nature human behaviour

Article

https://doi.org/10.1038/s41562-023-01592-y

Multiplexing working memory and time in the trajectories of neural networks

Received: 8 November 2022

Accepted: 22 March 2023

Published online: 20 April 2023

Check for updates

Shanglin Zhou[1,9], Michael Seay[2,9], Jiannis Taxidis[3,4], Peyman Golshani[5,6,7,8] & Dean V. Buonomano[1,2,5]

Working memory (WM) and timing are generally considered distinct cognitive functions, but similar neural signatures have been implicated in both. To explore the hypothesis that WM and timing may rely on shared neural mechanisms, we used psychophysical tasks that contained either task-irrelevant timing or WM components. In both cases, the task-irrelevant component influenced performance. We then developed recurrent neural network (RNN) simulations that revealed that cue-specific neural sequences, which multiplexed WM and time, emerged as the dominant regime that captured the behavioural findings. During training, RNN dynamics transitioned from low-dimensional ramps to high-dimensional neural sequences, and depending on task requirements, steady-state or ramping activity was also observed. Analysis of RNN structure revealed that neural sequences relied primarily on inhibitory connections, and could survive the deletion of all excitatory-to-excitatory connections. Our results indicate that in some instances WM is encoded in time-varying neural activity because of the importance of predicting when WM will be used.

Working memory (WM) refers to the ability to transiently store information, and subsequently use this information in a flexible manner for goal-oriented behaviours and decision-making[1,2]. Timing, here, refers to the ability to track elapsed time after a stimulus, to anticipate subsequent external events or generate appropriately timed motor responses[3–5]. Whereas it is widely recognized that the ability to transiently store information about the past and prospectively anticipate external events are among the most fundamental computations the brain performs[1–4,6–8], the fields of WM and timing have evolved mostly independently from each other because they have been seen as distinct cognitive functions with different underlying neural mechanisms. Yet, both share critical computational features. Both require transiently storing information, retrospective information in the case of WM and prospective information in the case of timing (for example, when a delayed reward will occur). In some cases, these properties are mirror images of each other. For example, a timer, such as a running hourglass, can be seen as encoding a transient memory that it was recently flipped over and of generating a prediction as to when an external event may occur.

Similar neural signatures—including ramping activity and neural sequences—have been associated with both WM and the encoding of time[3,5,9–14]. Although early groundbreaking studies suggested that WM is encoded in steady-state persistent neural activity[15–17], there is ongoing controversy regarding the neural encoding of WM[9,10,18]. Broadly speaking, in addition to steady-state persistent activity there are two additional broad classes of WM models[9,10,19]: (1) time-varying patterns of neural population activity, which can include low-dimensional ramping activity as well as high-dimensional neural trajectories (including, neural sequences) and (2) activity-silent mechanisms, in which short-term memory can be stored in the hidden state of neural networks—rather

[1]Department of Neurobiology, David Geffen School of Medicine, University of California, Los Angeles, CA, USA. [2]Department of Psychology, University of California, Los Angeles, CA, USA. [3]Program in Neurosciences and Mental Health, Hospital for Sick Children, Toronto, Ontario, Canada. [4]Department of Physiology, University of Toronto, Toronto, Ontario, Canada. [5]Department of Neurology, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA, USA. [6]Integrative Center for Learning and Memory, Brain Research Institute, University of California, Los Angeles, Los Angeles, CA, USA. [7]UCLA Semel Institute for Neuroscience and Behavioral Sciences, University of California, Los Angeles, Los Angeles, CA, USA. [8]West Los Angeles VA Medical Center, Los Angeles, CA, USA. [9]These authors contributed equally: Shanglin Zhou, Michael Seay. e-mail: dbuono@ucla.edu

than ongoing spiking activity—through mechanisms such as short-term synaptic plasticity (STSP). Ramping activity, neural sequences and STSP-based changes in the hidden state of networks have all been proposed to underlie timing as well[3–5,20].

The diversity of neural regimes implicated in WM may, in part, be dependent on the presence or absence of implicit timing components. The brain is always attempting to learn the temporal structure of the external world even if it is not explicitly relevant to the task at hand[7,21]. Implicit timing enables prediction of when events will take place, thus allowing for preparation and optimal allocation of cognitive resources. Indeed, previous studies have suggested a link between WM and timing[22,23], and recent human studies have demonstrated that WM can be impaired when information has to be retrieved at unexpected times[7,24]. Furthermore, some computational studies have implicitly linked the ability to encode elapsed time in a stimulus-specific manner[11,14,25,26].

We examine the hypothesis that WM and timing are, in some cases, essentially the same computation, that is, a given stimulus and stimulus-specific elapsed time can be encoded in the same dynamic pattern of neural activity. We first developed two psychophysical tasks that use the same stimulus structure but vary whether the WM or timing components are explicit (required to solve the task) or implicit (task irrelevant). Participants learned task-irrelevant WM information during an explicit timing task, and task-irrelevant timing information during an explicit WM task. Given the ongoing challenges in identifying brain regions causally responsible for both the encoding of time and WM, and the success of using artificial neural networks to examine the neural dynamic regimes underlying a diverse set of cortical computations[27], we trained recurrent neural networks (RNNs) on the same tasks the human participants performed. We show that cue-specific neural sequences emerge as the dominant regime for encoding memoranda and elapsed time from the onset of each memorandum, and that overall, training stages, task structure and hyperparameters captured much of the diversity of the experimentally observed neural dynamic regimes.

Results

The differential-delay-match-to-sample task

As a first step towards addressing a potential link between WM and timing, we developed variants of the standard delay-match-to-sample (DMS) WM task. In its simplest form, a DMS task presents either of two cues (in our case, a star or circle denoted by *A* or *B*, respectively), and following a delay period either of the two stimuli is presented again, resulting in four conditions (*AA*, *AB*, *BA*, *BB*). Participants are required to differentially respond to the match (*AA*, *BB*) versus non-match (*AB*, *BA*) conditions. Typically, the delay between the cue and probe is fixed or randomized, but in our differential-delay-match-to-sample (dDMS) task the cues predicted the delay duration (Fig. 1a). For example, the *AA* and *AB* trials might be associated with a 1 s delay, and *BA* and *BB* trials with a 2.2 s delay—but the delay itself is task irrelevant. To determine whether participants implicitly learned the temporal structure of the task and whether unexpected delays altered WM performance, the cue-delay contingency was reversed in 20% of the trials (Methods). The second task (Fig. 1a, right), termed interval-stimulus-association task (ISA), was based on the same exact stimulus structure as the dDMS task but framed differently: participants were explicitly instructed to press one key when there was a short delay followed by *A* (short-*A*) or a long delay followed by *B* (long-*B*) and another key after long-*A* or short-*B* trials. In the ISA task the interval (delay) is explicitly relevant, but the cue (the first stimulus) identity is irrelevant as it just serves as an indicator of $t = 0$ for the interval. Note that during standard trials, the dDMS and ISA tasks are isomorphic—that is, the correct responses could be produced with either strategy—the difference between the tasks lies in the reverse trials (Fig. 1a).

To determine whether participants implicitly learned the cue-delay associations we analysed the inverse efficiency (reaction time (RT)/accuracy), a measure designed to take into account between-participant differences in speed-accuracy trade-offs[28]. In the dDMS task (Fig. 1b) there was a main effect between standard and reverse trials (two-way ANOVA: $F_{1,26} = 9.05$, $P = 0.006$, effect size $\eta^2 = 0.071$), indicating that the violation of temporal expectation in the reverse trials altered performance. There was also a main effect of the actual delay as expected from the well-known hazard rate effect (two-way ANOVA, $F_{1,26} = 14$, $P = 0.0009$, effect size $\eta^2 = 0.098$) (ref. 29)—after the short interval had elapsed there was increased certainty that the probe will appear at the long delay thus decreasing RT. We also examined the raw RT and trial accuracy independently (Supplementary Fig. 1a,b), both exhibited a main effect of reversal (two-way ANOVA, RT $F_{1,26} = 7.41$, $P = 0.011$, effect size $\eta^2 = 0.043$; accuracy $F_{1,26} = 13.4$, $P = 0.001$, effect size $\eta^2 = 0.088$). To further validate the results of this new task we performed a replication study (Supplementary Fig. 1e–j), which confirmed a significant standard–reverse effect in inverse efficiency (two-way ANOVA, $F_{1,38} = 8.51$, $P = 0.006$, effect size $\eta^2 = 0.031$) and RT (two-way ANOVA, $F_{1,38} = 9.02$, $P = 0.005$, effect size $\eta^2 = 0.025$). There was no main effect of accuracy but there was an interaction between standard and reverse and the actual delay (two-way ANOVA, $F_{1,38} = 4.1$, $P = 0.05$, effect size $\eta^2 = 0.025$). These results establish that participants implicitly learned the task-irrelevant cue-delay association during a WM task, and that reversing the standard temporal contingency affected WM performance.

We next performed separate experiments using the explicit timing ISA task (Fig. 1c). Again, there was a significant main effect of reversal (two-way ANOVA, $F_{1,21} = 11$, $P = 0.003$, effect size $\eta^2 = 0.108$) on inverse efficiency, as well as on RT (two-way ANOVA, $F_{1,21} = 9.2$, $P = 0.006$, effect size $\eta^2 = 0.031$, Supplementary Fig. 1c) and accuracy (two-way ANOVA, $F_{1,21} = 12.9$, $P = 0.002$, effect size $\eta^2 = 0.141$, Supplementary Fig. 1d). A replication study (Supplementary Fig. 1k–p), further confirmed a significant main effect of standard–reverse trials on inverse efficiency (two-way ANOVA, $F_{1,24} = 8.81$, $P = 0.007$, effect size $\eta^2 = 0.093$), RT (two-way ANOVA, $F_{1,24} = 9.31$, $P = 0.006$, effect size $\eta^2 = 0.045$) and accuracy (two-way ANOVA, $F_{1,24} = 7.78$, $P = 0.01$, effect size $\eta^2 = 0.102$). These results establish that reversing the cue-delay contingency impairs performance on an explicit timing task in which the cue is task irrelevant, and thus that participants are, in effect, implicitly storing the cue in WM during a timing task.

Neural sequences in RNNs encoding WM and time

A large body of neurophysiological data across brain areas has revealed a multitude of neural signatures during WM and timing tasks, including neural sequences[12,30–35], and firing rate ramps[36–39]. Artificial neural networks, and RNNs in particular, have been invaluable in capturing the experimentally observed dynamics and explaining the dynamic regimes capable of storing WM and encoding time[11,40,41], but to date, with some exceptions[11], these attempts have primarily focused on either WM or timing tasks. Thus, anchored by our dDMS task, we next examined which dynamic regimes emerge in RNNs trained to encode both time and WM (Fig. 2). Having established that humans trained on the dDMS task implicitly learn its temporal structure, the RNNs were trained on a timing + WM (T + WM) task, in which the RNN had to learn both the WM and temporal expectation components. RNNs were also trained on two control tasks: a pure WM task without any timing requirements (WM task), and the ISA task that required the RNN to explicitly learn the cue-delay associations but not the match/non-match-to-sample component. Note that these tasks do not perfectly parallel the psychophysical studies because the standard distinction between explicit and implicit learning used in the animal literature does not apply to simple RNN models. Animals or humans, for example, are always attempting to learn the temporal structure of tasks (implicit timing), presumably because predicting the onset of upcoming stimuli optimizes attentional resources even if the temporal structure is irrelevant to the task. This is not the case with RNNs. Thus, to provide an approximation of the explicit or implicit distinction

https://doi.org/10.1038/s41562-023-01592-y

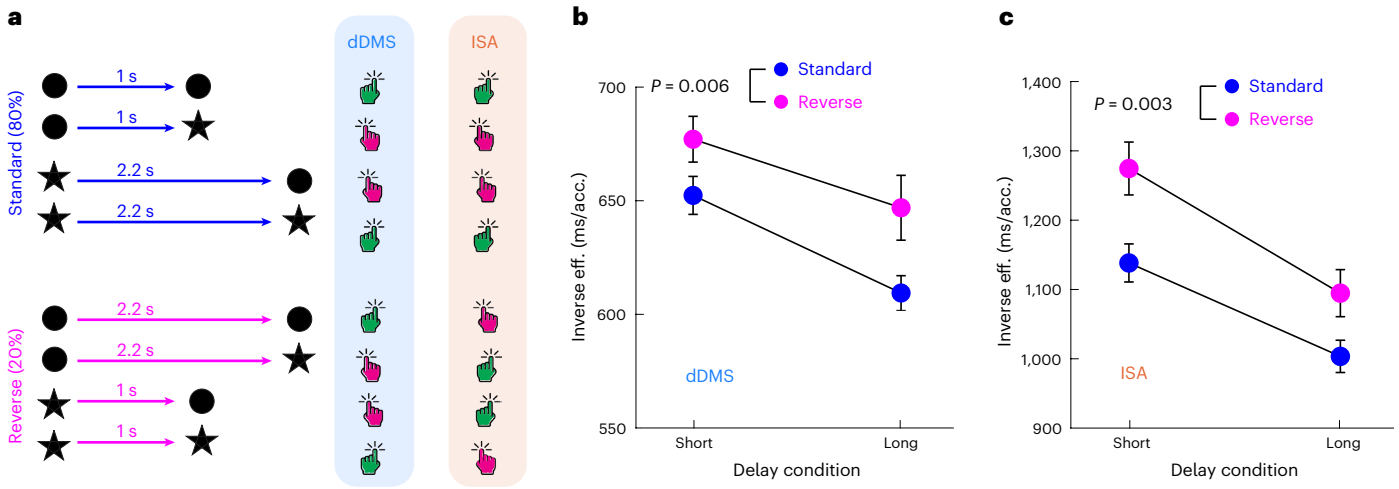

Fig. 1 | Humans implicitly learn the timing component of a WM task and the WM component of a timing task. a, Schematic of the dDMS task and the explicit timing ISA task. Note that the response patterns for the dDMS and ISA tasks only differ during the reverse trials. **b**, Inverse efficiency (eff.) (RT/accuracy (acc.)) of human participants on the dDMS task across the standard (cyan) and reverse (orange) trials. The short and long delays reflect the duration of the actual delay epochs (for example, a long delay on a standard trial is an 'expected' delay, and a long delay on a reverse trial corresponds to an 'unexpected' delay). There was a significant main effect of standard versus reverse conditions ($n = 27$ participants, two-way ANOVA: $F_{1,26} = 9.05$, $P = 0.006$, effect size $\eta^2 = 0.071$). **c**, Inverse efficiency in the ISA task across standard and reverse trials. There was a significant main effect of standard versus reverse conditions ($n = 22$ participants, two-way ANOVA, $F_{1,21} = 11$, $P = 0.003$, effect size $\eta^2 = 0.108$). Data are presented as mean values ± s.e.m. ms is for millisecond as in RT.

in the model, the cost weighting of the WM component was higher than that of the temporal component—that is, in the T + WM training prioritized learning of the WM component (Methods).

RNNs were composed of 256 units and had either one (WM, ISA) or two (T + WM) output units. The first output represented the motor response (for example, non-match detection), and the second output represented temporal expectation (implemented as a half ramp based on anticipatory licking data[33]). The network was composed of three weight matrices: W^{In}, the input to the RNN; W^{RNN}, the recurrent weights and W^{Out}, the connections from the RNN to the output. A number of steps were taken to enhance biological realism and improve our ability to dissect the mechanisms underlying the observed network dynamics (below): (1) Dale's law was implemented; (2) to capture the low spontaneous activity rates of most cortical neurons a rectified linear unit (ReLU) activation function with a bias of zero was used and (3) to focus our mechanistic analyses on the structure of W^{RNN}, biases of the RNN units and W^{In} were not trained (Methods).

The dynamics of the RNNs during the delay period were dramatically different between tasks. RNNs trained on the WM task primarily converged to persistent fixed-point activity during the delay (Fig. 2a,b), in which individual units exhibited cue-specific constant levels of activity during the delay epoch (Fig. 2c). RNNs trained on the T + WM exhibited dynamic activity during the delay that when sorted according to latency resembled neural sequences (Fig. 2d,e). Individual units in these RNNs often exhibited Gaussian-like time fields, in response to one cue and the absence of a response or a different time field in response to the second cue (Fig. 2f). The dynamics in the RNNs trained on the ISA task (Fig. 2g–i), was more mixed. Specifically, in the example shown in Fig. 2i–h both cues A and B triggered the same neural sequence ('erasing' WM information), followed by persistent activity after the initial 1 s period (corresponding to the short delay). This is an effective solution to the ISA task because a categorical encoding of short versus long intervals is sufficient to solve the task.

We compared the performance and dynamics of 17 RNNs trained on the three tasks. Performance was measured by the correct response of the motor unit. As expected, because the RNNs are trained on both the standard and reverse trials, performance was close to 100% on both conditions for all three tasks (Fig. 2j–k). To compare the dynamics of the RNNs across tasks we first quantified the effective dimensionality[14] and the sequentiality index[33] across the delay periods (Methods). The dimensionality was significantly higher in the T + WM task compared to both the WM and ISA task (Wilcoxon rank sum test, $n = 17$, $P = 3.58 \times 10^{-7}$ and $P = 4.44 \times 10^{-7}$, respectively), and there was a much smaller difference between the dimensionality for the WM and ISA tasks (Wilcoxon rank sum test, $n = 17$, $P = 1.9 \times 10^{-4}$). The selectivity index was also higher in the T + WM task compared to the WM and ISA tasks (Wilcoxon rank sum test, $P = 7.05 \times 10^{-7}$ and 7.46×10^{-5}, respectively), and between the ISA and WM tasks (Wilcoxon rank sum test, $P = 1.69 \times 10^{-6}$). These results are consistent with the interpretation that RNNs converge to fixed-point attractors when they only need to encode WM, but to neural sequences when they need to encode both WM and elapsed time, and to mixed dynamics when they need to encode elapsed time and the nature of the first cue is task irrelevant (ISA task). We note, however, that there is some variability in the solutions RNNs converged to in each task, particularly during the WM and ISA task. Specifically, sequences could emerge during the WM tasks, while ramping activity and mixed dynamics could emerge in the ISA task (Supplementary Fig. 2).

The T + WM task was designed to capture the human psychophysics data and, critically, in this task two distinct neural sequences generally encoded both WM and timing. There is no clear a priori reason that high-dimensional trajectories that approximate neural sequences, should emerge as the dominant solution to encode WM and time. Indeed, one might expect much lower dimensional cue-specific ramping activity to encode both WM and time (below and Discussion).

The transition from ramps to neural sequences over training

As stated above, either low-dimensional ramping activity or high-dimensional neural sequences can encode both WM and time. Furthermore, both types of dynamics have been observed experimentally during WM or timing tasks[12,30–34,36–39]. So far, ramping activity and neural sequences have been treated as fundamentally different dynamic regimes to encode WM or time[3,9,11,42]. To determine whether this is indeed the case, we analysed the development of neural sequence across training in the T + WM task.

Visualization of RNN dynamics from early to late training stages revealed a continuous shift from steady-state activity, to ramps, to

neural sequences, both at the population and single-unit level (Fig. 3a). Quantified across RNNs this transition was expressed as a progressive increase in the dimensionality of network dynamics (Fig. 3b). As expected—because the WM task is the 'explicit task' (higher cost weighting)—the WM performance (that is, discrimination of match versus non-match trials) peaked very early in training while the dimensionality was fairly low. At these early stages, timing, as measured by the loss function of the timing output (Fig. 3c) or by the ability to decode elapsed time from each cue (Fig. 3d), was poor but increased progressively over the course of training. Finally, there was a strong positive correlation between dimensionality and time decoding performance (Fig. 3e).

The transition from ramps to neural sequences was driven by the 'implicit' timing component of the task, because WM performance was high early in training. These results indicate that the same RNN can smoothly transition between ramping activity and neural sequences, and additionally that in both experiments and computational models dimensionality may be dependent on the degree of training and how well any implicit timing component has been learned.

Multiplexing of WM and elapsed time

To quantify the ability of the RNNs trained on the three tasks to encode both WM and elapsed time, we used a support vector machine (SVM) to decode both cue and time (cue-time)—that is, elapsed time from the onset of each cue based on population activity. Visual inspection of a sample confusion matrix (predicted versus actual cue-time bin) revealed robust cue-time decoding in the T + WM task—and thus that cue-specific elapsed time could be decoded during the delay (Fig. 4a). In contrast, relatively little temporal information was present in the WM task dynamics. And while relatively good decoding of time was possible in the ISA task, ISA-trained RNNs often confused the cue that signalled the start of each delay. Specifically, the secondary diagonal lines of the confusion matrix indicate that the decoder was confused about whether time bins of 0.5–1 s were associated with cue *A* or *B*. Across all tasks, the median performance, as measured by the correlation between predicted and actual cue-time bins, was above 80% (Fig. 4b, left), indicating that even the apparently persistent fixed-point activity in the WM retained a substantial amount of temporal information. But decoding was progressively better from the WM, to ISA, to T + WM task as measured both the performance and mean squared error (MSE) (Fig. 4b, right; all pairwise comparisons were significantly different with *P* values of at most $P < 10^{-6}$, Wilcoxon rank sum test). Consistent with the need for stimulus-specific encoding of time in the T + WM task and stimulus-independent encoding of time in the ISA task, cross-cue decoding of time revealed very poor performance in the T + WM (and WM) tasks, but good decoding in the ISA task (Supplementary Fig. 3).

We next asked whether WM memory and time were multiplexed at the level of individual units—as opposed to, for example, a modular strategy in which some units encoded WM and others time. Analysis of the correlation of the mean activity during the delay epochs revealed largely non-overlapping populations of units activated in response to the short and long cues in the WM and T + WM tasks but largely overlapping in the ISA task (Fig. 4c,d). The median Pearson correlation was close to zero across RNNs for the WM and T + WM tasks, but above 0.75 in the ISA task. These results reflect the preservation of cue-specific information during the delay period in the WM and T + WM tasks, but significant loss of cue-specific information during an explicit timing task (ISA). The multiplexing of WM and timing information in the T + WM task at the level of individual units was confirmed in the high levels of mutual information individual units contained about both cue and elapsed time, as well as the high degree of correlation between them (Supplementary Fig. 4).

Neural sequences rely heavily on inhibitory connectivity

Dynamic regimes that generate neural sequences have been observed in many brain areas[12,30–35,43], and produced in a number of artificial neural network models[13,14,25,44,45]. While these studies have led to critical insights into the potential circuit mechanisms underlying neural sequences, the mechanisms underlying regimes in which the same units participate in multiple sequences, as well as the contribution of different synapse classes, are not fully understood. To begin to dissect the circuit principles underlying the emergence of neural sequences in the T + WM task we partitioned W^{RNN} into its four submatrices ($W^{RNN}_{Ex\leftarrow Ex}$, $W^{RNN}_{Ex\leftarrow Inh}$, $W^{RNN}_{Inh\leftarrow Ex}$, $W^{RNN}_{Inh\leftarrow Inh}$), and ordered both the pre- and postsynaptic neurons of the matrix according to peak firing latency during the delay (the long delay was used for visualization purposes). The sorting was performed separately for the Ex and Inh populations (Fig. 5a). Next, to extract any general structure underlying the neural sequences we averaged the segmented and sorted weight matrices across all 17 RNNs into a master weight matrix (Fig. 5b–d). Note that the weights of the entire network are shown, including the units that never active during the long delay or whose peak was outside the delay—which were placed first in the sorting order—thus it is the weight structure of the latter units of the sorted sequence that are associated with the neural dynamics during the delay epoch. Additionally, the number of units participating in the delay dynamics varied considerably across RNNs. Despite these substantial sources of variability, a dominant diagonal component is visible in the synaptic structure of all four weight submatrices in the T + WM task.

To quantify the structure of weight submatrices, we averaged the weights according to the relative differences in peak activity latency of the presynaptic units (essentially the average of the diagonals in Fig. 5b–d), allowing for the visualization of the net synaptic relationships between a presynaptic unit and the postsynaptic units that fired before or after it (Fig. 5e–g). In the T + WM task, all four weight submatrices revealed peaks centred at approximately zero (corresponding to the main diagonal of the submatrices in Fig. 5c). The $W^{RNN}_{Ex\leftarrow Ex}$ submatrix reveals that presynaptic excitatory units provide above-average

Fig. 2 | Differential dynamics for the encoding of WM and time across in RNNs trained on three tasks. a, Schematic of the RNN architecture and the inputs and target outputs for the WM task during the control and reverse conditions. **b**, Neurograms during the *AA* (upper row) and *BA* (lower row) conditions (*A*, red line above neurogram; *B*, green), sorted according to the peak time during the short (left) or long (right) delays (standard trials), the images in left and right subpanels of each row is based on the same data, but differentially sorted. The self-sorted neurograms (top left and lower right) are cross-validated (average of even trials sorted on average of odd trials). Only the top 50% of units with the highest peak activity during the delay are shown. Heat map is normalized to one for each unit. The overlaid white line shows the 'motor' unit (right *y* axis). **c**, Mean activity of two sample units across all four trial types of the standard condition. **d**–**f**, Similar to **a**–**c** for the T + WM task: schematic (**d**), neurograms (**e**) and mean activity (**f**). In **e** the activity of the 'temporal expectation' output unit is shown in the overlaid grey lines. Note that while each cue elicits a neural sequence, each sequence is different, reflecting the embedding of multiple sequences within the RNN. **g**–**i**, Same as **a**–**c** for the ISA task: schematic (**g**), neurograms (**h**) and mean activity (**i**). Note that in this example both cues A and B elicit approximately the same neural sequence (**h**), and the units show similar time fields in response to cues *A* and *B* (**i**). **j**–**m**, Quantification across 17 RNNs trained on the three tasks. Performance (correct match/non-match responses) during standard (**j**) and reverse (**k**) trials, dimensionality during the delay epoch of the concatenated activity during the short and long trials (**l**, Wilcoxon rank sum test, $n = 17$, $P = 3.58 \times 10^{-7}$ and $P = 4.44 \times 10^{-7}$ for T + WM versus WM and ISA, respectively), and the sequentiality (Seq.) index during the long trials (**m**, Wilcoxon rank sum test, $n = 17$, $P = 7.05 \times 10^{-7}$ and 7.46×10^{-5} for T + WM versys WM and ISA, respectively). Whisker plots represent medians (circle centres), the interquartile range (boxes) and the most extreme values within 1.5× the interquartile range above or below the interquartile range (whiskers), dots represent 'outlier' values.

https://doi.org/10.1038/s41562-023-01592-y

input to excitatory units that fire shortly before and after it. The connectivity was biased in the forwards compared to the backwards direction. While there were differences in the $W^{RNN}_{Ex \leftarrow Ex}$ submatrix between the T + WM task compared to the other tasks, it was the inhibitory connections that were most distinct. For example, the $W^{RNN}_{Ex \leftarrow Inh}$ tuning was the broadest and much stronger in the T + WM task. Thus, there was a marked window of disinhibition from inhibitory to excitatory neurons with similar time fields (Fig. 5f)—in other words, the Inh units did not inhibit the excitatory neurons that activated them, but strongly inhibited 'past and future' excitatory (Ex) units (Discussion). To establish a causal link between the structure of each weight submatrix and the observed neural dynamics we examined RNN performance after

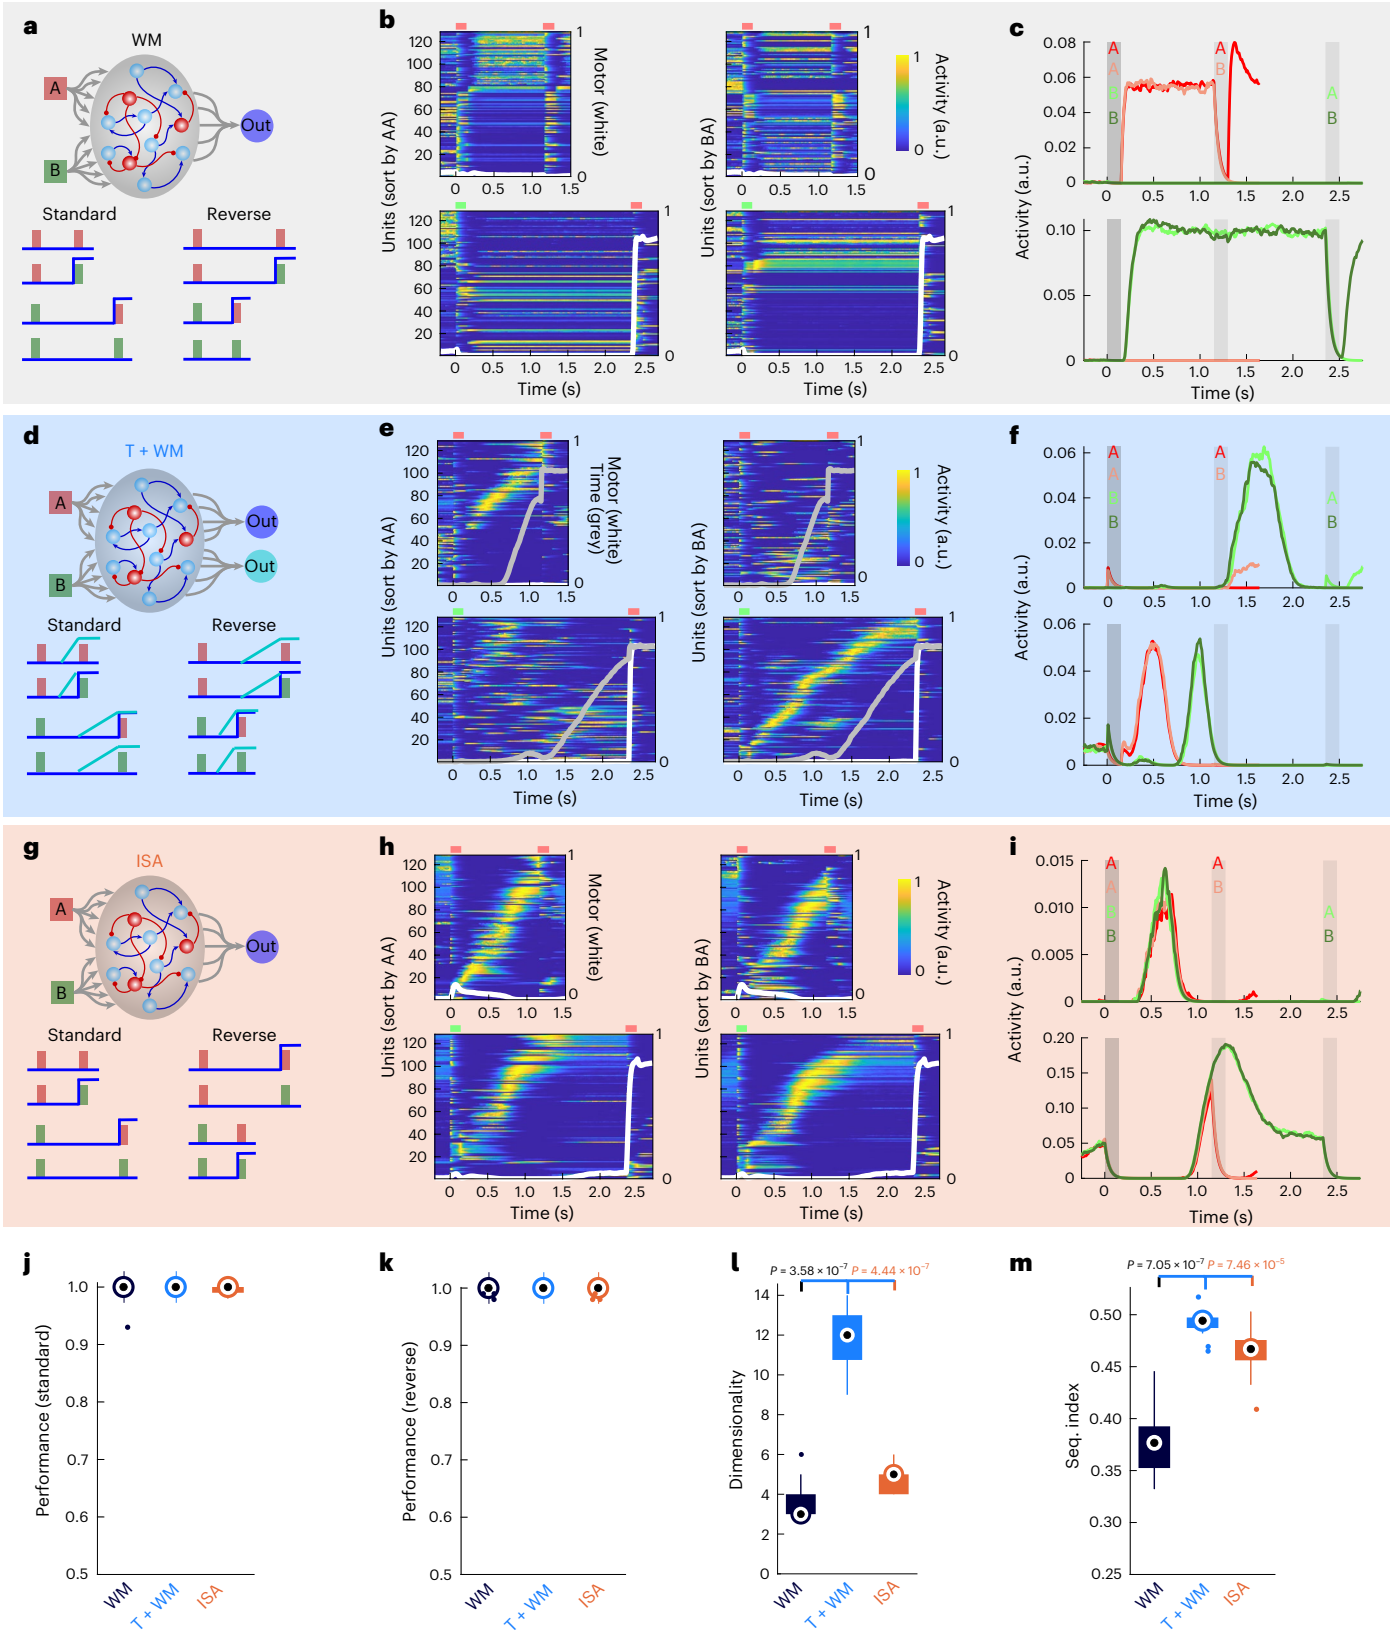

https://doi.org/10.1038/s41562-023-01592-y

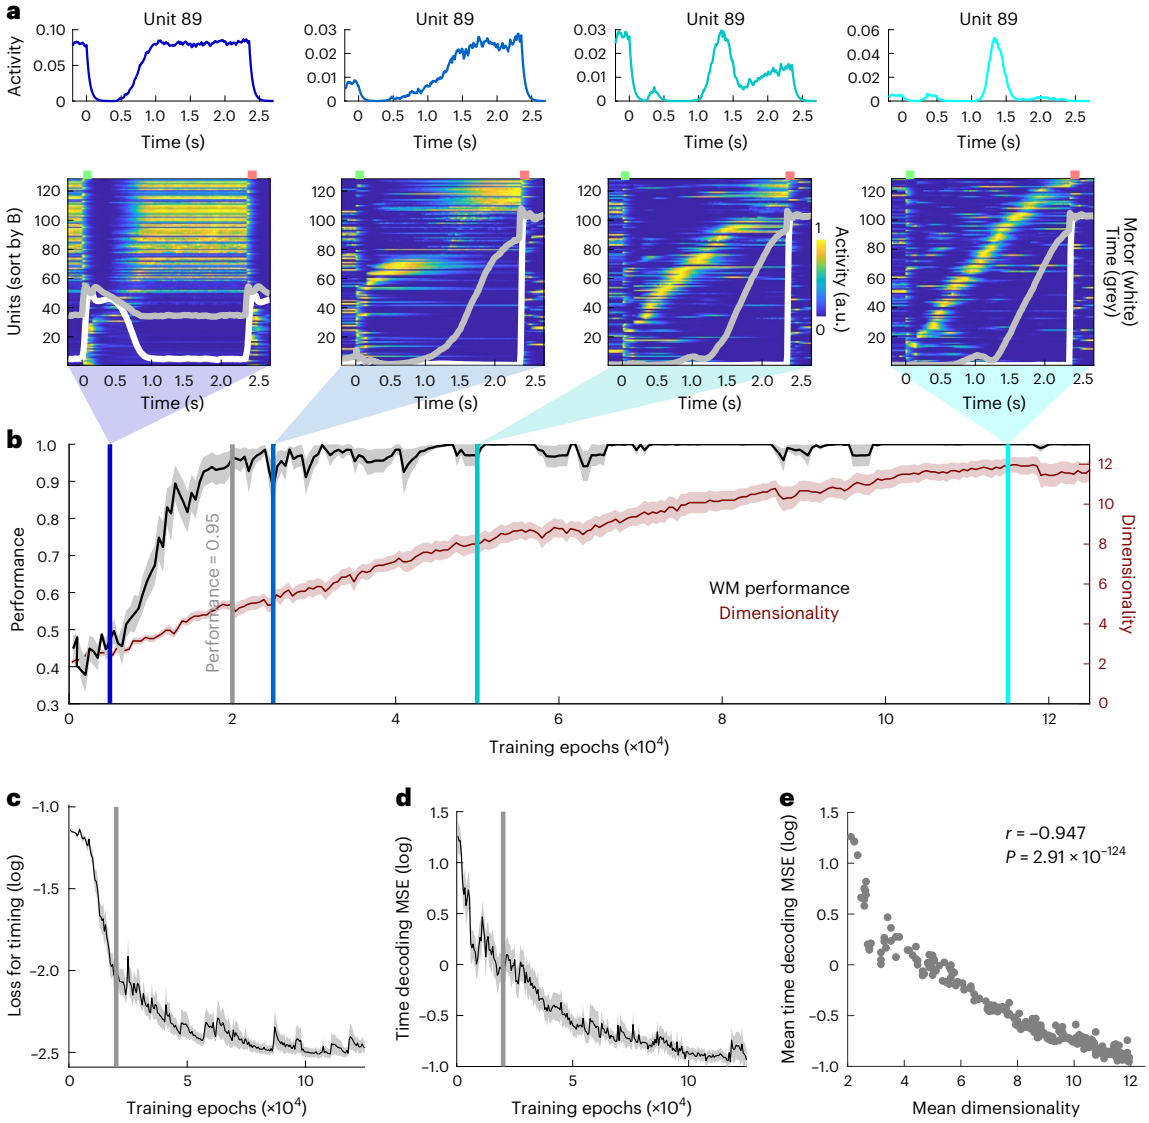

Fig. 3 | Transition from low-dimensional ramping to high-dimensional neural sequences over the course of training in the T + WM task. a, Temporal activity profile of a sample unit (top) and neurograms (bottom) of activity in an RNN trained on the T + WM task at the stages of training corresponding to the vertical coloured lines in **b**. Note that in the neurograms the sorting order of the units in the panels is different. **b**, Performance of WM and population dimensionality (during the delay period) across training. The grey vertical line denotes when the mean WM performance reached 0.95. **c**, Learning curve of the loss for the timing output. **d**, Same as **c** for MSE of the decoding of elapsed time from each cue. **e**, Relationship between the dimensionality and decoding time MSE averages across 17 RNNs (Pearson correlation, $r = 0.947$, $P = 2.91 \times 10^{-124}$). Data are presented as mean values ± s.e.m. for **b**–**d**.

shuffling all non-zero weights in each submatrix (Supplementary Fig. 5). Shuffling either of the excitatory or inhibitory matrices resulted in catastrophic drops in performance in the WM and ISA tasks. By contrast, for the T + WM task shuffling the excitatory weights ($W^{RNN}_{Ex\leftarrow Ex}$ or $W^{RNN}_{Inh\leftarrow Ex}$) led to median performance levels of approximately 75%, while shuffling the inhibitory weights ($W^{RNN}_{Ex\leftarrow Inh}$ or $W^{RNN}_{Inh\leftarrow Inh}$) resulted in near-chance performance. The distinct weight matrix structure and the presence of strong feed-forward modes in the RNNs trained on the T + WM task were also evident in the higher magnitude of the Schur modes of the weight matrix (Supplementary Fig. 6).

These results indicate that for the same RNN to generate multiple sequences with shared units, it relies more on the connectivity structure of inhibitory connections than excitatory connections (that is, both $W^{RNN}_{Ex\leftarrow Ex}$ and $W^{RNN}_{Inh\leftarrow Ex}$). Specifically, the partial resistance to shuffling Ex → Ex or Ex → Inh weights indicates that the specific weight values of the non-zero connections are not as important as those of the Inh → Ex or Inh → Inh weights. Indicating that the excitatory weights provide a non-specific excitatory drive, unit at any moment in time drive inhibitory units that specifically inhibit all other Ex units while opening a window of disinhibition for the current and next excitatory units in the sequence. These results generate the prediction that the most important site of plasticity for the generation of neural sequences is inhibitory plasticity onto excitatory neurons rather than between excitatory neurons.

Importance of hyperparameters

To determine whether the emergence of neural sequences was dependent on RNN hyperparameters we contrasted the RNN dynamics trained on all three tasks across different hyperparameter configurations (Methods), including learning rate, L2 activity regularization, noise in the recurrent units, W^{RNN} initialization, presence or absence of Dale's law, activation function and the profile of the temporal expectation function (half versus full ramp), and the weighting of the WM and timing components ('implicit' versus 'explicit'). Each hyperparameter

https://doi.org/10.1038/s41562-023-01592-y

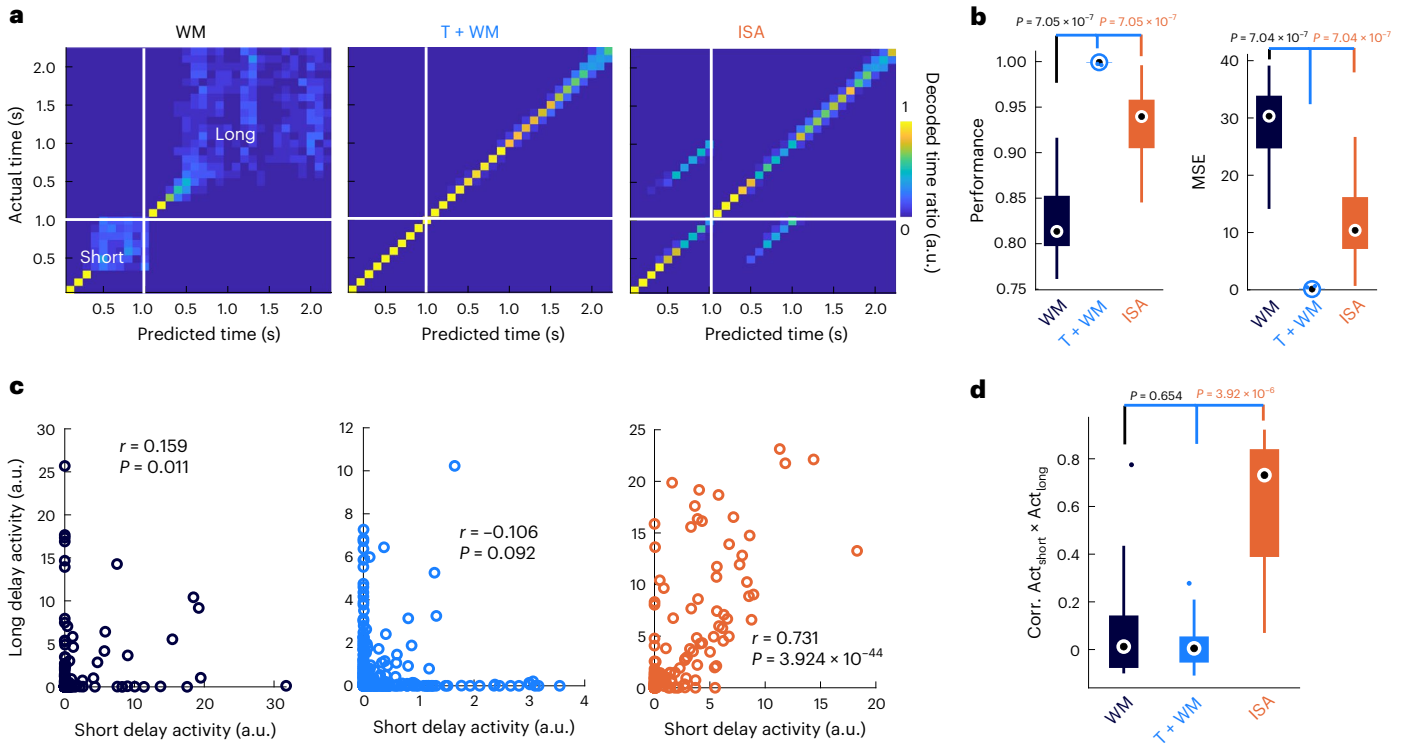

Fig. 4 | Multiplexing of time and WM. a, Confusion matrices of the decoding of cue and time bins (100 ms) for a sample RNN trained on each task. Note that in this example, the ISA network often confuses time bins from approximately 0.5–1 s during the short (cue *A*) and long (cue *B*) delays as reflected in the off-diagonal bands. **b**, Performance (left, Wilcoxon rank sum test, $n = 17$, $P = 7.07 \times 10^{-7}$ for T + WM versus WM and ISA) and MSE (right, Wilcoxon rank sum test, $P = 7.04 \times 10^{-7}$ for T + WM versus WM and ISA) of the decoders across 17 RNNs. **c**, Correlation between unit activity for sample RNNs during the short and long delays in the WM (left, Pearson correlation, $n = 17$, $r = 0.159$, $P = 0.011$), T + WM (middle, Pearson correlation, $r = 0.106$, $P = 0.092$) and ISA (right, Pearson correlation, $r = 0.731$, $P = 3.924 \times 10^{-44}$) tasks. **d**, For the WM and T + WM tasks there was little or no average correlation (Corr.) across all units within each RNN ($n = 17$ RNNs in each group). In the ISA task average correlation between unit activity (Act) in the short and long delay was high and significantly above the WM and T + WM tasks (Wilcoxon rank sum test, $n = 17$, $P = 5.58 \times 10^{-5}$ and $P = 3.92 \times 10^{-6}$). Additionally, only the ISA correlation was significantly different from 0 (two-sided Wilcoxon signed rank test, $P = 2.93 \times 10^{-4}$). Whisker plots represent medians (circle centres), the interquartile range (boxes) and the most extreme values within 1.5× the interquartile range above or below the interquartile range (whiskers), dots represent 'outlier' values.

took on values anchored around the default set of parameters used above, for a total of 366 RNNs over 12 seeds and the three tasks. Across all hyperparameters (Supplementary Fig. 7), with one exception, the dimensionality of the dynamics during the delay period was significantly higher in the T + WM task compared to the WM task (*P* values of at most 10^{-4}, Wilcoxon rank sum test) and the ISA task (*P* values of at most 0.001, Wilcoxon rank sum test). We also confirmed the generality of the observed increase in dimensionality across training (Fig. 3) for a subset of hyperparameters and multiplexing of time and WM for all hyperparameters studied (Supplementary Fig. 8a,b).

The clear exception to the formation of higher dimensional regimes in the T + WM task was the use of the softplus activation function versus the default ReLU function. Convergence was significantly worse for the softplus compared to ReLU activation function in the T + WM task ($n = 12$, $P = 3.66 \times 10^{-5}$, for the final loss value, Wilcoxon rank sum test). The dimensionality of the dynamics for the RNNs trained with the softplus function was uniformly low at a value of two for all tasks. These data indicate that the use of a softplus function shifts the encoding of time and WM in the T + WM task to ramps rather than population clocks (Supplementary Fig. 9). Indeed, by fitting the activity of the RNN units to both linear ramps and Gaussians during the delay, we observed a dramatic shift in the goodness of fit; while the units from RNNs trained with the ReLU activation function were on average very poorly fit by linear ramps, softplus units were fit well (Supplementary Fig. 8c,d). Consistent with the results above, indicating that low-dimensional ramps are not well suited to flexible timing, the ability of softplus RNNs to generate the half-ramp timing output was worse (the MSE of the timing units was significantly higher compared to the ReLU RNNs; $P = 4.69 \times 10^{-5}$, Wilcoxon rank sum test). The shift from high- to low-dimensional regimes with the softplus activation function, may be a result of its worst performance and an interaction between the learning algorithm and the continuous derivative of the softplus function (Discussion).

As mentioned above, depending on brain area, both the high-dimensional neural sequences and low-dimensional ramps are indeed observed in timing and WM tasks (Discussion). Thus our results establish that RNNs can account for both these experimentally observed regimes in a hyperparameter-dependent fashion. Raising the possibility that differential intrinsic neuronal properties in different areas could contribute to their observed dynamics.

Across all other hyperparameters, the dimensionality was higher in the T + WM task. Interestingly, some hyperparameters had surprising effects on the relative contribution of excitatory and inhibitory units. As described above, in the T + WM task the weights from the excitatory units contributed less to the dynamics than that of the inhibitory weights (Fig. 5 and Supplementary Fig. 5), this phenomenon was further amplified when the noise of the recurrent units (σ^{RNN}) was increased from 0.005 to 0.05 (Fig. 6). At the default noise level of $\sigma^{RNN} = 0.005$, deletion of all excitatory-to-excitatory weights (that is, zeroing the entire $W^{RNN}_{Ex \leftarrow Ex}$ submatrix after training), resulted in a catastrophic drop in performance. Note that even in the absence of excitatory-to-excitatory connections (and biases of zero) some activity is driven by the noise.

https://doi.org/10.1038/s41562-023-01592-y

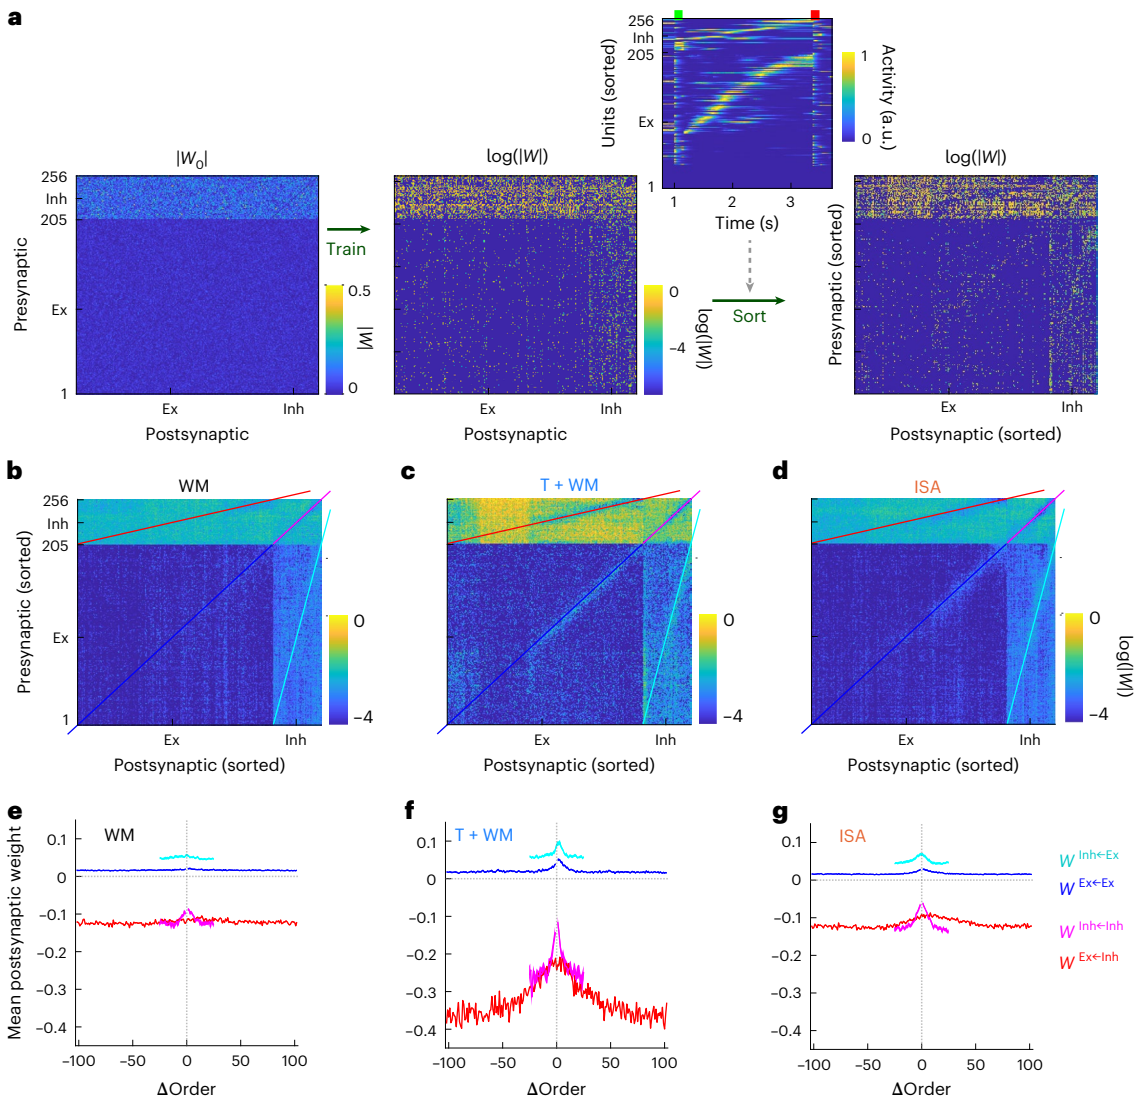

Fig. 5 | Circuit motifs underlying the generation of multiple neural sequences. a, Absolute weights of a sample recurrent weight matrix before training (left), log of weights after training on the T + WM task (middle) and the same post-training weight matrix sorted by the peak latency times of the excitatory (Ex) and inhibitory (Inh) units (inset) during the delay epoch (right). **b**–**d**, The mean of the sorted weight matrices across all 17 RNNs trained on the WM (**b**), T + WM (**c**) and ISA (**d**) tasks. All matrices are plotted on the same log scale. **e**–**g**, Mean weights (linear scale) from presynaptic units averaged across all RNNs trained on each task: WM (**e**), T + WM (**f**) and ISA (**g**). The presynaptic units are arranged according to their relative peak firing latency during the delay period. Thus ΔOrder = 0 captures the mean synaptic weights between pre- and postsynaptic units that have the same peak latency, ΔOrder values of 10 and −10 capture the mean weights of the connections from presynaptic units to the postsynaptic units whose peak activity is 10 ms after or before, respectively, the peak of their corresponding presynaptic units.

At noise levels of $\sigma^{RNN} = 0.05$, ablation of all excitatory-to-excitatory connections, only had a modest effect on the dynamics and performance (median greater than 90%) during the T + WM task (Fig. 6c,d). Critically, in the ISA task deletion of excitatory-to-excitatory weights decreased performance to chance. In other words, there is a fundamental shift in circuit architecture when RNNs encode a single sequence (ISA) versus two sequences (T + WM). In the former case, RNNs rely on excitatory-to-excitatory connections, but in the latter case RNNs can operate relatively well in the absence of recurrent excitation.

In the context of neurobiological circuits, our interpretation of these findings is that when encoding cue-specific elapsed time, in some contexts, including high noise, RNNs autonomously converge to circuits architectures that resemble the circuit motifs of the striatum, cerebellum and CA1: that is, circuits in which there are no excitatory-to-excitatory connections, which are driven by external input and negative—rather than positive—feedback loops[26,46].

Dynamic attractors

Whereas it is widely accepted that the neural dynamics generated by recurrent neural circuits play a fundamental computational role in WM and timing, the dynamics itself has generally been interpreted in the context of standard dynamical system regimes of fixed-point attractors, saddle points, line attractors and limit cycles. The high-dimensional trajectories and neural sequences observed in the T + WM tasks do not seem to neatly fit into these classes. But as with standard dynamic regimes, a critical question pertains to the stability of the trajectories: that is, when perturbed, do trajectories further diverge, remain approximately parallel or converge back to the original trajectory? To address this question, we performed perturbation experiments during the delay period.

We initially contrasted perturbed and unperturbed trajectories in the presence of frozen noise in RNNs trained on the T + WM task. At the level of individual units, the perturbation immediately altered

https://doi.org/10.1038/s41562-023-01592-y

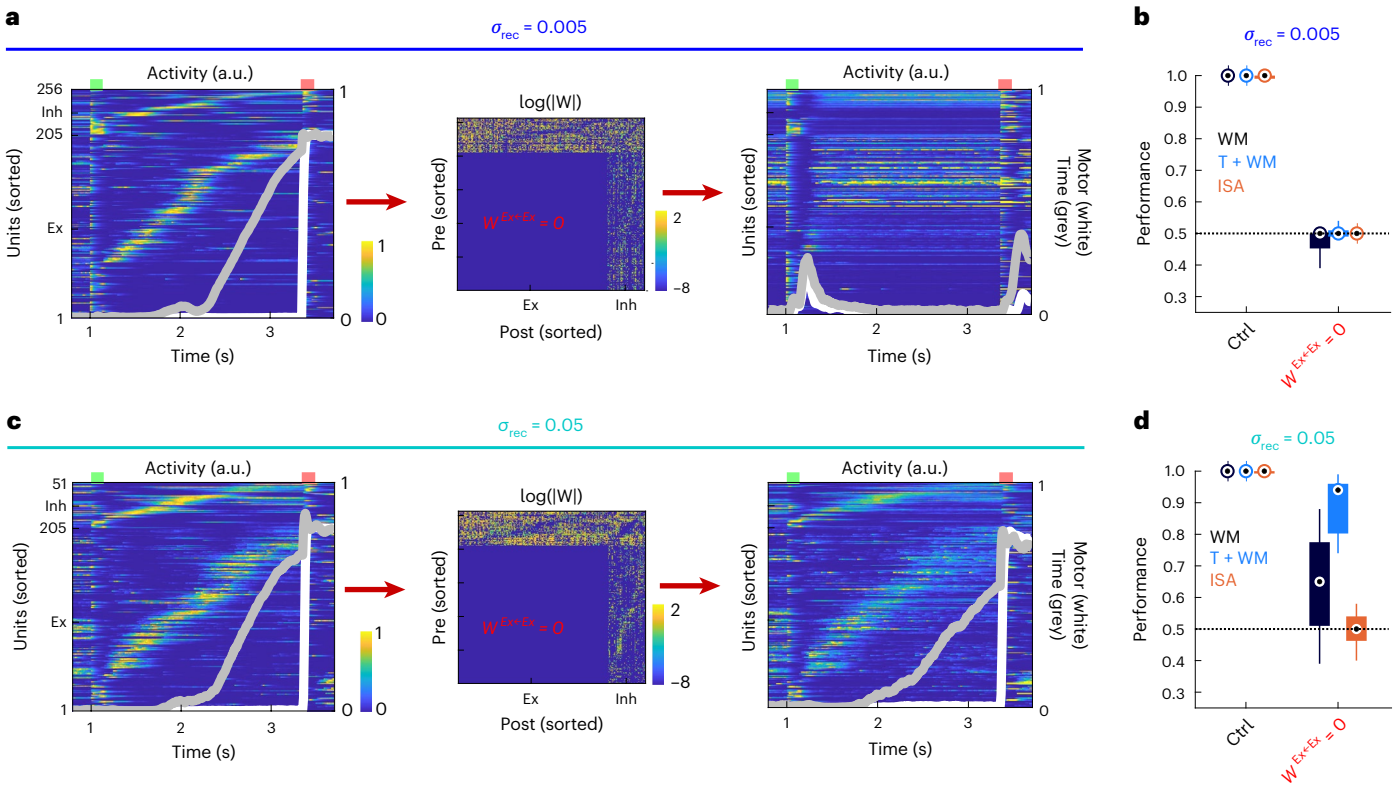

Fig. 6 | In the presence of high noise ablation of all Ex → Ex connections has little effect on performance in the T + WM task. a, Sample dynamics during the long (*BA*) trials with default low noise in the RNNs $\sigma = 0.005$ (left). Dynamics after ablation of all Ex → Ex connections (right). **b**, The mean performance for all three tasks fell from near perfect to chance on deleting all Ex → Ex connections. **c**, Dynamics corresponding to **a** in a sample RNN with noise $\sigma = 0.05$ before (left) and after (right) ablation of all Ex → Ex connections. **d**, In the presence of higher noise levels, the deletion of all Ex → Ex connections had a modest effect on the performance of the T + WM task and a moderate effect on the WM task. Activity scales are the same in the left and right neurograms of panels **a** and **c**. Data from the same set of stimulations in Supplementary Fig. 7 ($n = 12$ in each group). Whisker plots represent medians (circle centres), the interquartile range (boxes), and the most extreme values within 1.5× the interquartile range above or below the interquartile range (whiskers).

activity levels, but over the course of hundreds of milliseconds, not only did activity converge back to the unperturbed levels, but it also converged in a manner that preserved the original temporal alignment (Fig. 7a). The effect of the perturbation at the population level can be visualized in the cross-Euclidean distance matrix (Fig. 7b), which shows that after the perturbation the main diagonal seems to converge back to close to zero. If the trajectory remained parallel to the original, the diagonal would not return to zero; and if it converged back, but either ahead or behind in time, the minimal values would be off the main diagonal. To quantify the effects of perturbations across RNNs trained on all three tasks we plotted the distance of the main diagonal between the perturbed and unperturbed trajectories (Fig. 7c). While the distance does not generally converge to exactly zero, the perturbed trajectory always converges back towards the unperturbed trajectory. These results are consistent with the notion that RNNs trained on the T + WM task implement dynamic attractors, that is, locally stable transient channels[47,48] in which the dynamic attractor is a 'hypertube'. Within limits, perturbations of the trajectory can return to the hypertube in motion. The stability of the neural sequences in the T + WM task was quantitatively comparable to the fixed-point attractor-like dynamics in the RNNs trained on the WM task and both are significantly better than that trained on ISA tasks (Wilcoxon rank sum test, $P = 1.32 \times 10^{-6}$ and $P = 0.0012$ for WM versus ISA and T + WM versus ISA groups, respectively).

Discussion

WM and the ability to encode and tell time on the scale of seconds—and thus predict and anticipate external events—are critical to a wide range of cognitive and behavioural tasks. Here we propose that, in some instances, WM and the encoding of elapsed time may be two sides of the same coin, that is, both WM and elapsed time are represented in the same neural code. This link between WM and time is supported by our findings, and previous findings, that WM is impaired when information has to be recalled at unexpected times[7,24], and evidence that in many cases WM is encoded in time-varying patterns of activity[9,10,18,30,36,49].

We first showed that during a WM task (dDMS), participants implicitly learn the task-irrelevant cue-delay associations, and that the time at which the memorandum is accessed not only alters task RT, but accuracy as well. Conversely, during an explicit timing task (ISA), the task-irrelevant cue that marks $t = 0$, also influenced both RT and accuracy. These findings are consistent with results showing that when stimulus-specific temporal structure is present during WM memory tasks, that recall is 'dynamically prioritized'[24]. These psychophysical results, of course, do not establish that WM and time are multiplexed at the neural level, but demonstrate an interaction.

The observation that humans learn the temporal structure of the dDMS task justified training RNNs on a task that required learning of WM and elapsed time (the T + WM task). RNNs are not well suited to study certain psychophysical phenomenon, including implicit learning and behavioural RTs, as they are not bounded by evolutionarily cognitive strategies or resource optimization constraints. But they have consistently captured the dynamic regimes observed in the brain and provided insights into the biological circuit mechanisms underlying the dynamics[27,40,41,50]. Indeed, consistent with these previous studies, our results revealed RNN dynamics that mirrored a large range of experimental observations.

https://doi.org/10.1038/s41562-023-01592-y

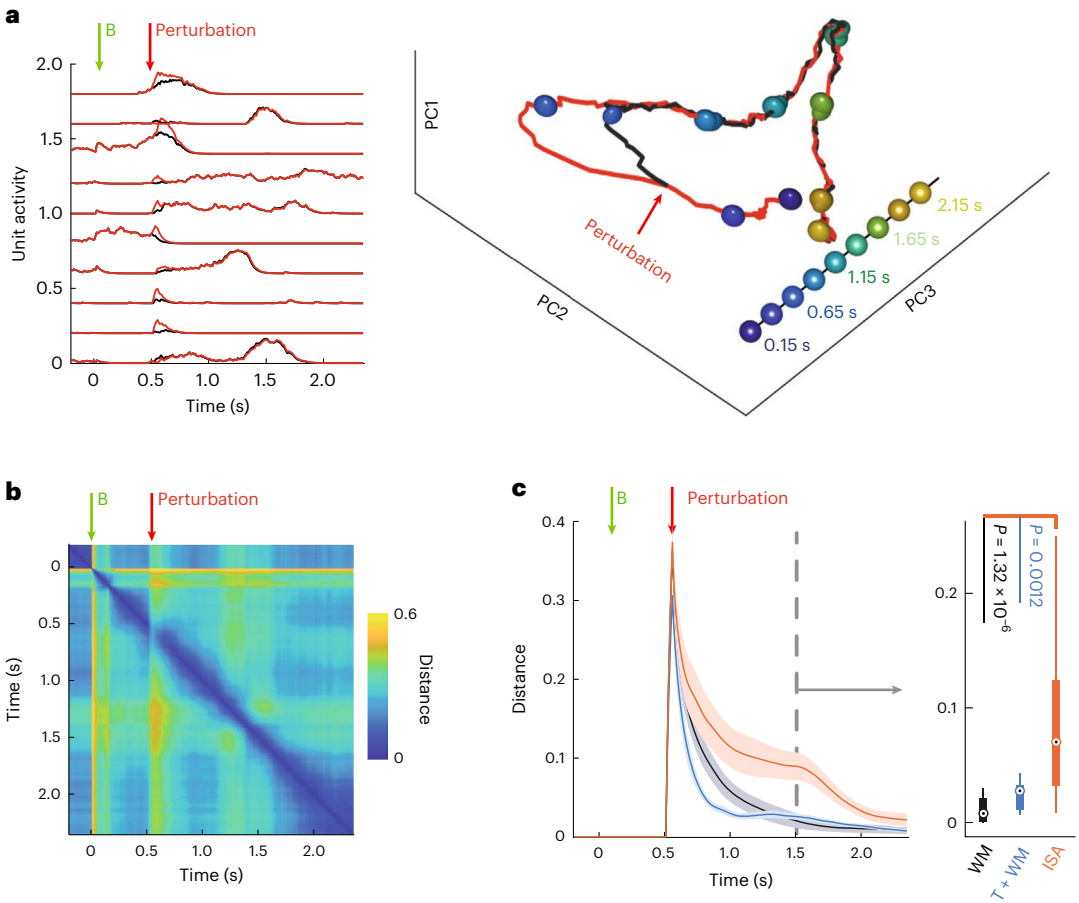

Fig. 7 | Neural sequences instantiate dynamic attractors. a, Activity profile (left) of a sample of 10 units in response to the *B* stimulus (long cue) in the absence (black) and presence (red) of a perturbation (frozen noise) for the T + WM task, and the unperturbed and perturbed neural trajectories in principal components (PC) analysis space (right). Note that the trajectory converges back to the unperturbed trajectory in an approximately time-aligned manner as indicated by the overlapping time marker spheres. **b**, Cross-Euclidean distance matrix between the unperturbed and perturbed trajectories. **c**, Mean Euclidean distance between the unperturbed and perturbed trajectories after the perturbation at 0.5 s averaged across 17 RNNs (left, data are presented as mean values ± s.e.m). Comparison of median Euclidean distance values averaged across the same 17 RNNs 1 s after the perturbation ($t = 1.5$, dashed line, right, Wilcoxon rank sum test, $P = 1.32 \times 10^{-6}$ and $P = 0.0012$ for WM versus ISA and T + WM versus ISA respectively). Whisker plots represent medians (circle centres), the interquartile range (boxes) and the most extreme values within 1.5× the interquartile range above or below the interquartile range (whiskers).

Fixed-point and dynamic representations

Early experimental[15,16] and computational[51,52] studies of WM focused on stable persistent activity, and provided an intuitive computational framework to transiently store memory information in a time-independent fashion. That is, precisely because information was encoded in a fixed-point attractor, a given memorandum could be retrieved at any time using the same encoding and decoding scheme. A counterintuitive aspect of storing WM in time-varying neural trajectories is that the downstream circuits must recognize that even though the code is changing in time, the memorandum is the same[11]. However, as long as WM is encoded in unique non-overlapping trajectories, downstream areas can either automatically generalize across time[42] or learn to recognize the trajectory at all time points—a process that could occur during memory consolidation. Here, even though the dynamics during the delay period of the T + WM task was time-varying and high-dimensional, performance was near perfect at the learned standard and reverse delays, and generalized to new intermediary intervals (Supplementary Fig. 10).

Dynamic activity has been observed during the delay period in many WM tasks, and it has been shown that the temporal structure of tasks can itself influence the observed dynamics. For example, a task with a fixed delay generated transient dynamics in a premotor area, while random delays resulted in more persistent and stable patterns[53]. However, in another study, the sequentiality of population activity in the prefrontal cortex was larger in a WM task with a random compared to a fixed delay[54]. One advantage of the dDMS task—in contrast to standard DMS tasks in which all stimuli share the same delay—is that if WM is encoded in persistent stable activity, both memoranda should elicit fixed-point dynamics. However, if WM is multiplexed with time, the duration or speed of the dynamics should be distinct across different memoranda.

Ramps versus sequences

There is ongoing debate regarding the neural regimes underlying both WM and the coding of elapsed time on the scale of seconds. Critically, however, with the exception of stable persistent activity, the candidate mechanisms for both WM and timing are largely overlapping. Ramping activity, neural sequences, complex neural trajectories and 'activity-silent' models have all been raised as possible mechanisms for both WM and timing. In the context of WM, activity-silent mechanisms have generally focused on STSP, which can maintain a memory of previous activity in the absence of ongoing neural activity[9,19,55]. Similarly, early computational models and subsequent experimental results indicate that STSP underlies some forms of sensory timing, by encoding elapsed time in the so-called 'hidden state' of neural circuits[20,56].

Neural sequences and high-dimensional activity have been observed across many brain areas during WM and timing tasks[12,30–34,54,57–59]. Conversely, low-dimensional ramping activity has also been observed across areas and tasks[36–39,60–62]. Furthermore, even within a single brain area different classes of excitatory neurons may differ in the degree to which they encode time and WM[37]. The diversity and complexity of the experimental findings make it challenging to develop area-specific computational models. But here we have shown that depending on task structure and hyperparameters RNNs converge to fixed-point attractors, ramping firing rate, or neural sequences. Critically, we show that from the perspective of the circuitry generating the dynamics, ramps and neural sequences may not represent fundamentally different regimes since we observed a transition from ramps to neural sequences over the course of training.

It is notable that although across all hyperparameters WM and timing were multiplexed, the softplus activation function dramatically shifted RNN dynamics from encoding time in neural sequences to ramping activity (Supplementary Figs. 8c,d and 9). Our results are consistent with the fact that ReLU activation function is often considered to lead to better convergence than softplus activation[63] as convergence time and final cost values were worse with the softplus function. Arguably, one might also consider the ReLu function to be more naturalistic in that it would seem to better capture the discrete nature of neuronal thresholds and can have output values of zero. However, there are counterarguments based on noise-induced 'smoothing' that continuous functions may be more naturalistic[64,65]. Future studies will have to examine the strong impact of the activation function on RNN dynamics, but our results reveal a strong interaction between the learning algorithm and the continuous derivative of the softplus function.

The current study does not speak to activity-silent models, but provides insights to the potential trade-offs between encoding information in neural sequences and ramping activity. Specifically, despite their apparent complexity and higher dimensionality, neural trajectories approximating neural sequences emerged in a highly robust manner across ReLU RNNs trained on the T + WM task. This is consistent with previous RNN models trained on WM or timing tasks, in which neural sequences were observed[41,66,67] and models of WM that rely on sequential dynamics[13]. A previous computational study has observed the emergence of low-dimensional ramping dynamics that can encode time during WM tasks[11], this study however, did not examine tasks that require stimulus-specific encoding of WM and time. The low-dimensional dynamics observed in that study are consistent with those obtained here with the softplus activation function, but a direct comparison is difficult because that study used smaller RNNs that did not implement Dale's rule (and tanh activation function). Here we show that task structure, hyperparameter choices and training stages, probably account for these differences. But a critical question across experimental and computational studies pertains to the computational trade-offs between the high- and low-dimensional representations of WM and time. One clear trade-off pertains to ease of generalization to new delays, and the use of the neural representations by downstream areas to create arbitrary time-varying outputs. While ramping activity in RNNs is a highly limited representation in terms of its ability to generate outputs other than ramps (including the half ramp used here), ramps are well suited to temporal generalization[42]. In contrast, neural sequences provide a robust high-dimensional set of basis functions that can drive arbitrarily complex temporal outputs including the default half ramp used here[33].

Conclusions and predictions

Internally generated high-dimensional neural trajectories, including neural sequences, have been reported in a large number of brain areas across many tasks[12,30–34,43,54,57–59] and present in many computational models[14,41,47,66,68]. We postulate that this is because neural sequences represent a canonical dynamic regime to encode WM, time and generate motor patterns. The relatively high dimensionality, stability and quasi-orthogonality of neural sequences, make them well suited for downstream areas to generate either simple or complex time-varying output patterns[33]. Our results predict that neural sequences observed in vivo are not solely the product of feed-forward architectures as proposed in some models[45,69], but require recurrent connectivity that can implement feed-forward dynamics. These results are in general agreement with those of Rajan et al.[14] and Orhan and Ma[66], which show that neural sequences emerge from non-symmetric but recurrent connectivity. In contrast to those studies, however, we implemented separate populations of excitatory and inhibitory neurons, and found that inhibitory (Inh → Ex and Inh → Inh) connections were more important than excitatory (Ex → Ex and Ex → Inh) connections (Supplementary Fig. 5). We also observed that in some hyperparameter regimes RNNs could perform well even after deletion of all Ex → Ex connections (Fig. 6). Overall our results indicate that sequence generation may rely on non-specific excitation (for example, reflected in the fact that shuffling the non-zero Ex → Ex weights did not dramatically impair performance in the default RNNs) that lead to suprathreshold activity through a transient window of disinhibition[26]. In other words, the active population of inhibitory units, disinhibit the currently active and to be activated excitatory neurons while blanketing the off-diagonal excitatory neurons with inhibition. Regarding the biological mechanisms underlying the emergence of neural sequences our results also predict that neural sequences are strongly dependent on inhibitory plasticity, and paradoxically, can be independent of structured excitatory-to-excitatory connections[50].

Consistent with the notion that memory serves both retrospective and prospective functions[22,70], we propose that when WM tasks contain temporal structure, WM and time are multiplexed either in neural sequences or ramping activity. Furthermore, the diversity of experimentally observed neural correlates in WM studies, is in part a reflection of temporal structure of the tasks used. The dDMS task provides a means to address the interaction between WM and implicit timing, as it allows for comparison of the neural dynamics in response to stimuli that have the same WM requirements, but different temporal requirements as to when items in WM will be used.

Multiplexing of WM and time impose additional challenges for downstream decoding[42], and is unlikely to be a universal encoding scheme for WM. However, multiplexing of WM and time may comprise an effective computational strategy in some instances, because, in addition to the need to transiently store retrospective information the brain is continuously attempting to predict when external events happen, including when WM will be used. Additionally, WM and time may be multiplexed because it provides a learned task-dependent manner to control how long items need to be stored, potentially implementing an expiration time on storage and optimizing cognitive resources.

Methods

Human psychophysics

All human psychophysics experiments were approved by the Institutional Review Board of UCLA (IRB no. 20-001801, issued on 21 September 2020). Participants provided informed consent before participating and were paid for their participation. Experiments were conducted online, with hosting provided by Gorilla (https://gorilla.sc/) and recruitment provided by Prolific (https://www.prolific.co/). The precision and accuracy of timing on the Gorilla platform (that is, of visual presentation and RT provides temporal precision with standard deviations of approximately 8–21 ms depending on the exact browser, operating system and device. Participants accessed the experiment using personal computers running Google Chrome or Mozilla Firefox. No other device types (that is, phones or tablets) or browsers were allowed. All analyses relied on within-participant statistics, thus decreasing the impact of cross-platform variability. Participants on the Prolific platform were only eligible for the study

if they were between the ages of 18 and 40, residing in the United States, fluent in English and had never participated in an online study from our laboratory on Prolific. Before beginning the task, participants read and signed an informed consent form that asked them to: (1) complete the study in a quiet place without distractions, (2) maximize their browser window and not adjust it during the experiment, (3) have normal or corrected-to-normal vision and (4) not participate if they had a history of seizures, epilepsy or stroke. After providing consent, participants completed a short demographics form including their age, handedness and gender. Participants were then given instructions on how to perform the task, which stressed the importance of both speed and accuracy. Participants were also informed that if they were faster and more accurate than the average of the other participants in a given sample of participants, they would receive a bonus payment. Across all experiments 130 participants (62 female, five left-handed, mean age of 29 years and range 18–40) participated in the study (each participant only took part in a single study). Data from 17 participants were excluded from analysis due to low accuracy (less than 70%) or consistently slow RTs such that too few trials met inclusion criteria (less than 50% of total possible trials in any reversal × delay condition remaining after RT exclusion, see below).

dDMS task

The background was always white, and all stimuli were black and presented in the centre of the screen. First, a 150 ms duration fixation cross was presented, which indicated the start of a new trial. Following a 500–1,000 ms interval, a 150 ms duration visual cue was presented, which could either be a black circle or black star, matched for area, with 50% probability. After the delay epoch (below), a 150 ms duration probe stimulus was presented that was either the same or remaining stimulus with 50% probability. Participants were instructed to press one of two buttons on their keyboards, F or J, to indicate whether they thought the cue and probe stimuli matched or did not match (counterbalanced across participants). The response period was unlimited in duration, and the task did not proceed unless a response was given. All incorrect responses were followed by negative feedback (a 'thumbs down' icon). After each response there was a 1,500–2,000 ms intertrial interval.

The critical manipulation was the delay time, that is, the interval between the cue (first stimulus) offset and probe (second stimulus) onset. When appearing as a cue, one stimulus (for example the circle) was followed by a delay of 1 s on 80% ('standard') of the trials, and a delay of 2.2 s on the remaining 20% ('reverse') of the trials. The other stimulus (for example, the star) was followed by a 2.2 s delay on 80% ('standard') of the trials and a 1 s delay on 20% ('reverse') of the trials. The mapping between the cue stimulus and the likely memory delay was counterbalanced across participants.

Five blocks of 80 trials (64 standard, 16 reverse) were presented for a total of 400 trials. In each block, trial order was pseudorandomized with the following constraints: (1) the first eight trials of each block were always standard trials. (2) A reverse trial could not immediately follow another reverse trial. Participants were given eight standard practice trials with each cue before the first block. Participants were given the opportunity to take short breaks between each block. Each block took approximately 8 min to complete, and participants finished the experiment in 45 min on average. A replication study of the dDMS task (Supplementary Fig. 1e–j) was preregistered (https://doi.org/10.17605/OSF.IO/XK3JH).

ISA task

The ISA task was identical in stimulus structure to the dDMS task, but rather than being instructed to compare the cue and probe stimuli to each other, participants were instructed to make a decision on the basis of the probe stimulus and the delay: for example, press the F key in response to a short delay followed by a circle or a long delay followed by a star (short–circle or long–star), and press the J key in response to a long delay followed by a circle or a short delay followed by a star (long–circle or short–star). The mapping between the response button and the pair of opposing interval-probe combinations was counterbalanced across participants. Participants were instructed that for any given trial the cue stimulus could be either a circle or star and that the cue stimulus was irrelevant to the task beyond indicating the onset of the delay interval. But as in the dDMS task, the cue stimulus identity (circle or star) predicted the delay interval on 80% of the trials (standard trials), while for the remaining 20% of the trials, the cue stimulus was followed by the other delay interval (reverse trials).

Analysis and statistics

Trials with RTs outside of the range of 100–3,000 ms were discarded. Three measures of performance were used: accuracy (percentage correct), RT and the inverse efficiency score. For each condition for every participant, trials with RT values larger than four standard deviations away from the mean were discarded. RTs were calculated as the median of the remaining trials for that condition. The inverse efficiency score, a combined measure of speed and accuracy in which larger values indicate worse performance, was calculated as the median RT divided by accuracy.

Statistics were based on within-participants 2 × 2 ANOVAs with a reversal (standard versus reversal trials; for example, circle → short/star → long versus circle → long/star → short) factor and the actual delays (short versus long) factor.

Unless otherwise specified all statistical analyses relied on two-sided tests.

RNN model

RNN architecture and training. RNNs were composed of 256 units, an input layer and an output layer composed of one (WM and ISA tasks) or two (T + WM task) units. The dynamics of the default RNN were described by:

$$\tau \frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} = -\mathbf{r} + \left[W^{\mathrm{RNN}}\mathbf{r} + W^{\mathrm{In}}\mathbf{u} + \mathbf{b}^{\mathrm{RNN}} + \varphi\right]_+$$

where $\mathbf{r}$ is the the firing rate vector of the recurrent units, $\mathbf{u}$ is the input vector and $\mathbf{b}^{\mathrm{RNN}}$ is the bias of the units. Each unit received private Gaussian noise $\varphi = \sqrt{2\tau}\sigma^{\mathrm{RNN}}N(0,1)$, where $N(0,1)$ represents a normal distribution with a mean of 0 and a standard deviation of 1. The threshold linear function $[]_+$ represents a ReLU function in which all negative values become 0.

The input layer was composed of 32 inputs representing a range of 0 to 2π. Stimuli *A* and *B* were represented by non-overlapping patterns of activity centred at 1 and 5.2 (corresponding to centre activation at units 6 and 28). These patterns can be interpreted as retinotopic or tonotopic activation of two visual or auditory stimuli. The output units ($\boldsymbol{z}$) were non-linear readouts of the recurrent network:

$$\boldsymbol{z} = \mathrm{sigmoid}\left(W^{\mathrm{Out}}\mathbf{r} + \mathbf{b}^{\mathrm{Out}}\right)$$

Dale's law was implemented by assigning 80% (205) of the units as excitatory and 20% (51) as inhibitory. After initializing W^{RNN} to a random orthogonal matrix (gain of 0.5), the absolute weights of all presynaptic excitatory units were multiplied by 1.0 and those from the inhibitory units were multiplied by 4.0 (to maintain an approximate excitatory and inhibitory balance). During training all weights were clipped at zero during any zero-crossings. The weight matrix was multiplied by a diagonal matrix composed of 1 s and −1 s, corresponding to the excitatory and inhibitory units, respectively.

To enhance our ability to dissect the circuit mechanisms underlying the observed dynamics, the recurrent biases were set to zero, and neither the recurrent biases or the W^{In} matrix was trained. This approach is consistent with the notion that synapses higher in the

processing hierarchy are more plastic. Additionally, this approach ensured that the differential RNN dynamics across all networks could be attributed to W^{RNN}, facilitating the synaptic structure analyses (Fig. 4). Thus, only the W^{RNN}, W^{Out} and $\mathbf{b}^{Out}$ parameters were trained.

During training a first-order Euler approximation was used with a τ of 50 ms and a dt of 10 ms, resulting in a discretization factor $\alpha = dt/\tau = 0.2$. RNNs were trained with ADAM and a batch size of 32. The loss function to be minimized was:

$$\mathcal{L} = \left\langle \left(m_{i,t} * \left(z_{i,t} - \hat{z}_{i,t} \right) \right)^2 \right\rangle_{i,t}$$

where $\mathcal{L}$ represents the time and output unit averaged loss function and the star represents element-wise product. $z_{i,t}$ and $\hat{z}_{i,t}$ represent the target and actual activity, respectively, of an output unit i at time t. m represents a cost mask that differentially weights the contribution of output units to the loss function during different points in time. The motor output target ($z_{1,t}$) was a step function from 0 to 0.8 at probe onset of the non-match trials, and when present, the temporal expectation output ($z_{2,t}$) was a linear 'half' ramp from 0 to 0.8 starting at 50% of the delay period until onset of the probe. For the motor output $m_{1,t}$ was equal to two from 250 ms before onset of the cue stimulus until the onset of the probe stimulus and five during the probe until 500 ms after probe offset (to place a higher weight on the match or non-match response), with a grace period 5dt ($m_{1,t} = 0$) during the onset of the probe. For the temporal expectation output, the cost mask was always $m_{2,t} = 1$. As with the psychophysics experiments, RNNs were trained on both the standard and reverse trials, but reversal trials comprised 10% of the total rather than 20%, to accelerate convergence (note that for the T + WM task the reverse trials impose a 'moving target' for the timing output pattern). Training was stopped when the loss reached 0.0015 for T + WM and 0.001 for the WM and ISA tasks, or reached a total 125,500 update epochs. Different stop criteria were used for T + WM task because the reverse trials make it impossible to fully converge to the same loss values. RNNs were implemented in TensorFlow and based on Yang et al.[40].

During training the onset time of the first stimuli was uniformly varied between 250 and 1,000 ms on each trial. The standard delays were 1.0 (short) and 2.2 s (long). During training these delays were jittered by ±10%, approximately corresponding to psychophysically observed Weber fractions. It is important to note that the presence of 'temporal noise' in the form of onset and delay time variability, contributes to the robustness of the solution, and that 'spurious' solutions that do not generalize to different onset times or delays can emerge in the absence of this 'temporal noise'.

Dimensionality. To estimate the dimensionality of the dynamics during the delay periods we first concatenated the average activity of all units during the short and long delays for a final matrix of 256 units × 320 time bins. Concatenation is important to distinguish regimes in which both cues elicit similar sequences (for example, in the ISA task) versus cases in which both cues elicit distinct sequences (for example, in the T + WM task). Effective dimensionality was defined as the minimum number of principal components that captured at least 95% of the variance of the activity across time bins[14].

SVM decoding. For the decoding of bue and elapsed time, the mean activity of each unit within 100 ms bin was used, comprising a total of 32 (10 and 22 bins for the short and long delays, respectively) input vectors of size 256 per trial. Thus target values represented bins 1:10 for cue A and bins 1:22 for cue B). SVMs were trained (SVMTRAIN from LIBSVM 1.2 for MATLAB) using multiclass classification, a linear kernel, and a cost parameter of 100. The short and long delay data sets consisted of 25 delay epoch trials of the AA and BA conditions, and testing relied on leave-one-out cross-validation across the 25 possible replications. Performance was quantified as the correlation between the predicted and target bins, as well as the MSE.

Mutual information. To estimate the per unit mutual information about the cue (that is, whether the first stimulus was A or B), the activity was averaged from 100 ms after the end of the cue (to allow for decay of stimulus-evoked activity) to the end of the delay period for each trial. Activity levels were categorically binned into ten bins from 0 to maximal activity for each unit. Mutual information was calculated across 25 trials of cues A and B. To calculate mutual information across time, activity was averaged across ten evenly spaced time bins across the delay period, and again activity was categorically binned into ten activity levels. Maximal mutual information was 1 and 3.32 bits for the cue and time mutual information estimates, respectively.

Schur decomposition. Similar to previous studies[13,66], we performed Schur decomposition analyses on the learned recurrent weight matrices in the WM, T + WM and ISA tasks. The schur decomposition function in MATLAB was used to obtain an upper triangular matrix, representing the interaction of Schur modes. Then we plotted the number of Schur modes with least one interaction of magnitude greater than a range of threshold values (from 0 to 4).

Perturbation experiments. To test the robustness of RNN dynamics to perturbations, we introduced activity to an input unit during the delay epoch to mimic a distraction produced by an irrelevant stimulus. Specifically, the perturbation input was at π for the input topological position (corresponding to centre activation at the 17th input) with random weights to the recurrent units similar to the standard inputs. Unless otherwise specified, the perturbation was introduced with the amplitude of 0.25, 500 ms after the onset of the first stimulus for a total duration of 50 ms. Control and perturbed trajectories were obtained using the same noise matrices across units and time ('frozen' noise).

Statistics. Comparison across RNN tasks relied on the non-parametric Wilcoxon rank sum test (ranksum command in MATLAB). All statistical tests are two-sided.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Human data and code for analysis are available at https://osf.io/HXSUG/.

Code availability

Code for the RNN simulations and analysis is available at https://github.com/BuonoLab/Timing-WM_RNN_2022.

References

1. D'Esposito, M. & Postle, B. R. The cognitive neuroscience of working memory. *Annu. Rev. Psychol.* **66**, 115–142 (2015).
2. Baddeley, A. D. & Hitch, G. J. in *Psychology of Learning and Motivation* (ed. Bower G. H.) 47–89 (Academic Press, 1974).
3. Paton, J. J. & Buonomano, D. V. The neural basis of timing: distributed mechanisms for diverse functions. *Neuron* **98**, 687–705 (2018).
4. Merchant, H., Harrington, D. L. & Meck, W. H. Neural basis of the perception and estimation of time. *Annu. Rev. Neurosci.* **36**, 313–336 (2013).
5. Issa, J. B., Tocker, G., Hasselmo, M. E., Heys, J. G. & Dombeck, D. A. Navigating through time: a spatial navigation perspective on how the brain may encode time. *Annu. Rev. Neurosci.* **43**, 73–93 (2020).

6. Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron* **14**, 477–485 (1995).
7. Nobre, A. C. & van Ede, F. Anticipated moments: temporal structure in attention. *Nat. Rev. Neurosci.* **19**, 34–48 (2018).
8. Coull, J. T., Cheng, R.-K. & Meck, W. H. Neuroanatomical and neurochemical substrates of timing. *Neuropsychopharmacology* **36**, 3–25 (2011).
9. Stokes, M. G. 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn. Sci.* **19**, 394–405 (2015).
10. Lundqvist, M., Herman, P. & Miller, E. K. Working memory: delay activity, yes! Persistent activity? Maybe not. *J. Neurosci.* **38**, 7013–7019 (2018).
11. Cueva, C. J. et al. Low-dimensional dynamics for working memory and time encoding. *Proc. Natl Acad. Sci. USA* **117**, 23021–23032 (2020).
12. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsaki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
13. Goldman, M. S. Memory without feedback in a neural network. *Neuron* **61**, 621–634 (2009).
14. Rajan, K., Harvey, C. D. & Tank D. W. Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
15. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* **61**, 331–349 (1989).
16. Fuster, J. M. & Alexander, G. E. Neuron activity related to short-term memory. *Science* **173**, 652–654 (1971).
17. Fuster, J. M., Bodner, M. & Kroger, J. K. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature* **405**, 347–351 (2000).
18. Constantinidis, C. et al. Persistent spiking activity underlies working memory. *J. Neurosci.* **38**, 7020–7028 (2018).
19. Mongillo, G., Barak, O. & Tsodyks, M. Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).
20. Buonomano, D. V. & Maass, W. State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* **10**, 113–125 (2009).
21. Coull, J. T. & Nobre, A. C. Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *J. Neurosci.* **18**, 7426–7435 (1998).
22. Rainer, G., Rao, S. C. & Miller, E. K. Prospective coding for objects in primate prefrontal cortex. *J. Neurosci.* **19**, 5493–5505 (1999).
23. Matell, M. S., Meck, W. H. & Lustig, C. Not 'just' a coincidence: frontal-striatal interactions in working memory and interval timing. *Memory* **13**, 441–448 (2005).
24. van Ede, F., Niklaus, M. & Nobre, A. C. Temporal expectations guide dynamic prioritization in visual working memory through attenuated α oscillations. *J. Neurosci.* **37**, 437–445 (2017).
25. Liu, J. K. & Buonomano, D. V. Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *J. Neurosci.* **29**, 13172–13181 (2009).
26. Murray, J. M. & Escola, G. S. Learning multiple variable-speed sequences in striatum via cortical tutoring. *eLife* **6**, e26084 (2017).
27. Yang, G. R. & Wang, X.-J. Artificial neural networks for neuroscientists: a primer. *Neuron* **107**, 1048–1070 (2020).
28. Townsend, J. & Ashby, F. in *Cognitive Theory* Vol. 3 (eds Castellan, J. & Restle, F.) 200–239 (Erlbaum, 1978).
29. Nobre, A. C., Correa, A. & Coull, J. T. The hazards of time. *Curr. Opin. Neurobiol.* **17**, 465–470 (2007).
30. Taxidis, J. et al. Differential emergence and stability of sensory and temporal representations in context-specific hippocampal sequences. *Neuron* **108**, 984–998 (2020).
31. MacDonald, C. J. et al. Hippocampal 'time cells' bridge the gap in memory for discontiguous events. *Neuron* **71**, 737–749 (2011).
32. Gouvea, T. S. et al. Striatal dynamics explain duration judgments. *eLife* **4**, e11386 (2015).
33. Zhou, S., Masmanidis, S. C. & Buonomano, D. V. Neural sequences as an optimal dynamical regime for the readout of time. *Neuron* **108**, 651–658.e655 (2020).
34. Heys, J. G. & Dombeck, D. A. Evidence for a subcircuit in medial entorhinal cortex representing elapsed time during immobility. *Nat. Neurosci.* **21**, 1574–1582 (2018).
35. Shimbo, A., Izawa, E.-I. & Fujisawa, S. Scalable representation of time in the hippocampus. *Sci. Adv.* **7**, eabd7013 (2021).
36. Brody, C. D., Hernandez, A., Zainos, A. & Romo, R. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb. Cortex* **13**, 1196–1207 (2003).
37. Bae, J. W. et al. Parallel processing of working memory and temporal information by distinct types of cortical projection neurons. *Nat. Commun.* **12**, 4352 (2021).
38. Emmons, E. B. et al. Rodent medial frontal control of temporal processing in the dorsomedial striatum. *J. Neurosci.* **37**, 8718–8733 (2017).
39. Jazayeri, M., Shadlen, & Michael, N. A neural mechanism for sensing and reproducing a time interval. *Curr. Biol.* **25**, 2599–2609 (2015).
40. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297–306 (2019).
41. Zhou, S., Masmanidis, S. C. & Buonomano, D. V. Encoding time in neural dynamic regimes with distinct computational tradeoffs. *PLoS Comput. Biol.* **18**, e1009271 (2022).
42. Murray, J. D. et al. Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc. Natl Acad. Sci. USA* **114**, 394–399 (2017).
43. Hahnloser, R. H. R., Kozhevnikov, A. A. & Fee, M. S. An ultra-sparse code underlies the generation of neural sequence in a songbird. *Nature* **419**, 65–70 (2002).
44. Hardy, N. F. & Buonomano, D. V. Encoding time in feedforward trajectories of a recurrent neural network model. *Neural Comput.* **30**, 378–396 (2018).
45. Fiete, I. R., Senn, W., Wang, C. Z. H. & Hahnloser, R. H. R. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* **65**, 563–576 (2010).
46. Medina, J. F., Garcia, K. S., Nores, W. L., Taylor, N. M. & Mauk, M. D. Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J. Neurosci.* **20**, 5516–5525 (2000).
47. Laje, R. & Buonomano, D. V. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* **16**, 925–933 (2013).
48. Monteforte, M. & Wolf, F. Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Phys. Rev. X* **2**, 041007 (2012).
49. Stokes, M. G. et al. Dynamic coding for cognitive control in prefrontal cortex. *Neuron* **78**, 364–375 (2013).
50. Kim, R. & Sejnowski, T. J. Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks. *Nat. Neurosci.* **24**, 129–139 (2021).
51. Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl Acad. Sci. USA* **79**, 2554–2558 (1982).
52. Wang, X.-J. Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.* **24**, 455–463 (2001).
53. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).

54. Park, J. C., Bae, J. W., Kim, J. & Jung, M. W. Dynamically changing neuronal activity supporting working memory for predictable and unpredictable durations. *Sci. Rep.* **9**, 15512 (2019).
55. Masse, N. Y., Yang, G. R., Song, H. F., Wang, X.-J. & Freedman, D. J. Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nat. Neurosci.* **22**, 1159–1167 (2019).
56. Motanis, H., Seay, M. J. & Buonomano, D. V. Short-term synaptic plasticity as a mechanism for sensory timing. *Trends Neurosci.* **41**, 701–711 (2018).
57. Bakhurin, K. I. et al. Differential encoding of time by prefrontal and striatal network dynamics. *J. Neurosci.* **37**, 854–870 (2017).
58. Tiganj, Z., Jung, M. W., Kim, J. & Howard, M. W. Sequential firing codes for time in rodent medial prefrontal cortex. *Cereb. Cortex* **27**, 5663–5671 (2017).
59. Tsao, A. et al. Integrating time from experience in the lateral entorhinal cortex. *Nature* **561**, 57–62 (2018).
60. Kim, J., Ghim, J.-W., Lee, J. H. & Jung, M. W. Neural correlates of interval timing in rodent prefrontal cortex. *J. Neurosci.* **33**, 13834–13847 (2013).
61. Saez, A., Rigotti, M., Ostojic, S., Fusi, S. & Salzman, C. D. Abstract context representations in primate amygdala and prefrontal cortex. *Neuron* **87**, 869–881 (2015).
62. Genovesio, A., Tsujimoto, S. & Wise, S. P. Feature- and order-based timing representations in the frontal cortex. *Neuron* **63**, 254–266 (2009).
63. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016).
64. Miller, K. D. & Troyer, T. W. Neural noise can explain expansive, power-law nonlinearities in neural response functions. *J. Neurophysiol.* **87**, 653–659 (2002).
65. Rauch, A., Camera, G. L., Lüscher, 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).
66. Orhan, A. E. & Ma, W. J. A diverse range of factors affect the nature of neural representations underlying short-term memory. *Nat. Neurosci.* **22**, 275–283 (2019).
67. Ghazizadeh, E. & Ching, S. Slow manifolds within network dynamics encode working memory efficiently and robustly. *PLoS Comput. Biol.* **17**, e1009366 (2021).
68. Bernacchia, A., Fiser, J., Hennequin, G. & Lengyel, M. Adaptive erasure of spurious sequences in sensory cortical circuits. *Neuron* **110**, 1857–1868 (2022).
69. Tupikov, Y. & Jin, D. Z. Addition of new neurons and the emergence of a local neural circuit for precise timing. *PLoS Comput. Biol.* **17**, e1008824 (2021).
70. Schacter, D. L., Addis, D. R. & Buckner, R. L. Remembering the past to imagine the future: the prospective brain. *Nat. Rev. Neurosci.* **8**, 657–661 (2007).

Acknowledgements

We thank A. Tanwar for assistance with the psychophysics experiments, J. Rissman for helpful advice and J. Fuster for his comments on an earlier version of this manuscript. This research was supported by National Institutes of Health grant no. NS116589 awarded to D.V.B. and P.G. The funders had no role in study design, data collection and analysis, the decision to publish or the preparation of the paper.

Author contributions

All authors participated in the conceptualizing of the study. M.S. collected the human psychophysics data. M.S. and D.V.B. analysed the human experimental data. S.Z. and D.V.B. performed the model simulations and analysed the modelling data. All authors wrote and revised the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41562-023-01592-y.

Correspondence and requests for materials should be addressed to Dean V. Buonomano.

Peer review information *Nature Human Behaviour* thanks Nicholas Myers, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

nature portfolio

Corresponding author(s): Dean Buonomano

Last updated by author(s): Feb 28, 2023

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
☐	☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒ The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒ A description of all covariates tested
☐	☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Human behavior experiments were conducted online, with hosting provided by Gorilla (https://gorilla.sc/) and recruitment provided by Prolific (https://www.prolific.co/). All computational models were simulated using custom code of Tensorflow (2.1.0) Matlab (R2022a) at https://github.com/BuonoLab/Timing-WM_RNN_2022.
Data analysis	Unless otherwise specified, data were analyzed by using custom code of Matlab (R2022a). For SVM decoding of timing information, SVMTRAIN from LIBSVM 1.2 for Matlab was used.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

Human data and code for analysis is available at https://osf.io/HXSUG/

Human research participants

Policy information about studies involving human research participants and Sex and Gender in Research.

Reporting on sex and gender	Because we were not interested in the effects of sex or gender, no analyses were planned a priori based on sex or gender. We randomly sampled from amongst all sexes and genders without bias; however, we asked participants to self-report their gender as being Male, Female, or Other. Post hoc analyses indicated that there were no significant effects of gender.
Population characteristics	Participants were aged 18-40 (mean = 29). There were 62 female, 65 male, and 3 other gender participants. There were 5 left-handed participants and 125 right-handed participants.
Recruitment	Participants were recruited using the online Prolific platform (https://www.prolific.co/). Participants on the Prolific platform were only eligible for the study if they were between the ages of 18 and 40, residing in the United States, fluent in English, and had never participated in an online study from our laboratory on Prolific.
Ethics oversight	The study was approved by the Institutional Review Board of UCLA.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Visual working memory psychophysics studies were performed online using a within group experimental design. Cue-delay-associations and response keys were counterbalanced.
Research sample	Participants were Prolific users (https://www.prolific.co/) aged 18-40 (mean = 29). There were 62 female, 65 male, and 3 "other gender" participants. There were 5 left-handed participants and 125 right-handed participants. This sample is representative. Eligibility criteria listed in "Recruitment" were chosen to sample from among adults in a single nationality that could understand the study instructions and had no prior bias.
Sampling strategy	Sampling was random. The approximate desired sample size was pre-determined using G*Power by estimating the effect size (0.45) and setting the target power (0.9) and alpha (0.05).
Data collection	Behavioral experiments were conducted online using the Gorilla platform, which records the buttons pressed and the timing of the button presses in order to determine reaction time with respect to stimuli displayed in the browser. Before beginning the task, participants read and signed an informed consent form that asked them to: 1) complete the study in a quiet place without distractions, 2) maximize their browser window and not adjust it during the experiment, 3) have normal or corrected-to-normal vision (i.e. to wear glasses or contacts if prescribed), and 4) not participate if they had a history of seizures, epilepsy, or stroke. Participants were randomly and equally assigned to one of four counter-balance conditions, but researchers were blind to this assignment during data collection. Subjects were instructed to perform the experiments in a quiet environment, but we did not track the environment the subjects did the experiment in.
Timing	Data collection began on June 18, 2021 and completed on January 25, 2022.
Data exclusions	As described in the manuscript data from seventeen participants were excluded from analysis due to low accuracy (less than 70%) or consistently slow reaction time (RT) such that too few trials met inclusion criteria (less than 50% of possible trials in any Reversal x Delay condition remaining after RT exclusion).
Non-participation	No participants dropped out or declined participation.
Randomization	The study had no between-subjects experimental factors; however, participants were randomly and equally assigned to one of four counter-balance conditions.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging