Washington University School of Medicine

Digital Commons@Becker

Open Access Publications

2016

Cortical integration of contextual information across objects

Tomer Livne

Washington University School of Medicine in St. Louis

Moshe Bar

Harvard University

Follow this and additional works at: https://digitalcommons.wustl.edu/open_access_pubs

Please let us know how this document benefits you.

Recommended Citation

Livne, Tomer and Bar, Moshe, "Cortical integration of contextual information across objects." *Journal of Cognitive Neuroscience*. 28, 7. 948-958. (2016).

https://digitalcommons.wustl.edu/open_access_pubs/5063

This Open Access Publication is brought to you for free and open access by Digital Commons@Becker. It has been accepted for inclusion in Open Access Publications by an authorized administrator of Digital Commons@Becker. For more information, please contact vanam@wustl.edu.

Cortical Integration of Contextual Information across Objects

Tomer Livne^{1,2,3} and Moshe Bar^{1,2,4}

Abstract

■ Recognizing objects in the environment and understanding our surroundings often depends on context: the presence of other objects and knowledge about their relations with each other. Such contextual information activates a set of medial lobe brain regions, the parahippocampal cortex and the retrosplenial complex. Both regions are more activated by single objects with a unique contextual association than by objects not associated with any specific context. Similarly they are more activated by spatially coherent arrangements of objects when those are consistent with their known spatial relations. The current study

tested how context in multiple-object displays is represented in these regions in the absence of relevant spatial information. Using an fMRI slow-event-related design, we show that the precuneus (a subpart of the retrosplenial complex) is more activated by simultaneously presented contextually related objects than by unrelated objects. This suggests that the representation of context in this region is cumulative, representing integrated information across objects in the display. We discuss these findings in relation to processing of visual information and relate them to previous findings of contextual effects in perception. ■

INTRODUCTION

Understanding our environment and coordinating our behavior accordingly can be enhanced and, in many instances, is based on the objects that surround us. Many objects are uniquely associated with a specific context and can point to the potential presence of other objects. The presence of a traffic light, for instance, will indicate the possibility of cars passing by and would suggest caution even if we do not see one at the moment. Similarly, seeing cars going by would (or at least should) prompt us to look for a traffic light if we wish to cross the street. Previous work has indicated that two regions in the medial lobe exhibit sensitivity to the contextual content of visual stimuli. The parahippocampal cortex (PHC) and the retrosplenial complex (RSC; composed of parts of the precuneus and the posterior cingulate) show an increased hemodynamic response when participants are presented with objects strongly associated with a single context compared with objects not associated with any context in particular (Bar & Aminoff, 2003). Similarly, this network shows increased hemodynamic response for images of coherent visual scenes compared with single objects (Epstein, 2008; O'Craven & Kanwisher, 2000; Epstein & Kanwisher, 1998). One possible explanation for these selective patterns of response is that during the recognition process (Bar, 2004) or following a successful recognition of individual objects (Park, Chun, & Johnson, 2010) the brain automatically retrieves informa-

tion regarding the contextual associations of each object. Such a mechanism could confer several benefits such as improving the identification of degraded element (Cheung & Bar, 2013; Bar, 2004) or facilitating judgment about the global scene by converging on the most probable setting, thereby guiding and facilitating detection of searched objects (Wolfe, Alvarez, Rosenholtz, Kuzmova, & Sherman, 2011; Droll & Eckstein, 2008; De Graef, Christiaens, & d'Ydewalle, 1990).

Disrupting the coherence of items in complex displays have been shown to have adverse effects on recognizing and manipulating objects (Mudrik, Lamy, & Deouell, 2010; Davenport & Potter, 2004; Ganis & Kutas, 2003; Bar & Ullman, 1996; Biederman, Glass, & Stacy, 1973; Biederman, 1972) on the recognition of scenes (Davenport, 2007; Joubert, Rousselet, Fize, & Fabre-Thorpe, 2007) and can affect the interpretation of ambiguous objects (Palmer, 1975). The disruption of spatial coherency in the arrangement of objects in a display can also affect the measured hemodynamic response in context selective regions (Gronau, Neta, & Bar, 2008). However, to detect spatial inconsistencies in the display, one has to first identify the objects composing it, then assess their relatedness, and only following both steps a representation of their expected spatial arrangements can be retrieved. It is unclear, however, whether objects are integrated into a unified representation when their contextual relatedness is first recognized or whether they are integrated only if their arrangement is consistent with the stored representation of the scene as previously suggested (Gronau et al., 2008).

¹Harvard Medical School, ²Massachusetts General Hospital, ³Washington University in St. Louis, ⁴Bar Ilan University, Israel

The current study tested the hypothesis that the brain actively integrates contextual information across objects regardless of spatial coherency (which is, as described above, a second-order contextual property). We therefore tested whether the tendency of two objects to appear in the same context is represented in the brain in a cumulative fashion, with converging evidence supporting a specific context translating to a stronger response in context- and scene-selective regions—specifically, asking whether related pairs would produce a larger hemodynamic response than unrelated pairs as an indication that contextual information is represented in a cumulative fashion. If, on the other hand, unrelated pairs would produce a stronger activation than related pairs, it would indicate that a representation of context is retrieved for each unique context (regardless of the number of objects in the display that are related to it). To verify that any finding related to contextual relations is the result of an active analysis of the stimuli and not due to focusing on co-occurrence as an abstract concept, we employed a cued two-task design (see Figure 1). This design enabled us to isolate processes related to the intention to perform a task from processes related to actual performance of that task.

Finally, unlike the majority of functionally defined higher-order visual brain regions such as the fusiform face area and the extrastriate body area (Downing, Jiang, Shuman, & Kanwisher, 2001; Kanwisher, McDermott, & Chun, 1997) that are mostly confined to a single brain structure, the RSC is composed of several distinct brain structures (Aminoff, Gronau, & Bar, 2007). Large portions of the RSC usually fall within the structurally defined precuneus and the cingulate (see Figure 2). This raises the possibility that the commonly used context and scene localizers might tap into multiple processes, each of those being performed by a different brain region. With the cingulate participating in the retrieval of information related to individual objects (e.g., Auger & Maguire, 2013) and the precuneus participating in retrieval of higher-order information such as objects relation via the retrieval of episodic memory information (e.g., Lundstrom, Ingvar, & Petersson, 2005). Another aim of the current study was therefore to investigate the potential different functions that subportions of the RSC perform in relation to contextual processing, as these processes necessarily involve retrieval of higher-order information from memory. This investigation would help extend our understanding of the structure of mental representations of multiple-objects displays.

METHODS

Participants

Fourteen students from the Boston area participated in the study (eight women; mean age = 24, *SD* = 3.9). None reported any neurological problems or use of medication. All had normal or corrected-to-normal vision.

Participants were compensated for their time. The study was approved by a departmental institutional review board.

Task

We tested participants' performance in two tasks: size comparison and contextual relatedness. In the size comparison task, participants were instructed to indicate whether the real-world size of the object presented on the left side of the screen was larger than that of the object presented on the right side of the screen (SIZE task). In the contextual relatedness task, they were asked to indicate whether they would expect to see both objects in the same place ("would you expect to see one if you saw the other"—PLACE task). In both tasks, participants were presented with color pictures of two everyday objects, located right and left from the center of the screen.

The tasks were randomly interleaved within a run, with a cue character appearing at the center of the screen for 1.5 sec, indicating which task the participant should perform on the coming targets. To isolate the response to the target from the response to the cue, there was an SOA of 8 sec between the cue and the targets. Target stimuli were positioned 3.8° of visual angle to the left and to the right of the center of the screen. Each object extended approximately 4.5° × 3.56° of visual angle (height and width, respectively). See Figure 1 for an illustration of a trial time course. Within each task, half of the pairs were drawn from the same environment (related), and half were drawn from a different environment (unrelated). In each condition (Task × Relation combination), half of the displays contained a larger left object, and half of the displays contained a larger right object.

Figure 1. Task design. Each trial started with a 2-sec blank screen, followed by a 1500-msec symbolic cue indicating which task the participant should perform in the following trial ("?" = place task, ">" = size task). Following a 6500-msec interstimulus blank screen interval, two objects appeared on the screen (one on each side of the center of the screen) for 750 msec. Participants were required to perform the task indicated by the preceding cue as quickly and accurately as possible. In the place task, they indicated with a key-press whether they would expect to see both objects at the same place. In the size task, they were required to indicate whether the real-world size of the left object was larger than that of the right object. The next trial started 7250 msec following the disappearance of the objects.

object. Each of the objects had an equal chance of appearing in each condition, and each object appeared only once throughout the experiment. There were six runs of this task, each containing four trials of each condition, presented in a random order. There was a 9.25-sec interval between the targets disappearance and the cue of the next trial (or 15.25 sec if a rest period was randomly assigned). Participants practiced the task before entering the magnet.

The experiment was programmed using Matlab (The MathWorks, Natick, MA) and the Psychtoolbox library (Brainard, 1997; Pelli, 1997). The stimuli were presented to the participants using a projector and a mirror mounted on the head coil.

Stimuli Norming

In a behavioral pilot, a separate group of college students ($n = 33$) was tested on an extended version of the stimuli set used in the main experiment, performing the same task, as well as a naming task and an explicit size estimation task for individual objects (each participant performed the latter two tasks on a subset of the objects). In the naming task, participants were asked to type in the name of the object and then to give a size estimation for this object on a scale from 1 (*smallest*) to 20 (*largest*). Only pairs that were found to be consistently indicated as belonging to the same environment (no more than one participant misclassified them) were used in the main study. In addition, objects that were incorrectly identified during the naming task were excluded from the main study. For the main study, we ended up with a list of 96 pairs (192 objects). Participants' estimation of size was highly correlated with those defined by us (Pearson $r^2 = .87$).

To verify that each of the pairs presented in the experiment conserved the relatedness assignment (related/unrelated) and the size assignment (left or right larger) in that trial, each object was defined using three variables: a paired object to be presented together with in case it was selected for a related trial, a context it belonged to used to exclude potential objects to be paired with if it was to be used in an unrelated trial, and a size estimate used to guide a paired object selection and the side of presentation.

Finally, to test whether relation strength affected the activation patterns for the related pairs, we obtained post hoc estimations of relatedness for each of these pairs from a new group of participants recruited through the Amazon Mechanical Turk platform. Thirty participants rated each pair by answering the question, "If you saw one of these objects, how often would you expect to see the other as well?" on a scale from 1 (*occasionally*) to 5 (*always*). We then averaged the ratings across participants and used the median (3.67) to split the pairs into two groups—strongly associated pairs and moderately associated pairs (48 pairs in each group).

Localizer Scans

To localize ROIs, participants performed two additional tasks upon completion of the main experiment. In two context-localizer scans, participants viewed a continuous stream of objects drawn from a strong context or a weak context pool of objects (see Bar & Aminoff, 2003). To ensure participants' processing of the individual objects, their task was to indicate for each object whether it could fit into a shoebox. Each object appeared once, and there were three blocks of each condition in each run.

To identify scene- and object-selective regions (Epstein & Kanwisher, 1998; Malach et al., 1995), participants performed a 1-back task on images of objects, outdoor scenes containing a building, and a scrambled version of the two previous image types. There were two such scans, each containing two blocks of each condition.

MRI Acquisition

The data were acquired using a Siemens (Erlingen, Germany) 3-T Trio Tim MR system, using a 32-channel RF head coil. Whole-brain echo-planar images were acquired using the following parameters: 33 interleaved slices, 2-sec repetition time, 30-msec echo time, 3-mm thickness. Voxel size was $3.1 \times 3.1 \times 3.0$. Two high-resolution T1-weighted multi-echo MPRAGE anatomical images (isotropic 1-mm voxels) were obtained and used to construct a high resolution image of each participant's brain.

MRI Initial Analysis

Structural and functional analyses were performed using the FreeSurfer and FS-FAST analysis tools and processing stream developed at the Martinos Center for Biomedical Imaging (surfer.nmr.mgh.harvard.edu). Frames were motion-corrected by aligning them to the middle frame in each run. Slice-timing correction was applied to account for the interleaved acquisition sequence. Functional data were coregistered to the high-resolution structural data and then segmented into left and right cortical surfaces. The intensity level of each frame was normalized. The raw time series was then resampled to the reconstructed surfaces. A 5-mm FWHM Gaussian smoothing was applied to the resampled data. This step also produced a cortical segmentation map of brain regions. Finally, all individual participants' reconstructed brains were registered to the 'fsaverage' brain to allow a group level analysis. To estimate the hemodynamic response during the different conditions, we used a gamma function defined by the following parameters: $\Delta = 2.25$ and $\tau = 1.25$. Motion correction parameters were used as nuisance regressors in the analysis. The first five frames (10 sec) were discarded from the analysis.

ROI Identification

Context-selective regions were defined as the significant clusters (after performing a signed cluster thresholding to correct for multiple comparisons) in which the strong contextual objects produced a stronger hemodynamic response than the weak contextual objects. This analysis produced three ROIs: bilateral PHC clusters and an RSC cluster in the left hemisphere (Figure 2A).

Scene-selective regions were defined as the significant clusters that responded stronger to the outdoor scene stimuli than to the object stimuli. These included bilateral scene-selective clusters corresponding to the previously described parahippocampal place area (PPA), RSC region, and the transverse occipital sulcus (TOS; Hasson, Harel, Levy, & Malach, 2003; O'Craven & Kanwisher, 2000; Haxby et al., 1999; Aguirre, Zarahn, & D'Esposito, 1998; Epstein & Kanwisher, 1998; Figure 2A, B).

Object-selective regions were defined as clusters where objects elicited stronger hemodynamic responses than scrambled images. These include bilateral regions in the lateral occipital cortex corresponding to area LOC (Malach et al., 1995; Figure 2B). All ROIs were defined at the group level on the average surface to which all participants were registered.

As we were interested in determining the potential different roles of the precuneus and the cingulate in relation to contextual processing, the RSC ROIs were further segmented based on the automatic anatomical FreeSurfer segmentation to a cingulate and a precuneus portion, henceforth referred to as cingulate and precuneus (Figure 2A).

Figure 2. Context-defined and scene-defined ROIs (see Methods for details). (A) The two main ROIs for this study were the RSC (containing portions of the precuneus and cingulate) and the PHC/PPA in the parahippocampal cortex. Both regions are usually found in context-related studies and scene perception-related studies. Yellow = regions exclusively defined by the context localizer, blue = regions exclusively defined by the scene localizer, red = regions where the activation pattern of the two localizers overlapped. (B) Lateral ROIs: LOC (red) and TOS (yellow). Boundary lines between regions correspond to the FreeSurfer automatic anatomical segmentation.

ROI Analysis of the Main Task

To analyze the main experiment, we recorded the FreeSurfer estimated beta values for each condition in each vertex on an average inflated brain. We then averaged these beta values within each ROI (for each condition and in each participant, separately) and used these average values as the dependent variable in the following analyses. As the overlap between the context and the medial lobe scene localizers was high (Figure 2A) and therefore the activation patterns recorded would not be independent, we conducted two separate analyses: one using the context localizer ROIs and the other using the scene and object localizer ROIs (this particular split was used as the object and the scene ROIs were defined using the same experimental data). For each set of localizers, we conducted a repeated-measures ANOVA with Participant as a random variable of no interest and the following effectors as independent variables: Task (place/size) $\times$ Relation (related/unrelated) $\times$ ROI (PPA/precuneus/cingulate/LOC/TOC or l-PHC/r-PHC/RSC) $\times$ Hemisphere (right/left—scene and object ROIs only).

Although our main interest was in the BOLD response to the target objects, we also conducted an analysis of the BOLD response to the cue to verify that any finding related to contextual relations is indeed the result of an active analysis of the stimuli and not due to focusing on co-occurrence as an abstract concept. To do so, we conducted the same analyses described above on the beta values relating to the cue, this time not including the relation factor as at the time of the cue the participants did not know of the actual contextual relation between the following two objects.

Whole-brain Analysis

To verify the findings related to performing the task (target response) obtained with the ROI analysis, we performed a 2×2 repeated-measures ANOVA (Task $\times$ Relation) for each vertex in the average inflated brain, with participant as a random variable of no interest.

Contextual Association Strength Test

To test whether the neural activation was affected by the strength of association between the related pairs, we conducted a 2×2 repeated-measures ANOVA with Task and Association strength (moderate/high) as the independent variables, for each of the ROIs.

RESULTS

Behavioral Performance

Participants performed the two tasks with a high level of accuracy (mean = 87% correct, $SE = 0.9$). A repeated-measures

Figure 3. Behavioral performance. There were no significant effects on accuracy (top graph). There was a significant interaction effect on RT (bottom graph). Participants were faster in responding to related pairs than unrelated pairs in the place task, but not in the size task.

2 × 2 ANOVA (Task × Relation) on accuracy indicated no differences between any of the conditions (main effects and interaction $p = ns$).

A repeated-measures 2 × 2 ANOVA on RT found a significant interaction effect ($F(1, 11) = 21.05, p < .001$). Post hoc t tests indicated a faster RT for related pairs than unrelated pairs in the place task only (1586 and 1752 msec, respectively). The two main effects were not significant (Figure 3).

fMRI Results

Two ANOVA tests were used to assess the effect of Contextual relation and Task in 13 localizer-defined ROIs (left/right PHC, left/right PPA, left RSC context, left/right precuneus, left/right cingulate, left/right TOS, and left/right LOC; see Figure 2 and the Methods section).

Context Localizer ROIs

We found a main effect of Task ($p = .041$), showing higher beta values for the place task than the size task. We also found a significant interaction between Relation and ROI ($p < .0001$). Post hoc t tests using a false discovery rate (FDR) correction for multiple comparisons (in which the significant p value is adjusted based on the distribution of all the p values smaller than the uncorrected p value set by the researcher; see Benjamini & Hochberg, 1995) with a p value of .05 indicated a significant effect only in the RSC ($p < .0001$), showing higher beta values when the objects were related than when they were unrelated. There was also a significant interaction between Task and ROI ($p = .017$). Post hoc paired-sample t tests using FDR with a p value of .05 indicated a significant Task effect in the RSC ($p = .005$, uncorrected) with higher beta values in the place task than in the size task.

Importantly, using these ROIs, we did not find a Task × Relation interaction ($p > .1$) or a Task × Relation × ROI interaction ($p > .1$). This indicates that the effect of the relation between the two objects was not affected by the task our participants were performing.

We also found a main effect of ROI ($p < .001$) which was of no interest to the current study and will not be discussed further. See Figure 4A for beta values in the individual ROIs in each condition.

Figure 4. Mean beta values in the context-defined ROIs. (A) Mean beta values during the presentation of the target objects. (B) Mean beta values during the presentation of the cue and preparation for the targets appearance. All error bars represent SEM.

None of the effects were significant when the ANOVA was conducted on the cue beta values, indicating that the effects we found with the target beta values were related to active performance of the task and not to preparation processes. See Figure 4B for beta values in the individual ROIs in each condition.

Scene and Object Localizer ROIs

The ANOVA conducted on the ROIs defined by the scene and object localizers indicated a significant interaction between Relation and ROI ($p = .0003$). Post hoc paired-sample t tests using FDR with a p value of .05 indicated a significant effect only in the precuneus ($p = .0002$, uncorrected), showing higher beta values when

the objects were related than when they were unrelated. There was a significant interaction between Task and ROI ($p = .0003$). Post hoc paired-sample t tests using FDR with a p value of .05 indicated a significant effect only in the cingulate ($p = .0014$, uncorrected), with higher beta values in the place task than in the size task.

The ANOVA also indicated a significant Relation $\times$ ROI $\times$ Hemisphere effect ($p = .048$). We therefore split the data into different hemispheres and conducted two post hoc repeated-measures ANOVA, with Relation and ROI as factors and Participant as a random variable of no interest. In both hemispheres, we found a significant Relation $\times$ ROI interaction ($p = .0006$, $p = .002$, left and right, respectively). In both hemispheres, post hoc t tests using FDR indicated a significant effect only in the precuneus

Figure 5. Mean beta values in the scene- and object-defined ROIs. (A) Mean beta values during the presentation of the target objects. (B) Mean beta values during the presentation of the cue and preparation for the targets appearance. All error bars represent *SEM*.

($p < .0001$, $p = .0028$, left and right, respectively). Both hemispheres also showed a main effect of ROI (both $ps < .0001$), which we will not discuss further.

The main ANOVA also indicated a significant Task $\times$ ROI $\times$ Hemisphere interaction ($p = .018$); we therefore conducted two secondary repeated-measures ANOVA: one for each hemisphere with Task and ROI as factors and Participant as a random variable of no interest. In the left hemisphere, we found a Task $\times$ ROI interaction ($p = .0001$); post hoc paired-sample t tests using FDR with a p value of .05 indicated a significant difference only in the cingulate ROI ($p = .0014$). In both hemispheres, the ROI factor was significant (see above).

Importantly—just as was the case with the context ROIs' ANOVA—the main ANOVA on the scene and object ROIs did not indicate any interaction between Task and Relation (with or without additional factors).

The main ANOVA indicated also a significant ROI effect ($p < .0001$) and an ROI $\times$ Hemisphere interaction ($p = .043$), which we will not discuss further. See Figure 5A for beta values in the individual ROIs in each condition.

Conducting the ANOVA on the cue beta values in these ROIs indicated a significant ROI effect ($p < .0001$), a significant Hemisphere effect ($p = .037$), and a significant ROI $\times$ Hemisphere interaction ($p = .003$). Importantly, the task main effect or its interaction effects were not significant (all $ps > .1$). See Figure 5B for beta values in the individual ROIs in each condition.

The results of all the analyses were therefore in agreement showing a relation effect in the precuneus, defined using the context localizer in the left hemisphere and using the scene localizers in both hemispheres (and by anatomical considerations). The results further indicated no relation effect in the cingulate portion of the scene RSC localizers (both hemispheres). The left hemisphere cingulate ROI did indicate a task effect, as did the context-defined RSC (left hemisphere).

Whole-brain Analysis

The results of the whole-brain analysis were in agreement with the ROI analysis indicating large bilateral clusters of a significant Relation effect overlapping the precuneus portions of our RSC ROIs (see Figure 6). These clusters were defined as contiguous vertices showing a significance level of $p \leq .02$, whose overall size was at least 140 mm^2 . Importantly, there was hardly any overlap between the cingulate RSC portions and these clusters (peak Talairach coordinates: left hemisphere: $-19, -56, 19$, right hemisphere: $16, -54, 23$). This lack of overlap supports the results obtained with the context localizer ROI in the left hemisphere and the anatomical parcellation of the scene RSC ROIs in both hemispheres motivated by the memory literature (Auger & Maguire, 2013; Lundstrom et al., 2005). Other regions showing large clusters of significant relation effect were observed bilaterally in the inferior parietal cortex (left: $-50, -56,$

Figure 6. Whole-brain analysis relation effect. These clusters were defined as contiguous vertices showing a significance level of $p \leq .02$ whose overall size was at least 140 mm^2 . See text for peak coordinates. Different colors represent different clusters.

31 , right: $38, -61, 47$), the right supramarginal ($36, -39, 34$), the right posterior cingulate (far away from the ROI portion of the cingulate: $5, -12, 37$), the left rostral-middle frontal region ($-44, 29, 24$), and bilaterally in the superior frontal cortex (left: $-12, 46, 34$, right: $15, 38, 41$).

Using the same cluster correction, we found bilateral precuneus clusters of Relation $\times$ Task interaction effect (Figure 7; peak Talairach coordinates: left hemisphere: $-17, -43, 33$ and $-42, 61, 44$, right hemisphere: $15, -43, 33$). These clusters partly overlapped the clusters of the Relation main effect but, importantly, did not overlap with the ROIs (which, as mentioned above, did not indicate any interaction effect). Overall, the interaction in vertices that exhibited an interaction but did not exhibit a Task main effect was due to a difference between the BOLD responses to related than to unrelated objects during performance of the place task with no difference found during the size task. These results therefore indicate a more widespread activation difference between the related and unrelated pairs when attention was focused on contextual relations than when focused on the objects' size. An additional cluster of interaction was found in the left hemisphere extending from the inferior parietal to the supramarginal ($-6, -65, 32$).

Contextual Associations' Strength

We conducted a 2×2 repeated-measures ANOVA with Task and Association strength (moderate/high) as the

Figure 7. Whole-brain analysis interaction effect. These clusters were defined as contiguous vertices showing a significance level of $p \leq .02$ whose overall size was at least 140 mm². See text for peak coordinates.

independent variables and Participant as a random variable of no interest. No ROI indicated a significant Association strength effect (all p s > .2, uncorrected) or a Task $\times$ Association strength interaction (all p s > .08, uncorrected). This suggests that, for related objects, the brain converges to the most probable and encompassing contextual interpretation, regardless of the strength of the associations. Although our experimental design was not intended to test this question and therefore the analysis was conducted on a lower number of trials than the main analysis and somewhat uneven number of trials in each association strength condition, this analysis did find a significant Task effect in the left RSC (context, $p = .008$; scene: cingulate, $p = .016$; precuneus, $p = .031$, all uncorrected) replicating the results of the main analysis, suggesting that the test did have sufficient power to detect existing effects.

DISCUSSION

In this study, we tested how the brain integrates contextual information from multiple objects in the absence of spatial coherency. Using a slow-event-related paradigm, we demonstrated that the brain exhibits a differential activation pattern when presented with displays consisting of contextually related objects compared with contextually unrelated objects. This differential activation was observed regardless of task demands, directing participants' attention toward or away from the objects' contextual relations. It was, however, more widespread when

attention was focused on contextual relations than when focused on the objects' size, as indicated by the interaction effect in the whole-brain analysis (Figure 7). This differential pattern of activation was observed in a previously defined context- and scene-selective region in the medial lobe, the precuneus, which is part of the RSC (Aminoff et al., 2007).

Previous research using pairs of objects to probe patterns of neural activation related to contextual relatedness typically manipulated the objects' relative spatial arrangement, comparing the activation patterns on trials in which these relations were consistent with real-world situations, with arrangement that violated them (Gronau et al., 2008; Cox, Meyers, & Sinha, 2004). Such a manipulation usually seeks to establish an active relationship between the objects (e.g., a lamp standing on a table or a bottle pouring wine into a glass; see also Aminoff et al., 2007, for a manipulation involving learned spatial relations between meaningless objects). In this study, we aimed at investigating the hemodynamic response to a nonspatial factor, specifically objects' tendency to appear within the same context. We used fixed generic presentation locations for our stimuli (left and right of fixation), which allowed us to contrast the hemodynamic response to related and to unrelated pairs without being confounded by a spatial interaction factor. We can therefore say that the posterior part of the precuneus, in both hemispheres, is sensitive, among other things, to objects' contextual co-occurrences. The selective response to such relationship seems to be task independent, as it was observed even when participants performed a size estimation task (however, it seems that this processing of contextual relations might not be an automatic response, as we discuss below). Because we used a mixed design, it might be possible that our participants also tried to evaluate the objects' contextual relationship even in the size trials or that they tried to imagine the objects inside a scene in those trials. However, given that we found a significantly shorter RT for related compared with unrelated objects in the place task, then assuming participants were consciously performing both tasks simultaneously or first trying to imagine a coherent scene and then compare sizes, one might expect a similar difference in the size task (shorter RT to related than unrelated objects) as well, which we did not observe (see Figure 3).

Despite a strong BOLD response to the stimuli in the PHC and the LOC ROIs, we found no significant difference between related and unrelated pairs. This result is in agreement with a previous report by Kim and Biederman (2011) that used a block design to probe a similar question. However, although Kim and Biederman did not use an RSC localizer, they did perform a whole-brain analysis and did not find a significant relation effect in the precuneus. A possible reason is their use of a block design versus the use of a slow-event-related design in the current study. Kim and Biederman had their participants

perform a repetition task during their experimental blocks. It is possible that when their participants were asked to compare the two objects in each trial to those of the previous trial they were actively integrating the two objects in each display, thereby bypassing the normal contextual integration process. This might also explain the lack of the expected difference between the related and the unrelated pairs in the interacting condition in the PHC (Gronau et al., 2008) in their study. Further studies will be required to determine under which tasks the default contextual integration process is used and under which tasks it is disrupted.

What is the context-related role of the precuneus and cingulate portions of the RSC? Within the scene-defined RSC ROI and in the whole-brain analysis, we found a division between the precuneus and the cingulate. The precuneus seems to be sensitive to contextual relatedness and coherence, whereas the cingulate does not. The fact that we observed no relation effect in the cingulate despite the overall strong activation in this region suggests that whether the cingulate is involved in the retrieval of objects information from memory (Auger & Maguire, 2013) or in the analysis of information retrieved by a different brain region, it likely does so independently for each object; otherwise, we would expect to see a difference between the activation patterns caused by related and unrelated pairs. This finding can potentially explain the stronger activation for buildings than for single objects used to define this ROI (not presented). The buildings in the building stimuli were embedded in a background, and most of them contained additional objects. Therefore, more activation would be expected relative to single objects assuming independent processing for each element in the stimulus. This explanation seems consistent with previous reports (e.g., Auger & Maguire, 2013), suggesting that the cingulate retrieves information independently for each object. Auger and Maguire additionally showed that the activation levels in this region are affected also by the object's tendency to appear in a particular position in the three-dimensional space (see also Troiani, Stigliani, Smith, & Epstein, 2014). This is, however, only one aspect of the information this region retrieves and responds to as evident by the strong hemodynamic response when processing nonpermanent objects in this study, even if to a lesser degree than would permanent objects. In this study, we demonstrated that this region was modulated by the task participants were performing, with stronger activation when performing the place task relative to the size task (left hemisphere). Because every object in our stimuli set could have appeared in any condition, the permanency property cannot be the only property this region is sensitive to. Similarly, our contextual relation effect was not confounded by individual object sizes as previously suggested (e.g., Troiani et al., 2014; Konkle & Oliva, 2012), because every object in the experiment was randomly assigned to any of the four conditions and therefore there were no consistent size differences between the conditions.

It is possible that when asked to perform the place task participants actively retrieved more relevant contextual information of each object causing the stronger activation. This would indicate that the level of processing and the breadth of the retrieved information are dependent on conscious intentions. As Auger and Maguire's results were obtained using a different method than the one used in this study (pattern classification and mean activation estimates, respectively), a direct comparison of the results is complicated, and it is possible that subsets of the neurons in this portion of the cingulate are dedicated to representing/retrieving spatial location information.

The precuneus, on the other hand, exhibited a stronger activation for related than unrelated pairs, suggesting that it is involved in evaluating and integrating information related to both objects. This explanation is consistent with previous reports suggesting that the precuneus is involved in the "generation and retrieval of rich episodic contextual association" (Lundstrom et al., 2005). And although we did not manipulate episodic memory, it is possible that such information is retrieved whenever familiar objects are recognized. This suggested that the role of the precuneus is also in-line with findings showing a strong activation level in this region when retrieving learned associated pairs from memory (Krause et al., 1999; Fletcher et al., 1995) and when perceiving familiar places (compared with unfamiliar places; Epstein, Parker, & Feiler, 2007). Related pairs in this study were expected to elicit a recognizable coherent context, whereas such context could not be activated with the unrelated pairs. An inspection of the set of regions that exhibited a relation effect in the whole-brain analysis offers another possible link to memory retrieval. In addition to the precuneus region, we also observed bilateral parietal (inferior parietal lobe) clusters and, to some extent, small frontal and temporal clusters. These regions partially overlap with the default mode network (DMN: Raichle et al., 2001; Shulman et al., 1997), which has been implicated in memory retrieval (Rugg & Vilberg, 2013; Buckner, Andrews-Hanna, & Schacter, 2008). A partial overlap between the context network and the medial portion of the DMN has been noted in Bar, Aminoff, Mason, and Fenske (2007) and has been discussed there in relation to memory and contextual associations' retrieval. Here we show that under the current experimental design we find also a partial overlap with the lateral and frontal portions of the DMN. The partial overlap of our task-defined clusters with the DMN is also consistent with previous reports of selective recruitment of DMN regions based on task demands (Sestieri, Corbetta, Romani, & Shulman, 2011). However, it should also be noted that the overlap with the DMN was only partial and that these regions have been implicated in many other mental processes. For instance, the inferior intraparietal sulcus with which our parietal clusters overlap has been suggested to participate in the integration of elements into unified

representations (Xu, 2008). This suggested role is also compatible with the current suggestion of integrating objects into a single context.

As indicated above, previous studies reported the effect of spatially contradictory information on brain activation patterns and on performance levels in behavioral tasks (Gronau et al., 2008; Joubert et al., 2007; Davenport & Potter, 2004; Biederman et al., 1973; Biederman, 1972). Our suggested role for the precuneus is not at odds with these findings. It is possible that the activation magnitude in the precuneus represents not only the fact that objects are related to each other but also the extent to which the existing evidence support a single coherent contextual interpretation. It is possible that reducing confidence in the interpretation by introducing spatial inconsistencies would scale down the magnitude of activation in this brain region. On the other hand, it is also possible that there is a fundamental difference between spatially incoherent and spatially coherent displays. Spatially incoherent displays, like the ones used in the current study, might give rise only to the retrieval of semantic information about the objects in the display and their relations, whereas spatially coherent displays might further elicit the construction of a mental representation of the scene these objects are a part of. Further studies are required to understand the interplay between contextual consistency and spatial coherency and the contribution of each type of information to the construction of a mental representation of a scene.

A possible interpretation of the current finding is that the brain represents contextual information in a winner-takes-all fashion. According to this explanation, a single context is always selected when several possibilities exists whereas the other possible contextual interpretations and the representation and retrieval of objects associated with them are inhibited. Although this study did not explicitly probe the content of the created representation, it does offer a partial support for the first part of this interpretation. We have found that moderately related pairs activated the brain to the same degree as strongly related pairs. This might suggest that the brain tends to converge to the most probable and encompassing contextual interpretation, whenever such interpretation is available. In future studies, we intend to explore how the brain deals with inconsistent information and the fate of discarded interpretations. If discarded interpretations are actively inhibited along with their associated objects, it could explain several past findings such as those of Davenport and Potter (2004) or Joubert et al. (2007) on extended recognition time required to identify objects and scenes in displays containing incongruent information.

Our results show that contextual relatedness alone is sufficient for the brain to integrate contextual information across individual objects and represent their shared context, even in the absence of spatial congruency between the presented and the real-world arrangement of these objects.

Acknowledgments

We thank Olivia Cheung for helpful discussions and for commenting on an early draft and Matthew Panichello for assistance in data acquisition. We thank the Harvard decision science laboratory for providing us with lab space and equipment to run the behavioral pilot part of the study. This work was supported by NSF grant BCS-0842947 and by the Israeli Center of Research Excellence in Cognitive Sciences (ICORE grant no. 51/11).

Reprint requests should be sent to Tomer Livne, Department of Radiology, Washington University School of Medicine in St. Louis, 4525 Scott Ave., St. Louis, MO 63110, or via e-mail: tlivne@hotmail.com.

REFERENCES

- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). An area within human ventral cortex sensitive to "building" stimuli: Evidence and implications. *Neuron*, 21, 373–383.
- Aminoff, E., Gronau, N., & Bar, M. (2007). The parahippocampal cortex mediates spatial and nonspatial associations. *Cerebral Cortex*, 17, 1493–1503.
- Auger, S. D., & Maguire, E. A. (2013). Assessing the mechanism of response in the retrosplenial cortex of good and poor navigators. *Cortex*, 49, 2904–2913.
- Bar, M. (2004). Visual objects in context. *Nature Reviews: Neuroscience*, 5, 617–629.
- Bar, M., & Aminoff, E. (2003). Cortical analysis of visual context. *Neuron*, 38, 347–358.
- Bar, M., Aminoff, E., Mason, M., & Fenske, M. (2007). The units of thought. *Hippocampus*, 17, 420–428.
- Bar, M., & Ullman, S. (1996). Spatial context in recognition. *Perception*, 25, 343–352.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B, Methodology*, 57, 289–300.
- Biederman, I. (1972). Perceiving real-world scenes. *Science*, 177, 77–80.
- Biederman, I., Glass, A. L., & Stacy, E. W. (1973). Searching for objects in real-world scenes. *Journal of Experimental Psychology*, 97, 22–27.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Buckner, R. L., Andrews-Hanna, J. L., & Schacter, D. L. (2008). The brain's default network. *Annals of the New York Academy of Sciences*, 1124, 1–38.
- Cheung, O., & Bar, M. (2013). The resilience of object predictions: Early recognition across viewpoints and exemplars. *Psychonomic Bulletin & Review*, 21, 682–688.
- Cox, D., Meyers, E., & Sinha, P. (2004). Contextually evoked object-specific responses in human visual cortex. *Science*, 303, 115–117.
- Davenport, J. L. (2007). Consistency effects between objects in scenes. *Memory & Cognition*, 35, 393–401.
- Davenport, J. L., & Potter, M. C. (2004). Scene consistency in object and background perception. *Psychological Science*, 15, 559–564.
- De Graef, P., Christiaens, D., & d'Ydewalle, G. (1990). Perceptual effects of scene context on object identification. *Psychological Research*, 52, 317–329.
- Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, 293, 2470–2473.
- Droll, J., & Eckstein, M. (2008). Expected object position of two hundred fifty observers predicts first fixations of seventy

- seven separate observers during search. *Journal of Vision*, 8, 320a.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392, 598–601.
- Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, 12, 388–396.
- Epstein, R. A., Parker, W. E., & Feiler, A. M. (2007). Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. *Journal of Neuroscience*, 27, 6141–6149.
- Fletcher, P. C., Frith, C. D., Baker, S. C., Shallice, T., Frackowiak, R. S., & Dolan, R. J. (1995). The mind's eye—Precuneus activation in memory-related imagery. *Neuroimage*, 2, 195–200.
- Ganis, G., & Kutas, M. (2003). An electrophysiological study of scene effects on object identification. *Cognitive Brain Research*, 16, 123–144.
- Gronau, N., Neta, M., & Bar, M. (2008). Integrated contextual representation for objects' identities and their locations. *Journal of Cognitive Neuroscience*, 20, 371–388.
- Hasson, U., Harel, M., Levy, I., & Malach, R. (2003). Large-scale mirror-symmetry organization of human occipito-temporal object areas. *Neuron*, 37, 1027–1041.
- Haxby, J. V., Ungerleider, L. G., Clark, V. P., Schouten, J. L., Hoffman, E. A., & Martin, A. (1999). The effect of face inversion on activity in human neural systems for face and object perception. *Neuron*, 22, 189–199.
- Joubert, O. R., Rousselet, G. A., Fize, D., & Fabre-Thorpe, M. (2007). Processing scene context: Fast categorization and object interference. *Vision Research*, 47, 3286–3297.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17, 4302–4311.
- Kim, J. G., & Biederman, I. (2011). Where do objects become scenes? *Cerebral Cortex*, 21, 1738–1746.
- Konkle, T., & Oliva, A. (2012). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, 74, 1114–1124.
- Krause, B. J., Schmidt, D., Mottaghy, F. M., Taylor, J., Halsband, U., Herzog, H., et al. (1999). Episodic retrieval activates the precuneus irrespective of the imagery content of word pair associates—A PET study. *Brain*, 122, 255–263.
- Lundstrom, B. N., Ingvar, M., & Petersson, K. M. (2005). The role of precuneus and left inferior frontal cortex during source memory episodic retrieval. *Neuroimage*, 27, 824–834.
- Malach, R., Reppas, J. B., Benson, R., Kwong, K. K., Jiang, H., Kennedy, W. A., et al. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, 92, 8135–8139.
- Mudrik, L., Lamy, D., & Deouell, L. Y. (2010). ERP evidence for context congruity effects during simultaneous object–scene processing. *Neuropsychologia*, 48, 507–517.
- O'Craven, K., & Kanwisher, N. (2000). Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience*, 12, 1013–1023.
- Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory and Cognition*, 3, 519–526.
- Park, S., Chun, M., & Johnson, M. (2010). Refreshing and integrating visual scenes in scene-selective cortex. *Journal of Cognitive Neuroscience*, 22, 2813–2822.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences, U.S.A.*, 98, 676–682.
- Rugg, M. D., & Vilberg, K. L. (2013). Brain networks underlying episodic memory retrieval. *Current Opinion in Neurobiology*, 23, 255–260.
- Sestieri, C., Corbetta, M., Romani, G. L., & Shulman, G. L. (2011). Episodic memory retrieval, parietal cortex, and the default mode network: Functional and topographic analyses. *Journal of Neuroscience*, 31, 4407–4420.
- Shulman, G. L., Fiez, J. A., Corbetta, M. B., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common blood flow changes across visual tasks: II. Decreases in cerebral cortex. *Journal of Cognitive Neuroscience*, 9, 648–663.
- Troiani, V., Stigliani, A., Smith, M. E., & Epstein, R. A. (2014). Multiple object properties drive scene-selective regions. *Cerebral Cortex*, 24, 883–897.
- Wolfe, J. M., Alvarez, G. A., Rosenholtz, R., Kuzmova, Y. I., & Sherman, A. M. (2011). Visual search for arbitrary objects in real scenes. *Attention, Perception, & Psychophysics*, 73, 1650–1671.
- Xu, Y. (2008). Representing connected and disconnected shapes in the human inferior intraparietal sulcus. *Neuroimage*, 40, 1849–1856.