CHAPTER 1

Top-down facilitation of visual object recognition: object-based and context-based contributions

Mark J. Fenske, Elissa Aminoff, Nurit Gronau and Moshe Bar*

MGH Martinos Center for Biomedical Imaging, Harvard Medical School, 149 Thirteenth Street, Charlestown, MA 02129, USA

Abstract: The neural mechanisms subserving visual recognition are traditionally described in terms of bottom-up analysis, whereby increasingly complex aspects of the visual input are processed along a hierarchical progression of cortical regions. However, the importance of top-down facilitation in successful recognition has been emphasized in recent models and research findings. Here we consider evidence for top-down facilitation of recognition that is triggered by early information about an object, as well as by contextual associations between an object and other objects with which it typically appears. The object-based mechanism is proposed to trigger top-down facilitation of visual recognition rapidly, using a partially analyzed version of the input image (i.e., a blurred image) that is projected from early visual areas directly to the prefrontal cortex (PFC). This coarse representation activates in the PFC information that is back-projected as "initial guesses" to the temporal cortex where it presensitizes the most likely interpretations of the input object. In addition to this object-based facilitation, a context-based mechanism is proposed to trigger top-down facilitation through contextual associations between objects in scenes. These contextual associations activate predictive information about which objects are likely to appear together, and can influence the "initial guesses" about an object's identity. We have shown that contextual associations are analyzed by a network that includes the parahippocampal cortex and the retrosplenial complex. The integrated proposal described here is that object- and context-based top-down influences operate together, promoting efficient recognition by framing early information about an object within the constraints provided by a lifetime of experience with contextual associations.

Keywords: object recognition; top-down; feedback; orbitofrontal cortex; low spatial frequencies; visual context; parahippocampal cortex; retrosplenial cortex; visual associations; priming

Successful interaction with the visual world depends on the ability of our brains to recognize visual objects quickly and accurately, despite infinite variations in the appearance of objects and the settings in which they are encountered. How does the visual system deal with all of this information in such a fluent manner? Here we consider the cortical mechanisms and the type of information that they rely on to promote highly efficient visual recognition through top-down processes. The evidence we review, from studies by our lab and others, suggests that top-down facilitation of recognition can be achieved through an object-based mechanism that generates predictions about an object's identity through rapidly analyzed, coarse information. We also review evidence that top-down facilitation of recognition can be achieved through the predictive information provided by contextual associations between an object or scene

*Corresponding author. Tel.: +1-617-726-7467; Fax: +1 617-726-7422; E-mail: bar@nmr.mgh.harvard.edu

and the other objects that are likely to appear together in a particular setting. In the following sections, we first consider each of these forms of top-down enhancement of object recognition separately, and then consider how these object- and context-based mechanisms might operate together to promote highly efficient visual recognition.

An object-based cortical mechanism for triggering top-down facilitation

The traditional view regarding visual processing is that an input image is processed in a bottom-up cascade of cortical regions that analyze increasingly complex information. This view stems from the well-defined functional architecture of the visual cortex, which has a clear hierarchical structure. However, several models propose that both bottom-up and top-down analyses can occur in the cortex simultaneously (Grossberg, 1980; Kosslyn, 1994; Ullman, 1995; Desimone, 1998; Engel et al., 2001; Friston, 2003; Lee and Mumford, 2003), and recent evidence suggests that top-down mechanisms play a significant role in visual processing (e.g., Kosslyn et al., 1993; Humphreys et al., 1997; Barceló et al., 2000; Hopfinger et al., 2000; Miyashita and Hayashi, 2000; Gilbert et al., 2001; Pascual-Leone and Walsh, 2001; Mechelli et al., 2004; Ranganath et al., 2004a). Nevertheless, the way in which such top-down processing is initiated remains an important outstanding issue. The crux of the issue concerns the fact that top-down facilitation of perception requires high-level information to be activated before some low-level information.

Recently, Bar (2003) proposed a model that specifically addresses the question of how top-down facilitation of visual object recognition might be triggered. The gist of this proposal concerns a cortical (or subcortical) "short-cut" of early information through which a partially analyzed version of the input image, comprising the low spatial frequency (LSF) components (i.e., a blurred image), is projected rapidly from early visual areas directly to the prefrontal cortex (PFC). This coarse representation is subsequently used to activate predictions about the most likely interpretations of the input
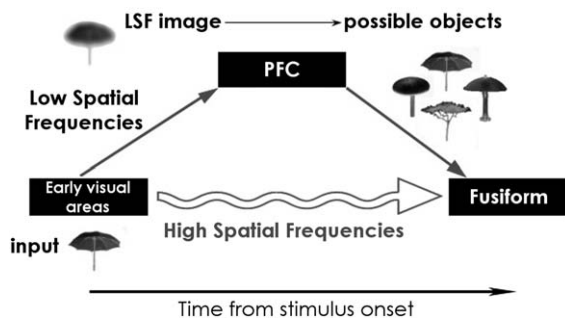

Fig. 1. Schematic illustration of a cortical mechanism for triggering top-down facilitation in object recognition (Bar, 2003). A low spatial frequency (LSF) representation of the input image is projected rapidly, possibly via the magnocellular dorsal pathway, from early visual cortex to the prefrontal cortex (PFC), in addition to the systematic and relatively slower propagation of information along the ventral visual pathway. This coarse representation is sufficient for activating a minimal set of the most probable interpretations of the input, which are then integrated with the relatively slower and detailed bottom-up stream of analysis in object-processing regions of the occipito-temporal cortex (e.g., fusiform gyrus).

image in the temporal cortex. For example, if an image of an oval blob on top of a narrow long blob were extracted from the image initially (Fig. 1), then object representations sharing this global profile would be activated (e.g., an umbrella, a tree, a mushroom, a lamp). Combining these top-down "initial guesses" with the systematic bottom-up analysis could thereby facilitate visual recognition by substantially limiting the number of object representations that need to be tested.

Orbitofrontal involvement in top-down facilitation of recognition

The proposal that object recognition might benefit from a cortical "short-cut" of early, cursory information about the input image implies that there should be an additional cortical region that shows recognition-related activity before other object-processing regions. The cortical regions most often associated with visual object recognition are situated in the occipito-temporal cortex (Logothetis et al., 1996; Tanaka, 1996). Within this region, the fusiform gyrus and the lateral occipital cortex are especially crucial for visual recognition in humans

(Kosslyn et al., 1995; Martin et al., 1996; Bar et al., 2001; Grill-Spector et al., 2001; Malach et al., 2002). The PFC has also been shown to be involved in visual recognition (e.g., Bachevalier and Mishkin, 1986; Wilson et al., 1993; Parker et al., 1998; Freedman et al., 2001), and recent evidence suggests that the orbitofrontal cortex (OFC) might be specifically related to top-down visual processing (Bar et al., 2001, 2006; Bar, 2003). Bar et al. (2001), for example, used functional magnetic resonance imaging (fMRI) to compare the cortical activity elicited by trials in which objects were successfully recognized with that elicited by the same pictures under identical conditions when they were not recognized. As expected, this contrast showed differential activity in the occipito-temporal regions previously associated with object recognition. However, successful recognition was also associated with increased activity in a site within the OFC (Fig. 2). The connection between activity in the OFC and successful object recognition makes it a prime candidate for being involved in top-down facilitation of visual recognition.

Rapid recognition-related activity in orbitofrontal cortex

If the recognition-related region of the OFC identified in the Bar et al. (2001) study is critical for early top-down facilitation, then recognition-
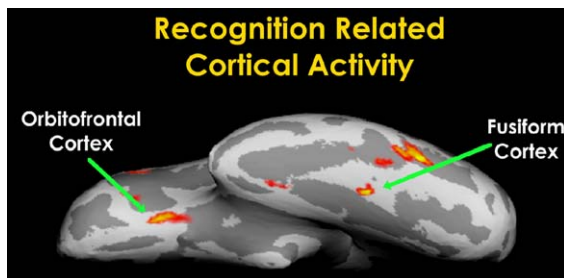

Fig. 2. A group averaged statistical activation map showing recognition-related cortical activity (adapted from Bar et al., 2001) on the left hemisphere (ventral view) of a brain that has been "inflated" to expose both sulci (dark gray) and gyri (light gray). Successful object recognition produced greater activity than unsuccessful recognition attempts in inferior regions of both the occipito-temporal (e.g., fusiform gyrus) and the frontal (e.g., orbitofrontal cortex) lobes.

related activity should develop in this site before other object-processing regions of occipito-temporal cortex. To test this prediction, Bar et al. (2006) used the same object recognition task as Bar et al. (2001) while obtaining magnetoencephalography (MEG) recordings, which provide millisecond-resolution measurements of cortical activity. As predicted, the contrast between recognized and not-recognized trials in this study revealed differential activation in the OFC 50 ms before it developed in the object-processing regions of the occipito-temporal cortex (Fig. 3). Moreover, a time-frequency, trial-by-trial covariance analysis of the MEG data demonstrated strong synchrony between occipital visual regions and the OFC at a relatively early stage (beginning at approximately 80 ms after stimulus onset), and a strong synchrony between the OFC and the fusiform gyrus activity at a relatively later stage (130 ms after stimulus onset). Taken together, these results provide strong converging support for the proposal that this frontal region plays a critical role in top-down facilitation of object recognition.

Early orbitofrontal activity is triggered by low spatial frequencies

The finding that early activity in the OFC is related to successful object recognition is consistent with the model of top-down facilitation proposed by Bar (2003). This model posits that early activity in the OFC is related to the direct projection from early visual areas of an LSF representation of the input image (Bar, 2003). The projection of this early and rudimentary information to the OFC would thereby allow it subsequently to sensitize the representation of the most likely candidate objects in the occipito-temporal cortex as "initial guesses." The possibility that early OFC activity is driven by LSF information is supported by physiological findings that the magnocellular pathway conveys LSF information (i.e., a blurred image) early and rapidly (Shapley, 1990; Merigan and Maunsell, 1993; Bullier and Nowak, 1995; Grabowska and Nowicka, 1996; Leonova et al., 2003), and by evidence from anatomical studies showing bidirectional connections between early visual
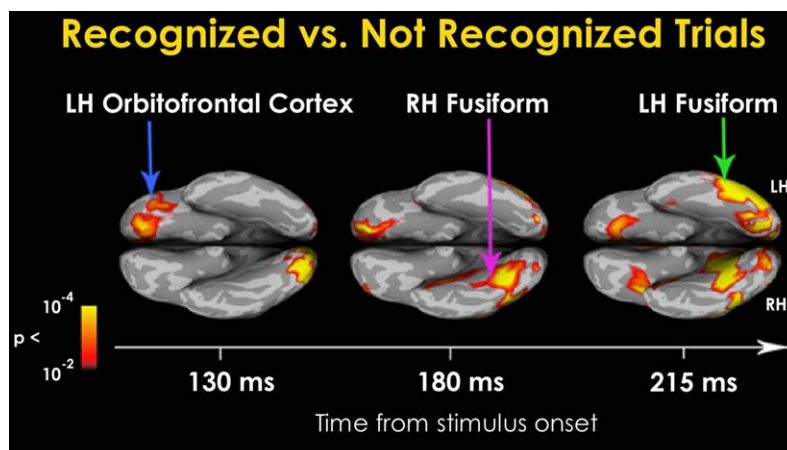
Fig. 3. Cortical dynamics of object recognition. Bar et al. (2006) found that recognition-related activity in the orbitofrontal cortex (OFC) precedes that in the temporal cortex. Anatomically (MRI) constrained group averaged statistical activation maps were calculated from MEG at three different latencies from stimulus onset. Recognition-related activation (recognized vs. not-recognized) peaked in the left OFC 130 ms from stimulus onset, 50 ms before it peaked in recognition-related regions in the temporal cortex.

areas and the prefrontal cortex (e.g., humans: Oenguer and Price, 2000; macaque: Rempel-Clower and Barbas, 2000) (although information about direct connections between occipital visual areas and the OFC is still lacking). These findings suggest that the neural infrastructure is indeed in place for LSF information to trigger prefrontal activity subserving top-down facilitation.

To test the prediction that early recognition-related activity in the OFC should depend on the availability of LSF information in the image, Bar et al. (in press) used both fMRI and MEG to compare activity in the OFC elicited by filtered images of objects (Fig. 4) containing predominantly LSFs or predominantly high spatial frequencies (HSFs). They found that the same region of the OFC that shows early recognition-related activity (Fig. 3) also shows early differential activity for LSF vs. HSF images. This differential activity associated with spatial frequency content was found to peak around 115 ms from stimulus onset. Moreover, a trial-by-trial covariance analysis of the MEG data indicated that there was a clear interaction between the occipital visual regions and the OFC, and between the OFC and the fusiform gyrus for LSF images, as suggested by phase synchrony, but no significant synchrony existed between these regions for the HSF stimuli.

Taken together, the results reviewed in this section provide strong support for the proposal that visual object recognition benefits from the rapid analysis and projection of coarse LSF information of an input image to orbitofrontal regions of prefrontal cortex. Importantly, this object-based mechanism allows early information about an object to act as the catalyst for top-down enhancement of its own bottom-up cortical processing.

A context-based cortical mechanism for triggering top-down facilitation

In addition to the top-down benefit provided by prefrontal analysis of cursory object information, recognition efficiency can be increased through processes that take advantage of naturally occurring regularities in the environment (Gibson, 1969). A lifetime of visual experience can thereby guide expectations about which objects are likely to appear in a given setting to aid subsequent recognition of those objects through their contextual associations (Biederman, 1972, 1981; Palmer, 1975; Biederman et al., 1982; Bar and Ullman, 1996; Davenport and Potter, 2004). A beach setting, for example, typically contains objects such as beach chairs and beach umbrellas. Recognizing a scene with ocean surf breaking on the

Fig. 4. Examples of the visual stimuli used in the Bar et al. (2006) MEG and fMRI studies. Unfiltered stimuli (middle) were used as intact controls for comparison with activity for low filtered (left) and high filtered (right) stimuli.

sand, combined with this prior knowledge about the scene's typical content, may act to rapidly sensitize the representations of these contextually related objects and facilitate their subsequent recognition. Indeed, several studies have shown that scenes can be recognized in a ''glance'' and that their content is extracted very rapidly (Biederman et al., 1974; Schyns and Oliva, 1994; Oliva and Torralba, 2001), lending support to the notion that early analysis of the surrounding contextual information may facilitate object recognition. Other studies have directly examined the interaction of scene and object processing, showing that objects are recognized faster when located in a typical environment than in an unusual setting. Palmer (1975), for instance, showed that recognition of briefly presented objects following the presentation of a contextual scene was significantly more accurate when the target object was contextually consistent with the preceding scene (e.g., kitchen — bread) than when it was inconsistent (e.g., kitchen — mailbox). In another series of experiments, Biederman et al. (1982) demonstrated that violating various types of contextual relations (e.g., relative position or size) between an object and its environment hinders subjects' ability to detect the object. Similar effects were reported by Davenport and Potter (2004). Taken together, these findings suggest that objects and their settings are processed interactively, supporting the view that contextual information directly facilitates the perceptual processes involved in visual object recognition (but see Hollingworth and Henderson, 1998, for a different perspective).

Cortical analysis of contextual information

To understand how cortical analysis of contextual associations might generally facilitate perception, and object recognition in particular, it is useful to first consider how the brain analyzes these associations. Bar and Aminoff (2003) recently addressed this issue in a series of fMRI studies. They identified those regions of the brain that are sensitive to contextual information by comparing cortical activity during recognition of individual objects with strong contextual associations with that for objects with weak contextual associations. For example, the association of a toaster with a kitchen setting is relatively stronger than the association of a camera with the various contexts in which it tends to occur. Bar and Aminoff found that such objects with strong contextual associations elicited greater activity than those with weak contextual associations, and did so primarily in the parahippocampal cortex (PHC) and the retrosplenial complex (RSC).[1] They found similar context-related activation in these regions for each subject and across several related experiments, showing clearly that the PHC and the RSC comprise a cortical "context

---

[1]The cortical region we refer to as the retrosplenial complex extends beyond the callosal sulcus, which is where recent cytoarchitectonic investigations have shown the retrosplenial cortex to be almost entirely located (Vogt et al., 2001). The retrosplenial region that is involved in contextual analysis (Bar and Aminoff, 2003) has a broader extent that is typically centered near the junction of the callosal sulcus, the precuneus, the caudal-ventral portion of the posterior cingulate gyrus, and the ventral-medial portion of the subparietal sulcus.

network'' that processes contextual associations during object recognition.

In addition to this newly established role in processing contextual associations (Bar and Aminoff, 2003), the PHC and RSC have both previously been shown to mediate the processing of spatial, place-related information (e.g., Aguirre et al., 1996, 1998; Epstein and Kanwisher, 1998; Levy et al., 2001). Indeed, a region within the PHC typically shows a preferential response to pictures of scenes and topographical landmarks, and has therefore been termed the parahippocampal place area (PPA, Epstein and Kanwisher, 1998). Interestingly, in Bar and Aminoff's (2003) studies, the same region of PHC showed robust context-related activity for individual objects presented in isolation (see Fig. 5). How could individual objects activate a region of the cortex that is typically sensitive to scenes and landmarks? One possibility is that the perception of a strongly contextual object (e.g., a bowling pin) indirectly activates the representation of the place with which it is associated (e.g., a bowling alley). However, it is also possible that the PHC and RSC process contextual associations more generally. Thus, while places and landmarks may often correspond with specific sets of associations, context-related activity in the

PHC and RSC might not be restricted to these highly spatial, place-specific contexts. This issue directly concerns the nature of the information that is processed by the context network.

The nature of the information processed by the context network

Does context-related activity in the PHC and RSC reflect the processing of contextual associations in general, or the processing of particular associations with spatial, place-specific contexts? To distinguish between these alternatives, Bar and Aminoff (2003) used objects with weak contextual associations as a baseline and then compared the relative increase in fMRI activation for objects with strong contextual associations with spatial, place-related contexts (e.g., farm) or more abstract nonspatial contexts (e.g., romance). As before, the PHC and RSC produced significantly greater activity for strongly contextual objects than for weakly contextual objects. Importantly, robust context-related activity was found for both types of contexts (Fig. 6). These results provide strong support for the hypothesis that the context network mediates the processing of contextual
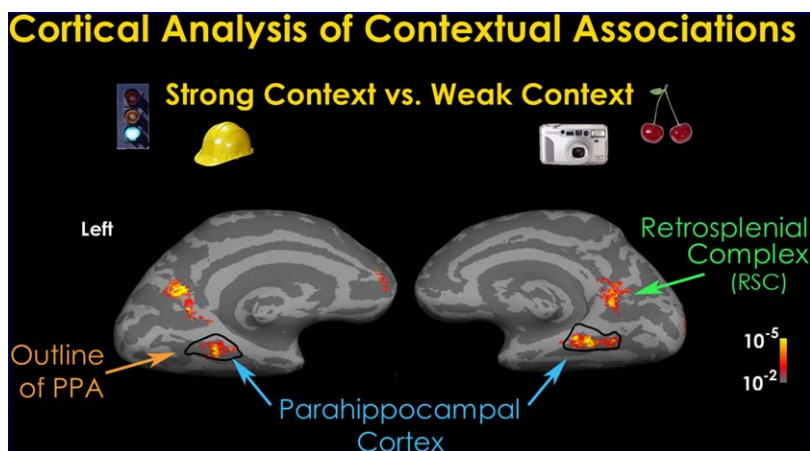

Fig. 5. Group averaged statistical activation maps (medial view) showing greater fMRI response in the parahippocampal cortex (PHC) and the retrosplenial complex (RSC) to individual objects with strong contextual associations (strong context) than to objects with weak contextual associations (weak context). The black outlines in the PHC show the boundaries of the parahippocampal place area (PPA), which was defined in a separate set of subjects by comparing cortical activation for indoor and outdoor scenes with that for objects, faces, and scrambled items (i.e., as in Epstein et al., 2003). Robust context-related activity occurs within the PPA for strong context vs. weak context individual objects.
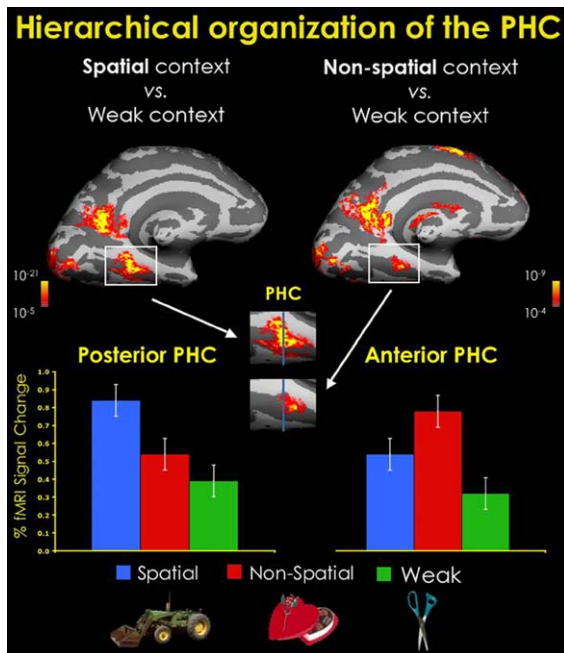
Fig. 6. Group averaged statistical activation maps (left hemisphere, medial view) showing greater fMRI response for objects strongly associated with spatial and nonspatial contexts than for objects with weak contextual associations. Bar graphs show the results of a region of interest analysis for the PHC. The magnitude of fMRI signal change was largest for spatial context objects in posterior PHC and for nonspatial context objects in anterior PHC.

associations in general, and that the representation of contextual associations in these regions is not confined to spatial, place-related contexts. Finding robust activation for objects from nonspatial contexts clearly shows that context-related activity in the PHC in response to individual objects is not merely a reflection of the indirect activation of the specific places with which they are associated.

Bar and Aminoff's (2003) comparison of cortical activity elicited by objects associated with spatial vs. nonspatial contexts revealed an interesting characteristic of the functional organization of the PHC. That is, the activity elicited by objects associated with nonspatial contexts was found to be confined to the anterior half of the PHC, while activity elicited by the objects associated with spatial contexts extended into the posterior half of the PHC (Fig. 6). This pattern of activation suggests that the representation of contextual associations

in the PHC is organized along a hierarchy of increasing spatial specificity for more posterior representations. Additional evidence for this type of hierarchical organization has since been found in subsequent studies showing similar differences in the sensitivity of posterior and anterior regions of the medial temporal lobe to spatial information (Düzel et al., 2003; Jackson and Schacter, 2004; Pihlajamaki et al., 2004). This characteristic pattern of cortical activity during contextual analysis indicates that the nature of the information processed by the context network involves both spatial and nonspatial associations.

The involvement of spatial information in contextual representations might go beyond mere associations, and extend to representing the typical spatial relations between objects sharing the same context. A traffic light, for example, is usually located relatively high above a street and is typically fixed to a standard that rises from the corner of a sidewalk, with crosswalk markings and pedestrians below it, and cars driving underneath or stopped in front of it. Are these spatial relations represented in the associations for spatial contexts? The results of several behavioral studies indicate that this type of spatial information can indeed be associated with an object's representation and thereby impact its recognition (Biederman, 1972, 1981; Palmer, 1975; Biederman et al., 1982; Bar and Ullman, 1996). Bar and Ullman (1996), for example, demonstrated that spatial contextual information is important when identifying an otherwise ambiguous object. Thus, a squiggly line that cannot be identified in isolation can readily be identified as a pair of eyeglasses when presented in the appropriate spatial configuration with a hat (Fig. 7). These results support the hypothesis that spatial relations between objects are part of contextual representations that bind information about typical members of a context as well as the typical spatial relations between them, termed *context frames*.

Disentangling the processing and representation of spatial and nonspatial contextual associations is complicated by the fact that real-world objects are encountered in specific places and spatial locations, even those that are otherwise most strongly associated with nonspatial contexts. For example, while a heart-shaped box of chocolates might be

Fig. 7. Example of the importance of spatial context in object recognition. An ambiguous line (A) cannot be interpreted unequivocally in isolation, but (B) is readily identified as a pair of eyeglasses when shown in the appropriate location below a drawing of a hat.
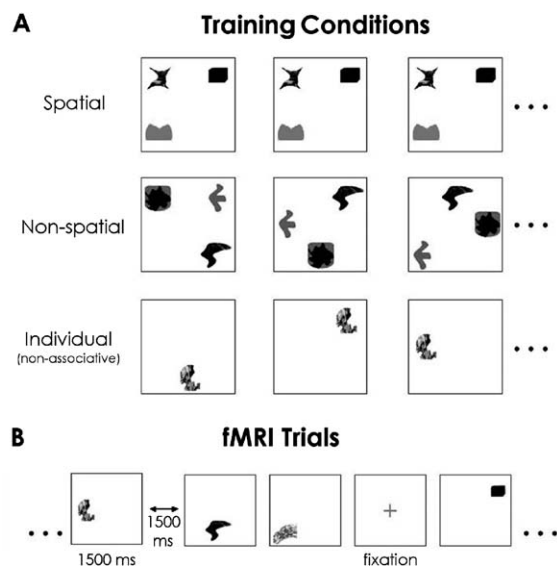

Fig. 8. Isolating spatial and nonspatial components in the formation of contextual associations. (A) During training, groups of three colored abstract shapes (examples shown in grayscale) are repeatedly presented together in the same specific spatial locations (spatial) or in random locations (nonspatial). A nonassociative control condition (individual) is formed by presenting single objects in random locations. (B) fMRI trials involve individual presentations of a single shape from each training-session display, regardless of condition.

most strongly related to the context of romance, people may also associate it with the store where it was sold or the room where the chocolates were eaten. Nevertheless, if the PHC is sensitive to both spatial and nonspatial associations through visual experience, such sensitivity should also develop in highly controlled situations using novel, unfamiliar contextual associations.

This type of experimental approach can maximize control over spatial and nonspatial contextual associations, and will be important for further examining the nature of the contextual information that is processed and represented in the PHC (Aminoff et al., submitted). For example, contextual associations could be formed through an extensive training phase in which subjects are repeatedly exposed to stimuli from each of three different conditions: spatial, nonspatial, and nonassociative (Fig. 8A). Under the spatial condition, individual groups of three shapes would be shown together in a consistent spatial configuration and with each shape always appearing in the same display location. Individual groups of three shapes would also be shown together under the nonspatial condition, but in a random spatial configuration on each trial. Finally, under a nonassociative condition, single shapes would be presented individually in a randomly determined location on each trial. These individual shapes would therefore not be associated with any other shape, nor any particular location. Consequently, these individual shapes could serve as control stimuli, analogous to the weak contextual objects used in the previous experiments.

After such a training session, fMRI could be used to detect differences in cortical activity for the contextual shapes relative to that for the individual, nonassociative shapes. To do this, a single shape from the previous training session could be presented during each trial in the scanner. Thus, the only difference for these individual shapes from the different conditions during scanning would be the prior experience the subject had with each shape (see Fig. 8B). In line with our previous results, we would expect the results of such a study to reveal a robust fMRI response in both the PHC and the RSC when subjects viewed shapes with newly formed contextual associations (spatial and nonspatial) relative to that for the nonassociative control stimuli. Importantly, the same anterior-to-posterior progression in spatial specificity in context-related activity of the PHC would be evident if the magnitude of the fMRI response were greatest in the posterior portion of the PHC for shapes from the spatial condition, and greatest in the anterior portion of the PHC for shapes from the nonspatial condition. Finding a similar anterior-to-posterior progression of activity in the PHC for nonspatial and spatial associations using otherwise

novel abstract stimuli would provide converging evidence that the representation of contextual information in the PHC is organized along a hierarchy of spatial specificity.

The general conclusion drawn from the results of the experiments reviewed in this section is that the PHC and RSC together mediate the processing of contextual associations, and that the nature of these associations involve both spatial and nonspatial information. Characterizing the role of these regions in associative processing provides a critical foundation for exploring the cognitive and neural mechanisms underlying contextual facilitation of object recognition. The evidence that the PHC and RSC mediate associative processing also provides a framework for bridging previous findings about the function of the PHC and RSC. Specifically, in addition to contextual analysis, the PHC and RSC have both been implicated in the processing of spatial or navigational information (Aguirre et al., 1998; Epstein and Kanwisher, 1998; Maguire, 2001), as well as in episodic memory (Valenstein et al., 1987; Brewer et al., 1998; Wagner et al., 1998; Davachi et al., 2003; Morcom et al., 2003; Wheeler and Buckner, 2003; Kirwan and Stark, 2004; Ranganath et al., 2004b). Thus, at first glance it would seem that the same cortical regions have been assigned several different and seemingly unrelated functions. This potential conflict is resolved, however, when one considers that all of these processes rely on associations as their building blocks. Place-related and navigational processing requires analysis of the associations between the objects and landmarks appearing in a specific place, as well as the spatial relations among them. Episodic memories also rely on familiar associations. For example, an episodic memory of last night's dinner is a conjunction of associated constituents (e.g., the location, the company, the food, the dishes, etc.). Thus, describing the key role of the PHC and RSC in terms of associative processing, rather than limiting this role to spatial or episodic information, provides a common foundation from which various functions can be derived. As we discuss below, the ability of these regions to rapidly process such associations is critical for triggering top-down contextual facilitation of visual object recognition.

Contextual facilitation of object recognition

In the preceding sections we have considered evidence that object recognition can be facilitated in the presence of contextually related information. We have also described the cortical network that is specifically involved in analyzing contextual associations. But how does the processing of contextual associations lead to facilitated visual recognition? We propose that the contextual associations analyzed during the recognition of an object or scene are represented and activated in a corresponding "context frame" in the PHC. Context frames can be considered to be contextual representations (or schema) that include information about the objects that typically appear within that context (see also Bar and Ullman, 1996; Bar, 2004). The precise nature of the information represented in context frames is still being explored, but the results reviewed above suggest that it includes both spatial and nonspatial aspects of contextual relations. In this way, the PHC might serve as a switchboard of associations between items that are represented in detail elsewhere (Bar, 2004). The specific associations mediated by an activated context frame might act to sensitize, or prime, representations of other contextually related objects (Fig. 9). Spreading activation of contextually related object representations might thereby help to facilitate subsequent recognition of such objects as a function of their contextual associations. Interestingly, while the functional role of such efficient representation of contextual associations is hypothesized to benefit perception, evidence from visual false memory studies suggests that certain memory distortions can arise as a byproduct of activating these associations. For example, when tested for their memory of a previously shown picture, subjects in these studies often "remember" having seen objects that are contextually related to the picture but that were not actually in the picture (Miller and Gazzaniga, 1998).

Bar and Aminoff's (2003) studies showed that robust context-related activity can be elicited by a single key object from a particular context. This finding suggests that recognizing a single object might be sufficient to activate a corresponding context frame, prime the associated representations
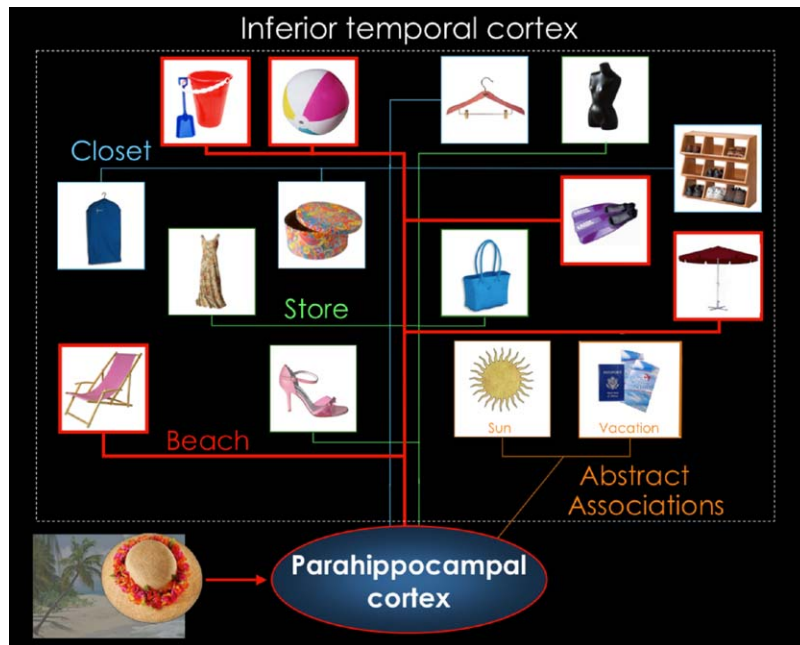
Fig. 9. A model of how contextual associations in the PHC might activate visual representations of contextually related objects in inferior temporal cortex to facilitate object recognition (adapted from Bar, 2004). Individual objects (e.g., a straw hat) can be associated with multiple context frames (e.g., a beach, a store, a closet). The experience-based set of associations represented in a specific context frame is activated as soon as the context has been established through recognition of a strongly contextual object (e.g., a palm tree) or through other contextual cues from the scene.

and thereby facilitate recognition of the other contextually related objects. The hypothesis that initial recognition of even a single strongly contextual object can facilitate subsequent recognition of a different, but contextually related, object was tested directly by Fenske et al. (submitted). While their study focused primarily on the neural substrates of contextual facilitation, it is interesting to note that no prior study had addressed even the behavioral characteristics of this type of object-to-object contextual priming using foveal presentations (cf., Henderson et al., 1987; Henderson, 1992).

Fenske et al. (submitted) presented individual objects with strong contextual associations within a visual priming paradigm, and asked subjects to make a simple size judgment about each object. Recognition improvement due to contextual priming effects was assessed by comparing response times (RTs) and the corresponding change in fMRI signal for objects that were immediately preceded by a contextually related object relative to that for objects from a new unprimed context

(Fig. 10). The important result of this study was robust contextual facilitation of recognition-related RTs and corresponding fMRI response reductions in both context- and object-processing regions. Cortical response reductions were found in bilateral PHC and in the left anterior fusiform, lateral occipito-temporal, and inferior frontal cortices. At the behavioral level, these findings replicate previous studies showing more efficient and rapid recognition of contextually related, than unrelated, items. At a neural level, the novel and important finding of the Fenske et al. study was the robust cortical response reduction obtained in the context network and object-processing regions. Such cortical response reductions are important, as traditional priming studies have shown them to be a hallmark of enhanced recognition following prior exposure to the same, or a related, stimulus (Schacter and Buckner, 1998; Koutstaal et al., 2001; Vuilleumier et al., 2002; Henson, 2003; Simons et al., 2003; Lustig and Buckner, 2004; Maccotta and Buckner, 2004; Zago et al., 2005).
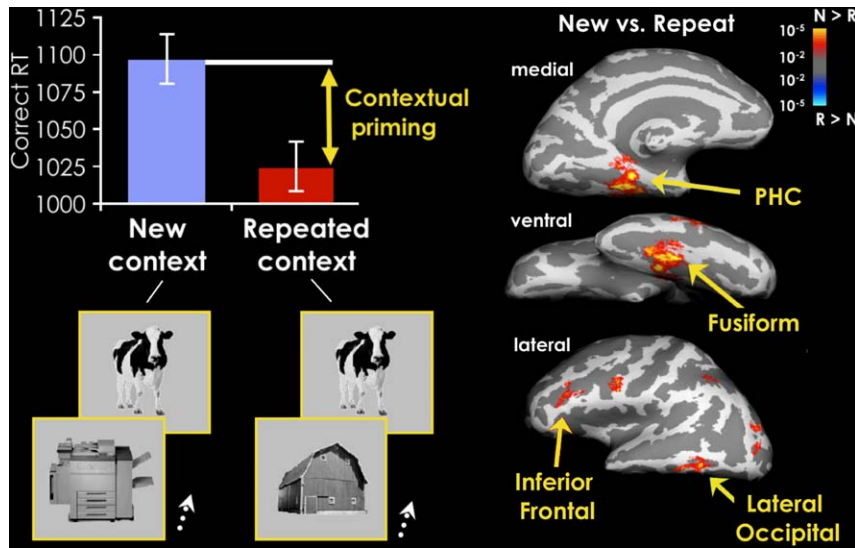
Fig. 10. Contextual facilitation of object recognition. Objects preceding repeated-context target items were always contextually related; objects preceding new-context items were never contextually related. After each picture presentation, participants were required to respond ''bigger than a shoebox,'' or ''smaller than a shoebox,'' by a key-press. Contextual priming was reflected by faster response times (RT) and reduced cortical fMRI response for repeated-context objects than for new-context objects. Contextual priming related fMRI response reductions occurred in the parahippocampal component of the context network, and in object-related regions of occipito-temporal and inferior frontal cortices.

Finding a response reduction in a primary component of the context network (i.e., PHC) supports the hypothesis that context-related facilitation of object recognition results from an enhanced processing of an object's contextual associations. Likewise, the response reduction seen in the left anterior fusiform gyrus implies that encountering a strongly contextual object activates representations of other contextually related objects (such as the contextually related targets), while the response reduction in the lateral occipital cortex, an area implicated in the processing of perceptual features, reflects activation of specific perceptual representations associated with these target objects. Finally, the response reduction in left inferior frontal cortex presumably reflects sensitization of semantic associations between context-related objects. The selective contextual priming effects found in these specific regions therefore provide important insight into the types of representations involved in the contextual facilitation of visual recognition. Importantly, these results suggest that there are many different types of information that are connected by the proposed context frames

stored in the PHC, but that the various types of representations (perceptual, conceptual, etc.) are stored elsewhere.

The results of the Fenske et al. (submitted) study provide support for our proposal that contextual analysis during initial recognition of a highly contextual object serves to activate a corresponding context frame and, through the inherent associative connections, sensitizes the representations of other contextually related objects. Recognition is therefore enhanced when these objects whose representations have been primed in this manner are subsequently encountered. However, objects rarely appear in isolation, so it makes sense that context frames might also be activated by the information contained in a scene image, per se. Indeed, just as a coarse, LSF representation of an input image is often sufficient for rapid object recognition, the LSF content in a scene image might also be sufficient for deriving a reliable guess about the context frame that needs to be activated. The work of Schyns, Oliva, and Torralba (Schyns and Oliva, 1994; Oliva and Torralba, 2001) supports the proposal that LSF information from a typical scene is

sufficient for a successful categorization of its context (e.g., a street, a beach), and the statistics of these images can help produce expectations about what and where is likely to appear in the scene. This work converges with psychophysical and computational evidence that LSFs are extracted from scene images earlier than HSFs (Palmer, 1975; Metzger and Antes, 1983; Bar, 2003). The proposal that emerges from this is that contextual facilitation of object recognition, just like object-based top-down facilitation, can be triggered rapidly on the basis of early cursory analysis of an input scene or object image.

While the results of the Fenske et al. (submitted) study clearly show that recognition of a strongly contextual object can facilitate the subsequent recognition of other contextually related objects, and that this contextual priming effect is associated with changes in activity in the PHC component of the context network, this study does not address the relative top-down influence of different types of contextual associations (e.g., spatial vs. nonspatial) on subsequent recognition. Nevertheless, the finding that contextual facilitation of object recognition is also associated with cortical response reductions in cortical regions associated with different aspects of object processing (e.g., perceptual, semantic, etc.) suggests that contextual facilitation may indeed be mediated through various types of object-related representations. To maximize contextual facilitation, the context frames that maintain the corresponding associations must therefore accommodate various types of representations. How are context frames organized for this to be achieved? We consider this question in the following section.

Contextual facilitation and the representation of spatial and nonspatial associations

The context network is proposed to mediate the processing and representation of various types of contextual associations through context frames. The nature of these context frames and how they represent different types of associations remains an open and exciting area for future research. Bar and Ullman (1996) suggested that contextual

frames contain information about not only the identity of objects that tend to co-occur in scenes but also their typical spatial relations. Indeed, our experiments have repeatedly shown that the PHC component of the context network is involved in the analysis of both spatial and nonspatial associations during recognition of individual objects (e.g., Bar and Aminoff, 2003). In addition, there is clear behavioral evidence that spatial contextual information can facilitate object recognition. Several studies have shown that objects are recognized more efficiently and accurately when located in contextually consistent (expected) locations than in contextually inconsistent (unexpected) locations (Biederman, 1972; Mandler and Johnson, 1976; Hock et al., 1978; Biederman, 1981; Biederman et al., 1982; Cave and Kosslyn, 1993; Bar and Ullman, 1996; Chun and Jiang, 2003). For instance, when viewing a scene of an office, a table lamp would typically be recognized faster when appearing on top of a desk than when appearing beneath it, suggesting that prior knowledge about the spatial constraints of objects in scenes contributes to object recognition.

However, it remains unclear whether nonspatial (e.g., semantic) information about the specific identities of objects in scenes is indeed integrated, as Bar and Ullman (1996) have suggested, with knowledge about the spatial relations between these objects. In other words, does context-based top-down facilitation of object recognition rely on spatial and nonspatial associations that are linked within unified context frames or that are represented independently? This question is critical for understanding the nature of contextual representation and the mechanisms through which associative information influences visual perception. Consider that while spatial analysis is an inherent part of visual processing and the interpretation of real-life scenes, it may nevertheless rely on different forms of cortical representation than those involved in the analysis of the nonspatial content of these scenes. That is, nonspatial and spatial context effects may in fact be mediated by separate, independent mechanisms. Thus, when viewing a desk in an office, a nonspatial context frame may generate predictions regarding the semantic information in the scene (e.g., the likelihood of certain

objects to appear with the desk), sensitizing the representation of a "lamp" regardless of its spatial position relative to the desk. At the same time, an independent spatial context frame might constrain the spatial relations between these objects, enhancing the representation of the lamp's *location*, regardless of any particular object identity (i.e., based on the knowledge that objects are typically positioned on top of desks and not beneath them).

An alternate view is that context frames contain unified representations for both nonspatial and spatial contextual information. Thus, when viewing a desk in an office, the representation of a desk lamp would *necessarily* include information about its expected location on top of the desk because of the strong connections between this type of spatial and nonspatial knowledge within the contextual schema of an "office." According to this view, the contextual representation of a particular type of scene is maintained as a single unit (i.e., a single context frame). Any violation of the spatial relations between objects in the scene should also impact analysis of the nonspatial associations between these objects (e.g., their expected identities) and the interpretation of the scene as a whole. Thus, seeing a desk lamp *under* a desk would not only violate expectations regarding the most probable *location* of objects in general (as objects typically appear on top of desks), but also violate expectations regarding the specific *identity* or role of the lamp (as desk lamps appear, by definition, on desks). In other words, in a unified context frame, both spatial and nonspatial knowledge jointly contribute to visual comprehension, and thus any spatial inconsistencies between objects may affect the interpretation of these objects' identity, or meaning, altogether.

The hypothesis that nonspatial and spatial contextual associations are represented independently is supported by evidence that an object's identity is typically analyzed through a ventral anatomical route including inferior occipito-temporal cortex, while its spatial location is typically analyzed through more dorsal fronto-parietal brain regions (Goodale and Milner, 1992). The existence of distinct neural systems for analysis of object's identity and spatial location raises the possibility that different cognitive mechanisms underlie facilitation

of object recognition through nonspatial and spatial contextual associations. Indeed, one interpretation of the results of Fenske et al. (submitted) study is that the recognition of contextually related objects is facilitated by a priming mechanism that sensitizes the nonspatial representations of objects that are strongly associated with a contextual cue. In contrast, the results of studies investigating spatial contextual facilitation suggest that recognition of "properly" positioned objects is facilitated by a spatial attention mechanism that enhances the processing of information at a specific location (Hollingworth and Henderson, 1998; Chun and Jiang, 2003; Gordon, 2004). Prior experience with a specific scene or fixed spatial configuration, for example, can guide spatial attention to the location that is most likely to contain a target of interest, thereby facilitating visual search and target recognition within that scene or configuration (Chun and Jiang, 2003).

The two views presented here concerning the underlying structure of context frames and the representation of spatial and nonspatial contextual associations (i.e., unified vs. independent representations) should not necessarily be considered to be mutually exclusive. Indeed, contextual facilitation of object recognition might ultimately be found to involve the top-down influence of both independent and unified representations of spatial and nonspatial contextual associations. Investigating this possibility will require a study that examines the effects of both nonspatial and spatial contextual factors on subsequent object recognition. To determine whether the two types of associative knowledge operate within separate contextual representations, one needs to test whether they can influence perception simultaneously without interacting. Thus, most importantly, an orthogonal design is required in which the two factors are manipulated independently, such that the unique effect of each factor, as well as their joint (interactive) effects on object recognition, can be directly assessed (cf., Sternberg, 2001). To the extent that the two types of associative knowledge are linked within a combined context frame, we would anticipate an interaction in reaction times to target objects. Specifically, one might expect the benefit of identifying an object in a contextually consistent

location relative to a contextually inconsistent location to be significantly greater for objects whose identities have been contextually primed than for contextually unprimed objects, suggesting a unified representation for both nonspatial and spatial factors. In this situation, a similar interaction might also be expected for corresponding cortical activity in brain regions associated with contextual processing (i.e., the PHC and/or RSC) and in object-related processing areas (e.g., fusiform gyrus), with the largest differential cortical activation when the target is consistent with both spatial and nonspatial contextual information. If, however, nonspatial and spatial contextual representations operate independently, their contribution to response latencies, as well as to brain activation, will be additive. In addition, a possible anatomical distinction within the contextual network may be found between the two types of representations (as found previously with spatial and nonspatial associations, Bar and Aminoff, 2003). With such an outcome, one could conclude that the two types of information reside in separate representational "stores," suggesting dissociable influences in context-based top-down facilitation of object recognition. Finally, because spatial and nonspatial contextual associations may involve the top-down influence of both independent and unified representations of spatial and nonspatial contextual associations, it is also possible that the effects of these representations will be additive in context- and object-related regions, but will interact at higher-level brain regions associated with postrecognition processes. Future research is needed to address these issues, and further increase our understanding of the nature of contextual representations and their effect on visual object recognition.

Integrated object- and context-based top-down facilitation of recognition

In this overview, we have described how top-down facilitation of recognition can be achieved either (1) through an object-based mechanism that generates "initial guesses" about an object's identity using rapidly analyzed coarse information about the input image or (2) through the predictive information provided by contextual associations between an object or a scene and the other objects that are likely to appear together in a particular setting. However, it is clear that objects do not appear in isolation, and that both object-based and context-related information is typically available to the visual system during recognition. It therefore makes sense that an optimized system might take advantage of both forms of top-down facilitation to maximize the efficiency of recognition processes. How is this achieved? In this section, we consider how both object-based and context-based influences might operate together to facilitate object recognition.

Central to this discussion is the observation that recognition of an object can be enhanced through neural mechanisms that sensitize the cortical representations of the most likely candidate interpretations of that particular object before information processed through the bottom-up stream has sufficiently accumulated. We propose two key mechanisms through which this "sensitization" is achieved. As reviewed in the first part of this paper, an object-based mechanism is proposed to capitalize on a rapidly derived LSF representation of the object itself to generate "initial guesses" about its identity. The back-projection of these candidate interpretations to object-processing regions thereby sensitizes the corresponding cortical representations in advance of the bottom-up information that continues to activate the correct representation. The context-based mechanism, reviewed in the second part of this chapter, is proposed to sensitize the representations of contextually related objects whose associations are activated through context frames stored in the PHC (see Fig. 9). Importantly, we propose that context frames can be activated following prior recognition of strongly contextual object (as in Fenske et al., submitted), or through early, coarse information about a scene or an object. This includes the possibility that contextual analysis begin promptly, even before a scene or an object is fully recognized. Considered together, the mechanisms for top-down facilitation of object recognition that we describe include one that is based on information from the "to-be-recognized" object itself and the other based on information about the context in which the object appears. Given the intricate links between

objects and the settings in which they appear, these mechanisms typically should not be expected to operate in isolation. Indeed, provided that sufficient information is available for activating the appropriate context frame, and that an LSF image of a single target object is sufficient for limiting its possible interpretations, the intersection of these two sources of information would result in a unique, accurate identification (Bar, 2004). An exciting direction for future research will be to assess how information about additional objects or the scene in which a target object appears may be processed in parallel to increase the opportunities for such valuable intersections of object- and context-based information.

The interactive nature of the two sources of top-down facilitation of object recognition that we have described emerges when the input to either the object- or context-based mechanism is ambiguous. For example, when the coarse object-related input to the prefrontal cortex is ambiguous (e.g., a blurred image of an umbrella also resembles a mushroom, lamp, and parasol), the benefit of having activated the appropriate context frame will be relatively greater than if the LSF profile of an object is completely diagnostic of the object's identity (e.g., a blurred image of a giraffe only resembles a giraffe). In addition, if ambiguous information is projected to the PHC, then this can result in the activation of multiple context frames. From this possibility emerges a rather counterintuitive prediction of our model. Consider, for instance, that a picture of a gun, when projected rapidly in a blurred (i.e., LSF) form to the PHC, may be interpreted also as a drill and a hairdryer (Fig. 11).
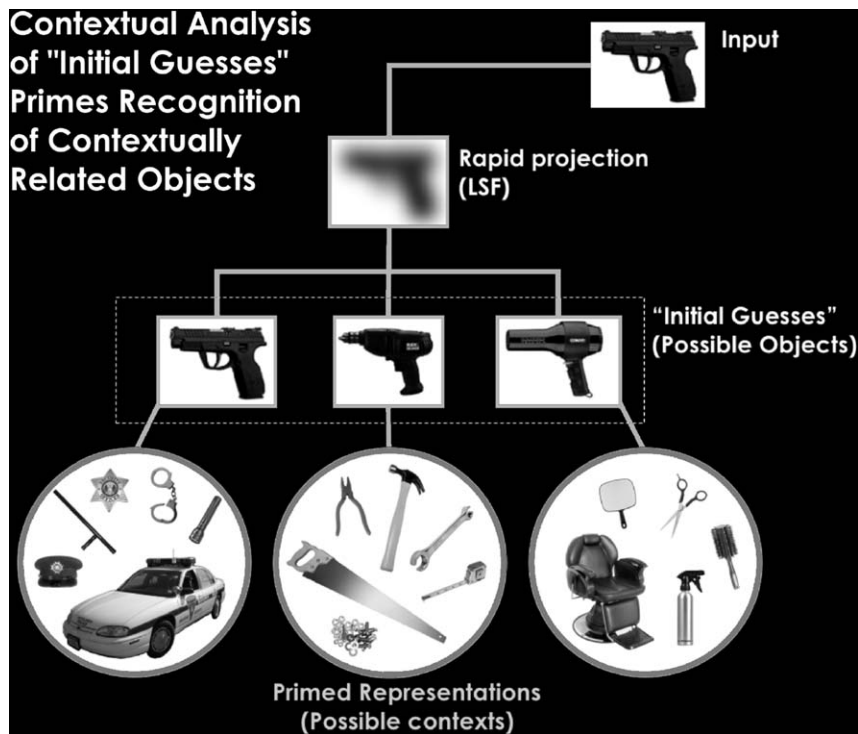

Fig. 11. In the proposed interactive model of object- and context-based top-down down facilitation, ''initial guesses'' about an object's identity are rapidly generated from cursory analysis of its LSF information. These ''initial guesses'' are processed by the context network to help determine the possible context, and thereby facilitate the recognition of other objects that may also be present. Contextual information helps to further constrain the ''initial guesses'' about an object's identity, just as early object-based information helps to determine which context frames should be activated. This interactive model accounts for the otherwise counterintuitive finding that a brief presentation of a gun can facilitate the subsequent recognition of a hairbrush (Fenske and Bar, submitted), despite the lack of any other perceptual, semantic, or contextual relation between these items.

These three objects are associated with three different context frames, and will subsequently trigger the activation of three sets of objects. Consequently, a gun will not only prime the recognition of a police car (i.e., contextual priming), but also the recognition of a hairbrush (i.e., a member of the context frame activated by the hairdryer). Importantly, because there is a complete lack of any other perceptual, semantic, or contextual relation between a gun and a hairbrush, finding significant priming for this condition is best explained through the interactive model of top-down facilitation that we propose. Indeed, we found significant priming for this condition (Fenske and Bar, submitted), and found further that this form of indirect priming existed for relatively short prime exposures (120 ms) but not for longer prime exposures (2400 ms). This finding underscores the interactive nature of the object- and context-based mechanisms we have described, and supports our notion that the arrival of additional information leaves active only the most relevant "initial guess" about an object's identity and only the most relevant context frame.

In conclusion, the models and research findings reviewed here emphasize the importance of top-down facilitation in visual object recognition. Building on the previous work (Bar, 2003), we have examined evidence for an object-based mechanism that rapidly triggers top-down facilitation of visual recognition using a partially analyzed version of the input image (i.e., a blurred image) that generates reliable "initial guesses" about the object's identity in the OFC. Importantly, we have also shown that this form of top-down facilitation does not operate in isolation. Work from our lab indicates that contextual associations between objects and scenes are analyzed by a network including the PHC and the RSC, and that the predictive information provided by these associations can also constrain the "initial guesses" about an objects' identity. We propose that these mechanisms operate together to promote efficient recognition by framing early information about an object within the constraints provided by a lifetime of experience with contextual associations.

Abbreviations

fMRI	functional magnetic resonance imaging
MEG	magnetoencephalography
OFC	orbitofrontal cortex
PFC	prefrontal cortex
PHC	parahippocampal cortex
PPA	parahippocampal place area
RSC	retrosplenial complex
RT	response time

References

Aguirre, G.K., Detre, J.A., Alsop, D.C. and D'Esposito, M. (1996) The parahippocampus subserves topographical learning in man. Cereb. Cortex, 6: 823–829.

Aguirre, G.K., Zarahn, E. and 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. and Bar, M. (submitted for publication). The parahippocampal cortex mediates spatial and non-spatial associations.

Bachevalier, J. and Mishkin, M. (1986) Visual recognition impairment follows ventromedial but not dorsolateral prefrontal lesions in monkeys. Behav. Brain Res., 20: 249–261.

Bar, M. (2003) A cortical mechanism for triggering top-down facilitation in visual object recognition. J. Cogn. Neurosci., 15: 600–609.

Bar, M. (2004) Visual objects in context. Nat. Rev. Neurosci., 5: 617–629.

Bar, M. and Aminoff, E. (2003) Cortical analysis of visual context. Neuron, 38: 347–358.

Bar, M., Kassam, K.S., Ghuman, A.S., Boshyan, J., Schmidt, A.M., Dale, A.M., Hamalainen, M.S., Marinkovic, K., Schacter, D.L., Rosen, B.R. and Halgren, E. (2006) Top-down facilitation of visual recognition. Proceedings of the National Academy of Science, 103: 449–454.

Bar, M., Tootell, R., Schacter, D., Greve, D., Fischl, B., Mendola, J., Rosen, B. and Dale, A. (2001) Cortical mechanisms of explicit visual object recognition. Neuron, 29: 529–535.

Bar, M. and Ullman, S. (1996) Spatial context in recognition. Perception, 25: 343–352.

Barceló, F., Suwazono, S. and Knight, R.T. (2000) Prefrontal modulation of visual processing in humans. Nat. Neurosci., 3: 399–403.

Biederman, I. (1972) Perceiving real-world scenes. Science, 177: 77–80.

Biederman, I. (1981) On the semantic of a glance at a scene. In: Kubovy, M. and Pomerantz, J.R. (Eds.), Perceptual Organization. Erlbaum, Hillsdale, NJ, pp. 213–253.

Biederman, I., Mezzanotte, R.J. and Rabinowitz, J.C. (1982) Scene perception: detecting and judging objects undergoing relational violations. Cogn. Psychol., 14: 143–177.

Biederman, I., Rabinowitz, J.C., Glass, A.L. and Stacy, E.W. (1974) On the information extracted from a glance at a scene. J. Exp. Psychol., 103: 597–600.

Brewer, J.B., Zhao, Z., Desmond, J.E., Glover, G.H. and Gabrieli, J.D. (1998) Making memories: brain activity that predicts how well visual experience will be remembered. Science, 281: 1185–1187.

Bullier, J. and Nowak, L.G. (1995) Parallel versus serial processing: new vistas on the distributed organization of the visual system. Curr. Opin. Neurobiol., 5: 497–503.

Cave, C.B. and Kosslyn, S.M. (1993) The role of parts and spatial relations in object identification. Perception, 22: 229–248.

Chun, M.M. and Jiang, Y. (2003) Implicit, long-term spatial contextual memory. J. Exp. Psychol. Learn. Mem. Cogn., 29: 224–234.

Davachi, L., Mitchell, J. and Wagner, A. (2003) Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. Proc. Natl. Acad. Sci. USA., 100: 2157–2162.

Davenport, J.L. and Potter, M.C. (2004) Scene consistency in object and background perception. Psychol. Sci., 15: 559–564.

Desimone, R. (1998) Visual attention mediated by biased competition in extrastriate visual cortex. Philos. Trans. R. Soc. Lond. B Biol. Sci., 353: 1245–1255.

Düzel, E., Habib, R., Rotte, M., Guderian, S., Tulving, E. and Heinze, H.J. (2003) Human hippocampal and parahippocampal activity during visual associative recognition memory for spatial and nonspatial stimulus configurations. J. Neurosci., 23: 9439–9444.

Engel, A.K., Fries, P. and Singer, W. (2001) Dynamic predictions: oscillations and synchrony in top-down processing. Nat. Rev. Neurosci., 2: 704–716.

Epstein, R., Graham, K.S. and Downing, P.E. (2003) Viewpoint-specific scene representations in human parahippocampal cortex. Neuron, 37: 865–876.

Epstein, R. and Kanwisher, N. (1998) A cortical representation of the local visual environment. Nature, 392: 598–601.

Fenske, M.J., Boshyan, J. and Bar, M. (submitted) Can a gun prime a hairbrush? The "initial guesses" that drive top-down contextual facilitation of object recognition. Paper presented at the 5th Annual Meeting of the Vision Sciences Society, Sarasota, FL.

Fenske, M.J. and Bar, M. (submitted for publication). Can A Gun Prime A Hairbrush? "Initial Guesses" that Mediate Contextual Facilitation of Object Recognition.

Freedman, D.J., Riesenhuber, M., Poggio, T. and Miller, E.K. (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. Science, 291: 312–316.

Friston, K. (2003) Learning and inference in the brain. Neural Networks, 16: 1325–1352.

Gibson, E.J. (1969) Principles of Perceptual Learning and Development. Appleton-Century-Crofts, New York.

Gilbert, C.D., Sigman, M. and Crist, R.E. (2001) The neural basis of perceptual learning. Neuron, 31: 681–697.

Goodale, M.A. and Milner, A.D. (1992) Separate visual pathways for perception and action. Trends Neurosci, 15: 20–25.

Gordon, R.D. (2004) Attentional allocation during the perception of scenes. J. Exp. Psychol. Hum. Percept. Perform., 30: 760–777.

Grabowska, A. and Nowicka, A. (1996) Visual-spatial-frequency model of cerebral asymmetry: a critical survey of behavioral and electrophysiological studies. Psychol. Bull., 120: 434–449.

Grill-Spector, K., Kourtzi, Z. and Kanwisher, N. (2001) The lateral occipital complex and its role in object recognition. Vision Res., 41: 1409–1422.

Grossberg, S. (1980) How does a brain build a cognitive code? Psychol. Rev., 87: 1–51.

Henderson, J.M. (1992) Identifying objects across saccades: effects of extrafoveal preview and flanker object context. J. Exp. Psychol. Learn. Mem. Cogn., 18: 521–530.

Henderson, J.M., Pollatsek, A. and Rayner, K. (1987) Effects of foveal priming and extrafoveal preview on object identification. J. Exp. Psychol. Hum. Percept. Perform., 13: 449–463.

Henson, R.N. (2003) Neuroimaging studies of priming. Prog. Neurobiol., 70: 53–81.

Hock, H.S., Romanski, L., Galie, A. and Williams, C.S. (1978) Real-world schemata and scene recognition in adults and children. Mem. Cogn., 6: 423–431.

Hollingworth, A. and Henderson, J.M. (1998) Does consistent scene context facilitate object perception? J. Exp. Psychol. Gen., 127: 398–415.

Hopfinger, J.B., Buonocore, M.H. and Mangun, G.R. (2000) The neural mechanisms of top-down attentional control. Nat. Neurosci., 3: 284–291.

Humphreys, G.W., Riddoch, M.J. and Price, C.J. (1997) Top-down processes in object identification: evidence from experimental psychology, neuropsychology and functional anatomy. Philos. Trans. R. Soc. Lond. B Biol. Sci., 352: 1275–1282.

Jackson, O. and Schacter, D.L. (2004) Encoding activity in anterior medial temporal lobe supports subsequent associative recognition. Neuroimage, 21: 456–462.

Kirwan, C.B. and Stark, C.E. (2004) Medial temporal lobe activation during encoding and retrieval of novel face-name pairs. Hippocampus, 14: 919–930.

Kosslyn, S.M. (1994) Image and Brain. MIT Press, Cambridge, MA.

Kosslyn, S.M., Alpert, N.M. and Thompson, W.L. (1995) Identifying objects at different levels of hierarchy: a positron emission tomography study. Hum. Brain Mapping, 3: 107–132.

Kosslyn, S.M., Alpert, N.M., Thompson, W.L., Chabris, C.F., Rauch, S.L. and Anderson, A.K. (1993) Visual mental imagery activates topographically organized visual cortex: PET investigations. J. Cogn. Neurosci., 5: 263–287.

Koutstaal, W., Wagner, A.D., Rotte, M., Maril, A., Buckner, R.L. and Schacter, D.L. (2001) Perceptual specificity in visual object priming: functional magnetic resonance imaging evidence for a laterality difference in fusiform cortex. Neuropsychologia, 39: 184–199.

Lee, T.S. and Mumford, D. (2003) Hierarchical Bayesian inference in the visual cortex. J. Opt. Soc. Am., 20: 1434–1448.

Leonova, A., Pokorny, J. and Smith, V.C. (2003) Spatial frequency processing in inferred PC- and MC-pathways. Vision Res., 43: 2133–2139.

Levy, I., Hasson, U., Avidan, G., Hendler, T. and Malach, R. (2001) Center-periphery organization of human object areas. Nat. Neurosci., 4: 533–539.

Logothetis, N.K., Leopold, D.A. and Sheinberg, D.L. (1996) What is rivalling during binocular rivalry? Nature, 380: 621–624.

Lustig, C. and Buckner, R.L. (2004) Preserved neural correlates of priming in old age and dementia. Neuron, 42: 865–875.

Maccotta, L. and Buckner, R.L. (2004) Evidence for neural effects of repetition that directly correlate with behavioral priming. J. Cogn. Neurosci., 16: 1625–1632.

Maguire, E.A. (2001) The retrosplenial contribution to human navigation: a review of lesion and neuroimaging findings. Scand. J. Psychol., 42: 225–238.

Malach, R., Levy, I. and Hasson, U. (2002) The topography of high-order human object areas. Trends Cogn. Sci., 6: 176–184.

Mandler, J.M. and Johnson, N.S. (1976) Some of the thousand words a picture is worth. J. Exp. Psychol. [Hum. Learn.], 2: 529–540.

Martin, A., Wiggs, C.L., Ungerleider, L.G. and Haxby, J.V. (1996) Neural correlates of category-specific knowledge. Nature, 379: 649–652.

Mechelli, A., Price, C.J., Friston, K.J. and Ishai, A. (2004) Where bottom-up meets top-down: neuronal interactions during perception and imagery. Cereb. Cortex, 14: 1256–1265.

Merigan, W.H. and Maunsell, J.H. (1993) How parallel are the primate visual pathways? Annu. Rev. Neurosci., 16: 369–402.

Metzger, R.L. and Antes, J.R. (1983) The nature of processing early in picture perception. Psychol. Res., 45: 267–274.

Miller, M.B. and Gazzaniga, M.S. (1998) Creating false memories for visual scenes. Neuropsychologia, 36: 513–520.

Miyashita, Y. and Hayashi, T. (2000) Neural representation of visual objects: encoding and top-down activation. Curr. Opin. Neurobiol., 10: 187–194.

Morcom, A.M., Good, C.D., Frackowiak, R.S. and Rugg, M.D. (2003) Age effects on the neural correlates of successful memory encoding. Brain, 126: 213–229.

Oenguer, D. and Price, J.L. (2000) The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. Cereb. Cortex, 10: 206–219.

Oliva, A. and Torralba, A. (2001) Modeling the shape of a scene: a holistic representation of the spatial envelope. Int. J. Comp. Vis., 42: 145–175.

Palmer, S.E. (1975) The effects of contextual scenes on the identification of objects. Mem. Cogn., 3: 519–526.

Parker, A., Wilding, E. and Akerman, C. (1998) The Von Restorff effect in visual object recognition memory in humans and monkeys. The role of frontal/perirhinal interaction. J. Cogn. Neurosci., 10: 691–703.

Pascual-Leone, A. and Walsh, V. (2001) Fast backprojections from the motion to the primary visual area necessary for visual awareness. Science, 292: 510–512.

Pihlajamaki, M., Tanila, H., Kononen, M., Hanninen, T., Hamalainen, A., Soininen, H. and Aronen, H.J. (2004) Visual presentation of novel objects and new spatial arrangements of objects differentially activates the medial temporal lobe subareas in humans. Eur. J. Neurosci., 19: 1939–1949.

Ranganath, C., DeGutis, J. and D'Esposito, M. (2004a) Category-specific modulation of inferior temporal activity during working memory encoding and maintenance. Brain Res. Cogn. Brain Res., 20: 37–45.

Ranganath, C., Yonelinas, A.P., Cohen, M.X., Dy, C.J., Tom, S.M. and D'Esposito, M. (2004b) Dissociable correlates of recollection and familiarity within the medial temporal lobes. Neuropsychologia, 42: 2–13.

Rempel-Clower, N.L. and Barbas, H. (2000) The laminar pattern of connections between prefrontal and anterior temporal cortices in the rhesus monkey is related to cortical structure and function. Cereb. Cortex, 10: 851–865.

Schacter, D.L. and Buckner, R.L. (1998) Priming and the brain. Neuron, 20: 185–195.

Schyns, P.G. and Oliva, A. (1994) From blobs to boundary edges: evidence for time- and spatial-dependent scene recognition. Psychol. Sci., 5: 195–200.

Shapley, R. (1990) Visual sensitivity and parallel retinocortical channels. Annu. Rev. Psychol., 41: 635–658.

Simons, J.S., Koutstaal, W., Prince, S., Wagner, A.D. and Schacter, D.L. (2003) Neural mechanisms of visual object priming: evidence for perceptual and semantic distinctions in fusiform cortex. Neuroimage, 19: 613–626.

Sternberg, S. (2001) Separate modifiability, mental modules, and the use of pure and composite measures to reveal them. Acta Psychol. (Amst.), 106: 147–246.

Tanaka, K. (1996) Inferotemporal cortex and object vision. Annu. Rev. Neurosci., 19: 109–139.

Ullman, S. (1995) Sequence seeking and counter streams: A computational model for bidirectional information flow in the visual cortex. Cereb. Cortex, 1: 1–11.

Valenstein, E., Bowers, D., Verfaellie, M., Heilman, K.M., Day, A. and Watson, R.T. (1987) Retrosplenial amnesia. Brain, 110: 1631–1646.

Vogt, B.A., Vogt, L.J., Perl, D.P. and Hof, P.R. (2001) Cytology of human caudomedial cingulate, retrosplenial, and caudal parahippocampal cortices. J. Comp. Neurol., 438: 353–376.

Vuilleumier, P., Henson, R.N., Diver, J. and Dolan, R.J. (2002) Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. Nat. Neurosci., 5: 491–499.

Wagner, A.D., Schacter, D.L., Rotte, M., Koutstaal, W., Maril, A., Dale, A.M., Rosen, B.R. and Buckner, R.L. (1998) Building memories: remembering and forgetting of verbal experiences as predicted by brain activity. Science, 281: 1188–1191.

Wheeler, M.E. and Buckner, R.L. (2003) Functional dissociation among components of remembering: control, perceived oldness, and content. J. Neurosci., 23: 3869–3880.

Wilson, F.A.W., Scalaidhe, S.P.O. and Goldman-Rakic, P.S. (1993) Dissociation of object and spatial processing domains in primate prefrontal cortex. Science, 260: 1955–1958.

Zago, L., Fenske, M.J., Aminoff, E. and Bar, M. (2005) The rise and fall of priming: how visual exposure shapes cortical representations of objects. Cereb. Cortex, 15: 1655–1665.