Neuronal Effects of Spatial and Feature Attention Differ Due to Normalization

Amy M. Ni^{1,2} and John H.R. Maunsell^{2,3}

¹Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, ²Department of Neurobiology, Harvard Medical School, Boston, Massachusetts 02115, and ³Department of Neurobiology, The University of Chicago, Chicago, Illinois 60637

Although spatial and feature attention have differing effects on neuronal responses in visual cortex, it remains unclear why. Response normalization has been implicated in both types of attention (Carandini and Heeger, 2011), and single-unit studies have demonstrated that the magnitude of spatial attention effects on neuronal responses covaries with the magnitude of normalization effects. However, the relationship between feature attention and normalization remains largely unexplored. We recorded from individual neurons in the middle temporal area of rhesus monkeys using a task that allowed us to isolate the effects of feature attention, spatial attention, and normalization on the responses of each neuron. We found that the magnitudes of neuronal response modulations due to spatial attention and feature attention are correlated; however, whereas modulations due to spatial attention are correlated with normalization strength, those due to feature attention are not. Additionally, spatial attention modulations are stronger with multiple stimuli in the receptive field, whereas feature attention modulations are not. These findings are captured by a model in which spatial and feature attention share common top-down attention signals that nonetheless result in differing sensory neuron response modulations because of a spatially tuned sensory normalization mechanism. This model explains previously reported commonalities and differences between these two types of attention by clarifying the relationship between top-down attention signals and sensory normalization. We conclude that similar top-down signals to visual cortex can have distinct effects on neuronal responses due to distinct interactions with sensory mechanisms.

Key words: feature attention; macaque; middle temporal; normalization; single unit; spatial attention

Significance Statement

Subjects use attention to improve their visual perception in several ways, including by attending to a location in space or to a visual feature. Prior studies have found both commonalities and differences between the effects of spatial and feature attention on neuronal responses in visual cortex, although it is unclear what mechanisms could explain this range of effects. Normalization, a computation by which neuronal responses are modified by stimulus context, has been implicated in many neuronal mechanisms throughout the brain. Here we propose that normalization provides a simple explanation for how spatial and feature attention could share common top-down attention signals that still affect sensory neuron responses differently.

Introduction

Both spatial and feature attention improve visual perception (Baluch and Itti, 2011). Spatial attention broadly improves perception at a specific location (Posner, 1980), whereas feature attention improves perception of a specific visual feature across all

locations in space (for review, see Carrasco, 2011). Many studies have described the effects of spatial and feature attention on the responses of neurons in visual cortex, but describe both similarities and differences in their effects, leaving uncertain whether they have common or distinct neuronal mechanisms (for review, see Maunsell, 2015).

The feature similarity gain model suggests a unified mechanism of attention in which space is treated as a visual feature (Treue and Martínez-Trujillo, 1999). Neuronal response modulations arising from spatial and feature attention combine additively suggesting that these forms of attention affect neurons in similar ways (Treue and Martínez-Trujillo, 1999; Ibos and Freedman, 2016). Further, spatial and feature attention are both associated with changes in the pairwise correlations between the firing rates of sensory neurons (Cohen and Maunsell, 2009;

Received Aug. 15, 2018; revised April 28, 2019; accepted April 30, 2019.

Author contributions: A.M.N. and J.H.R.M. designed research; A.M.N. performed research; A.M.N. analyzed data; A.M.N. and J.H.R.M. wrote the paper.

This work was supported by National Institutes of Health R01EY005911 (J.H.R.M.) and by a Grant from the Simons Foundation (award 500465 to A.M.N.). We thank Greg Horwitz, Thomas Luo, Patrick Mayo, and Bram Verhoef for helpful comments and discussions, and Anna Chambers, Steven Slebocka, and Vivian Imamura for technical assistance.

The authors declare no competing financial interests.

Correspondence should be addressed to John H. R. Maunsell at maunsell@uchicago.edu.

<https://doi.org/10.1523/JNEUROSCI.2106-18.2019>

Copyright © 2019 the authors

Mitchell et al., 2009; Verhoef and Maunsell, 2017), and the relationship between modulations of firing rate and of pairwise correlations is quantitatively indistinguishable for spatial versus feature attention (Cohen and Maunsell, 2011). Both forms of attention increase gamma frequency synchronization (for review, see Gregoriou et al., 2015) and, finally, spatial and feature attention may share a common top-down source of attention signals (Moore and Armstrong, 2003; Gregoriou et al., 2009; Zhou and Desimone, 2011; but see Paneri and Gregoriou, 2017).

However, other observations point to distinctions between the mechanisms that support spatial versus feature attention. In a task involving both types of attention, trial-to-trial neuronal response fluctuations due to spatial versus feature attention were uncorrelated (Cohen and Maunsell, 2011). Further, the effects of spatial and feature attention on neuronal firing rates have different time courses (Hayden and Gallant, 2005). Finally, feature attention effects spread globally to cortical responses across all locations in visual space (Saenz et al., 2002; Serences and Boynton, 2007) whereas spatial attention effects are retinotopically specific (Womelsdorf et al., 2006), matching the respective behavioral effects (Melcher et al., 2005; Maunsell and Treue, 2006).

How might these two forms of attention share a common top-down source of attention signals that result in differing response modulations at the level of sensory neurons? Models of attention have suggested that the sensory mechanism of normalization plays a role in both spatial and feature attention (Boynton, 2009; Lee and Maunsell, 2009; Reynolds and Heeger, 2009). Electrophysiological recordings of single-neuron responses have demonstrated that differences in normalization strength explain much of the difference in the strength of neuronal response modulations associated with spatial attention (Ni et al., 2012). Whether a normalization model can explain neuronal response modulations due to feature attention has yet to be tested.

Here we recorded from individual neurons in MT of rhesus monkeys while they performed a task that allowed us to independently measure neuronal response changes associated with spatial versus feature attention, as well as the amount of normalization demonstrated by each neuron. We found a neuron-by-neuron correlation between the effects of spatial and feature attention on neuronal responses. However, this correlation was not due to a shared link to normalization strength. Spatial attention modulations were correlated with normalization strength across neurons, but feature attention modulations were not. Additionally, although spatial attention modulations were stronger with multiple visual stimuli in the receptive field (Maunsell, 2015), a finding that has been attributed to normalization (Lee and Maunsell, 2010), this was not the case for feature attention. These findings are explained by a model in which spatial and feature attention share common top-down attention signals that nonetheless result in distinct sensory neuron effects due to a spatially tuned sensory normalization mechanism. The model clarifies why prior studies have found both similarities and differences between the sensory neuron effects of spatial and feature attention.

Materials and Methods

The Harvard Medical School Institutional Animal Care and Use Committee approved all animal procedures. Other findings based on portions of the data described here have been described previously (Ni et al., 2012; Ni and Maunsell, 2017). Specifically, the neuronal data that were collected during the normalization and spatial attention task variants are a subset of the data analyzed in a prior paper that explains why the equal-maximum suppression (EMS) spatially tuned normalization model used here is necessary to describe both across- and within-neuron differences

in normalization strength (Ni and Maunsell, 2017). The current study analyzes neuronal data from a feature attention variant of the behavioral task that have never been published, alongside the previously analyzed neuronal data from the normalization and spatial attention variants of the task (Ni and Maunsell, 2017). The purpose of the current study is to test whether the EMS spatially-tuned normalization model can explain the neuronal correlates of feature attention, and to explore the implication for the relationship between feature attention and normalization.

Behavioral task. Two male rhesus monkeys (*Macaca mulatta*) that weighed 8 and 12 kg. Each received a headpost and a scleral search coil implant under general anesthesia. After recovery, we trained each animal to do a change-detection task that manipulated spatial attention and feature attention.

During each trial of the task, the animal maintained fixation within 1° of a small white dot presented at the center of a video monitor (44 × 34°, 1024 × 768 pixels, 75 Hz refresh rate, gamma-corrected) on a gray background (42 cd/m²). After the animal fixated for 250 ms, an annulus-shaped cue appeared for 250 ms, instructing the animal to attend to one of three possible locations. Two of these locations were within the receptive field (RF) of the neuron being recorded. The third was a symmetric location on the opposite side of the fixation dot. All three locations were the same eccentricity from fixation. The cue was extinguished, and then up to three small drifting Gabors were flashed simultaneously for 200 ms (one Gabor per location), with a blank interstimulus period (varied randomly between each stimulus, from 158 to 293 ms) separating each flash of Gabors (Fig. 1A). The Gabors were flashed on and off until the Gabor at the cued location appeared with a direction that differed by >0° and <90° (a target). The animal was rewarded with juice for making a saccade directly to this target Gabor's location within 100–600 ms of its onset.

The appearance of the target Gabor was timed to follow an exponential distribution (a flat hazard function), encouraging the animal to maintain a constant level of attention across trial time. Direction changes of >0° and <90° also occurred at the two un-cued locations (distractors), with the same probability as the changes at the cued location. If the animal responded to a distractor, the trial was terminated without reward. In ~20% of the trials, the trial lasted 6 s without a target direction change occurring at the cued location (catch trials), in which case the animal received a reward for maintaining fixation for the 6 s.

We initially measured each recorded neuron for motion direction tuning (for RF mapping methods, see Single-unit electrophysiology) and assigned a preferred, a null (opposite), and an intermediate (orthogonal) direction of motion to each neuron. The purpose of the directional stimuli was to evoke differential neuronal responses from the preferred versus the null stimuli, which allowed measurements of normalization and attention effects.

The two Gabors presented at the locations inside the RF fell at locations separated by at least 5 times the SD of the Gabors (Gabor mean SD 0.45°, range 0.42–0.50°, mean separation of Gabor centers 4.2°, range 2.2–6.9°). Because RFs in the middle temporal visual area (MT) are large (Desimone and Ungerleider, 1986), the two stimuli fit easily within the borders of a single MT RF.

For the stimuli presented in the RF, one of six possible stimulus combinations was pseudorandomly assigned on each stimulus presentation flash of a trial (Fig. 1B; Ni and Maunsell, 2017). Three of the stimulus combinations presented the stimuli in Configuration A (preferred stimulus in Location 1, null stimulus in Location 2) and three of the combinations presented the stimuli in Configuration B (null stimulus in Location 1, preferred stimulus in Location 2). Finally, whereas the stimuli presented at the location outside of the RF were always presented at 100% (maximal) contrast, the stimuli presented in the RF were pseudorandomly assigned a contrast of 100, 50, or 0% (i.e., absent).

This task allowed us to measure the neuronal effects of normalization, feature attention, and spatial attention independently. For each recorded neuron, the animal performed three variants of the change-detection task (Fig. 1C–E). The three variants were run in alternating blocks. At least two complete blocks of each task variant were collected for each recorded neuron. The target direction change amount was always <90° and was adjusted independently for each task variant and each neuron

Figure 1. Task design. **A**, During each trial, the monkey was presented with a cue that directed its attention to one of three locations (2 within and 1 outside the RF of the MT neuron being recorded) while series of drifting Gabor stimuli (each having 0, 50 or 100% contrast) were presented on and off simultaneously at the three locations. **B**, On each stimulus presentation, the RF stimuli were pseudorandomly presented in one of six possible stimulus conditions. Three conditions were in Configuration A (the preferred stimulus in Location 1, the null stimulus in Location 2), and three conditions were in Configuration B (the null stimulus in Location 1, the preferred stimulus in Location 2). **C**, To measure normalization strength, attention was cued to stimuli outside of the RF, moving in the intermediate direction of motion. **D**, To measure feature attention modulation, attention was cued to stimuli outside of the RF, moving in the preferred direction of motion. **E**, To measure spatial attention modulation, attention was cued to a location inside of the RF. The monkey received a reward for ignoring changes between the preferred and null directions at each RF location and only detecting a direction change $<90^\circ$ at the cued location.

using an adaptive staircase procedure (Watson and Pelli, 1983) that maintained behavioral performance at $\sim 82\%$ correct.

In the normalization task variant (Fig. 1C), the animal was cued to attend to the location outside of the RF, to a stimulus that always drifted in the intermediate motion direction for the neuron (except during the target changes). By directing the monkey's attention to a distant spatial location in the opposite visual hemifield from the RF and to a constant, intermediate direction of motion, we minimized the effects of spatial and feature attention on our neuronal response measurements. For the measurements of normalization, attention could not differentially modulate the six different stimulus conditions (Fig. 1B), because we used random stimulus sequences and short (200 ms) stimulus flashes so that the animal would not be able to adjust its attention in response to the random content of each flashed set of stimuli (Williford and Maunsell, 2006).

The feature attention task variant (Fig. 1D) was identical to the normalization variant except that the attended stimuli outside of the RF always drifted in the neuron's preferred direction of motion. Comparing responses when the animal attended to the preferred direction of motion to those recorded while the animal attended to an intermediate direction of motion (in the normalization task) allowed us to measure neuronal modulations associated with attention to particular directions.

In the spatial attention task variant (Fig. 1E), the stimuli outside of the RF were presented in the neuron's intermediate direction of motion while the animal was cued to one of the locations within the RF in some blocks of trials, and to the other RF location in other blocks. Because one of six possible stimulus combinations was pseudorandomly assigned to the RF on each stimulus presentation flash (Fig. 1B), the monkey had to ignore direction changes between the preferred and null directions at the cued location of the block (180° changes), and respond only to direction changes $>0^\circ$ and $<90^\circ$ at the cued location. Importantly, the task was designed to prevent feature attention from varying systematically across different stimulus conditions during this spatial attention task variant. Because the animal could not anticipate which stimulus direction would be presented next at its attended location, and because the stimulus flashes were very brief (200 ms), the animal did not have time to adjust its feature attention in response to the stimulus condition pseudorandomly assigned per flash.

Monkey 1 was trained on the entire task at once, including all three task variants (normalization, feature attention, and spatial attention variants). On each day of training, training alternated between the three variants in blocks of trials. Monkey 2 was initially trained on the attend-out condition only (the feature attention and normalization variants). After mastering these variants, Monkey 2 was then trained on the spatial attention variant, but without the pseudorandom presentation of either the preferred or null stimulus at each stimulus location within a single trial. Instead, Monkey 2 was first trained with the preferred stimulus always presented at one location and the null stimulus always presented at the other. After mastering this version of the task variant, Monkey 2 was trained to ignore direction changes between the preferred and null directions at the cued location of the block (180° changes) and to respond only to direction changes $>0^\circ$ and $<90^\circ$ at the cued location.

changes) and to respond only to direction changes $>0^\circ$ and $<90^\circ$ at the cued location.

Single-unit electrophysiology. After the animal was trained on the behavioral task, a recording chamber was implanted to allow electrodes to reach MT from a posterior approach (axis $\sim 22\text{--}40^\circ$ from horizontal in a parasagittal plane). We recorded the activity of isolated neurons with glass-insulated platinum-iridium microelectrodes ($\sim 1\text{ M}\Omega$ at 1 kHz), using a guide tube and grid system (Crist et al., 1988) to penetrate the dura. Extracellular signals were filtered between 250 Hz and 8 kHz, amplified, and digitized at 40 kHz. We isolated action potentials from individual neurons using a window discriminator and recorded spike times with 1 ms resolution.

For each isolated neuron, we estimated the RF location with a hand-controlled visual stimulus. We then measured direction tuning (8 direc-

tions) and temporal frequency tuning (5 frequencies) with computer-controlled presentations of Gabor stimuli while the animal performed a fixation task. The direction that produced the strongest response was taken as the preferred direction, the direction 180° from the preferred direction was taken as the null direction, and one direction 90° from the preferred direction was used as the intermediate direction. All of the Gabors were presented at the same temporal frequency, the one that produced the strongest average response. The temporal frequency was rounded to a value that produced an integral number of cycles of drift during each stimulus presentation. In this way, the Gabors could start and end with odd spatial symmetry, so that the spatiotemporal integral of the luminance of each stimulus was the same as the background. The spatial frequency of all of the Gabors was fixed at 1 cycle per degree. A Gabor moving in the preferred direction and at the preferred temporal frequency was used to quantitatively map the RF (3 eccentricities, 5 polar angles) while the animal performed a fixation task. The RF eccentricities ranged from 5 to 14° for Monkey 1 and from 7 to 14° for Monkey 2. The two stimulus locations within the RF were chosen to give approximately equal responses. Direction selectivity was indistinguishable between the stimulated RF locations (Livingstone et al., 2001; Location 1 average direction selectivity modulation index = 1.1; Location 2 average = 1.1; paired *t* test: $t_{(56)} = 0.16$, $p = 0.87$).

Experimental design and statistical analysis. In total, we analyzed the responses of 57 individual MT neurons (42 from Monkey 1, 15 from Monkey 2). For each neuron, responses to the two stimulus configurations (Configuration A and Configuration B; Fig. 1B) were treated separately in all analyses, yielding two observations per neuron. This was done because normalization strength varies with stimulus configuration if stimuli presented at different RF locations contribute different weights to normalization (see Results; Ni and Maunsell, 2017). Thus, two “unit-configurations” were collected per neuron.

We included recorded neurons in the analyses if they had driven firing rates that were significantly greater than their spontaneous firing rates, and if we collected data from at least two blocks each of the normalization, feature attention, and spatial attention task variants. Additionally, we excluded three neurons from all analyses based on the application of the Tukey method if they had exceptional values for the modulation indices for normalization (NMI), feature attention (FMI), or spatial attention (SMI): we excluded one neuron with an NMI of -1.49 , one neuron with an FMI of 1.52 , and one neuron with an SMI of 3.41 . All excluded neurons were from Monkey 1. Only correct response trials were included in the analyses (correct change-detection trials and catch trials). Of the stimulus presentations on those trials, we excluded stimuli if they were presented at the same time as a target or distractor stimulus, if they appeared after the target presentation, or if they were presented within 400 ms of the start of the stimulus series. Excluding the first one or two stimulus presentations of each trial reduced variance arising from stronger responses to the initiation of the stimulus series. After stimulus exclusions, ~ 13 repetitions of the 6 stimulus conditions were collected per block.

Figure 2. Spatial and feature attention effects were correlated across unit-configurations. **A**, For unit-configuration Example 1, a PSTH illustrates the average response to the preferred and null stimuli together in the RF when attention was directed outside of the RF to the intermediate direction of motion (thick black line). Spatial attention directed inside of the RF resulted in a strong increase in neuronal response when the preferred stimulus was presented at the attended location (orange line; modulation index of 0.21). Feature attention directed to the preferred direction of motion also resulted in a strong increase in neuronal response when compared with attention to the intermediate direction of motion (green line; modulation index of 0.15). The orange shaded area represents the increase in neuronal response due to the animal directing spatial attention into the RF (quantified as SMI). The green shaded area represents the increase in neuronal response due to the animal directing its feature attention to the preferred feature (quantified as FMI). Both modulations due to spatial attention and to feature attention are compared against the same baseline response (thick black line). **B**, For unit-configuration Example 2, both spatial and feature attention had much smaller effects (SMI: 0.04; FMI: 0.01). **C**, Example 3 illustrates a unit-configuration that is strongly modulated due to spatial attention (SMI: 0.25) but that does not demonstrate any modulation due to feature attention (FMI: 0.00). **A–C**, Thick bars along the x-axis indicate the timing of the stimulus presentations. Thin gray lines indicate the spontaneous firing rate (both stimuli at 0% contrast during the normalization task). Each PSTH was smoothed by a Gaussian filter (SD: 10 ms). **D**, Feature and spatial attention modulations were correlated across all unit-configurations ($n = 114$) with two stimuli in the RF. **E**, Feature and spatial attention modulations were correlated with one stimulus in the RF. **D**, **E**, Monkey 1 unit-configurations are illustrated by gray circles, whereas Monkey 2 unit-configurations are illustrated by black circles.

We calculated neuronal firing rates using a 200 ms window starting 50 ms after stimulus presentation onset and ending 50 ms after stimulus presentation offset. Other than the peristimulus time histograms (PSTHs) plotted in Figure 2, which plot the neuronal firing rate, analyses used driven rate. Driven rate was calculated per stimulus presentation by subtracting the average baseline firing rate (the firing rate when both stimuli in the RF were presented at 0% contrast; that is, when no stimuli were presented in the RF) from the firing rate for that stimulus presentation. We calculated the average response rate per stimulus condition as the mean driven rate across all included stimulus repetitions.

P values were computed for Pearson's linear correlation coefficients using a Student's *t* distribution (“corr”; MATLAB R2015a, MathWorks). Confi-

dence intervals (95%) for partial correlation coefficients were determined using a Fisher's z transform, accounting for the reduction in degrees of freedom associated with the additional predictors. All t tests were two-sided. Mean matching of modulation indices was performed using previously described methods (Churchland et al., 2010). As outlined by Churchland et al. (2010), we computed the distribution of each attention condition, determined the greatest common denominator present during both conditions (with each bin of the common distribution having a height equal to the smallest value for that bin across all distributions), and randomly excluded unit-configurations from each bin until the height of that bin matched that of the common distribution. We repeated this process 100 times with different random seeds and calculated the mean across those 100 iterations. The spatial attention modulation indices of the mean-matched spatial attention subpopulation were corrected for the expected regression artifact from regression to the mean (The Research Methods Knowledge Base, <http://www.socialresearchmethods.net/kb/>).

We fit the model parameters for each of the proposed model equations via constrained nonlinear optimizing that minimized the sum-of-squares error ("fminsearch"; MATLAB R2015a, MathWorks). The model parameters α and β were constrained in the fit for the absolute values to be <10 (based on values previously described; Ni et al., 2012). The other model parameters were unconstrained. For each unit-configuration (see Experimental design and statistical analysis), we calculated the goodness of fit of each model as the total explained variance, which was determined by taking the square of the correlation coefficient between the estimated response rates from the model and the response rates of the unit-configuration across the stimulus conditions fit by the model. To quantify differences in explained variance between two tested models, we calculated an F statistic based on the residual sum of squares (RSS) and degrees of freedom of Models 1 and 2. We used the following equation to compare two models with the same number of parameters:

$$F = \text{RSS}_1 / \text{RSS}_2 \quad (1)$$

P values were computed per F statistic using the built-in function "fcdF" (MATLAB R2015a, MathWorks).

Additionally, we calculated the variance explained by the model for each modulation index. Per modulation index, we determined the correlation between the modulation indices for each unit-configuration based on the modeled neuronal responses and the modulation indices for each unit-configuration based on the actual neuronal responses. To assess the relative importance of each free parameter to the variance in the modulation indices explained by the model, we tested the effects of locking the parameters one at a time to their average fit when fit as a free parameter. We tested the significance of the reduction in variance explained due to locking a free parameter using Williams' procedure for comparing correlated correlation coefficients (Howell, 2007). We also used Williams' procedure to test the significance of the reduction in variance explained due to using the EMS stimulus-tuned normalization model instead of the EMS spatially-tuned normalization model.

Results

We analyzed the visual responses of 57 individual MT neurons (42 from Monkey 1, 15 from Monkey 2). Each neuron's responses were measured during a change-detection task (Fig. 1A) with two different stimulus configurations, Configuration A and Configuration B (Fig. 1B). Because normalization strength depends on stimulus configuration (Ni and Maunsell, 2017), the two stimulus configurations needed to be analyzed separately for this study. Thus, two unit-configurations were collected per neuron, resulting in an n of 114 unit-configurations.

Behavior

We used three variants of a direction change-detection task to measure the effects of normalization, feature attention, and spatial attention on neuronal response rates (see Materials and Methods). The size of the direction change to be detected was always $<90^\circ$ and was adjusted independently for each task vari-

ant using an adaptive staircase procedure to ensure that all the tasks were comparably challenging. The resulting detection thresholds for the normalization (Fig. 1C), feature attention (Fig. 1D), and spatial attention task variants (Fig. 1E) were 49, 50, and 55° for Monkey 1 and 42, 39, and 42° for Monkey 2.

Spatial and feature attention modulations were correlated

For each unit-configuration, we calculated a spatial attention modulation index (SMI) and a feature attention modulation index (FMI) using driven rates, both with a preferred and a null stimulus in the RF (Fig. 1B, Preferred + Null):

$$MI = \frac{A - B}{A + B} \quad (2)$$

For the SMI, A was the average neuronal driven rate (firing rate – baseline rate) with spatial attention directed inside of the RF, when the preferred stimulus was presented at the attended location. For the FMI, A was the average driven rate with feature attention directed to the preferred direction of motion outside of the RF. The reference for both measures (B) was the average driven rate with spatial attention directed outside of the RF to the intermediate direction of motion.

With two stimuli in the RF, there was a relationship between spatial and feature attention modulations across all unit-configurations. Figure 2A shows the average PSTHs of a representative unit-configuration that showed moderate modulations associated with feature attention (FMI: 0.15; green shaded area) and somewhat more modulation with spatial attention (SMI: 0.21; orange shaded area). Figure 2B illustrates a representative unit-configuration with weak spatial (SMI: 0.04) and feature attention (FMI: 0.01) modulations. Finally, Figure 2C illustrates a representative unit-configuration that is strongly modulated due to spatial attention (SMI: 0.25) but poorly modulated due to feature attention (FMI: 0.00). SMIs and FMIs were correlated across all 114 unit-configurations (Fig. 2D; Pearson's correlation coefficient: $r = 0.41$, $p = 7.5 \times 10^{-6}$; Monkey 1 only: $r = 0.40$, $p = 1.5 \times 10^{-4}$; Monkey 2 only: $r = 0.43$, $p = 0.018$), with spatial attention typically associated with stronger modulation.

As expected, in this dataset the two different stimulus configurations (Fig. 1B) resulted in two different SMIs per neuron (paired t test: $t_{(56)} = 4.9$, $p = 8.3 \times 10^{-6}$), although the stimulus presentations for the two configurations were pseudorandomly interleaved within each trial. SMIs and FMIs were also correlated when only analyzing responses collected using Configuration A ($n = 57$, $r = 0.42$, $p = 1.3 \times 10^{-3}$), or when only analyzing responses collected using Configuration B ($n = 57$, $r = 0.53$, $p = 2.5 \times 10^{-5}$).

This SMI/FMI correlation could reflect a commonality between the mechanisms supporting spatial versus feature attention. Alternatively, because each neuron was typically collected on a different day, the monkeys might have allocated more or less attentional effort while data were collected from different neurons. However, there are other, less interesting possibilities. The correlation could arise simply from variance in the strength of direction selectivity across unit-configurations, because those with little direction selectivity would show little attention-related modulation. A unit-configuration with little direction selectivity would not be modulated by directing feature attention to the preferred stimulus, and that unit-configuration would also not be modulated by the addition of a null stimulus to a RF containing a preferred stimulus resulting in little normalization, and little spatial attention modulation with attention directed to the

location of the preferred stimulus (Ni et al., 2012). On the other hand, a unit-configuration with strong direction selectivity would show much stronger modulations due to feature attention, normalization, and spatial attention. Thus, a range in direction selectivity strength alone could masquerade as a correlated range in attention modulation strengths. However, this scenario is unlikely to be the cause of the SMI/FMI correlation, because almost all MT neurons are highly direction selective with this selectivity preserved across the entire RF (Maunsell and van Essen, 1983; Albright, 1984; Tanaka et al., 1986; Richert et al., 2013; but see Cui et al., 2013), and as within-neuron differences in SMI between different stimulus configurations do not appear to depend on within-neuron differences in direction selectivity (Ni and Maunsell, 2017). Nevertheless, to test this possibility, we calculated a direction selectivity modulation index (DMI) for each unit-configuration using Equation 2. The DMI was based on the driven rates when attention was directed outside of the RF to the intermediate motion direction, comparing the average rate with the preferred stimulus alone in the RF (term *A*) to the average rate with the null stimulus alone in the RF (term *B*). The partial correlation between SMIs and FMIs was little affected when controlling for DMIs ($r = 0.39$, $p = 2.0 \times 10^{-5}$).

The correlation between SMIs and FMIs was similarly unaffected when controlling for differences in mean spontaneous firing rates (measured during stimulus intervals when both stimuli were assigned 0% contrast; $r = 0.41$, $p = 5.0 \times 10^{-6}$) or when controlling for differences in overall responsiveness (measured as the response of each unit-configuration to the preferred stimulus alone in the RF; $r = 0.40$, $p = 1.1 \times 10^{-5}$).

We also calculated response modulations due to spatial and feature attention with one stimulus in the RF (Fig. 1*B*, Preferred) using Equation 2. Even with one RF stimulus, spatial (SMI-1s) and feature attention modulation indices (FMI-1s) were correlated (Fig. 2*E*; $r = 0.53$, $p = 1.1 \times 10^{-9}$; Monkey 1 only: $r = 0.58$, $p = 2.6 \times 10^{-6}$; Monkey 2 only: $r = 0.65$, $p = 1.1 \times 10^{-4}$). SMI-1s and FMI-1s were also correlated when only analyzing the responses collected in Configuration A ($r = 0.58$, $p = 2.6 \times 10^{-6}$) or in Configuration B ($r = 0.47$, $p = 2.5 \times 10^{-4}$), or when controlling for variance in DMI ($r = 0.54$, $p = 8.8 \times 10^{-10}$), spontaneous firing rate ($r = 0.53$, $p = 1.9 \times 10^{-9}$), or neuronal responsiveness ($r = 0.51$, $p = 1.0 \times 10^{-8}$).

Normalization does not explain the correlation between spatial and feature attention effects

Prior electrophysiological studies have found that a large fraction of the neuron-to-neuron variance in spatial attention modulations depends on variance in the strength of normalization across neurons (Lee and Maunsell, 2009; Ni et al., 2012; Verhoef and Maunsell, 2016). Normalization is a mechanism that determines how neurons sum responses to multiple stimuli and has been observed throughout the brain (for review, see Carandini and Heeger, 2011). In the visual system, normalization is seen as a suppression of neuronal responses that increases with stimulus contrast (Heeger, 1992; Heeger et al., 1996), in cross-orientation inhibition (Morrone et al., 1982; Bonds, 1989; DeAngelis et al., 1992; Carandini et al., 1997), and in neuronal responses to multiple stimuli within a neuron's RF (Britten and Heuer, 1999; Heuer and Britten, 2002).

Modeling studies have suggested that both spatial and feature attention-related neuronal modulations may depend on the normalization mechanism (Boynton, 2009; Lee and Maunsell, 2009; Reynolds and Heeger, 2009). This suggests that the correlation between SMI and FMI might arise from a common relationship

to the strength of normalization. Here, we used electrophysiology to test whether the strength of normalization can account for the correlation between spatial and feature attention-related response modulations.

We measured the suppressive strength of normalization in response to increasing contrast by comparing the response to a single stimulus in the RF to the response to two stimuli in the RF (Morrone et al., 1982; DeAngelis et al., 1992; Heeger et al., 1996; Carandini et al., 1997). Typically, the response to a single, preferred stimulus in the RF is reduced when a weakly excitatory null stimulus is added to the RF. Although a linear model that sums inputs would predict a small increase in response from the addition of the weakly excitatory null stimulus, the divisive normalization model explains this response suppression with a divisive inhibition of the response that increases with total contrast in the RF.

To quantify the strength of normalization for each unit-configuration, we calculated an NMI using Equation 3:

$$NMI = \frac{(P + N) - PN}{(P + N) + PN} \quad (3)$$

Compared with the linear sum of the responses to the preferred (*P*) and null (*N*) stimuli when they are presented separately in the RF (*P*+*N*), normalization models predict a smaller neuronal response when both stimuli are presented together in the RF (*PN*), based on suppressive normalization factors associated with each stimulus (Morrone et al., 1982; DeAngelis et al., 1992; Heeger et al., 1996; Carandini et al., 1997). Unlike the NMI equation used in a prior study (Ni et al., 2012), first calculating the linear sum of the individually recorded driven rates in Equation 3 accounts for any potential suppression on the baseline rate from the null stimulus presented alone in the RF.

Studies in both V4 and MT have shown that the responses of some neurons are greatly reduced when a null stimulus is also in the RF, whereas the responses of other neurons are little reduced by this increase in contrast (Lee and Maunsell, 2009; Ni et al., 2012; Verhoef and Maunsell, 2016). This normalization variance has also been found between unit-configurations (Ni and Maunsell, 2017).

We found that normalization strength could not explain the correlation between spatial and feature attention modulation strengths across unit-configurations, both with two stimuli and one stimulus in the RF. With two stimuli, the correlation between SMIs and FMIs was unaffected when controlling for NMIs (partial correlation coefficient controlling for the variance in NMIs: $r = 0.41$, $p = 7.8 \times 10^{-6}$; Monkey 1 only: $r = 0.44$, $p = 3.7 \times 10^{-5}$; though the partial correlation coefficient was not significant for Monkey 2 only: $r = 0.27$, $p = 0.15$). Similarly, with one stimulus in the RF, the correlation between SMI-1s and FMI-1s was unaffected when controlling for NMIs (partial correlation coefficient controlling for the variance in NMIs: $r = 0.56$, $p = 1.8 \times 10^{-10}$; Monkey 1 only: $r = 0.56$, $p = 3.9 \times 10^{-8}$; Monkey 2 only: $r = 0.64$, $p = 1.7 \times 10^{-4}$). As expected (Ni and Maunsell, 2017), NMIs differed based on stimulus configuration within neurons (paired *t* test: $t_{(56)} = 5.7$, $p = 4.1 \times 10^{-7}$) although the monkeys did not have time to adjust their attention in response to the randomly selected configuration of each quickly flashed set of stimuli within each trial. The correlation between SMIs and FMIs for the responses collected in Configuration A alone was little affected when controlling for NMIs collected in that configuration ($r = 0.42$, $p = 1.4 \times 10^{-3}$), as was the correlation for responses collected in Configuration B alone ($r = 0.44$, $p = 6.0 \times 10^{-4}$). The same was true for the correlations between SMI-1s

and FMI-1s collected in Configurations A ($r = 0.59$, $p = 1.4 \times 10^{-6}$) and B ($r = 0.47$, $p = 2.8 \times 10^{-4}$).

Spatial and feature attention differ in their relationship to normalization

As shown previously (Lee and Maunsell, 2009; Ni et al., 2012; Verhoef and Maunsell, 2016; Ni and Maunsell, 2017), spatial attention modulation strength was correlated with normalization strength (Fig. 3A; $r = 0.52$, $p = 2.7 \times 10^{-9}$; Monkey 1 only: $r = 0.51$, $p = 6.5 \times 10^{-7}$; Monkey 2 only: $r = 0.54$, $p = 2.1 \times 10^{-3}$; Configuration A only: $r = 0.37$, $p = 4.2 \times 10^{-3}$; Configuration B only: $r = 0.61$, $p = 5.5 \times 10^{-7}$). Here we add that this correlation persists when controlling for FMIs ($r = 0.52$, $p = 3.0 \times 10^{-9}$; Monkey 1 only: $r = 0.54$, $p = 1.8 \times 10^{-7}$; Monkey 2 only: $r = 0.44$, $p = 0.016$; Configuration A only: $r = 0.38$, $p = 4.2 \times 10^{-3}$; Configuration B only: $r = 0.55$, $p = 1.3 \times 10^{-5}$).

On the other hand, feature attention modulation strength was not significantly correlated with normalization strength (Fig. 3B; $r = 0.12$, $p = 0.21$). When calculated per monkey, there was a modest correlation between FMI and NMI for Monkey 2 only (Monkey 2 only: $r = 0.40$, $p = 0.027$; Monkey 1 only: $r = 0.054$, $p = 0.63$); however, this relationship went away in partial correlations that controlled for SMI (Monkey 2 only: $r = 0.23$, $p = 0.24$; Monkey 1 only: $r = -0.19$, $p = 0.08$; all unit-configurations: $r = -0.12$, $p = 0.20$). Within configurations, there was a modest correlation between FMI and NMI in Configuration B only (Configuration B only: $r = 0.32$, $p = 0.016$; Configuration A only: $r = 0.077$, $p = 0.57$); however, this relationship went away in partial correlations that controlled for SMI (Configuration B only: $r = -5.1 \times 10^{-3}$, $p = 0.97$; Configuration A only: $r = -0.092$, $p = 0.50$).

Although feature attention modulations were significant at the population level (t test: $t_{(113)} = 4.6$, $p = 1.2 \times 10^{-5}$), they were weaker than SMIs (average 0.05 vs 0.20; paired t test: $t_{(113)} = 8.0$, $p = 9.5 \times 10^{-13}$), and might have failed to show dependence on NMIs only because of a low signal-to-noise. This would be surprising, because FMIs were correlated with SMIs and SMIs did not differ from NMIs in magnitude (paired t test: $t_{(113)} = 1.7$, $p = 0.094$). However, to address this possible explanation, we mean-matched SMIs to FMIs (Churchland et al., 2010). This analysis retained 43% of the unit-configurations. Averaging across 100 iterations of random unit-configuration exclusion from each bin of the SMI distribution until bin values were matched to that of the intersection of the SMI and FMI distributions, the SMIs of the mean-matched population were still correlated with NMIs ($r = 0.46$, $p = 1.0 \times 10^{-3}$). The SMIs of the mean-matched population were still correlated with NMIs when controlling for FMIs ($r = 0.42$, $p = 3.5 \times 10^{-3}$). The fact that the SMIs of weakly modulated unit-configurations were correlated with NMI suggests that the weak modulations caused by feature attention do not explain why FMIs were uncorrelated with NMIs.

The partial correlation between spatial attention modulation strength and normalization strength (controlling for feature attention modulation strength) was significantly different from the partial correlation between feature attention modulation strength and normalization strength (controlling for spatial attention modulation strength), and the same was true for the mean-matched population. Figure 3C shows that the 95% confidence interval for the partial correlation between SMIs and NMIs (either for the full population or for the mean-matched population), determined using a Fisher's z transform, did not overlap the 95% confidence interval for the partial correlation between FMIs and NMIs. In summary, spatial and feature attention-

A Spatial attention vs. normalization

B Feature attention vs. normalization

D Preferred vs. null Feature attention vs. normalization

Figure 3. Spatial but not feature attention modulations were correlated with normalization strength. **A**, Spatial attention and normalization modulation indices were correlated across all unit-configurations ($n = 114$), **B** whereas feature attention and normalization modulation indices were not. **C**, Partial correlation coefficients are plotted with 95% confidence interval error bars for the correlation between spatial attention and normalization modulation indices while controlling for the variance in feature attention modulation indices (Spat), for the correlation between mean-matched spatial attention and normalization modulation indices while controlling for the variance in feature attention modulation indices (MM Spat), and for the correlation between feature attention and normalization modulation indices while controlling for the variance in spatial attention modulation indices (Feat). **D**, Feature attention and normalization modulation indices were still not correlated across all unit-configurations when the sizes of feature attention modulation indices were maximized by calculating the difference in neuronal response when the monkey attended the preferred feature versus the null feature ($n = 113$; outlier with a Pref/Null FMI of 3.1 and an NMI of -0.10 was removed, but see text for analyses including the outlier). **A**, **B**, **D**, Monkey 1 unit-configurations are illustrated by gray circles, whereas Monkey 2 unit-configurations are illustrated by black circles.

related neuronal modulations, although correlated with each other, have differing relationships with tuned normalization.

Additionally, spatial and feature attention modulation indices were calculated by comparing the effects of both types of atten-

tion on the same stimulus condition: two stimuli in the RF with attention directed outside of the RF to stimuli moving in the intermediated direction between the preferred and null directions. By using the same baseline neuronal response for both types of attention, we were able to directly compare the effects of adding just spatial attention to the neuronal response to the effects of adding just feature attention to the same baseline. Adding spatial attention just required switching attention from outside to inside of the RF. This isolated the effects of spatial attention, because feature attention was not modulated with attention inside of the RF because the stimuli were flashed too quickly for the monkey to adjust its attention based on which stimulus was pseudorandomly presented at the attended RF location. Adding feature attention just required switching attention to the preferred stimulus, without moving the spatial location of attention. We could not use the condition when the monkey attended outside of the RF to the null stimulus as the baseline, because this would introduce feature attention into the SMIs.

However, our calculation of modulation indices did not maximize the strength of FMIs. Here, we tested whether maximized FMIs were correlated with NMIs. We calculated Pref/Null FMIs by comparing neuronal responses when the monkey attended out to the preferred stimulus to neuronal responses when the monkey attended out to the null stimulus (instead of to the intermediate stimulus). The modulation size increased significantly (from an average of 0.05 to an average of 0.13; paired t test: $t_{(113)} = 2.6$, $p = 0.011$). However, Pref/Null FMIs were still not correlated with NMIs: $r = -0.048$, $p = 0.61$ (Configuration A only: $r = -0.16$, $p = 0.23$; Configuration B only: $r = 0.025$, $p = 0.85$). Pref/Null FMIs were also not correlated with NMIs when controlling for the variance in SMIs: $r = -0.18$, $p = 0.053$ (Configuration A only: $r = -0.25$, $p = 0.065$; Configuration B only: $r = -0.080$, $p = 0.56$).

Pref/Null FMIs did include an outlier with a value of 3.1. This outlier had a corresponding NMI of only -0.10 , and may have affected the relationship between Pref/Null FMIs and NMIs. We removed the outlier and again tested whether there was any relationship between Pref/Null FMIs and NMIs. However, there was no detectable relationship, even with the outlier removed (Fig. 3D; $r = 0.14$, $p = 0.15$; Configuration A only: $r = 0.11$, $p = 0.40$; Configuration B only: $r = 0.025$, $p = 0.85$).

Spatial but not feature attention effects increase with multiple stimuli

An earlier study (Lee and Maunsell, 2010) hypothesized that spatial attention modulations are stronger with both a preferred and a non-preferred stimulus present in the RF because that configuration allows normalization to amplify modest changes in input associated with spatial attention (Morrone et al., 1982; DeAngelis et al., 1992; Heeger et al., 1996; Carandini et al., 1997). Although studies have indeed reported stronger spatial attention modulations with multiple stimuli in the RF (for review, see Bisley, 2011; Maunsell, 2015), this comparison has not been made for feature attention while maintaining task difficulty across different stimulus conditions. To keep task difficulty stable across different stimulus conditions, including conditions with one stimulus and with two stimuli in the RF (Fig. 1B), we used random stimulus sequences and short (200 ms) stimulus flashes (Williford and Maunsell, 2006).

Across all unit-configurations, spatial attention modulations with two stimuli in the RF were much larger than with one stimulus (average SMI 0.20, average SMI-1 0.10; paired t test: $t_{(113)} = 5.7$, $p = 1.2 \times 10^{-7}$; Fig. 4A; Monkey 1 only: $t_{(83)} = 5.9$, $p = 8.8 \times 10^{-8}$; though the difference was not significant for Monkey 2

A Spatial attention: 1 vs. 2 stim

B Feature attention: 1 vs. 2 stim

Figure 4. Attention modulations due to spatial but not feature attention were stronger with multiple stimuli in the RF. **A**, Spatial attention modulation indices were greater with two stimuli than with one stimulus in the RF. **B**, Feature attention modulation indices were not greater with two stimuli in the RF. **A**, **B**, Monkey 1 unit-configurations are illustrated by gray circles, whereas Monkey 2 unit-configurations are illustrated by black circles.

only: $t_{(29)} = 1.5$, $p = 0.14$). Whereas SMIs were correlated with NMIs, SMI-1s were not correlated with NMIs ($r = 0.067$, $p = 0.48$; Monkey 1 only: $r = 0.087$, $p = 0.43$; Monkey 2 only: $r = -0.089$, $p = 0.64$).

On the other hand, feature attention modulations were not stronger with two stimuli in the RF than with one stimulus (average FMI 0.05, average FMI-1 0.04; paired t test: $t_{(113)} = 1.2$, $p = 0.25$; Fig. 4B; Monkey 1 only: $t_{(83)} = 1.1$, $p = 0.28$; Monkey 2 only: $t_{(29)} = 0.47$, $p = 0.64$). Correspondingly, FMI-1s were not correlated with NMIs ($r = -0.18$, $p = 0.051$; Monkey 1 only: $r = -0.21$, $p = 0.055$; Monkey 2 only: $r = -0.15$, $p = 0.43$).

Spatial attention modulations may be stronger, whereas feature attention modulations are not because SMI-1s were larger than FMI-1s to begin with. To address this possibility, we mean-matched SMI-1s to FMI-1s (Churchland et al., 2010), retaining 52% of the unit-configurations. Even with the subpopulation in which the mean SMI-1 matched the mean FMI-1 of the whole population, spatial attention modulations were stronger with two stimuli in the RF (paired t test: $t_{(58)} = 2.9$, $p = 9.2 \times 10^{-3}$, after correction for expected regression to the mean).

Normalization model explains commonalities and differences between attention types

Here we demonstrate that the EMS spatially-tuned normalization model (Ni and Maunsell, 2017) that accounts for both across- and within-neuron differences in normalization strength also accounts for both the correlation between spatial and feature attention mod-

ulations and the difference in their relationship to normalization. We first briefly describe how the model adequately explains the effects of normalization, without attention (but see Ni and Maunsell, 2017 for a full explanation of this model). The normalization data analyzed here are a subset of the data analyzed by Ni and Maunsell (2017), and are only analyzed again here for explanatory purposes. We fit each unit-configuration's responses to eight stimulus conditions (the 3 stimulus conditions for each stimulus configuration illustrated in Fig. 1B; P , N , and PN , with all possible combinations of each stimulus appearing at either 50 or 100% contrast) presented during the normalization task variant (Fig. 1C), using Equation 4:

$$R_{P,N} = \frac{c_P L_P}{c_P + c_N \alpha_2 + \sigma} + \frac{c_N L_N}{c_P \alpha_1 + c_N + \sigma}. \quad (4)$$

The model has five free parameters: L_P , L_N , α_1 , α_2 , and σ . In the numerators, c_P and c_N are the contrasts of the preferred and null stimuli, and L_P and L_N set the excitatory drive associated with the preferred and null stimuli. In the denominators, c_P and c_N represent the contrast-dependent suppressive drive associated with each stimulus, and σ corresponds to the semi-saturation contrast when a single stimulus is presented. The suppression associated with each stimulus differs in magnitude between that stimulus and the distant stimulus (thus it is spatially tuned), as determined by α_1 and α_2 . This model is named the EMS spatially-tuned model because each stimulus generates equally strong suppression of its own excitatory drive but less suppression of stimuli in other locations.

The EMS spatially-tuned model accounted for an average of 97% of the data variance in the observed neuronal responses to the eight stimulus conditions, across all unit-configurations (Ni and Maunsell, 2017). We will use this explained variance calculation to compare the explanatory power of each equation. Additionally, we will calculate how well the model explains the variance in a particular modulation index to determine the relative importance of each free parameter to that modulation index. For Equation 4, we found that the model accounted for 89% of the variance in NMIs across all unit-configurations, based on the correlation between NMIs calculated using the modeled responses for each unit-configuration and NMIs calculated using the observed neuronal responses ($r = 0.94$, $p = 6.6 \times 10^{-56}$). To assess the relative importance of each of the five free parameters to the variance in NMIs explained by the model, we tested the effects of locking the parameters one at a time to their average value when fit as a free parameter. Locking the α_2 parameter had the most detrimental effect on how well the model explained the recorded variance in NMIs across the population. Locking this parameter reduced the variance explained from 89 to 17% (Williams' procedure for comparing correlated correlation coefficients: $t = 13.8$, $p = 2.2 \times 10^{-43}$). Locking the α_1 parameter also greatly reduced how well the model explained the recorded variance in NMIs, reducing the variance explained from 89 to 34% ($t = 12.0$, $p = 5.8 \times 10^{-33}$). Locking the other parameters also reduced the variance explained, though to a lesser extent (L_N : 66%, $t = 8.8$, $p = 9.5 \times 10^{-19}$; L_P : 74%, $t = 5.9$, $p = 4.1 \times 10^{-9}$; σ : 81%, $t = 3.2$, $p = 1.4 \times 10^{-3}$).

The model can describe spatial attention effects with the addition of a β parameter to represent spatially selective attention-related modulation (also previously described by Ni and Maunsell, 2017):

$$R_{P^{\text{Spat}},N} = \frac{\beta c_P L_P}{\beta c_P + c_N \alpha_2 + \sigma} + \frac{c_N L_N}{\beta c_P \alpha_1 + c_N + \sigma}. \quad (5)$$

Equation 5 has six free parameters: the five free parameters fit with Equation 4 (L_P , L_N , α_1 , α_2 , and σ) and the additional β parameter. It describes neuronal responses with spatial attention directed to the preferred stimulus, with the β parameter capturing the multiplicative effect of spatial attention (Ghose, 2009) on both the excitatory and the suppressive drives of the preferred stimulus. When attention is instead directed to the null stimulus, β modulates the excitatory and suppressive drives of the null stimulus.

To describe spatial attention effects in addition to normalization effects, we fit each unit-configuration's responses to 12 stimulus conditions: the eight stimulus conditions from the normalization task variant, plus four stimulus conditions collected during the spatial attention task variant (P^{Spat} , N^{Spat} , $P^{\text{Spat}}N$, PN^{Spat}). The model described by Equation 5 used the same six free parameter estimates for all 12 stimulus conditions per unit-configuration, with the exception of β set to equal 1 for the eight stimulus conditions collected during the normalization task variant (such that the normalization responses were essentially fit to Equation 4).

The model accounted for an average of 96% of the data variance in the observed neuronal responses to the 12 stimulus conditions, across all unit-configurations (Ni and Maunsell, 2017). Additionally, the model accounted for 85% of the variance in SMIs across all unit-configurations (and 85% of the variance in NMIs), based on the correlation between SMIs calculated using the modeled responses for each unit-configuration and SMIs calculated using the observed neuronal responses ($r = 0.92$, $p = 3.2 \times 10^{-47}$). Locking the α_2 and α_1 parameters had the most detrimental effects on how well the model explained the recorded variance in SMIs across the population, reducing the variance explained from 85 to 27% and 35%, respectively (α_2 : $t = 10.8$, $p = 2.8 \times 10^{-27}$; α_1 : $t = 9.0$, $p = 1.8 \times 10^{-19}$). Locking the β parameter also greatly reduced the variance explained, to 49% ($t = 7.4$, $p = 1.3 \times 10^{-13}$). Locking the other parameters also reduced the variance explained, to a lesser extent (L_P : 61%, $t = 5.4$, $p = 5.9 \times 10^{-8}$; L_N : 73%, $t = 3.2$, $p = 1.1 \times 10^{-3}$; σ : 81%, $t = 1.6$, $p = 0.11$).

An outstanding question from our prior paper (Ni and Maunsell, 2017) is whether the EMS-spatially tuned normalization model applies to the effects of feature attention. Feature attention has global effects, rather than the spatially specific effects of spatial attention (Saenz et al., 2002; Serences and Boynton, 2007). Whereas spatial attention effects are localized and thus modulate the spatially localized suppressive drive associated with tuned normalization (Ruff et al., 2016; Verhoef and Maunsell, 2016; Ni and Maunsell, 2017), feature attention to a particular direction should modulate localized suppressive drives across the whole visual field, eliminating differential normalization tuning between or within neurons. That is, the suppression arising from each stimulus will not vary with the attended feature. Thus, the suppressive terms associated with the preferred and null stimuli do not have a preferred direction (no L term), so attending to one direction or the other should not modulate the amount of suppression associated with either stimulus. We represent this hypothesis by allowing attention, β , to differentially affect only the excitation associated with the attended feature, modeled with Equation 6:

$$R_{P^{\text{Feat}},N} = \frac{\beta c_P L_P}{c_P + c_N \alpha_2 + \sigma} + \frac{c_N L_N}{c_P \alpha_1 + c_N + \sigma}. \quad (6)$$

Equation 6 has six free parameters (the same free parameters as Eq. 5, but with the β parameter only allowed to modulate the excitatory drive associated with the attended feature, and not the suppressive drive associated with that feature).

To test the hypothesis that spatial and feature attention differ only in how top-down attention signals interact with normalization, we fit each unit-configuration's responses to 16 stimulus conditions: the eight stimulus conditions from the normalization task variant, the four stimulus conditions collected during the spatial attention task variant (P^{Spat} , N^{Spat} , $P^{\text{Spat}}N$, PN^{Spat}), plus four stimulus conditions collected during the feature attention task variant (P^{Feat} , N^{Feat} , $P^{\text{Feat}}N$, PN^{Feat}). The model described by Equation 5 used the same six free parameter estimates for all 16 stimulus conditions per unit-configuration, with the exception of the β input in the numerator and the β inputs in the denominator set to equal 1 for the eight stimulus conditions collected during the normalization task variant (such that the normalization responses were essentially fit to Equation 4), and the β inputs in the denominator set to equal 1 for the four stimulus conditions collected during the feature attention task variant (such that the feature attention responses were essentially fit to Equation 6).

The model accounted for an average of 90% of the data variance in the observed neuronal responses to the 16 stimulus conditions, across all unit-configurations. The model also accounted for 45% of the variance in FMIs across all unit-configurations (and 77 and 85% of the variance in NMIs and SMIs, respectively), based on the correlation between FMIs calculated using the modeled responses for each unit-configuration and FMIs calculated using the observed neuronal responses ($r = 0.67$, $p = 2.0 \times 10^{-16}$). Figure 5A illustrates the model fits for the responses of a representative unit-configuration to each of the 16 stimulus conditions, based on this single model explanation of normalization, spatial attention, and feature attention effects. The model explained 95% of the variance across stimulus conditions for this example of a unit-configuration with large normalization (NMI: 0.23), spatial attention (SMI: 0.26), and feature attention modulations (FMI: 0.14). Figure 5B illustrates the model fits for a unit-configuration with small normalization (NMI: 0.06), spatial attention (SMI: 0.06), and feature attention modulations (FMI: 0.04; 97% explained variance across all stimulus conditions). Finally, Figure 5C illustrates the model fits for a unit-configuration with large normalization (NMI: 0.24) and spatial attention modulation indices (SMI: 0.25), but no apparent modulation due to feature attention (FMI: 0.00; 92% explained variance across all stimulus conditions).

Whereas locking the α_2 and α_1 parameters had detrimental effects on how well the model described NMIs and SMIs, this was not the case for FMIs. Instead, the parameter that most influenced how well the model explained the variance in FMIs across the population was β . Locking the β parameter reduced the variance explained from 45 to 2% ($t = 5.9$, $p = 3.4 \times 10^{-9}$). Locking the other parameters reduced the variance explained to a lesser extent (L_P : 21%, $t = 2.9$, $p = 3.2 \times 10^{-3}$; L_N : 37%, $t = 1.5$, $p = 0.14$; α_1 : 40%, $t = 0.80$, $p = 0.42$; α_2 : 41%, $t = 0.61$, $p = 0.54$; σ : 42%, $t = 0.54$, $p = 0.59$). In summary, whereas feature attention modulations were largely dependent on the β parameter, spatial attention modulations were dependent on both the β and α parameters because spatial attention interacts with a spatially tuned sensory normalization mechanism in a distinct manner.

The model supports the results from the data that demonstrate that normalization strength cannot explain the correlation between spatial and feature attention modulation strengths. In the model, a single β parameter per unit-configuration can inter-

Figure 5. Model fits to all 16 stimulus conditions for three example unit-configurations. **A**, The model fits (gray bars) for the responses of a representative unit-configuration to 16 stimulus conditions compared against the observed responses (black bars) of the same unit-configuration to the same stimulus conditions. The 16 stimulus conditions are illustrated in the order that follows. The first eight conditions were collected during the normalization task variant, with attention directed outside of the RF: N50 (null stimulus at 50% contrast), P50 (preferred stimulus at 50% contrast), P50N100 (preferred stimulus at 50% contrast with null stimulus at 100% contrast), P100N50, P50N50, N100, P100, P100N100. The next two conditions were collected during the spatial attention task variant (the stimulus receiving spatial attention input is illustrated with an "S" in the x-axis labels): attend preferred (P100^SN100), attend null (P100N100^S). The next two conditions were collected during the feature attention task variant (the stimulus receiving feature attention input is illustrated with an "F" in the x-axis labels): P100^FN100, P100N100^F. The last four conditions were collected during the spatial and feature attention task variants: P100^S, N100^S, P100^F, N100^F. The model explained 95% of the variance across stimulus conditions for this example unit-configuration with large normalization (NMI: 0.23), spatial attention (SMI: 0.26), and feature attention modulations (FMI: 0.14). **B**, The model explained 97% of the variance for this example unit-configuration with small normalization (NMI: 0.06), spatial attention (SMI: 0.06), and feature attention modulations (FMI: 0.04). **C**, The model explained 92% of the variance for this example unit-configuration with large normalization (NMI: 0.24) and spatial attention modulation indices (SMI: 0.25), but no apparent modulation due to feature attention (FMI: 0.00).

act with the normalization strength of that unit-configuration in the case of spatial attention, but not in the case of feature attention. In the case of spatial attention (Eq. 5), the β parameter can interact with the spatially tuned α parameters in the denominator. But in the case of feature attention (Eq. 6), that same β parameter cannot interact with the spatially tuned α parameters due to the spatially global effects of feature attention. In the case of feature attention, the β parameter is only present in the numerator. Thus, the model demonstrates how spatial and feature attention

can share common or similar top-down attention signals (represented by β), while still having differing interactions with a spatially-tuned normalization mechanism (represented by α).

Our prior study of attention modulation variance within neurons (Ni and Maunsell, 2017) rejected both our previously published stimulus-tuned normalization model (Ni et al., 2012) and an EMS stimulus-tuned normalization model, in favor of the EMS-spatially tuned normalization model tested here. However, that study only analyzed data collected during normalization and spatial attention variants of the behavioral tasks. An outstanding question from that paper is whether stimulus-tuned normalization would better describe data collected in the context of a feature attention behavioral task.

Our feature attention data allow us to address that question for the first time. Unlike the EMS-spatially tuned normalization model, in which each α term is associated with the suppression at a particular location in space, in the EMS stimulus-tuned normalization model, each α term is associated with the suppression from a particular stimulus (in this case, the preferred or the null stimulus), regardless of the location in which it appears:

$$R_{p^{Feat},N} = \frac{\beta c_p L_p}{c_p + c_N \alpha_N + \sigma} + \frac{c_N L_N}{c_p \alpha_p + c_N + \sigma}, \quad (7A)$$

$$R_{N,p^{Feat}} = \frac{c_N L_N}{c_N + c_p \alpha_p + \sigma} + \frac{\beta c_p L_p}{c_N \alpha_N + c_p + \sigma}. \quad (7B)$$

The EMS stimulus-tuned model uses six free parameters, just like the spatially-tuned model (Equation 6). In this model, Equation 7A describes data collected in Configuration A and Equation 7B describes data collected in Configuration B. Unlike EMS-spatially tuned normalization, in which α_1 always modulates suppression at the bottom location and α_2 always modulates suppression at the top location, EMS stimulus-tuned normalization has one α parameter (α_p) always modulating suppression from the preferred stimulus and one α parameter (α_N) always modulating suppression from the null stimulus.

We fit each unit-configuration's responses to 12 stimulus conditions only (because we already know that stimulus-tuned normalization models do not adequately describe spatial attention data; Ni and Maunsell, 2017): the eight stimulus conditions from the normalization task variant and the four stimulus conditions collected during the feature attention task variant (P^{Feat} , N^{Feat} , $P^{Feat}N$, PN^{Feat}). The model described by Equations 7A and 7B used the same six free parameter estimates for all 12 stimulus conditions per unit-configuration, with the exception of the β input set to equal 1 for the eight stimulus conditions collected during the normalization task variant.

The EMS stimulus-tuned normalization model only accounted for an average of 78% of the data variance across all tested stimulus configurations. This was significantly less than the variance explained by the EMS spatially-tuned model (F statistic = 3.5, $p = 0.024$). Additionally, the model only accounted for 21% of the variance in FMIs across unit-configurations. This was significantly less than the FMI variance explained by the spatially tuned model ($t = 3.5$, $p = 5.6 \times 10^{-4}$). In summary, the EMS spatially-tuned model does a superior job of describing neuronal responses modulated by feature attention as well, and suggests that spatial and feature attention differ only in how their top-down signals interact with a spatially tuned sensory normalization mechanism.

Discussion

Although all forms of top-down attention involve the feedback of attention-related signals from higher cortical areas (Knudsen, 2007; Baluch and Itti, 2011) that modulate the activity of specific cortical sensory neurons (Maunsell and Treue, 2006), it has been unclear whether two well studied forms of top-down attention, spatial attention and feature attention, engage common or distinct mechanisms in modulating the responses of sensory neurons (Kastner and Ungerleider, 2000; Maunsell, 2015). Although similarities in the ways that spatial and feature attention modulate sensory neuron firing rates and correlations between pairs of neighboring sensory neurons (Treue and Martínez-Trujillo, 1999; Martínez-Trujillo and Treue, 2004; Patzwahl and Treue, 2009; Cohen and Maunsell, 2011) support the feature similarity gain model of a common neuronal mechanism (Treue and Martínez-Trujillo, 1999; Maunsell and Treue, 2006), trial-to-trial neuronal response fluctuations due to spatial and feature attention have been found to be uncorrelated when measured within a single task (Cohen and Maunsell, 2009). Here we described a model that accounts for both the similarities and distinctions between the effects of spatial and feature attention on sensory neuron activity. These two forms of attention may share common top-down attention signals, or top-down inputs with similar effects on sensory neurons; however, these attention signals have differing interactions with a spatially tuned normalization mechanism, leading to differing effects on sensory neuron responses.

A previously described EMS-spatially tuned normalization model (Ni and Maunsell, 2017) accounts well for these results. This model captures the spatial localization of normalization strength observed in multiple sensory areas, including V1, V4, and MT (Ruff et al., 2016; Verhoef and Maunsell, 2016). The spatial localization of the suppressive drive of normalization was conceptualized in earlier normalization studies as spatially localized pools of neurons (DeAngelis et al., 1992; Carandini et al., 1997), with neurons in a normalization pool having overlapping RFs and the collective activity of a pool driving the suppressive effects of normalization. This spatial tuning of the normalization mechanism may be the key to differentiating the neuronal modulation effects of spatial versus feature attention reported here in area MT: whereas the localized spatial extent of spatial attention effects matches the localized spatial extent of normalization in MT neurons, the global spatial extent of feature attention effects does not match the spatial extent of normalization. Thus, only spatial attention can differentially modulate tuned normalization for each of the stimuli presented in the RF.

Although the difference in the spatial extent of the neural effects of spatial (Treue and Maunsell, 1996; Brefczynski and DeYoe, 1999) versus feature attention (Treue and Martínez-Trujillo, 1999; McAdams and Maunsell, 2000; Saenz et al., 2002; Serences and Boynton, 2007) has long been regarded as a potentially significant distinction between these two forms of attention, the consequences of this distinction have not been clarified. Our results demonstrate that interactions between attention and a spatially-tuned normalization mechanism depend critically on this distinction. They extend previous studies that used modeling (Reynolds and Heeger, 2009), psychophysics, and fMRI (Herrmann et al., 2010, 2012) to show that the visual field extent of spatial attention can greatly affect the way that human contrast response functions are altered by attention.

Earlier work by Rust et al. (2006) identified a range of tuned normalization in MT, and showed that it correlated with greater pattern versus direction selectivity. However, they assigned their

tuned normalization to V1 inputs and found it was strongest for inputs sharing the same direction preference as each MT neuron, whereas the tuned normalization we describe appears to be more selective for RF location (Ni and Maunsell, 2017). It is possible that normalization differs between brain areas, and that other brain areas do not share the spatial localization of normalization strength reported in areas V1, V4, and MT. In fact, of those three brain areas, MT could have the strongest normalization strength localization (Ruff et al., 2016). Spatial and feature attention may have different relationships with normalization in other areas, changing the relationship between the effects of these two forms of attention.

The results of this study reinforce the idea that the normalization mechanism plays an important role in the effect of attention on neuronal activity. Indeed, normalization provides a mechanism for the “biased competition” associated with attention (Desimone and Duncan, 1995; Maunsell, 2015). At the same time, our results underscore that attention and normalization are distinct. Prior electrophysiological studies of the relationship between attention and normalization focused mainly on the relationship between spatial attention and normalization (Lee and Maunsell, 2009; Ni et al., 2012; Verhoef and Maunsell, 2016; Ni and Maunsell, 2017; Ruff and Cohen, 2017). In those studies, the strong relationship between spatial attention and normalization made distinguishing the contributions of these two phenomena to the recorded neuronal modulations difficult. Here, the tests of the relationship between feature attention and normalization make the distinction between top-down attentional signals and sensory-level mechanisms clear: the magnitude of attention effects does not need to be correlated with normalization strengths.

Normalization has been linked not only to modulations of single-neuron firing rates (Carandini and Heeger, 2011) but also to modulations of the correlated variability between the firing rates of pairs of sensory neurons in response to repeated presentations of the same stimulus (for review, see Schmitz and Duncan, 2018). Growing evidence suggests that normalization plays an important role in the relationship between changes in correlated variability and behavioral improvements due to attention (Ruff and Cohen, 2016; Ruff et al., 2016; Verhoef and Maunsell, 2017). Although this evidence is mainly based on studies of spatial attention, studies have demonstrated that feature attention is also associated with changes in correlated variability (Cohen and Newsome, 2008; Cohen and Maunsell, 2011). Thus, although we demonstrated here that spatial and feature attention effects on single-neuron responses differ due to normalization, in the future it will be important to clarify and possibly distinguish how spatial and feature attention effects on correlated variability are affected by normalization.

Our model does not explain why spatial attention effects on neuronal responses in MT should be greater in magnitude than feature attention effects. It is possible that spatial attention is innately more potent than feature attention, perhaps because it is central to orienting responses and can involve multiple sensory modalities. Alternatively, the disparity might be due to a higher level of difficulty in the spatial attention task, which involved directing attention to one of the two stimuli in close proximity. Prior studies have demonstrated that task difficulty can affect firing rates (Boudreau et al., 2006; Chen et al., 2008). Although we used an adaptive staircase procedure that adjusted the difficulty of the change detection for each task variant individually, the spatial attention variant may have evoked more effort for the same level of performance than the others, and produced greater modulations for that reason. However, even if difficulty differed between task variants, that would not have affected the qualitative

findings of this study. Task difficulty has been shown to primarily affect the magnitude of modulations associated with spatial attention (Boudreau et al., 2006; Chen et al., 2008) and is likely to have an equivalent effect on modulations associated with feature attention. The analyses that matched modulation strength between spatial and feature attention demonstrate that there are significant differences between these two types of attention that are robust to the absolute strength of modulation.

Though the goal of this study was to devise an experimental paradigm that allowed independent measurements of spatial and feature attention, true independence is difficult to achieve. During the feature attention task variant, attention was always directed to a particular location in space. Although that location was as far away from the RF as possible, and feature attention modulation indices were calculated between two stimulus conditions in which the attended location never changed, spatial attention may have still affected the feature attention modulations. Similarly, during the spatial attention task variant, the monkeys were tasked with responding to a change in the stimulus features. Although the task was designed to prevent feature attention from varying systematically across different stimulus conditions, the monkey still had to actively ignore changes in direction between the preferred and null directions.

However, the behavioral paradigm allowed measurements of attention that were independent enough to suggest that spatial and feature do in fact differ in how they interact with normalization. In conclusion, although many previous studies of the relationship between attention and normalization have focused on spatial attention, understanding the relationship between feature attention and normalization may clarify not only our understanding of the relationship between spatial and feature attention effects on neuronal responses, but also of the relationship between top-down cognitive processes and sensory processes. Studying feature attention may similarly change our understanding of the relationship between normalization and correlated variability, and the role that normalization plays in processes that affect the amount of information available in the activity of a population of neurons in general.

References

- Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *J Neurophysiol* 52:1106–1130.
- Baluch F, Itti L (2011) Mechanisms of top-down attention. *Trends Neurosci* 34:210–224.
- Bisley JW (2011) The neural basis of visual attention. *J Physiol* 589:49–57.
- Bonds AB (1989) Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis Neurosci* 2:41–55.
- Boudreau CE, Williford TH, Maunsell JHR (2006) Effects of task difficulty and target likelihood in area V4 of macaque monkeys. *J Neurophysiol* 96:2377–2387.
- Boynton GM (2009) A framework for describing the effects of attention on visual responses. *Vision Res* 49:1129–1143.
- Brefczynski JA, DeYoe EA (1999) A physiological correlate of the “spotlight” of visual attention. *Nat Neurosci* 2:370–374.
- Britten KH, Heuer HW (1999) Spatial summation in the receptive fields of MT neurons. *J Neurosci* 19:5074–5084.
- Carandini M, Heeger DJ (2011) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51–62.
- Carandini M, Heeger DJ, Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17:8621–8644.
- Carrasco M (2011) Visual attention: the past 25 years. *Vision Res* 51:1484–1525.
- Chen Y, Martinez-Conde S, Macknik SL, Bereshpolova Y, Swadlow HA, Alonso JM (2008) Task difficulty modulates the activity of specific neuronal populations in primary visual cortex. *Nat Neurosci* 11:974–982.

- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378.
- Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600.
- Cohen MR, Maunsell JHR (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70:1192–1204.
- Cohen MR, Newsome WT (2008) Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60:162–173.
- Crist CF, Yamasaki DS, Komatsu H, Wurtz RH (1988) A grid system and a microsyringe for single cell recording. *J Neurosci Methods* 26:117–122.
- Cui Y, Liu LD, Khawaja FA, Pack CC, Butts DA (2013) Diverse suppressive influences in area MT and selectivity to complex motion features. *J Neurosci* 33:16715–16728.
- DeAngelis GC, Robson JG, Ohzawa I, Freeman RD (1992) Organization of suppression in receptive fields of neurons in cat visual cortex. *J Neurophysiol* 68:144–163.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222.
- Desimone R, Ungerleider LG (1986) Multiple visual areas in the caudal superior temporal sulcus of the macaque. *J Comp Neurol* 248:164–189.
- Ghose GM (2009) Attentional modulation of visual responses by flexible input gain. *J Neurophysiol* 101:2089–2106.
- Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324:1207–1210.
- Gregoriou GG, Paneri S, Sapountzis P (2015) Oscillatory synchrony as a mechanism of attentional processing. *Brain Res* 1626:165–182.
- Hayden BY, Gallant JL (2005) Time course of attention reveals different mechanisms for spatial and feature-based attention in area V4. *Neuron* 47:637–643.
- Heeger DJ (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9:181–197.
- Heeger DJ, Simoncelli EP, Movshon JA (1996) Computational models of cortical visual processing. *Proc Natl Acad Sci U S A* 93:623–627.
- Herrmann K, Montaser-Kouhsari L, Carrasco M, Heeger DJ (2010) When size matters: attention affects performance by contrast or response gain. *Nat Neurosci* 13:1554–1559.
- Herrmann K, Heeger DJ, Carrasco M (2012) Feature-based attention enhances performance by increasing response gain. *Vision Res* 74:10–20.
- Heuer HW, Britten KH (2002) Contrast dependence of response normalization in area MT of the rhesus macaque. *J Neurophysiol* 88:3398–3408.
- Howell DC (2007) Statistical methods for psychology. Belmont, CA: Thomson Wadsworth.
- Ibos G, Freedman DJ (2016) Interaction between spatial and feature attention in posterior parietal cortex. *Neuron* 91:931–943.
- Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci* 23:315–341.
- Knudsen EI (2007) Fundamental components of attention. *Annu Rev Neurosci* 30:57–78.
- Lee J, Maunsell JHR (2009) A normalization model of attentional modulation of single unit responses. *PLoS One* 4:e4651.
- Lee J, Maunsell JHR (2010) Attentional modulation of MT neurons with single or multiple stimuli in their receptive fields. *J Neurosci* 30:3058–3066.
- Livingstone MS, Pack CC, Born RT (2001) Two-dimensional substructure of MT receptive fields. *Neuron* 30:781–793.
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751.
- Maunsell JHR (2015) Neuronal mechanisms of visual attention. *Annu Rev Vis Sci* 1:373–391.
- Maunsell JHR, Treue S (2006) Feature-based attention in visual cortex. *Trends Neurosci* 29:317–322.
- Maunsell JHR, van Essen DC (1983) The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *J Neurosci* 3:2563–2586.
- McAdams CJ, Maunsell JHR (2000) Attention to both space and feature modulates neuronal responses in macaque area V4. *J Neurophysiol* 83:1751–1755.
- Melcher D, Papathomas TV, Vidnyánszky Z (2005) Implicit attentional selection of bound visual features. *Neuron* 46:723–729.
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888.
- Moore T, Armstrong KM (2003) Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421:370–373.
- Morrone MC, Burr DC, Maffei L (1982) Functional implications of cross-orientation inhibition of cortical visual cells: I. Neurophysiological evidence. *Proc R Soc Lond B Biol Sci* 216:335–354.
- Ni AM, Maunsell JHR (2017) Spatially tuned normalization explains attention modulation variance within neurons. *J Neurophysiol* 118:1903–1913.
- Ni AM, Ray S, Maunsell JHR (2012) Tuned normalization explains the size of attention modulations. *Neuron* 73:803–813.
- Paneri S, Gregoriou GG (2017) Top-down control of visual attention by the prefrontal cortex: functional specialization and long-range interactions. *Front Neurosci* 11:545.
- Patzwahl DR, Treue S (2009) Combining spatial and feature-based attention within the receptive field of MT neurons. *Vision Res* 49:1188–1193.
- Posner MI (1980) Orienting of attention. *Q J Exp Psychol* 32:3–25.
- Reynolds JH, Heeger DJ (2009) The normalization model of attention. *Neuron* 61:168–185.
- Richert M, Albright TD, Krekelberg B (2013) The complex structure of receptive fields in the middle temporal area. *Front Syst Neurosci* 7:2.
- Ruff DA, Cohen MR (2016) Stimulus dependence of correlated variability across cortical areas. *J Neurosci* 36:7546–7556.
- Ruff DA, Cohen MR (2017) A normalization model suggests that attention changes the weighting of inputs between visual areas. *Proc Natl Acad Sci U S A* 114:E4085–E4094.
- Ruff DA, Alberts JJ, Cohen MR (2016) Relating normalization to neuronal populations across cortical areas. *J Neurophysiol* 116:1375–1386.
- Rust NC, Mante V, Simoncelli EP, Movshon JA (2006) How MT cells analyze the motion of visual patterns. *Nat Neurosci* 9:1421–1431.
- Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nat Neurosci* 5:631–632.
- Schmitz TW, Duncan J (2018) Normalization and the cholinergic microcircuit: a unified basis for attention. *Trends Cogn Sci* 22:422–437.
- Serences JT, Boynton GM (2007) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55:301–312.
- Tanaka K, Hikosaka K, Saito H, Yukie M, Fukada Y, Iwai E (1986) Analysis of local and wide-field movements in the superior temporal visual areas of the macaque monkey. *J Neurosci* 6:134–144.
- Treue S, Martínez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579.
- Treue S, Maunsell JHR (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539–541.
- Verhoef BE, Maunsell JHR (2016) Attention operates uniformly throughout the classical receptive field and the surround. *eLife* 5:e17256.
- Verhoef BE, Maunsell JHR (2017) Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat Neurosci* 20:969–977.
- Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. *Percept Psychophys* 33:113–120.
- Williford T, Maunsell JHR (2006) Effects of spatial attention on contrast response functions in macaque area V4. *J Neurophysiol* 96:40–54.
- 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.
- Zhou H, Desimone R (2011) Feature-based attention in the frontal eye field and area V4 during visual search. *Neuron* 70:1205–1217.