Ann. N.Y. Acad. Sci. ISSN 0077-8923

ANNALS OF THE NEW YORK ACADEMY OF SCIENCES
Issue: *Competitive Visual Processing Across Space and Time*

Seeing without knowing: task relevance dissociates between visual awareness and recognition

Baruch Eitam, Roy Shoval, and Yaffa Yeshurun

Department of Psychology, University of Haifa, Haifa, Israel

Address for correspondence: Baruch Eitam, Department of Psychology, University of Haifa, Mount Carmel, Haifa 31905, Israel. beitam@psy.haifa.ac.il

We demonstrate that task relevance dissociates between visual awareness and knowledge activation to create a state of seeing without knowing—visual awareness of familiar stimuli without recognizing them. We rely on the fact that in order to experience a Kanizsa illusion, participants must be aware of its inducers. While people can indicate the orientation of the illusory rectangle with great ease (signifying that they have consciously experienced the illusion's inducers), almost 30% of them could not report the inducers' color. Thus, people can see, in the sense of phenomenally experiencing, but not know, in the sense of recognizing what the object is or activating appropriate knowledge about it. Experiment 2 tests whether relevance-based selection operates within objects and shows that, contrary to the pattern of results found with features of different objects in our previous studies and replicated in Experiment 1, selection does not occur when both relevant and irrelevant features belong to the same object. We discuss these findings in relation to the existing theories of consciousness and to attention and inattentional blindness, and the role of cognitive load, object-based attention, and the use of self-reports as measures of awareness.

Keywords: visual awareness; task relevance; selective attention; inattentional blindness

Introduction

But for those first affections,
Those shadowy recollections,
Which, be they what they may,
Are yet the fountain light of all our day,
Are yet a master light of all our seeing.

Wordsworth *Ode* in *Poems* II. 155

Notwithstanding Wordsworth's beautiful depiction of lucid (earthly) knowledge-less visual experiences, we intuitively believe that when we consciously perceive something familiar we also recognize it. That is, we know what it is. It seems almost nonsensical to argue otherwise. If you consciously perceive a green apple, you know, of course, that in front of your eyes is a green apple, a fruit, which tastes somewhat sour, that its peel has a smooth texture, and maybe even that it figures in that famous story of a forsaken garden. This statement can be put even stronger—seeing is knowing—awareness and knowing (given appropriate prior knowledge) are one and the same. But beyond our fallible intuition, why should we assume that visual awareness and knowing are so tightly bound? In fact, although visual awareness and recognition[a] usually cooccur, the relationship between them is far from clear either theoretically or empirically. Accordingly, the primary goal of this paper is to put to test this assumed equivalence of visual awareness and "knowing" by exploring whether a stimulus is necessarily recognized when a person is phenomenally aware of it.

[a]We use "recognition" of something or "to recognize" something to mean "to know what it is"—to identify something as being the same as or related to something previously known, to be able to name and associate it with prior knowledge.

doi: 10.1111/nyas.12673

Motivational relevance dissociates between visual awareness and knowledge activation

Substantial empirical evidence has accumulated to suggest that unintentional knowledge activation[b] can occur even when people are not visually aware of the inducing stimulus.[1–3] One such demonstration of knowledge activation without awareness is subliminal priming, which has been shown to briefly activate semantic knowledge even in lieu of a phenomenal experience of the inducing word.[4] Other demonstrations include priming during the attentional blink;[5] intact semantic priming in the neurological condition of unilateral neglect;[6] and more contentiously, semantic priming while suppressing awareness through continuous flash suppression (CFS).[7] In other words, it has been reasonably established that stimuli may lead to knowledge activation without being phenomenally experienced. However, here we focus on the complimentary possibility, namely that of seeing without knowing or awareness of a stimulus without corresponding activation of available knowledge about it.[c] This possibility has not yet been empirically explored.

Even if seeing without knowing is logically possible, can we speculate under what conditions it would be expected? We argue that visual awareness without recognition may occur with stimuli that are currently task irrelevant. Reviewing the social–cognitive literature on the effects of priming, Eitam and Higgins recently proposed that knowledge activation is modulated by its motivational relevance.[1,3,8] Specifically, they proposed that, although incoming stimuli may potentiate corresponding internal representations (e.g., semantic representations), whether these representations will eventually be available for cognitive processing depends on the outcome of a process that evaluates their degree of relevance (e.g., for the task at hand)—only sufficiently relevant representations become functionally active. Because this framework (named relevance of activated representations (ROAR)) assigns a cardinal role to (e.g., task) relevance, it predicts that irrelevant stimuli will have a negligible influence on current and future thought and action, even when they are presented at what appears to be the center of the attentional window. This prediction has been corroborated using measures of both implicit[9–11] and explicit[12] processing, and held whether processing resources were available for the processing of the irrelevant stimuli[10,12] or not.[10] In one of these demonstrations[12] that is most relevant for this study, we presented a colored circle surrounded by a larger, differently colored circle. The circles were presented in the center of the screen for 500 ms, without masking. One of the two circles was deemed relevant by asking the participants to concentrate on it. Following the circles' offset, participants were asked to identify the colors of both relevant and irrelevant circles (report order was counterbalanced). We found that one-fourth of the participants could not report the color of the irrelevant one. Given the simplicity of the visual display (two central, static circles), the long presentation duration, the lack of masking, and the extreme ease of the task (concentrate on one circle), we assumed that both cognitive and perceptual resources were available for the processing of both circles. Nevertheless, we also corroborated this assumption empirically by showing that recognition rates were near perfect when instructions made both circles relevant. Hence, the failure to report the color of the irrelevant circle was solely due to the fact that it was not relevant, as predicted by the ROAR framework. That is, even though the participants had enough resources, they often did not fully process the irrelevant circle merely because it was deemed irrelevant. We therefore termed this failure to report a clearly visible stimulus as pure irrelevance-induced "blindness."

In this study, we focus on another intriguing possibility afforded by the ROAR framework—specifically, that visual awareness can occur although knowledge activation is insufficient to allow recognition. Because, according to the ROAR framework, the degree of relevance determines the level of knowledge activation, not access to consciousness,

[b]We use "knowledge activation"—as it is commonly used in the social–cognitive literature—to refer to available (but latent) high-level (semantic, abstract) knowledge becoming functionally available to cognitive subsystems. Unlike how some priming studies are explained (i.e., as response priming), it enables explanation of more intricate priming effects (e.g., effects of priming on person perception, effects of priming on problem solving).

[c]Note that unlike previous usage of the word "seeing" to mean perception without awareness we use it to mean perception with awareness in the sense that what was seen was phenomenally experienced.

it suggests that experimentally lowering the relevance of information should lower the probability that information would be fully activated but does not entail that visual awareness of the stimulus conveying this information would be affected. In other words, ROAR is consistent with the possibility of seeing without knowing and, crucially, enables us to predict what should be seen but not known.[d]

The results of our previous study[12] are consistent with the prospect of seeing without knowing; given the simplicity of the stimuli, their central presentation, and lack of any form of processing load, it seems reasonable to assume that, although the participants often failed to report the color of the irrelevant circle, they were most likely aware of it (in the phenomenal sense of experiencing) while it was presented. In this study we aim to put this possibility to direct empirical test.

This study

To support the claim that people were indeed aware of the irrelevant stimuli, we would need to show that people who were unable to report the identity of the stimulus were nevertheless phenomenally aware of it. Demonstrating awareness that is coupled with a failure to report is obviously challenging, as we almost always rely on the ability to report as an indication of awareness. We therefore had to choose a less direct tactic. Our solution was to employ a task in which adequate performance depends on phenomenal awareness—if participants are able to perform such a task when it is based on irrelevant stimuli and also fail to report their identity, then we have taken a significant step toward establishing the possibility of seeing without knowing. This is the primary goal of this study and is examined in Experiment 1.

Another goal of this study was to further explore the relevance-based selection framework itself. While we have now established the "blinding" effect of irrelevance using different stimuli (e.g., colors, shapes) and tasks (e.g., passive looking, active search, *n*-back memory), we do not know anything about the boundary conditions for the effect. This is examined in Experiment 2.

Experiment 1

The goal of this experiment was to test the hypothesis that when capacity limitations are not met, and perceptual and cognitive resources are available for processing the entire visual display, irrelevant stimulation reaches awareness, but does not result in sufficient knowledge activation to allow recognition. According to this hypothesis, under conditions of minimal perceptual and cognitive load, people should be phenomenally aware of an irrelevant stimulus but fail to report its identity. The empirical challenge presented by this hypothesis lays in demonstrating that participants were aware of the irrelevant stimulus even though they fail to report its attributes—the gold standard measure of awareness. To this end, we took a novel approach; we build on evidence that the perception of illusory Kanizsa figures[13] (Fig. 1, inset) involves conscious processing.[14–16] Specifically, it has been shown that when the inducers of a Kanizsa triangle were made subjectively invisible through CFS, participants could not report the pointing direction of the illusory triangle.[15] This finding is strong evidence that conscious perception of a Kanizsa figure requires conscious perception of its inducers. Hence, the ability to report attributes of a task-relevant Kanizsa figure indicates awareness of its (task-irrelevant) inducers. If, at the same time, the participants are unable to report the attributes of the task-irrelevant inducers, this will be a demonstration of awareness without full recognition.

To test our hypothesis, we employed a Kanizsa rectangle with either a vertical or horizontal orientation, brought about by four same-color inducers (Fig. 1, inset). In the key condition, the rectangle-relevant condition, task instructions required identification of the rectangle's orientation, deeming the illusory rectangle relevant and the inducers irrelevant. In the control condition, the color-relevant condition, task instructions required identification of the inducers' color, deeming them relevant and the rectangle irrelevant. Similar to our previous study,[12] to keep load levels minimal, we presented the display for a relatively long duration (500 ms) without masking, and employed a simple discrimination task (Was the rectangle's orientation

[d]Note that unlike the framework of phenomenal overflow, in ROAR the assumption is not that relevance (there, attention) is only necessary for conscious access to the (otherwise) activated representation; rather, relevance is necessary for access to any cognitive subsystem. We elaborate this point further in the "General discussion" section.

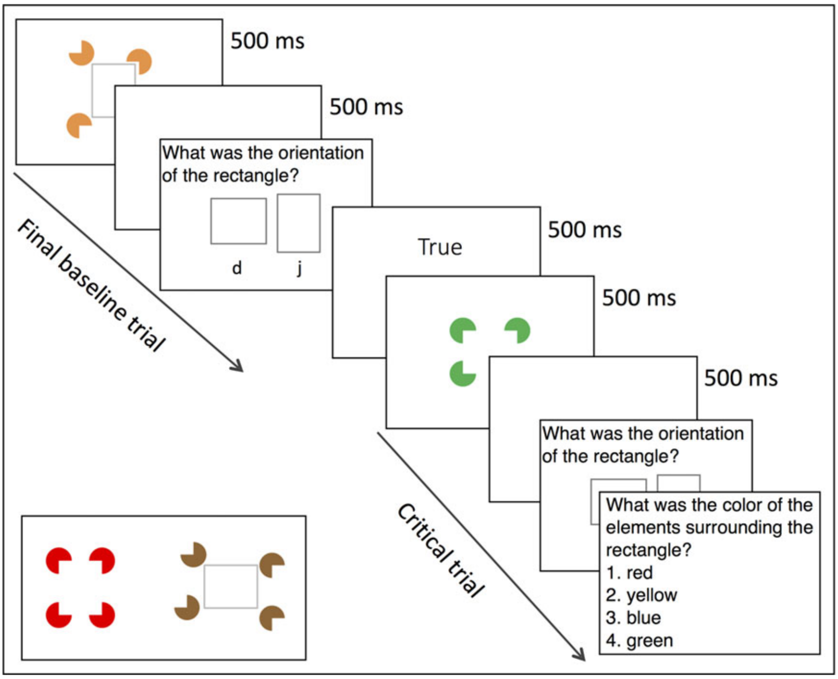

Figure 1. An illustration of the final trials in the rectangle-relevant condition of Experiment 1. The final baseline trial, depicted in the figure, was similar to all previous baseline trials (in the color-relevant condition, the question in the baseline trials was about the color of the inducers). The critical trial, in which participants were asked about both the rectangle's orientation and the inducers' color, was identical in both conditions (question order was counterbalanced). The inset depicts schematic examples of an illusory and real rectangle.

horizontal or vertical?). On the final (critical) trial and following the inducers' offset, participants of both conditions were asked to report both the rectangle's orientation and the inducers' color (order counterbalanced in both conditions).

In light of the finding that the perception of Kanizsa figures necessitates conscious perception of the inducers,[15] and given our hypothesis that when capacity limitations are not met, irrelevant information reaches awareness but does not result in recognition, we hypothesized that when the Kanizsa rectangle is deemed relevant but its inducers are not, the irrelevant inducers should reach awareness and therefore participants should be able to perceive the illusory rectangle and report its orientation; but because the inducers themselves are irrelevant, their activation should not suffice for recognition and the participants should find reporting their color difficult. Such a pattern would indicate that the irrelevant stimulation—the inducers—were consciously perceived but were not recognized. In other words, they were seen but not known.

Methods

Participants. One hundred and five University of Haifa students participated for class credit or pay. They were randomly assigned to one of two conditions: illusory rectangle-relevant ($n = 48$) and color-relevant ($n = 57$) conditions.[e]

Materials and procedure. Participants sat approximately 60 cm from the computer screen and

[e] We collected data in an additional rectangle-relevant condition in which the rectangle was never illusory. However, due to technical reasons, this condition ended up with fewer participants, so that after accuracy-based exclusion it included fewer than 40 participants ($n = 34$). We, therefore, did not include it in the general analysis. Importantly, the pattern of results in this condition is similar to that of the illusory rectangle-relevant condition: reporting the rectangle's orientation was not significantly different from perfect performance ($P = 0.12$), but report rate of the inducers' color was lower than perfect performance with marginal significance ($P = 0.057$).

viewed a (2 × 2.5 cm) rectangle surrounded by four colored Pac-Man–like shapes (henceforth referred to as inducers)—one at each corner of the rectangle (Fig. 1, inset). The inducers were created from a full circle (radius 0.5 cm) with one quarter missing. On each trial, the rectangle could appear with either horizontal or vertical orientation, chosen randomly, and the surrounding inducers all had the same color—one of eight possible colors (yellow, blue, light blue, brown, purple, orange, pink, or light green). The stimuli appeared at the center of the screen for 500 ms and were followed by a 500-ms blank screen and then by a question the participants had to answer (Fig. 1). In the baseline trials (the first eight trials in the rectangle-relevant condition and the first seven trials in the color-relevant condition), the rectangle had real contours (i.e., it was defined by a thin gray frame) and the inducers' color and individual orientation were chosen randomly. During these trials the participants' task in the rectangle-relevant condition was to indicate the orientation of the rectangle by choosing one orientation out of two options (vertical versus horizontal), and in the color-relevant condition to indicate the color of the inducers by choosing one color out of four options (the inducers' color and three others chosen randomly). An answer was coded as correct only if the actual color of the inducers was chosen. Visual feedback followed their response. The purpose of the baseline trials was to strengthen and stabilize the relevance of the task-relevant stimulation. Baseline trials were followed by the critical trial. In this trial, all inducers were colored either red or green and, critically, the rectangle was illusory (i.e., there was no gray frame; the rectangle's illusory contours were brought about by the inducers, whose individual orientation was chosen to generate an illusory rectangle).[13] On this final trial, the participants were presented with an additional question in which they were asked to report the identity of the task-irrelevant stimulation—the rectangle's orientation for the color-relevant condition or the inducers' color for the rectangle-relevant condition. The format of the questions was as described and their order (relevant versus irrelevant) was counterbalanced. No feedback was given. Once participants responded to the second question, they indicated on a scale of 1–100 how sure they were about that response and how they think they reached that answer: using memory, familiarity, imagery, intuition, or by guessing.[17]

Results and discussion

Four participants from the rectangle-relevant condition and six participants from the color-relevant condition were excluded from the analysis for committing more than a single error on the baseline trials (baseline accuracy <75%). The following analyses were conducted on performance in the critical trial of the remaining participants (rectangle relevant: $n = 44$; color relevant: $n = 51$). Whenever possible we performed a one-tailed z-test for proportions; when not possible (i.e., one or more of the cells included fewer than five cases) we used a one-tailed Fisher's exact test. Frequencies of correct versus incorrect responses and their corresponding proportions are presented in Table 1.

Testing the effect of relevance. First, we tested whether we replicated the "pure" relevance-based selection demonstrated in the past.[12] That is, we tested the prediction that, even when resource limitations are not met, the ability to report an irrelevant attribute of the stimulation is considerably worse compared to when it is relevant. To that end, we compared performance on the relevant and irrelevant conditions for both orientation and color attributes. To ensure that the participants attended the relevant attribute, we analyzed performance for the irrelevant attribute only for participants who answered the relevant question correctly. Corroborating our prediction, the proportion of correctly reporting the rectangle orientation when it was relevant (0.96; orientation question in the rectangle-relevant condition) was significantly higher ($P < 0.0001$) than when it was irrelevant (0.57; orientation question in the color-relevant condition). Similarly, the proportion of correctly reporting the inducers' color when it was relevant (0.86; color question in the color-relevant condition) was significantly higher ($z = 1.78$, $P < 0.04$) than when it was irrelevant (0.71; color question in the rectangle-relevant condition). Hence, this experiment replicates and generalizes our previous finding[12] in that mere irrelevance leads to strong selection.

Seeing without knowing. Most critical, however, for this study is the prediction that participants

Table 1. Frequencies of correct versus incorrect responses and their corresponding proportions (frequency/total)

	Report frequency			Report proportion		
Reported attribute	CR	ER	Total	CR	ER	Total
Inducers' color relevant (Exp. 1)	44	7	51	0.863	0.137	1
Inducers' color irrelevant (Exp. 1)	30	12	42	0.714	0.286	1
Rectangle orientation relevant (Exp. 1)	42	2	44	0.955	0.045	1
Rectangle orientation irrelevant (Exp. 1)	25	19	44	0.568	0.432	1
Object color relevant (Exp. 2)	49	2	51	0.961	0.039	1
Object color irrelevant (Exp. 2)	41	4	45	0.911	0.089	1
Object shape relevant (Exp. 2)	45	5	50	0.900	0.100	1
Object shape irrelevant (Exp. 2)	45	4	49	0.918	0.082	1

Note: In the irrelevant conditions, the total includes only those whose response to the corresponding relevant question was correct (e.g., the *n* of the color-irrelevant condition in Experiment 1 is 42 because 42 observers reported the rectangle's orientation correctly in the rectangle-relevant condition).
CR, correct; ER, incorrect.

should have little difficulty reporting the orientation of the rectangle (suggesting awareness of the inducers) but will have difficulties reporting the irrelevant inducers' color. We first examined whether the participants in the rectangle-relevant condition could report the illusory rectangle's orientation even though its inducers were deemed irrelevant. Report rate of the rectangle's orientation in this condition was very high. Only two out of the 44 participants of this condition failed to report the orientation of the rectangle (Fig. 2) correctly. Indeed, the proportion of correct responses on the rectangle's orientation question (0.96) was not significantly different from perfect performance ($P = 0.25$). Given the finding that recognition of a Kanizsa figure requires awareness of its inducers,[15] the fact that almost all of the participants in this condition could report the orientation of the illusory rectangle suggests that they were aware of the rectangle's inducers. Next, we examined their ability to report the color of the inducers. However, because we were only interested in the ability to report the irrelevant color of participants who were aware of the inducers (i.e., those who were able to report the rectangle orientation), we excluded the two participants who failed to report the orientation of the rectangle. This leaves us with 42 participants, of whom 12 could not report the color of the inducers. This report rate (0.71) was significantly lower than perfect performance ($P < 0.0001$). Still, because report rate of the inducers's color, when color was task relevant (0.86; color question in the color-relevant condition), was also

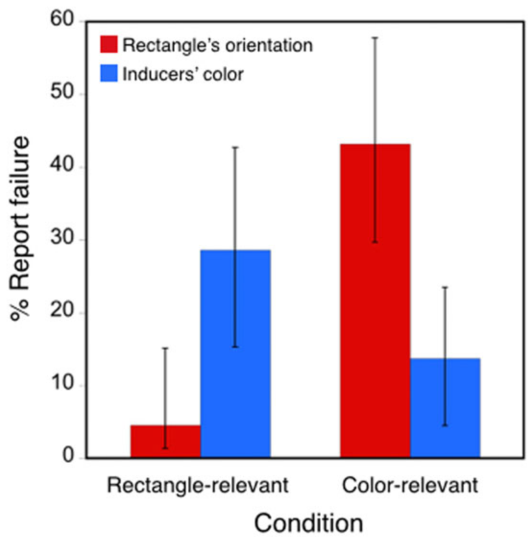

Figure 2. Erroneous recognition (percentage) in Experiment 1 as a function of the experimental condition (rectangle relevant versus color relevant) and the to-be-reported attribute (the orientation of the illusory rectangle versus the color of the inducers). The minimal recognition errors in the rectangle-relevant condition shows that people easily recognized the rectangle's orientation when it was relevant, suggesting that they were aware of the irrelevant inducers; at the same time, close to 30% of the same participants failed to report the inducers' color. This demonstrates seeing, in the sense of being phenomenally aware, without knowing, in the sense of recognizing what one saw. Recognition failure in the color-relevant condition replicates previous demonstrations of relevance-based selection: people easily recognized the inducers' color when it was relevant, but over 40% of them could not report the orientation of the irrelevant rectangle. Error bars represent 95% confidence intervals for binomial proportions.[42]

significantly lower than perfect performance ($P < 0.01$), the critical comparison is between report rate when the color was irrelevant and report rate when it was relevant. As described, this comparison was significant ($z = 1.78$, $P < 0.04$). Thus, while participants' performance for the relevant (illusory) rectangle's orientation was near perfect, their performance on the (irrelevant) color question was significantly lower. These results support our key prediction of seeing without knowing, as participants for whom the rectangle's orientation was relevant did phenomenally experience the inducers (a necessary condition for categorizing the illusory rectangle) while (many) fail to recognize their irrelevant color.

Experiment 2

As detailed above, we have demonstrated, using different stimuli and tasks,[12,13,17,18] that strong selection that is based solely on task relevance can occur even when there are no resource limitations (i.e., even when there are enough resources to process both relevant and irrelevant information). This finding was also replicated in Experiment 1. However, we know nothing about the boundary conditions of this pure relevance-based selection. The goal of this experiment was to examine one such potential boundary condition.

It is often suggested that the basic units on which attentional mechanisms operate are objects.[19–23] One of the most prominent theories associated with this statement is the integrated competition hypothesis.[24,25] According to this theory, directing attention to an object makes all its features (e.g., color, shape, orientation) available for verbal report. The theory is supported by behavioral findings demonstrating that when asked to report two different features, participants are considerably better when the two features belong to the same object than when they belong to different objects,[19,20,26] and by neurophysiological evidence of elevated neural activity for both relevant and irrelevant features of the selected object.[21,22] It seems reasonable, therefore, to assume that when the relevant and irrelevant aspects of the visual stimulation belong to the same object, pure relevance-based selection may not occur. If so, this will reveal a first boundary condition: pure relevance-based selection can only occur when the relevant and irrelevant attributes belong to different objects.

The experiments reported thus far cannot test this prediction because, in all, the display included more than one object. In our previous studies,[10,12] we presented two differently colored circles, with the larger one surrounding the smaller. Critical to the issue of objects segmentation, the participants were told that they were going to see two circles, and it was shown that for ambiguous object segmentation, informing the observers that they will see multiple objects is sufficient to ensure such segmentation.[27] Clearly then, Experiment 1 cannot test this prediction either, because its display is unambiguously segmented into five objects.[28] In Experiment 2, we tested this prediction by employing a design that is very similar to that employed in our previous study[12] apart from the fact that on each trial only an unambiguously single colored object was presented. Similar to our previous studies, to ensure minimal load levels, the object was presented for a relatively long duration (500 ms) without masking, and a simple task was employed: half of the participants had to report the color of the object and the other half had to report its shape. On the critical trial, the participants in both groups were asked to report both attributes. The key difference between this experiment and our previous studies is the fact that in this experiment only a single object is present and the relevant attribute (e.g., shape) and irrelevant attribute (e.g., color) both belong to that single unambiguous object. If relevance-based selection under minimal levels of load can take place even when relevant and irrelevant features belong to the same object, then we should replicate our previous findings: recognition rates should be lower when participants attempt to report the irrelevant feature than when they attempt to report the relevant feature. However, if relevance-based selection cannot operate on features of a single object, then similar recognition rates should be found for the relevant and irrelevant features.

Methods

Participants. One hundred and two students from the University of Haifa and Tel-Hai College participated for class credit or pay. They were randomly assigned to one of two conditions: shape relevant ($n = 51$) or color relevant ($n = 51$).

Materials and procedure. The materials and procedure were similar to Experiment 1 except for the

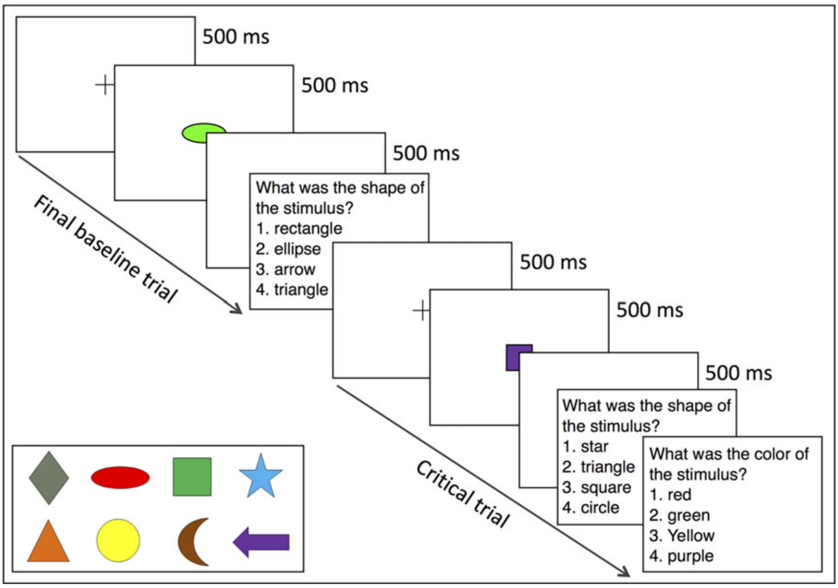

Figure 3. An illustration of the final trials in the shape-relevant condition of Experiment 2. The final baseline trial, depicted in the figure, was similar to all previous baseline trials (in the baseline trials of the color-relevant condition, the question was about the color of the stimulus). The critical trial, in which participants of both conditions were asked about both stimulus color and shape, was identical in both conditions (question order was counterbalanced). The inset depicts schematic examples of all possible shapes and colors.

following: on each trial only a single object was presented after a 500-ms fixation mark. The object had one shape out of eight possible shapes (square, triangle, circle, arrow, star, ellipse, crescent, or rhombus) and one color out of eight possible colors (red, green, yellow, blue, brown, purple, orange, or gray), both chosen randomly and independently (Fig. 3). The objects of different shapes varied somewhat in their size. The widest was the ellipse (4 cm) and the tallest were the star, crescent, and rhombus (2.8 cm). Critically, they were similar in their dimensions to the two-circles display we employed before,[12] with which relevance-based selection was observed. In the first seven trials—the baseline trials—participants were asked to report one attribute of the object by choosing one out of four options (the object's actual attribute and three other attributes alternatively and randomly chosen). Those in the shape-relevant condition were asked to indicate the object's shape, and those in the color-relevant condition were asked to indicate its color. No feedback was given. On the eighth and final trial, all participants were required to report both the color and shape of the object; the order of questions (relevant versus irrelevant) was counterbalanced between participants.

Results and discussion

One participant from the shape-relevant condition was excluded from the analysis for having more than a single error on the baseline trials (baseline accuracy $<75\%$). The following analyses were conducted on the remaining participants' performance in the critical trial (shape relevant: $n = 50$; color relevant: $n = 51$). Because for all comparisons at least one of the cells had fewer than five cases, a one-tailed Fisher's exact test was always used. Frequencies of correct versus incorrect responses and their corresponding proportions are presented in Table 1.

Testing the effect of relevance. We compared performance on the relevant and irrelevant questions of both attributes (Fig. 4). As in Experiment 1, for the irrelevant performance, we used only the data from participants who correctly answered the relevant question. Markedly different from the pattern obtained in both Experiment 1 and our previous studies,[12,18] report rate of the object's shape when it was relevant (0.9; shape question in shape-relevant

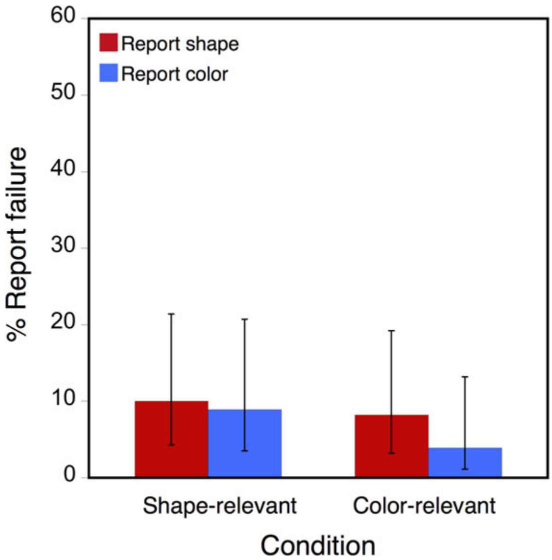

Figure 4. Erroneous recognition (percentage) in Experiment 2 as a function of the experimental condition (shape relevant versus color relevant) and the to-be-reported attribute (the object shape versus the object color). Contrary to all our previous studies, recognition errors are minimal in all conditions regardless of relevance instruction. This suggests that when both relevant and irrelevant attributes belong to a single object, pure relevance-based selection does not occur. Error bars represent 95% confidence intervals for binomial proportions.[42]

condition) was not significantly (or nominally) different from when it was irrelevant (0.92; shape question in color-relevant condition, $P = 0.51$). In the same manner, testing the report rates of the object's color yielded no significant differences between the condition in which the color was relevant (0.96; color question in color-relevant condition) and when it was irrelevant (0.91; color question in shape-relevant condition, $P = 0.28$). Thus, this is the first evidence that, at least under conditions of minimal load(s), selection on the basis of relevance cannot occur within an object.

General discussion

The primary goal of this study was to test the hypothesis that, under minimal levels of cognitive and perceptual load, irrelevant information reaches processing levels that allow phenomenal awareness but not consistent successful recognition. To that end, we relied in Experiment 1 on the findings that the conscious perception of a Kanizsa figure requires awareness of its inducers. We reasoned that if our hypothesis is correct, then when the Kanizsa figure is task relevant but its inducers are not, we should find successful recognition of the Kanizsa figure, indicating that the inducers were consciously perceived, but a considerably lower recognition rate of the inducers themselves. Indeed, we found that when the task required reporting the orientation of a Kanizsa rectangle, recognition of the rectangle's orientation was close to perfect performance. In contrast, recognition of the inducers' color was considerably worse, with close to 30% of the participants who could correctly report the orientation of the rectangle unable to report the color of the inducers. Beyond replicating our previous findings that strong relevance-based selection occurs even when capacity limitations are not met,[12] the findings also support the hypothesis that our participants were phenomenally aware of the irrelevant information but often could not report it. In more general terms, these findings demonstrate that seeing—in the sense of phenomenal experience—is possible without knowing—in the sense of recognizing what one is seeing. They further suggest that selection that is based on task relevance occurs for recognition but does not occur (or occurs to a lesser degree) for visual awareness. As detailed in the Introduction section, both the finding that task relevance is a powerful factor that mediates selection processes, and the finding that this selection affects recognition but not phenomenal awareness, are consistent with the ROAR framework.[1,8]

A possible interaction between relevance and load

It is important to note, however, that we do not argue that people are visually aware of everything. In particular, we do not argue that phenomenal awareness is always greater than its corresponding knowledge activation.[30,31] We argue only that such a situation may occur owing to the differential sensitivity of phenomenal awareness and knowledge activation to task relevance. Other factors may lead to different and opposite dissociations between knowledge activation and phenomenal awareness. Specifically, we speculate that the degree of available resources (or load levels) might be one of these factors. If we consider previous demonstrations of report failure or induced "blindness," it is evident in these demonstrations that load levels were high, likely depleting processing resources. Prominent examples are the classic inattentional blindness[32] in which participants were required to compare the lengths of subtly different lines, presented for a brief duration (200 ms) and followed by a mask—a difficult task as apparent from the high error rates (~25%[33]), and the dynamic inattentional blindness[34] that involved keeping tally of ball passes of one group and ignoring passes of another group, while the players constantly moved around. All of these cases were taken as examples of "true blindness" in the sense of a lack of phenomenal awareness. Ignoring the fact that in all these examples awareness failure was deduced from report failure, and as Experiment 1 demonstrates these do not necessarily overlap, it is plausible that under these conditions participants were indeed unaware. In comparison to these examples, this study and our previous demonstrations of irrelevance-induced blindness[12] involved a focal task that is seemingly very easy and also induced little perceptual load (i.e., few items, long exposure duration, and no mask). We argue that, under such minimal load conditions, our participants were aware of the irrelevant information. This leads us to suggest that the degree of available resources (or the level of load) interacts with task relevance to narrow visual awareness. That is, when participants are sufficiently taxed, the content of visual awareness and corresponding knowledge activation may both be limited to what is currently task relevant.[35] That is, under high load, seeing is indeed (consciously) knowing. In contrast, when load levels are low and processing resources are available for both task-relevant and task-irrelevant stimuli, only the level of knowledge activation is limited by task relevance.

Hence, like others,[36,37] our suggestion emphasizes the interplay between the two common interpretations of the term attention—as resource and as selection. Its unique contribution lies in the notion that the availability of resources does not necessarily entail lack of selection. Instead, the availability of resources determines whether selection operates on knowledge activation alone or also on visual awareness. Clearly, more research is needed to determine whether this is indeed the case. For instance, there is a need to establish whether the failure to report that was observed in the cases of inattentional blindness under high levels of load indeed also reflects failure of phenomenal experience; failure to recognize, as we found in Experiment 1; or both, as we suggest. Additionally, the manner by which resource limitation and selection interact may depend on the type of resources required for the task at hand, because when an additional memory task was added to the classical inattentional blindness paradigm, failure to report was surprisingly more severe with low than high working memory load.[38]

Relating the current results and ROAR to theories of consciousness

The results of Experiment 1 and the hypotheses inspired by ROAR also relate to influential theories of consciousness and its relation to attention. Both our predictions and results are closest to Block's theoretical differentiation between access and phenomenal consciousness.[30,31] Both our framework and that of Block's (extended by Lamme[39]) stipulate that people are phenomenally aware of more than they are able to report, and that phenomenal awareness is largely independent of pure endogenous selection. However, the reasons given for this phenomenal overflow differ. The difference is that within ROAR, it is not access consciousness (i.e., being consciously aware of what a stimulus is) that is modulated by relevance but cognitive access in general (i.e., knowledge activation). Specifically, Lamme argues that because visual awareness begins from the bottom up, attention's impact on phenomenal consciousness is minimal. Conversely, attention's impact on access consciousness is large, as it selects among competing neural assemblies at higher brain regions, thus gating entry to systems necessary for conscious reporting, among other things (mainly working memory). A consequent argument is that attention's effects on verbal report are mediated by the transient activation of unattended assemblies, which are forgotten and hence are unreportable. By ROAR, irrelevant stimulation simply fails to activate knowledge and hence is unavailable to mental processes at large. Thus, it is not the case that irrelevant knowledge is accessed and forgotten; irrelevant stimulation does not become *de facto* activated knowledge. These notions can be tested empirically, as done by Lamme; there is no reason why unattended stimuli should not lead to priming (at least for short stimulus onset asynchrony (SOA)) while ROAR precludes the possibility that irrelevant stimulation should lead to (at least semantic) priming effects.[40]

Orthogonally, the forgetting prediction is not corroborated by the data from the relevance-induced

Table 2. **Recognition rate as a function of the to-be-reported attribute, question order, and relevance condition with the corresponding relevance effect**

	Question order					
	First question			Second question		
Reported attribute	R	IR	R − IR	R	IR	R − IR
Inducers' color (Exp. 1)	1	0.84	0.16	0.73	0.61	0.12
Rectangle orientation (Exp. 1)	1	0.74	0.26	0.91	0.44	0.47
Object color (Exp. 2)	1	0.86	0.14	0.92	0.96	−0.04
Object shape (Exp. 2)	1	0.96	0.04	0.81	0.89	−0.08

R, relevant; IR, irrelevant; R – IR, relevant recognition rate − irrelevant recognition rate.

blindness studies, as pure relevance-based selection does not seem to be affected by question order. That is, if failure to report is due to forgetting (rather than nonactivation as ROAR suggests), then it should be more prominent when the question about the irrelevant object/attribute follows the question about the relevant object/attribute, simply because more time has passed between stimulation offset and report time. Our previous demonstration of pure relevance-based selection did not uncover such an effect for order; the effect of relevance was similar regardless of whether the irrelevant stimulus was reported first or second.[12] The number of participants in this study is too small to allow a reliable statistical analysis of order effects. However, inspection of the numerical values of the relevance effect (relevant recognition rate minus irrelevant recognition rate) for the various to-be-reported attributes as a function of report order does not support forgetting as the underlying mechanism (Table 2). Specifically, out of the four attributes examined in this study, a pattern consistent with forgetting appeared only for one—rectangle orientation. With all other attributes, the effect of relevance was in fact smaller for the second, compared to the first, question. Taken together, the results of our previous and current experiments do not support transient activation (forgetting) as the reason for the failure to report. Interestingly, although the order of questions did not modulate the relevance effect, it did matter, because for both relevant and irrelevant questions performance was better when a given question was first. Further research, with systematic manipulation of order, is required to better understand order's role.

A second, ostensibly relevant framework is that of partial awareness,[41] conceived to explain the dissociation between people feeling that they are phenomenally aware of the whole visual scene and clear failures to report seeing rather dramatic events. The framework's relevant tenet is that people's experience of complete awareness is illusory, as they can be conscious of only some aspects of a scene. Although the theory is seemingly successful in explaining people's self-reported visual experience, its relevance to our relatively simple experimental conditions (simple stimuli and easy task) is unclear. That is, given that in our experimental conditions capacity limitations were likely not met, how would this theory account for the fact that only some aspects of the display were consciously perceived? As such, although it is definitely possible, for example, that our participants phenomenally experienced the shape of the Kanisza inducers but not their color, it is unclear how such a prediction could be justified on the basis of the partial awareness framework.

Finally, as alluded to above, the findings of Experiment 1 also suggest that one should be cautious when relying on people's self-reports of the contents of their own awareness. This is because both verbal (e.g., "report what you have just seen") and nonverbal (e.g., "select from the following the stimulus you have just seen") manners of reporting depend on the participants having recognized the stimulus they had just seen. Experiment 1 shows that, at least under minimal load conditions, this may be an underestimation of the contents of awareness.

Relevance-based selection and objecthood

A second goal of this study was to examine whether "pure" relevance-based selection (i.e., relevance-based selection under minimal load conditions) occurs even when the relevant and irrelevant information belong to the same object. To this end, on each trial we presented a single object and asked the participants to either report its shape or color, deeming the to-be-reported attribute relevant and the other irrelevant. In the critical trial, all participants had to report both attributes. Unlike our previous findings[12] and Experiment 1, when a single unambiguous object was present, participants' report rates were similar (and high) regardless of whether the reported attribute was relevant. This finding suggests a boundary condition for "pure" relevance-based selection: it does not occur when the relevant and irrelevant information belongs to the same object. Furthermore, this finding is consistent with the assertion of the integrated competition hypothesis that the most basic unit of selection is a single object.[24,25] Accordingly, it is also consistent with various behavioral and neurophysiological studies demonstrating that the selection of one attribute of an object facilitates the processing of other attributes even if they are task irrelevant.[19–22,26] Interestingly, this boundary condition may also interact with the availability of resources, because a working-memory study[42] has demonstrated that enhanced neural activation for the irrelevant attribute of the selected object is only found with low levels of working-memory load (but see Ref. 43). Thus, further research is required to test whether relevance-based selection within a single object can occur when resource limitations are met. Finally, the fact that relevance-based selection is object based suggests that this paradigm may be useful to reveal what is considered an object by our perceptual system.[44] That is, if evidence of relevance-based selection is found for a given visual display, then it is likely that the relevant and irrelevant aspects of this display were not bound into a single object.

To conclude, this study replicates our previous demonstrations of strong selection that is based solely on task relevance, occurring even when resource limitations are not met. Critically, it demonstrates that this selection limits knowledge activation, or our ability to recognize what was the stimulus we saw, but does not equally limit phenomenal awareness. In other words, it demonstrates seeing without knowing. Additionally, this study suggests that relevance-based selection cannot occur when the relevant and irrelevant features belong to the same object, at least when load levels are minimal.

Acknowledgments

This research was supported by Grant 277/12 from the Israel Science Foundation (ISF) to B.E. We thank Nagham Ghantous, Neta Maymon, and Noam Karsh for collecting the data. All authors participated in the design, analysis, and writing of the reported research.

Conflicts of interest

The authors declare no conflicts of interest.

References

1. Eitam, B. & E.T. Higgins. 2010. Motivation in mental accessibility: relevance of a representation (ROAR) as a new framework. *Soc. Personal. Psychol. Compass* **4:** 951–967.
2. Higgins, E.T. 1996. "Knowledge activation: accessibility, applicability, and salience." In *Social Psychology: Handbook of Basic Principles.* E.T. Higgins & A.W. Kruglanski, Eds.: 133–168f. New York, NY: Guilford Press.
3. Higgins, E.T. & B. Eitam. 2014. Priming . . . Shmiming: it's about knowing when and why stimulated memory representations become active. *Soc. Cognition* **32**(Suppl.): 225–242.
4. Greenwald, A.G., S.C. Draine & R.L. Abrams. 1996. Three cognitive markers of unconscious semantic activation. *Science* **273:** 1699–1702.
5. Shapiro, K., J. Driver, R. Ward, *et al.* 1997. Priming from the attentional blink: a failure to extract visual tokens but not visual types. *Psychol. Sci.* **8:** 95–100.
6. Berti, A. & G. Rizzolatti. 1992. Visual processing without awareness: evidence from unilateral neglect. *J. Cogn. Neurosci.* **4:** 345–351.
7. Jiang, Y., P. Costello, F. Fang, *et al.* 2006. A gender and sexual orientation-dependent spatial attentional effect of invisible images. *Proc. Natl. Acad. Sci. U.S.A.* **103:** 17048–17052.
8. Eitam, B. & E.T. Higgins. 2014. What's in a goal? The role of motivational relevance in cognition and action. *Behav. Brain. Sci.* **37:** 141–142.
9. Eitam, B., Y. Schul & R.R. Hassin. 2009. Goal relevance and artificial grammar learning. *Q. J. Exp. Psychol.* **62:** 228–238.
10. Eitam, B., A. Glicksohn, R. Shoval, *et al.* 2013. Relevance-based selectivity: the case of implicit learning. *J. Exp. Psychol. Hum. Percept. Perform.* **39:** 1508–1515.
11. Eitam, B., R. Glass-Hackel, H. Aviezer, *et al.* 2014. Are task irrelevant faces unintentionally processed? Implicit learning

as a test case. *J. Exp. Psychol. Hum. Percept. Perform.* **40:** 1741–1747.
12. Eitam, B., Y. Yeshurun & K. Hassan. 2013. Blinded by irrelevance: pure irrelevance induced "blindness". *J. Exp. Psychol. Hum. Percept. Perform.* **39:** 611.
13. Kanizsa, G. 1976. Subjective contours. *Sci. Am.* **234:** 48–52.
14. Wokke, M.E., A.R.E. Vandenbroucke, H.S. Scholte, *et al.* 2013. Confuse your illusion: feedback to early visual cortex contributes to perceptual completion. *Psychol. Sci.* **24:** 63–71.
15. Harris, J.J., D.S. Schwarzkopf, C. Song, *et al.* 2011. Contextual illusions reveal the limit of unconscious visual processing. *Psychol. Sci.* **22:** 399–405.
16. Vandenbroucke, A.R.E., I.G. Sligte, J.J. Fahrenfort, *et al.* 2012. Non-attended representations are perceptual rather than unconscious in nature. *PLoS One* **7:** e50042.
17. Scott, R.B. & Z. Dienes. 2008. The conscious, the unconscious, and familiarity. *J. Exp. Psychol. Learn. Mem. Cogn.* **34:** 1264–1288.
18. Tapal, A., B. Eitam & Y. Yeshurun. In preparation. The role of suppression in Irrelevance induced blindness.
19. Blaser, E., Z.W. Pylyshyn & A.O. Holcombe. 2000. Tracking an object through feature space. *Nature* **408:** 196–199.
20. Duncan, J. 1984. Selective attention and the organization of visual information. *J. Exp. Psychol. Gen.* **113:** 501–517.
21. O'Craven, K.M., P.E. Downing & N. Kanwisher. 1999. fMRI evidence for objects as the units of attentional selection. *Nature* **401:** 584–587.
22. Schoenfeld, M.A., C. Tempelmann, A. Martinez, *et al.* 2003. Dynamics of feature binding during object-selective attention. *Proc. Natl. Acad. Sci. U.S.A.* **100:** 11806–11811.
23. Wannig, A., V. Rodriguez & W.A. Freiwald. 2007. Attention to surfaces modulates motion processing in extrastriate area MT. *Neuron* **54:** 639–651.
24. Duncan, J. 1998. Converging levels of analysis in the cognitive neuroscience of visual attention. *Phil. Trans. R. Soc. Lond. B. Biol. Sci.* **353:** 1307–1317.
25. Duncan, J., G. Humphreys & R. Ward. 1997. Competitive brain activity in visual attention. *Curr. Opin. Neurol.* **7:** 255–261.
26. Duncan, J. 1993. Coordination of what and where in visual attention. *Perception* **22:** 1261–1270.
27. Chen, Z. 1998. Switching attention within and between objects: the role of subjective organization. *Can. J. Exp. Psychol.* **52:** 7:adi
28. Kanizsa, G. 1979. *Organization in Vision.* New York, NY: Praeger.
29. Kellman, P.J. & T.F. Shipley. 1991. A theory of visual interpolation in object perception. *Cognitive Psychol.* **23:** 141–221.
30. Block, N. 2007. Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behav. Brain. Sci.* **30:** 481–499.
31. Block, N. 2011. Perceptual consciousness overflows cognitive access. *Trends. Cogn. Sci.* **15:** 567–575.
32. Mack, A. & I. Rock. 1998. *Inattentional Blindness.* The MIT Press. Cambridge, MA.
33. Rock, I., C.M. Linnett, P. Grant, *et al.* 1992. Perception without attention: results of a new method. *Cogn. Psychol.* **24:** 502–534.
34. Simons, D.J. & C.F. Chabris. 1999. Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception* **28:** 1059–1074.
35. Rees, G., C. Russell, C.D. Frith, *et al.* 1999. Inattentional blindness versus inattentional amnesia for fixated but ignored words. *Science* **286:** 2504–2507.
36. Kahneman, D. 1973. *Attention and Effort.* Englewood Cliffs, NJ: Prentice-Hall.
37. Lavie, N. 1995. Perceptual load as a necessary condition for selective attention. *J. Exp. Psychol. Hum. Percept. Perform.* **21:** 451–468.
38. de Fockert, J.W. & A.J. Bremner. 2011. Release of inattentional blindness by high working memory load: elucidating the relationship between working memory and selective attention. *Cognition* **121:** 400–408.
39. Lamme, V.A. 2003. Why visual attention and awareness are different. *Trends Cogn. Sci.* **7:** 12–18.
40. Kreitz, C., R. Schnuerch, P.A. Furley, H. Gibbons & D. Memmert. 2014. Does semantic preactivation reduce inattentional blindness? *Attention, Perception, & Psychophysics*, (Epub ahead of print)
41. Kouider, S., V. De Gardelle, J. Sackur & E. Dupoux. 2010. How rich is consciousness? The partial awareness hypothesis. *Trends Cogn. Sci.* **14**, 301–307.
42. Xu, Y. 2010. The neural fate of task-irrelevant features in object-based processing. *J. Neurosci.* **30:** 14020–14028.
43. Yin, J., J. Zhou, H. Xu, *et al.* 2012. Does high memory load kick task-irrelevant information out of visual working memory? *Psychon. Bull. Rev.* **19:** 218–224.
44. Scholl, B.J. 2001. Objects and attention: the state of the art. *Cognition* **80:** 1–46.