

NEUROSCIENCE

The threshold for conscious report: Signal loss and response bias in visual and frontal cortex

Bram van Vugt,^{1*} Bruno Dagnino,^{1*} Devavrat Vartak,^{1*} Houman Safaai,^{2,3,†} Stefano Panzeri,³ Stanislas Dehaene,^{4,5} Pieter R. Roelfsema^{1,6,7,†}

Why are some visual stimuli consciously detected, whereas others remain subliminal? We investigated the fate of weak visual stimuli in the visual and frontal cortex of awake monkeys trained to report stimulus presence. Reported stimuli were associated with strong sustained activity in the frontal cortex, and frontal activity was weaker and quickly decayed for unreported stimuli. Information about weak stimuli could be lost at successive stages en route from the visual to the frontal cortex, and these propagation failures were confirmed through microstimulation of area V1. Fluctuations in response bias and sensitivity during perception of identical stimuli were traced back to prestimulus brain-state markers. A model in which stimuli become consciously reportable when they elicit a nonlinear ignition process in higher cortical areas explained our results.

Understanding how conscious perception arises in the brain is a major challenge for neuroscience. Experimentally, one approach consists of comparing the neuronal activity evoked by identical weak stimuli, which are sometimes perceived and sometimes remain subliminal. Previous experiments have shown that subliminal stimuli elicit considerable activity in many brain areas, including the prefrontal cortex (1), raising the question of why this activity is insufficient for conscious report (2).

The classical model that describes how weak stimuli are perceived or missed is signal detection theory (SDT) (3). It posits that stimuli elicit a stochastic signal, which has to reach a threshold for perception (Fig. 1A). Stimuli that fail to reach the threshold are missed. In the absence of a stimulus, the signal usually stays below the threshold (correct rejection) but may cross the threshold on occasion, giving rise to a false alarm. According to SDT, a higher threshold decreases the number of false alarms but also increases the number of misses.

SDT does not specify the brain processes that determine the variability of the stimulus-induced

signal nor the mechanism that determines the threshold. By contrast, global neuronal workspace theory (GNWT) (4) proposes that stimuli reach awareness by propagating to the higher levels of the cerebral cortex, where they can lead to “ignition,” a nonlinear event that causes information about a brief stimulus to become sustained and broadcasted back through recurrent interactions between many brain areas (Fig. 1B) (5). According to GNWT, there are two reasons why a stimulus may fail to become consciously accessible. First, the propagation of activity to higher levels may be too weak. Second, global ignition may fail—for example, if the system is refractory because another stimulus caused ignition or if attention is diverted (1). Combining insights from SDT and GNWT (6, 7), we hypothesized that the SDT threshold might equal the amount of neural activity required for ignition. Furthermore, the stochasticity in signal strength might relate to variations in the propagation of activity from lower to higher cortical levels, possibly caused by fluctuations in prestimulus brain state.

We trained monkeys to detect low-contrast stimuli and recorded multiunit activity (MUA) in areas V1 and V4 of the visual cortex and in the dorsolateral prefrontal cortex (dlPFC) in order to examine the fate of identical subliminal and supraliminal stimuli. We asked the following questions: (i) Where in the visual hierarchy do subliminal signals get lost? (ii) Which neuronal mechanisms underlie the threshold for reporting a stimulus? (iii) What are the internal sources of fluctuations that allow a fixed stimulus to either cross or fail to cross the threshold?

The monkeys directed their gaze to a fixation point, and on half of the trials, we presented a 2° low-contrast circle as stimulus in the neurons' receptive field (RF) for 50 ms (Fig. 1C). After a delay of 450 ms, introduced to prevent reflexive eye movements (8), the monkey reported

the stimulus by making a saccade to its previous location. In the absence of a stimulus, the monkey made a saccade to another, smaller gray circle (the reject dot). Accuracy on such stimulus-absent trials was high (~5 to 10% of false alarms). We adjusted the contrast on stimulus-present trials close to the threshold of perception, at an accuracy of ~80% (supplementary materials, materials and methods). The contrast threshold (θ_{High} ; accuracy of 80%) varied with stimulus eccentricity between 2.5 and 7% (fig. S1, A to C). To examine perception of very weak stimuli, we also defined a second threshold, θ_{Low} , associated with an accuracy of 40% and categorized stimulus strength into three categories: easy (contrast > θ_{High}), intermediate (θ_{Low} < contrast < θ_{High}), and difficult (contrast < θ_{Low}) (Fig. 1D). We normalized the neuronal responses to the activity elicited by a high-contrast stimulus (supplementary materials, materials and methods).

Stimuli with higher contrasts elicited more activity than did stimuli with lower contrasts (time window 0 to 300 ms after stimulus onset, *t* tests, all $P < 10^{-3}$) (Fig. 1E). Within each strength category, we compared neuronal activity between hits and misses with identical stimulus contrast (supplementary materials, materials and methods). Hits elicited stronger activity in V1, V4, and the dlPFC than did misses, at every difficulty level (window 0 to 300 ms, paired *t* tests, all areas and categories $P < 0.01$) (Fig. 2A and fig. S2, example recording sites). Hence, during misses information is lost during the propagation of visual information to higher cortical areas. In the dlPFC and, to a lesser extent, areas V4 and V1, the extra neuronal activity for hits was maintained until the saccade (time window 300 to 500 ms, all areas and categories $P < 0.05$) (Fig. 2 and fig. S3).

To determine the locus of the information loss, we computed the miss fraction—the percentage of activity remaining for nonreported stimuli ($\text{Activity}_{\text{Miss}}/\text{Activity}_{\text{Hit}} \times 100\%$, time window 0 to 300 ms) (Fig. 2B). For the difficult stimuli, the miss fraction was 46% in area V1 and 14% in V4, implying a substantial loss of activity before V1 and a further loss between V1 and V4 [significant difference between miss fractions in V1 and V4, $t_{46} = 2.5$, $P < 0.01$, number of recording sites in V1 (N_{V1}) = 27, N_{V4} = 26]. For the intermediate and easy stimuli, the miss fractions were much higher (around 60 and 80%, respectively) and did not differ significantly between V1 and V4 (for both, $P > 0.05$), indicating substantial propagation of neural activity for misses. Now, however, extra activity was lost between V4 and the dlPFC, both for the intermediate (62 versus 33%, $t_{50} = 3.4$, $P < 0.05$, $N_{V4} = 33$, $N_{\text{dlPFC}} = 19$) and easy stimuli (83 versus 22%, $t_{49} = 4.0$, $P < 10^{-3}$, $N_{V4} = 34$, $N_{\text{dlPFC}} = 17$). Thus, subliminal stimuli can be lost at different stages of bottom-up signal propagation, depending on their strength.

The activity levels in V1 and V4 on easy miss trials (Fig. 2A, right, red curves) were at least as strong as those on difficult hit trials (Fig. 2A, left, green curves). Hence, the neuronal activity level in these areas does not fully predict stimulus detection. In the dlPFC, however, the activity

¹Department of Vision and Cognition, Netherlands Institute for Neuroscience, Meibergdreef 47, 1105 BA Amsterdam, Netherlands. ²Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA. ³Neural Computation Laboratory, Istituto Italiano di Tecnologia, 38068 Rovereto, Italy. ⁴Cognitive Neuroimaging Unit, Commissariat à l'Énergie Atomique et aux Énergies Alternatives, Direction des Sciences du Vivant/Institut d'Imagerie Biomédicale, INSERM, NeuroSpin center, Université Paris-Sud and Université Paris-Saclay, 91191 Gif-sur-Yvette, France. ⁵Collège de France, 75005 Paris, France. ⁶Department of Integrative Neurophysiology, Center for Neurogenetics and Cognitive Research, Vrije Universiteit, Amsterdam, Netherlands. ⁷Department of Psychiatry, Academic Medical Center, Amsterdam, Netherlands.

*These authors contributed equally to this work.

†Corresponding author. Email: p.roelfsema@nin.knaw.nl (P.R.R.); houman_safaai@hms.harvard.edu (H.S.)

elicited by hits was stronger at all contrast levels than that elicited by misses, implying that these neurons predicted conscious report in a more categorical manner, which is presumably related to the planning of an eye movement toward the

neurons' RF. These results were consistent between monkeys (fig. S4).

In the absence of a stimulus, neuronal activity also differed between false alarms and correct rejections (Fig. 2A). In the dlPFC and area V4,

neuronal activity was higher on false-alarm trials (time window 200 ms before saccade; dlPFC, $t_{27} = 4.5, P < 10^{-3}$; V4, $t_{36} = 4.8, P < 10^{-3}$), with a trend in the same direction in area V1 ($t_{34} = 1.8, P = 0.07$). The extra activity on false-alarm trials

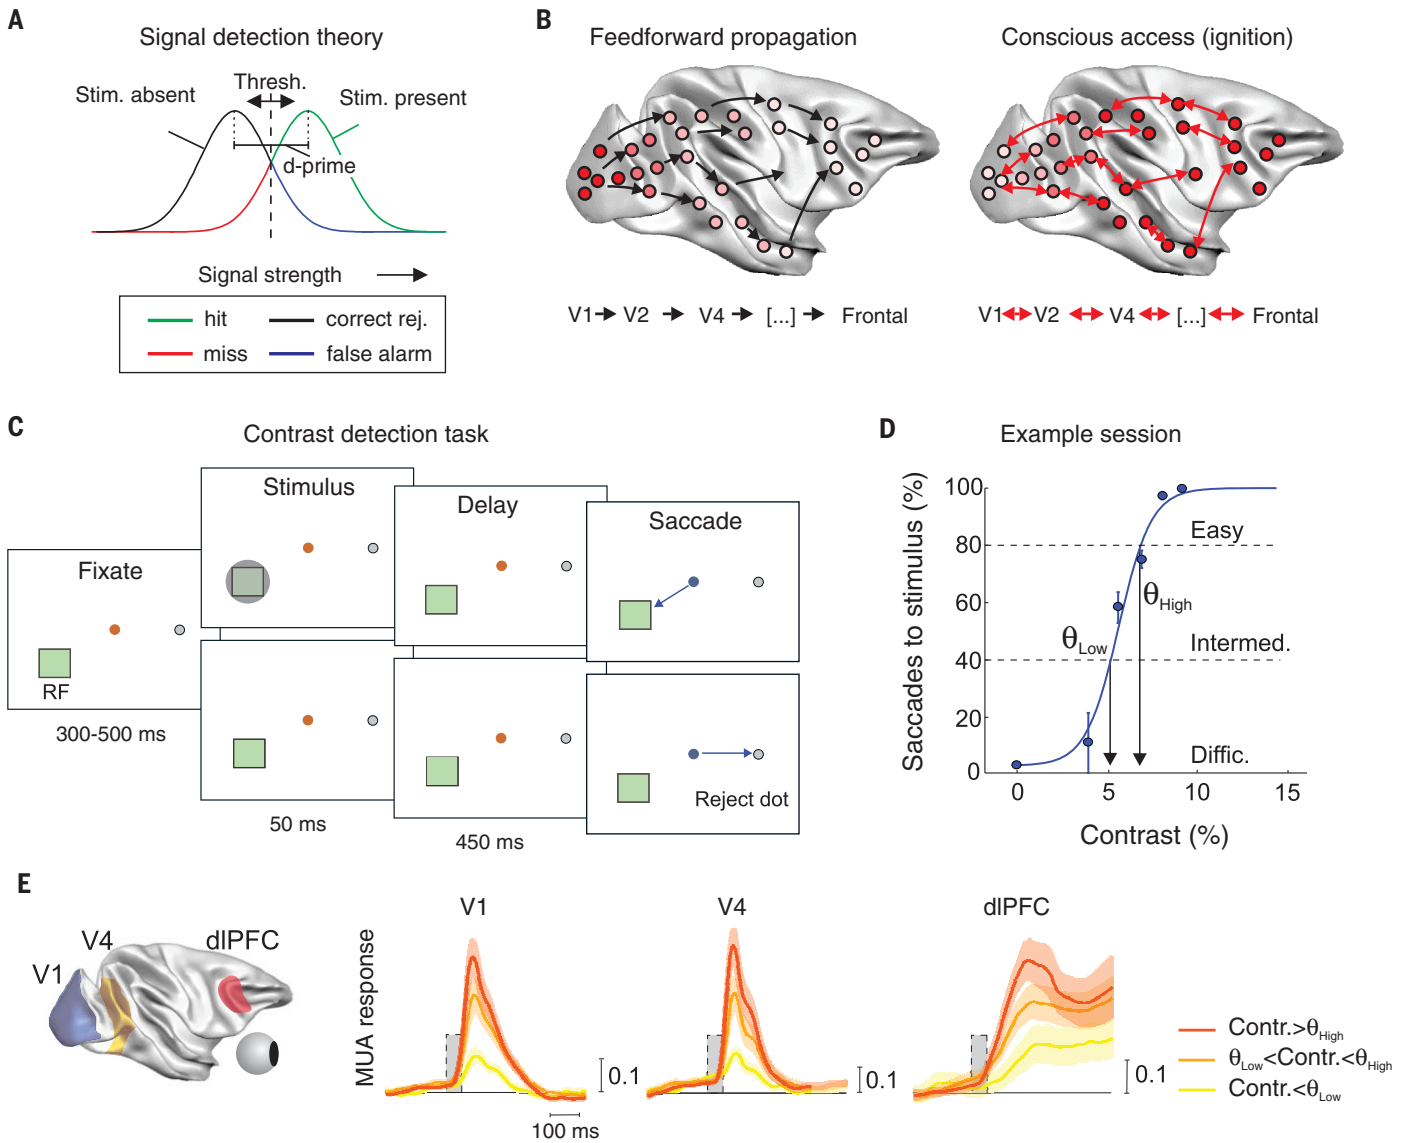


Fig. 1. Visual perception at low contrast. (A) SDT holds that stimuli need to cause an internal signal strength larger than a threshold to be perceived and reported. The internal signal strength is stochastic, and false alarms result if the signal crosses the threshold in the absence of a stimulus. In SDT, the subject's bias can be changed by shifting the threshold. A high threshold makes the subject more conservative, and a low threshold makes the subject less conservative, in his "target present" judgments. d -prime is a measure of the subject's sensitivity. It corresponds to the distance between the distributions of signal strength for stimulus-present and stimulus-absent trials, measured in units of the standard deviation. (B) According to GNWT, sensory activity first needs to be propagated to the higher stages of the cortical hierarchy. If it is strong enough, it can access awareness by causing "global ignition," a process that enables maintenance and sharing of information about the stimulus between cortical processors, manifested by an increase in activity. (C) Contrast detection task. On half of the trials, a low-contrast 2°

circle stimulus was presented for 50 ms after 300 to 500 ms of fixation. On the other half of the trials, there was no stimulus. After a delay of 450 ms, the monkey reported the stimulus by making a saccade to its previous location and the absence of the stimulus by making a saccade to a smaller gray circle (reject dot). (D) Psychometric detection curve for an example session in monkey B. We determined two thresholds, θ_{Low} (accuracy of 40%) and θ_{High} (accuracy of 80%), based on the psychometric function. (E) MUA elicited in V1, V4, and the dlPFC by means of easy (contrast $> \theta_{High}$), intermediate ($\theta_{Low} < \text{contrast} < \theta_{High}$), and difficult stimuli (contrast $< \theta_{Low}$). The activity was averaged across all recording sites per brain region. The numbers of sites— N_{Easy} , $N_{Intermed.}$, and $N_{Diff.}$ —in area V1 were 33, 25, and 23; in V4 were 36, 34, and 29; and in the dlPFC were 17, 20, and 14, respectively. The vertical scale bar is in units of normalized activity because the MUA at all sites was normalized to the response elicited by a high-contrast stimulus (supplementary materials, materials and methods, and fig. S8). The shaded areas represent \pm SEM.

was already present in a 300-ms time window before stimulus onset in all areas (for all, $P < 0.05$), suggesting that increased cortical excitation is a prestimulus brain-state marker that predicts false alarms. We therefore computed the area under the receiver-operating curve (AUROC). An AUROC of 0.5 indicates no predictive power, and a value of 1 indicates perfect prediction. We obtained AUROCs of 0.52, 0.52, and 0.57 in areas V1 and V4 and the dIPFC, respectively (300 ms time window) (Fig. 3A, striped green MUA bars). The AUROC values around the time that a stimulus could have been presented, known as “choice probabilities,” were higher (striped black bars). On stimulus-present trials, we focused on trials of intermediate difficulty level, for which we had enough trials, and obtained relatively low prestimulus AUROCs (V1, 0.52; V4, 0.50; and dIPFC, 0.51). Again, choice probabilities after stimulus presentation were higher (0 to 300 ms after stimulus onset; V1, 0.71; V4, 0.71; and dIPFC, 0.74; for all, $P < 0.01$) (Fig. 3A, solid black bars).

We next asked whether fluctuations in the prestimulus brain state were related to variations in stimulus detection. We therefore also evaluated several markers that might predict perceptual report, including the diameter of the pupil (Pu); its time derivative (Δ Pu) (9); the power in the α , β , and low or high γ bands of the local field potential (10, 11); and the time that the monkeys took to initiate a new trial, which is informative about their motivation. When considered individually, all markers gave weak predictions (Fig. 3A). We linearly combined prestimulus brain-state measures into a joint measure J (supplementary materials, materials and methods), which predicted perceptual outcome with an accuracy close to 60% (V1, 0.59; V4, 0.58; and dIPFC, 0.58; for all, $P < 0.001$). To examine the influence of J on neuronal activity, we selected all trials from the highest and lowest quintile of the J distribution across trials. A higher value for J was associated with higher prestimulus activity and a stronger visual response in all three areas (Fig. 3C). High J values also caused a slight increase in the false-alarm rate (Fig. 3B).

SDT distinguishes between fluctuations in signal strength and response bias. The bias determines the false alarm rate, and we computed another combination of prestimulus brain-state parameters, bias (B), which predicted false alarms with AUROC values around 0.6 (V1, 0.60; V4, 0.62; and dIPFC, 0.61; for all, $P < 0.01$) (Fig. 3, A and B). Higher B values were associated with extra prestimulus firing in all three areas in stimulus-present (not used to define B) (Fig. 3D) and stimulus-absent trials (fig. S5).

In SDT, accuracy also depends on d -prime, the distance between the stimulus-present and -absent distributions of signal strength (Fig. 1A). Our comparison between hit and miss trials (Fig. 2) suggested that there was variability in the propagation of neuronal activity to higher cortical levels. However, J did not have an isolated effect on d -prime because it also influenced the false-alarm rate. We therefore devised a third linear combination of prestimulus brain-state measures to index sen-

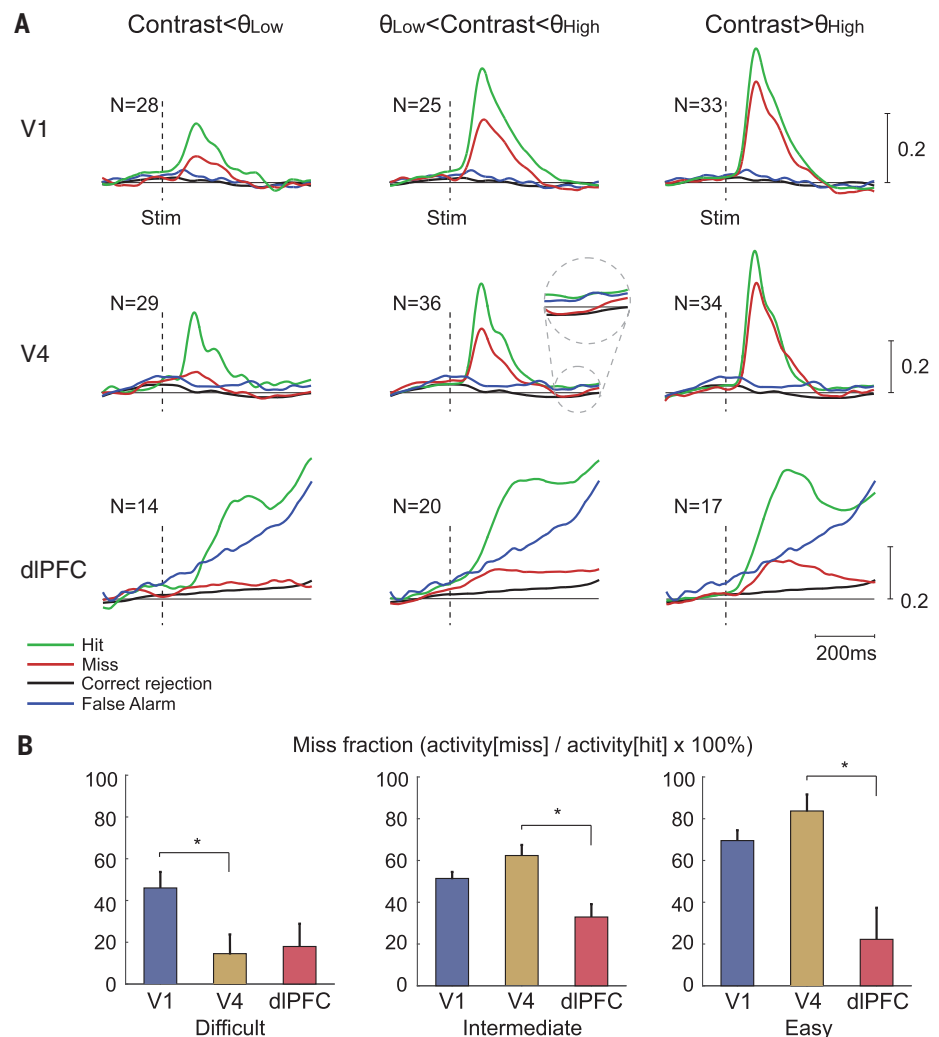


Fig. 2. Activity in areas V1 and V4 and the dIPFC in the contrast-detection task. (A) Activity averaged across all recording sites in (top) V1, (middle) V4, and (bottom) dIPFC by contrasts lower than θ_{Low} (difficult; left), between θ_{Low} and θ_{High} (intermediate; middle), and higher than θ_{High} (easy; right) for contrast-matched hits (green curves) and misses (red curves). The black curves indicate average activity on trials in which the monkeys correctly reported the absence of a stimulus, and the blue curves indicate activity on trials with false alarms. (Inset) The influence of choice on late V4 activity in one of the conditions. (B) Miss fraction ($Activity_{Miss}/Activity_{Hit} \times 100\%$) in V1 (blue bars), V4 (yellow bars), and the dIPFC (red bars) for the different stimulus categories (time window, 0 to 300 ms after stimulus onset).

sitivity (S), designed to discriminate hits from misses without influencing the false-alarm rate (Fig. 3, A and B, and supplementary materials, materials and methods). A high S value increased visually driven activity, especially in the higher areas, in accordance with an influence on the efficiency of activity propagation to higher levels (Fig. 3E). Hence, we identified separable influences of prestimulus brain state on the subject’s response bias and sensitivity. Bias B relates to an increase in ongoing activity, whereas sensitivity S relates to an increase in the efficiency of signal propagation (Fig. 3F).

We interpreted our results in terms of variation in signal propagation from area V1 to V4 and then onward to the PFC, but visual informa-

tion can reach higher areas through multiple routes, some of which bypass area V1 (12). To specifically demonstrate the role of V1-to-V4 propagation and its failures around the threshold of perception, we activated area V1 with electrical microstimulation while recording from V4. The monkeys reported a phosphene—an illusory light percept at the RF of the stimulated neurons (13)—elicited in V1 with five pulses (200 Hz) while we varied stimulation strength (Fig. 4A). We determined two thresholds for the stimulation current (Fig. 4B; the distribution of θ_{High} is shown in fig. S1D) and recorded MUA in V4 from neurons with RFs that overlapped with those of the stimulated V1 neurons (Fig. 4C, right) for a total of 84 V1-V4 pairs (58 in monkey B and 26 in monkey C).

Consistent with the propagation hypothesis, V1 microstimulation elicited V4 activity, with a temporal profile that resembled the response elicited by a visual stimulus, and higher V1 currents increased the V4 response (for all, $P < 0.05$) (Fig. 4C and fig. S6). Furthermore, V4 activity was larger on hit trials than on miss trials at all current strengths (time window from 0 to 150 ms after stimulus onset; paired t test, all $P_s < 10^{-6}$) (Fig. 4D and fig. S6C). Hence, the efficiency of activity propagation from V1 to V4 predicts perceptual report. The miss fraction increased from 16% in difficult trials to 42% in intermediate trials and to 83% in easy trials (Fig. 4E) (t tests; for all, $P_s < 10^{-3}$; $N_{Low} =$

47, $N_{Intermediate} = 41$, $N_{High} = 67$), implying that for the stronger electrical stimuli, information could still be lost at processing levels higher than that of area V4. Indeed, V4 activity on easy miss trials was stronger than on difficult hit trials, confirming that although correlated with the hit probability, the amplitude of the V4 response does not fully predict whether a stimulus will enable conscious report.

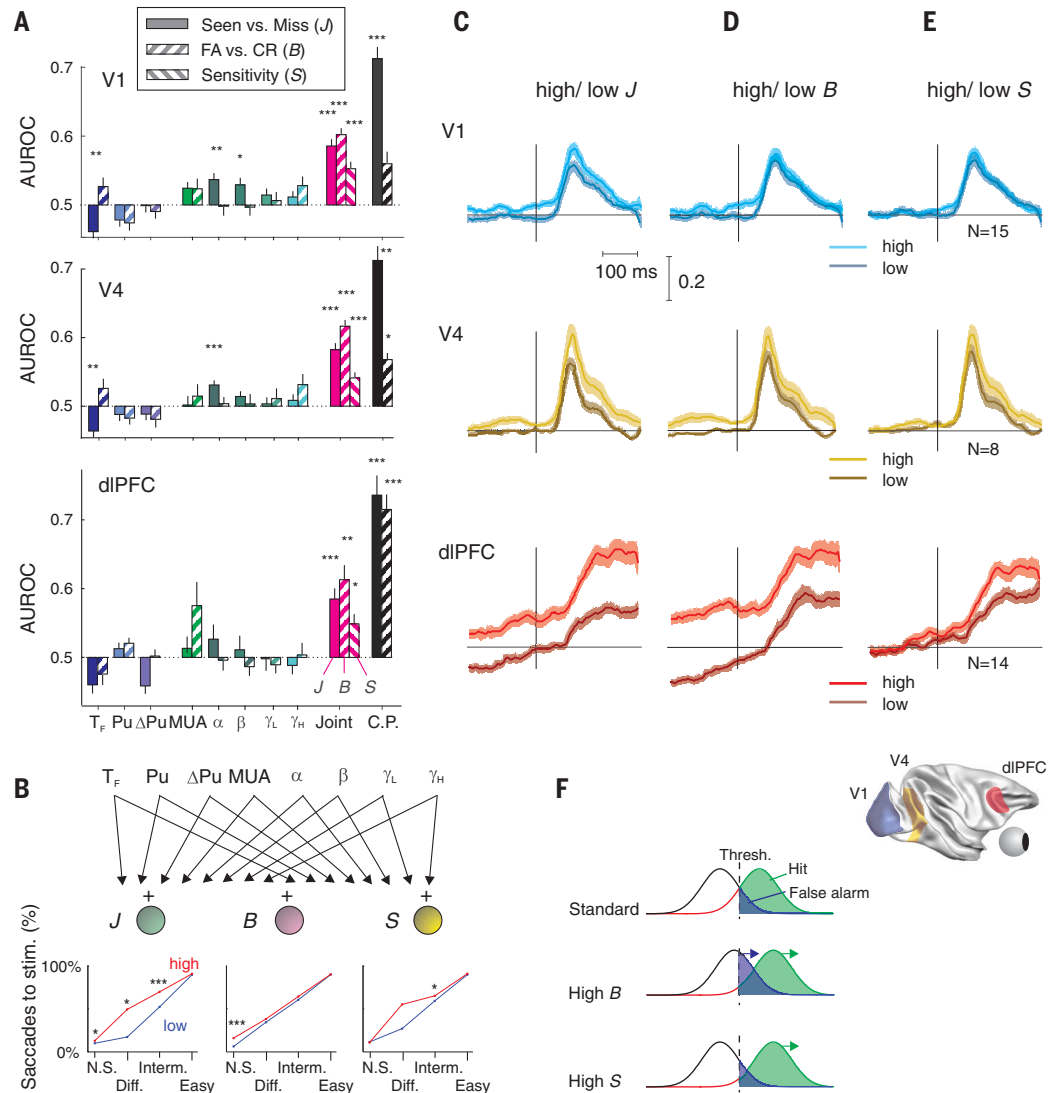
Can a simple mathematical model of hierarchically arranged areas reproduce these findings? The model's architecture was based on previous modeling studies (14) and contained the lateral geniculate nucleus and four hierarchically arranged cortical areas (Fig. 4F). We represented

the population of neurons in each area with a single, stochastic Ornstein-Uhlenbeck process so that only five variables described the evolution of the network state during simulated trials, one for each brain region. The model contained feedforward connections, self-connections within the areas, and feedback connections. The reciprocal connections between the parietal and frontal cortex were relatively strong so that activity exceeding a threshold in these areas became self-sustained (Fig. 4I, ignition, and supplementary materials). The model produced a realistic psychometric function, with increased accuracy for higher contrasts (Fig. 4G). The activity of the model units was

Fig. 3. Influence of prestimulus brain state on neuronal activity and choice.

(A) Behavioral and neurophysiological markers of prestimulus brain state that predict the animal's choice. Predictive value was quantified as the AUROC. Results of experiments in which we recorded activity in areas V1 and V4 and the dIPFC are shown at top, middle, and bottom, respectively. Positive AUROCs indicate that a higher value of the marker predicts a higher probability of hits or false alarms. T_f , time between the appearance of the fixation point and the moment that the monkey directed gaze to the fixation point; Pu, pupil diameter; Δ Pu, change in pupil diameter; MUA, prestimulus MUA; α , power from 5 to 15 Hz; β , 15 to 25 Hz; γ_L , 25 to 40 Hz; γ_H , 40 to 80 Hz; Joint, combination of markers best distinguishing between hits and misses (J), correct rejections and false alarms (B), and a measure that discriminates between hits and misses with a minimal influence on the false alarm rate (S) [schematic in (B)]; C.P., choice probability based on MUA in the stimulus-presentation time window. Solid bars indicate hits versus misses. Striped bars indicate false alarms versus correct rejections. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$, one-tailed t test with Holm-Bonferroni correction.

(B) J , B , and S were based on linear combinations of prestimulus brain-state markers. (Bottom) Influence of J , B , and S on the probability of reporting stimulus present on no-stimulus trials (N.S., false alarms) and stimulus-present trials (hits) at the three difficulty levels, comparing lowest (blue) and highest quintiles (red). **(C)** Neuronal activity (smoothed with a 40-ms window) on stimulus-present trials within the highest quintile (lighter shades) and lowest quintile (darker shades) of the distribution of J . The shaded regions indicate \pm SEM as determined with bootstrapping. **(D and E)** Activity on stimulus-present trials within higher and lower quintiles of B (D) and S (E) during the prestimulus epoch. B was



defined on the basis of stimulus-absent trials but is applied here to sort stimulus-present-trials. **(F)** High values of B , which increase the false-alarm rate, are associated with higher firing rates throughout the trial, causing neurons to be closer to the threshold of ignition. High values of S are associated with more efficient propagation of neuronal activity to higher processing levels, increasing the separation between the distributions of signal strength on target-present and target-absent trials.

remarkably similar to that recorded in the monkeys (Fig. 4I and fig. S7; compare with Fig. 2), and weak stimuli tended to get lost at lower levels than did stronger stimuli (fig. S7D). Once ignition occurred at the higher levels, the feedback connections from

frontal and parietal regions to the visual cortex caused a small increase in activity in areas V1 and V4 after the stimulus had disappeared, just as in the data in the epoch before the saccade (Fig. 2 and fig. S3A). The model also accounted

for the profile of neuronal activity on trials without a stimulus. In a fraction of these trials, stochastic fluctuations in activity caused spontaneous ignitions at variable time points, giving rise to false alarms. When averaged across trials, time-locked




Fig. 4. Phosphene perception and GNWT model. (A) Phosphene detection task. On half of the trials, a train of five microstimulation pulses was delivered to area V1 in order to evoke a phosphene at the RF of the stimulated cells (green rectangle). The other half of the trials were without microstimulation. After 480 ms, the monkey reported a phosphene by making an eye movement to its previous location and its absence with an eye movement to the gray circle. (B) Accuracy as function of current amplitude in an example recording session in monkey B. (C) V4 MUA on correct trials during V1 microstimulation with different current intensities, averaged across all recording sites. Shaded region indicates \pm SEM. (Right) Example RF of stimulated V1 neurons (white rectangle) and of an example V4 recording site (false color). (D) V4 multiunit responses averaged across recording sites that were elicited by V1 currents below θ_{Low} (difficult; left), between θ_{Low} and θ_{High} (intermediate; middle), and higher than θ_{High} (easy; right). Green curves indicate hits, and red curves indicate misses. The square below the x axis illustrates the microstimulation epoch (five pulses, with an interval of 5 ms). (E) Miss fraction in V4 ($Activity_{Miss}/Activity_{Hit} \times 100\%$) for the different performance categories (time window 0 to 150 ms; shorter than in the visual experiments because the V4 response evoked by V1 stimulation was more transient). (F) Structure of the GNWT model. The visual stimulus activated the lateral geniculate nucleus (LGN), and feedforward connections (FF) propagated activity from the visual cortex (areas V1 and V4) to the parietal and frontal cortex. Self-connections were within the areas, and feedback connections (FB) propagated activity from higher back to lower areas. (G) Probability of target-present response as function of stimulus strength. a.u., arbitrary units. (H) Miss fractions are lower at higher cortical levels, indicating that more

activity is lost on the miss trials. (I) Activity elicited in model V1, V4, and frontal cortex (compare with Fig. 2A). When the activity in the frontal cortex of the model reaches a threshold level of activity, it can sustain itself (ignition, red and green arrows) because of the strong reciprocal connectivity with the parietal cortex. Activity dies down if this threshold activity level is not reached and the stimulus is missed. The black curves indicate activity for correct rejections, and the blue curves indicate activity for false alarms.

Downloaded from <http://science.sciencemag.org/> on August 30, 2018

to trial onset, these spontaneous ignitions led to a ramping of activity, just as in the dlPFC of monkeys (Fig. 2A). On these trials, feedback to areas V1 and V4 also caused slightly higher activity levels than on trials with correct rejections, just as observed in the monkey visual cortex (Fig. 2A).

Overall, our results provide new insights into why a fixed stimulus sometimes leads to a conscious report and sometimes remains subliminal and inspire unification of SDT and GNWT. Both data and model support the concept of multiple bottlenecks for conscious access: Weak stimuli tend to get lost at early processing levels, whereas stronger stimuli may transiently activate frontal cortex but still fail to reach the threshold for reportability. For conscious detection to occur in the model, stimuli must elicit a threshold level of prefrontal activity that is sufficient for ignition; stimuli that fail to reach this level are missed. Ignition corresponds to a self-sustained pattern of neuronal activity at the higher processing levels. Because of the variability of neuronal activity, the ignition threshold is occasionally reached on stimulus-absent trials, so that a false alarm occurs. The model proposes that ignition is caused by strong reciprocal interactions between the parietal and frontal cortex, in accordance with studies demonstrating that anesthesia weakens these interactions (15).

In our study, V1 activity was weaker on miss than on hit trials, in accordance with previous work (16), but our findings differ from results in the primary somatosensory cortex of monkeys in which neuronal responses did not predict the perception of weak tactile stimuli (17). Here too, however, V1 and V4 activity also did not fully predict perceptual report because neuronal activity evoked by missed visual or electrical stimuli could be at least as strong as that elicited by stimuli that were reported, implying information loss in downstream areas. By contrast, the activity level in the dlPFC categorically predicted perceptual report, implying that dlPFC lies at or beyond the stage that determines the reporting threshold.

We note a few limitations of our experiments. First, they were not aimed at revealing all of the brain regions that contribute to conscious reportability; although we focused on dlPFC neurons that contribute to eye movement planning, a broader set of high-level regions, linked by bidirectional connections, is likely to contribute to the postulated global workspace. Other cortical regions—upstream from the dlPFC and including, for example, the temporal (18–20) and parietal cortex (21, 22)—also exhibit extra activity if a stimulus reaches awareness, suggesting that they take part in conscious perception. Second, our design did not dissociate the brain regions required for conscious experience from those involved in conscious access and reportability (23). Because our goal was to investigate the mechanism of reportability and its fluctuations, we needed an explicit behavioral report in order to sort otherwise identical trials according to their subjective detection. It has been proposed that

PFC activity mainly reflects conscious reporting (24, 25). Although this account may suffice for the present results, PFC neurons also selectively represent consciously perceived stimuli during binocular flash suppression even if there is no need for report (26).

The higher brain regions provided feedback to areas V1 and V4, where the initial responses were driven by the stimulus, but later activity was stronger if the animals reported the stimulus (Fig. 2). A previous functional magnetic resonance imaging (fMRI) study observed a comparable increase of neuronal activity in early visual cortex when subjects reported a stimulus (27), but the fMRI signal did not differentiate between hits and false alarms, as though it was blind to stimulus presentation. The difference with the present results may be caused by the nature of the fMRI signal, which is insensitive to fine-grained timing but sensitive to processes other than spiking activity, such as synaptic activity of feedback connections (28, 29).

SDT stipulates an undefined source of signal fluctuations across trials. We found that part of this variability arises before stimulus onset. Various markers of prestimulus brain state could be combined to predict stimulus detection with accuracies greater than 60%. Additional measures of prestimulus brain state may further increase predictive power, although part of the unpredictability may be caused by the intrinsic stochasticity of neuronal activity, which was an essential ingredient of the model (Fig. 4F). It proved possible to define linear combinations of prestimulus brain-state markers with independent information about the subject's response bias and sensitivity. A bias to report target present was associated with a higher baseline firing rate across different brain regions, bringing neurons closer to the threshold for ignition (Fig. 3). By contrast, a higher sensitivity was associated with an improved propagation of neuronal activity to higher processing levels, increasing the difference in activity levels between target-present and target-absent trials at the processing stage that determines the threshold for ignition. Although our study did not address the brain mechanisms that influence the prestimulus firing rate and the quality of signal propagation, previous studies have established relations between pupil size and the frequency bands of the electroencephalography and the tone of neuromodulators such as noradrenaline (9) and acetylcholine (30). Future studies could examine whether the activity of these neuromodulatory systems indeed exerts separable influences on the animals' bias and sensitivity when stimuli are near the threshold of conscious perception.

REFERENCES AND NOTES

1. S. Dehaene, J.-P. Changeux, *Neuron* **70**, 200–227 (2011).
2. L. Weiskrantz, *Consciousness Lost and Found* (Oxford Univ. Press, Oxford, 1998).
3. D. M. Green, J. A. Swets, *Signal Detection Theory and Psychophysics* (Wiley, 1966).
4. B. J. Baars, *Trends Cogn. Sci.* **6**, 47–52 (2002).

5. V. A. F. Lamme, P. R. Roelfsema, *Trends Neurosci.* **23**, 571–579 (2000).
6. J. R. King, S. Dehaene, *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **369**, 20130204 (2014).
7. Y. Ko, H. Lau, *Philos. Trans. R. Soc. London B Biol. Sci.* **367**, 1401–1411 (2012).
8. A. V. Belopolsky, A. F. Kramer, J. Theeuwes, *J. Cogn. Neurosci.* **20**, 2285–2297 (2008).
9. J. Reimer *et al.*, *Neuron* **84**, 355–362 (2014).
10. V. Wyart, C. Tallon-Baudry, *J. Neurosci.* **29**, 8715–8725 (2009).
11. L. Iemi, M. Chaumon, S. M. Crouzet, N. A. Busch, *J. Neurosci.* **37**, 807–819 (2017).
12. M. C. Schmid *et al.*, *Nature* **466**, 373–377 (2010).
13. E. M. Schmidt *et al.*, *Brain* **119**, 507–522 (1996).
14. S. Dehaene, J.-P. Changeux, *PLoS Biol.* **3**, e141 (2005).
15. U. Lee *et al.*, *Anesthesiology* **118**, 1264–1275 (2013).
16. C. Palmer, S.-Y. Cheng, E. Seidemann, *J. Neurosci.* **27**, 8122–8137 (2007).
17. V. de Lafuente, R. Romo, *Nat. Neurosci.* **8**, 1698–1703 (2005).
18. G. Avidan *et al.*, *J. Neurophysiol.* **87**, 3102–3116 (2002).
19. L. Fisch *et al.*, *Neuron* **64**, 562–574 (2009).
20. R. Q. Quiroga, R. Mukamel, E. A. Isham, R. Malach, I. Fried, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 3599–3604 (2008).
21. F. Siclari *et al.*, *Nat. Neurosci.* **20**, 872–878 (2017).
22. W. X. Herman *et al.*, *Cereb. Cortex* **10.1093/cercor/bhx327** (2018).
23. C. Koch, M. Massimini, M. Boly, G. Tononi, *Nat. Rev. Neurosci.* **17**, 307–321 (2016).
24. S. Frässle, J. Sommer, A. Jansen, M. Naber, W. Einhäuser, *J. Neurosci.* **34**, 1738–1747 (2014).
25. N. Tsuchiya, M. Wilke, S. Frässle, V. A. F. Lamme, *Trends Cogn. Sci.* **19**, 757–770 (2015).
26. T. I. Panagiotaropoulos, G. Deco, V. Kapoor, N. K. Logothetis, *Neuron* **74**, 924–935 (2012).
27. D. Ress, D. J. Heeger, *Nat. Neurosci.* **6**, 414–420 (2003).
28. A. Maier *et al.*, *Nat. Neurosci.* **11**, 1193–1200 (2008).
29. N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, A. Oeltermann, *Nature* **412**, 150–157 (2001).
30. L. Pinto *et al.*, *Nat. Neurosci.* **16**, 1857–1863 (2013).

ACKNOWLEDGMENTS

Funding: The work was supported by Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Brain and Cognition grant 433-09-208 and ALW grant 823-02-010) and the European Union Seventh Framework Program (Marie-Curie Action PITN-GA-2011-290011 “ABC,” grant agreement 7202070 “Human Brain Project,” and European Research Council grant agreement 339490 “Cortic_algorithms”) awarded to P.R.R.; S.P. was supported by Fondation Bertarelli, and S.D. was supported by the Canadian Institute for Advanced Research. **Author contributions:** P.R.R., B.v.v., and B.D. conceived the study; B.v.v., B.D. and D.V. collected the data; H.S., S.P., and P.R.R. designed the analysis of prestimulus brain state; S.D. constructed the model with input from P.R.R.; B.v.v. and P.R.R. wrote the paper; and all authors commented on it. **Competing interests:** The authors declare no competing financial interests. **Data and materials availability:** The data and the computer code used to analyze the data are available for download and curated at the Human Brain Project Joint Platform (https://object.cscs.ch/v1/AUTH_227176556f3c44bb38df9eea4b91200c/hbp-data-000789/index.html). Correspondence and requests for materials should be addressed to P.R.R. or H.S.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/360/6388/537/suppl/DC1
Materials and Methods
Table S1
Figs. S1 to S9
References (31–52)
Matlab Files S1 and S2
Instruction Files S1 and S2

10 December 2017; accepted 9 March 2018
Published online 22 March 2018
10.1126/science.aar7186

The threshold for conscious report: Signal loss and response bias in visual and frontal cortex

Bram van Vugt, Bruno Dagnino, Devavrat Vartak, Houman Safaai, Stefano Panzeri, Stanislas Dehaene and Pieter R. Roelfsema

Science **360** (6388), 537-542.
DOI: 10.1126/science.aar7186originally published online March 22, 2018

Setting conscious perception alight

What are the neuronal mechanisms that enable conscious perception? Why do some images remain subliminal? Van Vugt *et al.* trained monkeys to detect low-contrast images and compared neuronal activity in brain areas V1, V4, and the dorsolateral prefrontal cortex. Some stimuli made it into consciousness, and others were subliminal depending on their propagation, which can be variable for weak stimuli (see the Perspective by Mashour). Strongly propagated stimuli initiated a state in the higher brain areas called "ignition" that caused information about a brief stimulus to become sustained and broadcasted back through recurrent interactions between many brain areas.

Science, this issue p. 537; see also p. 493

| | |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ARTICLE TOOLS | http://science.sciencemag.org/content/360/6388/537 |
| SUPPLEMENTARY MATERIALS | http://science.sciencemag.org/content/suppl/2018/03/21/science.aar7186.DC1 |
| RELATED CONTENT | http://science.sciencemag.org/content/sci/360/6388/493.full http://stm.sciencemag.org/content/scitransmed/8/368/368re5.full http://stm.sciencemag.org/content/scitransmed/5/208/208ra148.full http://stm.sciencemag.org/content/scitransmed/5/198/198fs32.full http://stm.sciencemag.org/content/scitransmed/5/198/198ra105.full |
| REFERENCES | This article cites 47 articles, 16 of which you can access for free http://science.sciencemag.org/content/360/6388/537#BIBL |
| PERMISSIONS | http://www.sciencemag.org/help/reprints-and-permissions |

Use of this article is subject to the [Terms of Service](#)