

Object Recognition

Insights From Advances in fMRI Methods

Kalanit Grill-Spector and Rory Sayres

Stanford University

ABSTRACT—Recent advancements in imaging methods and analysis approaches have provided important insights about the neural bases of object recognition. We address the potential limitations of standard functional magnetic resonance imaging (fMRI) and discuss methodological advancements, including fMRI-adaptation, pattern analyses, and high-resolution fMRI, that may be more appropriate for studying object and face representations. fMRI-adaptation and high-resolution fMRI measure responses of neural subpopulations within standard fMRI voxels, and pattern analyses examine the information in the distributed activations across voxels, which may differ from the mean response across these voxels. These methods have provided evidence for a multitude of representations across the human ventral stream that provide empirical constraints for cognitive theories of recognition.

KEYWORDS—object recognition; face recognition; fMRI-adaptation; pattern analyses; high-resolution fMRI.

Object recognition is an amazing human feat. Humans effortlessly recognize objects within a fraction of a second. Variability in viewing conditions can cause changes in objects' location, size, viewpoint, illumination, occlusion, etc., yet recognition is largely invariant to these dramatic changes in objects' appearance. The underlying neural representations and computations that enable this to happen, however, are still mysterious.

Neuroimaging techniques, especially functional magnetic resonance imaging (fMRI) offer an opportunity to investigate neural mechanisms underlying object recognition. Object recognition is thought to involve the visual “ventral stream,” a processing pathway extending from the occipital lobe to the temporal lobe. Indeed, fMRI studies have revealed a constellation of object-selective regions in the lateral and ventral occipito-temporal cortex (known as the lateral occipital complex, or LOC; Malach et al., 1995; Fig. 1a) that respond more

strongly to objects than to noise patterns, textures, or scrambled objects (created by breaking object images into small squares and shuffling the locations of these squares) and several face-selective regions that respond more strongly to faces than to nonface objects (Fig. 1a). The most studied of these face-selective regions is the fusiform face area (FFA; Kanwisher, McDermott, & Chun, 1997). Activation in both the LOC and FFA correlates with subjects' recognition performance (Fig. 1b). When stimuli are ambiguous or presented close to the perceptual threshold (i.e., the border between seeing and not seeing), activity in face- and object-selective cortical areas is higher when subjects recognize the stimuli than when they do not (Grill-Spector, Knouf & Kanwisher, 2004; Grill-Spector, Kushnir, Hendler, & Malach, 2000; Tong, Nakayama, Vaughan, & Kanwisher, 1998). Further, prolonged experience with objects (Gauthier, Skudlarski, Gore, & Anderson, 2000; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999) and faces (Golarai et al., 2007) modifies these representations.

These studies provide fundamental knowledge about which brain regions are involved in object and face recognition. However, many questions about the nature of the underlying neural representations remain:

- How do these representations provide for invariant recognition without losing the acuity to distinguish between similar items?
- What mechanisms underlie rapid categorization (e.g., discrimination between a car and a face)?
- Are the same representations used for between-category discrimination (e.g., distinguishing between a car and a boat) and within-category discrimination (e.g., distinguishing between a sailboat and a motorboat)?
- Are representations of faces (or objects of expertise) fundamentally different from other object representations, and if so, how do they differ?
- How do experience and development shape these representations?

In addressing these questions, the inherent limitations of standard-resolution fMRI (SR-fMRI) pose significant challenges. SR-fMRI measures the pooled neural responses across a voxel

Address correspondence to Kalanit Grill-Spector, Department of Psychology, Jordan Hall, Stanford University, Stanford, CA 94305; e-mail: kalanit@psych.stanford.edu.

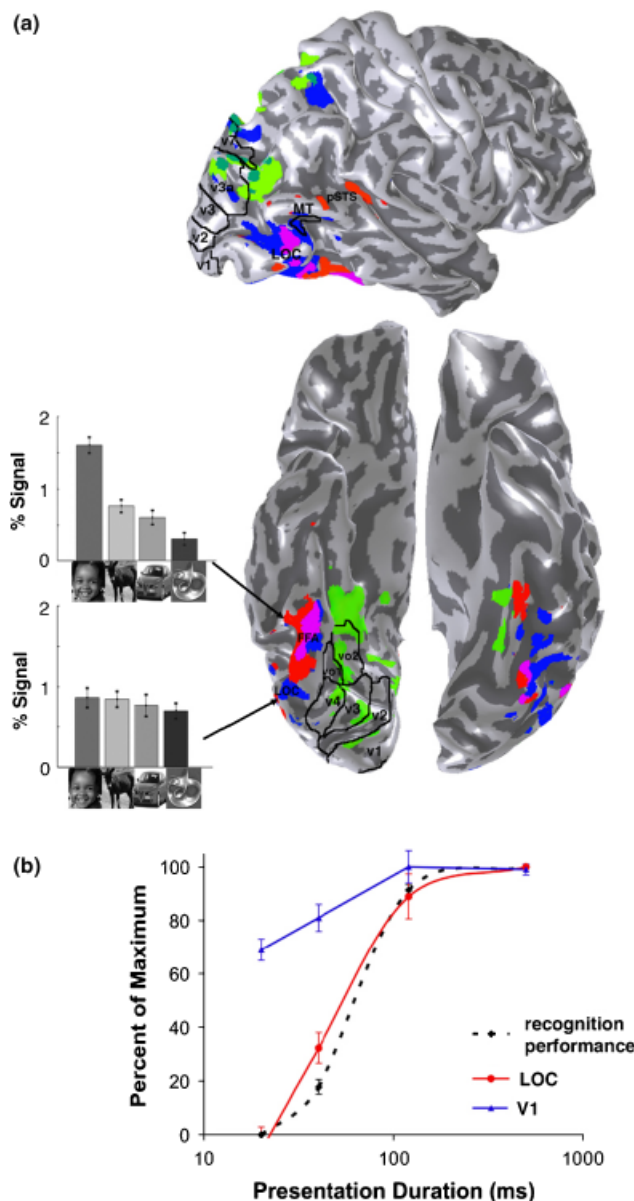


Fig. 1. Activation in the human ventral stream. Colored patches on the right lateral (a, top) and ventral (a, bottom) views of the cortical surface of one representative subject show object-, face-, and place-selective regions. Red = regions that respond more to faces than to nonface objects, green = regions that respond more to places than to nonface objects, blue = regions that respond to nonface objects more than to scrambled objects, pink = regions that respond both more to faces than to objects and more to objects than to places than to objects and more to objects than to scrambled objects. Boundaries of retinotopic visual areas are shown in black on the right hemisphere. The lateral occipital complex (LOC) includes object-selective regions (blue and pink regions) in the lateral and ventral occipital cortex. Face-selective regions include the fusiform face area (FFA, red and pink regions in the fusiform gyrus); a face-selective region (shown in red) in the posterior superior temporal sulcus (pSTS) and a face-selective region that overlaps the LOC (shown in pink). MT is a motion-selective region in the posterior end of the inferior temporal sulcus. Bar graphs on the left show the mean percentage signal from the LOC (excluding regions that respond more to faces than to objects; bottom bar graph) and FFA (top bar graph) from four kinds of stimuli, averaged across 10 subjects, relative to a scrambled-object baseline. Stimuli (indicated by icons) were faces, four-legged mammals, cars, and abstract sculptures. Object-selective regions respond to both objects and faces more than to the scrambled baseline, but they do not respond more strongly to faces than to objects. fMRI signal and recognition performance as a function of image exposure is shown in the bottom graph (b). During scanning, subjects viewed photographs of faces, animals, and common objects (e.g., cars, boats), which were followed by scrambled images. Subjects were asked to name pictures covertly at the basic level (e.g., face, dog, car, boat). In different blocks, images appeared for a different duration (20, 40, 120, or 500 milliseconds, ms) followed by a scrambled image completing a 500-ms trial (480, 460, 380, or 0-ms mask duration). Immediately after fMRI scans (but while in the scanner), subjects saw the same images at the same duration and participated in a behavioral experiment in which recognition performance was recorded. Recognition performance (dashed black line) is correlated with LOC activation (red) but not with V1 (primary visual cortex) activation (blue). Adapted from “The Dynamics of Object-Selective Activation Correlate With Recognition Performance in Humans,” by K. Grill-Spector, T. Kushnir, T. Hendler, & R. Malach, 2000, *Nature Neuroscience*, 3, pp. 837–843, Figure 3. Copyright 2000, Nature Publishing Group. Adapted with permission.

(a three-dimensional volume element analogous to a pixel in a two-dimensional digital image) or many voxels that constitute a brain region. Since a cubic millimeter (mm) of cortex contains about 10,000 to 50,000 neurons and fMRI signals are typically measured across regions that are in the range of 50 to 5,000 cubic mm in size, those regions reflect the activity of about a

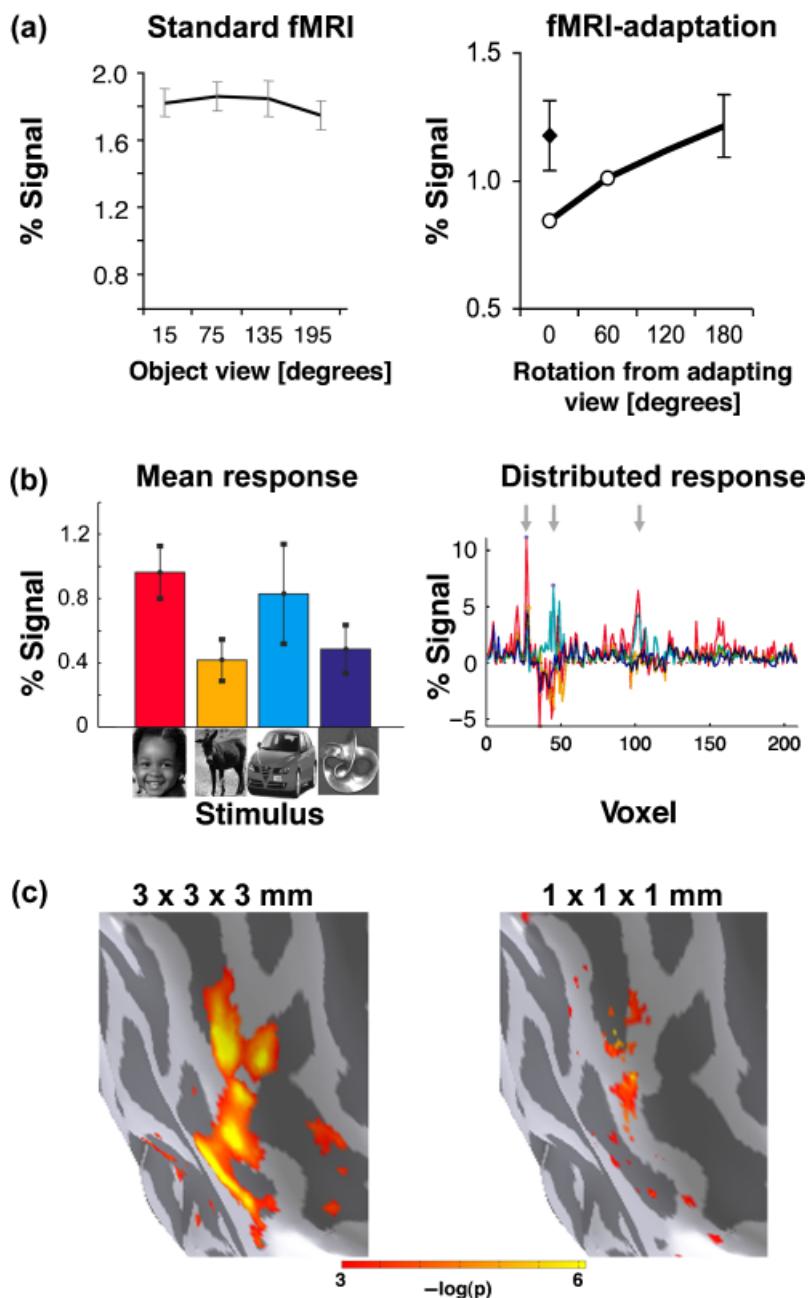
Fig. 2. fMRI methods used to study the properties of cortical regions involved in object recognition. The top (a) left graph shows the mean lateral occipital complex (LOC) response to animals presented in different degrees of rotation (vs. a scrambled baseline), averaged across 7 subjects, using standard fMRI. The top (a) right graph shows LOC responses, averaged across the same 7 subjects, during an fMRI-adaptation experiment in which we measured sensitivity of LOC responses to object rotation. In each fMRI-adaptation trial, two images of the same animal were shown either in the same view (0 rotation) or in different views that were rotated 60, 120, or 180 degrees away from the initial view. The black diamond indicates responses to trials containing two different animals presented in the same view (unadapted response); open circles indicate significant adaptation ($p < 0.05$) when two images of the same animal were shown compared to when two different animals were shown. Note that there was adaptation when the same animal and view were repeated and also when the same animal was shown at a view that was rotated 60 degrees away from the initial view, but that there was no significant adaptation when the same animal was shown rotated 120 or 180 away from the initial view (Andresen & Grill-Spector, 2006). The middle graphs (b) show pattern analyses. The mean LOC response from a single subject to different object categories (shown in the icons), relative to a scrambled-object baseline, is shown in the middle (b) left graph. The distributed response across the LOC to these categories, from the same subject, is shown in the middle (b) right graph. The x-axis represents voxels (in this subject, the LOC contained about 200 voxels); the y-axis represents percent signal change relative to the scrambled-object baseline. Each line represents the distributed response to one category: red = faces; yellow = animals; light blue = cars; and dark blue = abstract objects. Each category elicited a different pattern of response across the LOC. The bottom (c) images show face-selective activations in one subject on the cortical surface around the fusiform gyrus. The left panel shows activations measured with standard-resolution fMRI (voxels of $3 \times 3 \times 3$ mm); the right panel shows the same activations as measured with high-resolution fMRI (voxels of $1 \times 1 \times 1$ mm). The latter shows more localized face-selective activations. Activation maps were thresholded at a significance of $p < 0.001$.

million to a few hundred million neurons. Given results from electrophysiology studies, it is highly likely that such measurements pool across diverse neural populations that process different aspects of objects.

Evidence from electrophysiology suggests that a fine-grain neural organization (in the order of a half to a few mm) may be particularly relevant for understanding the neural representations of faces and objects. Electrophysiological recordings from the inferotemporal (IT) cortex of monkeys (a region that contains neurons responding to objects and shapes) suggest that neurons with similar properties—for example, similar shape preference (Fujita, Tanaka, Ito, Cheng, 1992) or face preference (Tsao, Freiwald, Tootell, & Livingstone, 2006)—are physically clustered

and may be organized in a columnar structure (Fujita et al., 1992). Columnar structure is prevalent throughout the visual cortex and is thought to be a basic computational unit in the cortex. In monkeys, columns are 50 to 500 micrometers (μm) in diameter, and in humans they are thought to be twice as large (Adams, Sincich, & Horton, 2007). Thus, directly imaging human brain structures at the millimeter or sub-millimeter scale may reveal the computational units underlying object and face recognition.

Here we review recent methodological advances in fMRI that elucidate the neural representations underlying recognition at a finer scale than standard fMRI. These include the method of fMRI-adaptation (fMRI-A), pattern analysis (PA), and high-resolution fMRI (HR-fMRI; see Fig. 2). Each of these approaches



examines the underlying representations in a different way: fMRI-A relies on adaptation of neural populations within a voxel by repeated stimulus presentations, PA capitalizes on small but systematic modulations across voxels, and HR-fMRI images the cortex directly at a substantially higher resolution than standard resolution does. In particular, we demonstrate how these methods have advanced cognitive theories of object recognition.

fMRI-ADAPTATION

Repeated presentation of objects or faces typically leads to a reduced fMRI response—or fMRI adaptation (fMRI-A)—in object- and face-selective cortical regions (Grill-Spector et al., 1999). This reduction is stimulus specific, as adaptation is maximal for the exact repetition of a stimulus but is lesser (and may disappear) for transformed versions of the item (such as the same object presented in different locations; Grill-Spector et al., 1999). fMRI-A is a robust phenomenon that is present even when repetitions of the same object occur after many intervening stimuli and across various time scales ranging from seconds to days. It is hypothesized that this reduced response measured by fMRI is related to the stimulus-specific decrements in the firing rates of IT neurons to repeated objects. However, the exact relation between neural adaptation and fMRI-A is not well understood (Grill-Spector et al., 2006).

fMRI-A allows tagging of specific neuronal populations within a voxel by adapting them with a repeated stimulus and then examining the sensitivity of the representation to different factors by testing which factors lead to a recovery from adaptation. For example, SR-fMRI of the LOC shows a similar level of response to different views of the same object (Fig. 2a, left). This may reflect a homogeneous population of neurons that responds equally across object views, or it may reflect a heterogeneous population of neurons within a voxel, each tuned to a specific object view. Using fMRI-A, these possibilities can be distinguished. Recovery from adaptation indicates that the LOC is sensitive to the object's view (as shown in Fig. 2a, right). Finding adaptation across object views would have indicated of view invariance.

fMRI-A has been used to examine theories of view-invariant object recognition. Some theories of object recognition suggest that objects are represented by a set of relatively simple, view-invariant features and their spatial relationships. Other theories suggest that object representations are view-dependent and that invariant recognition is accomplished by interpolation (or by a distributed code) across neurons that represent specific object views. Results of fMRI-A experiments provide evidence for view-dependent cortical representations (Grill-Spector et al., 1999; Vuilleumier, Henson, Driver, & Dolan, 2002). fMRI-A studies have also demonstrated a hierarchy of representations in the LOC, where more posterior regions recover from adaptation (and therefore are sensitive) to changes in object size, position, illumination, and view, and more anterior regions along the

occipito-temporal sulcus and fusiform gyrus are adapted when object position and size are changed but recover from adaptation when object view and illumination are changed (Grill-Spector et al., 1999). In addition, there is evidence for hemispheric differences: The left fusiform gyrus remains adapted when the same object is shown in a view different from the adapting view, whereas the right hemisphere recovers from adaptation (Vuilleumier et al., 2002), suggesting higher sensitivity to object view in the right hemisphere. Overall, results of fMRI-A studies suggest that any theory of invariant object recognition should include a hierarchy of processing stages involving intermediate stages that contain view-dependent representations and higher stages containing more abstract object representations.

Another focus of intensive study is on the neural mechanisms that enable face identification (i.e., distinguishing between individuals, such as Bill Clinton vs. George Bush). This research has focused on characterizing the response properties of face-selective regions in the fusiform gyrus (FFA) whose activity correlates with face-identification performance (Grill-Spector, Knouf, & Kanwisher, 2004). Rotshtein and colleagues (Rotshtein, Henson, Treves, Driver, & Dolan, 2005) showed that the FFA recovers from adaptation when the perceived identity of a face changes (e.g., Marilyn Monroe morphing into Margaret Thatcher) but remains adapted when the physical changes do not change the perceived identity, suggesting that the FFA's sensitivity to facial identity corresponds to face perception. Two other studies (Gilaie-Dotan & Malach, 2007; Jiang et al., 2006) reported a systematic recovery from fMRI-A as the morph level increased and the face became more dissimilar from the original face. However, they reported that small morphing levels that were not strong enough to be perceived as a different face nonetheless caused recovery from adaptation in the FFA. While these studies vary in their findings on the degree to which FFA's sensitivity to changes in face identity corresponds to perceived changes in identity, these three studies provide evidence for a population code for face identification—that is, a model in which the representation of a particular face is distributed across many neurons and the level of neural response depends on the similarity between the input face and the selectivity of particular neurons. However, a study by Loffler and colleagues (Loffler, Yourganov, Wilkinson, & Wilson, 2005) suggests a different code for face identification in which the FFA codes how much an input face is different from a mean (prototypical) face (i.e., a norm-based model), because they found that FFA responses were higher for atypical faces than for prototypical faces.

PATTERN ANALYSES

Many studies have reported regions in the brain specialized for processing specific object categories such as faces, places, body parts, animals, tools, and letter strings. These findings of regions

specialized for particular stimuli are highly reproducible and are consistent with previous neuropsychological evidence that damage to the human ventral stream can lead to specific deficits in object and face recognition.

The presence of category-preferring regions poses a theoretical problem: Is there a special region in the brain that codes for every category that humans recognize? Kanwisher and colleagues propose that the brain contains a few domain-specific modules, specialized for processing faces, places, and body parts, while other objects are processed and recognized by a general-purpose object-recognition system (e.g., the LOC). According to this model, activation of domain-specific modules is tightly linked to recognition of these specialized categories and items from these categories. Another model, posited by Haxby and colleagues (Haxby et al., 2001), suggests that a single system recognizes both faces and nonface objects. According to this model, objects and faces are represented by highly distributed and overlapping patterns of brain activity across the ventral stream and this distributed activity codes the categories of objects. This hypothesis builds upon connectionist models, which showed that a distributed code provides a larger representational capacity than does a domain-specific architecture and thus allows for the representation of a large number of objects and categories.

Proponents of each of the above models have provided evidence for their view by employing pattern analyses (PAs). PAs examine the distributed pattern of response across voxels instead of measuring the average response across voxels, as is typical in many fMRI studies (Fig. 2b). To support their model of distributed object representations, Haxby and colleagues performed PA of the distributed responses across face- and object-selective regions in the ventral stream (Haxby et al., 2001). They showed that the distributed response across the ventral stream to particular object categories is reproducible and that different categories produce distinct activation patterns. Therefore, it is possible to decode which category subjects viewed from the distributed response pattern. Subsequent studies showed that the distributed response for a certain category is consistent across sessions (Cox & Savoy, 2003) and extends to other exemplars of a category and to objects presented in a different format (e.g., line drawings vs. photographs; Spiridon & Kanwisher, 2002).

Nevertheless, since the method of PA is new, the theoretical foundations and limitations should also be considered. PAs are useful as long as there are reliable variations across voxels for the parameter being measured. Since each SR-fMRI voxel reflects the pooled response across millions of neurons, the variation of a voxel's response for a parameter may not be reliable or measurable. For example, there is evidence from fMRI-A studies that FFA responses are sensitive to face identity, but surprisingly there are no PA fMRI studies showing that face identity can be decoded from distributed FFA responses to faces. Thus, PAs of HR-fMRI data may be necessary to resolve the representations underlying face identification. Another caveat

is that while PAs reveal what information is present in the distributed response, they do not show whether the brain uses the entire distributed activity during object recognition or just uses information present in localized regions (but see Williams, Dang, & Kanwisher, 2007).

HIGH-RESOLUTION fMRI

More recently, researchers have used high-resolution fMRI (HR-fMRI) to directly image the fine-grain structure in object- and face-selective regions of the cortex. SR-fMRI studies have used voxels ranging between $3 \times 3 \times 3$ mm and $3.75 \times 3.75 \times 5$ mm to provide coverage of the entire brain. HR-fMRI uses voxels between $1 \times 1 \times 1$ mm and $1.6 \times 1.6 \times 2$ mm across parts of the brain, providing less coverage but a substantial increase in the imaging resolution by a factor of about 27 to 70 (Fig. 2c). One of the goals of HR-fMRI is to image the cortex at a higher level of spatial precision and to distinguish heterogeneous responses that may be averaged with SR-fMRI. We believe that the spatial scale of the representation matters. For example, examination of face-selective responses with HR-fMRI shows that there are several small clusters in the fusiform gyrus that respond to faces more than to objects; SR-fMRI, in contrast, shows a larger and more diffuse region with higher responses to faces than to objects (Fig. 2c).

HR-fMRI may improve the signal because it images more localized cortical responses and there is less mixing between different types of tissues such as gray matter and white matter (i.e., partial voluming). The hope is that, with HR-fMRI, brain measurements will be closer to the scale of cortical representations. For example, if there are columns tuned to objects or object features, as suggested by electrophysiology experiments, perhaps HR-fMRI can be used to directly image these representations. Whether or not this will happen is an empirical question and depends on several factors including

- The size of the point-spread function of the fMRI signal (i.e., how spatially diffuse the signal is relative to the localized neural response). Current estimates of the point-spread function of the fMRI signal are on the order of 2 mm.
- The spatial scale of the underlying representations. Are there clusters of neurons that code for features, objects, or categories? If so, what are the dimensions of these clusters?
- The cortical distance between neural clusters tuned to similar properties. Clusters may be less than 1 mm in size, but if the distance between clusters of similar properties is of the order of a few mm, it may be possible to resolve them with HR-fMRI.
- The signal-to-noise ratio (SNR) of HR-fMRI. SNR is thought to decrease as the voxel size decreases because the overall signal decreases yet the level of noise does not. Further, HR-fMRI may be particularly sensitive to noise factors such as subject motion. However, it is necessary to validate these

assumptions: The signal in HR-fMRI voxels may not decrease much compared to SR-fMRI, because there is less partial voluming with HR-fMRI and thus it potentially measures responses across more homogeneous neural populations. The SNR may be further improved by using higher-field scanners and by advancements in fMRI protocols and coil designs.

SUMMARY

Recent advancements in imaging methods and analysis approaches have provided fundamental constraints for cognitive theories of recognition and have offered important insights about the nature of cortical representations of faces and objects. We believe that studying these representations at the appropriate spatial scale is important. We addressed the limitations of SR-fMRI to study the neural bases of object recognition. First, SR-fMRI may not provide the optimal spatial resolution for investigating the underlying representations. We suggested methods for approaching a resolution that may be more appropriate for studying object and face representations. Second, methods that examine the mean response across a cortical region may overlook important components of the representation; information present in the distributed response may be different than that present in the average response. We believe that future directions that utilize multiple advancements, such as a combination of fMRI-A, HR-fMRI, and PA, will be the most fruitful for progress in understanding the neural representation of objects and faces and for providing empirical constraints for cognitive theories.

Recommended Reading

- Grill-Spector, K., Henson, R., & Martin, A. (2006). (See References). Discusses the relation between fMRI-A and neural responses.
- Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, *107*, 293–321. Discusses the method of fMRI-Adaptation.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., & Pietrini, P. (2001). (See References). Used PA to show distributed and overlapping responses to object categories across the human ventral stream.
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*, 679–685. An article showing why PA may provide for increased spatial resolution.
- Sawamura, H., Orban, G., & Vogels, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: A single cell study of the fMRI-adaptation paradigm. *Neuron*, *49*, 307–318. Examines the relation between fMRI-A and neural responses.
-

Acknowledgments—We thank Nicolas Davidenko, Golijeh Golarai, Nathan Witthoft, and Jennifer Yoon for their comments

on this manuscript and fruitful discussions. This research has been supported by National Institutes of Health Grant R21 EY016199-01 and Whitehall Foundation Grant 2005-05-111-RES to KGS.

REFERENCES

- Adams, D.L., Sincich, L.C., & Horton, J.C. (2007). Complete pattern of ocular dominance columns in human primary visual cortex. *Journal of Neuroscience*, *39*, 10391–10403.
- Andresen, D., & Grill-Spector, K. (2006, May). *View sensitivity of object representations in human object-selective visual cortex*. Paper presented at the annual meeting of the Vision Science Society, Sarasota, FL.
- Cox, D.D., & Savoy, R.L. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage*, *19*, 261–270.
- Fujita, I., Tanaka, K., Ito, M., & Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature*, *360*, 343–346.
- Gauthier, I., Skudlarski, P., Gore, J.C., & Anderson, A.W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, *3*, 191–197.
- Gauthier, I., Tarr, M.J., Anderson, A.W., Skudlarski, P., & Gore, J.C. (1999). Activation of the middle fusiform “face area” increases with expertise in recognizing novel objects. *Nature Neuroscience*, *2*, 568–573.
- Gilaie-Dotan, S., & Malach, R. (2007). Sub-exemplar shape tuning in human face-related areas. *Cerebral Cortex*, *2*, 325–338.
- Golarai, G., Ghahremani, D.G., Whitfield-Gabrieli, S., Reiss, A., Eberhardt, J.L., Gabrieli, J.D., et al. (2007). Differential development of high-level visual cortex correlates with category-specific recognition memory. *Nature Neuroscience*, *10*, 512–522.
- Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, *10*, 14–23.
- Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience*, *7*, 555–562.
- Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzhak, Y., & Malach, R. (1999). Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, *24*, 187–203.
- Grill-Spector, K., Kushnir, T., Hendler, T., & Malach, R. (2000). The dynamics of object-selective activation correlate with recognition performance in humans. *Nature Neuroscience*, *3*, 837–843.
- Grill-Spector, K., Sayres, R., & Ress, D. (2006). High-resolution imaging reveals highly selective nonface clusters in the fusiform face area. *Nature Neuroscience*, *9*, 1177–1185. Erratum in: *Nature Neuroscience*, *10*, 133.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 2425–2430.
- Jiang, X., Rosen, E., Zeffiro, T., Vanmeter, J., Blanz, V., & Riesenhuber, M. (2006). Evaluation of a shape-based model of human face discrimination using fMRI and behavioral techniques. *Neuron*, *50*, 159–172.

- 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.
- Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H.R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience*, *8*, 1386–1390.
- Malach, R., Reppas, J.B., Benson, R.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, USA*, *92*, 8135–8139.
- Rotshtein, P., Henson, R.N., Treves, A., Driver, J., & Dolan, R.J. (2005). Morphing marilyn into maggie dissociates physical and identity face representations in the brain. *Nature Neuroscience*, *8*, 107–113.
- Spiridon, M., & Kanwisher, N. (2002). How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron*, *35*, 1157–1165.
- Tong, F., Nakayama, K., Vaughan, J.T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, *21*, 753–759.
- Tsao, D.Y., Freiwald, W.A., Tootell, R.B., & Livingstone, M.S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, *311*, 670–674.
- Vuilleumier, P., Henson, R.N., Driver, J., & Dolan, R.J. (2002). Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*, *5*, 491–499.
- Williams, M.A., Dang, S., & Kanwisher, N.G. (2007). Only some spatial patterns of fMRI response are read out in task performance. *Nature Neuroscience*, *10*, 685–686.