

Adaptation in the Corticothalamic Loop: Computational Prospects of Tuning the Senses

Ulrich Hillenbrand
Institute of Robotics and Mechatronics
German Aerospace Center
Oberpfaffenhofen, 82234 Wessling, Germany

J. Leo van Hemmen
Physik Department
TU München
85747 Garching bei München, Germany

Abstract

The present article discusses computational hypotheses on corticothalamic feedback and modulation of cortical response properties. We have recently proposed the two phenomena to be related, hypothesizing that neuronal velocity preference in the visual cortex is altered by feedback to the lateral geniculate nucleus. We now contrast the common view that response adaptation to stimuli subserves a function of redundancy reduction with the idea that it may enhance cortical representation of objects. Our arguments lead to the concept that the corticothalamic loop is involved in reducing sensory input to behaviorally relevant aspects, a pre-attentive gating.

Keywords: Corticothalamic Loop, Feedback, Adaptation, Velocity Tuning, Object Segmentation, Pre-attentive Gating

Published as *Philosophical Transactions of the Royal Society: Biological Sciences* **357** (2002), pp. 1859–1867.

1 Computational Views of Corticothalamic Feedback

The thalamus is the major gate to the cortex for peripheral sensory signals, for input from various subcortical sources, and for reentrant cortical information. Cortical information processing thus fundamentally depends on its thalamic input. Thalamic nuclei, in turn, perform some operation on the signals they relay while being modulated by various transmitter systems (McCormick, 1992)

and under the influence of massive feedback from their cortical target areas (Guillery, 1995; Sherman, 1996; Sherman & Guillery, 1996). In view of the complexity of the thalamocortical network, one is compelled to believe that the thalamus not merely *delivers* information to the cortex but that cortical processing is somehow *entangled* with thalamic processing. Accordingly, it seems natural that cortical functioning cannot be properly understood without taking into account the corticothalamic feedback loop.

Whereas it is well established that the response characteristics of many neurons in primary sensory cortices are roughly shaped by different types of convergent thalamic inputs (Saul & Humphrey, 1992a; Saul & Humphrey, 1992b; Reid & Alonso, 1995; Alonso et al., 1996; Ferster et al., 1996; Jagadeesh et al., 1997; Murthy et al., 1998; Hirsch et al., 1998), the modulation effected by cortical feedback in thalamic response has been difficult to interpret. Nonetheless, experiments and theoretical considerations have produced a variety of views of corticothalamic function. Most of them are concerned with the primary visual pathway, that is, the lateral geniculate nucleus (LGN) and the primary visual cortex.

1.1 Gating and Enhancement of Neural Response

The ideas put forward to date are conceptually mainly of two kinds. On the one hand one offers an open-loop understanding of corticogeniculate feedback, proposing an operation of the visual cortex on the LGN, while deferring its control to some other, more complex process not considered. Examples of such computational concepts are attention-related gating of geniculate relay cells (GRCs) (Sherman & Koch, 1986), gain control of GRCs (Koch, 1987; Rivadulla et al., 2002), increasing transmitted information in GRC output (McClurkin et al., 1994), and switching GRCs from a detection to an analyzing mode (Godwin et al., 1996; Sherman, 1996; Sherman & Guillery, 1996). On the other hand one argues in favor of a positive or negative feedback signal for responses in the corticogeniculate loop. Its effect would be to amplify certain response characteristics such as length (Murphy & Sillito, 1987; Rivadulla et al., 2002) and orientation tuning (Sillito et al., 1994; Singer, 1994; Murphy et al., 1999). A positive feedback loop has also been proposed to underlie stabilization of learning in the geniculocortical pathways (Grunewald & Grossberg, 1998).

With rather vague open-loop and simple feedback concepts of corticothalamic operation, it seems likely that today, after almost thirty years of studying the subject (Singer, 1977), there is still a long way to go to arrive at an integrated view of thalamocortical functioning. We believe that computational modeling, although to some extent speculative, can play an important role in this endeavor, as it does in understanding brain function in general (van Hemmen & Sejnowski, 2003). Experimental data in neuroscience are interpreted functionally in the light of computational metaphors. It is thus crucial to continuously refine the repertoire of the metaphors we use. Computational modeling will lead to the development of more sophisticated functional concepts to be matched with empirical data.

1.2 Tuning of Neural Response Preference

As compared with corticogeniculate feedback, our understanding of receptive fields (RFs) in the primary visual cortex is relatively advanced. It may thus be a revealing perspective to try to understand corticogeniculate feedback functionally in terms of its modulation effects on cortical RFs.

Recently, we have elaborated the idea that the visual cortex controls, via feedback, the temporal response properties of GRCs in a way that alters, in turn, the preference of cortical cells for stimulus speed (Hillenbrand & van Hemmen, 2000; Hillenbrand & van Hemmen, 2001; Hillenbrand, 2001). Figure 1 shows the model of the primary visual pathway we have studied. In brief, we have investigated a model of cat GRCs that comprises 12 ionic membrane conductances, giving rise to the typical thalamic patterns of activity such as tonic and burst components of responses (Huguenard & McCormick, 1992; McCormick & Huguenard, 1992). We have subjected these model neurons to direct excitatory and indirect inhibitory synaptic input (cf. Fig. 1) as recorded from the retina during visual stimulation. We have found that, just by varying the relative strengths of excitation and inhibition, the temporal response of GRCs split into two distinct types, one resembling a nonlagged, the other a lagged response (Mastronarde, 1987; Humphrey & Weller, 1988; Saul & Humphrey, 1990); see Fig. 2 for some simulated responses. Moreover, the two response types react to changes in their membrane potential in an opposite manner. With hyperpolarization, the nonlagged-type neurons shift their peak response to a moving bar to earlier times, whereas the lagged-type neurons shift their peak response to later times; see Fig. 2. Numerous studies (Saul & Humphrey, 1992a; Saul & Humphrey, 1992b; DeAngelis et al., 1995; Jagadeesh et al., 1997; Murthy et al., 1998) indicate that GRCs of the two types converge, directly or indirectly, onto the same neurons in the primary visual cortex and lay out the basic spatiotemporal structure of simple cells' RFs. Consequently, the preference of cortical neurons for moving local features (bars, contrast edges, dots etc.) is predicted to shift to *lower speeds* as GRCs are *hyperpolarized* and to *higher speeds* as GRCs are *depolarized*; see Fig. 3.

The effect of both feedforward inhibition and resting membrane potential on GRCs' response timing is mediated largely by the low-threshold Ca^{2+} current I_T and its tendency to produce burst spikes after sufficient hyperpolarization (Huguenard & McCormick, 1992; McCormick & Huguenard, 1992). With only weak feedforward inhibition, nonlagged neurons respond to retinal input with immediate depolarization, eventually reaching the activation threshold for the Ca^{2+} current. If the Ca^{2+} current is in the de-inactivated state, it will boost depolarization and give rise to an early burst component of the visual response. The lower the resting membrane potential, the more de-inactivated and, hence, stronger the Ca^{2+} current will be, and the stronger the early burst relative to the late tonic response component. Lagged neurons, on the other hand, receive strong feedforward inhibition and, hence, initially respond to retinal input with hyperpolarization. Repolarization occurs when inhibition gets weaker. This may result either from cessation of retinal input or from adaptation, i.e., fatigue, of

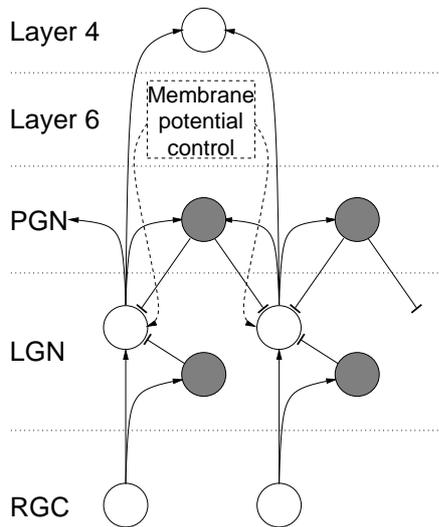


Figure 1: Network model of the primary visual X-pathway we have studied. Open/filled circles and arrow/bar heads indicate excitatory/inhibitory neurons and their respective synapses. A retinal X-ganglion cell (RGC) sends its axon to the lateral geniculate nucleus (LGN) and synapses excitatorily on an X-relay cell (open circle) and on an intrageniculate interneuron (filled circle), which in turn inhibits the same relay cell. This synaptic circuit is called ‘synaptic triad’. There is an inhibitory feedback loop via the perigeniculate nucleus (PGN). The relay cells’ output converges in cortical layer 4, where it lays out the basic spatiotemporal structure of simple cells’ receptive fields. The influence of cortical feedback, originating from cortical layer 6, has been modeled as a modulation of the relay cells’ membrane potential. For the complete set of equations describing the model see the appendices in Hillenbrand (2001).

the inhibitory input to GRCs. With the Ca^{2+} current I_T being de-inactivated by the excursion of the membrane potential to low values, lagged spiking starts with burst spikes as soon as the voltage reaches the Ca^{2+} -activation threshold. This will take longer, if the resting membrane potential is lower, leading to the shift in response timing with membrane polarization observed during the simulations.

Modulation of the GRCs’ membrane potential by cortical feedback can result directly from the action of metabotropic glutamate and NMDA receptors, leading to depolarization (McCormick & von Krosigk, 1992; Godwin et al., 1996; Sherman, 1996; Sherman & Guillery, 1996; von Krosigk et al., 1999), and indirectly, via the perigeniculate nucleus (PGN) or geniculate interneurons, from activation of GABA_B receptors, resulting in hyperpolarization of GRCs (Crunelli & Leresche, 1991; Sherman & Guillery, 1996; von Krosigk et al., 1999). Moreover, GRCs may be indirectly depolarized by the action of group-II metabotropic glutamate receptors on PGN neurons (Cox & Sherman, 1999).

For the predicted effect of feedback on cortical velocity tuning no direct experimental evidence is available yet. Dynamic velocity preference, however, has indeed been recently observed in the visual cortex (Li et al., 1999; Phinney

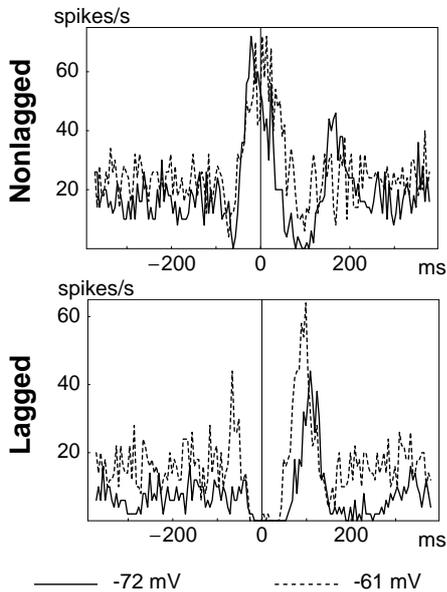


Figure 2: Computer simulation of moving-bar response of a single relay neuron at the two resting membrane potentials of -72 mV (solid lines) and -61 mV (dashed lines). Typical nonlagged responses are reproduced with strong feedforward excitation and weak feedforward inhibition, typical lagged responses with weak feedforward excitation and strong feedforward inhibition; cf. Fig. 1. The time of the retinal input peak has been set to zero. It turns out that the nonlagged bar-response peak shifts to earlier times as the membrane is hyperpolarized. Conversely, the lagged bar response shifts to later times. The rates are averaged over 100 bar sweeps. Figure adapted from Hillenbrand & van Hemmen (2001).

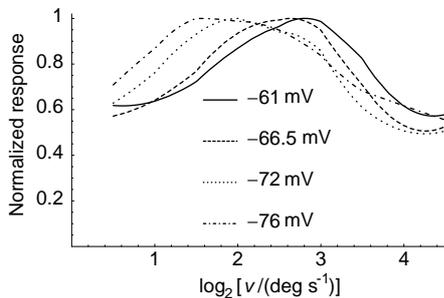


Figure 3: Computer simulation of geniculate moving-bar response and input to visual cortex at the geniculate resting membrane potentials indicated. We show second-order polynomial fits to the normalized mean response rate (30 bar sweeps) of the combined lagged and nonlagged neural populations as a function of the bar velocity v (logarithmic, base 2). It turns out that the total geniculate response peaks at progressively lower velocities as relay cells are hyperpolarized. Figure adapted from Hillenbrand (2001).

& Siegel, 2000).

As a result of adaptive velocity preference, the cortex would be able to tune into certain parts of a dynamic stimulus. We have developed a functional hypothesis of *velocity-preference tuning* as a mechanism to detect and represent coherent, object-related motion from the nonlocal velocity distribution of a stimulus (Hillenbrand & van Hemmen, 2000; Hillenbrand, 2001). In Sec. 2 we will present this tuning idea in the context of others about adaptive changes of cortical representation.

2 Computational Views of Adaptive Neural Responses

It is instructive to gain a somewhat broader perspective of adaptive cortical response properties. To this end, we consider possible scenarios of adaptation to certain stimulus conditions. We then go on to discuss the example of velocity-preference tuning in a specific functional manner. In particular, we contrast our proposal of *coherence representation* with the popular view of *redundancy reduction* as a short-term adaptive process.

2.1 Four Generic Scenarios of Response Adaptation

Perceptual repulsion is a phenomenon that is characteristically associated with adaptation to certain stimulus conditions. After having viewed a grating of a fixed orientation for an extended period of time (roughly a minute or longer), subsequently presented gratings of similar orientation are perceived as tilted away from the original, adapting orientation. This is the tilt after-effect (Gibson & Radner, 1937; Blakemore & Campbell, 1969; Campbell & Maffei, 1971). A similar effect of perceptual repulsion is induced by presenting a grating together with a second one at a slightly oblique angle, known as the tilt illusion (Blakemore et al., 1970; Smith & Wenderoth, 1999). Analogous repulsion is also known for size (Blakemore & Campbell, 1969), stereoscopic depth (Stevenson et al., 1991), and velocity perception (Clymer, 1973; Marshak & Sekuler, 1979; Smith & Edgar, 1994; Schrater & Simoncelli, 1998). Perceptual repulsion hence seems to be a rather universal pattern of interaction between stimuli, in both space and time. When analyzing adaptation of neural representation to stimulus conditions, it may thus be worthwhile to keep an eye on this effect.

We now consider four generic scenarios of changes of cortical representation in response to an adapting or inducing stimulus. We are interested in how they relate to perceptual repulsion. The scenarios are

- depression of responses of neurons preferring the adapting or inducing stimulus,
- facilitation of responses of neurons preferring the adapting or inducing stimulus,
- repulsion of neural preferences from the adapting or inducing stimulus, and
- attraction of neural preferences towards the adapting or inducing stimulus.

Let s be the orientation of a grating, the speed of a moving stimulus, or some other stimulus property within an appropriate range. Furthermore, let r_i be the response rate of neuron i belonging to a population $1 \leq i \leq n$. The neuron's response is determined by its tuning characteristic $r_i = g(s - p_i)$, where g is taken to be a Gaussian function and p_i is the neuron's stimulus preference. The

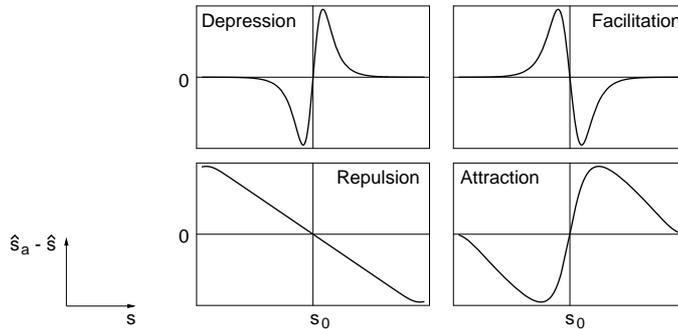


Figure 4: Four scenarios of response adaptation; see main text for details. For each scenario, the plot shows the difference $\hat{s}_a - \hat{s}$ between the perceived stimulus values for the adapted and unadapted neural populations as a function of the test-stimulus value s . The vertical line in each plot represents the adapting or inducing stimulus s_0 , the horizontal line indicates a perceptual difference of zero. Perceptual repulsion is seen to occur for the depression and attraction scenarios.

perceived stimulus \hat{s} is assumed to be recovered from a neural population code through

$$\hat{s}(s) = \frac{\sum_{i=1}^n r_i p_i}{\sum_{j=1}^n r_j} = \frac{\sum_{i=1}^n g(s - p_i) p_i}{\sum_{j=1}^n g(s - p_j)}, \quad (1)$$

the perceived value being the average of all neural preferences in the population, weighted by each neuron's response. In the case of a cyclic stimulus dimension such as orientation, s may be chosen to parameterize points on a unit circle. The qualitative results presented here, however, do not depend on this choice. Likewise, absolute values of s do not have any meaning and are thus omitted.

Let \hat{s} be the perceived value for an unadapted, \hat{s}_a for an adapted neural population. In Fig. 4 we show plots of the perceptual difference $\hat{s}_a - \hat{s}$ as a function of the test stimulus s for the four adaptation scenarios above. Evidently, response depression and attraction of neural preferences are both consistent with the phenomenon of perceptual repulsion whereas response facilitation and repulsion of neural preferences are not. Thus it would seem that recent reports of repulsive shifts in neural preferences away from an adapting orientation (Dragoi et al., 2000; Dragoi et al., 2001) do not directly relate to the mechanism underlying the tilt after-effect. Likewise, reported shifts of neural preferences away from a moving inducer's speed (Li et al., 1999) do not contribute to, but rather diminish perceptual velocity repulsion.

We note that there is a lot of free parameters in this very simple model, e.g., the shape of tuning curves, the number of neurons in the population, and the form of depression or of shifts in neural preferences. Moreover, although population coding is a common view of neural representation (Salinas & Abbott, 1994; Lewis, 1999), any particular form such as Eq. (1) is debatable. All these

parameters can influence the shape of the perceptual curves in Fig. 4. Here, however, we are only interested in the qualitative perceptual effect, which does not depend critically on these parameters.

It may be worth mentioning that another perceptual phenomenon that accompanies repulsion is in fact dependent on it. Discrimination of orientation is known to be enhanced near the adapting orientation, impaired further away from it, and back to normal even further. This is explained easily by taking the derivative of the perceptual curves of Fig. 4 for the two cases that show the repulsion effect. The derivative is the perceptual change per unit of stimulus change and hence describes discrimination of the stimulus. Discrimination is higher than normal where the curves of the figure have a positive slope and lower where it is negative.

Perceptual repulsion thus turns out to be in agreement with two generic types of adaptation that, interestingly, have a *contrary* effect on correlations between neural activity: response depression that *decorrelates* cortical responses, and preference attraction that *enhances correlations*. The decorrelation view is one of the major themes in understanding adaptation effects and has received some attention in the literature. We shall next discuss the rationale underlying this view and then turn to correlation enhancement by preference attraction.

2.2 Tuning Representations for Information Maximization

Over the past ten years, information maximization has become one of the leading paradigms to understand neural representation in primary sensory cortices, especially the primary visual cortex, in terms of an optimal coding strategy. According to this view, neural responses are transformed to reduce their redundancy (Barlow, 1989; Atick, 1992; Bell & Sejnowski, 1997), perhaps also to increase their sparseness (Field, 1994; Olshausen & Field, 1996), given the statistics of natural stimuli. The cortical network is usually believed to adapt to this statistics by long-term synaptic plasticity. The result would be a maximization of the information represented about a sensory stimulus, given certain bounds on neuronal resources. A similar approach has also been adopted for understanding transformations in subcortical pathways (Atick & Redlich, 1992; Dong & Atick, 1995; Haft & van Hemmen, 1998; Lewicki, 2002).

On a much shorter timescale, analogous principles of network plasticity or response modulation are often envisaged to underlie adaptation to specific stimuli as described in Sec. 2.1 (Barlow, 1990; Wainwright, 1999; Muller et al., 1999; Bednar & Miikkulainen, 2000). One possibility would be a strengthening of mutual inhibition between responsive neurons so as to depress any further coincident response and decrease correlations that are induced by that stimulus.

Maximization of the encoded information is a reasonable goal for a channel that transmits messages to a receiver. As plausible as this strategy is for a transmitting channel, it is also evident that this cannot be the ultimate goal of sensory processing since sensory processing must eventually lead to behavior. Behavior itself is unlikely to preserve the maximum amount of information on the total sensory input. At the end, only what is behaviorally relevant is coded,

hence, information is *selected for its behavioral relevance*. The same probably holds for the formation of memory. The principle of behavioral relevance, however, is not captured by maximization of information on elementary stimulus properties. Moreover, it has recently been shown empirically that the intrinsic complexity of a visual recognition task can increase significantly by transforming to statistically less dependent features, even without noise, leading to lower precision in recognition performance (Vasconcelos & Carneiro, 2002). It is, therefore, not obvious that neural representations anywhere in the brain are optimized for redundancy reduction or indeed for maximizing the information concerning a sensory stimulus.

In fact, whether elementary sensory information is maximized depends on the *separability* of the brain's processes in *universal representation stages* on the one hand, and more *behavior-related recognition stages* on the other. This two-step strategy is a common design principle for autonomous robotic systems, as it makes the design conceptually transparent. How far this strategy can explain the brain as a biologically evolved system remains an empirical question; cf. Cliff & Noble (1997). The success of explaining RF structure of neurons in the visual pathway up to simple cells by infomax-like coding principles suggests that some separability exists. One should not expect, however, such principles to prevail throughout the brain; see, e.g., Krahe et al. (2002).

There is a particular problem with viewing short-term adaptation as a redundancy-reducing process. The behaviorally relevant parts of a stimulus are usually marked by some sort of contingent or coherent structure in the input, that is, some redundancy that exceeds, for a short period of time, the long-term average level. This coherence is brought about by the coherence of objects in the outside world.

For instance, motion of a rigid body induces spatiotemporally correlated signals of local-motion encoding neurons and additional correlations with other stimulus dimensions such as color. Fading out the associated redundancy within typical object recognition times (hundreds of milliseconds) by rapidly adapting the neural code would render objects invisible in the neural representation. Hence, the behaviorally most relevant aspects of a stimulus would go unnoticed. It therefore seems evident that any redundancy-reducing adaptation can only operate on a longer timescale than detection of redundant patterns. Moreover, if we do not believe in the rapid formation of a dedicated, behaviorally effective 'grandmother' cell for each contingent conjunction of all correlated stimulus dimensions, it follows that somewhere in the brain a perceived object must be retained as a highly *correlated* representation.

2.3 Tuning Representations to Coherent Patterns

We now turn to the alternative scenario of adaptation from Sec. 2.1, the attraction of neural preferences towards certain stimulus parameters. We discuss the case of velocity-preference tuning and its possible role in detection and representation of coherent motion. In the light of the results presented by Hillenbrand & van Hemmen (2000, 2001) and Hillenbrand (2001) and reviewed in Sec. 1.2,

we propose that this function is supported by corticothalamic loops. As pointed out in Sec. 2.2, detection and representation of coherence should operate on a shorter timescale than any decorrelating mechanism of adaptation.

We argue that the computational goal of velocity-preference tuning is an enhanced representation of behaviorally relevant aspects of a stimulus, generally referred to as *objects* and characterized by coherence, and in conjunction with this suppression of less significant aspects, such as neuronal noise and incoherent background motion. The general idea is illustrated in Fig. 5. It agrees with the preference-attraction scenario of adaptation that is discussed in Sec. 2.1.

In Hillenbrand & van Hemmen (2000) and Hillenbrand (2001) we have analyzed a model of the corticogeniculate loop that produces the adaptive behavior that is sketched in Fig. 5. When the system is stimulated by a coherently moving object against a background of incoherent, directionally unbiased motion, the neural preferences approach the velocity of the object. This results in an enhanced representation of the object and suppression of background features moving at different velocities; see Fig. 6.

More specifically, without a background or with just a weak background component in the stimulus, the neural preferences settle in a stationary state close to the object's velocity. As more background is added to the stimulus, persistent oscillations of the neurons' preferences develop. It is important to notice that neuronal background activity, or noise, has the same effect on the adaptation dynamics as incoherent motion in the visual stimulus. The oscillatory time course of adaptation is associated with alternating phases of weak and strong population responses, the strong responses being restricted to object features; cf. Fig. 6. The dynamic neural preferences thus act as a pacemaker for distributed cortical activity. A periodic time structure is imposed that tends to *synchronize* the firing of cortical cells representing the object. We note that, due to superposition of depolarizing and hyperpolarizing effects of cortical feedback on GRCs, the oscillation period can be shorter than the duration and even the rise time of geniculate postsynaptic potentials. The rise time τ is a free scale parameter in the model. If τ is taken to be 100 ms, the oscillation frequency is roughly 25 Hz in the example shown in Fig. 6.

The model presents an integrated, closed-loop view of the geniculocortical system performing the high-level task of object segmentation. One may, however, ask: is there, apart from theoretical motivation and appeal, any *direct* evidence that similar mechanisms operate in the real system? As far as we know, there is none yet. The model rather offers a template for an alternative computational principle that one may look for when interpreting empirical data on response modulation in functional terms.

Indirect experimental evidence for the model may be seen in diverse phenomena such as visual-motion-induced oscillations of cortical synaptic potentials (Bringuier et al., 1997), visual-perceptual stabilization during fixation and the jitter after-effect (Murakami & Cavanagh, 1998), percepts of relative motion (Kaufman, 1974; Van Essen & Anderson, 1990), perceptual velocity repulsion (Clymer, 1973; Smith & Edgar, 1994; Schrater & Simoncelli, 1998), and interaction between stimulus density and speed perception (Watamaniuk et al., 1993),

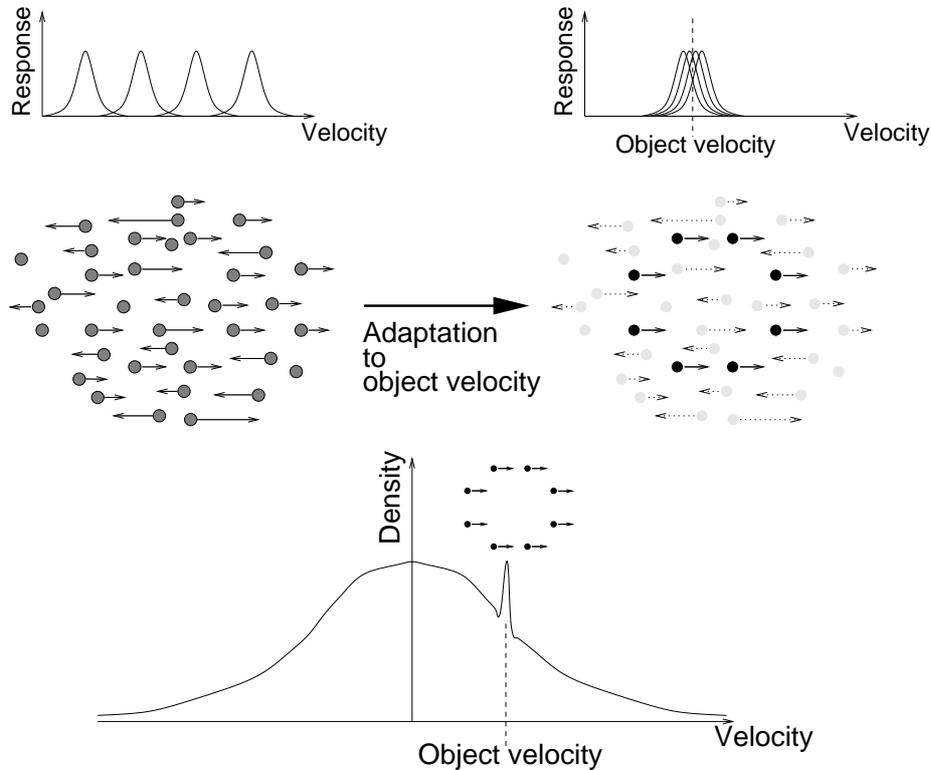


Figure 5: The rationale of object segmentation by velocity-preference tuning. The curves shown at the top of the figure schematically represent the response rates of four cortical neurons as a function of the velocity of a local feature (edge, texture element, etc.) passing their receptive field in an unadapted state (top left) and an adapted state (top right). Within some region of the visual field a stimulus (center) consists of a collection of local features (depicted as dots) moving from left to right and from right to left at various velocities (depicted as arrows). A subset of them is moving at a *common* velocity from left to right. The velocity density (bottom) of this type of stimulus consists of two components: one symmetrical and one asymmetrical with respect to the two directions of motion. The former derives from the incoherent, the latter from the coherent motion and is the statistical signature of a moving object. The adaptive motion system has to detect and tune in to the asymmetrical component of the velocity density. After adaptation of cortical velocity preferences (top right), object features are prominent in cortical representation, whereas other features are suppressed (center right). The shown stimulus scenario generalizes straightforwardly to motion in two dimensions. Figure adopted from Hillenbrand & van Hemmen (2000).

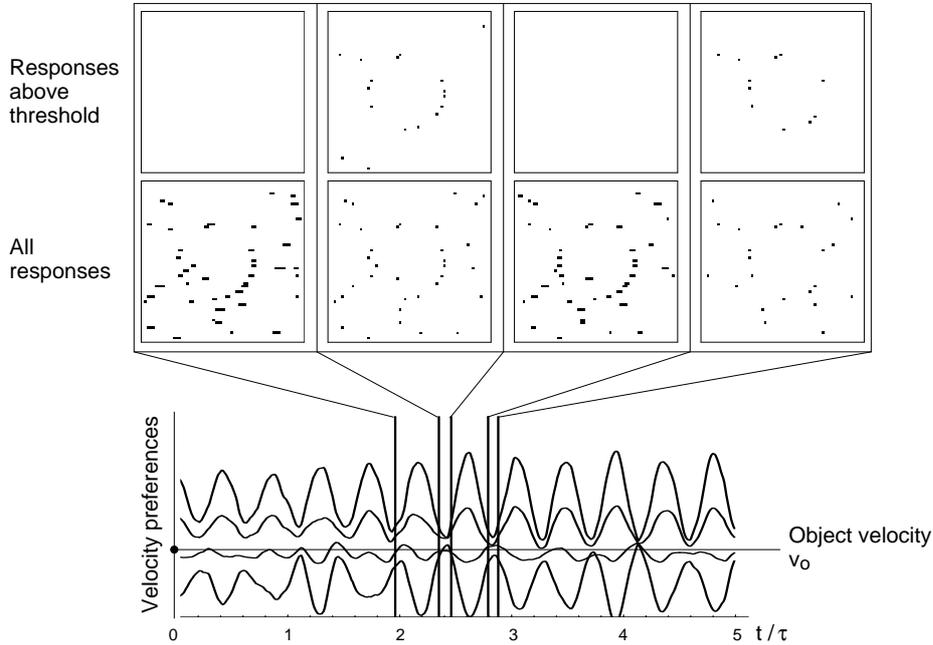


Figure 6: Adaptation dynamics and spatiotemporal pattern of cortical activity in a computer simulation of the corticogeniculate loop. The system is stimulated by 15 dots placed on a circle, all moving at a *common* velocity v_o to the right (object), and 30 dots moving at various random velocities, exceeding or less than the object velocity v_o , either to the left or to the right (background); cf. Fig. 5. The dynamics of velocity preferences of four representative cortical neurons are shown at the bottom of the figure (unit of time on the horizontal axis is τ , the rise time of corticothalamic synaptic potentials; the vertical velocity axis has been scaled nonlinearly). The adaptation time-course is oscillatory because of the strong stimulus background. Whenever the neurons' preferences get *close* to the object velocity v_o (horizontal line), object dots elicit a strong response in all cortical neurons that represent their actual retinal positions. The corresponding array of population activity is shown in the panels at the top of the figure. The lower row of panels displays *all* responses that occur within the indicated intervals of time; the upper row displays only responses that *exceed* a certain threshold. Nearly all suprathreshold activity is related to object features, demonstrating the segmentation of the object against the background. Spurious suprathreshold activity derives from responses to background dots that by chance have velocities very close to the object's velocity. The suprathreshold activity is oscillatory and synchronous like the adaptation time-course. Figure adapted from Hillenbrand & van Hemmen (2000).

all discussed in Hillenbrand (2001).

3 Conclusion

We may still have a long way to go to arrive at a computational understanding of the thalamocortical system that is scaled to its complexity and performance. The concept of object segmentation by tuning in neural preferences has been inspired, in part, by the need to get closer to the processing of behaviorally relevant information in an integrated thalamocortical model.

The modulation of cortical responses in the spatial and temporal context of a stimulus may have various effects on different timescales and levels of representation. The ultimate goal of sensory coding, however, is to support behavior. Accordingly, correlated, object-related patterns must be picked up from the sensory input and persistently represented somewhere in the brain. Hence, we have questioned the computational plausibility of redundancy reduction as an effect on a very short timescale or across all levels of sensory representation.

We have instead argued that responses are modulated so as to enhance object-related correlations and suppress incoherent stimulus aspects, initially at lower or universal levels and persistently at higher or behavior-related levels of representation. A modulation at the lower levels may be described as a pre-attentive gating. Object segmentation through velocity-preference tuning is a concrete realization, which we propose to be supported by corticothalamic loops.

Acknowledgments

We thank S. M. Sherman for stimulating discussions.

References

- Alonso, J. M., Usrey, W. M., & Reid, R. C. (1996). Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature*, 383, 815–819.
- Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing? *Network*, 3, 213–251.
- Atick, J. J. & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Comput.*, 4, 196–210.
- Barlow, H. B. (1989). Unsupervised learning. *Neural Comput.*, 1, 295–311.
- Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 363–375). New York: Cambridge University Press.
- Bednar, J. A. & Miikkulainen, R. (2000). Tilt aftereffects in a self-organizing model of the primary visual cortex. *Neural Comput.*, 12, 1721–1740.

- Bell, A. J. & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Res.*, 37, 3327–3338.
- Blakemore, C. & Campbell, F. (1969). On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images. *J. Physiol. (Lond.)*, 203, 237–260.
- Blakemore, C., Carpenter, R. H. S., & Georgeson, M. A. (1970). Lateral inhibition between orientation detectors in the human visual system. *Nature*, 228, 37–39.
- Bringuier, V., Fregnac, Y., Baranyi, A., Debanne, D., & Shulz, D. E. (1997). Synaptic origin and stimulus dependency of neuronal oscillatory activity in the primary visual cortex of the cat. *J. Physiol. (Lond.)*, 500, 751–774.
- Campbell, F. W. & Maffei, L. (1971). The tilt after-effect: a fresh look. *Vision Res.*, 11, 833–840.
- Cliff, D. & Noble, J. (1997). Knowledge-based vision and simple visual machines. *Phil. Trans. R. Soc. B*, 352, 1165–1175.
- Clymer, A. B. (1973). *The Effect of Seen Movement on the Apparent Speed of Subsequent Test Velocities: Speed Tuning of Movement*. PhD thesis, Columbia University, NY.
- Cox, C. L. & Sherman, S. M. (1999). Glutamate inhibits thalamic reticular neurons. *J. Neurosci.*, 19, 6694–6699.
- Crunelli, V. & Leresche, N. (1991). A role for gabab receptors in excitation and inhibition of thalamocortical cells. *Trends Neurosci.*, 14, 16–21.
- DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends Neurosci.*, 18, 451–458.
- Dong, D. W. & Atick, J. J. (1995). Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network*, 6, 159–178.
- Dragoi, V., Rivadulla, C., & Sur, M. (2001). Foci of orientation plasticity in visual cortex. *Nature*, 411, 80–86.
- Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron*, 28, 287–298.
- Ferster, D., Chung, S., & Wheat, H. (1996). Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature*, 380, 249–252.
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Comput.*, 6, 559–601.

- Gibson, J. & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines: I. Quantitative studies. *J. Exp. Psychol.*, 20, 453–467.
- Godwin, D. W., Vaughan, J. W., & Sherman, S. M. (1996). Metabotropic glutamate receptors switch visual response mode of lateral geniculate nucleus cells from burst to tonic. *J. Neurophysiol.*, 76, 1800–1816.
- Grunewald, A. & Grossberg, S. (1998). Self-organization of binocular disparity tuning by reciprocal corticogeniculate interactions. *J. Cog. Neurosci.*, 10, 199–215.
- Guillery, R. W. (1995). Anatomical evidence concerning the role of the thalamus in corticocortical communication: A brief review. *J. Anat.*, 187, 583–592.
- Haft, M. & van Hemmen, J. L. (1998). Theory and implementation of infomax filters for the retina. *Network*, 9, 39–71.
- Hillenbrand, U. (2001). *Spatiotemporal Adaptation in the Corticogeniculate Loop*. Doctoral thesis, Technische Universität München. <http://tumb1.biblio.tu-muenchen.de/publ/diss/ph/2001/hillenbrand.pdf>.
- Hillenbrand, U. & van Hemmen, J. L. (2000). Spatiotemporal adaptation through corticothalamic loops: A hypothesis. *Vis. Neurosci.*, 17, 107–118.
- Hillenbrand, U. & van Hemmen, J. L. (2001). Does corticothalamic feedback control cortical velocity tuning? *Neural Comput.*, 13, 327–355.
- Hirsch, J. A., Alonso, J.-M., Reid, R. C., & Martinez, L. M. (1998). Synaptic integration in striate cortical simple cells. *J. Neurosci.*, 18, 9517–9528.
- Huguenard, J. R. & McCormick, D. A. (1992). Simulation of the currents involved in rhythmic oscillations in thalamic relay neurons. *J. Neurophysiol.*, 68, 1373–1383.
- Humphrey, A. L. & Weller, R. E. (1988). Functionally distinct groups of X-cells in the lateral geniculate nucleus of the cat. *J. Comp. Neurol.*, 268, 429–447.
- Jagadeesh, B., Wheat, H. S., Kontsevich, L. L., Tyler, C. W., & Ferster, D. (1997). Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *J. Neurophysiol.*, 78, 2772–2789.
- Kaufman, L. (1974). *Sight and mind, an introduction to visual perception*. New York: Oxford University Press.
- Koch, C. (1987). The action of the corticofugal pathway on sensory thalamic nuclei: A hypothesis. *Neurosci.*, 23, 399–406.
- Krahe, R., Kreiman, G., Gabbiani, F., Koch, C., & Metzner, W. (2002). Stimulus encoding and feature extraction by multiple sensory neurons. *J. Neurosci.*, 22, 2374–2382.

- Lewicki, M. S. (2002). Efficient coding of natural sounds. *Nat. Neurosci.*, 5, 356–363.
- Lewis, J. E. (1999). Sensory processing and the network mechanisms for reading neuronal population codes. *J. Comp. Physiol. A*, 185, 373–378.
- Li, C. Y., Lei, J. J., & Yao, H. S. (1999). Shift in speed selectivity of visual cortical neurons: A neural basis of perceived motion contrast. *Proc. Natl. Acad. Sci. USA*, 96, 4052–4056.
- Marshak, W. M. & Sekuler, R. (1979). Mutual repulsion between moving visual targets. *Science*, 205, 1399–1401.
- Mastrorarde, D. N. (1987). Two classes of single-input X-cells in cat lateral geniculate nucleus. I. Receptive-field properties and classification of cells. *J. Neurophysiol.*, 57, 357–380.
- McClurkin, J. W., Optican, L. M., & Richmond, B. J. (1994). Cortical feedback increases visual information transmitted by monkey parvocellular lateral geniculate nucleus neurons. *Vis. Neurosci.*, 11, 601–617.
- McCormick, D. A. (1992). Neurotransmitter actions in the thalamus and cerebral cortex and their role in neuromodulation of thalamocortical activity. *Prog. Neurobiol.*, 39, 337–388.
- McCormick, D. A. & Huguenard, J. R. (1992). A model of the electrophysiological properties of thalamocortical relay neurons. *J. Neurophysiol.*, 68, 1384–1400.
- McCormick, D. A. & von Krosigk, M. (1992). Corticothalamic activation modulates thalamic firing through glutamate “metabotropic” receptors. *Proc. Natl. Acad. Sci. USA*, 89, 2774–2778.
- Muller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, 285, 1405–1408.
- Murakami, I. & Cavanagh, P. (1998). A jitter after-effect reveals motion-based stabilization of vision. *Nature*, 395, 798–801.
- Murphy, P. C., Duckett, S. G., & Sillito, A. M. (1999). Feedback connections to the lateral geniculate nucleus and cortical response properties. *Science*, 286, 1552–1554.
- Murphy, P. C. & Sillito, A. M. (1987). Corticofugal feedback influences the generation of length tuning in the visual pathway. *Nature*, 329, 727–729.
- Murthy, A., Humphrey, A. L., Saul, A. B., & Feidler, J. C. (1998). Laminar differences in the spatiotemporal structure of simple cell receptive fields in cat area 17. *Vis. Neurosci.*, 15, 239–256.

- Olshausen, B. A. & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.
- Phinney, R. E. & Siegel, R. M. (2000). Speed selectivity for optic flow in area 7a of the behaving macaque. *Cereb. Cortex*, 10, 413–421.
- Reid, R. C. & Alonso, J. M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378, 281–284.
- Rivadulla, C., Martínez, L. M., Varela, C., & Cudeiro, J. (2002). Completing the corticofugal loop: A visual role for the corticogeniculate type 1 metabotropic glutamate receptor. *J. Neurosci.*, 22, 2956–2962.
- Salinas, E. & Abbott, L. F. (1994). Vector reconstruction from firing rates. *J. Comput. Neurosci.*, 1, 89–107.
- Saul, A. B. & Humphrey, A. L. (1990). Spatial and temporal response properties of lagged and nonlagged cells in the cat lateral geniculate nucleus. *J. Neurophysiol.*, 64, 206–224.
- Saul, A. B. & Humphrey, A. L. (1992a). Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat. *J. Neurophysiol.*, 68, 1190–1207.
- Saul, A. B. & Humphrey, A. L. (1992b). Temporal-frequency tuning of direction selectivity in cat visual cortex. *Vis. Neurosci.*, 8, 365–372.
- Schrater, P. R. & Simoncelli, E. P. (1998). Local velocity representation: Evidence from motion adaptation. *Vision Res.*, 38, 3899–3912.
- Sherman, S. M. (1996). Dual response modes in lateral geniculate neurons: Mechanisms and functions. *Vis. Neurosci.*, 13, 205–213.
- Sherman, S. M. & Guillery, R. W. (1996). Functional organization of thalamocortical relays. *J. Neurophysiol.*, 76, 1367–1395.
- Sherman, S. M. & Koch, C. (1986). The control of retinogeniculate transmission in the mammalian lateral geniculate nucleus. *Exp. Brain Res.*, 63, 1–20.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, 369, 479–482.
- Singer, W. (1977). Control of thalamic transmission by corticofugal and ascending reticular pathways in the visual system. *Physiol. Rev.*, 57, 386–420.
- Singer, W. (1994). A new job for the thalamus. *Nature*, 369, 444–445.
- Smith, A. T. & Edgar, G. K. (1994). Antagonistic comparison of temporal frequency filter outputs as a basis for speed perception. *Vision Res.*, 34, 253–265.

- Smith, S. & Wenderoth, P. (1999). Large repulsion, but not attraction, tilt illusions occur when stimulus parameters selectively favour either transient (M-like) or sustained (P-like) mechanisms. *Vision Res.*, 39, 4113–4121.
- Stevenson, S. B., Cormack, L. K., & Schor, C. M. (1991). Depth attraction and repulsion in random dot stereograms. *Vision Res.*, 31, 805–813.
- Van Essen, D. C. & Anderson, C. H. (1990). Reference frames and dynamic remapping processes in vision. In E. L. Schwartz (Ed.), *Computational Neuroscience* (pp. 278–294). A Bradford Book, MIT Press.
- van Hemmen, J. L. & Sejnowski, T. J., Eds. (2003). *23 Problems in Systems Neuroscience*. New York: Oxford University Press. (to appear).
- Vasconcelos, N. & Carneiro, G. (2002). What is the role of independence for visual recognition? In *Proc. Europ. Conf. Comput. Vision*, volume 2350 of *Lecture Notes in Computer Science* (pp. 297–311). Berlin, Heidelberg: Springer.
- von Krosigk, M., Monckton, J. E., Reiner, P. B., & McCormick, D. A. (1999). Dynamic properties of corticothalamic excitatory postsynaptic potentials and thalamic reticular inhibitory postsynaptic potentials in thalamocortical neurons of the guinea-pig dorsal lateral geniculate nucleus. *Neurosci.*, 91, 7–20.
- Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. *Vision Res.*, 39, 3960–3974.
- Watamaniuk, S. N., Grzywacz, N. M., & Yuille, A. L. (1993). Dependence of speed and direction perception on cinematogram dot density. *Vision Res.*, 33, 849–859.