

Conscious Perception: Time for an Update?

Moti Salti¹, Asaf Harel¹, and Sébastien Marti²

Abstract

■ Understanding the neural mechanisms underlying conscious perception has become a central endeavor in cognitive neuroscience. In theories of conscious perception, a stimulus gaining conscious access is usually considered as a discrete neuronal event to be characterized in time or space, sometimes referred to as a conscious “episode.” Surprisingly, the alternative hypothesis according to which conscious perception is a dynamic process has rarely been considered. Here, we discuss this hypothesis and its implications. We show how it can reconcile

inconsistent empirical findings on the timing of the neural correlates of consciousness and make testable predictions. According to this hypothesis, a stimulus is consciously perceived for as long as it is recoded to fit an ongoing stream composed of all other perceived stimuli. We suggest that this “updating” process is governed by at least three factors (1) context, (2) stimulus saliency, and (3) observers’ goals. Finally, this framework forces us to reconsider the typical distinction between conscious and unconscious information processing. ■

INTRODUCTION

The main thrust of the scientific study of conscious perception is to explain the cognitive processes that make a piece of external or internal information available for subjective experience. In a seminal paper, Crick and Koch suggested that characterizing the neural activity that underlies conscious perception would advance our understanding of subjective experience (Crick & Koch, 1990). In a typical neural correlate of consciousness (NCC) experiment, degraded stimuli are presented to participants while their neural activity is recorded (note that the majority of NCC studies were done in the visual domain). Participants are asked whether they perceived the target or not. NCCs are isolated contrasting brain activity associated with each of these conditions. With this contrast, researchers are aiming to identify the “where” and “when” of consciousness. Namely, the goal is to find the latency and brain location of a discrete neuronal event that gives rise to conscious perception (for a review, see Dehaene & Changeux, 2011). However, during the almost three decades in which this approach has been practiced, there has been no consensus regarding the boundaries between conscious and unconscious perception and consequently regarding the neural activity underlying conscious and unconscious perception. Here, we point out an alternative hypothesis, which should be considered as well: Conscious perception might correspond to a dynamic process involving a series of brain processing stages. This hypothesis suggests that

subjective experience relies on ongoing brain dynamics rather than on a discrete event. We propose that an updating and recoding mechanism guided by perceptual context, stimulus saliency, and perceiver’s goals underlie subjective experience. In the following, we selectively review the empirical results that led to this hypothesis, consider its implications, confront it with existing data, and draw testable predictions.

“EARLY” VERSUS “LATE” NCCS: AN INSOLUBLE DEBATE?

The main endeavor in the field of consciousness research is to isolate the minimal spatial and temporal brain activity that gives rise to conscious perception. However, to date, there is still no consensus regarding the activity that gives rise to consciousness. An activation likelihood estimation meta-analysis recently performed by Bisenius, Trapp, Neumann, and Schroeter (2015) yielded diffused activity that included clusters of activation in inferior and middle occipital gyrus; fusiform gyrus; inferior temporal gyrus; caudate nucleus; insula; inferior, middle, and superior frontal gyri; precuneus; as well as in inferior and superior parietal lobules. In the temporal domain, as the explored space is more restricted, the P3b component (appearing ~350 msec after stimulus onset) and a negative early component, usually referred to as visual awareness negativity (VAN, appearing ~200 msec after stimulus onset), stand out as potential NCCs (Koivisto & Grassini, 2016; Koivisto & Revonsuo, 2003). According to the taxonomy put forward by Aru, Bachmann, Singer, and Melloni (2012), genuine NCCs should be distinguished from (i) events that precede and support consciousness

¹Ben-Gurion University of the Negev, ²Université Paris-Saclay

“NCC-pr” and (ii) events that follow consciousness and are a consequence of it “NCC-co.”

The empirical results in the temporal and spatial domain are utilized to support two main hypotheses. According to the early entry hypothesis, (visual) conscious perception relies on early activity in striate and extrastriate areas and is reflected by the VAN ERP component. The P3b is only thought to reflect response-related processes that follow conscious perception. Using the inattentive blindness paradigm, Pitts and colleagues showed that the conscious perception of an expected stimulus irrelevant to the task is associated with a larger amplitude of the VAN rather than the P3b. Only once the target stimulus became task-relevant did the P3b reflect conscious perception. This result was consistent with earlier findings in studies examining the effects of expectation on conscious perception. These studies revealed that when a target stimulus is seen and unexpected, a large P3b component is observed. By contrast, when the same stimulus is perceived but expected, a difference is observed at an earlier latency (e.g., on the N2 [Koivisto & Revonsuo, 2008] or during the P2 components [Melloni, Schwiedrzik, Müller, Rodriguez, & Singer, 2011]), but no effect is observed in the P3b time window.

In contrast, the late access hypothesis asserts that conscious perception relies on higher-level activity that includes parietal and prefrontal cortices and is reflected by the P3b ERP component. Many studies in humans and monkeys showed that visual processing up to ~250 msec, and thus including the VAN, is preserved even when participants deny having seen a target stimulus (Marti, Sigman, & Dehaene, 2012; Del Cul, Baillet, & Dehaene, 2007; Sergent, Baillet, & Dehaene, 2005; Lamme, Zipser, & Spekreijse, 2002; Rolls, Tovée, & Panzeri, 1999; Vogel, Luck, & Shapiro, 1998; Kovács, Vogels, & Orban, 1995). The VAN would therefore manifest processes that precede and support consciousness, like attention, for instance (Busch, Fründ, & Herrmann, 2010). This dispute on the exact timing of NCCs thus seems insoluble, and in fact, in many experiments “early” and “late” correlates of consciousness coexist (Railo, Koivisto, & Revonsuo, 2011, but see Koivisto, Revonsuo, & Lehtonen, 2006).

A BROADER LOOK AT NCC IN THE GENERAL ERP LITERATURE

One way to untie what seems to be a Gordian knot is to examine the role of these components as well as their dynamics and relationship in the general ERP literature. The P3b is one of the most studied ERP components, it can easily be elicited with an oddball paradigm in which participants are presented with low-probability target items that are mixed with high-probability nontarget (or “standard”) items, both perfectly visible. The P3b is clearly observed after the presentation of a target but is barely measurable after standard items. Prominent and universal theories attribute the P3b to “context updating”

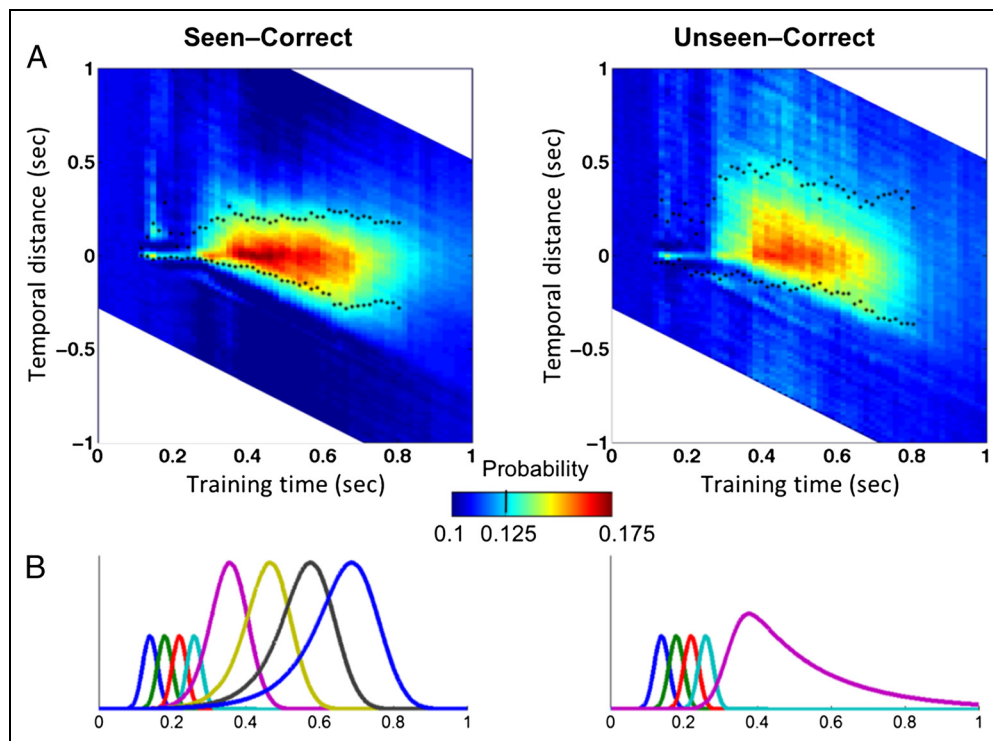
(Donchin & Coles, 1988), which indexes perceptual changes in one’s surroundings or the revision of a participant’s mental representation of the situation around him (or her) rather than acting as a call to action for a certain response, be it motor, verbal, or otherwise. The context-updating theory asserts that, after initial sensory processing, a comparison mechanism checks for changes from previous events in working memory. If there has been no change to any stimulus attribute, the current mental schema remains, evoking only the sensory evoked potentials (e.g., N1, P2, N2). Importantly, only if a novel stimulus (either in content or in characteristics) is detected and enough attentional resources are present, an “updating” of the stimulus representation occurs, eliciting the P3b potential (Polich, 2007).

A recent study conducted by Salti and colleagues may serve as a bridge between the interpretations of the P3b in the general ERP literature and what is observed in NCC studies. The authors used metacontrast masking to degrade the visibility of a target stimulus and trained multivariate classifiers to detect target features in M/EEG signals when it was consciously perceived or when it remained nonconscious (Salti et al., 2015). As expected, the detection of target features was superior in “Seen” trials than in “Unseen” trials 250–800 msec poststimulus. However, this improved classification did not stem from a recruitment of additional brain regions in the encoding of consciously perceived stimuli, nor was this information coded for a longer duration. Instead, the encoding of “Seen” and “Unseen” stimuli differed in their moment-to-moment dynamics, that is, the way information coding changed over time. An analysis aiming at testing the dynamics and transient stability of internal codes used the temporal generalization method (King & Dehaene, 2014). With this method, the efficiency with which an estimator trained at a certain time point t can generalize to other time points (t') is evaluated. If the representation is stable, the classifier should remain efficient even if applied at a different latency. This analysis showed that the encoding of “Unseen” stimuli was much more stable and exhibited a slow decay of ~350 msec. By contrast, the encoding of “Seen” trials consisted in a series of patterns of activity rapidly changing after ~160 msec. We suggest that these results reflect an MTM encoding mechanism: For stimuli reported as “Seen,” the related information was systematically recoded via a chain of short-lived processes. For stimuli reported as “Unseen,” however, this dynamic recoding was interrupted, resulting in a slow decay of the current stimulus coding (Figure 1).

CONSCIOUS PERCEPTION: AN MTM UPDATING PROCESS

Adopting the notion that an updating mechanism underlies conscious perception, we put forward the moment-to-moment (MTM) hypothesis: We suggest that the conscious perception of a stimulus is not a singular

Figure 1. The figure was reproduced from Figure 5 in Salti et al. (2015). The figure depicts the generalization of content decoding over time. Classifiers trained at a specific time were then tested on data from all other time points. (A) Average classification probability as a function of temporal distance from training time (positive or negative), with asterisks indicating classification endurance. Results show that classification endurance is larger for “Unseen” than for “Seen” trials. (B) Tentative model of a sequence of brain activations, which could yield the observed generalization matrices. Accordingly, encoding of “Seen” stimuli was systematically changing via a chain of short-lived processes. This dynamic, however, is different for “Unseen” stimuli. With “Unseen” stimuli, recoding was interrupted, resulting in a slow decay of the current stimulus coding.



event but a continuous process instead. According to this hypothesis, the stimulus-related information is encoded via a series of short-lived processes that allow its integration into the current perceptual context. Newly processed pieces of information are integrated to an existing perceptual context and therefore change it. We point out three putative parameters that might guide this updating as suggested by the existing data. The first parameter is the perceptual context to which a new piece of information is added. The context is a momentary representation of the world. If the processed information holds a very small change compared with the internal model, the induced update will be accordingly minimal (Melloni et al., 2011). The second parameter is the perceiver’s goal. The updating of a given stimulus within the context depends on its relevance to the perceiver’s momentary goal. Accordingly, irrelevant stimuli are less likely to be updated. An extreme example for that is the inattentional blindness paradigm in which a very salient stimuli is overlooked if it is irrelevant to the perceiver’s main task (Simons & Chabris, 1999). The examination of the ERPs elicited by the irrelevant stimuli revealed no measurable P3b (Shafto & Pitts, 2015). Considering that this component is thought to reflect the updating of the stimulus representation, an interpretation of these findings is that the stimulus representation is not updated because it is irrelevant to the participant’s task. Hence, the subjective experience of the irrelevant stimulus would fade away, thereby reducing the chance the participant will be able to report it. The last proposed

parameter is stimulus saliency, namely the physical properties of the stimulus that distinguish it from its background. In an oddball paradigm, for instance, at the single-trial level, more salient stimuli elicit a P3b component with larger amplitude (Teixeira et al., 2014). Saliency is intimately related to visibility: The more salient a stimulus is, the larger the chance it will be detected and seen. Del Cul and colleagues showed that the more visible a stimulus is, the larger the P3b is (Del Cul et al., 2007). These three parameters modulate the representations of the stimuli, prioritize them, and determine their availability for report, action, or other cognitive processes.

At the neuronal level, these parameters might then determine which neuronal substrates are recruited for the processing of the presented stimuli and may also affect the duration for which the stimuli echo in the system. Some of these contents stay activated in the perceptual system until they are exhausted whereas others fade away. By “exhausted,” we refer to a situation in which information of these contents served the goals for which they were prioritized. For the contents that endure, encoding evolves and changes to keep them coherent with the visual scene.

The duration of this continuous coding process or the lifespan of a stimulus in subjective experience is not constant and depends on the three parameters. This predicts that the temporal modulation of the context by the observer’s goals and stimulus saliency should affect the observer’s subjective experience. This leads to a key aspect of the hypothesis: Subjective perception evolves over

time. It does not correspond to a single neural event but instead is linked to a complex sequence of processes that vary with time. The underlying activity is not generic (in accordance with previous theories, see Dehaene, Sergent, & Changeux, 2003) but is specific to the stimulus and the general context it is updated in. During the processing of a stimulus and its integration into the current visual scene, multiple and distinct neuronal populations are progressively recruited, and each of these processing stages modifies subjective experience.

Previous studies showed that neural activity is more stable during conscious perception (Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene, 2015; Schurger, Pereira, Treisman, & Cohen, 2010). Although this finding might appear inconsistent with the proposed hypothesis, it is in fact not necessarily the case. Activity within an area of the brain can remain stable, whereas activity in other areas is changing over time. Salti et al. (2015), for instance, revealed that early stimulus-related activities located in the visual cortex were reactivated and remained stable over time, whereas series of processes were observed simultaneously in higher-order areas. This is consistent with other findings indicating that neural activity during conscious perception is rapidly evolving (Baria, Maniscalco, & He, 2017). The hypothesis of MTM updating is compatible with the idea of stability at the level of brain modules. However, it is not compatible with the idea of stability in neural activity over the entire brain as a mechanism for conscious perception.

Let's consider an example: A participant is presented with degraded stimuli and is instructed to produce a motor response according to the stimulus identity. The context is what the participant perceives right before the target stimulus is presented. The participant's goal is to identify the stimulus. The saliency of the stimulus is determined by its physical properties and the visual environment. The task is accomplished through a series of cognitive operations, from basic visual processing to motor planning and execution. For each of these steps, different populations of neurons are engaged. Every time a new population of neurons is recruited, the stimulus representation is updated and the subjective experience changes accordingly. Subjective experience of the stimulus may therefore vary in time, depending on how and for what purpose it is processed and as the successive stages of neuronal processes unfold. This does not suggest that the temporal evolution of subjective experience has to be gradual. In fact, drastic changes in subjective experience could suddenly occur if a multitude of distinct brain areas are recruited simultaneously.

A unique aspect of the MTM hypothesis is its attitude toward the conscious and unconscious perception dichotomy. The MTM hypothesis considers conscious perception as an ongoing process rather than a singular event in the brain. Accordingly, a stimulus is subjectively perceived as long as it is updated through a set of neuronal operations. If the processing chain is broken, then

the information is not updated, and it echoes in the last processor until it fades. This fading period could be considered unconscious as it might affect behavior but would not be subjectively experienced as the processed information would not be integrated into the current perceptual context. Hence, the MTM hypothesis does not exclude the possibility of unconscious perception, but it predicts that multiple forms of "degraded" perception can exist. These would result in distorted subjective experiences and variable effects on the observer's behavior. For instance, the participant could be partially aware of the stimulus, and this subjective experience may vary in time. In accordance with this prediction, a recent study showed that one feature of a visual stimulus (e.g., color) can be consciously perceived whereas another aspect of the very same object can remain invisible (e.g., shape; Elliott, Baird, & Giesbrecht, 2016). This prediction also fully agrees with the partial awareness hypothesis, which suggests that stimulus features of various complexities can be consciously accessed independently (Kouider, de Gardelle, Sackur, & Dupoux, 2010). The MTM hypothesis suggests that, because any part of the sequence of processes can fail, the remaining processes are not restrained to a specific level of complexity and can vary strongly from trial to trial.

Recent studies demonstrated that elaborate computations and manipulations could be done on stimuli even when they are rendered "Unseen." In these studies, participants had successfully categorized unseen pictures, words, and faces (Qiao & Liu, 2009; Dehaene et al., 2001) and managed to associate quantity to a number symbol (Van Opstal, Gevers, Osman, & Verguts, 2010; Naccache & Dehaene, 2001) and a valence to a word (Yeh, He, & Cavanagh, 2012; Kouider & Dehaene, 2007). It seems that even manipulations that demand long maintenance of perceived information like solving an arithmetic equation (Sklar et al., 2012) or maintaining information in working memory (Trübutschek et al., 2017) could be done on unconsciously perceived information (Dehaene, Charles, King, & Marti, 2014). These unexpected findings have dramatically extended the boundaries of unconscious processing, which were originally thought to be restrained to simple operations. The present hypothesis puts forward the possibility that the failure to report a stimulus means neither a complete absence of processing nor a complete absence of subjective experience. If stimulus updating is interrupted, the existent representation of the stimulus or task-related sequence of processes could still affect behavior in the absence of a clear subjective experience at the time of the report.

DISCUSSION

Theoretical Implications of the Current Hypothesis

The proposed hypothesis is results-driven, mainly inspired by NCC results in the temporal domain. The

outstanding recurrence of the P3b component in the NCC literature, its attribution to updating processes in the general ERP literature, and, finally, the absence of the P3b component in studies that used repetitive stimuli led us to suggest that the subjective experience of a stimulus relies on a continuous updating/recoding mechanism. Accordingly, the current hypothesis emphasizes the continuous and coherent nature of conscious perception as key aspects. The idea that subjective contents evolve over time has been proposed to explain the apparent discrepancy in the timing of the NCC (Bachmann, 2000). Albeit related, the hypothesis we are considering here is different: Here, the updating process is what makes and keeps the information conscious, whether the subjective experience is in an immature or highly detailed stage.

Multiple models of consciousness have been proposed over several years. Although they do not agree on the specific NCC and their role in the emergence of consciousness, they share a basic assumption, namely that conscious perception is a discrete event. Some propose a key role of early recurrent activity (<200 msec) between visual areas, which enables the exchange of information between areas and group the perceptual information (Koch, Massimini, Boly, & Tononi, 2016; Lamme, 2006). By contrast, other models attribute conscious perception to late activity (~the P3b time window) that involves prefrontal cortices (Lau & Rosenthal, 2011; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Dehaene et al., 2003). Moreover, even models that do agree on the timing of NCC disagree on the role of this activity for the emergence of conscious perception. The global neuronal workspace model, for example (Dehaene et al., 2003, 2006), asserts that this late component manifests the settling of brain activity into a temporary metastable state of global activity during which content-specific information is shared between distant areas. This metastable state is often characterized with a nonlinear transition between conscious and unconscious perception (Marti & Dehaene, 2017; Sergent et al., 2005). High-order theory of consciousness (Lau & Rosenthal, 2011) also propose a crucial role of frontal areas in consciousness but assigns it to a generic mental representation of the observer as being in a particular mental state (e.g., “I see a blue car”). Finally, the integrated information theory argues that the distributed activity (including parietal and frontal areas) represents an integration of various stimulus features (Tononi & Edelman, 1998). In this model, it is the causal properties of the brain, rather than a specific structure, that determines its level of consciousness.

The current hypothesis suggests that the question of which brain process has the unique property of supporting subjective perception is ill-posed and thus has no answer. The hypothesis highlights the importance of MTM encoding: Each brain area that is newly recruited modifies subjective experience. Accordingly, a stimulus is

updated with respect to an existing context, depending on its saliency and its importance in accomplishing one or more of the observer’s goals. Interestingly, rejecting the basic assumption that conscious perception is a discrete event allows accommodating the coexistence of some of the predictions of these models that seem contradicting. Conceptualizing conscious perception as an ongoing process puts the focus on the dynamics themselves, what they allow or deny, how they evolve over time, and what contributes or interferes with this evolution and so on.

In many aspects, this framing of subjective perception corresponds to the ideas that were elaborated in the more general framework of predictive coding. This framework suggests that the brain builds a model of the environment based on past sensory information and uses it to predict upcoming events. If the predictive model is accurate, the difference between the new sensory input and the internal model, the “prediction error,” is minimal. On the other hand, if the difference between the two is substantial, then the model is inaccurate and needs to be revised. Once the prediction error is minimized, then the brain has accurately generated a model inferring the causal structure of the external world (Friston, 2005, 2010). Predictive coding as a neuronal process was demonstrated with the MMN, a negative ERP component peaking 100–200 msec following a deviant stimulus in an oddball task. Wacongne, Changeux, and Dehaene (2012) for example, showed that the MMN reflected an active cortical prediction rather than passive synaptic habituation. Moreover, in minimally conscious patients, the presence of MMN in a passive auditory oddball task was a strong predictor for recovery (Morlet & Fischer, 2014). At the computational level, the MMN is thought to reflect a signal indicating a mismatch between the internal model and the external input (Näätänen, Paavilainen, Rinne, & Alho, 2007). According to the present hypothesis, the prediction error would be a prerequisite for the updating of the stimulus representation.

Considerations for Future Research on Consciousness

The MTM hypothesis opens new avenues for future research. In the following, we describe three of these topics that, so far, have been neglected in consciousness research:

- Previous empirical and theoretical studies focused on the perception of a single stimulus. However, in everyday life, multiple objects are present in the visual scene. Each object can enter or exit the observer’s consciousness, and this status can even fluctuate over time. Still, conscious experience appears to us as a continuous stream. The question of how such a stream is built has barely been explored.

The present hypothesis suggests (i) that the stimulus representation is continuously modified and updated and (ii) that this information is integrated into an existing perceptual context. Conceiving the stream of consciousness is therefore an integral part of the MTM hypothesis.

- The proposed hypothesis necessitates a respective methodological shift in the way neuroimaging data are handled. Neuroimaging techniques often rely on trial averaging to boost the signal-to-noise ratio. However, if indeed subjective experience relies on a chain of processes that can potentially break or be distorted at different stages from trial to trial, averaging would mask these subtle differences. The variability of conscious and nonconscious processes needs to be explored with great detail. Although technically challenging, the examination of single trials is a promising avenue for a better understanding of conscious and nonconscious processing.
- Another methodological shift concerns the critical “Seen”/“Unseen” contrast. According to the suggested hypothesis, it could no longer be considered as a gold standard in extracting the neuronal activity that gives rise to subjective experience. One could no longer expect a clear-cut comparison that would extract a neural substrate or pinpoint a specific moment in time where conscious and unconscious perception diverge. Patterns of activity isolated by contrasting seen and unseen trials are expected to vary between studies (at least to a certain extent) because the processing stages related to the performed task and the way they are disrupted on unseen trials are different. Future studies should instead focus on the content of subjective experience and how it relates to the evolution of stimulus representations in the brain. Recent advances in signal processing and machine learning techniques make this goal achievable, and some studies have started to track stimulus representations under conscious and nonconscious conditions (Marti & Dehaene, 2017; Trübutschek et al., 2017; Salti et al., 2015; Charles, King, & Dehaene, 2014; King et al., 2013).

CONCLUSION

We put forward a plausible parsimonious account for subjective perception. Based on empirical evidence, we propose that subjective experience is built through a dynamic updating process. Accordingly, subjective perception does not correspond to a single neural event but instead is linked to a complex sequence of processes that varies with time according to context, stimulus saliency, and participants’ goal. The suggested hypothesis also narrows the gap between subjective experience and neural processing. Previous models assumed that some brain processes had the unique capability to

produce subjective experience. The current hypothesis does not require such an assumption; it considers the possibility that successive processing steps directly contribute to subjective experience. This hypothesis bridges previously divergent interpretations of NCCs and provides new perspectives on the neural processes subtending subjective experience.

Reprint requests should be sent to Moti Salti, Zlotowski Center for Neuroscience, Ben-Gurion University of the Negev, Beer-Sheva, 84105 Israel, or via e-mail: motisalti@gmail.com.

REFERENCES

- Aru, J., Bachmann, T., Singer, W., & Melloni, L. (2012). Distilling the neural correlates of consciousness. *Neuroscience & Biobehavioral Reviews*, *36*, 737–746.
- Bachmann, T. (2000). *Microgenetic approach to the conscious mind*. Amsterdam: John Benjamins.
- Baria, A. T., Maniscalco, B., & He, B. J. (2017). Initial-state-dependent, robust, transient neural dynamics encode conscious visual perception. *PLoS Computational Biology*, *13*, e1005806.
- Bisenius, S., Trapp, S., Neumann, J., & Schroeter, M. L. (2015). Identifying neural correlates of visual consciousness with ALE meta-analyses. *Neuroimage*, *122*, 177–187.
- Busch, N. A., Fründ, I., & Herrmann, C. S. (2010). Electrophysiological evidence for different types of change detection and change blindness. *Journal of Cognitive Neuroscience*, *22*, 1852–1869.
- Charles, L., King, J.-R., & Dehaene, S. (2014). Decoding the dynamics of action, intention, and error detection for conscious and subliminal stimuli. *Journal of Neuroscience*, *34*, 1158–1170.
- Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, *2*, 263–275.
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, *70*, 200–227.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, *10*, 204–211.
- Dehaene, S., Charles, L., King, J.-R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, *25*, 76–84.
- 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.
- Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences, U.S.A.*, *100*, 8520–8525.
- Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biology*, *5*, e260.
- Donchin, E., & Coles, M. G. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences*, *11*, 357–427.
- Elliott, J. C., Baird, B., & Giesbrecht, B. (2016). Consciousness isn’t all-or-none: Evidence for partial awareness during the attentional blink. *Consciousness and Cognition*, *40*, 79–85.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *360*, 815–836.

- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*, 127–138.
- King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Sciences*, *18*, 203–210.
- King, J.-R., Faugeras, F., Gramfort, A., Schurger, A., El Karoui, I., Sitt, J. D., et al. (2013). Single-trial decoding of auditory novelty responses facilitates the detection of residual consciousness. *Neuroimage*, *83*, 726–738.
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, *17*, 307–321.
- Koivisto, M., & Grassini, S. (2016). Neural processing around 200 ms after stimulus-onset correlates with subjective visual awareness. *Neuropsychologia*, *84*, 235–243.
- Koivisto, M., & Revonsuo, A. (2003). An ERP study of change detection, change blindness and visual awareness. *Psychophysiology*, *40*, 423–429.
- Koivisto, M., & Revonsuo, A. (2008). Comparison of event-related potentials in attentional blink and repetition blindness. *Brain Research*, *1189*, 115–126.
- Koivisto, M., Revonsuo, A., & Lehtonen, M. (2006). Independence of visual awareness from the scope of attention: An electrophysiological study. *Cerebral Cortex*, *16*, 415–424.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*, 301–307.
- Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, *362*, 857–875.
- Kovács, G., Vogels, R., & Orban, G. A. (1995). Cortical correlate of pattern backward masking. *Proceedings of the National Academy of Sciences, U.S.A.*, *92*, 5587–5591.
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, *10*, 494–501.
- Lamme, V. A. F., Zipser, K., & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *Journal of Cognitive Neuroscience*, *14*, 1044–1053.
- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, *15*, 365–373.
- Marti, S., & Dehaene, S. (2017). Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature Communications*, *8*, 1955.
- Marti, S., Sigman, M., & Dehaene, S. (2012). A shared cortical bottleneck underlying attentional blink and psychological refractory period. *Neuroimage*, *59*, 2883–2898.
- Melloni, L., Schwiedrzik, C. M., Müller, N., Rodriguez, E., & Singer, W. (2011). Expectations change the signatures and timing of electrophysiological correlates of perceptual awareness. *Journal of Neuroscience*, *31*, 1386–1396.
- Morlet, D., & Fischer, C. (2014). MMN and novelty P3 in coma and other altered states of consciousness: A review. *Brain Topography*, *27*, 467–479.
- Näätänen, R., Paavilainen, P., Rinne, T., & Alho, K. (2007). The mismatch negativity (MMN) in basic research of central auditory processing: A review. *Clinical Neurophysiology*, *118*, 2544–2590.
- Naccache, L., & Dehaene, S. (2001). Unconscious semantic priming extends to novel unseen stimuli. *Cognition*, *80*, 215–229.
- Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, *118*, 2128–2148.
- Qiao, X., & Liu, Y. (2009). Adaptive weighted learning for unbalanced multicategory classification. *Biometrics*, *65*, 159–168.
- Railo, H., Koivisto, M., & Revonsuo, A. (2011). Tracking the processes behind conscious perception: A review of event-related potential correlates of visual consciousness. *Consciousness and Cognition*, *20*, 972–983.
- Rolls, E. T., Tovéé, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: Information analysis. *Journal of Cognitive Neuroscience*, *11*, 300–311.
- Salti, M., Monto, S., Charles, L., King, J.-R., Parkkonen, L., & Dehaene, S. (2015). Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. *eLife*, *4*, e05652.
- Schurger, A., Pereira, F., Treisman, A., & Cohen, J. D. (2010). Reproducibility distinguishes conscious from nonconscious neural representations. *Science*, *327*, 97–99.
- Schurger, A., Sarigiannidis, I., Naccache, L., Sitt, J. D., & Dehaene, S. (2015). Cortical activity is more stable when sensory stimuli are consciously perceived. *Proceedings of the National Academy of Sciences, U.S.A.*, *112*, E2083–E2092.
- Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, *8*, 1391–1400.
- Shafto, J. P., & Pitts, M. A. (2015). Neural signatures of conscious face perception in an inattention blindness paradigm. *Journal of Neuroscience*, *35*, 10940–10948.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, *28*, 1059–1074.
- Sklar, A. Y., Levy, N., Goldstein, A., Mandel, R., Maril, A., & Hassin, R. R. (2012). Reading and doing arithmetic nonconsciously. *Proceedings of the National Academy of Sciences, U.S.A.*, *109*, 19614–19619.
- Teixeira, M., Pires, G., Raimundo, M., Nascimento, S., Almeida, V., & Castelo-Branco, M. (2014). Robust single trial identification of conscious percepts triggered by sensory events of variable saliency. *PLoS One*, *9*, e86201.
- Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science*, *282*, 1846–1851.
- Trübtschek, D., Marti, S., Ojeda, A., King, J.-R., Mi, Y., Tsodyks, M., et al. (2017). A theory of working memory without consciousness or sustained activity. *eLife*, *6*, e23871.
- Van Opstal, F., Gevers, W., Osman, M., & Verguts, T. (2010). Unconscious task application. *Consciousness and Cognition*, *19*, 999–1006.
- Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1656–1674.
- Wacongne, C., Changeux, J.-P., & Dehaene, S. (2012). A neuronal model of predictive coding accounting for the mismatch negativity. *Journal of Neuroscience*, *32*, 3665–3678.
- Yeh, S.-L., He, S., & Cavanagh, P. (2012). Semantic priming from crowded words. *Psychological Science*, *23*, 608–616.