HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2020 March 23.

Published in final edited form as:

Nat Neurosci. 2019 October ; 22(10): 1687–1695. doi:10.1038/s41593-019-0487-z.

Layer-dependent activity in human prefrontal cortex during working memory

Emily S. Finn[1,*], **Laurentius Huber**[1,2], **David C. Jangraw**[1], **Peter J. Molfese**[1], **Peter A. Bandettini**[1]

[1]Section on Functional Imaging Methods, Laboratory of Brain and Cognition, National Institute of Mental Health; Bethesda, Maryland, USA [2]Present address: MR-Methods Group, Maastricht Brain Imaging Center, Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience; Maastricht University, Maastricht, The Netherlands

Abstract

Working memory involves storing and/or manipulating previously encoded information over a short-term delay period, which is typically followed by a behavioral response based on the remembered information. While working memory tasks often engage dorsolateral prefrontal cortex (dlPFC), few studies have investigated whether their sub-processes are localized to different cortical depths in this region, and none have done so in humans. Here, we use high-resolution functional MRI to interrogate the layer specificity of neural activity during different periods of a delayed-response task in dlPFC. We detect activity timecourses that follow the hypothesized patterns: namely, superficial layers are preferentially active during the delay period, and specifically in trials requiring manipulation (rather than mere maintenance) of information held in working memory, while deeper layers are preferentially active during the response. Results demonstrate that layer-specific fMRI can be used in higher-order brain regions to non-invasively map cognitive processing in humans.

INTRODUCTION

Working memory (WM) is the highly evolved mental capacity to store and manipulate information for short-term use. It is often probed with delayed-response tasks that require encoding a stimulus, sustaining a representation of the stimulus over a delay, and finally making a memory-guided behavioral response.

[*]**Corresponding author:** emily.finn@nih.gov.

AUTHOR CONTRIBUTIONS

E.S.F. conceptualized the study, designed the task paradigm, collected the data, analyzed the data, generated visualizations, and wrote the original manuscript. L.H. conceptualized the study, contributed to the task paradigm design, designed and optimized the data acquisition and analysis methodology, collected the data, analyzed the data, generated visualizations, and wrote portions of the original manuscript. D.C.J. contributed to the task paradigm design, generated visualizations, and provided comments on the manuscript. P.J.M. analyzed the data, generated visualizations, and provided comments on the manuscript. P.A.B. supervised study conceptualization, design, and interpretation, and provided comments on the manuscript.

COMPETING INTERESTS STATEMENT

The authors declare no competing interests.

The dorsolateral prefrontal cortex (dlPFC) has been linked to WM processes in both humans and non-human primates [1–4]. Like much of the cerebral cortex, dlPFC gray matter is organized into layers with distinct cytoarchitecture, connectivity and function. Early electrophysiological work in non-human primates suggested that in delayed-response tasks, different task periods are preferentially associated with activity in different cortical layers [5,6]. Specifically, delay-period activity is thought to be driven by recurrently connected networks of pyramidal cells in layer III [3], while response-related activity takes place predominantly in layer V [7] Two recent studies in macaques, which overcame the challenge of separating activity recorded from distinct cortical layers, provide direct evidence for this dissociation [8,9]

However, it remains unclear to what extent dlPFC exhibits homologous function between monkeys and humans. While dlPFC often appears active during WM tasks in human functional MRI studies, human dlPFC may not be strictly necessary for mere maintenance of information—that is, for sustaining the representation of a stimulus “as-is” without performing further operations on it. Instead, dlPFC may be necessary only when the task calls for rule-based manipulation of information stored in WM; for example, when items must be reordered or transformed in some other way. Indeed, disrupting dlPFC activity with lesions [10,11] or repetitive transcranial magnetic stimulation (rTMS) [12] impairs manipulation, but leaves maintenance largely intact.

To the extent that human dlPFC is specialized for manipulation rather than pure maintenance, the laminar specificity of these operations is unknown. Following an evolutionary progression, we hypothesize that manipulation in humans might recruit the same local recurrent excitatory networks of layer III pyramidal cells as maintenance does in non-human primates. This hypothesis is also supported by converging evidence from schizophrenia, which is associated with reduced dendritic spine density specifically in dlPFC layer III neurons [13,14] as well as behavioral deficits in manipulation (over and above maintenance)[15]. On the other hand, activity involved in response selection and action initiation may take place predominantly in infragranular layers, as has been observed in non-human primates [16,17,9]. To date there is no empirical evidence for such a dissociation in humans, largely because conventional neuroimaging techniques lack the sensitivity and specificity to resolve cortical layers.

Recent methodological advances in fMRI, including higher field strengths (i.e., 7 Tesla and above) combined with innovations in pulse sequences and contrast mechanisms, now allow for non-invasive, reliable measurements of cortical depth-dependent activity in humans. These advances have enabled layer-specific imaging in several primary cortices, including visual [18–20], auditory [21], and motor [22]. (Note that in the context of fMRI, the term ‘layer’ refers to estimates of different cortical depths, not necessarily to cortical layers as defined cytoarchitectonically.) While simulations suggest that fMRI should in principle be able to resolve laminar differences in more complex tasks [23], it is still unclear if these techniques are sensitive and robust enough to be applied outside primary cortices.

Here, by further developing layer-fMRI methods to move beyond unimodal cortex[22] into higher-order areas, we provide evidence for cortical depth-dependent processing during a

sophisticated cognitive task in one of the most highly evolved regions of human association cortex. Specifically, we use simultaneously acquired blood oxygen level-dependent (BOLD) and cerebral blood volume (CBV) images of human dlPFC during a working memory task to show that during the delay period, manipulation evokes greater activity than maintenance specifically in superficial layers, while during the response period, activity is localized to deeper layers. These results deepen our understanding of the laminar specificity of WM-based operations in humans, and demonstrate the promise of high-resolution fMRI for mapping cognitive cortical circuitry at the mesoscale.

RESULTS

Task paradigm

To test our hypotheses about layer-dependent activity during WM, we used a well-validated task paradigm that dissociates maintenance from manipulation during the delay period[24], and added a second contrast to separate action from non-action during the response period. See Fig. 1A for a schematic of the task. All trials are matched for sensory input, with the only difference being the nature of the mental activity during the delay for the first contrast, or the presence or absence of action selection and execution during the response period for the second contrast. (Note that an action-related signal can also be isolated from the first contrast by examining activity at the time of the response compared to all other timepoints; we exploit this in a second acquisition protocol described further below.)

Thus the main paradigm followed a 2×2×2 design, with trial type (manipulation/maintenance versus action/non-action), period (delay versus response), and cortical depth (superficial versus deep) as the three factors. We hypothesized a triple dissociation between trial type, period, and cortical depth, such that: (1) superficial layers would respond more strongly during the *delay* period of *manipulation* trials (as compared to the delay period of maintenance trials), and (2) deeper layers would respond more strongly during the *response* period of *action* trials (as compared to the response period of non-action trials). See Fig. 1b for a schematic of the hypothesis. The strength of this experimental design is that we control for each layer's timecourse of activity primarily by observing the same layer in a different condition, rather than directly comparing activity levels across layers; this avoids measurement biases associated with different cortical depths.

Data acquisition

Functional data are from n = 15 unique subjects scanned in a combined total of 20 imaging sessions. During each high-resolution functional run, we simultaneously measured changes in cerebral blood volume (CBV) and blood-oxygen-level dependent (BOLD) signal using the SS-SI-vascular space occupancy (VASO) method [25] with a 3D-EPI readout [26] on a 7 Tesla scanner. This method has been implemented to successfully demonstrate layer-specific activity in human motor cortex with good sensitivity and specificity [22]. The conventional BOLD signal has poor spatial specificity at high resolutions, since it tends to be dominated by large veins at the pial surface, and depends on non-linear interactions between physiological variables that can differ across cortical depths, making it difficult to quantitate. VASO, while it has a lower contrast-to-noise ratio, is a more quantitative measurement that

is less biased toward superficial depths. In short, BOLD is more sensitive, while VASO is more specific.

We used two different acquisition protocols over the course of the study. The first had a nominal voxel resolution of 0.9 × 0.9 × 1.1mm (referred to as the "axial [readout] protocol"). These data were used to quantitatively compare activity timecourses across two distinct cortical depths (superficial versus deep) at the group level. Later, we introduced a second, higher resolution protocol with nominal voxel resolution of 0.76 × 0.76 × 0.99mm (referred to as the "sagittal [readout] protocol"). These data were used to visualize activity across different layers in individual subjects. For both protocols, the field of view was not the whole brain but rather a slab centered on a region of interest within left dlPFC that was identified via on online functional localizer conducted at the start of each imaging session. (Due to restrictions on its MRI sequence parameter space, and the need to apply a slab-selective inversion pulse, VASO is currently limited in the spatial coverage that can be achieved at these resolutions.) See Methods and Fig. S1 for further details of our data acquisition and analysis pipeline.

Location of region of interest

Prefrontal cortex is large, and quite variable across individuals in terms of structure and functional anatomy. Unlike other cortical landmarks, such as the 'hand knob' of the primary motor cortex, functional subdivisions of prefrontal cortex are difficult to pinpoint in individual subjects using macroscale anatomical features. Therefore, regions of interest (ROIs) were selected for each subject on the basis of an online functional localizer conducted just prior to the experimental task runs (see Methods). Given that imaging parameters could only be optimized for one hemisphere at a time, we focused on left dlPFC in all subjects, based on previous reports as well as our own pilot experiments indicating that this task more strongly engages the left over the right hemisphere. (Because our stimuli, being letters, were verbal in nature, this lateralization may be due in part to a left-hemisphere dominance for language.)

Despite the variance in prefrontal cortex size and anatomy across subjects, the ROI location was highly consistent with respect to the subject-specific cortical folding structure that was visible in EPI space. In all subjects, the ROI was located in the ventral portion of the middle frontal gyrus corresponding approximately to Brodmann area 9/46[27]. To ensure that our ROI selection procedure was robust, we conducted test-retest scans separated by several weeks on two subjects. Results showed good overlap between ROIs derived from independent experimental sessions (Fig. S2), indicating that the functional region in question can be reliably localized within subjects. Fig. 1c shows the average ROI location across subjects computed from the whole-brain functional localizer (though note that this figure is a post-hoc visualization only; all analyses of the high-resolution experimental data were conducted in single-subject space to preserve spatial specificity). See Supplementary Videos 1–6 for slice-by-slice visualizations of the selected ROI in six individual subjects.

For each subject, two layers, superficial and deep, were each drawn manually within the selected ROI (see Fig. S3 for layer masks for all subjects scanned using the axial readout protocol). To better specify the position of our "superficial" and "deep" layers with respect

to cortical laminae defined cytoarchitectonically, we compared all available MRI-based anatomical contrasts with an existing histological image (Fig. S4). The boundary between our superficial and deeper layers fell approximately between layer III and layer IV.

Task performance

Subjects performed well on the task (overall mean accuracy = 0.82, s.d. = 0.13, range = 0.59 – 0.97; note that chance is approximately 0.2), including both manipulation trials (mean (s.d.), range: 0.79 (0.13), 0.54 – 0.96) and maintenance trials (mean (s.d.), range: 0.88 (0.15), 0.53 – 1.0). Subjects were less accurate on manipulation compared to maintenance trials (paired t-test, $t_{14} = -3.28$, $p = 0.01$), which is expected given previous reports using this task[24].

Overall mean reaction time (RT) was 2.37 s (s.d., range: 1.24, 1.05 – 5.17). Crucially, there was no difference between mean RT on manipulation versus maintenance trials (paired t-test, $t_{14} = 1.29$, $p = 0.22$). It is therefore unlikely that conditions differ in latency of peak response-related activity, allowing us to directly compare timecourses without deconvolution.

Activity timecourses

Using data from 15 experimental sessions (n = 13 unique subjects) scanned with the axial protocol, we observed layer-dependent activity timecourses that followed the hypothesized patterns: in superficial layers, activity was higher in manipulation relative to maintenance trials during the delay period, and in deeper layers, activity was higher in action versus non-action trials during the response period. These patterns were visible in both VASO and BOLD (Fig. 2, Fig. S5). Below we summarize characteristics of these depth-dependent timecourses during the two main periods of interest, delay and response.

Delay-related activity.—In superficial layers (Fig. 2a, top row), delay-period activity was uniformly high during manipulation trials. This is evident in trials labeled ‘alpha’, ‘action’ and ‘non-action’ (recall that both action and non-action trials call for alphabetizing, and they are indistinguishable from one another until the probe appears). Superficial delay-related activity was higher during manipulation than maintenance, although results from the more-sensitive BOLD contrast indicated that maintenance alone was also sufficient to evoke above-baseline activity (Fig. S5). In addition to the group-level results shown in Fig. 2, this effect was clearly visible in single-subject data (Fig. S6).

In contrast to superficial layers, deeper layers were markedly less active during the delay period (Fig. 2a, bottom row; although note that the BOLD data in particular suggest that their activity is still slightly above baseline during this period, Fig. S5). Thus, it seems that delay-related activity occurs predominantly, if not exclusively, in superficial layers, and particularly when task demands call for manipulation of information stored in WM rather than mere maintenance.

Response-related activity.—During the response period, we observe the opposite pattern: activity in deeper layers is high, but only in trials requiring an action. Deep-layer

activity peaks at the time of the response, which is expected at approximately 6-7 seconds after the probe comes onscreen (reflecting behavioral and hemodynamic delay). As expected, this peak is present in action but not non-action trials (Fig. 2a, bottom right). Again, this effect was also visible in most individual subjects (Fig. S7).

As for superficial layers, their activity is, if anything, suppressed at the response peak in both trial types (Fig. 2a, top right). This confirms our prediction that the response period is preferentially associated with activity in deeper cortical layers.

These same patterns were visible to some degree in the BOLD contrast (Fig. S5), although the strong superficial bias of BOLD make it difficult to draw firm conclusions from these data. (For example, the apparent difference between action and non-action trials in superficial layers visible in Fig. S5a, top right is likely an artifact of draining veins from the deeper layers, since this difference is not present at all in the VASO data shown in Fig. 2a, top right.) Due to the higher spatial specificity and more quantitative nature of VASO, we performed all statistical comparisons using this contrast as described in the following section.

Quantification of differential activity

To quantitatively compare activity within cortical depths, we performed a series of two-way, repeated-measures analyses of variance (ANOVAs) using representative signals from each trial type during each trial period. In each ANOVA, the two factors were trial type (either ‘alphabetize’ and ‘remember’, or ‘action’ and ‘non-action’) and trial period (delay and response), with subject as the repeated measure.

For superficial layers, we found a significant interaction between trial type (manipulation versus maintenance, or ‘alphabetize’ versus ‘remember’) and trial period ($F(1,14) = 34.7$, $p = 7.7e^{-5}$), such that activity was higher in manipulation trials, but only during the delay period (Fig. 2b, top left). As expected, the contrast between the second condition pair (action versus non-action) revealed a main effect of period ($F(1,14) = 123.0$, $p = 2.6e^{-8}$), such that activity was higher during the delay than during the response, but no interaction between period and trial type ($F(1,14) = 0.19$, $p = 0.68$; Fig. 2b, top right).

For deeper layers, as predicted, we found the opposite pattern of results. There was a significant interaction between trial type (action versus non-action) and trial period ($F(1,14) = 26.0$, $p = 0.002$), such that activity was higher in action trials during the response (Fig. 2b, bottom right). The contrast between the manipulation and maintenance conditions indicated an interaction such that activity was higher during the response than during the delay, but only in manipulation trials ($F(1,14) = 13.4$, $p = 0.004$; Fig. 2b, bottom left).

Another way to assess relevant differences is to subtract the average timecourse within each depth between the trial types of interest. Results indicate that for superficial layers, the difference between manipulation and maintenance peaks during the delay period (Fig. 3a, top and Fig. S5b, top), while for deeper layers, the difference between action and non-action trials peaks at the time of the response (Fig. 3a, bottom and Fig. S5b, bottom).

As a final quantification step, we statistically compared these differential activity levels by performing ANOVAs on representative signals from each period (delay and response) in each differential time course (manipulation–maintenance and action–non-action), again with subject as the repeated measure (Fig. 3b). While directly comparing superficial and deeper layers should be done with caution as results can be biased by cross-depth differences in baselines, scale factors and vascular cross-talk, in this case we use a difference-of-differences approach that helps mitigate some of these concerns. Results confirm that during both trial periods, there is an interaction between layer and condition pair such that during the delay period, superficial layers are more sensitive to the manipulation–maintenance contrast ($F(1,14) = 92.7$, $p = 6.9e^{-6}$; Fig. 3b, left), while during the response period, deeper layers are more sensitive to the action–non-action contrast ($F(1,14) = 30.5$, $p = 0.0003$; Fig. 3b, right).

Visualization of depth-dependent activity

To better visualize the depth-dependent distribution of activity associated with different periods within the trial, we used a second, higher-resolution imaging protocol in which the field of view was a sagittal slab centered on dlPFC with in-plane resolution of 0.76 × 0.76mm. In these experiments, the task consisted exclusively of manipulation/maintenance trials, all requiring an active response (i.e., the first contrast type shown in Fig. 1a, top). Functional signals during manipulation and maintenance trials were investigated across cortical depths.

We detected layer-dependent activity in all individual subjects imaged using this protocol (n = 5; Fig. 4). Manipulation evoked more activity than maintenance predominantly in superficial layers (green stripes), while signal associated with response (as compared to baseline; red stripes) was predominantly localized to deeper layers. These patterns were visible in both the BOLD (Fig. 4a) and VASO (Fig. 4b) contrasts (though note the different thresholds). Layer ROIs for each subject are shown in Fig. 4c, and a discussion of the observed variance in functional response across the cortical surface (i.e., across columns) is given in Fig. S8.

DISCUSSION

While working memory has been known to engage dlPFC for decades, the degree to which its sub-processes were layer-specific had been hypothesized [3] but demonstrated only a handful of times in non-human primates [9,8]. Furthermore, the extent of functional homology in this region between humans and non-human primates was unclear. Here, we interrogate layer-specific functionality directly and non-invasively in humans, shedding new light on the laminar specificity of WM processes in dlPFC. By developing and optimizing state-of-the-art techniques in high-resolution fMRI for cognitive brain areas, and using a task design for which we had hypotheses about the location, magnitude and timing of neural activity, we were able to detect timecourses at different cortical depths that followed the expected patterns. Namely, we observed delay-related manipulation activity that was predominantly localized to superficial layers, and response-related activity that was predominantly localized to deeper layers.

We interpret the observed laminar specificity of distinct working memory processes in light of what is known about underlying neural circuitry. First, superficial activity during the delay period may at least partially reflect recurrent excitatory connections. While in early parts of the cortical hierarchy, superficial layers give rise to feedforward connections, at the highest levels (i.e., PFC), laminar projections become more complex. Layer III expands and is the focus of extensive local, recurrent excitatory connections [28], as well as long-range recurrent connections with other regions that may be involved in storing items in working memory, e.g., parietal association cortex [7,29] Recurrent excitation among these cells is a feature of their unique molecular profile, notably their preferential expression of n-methyl-d-aspartate (NMDA) receptors and specifically the NR2B subunit, whose slower kinetics allow for persistent firing over long delays; this has been predicted by computational models [30] and confirmed experimentally in primates [31]. While our findings suggest that superficial layers are active specifically when the task calls for manipulating and not just storing information, with our current task design, we cannot fully rule out the possibility that superficial-layer activity depends somewhat on task difficulty or engagement more generally. In future work, designs that parametrically vary load under both manipulation and maintenance conditions will help define the precise functional role of superficial-layer cells in dlPFC.

Second, response-period activity in deeper layers likely reflects functions related to response selection, action execution, or both. In our task paradigm, a response could not be selected until the probe appeared onscreen. This is in keeping with typical delayed-response paradigms used in human neuroimaging, but different from those used with non-human primates, which are based on oculomotor responses to a single remembered item, meaning the animal can predict the upcoming response during the delay period. Human neuroimaging studies suggest a role for dlPFC in selecting and planning an appropriate task response [32–34], even in the absence of a working memory requirement [35]; this activity scales with factors affecting response selection even while eventual motor output is held constant [36], seeming to indicate response selection as the dominant process taking place in dlPFC. On the other hand, non-human primate electrophysiological studies, most notably those featuring laminar specificity [9,8], report deeper-layer activity that appears to track action execution (i.e., saccades) more directly. This activity might reflect one or a number of processes related to motor execution, such as initiating an action, suppressing prepotent responses, or a feedback mechanism such as corollary discharge. While dlPFC does not project directly to primary motor cortex (M1), it may influence motor behavior polysynaptically via higher-order cortical motor areas [37,38] or the striatum [39,40]. Like most delayed-response human fMRI paradigms, our task timing and temporal resolution do not allow us to separate response selection from action initiation itself, meaning future work will be necessary to dissociate these two processes and the extent to which they account for the layer-specific response profiles observed here.

Of note, schizophrenia is associated with altered genetics [7], morphology [14,13] and function [41] in this very dlPFC circuitry. Decreased delay-related activity in superficial layers, as well as disinhibition in deeper layers, may underlie the deficits in working memory and other cognitive functions seen in these patients. We expect that future studies using layer-fMRI in

populations with or at risk for schizophrenia will shed new light on the spatiotemporal dynamics of cognitive dysfunction in this illness.

From a methodological perspective, here we used advanced contrast mechanisms and balanced task design to offset differences in vascular physiology across cortical depths, which can introduce substantial biases and limit the interpretability of layer fMRI [42]. In contrast to gradient-echo BOLD (GE-BOLD), CBV-weighted fMRI signal acquired with VASO allows appropriate separation of microvascular responses at a layer-dependent level [43,44]. We avoid biases of different hemodynamic response functions (HRFs) across cortical depths [45,46] by refraining from using statistical general linear model (GLM) deconvolution with predefined HRFs, and by restricting our interpreting to quantitative signal differences that are obtained at the same latency within identical task blocks. Additionally, we collected conventional GE-BOLD fMRI concomitantly with VASO data. The near-simultaneous acquisition of BOLD and VASO data allowed us to obtain a clean BOLD-corrected, CBV-weighed VASO signal. The higher sensitivity of BOLD compared to VASO was helpful in selecting the correct ROI, while the higher spatial specificity of VASO was helpful for interpreting signal across cortical depths.

These methodological advances have exciting implications for non-invasive, *in vivo* mapping of input-output and feedforward-feedback connections in the human neocortex. Outstanding challenges include expanding spatial coverage without sacrificing resolution, which would allow for functional connectivity analyses to infer information flow between far-flung cortical areas. For example, simultaneous imaging of dlPFC, premotor and primary motor cortices would help characterize inter-region interactions during response selection and execution, while expanding coverage to parietal and sensory areas as well as neighboring prefrontal areas would help characterize interactions that support stimulus perception, information storage and manipulation during the encoding and delay periods.

Looking beyond working memory, these tools provide a starting point for future work mapping layer-specific connections within high-order association cortex, and between high-order and unimodal cortex, in the context of cognitive neuroscience. Many influential theories of brain function that posit top-down and bottom-up signals with origins and destinations in distinct cortical layers—e.g., predictive coding and related frameworks—may now be directly tested in humans[47]. This opens the door to investigating computational mechanisms behind any number of neuropsychological phenomena, such as selective attention, hallucinations and delusions, and even consciousness itself, to name a few[48]. We expect that the ever-advancing tools of high-resolution fMRI will ultimately transform our understanding of cognition in the awake, behaving human brain.

METHODS

Please refer to the Life Sciences Reporting Summary, published alongside the online version of this paper, to access a subset of this information in a standardized format.

Subjects

Seventeen healthy volunteers participated after granting informed consent under an NIH Combined Neuroscience Institutional Review Board-approved protocol (93-M-0170, ClinicalTrials.gov identifier:)) in accordance with the Belmont Report and US federal regulations that protect human subjects. Data from two subjects were excluded due to technical difficulties or experimenter error: in one subject, no clear activation was visible within the field of view (meaning the region of interest was likely outside the field of view), and in the second subject, an incorrect version of the task was used, resulting in altered event timings that made this subject's data incompatible with the rest of the dataset. Of the remaining 15 subjects (age 20-47 years at the time of the experiment) whose data entered into the analyses presented here, eight were male and seven were non-pregnant females.

The functional data presented here come from a total 40 hours of scan time collected in 20 two-hour scan sessions. Two different functional acquisition protocols were used over the course of the study: an "axial [readout] protocol" (n = 15 sessions) and a "sagittal [readout] protocol" (n = 5 sessions); these are described further in their respectively titled sections below. Of the 15 unique subjects, n = 8 were scanned only once using the axial protocol; n = 3 were scanned once using the axial protocol and once using the sagittal protocol; n = 2 were scanned only once using the sagittal protocol, and n = 2 were scanned twice on the axial protocol. Some overlap of subjects was by design, allowing us to assess test-retest reliability of our ROI location (see Fig. S2). No statistical methods were used to pre-determine sample sizes, but our sample size is consistent with or larger than those reported in previous layer fMRI studies [22,21,43,44,19,20].

All fifteen subjects were invited for a separate scan session to obtain high-resolution reference anatomical T_1-weighted data with an MPRAGE-based sequence. Five additional two-hour scan sessions were used as pilot experiments to optimize the task design and investigate motion limitations and sequence artifacts; data from these sessions are not shown.

Task paradigm

The task was created using PsychoPy2 software [49]. For the axial readout protocol (TR = 2 s, described below), each trial consisted of the following epochs (example, duration): letter string presentation (BDCAE, 2.5 s), fixation cross (+, 1.5 s), instruction cue (ALPHABETIZE or REMEMBER, 1 s), delay period with fixation cross (+, 9 s), probe (D? or *?, 2 s), inter-trial interval with fixation cross (+, 16 s). Subjects could register a response at any time following the appearance of the probe and before the start of the next trial (i.e., anytime during the inter-trial interval). Each trial thus lasted 32 s, and each run consisted of 20 trials plus brief (8 s) additional fixations at the beginning and end of the run, for a total of 10:56 min:sec per run. Runs alternated between two contrast types: (1) manipulation versus maintenance (consisting of a mix of ALPHABETIZE and REMEMBER trials, all requiring action), and (2) action versus non-action (consisting of a mix of action and non-action trials, all ALPHABETIZE). Within each run, the 10 trials of each type were presented in a fixed pseudorandom order that was the same for all runs, to facilitate averaging.

For the higher-resolution sagittal readout protocol (described below), all runs were of the first contrast type (manipulation versus maintenance), and trial epoch timings were adjusted to match the longer TR (2.5 s) by scaling the duration of each epoch by a multiplier of 1.25. Each trial thus lasted 40 s, and the duration of these runs was 13:40 min:sec. All other parameters, including the pseudorandom order, were kept the same as above.

Prior to the start of the experimental runs, we ran a 6-minute functional localizer that was conducted at standard resolution and analyzed in real time, allowing us to functionally define a region of interest within left dlPFC in each individual subject while the subject was in the scanner. This localizer consisted entirely of ALPHABETIZE trials and slightly altered timing. The length of all trial epochs was as described above except the inter-trial interval, which was shortened to 5 s to create a 10-s on, 10-s off paradigm. Delay-related activity (including cue plus delay-related fixation) was considered signal, while all other trial epochs were treated as baseline. The location of peak activity from the real-time general linear model (GLM) analysis was used to position the coverage of the subsequent sub-millimeter experiments.

Randomization and blinding

There were no experimental groups in this study; therefore, no randomization of subjects was necessary. As stated in the “Task paradigm” section above, within each run, the 20 trials (10 of each type) were presented in a fixed pseudorandom order that was the same for all subjects and all runs. This was done to facilitate averaging within subjects and to ensure a relatively even distribution of each trial type across the beginning, middle and end of runs (to mitigate concerns about signal drift that might differentially affect one trial type or the other).

Data collection and analysis were not performed blind to the conditions of the experiments. Subjects were not told the purpose of the study or specific hypotheses concerning differences between trial types and within-trial periods ahead of time, but were debriefed following data collection upon request.

Experimental setup

All imaging was performed on a MAGNETOM 7T scanner (Siemens Healthineers, Erlangen, Germany) with a single-channel-transmit/32-channel-receive heal coil (Nova Medical, Wilmington, MA, USA). Imaging sessions did not exceed 120 minutes. Imaging slice position and slice angle were adjusted individually for every subject on the basis of the functional localizer described above.

A 3rd-order B_0-shim was done with three iterations using vendor-provided tools. The shim volume covered the entire imaging field of view (FOV) and was extended down to the circle of Willis in order to obtain sufficient B_0-homogeneity to exceed the adiabaticity threshold of the inversion pulse.

Following the functional localizer, for the axial protocol, run type alternated between the first contrast (alphabetize/remember) and the second contrast (action/non-action). All subjects completed at least five runs (3 of the alphabetize/remember contrast and 2 of the

action/non-action) per imaging session. Therefore there were 30 ‘alphabetize’, 30 ‘remember’, 20 ‘action’ and 20 ‘non-action’ trials per subject per session. (Note that ‘alphabetize’ and ‘action’ trials are technically identical, although data were not pooled between these two conditions for analysis purposes given that they were acquired in different runs.) When time allowed (for n = 6 subject-sessions), a sixth run was acquired (action/non-action contrast); these sessions thus comprised 30 of each trial type.

For the sagittal protocol, all runs were of the first contrast type (alphabetize/remember), and also consisted of 10 trials of each type (20 total), although note each trial was scaled to be longer in duration in order to match the TR of this protocol. Most subjects scanned using this protocol (n = 3) completed four total runs, or 80 total trials (40 ‘alphabetize and 40 ‘remember’). One subject completed three total runs (60 total trials/30 of each type) and one subject completed five runs (100 total trials/50 of each type).

Axial readout protocol

The protocol parameters were as follows: Readout type: 3D-EPI with one segment per k-space plane [26], in-plane resolution 0.91 × 0.91 mm^2, slice thickness 1.1 mm, FLASH GRAPPA 3, partial Fourier in the first phase encoding direction: 6/8, no partial Fourier in the second phase encoding direction, TR_{VASO} = 2000 ms, $TR_{VASO+BOLD}$ = 4000 ms, FOV read and phase = 150 mm, matrix size = 162, TE = 20 ms, read bandwidth = 1144 Hz/Px, phase echo spacing = 0.98. Assuming a GM T_2^* = 28 ms, the expected T_2^* blurring for EPI-readout results in a signal leakage of 12% from one voxel into the neighboring voxels along the first phase-encoding direction. A more detailed list of scan parameters used can be found on GitHub: https://github.com/layerfMRI/Sequence_Github/blob/master/DLPFC_Emily/Emily_Intermediate_protocol.pdf.

Sagittal readout protocol

The protocol parameters are as follows: Readout type: 3D-EPI with one segment per k-space plane [26], in-plane resolution 0.75 × 0.75 mm^2, slice thickness 0.99 mm, FLASH GRAPPA 3, partial Fourier in the first phase encoding direction: 6/8, no partial Fourier in the second phase encoding direction, TR_{VASO} = 2500 ms, $TR_{VASO+BOLD}$ = 5000 ms, FOV read = 130 mm, FOV phase 98.8%, matrix size = 172, TE = 27 ms, read bandwidth = 908 Hz/Px, phase echo spacing = 1.23 (limited by peripheral nerve stimulation thresholds). Assuming a GM T_2* = 28 ms, the expected T_2^* blurring for EPI-readout results in a signal leakage of 14% from one voxel into the neighboring voxels along the first phase-encoding direction. A more detailed list of scan parameters used can be found on GitHub: https://github.com/layerfMRI/Sequence_Github/blob/master/DLPFC_Emily/DLPFC_high_res_076_0.76_1.pdf.

VASO-specific protocol parameters

Both readout protocols were acquired with the same VASO preparation module. The protocol parameters were: Inversion pulse type: TR-FOCI pulse with a bandwidth of 6.4 kHz, μ = 7, pulse duration: 10 ms, non-selective. The phase skip of the adiabatic inversion pulse was adjusted to 30 deg to achieve an inversion efficiency of 80%, shorter than the arterial arrival time in the dlPFC [50]. The inversion time was adjusted to match the blood-nulling time of 1100 ms as done in previous studies [22]. To account for the T_1-decay during

the 3D-EPI readout and potential related blurring along the segment direction, a variable flip angle was chosen. The flip angle of the first segment was adjusted to be 22 deg. The subsequent flip angles where exponentially increasing, until last k-space segment was excited with a desired flip angle of 90 deg.

Image reconstruction

Image reconstruction was done in the vendor-provided platform as done previously [22]. GRAPPA 3 kernel fitting was done on FLASH ACS data, using a 3 × 4 kernel, 48 reference lines, and regularization parameter $\chi = 0.001$. RF-channels were combined with the sum-of-squares. To minimize resolutions losses in the phase-encoding direction due to T_2^*-decay partial, Fourier reconstruction was done with POCS using 8 iterations.

Anatomical reference data

In separate scan sessions, 0.7 mm resolution T_1-maps where collected covering the entire brain with an MP2RAGE sequence [51] for every subject. These data were not used in the functional pipelines to analyze the layer-dependent activity changes. Instead, these images were used to investigate the reproducibility of location of activity across sessions (Figure S2) and across subjects (Figure 1c).

In four of the subjects that were invited for more than two 2-hour sessions, slab-selective isotropic 0.5 mm and 0.4 mm resolution anatomical data were collected with MP2RAGE and Multi-Echo FLASH, respectively. Those anatomical data were not used in the pipeline for generating cortical profiles. They are used to compare and validate the approximate position of the cyto-architectonically defined cortical layers of individual subjects to the 20 reconstructed cortical depths, in which the functional data are processed (Fig. S4).

Functional image preprocessing

This section describes processing steps that were common to both the axial and sagittal protocols. For a schematic overview of the analysis pipeline, see Fig. S1.

First, DICOM images were converted to NIFTI using the ISISCONV converter (Fig. S1a). Motion correction was performed using SPM software (Statistical Parametric Mapping; SPM12) [52] and was done separately for nulled and not-nulled frames (Fig. S1b). A 4th order spline function was used for spatial interpolation. Motion correction and registration across runs was done simultaneously. This minimized the effect of spatial resolution loss to one single resampling step [53]. Motion traces of nulled and not-nulled were visually inspected to ensure good overlap for the two contrasts (Fig. S1b).

Following these steps, frames were sorted into their respective contrast: not-nulled (BOLD) or nulled (VASO; Fig. S1c). Note that BOLD and VASO contrasts are kept separate from this point forward, and all analyses below were performed for each contrast individually.

Next, runs of the same contrast type were averaged (Fig. S1d), and within these average runs, trials of the same type were averaged (Fig. S1e). Because all runs have the same trial order, and all trials have the same epoch structure and timing, runs and trials can be averaged without deconvolving the hemodynamic response. This is an important feature of our

experimental design, since hemodynamic responses differ across cortical depths [46]. Following trial averaging, VASO data were BOLD corrected using the dynamic division method (Fig. S1e). Thus, for each contrast (BOLD and VASO), for the axial protocol, each subject had four average trials: alphabetize, remember, action, and non-action. For the sagittal protocol, each subject had two average trials: alphabetize and remember.

In a parallel analysis, a region of interest (ROI) in the left dlPFC was defined for each subject (Fig. S1f, right). The approximate location of the ROI was taken from the 6-minute functional localizer (Fig. S1f, left) following GLM analysis with FSL FEAT (Version 5.98) [54] . For the complete FEAT design protocol, please see (https://github.com/layerfMRI/repository/tree/master/DLPFC_Emily/Featdesign). The ROI was manually selected and drawn for every individual subject (see Fig. S3 for drawn ROIs in every subject). Rather than only acquire an additional T1-weighted image for anatomical reference, we used the functional EPI data itself to estimate the T1 contrast, and used this for manual delineation of two layers within this ROI, one superficial and one deep (Fig. S1f, right). The advantage of this approach is that it avoids the distortion correction and resampling steps necessary for registering EPI images to a separately acquired T1 image, preserving spatial specificity. See sections below for additional information about this layer-drawing procedure for both the axial and sagittal protocols.

Layering and timecourse extraction for axial protocol

This section describes the steps applied to data acquired using the axial protocol and shown in Figs. 2 and 3. The manual drawing of the layer masks was done according to the following guidelines: a) layers were drawn as connected collection of voxels without holes; b) the superficial layer was positioned such that there was no partial voluming with CSF; c) the deeper layer was positioned such that there was no partial voluming with WM; d) the superficial and deeper layers were eroded until there was no residual overlap of superficial and deeper layers; e.) the thickness of the superficial and deeper layers were kept similar along the cortical ribbon; f) the thickness of the superficial and deeper layers was chosen such that they fill as much of the cortex as possible without violating the guidelines above; and g) for consistency, the same person drew the layers for all subjects. The results of all drawings are shown in Fig. S3.

Next, at each timepoint, signal was averaged across all voxels within each layer to derive one average timecourse per layer in each of the four trial types. Thus, each subject had eight timecourses: one per layer (superficial, deeper) per trial type (alphabetize, remember, action, non-action; Fig. S2g).

Before pooling data across subjects, BOLD timecourses were normalized within subjects using the following steps. First, a per-layer (y) mean baseline BOLD signal ($\bar{b}_y$) was calculated by averaging signal during baseline timepoints across all four trial types (where “baseline timepoints” include the first timepoint, which is before the appearance of the stimulus, and the penultimate and ultimate timepoints, which are 18 and 22 seconds after the appearance of the probe, the point at which signal is expected to have returned to at or near

baseline). Next, the BOLD signal s for layer y at timepoint t was transformed to s' as follows, to yield values interpretable as percent signal change:

$$S_{y,t'} = \left(S_{y,t} / \bar{b}_y\right) * 100 - 100$$

Note that unlike BOLD, VASO is a quantitative measure that is proportional to a physical unit (mL per 100 mL tissue volume), meaning units can be directly interpreted and it is not necessary to convert to percent signal change. VASO data were instead transformed as follows. First, to facilitate interpretation, each subject's VASO signal v at each timepoint t was transformed from a negative to a positive contrast as:

$$v_t = v_t^* - 100$$

Following this, VASO signals were normalized within subjects by calculating a per-layer mean baseline VASO signal ($\bar{v}_y$) by averaging signal during baseline timepoints (same timepoints as for BOLD above) across all four trial types. This mean baseline signal was subtracted from each timepoint as follows:

$$v_{y,t'} = v_{y,t} - \bar{v}_y$$

All of the subsequent statistical contrasts were performed directly on these normalized signal timecourses. We refrained from using deconvolution or inferential statistical models (e.g., general linear models) to measure activation, to avoid biases of variable noise magnitudes and hemodynamical response functions across cortical depths.

For purposes of the two-way, repeated-measures analyses of variance (ANOVAs) depicted in Figs. 2b, 3b and S5b, the representative delay signal was the average of VASO measurements acquired at timepoints 4, 5 and 6 (corresponding to 12, 16 and 20 sec in trial time), and the representative response signal was the average of VASO measurements acquired at timepoints 7 and 8 (corresponding to 24 and 28 sec in trial time). While the repeated-measures ANOVA test is robust against violations of the assumption of normality, it does assume sphericity, which refers to the condition where the variances of the differences between all possible pairs of within-subject conditions (i.e., levels of the independent variable) are equal. Because there is currently no clear way to test for sphericity for the interaction term of a two-way repeated measures ANOVA (our main term of interest), here, we report the Greenhouse-Geisser-corrected p-value[55] for all tests, which is a conservative form of correction that is recommended when nothing is known about the sphericity of the data[56].

Layering for sagittal protocol

This section describes image processing for the single-subject layer-dependent activity profiles acquired using the sagittal protocol and shown in Fig. 4. Cortical depths were estimated directly in EPI space without alignment to so-called anatomical space. This minimizes the risk of resolution loss due to multiple spatial resampling steps and avoids any

potential errors in distortion correction and registration. An anatomical reference contrast was calculated from the functional data by calculating the inverse signal variability across nulled and not-nulled images, divided by the mean signal. This measure is called here T_1-EPI and provides a good contrast between white matter (WM), gray matter (GM) and cerebro-spinal fluid (CSF; see background images in Figure 4, S1 and S3). Borderlines between GM/WM and GM/CSF are manually drawn based on this contrast. The manual drawing was done as described in previous publications [57,22,58,59] according the following guidelines: a) borderlines were drawn as continuous lines without holes; b) the lines are drawn such that their curvature radius was kept smaller than the cortical thickness; c) the position of the GM/CSF border was drawn through voxels that were just above the GM, such that there was no GM partial voluming; d) the position of the GM/WM border was drawn through voxels that were just below the GM, such that there was no GM partial voluming—this means that the position of the voxels that are half filled with GM are in the respective upper-most and lower-most extracted layers. e) for consistency, the same person drew the layers for all subjects.

Manually drawn border lines are shown for all subjects in Figure 4c (bright yellow for GM/CSF and bright blue for GM/WM). Twenty-one layers were calculated between these borderlines with the LAYNII program LN_GROW_LAYERS (https://github.com/layerfMRI/LAYNII). In order to minimize partial volume effects and allow the calculation of smooth layers, the layering calculation was applied on a four-fold finer grid that the native functional resolution. This means that the number of layers is higher than the number of independent voxels sampled across the cortical depth. The number of layers should not be confused with the effective resolution across cortical depths. Given the cortical thickness of 3.5-4 mm in dlPFC [60,61], the resolution of 0.76 mm in-plane and 0.99 mm slice thickness is sufficient to sample 3-6 independent voxels across cortical depth. This is enough to estimate activity in superficial and deeper layers (red-yellow compared to blue-turquoise in Fig. 4c) with Nyquist sampling. The number 21 was chosen based on previous experience in finding a compromise between data size and smoothness (see Fig. S8 in [58] as well as [22,57]).

For best visibility, functional signals were smoothed along the tangential direction of the cortex (i.e., within “layers”) with a Gaussian kernel of 0.76 mm. In order to maintain the spatial specificity across layers, no smoothing was applied across cortical depths. This kind of layer smoothing can improve the detectability of fMRI signal changes without unwanted leakage of physiological noise above the cortical surface [62,22,58]. The application of such layer smoothing is based on the assumption that neighboring columnar structures are similarly engaged during the task. See Fig. S8 for a discussion of variance in the functional response across columns. Note that the batch of cortex investigated here is highly folded with respect to the external magnetic field. This means that the BOLD signal change can be substantially variable dependent on the columnar position along the sulcus[63,64].

Interpreting cortical depth-dependent results with respect to cytoarchitectonic layers

In order to interpret the fMRI results according to known input-output characteristics of different cortical layer groups II/III and V/VI, it is helpful to approximate the location of functional activity with respect to underlying layers as defined cytoarchitectonically. To

confirm the approximate borders and the different layers within these borders, we followed the approach outlined in earlier work[65,66]. This is a three step approach: First, we extracted layer signatures in high-resolution multi-modal post-mortem histology data of an individual cadaver brain sample from the Ding Atlas[67]. Second, we identified the MR-sensitive features and landmarks[68] in anatomical MRI scans from a subject from our study and estimated their relative position across the cortical thickness. Third, we used these features as markers of the cytoarchitectonic layers in the functional data from the same participant to confirm the relative depth-position of the functional responses. With this procedure, we can attempt to interpret the layer-origin of the functional signal solely based on the relative depth of the cortical thickness. The results of this procedure are shown in Fig. S4.

Note that this approach of comparing fMRI data with histology data is not conducted as part of the fMRI analysis pipeline. The time courses and layer profiles shown here are solely extracted based on relative distance to the GM/CSF and GM/WM borderlines. The comparison of the relative cortical depth in fMRI data and histology data is based on the assumption that the relative position of the cyto-architectonic layers and their relative thicknesses is the same across subjects (see insets in Fig. 4C).

Spatial alignment across sessions (within-subject)

Note that all layer data are taken from individual sessions, and are thus not susceptible to potential registration errors across days. However, it is important to ensure that the location of activity is generally consistent with a single subject across days and imaging sessions.

To investigate this consistency in the two subjects on whom we collected test-retest data (i.e., two imaging sessions separately by several days), each session's layer masks and the corresponding activation maps were transformed into subject-specific anatomical reference spaces. Registration was done with SyN in ANTs (Advanced Normalization Tools; [69]) with a spline interpolation. Since the imaging coverage of the functional data is significantly smaller than the whole brain, it was necessary to provide a manual starting point for the ANTs registration to converge on reasonable registration quality. The initial manual registration was done in ITK-SNAP. The registration from EPI-space to the subject specific anatomical space was done by means of the similar T_1 contrast of T_1-EPI and the MP2RAGE UNI-DEN image. The same spatial operation was applied to the layer masks and the functional activation maps. The resulting activation patterns where compared across days in the anatomical space of individual subjects (Fig. S2). Note that the registration quality here did not need to achieve accuracy levels at the sub-millimeter layer scale. Instead, the goal of this analysis was to demonstrate that the process of ROI selection (several millimeters large) was reproducible.

Spatial alignment across subjects (mean ROI location)

To verify placement of the ROI taken from the functional localizer, and to create the group-level image shown in Fig. 1c, we processed data from the localizer run in AFNI[70], using the standard "super-script" afni_proc.py. Each subject's high-resolution (T_1-MPRAGE) whole-brain anatomical data were registered to the MNI 152 template using a combined affine and nonlinear warp. To minimize interpolation, this transformation was concatenated with both

the affine transform used to register the echo-planar images to the individual-subject anatomical data, as well as the rigid (6 degrees of freedom) warp to account for subject motion. Data were then smoothed using a 4mm (2 voxels) Gaussian kernel, scaled to percent signal change, and submitted to a multiple regression. The standard boxcar block design was convolved with the HRF along with six motion parameters (3 translation, 3 rotation). Group analyses were conducted in 3dttest++, which yielded a cluster in left dlPFC with a whole-brain map at voxelwise $p < 0.01$. This cluster represents the approximate location where the higher-resolution layer slices were prescribed in the subsequent experimental runs, and is included here for convenience as a post-hoc visualization of the macroscale location of our region of interest.

DATA AVAILABILITY

Data are available via OpenNeuro at the following link: https://doi.org/10.18112/openneuro.ds002076.v1.0.1

CODE AVAILABILITY

All code is available in the following GitHub repository: https://github.com/layerfMRI/repository/tree/master/DLPFC_Emily

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

We thank A. Arnsten for guidance on experimental design and interpretation. We thank B. Poser and D. Ivanov for the 3D-EPI readout that is used in the VASO sequence used here. We thank A.H. Hall and K. Chung for administrative support of human volunteer scanning. We thank S. Kashyap for helpful tips on adjusting manual initial registration used to generate Fig. S2. We thank S. Marrett and D. Handwerker for technical advice and support, and four anonymous reviewers for helpful comments. We thank A. Thomas and T. Riddle for support with data formatting, organization, and sharing. Portions of this study used the high-performance computational capabilities of the Biowulf Linux cluster at the National Institutes of Health, Bethesda, MD (biowulf.nih.gov).

Funding: The research was funded by the National Institute of Mental Health Intramural Research Program (ZIAMH002783) to P.A.B., which also funded authors E.S.F., L.H., D.C.J. and P.J.M. During the latest periods of data analysis, author L.H. was funded from the NWO VENI project 016.Veni.198.032.

REFERENCES

1. Courtney SM, Ungerleider LG, Keil K & Haxby JV Transient and Sustained Activity in a Distributed Neural System for Human Working Memory. Nature 386, 608, (1997). [PubMed: 9121584]
2. Courtney SM, Petit L, Maisog JM, Ungerleider LG & Haxby JV An Area Specialized for Spatial Working Memory in Human Frontal Cortex. Science 279, 1347–1351, (1998). [PubMed: 9478894]
3. Goldman-Rakic P Cellular Basis of Working Memory. Neuron 14, 477–485, (1995). [PubMed: 7695894]
4. D'esposito M et al. The Neural Basis of the Central Executive System of Working Memory. Nature 378, 279, (1995). [PubMed: 7477346]
5. Sawaguchi T, Matsumura M & Kubota K Depth Distribution of Neuronal Activity Related to a Visual Reaction Time Task in the Monkey Prefrontal Cortex. J. Neurophysiol 61, 435–446, (1989). [PubMed: 2918365]

6. Sawaguchi T, Matsumura M & Kubota K Catecholaminergic Effects on Neuronal Activity Related to a Delayed Response Task in Monkey Prefrontal Cortex. J. Neurophysiol 63, 1385–1400, (1990). [PubMed: 2358882]
7. Arnsten AF, Wang MJ & Paspalas CD Neuromodulation of Thought: Flexibilities and Vulnerabilities in Prefrontal Cortical Network Synapses. Neuron 76, 223–239, (2012). [PubMed: 23040817]
8. Markowitz DA, Curtis CE & Pesaran B Multiple Component Networks Support Working Memory in Prefrontal Cortex. Proc. Natl. Acad. Sci. USA 112, 11084–11089, (2015). [PubMed: 26283366]
9. Bastos AM, Loonis R, Kornblith S, Lundqvist M & Miller EK Laminar Recordings in Frontal Cortex Suggest Distinct Layers for Maintenance and Control of Working Memory. Proc. Natl. Acad. Sci. USA, 201710323, (2018).
10. Barbey AK, Koenigs M & Grafman J Dorsolateral Prefrontal Contributions to Human Working Memory. Cortex 49, 1195–1205, (2013). [PubMed: 22789779]
11. Mackey WE, Devinsky O, Doyle WK, Meager MR & Curtis CE Human Dorsolateral Prefrontal Cortex Is Not Necessary for Spatial Working Memory. J. Neurosci 36, 2847–2856, (2016). [PubMed: 26961941]
12. Postle BR et al. Repetitive Transcranial Magnetic Stimulation Dissociates Working Memory Manipulation from Retention Functions in the Prefrontal, but Not Posterior Parietal, Cortex. J. Cogn. Neurosci 18, 1712–1722, (2006). [PubMed: 17014375]
13. Garey LJ et al. Reduced Dendritic Spine Density on Cerebral Cortical Pyramidal Neurons in Schizophrenia. Journal of Neurology, Neurosurgery & Psychiatry 65, 446–453, (1998).
14. Glantz LA & Lewis DA Decreased Dendritic Spine Density on Prefrontal Cortical Pyramidal Neurons in Schizophrenia. Arch. Gen. Psychiatry 57, 65–73, (2000). [PubMed: 10632234]
15. Cannon TD et al. Dorsolateral Prefrontal Cortex Activity During Maintenance and Manipulation of Information in Working Memory in Patients with Schizophrenia. Arch. Gen. Psychiatry 62, 1071–1080, (2005). [PubMed: 16203952]
16. Wang M, Vijayraghavan S & Goldman-Rakic PS Selective D2 Receptor Actions on the Functional Circuitry of Working Memory. Science 303, 853–856, (2004). [PubMed: 14764884]
17. Opris I, Hampson RE, Stanford TR, Gerhardt GA & Deadwyler SA Neural Activity in Frontal Cortical Cell Layers: Evidence for Columnar Sensorimotor Processing. J. Cogn. Neurosci 23, 1507–1521, (2011). [PubMed: 20695762]
18. Polimeni JR, Fischl B, Greve DN & Wald LL Laminar Analysis of 7 T Bold Using an Imposed Spatial Activation Pattern in Human V1. Neuroimage 52, 1334–1346, (2010). [PubMed: 20460157]
19. Kok P, Bains LJ, van Mourik T, Norris DG & de Lange FP Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. Current Biology 26, 371–376, (2016). [PubMed: 26832438]
20. Muckli L et al. Contextual Feedback to Superficial Layers of V1. Current Biology 25, 2690–2695, (2015). [PubMed: 26441356]
21. De Martino F et al. Frequency Preference and Attention Effects across Cortical Depths in the Human Primary Auditory Cortex. Proc. Natl. Acad. Sci. USA 112, 16036–16041, (2015). [PubMed: 26668397]
22. Huber L et al. High-Resolution Cbv-Fmri Allows Mapping of Laminar Activity and Connectivity of Cortical Input and Output in Human M1. Neuron 96, 1253–1263.e1257, (2017). [PubMed: 29224727]
23. Corbitt PT, Ulloa A & Horwitz B Simulating Laminar Neuroimaging Data for a Visual Delayed Match-to-Sample Task. NeuroImage 173, 199–222, (2018). [PubMed: 29476912]
24. D'Esposito M, Postle BR, Ballard D & Lease J Maintenance Versus Manipulation of Information Held in Working Memory: An Event-Related Fmri Study. Brain and cognition 41, 66–86, (1999). [PubMed: 10536086]
25. Lu H, Golay X, Pekar JJ & Van Zijl PC Functional Magnetic Resonance Imaging Based on Changes in Vascular Space Occupancy. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine 50, 263–274, (2003).

26. Poser BA, Koopmans PJ, Witzel T, Wald LL & Barth M Three Dimensional Echo-Planar Imaging at 7 Tesla. Neuroimage 51, 261–266, (2010). [PubMed: 20139009]
27. Petrides M & Pandya D Dorsolateral Prefrontal Cortex: Comparative Cytoarchitectonic Analysis in the Human and the Macaque Brain and Corticocortical Connection Patterns. Eur. J. Neurosci 11, 1011–1036, (1999). [PubMed: 10103094]
28. Melchitzky DS, Sesack SR, Pucak ML & Lewis DA Synaptic Targets of Pyramidal Neurons Providing Intrinsic Horizontal Connections in Monkey Prefrontal Cortex. Journal of Comparative Neurology 390, 211–224, (1998). [PubMed: 9453665]
29. Medalla M & Barbas H Diversity of Laminar Connections Linking Periarcuate and Lateral Intraparietal Areas Depends on Cortical Structure. European Journal of Neuroscience 23, 161–179, (2006). [PubMed: 16420426]
30. Wang X-J Synaptic Basis of Cortical Persistent Activity: The Importance of Nmda Receptors to Working Memory. J. Neurosci 19, 9587–9603, (1999). [PubMed: 10531461]
31. Wang M et al. Nmda Receptors Subserve Persistent Neuronal Firing During Working Memory in Dorsolateral Prefrontal Cortex. Neuron 77, 736–749, (2013). [PubMed: 23439125]
32. Bunge SA, Hazeltine E, Scanlon MD, Rosen AC & Gabrieli J Dissociable Contributions of Prefrontal and Parietal Cortices to Response Selection. Neuroimage 17, 1562–1571, (2002). [PubMed: 12414294]
33. Jiang Y & Kanwisher N Common Neural Substrates for Response Selection across Modalities and Mapping Paradigms. J. Cogn. Neurosci 15, 1080–1094, (2003). [PubMed: 14709228]
34. Curtis CE & D'Esposito M Persistent Activity in the Prefrontal Cortex During Working Memory. Trends in cognitive sciences 7, 415–423, (2003). [PubMed: 12963473]
35. Hadland KA, Rushworth MFS, Passingham RE, Jahanshahi M & Rothwell JC Interference with Performance of a Response Selection Task That Has No Working Memory Component: An Rtms Comparison of the Dorsolateral Prefrontal and Medial Frontal Cortex. J. Cogn. Neurosci 13, 1097–1108, (2001). [PubMed: 11784448]
36. Schumacher EH & D'Esposito M Neural Implementation of Response Selection in Humans as Revealed by Localized Effects of Stimulus–Response Compatibility on Brain Activation. Hum. Brain Mapp. 17, 193–201, (2002). [PubMed: 12391572]
37. Takada M et al. Organization of Prefrontal Outflow toward Frontal Motor-Related Areas in Macaque Monkeys. Eur. J. Neurosci 19, 3328–3342, (2004). [PubMed: 15217388]
38. Tomio A, Kyoko W & Kisou K Connections of Area 8 with Area 6 in the Brain of the Macaque Monkey. Journal of Comparative Neurology 277, 21–40, (1988). [PubMed: 2461971]
39. Arikuni T & Kubota K The Organization of Prefrontocaudate Projections and Their Laminar Origin in the Macaque Monkey: A Retrograde Study Using Hrp-Gel. Journal of Comparative Neurology 244, 492–510, (1986). [PubMed: 2420836]
40. Yeterian E & Pandya D Laminar Origin of Striatal and Thalamic Projections of the Prefrontal Cortex in Rhesus Monkeys. Experimental Brain Research 99, 383–398, (1994). [PubMed: 7957718]
41. Perlstein WM, Carter CS, Noll DC & Cohen JD Relation of Prefrontal Cortex Dysfunction to Working Memory and Symptoms in Schizophrenia. Am. J. Psychiatry 158, 1105–1113, (2001). [PubMed: 11431233]
42. Kay K et al. A Critical Assessment of Data Quality and Venous Effects in Ultra-High-Resolution Fmri. bioRxiv, 337667, (2018).
43. Goense J, Merkle H & Logothetis NK High-Resolution Fmri Reveals Laminar Differences in Neurovascular Coupling between Positive and Negative Bold Responses. Neuron 76, 629–639, (2012). [PubMed: 23141073]
44. Kim T & Kim S-G Cortical Layer-Dependent Arterial Blood Volume Changes: Improved Spatial Specificity Relative to Bold Fmri. Neuroimage 49, 1340–1349, (2010). [PubMed: 19800013]
45. Yacoub E, Ugurbil K & Harel N The Spatial Dependence of the Poststimulus Undershoot as Revealed by High-Resolution Bold-and Cbv-Weighted Fmri. Journal of Cerebral Blood Flow & Metabolism 26, 634–644, (2006). [PubMed: 16222242]
46. Petridou N & Siero JC Laminar Fmri: What Can the Time Domain Tell Us? NeuroImage, (2017).

47. Stephan KE et al. Laminar Fmri and Computational Theories of Brain Function. Neuroimage, (2017).
48. Lawrence SJD, Formisano E, Muckli L & de Lange FP Laminar Fmri: Applications for Cognitive Neuroscience. Neuroimage 197, 785–791, (2019). [PubMed: 28687519]

METHODS-ONLY REFERENCES

49. Peirce JW Psychopy—Psychophysics Software in Python. Journal of neuroscience methods 162, 8–13, (2007). [PubMed: 17254636]
50. Mildner T et al. Mapping of Arterial Transit Time by Intravascular Signal Selection. NMR in Biomedicine 27, 594–609, (2014). [PubMed: 24610794]
51. Marques JP et al. Mp2rage, a Self Bias-Field Corrected Sequence for Improved Segmentation and Tl-Mapping at High Field. Neuroimage 49, 1271–1281, (2010). [PubMed: 19819338]
52. Friston KJ et al. Statistical Parametric Maps in Functional Imaging: A General Linear Approach. Human brain mapping 2, 189–210, (1994).
53. Polimeni JR, Renvall V, Zaretskaya N & Fischl B Analysis Strategies for High-Resolution Uhf-Fmri Data. NeuroImage 168, 296–320, (2018). [PubMed: 28461062]
54. Worsley K Statistical Analysis of Activation Images. FunctionalMRI: An introduction to methods 14, 251–270, (2001).
55. Greenhouse SW & Geisser S On Methods in the Analysis of Profile Data. Psychometrika 24, 95–112, (1959).
56. Girden ER Anova: Repeated Measures. (Sage, 1992).
57. Huber L, Uludag K & Moller HE Non-Bold Contrast for Laminar Fmri in Humans: Cbf, Cbv, and Cmro2. Neuroimage, (2017).
58. Huber L et al. Ultra-High Resolution Blood Volume Fmri and Bold Fmri in Humans at 9.4 T: Capabilities and Challenges. Neuroimage 178, 769–779, (2018). [PubMed: 29890330]
59. Huber L et al. Techniques for Blood Volume Fmri with Vaso: From Low-Resolution Mapping Towards Sub-Millimeter Layer-Dependent Applications. Neuroimage 164, 131–143, (2018). [PubMed: 27867088]
60. Fischl B & Dale AM Measuring the Thickness of the Human Cerebral Cortex from Magnetic Resonance Images. Proceedings of the National Academy of Sciences 97, 11050–11055, (2000).
61. Williams SM & Goldman-Rakic PS Characterization of the Dopaminergic Innervation of the Primate Frontal Cortex Using a Dopamine-Specific Antibody. Cerebral cortex 3, 199–222, (1993). [PubMed: 8100725]
62. Blazejewska AI, Fischl B, Wald LL & Polimeni JR Intracortical Smoothing of Small-Voxel Fmri Data Can Provide Increased Detection Power without Spatial Resolution Losses Compared to Conventional Large-Voxel Fmri Data. Neuroimage 189, 601–614, (2019). [PubMed: 30690157]
63. Gagnon L et al. Quantifying the Microvascular Origin of Bold-Fmri from First Principles with Two-Photon Microscopy and an Oxygen-Sensitive Nanoprobe. J. Neurosci 35, 3663–3675, (2015). [PubMed: 25716864]
64. Fracasso A, Luijten PR, Dumoulin SO & Petridou N Laminar Imaging of Positive and Negative Bold in Human Visual Cortex at 7t. Neuroimage 164, 100–111, (2018). [PubMed: 28213112]
65. Geyer S, Weiss M, Reimann K, Lohmann G & Turner R Microstructural Parcellation of the Human Cerebral Cortex–from Brodmann's Post-Mortem Map to in Vivo Mapping with High-Field Magnetic Resonance Imaging. Front. Hum. Neurosci 5, 19, (2011). [PubMed: 21373360]
66. Turner R in Microstructural Parcellation of the Human Cerebral Cortex 197–220 (Springer, 2013).
67. Ding S-L et al. Comprehensive Cellular-Resolution Atlas of the Adult Human Brain. J. Comp. Neurol 524, 3127–3481, (2016). [PubMed: 27418273]
68. Stüber C et al. Myelin and Iron Concentration in the Human Brain: A Quantitative Study of Mri Contrast. Neuroimage 93, 95–106, (2014). [PubMed: 24607447]
69. Avants BB, Epstein CL, Grossman M & Gee JC Symmetric Diffeomorphic Image Registration with Cross-Correlation: Evaluating Automated Labeling of Elderly and Neurodegenerative Brain. Medical Image Analysis 12, 26–41, (2008). [PubMed: 17659998]

70. Cox RW Afni: Software for Analysis and Visualization of Functional Magnetic Resonance Neuroimages. Comput. Biomed. Res 29, 162–173, (1996) [PubMed: 8812068]

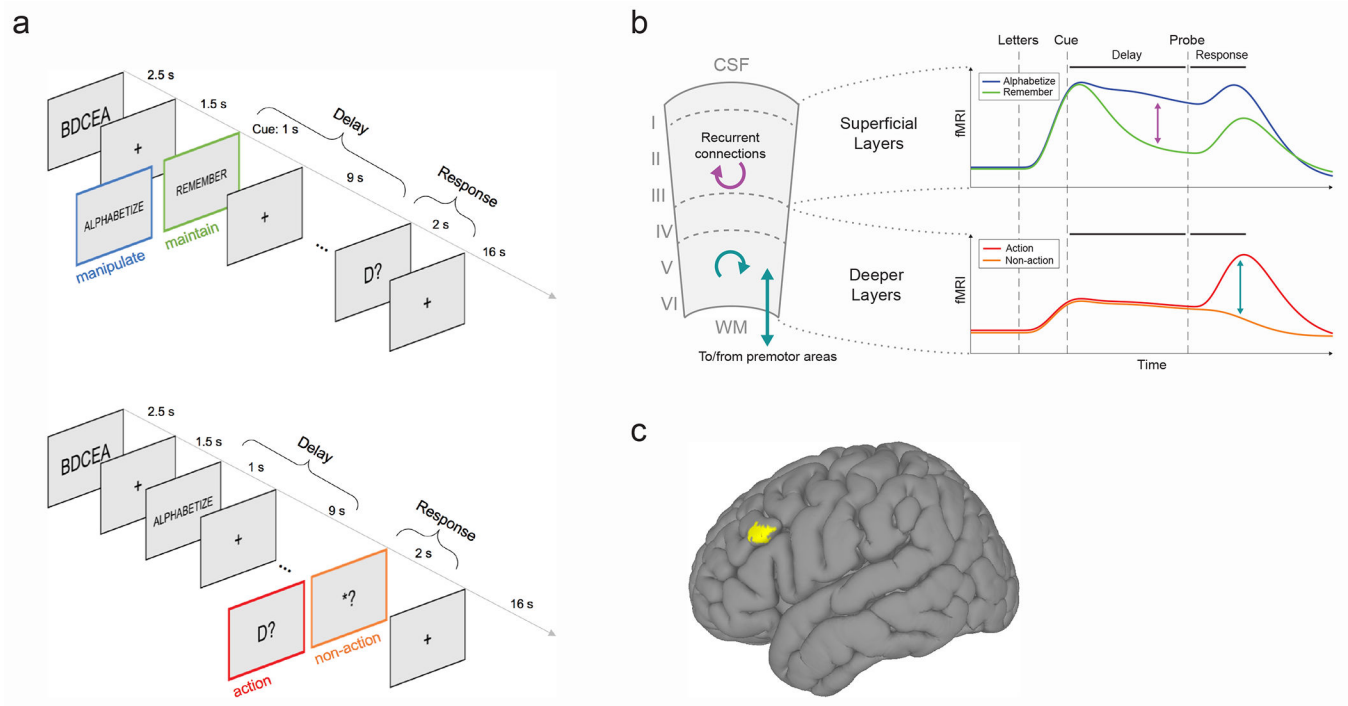

Fig. 1. Task, hypothesis and region of interest.

(A) Trial structure. Top panel: first contrast type, contrasting manipulation (‘alphabetize’) versus maintenance (‘remember’) during the delay period. Subjects see a string of five random letters (e.g., ‘BDCEA’), then a cue instructing them to either rearrange the letters in alphabetical order (‘ALPABETIZE’, manipulation condition) or to simply remember them in their original order (‘REMEMBER’, maintenance condition) over the course of a delay period, during which they see only a fixation cross. Finally, a probe letter comes onscreen (e.g, ‘D?’), and subjects make a response to indicate the alphabetical or ordinal position of the probed letter. Bottom panel: second contrast type contrasting action versus non-action during the response period. These trials are identical to the first until the response period, at which point subjects see either a true probe requiring a button press (e.g., ‘D?’, action condition), or a dummy probe (i.e., ‘*?’, non-action condition), which indicates that no response is required and they can forget the information associated with that trial. Colored frames are for schematic purposes only and were not seen by subjects. (B) Schematic of hypothesis. We hypothesized that (i) in superficial layers, manipulation trials would evoke more activity than maintenance trials specifically during the delay period due to recurrent excitation in layer III (purple arrows), and (ii) in deeper layers, action trials would evoke more activity than non-action trials due to action-selection and/or motor-related functions in layer V (teal arrows). WM, white matter; CSF, cerebro-spinal fluid. (C) Macroscale location of left dlPFC region of interest (MNI coordinates for center of mass: [x = +49, y = −21, z = +23], computed via group analysis of whole-brain functional localizer data (resulting in cluster displayed at voxelwise $p < 0.01$). For single-subject layer ROIs, see Fig. S3.

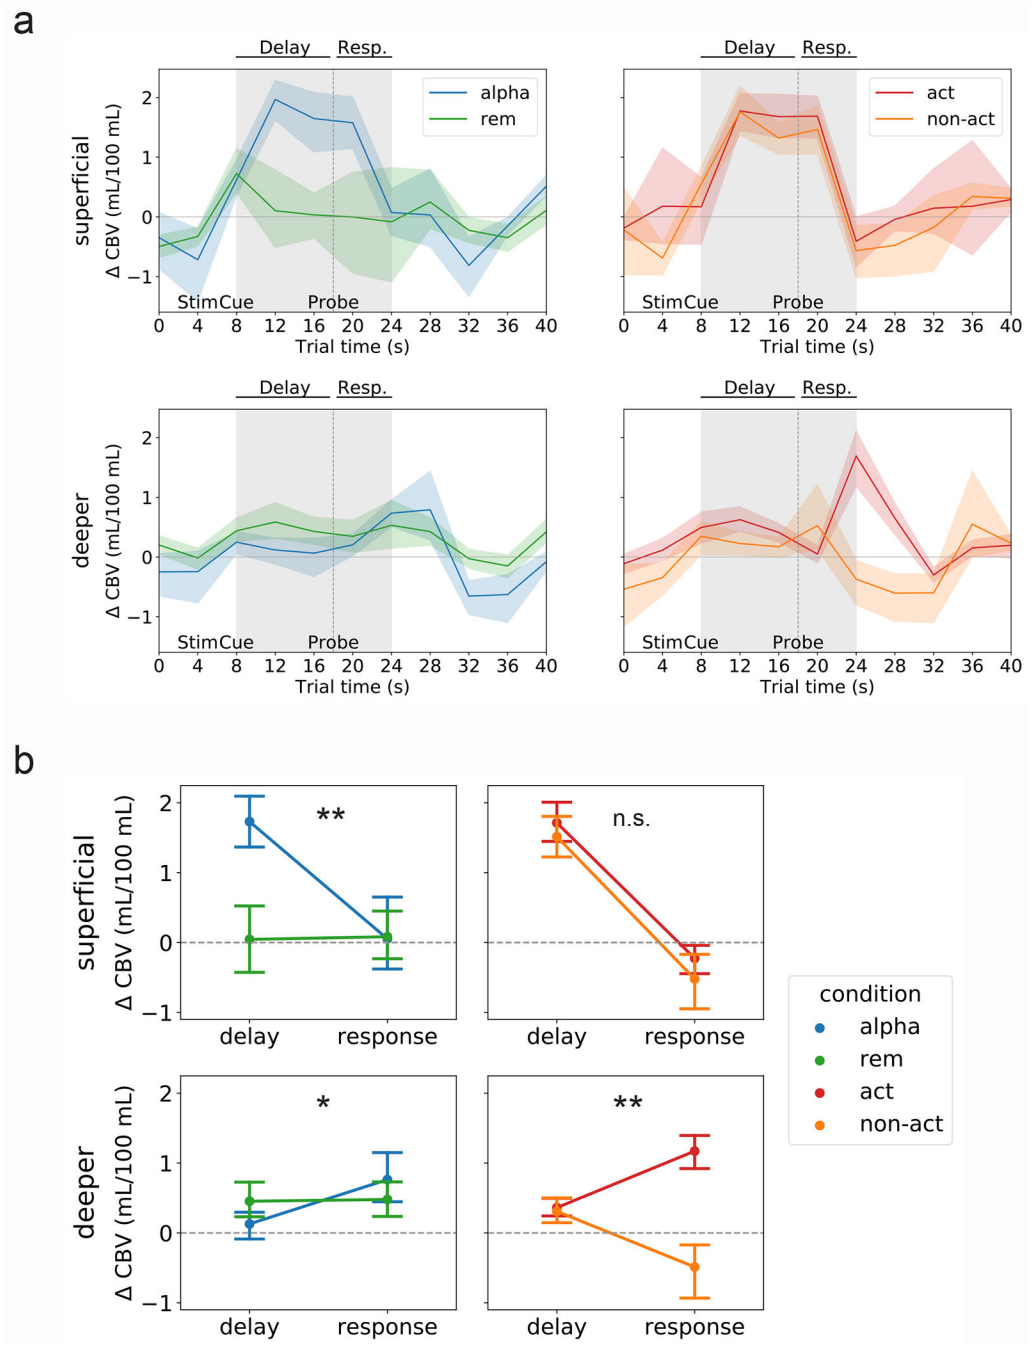

Fig. 2. Different trial types evoke distinct spatiotemporal patterns of activity.

(A) Left panel: mean VASO signal change (in units of mL/100 mL cerebral blood volume [CBV]) in superficial layers (top) and deeper layers (bottom) for the first contrast, manipulation trials (‘alpha’) versus maintenance trials (‘rem’). Right panel: mean VASO signal change in superficial layers (top) and deeper layers (bottom) for the second contrast, action trials (‘act’) versus non-action trials (‘non-act’). Lines represent mean and shaded area represents 95 percent confidence intervals for the mean (determined via bootstrapping with 1,000 iterations) across n = 15 sessions (13 unique subjects). See Figs. S6 and S7 for

single-subject timecourses, and Fig. S5a for mean BOLD timecourses. (B) Two-way analyses of variance (ANOVAs) with factors trial period (delay versus response) and trial type (either manipulation [‘alpha’] versus maintenance [‘rem’], or action versus non-action) in superficial (top) and deeper (bottom) layers. Panels as in (A). Dots represent mean and error bars reflect 95 percent confidence intervals for the mean. **, interaction significant at $p < 0.001$ ($p = 7.7e^{-5}$ and $p = 0.002$ for superficial alphabetize-versus-remember contrast and deeper action-versus-non-action contrast, respectively); *, interaction significant at $p < 0.01$ ($p = 0.004$ for deeper alphabetize-versus-remember contrast); n.s., interaction not significant ($p = 0.68$ for superficial action-versus-non-action contrast).

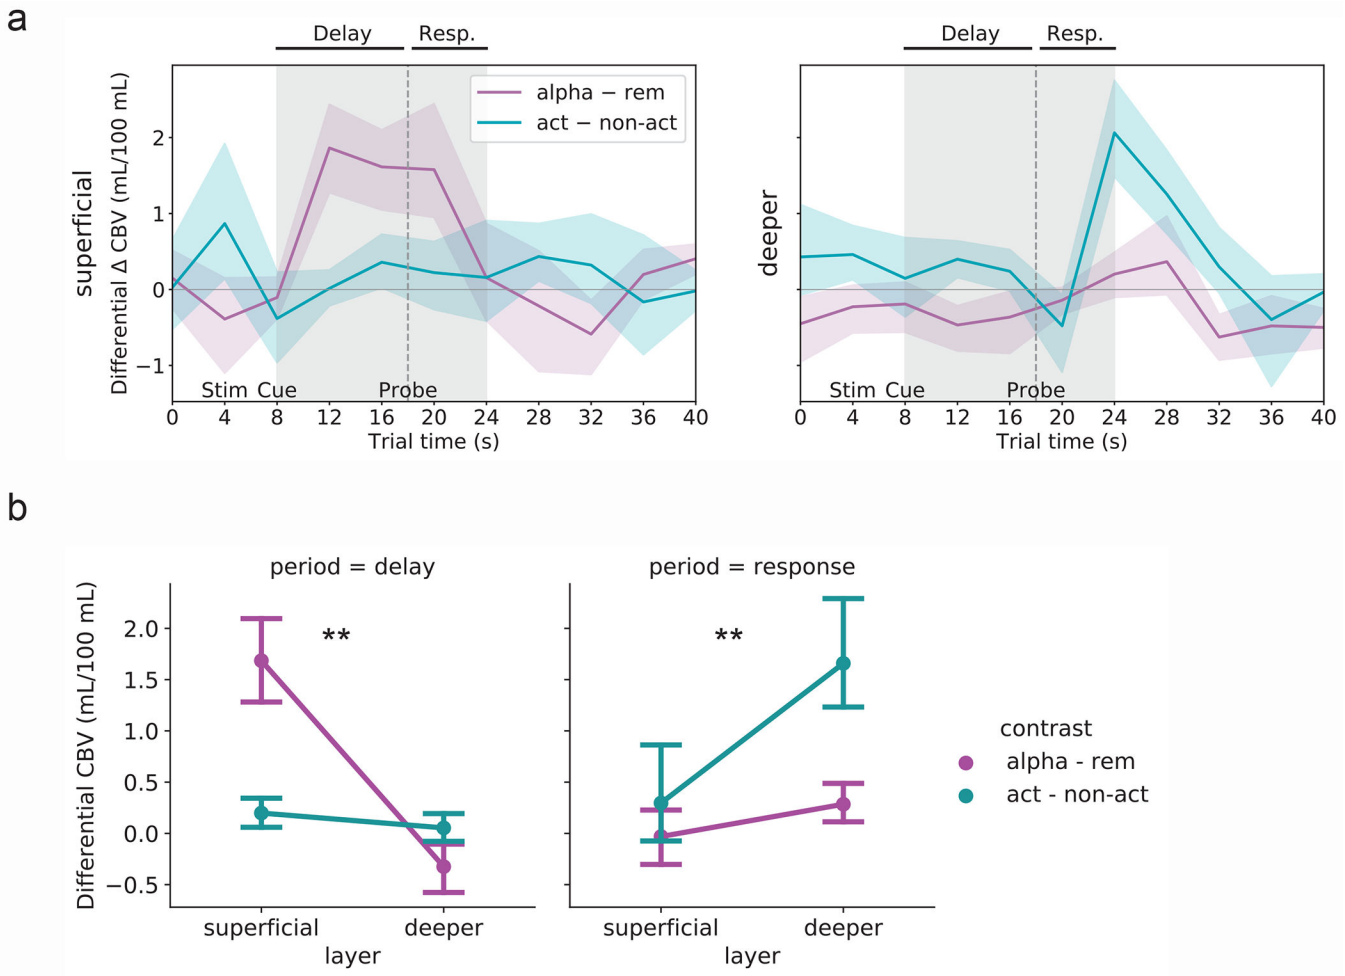

Fig. 3. Activity contrasts across layers and conditions of interest.

(A) Top: Superficial-layer VASO activity during maintenance ('rem') trials subtracted from activity during manipulation ('alpha') trials [purple line]. The largest difference can be seen during the delay period. Bottom: Deeper-layer VASO activity during non-action ('non-act') trials subtracted from activity during action ('act') trials [teal line]. The largest difference can be seen during the response period. Lines represent mean and shaded area represents 95 percent confidence intervals for the mean (determined via bootstrapping with 1,000 iterations) across n = 15 sessions (13 unique subjects; same data as in Fig. 2). See Fig. S5b for subtractions based on mean BOLD activity timecourses. (B) Two-way ANOVAs with factors layer (superficial versus deeper) and contrast (manipulation – maintenance ['alpha – rem', purple lines] versus action – non-action [teal lines]), for each trial period (delay and response). Dots represent mean and error bars reflect 95 percent confidence intervals for the mean (determined via bootstrapping with 1,000 iterations) across n = 15 sessions (13 unique subjects). **, interaction significant at $p < 0.001$ ($p = 6.9e^{-6}$ and $p = 3.0e^{-4}$ for the delay period [left] and response period [right], respectively).

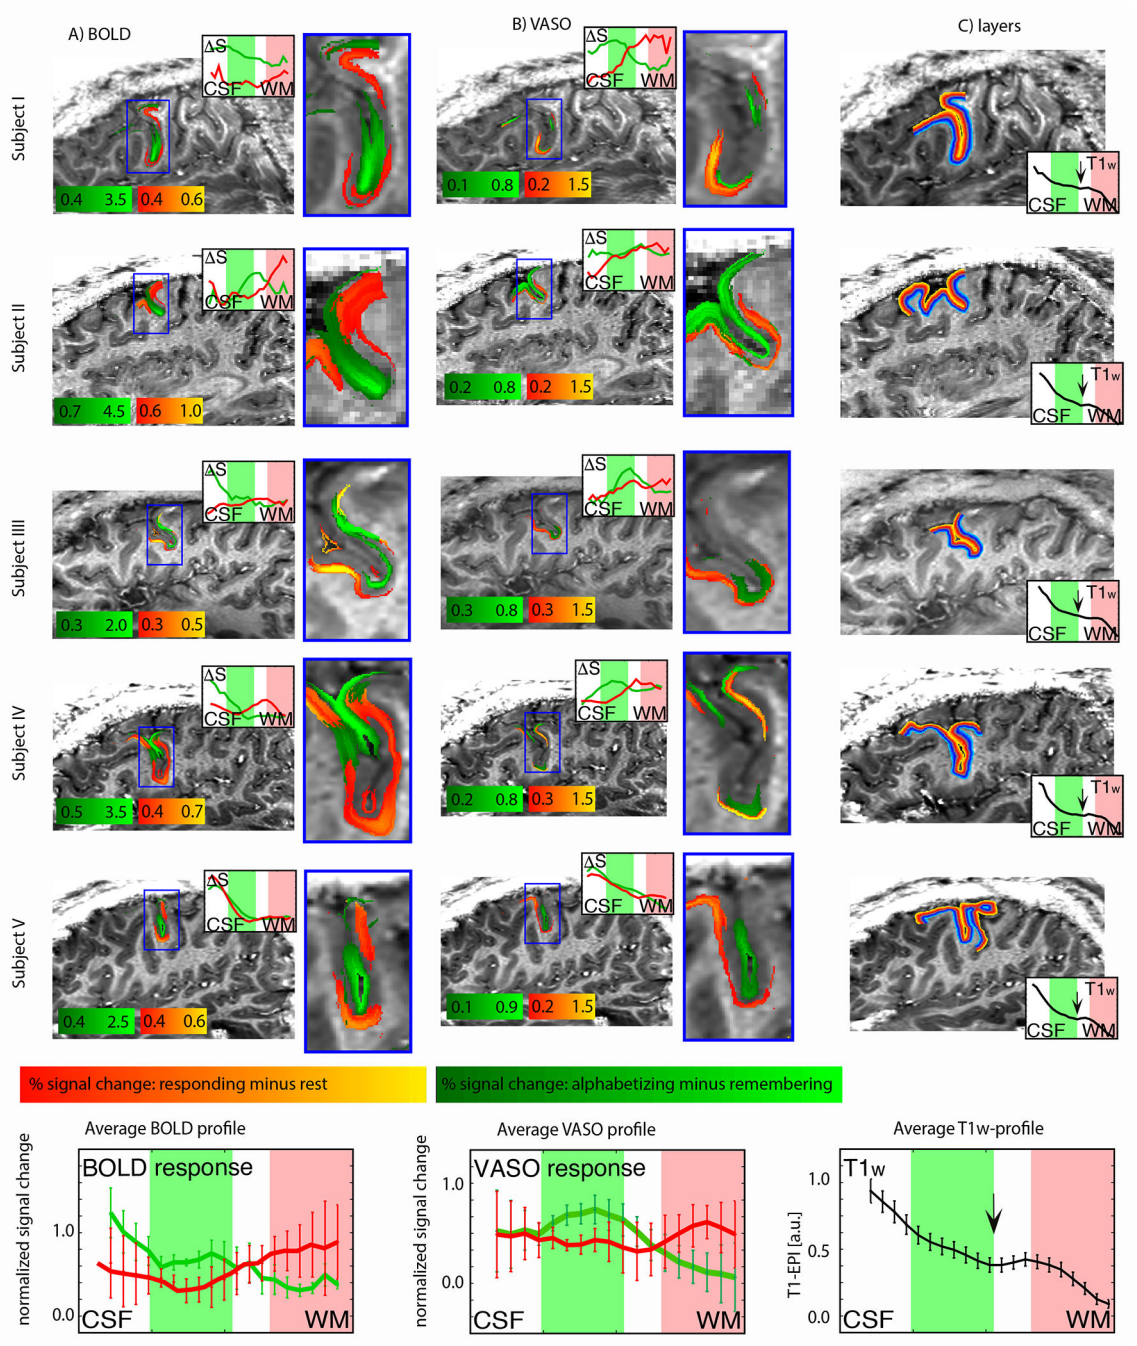

Fig. 4. Single-subject layer-dependent activity profiles.

Results from five subjects scanned using the sagittal protocol. Activity is shown in both functional contrasts, BOLD (A) and VASO (B). Signal changes for delay and response periods are smoothed within layers. No smoothing was applied across layers. Note the different color scales for BOLD and VASO. Color intensity indicates percent signal change. Red/orange reflects increased signal during the response period compared to baseline (inter-trial interval). Green represents increased signal during the delay period for manipulation compared to maintenance trials. Inset line graphs show the corresponding layer activity

profiles plotted across cortical depth. In VASO insets (B), note that the red line is always above the green line in the deeper layers (red shading), while the green line is always above the red line in the superficial layers (green shading), meaning that the task used here engages the layer superficial and deeper layers differently. This is consistent across subjects. Estimates of layers (cortical depths) for each subject are shown in (C). Insets in (C) are subject-specific layer profiles distribution of the T1-weighted EPI signal. The black arrow indicates the location of a myelin-related signal dip, which can be taken as a landmark for the transition region between cytoarchitectonic layer III and layer V (see Fig. S4). Error bars in average profiles (bottom row) reflect standard error of the mean across subjects.