Perceptual integration without conscious access

J. J. Fahrenfort, J. van Leeuwen, C.N.L. Olivers, H. Hogendoorn

Abstract

The visual system has the remarkable ability to integrate fragmentary and degraded visual input into a perceptually organized collection of surfaces, shapes and objects. Here we test the long-standing maxim that such integrative functions must be attributed to attentional selection and conscious access. Employing a full-factorial design of masking and the attentional blink, we show that behaviorally, both masking and the withdrawal of attention affect perceptual decisions about the presence of integrated surface structure from fragmented input. However, when using a multivariate classifier on electroencephalogram data, we are able to decode the presence of integrated percepts equally well regardless of the availability of attention. In contrast, masking destroys any evidence for integrated percepts, while leaving feedforward processing intact. Thus, there is a fundamental difference in the way masking and attention impact perceptual integration, despite having a similar impact on behavioral decisions. Whereas masking disrupts perceptual integration, the brain is able to organize fragmented visual input into perceptually meaningful wholes in the absence of conscious access.

Keywords:
Perceptual integration, consciousness, access, masking, attentional blink, perceptual organization, neuronal integration, perceptual inference, feature integration, top-down attention

Dating back to Helmholtz, conscious perception is thought to result from the unconscious integration of spatially scattered features, allowing the brain to make perceptual inferences about visual input[1]. Historically, attentional selection is thought to be required for this process of integration[2]. In this view, attention is crucial to promote unconscious representations to conscious ones, a position that echoes through in current theorizing about consciousness[3,4]. However, the link between perception and conscious access has been called into question in recent years, suggesting that perceptual structures may be formed despite not being consciously detected[5-7]. In this counterview, conscious access does not play a causal role in perception itself, so that integrated perceptual representations may exist without it.

The current study employs the Kanizsa illusion (see Fig. 1a), together with two well-known manipulations of consciousness, to assess whether neural representations can reach a state of integration in which features are combined to form perceptual entities, despite not being consciously detected. Kanizsa figures are similar to control figures in terms of physical input, but they have different perceptual outcomes, notably an illusory surface region with accompanying contours[8] and increased brightness[9]. These emergent properties are a primary demonstration of perceptual integration, as the constituent parts in isolation (the inducers) do not carry any of the effects that are brought about by their configuration. Earlier work has shown that Kanizsa configurations can facilitate detection of target stimuli, with and without competing objects[10-14]. However, in these studies conscious access has been implemented in various ways, while the dependent measure was always a behavioral response. The only study that has measured the neural substrate of perceptual integration in the absence of conscious report, postponed the behavioral response until after data collection[7], leaving open the possibility that subjects were consciously accessing the stimulus during scanning but had forgotten it at test time[15]. The level at which conscious access and perceptual integration interact thus remains unclear. The current study employed several electroencephalographic (EEG) measures to investigate the underlying neural substrate of perceptual integration under two very different types of manipulations known to affect consciousness: one in which the pattern was either masked or not masked and one in which attention was either available or not available (as induced by the attentional blink; AB) (see Fig. 1b for the factorial design). We expected masking to affect behavioral and neural measures of perceptual integration[14,16,17]. The crucial question was whether the same result would occur when conscious access was prevented by withdrawing attention, as would be predicted if access reflects or causes perceptual integration.

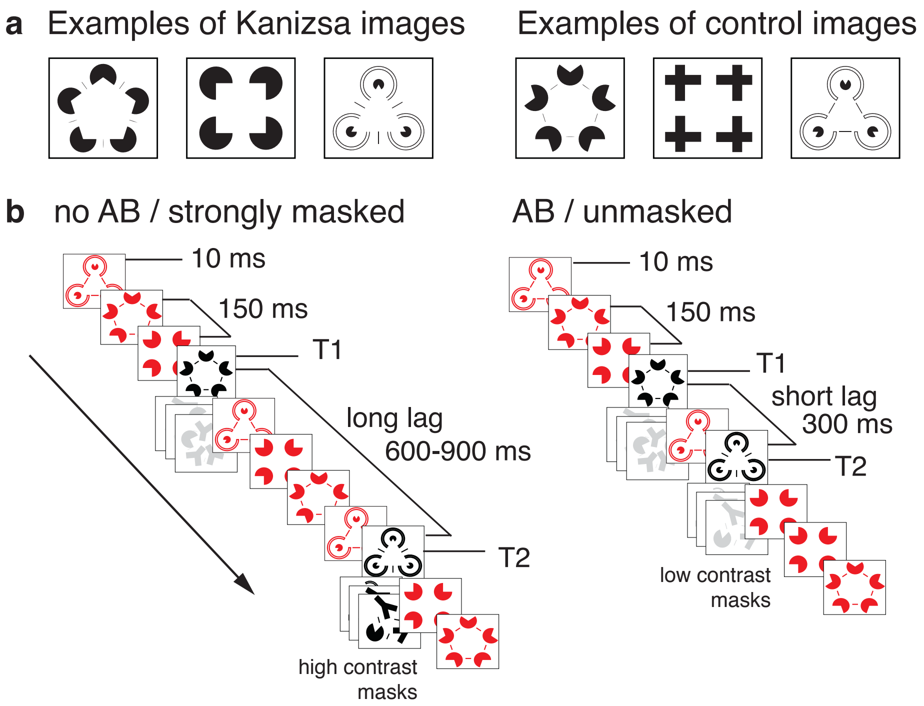

Figure 1. Experimental design. (a) Examples of different Kanizsa images and their controls as used in the experiment, see Fig. S1 for the complete stimulus set. (b) Examples of two of the four trial types in the factorial design: without an attentional blink (long lag) and strong masking (left) and with an attentional blink (short lag) and no masking (right).

Results

T1 classification accuracy reflects perceptual integration

We recorded 64-channel EEG data from human subjects in two EEG sessions. Two black target figures (T1 and T2) were shown in a rapid serial visual presentation (RSVP) containing red distractors. Each target could either be a Kanizsa or a control figure (Fig. 1a and Fig. S1). T1 and T2 lag was varied, inducing an attentional blink at short lags (300 ms) with recovery at long lags (≥600 ms). In half of the trials, T2 was strongly masked using high contrast masks. In the other half, low contrast masks were used, so that there was no effect of masking (see examples of masks in Fig. S2). Examples of two of the four trial types are shown in Fig. 1b. At the end of each trial, subjects indicated whether T1 and/or T2 contained a surface region (see online methods for details). Accuracies were computed as hit rate (HR) minus false alarm rate (FAR), reflecting how well subjects were able to distinguish surface from control figures, thus serving as a behavioral index of perceptual integration. T1 accuracy was high, at .90 (s.e.m. .02).

To establish a neural index for perceptual integration, we trained a linear discriminant to classify trials as either Kanizsa or control using EEG data. The training set was obtained from an *independent* RSVP task in which subjects pressed a button when a black figure would repeat (1-back task), while ignoring red distractors (see Fig. S3). To ensure that response-related processes would not contaminate the neural index for perceptual integration, response type (repeat or not) was independent from stimulus class (Kanizsa or control), and all response trials were excluded from the training set. Next, we used the resulting classifier on the experimental runs, computing classification accuracy (HR-FAR, as in the behavioral measure) for every time sample. Single trial classification accuracy (Kanizsa versus control) for T1 was well above chance, peaking at ~264 ms (Fig. 2a), and was strongly occipital in nature (see correlation/class-separability map[18] in Fig. 2b, online methods for details). To provide a direct link between peak classification accuracy and perceptual integration, we computed behavioral and classifier accuracy for 12 Kanizsa-control pairs (see Fig. S1 and online methods for characterization of the full stimulus set). A robust linear regression analysis[19] showed that T1 peak classification performance predicted T1 behavior with high accuracy (R^2=.61, p<.005, see Fig. 2c), confirming that peak classification accuracy reliably reflects perceptual integration.

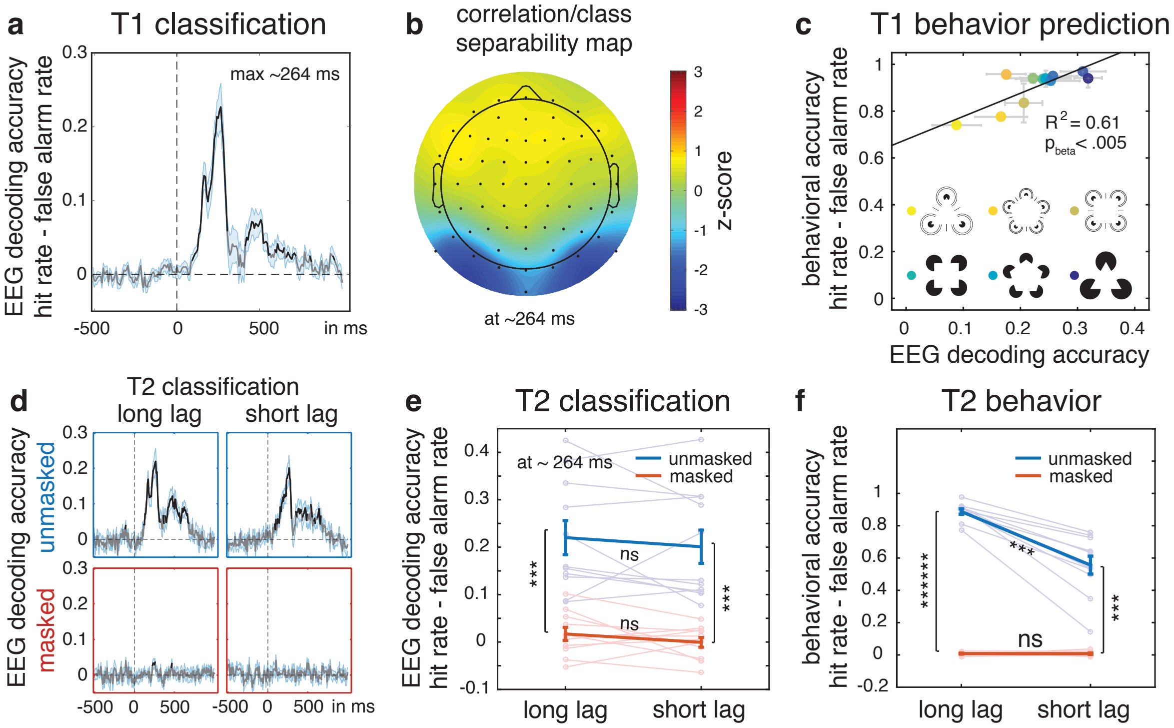

Fig 2. Peak classification accuracy reflects perceptual integration. (a) T1 EEG mean decoding accuracy of perceptual integration over time, black line reflects p<.05, +/- SEM in light blue. (b) and the correlation/class separability map reflecting the underlying neural sources for maximum decoding at ~264 ms, see methods. (c) The degree to which classification accuracy at ~264 ms predicts behavioral sensitivity to perceptual integration at T1. The legend shows six representative Kanizsa figures that correspond to the colored data points. Color follows the order of classification accuracy in T1, for the full legend see figure S1. (d) T2 EEG decoding accuracy over time for the four experimental conditions and (e) maximum decoding accuracy at ~264 ms for these conditions. (f) Behavioral sensitivity to perceptual integration for the four conditions (compare to e). Error bars are mean +/- SEM, individual data points are plotted in light in the background.

The attentional blink and masking differentially impact behavioral and neural measures of perceptual integration

We then established how this marker of perceptual integration is impacted by attention and masking. In terms of behavior, we observed the classic deleterious effects of both masking (mask vs. no mask, $F_{1,10}$=426.54, $P<10^{-8}$) and the attentional blink (short vs. long lag, $F_{1,10}$=51.89, $P<10^{-4}$) on accuracy (Fig. 2f). There was also an interaction ($F_{1,10}$=52.17, $P<10^{-4}$), which was entirely driven by the difference between unmasked long- and short-lag trials (post-hoc t-test, $P<10^{-4}$).

We hypothesized that if both masking and the attentional blink impact perceptual integration, they should affect neural markers of perceptual integration in similar ways. To enable a direct comparison with behavior, we extracted classification accuracy in the four experimental T2 conditions. Fig. 2d shows the entire time course, Fig. 2e shows peak classification accuracy at 264 ms (latency taken from T1). A 2×2 analysis of variance (ANOVA) showed a highly significant main effect of masking ($F_{1,10}$=37.68, P<.001), but no main AB effect (short vs long lag) ($F_{1,10}$=2.16, P=.172), and no significant interaction between masking and AB ($F_{1,10}$=0.02, P=.963). Post-hoc t-tests confirmed significant costs for masked versus unmasked stimuli for both long and short lag (both P<.001), but no significant differences between long lag and short lag (both P>.25).

Thus, while we observe a strong effect of masking in both brain and behavior, the classic effect of diminished performance on short lag (AB) trials only occurs in behavior. To further statistically underpin the differential effect of attentional selection on behavioral and neural measures of perceptual integration, we entered both measurements into a large 2×2×2 ANOVA with factors measure (normalized behavioral / normalized neural), AB (yes/no) and masking (yes/no). The validity of treating neural and behavioral HR-FAR data as repeated measures of the same thing (i.e. classification of a perceptual object) is discussed in the online methods section. In line with the other results, this analysis showed a three-way interaction effect driven by differences in behavioral and neural classification accuracies ($F_{1,10}$=9.30, P=.012), as well as a two-way interaction between measure and AB ($F_{1,10}$=10.92, P=.008) but no interaction between measure and masking ($F_{1,10}$=1.51, P=.247).

Feedforward processing remains intact during masking

These data support the idea that masking disrupts perceptual integration whereas the attentional blink does not. However, a concern might be that masking wiped out all processing of the stimulus, rather than specifically affecting perceptual integration, resulting in a floor effect. To test this, we selected a subset of the stimulus set that could be divided orthogonally according to its impact on contrast detection or perceptual integration. Fig. 3a illustrates this: the horizontal axis captures differences in perceptual integration (surface perception on the right but not on the left), while the vertical axis captures difference in bottom up feature contrast (high contrast at the top versus low contrast at the bottom, see Fig. S5 and online methods for a specification of the entire stimulus set). If masking wipes out all stimulus processing, we should no longer be able to classify high versus low contrast stimuli. We computed classification accuracy for feature contrast on the one hand and perceptual integration on the other, using a within-condition eight-fold cross validation scheme (see online methods for details). The results are shown in Fig. 3b-c. In an early time window ~80-90ms, both masked and unmasked stimuli showed highly significant classification accuracies for feature contrast (left panes, masked: t(10)=7.45, $p<10^{-4}$; unmasked: t(10)=8.82, $p<10^{-5}$, statistics at ~92 ms, T1 peak latency). Thus, despite strong masking, the bottom-up signal is processed up to the point of contrast detection. Conversely, masking does wipe out classification accuracy on the perceptual integration dimension (right panes, masked: t(10)=-.19, p=.852; unmasked: t(10)=6.82, $p<10^{-4}$). Note that for all analyses, the same type of masks would follow all stimulus classes (regardless of whether these were Kanizsa, control, high- or low contrast), such that the masks themselves could not bias classification accuracy. These results show that masking selectively abolishes perceptual integration, leaving feedforward processing largely intact (corroborating previous work[16,17]).

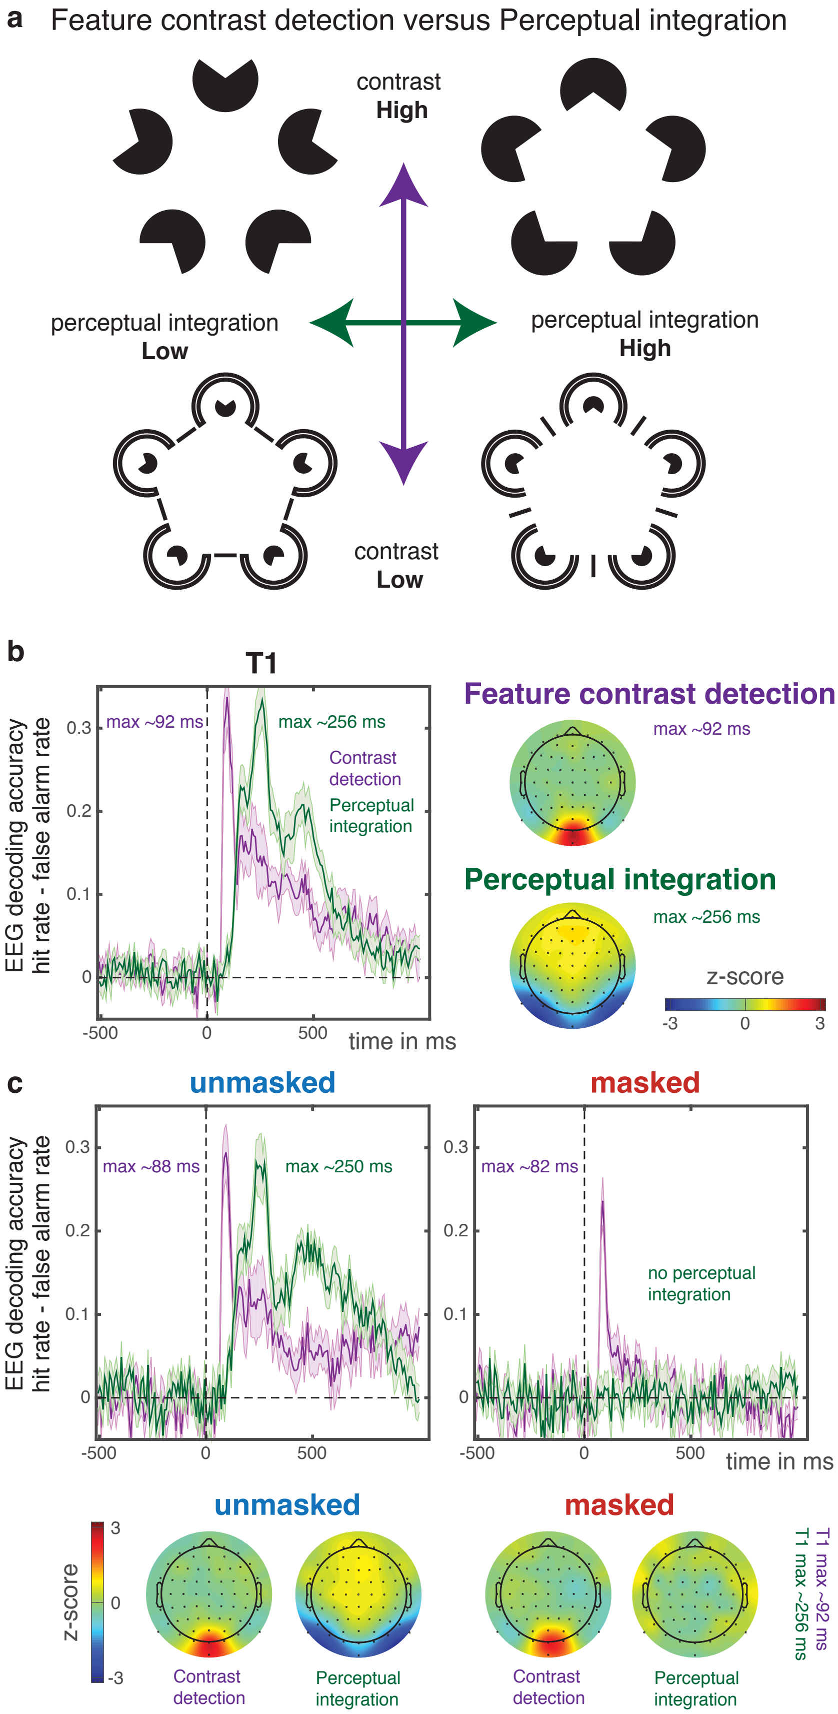

Fig 3. Separating out perceptual integration and feature contrast detection. (a) Example stimuli that were used to orthogonally classify feature contrast and perceptual integration on the same data. (b) Classification accuracies across time for contrast detection and perceptual integration (left) as well as correlation/-class separability maps (right) for T1, (c) and for unmasked (left) and strongly masked trials (right). Line graphs contain mean +/- SEM.

Masking selectively disrupts perceptual integration

Another concern might be that EEG classification accuracy is an all-or-none phenomenon whereas behaviour relies on graded evidence. In such a scenario, the behavioral effects on perceptual integration (Fig. 2f) might not be reflected in classification accuracy (Fig. 2e) due to a lack of sensitivity of the classifier to smaller effects such as those observed during the AB. To test this hypothesis, we conducted a control experiment in which we used a staircase to titrate mask contrast to get a weaker behavioral effect of masking, similar in magnitude to the effect of the attentional blink in Fig. 2f (see online methods for details). Fig. 4a shows the resulting behavioral effect of weak masking in this experiment. When computing classification accuracy on these data, we see that it nicely follows behavior (Fig. 4b-c), t(5)=3.82, *P*=.012. Together, these results show that the drops in behavioral accuracy caused by masking and the attentional blink have different root causes: masking impacts perceptual integration directly, whereas the attentional blink leaves it intact.

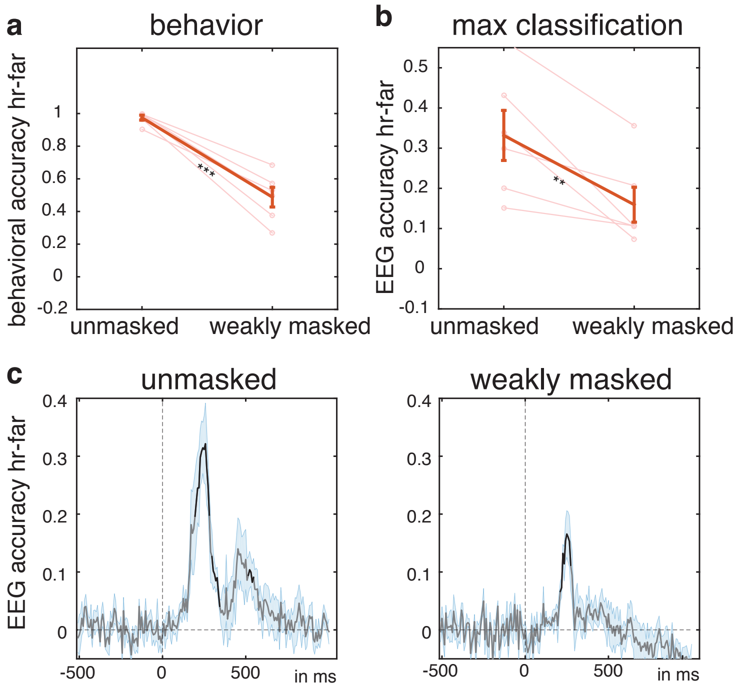

Fig 4. Masking control experiment. (a) behavioral results (b) max decoding accuracy. Error bars are mean +/- SEM, individual data points are plotted in light in the background. (c) Raw decoding accuracies over time for unmasked and weakly masked conditions. Line graphs contain mean +/- SEM , black line reflects p<.05.

Perceptual integration predates conscious access

So what neural process causes the dip in behavioral accuracy during the attentional blink? A natural hypothesis would be that the attentional blink interferes with conscious access after perceptual integration has already taken place. If true, we should be able to observe evidence of a selection process that results in conscious access at a later point in time. Investigating this issue requires a classifier that is sensitive to such a selection mechanism. Since the independent training runs that we used for training the classifier in the first analysis were designed to control for the direct influence of decision-related processes, these would not capture such a selection mechanism. The neural response to T1 however, does involve a conscious decision about the presence of a Kanizsa. We therefore trained a classifier on T1 data, and tested it on T2 data (see online methods for details). Fig. 5a and 5b show classification accuracies for the four experimental conditions when using this T1 classifier.

We again find the initial peak at 264 ms that was described before. Despite the potential contribution of decision mechanisms to classification accuracy when training on T1, this peak follows a pattern that is similar to the pattern that we observed when training on the independent training runs (cf. Fig. 2e and Fig 5b, top), and which is not in line with behavioral accuracy (Fig. 2f, see caption of Fig. 5b for statistical tests). So at what point in time is the behavioral effect of the attentional blink reflected in the neural data? The most notable difference when training on T1, is a second peak in classification accuracy occurring around 406 ms, and which is heavily modulated by AB (see top of Fig. 5a, online methods and Fig. S7). At this time point, the pattern of results is identical to that obtained in behavior (cf. Fig. 2f and Fig. 5b, bottom). All manipulations had highly significant effects on classification accuracy: a main effect of AB ($F_{1,10}$=7.96, P=.018), a main effect of masking ($F_{1,10}$=130.19, $P<10^{-6}$), as well as a strong interaction effect ($F_{1,10}$=14.92, P=.003).

To again directly compare behavioral to neural data at 406 ms, we once more entered the normalized measurements into a large 2×2×2 ANOVA with factors measure (behavioral/neural), AB (yes/no) and masking (yes/no). The results show highly significant main effects of AB ($F_{1,10}$=23.65, P<.001), masking ($F_{1,10}$=528.18, $P<10^{-9}$), as well as a strong interaction effect between AB and masking ($F_{1,10}$=51.55, $P<10^{-4}$), but importantly, now no two- or three-way interaction effects with measurement (neural/behavioral, all $F_{1,10}$<3.08, all P>.110), underpinning the similarity between behavioral and neural data pattern at this time point. The correlation/class separability map at 406 ms (Fig 5b, bottom) has the same topology and as that of a classical P300 (or P3b), which has often been associated with conscious access and perceptual decision-making[20-22]. Our data provide converging evidence that neural signals around the time frame of the P300 reflect a post-perceptual signal that is involved in conscious access, rather than perceptual integration itself. What we unambiguously show is that perceptual integration precedes such conscious access.

In a statistical sense we have so far regarded behavioral and classification accuracy data as repeated measures of the same underlying perceptual object. Another approach would be to view neural mechanisms as the cause of behavioral outcomes, by assessing the degree to which the neural data are able to serve as a model for behavior across time. To do this, we used normalized classification accuracies as reference points to determine the goodness of fit (GOF) with normalized behavioral accuracies as test data. As a measure of goodness of fit, we used the normalized root mean square error (NRMSE) cost function given by:

$$fit(t) = 1 - \frac{||xref(:,t) - x(:)||}{||xref(:,t) - mean(xref(:,t))||}$$

where x denotes the test data (behavioral accuracy), $xref$ denotes the neural data (classification accuracy), || indicates the 2-norm (Euclidean length) of a vector, fit is a row vector of length Nt and $t = 1,...,Nt$, where Nt is the number of time points. NRMSE costs vary between -Infinity (bad fit) to 1 (perfect fit). If the GOF cost function is equal to zero, then x is no better than a straight line at matching $xref$. We obtained this fitness measure separately for the different factors by collapsing the neural and behavioral data either across the masking factor, across the AB factor, or

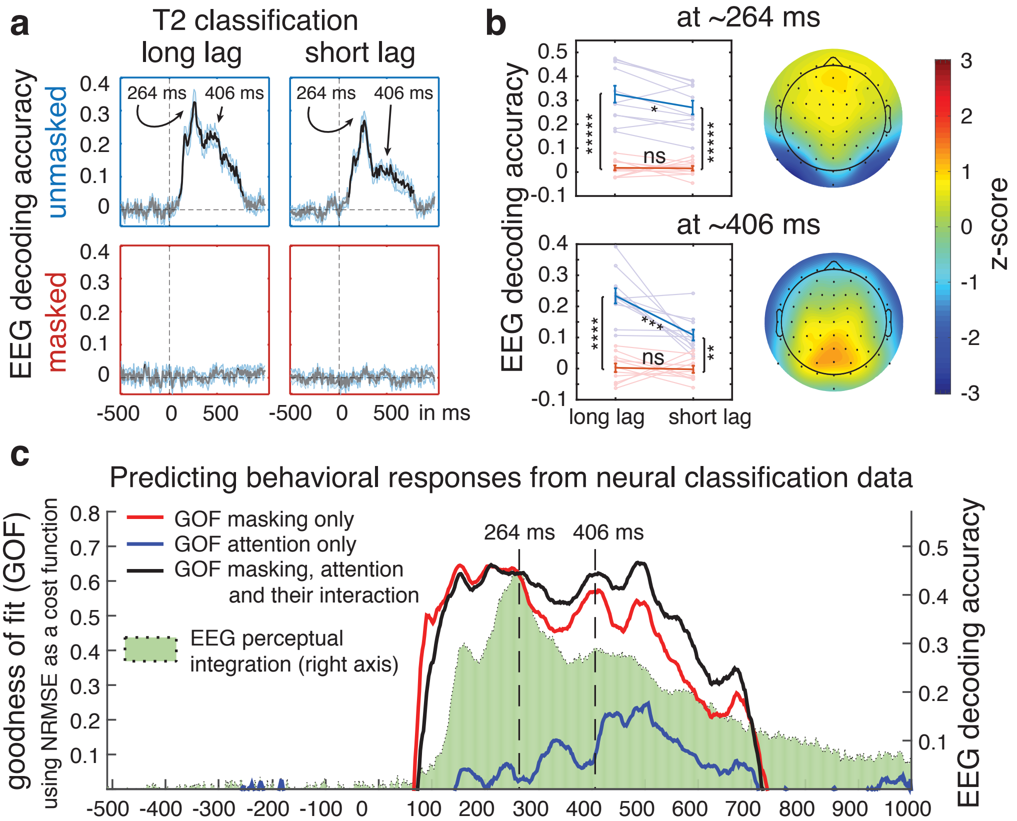

Fig 5. The impact of masking and attention on perceptual integration over time (a) EEG decoding accuracy for the four experimental T2 conditions when training on T1. (b) EEG decoding accuracies and correlation/class separability maps plotted at peak classification performance 264 ms (top) and at the second peak 406 ms (bottom). Blue lines represent the unmasked condition, red lines represent the masked condition. At the 264 ms time point, there was a strong main effect of masking ($F_{1,10}$=91.63, P<10-5), a main effect of AB ($F_{1,10}$=8.22, P=.017), and a trending interaction between masking and AB ($F_{1,10}$=4.06, P=.071). To test directly whether the measurement source (neural or behavioral) at 264 ms results in a differential effect on classification accuracy, we entered the normalized measurements into a large 2×2×2 ANOVA with factors measure (behavioral/neural), AB (yes/no) and masking (yes/no), see online methods. There was no interaction between measure and masking ($F_{1,10}$=.274, P=.61), but there was an interaction between measure and AB ($F_{1,10}$=6.75, P=.027), as well as a trending three-way interaction ($F_{1,10}$=4.50, P=.060), confirming that even when decision mechanisms are allowed contribute to classifier performance, the neural data at 264 ms cannot explain the pattern of results that is observed in behavior. The 406 ms time point on the other hand follows the same pattern as behavioral accuracy (see main text for statistics) and has a spatial distribution that is homologous to that of a classical P300. (c) An estimation of the goodness of fit when using the normalized EEG classification accuracy data as a model for the normalized behavioral detection data (left axis). Datasets are either collapsed over the AB dimension (GOF masking), over the masking dimension (GOF attention) or without collapsing over either dimension (GOF masking, attention and their interaction). T1 decoding accuracy is plotted as a green shade in the background for reference (right axis). Not until after the perceptual integration signal has peaked at 264 ms does the black line overtake the red line, showing a postperceptual contribution of attention to behavioral accuracy.

without regard to either factor (see online methods for details). The results are shown in Fig. 5c, where we also plot T1 classification accuracy as a reference for the time course of perceptual integration. Fig. 5c confirms that up to 264 ms, the masking manipulation uniquely models (predicts) behavior, indeed even better than when AB is also allowed to contribute to the fit. Only after 264 ms does attention start to contribute to behavioral outcomes, trailing the perceptual integration signal itself and in line with prior analyses.

Discussion

We show that EEG can be used to decode the presence of integrated percepts in visual cortex. Furthermore, we show that masking obliterates behavioral accuracy and classifier performance while retaining the ability to decode feature contrast, showing that it uniquely disrupts perceptual integration while leaving feedforward signals intact[16,17]. Interestingly however, peak classification performance on integration remains unchanged during the attentional blink, despite causing a marked dip in behavioral accuracy. This shows that the brain is able to integrate features into perceptual objects in the absence of conscious access.

This conclusion is seemingly at odds with experiments on object-based attention. For example, in an experiment by Roelfsema and colleagues[23], monkeys were trained to perform a curve-tracing task in a display with overlapping curves. Attention to the task-relevant curve resulted in a spreading activation across V1 neurons that coded the features belonging to the curve, thus binding the constituent elements of the curve together. This suggests that attention is the glue that unites an object, in line with the classical framework put forward by Treisman[2], and inconsistent with the position that attentional selection is not required for perceptual integration. Other studies have shown that such spreading activation follows Gestalt rules[24], and encapsulates task irrelevant features as long as they are part of a task relevant object[25,26]. However, with few exceptions[e.g.27], task relevance and selection are intertwined in experiments on object-based attention. The relationship between object-based attention and perceptual integration might therefore be caused by task relevance, rather than by conscious access per se.

Here we show that Kanizsa figures can be integrated in visual cortex despite not being promoted to a consciously accessible state through attentional selection. In contrast, masking destroys perceptual integration regardless of task demands. Naturally, this difference must be reflected in neural mechanisms. Dynamic feature grouping that underlies perceptual integration is thought to rely on cortico-cortical feedback[16,28-35]. While much remains to be learned about the origin of these feedback signals, evidence suggests that they originate from within visual cortex[29,30,33-38]. Although attentional selection also involves feedback, this feedback originates from frontoparietal cortex[20,39-44]. In the consciousness literature, such long-range integration is often referred to as ‘global ignition’[3,45]. The current data suggest that global ignition is not required to effectuate perceptual integration within visual cortex.

Our results also speak to a current debate about whether consciousness overflows cognitive access mechanisms[46,47]. In this debate, the question is whether access causes representational content to be extracted, or whether attention acts to select from a rich representational set that cannot be accessed in its entirety, reminiscent of the debate on early versus late selection[48]. In support of the latter position, a number of retro-cueing studies show that the representational capacity in early visual cortex is much larger than what can be accessed at any given moment, and that the extraction of this rich set from retinal input does not require conscious access[49,50]. A recent study has questioned such results, suggesting that a retro-cue might serve to postdictively impact perception after the display has already disappeared, dismissing the idea that retro-cue experiments are able to convincingly show that perceptual representations can exist without access[51]. The current experiment resolves this issue by providing a direct neural measure of perceptual integration, showing that neural representations in visual cortex can reach a perceptually integrated state in the absence of conscious access.

We would like to thank Anouk van Loon for very helpful comments during the materialization of this project.

Author contributions

J.J.F.and J.v.L. designed the experiment, J.v.L. collected the data, J.J.F., J.v.L. and H.H. analyzed the data, J.J.F., J.v.L. and C.N.L.O. wrote the paper.

References

1. Helmholtz, Von, H. *Handbuch der physiologischen Optik - Hermann von Helmholtz - Google Books*. (1867).
2. Treisman, A. M. & Gelade, G. A feature-integration theory of attention. *Cognit Psychol* **12,** 97–136 (1980).
3. Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J. & Sergent, C. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn Sci* **10,** 204–211 (2006).
4. Baars, B. J. *A Cognitive Theory of Consciousness*. (Cambridge University Press, 1988).
5. Block, N. Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behav Brain Sci* **30,** 481–99; discussion 499–548 (2007).
6. Lamme, V. A. F. How neuroscience will change our view on consciousness. *Cogn Neurosci* **1,** 204–220 (2010).
7. Vandenbroucke, A. R. E., Fahrenfort, J. J., Sligte, I. G. & Lamme, V. A. F. Seeing without knowing: neural signatures of perceptual inference in the absence of report. *J Cognitive Neurosci* **26,** 955–969 (2014).
8. Heydt, von der, R., Peterhans, E. & Baumgartner, G. Illusory contours and cortical neuron responses. *Science* **224,** 1260–1262 (1984).
9. Laeng, B. & Endestad, T. Bright illusions reduce the eye's pupil. *P Natl Acad Sci USA* **109,** 2162–2167 (2012).
10. Davis, G. & Driver, J. Parallel detection of Kanizsa subjective figures in the human visual system. *Nature* **371,** 791–793 (1994).
11. Vuilleumier, P. & Landis, T. Illusory contours and spatial neglect. *Neuroreport* **9,** 2481–2484 (1998).
12. Mattingley, J. B., Davis, G. & Driver, J. Preattentive filling-in of visual surfaces in parietal extinction. *Science* **275,** 671–674 (1997).
13. Conci, M. *et al.* Preattentive surface and contour grouping in Kanizsa figures: evidence from parietal extinction. *Neuropsychologia* **47,** 726–732 (2009).
14. Harris, J. J., Schwarzkopf, D. S., Song, C., Bahrami, B. & Rees, G. Contextual illusions reveal the limit of unconscious visual processing. *Psychol Sci* **22,** 399–405 (2011).

15. Wolfe, J. M. *Inattentional amnesia*. (Fleeting memories, 1999).
16. Fahrenfort, J. J., Scholte, H. S. & Lamme, V. A. F. Masking disrupts reentrant processing in human visual cortex. *J Cognitive Neurosci* **19,** 1488–1497 (2007).
17. Lamme, V. A. F., Zipser, K. & Spekreijse, H. Masking interrupts figure-ground signals in V1. *J Cognitive Neurosci* **14,** 1044–1053 (2002).
18. Haufe, S. *et al.* On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage* **87,** 96–110 (2014).
19. Rousseeuw, P. J. & Leroy, A. M. *Robust Regression and Outlier Detection*. (John Wiley & Sons, 2005).
20. Sergent, C., Baillet, S. & Dehaene, S. Timing of the brain events underlying access to consciousness during the attentional blink. *Nat Neurosci* **8,** 1391–1400 (2005).
21. O'Connell, R. G., Dockree, P. M. & Kelly, S. P. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci* **15,** 1729–1735 (2012).
22. Nieuwenhuis, S., Aston-Jones, G. & Cohen, J. D. Decision making, the p3, and the locus coeruleus-norepinephrine system. *Psychol Bull* **131,** 510–532 (2005).
23. Roelfsema, P. R., Lamme, V. A. F. & Spekreijse, H. Object-based attention in the primary visual cortex of the macaque monkey. *Nature* **395,** 376–381 (1998).
24. Wannig, A., Stanisor, L. & Roelfsema, P. R. Automatic spread of attentional response modulation along Gestalt criteria in primary visual cortex. *Nat Neurosci* **14,** 1243–1244 (2011).
25. O'Craven, K. M., Downing, P. E. & Kanwisher, N. fMRI evidence for objects as the units of attentional selection. *Nature* **401,** 584–587 (1999).
26. Schoenfeld, M. A. *et al.* Dynamics of feature binding during object-selective attention. *P Natl Acad Sci USA* **100,** 11806–11811 (2003).
27. Peelen, M. V., Fei-Fei, L. & Kastner, S. Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature* **460,** 94–U105 (2009).
28. Lamme, V. A. F. The neurophysiology of figure-ground segregation in primary visual cortex. *The Journal of Neuroscience* (1995).
29. Fahrenfort, J. J. *et al.* Neuronal integration in visual cortex elevates face category tuning to conscious face perception. *P Natl Acad Sci USA* **109,** 21504–21509 (2012).
30. Roelfsema, P. R. Cortical algorithms for perceptual grouping. *Annu Rev Neurosci* **29,** 203–227 (2006).
31. Kok, P. & de Lange, F. P. Shape perception simultaneously up- and downregulates neural activity in the primary visual cortex. *Curr. Biol.* **24,** 1531–1535 (2014).
32. Kok, P., Bains, L. J., van Mourik, T., Norris, D. G. & de Lange, F. P. Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Curr. Biol.* **26,** 371–376 (2016).
33. Wokke, M. E., Vandenbroucke, A. R. E., Scholte, H. S. & Lamme, V. A. F. Confuse Your Illusion: Feedback to Early Visual Cortex Contributes to Perceptual Completion. *Psychol Sci* (2012). doi:10.1177/0956797612449175
34. Murray, M. M. *et al.* The spatiotemporal dynamics of illusory contour

processing: combined high-density electrical mapping, source analysis, and functional magnetic resonance imaging. *J Neurosci* **22,** 5055–5073 (2002).

35. Halgren, E., Mendola, J., Chong, C. D. R. & Dale, A. M. Cortical activation to illusory shapes as measured with magnetoencephalography. *Neuroimage* **18,** 1001–1009 (2003).
36. Hupe, J. M. *et al.* Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* **394,** 784–787 (1998).
37. Lamme, V. A. F., Super, H. & Spekreijse, H. Feedforward, horizontal, and feedback processing in the visual cortex. *Curr Opin Neurobiol* **8,** 529–535 (1998).
38. Roelfsema, P. R., Lamme, V. A. F. & Spekreijse, H. The implementation of visual routines. *Vision Res* **40,** 1385–1411 (2000).
39. Corbetta, M. M. & Shulman, G. L. G. Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* **3,** 201–215 (2002).
40. Pooresmaeili, A., Poort, J. & Roelfsema, P. R. Simultaneous selection by object-based attention in visual and frontal cortex. *P Natl Acad Sci USA* (2014). doi:10.1073/pnas.1316181111
41. Gregoriou, G. G., Gotts, S. J., Zhou, H. & Desimone, R. High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* **324,** 1207–1210 (2009).
42. Gregoriou, G. G., Gotts, S. J., Zhou, H. & Desimone, R. Long-range neural coupling through synchronization with attention. *Prog. Brain Res.* **176,** 35–45 (2009).
43. Baldauf, D. & Desimone, R. Neural mechanisms of object-based attention. *Science* **344,** 424–427 (2014).
44. Gross, J. *et al.* Modulation of long-range neural synchrony reflects temporal limitations of visual attention in humans. *P Natl Acad Sci USA* **101,** 13050–13055 (2004).
45. Dehaene, S., Sergent, C. & Changeux, J. P. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *P Natl Acad Sci USA* **100,** 8520–8525 (2003).
46. Block, N. Perceptual consciousness overflows cognitive access. *Trends Cogn Sci* **15,** 567–575 (2011).
47. Phillips, I. No watershed for overflow: Recent work on the richness of consciousness. *Philosophical Psychology* 1–14 (2015). doi:10.1080/09515089.2015.1079604
48. Deutsch, J. A. & Deutsch, D. Attention - Some Theoretical Considerations. *Psychol Rev* **70,** 80–90 (1963).
49. Sligte, I. G., Vandenbroucke, A. R. E., Scholte, H. S. & Lamme, V. A. F. Detailed sensory memory, sloppy working memory. *Front. Psychology* **1,** 175 (2010).
50. Vandenbroucke, A. R. E., Sligte, I. G., Fahrenfort, J. J., Ambroziak, K. B. & Lamme, V. A. F. Non-Attended Representations are Perceptual Rather than Unconscious in Nature. *Plos One* **7,** e50042 (2012).
51. Sergent, C. *et al.* Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Curr. Biol.* **23,** 150–155 (2013).