

*SHORT REPORT*

## From sensation to perception: Using multivariate classification of visual illusions to identify neural correlates of conscious awareness in space and time

---

Hinze Hogendoorn

Helmholtz Institute, Neuroscience & Cognition Utrecht, Experimental Psychology Division, Utrecht University, Heidelberglaan 1, 3584 CS Utrecht, The Netherlands;  
e-mail: j.h.a.hogendoorn@uu.nl

Received 5 August 2014, in revised form 26 November 2014, published online 12 January 2015

---

**Abstract.** An important goal of cognitive neuroscience is understanding the neural underpinnings of conscious awareness. Although the low-level processing of sensory input is well understood in most modalities, it remains a challenge to understand how the brain translates such input into conscious awareness. Here, I argue that the application of multivariate pattern classification techniques to neuroimaging data acquired while observers experience perceptual illusions provides a unique way to dissociate sensory mechanisms from mechanisms underlying conscious awareness. Using this approach, it is possible to directly compare patterns of neural activity that correspond to the contents of awareness, independent from changes in sensory input, and to track these neural representations over time at high temporal resolution. I highlight five recent studies using this approach, and provide practical considerations and limitations for future implementations.

**Keywords:** perception, awareness, classification, decoding, illusions

One of the important ambitions of research into perception is to identify the mechanisms involved in generating conscious awareness of stimulus properties. For most modalities, the physiological mechanisms underlying detection of a stimulus itself (in terms of converting stimulus energy into neural signals) are relatively well understood. What is much less well understood is how those sensory signals produce a sensory experience in our conscious awareness. This is currently the topic of considerable scientific interest.

The advent of noninvasive neuroimaging techniques such as electro- and magnetoencephalography (EEG and MEG) and functional magnetic resonance imaging (fMRI) has given us powerful tools to investigate perceptual mechanisms in the human brain (for a review of multivariate neuroimaging approaches to the study of consciousness, see Haynes, 2009). However, isolating mechanisms related to the processing of low-level stimulus information from the mechanisms involved in realizing conscious awareness of that information is made difficult by the fact that the two are naturally highly correlated. Because it is difficult to influence what an observer perceives without also changing the actual stimulus presented to the observer, it is difficult to compare the neural signatures of two or more percepts without introducing differences in the processing of low-level information.

Nevertheless, a number of paradigms have been devised to study changes in an observer's conscious percept without concurrent changes in the stimulus. One such approach is to use bistable stimuli, in which two (or more) stable perceptual interpretations are possible for a given stimulus. The Necker cube (Necker, 1832) and binocular rivalry are two classic examples. In this paradigm an observer views a static image, and the observer's percept switches between possible interpretations. This makes it possible to look for patterns of neural activity that correlate to each of the different interpretations. For example, Tong, Nakayama, Vaughan, and Kanwisher (1998) showed that the timecourse of the BOLD signal in the fusiform face

---

area and the parahippocampal place area closely followed the perceptual interpretation of an observer viewing images of faces and houses during binocular rivalry. Although the conclusion that these two relatively high-level brain areas are involved in the conscious perception of faces and houses is consistent with other experimental paradigms (eg Haynes, 2009), more generally there are limitations to the use of bistable stimuli for studying awareness. First of all, it is as yet unclear at which point(s) in the visual pathway switches between bistable stimuli take place. This means that the activity of any area downstream of such switches will correlate to the percept of the observer, even if it is not actually involved in realizing awareness. In fact, there is evidence that interocular dominance is resolved as early as V1 (eg Polonsky, Blake, Braun, & Heeger, 2000), suggesting that substantial low-level processing will fluctuate along with the observer's interpretation. In the case of Tong et al.'s face-house rivalry, this could lead to the misleading conclusion that V1 is already involved in the conscious perception of faces and houses. Furthermore, because perceptual switches occur stochastically, it is difficult to know or control which interpretation the observer perceives without explicitly asking the observer. Explicitly asking the observer requires the observer to prepare some kind of response, and because the response necessarily correlates with the percept, the neural activity associated with the preparation of the response will confound the recording of activity related to the percept. This again makes the isolation of neural activity corresponding to visual awareness *per se* difficult.

Another approach is to use stimuli presented at or near detection threshold, so that identical sensory input leads to conscious perception on some trials but not on others. In this paradigm trials on which a target is detected can be compared with trials where a target is missed. However, an important limitation of this paradigm is that trial-by-trial fluctuations in an observer's attentiveness covary with detection performance. These could be random (as in lapses or fatigue) or rhythmic (recent evidence suggests that, in the visual domain, attention samples visual information at a rate of 4–8 Hz; eg VanRullen, Carlson, & Cavanagh, 2007). In either case a contrast of hits and misses is likely to be confounded by attentional mechanisms or the effects of attentional availability. Although attention is inextricably linked to awareness, the two are both conceptually and experimentally dissociable (eg Lamme, 2003).

A third approach is to exploit trial-by-trial fluctuations in timing, rather than performance. For example, Kutas, McCarthy, and Donchin (1977) demonstrated several decades ago that the latency of the P300 event-related potential component covaried with response latency on a decision task. However, when correlating neural latencies to response latencies, any activity related to response preparation will also covary, hampering the isolation of only those mechanisms involved in realizing awareness. We recently modified this paradigm such that a speeded behavioural response was not required, and still identified two scalp locations (contralateral parieto-occipital and fronto-central) where the latency of the evoked activity covaried with attentional selection latency (Hogendoorn, Carlson, & Verstraten, 2011). Nevertheless, the limited spatial resolution and the necessary temporal variability make it difficult to isolate which neural structures or processes underlie such findings.

Here, I advocate a new approach to studying the neural mechanisms involved in realizing conscious awareness: applying multivariate pattern classification techniques to visual illusions. Multivariate pattern classification is an approach first developed to exploit subtle spatial structure in fMRI data (eg Haxby et al., 2001), but has since also been applied to neuroimaging data from other techniques, in both humans and other animals (see King & Dehaene, 2014, for a review). This approach is especially suitable for the study of conscious awareness because conscious awareness is often thought to involve distributed, rather than focal, patterns of activation (eg Rees, 2007; Zeki & Bartels, 1998). I will highlight examples from the visual domain, but the principle could unreservedly be adapted to any sensory modality. Accordingly, the aim of this paper is not to review the state of the art in any of

---

the fields from which these examples are drawn. Rather, they serve to illustrate the kinds of questions that the technique can be used to answer, and the unique insight that the technique can provide into the temporal dynamics of the underlying representations.

The idea of pattern classification is that a subset of a neuroimaging dataset is used to train a classifier—for example, to discriminate the neural activity patterns associated with faces, on the one hand, and houses, on the other. Essentially, the patterns of activity evoked by individual trials or time points are treated as points in  $N$ -D space, where  $N$  is the number of available voxels, electrodes, SQUIDS, or other spatially separate acquisition sites. The trials used to train the classifier therefore form two clouds of points in this  $N$ -D space, and the classifier searches for a hyperplane that optimally separates these two clouds of points. Subsequently, the classifier is tested on the remaining data, essentially asking for each trial or time point on which side of the hyperplane that point falls—in the above example, is the pattern of activity more similar to the pattern associated with a face or a house? Above-chance classification performance on this task is indicative of the availability of systematic information in the dataset, discriminating between the two types of trial being compared. The important advantage of this approach over some other approaches is that it does not require averaging over space, time, trials, or observers. All available systematic information is combined across all acquisition sites without any assumption about the spatial distribution or polarity of that information, or about the regularity of that distribution across different observers. Furthermore, it can be applied to individual time points, which gives unique insight into the temporal structure of evolving neural representations when combined with high-temporal resolution neuroimaging techniques such as EEG and MEG. Finally, classification techniques can reveal systematic information in individual trials even in the case of EEG, where the signal-to-noise ratio is so low that averaging over trials is often thought to be essential (eg Hogendoorn & Verstraten, 2013a).

As such, pattern classification is a powerful technique to demonstrate the availability of discriminating information, without requiring many assumptions about the spatial or temporal extent of that information. More importantly, a classifier can be trained to discriminate two classes of stimuli (eg faces vs houses) and then tested on a third class of stimuli (eg cars). Essentially, this cross-classification approach reveals to which of the first two stimuli the third stimulus is most similar, in terms of its neural representation. This idea has recently been used very effectively to study object-space—the organization of neural representations of a wide range of visual objects. By plotting the similarity of patterns of neural activation for a wide range of different visual stimuli, it is possible to explore which kinds of objects are similar in terms of their neural representation, and which are different. This allows the experimenter to construct a representational distance matrix, which visualizes what kinds of object categories are independently represented in the visual brain (Carlson, Tovar, Alink, & Kriegeskorte, 2013; Kriegeskorte & Kievit, 2013).

However, the cross-classification approach is particularly valuable when applied to (visual) illusions. In an illusion a sensory stimulus is by definition perceived as something different from what it actually is. However, the actual low-level processing of incoming information is presumably veridical until a certain stage in the processing hierarchy. Using cross-classification, it is possible to ask where and when the neural representation of a given illusory stimulus corresponds to the veridical stimulus, and where and when it corresponds to what the observer actually perceives. For instance, one might train a classifier to discriminate between the presentation of a small disc and the presentation of a large disc. Subsequently, this classifier could be tested on trials in which a large disc is made to appear smaller (for example, using the Ebbinghaus illusion). One might expect that at early time points the classifier assigns these illusion trials to the ‘large disc’ category (since that most closely matches the afferent sensory information). At later time points it might instead assign these trials to the ‘small disc’

---

condition, since that matches the conscious percept. In this way, it allows sensory processes to be dissociated from processes associated with awareness, in a way that no other technique allows.

Using fMRI, Maus and colleagues (Maus, Fischer, & Whitney, 2013) recently used this approach to study the flash-drag effect—an illusion where the perceived position of a flashed stationary object is shifted in the direction of nearby motion (Fukiage, Whitney, & Murakami, 2011). They analyzed the patterns of fMRI activity in five different visual areas, and showed that in two of these areas the pattern of neural activity evoked by the illusory stimulus was more similar to a control nonillusion stimulus shifted in the same direction than to a control stimulus shifted in the opposite direction. In other words, the pattern of activity in these two areas (V3A and MT) corresponded to the perceived location of the stimulus, rather than the actual location. The same was not true for three earlier visual areas (V1, V2, and V3)—activity in these areas corresponded to the veridical location of the stimulus. This finding is not only important for our understanding of position coding, but also points to these brain areas as neural correlates of conscious awareness—at least, awareness of visual position.

We recently approached a different illusion, the motion aftereffect (MAE), in a similar way (Hogendoorn & Verstraten, 2013b). In this famous illusion viewing motion in a given direction for a sustained period of time causes a subsequent stationary object or texture to appear to move in the opposite direction (Anstis, Verstraten, & Mather, 1998). A number of models of motion processing have been put forth to explain this illusion, based largely on animal physiology and human psychophysical experiments (eg Van Wezel & Britten, 2002). We asked the question whether (and if so, where) in the human brain the pattern of neural activity associated with illusory motion in a given direction corresponds to the pattern of neural activity recorded when actual motion is presented in the same direction. Using multivariate pattern classification, we showed that both motion direction and the observer's adaptation state can be decoded from several visual areas (V1, V2, V3, V4, and MT) but that only MT showed a change in classification performance during the actual MAE. However, in terms of neural activity, MAE motion did not resemble real motion in the same direction in any visual area; instead, MT showed a suppression of the population response coding for the adapted direction. This pattern of results indicated that motion direction is encoded in MT as a shift in the response of large populations of neurons, rather than in the isolated activity of only those neurons tuned for the perceived direction. In this way, applying pattern classification to a visual illusion gives insight into the organization of neural representations.

Interestingly, Crowe, Leuthold, and Georgopoulos (2010) applied multivariate techniques to MEG data recorded whilst observers viewed an annulus with randomly moving dots. They found that a classifier trained on trials with coherently moving dots (true coherent motion) could accurately decode the perceived direction of motion on trials where observers reported coherent motion even though no coherent motion was actually presented (illusory coherent motion). Moreover, when decoding trials *with* coherent motion, classification performance was independent of the degree of motion coherence. This further reinforces the idea that the neural activity underlying classification performance in this study is associated with the conscious percept (which remains comparable as motion coherence increases beyond the initial detection threshold) rather than with the processing of incoming sensory evidence (which would be expected to depend on the degree of coherence).

Because the temporal resolution of fMRI is limited by the hemodynamic response function, information about the temporal dynamics of evolving neural representations is more easily recovered from EEG and MEG, and other more invasive electrophysiological measures. Indeed, King and Dehaene (2014) recently put forth a compelling account of how cross-classification in the time domain (training a classifier on one time point in a trial, and testing it on a different time point) can give valuable insight into the temporal dynamics

---

of neural representations. For example, phasic and sustained patterns of activity have been argued to be differentially involved in realizing and maintaining conscious awareness (Tononi & Koch, 2008). By evaluating all combinations of training and test time points, one can construct a temporal generalization matrix, which can reveal whether neural representations in a given task might be sequential, oscillatory, sustained, re-entrant, or follow some other temporal structure. The high temporal resolution of these methods is also very valuable when applied to visual illusions, because it makes it possible to track the evolution of the neural response to a stimulus over time. This in turn makes it possible to identify time windows in which neural responses correspond to bottom-up sensory input, and time windows in which neural responses correspond to the actual percept that the observer becomes aware of.

We recently applied multivariate classification to EEG data recorded while observers viewed a variant of the flash-grab illusion (Cavanagh & Anstis, 2013; Hogendoorn & Verstraten, 2013a). In this illusion a flashed object is presented on top of a moving texture at the moment that the texture reverses direction, which causes the perceived position of the flashed object to shift in the direction of the subsequent motion of the texture. This allowed us to create situations where two stimuli presented in identical positions were perceived in different positions, and conversely where two stimuli presented in different positions were perceived in similar positions. The pattern of classification performance strongly corresponded to the observer's illusory percept. For example, the classifier was better able to discriminate between two discs presented in adjacent locations when the flash-grab illusion shifted the perceived positions of the two discs apart than when it shifted them together. Furthermore, a classifier trained to discriminate between two discs presented in different locations performed above chance at classifying trials in which a disc was actually presented in a third location, but appeared shifted in the direction of one of the first two. Overall, independent of the actual similarity of sensory input, the similarity of neural responses strongly covaried with the similarity of the percept that the observer became aware of. Most surprisingly, these effects were already visible in the very earliest stimulus-related EEG activity that could be recorded from the scalp, but were not localized to primary visual cortex. Instead, the effect of motion on perceived position seems part of a rapid, predominantly feed-forward mechanism in motion area MT, which responds very rapidly to visual input despite classically being seen as a mid-level visual area. One interpretation is that the illusion therefore reflects anticipatory coding of an object's future position. Indeed, in animal physiology experiments such anticipatory coding has been reported as early as LGN (in cats; Sillito, Jones, Gerstein, & West, 1994) and even in the retina itself (in salamanders and rabbits; Berry, Brivanlou, Jordan, & Meister, 1999; Schwartz, Taylor, Fisher, Harris, & Berry, 2007). This means that, in awareness, moving objects are actually perceived ahead of where they were when their visual image landed on the retina. However, owing to the processing delays incurred in *reaching* awareness, where we perceive the object to be *at the moment that we perceive it* would be made more accurate by such a mechanism.

Finally, Vandenbroucke, Fahrenfort, Sligte, and Lamme (2014) showed that the classification approach can be used to study the neural processes involved in the construction of an illusory percept even when that percept is rendered invisible. On the basis of patterns of fMRI activation, they trained a classifier to discriminate between trials in which an array of shapes was arranged to either induce an illusory Kanizsa surface (Kanizsa, 1976) or not. Importantly, in one set of conditions the inducers were not consciously perceived due to an attentionally demanding concurrent task. Nevertheless, the performance of a classifier trained on the invisible condition generalized to a visible condition in which the illusory surface *could* be perceived. This indicates that some perceptual inference processes involved in the construction of the conscious percept take place even when the stimulus does not eventually achieve awareness.

---

As examples, I have highlighted five recent studies that have effectively applied this technique to different visual illusions using different neuroimaging techniques. However, the paradigm is by no means constrained to the visual modality, and would work equally well for other sensory modalities or multimodal paradigms. One of the only practical restrictions is that a sufficient number of trials needs to be available within individual observers. Sometimes known as the ‘curse of dimensionality’, the practical implication is roughly that, in order to train a classifier to discriminate two conditions, at least as many data points need to be available in each condition being compared as the acquisition space has dimensions, plus one. For example, to classify data points on the basis of 64-channel EEG data, at least 65 data points need to be available for each condition per observer (after artefact rejection). This also means that classification on the basis of fMRI data is usually limited to a subset of voxels, often specific functionally or anatomically defined regions of interest. Note that if the time course is not of specific interest—for example, in a block design—then all individual acquisitions across trials can be aggregated to train the classifier. For example, in our study of the MAE (Hogendoorn & Verstraten, 2013b) classifiers were trained on 27 blocks, each with 2 repetitions per condition and with 15 acquisitions per trial, for a total of 810 data points per condition. This was higher than the number of voxels in each of the visual areas we considered, which ranged from 85 to 368.

Even more so than is generally true for perception experiments, it is advisable to acquire as many trials as possible, because the outcome of a classification procedure (proportion-correct classification) is dependent on the number of available trials in two ways. Firstly, increasing the number of trials improves the signal-to-noise ratio in the training set, which means the classifier will use a better estimate of the true hyperplane separating the two stimulus classes, thereby increasing performance. Secondly, with more trials to classify, the granularity of the final proportion correct is lower, meaning that the variability that can be expected based only on noise is reduced. Because the classifier is trained and tested within observers, this means that this imposes a lower limit on the amount of time required per observer.

There are different ways to draw training and test sets from the complete dataset. Maximum power is achieved for classification purposes by removing a single trial from the complete dataset, training the classifier on all the remaining data (ie the  $N-1$  remaining trials), then testing the classifier on the trial that was removed, and finally repeating the entire procedure for each of the remaining individual trials (also known as  $N-1$  jackknife). Another approach is to split the dataset in two, using one half to train and the other to test (also known as split-half or cross-validation). Finally, it is possible to generate training sets using a reduced number of trials or acquisition sites. This is most useful when comparing two brain areas of different sizes (such that the number of available voxels differs) or to compare two different classifiers that differ in the number of available trials.

The output of a classification procedure is a proportion-correct classification performance for each observer. Statistically evaluating whether this proportion differs significantly from chance is possible using classical parametric statistics across observers (eg one-tailed  $t$ -tests against chance), but a more powerful approach is to use bootstrap methods (Efron & Tibshirani, 1994). There are two ways to statistically evaluate the relevance of a classification outcome. The first approach is to use bootstrapping to estimate confidence intervals of the outcome variable. To do this, numerous new samples are repeatedly drawn from the complete dataset (with replacement), and analyzed as per the analysis of the original data. Repeatedly drawing new training and test sets in this way yields a distribution of the outcome variable: classification performance. Using this procedure, it is therefore possible to estimate confidence intervals of classification performance for a particular comparison (for example, by taking the 2.5th and 97.5th percentiles of this distribution as the limits of the 95% confidence interval).

---

This approach is not only useful for illustration purposes, but can also be used to compare performance across different conditions within observers—for example, using independent-samples *t*-tests.

Importantly, it is also possible to estimate the distribution of classification performance under the null hypothesis—that is, that no systematic information relevant to the comparison being tested is actually available in the training data. An effective way to test the null hypothesis that classification performance is merely due to chance is to repeatedly create new samples from the complete original dataset by reshuffling trial condition labels in the training (or test) set. Removing systematic variation between the two categories used to train the classifier in this way yields a distribution of the outcome variable under the null hypothesis. As per conventional parametric statistics, the observed outcome variable in the original dataset can be compared against this dataset to establish a *p*-value [for example, if the observed value lies in the 99th percentile, then  $p < 0.01$  (one-tailed)]. Using bootstrapping approaches to generate distributions of the outcome variable in this way has the important advantage of being free from the assumptions that need to be satisfied in order to apply parametric statistics to the outcome variable directly.

Altogether, multivariate pattern classification techniques have proven to be a valuable tool in the toolbox of the cognitive neuroscience researcher. Here, I highlight five studies that applied pattern classification techniques to neuroimaging data acquired while observers experienced perceptual illusions. I believe this approach to perceptual illusions provides a powerful opportunity to directly probe the contents of perceptual awareness without concurrent confounding changes in bottom-up sensory input, giving a unique window into the neural processes underlying conscious awareness.

## References

- Anstis, S. M., Verstraten, F. A. J., & Mather, G. (1998). The motion aftereffect: A review. *Trends in Cognitive Science*, **2**, 111–117.
- Berry, M. J., II., Brivanlou, I. H., Jordan, T. A., & Meister, M. (1999). Anticipation of moving stimuli by the retina. *Nature*, **398**, 334–338.
- Carlson, T. A., Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, **13**(10):1, 1–19.
- Cavanagh, P., & Anstis, S. (2013). The flash grab effect. *Vision Research*, **91**, 8–20.
- Crowe, D. A., Leuthold, A. C., & Georgopoulos, A. P. (2010). Differential brain activity states during the perception and nonperception of illusory motion as revealed by magnetoencephalography. *Proceedings of the National Academy of Sciences of the USA*, **107**, 22677–22681.
- Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. Boca Raton, FL: CRC Press.
- Fukiage, T., Whitney, D., & Murakami, I. (2011). A flash-drag effect in random motion reveals involvement of preattentive motion processing. *Journal of Vision*, **11**(13):12, 1–13.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, **293**, 2425–2430.
- Haynes, J. D. (2009). Decoding visual awareness from human brain signals. *Trends in Cognitive Sciences*, **13**, 194–202.
- Hogendoorn, H., Carlson, T. A., & Verstraten, F. A. J. (2011). Mapping the route to visual awareness. *Journal of Vision*, **3**(13):4, 1–10.
- Hogendoorn, H., & Verstraten, F. A. J. (2013a). Decoding the motion aftereffect in human visual cortex. *NeuroImage*, **82**, 426–432.
- Hogendoorn, H., & Verstraten, F. A. J. (2013b). Decoding sensation and perception over time with EEG pattern cross-classification. *Perception*, **42** (ECPV Supplement), 97. (Abstract)
- Kanizsa, G. (1976). Subjective contours. *Scientific American*, **234**, 48–52.
- King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Science*, **18**, 203–210.

- 
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Science*, **17**, 401–412.
- Kutas, M., McCarthy, G., & Donchin, E. (1977). Augmenting mental chronometry: The P300 as a measure of stimulus evaluation time. *Science*, **197**(4305), 792–795.
- Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, **7**, 12–18.
- Maus, G. W., Fischer, J., & Whitney, D. (2013). Motion-dependent representation of space in area MT+. *Neuron*, **78**, 554–562.
- Necker, L. A. (1832). Observations on some remarkable optical phaenomena seen in Switzerland; and on an optical phaenomenon which occurs on viewing a figure of a crystal or geometrical solid. *London and Edinburgh Philosophical Magazine and Journal of Science*, **1**, 329–337.
- Polonsky A., Blake R., Braun J., & Heeger D. J. (2000). Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nature Neuroscience*, **3**, 1153–1159.
- Rees, G. (2007). Neural correlates of the contents of visual awareness in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **362**, 877–886.
- Schwartz, G., Taylor, S., Fisher, C., Harris, R., & Berry, M. J., II. (2007). Synchronized firing among retinal ganglion cells signals motion reversal. *Neuron*, **55**, 958–969.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. E. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, **369**, 479–482.
- Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, **21**, 753–759.
- Tononi, G., & Koch, C. (2008). The neural correlates of consciousness: An update. *Annals of the New York Academy of Sciences*, **1124**, 239–261.
- Vandenbroucke, A. R. E., Fahrenfort, J. J., Sligte, I. G., & Lamme, V. A. F. (2014). Seeing without knowing: Neural signatures of perceptual inference in the absence of report. *Journal of Cognitive Neuroscience*, **26**, 955–969.
- VanRullen, R., Carlson, T. A., & Cavanagh, P. (2007). The blinking spotlight of attention. *Proceedings of the National Academy of Sciences of the USA*, **104**, 19204–19209.
- Van Wezel, R. J. A., & Britten, K. H. (2002). Motion adaptation in area MT. *Journal of Neurophysiology*, **88**, 3469–3476.
- Zeki, S., & Bartels, A. (1998). The autonomy of the visual systems and the modularity of conscious vision. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **353**, 1911–1914.