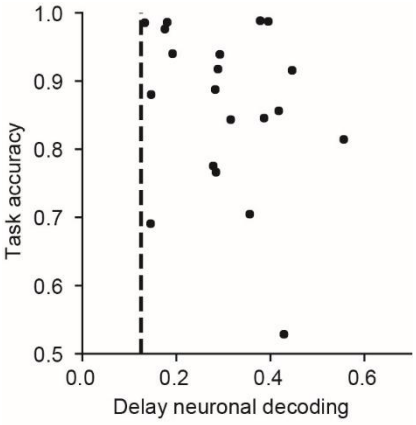

Figure S1

Figure S1 Silencing neuronal activity during at the end of the delay. To confirm that most networks maintained information solely in synaptic efficacies during the DMS task, we measured task performance on trials during which we set neuronal activity for all recurrently connected neurons to zero for the last 50 ms of the delay period. Scatter plot shows the mean neuronal decoding accuracy (measured without suppressing activity) during the last 100 ms of the delay period (x-axis) against the task accuracy after suppressing all neuronal activity during the last 50 ms of the delay period (y-axis) for all networks that solved the DMS task. Each dot represents the results from one of the 20 networks. Suppressing neuronal activity during the end of the delay period had little effect on many networks, with accuracy remaining >90% for 16 out of 20 networks. Thus, the sample direction could be silently maintained in the synaptic efficacies across the delay period for most networks, which could then be read out when the network was driven by upcoming test stimulus.

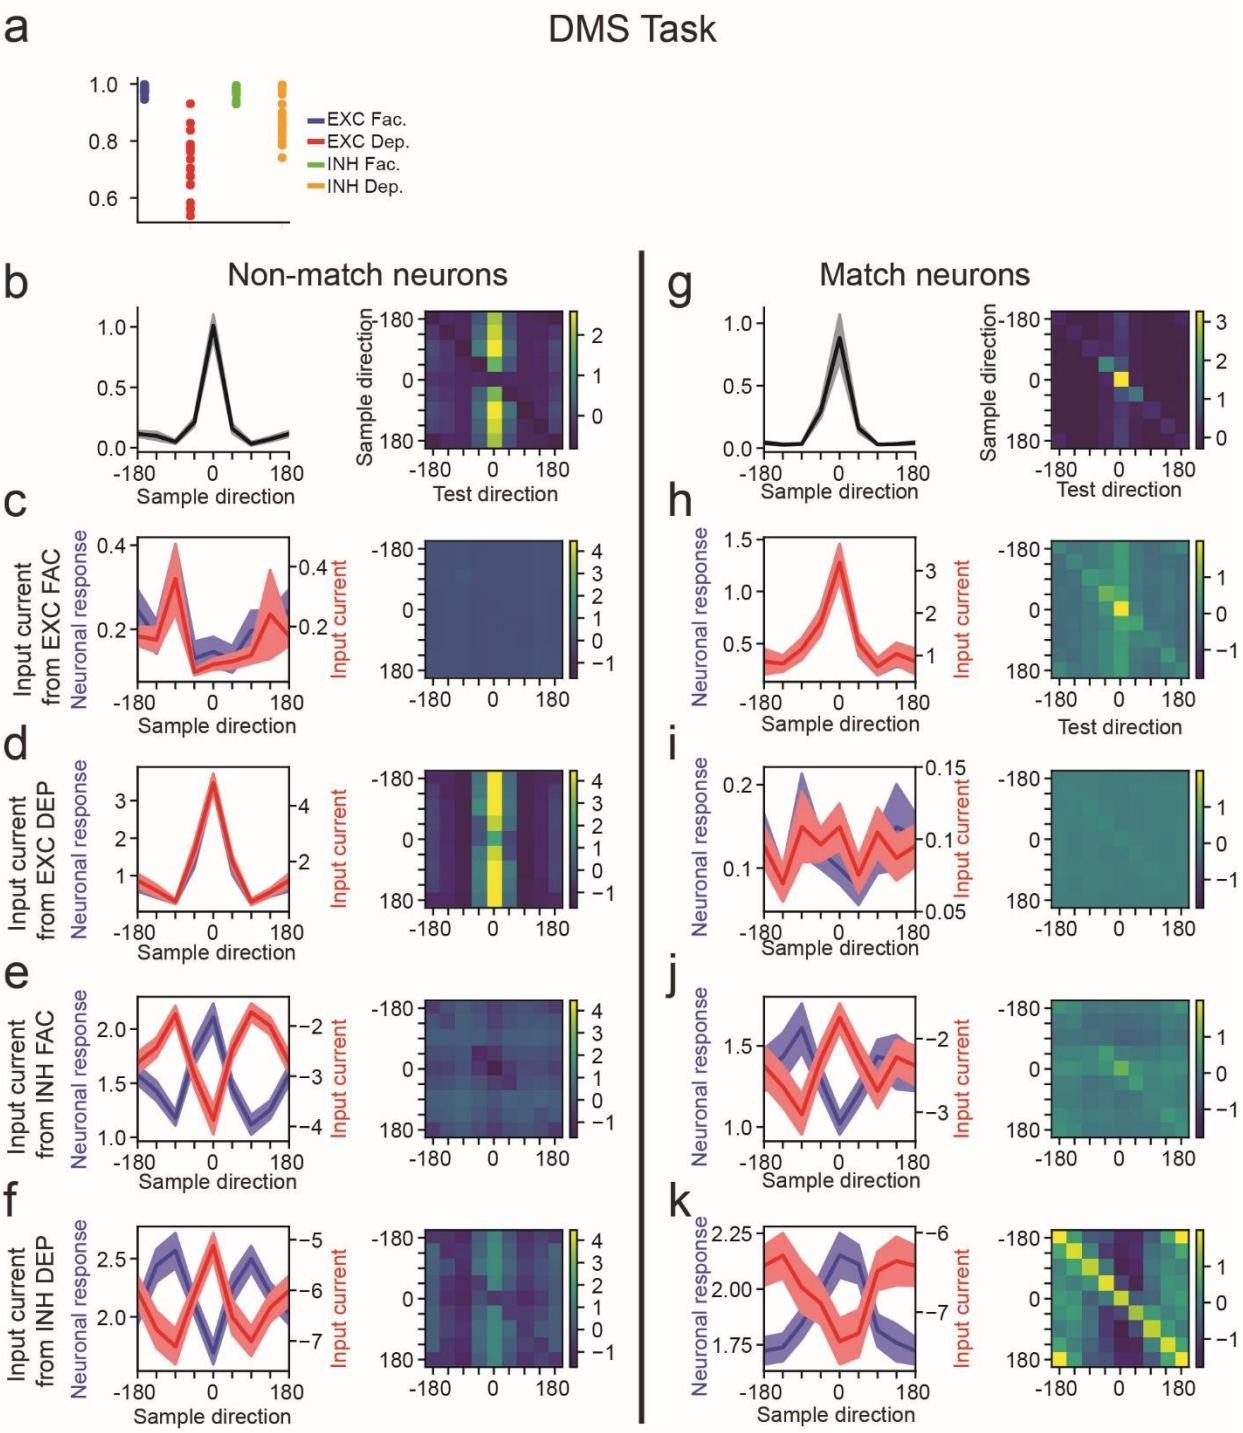

Figure S2

Figure S2 Match/non-match decision-making for the DMS task. (**a**) Task accuracy after shuffling synaptic efficacies at the end of the delay period for each of the four groups of neurons

(blue = excitatory neurons with facilitating synapses (EXC FAC); red = excitatory neurons with depressing synapses (EXC DEP); green = inhibitory neurons with facilitating synapses (INH FAC); orange = inhibitory neurons with depressing synapses (INH DEP)). When shuffling the synaptic efficacies of one group of neurons, the efficacies of the other three groups was left unchanged. Each dot represents the accuracy from one of the 20 networks. This panels suggests that synaptic efficacies from EXC DEP and INH DEP neurons were especially important for solving the task. (**b**) Left panel: mean neuronal activity averaged across the sample period for all neurons that preferentially project to the non-match output neuron (i.e. the connection weight to the non-match output neuron was greater than the connection weight to the match and the fixation output neurons). For each neuron, we determined the "preferred" sample direction that generated the greatest neuronal response, and then aligned all directions so that the preferred direction was at 0°. Right panel: heat map showing mean neuronal activity from the same group of neurons averaged across the test period for all sample-test direction pairs. The test motion direction is shown along the x-axis, and the sample direction is shown along the y-axis. Responses were aligned in the same manner as in the left panel. This panel shows that non-match projecting neurons strongly responded to test stimuli that matched its preferred sample direction (shown by bright vertical band at 0°) *except* when the test stimulus matched the sample stimulus. (**c**) Left panel: blue curve shows the mean neuronal response measured at the end of the sample stimulus, averaged across all EXC FAC neurons, weighted by the connection weights between these neurons and the non-match projecting neurons defined in (b). The red curve shows the mean synaptic current (neuronal activity X connection weight X synaptic efficacy) from EXC FAC neurons onto non-match projecting neurons. Error bars show the SEM. Both the neuronal responses and the synaptic currents have been aligned to the preferred sample directions of the non-match projecting neurons. Note that the scale of the y-axis differs between panels d-f, and shows that the mean neuronal response of EXC FAC neurons was weaker than the other three neuronal groups. Right panel: heat map showing the mean synaptic current flowing into the non-match projecting neurons from EXC FAC neurons for all sample/test pairs. For visualization purposes, the mean synaptic current was centered so that its mean was 0. Given that the mean synaptic current across all sample/test pairs was fairly uniform, this panel suggests that EXC FAC had relatively little impact upon the neuronal responses of non-match projecting neurons. (**d**) Same as (c), except for EXC DEP neurons. The left panel suggests that EXC DEP neurons preferentially connect to non-match projecting neurons if their preferred stimulus directions align. Thus, when the test stimulus is presented at the preferred stimulus direction of the non-match neurons, they will receive strong excitation from the EXC DEP neurons (shown by the vertical yellow band at 0° in the right panel). However, when the sample and test stimulus were both presented at the preferred stimulus direction (and thus was a match trial), EXC DEP neurons strongly respond to the sample stimulus, decreasing the efficacy of their synapses, reducing the excitation they project onto the non-match projecting neuron. Thus, EXC DEP decrease their effective excitation of non-match projecting neurons during match trials. (**e-f**) Same as (c), except for INH FAC and INH DEP neurons, respectively. (**g-k**) Same as (b-f), except for neurons that preferentially project to the match output neurons. For the INH DEP neurons (k), the left panel suggests that these neurons preferentially connect to match projecting neurons if their preferred stimulus directions align. Thus, when the test stimulus is presented at the preferred stimulus direction of the match

neurons, they will receive strong inhibition from the INH DEP neurons (shown by the vertical dark blue band at 0° in the right panel). However, when the sample and test stimuli were both presented at the preferred stimulus direction (and thus was a match trial), INH DEP neurons strongly respond to the sample stimulus, decreasing the efficacy of their synapses, reducing the inhibition they project onto the match output unit. Thus, INH DEP decrease their effective inhibition of match projecting neurons during match trials. (**b-k**) Tuning curves are mean values across n = 20 networks. Shaded error bars indicate one SEM.

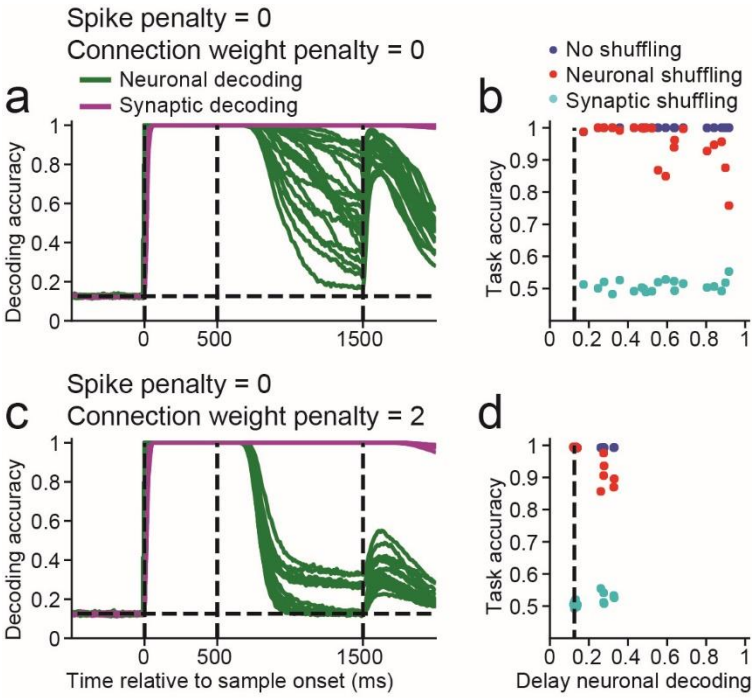

Figure S3

Figure S3 Varying the spike and connection weight penalties. (a) Sample decoding accuracy, calculated using neuronal activity (green curves) or synaptic efficacy (magenta curves) for n = 20 networks that successfully solved the DMS task with spike cost set to zero. The dashed vertical lines, from left to right, indicate the sample onset, offset, and end of the delay period. **(b)** Scatter plot showing the neuronal decoding accuracy measured at the end of the delay (x-axis) versus the task accuracy (y-axis) for all 20 networks trained on this task (blue circles), the task accuracy for the same 20 networks after neuronal activity was shuffled right before test onset (red circles) or synaptic efficacies were shuffled right before test onset (cyan circles). Thus, for each blue circle, there is a corresponding red and cyan circle with matching neuronal decoding accuracy. The dashed vertical line indicates chance level decoding accuracy. (**c&d**) Similar (a&b), except that results from 20 networks trained to solve the DMS task with a penalty on connection weights (see Methods) and zero spike cost are shown.

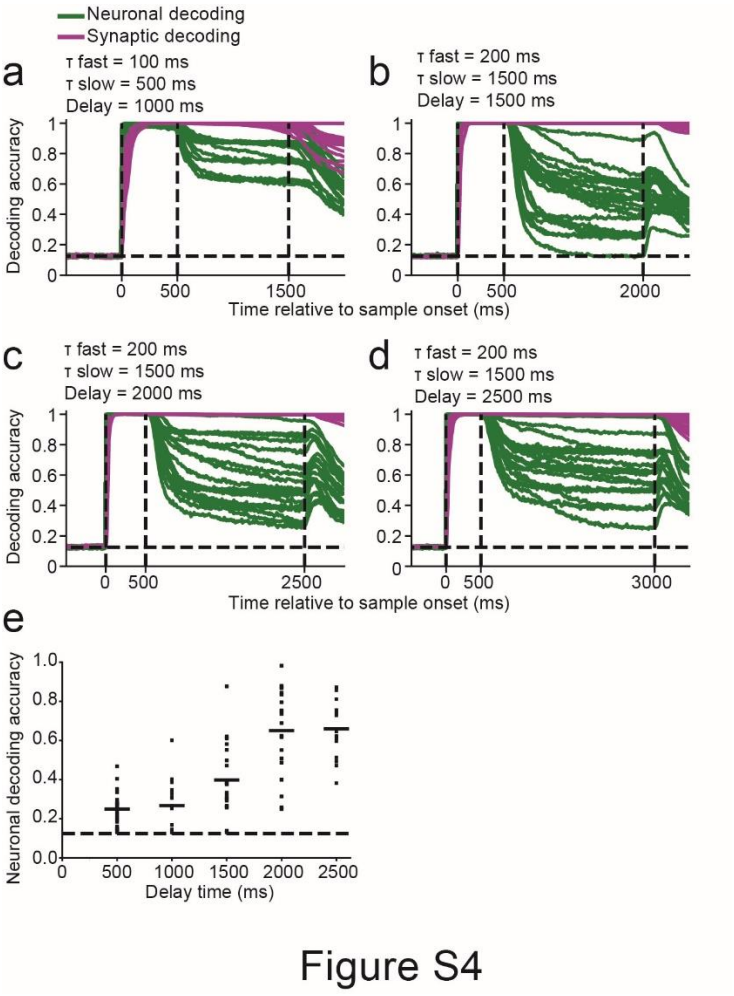

Figure S4

Figure S4 Varying delay times and synaptic time constants. (a) Sample decoding accuracy, calculated using neuronal activity (green curves) or synaptic efficacy (magenta curves for n = 20 networks that successfully solved the DMS task with the STSP time constants of 100 and 500 ms. Neuronal decoding accuracy was significantly greater than chance (P>0.05, bootstrap) during the last 100 ms of the delay period for all 20 networks. **(b-d)** Same as (a), with STSP time constants set to 200 and 1500 ms, but with a delay period of 1500 (b), 2000 (c), and 2500 ms (d), respectively. Neuronal decoding accuracy was not significantly greater than chance (P>0.05, bootstrap) during the last 100 ms of the delay period for 1 network out of 20 when the delay period was 1500 ms (B), but when the delay period was 2000 or 2500 ms, neuronal decoding accuracy was significantly greater than chance for all networks. Thus, networks did not solve the DMS task with a delay period of 2000 ms or greater using activity-silent memory. **(e)** Scatter plot showing the neuronal decoding accuracy measuring during the last 100 ms of the delay (y-axis) versus the delay period duration of the DMS task (x-axis). Each dot represents the values for one of the 20 networks trained for each delay period duration. The horizontal bars for each delay period duration indicate the mean neuronal decoding accuracy across the 20 networks. The dashed line indicates chance level performance.

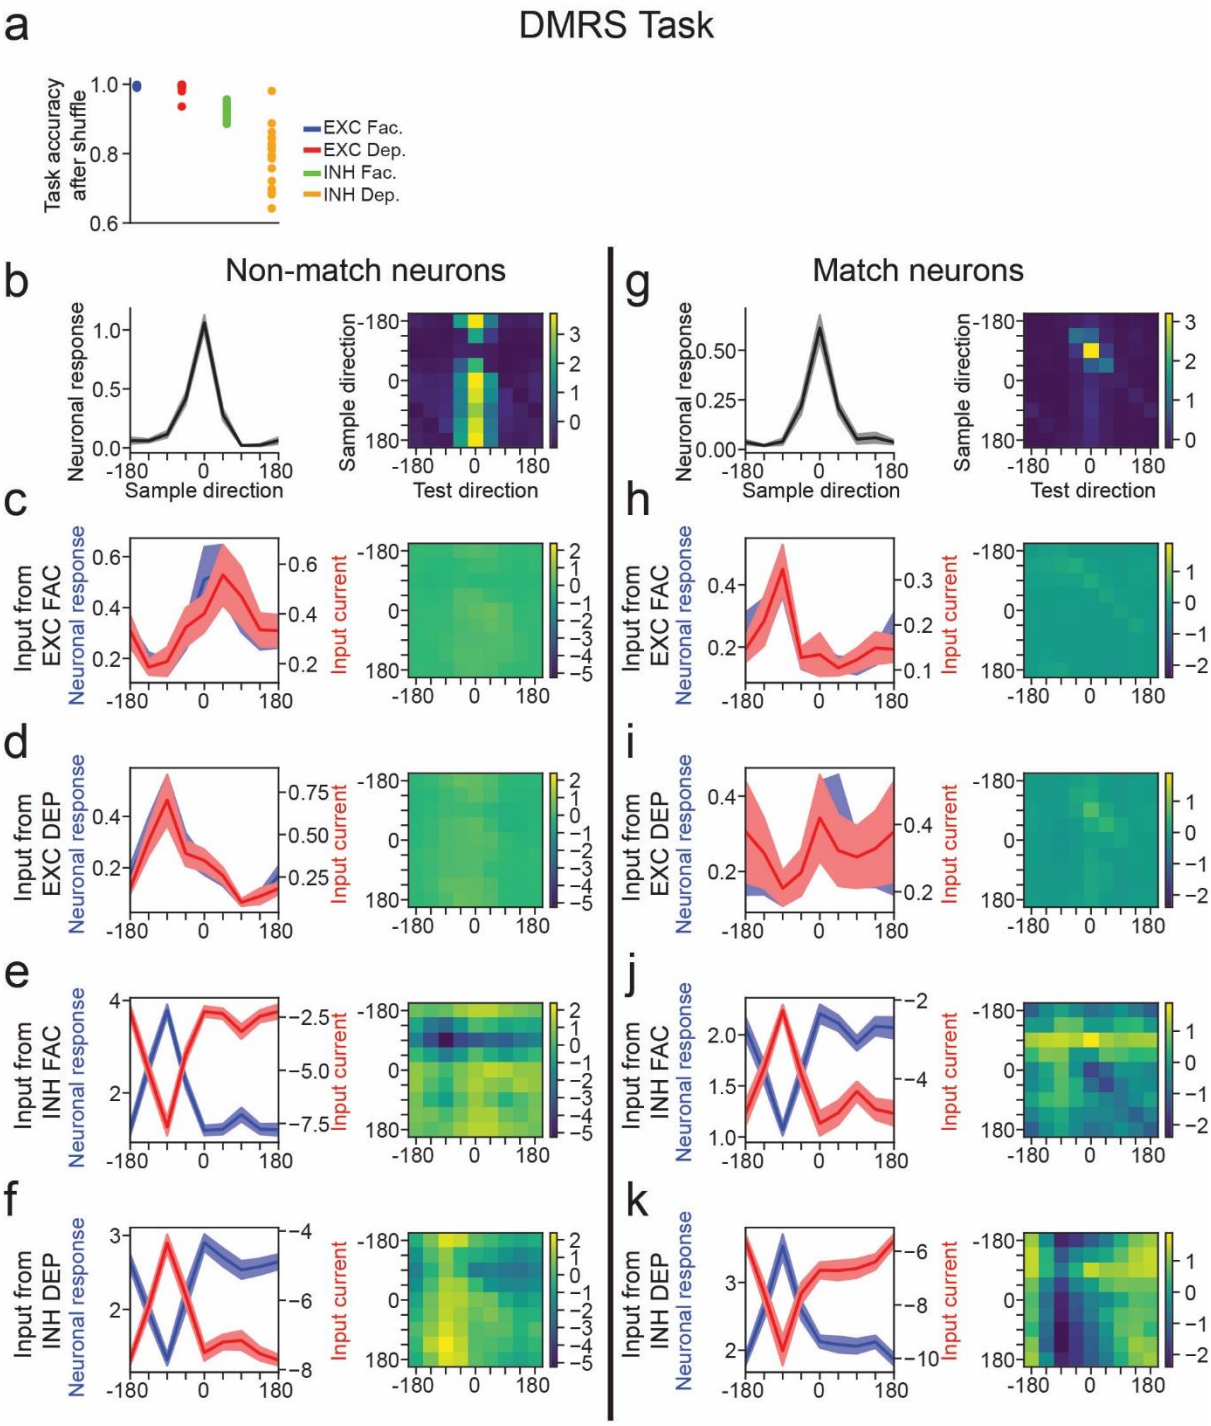
a
DMRS Task
Task accuracy after shuffle
EXC Fac.
EXC Dep.
INH Fac.
INH Dep.
b
Non-match neurons
Neuronal response
Sample direction
Test direction
g
Match neurons
c
h
Input from EXC FAC
d
i
Input from EXC DEP
e
j
Input from INH FAC
f
k
Input from INH DEP
Neuronal response
Input current

Figure S5

Figure S5 Match/non-match decision-making for the DMRS task. (**a**) Task accuracy after shuffling synaptic efficacies at the end of the delay period for each of the four groups of neurons (blue = EXC FAC; red = EXC DEP; green = INH FAC; orange = INH DEP). This panels suggests that synaptic efficacies from INH DEP, and INH FAC neurons to a lesser extent, were especially important for solving the task. This panel differs slightly from the shuffle analysis shown in Figure 3g, in which synaptic efficacies were shuffled at the end of the sample period, instead of the end of the delay as done here. (**b**) Left panel: mean neuronal activity averaged across the sample period for all neurons that preferentially project to the non-match output neuron (i.e. the connection weight to the non-match output neuron was greater than the connection weight to the match and the fixation output neurons). For each neuron, we determined the "preferred" sample direction that generated the greatest neuronal response, and then aligned all directions so that the preferred direction was at 0°. Right panel: heat map showing mean neuronal activity from the same group of neurons averaged across the test period for all sample-test direction pairs. The test motion direction is shown along the x-axis, and the sample direction is shown along the y-axis. Responses were aligned in the same manner as in the left panel. This panel shows that non-match projecting neurons strongly responded to test stimuli that matched its preferred sample direction (shown by strong vertical band at 0°) *except* when the test stimulus was 90° clockwise from the sample stimulus (a match trial). (**c-f**) Left panel: blue curve shows the mean neuronal response measured at the end of the sample stimulus, averaged across all EXC FAC (panel c) neurons, across all EXC DEP (panel d) neurons, across all INH FAC (panel e) neurons, and across all INH DEP (panel g) neurons, respectively. Responses were weighted by the connection weights between these neurons and the non-match projecting neurons defined in (b). The red curve shows the mean synaptic current (neuronal activity X connection weight X synaptic efficacy) from each neuronal group onto non-match projecting neurons. Both the neuronal responses and the synaptic currents have been aligned to the preferred sample directions of the non-match projecting neurons. Note that the scale of the y-axis differs between panels c-f. Right panel: heat map showing the mean synaptic current flowing into the non-match projecting neurons from each neuronal group for all sample-test pairs. For visualization purposes, the mean synaptic current was centered so that its mean was 0. (**g-k**) Same as (b-f), except for neurons that preferentially project to match output neurons. (**b-k**) Tuning curves are mean values across n = 20 networks. Shaded error bars indicate one SEM.

Interpretation of Figure S5

INH DEP neurons (panels f&k), which appear to play a significant role in the match/non-match decision based on the shuffle analysis in panel a, preferentially projected onto non-match projecting neurons when their preferred stimulus directions aligned, such that the weakest INH DEP neuronal response was 90° counterclockwise. Thus, for match trials in which the test direction was presented at the preferred stimulus direction of the INH DEP neurons, neuronal responses will be at their maximum, with minimum synaptic depression, increasing the inhibition onto non-match projecting neurons. In contrast, INH DEP neurons preferentially projected onto match projecting neurons when the preferred stimulus direction of the INH DEP neuron was 90°

counterclockwise from the preferred stimulus direction of the match projecting neuron. Thus, for match trials in which the sample direction was presented at the preferred stimulus direction of the INH DEP, synaptic depression will be at its maximum, reducing the level of inhibition these neurons project onto match projecting neurons. That said, the mechanism supporting the match/non-match decision for the DMRS task appears more nuanced than those observed for the DMS task (Figure S2), and INH FAC neurons also appear to play a complimentary role in this process. While this analysis suggests how connectivity patterns from different neuronal groups support that match/non-match decision, we do not claim this analysis fully maps out all of the mechanisms involved.

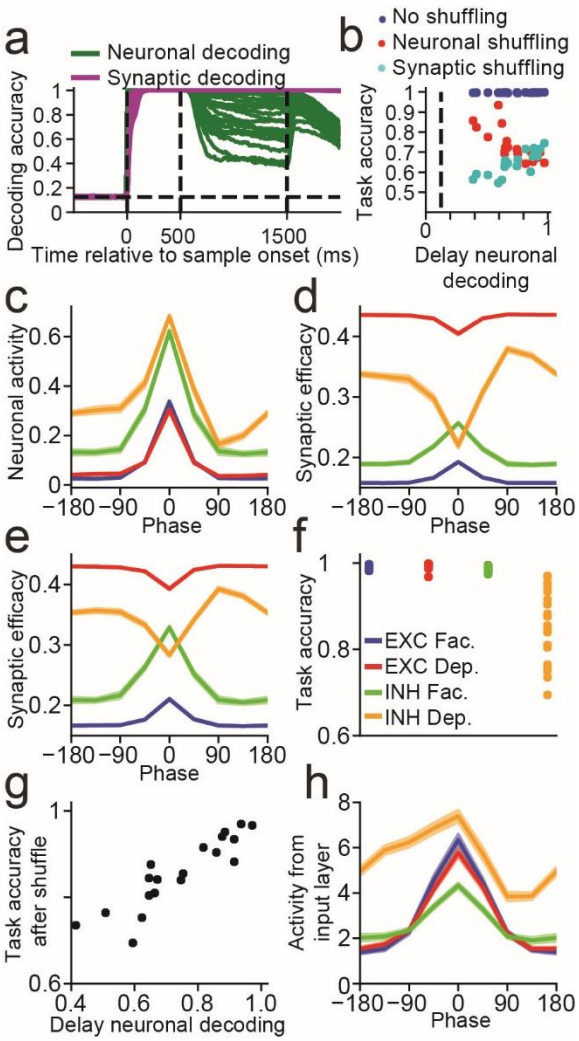

Figure S6

Figure S6 DMRS task with a 90° counterclockwise rotation. Same as Figure 3, except that networks were trained to solve a DMRS task in which the matching test stimulus was rotated 90° counterclockwise from the sample direction. Results are similar except now INH DEP neuronal tuning is asymmetric in that sample stimuli 90º *clockwise* of the neuron's preferred sample direction generate the weakest response. (**c-e,h**) Tuning curves are mean values across n = 20 networks. Shaded error bars indicate one SEM.

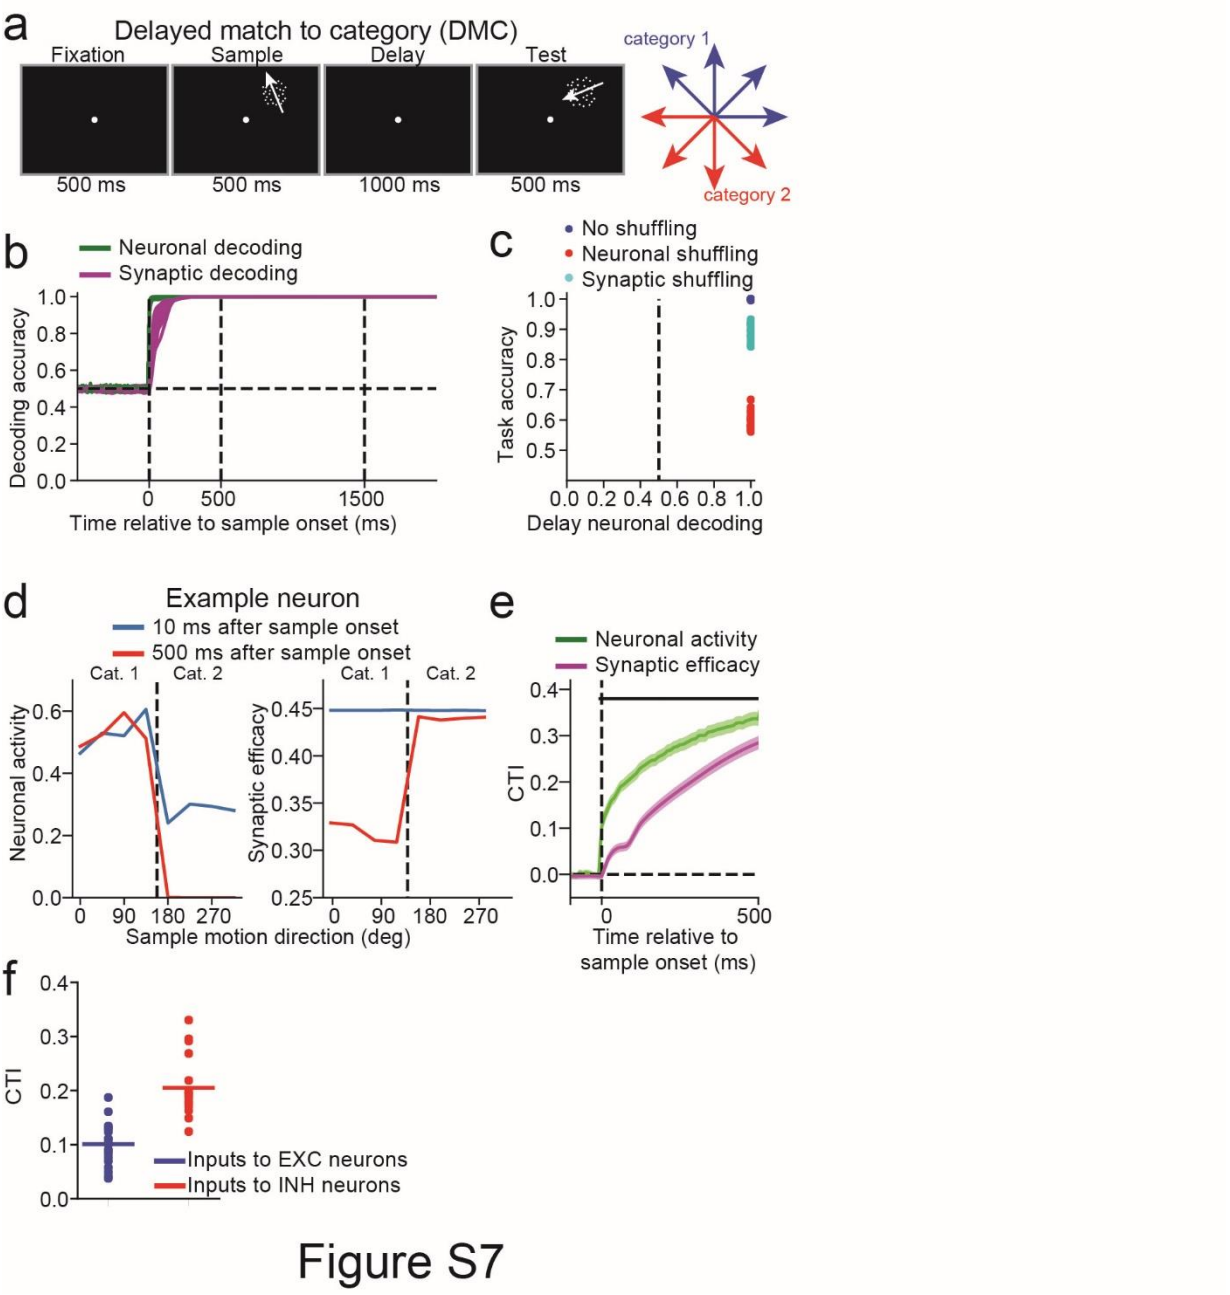

Figure S7

Figure S7 Delayed match-to-category task. (a) The delayed match-to-category task is similar to the DMS task, except that the network was trained to indicate whether the category membership of the sample and test stimuli matched. The eight motion directions were divided into two categories, indicated by the red and blue motion directions. **(b)** Sample decoding accuracy, calculated using neuronal activity (green curves) or synaptic efficacy (magenta curves) are shown for all 20 networks. The neuronal decoding accuracy, measuring during the last 100 ms of the delay period, significantly greater than chance ($P<0.01$, bootstrap) for all 20 networks. **(c)** Scatter plot showing the neuronal decoding accuracy measured at the end of the delay (x-axis) versus the task accuracy (y-axis) for all 20 networks trained on this task (blue circles), the task accuracy for the same 20 networks after neuronal activity was shuffled right

before test onset (red circles) or synaptic efficacies were shuffled right before test onset (cyan circles). **(d)** Motion direction tuning of a single example neuron showing the mean neuronal activity (left panel) or mean synaptic efficacy (right panel) as a function of the motion direction index (x-axis). The dashed vertical line divides the motion directions into two categories. The blue curves show the tuning function calculated 10 ms after sample onset, and the red curves shows the tuning function calculated at the end of the sample period. This neuronal activity for this example neuron is category-selective 10 ms after sample onset, showing high, mostly uniform, activity for category 1 and uniformly weaker activity for category 2. The synaptic efficacies for this example neuron were not category selective 10 ms after sample onset. **(e)** The time course of the mean category-tuning index (CTI, see Methods) using neuronal activity (green curve) and synaptic efficacy (magenta curve) averaged across n = 20 networks. Shaded error bars indicate one SEM. **(f)** The mean CTI of the activity projected from the neurons in the input layer into excitatory (blue) and inhibitory neurons (red). Each dot shows the mean value for each of the 20 networks. The blue and red horizontal lines shows the mean across all 20 sessions. Both values are significantly above zero ($P < 10^{-6}$ for both, two-sided, t-test), indicating that connection weights from the input neurons are at least partially responsible for the category selectivity observed in (e).

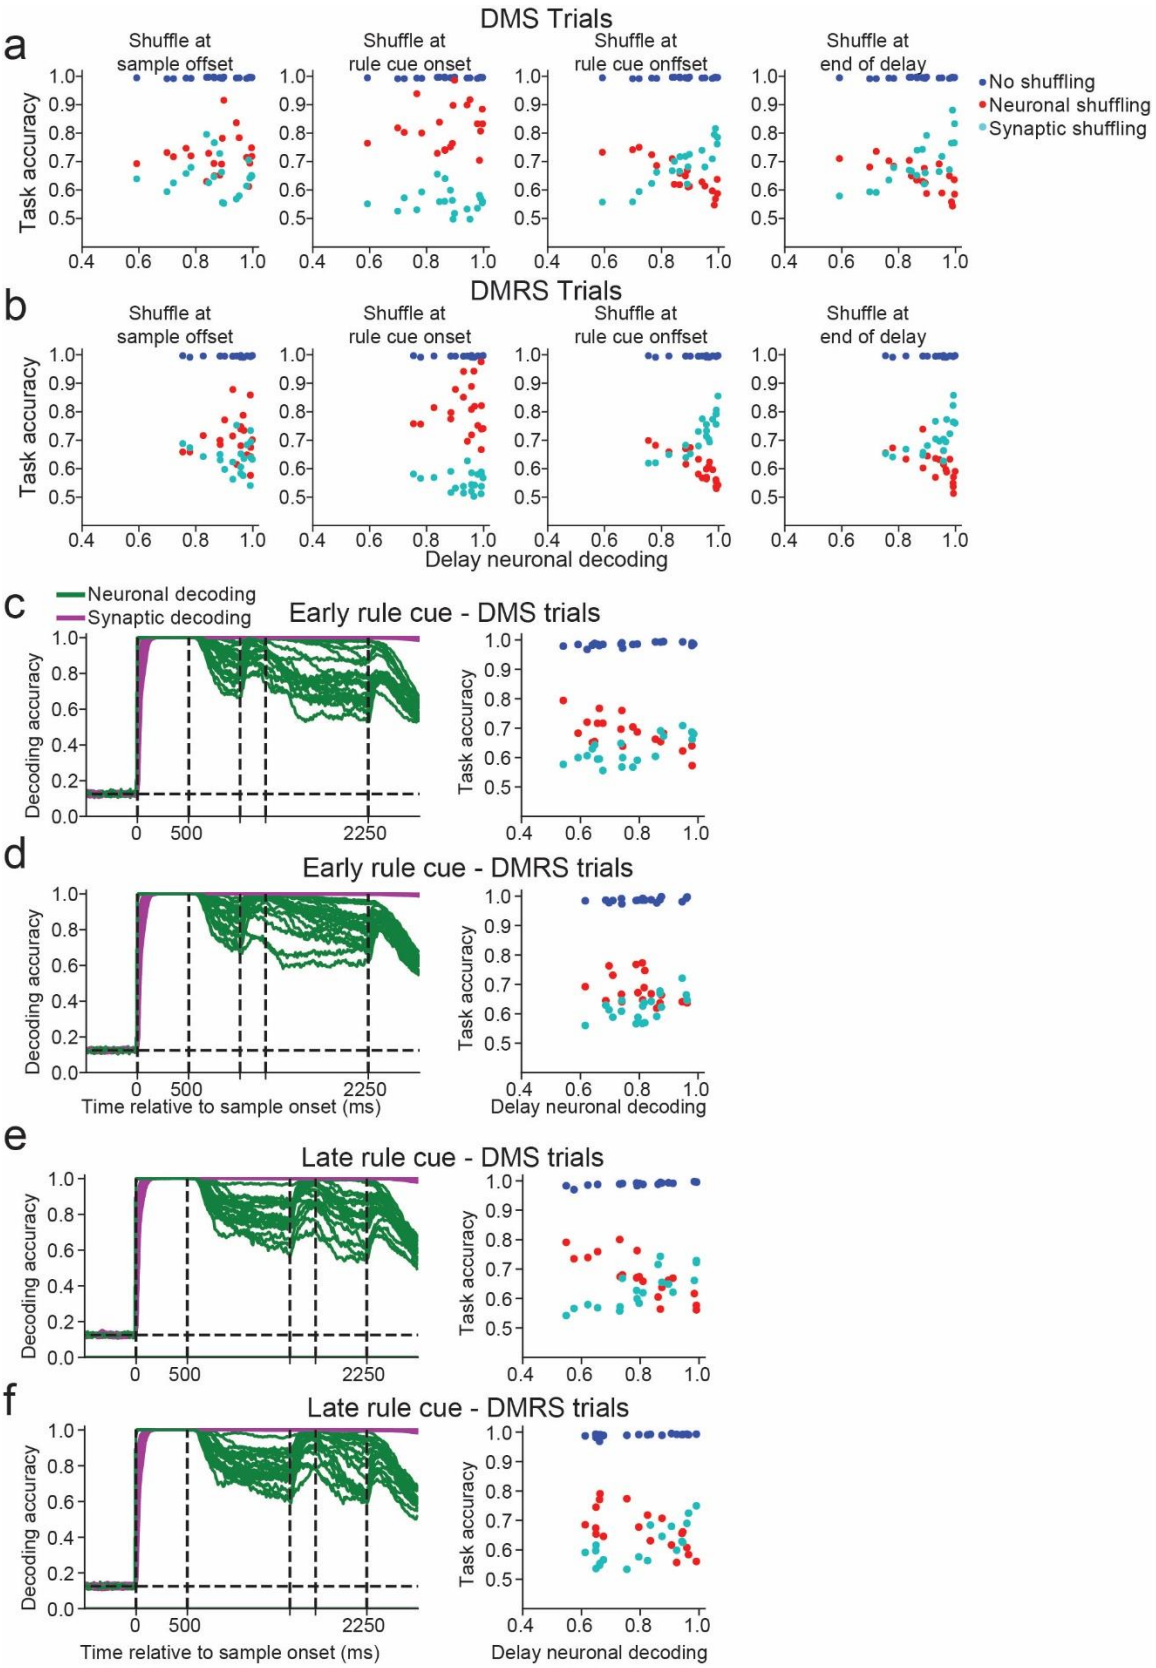
a
DMS Trials
Shuffle at sample offset
Shuffle at rule cue onset
Shuffle at rule cue onffset
Shuffle at end of delay
Task accuracy
No shuffling
Neuronal shuffling
Synaptic shuffling
b
DMRS Trials
Delay neuronal decoding
c
Neuronal decoding
Synaptic decoding
Early rule cue - DMS trials
Decoding accuracy
d
Early rule cue - DMRS trials
Time relative to sample onset (ms)
e
Late rule cue - DMS trials
f
Late rule cue - DMRS trials

Figure S8

Figure S8 Delayed rule task. (**a**) Similar to Figure 4c, except that we shuffled neuronal activity and synaptic efficacies at four different time points during DMS trials. Similar to before, each scatter plot shows the neuronal decoding accuracy measured at the end of the delay (x-axis) versus the task accuracy (y-axis) for all 20 networks trained on this task (blue circles), the task accuracy for the same 20 networks after neuronal activity (red circles) or synaptic efficacies (cyan circles) were shuffled at sample offset (left panel), rule cue onset (second from left), rule cue offset (third from left), or end of the delay (right panel). Mean accuracy after shuffling neuronal activity was 0.72, 0.82, 0.65 and 0.65 for each time point described above. Mean accuracy after shuffling synaptic efficacies was 0.65, 0.56, 0.69 and 0.70 for each time point described above. (**b**) Similar to (a), except for DMRS trials. Mean accuracy after shuffling neuronal activity was 0.71, 0.81, 0.60 and 0.61 for each time point described above. Mean accuracy after shuffling synaptic efficacies was 0.64, 0.55, 0.72 and 0.71 for each time point described above. (**c&d**) We trained 20 networks to perform a similar delayed rule task, except with a delay of 1750 ms, and in which the test onset was 1000 ms after rule offset. DMS trials shown in (c), and DMRS trials shown in (d). Left panels show sample decoding accuracies using neuronal activity (green curves), and synaptic efficacies (magenta curves) for all 20 networks. Neuronal decoding accuracy was above chance levels ($P < 0.01$, bootstrap) at all time points after sample onset for all networks. Right panels show shuffle analysis similar to (a) and (b), where all shuffling was performed at the end of the delay period. (**e&f**) We trained 20 networks to perform a similar delayed rule task, also with a delay of 1750 ms, but in which the rule cue onset was 1000 ms after the sample offset. DMS trials shown in (e), and DMRS trials shown in (f). Neuronal decoding accuracy was above chance levels ($P < 0.01$, bootstrap) at all time points after sample onset for all 20 networks.

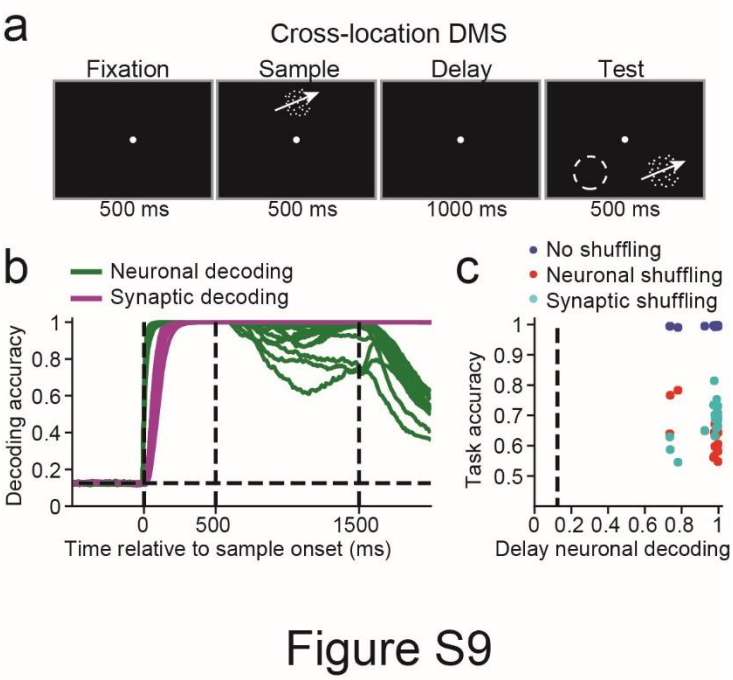

Figure S9

Figure S9 Cross-location delayed match to sample task. (**a**) Similar to the DMS task, except that the sample stimulus was presented in one spatial location (which was encoded by one group of 24 motion direction tuned input neurons), and the test stimulus was presented in one of two different spatial locations (where each location was encoded by a separate group of 24 motion direction tuned input neurons). The sample stimulus was always presented in the same location. (**b**) Sample decoding accuracy using neuronal activity (green curves) and using synaptic efficacies (magenta curves) for all 20 networks. Neuronal decoding accuracies measured at the end of the delay period was significantly greater than chance ($P < 0.01$, bootstrap) for all 20 networks. (**c**) Scatter plot showing the neuronal decoding accuracy measured at the end of the delay (x-axis) versus the task accuracy (y-axis) for all 20 networks trained on this task (blue circles), the task accuracy for the same 20 networks after neuronal activity was shuffled right before test onset (red circles) or synaptic efficacies were shuffled right before test onset (cyan circles).

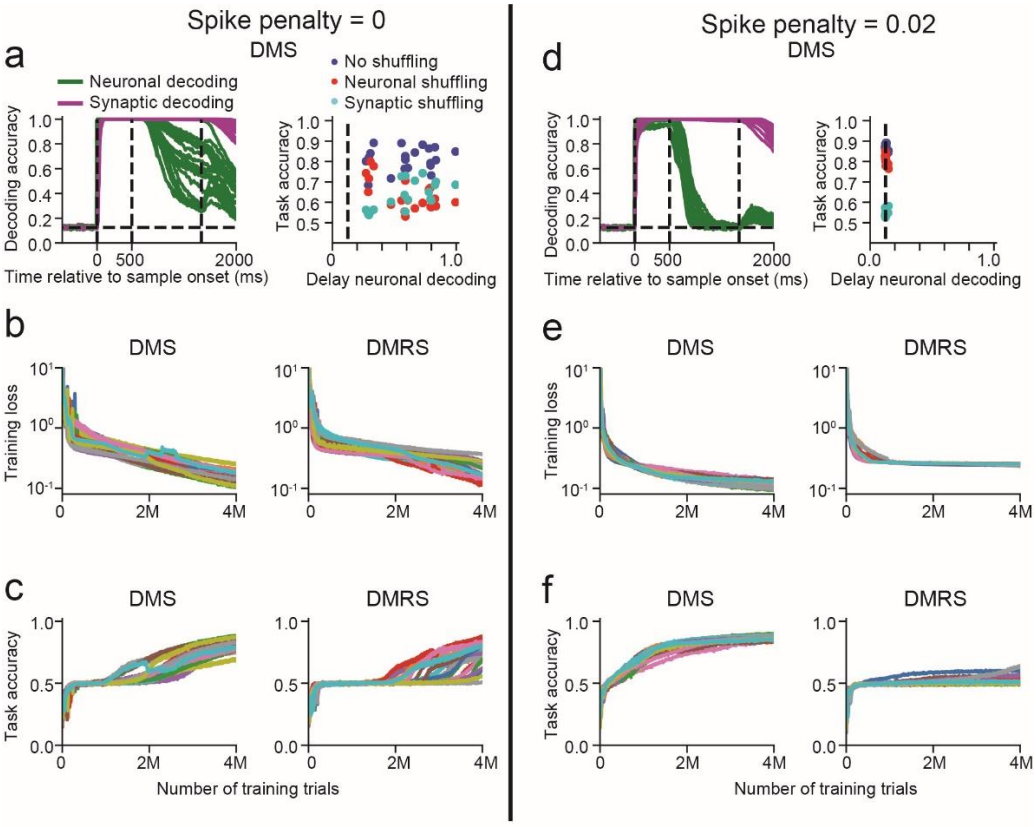

Figure S10

Figure S10 Networks in which all excitatory neurons have facilitating synapses. To facilitate learning, we increased the number of training batches from 2000 to 4000, and we lowered the learning rate from 0.02 to 0.01 which improved learning (data not shown). (**a**) Similar analysis as shown in Figure 2 and elsewhere applied to networks trained on the DMS task without a penalty on neuronal activity. The left panel shows the sample decoding accuracy using neuronal activity (green curves) and using synaptic efficacies (magenta curves) is shown for all 20 networks trained on the task. The right panel shows scatter plot showing the neuronal decoding accuracy measured at the end of the delay (x-axis) versus the task accuracy (y-axis) for all 20 networks trained on this task (blue circles), the task accuracy for the same 20 networks after neuronal activity was shuffled right before test onset (red circles) or synaptic efficacies were shuffled right before test onset (cyan circles). (**b**) The logarithm of the loss value during DMS (left panel) and DMRS training (right panel). Each trace indicates the loss value of one of the 20 networks. All curves were smoothed with a boxcar filter of size 5120 trials (5 training batches). (**c**) Task accuracy during DMS (left panel) and DMRS (right) training. 6 and 4 networks out of 20 achieved an >85% accuracy by the end of DMS and DMRS training, respectively. (**d-f**) Same as (a), except for networks trained on the DMS task with the penalty on high neuronal activity. 17 and 0 networks out of 20 achieved an >85% accuracy by the end of DMS and DMRS training, respectively.

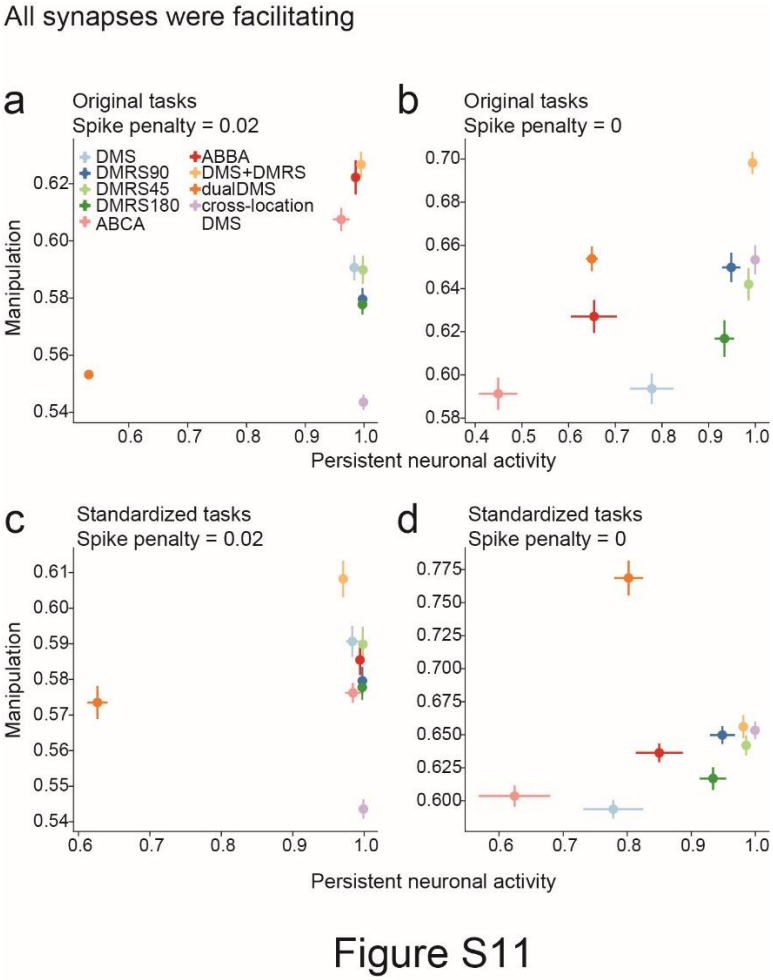

Figure S11

Figure S11 Networks in which all synapses were facilitating. The same analysis as used in Figure 7 was applied. (**a**) Scatter plot shows the level of stimulus-selective persistent neuronal activity, measured as the neuronal decoding accuracy during the last 100 ms of the delay period (x-axis), versus the level of manipulation (y-axis). Each dot represents the average values across 20 networks trained on one specific task. Persistent activity saturated near 1.0 for most tasks, except for the dual DMS task. As a result, there was no significant correlation between manipulation and persistent neuronal activity (Spearman correlation R = -0.32, P = 0.40, N = 9). (**b**) Same as (a), except that networks were trained without the penalty on high neuronal activity. There was a positive, but non-significant correlation between manipulation and persistent activity (Spearman correlation R = 0.55, P = 0.12, N = 9). (**c**) Same as (a), except that persistent activity and task manipulation were measured by presenting all networks a 500 ms motion stimulus followed by a 1000 ms delay. As with (a), persistent activity saturated near 1.0 for most tasks, except for the dual DMS task. As a result, there was no significant correlation between manipulation and persistent neuronal activity (Spearman correlation R = -0.25, P = 0.51, N = 9). (**d**) Same as (c), except for networks trained without the penalty on high neuronal activity. There was a positive, but non-significant correlation between manipulation and persistent activity (Spearman correlation R = 0.53, P = 0.14, N = 9). (**a-d**) Center of each dot represents the mean value across n = 20 networks trained on one specific task. Error bars represent one SEM.

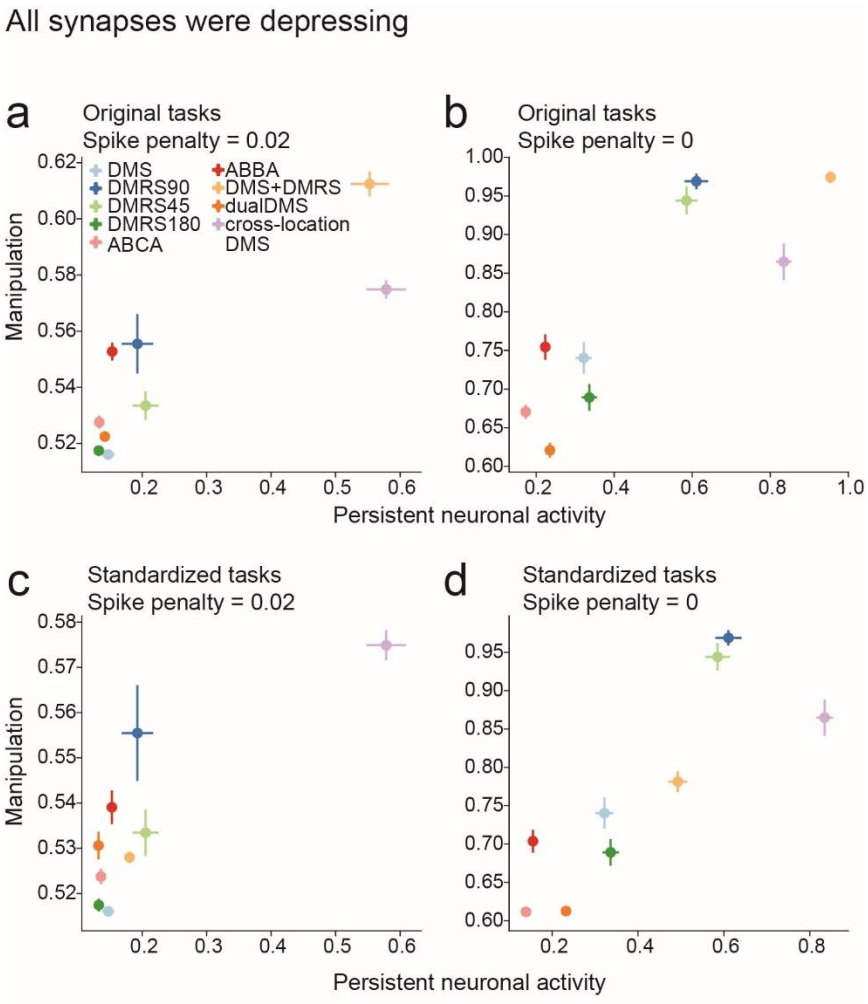

Figure S12

Figure S12 Networks in which all synapses were depressing. The same analysis as used in Figure 7 was applied. (**a**) Scatter plot shows the level of stimulus-selective persistent neuronal activity, measured as the neuronal decoding accuracy during the last 100 ms of the delay period (x-axis), versus the level of manipulation (y-axis). Each dot represents the average values across 20 networks trained on one specific task. There was a significant positive correlation between manipulation and persistent activity (Spearman correlation R = 0.82, P = 0.007, N = 9). (**b**) Same as (a), except that networks were trained without the penalty on high neuronal activity. There was also a significant positive correlation between manipulation and persistent activity (Spearman correlation R = 0.80, P = 0.01, N = 9). (**c**) Same as (a), except that persistent activity and task manipulation were measured by presenting all networks a 500 ms motion stimulus followed by a 1000 ms delay. Again, there was a significant positive correlation between manipulation and persistent activity (Spearman correlation R = 0.68, P = 0.042, N = 9). (**d**) Same as (c), except for networks trained without the penalty on high neuronal activity. There was a significant positive correlation between manipulation and persistent activity (Spearman correlation R = 0.87, P = 0.002, N = 9). (**a-d**) Center of each dot represents the mean value across n = 20 networks trained on one specific task. Error bars represent one SEM.

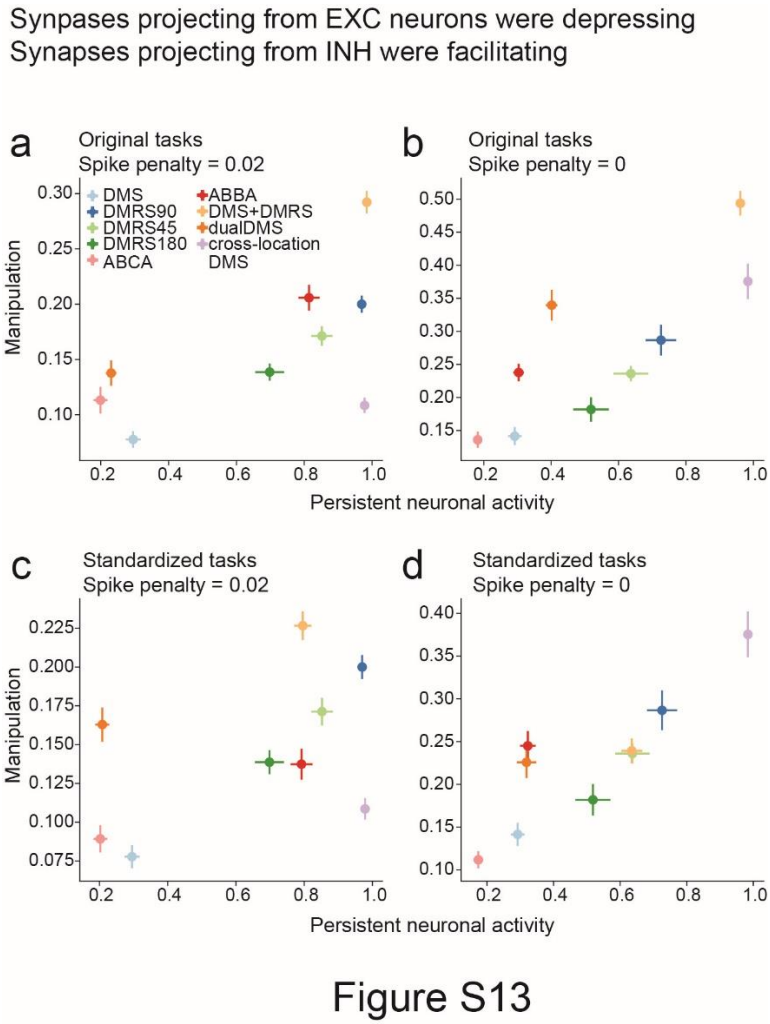

Figure S13

Figure S13 Networks in which synapses projecting from excitatory neurons were depressing, and synapses projecting from inhibitory neurons were facilitating. The same analysis as used in Figure 7 was applied. (**a**) Scatter plot shows the level of stimulus-selective persistent neuronal activity, measured as the neuronal decoding accuracy during the last 100 ms of the delay period (x-axis), versus the level of manipulation (y-axis). Each dot represents the average values across 20 networks trained on one specific task. There was a positive, but non-significant, correlation between manipulation and persistent activity (Spearman correlation R = 0.52, P = 0.15, N = 9). (**b**) Same as (a), except that networks were trained without the penalty on high neuronal activity. There was a significant positive correlation between manipulation and persistent activity (Spearman correlation R = 0.80, P = 0.01, N = 9). (**c**) Same as (a), except that persistent activity and task manipulation were measured by presenting all networks a 500 ms motion stimulus followed by a 1000 ms delay. There was a positive, but non-significant, correlation between manipulation and persistent activity (Spearman correlation R = 0.43, P = 0.24, N = 9). (**d**) Same as (c), except for networks trained without the penalty on high neuronal activity. There was a significant positive correlation between manipulation and persistent activity (Spearman correlation R = 0.85, P = 0.004, N = 9). (**a-d**) Center of each dot represents the mean value across n = 20 networks trained on one specific task. Error bars represent one SEM.

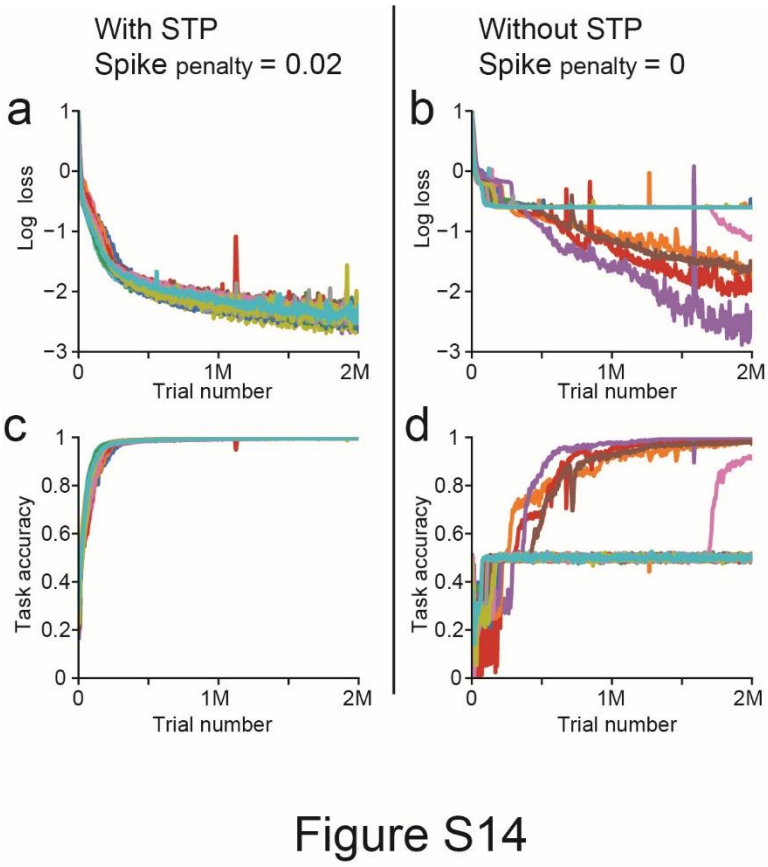

Figure S14

Figure S14 Network training with and without STSP. (a) The logarithm of the loss value during training of the networks with STSP on the DMS task. These 20 networks were the same used for Figure 2, and each trace indicates the loss value of one of the 20 networks. All curves were smoothed with a boxcar filter of size 5120 trials (5 training batches). As expected, the loss value decreased throughout training. **(b)** Similar to (a), except for networks without STSP and without a penalty on neuronal activity. The loss value initially decreases at the start of training, but then stalls for many networks. **(c)** Task accuracy during training of networks with STSP. All 20 networks successfully learned to perform the task with accuracy >98% after 400,000 trials. **(d)** Similar to (c), except for networks without STSP and without the penalty on neuronal activity. Only 5 out of the 20 networks successfully learned to perform the task with accuracy >90% after 2,000,000 trials, and 2 out of 20 obtained an accuracy > 98%.

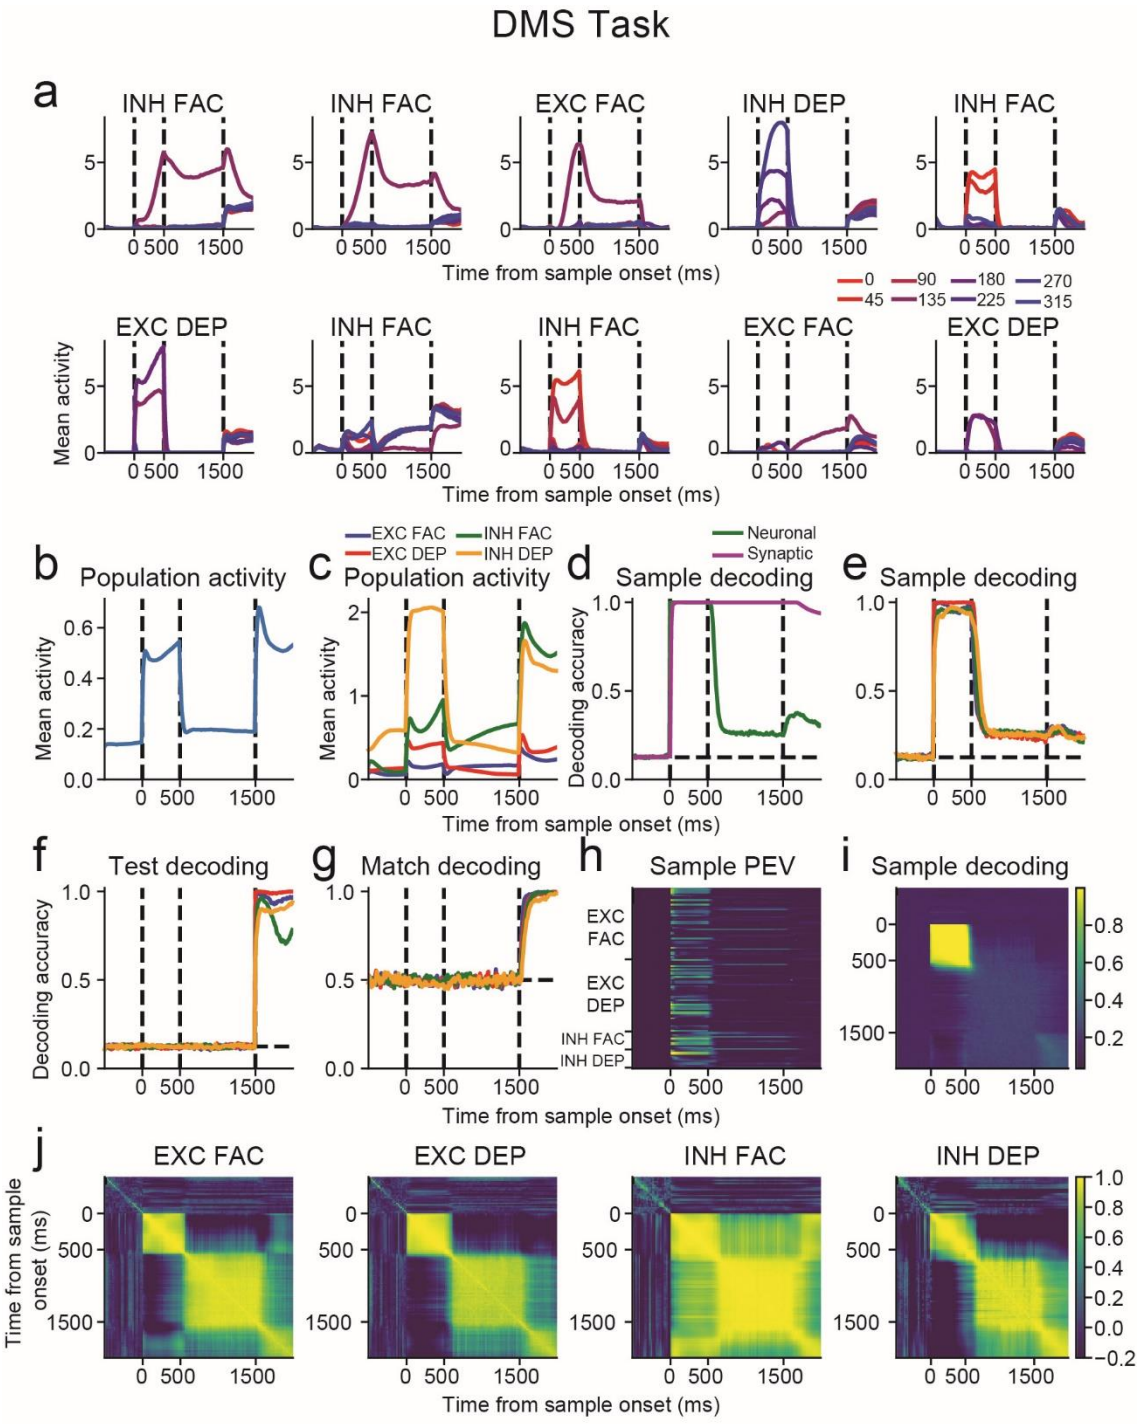

Figure S15

Figure S15 Properties of an example network trained on the DMS task. (**a**) Each panel shows the mean neuronal activity of the 10 most selective neurons (according to their proportion of explained variance, PEV (see Methods), measured between sample onset and test onset).

Each of the eight colored traces shows the mean response to each of the eight different sample directions. (**b**) Mean neuronal activity averaged across all n = 100 neurons and all sample directions. (**c**) Mean neuronal activity averaged across all sample directions for each of the four neuronal subgroups (blue, EXC FAC, n = 40 neurons; red, EXC DEP, n = 40 neurons; green, INH FAC, n = 10 neurons; orange, INH DEP, n = 10 neurons). (**d**) Sample decoding accuracy from all neurons measured using neuronal activities (green curve) or synaptic efficacies (magenta curve). (**e**) Neuronal sample decoding accuracy for each of the four neuronal subgroups. (**f**) Neuronal test decoding accuracy for each of the four neuronal subgroups. (**g**) Neuronal match/non-match decoding accuracy for each of the four neuronal subgroups. (**h**) Heat map of neuronal sample PEV values across all 100 neurons sorted by neuronal subgroup. (**i**) Neuronal sample encoding stability is shown by training and testing linear SVM classifiers at different time points. The y-axis indicates the time point from which data was used to train the classifier, and the x-axis indicates the time at which the classifier was tested. All neurons were used to train the classifier. Periods of stable encoding are represented by blocks of high decoding accuracy (i.e. high decoding accuracy when training and testing the decoder at different time points), whereas periods of more dynamic encoding are represented by high decoding accuracy solely on the diagonal (i.e. low decoding accuracy when training and testing the decoder at different time points). (**j**) Neuronal sample encoding stability is shown for the four different neuronal subgroups by calculating the sample tuning similarity index (TSI, see Methods) measured between different points in the trial. Specifically, the value at each location on the heat map is the mean dot-product between neurons' preferred sample directions, calculated at two different times in the trial indicated by the x and y-axis, weighted by their sample PEV values. As with panel i, periods of stable encoding are represented by blocks of high similarity values (i.e. preferred sample directions are similar between different points in the trial), whereas periods of more dynamic encoding are represented by high similarity values solely on the diagonal (i.e. preferred sample directions differ between different points in the trial).

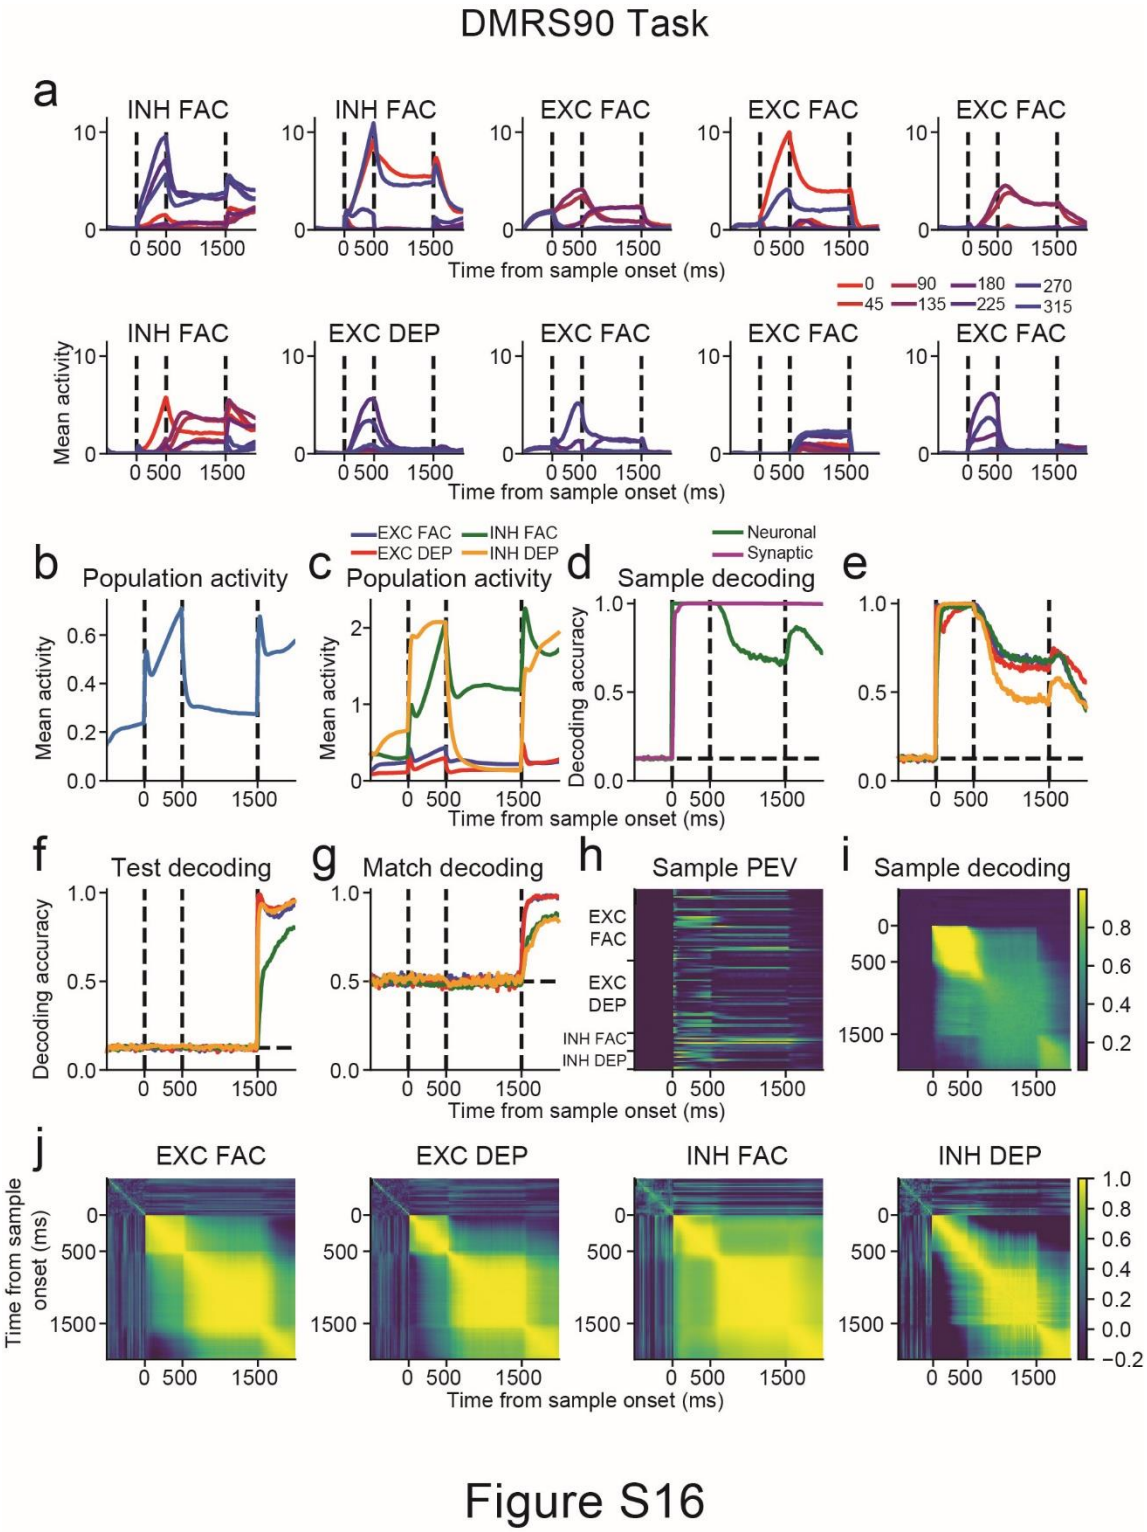

Figure S16

Figure S16 Properties of an example network trained on the DMRS90 task. Similar to Figure S15, except showing the properties of an example network that was trained on the DMRS90 task.

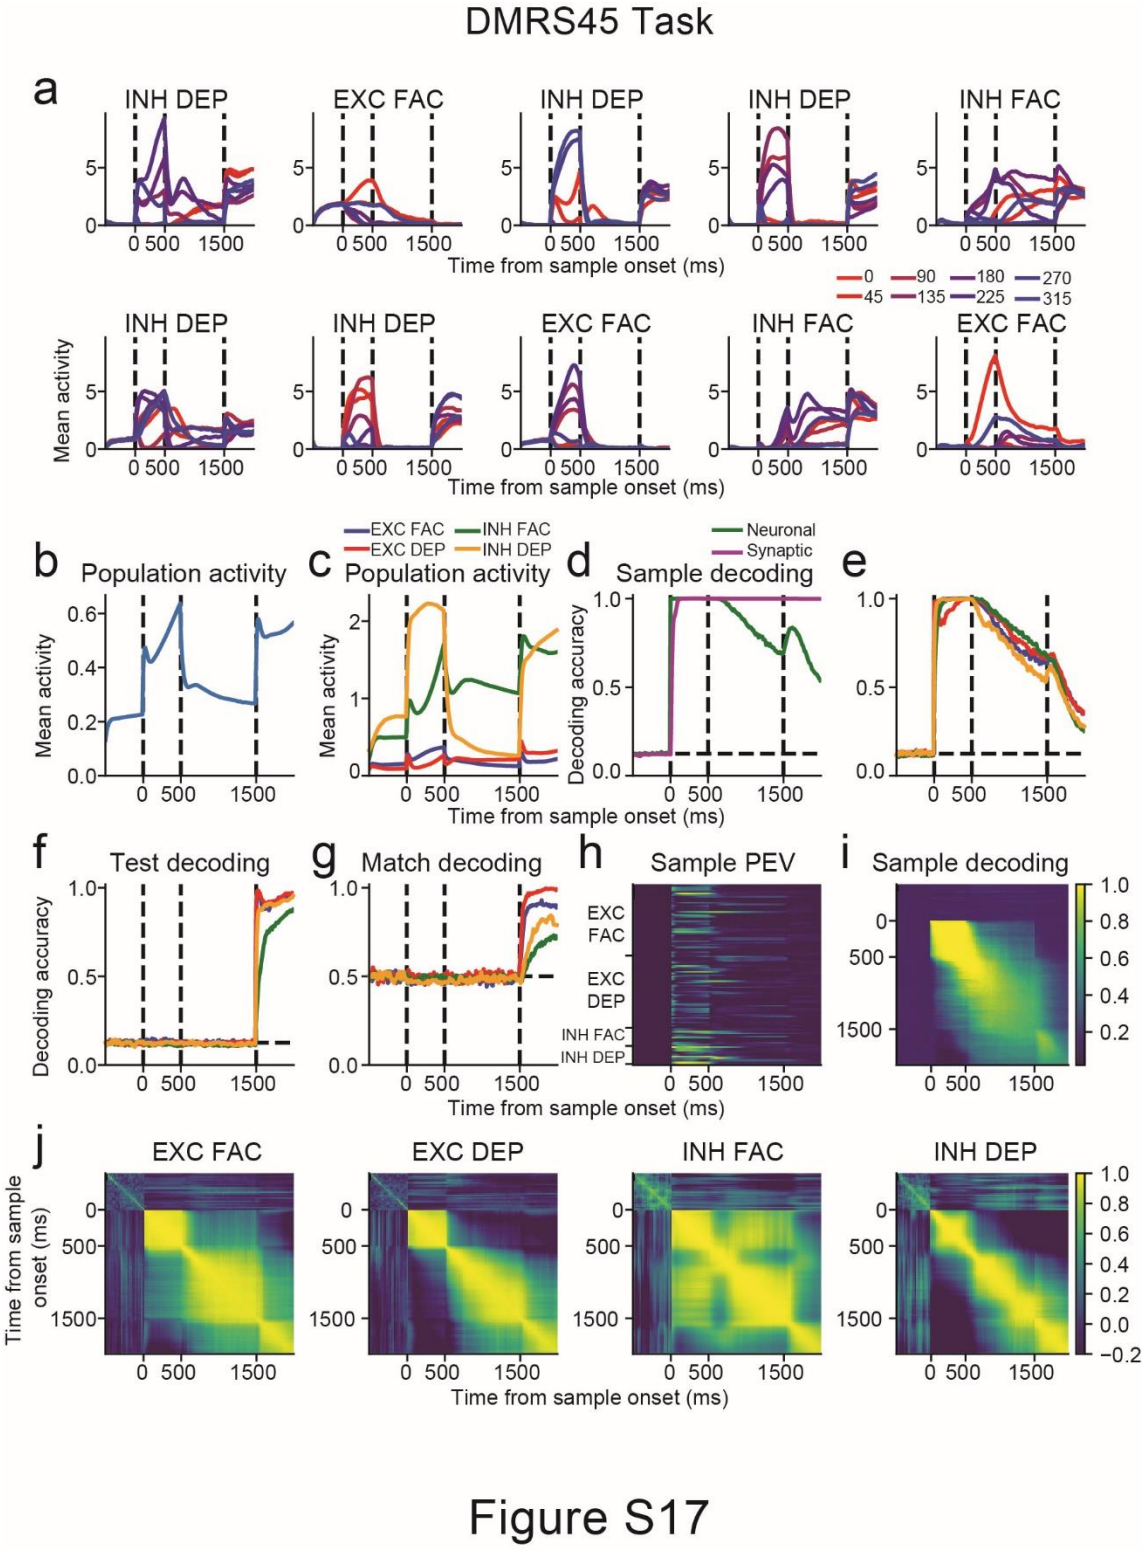

Figure S17

Figure S17 Properties of an example network trained on the DMRS45 task. Similar to Figure S15, except showing the properties of an example network that was trained on the DMRS45 task.

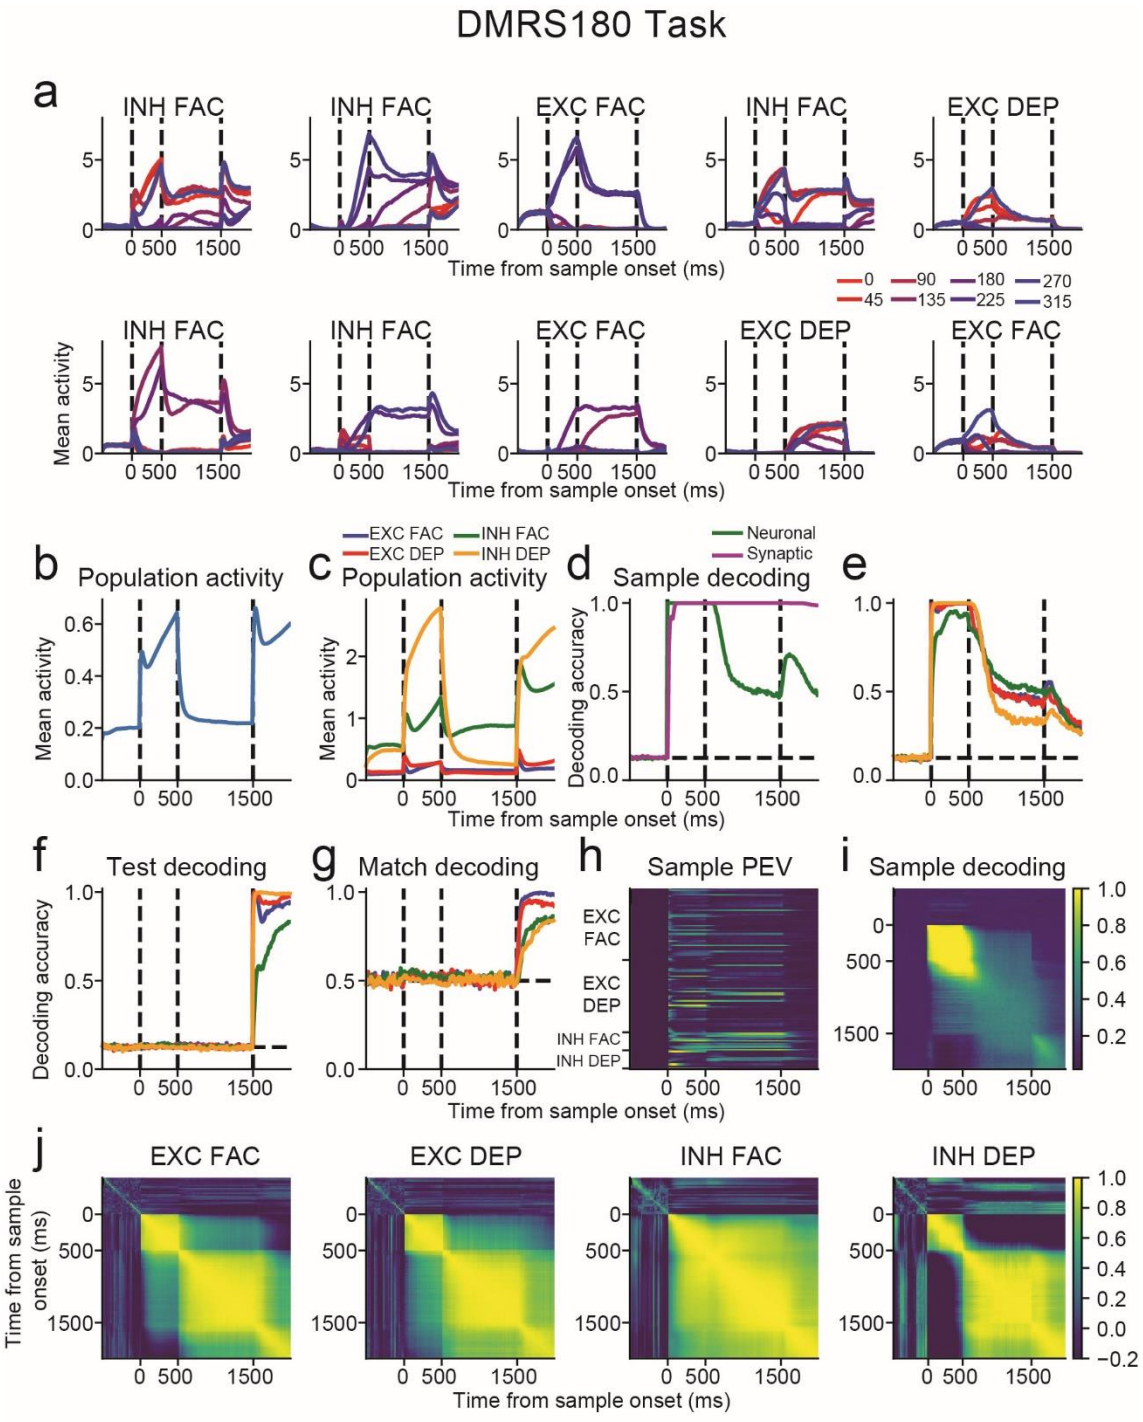

Figure S18

Figure S18 Properties of an example network trained on the DMRS180 task. Similar to Figure S15, except showing the properties of an example network that was trained on the DMRS180 task.

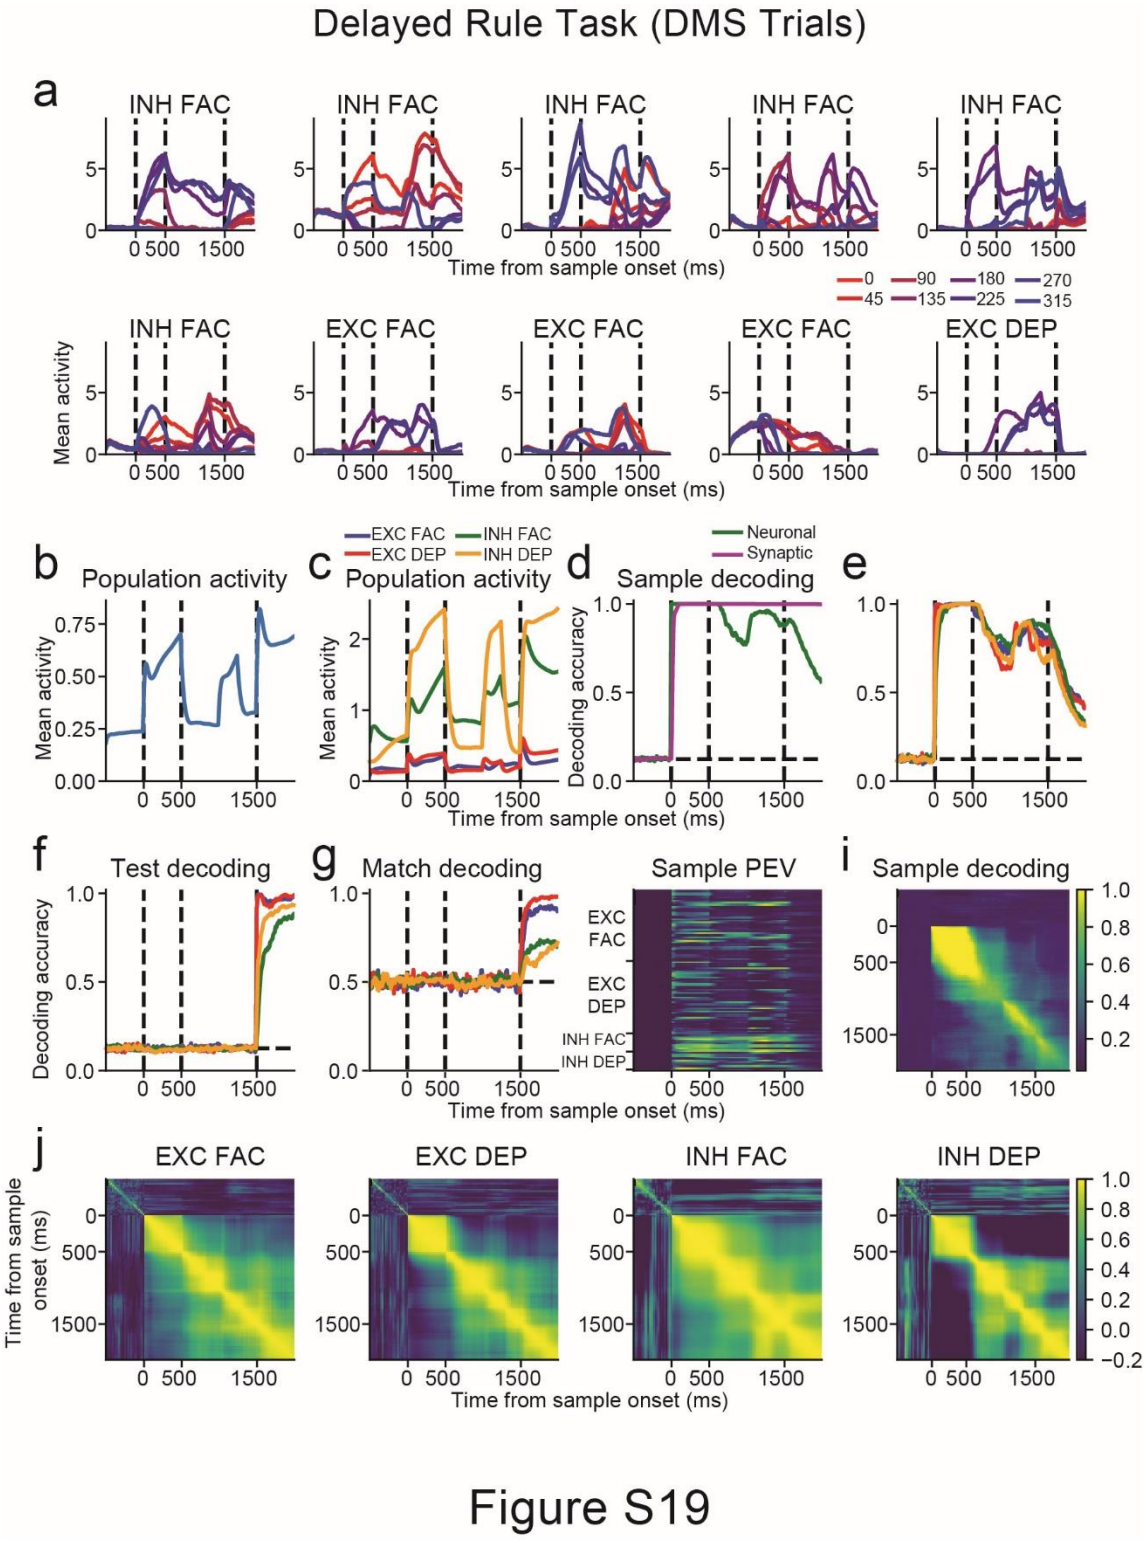

Figure S19

Figure S19 Properties of an example network trained on the delayed rule task – DMS trials only. Similar to Figure S15, except showing the properties of an example network that was trained on the delayed rule task. Showing results for DMS trials only.

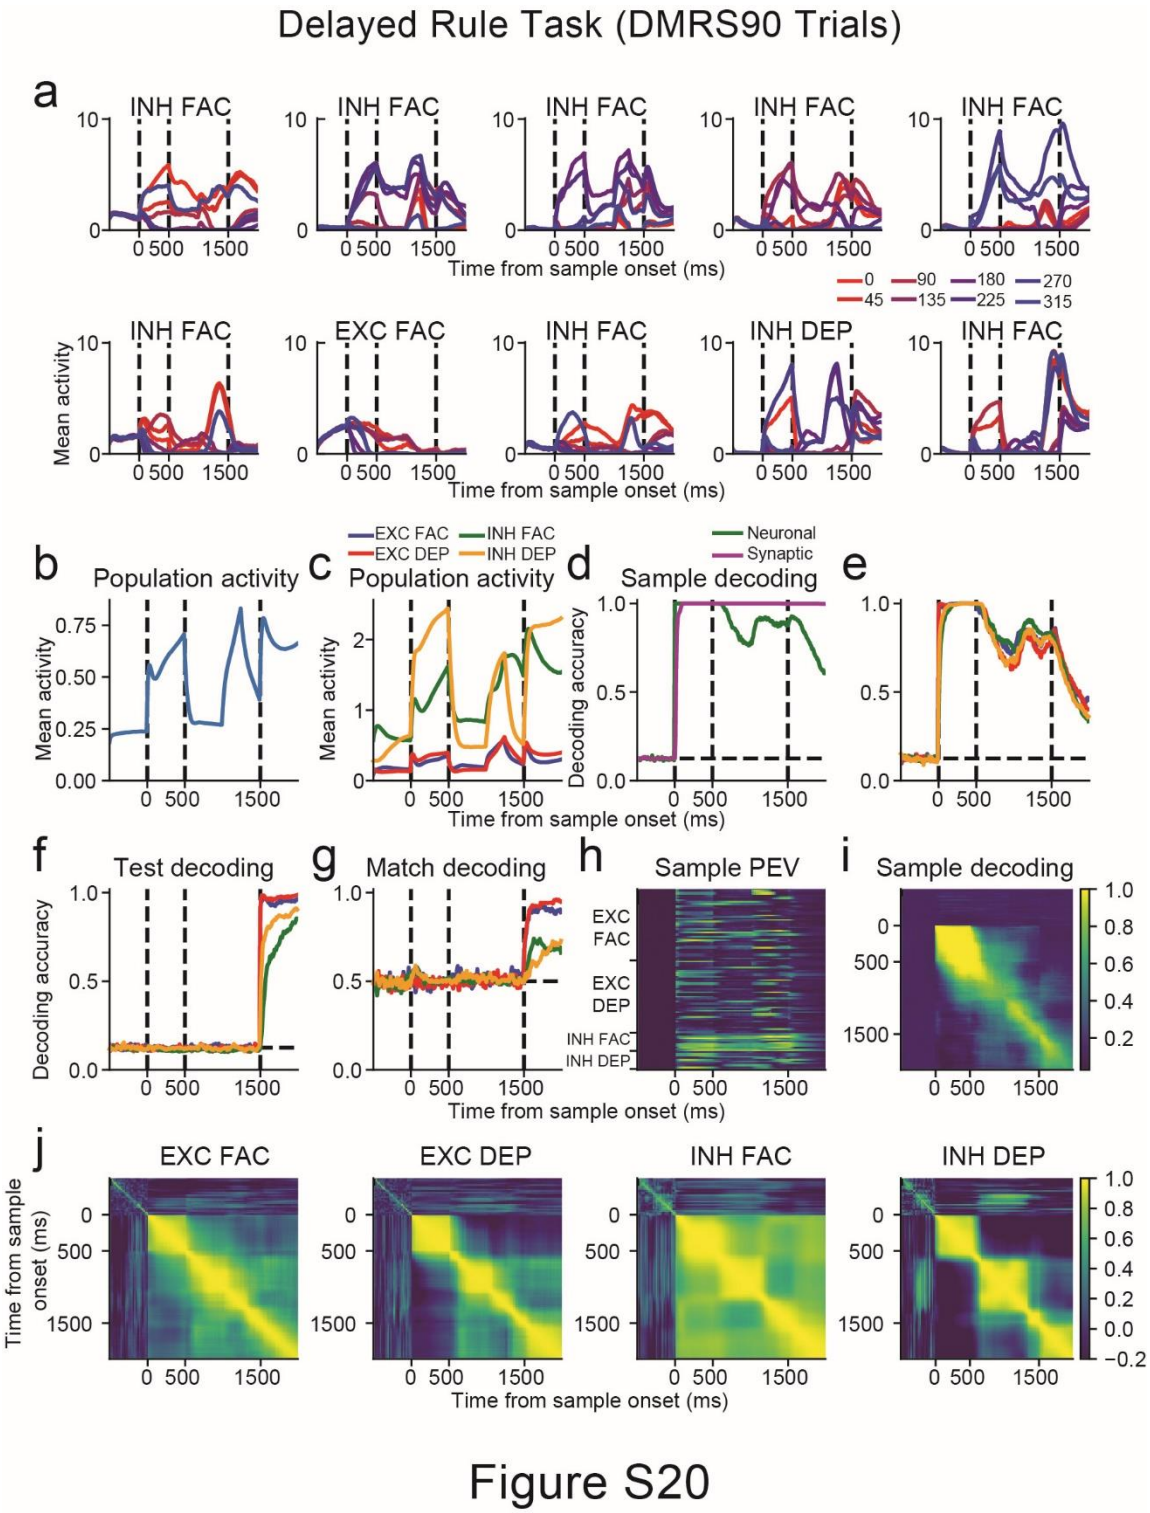

Figure S20

Figure S20 Properties of an example network trained on the delayed rule task – DMRS90 trials only. Similar to Figure S15, except showing the properties of an example network that was trained on the delayed rule task. Showing results for DMRS90 trials only.

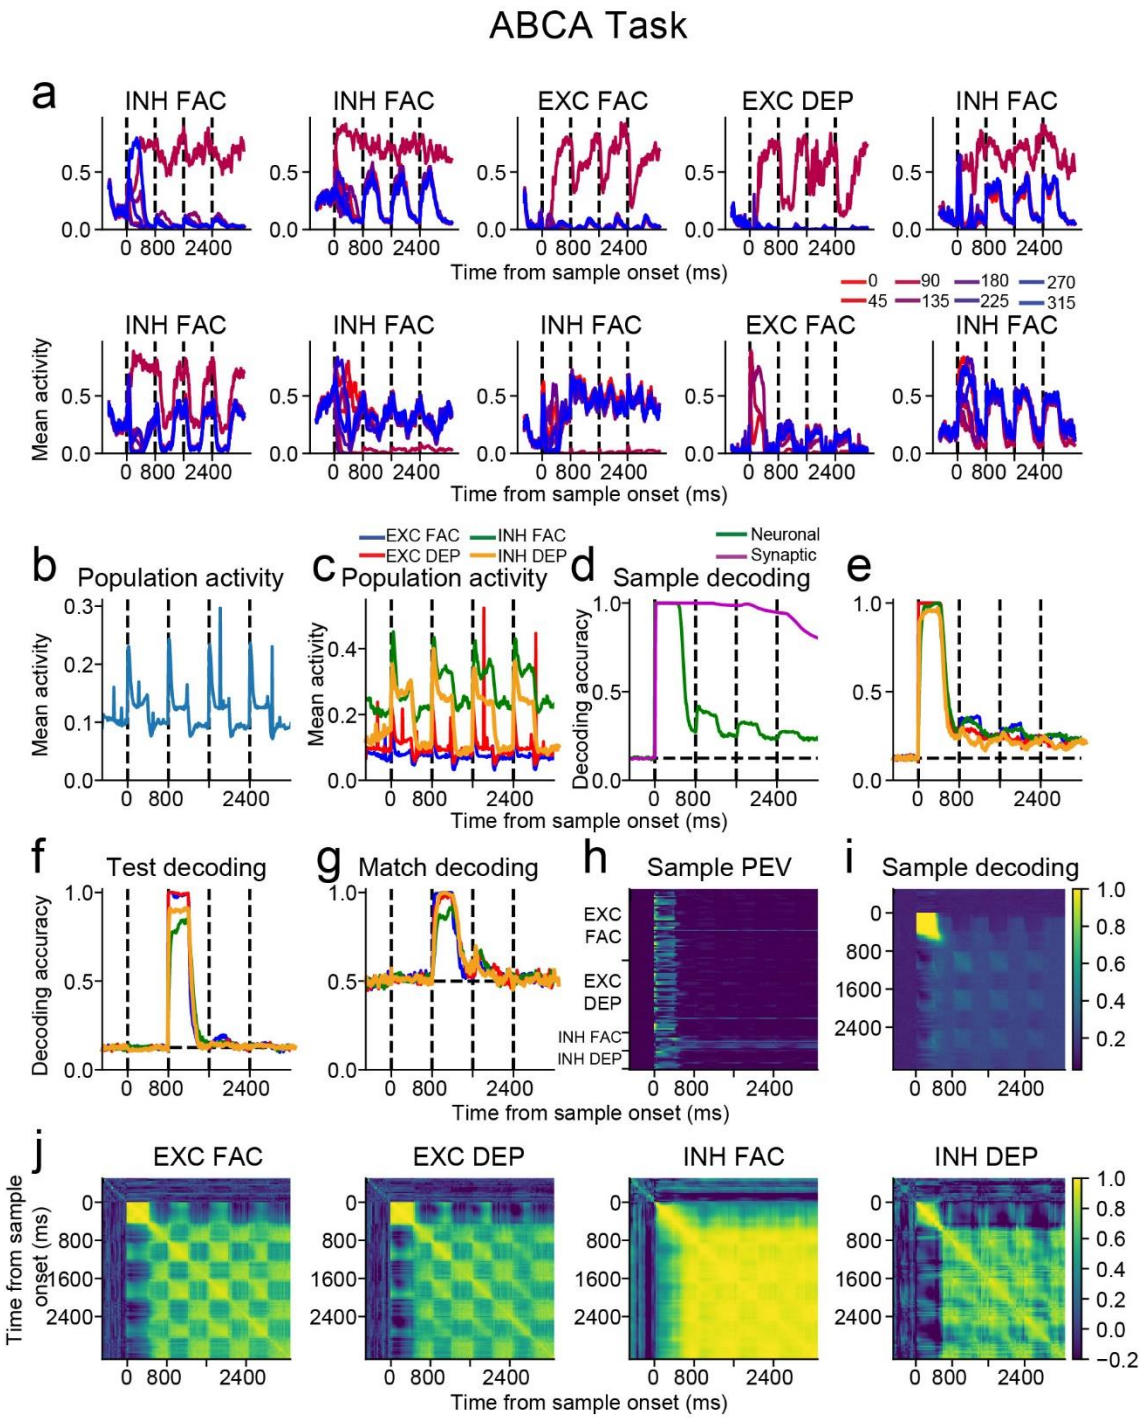

Figure S21

Figure S21 Properties of an example network trained on the ABCA task. Similar to Figure S15, except showing the properties of an example network that was trained on the ABCA task.

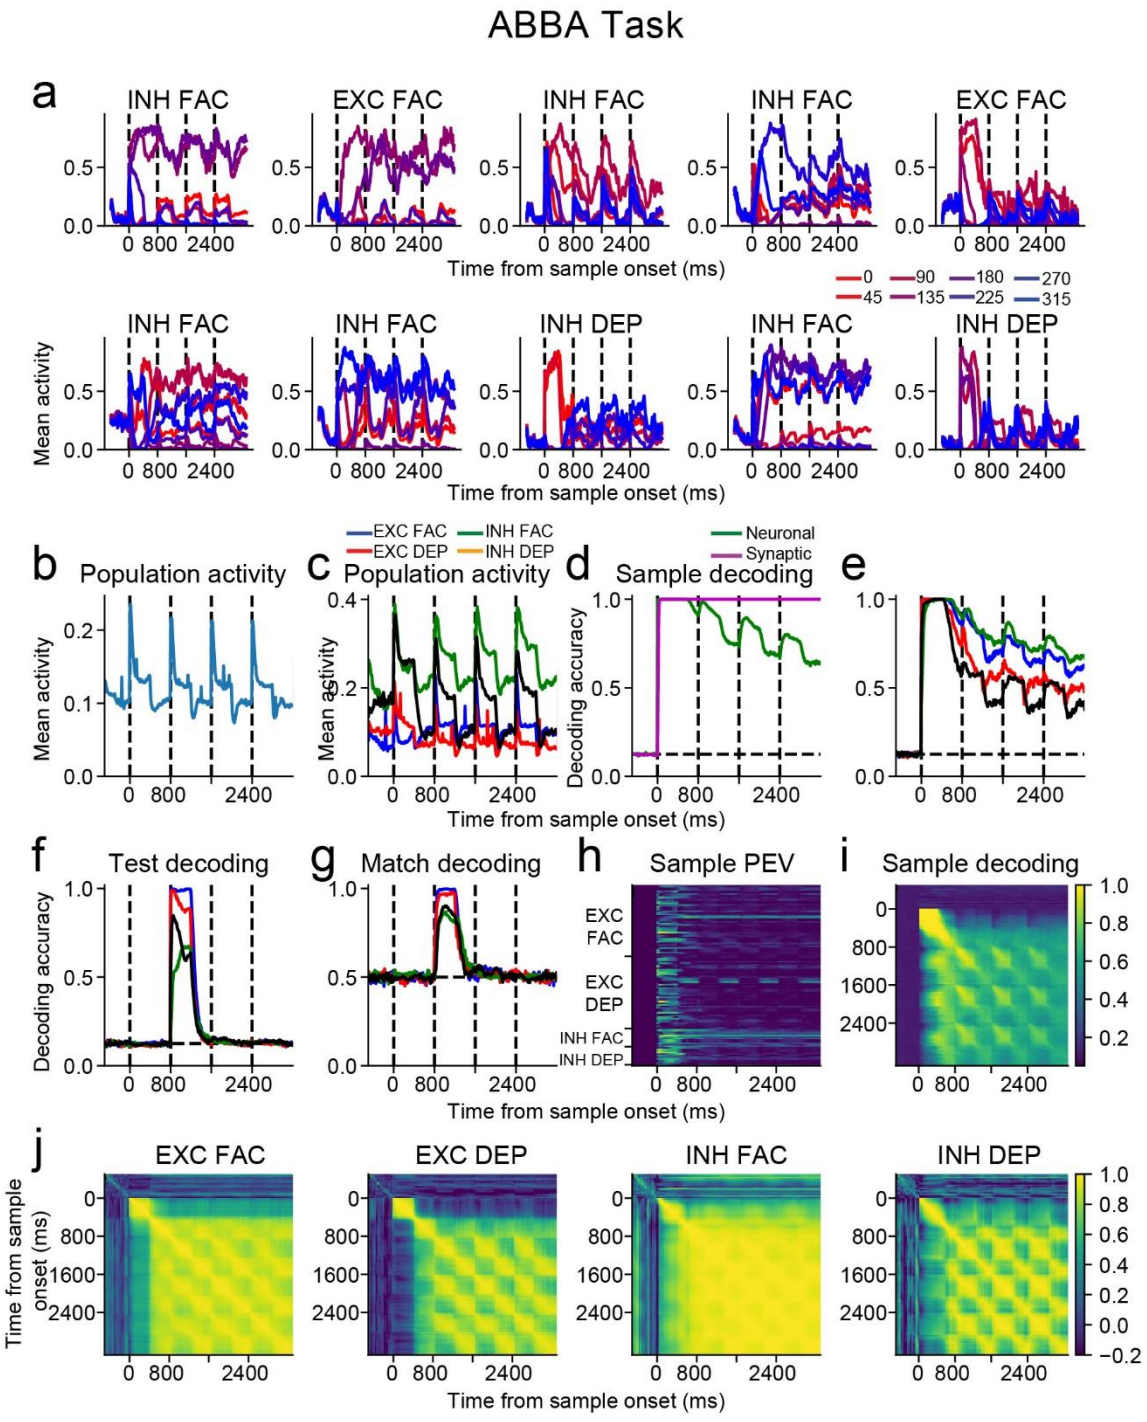

Figure S22

Figure S22 Properties of an example network trained on the ABBA task. Similar to Figure S15, except showing the properties of an example network that was trained on the ABBA task.

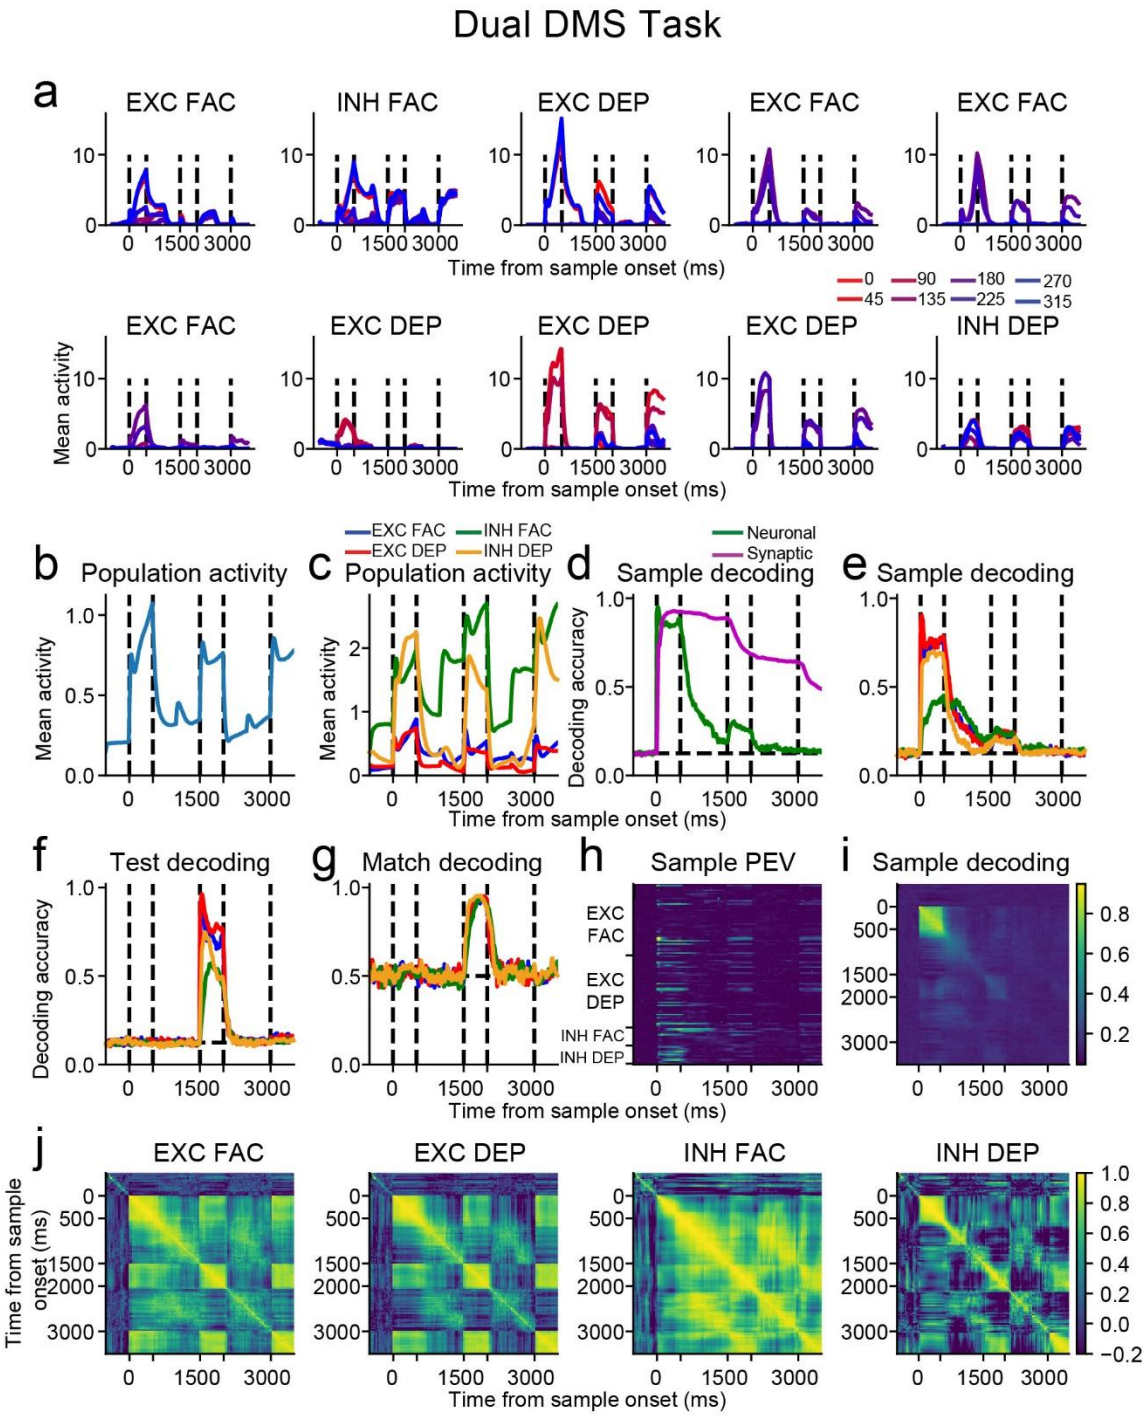

Figure S23

Figure S23 Properties of an example network trained on the dual DMS task. Similar to Figure S15, except showing the properties of an example network that was trained on the dual DMS task.

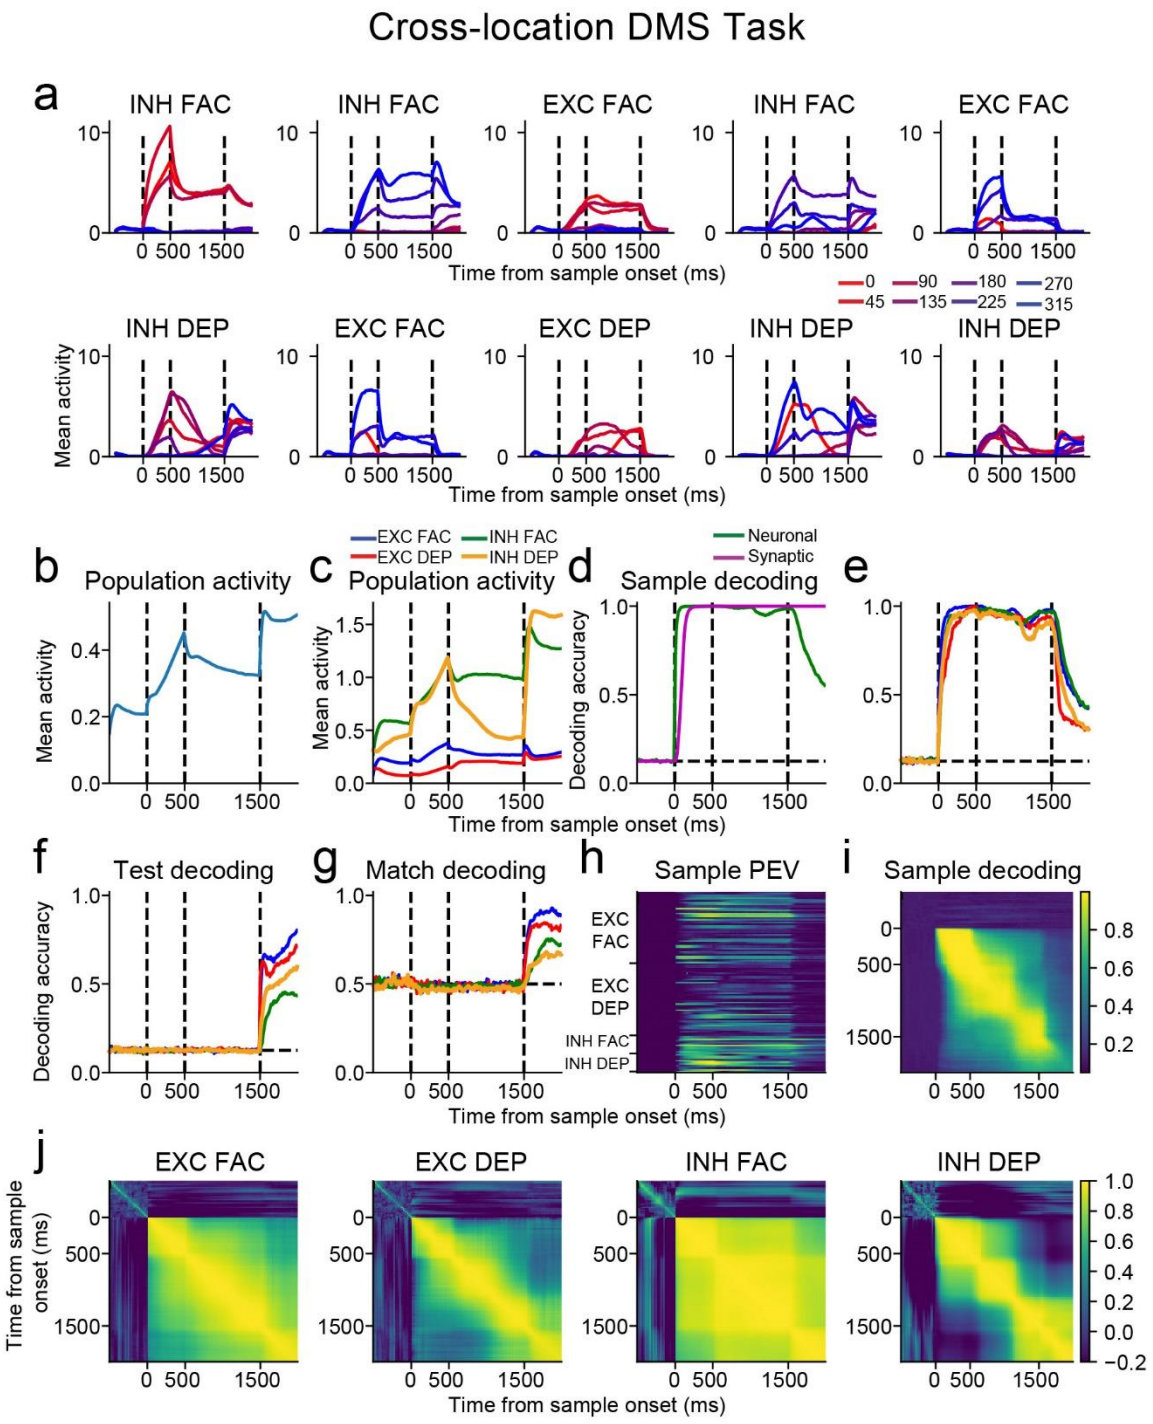

Figure S24

Figure S24 Properties of an example network trained on the cross-location DMS task. Similar to Figure S15, except showing the properties of an example network that was trained on the cross-location DMS task.

Modelling Notes

Effect of noise and regularization

In the brain, neurons receive inputs from a wide range of cortical sources, from both local and distant areas, as well from a variety of subcortical structures. To model the impact of these diverse inputs, in which many might be unrelated to the task at hand, we added Gaussian noise directly to the RNN neurons, and to the input neurons (see Methods). We noticed that increasing the amount of noise decreases the reliability of decoding motion direction from neuronal activity, which encourages RNNs to rely more heavily upon information maintained in synaptic efficacies to solve the task. Additionally, we added a penalty term on the mean squared neuronal activity values of the recurrent neurons to encourage networks to solve the task using lower levels of neuronal activity. Increasing this penalty term also encourages RNNs to rely more heavily upon information maintained in synaptic efficacies, and less upon neuronal activity, to solve the task.

We obviously do not have access to effective noise levels in the brain, nor on how much pressure they are under to minimize neuronal activity or other various constraints. However, it is worth noting that how information maintained in WM is distributed between persistent activity compared to synaptic efficacies depends on these values, which might vary between species, cortical areas, etc.

Varying network regularization, delay times

We wanted to know whether networks maintained information using STSP only because we imposed a cost on neuronal activity (see Methods). Thus we repeated our analysis in Figure 2 on 20 networks that solved the same task with zero penalty on neuronal activity (Figure S3a). For all 20 networks, synaptic decoding accuracy was significantly greater than neuronal decoding measured at the end of the delay (neuronal decoding = 0.56, synaptic decoding = 1.00, $t(19) = 8.67$, $P < 10^{-7}$, two-sided paired t-test, $N = 20$), and shuffling synaptic efficacies produced a greater decrease in task accuracy compared to shuffling neuronal activity (accuracy after neuronal shuffling = 0.95, after synaptic shuffling = 0.51, $t(19) = 24.81$, $P < 10^{-15}$). Thus, all 20 networks primarily used information maintained in synaptic efficacies to solve the task despite having no penalty on neuronal activity and greater overall neuronal activity (mean activity without penalty = 4.05, with penalty = 0.34, $t(38) = 24.54$, $P < 10^{-24}$, two-sided, unpaired t-test).

We also experimented with an L1 penalty term on neuronal activity as opposed the L2 term we used in this study. We found that with a sufficiently high L1 penalty, RNNs were equally able to solve the DMS task in an activity-silent manner (data not sown).

In addition to constraints on neuronal activity, the brain is also likely under pressure to minimize the amount of wiring connecting different neurons[1]. Thus, we next trained 20 networks to solve the DMS task with a penalty on the connectivity between the recurrent neurons (see Methods:

Network training) instead of a penalty on neuronal activity (Figure S3b). The average connectivity decreased in networks with the weight penalty (mean connection weight without penalty = 0.23, with penalty = 0.02, $t(38) = 82.97$, $P < 10^{-43}$, two-sided, unpaired t-test). More importantly, 14 out of 20 networks with the weight penalty solved the task using activity-silent WM, in which neuronal activity during the last 100 ms of the delay was not significantly greater than chance ($P > 0.05$, bootstrap), and shuffling synaptic efficacies produced a greater decrease in task accuracy compared to shuffling neuronal activity (accuracy after neuronal shuffling = 0.97, after synaptic shuffling = 0.51, $t(19) = 33.70$, $P < 10^{-17}$). Thus, wiring costs in the brain might also pressure the brain to maintain information in an activity-silent manner.

Finally, we reasoned that the networks' ability to silently maintain information in WM should depend upon the ratio between the delay duration and the values of STSP time constants. We found that networks were unable to solve the DMS task in an activity-silent manner when we decreased the STSP time constants to 100 and 500 ms (Figure S4a), or increased the delay to 2000 or 2500 ms (Figure S4c&d). In contrast, some networks were still able to solve in activity-silent manner for a delay period of 1500 ms (Figure S4b). This suggests that our networks are only capable of activity-silent mnemonic encoding when the delay period is not significantly greater than the slow STSP time constant.

Synaptic decoding accuracy for the A-B-C-A task

As noted in the Results, the synaptic decoding accuracy for the A-B-C-A task was relatively lower compared to other tasks, which was because of how we calculated decoding accuracy (see Methods). When training our networks on the A-B-C-A task, matching test stimuli, which can boost the sample information maintained in short-term memory, occur at least once during a trial with 87.5% probability ($1 - \frac{1}{2}^3$). However, when calculating decoding accuracy, we randomly sampled test motion direction independently of sample directions, so that matching test stimuli occur at least once during a trial with 30.0% probability ($1 - \frac{7}{8}^3$). We sampled independently when calculating decoding accuracy because a sample stimulus and a matching test stimulus generate similar patterns of neural and synaptic activity, meaning that sample decoding accuracy would be artificially high if matching test stimuli occurred more frequently than chance. Thus, synaptic decoding accuracy in A-B-C-A task likely appears lower than expected because networks receive less matching test stimuli than they receive during training.

The correlation between task manipulation and persistent activity for different configurations of STSP

In all the networks discussed so far, half the neurons had facilitating projections, and the other half had depressing projections. We wanted to know whether the correlation between persistent neuronal activity and manipulation existed for different configurations of STSP (similar to Figure 7). We tried training networks in which synapses projecting from all excitatory neurons were facilitating. While it was possible to train networks to achieve close to 90% task accuracy on the DMS task, either with or without the penalty on neuronal activity, many networks could not learn

the DMRS task (Figure S10), consistent with the important role of INH DEP neurons in this task Figure 3). We thus did not attempt to measure the correlation between manipulation and persistent activity for this STSP configuration. However, we successfully trained networks in which all synapses were facilitating (Figure S11), all synapses were depressing (Figure S12), and all synapses projecting from excitatory neurons were depressing and all synapses projecting from inhibitory neurons were facilitating (Figure S13). We observed significant positive correlations between manipulation and persistent activity when all synapses were depressing (Figure S12) and when all synapses projecting from excitatory neurons were depressing and all synapses projecting from inhibitory neurons were facilitating and there was penalty on high neuronal activity (Figure S13). We observed positive, but non-significant correlations in other cases.

References

1. Bullmore, E. & Sporns, O. The economy of brain network organization. *Nat. Rev. Neurosci.* **13,** 336–349 (2012).