EJN European Journal of Neuroscience

FENS Federation of European Neuroscience Societies

European Journal of Neuroscience, Vol. 46, pp. 2844–2858, 2017 doi:10.1111/ejn.13764

COGNITIVE NEUROSCIENCE

Distracter suppression dominates attentional modulation of responses to multiple stimuli inside the receptive fields of middle temporal neurons

Nour Malek,[1] Stefan Treue,[2,3] Paul Khayat[4,5] and Julio Martinez-Trujillo[1,6]

[1]Department of Physiology, McGill University, Montreal, QC, Canada
[2]Cognitive Neuroscience Laboratory, German Primate Centre – Leibniz Institute for Primate Research, Göttingen, Germany
[3]Faculty for Biology and Psychology, University of Göttingen, Göttingen, Germany
[4]McGill Vision Research Unit, Department of Ophthalmology, McGill University, Montreal, QC, Canada
[5]École d'Optométrie, Université de Montréal, Montreal, QC, Canada
[6]Departments of Psychiatry, Physiology and Pharmacology, Robarts Research Institute, University of Western Ontario, London, ON, N6A 5B7, Canada

Keywords: attention, distracter suppression, filter model, macaque monkey, middle temporal

Abstract

Single-cell studies in macaques have shown that attending to one of two stimuli, positioned inside a visual neuron's receptive field (RF), modulates the neuron's response to reflect the features of the attended stimulus. Such a modulation has been described as a 'push–pull' effect relative to a reference response: a neuron's response increases when attention is directed to a preferred stimulus, and decreases when attention is directed to a non-preferred stimulus. It has been further suggested that the response increase when attending to a preferred stimulus is the predominant effect. Here, we show that the observed attentional modulation depends on the reference response. We recorded neuronal responses in motion processing area middle temporal (MT) of macaques to two moving random dot patterns positioned inside neurons' RF. One pattern always moved in the neuron's antipreferred direction (null pattern), while the other moved in one of 12 directions (tuning pattern). At the beginning of a trial, a cue indicated the location and direction of the target. The animal was required to release a lever when a change in the target direction occurred, and to ignore changes in the distracter. Relative to neurons' initial responses to the dual stimuli (when attention was less likely to modulate responses), attending to the tuning pattern did not significantly modulate responses over time. However, attending to the null pattern progressively decreased responses over time. These results were quantitatively described by filter and input gain models, characterising a predominant response suppression relative to a reference response, rather than response enhancement.

Introduction

Attention modulates the responses of neurons across visual cortical areas of primates (Maunsell & Cook, 2002; Noudoost *et al.*, 2010). Previous attention investigations in macaque monkeys recorded single unit responses to two different stimuli inside the visual neuron's receptive field (RF) (Reynolds *et al.*, 1999; Treue & Martinez-Trujillo, 1999; Ghose & Maunsell, 2008). These studies reported a response increase when macaques attend to the neuron's preferred stimulus (e.g., a random dot pattern (RDP) moving in the neuron's preferred direction), and a relative response decrease when macaques attend to a less preferred stimulus (e.g., an RDP moving in the neuron's antipreferred direction). This has been described as a 'push–pull' effect, or a symmetric modulation, relative to a reference response when attention is directed outside the RF (Treue & Martinez-Trujillo, 1999). Another study, which used a similar reference response, reported a greater enhancement of responses when attending to a preferred stimulus than suppression of responses when attending to a less preferred stimulus (Ni *et al.*, 2012). We hypothesise that the effect of attentional modulation (e.g., 'push–pull' or enhancement bias) depends on the selected reference response.

To better illustrate this issue, we will consider predictions of three different models of attentional modulation: spotlight, filter and input gain (Britten & Heuer, 1999; Ghose & Maunsell, 2008). These

Correspondence: Dr J. Martinez-Trujillo, [5]Departments of Psychiatry, Physiology and Pharmacology, as above.
E-mail: julio.martinez@robarts.ca

Received 4 April 2017, revised 21 October 2017, accepted 25 October 2017

Edited by Guillaume Rousselet

Reviewed by Leonardo Chelazzi, University of Verona, Italy; and Farran Briggs, Dartmouth College, USA

The associated peer review process communications can be found in the online version of this article.

models do not describe the mechanism of the modulation (e.g., changes in response normalisation), rather they quantify the modulation of responses over time relative to a reference response. Figure 1A shows a cartoon of a hypothetical attention paradigm. The stimulus configuration consists of two RDPs positioned within a neuron's RF. One RDP is moving in the neuron's preferred direction, and the other in its antipreferred direction. A third RDP, of no specific direction, has been positioned outside the RF and in the opposite hemifield. In this paradigm, a cue is presented before dual stimuli onset to indicate the target stimulus to be attended. Moreover, target direction changes are biased to occur towards the end of a trial so that attention allocation is strengthened over time. The scenarios on the right depict predicted responses based on the three different models of attentional modulation. The green traces are the predicted responses when attending to the preferred stimulus, and the red traces are the predicted responses when attending to the antipreferred stimulus. The grey dashed line represents a common reference response when attention is allocated outside the RF (i.e., to the third RDP in the opposite hemifield), while the black dashed line illustrates an alternative reference response during initial dual stimuli presentation before significant attentional modulation effects.

Relative to the reference response during initial dual stimuli presentation (black dashed line), the input gain model predicts that attending to the preferred pattern increases a neuron's responses over time (green solid line), and attending to the antipreferred pattern decreases a neuron's responses over time (red solid line) (Fig. 1A, Scenario 1). This is the symmetric 'push–pull' effect alluded to earlier. The spotlight model predicts an increase in responses over time when the preferred pattern is attended (green solid line), but no modulation when the antipreferred pattern is attended (red solid line) (Fig. 1A, Scenario 2). The filter model predicts that attending to the preferred stimulus does not substantially change a neuron's response (green solid line), while attending to the antipreferred stimulus suppresses a neuron's response (red solid line) (Fig. 1A, Scenario 3). Critically, these predictions would change if a different reference response was used.

A common reference response that has been used to assess attentional modulation is a neuron's response when attending to a pattern

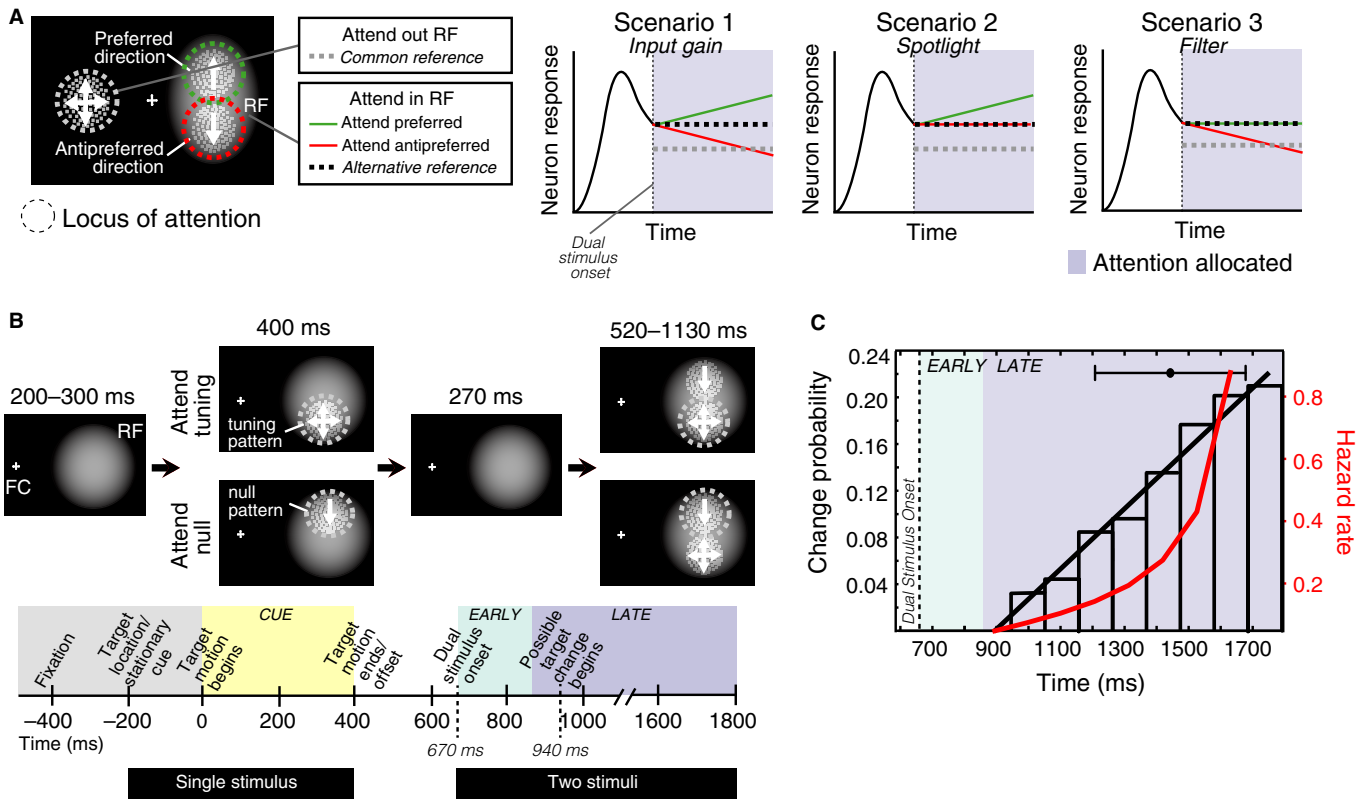

FIG. 1. Experiment 1 task paradigm. (A) A hypothetical attention paradigm and three scenarios exemplifying neuronal response predictions based on proposed models of attentional modulation. The black rectangle represents the computer screen, and the central cross the fixation cross (FC). The white blobs illustrate the neuron's receptive field (RF). Dashed circles demarcate the locus of attention, and the arrows indicate the motion direction of the dots in the random dot pattern (RDP). The scenarios portray schematics of modulation during attention (i.e., time shaded in muted grey) when attending to the preferred direction (RDP moving up in the panel; green line) and when attending to the antipreferred direction (RDP moving down in the panel; red line). The modulation is sketched in relation to two possible reference responses: (i) neuronal response when attention is allocated outside of the RF (grey dashed line) and (ii) response to the onset of dual stimuli in the RF during early attention allocation (black dashed line). For purposes of illustration, we assume that increases and decreases in response follow a linear trend; however, this does not need to be the case. (B) Experiment 1 task details. The timeline provides a representation of when trial events occurred and how each trial was broken down for data analysis. The numbers above each stimulus illustration represent durations. (C) Distributions of times at which the target changed motion direction (event time) normalised to the total number of hit trials to exhibit its probability density function (i.e., change probability); the bar above the histogram shows the mean time from the start of trials at which target direction changes occurred and its associated standard deviation. A best-fit line is shown in black. In red is the exponential hazard rate (Eqn 1) to demonstrate the rate of an event change occurring at a specific time point with no event change taking place beforehand. Note, the hazard rate at the last time bin is infinity and hence not drawn in.

outside the neuron's RF and in the opposite hemifield (grey dashed line). Previous studies have shown that a neuron's response decreases when directing attention outside the RF as compared to when attention is allocated inside the RF (Luck *et al.*, 1997; McAdams & Maunsell, 1999, 2000; Seidemann & Newsome, 1999; Treue & Martinez-Trujillo, 1999; Niebergall *et al.*, 2011). As can be seen in Fig. 1A, this reference response would predict an increased response for all models of attentional modulation when attending to the preferred direction (i.e., the green solid line is above the grey dashed line in all three scenarios). When attending to the antipreferred direction, the input gain and filter model—which are meant to predict a response suppression—show either a small enhancement or no modulation, and the spotlight model predicts a response enhancement. Essentially, this reference response would suggest a predominant enhancement bias, and the spotlight model or a biased input gain model would better describe the data. Extrapolating beyond this hypothetical paradigm, one could easily recognise that using a different reference response would again produce different results.

The goal of our study was to quantify attentional modulation in the visual motion processing area middle temporal (MT) of macaque monkeys. We investigated attentional modulation relative to neurons' responses within the RF before attention is strongly allocated (Fig. 1A, black dashed line), rather than when attention is allocated outside the RF (grey dashed line). Moreover, we characterised the time course of this attentional modulation. We found that the filter and input gain models could account for attentional response modulations in MT, but the spotlight model could not. The predominant effect of attentional modulation was a progressive decrease in single neuron responses over time when the animals attended to the antipreferred stimulus. This stronger tendency to suppress distracters also correlated with behavioural measurements of reaction time.

Materials and methods

Experiment 1

The responses of 77 direction-selective neurons were recorded in area MT of two behaving male monkeys (*Macaca mulatta*) when two RDPs were positioned inside their RFs (Treue & Martinez-Trujillo, 1999; Martinez-Trujillo & Treue, 2002). All procedures abided by local and national regulations and were approved by the appropriate regional government office (Regierungspraesidium Tuebingen, Germany).

Stimuli

Random dot patterns were made of small bright dots at a density of 5 dots per degree^2 and average luminance of 55 cd/m^2. RDPs were plotted within a stationary, virtual and circular aperture on a computer monitor (viewing distance 57 cm). The diameter of the aperture varied depending on the size of the recorded neuron's RF. Prior to an experimental session, the RF size, preferred speed and preferred direction of the recorded neuron were estimated (Treue & Martinez-Trujillo, 1999). For every trial, two RDPs of equal size were presented at two separate positions within the RF. The size of the patterns was adjusted so that they both fell within the recorded neuron's classical RF. The two positions were chosen in such a manner that different directions of a single RDP elicited similar responses when presented at either position and that both positions' centres were approximately equidistant from the fixation cross. One of the two RDPs (the null pattern) always moved in the neuron's antipreferred direction (180° from the preferred direction), while the second (the tuning pattern) could move in one of 12 directions, spaced every 30° from the neuron's preferred direction.

Behavioural task

The monkeys were trained to attend to one of two moving RDPs (the target), while ignoring the presence of the other RDP (the distracter) and maintaining their gaze on a stationary cross for the entire trial duration (Fig. 1B); 200 ms after the animal foveated the cross, a stationary RDP (the cue) was presented—marking the location of the upcoming target. Once the monkey pressed a lever, the cue moved in the target direction for 400 ms (cue epoch) and then disappeared. After an interval of 270 ms, with only the fixation cross remaining on the screen, two moving RDPs appeared: one (the target) at the same location and moving in the same direction as the cue and the other (the distracter) at the other location. A direction change (15°–25° for 110 ms), in either target or distracter, occurred 270–1130 ms after target and distracter onset. The monkey was required to release the lever only in response to a direction change in the target within a reaction time window of 250–700 ms after the direction change. The reaction time window and the magnitude of the detected direction change were chosen based on animals' performances during training sessions to maximise the number of hit trials. The reaction time window depended on the magnitude of the direction change and proximity between target and distracter stimuli. While kept constant within an experiment session, these parameters could slightly vary from session to session. The monkey was also required to ignore direction changes that occurred in the distracter (distracter trials). Half of an experimental session's trials were distracter trials. These trials also contained behaviourally relevant direction changes in the target occurring 500–600 ms after the distracter change. Only correctly completed trials were rewarded with a drop of juice and considered for further data analysis. Trials where the monkey broke fixation by deviating its gaze more than 1° from the fixation cross or responded outside the reaction time window were considered errors and were aborted without reward.

Change probability distribution and hazard rate

The large time window over which a direction change could occur in the target (270–1130 ms after target and distracter onset) was chosen to reduce monkeys' expectations of when a target change would occur. Target direction changes were more likely to occur later in a trial so that the monkey would allocate attention for longer and more strongly over time. The timings of direction changes for hit trials were binned and normalised to the total number of hit trials (41 140 trials) (Fig. 1C). The change probability distribution was computed by fitting a straight line to the top of the bars comprising this histogram; the slope parameter (m) of the linear fit for the change probability distribution was 0.00027 [R^2 = 0.99, lower confidence interval (LCI) = 0.00024, upper confidence interval (UCI) = 0.00029]. We also obtained the hazard rate function, $H(t)$, (Eqn 1) to indicate the proportion of direction changes that would occur at a specific time point given that no change had occurred prior to that time point (Luce, 1986; Ghose & Maunsell, 2002):

$$H(t) = \frac{p(t)}{1 - \int_0^t p(t)\mathrm{d}t}, \quad (1)$$

where $p(t)$ refers to the change probability, as determined earlier by the linear fit, at binned time point t, and $\int_0^t p(t)\mathrm{d}t$ is the integral from the starting time point that a direction change may occur until the binned

time point t (i.e., the cumulative distribution). The denominator expression $(1 - \int_0^t p(t)\mathrm{d}t)$ represents the survival function, or the probability that the direction change has not yet occurred by time point t. The computed hazard rate exhibits an exponential increase (red line in Fig. 1C represents the exponential fit of the hazard rate (HF(t)): HF(t) = 0.0097 × ($e^{0.0058(t)}$), R^2 = 0.99). To determine whether the hazard rate biased attention towards the end of the trial, we assessed the monkeys' performance as a function of time from trial onset.

The histogram of the timings of target direction changes was used when determining an appropriate time window for the reference response. This is discussed in more detail below.

Recordings

Chambers were implanted in both monkeys during stereotaxic surgeries on top of a craniotomy of the left hemisphere's parietal bone, providing access to MT along a vertical approach (Martinez-Trujillo & Treue, 2002). Extracellular signals were recorded using tungsten microelectrodes (diameter 150 μm, impedance 0.5–2 mΩ; Microprobe Inc. and FHC Inc.). Single units were isolated online using a window discriminator (Bak Electronics Corp., USA). Only units that demonstrated a fourfold response modulation when presented the preferred direction vs the antipreferred direction were considered direction-selective and recorded from Martinez-Trujillo & Treue (2002). These units were verified to be from MT based on RF size, eccentricity and direction selectivity (Treue & Martinez-Trujillo, 1999).

Analysis

Spike density functions for each neuron were obtained by convolving each spike with a Gaussian function (1 ms resolution, Gaussian kernel with σ = 30 ms). The single neuron responses to all 12 combinations of null and tuning patterns were then normalised to the neuron's response during the cue presentation when the stimulus moved in the neuron's preferred direction. Tuning curves were derived for both attentional conditions, when the target was either a tuning pattern (attend-tuning conditions) or the null pattern (attend-null conditions), by fitting neuronal responses as a function of the direction of the tuning pattern throughout a trial:

$$\text{Response (direction)} = \text{baseline} + \text{gain} \times \exp^{-0.5\times\frac{(\text{direction-center})^2}{(\text{width})^2}} \quad (2)$$

Response (direction) refers to the normalised neuronal response as a function of the tuning RDP's direction, baseline is the neuron's response to the antipreferred direction, gain (or amplitude) is the height of the tuning curve or difference between the baseline and maximal response, centre is the preferred direction in degrees, and width is the range of directions the neuron responds to Treue & Martinez-Trujillo (1999). To reveal the temporal dynamics of these tuning functions, curves were obtained by integrating the responses over 50 ms in steps of 10 ms. The resulting tuning curves were then aligned to the preferred direction (so that 0° represents the neuron's preferred direction and 180° the neuron's antipreferred direction).

The fit tuning curves for all presented directions in both condition types generated two 3D matrices of normalised responses for each neuron. These were plotted as heat maps to demonstrate each neuron's normalised responses as a function of time and across all stimuli direction combinations for both conditions. A heat map representing the population's activity was also generated that averaged these normalised responses across all recorded neurons.

The neuronal firing rates for each condition were computed during three time windows: cue, dual stimulus early, and dual stimulus late (Fig. 1B). The cue epoch is the 400-ms time window of the cue motion. The dual stimulus early epoch is the initial 200 ms after target and distracter onset. According to the change probability distribution, no change in the stimuli occurred during this time nor during the following 70 ms (i.e., there was no stimulus change during the 270 ms following target–distracter onset; Fig. 1C). This time window did not require a monkey's response; therefore, during this time, the monkey did not need to and most likely did not strongly allocate attention. Yet, this time window did provide the same paired stimulus presentation in the RF as was used for the remainder of the trial. As such, we utilised the responses elicited during this epoch as the reference response. Finally, the dual stimulus late epoch is the 930-ms time interval following the dual stimulus early epoch (i.e., the dual stimulus late epoch refers to the 200–1130 ms after target–distracter onset), during which a direction change occurred with an increasing change probability and hazard rate towards the end of the trial. During this epoch, animals likely increased attention as a function of time to better respond to direction changes in the target and ignore direction changes in the distracter.

Accounting for potential adaptation effects. The 400-ms moving cue may have led to motion sensory adaptation (Solomon & Kohn, 2014), potentially influencing the reference response that we intended to utilise. To control for this potential confounder, we excluded neurons that exhibited significantly different responses during the dual stimulus early epoch of the two condition types based on Wilcoxon rank-sum tests. This was based on the rationale that during the attend-tuning condition, where preferred and other excitatory stimuli are shown in the RF as the cue, sensory adaptation of the classical RF may occur. This may cause a decreased response to the two stimuli during the dual stimulus early period. On the other hand, during the attend-null condition, adaptation of the surrounding normalisation (inhibitory) component may occur, particularly in neurons where the cue moved in the antipreferred direction and led to a response decrease below baseline. This may enhance responses during the dual stimulus early epoch. Therefore, we reasoned that neurons that do not show a difference in response to the two stimuli between the two conditions during the dual stimulus early period were less likely affected by adaptation or underwent adaptation to a lesser extent. Eliminating neurons that had significantly different responses during the dual stimulus early epoch bolstered this epoch's response as a more appropriate reference when modelling spatial integration and comparing between attention conditions.

We confirmed which neurons were more strongly affected by adaptation by additionally computing spatial summation indices that compared how each neuron responded to the onset of a preferred motion direction with a null stimulus vs the onset of either stimulus alone when attending to the preferred direction (SSI_P) or the antipreferred direction (SSI_N):

$$SSI_P = \frac{\text{Preferred}_{\text{early}}}{\text{Preferred}_{\text{cue}} + \text{Null}_{\text{cue}}}, \quad (3)$$

$$SSI_N = \frac{\text{Null}_{\text{early}}}{\text{Preferred}_{\text{cue}} + \text{Null}_{\text{cue}}}, \quad (4)$$

$\text{Preferred}_{\text{early}}$ and $\text{Null}_{\text{early}}$ refer to the neuron's activity during the dual stimulus early epoch when the paired preferred and null stimuli first appear during the attend-tuning and attend-null conditions, respectively. $\text{Preferred}_{\text{cue}}$ and Null_{cue} refers to the neuron's response during the first 200 ms at the onset of the cue's motion (when the

preferred or antipreferred direction, respectively, is first presented alone). Neurons whose SSI_N exceeded '1' showed a greater transient during the dual stimulus early epoch in the attend-null relative to the attend-tuning condition; these neurons were removed from further analysis. A heat map averaging the normalised responses of the 53 included neurons was regenerated.

Modelling spatial integration. Spatial integration models, based on Britten & Heuer's (1999) power-law summation rules, were fit to the data to relate the dual stimulus early response during the attend-tuning and attend-null conditions to the response of each individual stimulus:

$$\text{Tuning}_{\text{early}} = \alpha(\text{Null}_{\text{cue}}^{n} + \text{Tuning}_{\text{cue}}^{n})^{\frac{1}{n}}, \quad (5)$$

$$\text{Null}_{\text{early}} = \alpha(\text{Null}_{\text{cue}}^{n} + \text{Tuning}_{\text{cue}}^{n})^{\frac{1}{n}}, \quad (6)$$

Observe that the distinction between Eqns (5) and (6) is that the 200-ms dual stimulus early response could come from two sources depending on the attention condition type—$\text{Tuning}_{\text{early}}$ and $\text{Null}_{\text{early}}$. $\text{Tuning}_{\text{cue}}$ and Null_{cue} refer to the single stimulus response during the initial 200 ms of cue motion. Tuning takes into account the responses not only when the preferred pattern was presented with the null, but also for stimuli moving 30° from preferred. More importantly, the scale factor, α, and exponent, n, were free to vary. This generalised scaled power model is known to provide a reasonable fit to similar data as ours (Ghose & Maunsell, 2008). The distributions of neurons' fitted α and n were then assessed to determine which model of spatial integration best described the recorded response of MT neurons [e.g., winner-take-all (large n, $\alpha = 1$), averaging ($n = 1$, $\alpha = 0.5$), or the normalisation model proposed by Simoncelli & Heeger (1998) ($n = 0.5$, varying α)].

Once the α and n were determined, additional β parameters were also fit to model attentional modulation in three different ways (Ghose & Maunsell, 2008): (i) limiting the attentional effect to an enhancement of responses to the attended stimulus (the spotlight model; Eqns 7 and 8), (ii) limiting the attention effect to a suppression of responses to the unattended stimulus (the filter model; Eqns 9 and 10) or (iii) allowing attention to both enhance responses to the attended stimulus and suppress responses to the unattended stimulus (the input gain model; Eqns 11 and 12).

The spotlight model:

$$\text{Tuning}_{\text{late}}(i) = \alpha(\text{Null}_{\text{cue}}^{n} + (\beta_{\text{tuning}}(i) \times \text{Tuning}_{\text{cue}})^{n})^{\frac{1}{n}}, \quad (7)$$

$$\text{Null}_{\text{late}}(i) = \alpha((\beta_{\text{null}}(i) \times \text{Null}_{\text{cue}})^{n} + \text{Tuning}_{\text{cue}}^{n})^{\frac{1}{n}}, \quad (8)$$

The filter model:

$$\text{Tuning}_{\text{late}}(i) = \alpha((\beta_{\text{tuning}}(i) \times \text{Null}_{\text{cue}})^{n} + \text{Tuning}_{\text{cue}}^{n})^{\frac{1}{n}}, \quad (9)$$

$$\text{Null}_{\text{late}}(i) = \alpha(\text{Null}_{\text{cue}}^{n} + (\beta_{\text{null}}(i) \times \text{Tuning}_{\text{cue}})^{n})^{\frac{1}{n}} \quad (10)$$

The input gain model:

$$\text{Tuning}_{\text{late}}(i) = \alpha((\beta_{\text{tuning}}(i) \times \text{Null}_{\text{cue}})^{n} + (\beta_{\text{tuning}}(i) \times \text{Tuning}_{\text{cue}})^{n})^{\frac{1}{n}}, \quad (11)$$

$$\text{Null}_{\text{late}}(i) = \alpha((\beta_{\text{null}}(i) \times \text{Null}_{\text{cue}})^{n} + (\beta_{\text{null}}(i) \times \text{Tuning}_{\text{cue}})^{n})^{\frac{1}{n}} \quad (12)$$

The fits were once again performed on responses when the tuning pattern moved in the preferred direction or 30° from the preferred direction. The α and n were fixed to what was determined beforehand for each neuron using Eqns (5) and (6), and the additional fits were performed on the paired response during 200-ms time bins (i refers to the analysed bin number) throughout the last 70 ms of the dual stimulus early and the entire dual stimulus late epoch—allowing us to examine changes in attentional modulation via parameters β_{Tuning} and β_{null} over 1000 ms (i.e., 5 time bins).

The goodness of fit was then compared for the spotlight vs filter model using the correlation coefficient (R^2). Additionally, the goodness of fit for the input gain model was compared to the filter model using the Akaike information criterion (AIC; Akaike, 1973) to account for unequal number of varying parameters (two in the input gain model and one in the filter model). Finally, a Wilcoxon rank-sum test was used to determine which model provided the best fit. Average parameters were also computed across the population of analysed MT neurons for each model, and predicted paired responses during the time window of the dual stimulus late epoch were calculated and plotted for better visual comparison with the real, averaged recorded responses.

Experiment 2

Data from a second experiment in two different male monkeys, which also investigated the effects of attention in area MT, underwent similar analysis (Khayat & Martinez-Trujillo, 2015). Procedures complied with Canadian Council of Animal Care guidelines and were pre-approved by the McGill Faculty of Medicine's animal care committee.

Stimuli

Moving RDPs containing small bright dots at a density of 4 dots per degree^2 (viewing distance = 57 cm) were utilised. The sizes of the RDPs were chosen to fit within the excitatory RF boundaries of each neuron. The RDP's speed and direction were chosen based on the recorded neuron's preferred speed and direction.

Two conditions, attend-preferred alone and attend-null alone, consisted of two moving RDPs located in opposite hemifields and both moving either in the neuron's preferred or antipreferred direction. Two other conditions, attend-null pair and fixation, comprised two pairs of RDPs: two were located inside the neuron's RF, and the other two in the opposite hemifield relative to a central fixation spot. Each pair contained one RDP moving in the neuron's antipreferred direction (null pattern), while the other RDP moved in the neuron's preferred direction.

Behavioural task

A trial began once a monkey pressed a lever and fixated on a fixation spot. After 470 ms, all RDPs (two for the attend-preferred alone and attend-null alone conditions; four for the attend-null pair and fixation conditions) were presented and remained on the screen for the duration of the trial; 350 ms after stimuli onset, a small line (1° length) would either appear next to the fixation spot to cue which of the patterns to attend to, or no cue line was presented and the monkey would have to maintain attention to the fixation spot. For trials that presented the small line, the monkey had to release the lever in response to a brief motion direction change (30° for 118 ms) in the cued RDP. A direction change in the target would occur between 660 and 2900 ms from cue onset, and the monkey was required to respond within a reaction time window of 150–500 ms to receive a juice reward (Khayat & Martinez-Trujillo, 2015).

The uncued null pattern (distracter), in the opposite hemifield, changed motion direction in half of the experimental session's trials and had to be ignored while waiting for the target to change (which took place at least 550 ms after the distracter null pattern underwent a direction change). Note that during the fixation condition, the monkeys were required to maintain attention on the fixation spot until the end of the trial (i.e., no lever release). Only correct trials were analysed. Trials in which the monkey responded before the target direction change, responded incorrectly to the null distracter, or broke fixation were aborted without reward.

Recordings

Each animal was implanted with a titanium head post and a recording chamber (20 mm diameter; Crist Instruments, Hagerstown, MD, USA). The chamber was positioned over a craniotomy of the parietal bone providing access to area MT (Khayat *et al.*, 2010). Extracellular recordings from area MT were collected from two male monkeys using tungsten electrodes (1–2 MΩ at 1 kHz; FHC, ME, USA) and a Plexon data acquisition system. Signals were amplified, filtered between 250 Hz and 8 kHz and digitised at 40 kHz. All single units were sorted offline. Cells were determined as belonging to area MT according to response properties and electrode positions visualised through magnetic resonance imaging scans (Khayat *et al.*, 2010).

Analysis

The spike density function (1 ms resolution, Gaussian kernel with σ = 20 ms) for each neuron was computed. Each neuron's responses were normalised to the neuron's maximum response and then averaged across neurons to obtain the population's activity. The maximum response was obtained from the attend-preferred alone condition after stimuli onset but before the presentation of the cue line.

Results and statistical analysis

Experiment 1

Behavioural performance

We recorded the responses of 77 MT neurons from two male monkeys (F and M) during the task illustrated in Fig. 1B. Monkey F correctly detected the change in the target and ignored the distracter (hits) on 88% of the trials, incorrectly responded to a direction change in the distracter on 5% of trials, and responded too early or withheld the response on 7% of trials. Monkey M performed similarly at 89, 5 and 6%, respectively. A proxy for the behavioural effects of attention can be observed when considering the relative success of each monkey during conditions in which distracters were ignored. Of the 88% hit trials performed by Monkey F, 40% was due to correctly ignoring distracters. Monkey M exhibited a similar success rate with 41% of the 89% hit trials attributed to correctly ignoring distracters.

The monkeys' reaction times follow a linearly descending slope of −0.047 (R^2 = 0.11, LCI = −0.053, UCI = −0.041) as a function of time from the beginning of the trial (black line in Fig. 2A). The slope was statistically different from zero (as indicated by its associated confidence intervals), indicating that the reaction times slightly but significantly decreased as the probability of the change increased during the trial.

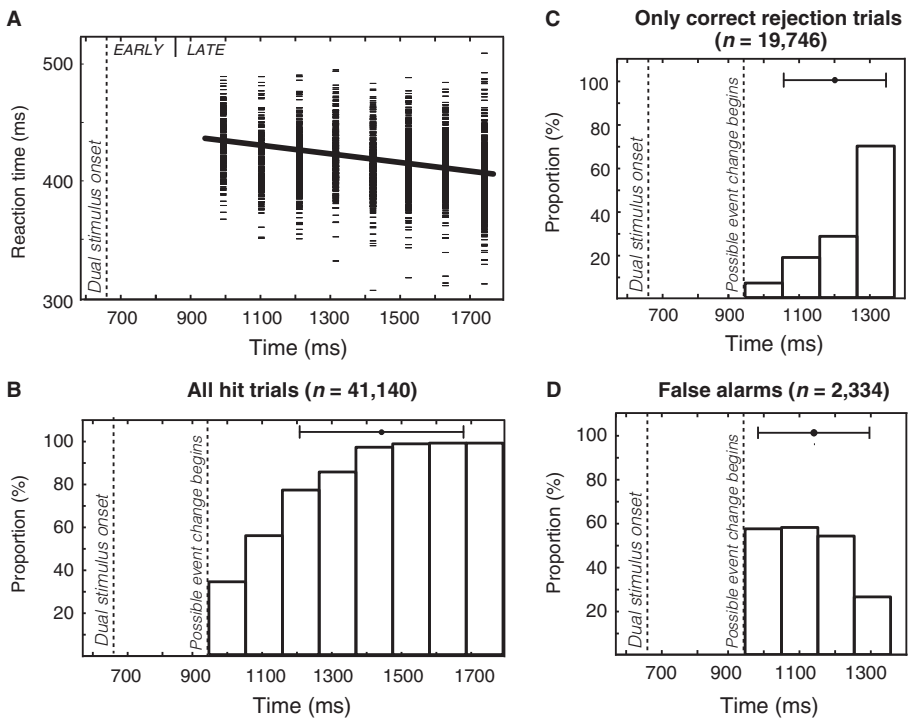

FIG. 2. Reaction times and performance rates during Experiment 1. (A) The short, thin black lines are the reaction times plotted with a similar bin size as the event time histogram depicted in Fig. 1C. A best-fit line, as a function of time since the start of at trial, is shown by the thick black line. (B, C) The improved performance as a function of time for all hit trials and exclusively distracter trials, respectively. (D) The decreased error over time during distracter trials. All performance rate histograms are binned with a similar bin size as the reaction times in panel A. The mean and standard deviation are also illustrated at the top of each histogram. Note, the histogram for trials with distracters (i.e., C and D) only extends until 1400 ms, as distracter changes could only occur at 940–1300 ms to leave adequate time (500–600 ms) before a target motion change.

To determine whether the hazard rate satisfied stronger attention towards the end of the trial, we assessed the monkeys' performance as a function of time from trial onset. We binned performance for all hit trials (Fig. 2B), correct rejection trials (Fig. 2C) and false alarm trials (Fig. 2D) with a similar bin size as the reaction time histogram in Fig. 2A. The binned performance appeared to also follow an exponentially increasing pattern like that observed for the hazard rate. Therefore, attention seems to get stronger over the course of a trial, as is evident in the monkeys' improved performance over time. This is likely owing to the monkeys' progressively stronger commitment as a function of change probability.

Accounting for potential adaptation effects. Figure 3A and B displays the heat maps across all 77 recorded neurons for both attention conditions. Of the 77 neurons, 24 exhibited a significantly greater response in the attend-null than in the attend-preferred direction during the dual stimulus early epoch ($P < 0.029$, Wilcoxon rank-sum test). This significantly different response during the dual stimulus early epoch indicates that these neurons' responses were likely influenced by motion adaptation. Heat maps for these 24 neurons in both attention conditions are illustrated in Fig. 3C and D. The remaining 53 neurons did not show a significant difference in response between the two attention conditions during the dual stimulus early epoch ($P > 0.056$, Wilcoxon rank-sum test) (Fig. 3E and F).

While we presented the cue for a relatively short period (400 ms) to avoid strong effects of motion adaptation, the stronger initial response during the attend-tuning condition's cue period in response to preferred motion direction (0°) and directions close to it (e.g., ±30°, ±60°) may have led to response adaptation of the classical excitatory RF component—potentially decreasing neuron response during the dual stimulus early epoch. It is also possible that presentation of the null cue alone may have adapted the 'gain control' or normalisation signal component (Solomon & Kohn, 2014), as this stimulus was observed to decrease the firing rates of some neurons below baseline during the cue epoch. Adaptation of the normalisation signal potentially increased neuron response during the dual stimulus early epoch (Fig. 3F).

When computing the SSI_P and SSI_N indices, the 24 neurons that had significantly different responses during the dual stimulus early epoch also demonstrated an SSI_N greater than 1 (Eqns 3 and 4) (Fig. 3G). This suggests that the initial transient during the dual stimulus early epoch of the attend-null condition was even greater than the neuron's response to the preferred stimulus alone, leading us to suspect that the normalisation signal component in these neurons was likely adapted upon exposure to the antipreferred cue and may have caused a strong 'rebound' response (Solomon & Kohn, 2014). These spatial summation indices supported the exclusion of these 24 neurons from further analysis.

Single neuron example

The spike density functions of one example neuron to the 12 different combinations of the null and tuning patterns in both attentional conditions are shown in Fig. 4A and C. During the cue epoch (yellow background), the different directions of the tuning pattern presented inside the RF evoked greater responses compared to presentation of the null pattern. The stationary cue presentation (that was identical across all conditions) evoked an initial transient response, followed by a response to the cue motion that varied depending on the cue's direction. For cue directions ±30° from the antipreferred direction (see colour legend), the cell did not respond beyond the transient cue onset response and showed a transient cue offset response. For the other cue directions, the response magnitude reflected the cell's direction tuning. Following the cue offset, the response decreased for the short time interval when only the fixation cross was on the screen. During the dual stimulus early epoch, the responses in both attentional conditions were very similar. Shortly after, during the dual stimulus late epoch, the responses diverged substantially between the conditions. This can be better visualised in the heat maps for this neuron in Fig. 4B and D. The response when attending to the tuning pattern remained at similar levels until the end of the trial, while responses when attending to the null pattern strongly decreased over time. This suggests that allocating attention to the null pattern strongly suppressed the response component evoked by the tuning distracter pattern.

Population activity

The Gaussian functions that were fit to the responses of the 53 neurons across epochs were pooled to construct a population tuning function for each condition (heat maps in Fig. 5A and B). These profiles were similar to those obtained for the single neuron example. We also obtained the time series of the four Gaussian parameters (gain, centre, width and baseline; refer to Eqn 2) for each neuron and pooled them across neurons to obtain average population tuning curve parameters (Fig. 5C–J).

During the dual stimulus early and dual stimulus late epochs of both the attend-tuning and attend-null conditions, the centre, width and baseline did not show a clear trend across time (Fig. 5F, H and J). Additionally, during the dual stimulus early epoch, the gain was not significantly different in the attend-null (blue) relative to the attend-tuning (red) condition ($P = 0.45$, Wilcoxon rank-sum test) (Fig. 5D). During the dual stimulus late epoch, the gain of the population tuning curves in the two conditions differed (Fig. 5D). During the attend-tuning condition, the response gain stays relatively constant and can be best represented by a straight line with a slope that does not differ from 0 ($m = 9.55 \times 10^{-5}$, $R^2 = 0.73$, LCI = −0.0014, UCI = 0.078; AIC comparisons between all models resulted in $P > 0.98$, rendering the linear fit sufficient) (Fig. 5D). On the other hand, during the attend-null condition, the tuning curve gain progressively decreased over time. This decrease was best modelled as a sum of two exponential functions with an earlier rapid and later slow decay rate ($R^2 = 0.94$; AIC comparisons for the double-exponential were significantly better than the linear fit ($P = 0.0043$, Wilcoxon rank-sum test) and the single exponential fit ($P = 0.034$, Wilcoxon rank-sum test)); see equations in Fig. 5D. This result indicates that, during the dual stimulus late epoch, attending to the tuning pattern did not significantly modulate the gain of responses over time, while attending to the null pattern progressively decreased the gain of responses over time from target–distracter onset.

Previous studies that employed a change detection task with hazard functions similar to the one used in our study (Fig. 1C) have shown that the hazard rate utilised determines the shape of the attentional modulation of responses (Ghose & Maunsell, 2002). To investigate this, we used the robust correlation toolbox (Pernet *et al.*, 2012) to do a bootstrap correlation of the hazard rate (Fig. 1C) to the population neural response gain during the attend-null condition (Fig. 5D). We observed that the longer it takes a null target to change direction, the more significant the suppression of the distracter tuning pattern (Pearson's correlation coefficient (r) = −0.86, $P = 0.0067$, LCI = −0.99, UCI = −0.83).

Using the same analysis with the robust correlation toolbox, we related the observed neuronal trends to behaviour. As mentioned earlier, monkeys' reaction times slightly but significantly decreased

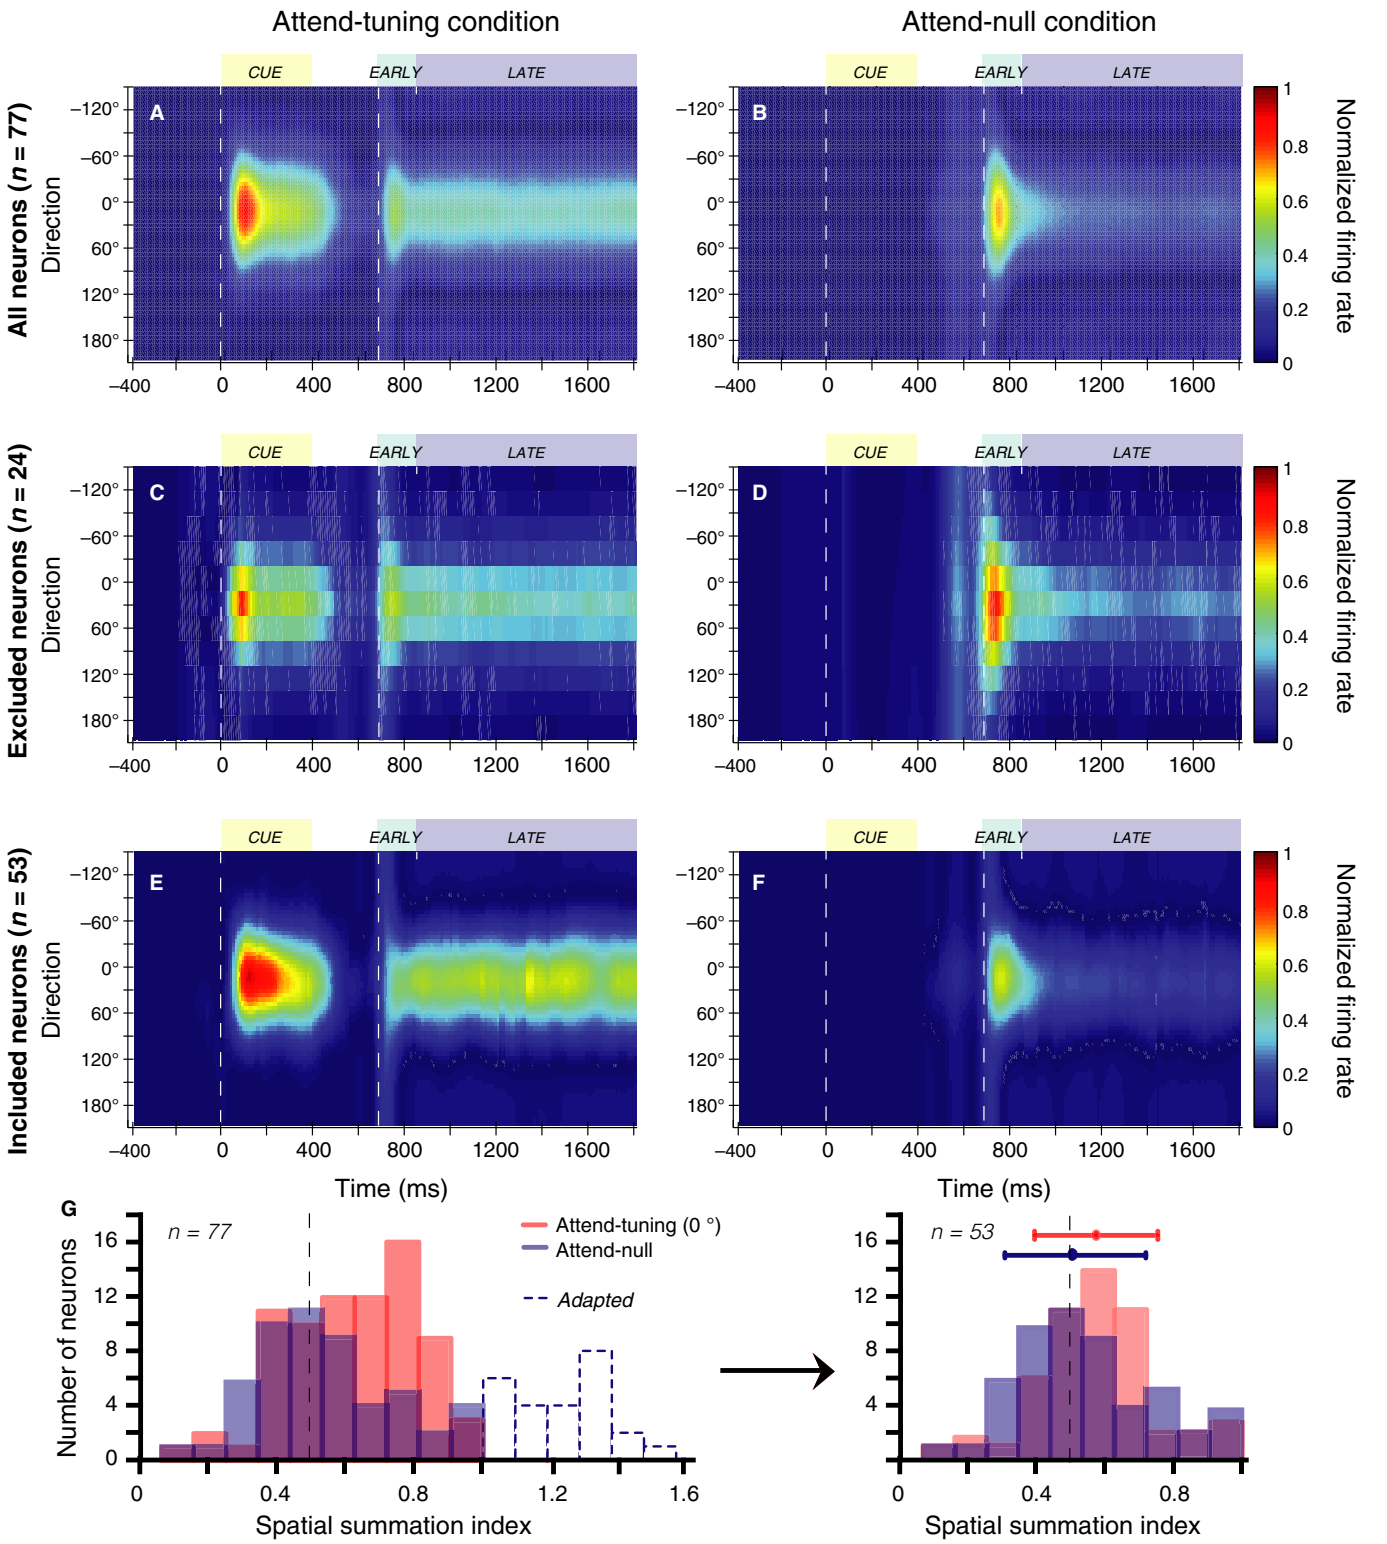

FIG. 3. Population heat maps with different inclusion criteria. Heat maps represent averaged normalised neuronal firing rates for all pairs of the null pattern and tuning patterns (*y*-axis) across time (*x*-axis). The different analysis epochs are highlighted: yellow indicates the cue epoch; light blue, the dual stimulus early epoch; muted grey, the dual stimulus late epoch. The normalised neuronal activity of all recorded neurons (n = 77) during (A) the attend-tuning condition and (B) the attend-null condition. (C and D) Heat maps for the 24 neurons that were excluded from further analysis. These neurons exhibited suppression below baseline during the cue period when the antipreferred pattern was presented (attend-null condition); these same neurons had a significantly greater response during the dual stimulus early epoch of the attend-null condition as compared to attend-tuning condition. (E and F) Heat maps for the remaining 53 neurons. These neurons show similar neuronal trends to panels A and B and were used for all Experiment 1 analysis. (G) The distribution of spatial summation indices (SSI; Eqns 3 and 4) before (on left) and after (on right) excluding the 24 neurons that showed an $SSI_N > 1$ during the attend-null condition.

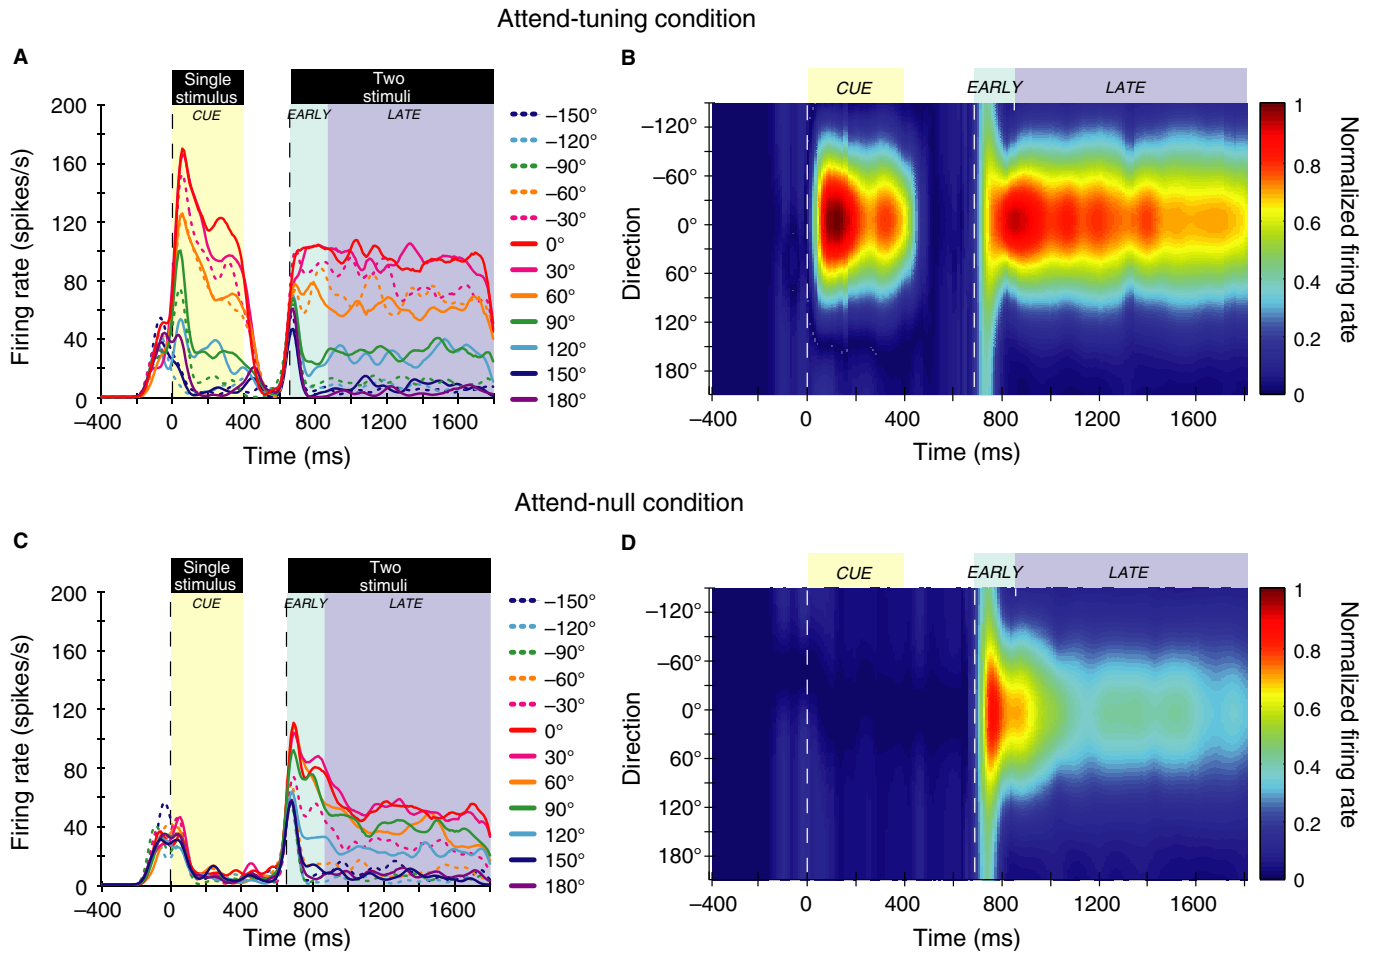

FIG. 4. Experiment 1 single neuron example. (A) Spike density functions during the attend-tuning condition were obtained by convolving each spike with a Gaussian kernel of 30 ms (Martinez-Trujillo & Treue, 2004). Responses were aligned so that the neuron's preferred direction is 0°; each colour denotes the spike density function for a different motion direction. (B) Heat maps of the normalised firing rates for the attend-tuning condition. (C and D) Spike density functions and heat map for the attend-null condition. Note, the cue epoch refers to the period during which the random dot pattern (RDP) was moving in the target direction for 400 ms; the slight activity prior to the start of the cue period is due to the earlier presentation of the stationary RDP (which marks target location) for 200 ms.

as the probability of the change increased during a trial. We found that the decrease in reaction times with longer attention allocation correlated with the increased suppression of the tuning stimulus over time during the attend-null condition ($r = 0.9968$, $P = 1.2 \times 10^{-77}$, LCI = 0.9961, UCI = 0.9975). This correlation was not significant when relating the reaction times to the maintained neuronal response during the attend-tuning condition ($r = -2.7 \times 10^{-15}$, $P = 1.00$, LCI = -2.0×10^{-14}, UCI = 1.8×10^{-14}). This strongly suggests that suppressing the response to a preferred distracter stimulus, when attending to an antipreferred stimulus, was strongly linked to behavioural performance. On the other hand, the time course of the response when attending to a preferred stimulus did not correlate with performance changes.

Modelling spatial integration. We fit the power-law summation model suggested by Britten & Heuer (1999) (Eqns 5 and 6) to predict responses to both stimuli during the dual stimulus early period from the responses to the cue patterns. The summation model has parameters α (scale factor) and n (exponent). We included pairings of the preferred and null pattern as well as of patterns moving in a direction away from the preferred by 30° and the null pattern. This was because stimulus pairs with tuning patterns moving away from the preferred direction by more than 30° led to low firing rates and model fits that substantially deteriorated. This relatively simple model was initially used to obtain 'seed' parameters for the next step—modelling attentional modulation over time.

FIG. 5. Differences in population Gaussian tuning curve parameters over time for Experiment 1. (A) and (B) exhibit average heat maps for the attend-tuning and attend-null conditions, respectively, across the 53 analysed middle temporal (MT) neurons (repeated from Fig. 3E and F). The maps resemble those observed in the single neuron. (C–J) Average Gaussian tuning curve parameters across neurons—gain (cue epoch: panel C, and dual stimulus early + dual stimulus late epochs: panel D), centre (panels E and F), width (panels G and H) and baseline (panels I and J), across time for each condition type. The red trace refers to the parameters averaged across neurons obtained for the attend-tuning conditions, and the blue for the attend-null conditions. The fainter red and blue plots surrounding each trace represent standard error. The best-fit linear model (red, dashed line) and double-exponential model (blue, dashed line) during the dual stimulus late epoch for the gain parameter ($G(t)$) of the attend-tuning and attend-null conditions, respectively, are also depicted.

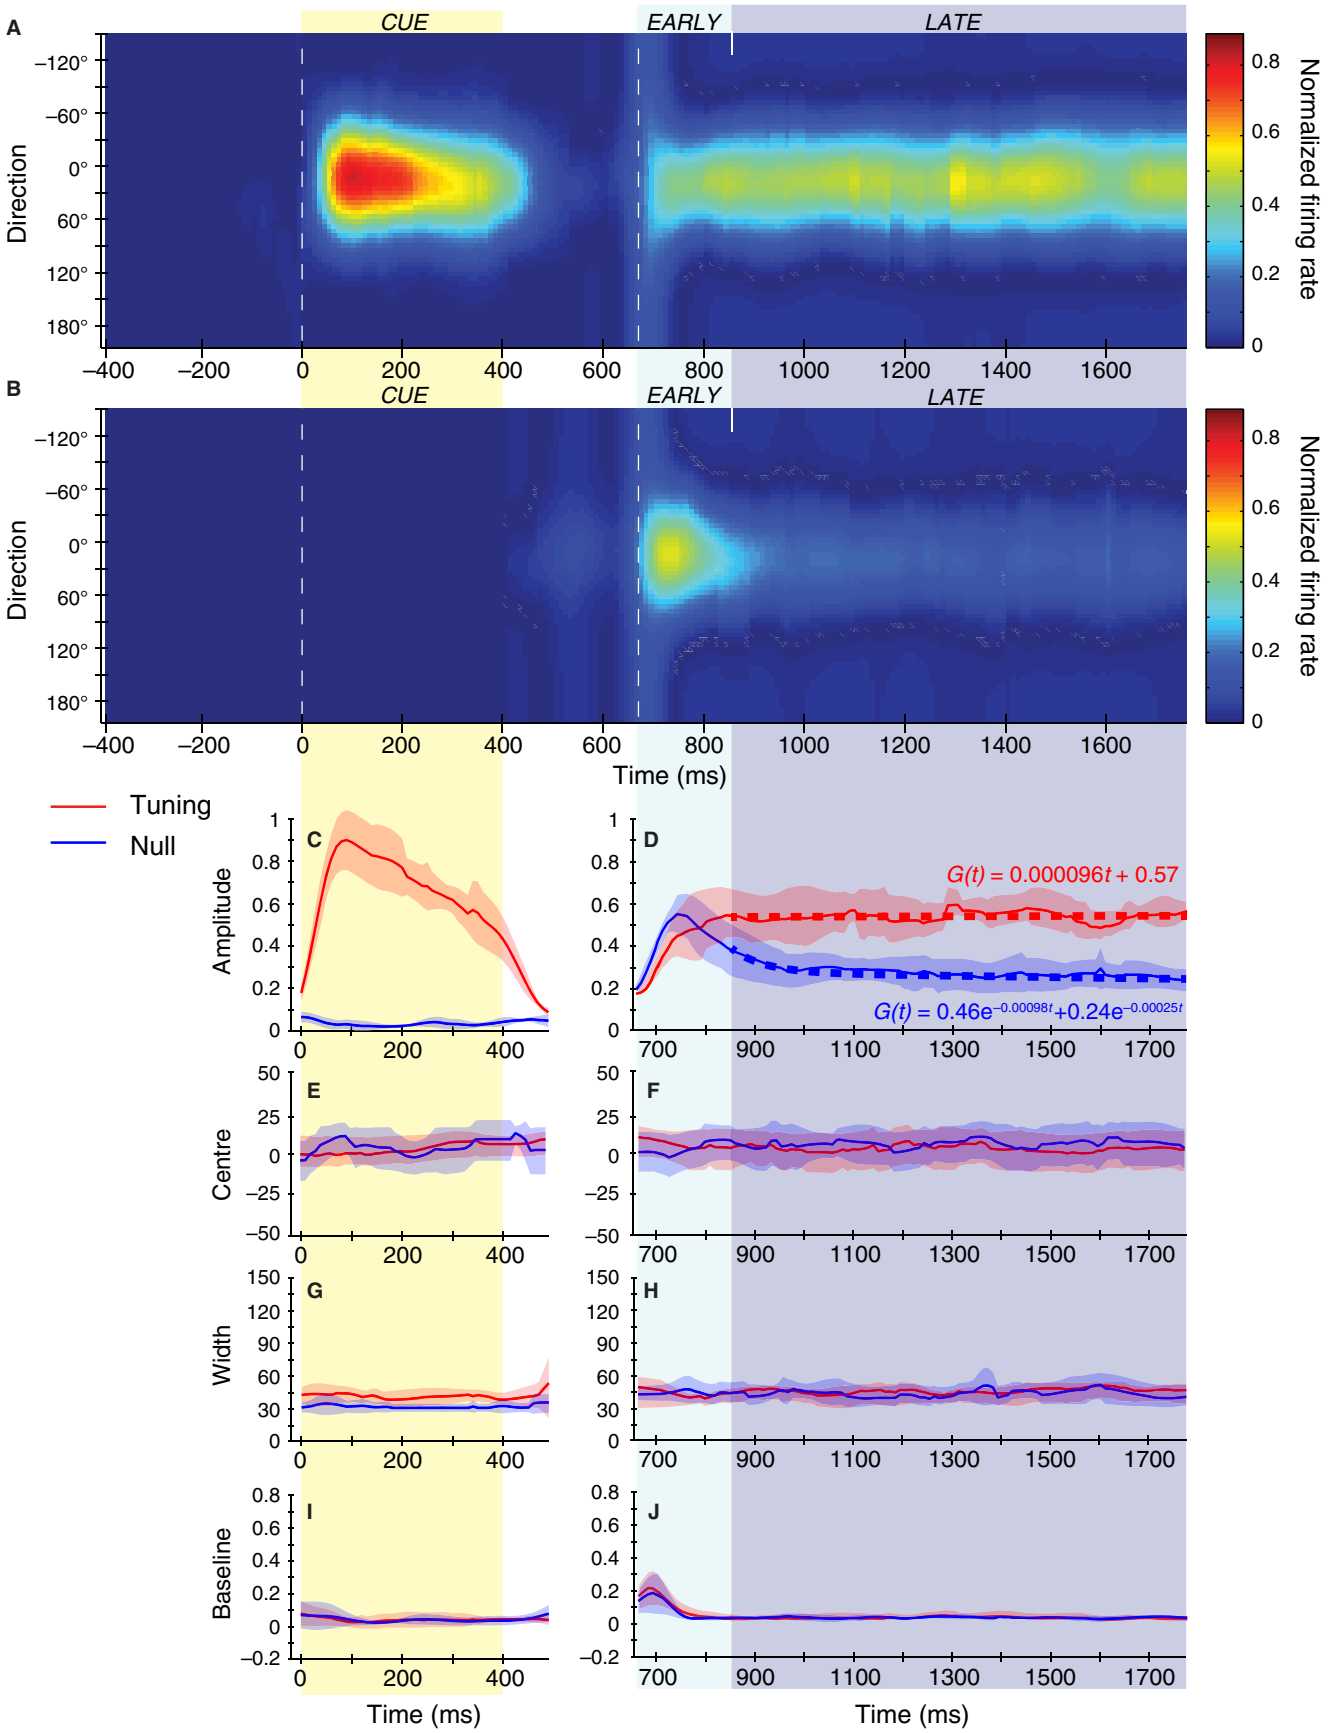

Histograms of fitted α and *n* for all neurons are plotted in Fig. 6A and B, respectively. The mean α across neurons was 0.56 [standard deviation (SD) = 0.094], and the mean *n* was 1.27 (SD = 0.70). The models for each neuron had consistently good fits with a mean correlation coefficient (R^2) of 0.82 (SD = 0.11). Given these results, an averaging model appears to best represent spatial integration in MT when presented with two coherently moving, high-contrast RDPs. This was also previously suggested in the studies by Recanzone *et al.* (1997) and Snowden *et al.* (1991), which also investigated spatial integration in MT with RDPs varying in motion direction. However, if spatial integration were solely explained by the averaging of inputs to a cell, then attention should influence responses equally when directed to a preferred or an antipreferred stimulus. Considering our results differed in the attend-tuning and attend-null conditions, an additional parameter to α and *n* that biases how inputs to a neuron are combined over time during attention is needed.

Three models of attention were tested following similar procedures as described by Ghose & Maunsell (2008): spotlight (Eqns 7 and 8), filter (Eqns 9 and 10) and input gain (Eqns 11 and 12). These models also correspond to the scenarios illustrated in Fig. 1A. These models incorporate either one or two additional parameters to the power-law summation model so that gains imposed by attentional modulations can be selectively applied to stimuli that are simultaneously presented and differentially attended (Ghose & Maunsell, 2008). Note, when testing each model, the computed α and *n* parameters were incorporated and fixed across time. The only parameter in the new equations was the gain parameter(s), β, that was multiplied by the responses to individual stimuli. By examining how the β parameters change across time, we obtain insight into MT's attentional modulation dynamics during the different attention conditions.

In the spotlight model, the gain parameter, β, was only applied to the attended stimulus. We find that during the attend-tuning condition, when the gain β is applied to the responses of the tuning pattern as the attended target (Eqn 7), β stays relatively constant across time and is more often less than or not different from 1 (mean β = 0.83, SD = 0.037, R^2 = 0.48 ± 0.041) (red line in the left most panel of Fig. 6C). This rather constant β over time is best fit

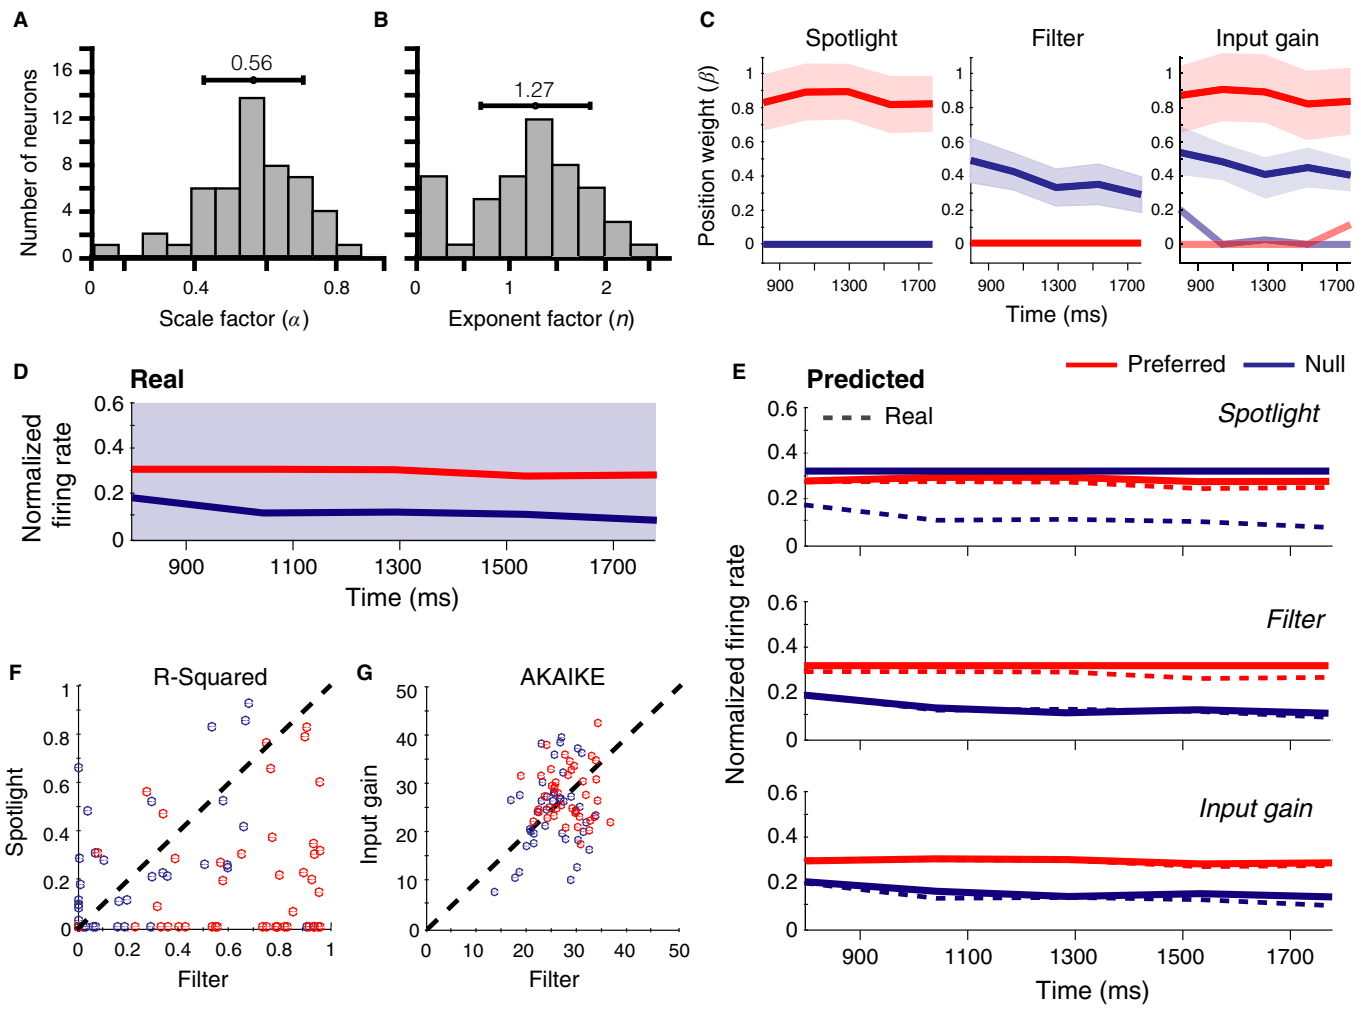

FIG. 6. Models of spatial integration. (A, B) Spatial integration of the 53 middle temporal (MT) neurons was modelled based on Eqns (5) and (6) for the attend-tuning and attend-null conditions, respectively. (A) Histograms of α and (B) of *n* along with their means and standard deviations are illustrated. (C) Weights associated with the attended stimulus, also known as β, were additionally fit to assess attentional modulation. β changes over time when fit by each of the three models for both attend-tuning conditions, in red, and attend-null conditions, in blue. (D) The normalised firing rate across the 53 included neurons obtained during the dual stimulus late epoch. (E) The predicted normalised firing rate during the dual stimulus late epoch using the parameters fit for all three models (solid lines). The neurons' responses from (D) are redrawn in as dashed lines for comparative purposes. (F) Scatter plot of the R^2-values corresponding to the spotlight vs filter model for the attend-tuning (red) and attend-null (blue) conditions. (G) Scatter plot of Akaike information criterion (AIC) of the input gain vs the filter model.

by a horizontal line with a slope not different from 0 ($m = -0.0084$, LCI = −0.048, UCI = 0.031, $R^2 = 0.67$). When the gain β is applied to the null pattern responses in the attend-null condition (Eqn 8), it yields values close to 0 (mean β = 0.010, SD = 1.07×10^{-13}, $R^2 = 0.14 \pm 0.013$) (blue line in Fig. 6C). The spotlight model fails to fit the attend-null condition data. As demonstrated in the top panel of Fig. 6E, the predicted values obtained using the spotlight model, especially for the attend-null condition, do not reflect the real values observed (dashed blue vs solid blue lines).

In the filter model, the β gain parameter is multiplied by the responses to the unattended distracter. In the attend-tuning condition (Eqn 9), β assumes values close to 0 (mean β = 0.010, SD = 4.11×10^{-12}, $R^2 = 0.59 \pm 0.057$) (red line in the middle panel of Fig. 6C), which allows the neuronal response to be dominated by the contribution of the responses to the tuning pattern alone. On the other hand, during the attend-null condition (Eqn 10), β assumes values lower than 1 (mean β = 0.43, SD = 0.057, $R^2 = 0.87 \pm 0.070$) (blue line in the middle panel of Fig. 6C), filtering out the contribution of the distracter tuning pattern to the response. Moreover, β grows smaller over time as was best modelled by a line with a significantly negative slope ($m = -0.031$, LCI = −0.048, UCI = −0.013, $R^2 = 0.71$). Thus, as the trial progresses, the suppression of the tuning pattern grows larger. As shown in the middle panel of Fig. 6E, when comparing the obtained values and the filter model's predicted values, this model provides a good description of the data.

The third model, input gain, includes two β parameters so that attention affects responses to both the attended and unattended stimuli. During the attend-tuning condition (Eqn 11), the β parameter associated with the tuning pattern remained relatively close to 1 (mean β = 0.87, SD = 0.036) and was relatively constant across time (as modelled by a straight line: $m = -0.015$, LCI = −0.046, UCI = 0.016, $R^2 = 0.59$) (bright red line in the rightmost panel of Fig. 6C), while that associated with the null pattern was again close to 0 (mean β = 0.033, SD = 0.052) ($R^2 = 0.94 \pm 0.059$) (light red line). In the attend-null condition (Eqn 12), the β parameter associated with the tuning pattern remained relatively close to 0 (mean β = 0.056, SD = 0.089), while that associated with the null pattern was consistently less than 1 (mean β = 0.46, SD = 0.055) ($R^2 = 0.96 \pm 0.039$) (light blue line) and decreased significantly over time (as modelled by a linear regression with a significantly negative slope: $m = -0.041$, LCI = −0.064, UCI = −0.017, $R^2 = 0.73$). This model also provides a good fit of the data (bottom panel of Fig. 6E).

We compared the goodness of fit of the three models. First, we compared the filter vs the spotlight model, which share a similar number of parameters. When comparing the R^2 across neurons and for both attentional conditions (Fig. 6F; red dots refer to attend-tuning condition fits and blue dots to attend-null condition fits), most dots fall below the diagonal. Therefore, neurons consistently have a significantly better fit with the filter than with the spotlight model (attend-tuning: $P = 2.45 \times 10^{-7}$; attend-null: $P = 1.38 \times 10^{-7}$, Wilcoxon rank-sum test). Next, we compared the filter model against the input gain model. Because the models have a different number of free parameters (one parameter in the filter and two parameters in the input gain), we used the AIC to compensate for such a difference (Fig. 6G). Most dots fall along the diagonal, representing equivalent AIC values across neurons for both the filter and input gain models (attend-tuning: $P = 0.41$; attend-null: $P = 0.77$; Wilcoxon rank-sum test). The additional parameter in the input gain model (relative to the filter model) does not significantly improve the goodness of fit. Taken together, our data were not explained by a spotlight model, but could be explained by either the filter or input gain models. Furthermore, the filter model explained the data with one fewer parameter than the input gain model.

Experiment 2

Data from a second attention experiment, in two different male monkeys, were also analysed. This experiment did not have the attend-tuning/preferred condition; rather, it was designed to explore neuronal response suppression when attending to the null pattern while excluding the effects of adaptation. The experiment included a fixation condition, in which no cue was presented and the animals were required to maintain fixation until they detected a change in the colour of the fixation spot. Refer to Fig. 7A for further details on the task paradigm. In this experimental design, the cue was stationary, did not precede the onset of the two patterns and was presented well outside the RF. No motion adaptation to the cue should have taken place.

Behavioural performance

Monkey S performed correctly in 90% of trials, incorrectly responded to distracters in 6% of trials and responded too early or withheld the response in 4% of trials. Monkey L performed similarly with 90, 5 and 5%, respectively.

Single neuron example

The responses of an example neuron in the different conditions are illustrated in Fig. 7B. As anticipated, presentation of the preferred stimulus alone in the RF evoked the greatest response upon stimulus onset followed by a sustained firing until detection of the motion change (red line, Fig. 7B). The null stimulus alone in the RF evoked the weakest response (dark blue line). During the fixation condition, the firing rate of the neuron increased following stimuli onset and remained at a similar level for the rest of the trial (black line). On the other hand, the firing rate during the attend-null pair condition initially increased but progressively dropped after cue presentation towards the end of the trial (purple line). While pronounced in this particular neuron, the neural pattern resembled that observed in Experiment 1 for the attend-null condition.

Population activity

The responses of 51 neurons during all conditions were normalised and averaged (Fig. 7C). To estimate attentional response modulation, we fitted linear, exponential or double-exponential equations to the responses as a function of time after the presentation of the cue. Responses in the fixation condition initially increased, with a profile typical of MT units, and later reached a plateau that remained after the cue presentation and until the end of the trial (black line). The population response during the fixation condition was best fit by a flat line with a slope that was not significantly different from 0 ($m = -6.19 \times 10^{-6}$, $R^2 = 0.71$; LCI = -1.46×10^{-5}, UCI = 2.13×10^{-6}; AIC comparisons $P > 0.11$, Wilcoxon rank-sum test, rendering the linear-fit sufficient) (dashed white line). On the other hand, responses during the attend-null pair condition split from fixation responses after cue onset and continued to decrease until the end of the trial. These responses were best fit by a sum of two exponential functions, with an initial fast decay followed by a slower decay

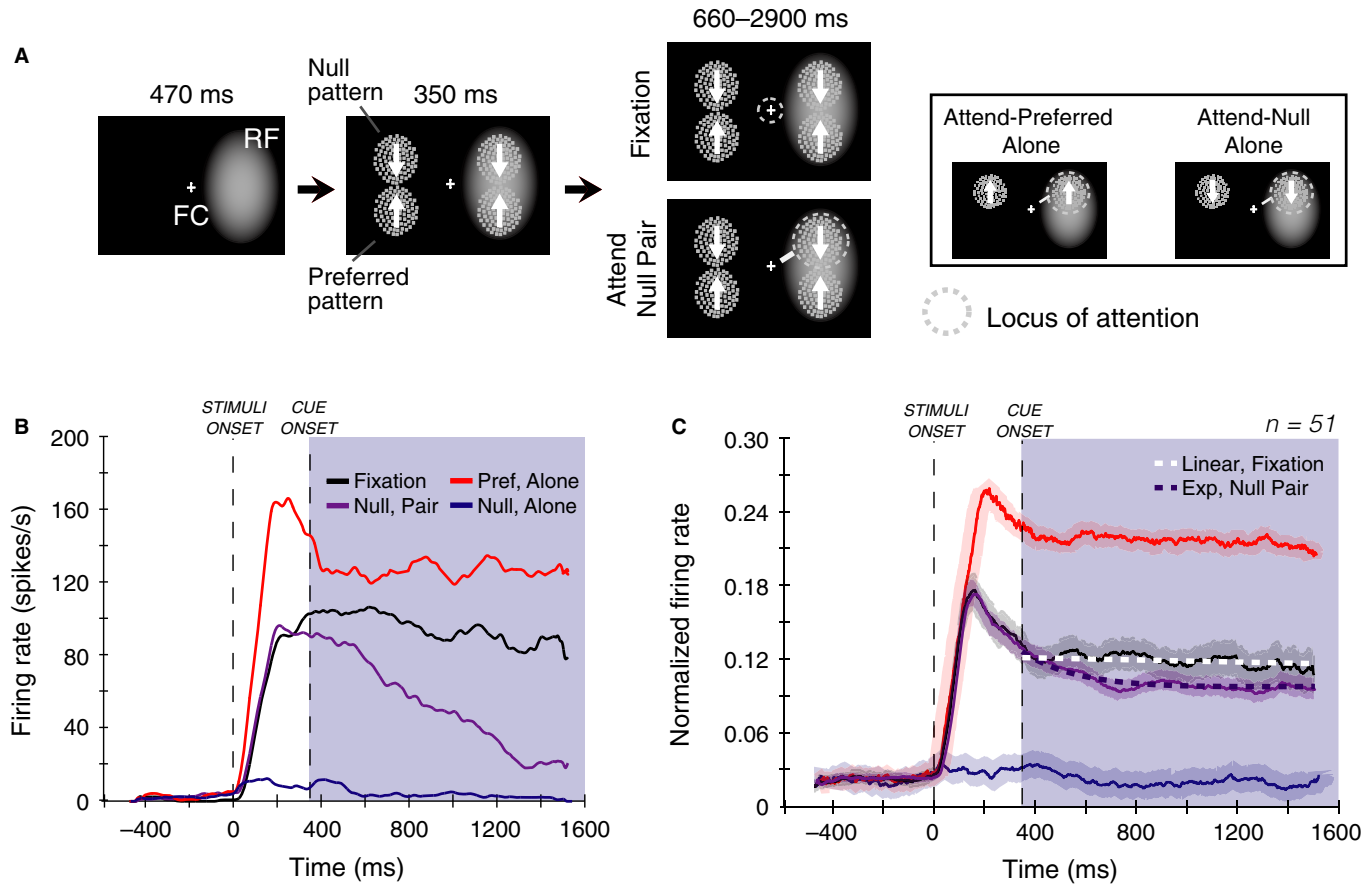

FIG. 7. Experiment 2 task paradigm and results. (A) As in Fig. 1, the black rectangles represent the computer screen on which the stimuli were presented, the cross is the fixation spot (i.e., fixation cross (FC)) and the random dot patterns (RDPs) the stimuli. The monkey maintained fixation on the FC during the trial length. RDPs only moved in one of two directions, the preferred or the antipreferred direction of the recorded neuron. The monkey was cued which pattern was the target by a bar that appeared 350 ms after foveating on the fixation cross and pushing down on the lever. The monkey was either exposed to four stimuli—with two in the receptive field (RF; in grey), during the fixation and attend-null conditions, or two stimuli—with one in the RF, during the attend-preferred alone or attend-null alone conditions (portrayed in the inset). (B) A single neuron example response during all four conditions, attend-null (purple), fixation (black), preferred alone (red) and antipreferred alone (blue). (C) Population firing activity across the 51 analysed neurons of Experiment 2, during all tested conditions. Models of best fit for normalised firing rate across time ($FR(t)$) are also illustrated, specifically linear regression for the fixation condition ($FR(t) = -0.0000062t + 0.12$; $R^2 = 0.71$) and a sum of exponential functions ($FR(t) = 0.14e^{-0.0092t} + 0.10e^{-0.0000059t}$; $R^2 = 0.86$) for the attend-null condition.

($R^2 = 0.86$; AIC comparison was significantly greater for the double-exponential than for the linear fit ($P = 0.00060$, Wilcoxon rank-sum test) and the single exponential fit ($P = 0.00057$)) (dashed purple line). The trends during the fixation condition and attend-null pair condition of Experiment 2 were similar to the fits of the gain parameters during the attend-tuning and attend-null conditions of Experiment 1, respectively.

Discussion

We measured the modulation of MT neuron responses when attention was directed to one of two RDPs, which were simultaneously presented within a neuron's RF. Using the neuron's response before strong attentional modulation as the reference, we demonstrated progressively decreasing population activity over time when directing attention to a null stimulus. This remained the predominant effect of attentional modulation even after accounting for possible adaptation. There were minimal changes in response when directing attention to a tuning stimulus. In contrast to previous studies, our data were best described by the filter and input gain models, but not by the spotlight model.

Reference response and attentional modulation

Previous studies have proposed that when two stimuli fall inside the RF of a visual neuron, attending to the preferred stimulus increases the response, while attending to a less preferred stimulus decreases the response (Moran & Desimone, 1985; Treue & Maunsell, 1996; Reynolds *et al.*, 1999; Martinez-Trujillo & Treue, 2004; Boynton, 2009). A common feature in these studies is that response increases/decreases with attention are often described relative to the responses of neurons to the stimulus pair when attention is allocated outside the RF. However, attending outside the RF decreases a neuron's response relative to attending inside the RF (Treue & Martinez-Trujillo, 1999; McAdams & Maunsell, 2000). Such a reference response may bias the magnitude of the response modulation when directing attention to one of two stimuli inside a neurons' RF (see Fig. 1A).

We propose the use of a reference response obtained earlier in a trial, when attentional modulation is less likely to have taken place. In Experiment 1, the cue was presented at the beginning of the trial and was followed by a pause until the dual stimuli (target and distracter) were presented. We used the neuronal response when the dual stimuli were presented and did not yet significantly differ

between attentional conditions as the reference response. It is possible that the animal's attention was already on the target or maintained on the fixation point during this time window; however, three findings would suggest otherwise. First, no direction changes happened in the target during this period. Second, the hazard function and performance rates indicate that attention was increasingly allocated to the target as a trial advanced. And third, no attentional modulation was observed during this dual stimulus early period. In Experiment 2, the dual stimuli were presented before the cue. The reference response utilised was the neuronal response once the cue was presented.

Even after accounting for the potential effects of adaptation, similar attentional modulations were observed in both experiments when using the reference response suggested here: a progressive decrease, with an initial fast decay and later slow decay, when attending to the null pattern in the presence of a tuning pattern; a maintained response when attending to the tuning pattern in the presence of a null pattern or when maintaining fixation.

Models of attentional modulation

For Experiment 1, we fit our data with three different models to investigate the attentional modulation of responses: spotlight, filter and input gain (Ghose & Maunsell, 2008). We found that the spotlight model did not fit our data well. Rather, the filter model provided a good fit similar to that from the input gain model, and did so with one fewer parameter.

In a similar experiment, Ghose & Maunsell (2008) reported that their data were best described by the input gain model followed by the spotlight and finally the filter model. We consider at least two explanations for these differences. First, it is possible that the Gabor stimuli used by Ghose and Maunsell did not drive the neurons into the saturation part of the contrast response functions, while the high-contrast RDP used in our study did. When high-contrast stimuli saturate a cell's response, attention can no longer increase the responses to stimuli in the RF (Reynolds *et al.*, 2000), but it can decrease the responses to distracters (Martinez-Trujillo & Treue, 2002). Second, to model spatial integration they used the response to the single patterns when attending outside the RF, while we used a response within the RF. Using the attend-out response could underestimate the response to either single pattern relative to when attention is directed inside the RF, and this underestimation may more extensively affect observed responses to the preferred pattern (i.e., yield higher firing rates). That said, the input gain model also fit our data well, rendering it an appropriate description of our data.

Another study reported an asymmetry in the attentional modulation of responses to the simultaneous presentation of a preferred and an antipreferred stimulus in the RF of MT neurons (Ni *et al.*, 2012). When attention was directed to the preferred stimulus, the neurons' responses increased—resembling the response to that stimulus when presented alone with attention outside the RF. On the other hand, when attention was directed to the antipreferred stimulus, the neurons' responses changed by only a small amount; their results were more compatible with the spotlight and input gain models than with the filter model. Ni *et al.* (2012) termed this model the tuned normalisation model. Again, these authors used Gabor stimuli with different contrast levels; thus, a similar explanation for the difference in results as that for between Ghose & Maunsell (2008) and our study also applies here. Moreover, they used short stimulus presentations (200 ms), focusing on the initial period of MT neurons' responses to the two stimuli in the RF. In contrast, we used longer stimulus presentations to focus on the later (sustained) response phase. The dynamics of response normalisation and attentional modulation may differ between these two phases (Khayat *et al.*, 2010).

One may consider the input gain model as a generalisation of the three models: fixing one parameter of the input gain model (Eqns 11 or 12) to equal 1 leads to either the filter or spotlight model. Particularly with this study, the input gain model predicted a suppression of neuronal firing rate when attention was directed towards a null stimulus (i.e., $\beta_{null} \neq 1$) and fit a flat line when attention was directed towards a preferred stimulus or a stimulus moving $\pm 30°$ from the preferred direction (i.e., $\beta_{tuning} \approx 1$). This is equivalent to the prediction using the filter model, only the filter model does so with one parameter. Our findings and their relation to those from previous research have led us to believe that the model that best fits data—whether from area MT or other visual extrastriate areas—may depend on the task and stimuli used. For example, low-contrast stimuli, which do not drive the neurons to saturation, tend to favour enhancement of responses to targets (spotlight and input gain/tuned normalisation), while high-contrast stimuli, which drive cells to response saturation, favour suppression of responses to distracters (filter model and input gain/tuned normalisation).

The results presented here suggest that conclusions drawn about attentional modulation change depending on the reference response chosen. Thus, the degree of target enhancement vs distracter suppression as well as other inferences made when modelling attentional modulation may be contingent on several factors. Besides the reference response, such factors also include the type and contrast of stimuli presented in the RF, the task and its difficulty, the recorded area, the properties of the individual recorded neurons and duration of time over which attention is analysed. Overall, attention is a flexible mechanism that interacts with several variables to shape visual circuit dynamics and, consequently, the processing of visual stimuli. Attention may act as a spotlight (Ni *et al.*, 2012), as a filter (this study), or as both (Ghose & Maunsell, 2008) depending on the experiment's conditions and the experimenter's choices when measuring its effects.

Acknowledgements

Thank you to Lyndon Duong for his help during analysis, and to Cameron Ellis for his input while revising the manuscript. This work was supported by grants to J.M.-T. from CIHR and to S.T. from the Deutsche Forschungsgemeinschaft (DFG): Research Unit 1847-A1 'Physiology of Distributed Computing Underlying Higher Brain Functions in Non-Human Primates'. The NSERC Canada Graduate Scholarship supported N.M. Correspondence concerning this article should be addressed to Dr. Julio Martinez-Trujillo, Department of Physiology & Pharmacology, Western University, London, ON, CA.

Conflict of interest

The authors declare no competing interests.

Author contributions

J.M.-T. trained and recorded from the animals for the first experiment under S.T.'s supervision, and P.K. did the same for the second experiment under J.M.-T's supervision. N.M. conducted the analyses. N.M. and J.M.-T. wrote the manuscript, with assistance from S.T. All authors revised and approved the final version of the manuscript.

Data accessibility

Neuronal data, behavioural data and their accompanying MatLab analyses scripts can be made available upon request.

Abbreviations

AIC, Akaike information criterion; FC, fixation cross; LCI, lower confidence interval; MT, middle temporal; *m*, slope; *r*, Pearson's correlation coefficient; R^2, correlation coefficient; RDP, random dot pattern; RF, receptive field; SD, standard deviation; SSI_N, spatial summation index for the attend-null condition; SSI_P, spatial summation index for the attend-tuning condition; UCI, upper confidence interval.

References

Akaike, H. (1973) Information theory and an extension of the maximum likelihood principle. In Petrov, B.N. & Csáki, F. (Eds), 2nd International Symposium on Information Theory, Tsahkadsor, Armenia, USSR, September 2–8, 1971. Akadémiai Kiadó, Budapest, pp. 267–281.

Boynton, G.M. (2009) A framework for describing the effects of attention on visual responses. *Vision. Res.*, **49**, 1129–1143.

Britten, K.H. & Heuer, H.W. (1999) Spatial summation in the receptive fields of MT neurons. *J. Neurosci.*, **19**, 5074–5084.

Ghose, G.M. & Maunsell, J.H.R. (2002) Attentional modulation in visual cortex depends on task timing. *Nature*, **419**, 616–620.

Ghose, G.M. & Maunsell, J.H.R. (2008) Spatial summation can explain the attentional modulation of neuronal responses to multiple stimuli in area V4. *J. Neurosci.*, **28**, 5115–5126.

Khayat, P.S. & Martinez-Trujillo, J.C. (2015) Effects of attention and distractor contrast on the responses of middle temporal area neurons to transient motion direction changes. *Eur. J. Neurosci.*, **41**, 1603–1613.

Khayat, P.S., Niebergall, R. & Martinez-Trujillo, J.C. (2010) Attention differentially modulates similar neuronal responses evoked by varying contrast and direction stimuli in area MT. *J. Neurosci.*, **30**, 2188–2197.

Luce, R.D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford Univ. Press, New York, NY.

Luck, S.J., Chelazzi, L., Hillyard, S.A. & Desimone, R. (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.*, **77**, 24–42.

Martinez-Trujillo, J.C. & Treue, S. (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron*, **35**, 365–370.

Martinez-Trujillo, J.C. & Treue, S. (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.*, **14**, 744–751.

Maunsell, J.H.R. & Cook, E.P. (2002) The role of attention in visual processing. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, **357**, 1063–1072.

McAdams, C.J. & Maunsell, J.H.R. (1999) Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron*, **23**, 765–773.

McAdams, C.J. & Maunsell, J.H.R. (2000) Attention to both space and feature modulates neuronal responses in macaque area V4. *J. Neurophysiol.*, **83**, 1751–1755.

Moran, J. & Desimone, R. (1985) Selective attention gates visual processing in the extrastriate cortex. *Science*, **229**, 782–784.

Ni, A.M., Ray, S. & Maunsell, J.H.R. (2012) Tuned normalisation explains the size of attention modulations. *Neuron*, **73**, 803–813.

Niebergall, R., Khayat, P.S., Treue, S. & Martinez-Trujillo, J.C. (2011) Multifocal attention filters targets from distracters within and beyond primate MT neurons' receptive field boundaries. *Neuron*, **72**, 1067–1079.

Noudoost, B., Chang, M.H., Steinmetz, N.A. & Moore, T. (2010) Top-down control of visual attention. *Curr. Opin. Neurobiol.*, **20**, 183–190.

Pernet, C.R., Wilcox, R. & Rousselet, G.A. (2012) Robust correlation analyses: false positive and power validation using a new open source Matlab toolbox. *Front. Psychol.*, **3**, 606.

Recanzone, G.H., Wurtz, R.H. & Schwarz, U. (1997) Responses of MT and MST neurons to one and two moving objects in the receptive field. *J. Neurophsiol.*, **78**, 2904–2915.

Reynolds, J.H., Chelazzi, L. & Desimone, R. (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *J. Neurosci.*, **19**, 1736–1753.

Reynolds, J.H., Pasternak, T. & Desimone, R. (2000) Attention increases sensitivity of V4 neurons. *Neuron*, **26**, 703–714.

Seidemann, E. & Newsome, W.T. (1999) Effect of spatial attention on the responses of area MT neurons. *J. Neurophysiol.*, **81**, 1783–1794.

Simoncelli, E.P. & Heeger, D.J. (1998) A model of neuronal responses in visual area MT. *Vision. Res.*, **38**, 743–761.

Snowden, R.J., Treue, S., Erickson, R.G. & Andersen, R.A. (1991) The response of area MT and V1 neurons to transparent motion. *J. Neurosci.*, **11**, 2768–2785.

Solomon, S.G. & Kohn, A. (2014) Moving sensory adaptation beyond suppressive effects in single neurons. *Curr. Biol.*, **24**, R1012–R1022.

Treue, S. & Martinez-Trujillo, J.C. (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, **399**, 575–579.

Treue, S. & Maunsell, J.H.R. (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*, **382**, 539–541.