



Research Report

Multisensory perceptual awareness: Categorical or graded?



Jean-Paul Noel ^{a,b,*,1}, Nathan Faivre ^{c,d,**,1}, Elisa Magosso ^e, Olaf Blanke ^{c,f}, David Alais ^g and Mark Wallace ^{a,h,i}

^a Vanderbilt Brain Institute, Vanderbilt University, Nashville, USA

^b Center for Neural Science, New York University, New York City, NY, USA

^c Laboratory of Cognitive Neuroscience, Brain Mind Institute, Faculty of Life Sciences, Swiss Federal Institute of Technology (EPFL), Geneva, Switzerland

^d Laboratoire de Psychologie et Neurocognition, LPNC CNRS 5105, Université Grenoble Alpes, France

^e Department of Electrical, Electronic, and Information Engineering "Guglielmo Marconi", University of Bologna, Cesena, Italy

^f Department of Neurology, University of Geneva, Geneva, Switzerland

^g School of Psychology, The University of Sydney, Sydney, Australia

^h Department of Hearing and Speech, Vanderbilt University, Nashville, USA

ⁱ Department of Psychology, Vanderbilt University, Nashville, USA

ARTICLE INFO

Article history:

Received 29 December 2018

Reviewed 06 February 2019

Revised 31 March 2019

Accepted 30 May 2019

Action editor Angela Sirigu

Published online 27 June 2019

Keywords:

Multisensory

Consciousness

Neural networks

Psychophysics

Global neuronal workspace

ABSTRACT

Neural evidence suggests that mechanisms associated with conscious access (i.e., the ability to report on a conscious state) are “all-or-none”. Upon crossing some threshold, neural signals are globally broadcast throughout the brain and allow conscious reports. However, whether subjective experience (*phenomenal* consciousness) is categorical (i.e., transitioning abruptly from unconscious to conscious states) or graded (i.e., characterized by multiple intermediate states) remains an open question. To address this issue, we built a series of artificial neural networks containing distinct feedback connectivity from “multisensory” to “unisensory” cortices. In line with consciousness theories, we operationalized perceptual consciousness by the presence of feedback from higher-order nodes back to unisensory nodes which allow ‘neural ignition’ – a rapid, non-linear boost in response putatively leading to phenomenal consciousness. When simulating how these networks responded to unisensory and multisensory inputs, we found the fastest responses for multisensory presentations associated with multisensory feedback, and the slowest responses for multisensory presentations without feedback. Most interestingly, despite being built in line with “all-or-none” models of consciousness, multisensory stimuli associated with unisensory feedback (i.e., auditory or visual), and hence consistent with unisensory phenomenology according to theories of consciousness, generated

* Corresponding author. 7110 MRB III BioSci Bldg, 465, 21st Ave South, Nashville, TN, 3721, USA.

** Corresponding author. Laboratoire de Psychologie et NeuroCognition, CNRS UMR 5105, UGA BSHM, 1251 Avenue Centrale, 38058, Grenoble, Cedex 9, France.

E-mail addresses: Jean-paulnc@gmail.com (J.-P. Noel), nathanfaivre@gmail.com (N. Faivre).

¹ These authors contributed equally to the work.

<https://doi.org/10.1016/j.cortex.2019.05.018>

0010-9452/© 2019 Elsevier Ltd. All rights reserved.

intermediate reaction times. To extend these models to human perception and performance, we conducted extensive psychophysical testing in 29 subjects who each completed 10 h of a multisensory cue-congruency task. Consistent with the modeling results, we found that reaction times to multisensory cues reported as unisensory were intermediate between those of fully aware and fully unaware cues. These results support the existence of graded forms of phenomenological consciousness that can be instantiated by simple neural networks built in line with “all-or-none” models of consciousness.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

Theories of consciousness such as the Global Neuronal Workspace (GNW; Dehaene & Changeux, 2001) and Local Recurrence (LC; Lamme, 2006) emphasize the importance of global processing, and posit that recurrent activity among neural structures leads to consciousness. These theories state that stimuli reach consciousness when ascending sensory input is either strong enough in isolation or amplified via top-down attentional signals to result in “neural ignition” – the step-function broadcasting of sensory information throughout the neocortical mantle (Dehaene & Changeux, 2011). Feedback projections from prefrontal and parietal associative areas (in GNW; Dehaene & Changeux, 2011) or from higher-order sensory areas (in LC; Lamme, 2006) to lower sensory regions are considered crucial in supporting recurrent neural activity and thus consciousness. This all-or-none “neural ignition” account of neural processing in perceptual consciousness comes with a wealth of empirical support (van Vugt et al., 2018). Supporting data generally come from a comparison of neural states when participants report being aware (i.e., conscious) versus unaware of stimuli. Studies using approaches ranging from functional magnetic resonance imaging (fMRI; Grill-Spector, Kushnir, Hendler, & Malach, 2000; Dehaene et al., 2001; Rees, Kreiman, & Koch, 2002), to magneto- and electro-encephalography (M/EEG; Del Cul et al., 2006, 2007; Fahrenfort, Scholte, & Lamme, 2007; Gaillard et al., 2009; Noel, Simon et al., 2018), and single-unit recordings (Leopold & Logothetis, 1996; Logothetis & Schall, 1989; Noel, Ishizawa, Patel, Eskandar, & Wallace, 2018; Quiroga, Mukamel, Isham, Malach, & Fried, 2008; van Vugt et al., 2018), all agree that neural activity is more sustained when participants report being aware compared to unaware of the stimuli (Dehaene & Changeux, 2011; however see Frässle, Sommer, Jansen, Naber, & Einhäuser, 2014; Tsuchiya, Wilke, Frässle, & Lamme, 2015; for recent objections to contrasting reports of perceived vs non-perceived in consciousness studies).

Nevertheless, while access to consciousness (i.e., the ability to report) appears likely to be supported by categorical neural transitions, it is still a matter of debate whether phenomenological experience is categorical or graded (Block, 1995; Cleeremans & Jimenez, 2002; Kouider, de Gardelle, Sackur, & Dupoux, 2010). Take for example the case of multisensory consciousness, which may be described either as cohesive and singular (i.e., categorical change from no

experience in dreamless sleep to multisensory experience), or conversely, as a collection of unisensory experiences dynamically mixing throughout daily life in a graded fashion (Deroy, Chen, & Spence, 2014; Noel, Wallace, & Blake, 2015; O’Callaghan, 2017). Interestingly, theories of consciousness that start from phenomenological introspection – as opposed to empirical data (e.g., Integrated Information Theory; Tononi & Koch, 2015; Tononi, Boly, Massimini, & Koch, 2016) – suggest that phenomenological consciousness is graded. This work has derived consciousness indices placing patients with different disorders of consciousness (Casali et al., 2013; Sarasso et al., 2015) and in different sleep/awake cycles (Andrillon, Poulsen, Hansen, Léger, & Kouider, 2016; Schartner et al., 2015; 2017) along a continuum of consciousness. Thus, while theories based on neural data suggest that consciousness is all-or-none, phenomenological introspection suggests that consciousness is graded.

In the current work, we seek to reconcile this paradox by testing whether circuit architectures that support categorical neural steps can produce graded patterns of response. To address this “categorical versus graded” debate (Kouider et al., 2010), we structured neural networks that formalize the role of feedback connections and neural non-linearities (Fig. 1, upper row) in biologically-inspired, non-spiking neural networks. In line with theoretical postulates supporting “all-or-none” consciousness (Dehaene & Changeux, 2011; Lamme, 2006; Joglekar, Mejias, Yan, & Wang, 2018), each sensory area processes information via a sigmoidal activation function. Furthermore, multisensory consciousness was operationalized by the presence of feedback from an audio-visual (AV) network node back to unisensory regions (Fig. 1, rightmost panel). In such architecture, given the non-linear input–output function, as well as the connectivity patterns, neural ignition follows once sensory input is sufficiently robust. In contrast, a network was considered to not be able to support consciousness – not even “in principle” – when no feedback was present (second row, leftmost panel). We also built networks that, arguably and in principle, may be capable of supporting consciousness in a single sensory modality (e.g., audio consciousness when there is a feedback connection from AV to A but not V, and visual consciousness when there is a feedback connection from AV to V but not A; Figure 1, 1st row, middle panels). These networks are not full models of consciousness, but best understood as computational instantiations of consciousness theory that postulates that feedback connectivity is central to phenomenology. We then

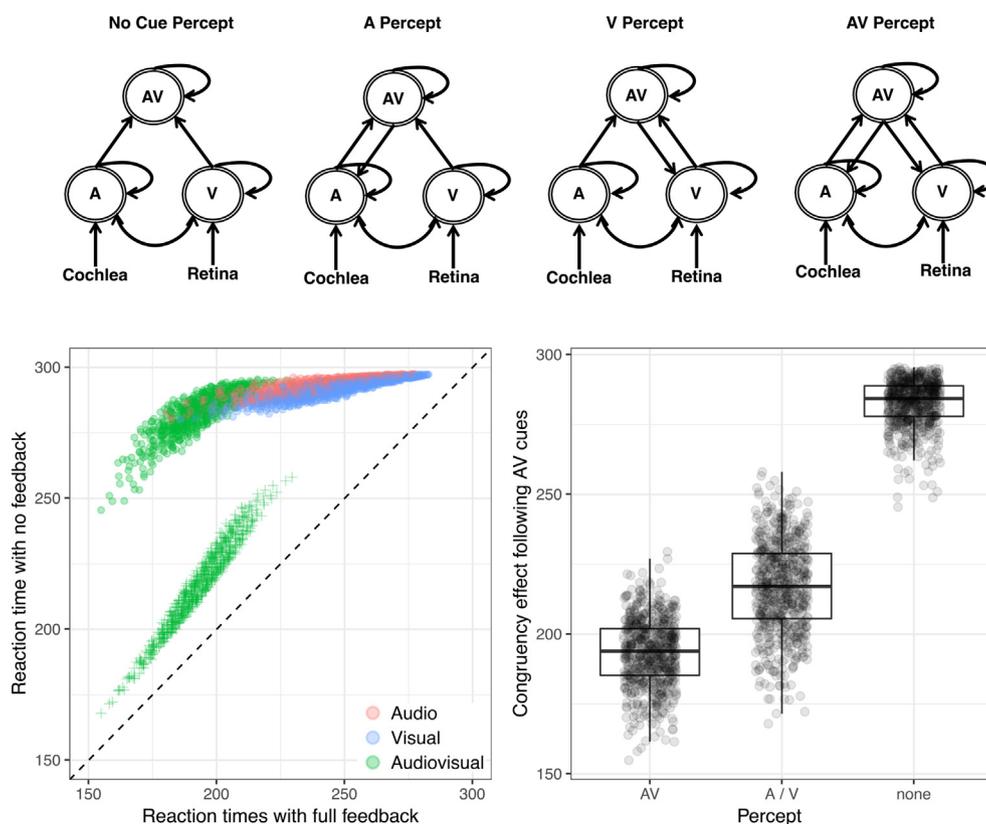


Fig. 1 – Neural Network Implementation and Results. Top: According to the Global Neural Workspace (GNW) and Local Recurrence (LC) theories of consciousness, a stimulus becomes conscious and reportable when it engages both feedforward and feedback projections, hence establishing a recurrent network of sustained activity and neural information is widely broadcast throughout the brain. Hence, when feedback projections are engaged and a recurrent network is formed, the system is capable of auditory-alone (top row, second column; no feedback projection to V), visual-alone (top row, third column; no feedback projection to A), or audiovisual consciousness. Bottom: Observed pattern of reaction times as a function of auditory, visual, and audiovisual cues, and whether the network is capable of no consciousness and full consciousness (left-hand plot; no feedback or full feedback; results plotted are from the “general” model). Left panel equally plots AV presentations that are partially perceived (e.g., AV perceived as V or A, green crosses) versus fully perceived trials (x-axis). A total of 625 dots/crosses are plotted, each being the average of 100 trials simulated for a given parameter set (625 sets of parameters). Right panel emphasizes the reaction times to AV stimulation (green in left panel) when they were fully perceived as AV, partially perceived as A or V, and not perceived at all.

ran simulations in these networks to characterize the graded versus categorical relation between putative states of consciousness, and estimated task performance based on simulated reaction times. In a second modeling effort, we left all feedback projections intact, as in the AV-consciousness model. We hypothesized that while conscious access (AV, A, V, none) is dependent on threshold crossings (Dehaene & Changeux, 2001; Lamme, 2006; van Vugt et al., 2018), graded phenomenology involves an integrative function. The modeling results show that a network with non-linear input–output functions can give rise to graded reaction times, which we argue could reflect graded phenomenology.

In the second element of this work, we then assessed to what extent these simple models may capture human behavior under distinct states of phenomenal consciousness. For this purpose, we presented participants with unisensory (i.e., V alone, A alone) and multisensory (i.e., VA) spatial cues near their threshold for detection, and asked them to report

the location of a subsequent supra-threshold tactile target via a speeded button press. We compared reaction times to the tactile targets depending on whether perception of the preceding cues was self-reported as being multisensory, unisensory (i.e., partial consciousness) or unperceived. The assumption here is that strength of congruency priming indicates the vividness of phenomenological experience (see Faivre, Mudrik, Schwartz, & Koch, 2014; Gelbard-Sagiv, Faivre, Mudrik, & Koch, 2016 for related findings). Of particular interest was whether reaction times following partially perceived cues would more resemble those following fully perceived versus unperceived cues. We argue that a graded account of consciousness would be supported by a specific pattern of reaction times following multisensory cue presentations; If multisensory cues can be partially conscious, such cues should generate response times intermediate between response times for fully perceived versus non-perceived cues.

2. Methods

2.1. Neural network modeling

2.1.1. Network description

The network is built to generate artificial reaction times given a visual, auditory, or audiovisual stimulus. In the case of audiovisual stimulation, auditory and visual stimuli are temporally and spatially coincident. In its general form, each neuron in the network receives a net input and responds to it via its intrinsic temporal dynamics (i.e., a low-pass first order dynamics) and a sigmoidal activation function. The net input activity is the sum of external network-based inputs as well as auto-excitatory inputs. The neuron's response generates an output activity, which represents the neuron's firing rate. Due to the sigmoidal activation function, each node's output is constrained between 0 (neuron's spontaneous activity) and 1 (neuron's maximal activation) and the input–output function is non-linear.

The networks developed here has five neural areas (see Fig. 1): a region representing the cochlea, area A representing a cortical auditory area, a region representing the retina, region V representing a cortical visual area, and finally region AV representing a cortical audiovisual multisensory area. For simplicity, each cortical area is simulated by a single node, representing an ensemble of cells sharing similar properties. In the networks, the cochlea and the retina just replicate the external input (i.e., they simulate the presence of a sensory auditory and visual stimulus respectively, applied for a given time interval) with the addition of a noise drawn randomly on each trial from a normal distribution. From these peripheral regions, the input is then propagated to cortical areas. The cochlea projects to area A with weight, W^{ac} while the retina projects to area V with weight W^{vr} . In turn, unisensory cortical areas project to the multisensory area with feedforward strengths W^{ma} and W^{mv} , respectively, in the case of audition and vision. Further, unisensory areas are reciprocally connected via inhibitory connections of strength I^{av} (from visual area to auditory area) and I^{va} (from auditory to visual area). Direct connections between unisensory areas are well established anatomically in biological systems (see Kayser, Petkov, & Logothetis, 2009, for review) and the mutual inhibitory pattern is routinely employed in neural networks in order to instantiate competition between areas (Cisek, 2007; Seely & Chow, 2011). Further, these inhibitions contribute to prevent the formation of “phantom activation” (e.g., activation of area V under auditory stimulation alone) due to feedback connections. The multisensory area, in turn, depending on the particular model (Fig. 1), may or may not send feedback connections in return to unisensory areas; to A with weight B^{am} and to V with weight B^{vm} . Lastly, in order to simulate the role of lateral synapses within each area, regions A, V, and AV receives excitatory self-connections with weight L^a , L^v , L^m , respectively. See Fig. 1 for network diagrams and Table 1 for the values of model parameters. It is important to note that an array of 625 different parameter values (parametrically varied) were used, indicating that the reported results are specific to network architectures and not the particular parameters we use.

2.1.2. Network implementation and readout

2.1.2.1. “A PRIORI” MODELS. All equations were implemented and numerically solved within the MATLAB (MathWorks Inc., Natick, USA) software environment using a simulation time step $\Delta t = .1$ msec. Simulations started at $t = 0$ from the resting condition (unperturbed network). Then, an external input, representing a visual and/or auditory stimulus, was applied after 20 msec of simulation (i.e., 200 time steps, hence onset of stimulus application is $t_{on} = 20$ msec) for a duration of 10 msec (100 time steps, as in the behavioral experiment below, hence offset of stimulus application is $t_{off} = 30$ msec). Simulations terminated after 1200 iteration steps, corresponding to $t_{end} = 120$ msec. This simulation length was utilized as it was largely sufficient for the network to exhaust its initial transient response and reach a new steady state in response to the stimulation.

The overall input (i.e., $u(t)$) to a generic neuron in the network is processed via functions governing first-order temporal dynamics Eq. (1) and sigmoidal activation Eq. (2), generating the neuron's output activity (i.e., $y(t)$):

$$\tau \frac{dq_s(t)}{dt} = -q_s(t) + u_s(t) \quad (1)$$

$$y_s(t) = F(q_s(t)) = \frac{1}{1 + \exp(-(q_s(t) - \varphi_s)\xi_s)} \quad (2)$$

where subscript s may assume value a , v , or m (auditory area A, visual area V or multisensory area AV). In Eq. (1) τ represents the internal time constant of the neuron (we assume the same time constant for all neurons), and ξ_s and φ_s are parameters which establish respectively the slope and the central position of the sigmoidal relationship (see Table 1 for parameter values). According to Eq. (2), the output activity of each neuron is constrained between ~ 0 (i.e., neuron's spontaneous activity) and 1 (i.e., neuron's maximal activity).

The net input $u(t)$ that reaches a neuron may be generally written as the sum of three contributions: an external input $e(t)$ due to a stimulus being presented (auditory, visual, or audiovisual), a lateral input $L(t)$ mimicking lateral synapses (self connections), and network inputs coming from neurons in other areas via inter-area synapses (e.g., I^{av} and I^{va} in Equations (3) and (4), and W^{mv} and W^{ma} in Equation (5)). Hence, given the network architecture and synaptic weighting, we can more precisely describe these relationships in the following equations (in this case for a network with no feedback connections):

$$u_a(t) = e_a(t) \cdot W^{ac} + y_a(t) \cdot L^a - y_v(t) \cdot I^{av} \quad (3)$$

$$u_v(t) = e_v(t) \cdot W^{vr} + y_v(t) \cdot L^v - y_a(t) \cdot I^{va} \quad (4)$$

$$u_m(t) = y_v(t) \cdot W^{mv} + y_a(t) \cdot W^{ma} + y_m(t) \cdot L^m \quad (5)$$

$e_a(t)$ and $e_v(t)$ in Eqs. (3) and (4) represent the signals from the cochlea and the retina having the following expression:

$$e_s(t) = \begin{cases} 0, & 0 < t < t_{on} \\ E_s, & t_{on} \leq t \leq t_{off} \\ 0, & t_{off} < t < t_{end} \end{cases} \quad (6)$$

In Eq. (6), s holds for a and v . E_s is a constant value

Table 1 – Neural network parameters.

| Dynamics and sigmoidal activation function within nodes | | | | |
|---|-------------------------------|---------------------------------|----------------------------|-----------|
| $\Delta t = .1$ ms | $\tau = 3$ msec | $T = 120$ msec | | |
| $\varphi_a = [8 : .5 : 10]^*$ | $\varphi_v = [8 : .5 : 10]^*$ | $\varphi_m = [5 : .5 : 7]^{**}$ | | |
| $\xi_a = .75$ | $\xi_v = .75$ | $\xi_m = .75$ | | |
| External auditory and visual stimuli (Mean μ and standard deviation δ of the Gaussian from which the random input was drawn) | | | | |
| $\mu = 10$ | | $\delta = 2.5$ | | |
| Inter-area and lateral synapses | | | | |
| $L^a = 1$ | $W^{ac} = 1$ | $L^v = 1$ | $W^{vr} = 1$ | $L^m = 1$ |
| | $W^{ma} = [8:.5:10]^{***}$ | | $W^{mv} = [8:.5:10]^{***}$ | |
| | $I^{av} = 4$ | | $I^{va} = 4$ | |
| | $B^{am} = [4:.5:6]^{****}$ | | $B^{vm} = [4:.5:6]^{****}$ | |
| Consciousness decoding (“general model”) | | | | |
| $T_u = .2$ | | $T_m = .8$ | | |

* Range utilized during different simulations. Central point of the sigmoidal activation function in area A and V were always yoked.
**Range of utilized values for the central point of the multisensory area's sigmoidal activation function.
*** Weight of feedforward projection from unisensory to multisensory areas was modulated in different simulations yet A and V projections were always the same.
**** Weight of feedback projections from the multisensory to unisensory areas was modulated in different simulations yet projection weights were always the same to both A and V areas.

representing input strength drawn randomly on each trial from the normal distribution, $N(\mu, \delta)$ (of course in the absence of an external auditory or visual stimulus, E_a or E_v are set to 0). Auditory and visual noise constants are drawn independently on each trial. Contrarily, in a network where both auditory and visual unisensory areas receive feedback, equations (3)–(5) above would take on the following form:

$$u_a(t) = e_a(t) \cdot W^{ac} + y_a(t) \cdot L^a - y_v(t) \cdot I^{av} + y_m(t) \cdot W^{am} \quad (7)$$

$$u_v(t) = e_v(t) \cdot W^{vr} + y_v(t) \cdot L^v - y_a(t) \cdot I^{va} + y_m(t) \cdot W^{vm} \quad (8)$$

$$u_m(t) = y_v(t) \cdot W^{mv} + y_a(t) \cdot W^{ma} + y_m(t) \cdot L^m \quad (9)$$

Four different models were implemented by keeping all parameters equal except for the feedback projections existing from area AV to unisensory areas. As a short hand, we refer to these models as in theory putatively supporting no consciousness, A-consciousness, V-consciousness, and AV-consciousness. Of course, none of these models actually instantiates any form of subjective experience and their architectures are too simple to provide insight into the biological neural processing supporting consciousness. They are simply meant to reflect the theoretical position (Dehaene & Changeux, 2011; Lamme, 2006) that feedback connectivity may, in principle, be important in consciousness. That is, they are best understood as simple computational instantiation of well-established theories of consciousness (Dehaene & Changeux, 2011; Lamme, 2006).

Each of these models was initialized with a set of 625 different parameters, parametrically manipulating the key variables φ_a and φ_v (slope of sigmoidal activation function in areas A and V, which were always kept the same), φ_m (slope of the sigmoidal activation function in area AV), W^{am} and W^{vm} (strength of feedforward projection from unisensory areas to the multisensory region, which were always kept the same), and B^{ma} and B^{mv} (strength of feedback projection from the

multisensory area to unisensory areas, which were always kept the same; see Table 1 for details). For each of these instantiations, 100 trials were simulated with varying input levels according to the normal distribution, $N(\mu, \delta)$. For each trial, the area under the activation curve in the multisensory region was a priori considered to be proportional to reaction time, well in line with panoply of evidence suggesting perceptual decisions are reached following an accumulation process (Shadlen, Hanks, Churchland, Kiani, & Yang, 2006; Usher & McClelland, 2001). The area under the activation curve in the multisensory area was estimated via the trapezoidal method, and this area was multiplied by 10 (ms) and subtracted from a baseline reaction time of 300 msec. In other words, for a given trial, the area under the activation curve of the multisensory region was inversely proportional to the simulated reaction time and had a slope of 10 msec per unit. The 100 trials for a given parameter set were averaged, and subsequently a grand mean reaction time for all 625 parameter set was calculated for each model.

2.2. General model

The separate “a priori” models described above most clearly illustrate cases in which, according to the GNW and LC theories (Dehaene & Changeux, 2011; Lamme, 2006), networks may model a system supporting consciousness of different perceptual elements (none, A, V, or AV). Nevertheless, as humans count with a single brain, and as feedback projections do not appear and disappear to allow for uni versus multisensory consciousness, we developed a more general model accounting for all types of unisensory and multisensory behavior with a common architecture. Namely, we left all feedback projections intact (as in AV-consciousness model) and assumed that conscious access (AV, A, V, none) was dependent upon threshold crossings, in line with GNW/LC (Dehaene & Changeux, 2001; Lamme, 2006) and empirical

observations (e.g., van Vugt et al., 2018). Inspired by neuroimaging studies (see Tong, 2003, for review), we assumed that consciousness is extracted based on the peak activation level in both unisensory and multisensory areas. That is, an input is taken to evoke a conscious perceptual experience when it elicits a neural activation peak above a given threshold in both unisensory (T_u) and multisensory (T_m) areas (see Table 1 for details). When the peak activation in area AV surpasses T_M and peak activation in areas A and V surpass T_u , the stimulus is considered to be AV-conscious (all areas are supra-threshold for consciousness). If area AV surpasses T_M but one of the unisensory areas (e.g., area V) does not peak above the T_u threshold, the stimulus is considered to be perceived as audio alone. With the exception of the decoding of perceptual consciousness (i.e., no-consciousness, A-consciousness, V-consciousness, and AV-consciousness), implementation and readout in the “general” model followed exactly as described for the “a priori” models.

2.3. Psychophysics

2.3.1. Rationale

We aimed to query whether human behavior feel in line with the modeling results. Inspired by previous work on unconscious multisensory congruency priming (Faivre et al., 2014), we devised a task where multisensory cues predicted the location of subsequent tactile targets. Given that our model captured how multisensory cues are processed under distinct putative conditions of conscious access, we expected partially perceived cues to generate responses to the target location with intermediate reaction times, lying between responses following fully perceived and unperceived cues. We report how we determined our sample size, all data exclusions, all inclusion/exclusion criteria, whether inclusion/exclusion criteria were established prior to data analysis, all manipulations, and all measures in the study.

2.3.2. Participants

Twenty-nine healthy volunteers (14 females, mean age 20.9 ± 2.3 years), including one of the authors (JPN) took part in this experiment. All participants self-reported normal hearing and touch, and had normal or corrected-to-normal visual acuity. Participants gave written, informed consent to partake in the study, which was approved by Vanderbilt University Medical Center's ethics board. No part of the study procedures or analyses was pre-registered in a time-stamped, institutional registry prior to the research being conducted. The data and code is publicly available at: <https://osf.io/45gch/>. A sample size commensurate with previous reports (Faivre et al., 2014) was collected – however in this case we collected approximately 10 h of data per subject, to assure not only between-subject but also within-subject reproducibility.

2.3.3. Materials and apparatus

In a congruency-priming task, participants were prompted toward either their left or right hand via auditory, visual, or audiovisual cues located at the fingertips before the delivery of a target tactile stimulation to one of the index fingers. Auditory stimuli consisted of a pure tone presented for 10 ms at a frequency of 2.3 kHz and 50 dB SPL intensity, via a Piezo

Buzzer (F/UCW 06 Piezo Buzzer, Digisound, Norderstedt, Germany). Participants wore headphones (HD 558, Sennheiser, Wedemark, Germany) over which broadband noise was continuously delivered (range: 57 dB SPL – 91 dB SPL). The ratio between the tone and noise was determined using a staircase procedure, in order to titrate detection performance to perceptual threshold (see below). Visual stimuli consisted of a transient luminance increase produced by a white LED (3 mm diameter, Adafruit, New York, United States; background luminance: 200 mcd; increased by 2 mcd, as determined by a staircase procedure; see below). Audiovisual cues consisted of the synchronous and co-localized presentation of auditory and visual cues at their predetermined threshold level. The target stimuli consisted of a 100 msec supra-threshold vibrotactile stimulation (model 312-101, Precision MicroDrives, London, United Kingdom) presented 600 msec after the cue offset. A vibrator was attached with medical tape to the left and right index finger (medial phalange) of participants. Auditory, visual, and tactile stimulation were controlled via a micro-controller (clock rate: 16 MHz; ATmega1280, Arduino, Italy) in serial communication with a PC (Dell Vostro, 7000 Series, Round Rock, TX, USA; Baud rate: 115200 Hz). LEDs were connected to Pulse-Width-Modulated (PWD) pins, which allowed dividing input voltages into 250 steps. General experimental procedures were controlled via purpose-made MATLAB (MathWorks Inc., Natick, USA) scripts in conjunction with the Psychophysics Toolbox extension (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997). Timing of all experimental components was verified with an oscilloscope (HM507, Hameg, Germany).

2.3.4. Procedure

Participants sat in an unlit and sound-attenuated room (WhisperRoom™) and rested their chin on a chinrest approximately 60 cm from a CRT monitor (Mitsubishi Diamond Pro 2070SB, Sydney, Australia). Before undertaking the main experiment, auditory threshold in noise and visual difference threshold were measured. In separate blocks for audition and vision, an adaptive 1-up-1-down staircase approach was undertaken wherein auditory noise was initially delivered at 65 dB SPL and adjusted in steps of .5 dB SPL, while visual targets had an initial differential of 30 mcd with respect to a baseline of 200mcd and were adjusted in steps of 2 mcd. In the case of the auditory staircase, participants performed a two alternative forced choice (2AFC) task wherein they were asked to indicate whether a 10 msec beep had occurred leftward or rightward of fixation. Similarly, in the visual staircase they were asked in a 2AFC task to indicate whether the leftward or rightward LED (10 msec) had flashed. The locations of auditory and visual signals were the same during the staircase procedure and the main experiment. The staircases were continued for a total of 40 trials yielding a final split of approximately 50% of targets correctly discriminated. The final auditory noise and LED intensity differentials were set as parameters to be utilized in the main experiment. Once auditory and visual parameters were set, participants were asked to place their left and right index fingers on response keys situated immediately in front of them, parallel to their left and right shoulders, respectively. The Piezo Buzzer delivering the auditory cues and LEDs emitting the visual cues were

placed at shoulder width and three centimeters further in depth from the response keys (see Fig. 2, top row). Participants were informed that a peri-threshold auditory, visual, or audiovisual cue would occasionally precede a tactile target stimulation they ought to respond to as quickly as possible by button press. In 80% of cued trials the cue was congruent with the tactile target (e.g., left hand cued correctly indicated target tactile stimulation on the left hand), while on the remaining 20% of cued trials the cue was incongruent (e.g., left hand cued and target tactile stimulation on the right hand). In addition, participants were informed that no-cue trials would also occur (10% of all trials). After responding to tactile stimulation by button press, a fixation point on the CRT monitor turned from red to green, indicating that the tactile response was registered and prompting the participants to answer a second question regarding the sensory nature of the cue; auditory, visual, audiovisual, or no cue. That is, on every trial the

participant provided a speeded response to tactile stimulation, and a report regarding their subjective experience of the cue. The session was divided into 6 blocks, each block composed of 12 tactile stimulations preceded by no-cue (6 left and 6 right index finger), 16 audio congruent cues, 16 visual congruent cues, 16 audiovisual congruent cues, 4 audio incongruent cues, 4 visual incongruent cues, and 4 audiovisual incongruent cues. Thus, every block was 56 trials long. In total, during each session each cue was repeated 96 times in the congruent condition, and 24 times in the incongruent condition. All trial types were randomized within-blocks and inter-trial interval lasted between 1.0 and 2.5 sec (randomly sampled from a uniform distribution). Participants were allowed to take self-paced breaks between blocks, and before initiation of the main experiment, half a block of practice trials was administered. The entire duration of one experimental session was about 60 min. In total participants took

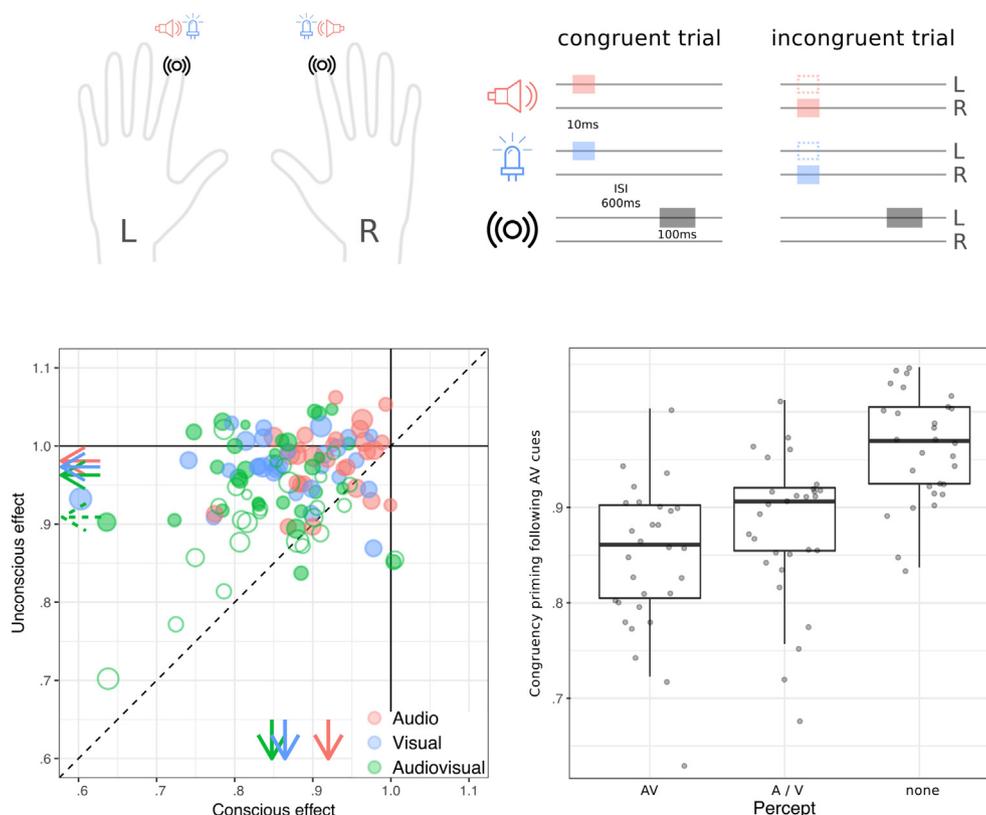


Fig. 2 – Experimental Paradigm and Psychophysical Results. Top: Participants were given a valid (80% of trials, middle column top panel) or invalid (20% of trials, rightmost panel top panel), auditory (red), visual (blue), or audiovisual (both red and blue, as illustrated) cues as to the location of a subsequent tactile target (black) they had to respond to as fast as possible. Spatial cues were titrated before each experimental session and for each subject in order to be at detection threshold. Bottom: Behavioral data showing the effect of auditory, visual, and audiovisual cues to a lateralized tactile detection task with cues presented at detection threshold. Left panel shows the ratio of mean reaction times for valid relative to invalid cues where values less than 1 indicate a congruency priming effect for tactile targets when preceded by an auditory (red), visual (blue), or audiovisual (green) cue. The data are subdivided by whether participants perceived the cue consciously (x-axis) or reported no-cue (y-axis). Additionally, open green circles (and dashed green arrow on the y-axis) represent audiovisual trials where participants were only partially aware; reporting either A or V percept, but not AV. Every dot represents a participant and the colored arrows on the axes of the scatter plot show the mean of the group. Right panel shows boxplot and individual subject data for the congruency effect associated with AV presentation when participants reported the full percept (left-most), a partial percept (middle), and no percept (right-most), demonstrating a graded effect of consciousness on multisensory performance.

part in 10 different 60-minute sessions (maximum of 2 session per day separated by at least a 15 min break), and hence in total completed approximately 50 blocks (over 5000 trials per participant). This high number of repetitions permitted us to collect a sufficiently large sample of trials with partially perceived stimuli (>300 trials/subject and trial type).

2.3.5. Analyses

The cue in a given trial was defined as conscious if it was properly classified by the participant's subjective report (e.g., a 'visual' report following a visual cue, an 'audiovisual' report following an audiovisual cue), and as unconscious if it was missed (i.e., report of no sensory experience following auditory, visual, or audiovisual cues). Trials were excluded if the cue was misclassified (e.g., visual report following an audio cue). Partially conscious trials were defined as those in which AV stimuli were presented, but solely A or V was reported. Trials with incorrect tactile discrimination, or with reaction times faster than the 2.5th percentile or slower than the 97.5th percentile for a given participant and condition of conscious access were excluded (amounting to 1.7% and 5.0% of total trials, respectively). The remaining reaction times were analyzed with linear mixed effects models (`lme4` and `lmerTest` packages; Kuznetsova, Bruun Brockhoff, & Haubo Bojesen Christensen, 2013; Bates et al., 2014), with the fixed effects being cue (visual, auditory, or audiovisual) and congruency (congruent or incongruent), and random intercepts for subjects. Random slopes for each fixed effect were included in the model, following model selection based on Bayesian information criterion (BIC). Significance of fixed effects was estimated using Satterthwaite's approximation for degrees of freedom of F statistics (Satterthwaite, 1946). Conditions of conscious access of cues (i.e., conscious, unconscious, or partially conscious) were analyzed separately. All behavioral analyses were performed with R (2016) with `ggplot2` (Wickham, 2009) for graphical representations.

3. Results

3.1. Neural network modeling

Results from the neural network simulations must be assessed qualitatively, as statistical significance in these networks is heavily influenced by the number of trials simulated within a parameter set (analogous to within-subjects trials in psychophysics), the number of parameter sets utilized (analogous to number of subjects in psychophysics), and the RT decoding scheme (differences between conditions may be accentuated or diminished by the arbitrary relation between neural activation and the reaction time output).

Regarding the separate "a priori" models, on average the simulated reaction time under a neural architecture with no feedback projections was 272.1 msec to audio input, 277.6 msec to visual input, and 256.3 msec to audio-visual input. Thus, multisensory facilitation, or the difference between the fastest unisensory condition and the multisensory condition, was on the order of 15.7 msec for a network with no feedback projections. When feedback projections were introduced between the multisensory node and a single unisensory

node (either A or V), the average reaction time to AV stimuli was 243.8 msec. With such a network structure, the average reaction time to the faster unisensory condition (usually A when feedback from AV to A was present, and V when feedback from AV to V was present) was 264.8 msec, and the average reaction time was 274.7 msec for the slower unisensory condition (V when feedback was from AV to A, and A when feedback was from AV to V). Thus, multisensory facilitation of reaction time predicted by the network modeling partial consciousness was on the order of 21 msec. Lastly, when the network was equipped with feedback projections from the multisensory area to both unisensory areas, the average reaction times were 261.3 msec for audio input, 268.1 msec for visual input, and 234.8 for audiovisual input. Hence, in the network modeling full, multisensory facilitation was on the order of 26.5 msec.

To summarize, architectures with no feedback, partial feedback, or full feedback produced facilitated reaction times (unisensory vs multisensory) of 15.7, 21.0 and 26.5 msec, respectively. A linear interpolation between AV performance in the no-feedback and complete feedback networks suggests that partial feedback should result in reaction times of about 245.5 msec. The mid-point between this linear interpolation (245.5 msec) and the full feedback condition is 240.1 msec; this latter value therefore demarcates the boundary where simulated partially perceived trials would be closer to either full consciousness (less than 240.1 msec) or to the linear interpolation between full consciousness and no consciousness (more than 240.1 msec). As stated above, the average reaction time to partially perceived multisensory presentation was 243.8 msec, suggesting that the network with partial feedback lies closer to the linear interpolation than to the fully connected or disconnected networks.

A similar pattern of results emerges when considering the "general" model in which putative reports of consciousness (i.e., access consciousness) are decoded relative to an activation threshold (results plotted in Fig. 1, bottom row), and not assumed given connectivity patterns (as in the "a priori" models, where solely RTs are decoded, and not access consciousness). The average reaction time when both components of an audiovisual stimulus were deemed to be consciously perceived was 193.2 msec. The faster average unisensory reaction time that was fully perceived (A-consciousness when A was presented or V-consciousness when V was presented) was 238.8 msec, while the slower average was 248.0 msec. Thus, multisensory facilitation under full consciousness in this model was on the order of 45.5 msec (see Fig. 1; contrast between red/blue and green dots projected on the x-axis). In contrast, when stimuli were not perceived, the average reaction times were considerably slower (292.7 msec for audio stimuli, 291.8 msec for visual stimuli, and 282.6 msec for audiovisual stimuli; see Fig. 1) and multisensory facilitation was reduced to approximately 9 msec (see Fig. 1; contrast between red/blue and green dots projected on the y-axis). Lastly, when the consciousness threshold was surpassed in the multisensory node and one of the unisensory nodes (operationally resulting in partial consciousness when audiovisual stimuli were presented or full consciousness when a unisensory stimulus was presented), the average reaction time to multisensory stimuli (yet experienced as unisensory)

was 228.6 msec. The average reaction time to the faster unisensory stimulus (i.e., the perceived one) was 243.4 msec, while it was 292.3 msec to the slower unisensory stimuli. Thus, under conditions of partial consciousness, multisensory facilitation was 14.8 msec. In order to determine whether the partially perceived trials are indicative of an abrupt transition (where partially perceived RTs would lie closer to full consciousness than the mid-point between full consciousness and no consciousness) or a graded transition (where partially perceived RTs would lie closer to the linear interpolation between full and no consciousness, as opposed to closer to one of the extremes) we performed a linear interpolation. The interpolation between AV aware and unaware performance suggests that the partial consciousness condition should result in reaction times of about 237.9 msec. The mid-point between this linear interpolation and the aware condition is 215.5 msec (this last number thus demarking the boundary between evidence for categorical vs graded transition), which is quicker than the partially aware condition, suggesting that partially aware performance is closest to the linear interpolation between fully aware and fully unaware cases (vs closer to the fully aware case; see Fig. 1, bottom right panel; A/V percept lies closer to AV than “none”, but closer yet to the linear interpolation between AV and “none”).

In summary, whether building distinct neural models formalizing the hypothetical relationship between consciousness and feedback connections postulated by GNW and LC theories (“a priori models”; Dehaene & Changeux, 2001; Lamme, 2006), or building a “general” model in which both reaction times (which here serves as a proxy for putative phenomenal consciousness) and reports of consciousness (access consciousness) are decoded, multisensory performance appears to follow a graded pattern, being best (i.e., fastest) under full consciousness, intermediate under partial consciousness, and worst when not perceived. The partially perceived trials are closer to a true interpolation between fully conscious and unconscious cases, than closer to either of these. That is, there is no abrupt step, but a gradual transition.

3.2. Psychophysics

A linear mixed-effect analysis on RTs in trials where cues were consciously perceived revealed a main effect of congruency ($F(1,29) = 54.97, p < .001$), and an interaction between cue modality (A, V, or AV) and congruency ($F(2,32148) = 45.21, p < .001$). It was found that AV cues induced a larger congruency effect (median of normalized congruency effect [RT congruent/RT incongruent] = $.861 \pm .03$; a value of 1 indicates no effect) when compared with A cues ($.925 \pm .02, p < .001$; see Fig. 2, bottom left panel, arrow (medians) on the x-axis). The difference between AV and V cues did not reach significance ($.857 \pm .03, p = .40$). The same RT analysis for trials where the cue was not perceived (Fig. 2, bottom left panel, y-axis) revealed a main effect of congruency ($F(1, 28.5) = 22.70, p < .001$), but importantly, no interaction between cue modality and congruency ($F(2,22990) = 1.93, p = .15$). Thus, unconscious cues in different modalities elicited equivalent congruency effects (AV: $.969 \pm .02$; A: $.988 \pm .02$; V: $.974 \pm .01$). Taken together, these findings indicate that congruency priming is present and modality-dependent when subjects

consciously perceive cues, but these effects disappear when cues are not perceived (Fig. 2, bottom left panel). Perhaps most importantly, and in line with the modeling results, when including in the analyses AV cues that were partially perceived (as either A or V; Fig. 2, bottom left panel dashed green arrow on y-axis and A/V percept on right panel), congruency effects suggest a graded pattern. Thus, effects are smallest when AV cues are not perceived (median of normalized congruency effect = $.96 \pm .02$; t-test comparing against 1.0, $t(28) = 3.56, p = .001$), intermediate when they are partially perceived ($.91 \pm .03, t(28) = 8.52, p < .001$), and largest when they are fully perceived ($.86 \pm .03; t(28) = 10.84, p < .001$; one-way ANOVA, $F(1.81, 50.67) = 30.53, p < .001$; see Fig. 2 bottom right panel A/V lying between “AV” and “none”).

4. Discussion

The vast majority of studies regarding perceptual consciousness have contrasted neural states in which participants report being aware versus unaware of stimuli in the environment (Baars, 2002; Koch, Massimini, Boly, & Tononi, 2016). These investigations agree that when a conscious report is made, neural ignition brings online a large neural network composed of early sensory areas and associative frontoparietal areas (Dehaene & Changeux, 2001). However, some (Aru, Bachmann, Singer, & Melloni, 2012; Pitts, Metzler, & Hillyard, 2014, 2012; Tsuchiya et al., 2015) have pointed out that dichotomizing reports based on the presence versus absence of consciousness does not directly address the question of phenomenal consciousness (i.e., “what does it feel like?”; Nagel, 1974), but rather focuses largely on the nature of the report (access consciousness; Block, 1995). Thus, it remains unclear whether phenomenal experience is graded, or categorical, as would be suggested by the neural findings pertinent to access consciousness.

Here we show that neural networks built in line with “all-or-none” theories of consciousness may nonetheless produce a graded relationship between the level of access and behavioral performance. We hypothesized that access to consciousness may be indexed by threshold crossings, as put forward in neuroimaging studies demonstrating differences in neural activation for aware versus unaware conditions (i.e., neural ignition; see Dehaene & Changeux, 2011 for review). On the other hand, phenomenal consciousness may be driven by a distinct operation – perhaps one more akin to decision-making processes (i.e., area under the activation function; integration operation). Such a view has been argued from a philosophical perspective (Block, 1995), which states that access and phenomenal consciousness are clearly related, but may be dissociable. In a neural network designed to encapsulate this hypothesis we show that multisensory task performance is best when cues meet an operational definition of consciousness in line with theory (i.e., threshold crossing; Dehaene & Changeux, 2011; Lamme, 2006), worst when the stimuli are operationally not accessed at all, and intermediate when they are only partially accessed. That is, while the networks are built in line with all-or-none indices of consciousness, they have the capacity to produce a graded pattern of reaction times – and thus potentially a graded pattern of

phenomenology (assuming an over-simplified identify relation between the latter two). We conceive of the modeling results not as evidence in favor of the fact that phenomenal consciousness is graded – but as demonstrating that GNW (Dehaene & Changeux, 2011) and LC (Lamme, 2006) and not incongruent with this possibility.

In a final step, we queried whether the performance of human observers would support a conceptual model where access to consciousness is based on threshold crossings (as demonstrated in neuroimaging; Dehaene & Changeux, 2011), yet phenomenal consciousness may still be graded. Support for this possibility came in the form of graded behavioral performance dependent upon conscious access to the priming cues. In more detail, we explicitly make the assumption that phenomenological richness is related to task performance. In these analyses, we employ both a direct measure of phenomenological experience (i.e., participant's reports of what they perceived) as well as an indirect one (i.e., the degree to which multisensory presentations elicited location congruency priming). Results demonstrate that the indirect measure places multisensory presentation experienced as unisensory closest to the arithmetic mean of full (audio-visual content) and null (no content) perception, and not closer to either of these extremes. These results highlight the fact that while the majority of studies on perceptual awareness may ask participants to categorize their conscious content in discrete bins (perceived *vs* not perceived), indirect measures of consciousness may suggest a graded pattern.

Taken together, the current findings argue that while the vast majority of previous research regarding perceptual consciousness has emphasized an abrupt transition between consciousness and absence of consciousness, this does not preclude the presence of a gradual phenomenology. Here we demonstrate via neural network modeling that current theories of consciousness based on “all-or-none” threshold crossings are not incongruent with the possibility for graded phenomenology, and show that indirect behavioral measures of consciousness suggest a graded pattern from no consciousness, to partial consciousness, and finally full consciousness. The current work is novel in that instead of indexing neural activity, we provide *in silico* evidence that a network built in line with “all-or-none” accounts of access consciousness may still support gradual phenomenology. However, to provide additional support for these findings, future work should perform neuroimaging in conjunction with the current psychophysical task in order to explicitly index the neural non-linearity we infer to be occurring (according to previous research and theory; Dehaene & Changeux, 2011) when participants report being aware (*vs* unaware) of the different cues. By combining the current computational-psychophysical approach with neuroimaging, the neural networks could putatively equally be extended in order to provide a grasp on neural processing. The models used here are admittedly simplistic, and are not intended to mimic actual neural processes; in fact, they do not even spike. The modeling effort here is a simple conceptual rather than neurobiological account suggesting that all-or-none threshold crossing within a node can nonetheless

engender graded patterns in the network as a whole (or in a different node). Further, it must be acknowledged that here we equate behavioral performance with the richness of phenomenal content, a widely held assumption (Lau & Passingham, 2006), yet nonetheless an assumption. Lastly, while we do consider the current results to argue against consciousness being categorical, it is true that only three locations in the phenomenological “state-space” were indexed here (i.e., full, intermediate, and null). These are sufficient to address whether consciousness is graded or categorical, but constitute a minimal case. Future research adding a third sense (e.g., touch) or graded levels of cue stimuli could be employed to draw a more detailed picture of the phenomenological landscape.

Declaration of conflicting interests

The authors declare no conflict of interest.

Funding

JPN was supported by a NSF GRF (14590), NF and OB were supported by the Swiss National Science Foundation and the Bertarelli Foundation. MTW is supported by the NIH (grants HD083211 and MH109225).

Open practices

The study in this article earned Open Materials and Open Data badges for transparent practices. Materials and data for the study are available at <https://osf.io/45gch/>.

CRediT authorship contribution statement

Jean-Paul Noel: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Nathan Faivre:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Elisa Magosso:** Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Writing - review & editing. **Olaf Blanke:** Project administration, Funding acquisition. **David Alais:** Conceptualization, Project administration, Funding acquisition, Writing - original draft, Writing - review & editing. **Mark Wallace:** Conceptualization, Project administration, Funding acquisition, Writing - original draft, Writing - review & editing.

Acknowledgments

The authors acknowledge Kahan Modi, Nicholas Lazzara, and Andrew Rosenzweig for help in data collection.

REFERENCES

- Andrillon, T., Poulsen, A. T., Hansen, L. K., Léger, D., & Kouider, S. (2016). Neural markers of responsiveness to the environment in human sleep. *Journal of Neuroscience*, 36, 6583–6596.
- Aru, J., Bachmann, T., Singer, W., & Melloni, L. (2012). Distilling the neural correlates of consciousness. *Neuroscience and Biobehavioral Reviews*, 36, 737–746.
- Bates, D., Machler, M., Bolker, B., & Walker, S. (2014). *Fitting linear mixed-effects models using lme4*. arXiv preprint arXiv:1406.5823.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2).
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436.
- Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., et al. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. *Science Translational Medicine*, 5, 198ra105.
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 1585–1599.
- Cleeremans, A., & Jimenez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. In R. M. French, & A. Cleeremans (Eds.), *Implicit learning and consciousness* (pp. 1–40). Hove, England: Psychology Press.
- Dehaene, S., & Changeux, J. P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70, 200–227.
- Dehaene, S., Naccache, L., Cohen, L., Bihan, D. L., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, 4, 752–758.
- Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biology*, 5, e260.
- Del Cul, A., Dehaene, S., & Leboyer, M. (2006). Preserved subliminal processing and impaired conscious access in schizophrenia. *Archives of General Psychiatry*, 63, 1313–1323.
- Deroy, O., Chen, Y. C., & Spence, C. (2014). Multisensory constraints on awareness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20130207.
- Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. (2007). Masking disrupts reentrant processing in human visual cortex. *Journal of Cognitive Neuroscience*, 19, 1488–1497.
- Faivre, N., Mudrik, L., Schwartz, N., & Koch, C. (2014). Multisensory integration in complete unawareness: Evidence from audiovisual congruency priming. *Psychological Science*, 1–11.
- Frässle, S., Sommer, J., Jansen, A., Naber, M., & Einhäuser, W. (2014). Binocular rivalry: Frontal activity relates to introspection and action but not to perception. *Journal of Neuroscience*, 34, 1738–1747.
- Gaillard, R., Dehaene, S., Adam, C., Clémenceau, S., Hasboun, D., Baulac, M., et al. (2009). Converging intracranial markers of conscious access. *PLoS Biology*, 7, e61.
- Gelbard-Sagiv, H., Faivre, N., Mudrik, L., & Koch, C. (2016). Low-level awareness accompanies unconscious high-level processing during continuous flash suppression. *Journal of Vision*, 16(1), 1–16. <https://doi.org/10.1167/16.1.3>.
- Grill-Spector, K., Kushnir, T., Hendler, T., & Malach, R. (2000). The dynamics of object-selective activation correlate with recognition performance in humans. *Nature Neuroscience*, 3, 837–843.
- Joglekar, M. R., Mejias, J. F., Yan, G. R., & Wang, X.-J. (2018). Inter-areal balanced amplification enhances signal propagation in a large-scale circuit model of the primate cortex. *Neuron*, 98, 222–234.
- Kayser, C., Petkov, C. I., & Logothetis, N. K. (2009). Multisensory interactions in primate auditory cortex: fMRI and electrophysiology. *Hearing Research*, 258, 80–88.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. *Perception*, 36(14), 1–16.
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, 17, 307. e21.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14(7), 301–307.
- Kuznetsova, A., Bruun Brockhoff, P., & Haubo Bojesen Christensen, R. (2013). *lmerTest: Tests for random and fixed effects for linear mixed effect models (lmer objects of lme4 package)*. r package version 2.0-0.
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10, 494–501.
- Lau, H. C., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences*, 103, 18763–18768.
- Leopold, D. A., & Logothetis, N. K. (1996). Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature*, 379, 549–553.
- Logothetis, N. K., & Schall, J. D. (1989). Neuronal correlates of subjective visual perception. *Science*, 245, 761–763.
- Nagel, T. (1974). What is it like to be a bat. *The Philosophical Review*, 83, 435–450. <https://doi.org/10.2307/2183914>.
- Noel, J. P., Ishizawa, Y., Patel, S., Eskandar, E., & Wallace, M. (2018). Leveraging non-human primate multisensory neurons and circuits in assessing consciousness theory. *BioRxiv*. <https://doi.org/10.1101/584516>.
- Noel, J. P., Simon, D., Thelen, A., Maier, A., Blake, R., & Wallace, M. T. (2018). Probing electrophysiological indices of perceptual awareness across unisensory and multisensory modalities. *Journal of Cognitive Neuroscience*, 30, 814–828.
- Noel, J. P., Wallace, M., & Blake, R. (2015). Cognitive neuroscience: Integration of sight and sound outside of awareness? *Current Biology*, 25(4), R157–R159. <https://doi.org/10.1016/j.cub.2015.01.007>.
- O'Callaghan, C. (2017). Grades of multisensory awareness. *Mind and Language*, 32, 155–181.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pitts, M. A., Martinez, A., & Hillyard, S. A. (2012). Visual processing of contour patterns under conditions of inattention blindness. *Journal of Cognitive Neuroscience*, 24, 287–303.
- Pitts, M. A., Metzler, S., & Hillyard, S. A. (2014). Isolating neural correlates of conscious perception from neural correlates of reporting one's perception. *Frontiers in Psychology*, 5, 1078.
- Quiroga, R. Q., Mukamel, R., Isham, E. A., Malach, R., & Fried, I. (2008). Human single-neuron responses at the threshold of conscious recognition. *Proceedings of the National Academy of Sciences of the United States of America*, 105, 3599–3604.
- Rees, G., Kreiman, G., & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews Neuroscience*, 3, 261–270.
- Sarasso, S., Boly, M., Napolitani, M., Gosseries, O., Charland-Verville, V., Casarotto, S., et al. (2015). Consciousness and complexity during unresponsiveness induced by propofol, xenon, and ketamine. *Current Biology*, 25, 3099–3105.
- Satterthwaite, F. E. (1946). An approximate distribution of estimates of variance components. *Biometrics Bulletin*, 2(6), 110–114.

- Schartner, M. M., Carhart-Harris, R. L., Barrett, A. B., Seth, A. K., & Muthukumaraswamy, S. D. (2017). Increased spontaneous MEG signal diversity for psychoactive doses of ketamine, LSD and psilocybin. *Scientific Reports*, 7, 46421.
- Schartner, M. M., Seth, A., Noirhomme, Q., Boly, M., Bruno, M. A., Laureys, S., et al. (2015). Complexity of multi-dimensional spontaneous EEG decreases during propofol induced general anaesthesia. *PLoS One*, 10, e0133532.
- Seely, J., & Chow, C. C. (2011). Role of mutual inhibition in binocular rivalry. *Journal of Neurophysiology*, 106(5), 2136–2150. <https://doi.org/10.1152/jn.00228.2011>.
- Shadlen, M. N., Hanks, T. D., Churchland, A. K., Kiani, R., & Yang, T. (2006). The speed and accuracy of a simple perceptual decision: A mathematical primer. In K. Doya, S. Ishii, R. Rao, & A. Pouget (Eds.), *Bayesian brain: Probabilistic approaches to neural coding* (pp. 209–237). Cambridge: MIT Press.
- Tong, F. (2003). Primary visual cortex and visual awareness. *Nature Reviews Neuroscience*, 4, 219–229.
- Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17, 450–461.
- Tononi, G., & Koch, C. (2015). Consciousness: Here, there and everywhere? *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 370, 20140167.
- Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. (2015). No-report paradigms: Extracting the true neural correlates of consciousness. *Trends in Cognitive Science*, 19, 757–770.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592.
- van Vugt, B., Dagnino, B., Vartak, D., Safaai, H., Panzeri, S., Dehaene, S., et al. (2018). The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science*. <https://doi.org/10.1126/science.aar7186>. eaar7186.
- Wickham, H. (2009). *Ggplot2: Elegant graphics for data analysis*. New York, New York: Springer.