



UNIVERSITY
OF SKÖVDE

ERP AND MEG CORRELATES OF VISUAL CONSCIOUSNESS An Update

Master Degree Project in Cognitive Neuroscience
One year Advanced level 30 ECTS
Spring term 2019

Jona Förster

Supervisor: Antti Revonsuo
Examiner: Sakari Kallio

Abstract

Two decades of event-related potential (ERP) research have established that the most consistent correlates of the onset of visual consciousness are the early visual awareness negativity (VAN), a negative component in the N2 time range over posterior electrode sites, and the late positivity (LP), a positive component in the P3 time range over fronto-parietal electrode sites. A review by Koivisto & Revonsuo (2010) had looked at 39 studies and concluded that the VAN is the earliest and most reliable correlate of visual phenomenal consciousness, whereas the LP probably reflects later processes associated with reflective/access consciousness. However, an “early” vs. “late” debate still persists. This thesis provides an update to that earlier review. All ERP and MEG studies that have appeared since 2010 and directly compared ERPs of aware and unaware conditions are considered. The result corroborates the view that VAN is the earliest and most consistent signature of visual phenomenal consciousness, and casts further doubt on the LP as an ERP correlate of consciousness. Important new methodological, empirical, and theoretical developments in the field are described, and the empirical results are related to the theoretical background debates.

Keywords: neural correlates of consciousness, event-related potentials, event-related fields, visual awareness negativity, late positivity, Global Neuronal Workspace theory, Recurrent Processing theory

Contents

Abstract.....	1
List of Abbreviations.....	3
1. Introduction.....	4
2. Preliminaries: Concepts, Theories, and the “Early” vs. “Late” Debate.....	8
2.1. Concepts of Consciousness.....	8
2.2. Relevant Empirical Theories of Consciousness.....	10
2.2.1. The Global Neuronal Workspace Theory.....	11
2.2.2. The Recurrent Processing Theory.....	12
2.3. The “Early” vs. “Late” Debate in the Context of the Science of Consciousness.....	13
2.3.1. The “Early” Component: Visual Awareness Negativity.....	14
2.3.2. The “Late” Component: P3/LP.....	17
3. Electrophysiological Evidence for “Early” and “Late”: A Review of ERP Studies.....	19
3.1. Results of the Review by Koivisto & Revonsuo (2010).....	19
3.2. An Update on the Review by Koivisto & Revonsuo (2010).....	22
3.2.1. Objective Task-Performance as a Potential Confound.....	23
3.2.2. A New Contrastive ERP Paradigm: Inattentional Blindness.....	25
3.2.3. Task-Relevance: An Important Confound in Studies of Awareness.....	28
3.2.4. No-Report Paradigms: Wining the NCC from Its Prerequisites and Consequences.....	31
3.2.5. The Relationship Between Consciousness and Attention.....	33
3.2.6. The Role of Expectations and Prior Beliefs.....	37
3.2.7. Awareness: Graded, Dichotomous, or Both?.....	39
3.3. A Review of Recent MEG Studies on the Correlates of Visual Awareness.....	45
4. Conclusion.....	48
References.....	53

List of Abbreviations

EEG	Electroencephalography
ERF	Event-related field (in magnetoencephalography)
ERP	Event-related potential (in electroencephalography)
FFS	Fast feedforward sweep (in Recurrent Processing theory)
GNWT	Global Neuronal Workspace theory
LP	Late positivity (ERP component associated with visual awareness in many circumstances; occurs in the P3(00) time window)
MEG	Magnetoencephalography
NCC	Neural correlates of consciousness
PAS	Perceptual awareness scale (graded, usually four-point scale for assessments of the subjective quality of visual awareness)
RN	Reversal negativity (ERP component associated with perceptual reversals in bistable perception)
RPT	Recurrent Processing theory
SN	Selection negativity (ERP component associated with selective attention)
SOA	Stimulus onset asynchrony (the time between the onset of a stimulus and the onset of a subsequent stimulus, e.g., a mask; usually measured in milliseconds)
VAN	Visual awareness negativity (ERP component associated with visual awareness in many circumstances)

1. Introduction

In the 21st century, consciousness remains one of the most exciting and, at the same time, most elusive topics in science. After nearly a century of neglect on the part of behavioristic psychology and cognitive science, it was only in the late 1980s that consciousness became a respectable topic (Baars, 1988) that was taken up by the then-emerging new discipline of cognitive neuroscience (Crick & Koch, 1990). The dominating approach ever since has been the search for the so-called “neural correlates of consciousness” (NCC), that is (roughly), the set of parts of and processes in the brain that is minimally sufficient for bringing about consciousness.¹ However, despite some undeniable progress (Koch, Massimini, Boly, & Tononi, 2016a), the field still lacks agreement upon a common research program and therefore remains disunified (Revonsuo, 2006), or even “chaotic” (Mudrik & Lamy, 2007, p. 380), and it is characterized by a “wild variety of views, theories, and evidence” (Revonsuo, 2006, p. 5). This situation limits the prospects for empirical progress, since background assumptions and theories inevitably influence the interpretation of the evidence: not even state-of-the-art brain imaging technology can discover the NCC when there is not even disagreement about what consciousness *is*. On the other hand, it is reasonable to hope that empirical evidence can at least constrain the space of possible background theories, or that the explanation of certain persistent experimental results will require an inference to the best explanation (Lipton, 2004) that favors one theory of consciousness over its rivals.

The scope of this thesis is restricted to *visual* consciousness. Visual consciousness (or, synonymously, visual awareness)², is the most prominent “model system” (Revonsuo, 2006, p. 73–5) for the empirical study of consciousness, mostly because the visual system is one of the best-studied parts of the central nervous system, and because “the visual input is often highly structured yet easy to control” (Crick & Koch, 1998, p. 97; see also Crick & Koch, 1990; Revonsuo, 2006, p.

1 One well-known and oft-cited definition of “NCC” is the following: “In round terms, the NCC is the minimal set of neuronal events that gives rise to a specific aspect of a conscious percept” (Crick & Koch, 2003, p. 119).

2 In congruence with the literature reviewed here, I treat the terms “consciousness” and “awareness”, as well as the redundant “conscious awareness”, as synonymous; the sole reason I employ them all is for the sake of stylistic variation.

74). The standard technique is to present a subject with a visual stimulus and then have her report whether or not she has consciously perceived it. In the past decades, a number of ingenious methods have been developed which allow the researcher to establish a “minimal contrast” (Dehaene, 2014, p. 26) between “seen” (in the sense of “consciously perceived”) and “unseen” stimuli. The general strategy is to vary the actually presented physical stimulus as minimally as possible while generating differences in subjective awareness of the stimulus.³ For example, in *binocular rivalry*, the subjects’ eyes are constantly presented with a different stimulus each, which leads to alternating subjective percepts despite constant sensory input (Miller, 2013). In *subliminal perception*, subjects are presented with stimuli the strength of which varies around the subjective threshold of awareness (Dehaene, 2014). This is effectively achieved by *masking* the stimulus, i.e., by presenting a different, often nonsensical, stimulus shortly after the first (Bachmann & Francis, 2013).⁴ If applied at the right time, this has the effect of rendering the target stimulus invisible to the subject (the threshold for healthy, adult individuals is usually around 50 ms after presentation of the target (Dehaene, 2014, p. 40)).⁵

3 Dehaene (2014, p. 26) notes that this “minimal contrast” strategy is due to Bernard Baars (1988). It should be noted that this strategy is based on a strong, but experimentally indispensable simplification: since the subject is, in the most common setting, seated in front of a computer screen and presented stimuli on it, the contrast is in fact not between “no visual awareness at all” and “visual awareness of a certain stimulus and nothing else”, but really between “visual awareness of the screen, the background of the room, etc.” and “visual awareness of the screen, the background of the room, etc., *plus* visual awareness of the presented stimulus”. In other words, studies employing this standard setting are in reality looking for the correlates of a certain *difference* in the contents of visual awareness, namely, the difference that the presentation of the stimulus makes. It is difficult to escape the confines of this standard setting, and any such attempt will face heavy drawbacks in terms of the available stimulus material. The classic study by Johansson and colleagues, who attached light points to the joints of actors in a dark room in order to investigate biological motion perception (Johansson, 1973), may be considered an example of the kind of stimulus material that would be required for establishing a contrast between “nothing” and “something” (although Johansson did not study consciousness specifically, nor did he employ the contrastive method).

4 There are many different kinds of masking (see Bachmann & Francis, 2013, p. 4, for an overview). The most common is *backward masking*, where a mask (often, a random display, such as a scrambled image, or a black-white pattern) is presented some time after the target stimulus. The difference between the onset times of the stimulus and the mask is called the “stimulus onset asynchrony” (SOA). *Forward masking* is also possible, but usually less effective. In *metaccontrast masking*, the mask is presented after the stimulus, but (unlike standard backward masking), not in the same location as the stimulus. In *object substitution masking*, mask (often, “four dots surrounding the target”) and target stimulus are presented simultaneously, but “the mask persists after the offset of the target” (Koivisto, Kastrati, & Revonsuo, 2013, p. 224), preventing the target from becoming aware.

5 There are many further techniques besides binocular rivalry and subliminal presentation. These include, among others, so-called “continuous flash suppression”, and techniques of “preconscious presentation”, which exploit phenomena such as “change blindness”, “inattention blindness” and the “attentional blink.” See e.g. Dehaene &

Using these methods in combination with brain imaging, the NCC can be investigated. The underlying reasoning is that the physical similarity (or, ideally, identity) of the stimuli in seen and unseen trials should minimize (or even eliminate) the differences in preconscious sensory processes and leave only the difference between conscious and nonconscious processing for brain imaging to detect. With functional magnetic resonance imaging (fMRI), areas of heightened activation during consciousness can be identified; unfortunately, the temporal resolution of this method is inherently limited, because the blood oxygen level dependent (BOLD) signal picked up by fMRI occurs with a delay on the order of seconds relative to the event of interest (Gazzaniga, Ivry, & Mangun, 2019, p. 110), which makes it difficult to relate the results to the very small timing differences involved, e.g., in masking. With electroencephalography (EEG), the recording of brain-generated electrical activity from the scalp with electrodes, the temporal dynamics of the processes generating consciousness can be captured and made visible in event-related potentials (ERPs), an averaging technique that isolates the waveform of electrical activity time-locked to an event (usually, the onset of a stimulus, or a subject's response to it) (Luck, 2014); however, EEG does not permit a precise localization of the neural generators responsible for the scalp activation. Nevertheless, combining the strengths of these and other methods, and integrating constraints from anatomical, neuropsychological, computational, and other investigations, it is possible to theorize about the NCC, and this has led to a burgeoning empirical literature.

This literature is characterized by enduring controversies regarding both the location and the timing of consciousness. As to their location, it is currently debated whether the NCC are located in a "broad fronto-parietal network", or rather in a "posterior hot zone" (Koch et al., 2016a). With respect to timing, there is an older but still ongoing debate about the proper electrophysiological signature of the onset of consciousness. Some groups of researchers suggest that an ERP component

Changeux (2011, p. 201), for a description of these paradigms, and de Graaf & Sack (2015) for a taxonomy of these and other paradigms in consciousness research. Another, quite different body of information about the NCC comes from neuropsychology, the study of perceptual and cognitive deficits resulting from neurological damage. A good overview of the main findings of this discipline in relation to consciousness science is provided by Revonsuo (2010, Chapters 4 & 5).

called the P3, which occurs relatively late (about 300 ms after stimulus presentation) embodies this signature (e.g., Del Cul, Baillet, & Dehaene, 2007; Lamy, Salti, & Bar-Haim, 2008; Salti, Bar-Haim, & Lamy, 2012), while other groups have found earlier signatures, the most consistently observed of which is the so-called “visual awareness negativity”, arising already after about 200 ms (Koivisto & Revonsuo, 2010). (The VAN and LP are described in detail below, in sections 2.3.1. and 2.3.2.) These controversies are interrelated, as they speak to different philosophical and theoretical background assumptions, as well as to conceptual and methodological difficulties with the notion of the “NCC” (see section 2.1).

In this thesis, I will focus on the “early” vs. “late” controversy in order to shed light on this tangle of complications. In section 2, I will first introduce some important conceptual distinctions and the most relevant background theories, before I present that debate in more detail and situate it within the broader background of consciousness research. In section 3, the relevant empirical literature will be thoroughly examined. Given the importance of the temporal dynamics in the “early” vs. “late” debate, this section is essentially a comprehensive review of ERP and magnetoencephalography (MEG) studies on visual awareness of the past decade. A number of subsections will highlight important empirical, methodological, and conceptual developments that currently shape the field. Section 4 summarizes the results of the previous sections and relates them to the conceptual and theoretical questions described in section 2.

2. Preliminaries: Concepts, Theories, and the “Early” vs. “Late” Debate

2.1. Concepts of Consciousness

As noted above, any science of consciousness should be as clear as possible about what it is talking about, since “[o]ur philosophical commitments tend to guide the empirical science we make” (Revonsuo, 2006, p. 138)—even if they remain implicit. One of the most central and momentous distinctions is that between *phenomenal consciousness* and *access consciousness*, first

introduced by Block (1995). “Phenomenal consciousness” refers to subjective experience, the “what-it-is-like”-ness (Nagel, 1974) of our every sensation and thought. Owing to its private nature (everyone experiences it, but no one can directly know anything about *others*’ experience), phenomenal consciousness seems at the same time indubitable and difficult to tract scientifically. The standard way to establish its presence in experimental contexts is to rely on subjects’ reports on whether and what they have experienced, but recently, so-called no-report paradigms (see section 3.2.4.) have been explored as well (Tsuchiya, Wilke, Frässle, & Lamme, 2015). “Access consciousness”—or “reflective consciousness”, as some authors prefer to call it (Koivisto & Revonsuo, 2010; Revonsuo, 2006)—does not refer to subjective experience at all, but rather to a system’s central access to information, which then “can be used to control reasoning and behavior” (Block, 1995, p. 229). It is defined by Block as a purely cognitive, information-processing notion, and as such it is arguably possible even in systems that lack all phenomenal consciousness. However, some empirical scientists equate these two notions of consciousness and identify consciousness with global access to information (see section 2.2). In this thesis I will be concerned with the subjective aspect of experience, and hence with *phenomenal* consciousness—even if it should eventually turn out to be identical with reflective/access consciousness.

Another important distinction is that between the *state* (or *level*) of consciousness and the *contents* of consciousness (e.g. Hohwy, 2009; Koch et al., 2016). On the one hand, there is a difference between various states of consciousness (such as coma, the vegetative state, and the fully conscious wakeful state) that can be arranged on a continuum of different levels of awareness (e.g. Di Perri et al., 2014; Laureys, 2005). On the other hand, there are differences between the contents of consciousness *within* each state that can count as at least minimally conscious. Normally, our ever-changing subjective visual experience features a multitude of different qualities, objects, and scenes—i.e., of different contents. The search for NCC can focus on either of these two aspects of consciousness; accordingly, Koch et al. (2016a, p. 308) distinguish between the “full NCC” and the

“content-specific NCC”. Of course, the two aspects are interrelated: for example, a minimally conscious patient presumably enjoys only a greatly reduced form of subjective experience. This entanglement raises some methodological issues, the implications of which have recently been discussed by a number of authors (Aru, Bachmann, Singer, & Melloni, 2012; Bachmann & Hudetz, 2014; Hohwy, 2009). However, since the majority of studies pertaining to the “early” vs. “late” debate have focused on the contents of consciousness and consequently on the content-specific NCC, and have experimented only with healthy, awake, fully conscious subjects, I will ignore these complications and likewise focus exclusively on the contents of visual consciousness, assuming a “normal” state of alert, conscious wakefulness.

Even with this restriction in place, the concept of NCC raises some issues worthy of consideration. Presumably, consciousness has both *prerequisites* and *consequences* (Aru et al., 2012), that is, there are processes that regularly, or even necessarily, precede consciousness, and others that follow it. Now if the task is to isolate the NCC proper, rather than the correlates of these various processes, it is absolutely crucial to control for such confounds. A prime example of a confounding process are the varieties of attention (Koch & Tsuchiya, 2007, 2012); another one are the processes associated with subjects’ reports about their awareness (or lack thereof). Both these examples are important also in the context of the “early” vs. “late” debate, and will accordingly be discussed in this thesis.

A final clarification is in place: practically all research on the NCC tries to find the correlates of the *onset* of consciousness, that is, the moment when a given content first becomes conscious. This is hardly ever made explicit; a rare exception is the study by Andersen, Pedersen, Sandberg, & Overgaard (2016), who distinguish the “becoming conscious” of a stimulus from its “remaining conscious”, and clearly state that they are interested in the former. Of course, a precise moment in time when a stimulus suddenly becomes conscious is an—albeit useful—abstraction, and there is also an important tradition studying the “microgenesis” of consciousness (Bachmann,

2000). This is not unrelated to the question of different degrees of consciousness, which will be discussed in detail in section 3.2.7.

2.2. Relevant Empirical Theories of Consciousness

It is useful to distinguish empirical theories from research programs. “A research program is a set of background assumptions”, which “involve significant ontological and methodological commitments” (Revonsuo, 2006, p. 7). Insofar as it specifies which phenomena fall in its domain and what kinds of “substantive entities and processes” it assumes, a research program includes a *metaphysical* theory. In consciousness research, an example of a proposed research program would be Antti Revonsuo’s “Biological Realism”, powerfully laid out in (Revonsuo, 2006). *Empirical* theories, in the sense in which I use the term here, are more specific attempts at *explanation* of the phenomena in question: where, when, why, and how does consciousness arise from the workings of the brain? At the same time, empirical theories often remain implicit about their metaphysical commitments, so that it is not always obvious whether they are actually theories of the same phenomenon. The two most relevant empirical theories in the context of the “early” vs. “late” debate are the “Global Neuronal Workspace” theory (Dehaene, 2014; Dehaene & Changeux, 2011; Dehaene, Kerszberg, & Changeux, 1998; Dehaene & Naccache, 2001) and the “Recurrent Processing” theory (Lamme, 2000, 2006, 2010; Lamme & Roelfsema, 2000). This is because they are the empirical theories that make the most detailed predictions as to *when* consciousness should arise in the brain.⁶

6 There are, of course, other theories of consciousness, but they are usually hard to evaluate in terms of the early vs. late debate. For example, Zeki’s microconsciousness theory (Zeki & Bartels, 1999) postulates that each specialized visual area (such as V4 for color, V5/MT for motion, etc.) generates its own “microconsciousness”, which are somehow bound together into a “macroconsciousness” at a later stage, and into a “unified consciousness” later still. Taken at face value, Zeki’s theory thus seems to suggest *various* onsets of *various* “consciousnesses”; however, the relations between these hypothesized consciousnesses are so unclear (cf. (Revonsuo, 2010, pp. 217–218) that it is difficult to derive testable predictions from the theory. Another family of theories are the so-called “higher-order thought” (or HOT) theories of consciousness. HOT are philosophical in origin (see e.g. Gennaro (2018)), and as such at one remove from the more directly empirical questions of consciousness science. Nevertheless, since HOT theories identify conscious states with higher-order mental states, they are usually interpreted to imply prefrontal activation as necessary for consciousness (Lau & Rosenthal, 2011), and are often placed in the “late” camp (Andersen, Pedersen, Sandberg, & Overgaard, 2016; Koivisto & Grassini, 2016). However, this interpretation does not seem inevitable. Information integration theory (or IIT) is likewise ambiguous with respect to timing: on the

2.2.1. The Global Neuronal Workspace Theory

The Global Neuronal Workspace theory (GNWT) is a neuroscientifically updated version of the older, purely cognitive “Global Workspace theory” (Baars, 1988). According to GNWT, consciousness arises when the outputs of multiple specialized neuro-cognitive processes gain access to a dedicated “workspace” system that distributes them so that they become globally available to the organism; indeed, “[a]ccording to this theory, consciousness is just brain-wide information sharing” (Dehaene, 2014, p. 165). This “broadcasting”, and hence consciousness, is required by a number of higher cognitive processes, such as “durable and explicit information maintenance, novel combinations of operations, and intentional behavior” (Dehaene & Naccache, 2001, p. 9). Anatomically, the workspace system is thought to consist of long-distance connections distributed throughout and across prefrontal and parietal brain regions, predominantly arising from pyramidal neurons in cortical layers II and III, and additionally involving “the nonspecific thalamic nuclei, the basal ganglia, and some cortical nodes” (Dehaene & Changeux, 2011, p. 214). Crucially, in order to get access to the workspace system, inputs must cross a threshold for “ignition”, and this requires that they be selected by attention. According to GNWT, attention is thus “a necessary prerequisite for conscious perception” (Dehaene & Naccache, 2001, p. 8). It follows that the onset of consciousness cannot occur prior to the operations of attentional processes, and not before large-scale activation of the fronto-parietal workspace regions has occurred. This suggests the prediction that the earliest physiological correlates of consciousness should have a relatively large latency relative to stimulus onset—and indeed, the proponents of GNWT tend to regard the P3b component, which begins approximately 300–500 ms post stimulus, as the earliest ERP signature of

one hand, IIT suggests “that consciousness is associated with prefrontal activation”, and is sometimes placed in the “late” camp on these grounds (Andersen et al., 2016); on the other hand, Giulio Tononi has suggested that his Φ value should be sufficiently large for consciousness to arise already after 100 ms (Tononi, 2004), which would make IIT compatible with an “early” onset as well. It seems that, at present, the empirical relevance of IIT is greater for clinical investigations into the *level* of consciousness than for experimental studies of its *contents* (Casali, Gosseries, Rosanova, Boly, Sarasso, Casali, ...Massimini, 2013).

consciousness (e.g., Dehaene, 2014; Dehaene & Changeux, 2011). In the “early” vs. “late” debate, GNWT is thus clearly on the “late” side.

2.2.2. The Recurrent Processing Theory

The Recurrent Processing theory (RPT) was developed mainly by Victor Lamme (Lamme, 2000, 2006, 2010; Lamme & Roelfsema, 2000). Its main tenet is that purely feedforward activity in the brain is not sufficient for consciousness, but that feedback activity, or “recurrent processing” from “higher” to “lower” areas in the brain is necessary for awareness to occur. When a stimulus is presented to the visual system, it causes activation that spreads from the retina via the lateral geniculate nucleus of the thalamus to the primary visual cortex (V1), and from there to other cortical visual and motor areas, up to the most prefrontal areas. This forward spread of activation throughout the brain is completed within 100-150 ms after stimulus presentation, and is called the “fast feedforward sweep” (FFS) by Lamme. Already the FFS “enables a rapid extraction of complex and meaningful features from the visual scene, and lays down potential motor responses to act on the incoming information” (Lamme, 2006, p. 497). However, as soon as the feedforward activity reaches a given area, activation also spreads back to lower areas via feedback connections; and according to RPT, it is only during this recurrent processing that visual awareness of the stimulus can arise.⁷ In the latest formulation of his theory, Lamme distinguishes four stages of visual neural processing (Lamme, 2010). Both the FFS and recurrent processing are subdivided into “superficial” and “deep” (or “widespread”) stages. This subdivision is intended to capture the dimension of attention: only attended stimuli reach the “deep” stages of either feedforward or recurrent processing and make it to the higher levels of the visual hierarchy, while both unattended and attended stimuli reach awareness, provided that there is recurrent processing, be it superficial or deep. In other words, the dimension of attention is orthogonal to that of awareness, and both

⁷ RPT also yields a possible explanation of visual masking: a backward mask that is presented within 100 ms after stimulus presentation disables visual awareness because it interferes with recurrent processing, but it does not preclude certain cognitive feats, such as above-chance detection of stimuli, that can be achieved unconsciously, because it does not interfere with the FFS (Lamme & Roelfsema, 2000, p. 577)

attention and awareness can occur independently of each other. Thus, unlike in GNWT, attention is *not* a necessary prerequisite for awareness in RPT. And since recurrent processing does not have to wait until the FFS is completed, but begins as soon as the latter has reached the first areas that feature feedback connections, RPT predicts an “early” onset of consciousness, and situates itself accordingly in the early vs. late debate.

2.3. The “Early” vs. “Late” Debate in the Context of the Science of Consciousness

As the preceding discussion indicates, the “early” vs. “late” debate does not occur in a vacuum. The question of timing is systematically related to that of location. While neither theory strictly localizes consciousness in any particular area of the brain, GNWT stresses the crucial importance of frontally and parietally located areas in the process of “broadcasting”, while according to RPT, consciousness arises already with the occurrence of relatively local, occipito-temporal recurrent processing. Moreover, while not devoid of intrinsic interest, the answers to the “where” and “when” questions are ultimately instrumental in determining the answers to the far more interesting and vexing “hard” questions (Chalmers, 2007) of why and how consciousness arises in and from the brain. Note that RPT does not address these questions (it does not tell us why and how consciousness should arise from recurrent processing), and that GNWT answers them only at the price of identifying phenomenal consciousness with certain cognitive operations that seem to have no necessary connection to it. Nevertheless, if these problems are to find an answer at all, the search for the NCC, if successful, will probably make an important contribution to it. In the rest of this section, I will introduce the opposing positions and the ERP correlates at stake in some more detail, before I will undertake a thorough review of the relevant evidence from ERP studies in section 3. It should also be noted from the outset that the neat mapping of the “early” position on RPT and the “late” position on GNWT, while roughly true and certainly tempting, is really a simplification. In section 4, their relationship will be discussed in a more nuanced way.

The “early” vs. “late” controversy has largely focused on two candidates for ERP correlates of visual consciousness: the visual awareness negativity in the N2 time range, and the P3, to be introduced in the upcoming sections. It should be noted, however, that some researchers used to defend a third, “very early”, position in the debate. A fraction of the ERP studies reviewed by Koivisto & Revonsuo (2010) had identified an enhanced P1 component around 100 ms as correlated with awareness. However, many studies have failed to find awareness-related P1, and those that did were usually open to attention and arousal confounds (Railo, Koivisto, & Revonsuo, 2011). By now, even the Estonian researcher Talis Bachmann and his research group, who used to be the most outspoken proponents of the “very early” view (Aru & Bachmann, 2009; Bachmann, 2009), have abandoned that position (Rutiku, Aru, & Bachmann, 2016; Rutiku, Martin, Bachmann, & Aru, 2015; Rutiku & Bachmann, 2017), and now hold “that such early responses do not seem to be markers of direct conscious perception of near-threshold stimuli” (Rutiku et al., 2016). It seems, then, that the case on the “very early” position in the debate is closed with a negative answer.

2.3.1. The “Early” Component: Visual Awareness Negativity

Early on, the research group of Mika Koivisto and Antti Revonsuo had, over the course of several experiments, consistently identified a negative deflection around 200 ms after stimulus onset in the ERPs for conditions in which the subjects were aware of the presented stimuli, when compared to the ERPs for conditions in which they were not. This ERP difference can elegantly be visualized in a so-called difference wave (the waveform that results when the ERP of the “aware” condition is subtracted from the ERP of the “unaware” condition; see figure 1), and has been termed the “visual awareness negativity” (or VAN) by its discoverers (Ojanen, Revonsuo, & Sams, 2003; Wilenius-Emet, Revonsuo, & Ojanen, 2004). The VAN is typically observed over posterior scalp electrode sites, in particular occipital and posterior temporal ones (see figure 2). There are relatively few attempts to localize the cortical generators of the VAN; no doubt this is partly because EEG is notoriously ill-suited for localizing the origins of the waveforms it detects. MEG permits slightly

more reliable source reconstructions. An early MEG study identified awareness-related activity in the right lateral occipital cortex, indicating that VAN may be generated along the ventral visual stream (Vanni, Revonsuo, Saarinen, & Hari, 1996); consistent with this result, a newer MEG study found posterior activity between 190 and 350 ms “bilaterally on the lateral convexity of the occipito-temporal region, in the Lateral Occipital (LO) complex, as well as in the right posterior infero-temporal region” (Liu, Paradis, Yahia-Cherif, & Tallon-Baudry, 2012). Similarly, in a source reconstruction on the data of an ERP experiment on awareness (Koivisto, Kainulainen, & Revonsuo, 2009) with low resolution electromagnetic tomography (LORETA), Koivisto & Revonsuo found “that contralateral occipital and temporal lobes were sensitive to the manipulation of awareness” (Koivisto & Revonsuo, 2010, p. 931); and another group, using local autoregressive averaging (LAURA), again identified the lateral-occipital complex in the ventral stream as the source of VAN (Pitts, Martínez, & Hillyard, 2011). The converging evidence from different source localization techniques, performed over data from a variety of experiments using different paradigms, permits a relatively confident judgment as to the occipito-temporal origin of the VAN.

More recent studies have hypothesized and shown that the VAN is lateralized, i.e., its amplitude is considerably stronger over the hemisphere contralateral to the side of the visual field where the stimulus is presented (Eklund & Wiens, 2018; Koivisto & Grassini, 2016). Earlier studies found the VAN to be strongest sometimes over the right hemisphere (e.g., Wilenius & Revonsuo, 2007) and sometimes over the left hemisphere (e.g., Koivisto et al., 2008). In these studies the stimuli were shown in the center of the screen rather than lateralized; a possible explanation for the hemispheric variation is that the stimuli might have overlapped a larger part of one half of the visual field than of the other.

Regarding the timing of VAN, the 200 ms are a rough value; more precisely, the *onset* of VAN (i.e., the earliest time point when the ERP of the “aware” condition becomes statistically significantly different from that of the “unaware” condition) often begins already after 100 ms,

while its *peak latency* usually lies between 200 and 250 ms post-stimulus. Under some conditions, considerably later VAN onset and peak latencies have also been observed (Koivisto & Revonsuo, 2010); this often happens in studies using low-contrast stimuli (Wilenius & Revonsuo, 2007) and, more generally, low stimulus visibility (Railo & Koivisto, 2009b). In many studies by other research groups, the VAN has been called “N2”, because it occurs during the second large negative deflection visible in the ERP. Several other ERP components regularly occur in the N2 time range under certain conditions, such as the face-related N170, the attention-related N2pc and selection negativity (SN), and the reversal negativity (RN), which is associated with perceptual reversals during bistable perception (Koivisto & Revonsuo, 2010). The term “visual awareness negativity” is intended to refer to only the part of the N2 that is specifically related to visual awareness. Evidently, it is of great importance to experimentally isolate the VAN from all the other N2 components, and thereby to confirm that it is not a prerequisite of consciousness, or some regularly co-occurring process. For a while, the research group around Michael A. Pitts in the USA preferred to refer to the “N2”, because they felt that this isolation has not yet been achieved to a sufficient degree (Pitts et al., 2011; Pitts, Padwal, Fennelly, Martínez, & Hillyard, 2014); however, in spite of persisting doubts, by now they have adopted the term “VAN” (Pitts, Metzler, & Hillyard, 2014; Shafto & Pitts, 2015).

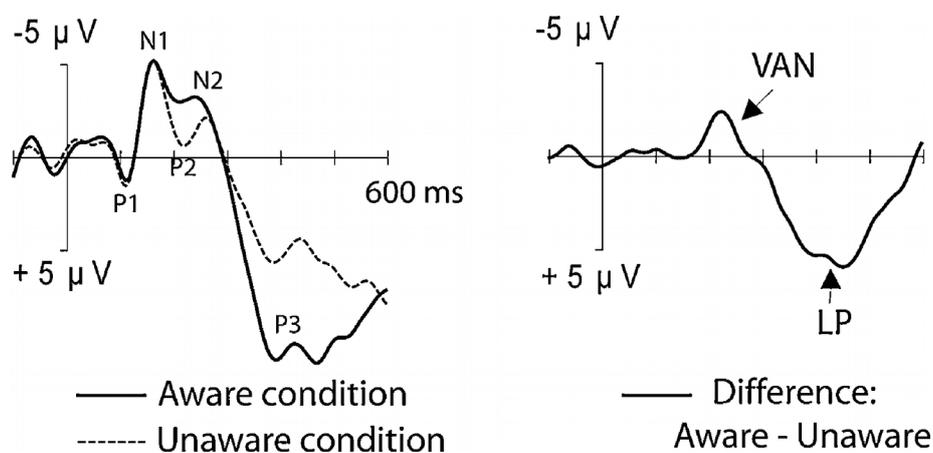


Figure 1: Left: the typical time course of an ERP over occipital sites in response to a visual stimulus; right: the difference wave resulting from subtraction of the unaware from the aware condition. Reprinted from *Neuroscience and Biobehavioral Reviews* 34, Koivisto & Revonsuo, *Event-related potential correlates of visual awareness*, p. 923, 2010, with permission from Elsevier.

2.3.2. The “Late” Component: P3/LP

The second candidate for an awareness-related ERP component is the P3 (i.e., the third positive peak clearly visible in the ERP), often also called the “P300” (because it peaks around 300 ms after stimulus onset). More precisely, the candidate is the *difference* in P3 amplitude between aware and unaware conditions (the amplitude is typically greater in aware conditions), often called the “late positivity” (or LP; see figure 1). Like the N2, the P3 in fact denotes a whole family of components occurring in the same time range, each of which occurs under specific conditions (and fails to occur under others). Although the P3 has been discovered already more than 50 years ago, both its cerebral origins and the function or functions it reflects have not been clearly determined. It has become clear, however, that the P3 reflects several different cognitive processes, and therefore in all likelihood has multiple cerebral sources (Picton, 1992; Polich, 2007). A clear distinction can be made between a P3a component with a frontal scalp distribution, peaking around 250 ms after stimulus onset, and more posterior P3b component which peaks around 350 ms (Picton, 1992). A recent LORETA study localized the P3a generators “in cingulate, frontal and right parietal areas”, and the P3b generators in “bilateral frontal, parietal, limbic, cingulate and temporo-occipital regions” (Volpe et al., 2007, p. 220). Hippocampal sites are also clearly involved (Ludowig, Bien, Elger, & Rosburg, 2010). As the amplitude of the P3 varies with stimulus probability, it has been associated with context-updating in working memory (Donchin, 1981), and also, quite early on, to the transfer of information to ““controlled” (“conscious” or “aware”) processing” (Picton, 1984, p. 174). Today, the P3a component is commonly associated with bottom-up, stimulus-driven attention mechanisms (Polich, 2007) that can occur both consciously and nonconsciously (Dehaene & Changeux, 2011), while the P3b is related to “the effortful processing of task-relevant events” (Volpe et al., 2007), and to the memory processing of such events (Polich, 2007). Proponents of the “late” camp view the P3b as strongly correlated with subjective awareness, and as a reliable signature of consciousness (Dehaene, 2014; Dehaene & Changeux, 2011; Del Cul et al., 2007;

Naccache, Marti, Sitt, Trübutschek, & Berkovitch, 2016; Sergent, Baillet, & Dehaene, 2005; Sergent & Naccache, 2012).

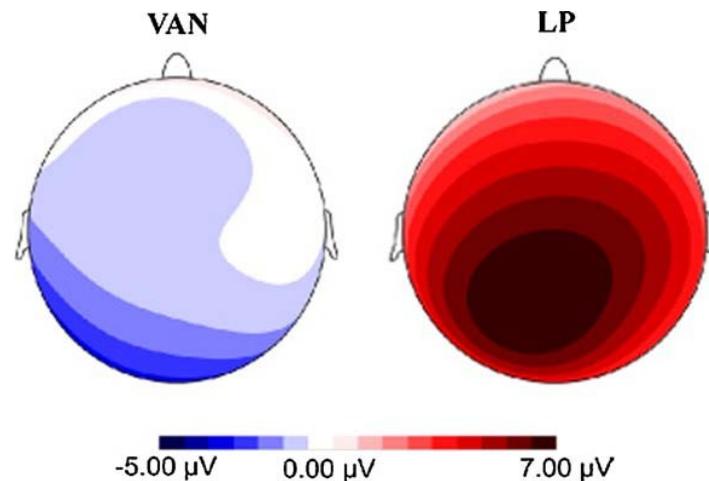


Figure 2: Typical scalp distribution of the VAN and LP components. Reprinted from *Neuroscience and Biobehavioral Reviews* 34, Koivisto & Revonsuo, Event-related potential correlates of visual awareness, p. 928, 2010, with permission from Elsevier.

It should be noted that the proponents of the “early” view usually also report an enhanced P3 (i.e., an LP) in addition to the VAN (Koivisto & Revonsuo, 2010), so the issue is not whether or not it occurs, but whether it is correlated specifically with consciousness as opposed to certain post-perceptual processes, such as decision-making, or report-related processes. Koivisto and colleagues view these latter processes as aspects of reflective/access consciousness, and propose to regard the VAN as the correlate of *phenomenal* consciousness, and the LP as the correlate of *reflective/access* consciousness. In other words, in their view, there are *two* NCC, and VAN is the earlier of the two. In light of the recent debate about the prerequisites and consequences of consciousness, the issue may be reframed, so that the question with respect to the LP (specifically, the P3b) becomes whether it reflects a *proper* NCC or a regular *consequence* of consciousness. Conversely, the “late” proponents have not always found VAN (Lamy et al., 2008; Salti et al., 2012), leaving the possibility that VAN reflects some non-necessary *prerequisite* of awareness rather than awareness itself. In the next section, the relevant ERP evidence pertaining to these questions will be thoroughly examined.

3. Electrophysiological Evidence for “Early” and “Late”: A Review of ERP Studies

In the past two decades, a large body of ERP studies on the question of the timing of consciousness has accumulated. In 2010, a review summarized the findings that had appeared until then, reporting for each study whether they had found NCC in the “very early” (enhanced positivity around 100 ms after stimulus onset), the “early” (enhanced negativity around 200 ms), and the “late” (enhanced positivity around 400 ms) time range (Koivisto & Revonsuo, 2010). In the following two sections, I will first summarize the results and open questions of that review, and then provide an update by examining all relevant ERP studies that appeared since 2010. Since the “very early” view is not at issue anymore, this update will focus exclusively on the remaining “early” and “late” time ranges. Like in the earlier review, other electrophysiological measures besides ERPs, like time-frequency analyses, will not be taken into account.

3.1. Results of the Review by Koivisto & Revonsuo (2010)

The review by Koivisto & Revonsuo (2010) included ERP studies from 1999 onward that compared the ERPs of aware and unaware conditions across a variety of paradigms, namely “different forms of masking, contrast level, attentional blink, change blindness, and bistable perception” (Koivisto & Revonsuo, 2010, p. 923). Of 39 reviewed studies, 13 reported enhanced very early negativity in the P1 range, 32 reported enhanced negativity in the N1-N2 range (VAN), and 36 reported enhanced late positivity (LP, or enhanced P3) in at least one aware-unaware comparison condition. However, only 29 studies reported P3 for *all* aware conditions, while 30 studies reported VAN for all aware conditions (not counting the studies reporting “attenuated” findings in one or more conditions). Going by the sheer numbers, VAN is thus slightly more reliable as an NCC than LP. Apart from an experiment on bistable perception that was not further discussed in the review, the only study that reported lack of VAN in one of its aware conditions was an experiment that attempted to disentangle awareness from spatial attention, and found that awareness

of a stimulus seems to depend on the presence of spatial attention to the region in which the stimulus occurs (Koivisto et al., 2009). The seven studies that failed to find LP for one or more of the aware conditions were all from a series of experiments by Koivisto and colleagues that controlled for various aspects of attention, and some other possible confounds. While VAN was found across the board, Koivisto, Revonsuo, & Salminen (2005) found only attenuated LP when controlling for object-based attention; Koivisto, Revonsuo, & Lehtonen (2006) found no LP in the local attention condition when controlling for the scope (local/global) of attention; Koivisto, Lähteenmäki, Sørensen, Vangkilde, Overgaard, & Revonsuo (2008) found only attenuated LP for low-contrast stimuli when they used random stimulus onsets; Koivisto & Revonsuo (2007) and Koivisto et al. (2009) found no LP for either target or non-target stimuli that were not spatially attended; and Koivisto & Revonsuo (2008b) found no LP when controlling for selective feature-based attention. All but one of these experiments used masking paradigms (one used low-contrast stimuli instead). In a combined attentional blink/repetition blindness experiment, Koivisto & Revonsuo (2008a) found no P3 difference between recognized (aware) and unrecognized (unaware) targets. Moreover, in some of the studies that reported LP but no VAN, the explanation is that the authors simply did not bother to look in the N2 time range, while others failed to differentiate between the correct localization of a stimulus and awareness of it (in particular, Babiloni, Vecchio, Miriello, Romani, & Rossini, 2006; Lamy et al., 2008). Since the VAN is a relatively small deflection compared to the large P3, a lack of statistical sensitivity may be another reason for these null findings (Koivisto & Revonsuo, 2010, p. 927).

Overall, the 2010 review seems to point towards the VAN as the earliest reliable ERP correlate of consciousness. Besides the mentioned paradigms, VAN was also found in studies using low-contrast stimuli (e.g., Ojanen et al., 2003; Pina & ffytche, 2003), even when controlling for physical stimulus differences (Wilenius & Revonsuo, 2007); in studies of change blindness (e.g., Koivisto & Revonsuo, 2003, 2005); in metacontrast masking (Railo & Koivisto, 2009b) and in

studies of bistable perception (e.g., Kaernbach, Schröger, Jacobsen, & Roeber, 1999). Indeed, even two of the most important studies carried out by the “late” camp found VAN (Del Cul et al., 2007; Sergent et al., 2005), but they concluded that it cannot reflect consciousness because, unlike the P3, it did not follow the characteristic non-linear all-or-none pattern in activation increase that GNWT predicts for stimuli that make it beyond the threshold for “global ignition”. Of course, this reasoning presupposes what it purports to show, namely, the truth of GNWT, or at least, of one of its central predictions (Koivisto & Revonsuo, 2010). Furthermore, the results of Del Cul et al. (2007), who used a quasi-continuous variation of stimulus onset asynchrony (SOA), were probably afflicted by floor and ceiling effects (Railo & Koivisto, 2009b, p. 795; Railo et al., 2011, p. 979), while those of Sergent et al. (2005) “might have artificially dichotomized the visibility ratings” (Railo & Koivisto, 2009b, p. 795).

As mentioned, a number of confounds have been addressed in the literature on VAN and P3/LP as correlates of visual awareness until 2010. Besides controls for physical stimulus differences and the effects of masks in various ways (Del Cul et al., 2007; Koivisto et al., 2009; Koivisto & Revonsuo, 2008b; Railo & Koivisto, 2009b, 2009a; Wilenius & Revonsuo, 2007; Wilenius-Emet et al., 2004), the series of experiments on attentional confounds of VAN is particularly noteworthy. VAN has been successfully dissociated from the SN component associated with feature- and object-based attention, and it has been shown that VAN is independent from the scope of attention, while the allocation of spatial attention to the stimulus region seems necessary for VAN to occur (see Koivisto et al., 2009, for a review). Moreover, there are some grounds for assuming that VAN is distinct from the attention-related N2pc (Koivisto & Revonsuo, 2010, pp. 926–927).

3.2. An Update on the Review by Koivisto & Revonsuo (2010)

This new review uses the same inclusion criteria for studies as did the old one: all ERP studies that have appeared since 2010 and directly compared target-locked ERPs of aware and unaware conditions are included. Table 1 lists all reviewed studies, sorted by the type of

Table 1

Results of the Review of ERP Studies for VAN (N2 range) and LP (P3 range)

Manipulation	Study	Enhanced early negativity (N2 range)	Enhanced late positivity (P3 range)
Contrast	Chica et al. (2010)		Yes
	Eklund & Wiens (2018)	Yes/Yes	Yes/Yes
	Koivisto et al. (2016)	Yes/Yes	Yes/Attenuated
	Koivisto & Grassini (2016)	Yes	Yes
	Koivisto et al. (2017)	Yes/No	Yes/Yes
	Koivisto et al. (2018)	Yes/Yes/Yes	Yes/Yes/Attenuated
	Melloni et al. (2011)	Yes/Yes	Yes/No
	Rutiku et al. (2016)	Yes/No*	Yes/Yes*
	Tagliabue et al. (2016)	Yes**	Yes**
	Ye & Lyu (2019)	Yes/Yes	Yes/Attenuated
Masking	Babiloni et al. (2016)	No/No/Yes	Attenuated/No/No
	Davoodi et al. (2015)	No/No	Yes/Attenuated
	Del Zotto & Pegna (2015)	Yes/Yes	Yes/Attenuated
	Fu et al. (2017)	Yes**	Yes**
	Jimenez et al. (2018)	Yes**	Yes**
	Koivisto et al. (2013)	Yes/Attenuated	Yes/Attenuated
	Pitts, Metzler, et al. (2014)	Yes/Yes	Yes/No
	Rutiku et al. (2015)	Yes/Yes	Yes/Yes
	Salti et al. (2012)	No	Yes
Inattentional Blindness	Pitts et al. (2011)	Yes/Yes	Yes/No
	Schelonka et al. (2017)	Yes/Yes	Yes/No
	Shafto & Pitts (2015)	Yes/Yes	Yes/No
Attentional Blink	Batterink (2012)	No/No	No/Yes
	Weller et al. (2019)		
Change Blindness	Scrivener et al. (2019)	No/No	Yes/Yes
Other	Boncompete & Cosmelli (2018)	No	Yes
	Pitts, Padwal, et al. (2014)	Yes/Yes/Yes/Yes	Yes/No/Yes/Attenuated

Note. For each study, I report which of the two relevant ERP correlates (VAN and LP) have been found. In the case of two or more entries, the studies have employed more than one experimental condition. *) The study by Rutiku et al. (2016) used over 70 different stimuli, and carried out 100 contrastive analyses in total, of which 81 found VAN, and 100 found LP. **) The studies by Fu et al. (2017), Jimenez et al. (2018), and Tagliabue et al. (2016) have compared different Perceptual Awareness Scale (PAS) contrasts (see section 3.2.7.); the amplitudes of the components they found varied with PAS rating (either gradually or dichotomously).

manipulation used, and reports their findings in the N2 range (around 200 ms after stimulus onset) and the P3 range (around 400 ms). A “yes” entry means that activity in the respective time range was found and reported; a “no” entry means that no activity in the respective time range was found and reported; and no entry (an empty field) means that the respective time range was not

investigated, and/or no results for that time range were reported. “Attenuated” means that the VAN or LP amplitude varied in one or more conditions, indicating an experimentally induced modulation of the respective component. Special cases that are difficult to fit into this framework are marked with one or more asterisks (*), and more information about them is provided in the legend of table 2. In total, 27 studies were reviewed, of which 16 found VAN and 12 found LP in all tested aware conditions (again, not counting the studies reporting “attenuated” findings). The result corroborates the one reached by Koivisto & Revonsuo (2010) that VAN is a more reliable ERP correlate of visual awareness than LP. At the same time, the informativeness of these numbers, and of table 1, is limited, as they conceal important new questions that have emerged in ERP research on visual consciousness, and new approaches that have been explored to tackle them. Consequently, table 1 glosses over important differences between the studies. Below, these will be discussed in detail, and the new developments will be described. Furthermore, a number of MEG studies highly relevant to the “early” vs. “late” have been carried out over the past ten years. They are listed in a separate table 2, and likewise will be reviewed below (in section 3.3.).

3.2.1. Objective Task-Performance as a Potential Confound

The study by Lamy et al. (2008) had drawn attention to an important confound that may have distorted earlier results: previous studies had failed to control for objective task-performance, and not taken into account the factor of whether participants performed a given task correctly or not. In effect, mixing up “unaware-correct” and “unaware-incorrect” trials in the analysis can lead to drastically varying baselines of the “unaware” condition, and consequently of the magnitude of its difference to the “aware” condition. Potentially, this can dramatically distort the results, leading to false positive and/or false null findings in different experiments. Lamy et al. (2008) controlled for this confound by comparing only ERPs for “aware-correct” and “unaware-correct” trials, while discarding, for the purposes of this comparison, all trials on which subjects had performed incorrectly. Since they found LP but no VAN for the “aware” condition when comparing ERPs in

this way, they concluded that previous VAN findings were probably due to the failure to control for task-performance. In a follow-up study, Salti et al. (2012) reproduced this result (LP but no VAN), while simultaneously controlling for the level of confidence of the subjects' awareness report.

This challenge has been taken up by at least two studies. Koivisto & Grassini (2016) suspected that the reason for the null findings of Lamy et al. (2008) and Salti et al. (2012) was a lack of sensitivity for the comparatively small VAN component. To make their study more sensitive, they used a larger stimulus, low-contrast Gabor patches instead of masking (in order to eliminate any possible noise introduced by the masks), and they controlled for lateralization. Because the VAN is hypothesized to reflect local recurrent processing in visual cortex, it is assumed to emerge in the hemisphere contralateral to the visual field in which the stimulus occurs. With these modifications in place, Koivisto & Grassini (2016) found both VAN and LP for the "aware" condition while controlling for task performance. A signal detection analysis revealed that LP was reduced when subjects adopted a conservative response criterion in the forced-choice localization task that the experiment employed, suggesting that the P3 is associated with post-perceptual processes such as "participants' meta-cognitive evaluations concerning their awareness" (Koivisto & Grassini, 2016, p. 242) rather than with awareness *per se*. A preregistered further study tried to determine which of the three modifications in this experiment had been the one that was successful in increasing the sensitivity sufficiently for VAN to be detected. Eklund & Wiens (2018) closely followed Koivisto & Grassini (2016), but used two different Gabor patch sizes. They found VAN and LP for aware trials, and no evidence for an effect of stimulus size. Presumably, then—apart from the already mentioned possibility of a lack of sensitivity—either unwanted mask effects or the lack of control for lateralization were responsible for the VAN null findings in the studies of Lamy, Salti, and colleagues. These possibilities could be investigated with new experiments, and via reanalysis of their data with hemisphere as an additional factor.

3.2.2. A New Contrastive ERP Paradigm: Inattentional Blindness

An interesting novelty relative to the 2010 review is the introduction of the inattentional blindness paradigm, which has been adapted for use in ERP studies by Michael Pitts and colleagues. In inattentional blindness, subjects can miss even quite salient stimuli right in the middle of their visual field, if their attention is focused on a distracter task (Mack & Rock, 1998). In a classical experiment, subjects failed to detect even a man dressed up as a gorilla walking into the scene and pounding his chest, when they were distracted with counting the passes of a handball team (Simons & Chabris, 1999). The problem with inattentional blindness as a paradigm in ERP research is that in ERP, many trials have to be averaged in order to distill a visible, statistically significant effect. However, in inattentional blindness, once a subject has become aware of the “hidden” stimulus, she will notice it in the future. Requiring her to report on her awareness of a hidden stimulus will inevitably induce her to look out for such stimuli in all further trials. Pitts et al. (2011) got around this problem by requiring reports not after each single trial, but only after an entire block (see figure 3). During a first phase, they presented a grid of line segments, which could either be randomly configured or produce patterns (frequently, squares, and rarely, diamonds). The grid was surrounded by a ring of red discs, one of which appeared fainter than the others. The subjects’ task was to detect which of the discs was faint on each trial, and this was supposed to distract their attention from the grid shapes that would occur on some of the trials. After the first phase, subjects were first asked whether they saw any patterns, and then they were presented with example patterns (including squares and diamonds) and asked to rate their confidence in having seen each of them. About half of the subjects were oblivious to any patterns presented. The same procedure was then repeated, followed by the same questionnaires, and this time, all subjects reported that they had seen the squares. In a third phase, subjects were asked to ignore the ring of red discs and focus on the grid in order to detect the—rare—diamond patterns, which they easily did. Pitts and colleagues then compared the ERPs of phase-one trials of subjects who were “inattentionally blind” according to the

questionnaires with the ERPs of phase-two trials, thereby obtaining an aware/unaware contrast. They also compared the ERPs of subjects who had been aware of the squares during phase one with the ERPs of those who had not, producing another aware/unaware contrast. Finally, they compared the ERPs for squares versus random patterns within phase three, when the squares were task-relevant (as they had to be discriminated from the target diamonds). They found a negative component that strongly resembled the VAN (except for a somewhat late timing) for aware conditions, irrespective of task-relevance, as well as an SN component that was present only in phase three, thereby confirming the dissociation of SN and VAN reported by Koivisto et al. (2009). The most outstanding finding in the context of the “early” vs. “late” debate, however, was that LP occurred only when the stimulus was both aware and task-relevant (see figure 3). This strongly suggests that the LP component reflects postperceptual processing of task-relevant stimuli rather than conscious awareness as such.

Shafto & Pitts (2015) modified this paradigm by using a grid of line segments that would align to form female faces in some trials (some with missing features), on which they superimposed concentric rings with moving green dots, some of which sometimes brightened up for a short time. The distracter task consisted in detecting when these changes in brightness happened; by superimposing the distracter stimuli on the “hidden” stimuli of interest (i.e., the faces), Shafto & Pitts controlled for spatial attention. The procedure closely followed that of the previous experiment, with phase one and two followed by questionnaires analogous to those of Pitts et al. (2011). In phase three, subjects were required to detect the faces with missing features. ERPs were then compared as in Pitts et al. (2011). Apart from a face-specific N170 component during aware conditions, Shafto & Pitts again found VAN to correlate with awareness irrespective of task-relevance, while SN and P3b appeared only in phase three, when the faces became task-relevant.

Schelonka et al. (2017) carried out a further variation of the paradigm, this time focusing on orthographic and lexical processing. Using a similar grid of line segments and similar distracter

tasks as Shafto & Pitts (2015), the “hidden” stimuli were word forms (either proper words or consonant strings). Once again, they found VAN to correlate with awareness and P3b to correlate with task-relevance.

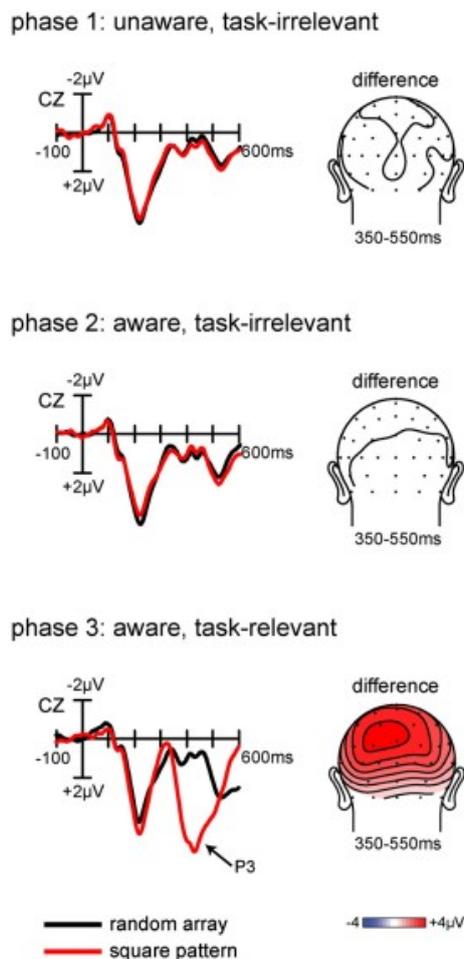


Figure 3: The three phases of the inattentional blindness paradigm by Pitts and colleagues. P3 occurs only in phase 3, when the hidden target is task-relevant. Reprinted from *NeuroImage* 101, Pitts, Padwal, et al., Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness, p. 340, 2010, with permission from Elsevier.

Inattentional blindness is a welcome addition to the established corpus of paradigms in experimental consciousness research. Its decisive advantage over other paradigms is that it permits the researcher to isolate awareness-related ERP activity from all post-perceptual processing, as there is no requirement to memorize, report, or otherwise act on, or immediately “access” aware stimuli during phases one and two. However, there are also some disadvantages, most of which have already been mentioned by Pitts, Padwal, et al. (2014): since awareness-reports can be obtained

only after the presentation of entire blocks in order to maintain the inattentive blindness effect, it is impossible to know on which, and on how many, trials subjects were actually aware of the “hidden” stimuli. Subjects who were aware on many trials will thus be grouped together with subjects who were aware only on few trials, which will affect the ERPs obtained from this group; and it will presumably affect them differently in different experiments, because the “hidden” stimuli can be expected to differ in salience (e.g., rectangle shapes versus human faces), which is likely to have an influence on the number of trials that become aware especially during phase one. Pitts, Padwal, et al. (2014) tried to alleviate such concerns by having subjects estimate their number of aware trials. These subjective estimates correlated with VAN amplitude, and therefore are probably reliable to some degree; but an imperfection remains, and VAN amplitude should be expected to vary considerably between experiments using this paradigm. Another disadvantage is that counterbalancing of the blocks is not possible, since the order of phases is fixed. However, this is not a problem for the within-phase ERP comparisons, which are the most important ones in this paradigm. Finally, it is by definition impossible to manipulate attention and awareness independently in inattentive blindness paradigms, although this drawback is shared by many other paradigms (see Pitts, Padwal, et al., 2014, p. 348).

3.2.3. Task-Relevance: An Important Confound in Studies of Awareness

The results by Pitts and colleagues on P3 and task-relevance are particularly important because the stimuli in the vast majority of previous studies were task-relevant, and consequently, their LP/P3 findings suffer from a potentially devastating confound (this concerns, for example, Del Cul et al., 2007; Lamy et al., 2008; Salti et al., 2012; Sergent et al., 2005; and most of the studies by Koivisto and colleagues). One exception is the study by Koivisto & Revonsuo (2008b), which featured a passive viewing condition without any task. In striking agreement with the results just discussed, they had found that LP was completely eliminated in this condition.

Pitts and colleagues reproduced and extended their findings in further studies. Pitts, Padwal, et al. (2014) abandoned the ring of discs and presented patches of colored line segments in the same grid region in which they presented horizontal and vertical rectangle shapes, thereby controlling for spatial attention and competition between objects. This study focused on task-relevance in aware trials, so there was no need to induce inattention blindness. Instead, subjects were shown both types of stimuli in advance, and then simply required to respond to either the color patches or the rectangle shapes. ERPs were then compared for rectangle shapes vs. random stimuli, both when they were relevant and when they were irrelevant. As they had predicted, Pitts et al. found LP for the shapes only when they were task-relevant. In another experiment reported in the same study, they manipulated the *degree* of task-relevance, operationalized as the degree of similarity between the target and other stimuli. They found that LP amplitude correlated positively with the degree of task-relevance, and that LP was again only present for stimuli that had at least some degree of task-relevance.

A drawback of this study and the inattention blindness studies is that they did not feature a task-relevant but unaware condition, which would be required for a complete double dissociation of awareness and task-relevance. Pitts, Metzler, et al. (2014) realized a 2×2 design in a backward masking paradigm with two SOAs, one of which was sufficiently short (16 ms) to ensure that the target became unaware, and the other one of which (300 ms) guaranteed that the mask did not interfere with stimulus visibility.⁸ They used variations of stimuli from their earlier experiments: a grid of line segments, which aligned to form square or diamond shapes on some trials, and either

8 The same method of manipulating stimulus visibility (awareness) by using different SOAs has been used in several earlier backward masking studies, including (Koivisto, Kainulainen, & Revonsuo, 2009; Koivisto, Revonsuo, & Salminen, 2005; Koivisto & Revonsuo, 2007, 2008b). It has been criticized (albeit in the context of metacontrast masking) by Talis Bachmann (2009) for introducing possible interaction confounds between stimuli and masks at the different SOAs (but see Railo & Koivisto, 2009a, for a reply). Whereas these earlier studies had employed control experiments with constant SOAs to rule out ERP differences caused by the different timing of the masks, Pitts, Metzler, et al. (2014) presented masked control stimuli at both SOAs and subtracted these control ERPs from the ERPs of interest before contrasting aware with unaware conditions. This was supposed to remove all mask-elicited activity. The fact that Pitts, Metzler, et al. (2014) arrived at similar results as Koivisto and colleagues provides convergent evidence for the validity of the masking procedure with different SOAs.

three or four patches of colored line segments on others. The task on each block was to detect either the patches with three lines colored or the diamond stimuli, and report them via button pressure. The twist here was that all target stimuli were discarded for the analysis (thereby removing any motor preparation and execution confounds), and only the ERPs of the non-target stimuli used for comparisons. In each task, the non-target of the same category (in the diamond-detection task: the square) was deemed task-relevant, and the stimuli of the other category task-irrelevant. In combination with masking, this yielded four conditions, among them the sought-after “unaware/task-relevant” condition. The ERP analysis revealed VAN for both aware conditions (and for neither of the two unaware conditions). In contrast, “the P3b was present in some of the unaware, task-relevant conditions, and absent in some of the aware, task-irrelevant conditions” (Pitts, Metzler, et al., 2014, p. 11). This is particularly interesting, because it goes beyond all the results reviewed here so far. The previous ERP studies by Pitts et al. had questioned the *necessity* of the P3b component for awareness by showing that awareness could occur without it (namely, in task-irrelevant conditions). The fact that the P3b can also occur in *unaware* conditions shows that it is not *sufficient* for awareness, either.

Silverstein et al. (2015) provided independent evidence for this insufficiency by employing a subliminal oddball paradigm. They presented rare and frequent words for extremely brief periods, followed by masks, to achieve complete stimulus invisibility, and asked subjects to closely attend to the subliminal stimuli (even though they protested that they could not see anything to focus attention on). Silverstein and colleagues employed a further condition including a detection task, and applied signal detection theory to ensure that subjects were unable to detect the subliminal stimuli above chance. What they found was that, just like in supraliminal oddball experiments, the rare stimuli elicited larger P3, followed by a sustained late slow wave, than the frequent stimuli—although all stimuli were unconscious and undetectable. A principal component analysis (PCA)

identified separate P3a and P3b components, and thus verified that they had indeed found unconscious P3b.⁹

3.2.4. No-Report Paradigms: Wining the NCC from Its Prerequisites and Consequences

These results underscore how vitally important it is to distinguish the NCC from their prerequisites and consequences. Many of the results reviewed here depict the P3b as a prime example of an ERP component reflecting a post-perceptual process (or a number thereof) that has been systematically conflated with the proper NCC because of insufficient controls for task-related activity. A recent study by Koivisto, Ruohola, Vahtera, Lehmusvuo, & Intaite (2018) using low-contrast stimuli and a dual-task paradigm may serve as yet another example. The authors have dissociated the VAN from the effects of working memory—another important potential confound stemming from an arguably post-perceptual process—under three different conditions; even the LP was modulated by working memory only under heavy executive load.

As a way of dealing specifically with report-related confounds (which may include all kinds of post-perceptual processes, such as working memory, response selection, motor preparations and executions, etc.), no-report paradigms have recently been discussed extensively (Koch et al., 2016a; Overgaard & Fazekas, 2016; Tsuchiya, Frässle, Wilke, & Lamme, 2016; Tsuchiya et al., 2015). Such paradigms may replace requirements of subjective awareness reports with objective measurements, e.g., eye movement or pupil dilation in response to perceptual shifts in binocular rivalry (Tsuchiya et al., 2015). Another way is the “inclusion of no-report control conditions” into classical report-based paradigms (Tsuchiya et al., 2016, p. 242). Just such a combination of report and no-report conditions has been realized in a recent ERP experiment (Koivisto, Salminen-Vaparanta, Grassini, & Revonsuo, 2016). The authors of this study presented low-contrast Gabor patches in a “go”/“no go” task and counterbalanced the response requirements: in one condition,

⁹ Naccache et al. (2016) raised some methodological concerns about this study, but Silverstein et al. (2016) convincingly rejected most of them as either irrelevant or unwarranted. For example, Naccache et al. (2016) had completely ignored the results of the PCA in their critique.

subjects had to respond when they saw the stimulus and not respond when they saw nothing, whereas in the second condition, they had to respond when they did *not* see the stimulus (and *not* respond when they did see it). This enabled the Koivisto and colleagues to compare the ERPs of four conditions: “aware-GO, aware-NOGO, unaware-GO, and unaware-NOGO” (Koivisto et al., 2016, p. 3). They found VAN to be present in all aware (and absent in all unaware) conditions, and no difference in VAN between the aware-GO and aware-NOGO conditions, suggesting that VAN was not changed by response requirements. In contrast, while LP was likewise present only in aware conditions, its amplitude was larger in the aware-GO than in the aware-NOGO condition. While it remains open why such a difference was not observable between the unaware GO and NOGO conditions, this result clearly suggests that P3 is at least in part caused by response-related processes occurring only after the stimulus has already become conscious.

The ERPs obtained in this experiment may still include some activity that is indirectly related to reports; for example, in principle even the VAN found here might still include activity related to perceptual decisions, or, during NOGO conditions, to response inhibition, that remained undetected, perhaps for a lack of sensitivity. Other paradigms may be even more powerful in isolating the NCC from its consequences. In fact, the inattention blindness paradigm discussed in a previous section is a particularly powerful instance of a no-report paradigm, that can arguably avoid “any and all post-NCC confounds”, so that “a potential NCC can be fully dissociated from cognitive access and attention” (Tsuchiya et al., 2016, p. 243). It should also be noted that no-report paradigms have occasionally been employed already *avant la lettre*; an example is the passive viewing task employed by Koivisto & Revonsuo in an already mentioned experiment (2008b).¹⁰

No-report paradigms are not without their own problems and limitations. Notably, while they can exclude certain confounds, and thereby help to narrow down the space of possible NCC, they may still overestimate the NCC “by including certain pre-NCC activity” (Overgaard &

10 One problem with passive viewing conditions, however, is that they might induce mind-wandering and suffer from a “lack of control over closely related processes such as attention and working memory” (Pitts, Lutsyshyna, & Hillyard, 2018, p. 8), which arguably renders them less than ideal for NCC research.

Fazekas, 2016, p. 241). The most reasonable way to proceed is surely to combine no-report paradigms with conventional ones, to include no-report conditions in the latter, and to employ both kinds of paradigms across the whole range of available techniques for manipulating awareness and other parameters of interest (Tsuchiya et al., 2016).

3.2.5. The Relationship Between Consciousness and Attention

Another major potential confound in NCC research is attention, and the relationship between awareness and attention is a fundamental issue in the science of consciousness (Aru & Bachmann, 2013; Cohen, Cavanagh, Chun, & Nakayama, 2012b, 2012a; Koch & Tsuchiya, 2007, 2012; Lamme, 2003, 2004; Noah & Mangun, in press; Tsuchiya, Block, & Koch, 2012; Tsuchiya & Koch, 2016; Van Boxtel & Tsuchiya, 2015; Van Boxtel, Tsuchiya, & Koch, 2010). As noted in section 2.2.1., according to GNWT, attention is a necessary prerequisite for conscious access, and if this is true, it obviously affects the search for the NCC. In particular, any claim about a putative NCC must be substantiated by showing that it is not, in fact, merely a correlate of a necessary prerequisite of consciousness, namely, attention.

The issue is complicated further by the fact that attention is no more unitary a concept than consciousness. A basic distinction separates exogenous, bottom-up attention from endogenous, top-down attention. The former is automatically attracted by salient (e.g., by particularly strong, and often, by evolutionarily relevant) stimuli, whereas the latter can be employed voluntarily, according to the task at hand. Both forms of attention can be feature- or object-based; they can be employed to global or to local features of an object or a scene; and also to particular spatial or temporal regions (see Eysenck & Keane, 2015, Chapter 5, for an overview). These varieties of attention constitute subsystems that can operate quite independently of each other (Cohen et al., 2012b; Tsuchiya & Koch, 2016), although it is not yet clear how (dis-)integrated their mechanisms are at the neural level. Since the mid-2000s, a number of researchers have begun to state and defend a double dissociation between top-down attention and visual awareness (Koch & Tsuchiya, 2007, 2012;

Tsuchiya et al., 2012; Tsuchiya & Koch, 2016; Van Boxtel et al., 2010), while others insist that the dissociation runs only one way, and top-down attention remains necessary for awareness (Cohen et al., 2012b, 2012a; Noah & Mangun, in press).

In the field of ERP research, the most pressing question is whether the ERP correlates of the various forms of attention can be separated from those of awareness. As noted in section 3.1., a number of successful dissociations of the VAN from several attention-related components, such as the SN, had already been achieved prior to 2010 by Koivisto and colleagues (reviewed in Koivisto et al., 2009). Since then, further ERP studies have been carried out.

While most work on the relationship of attention and awareness focuses on endogenous, top-down attention, Chica et al. (2010) investigated how exogenous, bottom-up attention interacts with visual awareness. This is usually considered difficult, because manipulating the salience of a stimulus almost inevitably implies manipulating its physical properties, thereby introducing an unwanted confound (Van Boxtel et al., 2010). Chica and colleagues used low-contrast gratings with slightly different orientations as stimuli, while they presented peripheral salient, but uninformative cues in order to manipulate exogenous attention. They reasoned that “[i]f exogenous attention improved conscious perception, more targets should be consciously reported at the exogenously attended location” (Chica et al., 2010, p. 1206), which is what they found. In an analysis of ERPs locked to the appearance of the cue, they found a strong P100 component both for seen targets when the cue was valid, and for unseen targets when the cue was invalid. Chica et al. interpreted this component to reflect the capture of exogenous attention by the cue, and concluded that both the successful attraction of exogenous attention to the relevant location by a valid cue and the failure of a misleading cue to attract exogenous attention to the wrong location facilitate conscious perception. Comparing the target-locked ERPs of the aware and unaware conditions, they found an LP that had a slightly earlier onset latency for seen targets at invalid locations, which they interpreted as a sign of a re-orientation of exogenous attention from the invalid cue to the target.

This study reported no VAN, and apparently the N2 time window has not been scrutinized by the authors. Accordingly, nothing can be learned from this study about possible effects of exogenous attention on the VAN. Studies explicitly inquiring about this relationship remain to be carried out.

Del Zotto & Pegna (2015) investigated “the interactions between non-spatial selective attention, awareness and emotion processing” (Del Zotto & Pegna, 2015, p. 1) in a backward masking paradigm, using faces with different emotional expressions as stimuli, and a categorization task supposed to engage selective attention. Their results reproduced the dissociation of VAN and SN found in earlier studies (Koivisto et al., 2009, 2005; Koivisto & Revonsuo, 2007, 2008b); moreover, P3 depended on selective attention in that it was larger for targets than for non-targets, irrespective of awareness, suggesting to the authors “that the P300 should rather be seen as a consequence of consciousness, related for example to post-perceptual processing, rather than to awareness *per se*” (Del Zotto & Pegna, 2015, p. 10).

Davoodi et al. (2015) presented “masked” (short SOA) and “unmasked” (long SOA) pictures of human faces and fruits, and the participants were required to attend to the pictures on some blocks, and to count the number of targets across trials on other blocks. Their results are somewhat anomalous, in that the only ERP components that correlated with awareness in both the attentive and inattentive conditions were an increased early P100 over occipital and parietal electrode sites (which they baptized “early consciousness positivity”, ECP), and an increased negativity in the N150 time range over frontal sites F3 and F4. While the time range and polarity of this component resemble the VAN, its highly untypical scalp distribution makes it clear that it must be some other component. The P300 correlated with attention irrespective of awareness in occipital, occipito-parietal, and parietal electrodes. The authors provide no information about the measures of awareness used in their study, except that the “participants were asked to be attentive and to respond to the specified targets” (Davoodi et al., 2015, p. 146). It is not clear whether and how they verified that subjects were indeed aware of the targets in the inattentive condition (for example, whether

they performed above chance), or that the different mask SOAs successfully enabled or shut off awareness. In the absence of such information, it is difficult to critically assess their results, which should therefore be treated with caution. Moreover, Davoodi and colleagues' experimental control over attention seems rather casual. While their conclusion that "[w]hile bottom-up attention seems to be strictly connected to awareness, top-down attention is not necessary for consciousness" (Davoodi et al., 2015, p. 153) is in line with the results of Chica et al. (2010), it is not clear how exactly their procedure and results are supposed to justify it.

Chen et al. (2017) investigated the relation of exogenous feature-based attention and awareness. Since their ERPs were cue- rather than target-locked, resulting in different baselines and component latencies, their results are not directly comparable to the other studies reviewed here, and therefore do not appear in table 1. They are interesting nonetheless, as they suggest that exogenous attention can be dissociated from awareness just as well as endogenous attention. More studies are of course needed to corroborate this isolated finding.

Overall, the relationship between attention and consciousness is still perplexing, and the effects of different forms of attention on putative ERP correlates of visual awareness are, in large part, far from established. The results reviewed by Koivisto et al. (2009) are an important first foray into this territory, but they need to be replicated under different conditions, taking into account the possible confounds that have been discussed in the past ten years, and under review here: objective task-performance, task-relevance, expectations, etc. The exciting new results on exogenous attention and awareness should explicitly be investigated in their relationship to the VAN in future studies in order to corroborate or undermine, as the case may be, its status as a proper ERP correlate of visual awareness.

3.2.6. The Role of Expectations and Prior Beliefs

The recent decade of cognitive and computational neuroscience has seen an ever-increasing interest in the role of expectations in perception, and perception has frequently been treated as,

fundamentally, a problem of (Bayesian) inference (Kersten & Yuille, 2014; Olshausen, 2014). In a—somewhat surprisingly, given the importance of the topic—quite singular ERP study, Melloni, Schwiedrzik, Müller, Rodriguez, & Singer (2011) have investigated the effect of expectations on visual awareness. To manipulate expectations, they presented stimuli (letters, numbers, or symbols) in a “ladder” procedure, repeating the same stimulus eleven times, with contrast (visibility) increasing step-wise over the first six trials (the “ascending” part of the sequence), and decreasing contrast over the remaining five (the “descending” part), as visualized in figure 4. Participants were to report their level of subjective visibility on a graded perceptual awareness scale. They became aware of the stimuli more often during the descending part of the sequence, showing that conscious perception changed non-linearly with degradation level, an instance of an effect called “hysteresis” (Kleinschmidt, Büchel, Hutton, Friston, & Frackowiak, 2002). Plausibly, this effect was due to the presence of expectations in the subjects in the descending part, built up during the ascending part.

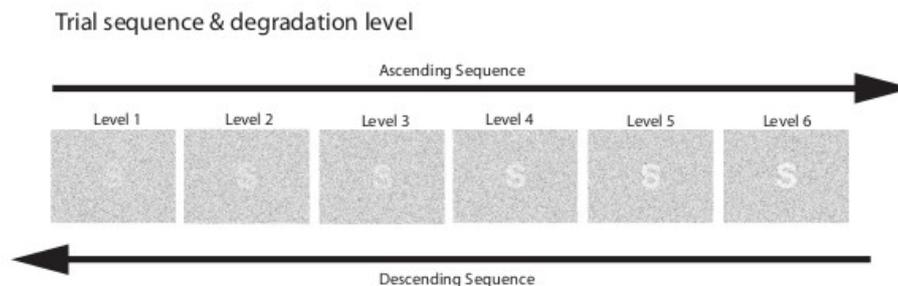


Figure 4: Visualization of the stimulus sequence used by Melloni and colleagues. Stimulus contrast increased over six degradation levels, then decreased over the same six levels. Reprinted from *The Journal of Neuroscience* 31(4), Melloni et al., *Expectations Change the Signature and Timing of Electrophysiological Correlates of Perceptual Awareness*, p. 1387, 2011.

Melloni and colleagues found that the only ERP component that differentiated between seen and unseen trials in both the ascending and descending parts of the sequence was the P2, which was significantly reduced in amplitude in the seen trials (i.e., VAN). The P3 component was higher in amplitude for seen trials, but only in the ascending part of the sequence, suggesting that it is an ERP correlate of visual awareness “only in the absence of expectations” (Melloni et al., 2011, p. 1395). However, in a control experiment, in which different stimuli were used on each trial during a

sequence in order to control for possible sequence effects, and in which, consequently, no expectations could be built up by the brain, the P3 reflected “visibility more closely than P2” (Melloni et al., 2011, p. 1395). The authors interpreted this to show that the onset of awareness is neither “early” nor “late” *per se*, but that the latency of the electrophysiological NCC depends on the presence or absence of expectations in the perceiving subject. This interpretation is most interesting in the present context, but it runs counter to the robust VAN findings in most of the literature. It is conceivable that some form of expectations was always present in these earlier studies, but it is impossible to know, since this factor has never been systematically controlled for. Clearly, more studies addressing expectancy, and its interaction with other factors such as task-relevance and attention, are urgently called for.

Melloni et al. (2011) explicitly relate their study to a particularly important inferential account of perception, namely “predictive coding”, also known as “prediction error minimization” (Hohwy, 2013, 2015) or “predictive processing” (Clark, 2013, 2016), and intimately related to Karl Friston’s “free-energy principle” (Friston, 2005, 2009). According to such accounts, the brain is fundamentally engaged in generating predictions about the expected sensory input, and optimizing its prior beliefs, which are implemented in a generative model of the external world. Melloni et al. (2011, p. 1394) note the seemingly high compatibility of predictive coding and RPT: while the latter associates recurrent feedback activity with conscious awareness, the former identifies it with the predictions generated by the brain. If one assumes, as is done in leading formulations of predictive coding, that “perceptual [read: phenomenal, J.F.] content is the predictions of the currently best hypothesis about the world” (Hohwy, 2013, p. 48, italics in original), then one might even argue that predictive coding, by assigning a specific functional role to recurrent processing, goes some way toward explaining *why* recurrent processing should be accompanied by phenomenal consciousness (although, of course, it remains equally mysterious why phenomenal contents should be identical to neural predictions). However, predictive coding is probably also compatible with GNWT; at any

rate, Stanislas Dehaene tries to conciliate the two in recent publications (e.g., Dehaene, 2014, pp. 93, 96, 257). A thorough investigation of the exact relations between the various empirical theories of consciousness and the predictive coding scheme remains to be carried out, and is obviously beyond the scope of a single review. One of the topics that may look quite differently from a predictive coding perspective is the relationship between consciousness and attention.

ERP research need not await these theoretical developments, however, as the study by Melloni et al. (2011) demonstrates. At some distance from research specifically related to conscious awareness, an electrophysiological literature looking for predictions in the brain is already developing (e.g., Bendixen, SanMiguel, & Schröger, 2012; Winkler & Czigler, 2012), and it will be most fruitful to make the connection to consciousness research and the literature reviewed here.

3.2.7. Awareness: Graded, Dichotomous, or Both?

Yet another intensely debated topic in the science of consciousness is the nature of awareness as either graded or dichotomous. In line with GNWT's hypothesis that conscious access to a stimulus requires the neural activity induced by it to cross a threshold of "global ignition", the studies carried out by this camp generally found a non-linear increase in neural activation for stimuli that became conscious (Del Cul et al., 2007; Sergent et al., 2005; Sergent & Dehaene, 2004). These studies used a quasi-continuous scale to capture subjects' conscious reports, where visibility ratings fell on one of 21 contiguous positions, with only the extreme ends of the scale labeled as "nothing seen", and "maximal visibility", respectively. The distribution of the visibility scores turned out to be bimodal in these experiments, with almost all of the ratings falling on the extreme ends. From this result, the authors concluded that visual consciousness is an all-or-none phenomenon, just like GNWT predicts.

However, another group of researchers produced diverging evidence. Exploring different scales of visibility ratings, Ramsøy & Overgaard (2004) found that their subjects, when encouraged to rate visibility in whatever way seemed most convenient to them, all converged on a four-point

scale with the categories “no experience”, “brief glimpse”, “almost clear image”, and “absolutely clear image”. The score on this “perceptual awareness scale” (PAS) co-varied linearly with both stimulus duration and probability of a correct answer in a variety of tasks, including shape and color discrimination (Ramsøy & Overgaard, 2004) and stimulus localization (Overgaard, Rote, Mouridsen, & Ramsøy, 2006; Ramsøy & Overgaard, 2004), across different report modalities. To the authors, this suggested that visual consciousness is a phenomenon that admits of different degrees, rather than being simply “present” or “absent”.

The PAS, as well as modified versions of it, have been used in a number of ERP studies of visual consciousness, including some of the already mentioned ones (Eklund & Wiens, 2018; Koivisto & Grassini, 2016; Melloni et al., 2011). In these studies, the PAS has usually been artificially dichotomized (i.e., the lowest and the highest two response categories have been collapsed) for the purposes of analysis, mainly because of difficulties in obtaining enough responses in each category. For example, Koivisto, Kastrati, & Revonsuo (2013) recorded ERPs during a scene categorization task using object substitution masking (see footnote 4 of this thesis) and found that masking influenced the visibility ratings on the PAS scale, but did not use their PAS data for ERP analysis, comparing instead only “masked” and “unmasked” trials.

One study explicitly set out to compare ERPs for all four PAS categories, using low-contrast stimuli at detection threshold (Tagliabue et al., 2016). Holding the ERPs for each category against the ERP of the lowest possible rating (“no experience”), they found that both the VAN and LP amplitudes correlated linearly and positively with PAS ratings (see figure 5). Given that the P3, as evidenced by many of the studies reviewed here, is likely to reflect post-perceptual processes rather than consciousness as such, the authors suggest that the increase in LP amplitude might index “different levels of accumulation of sensory evidence”, in line with Melloni et al. (2011), or “working memory update”, in line with Polich (2007) (Tagliabue et al., 2016, p. 8).

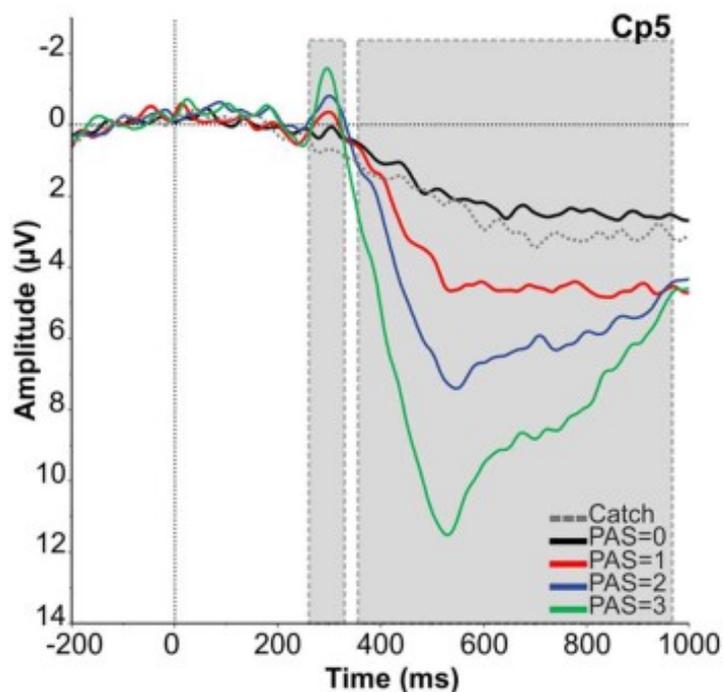


Figure 5: Both VAN and LP amplitudes co-vary linearly with PAS score. Reprinted from *Frontiers in Psychology* 7, Tagliabue et al., *Early Local Activity in Temporal Areas Reflects Graded Content of Visual Perception*, p. 4, 2016.

Bert Windey and colleagues observed that studies producing results that favor a graded view of consciousness have usually used simple stimuli (such as simple shapes and colors) and tasks (e.g., simple detection), while studies supporting a dichotomous notion of consciousness typically employed more complex stimuli (e.g., letters and words) and tasks (e.g., word identification and object categorization tasks). They suspected that what they call the “level of processing” affects the distribution of visibility ratings on PAS and similar scales. In a masking study, they gave their subjects both a low-level and a high-level task (color identification and numerical judgments, respectively), but used identical stimuli (colored numbers) in both tasks (Windey, Gevers, & Cleeremans, 2013). Visibility ratings were obtained using the PAS. The results confirmed their suspicion and led them to put forward a “levels of processing hypothesis”: low-level stimuli and tasks exhibit “a graded relationship between stimulus duration and visibility” and yield a linear psychophysical curve when plotting these two parameters against each other, while high-level stimuli and tasks lead to non-linear curves and dichotomous visibility ratings (Windey et

al., 2013, p. 404). The “levels of processing hypothesis” integrates the conflicting views in that, according to it, consciousness is neither exclusively graded nor exclusively dichotomous, but both, depending on the level of processing that a stimulus or task demands.

Building upon this insight, a number of researchers have realized that stimulus and task complexity have not been taken into account to a sufficient degree in ERP research on visual awareness, and attempted to remedy this situation. Koivisto, Grassini, Salminen-Vaparanta, & Revonsuo (2017) presented low-contrast stimuli, and used a modified PAS to define two distinct thresholds for detection and identification, reasoning that the latter is a more complex task requiring a higher level of processing. The difference between the lowest two PAS ratings (“I did not see any stimulus” vs. “I saw something (but could not identify the stimulus)”) defined the detection threshold, while the difference between the latter and the next-higher rating (“I saw the stimulus almost clearly (and could identify it)”) defined the identification threshold. Expecting to find both VAN and LP in all aware trials (regardless in which of these two ways “awareness” was operationalized) when comparing ERPs according to these PAS contrasts, to their surprise, the authors found VAN to correlate only with detection, but not with identification, whereas LP correlated with both. Although this result does not directly speak to the “levels of processing hypothesis” in that it does not compare PAS distributions at different levels of processing (but instead uses different PAS contrasts to define such levels), it is clearly compatible with it, and pushing in the same direction in that it identifies different ERP correlates for awareness at different levels of processing.

A recent study explored the “levels of processing hypothesis” more directly by investigating whether and how the amplitudes of VAN and LP were affected by different levels of processing (Jimenez, Grassini, Montoro, Luna, & Koivisto, 2018). Using backward masking and a low-level target detection task, as well as a high-level number/letter identification task, they found that VAN amplitude varied with PAS ratings in a graded manner only in the low-level task, whereas LP

amplitude varied gradually with PAS ratings in both the low- and the high-level tasks. The fact that only VAN was affected by the level of processing suggested to the authors that the level of processing manipulation effects recurrent processing in the visual cortex, maybe modulated by attention. The result that LP varied gradually and never dichotomous is in conflict with the “levels of processing hypothesis”; Jimenez and colleagues propose the interpretation that their “high-level task” was in fact not high enough to produce dichotomous LP amplitude variation, as it involved stimulus identification, but, unlike, e.g., the study by Del Cul et al. (2007), no “higher cognitive and decision-making processes (such as comparison of meanings)” (Jimenez et al., 2018).

Fu, Liu, Dienes, Wu, Chen, & Fu (2017) used complex natural scene stimuli in both color photograph and line-drawing versions in a scene categorization task and backward masking at different SOAs, and found that while the VAN amplitude increased variably with PAS ratings, while “LP (or P3) amplitudes seemed to change non-linearly with subjective ratings” (Fu et al., 2017, p. 10). They interpret this along the lines of Koivisto and colleagues (see section 2.3.2.), and suggest that VAN reflects phenomenal consciousness, while LP reflects reflective/access consciousness. When taken together with the study by Tagliabue et al. (2016), the results directly contradict those of the study by Jimenez et al. (2018). In the latter, only the VAN was affected by the level of processing, whereas both Tagliabue et al.’s study (which used simple Gabor stimuli) and Fu et al.’s study (which used complex natural scene stimuli) found graded VAN, while the LP was apparently affected by the level of processing (graded for simple stimuli in Tagliabue et al.; dichotomous for complex stimuli in Fu et al.).

It has rightly been noted that one has to be careful in specifying precisely which aspect or aspects of consciousness one takes to be graded or dichotomous (Windey & Cleeremans, 2015; Windey, Vermeiren, Atas, & Cleeremans, 2014). A graded notion of consciousness is often associated with RPT (e.g., Windey et al., 2013, p. 405), while GNWT is taken to be committed to a dichotomous concept of awareness. However, such a view is actually too simplistic. The proponents

of GNWT have refined their notion of access and put forward the “partial awareness hypothesis” (Kouider, de Gardelle, Sackur, & Dupoux, 2010). According to this hypothesis, the *mechanisms* of consciousness operate in an all-or-none fashion, while what they provide access to is a hierarchy of representations of a given stimulus on different levels of coarser or finer detail. These levels can be accessed independently of each other (depending on signal strength and other factors, such as expectations), giving rise to a *phenomenology* of conscious access that is gradual rather than dichotomous. As for RPT, contrary to what, e.g., Timmermans, Windey, & Cleeremans (2010) read into Lamme’s account, there does not seem to be an intrinsic connection between RPT and graded awareness; in fact, Lamme can be read as advocating something very close to the “partial awareness hypothesis”.¹¹ Overall, it seems that more work needs to be done before the issue of graded vs. dichotomous awareness can be settled—both empirically, as the ERP results reviewed here do not speak unequivocally for either view (nor is it clear just how they should be interpreted with regard to the “levels of processing hypothesis”), and theoretically, as neither the exact relationships between the “levels of processing” and the “partial awareness” hypotheses, nor the relationships of both RPT and GNWT to either of them have been explicated in sufficient detail. The conflicting results of two MEG studies (Andersen et al., 2016; Sekar, Findley, Poeppel, & Llinás, 2013), to be reviewed in the next section, add to this ambiguity.

In sum, new developments in research on ERP correlates of visual consciousness abound and raise more questions than they answer. However, among the questions that they *do* answer is the one about the earliest and most reliable correlate: the VAN, or visual awareness negativity, is the most promising candidate for a genuine NCC, while the occurrence of the later LP is far more contingent on reflective, post-perceptual processes.

11 “The theory predicts that as soon as you have modules representing different aspects of a scene engage in recurrent interactions, you will have a core of conscious sensation linked to whatever is represented by these modules. Adding more modules—of a lower or higher level—makes the conscious sensation of a richer content, but not more or less conscious” (Lamme, 2010b, p. 235). If the outputs of these “modules” are understood as representations at different levels of detail, as described by Kouider et al. (2010), then the only obvious difference to the latter’s “partial awareness hypothesis” lies in the mechanisms taken to be responsible for conscious awareness—local recurrent processing vs. distributed global workspace activation.

3.3. A Review of Recent MEG Studies on the Correlates of Visual Awareness

This section reviews the five MEG studies that have been carried out on the correlates of visual consciousness since 2010. The inclusion criteria are the same as above for the ERP studies. As can be ascertained from table 2 below, the results are even more univocal than the ERP results: all five studies found VAN in all aware conditions, and none of them found LP in all aware conditions.¹² This further corroborates the main result of this thesis. The individual studies are briefly discussed below.

Table 2

Results of the Review of MEG Studies for VAN (N2 range) and LP (P3 range)

Manipulation	Study	Enhanced early negativity (N2 range)	Enhanced late positivity (P3 range)
Contrast	Liu et al. (2012)	Yes	No
Masking	Andersen et al. (2016)	Yes/Yes/Yes	No/No/No
Threshold duration	Sekar et al. (2013)	Yes/Yes	Yes/Attenuated
Bistable perception	Sandberg et al. (2013)	Yes	Attenuated
	Sandberg et al. (2014)	Yes	Attenuated

Note. For each study, I report which of the two relevant MEG correlates (the equivalents of the ERP components VAN and LP in terms of both temporal occurrence and scalp distribution) have been found. In the case of two or more entries, the studies have employed more than one experimental condition.

The already mentioned study by Liu et al. (2012) used low-contrast colored gratings in two different orientations and an objective (two-alternative forced choice about stimulus orientation) as well as a subjective task (deciding whether a grating was present). They found VAN and localized it in the lateral-occipital complex (see also section 2.3.1.). Additionally, they found gamma-band oscillations associated with aware trials; as Pitts, Padwal, et al. (2014) have pointed out, however, these probably reflect post-perceptual processes rather than visual awareness. No frontal or parietal activity whatsoever was found to be related to awareness.

Sandberg and colleagues carried out two studies on bistable perception, the only ones in this review. Sandberg, Bahrami, Kanai, Barnes, Overgaard, & Rees (2013) studied face perception

¹² Note that “VAN”, “LP”, and “P3” are primarily the names of ERP components, and that their event-related field (ERF) equivalents are usually given different names. In this review, I sometimes use “VAN” and “LP” to refer to awareness-related amplitude differences of both ERP and ERF components.

during intermittent binocular rivalry (using a face on one eye, and a grating on the other).¹³ In the classical univariate analysis of event-related fields (ERFs), they found the M170 (the magnetoencephalographic equivalent of the face-specific N170 ERP component) and the P2m (the equivalent of the P3 component in ERP) to correlate with conscious face perception. In an additional multivariate prediction analysis, they applied a machine learning algorithm to detect the time ranges in which the evoked responses were most predictive of conscious perception (in other words, they “decoded” conscious perception from brain activity). This analysis confirmed that activity in the N2 and early P3 time range best predicted conscious face perception, while activity after 300 ms post-stimulus was not predictive at all. Interestingly, these predictions were possible even across subjects, i.e., the time course of brain activity of one subject could successfully predict the onset of conscious face perception in other subjects. Source analysis revealed that fronto-parietal activation was far worse at predicting consciousness than posterior activation, contrary to the findings of Del Cul et al. (2007) and Sergent et al. (2005). The authors note the possibility that this was because these earlier studies compared awareness of a stimulus with no awareness at all, whereas binocular rivalry studies, by their nature, compare awareness of a stimulus with awareness of another stimulus. In a follow-up study, Sandberg, Barnes, Bahrami, Kanai, Overgaard, & Rees (2014) used the same stimuli to investigate the MEG correlates of conscious face perception, perceptual reversals, and stabilization. Regarding conscious perception, they reproduced the results of their previous study (M170 and P2m).¹⁴

The remaining two studies to be reviewed here addressed the graded/dichotomous controversy. Sekar et al. (2013) presented numbers (0 to 9) at threshold contrast and obtained visibility ratings using a four-point scale similar to the PAS, but inquiring the subjects’ perceptual

13 Usually, in binocular rivalry subjects are continually presented with a different stimulus to each eye. In intermittent binocular rivalry, however, the presentation is interrupted by periods of blank displays of arbitrary duration, up to several seconds. This leads to a drastic reduction of the rate at which the subject’s percept alternates, an effect called “stabilization” (Leopold, Wilke, Maier, & Logothetis, 2002).

14 In the context of bistable perception, it should be noted for the sake of completeness that there is a study which has successfully dissociated the VAN from the reversal negativity (RN) component (Intaitė, Koivisto, Rukšėnas, & Revonsuo, 2010).

certainty rather than directly the quality of their percepts. While they acknowledged the potentially important distinction between detection and identification, it made no difference in their ERF analysis: they found both a mid-latency and a late-latency component (i.e., VAN and LP) for aware trials, regardless of the degree of awareness according to the certainty rating. Although the mid-latency component increased gradually with certainty rating, this increase remained insignificant in pairwise comparisons, which the authors interpreted as evidence that consciousness is all-or-none rather than gradual. However, they note the possibility that their study was not sensitive enough to detect more graded nuances, the existence of which is, after all, suggested by the very fact that subjects used the full range of options when rating visibility. The late-latency component had a similar peak amplitude for lowest (“didn’t see”) and the next-higher (“couldn’t identify”) visibility rating condition and was thus less reliable as an NCC than the mid-latency component, in which this contrast was significant.

Andersen et al. (2016) used an identification task (diamond vs. square shapes) in combination with the PAS in a backward masking paradigm and applied multivariate pattern analysis to identify the time range and cortical location in which activity was most predictive of graded visual awareness. They compared all three possible conditions resulting from contrasts of contiguous PAS categories, i.e., “no experience”/“weak experience”, “weak/almost clear experience”, and “almost clear/clear experience”. The results were overwhelmingly clear: occipital sources were the only ones to contain information capable of decoding all three PAS contrasts above chance, and “only the VAN time range was associated with a significant increase in classification accuracy” (Andersen et al., 2016, p. 2682). Frontal activity was utterly powerless to predict any but one of the PAS contrasts, as was any activity in the P3 time range.

One noteworthy development visible in recent MEG (and, sometimes, ERP) research is the predictive approach, i.e., the use of multivariate pattern analysis with the help of machine learning algorithms. One strength of this approach is its ability to detect distributed patterns of information

that would be difficult to detect by conventional means of univariate analysis of summary statistics such as mean peak amplitudes or latencies (e.g., Sandberg et al., 2013, p. 977). In the study of Andersen et al. (2016), univariate analysis did not lead to the results revealed by the predictive approach. Another possible advantage follows from the reasoning that the activity most predictive of conscious awareness is also the activity most likely to reflect consciousness *per se*, as opposed to its prerequisites or consequences (Andersen et al., 2016, p. 2679). However, this reasoning is, of course, not valid for processes that precede (or follow) consciousness *necessarily*, or at least with great regularity.

Overall, the MEG studies reviewed here confirm the view that VAN is the earliest and most reliable electro- and magnetophysiological correlate of visual consciousness. It seems, then, that the “early” vs. “late” debate is settled—at least as far as this question is concerned. The next section will evaluate the importance of this finding in the broader context of consciousness research.

4. Conclusion

The main empirical result of this review seems clear: no ERP/MEG correlate is more closely related to visual awareness than the VAN. The occurrence and modulation of the LP is far more contingent on a variety of factors, such as task-relevance, expectations, and reports, making it more likely to reflect post-perceptual processes. Furthermore, the P3b can occur during unconscious processing, undermining the claims by proponents of GNWT that it is sufficient to demonstrate the presence of consciousness.

The question of what this means for the broader theoretical context of consciousness research is, however, more complicated than it may seem at first glance. It is clear that the empirical evidence clashes with some of the most central predictions of GNWT: the P3b is no reliable marker of consciousness at all; frontal, allegedly workspace-associated activation does not seem necessary at all for a stimulus to become conscious; there is evidence that consciousness might be gradual

rather than an all-or-none phenomenon; and the demonstrably most reliable ERP correlate of consciousness occurs much earlier than GNWT hypothesizes.

This picture appears even more substantiated if one looks beyond the “early” vs. “late” debate. As mentioned in the introduction, and as suggested by the central tenets of the two major conflicting empirical theories of consciousness—GNWT and RPT—there is a corresponding “front of the head” vs. “back of the head” debate (Boly et al., 2017; Koch et al., 2016a; Koch, Massimini, Boly, & Tononi, 2016b; Naber & Brascamp, 2015; Odegaard, Knight, & Lau, 2017; Safavi, Kapoor, Logothetis, & Panagiotaropoulos, 2014; Sandberg, Frässle, & Pitts, 2016), and it is becoming increasingly clear, not least due to the new no-report paradigms, that frontal activation is in fact not necessary for consciousness (Hurme, Koivisto, Revonsuo, & Railo, 2017; Overgaard, 2017), but rather located in a “posterior hot zone” (Koch et al., 2016a). Evidently, this is in formidable congruence with the source localization results from the ERP and MEG studies reviewed in this thesis.

On the other hand, it is not inconceivable that GNWT can be rescued by modifying it sufficiently to accommodate its threatened predictions. As the “partial awareness hypothesis” (Kouider et al., 2010) testifies, GNWT is versatile enough to integrate findings that appear to be incompatible with it at first. Maybe it is possible, then, to drop the P3b as a central signature of consciousness, and fronto-parietal cortex as the central seat of workspace neurons, and still maintain some viable version of GNWT. However, it is doubtful whether this new theory would have enough in common with GNWT to be recognizable as a version of it. The “partial awareness hypothesis” might be a case in point here. On the one hand, it makes room for findings apparently contradicting GNWT’s view that consciousness is dichotomous. On the other hand, it still seems committed to broadly distributed fronto-parietal activation even for degraded, lower-level representations to become consciously accesses—otherwise, the concept of a “global workspace” becomes almost empty. If local, posterior activation were sufficient for conscious access, there

would remain no meaningful distinction (in terms of localization) between GNWT and other “local” theories like RPT. It seems, then, that the star of GNWT is waning indeed.

However, this does not mean that the case for RPT has been made conclusively. Although it has a lot of evidence in its favor—not least from many of the studies reviewed here—, there are many ambiguities and unresolved issues, too. For example, it is not yet clear precisely how to interpret the findings that speak for the “levels of processing hypothesis” in the context of RPT. One possibility, favored by some authors, is that recurrent processing can occur in degrees and coincides with the onset of phenomenal consciousness, while the more dichotomous awareness ratings occurring at higher levels of (e.g., semantic) processing reflect the workings of reflective/access consciousness (Jimenez et al., 2018; Koivisto et al., 2017). Windey and colleagues themselves propose the “levels of processing hypothesis” as an integration of central elements of both RPT and GNWT, and thus as a first step toward a new theory that transcends both (Windey & Cleeremans, 2015; Windey et al., 2013, 2014).

For another thing, while the VAN is often taken to reflect recurrent processing, and does indeed correlate well with conscious awareness, even here there remains considerable uncertainty. Many authors are still cautious about addressing VAN as an NCC. In particular, Pitts and colleagues are always careful to point out that “it remains possible that the VAN reflects a particular type of attentional process that is necessary for awareness but not identical (in this case the label “visual awareness negativity” would be premature and potentially misleading)” (Schelonka et al., 2017, p. 68). In fact, they argue that VAN might reflect workspace access, and they still prefer theories that hold attention to be necessary for awareness, such as GNWT, to theories that do not, such as RPT. However, their claim that “[u]nfortunately, the VAN as currently understood cannot distinguish between existing theories of consciousness” (Schelonka et al., 2017, p. 69) seems exaggerated; as I argued above, GNWT would have to be modified almost beyond recognition in order to accommodate all its anomalies.

Finally, it is as yet unclear how to integrate the new theoretical developments with the existing theories. Notably, the predictive coding framework awaits its theoretical connection with both RPT and GNWT. On the empirical side, the role of expectations and its influence on ERP correlates of consciousness, in particular on the VAN, need to be investigated further. Before these steps are taken, it will be premature to pass the final judgment over GNWT and RPT. It is well possible that some new theory will at once integrate and transcend them.

The science of consciousness is a rich and burgeoning field, and a single review will always be able to illuminate only a modest portion of this richness. Among the many neighboring topics not covered in this thesis are some that deserve a brief mentioning here.

The theoretical debates about consciousness are vast and deep. A more thorough theoretical investigation of GNWT and RPT, and of the very distinctions of phenomenal and access consciousness, would have discussed the debate about the “overflow” of the former over the latter (Block, 2011; Brown, 2012; Cohen & Dennett, 2011), and the methodological and philosophical problems associated with both sides of that debate.

Empirically, transcranial magnetic stimulation (TMS) is an immensely important approach that has been neglected here, even for the studies in which it was combined with ERP research (e.g., Railo, Revonsuo, & Koivisto, 2015). The effects of TMS are too complex to include them in an unqualified way as merely another way of manipulating conscious awareness, and therefore such studies have not been included in this thesis. TMS would require a review of its own (see Rutiku & Bachmann (2017) for a brief review of “perturbational approaches”). The same goes for studies using other imaging methods, such as fMRI and positron emission tomography (PET).

Overall, this thesis confirms the VAN component as the earliest and most reliable marker of visual awareness, and at the same time it provides an overview of the current state of ERP and MEG research on visual consciousness. However, it has also become clear that the issue may be far more

complex than that. New paradigms have unearthed a number of complications, and it seems that the “early” vs. “late” debate has been transformed rather than settled. For example, maybe, consciousness is sometimes early and sometimes late, depending on the presence or absence of expectations (Melloni et al., 2011). Maybe consciousness is early for detection, but late for identification (Koivisto et al., 2017). Or maybe a combination of many such factors influences the timing of the onset of consciousness. Many are the developments, and there is no shortage of experimental work to be done. New questions abound, and old studies have to be reevaluated in the light of new findings and possible confounds. Hopefully, by discussing a large number of studies within the same context, and presenting open questions in an accessible synthesis, this thesis will help to facilitate new research and spawn new experiments.

References

- Andersen, L. M., Pedersen, M. N., Sandberg, K., & Overgaard, M. (2016). Occipital MEG Activity in the Early Time Range (<300 ms) Predicts Graded Changes in Perceptual Consciousness. *Cerebral Cortex*, *26*(6), 2677–2688. <https://doi.org/10.1093/cercor/bhv108>
- Aru, J., & Bachmann, T. (2009). Occipital EEG correlates of conscious awareness when subjective target shine-through and effective visual masking are compared: Bifocal early increase in gamma power and speed-up of P1. *Brain Research*, *1271*, 60–73. <https://doi.org/10.1016/j.brainres.2008.12.085>
- Aru, J., & Bachmann, T. (2013). Phenomenal awareness can emerge without attention. *Frontiers in Human Neuroscience*, *7*. <https://doi.org/10.3389/fnhum.2013.00891>
- Aru, J., Bachmann, T., Singer, W., & Melloni, L. (2012). Distilling the neural correlates of consciousness. *Neuroscience & Biobehavioral Reviews*, *36*(2), 737–746. <https://doi.org/10.1016/j.neubiorev.2011.12.003>
- Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
- Babiloni, C., Marzano, N., Soricelli, A., Cordone, S., Millán-Calenti, J. C., Del Percio, C., & Buján, A. (2016). Cortical Neural Synchronization Underlies Primary Visual Consciousness of Qualia: Evidence from Event-Related Potentials. *Frontiers in Human Neuroscience*, *10*. <https://doi.org/10.3389/fnhum.2016.00310>
- Babiloni, C., Vecchio, F., Miriello, M., Romani, G. L., & Rossini, P. M. (2006). Visuo-spatial Consciousness and Parieto-occipital Areas: A High-resolution EEG Study. *Cerebral Cortex*, *16*(1), 37–46. <https://doi.org/10.1093/cercor/bhi082>
- Bachmann, T. (2000). *Microgenetic Approach to the Conscious Mind*. Amsterdam/Philadelphia: John Benjamins Publishing.
- Bachmann, T. (2009). Finding ERP-signatures of target awareness: Puzzle persists because of experimental co-variation of the objective and subjective variables. *Consciousness and Cognition*, *18*(3), 804–808. <https://doi.org/10.1016/j.concog.2009.02.011>
- Bachmann, T., & Francis, G. (2013). *Visual Masking: Studying Perception, Attention, and Consciousness*. Oxford/Waltham: Academic Press.

- Bachmann, T., & Hudetz, A. G. (2014). It is time to combine the two main traditions in the research on the neural correlates of consciousness: $C = L \times D$. *Frontiers in Psychology*, 5. <https://doi.org/10.3389/fpsyg.2014.00940>
- Batterink, L., Karns, C. M., & Neville, H. (2012). Dissociable Mechanisms Supporting Awareness: The P300 and Gamma in a Linguistic Attentional Blink Task. *Cerebral Cortex*, 22(12), 2733–2744. <https://doi.org/10.1093/cercor/bhr346>
- Bendixen, A., SanMiguel, I., & Schröger, E. (2012). Early electrophysiological indicators for predictive processing in audition: A review. *International Journal of Psychophysiology*, 83(2), 120–131. <https://doi.org/10.1016/j.ijpsycho.2011.08.003>
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227–247. <https://doi.org/10.1017/S0140525X00038188>
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 15(12), 567–575. <https://doi.org/10.1016/j.tics.2011.11.001>
- Boly, M., Massimini, M., Tsuchiya, N., Postle, B. R., Koch, C., & Tononi, G. (2017). Are the Neural Correlates of Consciousness in the Front or in the Back of the Cerebral Cortex? Clinical and Neuroimaging Evidence. *Journal of Neuroscience*, 37(40), 9603–9613. <https://doi.org/10.1523/JNEUROSCI.3218-16.2017>
- Boncompagni, G., & Cosmelli, D. (2018). Neural Correlates of Conscious Motion Perception. *Frontiers in Human Neuroscience*, 12. <https://doi.org/10.3389/fnhum.2018.00355>
- Brown, R. (2012). The myth of phenomenological overflow. *Consciousness and Cognition*, 21(2), 599–604. <https://doi.org/10.1016/j.concog.2011.06.005>
- Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., ... Massimini, M. (2013). A Theoretically Based Index of Consciousness Independent of Sensory Processing and Behavior. *Science Translational Medicine*, 5(198), 198ra105. <https://doi.org/10.1126/scitranslmed.3006294>
- Chalmers, D. (2007). The Hard Problem of Consciousness. In M. Velmans & S. Schneider (Eds.), *The Blackwell Companion to Consciousness* (pp. 223–235). <https://doi.org/10.1002/9780470751466.ch18>

- Chen, Y., Wang, X., Yu, Y., & Liu, Y. (2017). Dissociable Electroencephalograph Correlates of Visual Awareness and Feature-Based Attention. *Frontiers in Neuroscience, 11*.
<https://doi.org/10.3389/fnins.2017.00633>
- Chica, A. B., Lasaponara, S., Lupiáñez, J., Doricchi, F., & Bartolomeo, P. (2010). Exogenous attention can capture perceptual consciousness: ERP and behavioural evidence. *NeuroImage, 51*(3), 1205–1212. <https://doi.org/10.1016/j.neuroimage.2010.03.002>
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181–204.
<https://doi.org/10.1017/S0140525X12000477>
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford/New York: Oxford University Press.
- Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012a). Response to Tsuchiya et al.: considering endogenous and exogenous attention. *Trends in Cognitive Sciences, 16*(11), 528.
<https://doi.org/10.1016/j.tics.2012.09.002>
- Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012b). The attentional requirements of consciousness. *Trends in Cognitive Sciences, 16*(8), 411–417.
<https://doi.org/10.1016/j.tics.2012.06.013>
- Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences, 15*(8), 358–364. <https://doi.org/10.1016/j.tics.2011.06.008>
- Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences, 2*, 263–275.
- Crick, F., & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex, 8*(2), 97–107.
<https://doi.org/10.1093/cercor/8.2.97>
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience, 6*, 119–126.
<https://doi.org/10.1038/nn0203-119>
- Davoodi, R., Moradi, M. H., & Yoonessi, A. (2015). Dissociation Between Attention and Consciousness During a Novel Task: An ERP Study. *Neurophysiology, 47*(2), 144–154.
<https://doi.org/10.1007/s11062-015-9511-4>

- de Graaf, T. A., & Sack, A. T. (2015). On the various neural correlates of consciousness: are they distinguishable? In S. M. Miller (Ed.), *The Constitution of Phenomenal Consciousness. Toward a Science and Theory* (pp. 177–204). Amsterdam/Philadelphia: John Benjamins Publishing Company.
- Dehaene, S. (2014). *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. New York: Penguin.
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70(2), 200–227. <https://doi.org/10.1016/j.neuron.2011.03.018>
- Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95(24), 14529–14534. <https://doi.org/10.1073/pnas.95.24.14529>
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1), 1–37. [https://doi.org/10.1016/S0010-0277\(00\)00123-2](https://doi.org/10.1016/S0010-0277(00)00123-2)
- Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain Dynamics Underlying the Nonlinear Threshold for Access to Consciousness. *PLOS Biology*, 5(10), e260. <https://doi.org/10.1371/journal.pbio.0050260>
- Del Zotto, M., & Pegna, A. J. (2015). Processing of masked and unmasked emotional faces under different attentional conditions: an electrophysiological investigation. *Frontiers in Psychology*, 6. <https://doi.org/10.3389/fpsyg.2015.01691>
- Di Perri, C., Thibaut, A., Heine, L., Soddu, A., Demertzi, A., & Laureys, S. (2014). Measuring consciousness in coma and related states. *World Journal of Radiology*, 6(8), 589–597. <https://doi.org/10.4329/wjr.v6.i8.589>
- Donchin, E. (1981). Surprise! ... Surprise? *Psychophysiology*, 18(5), 493–513. <https://doi.org/10.1111/j.1469-8986.1981.tb01815.x>
- Eklund, R., & Wiens, S. (2018). Visual awareness negativity is an early neural correlate of awareness: A preregistered study with two Gabor sizes. *Cognitive, Affective, & Behavioral Neuroscience*, 18(1), 176–188. <https://doi.org/10.3758/s13415-018-0562-z>

- Eysenck, M. W., & Keane, M. T. (2015). *Cognitive Psychology: A Student's Handbook: Seventh Edition*. New York: Psychology Press.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, *13*(7), 293–301. <https://doi.org/10.1016/j.tics.2009.04.005>
- Fu, Q., Liu, Y.-J., Dienes, Z., Wu, J., Chen, W., & Fu, X. (2017). Neural Correlates of Subjective Awareness for Natural Scene Categorization of Color Photographs and Line-Drawings. *Frontiers in Psychology*, *8*. <https://doi.org/10.3389/fpsyg.2017.00210>
- Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (2019). *Cognitive Neuroscience: The Biology of the Mind: Fifth Edition*. New York/London: W. W. Norton.
- Gennaro, R. J. (2018). Representational Theories of Consciousness. In R. J. Gennaro (Ed.), *The Routledge Handbook of Consciousness* (pp. 107–121). New York/London: Routledge.
- Hohwy, J. (2009). The neural correlates of consciousness: New experimental approaches needed? *Consciousness and Cognition*, *18*(2), 428–438. <https://doi.org/10.1016/j.concog.2009.02.006>
- Hohwy, J. (2013). *The Predictive Mind*. Oxford/New York: Oxford University Press.
- Hohwy, J. (2015). The Neural Organ Explains the Mind. In T. Metzinger & J. M. Windt (Eds.), *Open MIND*. <https://doi.org/10.15502/9783958570016>
- Hurme, M., Koivisto, M., Revonsuo, A., & Railo, H. (2017). Early processing in primary visual cortex is necessary for conscious and unconscious vision while late processing is necessary only for conscious vision in neurologically healthy humans. *NeuroImage*, *150*, 230–238. <https://doi.org/10.1016/j.neuroimage.2017.02.060>
- Intaitė, M., Koivisto, M., Rukšėnas, O., & Revonsuo, A. (2010). Reversal negativity and bistable stimuli: Attention, awareness, or something else? *Brain and Cognition*, *74*(1), 24–34. <https://doi.org/10.1016/j.bandc.2010.06.002>
- Jimenez, M., Grassini, S., Montoro, P. R., Luna, D., & Koivisto, M. (2018). Neural correlates of visual awareness at stimulus low vs. high-levels of processing. *Neuropsychologia*, *121*, 144–152. <https://doi.org/10.1016/j.neuropsychologia.2018.11.001>

- Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, *14*(2), 201–211. <https://doi.org/10.3758/BF03212378>
- Kaernbach, C., Schröger, E., Jacobsen, T., & Roeber, U. (1999). Effects of consciousness on human brain waves following binocular rivalry. *NeuroReport*, *10*(4), 713.
- Kersten, D., & Yuille, A. (2014). Inferential Models of the Visual Cortical Hierarchy. In M. S. Gazzaniga & G. R. Mangun (Eds.), *The Cognitive Neurosciences: Fifth Edition* (pp. 389–404). Cambridge/London: MIT Press.
- Kleinschmidt, A., Büchel, C., Hutton, C., Friston, K. J., & Frackowiak, R. S. J. (2002). The Neural Structures Expressing Perceptual Hysteresis in Visual Letter Recognition. *Neuron*, *34*(4), 659–666. [https://doi.org/10.1016/S0896-6273\(02\)00694-3](https://doi.org/10.1016/S0896-6273(02)00694-3)
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016a). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, *17*(5), 307–321. <https://doi.org/10.1038/nrn.2016.22>
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016b). Posterior and anterior cortex — where is the difference that makes the difference? *Nature Reviews Neuroscience*, *17*(10), 666. <https://doi.org/10.1038/nrn.2016.105>
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, *11*(1), 16–22. <https://doi.org/10.1016/j.tics.2006.10.012>
- Koch, C., & Tsuchiya, N. (2012). Attention and consciousness: related yet different. *Trends in Cognitive Sciences*, *16*(2), 103–105. <https://doi.org/10.1016/j.tics.2011.11.012>
- Koivisto, M., & Grassini, S. (2016). Neural processing around 200ms after stimulus-onset correlates with subjective visual awareness. *Neuropsychologia*, *84*, 235–243. <https://doi.org/10.1016/j.neuropsychologia.2016.02.024>
- Koivisto, M., Grassini, S., Salminen-Vaparanta, N., & Revonsuo, A. (2017). Different Electrophysiological Correlates of Visual Awareness for Detection and Identification. *Journal of Cognitive Neuroscience*, *29*(9), 1621–1631. https://doi.org/10.1162/jocn_a_01149
- Koivisto, M., Kainulainen, P., & Revonsuo, A. (2009). The relationship between awareness and attention: Evidence from ERP responses. *Neuropsychologia*, *47*(13), 2891–2899. <https://doi.org/10.1016/j.neuropsychologia.2009.06.016>

- Koivisto, M., Kastrati, G., & Revonsuo, A. (2013). Recurrent Processing Enhances Visual Awareness but Is Not Necessary for Fast Categorization of Natural Scenes. *Journal of Cognitive Neuroscience*, *26*(2), 223–231. https://doi.org/10.1162/jocn_a_00486
- Koivisto, M., Lähteenmäki, M., Sørensen, T. A., Vangkilde, S., Overgaard, M., & Revonsuo, A. (2008). The earliest electrophysiological correlate of visual awareness? *Brain and Cognition*, *66*(1), 91–103. <https://doi.org/10.1016/j.bandc.2007.05.010>
- Koivisto, M., & Revonsuo, A. (2003). An ERP study of change detection, change blindness, and visual awareness. *Psychophysiology*, *40*(3), 423–429. <https://doi.org/10.1111/1469-8986.00044>
- Koivisto, M., & Revonsuo, A. (2005). Prechange event-related potentials predict change blindness in various attention conditions. *NeuroReport*, *16*(8), 869.
- Koivisto, M., & Revonsuo, A. (2007). Electrophysiological correlates of visual consciousness and selective attention. *NeuroReport*, *18*(8), 753.
<https://doi.org/10.1097/WNR.0b013e3280c143c8>
- Koivisto, M., & Revonsuo, A. (2008a). Comparison of event-related potentials in attentional blink and repetition blindness. *Brain Research*, *1189*, 115–126.
<https://doi.org/10.1016/j.brainres.2007.10.082>
- Koivisto, M., & Revonsuo, A. (2008b). The role of selective attention in visual awareness of stimulus features: Electrophysiological studies. *Cognitive, Affective, & Behavioral Neuroscience*, *8*(2), 195–210. <https://doi.org/10.3758/CABN.8.2.195>
- Koivisto, M., & Revonsuo, A. (2010). Event-related brain potential correlates of visual awareness. *Neuroscience & Biobehavioral Reviews*, *34*(6), 922–934.
<https://doi.org/10.1016/j.neubiorev.2009.12.002>
- Koivisto, M., Revonsuo, A., & Lehtonen, M. (2006). Independence of Visual Awareness from the Scope of Attention: an Electrophysiological Study. *Cerebral Cortex*, *16*(3), 415–424. <https://doi.org/10.1093/cercor/bhi121>
- Koivisto, M., Revonsuo, A., & Salminen, N. (2005). Independence of visual awareness from attention at early processing stages. *NeuroReport*, *16*(8), 817.

- Koivisto, M., Ruohola, M., Vahtera, A., Lehmusvuo, T., & Intaite, M. (2018). The effects of working memory load on visual awareness and its electrophysiological correlates. *Neuropsychologia*, *120*, 86–96. <https://doi.org/10.1016/j.neuropsychologia.2018.10.011>
- Koivisto, M., Salminen-Vaparanta, N., Grassini, S., & Revonsuo, A. (2016). Subjective visual awareness emerges prior to P3. *European Journal of Neuroscience*, *43*(12), 1601–1611. <https://doi.org/10.1111/ejn.13264>
- 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. <https://doi.org/10.1016/j.tics.2010.04.006>
- Lamme, V. A. F. (2000). Neural Mechanisms of Visual Awareness: A Linking Proposition. *Brain and Mind*, *1*(3), 385–406. <https://doi.org/10.1023/A:1011569019782>
- Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*(1), 12–18. [https://doi.org/10.1016/S1364-6613\(02\)00013-X](https://doi.org/10.1016/S1364-6613(02)00013-X)
- Lamme, V. A. F. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks*, *17*(5), 861–872. <https://doi.org/10.1016/j.neunet.2004.02.005>
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, *10*(11), 494–501. <https://doi.org/10.1016/j.tics.2006.09.001>
- Lamme, V. A. F. (2010a). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, *1*(3), 204–220. <https://doi.org/10.1080/17588921003731586>
- Lamme, V. A. F. (2010b). What introspection has to offer, and where its limits lie. *Cognitive Neuroscience*, *1*(3), 232–240. <https://doi.org/10.1080/17588928.2010.502224>
- Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*(11), 571–579. [https://doi.org/10.1016/S0166-2236\(00\)01657-X](https://doi.org/10.1016/S0166-2236(00)01657-X)
- Lamy, D., Salti, M., & Bar-Haim, Y. (2008). Neural Correlates of Subjective Awareness and Unconscious Processing: An ERP Study. *Journal of Cognitive Neuroscience*, *21*(7), 1435–1446. <https://doi.org/10.1162/jocn.2009.21064>

- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, *15*(8), 365–373.
<https://doi.org/10.1016/j.tics.2011.05.009>
- Laureys, S. (2005). The neural correlate of (un)awareness: lessons from the vegetative state. *Trends in Cognitive Sciences*, *9*(12), 556–559. <https://doi.org/10.1016/j.tics.2005.10.010>
- Leopold, D. A., Wilke, M., Maier, A., & Logothetis, N. K. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience*, *5*(6), 605. <https://doi.org/10.1038/nn0602-851>
- Lipton, P. (2004). *Inference to the Best Explanation*. London/New York: Routledge.
- Liu, Y., Paradis, A.-L., Yahia-Cherif, L., & Tallon-Baudry, C. (2012). Activity in the lateral occipital cortex between 200 and 300 ms distinguishes between physically identical seen and unseen stimuli. *Frontiers in Human Neuroscience*, *6*. <https://doi.org/10.3389/fnhum.2012.00211>
- Luck, S. J. (2014). *An Introduction to the Event-Related Potential Technique: Second Edition*. Cambridge/London: MIT Press.
- Ludowig, E., Bien, C. G., Elger, C. E., & Rosburg, T. (2010). Two P300 generators in the hippocampal formation. *Hippocampus*, *20*(1), 186–195. <https://doi.org/10.1002/hipo.20603>
- Mack, A., & Rock, I. (1998). *Inattentional Blindness*. Cambridge/London: MIT Press.
- 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*(4), 1386–1396.
<https://doi.org/10.1523/JNEUROSCI.4570-10.2011>
- Miller, S. M. (2013). *The Constitution of Visual Consciousness: Lessons from Binocular Rivalry*. Amsterdam/Philadelphia: John Benjamins Publishing.
- Mudrik, L., & Lamy, D. (2007). Antti Revonsuo, Inner Presence: Consciousness as a Biological Phenomenon. *Pragmatics & Cognition*, *15*(2), 379–387.
<https://doi.org/10.1075/pc.15.2.09mud>
- Naber, M., & Brascamp, J. (2015). Commentary: Is the Frontal Lobe Involved in Conscious Perception? *Frontiers in Psychology*, *6*. <https://doi.org/10.3389/fpsyg.2015.01736>

- Naccache, L., Marti, S., Sitt, J. D., Trübtschek, D., & Berkovitch, L. (2016). Why the P3b is still a plausible correlate of conscious access? A commentary on Silverstein et al., 2015. *Cortex*, 85, 126–128. <https://doi.org/10.1016/j.cortex.2016.04.003>
- Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>
- Noah, S., & Mangun, G. R. (in press). Recent evidence that attention is necessary, but not sufficient, for conscious perception. *Annals of the New York Academy of Sciences*. <https://doi.org/10.1111/nyas.14030>
- Odegaard, B., Knight, R. T., & Lau, H. (2017). Should a Few Null Findings Falsify Prefrontal Theories of Conscious Perception? *Journal of Neuroscience*, 37(40), 9593–9602. <https://doi.org/10.1523/JNEUROSCI.3217-16.2017>
- Ojanen, V., Revonsuo, A., & Sams, M. (2003). Visual awareness of low-contrast stimuli is reflected in event-related brain potentials. *Psychophysiology*, 40(2), 192–197. <https://doi.org/10.1111/1469-8986.00021>
- Olshausen, B. A. (2014). Perception as an Inference Problem. In M. S. Gazzaniga & G. R. Mangun (Eds.), *The Cognitive Neurosciences: Fifth Edition* (pp. 295–304). Cambridge/London: MIT Press.
- Overgaard, M. (2017). The Status and Future of Consciousness Research. *Frontiers in Psychology*, 8. <https://doi.org/10.3389/fpsyg.2017.01719>
- Overgaard, M., & Fazekas, P. (2016). Can No-Report Paradigms Extract True Correlates of Consciousness? *Trends in Cognitive Sciences*, 20(4), 241–242. <https://doi.org/10.1016/j.tics.2016.01.004>
- Overgaard, M., Rote, J., Mouridsen, K., & Ramsøy, T. Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and Cognition*, 15(4), 700–708. <https://doi.org/10.1016/j.concog.2006.04.002>
- Picton, T. W. (1984). Report of Panel II: The ERP and Decision and Memory Processes. In E. Donchin (Ed.), *Cognitive Psychophysiology: Event-Related Potentials and the Study of Cognition* (1st ed., pp. 139–177). Hillsdale/London: Lawrence Erlbaum Associates.

- Picton, T. W. (1992). The P300 Wave of the Human Event-Related Potential. *Journal of Clinical Neurophysiology*, 9(4), 456–479.
- Pins, D., & ffytche, D. (2003). The Neural Correlates of Conscious Vision. *Cerebral Cortex*, 13(5), 461–474. <https://doi.org/10.1093/cercor/13.5.461>
- Pitts, M. A., Lutsyshyna, L. A., & Hillyard, S. A. (2018). The relationship between attention and consciousness: an expanded taxonomy and implications for ‘no-report’ paradigms. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170348. <https://doi.org/10.1098/rstb.2017.0348>
- Pitts, M. A., Martínez, A., & Hillyard, S. A. (2011). Visual Processing of Contour Patterns under Conditions of Inattentive Blindness. *Journal of Cognitive Neuroscience*, 24(2), 287–303. https://doi.org/10.1162/jocn_a_00111
- 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. <https://doi.org/10.3389/fpsyg.2014.01078>
- Pitts, M. A., Padwal, J., Fennelly, D., Martínez, A., & Hillyard, S. A. (2014). Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage*, 101, 337–350. <https://doi.org/10.1016/j.neuroimage.2014.07.024>
- Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, 118(10), 2128–2148. <https://doi.org/10.1016/j.clinph.2007.04.019>
- Railo, H., & Koivisto, M. (2009a). Reply to Bachmann on ERP correlates of visual awareness. *Consciousness and Cognition*, 18(3), 809–810. <https://doi.org/10.1016/j.concog.2009.05.005>
- Railo, H., & Koivisto, M. (2009b). The electrophysiological correlates of stimulus visibility and metacontrast masking. *Consciousness and Cognition*, 18(3), 794–803. <https://doi.org/10.1016/j.concog.2009.01.006>
- 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(3), 972–983. <https://doi.org/10.1016/j.concog.2011.03.019>

- Railo, H., Revonsuo, A., & Koivisto, M. (2015). Behavioral and electrophysiological evidence for fast emergence of visual consciousness. *Neuroscience of Consciousness*, 2015(1).
<https://doi.org/10.1093/nc/niv004>
- Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, 3(1), 1–23.
<https://doi.org/10.1023/B:PHEN.0000041900.30172.e8>
- Revonsuo, A. (2006). *Inner Presence: Consciousness as a Biological Phenomenon*. Cambridge/London: MIT Press.
- Revonsuo, A. (2010). *Consciousness: The Science of Subjectivity*. Hove/New York: Psychology Press.
- Rutiku, R., Aru, J., & Bachmann, T. (2016). General Markers of Conscious Visual Perception and Their Timing. *Frontiers in Human Neuroscience*, 10.
<https://doi.org/10.3389/fnhum.2016.00023>
- Rutiku, R., & Bachmann, T. (2017). Juxtaposing the real-time unfolding of subjective experience and ERP neuromarker dynamics. *Consciousness and Cognition*, 54, 3–19.
<https://doi.org/10.1016/j.concog.2017.05.003>
- Rutiku, R., Martin, M., Bachmann, T., & Aru, J. (2015). Does the P300 reflect conscious perception or its consequences? *Neuroscience*, 298, 180–189.
<https://doi.org/10.1016/j.neuroscience.2015.04.029>
- Safavi, S., Kapoor, V., Logothetis, N. K., & Panagiotaropoulos, T. I. (2014). Is the frontal lobe involved in conscious perception? *Frontiers in Psychology*, 5.
<https://doi.org/10.3389/fpsyg.2014.01063>
- Salti, M., Bar-Haim, Y., & Lamy, D. (2012). The P3 component of the ERP reflects conscious perception, not confidence. *Consciousness and Cognition*, 21(2), 961–968.
<https://doi.org/10.1016/j.concog.2012.01.012>
- Sandberg, K., Bahrami, B., Kanai, R., Barnes, G. R., Overgaard, M., & Rees, G. (2013). Early Visual Responses Predict Conscious Face Perception within and between Subjects during Binocular Rivalry. *Journal of Cognitive Neuroscience* 25(6), 969–985.
https://doi.org/10.1162/jocn_a_00353

- Sandberg, K., Barnes, G. R., Bahrami, B., Kanai, R., Overgaard, M., & Rees, G. (2014). Distinct MEG correlates of conscious experience, perceptual reversals and stabilization during binocular rivalry. *NeuroImage*, *100*, 161–175.
<https://doi.org/10.1016/j.neuroimage.2014.06.023>
- Sandberg, K., Frässle, S., & Pitts, M. (2016). Future directions for identifying the neural correlates of consciousness. *Nature Reviews Neuroscience*, *17*(10), 666.
<https://doi.org/10.1038/nrn.2016.104>
- Schelonka, K., Grafton, C., Canseco-Gonzalez, E., & Pitts, M. A. (2017). ERP signatures of conscious and unconscious word and letter perception in an inattentive blindness paradigm. *Consciousness and Cognition*, *54*, 56–71.
<https://doi.org/10.1016/j.concog.2017.04.009>
- Scrivener, C. L., Malik, A., Marsh, J., Lindner, M., & Roesch, E. (2019). An EEG study of detection without localisation in change blindness. *BioRxiv*, 513697. <https://doi.org/10.1101/513697>
- Sekar, K., Findley, W. M., Poeppel, D., & Llinás, R. R. (2013). Cortical response tracking the conscious experience of threshold duration visual stimuli indicates visual perception is all or none. *Proceedings of the National Academy of Sciences*, *110*(14), 5642–5647.
<https://doi.org/10.1073/pnas.1302229110>
- Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, *8*(10), 1391–1400.
<https://doi.org/10.1038/nn1549>
- Sergent, C., & Dehaene, S. (2004). Is Consciousness a Gradual Phenomenon?: Evidence for an All-or-None Bifurcation During the Attentional Blink. *Psychological Science*, *15*(11), 720–728.
<https://doi.org/10.1111/j.0956-7976.2004.00748.x>
- Sergent, C., & Naccache, L. (2012). Imaging neural signatures of consciousness: ‘What’, ‘When’, ‘Where’ and ‘How’ does it work? *Archives Italiennes de Biologie*, *150*, 91–106.
- Shafto, J. P., & Pitts, M. A. (2015). Neural Signatures of Conscious Face Perception in an Inattentive Blindness Paradigm. *Journal of Neuroscience*, *35*(31), 10940–10948.
<https://doi.org/10.1523/JNEUROSCI.0145-15.2015>

- Silverstein, B. H., Snodgrass, M., Shevrin, H., & Kushwaha, R. (2015). P3b, consciousness, and complex unconscious processing. *Cortex*, *73*, 216–227.
<https://doi.org/10.1016/j.cortex.2015.09.004>
- Silverstein, B. H., Snodgrass, M., Shevrin, H., & Kushwaha, R. (2016). Unconscious P3b and complex unconscious processing: Reply to Naccache et al., 2016. *Cortex*, *85*, 129–132.
<https://doi.org/10.1016/j.cortex.2016.09.007>
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception*, *28*(9), 1059–1074. <https://doi.org/10.1068/p281059>
- Tagliabue, C. F., Mazzi, C., Bagattini, C., & Savazzi, S. (2016). Early Local Activity in Temporal Areas Reflects Graded Content of Visual Perception. *Frontiers in Psychology*, *7*.
<https://doi.org/10.3389/fpsyg.2016.00572>
- Timmermans, B., Windey, B., & Cleeremans, A. (2010). Experiencing more complexity than we can tell. *Cognitive Neuroscience*, *1*(3), 229–230.
<https://doi.org/10.1080/17588928.2010.497586>
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, *5*, 42–22. <https://doi.org/10.1186/1471-2202-5-42>
- Tsuchiya, N., Block, N., & Koch, C. (2012). Top-down attention and consciousness: comment on Cohen et al. *Trends in Cognitive Sciences*, *16*(11), 527.
<https://doi.org/10.1016/j.tics.2012.09.004>
- Tsuchiya, N., Frässle, S., Wilke, M., & Lamme, V. A. F. (2016). No-Report and Report-Based Paradigms Jointly Unravel the NCC: Response to Overgaard and Fazekas. *Trends in Cognitive Sciences*, *20*(4), 242–243. <https://doi.org/10.1016/j.tics.2016.01.006>
- Tsuchiya, N., & Koch, C. (2016). The Relationship Between Consciousness and Top-Down Attention. In S. Laureys, O. Gosseries, & G. Tononi (Eds.), *The Neurology of Consciousness* (2nd ed., pp. 71–91). Amsterdam: Academic Press.
- Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. F. (2015). No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences*, *19*(12), 757–770. <https://doi.org/10.1016/j.tics.2015.10.002>

- Van Boxtel, J. J. A., & Tsuchiya, N. (2015). De-confounding the neural constitution of phenomenal consciousness from attention, report and memory. In S. M. Miller (Ed.), *The Constitution of Phenomenal Consciousness. Toward a Science and Theory* (pp. 81–103). Amsterdam/Philadelphia: John Benjamins Publishing Company.
- Van Boxtel, J. J. A., Tsuchiya, N., & Koch, C. (2010). Consciousness and Attention: On Sufficiency and Necessity. *Frontiers in Psychology, 1*. <https://doi.org/10.3389/fpsyg.2010.00217>
- Vanni, S., Revonsuo, A., Saarinen, J., & Hari, R. (1996). Visual awareness of objects correlates with activity of right occipital cortex. *NeuroReport, 8*(1), 183.
- Volpe, U., Mucci, A., Bucci, P., Merlotti, E., Galderisi, S., & Maj, M. (2007). The cortical generators of P3a and P3b: A LORETA study. *Brain Research Bulletin, 73*(4), 220–230. <https://doi.org/10.1016/j.brainresbull.2007.03.003>
- Weller, P., Rabovsky, M., & Rahman, R. A. (2019). Semantic Knowledge Enhances Conscious Awareness of Visual Objects. *Journal of Cognitive Neuroscience, 1–11*. https://doi.org/10.1162/jocn_a_01404
- Wilenius, M. E., & Revonsuo, A. T. (2007). Timing of the earliest ERP correlate of visual awareness. *Psychophysiology, 44*(5), 703–710. <https://doi.org/10.1111/j.1469-8986.2007.00546.x>
- Wilenius-Emet, M., Revonsuo, A., & Ojanen, V. (2004). An electrophysiological correlate of human visual awareness. *Neuroscience Letters, 354*(1), 38–41. <https://doi.org/10.1016/j.neulet.2003.09.060>
- Windey, B., & Cleeremans, A. (2015). Consciousness as a graded and an all-or-none phenomenon: A conceptual analysis. *Consciousness and Cognition, 35*, 185–191. <https://doi.org/10.1016/j.concog.2015.03.002>
- Windey, B., Gevers, W., & Cleeremans, A. (2013). Subjective visibility depends on level of processing. *Cognition, 129*(2), 404–409. <https://doi.org/10.1016/j.cognition.2013.07.012>
- Windey, B., Vermeiren, A., Atas, A., & Cleeremans, A. (2014). The graded and dichotomous nature of visual awareness. *Philosophical Transactions of the Royal Society B: Biological Sciences, 369*(1641), 20130282. <https://doi.org/10.1098/rstb.2013.0282>

- Winkler, I., & Czigler, I. (2012). Evidence from auditory and visual event-related potential (ERP) studies of deviance detection (MMN and vMMN) linking predictive coding theories and perceptual object representations. *International Journal of Psychophysiology*, *83*(2), 132–143. <https://doi.org/10.1016/j.ijpsycho.2011.10.001>
- Ye, M., & Lyu, Y. (2019). Later Positivity Reflects Post-perceptual Processes: Evidence From Immediate Detection and Delayed Detection Tasks. *Frontiers in Psychology*, *10*. <https://doi.org/10.3389/fpsyg.2019.00082>
- Zeki, S., & Bartels, A. (1999). Toward a Theory of Visual Consciousness. *Consciousness and Cognition*, *8*(2), 225–259. <https://doi.org/10.1006/ccog.1999.0390>