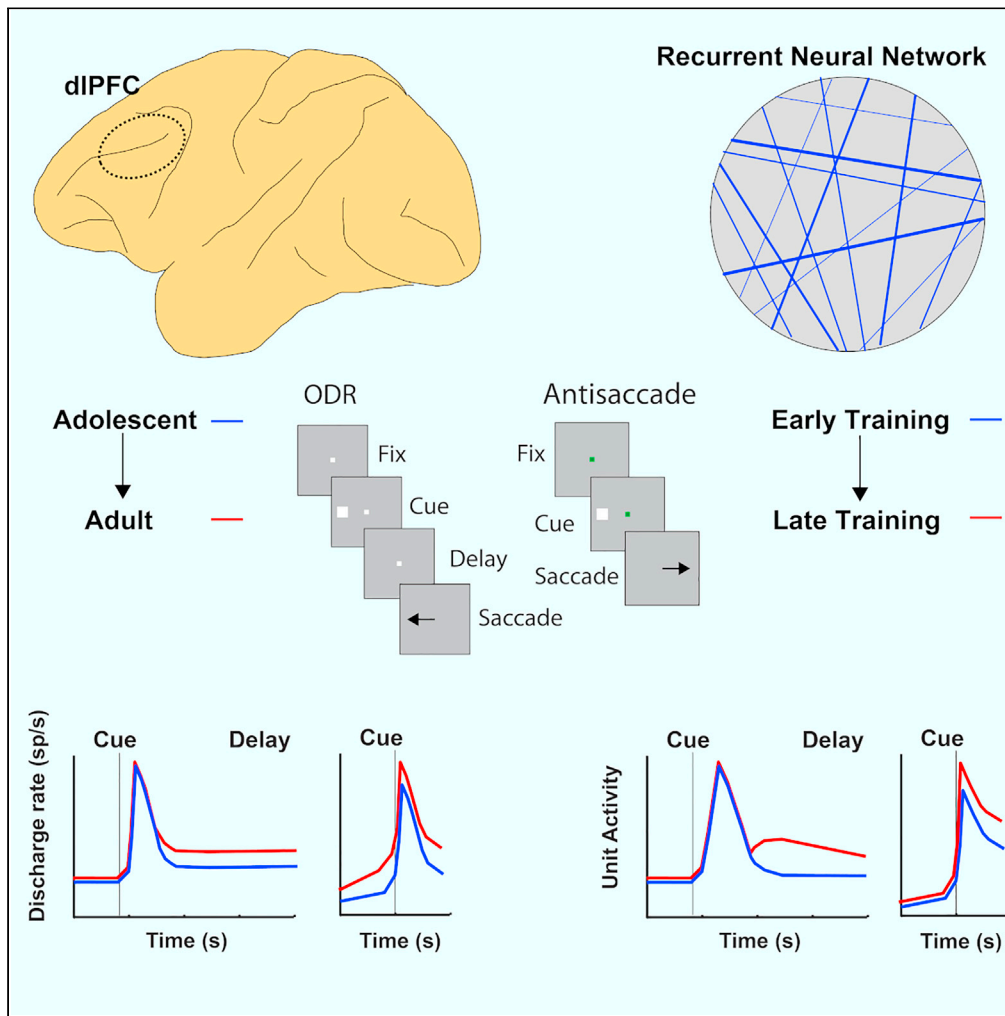


Article

Emergence of prefrontal neuron maturation properties by training recurrent neural networks in cognitive tasks



Yichen Henry Liu,
Junda Zhu,
Christos
Constantinidis, Xin
Zhou

maizie.zhou@vanderbilt.edu

Highlights

Properties of RNN networks during training offer insights in prefrontal maturation

Fully trained networks exhibit higher levels of activity in working memory tasks

Trained networks also exhibit higher activation in antisaccade tasks

Partially trained RNNs can generate accurate predictions of immature PFC activity

Liu et al., iScience 24, 103178
October 22, 2021 © 2021 The Author(s).
<https://doi.org/10.1016/j.isci.2021.103178>



Article

Emergence of prefrontal neuron maturation properties by training recurrent neural networks in cognitive tasks

Yichen Henry Liu,¹ Junda Zhu,² Christos Constantinidis,^{2,3,4} and Xin Zhou^{1,3,5,6,*}

SUMMARY

Working memory and response inhibition are functions that mature relatively late in life, after adolescence, paralleling the maturation of the prefrontal cortex. The link between behavioral and neural maturation is not obvious, however, making it challenging to understand how neural activity underlies the maturation of cognitive function. To gain insights into the nature of observed changes in prefrontal activity between adolescence and adulthood, we investigated the progressive changes in unit activity of recurrent neural networks as they were trained to perform working memory and response inhibition tasks. These included increased delay period activity during working memory tasks and increased activation in antisaccade tasks. These findings reveal universal properties underlying the neuronal computations behind cognitive tasks and explicate the nature of changes that occur as the result of developmental maturation.

INTRODUCTION

Cognitive functions necessary for executive control such as working memory (the ability to maintain information in mind over a period of seconds) and response inhibition (the capacity to resist immediate responses and plan appropriate actions) mature relatively late in life, after adolescence (Davidson et al., 2006; Fry and Hale, 2000; Gathercole et al., 2004; Kramer et al., 2005; Mischel et al., 1989; Ordaz et al., 2013; Ullman et al., 2014). This prolonged cognitive enhancement that persists after the onset of puberty parallels the maturation of the prefrontal cortex (PFC) (Chugani et al., 1987; Jernigan et al., 1991; Pfeifferbaum et al., 1994; Sowell et al., 2001; Ullman et al., 2014; Yakovlev and Lecours, 1967). Neurodevelopmental and psychiatric conditions such as attention deficit hyperactivity disorder, bipolar disorder, and schizophrenia are characterized by poor working memory and/or inhibitory control and they also manifest themselves in early adulthood (Diler et al., 2013; McDowell et al., 2002; Smyrnis et al., 2004). Changing patterns of prefrontal activation between childhood and adulthood have been well documented in human imaging studies for tasks that require working memory (Bunge et al., 2002; Burgund et al., 2006; Klingberg et al., 2002; Kwon et al., 2002; Luna et al., 2001; Olesen et al., 2003, 2007) and response inhibition (Bunge et al., 2002; Luna et al., 2001, 2008; Ordaz et al., 2013; Satterthwaite et al., 2013). Recordings from single neurons in these tasks have been recently obtained in nonhuman primate models, at different developmental stages (Zhou et al., 2013, 2014, 2016b). Monkey maturation of working memory and response inhibition parallels that of humans (Constantinidis and Luna, 2019). The neurophysiological results are therefore highly informative about how mature cognitive functions are achieved in the adult brain. Many empirical changes observed in neurophysiological recordings between developmental states, however, remain difficult to interpret, and it is unclear if systematic differences in neural activity at different developmental stages are causal to cognitive improvement or incidental.

A potential means of understanding the nature of computations performed by neural circuits is to rely on deep learning methods (Cichy and Kaiser, 2019; Yang and Wang, 2020), which have been very fruitful in the analysis of the functions of neurons along the cortical visual pathways. Convolutional neural networks have had remarkable success in artificial vision, and the properties of units in their hidden layers have been found to mimic the properties of real neurons in the primate visual pathway (Bashivan et al., 2019; Cadena et al., 2019; Khaligh-Razavi et al., 2017; Rajalingham et al., 2018; Yamins and DiCarlo, 2016). It is possible to directly compare the activation profile of units in the hidden layers of artificial networks with neurons in the ventral visual pathway (Pospisil et al., 2018). Deep learning models are thus being used to understand

¹Department of Computer Science, Vanderbilt University, Nashville, TN 37235, USA

²Neuroscience Program, Vanderbilt University, Nashville, TN 37235, USA

³Department of Biomedical Engineering, Vanderbilt University, Nashville, TN 37235, USA

⁴Department of Ophthalmology and Visual Sciences, Vanderbilt University Medical Center, Nashville, TN 37232, USA

⁵Data Science Institute, Vanderbilt University, Nashville, TN 37235, USA

⁶Lead contact

*Correspondence: maizie.zhou@vanderbilt.edu
<https://doi.org/10.1016/j.isci.2021.103178>



the development, organization, and computations of the sensory cortex (Bashivan et al., 2019; Rajalingham et al., 2018; Yamins and DiCarlo, 2016).

Another class of artificial networks models, recurrent neural networks (RNNs) has been used recently to model performance of cognitive tasks and to study cortical areas involved in cognitive function (Mante et al., 2013; Song et al., 2017). RNN units are capable of exhibiting temporal dynamics resembling the time course of neural activity. RNNs can be trained to simulate performance of working memory and response inhibition tasks (Kim and Sejnowski, 2021; Masse et al., 2019; Yang et al., 2019) and have revealed, for example, how different neuron clusters represent different cognitive tasks (Yang et al., 2019). It has been hypothesized that structured neural representations necessary for complex behaviors might emerge from a limited set of computational principles (Saxe et al., 2021). RNNs have been used successfully to simulate the changes that occur in real neural circuits during learning, providing valuable insights on the principles guiding plasticity (Cao et al., 2020; Feulner and Clopath, 2021; Warnberg and Kumar, 2019; Wenliang and Seitz, 2018). We were motivated, therefore, to use RNNs achieving different levels of performance (by virtue of controlling the number of parameter updates) as a way of probing the changes that characterize maturation of real brain networks, which achieve different levels of asymptotic performance at different developmental stages due to neural circuit development. Our analysis compared directly the levels of activity and temporal dynamics of single RNN units with those of single PFC neurons. We further sought to understand whether changes in activity we observed in PFC during brain maturation also emerged in the RNN network. Our work provides a framework for understanding the computations underlying cognitive maturation, simulating a wide range of cognitive tasks that may exhibit developmental differences and generating predictions for further experimental study.

RESULTS

Our approach compared the activity of real neurons recorded in the PFC of adolescent and adult monkeys as they performed working memory and response inhibition tasks with RNN units trained to perform the same tasks, at networks that achieved comparable performance. Neural data were obtained from four macaque monkeys (*Macaca mulatta*) trained to perform two variants of the oculomotor delayed response task (ODR and ODR plus distractor task - Figure 1A, B) and three different variants of the antisaccade task differing in the timing of the cue onset relative to the fixation point offset (Figure 1D-F). Stimuli in all tasks could appear at 8 possible locations (Figure 1C). Behavioral data and neurophysiological recordings from areas 8a and 46 of the dorsolateral PFC were obtained around the time of puberty (the “young” stage henceforth) and after the same animals had reached full maturity (the “adult” stage). The monkeys generally mastered the basic rules of the tasks very easily, in a matter of a few sessions. Young animals still made a lot of errors when a delay more than 1 s was imposed in the working memory task, and this performance level improved little with further training. Neurophysiological recordings were obtained once asymptotic performance was reached (Zhou et al., 2013, 2016b, 2016c). In analogy, we examined responses of RNNs at different stages of network performance, during execution of the exact same tasks. In all results that follow, RNNs were trained to perform all task variants.

We used an approach analogous to neurophysiological recordings to identify RNN units that respond to the task and study their responses to different task conditions (Figure 2). Such analysis of units in artificial neural networks has been called “artphysiology” and offers direct insights on the role of neuronal activity in computation. Use of RNNs allowed us to plot the entire time series of the unit’s activation during a trial in the task and compare it with the Peri-stimulus Time Histogram (PSTH) of real prefrontal neurons. For each recurrent unit, we analyzed trials with the stimulus appearing at each of the eight locations and calculated its firing rate during the cue and delay period. We identified individual units with responsiveness to the task, evidenced by a significant increase in activity during stimulus presentation or delay period compared with the baseline period (t test, $p < 0.05$) and selectivity for different stimulus locations (one-way ANOVA test, $p < 0.05$). This selection procedure was meant to mirror the analysis of PFC neurons; RNN units that did not meet the criterion may still have exhibited weak task-related responses and contributed to the functionality of the network. The time course of their firing rate could then be plotted, in a manner similar to what we have done for real neurons (Figure 3). Responses from responsive RNN units were averaged together to produce population activity. Of importance, we performed this analysis at different phases of the RNN network training, simulating different stages of the brain development. We identified three stages (early, middle, and mature), demarcated by transition points at which the network performed each task achieving

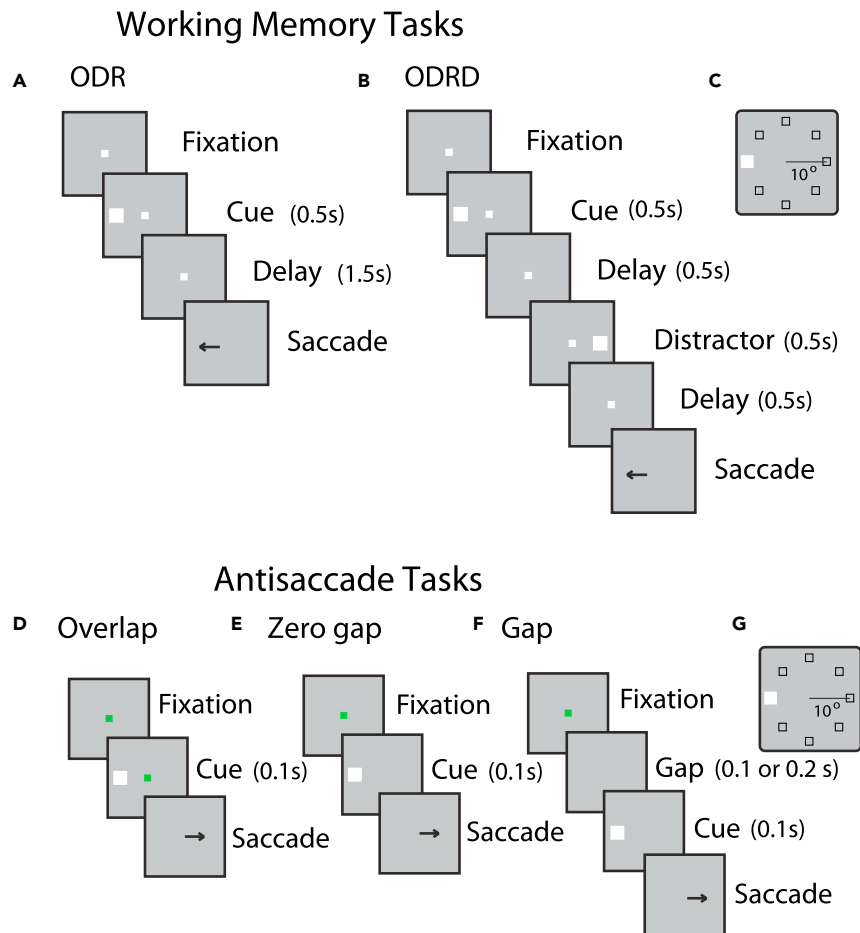


Figure 1. Behavioral tasks

Successive frames illustrate the sequence of events in the working memory and antisaccade tasks simulated.

(A) Oculomotor delayed response (ODR) task. The monkey needs to observe a cue and after a delay period to make an eye movement to the remembered direction of the cue.

(B) ODR + distractor task. A distractor stimulus is presented after the cue, which the monkey needs to ignore.

(C) Possible locations of the cue in the ODR and ODRD tasks.

(D) Overlap variant of the antisaccade task. The cue appears while the fixation point is still present.

(E) Zero-gap variant of the antisaccade task. The cue appears simultaneously with the fixation point disappearing.

(F) Gap variant of the antisaccade task. The cue appears after the fixation point was turned off.

(G) Schematic illustration of possible stimulus locations in the antisaccade tasks, which are the same as in the working memory tasks.

a percentage of correct trials of <35%, 35–65%, and >65%, with the middle interval selected as to approximate the range of performance of adolescent monkeys in these tasks (Zhou et al., 2013, 2016b).

RNN activity in working memory tasks

We first examined RNN activity in the simplest working memory task, the ODR task. As we have reported previously, transition from the adolescent to the adult stage of maturation in monkeys is characterized primarily by an increase in activity of prefrontal neurons during the delay period of the ODR task, after the stimulus was no longer present. We wished to test whether elevated activity during the delay period is a property of fully trained RNNs, particularly because the necessity of persistent neural activity in working memory itself has been a matter of debate in recent years (Constantinidis et al., 2018; Lundqvist et al., 2018).

To avoid a sampling bias in the analysis of neuronal data, we identified neurons with stimulus-driven activity, defined as firing rate during the cue presentation period that was elevated relative to the fixation period

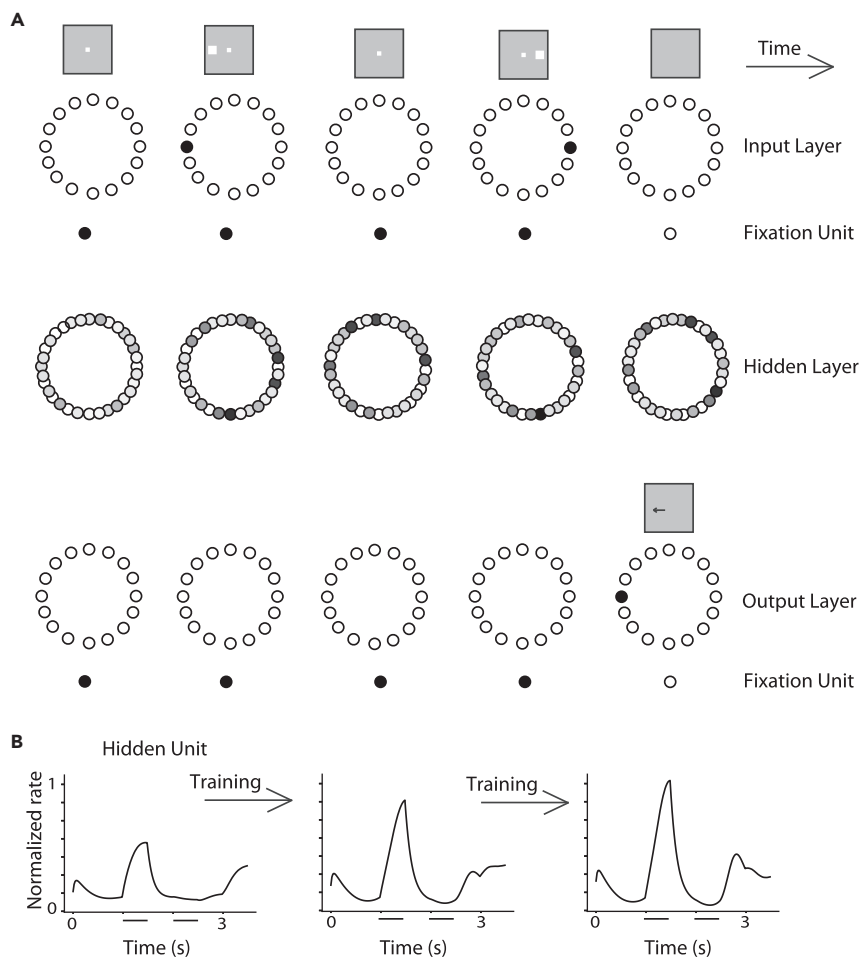


Figure 2. RNN network

(A) Schematic architecture of input, hidden, and output layers of the network. Panels are arranged as to indicate successive events in time, in a single trial, across the horizontal axis. Appearance of the fixation point in the screen (top left panel) is simulated by virtue of activation of fixation units, a subset of the input units. Appearance of the visual stimulus to the left of the fixation activates input units in representing this location. Input units are connected to hidden layer units and the latter are connected to output layer units. In the ODRD task, the trained network generates a response to the remembered location of the cue by virtue of activation of the corresponding output unit.

(B) The firing rate of a single hidden unit is plotted as a function of time, during the duration of one trial. Successive panels now represent evolution of the unit's response as the training progresses. The horizontal lines represent the times of stimulus appearance in the ODRD task.

and then compared their delay period activity, even though the latter was not a criterion for selection. Approximately half of the PFC neurons exhibited stimulus-driven activity thus defined (309/607 PFC neurons in the young dataset and 324/830 in the adult dataset). Results of this analysis, replicated in [Figure 4A](#), revealed a significant increase of firing rate during the delay interval. In contrast, responses to the visual stimulus itself remain essentially unchanged, at least relative to the baseline fixation period. Delay period in RNN units similarly increased between stages ([Figure 4B](#)). An increase in stimulus responses was also evident early on, but this had essentially matured by the mid-trained stage. What continued to change until the fully trained stage was RNN unit activation that persisted in the delay period. The mean firing rate in the delay period differed significantly between phases (one-way ANOVA, $F_{2,331} = 7.59$, $p = 6.0 \times 10^{-4}$ for the network units on which [Figure 4](#) was based). The average rate in each phase can be seen in [Figure S1](#). The full time course of changes in activity can be seen in [Figures S2A](#) and [S2B](#). Results based on the activity of all PFC neurons and all RNN units were essentially identical ([Figure S3](#)). It was also possible to examine the activity of the exact same RNN units as training progressed. This analysis

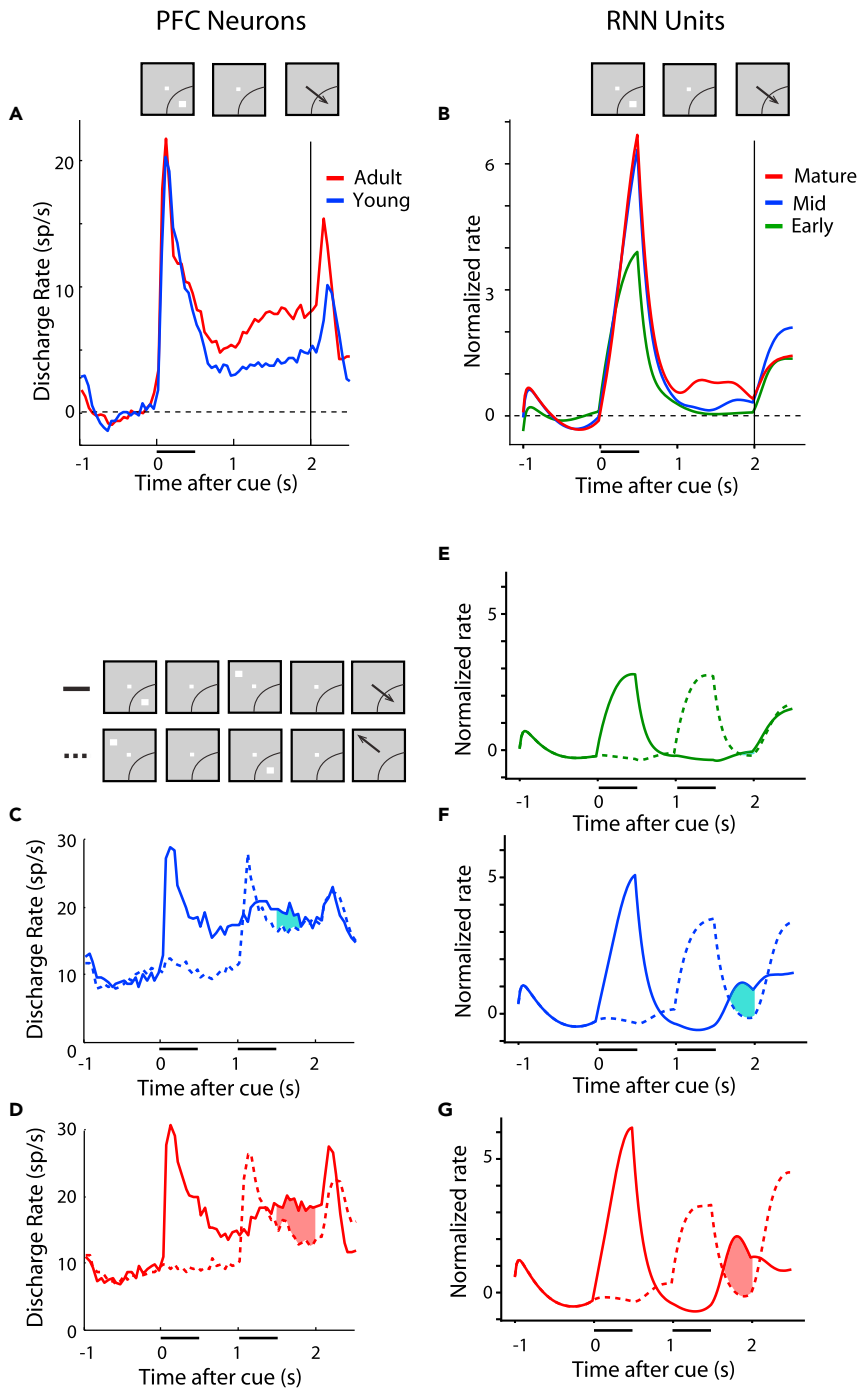


Figure 4. RNN and PFC activity in working memory tasks

(A) Population discharge rate from young and adult PFC stages in the ODR task. Mean firing rate is plotted across neurons ($n = 309$ for young, $n = 324$ for adult), after subtracting the respective baseline firing rate. Insets on top of the figure represent the sequence of events in the ODR task. Arc is meant to represent the neuron's receptive field, indicating that the stimulus was presented at each neuron's most responsive location.

(B) Mean rate of responsive RNN units, during three developmental stages ($n = 114$ for early, $n = 131$ for mid, $n = 144$ for late, among 256 total units in the network).

(C) Population discharge rate from the young PFC stage in the ODR + distractor task. The solid line represents appearance of a stimulus in the receptive field, followed by a distractor out of the receptive field. Insets represent the sequence of events in the ODRD task. The dotted line represents the appearance of a stimulus out of the receptive field,

Figure 4. Continued

followed by a distractor in the receptive field. The difference between conditions indicated by the shaded line represents a measure of how well the distracting stimulus is filtered in the second delay period.

(D) Same as in C for adult data.

(E–G) Mean rate of RNN units plotted using the same conventions for three developmental stages. Neural data in (A), (C), and (D), from Zhou et al. (2016a).

Other RNN characteristics related to persistent activity also resembled neuronal responses. The number of selective RNN units for the different stimulus locations increased modestly between the different stages of training: 45% in the early phase, 51% in the middle, 56% in the mature phase (chi-square test, $p = 0.029$). Across the population, RNN units trained to perform the ODR task also exhibited a population tuning profile that resembled that of the population of PFC neurons (Figure S5).

RNN dynamics during working memory

To gain insight into the changes of the temporal dynamics of unit activity as training progressed, we relied on a dimensionality reduction technique, termed dPCA, which has been used to understand dynamics of neuronal activity (Kobak et al., 2016), separating components relating to specific aspects of the task used. We processed the activity of RNN units as we did for neuronal activity and identified the top components capturing the temporal envelope of the response, selectivity for the visual stimulus location, and for mixture of time and stimulus (Figure 5). The transition from the adolescent to the adult PFC was characterized by more robust separation of stimulus components during the delay period of the task, by themselves (Figures 5A and 5B, middle panels) or in mixtures with time components (Figures 5A and 5B, right panels). A very similar progression was observed in the RNN activity (Figures 5C–5E, middle panels). Whereas early in training, stimulus components were separable only shortly after the cue appearance (Figure 5C), in the mature network different stimuli were represented robustly through the end of the trial. However, some differences were also present; transient stimulus representation in the delay interval also emerged in the mature RNN (see Figure 5E, right panel), in agreement with previous studies (Orhan and Ma, 2019), although this was absent in the PFC neural data. These conclusions were confirmed when we plotted the timing of RNN activity across units in the network (Figure S6). Many RNN units exhibited transient responses; however, the predominant effect of training is the increase in stability of delay-period activity. Although the majority of neurons in the early network responded only to the cue presentation and ceased to be active in the delay period (Figure S6A), the activity of more units in the mature network spanned the cue and at least part of delay period (Figure S6C). Cross-temporal decoding similarly revealed that stability of the stimulus representation across the trial improved for the adult PFC (Figures S7A and S7B), as well as for RNN units across training (Figures S7C–S7E).

We also used a more challenging working memory task to probe the maturation of working memory ability, the ODRD task. In that task, after the original cue presentation that needed to be remembered, a distractor presentation followed that needed to be ignored over a second delay period (Figure 1B). In neural activity, appearance of the cue in the receptive field generated a transient response, followed by persistent activity in the delay period of the task, which continued to be present, even after the distractor was presented (Figures 4C and 4D). Appearance of the distractor in the receptive field also generated activation; however, the relative difference between the activity generated by the cue and by the distractor in the second delay period of the task characterized the maturation of prefrontal neuronal activity (shaded area in Figures 4C and 4D). This can be viewed as a measure of the prefrontal network ability to filter distracting activation. In this case, too, we wished to test whether persistent activation of the RNN network was a result of training. In the early training stage, virtually no difference was present in the second delay period between cue and distractor-elicited activation (Figure 4E). A difference between the two conditions began to emerge in the mid networks (Figure 4F) and further grew in the fully trained ones (Figure 4G). Tracking the activity of the same RNN units across training produced very similar results (Figures S4B–S4D). The mean firing rate across units in the second delay period differed significantly between phases (one-way ANOVA, $F_{2,579} = 5.44$, $p = 0.004$; see also Figure S1B). In this case, too, the behavior of RNN units replicated the empirical pattern we have reported before in the PFC.

RNN activity in response inhibition tasks

We proceeded to examine response inhibition, the ability to resist a response toward a prepotent stimulus, which is another cognitive domain that matures markedly between adolescence and adulthood. We relied

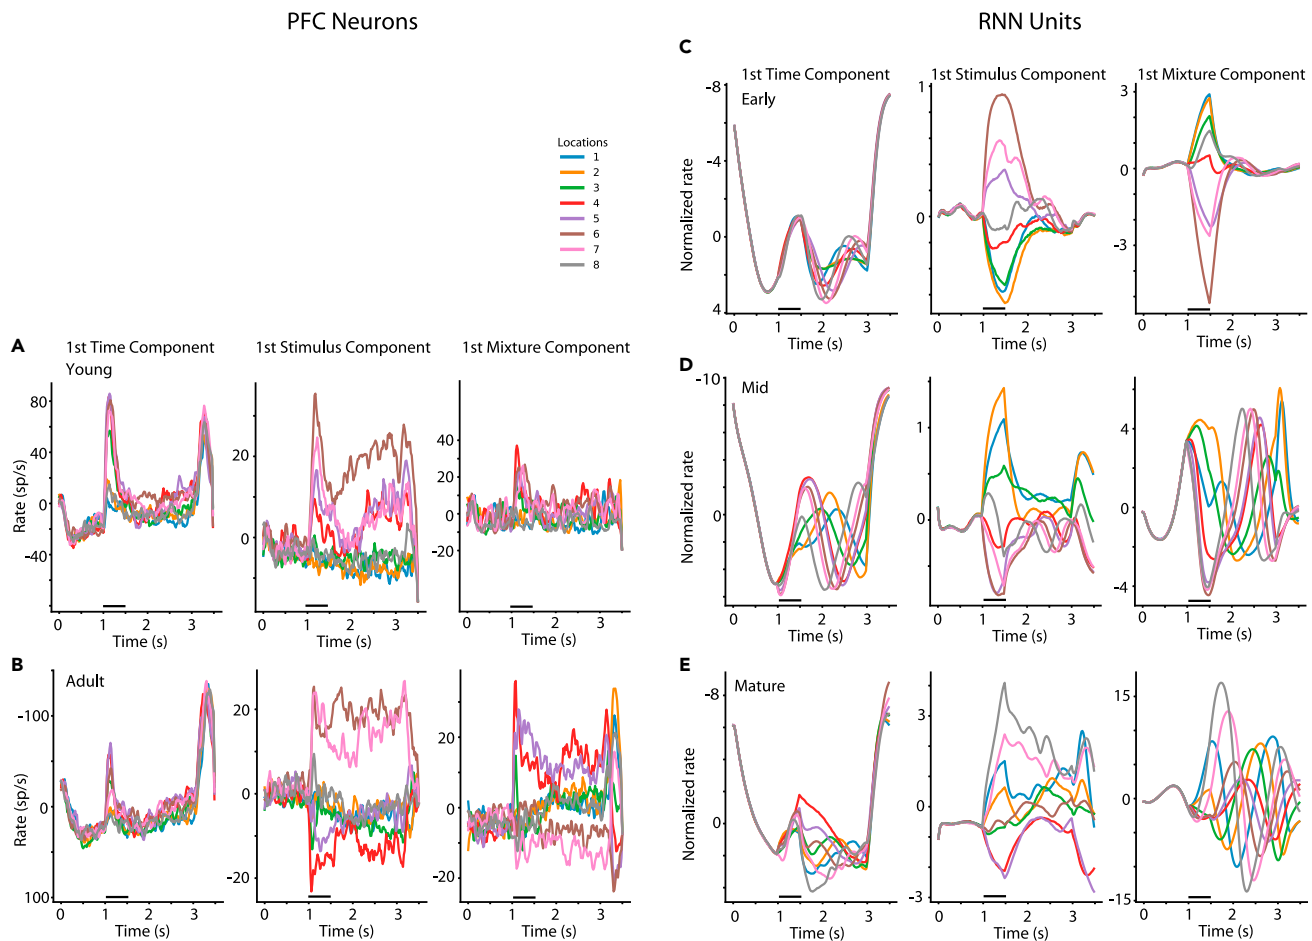


Figure 5. Demixed PCA analysis

(A) Data are plotted for PFC neurons recorded from the young monkeys. Panels from left to right represent the first time (condition-independent) component of dPCA analysis, the first stimulus-related dPCA component, and the first stimulus/time mixture dPCA component. Different color lines represent the eight stimulus locations.

(B) Data from the Adult PFC plotted as in (A).

(C–E) Same analysis performed for RNN units in the early, mid-trained, and mature networks.

on the antisaccade task that involves presentation of a visual stimulus that subjects should resist looking at but requires them to make an eye movement in the opposite direction (Figures 1D–1G). We have trained adolescent and adult monkeys in three different variants of the task of varying difficulty, by manipulating the relative timing of the cue appearance and fixation point turning off. The overlap variant was the easiest, as it allowed the monkeys to view the cue stimulus for 100 ms before both the fixation point and cue turned off, which was the signal to the monkey to initiate the saccade (Figure 1D). The gap variant (Figure 1F) was the hardest as the fixation point turned off for 100 ms before the cue appeared (creating a “gap” of 100 ms when the screen was blank). In the absence of a fixation point where the subject can hold its gaze, it is much more difficult to resist making an erroneous saccade to the stimulus than away from it. The zero gap variant involved turning off the fixation point simultaneously with the cue appearance and was intermediate in difficulty (Figure 1D).

Neural responses in the antisaccade task mature considerably between the time of adolescence and adulthood (Zhou et al., 2016b). In principle, improvement in performance in the antisaccade task may be achieved by a lower level of activation following the presentation of the visual stimulus, in essence filtering the representation of the prepotent stimulus toward which a saccade should be avoided or by an increase of activity representing the target. The former alternative received some support, at least on the surface, by imaging studies that show decrease of prefrontal cortical activation between the time of adolescence and

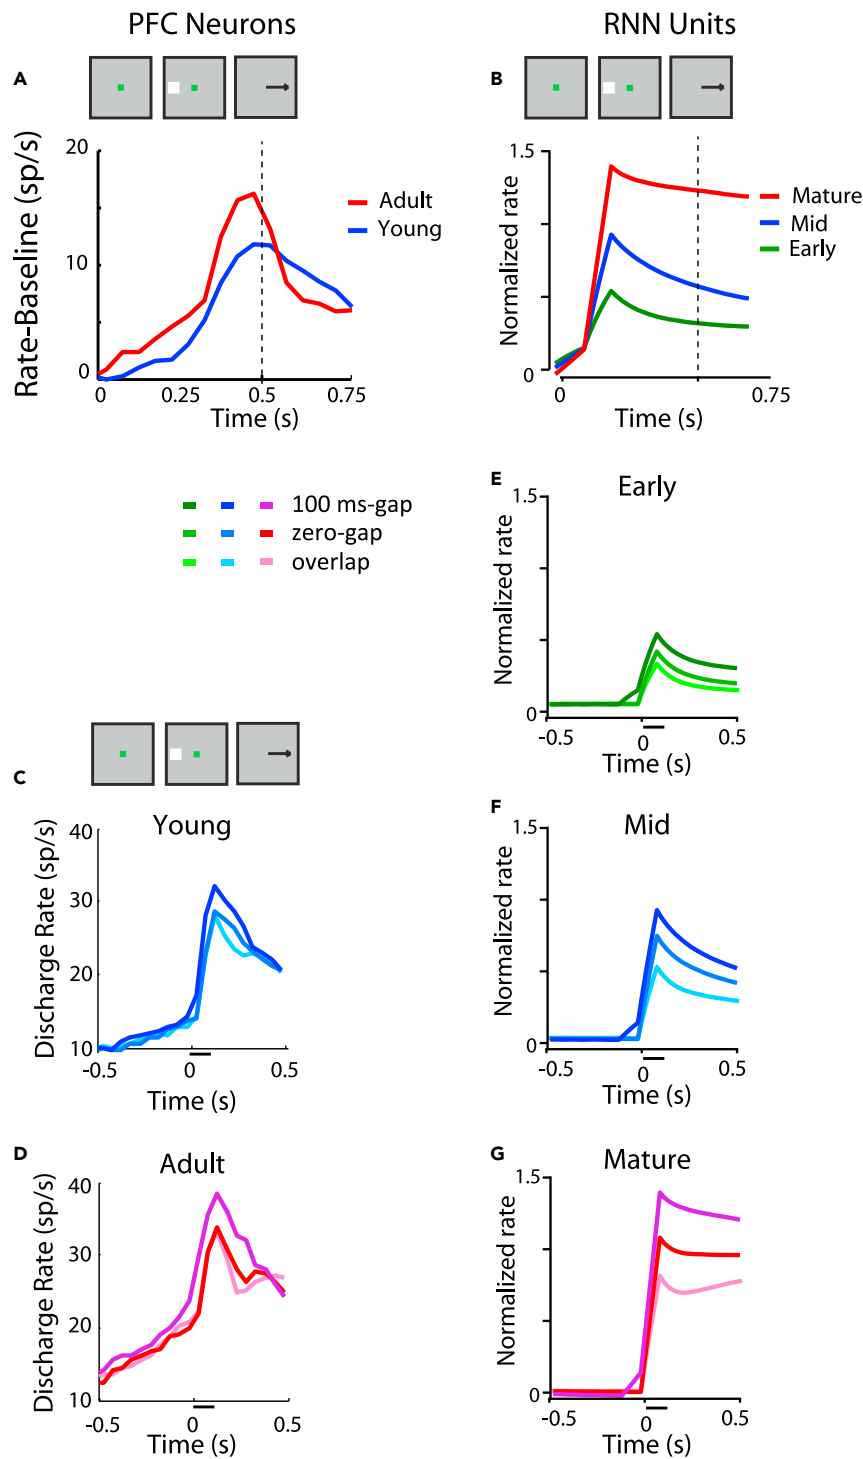


Figure 6. RNN and PFC activity in response inhibition tasks

(A) Population discharge rate from young and adult PFC stages in the overlap variant of the antisaccade task, after subtracting the respective baseline firing rate. Data have been synchronized to the onset of the saccade, which is plotted here, at time 0.5 (indicated by vertical line).

(B) Mean rate of the RNN units responsive in the antisaccade task, during three developmental stages.

(C) Population discharge rate from the young PFC stage three variants of the antisaccade task: gap, zero-gap, and overlap. The horizontal line represents the appearance of the cue stimulus in the receptive field, at time 0.

Figure 6. Continued

(D) Same as in C for adult data.

(E–G) Mean rate of RNN units plotted using the same conventions for three developmental stages. Neural data in (A), (C), and (D), from Zhou et al. (2016b).

adulthood (Simmonds et al., 2017). Somewhat unexpectedly, neurophysiological results supported the latter outcome. The most salient difference in neurophysiological recordings is an increase in activity in the adult stage for results synchronized on the onset of the saccade, even after subtracting the baseline firing rate (Figure 6A). We wished to test therefore if RNNs would capture this property. This was indeed the case (Figure 6B). A progressive increase in the peak of activation was observed from early, to mid, to mature networks (Figure 6B). This difference between stages was highly significant (one-way ANOVA for firing rate in 250 ms after cue at early, mid, mature stage of gap task, $F_{2,638} = 12.2$, $p = 6.0 \times 10^{-6}$).

Interestingly, differences in firing rate elicited by different variants of the task at each training phase were also mirrored in the activity of RNN units. A higher firing rate was observed as a function of the task variant difficulty, with the gap variant eliciting the highest firing rate among prefrontal neurons. This was true for both the young stage (Figure 6C) and the adult stage (Figure 6D). The exact same pattern has been observed for RNN unit activation in the early (Figure 6E), mid (Figure 6F), and mature RNN networks (Figure 6G). No neural theory to date has been proposed to explain how these differences relate to task difficulty and performance. The current results demonstrate that emergent processes during the optimization process of artificial neural networks capture actual neural properties.

Task learning and generalization

Results presented so far summarized responses of RNN units averaged over long training periods. To understand how unit activity evolved over time as the RNN mastered performance of the trained tasks, we analyzed the responses of the same networks across training steps. Analysis of activity and task performance during RNN network training revealed that networks typically went through a long period of weight adjustment with little change in either performance or unit activity, before they began to master the task over a relatively shorter period of training, which was also characterized by rapid changes in unit activity. This was evident in the activity of the network already discussed (Figures S2A and S2B). The delay period activity in other network instances is shown in Figures S2C and S2D. In all cases, mastery of the task (transition from the early to the mid and late performance stages) occurred relatively rapidly, over a period of parallel increase in unit activity. We examined unit activity around the time of this rapid transition, to determine if a qualitative change in activity characterized an inflexion point in task improvement (Figure S8). Unit activity around the inflexion point was generally characterized by more rapid changes in the firing rate in the direction defined by the transition from the early to the mature networks (Figure S8D).

We also examined changes in the connection weights during the course of training (Figure S9). To better visualize changes in weights, we computed each unit's variance across stimulus conditions at each time point and then averaged the variance across all time points within each time epoch to get the final variance of each epoch for this unit. We used this variance value to group recurrent units into clusters. The weight matrix obtained this way for the ODR task is shown in Figure S9. Matrices from the same model but with the position of units rearranged based on clusters obtained for the ODRD task is shown in Figure S10. Change of weights during training was characterized by the emergence of additional, distinct clusters as training progressed.

Another important question for understanding the behavior of RNN networks is whether the network generalizes across tasks during training, i.e., represents shared task elements in the same manner across all tasks, so that training in one task aids performance in another, or creates specialized subnetworks for each task. Previous work supports the idea that RNNs exhibit compositionality and perform similar task by activation of overlapping sets of units (Yang et al., 2019). In the context of our tasks, we wished to understand whether units were tuned to the same stimulus location in the ODR and ODRD task. This was clearly the case, and this property persisted throughout the training regime (Figure S11). A weaker correlation was observed between the best stimulus location that a unit exhibited in the ODR task and the antisaccade task (Figure S11B), also in agreement with neurophysiological data. This can be understood when one considers that the antisaccade task is fundamentally different from the ODR task than the ODRD is. It was notable that this divergence was amplified late in training, when the network mastered the task performance.

We also tested the RNN's ability to generalize for different delay intervals, in two ways. First, we trained a network with a fixed delay period of 1.5 or 3.0 s and tested it with a delay duration that was not trained (3.0 and 1.5 s, respectively). Networks trained with 1.5-s period delay exhibited some ability to perform the task with the longer interval. On the other hand, networks fully trained with the more challenging 3.0-s delay period activity achieved near-perfect performance (96%) when tested with the shorter, 1.5-s delay period activity. Activity of units in the fully trained networks tested with a different delay period than that they were trained with was also revealing (Figure S12); networks trained with a shorter delay were able to continue generating activity if the delay was prolonged but often produced oscillatory dynamics (Figure S12). Plotting performance of networks trained with a fixed, 1.5 delay interval revealed that performance tended to be relatively high for untrained delay intervals that were approximate multiples of the trained interval (Figure S12E). We saw a similar, limited ability of networks trained with the ODRD task to generalize across delay intervals as well (Figure S12F).

In a second approach, we trained networks with a variable delay from the outset, between 0.5 and 2.5 s. These networks generally required longer time periods to train, but their activity tended to be much more robust and remained stable across a range of delay periods in both the ODR task (Figures S13A–S13C) and the ODRD task (Figures S13D–S13F). Networks trained with a variable delay period still captured qualitatively many of the properties of networks trained with fixed intervals. For example, they tended to generate higher levels of activity representing the cue stimulus than a distractor (Figures S13G–S13I). These networks also achieved much more robust performance across a range of delay intervals (Figures S13J and S13K). The stability of networks trained with variable delay intervals could be best exemplified in terms of cross-temporal decoding performance. Mature RNN networks exhibited cross-temporal generalization that now more strongly resembled that of PFC units (Figures S7F–S7H).

Analysis of the training phases revealed some “unnatural” behaviors as well. When all tasks were trained simultaneously, the RNN network mastered the antisaccade tasks much faster (Figure S14). The variants of the antisaccade task are conceptually very simple and require a straightforward mapping of weights in a naive network: when a stimulus activates input unit θ , prepare a saccade at output position $\theta + \pi$. In the absence of a delay period, which the ODR imposes, this can be implemented easily. However, the antisaccade mapping is much more difficult to achieve in real brain networks, which are wired to perform saccades to the location of the visual stimulus.

Network hyper-parameters

For the simulations we performed, we relied on parameters that were found to be effective for RNN networks learning to perform similar tasks (Masse et al., 2019; Yang et al., 2019). However, we wished to consider whether the results were specific to this choice of hyper-parameters. We thus replicated these experiments by systematically varying the parameters of the RNNs (Figure S15). In networks that successfully acquired the tasks with different hyper-parameters, results were qualitatively similar and also resembled neuronal responses. We generally confirmed, however, that the hyper-parameters previously identified in the literature most closely resembled real neuronal responses for our results, too. In terms of weight initialization, randomized orthogonal weights produced best results, in contrast to diagonal initialization (Figure S15B), but random weights drawn from a Gaussian distribution also produced comparable results (Figure S15C). The number of the units of the network also affected network responses (Figures S15D–S15G), with results achieved by networks in the range of 128–512 units, most closely resembling neural results. Finally, the softplus activation function worked best compared with Relu, tanh, retanh, and power functions (Figures S15H–S15K), as it most closely resembles the nonlinearity function of real neurons.

DISCUSSION

The use of artificial neural networks has exploded in recent years (LeCun et al., 2015). Convolutional neural networks have had remarkable success in artificial vision, and the properties of hidden layers of artificial neural networks have been found to mimic the properties of real neurons in the primate visual pathway (Bashivan et al., 2019; Cadena et al., 2019; Khaligh-Razavi et al., 2017; Rajalingham et al., 2018; Yamins and DiCarlo, 2016). These results suggest that the same fundamental operations performed by the human brain are captured by artificial neural networks, allowing the use of such networks as scientific models (Cichy and Kaiser, 2019; Saxe et al., 2021). The activity of individual units of these networks can be dissected, with techniques inspired from Neuroscience (Pospisil et al., 2018). Neuroscience principles have also been instructive for the design of more efficient networks and learning algorithms (Sinz et al., 2019). In addition to convolutional networks, other architectures have had

practical applications in Neuroscience questions, for example, to uncover neuron spike dynamics, or encoding of elapsed time (Bi and Zhou, 2020; Pandarinath et al., 2018). The activity of the PFC has been investigated successfully with RNN frameworks. Fully trained RNN models capture many properties of the PFC, including its ability to maintain information in memory and to perform multiple cognitive tasks. Examining RNN activity during the course of training the network has also shed insights on processes such as perceptual and motor learning (Feulner and Clopath, 2021; Wenliang and Seitz, 2018) and the emergence, after training, of properties encountered in neural circuits such as grid-like representations, low-dimensional manifolds, and illusory correlations (Cao et al., 2020; Warnberg and Kumar, 2019).

Here, we capitalized on this approach to investigate a relatively unexplored aspect of prefrontal cortical function, its developmental maturation. In a series of recent studies, we have addressed for the first time the nature of changes that the PFC undergoes between the time of puberty and adulthood (Zhou et al., 2013, 2014, 2016b, 2016c). These studies revealed a number of changes in activity between young and adult subjects; however, the significance of changes at some task epochs relative to others and their relationship to behavior has been difficult to interpret. By using artificial RNNs, not explicitly modeled on any brain function, we show that changes related to maturation of the primate PFC also allow performance of cognitive tasks that require working memory and response inhibition by artificial neural networks.

Parallel neural and RNN changes

We conceptualize neural plasticity in the context of our experiment as a process that operates at two timescales: (1) the developmental maturation timescale, with a time constant of years and (2) the training timescale with a time constant in the order of weeks. This model implies that monkeys can reach asymptotic performance after practicing the task for a few dozen sessions, at each developmental stage. At that point further improvements may only be gained very slowly, as the brain matures. We use the RNN to model only the former process. By comparing activity in the network at states that achieve different levels of performance, we can determine what changes allow mature networks to achieve higher performance than immature ones, in analogy to the brain changes that allow the mature PFC to achieve higher performance than adolescent ones. In principle, the same approach can be used to compare the neural circuits of different species that have evolved different levels of cognitive abilities (Hasson et al., 2020).

In the working memory tasks, the most salient change that characterized networks with improved performance was an increase in delay period activity across the network, with little change in the cue presentation period itself. The activity of some single units spanned the entire delay period, and its magnitude increased as performance improved. The increase in activity was most striking in averages across networks (Figure S4F), in networks fully trained to achieve near perfect performance (Figure S12A), and networks trained with variable delay intervals (Figure S13). Stimulus representation was also more separable during the delay period (Figures 5C–5E). This change strongly paralleled the primate prefrontal development between the stage of adolescence to adulthood. In the ODRD task, the difference between delay period corresponding to the cue and distractor was also characteristic of development. The results also inform the debated of the basis of working memory and whether it depends on persistent discharges or not (Constantinidis et al., 2018; Lundqvist et al., 2018; Miller et al., 2018; Riley and Constantinidis, 2016). We saw that RNN units spontaneously develop working memory-related activity, in agreement with previous studies (Masse et al., 2019; Yang et al., 2019). Although it is possible for some RNN networks to perform working memory tasks without an overt increase in firing rate, more complex tasks do require persistent activity (Masse et al., 2019). It is also important to emphasize that units with a diversity of neuronal activity time courses emerged in the network, including units with delay period activity that decreased in the delay period relative to the baseline. The diversity of neuronal responses is a well-established property in the PFC, and neurons with decreasing delay period activity have also been described (Wang et al., 2004; Zhou et al., 2012). An increase in delay period activity was therefore not a monolithic change of RNN training and neurons with decreasing activity likely performed important computations as part of the network, as in fact has been suggested by prior computational studies (Barak et al., 2013; Cao et al., 2020; Cueva et al., 2020; Michaels et al., 2020; Sussillo et al., 2015). Nonetheless, emergence of persistent activity with increasing magnitude is a developmental property that characterizes improved performance in mature networks.

Analysis of RNN activity also provided important insights about the neural basis of response inhibition. Prefrontal responses in the antisaccade task are characterized by differential levels of activity in task variants depending on difficulty and parallel increases in maturation (Zhou et al., 2016b). We now found that

this pattern of responses emerges spontaneously in RNN networks, providing evidence again that the prefrontal activity optimizes the execution of the task. Improved performance in adulthood was not the result of more efficient filtering of the response to the stimulus but rather stronger representation of the saccade goal.

Limitations of study

Although we have pointed out similarities in real and simulated networks, we do not wish to overstate this analogy. We considered RNNs to be “mature” and “immature” by virtue of the performance level they reached; however, the same network architecture was ultimately capable of reaching higher performance with additional training. In contrast, the immature brain was characterized by lower asymptotic performance in cognitive tasks, even when the subjects were exposed to additional training. This comparison was still meaningful because differences between young and adult neural responses were generally subtle, suggesting only small changes in network properties, not unlike those exhibited by RNN networks achieving similar behavioral performance. The task themselves were simple for the monkeys to acquire so the updating of RNN weights through supervised learning better paralleled the maturation of synaptic connections between neurons rather than the acquisition of task rules.

The RNN networks exhibited a number of “unnatural” behaviors, owing to the unconstrained nature of their weights and the complete lack of prior experience. For example, they were able to master the anti-saccade task very easily, much faster than the working memory tasks. Although mapping an association between a stimulus and the opposite location is computationally straightforward, the brain is wired to direct saccades toward objects and not in diametric locations to them, which poses substantial challenges, particularly for the immature brain (Constantinidis and Luna, 2019). In essence, the real PFC has already undergone a substantial shaping of its weights, which cannot always be captured by the RNN networks. In the working memory tasks, in addition to exhibiting persistent activity, they also maintained information in the delay period in a transient fashion, which is a documented property of RNNs trained to maintain information in short-term memory (Orhan and Ma, 2019). This was particularly prominent in the distractor task, where the representation of the cue “reactivated” late in the trial, after the appearance of the distractor (Figures 4F and 4G). Oscillatory dynamics were also much more prominent in the RNNs than in prefrontal neurons. We should note, however, that such dynamics have been observed in other neural datasets, and arguments have been made of a significant role in working memory maintenance (de Vries et al., 2020; Lundqvist et al., 2016; Roux and Uhlhaas, 2014). Networks trained with fixed intervals were much more likely to develop oscillatory dynamics rather than stable persistent activity, and their ability to generalize across untrained delay intervals was limited. These examples illustrate that RNNs are not expected to be a precise replica of the brain. The critical question, however, was regarding the changes in neural activity across training stages, and in that respect, RNNs followed a very similar trajectory as the maturing PFC. Our results document several aspects of activity that do in fact characterize improved performance in both the maturing brain and trained neural networks. We do not wish to overstate this finding, either. Although we have explored a number of parameters, we cannot rule out that networks starting with a high initial activity state may instead mature by virtue of decreasing activity.

Use of RNNs to study neural processes opens up a wide array of opportunities for the study of cognitive development. In all simulations presented in the paper, we included an “early” training phase, meant to simulate the pre-adolescent state of the PFC. By definition, the immature brain is unable to execute cognitive tasks at full capacity. For this reason, teaching young subjects (humans and monkeys) tasks presents challenges. Scant neurophysiological findings are available from preadolescent subjects for this reason. Our findings make a number of predictions for immature PFC activity to be tested in future experiments: We posit that, in the very young PFC, levels of delay period activity are even lower than adolescent ones (Figure 4), and it is less able to represent different stimuli (Figure 5); distractors interrupt more severely activity representing remembered stimuli (Figure 4E); and the antisaccade task elicits even lower levels of activity than in the adolescent state (Figure 6). Furthermore, the analysis of Figure S2 suggests an abrupt transition in working memory abilities, rather than a continuous improvement.

Further use of the RNNs training paradigm that we introduce here will also allow us to simulate a wide range of cognitive tasks and to identify ones that exhibit the widest contrasts between early and late training

phases. These can then be implemented experimentally, and probed in terms of behavior and neurophysiology, allowing us to uncover hitherto unknown processes related to cognitive maturation.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
 - Behavioral and neurophysiological data
- **METHOD DETAILS**
 - Recurrent neural networks
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Response dynamics
 - Training and testing with different delay intervals

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2021.103178>.

ACKNOWLEDGMENTS

Research reported in this paper was supported by the National Institute of Mental Health of the US National Institutes of Health under award numbers R01 MH117996 and R01 MH116675. We wish to thank Yuanqi Xie for technical help and Balbir Singh and Jiajie Xiao for helpful comments on the manuscript.

AUTHOR CONTRIBUTIONS

X.Z. and C.C. designed research; Y.H.L. and X.Z. performed simulations; J.Z. and X.Z. performed analysis of neural data; Y.H.L., X.Z., and C.C. wrote the paper.

DECLARATION OF INTEREST

The authors declare no competing interests.

Received: April 14, 2021

Revised: June 16, 2021

Accepted: September 22, 2021

Published: October 22, 2021

REFERENCES

- Barak, O., Sussillo, D., Romo, R., Tsodyks, M., and Abbott, L.F. (2013). From fixed points to chaos: three models of delayed discrimination. *Prog. Neurobiol.* *103*, 214–222. <https://doi.org/10.1016/j.pneurobio.2013.02.002>.
- Bashivan, P., Kar, K., and DiCarlo, J.J. (2019). Neural population control via deep image synthesis. *Science* *364*, eaav9436. <https://doi.org/10.1126/science.aav9436>.
- Bi, Z., and Zhou, C. (2020). Understanding the computation of time using neural network models. *Proc. Natl. Acad. Sci. U. S. A.* *117*, 10530–10540. <https://doi.org/10.1073/pnas.1921609117>.
- Bunge, S.A., Dudukovic, N.M., Thomason, M.E., Vaidya, C.J., and Gabrieli, J.D. (2002). Immature frontal lobe contributions to cognitive control in children: evidence from fMRI. *Neuron* *33*, 301–311.
- Burgund, E.D., Lugar, H.M., Miezin, F.M., Schlaggar, B.L., and Petersen, S.E. (2006). The development of sustained and transient neural activity. *Neuroimage* *29*, 812–821.
- Cadena, S.A., Denfield, G.H., Walker, E.Y., Gatys, L.A., Tolia, A.S., Bethge, M., and Ecker, A.S. (2019). Deep convolutional models improve predictions of macaque V1 responses to natural images. *Plos Comput. Biol.* *15*, e1006897. <https://doi.org/10.1371/journal.pcbi.1006897>.
- Cao, Y., Summerfield, C., and Saxe, A. (2020). Characterizing Emergent Representations in a Space of Candidate Learning Rules for Deep Networks (Vancouver).
- Chugani, H.T., Phelps, M.E., and Mazziotta, J.C. (1987). Positron emission tomography study of human brain functional development. *Ann. Neurol.* *22*, 487–497.
- Cichy, R.M., and Kaiser, D. (2019). Deep neural networks as scientific models. *Trends Cogn. Sci.* *23*, 305–317. <https://doi.org/10.1016/j.tics.2019.01.009>.
- Constantinidis, C., Funahashi, S., Lee, D., Murray, J.D., Qi, X.L., Wang, M., and Arnsten, A.F.T. (2018). Persistent spiking activity underlies working memory. *J. Neurosci.* *38*, 7020–7028. <https://doi.org/10.1523/JNEUROSCI.2486-17.2018>.
- Constantinidis, C., and Luna, B. (2019). Neural substrates of inhibitory control maturation in

- adolescence. *Trends Neurosci.* 42, 604–616. <https://doi.org/10.1016/j.tics.2019.07.004>.
- Cueva, C.J., Saez, A., Marcos, E., Genovesio, A., Jazayeri, M., Romo, R., Salzman, C.D., Shadlen, M.N., and Fusi, S. (2020). Low-dimensional dynamics for working memory and time encoding. *Proc. Natl. Acad. Sci. U. S. A.* 117, 23021–23032. <https://doi.org/10.1073/pnas.1915984117>.
- Davidson, M.C., Amso, D., Anderson, L.C., and Diamond, A. (2006). Development of cognitive control and executive functions from 4 to 13 years: evidence from manipulations of memory, inhibition, and task switching. *Neuropsychologia* 44, 2037–2078. <https://doi.org/10.1016/j.neuropsychologia.2006.02.006>.
- de Vries, I.E.J., Slagter, H.A., and Olivers, C.N.L. (2020). Oscillatory control over representational states in working memory. *Trends Cogn. Sci.* 24, 150–162. <https://doi.org/10.1016/j.tics.2019.11.006>.
- Diler, R.S., Segreti, A.M., Ladouceur, C.D., Almeida, J.R., Birmaher, B., Axelson, D.A., Phillips, M.L., and Pan, L. (2013). Neural correlates of treatment in adolescents with bipolar depression during response inhibition. *J. Child Adolesc. Psychopharmacol.* 23, 214–221. <https://doi.org/10.1089/cap.2012.0054>.
- Feulner, B., and Clopath, C. (2021). Neural manifold under plasticity in a goal driven learning behaviour. *Plos Comput. Biol.* 17, e1008621. <https://doi.org/10.1371/journal.pcbi.1008621>.
- Fry, A.F., and Hale, S. (2000). Relationships among processing speed, working memory, and fluid intelligence in children. *Biol. Psychol.* 54, 1–34.
- Gathercole, S.E., Pickering, S.J., Ambridge, B., and Wearing, H. (2004). The structure of working memory from 4 to 15 years of age. *Dev. Psychol.* 40, 177–190.
- Hasson, U., Nastase, S.A., and Goldstein, A. (2020). Direct fit to nature: an evolutionary perspective on biological and artificial neural networks. *Neuron* 105, 416–434. <https://doi.org/10.1016/j.neuron.2019.12.002>.
- Jernigan, T.L., Archibald, S.L., Berhow, M.T., Sowell, E.R., Foster, D.S., and Hesselink, J.R. (1991). Cerebral structure on MRI, Part I: localization of age-related changes. *Biol. Psychiatry* 29, 55–67.
- Khaligh-Razavi, S.M., Henriksson, L., Kay, K., and Kriegeskorte, N. (2017). Fixed versus mixed RSA: explaining visual representations by fixed and mixed feature sets from shallow and deep computational models. *J. Math. Psychol.* 76, 184–197. <https://doi.org/10.1016/j.jmp.2016.10.007>.
- Kim, R., and Sejnowski, T.J. (2021). Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks. *Nat. Neurosci.* 24, 129–139. <https://doi.org/10.1038/s41593-020-00753-w>.
- Kingma, D.P., and Ba, J. (2015). Adam: A Method for Stochastic Optimization, arXiv, 1412.6980.
- Klingberg, T., Forssberg, H., and Westerberg, H. (2002). Training of working memory in children with ADHD. *J. Clin. Exp. Neuropsychol.* 24, 781–791.
- Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C.E., Kepecs, A., Mainen, Z.F., Qi, X.L., Romo, R., Uchida, N., and Machens, C.K. (2016). Demixed principal component analysis of neural population data. *Elife* 5, e10989. <https://doi.org/10.7554/eLife.10989>.
- Kramer, A.F., de Sather, J.C., and Cassavaugh, N.D. (2005). Development of attentional and oculomotor control. *Dev. Psychol.* 41, 760–772.
- Kwon, H., Reiss, A.L., and Menon, V. (2002). Neural basis of protracted developmental changes in visuo-spatial working memory. *Proc. Natl. Acad. Sci. U. S. A.* 99, 13336–13341.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. <https://doi.org/10.1038/nature14539>.
- Luna, B., Thulborn, K.R., Munoz, D.P., Merriam, E.P., Garver, K.E., Minshew, N.J., Keshavan, M.S., Genovese, C.R., Eddy, W.F., and Sweeney, J.A. (2001). Maturation of widely distributed brain function subserves cognitive development. *Neuroimage* 13, 786–793.
- Luna, B., Velanova, K., and Geier, C.F. (2008). Development of eye-movement control. *Brain Cogn.* 68, 293–308. <https://doi.org/10.1016/j.bandc.2008.08.019>.
- Lundqvist, M., Herman, P., and Miller, E.K. (2018). Working memory: delay activity, yes! Persistent activity? Maybe not. *J. Neurosci.* 38, 7013–7019. <https://doi.org/10.1523/JNEUROSCI.2485-17.2018>.
- Lundqvist, M., Rose, J., Herman, P., Brincat, S.L., Buschman, T.J., and Miller, E.K. (2016). Gamma and beta bursts underlie working memory. *Neuron* 90, 152–164. <https://doi.org/10.1016/j.neuron.2016.02.028>.
- Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. <https://doi.org/10.1038/nature12742>.
- Masse, N.Y., Yang, G.R., Song, H.F., Wang, X.J., and Freedman, D.J. (2019). Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nat. Neurosci.* 22, 1159–1167. <https://doi.org/10.1038/s41593-019-0414-3>.
- McDowell, J.E., Brown, G.G., Paulus, M., Martinez, A., Stewart, S.E., Dubowitz, D.J., and Braff, D.L. (2002). Neural correlates of refixation saccades and antisaccades in normal and schizophrenia subjects. *Biol. Psychiatry* 51, 216–223.
- Meyer, T., and Constantinidis, C. (2005). A software solution for the control of visual behavioral experimentation. *J. Neurosci. Methods* 142, 27–34.
- Michaels, J.A., Schaffelhofer, S., Agudelo-Toro, A., and Scherberger, H. (2020). A goal-driven modular neural network predicts parietofrontal neural dynamics during grasping. *Proc. Natl. Acad. Sci. U. S. A.* 117, 32124–32135. <https://doi.org/10.1073/pnas.2005087117>.
- Miller, E.K., Lundqvist, M., and Bastos, A.M. (2018). Working memory 2.0. *Neuron* 100, 463–475. <https://doi.org/10.1016/j.neuron.2018.09.023>.
- Mischel, W., Shoda, Y., and Rodriguez, M.I. (1989). Delay of gratification in children. *Science* 244, 933–938.
- Olesen, P.J., Macoveanu, J., Tegner, J., and Klingberg, T. (2007). Brain activity related to working memory and distraction in children and adults. *Cereb. Cortex* 17, 1047–1054.
- Olesen, P.J., Nagy, Z., Westerberg, H., and Klingberg, T. (2003). Combined analysis of DTI and fMRI data reveals a joint maturation of white and grey matter in a fronto-parietal network. *Brain Res. Cogn. Brain Res.* 18, 48–57.
- Ordaz, S.J., Foran, W., Velanova, K., and Luna, B. (2013). Longitudinal growth curves of brain function underlying inhibitory control through adolescence. *J. Neurosci.* 33, 18109–18124. <https://doi.org/10.1523/JNEUROSCI.1741-13.2013>.
- Orhan, A.E., and Ma, W.J. (2019). A diverse range of factors affect the nature of neural representations underlying short-term memory. *Nat. Neurosci.* 22, 275–283. <https://doi.org/10.1038/s41593-018-0314-y>.
- Pandarinath, C., O’Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* 15, 805–815. <https://doi.org/10.1038/s41592-018-0109-9>.
- Pfefferbaum, A., Mathalon, D.H., Sullivan, E.V., Rawles, J.M., Zipursky, R.B., and Lim, K.O. (1994). A quantitative magnetic resonance imaging study of changes in brain morphology from infancy to late adulthood. *Arch. Neurol.* 51, 874–887.
- Pospisil, D.A., Pasupathy, A., and Bair, W. (2018). ‘Artphysiology’ reveals V4-like shape tuning in a deep network trained for image classification. *Elife* 7, e38242. <https://doi.org/10.7554/eLife.38242>.
- Rajalingham, R., Issa, E.B., Bashivan, P., Kar, K., Schmidt, K., and DiCarlo, J.J. (2018). Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *J. Neurosci.* 38, 7255–7269. <https://doi.org/10.1523/JNEUROSCI.0388-18.2018>.
- Riley, M.R., and Constantinidis, C. (2016). Role of prefrontal persistent activity in working memory. *Front. Syst. Neurosci.* 9, 181. <https://doi.org/10.3389/fnsys.2015.00181>.
- Roux, F., and Uhlhaas, P.J. (2014). Working memory and neural oscillations: alpha-gamma versus theta-gamma codes for distinct WM information? *Trends Cogn. Sci.* 18, 16–25. <https://doi.org/10.1016/j.tics.2013.10.010>.
- Satterthwaite, T.D., Wolf, D.H., Erus, G., Ruparel, K., Elliott, M.A., Gennatas, E.D., Hopson, R., Jackson, C., Prabhakaran, K., Bilker, W.B., et al. (2013). Functional maturation of the executive system during adolescence. *J. Neurosci.* 33,

16249–16261. <https://doi.org/10.1523/JNEUROSCI.2345-13.2013>.

Saxe, A., Nelli, S., and Summerfield, C. (2021). If deep learning is the answer, what is the question? *Nat. Rev. Neurosci.* 22, 55–67. <https://doi.org/10.1038/s41583-020-00395-8>.

Simmonds, D.J., Hallquist, M.N., and Luna, B. (2017). Protracted development of executive and mnemonic brain systems underlying working memory in adolescence: a longitudinal fMRI study. *Neuroimage* 157, 695–704. <https://doi.org/10.1016/j.neuroimage.2017.01.016>.

Sinz, F.H., Pitkow, X., Reimer, J., Bethge, M., and Tolias, A.S. (2019). Engineering a less artificial intelligence. *Neuron* 103, 967–979. <https://doi.org/10.1016/j.neuron.2019.08.034>.

Smyrnis, N., Malogiannis, I.A., Evdokimidis, I., Stefanis, N.C., Theleritis, C., Vaidakis, A., Theodoropoulou, S., and Stefanis, C.N. (2004). Attentional facilitation of response is impaired for antisaccades but not for saccades in patients with schizophrenia: implications for cortical dysfunction. *Exp. Brain Res.* 159, 47–54. <https://doi.org/10.1007/s00221-004-1931-0>.

Song, H.F., Yang, G.R., and Wang, X.J. (2017). Reward-based training of recurrent neural networks for cognitive and value-based tasks. *Elife* 6, e21492. <https://doi.org/10.7554/eLife.21492>.

Sowell, E.R., Delis, D., Stiles, J., and Jernigan, T.L. (2001). Improved memory functioning and frontal lobe maturation between childhood and adolescence: a structural MRI study. *J. Int. Neuropsychol. Soc.* 7, 312–322.

Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of

muscle activity. *Nat. Neurosci.* 18, 1025–1033. <https://doi.org/10.1038/nn.4042>.

Ullman, H., Almeida, R., and Klingberg, T. (2014). Structural maturation and brain activity predict future working memory capacity during childhood development. *J. Neurosci.* 34, 1592–1598. <https://doi.org/10.1523/JNEUROSCI.0842-13.2014>.

Wang, X.J., Tegner, J., Constantinidis, C., and Goldman-Rakic, P.S. (2004). Division of labor among distinct subtypes of inhibitory neurons in a cortical microcircuit of working memory. *Proc. Natl. Acad. Sci. U. S. A.* 101, 1368–1373. <https://doi.org/10.1073/pnas.03053371010305337101>.

Warnberg, E., and Kumar, A. (2019). Perturbing low dimensional activity manifolds in spiking neuronal networks. *Plos Comput. Biol.* 15, e1007074. <https://doi.org/10.1371/journal.pcbi.1007074>.

Wenliang, L.K., and Seitz, A.R. (2018). Deep neural networks for modeling visual perceptual learning. *J. Neurosci.* 38, 6028–6044. <https://doi.org/10.1523/JNEUROSCI.1620-17.2018>.

P.I. Yakovlev, and A.R. Lecours, eds. (1967). *The Myelogenetic Cycles of Regional Maturation of the Brain* (Blackwell).

Yamins, D.L., and DiCarlo, J.J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* 19, 356–365. <https://doi.org/10.1038/nn.4244>.

Yang, G.R., Joglekar, M.R., Song, H.F., Newsome, W.T., and Wang, X.J. (2019). Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* 22, 297–306. <https://doi.org/10.1038/s41593-018-0310-2>.

Yang, G.R., and Wang, X.J. (2020). Artificial neural networks for neuroscientists: a primer. *Neuron*

107, 1048–1070. <https://doi.org/10.1016/j.neuron.2020.09.005>.

Zhou, X., Katsuki, F., Qi, X.L., and Constantinidis, C. (2012). Neurons with inverted tuning during the delay periods of working memory tasks in the dorsal prefrontal and posterior parietal cortex. *J. Neurophysiol.* 108, 31–38. <https://doi.org/10.1152/jn.01151.2011>.

Zhou, X., Qi, X.L., and Constantinidis, C. (2016a). Distinct roles of the prefrontal and posterior parietal cortices in response inhibition. *Cell Rep* 14, 2765–2773. <https://doi.org/10.1016/j.celrep.2016.02.072>.

Zhou, X., Zhu, D., Katsuki, F., Qi, X.L., Lees, C.J., Bennett, A.J., Salinas, E., Stanford, T.R., and Constantinidis, C. (2014). Age-dependent changes in prefrontal intrinsic connectivity. *Proc. Natl. Acad. Sci. U. S. A.* 111, 3853–3858. <https://doi.org/10.1073/pnas.1316594111>.

Zhou, X., Zhu, D., King, S.G., Lees, C.J., Bennett, A.J., Salinas, E., Stanford, T.R., and Constantinidis, C. (2016b). Behavioral response inhibition and maturation of goal representation in prefrontal cortex after puberty. *Proc. Natl. Acad. Sci. U. S. A.* 113, 3353–3358. <https://doi.org/10.1073/pnas.1518147113>.

Zhou, X., Zhu, D., Qi, X.L., Lees, C.J., Bennett, A.J., Salinas, E., Stanford, T.R., and Constantinidis, C. (2013). Working memory performance and neural activity in the prefrontal cortex of peri-pubertal monkeys. *J. Neurophysiol.* 110, 2648–2660. <https://doi.org/10.1152/jn.00370.2013>.

Zhou, X., Zhu, D., Qi, X.L., Li, S., King, S.G., Salinas, E., Stanford, T.R., and Constantinidis, C. (2016c). Neural correlates of working memory development in adolescent primates. *Nat. Commun.* 7, 13423. <https://doi.org/10.1038/ncomms13423>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
<i>Macaca mulatta</i>	Worldwide primates	N/A
Software and algorithms		
Tensorflow	https://www.tensorflow.org/	v. 1.9 or higher
MATLAB	https://www.mathworks.com/	R2021a
Python	https://www.python.org/	v. 3.0 or higher
RNN code for analysis	Zenodo	https://doi.org/10.5281/zenodo.5518223

RESOURCE AVAILABILITY

Lead contact

- Further information and requests for resources and code should be directed to and will be fulfilled by the lead contact, Xin Zhou (maizie.zhou@vanderbilt.edu).

Materials availability

N/A.

Data and code availability

- All data reported in this paper will be shared by the lead contact upon request.
- All original code has been deposited at Zenodo and is publicly available as of the date of publication. DOIs are listed in the key resources table.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Behavioral and neurophysiological data

We relied on analysis of behavioral and neurophysiological results from monkeys performing working memory and response inhibition tasks, as these have been described in detail previously (Zhou et al., 2016a, 2016b, 2016c). Briefly, the dataset was collected from four male rhesus monkeys (*Macaca mulatta*) with a median age of 4.3 years at the adolescent stage (range: 4.0–5.2 years), and 6.3 years at the adult stage (range: 5.6–7.3). All surgical and animal use procedures were reviewed and approved by the Wake Forest University Institutional Animal Care and Use Committee, in accordance with the U.S. Public Health Service Policy on humane care and use of laboratory animals and the National Research Council's Guide for the care and use of laboratory animals. Development was tracked on a quarterly basis before, during, and after neurophysiological recordings allowing us to ascertain that data were collected during two stages: in adolescence, and adulthood (Zhou et al., 2016b, 2016c).

The monkeys were trained to perform two variants of the Oculomotor Delayed Response (ODR) task and three different variants of the antisaccade task (Figure 1). The ODR task is a spatial working memory task requiring subjects to remember the location of a cue stimulus flashed on a screen for 0.5 s. The cue was a 1° white square stimulus that could appear at one of eight locations arranged on a circle of 10° eccentricity. After a 1.5 s delay period, the fixation point was extinguished and the monkey was trained to make an eye movement to the remembered location of the cue within 0.6 s. In the ODR with distractor variant (ODRD), a second stimulus appeared after a 0.5 s delay, followed by a second delay period. The monkey still needed to saccade to the original remembered stimulus. In the antisaccade task, each trial starts with the monkey fixating a central green point on the screen. After 1s fixation, the cue appears, consisting of a 1° white square stimulus that could appear at one of four locations arranged on a circle of 10° eccentricity for

0.1 s. The monkey is required to make a saccade at the location diametric to the cue. The saccade needed to terminate on a 5–6° radius window centered on the stimulus (within 3–4° from the edge of the stimulus), and the monkey was required to hold fixation within this window for 0.1 s. We used three different variants for the antisaccade task: overlap, zero gap, and gap, differing in the sequence of the cue onset relative to the fixation point turning off (Figure 1A). In the overlap condition, the cue appears first, and then fixation point and cue are simultaneously extinguished. In the zero gap condition, the fixation turns off and the cue turns on at the same time. In the gap condition, the fixation turns off and a 100 or 200 ms blank screen is inserted before the cue onset. Visual stimuli display, monitoring of eye position, and the synchronization of stimuli with neurophysiological data were performed with in-house software (Meyer and Constantinidis, 2005) implemented on the MATLAB environment (Mathworks, Natick, MA). Behavioral and neural results were collected from the same animals in the young and adult stage.

Neural recordings were collected with epoxy-coated Tungsten electrodes with a diameter of 250 μm and an impedance of 4 MΩ at 1 KHz (FHC Bowdoin, ME). Electrical signals recorded from the brain were amplified, band-pass filtered between 500 and 8 kHz, and stored through a modular data acquisition system at 25 μs resolution (APM system, FHC, Bowdoin, ME). Recordings analyzed here were obtained from areas 8a and 46 of the dorsolateral prefrontal cortex. Recorded spike waveforms were sorted into separate units using an automated cluster analysis method based on the KlustaKwik algorithm. Firing rate of units was then determined by averaging spikes in each task epoch. A total of 830 neurons from the adult stage, and 607 neurons from the young stage were used for subsequent analysis. In the ODR task, we identified neurons with significant elevation of firing rate in the 500 ms presentation of the cue, the 1500 ms delay period, and the 250 ms response epoch, after the fixation point turned off. Firing rate in this period was compared to the 1 s baseline fixation period, prior to the presentation of the cue, and neurons with significant difference in firing rate were identified (paired t test, $p < 0.05$). Responsive neurons in the antisaccade task were identified based on significantly elevated responses in the 250 ms window following the onset of the cue compared to the fixation interval.

METHOD DETAILS

Recurrent neural networks

We trained Leaky, Recurrent Neural Networks (RNNs) to perform multiple tasks simultaneously, following established procedures used to train RNNs to perform tasks used in neurophysiological experiments (Yang et al., 2019). Implementation was based in Python 3.6, using the TensorFlow package. We used time-discretized RNNs with positive activity, modelled with the following equation (before descretization):

$$\tau \frac{dr}{dx} = -r + f\left(W^{rec}r + W^{in}u + \mathbf{b} + \sqrt{2\tau\sigma_{rec}^2}\xi\right)$$

Where τ is the neuronal time constant (set to 100 ms in our simulations), u the input to the network, \mathbf{b} the background input, f the neuronal nonlinearity, ξ a vector of independent white noise process with zero mean and σ_{rec} the strength of noise (set to 0.05). Each time step (iteration) in our implementation represented 20 ms. That mean that a delay of 1.5 s was represented in 75 timesteps. Networks typically contained 256 units (results shown in main figures). Results from networks with numbers of units ranging from 64 to 1024 are shown in the supplementary material. We modeled the neuron nonlinearity based on the Softplus function

$$f(x) = \log(1 + e^x)$$

Results of networks with other functions are shown in the supplementary material, including ReLU (rectified, linear function),

$$f(x) = \max(x, 0)$$

tanh (hyperbolic tangent),

$$f(x) = \tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

and retanh (rectified hyperbolic tangent):

$$f(x) = \max(\tanh(x), 0)$$

Output units, \mathbf{z} read out the non-linearity from the network as:

$$\mathbf{z} = g(W^{out}r),$$

where $g(x)$ is the logistic function

$$g(x) = \frac{1}{1 + e^{-x}}$$

and W^{out} the weights of units connected to the output units.

Our networks received three types of noisy input: fixation, visual stimulus location, and task rule. The weights of the recurrent unit matrix (W^{rec}) were initialized with random orthogonal initialization (main figures), as well as diagonal, and random Gaussian values (supplementary material). Initial input weights (W^{in}) were drawn from a standard normal distribution divided by the square root of the unit's number. Initial output weights (W^{out}) were initialized with the Xavier uniform initializer. The background input vector (\mathbf{b}) was initialized with zero values.

To train an RNN to perform the working memory tasks and anti-saccade tasks, we used a three-dimensional tensor as the input to the network. This fully described the sequence of events in the six behavioral tasks used: ODR, ODRD, overlap, zero-gap, and two gap tasks (100 and 500 ms). The first dimension of the tensor encodes the noisy inputs of three types: fixation, stimulus location, and task rule. Fixation input was modeled as a binary input of either 1 (meaning the subject needs to fixate) or 0, otherwise. The stimulus is considered to appear at a ring of fixed eccentricity, and its location is fully determined by the angular dimension. Therefore, stimulus inputs consist of a ring of 8 units, with preferred directions uniformly spaced between 0 and 2π . The rule of the task was represented as a one-hot vector with a value of 1 representing the current task the subject is required to perform and 0 for all other possible tasks. The rule input unit corresponding to the current task was activated throughout the whole trial. Simulations involved 6 task inputs, therefore the first dimension of the tensor had a total of $1 + 8 + 6 = 15$ inputs. The second dimension of the tensor encoded the batch size (number of trials). The third dimension encoded the time series for each trial.

A ring of 8 output units (plus one fixation output unit) similarly indicated the direction of gaze at each time point in the trial. While the fixation point was on, the fixation output unit should produce high activity. Once the fixation input was off, the subject had to make an eye movement (e.g. in the direction of the original stimulus in the ODR, or in the opposite direction in the antisaccade task), which was represented by activity in the network of tuned output units. The response direction of the network was read out using a population vector method. A trial is considered correct only if the network correctly maintained fixation (fixation output unit remained at a value >0.5) and the network responded within 36° of the target direction.

Each task can be separated into distinct epochs which duration was equal to the duration of the corresponding task in the neurophysiological experiments, described above. Fixation (fix) epoch is the period before any stimulus is shown, and lasted for 1 s. It was followed by the stimulus epoch 1 (stim1), which was equal to 0.5 s for the ODR and ODRD tasks, and 0.1 s for the antisaccade tasks. If there are two stimuli separated in time, or if the stimulus and response are separated in time, then the period between the two is the delay epoch (1.5 s in the ODR task, 0.5 s in the ODRD task) and the second stimulus is epoch 2 (stim2, equal to 0.5 s). The period when the network should respond is the go epoch.

The RNNs are trained with supervised learning, based on variants of stochastic gradient descent, which modifies all connection weights (input, recurrent and output) to minimize a cost function L representing the difference between the network output and a desired (target) output (Yang and Wang, 2020). We relied on the Adam optimization algorithm (Kingma and Ba, 2015) to update network weights iterative based in training data. For each step of training, the loss is computed using a small number M of randomly selected training examples, or minibatch. Trials representing all six tasks were included in a single minibatch during training of our networks. Trainable parameters, collectively denoted as θ are updated in the opposite direction of the gradient of the loss, with a magnitude proportional to the learning rate η :

$$\Delta\theta = -\eta \frac{\partial L}{\partial \theta}$$

We found that the ability of the networks to master the task was quite sensitive to the value of η . This was set to 0.001 for most of the simulations included in the paper.

The activity of recurrent units was read out at discrete time points representing 20 ms bins. These can be compared with “Peri-stimulus Time histograms” of real neurons. Different stages of network training were used to simulate different developmental stages or task-training phases. For most analyses, we used both mixed training (the six tasks were randomly interleaved during training) and independent training (the network learned one task rule each time). We defined three “developmental” stages as the trial in which the network achieved <35%, 35-65% and >65% correct performance, respectively. Networks were allowed to continue training, until they reached 100% performance, when no errors were detected and hence weights would no longer be updated. Firing rates of individual units and averages across units are typically presented as normalized rates, obtained by subtracting the unit’s baseline firing rate (obtained in the 1 s period prior to the appearance of the cue) from the unit’s raw firing rate during the course of the trial, and dividing by the same baseline firing rate.

QUANTIFICATION AND STATISTICAL ANALYSIS

Response dynamics

Demixed Principal Component Analysis (dPCA) was performed, as we have described elsewhere (Kobak et al., 2016). The input is a four-dimensional array which represents number of neurons, discharge rate within a 20ms bin, stimulus locations, and trials for each stimulus condition. The method treats the responses of each neuron across time and stimulus conditions as one dimension, and then performs dimensionality reduction to determine components that correspond to stimulus and task variables. dPCA was performed independently on groups of RNN units and PFC neurons recorded from different maturation stages.

Discharge rate across the length of the trial was calculated at every 200 ms window bins with an incremental step of 20 ms, and data from eight trials for the eight stimulus locations were then used for the analysis. Data from 256 RNN units were always used for RNN simulations. Decoding analysis was performed on the same dataset relying on an SVM classifier. The classifier was trained with a subset of trials at one window bin, and then tested with the remaining trials at the same window bin, and other window bins.

To visualize the timing of unit activation across the network, we sorted units by the time bin in which they achieved their peak firing rate, then plotted their rate in a uniform scale by dividing their firing rate in each time bin divided by each unit’s peak firing rate.

Training and testing with different delay intervals

Fully trained networks with the standard 1.5 s delay period for the ODR task were tested with a 3.0 s delay period and vice versa. For these simulations, only the duration of the delay interval during the test runs was altered, while the structure of the network and representation of stimuli and output remained the same. Additionally, we trained RNNs with a variable delay interval. In these networks, each minibatch of the ODR task contained 5 delay intervals varying from 0.5 to 2.5 seconds, spaced 0.5 s apart. Once the network was fully trained, it was then tested with different delay periods, including a 3.0 s delay, which was longer than any interval trained with.