# Neural networks and perceptual learning

Misha Tsodyks<sup>1</sup> & Charles Gilbert<sup>2</sup>

<sup>1</sup>Department of Neurobiology, Weizmann Institute, Rehovot 76100, Israel (e-mail: misha@weizmann.ac.il)
<sup>2</sup>The Rockefeller University, 1230 York Avenue, New York, New York 10021, USA (e-mail: gilbert@mail.rockefeller.edu)

Sensory perception is a learned trait. The brain strategies we use to perceive the world are constantly modified by experience. With practice, we subconsciously become better at identifying familiar objects or distinguishing fine details in our environment. Current theoretical models simulate some properties of perceptual learning, but neglect the underlying cortical circuits. Future neural network models must incorporate the top-down alteration of cortical function by expectation or perceptual tasks. These newly found dynamic processes are challenging earlier views of static and feedforward processing of sensory information.



In perceptual learning, the improvement develops progressively over many trials, as opposed to other, more explicit types of learning which may require only a single exposure to a stimulus. Perceptual learning is implict, subjects are not consciously aware of it and it progresses even in the absence of a reward for correct responses. Perceptual learning unfolds automatically on repeated exposures to the sensory stimulus, and from integrated efforts at discrimination over a long time. Perceptual learning has important advantages as a brain process amenable to scientific study. First, the behaviour can be quantified relatively accurately under well-defined experimental conditions. Second, there are good reasons to believe that perceptual learning is mediated by neuronal processes that occur at the level of the primary sensory cortex. These areas are the first to receive information from the sensory organs and their circuitry is the best understood of that in the cerebral cortex. Perceptual learning can therefore be quantitatively assessed using three approaches: psychophysical measurement of behaviour, physiological recording of living cortical neurons and computer modelling of well-defined neuronal networks.

Any model of perceptual learning must include at least two components. First, it has to describe the way the sensory world is represented by neuronal activity in the sensory areas of the brain. Second, it has to describe the changes that occur in the sensory pathways when perceptual learning occurs. The current consensus stipulates that every sensory attribute is represented by population activities in the early sensory areas that are dedicated to this attribute. For example, the orientation of visual stimuli is represented by a population of orientation-sensitive neurons in the primary visual areas; the pitch of tonic sounds is represented by a population of frequency-selective neurons in the primary auditory cortex, and so on. The output of such population activity is then interpreted by higher-order cortical areas, which make perceptual decisions. Much less is known about the nature and location of the changes that underlie the improved performance in a sensory task, although evidence is accumulating that the same early stages in sensory processing that initially represent an attribute also mediate the changes involved in improving the discrimination of that attribute. A daunting challenge posed by this picture is to understand how primary circuits can undergo repeated changes that result from learning, but simultaneously be able to operate in tasks that have already been learned.

Here, we review a few representative models of neural networks and assess their performance in terms of perceptual learning. 'Feedforward networks', although based on a very limited number of input units, provide specific read-outs that improve very specifically and quickly during training. The main drawback of feedforward networks, however, is that they rely on a feedback teaching signal, which does not fit with known brain neuroanatomy. By contrast, 'recurrent networks' rely on more realistic horizontal connections, which allows them to learn without the need for any reinforcement signals. Recurrent network models, however, perform relatively poorly on specific perceptual tasks. Models that combine both feedforward and recurrent architectures address some of these problems, but current models are a long way from matching biological circuits.

In the second section of this review, we discuss in more detail the defining characteristics of perceptual learning, as it occurs in real brains. For each property, we consider the challenges it presents for future modellers. In particular, models must accommodate the effect of top-down influences of attention, expectation and perceptual task on the operation of intrinsic cortical circuits. It is becoming increasingly clear that both the encoding and the retrieval of learned information is dependent on feedback interactions between higher- and lower-order cortical areas in sensory cortex. Models should allow for learning in the absence as well as in the presence of reward feedback. They need to account for the high degree of specificity that perceptual learning is known to have. They also need to allow the same circuits to undergo the changes required to encode learned information without this disrupting their existing role in the analysis of the sensory environment. Finally, the rules and characteristics of cellular plasticity have to be integrated at the synaptic, neuronal and network levels (see review in this issue by Abbott and Regehr, page 796) to fully account for the mechanisms underlying perceptual learning.

#### Neural network models of perceptual learning

Models of perceptual learning can be broadly divided into two classes: feedforward versus feedback or recurrent network models. These differ in: (1) network architecture; and (2) the location of the functional changes (output versus input levels, respectively<sup>3</sup>). In feedforward networks (for example, Fig. 2) neurons are located in distinct consecutive layers such that information flows unidirectionally from one layer to another, and learning is implemented by appropriate changes in the relative strengths of feedforward connections. The trigger for changing the connections is usually a discrepancy between the activity at the upper layer (output) and a 'desired' output, which has to be provided to the network during learning ('supervised learning'). In the feedback networks (for example, Fig. 3) information can propagate in loops within a layer or be transferred from higher to lower layers. Such networks allow learning without the



**Figure 1** Example of perceptual learning involving a three-line bisection task. **a**, The subject has to determine whether the central line of three parallel lines is nearer the line on the left or on the right. FP indicates the position of the fixation point. Horizontal line shows one degree of viewing angle. **b**, Practicing this task over thousands of trials for many weeks produces a threefold improvement in the "threshold" — the amount of offset from the central position required to correctly judge the direction of offset. Minute of arc is a 60th of a degree of a viewing angle. The task is practiced in one visual field position, and the improvement is relatively specific to the trained location and orientation, suggesting the involvement of early stages in the visual cortical pathway, where receptive fields are smallest and orientation tuning sharpest. **c**, Importantly, the training is specific for context or the spatial configuration of the stimulus; improvement in the discrimination of the position of a line relative to two parallel lines (three-line bisection) does not transfer to discrimination]. Adapted from ref. 21.

need for a reward or any 'teaching signal' ('unsupervised learning'). A combination of both architectures has been introduced in some models<sup>4</sup>. Indeed, in complete brains most of the sensory areas have the role of read-out for the previous levels and input representation for the subsequent levels of processing.

#### **Feedforward networks**

The best known model of the feedforward type is that conceived by Poggio *et al.*<sup>5</sup> on visual hyperacuity. Poggio *et al.*<sup>5</sup> proposed a threelayers feedforward network (Fig. 2b), the input layer of which consists of a small number of gaussian filters (receptive fields) that transform any input pattern into a vector of activity levels by convolving the input with the corresponding receptive field profiles. The next layer of the network is a set of radial basis functions, each computing the weighted distance between the input vector and a certain template vector that is unique for each function. Finally, the output of the module is computed as a linear combination of the radial basis functions. In models of vernier discrimination, where subjects determine the direction of offset of one line relative to a nearly collinear reference line, the output value of the model determines the perceptual decision, with positive and negative values of output unit being interpreted as the direction to which the target is shifted relative to the reference.

When the model is trained on a set of example inputs with known outputs, the input receptive fields do not change, but the number of radial basis functions and internal parameters of the network are updated. Surprisingly, with only a very limited number of input receptive fields (eight), the model reproduces some salient properties of perceptual learning with high fidelity. The model's hyperacuity level of performance is similar to the experimentally measured one. This increases with the length of the two segments of the vernier stimulus and is specific to its orientation, all in accordance with psychophysical observations.

The main appeal of Poggio *et al*.'s<sup>5</sup> model is that it raises the possibility that when a certain perceptual task is practiced, the brain quickly synthesizes a specialized neural module that reads out the responses in the primary sensory areas of the brain in a way that is optimal for this particular task. Because the responses of sensory neurons are not affected by learning and the synthesized module is not involved in any other tasks, the obtained improvement in the performance is highly specific to the task that was practiced. The model also very successfully replicates observed performance under various stimulus manipulations. However, this class of model has some drawbacks. First, because different elements of the input pattern do not interact with each other directly, the ability of the read-out module to discriminate between different inputs does not strongly depend on the spatial shape of the inputs. For example, the model learns to estimate the horizontal displacement of a single vertical bar relative to an arbitrary reference position, with an absolute precision that is similar to that obtained for a corresponding vernier stimulus (Sahar-Pikielny et al., unpublished data). The fact that spatial features of vernier stimuli seem to be crucial for hyperacuity indicates the involvement of lateral interactions between the receptive fields that respond to different components of the stimulus. Second, most of the learning algorithms in feedforward networks, including the one used by Poggio et al.5, require a teaching signal. Yet, perceptual learning does not require a feedback to proceed, although without it learning proceeds at a slower pace<sup>6</sup>. Third, for a vernier task, human subjects show hyperacuity on the very first trials, which may not leave time for the synthesis of a specialized module.

#### **Recurrent networks**

Several recent observations indicating that perceptual learning results in specific changes in the corresponding primary sensory areas, both on the functional level and on the level of single neuron response properties, provide strong support for models based on recurrent networks. An example of this type of model is that proposed by Adini and colleagues<sup>7,8</sup> which describes perceptual learning in the



**Figure 2** Three-layer feedforward network model of a vernier hyperacuity task. Subjects are required to detect the spatial displacement between the two line segments shown in **a**, superimposed on the receptive fields of the input gaussian filters. The network is shown in **b**. Gaussian filters transform any input pattern into a vector of activity levels by convolving the input with the corresponding receptive field profiles:  $x_i = G_{\alpha}(r - r_i)^* l(r)$ , where l(r) is an amplitude of a visual input at retinal location *r* and  $r_i$  is the centre of the corresponding receptive field. The next layer of the network is a set of radial basis functions (RBFs), each computing the weighted distance between the input vector and a certain template vector that is unique for each function:  $y_a = G(||x - t_a||_{\mathbf{W}})$ , where  $||x - t_a||_{\mathbf{W}} = (x - t_a)^T \mathbf{W}^T \mathbf{W}(x - t_a)$  is the weighted distance between the input vector and the template for the function. **W** denotes the vector of corresponding weights. Finally, the output of the module is computed as a linear combination of the radial basis functions, for example,

positive and negative values of *z* are interpreted as opposite senses of the vernier displacement. Adapted from ref. 5.

case of contrast discrimination. Adini *et al.*<sup>7</sup> assume that perceptual learning is mediated by an increase in contrast sensitivity. This, in turn, results from stimulus-evoked modifications to recurrent connections in the local network in the primary visual cortex.

The model assumes that contrast discrimination is mediated by a local cortical column consisting of two interconnected subpopulations of excitatory and inhibitory neurons<sup>9</sup> (Fig. 4). The activity of the excitatory (*E*) and the inhibitory (*I*) subpopulations is determined by the external feedforward inputs (*e* and *i*, respectively), which increase with the stimulus contrast (*C*), and by the strength of recurrent interactions in the local network (*Is*). Moreover, sensory input from the eye is divided by a fixed proportion between the two populations ( $i \equiv ke$ ; where the constant *k* does not depend on the contrast).

The contrast discrimination threshold is controlled by the steepness of the relationship between the activity (E) and the contrast; that is, by contrast sensitivity. The synaptic learning rule chosen guarantees the convergence of the synaptic strengths to an equilibrium level after repeated presentations of the stimulus. This equilibrium depends on the way the inputs are divided between the populations (that is, on the value of the constant k), but not on the contrast of the stimulus. So, after the stimulus is seen many times, the network adapts to its configuration and terminates the synaptic modifications. However, surrounding the target stimulus with flankers may rekindle the modifications if the additional input to the target, mediated by intracortical connections, is divided differently between the two populations (that is, if it has a different value for k). To explain the psychophysical results Adini et al.<sup>7</sup> assumed that in the presence of flankers the intracortical input is biased in favour of the inhibitory component more than the feedforward input. If this is the case, practicing the contrast discrimination task in the presence of flankers leaves the local network with higher contrast sensitivity than before practice.

An attractive feature of Adini *et al*: $s^7$  model is that it does not require a feedback teaching signal because synaptic modifications are activity-dependent in a hebbian sense. (A hebbian rule of synaptic modification refers to the idea that synapses between neurons that are simultaneously active become stronger.) However, the model cannot easily account for the task-specificity of perceptual learning.

The problem of having perceptual learning affect general processing mechanisms is shared by any model of perceptual learning based on activity-dependent modifications in the lateral connections in the primary sensory areas<sup>10</sup>. A further example is Teich and Qian's<sup>11</sup> model of learning orientation discrimination. The goal of this model was to propose a mechanism for experimentally observed changes in orientation tuning of monkey V1 cells that are specific to the trained orientation<sup>12</sup>. (The model is based on the well-studied recurrent model of orientation selectivity proposed in refs 13-15.) Teich and Qian<sup>11</sup> demonstrate that observed changes in orientation tuning are reproduced in the model if intracortical excitatory connections to cells at and near the trained orientation weaken slightly as a result of learning. In particular, the tuning curves of cells (neurons' responses as a function of a change in the stimulus) whose preferred orientation is near the trained one becomes sharper, in contrast to the broadened tuning curves of cells whose preferred orientation is farther away from the trained one. Similar manipulations, but ones that involve weakening of both excitatory and inhibitory connections around the trained orientation, lead to effects that are observed during the so-called tilt illusion and adaptation experiments<sup>16,17</sup>, including iso-orientation inhibition and changes in orientation tuning bandwidth. These two modifications in tuning lead to opposite effects on the orientation discrimination at the trained or adapted orientation (improvement for learning and deterioration for adaptation). An important issue for future studies on synaptic plasticity and its relationship to perceptual learning is the incorporation of mechanisms that guarantee synaptic modifications that lead to improvement in performance during training. Indeed, we know of only one report of practice-induced deterioration in performance<sup>18</sup>, which indicates that in general, practice leads to an improvement in performance.



**Figure 3** A schematic representation of a cortical column, consisting of two interconnected excitatory (*E*) and inhibitory (*I*) subpopulations, with modifiable intrinsic connections — used here to model contrast discrimination. Both *E* and *I* populations receive external input from the thalamus (*e*, *i* respectively) and from within the cortex ( $\Delta e$ ,  $\Delta i$ ) when surrounding stimuli are present. Thalamic input increases with contrast but the ratio between *e* and *i* remains fixed (*k* = constant). The resulting network activity also increases with contrast. Assuming the threshold-linear gain functions for both subpopulations, their activity is given by:

$$E = e \frac{1 + J_{ii} - kJ_{ei}}{J_{ei}J_{ie} - (J_{ee} - 1)(J_{ii} + 1)}; I = e \frac{J_{ie} - k(J_{ee} - 1)}{J_{ei}J_{ie} - (J_{ee} - 1)(J_{ii} + 1)}$$

(see ref. 8 for more details). *J* refers to a strength of corresponding connections. The form of the contrast sensitivity in the network is therefore determined by the feedforward input *e*, with network interactions providing additional scaling factors. The equilibrium strength of intrinsic connections reached after repeated presentation of the central stimulus alone depends on *k* but not on the stimulus contrast.

#### **Combined models**

Zhaoping et al.<sup>4</sup> proposed a model that combines both recurrent and feedforward learning. This model aims to explain the ability of observers to solve a bisection task with very high precision. Zhaoping and colleagues<sup>4</sup> demonstrate that a linear feedforward mechanism can account for the observed performance provided that the retinal position of the stimulus array is fixed. This condition, however, is too restrictive. First, experimentally, the effects of learning persist when the stimulus is presented up to several degrees away from the trained position. Second, fixation errors, eye tremor and microsaccades are inevitable over the course of the experiment. As shown by Zhaoping et al.4, these uncertainties in the exact position of the stimulus lead to a very poor performance of the purely feedforward read-out mechanism. Zhaoping and colleagues propose that this problem can be rectified if the stimulus undergoes recurrent pre-processing based on the horizontal connections in the primary visual cortex. The pattern of this connection has to be chosen in a way that is highly specific to the particulars of the task, such as the range of stimulus array positions and the distance between the stimulus components. If this is done, the bumps of activity that are evoked by each bar of the stimulus are shifted laterally in such a way as to facilitate the consequent perceptual decision mediated by the feedforward mechanism that reads out the activity in the primary visual cortex.

Although recurrent networks provide a more realistic setting as a substrate for perceptual learning, training them to produce an optimal performance on a task is in general an unsolved problem. An interesting approach has recently been proposed by Seung<sup>19</sup>, which applies the well-known reinforcement learning algorithm<sup>20</sup> to biologically realistic neural networks. The learning algorithm derived by Seung<sup>19</sup> uses the stochastic nature of the synaptic transmission,





which is mediated by probabilistic release of neurotransmitter. According to this algorithm, connections that show a consistent correlation between the neurotransmitter release and good overall performance of the network on the task are 'rewarded' by having their release probabilities increased. Importantly, this idea can be applied to networks with arbitrary architectures, having both feedforward and recurrent elements. However, reaching an optimal performance is crucially dependent on the global evaluation signal (reward) that is available to the synaptic connections in the network.

#### **Real brains' challenges to models** Perceptual learning is highly specific

If a subject is trained on a discrimination task at one location in space, the improvement in performance is relatively specific for that location and does not transfer to other locations in the sensory map. For example, training a subject on a three-line bisection task leads to improvement at the trained location, but the degree of transfer drops off as the visual stimulus is shifted to locations up to 8° away, and there is no transfer when the stimulus is shifted to the opposite hemifield. The training is also specific to the orientation of the trained stimulus. This suggests the involvement of early stages in cortical processing (such as primary visual cortex, V1), where the receptive fields are smallest, the visuotopic maps most highly organized, and the orientation selectivity sharpest<sup>21</sup>. However it is interesting to note that the degree of transfer observed is larger than the receptive fields in V1. This amount of spread of learned information should inform the implementation of computational models of learning.

Perceptual learning is also specific for context and the configuration of the stimulus learned in training. For example, training on a three-line bisection task (Fig. 1) does not transfer to a

vernier discrimination task. In both tasks the target of the discrimination has the same visual field position and orientation, and the trained attribute (position) is also the same. But in one task the context is two sideby-side flanking parallel lines and in the other it is two lines that are collinear with the target. Contextual specificity has been seen in other forms of perceptual learning, such as depth discrimination<sup>22</sup>. It is worth noting, however, that nonspecific effects of perceptual learning on the basic representations within an adult visual system have recently been reported for amblyopic patients<sup>23</sup>. But the more general rule is that learning on one task only shows transfer to another task to the degree that both tasks have elements in common. Further work is needed to determine - when training in discriminating multiple stimuli - which components of these stimuli are employed for making the discrimination. Models will assume greater importance in guiding these studies by showing which features are most useful for recognition systems to generalize to novel stimuli<sup>24</sup>.

The observed task specificity of perceptual learning poses a serious challenge to models based on changes in the wiring of neural circuits in the primary sensory areas. This is because task specificity should lead to some general effects on sensory processing in the particular domain that is affected by training. An exciting possibility that could explain the relative absence of cross-talk could be a task-dependence of the lateral interactions in the sensory areas. Indeed, after a monkey was trained on a three-line bisection task, the modulation of the cell's response to a line segment within the receptive field by a second parallel line, placed outside the receptive field, differed depending on whether the monkey was tested on the trained task or on an unrelated fixation or vernier discrimination task<sup>25,26</sup>.

This degree of specificity also has important implications for the way in which acquired information is represented in the cortex. A suggested mechanism is referred to as 'cortical recruitment'. This involves an increase in the area of cortex representing the trained location. Experiments demonstrating this phenomenon were done in the somatosensory and auditory systems<sup>27,28</sup>. However, even here other cortical changes seem to correlate better with the improvement in performance. These include a change in the temporal characteristics of neuronal responses, with a more reliable entrainment of the response to the periodicity of the stimulus<sup>29</sup>. In the visual system no such cortical magnification has been observed<sup>25</sup>. It is still unclear whether there are differences between results from the visual compared to other sensory systems, although there are some differences in the experimental designs used. For example, in the visual studies emergent properties of cortex are associated with training, whereas in the somatosensory and auditory systems properties of the cortical inputs are involved. Modelling the cortical changes underlying perceptual learning must allow for the specificity of learning for the trained stimulus. Moreover, these models must be consistent with the finding that training on one stimulus at one location does not produce a degradation in performance when discriminating other stimuli at other locations.

#### Time course of perceptual learning

An important component of models of perceptual learning is the rate at which learning occurs. As shown below, in some experiments there is an initial period of fast learning, which is then followed by a much slower rate of improvement (see Fig. 4). Several neural network models are able to reproduce this behaviour, albeit by using different mechanisms. In the feedforward network of Poggio *et al.*<sup>5</sup> (Fig. 3), during the first phase of learning new units are added to the intermediate layer of the network, ensuring the coverage of all the space of possible input patterns. As a result, the classification error rate comes within 10% of its asymptotic value after just several examples. This is followed by a later, slower phase of learning during which the architecture of the network remains fixed but the parameters of the network slowly adapt to their optimal values. This leads to incremental improvement in performance. In neural terms, the first phase could correspond to the recruitment of neurons in intermediate levels of visual processing which would represent the stimuli encountered by observers at the beginning of practice. In Zhiaoping et al.'s<sup>4</sup> model (which combines both recurrent and feedforward mechanisms) two phases of learning could correspond to differing speeds of modification in the corresponding connections.

#### Perceptual learning requires repetition but not feedback

The improvement in performance seen in perceptual learning is proportional to the number of trials taken, although performance eventually asymptotes to a point beyond which additional trials make no further difference. During a discrimination task improvement is seen even in the absence of a reward or any indication that the correct response was made. Nevertheless, brain reward systems have been shown to have a role in perceptual learning. One of the sources of reward in the brain is thought to be the cholinergic input from the nucleus basalis. Pharmacological blockade of the cholinergic input can inhibit, and stimulation of the nucleus basalis can promote, perceptual learning<sup>30,31</sup>. So, it is possible that mere performance of the task has an implicit reward associated with it, even when a reward is not given in every trial. Although learning can occur in the absence of feedback, feedback can facilitate learning. Moreover, feedback that is uncorrelated with the response disrupts learning. But the nature of effective feedback is interesting, because block feedback (that is, feedback after several trials, so in response to a certain percentage correct after a number of presentations) is as effective as trial-bytrial feedback6.

These observations put obvious constraints on the feedforward networks with supervised learning, in which feedback is usually implemented as a 'teaching' signal that is required for the correct change in the strength of synaptic connections<sup>32</sup>. An interesting modification of supervised models of perceptual learning is that proposed by Herzog and Fahle<sup>33</sup>. The main innovation of this model is that an internal evaluating feedback signal is used to guide selective connections between the input units of the model and the next network layer. Internal feedback is estimated as a difference between the responses of the output units to inputs that have to be discriminated. Learning then selectively inhibits the feedforward connections that are not providing the signal required to increase the evaluated performance (a process called gating; see refs 34, 35 for similar ideas). Unsupervised learning algorithms in feedforward networks have also been proposed<sup>36,37</sup>.

#### Longevity of perceptual learning

A striking long-term stability of the improvement in performance is observed in certain tasks. For example, in Karni and Sagi's experiments on texture discrimination<sup>38</sup> subjects achieved a significant improvement in performance over four to five days. However, once subjects learned the task, they maintained their improved level of performance for at least three years without needing further practice. This observation particularly challenges any model that is based on activity-dependent synaptic plasticity in the sensory areas. Obviously, neurons in these areas are constantly responding to a continuous stream of sensory inputs that should, with time, wipe out specific traces produced by training. A possible explanation for long-term improvement could be that a certain fraction of synaptic connections becomes resilient to modification as a result of perceptual learning (see review in this issue by Abbott and Regehr, page 796). Alternatively, training on one task could affect a small subset of inputs that are only engaged when that task is performed. Even if the same cells participate in different tasks they may engage different inputs. This would minimize negative interference in the traces produced by training on the different tasks. Understanding the causes for the striking longevity of perceptual learning and its dependence on the parameters of practice protocols may be an important step towards elucidating the process of consolidation of long-term memories in general.

#### Perceptual learning involves top-down influences

In most instances of perceptual learning the subject must attend to the trained stimulus for improvement to occur<sup>39,2</sup>, although some studies have suggested that learning can occur in the absence of attention<sup>40,41</sup>. This is one form of evidence for the role of top-down influences in learning; that is, for the regulation of informationencoding in lower-order cortical areas by higher-order areas. The top-down signal may be carried by cortical feedback connections. A generally accepted view of pathways of connectivity between cortical areas is that of a hierarchy which starts from primary sensory cortex and proceeds up to the highest areas encoding the most complex information. For every forward connection in this pathway, however, there is a reverse or feedback connection. The role of cortical feedback is little understood, but increasing evidence for attentional influences at early stages suggests that feedback may be involved in transmitting this kind of cognitive control. The attentional signal may have a role both in the ongoing processing of sensory information and in the encoding of learned information. A recent study has shown that the top-down influence can be extremely specific to different discrimination tasks at the same visual location. In these experiments, neurons in V1 changed their functional properties according to the task being performed, and these properties were only present when the animal was performing the trained task<sup>26</sup>. So, there is a long-term change in function associated with the period of training (which can take place over many weeks), and a short-term switching between different functional states as the subject shifts from one trained task to another. The same top-down influences or feedback circuits involved in regulating the encoding of learned information may also be involved in its recall.

An appealing hypothesis, from a theoretical point of view, assigns to the feedback influences the role of transmitting to primary cortical areas signals that reflect the expectations of the sensory inputs. These signals are based on the internal representation of the sources of these inputs<sup>42</sup>. The neurons in the lower areas then respond to the deviations of the actual sensory inputs from the predicted ones. For this predictive coding system to work well it has to learn the statistical regularities in the sensory environment. This kind of model has not yet been directly applied to perceptual learning, and no direct experimental evidence for the effect of internally generated expectations on the neural responses in the primary sensory areas is currently available. Although there is no evidence that early sensory areas respond to deviations - instead they carry information more fully related to the stimulus - their tuning is clearly modulated by top-down influences of attention, expectation and perceptual task. Models that incorporate top-down interactions for both encoding and recall will assume increasing importance as experimental results provide further evidence for these interactions.

#### Cortical representation associated with perceptual learning

There is considerable debate concerning which cortical areas represent the higher order properties associated with contextual influences, and which circuits carry these influences (including intrinsic circuits within individual cortical areas and feedback connections to those areas). Even so, it is becoming increasingly clear that many areas, including primary sensory cortex, show functional changes that reflect learned information. The notion that different cortical areas specialize in particular kinds of information will probably change, for several reasons. As supported by both psychophysical and fMRI studies, the strategies that the brain uses for object recognition change depending on the degree of familiarity the subject has with the object. Learning to identify an object is associated with global changes in the representation of information across the cortical sensory pathway (Sigman et al., submitted). Moreover, the act of object recognition does not involve a single cortical area but an interaction between multiple cortical areas and between forward, intrinsic and feedback circuits<sup>43,26</sup> (also Sigman et al., submitted). This of course creates a formidable challenge in terms of creating models that can replicate the multiple levels at which information can be represented in the cortex.

#### **Rules of plasticity**

The most generally accepted view of plasticity at the synaptic level is that, with coincidence in the activation of the presynaptic terminal and the postsynaptic cells, the synapses involved become strengthened<sup>44</sup>. This hebbian rule is dealt with in more detail elsewhere in this issue (see review in this issue by Abbott and Regehr, page 796). But this rule has profound implications at the systems level, although some experimental results suggest that this rule does not operate exclusively. A fundamental question is whether sensory systems are designed to pick out exceptions in the environment, or to identify common coincidences. Coincidences or correlations in the environment can be represented at the level of correlations in neuronal firing, which then leads to synaptic strengthening. Information theoretic considerations, on the other hand, suggest that sensory systems are designed to pick up exceptions or changes in the environment (for example, the presence of a surface edge as opposed to the continuation of a uniform surface). Similar to JPEG compression, this would suggest that to carry the maximum amount of information along a limited number of channels (the optic nerve, for example), the functional properties of neurons have to be as distinct from one another as possible. This requires 'decorrelation' of their function, which suggests the need for an 'anti-hebbian' rule of plasticity45.

A particular form of synaptic plasticity combining hebbian and anti-hebbian rules in a way that is motivated by recent studies on spike-time-dependent synaptic plasticity<sup>46,47</sup> was proposed by Adini and colleagues<sup>37</sup>. Here, the learning rule chosen guarantees the convergence of the synaptic strengths to an equilibrium level after repeated presentations of the stimulus. However, synaptic modification restarts when the stimulus changes. This property could account for the saturation of perceptual learning after repeated practice. The perceptual task that was studied — contrast discrimination — seems to be saturated when it is performed on a wide range of contrasts, but not when a single contrast, or a few contrast levels in a fixed order are presented (see refs 48, 49 for a recent debate on this issue). When the stimulus configuration is changed during practice by adding surrounding components, Adini *et al.*<sup>7,49</sup> observed an improvement in performance that was to a large degree independent of uncertainty in the stimulus contrast, in accordance with the above-mentioned feature of the learning rule.

#### Neuronal changes associated with perceptual learning

Various experimental observations, and computational models, have revealed changes in functional properties at the neuronal level that are associated with perceptual learning. These include changes in the tuning of neurons to the trained attribute. Steepening of the slope of the tuning curve reduces the threshold at which neurons show significant differences in response, and therefore the threshold required for discrimination<sup>11,12</sup> (although others have failed to find such a change<sup>50</sup>). This in effect leads to a reduction in the number of neurons responding to the stimulus, contrary to the observed increase in response in the cortical area representing the trained stimulus. Changes have been observed in the timing and reliability of neuronal responses, which represents an increase in signal to noise. This also leads to a reduction in the threshold at which there are significant changes in response. Along with improvement in discrimination of more complex forms, neurons show changes in contextual tuning. This is the way in which the elements of a complex stimulus interact, in terms of neurons' responses, as the stimulus configuration is changed.

#### Outlook

Many facets of perceptual learning have been successfully reproduced in simple, but plausible, neural network models. These models provide important insights into possible mechanisms which can then be tested experimentally. But, so far, these models are far too specific to provide a full account of the phenomenology of perceptual learning. As a result, they do not support a more general understanding of the neuronal processes underlying early stages of information processing. What is needed is a combination of feedforward models with models based on lateral feedback and top-down influences representing the task, expectations, attention and signals controlling synaptic modification. Future work will have to increasingly include details of the biophysical mechanisms of synaptic learning in cortical networks.

#### doi:10.1038/nature03013

- 1. Wang, Q., Cavanagh, P. & Green, M. Familiarity and pop-out in visual search. *Percept. Psychophys.* 56, 495–500 (1994).
- 2. Sigman, M. & Gilber, C. D. Learning to find a shape. Nature Neurosci. 3, 264-269 (2000).
- Hertz, J., Krogh, A. & Palmer, R. G. Introduction to the Theory of Neural Computation (Perseus Publishing, Cambridge, Massachussetts, 1991).
- Zhaoping, L., Herzog, M. & Dayan, P. Nonlinear observation and recurrent preprocessing in perceptual learning. *Network* 14, 233–247 (2003).
- Poggio, T., Fahle, M. & Edelman, S. Fast perceptual learning in visual hyperacuity. Science 256, 1018–1021 (1992).
- Herzog, M. H. & Fahle, M. The role of feedback in learning a vernier discrimination task. Vision Res. 37, 2133–2141 (1997).
- Adini, Y., Sagi, D. & Tsodyks, M. Context enabled learning in human visual system. Nature 415, 790–794 (2003).
- Tsodyks, M., Adini, Y. & Sagi, D. Associative learning in early vision. *Neural Netw.* 17, 823–832 (2004).
- Wilson, H. R. & Cowan, J. D. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J.* 12, 1–24 (1972).
- Hoshino, O. Neuronal bases of perceptual learning revealed by a synaptic balance scheme. *Neural Comput.* 16, 563–594 (2004).
- Teich, A. & Qian, N. Learning and adaptation in a recurrent model of V1 orientation selectivity. J. Neurophysiol. 89, 2086–2100 (2003).

- Schoups, A., Vogels, R., Qian, N. & Orban, G. Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412, 549–553 (2001).
- Ben-Yishai, R., Bar-Or, R. L. & Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proc. Natl Acad. Sci. USA* 92, 3844–3848 (1995).
- Douglas, R., Koch, C., Mahowald, M., Martin, K. & Suarez, H. Recurrent excitation in neocortical circuits. *Science* 269, 981–985 (1995).
- Somers, D., Nelson, S. & Sur, M. An emergent model of orientation selectivity in cat visual cortical simple cells. J. Neurosci. 15, 5448–5465 (1995).
- Gilbert, C. D. & Wiesel, T. N. The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res.* 30, 1689–1701 (1990).
- Dragoi, V., Sharma, J. & Sur, M. Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron* 28, 287–298 (2000).
- Polat, U. & Sagi, D. Spatial interactions in human vision: from near to far via experience-dependent cascade of connections. *Proc. Natl Acad. Sci. USA* 91, 1206–1209 (1994).
- 19. Seung, H. S. Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* **40**, 1063–1073 (2003).
- Williams, R. Simple statistical gradient-following algorithms for connectionst reinforcement learning. Mach. Learn. 8, 229–256 (1992).
- Crist, R. E., Kapadia, M., Westheimer, G. & Gilbert, C. D. Perceptual learning of spatial localization: specificity for orientation, position and context. J. Neurophysiol. 78, 2889–2894 (1997).
- McKee, S. P. & Westheimer, G. Improvement in vernier acuity with practice. *Percept. Psychophys.* 24, 258–262 (1978).
- 23. Polat, U., Ma-Naim, T., Belkin, M. & Sagi, D. Improving vision in adult amblyopia by perceptual learning. *Proc. Natl Acad. Sci. USA* 101, 6692–6697 (2004).
- Ullman, S. & Bart, E. Recognition invariance obtained by extended and invariant features. *Neural* Netw. 17, 833–848 (2004).
- 25. Crist, R., Li, W. & Gilbert, C. Learning to see: experience and attention in primary visual cortex. *Nature Neurosci.* **4**, 519–525 (2001).
- Li, W. Piech, V. & Gilber, C. D. Perceptual learning and top-down influences in primary visual cortex. Nature Neurosci. 7, 651–657 (2004).
- Recanzone, G. H., Merzenich, M. M. & Jenkins, W. M. Frequency discrimination training engaging a restricted skin surface results in an emergence of a cutaneous response zone in cortical area 3a. *J. Neurophysiol.* 67, 1057–1070 (1992).
- Recanzone, G. H., Schreiner, C. E. & Merzenich, M. M. Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *J. Neurosci.* 13, 87–103 (1993).
- Recanzone, G. H., Merzenich, M. M. & Schreiner, C. E. Changes in the distributed temporal response properties of SI cortical neurons reflect improvements in performance on a temporally based tactile discrimination task. J. Neurophysiol. 67, 1071–1091 (1992).
- Bakin, J. S. & Winberger, N. M. Induction of a physiological memory in the cerebral cortex by stimulation of the nuclear basalis. *Proc. Natl Acad. Sci. USA* 93, 11219–11224 (1996).
- Kilgard, M. P. & Merzenich, M. M. Cortical map reorganization enabled by nucleus basalis activity. Science 279, 1714–1718 (1998).

- Seung, H. S., Sompolinsky, H. & Tishby, N. Statistical mechanics of learning from examples. *Phys. Rev. A* 45, 6056–6091 (1992).
- Herzog, M. H. & Fahle, M. Modeling perceptual learning difficulties and how they can be overcome. Biol. Cybern. 78, 107–117 (1998).
- Mato, G. & Sompolinsky, H. Neural network models of perceptual learning of angle discrimination. Neural Comput. 8, 270–299 (1996).
- Dosher, B. A. & Lu, Z. L. Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proc. Natl Acad. Sci. USA* 95, 13988–13993 (1998).
- Moses, Y., Schechtman, G. & Ullman, S. Self-calibrated collinearity detector. *Biol. Cybern.* 63, 463–475 (1990).
- Weiss, Y., Edelman, S. & Fahle, M. Models of perceputal learning in vernier hyperacuity. *Neural Comput.* 5, 695–718 (1993).
- 38. Karni, A. & Sagi, D. The time course of learning a visual skill. Nature 365, 250–252 (1993).
- Ahissar, M. & Hochstein, S. Attentional control of early perceptual learning. Proc. Natl Acad. Sci. USA 90, 5718–5722 (1993).
- Watanabe, T., Nanez, J. E. & Sasaki, Y. Perceptual learning without perception. Nature 413, 844–848 (2001).
- Seitz, A. R. & Watanabe, T. Psychophysics: is subliminal learning really passive? Nature 422, 36 (2003).
- Rao, R. P. & Ballard, D. H. Dynamic model of visual recognition predicts neural response properties in the visual cortex. *Neural Comput.* 9, 721–763 (1997).
- Ullman, S. Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb. Cortex* 5, 1–11 (1995).
- 44. Hebb, D. O. Organization of Behavior (John Wiley & Sons Inc, 1949).
- 45. Barlow, H. B. & Foldiak, P. in *The Computing Neuron* (eds Durbin, R. Miall, C. & Mitchison, G.) 54–72 (Addison-Wesley, Workingham, England, 1989).
- Markram, H., Lubke, J., Frotscher, M. & Sakmann, B. Regulation of synaptic efficacy by coincidence of postsynaptic Aps and EPSPs. *Science* 275, 213–215 (1997).
- Senn, W., Markram, H. & Tsodyks, M. An algorithm for modifying neurotransmitter release probability based on pre- and post-synaptic spike timing. *Neural Comput.* 13, 35–67 (2001).
- Yu, C., Levi, D. M. & Klein, S. A. Perceptual learning in contrast discrimination and the (minimal) role of context. J. Vision 4, 169–182 (2004).
- Adini, Y., Wilkonsky, A., Haspel, R., Tsodyks, M. & Sagi, D. Perceptual learning in contrast discrimination: the effect of contrast uncertainty. J. Vision, in the press.
- Ghose, G. M., Yang, T. & Maunsell, J. H. R. Physiological correlates of perceptual learning in monkey V1 and V2. J. Neurophysiol. 87, 1867–1888 (2002).

**Acknowledgements** We thank W. Li, V. Piech, D. Sagi and K. Pawelzik for their suggestions on the manuscript. M.T. is supported by Israeli Science Foundation and Irving B. Harris foundation. C.G. is supported by NIH.

**Competing interests statement** The authors declare that they have no competing financial interests.

## insight review articles