

Seeing without Knowing: Neural Signatures of Perceptual Inference in the Absence of Report

Annelinde R. E. Vandenbroucke¹, Johannes J. Fahrenfort²,
Ilja G. Sligte^{1,3}, and Victor A. F. Lamme¹

Abstract

■ Every day, we experience a rich and complex visual world. Our brain constantly translates meaningless fragmented input into coherent objects and scenes. However, our attentional capabilities are limited, and we can only report the few items that we happen to attend to. So what happens to items that are not cognitively accessed? Do these remain fragmentary and meaningless? Or are they processed up to a level where perceptual inferences take place about image composition? To investigate this, we recorded brain activity using fMRI while participants viewed images containing a Kanizsa figure, an illusion in which an object is perceived by means of perceptual inference. Participants were presented with the Kanizsa figure and three matched non-illusory control figures while they were engaged in an attention-

ally demanding distractor task. After the task, one group of participants was unable to identify the Kanizsa figure in a forced-choice decision task; hence, they were “inattentionally blind.” A second group had no trouble identifying the Kanizsa figure. Interestingly, the neural signature that was unique to the processing of the Kanizsa figure was present in both groups. Moreover, within-subject multivoxel pattern analysis showed that the neural signature of unreported Kanizsa figures could be used to classify reported Kanizsa figures and that this cross-report classification worked better for the Kanizsa condition than for the control conditions. Together, these results suggest that stimuli that are not cognitively accessed are processed up to levels of perceptual interpretation. ■

INTRODUCTION

Perception does not directly emerge from the physical stimulation of photoreceptor cells in the retina. Rather, the brain continuously interprets incoming information to make sense of it: Through perceptual inference, visual input is translated from meaningless fragmented input into bound objects and scenes. For example, when we see a pen lying on top of a paper, we do not perceive the paper as having a pen-shaped hole in it. Instead, the paper is filled in underneath the pen, and we perceive the paper as an uninterrupted rectangle. In this study, we investigated whether this type of inference depends on the ability to attend to and cognitively access visual percepts. When a part of the visual field is neither attended nor reported, does vision represent its constituent parts as consisting of bound and completed objects? Or do they remain fragmentary and meaningless? The answer to this question has important implications for understanding the nature of vision and may ultimately change our view on conscious perception.

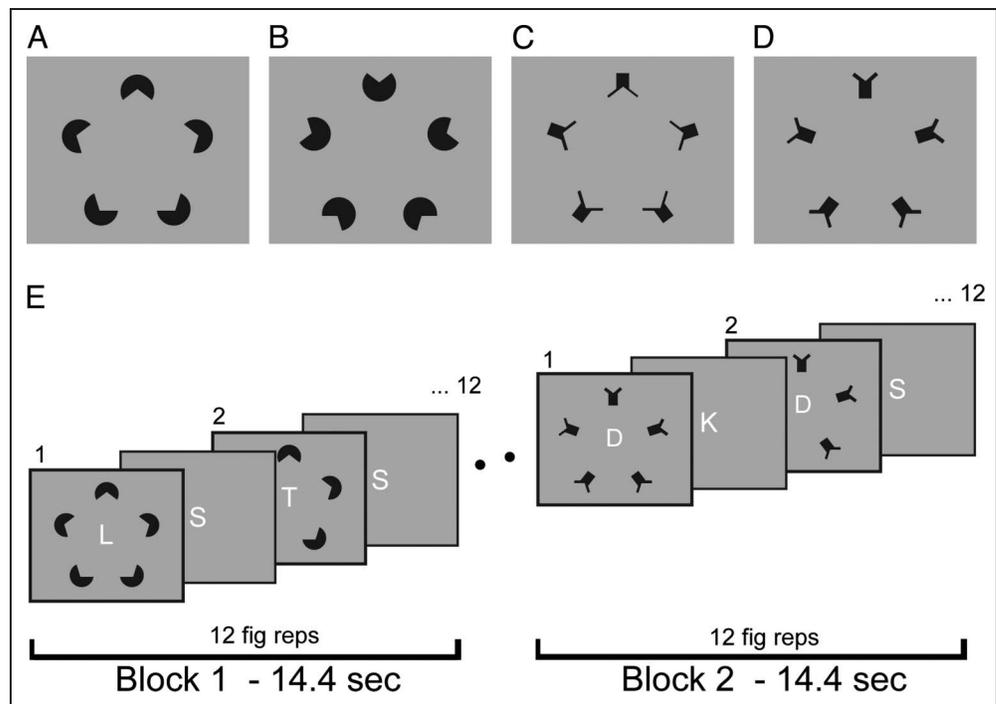
A prime example of perceptual inference is the Kanizsa illusion (Kanizsa, 1976), in which a set of inducers is aligned in such a way that observers perceive an occluding sur-

face lying on top of black disks (Figure 1A). This occluding object is defined by illusory contours and by the illusory contrast difference between surface and background. The illusory contours and illusory contrast difference do not emerge when the inducers are not properly aligned (Figure 1B and D) or when the inducers are not likely to be completed as occluded objects (Figure 1C). The formation of the illusion involves feedback from higher-level visual areas such as the lateral occipital complex (LOC) to lower visual areas V1/V2 (Knebel & Murray, 2012; Maertens, Pollmann, Hanke, Mildner, & Möller, 2008; Halgren, Mendola, Chong, & Dale, 2003; Lee & Nguyen, 2001; Seghier et al., 2000). Moreover, the perceptual nature of the Kanizsa figure has been shown to depend on activation in these regions in a reverse hierarchical manner (Wokke, Vandenbroucke, Scholte, & Lamme, 2013). These studies suggest that the inference mechanisms at play in the Kanizsa illusion depend on interactions between functionally divergent visual areas.

A remarkable observation is that when the inducers of the figure are rendered invisible by continuous flash suppression, the illusion is not perceived, even when attended, showing that perception of the Kanizsa illusion requires conscious processing of the inducers (Harris, Schwarzkopf, Song, Bahrami, & Rees, 2011). This is in contrast with the simultaneous brightness illusion—an illusion of a white disc seeming brighter when presented on a black background

¹University of Amsterdam, ²University of Utrecht, ³University of Birmingham

Figure 1. Stimuli and task design. The four stimuli used in the experiment were (A) A Kanizsa figure, (B) A Kanizsa control figure, (C) A line figure (same spatial layout and center of gravity as Kanizsa figure), and (D) A line control figure. (E) Participants performed a 2-back task in which they detected the repetition of a letter that was presented two serial positions before. Letters were presented for 600 msec each and every 1200 msec; one of the four figures was presented around the letters for 400 msec. Each figure was presented 12 times in a row, resulting in a blocked design with blocks lasting 14.4 sec.



than on a gray background—that persists even when the background is not perceived because of flash suppression. Perceptual inference in the Kanizsa illusion thus occurs at a higher level of processing and is susceptible to manipulations that selectively interfere with conscious perception. Studying to what extent the processes underlying the Kanizsa illusion require conscious access or reportability is therefore of direct relevance to the question whether access is necessary for the formation of a full perceptual representation.

To investigate whether the neural correlates of perceiving a Kanizsa figure are present when participants do not cognitively access the figure, we combined an inattentional blindness paradigm (Scholte, Witteveen, Spekreijse, & Lamme, 2006; Rees, Russell, Frith, & Driver, 1999; Rensink, O’Regan, & Clark, 1997) with fMRI measurements. Inattentional blindness occurs when a participant is attentionally engaged in another task, rendering a non-target stimulus unnoticed and unreported even when explicitly asked about it. When participants are informed about the presence of a certain stimulus during the task, however, they have no trouble seeing the stimulus, even when engaged in the distractor task. This suggests that their inability to report these stimuli occurs because they simply did not access them, not because of perceptual load (Lavie, 2005; Yi, Woodman, Widders, Marois, & Chun, 2004). The paradigm of inattentional blindness formalizes the common intuition that many stimuli in plain sight remain unnoticed and are therefore never accessed, although they are potentially accessible.

In an fMRI scanner, participants performed an attentionally demanding 2-back letter task, whereas the Kanizsa

and three control figures (Figure 1A–D) were presented surrounding the letters (Figure 1E). Participants were instructed that black “distractor” stimuli would be flashed around the letters, but that they should focus on the letter task to maximize their score on the 2-back task. After three runs of the task, participants were unexpectedly asked whether they had seen which figure was presented surrounding the letters. The participants that were unable to select the correct option were labeled as “inattentionally blind” (IB). If participants selected the Kanizsa figure, they were labeled as “not inattentionally blind” (NIB). We employed univariate and multivariate analysis techniques to compare the Kanizsa illusion to the three control figures. This allowed us to determine a neural signature that is unique to the illusion. The presence of a unique signature for the Kanizsa illusion for both IB and NIB participants would indicate that access is not required to process the illusion. If, however, only NIB participants show a unique signature of the illusion, this would suggest that processing the illusion requires access mechanisms.

METHODS

fMRI Acquisition

Forty-two students (three men) from the University of Amsterdam participated in the experiment for course credit or a monetary reward. The experiment was approved by the local ethics committee, and participants gave their written informed consent. All participants had normal or corrected-to-normal vision and were screened on the possibility of metal in their bodies and other risk

factors precluding participation in MRI studies. Scanning was performed on a 3T Philips Achieva MRI scanner at the Spinoza Center in Amsterdam. A high-resolution T1-weighted anatomical image (repetition time [TR] = 8.21 msec, echo time = 3.81 msec, field of view = $256 \times 256 \times 160$) was recorded for each participant. fMRI was recorded using a gradient-echo, echo-planar pulse sequence (TR = 2000 msec, echo time = 27.63 msec, flip angle = 90° , 27 slices with interleaved acquisition, voxel size = $2 \times 2 \times 3$ mm, 96×96 matrix, field of view = $89 \times 89 \times 192$) centered around the calcarine sulcus. Stimuli were back-projected on a 61×36 cm LCD screen using Presentation software (Neurobehavioral Systems, Inc., Albany, CA) and viewed through a mirror attached to the head coil.

Stimuli

To isolate the neural signature associated with the Kanizsa figure, the figure was compared with three control figures (Figure 1B–D). In addition to the classical control figure in which the inducers are rotated (Figure 1B; Mendola, Dale, Fischl, Liu, & Tootell, 1999), we used two additional figures (Figure 1C–D) to control for potential confounds. The Kanizsa illusion has both cognitive and perceptual attributes when compared with its traditional control figure. One can cognitively infer a pentagon from the layout of the Kanizsa figure by connecting the lines between the “pacmen” inducers; similar to, for example, knowing that a car has moved because the second time you see it, it is in a different location versus actually perceiving the movement of the car (Pylyshyn, 1999). Importantly, the illusion also has perceptual attributes: The surface of the pentagon seems to be “real,” as the surface of the pentagon is perceived as slightly brighter than the gray background and as lying on top of full disks instead of pacmen. Completion of the pacmen to full disks and the increased brightness are perceptually inferred (Nakayama & Shimojo, 1992; Kanizsa, 1974). Using the traditional control figure in which the inducers are rotated outward does not allow one to tease apart the perceptual aspects of the illusion (completion of the disks, illusory contours, and the brightness illusion) from its cognitive aspects (presence or absence of a pentagon). Therefore, we devised two additional control figures that isolate the cognitive attributes of the illusion (Figure 1C–D) using “hats” instead of pacmen inducers. In these control figures, a pentagon shape is either present or absent, yet only by cognitive inference and without the perceptual inferences that characterize the Kanizsa illusion.

The additional control figures also allowed us to circumvent other confounds in the traditional control condition. Particularly, it could be argued that the Kanizsa with rotated inducers is not a perfect control: Inwardly rotated inducers create an image with a different low spatial frequency content compared with an image with outwardly rotated inducers (there is a larger “white gap”

in the former; see Davis & Driver, 1998). If this is an important factor driving the neural signals we record, a difference found between the Kanizsa figure and its outward rotated counterpart should also be found between the line figure and its outward rotated control that contain a similar spatial frequency difference. Together, the three controls (Figure 1B–D) allowed us to assess the unique influence of the illusory nature of the Kanizsa (Figure 1A) on cortical processing when compared with more cognitive or low-level influences on stimulus processing.

The four figures (one Kanizsa, three control) consisted of five black inducers placed on a gray background (Figure 1A–D). The Kanizsa figure (Figure 1A, total span = $12.6^\circ \times 12.7^\circ$) and its control (Figure 1B, total span = $13.1^\circ \times 13.2^\circ$) consisted of pacmen-like inducers; five black circles (diameter = 2.6°) with a gap taken out (width = 2.1° , angle = 78°). The Line figure (Figure 1C, total span = $12.6^\circ \times 12.5^\circ$) and its control (Figure 1D, total span = $13.4^\circ \times 13.2^\circ$) consisted of hat-like inducers ($2.3^\circ \times 2.4^\circ$): line elements with a rectangle on top that had the same gap angle. In the Kanizsa figure, the inducers were aligned in such a way that a pentagon could be inferred lying on top of black discs. The support ratio (the length of the real contours relative to the illusory contour) of the Kanizsa figure is strongly related to the perception of the illusion. In this study, we used a support ratio of 0.42, which is sufficient to produce the illusion (Wokke et al., 2013; Seghier & Vuilleumier, 2006). In the classical Kanizsa control condition, the inducers were rotated 180° around their center of gravity—that was calculated by taking the pixel for which the amount of black pixels surrounding it was equal above and below, left and right. This was done to evenly divide the mass of the figures over the visual field and prevented the Kanizsa control figure from having more mass toward the center of the screen. Differences in the distribution of contrast discontinuities across the visual field can easily be picked up by simple spatial frequency filters (Becker & Knopp, 1978; Ginsburg, 1975), which might result in different responses to a Kanizsa and a classical control Kanizsa that has nothing to do with the perception of an illusion. The Line figure, which constituted the second control figure, resembled the Kanizsa figure in its spatial layout, and its inducers had the same center of gravity as the Kanizsa inducers. In this figure, a pentagon could be inferred; however, the illusion of contours and a contrast difference between surface and background was not present. The Line figure was compared with a third control figure, which contained the same elements as the Line figure, but now rotated 180° around their center of gravity, just as the Kanizsa control figure.

Behavioral Task

Participants performed a 2-back task on letters that were presented in a Rapid Serial Visual Presentation. They were instructed to press a button when the same letter was presented as two serial positions before. Letters

(0.5°) were presented in the center of the screen for 600 msec each. In every sequence of eight letters, a repetition occurred (jittered between Locations 3 and 8), and a total of 78 sequences was presented. At the end of each run, participants received feedback about the percentage of correctly detected targets.

Procedure

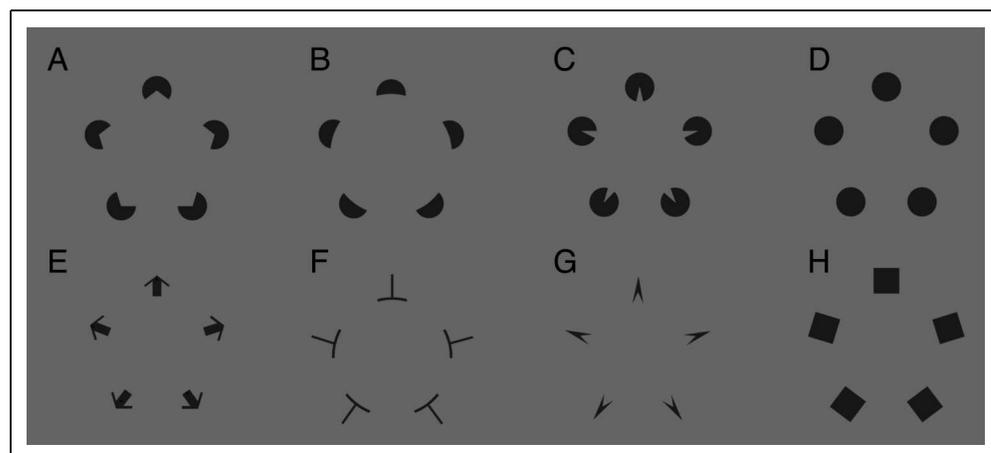
Participants were informed that they were participating in a study on the ability to perform a memory task while visually distracted. Before the start of the MRI session, they practiced the behavioral task for blocks of 2 min until they reached a performance of at least 80%. Then, they performed a block of around 6 min, the same length as in the MRI scanner. During this block, distracter stimuli were presented surrounding the letters. These distracter stimuli were similar to those used in the actual experiment, consisting of black stimuli (rectangles and half circles) at the same position and changing every 14.4 sec. However, they did not form a Kanizsa figure and were intended to get participants used to the flashing stimuli while focusing on the letter task.

Each functional run started with 10-sec fixation. Subsequently, two letter sequences (9.6 sec) were presented. After two sequences, the first surrounding figure was presented for four sequences (18.8 sec). It was expected that the first presentation of a figure would elicit a heightened activity because of its sudden onset regardless of figure type. We therefore presented a figure that was different from the four experimental figures at the start of each run (consisting of five half circles presented at the same position as the inducers of the four experimental figures). This way, none of the experimental figures had the advantage of being the first figure that was presented. After the presentation of the start figure, the four experimental figures were presented surrounding the letters in blocks of 14.4 sec (Figure 1E). In each block, one of the four figures was flashed around the letters with a duration of

400 msec and an ISI of 800 msec, making a total of 12 presentations per block. There was no rest between blocks, and the same figure was never repeated. The blocks were counterbalanced such that each figure was followed by one of the other figures for an equal amount of times. In each run, six blocks of each stimulus were presented, resulting in a total of 24 blocks per run. Each run ended with a 16-sec rest period, making the total runtime 400 sec (200 volumes).

After three runs in the scanner, participants were presented with a surprise question. In this question, participants were informed that the black stimuli surrounding the letters had formed figures and that they should choose the figure they thought had been presented during the three runs. After reading this question, they received eight options (Figure 2), of which only one contained the illusory Kanizsa figure that was presented during the experimental runs. They were asked to choose one of these options, even if they had to guess. We embedded two other Kanizsa-type figures (Figure 2B and C) to prevent participants from guessing the Kanizsa figure because it was the figure that popped out compared with the other figures. All participants that selected the correct figure were categorized as NIB, as they might have either explicit or implicit knowledge or familiarity with the figure, even if they felt they were guessing. All participants that selected an incorrect figure were labeled IB. After participants answered the surprise question, the correct answer was not given to them. Instead, they were asked to perform the exact same run again and, while performing the letter task, to try to detect which of the eight options was shown. The control run was identical to the experimental runs. After the control run, participants were asked again to identify the correct figure from the same eight options, after which the correct answer was revealed. Only participants that had answered the second question correctly were included in further analyses, ensuring their ability to perform the letter task and detect the Kanizsa figure at the same time.

Figure 2. The eight options presented in the multiple choice question after Experimental Runs 1–3 and the control run. For the NIB group, option A was chosen by all participants after Runs 1–3 and the control run. For the IB group, option A was never chosen after Runs 1–3. Instead, option B was chosen five times, option C was chosen two times, option F was chosen two times, option H was chosen three times, and options D, E, and G were never chosen. After the control run, all participants chose option A.



ROI Localization

For each participant, V1, V2, V3, V3AB, and V4 (Figure 3) were localized using a polar angle mapping, an eccentricity mapping, and a study specific localizer. For polar and eccentricity mapping, we used standard procedures such as described by Wandell and Winawer (2011). For polar angle mapping, a checkerboard (red–green, flickering at 8 Hz) wedge rotated around fixation (complete revolution in 30 sec, eight repetitions), and for eccentricity mapping, a checkerboard ring (red–green, flickering at 8 Hz) expanded from center to periphery (complete revolution in 30 sec, eight repetitions). During these two runs, participants fixated at the center while detecting blue squares presented in the red–green checkerboard stimuli. The TR of these two runs was set to 2500 msec as the phase of the wedge, and expanding ring was set at 2500 msec (six phases resulted in one cycle of 15 sec). In addition to the retinotopic mapping, a study-specific localizer was used in which a black circle (diameter = 2.6°; flickering at 2 Hz in 16-sec blocks) presented in the center of the screen was alternated with five black circles presented at the inducer positions (total figure 12.7° × 12.7°, flickering at 2 Hz in 16-sec blocks). These two conditions were separated by a 16-sec rest period and were repeated five

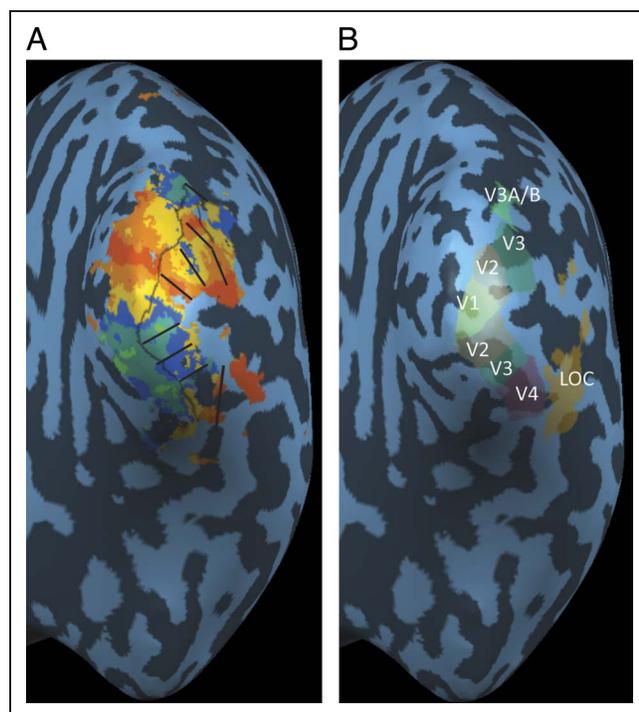


Figure 3. ROIs. Each individual's cortex was inflated and a retinotopic map (A, depicted here on a representative participant), eccentricity map, study-specific localizer, and a mapping for LOC (also shown here) were projected onto the inflated surface. On the basis of these mappings, six ROIs (V1, V2, V3, V3A/B, V4, LOC) were identified for each participant (B, depicted here on the representative participant in A). The gray line in A represents the border between the central visual region and the periphery where the inducers were presented.

times. Throughout the run, participants maintained fixation and performed a fixation task in which they had to detect a rotation of the fixation cross. Data of these runs were projected onto an inflated surface reconstruction, and ROIs were defined for each participant (see Figure 3). To define our ROIs, we followed the mapping procedure as described by Wandell and Winawer (2011). Because processing in the peripheral part of the visual regions might differ because of low-level feature differences in the inducers, we excluded these regions based on the stimulus-specific mapper. Activity for the inner circle was contrasted with activity for the outer circles, and a border was drawn from which only the inner part was taken as ROI (Figure 3A). Also, we did not include the foveal part of V1, V2, and V3, as foveal confluence makes it difficult to distinguish between these regions. V3A and V3B have a shared center and were in some participants hard to distinguish. We therefore took the central part of these regions together as one ROI.

In addition to these lower visual areas, LOC was localized using a mapper in which blocks of houses, faces, objects (chairs, scissors, bottles), and phase-scrambled versions of these objects were presented (Scholte, Jolij, Fahrenfort, & Lamme, 2008). Each block lasted 16 sec (eight presentations of 1000 msec per block) with a rest of 12 sec between each block. Each stimulus category was repeated four times. Activity for the contrast between objects and scrambled pictures was mapped on an inflated surface reconstruction, and LOC was defined for each participant. When contrasting objects against nonobjects, often find a more dorsal and a more ventral cluster are found (for an overview, see Wandell & Winawer, 2011). As in our participant pool, the ventral LOC cluster was most pronounced, we chose to take the ventral region as ROI (Figure 3). All localizer runs were performed during the same session as the experimental runs.

Univariate fMRI Analysis

Data were analyzed using Brainvoyager 2.1 (Brain Innovation, Maastricht, The Netherlands; Goebel, Esposito, & Formisano, 2006) and Matlab 2010 (MathWorks, Inc., Natick, MA). Functional scans were slice-time corrected, motion corrected, spatially smoothed with a Gaussian of 2 mm FWHM, and high-pass filtered using a general linear model (GLM) with Fourier basis set (three cycles). All functional scans were aligned to the first functional scan, which was coregistered to the T1-weighted anatomical image. Structural images were transformed to Talairach space using an ACPC transform (Talairach & Tournoux, 1988).

A GLM with five predictors (four experimental figures and a start figure) was defined for each participant. A whole-brain analysis (correcting for multiple comparisons using a false discovery rate [FDR] of 0.05) was performed for the two groups separately, combining the

three experimental runs (z -transformed) for each participant. For the ROI analyses, the GLM was modeled in each participant and ROI separately for the three experimental runs combined. To test the effect of Figure, ROI, and Group, a 4 (Figure: Kanizsa, Kanizsa Control, Lines, Lines Control) \times 6 (ROI: V1, V2, V3, V3AB, V4, LOC) \times 2 (Group: IB vs NIB) mixed repeated-measures ANOVA on the beta values was performed.

Multivariate fMRI Analysis

The same preprocessing steps as described for the univariate analyses were performed. For each block separately, the response for each voxel in each ROI was calculated. This was done by first z -transforming the whole time series and then averaging over six volumes following the first stimulus presentation in a block, with a 2 volume delay to account for the hemodynamic lag. The response for each block of the three experimental runs was fed into a training algorithm implemented in the Princeton MVPA Toolbox (code.google.com/p/princeton-mvpa-toolbox) using the backpropagation algorithm of the Netlab Neural Network Toolbox (www1.aston.ac.uk/eas/research/groups/ncrg/resources/netlab/). This yielded a specific voxel pattern for each figure in each ROI (all voxels in each ROI were used). First, the four figures were classified within experimental runs by training on two experimental runs (while participants were IB) and testing the remaining experimental run in a leave-one-out procedure. Then, the patterns of the blocks in the control run (while participants no longer suffered IB) were classified based on the three experimental runs together. For each iteration, the training runs were used to maximally distinguish the four neural patterns underlying each figure. Then, the labels that were given to figure presentations in the test run were classified as correct or not (1 or 0). This yielded a percentage correct for each figure presentation.

For both the experimental and the control run, we thus obtained a classification score (percentage correct) for each figure per participant per ROI. Chance performance was calculated for each individual and each training set by shuffling the test labels 1000 times and calculating the baseline performance on each figure for each specific training set. This allowed us to determine whether the training sets had a bias toward favoring one stimulus, thus creating an actual chance performance that differed from the expected 25% (based on four figures). On average, chance performance was 25% correct for each figure. To investigate whether classification of the Kanizsa figure was better than classification of the three control figures and whether there were any differences between the two groups, a 4 (Figure: Kanizsa, Kanizsa Control, Lines, Lines Control) \times 6 (ROI: V1, V2, V3, V3AB, V4, LOC) \times 2 (Group: IB vs. NIB) mixed repeated-measures ANOVA was performed.

RESULTS

Behavioral Performance

Of 42 participants, 5 participants did not select the correct answer after the control run, and these were excluded from all analyses. Twelve participants did not select the correct answer after the three experimental runs but did select the correct answer after the control run (IB group). From the participants that were correct on both questions, 12 participants (NIB group) were matched to the IB group on age (IB: 20.3 years, NIB: 20.2 years), sex (all women) and overall performance on the letter task (IB: 88%, NIB: 83%, $F(1, 22) = 1.345, p = .259$). There was a main effect of task, showing that n -back performance during the experimental runs was higher than performance during the control run (87% vs. 83%), $F(1, 22) = 6.703, p = .017$. This was probably because of the fact that participants were slightly engaged in detecting the correct figure during the control task. However, this slight performance decrement was the same for both groups, $F(1, 22) = .016, p = .900$.

Univariate fMRI Analysis

We examined visual areas that are known to be involved in generating the Kanizsa illusion (for an overview, see Seghier & Vuilleumier, 2006). We defined ROIs for V1, V2, V3, V3AB, V4, and LOC based on retinotopic and object-specific localizers (see Figure 3 and Methods). For each of these regions, a GLM was fitted to determine the activity corresponding to each condition (Figure 4). We compared the mean activity for each figure and the activity for the two groups by performing a 4 (Figure: Kanizsa, Kanizsa Control, Lines, Lines Control) \times 6 (ROI: V1, V2, V3, V3AB, V4, LOC) \times 2 (Group: NIB vs. IB) mixed repeated-measures ANOVA. A significant effect of Figure was found, $F(3, 66) = 13.8, p < .001$, showing that the Kanizsa illusion resulted in stronger activations of both lower- and higher-tier visual areas compared with the three control figures, post hoc t tests: all $p < .002$, Bonferroni-corrected. Moreover, the amount of activity for the Line figure and its control did not significantly differ, post hoc t test: $p > .999$, Bonferroni-corrected, suggesting that these two figures could not be dissociated based on activity in these visual areas. This shows that, although cognitive inference is possible for the Line figure (i.e., a pentagon can be inferred because logically, it is the only possible configuration), there is no specific visual neural signature that accompanies the figure. This confirms that there is no perceptual inference for the Line figure and the heightened activity associated with the Kanizsa illusion was because of the perceptual characteristics of the illusion and not because of other visual properties such as layout, spatial frequency, collinear contours, or the “cognitive binding” of the inducers into a single object.

The Kanizsa illusion elicited heightened activity across all of visual cortex, regardless of whether participants

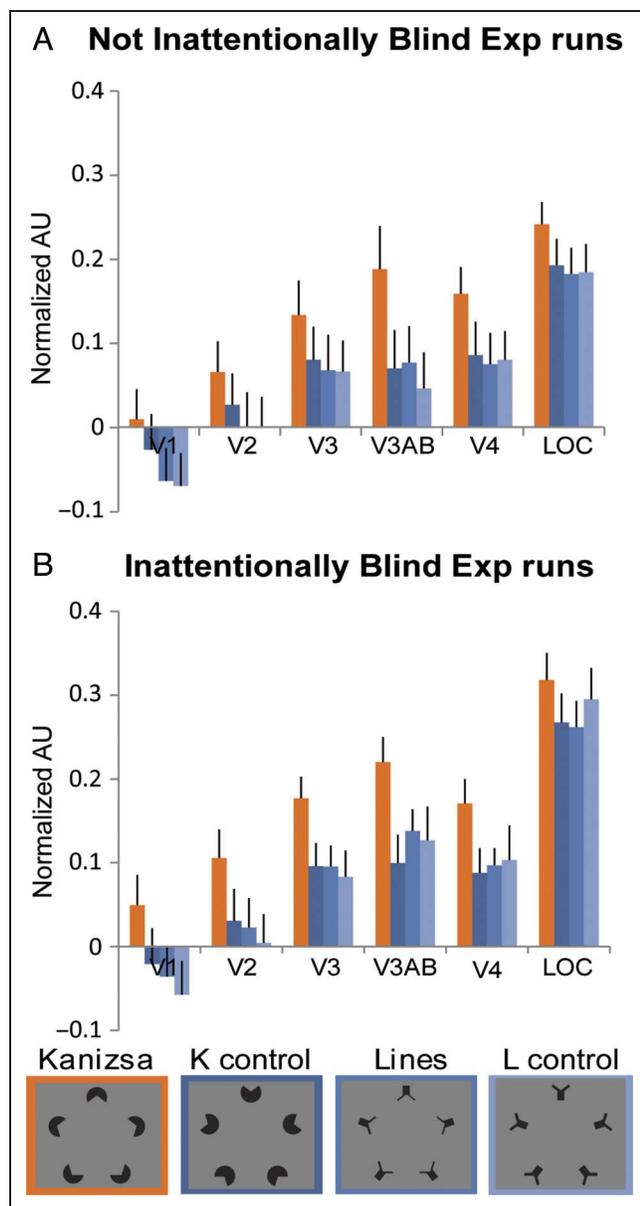


Figure 4. Univariate results for the experimental runs. Activity (in normalized arbitrary units) associated with the four figures in each ROI for the NIB group (A) and the IB group (B). In both groups, the Kanizsa figure elicited more activity across lower and higher visual areas. Error bars denote *SE*.

reported the figure. The lack of an interaction effect showed that the pattern of activation for the four figures was the same for the NIB and the IB group, $F(3, 66) = .2, p = .90$. This suggests that the Kanizsa illusion is processed even when the percept is not accessed or reported. There were seven participants in the IB group who chose one of the two Kanizsa-type foils in the experimental question (Figure 2). Although these foils have a different perceptual appearance than the target Kanizsa, it could be that these participants were able to report that they saw a Kanizsa-type figure but were not able to report the details of the figure they saw. To investigate

whether the results of the IB group were determined by these participants, we analyzed the five participants who chose option D–H to see whether the pattern of results was the same. Although statistical testing with five participants yields too little power, the pattern for these five participants was similar to the pattern for the whole group (Figure 5).

In addition to a main effect of Figure, there was a main effect of ROI, $F(2.05, 45.06) = 30.5, p < .001$, Greenhouse–Geisser corrected, and an interaction between Figure and ROI, $F(6.71, 147.78) = 5.6, p < .001$, Greenhouse–Geisser corrected. The main effect of ROI was driven by the fact

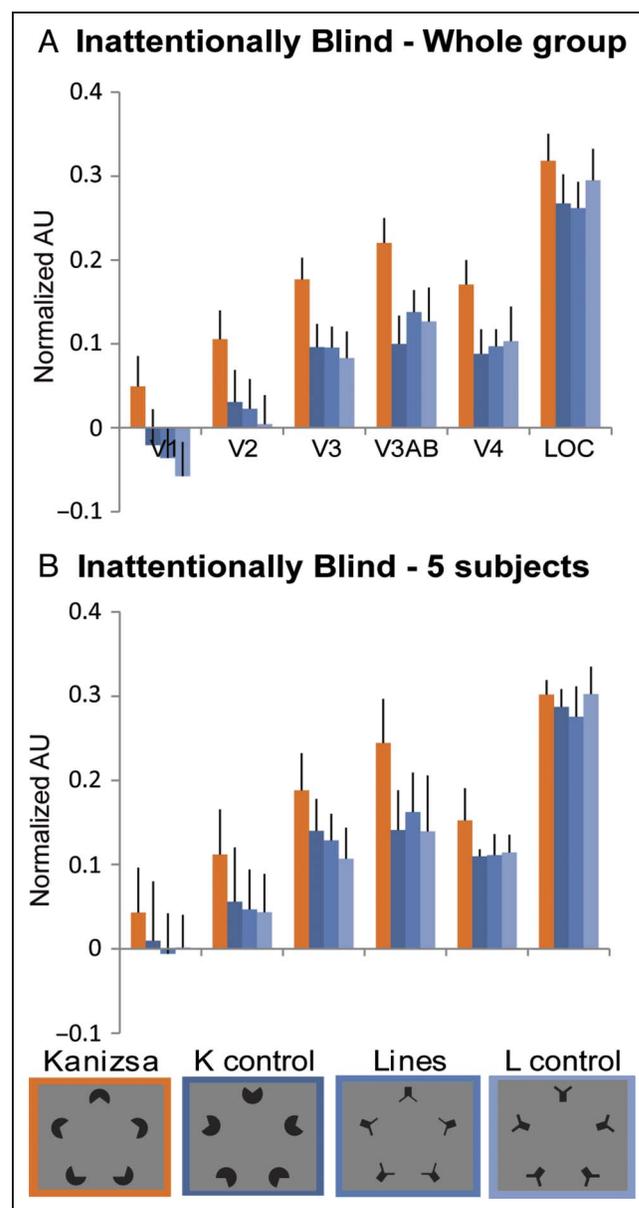


Figure 5. Univariate results for the whole IB group (A) and the five participants that chose a non-Kanizsa-type figure in the experimental question (B, see Figure 2). The activity pattern for the whole group and these five participants is similar, showing that the results are not dominated by the participants that selected a Kanizsa-type foil.

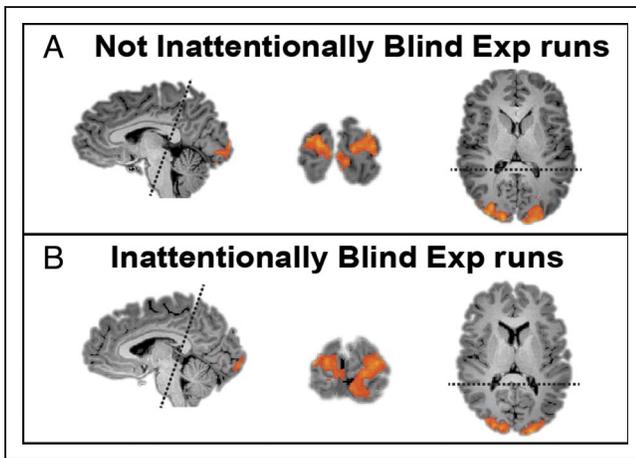


Figure 6. A whole-brain analysis was performed for the NIB (A) and IB groups (B), in which the Kanizsa figure was contrasted with the three control figures (Kanizsa > Controls). Activity is in lower and higher visual areas and does not differ between the two groups. Note that the scans did not cover the whole brain but only recorded visual cortex and extended into a part of the temporal and parietal cortex. Dotted lines indicate upper and lower limits of epi acquisition.

that overall activity in V1 was lower and activity in LOC was higher than the other ROIs. The interaction effect revealed that the difference between the three figures was largest in V3AB and smaller in LOC. The involvement of V3AB in addition to LOC has been found in previous literature as well (Mendola et al., 1999). No interaction effects with Group were found, all $p > .115$, Greenhouse–Geisser corrected, suggesting that the pattern of activity across ROIs was similar for the IB and NIB group.

To further investigate whether there were any differences between the NIB and the IB group apart from the ROIs specified, a multisubject whole-brain analysis was performed (note that the scans did not cover the front of the brain). Figure 6 shows the multisubject whole-brain analyses for the two groups separately (FDR = .05), which show that mainly the lower visual areas are involved, similarly for the NIB and IB group. There were no regions that were significantly more activated for the NIB group for the Kanizsa figures versus the other figures (FDR = .05).

We found a very clear and statistically strong modulation of activity in lower and higher visual areas for the Kanizsa figures compared with its control. That there was no difference in this pattern between the NIB and IB group suggests that the Kanizsa figure is processed similarly for the two groups. However, these conclusions are based on a null-result and therefore hard to interpret. Perhaps, if we would have tested more participants, a difference between groups would have become evident. Nevertheless, in both groups, the Kanizsa figure elicited more activity compared with its controls was clearly manifested, and testing more participants would not diminish this effect. Therefore, a potential difference between groups could only reveal a difference in the strength of the modulation.

Our univariate analyses were based on a between-subject design, making it more difficult to compare the neural patterns for reported and unreported conditions. Therefore, we performed multivariate analyses within participants. To define the consistency between the condition in which the Kanizsa figure was unreported and the condition in which it was reported, we compared the neural patterns for the experimental runs with those of the control run. Although from the mean activity change it seems that the pattern in the control run was similar to that in the experimental runs for both the NIB and the IB group (Figure 7; activity in all regions was higher for the Kanizsa figures versus the control figures), we

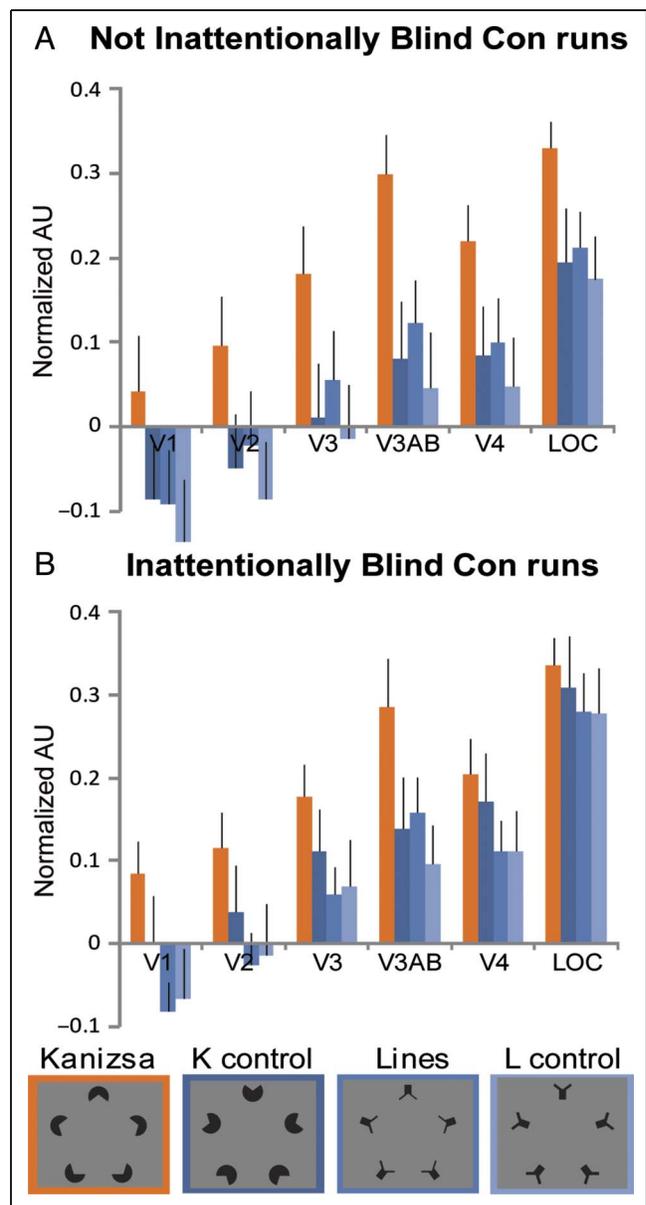


Figure 7. Univariate results for the control run. Activity (in normalized arbitrary units) associated with the four figures in each ROI for the NIB group (A) and the IB group (B). Error bars denote SE.

directly compared the two states by performing multivariate pattern analyses (MVPA).

MVPA

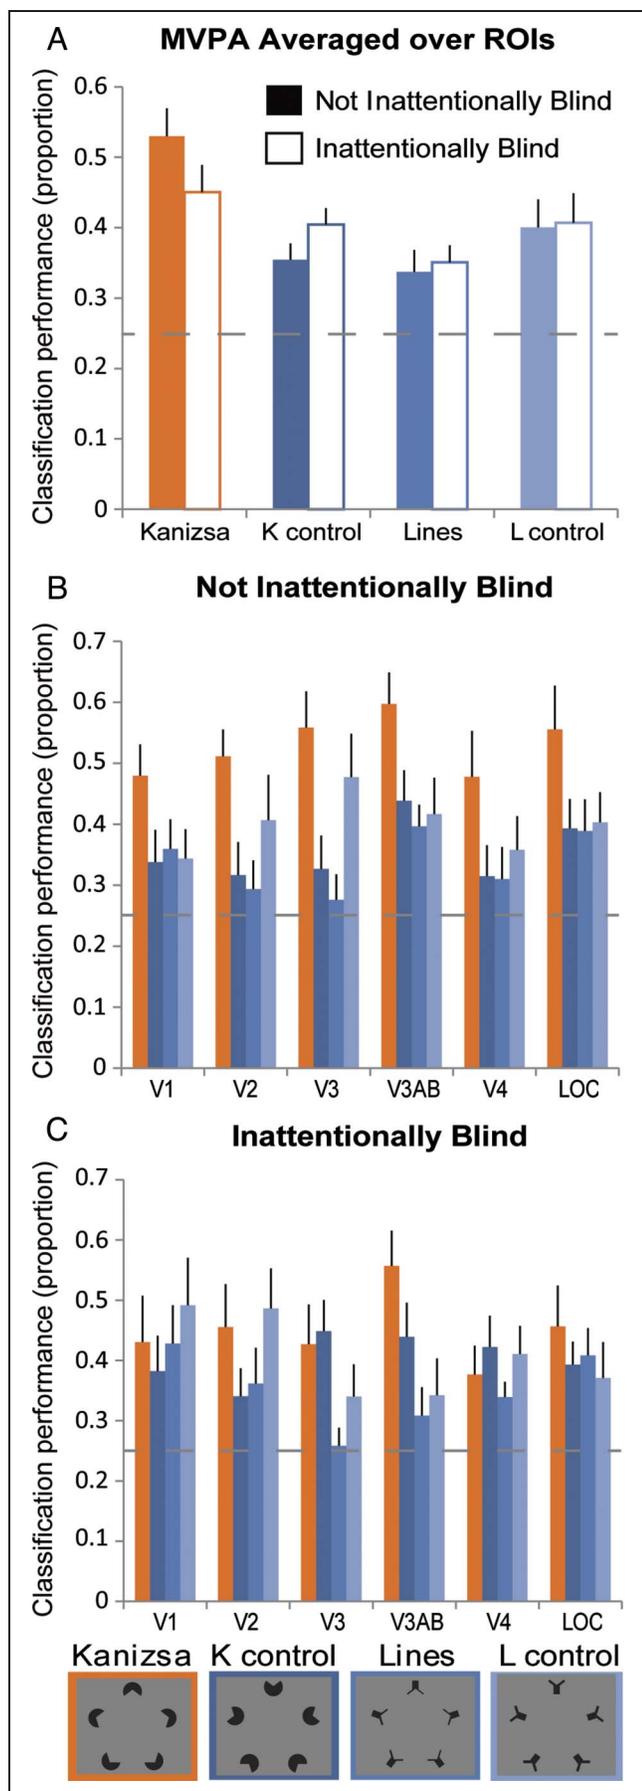
To compare the neural response of the Kanizsa figure during reported and unreported conditions, we used MVPA. With MVPA, we could predict which stimulus was seen in one set of trials based on the neural patterns obtained from an independent set of trials. This allowed us to directly compare the neural response during the experimental runs with those in the control run. First, we tested whether MVPA worked on our data set by performing the analysis within the three experimental runs. We hypothesized that because only the Kanizsa figure elicited a perceptually integrated percept, its multivoxel pattern should be more consistent than the voxel patterns underlying the three control figures and therefore result in higher decoding performance. The logic behind this reasoning stems from the conclusion drawn from the univariate results: Visual areas processing the illusion are more strongly activated, because there is more information present in the Kanizsa illusion versus the control conditions. When extending this conclusion to multivariate analyses, one would expect that the neural pattern underlying the Kanizsa illusion is most consistent because of the need to encode this information more precisely. To test this hypothesis, we used ROI voxel patterns for the four stimulus figures within each participant to classify each of the figures during the experimental runs. The two training runs were used to maximally distinguish the four neural patterns underlying each figure. Then the labels that were given to the test run were classified as correct or not (1 or 0), and a percentage correct for each figure presentation was obtained. We confirmed that classification performance was highest for the Kanizsa figure compared with the control figures, $F(3, 66) = 10.6$, $p < .001$, and there was no interaction between performance for the NIB and IB group, $F(3, 66) = 1.3$, $p = .282$. This shows that processing a Kanizsa figure is reflected in a more consistent multivoxel representation, regardless of whether participants were able to report the figure. Next, we wanted to examine whether this more consistent multivoxel representation for Kanizsa figures persists across runs in which the figure was reported compared with the runs in which it was not reported. To do so, we tested whether the patterns elicited during the experimental condition could be used to classify the patterns elicited in the control condition. We trained a neural pattern classifier on the three experimental runs and classified the patterns from the control run. Again, we calculated the classification (percentage correct) for each figure separately. If the patterns underlying the Kanizsa illusion remained the same and reportability has no influence on its neural representation, classification between the experimental and control runs is predicted to be better for the Kanizsa figures than for the control figures. If,

however, the underlying neural pattern changed because of access to the figure, classification should not work between the experimental and the control runs, resulting in similar or even lower classification performance for the Kanizsa figures than for the control figures.

Pattern classification was obtained for each ROI separately. Classification performance for all ROIs averaged and all ROIs separately is shown in Figure 8 (average chance performance = 25%, see Methods). We were able to determine which figure was presented during the control run based on the patterns resulting from the experimental run, as all figures could be classified well above chance. Again, classification worked best for the Kanizsa figure, $F(2.25, 49.50) = 11.6$, $p < .001$, Greenhouse–Geisser corrected; post hoc t tests compared with controls, all $p < .002$, Bonferroni-corrected. This shows that the multivoxel pattern underlying the Kanizsa figure is more consistent even across the experimental and control runs. There was no interaction between figure type and group, suggesting that average classification performance for the NIB and the IB group was the same (Figure 8A), $F(2.25, 49.50) = 2.2$, $p = .112$, Greenhouse–Geisser corrected. However, there was a trend toward significance. Although for both the NIB and the IB group the Kanizsa figure was classified best, $F(3, 33) = 8.9$, $p < .001$ and $F(3, 33) = 3.4$, $p = .029$, respectively, the Kanizsa figure might be better classified for the NIB than for the IB group. To test whether the lack of an interaction effect may have been because of a lack of power, we maximized power by averaging classification performance for the three control figures and tested against the Kanizsa figure. Now, we found an interaction between Figure and Group $F(1, 22) = 8.1$, $p = .01$, showing that the difference between the Kanizsa figure and its control figures was larger for the NIB group than for the IB group. This suggests that the neural pattern underlying the Kanizsa figure is more similar between the experimental run and the control run for the NIB group than for the IB group. As can be seen from Figure 8, the difference between the Kanizsa figure and its control figures for the NIB group was present in all ROIs, whereas for the IB group it was most pronounced in area V3A and LOC. Therefore, it seems that although overall activity in areas V1, V2, V3, and V4 was higher for the Kanizsa figure in both the unreported and reported conditions (see Figures 4 and 7), the consistency of activated voxels within these regions varied between states. The implications of these results are discussed below.

DISCUSSION

In this study, we examined the influence of access on the perceptual processing of the Kanizsa illusion using an inattentional blindness paradigm (Scholte et al., 2006; Rees et al., 1999; Simons & Chabris, 1999) in combination with fMRI. Although a large group of participants was not able



to select the Kanizsa figure after the experimental runs, these participants still displayed a unique neural pattern associated with processing a Kanizsa figure. The illusory figure elicited heightened activity in the areas that are critical for the perception of the illusion (Seghier & Vuilleumier, 2006). Also, MVPA showed that the neural signature of the Kanizsa during the unreported state could be used to classify its neural signature during the reported state, and classification performance was better for the Kanizsa figure than for any of the nonillusory conditions. Together, these results suggest that access is not necessary for the type of perceptual inference that underlies the Kanizsa illusion to take place in visual cortex.

The perceptual inference underlying the Kanizsa illusion is a complex process that involves the grouping of elements, surface segmentation, modulation of perceived brightness and depth, and the creation of illusory contours. Together, these processes lead to the perception of a figure lying on top of black disks. On the basis of the neural correlates found in this study, we cannot exclusively confirm which of these processes took place in the absence of access to the figure. In the paragraphs below, we will reflect on which processes probably occurred based on the presence of this specific neural pattern.

Mechanisms Underlying Kanizsa Processing

The neural correlates that accompanied Kanizsa figure processing in this study were related to the illusory nature of the percept and not to physical or cognitive stimulus attributes. By using two additional control figures (Figure 1C and D), we were able to dissociate perceptual inference from cognitive inference; although the Line figure had the same inducer layout and a pentagon could be cognitively inferred, the perception of illusory contours and a contrast difference between figure and background were absent. The absence of this illusory percept was reflected in the neural signature, such that the Line figure (Figure 1C) showed the same modulation as the Line control figure (Figure 1D), both in the accessible and unreported states. Moreover, the spatial frequency difference for the Line figure and its control was comparable with the spatial frequency difference for the Kanizsa figure and its control. Also, the elements of the Line figure could be grouped to form a coherent whole more easily compared with the Line control figure. The similarity between the neural signature for the Line figure and its control confirm that there was no perceptual characteristic added to this figure

Figure 8. Multivariate results. Classification performance in percentage correct averaged over all ROIs for the NIB group and the IB group (A) and for each separate ROI (B and C). The classification for Kanizsa figures is better than classification for the three control figures. Average chance performance was 25%, corrected for each figure separately (see Methods). Error bars denote *SE*.

and that the neural modulation observed for the Kanizsa figure was not because of differences in spatial frequency, collinearity, or grouping mechanisms. Therefore, we suggest that higher-level inference processes such as surface segmentation, modulation of contrast and depth perception, or the formation of illusory contours drove the neural signature that was found for the Kanizsa figure.

The Kanizsa illusion is a prime example of perceptual inference, a process that is linked to conscious processing (Wokke et al., 2013; Harris et al., 2011; Fahrenfort, Scholte, & Lamme, 2007). It has been shown that the Kanizsa illusion is not perceived when its inducers are masked (Harris et al., 2011), suggesting that conscious processing of the inducers is necessary to perceive the figure. However, other studies have shown that the Kanizsa illusion survives crowding (Lau & Cheung, 2012) and breaks through interocular suppression more easily (Wang, Weng, & He, 2012), suggesting that the Kanizsa illusion can be processed unconsciously. Although these findings seem to contradict each other, the perception of the Kanizsa illusion may depend on multiple mechanisms.

To process a Kanizsa figure, its inducers should be grouped and processed as one object. This process might be driven by grouping mechanisms that depend on fast, feedforward activity and could be performed unconsciously (Roelfsema, 2006). Then, the details of the figure—the specific illusory shape that is seen—are filled in by feedback mechanisms. In a recent TMS experiment, it was shown that the critical time window for V1/V2 in which discrimination of Kanizsa figures was affected occurred after the critical time window in which the LOC (Wokke et al., 2013)—an area that sits higher up the visual hierarchy and is involved in object detection (Malach, Levy, & Hasson, 2002)—was involved. Moreover, this effect occurred only when the support ratio of the Kanizsa inducers was large enough to clearly cause an illusory percept. Critically, for all the support ratios that were used, also for those not evoking an illusory percept, the inducers could be grouped. This suggests that V1 is causally involved in the shape formation of the illusion and not in the initial grouping of elements. These findings match with the behavioral findings of Wang et al. (2012) on the one hand and Harris et al. (2011) on the other hand. The Kanizsa figure may break through interocular suppression easier than a control figure: If the grouping of elements occurs unconsciously, a Kanizsa figure will be seen more easily than a control figure that cannot be grouped. In the study in which the Kanizsa inducers were masked, however, the critical manipulation was for participants to perceive which direction the illusory triangle was facing, and thus, the shape of the figure should be processed. If shape processing depends on feedback interactions, masking should indeed inhibit the formation of the shape (Harris et al., 2011; Fahrenfort et al., 2007). Together, these results suggest that perceiving the Kanizsa illusion depends on unconscious grouping mechanisms and conscious figure formation, which are

supported by feedforward and feedback mechanisms, respectively.

Although in the current study we were not able to directly test the involvement of feedforward and feedback mechanisms because of the temporal resolution of the fMRI signal, we found heightened activity in V1 and V2, although the receptive field sizes in V1 and V2 are an order of 6–12 times too small to encapsulate the entire Kanizsa figure (Smith, Singh, Williams, & Greenlee, 2001). This suggests that, in our study, feedback from higher areas modulated activity in lower visual areas, thereby suggesting that processes underlying shape formation that accompanies the perception of the pentagon itself had occurred (Wokke et al., 2013; Harris et al., 2011).

Contrary to the fact that we found heightened activity in V1/V2 for the Kanizsa figure both in the unreported and reported conditions, we did not find evidence for a more consistent voxel pattern within V1/V2 between unreported and reported conditions. Possibly, the modulation in lower visual areas should not be attributed to feedback from higher areas, but to a more general attention mechanisms: It could be that the Kanizsa figure drew more attention away from the central letter task, and therefore, participants had to put more resources into staying focused on the central task compared with when the control figures were presented. However, if a larger attentional demand could explain the neural modulation in lower visual areas, then this would still show that the participants who were IB also processed this figure to a certain extent. As our control figures rule out any lower level or cognitive inference mechanisms such as spatial frequency, collinearity, or grouping mechanisms, the IB group must have processed the Kanizsa figure at least up to a level where surface segmentation, a modulation of contrast/depth, or illusory contour formation took place.

There are more explanations for the higher mean activity, but not necessarily higher consistency for the Kanizsa figure between unreported and reported conditions. It might be that surface segmentation and perhaps a modulation of contrast/depth occurred through mechanisms of feedback, however, the illusory contour formation itself was less clear. Participants might have experienced a “blurrier” representation. Indeed, attention has been shown to optimize signal-to-noise levels by both signal enhancement and noise reduction (for an overview, see Carrasco, 2011). Reducing the noise and thus the variability in the signal might lead to a voxel pattern that is more consistent over time. Therefore, the neural pattern for the NIB group might have been more consistent between the experimental and control runs, whereas for the IB group, although higher-level inference processes took place, the neural patterns were less well defined. Also, the fact that the NIB group already reported about the figure after the experimental runs made their task during the experimental runs much more similar than for the IB group; they did not have to search for a specific figure. Perhaps the way in which the task was performed influenced the

neural patterns underlying the processing of the figures as well.

Inattentive Blindness and the Kanizsa Figure

To create inattentive blindness, we specifically manipulated cognitive load and not perceptual load (Yi et al., 2004). We chose to manipulate cognitive load to test the hypothesis that cognitive access is necessary for perceptual inference. IB participants could not identify the presented figure when they were uninformed about the configuration of the distracting stimuli during the experimental runs, although when they were informed about the configuration during the control run, they were able to select the correct figure although they maintained similar performance on the *n*-back task. This warrants the conclusion that the figure was potentially accessible, yet not accessed during its presentation. It may well be that if we had manipulated perceptual instead of cognitive load, results would have been different. Outcomes of neuronal modeling support the prediction that during inattentive blindness incoming sensory information might be processed but blocked from access because the network is engaged in processing distracting information (Dehaene & Changeux, 2005). In that sense, the paradigm of inattentive blindness formalizes the common intuition that many stimuli in plain sight remain unnoticed and are therefore never accessed, although they are potentially accessible, and forms the most rigorous test of the fate of unaccessed visual stimuli.

A problem of inattentive blindness paradigms, in which cognitive load and not perceptual load is manipulated, is that it is inherently a between-subject design: If a participant knows that a certain figure is present during the task, they will be much more prone to notice it on a next run. To obtain a sufficient amount of data, we therefore chose to present participants with three runs before asking them about the presence of the Kanizsa figure. Moreover, during runs, participants were presented with three control figures as well. A potential alternative explanation for the participants' behavior is that the inability to select the correct figure was not a result of inattentive blindness, but of inattentive amnesia (Wolfe, 1999). It could be the case that participants were able to access the figure at the moment of presentation, but a memory failure—perhaps because of overwriting of the subsequently presented control figures—prevented them from selecting the correct figure when asked about it. However, in comparison with studies where the target figure was presented just once (Thakral, 2011; Simons & Chabris, 1999), in our study, the figure was presented 18 times, each for a period of 14.4 sec. The question about these figures was then asked within 1 min after the last stimulus presentation. This makes it improbable that the failure to select the correct figure was because of simple memory failure. Moreover, the same participants were perfectly able to select the figure when in the con-

trol run their task instruction was to pay attention to the figures; thus, it is not the case that these participants simply had bad visual memory.

Another possibility might be that the IB participants had more imprecise memory than the NIB participants and therefore chose the wrong figure. Indeed, this could have been the case for the seven participants in the IB group that chose a Kanizsa-type foil. However, the five participants that did not choose a Kanizsa-type foil displayed the same neural pattern, suggesting that the results for the IB group were not driven by participants who had possibly caught a glimpse of the Kanizsa figure.

Although we cannot be certain that participants in the IB group had absolutely no access to the figure during the presentation itself, the main point is that in both frameworks—inattentive blindness or amnesia—the participant was not able to report or recognize a recently presented figure even when forced to make a decision. One of the functions of cognitive access is storage in working memory, and one could say that access has failed when participants cannot report about objects that were presented multiple times only a few seconds ago. At the same time, the neural signature coding for the processing of the figure was clearly present. This confirms the importance of using brain measurements when investigating the nature of visual representations instead of relying on behavioral measures of report only (Kanai & Tsuchiya, 2012; Lamme, 2010).

In this study, the neural correlates that we found for the NIB and IB groups were similarly present in lower and higher visual areas. This suggests that the areas involved in processing the figure itself did not differ between conditions. Possibly, whether the percept could be reported or not depended on neural patterns in frontal or frontoparietal areas (not included in the current EPI sequence; Carmel, Lavie, & Rees, 2006; Lumer & Rees, 1999; although see Thakral, 2011). However, the neural patterns that were unique to the Kanizsa illusion were present in both lower- and higher-level visual areas for the IB group. Therefore, even if the difference in reportability depended on activity in more frontal areas, this did not change the neural patterns associated with Kanizsa processing such as surface segmentation, the modulation of contrast difference, and perhaps illusory contour formation such as isolated in this study and other studies investigating Kanizsa processing (Seghier & Vuilleumier, 2006). This suggests that access to a stimulus does not alter the neural mechanisms that are involved in processing the perceptual characteristics of that stimulus. Instead, it might make the representation globally available for cognitive manipulation and report. Possibly, participants in the NIB group had some leftover attention that they allocated to the distracter stimuli, thereby gaining access to the Kanizsa figure: As discussed above, attention might have improved the neural signature, making the neural pattern better defined. This might explain why the classification scores for the Kanizsa figure in the NIB group

trended to be higher than those in the IB group: Access might make the neural patterns more pronounced and consistent over time. At the same time, this does not weaken or disqualify our main finding that processes of perceptual inference associated with the Kanizsa illusion are present without access (or attention).

Converging Evidence with Patient Studies

The results found in this study are in line with previous literature on the perception of Kanizsa figures for patients with hemispatial neglect or parietal extinction (Conci et al., 2009; Vuilleumier, Valenza, & Landis, 2001; Mattingley, Davis, & Driver, 1997). These patients have a deficit in perceiving stimuli in their left visual field when they are simultaneously presented with stimuli in their right visual field. Presumably, this deficit stems from their unilateral parietal brain damage that causes an inability to attend to the left visual field when a stimulus is presented in the right visual field. When presented with a bilateral Kanizsa figure, patients are unable to judge the presence or similarity of Kanizsa inducers on the left side of their visual field, but they are nevertheless able to make judgments about the surface of the Kanizsa figure. This shows that, although half of the inducers are not well perceived because of a deficit in paying attention to these stimuli, surface segmentation remains intact. In the current study, we showed that, indeed, the neural signature in lower visual areas underlying processes such as surface segmentation remain intact despite a lack of attention. This confirms that attention and thus the ability to report about (the details of) a figure do not alter the way in which this figure is processed on a visual level.

Implications for Consciousness Research

One of the main questions in consciousness research nowadays is whether attention—and therefore access—to a percept is crucial for consciousness to occur (Block, 2007, 2011; Cohen & Dennett, 2011; Lau & Rosenthal, 2011; Kouider, De Gardelle, Sackur, & Dupoux, 2010; Lamme, 2006, 2010; Koch & Tsuchiya, 2007; Fahrenfort & Lamme, 2012). In the current study, we showed that (at least some of the) neural processes underlying perceptual inference occur even when participants are not able to report about their percept. As perceiving the Kanizsa illusion is highly likely to depend on feedback from higher to lower visual areas (Wokke et al., 2013; Harris et al., 2011; Knebel & Murray, 2012) and we found modulation of lower visual areas, this suggests that higher-level integration can take place in the absence of access to the percept. If one takes this type of perceptual inference as an indicator of conscious processing, one could conclude that consciousness can occur in the absence of attention. The fact that such processes can occur in the absence of direct access appeals to the rich visual world we experience when we

look around: We do not have to pay attention to every object to be able to experience the world as a whole. However, more work is needed to disentangle which processes can and cannot occur in the absence of attention and whether these processes qualify as being termed conscious or not.

Conclusions

This study extends previous work showing that several perceptual processes such as figure ground segregation (Scholte et al., 2006), feature grouping (Pitts, Martínez, & Hillyard, 2012; Moore & Egeth, 1997), and visual context effects (Lathrop, Bridgeman, & Tseng, 2011) occur during inattentive blindness. In this study, we show that the neural Kanizsa signature does not subside when participants are not able to report about their percept. Importantly, the Kanizsa figure is accompanied by a unique signature that is absent when controlling for confounding factors such as collinearity, spatial frequency, grouping, and cognitive inference. Including these controls strongly suggests that the signature is a correlate of the illusory percept itself and not of something else. This implies that one or more processes underlying this type of perceptual inference occurs in the absence of access to the percept, potentially putting nonaccessed states in the realm of conscious rather than unconscious processing (Lamme, 2010).

Acknowledgments

This work was made possible by an advanced ERC grant to V. A. F. L. We thank Ned Block for his helpful comments on this manuscript.

Reprint requests should be sent to Annelinde R. E. Vandembroucke, Department of Psychology, University of Amsterdam, Weesperplein 4, 1018 XA, Amsterdam, the Netherlands, or via e-mail: a.r.e.vandembroucke@uva.nl.

REFERENCES

- Becker, M. F., & Knopp, J. (1978). Processing of visual illusions in the frequency and spatial domains. *Perception & Psychophysics*, *23*, 521–526.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral Brain Sciences*, *30*, 481–499; discussion 499–548.
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*, 567–575.
- Carmel, D., Lavie, N., & Rees, G. (2006). Conscious awareness of flicker in humans involves frontal and parietal cortex. *Current Biology*, *16*, 907–911.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*, 1484–1525.
- Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, *15*, 358–364.
- Conci, M., Böbel, E., Matthias, E., Keller, I., Müller, H. J., & Finke, K. (2009). Preattentive surface and contour

- grouping in Kanizsa figures: Evidence from parietal extinction. *Neuropsychologia*, *47*, 726–732.
- Davis, G., & Driver, J. (1998). Kanizsa subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 169–184.
- Dehaene, S., & Changeux, J.-P. (2005). Ongoing spontaneous activity controls access to consciousness: A neuronal model for inattentional blindness. *PLoS Biology*, *3*, e141.
- Fahrenfort, J. J., & Lamme, V. A. F. (2012). A true science of consciousness explains phenomenology: Comment on Cohen and Dennett. *Trends in Cognitive Sciences*, *16*, 138–139.
- Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2007). Masking disrupts reentrant processing in human visual cortex. *Journal of Cognitive Neuroscience*, *19*, 1488–1497.
- Ginsburg, A. P. (1975). Is the illusory triangle physical or imaginary? *Nature*, *257*, 219–220.
- Goebel, R., Esposito, F., & Formisano, E. (2006). Analysis of functional image analysis contest (FIAC) data with brainvoyager QX: From single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. *Human Brain Mapping*, *27*, 392–401.
- Halgren, E., Mendola, J., Chong, C. D. R., & Dale, A. M. (2003). Cortical activation to illusory shapes as measured with magnetoencephalography. *Neuroimage*, *18*, 1001–1009.
- Harris, J. J., Schwarzkopf, D. S., Song, C., Bahrami, B., & Rees, G. (2011). Contextual illusions reveal the limit of unconscious visual processing. *Psychological Science*, *22*, 399.
- Kanai, R., & Tsuchiya, N. (2012). Qualia. *Current Biology*, *22*, R392–R396.
- Kanizsa, G. (1974). Contours without gradients or cognitive contours. *Italian Journal of Psychology*, *1*, 93–113.
- Kanizsa, G. (1976). Subjective contours. *Scientific American*, *234*, 48–52.
- Knebel, J.-F., & Murray, M. M. (2012). Towards a resolution of conflicting models of illusory contour processing in humans. *Neuroimage*, *59*, 2808–2817.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, *11*, 16–22.
- Kouider, S., De Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*, 301–307.
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, *10*, 494–501.
- Lamme, V. A. F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience, iFirst*, 1–21.
- Lathrop, W. B., Bridgeman, B., & Tseng, P. (2011). Perception in the absence of attention: Perceptual processing in the Roelofs effect during inattentional blindness. *Perception*, *40*, 1104–1119.
- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, *15*, 365–373.
- Lau, J. S. F., & Cheung, S. (2012). Illusory contour formation survives crowding. *Journal of Vision*, *12*, 15.
- Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences*, *9*, 75–82.
- Lee, T. S., & Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences*, *98*, 1907.
- Lumer, E. D., & Rees, G. (1999). Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proceedings of the National Academy of Sciences*, *96*, 1669–1673.
- Maertens, M., Pollmann, S., Hanke, M., Mildner, T., & Möller, H. (2008). Retinotopic activation in response to subjective contours in primary visual cortex. *Frontiers in Human Neuroscience*, *2*.
- Malach, R., Levy, I., & Hasson, U. (2002). The topography of high-order human object areas. *Trends in Cognitive Sciences*, *6*, 176–184.
- Mattingley, J. B., Davis, G., & Driver, J. (1997). Preattentive filling-in of visual surfaces in parietal extinction. *Science*, *275*, 671–674.
- Mendola, J. D., Dale, A. M., Fischl, B., Liu, A. K., & Tootell, R. B. H. (1999). The representation of illusory and real contours in human cortical visual areas revealed by functional magnetic resonance imaging. *The Journal of Neuroscience*, *19*, 8560–8572.
- Moore, C. M., & Egeth, H. (1997). Perception without attention: Evidence of grouping under conditions of inattention. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 339–352.
- Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science*, *257*, 1357–1363.
- Pitts, M. A., Martínez, A., & Hillyard, S. A. (2012). Visual processing of contour patterns under conditions of inattentional blindness. *Journal of Cognitive Neuroscience*, *24*, 287–303.
- Pylshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral Brain Sciences*, *2*, 341–365. Open peer commentary 366–423.
- Rees, G., Russell, C., Frith, C. D., & Driver, J. (1999). Inattentional blindness versus inattentional amnesia for fixated but ignored words. *Science*, *286*, 2504–2507.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*, 368–373.
- Roelfsema, P. R. (2006). Cortical algorithms for perceptual grouping. *Annual Reviews Neuroscience*, *29*, 203–227.
- Scholte, H. S., Jolij, J., Fahrenfort, J. J., & Lamme, V. A. F. (2008). Feedforward and recurrent processing in scene segmentation: Electroencephalography and functional magnetic resonance imaging. *Journal of Cognitive Neuroscience*, *20*, 2097–2109.
- Scholte, H. S., Witteveen, S. C., Spekrijse, H., & Lamme, V. A. F. (2006). The influence of inattention on the neural correlates of scene segmentation. *Brain Research*, *1076*, 106–115.
- Seghier, M., Dojat, M., Delon-Martin, C., Rubin, C., Warnking, J., Segebarth, C., et al. (2000). Moving illusory contours activate primary visual cortex: An fMRI study. *Cerebral Cortex*, *10*, 663–670.
- Seghier, M., & Vuilleumier, P. (2006). Functional neuroimaging findings on the human perception of illusory contours. *Neuroscience & Biobehavioral Reviews*, *30*, 595–612.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, *28*, 1059–1074.
- Smith, A. T., Singh, K. H., Williams, A. L., & Greenlee, M. W. (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cerebral Cortex*, *11*, 1182–1190.

- Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain: 3-Dimensional proportional system: An approach to cerebral imaging*. New York: Thieme.
- Thakral, P. P. (2011). The neural substrates associated with inattention blindness. *Consciousness and Cognition, 20*, 1768–1775.
- Vuilleumier, P., Valenza, N., & Landis, T. (2001). Explicit and implicit perception of illusory contours in unilateral spatial neglect: Behavioral and anatomical correlates of preattentive grouping mechanisms. *Neuropsychologia, 39*, 597–610.
- Wandell, B. A., & Winawer, J. (2011). Imaging retinotopic maps in the human brain. *Vision Research, 51*, 718–737.
- Wang, L., Weng, X., & He, S. (2012). Perceptual grouping without awareness: Superiority of Kanizsa triangle in breaking interocular suppression. *PLoS One, 7*, e40106.
- Wokke, M. E., Vandenbroucke, A. R. E., Scholte, H. S., & Lamme, V. A. F. (2013). Confuse your illusion: Feedback to early visual cortex contributes to perceptual completion. *Psychological Science, 24*, 63–71.
- Wolfe, J. M. (1999). Inattention blindness. In V. Coltheart (Ed.), *Fleeting memories* (pp. 71–94). Cambridge, MA: MIT Press.
- Yi, D.-J., Woodman, G., Widders, D., Marois, R., & Chun, M. M. (2004). The neural fate of ignored stimuli: Dissociable effects of perceptual and working-memory load. *Nature Neuroscience, 7*, 992–996.