, the time of peak visual burst was identified from the population mean (average across channels)
690 firing rate on individual trials and was used to align the traces across trials before computing the trial
691 average.

692 For the mean-matched control in Figure 1E&F, we performed the temporal stability analysis only on the
693 subset of trials in which the peak of the population visual response was greater than or equal to the
694 population activity at the time of saccade initiation, after accounting for an efferent delay. We chose 15
695 ms before saccade onset as the latest time at which SC activity could influence saccade initiation,
696 because it is consistent with a range of putative efferent delays reported using various approaches[4,5,42].

697

698 *Patterned microstimulation*

699 The patterned microstimulation experiments allowed us to causally evaluate and compare the temporal
700 stability model to other models of movement initiation. Given the ability to control the delivery of
701 individual stimulation pulses to each contact on the laminar electrode, we used stimulation patterns
702 with specific spatiotemporal features to evaluate their relative efficacy in evoking movements. We
703 designed these patterns as follows. All individual stimulation pulses were biphasic with a leading
704 cathodic phase, pulse widths of 250 µs, and an interphase interval of 100 µs. Our first step was to
705 determine the appropriate range of current intensities and frequencies for each experiment. During
706 each session (i.e., for each electrode penetration into SC), we first stimulated from individual contacts
707 with standard suprathreshold parameters (40 µA, 400 Hz) to identify the set of contacts that elicited a
708 movement. We restricted the experiment to these contacts for the rest of that session. We then
709 stimulated simultaneously across all these contacts at the same current and frequency, starting from 4
710 µA, 100 Hz, stepping up 1 µA and 50 Hz, to determine the threshold for evoking a movement with multi-
711 channel stimulation. The threshold, defined as the current/frequency combination for which stimulation
712 generated movements on approximately half the trials, ranged from 5-9 µA and 150-200 Hz across
713 sessions. For each session, we used the metrics of the saccades evoked at these threshold parameters
714 from constant frequency stimulation in lieu of "fixed vector saccades" as a control against which to
715 compare the patterned microstimulation results (EDF 5).

716 Once we identified the threshold parameters, we designed stimulation patterns at that current intensity
717 but with time-varying frequencies intended to simulate a burst of activity. Each stable pattern was
718 created with a linearly decreasing inter-pulse interval (IPI) sequence for one contact (to simulate a burst)
719 and scaling the IPIs for other contacts by a uniformly spaced factor. The 150 ms duration burst was
720 composed of a flat baseline (40 ms), a rising phase (80 ms), and a falling phase (30 ms). The peak
721 frequency for individual channels ranged uniformly between 20 Hz and 400 Hz with the mean peak
722 frequency across contacts pegged to the threshold frequency determined from the constant frequency
723 multi-channel stimulation in an earlier step. At these parameters, stimulation at no individual channel
724 evoked movements. Each stable stimulation pattern (e.g., blue pulse pattern in Figure 4A) can be
725 envisioned as a $n \; x \; m$ matrix, when n is the contact number and m is stimulation duration. The matrix
726 contains ones and zeros, identifying the time of pulse delivered to each contact. To create additional
727 realizations of the stable stimulation patterns (e.g., blue pulse patterns in EDF 4B & C), the assignment
728 of specific peak frequencies to individual channels was randomly shuffled across trials.

For each stable stimulation pattern, we created a corresponding unstable pattern/pair by jittering stimulation pulses within a 20 ms window for each channel and randomly shuffling pulses across channels (e.g., red pulse patterns in Figure 1A and EDF 4B & C). This step created instability in the population pulse pattern by destroying the relative scaling of pulse rates across neurons while also ensuring that both total pulse counts per channel and the mean instantaneous pulse rates across channels were preserved between the stable and unstable patterns. Thus, all subsequent analyses were performed on pairs of trials with stable and unstable stimulation patterns matched in these aspects. We also ensured that the inter-pulse interval was never less than 2 ms (i.e., the peak frequency never exceeded 500 ms) for any stimulation train. The stimulation patterns were generated offline and the pattern trains corresponding to each trial were introduced 100 ms into the gap period in a gap task. Stable-unstable pairs were randomly interleaved within a block in which roughly 80% of all trials were stimulation trials.

We used the latency of the first saccade after stimulation onset but before stimulation offset as an indicator of the occurrence of a stimulation-evoked saccade (Figure 3D); such saccades exhibited latencies <150 ms. If the microstimulation was ineffective at evoking a saccade, the first movement was typically directed to the target presented after the gap period ended; in such cases, the saccade was produced >400 ms after stimulation onset. We estimated the relative likelihood of evoking a movement by examining trials in which only one of the stable or unstable pulse patterns in a pair yielded a stimulation-evoked saccade (also see Figure 4E and next section). This was quantified based on the number of trials pairs in which only the stable pattern evoked a movement during the window of stimulation ($= \mathrm{N}_{\mathrm{s+us-}}$, number of points in the blue shaded region in Figure 3D) and the number of trial pairs in which only the unstable pattern evoked a movement ($= \mathrm{N}_{\mathrm{us+s-}}$, number of points in the red shaded region in Figure 3D). The relative likelihood of evoking a movement with the stable pattern was defined as

$$\mathrm{RL_S} = \frac{\mathrm{N}_{s+us-}}{\mathrm{N}_{s+us-}+\mathrm{N}_{us+s-}}.$$

$\mathrm{RL_S}$ ranges from 0 to 1 and is symmetric with a neutral value of 0.5. To estimate the significance of this estimate, we performed permutation tests with respect to a surrogate dataset in which trials were randomly assigned stable/unstable labels. This lowered the expected estimate of relative likelihood to purely chance levels (Figure 3E).

Discriminability of population patterns – neural activity and microstimulation

To determine whether the visual and premotor bursts are discriminable using static, non-temporal features of population activity, we performed factor analysis (FA) on 400 ms snippets of activity from these two epochs (visual epoch: 0-400 ms from target onset, motor epoch: -200:200 ms relative to saccade onset). We used DataHigh[19], a dimensionality reduction and visualization toolbox written in Matlab, to perform FA on the 16- or 24-dimensional laminar recordings and visualize the projections onto the top latent dimensions. The top 3 dimensions accounted for >95% of the variance in neural states for all our datasets (eigenspectrum not shown), and a 3-dimensional projection is shown for an example dataset in Figure 4A, with the visual and premotor snippets coded in different colours. We also performed FA on the microstimulation patterns (the "states" were computed across the entire 150-ms

stimulation duration). Almost all dimensions were needed to account for >95% of the variance in the stimulation patterns (eigenspectrum inset in Figure 4C), owing to the randomized assignment of frequencies across contacts on different trials. Thus, we chose an arbitrary 3-dimensional projection for visualization in Figure 4C; however, note that the qualitative result did not depend on the projection. Since we used this step purely for visualization, we did not perform any subsequent analysis on the estimated FA dimensions.

Next, we used a linear discriminant analysis (Fisher's LDA)[43] to assess a decoder's ability to discern whether its inputs signify a movement command at the population level, based on static population features alone, and how this discriminability relates to temporal stability. We first trained a binary LDA classifier on the visual and premotor snippets described above in the native neural space. LDA finds a projection defined by a hyperplane that maximizes the separation between the two classes:

$$\boldsymbol{w} \propto \Sigma^{-1}(\boldsymbol{\mu_1} - \boldsymbol{\mu_0}) ,$$

where the vector $\boldsymbol{w}$ is normal to the hyperplane defining the class boundary, $\Sigma = \frac{\mathrm{n}_0\Sigma_0 + n_1\Sigma_1}{n_0+n_1}$, the pooled covariance matrix derived from the within-class covariance matrices Σ_0 and Σ_1, n_0 and n_1 are class occupancies, and $\boldsymbol{\mu_0}$ and $\boldsymbol{\mu_1}$ are the class means. We subjected the linear scores

$$\boldsymbol{s_i} = \boldsymbol{w}.\left(\boldsymbol{x} - \frac{\boldsymbol{\mu_i}}{\mathbf{2}}\right)$$

obtained from the LDA projections to a softmax transformation to compute class probabilities:

$$P(\boldsymbol{x} \in i) = \frac{e^{\boldsymbol{s}_i}}{\sum_i e^{\boldsymbol{s}_i}}$$

In all cases, we used k-fold crossvalidation to test LDA performance (on the classification of visual and premotor population states, and on stimulation patterns described below), with the choice of k dictated by the total number of trials in the dataset such that each fold (test set) had at least 10 trials. We quantified classifier performance using the Matthews Correlation Coefficient (MCC)[44], since accuracy can be significantly biased and have high variance for classification problems with unequal number of training or test points[45], which was the case with the stimulation patterns below. Instead, MCC takes the complete confusion matrix into account to quantify binary classifier performance as:

$$\mathrm{MCC} = \frac{\mathrm{N}_{00}\mathrm{N}_{11} - \mathrm{N}_{01}\mathrm{N}_{10}}{\sqrt{(N_{00}+N_{01})(N_{00}+N_{10})(N_{11}+N_{01})(N_{11}+N_{10})}},$$

where N_{ij} represents the number of instances where a trial in class *i* is classified to be in class *j* in the classification confusion matrix. Similar to a standard correlation coefficient, MCC values range from -1 to 1 (-1, 0, and 1 respectively indicate poor, average, and perfect classification).

To evaluate the discriminability of microstimulation patterns, we split each dataset in 3 different ways. The first split was based on the predefined stable-unstable pattern pairs and was used to classify pulse patterns as stable or unstable based on their population states alone (Figure 4D). The next two splits were used to classify stimulation patterns as movement-evoking or non-evoking, in order to identify other features of the stimulation patterns that potentially could have impacted movement occurrence. We split the stable-unstable trial pairs into two subsets – the neutral subset (NS), which was made up of trial pairs where either both or neither pattern in the stable-unstable pair evoked a movement, and the

stability subset (SS), made up of trial pairs where only one of the stable-unstable pair evoked a movement (Figure 4E). The rationale was that a classifier trained on NS must identify features *other* than temporal stability in order to classify a trial as (stimulation-evoked) movement or non-movement because of the matched numbers of stable-unstable trials regardless of movement evoked. On the other hand, a classifier trained on SS may potentially be able to identify temporal stability as a predictor because of the differential relationship between stability and movement occurrence. We further evaluated whether a dynamic feature of population activity such as temporal stability can be utilized for classification based on population states alone, by re-training the classifier with explicit addition of temporal stability as an input dimension and comparing it with the performance of one without this addition.

Correlation analysis

In order to compare against pooled accumulator mechanisms[20], we estimated the correlation structure in our population of neurons and microstimulation patterns. We first computed the spike count correlation across trials for every pair of neurons (or multi-unit clusters, n = 1963 pairs) during the visual and motor epochs as defined above. For the microstimulation patterns, we computed the pulse count correlations for every pair of stimulation channels (n = 899 pairs) during the stimulation window for the stable and unstable patterns separately. Note that in this case, pulse count correlation is equivalent to computing the correlation between accumulation rates (and therefore directly comparable to the accumulator mechanism mentioned above), since the duration of stimulation was fixed and the baseline and peak rates were matched between the stable and unstable conditions.

Spiking neuron model

We developed a biophysically realistic computational model with spiking neurons to discriminate and read out temporal structure in population activity. The core component of the model was a recurrently connected module comprising a decoder element and a disinhibitor element, putatively representing pontine high-frequency burst neurons and OPNs, respectively[25]. The decoder and disinhibitor were both modelled as leaky integrate-and-fire (LIF) spiking neurons:

$$\tau_m \frac{dV_{neu}}{dt} = -(V_{neu}(t) - V_{rest}) + R_m I_{neu}(t)$$

where $V_{neu}(t)$ is the membrane potential of the neuron, V_{rest} is the resting potential, τ_m is the membrane time constant, R_m is the membrane resistance, $I_{neu}(t)$ is the net time-varying input to the neuron, and *neu* represents the decoder (*dec*) or disinhibitor (*dis*). The spiking followed standard threshold crossing rules, i.e., if $V_{neu}(t = t^k) > V_{th}$, where V_{th} represents the action potential threshold, the membrane potential was reset to $V_{neu}(t = t^k) \equiv V_{reset}$, and the neuron's spike times $[t^i]$ were updated with the k^{th} spike time t^k.

The current input to both the decoder and disinhibitor was composed of excitatory and inhibitory components:

$$I_{neu}(t) = G_{neu}^{exc}(t)\left(E_{exc} - V_{neu}(t)\right) + G_{neu}^{inh}(t)\left(E_{inh} - V_{neu}(t)\right) + I_{neu}^{0}$$

where $G_{neu}^{exc}(t)$ and $G_{neu}^{inh}(t)$ are the respective net excitatory and inhibitory post-synaptic conductances, and E_{exc} and E_{inh} are the excitatory and inhibitory reversal potentials representing the contribution of AMPA and GABA channels, respectively. I_{neu}^{0} is a constant current term that was set to zero for the decoder and to a fixed value of 4 for the disinhibitor in order to simulate tonic firing without external input.

The decoder received excitatory inputs from the SC in the form of stimulation pulse trains, and inhibitory input from the disinhibitor. The disinhibitor received both excitatory and inhibitory inputs from the SC and reciprocal inhibition from the decoder[21]. The net time-varying conductances were modelled as follows:

$$G_{dec}^{exc}(t) = \sum_{i=1}^{N} w_i^{exc-dec} g_i(t) \text{ , and } g_i(t) = \sum_{k=1}^{n_i} \epsilon\left(t - t^k\right)\left(1 + g_i\left(t^k\right)\right),$$

$$G_{dec}^{inh}(t) = w^{dis-dec} \sum_{k=1}^{n_{dis}} \epsilon\left(t - t^k\right),$$

$$G_{dis}^{exc}(t) = \sum_{i=1}^{N} w_i^{exc-dis} g_i(t) \text{ , and } g_i(t) = \sum_{k=1}^{n_i} \epsilon\left(t - t^k\right)\left(1 + g_i\left(t^k\right)\right),$$

$$G_{dis}^{inh}(t) = w^{dec-dis} \sum_{k=1}^{n_{dec}} \epsilon\left(t - t^k\right) + \sum_{i=1}^{N} w_i^{inh-dis} g_i(t) \text{ , and } g_i(t) = \sum_{k=1}^{n_i} \epsilon\left(t - t^k\right)\left(1 + g_i\left(t^k\right)\right).$$

In each equation they appear, N is the number of input units (i.e., number of channels with pulse trains), $w_i^{exc/inh-neu}$ is the excitatory/inhibitory weight from the i^{th} input unit to the spiking neuron, n_i is the number of pulses in the i^{th} input unit, $w^{dis-dec}$ is the weight from the disinhibitor to the decoder, $w^{dec-dis}$ is the weight from the decoder to the disinhibitor, and t^k is the time of the k^{th} input pulse or output spike (depending on which summation term it appears in). The g_i term, when it appears in the sum that produces g_i (as in all but the second conductance equation), represents the state-dependent influence of each incoming pulse input. $\epsilon(t)$ is the post-synaptic conductance kernel defined as:

$$\epsilon\left(t - t^k\right) = \begin{cases} -e^{-\frac{t}{\tau_1}} + e^{-\frac{t}{\tau_2}} \; if \; t > t^k \\ 0 \; if \; t < t^k \end{cases},$$

The conductance and input current terms that were dependent solely on the pulse input patterns were pre-computed and fed into a Eulerian solver (time step = 0.1 ms) for the differential equations governing the membrane potential of the spiking decoder and disinhibitor units. Extended Data Table 1 shows the values of the parameters and constants used for the model simulations.

Extended Data Figures

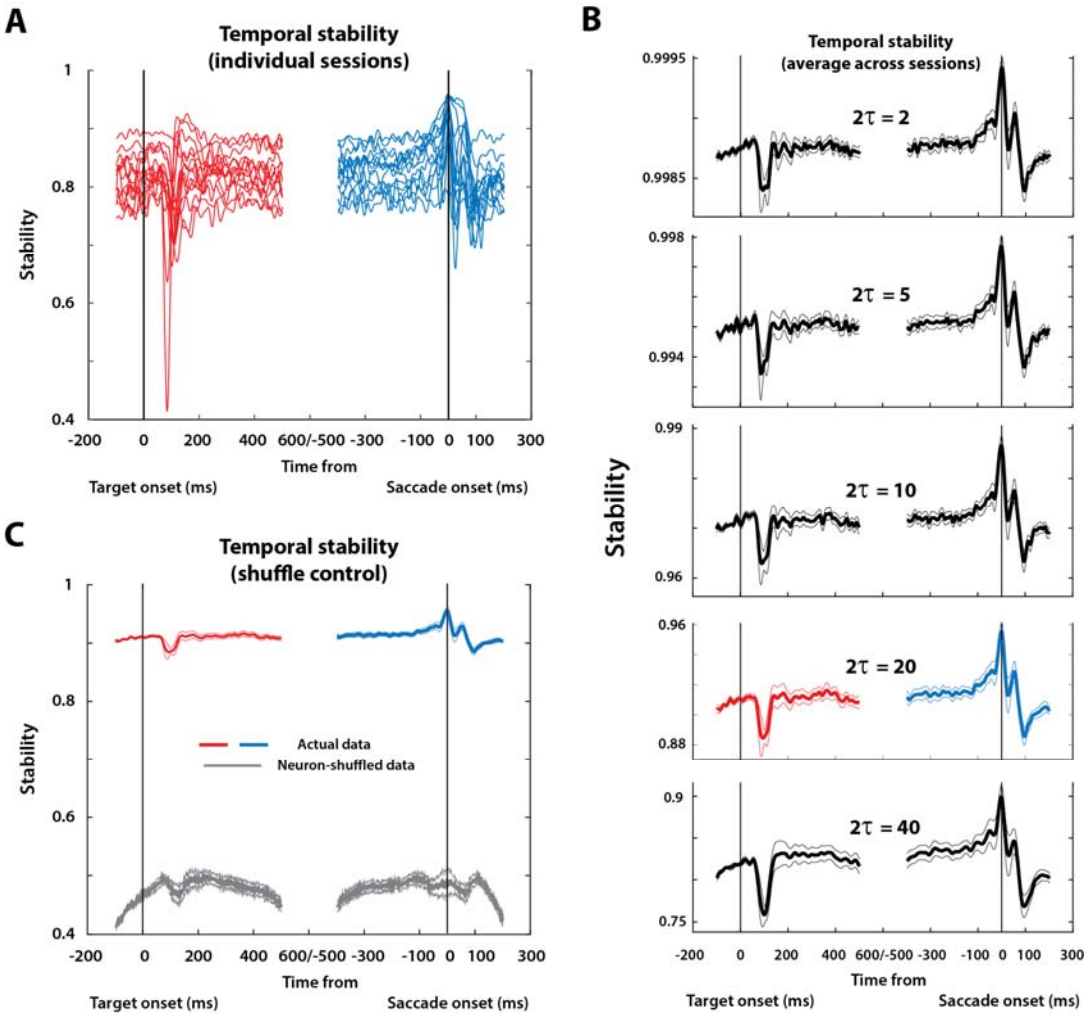

Extended Data Figure 1. Temporal stability profiles within and across sessions collected with laminar probe. A. Temporal stability averaged across trials for individual sessions (thin traces) have different baselines, possibly due to varying noise levels, spike isolation differences, etc. The baselines were shifted to their mean level before averaging across sessions in Figure 1. **B.** In order to ensure that the choice of τ did not have an undue effect on our results, we computed stability for several values of τ. The panels show temporal stability profiles for (from top to bottom), τ = 1, 2.5, 5, 10, and 20, respectively. The coloured traces for τ =10 are the ones shown in the main Figure 1. As evident, the absolute magnitude of stability was inversely related to τ, a consequence of the smooth and continuous nature of the trajectory. In other words, since the state of the neural population evolves smoothly, it must traverse intermediate states in order to move from one state to another, resulting in greater similarity between state vectors close together in time than those further apart. However, the relative shape of stability profile was largely preserved across τ. Thus, the relative instability during the visual epoch and stability during the premotor epoch persists across a broad range of time separations. **C.** Temporal stability averaged across sessions for shuffled data (gray traces, mean +/- s.e.m.). For each trial in each session's dataset, the activities of neurons were shuffled at each time point so that the instantaneous population average was preserved. The coloured traces show the stability profiles for actual data from Figure 1E.

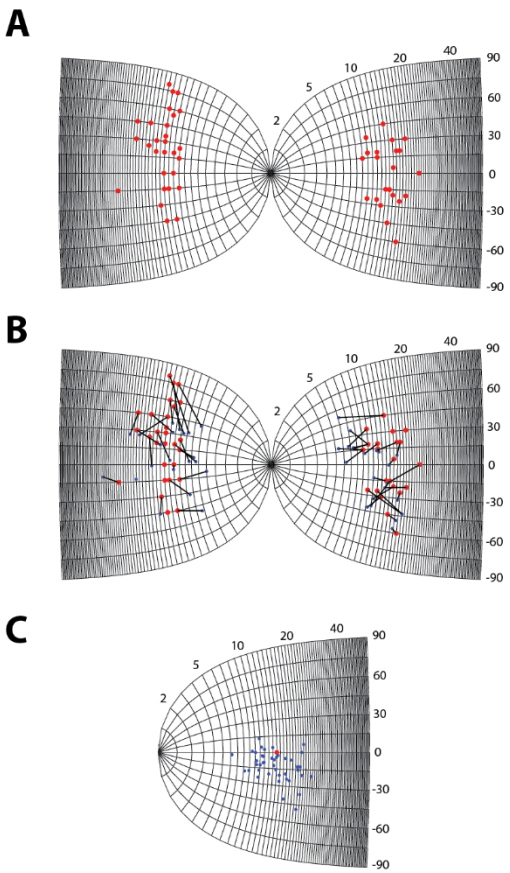

883

884 **Extended Data Figure 2. Inferring population dynamics from single-unit recordings. A.** Point images (red) of
885 target locations on the SC map across all experimental sessions. The locations on the SC were computed using
886 known transformations between visual space and SC tissue coordinates (see Methods). **B.** Same as in A, with the
887 endpoint of the stimulation-evoked saccade vector at the recording site shown in blue. The stimulation vector
888 provides a proxy for the RF center of the recorded neuron. Neuron-target pairs (blue-red) from individual sessions
889 are connected using black lines (unconnected points did not have the corresponding stimulation/target data). **C.**
890 The active pseudo-population on the SC map. The red locations from A and B were referenced to one arbitrarily
891 selected location on the SC map (here, R=15, theta = 0) and the blue locations appropriately translated. Points
892 from both colliculi are shown on a single SC for the sake of clarity.

893

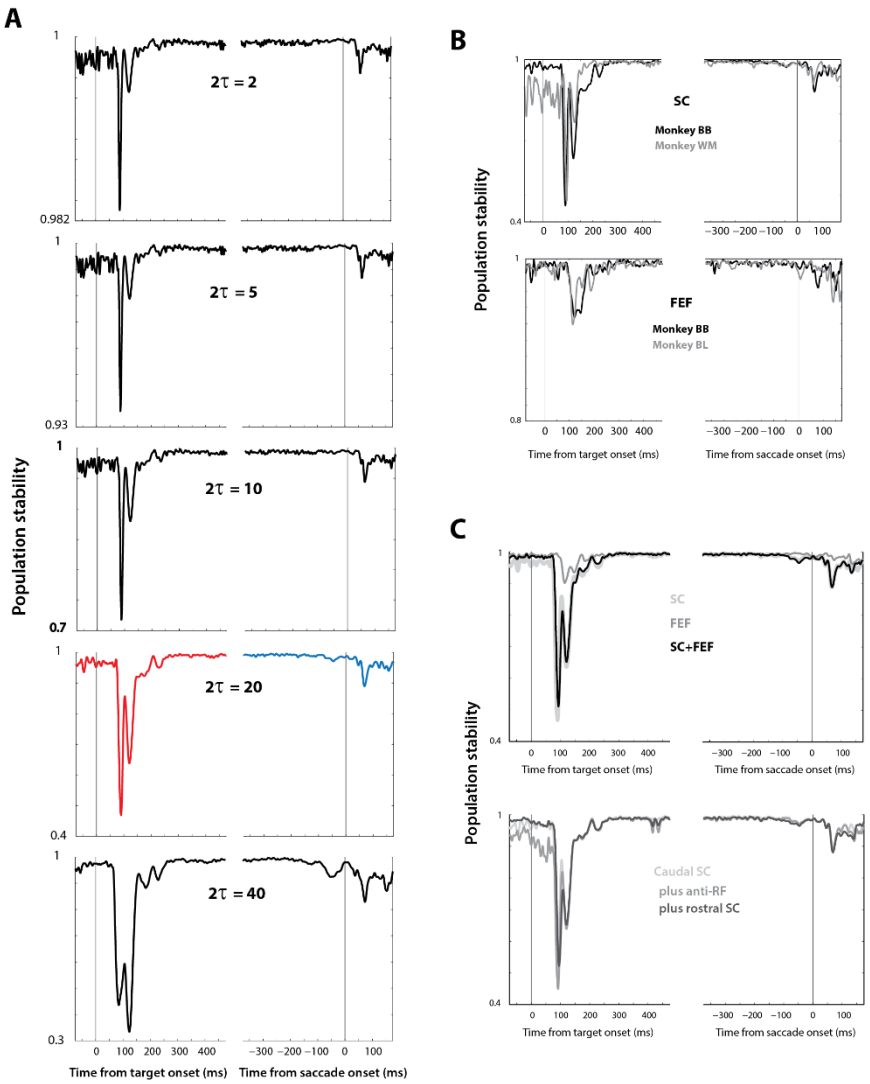

894

895 **Extended Data Figure 3. Temporal stability profiles based on pseudo-population analyses. A.** Effect of τ on
896 temporal stability for the pseudo-population data (similar to EDF 1B). **B.** Temporal stability profiles in SC and FEF
897 for individual subjects are qualitatively similar to the combined result in Figure 2C. **C.** Temporal stability profiles of
898 various combinations of pseudo-populations.

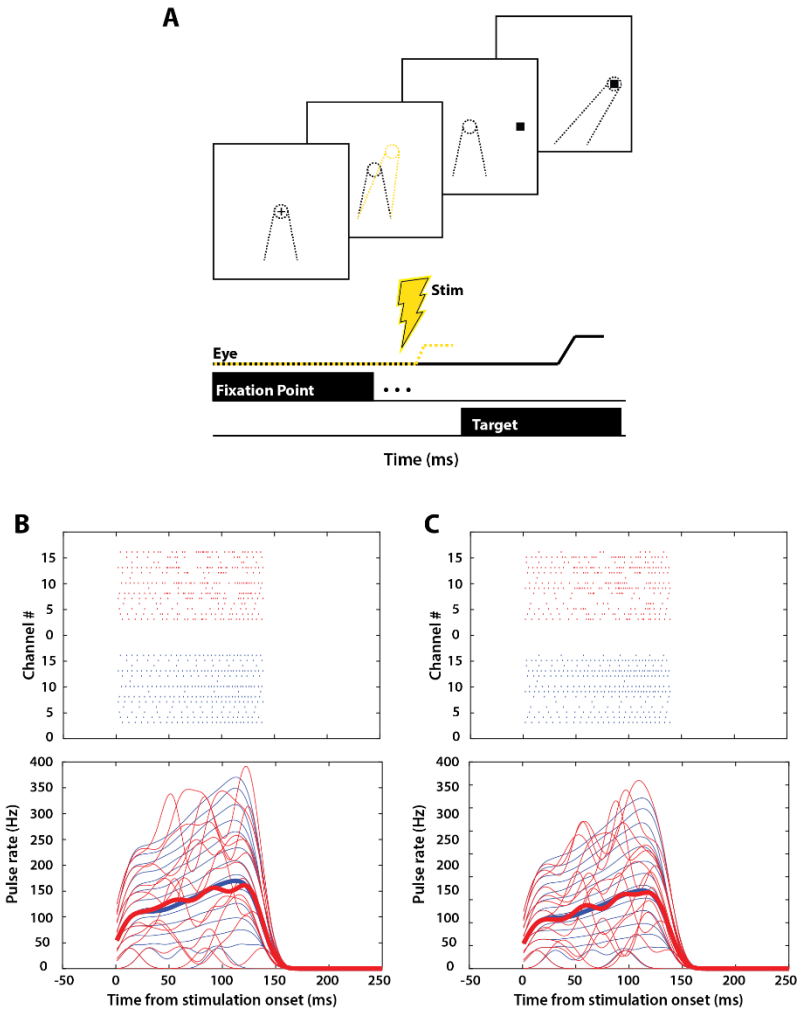

Extended Data Figure 4. Task and pulse pattern configurations for microstimulation experiments. A. Schematic of the gap task used in the patterned microstimulation experiments. **B** and **C.** Two more examples of pulse rasters and pulse rates for individual channels for stable (blue) and unstable (red) pairs of stimulation patterns (see Figure 3A).

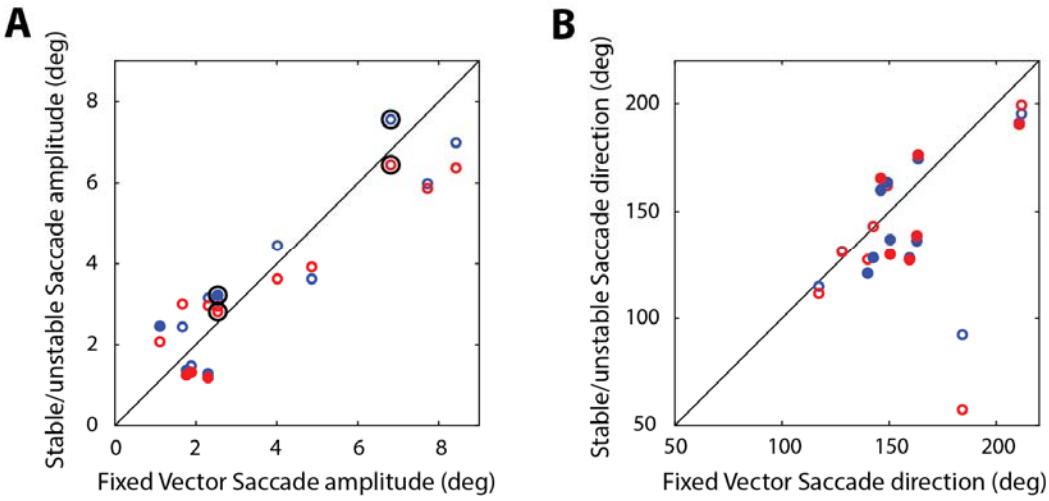

Extended Data Figure 5. Features of saccade vectors evoked by stable and unstable pulse patterns. A. Mean stimulation-evoked saccade amplitude for stable (blue points) and unstable (red points) stimulation patterns plotted against the fixed vector saccade amplitude obtained from near-threshold constant frequency stimulation (see Methods). Each point represents one session for that condition. The diagonal represents the unity relationship. Filled circles denote a significant difference of the stable/unstable saccade amplitude from the fixed vector saccade ($p < 0.01$, Wilcoxon rank-sum test). Points circumscribed by black circles denote sessions where the saccade amplitudes on stable and unstable trials were significantly different from each other ($p < 0.01$, Wilcoxon signed-rank test). **B.** Similar to A (including criteria for statistical significance), but for stimulation-evoked saccade direction.

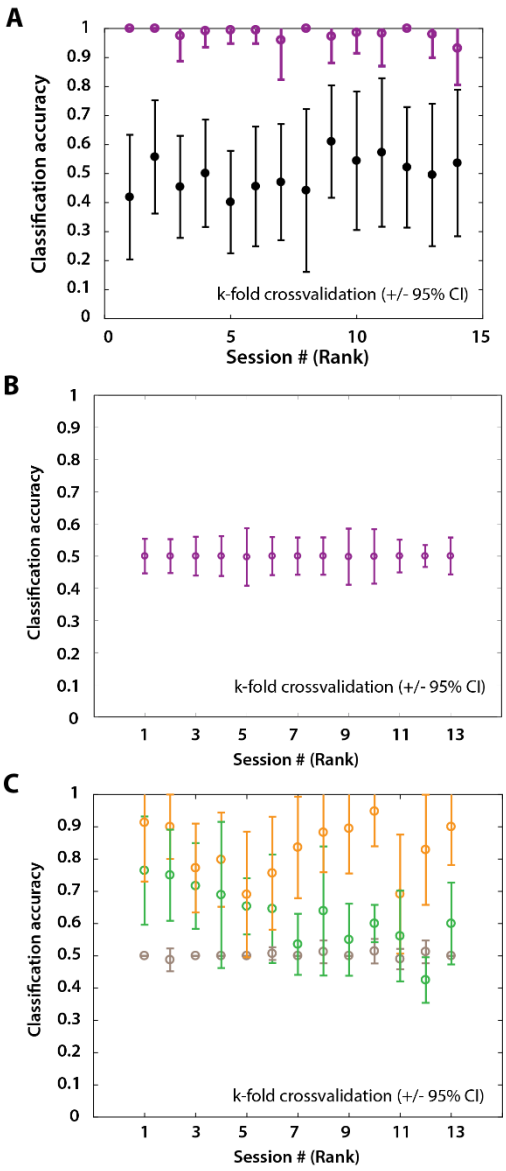

916

917 **Extended Data Figure 6. Linear discriminability of population activity and microstimulation patterns. A-C.** Similar
918 to panels B, D, and F, respectively, from Figure 4, but for uncorrected classification percentage accuracy instead of
919 Matthews correlation coefficient (MCC).

920

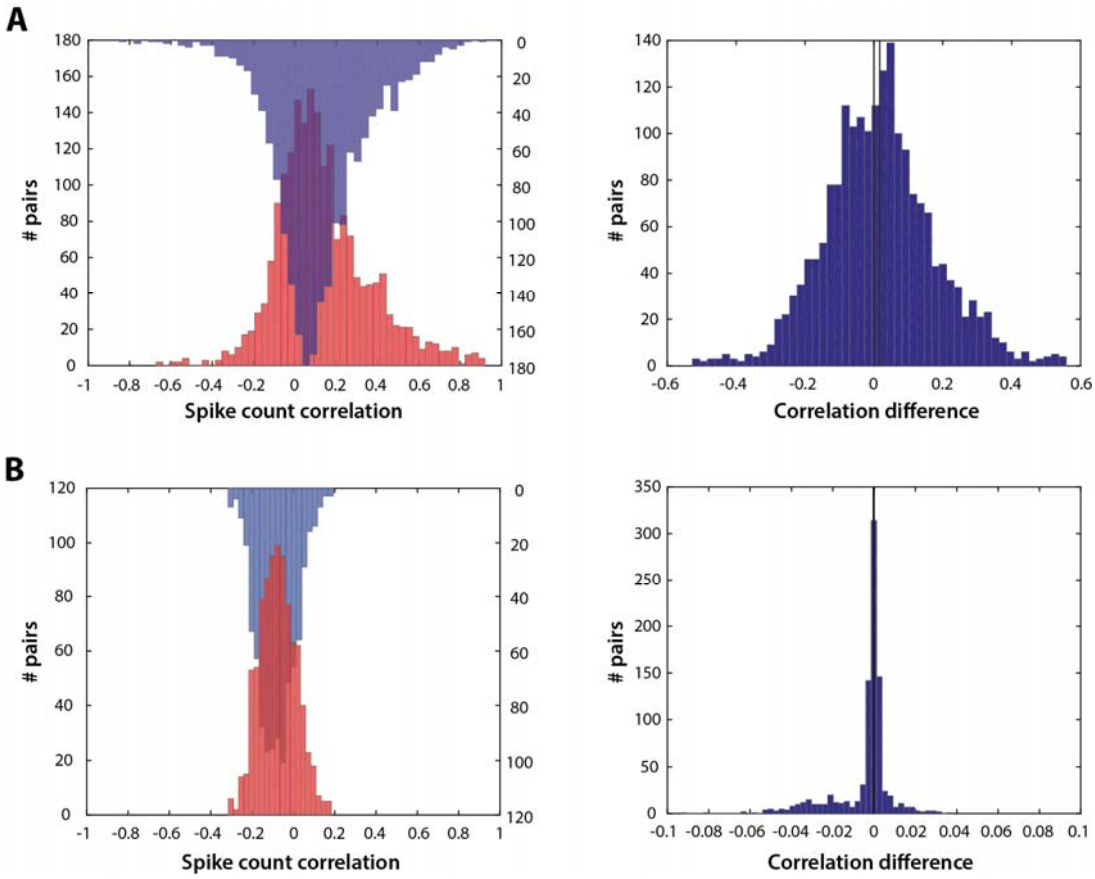

Extended Data Figure 7. Across-trial spike and pulse count correlations. A. Left: Histograms of spike count correlations for all pairs of channels (n = 1963) for the visual and premotor epochs (red and blue histograms, respectively). One of the histograms is shown inverted for the sake of clarity. Right: Distribution of differences between visual and premotor epoch correlations computed for each pair. The two vertical lines represent zero and the mean of the distribution (left and right lines, respectively). **B.** Same as A for pulse counts correlations for all pairs of stimulation channels (n = 899) for the stable (blue) and unstable (red) patterns, with the corresponding difference histogram shown on the right.

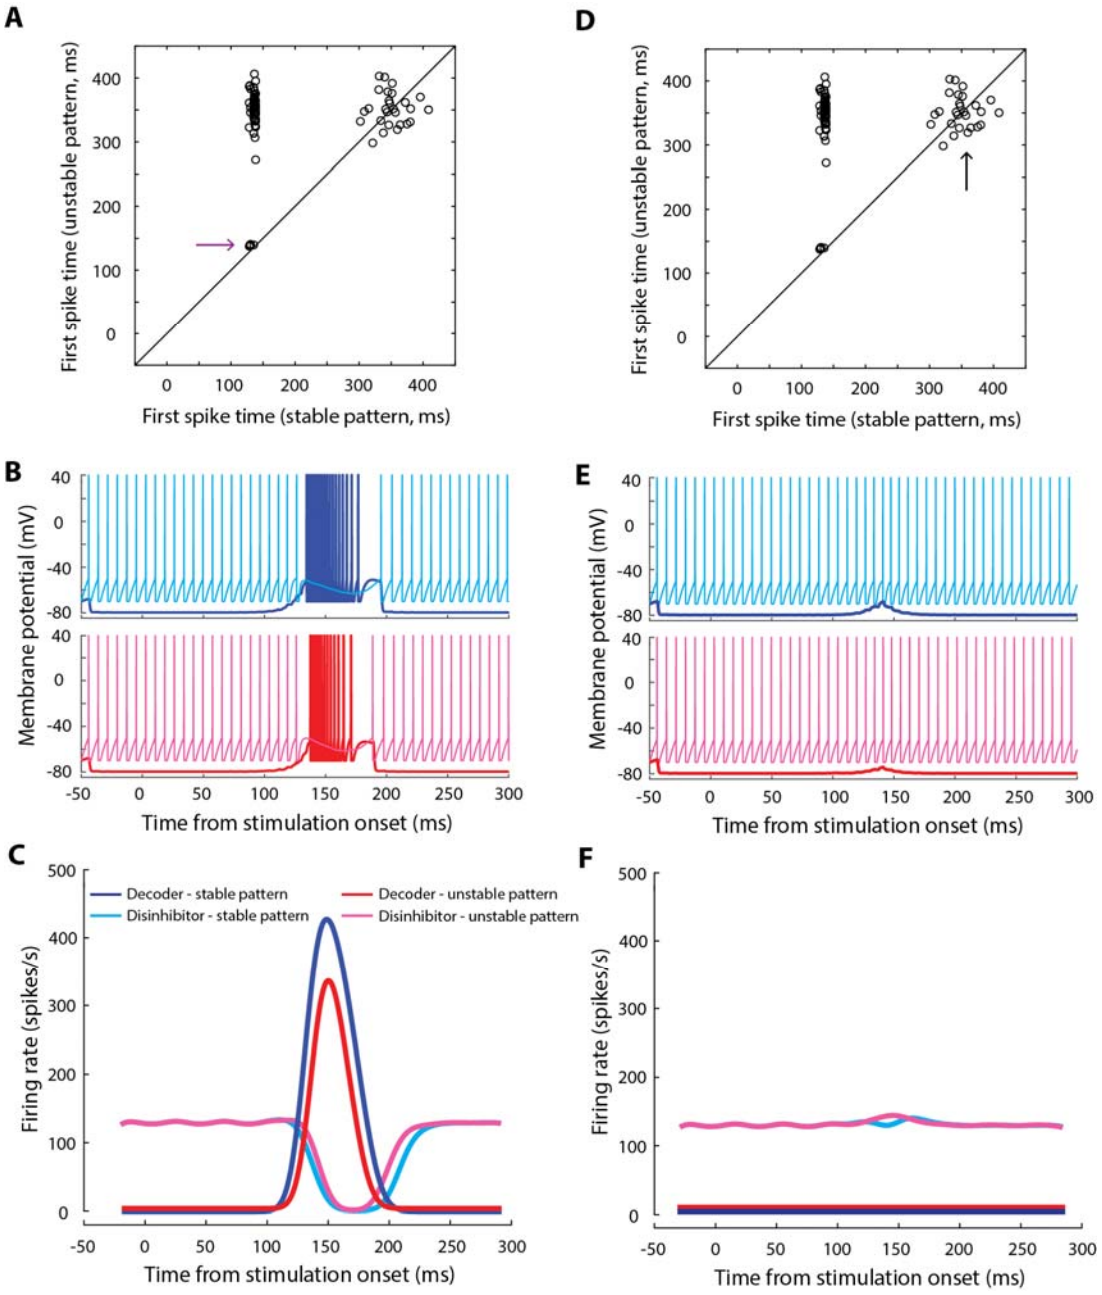

930

931 **Extended Data Figure 8. Performance of spiking neuron model for various simulation outcomes with stable and**
932 **unstable pulse patterns. A.** Scatter plot of first spike latencies in the decoder (putatively representing saccade
933 initiation) for stable versus unstable inputs from matched pairs (same as Figure 5C). The next two panels are based
934 on the trial pairs highlighted by the arrow (in which both stable and unstable patterns produced a saccade). **B.**
935 Simulated membrane potential of the decoder (blue and red traces) and disinhibitor (cyan and magenta traces) for
936 an example matched trial pair from the subset in panel A, for stable (top row) and unstable (bottom row) input
937 patterns. **E.** Average activity of the disinhibitor (cyan and magenta traces, stable and unstable inputs, respectively)
938 and decoder (blue and red traces, stable and unstable inputs, respectively) for all trial pairs in which both stable
939 and unstable input patterns produced a movement. **D.** Same as panel A, except for the focus on trial pairs (arrow).
940 The next two panels are based on the trial pairs in which neither stable nor unstable patterns produced a saccade.
941 **E.** Simulated membrane potential of the decoder and disinhibitor for an example matched trial pair from the

942 subset in panel D. Layout and colour scheme as in panel B. **F.** Average activity of the disinhibitor and decoder for all
943 trial pairs in which neither stable nor unstable input patterns produced a movement. Colour scheme as in panel C.
944 Note the small increase in disinhibitor activity around the time when a movement would have normally occurred,
945 reminiscent of the slight increase in OPN activity during visual input[26].

946

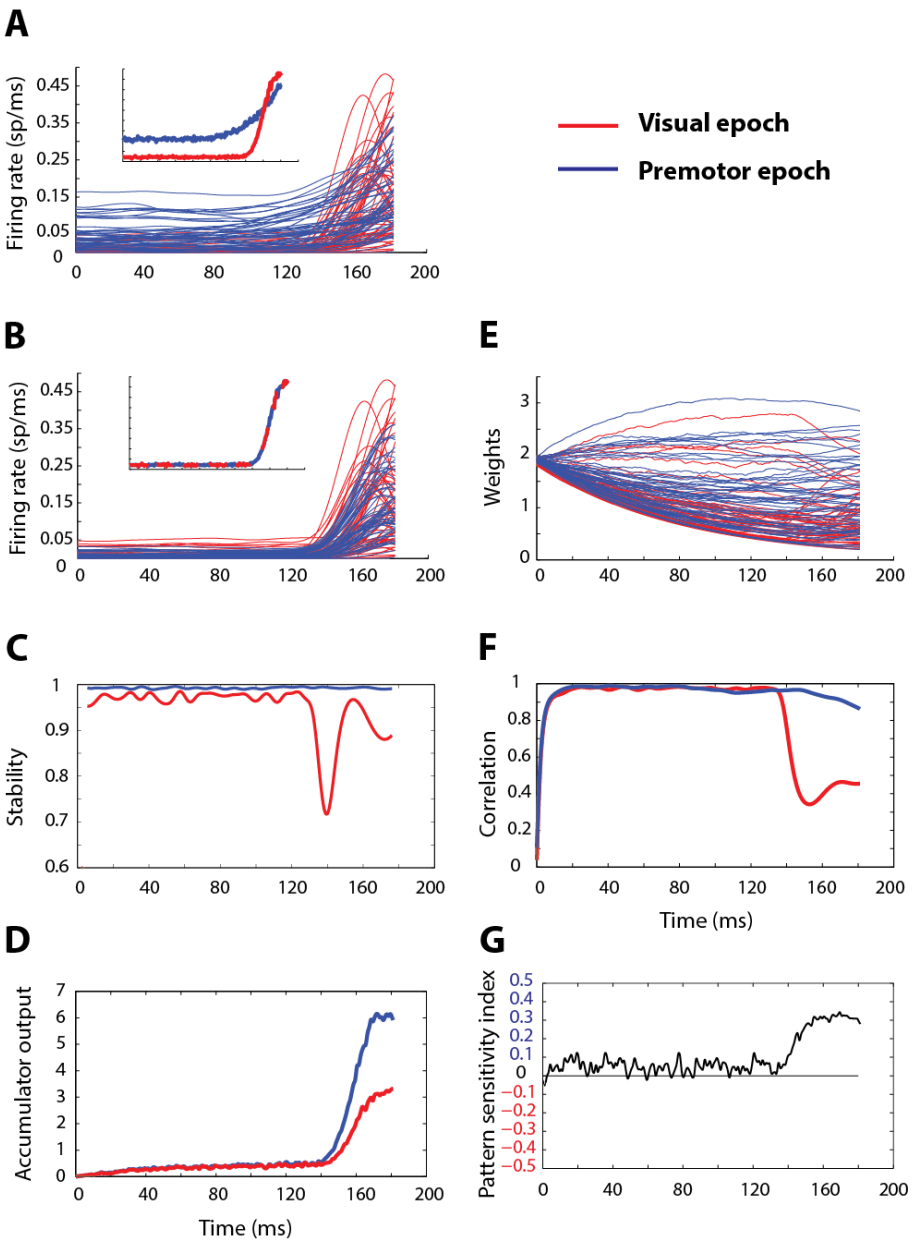

Extended Data Figure 9. Leaky accumulator with facilitation (LAF) model can discriminate population temporal structure. A. Inputs to the model from SC visuomovement neurons. Raw unstable (red) and stable (blue) input profiles (inset – population means). The two sets of inputs are 180 ms snippets taken from the visual and premotor bursts, respectively, in the spike density profiles shown in Figure 2A. **B.** Mean-matched input profiles and population means. **C.** Temporal stability of the two mean-matched populations. **D.** Output of the LAF accumulator in response to the stable and unstable model inputs. **E.** Evolution of synaptic weights for the two conditions. **F.** Correlation between instantaneous weights and input rates for the two conditions. **G.** Pattern sensitivity index of the model's ability to discriminate between the two types of inputs. Values in the top half of the plot indicate higher sensitivity (faster accumulation) to stable population input.

958 **Extended Data Table 1. Values of the constants and parameters used for the spiking neuron model.**

Biophysical constants	**Values**
τ_m	10 ms
R_m	10 Mohms
V_{th}	-50 mV
V_{rest}	-65 mV
V_{reset}	-70 mV
E_{exc}	50 mV
E_{inh}	-80 mV
Model parameters	
τ_1	2 ms
τ_2	10 ms
$\boldsymbol{w}^{exc-dec}$	$5 \times 10^{-6}\ uniform[0,2]$
$w^{dis-dec}$	5
$\boldsymbol{w}^{exc-dis}$	$0.2\ \boldsymbol{w}^{exc-dec}$
$\boldsymbol{w}^{inh-dis}$	$0.8\ \boldsymbol{w}^{exc-dec}$
$w^{dec-dis}$	0.1
Numerical parameters	
dt (time step)	0.05

959

S1. Inferring population dynamics from single-unit recordings

The main text and methods describe how we reconstructed a pseudo-population from unit recordings. Using knowledge of the RF centre of a given session's local population obtained from microstimulation, and known transformation from visual space to SC tissue coordinates, here we reconstruct the pseudo-population on the SC map.

We first transferred target locations (R_T, θ_T) inside the RF for each neuron onto the SC map using the following transformations[41] –

$x_T = \frac{B_x}{2} \ln \frac{(H+A)^2+V^2}{A^2}$, $y_T = B_y \tan^{-1} \frac{V}{H+A}$, where,

$H = R_T \cos\theta_T$, $= R_T \sin\theta_T$, and $A = 3$, $B_x = 1.4$, $B_y = 1.8$.

The transformed locations (x_T, y_T) on the bilateral SC map are shown in EDF 2, top row. In order to move these target locations to one pseudo-population, we need to identify where these points reside in the RF of each neuron or local population. We used the endpoints of the site-specific microstimulation-evoked saccades as a proxy for the respective RF centres. The transformed endpoints, with their locations relative to the corresponding target locations, are shown in EDF 2, middle row. We picked an arbitrary location in the visual field as the RF centre of the pseudo-population, $(R_c, \theta_c) = (15,0)$. We used the fact that the size of the active SC population is invariant regardless of the encoded saccade to preserve relative distances in tissue space between a site's target location and RF centre, and translated the transformed microstimulation endpoints to the common pseudo-population centre, along with the respective target locations. To construct one active population from sites gleaned from both hemifields (and therefore both colliculi), we reflected the coordinates from one colliculus onto the other. The resulting pseudo-population (EDF 2, bottom row) shows a fairly representative sampling of neurons from the pseudo-population, with its extent consistent with a large (25% of SC tissue) active population, albeit one that is biased to one side of the population.

S2. Leaky accumulator with facilitation (LAF) model

We developed an alternative computational model to demonstrate how the temporal structure of population activity could be used by neural circuits to gate decisions such as movement initiation. Similar to the spiking neuron model presented in the main text, a core principle behind this model is the hypothesis that signal integration is stronger when the temporal structure in input activity is stable, which allows the decoder neuron to reach threshold during the high frequency burst that triggers the movement. We constructed an accumulator as an abstraction of a neuron receiving population inputs through its network of dendrites. The total synaptic current $I(t)$ and the firing rate $v(t)$ of the accumulator were defined as

$$\tau_s \dot{I} = -I + \sum_1^n w_i u_i$$

and

$$v = F(I)$$

where u_i and w_i are the instantaneous firing rate and synaptic gain of the i^{th} input neuron, and τ_s is the time constant of the synaptic current. The output firing rate of the single decoder neuron $v(t)$ can be described by a standard monotonic function (e.g., linear or sigmoid) applied to the net current[46]. The family of stochastic, leaky accumulator models that integrate the firing rate of neurons toward a threshold criterion has been commonly used to explain reaction times, decision making, and perception[20,47-49]. We used the following heuristic in order to extend this framework to incorporate temporal structure. In order to assess stability, the decoder neuron must keep track of the short term history of the population activity, use this "memory" to evaluate stability, and respond selectively when the activity pattern is deemed stable over some time scale. We implemented these requirements by introducing short-term plasticity in the form of facilitatory connections[50,51] from the input population onto the output unit (accumulator). The change in connection strength or gain of each neuron on the decoder neuron w_i can be defined by a differential equation that incorporates a Hebbian-like learning rule and leak current:

$$\tau_w \dot{w}_i = -w_i + \frac{u_i v}{wnf_i}$$

The Hebbian-learning component, $u_i v$ describes the change in weight coupled to the firing rates of the i^{th} pre-synaptic neuron and the post-synaptic accumulator neuron. τ_w is the time constant of the weight parameter. Since this version of the model contains only excitatory connections, the weight parameter can exhibit unbounded accumulation, which can be controlled by incorporating normalization. We normalized the Hebbian-learning component $u_i v$ with the contribution of that neuron to the total resource pool. The resource pool available for facilitation at any given time was defined as the sum of the Hebbian-learning component across all input units $\sum_1^n u_i v$. The contribution of a neuron to the output rate $v(t)$ then determines its contribution to the overall pool, giving rise to the weight normalization factor:

$$wnf_i = \sum_1^n u_i w_i u_i$$

For this model simulation, we used visuomovement neuron activity drawn from the SC pseudo-population as inputs (EDF 8A). We created two sets of input snippets from the visual and premotor bursts, each 180 ms in length. For the visual burst, we used the activity upto the peak in the population response. For the premotor burst, we used activity upto 35 ms before saccade onset (which was the time when the response magnitude reached the same level as the peak of the visual burst). In order to control for the effect of average firing rate on the accumulation, we mean-matched the input profiles as follows. We divided the activations in the premotor input set at each instant by the mean activation of the visual inputs at the corresponding instant. That is,

$$u_i^{pre}(t)_{mm} = \frac{u_i^{pre}(t)}{\frac{1}{n}\sum_1^n u_i^{vis}(t)}$$

where $u_i^{vis}(t)$ and $u_i^{pre}(t)$ are the activity of neuron *i* at time *t* in the visual and premotor input sets, respectively, and $u_i^{pre}(t)_{mm}$ is the premotor input instantaneously rescaled to match the mean of the visual input (EDF 8B). This ensured that we isolated the effect of temporal structure (EDF 8C), independent of firing rate, on the model's response.

The accumulator increased its activity at a faster rate in response to the stable pattern compared to the unstable pattern (EDF 8D), suggesting that the network was able to discriminate between two types of population patterns even though the net input drive was identical. How critical is facilitation to this function? Like the inputs themselves, the weights also showed greater fluctuation during the visual burst (EDF 8E). Consistent with our heuristic, the weights tracked the input rates when the input was stable, but this correlation dropped away when the input was unstable (EDF 8F). Note that the shape of the correlation profile is not unlike the stability profile in EDF 8C. Therefore, facilitation allows the weights to retain the memory of a pattern over the time scale of tens of milliseconds, but the memory can be scrambled by a fluctuating input pattern.

We also quantified the ability of the synaptic weights to track the inputs by computing the Pearson's correlation between the weight and input vectors at each time point (EDF 8G). To quantify the accumulator's ability to discriminate temporal pattern in population input, we computed a pattern sensitivity index (*psi*) as

$$psi(t) = \frac{v^{pre}(t) - v^{vis}(t)}{v^{pre}(t) + v^{vis}(t)}$$

where $v(t)$ is the output of the accumulation for the corresponding inputs.

The results here suggest that the accumulator is driven more strongly by the stable premotor burst, even when other features of population activity remain the same, providing yet another biophysical mechanism by which temporal structure could be read out.

References

1 Wurtz, R. H., Sommer, M. A., Paré, M. & Ferraina, S. Signal transformations from cerebral cortex to superior colliculus for the generation of saccades. *Vision Res* **41**, 3399-3412 (2001).

2 Buneo, C. A., Jarvis, M. R., Batista, A. P. & Andersen, R. A. Direct visuomotor transformations for reaching. *Nature* **416**, 632-636 (2002).

3 Gallese, V., Fadiga, L., Fogassi, L. & Rizzolatti, G. Action recognition in the premotor cortex. *Brain : a journal of neurology* **119 (Pt 2)**, 593-609 (1996).

4 Miyashita, N. & Hikosaka, O. Minimal synaptic delay in the saccadic output pathway of the superior colliculus studied in awake monkey. *Experimental brain research* **112**, 187-196 (1996).

5 Hanes, D. P. & Schall, J. D. Neural control of voluntary movement initiation. *Science* **274**, 427-430 (1996).

6 Jantz, J. J., Watanabe, M., Everling, S. & Munoz, D. P. Threshold mechanism for saccade initiation in the frontal eye field and superior colliculus. *J Neurophysiol* **109**, 2767-2780 (2013).

7 Rodgers, C. K., Munoz, D. P., Scott, S. H. & Pare, M. Discharge properties of monkey tectoreticular neurons. *J Neurophysiol* **95**, 3502-3511 (2006).

8 Segraves, M. A. Activity of monkey frontal eye field neurons projecting to oculomotor regions of the pons. *J Neurophysiol* **68**, 1967-1985 (1992).

9 Branco, T., Clark, B. A. & Hausser, M. Dendritic discrimination of temporal input sequences in cortical neurons. *Science* **329**, 1671-1675 (2010).

10 Carter, A. G., Soler-Llavina, G. J. & Sabatini, B. L. Timing and location of synaptic inputs determine modes of subthreshold integration in striatal medium spiny neurons. *J Neurosci* **27**, 8967-8977 (2007).

11 Harvey, M. A., Saal, H. P., Dammann, J. F., 3rd & Bensmaia, S. J. Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. *PLoS biology* **11**, e1001558 (2013).

12 Bruce, C. J. & Goldberg, M. E. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* **53**, 603-635 (1985).

13 Munoz, D. P. & Wurtz, R. H. Fixation cells in monkey superior colliculus. I. Characteristics of cell discharge. *J Neurophysiol* **70**, 559-575 (1993).

14 Moschovakis, A. K. *et al.* An anatomical substrate for the spatiotemporal transformation. *J Neurosci* **18**, 10219-10229 (1998).

15 Hafed, Z. M., Goffart, L. & Krauzlis, R. J. A neural mechanism for microsaccade generation in the primate superior colliculus. *Science* **323**, 940-943 (2009).

16 Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51-56 (2012).

17 Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nature neuroscience* **17**, 440-448 (2014).

18 Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun* **7**, 13239, doi:10.1038/ncomms13239 (2016).

19 Cowley, B. R. *et al.* DataHigh: graphical user interface for visualizing and interacting with high-dimensional neural activity. *Journal of neural engineering* **10**, 066012, doi:10.1088/1741-2560/10/6/066012 (2013).

20 Zandbelt, B., Purcell, B. A., Palmeri, T. J., Logan, G. D. & Schall, J. D. Response times from ensembles of accumulators. *Proceedings of the National Academy of Sciences of the United States of America* **111**, 2848-2853, doi:10.1073/pnas.1310577111 (2014).

21 Takahashi, M., Sugiuchi, Y., Izawa, Y. & Shinoda, Y. Commissural excitation and inhibition by the superior colliculus in tectoreticular neurons projecting to omnipause neuron and inhibitory burst neuron regions. *J Neurophysiol* **94**, 1707-1726 (2005).

22 Gandhi, N. J. & Keller, E. L. Activity of the brain stem omnipause neurons during saccades perturbed by stimulation of the primate superior colliculus. *J Neurophysiol* **82**, 3254-3267 (1999).

23 Deisz, R. A., Fortin, G. & Zieglgansberger, W. Voltage dependence of excitatory postsynaptic potentials of rat neocortical neurons. *J Neurophysiol* **65**, 371-382, doi:10.1152/jn.1991.65.2.371 (1991).

24 Stuart, G. & Sakmann, B. Amplification of EPSPs by axosomatic sodium channels in neocortical pyramidal neurons. *Neuron* **15**, 1065-1076 (1995).

25 Keller, E. L. Participation of medial pontine reticular formation in eye movement generation in monkey. *J Neurophysiol* **37**, 316-332 (1974).

26 Everling, S., Paré, M., Dorris, M. C. & Munoz, D. P. Comparison of the discharge characteristics of brain stem omnipause neurons and superior colliculus fixation neurons in monkey: implications for control of fixation and saccade behavior. *J Neurophysiol* **79**, 511-528 (1998).

27 Hauser, C. K., Zhu, D., Stanford, T. R. & Salinas, E. Motor selection dynamics in FEF explain the reaction time variance of saccades to single targets. *Elife* **7**, doi:10.7554/eLife.33456 (2018).

28 Jagadisan, U. K. & Gandhi, N. J. Disruption of Fixation Reveals Latent Sensorimotor Processes in the Superior Colliculus. *J Neurosci* **36**, 6129-6140, doi:10.1523/JNEUROSCI.3685-15.2016 (2016).

29 Corneil, B. D., Olivier, E. & Munoz, D. P. Visual responses on neck muscles reveal selective gating that prevents express saccades. *Neuron* **42**, 831-841 (2004).

30 Gu, C., Wood, D. K., Gribble, P. L. & Corneil, B. D. A Trial-by-Trial Window into Sensorimotor Transformations in the Human Motor Periphery. *J Neurosci* **36**, 8273-8282, doi:10.1523/JNEUROSCI.0899-16.2016 (2016).

31 Rizzolatti, G., Riggio, L., Dascola, I. & Umilta, C. Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* **25**, 31-40 (1987).

32 Vigneswaran, G., Philipp, R., Lemon, R. N. & Kraskov, A. M1 corticospinal mirror neurons and their role in movement suppression during action observation. *Curr Biol* **23**, 236-243 (2013).

33 Tan, A. Y., Chen, Y., Scholl, B., Seidemann, E. & Priebe, N. J. Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature* **509**, 226-229, doi:10.1038/nature13159 (2014).

34 Hauptmann, C. & Tass, P. A. Restoration of segregated, physiological neuronal connectivity by desynchronizing stimulation. *Journal of neural engineering* **7**, 056008, doi:10.1088/1741-2560/7/5/056008 (2010).

35 Popovych, O. V. & Tass, P. A. Control of abnormal synchronization in neurological disorders. *Front Neurol* **5**, 268, doi:10.3389/fneur.2014.00268 (2014).

36 Softky, W. Sub-millisecond coincidence detection in active dendritic trees. *Neuroscience* **58**, 13-41 (1994).

37 Bryant, C. L. & Gandhi, N. J. Real-time data acquisition and control system for the measurement of motor and neural data. *J Neurosci Methods* **142**, 193-200 (2005).

38 Sparks, D. L., Holland, R. & Guthrie, B. L. Size and distribution of movement fields in the monkey superior colliculus. *Brain Res* **113**, 21-34 (1976).

39 Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78-84 (2013).

40 Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585-590 (2013).

41 Ottes, F. P., Van Gisbergen, J. A. & Eggermont, J. J. Visuomotor fields of the superior colliculus: a quantitative model. *Vision Res* **26**, 857-873 (1986).

42 Jagadisan, U. K. & Gandhi, N. J. Removal of inhibition uncovers latent movement potential during preparation. *Elife* **6**, doi:10.7554/eLife.29648 (2017).

43 Fisher, R. A. The use of multiple measurements in taxonomic problems. *Ann Eugenic* **7**, 179-188, doi:DOI 10.1111/j.1469-1809.1936.tb02137.x (1936).

44 Matthews, B. W. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochim Biophys Acta* **405**, 442-451 (1975).

45 Boughorbel, S., Jarray, F. & El-Anbari, M. Optimal classifier for imbalanced data using Matthews Correlation Coefficient metric. *PLoS One* **12**, e0177678, doi:10.1371/journal.pone.0177678 (2017).

46 Rauch, A., La Camera, G., Luscher, H. R., Senn, W. & Fusi, S. Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *J Neurophysiol* **90**, 1598-1612 (2003).

47 Usher, M. & McClelland, J. L. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* **108**, 550-592 (2001).

48 Purcell, B. A. *et al.* Neurally constrained modeling of perceptual decision making. *Psychol Rev* **117**, 1113-1143, doi:10.1037/a0020311 (2010).

49 Wong, K. F., Huk, A. C., Shadlen, M. N. & Wang, X. J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front Comput Neurosci* **1**, 6, doi:10.3389/neuro.10.006.2007 (2007).

50 Tsujimoto, T., Jeromin, A., Saitoh, N., Roder, J. C. & Takahashi, T. Neuronal calcium sensor 1 and activity-dependent facilitation of P/Q-type calcium currents at presynaptic nerve terminals. *Science* **295**, 2276-2279 (2002).

51 Borst, J. G. & Sakmann, B. Facilitation of presynaptic calcium currents in the rat brainstem. *The Journal of physiology* **513 (Pt 1)**, 149-155 (1998).