

Published in final edited form as:

Trends Cogn Sci. 2010 February ; 14(2): 64–71. doi:10.1016/j.tics.2009.11.005.

Perceptual learning rules based on reinforcers and attention

Pieter R. Roelfsema^{1,2}, Arjen van Ooyen², and Takeo Watanabe³

¹Department of Vision & Cognition, Netherlands Institute for Neurosciences, an institute of the Royal Netherlands Academy of Arts and Sciences (KNAW), Meibergdreef 47, 1105 BA, Amsterdam, The Netherlands. ²Department of Integrative Neurophysiology, Centre for Neurogenomics and Cognitive Research, VU University, De Boelelaan 1085, 1081 HV Amsterdam, Amsterdam, The Netherlands. ³Department of Psychology, Boston University, 64 Cummington St., Boston, MA 02215, USA

Abstract

How does the brain learn those visual features that are relevant for behavior? In this article, we focus on two factors that guide plasticity of visual representations. First, reinforcers cause the global release of diffusive neuromodulatory signals that gate plasticity. Second, attentional feedback signals highlight the chain of neurons between sensory and motor cortex responsible for the selected action. We here propose that the attentional feedback signals guide learning by suppressing plasticity of irrelevant features while permitting the learning of relevant ones. By hypothesizing that sensory signals that are too weak to be perceived can escape from this inhibitory feedback, we bring attentional learning theories and theories that emphasized the importance of neuromodulatory signals into a single, unified framework.

Perception improves with training

Visual perception improves with practice. A birdwatcher sees differences between birds that are invisible to the untrained eye. To gain an understanding of perceptual learning one can compare the perception of bird experts to the perception of subjects with other interests [as explained in detail in 1]. Alternatively, one can study perceptual learning in the laboratory, which has produced many important insights. Training improves perception, even in adult observers, provided they are willing to invest some effort in the task. Subjects typically have to train for a few hundred trials per day over a few days before perceptual improvements are noticeable. Under these conditions, subjects can even become better in discriminating between basic features, for example, between subtle variations in the orientation or motion direction of a visual stimulus [2]. Once perceptual learning has occurred, it is persistent and can last for many months [3] or years [4]. The learning effects are often specific so that perceptual improvements in a particular version of a task do not generalize to other versions. Performance improvements are not observed, for example, if the test stimulus has a different orientation [4–6], motion direction [7–8] or contrast [9–10] than the trained stimulus. Moreover, the training effects are often retinotopically specific. After training in one region of the visual field, the improvement in performance does not transfer to other visual field locations [4–6–9–11–12], although special learning procedures can cause better generalization [13–14].

© 2009 Elsevier Ltd. All rights reserved.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

What are the mechanisms that determine perceptual learning? The improvements in perception could be the result of changes in sensory representations, but they could also be the result of the way that sensory representations are read out by decision-making areas [15-17]. Furthermore, the role of selective attention in learning is unclear. It appears to be important for some forms of learning [10-18-22], but not for others [8]. Similarly, in some cases learning takes place without giving explicit feedback to the subject about the accuracy of the responses [23], while in other cases such feedback facilitates learning [24]. The goal of the present review is to put together recent neurophysiological and psychological findings into a coherent theoretical framework for perceptual learning. In the context of this framework, we will discuss attention-gated reinforcement learning (AGREL), a model that posits that selective attention and neuromodulatory systems jointly determine the plasticity of sensory representations [25]. We will proceed by proposing a generalization of this model that provides a new, neurophysiologically plausible demarcation between the conditions where attention is required for learning and the conditions where it is not.

Neuronal correlates of perceptual learning

How does the visual cortex change as a result of perceptual learning? The available evidence implicates early sensory representations but also higher association areas in perceptual learning, although the relative contribution of low-level and high-level mechanisms is under debate [15-16]. With regard to early visual representations, functional imaging studies have revealed increases in neuronal activity at the representation of the location in the visual field that is trained [26], often with particularly strong effects in the primary visual cortex [27-31]. Moreover, neurophysiological studies have demonstrated that neurons in early visual areas change their response properties during perceptual learning [5-32-34]. For example, when monkeys are trained to judge the orientation of a stimulus, neurons in the primary visual cortex (area V1) become sensitive to small variations around the trained orientation. In addition to area V1, increased sensitivity of neurons to category boundaries is also observed in higher visual areas, including area V4 [32-35], the inferotemporal cortex [36], area LIP [37] and the prefrontal cortex [38]. From a functional point of view, the amplified representation of feature values close to a category boundary is useful, because small changes in input should lead to categorically different behavioral responses [39-40]. If the stimuli that need to be classified differ in multiple feature dimensions, the increases in sensitivity caused by training are strongest for the features that distinguish between categories and weaker for features that do not [41-42]. Thus, neurons in many areas of the visual cortex change their tuning in accordance with the arbitrary categories imposed by a task.

A few studies have directly compared plasticity in lower and higher visual areas and found stronger effects in the higher areas [15-36-43], although there are also tasks where the changes in lower areas predominate [27-28]. Some tasks require a categorization on the basis of features that are easy to discriminate, and subjects can learn these tasks as new stimulus-response mappings that may not depend on plasticity in the visual cortex. Genuine perceptual learning paradigms, on the other hand, train subjects to perceive small variations in features invisible to the untrained eye, and may therefore engage plasticity in the visual cortex. Ahissar and Hochstein [19-44] demonstrated that learning in easy (or low precision [45]) tasks generalizes across locations and feature values, suggestive of plasticity at high representational levels, while training in higher precision tasks is more specific to the trained stimulus implicating lower representational levels.

Theories of perceptual learning therefore have to explain where, when and why plastic changes occur. A particularly challenging question for these theories is how learning effects occur in early sensory areas, remote from the areas where perceptual decisions are made and task performance is monitored. What signal informs the sensory neurons to become tuned to the

feature variations that matter? In what follows we consider two important routes for these effects to reach sensory areas: diffuse neuromodulatory systems and feedback connections propagating attentional signals from higher to lower areas.

Global neuromodulatory systems that gate plasticity

There are a number of neuromodulatory systems that project broadly to most areas of the cerebral cortex and deliver information about the relevance of stimuli and the association between stimuli and rewards. The two neuromodulatory systems that are most often implicated in neuronal plasticity are dopamine and acetylcholine. The substantia nigra and the ventral tegmental area are dopaminergic structures in the midbrain that project to the basal ganglia and the cerebral cortex. In an elegant series of studies (reviewed in [46]), Schultz and co-workers demonstrated that dopamine neurons code deviations from reward expectancy. They respond if a reward is given when none was expected and also to stimuli that predict rewards, causing a surge of dopamine in the basal ganglia and cerebral cortex [47]. Because the increase in the dopamine concentration signals that the outcome of a behavioral choice is better than expected, it is beneficial to potentiate active synapses and thereby increase the probability that the same choice will be made again in the future. In slice preparations of the basal ganglia, dopamine has indeed been shown to control synaptic plasticity [48]. Moreover, there is *in vivo* evidence for the control of plasticity by dopamine. If transient dopamine signals are paired with an auditory tone, the representation of this tone is expanded in the auditory cortex [49].

Acetylcholine is another neuromodulator that has been linked to synaptic plasticity. Neurons in the basal forebrain project to the cortex to supply acetylcholine. These neurons also respond to rewards [50], although the relation between their activity and reward prediction is not as well understood as for the dopamine neurons. If artificial stimulation of the basal forebrain is paired with an auditory tone, then the representation of this tone in the auditory cortex increases [51-52]. Thus acetylcholine promotes neuronal plasticity *in vivo* and it also influences synaptic plasticity in cortical slice preparations [53]. Other studies have shown that acetylcholine is necessary for plasticity, because a reduction of the cholinergic input reduces cortical plasticity [54] and impairs learning [55-57]. These results, taken together, provide strong support for the idea that learning and plasticity of cortical representations are controlled by neuromodulatory systems that change their activity in relation to rewarding stimuli or stimuli that predict reward.

Role of selective attention in learning

Visual attention provides a second route for signals about behavioral relevance to reach the visual cortex. There is substantial evidence for a role of selective attention (here we will not consider the effects of ‘general attention’ or arousal [58] and do not use the word ‘attention’ for the effects of neuromodulators) in determining what is learned and what not. One powerful approach for studying the role of attention in learning is given by the ‘redundant relevant cues’ method [18-59]. The subjects have to learn to associate stimuli with responses and can use multiple features of a stimulus, e.g. color and shape, to determine the correct response. The critical manipulation is that the subjects are cued to direct their attention selectively to one of the features and not to the other. In these situations, they usually learn to use the attended feature and even exhibit an increase in perceptual sensitivity for this feature [18] while they do not learn to use the other, redundant feature even though it is presented and rewarded equally often. As a result, the subjects cannot perform the task after the training phase if the attended feature is taken away so that they are forced to use the redundant feature. Thus, in these cases attention to a feature determines which representations undergo plasticity and which do not. Similar effects occur for spatial attention, because perceptual learning is particularly pronounced for stimuli at attended locations [e.g. 22].

How do these feature-based and spatial attentional effects reach the early levels of the visual system? The most likely route is through feedback connections that run from the higher areas back to lower areas of the visual cortex [60-61]. Cortical areas involved in response selection feed back to sensory areas so that objects relevant for behavior are represented more strongly than irrelevant ones. Such a direct relationship between behavioral relevance and visual selection was made very explicit in the 'premotor theory' of attention [62-63]. When a stimulus is selected for a behavioral response, the relevant features automatically receive attention. This theory is supported by experiments on eye movements, as attention is invariably directed to items that are selected as target for an eye movement [64-66].

Neurophysiological findings provide support for the coupling between movement selection and spatial as well as feature-based attention. During visual search, for example, the representation of the features of the item that is searched as well its spatial location are enhanced in the visual cortex [67] and also in the frontal eye fields (area FEF) [68], and the frontal eye fields may feed back to cause attentional selection in the visual cortex [69]. The same is true in other tasks where visual stimuli compete for selection. Figure 1 shows data from a study where monkeys were trained to select the circle at the end of a curve that was connected to a fixation point as the target for an eye movement (Figure 1b), while ignoring a distractor curve [70]. Neurons in area FEF responded to the appearance of the stimulus in their receptive field, although their initial response did not discriminate between the target curve and the distractor (Figure 1c). After a short delay, however, responses evoked by the target curve became much stronger than those evoked by the distractor (striped bar), and this enhanced activity is a neuronal correlate of target selection [71]. A similar selection signal is observed in area V1 (Figure 1c), where responses evoked by the relevant, attended curve are enhanced over the responses evoked by a distractor curve, in a later phase of the response.

These results support the idea that the appearance of a visual stimulus, be it a target or distractor, initially triggers the rapid propagation of activity from lower to higher areas of the visual cortex through feedforward connections (Figure 1a) [72-73]. This phase is followed by an epoch where neurons in the frontal cortex that code different actions engage in a stochastic competition. The cells that code the action that wins the competition have stronger responses than the neurons that lose, and feed back to the representation of the selected object in the visual cortex [74-75], causing a response enhancement that is a correlate of selective attention [76]. Such a counterstreams model [77] requires reciprocal connections between the visual and frontal cortex, so that actions that are selected in the frontal cortex provide feedback to neurons that gave input for this particular choice (orange neurons in Figure 1a), thus highlighting the circuits in the visual cortex that determine the course of action. Reinforcement learning theories (like AGREL, see below) hold that actions are selected stochastically, so that the same visual stimulus can give rise to different actions and therefore also different patterns of feedback [25] (see Box 1).

Such a coupling between motor selection in the frontal cortex and attentional effects in the visual cortex during action selection is useful for guiding plasticity because plasticity occurs for connections between neurons that are important for the selected response. This computational idea is supported by psychological studies showing that attention gates learning [18-59-78-80]. Moreover, a recent pharmacological study that investigated the role of different receptors in feedforward and feedback processing demonstrated that feedback connections have a larger proportion of NMDA-receptors than feedforward connections [81]. This result suggests that the feedback connections might gate perceptual learning [82] by activating NMDA receptors [83]. Another hypothetical route for feedback connections to gate plasticity involves acetylcholine receptors that are involved in selective attention [84] and also play a role in the gating of synaptic plasticity, as was discussed above.

Interactions between attention and diffuse reinforcement learning signals

So far we have reviewed evidence for the gating of plasticity by neuromodulatory systems as well as selective attention. Roelfsema and van Ooyen proposed a framework called ‘attention-gated reinforcement learning’ (AGREL) [25] that holds that these two signals are complementary and jointly determine plasticity (Box 1). The network receives a reward for a correct choice, while it receives nothing if it makes an error. After the action, neuromodulators are released into the network to indicate whether the rewarded outcome is better or worse than expected (Figure 1) [46]. If the network receives more reward than expected, the neuromodulators cause an increase in strength of the connections between active cells, so that this action becomes more probable in the future; the opposite happens for actions with a disappointing outcome. The second signal is the attentional feedback during action selection that ensures the specificity of synaptic changes. Although the neuromodulators are released globally, the synaptic changes occur only for units that received the attentional feedback signal from the response selection stage during action selection. AGREL causes feedforward and feedback connections to become reciprocal, in accordance with the anatomy of the cortico-cortical connections. Consequently, the neurons that give most input to the winning action also receive most feedback. As a result, only sensory neurons involved in the perceptual decision change their tuning, while the tuning of other neurons remains the same. The attentional feedback signal thereby acts as a credit assignment signal, highlighting those neurons and synapses that are responsible for the outcome of a trial, thus increasing the efficiency of the learning process substantially [85]. A remarkable result is that under some conditions, the global neuromodulatory signal combined with the attentional feedback signal gives rise to learning rules that are as powerful as supervised learning schemes, like error-backpropagation, although the learning scheme operates by trial and error and is plausible from a neurophysiological point of view. Learning rules that combine the two factors, like AGREL, can reproduce the effects of categorization learning if there is a direct mapping of stimuli onto responses. They steepen the tuning curve of sensory neurons at the boundary between categories and cause a selective representation of ‘diagnostic’ features that matter for the task. It is still an open question whether these reinforcement learning models can be adapted to explain perceptual learning in tasks that require a comparison between stimuli presented at different times, for example in delayed match-to-sample tasks where subjects have to judge whether two sequentially presented stimuli are the same. These tasks require the comparison between a memory trace of the first stimulus and the perceptual representation of the second stimulus, while the existing reinforcement learning models do not have such a working memory.

Perceptual learning without attention

Although the studies reviewed above demonstrate that attention gates learning, there are also forms of perceptual learning that occur without attention. Watanabe and his colleagues [8] demonstrated that perceptual learning can occur for stimuli too weak to be perceived, if they are paired with the detection of another stimulus. In one of their experiments [86], subjects monitored an RSVP (rapid serial visual presentation) stream for target digits that were presented on a background of moving dots (Figure 2a). Unbeknownst to the subjects, the target digits were consistently paired with a very weak motion stimulus in one direction and, remarkably, the subjects became better in detecting motion in the paired direction. It is unlikely that they directed their attention to this subthreshold motion stimulus, and yet they learned.

Seitz and Watanabe [87] proposed that neuromodulatory signals can explain these findings if the successful detection of a target letter in the RSVP stream generates an internal reward. In accordance with this view, the pairing of subliminal stimuli only results in learning if they are paired with successfully detected target digits and not if they are paired with targets that are

missed [88]. Moreover, it is possible to replace the internal reward by an external one. A recent experiment tested subjects who were deprived of water and food for several hours and then exposed to an orientation that was paired with water as reward, and also to another orientation not paired with water [89]. The subjects became better in discriminating between orientations around the paired orientation. These findings, taken together, indicate that task-irrelevant learning occurs for subliminal, task-irrelevant features if they are paired with external or internal rewards, which presumably cause release of neuromodulatory factors such as dopamine and acetylcholine [87].

Reconciliation of new results and theories

It is evident that the studies reviewed so far agree about the role of neuromodulatory signals but also that they appear to contradict each other regarding the role of selective attention. Some studies demonstrated an important role for attention in learning while others demonstrated learning for unattended, irrelevant and even imperceptible stimuli. Theories about the mechanisms for learning may appear to be equally contradictory. While AGREL [25] stresses the importance of attention, the model by Seitz and Watanabe [87] indicates that the coincidence of a visual feature and an internal or external reward is sufficient for learning.

To resolve these apparent contradictions, we propose that the attentional feedback signal that enhances the plasticity of task-relevant features in the visual cortex also causes the inhibition of task-irrelevant features so that their plasticity is switched off. We further propose that stimuli that are too weak to be perceived escape from the inhibitory feedback signal so that they are learned if consistently paired with the neuromodulatory signal. This proposal can explain why studies using stimuli close to or below the threshold for perception observed task-irrelevant perceptual learning while studies using supra-threshold stimuli invariably implicate selective attention in learning. A recent study [90] directly compared task-irrelevant learning for a range of stimulus strengths and indeed observed that learning only occurred for motion strengths at or just below the threshold for perception but not for very weak or strong stimuli (Figure 2b). It is easy to understand why very weak motion signals are not learned because they hardly activate the sensory neurons (Figure 2c, left). According to our proposal, the strong motion signals could interfere with the primary letter detection task and are therefore suppressed by the attentional feedback that also blocks plasticity (Figure 2c, right). Threshold stimuli, however, might stay ‘under the radar’ of this attentional inhibition mechanism so that they are not suppressed and can be learned if consistently paired with the neuromodulatory signal (Figure 2c, middle).

Recent results of Tsushima, Sasaki and Watanabe [91] provide further support for this view. They measured the interference caused by irrelevant motion stimuli in the letter detection task of Figure 2 and found an unexpected dependence on signal strength. Weak motion stimuli interfered more than suprathreshold motion stimuli, and an fMRI experiment revealed that they caused stronger activation of motion sensitive area MT+. The reason for the enhanced activation of MT+ by threshold stimuli was observed in the dorsolateral prefrontal cortex (DLPFC), a region that generates attentional inhibition signals. The suprathreshold stimulus activated DLPFC, which then suppressed MT+, while the threshold stimulus did not (Figure 2c). These psychophysical and fMRI results, taken together, indicate that the weak motion signals can indeed escape from the attentional control system so that they can be learned [91].

Concluding remarks

We conclude that there is substantial evidence for an important role for neuromodulatory reward signals and selective attention in the control of perceptual learning. These two factors

can act in concert to implement powerful and neurobiologically plausible learning rules in the cortex. The neuromodulatory signals reveal whether the outcome of a trial is better or worse than expected, while the attentional feedback signal highlights the chain of neurons between sensory and motor cortex responsible for the selected action.

We here proposed that the attentional feedback signals guide learning by suppressing plasticity of irrelevant features while permitting the learning of relevant ones. By hypothesizing that sensory signals that are too weak to be perceived can escape from this inhibitory feedback, we have brought attentional learning theories, like AGREL [25], and theories that emphasized the importance of neuromodulatory signals, like the model of Seitz and Watanabe [87], into a single unified framework.

In most studies on task-irrelevant perceptual learning, attention was focused on the primary RSVP task which was in close proximity to the threshold stimulus to be learned. If task-irrelevant learning and attention-dependent learning are manifestations of the single unifying learning rule, it should even be possible to influence task-irrelevant learning by shifts of selective attention. A recent study [92] manipulated spatial attention by presenting two RSVP streams while instructing subjects to attend only one of them. Task-irrelevant learning occurred for subthreshold stimuli close to the relevant RSVP stream but not for stimuli close to the irrelevant one. In this case the selective attentional signal that gates plasticity has a different origin than the attentional signal in the AGREL model: it could now either come from the instruction to attend one of the RSVP streams or from the response selection stage of the RSVP task. Thus, even the learning of subliminal task-irrelevant stimuli can be brought under attentional control by changing the relevance of nearby suprathreshold stimuli.

In the introduction we asked how sensory neurons can be informed about the relevance of stimuli so that they can sharpen their tuning for features that are important for behavior. The present framework requires two such signals: a global, neuromodulatory signal that signals the rewarded outcome of a trial and an attentional credit assignment signal that restricts plasticity to those sensory neurons that matter in the decision. If acting in concert, these factors can give rise to biologically realistic learning rules that are as powerful as error-backpropagation. Future studies could test the predictions of this new perceptual learning theory, and unravel the mechanisms underlying the interactions between learning, selective attention and reward signaling at the systems level as well as at the cellular and molecular level (Box 2).

Box 1

Attention-gated reinforcement learning (AGREL)

Attention-gated reinforcement learning (AGREL) is a reinforcement learning model that learns by trial and error. AGREL uses two complementary factors to guide learning: (1) a global neuromodulatory signal that informs all synapses whether the outcome of a trial is better or worse than expected; (2) feedback from the response selection stage restricts plasticity to those synapses that were responsible for the behavioral choice.

On every trial, a pattern is presented to the input layer (left in Figure 1a), activity is propagated by feedforward connections to the intermediate layer and from there to the output layer. Output units engage in a stochastic competition, where units with stronger input have a higher likelihood to win. During the action selection process, the winning output unit gives an attentional feedback signal to those lower level units that provided the input during the competition and helped it win. This feedback signal enables plasticity for a selected subset of the connections (thick lines on the right side of Figure 1a). After the response, a reward is delivered or not, and the network computes a global error signal that depends on the difference between the amount of reward that was expected and that is

obtained. If the outcome is better than expected, globally released neuromodulators (green in Figure 1a) cause synapses between active units to increase in strength. In case of a wrong response, the neuromodulators produce a decrease in synapse strength (red in Figure 1a).

The problem of adjusting the deep weights v_{ij} in the network proceeds according to the following rule:

$$\Delta v_{ij} = \beta \cdot X_i \cdot f(Y_j) \cdot g(\text{reward}) \cdot FB_{sj} \quad (1)$$

Here Δv_{ij} is the change in synaptic strength, β is the learning rate, X_i and Y_j are pre- and postsynaptic activity, and FB_{sj} is the attentional feedback signal from the winning output unit s . Thus, the plasticity depends on the product of pre- and postsynaptic activity and is gated by the reward through the function g and by the attentional feedback FB_{sj} that conveys whether unit Y_j is responsible for selected action s . This feedback signal could operate by enhancing the plasticity of unit j or by blocking plasticity of other units $Y_l, l \neq j$, by providing inhibition directly or indirectly (Figure 1b). Note that these signals are all available locally, at the synapse. The learning rule for the connections w_{jk} to the output layer is even simpler because it does not depend on the attentional feedback signal. On average, AGREL changes the connections in the same way as error-backpropagation, a learning rule that is popular and efficient but usually considered to be biologically implausible [93]. Further details about the AGREL model can be found in [25].

Box 2

Questions for future research

- Which neuromodulatory signals determine learning?
- When are the neuromodulatory signals released?
- How do neuromodulatory signals and selective attention interact with each other?
- Which neurotransmitter receptors convey the reward signals and which ones the effect of attention?
- How do lateral inhibition and top-down inhibition contribute to the gating of plasticity?
- Can the AGREL learning rules be generalized to tasks other than categorization?
- How should perceptual learning in delayed match to sample tasks be modeled?
- Can purely bottom-up signals lead to plasticity without attention or reward signals?

Acknowledgments

The work on AGREL was supported by an NWO-Exact grant. PRR was supported by an NWO-VICI and an HFSP grant and TW by NIH R21 EY017737, NIH R21 EY018925, NIH R01, EY015980-04A2, NIH R01 EY019466, NSF BCS-PR04-137, NSF BCS-0549036, and HFSP-RGP0018.

Glossary

AGREL	attention-gated reinforcement learning
Feedforward connection	propagates information from lower to higher levels
Feedback connection	propagates information from higher back to lower levels

Perceptual learning	improvement of perception through learning
Selective attention	behavioral selection of one representation over another one
Neuromodulatory systems	systems that release neuromodulators to code the rewarded outcome of a trial
RF	receptive field

Reference List

- Gauthier I, et al. Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neurosci* 2000;3:191–197. [PubMed: 10649576]
- Fahle, M.; Poggio, T. *Perceptual learning*. The MIT Press; 2002.
- Watanabe T, et al. Greater plasticity in lower-level than higher-level visual motion processing in a passive perceptual learning task. *Nature Neurosci* 2002;5:1003–1009. [PubMed: 12219093]
- Karni A, Sagi D. The time course of learning a visual skill. *Nature* 1993;365:250–252. [PubMed: 8371779]
- Schoups A, et al. Practising orientation identification improves orientation coding in V1 neurons. *Nature* 2001;412:549–553. [PubMed: 11484056]
- Saarinen J, Levi DM. Perceptual learning in vernier acuity: what is learned? *Vision Res* 1995;35:519–527. [PubMed: 7900292]
- Ball K, Sekuler R. Direction-specific improvement in motion discrimination. *Vision Res* 1987;27:953–965. [PubMed: 3660656]
- Watanabe T, et al. Perceptual learning without perception. *Nature* 2001;413:844–847. [PubMed: 11677607]
- Adini Y, et al. Context-enabled learning in the human visual system. *Nature* 2002;415:790–793. [PubMed: 11845209]
- Yu C, et al. Perceptual learning in contrast discrimination and the (minimal) role of context. *J. Vision* 2004;4:169–182.
- Schoups AA, et al. Human perceptual learning in identifying the oblique orientation: retinotopy, orientation specificity and monularity. *J Physiol* 1995;483:797–810. [PubMed: 7776259]
- Crist RE, et al. Perceptual learning of spatial localization: specificity for orientation, position, and context. *J Neurophysiol* 1997;78:2889–2894. [PubMed: 9405509]
- Liu Z. Perceptual learning in motion discrimination that generalizes across motion directions. *Proc. Natl. Acad. Sci. USA* 1999;96:14085–14087. [PubMed: 10570202]
- Xiao LQ, et al. Complete transfer of perceptual learning across retinal locations enabled by double training. *Curr. Biol* 2008;18:1922–1926. [PubMed: 19062277]
- Law CT, Gold JI. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci* 2008;11:505–513. [PubMed: 18327253]
- Petrov AA, et al. The dynamics of perceptual learning: an incremental reweighting model. *Psychol Rev* 2005;112:715–743. [PubMed: 16262466]
- Chowdhury SA, DeAngelis GC. Fine discrimination training alters the causal contribution of macaque area MT to depth perception. *Neuron* 2008;60:367–377. [PubMed: 18957227]
- Ahissar M, Hochstein S. Attentional control of early perceptual learning. *Proc. Natl. Acad. Sci. USA* 1993;90:5718–5722. [PubMed: 8516322]
- Ahissar M, Hochstein S. Task difficulty and the specificity of perceptual learning. *Nature* 1997;387:401–406. [PubMed: 9163425]
- Yotsumoto Y, Watanabe T. Defining a link between perceptual learning and attention. *PLoS Biol* 2008;6:e221–e223. [PubMed: 18752357]
- Zhang J-Y, et al. Stimulus coding rules for perceptual learning. *PLoS Biol* 2008;6:1651–1660.
- Gutnisky DA, et al. Attention alters visual plasticity during exposure-based learning. *Curr. Biol* 2009;19:555–560. [PubMed: 19268592]

23. Poggio T, et al. Fast perceptual learning in visual hyperacuity. *Science* 1992;256:1018–1021. [PubMed: 1589770]
24. Herzog MH, Fahle M. The role of feedback in learning a vernier discrimination task. *Vision Res* 1997;37:2133–2141. [PubMed: 9327060]
25. Roelfsema PR, van Ooyen A. Attention-gated reinforcement learning of internal representations for classification. *Neural Comp* 2005;17:2176–2214.
26. Kourtzi Z, et al. Distributed neural plasticity for shape learning in the human visual cortex. *PLoS Biol* 2005;3:1317–1327.
27. Schwartz S, et al. Neural correlates of perceptual learning: a functional MRI study of visual texture discrimination. *Proc. Natl. Acad. Sci. USA* 2002;99:17137–17142. [PubMed: 12446842]
28. Furmanski CS, et al. Learning strengthens the response of primary visual cortex to simple patterns. *Curr. Biol* 2004;14:573–578. [PubMed: 15062097]
29. Yotsumoto Y, et al. Location-specific cortical activation changes during sleep after training for perceptual learning. *Curr. Biol* 2009;19:1278–1282. [PubMed: 19576772]
30. Yotsumoto Y, et al. Different dynamics of performance and brain activation in the time course of perceptual learning. *Neuron* 2008;57:827–833. [PubMed: 18367084]
31. Walker MP, et al. The functional anatomy of sleep-dependent visual skill learning. *Cereb. Cortex* 2005;15:1666–1675. [PubMed: 15703253]
32. Yang T, Maunsell JHR. The effect of perceptual learning on neuronal responses in monkey visual area V4. *J. Neurosci* 2004;24:1617–1626. [PubMed: 14973244]
33. Li W, et al. Learning to link visual contours. *Neuron* 2008;57:442–451. [PubMed: 18255036]
34. Rainer G, et al. The effect of learning on the function of monkey extrastriate cortex. *PLoS Biol* 2004;2:275–283.
35. Raiguel S, et al. Learning to see the difference specifically alters most informative V4 neurons. *J. Neurosci* 2006;26:6589–6602. [PubMed: 16775147]
36. Freedman DJ, et al. A comparison of primate prefrontal cortex and inferior temporal cortices during visual categorization. *J. Neurosci* 2003;23:5235–5246. [PubMed: 12832548]
37. Freedman DJ, Assad JA. Experience dependent representation of visual categories in parietal cortex. *Nature* 2006;443:85–88. [PubMed: 16936716]
38. Freedman DJ, et al. Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 2001;291:312–316. [PubMed: 11209083]
39. Oristaglio J, et al. Integration of visuospatial and effector information during symbolically cued limb movements in monkey lateral intraparietal cortex. *J. Neurosci* 2006;26:8310–8319. [PubMed: 16899726]
40. Mirabella G, et al. Neurons in area V4 of the macaque translate attended visual features into behaviorally relevant categories. *Neuron* 2007;54:303–318. [PubMed: 17442250]
41. Sigala N, Logothetis NK. Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature* 2002;415:318–320. [PubMed: 11797008]
42. De Baene W, et al. Effects of category learning on the stimulus selectivity of macaque inferior temporal neurons. *Learning & Memory* 2008;15:717–727. [PubMed: 18772261]
43. Li S, et al. Learning shapes the representation of behavioral choice in the human brain. *Neuron* 2009;62:441–452. [PubMed: 19447098]
44. Hochstein S, Ahissar M. Views from the top: hierarchies and reverse hierarchies in the visual system. *Neuron* 2002;36:791–804. [PubMed: 12467584]
45. Jeter PE, et al. Task precision at transfer determines specificity of perceptual learning. *J. Vision* 2009;9:1–13.
46. Schultz W. Getting formal with dopamine and reward. *Neuron* 2002;36:241–263. [PubMed: 12383780]
47. Stuber GD, et al. Reward-predictive cues enhance excitatory synaptic strength onto midbrain dopamine neurons. *Science* 2008;321:1690–1692. [PubMed: 18802002]
48. Shen W, et al. Dichotomous dopaminergic control of striatal synaptic plasticity. *Science* 2008;321:848–851. [PubMed: 18687967]

49. Bao S, et al. Cortical remodelling induced by activity of ventral tegmental dopamine neurons. *Nature* 2001;412:79–81. [PubMed: 11452310]
50. Richardson RT, DeLong MR. Nucleus basalis of Meynert neuronal activity during a delayed response task in monkey. *Brain Res* 1986;399:364–368. [PubMed: 3828770]
51. Bakin JS, Weinberger NM. Induction of a physiological memory in the cerebral cortex by stimulation of the nucleus basalis. *Proc. Natl. Acad. Sci. USA* 1996;93:11219–11224. [PubMed: 8855336]
52. Kilgard MP, Merzenich MM. Cortical map reorganization enabled by nucleus basalis activity. *Science* 1998;279:1714–1718. [PubMed: 9497289]
53. Seol GH, et al. Neuromodulators control the polarity of spike-timing-dependent plasticity. *Neuron* 2007;55:919–929. [PubMed: 17880895]
54. Juliano SL, et al. Cholinergic depletion prevents expansion of topographic maps in somatosensory cortex. *Proc. Natl. Acad. Sci. USA* 1991;88:780–784. [PubMed: 1992469]
55. Winkler J, et al. Essential role of neocortical acetylcholine in spatial memory. *Nature* 1995;375:484–487. [PubMed: 7777056]
56. Easton A, et al. Unilateral lesions of the cholinergic basal forebrain and fornix in one hemisphere and inferior temporal cortex in the opposite hemisphere produce severe learning impairments in rhesus monkeys. *Cereb. Cortex* 2002;12:729–736. [PubMed: 12050084]
57. Warburton EC, et al. Cholinergic neurotransmission is essential for perirhinal cortical plasticity and recognition memory. *Neuron* 2003;38:987–996. [PubMed: 12818183]
58. Fan J, et al. Testing the efficiency and independence of attentional networks. *J. Cognit. Neurosci* 2002;14:340–347. [PubMed: 11970796]
59. Trabasso, T.; Bower, GH. *Attention in Learning: Theory and Research*. Krieger Pub Co; 1968.
60. Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1991;1:1–47. [PubMed: 1822724]
61. Desimone R, Duncan J. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci* 1995;18:193–222. [PubMed: 7605061]
62. Rizzolatti G, et al. Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* 1987;25:31–40. [PubMed: 3574648]
63. Allport, A. Visual attention. In: Posner, MI., editor. *Foundations of cognitive science*. 0 edn. MIT Press; 1989. p. 631–682.
64. Hoffman JE, Subramaniam B. The role of visual attention in saccadic eye movements. *Percept. Psychophys* 1995;57:787–795. [PubMed: 7651803]
65. Kowler E, et al. The role of attention in the programming of saccades. *Vision Res* 1995;35:1897–1916. [PubMed: 7660596]
66. Deubel H, Schneider WX. Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Res* 1996;36:1827–1837. [PubMed: 8759451]
67. Chelazzi L, et al. Responses of neurons in macaque area V4 during memory-guided visual search. *Cereb. Cortex* 2001;11:761–772. [PubMed: 11459766]
68. Schall JD, Thompson KG. Neural selection and control of visually guided eye movements. *Annu. Rev. Neurosci* 1999;22:241–259. [PubMed: 10202539]
69. Gregoriou GG, et al. High-frequency, long-range coupling between prefrontal and visual cortex during attention (324 edn) 2009:1207–1209.
70. Khayat PS, et al. Time course of attentional modulation in the frontal eye field during curve tracing. *J. Neurophysiol* 2009;101:1813–1822. [PubMed: 19176609]
71. Thompson KG, et al. Neuronal basis of covert spatial attention in the frontal eye field. *J. Neurosci* 2005;25:9479–9487. [PubMed: 16221858]
72. Thorpe S, et al. Speed of processing in the human visual system. *Nature* 1996;381:520–522. [PubMed: 8632824]
73. Lamme VAF, Roelfsema PR. The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci* 2000;23:571–579. [PubMed: 11074267]
74. Moore T, Armstrong KM. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 2003;421:370–373. [PubMed: 12540901]

75. Ekstrom LB, et al. Bottom-up dependent gating of frontal signals in early visual cortex. *Science* 2008;321:414–417. [PubMed: 18635806]
76. Scholte HS, et al. The spatial profile of visual attention in mental curve tracing. *Vision Res* 2001;41:2569–2580. [PubMed: 11520504]
77. Ullman S. Sequence seeking and counterstreams: a computational model for bidirectional information flow in the visual cortex. *Cereb. Cortex* 1995;5:1–11. [PubMed: 7719126]
78. Jiang Y, Chun MM. Selective attention modulates implicit learning. *Q. J. Exp. Psychol* 2001;54:1105–1124.
79. Baker CI, et al. Role of attention and perceptual grouping in visual statistical learning. *Psychol. Sci* 2004;15:460–466. [PubMed: 15200630]
80. Turk-Browne NB, et al. The automaticity of visual statistical learning. *J. Exp. Psychol. : Gen* 2005;134:552–564. [PubMed: 16316291]
81. Self M, et al. Feedforward and feedback processing utilise different glutamate receptors 2008;769.14.
82. Dinse H, et al. Pharmacological modulation of perceptual learning and associated cortical reorganization. *Science* 2003;301:91–94. [PubMed: 12843392]
83. Sheng M, Kim MJ. Postsynaptic signaling and plasticity mechanisms. *Science* 2002;298:776–780. [PubMed: 12399578]
84. Herrero J, et al. Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 2008;454:1110–1114. [PubMed: 18633352]
85. Watling L, et al. The attention-gated reinforcement learning model: performance and predictions. *Front. Syst. Neurosci.* 2009 Abstract, doi: 10.3389/conf.neuro.06.2009.03.010.
86. Seitz AR, Watanabe T. Is subliminal learning really passive? *Nature* 2003;422:36. [PubMed: 12621425]
87. Seitz AR, Watanabe T. A unified model for perceptual learning. *Trends Cogn. Sci* 2005;9:329–334. [PubMed: 15955722]
88. Seitz AR, et al. Requirement for high-level processing in subliminal learning. *Curr. Biol* 2005;15:R1–R3. [PubMed: 15649341]
89. Seitz AR, et al. Rewards evoked learning of unconsciously processed visual stimuli in adult humans. *Neuron* 2009;61:700–707. [PubMed: 19285467]
90. Tsushima Y, et al. Task-irrelevant learning occurs only when the irrelevant feature is weak. *Curr. Biol* 2008;18:R516–R517. [PubMed: 18579092]
91. Tsushima Y, et al. Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science* 2006;314:1786–1788. [PubMed: 17170308]
92. Nishina S, et al. Effect of spatial distance to the task stimulus on task-irrelevant perceptual learning of static gabors. *J. Vision* 2007;7:1–10.
93. Crick F. The recent excitement about neural networks. *Nature* 1989;337:129–132. [PubMed: 2911347]

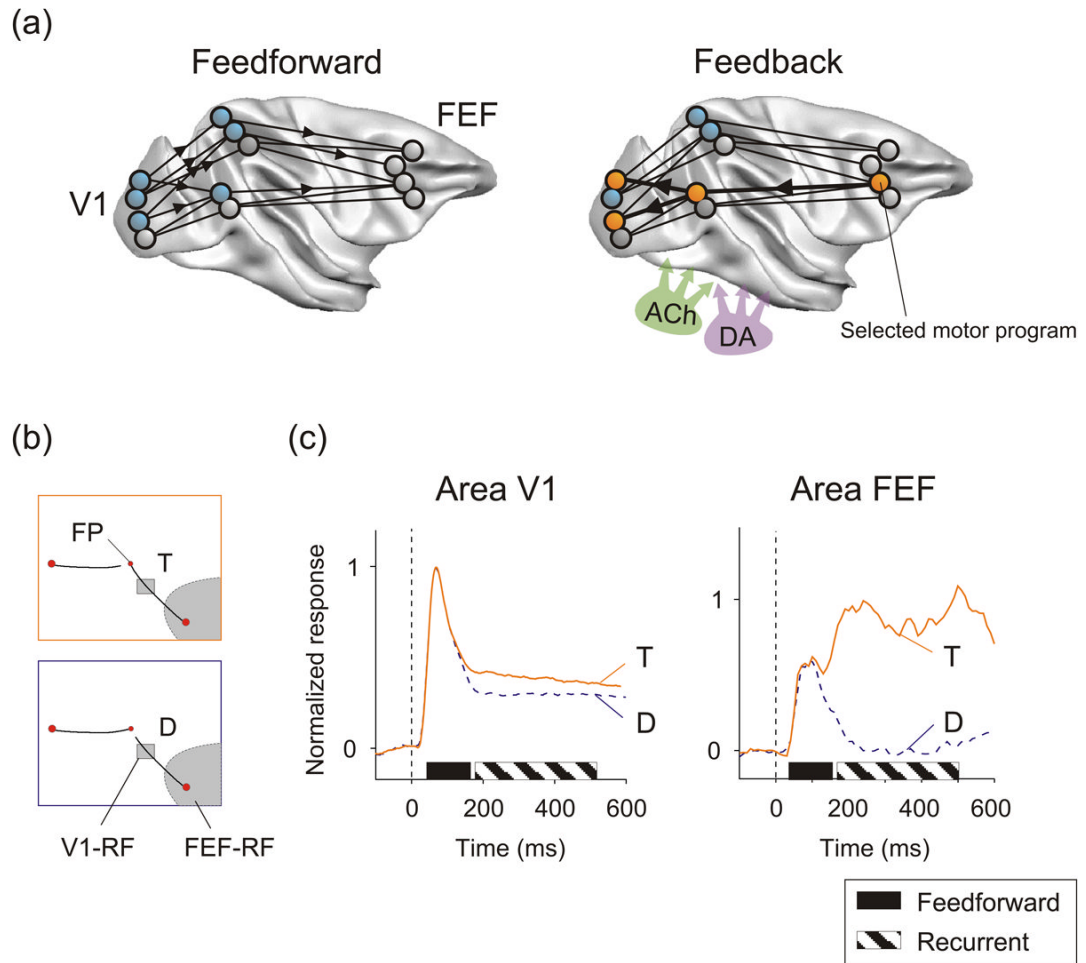


Figure 1. Factors that modulate visual cortical plasticity

(a) Left, a visual stimulus is initially registered by a fast feedforward wave of activity that activates feature selective neurons in the many areas of the visual cortex. Right, neurons in the frontal cortex engage in a competition to select a behavioral response, and the neurons that win the competition feed attentional signals back to the visual cortex. Neuromodulatory systems, including acetylcholine (green) and dopamine (purple), modulate the activity as well as the plasticity of sensory representations. (b) Contour grouping task where monkeys have to trace a target curve (T) that is connected to a fixation point (FP). A red circle at the end of this curve is the target for an eye movement (Figure 1b). The animals have to ignore a distractor curve (D). The two stimuli differ so that the RFs of neurons in areas V1 and FEF are either on the target curve (upper panel) or on the distractor curve (lower panel). (c) Neurons in the frontal cortex (area FEF) and the visual cortex (area V1) enhance their response when the target curve falls in their receptive field. Note that this attentional modulation comes at a delay (striped bar), while the initial neuronal responses do not discriminate between the relevant and irrelevant curve (black bar). Modified from Khayat et al. [70].

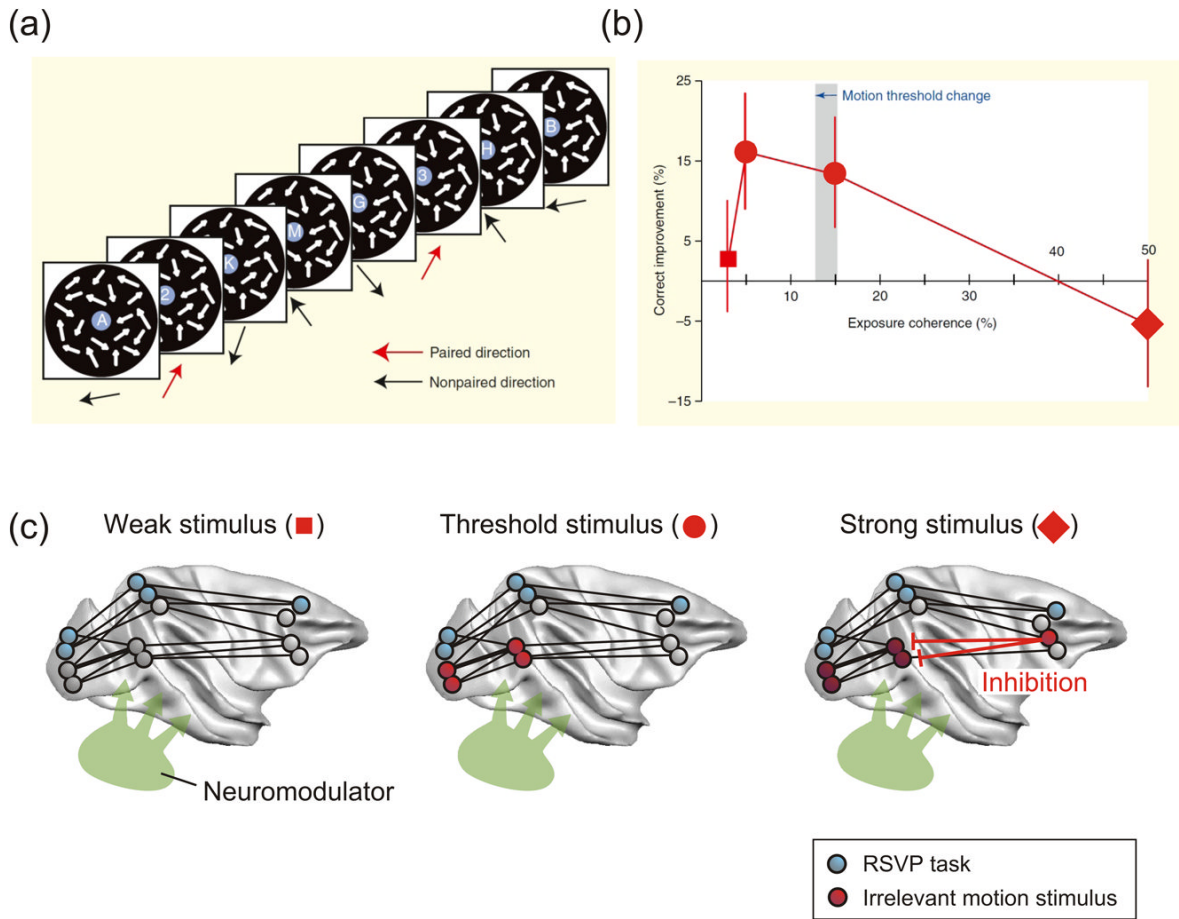


Figure 2. Task-irrelevant perceptual learning

(a) Subjects monitor a central RSVP stream to detect target digits that are presented among letter distractors. In the background, motion signals are presented that are irrelevant for the subject's task. (b) A separate test determines the motion sensitivity of the subjects. This test reveals that the sensitivity for the motion direction paired with the target digits (red arrows in a) increases. Learning occurs for paired motion stimuli that are at or below the visibility threshold, but not for very weak or strong motion stimuli. (c) Blue circles denote neurons involved in the RSVP task. Red circles denote the representation of the irrelevant motion signals. **Left**, weak motion stimuli are not registered well by the visual cortex and are not learned. **Middle**, threshold motion stimuli are represented by motion sensitive neurons and are learned if paired with the neuromodulatory signal. **Right**, strong motion signals might cause interference and therefore receive top-down inhibition from the frontal cortex. These suppressive signals block perceptual learning. Adapted from Tsushima et al. [90].

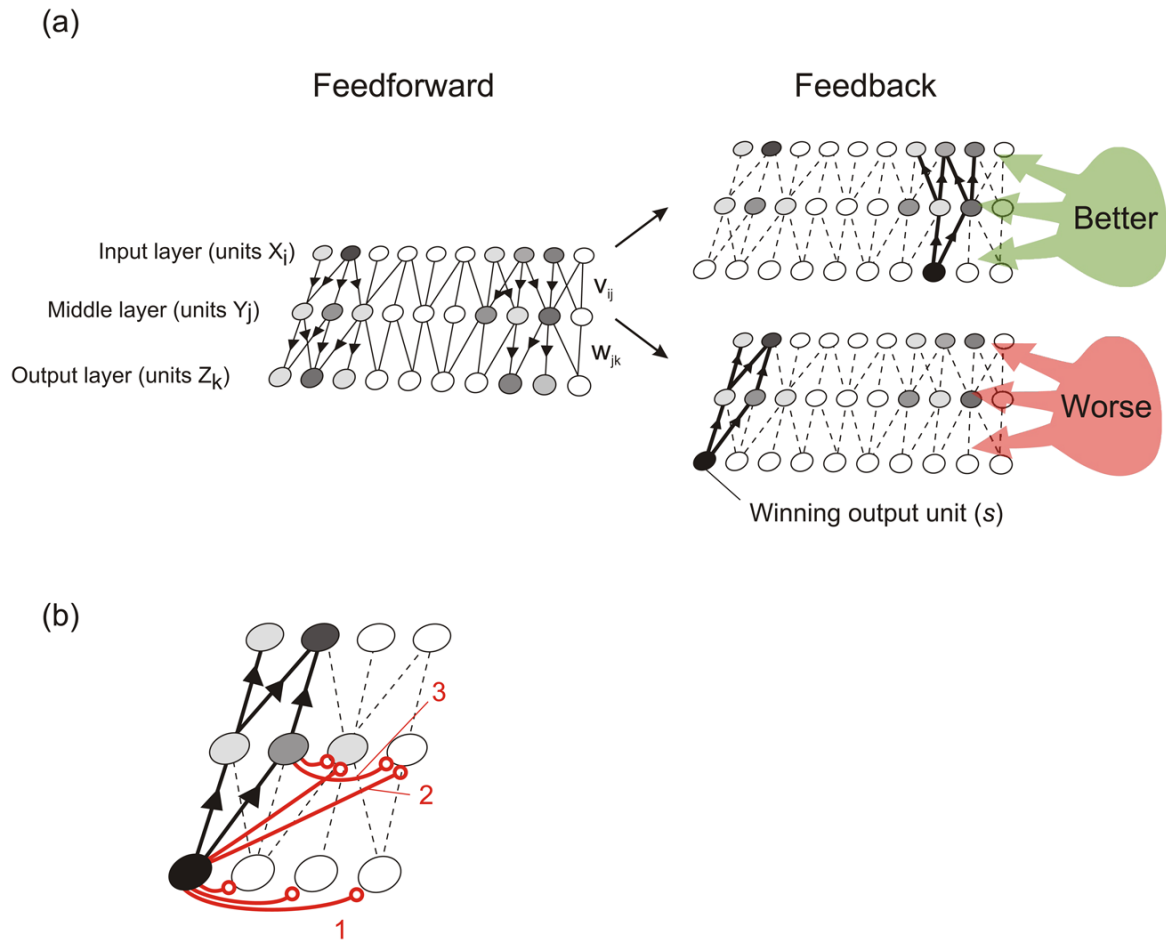


Figure I. Attention-gated reinforcement learning

(a) **Left**, activity is propagated from lower to higher layers through feedforward connections. The output units engage in a competition to determine the selected action. **Right**, if an action has been selected, the winning output units provide an attentional feedback signal that highlights the lower level units responsible for the selected action (thick connections) enabling their plasticity. The plasticity of other connections is blocked by inhibition (dashed connections). Different actions thereby enable plasticity for different sensory neurons. Neuromodulators indicate whether the rewarded outcome was better (green) or worse (red) than expected. (b) Inhibition (red connections) blocks the plasticity of connections that are not involved in the selected action. At the output level, the actions compete through inhibitory interactions (connection type 1). The winning output unit could directly inhibit neurons at lower levels (connection type 2) or indirectly through inhibitory lateral interactions between the lower level units (connection type 3). These inhibitory effects do not occur for stimuli too weak to be consciously perceived.