



NEUROSCIENCE

Unified control of temporal and spatial scales of sensorimotor behavior through neuromodulation of short-term synaptic plasticity

Shanglin Zhou^{1,2,3,4,*} and Dean V. Buonomano^{5,6}

Neuromodulators have been shown to alter the temporal profile of short-term synaptic plasticity (STP); however, the computational function of this neuromodulation remains unexplored. Here, we propose that the neuromodulation of STP provides a general mechanism to scale neural dynamics and motor outputs in time and space. We trained recurrent neural networks that incorporated STP to produce complex motor trajectories—handwritten digits—with different temporal (speed) and spatial (size) scales. Neuromodulation of STP produced temporal and spatial scaling of the learned dynamics and enhanced temporal or spatial generalization compared to standard training of the synaptic weights in the absence of STP. The model also accounted for the results of two experimental studies involving flexible sensorimotor timing. Neuromodulation of STP provides a unified and biologically plausible mechanism to control the temporal and spatial scales of neural dynamics and sensorimotor behaviors.

INTRODUCTION

A universal feature of motor behavior is the ability to flexibly adjust the temporal and spatial scales of motor outputs. In the temporal domain, it is possible to produce very similar motor output patterns at different speeds or overall durations (1–4). For example, people can flexibly control the tempo of a musical piece or the duration it takes to sign their names by altering their writing speed. Analogously, in the spatial domain, we can also flexibly change the size of one's handwriting depending on the writing surface area available (5, 6). Similarly, in the sensory timing domain, the encoding of temporal intervals can be flexibly modulated by a range of factors, including dopamine levels (7–9).

It is increasingly clear that motor control and its spatial and temporal flexibility are, in part, governed by the neural dynamics of recurrent neural networks (RNNs) (10–18), suggesting that the neural dynamics of RNNs themselves may undergo transformations that underlie both temporal and spatial scaling. However, the neural circuit mechanisms underlying flexible temporal and spatial transformations remain largely unknown. Although, some neurocomputational models have demonstrated that it is possible to temporally scale RNN dynamics—that is, speed up and slow down the speed at which neural dynamics unfolds—by providing a “speed” input (1, 2, 11, 13, 19, 20) or adjusting the neural input-output gains (11, 21).

Here, we propose a unified and biologically inspired mechanism based on the neuromodulation of STP, to flexibly govern both the temporal and spatial scales of RNN dynamics and sensorimotor behaviors. STP refers to a universal form of use-dependent synaptic plasticity that operates on the subsecond timescale (22–24). Despite its presence at almost all synapses in the brain, the computational functions of STP remain poorly understood. One experimentally characterized feature of STP is that it can be flexibly modulated by neuromodulators such as dopamine (25–33) as well as γ -aminobutyric

acid type b (GABA_b) receptors (34–37). Specifically, neuromodulators can alter the temporal profile of STP by governing release probability: Enhancing initial release can more rapidly exhaust neurotransmitter vesicles from the readily releasable pool and favor short-term depression; in contrast, decreasing release probability can decrease short-term depression and favor short-term facilitation.

Even though STP is universally present at cortical synapses, most neural network models do not incorporate STP [for some exceptions, see (38–41)]. In addition, to the best of our knowledge, no previous neural network models have examined the computational role of the neuromodulation of STP. Here, we demonstrate that the incorporation of STP, and its neuromodulation, into RNN models provides a powerful and flexible mechanism to temporally and spatially modulate RNN dynamics and thus sensorimotor control. We show that neuromodulation of STP accounts for experimental results on scaling tasks (1, 9) and establish that while conventional RNNs can learn to temporally and spatially scale their dynamics, the incorporation of STP significantly enhances the ability of networks to generalize across temporal and spatial scales. Our results provide a hypothesis as to why synapses may exhibit STP and its neuromodulation, and provide a computational mechanism for unified spatial and temporal control of sensorimotor behavior.

RESULTS

Firing-rate-based RNN models have successfully been used to capture neural dynamics of cortical circuits and account for how biological neural networks can perform a range of complex cognitive tasks (42–46). With a few exceptions (38, 47), these RNN models generally do not incorporate STP. Here, we incorporate STP in all the synapses of an RNN using the standard Tsodyks-Markram model composed of two variables—the depression variable x and facilitation variable u . In this model, three parameters define the evolution of these two variables (48): U , which can be interpreted as initial release probability or proportion of vesicles released from the readily releasable pool; τ_x , the time constant of recovery from depression; τ_u , the time constant of facilitation (Fig. 1A). The actual synaptic strength is obtained by multiplying the postsynaptic recurrent weight W^{rec} by the presynaptic variables $x \times u$ and the presynaptic neuron firing rate

¹Institute for Translational Brain Research, Fudan University, Shanghai, China. ²State Key Laboratory of Medical Neurobiology, Fudan University, Shanghai, China. ³MOE Frontiers Center for Brain Science, Fudan University, Shanghai, China. ⁴Zhongshan Hospital, Fudan University, Shanghai, China. ⁵Department of Neurobiology, University of California, Los Angeles, Los Angeles, CA, USA. ⁶Department of Psychology, University of California, Los Angeles, Los Angeles, CA, USA.

*Corresponding author. Email: zhoushanglin@fudan.edu.cn

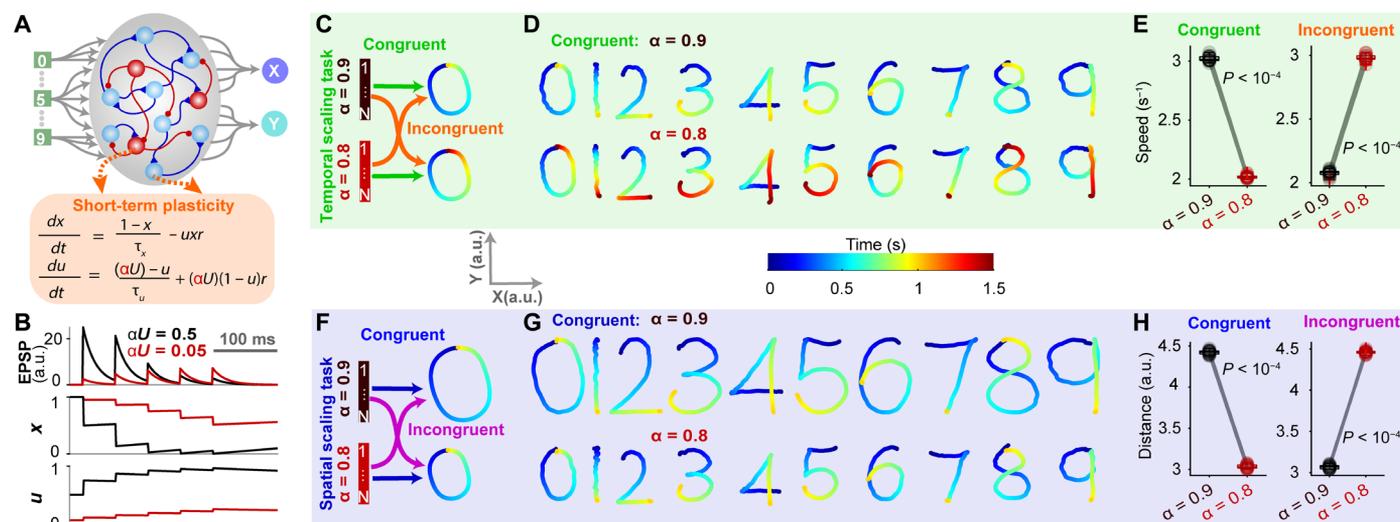


Fig. 1. Temporal and spatial control of motor trajectories through neuromodulation of STP. (A) Schematic of the RNN. Transient activation of either of 10 inputs triggers the production of a digit. STP was implemented through the synaptic depression variable (x) and facilitation variable (u). For each trial, the constant U was scaled by α to signal temporal or spatial scale. X, Y are two output units corresponding to digit coordinates. (B) Example of neuromodulation of STP on excitatory postsynaptic potential (EPSP). In the extreme, different αU can result in short-term depression (black) or facilitation (red) as shown by dynamics of EPSP (top), x (middle), and u (bottom). τ_u and τ_x are both 1 s. (C) Schematic of the temporal scaling task. For each trial, α was 0.9/0.8, corresponding to a given digit production duration of 1 s/1.5 s in the congruent condition or 1.5 s/1 s in the incongruent condition. The size of the target was the same for both α levels. (D) Example output traces under $\alpha = 0.9$ (top) and $\alpha = 0.8$ (bottom) for the congruent temporal scaling condition. (E) Summary of the output speed averaged across digits for the congruent (left) and incongruent (right) conditions for the temporal scaling task ($n = 20$ RNNs; $P < 0.0001$ two-sided Wilcoxon signed-rank test). (F) Schematic of the spatial scaling task. For each trial, α was 0.9/0.8 corresponding to a spatial scale of $1.5 \times / 1 \times$ in the congruent condition or $1 \times / 1.5 \times$ in the incongruent condition. The duration of the target was always 1 s. (G) Example output traces for the congruent spatial scaling task. (H) Summary of the output distance averaged across digits for the congruent (left) and incongruent (right) conditions for the spatial scaling task ($n = 20$ RNNs; $P < 0.0001$ two-sided Wilcoxon signed-rank test).

r (Eq. 1 in Materials and Methods). We further implemented neuromodulation of STP via a factor α that modulated U (Fig. 1A). This factor α represents the level of a neuromodulator such as dopamine at the beginning of a trial. By modulating U through α , STP can range from short-term depression with high U values (where depression variable x dominates) to short-term facilitation (where the facilitation variable u dominates) with low U values in the extreme case (Fig. 1B).

To address whether neuromodulation of STP can be used to control temporal and spatial scales, we trained RNNs on temporal and spatial scaling tasks, respectively. Each RNN was trained to produce 10 complex motor trajectories—handwritten digits 0 to 9—in response to 1 of 10 brief inputs. In the temporal scaling task (Fig. 1C), RNNs were trained to produce each digit at a fast or slow speed (corresponding to a total duration of 1 or 1.5 s, respectively). The different speeds were cued by the values of α . Note that we can intuit this setting as a behavioral trial (fast/slow) cued by context cues or block design, which corresponds to different neuromodulatory activity or α levels. Two conditions were then tested: In the “congruent condition,” a higher value of α (0.9) corresponded to higher speed, and a lower value of α (0.8) to the slower speed; and an “incongruent condition,” in which these relationships were reversed, i.e., higher and lower α values cued slower and faster speeds. In both the congruent (Fig. 1D) and incongruent (fig. S1A) conditions, RNNs learned the temporal scaling task equally well, as quantified by the speeds of trained output trajectories, which was 1.5 times faster at $\alpha = 0.9$ than $\alpha = 0.8$ in the congruent condition and vice versa for the incongruent condition (Fig. 1E). Training to criterion on the temporal scaling task was successful

across a diverse range of hyperparameters including the mean time constants of depression and facilitation, the α levels, and scaling factors (fig. S2, A, C, and E).

For the spatial scaling task, RNNs were trained to generate digits with the same duration but with different spatial scales ($1 \times$ and $1.5 \times$). As in the temporal scaling task, the relationship between α and the scaling factor could be congruent ($0.8 \rightarrow 1 \times$ and $0.9 \rightarrow 1.5 \times$) or incongruent ($0.9 \rightarrow 1 \times$ and $0.8 \rightarrow 1.5 \times$) (Fig. 1F). In both the congruent (Fig. 1G) and incongruent (fig. S1B) conditions, RNNs learned the spatial scaling task well, as quantified by the Euclidian distance traversed by the output trajectories, which was 1.5 times more at $\alpha = 0.9$ than $\alpha = 0.8$ in the congruent condition and vice versa for the incongruent condition (Fig. 1H). Again, the training for the spatial scaling task was robust across a diverse range of hyperparameters (fig. S2, B, D, and F).

Although RNNs can learn equally well in both the congruent and incongruent conditions, the number of epochs needed to reach the same criterion for the congruent was significantly lower than that for the incongruent conditions in both the temporal and spatial scaling tasks (fig. S1C), suggesting that the congruent condition may offer intrinsic computational advantages (see below).

To account for the natural biological variability of neuromodulation, we also tested our approach by selecting α values from a Gaussian distribution with a mean of 0.9 or 0.8 (for short/long and large/small scales) (fig. S3A). Similar to the standard congruent temporal and spatial scaling tasks, RNNs can learn both temporal and spatial scaling tasks under dynamical α values as shown by the example traces and summary speed and distance plots (fig. S3, B and C). For simplicity,

unless stated otherwise, in the subsequent simulations, we focus on the standard case using two constant α values for different scales.

These results demonstrate that in principle, the α levels, which modulate STP through the initial “release ratio” can control either temporal or spatial scales, under both congruent and incongruent conditions. These results, however, do not address the more important question of how temporal and spatial scaling generalizes to novel values of α .

Congruent modulation of STP generalizes better to novel scales in both temporal and spatial scaling tasks

To determine whether temporal and spatial scaling generalizes to novel values of α , we tested RNN performance under interpolated ($\alpha = 0.8$ to 0.9) and extrapolated ($\alpha < 0.8$, $\alpha > 0.9$) conditions by varying α values uniformly from 0.95 to 0.75 in both tasks. Optimal generalization would consist of output patterns that scaled time/space linearly with α . We quantified generalization as the root mean square error (RMSE) between the actual outputs and the linearly scaled targets (see Materials and Methods). For the temporal scaling task, RNNs generalized much better under the congruent, compared to the incongruent, condition (Fig. 2, A and C). In addition, the speed of the output trajectory scaled much more linearly in the congruent condition (Fig. 2D). In the spatial scaling task, the generalization was also significantly better (but not as markedly so) in the congruent

condition (Fig. 2, B, E, and F). We also confirmed the good generalization performance in the congruent conditions with α values from Gaussian distributions (fig. S3, D to G).

These results establish that while RNNs can be trained to perform temporal or spatial scaling in both the congruent and incongruent conditions, the congruent neuromodulation of STP is inherently superior in modulating RNN dynamics, suggesting, for example, that in the case of the temporal scaling task, increasing release probability (favoring short-term depression) might be intrinsically better suited to accelerate, than decelerate, the internal dynamics of RNNs.

Temporal and spatial profiles of recurrent dynamics

To begin to understand how neuromodulation of STP drives the scaling of RNN dynamics across temporal and spatial scales, we first analyzed the dynamics of the recurrent networks under different values of α . We plotted the population activity normalized to the maximal activity across α levels for each unit and sorted units by latency at different α values for the congruent conditions. For the temporal scaling task, the sequential order of the dynamics at $\alpha = 0.8$ in the congruent condition was largely preserved with respect to the dynamics at $\alpha = 0.9$ but shifted to the right (Fig. 3A), consistent with slower internal dynamics of the RNN. In temporal scaling task, the maximal firing rate of the units (related to the amplitude of RNN dynamics) appeared to be higher at $\alpha = 0.9$ compared to $\alpha = 0.8$

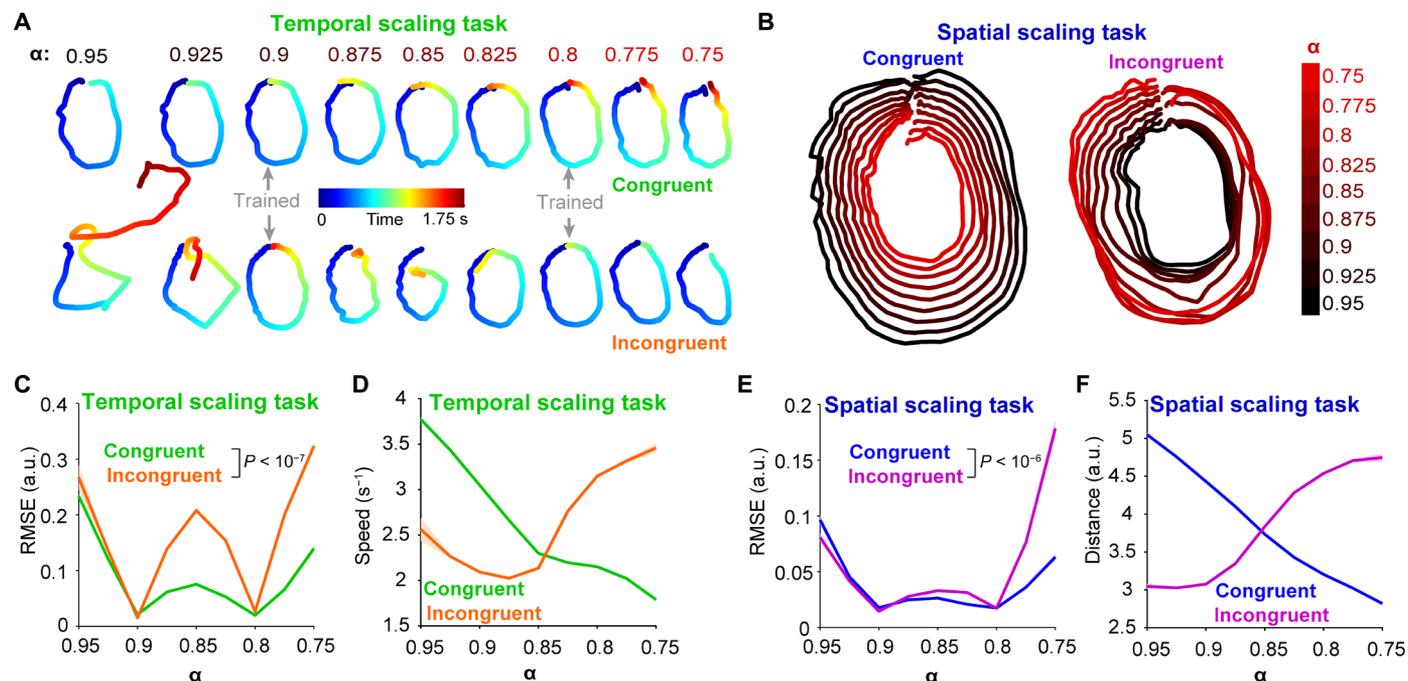


Fig. 2. Congruent modulation of STP produces better generalization to novel temporal and spatial scales. (A) Example output traces of digit 0 under novel α levels for congruent (top) and incongruent (bottom) conditions in the temporal scaling task. Gray arrows denote the α levels used for training. (B) Similar to (A) but for the spatial scaling task under congruent (left) and incongruent (right) conditions. Color codes different α levels. (C) Summary of the generalization performance for the temporal scaling task as measured by RMSE between the actual output and targets linearly warped according to the corresponding α level. Note that RMSE at the novel (untrained) α values for the congruent (green) condition was significantly lower than that for the incongruent (orange) condition ($n = 20$ RNNs; $P < 10^{-7}$, two-sided Wilcoxon rank sum test). (D) Summary of speed versus α levels for the congruent (green) and incongruent (orange) conditions in the temporal scaling task. Note that the relation for the congruent condition was more linear. (E) Same as (C) but for the spatial scaling task. RMSE for the congruent (blue) condition was significantly lower than that for the incongruent (purple) condition ($n = 20$ RNNs; $P < 10^{-6}$, two-sided Wilcoxon rank sum test). (F) Same as (D) but for distances with the congruent (blue) and incongruent (purple) conditions in the spatial scaling task. Data were presented as means \pm SEM (light overlay).

(note the change in yellow intensity). These observations were confirmed by plotting and comparing the two congruent RNN trajectories (non-normalized) in principal components analysis (PCA) space (Fig. 3C). In the spatial scaling task, the sequential order was largely preserved in the congruent condition (Fig. 3B), while the size of the trajectories for the different α values also existed (Fig. 3D).

To quantify the changes in RNN dynamics produced by the neuromodulation of STP, we developed three interrelated measures: Temporal Scaling Factor (TSF), Spatial Scaling Factor (SSF), and a Scale-Specific Index (SSI), all calculated from the same algorithm (schematized in Fig. 3E). To compare the trajectories generated with α values of 0.9 and 0.8 ($\mathbf{r}_{0.9}$ versus $\mathbf{r}_{0.8}$) we searched for the best temporal

(TSF) and spatial (SSF) warping factors that create the best match of $\mathbf{r}_{0.9}$ to $\mathbf{r}_{0.8}$. The Euclidean distance between $\mathbf{r}_{0.9}$ and $\mathbf{r}_{0.8}$ at the best temporal (TSF) and spatial (SSF) warping factors is then used to obtain an SSI value (see Materials and Methods). Intuitively, the lower the SSI, the better $\mathbf{r}_{0.8}$ can be fit by warping $\mathbf{r}_{0.9}$ temporally by TSF and spatially by SSF. Applying these measures to the temporal scaling task, we found that the TSF for the congruent condition was significantly higher than 1 and close to the target value of 1.5, suggesting a robust temporal scaling of the RNN dynamics (Fig. 3F, left). The SSI was also significantly lower than 1 (Fig. 3F, right) suggesting that in the congruent condition, RNN trajectories at $\alpha = 0.8$ were a linearly warped version of the $\alpha = 0.9$ trajectory. Note that

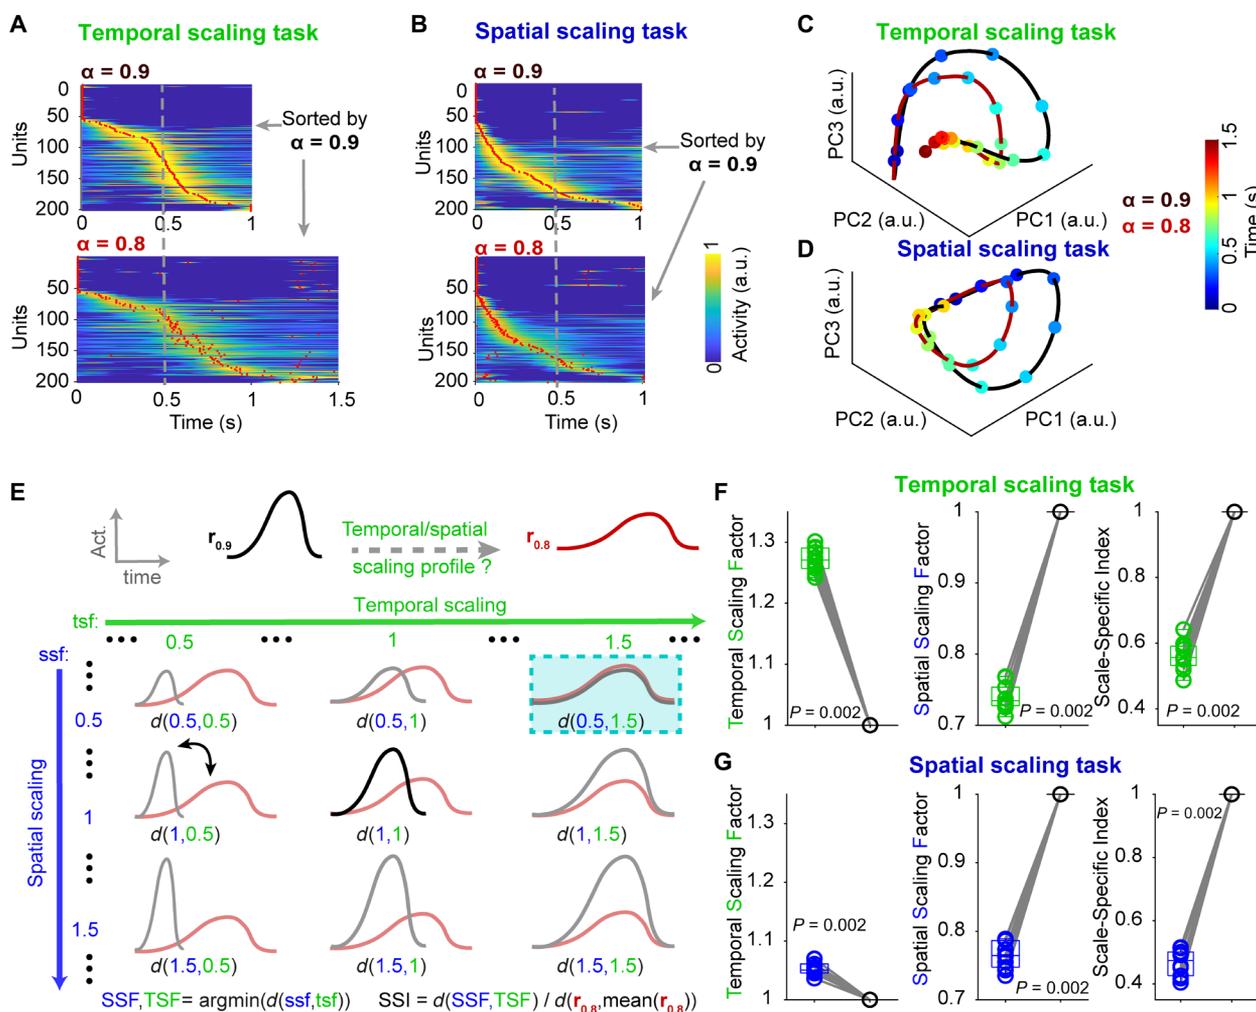


Fig. 3. Temporal and spatial scaling of recurrent dynamics. (A) Normalized recurrent population activity at $\alpha = 0.9$ (top) and $\alpha = 0.8$ (bottom) sorted according to the peak activity latency at $\alpha = 0.9$ for congruent temporal scaling tasks. The red dots denote the peak activity time for each unit. The gray dashed line denotes 0.5 s. The activity of each unit is normalized to its maximal activity across α levels. (B) Same as (A) but for spatial scaling task. (C) Plot of the first three principal components (PCs) of population activity at $\alpha = 0.9$ (black) and 0.8 (dark red) in temporal scaling task. The color bar codes for time. (D) Same as (C) but for the spatial scaling task. (E) Schematic of the TSF, SSF, and SSI. For two hypothetical neural trajectories $\mathbf{r}_{0.9}$ (black) and $\mathbf{r}_{0.8}$ (red), $\mathbf{r}_{0.9}$ is temporally warped by either linear interpolating or subsampling a range of candidate temporal scaling factors (tsf). Each time-warped trajectory is multiplied with a candidate spatial scaling factor (ssf), resulting in a grid of temporally spatially warped dynamics of $\mathbf{r}_{0.9}$ (gray). The Euclidean distance between these warped dynamics and $\mathbf{r}_{0.8}$ (light red traces on top) is computed. The tsf and ssf leading to the minimal distance are defined as TSF and SSF, respectively. SSI is defined as the distance at TSF/SSF divided by the distance between $\mathbf{r}_{0.8}$ and its mean. (F) Comparison of the average TSF (left), SSF (middle), and SSI (right) values against 1 for the congruent conditions across 20 RNNs in the temporal scaling task ($n = 10$ digits; $P = 0.002$, two-sided Wilcoxon signed-rank test for TSF, SSF, and SSI). (G) Same as (F) but for the spatial scaling task ($n = 10$ digits; $P = 0.002$, two-sided Wilcoxon signed-rank test for TSF, SSF, and SSI).

the SSF for the temporal scaling task is below 1 (Fig. 3F, middle), indicating that the faster trajectory is accompanied by higher-amplitude firing rates of the RNN units.

Quantification of RNN trajectories in the spatial scaling task revealed that the SSF and SSI for the congruent condition were both significantly lower than 1 (Fig. 3G, middle, right), suggesting a robust spatial scaling of the RNN dynamics across different α . As expected, because the duration of both the large and small outputs is similar, the TSF for the spatial scaling task was close to 1 (Fig. 3G, left).

In the incongruent conditions, the temporal and spatial scaling profiles in both the temporal and spatial scaling tasks are significantly weaker as quantified by the same TSF, SSF, and SSI measures compared to that in the congruent condition (fig. S4). To further corroborate the above measures we performed time and unit shuffled controls (fig. S5). As expected, shuffling resulted in significantly higher SSI values and disrupted the temporal and spatial profiles as quantified by the TSF and SSF. We performed the same analyses based on the temporal profile of the synaptic efficacy as defined by the product of \mathbf{x} and \mathbf{u} of STP. This revealed similar temporal and spatial scaling profiles as the dynamics of activity (fig. S6).

These results demonstrate that in the temporal scaling task, the congruent neuromodulation of STP produced temporal scaling of output trajectories, by the temporally scaling of RNN trajectories. That is, while an increase in neuromodulator concentration could potentially increase or decrease release probability, these results show that congruent modulation—where higher synaptic release probabilities (favoring depression) correspond to faster speeds—is intrinsically better at temporally scaling RNN dynamics. Similarly, but to a lesser degree, in the spatial scaling task congruent modulation—where higher synaptic release probabilities correspond to larger sizes—is better suited to produce a uniform modulation of firing rate amplitude and thus spatial scaling of the output.

To dissect the mechanisms underlying the differential scaling of RNN trajectories during congruent and incongruent neuromodulation of STP, we focused on the RNN state \mathbf{s} (Eq. 1 in Materials and Methods) trajectories at $\alpha = 0.9$ and $\alpha = 0.8$ ($\mathbf{s}_{0.9}$ and $\mathbf{s}_{0.8}$, respectively, in fig. S7A). At any given time point in $\mathbf{s}_{0.9}$ and $\mathbf{s}_{0.8}$, there are corresponding velocity vectors $\mathbf{v}_{0.9}$ and $\mathbf{v}_{0.8}$, that can be decomposed into a decay component ($\mathbf{d}_{0.9}$ and $\mathbf{d}_{0.8}$) and a recurrent component ($\mathbf{rec}_{0.9}$ and $\mathbf{rec}_{0.8}$). In addition, there is a vector, \mathbf{p} , denoting the direction from $\mathbf{s}_{0.8}$ to $\mathbf{s}_{0.9}$ (normalized time). We defined the angle between $\mathbf{rec}_{0.8}$ and \mathbf{p} as θ and the angle between $\mathbf{rec}_{0.8}$ and $\mathbf{v}_{0.9}$ as μ . Intuitively, increasing α from 0.8 to 0.9 should increase the recurrent drive. For $\mathbf{s}_{0.9}$ to speed up compared to $\mathbf{s}_{0.8}$, we should expect μ to be smaller than 90° . If $\mathbf{s}_{0.9}$ is larger than $\mathbf{s}_{0.8}$, we should expect the angle θ to be small to increase drive in the \mathbf{p} direction.

In the temporal scaling task, the mean μ across time in the congruent condition was less than 90° , in contrast to the incongruent condition in which it was larger than 90° (fig. S7B, left), indicating a faster and slower $\mathbf{s}_{0.9}$ trajectory in the congruent and incongruent conditions, respectively. θ in the congruent condition was significantly lower than in the incongruent condition (fig. S7B, right), indicating that the difference in the size (amplitude) between the $\mathbf{s}_{0.8}$ and $\mathbf{s}_{0.9}$ trajectories was larger in the congruent condition (consistent with SSF in the congruent condition being significantly lower as in Fig. 3F, middle).

For the spatial scaling task, θ is significantly lower in the congruent compared to the incongruent condition (fig. S7C, right), implying that the difference in size between the $\mathbf{s}_{0.8}$ and $\mathbf{s}_{0.9}$ trajectories was

larger in the congruent condition (consistent with SSF in the congruent condition being significantly lower as in Fig. 3G, middle). In both the congruent and incongruent conditions, μ was close to 90° (fig. S7C, left), suggesting that the changes in speed from $\mathbf{s}_{0.8}$ and $\mathbf{s}_{0.9}$ trajectories were not as marked as in the temporal scaling task.

As shown in Fig. 3, for both tasks the SSF values of the recurrent dynamics in both the congruent and incongruent conditions are less than 1, which indicates that the size of the recurrent dynamics at $\alpha = 0.9$ is larger than $\alpha = 0.8$. In the spatial scaling task, the larger recurrent dynamics at $\alpha = 0.9$ in the congruent condition would be appropriate for generating larger output as the task requires. However, in the incongruent condition, the larger trajectory paradoxically generates a smaller output (as required by training). The solution to this paradox can be understood by using a light projection analogy (fig. S7D). In light projection, to get a smaller shadow from a larger trajectory, one can arrange the larger trajectory to have a larger angle with the ground, the plane on which the shadow is located. Analogously in our model, the output trajectory is the projection of the recurrent dynamics onto the output space governed by the output weights. Thus, in the incongruent spatial scaling task, we would expect the angle between larger recurrent dynamics [subspace of the first two principal components (PCs)] and output space to be larger than that between smaller recurrent dynamics and output space. In the incongruent spatial scaling task, the angle between recurrent space at $\alpha = 0.9$ (which has larger dynamics) and the output space was slightly but significantly higher than at $\alpha = 0.8$ (fig. S7E, right). As expected, there was no such angle difference in the congruent spatial scaling (fig. S7E, left).

Relatedly, in the temporal scaling task, the α -induced change in trajectory amplitude was larger in the congruent condition (Fig. 3F, middle). Specifically, SSF in the congruent condition was smaller (reflecting a smaller trajectory at $\alpha = 0.8$), even though it generated an output of the same size. Again, this paradox is resolved by the smaller projection angle φ when $\alpha = 0.8$ (fig. S7E, left). Last, these findings are robust when quantifying the angles for higher dimensional recurrent space expanded by more PCs (fig. S8).

α -cued scaling outperforms input amplitude-cued scaling for generalization performance in temporal scaling task

In addition to the current approach of using neuromodulation of STP to govern different temporal and spatial scales, a series of previous studies have used the amplitude of external input in different scales (1, 2, 13, 19, 20)—e.g., by incorporating a separate speed or “size” input to cue different speeds or sizes. To directly compare the generalization performance between these two distinct approaches, we trained the same RNNs without STP to perform the same temporal and spatial scaling but used the amplitude of additional input to cue different scales in the congruent condition: i.e., higher input amplitude corresponds to faster speed or larger size (Fig. 4A). We trained RNNs under a series of amplitude pairs ranging from 0.9/0.85 to 0.9/0.1 to match the values used by most of the previous studies (Fig. 4, B and C). We then performed the same generalization analysis as above and found that amplitude pair 0.9/0.7 reached the best generalization performance in both temporal and spatial scaling tasks for the input amplitude-cued approach (Fig. 4, D and E). The generalization performance for the STP-based scaling approach is significantly better than the input amplitude-cued approach under all the tested amplitude pairs in the temporal scaling task (Fig. 4D) but comparable or slightly worse in the spatial scaling task (Fig. 4E). To further

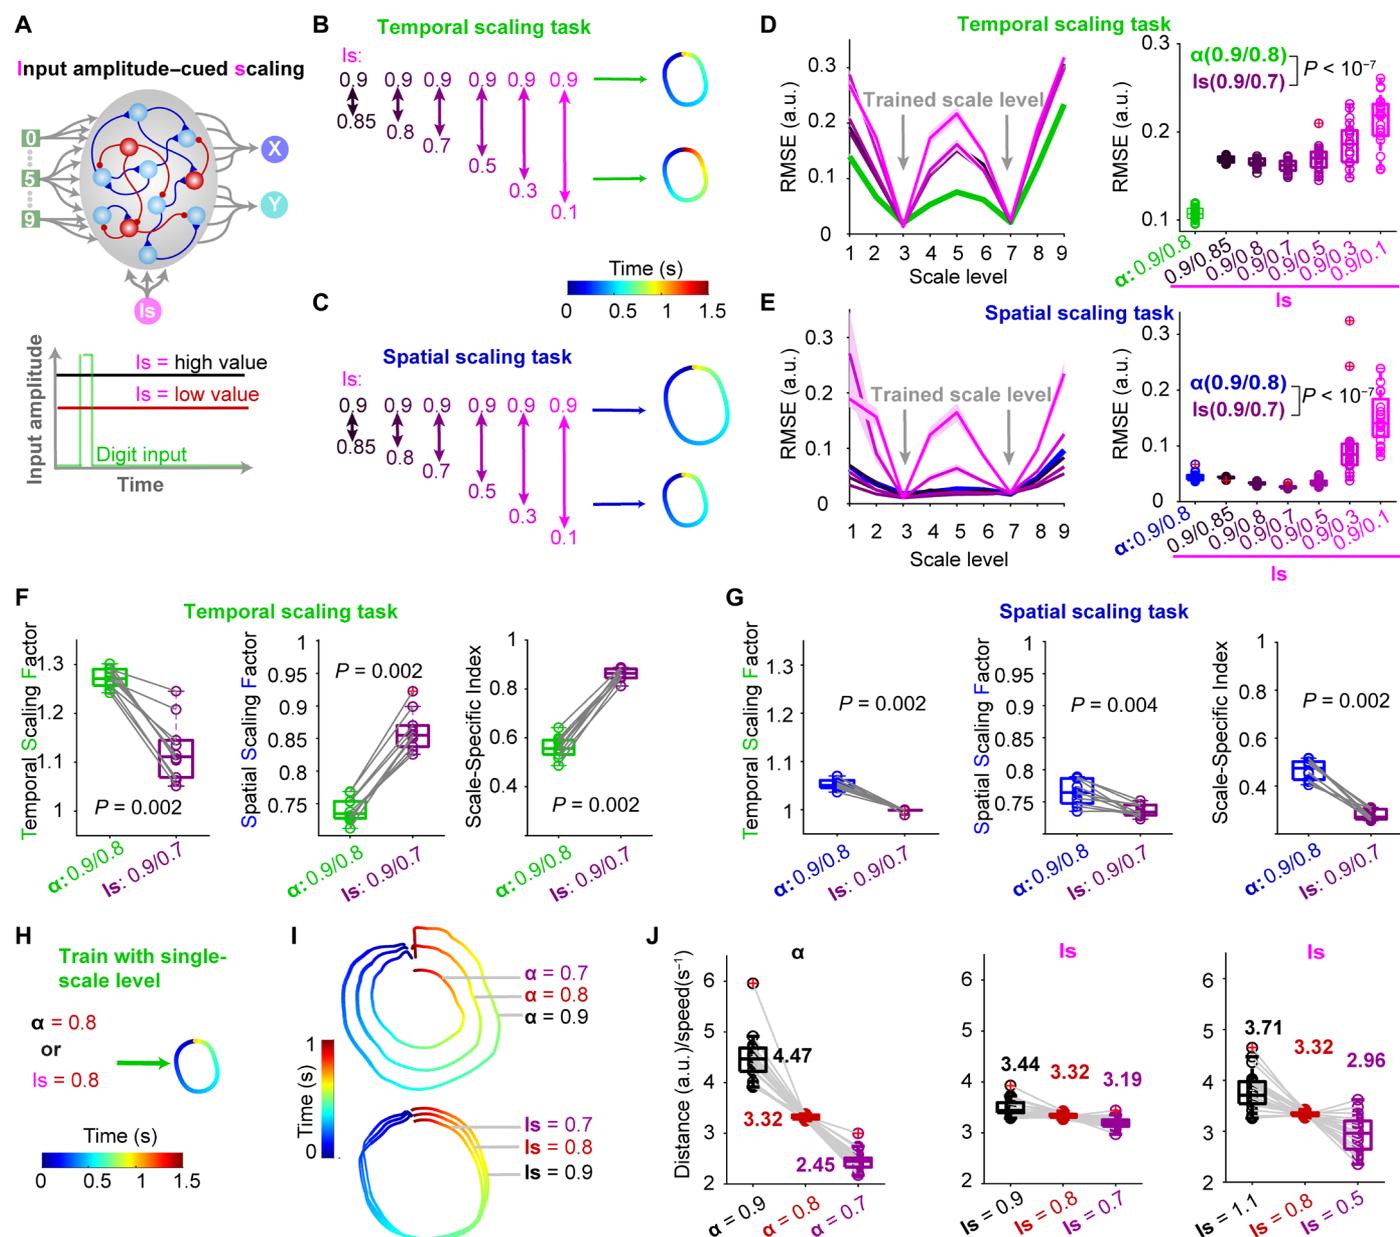


Fig. 4. RNNs trained with temporal and spatial scales cued by input amplitude. (A) Schematic of RNNs trained with input amplitude scaling (Is), where an external input that cues scale is continuously presented during the whole trial. (B) Similar to the congruent settings for the standard α -cued scaling task, a series of pairs of high/low inputs correspond to the short/long duration. (C) Similar to (B) but for the spatial scaling task. (D) Left: The generalization performance (RMSE) across linearly spaced novel levels for the two approaches on the temporal scaling task: input amplitude–cued (graded magenta) and α -cued (green). Right: Comparison of averaged RMSE across all generalization levels for α (0.9/0.8) and Is (0.9/0.7) ($n = 20$ RNNs; $P < 10^{-7}$, Wilcoxon rank sum test). (E) Same as (D) but for the spatial scaling task ($n = 20$ RNNs; $P < 10^{-7}$, Wilcoxon rank sum test). (F) Comparison of the average TSF (left), SSF (middle), and SSI (right) values for α (0.9/0.8) and Is (0.9/0.7) across 20 RNNs in the input amplitude–cued temporal scaling task ($n = 10$ digits; $P = 0.002$, two-sided Wilcoxon signed-rank test for TSF, SSF, and SSI). (G) Same as (F) but for the input amplitude–cued spatial scaling task ($n = 10$ digits; $P = 0.002$, 0.004, and 0.002, two-sided Wilcoxon signed-rank test for TSF, SSF, and SSI, respectively). (H) Schematic of training RNNs with only one-scale level 0.8 in both α -cued and input amplitude–cued scaling task. (I) Example output traces of digit 0 on scale level of 0.7, 0.8, and 0.9 in both α -cued and input amplitude–cued scaling tasks. (J) The averaged output distance during 1 s for RNNs trained with single-scale level 0.8 and tested with $\alpha = 0.9/0.7$ (left), Is = 0.9/0.7 (middle), and Is = 1.1/0.5 for a wider range (right).

understand why input amplitude–cued scaling task is worse in temporal generalization and the underlying temporal-spatial profile of neural dynamics across different scales, we performed the same TSF, SSF, and SSI analysis as above. The TSF for RNNs with input amplitude–cued temporal scale is above 1 (Fig. 4F, left), and SSI is below 1 (Fig. 4F, right), implying that recurrent dynamics for different scales are temporally

scaled to some extent, which is consistent with previous studies (1, 2, 13, 19, 20). However, input amplitude–cued scaling led to lower TSF and higher SSI compared to those for α -cued scaling, which implies that the α -cued scaling approach gives rise to better temporal scaling for dynamics across different scales and thus leads to a better generalization performance. This reasoning can be applied to explain

the difference in the spatial scaling tasks, in which the difference between SSF and SSI is slight (Fig. 4G), which might lead to the minor difference in the generalization performance.

To further test the intrinsic generalization performance for both approaches, we trained RNNs with a single-scale level (either single α or single input amplitude) and tested with novel-scale levels (Fig. 4H). We found that RNNs trained with the α -cued scaling approach tended to generalize better to novel-scale levels as shown by more temporally and spatially scaled output example traces (Fig. 4I) and more distance traversed in 1 s (equal to the average speed) (Fig. 4J).

Note that for all the comparisons between the STP-controlled (α -cued) and input amplitude-controlled scaling approach, we implemented the same RNNs with the same number of units and synapses. Thus, generally, the α -cued approach would have extra dynamical variables as in the STP model, and it requires more time to train. However, because the STP parameters were not trained, the numbers of trained variables in both approaches were the same, which makes the comparison fair to some extent.

In sum, compared to previous approaches—such as using input amplitude to signal the scales, neuromodulation of STP (α -cued scaling approach) is better at generalization to novel scales in temporal scaling tasks but less so in spatial scaling tasks when trained with two-scale levels. Furthermore, when trained with only one scale, the α -cued scaling approach intrinsically generalizes better in both temporal and spatial scaling.

Short-term plasticity enhances generalization and speeds up training

The above results demonstrate that we can modulate the temporal or spatial scale by changing the initial synaptic release probability controlled by α . Although α controlled release probability U , it did not directly address whether STP is actually contributing to the results. That is, does simply adjusting the synaptic release probability in the absence of STP (i.e., simply scaling all synaptic weights) result in similar performance? To address that question, we ran control simulations for RNNs without STP but still included an α term. Specifically, we fixed x at 1 and u at αU during the whole trial. These modifications removed the synaptic dynamics but allowed α to scale the strength of the nondynamical synapses in the RNN. We then trained and tested the RNNs without STP using the same task in the congruent condition as in the standard model. The generalization performance for the RNNs without STP markedly decreased for both the temporal (Fig. 5, A and C) and spatial scaling tasks (Fig. 5, B and C). Furthermore, the absence of STP markedly slowed training as shown by the increase in training epochs needed to reach the same criterion (Fig. 5D). These results suggest that STP does indeed provide a unified mechanism to effectively scale temporal and spatial neural dynamics.

Joint control of temporal and spatial scales and shape via neuromodulation of STP

Up to now, we have demonstrated that adjusting α can control either temporal or spatial scales separately. We next ask whether α can jointly control temporal and spatial scales in a single RNN (here, the relationship between α and the temporal and spatial scales was always congruent). To explore this possibility, we arbitrarily divided the recurrent units into two groups. One group received a neuromodulatory signal that controls temporal scale, and another group receives a signal that controls spatial scale (Fig. 6A). RNNs learn the task well (Fig. 6B), and in this joint control task, RNNs generalize to all the

combinations of novel α levels for the temporal and spatial scales (Fig. 6C, left). Note that speed (defined by the mean distance traversed per second) progressively increases toward the top left part of the speed plot, which corresponds to the fastest speed required to generate the largest output (spatial $\alpha = 0.95$ for large size) in the shortest time (temporal $\alpha = 0.9$ for 1 s). Similarly, the distance progressively increases toward the bottom left (Fig. 6C, right), which corresponds to the longest distance required to generate the largest output in the longest time.

We next asked whether α can jointly control an additional task dimension: the output shape, i.e., rather than cueing digit identity by distinct inputs as in the above implementations, the digit identity was cued by changing α in a subset of units. To achieve this, we divided the recurrent units into three groups (50, 25, and 25%): 50% for digit shape, 25% for temporal scale, and 25% for spatial scale (Fig. 6D). The digit group was further subdivided into 10 subgroups corresponding to each of 10 digits, and the α value of each set corresponding to 0.6 or 1 (Fig. 6D). With this architecture, RNNs learned to generate all digits at all trained durations and spatial scales. Again, the RNNs generalized well to novel joint values α used to control temporal and spatial scales (Fig. 6, E and F).

Cueing digit identity by using either different inputs or α values resulted in similar learning and generalization performance. PCA plots of RNN dynamics revealed that using α to signal digit identity led to recurrent dynamics that was more similar across digits (fig. S9A). These observations are further confirmed by the cross-digit correlations for the two digit encoding strategies (fig. S9, B and C)—reflecting lower dimensional RNN dynamics across all digits when digit identity is cued by α . These findings suggest that similar output can be generated from different recurrent dynamical regimes.

Neuromodulation of STP captures the temporal scaling observed in two sensorimotor tasks

Above, we mainly focused on how α can control the temporal and spatial scales in a motor task—generating handwritten digits. To investigate whether modulating α can account for experimental findings on temporal scaling, we next simulated two experimental studies: one from rodents (9) and another from nonhuman primates (1). First, we trained RNNs to solve an interval-alternative-forced-choice (IAFC) task where rats needed to classify intervals as short or long (9). In this study, optogenetically increasing dopamine levels selectively shifted the decision toward the short intervals (i.e., the psychometric curve was shifted right). To simulate this experiment, we first trained RNNs to discriminate the same intervals used in the experimental task using a single α value of 0.8 for all units (Fig. 7A). At this trained condition, the RNNs replicated the behavioral results as shown by the output traces and psychometric function (Fig. 7, B and C). We then sought to simulate the dopamine manipulation experiments. Multiple studies have demonstrated that dopamine decreases the synaptic release probability in both excitatory and inhibitory cortical synapses (29–33). We thus simulated low or high dopamine levels in the IAFC task by changing α to 0.9 or 0.7, respectively. Decreasing α from 0.8 to 0.7 shifts the psychometric curve to the right, as in the dopamine manipulation experiments, and vice versa when α was increased to 0.9 (Fig. 7, D and E). It is important to note that, here, the RNNs were only trained with $\alpha = 0.8$; thus, the left and rightward shifts in the psychometric function with α values of 0.9 and 0.7, respectively, occur naturally—i.e., they reflect intrinsic properties of RNN dynamics. This intrinsic relationship between increasing RNN

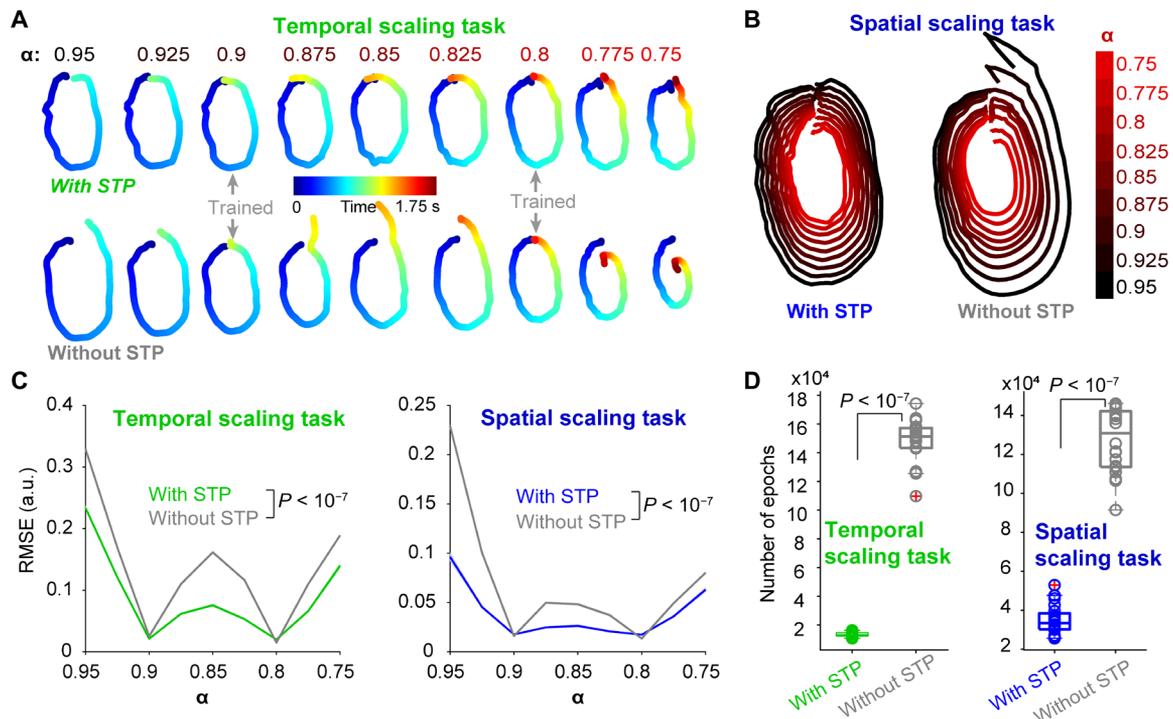


Fig. 5. STP enhances generalization and speeds up learning. (A) Example output traces of digit 0 under different α levels for the congruent temporal scaling task with STP (top) and congruent condition without STP (bottom)—in which case scale was cued by scaling all weights in the absence of STP. Gray arrows denote the α level used for training. (B) Same as (A) but for the spatial scaling task. Color codes for different α levels. (C) Summary of the generalization performance as measured by the average RMSE between output and the linearly scaled targets. Note that in both the temporal (left) and spatial (right) scaling tasks, the RMSE for the novel α levels for RNNs with STP (green or blue) was significantly lower than for RNNs without STP (gray) ($n = 20$ RNNs; $P < 10^{-7}$, two-sided Wilcoxon rank sum test for the temporal and spatial scaling tasks). (D) Comparison of the number of training epochs to reach criteria for RNNs with STP and without STP in temporal scaling task (left) and spatial scaling tasks (right) ($n = 20$ RNNs; $P < 10^{-7}$, two-sided Wilcoxon rank sum test for both tasks).

trajectory speed and higher values of α also accounts for the superior performance of the RNNs above in the congruent conditions.

We next simulated an experimental task that used the Ready-Set-Go task in monkeys (1). Here, subjects were presented with temporal intervals demarcated by ready and set cues, and they had to produce a set-go interval that reproduced the ready-set interval. In this study, there was an additional scaling cue, which determined whether subjects were required to produce the exact ready-set interval (1 \times) or scale the interval by 1.5 \times . To simulate this flexible sensorimotor timing task, we trained RNNs to do the task using $\alpha = 0.9/0.8$ to cue the 1 \times /1.5 \times context, respectively (Fig. 7, F and G). RNNs learned this task well and captured some important features of the behavior, such as the regression to the mean effect—i.e., a bias of the produced intervals toward the mean interval (Fig. 7, G and H). Similar to the digit production task (Fig. 2), RNNs generalized well to novel α levels (ranging from 0.75 to 0.95), resulting in scaling values of just below 1 \times to just above 1.5 \times (Fig. 7I).

DISCUSSION

Here, we have proposed, and provided support, for the hypothesis that neuromodulation of STP provides a mechanism for RNNs to scale their dynamics in both time and space. RNNs that incorporated STP and used neuromodulation of STP to signal changes in temporal and/or spatial scale exhibited better generalization and performance

than RNN models in which temporal and spatial scales were signaled by changes in absolute synaptic weights alone (Fig. 5) or distinct inputs (Fig. 4, comparable for spatial scale). Furthermore, neuromodulation of STP allowed RNNs to capture the results of two experimental studies based on distinct sensorimotor timing tasks. While neuromodulation of STP is a well-established experimental phenomenon in cortical and subcortical circuits alike (25, 27, 28, 35, 49–52), its potential role in neurocomputation has not been addressed. Here, we demonstrate that it provides a unified mechanism for the flexible regulation of neural dynamics and thus of the control of temporal and spatial scales for diverse sensorimotor tasks.

Temporal scaling of sensorimotor behavior through neuromodulation of STP

Neuromodulators such as dopamine have been implicated in a large range of cognitive functions including reinforcement learning (53) and timing. In the case of timing, it has been proposed that dopamine may alter “clock speed” (7, 9, 54). How dopamine could alter the speed of timers at the neural level, however, has not been addressed. Here, we propose that dopamine’s ability to modulate STP provides a mechanism to link findings at the neural and cognitive levels. Specifically, in cortical circuits, synaptic transmission studies indicate that dopamine can decrease excitatory postsynaptic potential amplitude by decreasing synaptic release probability (25, 28, 29, 35). By incorporating STP into RNNs and emulating dopaminergic inhibition of

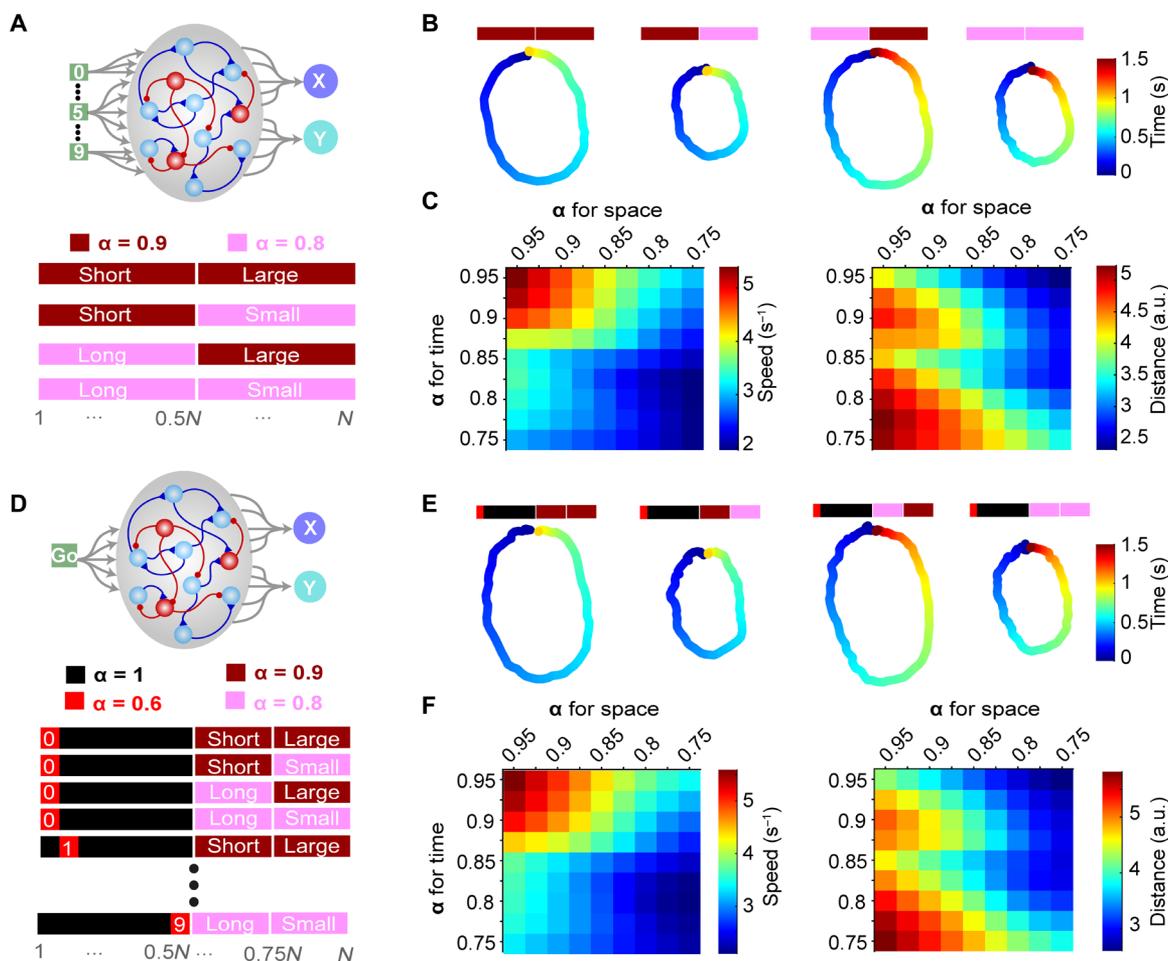


Fig. 6. Joint control of temporal and spatial scales in RNNs through differential modulation of α in distinct subpopulations. (A) Schematic of the joint control of temporal and spatial scales. The α level of 50% of RNN units controlled temporal scaling, and the other half spatial scaling. (B) Example output traces for digit 0 in four cases: short-large, short-small, long-large, and long-small. (C) Summary of average speed (left) and distance (right) across different α levels. (D) Schematic of the encoding temporal and spatial scales, as well as digit identity through α . Fifty percent (of $N = 400$) units to signaled digit identity, while half of the remaining cued duration and spatial scale. (E) Same as (B) but for the model in (D). (F) Same as (C) but for the model in (D).

release probability through decreases in the variable α , we were able to link cellular-level observations with previous systems and behavioral-level results (9).

The finding that neuromodulation of STP may serve as a neural mechanism for temporal scaling is consistent with the fact that dopamine has been linked to timing for decades (54–58). However, there is a long-standing debate as to the direction of this relationship, which has remained a point of controversy (59). Specifically, early reports suggested that neuropharmacologically enhancing dopamine levels accelerated the neural clock (7, 60) or that there was no consistent effect (8). In contrast, more recent optogenetic studies suggested that dopamine slowed the neural clock (9). Assuming that dopamine acts, in part, by decreasing release probability—as in the canonical case of dopaminergic neuromodulation (25, 35)—our results strongly predict that increasing dopamine (decreasing release probability) should slow the internal clock. Specifically, as shown in Fig. 2, the congruent relationship where high values of α (low dopamine) increase speed results in better performance compared to the incongruent condition where decreasing α increases speed. Furthermore, when RNNs are trained on an interval discrimination task at only

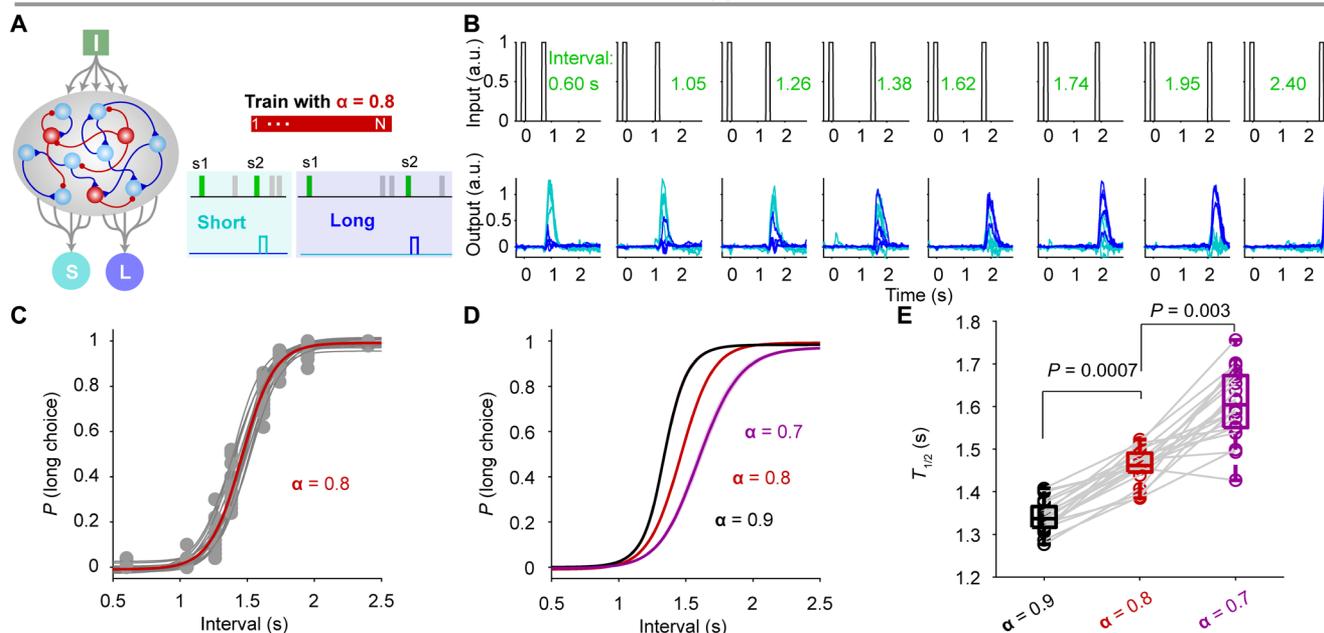
one α level, subsequently decreasing α (high dopamine) reveals an intrinsic slowing of RNN dynamics (Fig. 7, A to E).

We emphasize, however, that dopaminergic modulation of synaptic transmission is complex and dependent on brain areas and synapse classes (28, 29, 35, 49). Dopamine can affect synaptic plasticity, and intrinsic excitability, in addition to synaptic strength and STP (49, 61). Similarly, at the behavioral and cognitive levels, the effects of dopamine are also highly complex. In the motor domain, for example, dopamine has been demonstrated to increase movement amplitude through modulation of the striatal activity (62), potentially, reflecting differences in dopaminergic neuromodulation in different brain areas. Thus, future directions should explore how the additional modes of action of dopamine would shape the temporal and spatial scales of neural dynamics.

Temporal versus spatial scaling

Previous studies have shown that RNNs can account for temporal scaling by cueing speed through the amplitude of a tonic speed input (1, 2, 13, 19, 20, 63) or by altering the gain through changes in intrinsic excitability (11, 21). As in the current study, in these previous computational models, as well as in experimental studies (1, 19, 64,

Soares et al. (9)



Remington et al. (1)

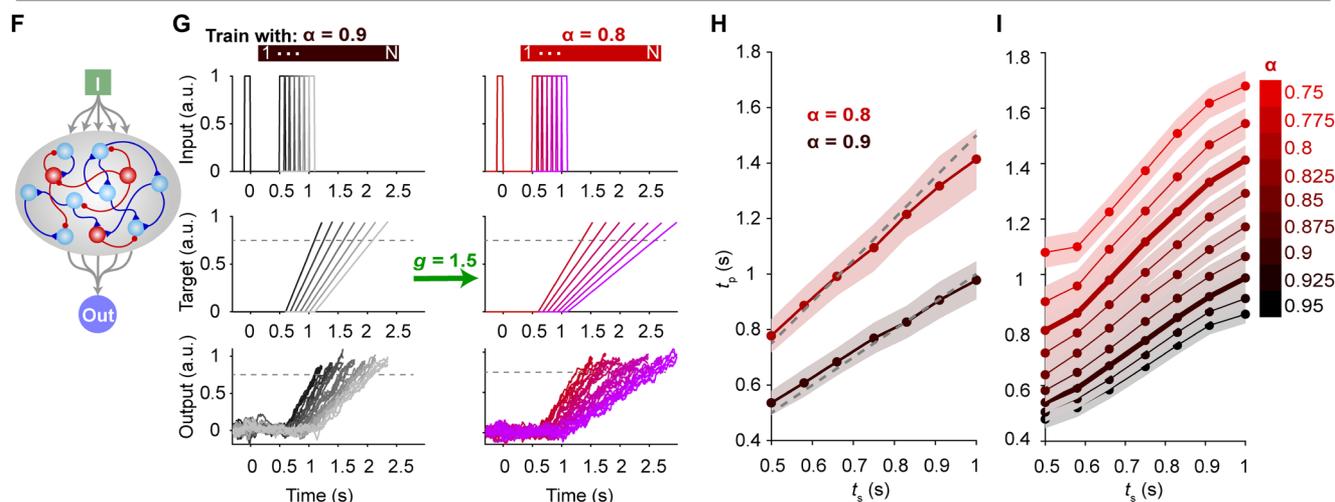


Fig. 7. Simulations of two sensorimotor timing tasks. (A) Schematic of RNN used to simulate an interval discrimination task. RNNs were composed of one input for delivering two events demarcating a range of intervals between 0.6 and 2.4 s and two outputs corresponding to a short or long decision (short <1.5 s; long >1.5 s). RNNs were trained only with $\alpha = 0.8$. (B) Output traces of an example RNN. (C) Sigmoidal fits of the long choice probability tested at $\alpha = 0.8$. (D) Similar to (C) but at $\alpha = 0.9$ (black), $\alpha = 0.7$ (purple), and $\alpha = 0.8$ (red) for comparison. (E) Summary of the time “point-of-subjective equality” ($T_{1/2}$) for the sigmoid fits in (D). Changing α significantly changed the $T_{1/2}$ [$n = 20$ RNNs, Kruskal-Wallis test, $P < 10^{-10}$, $\chi^2_{(2,57)} = 48.0$], and $T_{1/2}$ for $\alpha = 0.9$ and 0.7 was significantly lower and higher than for $\alpha = 0.8$, respectively ($P = 0.0007$ and $P = 0.003$, Dunn’s multiple comparison test). (F) Schematic of RNN used to simulate a Ready-Set-Go task. RNNs were composed of one input that delivered two events (demarcating the Ready-Set interval) and one output. On the basis of the context, cued by $\alpha = 0.9$ or 0.8 , the output unit should generate an interval of $1\times$ or $1.5\times$ the Ready-Set interval. (G) Plot of the input (top), target (middle), and output traces for $\alpha = 0.9$ (left) and 0.8 (right) for an example RNN. The dashed line denoted the crossing time threshold. (H) Plot of the production time (mean \pm SD) versus the sensory time for $\alpha = 0.9$ (black) or $\alpha = 0.8$ (red) in an example. (I) Summary of the average production time (mean \pm SEM) for novel α levels with trained α shown thicker lines, across RNNs ($n = 20$).

65), temporal scaling was achieved by creating parallel neural trajectories that flowed at different speeds (Fig. 3). Thus, there appear to be multiple mechanisms to flexibly control the temporal scale of outputs while largely preserving the geometry of the RNN dynamics.

We also directly compared previous approaches with neuromodulation of STP by implementing the input amplitude–cued mechanism

in RNNs but in the absence of STP (Fig. 4) as well as directly modulating synaptic strength in the absence of STP. Neuromodulation of STP led to significantly better generalization in both the temporal and spatial scaling tasks, compared to weight modulation alone—i.e., modulating the weights of synapses that did not exhibit STP (Fig. 5). Neuromodulation of STP also led to markedly better generalization in

the temporal scaling task compared to when scaling was cued by the value of a speed input (Fig. 4D). In the case of the spatial scaling task, there was not a marked difference between generalization through neuromodulation of STP neuromodulation and input amplitude approaches (Fig. 4E)—and the latter led to slightly better generalization. However, after training with a single α or input level, changing α led to more robust spatial scaling than changing the input level (Fig. 4, I and J). Furthermore, STP might intrinsically speed up the training (Fig. 5) by potentially stabilizing the dynamics (66).

These results suggest that while neuromodulation of STP can underlie both temporal and spatial scaling, neuromodulation of STP is particularly well suited to control temporal dynamics (Figs. 2 and 5). We note that this is consistent with the experimental observation that dopamine—a well-known modulator of STP—has been strongly associated with temporal scaling rather than spatial scaling.

Conclusions and predictions

Even though STP is universally present at cortical synapses, the potential computational function of STP and its neuromodulation has remained largely unaddressed. Here, we propose that neural modulation of STP provides a unified biological mechanism for controlling temporal scaling, and to a lesser degree spatial scaling, of neural dynamics and sensorimotor behaviors. We found that when STP is incorporated into the synapses of RNNs, the neural dynamics of these RNNs are naturally accelerated by increasing α (i.e., shifting STP toward short-term depression by increasing a variable related to release probability). Because increasing release probability is equivalent to increasing synaptic weights when a synapse is first activated, we performed control experiments in which temporal scaling was associated with changes in absolute synaptic strength in the absence of STP. The superior generalization and marked decrease in training time in the presence of STP indicate that there is an intrinsic interaction between the short-term dynamics of individual synapses and the global neural dynamics of RNNs. Multiple factors are likely to contribute to this interaction, including that short-term depression helps counteract the positive feedback within RNNs and that shifting the effective strength of dynamic synapses to the onset of neural activity (high α) naturally accelerates the speed at which activity flows through RNNs. The nature of this interaction provides a link between the synaptic actions of dopamine on STP and long-standing data linking dopamine to alterations in timing and time perception. In addition, we provide a theoretical support for the long-standing debate as to whether dopamine increases or decreases neural clock speed; specifically, we predict that an increase in dopamine should slow down the neural trajectories of cortical circuits.

Overall, our results demonstrate that the incorporation of STP and its neuromodulation into RNN models provides a powerful and flexible mechanism to implement temporal and spatial scaling of the dynamics of RNNs. These results thus provide a hypothesis as to why synapses may exhibit STP, and provide a computational mechanism for the neural dynamics, and thus temporal and spatial control of sensorimotor tasks.

MATERIALS AND METHODS

RNN model

Network architecture and STP

RNNs were based on firing-rate units that obeyed Dale’s law ($N = 200$ unless otherwise specified, 80/20% excitatory/inhibitory). RNN dynamics was described by the following equations

$$\tau \frac{ds}{dt} = -s + W^{rec*}(\mathbf{r} \cdot \mathbf{x} \cdot \mathbf{u}) + W^{in*} \mathbf{I} + \sigma \mathbf{N}(0, 1) \sqrt{2\tau} \quad (1)$$

$$\frac{dx}{dt} = \frac{1-x}{\tau_x} - \mathbf{u} \cdot \mathbf{x} \cdot \mathbf{r} \quad (2)$$

$$\frac{du}{dt} = \frac{\alpha \cdot \mathbf{U} - \mathbf{u}}{\tau_u} + \alpha \cdot \mathbf{U} \cdot (1 - \mathbf{u}) \cdot \mathbf{r} \quad (3)$$

$$\mathbf{o} = W^{out*} \mathbf{r} + \mathbf{b}^{out} \quad (4)$$

$$\mathbf{r} = \text{relu}(\mathbf{s}) \quad (5)$$

where $\mathbf{s} \in \mathbb{R}^{N \times 1}$ represents the state of the RNN units, and the firing rate vector \mathbf{r} corresponds to the rectified linear activation function on \mathbf{s} . The time constant τ was 100 ms for all units. $W^{in} \in \mathbb{R}^{N \times Ni}$ is the input weight matrix, and Ni represents the number of inputs. \mathbf{I} represents the external inputs. Each unit received independent Gaussian noise $\mathbf{N}(0,1)$ with an SD of $\sigma\sqrt{2\tau}$. Unless otherwise specified, $\sigma = 0.01$. $W^{rec} \in \mathbb{R}^{N \times N}$ is the recurrent weight matrix. Self-connections were absent in the network. Asterisks represent matrix multiplication and dots elementwise multiplication.

STP was incorporated as in previous firing-rate models (38, 41). Specifically, cell-specific STP was implemented in the recurrent units as described in Eqs. 2 and 3: The depression variable \mathbf{x} and facilitating variable \mathbf{u} were shared for all synapses from the same presynaptic neuron. The vector \mathbf{U} corresponds to the initial synaptic release probability or baseline percentage of available transmitter released. To implement neuromodulation of STP, we scaled \mathbf{U} with a factor α in the range of 0 to 1.

The output (\mathbf{o}) of the network was computed linearly from the output weights W^{out} and \mathbf{r} with a bias term \mathbf{b}^{out} . RNNs were implemented and trained in Tensorflow 2.3 based on the code from a previous study (67).

Training

Networks were trained using the adaptive moment estimation stochastic gradient descent algorithm (Adam) implemented in Tensorflow2 to minimize the RMSE between network output \mathbf{o} and target \mathbf{z}

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=0}^T [o(t) - z(t)]^2} \quad (6)$$

where T is the total length of a given trial. The target is task dependent as described below. The learning rate was 0.001, and other TensorFlow default values were used. A discretization step of 10 ms was used for the simulations.

W^{rec} was initialized as a random matrix with full connectivity from a Gamma distribution with a shape parameter of 0.1 and a scale parameter of 1.0, multiplied by a gain factor of 0.5. To start from an approximately balanced regime, the inhibitory weights were multiplied by 4. To respect Dale’s law during training, a rectified linear operation was applied on W^{rec} to clip the weights at zero, and then excitation and inhibition were implemented by multiplying the clipped W^{rec} with a diagonal matrix of 1 and -1 representing excitatory and inhibitory units, respectively. Initial W^{in} was drawn

Downloaded from https://www.science.org at University of California Los Angeles on May 07, 2024

from the same Gamma distribution clipped to zero during training the same way as \mathbf{W}^{rec} . \mathbf{W}^{out} and \mathbf{b}^{out} were initialized at zero.

\mathbf{U} was drawn from a Gaussian distribution with a mean of 0.5 and SD of 0.17 (mean/3) and cut off at 0.001 and 0.99. Unless otherwise specified, τ_x and τ_u were drawn from a Gaussian distribution with a mean of 1 s and SD of 0.33 s (mean/3) and cutoff at 0.1 s and 3 s to ensure numerical stability. α was task specific as described below.

\mathbf{W}^{in} , \mathbf{W}^{rec} , \mathbf{W}^{out} , and \mathbf{b}^{out} were trained. Parameters were updated after each batch of 16 trials. After every 100 batches of training, the network was tested for 20 batches to compute the task performance and mean error. For all the spatial and temporal scaling motor trajectory tasks, the training was considered a success and stopped when the mean error was lower than 0.02; for the IAF and flexible sensorimotor timing tasks, the stop criterion was task performance above 90 or 98%, respectively, to capture the experimental features.

Temporal and spatial scaling motor trajectory tasks

We trained RNNs to generate a series of complex motor trajectories: 10 handwritten digits from 0 to 9 (68). Unless otherwise specified, each of the 10 input signals was presented for 0.1 s with onset time randomly drawn from a uniform distribution (0.2 to 0.6 s) to cue digit identity. Following the input, the network evolved freely for a specific duration to match the corresponding targets warped temporally and/or spatially according to the task requirements.

Temporal scaling task

During training, for the congruent condition, $\alpha = 0.9$ corresponded to the standard target with a duration of 1 s, while $\alpha = 0.8$ trials corresponded to the target with a duration of 1.5 s uniformly interpolated from the standard target. For the incongruent condition, the association between the α and target duration switched, namely, 0.9 and 0.8 corresponded to 1.5 and 1 s, respectively.

Spatial scaling task

During training, for the congruent condition, $\alpha = 0.8$ corresponded to the standard size target with a duration of 1 s, and $\alpha = 0.9$ corresponded to the larger target also with a duration of 1 s but with the amplitude of the output target multiplied by a factor of 1.5, namely, the size of target was 1.5× larger than the standard one. For the incongruent condition, the association between the α and target size was reversed.

In simulations in fig. S3, α values in each trial were randomly selected from Gaussian distributions with SD = 0.0075 and means = 0.9/0.8 for short/long or large/small scales, respectively, in temporal and spatial congruent scaling tasks.

Joint temporal-spatial scaling task

For the joint control of temporal and spatial scales in the same network (Fig. 6), we increased N to 400. Fifty percent of recurrent units were used for temporal scaling and the other 50% for spatial scaling (Fig. 6A).

To examine joint control of shape (digit identity) and temporal and spatial scales in a single network, we randomly divided the 400 recurrent units into three groups: shape (1 to 200), temporal scale (201 to 300), and spatial scale (301 to 400). The temporal and spatial scales were cued the same way as the joint task, but we further divided the shape group into 10 subgroups corresponding to the 10 digits. To cue a given digit, we set the α in the corresponding subgroup to 0.6 while leaving the α in the rest of the shape group at 1 (Fig. 6D).

Generalization performance

To see whether RNNs trained with two α levels can generalize to other temporal or spatial scales, we tested the RNNs with α level in

between (interpolation) and outside (extrapolation) of the trained levels. Specifically, for trained α levels of 0.8 and 0.9, we tested $\alpha = 0.75, 0.775, 0.825, 0.85, 0.875, 0.925,$ and 0.95 . The generalization performance was quantified as the RMSE between the output under the testing α and the target $z_\alpha(t)$ linearly warped to the trained α levels. For instance, the length of $z_\alpha(t)$, T_α for the congruent temporal scaling task with the trained α_1 and α_2 values corresponding to the digit lengths T_1 and T_2 (1, 1.5 s) would be

$$T_\alpha = T_1 + \frac{T_2 - T_1}{\alpha_2 - \alpha_1}(\alpha - \alpha_1)$$

Therefore

$$z_\alpha(t) = z_{\alpha_1}\left(t^* \frac{T_1}{T_\alpha}\right)$$

Similarly, the size of $z_\alpha(t)$, S_α for the congruent spatial scaling task with the trained α_1/α_2 corresponding to the digit size S_1/S_2 (1.5/1) would be

$$S_\alpha = S_2 + \frac{S_1 - S_2}{\alpha_1 - \alpha_2}(\alpha - \alpha_2)$$

Therefore

$$z_\alpha(t) = z_{\alpha_2}(t) * \frac{S_\alpha}{S_2}$$

The incongruent conditions were modified accordingly.

Temporal Scaling Factor, Spatial Scaling Factor, and Scale-Specific Index

To quantify temporal and spatial scaling of the recurrent dynamics, we extended a previously described method (63, 64) to define three measures: TSE, SSF, and SSI. As in Fig. 3E, for two given population trajectories: $\mathbf{r}_{0.9}$ ($N \times T_1$) and $\mathbf{r}_{0.8}$ ($N \times T_2$) with $T_1 < T_2$, the goal of the algorithm is to find the best temporal and spatial scaling factors, by which warping $\mathbf{r}_{0.9}$ gives the best match to $\mathbf{r}_{0.8}$. Specifically, we searched among a range of temporal scaling factors (tsf, 0.5 to 2), and spatial scaling factors (ssf, 0.5 to 2). We then warped $\mathbf{r}_{0.9}$ temporally and spatially as follows

$$\mathbf{r}_{0.9}^{\text{warp}}(t) = \begin{cases} \mathbf{r}_{0.9}\left(\frac{t}{\text{tsf}}\right) * \text{ssf}, & t \leq T_1 * \text{tsf} \\ \text{mean}(\mathbf{r}_{0.9}) * \text{ssf}, & t > T_1 * \text{tsf} \end{cases}$$

where the mean() function is applied to each unit.

To compare with $\mathbf{r}_{0.9}^{\text{warp}}(t)$, we extended the $\mathbf{r}_{0.8}$ dynamics as follows

$$\mathbf{r}_{0.8}'(t) = \begin{cases} \mathbf{r}_{0.8}(t), & t \leq T_2 \\ \text{mean}(\mathbf{r}_{0.8}), & t > T_2 \end{cases}$$

We then obtained the maximal length, T_{max} between T_2 and $T_1 * \text{tsf}$

$$T_{\text{max}} = \max(T_2, T_1 * \text{tsf})$$

We next compute the mean Euclidean distance $d(\text{ssf}, \text{tsf})$ between $\mathbf{r}_{0.9}^{\text{warp}}(t)$ and $\mathbf{r}_{0.8}'(t)$ for each pair of ssf and tsf as

$$d(\text{ssf}, \text{tsf}) = \frac{1}{T_{\text{max}}} \sum_{t=0}^{T_{\text{max}}} \|\mathbf{r}_{0.9}^{\text{warp}}(t) - \mathbf{r}_{0.8}'(t)\|$$

$\|\cdot\|$ denotes Euclidean norm of a given vector. Then the SSF and TSF were defined as the tsf and ssf that gives the minimal $d(\text{ssf}, \text{tsf})$

$$\text{SSF, TSF} = \operatorname{argmin}_{\text{ssf}, \text{tsf}} [d(\text{ssf}, \text{tsf})]$$

Last, we defined the SSI as

$$\text{SSI} = \frac{d(\text{SSF}, \text{TSF})}{\frac{1}{T_{\max}} \sum_{t=0}^{T_{\max}} \|\operatorname{mean}(\mathbf{r}_{0.8}') - \mathbf{r}_{0.8}'(t)\|}$$

Intuitively, SSI provides a measure of how well the relation from $\mathbf{r}_{0.9}$ to $\mathbf{r}_{0.8}$ can be explained by temporal and spatial scaling profiles, namely, the smaller the SSI, the better $\mathbf{r}_{0.8}$ can be fitted by warping $\mathbf{r}_{0.9}$ temporally and spatially—e.g., SSI = 0 indicates that warping of $\mathbf{r}_{0.9}$ perfectly fits $\mathbf{r}_{0.8}$, and SSI > 1 indicates that the fit is worse than simply using the mean of $\mathbf{r}_{0.8}$.

Trajectory decomposition analysis

To understand the transitions between trajectories at different α levels, we started from the state trajectories (Eq. 1) at $\alpha = 0.9$ ($\mathbf{s}_{0.9}$) and $\alpha = 0.8$ ($\mathbf{s}_{0.8}$) (fig. S7A). For the spatial scaling task, $\mathbf{s}_{0.9}$ and $\mathbf{s}_{0.8}$ naturally have the same duration, while for the temporal scaling task, we uniformly subsampled the trajectory with the longer duration to the same as the short trajectory. For a given time point on $\mathbf{s}_{0.8}$ and its corresponding time point on $\mathbf{s}_{0.9}$ with direction \mathbf{p} from $\mathbf{s}_{0.8}$ to $\mathbf{s}_{0.9}$, there are velocity vectors $\mathbf{v}_{0.8}$ and $\mathbf{v}_{0.9}$, respectively. In general, $\mathbf{v}_{0.8}$ can be decomposed into the recurrent component, $\mathbf{rec}_{0.8}$, and decay component $\mathbf{d}_{0.8}$. We then sought to compute the angle between $\mathbf{rec}_{0.8}$ and \mathbf{p} or the angle between $\mathbf{rec}_{0.8}$ and $\mathbf{v}_{0.9}$ at each corresponding time point on $\mathbf{s}_{0.9}$ and $\mathbf{s}_{0.8}$. Last, mean angle across time was obtained for comparison.

Subspace angle analysis

In fig. S7 (D to F), we computed the angle between the subspace of the recurrent dynamics at different α levels with the subspace of the output. Specifically, for a given trajectory \mathbf{r} , we performed the PCA, then the recurrent space was expanded by the first n PCs. The output space was expanded by the learned output weights which led to a two-dimensional space. Last, the angle between the recurrent space and output space is computed by the MATLAB function `subspace()` between these two spaces. Higher dimensions of recurrent space expanded by more principle components were also tested.

RNNs without STP

To study whether STP affects the training and generalization in the temporal or spatial scaling tasks (Fig. 5), we modified the standard congruent temporal or spatial scaling task by removing the STP dynamics during training and testing. Specifically, we trained and tested RNNs with $\mathbf{x} = \mathbf{1}$ (Eq. 2) and $\mathbf{u} = \alpha\mathbf{U}$ (Eq. 3) during the whole trial, while other variables and parameters remained the same.

RNNs with scale cued by a speed or size input

To compare the neuromodulation of STP strategy to using the input amplitude to cue different scales (Fig. 4): (i) we removed STP by fixing the variable \mathbf{u} at $0.85*\mathbf{U}$ and $\mathbf{x} = \mathbf{1}$ across whole trials; and (ii) added a tonic input continuously presented across whole trials, the amplitude of which, cued either the length of the trials in the temporal scaling task or the size of the digit in the spatial scaling task. Specifically, 0.9/0.85, 0.9/0.8, 0.9/0.7, 0.9/0.5, 0.9/0.3, or 0.9/0.1 corresponded to

either 1/1.5 s or 1.5×/1× size, respectively. Generalization performance was tested similarly to the α -cued scale models. To study the intrinsic generalization of the input amplitude-cued model, we also trained the RNNs with a single-input amplitude level or single α value and tested them with different novel levels.

Simulations of the experimental sensorimotor timing task IAFC task

Task parameters were the same as the experimental conditions for rats in the IAFC task (9). RNNs were composed of one input for delivering two stimuli lasting 150 ms with a range of intervals, 0.6, 1.05, 1.26, 1.38, 1.62, 1.74, 1.95, and 2.4 s (short if interval < 1.5 s and long if > 1.5 s), and two outputs corresponding to short or long choices, respectively. The target of the correct output was set to 1 for a response period of 200 ms right after the second-stimulus offset. The decision was made on the basis of the mean activity during the response period for the two outputs in a winner-take-all manner, and performance was defined as the percentage of the correct trials. RNNs were trained with $\alpha = 0.8$ for all units and tested with 0.9 and 0.7 to simulate the effect of optogenetic inactivation or activation of dopamine activity, respectively. We set σ (in Eq. 1) to 1 to match the noise level of the experiments. Same as the experiment, we fitted the long choice probability by a sigmoid function

$$P = A \frac{\exp\left(\frac{y - \text{bias}}{\text{slope}}\right)}{1 + \exp\left(\frac{y - \text{bias}}{\text{slope}}\right)} + \text{offset}$$

where y represents the input intervals; P represents the probability of choice of long interval. A , bias , slope , and offset are the four fitting variables.

Flexible sensorimotor timing task

For the simulation of the Ready-Set-Go task (1), one input to the RNN delivered two stimuli lasting 100 ms with the intervals from a pool of 7 uniformly spaced intervals in the range of 0.5 to 1 s (sensory time, t_s). On the basis of the context cued by $\alpha = 0.9$ or $\alpha = 0.8$, the output unit should generate a linear ramp (0 to 1) crossing threshold (0.75) at 1× or 1.5× the duration of sensory interval (the target time t_t) beginning at the offset of the second stimulus (production time, t_p). As in the experimental study, we defined a trial with t_p as correct if the error = $|t_p - t_t|$ was smaller than $0.2*t_t + 0.025$ s. Again, performance was defined as the percentage of the correct trials.

Statistical test

Unless otherwise stated, all the statistical tests were nonparametric and performed in MATLAB (MathWorks); the specific tests are stated in the figure captions; the data were presented in with Boxplot: central lines, median; bottom and top edges, lower and upper quartiles; whiskers, extremes; red cross, outliers.

Supplementary Materials

This PDF file includes:

Figs. S1 to S9

REFERENCES AND NOTES

1. E. D. Remington, D. Narain, E. A. Hosseini, M. Jazayeri, Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019.e5 (2018).
2. N. F. Hardy, V. Goudar, J. L. Romero-Sosa, D. V. Buonomano, A model of temporal scaling correctly predicts that motor timing improves with speed. *Nat. Commun.* **9**, 4732 (2018).

3. G. M. Cicchini, R. Arrighi, L. Cecchetti, M. Giusti, D. C. Burr, Optimal encoding of interval timing in expert percussionists. *J. Neurosci.* **32**, 1056–1060 (2012).
4. G. B. M. Mello, S. Soares, J. J. Paton, A scalable population code for time in the striatum. *Curr. Biol.* **9**, 1113–1122 (2015).
5. D. A. Rosenbaum, *Human Motor Control* (Elsevier Inc, Amsterdam ; Boston, MA, ed. 2, 2010), pp. xvi, 505 p.
6. N. K. Harpaz, T. Flash, I. Dinstein, Scale-invariant movement encoding in the human motor system. *Neuron* **81**, 452–462 (2014).
7. C. V. Buhusi, W. H. Meck, Differential effects of methamphetamine and haloperidol on the control of an internal clock. *Behav. Neurosci.* **116**, 291–297 (2002).
8. M. R. Drew, S. Fairhurst, C. Malapani, J. C. Horvitz, P. D. Balsam, Effects of dopamine antagonists on the timing of two intervals. *Pharmacol. Biochem. Behav.* **75**, 9–15 (2003).
9. S. Soares, B. V. Atallah, J. J. Paton, Midbrain dopamine neurons control judgment of time. *Science* **354**, 1273–1277 (2016).
10. M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, K. V. Shenoy, Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
11. J. P. Stroud, M. A. Porter, G. Hennequin, T. P. Vogels, Motor primitives in space and time via targeted gain modulation in cortical networks. *Nat. Neurosci.* **21**, 1774–1783 (2018).
12. G. Hennequin, T. P. Vogels, W. Gerstner, Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* **82**, 1394–1406 (2014).
13. S. Saxena, A. A. Russo, J. Cunningham, M. M. Churchland, Motor cortex activity across movement speeds is predicted by network-level strategies for generating muscle activity. *eLife* **11**, e67620 (2022).
14. H. Merchant, O. Pérez, R. Bartolo, J. C. Méndez, G. Mendoza, J. Gámez, K. Yc, L. Prado, Sensorimotor neural dynamics during isochronous tapping in the medial premotor cortex of the macaque. *Eur. J. Neurosci.* **41**, 586–602 (2015).
15. D. A. Crowe, W. Zarco, R. Bartolo, H. Merchant, Dynamic representation of the temporal and sequential structure of rhythmic movements in the primate medial premotor cortex. *J. Neurosci.* **34**, 11972–11983 (2014).
16. S. Vyas, M. D. Golub, D. Sussillo, K. V. Shenoy, Computation through neural population dynamics. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
17. L. Loggiaco, L. F. Abbott, S. Escola, Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell Rep.* **35**, 109090 (2021).
18. J. M. Murray, Local online learning in recurrent networks with random feedback. *eLife* **8**, e43299 (2019).
19. J. Wang, D. Narain, E. A. Hosseini, M. Jazayeri, Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
20. M. Beiran, N. Meirhaeghe, H. Sohn, M. Jazayeri, S. Ostojic, Parametric control of flexible timing through low-dimensional neural manifolds. *Neuron* **111**, 739–753.e8 (2023).
21. H. Lindén, P. C. Petersen, M. Vestergaard, R. W. Berg, Movement is governed by rotational neural dynamics in spinal motor networks. *Nature* **610**, 526–531 (2022).
22. R. S. Zucker, W. G. Regehr, Short-term synaptic plasticity. *Annu. Rev. Physiol.* **64**, 355–405 (2002).
23. L. F. Abbott, W. G. Regehr, Synaptic computation. *Nature* **431**, 796–803 (2004).
24. H. Motanis, M. J. Seay, D. V. Buonomano, Short-term synaptic plasticity as a mechanism for sensory timing. *Trends Neurosci.* **41**, 701–711 (2018).
25. J. K. Seamans, D. Durstewitz, B. R. Christie, C. F. Stevens, T. J. Sejnowski, Dopamine D1/D5 receptor modulation of excitatory synaptic inputs to layer V prefrontal cortex neurons. *Proc. Natl. Acad. Sci. U.S.A.* **98**, 301–306 (2001).
26. S. Kroener, L. J. Chandler, P. E. M. Phillips, J. K. Seamans, Dopamine modulates persistent synaptic activity and enhances the signal-to-noise ratio in the prefrontal cortex. *PLoS ONE* **4**, e6507 (2009).
27. F. Tecuapetla, L. Carrillo-Reid, J. Vargas, E. Galarraga, Dopaminergic modulation of short-term synaptic plasticity at striatal inhibitory synapses. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 10258–10263 (2007).
28. J. M. Leyrer-Jackson, M. P. Thomas, Layer-specific effects of dopaminergic D1 receptor activation on excitatory synaptic trains in layer V mouse prefrontal cortical pyramidal cells. *Physiol. Rep.* **6**, e13806 (2018).
29. N. X. Tritsch, B. L. Sabatini, Dopaminergic modulation of synaptic transmission in cortex and striatum. *Neuron* **76**, 33–50 (2012).
30. C. Q. Chiu, N. Puente, P. Grandes, P. E. Castillo, Dopaminergic modulation of endocannabinoid-mediated plasticity at GABAergic synapses in the prefrontal cortex. *J. Neurosci.* **30**, 7236–7248 (2010).
31. W. J. Gao, L. S. Krimer, P. S. Goldman-Rakic, Presynaptic regulation of recurrent excitation by D1 receptors in prefrontal circuits. *Proc. Natl. Acad. Sci. U.S.A.* **98**, 295–300 (2001).
32. W. J. Gao, Y. Wang, P. S. Goldman-Rakic, Dopamine modulation of perisomatic and peridendritic inhibition in prefrontal cortex. *J. Neurosci.* **23**, 1622–1630 (2003).
33. J. K. Seamans, N. Gorelova, D. Durstewitz, C. R. Yang, Bidirectional dopamine modulation of GABAergic inhibition in prefrontal cortical pyramidal neurons. *J. Neurosci.* **21**, 3628–3638 (2001).
34. M. M. Hoskison, J. A. Connor, C. W. Shuttleworth, GABA(B)-receptor modulation of short-term synaptic depression at an excitatory input to murine hippocampal CA3 pyramidal neurons. *Neurosci. Lett.* **365**, 48–53 (2004).
35. K. J. Burke Jr., C. M. Keeshen, K. J. Bender, Two forms of synaptic depression produced by differential neuromodulation of presynaptic calcium channels. *Neuron* **99**, 969–984.e7 (2018).
36. J. Urban-Ciecko, E. E. Fanselow, A. L. Barth, Neocortical somatostatin neurons reversibly silence excitatory transmission via GABA_B receptors. *Curr. Biol.* **25**, 722–731 (2015).
37. D. V. Buonomano, M. M. Merzenich, Net interaction between different forms of short-term synaptic plasticity and slow-IPSPs in the hippocampus and auditory cortex. *J. Neurophysiol.* **80**, 1765–1774 (1998).
38. N. Y. Masse, G. R. Yang, H. F. Song, X. J. Wang, D. J. Freedman, Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nat. Neurosci.* **22**, 1159–1167 (2019).
39. J. M. Murray, G. S. Escola, Learning multiple variable-speed sequences in striatum via cortical tutoring. *eLife* **6**, e26084 (2017).
40. D. V. Buonomano, M. M. Merzenich, Temporal information transformed into a spatial code by a neural network with realistic properties. *Science* **267**, 1028–1030 (1995).
41. G. Mongillo, O. Barak, M. Tsodyks, Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).
42. R. Laje, D. V. Buonomano, Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* **16**, 925–933 (2013).
43. V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
44. D. Sussillo, L. F. Abbott, Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544–557 (2009).
45. W. Chaisangmongkon, S. K. Swaminathan, D. J. Freedman, X.-J. Wang, Computing by robust transience: How the fronto-parietal network performs sequential, category-based decisions. *Neuron* **93**, 1504–1517.e1504 (2017).
46. J. D. Murray, A. Bernacchia, N. A. Roy, C. Constantinidis, R. Romo, X.-J. Wang, Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 394–399 (2017).
47. O. Barak, M. Tsodyks, Working models of working memory. *Curr. Opin. Neurobiol.* **25**, 20–24 (2014).
48. H. Markram, Y. Wang, M. Tsodyks, Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Natl. Acad. Sci. U.S.A.* **95**, 5323–5328 (1998).
49. F. Nadim, D. Bucher, Neuromodulation of neurons and synapses. *Curr. Opin. Neurobiol.* **29**, 48–56 (2014).
50. A. M. Rush, J. Kilbride, M. J. Rowan, R. Anwyl, Presynaptic group III mGluR modulation of short-term plasticity in the lateral perforant path of the dentate gyrus in vitro. *Brain Res.* **952**, 38–43 (2002).
51. D. V. Baimoukhametova, S. A. Hewitt, C. A. Sank, J. S. Bains, Dopamine modulates use-dependent plasticity of inhibitory synapses. *J. Neurosci.* **24**, 5162–5171 (2004).
52. C. Gonzalez-Islas, J. J. Hablitz, Dopamine enhances EPSCs in layer II–III pyramidal neurons in rat prefrontal cortex. *J. Neurosci.* **23**, 867–875 (2003).
53. W. Schultz, P. Dayan, P. R. Montague, A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
54. W. H. Meck, Neuropharmacology of timing and time perception. *Cog. Brain Res.* **3**, 227–242 (1996).
55. A. V. Maricq, R. M. Church, The differential effects of haloperidol and methamphetamine on time estimation in the rat. *Psychopharmacology (Berl)* **79**, 10–15 (1983).
56. B. J. Fung, E. Sutilif, M. G. Hussain Shuler, Dopamine and the interdependency of time perception and reward. *Neurosci. Biobehav. Rev.* **125**, 380–391 (2021).
57. T. H. Rammesayer, Neuropharmacological evidence for different timing mechanisms in humans. *Q. J. Exp. Psychol. B* **52**, 273–286 (1999).
58. Y.-C. Kim, S.-W. Han, S. L. Alberico, R. N. Ruggiero, B. De Corte, K.-H. Chen, N. S. Narayanan, Optogenetic stimulation of frontal D1 neurons compensates for impaired temporal control of action in dopamine-depleted mice. *Curr. Biol.* **27**, 39–47 (2017).
59. P. Simen, M. Matell, Why does time seem to fly when we're having fun? *Science* **354**, 1231–1232 (2016).
60. J. I. Lake, W. H. Meck, Differential effects of amphetamine and haloperidol on temporal reproduction: Dopaminergic regulation of attention and clock speed. *Neuropsychologia* **51**, 284–292 (2013).
61. K. He, M. Huertas, S. Z. Hong, X. Tie, J. W. Hell, H. Shouval, A. Kirkwood, Distinct eligibility traces for LTP and LTD in cortical synapses. *Neuron* **88**, 528–538 (2015).
62. B. Panigrahi, K. A. Martin, Y. Li, A. R. Graves, A. Vollmer, L. Olson, B. D. Mensh, A. Y. Karpova, J. T. Dudman, Dopamine is required for the neural representation and control of movement vigor. *Cell* **162**, 1418–1430 (2015).
63. S. Zhou, S. C. Masmanidis, D. V. Buonomano, Encoding time in neural dynamic regimes with distinct computational tradeoffs. *PLoS Comput. Biol.* **18**, e1009271 (2022).
64. S. Zhou, S. C. Masmanidis, D. V. Buonomano, Neural sequences as an optimal dynamical regime for the readout of time. *Neuron* **108**, 651–658.e5 (2020).
65. A. Shimbo, E.-I. Izawa, S. Fujisawa, Scalable representation of time in the hippocampus. *Sci. Adv.* **7**, eabd7013 (2021).
66. Y. K. Wu, F. Zenke, Nonlinear transient amplification in recurrent neural networks with short-term plasticity. *eLife* **10**, e71263 (2021).

67. R. Kim, Y. Li, T. J. Sejnowski, Simple framework for constructing functional spiking recurrent neural networks. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 22811–22820 (2019).
68. V. Goudar, D. V. Buonomano, Encoding sensory and motor patterns as time-invariant trajectories in recurrent neural networks. *eLife* **7**, e31134 (2018).

Acknowledgments: We thank S. Soldado-Magraner for reading the paper and helpful discussions. **Funding:** S.Z. was supported by Fudan University Strategic Talent Start-up Funds for Academic Research JIH2641028Y, and D.V.B. was supported by National Science Foundation grant RI-2008741. **Author contributions:** Conceptualization: S.Z. and D.V.B. Methodology: S.Z. and D.V.B. Simulations: S.Z. Analysis: S.Z. and D.V.B. Supervision: D.V.B. Writing—original draft: S.Z. and D.V.B. Writing—review and editing: S.Z. and D.V.B. **Competing interests:** The authors

declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Codes used for the simulations in this paper are available at <https://zenodo.org/records/10405269> (DOI: 10.5281/zenodo.10405269) and https://github.com/ShanglinZhou/Temporal_Spatial_Scale_STP.

Submitted 11 September 2023

Accepted 3 April 2024

Published 3 May 2024

10.1126/sciadv.adk7257