ASSOCIATION FOR
PSYCHOLOGICAL SCIENCE

Research Article

Psychological Science
2020, Vol. 31(6) 663–677

DOI: 10.1177/0956797620915887
www.psychologicalscience.org/PS
SAGE

B or 13? Unconscious Top-Down Contextual Effects at the Categorical but Not the Lexical Level

Dan Biderman[1,2], Yarden Shir[1], and Liad Mudrik[1,3]
[1]School of Psychological Sciences, Tel Aviv University; [2]Department of Neuroscience and the Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University; and [3]Sagol School of Neuroscience, Tel Aviv University

Abstract
Contextual effects require integration of top-down predictions and bottom-up visual information. Given the widely assumed link between integration and consciousness, we asked whether contextual effects require consciousness. In two experiments (total $N = 60$), an ambiguous stimulus (which could be read as either B or 13) was presented alongside masked numbers (12 and 14) or letters (A and C). Context biased stimulus classification when it was consciously and unconsciously perceived. However, unconsciously perceived contexts evoked smaller effects. This finding was replicated and generalized into another language in a further experiment ($N = 46$) using a different set of stimuli, strengthening the claim that symbolic contextual effects can occur without awareness. Moreover, four experiments (total $N = 160$) suggested that these unconscious effects might be limited to the categorical level (numbers context vs. letters context) and do not extend to the lexical level (words context vs. nonwords context). Taken together, our results suggest that although consciousness may not be necessary for effects that require simple integration or none at all, it is nevertheless required for integration over larger semantic windows.

Keywords
consciousness, top-down, perception, ambiguous objects, Bayesian inference, integration, open data, open materials, preregistered

Received 7/26/19; Revision accepted 1/10/20

How do we perceive and understand visual objects? One hypothesis is that the brain hierarchically builds objects from their lower level visual features in a bottom-up manner (e.g., Hubel & Wiesel, 1962; Riesenhuber & Poggio, 1999). This process is held to be mediated by a set of interconnected cortical areas both within (Grill-Spector, Kourtzi, & Kanwisher, 2001) and outside (Georgieva, Peeters, Kolster, Todd, & Orban, 2009) the ventral visual stream. Such bottom-up vision models also enjoy prominence in the field of computer vision and computational neuroscience (Kriegeskorte, 2015).

Yet a growing body of literature suggests that higher level factors—such as context, motivation, language, and emotional state—influence perception in a top-down manner (e.g., Balcetis, 2016; Bar, 2004; but see Firestone & Scholl, 2016). A classic example is the perception of ambiguous objects in different contexts. Bruner and Minturn's (1955) famous B/13 object is judged as 13 when embedded between 12 and 14 but as B when between A and C. Thus, subjects' context-based expectations are integrated with the visual features of the ambiguous object (Panichello, Cheung, & Bar, 2013). Such contextual effects are presumably supported by feedback to visual areas from higher level cortices (Bar, 2004).

Do such context effects depend on consciousness? Arguably, one might suspect that they do because they involve integration of top-down and bottom-up processes; the close tie between consciousness and integration has

---

Corresponding Author:
Liad Mudrik, Tel Aviv University, School of Psychological Sciences and Sagol School of Neuroscience, Ramat Aviv, P.O. Box 39040, Tel Aviv 69978, Israel
E-mail: mudrikli@tauex.tau.ac.il

been advocated since the times of Descartes, Kant, and James (Mudrik, Faivre, & Koch, 2014). Prominent theories of consciousness assign a crucial role to integration (e.g., integrated-information theory; Tononi, Boly, Massimini, & Koch, 2016) or explicitly predict that global integration across brain regions should occur only during conscious processing (e.g., global-neuronal-workspace theory; Baars, 1997; Dehaene & Naccache, 2001). This topic has attracted substantial scientific attention, with repeated attempts to examine whether humans can integrate stimuli that they do not consciously perceive (Moors, Hesselmann, Wagemans, & van Ee, 2017; Sklar, Deouell, & Hassin, 2018). Whereas some findings implied that they can, even when high-level integration is involved (e.g., Fahrenfort, van Leeuwen, Olivers, & Hogendoorn, 2017; Sklar et al., 2018), other researchers attributed these effects to lower level factors or failed to replicate them (Moors & Hesselmann, 2018; Moors, Wagemans, Van Ee, & de-Wit, 2016; Rabagliati, Robertson, & Carmel, 2018). Thus, the scope of unconscious integration is still a matter of ongoing controversy and requires further scrutiny. One suggestion is that although some unconscious integration can be obtained without awareness, it is nevertheless limited. According to the windows-of-integration hypothesis (Mudrik et al., 2014), integration can be defined over different window sizes (in temporal, spatial, or semantic spaces), and consciousness might be required only for integrating over bigger windows.

Here, we focused on the possible role of consciousness in contextual effects, which are held to involve integration between top-down semantic context and bottom-up visual information. Thus far, such effects have mostly been shown using fully visible contexts and objects (for a review and criticism, see Firestone & Scholl, 2016). The few studies using subliminal contextual cues focused on low-level contextual visual illusions (e.g., tilt: Clifford & Harris, 2005; Mareschal & Clifford, 2012; brightness: Harris, Schwarzkopf, Song, Bahrami, & Rees, 2011; size contrast: Nakashima & Sugita, 2018). Conflicting findings were found for higher level contextual effects, such as the Kanizsa illusion (Kanizsa, 1979), in which a specific spatial organization of separate cues evokes the experience of an illusory surface region: Some researchers found that it could be induced by invisible cues (e.g., Fahrenfort et al., 2017; Persuh, Emmanouil, & Ro, 2016; Wang, Weng, & He, 2012), yet others suggested that these effects might be driven by low-level properties of the stimulus (Moors et al., 2016). Moreover, others did not find such an effect at all (e.g., Banica & Schwarzkopf, 2016; Harris et al., 2011), leaving the matter of unconscious contextual effects unresolved.

Accordingly, we asked whether visible and invisible semantic contexts—letters, numbers, and words—can bias the classification of a clearly visible symbolic ambiguous object. In two experiments (preregistration: osf.io/5dpvs), we found such effects on the B/13 stimulus (Bruner & Minturn, 1955). In another experiment, this finding was replicated and further generalized into a different language using an ambiguous Hebrew letter (/tet/) that could also be perceived as "6" (preregistration: osf.io/5jnkg). Thus, the effect of context on ambiguous stimuli that could be classified as either letters or digits was demonstrated in three different experiments, across two languages. In four additional experiments (preregistrations: osf.io/3ccv3, osf.io/hvh8f, osf.io/kzgtr, osf.io/5jnkg), we asked whether this effect extends also to the lexical domain, which requires a finer grained processing of words than of nonwords or, put differently, integration over bigger semantic windows. We accordingly presented participants with ambiguous Hebrew letters embedded in either familiar three-letter words or nonwords (Eldar, Niv, & Cohen, 2016) and found that, as opposed to categorical contextual effects, lexical effects were evoked only by conscious inducers and not during unconscious ones. Data and materials from all experiments, along with data-analysis and experimental scripts, are available at osf.io/7236z.

Method

The methodology below pertains to all experiments included in this study. Because the general methodology was similar, most changes between experiments relate to (a) manipulating participants' conscious perception of the stimuli, (b) the stimuli themselves, and (c) sample size. Below, we first outline the common methodology across experiments and then describe experiment-specific details.

Participants

Overall, 152 participants took part in this study for course credit or payment (~$10): There were 12 participants in Experiment 1a (3 women; 11 right-handed; age: *M* = 25.91 years, *SD* = 2.02), 24 in Experiment 1b (15 women; 20 right-handed; age: *M* = 24.52 years, *SD* = 1.73), 24 in Experiment 2 (18 women; 21 right-handed; age: *M* = 24.62 years, *SD* = 3.13), 12 in Experiment 3a (11 women; 8 right-handed; age: *M* = 23.17 years, *SD* = 1.99; all native Hebrew speakers), 34 in Experiment 3b (31 women; 31 right-handed; age: *M* = 23.06 years, *SD* = 1.46; all native Hebrew speakers), 12 in Experiment 4a (8 women; 11 right-handed; age: *M* = 23.58 years, *SD* = 1.83), and 34 in Experiment 4b (29 women; 31 right-handed; age: *M* = 23.06 years, *SD* = 1.52). Experiment 1 was exploratory, whereas Experiments 2 to 4 were all preregistered. In the first two experiments, we chose to double the sample size for the unconscious-processing group (Experiments 1b and 2), given that the unconscious effects are commonly weaker (Greenwald, Draine, & Abrams, 1996). From Experiment 3 onward, sample size was determined so as to reach a power of .9 for the context effect observed in Experiment 2 on the basis of a simplified, no-pooling version of the logistic regression model (we fitted a logistic function for each of the participants on noncentered data and then tested the mean context coefficient against zero), $t(23) = 2.82$, $p < .01$; Cohen's $d = 0.58$. In all experiments, data collection stopped when we reached the predefined sample size.

In all experiments, additional participants were excluded if they met at least one of the predefined exclusion criteria: (a) too few visibility-1 trials (see below) in each experimental cell in the main session or in the posttest session, (b) high performance in the posttest (mean $d' > 1.5$ or mean accuracy > 65%), or (c) mean reaction time (RT) in the main session deviating by more than 3 standard deviations from the group mean. Overall, 9 participants were excluded in the exploratory Experiment 1b (4 because of criterion a, 5 because of criterion b), and 21 participants were excluded in the preregistered Experiments 2 to 4 (Experiment 2: 3 because of criterion a, 3 because of criterion b; Experiment 3b: 2 because of criterion a, 1 because of criterion c; Experiment 4a: 3 because of criterion a; Experiment 4b: 8 because of criterion a, 1 for not showing up for the second session of the experiment). Notably, we excluded these participants prior to analyzing the effects of interest, and no outliers were further excluded on the basis of participants' effect-related result. For all experiments, the results remained unchanged when we included all participants in the analyses. And, importantly, our effects were found for categorical but not lexical contexts, even when probed in the same participants, for whom regression to the mean presumably should have had an equal effect. Therefore, our findings are less likely to be driven by regression to the mean resulting from extreme group analysis (Shanks, 2017). Exclusion criteria and analysis plans were preregistered on the Open Science Framework (for deviations from analysis plans, see Supplemental Note and Tables S1 and S2 in the Supplemental Material available online).

Apparatus and stimuli

Stimuli were presented on an LCD monitor (23-in. ASUS SyncMaster; 1,920 × 1,080 resolution; 60-Hz refresh rate) using MATLAB (The MathWorks, Natick, MA) and the Psychophysics Toolbox Version 3 (Brainard, 1997). Participants sat in a dimly lit room, and their heads were stabilized using a chin rest located 60 cm from the screen.

Stimuli varied between the experiments, but all were built using the same logic (see Fig. 1 and osf.io/7236z): First, we created an ambiguous stimulus that could be classified in one of two ways (Experiments 1, 2, and 4: B/13; Experiments 3 and 4: ambiguous letters; see details below). Then, different versions of that stimulus were made by manipulating its physical properties to make it more biased toward each of the two possible classifications in a bottom-up manner. In addition, we selected context inducers that are likely to affect classification in a top-down manner. Finally, we constructed a mask display by placing two masks in the same location as the context inducers, separated by a central blank space (in the location of the ambiguous stimulus). Each mask was composed of various superimposed keyboard symbols (Experiments 1, 2, and 4) or letters (Experiments 3 and 4). All stimuli were colored white and presented on a gray background (red, green, blue, or RGB values: 128, 128, 128).

Experiments 1 and 2 (categorical context). In Experiments 1 and 2, which probed categorical effects on the classification of the B/13 stimulus, the central ambiguous object was composed of two parts—"1" and "3"—so that it could be interpreted as either the letter "B" or the number "13." The distance between these two parts was manipulated (ranging between ~0.03° and ~0.18° of visual angle in intervals of ~0.03°) to create six versions of the central object, all still somewhat ambiguous yet with the shortest distance more resembling B and the longest distance more resembling 13 (see Fig. 1a). The context inducers were the letters "A" and "C," the numbers "12" and "14," and the signs "@" and "%" (in Experiment 1 only), which were presented on either side of the central object (with A, 12, and @ to the left of the ambiguous stimulus and C, 14, and % to its right). The entire stimulus ensemble (context inducers and ambiguous stimulus) subtended 5.7° to 6.2° (width, depending on the version of the central object and the context inducers) and 2.4° (height) of the visual field. The composite mask was of the same size.

Experiment 3 (lexical context). In Experiment 3, which focused on lexical effects on the classification of an ambiguous letter, the central stimulus was an ambiguous hybrid of two Hebrew letters, created in Adobe Photoshop and Adobe Illustrator by a professional designer. The adjacent letters formed a three-letter word in which only one of the two possible interpretations of the ambiguous stimulus was possible, so that a pseudoword was formed with the second letter (congruent and incongruent conditions with respect to one of the two letters). The following example provides an English parallel. An ambiguous letter can be generated from the letters A and H and the context can be induced by T-E or C-T (Eldar et al., 2016). The new three-letter strings can form either a word (THE or CAT) or a nonword (TAE or CHT; note that in

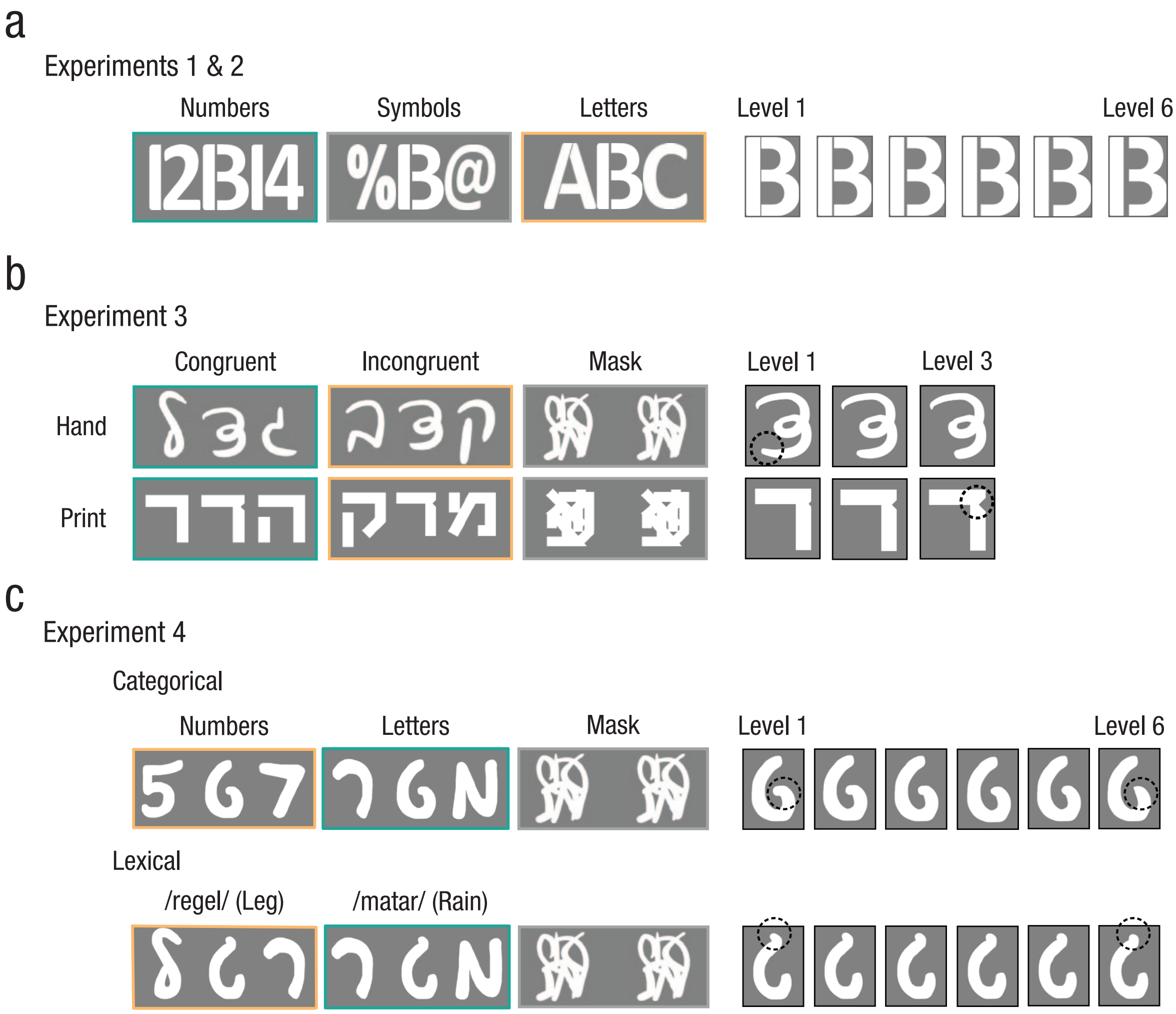

Fig. 1. Stimuli in all experiments. The stimuli in Experiments 1 and 2 (a, left) consisted of an ambiguous object, which could be interpreted as either B or 13, and context inducers, which flanked the ambiguous object. The symbols context inducers were used in Experiment 1 only. The numbers and letters context inducers were used in Experiments 1 and 2. Colored borders indicate whether the context was congruent with classification of the object as the number 13 (blue) or incongruent with it (red). Distance between the parts of the ambiguous stimulus (a, right) was increased in equal steps to create six levels of distance. Example congruent (Hebrew word) and incongruent (Hebrew pseudoword) stimuli used in Experiment 3 (b) are shown separately for trials in which the ambiguous letter appeared in handwritten style and print style. Associated masks are shown in the middle column. Three distance versions of an ambiguous object in handwritten style and print style are shown on the right. Dashed circles denote the loci of physical changes. Example stimuli for trials with number and letter context inducers in Experiment 4 (c) are shown separately for the categorical and lexical conditions. In the categorical condition, the ambiguous object ranged from the number 6 to the Hebrew letter ט /tet/, forming the sequence 5-6-7 or the Hebrew word /matar/ (rain), respectively. In the lexical condition, the ambiguous object ranged from the Hebrew letter ג (/gimel/) to ט (/tet/), forming Hebrew words that visually resemble the categorical stimuli. Associated masks are shown in the middle column. Six distance versions of an ambiguous object are shown at the right for each of the conditions. The dashed circles mark the changes in the stimuli.

our case, and as opposed to CHT, these were pseudowords and not nonwords; that is, they always satisfied the orthophonological constrains of the Hebrew language). Three versions of the ambiguous stimulus were created to make it more similar to each of the two letters it could be classified as (see Fig. 1b). The context inducers were chosen so that they formed a word with only one of the two letters and a pseudoword with the other. Presented together, the string composed of the ambiguous letter and the context inducers subtended a maximum of 6.0° (width) and 2.3° (height) of the visual field.

The stimuli (object and inducers) were presented in two different writing styles, which appeared in separate blocks: Seven stimuli were in print style (an adaptation of Haim Hebrew font) and seven in handwritten style (written manually by the designer using a graphics pad).

Some stimuli appeared in both writing styles, yet the context inducers were unique to each letter pair. For each stimulus, we arbitrarily defined a reference letter, in relation to which we term the context as either congruent or incongruent. The following example provides an English parallel. For the ambiguous letter A/H, when A is chosen as the reference letter, C-T would be the resulting congruent context (forming the congruent word CAT) and T-E the incongruent one (forming the incongruent nonword TAE). Thus, we had two sets of words, congruent and incongruent, defined according to the randomly chosen reference letter for each pair. The congruent-words set contained 10 nouns and two verbs, whereas the incongruent-words set included 8 nouns, one verb, one adverb, and two adjectives. The two sets of words were comparable in word frequency and concreteness measures. We also validated that there were no consistent visual differences between the two sets of words, which appeared in the same font, size, and colors and comprised an equivalent number of segments (for details and analyses, see the Supplemental Material).

Experiment 3a (conscious processing) also served as a pretest for the stimuli used in Experiment 3b (unconscious processing). There were two stimuli that generated weak context effects in Experiment 3a (defined by calculating the proportion of classification as the reference letter in congruent and incongruent conditions and then subtracting the latter from the former; difference in proportion below .1 was defined as a reason to exclude the stimuli). These stimuli were accordingly excluded from Experiment 3b. To allow a clearer comparison between Experiment 3a and Experiment 3b, we omitted from analysis all trials in which these pairs were presented in Experiment 3a (14.42% of the trials; results remained unchanged when all trials were included).

Experiment 4 (categorical and lexical contexts). Finally, in Experiment 4, which probed both categorical and lexical processing, two stimulus sets (each including a central ambiguous letter, two contexts, and masks) were created. The first probed categorical contexts, as in Experiments 1 and 2, and the second probed lexical contexts, similarly to Experiment 3. For the former, the critical ambiguous stimulus could look like the number 6 and like the handwritten Hebrew letter ט (/tet/; see Fig. 1c). Six versions of the sign were created, ranging in resemblance from 6 to /tet/. Context inducers were either the digits 5 and 7 or the Hebrew letters מ (/mem/) and ר (/reish/), also forming the Hebrew word מטר (/matar/, the noun *rain*) with the ambiguous sign perceived as /tet/.

In the lexical stimulus set, the critical ambiguous stimulus could again look like the Hebrew letter ט (/tet/) but also the letter ג (/gimel/). Six versions of the ambiguous letter were created, ranging between the two letters (see Fig. 1c). Congruent and incongruent words were also created: To keep the stimuli similar between conditions, we again used the Hebrew letters מ (/mem/) and ר (/reish/) with the ambiguous sign perceived as ט (/tet/). And the Hebrew letters ר (/reish/) and ל (/lamed/) formed the Hebrew word רגל (/regel/, the noun *leg*) with the ambiguous sign perceived as ג (/gimel/).

Procedure

Experimental sessions. Experiments 1a, 3a, and 4a probed conscious processing and, accordingly, consisted of one session only. To make sure that the context inducers were indeed unconsciously processed in Experiments 1b, 2, 3b, and 4b, we added an additional posttest session that followed the main experiment. In this session, we obtained an objective visibility measure for each participant. Finally, in Experiment 4b, we wanted to compare lexical and categorical effects within the same participants. Thus, participants came to the lab twice (0–3 days apart): once to conduct the categorical task and again to conduct the lexical task (each included a main session and a posttest; task order was counterbalanced). The number of trials in each session of each experiment is detailed in the Supplemental Material.

Sequence of events. The trial sequence in Experiments 1a, 3a, 4a, which probed conscious processing, and in Experiments 1b, 2, 3b, 4b, which probed unconscious processing, included the exact same events, but their order was changed to manipulate participants' awareness of the context inducers (see Figs. 2a and 2b). Each trial began with a 500- to 700-ms fixation cross. Then, the composite mask (50 ms) and a blank screen (100 ms) appeared. In experiments probing conscious processing, the composite mask was presented first and the blank screen followed, and in experiments probing unconscious processing, the order was reversed: the blank screen first and the mask second. The ambiguous object and the context inducers then appeared together for 33 ms and were followed by another blank screen (100 ms) and mask (50 ms): In experiments probing conscious processing, the blank screen appeared first, and in experiments probing unconscious processing, the mask preceded the blank screen. Thus, in the former case, the context inducers were separated from the mask by blank screens, so they were visible, whereas in the latter case, they were immediately preceded and followed by the masks, rendering them invisible. Note that because the masks did not extend to the space in which the ambiguous stimulus appeared, the ambiguous stimulus was visible in both experiments: Only the contextual inducers were masked. In experiments including a posttest session, the sequence of events in the posttest was exactly the same as the one used in the main session.

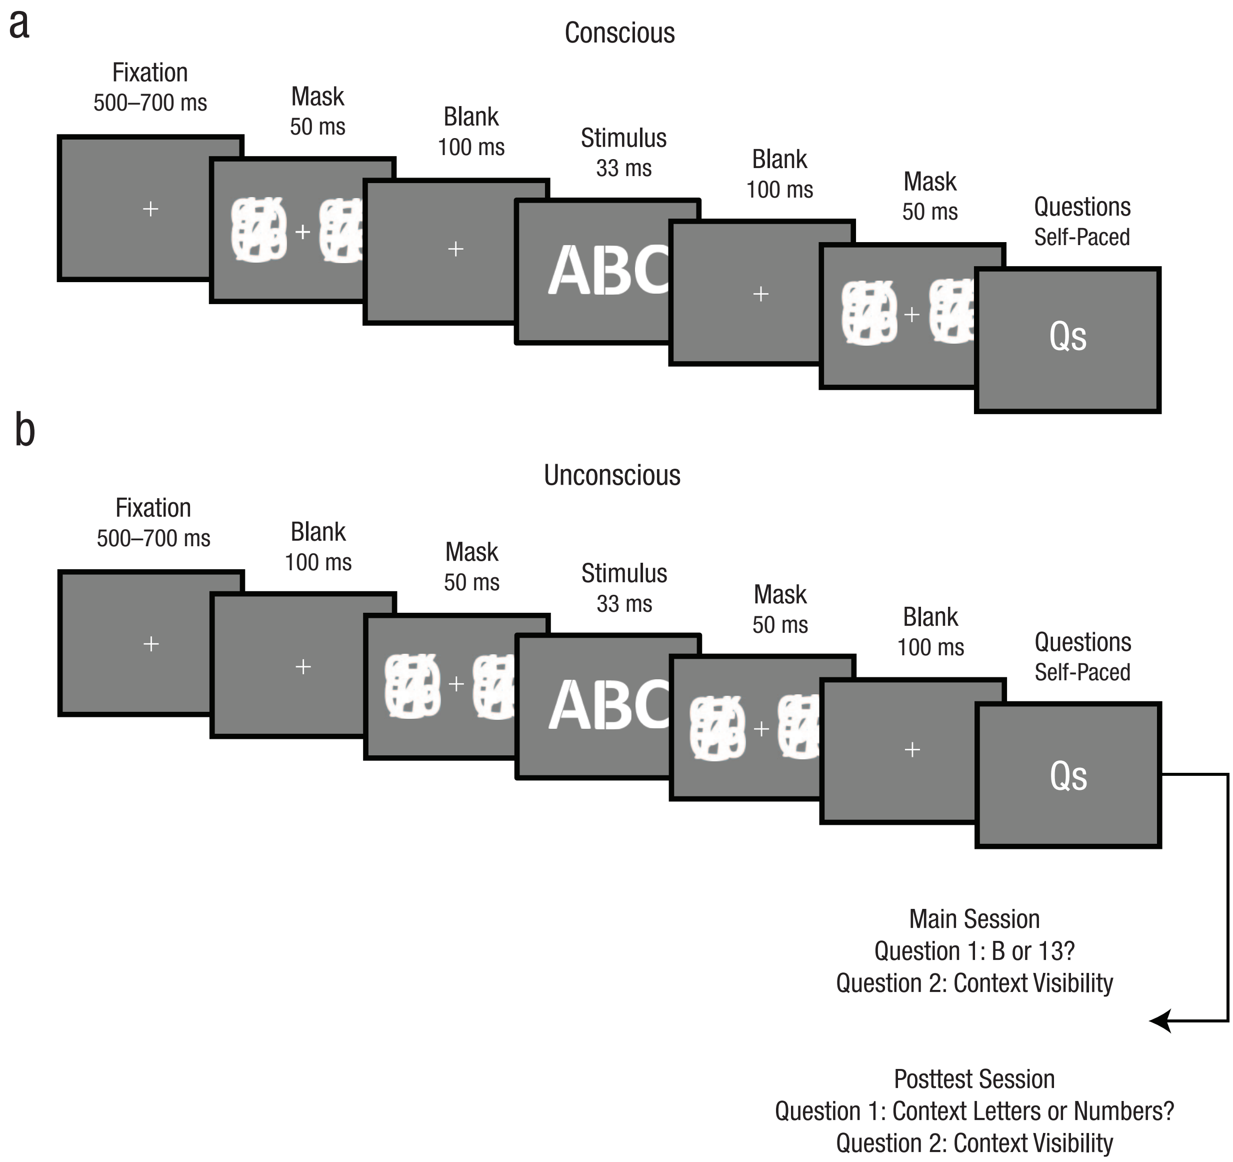

Fig. 2. Sequence of events in Experiments 1a, 1b, and 2. In the conscious sequence used in Experiment 1a (a), the critical stimulus display (here, B/13 flanked by A and C) was preceded and followed by blank screens; the ambiguous object (which could be interpreted as either B or 13) and the context inducers flanking it (A and C) were thus fully visible. In the unconscious sequence used in Experiments 1b and 2 (b), the critical stimulus display was sandwiched between two masks, so the context inducers, but not the ambiguous object, were rendered invisible. The experimental sequence in both conditions was followed by two questions (Qs). In the main session, participants were first asked whether the ambiguous object was B or 13 and then asked to rank the visibility of the context inducers on a 4-point scale. In the posttest, which followed the main session of Experiments 1b and 2, the first question asked whether the context inducers were letters or numbers, and the second asked about context visibility. (The same trial structure was used in subsequent experiments but with different stimuli).

Experimental tasks. In all experiments, the primary task in the main session was to classify the ambiguous stimulus as one of its two possible interpretations by pressing either the left or the right arrow keys. Question onset was indexed by a question mark in Experiments 1, 2, and 4 (categorical task). In Experiments 3 and 4 (lexical task), the question was presented by introducing the two letters from which the ambiguous stimulus was created, separated by a hyphen (e.g., for the English example, A-H; letter order was counterbalanced), in their corresponding writing style (print: Arial font; handwritten: Tamir font; both font sizes = 70). Then, participants were instructed to report how well they perceived the context inducers (subjective measure). Ratings were made on a 4-point scale

(Perceptual Awareness Scale; Ramsøy & Overgaard, 2004) from 1 (“I saw nothing”), 2 (“I had a brief glimpse of the inducers”), and 3 (“I saw the inducers almost clearly”), to 4 (“I saw the inducers clearly”) using the numeric keys 1 to 4.

In the posttest session (Experiments 1b, 2, 3b, and 4b), on the other hand, participants’ primary task was to discriminate between two context inducers and then rank their visibility on the same 4-point scale. In Experiments 1, 2, and 4 (categorical task), the task was to determine whether the inducers were letters or numbers. In Experiments 3 and 4 (lexical task), they were asked to determine which of the two possible pairs of context inducers for the current ambiguous letter were presented; the two pairs were separated by a slash, whereas the two letters inside each letter pair were separated by a hyphen (e.g., C-T/T-E).

Analysis

Trial-exclusion criteria. Trials that met at least one of the following predefined criteria were excluded from the analyses: (a) Participants erroneously pressed a nonassigned button, (b) RTs were shorter than 250 ms or longer than 4 s, and (c) RTs deviated by 3 standard deviations or more from the mean for each experimental condition (for number of excluded trials in each experiment, see Supplemental Note in the Supplemental Material). After this initial clean-up, we excluded trials on the basis of context-visibility ratings. Namely, in Experiments 1a, 3a, and 4a, in which we aimed to assess conscious processing of the context inducers, we used only two to four visibility trials, in which participants reported seeing something on a scale ranging from *vaguely* (2) to *clearly* (4). In Experiments 1b, 2, 3b, 4b, Control Experiment 1, and Control Experiment 2, all of which tapped unconscious processing of the context inducers, we used only visibility-1 trials, in which participants reported not seeing the context inducers at all (for a breakdown of visibility ratings for each experiment, see Table S3 in the Supplemental Material). Note that for all experiments, we excluded relatively few trials because of visibility reports and found that the results remained unchanged when we included trials from all visibility ratings (Hesselmann, Darcy, Ludwig, & Sterzer, 2016), alleviating the concern that our findings were influenced by post hoc selection of trials (Shanks, 2017).

Multilevel Bayesian modeling. The results were analyzed taking a Bayesian approach (Gelman et al., 2013; Kruschke, 2014; for a detailed explanation and rationale, see Supplemental Note in the Supplemental Material). In a nutshell, we modeled participants’ responses using hierarchical regression models that included participant-specific and group-level coefficients. We approximated a posterior distribution over these coefficients using Markov chain Monte Carlo methods. To interpret the results, we computed the median and the 95% highest-density interval (HDI) of the posterior distribution of each group-level regression coefficient, under the assumption that for a reliable effect, the HDI would exclude zero (Gelman et al., 2013; Kruschke, 2014). The estimated parameters included the independent variables—context, which was either categorical (Experiment 1: numbers, letters, and symbols; Experiments 2 and 4: numbers and letters) or lexical (congruent, incongruent)—and distance (continuous; Levels 1–6 in Experiments 1, 2, and 4 and Levels 1–3 in Experiment 3) and their interactions.

In all experiments, the dependent variable was participants’ binary classifications of the ambiguous object as one of the two possible interpretations. In Experiments 1, 2, and 4 (categorical task), the ambiguous object could be classified as either a number or a letter, and in Experiments 3 and 4 (lexical task) as one of two possible letters (a reference letter). Finally, we further conducted an analysis on RTs that was of secondary importance and did not yield any meaningful differences between conditions (see Tables S4–S8 in the Supplemental Material).

Results

Visibility of contextual inducers

Subjective and objective visibility measures confirmed that participants were indeed unaware of the contextual inducers in Experiments 1b, 2, 3b, and 4b. This was confirmed also by a meta-analysis across experiments and further corroborated by the lack of correlation between the observed effect and participants’ awareness scores, as defined by their performance on the objective task (see Fig. S1 in the Supplemental Material). All results, as well as group-level coefficients for classification analysis (see Tables S9–S14), are reported in the Supplemental Material.

Experiments 1 and 2: categorical context effects on B/13 classification

Experiment 1a (conscious processing). As expected, the distance between the 1 and 3 components of the B/13 object influenced its classification. That is, participants increasingly tended to classify the object as 13 when the distance between its 1 and 3 components increased. Letters context and numbers context influenced participants’ classification in the expected direction: Participants classified the object more as 13 when it appeared in a numbers context (12-14) compared with both the baseline symbol context (%-@) and the letters

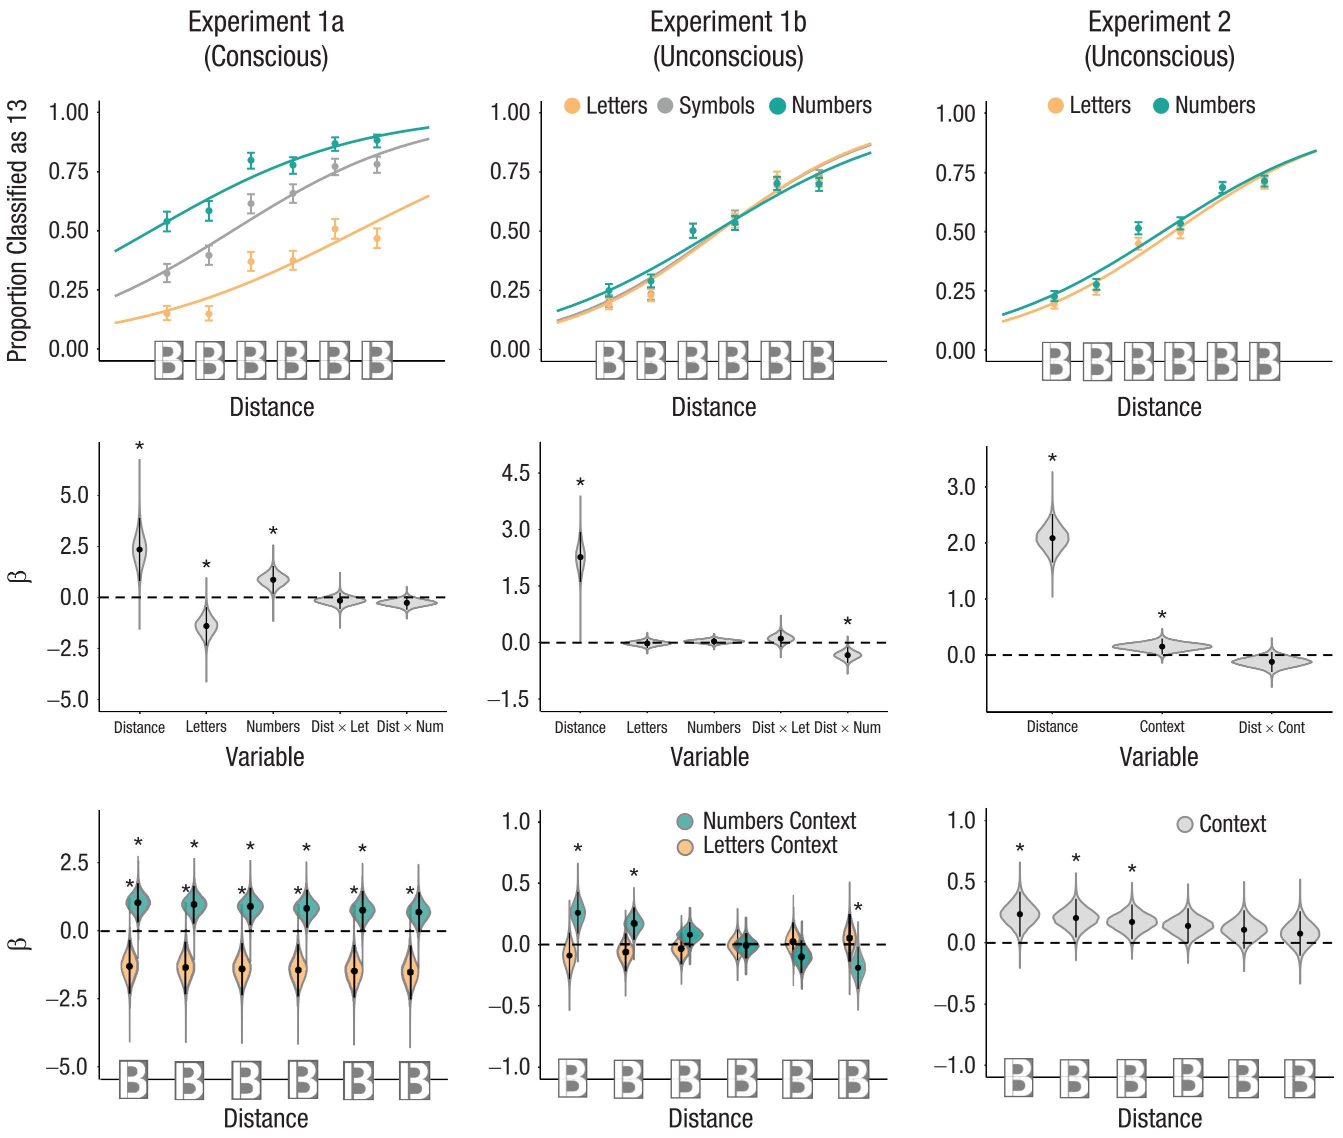

Fig. 3. Results of Experiments 1 and 2. The top row shows mean aggregate classifications of the ambiguous stimulus as 13 (dots) and psychometric-function fits as a function of distance between the two parts of the ambiguous stimulus and whether the context inducers were letters, symbols, or numbers (lines). Error bars denote Wilson 95% binomial confidence intervals. The middle row shows group-level posterior coefficient estimates for the key variables in each experiment (estimates were obtained from a multilevel logistic regression model). The bottom row shows posterior coefficients for numbers and letters for each distance. In the middle and bottom rows, dots indicate medians, and the width of the violin plots denotes the probability density in the full range of the posterior (i.e., depicting every Markov chain Monte Carlo sample of the posterior). Error bars denote 95% highest-density intervals (HDIs). The dashed lines indicate the point at which the regression coefficient equals 0 and is therefore not predictive of the outcome variable. Effects for which the 95% HDI excluded 0 are marked with an asterisk. Model predictors were standardized so the regression coefficients would all lie on the same scale.

context (A-C); classifications as 13 were lowest in the letters context. Both interaction terms (Distance × Letters, Distance × Numbers) were negligible. Calculation of the context effect in each distance level showed that the 95% HDI of the letters context was below 0 in all six distances, and for the numbers context, it was reliably positive in all distances but the last (see Fig. 3; see also Table S9).

Experiment 1b (unconscious processing). In Experiment 1b as well, the distance credibly influenced participants' classifications: There were more "13" classifications as the distance grew. As for the masked context inducers, only a negative interaction between distance and numbers context was reliable. Analysis of the numbers-context coefficient separately for each distance level revealed reliable positive effects for the first two distances, in which participants were affected by the numbers context and classified the stimulus more as 13 (see Fig. 3; see also Table S9). An unexpected effect in the opposite direction was found in the last distance (i.e., tendency to choose B when the numbers context was presented). To ensure that these effects were not driven by the neutral condition (because we compared the neutral condition with either the letters or the numbers condition), we conducted an additional analysis in which we excluded the neutral-symbols trials and ran a model with distance, context (letters = 0, numbers = 1), and their interaction as predictors. The results remained unchanged.

Experiment 2 (unconscious processing). Experiment 1b, which was exploratory, revealed an evidently weak unconscious contextual effect. Experiment 2, which was preregistered, was accordingly devised to assess the reproducibility of the effect in a new group of participants. In line with our expectations, Experiment 2 revealed both a reliable distance effect and, critically, a reliable context effect. Overall, participants tended to classify the B/13 object as 13 when it appeared with numbers and as B when it appeared with letters, despite not seeing the numbers and the letters. The two-way interaction was negligible, tending negatively, and the context effect was reliable in the first three distance levels,[1] in line with our hypothesis. Note that the negative context effect in the last distance of Experiment 1b was not replicated in this experiment (see Fig. 3; see also Table S10 in the Supplemental Material).

As in Experiment 1a, participants' classifications of the ambiguous letter were affected both by the distance (i.e., classifying the ambiguous letter more as the letter to which it was physically similar) and by the context (i.e., classifying the ambiguous letter more as the letter that formed a word with the context inducers than as the letter that did not form such a word), with the latter seemingly being more pronounced. The two-way interaction was negligible, and the 95% HDI of the context coefficient excluded 0 in the first three distances.

Experiment 3: lexical context effects on letter classification

Experiment 3a (conscious processing). The results of Experiment 3a mirrored those of Experiment 1a: Both distance and context influenced letter classification. That is, participants tended to classify the ambiguous letter as the letter that it most resembled as well as the letter supported by the context. There was no interaction between context and distance (see Fig. 4a; see also Table S11 in the Supplemental Material).

Experiment 3b (unconscious processing). The results of Experiment 3b were similar to the results of Experiments 1 and 2 with respect to distance but not to context: Distance influenced letter classification to the same extent when the context was not consciously perceived. Under invisible conditions, the context neither influenced participants' classifications of the central letters nor interacted with distance: No contextual effects were observed in any of the distances (see Fig. 4b; see also and Table S11). These null results could suggest a qualitative difference in the role of conscious processing between categorical and lexical contextual effects, implying that conscious processing is needed for the latter but not the former. Alternatively, it might also be explained by three methodological points; in Experiments 1 and 2, there were only two possible context inducers, whereas here, a broader range of contextual inducers was used (12 pairs), with fewer repetitions for each, which may have reduced the chances of obtaining an unconscious effect (Van den Bussche, Van den Noortgate, & Reynvoet, 2009). In addition, Experiment 3 had 12 response displays, whereas in Experiments 1 and 2, the response display remained fixed across trials. Finally, in Experiment 3, only three distance levels were used, as opposed to six used before.

We also conducted two preregistered control experiments (Control Experiment 1: $N = 34$, osf.io/hvh8f; Control Experiment 2: $N = 34$, osf.io/kzgtr; for further details, see Supplemental Note, Fig. S2, and Table S12 in the Supplemental Material). These experiments showed that the lack of an unconscious lexical context effect did not stem from these methodological points. In Control Experiment 1 and its replication Control Experiment 2, we adapted the design of Experiment 3b to be as similar as possible to the one used in Experiments 1b and 2: A single stimulus set and response display were used, and there were six levels of distance for the ambiguous object. We further selected the ambiguous letters that induced the strongest (Control Experiment 1) and the second strongest (Control Experiment 2) conscious contextual effects to maximize the chances of finding an effect. Yet lexical inducers still did not evoke contextual effects in the expected direction (see Table S12; if anything, participants sometimes tended to classify the ambiguous letter in the opposite direction, but this tendency was not found in Experiment 4, described below).

Experiment 4: within-participants comparison of categorical and lexical context effects

Taken together, the results seem to imply that top-down contextual effects can occur even without awareness but possibly only at the categorical level (letter/number classification) and not at the lexical level. Because the positive finding at the categorical level goes substantially beyond previous findings, in Experiment 4, we

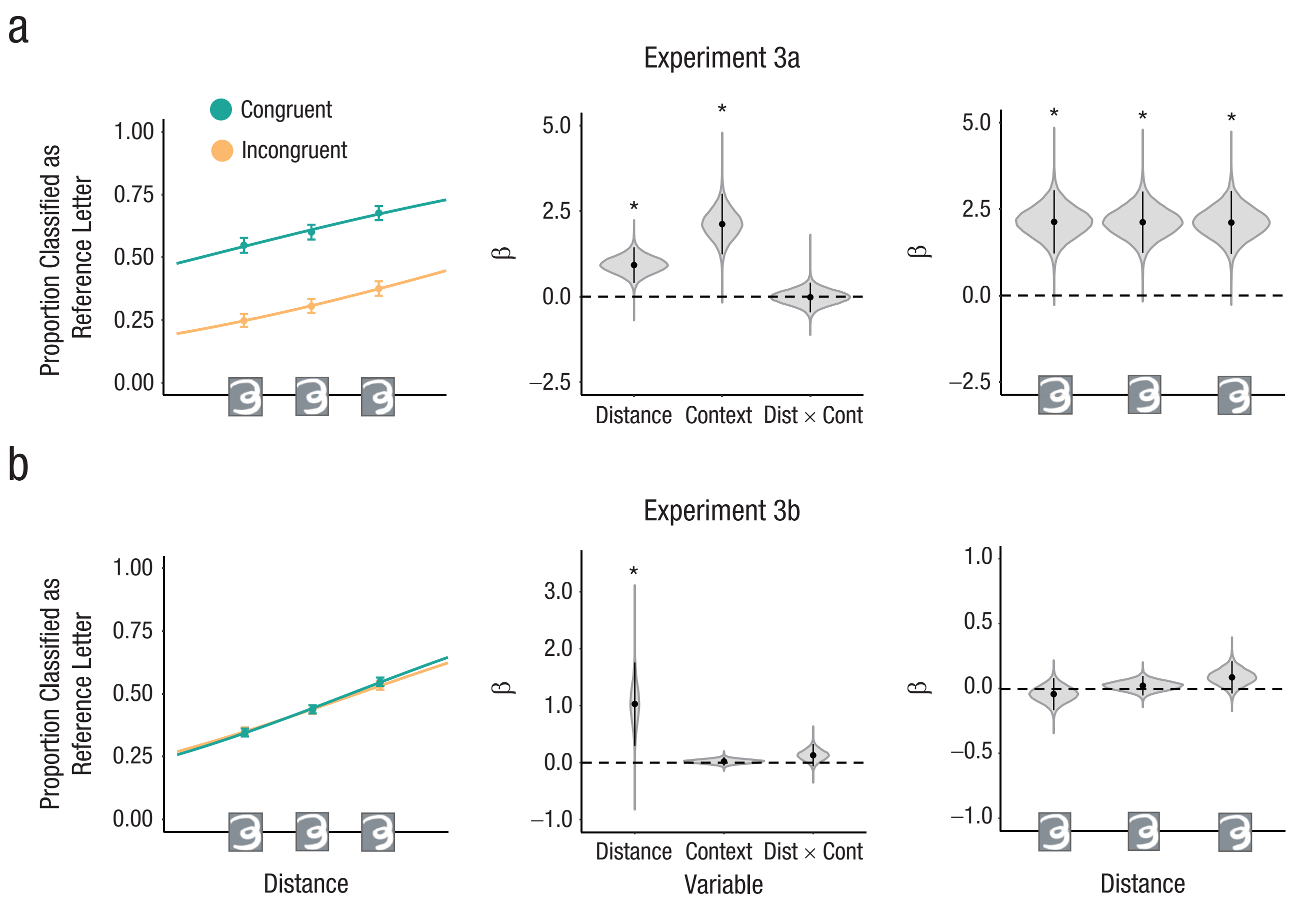

Fig. 4. Results of Experiment 3a (a) and Experiment 3b (b). The graphs in the left column show mean aggregate classifications of the ambiguous stimulus as the reference letter (dots) overlaid with psychometric-function fits (lines). The proportion of classifications is shown as a function of distance between the two parts of the ambiguous stimulus and trial type (congruent vs. incongruent). Error bars denote Wilson 95% binomial confidence intervals. The middle column shows the regression coefficients for the distance, context, and Distance × Context effects, and the right column shows the context coefficient for each distance. In the violin plots, dots indicate medians, and the width of the violin plots denotes the probability density in the full range of the posterior (i.e., depicting every Markov chain Monte Carlo sample of the posterior). Error bars denote 95% highest-density intervals (HDIs). The dashed lines indicate the point at which the regression coefficient equals 0 and is therefore not predictive of the outcome variable. Effects for which the 95% HDI excluded 0 are marked with an asterisk. Model predictors were standardized so the regression coefficients would all lie on the same scale.

wanted to further test its reproducibility and ask whether it is confined to the specific B/13 stimulus. An additional goal was to compare categorical and lexical effects within the same experiment and same participants, using stimuli that are more similar to each other in the two conditions. To that end, we devised a new stimulus set, now using the Hebrew letter ט (/tet/), which in handwriting looks both like the number 6 (see Fig. 1c) and like the handwritten letter ג (/gimel/). This enabled us to create a unique comparison between the context types using very similar stimuli.

Experiment 4a (conscious processing). Context influenced classifications of the ambiguous stimulus in both the categorical and the lexical conditions, yet in this experiment, the categorical effect was greater than the lexical one (and also than the previous contextual effects that we found during conscious processing in Experiments 1–3). The distance coefficient was reliably positive, and the Distance × Condition coefficient was reliably negative, suggesting that distance reliably influenced participants' classifications in the categorical session but not in the lexical session. This might reflect a higher overall ambiguity of the stimuli in the lexical session (i.e., the slope of the psychometric curve was flatter for the lexical relative to the categorical experiment; see Fig. 5a and 5b; also see Table S13 in the Supplemental Material) and might result from our attempt to minimize the visual differences between the stimuli in the categorical and lexical tasks (as for both tasks, we used the same Hebrew letter—ט, /tet/—and tried to make it similar to both the letter ג /gimel/ and the number 6).

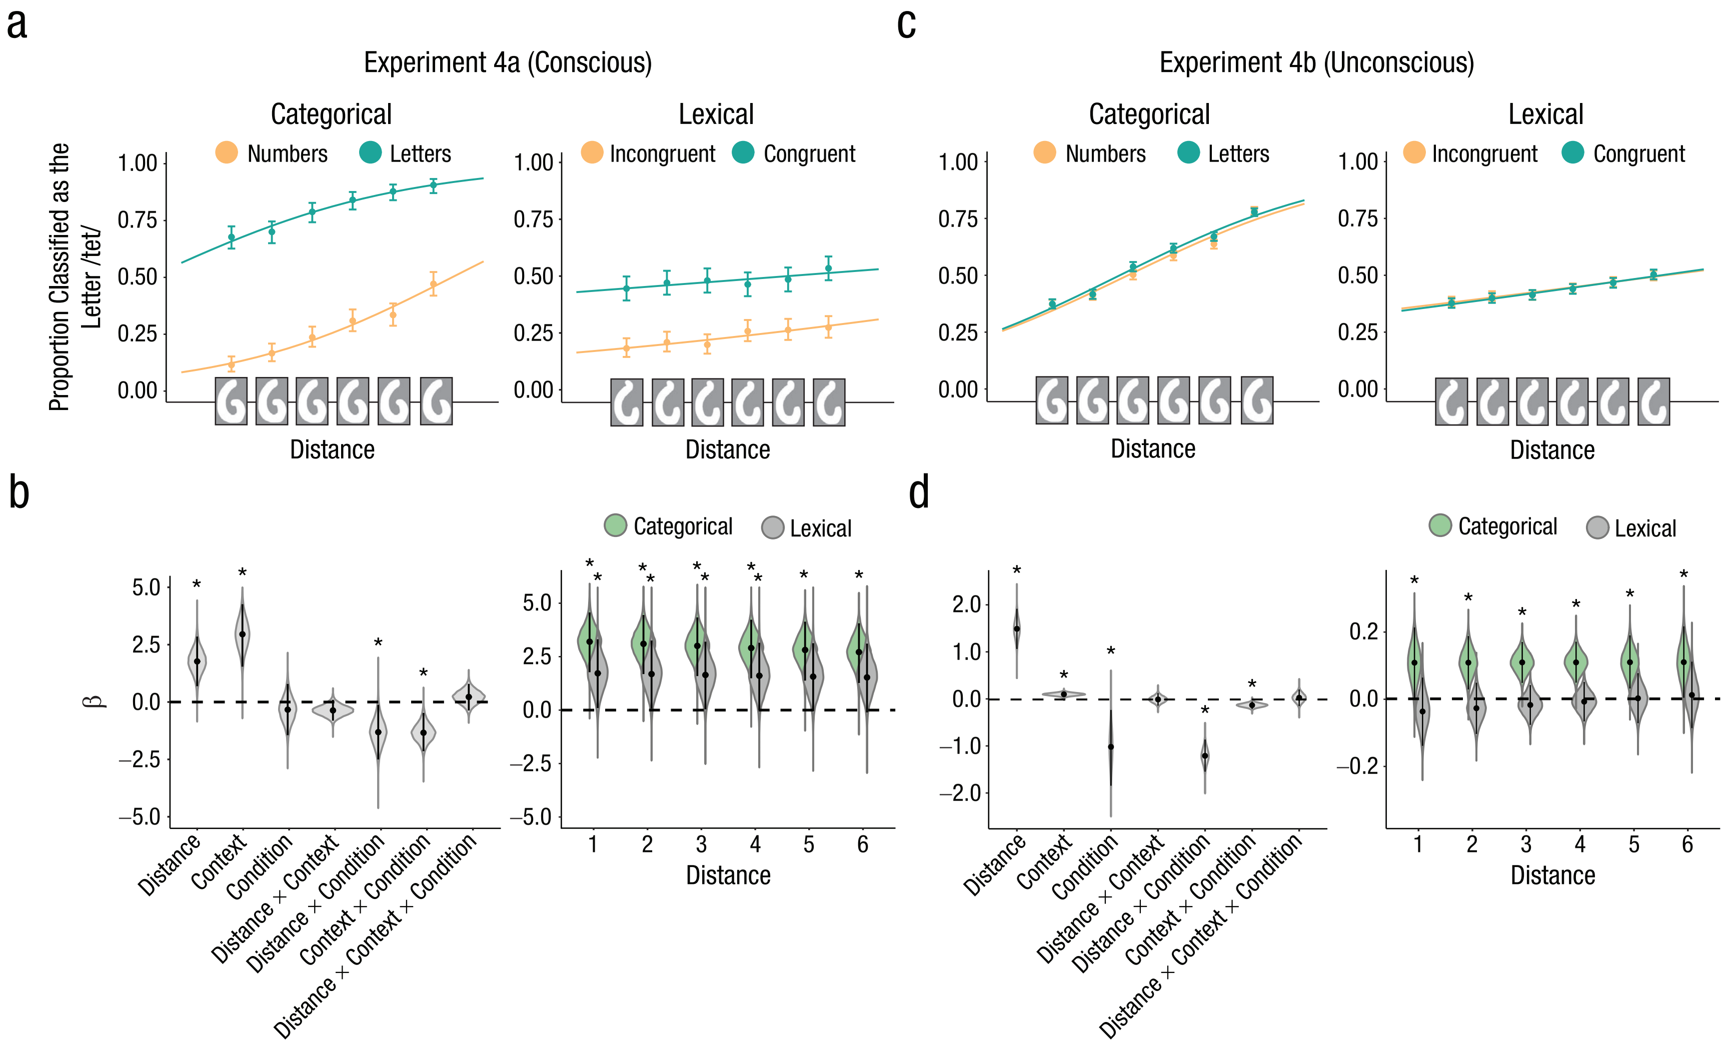

Fig. 5. Results of Experiment 4a (a, b) and Experiment 4b (c, d). The graphs in (a) and (c) show mean aggregate classifications of the ambiguous stimulus as the letter /tet/ (dots) and psychometric-function fits (lines), separately for the categorical and lexical conditions. The proportion of classifications is shown as a function of distance in the ambiguous stimulus and whether the context was of numbers or letters (categorical condition) or whether the context was congruent or incongruent (lexical condition). Error bars denote Wilson 95% binomial confidence intervals. The violin plots in (b) and (d) show group-level posterior coefficient estimates for the key variables in each experiment as well as the context coefficient for each distance (shown separately for the categorical and lexical conditions). Dots indicate medians, and the width of the violin plots denotes the probability density in the full range of the posterior (i.e., depicting every Markov chain Monte Carlo sample of the posterior). Error bars denote 95% highest-density intervals (HDIs). The dashed lines indicate the point in which the regression coefficient equals 0 and is therefore not predictive of the outcome variable. Effects for which the 95% HDI excluded 0 are marked with an asterisk. Model predictors were standardized so the regression coefficients would all lie on the same scale.

Experiment 4b (unconscious processing). In the categorical task, the context coefficient was reliably positive for all six distances. This confirmed the reproducibility of this effect and demonstrated its generalizability across stimuli and languages. Results were akin to those of Experiment 3 and the two control experiments: In the lexical task, the context coefficient was around zero for all distances. That is, in this within-participants design, only categorical context influenced participants' classifications. As for the distance effect, results were similar to those of Experiment 4a: It was reliable only in the categorical session. We also found a negative condition coefficient, suggesting that the average height of the psychometric curve was lower in the lexical condition (see Figs. 5c and 5d and Table S13).

Meta-analysis

Finally, using a meta-level Bayesian logistic regression (a form of Bayesian meta-analysis), we analyzed the results of all six experiments (including the two control experiments) in which unconscious processing was probed (see Fig. 6; see also the full model estimates in Table S14). We found a reliably negative interaction between context and condition (*Mdn* = −0.13, 95% HDI = [−0.23, −0.03]), indicating that the context effect was larger for categorical relative to lexical experiments. Breaking down the Context × Condition interaction shows that at the meta level, context reliably biased classifications of the ambiguous object only in the categorical experiments (*Mdn* = 0.10, 95% HDI = [0.02, 0.17]) but not in the lexical experiments (*Mdn* = −0.03, 95% HDI = [−0.10, 0.03]). The experiment-level context coefficients mirrored the context coefficients obtained when analyzing each experiment separately (note that for Experiment 1b, the context coefficient was now reliably positive, potentially because of the hierarchical pooling to the other categorical experiments and the exclusion of the symbols trials). Moreover, the meta-level distance coefficient reliably excluded zero only in the lexical experiments and was noisy overall, potentially because the magnitude of the distance effects varied considerably between participants and experiments. The condition coefficient was reliably negative, implying that the overall percentage of responses as the reference letter was lower in the lexical experiments. However, the condition effect was completely dependent on the arbitrary definition of the reference letter per experiment and is therefore noninformative in itself. Finally, the meta-level Distance × Context interaction was negligible in both conditions.

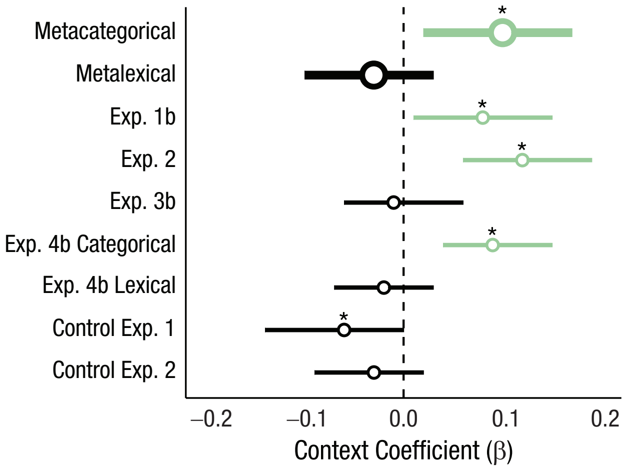

Fig. 6. Context coefficients in the Bayesian meta-analysis of all experiments in which unconscious processing was probed. Categorical-condition coefficients are depicted in green, and lexical-condition coefficients are depicted in black. Thick lines and larger circles depict meta-level context coefficients; thin lines and smaller circles depict the experiment-level coefficients. Circles indicate the median values, and lines signify the 95% highest-density intervals (HDIs) of the posterior distributions. Effects for which the 95% HDI excluded 0 are marked with an asterisk. Model predictors were standardized so the regression coefficients would all lie on the same scale.

Discussion

In six experiments, we tested the role of consciousness in top-down contextual effects on symbolic object processing. To this end, we presented an ambiguous object at varying levels of ambiguity alongside visible or invisible (i.e., masked) context inducers. Visible categorical contexts (numbers, letters, or symbols) greatly modulated the classification of the ambiguous B/13 object at all levels of ambiguity (Experiment 1a). When the context was rendered invisible, it still biased participants' classifications, although to a considerably lower extent: The effects were restricted to the stimulus versions that most resembled the letter B (Experiments 1b and 2). This was further replicated using a completely different stimulus set in a different language (Experiment 4), in which the effect was found for all versions of the ambiguous object. Importantly, invisibility of the context inducers was confirmed in these experiments using both subjective and objective measures. We further corroborated their invisibility by conducting a meta-analysis across objective tasks as well as showing that the strength of the effect was not correlated with participants' level of awareness of the inducers (as reflected by their *d′* performance in the objective task). Thus, the main finding of this study is that unconscious semantic information can indeed affect symbolic object processing in a top-down manner.

Interestingly, such unconscious contextual effects were not found for lexical contexts that formed either words or nonwords with the central ambiguous letter. There, although a clear and seemingly stronger contextual effect was found during conscious processing of the context (Experiment 3a), no effects were found when the inducers were invisible (Experiment 3b), even when a single pair of context inducers was heavily repeated—Control Experiment 1, Control Experiment 2, and Experiment 4b (lexical task)—akin to the design of the B/13 experiments. This dissociation was further demonstrated in a Bayesian meta-analysis showing that across experiments, categorical—but not lexical—context unconsciously biased participants' classification of the ambiguous object. Taken together, the results seem to suggest that lexical contextual effects might require conscious processing. However, as this is based on a series of null results, we cannot make such a conclusive claim at this point.

The unconscious contextual effects we observed at the categorical level (numbers vs. letters) go beyond previous findings of contextual visual illusions evoked by invisible inducers (Banica & Schwarzkopf, 2016; Clifford & Harris, 2005; Fahrenfort et al., 2017; Harris et al., 2011; Mareschal & Clifford, 2012; Moors et al., 2016; Nakashima & Sugita, 2018; Wang et al., 2012): We used semantic, symbolic contexts instead of purely visual ones and tested their influence on symbolic object classification rather than on low-level visual judgments. The use of symbolic stimuli such as numbers or letters plausibly tapped higher level areas that are involved in reading (e.g., visual word-form area; Cohen et al., 2000) and number–symbol processing (e.g., prefrontal cortex and intraparietal sulcus; Nieder & Dehaene, 2009). For such high-level brain areas to influence object classification, feedback to the visual cortex—the hallmark of top-down effects on perception—is most likely needed (Bar, 2004). Our findings thus suggest that such top-down effects might take place without awareness.

At face value, our findings seem to go against the traditional view that consciousness is required for integration. Because participants were affected by the categorical context inducers, and assuming that this requires integration between the context and the ambiguous stimulus, one could claim that this is a case of unconscious integration of the types reported by other studies (e.g., temporal: Reber & Henke, 2012; multisensory: Arzi et al., 2012; Faivre, Mudrik, Schwartz, & Koch, 2014; semantic, knowledge-based integration: e.g., Lin & Murray, 2014; for a review, see Mudrik et al., 2014).

Yet the fact that this effect was limited to categorical contexts only, and not lexical ones, gives rise to two alternative interpretations of the results: The first suggests that our findings actually highlight the *limits* of unconscious integration. This interpretation is in line with the windows-of-integration hypothesis (Mudrik et al., 2014), which was thus far mostly tested in the temporal domain (e.g., Faivre & Koch, 2014; Tu et al., 2019). Under this interpretation, unconscious categorical contextual effects reflect integration that takes place over small integration windows (here, semantically defined), which can be performed without awareness. The lexical contextual effects, on the other hand, might imply bigger integration windows that involve a finer lexical analysis of words and nonwords, hereby requiring conscious processing. Thus, our results can be taken as showing that unconscious integration can occur, but it does so only for processes that involve small integration windows.

The second interpretation goes further, suggesting that although the results convincingly demonstrate unconscious top-down processes, the latter need not involve integration in the sense of forming a unified percept/concept (Mudrik et al., 2014). That is, in the categorical experiments, participants did not have to form a percept of the entire three-character string (e.g., A-B-C) but could instead detect whether the inducers were numbers or letters and infer that the middle character should belong to the same category. Participants could even rely on the category of just one of the inducers and perform equally well. Arguably, then, the bias found in classifying the B/13 stimulus and its Hebrew counterpart was based on the coactivation of that stimulus with the contextual inducers, creating some form of "on-line priming," and not on the formation of an integrated percept of A-B-C or 12-13-14. Such on-line priming cannot drive lexical contextual effects, as they depend on the integration of the three letters into a unified word. According to this account, then, the current study does not show that integration can occur unconsciously but is limited to smaller integration windows, such as the categorical level. Rather, it shows that contextual effects can occur only when integration is not actually needed. And they may fail to exert an effect when integration is required (i.e., when combining invisible letters with a visible one). Future studies are needed to arbitrate between these two alternative interpretations—small size versus no integration—as well as to better delineate the boundary between contexts that can affect perception and those that cannot, for example, by manipulating saliency in the emotional, social, or value domains.

In summary, we demonstrated that top-down contextual effects on the processing of ambiguous objects could be induced by invisible context at the categorical level. These effects bias not only the way that some features of the stimulus are processed (e.g., brightness, tilt) but also the way that the entire stimulus is interpreted and classified. Our findings thus imply that consciousness may not be necessary for top-down contextual effects. As for the role of consciousness in integration, this study speaks for such a role, with two possible interpretations: Either no genuine integration occurs in the absence of awareness, or it is limited to small integration windows, in line with the windows-of-integration hypothesis (Mudrik et al., 2014). When these processes become more and more demanding, consciousness might still be required for them to occur.

Transparency

Action Editor: Alice J. O'Toole
Editor: D. Stephen Lindsay
Author Contributions

D. Biderman and L. Mudrik designed Experiments 1, 2, and 3 and Control Experiment 1. All of the authors designed Control Experiment 2 and Experiment 4. D. Biderman ran Experiments 1, 2, and 3 and Control Experiment 1, and Y. Shir ran Control Experiment 2 and Experiment 4. D. Biderman analyzed the data and wrote the manuscript, which was modified by all of the authors. All the authors approved the final manuscript for submission.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding

This study was supported by the Israel Science Foundation (Grant No. 1847/16).

Open Practices

All data and materials have been made publicly available via the Open Science Framework and can be accessed at osf.io/7236z. The design and analysis plans for the experiments were preregistered (Experiment 2: osf.io/5dpvs, Experiment 3: osf.io/3ccv3, Experiment 4: osf.io/5jnkg, Control Experiment 1: osf.io/hvh8f, Control Experiment 2: osf.io/kzgtr). Deviations from our preregistered analysis plans, as well as the results of the original analyses, are detailed in Supplemental Note and Tables S1 and S2 in the Supplemental Material available online. The complete Open Practices Disclosure for this article can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797620915887. This article has received the badges for Open Data, Open Materials, and Preregistration. More information about the Open Practices badges can be found at http://www.psychologicalscience.org/publications/badges.

TC

ORCID iD

Dan Biderman https://orcid.org/0000-0003-2054-8439

Acknowledgments

We thank Dorit Inbar and Uri Koriski for their help in programming Experiments 1 and 2, and Shir Levin and Alex Manevich for their help in data collection. We especially thank Natalie Biderman for her insightful comments on the manuscript and for help in programming some of the supplementary analyses and in designing the figures. Finally, we thank the members of the Mudrik lab for fruitful discussions throughout this project.

Supplemental Material

Additional supporting information can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797620915887

Note

1. When we included all participants in the analysis, the 95% HDI of the context effect for the fourth distance level reliably excluded zero as well.

References

Arzi, A., Shedlesky, L., Ben-Shaul, M., Nasser, K., Oksenberg, A., Hairston, I. S., & Sobel, N. (2012). Humans can learn new information during sleep. *Nature Neuroscience*, *15*, 1460–1465. doi:10.1038/nn.3193

Baars, B. J. (1997). In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies*, *4*, 292–309.

Balcetis, E. (2016). Approach and avoidance as organizing structures for motivated distance perception. *Emotion Review*, *8*, 115–128. doi:10.1177/1754073915586225

Banica, T., & Schwarzkopf, D. S. (2016). Induction of Kanizsa contours requires awareness of the inducing context. *PLOS ONE*, *11*(8), Article e0161177. doi:10.1371/journal.pone.0161177

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, *5*, 617–629. doi:10.1038/nrn1476

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436. doi:10.1163/156856897X00357

Bruner, J. S., & Minturn, A. L. (1955). Perceptual identification and perceptual organization. *The Journal of General Psychology*, *53*, 21–28. doi:10.1080/00221309.1955.9710133

Clifford, C. W. G., & Harris, J. A. (2005). Contextual modulation outside of awareness. *Current Biology*, *15*, 574–578. doi:10.1016/j.cub.2005.01.055

Cohen, L., Dehaene, S., Naccache, L., Lehéricy, S., Dehaene-Lambertz, G., Hénaff, M. A., & Michel, F. (2000). The visual word form area: Spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain*, *123*, 291–307. doi:10.1093/brain/123.2.291

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37. doi:10.1016/S0010-0277(00)00123-2

Eldar, E., Niv, Y., & Cohen, J. D. (2016). Do you see the forest or the tree? Neural gain and breadth versus focus in perceptual processing. *Psychological Science*, *27*, 1632–1643. doi:10.1177/0956797616665578

Fahrenfort, J. J., van Leeuwen, J., Olivers, C. N. L., & Hogendoorn, H. (2017). Perceptual integration without conscious access. *Proceedings of the National Academy of Sciences, USA*, *114*, 3744–3749. doi:10.1073/pnas.1617268114

Faivre, N., & Koch, C. (2014). Temporal structure coding with and without awareness. *Cognition*, *131*, 404–414. doi:10.1016/j.cognition.2014.02.008

Faivre, N., Mudrik, L., Schwartz, N., & Koch, C. (2014). Multisensory integration in complete unawareness: Evidence from audiovisual congruency priming. *Psychological Science*, *25*, 2006–2016. doi:10.1177/0956797614547916

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral & Brain Sciences*, *39*, Article e229. doi:10.1017/S0140525X15000965

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Boca Raton, FL: CRC Press.

Georgieva, S., Peeters, R., Kolster, H., Todd, J. T., & Orban, G. A. (2009). The processing of three-dimensional shape from disparity in the human brain. *The Journal of Neuroscience*, *29*, 727–742. doi:10.1523/jneurosci.4753-08.2009

Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science*, *273*, 1699–1702. doi:10.1126/science.273.5282.1699

Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, *41*, 1409–1422. doi:10.1016/S0042-6989(01)00073-6

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*, 399–405. doi:10.1177/0956797611399293

Hesselmann, G., Darcy, N., Ludwig, K., & Sterzer, P. (2016). Priming in a shape task but not in a category task under continuous flash suppression. *Journal of Vision*, *16*(3), Article 17. doi:10.1167/16.3.17

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*, 106–154.

Kanizsa, G. (1979). *Organization in vision: Essays on gestalt perception*. New York, NY: Praeger.

Kriegeskorte, N. (2015). Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, *1*, 417–446. doi:10.1146/annurev-vision-082114-035447

Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). San Diego, CA: Academic Press.

Lin, Z., & Murray, S. O. (2014). Unconscious processing of an abstract concept. *Psychological Science*, *25*, 296–298. doi:10.1177/0956797613504964

Mareschal, I., & Clifford, C. (2012). Dynamics of unconscious contextual effects in orientation processing. *Proceedings of the National Academy of Sciences, USA*, *109*, 7553–7558. doi:10.1073/pnas.1200952109 10.1167/12.9.86

Moors, P., & Hesselmann, G. (2018). A critical reexamination of doing arithmetic nonconsciously. *Psychonomic Bulletin & Review*, *25*, 472–481. doi:10.3758/s13423-017-1292-x

Moors, P., Hesselmann, G., Wagemans, J., & van Ee, R. (2017). Continuous flash suppression: Stimulus fractionation rather than integration. *Trends in Cognitive Sciences*, *21*, 719–721. doi:10.1016/j.tics.2017.06.005

Moors, P., Wagemans, J., Van Ee, R., & de-Wit, L. (2016). No evidence for surface organization in Kanizsa configurations during continuous flash suppression. *Attention, Perception, & Psychophysics*, *78*, 902–914. doi:10.3758/s13414-015-1043-x

Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive Sciences*, *18*, 488–496. doi:10.1016/j.tics.2014.04.009

Nakashima, Y., & Sugita, Y. (2018). Size-contrast illusion induced by unconscious context. *Journal of Vision*, *18*(3), Article 16. doi:10.1167/18.3.16

Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual Review of Neuroscience*, *32*, 185–208. doi:10.1146/annurev.neuro.051508.135550

Panichello, M. F., Cheung, O. S., & Bar, M. (2013). Predictive feedback and conscious visual experience. *Frontiers in Psychology*, *3*, Article 620. doi:10.3389/fpsyg.2012.00620

Persuh, M., Emmanouil, T. A., & Ro, T. (2016). Perceptual overloading reveals illusory contour perception without awareness of the inducers. *Attention, Perception, & Psychophysics*, *78*, 1692–1701. doi:10.3758/s13414-016-1146-z

Rabagliati, H., Robertson, A., & Carmel, D. (2018). The importance of awareness for understanding language. *Journal of Experimental Psychology: General*, *147*, 190–208. doi:10.1037/xge0000348

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, *3*, 1–23. doi:10.1023/b:phen.0000041900.30172.e8

Reber, T. P., & Henke, K. (2012). Integrating unseen events over time. *Consciousness and Cognition*, *21*, 953–960. doi:10.1016/j.concog.2012.02.013

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*, 1019–1025. doi:10.1038/14819

Shanks, D. R. (2017). Regressive research: The pitfalls of post hoc data selection in the study of unconscious mental processes. *Psychonomic Bulletin & Review*, *24*, 752–775. doi:10.3758/s13423-016-1170-y

Sklar, A. Y., Deouell, L. Y., & Hassin, R. R. (2018). Integration despite fractionation: Continuous flash suppression. *Trends in Cognitive Sciences*, *22*, 956–957. doi:10.1016/j.tics.2018.07.003

Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience*, *17*, 450–461. doi:10.1038/nrn.2016.44

Tu, S., Wan, S., Jou, J., Ma, Y., Zhao, G., & Pan, W. (2019). Can unconscious sequential integration of semantic information occur when the prime Chinese characters are displayed from left to right? *Attention, Perception, & Psychophysics*. Advance online publication. doi:10.3758/s13414-019-01816-2

Van den Bussche, E., Van den Noortgate, W., & Reynvoet, B. (2009). Mechanisms of masked priming: A meta-analysis. *Psychological Bulletin*, *135*, 452–477. doi:10.1037/a0015329

Wang, L., Weng, X., & He, S. (2012). Perceptual grouping without awareness: Superiority of Kanizsa triangle in breaking interocular suppression. *PLOS ONE*, *7*(6), Article e40106. doi:10.1371/journal.pone.0040106