

Gain Modulation as a Mechanism for Switching Reference Frames, Tasks, and Targets

Emilio Salinas and Nicholas M. Bentley

Abstract In the mammalian brain, gain modulation is a ubiquitous mechanism for integrating information from various sources. When a parameter modulates the gain of a neuron, the cell's overall response amplitude changes, but the relative effectiveness with which different stimuli are able to excite the cell does not. Thus, modulating the gain of a neuron is akin to turning up or down its "loudness". A well-known example is that of visually sensitive neurons in parietal cortex, which are gain-modulated by proprioceptive signals such as eye and head position. Theoretical work has shown that, in a network, even relatively weak modulation by a parameter P has an effect that is functionally equivalent to turning on and off different subsets of neurons as a function of P . Equipped with this capacity to switch, a neural circuit can change its functional connectivity very quickly. Gain modulation thus allows an organism to respond in multiple ways to a given stimulus, so it serves as a basis for flexible, nonreflexive behavior. Here we discuss a variety of tasks, and their corresponding neural circuits, in which such flexibility is paramount and where gain modulation could play a key role.

The Problem of Behavioral Flexibility

The appropriate response to a stimulus depends heavily on the context in which the stimulus appears. For example, when James Brown yells "Help me!" in the middle of a song, it calls for a different response than when someone yells the same thing from inside a burning building. Normally, the two situations are not confused: concert goers don't rush onstage to try to save James Brown and firemen don't take to dancing outside of burning buildings. This is an extreme example, but to a certain degree, even the most basic sensory stimuli are subject to similar interpretations that depend on the ongoing behavioral context. For instance, the sight and smell

E. Salinas (✉)

Department of Neurobiology and Anatomy, Wake Forest University School of Medicine,
Winston-Salem, NC 27157, USA

e-mail: esalinas@wfubmc.edu

of food may elicit very different responses depending on how hungry an animal is. Such context sensitivity is tremendously advantageous, but what are the neuronal mechanisms behind it?

Gain modulation is a mechanism whereby neurons integrate information from various sources [47, 48]. With it, contextual and sensory signals may be combined in such a way that the same sensory information can lead to different behavioral outcomes. In essence, gain modulation is the neural correlate of a switch that can turn on or off different subnetworks in a circuit, and thus provides, at least in part, a potential solution to the problem of generating flexible, context-sensitive behavior illustrated in the previous paragraph. This chapter will address three main questions: What is gain modulation? What are some of its experimental manifestations? And, what computational operations does it enable? We will show that relatively weak modulatory influences like those observed experimentally may effectively change the functional connectivity of a network in a drastic way.

What is Gain Modulation?

When a parameter modulates the gain of a neuron, the cell's response amplitude changes, but its selectivity with respect to other parameters does not. Let's elaborate on this. To describe gain modulation, we must consider the response of a neuron as a function of two quantities; call them x and y . To take a concrete example, suppose that x is the location of a visual stimulus with respect to the fixation point and y is the position of the eye along the horizontal axis. The schematic in Fig. 1 shows the relationship between these variables. First, fix the value of y , say it is equal to 10° to the left, and measure the response of a visual neuron as a function of stimulus location, x . Assume that the response to a stimulus is quantified by the number of action potentials evoked by that stimulus within a fixed time window, which is a standard procedure. When the evoked firing rate r is plotted against x , the result will be a curve like those shown in Fig. 2, with a peak at the preferred stimulus location of the cell. Now repeat the measurements with a different value of y , say 10° to the right this time. How does the second curve (r vs. x) compare with the first one?


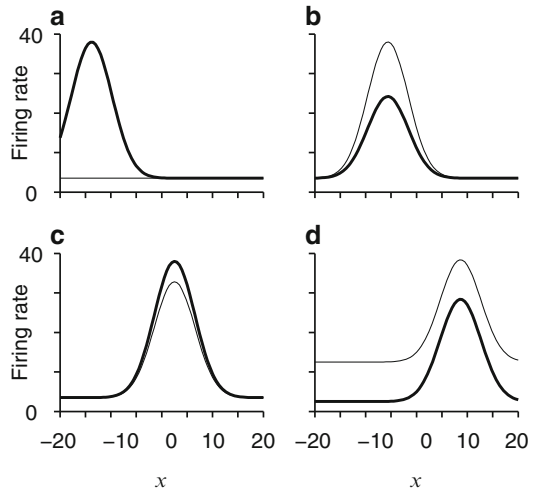


Fig. 1 Reference points for locating an object. The *cross* indicates the fixation point and the *star* indicates a visual stimulus. The angle of the eye in the orbit is y . The position of the object in eye-centered coordinates (i.e., relative to the fixation point) is x

Fig. 2 Hypothetical responses of four neurons that are sensitive to two quantities, x and y . Each panel shows the firing rate as a function of x for two values of y , which are distinguished by *thin* and *thick* lines. (a–c), Three neurons that are gain-modulated by y . (d), A neuron that combines x - and y -dependencies linearly, and so is not gainmodulated



If the neuron is not sensitive to eye position at all, then the new curve will be the same as the old one. This would be the case, for instance, with the ganglion cells of the retina, which respond depending on where light falls on the retina. On the other hand, if the neuron is gain-modulated by eye position, then the second curve will be a scaled version of the first one. That is, the two curves will have the same shapes but different amplitudes. This is illustrated in Fig. 2b and c, where thin and thick lines are used to distinguish the two y conditions. This is precisely what happens in the posterior parietal cortex [1, 2, 12], where this phenomenon was first documented by Andersen and Mountcastle [3]. To describe the dependency on the modulatory quantity y , these authors coined the term “gain field,” in analogy with the receptive field, which describes the dependency on the sensory stimulus x . Of course, to properly characterize the gain field of a cell, more values of y need to be tested [2], but for the moment we will keep discussing only two.

To quantify the strength of the modulation, consider the decrease in response amplitude observed as a neuron goes from its preferred to its nonpreferred y condition. In panels a–c of Fig. 2, the resulting percent suppression is, respectively, 100%, 40%, and 15%. In the parietal cortex, and in experimental data in general, 100% suppression is very rarely seen; on the contrary, the effects are often quite subtle (more on this below). More importantly, however, the firing rate r of the neuron can be described as a product of two factors, a function of x , which determines the shape of the curves, times a gain factor that determines their amplitude:

$$r_j(x, y) = g_j(y)f_j(x) \quad (1)$$

where the index j indicates that this applies to neuron j . The functions f and g may have different forms, so this is a very general expression. Much of the data reported by Andersen and colleagues, as well as data from other laboratories, can be fit quite well using (1) [2, 12, 30, 31, 57]. Note, however, that for the theoretical results

discussed later on, the key property of the above expression is that it combines x and y in a nonlinear way. Thus, in general, the theoretical results are still valid if the relationship between x - and y -dependencies is not exactly multiplicative, as long as it remains nonlinear. A crucial consequence of this is that a response that combines x and y dependencies linearly, for instance,

$$r_j(x, y) = g_j(y) + f_j(x) \quad (2)$$

does *not* qualify as a gain-modulated response. Note how different this case is from the multiplicative one: when y changes, a neuron that responds according to (2) increases or decreases its firing rate by the same amount for all x values. This is illustrated in Fig. 2d, where the peaked curve moves up and down as y is varied but nothing else changes other than the baseline. Such a neuron is sensitive to both x and y , but (1) it cannot be described as gain-modulated by y , because it combines f and g linearly, and (2) it cannot perform the powerful computations that gain-modulated neurons can.

Experimental Evidence for Gain Modulation

To give an idea of the diversity of neural circuits in which gain modulation may play an important functional role, in this section we present a short survey of experimental preparations in which gain modulation has been observed. In all of these cases the reported effects are roughly in accordance with (1). Note that, although the x and y variables are different in all of these examples, all experiments are organized like the one discussed in the previous section: the firing rate of a neuron is plotted as a function of x for a fixed value of y ; then y is varied and the response function or tuning curve (r vs. x) of the cell is plotted again; finally, the results with different y values are compared.

Modulation by Proprioceptive Information

To begin, we mention three representative examples of proprioceptive signals that are capable of modulating the visually evoked activity of cortical neurons. First of all, although initial work described responses tuned to stimulus location (x) with gain fields that depend on eye position, an additional dependence on head position was discovered later, such that the combination of head and eye angles – the gaze angle – seems to be the relevant quantity [12]. Thus, the appropriate modulatory parameter y is likely to be the gaze angle.

Another example of modulation by proprioceptive input is the effect of eye and head velocity (which now play the role of y) on the responses of neurons that are sensitive to heading direction (which now corresponds to x). Neurons in area MSTd respond to the patterns of optic flow that are generated by self-motion; for instance,

when a person is walking, or when you are inside a moving train and look out the window. Under those conditions, the MSTd population encodes the direction of heading. Tuning curves with a variety of shapes are observed when the responses of these neurons are plotted as functions of heading direction, and crucially, for many units those curves are gain-modulated by eye and head velocity [50].

A third example of proprioceptive modulation occurs in an area called the parietal reach region, or PRR, where cells are tuned to the location of a reach target. The reach may be directed toward the location of a visible object or toward its remembered location; in either case, PRR cells are activated only when there is an intention to reach for a point in space. The response of a typical PRR cell as a function of target location (the x parameter in this case) is peaked. A key modulatory quantity for these cells is the initial location of the hand (this is y now): the amplitude of the peaked responses to target location depends on the position of the hand *before* the reach [7, 13]. We will come back to this example later on.

Attentional Modulation

Another modulatory signal that has been widely studied is attention, and its effects on sensory-driven activity are often very close to multiplicative. For example, visual neurons in area V4 are sensitive to the orientation of a bar shown in their receptive field. The corresponding tuning curves of firing rate vs. bar orientation change depending on where attention is directed by the subject. In particular, the tuning curves obtained with attention directed inside vs. outside the receptive field are very nearly scaled versions of each other, with an average suppression of about 20% in the non-preferred condition [31]. The magnitude of this effect is not very large, but this is with rather simple displays that include only two stimuli, a target and a distractor; results reported using more crowded displays and more demanding tasks are two-to-three times as large [14, 21].

Another cortical area where attentional effects have been carefully characterized is area MT, which processes visual motion. There, neurons are sensitive to the direction in which an object, or a cloud of objects, moves across their receptive fields. Thus, for each neuron the relevant curve in this case is that of firing rate vs. movement direction. What the experiments show is that the location where attention is directed has an almost exactly multiplicative effect on the direction tuning curves of MT cells. But interestingly, the amplitude of these curves also depends on the specific movement direction that is being attended [30, 57]. Thus, as in V4 [21], stronger suppression (of about 50%) can be observed when an appropriate combination of attended location and attended feature is chosen.

Nonlinear Interactions between Multiple Stimuli

In addition to studies in which two distinct variables are manipulated, there is a vast literature demonstrating other nonlinear interactions that arise when multiple

stimuli are displayed simultaneously. For instance, the response of a V1 neuron to a single oriented bar or grating shown inside its receptive field may be enhanced or suppressed by neighboring stimuli outside the receptive field [20, 25, 28, 51]. Thus, stimuli that are ignored by a neuron when presented in isolation may have a strong influence when displayed together with an effective stimulus. These are called “extra classical receptive field” effects, because they involve stimuli outside the normal activation window of a neuron.

Nonlinear summation effects, however, may also be observed when multiple stimuli are shown inside the receptive field of a cell. For example, suppose that a moving dot is presented inside the receptive field of an MT neuron and the evoked response is r_1 . Then another moving dot is shown moving in the same direction but at a different location in the receptive field, and the response is now r_2 . What happens if the two moving dots are displayed simultaneously? In that case, the evoked response is well described by

$$r = \alpha (r_1^n + r_2^n)^{1/n} \quad (3)$$

where α is a scale factor and n is an exponent that determines the strength of the nonlinearity [11]. This expression is interesting because it captures a variety of possible effects, from simple response averaging ($\alpha = 0.5$, $n = 1$) or summation ($\alpha = 1$, $n = 1$), to winner-take-all behavior ($\alpha = 1$, n large). Neurons in MT generally show a wide variety of n exponents [11], typically larger than 1, meaning that the presence of one stimulus modulates nonlinearly the response to the other. Neurons in V4 show similar nonlinear summation [21].

In general, many interactions between stimuli are well described by divisive normalization models [52, 53], of which (3) is an example. Such models implement another form of gain control where two or more inputs are combined nonlinearly.

Context- and Task-Dependent Modulation

Finally, there are numerous experiments in which more general contextual effects have been observed. This means that activity that is clearly driven by the physical attributes of a sensory stimulus is also affected by more abstract factors, such as the task that the subject needs to perform, or the subject’s motivation. These phenomena are often encountered in higher-order areas, such as the prefrontal cortex. For example, Lauwereyns and colleagues designed a task in which a monkey observed a cloud of moving dots and then, depending on a cue at the beginning of each trial, had to respond according to either the color or the direction of movement of the dots [27]. Prefrontal neurons were clearly activated by the moving dots; however, the firing rates evoked by identical stimuli varied quite dramatically depending on which feature was important in a given trial, color or movement direction.

Another striking case of task-dependent modulation was revealed by Koida and Komatsu [26] using color as a variable. They trained monkeys to perform two color-based tasks. In one of them, the monkeys had to classify a single color patch into

one of two categories, reddish or greenish. In the other task, the monkeys had to indicate which of two color patches matched a previously shown sample. Neuronal responses in inferior-temporal cortex (IT) were recorded while the monkeys performed these two tasks plus a passive fixation task, in which the same stimuli were presented but no motor reactions were required. The responses of many IT neurons were strongly gain-modulated by the task: their color tuning curves had similar shapes in all cases, but the overall amplitude of their responses depended on which task was performed. This means that their activity was well described by (1), with x being color and y being the current task. Note that, in these experiments, the exact same stimuli were used in all tasks – what varied was the way the sensory information was processed. Furthermore, the classification and discrimination tasks were of approximately equal difficulty, so it is unlikely that the observed modulation was due to factors such as motivation or reward rate. Nevertheless, these two variables are known to modulate many types of neurons quite substantially [36, 49], so the extent of their influence is hard to determine.

A last example of modulation due to task contingencies comes from neurophysiological work in the supplementary and presupplementary motor areas, which are involved in fine aspects of motor control such as the planning of movement sequences. Sohn and Lee made recordings in these areas while monkeys performed sequences of joystick movements [54]. Each individual movement was instructed by a visual cue, but importantly, the monkey was rewarded only after correct completion of a full sequence, which consisted of 5–10 movements. The activity of the recorded neurons varied as a function of movement direction, as expected in a motor area, but interestingly, it also varied quite strongly according to the number of remaining movements in a sequence. It is possible that the true relevant variable here is the amount of time or the number of movements until a *reward* is obtained. Regardless of this, remarkably, the interaction between movement direction (the x variable in this case) and the number of remaining movements in a sequence (the y variable) was predominantly multiplicative, in accordance with (1).

In summary, there is an extremely wide range of experimental preparations and neuronal circuits that reveal strongly nonlinear modulation of evoked activity. The next sections discuss (1) some computational operations that neuronal networks can implement relatively easily based on such nonlinearities, and (2) a variety of behaviors and tasks where such operations could play crucial functional roles.

Computations Based on Gain Modulation

The previous section reviewed a large variety of experimental conditions under which gain modulation – understood as a nonlinear interaction between x - and y -dependencies – is observed. What is interesting about the chosen examples, and is not immediately obvious, is that all of those effects may be related to similar computational operations. That is precisely what theoretical and modeling studies have shown, that gain modulation serves to implement a type of mathematical

transformation that arises under many circumstances. It is relatively well established that gain modulation plays an important role in carrying out coordinate transformations [4, 38, 46, 48], so that is the first type of operation that we will discuss. However, the applicability of gain modulation goes way beyond the traditional notion of coordinate transformation [41].

Coordinate Transformations

The first clue that gain modulation is a powerful computing mechanism was provided by Zipser and Andersen [63], who trained an artificial neural network to transform the neural representation of stimulus location from one reference frame to another. Their network had three layers of model neurons. The activities in the input and output layers were known – they were given to the model – and the job of the model was to determine the responses of the units in the middle layer. The input layer had two types of neurons, ones that encoded the location of a visual stimulus in eye-centered coordinates, i.e., with respect to the fixation point, and others that encoded the gaze angle. Thus, stimulus location was x and gaze angle was y , exactly as discussed in the section “What is Gain Modulation.” In contrast, the neurons in the output layer encoded the location of the stimulus in head-centered coordinates, which for this problem means that they responded as functions of $x + y$. It can be seen from Fig. 1 that $x + y$ in this case gives the location of the stimulus relative to the head.

Adding x plus y seems deceptively simple, but the problem is rather hard because x , y , and $x + y$ are encoded by populations of neurons where each unit responds nonlinearly to one of those three quantities. The middle neurons in the model network had to respond to each pattern of evoked activity (in the first layer) and produce the correct output pattern (in the third layer). The way the authors were able to “train” the middle layer units was by using a synaptic modification rule, the backpropagation algorithm, which modified the network connections each time an input and a matching output pattern were presented, and thousands of input and output examples were necessary to learn the mapping with high accuracy. But the training procedure is not very important; what is noteworthy is that, after learning, the model units in the middle layer had developed gaze-dependent gain fields, much like those found in parietal cortex. This indicated that the visual responses modulated by gaze angle are an efficient means to compute the required coordinate transformation.

Further intuition about the relation between gain fields and coordinate transformations came from work by Salinas and Abbott [43], who asked the following question. Consider a population of parietal neurons that respond to a visual stimulus at retinal location x and are gain-modulated by gaze angle y . Suppose that the gain-modulated neurons drive, through synaptic connections, a downstream population of neurons involved in generating an arm movement toward a target. To actually reach the target with any possible combination of target location and gaze angle, these

downstream neurons must encode target location in body-centered coordinates, that is, they must respond as functions of $x + y$ (approximately). Under what conditions will this happen? Salinas and Abbott developed a mathematical description of the problem, found those conditions, and showed that, when driven by gain-modulated neurons, the downstream units can explicitly encode the sum $x + y$ or any other linear combination of x and y . In other words, the gain-field representation is powerful because downstream neurons can extract *from the same modulated neurons* any quantity $c_1x + c_2y$ with arbitrary coefficients c_1 and c_2 , and this can be done quite easily – “easily” meaning through correlation-based synaptic modification rules [43, 44]. Work by Pouget and Sejnowski [37] extended these results by showing that a population of neurons tuned to x and gain-modulated by y can be used to generate a very wide variety of functions of x and y downstream, not only linear combinations.

A crucial consequence of these theoretical results is that when gain-modulated neurons are found in area A , it is likely that neurons in a downstream area B will have response curves as functions of x that shift whenever y is varied. This is because a tuning curve that is a function of $c_1x + c_2y$ will shift along the x -axis when y changes. To illustrate this point, consider another example of a coordinate transformation; this case goes back to the responses of PRR cells mentioned in the section “Modulation by Proprioceptive Information.”

The activity of a typical PRR cell depends on the location of a target that is to be reached, but is also gain-modulated by the initial position of the hand before the reach [7, 13]. Thus, now x is the target’s location and y is the initial hand position (both quantities in eye-centered coordinates). Fig. 3a–c illustrates the corresponding experimental setup with three values of y : with the hand initially to the left, in the middle, or to the right. In each diagram, the small dots indicate the possible locations of the reach target, and the two gray disks represent the receptive fields of two neurons. The response of a hypothetical PRR neuron is shown in Fig. 3e, which includes three plots, one for each of the three initial hand positions. In accordance with the reported data [7, 13], the gain of this PRR cell changes as a function of initial hand position. The other two cells depicted in the figure correspond to idealized responses upstream and downstream from the PRR. Figure 3d shows the firing rate of an upstream cell that is insensitive to hand position y , so its receptive field (light disk in Fig. 3a–c) does not change. This cell is representative of a neural population that responds only as a function of x , and simply encodes target location in eye-centered coordinates. In contrast, Fig. 3f plots the tuning curve of a hypothetical neuron downstream from the PRR. The curve shifts when the initial hand position changes because it is a function of $x - y$ (as explained above, but with $c_1 = 1$ and $c_2 = -1$). For such neurons, the receptive field (dark disk in Fig. 3a–c) moves as a function of hand position. This cell is representative of a population that encodes target location in hand-centered coordinates: the firing responses are the same whenever the target maintains the same spatial relationship relative to the hand. Neurons like this one are found in parietal area 5, which is downstream from the PRR [13]. They can be constructed by combining the responses of multiple gain-modulated neurons like that in Fig. 3e, but with diverse target and hand-location preferences.

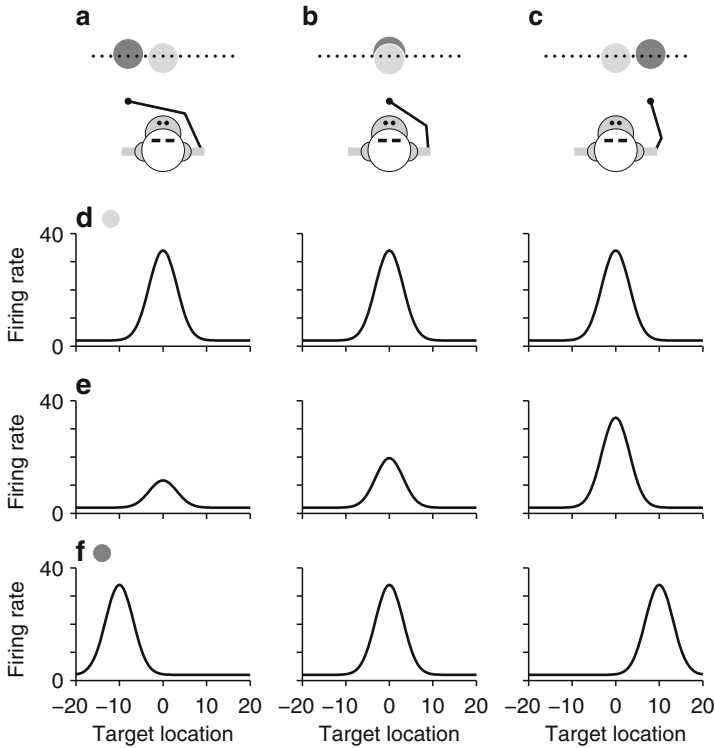


Fig. 3 Hypothetical responses of three neurons that play a role in reaching a target. (a–c), Three experimental conditions tested. *Dots* indicate the possible locations of a reach target. Conditions differ in the initial hand position (y). *Circles* indicate visual receptive fields of two hypothetical neurons, as marked by gray disks below. (d–f), Firing rate as a function of target location (x) for three idealized neurons. Each neuron is tested in the three conditions shown above. Gaze is assumed to be fixed straight ahead in all cases

Why are such shifting tuning curves useful? Because the transformation in the representation of target location from Fig. 3d to Fig. 3f partially solves the problem of how to acquire a desired target. In order to reach an object, the brain must combine information about target location with information about hand position to produce a vector. This vector, sometimes called the “motor error” vector, encodes how much the hand needs to move and in what direction. Cells like the one in Fig. 3d do not encode this vector at all, because they have no knowledge of hand position. Cells that are gain modulated by hand position, as in Fig. 3e, combine information about x and y , so they do encode the motor vector, but they do so implicitly. This means that actually reading out the motor vector from their evoked activity requires considerable computational processing. In contrast, neurons like the one in Fig. 3f encode the motor vector explicitly: reading out the motor vector from their evoked activity requires much simpler operations.

To see this, draw an arrow from the hand to the center of the shifting receptive field (the dark disk) in Fig. 3a–c. You will notice that the vector is the same regardless of hand position; it has the same length L and points in the same direction, straight ahead, in the three conditions shown. Thus, a strong activation of the idealized neuron of Fig. 3f, which is elicited when x is equal to the peak value of the tuning curve, can be easily translated into a specific motor command; such activation means “move the hand L centimeters straight ahead and you’ll reach the target.” Conversely, for this command to be valid all the time, the preferred target position of the cell must move whenever the hand moves.

In reality, area 5 neurons typically do not show full shifts [13]; that is, their curves do shift but by an amount that is smaller than the change in hand position. This is because their activity still depends to a certain extent on eye position. It may be that areas further downstream from area 5 encode target position in a fully hand-centered representation, but it is also possible that such a representation is not necessary for the accurate generation of motor commands. In fact, although full tuning curve shifts associated with a variety of coordinate transformations have been documented in many brain areas, partial shifts are more common [6, 19, 22, 24, 50, 55]. This may actually reflect efficiency in the underlying neural circuits: according to modeling work [16, 62], partial shifts should occur when more than two quantities x and y are simultaneously involved in a transformation, which is typically the case. Anyhow, the point is that a gain-modulated representation allows a circuit to construct shifting curves downstream, which encode information in a format that is typically more accessible to the motor apparatus.

A final note before closing this section. Many coordinate transformations make sense as a way to facilitate the interaction with objects in the world, as in the examples just discussed. But the shifting of receptive fields, and in general the transformation of visual information from an eye-centered representation to a reference frame that is independent of the eye may be an important operation for perception as well. For instance, to a certain degree, humans are able to recognize objects regardless of their size, perspective, and position in the visual field [15, 18]. This phenomenon is paralleled by neurons in area IT, which respond primarily depending on the type of image presented, and to some extent are insensitive to object location, scale or perspective [17, 56, 64]. For constructing such invariant responses, visual information must be transformed from its original retinocentered representation, and modeling work suggests that gain modulation could play an important role in such perceptual transformations too, because the underlying operations are similar. In particular, directing attention to different locations in space alters the gain of many visual neurons that are sensitive to local features of an image, such as orientation [14]. According to the theory [44, 45], this attentional modulation could be used to generate selective visual responses similar to those in IT, which are scale- and translation-invariant. Such responses would represent visual information in an attention-centered reference frame, and there is evidence from psychophysical experiments in humans [29] indicating that object recognition indeed operates in a coordinate frame centered on the currently attended location.

Arbitrary Sensory-Motor Remapping

In the original modeling studies on gain modulation, proprioceptive input combined with spatial sensory information led to changes in reference frame. It is also possible, however, to use gain modulation to integrate sensory information with other types of signals. Indeed, the mechanism also works for establishing arbitrary associations between sensory stimuli and motor actions on the basis of more abstract contextual information.

In many tasks and behaviors, a given stimulus is arbitrarily associated with two or more motor responses, depending on separate cues that we will refer to as “the context.” Figure 4 schematizes one such task. Here, the shape of the fixation point indicates whether the correct response to a bar is a movement to the left or to the right (arrows in Fig. 4a and b). Importantly, the same set of bars (Fig. 4c) can be partitioned in many ways; for instance, according to orientation (horizontal vs. vertical), color (filled vs. not filled), or depending on the presence of a feature (gap vs. no gap). Therefore, there are many possible maps between the eight stimuli and the two motor responses, and the contextual cue that determines the correct map is rather arbitrary.

Neurophysiological recordings during such tasks typically reveal neuronal responses that are sensitive to both the ongoing sensory information and the current context or valid cue, with the interaction between them being nonlinear [5, 27, 58, 59]. Some of these effects were mentioned in the section “Context- and Task-Dependent Modulation.” The key point is that, as the context changes in such tasks, there is (1) a nonlinear modulation of the sensory-triggered activity, and (2) a functional reconnection between sensory and motor networks that must be very fast. How is this reconnection accomplished by the nervous system?

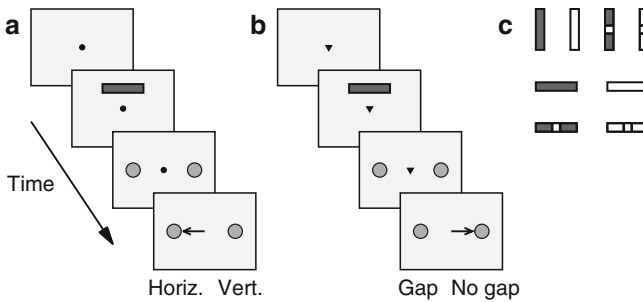


Fig. 4 A context-dependent classification task. The shape of the fixation point determines the correct response to a stimulus. (a), Events in a single trial in which a bar is classified according to orientation. (b), Events in a single trial in which a bar is classified according to whether it has a gap or not. (c), Stimulus set of eight bars

How the Contextual Switch Works

The answer is that such abstract transformations can be easily generated if the sensory information is gain-modulated by the context. According to modeling work [40,41], when this happens, multiple maps between sensory stimuli and motor actions are possible but only one map, depending on the context, is implemented at any given time. Let's sketch how this works.

In essence, the mechanism for solving the task in Fig. 4, or other tasks like it, is very similar to the mechanism for generating coordinate transformations (i.e., shifting receptive fields) described in section "Coordinate Transformations," except that x now describes the stimulus set and y the possible contexts. Instead of continuous quantities, such as stimulus location and eye position, these two variables may now be considered indices that point to individual elements of a set. For instance, $x = 1$ may indicate that the first bar in Fig. 4c was shown, $x = 2$ may indicate that the second bar was shown, etc. In this way, the expression $f_j(x)$ still makes sense: it stands for the firing rate of cell j in response to each stimulus in the set. Similarly, $y = 1$ now means that the current context is context 1, $y = 2$ means that the current context is context 2, and so on. Therefore, the response of a sensory neuron j that is gain-modulated by context can be written exactly as in (1),

$$r_j(x, y) = g_j(y) f_j(x) \quad (4)$$

except that now $r_j(x = 3, y = 1) = 10$ means "neuron j fires at a rate of 10 spikes/s when stimulus number 3 is shown and the context is context 1." Although x and y may represent quantities with extremely different physical properties, practically all of the important mathematical properties of (1) remain the same regardless of whether x and y vary smoothly or discretely. As a consequence, although the coordinate transformations in Figs. 1 and 3 may feel different from the remapping task of Fig. 4, mathematically they represent the very same problem, and the respective neural implementations may thus have a lot in common.

Now consider the activity of a downstream neuron that is driven by a population of gain-modulated responses through a set of synaptic weights,

$$R(x, y) = \sum_j w_j g_j(y) f_j(x) \quad (5)$$

where w_j is the connection from modulated neuron j to the downstream unit. Notice that the weights and modulatory factors can be grouped together into a term that effectively behaves like a context-dependent synaptic connection. That is, consider the downstream response in context 1,

$$R(x, y = 1) = \sum_j u_j^{(1)} f_j(x) \quad (6)$$

where we have defined $u_j^{(1)} = w_j g_j(y = 1)$. The set of coefficients $u_j^{(1)}$ represents the effective synaptic weights that are active in context 1. Now comes the crucial result: because the modulatory factors g_j change as functions of context, those effective weights will be different in context 2, and so will be the downstream response,

$$R(x, y = 2) = \sum_j u_j^{(2)} f_j(x) \quad (7)$$

where $u_j^{(2)} = w_j g_j(y = 2)$. Thus, the effective network connections are $u_j^{(1)}$ in one context and $u_j^{(2)}$ in the other. The point is that the downstream response in the two contexts, given by (6) and (7), can produce completely different functions of x . That is exactly what needs to happen in the classification task of Fig. 4a and b – the downstream motor response to the same stimulus (x) must be able to vary arbitrarily from one context to another. The three expressions above show why the mechanism works: the multiplicative changes in the gain of the sensory responses act exactly as if the connectivity of the network changed from one context to the next, from a set of connections $u_j^{(1)}$ to a set $u_j^{(2)}$.

Figure 5 illustrates this with a model that performs the task in Fig. 4, classifying a stimulus set differently in three possible contexts. In contexts 1 and 2, the stimuli trigger left and right movements according, respectively, to their orientation or to whether they have a gap or not, as in Fig. 4a and b. In contrast, in context 3 the correct response is to make no movement at all; this is a no-go condition. The model network used to generate Fig. 5 included 160 gain-modulated sensory neurons in the first layer and 50 motor neurons in the second, or output layer. Each sensory neuron responded differently to a set of eight distinct stimuli like those shown in Fig. 4c. For each cell, the maximum suppression across the three contexts varied between 0 and 30%, with a mean of 15%. The figure shows the responses of all the




Fig. 5 A model network that performs the remapping task of Fig. 4. Gray-scale images represent the responses of 160 gain-modulated sensory neurons to the combinations of x (stimulus) and y (context) values shown above. White and black correspond to firing rates of 3 and 40 spikes/s, respectively. Plots at the bottom show the responses of 50 motor neurons driven by the gain-modulated units through synaptic connections. Identical stimuli lead to radically different motor responses depending on the current context

neurons in the network in four stimulus-context combinations. The firing rates of the gain-modulated neurons of the first layer are indicated by the gray-level maps. In these maps, the neurons are ordered such that their position along the x -axis roughly reflects their preferred stimulus, which goes from 1 to 8. The firing rates of the 50 motor units of the second layer are plotted below. For those cells, the location of the peak of activity encodes the location of a motor response.

There are a couple of notable things here. First, that the sensory activity changes quite markedly from one stimulus to another (compare the two conditions with $y = 1$), but much less so from one context to another (compare the two conditions with $x = 2$). This is because the modulatory factors do not need to change very much in order to produce a complete change in connectivity (more on this below). Second, what the model accomplishes is a drastic change in the motor responses of the second layer *to the same stimulus*, as a function of context. This includes the possibility of silencing the motor responses altogether or not, as can be seen by comparing the two conditions with $x = 7$.

In summary, in this type of network a set of sensory responses are nonlinearly modulated by contextual cues, but these do not need to be related in any way to the physical attributes of the stimuli. The model works extremely well, in that accuracy is limited simply by the variability (random fluctuations) in the gain-modulated responses and by the number of independent contexts [40]. In theory, it should work in any situation that involves choosing or switching between multiple functions or maps, as long as the modulation is nonlinear – as mentioned earlier, an exact multiplication as in (4) is not necessary, as long as x and y are combined nonlinearly. With the correct synaptic connections driving a group of downstream neurons, such a network can perform a variety of remapping tasks on the basis of a single transformation step between the two layers [40–42].

Switching as a Fundamental Operation

The model network discussed in the previous section can be thought of as implementing a context-dependent switch that changes the functional mapping between the two layers in the network. In fact, it can be shown that gain modulation has an effect that is equivalent to turning on and off different subsets of neurons as a function of context, or more generally, of the modulatory quantity y [41].

To illustrate this, consider the problem of routing information from a network A to two possible target networks B and C , such that either the connections $A \rightarrow B$ are active or the connections $A \rightarrow C$ are active, but not both. An example of this is a task in which the subject has to reach for an object that may appear at a variety of locations, but depending on the shape of the fixation point, or some other cue, he must reach with either the left or the right hand. In this situation, the contextual cue must act as a switch that activates one or the other downstream motor circuit, but it must do so for any stimulus value.

Fig. 6 Two possible schemes for routing information to either of two motor networks. One motor network should respond in context 1 and the other in context 2. **(a)** With two subpopulations of “switching neurons,” which are 100% suppressed in their nonpreferred contexts. **(b)** With one population of partially modulated neurons, which show mild suppression in their nonpreferred contexts

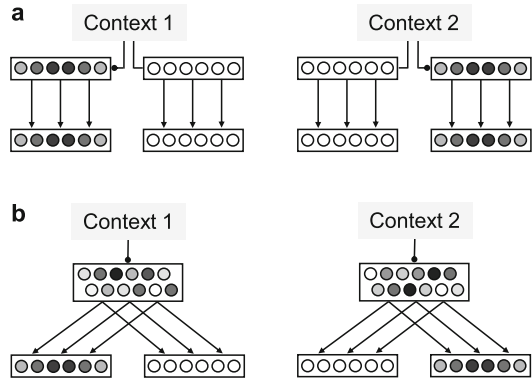


Figure 6 schematizes two possible ways in which such a switch could be implemented. In the network of Fig. 6a, there are two separate sets of sensory responses that are switched on and off according to the context, and each subpopulation drives its own motor population downstream. This scheme requires true “switching neurons” that can be fully suppressed by the nonpreferred context. This type of neuron would need to respond as in Fig. 2a. This, admittedly, is a rather extreme and unrealistic possibility, but it nonetheless provides an interesting point of comparison: clearly, with such switching neurons it is trivial to direct the sensory information about target location to the correct motor network; and it is also clear that the two motor networks could respond in different ways to the same sensory input. In a way, the problem has simply been pushed back one step in the process.

An alternative network organization is shown in Fig. 6b. In this case, one single set of gain-modulated sensory neurons is connected to the two motor populations. What can be shown theoretically [41], is that any sensory-motor mapping that can be achieved with switching neurons (as in Fig. 6a), can also be implemented with standard, partially modulated neurons (as in Fig. 6b). Of course the necessary synaptic weights will be different in the two cases, but in principle, under very general conditions, if something can be done with 100% suppression, it can also be done with a much smaller amount of suppression (there may be a cost in terms of accuracy, but this depends on the level of noise in the responses and on the size of the network, among other factors [41]). In particular, a suppression of 100% is not necessary to produce a switch in the activated downstream population, as required by the task just discussed in which either the left or the right hand is used. This result offers a unique intuition as to why gain modulation is so powerful: modulating the activity of a population is equivalent to flipping a switch that turns on or off various sets of neurons.

This result is illustrated in Fig. 7, which shows the responses of a model network that has the architecture depicted in Fig. 6b. The x and y variables here are, respectively, the location of a stimulus, which takes a variety of values between -16 and 16 , and the context, which takes two values. Figure 7a, c shows the activity of the 100 gain-modulated sensory neurons in the network in contexts 1 and 2,

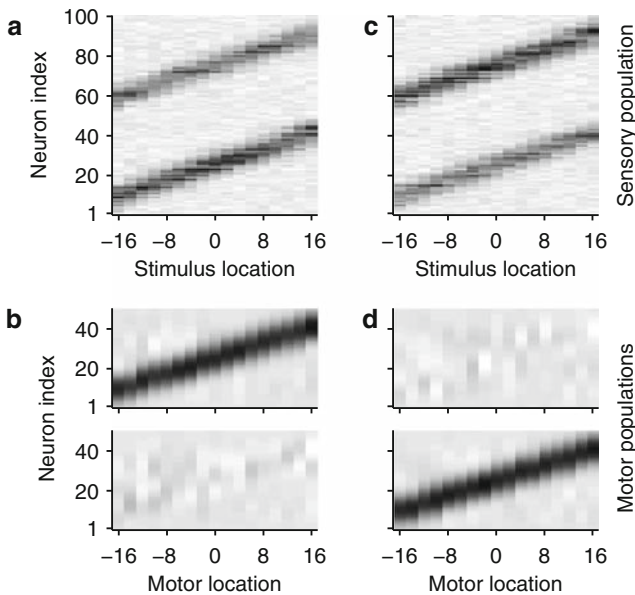


Fig. 7 Responses in a model network that routes information to two separate motor populations, as in Fig. ch07:fig6b. In context 1, one motor population is active and responds to all stimulus locations (x); in context 2, the other motor population is active. **a, c** Responses of the gain-modulated sensory neurons in contexts 1 and 2, respectively. **b, d** Responses of the motor neurons in contexts 1 and 2, respectively

respectively. These units are tuned to stimulus location, so they have peaked tuning curves like those in Fig. 2b, c. The 100 units have been split into two groups, those that have a higher gain or response amplitude in context 1 (units 1–50) and those that have a higher gain in context 2 (units 51–100). The neurons in each group are arranged according to their preferred stimulus location. This is why Fig. 7a and c show two bands: for each value of x , the neurons that fire most strongly in each group are those that have preferred values close to x . Notice that the activity patterns in Fig. 7a, c are very similar. This is because the differences in gain across the two contexts are not large: the percent suppression in this case ranged between 0 and 65%, with a mean of 32%.

What is interesting about this example is the behavior of the motor populations downstream. In context 1, one population responds and the other is silent (Fig. 7b), whereas in context 2 the roles are reversed, the second population responds and the first is silent (Fig. 7d). Importantly, this does not occur just for one value or a small number of values of x , but for any value. This switch, controlled by a contextual cue y , is generated on the basis of relatively small modulatory effects in the top layer. Thus, the modeling results reviewed here provide another interesting insight: small variations in the response gain of neurons in area A may be a signature of dramatically large variations in the activity of a downstream area B .

In the example of Fig. 7, the activated motor population simply mirrors the activation pattern in the sensory layer; that is, the sensory-motor map is one-to-one. This simple case was chosen to illustrate the switching mechanism without further complications, but in general, the relationship between stimulus location and the encoded motor variable (motor location) may be both much more complicated and different for the two contexts.

Flexible Responses to Complex Stimuli


To recapitulate what has been presented: coordinate transformations and arbitrary sensory-motor associations are types of mathematical problems that are akin to implementing a switch from one function of x to another, with the switch depending on a different variable y . Thus, in the mammalian brain, performing coordinate transformations and establishing arbitrary sensory-motor maps should lead to the deployment of similar neural mechanisms. Modeling work has shown that the non-linear modulation of sensory activity can reduce the problem of implementing such a switch quite substantially, and this is consistent with experimental reports of nonlinear modulation by proprioceptive, task-dependent, and contextual cues in a variety of tasks (“Experimental Evidence for Gain Modulation”).

There is, however, an important caveat. One potential limitation of gain modulation as a mechanism for performing transformations is that it may require an impossibly large number of neurons, particularly in problems such as invariant object recognition [18, 39, 44, 64] in which the dimensionality of the relevant space is high. The argument goes like this. Suppose that there are N stimuli to be classified according to M criteria, as a generalization of the task in Fig. 4. To perform this task, a neuronal network will need neurons with all possible combinations of stimulus and context sensitivities. Thus, there must be neurons that prefer stimuli 1, 2, \dots N in context 1, neurons that prefer stimuli 1, 2, \dots N in context 2, and so on. Therefore, the size of the network should grow as $N \times M$, which can be extremely large. This is sometimes referred to as “the curse of dimensionality.”

Although this is certainly a challenging problem, its real severity is still unclear. In particular, O’Reilly and colleagues have argued [34, 35] that the combinatorial explosion may be much less of a problem than generally assumed. In essence, they say, this is because neurons are typically sensitive to a large range of input values, and the redundancy or overlap in this coarse code may allow a network to solve such problems accurately with many fewer than $N \times M$ units. At least in some model networks, it has indeed been demonstrated that accurate performance can be obtained with many fewer than the theoretical number of necessary units [10, 34, 35]. In practice, this means that it may be possible for real biological circuits to solve many remapping tasks with relatively few neurons, if they exhibit the appropriate set of broad responses to the stimuli and the contexts.

So, assuming that the curse of dimensionality does not impose a fundamental limitation, does gain modulation solve the flexibility problem discussed at the

Fig. 8 Two examples of scenes with multiple objects. Each display can be considered as a single complex stimulus



beginning of the chapter? No, it does not, and this is why. All of the tasks and behaviors mentioned so far have been applied to single, isolated stimuli. If an organism always reacted to one stimulus at a time, then gain modulation would be enough for selecting the correct motor response under any particular context, assuming enough neurons or an appropriate coarse code. However, under natural conditions, a sensory scene in any modality contains multiple objects, and choices and behaviors are dictated by such crowded scenes. Other neural mechanisms are needed for (1) selecting one relevant object from the scene, or (2) integrating information across different locations or objects within a scene (note that similar problems arise whenever a choice needs to be made on the basis of a collection of objects, even if they appear one by one spread out over time). Thus, to produce a flexible, context-dependent response to a complex scene, gain modulation would need to be combined at least with other neural mechanisms for solving these two problems.

To put a simple example, consider the two displays shown in Fig. 8. If each display is taken as one stimulus, then there are many possible ways to map those two stimuli into “yes” or “no” responses. For instance, we can ask: Is there a triangle present? Is an open circle present on the right side of the display? Are there five objects in the scene? Or, are all objects different? Each one of these questions plays the role of a specific context that controls the association between the two stimuli and the two motor responses. The number of such potential questions (contexts) can be very large, but the key problem is that answering them requires an evaluation of the whole scene. Also, it is known from psychophysical experiments that subjects can change from one question to another extremely rapidly [61]. Gain modulation has the advantage of being very fast, and could indeed be used to switch these associations, but evidently the mechanisms for integrating information across a scene must be equally important.

As a first step in exploring this problem, we have constructed a model network that performs a variety of search tasks [8]. It analyzes a visual scene with up to eight objects and activates one of two output units to indicate whether a target is present or absent in the scene. As in the models described earlier, this network includes sensory neurons, which are now modulated by the identity of the search target, and output neurons that indicate the motor response. The target-dependent modulation allows the network to switch targets, that is, to respond to the question, is there a red vertical bar present, or to the question, is there a blue bar present, and so on. Not surprisingly, this model does not work unless an additional layer is included that integrates information across different parts of the scene. When this is done, gain modulation indeed still allows the network to switch targets, but the properties of the integration layer become crucial for performance too. There are other

models for visual search that do not switch targets based on gain modulation; these are generally more abstract models [23, 60]. Our preliminary results [8] are encouraging because neurophysiological recordings in monkeys trained to perform search tasks have revealed sensory responses to individual objects in a search array that are indeed nonlinearly modulated according to the search target or the search instructions [9, 33]. Furthermore, theoretical work by Navalpakkam and Itti has shown that the gains of the sensory responses evoked by individual elements of a crowded scene can be set for optimizing the detection of a target [32]. This suggests that there is an optimal gain modulation strategy that depends on the demands of the task, and psychophysical results in their study indicate that human subjects do employ this optimal strategy [32].

In any case, the analysis of complex scenes should give a better idea of what can and cannot be done using gain modulation, and should provide ample material for future studies of the neural basis of flexible behavior.

Acknowledgments Research was partially supported by grant NS044894 from the National Institute of Neurological Disorders and Stroke.

References

1. Andersen RA, Bracewell RM, Barash S, Gnadt JW, Fogassi L (1990) Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10:1176–1198
2. Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230:450–458
3. Andersen RA, Mountcastle VB (1983) The influence of the angle of gaze upon the excitability of light-sensitive neurons of the posterior parietal cortex. *J Neurosci* 3:532–548
4. Andersen RA, Snyder LH, Bradley DC, Xing J (1997) Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annu Rev Neurosci* 20:303–330
5. Asaad WF, Rainer G, Miller EK (1998) Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21:1399–1407
6. Avillac M, Deneve S, Olivier E, Pouget A, Duhamel JR (2005) Reference frames for representing visual and tactile locations in parietal cortex. *Nat Neurosci* 8:941–949
7. Batista AP, Buneo CA, Snyder LH, Andersen RA (1999) Reach plans in eye-centered coordinates. *Science* 285:257–260
8. Bentley N, Salinas E (2007) A general, flexible decision model applied to visual search *BMC Neuroscience* 8(Suppl 2):P30, doi: 10.1186/1471-2202-8-S2-P30
9. Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529–534
10. Botvinick M, Watanabe T (2007) From numerosity to ordinal rank: a gain-field model of serial order representation in cortical working memory. *J Neurosci* 27:8636–8642
11. Britten KH, Heuer HW (1999) Spatial summation in the receptive fields of MT neurons. *J Neurosci* 19:5074–5084.
12. Brotchie PR, Andersen RA, Snyder LH, Goodman SJ (1995) Head position signals used by parietal neurons to encode locations of visual stimuli. *Nature* 375:232–235
13. Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. *Nature* 416:632–636

14. Connor CE, Preddie DC, Gallant JL and Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214
15. Cox DD, Meier P, Oertelt N, DiCarlo JJ (2005) ‘Breaking’ position-invariant object recognition. *Nat Neurosci* 8:1145–1147
16. Deneve S, Latham PE, Pouget A (2001) Efficient computation and cue integration with noisy population codes. *Nat Neurosci* 4:826–831
17. Desimone R, Albright TD, Gross CG, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4:2051–2062
18. DiCarlo JJ, Cox DD (2007) Untangling invariant object recognition. *Trends Cogn Sci* 11:333–341
19. Duhamel J-R, Colby CL, Goldberg ME (1992) The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255:90–92
20. Fitzpatrick D (2000) Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol* 10:438–443.
21. Ghose GM, Maunsell JH (2008) Spatial summation can explain the attentional modulation of neuronal responses to multiple stimuli in area V4. *J Neurosci* 28:5115–5126
22. Graziano MSA, Hu TX, Gross CG (1997) Visuospatial properties of ventral premotor cortex. *J Neurophysiol* 77:2268–2292
23. Grossberg S, Mingolla E, Ross WD (1994) A neural theory of attentive visual search: interactions of boundary, surface, spatial, and object representations. *Psychol Rev* 101:470–489
24. Jay MF, Sparks DL (1984) Auditory receptive fields in primate superior colliculus shift with changes in eye position. *Nature* 309:345–347
25. Knierim JJ, Van Essen DC (1992) Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J Neurophysiol* 67:961–980.
26. Koida K, Komatsu H (2007) Effects of task demands on the responses of color-selective neurons in the inferior temporal cortex. *Nat Neurosci* 10:108–116
27. Lauwereyns J, Sakagami M, Tsutsui K, Kobayashi S, Koizumi M, Hikosaka O (2001) Responses to task-irrelevant visual features by primate prefrontal neurons. *J Neurophysiol* 86:2001–2010
28. Levitt JB, Lund JS (1997) Contrast dependence of contextual effects in primate visual cortex. *Nature* 387:73–76
29. Mack A, Rock I (1998) *Inattentive Blindness*. MIT Press, Cambridge, Massachusetts.
30. Martínez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751
31. McAdams CJ, Maunsell JHR (1999) Effects of attention on orientation tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441
32. Navalpakkam V, Itti L (2007) Search goal tunes visual features optimally. *Neuron* 53:605–617
33. Ogawa T, Komatsu H (2004) Target selection in area V4 during a multidimensional visual search task. *J Neurosci* 24:6371–6382
34. O’Reilly RC, Busby RS (2002) Generalizable relational binding from coarse-coded distributed representations. In: Dietterich TG, Becker S, Ghahramani Z (eds) *Advances in neural information processing systems (NIPS)*. MIT Press, Cambridge, MA
35. O’Reilly RC, Busby RS, Soto R (2003) Three forms of binding and their neural substrates: alternatives to temporal synchrony. In: Cleermans A (ed) *The unity of consciousness: binding, integration and dissociation*. OUP, Oxford, pp. 168–192
36. Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238
37. Pouget A, Sejnowski TJ (1997) Spatial transformations in the parietal cortex using basis functions. *J Cogn Neurosci* 9:222–237
38. Pouget A, Snyder LH (2000) Computational approaches to sensorimotor transformations. *Nat Neurosci* 3:1192–1198
39. Riesenhuber M, Poggio T (1999) Hierarchical models of object recognition in cortex. *Nat Neurosci* 2:1019–1025
40. Salinas E (2004) Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *J Neurosci* 24:1113–1118

41. Salinas E (2004) Context-dependent selection of visuomotor maps. *BMC Neurosci* 5:47, doi: 10.1186/1471-2202-5-47.
42. Salinas E (2005) A model of target selection based on goal-dependent modulation. *Neurocomputing* 65-66C:161–166
43. Salinas E, Abbott LF (1995) Transfer of coded information from sensory to motor networks. *J Neurosci* 15:6461–6474
44. Salinas E, Abbott LF (1997) Invariant visual responses from attentional gain fields. *J Neurophysiol* 77:3267–3272
45. Salinas E, Abbott LF (1997) Attentional gain modulation as a basis for translation invariance. In: Bower J (ed) *Computational neuroscience: trends in research 1997*. Plenum, New York, pp. 807–812
46. Salinas E, Abbott LF (2001) Coordinate transformations in the visual system: how to generate gain fields and what to compute with them. *Prog Brain Res* 130:175–190
47. Salinas E, Sejnowski TJ (2001) Gain modulation in the central nervous system: where behavior, neurophysiology, and computation meet. *Neuroscientist* 7:430–40
48. Salinas E, Thier P (2000) Gain modulation: a major computational principle of the central nervous system. *Neuron* 27:15–21
49. Sato M, Hikosaka O (2002) Role of primate substantia nigra pars reticulata in reward-oriented saccadic eye movement. *J Neurosci* 22:2363–2373
50. Shenoy KV, Bradley DC, Andersen RA (1999) Influence of gaze rotation on the visual response of primate MSTd neurons. *J Neurophysiol* 81:2764–2786
51. Sillito AM, Grieve KL, Jones HE, Cudeiro J, Davis J (1995) Visual cortical mechanisms detecting focal orientation discontinuities. *Nature* 378:492–496
52. Simoncelli EP, Heeger DJ (1998) A model of neuronal responses in visual area MT. *Vision Res* 38:743–761
53. Schwartz O, Simoncelli EP (2001) Natural signal statistics and sensory gain control. *Nat Neurosci* 4:819–825
54. Sohn JW, Lee D (2007) Order-dependent modulation of directional signals in the supplementary and presupplementary motor areas. *J Neurosci* 27:13655–13666
55. Stricanne B, Andersen RA, Mazzoni P (1996) Eye-centered, head-centered, and intermediate coding of remembered sound locations in area LIP. *J Neurophysiol* 76:2071–2076
56. Tovee MJ, Rolls ET, Azzopardi P (1994) Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. *J Neurophysiol* 72:1049–1060
57. Treue S, Martínez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579
58. Wallis JD, Miller EK (2003) From rule to response: neuronal processes in the premotor and prefrontal cortex. *J Neurophysiol* 90:1790–1806
59. White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res* 126:315–335
60. Wolfe JM (1994) Guided Search 2.0 – a revised model of visual search. *Psychonomic Bul Rev* 1:202–238
61. Wolfe JM, Horowitz TS, Kenner N, Hyle M, Vasan N (2004) How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Res* 44:1411–1426
62. Xing J, Andersen RA (2000) Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *J Cogn Neurosci* 12:601–614
63. Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684
64. Zoccolan D, Kouh M, Poggio T, DiCarlo JJ (2007) Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. *J Neurosci* 27:12292–12307