

Mechanisms of learning a visual motion discrimination task

Application to The Robert J. Glushko Dissertation Prize in Cognitive Science

Chi-Tat (Jeff) Law

Stanford University

A. INTRODUCTION

This précis summarizes my dissertation research in Professor Joshua Gold's laboratory at the University of Pennsylvania, which focused on understanding the long-standing question of how experience shapes our perceptual abilities. Our ability to make sense of what we see is central to how we perceive and interact with the world. This ability is regularly challenged, like when we drive, play video games or otherwise receive fleeting or ambiguous sensory input, but can improve with practice. Using a combination of psychophysics, electrophysiology and computational modeling techniques, I showed that improvements in perceptual ability in visual motion discrimination involved changes not in how the brain represents the sensory information but rather how it interprets the sensory representation to form the perceptual decision. Furthermore, a feedback-reinforcement signal appears to drive the changes in the decision process (§B). My results thus identified mechanisms central to perception, learning and decision-making and have potential significance to psychology, neuroscience, computer science and education (§C). My dissertation research was published in two Nature Neuroscience papers and has won the Saul Winegrad Award for Outstanding Dissertation. I hope my interdisciplinary approach in studying the neural mechanisms of visual perceptual learning align with the core principles of The Robert J. Glushko Dissertation Prize in Cognitive Science. I would like to thank the committee for considering my application.

B. Research Overview:

Training can cause substantial, long-lasting improvements in perceptual ability for both children and adults^{1,2}. This phenomenon, known as perceptual learning, implies a persistent capacity for plastic changes in the nervous system³. Even the earliest reports of perceptual learning, which involved improved acuity for tactile stimuli after only hours of practice, seemed to imply that the changes did not involve new receptors in the periphery but rather improved sensory processing within the central nervous system⁴. Thus, a long-standing goal of researchers in this field has been to identify where and how in the brain such changes occur⁵.

One prevalent hypothesis implicates primary sensory areas of cortex. Perceptual learning is often specific to the stimulus configuration used during training⁶. This specificity helps to distinguish perceptual learning from other forms of learning like adjustments of strategy that are more likely to generalize. Moreover, this specificity has been used to argue that the underlying changes likely occur in early sensory areas where the specificity of neuronal tuning is comparable to the specificity of learning⁷⁻¹⁰. Accordingly, changes in early sensory areas have been identified in monkeys trained on auditory¹¹ and tactile¹² discrimination tasks. However, in visual-discrimination learning, changes in early sensory areas are limited and in some cases insufficient to account for the behavioral improvements^{10,13-16}.

Changes in higher stages of processing, including those that contribute to decision-making and attention, have also been inferred from psychophysical studies^{13,17-21}. The resulting changes in top-down feedback are thought to be responsible for at least some of the training-induced changes found in primary visual cortex²¹⁻²³. Moreover, higher-order areas can also have narrow tuning and in principle can account for the stimulus specificity of perceptual learning¹⁹. However, changes in higher-order areas have not been identified directly in the brain, in part because for most perceptual tasks, little is known about the neural mechanisms underlying high-level processing, even in trained subjects.

I exploited recent advances in our understanding of the neural circuits underlying visual motion processing and perceptual decision-making in a random-dot motion

discrimination task (**Fig. 1a**) to examine *where* and *how* neural plasticity occurs during visual perceptual learning. In this task, neurons in the middle temporal area (MT) of extrastriate visual cortex are thought to provide the sensory evidence necessary for solving the task²⁴⁻²⁷. Neurons in the lateral intraparietal area (LIP) are thought to represent a decision process that reads out the sensory evidence from area MT to form a categorical judgment (leftward or rightward motion)²⁸⁻³¹ (**Fig. 1b**). I recorded the activity of individual MT and LIP neurons while two naïve monkeys were trained on the motion-discrimination task in order to identify experience-dependent changes in sensory representation and decision-making with training (**§B.1** and **§B.2**). In addition, I performed computational modeling to examine the signals that may drive these neural changes (**§B.3**). Below is a summary of the results from my **electrophysiology**, **psychophysics** and **computational modeling** studies:

B.1 Electrophysiology (Chapter 2 of Dissertation). I recorded the activity of individual neurons in areas MT and LIP of monkeys learning the motion discrimination task to identify experience-dependent changes in visual processing that accompanied

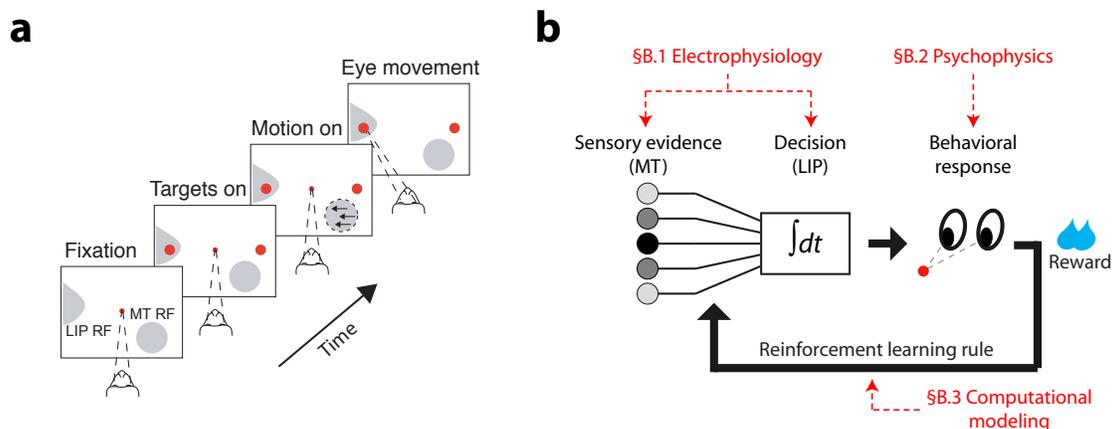


Figure 1. Behavioral task and neural substrates. **a**, Monkeys were trained to decide the direction of random-dot motion and indicate their decision with an eye movement. The motion direction (leftward or rightward), duration (0.1-1.4s) and strength (percentage of dots moving coherently to one direction; range from 0–100%) were varied randomly from trial-to-trial. **b**, Model of perceptual processing for the motion discrimination task. Area MT of the primate cortex provides the sensory information necessary for solving the task. Area LIP represents a decision process that reads out the motion information from MT to instruct the behavioral (eye movement) response. I recorded the activity of individual neurons in areas MT and LIP during training to identify neural changes in sensory and decision areas that corresponded to improvements in sensitivity to visual motion (**§B.1**). As a complementary study, I inferred changes in sensory readout from sensory to decision areas from the behavioral responses using the response-classification method (**§B.2**). Finally, I examined the computational principles that guide behavioral and neural changes with perceptual learning using modeling (**§B.3**).

their performance gains. I found that training had little effect on motion-driven responses in MT. Specifically, the ability of an ideal observer to determine the direction of motion based only on the responses of an individual MT neuron did not change systematically across sessions either before or during training (**Fig. 2a,b**, left panels). This result also likely implies no systematic changes in primary visual cortex, which is the primary, ascending source of motion selectivity in MT³².

In contrast, motion-driven responses of neurons in area LIP changed dramatically with training (**Fig. 2a,b**, middle panels). At the beginning of training, individual LIP neurons were selective for the monkey's eye movement response but were relatively insensitive to the motion stimulus. As training progressed and the monkeys' perceptual sensitivity improved (**Fig. 2a,b**, right panels), LIP neurons became increasingly sensitive to the strength and duration of the motion stimulus. Both the time course and magnitude of changes in LIP matched the behavioral improvements (Fig. 2.6 in Dissertation). Moreover, like behavioral performance, these changes in LIP tended to be somewhat specific to the stimulus configurations used during training, with slightly higher

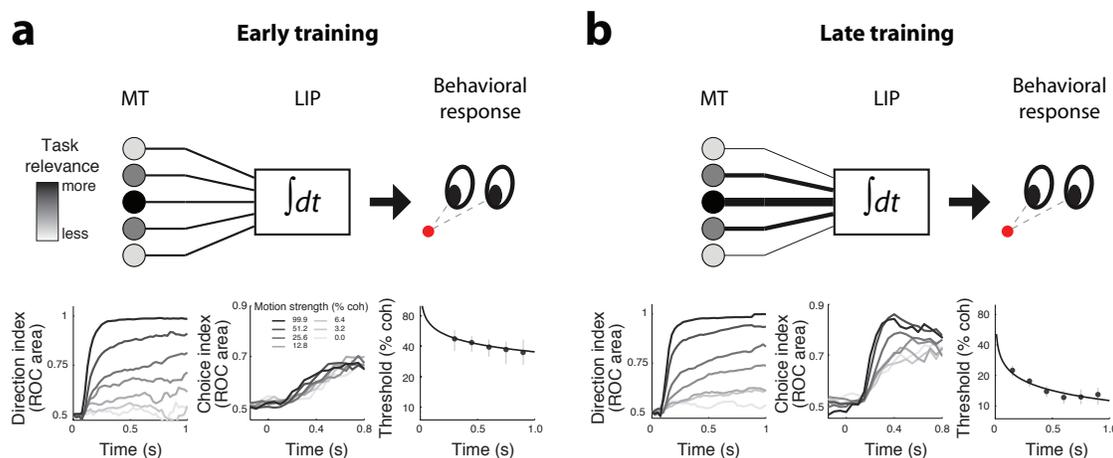


Figure 2. Training-induced plasticity in decision (area LIP) but not sensory (area MT) neurons. **a,b**, Training had little effect on motion-driven responses in MT but a dramatic effect on both the coherence- (i.e., motion strength) and time-dependent responses in LIP and behavioral performance. The lower panels show summary data from a single monkey for the first (panel **a**) and last (panel **b**) one-third of training sessions. Population MT data are shown as an ideal-observer (ROC) index of the ability to prediction motion direction from the neural responses plotted as a function of viewing time for different motion strengths. Population LIP data are shown as an ideal-observer (ROC) index of the ability to prediction the monkey's choices from the neural responses plotted as a function of viewing time for different motion strengths. Behavioral data are shown as discrimination threshold as a function of viewing time. These data are consistent with an increasingly selective readout (depicted as thicker lines connecting MT to LIP) of the most informative MT neurons (depicted as darker shades) to form the decision that guides behavior.

sensitivity for motion directions used in previous training sessions (Fig. 2.7 in Dissertation). Thus, for this task, perceptual learning appears to involve improvements not in how sensory information is represented in the brain but rather how the sensory representation is interpreted to form the decision that guides behavior. Results from this study were **published in the April 2008 issue of Nature Neuroscience**³³.

B.2 Psychophysics (Chapter 3 of Dissertation). I used a response-classification method to measure correlations between trial-by-trial, random fluctuations of stimulus motion directions and the subject's choices to characterize the directions of motion that exerted the greatest influence on behavior³⁴⁻³⁷. I computed the motion strength (called motion energy) along different direction axes using a set of motion filters for trials using ambiguous stimuli (0% coherence) and examined how random fluctuations of stimulus motion directions would influence the monkey's subsequent decisions to choose the leftward (**Fig. 3**, blue symbols) and rightward (red symbols) targets.

Training increased the influence of motion signals near the stimulus's directions on behavioral choice. Early in training (**Fig. 3**, left panel), the average motion energy was similar for leftward and rightward decisions. In contrast, late in training (right panel), the average motion energy was more positive for rightward choices and more negative for leftward choices for stimulus directions around the 0° direction axis. One explanation for this result is that the subject learned to use task-relevant motion signals around the 0° direction axis more selectively with training. Therefore, stimuli with slightly

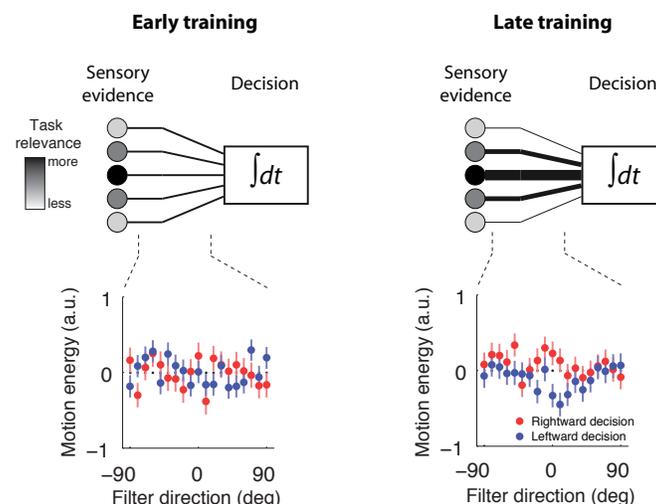


Figure 3. Changes in sensory readout inferred from the response-classification method. **(Top)** Schematics depicting changes in readout from sensory to decision areas with training. **(Bottom)** Average motion energy as a function of direction tuning of the motion energy filter (-90 to 90° in 10° intervals) for rightward (red symbols) and leftward (blue symbols) decisions early (left) and late (right) in training. Positive and negative values of motion energy represent net rightward and leftward motion signals, respectively. Error bars are SEM across 0% coherence trials.

more positive motion energy near 0° tended to produce rightward decisions, and stimuli with slightly more negative motion energy around 0° tended to produce leftward decisions. Together, results from my electrophysiology and psychophysics experiments suggested that improvements in perceptual sensitivity in the motion discrimination task corresponded to an increasingly selective read-out of highly sensitive MT neurons (i.e., neurons tuned near the stimulus directions and/or with higher sensitivity; **Fig. 3** and Fig. 2.3c in Dissertation) by a decision process, represented in LIP that instructed the behavioral response.

An implication from these results is a close relationship between mechanisms of perceptual and associative learning³⁸. Associations can be formed by establishing functional connectivity from sensory neurons representing a particular visual feature to decision neurons that generate a behavior based on that feature. The results presented in **§B.1** and **§B.2** suggest that perceptual learning can involve refining this connectivity so that the readout scheme is more appropriate for the particular task demands, allowing the decision process to more effectively distinguish between the alternatives³⁹⁻⁴³.

This link between associative and perceptual learning suggests that they might also share signals to guide the changes in neural connectivity. One potential signal is based on feedback reinforcement⁴⁴, which has extensively documented relationships to the dopaminergic system in the brainstem⁴⁵. Models of plasticity based on feedback reinforcement signals are consistent with plausible synaptic mechanisms and can account for many forms of associative learning⁴⁶⁻⁴⁸. In addition, feedback reinforcement is thought to play a role in some forms of perceptual learning^{49,50}. Therefore, I examined whether a reinforcement-learning rule based on a reward-prediction error could account for both associative and perceptual learning on the motion discrimination task⁵¹.

B.3 Computational modeling (Chapter 4 of Dissertation). I modeled performance based on the readout of simulated responses of direction-selective neurons in MT. A reward prediction error guided changes in connections between these sensory neurons and the decision process (**Fig. 1b** and Fig. 4.1 in Dissertation), first help to establish functional connections from the population of MT-like sensory neurons to a population of LIP-like decision neurons that interpret the sensory information to determine the saccadic

response. The same mechanisms then further refine these connections to weigh more strongly inputs from the most informative sensory neurons, thereby improving perceptual sensitivity (Fig. 4.3 in Dissertation). I simulated trial-by-trial performance of the model using the exact sequences of stimulus conditions (i.e. motion directions, coherences, and viewing durations for each monkey) experienced by the monkeys and compared the simulated results to the real behavioral, MT and LIP data.

The model can account for both associative and perceptual learning of the monkeys (**Fig. 4**). I computed the lapse rate (gray symbols, a measure of knowledge of task rules including the sensory-motor association) and discrimination threshold (black symbols, a measure of perceptual sensitivity) from the simulated behavioral data and compared that with the real data of the monkeys. The model can reproduce the time courses and asymptotic values of both the associative (lapse rate) and perceptual (discrimination threshold) learning of the monkeys. Like for the monkey data, lapse rates declined rapidly to near zero. Discrimination threshold improved more gradually, eventually reaching lower asymptotes comparable to those reached by the monkeys. These results were robust to a variety of parameters, pooling schemes and reinforcement rules used in the model (Figs. 4.2c, S4.1, S4.2 and Table 4.1 in Dissertation). In addition, the model can also account for neural changes in MT related to the readout of the sensory representation (Figs. 2.3c and 4.4 in Dissertation), and progression of motion-sensitive responses in neurons that form the decision in LIP (**Fig. 5**; Fig. 4.5 in Dissertation). Thus, the results suggest that reinforcement learning might play a general role in both establishing and shaping patterns of connectivity critical for forming perceptual

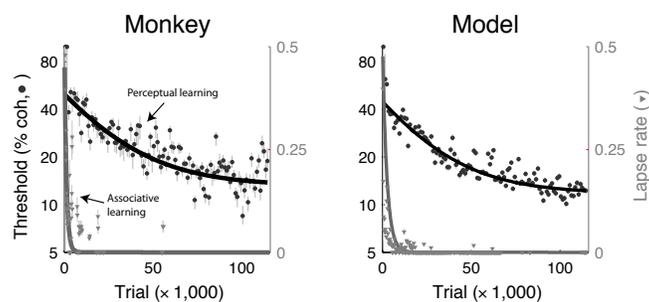


Figure 4. Discrimination performance of the monkey (left) and model (right) with training. Discrimination threshold (●; logarithmic scale on the left ordinate) and lapse rate (▼; error rate at 99.9% coherence, linear scale on the right ordinate) with 68% CIs plotted as a function of trial number (computed using 1000-trial bins for discrimination threshold, and 250-trial bins for lapse rate) computed from psychophysical data of monkey C during training. Solid lines are best fits of a decaying single-exponential function.

decisions. Results from this study were **published in the May 2009 issue of Nature Neuroscience**⁵¹.

To summarize, my dissertation combined the use of electrophysiology, psychophysics and computational modeling to examine where and how neural plasticity that underlies perceptual improvements occurs. My results provide the clearest and most complete description of the neural changes that accompany visual perceptual learning published to date. I identified changes at the level of decision, but not sensory, circuits, which likely explains the limited changes found previously in studies that focused on sensory representations^{10,14}. In addition, I showed that a reinforcement learning rule that shapes functional connectivity between sensory and decision neurons is sufficient to account for the behavioral and neural improvements with training. These results have led to a novel hypothesis that perception learning and associative learning might share a common, reinforcement-driven mechanism (an idea reviewed in Ref. 43).

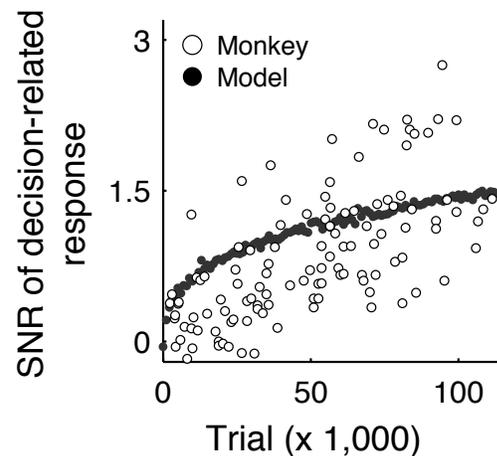


Figure 5. Changes in decision-related responses with training. The signal-to-noise ratio (i.e., the difference in mean responses to the two directions of motion divided by their common standard deviation; SNR) of the simulated decision neurons (black symbols; 51.2% coherence at 0.4-s viewing time; 1,000-trial bins) as a function of training. Open symbols are the SNR of LIP responses for Monkey C at the same motion coherence and viewing time for motion into versus out of the neuron's response field. See Fig. 4.5 in Dissertation for comparison at other stimulus conditions.

C. Potential significance to cognitive science research:

C.1 Significance to psychology. My dissertation established a framework to relate three areas of psychology research – associative, perceptual and reinforcement learning. Associative learning is a process to map situations to actions. Perceptual learning refers to the long-lasting improvements in our ability to interpret and respond to sensory signals in the environment. My results suggested that under certain circumstances, both forms of learning can be driven by the same feedback-reinforcement mechanism^{38,43}. Although

further works is need established the validity of this scheme, this research direction would potentially reveal common neural substrates for major forms of non-declarative memories in the brain.

C.2 Significance to neuroscience. Changes in functional connectivity within neural circuits are thought to be responsible for many forms of learning and memory⁵². However, there remains a substantial divide between studies of molecular and cellular mechanisms of learning and studies that characterize learning behavior, in large part because of a lack of model systems for studying mechanisms of plasticity in a functional context. My dissertation showed for the first time that changes in functional connections between sensory and decision areas might be the primary site of plasticity that accompanies visual perceptual learning. Thus, my results have established a model system for studying the molecular and cellular mechanisms underlying this neural plasticity during learning behaviors. The results would likely reveal broadly applicable principles of plasticity in neural circuits and might lead to important clinical applications⁵³.

C.3 Significance to computer science. My modeling study compared real and simulated behavioral and neural data using a wide range of model parameters, pooling schemes and reinforcement rules, thus establishing constraints on the learning algorithms that the brain might use and circumstances under which perceptual learning could be generalized to other behavioral contexts (Fig. 4.7 in Dissertation). Both the principles and generalization of learning are important issues in computer science. Therefore, my results would potentially inspire new design principles for machine learning and machine perception^{54,55}.

C.4 Significance to education. A key challenge in science and mathematics education is to design effective methods to convey abstract concepts and relationships to the students. Interestingly, recent evidence suggested that perceptual learning might facilitate this kind of intellectual reasoning^{56,57}. For example, learning to recognize symbolic patterns in mathematical equations (without solving them) improved algebraic proficiency. An explanation for these results is that perceptual recognition and cognitive reasoning both shared common mechanisms that select and interpret the relevant sensory

signals. My research identified a neural substrate for improvements in sensory selection in the parietal cortex⁵⁸⁻⁶¹. Moreover, I showed that these improvements could be driven by a feedback reinforcement-learning rule. These biological principles could be used to design better teaching methods for science and mathematics education.

D. Conclusion:

Training can induce long-lasting changes in our ability to detect, discriminate or identify a sensory stimulus, even in adult. My research aimed to understand the neural mechanisms of this form of perceptual learning. Using an interdisciplinary approach, I identified novel changes at the level of decision-making, but not sensory, circuits that accompanied improvements in visual discrimination. My research has potential significance to multiple fields in cognitive science, including psychology, neuroscience, computer science and education. I believe this work closely aligns the core belief of The Robert J. Glushko Dissertation Prizes in Cognitive Science that “*understanding how minds work requires the synthesis of many different empirical methods, formal tools and analytic theories*” to “*bridge between the areas of cognitive science and create theories of general interest to the multiple fields concerned with the nature of minds and intelligent systems*”. I would like to thank the committee again for considering my application.

E. References:

- 1 Goldstone. *Annu Rev Psychol* **49**, 585-612, (1998).
- 2 Gibson. *Annu Rev Psychol* **14**, 29-56, (1963).
- 3 Sur et al. *Curr Biol* **12**, R168-170, (2002).
- 4 Volkman. *Math.-phys. Classe* **10**, 38-69, (1858).
- 5 Gilbert et al. *Neuron* **31**, 681-697, (2001).
- 6 Fahle. *Curr Opin Neurobiol* **15**, 154-160, (2005).
- 7 Karni & Sagi. *Proc Natl Acad Sci U S A* **88**, 4966-4970, (1991).
- 8 Berardi & Fiorentini. *Journal of Physiology* **Vol. 384**, 633-647, (1987).
- 9 Poggio et al. *Science* **256**, 1018-1021, (1992).
- 10 Schoups et al. *Nature* **412**, 549-553, (2001).
- 11 Recanzone et al. *J Neurosci* **13**, 87-103, (1993).
- 12 Recanzone et al. *J Neurophysiol* **67**, 1031-1056, (1992).
- 13 Crist et al. *Nat Neurosci* **4**, 519-525, (2001).

- 14 Ghose *et al.* *J Neurophysiol* **87**, 1867-1888, (2002).
15 Yang & Maunsell. *J Neurosci* **24**, 1617-1626, (2004).
16 Raiguel *et al.* *J Neurosci* **26**, 6589-6602, (2006).
17 Ahissar & Hochstein. *Proc Natl Acad Sci U S A* **90**, 5718-5722, (1993).
18 Saarinen & Levi. *Vision Res* **35**, 519-527, (1995).
19 Mollon & Danilova. *Spat Vis* **10**, 51-58, (1996).
20 Doshier & Lu. *Vision Res* **39**, 3197-3221, (1999).
21 Li *et al.* *Nat Neurosci* **7**, 651-657, (2004).
22 Sigman *et al.* *Neuron* **46**, 823-835, (2005).
23 Ito *et al.* *Neuron* **20**, 1191-1197, (1998).
24 Newsome & Pare. *J Neurosci* **8**, 2201-2211, (1988).
25 Britten *et al.* *Vis Neurosci* **13**, 87-100, (1996).
26 Salzman *et al.* *Nature* **346**, 174-177, (1990).
27 Britten *et al.* *J Neurosci* **12**, 4745-4765, (1992).
28 Shadlen & Newsome. *J Neurophysiol* **86**, 1916-1936, (2001).
29 Roitman & Shadlen. *J Neurosci* **22**, 9475-9489, (2002).
30 Hanks *et al.* *Nat Neurosci* **9**, 682-689, (2006).
31 Huk & Shadlen. *J Neurosci* **25**, 10420-10436, (2005).
32 Movshon & Newsome. *J Neurosci* **16**, 7733-7741, (1996).
33 Law & Gold. *Nat Neurosci* **11**, 505-513, (2008).
34 Gold *et al.* *Nature* **402**, 176-178, (1999).
35 Gold *et al.* *Current Biology* **10**, 663-666, (2000).
36 Gold *et al.* *Cognitive Science* **28**, 167-207, (2004).
37 Li *et al.* *Nat Neurosci* **7**, 178-183, (2004).
38 Hall. (Clarendon Press; Oxford University Press, 1991).
39 Pouget *et al.* *Annu Rev Neurosci* **26**, 381-410, (2003).
40 Geisler & Albrecht. *Vis Neurosci* **14**, 897-919, (1997).
41 Seung & Sompolinsky. *Proc Natl Acad Sci U S A* **90**, 10749-10753, (1993).
42 Green *et al.* *Curr Biol* **20**, 1573-1579, (2010).
43 Law & Gold. *Topics in Cognitive Science* **2**, 226-238, (2010).
44 Sutton & Barto. (MIT Press, 1998).
45 Schultz. *Annu Rev Neurosci* **30**, 259-288, (2007).
46 Seung. *Neuron* **40**, 1063-1073, (2003).
47 Dayan & Balleine. *Neuron* **36**, 285-298, (2002).
48 Fusi *et al.* *Neuron* **54**, 319-333, (2007).
49 Seitz & Dinse. *Curr Opin Neurobiol* **17**, 148-153, (2007).
50 Bao *et al.* *Nature* **412**, 79-83, (2001).
51 Law & Gold. *Nat Neurosci* **12**, 655-663, (2009).
52 Lewis *et al.* *Proc Natl Acad Sci U S A* **106**, 17558-17563, (2009).
53 Levi. *Dev Psychobiol* **46**, 222-232, (2005).
54 Bredeche *et al.* *Proceedings the Second IEEE International Conference on Cognitive Informatics*, 18-25|x+227, (2003).
55 Taylor. *Acad Radiol* **14**, 1252-1263, (2007).
56 Kellman *et al.* *Topics in Cognitive Science* **2**, 285-305, (2010).
57 Goldstone *et al.* *Topics in Cognitive Science* **2**, 265-284, (2010).
58 Colby & Goldberg. *Annu Rev Neurosci* **22**, 319-349, (1999).
59 Platt & Glimcher. *Nature* **400**, 233-238, (1999).
60 Snyder *et al.* *Vision Res* **40**, 1433-1441, (2000).
61 Gold & Shadlen. *Annu Rev Neurosci* **30**, 535-574, (2007).