

Cognition and Behavior

Attention Effects on Neural Population Representations for Shape and Location Are Stronger in the Ventral than Dorsal Stream

Anne B. Sereno,^{1,2} and Sidney R. Lehky³DOI:<http://dx.doi.org/10.1523/ENEURO.0371-17.2018>

¹Department of Psychological Sciences, Purdue University, West Lafayette, IN 47907, ²Weldon School of Biomedical Engineering, Purdue University, West Lafayette, IN 47907, and ³Computational Neurobiology Laboratory, The Salk Institute, La Jolla, CA 92037

Abstract

We examined how attention causes neural population representations of shape and location to change in ventral stream (AIT) and dorsal stream (LIP). Monkeys performed two identical delayed-match-to-sample (DMTS) tasks, attending either to shape or location. In AIT, shapes were more discriminable when directing attention to shape rather than location, measured by an increase in mean distance between population response vectors. In LIP, attending to location rather than shape did not increase the discriminability of different stimulus locations. Even when factoring out the change in mean vector response distance, multidimensional scaling (MDS) still showed a significant task difference in AIT, but not LIP, indicating that beyond increasing discriminability, attention also causes a nonlinear warping of representation space in AIT. Despite single-cell attentional modulations in both areas, our data show that attentional modulations of population representations are weaker in LIP, likely due to a need to maintain veridical representations for visuomotor control.

Key words: macaque monkey; object attention; population coding; spatial attention; vision

Significance Statement

Disruption of attention is a fundamental characteristic of many human disorders. Although much work has focused on how attention to shape and space modulates responses at a single neuron level, that has provided little insight into how these modulations translate into effects at the population level. Similar to prior findings, we show attentional modulations at the single-cell level in both ventral and dorsal cortical areas. However, we also show that attention alters population representations of both shape and space more strongly in the ventral stream (AIT) than in the dorsal stream (LIP). This study is a first demonstration of how attention modulates the organization of high-dimensional neurophysiological representations, findings that help link cellular attentional modulations more tightly with behavior.

Introduction

A disruption of attention is considered a key characteristic in a range of human neurologic, psychiatric, and developmental diseases such as Parkinson's disease,

schizophrenia, and autism. Attention has a variety of effects on neural activity. First, there are competitive interactions that occur when directing attention to one of multiple stimuli presented simultaneously within a single

Received October 29, 2017; accepted April 19, 2018; First published April 30, 2018.

The authors declare no competing financial interests.

Author contributions: A.B.S. designed research; A.B.S. performed research; A.B.S. and S.R.L. contributed unpublished reagents/analytic tools; A.B.S. and S.R.L. analyzed data; A.B.S. and S.R.L. wrote the paper.

This work was supported in part by Purdue University funds.

Acknowledgment: We thank Margaret E. Sereno for helpful comments on this manuscript.

Correspondence should be addressed to Anne B. Sereno, Department of Psychological Sciences, Purdue University, 703 Third Street, West Lafayette, IN 47907, E-mail: asereno@purdue.edu.

receptive field (Moran and Desimone, 1985; Motter, 1993; Luck et al., 1997; Reynolds et al., 1999; Reddy et al., 2009; Lee and Maunsell, 2010; Zhang et al., 2011). Second, attention modulates the gain of neural responses to a single stimulus within a receptive field (Colby et al., 1996; McAdams and Maunsell, 1999, 2000; Treue and Martínez-Trujillo, 1999; Reynolds et al., 2000; Martínez-Trujillo and Treue, 2002; Williford and Maunsell, 2006). Third, attention changes the statistics of stimulus responses. These include a decrease in the Fano factor (Mitchell et al., 2007; Cohen and Maunsell, 2009), a decrease in noise correlation between neurons (Cohen and Maunsell, 2009, 2011; Mitchell et al., 2009), and an increase in response correlations (e.g., γ -synchronization; Fries et al., 2001, 2008; Bichot et al., 2005; Lakatos et al., 2008; Wrobel, 2014).

In all three cases, attention increases the signal-to-noise ratio, or in other words, increases stimulus salience. Competitive effects reduce perturbations or “noise” caused by irrelevant stimuli. Increases in gain or synchrony increase the numerator of the signal-to-noise ratio, while noise attenuation decreases the denominator. The idea that attention increases stimulus salience dominates current thinking.

However, there is also a fourth class of attentional effects, one that does not involve modulating the signal-to-noise ratio, which has received less consideration. That is a shift in tuning curves caused by attention. Attention can shift receptive fields (spatial tuning curves) toward the attended location (Connor et al., 1996, 1997; Tolias et al., 2001; Womelsdorf et al., 2006, 2008; Anton-Erxleben et al., 2009; Kay et al., 2015). Attention can also cause shifts in feature tuning curves toward the attended feature characteristic (David et al., 2008; Ipata et al., 2009, 2012; Ibo and Freedman, 2014, 2016). These features include color, motion direction, shape, and orientation. Similarly, attention can also cause a shift in tuning for object categories (Çukur et al., 2013). Tuning curve shifts caused by attention change the neural representation of a stimulus, rather than increase the salience for a particular representation.

This study focuses on how attention changes stimulus representations rather than how attention increases stimulus salience. We look at these attentional effects at the population level, examining how attention to stimulus shape or attention to stimulus location across populations of cells affects the organization of multidimensional representation spaces defined by this population activity. In particular we examine (1) how attention to shape or location alters the representational discriminability or distinctiveness of different shapes or different locations (e.g., does attention magnify the distances between shapes in the neural representation space) and (2) whether attention to shape or location warps the organization of the repre-

sentation space in a way that goes beyond simple magnification. Furthermore, we make a comparison of these attentional effects in a dorsal stream structure (LIP) and a ventral stream structure (AIT) under identical conditions. Such findings would provide the first insight into how single-cell attentional modulations combine to alter the population representation of shape and space and provide a first comparison across the two cortical streams of attentional changes in the representation of shape and location at the level of neural populations.

Each stimulus is represented by a point in the neural representation space. The distance between different stimuli within the representation space is an indication of how similar or how distinct the stimuli are. The goal of this study is to examine how attention modifies the relative positions of stimuli in neural representation spaces. Our analysis of the data centers on multidimensional scaling (MDS), which can be considered a form of intrinsic population decoding (Lehky et al., 2013). MDS has previously been used to visualize shape representations in neural representation space for both the dorsal (Murata et al., 2000; Romero et al., 2013) and ventral visual streams (Rolls and Tovee, 1995; Op de Beeck et al., 2001; Vogels et al., 2001; Eifuku et al., 2004; Kiani et al., 2007). In our own work we have used MDS to construct neural representation spaces for stimulus shape (Lehky and Sereno, 2007), stimulus retinal position (Sereno and Lehky, 2011a), eye position (gaze angle; Sereno et al., 2014), as well as modeling of stimulus retinal position (Lehky and Sereno, 2011) and eye position (Lehky et al., 2016). Here, we extend this work to examine how attention modulates representation spaces for stimulus shape and retinotopic location. This work also extends a previous analysis of single-cell effects of attention (Sereno and Amador, 2006).

To do this, we recorded while the monkey performed two attention tasks. In one, attention was directed to stimulus shape and in the other to stimulus location. Stimuli were physically identical in both cases. Passive fixation conditions provided baseline responses to the same stimuli for comparison with the attentional tasks.

Materials and Methods

Animals, animal care, and surgical procedures

Two male macaque monkeys were trained to perform the behavioral tasks described below. A scleral search eye coil, implanted before training, monitored eye position. We recorded from both LIP and AIT of each monkey. The chambers for LIP were implanted first and centered 10–12 mm lateral and 3–5 mm posterior, and the chambers for AIT were implanted after recording from LIP and were centered 18 mm anterior and 18–21 mm lateral. Details of the surgical procedures have been described earlier (Sereno and Amador, 2006; Lehky and Sereno, 2007). All experimental protocols were approved by the Baylor College of Medicine, Rutgers University, and University of Texas Animal Welfare Committees and complied with the National Institutes of Health Guidelines.

DOI:<http://dx.doi.org/10.1523/ENEURO.0371-17.2018>

Copyright © 2018 Sereno and Lehky

This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International license](#), which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

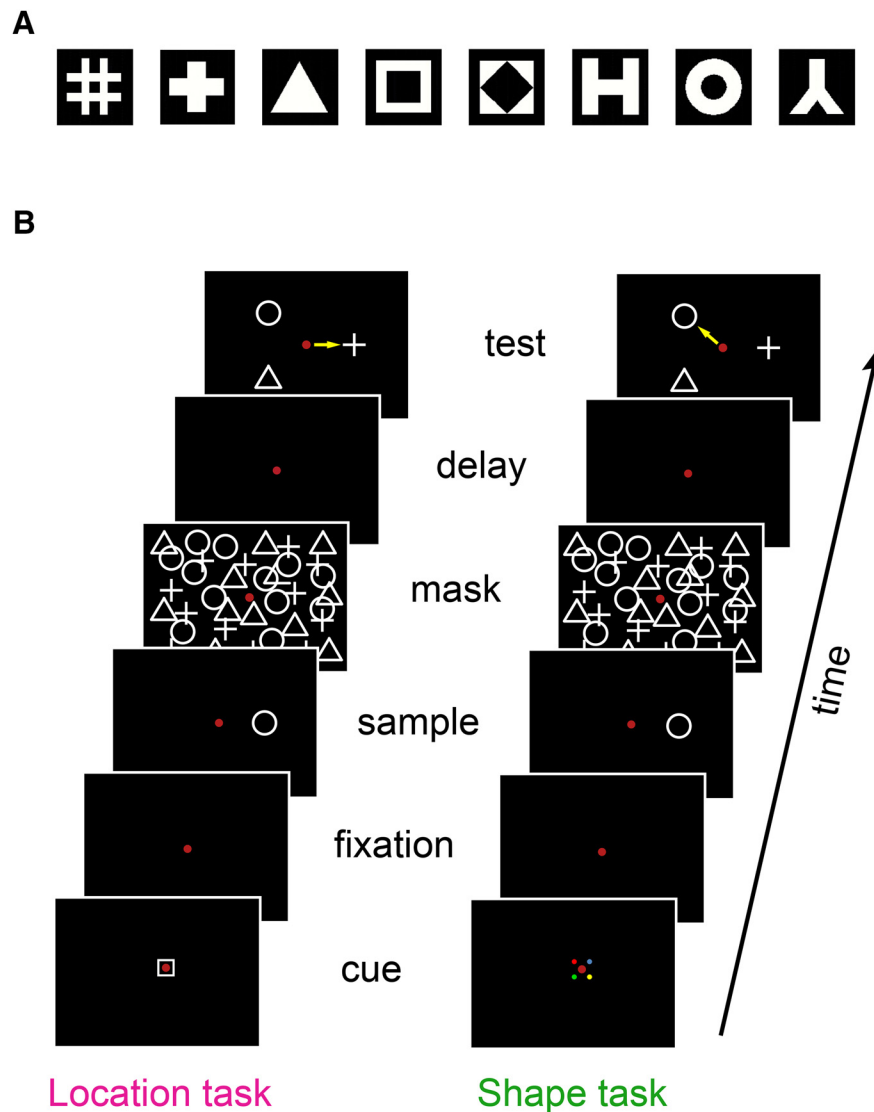


Figure 1. Stimuli and task. **A**, Eight stimulus shapes used in the task. **B**, Task sequence. A cue at start of trial indicated to the monkey which of two attention tasks to perform, either a shape attention task or a location attention task. Following the cue, a sample stimulus was shown, which consisted of a single shape presented at a particular location. Following a delay period, the test stimulus was shown, which consisted of three shapes presented at three locations. For the shape attention task, the monkey made a saccade to the test shape matching the sample shape, ignoring location. For the location attention task, the monkey made a saccade to the location matching the location of the sample stimulus, ignoring shape.

Stimuli

Stimuli for all tasks were chosen from among eight possible shapes (Fig. 1A; Sereno and Maunsell, 1998). Each shape was a simple 2D geometric form, black-and-white in color and presented on the screen against a black background. All shapes had an equal number of white pixels and fit within a square region of the same size.

Behavioral tasks

We recorded neural responses while the monkeys performed four different behavioral tasks. These included two delayed-match-to-sample (DMTS) attention tasks, namely a shape attention DMTS task and a location attention DMTS task. A comparison of neural responses under these two tasks at a population level is the focus of this report. The other two tasks were a passive shape task

(all eight shapes presented at a single location) and a passive location task (a single shape presented at eight locations), both used to characterize general response properties of each neuron without active attentional modulations.

Attention tasks

The two DMTS attention tasks have been described in detail in Sereno and Amador (2006) and are summarized in Figure 1B.

At the start of each trial the monkey was centrally fixated. A small cue surrounding the fixation point then informed the monkey which attention task to perform. Following the cue a sample stimulus was presented. The sample stimulus consisted of a single shape (selected from eight possible shapes) presented at one of eight

possible eccentric locations arranged in a circle around the fixation spot. For each cell the stimulus eccentricity was always constant. For different cells across the recorded population eccentricity varied (AIT mean eccentricity 4.0° , range 10.0 – 2.0° ; LIP mean eccentricity 11.2° , range 17.7 – 7.4°). Following a delay period (variable from 600 to 2700 ms), an array of test stimuli was presented. The test stimuli consisted of three shapes (again selected from the eight possible shapes) presented at three locations around the fixation spot, and at the same eccentricity as the sample.

For the shape attention task, the monkey made a saccade to the test shape matching the sample shape, ignoring location. For the location attention task, the monkey made a saccade to the location matching the location of the sample stimulus, ignoring shape. We were interested in examining attentional effects during the sample presentation period (shifted by visual latency).

The stimuli were identical for the two attention tasks. The cue at the start of the trial indicated to the monkey whether to direct attention to stimulus shape or stimulus location to successfully perform the task for a juice reward. Stimulus size ranged from 0.65° to 2.00° , increasing with eccentricity to maintain discriminability.

Catch trials typically comprised 20% of trials. Catch trials for the shape attention task did not include the sample shape in the test array, and those for the location attention task did not include the sample location. During catch trials, the monkey was rewarded for maintaining fixation on the central spot until it disappeared after an additional 1.3–1.9 s.

Passive shape task

Sereno and Amador (2006) have previously described this task. For this task, the stimulus shape was placed at the receptive field location producing the most effective response during the attention tasks. In each trial, the stimulus shape was randomly selected from among the same set of eight shapes used in the attention tasks. The selected shape was then presented four times in rapid succession (250–300 ms each time) before extinguishing the central fixation spot. The monkeys were rewarded for maintaining fixation on the central spot until it disappeared. Thus, over the course of multiple trials, the response of the neuron to all eight shapes was determined at the most effective location.

Passive location task

Sereno and Amador (2006) have also previously described this task. Each trial began with the presentation of a fixation point at the center of the visual display. Then the most effective stimulus shape for the neuron (out of the eight possibilities) was presented at one of eight peripheral locations arranged in a ring around the fixation spot, at the same eccentricities used in the attention tasks. The animal was required to make a saccade to the stimulus to obtain a juice reward. We were interested in the period after the stimulus appeared on the screen and before the monkey made a saccade toward it, when the stimulus was at one of eight possible retinal positions. Thus, over the course of multiple trials the response of the neuron at

eight locations was determined for the most effective shape.

Recording procedures

Stimuli were displayed on a 20-inch, 75-Hz CRT monitor with a resolution of 1152×864 pixels, placed 65 cm in front of the animal. The monitor subtended a visual angle of 27° height and 36° width. Beyond the monitor was a featureless 45×60 -cm black screen ($40^\circ \times 54^\circ$ visual angle), which supported an electromagnetically shielded window in its center through which the animal viewed the monitor. The monkeys viewed the stimuli binocularly. Experiments were conducted in a darkened room. Action potentials were recorded extracellularly while animals performed the behavioral tasks, using either tungsten (1–2 M Ω) or Pt/Ir single microelectrodes (1–2 M Ω , MicroProbes for Life Science). Electrode signals were amplified, filtered, and transformed into pulses by a window discriminator. Spike times were recorded with a resolution of 1 ms.

Histologic reconstruction in one monkey showed that the units recorded in posterior parietal cortex lay within area LIP in the lateral bank of the intraparietal sulcus. We did not distinguish between LIPd and LIPv, although cells were generally closer to the fundus than the lip of the sulcus. Units recorded in AIT were in the lower bank of the superior temporal sulcus (STS) and convexity of the middle temporal gyrus. A few perirhinal cells were included at our most anterior recording positions. Further details are given in Sereno and Maunsell (1998); Lehy and Sereno (2007)

Monkeys performed the two attention tasks while we slowly advanced the recording electrode with a hydraulic micromanipulator in search of neurons. We recorded from any neuron that was well isolated and appeared stable. After isolating a neuron, we presented a stimulus at various locations to determine the receptive field location that produced the most effective response. Typically, we tested eight locations equally spaced in a circle around the fixation spot. Then we selected three of those locations for use in the attention tasks. The three possible stimulus locations were selected so that at least one fell within the receptive field at the location that produced the most effective response at a given eccentricity. The three chosen locations were typically equally spaced at 120° intervals. We then presented all eight possible shape stimuli at the preferred location and selected three of them for the attention tasks. For most units, we chose the two shapes that produced the strongest response and the one that produced the weakest response. For most units, we then collected data with the two attention task trials intermixed. Shape and location attention task trials were intermixed either by presenting them in nine-trial blocks or randomly interleaved on a trial-by-trial basis. However, for a few units, data were collected sequentially for the two attention tasks. As long as the isolation could be maintained, we recorded with a goal of obtaining data from 216 correct matching trials [12 trials for each of nine sample conditions (three shapes by three locations) for each of two matching tasks], requiring ~25–30 min to

collect. For each cell included in data analysis we recorded a minimum of five repetitions (median 12) of each combination of sample location, sample shape, and task.

Following completion of the two attention tasks we also recorded responses during the passive shape and passive location tasks, provided the cell remained well isolated. In the passive shape task, there was a median of four repeat trials for each of the eight shapes, with four repeats of the same shape within each trial. Trial response was the average of the four repeats. For the passive location task, there was a median of 12 repeat trials of each of the eight locations.

Data analysis

We examined attentional effects during the sample period of the DMTS tasks. As DMTS is a working memory task, our data analyses thus focused on the stimulus encoding phase of working memory, and not the retention and retrieval phases of the task. As the sample stimulus appeared by itself on an otherwise black screen (aside from the fixation spot), it was a singleton stimulus.

Data analysis focused on using MDS to compare population responses during different tasks. However, using MDS mathematically requires that every neuron in the recorded population receives an identical stimulus set, and that was not the case for the two attention tasks. Although for the two tasks both the stimulus shapes and stimulus locations were identical for each particular cell, stimulus shapes and locations varied from cell to cell. During the attention tasks, each neuron was tested with a random subset of three stimulus shapes from the eight shapes used in the passive shape task due to time constraints. Similarly, during the attention tasks each neuron was tested at a different set of only three locations roughly 120° apart along in circle centered on the fixation point, with those locations adjusted to reflect the position of that neuron's receptive field. Completing both attention tasks with all combinations of shapes in eight locations would have required stable recording for over three hours for each cell.

The three locations in the attentional tasks approximately corresponded to three of the eight locations in the passive tasks, although they were not exactly the same. We used interpolation to estimate responses at exactly matching locations, as described below. The response interpolation was done independently for each neuron.

We calculated the interpolation as follows. In the passive task, there were eight stimulus locations at constant eccentricity arranged along a circle around fixation. In the active task, there were three stimulus locations arranged along a circle at the same eccentricity. The polar angle positions of the three active locations did not exactly match any of the polar angle positions of the eight passive locations, although they matched in eccentricity. A 1D spline interpolation was calculated along the circle of constant eccentricity using responses from the eight passive-task points to estimate passive responses at polar angles corresponding to the three points in the active task. For AIT the mean shift in polar angle during interpolation was 4.9°. In other words, on average, we estimated

the neural response of a cell at a location shifted 4.9° in polar angle and 0° in eccentricity from the location of a data point. For LIP the median shift in polar angle during interpolation was 5.2°.

Attentional gain factors

We wanted to compare responses under the attentional tasks and the passive tasks using MDS. We therefore had to estimate attentional responses for the eight shapes and eight locations used in the passive tasks, based on the three shapes and three locations actually used in the attentional tasks, all collected from the same cell. To do that, for each cell we calculated an attentional gain factor based on responses to the three stimuli (shapes or locations) that were used in both the attentional and passive tasks. Then that gain factor was used to estimate attentional responses to the remaining stimuli that were used in the passive task but not used in the attentional task. For each cell, we determined four gain factors: (1) shape attentional gain factor when decoding population data for stimulus shape; (2) location attentional gain factor when decoding population data for stimulus shape; (3) shape attentional gain factor when decoding population data for stimulus location; and (4) location attentional gain factor when decoding population data for stimulus location.

The shape attention task required endogenous attention directed to stimulus shape, while the location attention task required endogenous attention directed to stimulus location. In contrast, the passive shape and passive location tasks were passive in the sense that they did not require deployment of endogenous object or location selective attention. Thus, taking the ratio of responses to the same stimuli in the same neuron during an attention task and a passive task allows one to define an attentional gain factor for the attention task. As all stimuli had fixed contrast at maximum value (contrast = 1.0), we treated the modulation of neural activities by attention as response gain (McAdams and Maunsell, 1999; Treue and Martínez-Trujillo, 1999) rather than contrast gain (Reynolds et al., 2000; Martínez-Trujillo and Treue, 2002).

After determining the attentional gain factors for a neuron based on the three overlapping stimuli between the attention and passive tasks, all eight stimuli (shape or location) in the passive task were then multiplied by that gain factor to estimate their responses during the attention task.

For the calculation of the attentional gain factors, a linear regression was calculated between responses to the three stimuli in the attention task and the three corresponding stimuli (out of eight stimuli) in the passive task. The regression was subject to the constraint that it passed through the origin. The slope of this regression line defined the attentional gain factor. The assumption underlying this procedure was that the effect of attention was to linearly scale the passive responses of neurons by some multiplicative factor. This assumption is supported by extensive experimental evidence for both spatial and feature attention (McAdams and Maunsell, 1999; Maunsell and Treue, 2006; Maunsell, 2015).

Neural representation of shape space We derived two attentional gain factors for a given cell relative to the

passive shape task. One was based on responses during the shape attention task relative to the passive shape task and the other based on responses during the location attention task relative to the passive shape task. Multiplying the passive shape data by these two attentional gain factors gave two estimates of responses to the eight shapes, one when directing attention to stimulus shape and the other when directing attention to stimulus location. Following this procedure for every cell in the recorded population produced two MDS population estimates of the neural shape representation space under different attentional conditions (shape or location).

Neural representation of location space We also derived two attentional gain factors for that cell relative to the passive location task. One was based on responses during the shape attention task relative to the passive location task and the other based on responses during the location attention task relative to the passive location task. Multiplying the passive location data by these two attentional gain factors gave two estimates of responses to the eight locations, one when directing attention to stimulus shape and the other when directing attention to stimulus location. Following this procedure for every cell in the recorded population produced two MDS population estimates of the neural location representation space under the two attentional conditions.

Statistical power analysis We performed a power analysis on the difference between mean vector response distances under the shape attention task and the location attention task. The power analysis was done using the MATLAB (MathWorks Inc.) `sampsizepwr` command in the Statistics and Machine Learning Toolbox with `testtype` set to `z-test`. Sample size was 28, which is the number of paired comparisons $n(n-1)/2$ for $n = 8$ shapes or locations. The null hypothesis was that there was zero mean response difference between the two attentional conditions with a standard deviation as given by our data. For the alternative hypothesis we set the mean response distance to a value that produced a statistical power of 0.8 at $p = 0.05$.

On decoding neural population data for stimulus shape, in AIT an attentional effect of 11.2% produced a power of 0.8, while in LIP an attentional effect of 6.1% produced a power of 0.8. When decoding the data for stimulus location, an attentional effect of 5.2% in AIT and 3.6% in LIP produced a power of 0.8.

Published data shows attentional effects are strong in extrastriate areas, and it is expected in mid to late stages to have attentional modulations of 20–50% (McAdams and Maunsell, 1999; Martínez-Trujillo and Treue, 2002; Williford and Maunsell, 2006). Thus, this analysis indicates that our sample size was sufficient to reveal significant statistical effects in the areas from which we were recording.

Multidimensional scaling (MDS)

Using MDS, it was possible to decode the data to establish a shape representation space using responses from eight shapes or decode the data to establish a location representation space using responses from eight locations. These provided estimates of neural shape rep-

resentation space and neural location representation space, respectively.

The shape representation space (using MDS to decode for shape) was determined when directing attention either to stimulus shape or stimulus location. The location representation space (using MDS to decode for location) was also determined when directing attention to stimulus shape or stimulus location. Therefore, four MDS analyses under attentional conditions were performed.

Stimulus eccentricities had been adjusted for each recorded neuron to reflect the position of that neuron's receptive field. To make these data mathematically suitable for MDS analyses when decoding for stimulus location, we only included those cells whose stimuli fell in a narrow band of eccentricities (AIT: 2–6°; LIP: 7–12°). Then, for purposes of MDS, we mathematically treated the stimulus eccentricities for all cells as if they were identically located at the average eccentricity of the population. We have previously used this averaging method for MDS analysis in (Lehky and Sereno, 2007; Sereno and Lehky, 2011a; Sereno et al., 2014).

At this point, the MDS procedure was the same as what we have done in previous studies for both stimulus shape (Lehky and Sereno, 2007) and stimulus location (Lehky and Sereno, 2011; Sereno and Lehky, 2011a; Sereno et al., 2014; Lehky et al., 2016). The first step was to define a population response vector for each stimulus. The population response vector was an ordered list of the firing rates of all neurons in the recorded population to the same stimulus. The set of neural responses forming a single response vector pooled data from multiple recording sessions.

There was one population response vector for each stimulus. When decoding for shape we had eight population response vectors, one for each of the eight shapes in our stimulus set whose responses were measured at the most effective location. When decoding for location we also had eight population response vectors, one for each of the eight locations in our stimulus set whose responses were measured using the most effective shape.

Response distances To test how attention alters the response discriminability or distinctiveness of stimuli, we calculated the mean distance between population response vectors to all eight shapes or all eight locations. Each element in a response vector represented the mean response of a single neuron to the stimulus. A larger mean response distance between stimuli in the representation space indicated that the stimuli were more discriminable or distinctive. A larger mean would thus be an indication that attention caused a general expansion or magnification of the representation space.

To calculate mean response distance, we first measured the distances between all the response vectors. With our $n = 8$ stimuli (shape or location) this produced an 8×8 distance matrix containing $n(n-1)/2 = 28$ unique distances. We used a correlation distance metric, defined as $d = 1-r$ where r is the correlation coefficient between the elements of two response vectors. Thus, response vectors to different stimuli (e.g., different shapes or differ-

ent locations) that had less correlated responses across the population would have larger response distances within the representation space and be more discriminable or distinctive. Mean response distance was then determined by averaging all response distances between different stimuli. Note we measured the correlation between response vectors for the same neural population under two behavioral conditions, not correlations between individual neurons within that population.

MDS plots To test whether attention could alter the organization of the neural representational space (i.e., relative positions of individual stimuli), we compared results from MDS analyses under different attentional conditions. A question here was whether attention could cause a warping of the representation space, apart from any magnification effects on the space. The distance matrix, described above in *Response distances*, also serves as the input to the MDS algorithm proper. To do MDS we used the `cmdscale` command in the MATLAB Statistics and Machine Learning Toolbox. The MDS algorithm served to reduce the dimensionality of the space representing each stimulus shape from potentially up to n dimensions, where n is the neural population size, to a smaller number of dimensions that capture most of the variance in the data. Such a low-dimensional representation allows easier visualization of patterns within the data. In presenting the MDS results we plot two dimensions. The two dimensions capture 70–80% of the data variance for shape and over 85% of the variance for location. There are eight points plotted in the low-dimensional space resulting from the MDS procedure because there are eight stimuli (either shapes or locations).

Stress We used an error measure called stress to quantitatively compare MDS results from one condition with some other condition (either another MDS result or some aspect of the physical world). The configuration of points in low-dimensional space that results from MDS will have particular positions with respect to each other, but the configuration will be linearly transformed in an indeterminate way, meaning that the configuration as a whole can be linearly scaled, translated, rotated, or reflected. Because we wanted to compare the configuration of points produced by MDS with a set of corresponding points from some other condition, we removed this linear indeterminacy by applying the Procrustes transform (using the `procrustes` command in the MATLAB Statistics and Machine Learning Toolbox). The Procrustes transform takes two configurations of points and applies a linear transform that minimizes the difference between them, as indicated by the stress error measure:

$$\text{Stress} = \sqrt{\frac{\sum_i \sum_j (d_{ij}^1 - d_{ij}^2)^2}{\sum_i \sum_j (d_{ij}^1 - \langle d_{ij}^1 \rangle)^2}} \quad (1)$$

In the equation, d_{ij}^1 is Euclidean distance between points i and j in the first configuration and d_{ij}^2 is the Euclidean distance between the corresponding points in the second configuration. The residual stress that remains after the

best-fit linear transform is a global measure of how closely the positions of two sets of points match each other.

Stress is therefore a global measure of the difference between two neural response spaces. To measure the statistical significance of that difference, we used bootstrap resampling of the data. We started with a null hypothesis that there was no statistical difference between the data under two attentional conditions (i.e., data from the two conditions were two samples from an identical distribution). To test that hypothesis, we pooled firing rates from individual trials for the two attentional conditions into a single array, shuffled the array, and then split the shuffled array into two subarrays. We then separately did the MDS analysis on these two subarrays and calculated the stress between the MDS results for the two. The MDS analysis was repeated 10,000 times, so we had 10,000 stress values. The p value for the hypothesis that the two attentional conditions were significantly different was given by the fraction of resampled stress values greater than the stress value observed in the original data.

Results

Behavioral response times and error rates for the different tasks have previously been presented in [Sereno and Amador \(2006\)](#); their Fig. 16).

Single-cell analyses

To allow direct comparisons of our data to much previous work on the neural basis of attention, we first performed single-cell analyses, separately examining responses of each cell to all stimulus shapes (at the most effective location) and to all stimulus locations (using the most effective shape).

Attentional gain factors for stimulus shape

For each cell we calculated an attentional gain factor for the two attentional tasks relative to the passive shape task. In AIT, there were 85 cells held long enough to obtain sufficient data for all three tasks, while in LIP there were 53 cells. Histograms of these gain factors for AIT are shown in [Figure 2A](#), and for LIP in [Figure 2B](#). On average, at the single-cell level, attentional demands (both attention to shape or attention to location) increased responsiveness to the different shape stimuli in both AIT and LIP compared to the passive shape task. At the single-cell level, LIP showed greater attentional modulations than AIT (difference between LIP and AIT significant at $p = 0.010$ for the shape attention task and $p = 0.0047$ for the location attention task under the Wilcoxon rank sum test). Noteworthy is that in both brain areas there was a broad range of attentional gain factors for stimulus shape across the population.

Attentional gain factors for stimulus location

Analogous to shape, for each cell we calculated an attentional gain factor for the two attentional tasks relative to the passive location task. In AIT, there were 76 cells held long enough to complete the two attention tasks and passive location task, while in LIP, there were 47 cells. Histograms of these gain factors for AIT are shown in [Figure 2C](#), and for LIP in [Figure 2D](#). In LIP, both attention to shape and to location actually decreased mean respon-

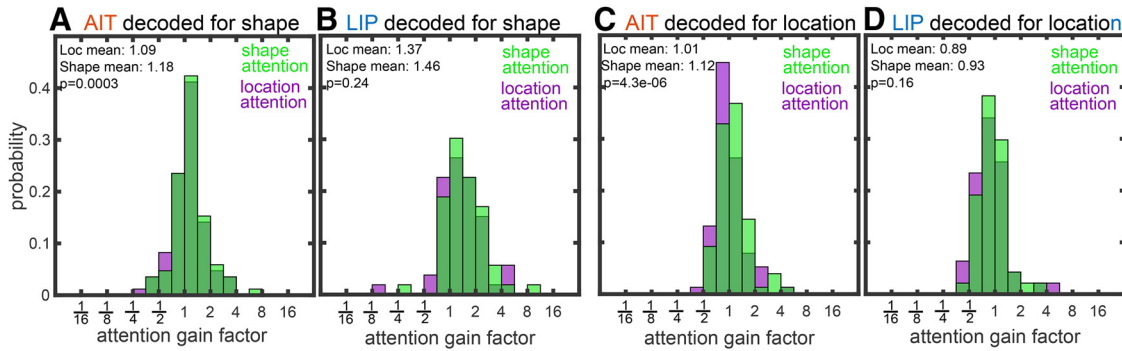


Figure 2. Attentional gain factors associated with the two attention tasks (shape attention task and location attention task). **A, B,** Attentional gain factors of AIT and LIP neurons when neural responses were decoded for stimulus shape. Gains for the two attention tasks (attention to shape in green and location in purple) were measured relative to neural activities during the passive shape task. **C, D,** Attentional gain factors of AIT and LIP neurons when neural responses were decoded for stimulus location. Gains for the two attention tasks were again measured relative to neural activities during the passive location task. Significance of the difference between attention tasks was calculated using the Wilcoxon sign rank test.

siveness to different locations compared to the passive task (mean gain factor < 1). In AIT, there was little change in mean responses during attention to location and a slight enhancement for attention to shape compared to the passive task (mean gain > 1). The *p* values for the difference between AIT and LIP was *p* = 0.017 for the shape attention task and *p* = 0.21 for the location attention task under the Wilcoxon rank sum test. In both brain areas, stimulus location shows a broad range of attentional gain factors across the population, as was the case for stimulus shape.

Population analyses

We refer to neural population representations for a set of stimulus shapes as collectively forming a multidimensional “shape space.” Each shape will have a position in the shape space defined by the population response to that shape. The representational distance from other shapes is defined by the difference in population responses as specified by some metric (in our case, the correlation distance, 1-*r*). Analogously, population representations for a set of stimulus locations collectively form a “location space.” Below we examine attentional effects on the shape space and location space.

Response distances of shapes and locations

For each cortical area, Table 1, columns AIT and LIP under shape space, shows the mean response distances between different stimulus shapes for each of the three tasks (shape attention, location attention, and passive), as well as the average distance collapsed across the three tasks. For each cortical area, Table 1, columns AIT and LIP under location space, shows the mean response distances between different stimulus locations for each of the three tasks, as well as the average distance collapsed across the three tasks. Figure 3A (plotting the bottom row of Table 1) illustrates the average distance of the three tasks in area AIT and LIP for different shapes and different locations. These data demonstrate that response distances between shapes were greater (greater representational discriminability between shapes) in AIT than LIP, and furthermore, distances between locations were

greater (greater representational discriminability between locations) in LIP than AIT.

Figure 3B (based on Table 2) breaks down the averages shown in Figure 3A according to task, plotting the effect of attention to shape versus attention to location on mean distances for each cortical area. These attentional effects are plotted separately for shape space (Fig. 3Bi,ii) and location space (Fig. 3Biii,iv). Interestingly, the attentional condition (attention to shape vs location) did not significantly alter mean distances in LIP in either shape space (Fig. 3Bii; Table 2; -4.5%, *p* = 0.08) or location space (Fig. 3Biv; Table 2; -0.47%, *p* = 0.31). These data indicate that at the population level LIP responses were relatively constant under changing voluntary attentional demands. In contrast, attentional condition significantly affected mean distances in AIT for both shape space (Fig. 3Bi) and location space (Fig. 3Biii).

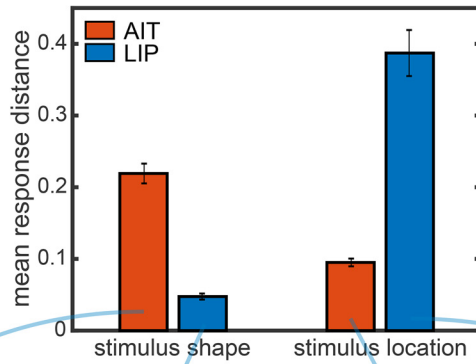
We next compared the mean response distance for each attentional task (shape or location) with the response distance for its corresponding passive condition (dashed horizontal line in each plot in Fig. 3B, with significance indicated immediately above the histogram bar; see also Table 2).

Table 1. Population response vector distances

Task	Shape Space		Location Space	
	i. AIT	ii. LIP	iii. AIT	iv. LIP
a. Passive (shape/location)	0.2303	0.0481	0.0858	0.3800
b. Shape Attention	0.2329	0.0461	0.1080	0.3899
c. Location Attention	0.1943	0.0483	0.0914	0.3917
d. 3 Task Average	0.2192	0.0475	0.0951	0.3872

i, ii, Mean distances between population response vectors for different stimulus shapes under three tasks, the shape attention task, the location attention task, and the passive shape task. Results for AIT are based on the distance matrices underlying the MDS results in Figure 4. Results for LIP are based on the distance matrices underlying the MDS results in Figure 5. iii, iv, Mean distances between population response vectors for different stimulus locations under three tasks, the shape attention task, the location attention task, and the passive location task. Results for AIT are based on the distance matrices underlying the MDS results in Figure 6. Results for LIP are based on the distance matrices underlying the MDS results in Figure 7.

A



B

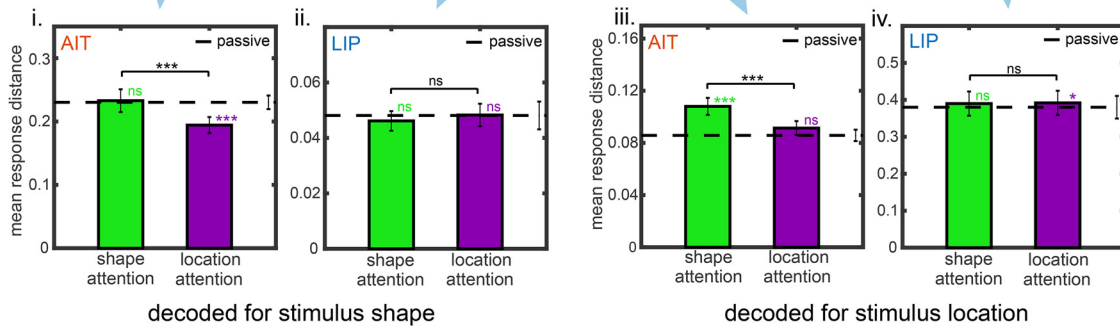


Figure 3. Mean distance between population response vectors. **A**, Mean distances between population responses pooled over all stimulus shapes or all stimulus locations in areas AIT (red bars) and LIP (blue bars), and also pooled across all attentional conditions. This plots the “3 task average” distances in Table 1, bottom row. **Bi, Bii**, Attentional effects on mean response distance pooled over all stimulus shapes. The dashed lines indicate the response distances under the passive shape task. **Biii, Biv**, Attentional effects on mean response distance pooled over all stimulus locations. The dashed lines indicate the response distances under the passive location task. Asterisks indicate statistically significant differences under the Wilcoxon rank sum test. Error bars indicate SE.

Table 2. Comparison of population response vector distances under different attentional conditions

Task Comparison	Shape Space		Location Space	
	i. AIT	ii. LIP	iii. AIT	iv. LIP
a. Shape vs Location Attention				
% change	18.0%	-4.5%	16.7%	-0.47%
p-value	p=1.7x10 ⁻⁵	p=0.08	p=3.8x10 ⁻⁶	p=0.31
b. Shape Attention vs Passive				
% change	1.11%	-4.2%	22.9%	2.56%
p-value	p=0.48	p=0.96	p=4.2x10 ⁻⁵	p=0.06
c. Location Attention vs Passive				
% change	-16.9%	0.34%	6.31%	3.04%
p-value	p=2.1x10 ⁻⁴	p=0.51	p=0.12	p=0.05

a, Shape attention versus location attention. Positive % change values indicate shape attention larger. b, Shape attention versus passive. Positive % change values indicate shape attention larger. c, Location attention versus passive. Positive % change values indicate location attention larger. For each comparison, i, ii shows population response distances between different shapes, and iii, iv shows population response distances between different locations. Red shading shows statistically significant attention effects at the $p = 0.05$ level. In AIT (but not LIP), attention to shape significantly increased response distances to both shape and location, making the shapes and the locations of the shapes more discriminable or distinctive in their neural representations. The comparison between the two attention tasks is highlighted in gray as being of primary interest here.

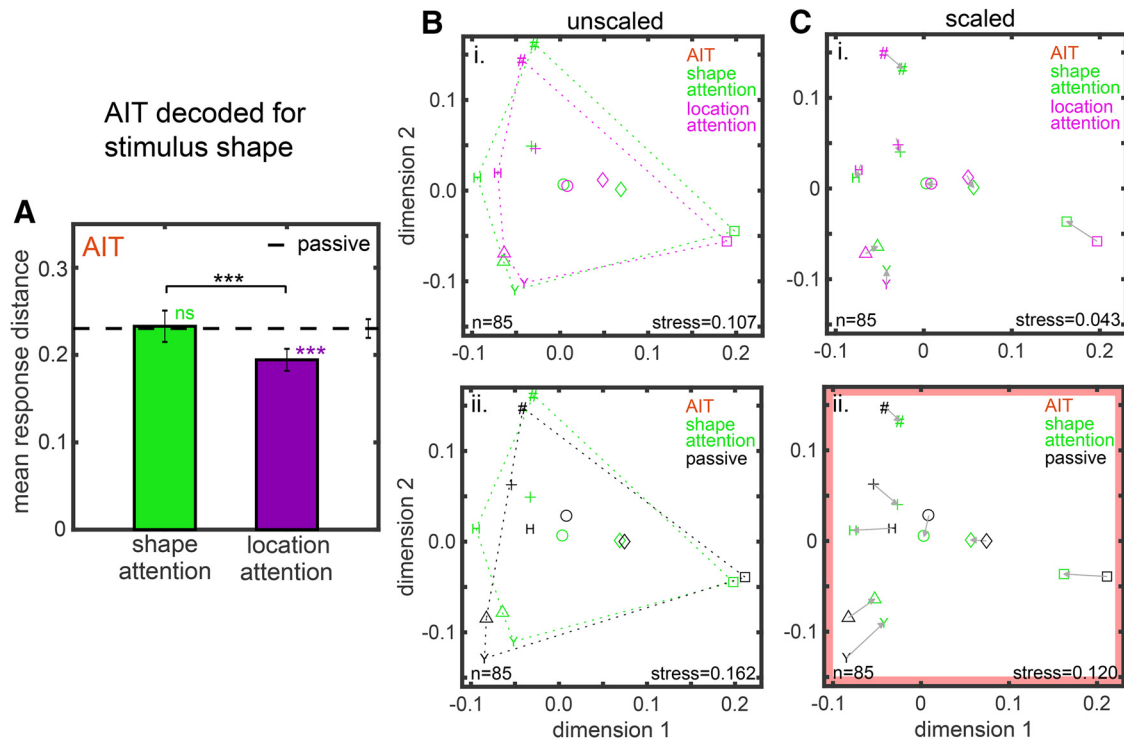


Figure 4. Attentional effects on shape space in AIT based on MDS analysis of neural population responses. **A**, Mean response distance pooled over all stimulus shapes under the shape attention and location attention tasks (from Fig. 3*B*). **B**, MDS shape space compared under two attentional conditions. **Bi**, Shape attention task versus location attention task. **Bii**, Shape attention task versus passive shape task. In both cases, the Procrustes transform was applied to align results from the two attentional conditions. Here, the scaling aspect of the Procrustes transform was turned off to retain differences in mean distance between attentional conditions as seen in **A**. Dotted lines indicate the convex hull around the eight points in the shape space for each attentional condition. The areas enclosed by the convex hulls gives a rough measure of the relative mean response distance for the two attentional conditions, in accord with **A**. **Ci**, **Cii**, Same as **B**, except that the full Procrustes transform was performed here, including scaling, to minimize stress between shape spaces under the two attentional conditions. Arrows link corresponding points in the two shape spaces. Significances of attentional differences between scaled shape spaces were examined by bootstrap resampling, and those significant at $p = 0.05$ level have the panel outlined in red.

For shape space, when comparing the attentional tasks with the passive tasks, LIP showed little change in mean response distances between different stimulus shapes (no differences; Fig. 3*Bii*; Table 2; -4.2%, $p = 0.96$ for shape attention, and +0.34%, $p = 0.51$ for location attention). Similarly, for location space, LIP showed little change in mean response distances between different stimulus locations (marginal or small differences; Fig. 3*Biv*; Table 2; +2.56%, $p = 0.06$ for shape attention, and +3.04%, $p = 0.05$ for location attention).

AIT showed larger modulations than LIP in mean response distances between the attentional and passive tasks in both shape and location space. Specifically, for shape space (Fig. 3*Bi*; Table 2), attention to location reduced shape response distance or shape discriminability (-16.9%, $p = 0.0002$) but with little change for attention to shape (+1.11%, $p = 0.48$). For location space (Fig. 3*Bii*; Table 2), attention to shape increased location response distance or spatial discriminability (+22.9%, $p = 0.00004$) with little change for attention to space (+6.31%, $p = 0.12$).

MDS plots and changes in the organization of shape and location space

In addition to quantifying attentional effects on mean response distances pooled over all stimuli, as in Figure 3,

we examined how attention affects relative responses of individual stimuli by plotting results from MDS. Multidimensional representation spaces for stimulus shape are shown for AIT (Fig. 4) and LIP (Fig. 5). These show response distances between individual shapes and their relative positions in shape space. The multidimensional representation spaces for location are shown for AIT (Fig. 6) and LIP (Fig. 7), showing response distances between individual stimulus locations and their relative positions in location space.

Examining results for shape in AIT further (Fig. 4), Figure 4*A* shows the respective mean response distances, replotted from Figure 3. Figure 4*B* shows attentional effects on MDS shape spaces, comparing the shape attention task versus location attention task (Fig. 4*Bi*) and the shape attention task versus passive task (Fig. 4*Bii*). For each of the two comparisons, the shape spaces have been aligned using the Procrustes transform, but with the scaling aspect of the Procrustes transform turned off. (The Procrustes transform normally aligns two MDS maps using linear scaling, translation, rotation, and reflection to minimize stress between them and thus facilitate comparison of the configuration of points in the two maps.) Because scaling has been turned off, the MDS plots in

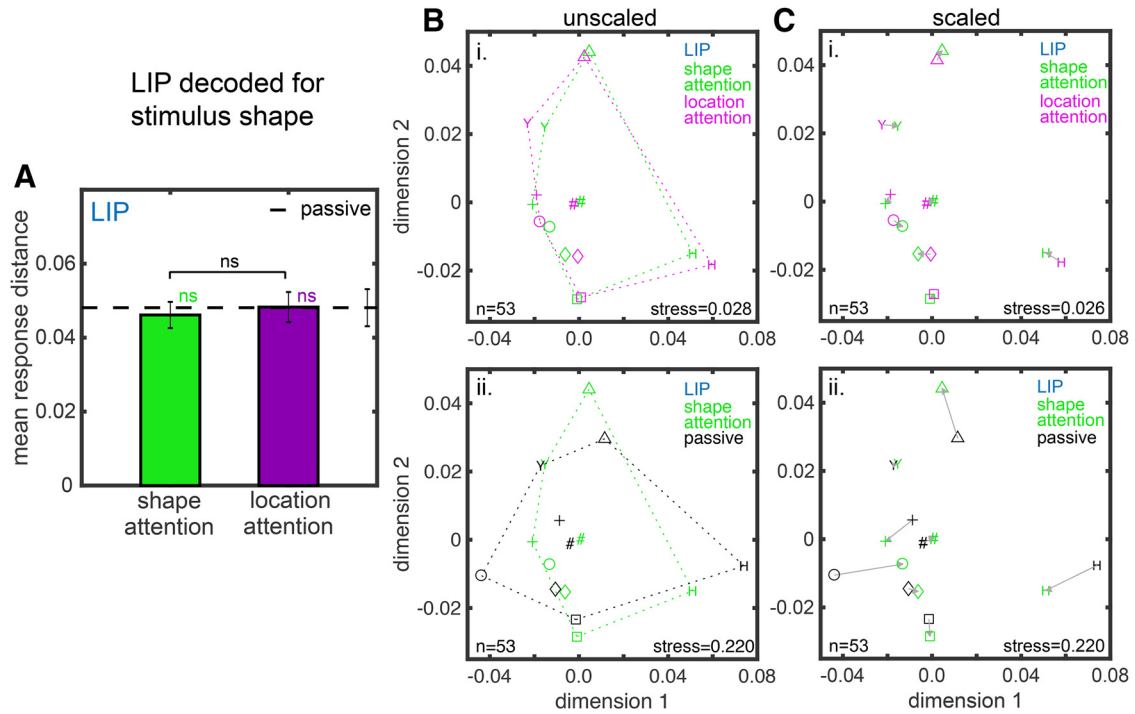


Figure 5. Attentional effects on shape space in LIP based on MDS analysis of neural population responses. Panels are analogous to Figure 4.

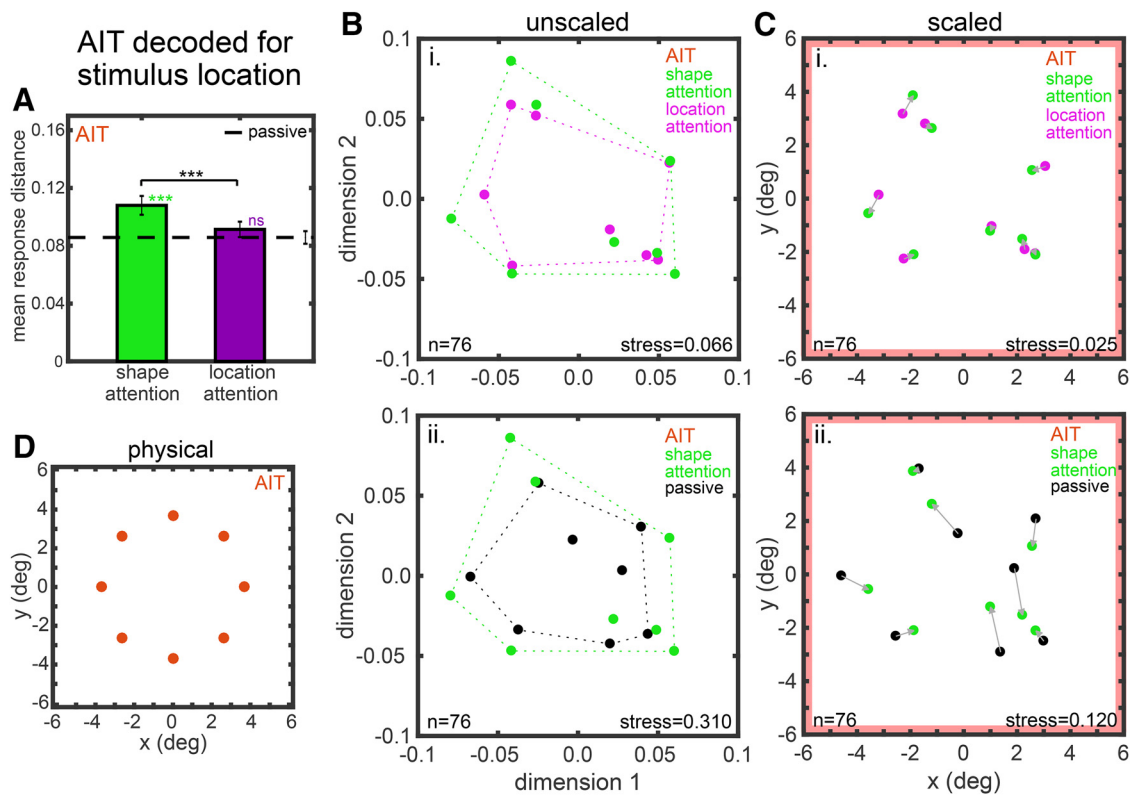


Figure 6. Attentional effects on location space in AIT based on MDS analysis of neural population responses. **A–C**, Analogous to Figure 4. **D**, Physical locations of stimuli.

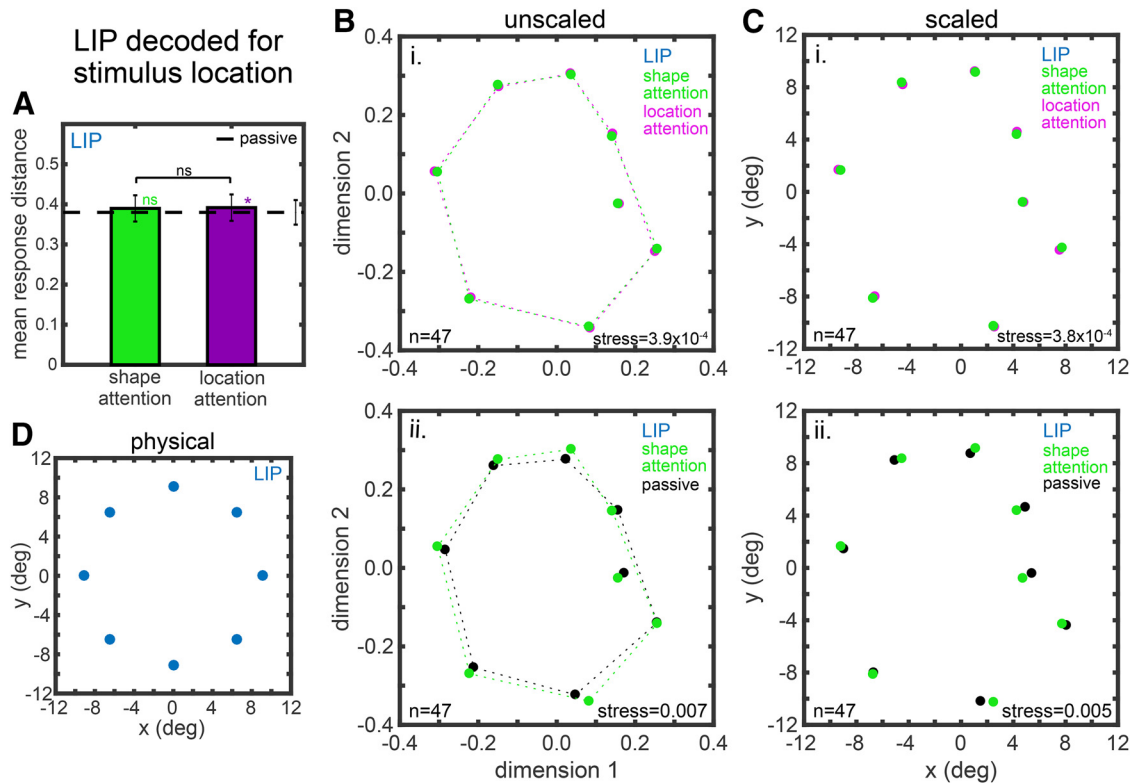


Figure 7. Attentional effects on location space in LIP based on MDS analysis of neural population responses. **A–C**, Analogous to Figure 4. **D**, Physical locations of stimuli.

Figure 4B retain information about differences in mean response distance as shown in Figure 4A. In each MDS plot in Figure 4B the convex hull (enclosing perimeter) around all eight points in the shape space has been drawn (dotted lines). Comparing the sizes (radii) of the convex hulls under two attentional conditions gives a rough indication of the change in mean response distance between the two conditions. In Figure 4Bi, the convex hull for the shape attention task is larger than that for the location attention task, consistent with the mean distance histogram in Figure 4A. In Figure 4Bii, the convex hull for the shape attention task is about the same as the convex hull for the passive task, again consistent with Figure 4A. Thus, we see consistency between the mean distance analysis in Figure 3 and the MDS analysis in Figure 4.

Figure 4C shows attentional effects on MDS shape representation maps in AIT, as in Figure 4B, but this time with the maps for each pair of attentional conditions aligned using the full Procrustes transform, including scaling. Using the full Procrustes transform minimizes the stress between maps. By using the Procrustes transform, we have removed linear effects of attention on the spatial maps, leaving nonlinear effects. By nonlinear attentional effects we mean that attention causes some shape response distances to increase and other shape response distances to decrease or causes response distances to increase or decrease non-uniformly for different shapes. In other words, in addition to linear scaling effects, attention has the possibility of nonlinearly distorting or warping the multidimensional shape representation space. To de-

termine whether the attentional warping of the shape spaces was significant, we performed bootstrap resampling of the data as described in the Methods section. We found the shape spaces for the shape attention task versus location attention task were not significantly different at the $p = 0.05$ level (Fig. 4Ci). However, there was a significant difference in the shape attention task versus passive comparison, indicated by a red border around Figure 4Cii.

How attention affects the configuration (relative positions) of responses for individual stimuli within the shape space for LIP is shown in Figure 5, with all the conventions as in Figure 4. For LIP there were no significant attentional effects in the MDS maps for either the shape attention task versus location attention task (Fig. 5Ci) or the shape attention versus passive task (Fig. 5Cii) comparisons (hence no red borders in the panels).

The representation space for stimulus location for AIT is shown in Figure 6 and for LIP in Figure 7. Figures 6A–C, 7A–C are again analogous to Figure 4. Panel d shows the physical locations of stimuli (this panel has no analog for shape in Figures 4 or 5, because there is no canonical or universally accepted physical shape space). The attentional effects on AIT location space are significant at the $p = 0.05$ level for both the shape attention versus location attention comparison (Fig. 6Ci) and the shape attention versus passive comparison (Fig. 6Cii), as indicated by the red borders of those panels. In contrast, attentional effects on LIP location space were not significant for either attentional comparison (Fig. 7Ci,ii).

Table 3. Stress between pairs of MDS stimulus representation spaces corresponding to different attentional tasks

Task Comparison	Shape Space		Location Space		
	i. AIT	ii. LIP	iii. AIT	iv. LIP	
a. Shape vs Location Attention	stress	0.0427	0.0256	0.0251	0.0004
	p-value	p=0.1570	p=0.7909	p=0.0019	p=0.8335
b. Shape Attention vs Passive	stress	0.1197	0.2200	0.1196	0.0047
	p-value	p=0.0352	p=0.0770	p<1x10 ⁻⁴	p=0.9323
c. Location Attention vs Passive	stress	0.0658	0.1696	0.1657	0.0052
	p-value	p=0.2647	p=0.1858	p<1x10 ⁻⁴	p=0.9171

Stress (Eq. 1) quantifies global differences between the shape spaces or location spaces under the different attentional conditions. Smaller stress indicates that representation spaces for two tasks are more similar; *p* values of the task differences are based on bootstrap resampling. Red shading shows statistically significant differences at the *p* = 0.05 level. For shape space, results for AIT correspond to MDS plots in Figure 4 and results for LIP correspond to MDS plots in Figure 5. For location space, results for AIT correspond to MDS plots in Figure 6 and results for LIP correspond to MDS plots in Figure 7. The comparison between the two attention tasks is highlighted in gray as being of primary interest here.

Table 3 summarizes how attention distorts MDS shape representation spaces (Figs. 4, 5) and location representation spaces (Figs. 6, 7) as measured by stress (Eq. 1). It shows stress values for different attentional task comparisons, as were presented in Figures 4–7C, plus the *p* values for the comparisons. The table also includes results for the location attention versus passive comparisons, whose MDS plots were not shown in Figures 4–7.

Significant attentional task effects at the *p* < 0.05 level are indicated by red shading in Table 3. We see significant task effects in both the AIT shape space and AIT location space. None of the attentional task comparisons were significant for either LIP shape space or, interestingly, LIP location space. These findings show that attentional effects can produce changes in the population encoding of individual stimuli such as to shift the relative positions of those stimuli non-uniformly within a multidimensional representation space. In other words, these data show attentional warping of representation spaces. Such warping was substantially larger in AIT than in LIP.

Finally, in Table 4, we see stress values when comparing stimulus locations recovered by MDS with physical stimulus locations. As we have reported previously (Serenó and Lehky, 2011b), stress values were much

smaller in LIP than AIT, indicating that stimulus location was represented more accurately by neural population representations in LIP. Further, attention directed to shape or location did not have a major effect on the accuracy with which LIP represented physical locations.

Discussion

Studies of attention generally focus on its ability to increase the salience of a target through a variety of means, including competitive effects as well as changes in gain, noise, correlation, and synchrony. All these mechanisms essentially increase the signal-to-noise ratio of the stimulus without indicating whether attention affects the representation or encoding of stimulus shape or location itself. While enhancing salience is certainly an important mechanism for attention, here we demonstrate how attention to stimulus shape or location differentially affects the organization of neural population representation spaces for shape and location in ventral and dorsal cortical brain regions. Such findings indicate that attention may alter population representations of stimuli, and not just enhance a fixed, pre-existing stimulus representation. Most importantly, despite similar attentional gain effects at the single-cell level across the two visual streams, we demonstrate striking differences in attentional effects at the level of population representation within different cortical areas.

Specifically, we report two changes in high-dimensional neural representation spaces due to attention. The first was a general magnification of the representation space, leading to increased discriminability between stimuli because their representations were further apart in the representation space. This increase in discriminability was measured as an increase in distances between population response vectors for different stimuli (Fig. 3). The second was a global, nonlinear warping of the representation space (Fig. 4C). Both these effects were much larger in AIT than in LIP.

Table 4. Stress between MDS representational spaces and physical stimulus spaces

Comparison	Shape Space		Location Space	
	i. AIT	ii. LIP	iii. AIT	iv. LIP
a. Shape Attention vs Physical	n/a	n/a	0.1792	0.0856
b. Location Attention vs Physical	n/a	n/a	0.1791	0.0812
c. Passive vs Physical	n/a	n/a	0.1332	0.0603

Since a canonical physical shape space does not exist, no values are given for shape space. The values for location space indicate the stress between physical locations of stimuli (Figs. 6D, 7D) and MDS results under different attentional conditions.

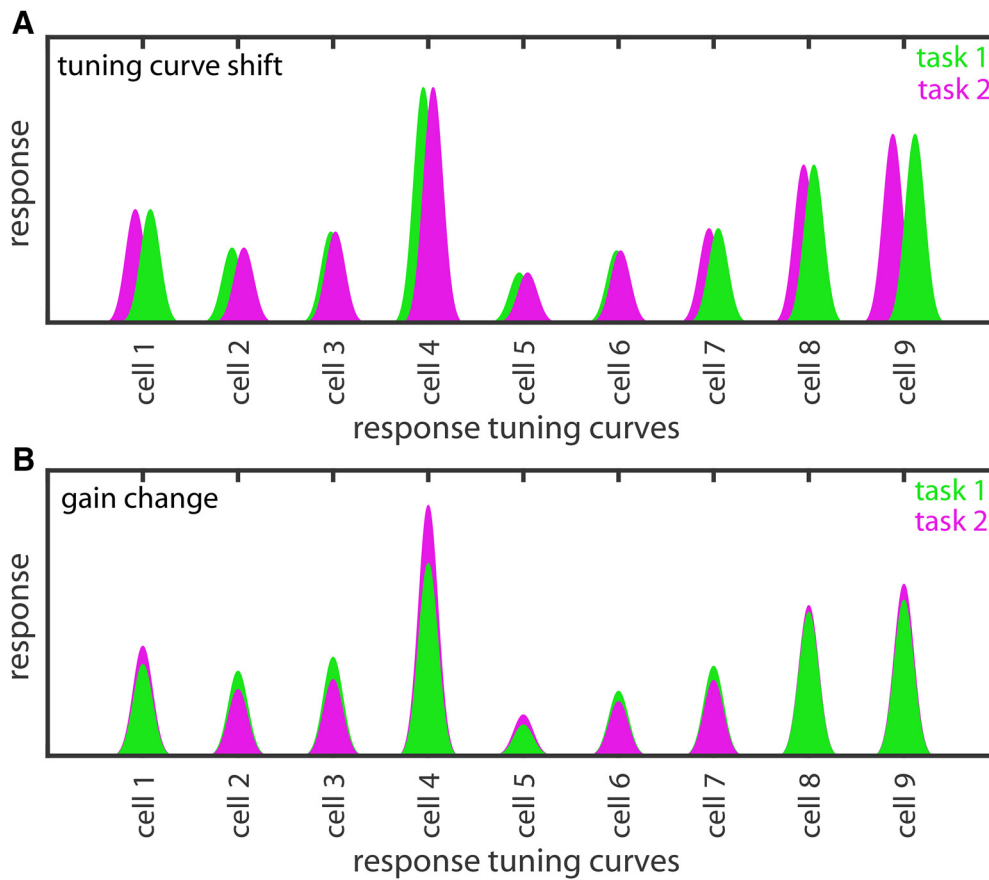


Figure 8. Schematic diagram showing two possible ways how attentional modulations of neural properties across a population could lead to a change in the representation of a stimulus. **A**, Attention dependent shifts in neural tuning curves. **B**, Attention dependent changes in neural gain.

Previous investigations using fMRI (Jiang et al., 2013; Jackson et al., 2017) or evoked potentials (Rotermund et al., 2009) have suggested that attention can change the representation of stimuli, as shown by increased discriminability between pairs of stimuli or pairs of stimulus categories. However, none of these studies went beyond pairwise analyses to perform a full multivariate analysis over a larger set of stimuli to reveal the global organization of perceptual representation spaces. Çukur et al. (2013) have demonstrated warping of semantic representations using multivariate analysis of fMRI data. This study extends those observations by demonstrating the effect in perceptual representations rather than semantic representations, by demonstrating the effect in populations of individual neurons rather than brain imaging data, and by demonstrating differences in these attentional effects in different cortical areas.

Two possible ways that attention could change the population representation of a stimulus are: (1) shifting tuning curves or (2) changing gains (Fig. 8). As reviewed in the introduction, there is evidence for both attentional shifts in tuning curves and gain changes at the single-cell level. Note that either non-uniformity in attentional gain across a population or non-uniformity in tuning curve shifts can cause warping in a representation space. The relative extent to which these two mechanisms actually

make significant contributions to changes in representation caused by attention remains to be investigated.

At the single-cell level, attentional effects were apparent in both AIT and LIP and comparable in the two brain areas, as seen in the attentional gain factors in Figure 2. Observation of such single-cell attentional effects in both visual streams is consistent with previous reports, including those in AIT (REFS here) and those in LIP (REFS here), among many others. (Desimone, 1998; McAdams and Maunsell, 1999; Reynolds et al., 2000; Williford and Maunsell, 2006) in AIT and (Treue and Maunsell, 1996; Ben Hamed et al., 2002; Martínez-Trujillo and Treue, 2002; Womelsdorf et al., 2008; Ibos and Freedman, 2014, 2016) in LIP, among many others. None of these studies examined attentional effects at the population level, as we have done here. Why attentional modulations at the single-cell level do not necessarily translate into strong attentional modulations at the population level may be an interesting topic for future modeling.

We found that attentional effects at the population level were substantially weaker in LIP than AIT. That was true for both population vector distance (Table 2) and the MDS analysis (Table 3). Although at the population level we found that all the significant attentional task comparisons (shape attention vs location attention or attention vs passive) were in AIT and none in LIP at the $p = 0.05$ level,

some LIP comparisons were borderline for significance (in the range $p = 0.05\text{--}0.07$). Given that some population attentional effects in LIP were marginal, we do not feel our findings justify stating that population attentional effects occurred only in AIT but not LIP. It is possible that under different conditions, perhaps boosting signal or reducing noise, one would see small but significant attentional effects in LIP as well. Nevertheless, our findings do indicate that under identical task conditions there were stronger attentional effects at the population level in AIT than in LIP.

As might be expected, in AIT attention to shape versus attention to location increased response distances between shapes, making them more discriminable or distinctive (Fig. 3Bi). Counterintuitively, however, in AIT attention to shape versus location also increased response distances between locations, indicating different spatial locations were more discriminable as well (Fig. 3Biii). In contrast, attention to shape versus location did not increase response distances (discriminability) between locations in LIP, making LIP population changes an unlikely source for the effects in AIT. Instead, these findings raise the possibility that the changes in AIT were specific to AIT and dependent on the object properties themselves. Such an interpretation suggests that both spatial and shape representations as well as attentional effects in one cortical stream are independent of those in the other cortical stream, and are developed, organized, and independently constrained by different goals, attentional demands, or functions. These findings suggest that attention effects might better be conceptualized as more local, non-homogenous systems rather than a centralized system (cf. Patel et al., 2015).

The spatial relationships between objects are important for defining scenes, and that information is present in AIT (Hayworth et al., 2011; Kim and Biederman, 2011). Small spatial translations in objects or scenes are often important and meaningful, as is the spatial arrangement of parts within objects (e.g., open vs closed lock or window). Our findings suggest that attending to shape in AIT, in addition to enhancing shape representations, may directly enhance the representational discriminability or distinctiveness of different spatial arrangements of stimuli, such as objects within scenes or parts within an object.

Our data had a number of limitations. First, population data were recorded serially, one cell at a time, using single electrodes, rather than in parallel with multielectrodes. Thus, cells recorded at different times were likely to be modulated by randomly varying states of general arousal in the monkey. The effect of constructing populations from serial recordings can be considered equivalent to adding uncorrelated noise to our data across different attentional conditions. Second, we had to estimate some missing values to make the data mathematically suitable for population analyses. The existence of estimation errors was again equivalent to adding uncorrelated noise to our data. We believe the net effect of both these limitations was to make our data noisier, thereby indiscriminately reducing statistical p values under all conditions and making it less likely to observe significant effects rather than make it more likely. Nevertheless, we did

observe highly significant p values under some conditions, and just barely missed significance in a few other conditions.

Task difficulty is a factor that can affect neural responses (Spitzer et al., 1988). Previous analysis of behavioral data (Serenio and Amador, 2006) shows significant increases in response time in the shape attention task compared to the location attention task under identical stimulus conditions. Hence, it is possible that there are differences in task difficulty across the two attention tasks (e.g., shape attention task more difficult than location attention task). However, in these data it is unlikely that differences in task difficulty are driving observed changes in stimulus encoding at the population level. First, task difficulty differences were identical regardless where we were recording from and cannot in themselves account for cortical area differences. Second, the passive tasks are much simpler tasks than either attention task. Despite that, in AIT results from the passive task either aligned with the “more difficult” shape task for shape space (Fig. 3Bi) or with the “less difficult” location task for location space (Fig. 3Biii). These findings are difficult to explain by task difficulty. Finally, in LIP, there is not much difference among any of the tasks regardless of task difficulty.

We measured attentional effects during the task sample period when there was only a solitary stimulus present on an otherwise blank screen (aside from the fixation spot). Possibly such a singleton stimulus may have resulted in a ceiling effect for attentional modulations in LIP, contributing to weak attentional modulations under different attentional conditions in the population analyses of LIP data. However, looking at the single-cell data, there was nothing in the gain factor plots (Fig. 2) to suggest that the singleton stimulus resulted in a ceiling effect for attention in LIP but not AIT. In both areas there was a wide range of gain factor values. Also, in both animals, there were significant behavioral differences under different attention tasks (Serenio and Amador, 2006). Nevertheless, further work is needed to examine the effects of distractors on population coding and cortical area differences during attentional modulations.

Our analyses have focused solely on attentional effects on mean firing rates. Attention also has effects on other neural firing statistics. As was mentioned in the introduction, those effects include a decrease in the Fano factor, a decrease in noise correlations between neurons, and an increase in correlated oscillations across a population. We can qualitatively estimate the effects those other statistical modulations had on our results.

The effect of Fano factors was already implicitly included in our MDS analysis, which involves trial-by-trial bootstrap resampling of the data to calculate p values. Lower Fano factors, as occurs during attention, would lead to more precise estimates of mean firing, and in turn larger p values for differences between attentional conditions in the MDS analysis.

Noise correlation between neurons decreases during attention. Looking at neural responses under two attentional conditions meant looking at neural responses during two different trials. Anything that reduced noise

correlation between trials would increase the difference between population activities during the two trials for the two attentional conditions. Therefore, p values for differences between attentional conditions in the MDS analysis would again have increased.

Attention-related correlated oscillations within the cortical inputs to AIT or inputs to LIP would be expected to increase firing rates in AIT or LIP due to increased effectiveness of synaptic transmission from those input areas (Salinas and Sejnowski, 2001; Azouz and Gray, 2003). Correlated oscillations within AIT or LIP themselves would further increase the effectiveness of input synaptic transmission (Fries, 2015; Ni et al., 2016). Any increased firing rates in AIT and LIP caused by attentionally induced correlated oscillations would have already been included in our analyses, leading to increased p values when comparing attentional and passive conditions.

Warping of a representation space is the expected consequence whenever attentional modulations are non-uniform across a population. While we have qualitatively estimated the effect of various firing rate statistics on the mean size of attentional modulations, it is more difficult to predict their effects on the variance of attentional modulations across a population, which is the relevant parameter for determining warping.

Weaker attentional modulations in LIP at the population level might be understood functionally in the context of the distinction made by Goodale and Milner (1992) between vision-for-perception in AIT and vision-for-action in LIP. Perhaps dorsal representations of shape and location are more resistant to cognitive modulations, such as attention, to the extent that these representations are engaged in visuomotor control and therefore need to accurately guide interactions of the body with the physical world. Reports that visuomotor control of grasping is resistant to visual illusions (Haffenden et al., 2001; Aglioti et al., 1995), although controversial (Carey, 2001; Franz, 2001), are consistent with greater cognitive impenetrability in LIP visual processing. Our findings of reduced attentional modulations of population representations in LIP compared to AIT, including alterations of either the representational discriminability of different stimuli or the warping of representation spaces, are consistent with the idea of relative cognitive impenetrability in LIP.

In conclusion, we demonstrate that attention is capable of producing changes in stimulus representations at the population level, going beyond well-studied effects of attention in increasing the salience (or signal-to-noise ratio) of stimuli. We found the strongest population effects of attention in altering representations to be in AIT, where we found that attention can alter both shape and spatial representations, enhancing stimulus discriminability or distinctiveness and also warping representational geometry. In contrast, in LIP, we found relative stability of both shape and location representations under varying attentional conditions. We suggest that reorganizations in the geometry of multidimensional representational spaces in particular brain areas may be an important effect of attention. Our findings suggest that attention will modulate the multidimensional representational spaces in AIT in

complex ways that may expand or shrink certain regions of object space or location space to enhance performance of the task at hand (for example grouping stimuli in different ways during a particular categorization task). Together these findings demonstrate that single-cell findings do not necessarily relate to what is happening at a population level (as both areas show comparable single-cell modulations whereas one area shows substantial and stronger changes in the representations of shape and space compared to the other). Our findings also demonstrate that attention can have independent and locally-specific effects on different cortical areas. Importantly, we demonstrate distinctive and counterintuitive attentional effects in cortical representations, with attention to shape improving the spatial representation in AIT. Finally, these findings reveal plasticity in both population representations of shape and location in AIT due to cognitive factors, findings that help link cellular attentional modulations more tightly with behavior and which are certain to stimulate further investigations.

References

- Aglioti S, Goodale MA, DeSouza JFX (1995) Size-contrast illusions deceive the eye but not the hand. *Curr Biol* 5:679–685. [Medline](#)
- Anton-Erxleben K, Stephan VM, Treue S (2009) Attention reshapes center-surround receptive field structure in macaque cortical area MT. *Cereb Cortex* 19:2466–2478. [CrossRef](#)
- Azouz R, Gray CM (2003) Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron* 37:513–523. [Medline](#)
- Ben Hamed S, Duhamel JR, Bremmer F, Graf W (2002) Visual receptive field modulation in the lateral intraparietal area during attentive fixation and free gaze. *Cereb Cortex* 12:234–245. [Cross-Ref](#)
- Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529–534. [CrossRef](#)
- Carey DP (2001) Do action systems resist visual illusions? *Trends Cogn Sci* 5:109–113. [CrossRef](#)
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600. [CrossRef](#)
- Cohen MR, Maunsell JH (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70:1192–1204. [CrossRef](#)
- Colby CL, Duhamel J-R, Goldberg ME (1996) Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J Neurophysiol* 76:2841–2852. [CrossRef](#)
- Connor CE, Gallant JL, Preddie DC, Van Essen DC (1996) Responses in area V4 depend on the spatial relationship between stimulus and attention. *J Neurophysiol* 75:1306–1308. [CrossRef](#)
- Connor CE, Preddie DC, Gallant JL, Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214. [CrossRef](#)
- Çukur T, Nishimoto S, Huth AG, Gallant JL (2013) Attention during natural vision warps semantic representation across the human brain. *Nat Neurosci* 16:763–770. [CrossRef](#)
- David SV, Hayden BY, Mazer JA, Gallant JL (2008) Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron* 59:509–521. [CrossRef](#)
- Desimone R (1998) Visual attention mediated by biased competition in extrastriate visual cortex. *Philos Trans R Soc Lond B Biol Sci* 353:1245–1255. [CrossRef](#) [Medline](#)
- Eifuku S, De Souza WC, Tamura R, Nishijo H, Ono T (2004) Neuronal correlates of face identification in the monkey anterior temporal cortical areas. *J Neurophysiol* 91:358–371. [CrossRef](#)

- Franz VH (2001) Action does not resist visual illusions. *Trends Cogn Sci* 5:457–459. [Medline](#)
- Fries P (2015) Rhythms for cognition: communication through coherence. *Neuron* 88:220–235. [CrossRef](#)
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563. [CrossRef](#)
- Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci* 28:4823–4835. [CrossRef](#)
- Goodale MA, Milner AD (1992) Separate visual pathways for perception and action. *Trends Neurosci* 15:20–25. [CrossRef](#)
- Haffenden AM, Schiff KC, Goodale MA (2001) The dissociation between perception and action in the Ebbinghaus illusion: nonillusory effects of pictorial cues on grasp. *Curr Biol* 11:177–181. [CrossRef](#)
- Hayworth K, Leszczoart MD, Biederman I (2011) Neural encoding of relative position. *J Exp Psychol Hum Percept Perform* 37:1032–1050. [CrossRef](#) [Medline](#)
- Ibos G, Freedman DJ (2014) Dynamic integration of task-relevant visual features in posterior parietal cortex. *Neuron* 83:1468–1480. [CrossRef](#) [Medline](#)
- Ibos G, Freedman DJ (2016) Interaction between spatial and feature attention in posterior parietal cortex. *Neuron* 91:931–943. [CrossRef](#)
- Ipata AE, Gee AL, Bisley JW, Goldberg ME (2009) Neurons in the lateral intraparietal area create a priority map by the combination of disparate signals. *Exp Brain Res* 192:479–488. [CrossRef](#)
- Ipata AE, Gee AL, Goldberg ME (2012) Feature attention evokes task-specific pattern selectivity in V4 neurons. *Proc Natl Acad Sci USA* 109:16778–16785. [CrossRef](#) [Medline](#)
- Jackson J, Rich AN, Williams MA, Woolgar A (2017) Feature-selective attention in frontoparietal cortex: multivoxel codes adjust to prioritize task-relevant information. *J Cogn Neurosci* 29:310–321. [CrossRef](#)
- Jiang J, Summerfield C, Eger T (2013) Attention sharpens the distinction between expected and unexpected percepts in the visual brain. *J Neurosci* 33:18438–18447. [CrossRef](#)
- Kay KN, Weiner KS, Grill-Spector K (2015) Attention reduces spatial uncertainty in human ventral temporal cortex. *Curr Biol* 25:595–600. [CrossRef](#)
- Kiani R, Esteky H, Mirpour K, Tanaka K (2007) Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *J Neurophysiol* 97:4296–4309. [CrossRef](#)
- Kim JG, Biederman I (2011) Where do objects become scenes? *Cereb Cortex* 21:1738–1746. [CrossRef](#) [Medline](#)
- Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320:110–113. [CrossRef](#)
- Lee J, Maunsell JH (2010) Attentional modulation of MT neurons with single or multiple stimuli in their receptive fields. *J Neurosci* 30:3058–3066. [CrossRef](#) [Medline](#)
- Lehky SR, Sereno AB (2007) Comparison of shape encoding in primate dorsal and ventral visual pathways. *J Neurophysiol* 97:307–319. [CrossRef](#) [Medline](#)
- Lehky SR, Sereno AB (2011) Population coding of visual space: modeling. *Front Comput Neurosci* 4:155. [CrossRef](#)
- Lehky SR, Sereno ME, Sereno AB (2013) Population coding and the labeling problem: extrinsic versus intrinsic representations. *Neural Comput* 25:2235–2264. [CrossRef](#)
- Lehky SR, Sereno ME, Sereno AB (2016) Characteristics of eye-position gain field populations determine geometry of visual space. *Front Integr Neurosci* 9:72. [CrossRef](#)
- Luck SJ, Chelazzi L, Hillyard SA, Desimone R (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J Neurophysiol* 77:24–42. [CrossRef](#)
- Martínez-Trujillo J, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron* 35:365–370. [CrossRef](#)
- Maunsell JH (2015) Neuronal mechanisms of visual attention. *Ann Rev Vis Sci* 1:373–391. [CrossRef](#) [Medline](#)
- Maunsell JH, Treue S (2006) Feature-based attention in visual cortex. *Trends Neurosci* 29:317–322. [CrossRef](#) [Medline](#)
- McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441. [CrossRef](#)
- McAdams CJ, Maunsell JH (2000) Attention to both space and feature modulates neuronal responses in macaque area V4. *J Neurophysiol* 83:1751–1755. [CrossRef](#)
- Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* 55:131–141. [CrossRef](#)
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888. [CrossRef](#)
- Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782–784. [CrossRef](#)
- Motter BC (1993) Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J Neurophysiol* 70:909–919. [CrossRef](#)
- Murata A, Gallese V, Luppino G, Kaseda M, Sakata H (2000) Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. *J Neurophysiol* 83:2580–2601. [CrossRef](#)
- Ni J, Wunderle T, Lewis CM, Desimone R, Diester I, Fries P (2016) Gamma-rhythmic gain modulation. *Neuron* 92:240–251. [CrossRef](#)
- Op de Beeck H, Wagemans J, Vogels R (2001) Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nat Neurosci* 4:1244–1252. [CrossRef](#)
- Patel SS, Red S, Lin E, Sereno AB (2015) Single canonical model of reflexive memory and spatial attention. *Sci Rep* 5:15604. [CrossRef](#)
- Reddy L, Kanwisher NG, VanRullen R (2009) Attention and biased competition in multi-voxel object representations. *Proc Natl Acad Sci USA* 106:21447–21452. [CrossRef](#)
- Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *J Neurosci* 19:1736–1753. [CrossRef](#)
- Reynolds JH, Pasternak T, Desimone R (2000) Attention increases sensitivity of V4 neurons. *Neuron* 26:703–714. [Medline](#)
- Rolls ET, Tovee MJ (1995) Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J Neurophysiol* 73:713–726. [CrossRef](#) [Medline](#)
- Romero MC, Van Dromme IC, Janssen P (2013) The role of binocular disparity in stereoscopic images of objects in the macaque anterior intraparietal area. *PLoS ONE* 8:e55340. [CrossRef](#)
- Rotermund D, Taylor K, Ernst UA, Kreiter AK, Pawelzik KR (2009) Attention improves object representation in visual cortical field potentials. *J Neurosci* 29:10120–10130. [CrossRef](#)
- Salinas E, Sejnowski TJ (2001) Correlated neuronal activity and the flow of neural information. *Nature Rev Neurosci* 2:539–550. [CrossRef](#)
- Sereno AB, Maunsell JHR (1998) Shape selectivity in primate lateral intraparietal cortex. *Nature* 395:500–503. [CrossRef](#)
- Sereno AB, Amador SC (2006) Attention and memory-related responses of neurons in the lateral intraparietal area during spatial and shape-delayed match-to-sample tasks. *J Neurophysiol* 95:1078–1098. [CrossRef](#)
- Sereno AB, Lehky SR (2011a) Population coding of visual space: comparison of spatial representations in dorsal and ventral pathways. *Front Comput Neurosci* 4:159. [CrossRef](#)
- Sereno AB, Lehky SR (2011b) Population decoding of visual space: dorsal versus ventral. Presented at the Gordon research conference on eye movements, Biddeford, ME.
- Sereno AB, Sereno ME, Lehky SR (2014) Recovering stimulus locations using populations of eye-position modulated neurons in dorsal and ventral visual streams of non-human primates. *Front Integr Neurosci* 8:28. [CrossRef](#)

- Spitzer H, Desimone R, Moran J (1988) Increased attention enhances both behavioral and neuronal performance. *Science* 240:338–340. [CrossRef](#)
- Tolias AS, Moore T, Smirnakis SM, Tehovnik EJ, Siapas AG, Schiller PH (2001) Eye movements modulate visual receptive fields of V4 neurons. *Neuron* 29:757–767. [CrossRef](#)
- Treue S, Maunsell JH (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539–541. [CrossRef](#) [Medline](#)
- Treue S, Martínez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579. [CrossRef](#)
- Vogels R, Biederman I, Bar M, Lorincz A (2001) Inferior temporal neurons show greater sensitivity to nonaccidental than to metric shape differences. *J Cogn Neurosci* 13:444–453. [CrossRef](#)
- Williford T, Maunsell JH (2006) Effects of spatial attention on contrast response functions in macaque area V4. *J Neurophysiol* 96:40–54. [CrossRef](#) [Medline](#)
- Womelsdorf T, Anton-Erxleben K, Pieper F, Treue S (2006) Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nat Neurosci* 9:1156–1160. [CrossRef](#)
- Womelsdorf T, Anton-Erxleben K, Treue S (2008) Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *J Neurosci* 28:8934–8944. [CrossRef](#)
- Wróbel A (2014) Attentional activation in corticothalamic loops of the visual system. In: *The new visual neurosciences* (Werner JS, Chalupa LM, eds), pp 339–349. Cambridge, MA: MIT Press.
- Zhang Y, Meyers EM, Bichot NP, Serre T, Poggio TA, Desimone R (2011) Object decoding with attention in inferior temporal cortex. *Proc Natl Acad Sci USA* 108:8850–8855. [CrossRef](#) [Medline](#)