NEUROSCIENCE of Consciousness

Neuroscience of Consciousness, 2022, **2022(1)**, 1–12

DOI: https://doi.org/10.1093/nc/niac005

Research Article

Special Issue: Consciousness science and its theories

Decoding rapidly presented visual stimuli from prefrontal ensembles without report nor post-perceptual processing

Joachim Bellet[1,*], Marion Gay[1], Abhilash Dwarakanath[1], Bechir Jarraya[1,2,3], Timo van Kerkoerle[1], Stanislas Dehaene[1,4], and Theofanis I. Panagiotaropoulos[1,†]

[1]Cognitive Neuroimaging Unit, Commissariat à l'Énergie Atomique et aux Énergies Alternatives, INSERM, Université Paris-Saclay, NeuroSpin, Gif-Sur-Yvette 91191, France; [2]University of Versailles Saint-Quentin-en-Yvelines, Versailles 78000, France; [3]Neuromodulation Unit, Foch Hospital, Suresnes 92150, France;[4]Collège de France, Université Paris-Sciences-Lettres (PSL), 11 Place Marcelin Berthelot, Paris 75005, France

[†]Theofanis I. Panagiotaropoulos, http://orcid.org/0000-0002-4621-1183

*Correspondence address. Cognitive Neuroimaging Unit, Commissariat à l'Énergie Atomique et aux Énergies Alternatives, INSERM, Université Paris-Saclay, NeuroSpin, Gif-Sur-Yvette 91191, France. E-mail: joachim.bellet@cea.fr

Abstract

The role of the primate prefrontal cortex (PFC) in conscious perception is debated. The global neuronal workspace theory of consciousness predicts that PFC neurons should contain a detailed code of the current conscious contents. Previous research showed that PFC is indeed activated in paradigms of conscious visual perception, including no-report paradigms where no voluntary behavioral report of the percept is given, thus avoiding a conflation of signals related to visual consciousness with signals related to the report. Still, it has been argued that prefrontal modulation could reflect post-perceptual processes that may be present even in the absence of report, such as thinking about the perceived stimulus, therefore reflecting a consequence rather than a direct correlate of conscious experience. Here, we investigate these issues by recording neuronal ensemble activity from the macaque ventrolateral PFC during briefly presented visual stimuli, either in isolated trials in which stimuli were clearly perceived or in sequences of rapid serial visual presentation (RSVP) in which perception and post-perceptual processing were challenged. We report that the identity of each stimulus could be decoded from PFC population activity even in the RSVP condition. The first visual signals could be detected at 60 ms after stimulus onset and information was maximal at 150 ms. However, in the RSVP condition, 200 ms after the onset of a stimulus, the decoding accuracy quickly dropped to chance level and the next stimulus started to be decodable. Interestingly, decoding in the ventrolateral PFC was stronger compared to posterior parietal cortex for both isolated and RSVP stimuli. These results indicate that neuronal populations in the macaque PFC reliably encode visual stimuli even under conditions that have been shown to challenge conscious perception and/or substantially reduce the probability of post-perceptual processing in humans. We discuss whether the observed activation reflects conscious access, phenomenal consciousness, or merely a preconscious bottom-up wave.

Keywords: conscious perception; prefrontal cortex; decoding; neuronal populations; electrophysiology; non-human primate; global neuronal workspace theory; integrated information theory

Introduction

During the last decades, there has been significant progress in our understanding of the neural correlates of consciousness (NCC), in particular, conscious visual perception (Koch 2004; Dehaene and Changeux 2011; Dehaene 2014; Koch *et al.* 2016). The field has advanced thanks to macroscopic brain imaging methods such as the blood-oxygen-level-dependent (BOLD) signal (measured with functional magnetic resonance imaging—fMRI), electroencephalography, magnetoencephalography, and electrocorticography signals (EEG, MEG, EcoG) that measure electric and magnetic field potentials. Significant information about localized single neurons and small neuronal populations also came from a handful of studies in human patients undergoing neural surgery and consenting for intracortical recordings (Kreiman *et al.* 2002;

Received: 19 January 2021; **Revised:** 9 December 2021; **Accepted:** 27 January 2022

Quiroga et al. 2008; Reber et al. 2017; Gelbard-Sagiv et al. 2018). The function of isolated neurons has also been thoroughly described in several parts of the visual system using monkeys trained to report their subjective percepts (Logothetis 1998). These studies show activation of neurons and modulation of signals in many cortical areas (and to various degrees) during conscious perception. However, due to a number of variables that could conflate perceptual signals, including the motor reports often used by subjects to report their percepts, it is still a matter of debate which mechanisms and which specific brain regions are involved in perceptual awareness. This is because the measured signals could potentially reflect subliminal processes preceding conscious experience or being a consequence of it such as introspection or report.

The potential implication of the prefrontal cortex (PFC) in consciousness is a pivotal element in these discussions, because of its importance in theories of consciousness. According to the global neuronal workspace (GNW) theory, PFC should be a key node in a network of global-workspace neurons whose population activity should encode, at any given moment, the details of the current conscious experience (Mashour et al. 2020). This role of PFC is questioned by Information Integration Theory (IIT), according to which, the major complex encoding perceptual experience should be primarily located in posterior areas, including sensory and parietal cortices (Tononi et al. 2016). Finally, higher-order thought (HOT) theories consider PFC important, not because PFC neurons are directly encoding the experience, but because a subset of PFC regions might encode a HOT about the experience (e.g. confidence in seeing a stimulus) (Lau and Rosenthal 2011; Brown et al. 2019). Thus, whether PFC populations contain a detailed code for the current perceptual experience is a central question for consciousness research.

Empirically, PFC activity has classically been associated with consciousness using fMRI, EEG, MEG, and intracranial recordings (Lumer and Rees 1999; Dehaene et al. 2001; Lau and Passingham 2006; Del Cul et al. 2009; Gaillard et al. 2009; Rounis et al. 2010; Li et al. 2014; Levinson et al. 2021). Also, electrical stimulation of the prefrontal cortex can elicit vivid hallucinations (Blanke et al. 2000; Vignal et al. 2000) and intrusive thoughts (Popa et al. 2016; Liu et al. 2020). Furthermore, in monkeys, PFC firing is modulated shortly before the report of a conscious percept of motion (Libedinsky and Livingstone 2011) and shows a characteristic ignition only in trials where a stimulus near the threshold of perception is perceived by the animal (Thompson and Schall 1999; Van Vugt et al. 2018). However, due to the role of PFC in higher-order cognitive processes, it has been argued that this activity could potentially reflect reporting, self-monitoring, and introspection processes related to the report requirements of the tasks, rather than perception per se (Frassle et al. 2014; Farooqui and Manly 2018).

Recently, there was evidence that the neuronal activity in the ventrolateral prefrontal cortex (VLPFC) continues to correlate with conscious content even in the absence of any need to report. Two no-report paradigms were used: binocular flash suppression (BFS) and binocular rivalry (BR) combined with eye-tracking. These paradigms are called 'no-report' because they allow inferring a conscious percept without requiring the subject to report it. In the BFS experiment, an image is first projected to one eye and it stays physically present when a second image is then suddenly presented in the second eye. The appearance of the second image irreversibly causes the conscious percept to switch from the first image to the second, although both images are displayed. The percept of the subject is identical to when the first image is removed from the screen at the same time the second image is flashed (Wolfe 1984). These BFS experiments enabled one of us to reveal PFC neurons specifically activated when a stimulus is consciously perceived rather than unconsciously processed (Panagiotaropoulos et al. 2012; Kapoor et al. 2018). The second approach, the BR experiments, consisted of presenting simultaneously a drifting grating to one eye and another grating to the other eye, drifting in the opposite direction compared to the first grating. Under these conditions, the content of visual consciousness alternates spontaneously between the two opposing stimuli. These stimuli cause involuntary optokinetic nystagmus, with the direction of the smooth pursuit phase correlating with the content of conscious perception in humans (Naber et al. 2011). In these experiments, further evidence was provided revealing that PFC neural activity correlates with the contents of consciousness in the absence of explicit reports (Dwarakanath et al. 2020; Kapoor et al. 2020). Importantly, the code for the direction of motion could generalize to trials without rivalry (and vice versa) using a controlled alternation of stimulus direction.

However, because they contrast only two stimuli, these no-report experiments provide a limited test of the GNW theory, according to which the PFC should participate in the encoding of every conscious percept. Moreover, it was recently pointed out that these PFC correlates of consciousness might just be post-perceptual signals, that is, a consequence but not a cause of consciousness even in such no-report paradigms (Block 2020, but see Panagiotaropoulos et al. 2020). According to this view, the PFC activity is related only to what can be reported (access consciousness), which (so the argument goes) is different from what is actually felt or experienced by the subject (phenomenal consciousness). Phenomenal consciousness is suggested to occur before access consciousness, in hierarchically earlier cortical areas, and to contain richer sensory information than what can be reported by the subject—the 'overflow' argument (Block 2011).

In the present study, we provide evidence that neural representations at the level of the VLPFC are rich and occur shortly after stimulus onset. We show that the identity of each of the presented images can be decoded very shortly after stimulus onset, at timescales that correspond to when feedforward computations are still being processed in most visual areas (Maunsell and Gibson 1992; Nowak et al. 1995; Schmolesky et al. 1998; Pouget et al. 2005; Self et al. 2013; Chen et al. 2018; Zamarashkina et al. 2020). We further show that significant information is still present during rapid serial visual presentation (RSVP) when images are presented in streams of 10 images per second, a rate at which access consciousness and post-perceptual processing is challenged (Lawrence 1971; Broadbent and Broadbent 1987; Raymond et al. 1992; Sergent et al. 2005; Marti and Dehaene 2017) but the phenomenal percept is still one of seeing a stream of very brief images, without necessarily being able to report the image identity. In this procedure of RSVP, if the images are consecutive words forming a sentence, human subjects are able to understand it and repeat it, even when an occasional noun is replaced by the corresponding picture (Potter et al. 1986). However, when having to report, and especially if the stream of stimuli cannot be integrated into a single representation such as a sentence, humans typically fail to report at least some of them (Broadbent and Broadbent 1987; Raymond et al. 1992). This is true even in conditions avoiding the so-called attentional blink, where the active report of a target renders a subsequent target invisible (Olivers et al. 2007). Thus, conscious access and post-perceptual processes are unlikely to happen for every stimulus of a RSVP sequence. Yet, we found that the decoding accuracy of each consecutive stimulus, at its peak, was of the same magnitude during RSVP as when

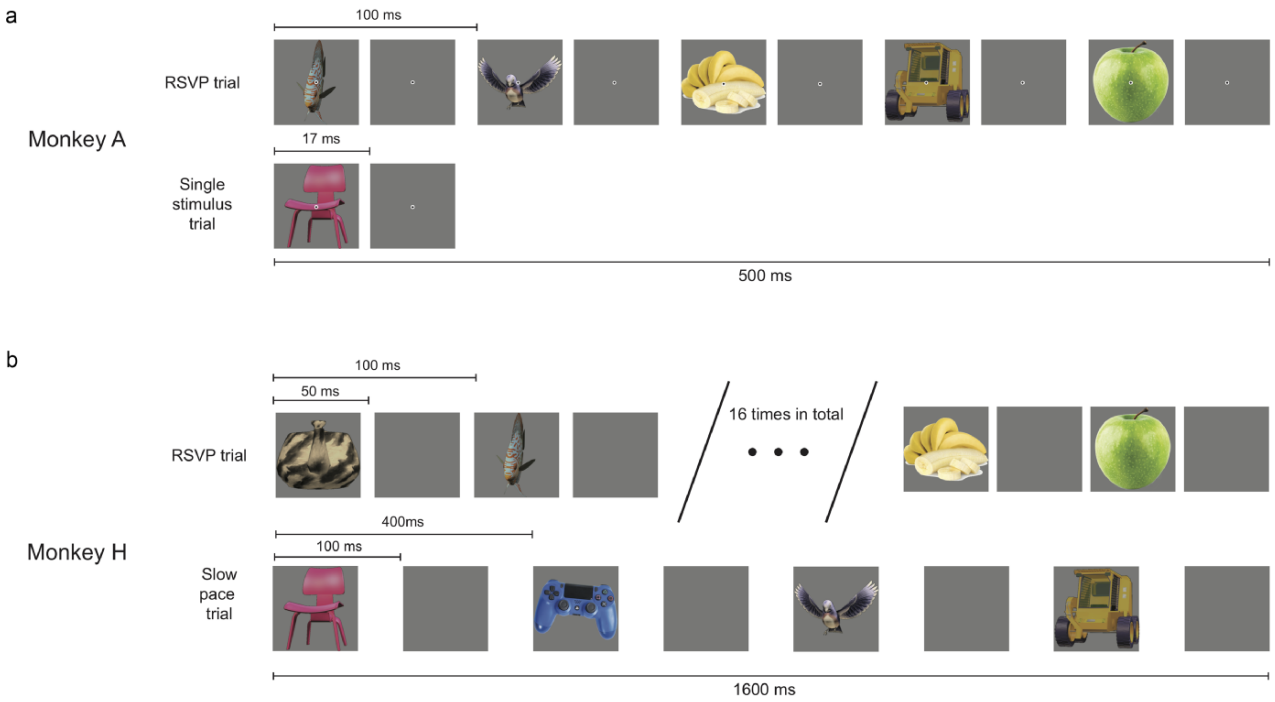

Figure 1. Task design. (a) Task performed by monkey A. After an initial fixation of 300 ms, either a stream of five random stimuli would appear for one frame each within 500 ms (RSVP trial) or a single random stimulus would appear for one screen frame (single stimulus trial). For both types of trials, a juice reward was delivered at a random time ranging between 500 and 700 ms relative to the first stimulus onset. Note that for the sake of figure readability, the time duration depicting a single screen frame is not to scale with the inter-stimulus interval and the duration of a trial. (b) Task performed by monkey H. After an initial fixation of 300 ms, a stream of stimuli would be presented over 1.6 s. Some trials contained 16 stimuli with 100 ms SOA (RSVP trial), others contained four stimuli with 400 ms SOA

the same stimulus was presented in isolation and clearly visible. Therefore, encoding of visual information in the VLPFC can occur prior to post-perceptual processes and access consciousness. We also report a stable and sustained code emerging after the peak of maximal information at 150 ms when stimuli were presented in isolation. Such a stable representation is a hallmark of the reverberating activity constituting NCCs in the GNW theory (Dehaene *et al.* 2003) and, interestingly, was reduced in the RSVP condition. We later discuss how our observations raise challenges for both GNW and IIT theories and may help refine our understanding of the role of PFC in the NCCs.

Results

Visual information allows linear classification of every stimulus and reaches VLPFC rapidly

We recorded multi-unit spiking activity (MUA) using Utah arrays composed of 96 electrodes implanted in the VLPFC (inferior convexity, area 45a) of two macaque monkeys, monkey A and monkey H. The main figures of this article display the results obtained from monkey A. In monkey H, the quality of the VLPFC recording was lower with only 32 out of 96 channels sensing more than one spike per second. Yet, all results obtained in monkey A were replicated in monkey H (those results are presented as Supplementary Material). Monkey A and monkey H performed similar tasks, passively viewing stimuli presented either at a slow pace or in a RSVP sequence of 10 images per second (Fig. 1).

We trained linear classifiers to predict the displayed image from neural activity. Each of the 18 images used in our experiment could be decoded with high accuracy (Fig. 2 for results of monkey A and Supplementary Fig. S1 for results of monkey H). Note that for all the timing information provided in our study, the decoding score at a given time point rests upon all the spikes collected from the active channels in a window of 50-ms duration preceding that time point. Decodable information reached the VLPFC already by 75 ms for the first image (representing spikes collected in the interval 25–75 ms after image onset; Fig. 3b).

To ensure that the above chance decoding accuracy was not due to a single image or a few images leading to unusually fast response in the VLPFC, we performed fine-grained measurement of the timing of decoding accuracy between every 153 possible pairs of images (Supplementary Fig. S4). The media first time of significant decoding was 60 ms in monkey A and 90 ms in monkey H (Supplementary Fig. S4). Information was maximal at 150 ms and slowly decayed afterwards (Fig. 3b).

We compared the decoding accuracy of VLPFC to a portion of the lateral surface of the posterior parietal cortex (PPC), approximately over areas 7a/7b (Petrides and Pandya 2009), which is connected directly to the VLPFC, and is likely to be involved in conscious processing according to both IIT and GNW theories. Despite PPC being part of the dorsal pathway, decoders could classify the images above chance level. However, the decoding accuracy was higher in VLPFC than in PPC when keeping the number of training samples and the number of active channels equal between both areas (Supplementary Fig. S5). This result shows that the visual information available to VLPFC is superior to that of PPC, which is part of the posterior cortical 'hot zone' postulated in IIT to be especially important for consciousness.

We also assessed the sparseness of the VLPFC neural population. Decoding remained possible when up to 50% of the most informative channels were excluded from the analysis

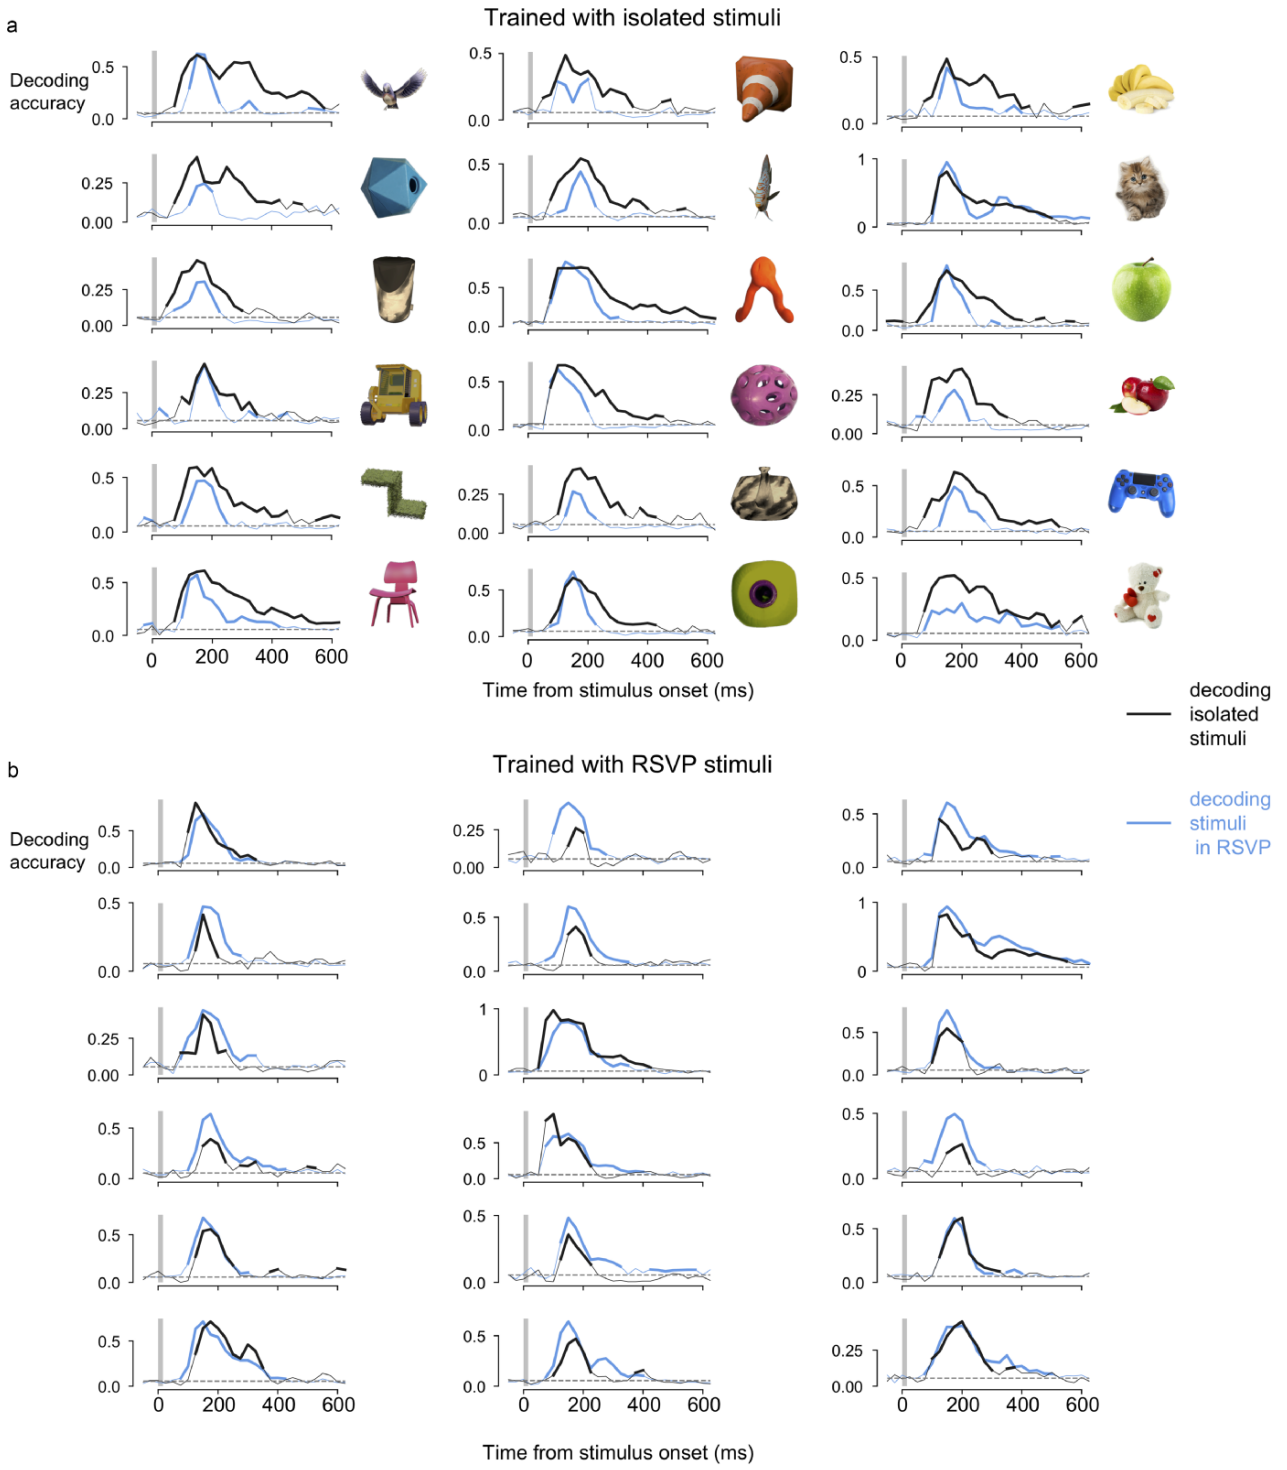

Figure 2. Every tested image is decodable from the neural activity in the VLPFC of monkey A. (a) Time course of the accuracy of decoder relative to stimulus onset for decoders trained with isolated stimuli. Black lines are accuracy levels averaged across all occurrences of a given stimulus in isolation. Blue lines are accuracy levels averaged across stimuli occurring in positions 2, 3, or 4 of the RSVP sequence. Thick lines depict groups of time points for which the accuracy is significantly higher than the chance level ($P < 0.001$, random permutations with cluster correction for multiple comparisons). Gray vertical bars depict the duration of the physical presence of the stimulus on the screen. Gray dotted lines indicate chance level. (b) Same as (a) but with decoders trained with stimuli in positions 2, 3, or 4 of the RSVP sequence

(Supplementary Fig. S6), indicating that the representation of stimulus identity was distributed over a large VLPFC neural population. However, conversely, keeping 10% of the most informative channels sufficed to reach an asymptotic level of decoding. Those results suggest that the representation is based on both a relatively sparse set of highly discriminating units (~10%) and a more distributed set of weakly informative ones, representing ~50% of VLFPC units.

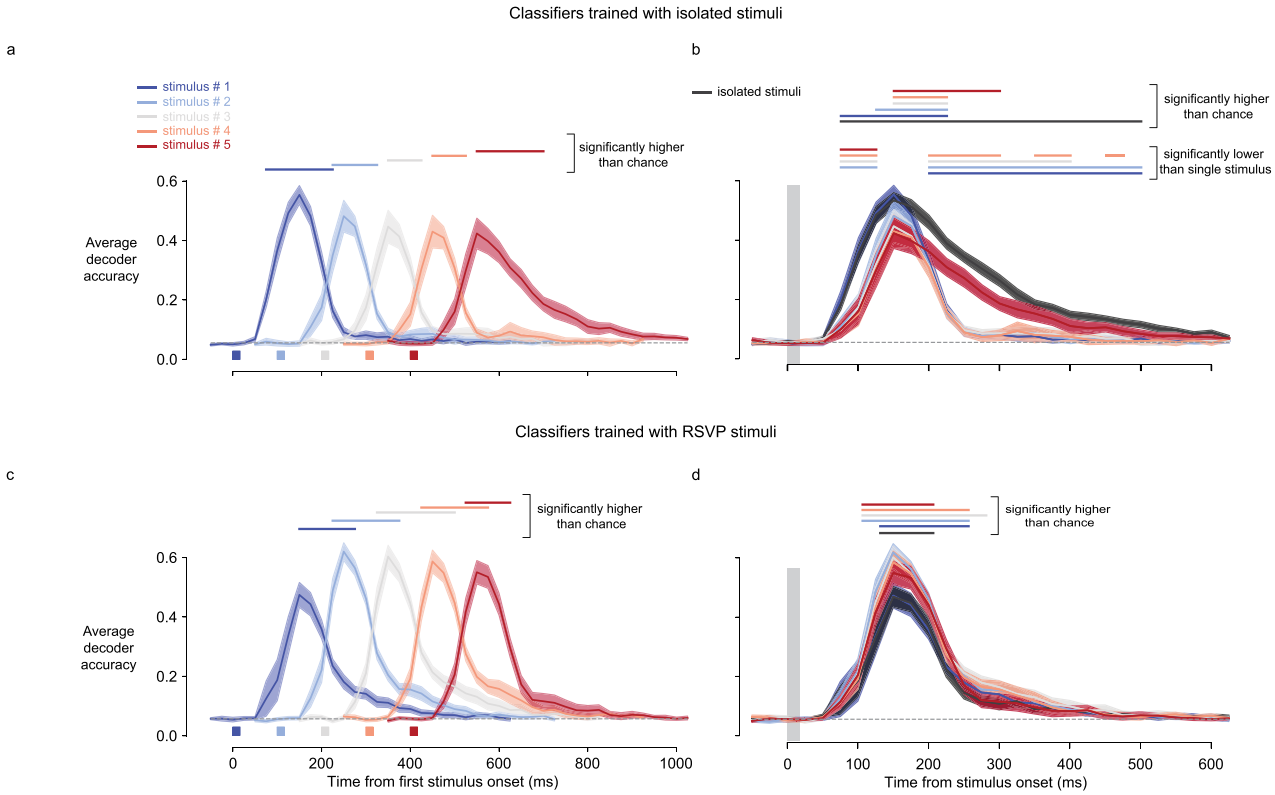

Figure 3. Decoder accuracy as a function of time for stimuli presented in RSVP and stimuli presented in isolation. (a) Grand average decoder accuracy across the 18 accuracy levels corresponding to every stimulus. Accuracy is averaged independently for each stimulus of the RSVP sequence as a function of time since the first stimulus onset. Error bars are the standard error of the mean across the 18 stimuli. Horizontal bars above the curves indicate time points with decoder accuracy higher than chance level ($P < 0.001$, single-sample T-tests corrected for multiple comparisons using false discovery rate procedure). Rectangles under the curves represent the time of physical presence of the stimuli on the screen. (b) Same as in (a) but aligned to stimulus onset and adding the curve for stimuli presented in isolation (black curve). Uppermost horizontal lines display the same statistics as in (a) but adding the curve for the isolated stimuli. The other horizontal lines represent the time for which stimuli presented in the RSVP sequence are less decodable than stimuli presented in isolation ($P < 0.001$, dependent T-tests corrected for multiple comparisons using false discovery rate procedure). The gray vertical rectangle represents the time of physical presence of the stimuli on the screen. (c) Same as in (a) but with decoders trained with stimuli in positions 2, 3, or 4 of the RSVP sequence. (d) Same as in (b) but with decoders trained with stimuli in positions 2, 3, or 4 of the RSVP sequence

VLPFC represents visual stimuli also in RSVP conditions

Our experimental design included RSVP trials with stimuli presented only 100 ms apart. When we applied the decoders trained on single-image trials to RSVP data, we found that every stimulus could still be decoded in monkey A (Fig. 2a, blue traces) and 17 out of 18 stimuli in monkey H (Supplementary Fig. S1a, blue traces). On average, the peak decoding accuracy for stimuli presented at any of the positions in the RSVP sequence was of the same magnitude (Fig. 3a and b and Supplementary Fig. S2a and b). When compared to stimuli presented in isolation, the RSVP altered decodability in a time-dependent manner. Information was decodable only starting from 125 ms following stimulus onset when it occurred in positions 2, 3, 4, or 5 of the RSVP sequence (Fig. 3b, Supplementary Fig. S2b). At 150 ms, the magnitude of the average accuracy was similar to that of isolated stimuli. At that time, the difference with the accuracy for the isolated stimulus condition was rather weak (P-values corrected for multiple comparisons using the Benjamini–Hochberg false discovery rate (FDR) procedure are: 0.213, 0.124, 0.009, 0.002, 0.001 for stimuli presented in positions 1, 2, 3, 4, and 5, respectively; Fig. 3b). However, at 250 ms, the decodability was drastically reduced for stimuli that were followed by another one. The difference with the accuracy in isolated stimulus condition was substantial (P-values FDR corrected: 4.02×10^{-08}, 2.01×10^{-08}, 1.92×10^{-08}, 1.18×10^{-08}, 0.050 for stimuli presented in positions 1, 2, 3, 4, and 5, respectively; Fig. 3b). Thus, consecutive images had a suppressive effect on the neural code for the earlier images, potentially reflecting masking, as previously reported in the inferotemporal cortex (Kovacs *et al.* 1995). In summary, RSVP both delayed the peak time at which information was decodable and impaired the late sustained neuronal activity but had little effect on the amount of information carried by VLPFC at its peak 150 ms after stimulus onset.

When decoders were trained and tested on neural activity evoked by stimuli in the RSVP sequence, information could be retrieved for up to 275 ms after stimulus onset (Fig. 3c and d) compared to 225 ms with classifiers trained on the isolated-stimuli condition (Fig. 3a) and the decoding accuracy peaked slightly higher than with classifiers trained on isolated stimuli. However, the classifier trained with the RSVP condition did not reveal the sustained information of stimuli presented in isolation and of the fifth stimulus in the sequence (Fig. 3d). The discrepancy between the behavior of classifiers trained with isolated stimuli and classifiers trained with RSVP stimuli suggests that a neural code that was present in the late time period after the presentation of an isolated image was reduced or shortened in the RSVP condition.

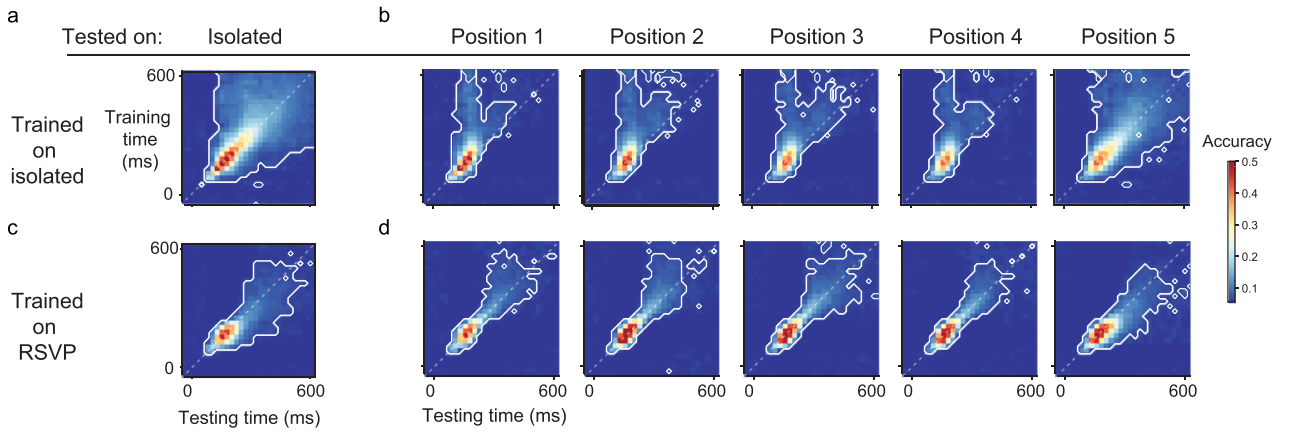

Figure 4. Temporal generalization matrices. (a) We trained independent decoders for each time bin around the onset of isolated stimuli. These decoders were then tested on the other time bins and we display here their accuracy. The gray dashed line shows times at which training and testing times coincide. White contours delineate clusters of pixels for which accuracy is significantly higher than chance ($P < 0.001$, random permutations corrected for multiple comparisons using the Benjamini–Hochberg false discovery rate procedure). Note that after an early period with time-dependent representations, the code becomes stable and generalizes to other time bins. (b) Same as in (a). The classifiers are still trained with isolated stimuli but tested on stimuli presented in RSVP sequences. Note that for stimuli in positions 1–4, late decoders can decode early activity but the reverse is not true. (c) Same as in (a) but with decoders trained with stimuli in positions 2, 3, or 4 of the RSVP sequence. Note that early decoders can decode late activity but the reverse is not true. (d) Same as in (b) but with decoders trained with stimuli in positions 2, 3, or 4 of the RSVP sequence

A dynamic code coexists with a sustained representation of stimuli

To clarify the previous decoding observations, we examined how the decoders generalize over time. Recent MEG experiments in humans suggest that, once a stimulus has been consciously accessed, regardless of its original duration, it is internally represented by a metastable code that remains relatively fixed for a period of several hundreds of milliseconds (King *et al.* 2014, 2016; Schurger *et al.* 2015a; Baria *et al.* 2017; Marti and Dehaene 2017). Because of the limited spatial resolution of MEG, the brain areas contributing to these stable representations are unknown. Therefore, we analyzed the generalization-across-time (GAT) matrices of our decoders (King and Dehaene 2014). A square GAT matrix indicates a metastable neural representation, as previously reported for conscious visual and auditory stimuli (King *et al.*, 2014, 2016; Baria *et al.* 2017; Marti and Dehaene 2017). We investigated whether the neural representations in VLPFC also exhibit this property. To this end, using only trials from the single-image condition, in which the stimulus was clearly visible, we trained the decoders at specific time bins after stimulus onset and tested performance at this and other time bins. In the single-image condition, a dynamic code unfolded with time from stimulus onset, first showing a narrow diagonal and then an increasingly longer window of temporal generalization, and ~200 ms with a nearly square GAT matrix indicating a metastable representation (Fig. 4a). When testing the same decoders on neural activity evoked by the RSVP stimuli, we still observed the diagonal component and early activity being accurately classified by late decoders (Fig. 4b, diagonal, and above). However, for stimuli presented in positions 1–4, the late decoders showed reduced performance, and the corresponding period of significant decoding was shortened. Still, the decoders that, for a single image, achieved stable and significant decoding from 150 to 600 ms remained significant from ~175 to 400 ms after RSVP picture onset. Significant off-diagonal decoding (King and Dehaene 2014) indicated that the late wave of PFC neural activity was compressed in time, but did not totally vanish—for instance, even the decoders at 400–500 ms after single-image onset continued to provide above chance decoding in the RSVP context up to 400 ms or more. Finally, the 5th RSVP image, which was not followed by any image, could be decoded with a longer-duration GAT matrix, similar to a single image (Fig. 4b). In summary, in the RSVP condition, virtually all of the successive neural representations that normally unfold until ~500 ms after the stimulus remained present for every picture: they were shorter-lived and compressed, but still overlapped partially with each other.

Classifiers trained exclusively with stimuli presented during the RSVP sequence did not reveal a square GAT matrix when tested on isolated stimuli (Fig. 4d and Supplementary S3d). Yet, in the isolated condition, they could also classify above chance neural activity occurring later than the time interval they were trained on (Fig. 4c and Supplementary Fig. S3c). This could be interpreted as the dynamic neural code being compressed in time during the RSVP sequence compared to the isolated stimulus condition. It is also likely that the square GAT matrix in the isolated condition results from the maintenance of the neural ensembles successively activated during the dynamic phase. The asymmetries when testing the generalization of classifiers from one condition to the other would be a consequence of a lack of sustained activity of the sequentially activated neural ensembles during the RSVP condition.

Discussion

In the current experiment, we studied the response of populations of VLPFC neurons to isolated and rapidly presented visual stimuli (RSVP) in two macaque monkeys. Using linear decoders of spiking activity at the neuronal population level, we found that rich information was available as early as 75 ms after the onset of a visual stimulus. We further showed that the peak in decodability occurring at 150 ms is unaltered in an RSVP condition where the monkey is unlikely to have any time to introspect on every stimulus, due to the rapid change in visual input.

Because timing can be crucial for brain processing (Singer 1999; Oram *et al.* 2002), one can only grasp the importance of the present results in light of the latencies previously described in other visual areas. At 75 ms, most of the feedforward visual processing is already completed in the lateral geniculate nucleus

(Schmolesky *et al.* 1998), superior colliculus (Chen *et al.* 2018), and area V1 (Maunsell and Gibson 1992; Nowak *et al.* 1995; Schmolesky *et al.* 1998; Self *et al.* 2013). However, in anesthetized monkeys, the proportion of neurons having emitted a first spike in response to a visual stimulus is only ∼60% in areas of the dorsal stream (V3, MT, MST, and FEF) and, in the ventral stream, 30% and 10% for areas V2 and V4, respectively (Schmolesky *et al.* 1998). Those timings are shorter in awake animals: at 75 ms, 70% of FEF neurons have started to respond (Pouget *et al.* 2005) and close to 50% in V4 (Zamarashkina *et al.* 2020). Thus, early areas in the visual hierarchy are still processing information when it reaches VLPFC. Moreover, object identity can only be reliably decoded after 100 ms in the inferotemporal cortex (Hung *et al.* 2005). Thus, the VLPFC is likely to be actively involved in the process of object recognition even before the feedforward sweep of visual information processing has been completed. This interpretation is supported by a recent study showing that pharmacological inactivation of VLPFC impairs category decoding in inferotemporal cortex (Kar and DiCarlo 2021). Our results are in line with the hypothesis that an early wave of PFC activity provides facilitatory top-down inputs to the ventral object recognition pathway (Bar *et al.* 2006). In the present study, half of the pairs of images could be discriminated from each other before 60 ms based on our decoding results (Supplementary Fig. S5). Besides the classical ventral occipito-temporal pathway, another fast anatomical route for visual information to reach VLPFC is through the superior colliculus (Benevento and Fallon 1975; Chen *et al.* 2018) via the mediodorsal thalamus (Goldman-Rakic and Porrino 1985; Sommer and Wurtz 2004). Given that at 60 ms the superior colliculus mainly responds to low spatial frequencies (Chen *et al.* 2018) and that the magno-cellular pathway, sensitive to low-spatial spatial frequencies, dominates over the parvocellular pathway (Schmolesky *et al.* 1998), it seems likely that VLPFC gains first access to coarse visual representations and integrates later information processed by the ventral stream. Therefore, the earliest responses of the VLPFC are unlikely to constitute a full code for the contents of consciousness, but only a coarse first pass.

Our results also support a prediction of the GNW model, according to which populations of PFC neurons encode every aspect of conscious content, even in the absence of report or of sufficient time for complex post-perceptual operations. It is remarkable that as many as 18 different images could be decoded way above chance levels from as few as 96 channels in monkey A and 32 channels in monkey H. Since the images depicted three-dimensional (3D) objects that differed in multiple ways, including identity, color, and orientation, future work, using stimuli that vary these features in a systematic manner, will be required to understand whether and how the PFC code can be 'factorized' into these different dimensions. With regard to theories of consciousness, it will be essential to probe to what extent the PFC code reflects every aspect of perceptual experience, but also fails to reflect dimensions that are not perceived consciously, as predicted by GNW. It has recently been shown that the inferotemporal cortex ensembles, from which VLPFC receives direct inputs, encode object properties usually attributed to the dorsal stream (e.g. size and position of an object) (Hong *et al.* 2016). Future experiments need to be designed to investigate whether this is also true in VLPFC.

These results, taken together with previous work (Ó Scalaidhe *et al.* 1999; Constantinidis *et al.* 2001; Freedman *et al.* 2001; Fuster 2015; Riley *et al.* 2017), establish that PFC contains a rich visual code, richer than the binary code used in previous experiments studying visual consciousness using perceptual suppression (Panagiotaropoulos *et al.* 2012; Kapoor *et al.* 2018, 2020; Dwarakanath *et al.* 2020), and does so in a passive fixation task, and even under an RSVP protocol that challenges conscious perception of stimuli at least according to human data. The finding of a rich PFC code during RSVP challenges the view that the prefrontal cortex is only involved in cognitive or post-perceptual processes, secondary to conscious perception of visual stimuli, and that most of the details of conscious experience are specified by posterior areas (Tononi *et al.* 2016; Boly *et al.* 2017). In fact, we found that decoding accuracy was higher in the VLPFC than in the lateral PPC, two areas known to be directly connected through the superior longitudinal fasciculus (Petrides and Pandya 2009), even during the early phase of response. Thus, VLPFC could be as much a valid candidate for processing phenomenal consciousness as PPC or other areas upstream in the visual hierarchy. We should stress that the PPC still enabled significant decoding of image identity although it is part of the dorsal stream traditionally described for processing location and spatial relationships as opposed to the ventral stream involved in object recognition (Mishkin *et al.* 1983). This observation supports both the GNW theory and IIT according to which neural correlates of consciousness are integrated in large portions of the cortex, including PPC. It is likely that decoding accuracy in PPC could be on par with other parts of the PFC less connected to the ventral stream than is VLPFC or that other PPC areas exhibit a stronger perceptual effect. Still, recordings of single neurons in the lateral intraparietal area, dorsal to the parietal areas we recorded, showed that only a minority of neurons exhibited stimulus and perceptual selectivity during BFS and these responses were modulated by association of the visual stimulus with reward (Bahmani *et al.* 2019).

In IIT, the 'major complex' of brain regions normally involved in phenomenology entails the inferotemporal cortex (Tononi *et al.* 2016; Boly *et al.* 2017). In this region, the decoding of image identity is possibly the highest for a given number of channels (Majaj *et al.* 2015; Hong *et al.* 2016). However, NCCs are likely composed of populations that are orders of magnitude bigger than those measured experimentally. Therefore, the magnitude of decoding accuracy, alone, should be discussed with caution when considering an area as being involved in NCC or not. In general, any part of the brain exhibiting unique states of activity corresponding to any perceptually differentiable stimuli should be considered as a potential candidate for processing phenomenal consciousness.

Due to the lack of subjective report under the present RSVP conditions, two interpretations of the data remain tenable: macaques could either be very fast in perceiving the stimuli or most of the stimuli could have remained unconscious. The interpretation of the results could therefore change depending on the subjective experience of the animals, to which we did not have access. In any case, the current data pose new challenges and questions in the discussion about the neural correlates of phenomenal and access consciousness and the nature of visual information processing related to consciousness in the PFC.

Furthermore, our results also contradict the interpretation of a recent study investigating extensively the causal effect of intracranial electrical stimulation on conscious perception in humans (Fox *et al.* 2020). This study and a recent review suggest that electrical stimulation of PFC rarely elicits exogenous sensations (Raccah *et al.* 2021, but see Blanke *et al.* 2000; Vignal *et al.* 2000). However, electrical microstimulation activates neurons within a volume of about 1 mm in diameter around the tip of the electrode, independently of the current used (Histed *et al.* 2009). Thus, it is not

surprising that well-defined percepts can be elicited by stimulating sensory areas with topological organization (e.g. retinotopy, tonotopy, etc.) but that it is more difficult to elicit such percepts in areas like PFC that seem to lack topography, but rely on highly distributed and overlapping neural representations.

These decoding results also indicate that the VLPFC code for stimulus identity achieves stability after about 200 ms, which agrees with a lower estimate of the correlates of access consciousness reported in humans (King *et al.*, 2014, 2016; Schurger *et al.* 2015b; Baria *et al.* 2017; Marti and Dehaene 2017) and the ignition postulated in the GNW theory of consciousness (Mashour *et al.* 2020). Interestingly, we found that the RSVP condition did not fully abolish this late sustained activity, but only significantly shortened and reduced it. Although the early feedforward wave was dominant, weak but significant later responses could also be detected. Since the present work arises from an averaging of about 3500 trials per condition, future work will be needed to establish if the responses vary on a trial-by-trial basis or if they are truly present on every trial. At any rate, they replicate and extend, within a single brain area (VLPFC) in the apex of the cortical visual processing hierarchy, previous human MEG findings that all successive pictures in an RSVP stream can be encoded at several successive stages of processing (Marti and Dehaene 2017).

Given that human observers declare seeing most if not all stimuli of a RSVP stream, yet are unable to report them all, and in the eventuality that the stable code observed in this no-report study describes the same phenomenon as the stable code previously attributed to conscious access (King *et al.* 2016), its partial suppression under RSVP can be interpreted in several ways. First, RSVP stimuli could be consciously perceived, but too flimsily to lay down a short-term memory that would persist long enough to be later reportable in full. Alternatively, stimuli in the RSVP may not all be consciously perceived during their presentation—the impression of seeing them all would be, in part, an illusion, a consequence of our ability to consciously access any of them later on, as demonstrated by the phenomenon of retro-cueing (Sperling 1960; Sergent *et al.* 2013; Thibault *et al.* 2016; Xia *et al.* 2016). Under this interpretation, the transient neural code that we observed would correspond to what has been termed a 'pre-conscious' representation, and only some of them (perhaps variable across trials) would lead to genuine, stable conscious access. Such non-conscious representations of sensory input have been previously reported in the PFC using fMRI (Van Gaal *et al.* 2010; Levinson *et al.* 2021; Mei *et al.* 2021).

The main limitation of the present study is the absence of behavioral reports that could have provided insight as to whether any of the signals we have decoded are correlates of consciousness. A report version of the RSVP paradigm would be necessary to disentangle the potential contribution of the different representations to access consciousness. Alternatively, training animals to report their subjective feeling of similarity between stimuli could serve as a basis for representational similarity analysis (Kriegeskorte *et al.* 2008) revealing which of the successive neural representations display a geometry that reflects the subjective distance between stimuli. Such experiments could help refine theories of consciousness. For instance, the GNW theory, which makes no distinction between phenomenal and access consciousness (Naccache and Dehaene 2007), predicts that the late sustained and stable representation observed in our study corresponds to NCCs. Thus, according to this theory, the stable representation should be specifically reduced in trials when an animal fails to report a target and should have a geometry reflecting the subjective proximity between stimuli. If, on the other hand, subjective similarity between stimuli is reflected by the early phase of PFC activity during which the code does not generalize to later neural activity, then this early PFC response could contribute to pure phenomenal consciousness in the framework of the IIT theory. Still, if the latter is true, PFC could be involved in phenomenal consciousness, suggesting a revision of both GNW theory and IIT.

Methods

All procedures were conducted following the European convention for animal care (86-406), the National Institutes of Health's Guide for the Care and Use of Laboratory Animals, and were approved by the institutional Ethical Committee (CETEA 10-003). Our license number for animal experiments was A19_068.

Experiment

Two adult male macaque monkeys, monkey A and monkey H, participated in the experiment. The monkeys were never trained to report their perception of the stimuli used or perform any kind of report in other tasks. Monkeys A's task was to fixate a dot at the center of the screen for a duration uniformly distributed between 900 and 1100 ms to obtain a liquid reward composed of water and apple juice. During fixation, 18 pictures subtending 8 degrees of visual angle appeared briefly on the screen. The nature of the images was irrelevant to the task (neither the amount of reward nor the total duration of fixation depended on it). Every image used in the experiment is displayed in Fig. 2. The 12 images in the first two columns are 3D objects generated in Blender. The six images from the third column were obtained from https://pixabay.com/ and are under CC0 license. After 300 ms following fixation onset, images were flashed foveally for 17 ms each. On half of the trials (isolated stimulus trials), only a single image appeared. In the remaining trials, five images were presented with a stimulus onset asynchrony (SOA) of 100 ms (RSVP trials). Throughout an experiment, we presented in alternation pairs of RSVP trials and pairs of isolated stimulus trials. On each trial, the stimuli were randomly chosen without replacement among 18 colored images of single objects in a gray background. A total of 3265 successful RSVP trials and 3497 isolated stimulus trials were collected over three consecutive sessions.

Monkey H performed a similar task than monkey A but with slight modifications. After 300 ms following fixation onset, the fixation spot disappeared and the monkey was free to move his eyes in a 5-degree window around the center of the images displayed. The monkey had to keep his gaze on the stream of images for 1.6 s after the onset of the first image in order to obtain the liquid reward. Three types of trials could occur. Either 16 stimuli with 100 ms SOA each presented for 50 ms on the screen (RSVP condition). Or 4 stimuli with 400 ms SOA each presented for 100 ms on the screen (isolated stimuli condition). Or 8 stimuli with 200 ms SOA each presented for 100 ms on the screen. The latter condition was not included in the analysis of this paper.

Apparatus

During the experiment, the monkeys sat in a primate chair. Head movements were restricted by employing a skull-form-specific, surgically implanted head post. Eye movements were recorded using an infrared eye tracker. We recorded neural activity in the VLPFC using a Utah array connected to a BlackRock acquisition system. Wide-band neural signals were collected at 30 kHz and aligned in time with eye-tracking and photodiode signals sampled at 2 kHz. Stimuli were displayed with Matlab using the Psychtoolbox extension (Brainard 1997; Kleiner *et al.* 2007).

Data preprocessing

We chose to study the MUA because it is a signal that varies little on consecutive days of recording (Dai *et al.* 2019). MUA data have been shown to lead to equivalent results than populations of sorted single unit data when studying dynamics of neural ensembles (Trautmann *et al.* 2019). For the data collected from monkey A, we first extracted MUA activity for each channel. To this end, we band-passed the raw signal between 500 and 6000 Hz using a finite impulse response filter of order 200. We then selected any spike crossing a negative threshold determined independently for each channel. This threshold is given by the formula $-4 \cdot median(|X|)/0.6745$, where X is the filtered signal. Although rare, we excluded events occurring simultaneously (within 0.03 ms) in any two channels because they were likely the result of a correlated artifact during the experiment. The MUA was then aligned to the time of every stimulus onset and averaged for each channel in bins of 50 ms with 25-ms overlap. The average MUA was normalized independently for each channel and each day using the formula $\frac{Y-\bar{Y}}{SD}$, where Y are all the averaged bins and SD the standard deviation of Y.

For the data collected from monkey H, the MUA extraction was similar to monkey A except that the filter used was a forward and reverse Butterworth band-pass filter between 600 and 3000 Hz. The threshold for spike selection was given by the formula $-5 \cdot median(|X|)/0.6745$, where X is the filtered signal. The correlated artifact was removed when more than half of the recording channels detected a spike within 3 ms. The alignment of the MUA to stimulus onset and the normalization of the firing rate was obtained using the same analysis steps as monkey A.

To obtain a finer resolution of the time course of information reaching VLPFC, we also performed the same data preprocessing but using a 50-ms running window with overlapping steps of 5 ms (Supplementary Fig. S4).

Decoding

We trained independent linear decoders using the normalized firing rate pooled from all recording days. To this end, we used multinomial logistic regression. The probability P for an image class Y is given by:

$$P(Y_i = j) = \frac{expexp(\beta_j X_i)}{1 + \sum_k^{K-1} exp(\beta_k X_i)} \quad (1)$$

where β is the coefficient vector and X the vector of firing rates from all channels. The subscript i indicates the trial number and j the image identity from a total of K images. As the probabilities of all K images sum to 1, the chance level equals $1/K$ (1/18 in our study). To estimate the coefficients, we used the LogisticRegression function from the Scikit-learn toolbox in python (Pedregosa *et al.* 2011) with the arguments: fit_intercept = False, solver = 'lbfgs', multi_class='auto', and max_iter = 10 000. We adopted a 10-fold cross-validation strategy.

For Supplementary Fig. S4, we used binomial logistic regression classifiers contrasting every 153 possible pairs of stimuli.

Comparison between VLPFC and PPC

To compare the amount of information linearly decodable from the neural ensembles of VLPFC and PPC, we ensured that the number of visually responsive channels was the same between both areas. To this end, we selected every channel for which the average firing rate in the period 100–200 ms after stimulus onset differed significantly, by an increase or a decrease of firing rate, from the baseline activity (–100 to 50 ms around stimulus onset). In monkey A, this procedure selected 67 channels from the VLPFC array and 58 channels from the PPC array. Classifiers were then trained with a bootstrap strategy, which consisted of randomly choosing 52 visually responsive channels in each array, dropping out 15 channels from the PFC array and 6 from the PPC array. This bootstrap procedure was run for 50 iterations, thus enabling us to estimate a confidence interval of the median accuracy in each area.

In monkey H, the PPC data were collected while the animal was performing the same task as monkey A (Fig. 1a). However, unlike in monkey A, PPC and VLPFC data were recorded in separate sessions. To obtain the fairest comparison between those two datasets collected in different conditions, the number of stimuli presentation used to train the decoders was downsampled from the VLPFC data to match that of the PPC data. Thus, in monkey H, the decoding comparison between VLPFC and PPC is based on 20 visually responsive channels, 2498 trials from the RSVP condition, and 923 trials from the slow-paced/single-stimulus condition.

Statistics

In Fig. 2, the clusters of time points with significant decoding accuracy were revealed by a random-permutations cluster-level statistic corrected for multiple comparisons (Maris and Oostenveld 2007). We first obtained 2000 surrogate classifications by randomly permuting the class labels given by the classifiers. Specifically, for a given stimulus in a given trial, all the predictive probabilities attributed to class #1 by our classifiers could be attributed randomly to any of the 18 other classes. This way, we estimated the likelihood of obtaining by chance an average classification accuracy of a certain magnitude. This resulted in P-values for every time point relative to stimulus onset. We then selected all clusters of points with P-values <0.05 and computed within each cluster the sum of the t-values relative to the surrogate distribution. We then displayed in Fig. 2 all clusters for which the sum of t-values was superior to 99.9% of the sum of t-values obtained from clusters from the surrogate dataset (i.e. $P < 0.001$).

In Fig. 3, single-sample T-tests were conducted to evaluate the likelihood that any stimulus could be decoded by chance. The mean distribution of accuracy averaged independently across all 18 stimuli was tested against 1/18, and the P-values were corrected for multiple comparisons using the Benjamini–Hochberg false discovery rate (FDR) procedure (Benjamini and Hochberg 1995). We also used dependent T-tests corrected for multiple comparisons using FDR in order to evaluate the difference between the decodability in RSVP trials compared to corresponding isolated stimulus trials.

In Fig. 4, we used 2000 random permutations to obtain surrogate average classification accuracies. These surrogate data enabled us to estimate P-values for every combination of test time and training time in the time generalization matrices, with FDR correction. We chose to use the Benjamini–Hochberg FDR rather than cluster-based corrections for multiple comparisons to ensure that the significance of a pixel is solely based on the strength of its corresponding accuracy and not influenced by the accuracy associated with neighboring pixels.

In Supplementary Fig. S4, we obtained P-values of the accuracy by performing a binomial test with a null hypothesis set to a probability of 0.5. P-values were corrected for multiple comparisons with the FDR procedure.

In Supplementary Fig. S5, the *P*-values corresponding to the difference in median accuracy between VLPFC and PPC were calculated by means of 2000 permutations and cluster-corrected for multiple comparisons.

In Supplementary Fig. S6, the mean distribution of accuracy averaged independently across all 18 stimuli was tested against 1/18 using single sample T-tests and the *P*-values were FDR corrected.

Supplementary data

Supplementary data is available at *NCONSC* online.

Data availability

All the data used in this paper are available in Figshare: 10.6084/m9.figshare.14159753.
All analysis scripts are available on the github link: https://github.com/jobellet/fast_and_rich_decoding_in_VLPFC.

Acknowledgements

We thank Marie Bellet for useful feedback concerning decoding analysis and early versions of the manuscript, and Joël Cotton, Jean-Robert Deverre, and the NeuroSpin support teams for their help at various stages of this project.

Funding

This work received funding by INSERM, CEA, Collège de France, the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3), a grant from Templeton World Charity Foundation, Inc (TWCF0562) to Theofanis Panagiotaropoulos and the Bettencourt-Schueller Foundation.

Conflict of interest statement

None declared.

Author contributions

J.B., T.v.K., S.D., and T.I.P. designed the study and wrote the paper. J.B. collected the data from monkey A. M.G. trained the monkeys and collected the prefrontal data from monkey H. A.D. collected the parietal data from monkey H. T.I.P., B.J., A.D., and M.G. implanted the Utah arrays. J.B. analyzed the data.

References

Bahmani H, Qinglin L, Logothetis NK *et al.* Responses of neurons in lateral intraparietal area depend on stimulus-associated reward during binocular flash suppression. *Front Syst Neurosci* 2019;**13**:9.

Bar M, Kassam KS, Ghuman AS *et al.* Top-down facilitation of visual recognition. *Proc Natl Acad Sci* 2006;**103**:449–54.

Baria AT, Maniscalco B, He BJ. Initial-state-dependent, robust, transient neural dynamics encode conscious visual perception. *PLoS Comput Biol* 2017;**13**:1–29.

Benevento LA, Fallon JH. The ascending projections of the superior colliculus in the rhesus monkey (Macaca mulatta). *J Comp Neurol* 1975;**160**:339–61.

Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc* 1995;**57**:289–300.

Blanke O, Landis T, Seeck M. Electrical cortical stimulation of the human prefrontal cortex evokes complex visual hallucinations. *Epilepsy Behav* 2000;**1**:356–61.

Block N. Perceptual consciousness overflows cognitive access. *Trends Cogn Sci* 2011;**15**:567–75.

Block N. Finessing the bored monkey problem. *Trends Cogn Sci* 2020;**24**:167–8.

Boly M, Massimini M, Tsuchiya N *et al.* Are the neural correlates of consciousness in the front or in the back of the cerebral cortex? Clinical and neuroimaging evidence. *J Neurosci* 2017;**37**:9603–13.

Brainard DH. The psychophysics toolbox. *Spat Vis* 1997;**10**:433–6.

Broadbent DE, Broadbent MHP. From detection to identification: response to multiple targets in rapid serial visual presentation. *Percept Psychophys* 1987;**42**:105–13.

Brown R, Lau H, Ledoux JE. Understanding the higher-order approach to consciousness. *Trends Cogn Sci* 2019;**23**:754–68.

Chen CY, Sonnenberg L, Weller S *et al.* Spatial frequency sensitivity in macaque midbrain. *Nat Commun* 2018;**9**:1–13.

Constantinidis C, Franowicz MN, Goldman-Rakic PS. The sensory nature of mnemonic representation in the primate prefrontal cortex. *Nat Neurosci* 2001;**4**:311–6.

Dai J, Zhang P, Sun H *et al.* Reliability of motor and sensory neural decoding by threshold crossings for intracortical brain–machine interface. *J Neural Eng* 2019;**16**:036011.

Dehaene S. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. New York: Viking Penguin, 2014.

Dehaene S, Changeux J-P. Experimental and theoretical approaches to conscious processing. *Neuron* 2011;**70**:200–27.

Dehaene S, Naccache L, Cohen L *et al.* Cerebral mechanisms of word masking and unconscious repetition priming. *Nat Neurosci* 2001;**4**:752–8.

Dehaene S, Sergent C, Changeux JP. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proc Natl Acad Sci* 2003;**100**:8520–5.

Del Cul A, Dehaene S, Reyes P *et al.* Causal role of prefrontal cortex in the threshold for access to consciousness. *Brain* 2009;**132**:2531–40.

Dwarakanath A, Kapoor V, Werner J *et al.* Prefrontal state fluctuations control access to consciousness. *bioRxiv* 2020. 10.1101/2020.01.29.924928.

Farooqui AA, Manly T. When attended and conscious perception deactivates fronto-parietal regions. *Cortex* 2018;**107**:166–79.

Fox KC, Shi L, Baek S *et al.* Intrinsic network architecture predicts the effects elicited by intracranial electrical stimulation of the human brain. *Nature Human Behaviour* 2020;**4**:1039–52.

Frassle S, Sommer J, Jansen A *et al.* Binocular rivalry: frontal activity relates to introspection and action but not to perception. *J Neurosci* 2014;**34**:1738–47Av.

Freedman DJ, Riesenhuber M, Poggio T *et al.* Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 2001;**291**:312–6.

Fuster JM. *The Prefrontal Cortex*. New York: Raven Press, 1989.

Gaillard R, Dehaene S, Adam C *et al.* Converging intracranial markers of conscious access. *PLoS Biol* 2009;**7**:e1000061.

Gelbard-Sagiv H, Mudrik L, Hill MR *et al.* Human single neuron activity precedes emergence of conscious perception. *Nat Commun* 2018;**9**:1–13.

Goldman-Rakic PS, Porrino LJ. The primate mediodorsal (MD) nucleus and its projection to the frontal lobe. *J Comp Neurol* 1985;**242**:535–60.

Histed MH, Bonin V, Reid RC. Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. *Neuron* 2009;**63**:508–22.

Hong H, Yamins DL, Majaj NJ *et al.* Explicit information for category-orthogonal object properties increases along the ventral stream. *Nat Neurosci* 2016;**19**:613.

Hung CP, Kreiman G, Poggio T *et al.* Fast readout of object identity from macaque inferior temporal cortex. *Science* 2005;**310**:863–6.

Kapoor V, Besserve M, Logothetis NK *et al.* Parallel and functionally segregated processing of task phase and conscious content in the prefrontal cortex. *Commun Biol* 2018;**1**:1–12.

Kapoor V, Dwarakanath A, Safavi S *et al.* Decoding the contents of consciousness from prefrontal ensembles. *bioRxiv* 2020. 10.1101/2020.01.28.921841 11.

Kar K, DiCarlo JJ. Fast recurrent processing via ventrolateral prefrontal cortex is needed by the primate ventral stream for robust core visual object recognition. *Neuron* 2021;**109**.

King J-R, Dehaene S. Characterizing the dynamics of mental representations: the temporal generalization method. *Trends Cogn Sci* 2014;**18**:203–10.

King J-R, Gramfort A, Schurger A *et al.* Two distinct dynamic modes subtend the detection of unexpected sounds. *PLoS One* 2014;**9**:e85791.

King JR, Pescetelli N, Dehaene S. Brain mechanisms underlying the brief maintenance of seen and unseen sensory information. *Neuron* 2016;**92**:1122–34.

Kleiner M, Brainard D, Pelli D. What's new in Psychtoolbox-3? 2007.

Koch C. The quest for consciousness a neurobiological approach. 2004.

Koch C, Massimini M, Boly M *et al.* Neural correlates of consciousness: progress and problems. *Nat Rev Neurosci* 2016;**17**:307–21.

Kovacs G, Vogels R, Orban GA. Cortical correlate of pattern backward masking. *Proc Natl Acad Sci* 1995;**92**:5587–91.

Kreiman G, Fried I, Koch C. Single-neuron correlates of subjective vision in the human medial temporal lobe. *Proc Natl Acad Sci U S A* 2002;**99**:8378–83.

Kriegeskorte N, Mur M, Bandettini PA. Representational similarity analysis-connecting the branches of systems neuroscience. *Front Syst Neurosci* 2008;**2**:4.

Lau H, Rosenthal D. Empirical support for higher-order theories of conscious awareness. *Trends Cogn Sci* 2011;**15**:365–73.

Lau HC, Passingham RE. Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proc Natl Acad Sci* 2006;**103**:18763–8.

Lawrence DH. Two studies of visual search for word targets with controlled rates of presentation. *Percept Psychophys* 1971;**10**:85–9.

Levinson M, Podvalny E, Baete SH *et al.* Cortical and subcortical signatures of conscious object recognition. *Nat Commun* 2021;**12**:1–16.

Li Q, Hill Z, He BJ. Spatiotemporal dissociation of brain activity underlying subjective awareness, objective performance and confidence. *J Neurosci* 2014;**34**:4382–95.

Libedinsky C, Livingstone M. Role of prefrontal cortex in conscious visual perception. *J Neurosci* 2011;**31**:64–9.

Liu A, Friedman D, Barron DS *et al.* Forced conceptual thought induced by electrical stimulation of the left prefrontal gyrus involves widespread neural networks. *Epilepsy Behav* 2020;**104**:106644.

Logothetis NK. Single units and conscious vision. *Philos Trans R Soc B* 1998;**353**:1801–18.

Lumer ED, Rees G. Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc Natl Acad Sci* 1999;**96**:1669–73.

Majaj NJ, Hong H, Solomon EA *et al.* Simple learned weighted sums of inferior temporal neuronal firing rates accurately predict human core object recognition performance. *J Neurosci* 2015;**35**:13402–18.

Maris E, Oostenveld R. Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* 2007;**164**:177–90.

Marti S, Dehaene S. Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nat Commun* 2017;**8**.

Mashour GA, Roelfsema P, Changeux JP *et al.* Conscious processing and the global neuronal workspace hypothesis. *Neuron* 2020;**105**:776–98.

Maunsell JHR, Gibson JR. Visual response latencies in striate cortex of the macaque monkey. *J Neurophysiol* 1992;**68**:1332–44.

Mei N, Santana R, Soto D. Informative neural representations of unseen objects during higher-order processing in human brains and deep artificial networks. *bioRxiv* 2021. 2021.01.12.426428.

Mishkin M, Ungerleider LG, Macko KA. Object vision and spatial vision: two cortical pathways. *Trends Neurosci* 1983;**6**:414–7.

Naber M, Frässle S, Einhäuser W *et al.* Perceptual rivalry: reflexes reveal the gradual nature of visual awareness. *PLoS One* 2011;**6**:e20910.

Naccache L, Dehaene S. Reportability and illusions of phenomenality in the light of the global neuronal workspace model. *Behav Brain Sci* 2007;**30**:518–20.

Nowak LG, Munk MHJ, Girard P *et al.* Visual latencies in areas V1 and V2 of the macaque monkey. *Vis Neurosci* 1995;**12**:371–84.

Ó Scalaidhe SP, Wilson FAW, Goldman-Rakic PS. Face-selective neurons during passive viewing and working memory performance of rhesus monkeys: evidence for intrinsic specialization of neuronal coding. *Cereb Cortex* 1999;**9**:459–75.

Olivers CNLL, Van Der Stigchel S, Hulleman J *et al.* Spreading the sparing: against a limited-capacity account of the attentional blink. *Psychol Res* 2007;**71**:126–39.

Oram MW, Xiao D, Dritschel B *et al.* The temporal resolution of neural codes: does response latency have a unique role? *Philos Trans R Soc B Biol Sci.* 2002;**357**:987–1001.

Panagiotaropoulos TI, Deco G, Kapoor V *et al.* Neuronal discharges and gamma oscillations explicitly reflect visual consciousness in the lateral prefrontal cortex. *Neuron* 2012;**74**:924–35.

Panagiotaropoulos TI, Dwarakanath A, Kapoor V. Prefrontal cortex and consciousness: beware of the signals. *Trends Cogn Sci* 2020;**24**:343–4.

Pedregosa F, Varoquaux G, Gramfort A *et al.* Scikit-learn: machine learning in python. *J Mach Learn Res* 2011;**12**:2825–30.

Petrides M, Pandya DN. Distinct parietal and temporal pathways to the homologues of Broca's area in the monkey. *PLoS Biol* 2009;**7**:e1000170.

Popa I, Donos C, Barborica A *et al.* Intrusive thoughts elicited by direct electrical stimulation during stereo-electroencephalography. *Front Neurol* 2016;**7**:1–6.

Potter MC, Kroll JF, Yachzel B *et al.* Pictures in sentences: understanding without words. *J Exp Psychol Gen* 1986;**115**:281.

Pouget P, Emeric EE, Stuphorn V *et al.* Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *J Neurophysiol* 2005;**94**:2086–92.

Quiroga RQ, Mukamel R, Isham EA *et al.* Human single-neuron responses at the threshold of conscious recognition. *Proc Natl Acad Sci* 2008;**105**:3599–604.

Raccah O, Block N, Fox KC. Does the prefrontal cortex play an essential role in consciousness? Insights from intracranial electrical stimulation of the human brain. *J Neurosci* 2021;**41**:2076–87.

Raymond JE, Shapiro KL, Arnell KM *et al.* Temporary suppression of visual processing in an RSVP task: an attentional blink? *J Exp Psychol Hum Percept Perform* 1992;**18**:849.

Reber TP, Faber J, Niediek J *et al.* Single-neuron correlates of conscious perception in the human medial temporal lobe. *Curr Biol* 2017;**27**:1–8.

Riley MR, Qi XL, Constantinidis C. Functional specialization of areas along the anterior-posterior axis of the primate prefrontal cortex. *Cereb Cortex* 2017;**27**:3683–97.

Rounis E, Maniscalco B, Rothwell JC *et al.* Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cogn Neurosci* 2010;**1**:165–75.

Schmolesky MT, Wang Y, Hanes D *et al.* Signal timing across the macaque visual system. *J Neurophysiol* 1998;**79**:3272–8.

Schurger A, Sarigiannidis I, Naccache L *et al.* Cortical activity is more stable when sensory stimuli are consciously perceived. *Proc Natl Acad Sci U S A* 2015a;**112**:E2083–92.

Schurger A, Sarigiannidis I, Naccache L *et al.* Cortical activity is more stable when sensory stimuli are consciously perceived. *Proc Natl Acad Sci* 2015b;**112**:E2083–92.

Self MW, van Kerkoerle T, Supèr H *et al.* Distinct roles of the cortical layers of area V1 in figure-ground segregation. *Curr Biol* 2013;**23**:2121–9.

Sergent C, Baillet S, Dehaene S. Timing of the brain events underlying access to consciousness during the attentional blink. *Nat Neurosci* 2005;**8**:1391–400.

Sergent C, Wyart V, Babo-Rebelo M *et al.* Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Curr Biol* 2013;**23**:150–5.

Singer W. Time as coding space? *Curr Opin Neurobiol* 1999;**9**:189–94.

Sommer MA, Wurtz RH. What the brain stem tells the frontal cortex. I. Oculomotor signals sent from superior colliculus to frontal eye field via mediodorsal thalamus. *J Neurophysiol* 2004;**91**:1381–402.

Sperling G. The information available in brief visual presentations. *Psychol Monogr Gen Appl* 1960;**74**:1–29.

Thibault L, Van Den Berg R, Cavanagh P *et al.* Retrospective attention gates discrete conscious access to past sensory stimuli. *PLoS One* 2016;**11**:1–13.

Thompson KG, Schall JD. The detection of visual signals by macaque frontal eye field during masking. *Nat Neurosci* 1999;**2**:283.

Tononi G, Boly M, Massimini M *et al.* Integrated information theory: from consciousness to its physical substrate. *Nat Rev Neurosci* 2016;**17**:450–61.

Trautmann EM, Stavisky SD, Lahiri S *et al.* Accurate estimation of neural population dynamics without spike sorting. *Neuron* 2019;**103**:292–308.

Van Gaal S, Ridderinkhof KR, Scholte HS *et al.* Unconscious activation of the prefrontal no-go network. *J Neurosci* 2010;**30**:4143–50.

Van Vugt B, Dagnino B, Vartak D *et al.* The threshold for conscious report: signal loss and response bias in visual and frontal cortex. *Science* 2018;**360**:537–42.

Vignal JP, Chauvel P, Halgren E. Localised face processing by the human prefrontal cortex: stimulation-evoked hallucinations of faces. *Cogn Neuropsychol* 2000;**17**:281–91.

Wolfe JM. Reversing ocular dominance and suppression in a single flash. *Vision Res* 1984;**24**:471–8.

Xia Y, Morimoto Y, Noguchi Y. Retrospective triggering of conscious perception by an interstimulus interaction. *J Vis* 2016;**16**:1–8.

Zamarashkina P, Popovkina DV, Pasupathy A. Timing of response onset and offset in macaque V4: stimulus and task dependence. *J Neurophysiol* 2020;**123**:2311–25.