



## Perceptual learning, roving and the unsupervised bias

Michael H. Herzog<sup>a,\*</sup>, Kristoffer C. Aberg<sup>a,b</sup>, Nicolas Frémaux<sup>c</sup>, Wulfram Gerstner<sup>c</sup>, Henning Sprekeler<sup>c,d</sup>

<sup>a</sup>Laboratory of Psychophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

<sup>b</sup>Laboratory for Neurology and Imaging of Cognition, Department of Neurosciences, University of Geneva, Switzerland

<sup>c</sup>Laboratory of Computational Neuroscience, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

<sup>d</sup>Institute for Theoretical Biology, Humboldt-University Berlin, Invalidenstr. 43, 10115 Berlin, Germany

### ARTICLE INFO

#### Article history:

Received 1 June 2011

Received in revised form 13 October 2011

Available online 19 November 2011

#### Keywords:

Perceptual learning

Neural networks

Roving

Bisection stimuli

### ABSTRACT

Perceptual learning improves perception through training. Perceptual learning improves with most stimulus types but fails when *certain* stimulus types are mixed during training (roving). This result is surprising because classical supervised and unsupervised neural network models can cope easily with roving conditions. What makes humans so inferior compared to these models? As experimental and conceptual work has shown, human perceptual learning is neither supervised nor unsupervised but reward-based learning. Reward-based learning suffers from the so-called unsupervised bias, i.e., to prevent synaptic “drift”, the average reward has to be exactly estimated. However, this is impossible when two or more stimulus types with different rewards are presented during training (and the reward is estimated by a running average). For this reason, we propose no learning occurs in roving conditions. However, roving hinders perceptual learning only for combinations of similar stimulus types but not for dissimilar ones. In this latter case, we propose that a critic can estimate the reward for each stimulus type separately. One implication of our analysis is that the critic cannot be located in the visual system.

© 2011 Elsevier Ltd. All rights reserved.

### 1. Perceptual learning, roving and neural overlap

Perceptual learning is the ability to improve perception through training. Perceptual learning occurs with motion (Ball & Sekuler, 1982; Koyama, Harner, & Watanabe, 2004; Kuai et al., 2005), vernier (Herzog & Fahle, 1997), bisection (Crist et al., 1997), Gabor (Yu, Klein, & Levi, 2004; Zhang et al., 2008), and many other stimuli (for reviews: Fahle & Poggio, 2002; Sagi, 2011). In perceptual learning experiments, usually one out of two stimulus alternatives of one stimulus type is presented per trial, for example, a bisection stimulus which is either offset to the left or right. With this setup, performance improves (Fig. 1A; Aberg & Herzog, 2010; Otto et al., 2006; Tartaglia, Aberg, & Herzog, 2009).

Interestingly, perceptual learning does not occur when certain stimulus types are combined during training, i.e. under so called roving conditions (Banai et al., 2009; Kuai et al., 2005; Otto et al., 2006; Parkosadze et al., 2008; Tartaglia, Aberg, & Herzog, 2009; Yu, Klein, & Levi, 2004; Zhang et al., 2008). For example, there is no improvement of performance when bisection stimuli with outer line distances of 20' (arcmin) and 30' are presented randomly interleaved (Fig. 1D). This result surprises because perceptual learning occurs when each of the two types of bisection stimuli is trained in separate sessions (Fig. 1B for the 20' bisection stimu-

lus). It is the combination of the stimulus types which makes perceptual learning impossible.

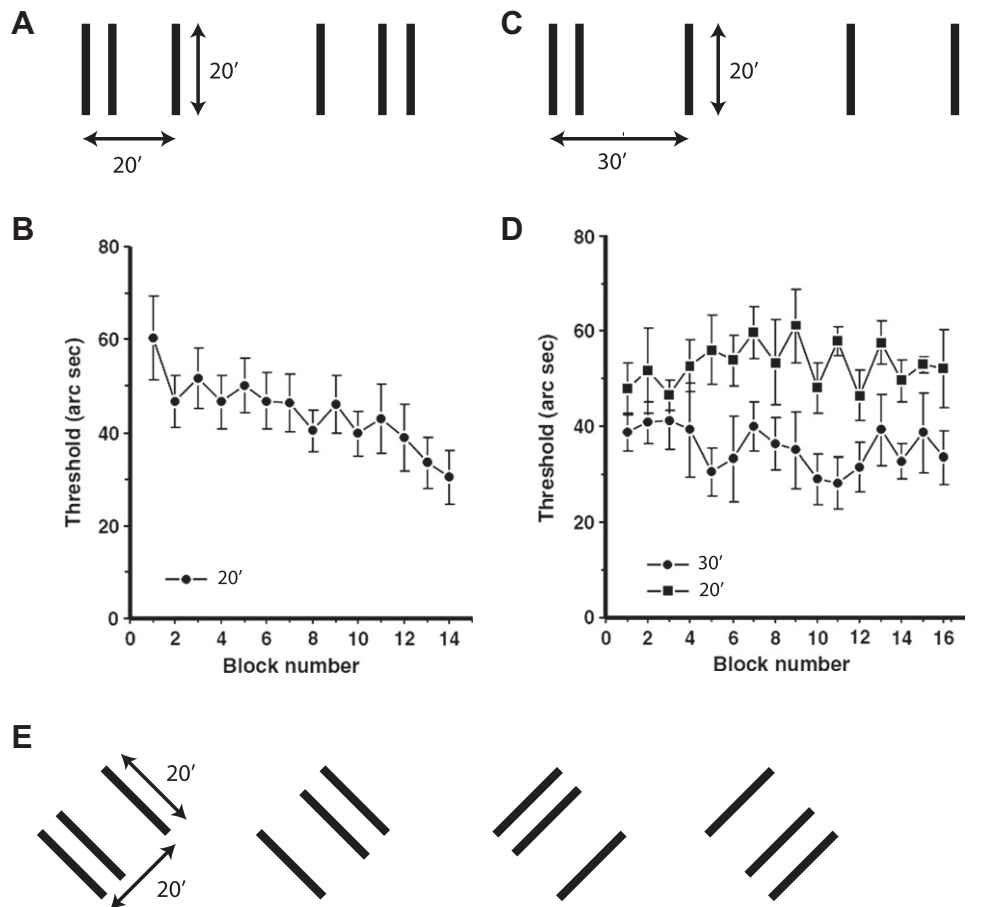
However, not all combinations of stimulus types hinder perceptual learning. For example, perceptual learning occurs when bisection stimuli of different orientations are roved (Fig. 1E; Tartaglia, Aberg, & Herzog, 2009). We proposed that roving hinders perceptual learning when stimulus types are sufficiently but not too much different involving synaptic changes in overlapping but not identical populations of neurons (Tartaglia, Aberg, & Herzog, 2009). When stimulus types are almost identical, e.g. bisection stimuli with outer line distances of 20' and 20.01', perceptual learning must occur because the stimuli are nearly identical and the very same neurons are involved in the processing (small head movements towards the screen yield variations in stimulus size much larger than 0.01'). On the other extreme, roving visual and acoustical stimuli does not yield interference simply because stimuli are too different (otherwise, any interleaved learning would be impossible in general). It should be mentioned that roving does not hinder perceptual learning in principle, as previously claimed (Yu, Klein, & Levi, 2004), but just slows it down by a factor of about 15 (Parkosadze et al., 2008).

### 2. Models, feedback and the unsupervised bias

Learning models are supervised, unsupervised, or reward-based. Most models of perceptual learning are supervised or unsupervised models of the neural network type (supervised

\* Corresponding author at: Ecole Polytechnique Fédérale de Lausanne (EPFL), EPFL-SV, Station 15, 1015 Lausanne EPFL, Switzerland. Fax: +41 216939645.

E-mail address: [michael.herzog@epfl.ch](mailto:michael.herzog@epfl.ch) (M.H. Herzog).



**Fig. 1.** Roving bisection stimuli. (A) A bisection stimulus comprises two outer lines and a central line which bisects the interval between the outer lines (here, the outer line distance is 20'). The central line is randomly either closer to the left or right outer line. Observers indicate this offset direction. (B) Performance improves with training, i.e. thresholds decrease. (C) A bisection stimulus with a larger outer line spacing of 30'. (D) In roving experiments, all four stimulus alternatives from (A) and (C) are presented randomly interleaved. Performance does not improve with training. (E) Orthogonal bisection stimuli. When all four alternatives are presented randomly interleaved, performance improves with training (results not shown). Figures adapted with permission from Otto et al. (2006).

perceptron models, e.g. Doshier & Lu, 1998; Petrov, Doshier, & Lu, 2005; Poggio, Fahle, & Edelman, 1992; Sotiropoulos, Seitz, & Series, 2011; unsupervised Hebbian learning type: Sundareswaran & Vaina, 1996). A stimulus is presented which activates neurons in an input layer. This activation is propagated through the network to an output neuron. Synaptic weights within the network are changed in order to minimize the error of the desired vs. actual output (supervised learning) or by local Hebbian rules (unsupervised learning).

In Hebbian learning, the synaptic weight  $w_{ij}$  between a pre-synaptic neuron  $i$  and a post-synaptic neuron  $j$  increases when the activities of both neurons are high. In its simplest form  $w_{ij}$  is updated by the following learning rule:

$$\Delta w_{ij} = pre_i \times post_j. \quad (1)$$

$\Delta w_{ij}$  is the change of the synaptic weight  $w_{ij}$ ,  $pre_i$  and  $post_j$  are the activities of the pre- and post-synaptic neurons, respectively.

In supervised learning, the network learns to assign a desired class label  $F(x)$  to each stimulus  $x$ , e.g. the networks learn to return  $-1$  for a left and  $+1$  for a right offset bisection stimulus, respectively, independent of the absolute offset size. During training, the actual class label  $F(x)^t$  in each trial  $t$  is compared with the desired label  $F(x)$  and an error term is computed, e.g.  $E = F(x)^t - F(x)$ .  $E$  is "propagated" to each synapse as an "individual" error term  $E_{ij}$  which determines whether  $w_{ij}$  is up- or down regulated

$$\Delta w_{ij} = pre_i \times E_{ij}. \quad (2)$$

Models of supervised learning can reproduce an amazing set of empirical data. However, all models fail to reproduce the results of roving experiments – paradoxically because models learn but humans do not. This is even more surprising since the roved bisection stimuli are linearly separable and, hence, can be learned by the simplest models (interestingly, these models can cope only with linearly separable stimuli). Hence, humans fail in the simplest learning situation. What makes humans so inferior compared to these models?

Perceptual learning is neither supervised nor unsupervised because perceptual learning can occur without feedback, ruling out purely supervised models. On the other hand, feedback speeds up learning compared to no-feedback conditions ruling out purely unsupervised models (Herzog & Fahle, 1997). Hence, we propose that perceptual learning is reward-based learning. Reward signals are sluggish and often based on an average performance measure such as tomatoes of sort A taste better on *average* than tomatoes of sort B. This is in accordance with experiments showing that observers can use block feedback almost as efficiently as trial-by-trial error feedback (Herzog & Fahle, 1998; Shibata et al., 2009). In block feedback conditions, a performance score is delivered after a number of trials. For example, the percentage of correct responses for 40 trials is shown on the computer screen. By comparing consecutive scores, internal feedback can be computed which can be used for learning (Herzog & Fahle, 1998). However, block feedback can clearly not be used in supervised learning which requires exact feedback in *each* trial.

Reward-based learning is usually modeled as Hebbian learning augmented by a reward rather than a teacher term. There are two major classes of reward-based learning rules (Frémaux, Sprekeler, & Gerstner, 2010). In the  $R_{max}$  learning rule, synaptic changes rely on reward  $R$  and on pre- and post-synaptic activity:

$$\Delta w_{ij} = pre_i \times (post_j - \overline{post_j}) \times R. \quad (3)$$

$R$  is a reward signal which is identical for all synapses whereas  $E_{ij}$  in Eq. (2) is computed for each synapse individually. The idea of the  $R_{max}$  rule is straightforward. The synaptic weight  $w_{ij}$  increases if higher than average activity  $\overline{post_j}$  leads to more reward.

However,  $R_{max}$  learning rules, as supervised learning rules, do not suffer from roving conditions. Hence,  $R_{max}$  rules fail, as supervised learning rules, because they are too powerful compared to human learning.

The second reward-based learning rule  $R_{av}$  augments Hebbian learning by a reward prediction error and, hence, relies on reward prediction  $\bar{R}$  rather than on activity prediction  $\overline{post_j}$  (Farries & Fair-

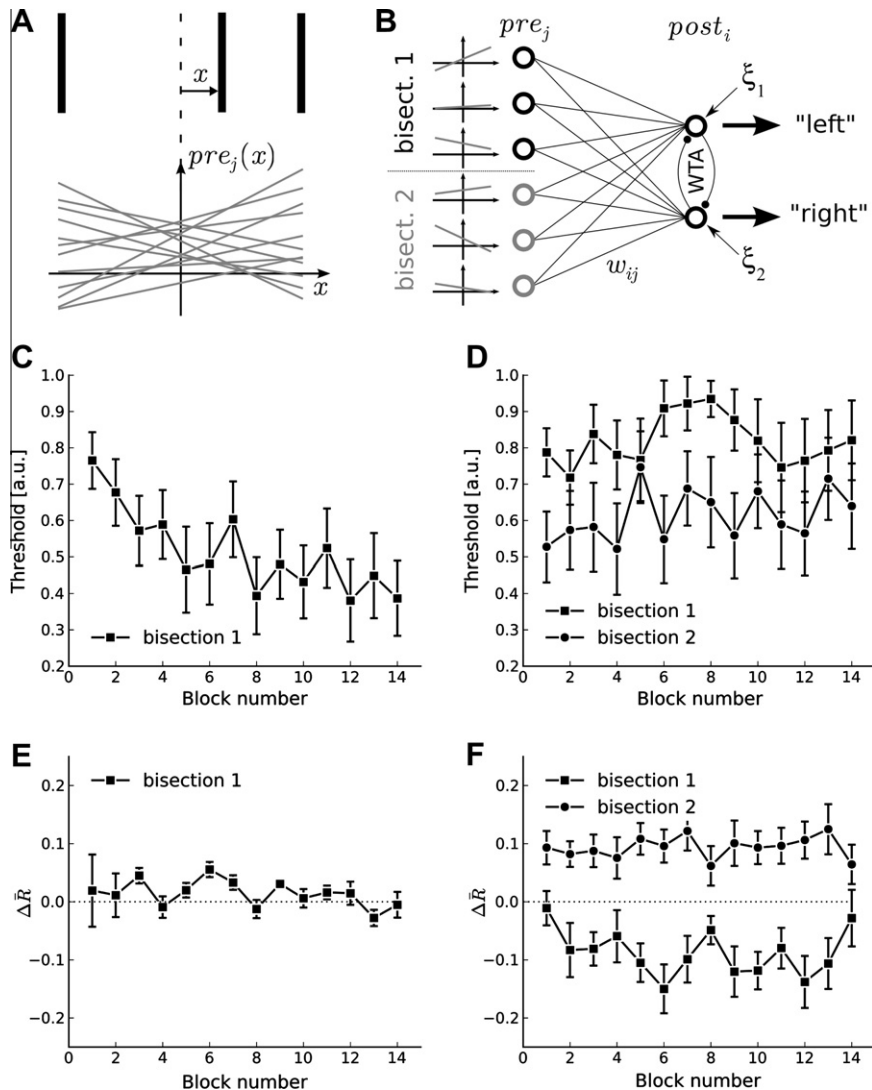
hall, 2007; Frémaux, Sprekeler, & Gerstner, 2010; Izhikevich, 2007; Legenstein, Pecevski, & Maass, 2008):

$$\Delta w_{ij} = pre_i \times post_j \times (R - \bar{R}). \quad (4)$$

This learning rule increases reward when the prediction of the mean reward  $\bar{R}$  is correctly estimated. If reward prediction is not exact, this learning rule suffers from an additional *unsupervised bias* term (Frémaux, Sprekeler, & Gerstner, 2010; Loewenstein, 2008):

$$\Delta w_{ij} = pre_i \times post_j \times (R - \bar{R}) + pre_i \times post_j \times \Delta \bar{R}, \quad (5)$$

where  $\Delta \bar{R}$  denotes the error of the predicted mean reward. The first part of this equation is the “good guy” increasing the reward by improving performance. The unsupervised bias  $pre_i \times post_j \times \Delta \bar{R}$  is the “bad guy”. It depends on reward averages, but not on the correlation of actual reward and activity in a given trial. Therefore, it cannot relate feedback with performance and – instead of increasing the reward – causes a drift of the synaptic weight  $w_{ij}$  in a potentially



**Fig. 2.** Simulation experiment. (A) Presynaptic neurons have randomly chosen linear tuning functions. (B) Presynaptic neurons respond either to the bisection stimulus with smaller or wider outer lines distance (corresponding to the 20' and 30' bisection stimuli in the psychophysical experiments, respectively). A simple neural network, with a winner-takes-all mechanism, determines the offset direction of the center line. The network is trained with the  $R_{av}$  learning rule. (C and D) Simulation results corresponding to the experimental results of Fig. 1C and D. (C) Training with the 20' bisection stimulus improves performance. (D) No learning occurs when both types of bisection stimuli are presented interleaved during training. (E and F) Show the reward prediction error  $\Delta \bar{R}$  corresponding to (C) and (D), respectively. (E) The reward prediction is roughly correct when only one bisection stimulus is used. (F) Reward predictions for the two-bisection task are consistently different from zero explaining why no learning occurs when both types of bisection stimuli are presented during training.

unrewarding direction. This unsupervised bias effect can hinder or strongly slow down learning, even when the bias is small (Frémaux, Sprekeler, & Gerstner, 2010; Loewenstein, 2008).

A simple method to obtain an average reward is a running average. A running average, as an estimate of  $\bar{R}$ , works well when only one stimulus type is presented. However, no or slowed learning occurs when two stimulus types are roved and reward for each stimulus type is different. Different tasks come usually with different rewards because difficulty varies strongly amongst tasks. For example, the offsets of the 20' bisection stimulus type are much easier to discriminate than the ones of the 30' bisection stimulus type. Learning does not occur because the reward cannot be estimated exactly for each type of bisection stimulus separately and, hence, the unsupervised bias is not zero. This is illustrated by simulations with a simple neural network. In Fig. 2C, performance improves when only one bisection stimulus is presented. The unsupervised bias fluctuates around zero. When two types of bisection stimuli are roved, the unsupervised bias is consistently different from zero for both stimulus types. No learning occurs (Fig. 2D and F). Details of the simulation are presented in Appendix A.

The  $R_{av}$  rule, based on Eq. (4) with a simple running average for  $\bar{R}$ , predicts that learning never occurs whatever stimulus types are roved because  $R_{av}$  cannot relate reward to different stimuli differently. However, humans normally can clearly distinguish the reward of various events. Indeed, roving hinders perceptual learning for certain stimulus combinations but not for others (Tartaglia, Aberg, & Herzog, 2009). For this reason, we propose that the human brain contains a critic which assigns reward signals to stimulus types separately. With such a critic, the unsupervised bias problem vanishes because for each single stimulus type  $\Delta\bar{R} = 0$  and hence learning can occur.

This model comes with a strong counterintuitive implication. Roving hinders perceptual learning with the 20' and 30' bisection stimulus types. However, these two bisection stimulus types are clearly discriminable from each other. For this reason, we propose that the critic has a much lower resolution and is therefore not located in the visual system but rather in the limbic and dopaminergic systems (Schultz, 2007, 2010; Schultz, Dayan, & Montague, 1997).

### 3. Discussion

It was shown that perceptual learning can occur when stimulus types are presented in an alternated fashion even though no learning occurs when the very same stimulus types are roved, i.e. randomly interleaved rather than alternated presentation (Kuai et al., 2005; Zhang et al., 2008). Also, pre-cueing the stimulus type can obliterate the deleterious effects of roving on perceptual learning (Zhang et al., 2008). We propose that in these situations, the critic learns to assign reward properly to the two stimulus types – in line with cognitive explanations on stimulus predictability (Kuai et al., 2005; Zhang et al., 2008). The predictable alternated order or the precue may act as a signal which enables learning of the critic. It should be mentioned that alternated presentations of stimulus types not always enable perceptual learning (Aberg & Herzog, 2009). Why and for which stimulus types this is the case has to be investigated in future research.

Roving hinders perceptual learning on the typical time scales where it otherwise occurs in non-roving conditions, i.e. in the range of 1000–3000 trials. However, roving can be overcome with about 18,000 trials (Parkosadze et al., 2008). Also in this case, we propose that the critic improves its resolution and learns to assign reward properly.

Perceptual learning under normal, non-roving conditions requires a minimal number of stimulus presentation per session

(Aberg, Tartaglia, & Herzog, 2009; Wright & Sabin, 2007; Wright et al., 2010). For example, three groups of observers trained with a chevron discrimination task with the same number of trials (1600 trials) split up in various sessions. One group trained with two sessions on consecutive days with 800 trials per session. Performance improved. The second group trained with four daily sessions with 400 trials each. Performance improved. The third group trained with 160 trials in 10 daily sessions – without success (Aberg, Tartaglia, & Herzog, 2009). The involvement of a critic is in accordance with these findings because a critic needs to average over a certain, statistically large sample of trials to estimate mean and variance well (Frémaux, Sprekeler, & Gerstner, 2010).

Perceptual learning shares a remarkable number of common characteristics with long-term potentiation (LTP). In particular, LTP needs a minimal stimulation (Bliss & Collingridge, 1993; Frey & Morris, 1997; Malenka & Bear, 2004) which mimicks the minimal number of trials needed for perceptual learning (Aberg & Herzog, submitted for publication). As previously shown, simple LTP based on Hebbian *pre- × post*-rules cannot explain learning. For this reason, a third factor related to dopamine (Schultz, 2010) is often added (Frémaux, Sprekeler, & Gerstner, 2010). We suggest that this factor is related to the critic and needed for all kinds of reward-based learning including perceptual learning.

### 4. Summary

Perceptual learning cannot be explained with purely supervised and unsupervised models. Perceptual learning is reward-based learning. There are two types of reward-based learning rules. The  $R_{max}$  rule is insensitive to roving and, hence, must be ruled out because roving *can* hinder human perceptual learning. The  $R_{av}$  rule is sensitive to roving. The  $R_{av}$  rule faces the unsupervised bias problem, i.e. reward must be exactly estimated. However, reward estimation is impossible for two stimulus types with different reward and, hence, no learning occurs. However, the  $R_{av}$  rule predicts that no learning occurs for all combinations of stimulus types which is clearly not true. The introduction of a critic, analyzing reward separately for each stimulus types, overcomes this problem. The spatial resolution of the critic must be, however, below perceptual resolution because roving blocks learning, for example, for bisection stimuli with outer line distances of 20' and 30'. Observers can clearly discriminate the two stimulus types. For this reason, we propose that the critic is, as in most reinforcement learning models, located in the limbic system but not in the visual system.

### Acknowledgments

This work is part of the sinergia project “State representations in reward-based learning” of the Swiss National Science Foundation (SNF).

### Appendix A

In the computer simulations, the two types of bisection stimuli, with larger or smaller outer line spacing, excite two different populations of neurons, each containing  $N = 50$  neurons. Presynaptic neurons respond to the random offset  $x$  of the center line with  $pre_i(x) = a_i + b_i x$ . The tuning parameters  $a_i$  and  $b_i$  are drawn from normal distributions (Fig. 2A),  $a_i \sim \mathcal{N}(2, 0.5)$  and  $b_i \sim \mathcal{N}(0, \beta)$ . Two postsynaptic neurons, coding for the decision (“left”/“right”), receive input from the presynaptic neurons, and noise  $\xi$  is added:  $post_j = \sum_i w_{ij} \times pre_i + \xi_j$ , with  $\xi_j \sim \mathcal{N}(0, 10)$ . A “winner-takes-all” mechanism suppresses the activity of the less active postsynaptic neuron,  $post_j \rightarrow 0, j = \operatorname{argmin}_j \{post_j\}$ , enforcing the decision. Task difficulty is modeled by changing the standard deviation  $\beta$  of the



randomly chosen  $b_i$ , with  $\beta = 0.25$  for the smaller outer bar spacing and 0.375 larger outer line spacing. The larger  $\beta$ , the stronger the effect of  $x$  on the postsynaptic neuron with respect to noise and the easier the learning.

Human subject's initial performance level is better than chance. To reproduce this effect in the model, we "pre-wired" the network with random starting synaptic weights biased in the right direction:  $w_{ij}^0 \sim \mathcal{U}(0, 1) + 2b_i \times (-1)^j$ , where  $\mathcal{U}$  is the uniform distribution.

The reward  $R$  is 1 if  $x > 0$  and "right" was chosen or  $x < 0$  and "left" was chosen, and  $-1$  otherwise. The mean reward at trial  $n$  is computed recursively:  $\bar{R}_n = [(\tau(n) - 1)\bar{R}_{n-1} + R_n] / \tau(n)$ , with  $\tau(n) := \min(\tau_R, n)$ , where  $\tau_R = 50$  is the reward averaging "time constant". We apply the learning rule  $R_{av}$  to the synaptic weights:  $\Delta w_{ij} = \eta \times pre_i \times post_j \times (R - \bar{R})$ , with  $\eta = 0.002$ . The weights are algorithmically constrained to a finite interval  $-10 \leq w_{ij} \leq 10$  to avoid runaway postsynaptic activity. Each simulation consisted of 14 blocks of 80 trials by distinct bisection spacing, randomly interleaved within each block. Fig. 2C–F shows the mean of 10 simulated subjects with errorbars representing the standard error of the mean.

## References

- Aberg, K. C., & Herzog, M. H. (2009). Interleaving bisection stimuli – randomly or in sequence – does not disrupt perceptual learning, it just makes it more difficult. *Vision Research*, 49, 2591–2598.
- Aberg, K., & Herzog, M. (2010). Does perceptual learning suffer from retrograde interference? *PLoS One*, 5, e14161.
- Aberg, K. C., & Herzog, M. H. (submitted for publication). About similar characteristics of visual perceptual learning and LTP.
- Aberg, K. C., Tartaglia, E. M., & Herzog, M. H. (2009). Perceptual learning with chevrons requires a minimal number of trials, transfers to untrained directions, but does not require sleep. *Vision Research*, 49, 2087–2094.
- Ball, K., & Sekuler, R. (1982). Specific and enduring improvement in visual motion discrimination. *Science*, 218, 697–698.
- Banai, K., Ortiz, J. A., Oppenheimer, J. D., & Wright, B. A. (2009). Learning two things at once: Differential constraints on the acquisition and consolidation of perceptual learning. *Neuroscience*, 165, 436–444.
- Bliss, T. V. P., & Collingridge, G. L. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, 361, 31–39.
- Crist, R. E., Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1997). Perceptual learning of spatial localization: Specificity for orientation, position, and context. *Journal of Neurophysiology*, 78, 2889–2894.
- Dosher, B. A., & Lu, Z. L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences of the United States of America*, 95(23), 13988–13993.
- Fahle, M., & Poggio, T. (Eds.). (2002). *Perceptual learning*. MIT Press.
- Farries, M. A., & Fairhall, A. L. (2007). Reinforcement learning with modulated spike timing dependent synaptic plasticity. *Journal of Neurophysiology*, 98(6), 3648–3665.
- Frémaux, N., Sprekeler, H., & Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. *Journal of Neuroscience*, 30(40), 13326.
- Frey, U., & Morris, R. G. M. (1997). Synaptic tagging and long-term potentiation. *Nature*, 385, 533–536.
- Herzog, M., & Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Research*, 37, 2133–2141.
- Herzog, M., & Fahle, M. (1998). Modelling perceptual learning: Difficulties and how they can be overcome. *Biological Cybernetics*, 78, 107–117.
- Izhikevich, E. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex*, 17, 2443–2452.
- Koyama, S., Harner, A., & Watanabe, T. (2004). Task-dependent changes of the psychophysical motion-tuning functions in the course of perceptual learning. *Perception*, 33, 1139–1147.
- Kuai, S.-G., Zhang, J.-Y., Klein, S., Levi, D. M., & Yu, C. (2005). The essential role of stimulus temporal patterning in enabling perceptual learning. *Nature Neuroscience*, 8, 1497–1499.
- Legenstein, R., Pecevski, D., & Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Computational Biology*, 4, e1000180.
- Loewenstein, Y. (2008). Robustness of learning that is based on covariance-driven synaptic plasticity. *PLoS Computational Biology*, 4(3), e1000007.
- Malenka, R. C., & Bear, M. F. (2004). LTP and LTD: An embarrassment of riches. *Neuron*, 44, 5–21.
- Otto, T. U., Herzog, M. H., Fahle, M., & Zhaoping, L. (2006). Perceptual learning with spatial uncertainties. *Vision Research*, 46, 3223–3233.
- Parkosadze, K., Otto, T. U., Malania, M., Kezeli, A., & Herzog, M. H. (2008). Perceptual learning of bisection stimuli under roving: Slow and largely specific. *Journal of Vision*, 8, 1–8.
- Petrov, A. A., Doshier, B. A., & Lu, Z. L. (2005). The dynamics of perceptual learning: An incremental reweighting model. *Psychological Review*, 112(4), 715–743.
- Poggio, T., Fahle, M., & Edelman, S. (1992). Fast learning in visual hyperacuity. *Science*, 256, 1018–1021.
- Sagi, D. (2011). Perceptual learning in Vision Research. *Vision Research*, 51(13), 1552–1566.
- Schultz, W. (2007). Behavioral dopamine signals. *Trends in Neurosciences*, 30, 203–210.
- Schultz, W. (2010). Dopamine signals for reward value and risk: Basic and recent data. *Behavioral and Brain Functions*, 6, 24.
- Schultz, W., Dayan, P., & Montague, R. R. (1997). A neural substrate for prediction and reward. *Science*, 275, 1593–1599.
- Shibata, K., Yamagishi, N., Ishii, S., & Kawato, M. (2009). Boosting perceptual learning by fake feedback. *Vision Research*, 49, 2574–2585.
- Sotiropoulos, G., Seitz, A. R., & Series, P. (2011). Perceptual learning in visual hyperacuity: A reweighting model. *Vision Research*, 51(6), 585–599.
- Sundareswaran, V., & Vaina, L. M. (1996). Adaptive computational models of fast learning of motion direction discrimination. *Biological Cybernetics*, 74, 319–329.
- Tartaglia, E., Aberg, K., & Herzog, M. (2009). Perceptual learning and roving: Stimulus types and overlapping neural populations. *Vision Research*, 49, 1420–1427.
- Wright, B. A., & Sabin, A. T. (2007). Perceptual learning: How much daily training is enough? *Experimental Brain Research*, 180, 727–736.
- Wright, B., Sabin, A., Zhang, Y., Marrone, N., & Fitzgerald, M. (2010). Enabling perceptual learning by alternating practice with sensory stimulation alone. *Journal of Neuroscience*, 30, 12868–12877.
- Yu, C., Klein, S. A., & Levi, D. M. (2004). Perceptual learning in contrast discrimination and the (minimal) role of context. *Journal of Vision*, 4, 169–182.
- Zhang, J.-Y., Kuai, S.-G., Xiao, L.-Q., Klein, S. A., Levi, D. M., & Yu, C. (2008). Stimulus coding rules for perceptual learning. *PLoS Biology*, 6, 1651–1660.