Chapter 6
Effects of Attention in Visual Cortex: Linking Single Neuron Physiology to Visual Detection and Discrimination

Vincent P. Ferrera

6.1 Introduction

Studies of neuronal activity in visual cortex have relied heavily on macaque monkeys as a model system. Macaques, like humans, are old world primates and range throughout Asia and North Africa. The macaque genus comprises 23 species, including *Macaca mulatta* (rhesus monkey), *Macaca fascicularis* (cynomolgus or "crab-eating" monkey), and *Macaca fuscata* (Japanese snow monkey). The most recent common ancestor of humans and macaques lived roughly 25 million years ago. Macaques are largely diurnal animals that have trichromatic color vision and a retina that is anatomically almost identical to humans. In particular, the macaque retina has a distinct fovea for high-acuity central vision.

Macaques explore their visual environment in much the same way as humans. They have forward-looking eyes whose monocular visual fields are largely overlapping, providing a large binocular field with excellent stereoscopic depth perception [1]. Their oculomotor behavior is similar to humans, particularly with regard to voluntary eye movements. Macaques have *vergence* eye movements that align the foveae of the two eyes on targets at a particular distance. They make rapid and frequent *saccades* to foveate objects of interest. They can track moving targets with *smooth pursuit*, a behavior that appears to be unique to primates (at least among mammals.)

V.P. Ferrera (✉)
Department of Neuroscience, Columbia University, 1051 Riverside Drive, Unit 87, New York, NY 10032, USA
e-mail: vpf3@cumc.columbia.edu; vincent.ferrera@gmail.com

M. Mancas et al. (eds.), *From Human Attention to Computational Attention*,
Springer Series in Cognitive and Neural Systems 10,
DOI 10.1007/978-1-4939-3435-5_6

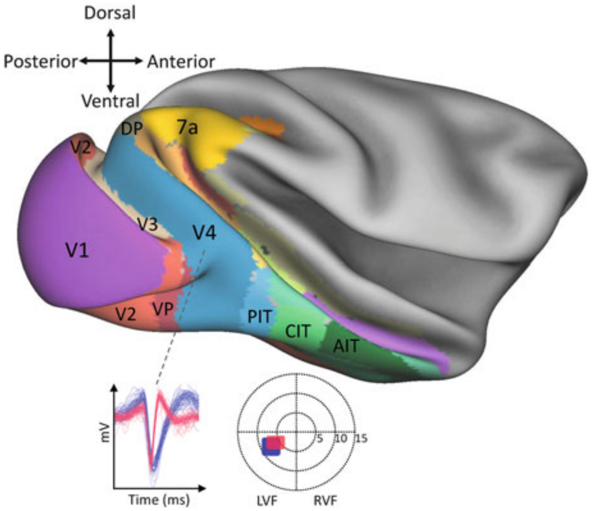

Fig. 6.1 Macaque cerebral cortex (lateral view, partially inflated) showing visual cortical areas *V1* (primary), *V2*, *V3*, *V4*, *DP* (dorsal posterior), *7a*, *VP* (ventral posterior), *PIT* (posterior inferotemporal), *CIT* (central inferotemporal), and *AIT* (anterior inferotemporal; after [2]). A microelectrode can be used to record neuronal activity at a precise cortical location (*dashed line*). Extracellular action potential waveforms for two simultaneously recorded neurons are shown in the *lower left* panel. The visual field locations and sizes of the receptive fields (*blue* and *red* squares) of the neurons are shown in the *lower right* panel. *LVF* left visual field, *RVF* right visual field

In the macaque monkey brain, there are 32 cortical areas that are involved in vision and visuomotor function [2]. For many of these areas, human homologues have been identified [3]. Macaques can be trained to perform simple tasks that involve visual detection, discrimination, and eye movements. The electrical activity of individual neurons can be recorded by fine metal microelectrodes inserted into the cerebral cortex while the animal is performing a visual task. Of the 32 visual areas in macaques, several have been studied extensively in behavioral paradigms that manipulate selective attention (Fig. 6.1). These studies have examined how attention affects receptive field properties as well as the sensitivity and reliability of neuronal responses. The current state of knowledge makes it possible to relate these neuronal response properties to psychophysical performance using simple computational models. The goal of the present chapter is to understand how attention alters the representation of information in visual cortex and thus affects an observer's ability to detect weak stimuli and to discriminate between similar stimuli.

Visual neurons are those that receive information directly or indirectly from the retina. The part of the environment that gives rise to light that falls onto the retina defines the visual field. Visual neurons typically do not respond to light that arises from anywhere in the visual field, but are sensitive to only a small region, called the *receptive field*. The receptive field for an individual neuron is the part of the retina within which changes in illumination cause changes in the electrical activity (typically, the firing rate) of the cell. If the eyes are not moving, the receptive field corresponds to a fixed region of visual space. Every cell in the visual system, from retina photoreceptors to cortical neurons, has a receptive field. The size of receptive fields generally increases along the visual hierarchy from retina to lateral geniculate to cortex and also with retina eccentricity (distance from the fovea). If a monkey is trained to fixate its gaze on a small target presented on a video display, then the borders of the receptive field can be easily mapped. This may be done by moving a spot or bar of light through the visual field and outlining the region where the stimulus causes a change in firing rate of the cell. Firing rate can be monitored qualitatively by amplifying the action potentials and playing them through an audio speaker. As long as the monkey is fixating, a particular stimulus other stimuli presented in the visual field will have a known spatial relationship with respect to the receptive field of a given neuron. Controlling the retina stimulus in this manner makes it possible to study the influence of *extraretinal* factors, such as attention, on the activity of visual neurons.

The receptive field of a neuron can be modeled mathematically as a spatial weighting function, which specifies the neuron's firing rate as a function of the retina position of a small spot of light. A visual neuron's sensitivity to light within the receptive field is not necessarily uniform, but may have subregions that are excited or inhibited by light. The spatiotemporal structure of the receptive field may confer selectivity for orientation and direction of motion. Different parts of the receptive field may also be sensitive to different wavelengths of light, giving rise to color selectivity. For current purposes, we will ignore the internal structure of visual receptive fields and simply model sensitivity within the spatial receptive field (RF) as a two-dimensional Gaussian:

$$\mathrm{RF}\,(x,y) = A + B \times \exp\left[-\left(\left(x - x'\right)^2 + C \times \left(y - y'\right)^2\right)/s^2\right] \qquad (6.1)$$

Here, (x', y') is the center of the receptive field, s is the spread or size of the RF (otherwise known as the *space constant*), B is the overall gain or sensitivity, C determines the aspect ratio (length/width), and A is a constant that accounts for the baseline firing of the cell in the absence of a stimulus. Many studies of the effects of attention on the activity of visual neurons have examined changes in spatial parameters that correspond to shrinking or expanding of the receptive field. These are modeled as changes in the space constant, s. Other studies have documented shifts of the RF center (x', y') and changes in overall sensitivity (B) and background firing (A).

To understand how attention-related changes in receptive field properties affect stimulus detectability and discriminability, it is necessary to consider the statistics of neuronal responses, i.e., the variability in neuronal firing when the same stimulus is presented repeatedly under the same conditions.

To a first approximation, cortical neurons fire at purely random times. Their firing can be modeled as a Poisson process where the probability of an action potential at any given time is determined by a rate parameter, r, and is independent of the time of occurrence of any other action potentials. The interspike intervals (times between two successive action potentials) follow a Poisson distribution. The number of action potentials in a fixed time window (spike count) is also Poisson distributed. Spike count variability can be quantified by the Fano factor [4], which is the variance in spike count divided by its mean. For a Poisson process, the Fano factor is always around 1.0 as the variance scales in direct proportion to the mean spike count.

Poisson firing statistics represent an ideal case that is never achieved in reality. In particular, a purely Poisson neuron could have infinitely small interspike intervals, which are biophysically impossible. Real neurons have refractory periods – a short window of time following a spike during which the cell is unable to fire another spike (absolute refractory period) or has an elevated threshold for firing (relative refractory period). Refractory periods are easy to incorporate into simulations that generate pseudo-Poisson spike trains using random number generators [5]. Refractory periods cause neuronal firing to become more regular (lower variance in interspike intervals). Any finite refractory period therefore reduces the Fano factor below 1.0. A number of studies have documented sub-Poisson variability in macaque visual cortex and in higher, attention-related cortical areas [6–10].

Figure 6.2 shows simulated Poisson-like spike trains generated by an algorithm that incorporates an absolute refractory period. In the top-left panel are spike trains where the refractory period is equal to 0 and below that the spike count histogram and Fano factor versus mean spike count. The right column shows spike trains with the same average rate, but a longer refractory period, making both the interspike intervals and spike counts much more regular.

It has been found empirically that attention can reduce neuronal variability [11], but the reduction is small and not always statistically significant [12]. It seems intuitive that reduced variability should improve the ability to detect and discriminate stimuli. One of the goals of the models presented below is to test whether this is indeed the case.

A simple model of the response of an individual visual neuron can be obtained by using Eq. 6.1 to provide the input to a Poisson spike generating process. This is illustrated in Fig. 6.3 which shows the mean rate according to a one-dimensional reduction of Eq. 6.1 (Fig. 6.3, top) and the Poisson spike counts (Fig. 6.3, bottom) generated when a stimulus is present ($A = 5$, $B = 10$) or absent ($A = 5$, $B = 0$). Detectability can be computed for each stimulus position as the overlap (area under ROC) of the stimulus-present and stimulus-absent spike count distributions. Note that this is not a complete neuronal model as it does not include contrast nonlinearities, adaptation, or other factors that affect firing. Real visual neurons tend to have sigmoidal contrast response functions, and their contrast sensitivity may be modulated by attention [13, 14]. However, the current model is adequate for testing effects of changes in sensitivity or variability for briefly presented stimuli of fixed contrast.

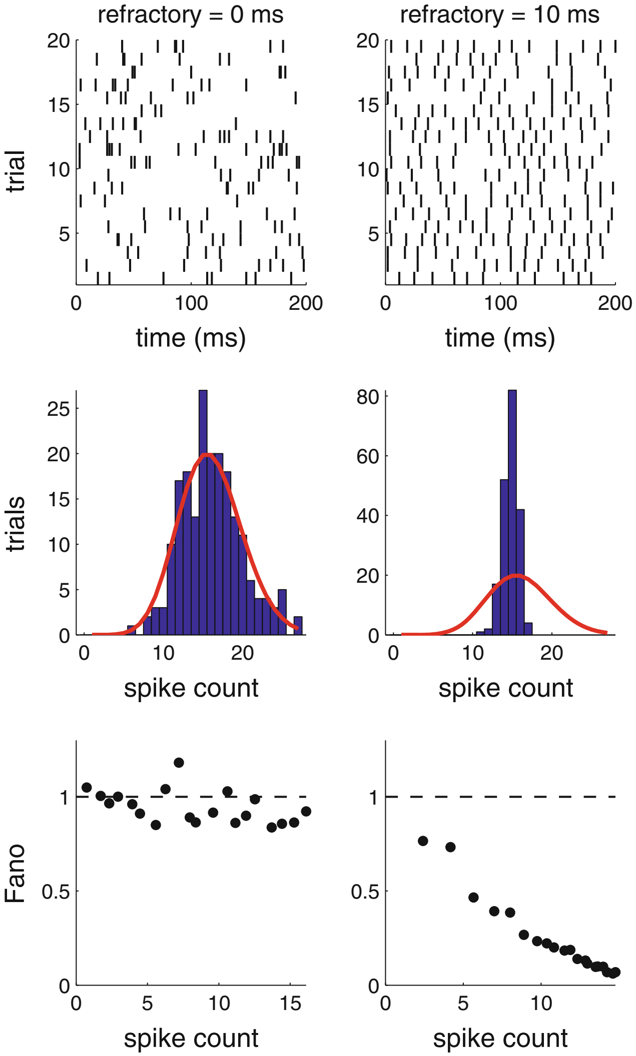

Fig. 6.2 Neuronal firing statistics. (*Left*) Pure Poisson process. (*Right*) Poisson with refractory period. *Top row* shows 20 spike trains for each model. *Middle row* shows spike count distributions for several hundred trials. *Bottom row* shows Fano factor versus spike count

6.2 Effects of Attention on Neuronal Responses

Moran and Desimone [15] published one of the first studies of the effect of attention on neurons in macaque visual cortex. They trained monkeys to fixate a small spot presented in the center of a video display. Eye movements were monitored so that visual stimuli could be presented at known positions on the retina. While

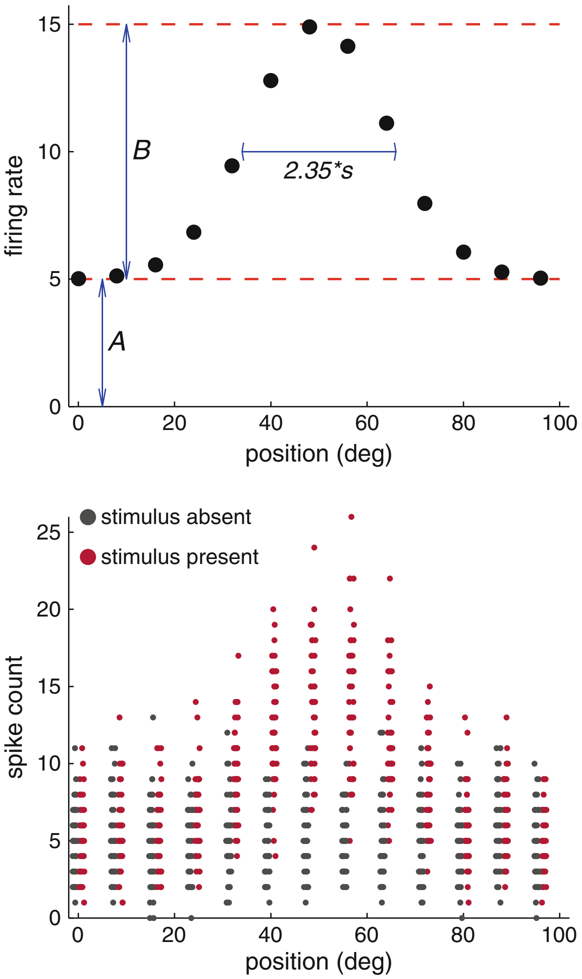

Fig. 6.3 Minimal model of a single neuron. (*Top*) Mean firing rate as a function of position. (*Bottom*) Distributions of spike counts for stimulus-absent and stimulus-present conditions as a function of stimulus position in receptive field

the monkeys fixated, two stimuli were presented, and the monkeys were rewarded for responding to one stimulus. The monkeys are presumed to have attended to rewarded stimulus and to have ignored the other.

Moran and Desimone recorded from neurons in visual area V4 and in the inferior temporal (IT) cortex. Neural responses were quantified as changes in firing rate (action potentials per second), while visual stimuli were presented to the animal. The receptive fields of the neurons were in the peripheral visual field and were large enough that two stimuli could be presented inside the receptive field and the monkey could still discriminate them. If both stimuli were in the receptive field of the neuron, the cell responded well to the attended stimulus, but weakly to the unattended stimulus. The experimenters could therefore compare the response to the same stimulus when it was attended or unattended. Generally, the response to the

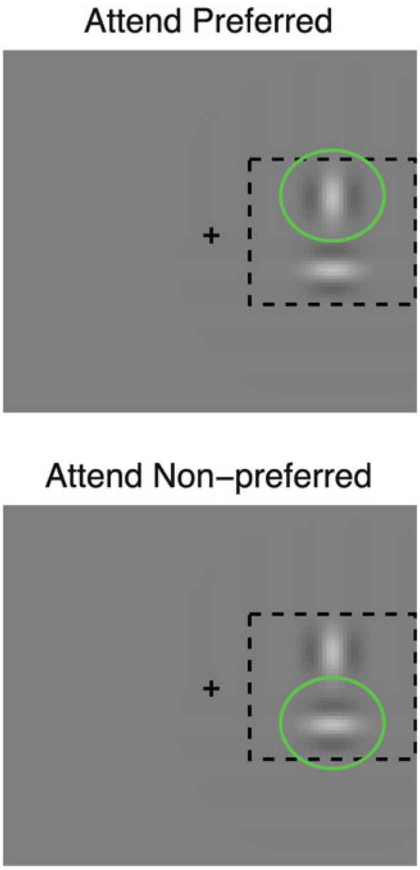

Fig. 6.4 Attention task used by Reynolds et al. [16]. Monkeys were trained to fixate their gaze at the center of the display (+) while the activity of a visual neuron was recorded. The receptive field of the neuron is indicated by the *dashed box*. Two stimuli were presented inside the receptive field, and the monkey was rewarded for responding to one or the other. The attended stimulus is indicated by the *green circle* (this cue was not presented to the animal)

stimulus was greater when it was attended. If one stimulus was inside the receptive field and the other was outside, the effect of attention was reduced as compared to when both were inside the receptive field.

A later study by Reynolds et al. [16] expanded on this result. Reynolds' study used the strategy of placing two oriented bar stimuli in the receptive field of a V4 neuron (Fig. 6.4). Neurons in V4 tend to be selective for stimulus orientation. The orientation of one of the bars was matched to the cell's preferred orientation and evoked a strong response. The other stimulus was at a non-preferred orientation. In the absence of attention, the neuronal response when both stimuli were presented together was the average of the response to either stimulus alone. When the monkey was rewarded for attending to one stimulus or the other, the cell behaved as if there was only one stimulus in the receptive field; if the monkey attended the preferred stimulus, the response was greater than the average; if he attended the non-preferred stimulus, the response was less than the average. Thus, attention caused the cells to shift from a response-averaging mode to a winner-take-all mode. These results are consistent with a shrinking of the receptive field around the attended stimulus

[17, 18]. Since Moran and Desimone's [15] paper, a large number of studies have demonstrated changes in the receptive field spatial weighting function that are correlated with attention.

Attention can change the overall gain of visual responses. This was demonstrated for neurons in visual area V4 by McAdams and Maunsell [12, 19]. This study examined orientation tuning. Orientation selectivity is reasonably described by Gaussian-shaped tuning function. McAdams and Maunsell placed an oriented grating pattern in the receptive field of a V4 neuron (Fig. 6.5). They then recorded responses to stimuli of various orientations and compared the orientation-tuned responses when attention was directed toward the stimulus inside the RF or to a similar stimulus well outside the receptive field. They modeled V4 responses using

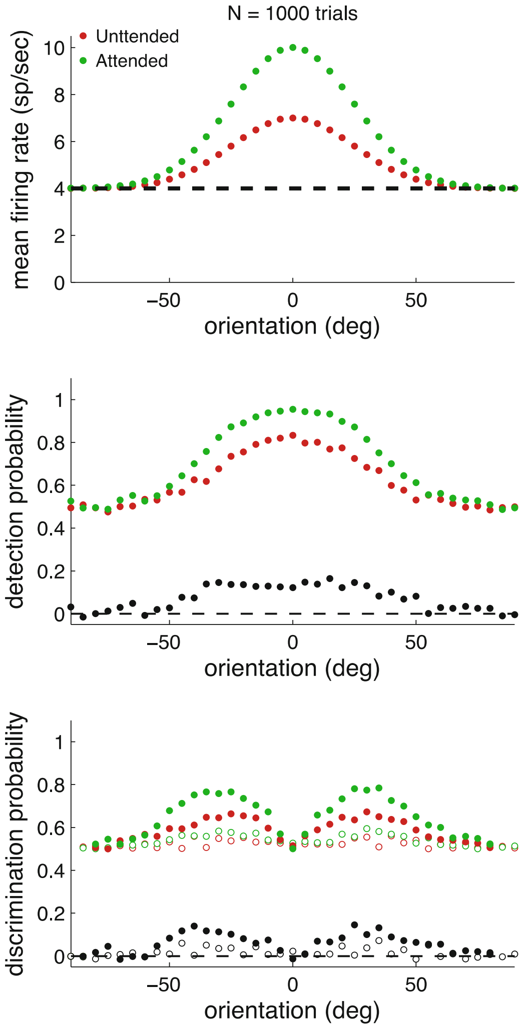

Fig. 6.5 Effects of attention on response of a single neuron. (*Top*) Mean rates when attention is directed outside the neuron's receptive field (*red*) and inside the RF (*green*). (*Middle*) Effect of attention on detectability as a function of stimulus position. *Black dots* are difference attended (*green*) – unattended (*red*). (*Bottom*) Effect of attention on stimulus discriminability

an equation similar to Eq. 6.1 and concluded that attention mainly affects the overall gain (B). Importantly, they also measured activity in the absence of a stimulus and found that attention did not affect the baseline (undriven) firing rate (A).

These few studies provide enough information to simulate the effects of attention in the model introduced previously. In this model, the receptive field equation (Eq. 6.1) is used to determine the mean firing rate for a small spot of light presented at any position in the receptive field. This mean rate is then fed into a function that generates a pseudo-Poisson spike train for a fixed time interval (1.0 s). This spike train can be purely Poisson (refractory period = 0), or can have a finite refractory period, resulting in sub-Poisson variability. For each stimulus, a large number of spike trains are generated and the total spike count for each train is used as the measure of neuronal response. One can then use principles of signal detection theory to determine the ability of one or more simulated neurons to detect or discriminate visual stimuli, given the trial-to-trial variability in the neuronal responses. Figure 6.5 (top) shows the mean response of a single visual neuron (reduced to one dimension) as a function of stimulus location. The effect of attention is modeled as an overall gain factor, G, applied to the stimulus-driven response, so that

$$\mathrm{RF}\,(\mathrm{x}) = \mathrm{G} \times \left\{\mathrm{A} + \mathrm{B} \times \exp\left[-\left(\left(\mathrm{x} - \mathrm{x}'\right)^2\right)/\mathrm{s}^2\right]\right\} \tag{6.2}$$

The responses in Fig. 6.5 show the cases where G = 1.0 (red, attention outside receptive field) and G = 2.0 (green, attention inside RF). In the absence of a stimulus, the response is simply RF(x) = G*A, where A is the baseline firing rate.

Detectability and discriminability are computed by applying signal detection theory to the spike count distributions for each stimulus. Detectability is defined as the area under the ROC curve computed with stimulus-present and stimulus-absent trials. The effect of attention on detectability is shown in Fig. 6.5 (middle). Even though attention increases the driven firing rate by twofold, the maximum change in detection probability is only 0.1. It should be noted that an attentional gain of 2.0 is unusual. Typically, attention enhances neuronal responses by increasing mean firing rate from 20 % to 40 %. For many cells, attention actually reduces responses.

In the simulation shown in Fig. 6.5, the baseline firing rate in the unattended condition was 4 spikes/sec, and the maximum firing rate was 10 spikes/sec. This value for maximum firing rate is on the low end of the range for cortical neurons. Values of 30 spikes/sec or greater are more typical for responses to optimal stimuli. Hence, the low ratio of max firing rate to baseline can be thought of as representing the response to suboptimal or weak stimuli. Detection probability in the unattended condition starts to saturate at 1.0 (perfect performance) when the maximum firing rate is about 3 times the baseline rate. Attention cannot improve performance when detection rates in the unattended condition are already optimal. Thus, attention should have the greatest effect on detectability for weak or suboptimal stimuli or cells that simply have low signal to noise even for optimal stimuli.

Stimulus discriminability is defined as the ROC area computed for pairs of similar stimuli. In these simulations, the difference between neighboring orientations was 5°. For the parameters used in the simulations of Fig. 6.5, attention had little effect on discriminability for neighboring orientations (open symbols). However, for cruder discriminations (stimuli separated by 20°.), discrimination performance was better overall (Fig. 6.5, bottom, filled symbols) and was more strongly enhanced by attention. Further simulations showed that as the maximum firing rate is increased, attention had a greater effect on improving discrimination performance. For example, at a maximum firing rate of 50 spikes/sec (keeping all other parameters the same), the best discrimination for neighboring orientations improved from 74 % correct to 82 %. Hence, while attention improves detection performance mainly for neurons with low signal to noise, it improves discrimination for neurons with high signal strength. This suggests that different tasks might reveal attention effects on different subpopulation of neurons.

The effects of attention on psychophysical performance predicted by the model are fairly modest. Using realistic parameters, attention improves detection and discrimination rates by a maximum of about 10 %. Such small changes in performance are far below what is typically reported in the literature. For example, [20] found that attention could produce up to fourfold improvements in contrast sensitivity. We will consider two factors that could bridge this gap. The first is the effect of attention on undriven (stimulus-absent) activity. The second is the effect of attention on spike count variability.

So far, we have assumed that attention affects firing rates proportionately for both stimulus-present (driven activity) and stimulus-absent (undriven or baseline activity) conditions. This point is disputed. Some studies report that attention affects background firing rates [21], while others [12, 19] reported that attention did not affect undriven activity. When undriven activity is held constant in the model, so that attention enhances activity only in the presence of a stimulus, the affects of attention on detection are greatly increased. Repeating the simulations of Fig. 6.5 with a constant baseline, the improvement in detection probability goes from 0.1 to 0.3. This is a large enough improvement to account for actual psychophysical performance. Thus, the issue of whether attention affects baseline activity is critical for understanding improvements in detection performance. However, in the model, baseline firing rate plays no role in discrimination performance.

Now we can address the issue of attention-related changes in spike count variability. As noted above, some studies have reported that attention can reduce trial-to-trial variability in firing activity [11]. Here, we reduce spike count variability by introducing a refractory period. There is a caveat to this approach: for any two spike trains with the same underlying rate, the one with the longer refractory period will have a lower spike count. Thus, it is important to equalize spike count when assessing the effects of regularity. Figure 6.6 (top) shows Fano factor as a function of mean spike count for spike trains with no refractory period (red) and with a refractory period of 10 ms (green). Note that refractoriness causes Fano factor to decrease with mean spike count, being reduced by about half for the highest firing rate.

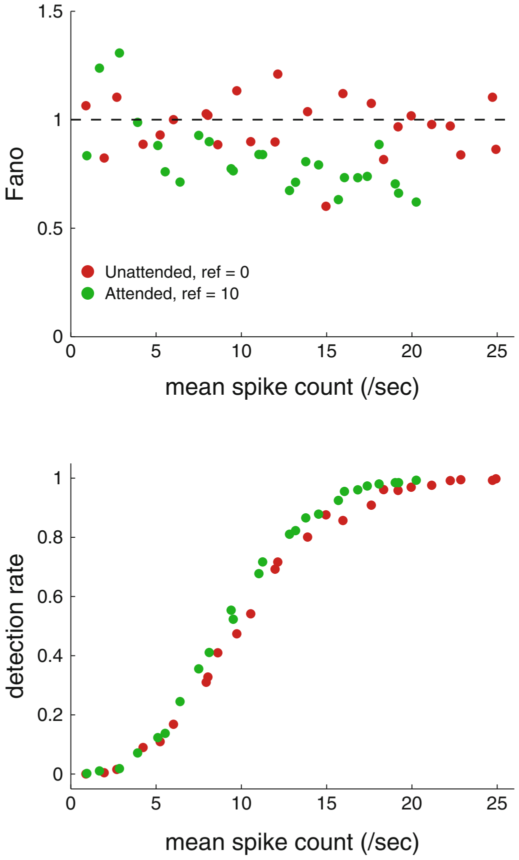

Fig. 6.6 Effects of attention modeled as changes in trial-to-trial spike count variability. (*Top*) Fano factor as a function of mean spike count for refractory period = 0 (*red*) or 10 (*green*). (*Bottom*) Detection rate as a function of mean spike count. Same convention as *top panel*

The effect of spiking regularity on detection rate is shown in Fig. 6.6 (bottom). Again, the red dots are for spike trains with zero refractory period; the green are for a refractory period of 10 ms. The same refractory period was used for both stimulus-present and stimulus-absent conditions, although there is some evidence that stimulus onset itself is accompanied by a reduction in spike count variability [22]. What is evident from Fig. 6.6 is that a reduction in variability improves detection rates, but only by about 5–10 %. The improvement is greatest when the signal to noise is relatively weak, such that the maximum firing rate is about twice the baseline firing. When the maximum firing rate increases beyond this, detection rates saturate and spike count regularity has no effect. The simulations were carried out with a baseline of 10 spikes/sec. Changing the baseline firing rate shifts the curves left and right, but the same principles apply.

While spike count regularity alone results in some enhancement of detectability, it has a smaller effect on discrimination performance. Figure 6.7 shows simulations of a neuron whose receptive field is modeled as a one-dimensional Gaussian function of position, with preferred position at 50°. The left panel shows Fano

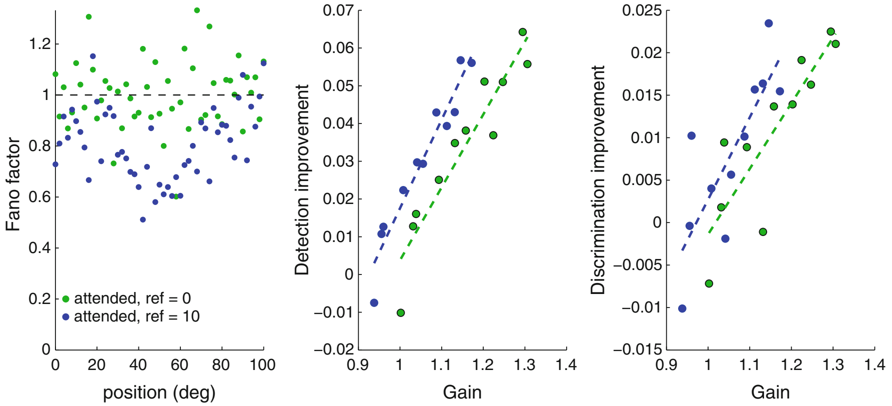

Fig. 6.7 Effects of spike count variability on discrimination performance. (*Left*) Fano factor as a function of stimulus position for unattended (*green*, refractory period = 0) and attended (*green*, refractory period = 1). (*Middle*) Effects of variability on detection. Gain is attended response/unattended response. Detection improvement is the difference in detection rate for each position (attended – unattended) averaged over all positions. *Dashed lines* are best fit linear regressions. (*Right*) Effects of variability on detection. Same conventions as *middle* panel

factors as a function of stimulus position in the receptive field for no refractory period (green) or a refractory period of 10 ms (blue). There is a substantial, stimulus position-dependent decrease in Fano factor. However, the decrease in variability is accompanied by a proportional decrease in mean spike count due to the refractoriness of the cell. In other words, the overall response of the cell is scaled, including the mean spike count and variance. The decrease in variability leads to an improvement in detectability of a few percent. Detection rate is calculated by computing the ROC for stimulus-present versus stimulus-absent conditions and assuming that attention does not affect baseline firing (either average rate or variability) in the absence of the stimulus. Discriminability is based on the spike count distributions for neighboring stimuli. Since refractoriness scales both distributions proportionately, firing regularity only has a small effect on ROC area. When the attention-related improvement in discrimination performance is plotted as a function of actual gain (Fig. 6.7, right), the improvement in performance is quite small (green, refractory period = 0; blue, refractory period = 10 ms).

To summarize, attention can affect the gain of neuronal responses as well as their reliability. Large changes in response gain lead to only modest improvements in detection and discrimination rates. If baseline activity is unaffected by the gain change, then much larger increases in detection rates are achievable, but there is no effect on discrimination. Improving reliability by incorporating a refractory period into the spike train generator has a small effect on detection and an even smaller effect on discrimination. One caveat is that refractoriness always reduces both the variance and mean of the spike counts. Other methods that reduce variability without changing mean rate were not explored.

6.3 Effects of Attention Across Multiple Neurons

When considering the effects of attention across multiple neurons, there is a general expectation that such effects will be stronger and/or more reliable. This expectation may be frustrated for several reasons. Having more neurons can improve signal processing, but it also means that there will be more noise due to random firing from neurons that are not sensitive to the stimulus. Indeed the problem of selective attention is not only one of selecting the most relevant stimulus, but, perhaps more importantly, selecting the most relevant neurons.

To model the effects of attention across multiple neurons, consider an array of neurons that are identical except for the location of their receptive field centers. Instead of the scalar attentional gain factor in the single neuron model described above, attention is modeled as a gain field [G(x), [18]] that ranges over the entire visual field:

$$\mathrm{G}(\mathrm{x}) = 1.0 + \mathrm{a} \times \exp\left(-\left(\mathrm{x}-\mathrm{x}'\right)^2/\mathrm{s}^2\right) \tag{6.3}$$

where a is the attentional enhancement, x is visual field location, x′ is the focus of attention, and s is the spread of attention. Figure 6.8 shows the effect of attentional gain enhancement on an array of neurons that have identical tuning width and sensitivity, but different receptive field centers. Tuning curves for the unattended case are shown in the left. On the right are tuning curves with a maximum attentional gain of 2.0. Attention does not enhance the activity of cells whose preferred locations are remote from the focus of attention. It should be noted that the attentional gain field not only enhances the response of cells at the focus of attention, but it also distorts the tuning functions. Such shifting of receptive fields has been documented in visual areas V4 [23] and MT [18].

In the case of a single neuron, it was found that if attention enhanced both undriven and stimulus-driven activity, there was little improvement in detection or discrimination performance. Here, we test what happens if attention does not affect baseline firing in either the stimulus-absent or stimulus-present conditions. When we simulate this condition, it turns out that attention has little effect on detection (Fig. 6.8, middle left) and no effect on discrimination (Fig. 6.8, middle right). This result holds over a wide range of signal strengths (maximum firing rate re: baseline). To obtain even a small increase in detectability requires an attention gain of about 4x or greater. Discriminability does not improve for any gain level. There was no effect of refractoriness on detection or discrimination rates.

The finding that attentional gain has little effect on detection may seem counter-intuitive. However, it makes perfect sense. There are nine neurons in the simulation, and, as can be seen in Fig. 6.8, attention only affects 3 of them. For any given stimulus, most of the cells do not respond at all. Yet, all of the cells must be included when computing detectability, even if they are unmodulated by attention, or not even driven by the stimulus. The reason for this is that the stimulus has an equal probability of occurring at any location and this location is not known in advance.

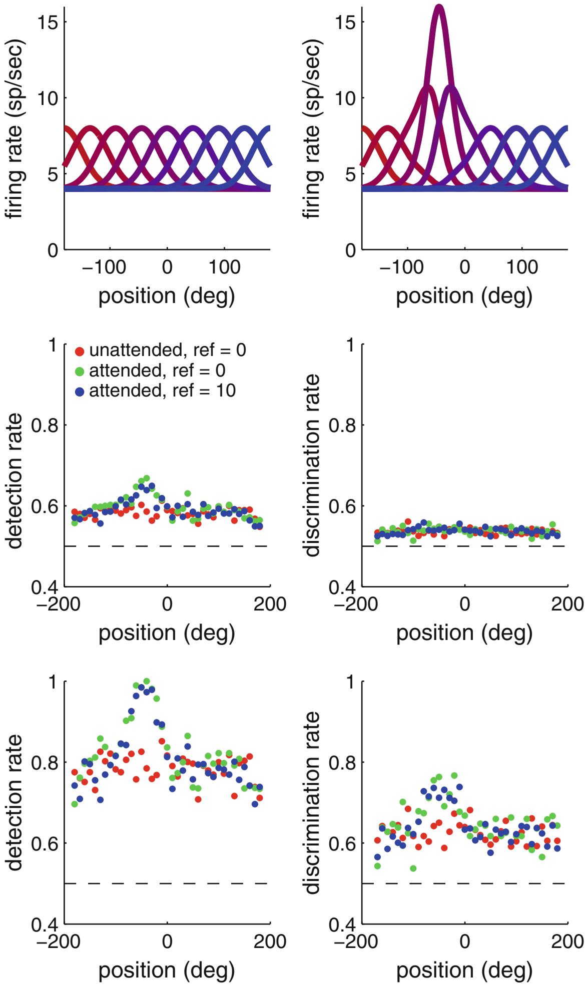

Fig. 6.8 (*Top*) Tuning functions for an array of model neurons in the unattended condition (*left*) and with attention focused at −50°. (*Middle*) detection and discrimination rates. (*Bottom*) Detection and discrimination rates when responses are pooled across neurons

Thus, at any given time, most of the cells are simply contributing noise. This not only dilutes the effect of attention; it can negate the effect altogether.

However, we have yet to consider the issue of pooling activity across neurons. By this, we mean how signals from different neurons are combined when computing the joint ROC. None of the multi-neuron simulations discussed above included any pooling; each response was considered as an independent observation and was weighted equally in the ROC analysis. One way to pool responses is to compute the mean spike count across all neurons in the model on each trial. Thus, the data are reduced from 3 dimensions (neuron × stimulus × trial) to only 2 (stimulus × trial). This averaging is done before the ROC area is computed. The effect of this kind of pooling is that the neurons that are sensitive to the stimulus tend to pull up the average response of the ensemble. On the other hand, when there is no stimulus, averaging across neurons has little effect because they all have the same baseline activity.

We thus consider a model in which attention has no effect on baseline activity, but activity on every trial is pooled by averaging across all neurons. The results are shown in the bottom row of Fig. 6.8. For this model, attention enhances both detection (Fig. 6.8 bottom left) and discrimination (Fig. 6.8, bottom right). As in previous simulations, reducing trial-to-trial variability has no effect (compare green dots, refractory period = 0, to blue, refractory period = 10 ms).

Averaging over all neurons is an extreme form of pooling that is not physiologically or anatomically plausible. It requires that all of the sensory neurons converge onto a single decision neuron. However, one can imagine a pooling function that computes a weighted average of responses over a limited spatial extent so that only cells with similar receptive field locations are combined. This agrees well with how the visual cortex is wired and the fact that receptive fields get larger as one traverses the cortical hierarchy from primary visual cortex (V1) to V2, V3, V4, and IT.

To appreciate how attention affects the representation of information in visual cortex, we can use some of the aforementioned ideas to construct "neural" images of simple stimuli. Figure 6.9 shows simulations of a 2D array of model neurons. The input image consists of two vertically oriented Gabor patterns embedded in random noise. The green circle (Fig. 6.9, left) indicates the focus of attention, but was not present in the image used for the simulations. Each model neuron comprised a Gaussian spatial weighting function that represented the neuron's receptive field. Each receptive field was approximately 1/20th the size of the image in linear dimension. There were approximately 200 × 200 neurons whose RF centers were distributed to cover the entire image. The response of each neuron was computed by calculating the inner product of the weighting function and the part of the image within the receptive field. This number was used as the rate parameter for a Poisson spike generation function. Each pixel in Fig. 6.9 (middle and right) represents the resulting spike count for a single neuron. The middle panel of Fig. 6.9 illustrates a condition where attention increased the gain of the response at the attended location. The right panel shows a condition where the gain was constant across the image,

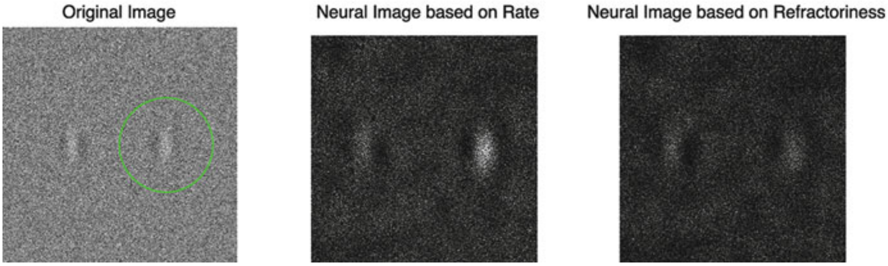

Fig. 6.9 Neural images created by computing the responses of a 2D array of model neurons. (*Left*) Original stimulus. The *green circle* indicates the focus of attention and was not present in the image used for model simulations. (*Middle*) A 2D array of model neurons. Attention increases the gain of the response in the attended region. Pixel intensity represents firing rate. (*Right*) A 2D array with a constant response gain across location, but increased refractoriness at the attended location

but the refractoriness of the cells was increased at the attended location. The result of increasing refractoriness is that there is less variability across cells that have the same input. The simulations suggest that increasing the gain has a pronounced effect on salience, whereas reducing variability through refractoriness has little effect. These neural images can be converted to detectability maps by running multiple trials with and without the stimulus and computing ROC functions for each neuron.

The simulations in this chapter have explored attentional gain control and reliability and how these affect detection and discrimination performance. Some features of the model that turned out to be important are (1) that attention enhances stimulus-driven responses but not baseline activity and (2) that responses are pooled over multiple neurons. Pooling of responses across neurons reduces variability and can have a pronounced effect on performance. One feature that was of only modest importance was trial-to-trial spike count variability; when variability is reduced by refractoriness there is little effect on detection or discrimination performance. Relatively, few empirical studies have investigated effects of attention on neural detection and discrimination thresholds [12, 13, 24, 25]. Fewer still have related changes in neural responses to behavioral thresholds [26]. This is an area that warrants further investigation and can profit from approaches that combine computational modeling and neurophysiological experimentation.

Acknowledgements This work was supported by NIH grant MH059244. Dr. Gaurav Patel assisted with the preparation of Fig. 6.1.

Literature Cited and Further Reading

1. DeAngelis, G. C., Cumming, B. G., & Newsome, W. T. (1998). Cortical area MT and the perception of stereoscopic depth. *Nature, 394*(6694), 677–680.
2. Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex, 1*(1), 1–47.
3. Van Essen, D. C., Lewis, J. W., Drury, H. A., Hadjikhani, N., Tootell, R. B., Bakircioglu, M., et al. (2001). Mapping visual cortex in monkeys and humans using surface-based atlases. *Vision Research, 41*(10–11), 1359–1378.
4. Fano, U. (1947). Ionization yield of radiations. II. The fluctuations of the number of ions. *Physical Review, 72*(1), 26.
5. Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge: MIT Press.
6. Bair, W., & Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving monkey. *Neural Computation, 8*(6), 1185–1202.
7. Huang, X., & Lisberger, S. G. (2013). Circuit mechanism revealed by spike-timing correlations in macaque area MT. *Journal of Neurophysiology, 109*(3), 851–866.
8. Maimon, G., & Assad, J. A. (2009). Beyond poisson: Increased spike-time regularity across primate parietal cortex. *Neuron, 62*(3), 426–440.
9. Osborne, L. C., Bialek, W., & Lisberger, S. G. (2004). Time course of information about motion direction in visual area MT of macaque monkeys. *The Journal of Neuroscience, 24*(13), 3210–3222.
10. Ferrera, V. P. (2015). Smooth pursuit preparation modulates neuronal responses in visual areas MT and MST. *Journal of Neurophysiology, 114*(1), 638–649.

11. Mitchell, J. F., Sundberg, K. A., & Reynolds, J. H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron, 55*(1), 131–141.
12. McAdams, C. J., & Maunsell, J. H. R. (1999). Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron, 23*(4), 765–773.
13. Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron, 26*, 703–714.
14. Williford, T., & Maunsell, J. H. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology, 96*(1), 40–54.
15. Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science, 229*(4715), 782–784.
16. Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience, 19*(5), 1736–1753.
17. Roberts, M., Delicato, L. S., Herrero, J., Gieselmann, M. A., & Thiele, A. (2007). Attention alters spatial integration in macaque V1 in an eccentricity-dependent manner. *Nature Neuroscience, 10*(11), 1483–1491.
18. Womelsdorf, T., Anton-Erxleben, K., & Treue, S. (2008). Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *Journal of Neuroscience, 28*(36), 8934–8944.
19. McAdams, C. J., & Maunsell, J. H. R. (1999). Effects of attention on orientation-tuning functions of single neurons in Macaque cortical area V4. *Journal of Neuroscience, 19*(1), 431–441.
20. Zenger, B., Braun, J., & Koch, C. (2000). Attentional effects on contrast detection in the presence of surround masks. *Vision Research, 40*, 3717–3724.
21. 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. *Journal of Neurophysiology, 77,* 24–42.
22. Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., Priebe, N. J., Finn, I. M., Ferster, D., Ryu, S. I., Santhanam, G., Sahani, M., Shenoy, K. V. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience, 13*(3), 369–378.
23. Connor, C. E., Preddie, D. C., Gallant, J. L., & Van Essen, D. C. (1997). Spatial attention effects in macaque area V4. *Journal of Neuroscience, 17*(9), 3201–3214.
24. Boudreau, C. E., Williford, T. H., & Maunsell, J. H. (2006). Effects of task difficulty and target likelihood in area V4 of macaque monkeys. *Journal of Neurophysiology, 96*(5), 2377–2387.
25. Moore, T., & Chang, M. H. (2009). Presaccadic discrimination of receptive field stimuli by area V4 neurons. *Vision Research, 49*(10), 1227–1232.
26. Cook, E. P., & Maunsell, J. H. R. (2002). Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of Macaque Monkey. *Journal of Neuroscience, 22*(5), 1994–2004.
27. Bushnell, C., Goldberg, M. E., & Robinson, D. L. (1981). Behavioral enhancement of visual responses in monkey cerebral cortex. I. Modulation in posterior parietal cortex related to selective visual attention. *Journal of Neurophysiology, 46*, 755–772.
28. Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience, 7*, 308–313.
29. Connor, C. E., Gallant, J. L., Preddie, D. C., & Van Essen, D. C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. *Journal of Neurophysiology, 75*, 1306–1309.
30. Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate cortex. *Philosophical Transactions of the Royal Society of London, Series B, 353*, 1245–1255.
31. Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18*, 193–222.

32. Fries, P., Reynolds, J. H., Rorle, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective attention. *Science, 291*, 1560–1563.
33. Martinez-Trujillo, J. C., & Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron, 35*, 365–370.
34. Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology, 14*, 744–751.
35. McAdams, C. J., & Maunsell, J. H. R. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *Journal of Neurophysiology, 83*, 1751–1755.
36. Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology, 70*(3), 909–919.
37. Motter, B. C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of Neuroscience, 14*(4), 2178–2189.
38. Reynolds, J. H., & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience, 27*, 611–647.
39. Roelfsema, P. R., Lamme, V. A. F., & Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature, 395*, 376–381.
40. Seidemann, E., & Newsome, W. T. (1999). Effect of spatial attention on the responses of area MT. *Journal of Neurophysiology, 81*, 1783–1794.
41. Skottun, B. C., De Valois, R. L., Grosof, D. H., Movshon, J. A., Albrecht, D. G., & Bonds, A. B. (1991). Classifying simple and complex cells on the basis of response modulation. *Vision Research, 31*(7–8), 1079–1086.
42. Spitzer, H., Desimone, R., & Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science, 240*(4850), 338–340.
43. Treue, S., & Maunsell, J. H. R. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature, 382*(6591), 539–541.
44. Treue, S., & Martinez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399*(6736), 575–579.
45. Treue, S., & Martinez Trujillo, J. C. (1999). Reshaping neuronal representations of visual scenes through attention. *Current Psychology of Cognition, 18*(5–6), 951–972.
46. Treue, S., & Maunsell, J. H. R. (1999). Effects of attention on the processing of motion in macaque visual cortical areas MT and MST. *Journal of Neuroscience, 19*(17), 7603–7616.
47. Treue, S. (2001). Neural correlates of attention in primate visual cortex. *TINS, 24*(5), 295–300.