Neuron, Vol. 39, 1031–1042, September 11, 2003, Copyright ©2003 by Cell Press

Active Reversal of Motor Memories Reveals Rules Governing Memory Encoding

Edward S. Boyden and Jennifer L. Raymond*
Department of Neurobiology
Stanford University
Stanford, California 94305

Summary

Learning systems must be able to store memories reliably, yet be able to modify them when new learning is required. At the mechanistic level, new learning may either reverse the cellular events mediating the storage of old memories or mask the old memories with additional cellular changes that preserve the old cellular events in a latent form. Behavioral evidence about whether reversal or masking occurs in a particular circuit can constrain the cellular mechanisms used to store memories. Here we examine these constraints for a simple cerebellum-dependent learning task, motor learning in the vestibulo-ocular reflex (VOR). Learning can change the amplitude of the VOR in two opposite directions. Contrary to previous models about memory encoding by the cerebellum, our results indicate that these behavioral changes are implemented by different plasticity mechanisms, which reverse each other with unequal efficacy.

Introduction

Learning systems must, by their very nature, be capable of encoding different content throughout life. As circumstances change and new memories are learned, expression of old memories must in some cases be suppressed to prevent interference. There are two general strategies for suppressing old memories: the cellular events encoding them may be reversed by the mechanisms supporting the new memories, or their expression may be masked by cellular events elsewhere, thus preserving the old memories in a latent form. These two strategies endow learning systems with different capacities. The first strategy, reversal, could free up resources for further memory encoding. This strategy has been suggested for areas such as the hippocampus, which has been modeled as a "scratchpad" for temporary storage of memories (for review, see Squire and Alvarez, 1995). The second strategy, masking, could allow an animal to relearn an old behavior faster, since part of the information is quiescently stored in the brain. This latter strategy seems to play an important role in extinction of fear conditioning, where the expression of fear memory may be inhibited by additional changes in the circuit (for review, see Myers and Davis, 2002).

Experimental determination of which strategy is used for a given behavioral task can provide insight regarding the cellular mechanisms used to store memories in a circuit over time, especially if the anatomy is well understood and the plasticity mechanisms in the circuit are well characterized. Of particular relevance are the reversal properties of plasticity mechanisms found in the circuit. In vitro studies are now characterizing these properties for a variety of plasticity mechanisms, including long-term potentiation (LTP) and long-term depression (LTD), as well as changes in intrinsic excitability (Caria et al., 2001; Fujii et al., 1991; Han et al., 2000; Montgomery and Madison, 2002; Smith et al., 2002; Staubli and Chun, 1996). These experiments are revealing a number of contingencies under which reversal of plasticity can occur, such as limited temporal windows for reversal of LTP. Ultimately the strategies for reversal of old memories must be understandable in terms of such constraints, when operating in the context of a functional circuit.

We used a simple cerebellum-dependent behavior to explore the interaction between old and new memories, at a level that permits us to deduce rules governing how their underlying cellular mechanisms must interact. The vestibulo-ocular reflex (VOR) provides an ideal arena for analyzing how new learning and old memories interact. During the VOR, head motion in one direction elicits eye motion in the opposite direction so as to stabilize images on the retina during head movement. The amplitude, or gain, of this reflex can be adaptively increased or decreased by exposing an animal to particular combinations of visual and vestibular stimuli (Figures 1A and 1B; Gonshor and Jones, 1973; Ito et al., 1974; Miles and Fuller, 1974). By concatenating stimuli that increase and decrease the gain of the VOR, we analyzed whether reversal or masking of old memories occurred during these oppositely directed behavioral changes. We also analyzed how the degree of masking versus reversal depends on the quantity and temporal distribution of training to be suppressed. Our results suggest that (1) memories for increases and decreases in VOR gain rely on different plasticity mechanisms, and (2) these plasticity mechanisms reverse each other with unequal efficacy. Our findings challenge a longstanding model of cerebellum-dependent learning and reveal a complex interaction between old and new cerebellum-dependent motor memories.

Results

Acute Training Induces Persistent Changes in the VOR

We induced motor learning in the VOR of mice by pairing head rotation with rotation of an optokinetic drum. Moving the drum in the opposite direction from the head (gain-up stimulus; Figure 1A) caused the gain of the VOR to adaptively increase, whereas moving the drum in the same direction as the head (gain-down stimulus; Figure 1B) caused the gain of the VOR to decrease. We chose the speed of the drum so that the magnitudes of the increase and decrease in VOR gain would be similar (see Experimental Procedures). Eye velocity traces, recorded during head rotation in the dark before and after training, illustrate the learned changes in VOR gain induced by

*Correspondence: jenr@stanford.edu

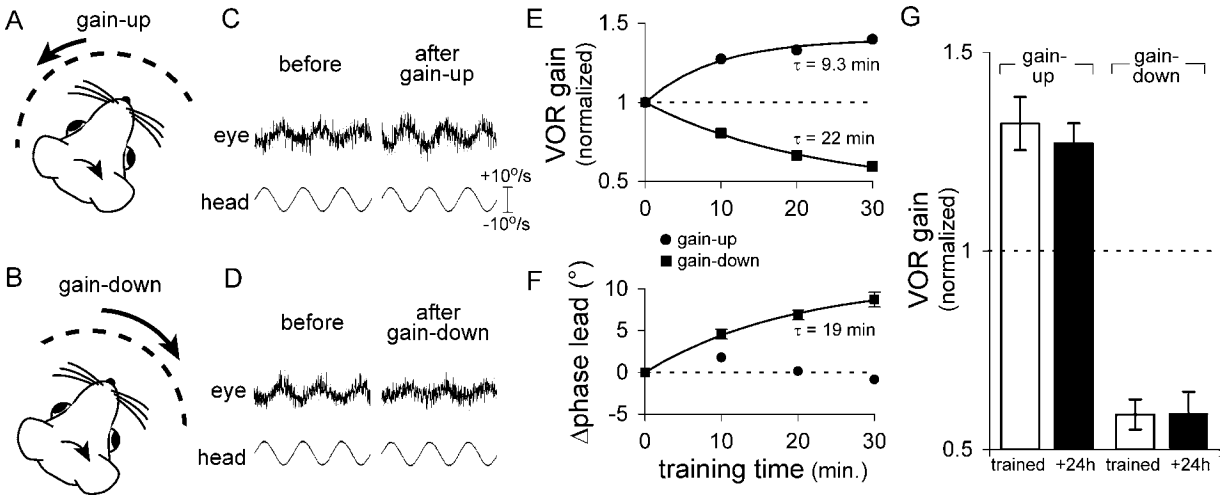

Figure 1. Persistent Changes in VOR Induced by Gain-Up and Gain-Down Stimuli

(A) Gain-up stimulus. Optokinetic drum (dashed circle) and head move sinusoidally in opposite directions, with peak drum velocity equal to 50% of peak head velocity.

(B) Gain-down stimulus. The drum moves with the head.

(C) Representative eye (top) and head (bottom) velocity traces recorded in darkness before and after 30 min of gain-up training.

(D) Representative eye and head velocity traces as in (C), but for a gain-down training session.

(E) Changes in VOR gain, induced by 30 min of gain-up (circles; n = 30) or gain-down (squares; n = 30) training. VOR was measured in the dark every 10 min, during a 30 min training session. In all figures, data points are plotted as mean ± SEM, and n refers to the number of mice (each used only once). Throughout the figures, VOR gain is normalized by the initial VOR gain, and error bars are omitted whenever they are smaller than the symbols. τ, time constant for learning.

(F) Changes in VOR phase, induced by 30 min of gain-up and gain-down stimuli. Positive values represent an increase in phase lead. Symbols as in (E).

(G) VOR gain measured at the end of a 30 min gain-up or gain-down training session (trained, white bars) and after an additional 24 hr in complete darkness (+24h, dark bars) (n = 9).

gain-up and gain-down stimuli (Figures 1C and 1D). The gain of the VOR, defined as the ratio of eye to head velocity, increased with 30 min of gain-up training from an initial value of 0.40 ± 0.03 (mean ± SEM) to a final value of 0.56 ± 0.03, a change of +40% (n = 30, Figure 1E). With 30 min of gain-down training, the gain of the VOR decreased from 0.42 ± 0.02 to a final value of 0.25 ± 0.01, a change of −40% (n = 30). In each case, the time course of learning was well fit by a single exponential with a rapid time constant (9.3 min for gain-up, 22 min for gain-down).

The dynamics of the VOR were affected differently by gain-down and gain-up training. The phase of the VOR is a measure of the timing of peak eye velocity relative to peak head velocity (see Experimental Procedures). Naive animals had a VOR phase lead of 22.9° ± 0.5° ahead of perfect compensation, meaning that peak eye velocity occurred earlier than peak head velocity (n = 60; mean ± SEM). Training with gain-up stimuli for 30 min caused no significant change in the phase of the VOR ($p > 0.05$, Wilcoxon signed-rank test [WSRT]; Figure 1F). In contrast, gain-down training increased the phase lead of the VOR significantly ($p < 0.05$). This increase in phase lead proceeded with a time constant similar to that of the decrease in gain.

Previously, it was shown that changes in the VOR induced by extended training persisted, in the absence of visuovestibular stimuli, for days after induction (Miles and Eighmy, 1980; Robinson, 1976). However, persistence of learning has not previously been characterized after acute training. In order to measure the persistence of learned changes in the VOR, we remeasured the VOR 24 hr after a 30 min training period. During this 24 hr period, the mice were allowed to roam freely in their cages, which were kept in a completely dark chamber to prevent further adaptive changes in VOR gain. As shown in Figure 1G, the changes in VOR gain induced by gain-up or gain-down training did not decay during 24 hr in darkness ($p > 0.05$, WSRT; n = 9 each training condition). Thus, the rapidly induced changes produced by gain-up or gain-down training are persistent motor memories.

Saturation of the Increase in VOR Gain

We next determined whether motor learning in the VOR could be saturated with acute training. Mice were exposed to three 30 min training sessions with either gain-up or gain-down stimuli, with 2 hr rest periods in darkness between sessions. These rests allowed us to measure the persistence of changes induced during each training session. When mice were trained with gain-up stimuli, each additional training session increased the VOR gain, but the additional increases induced by the second and third training sessions did not persist through a 2 hr rest (Figure 2A, filled symbols). During each additional training session, the VOR gain increased significantly ($p < 0.05$ for both, WSRT; n = 14 for second session, 9 for third session) but decayed during the 2 hr rest to the value at the beginning of the previous session ($p > 0.05$ for both). Thus, although the original gain-up training session resulted in a persistent increase in VOR gain, additional gain-up training sessions resulted only in

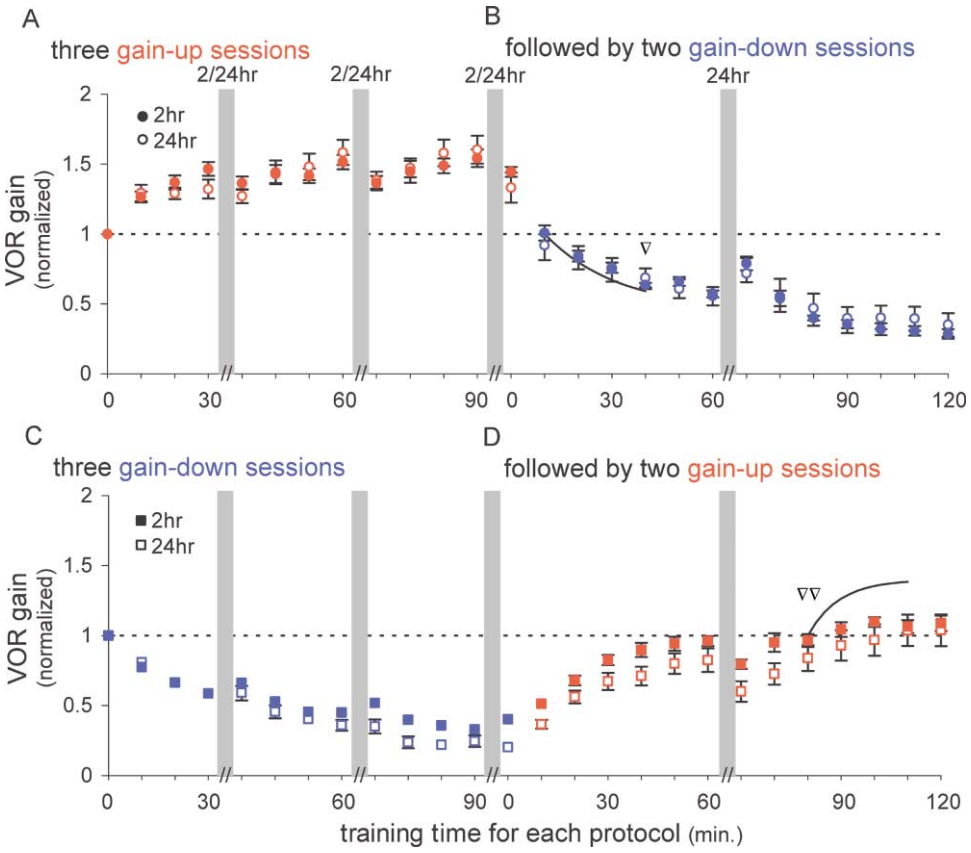

Figure 2. Changes in VOR Gain Induced when Gain-Up and Gain-Down Training Sessions Are Concatenated

(A) Increases in VOR gain induced by three 30 min gain-up training sessions (red symbols). Between training sessions, there were either 2 hr (filled circles) or 24 hr (open circles) rest periods in darkness, indicated by shaded bars (n = 14, n = 9, respectively). The *x* axis indicates the cumulative time of gain-up training experienced.

(B) Oppositely directed changes in VOR gain induced by gain-down training (blue symbols) following gain-up training (n = 10, n = 9, respectively, through 1 hr of gain-down training; n = 5, n = 5, respectively, through 2 hr). The *x* axis time is reset to zero when the gain-down protocol begins. The solid curve starting at the 10 min point in gain-down training is the exponential fit to gain-down training from the naive state (Figure 1E) for comparison to the time courses shown here. ∇ marks the time point used for comparison to naive mice undergoing gain-down in the text.

(C) Decreases in VOR gain induced by three 30 min gain-down training sessions (blue symbols), with either 2 hr (filled squares) or 24 hr (open squares) rest periods in between training sessions (n = 15, n = 9, respectively).

(D) Oppositely directed changes in VOR gain induced by gain-up training (red symbols) following gain-down training with either 2 hr (filled squares) or 24 hr (open squares) rest periods (n = 14, n = 9, respectively, through 1 hr of gain-up training; n = 12, n = 6, respectively through 2 hr). The solid curve beginning at the 80 min point in gain-up training is the exponential fit to gain-up training from the naive state (Figure 1E). ∇∇ marks the time point at which the VOR is restored to its original gain, for mice pretrained with gain-down sessions with 2 hr rests.

transient increases. This saturation of the persistent increase in VOR gain occurred even though the gain required for image stabilization had not been achieved. Retinal image slip, defined as the movement of the visual image relative to the eye, is believed to be an error signal controlling motor learning in the VOR. The velocity of retinal image slip during exposure to the gain-up stimulus decreased slightly during the three training sessions but was not reduced dramatically (Figure 3A).

Decreases in VOR gain did not saturate like the increases in gain did (Figure 2C, filled symbols). When mice were trained with three gain-down sessions, the gain returned toward the initial gain during the 2 hr rest following each session but remained lower than the gain at the beginning of that training session ($p < 0.05$; n = 15 for second session, 9 for third session). Thus, additional gain-down training sessions caused lasting decreases in VOR gain, whereas additional gain-up training after the first 30 min training session caused only transient increases in VOR gain.

The changes in VOR phase after multiple training sessions paralleled the changes in VOR gain. During three gain-up sessions, only a small (but significant, $p < 0.05$) reduction in phase lead accumulated (Figure 4A, filled symbols). However, during three gain-down training sessions, the phase lead increased steadily (Figure 4C, filled symbols). Thus, additional training sessions changed the VOR phase in ways that continued the trends observed during the first 30 min of training (Figure 1F).

Reversibility of Increases and Decreases in VOR Gain

In order to examine how the mechanisms for increasing and decreasing the gain of the VOR interact, we exposed

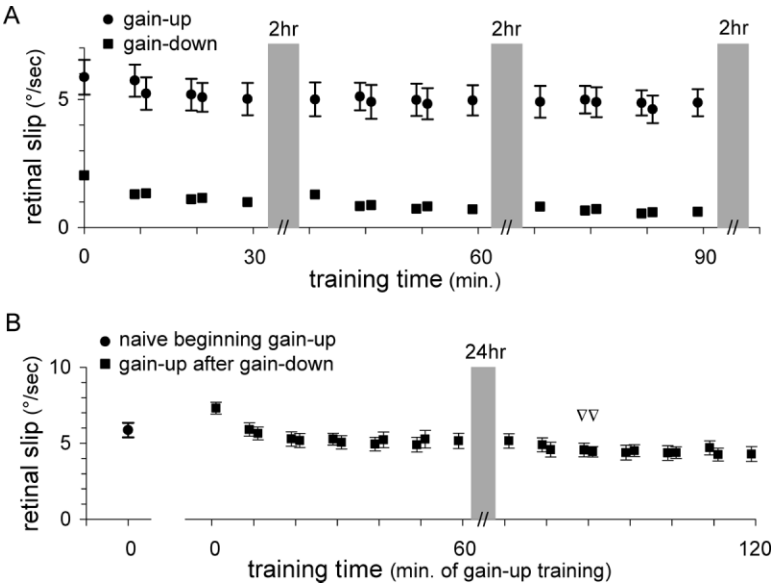

Figure 3. Retinal Image Slip during Training

(A) Retinal image slip velocity, measured during gain-up (circles) and gain-down (squares) training sessions. Measurements were made at the beginning and end of every 10 min training block.

(B) Retinal image slip velocity during gain-up training, for both naive mice (filled circle; $n = 30$) and mice pretrained with gain-down stimuli (squares; $n = 12$). $\nabla\nabla$ marks the time point when the VOR is restored to its original gain.

mice to sequential sessions of gain-up and gain-down training. Previous experiments in primates have shown that several days of normal experience can reverse motor learning in the VOR induced by several weeks of wearing magnifying or miniaturizing spectacles (Miles and Eighmy, 1980). However, it has not been determined whether acute protocols that rapidly increase or decrease VOR gain can reverse each other over short timescales. To examine this, we trained mice with three 30 min gain-up or gain-down training sessions, separated by 2 hr rests. Immediately following the last rest, we trained the same mice with two 1 hr sessions using the oppositely directed training stimulus, separated by a 24 hr rest.

When mice were initially trained with three gain-up training sessions (Figure 2A), the two subsequent gain-down sessions decreased the VOR gain (Figure 2B) to values that were comparable to those induced by gain-down training in naive mice (compare blue symbols in Figures 2B and 2C). The decrease induced by gain-down training (Figure 2B) could be divided into two components with different kinetics. During the first 10 min of gain-down training, the elevated VOR gain decreased back to its initial value. This fast reversal of the effects of gain-up training was followed by a slow decrease in VOR gain, which was comparable to the one induced by gain-down training in naive mice. After 30 min of this slow decrease (Figure 2B, ∇), the gain of mice pretrained with gain-up sessions was indistinguishable from that of naive mice after 30 min of gain-down training ($p > 0.05$, Mann-Whitney U test [MWUT]; $n = 10$). This indicates that mice had not only reattained their initial VOR gain within 10 min but were capable of proceeding with normal learning in response to the gain-down stimulus.

The degree of reversibility observed was quite different when mice were trained in the opposite order. When gain-down training (Figure 2C) was followed by gain-up training (Figure 2D), the VOR gain slowly increased during the two 1 hr gain-up sessions, reaching values much lower than those reached by naive mice undergoing gain-up training (compare red symbols in Figures 2D and 2A). When mice were initially trained with three gain-down training sessions, it took 80 min of gain-up training for the VOR gain to return to its initial value ($n = 12$; Figure 2D, $\nabla\nabla$). During the next 40 min, the VOR gain reached a value significantly less than that reached by naive mice with just 30 min of gain-up training (Figure 1E; $p < 0.05$, MWUT). This indicated that although the mice had reattained their initial VOR gain after 80 min of gain-up training, they were not capable of normal learning in response to the gain-up stimulus. Residual effects of gain-down training must have persisted in the VOR circuit and limited the capacity for increasing the gain, even after the VOR gain was restored to its initial value.

One factor that could limit the reversal of learning in the VOR would be a deficit in the sensorimotor variables that provide the error signals guiding learning. In particular, if pretraining with gain-down stimuli affected retinal slip during subsequent gain-up training, this could explain the impairment of learning relative to naive mice. However, when mice pretrained with gain-down sessions had reattained their initial VOR gain, the retinal slip they experienced during gain-up training was not significantly different from that of naive mice beginning gain-up training, either in amplitude or phase ($p > 0.05$, MWUT; Figure 3B, $\nabla\nabla$). Thus, although the VOR gain of gain-down-pretrained mice was restored to a value similar to that of naive mice by the gain-up stimulus, the capacity for learning in response to the gain-up stimulus was not restored to the naive state.

Reversibility of Changes in VOR Phase

The gain data described above indicated that the changes in the VOR circuit induced by gain-down training were not fully reversed by subsequent gain-up training. However, from these data alone it was not clear whether the changes induced by gain-up training reflected solely a masking of the effects of prior training, or whether gain-up training could reverse any of the

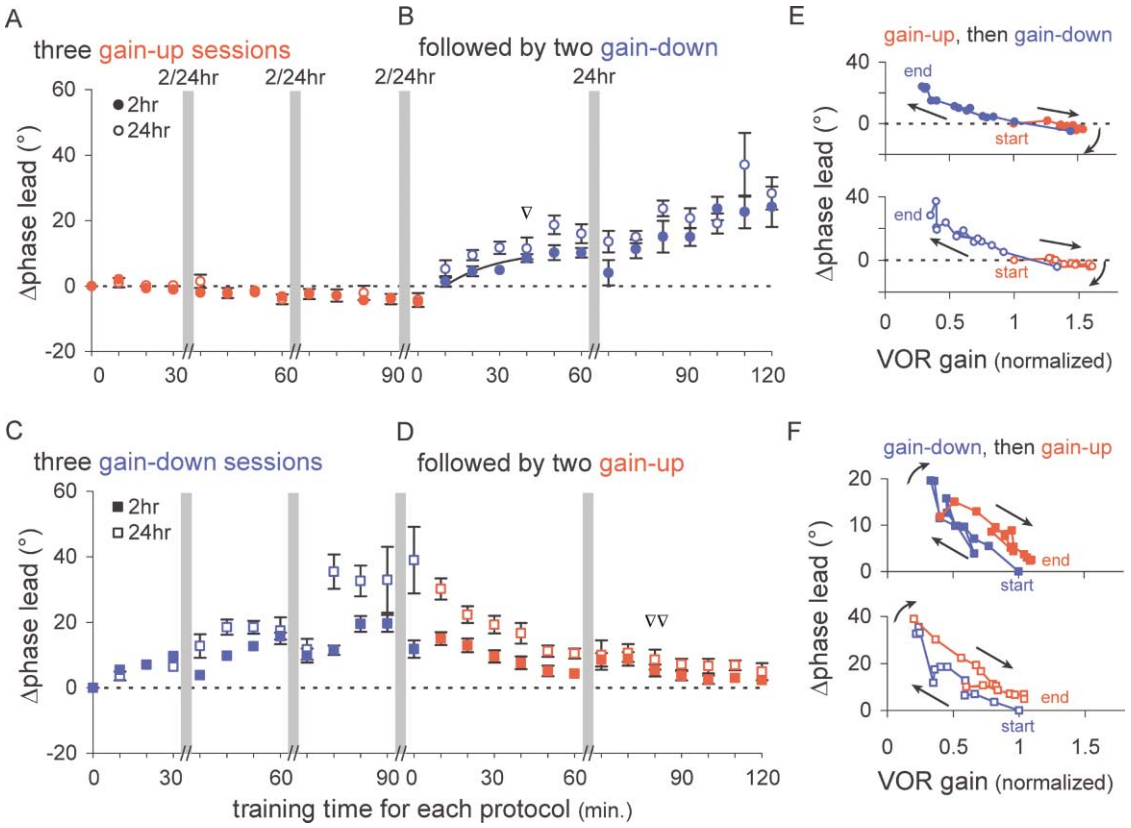

Figure 4. Changes in VOR Phase Induced when Gain-Up and Gain-Down Training Sessions Are Concatenated

(A–D) Changes in VOR phase accompanying the changes in VOR gain shown in Figures 2A–2D, respectively. Symbols are as in Figure 2.

(E) Phase change plotted versus gain change, during the experiments where gain-up training (red), with either 2 hr (top) or 24 hr (bottom) rests, is followed by gain-down training (blue). "Start" and "end" indicate the beginning and end of the entire protocol. Arrows indicate progression of time.

(F) Phase change plotted versus gain change, during the experiments where gain-down training (blue), with either 2 hr (top) or 24 hr (bottom) rests, is followed by gain-up training (red). Hysteresis is prominent in the gain-phase relation.

effects of prior training on the VOR circuit. Since the phase of the VOR was differentially affected by gain-up and gain-down training (Figure 1F), this provided an additional variable for assessing whether reversal or masking was occurring.

From the naive state, gain-down training induced a large increase in phase lead, but gain-up training had little effect on phase (Figures 4A and 4C). Thus, if at the neural level gain-up training simply masked the effects of prior gain-down training by superimposing additional changes in the VOR circuit, then gain-up training should have little effect on the phase lead in the mice pretrained with gain-down stimuli. However, if at the neural level gain-up training reversed the effects of gain-down training, then gain-up training should reduce the phase lead induced by gain-down training. Consistent with the latter scenario, gain-up training caused a reduction in phase lead in mice pretrained with gain-down stimuli (Figure 4D; $p < 0.05$, WSRT). However, this decline in phase lead was slower than one would expect if the changes during gain-up training were mediated exclusively by reversal of the neural events induced by prior gain-down training. For example, when mice pretrained with gain-down stimuli had reattained their initial VOR gain (Figure 2D, $\nabla\nabla$), their phase lead was still larger than its initial value ($p < 0.05$; Figure 4D, $\nabla\nabla$). This partial phase restoration suggests that gain-up training both partially masked and partially reversed the effects of prior gain-down training.

The degree of masking and reversal can be visualized by plotting the VOR phase change versus the VOR gain change throughout the entire experiment. If complete reversal occurred, gain-up training would change gain and phase along the same trajectory as prior gain-down training, but in the opposite direction. If complete masking occurred, gain-up training would change gain without changing phase, as gain-up training does in naive mice (a horizontal excursion on the plot). In our experiments, there was hysteresis in the gain-phase curve for gain-down training followed by gain-up training (Figure 4F, top), with phase leads for any particular gain being larger during the gain-up training part of the experiment than during the initial gain-down training. This trajectory is consistent with masking and reversal occurring in parallel throughout the gain-up training.

In contrast, the phase data for training in the opposite order were consistent with complete reversal of changes induced by gain-up training, as suggested by the gain data. In these mice, gain-down training induced a change in phase lead similar to the change seen in naive

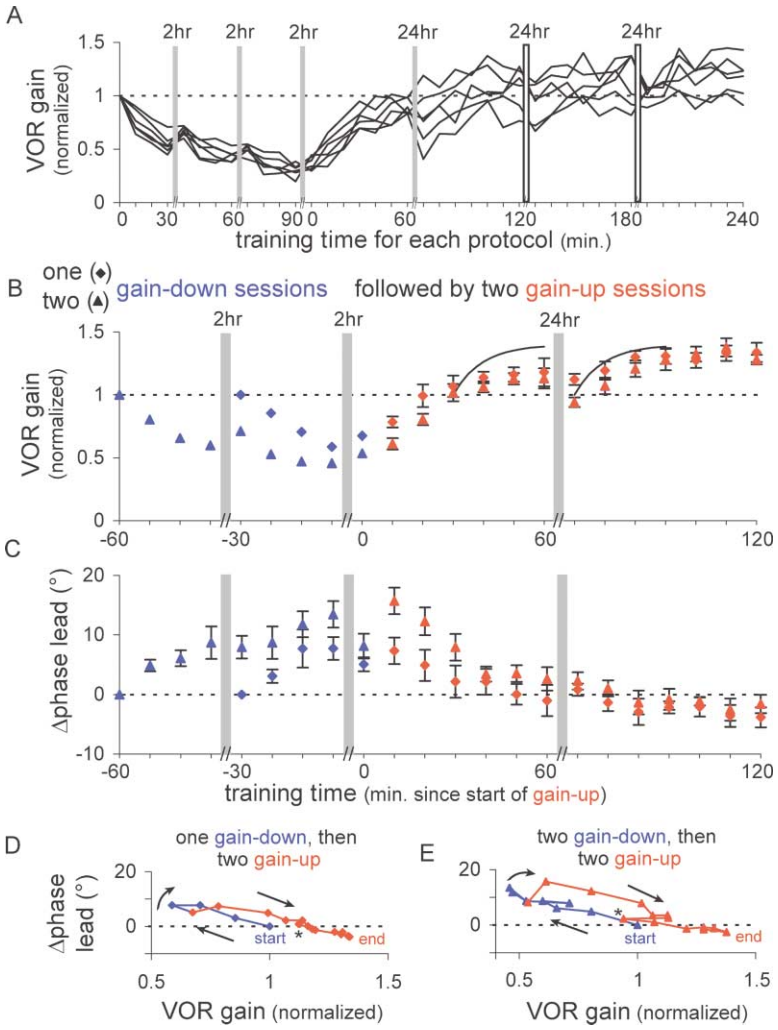

Figure 5. Dependence of Reversal of Decreases in VOR Gain on Amount of Training

(A) Effects of extensive gain-up training and periods of normal visual experience following three gain-down sessions (n = 6). Time courses are shown for each of the six individual mice run on this protocol. Shaded bars indicate periods of darkness, whereas open bars indicate periods of normal light-dark cycles.

(B) Changes in VOR gain observed during one (diamonds; n = 7) or two (triangles; n = 7) gain-down training sessions (blue symbols), followed by gain-up training (red symbols). The solid curves beginning at the 30 and 60 min points in gain-up training are the exponential fit to gain-up training from the naive state (Figure 1E), for comparison.

(C) Changes in VOR phase, for one or two gain-down training sessions, followed by gain-up training.

(D and E) Phase change plotted versus gain change, during the experiments in which one (D) or two (E) gain-down training sessions (blue symbols) are followed by gain-up training (red symbols). Asterisk marks the beginning of the second gain-up training session.

mice undergoing gain-down training (compare Figures 4B and 4C). At time points during gain-down training when naive mice and mice pretrained with gain-up had similar gains, the phase leads were also indistinguishable ($p > 0.05$; Figure 4B, ∇). Reversal can be seen when gain is plotted versus phase for both the gain-up and gain-down parts of the experiment (Figure 4E, top). The single trajectory of this curve is consistent with complete reversal of the effects of gain-up training by gain-down training.

Reversibility Depends on Amount of Training

Permanent irreversibility of the effects of gain-down training would be maladaptive, limiting the capacity for further learning in the VOR. To determine whether the irreversibility we observed could be overcome by more extended training, we subjected a subset of the mice to additional 1 hr gain-up training sessions, separated by 24 hr periods of normal light-dark cycles in the home cage. After two additional gain-up training sessions, mice began to show learning curves more similar to those in naive mice (Figure 5A). Thus, the irreversibility of the effects of gain-down training was present primarily at short timescales and either disappeared with time or was overcome by more extensive gain-up training.

We next tested whether the ability of gain-up training to reverse the effects of gain-down training depended on the amount of gain-down pretraining. We trained mice with one or two 30 min gain-down training sessions, followed by two 1 hr sessions of gain-up training separated by a 24 hr rest. The changes induced by one or two gain-down training sessions (Figure 5B) were more readily reversed than those induced by three gain-down training sessions (Figure 2C). From the beginning of the second gain-up training session, mice increased their VOR gains in a way that paralleled gain-up training in naive mice (Figure 5B; n = 7 each), reaching values comparable to those reached by naive mice with 30 min of gain-up training ($p > 0.05$ for either one or two gain-down pretraining sessions, MWUT). This suggests that by the start of the second gain-up training session, these mice had not only reattained their initial VOR gain but were capable of normal learning in response to the gain-up stimulus. The phase changes during gain-up training also were consistent with reversal of changes induced by gain-down pretraining. During gain-up training, phase leads for mice previously trained with one or two gain-down sessions declined, indicating that reversal of the changes induced by gain-down training had occurred (Figure 5C). When these mice had reattained their

initial VOR gain at the beginning of the second gain-up training session, their phase lead was also restored to its initial value ($p > 0.05$; WSRT).

Plots of the changes in gain versus the changes in phase in these experiments suggest that reversal of the effects of gain-down training on the circuit for the VOR had occurred by the beginning of the second gain-up session (Figures 5D and 5E). During the first hour of gain-up training, hysteresis was apparent in the gain-phase plot, especially for mice pretrained with two gain-down sessions. The shape of this curve suggests that both masking and reversal were occurring in the circuit. However, during the second gain-up session, which resulted in an increase in gain similar to that in naive mice, the hysteresis diminished as the gain-phase relation converged on the initial path (Figure 5E, asterisk). This is consistent with complete reversal, suggesting that the capacity for learning had been restored to a state similar to that in naive mice. Thus, the effects of a limited amount of gain-down training were more readily reversible than the effects induced by more extended gain-down training. However, this reversal still required a greater amount of training, compared to reversal of the effects of gain-up training.

Finally, we modulated the dose of training by a manipulation that did not modify total training time, but instead increased the amount of resting time after each of the first three training sessions from 2 hr ("massed" training) to 24 hr ("spaced" training). In many learning systems, spaced training results in larger or more robust changes than massed training. Accordingly, the gains of mice undergoing spaced gain-down training decreased more than those undergoing massed gain-down training (Figure 2C, open symbols for spaced training). Some of this difference was due to the VOR gain returning toward its initial value when observed 2 hr, but not 24 hr, after training. For mice undergoing massed gain-down training (Figure 2C, filled symbols), the VOR returned toward the initial gain during the first 2 hr rest ($p < 0.05$, WSRT). However, there was negligible change ($p > 0.05$; $n = 9$) during the first 24 hr rest for mice undergoing spaced gain-down training (Figure 2C, open symbols). Similarly, during the second and third rests, the VOR gain returned toward its initial value for mice undergoing training with massed, but not spaced, gain-down stimuli ($p < 0.05$ for massed; $p > 0.05$ for spaced). After three gain-down training sessions, the VOR gains (Figure 2C) and phases (Figure 4C) of mice undergoing massed and spaced training were quite different ($p < 0.05$ for both, MWUT). Thus, spaced training facilitates the gain and phase changes induced by gain-down training.

When we subsequently exposed these mice to gain-up training sessions, the VOR gain of mice pretrained with spaced gain-down sessions was consistently lower, at each time point, than that of mice pretrained with massed gain-down sessions (Figure 2D), and the phase consistently showed greater lead (Figure 4D). Thus, the enhanced decrease in VOR gain induced by spaced gain-down training had a residual effect that persisted throughout subsequent gain-up training.

In contrast to the results with gain-down training, we saw no difference between massed and spaced gain-up training sessions ($p > 0.05$ at each time point, MWUT; $n = 9$ for spaced condition; Figure 2A). Likewise, there was no difference between the effects of massed and spaced gain-up training on VOR phase (Figure 4A). During subsequent gain-down training, the VOR gains and phases of mice pretrained with the spaced and massed gain-up stimuli changed in similar ways, consistent with complete reversal in both cases (Figures 2B and 4B). Thus, whereas the enhancement of gain-down training by longer rests retarded reversal by gain-up training, there was no effect of spaced training either on learning in response to the gain-up stimulus or on the vulnerability of these changes to reversal by gain-down training.

Discussion

The simple anatomy and well-characterized plasticity mechanisms of the cerebellum make it a good system for the study of memory. Motor learning in the VOR is a cerebellum-dependent task well suited for studying how new and old memories interact, since the VOR can undergo bidirectional changes in gain. One straightforward idea is that oppositely directed changes in VOR gain could be implemented in the brain by inverse plasticity mechanisms, e.g., LTP and LTD at a particular synaptic site. However, most previous studies have focused on a model in which an increase and a decrease in VOR gain are each implemented using the same synaptic plasticity mechanism, namely LTD at parallel fiber-Purkinje cell synapses ("cerebellar LTD") (Ito, 1972).

In its most general form, the cerebellar LTD model suggests that the diversity of signals carried by parallel fibers would enable a single plasticity mechanism, applied to different sets of parallel fiber synapses, to mediate a diverse set of stimulus-response associations (Albus, 1971; Marr, 1969). As applied to the VOR, this model attributes both an increase and a decrease in VOR gain to a single synaptic plasticity mechanism by suggesting that cerebellar LTD operates independently on parallel fibers that are active at different times during the VOR (Figure 6A; Ito, 1972; Ito, 1982). More specifically, it has been proposed that LTD of parallel fibers firing during ipsiversive head turns would induce an increase in VOR gain, whereas LTD of parallel fibers firing during contraversive head turns would induce a decrease in VOR gain. This model predicts that (1) increases and decreases in VOR gain would have similar properties due to their shared plasticity mechanism, and (2) increases and decreases in VOR gain would not reverse each other at the mechanistic level. Neither prediction is borne out by our results. We find that increases and decreases in VOR gain exhibit key differences, suggesting that (1) they depend upon different cellular plasticity mechanisms, and (2) these plasticity mechanisms reverse each other with unequal efficacy.

Different Time Courses for Increases and Decreases in VOR Gain

Increases and decreases in VOR gain possess different temporal properties. An early study using long training periods in primates found that in the absence of visuovestibular stimuli, increases in VOR gain decayed more than decreases over the course of several days (Miles and Eighmy, 1980). Our study found additional differences in the time courses of increases and decreases

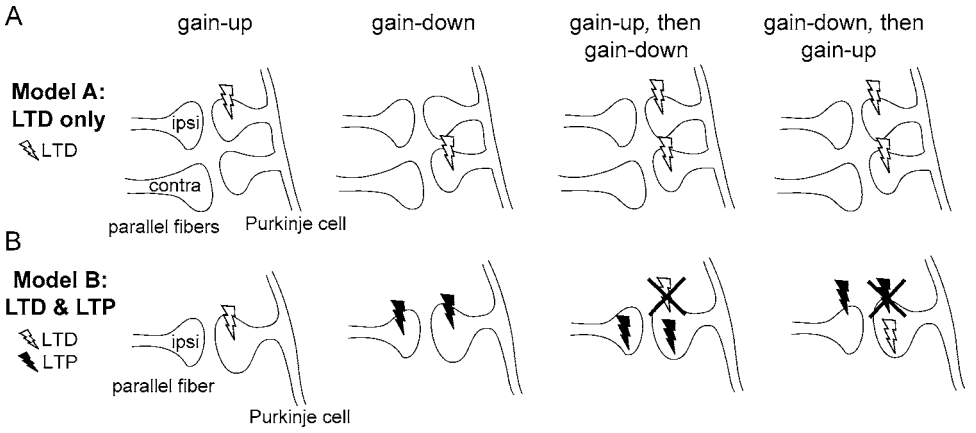

Figure 6. Two Classes of Model for How an Increase and a Decrease in VOR Gain Are Each Encoded

(A) A model using the same plasticity mechanism for both an increase and a decrease in VOR gain. For each behavioral training paradigm (gain-up, gain-down, gain-up then gain-down, and gain-down then gain-up), two parallel fibers are shown synapsing onto a single Purkinje cell dendrite. One parallel fiber fires during ipsiversive head turns (toward the side of the Purkinje cell), and the other fires during contraversive head turns (away from the side of the Purkinje cell). Lightning bolts indicate plasticity mechanisms being considered at the parallel fiber-Purkinje cell synapse. It has been proposed that LTD (open lightning bolts) of ipsiversive-responding parallel fibers induces an increase in VOR gain (gain-up), and LTD of contraversive-responding parallel fibers induces a decrease in VOR gain (gain-down) (Ito, 1972; Ito, 1982). When an increase and a decrease in VOR gain are concatenated in either order (gain-up then gain-down, gain-down then gain-up), the predicted changes mask but do not reverse each other. LTD is induced in both synapses, in either case.

(B) A model using different plasticity mechanisms for an increase and a decrease in VOR gain. Parallel fiber-Purkinje LTD is expressed postsynaptically, but there are both pre- and postsynaptically expressed forms of LTP at this synapse (Hemart et al., 1994; Lev-Ram et al., 2002; Linden et al., 1991; Salin et al., 1996). Only ipsiversive-responding parallel fibers are needed, consistent with the observation that the majority of Purkinje cells increase their firing during ipsiversive head motion, when the eyes are held still (Lisberger and Fuchs, 1978). An increase in VOR gain is implemented as in the original cerebellar LTD model (gain-up), but in this model a decrease in VOR gain is mediated by both pre- and postsynaptic forms of LTP (gain-down). When an increase in VOR gain is followed by a decrease in gain, first postsynaptically expressed LTD is induced, followed by pre- and postsynaptically expressed LTP (gain-up then gain-down). Reversal is complete if the postsynaptic component of LTP erases postsynaptic LTD. When a decrease is followed by an increase, however, first pre- and postsynaptically expressed LTP are induced, followed by postsynaptically expressed LTD (gain-down then gain-up). Even if the postsynaptically expressed plasticity components reverse each other, the presynaptic component of LTP is not reversed. This general class of models could explain the asymmetry in reversibility of increases and decreases in VOR gain.

in VOR gain, using acute training protocols in mice. In our study, we found that the increase in gain saturated quickly despite the continued presence of significant tracking error, whereas the decrease in gain did not saturate despite reduction of tracking error. Further examination of the effects of gain-up training revealed two distinct temporal components of memory expression, one of which lasted less than 2 hr, and one of which persisted for at least 24 hr. These separate components are reminiscent of the short- and long-lasting components of expression seen in vitro for many plasticity mechanisms, including cerebellar LTD (Ahn et al., 1999; Murashima and Hirano, 1999). After repeated gain-down training, however, we did not find any decay of the resultant motor memories over 24 hr. When we compared the effects of massed and spaced training on gain-down stimuli, we found reduced expression of decreased VOR gain after 2 hr, but not 24 hr, rests. The slower component of memory expression after gain-down training reflects the delayed expression of memories induced 24 hr previously. Delayed components of memory processes have been described in several learning systems, such that expression of plasticity is reduced at an intermediate time point, only to return to higher levels of expression later (Schulz et al., 1999; Sutton et al., 2001). The differences in the time course of acquisition, expression, and decay of the motor memories for an increase versus decrease in VOR gain are consistent with different plasticity mechanisms being responsible for each, in contrast to the cerebellar LTD model, which would predict similar temporal components of expression for these two memories.

A pharmacological study in goldfish provides support for the idea that increases and decreases in VOR gain depend on different plasticity mechanisms. Induction of LTD can be blocked by inhibiting nitric oxide (NO) activity in Purkinje cells (Crepel and Jaillard, 1990; Shibuki and Okada, 1991). It has been reported that blocking NO signaling in the cerebellum of goldfish affects increases but not decreases in VOR gain (Li et al., 1995). This is consistent with the idea that increases in VOR gain depend, more than decreases in VOR gain, upon an NO-dependent process such as cerebellar LTD.

A Rule Governing Models of Memory Storage Mechanisms: Asymmetric Reversibility

Our behavioral results provide insight into the sets of plasticity mechanisms that could mediate oppositely directed cerebellum-dependent motor memories. In particular, our study of the interaction between increases and decreases in VOR gain constrains how the different plasticity mechanisms mediating these changes must interact, namely that they must reverse each other in an asymmetric fashion.

At the behavioral level, we found striking asymmetries in the reversal of increases and decreases in VOR gain.

Gain-down training after gain-up training not only restored the VOR gain to its initial state, but also apparently restored the capacity for learning in response to the gain-down stimulus. In contrast, acute gain-up training only partially reversed the effects of gain-down training; even when the VOR gain was restored to its initial value, the capacity for learning in response to the gain-up stimulus was not restored to its initial state. The time course of reversal was also different. Gain-down training reversed the effects of gain-up training with a rapid time course (minutes), whereas the reversal of the effects of gain-down training by gain-up training was slow (hours). This difference in the time course during reversal training was observed despite more similar time constants of learning for increases and decreases in gain from the naive state, and it suggests a mechanistic difference in the reversal of these two behavioral states.

These differences in amplitude and time course of reversal are difficult to explain with a model relying on a single plasticity mechanism for increases and decreases in VOR gain. In the cerebellar LTD model, for example, both an increase and a decrease in VOR gain rely on LTD of separate sets of parallel fiber synapses. Whether an increase in gain is followed by a decrease, or a decrease is followed by an increase, the end state predicted by the model would be the same: both sets of synapses would undergo LTD, potentially to the point of saturation (Figure 6A). This kind of model cannot readily account for the differences we observed between the reversal of increases and decreases in gain.

The cerebellar LTD model also would predict that at the mechanistic level, the changes mediating increases and decreases in VOR gain would mask rather than reverse each other (Lisberger, 1996; Sejnowski, 1977). In contrast to this prediction of pure masking, several observations support the idea that the initial neural changes are being reversed. First, when the VOR is restored to the basal state by gain-down training after gain-up training, the VOR is capable of normal changes in response to the gain-down stimulus. This is consistent with the circuit being restored to the naive state. Second, gain-up training does not alter the phase of the VOR from the naive state, but it can partially reverse the phase changes induced by gain-down training. This is consistent with a partial reversal of the neural events mediating the previous decrease in VOR gain, as opposed to pure superposition of the effects of gain-up training upon the previous changes. Particularly telling was the observation that when the phase changes induced by prior gain-down training were completely reversed, normal learning was observed in response to the gain-up stimulus (Figures 5D and 5E). This suggests that the reversal of the change in phase reflected a return of the circuit to the naive state. Thus, our data are most consistent with some active reversal of the neural changes mediating both increases and decreases in VOR gain, yet there is clearly an asymmetry. Whereas the changes induced by gain-down training are clearly not fully reversed by 2 hr of gain-up training, it is possible that the changes induced by gain-up training are fully reversed by gain-down training, perhaps within 10 min.

We found that the ability of gain-up training to reverse the effects of gain-down training depended on the dose of gain-down training experienced. This suggests a threshold beyond which the effects of gain-down training become significantly less reversible. Longer periods of gain-down training must result in changes that are less subject to reversal by the plasticity mechanisms engaged by gain-up training.

How Might Asymmetric Reversibility Be Implemented at the Neural Level?

The asymmetrically reversible plasticity mechanisms predicted by our results could take several forms. One mechanistic difference that could result in asymmetric reversibility at the behavioral level is that different sites in the circuit could be used for storing increases and decreases in gain. In the VOR, evidence from in vivo recordings suggests that learning produces changes that are distributed between the cerebellum and brainstem (Lisberger et al., 1994a, 1994b; Miles et al., 1980; Partsalis et al., 1995a). Nevertheless, it seems unlikely that memories of increases and decreases in gain are stored in separate parts of the brain, since posttraining lesions of the cerebellum have similar effects on both learned increases and decreases in VOR gain (Luebke and Robinson, 1994; McElligott et al., 1998; Michnovicz and Bennett, 1987; Partsalis et al., 1995b; Pastor et al., 1994).

Another possible mechanistic difference is that increases and decreases in VOR gain could be mediated by oppositely directed synaptic changes, which asymmetrically reverse each other. The reversibility of plasticity has been examined at a few sites in the vestibulocerebellar circuit. Full reversibility has been reported for LTP and LTD of vestibular inputs to the vestibular nuclei and for changes in intrinsic excitability of Purkinje target neurons in the vestibular nucleus (Caria et al., 2001; Smith et al., 2002). However, there are many sites in the VOR circuit for which reversibility of plasticity mechanisms has not been characterized. Reversibility at the parallel fiber-Purkinje cell synapse has not been explicitly analyzed, although the differential pre- and postsynaptic localization of LTP and LTD at this synapse constrains how the reversal of plasticity must operate at this site. LTP at this synapse seems to have pre- and postsynaptically expressed components (Lev-Ram et al., 2002; Salin et al., 1996), but LTD seems to be postsynaptically expressed (Linden et al., 1991). A model that explores how differential localization of plasticity mechanisms within a synapse could lead to asymmetric reversibility at the behavioral level is described in Figure 6B.

A second example of how asymmetric reversibility of LTP and LTD might arise at a synapse is if there were specific temporal requirements for reversal of synaptic changes. It has been reported that the ability of low-frequency stimulation to reverse hippocampal LTP can either increase or disappear in the 30 min following LTP induction, depending on the precise history of synaptic strength (Montgomery and Madison, 2002; Staubli and Chun, 1996). Since the reversal power of gain-up training depends on the duration and timing of prior gain-down training, this may suggest a constraint on the temporal requirements for reversal of the underlying plasticity mechanisms. This issue has not been explored for plasticity mechanisms in the cerebellum, and our results

suggest that in vitro experiments on the dose and timing dependence of plasticity in this circuit would be highly informative.

A third asymmetry that could account for the asymmetric reversibility constraint we find would be a difference in the spread of LTP and LTD to nonactivated synapses. Some plasticity mechanisms may operate on "volumes" of colocalized synapses (Montague and Sejnowski, 1994). Memories encoded by plasticity mechanisms capable of spreading to nonactivated synapses may be less reversible than those encoded by plasticity mechanisms not capable of spread. For cerebellar LTD, it has been shown that pairing parallel fiber and climbing fiber activation causes a decrease in synaptic strength that is also expressed at nearby unpaired synapses (Reynolds and Hartell, 2000; Wang et al., 2000). However, the spread of other plasticity mechanisms in the cerebellum, such as LTP at this synapse, has not been characterized. Experimental measurement of this property in the context of reversal may illuminate how reversal of plasticity operates over ensembles of synapses, which is important for understanding how plasticity operates in the behaving animal.

With further knowledge of the reversal properties of plasticity mechanisms in the brain, it will be possible to understand the strategies used by learning systems when circumstances change. Ultimately, the operation of a circuit capable of storing memories must be understood not only in terms of the set of plasticity mechanisms available for learning, but the entire sequence of mechanisms that are engaged to encode a history of memories as they are consolidated or erased.

Experimental Procedures

Experiments were performed on 78 male C57BL/6 mice, 9–12 weeks old, from Charles River Labs (Wilmington, MA). All procedures were approved by the Stanford University Administrative Panel for Laboratory Animal Care (APLAC).

Surgical Procedure

Each mouse was anesthetized with ketamine/medetomidine, followed by isoflurane. After making a midline incision along the scalp, three screws were embedded in the skull. Using forceps, a pocket was blunt-dissected beneath the conjunctiva of the temporal portion of the right eye. An 80-turn copper scleral search coil (IET, Marly, Switzerland), 1 mm in diameter, was glued into the pocket with Vetbond (3M Animal Care, St. Paul, MN). The twisted wire leads were threaded through the top of the eye, emerging from under the scalp near bregma. A few millimeters of the wire were tucked under the skin just posterior to the eye to provide slack for eye motion. The ends of the wires were soldered to a 2-pin connector. This connector and a plastic headpost (placed approximately over lambda) were cemented with dental acrylic to the three anchor screws (Henry Schein, Melville, NY).

Experimental Equipment

During each behavioral experiment, the head of the mouse was immobilized by placing it in a custom-made restrainer to which its headpost was fixed. Vestibular stimuli were applied to the mouse by rotating this restrainer, mounted on a computer-controlled turntable (Carco IGTS, Pittsburgh, PA). Optokinetic stimuli were applied by rotating a hemispherical drum, 30 cm in diameter, mounted on a motor with a shaft encoder (Gurley Precision Instruments, Troy, NY). This motor was driven from a PC by a controller card (Precision MicroControl, Carlsbad, CA) and an analog amplifier. The drum was made of white translucent plastic with black vertical stripes subtending 7.5° visual angle. The drum was backlit by two 60-watt bulbs placed 6 inches outside the drum. A silvered acrylic plate (McMaster Carr) was placed under the mouse to provide nearly full-field visual motion. A set of 18-inch magnetic coils (CNC Engineering, Seattle, WA), fixed to the turntable, provided the signals for measuring eye position using the mouse's scleral search coil (Judge et al., 1980; Robinson, 1963). Eye velocities were calculated from eye position with an analog differentiator and filter (corner frequency 300 Hz; designed by S.G. Lisberger). Signals were digitized at 500 Hz.

Behavioral Protocols

The VOR gain was measured by delivering 1 Hz, ±10°/s peak velocity sinusoidal turntable rotations in the dark. Measurements were taken in 30 s blocks. Any cycle containing a saccade or motion artifact was deleted. Head and eye velocity traces were aligned on the zero crossings of head velocity, then averaged (8–15 artifact-free traces per block). Fourier analysis was then used to extract amplitude and phase from the averaged traces. The VOR gain was calculated to be the ratio of the eye to the head velocity amplitudes, and the VOR phase was calculated to be the eye velocity phase minus the head velocity phase, minus 180°. A perfectly compensatory VOR would thus have a phase of zero. The naive VOR gain for mice, measured on the sixth day after surgery, was 0.43 ± 0.12 ($n = 70$; mean ± SD). The optokinetic reflex (OKR) gain was measured by delivering 1 Hz, ±10°/s peak velocity sinusoidal illuminated drum rotation, and calculated as the ratio of averaged eye velocity amplitude to averaged drum velocity amplitude. The gain-up stimulus consisted of 1 Hz, ±10°/s sinusoidal turntable rotation paired with oppositely directed 1 Hz, ±5°/s sinusoidal drum rotation (Figure 1A). For the gain-down stimulus, the illuminated drum was held stationary relative to the mouse, while he experienced ±10°/s sinusoidal turntable rotation (Figure 1B). For each training session, mice were trained in three or six 10 min periods. After each 10 min period, the VOR was measured during two 30 s blocks of turntable rotation in the dark. Between each block of a multiple-block measurement, an attention-normalizing stimulus (usually a flash of light, but a sharp noise like a clap worked equivalently) was given, followed by an 8 s pause before beginning the eye movement measurement. Measurements were also made of the eye movements in the presence of the gain-up or gain-down stimuli, at the beginning and end of each 10 min training session. Retinal image slip was calculated by extracting the amplitude and phase from the averaged difference between drum velocity and eye movement velocity. During experiments with multiple training sessions, night vision goggles were used to transfer the animals to and from their cages, which were kept in a completely dark chamber during the rest periods.

Schedule of Acclimatization and Calibration

On the sixth day after surgery, each mouse was acclimatized to head restraint for two 15 min sessions. During the first of these sessions, the mouse's scleral search coil was calibrated by rotating the magnetic field coils sinusoidally (±10°/s peak velocity) around the mouse, which was held stationary in darkness. During the second 15 min acclimatization session, the VOR gain, OKR gain, and eye movement responses to the gain-up and gain-down stimuli were measured. To minimize possible learning effects during these measurement sessions, only one block of data was taken in each condition. Six mice out of the 78, with obvious eye damage or impaired visuomotor ability (defined as VOR or image tracking ability lower than two standard deviations below the mean), were not experimented upon. Two additional mice were not experimented upon because of irregular basal VOR gain. All experiments began on day 7 after surgery, at which time the eye healing process appeared to be complete (van Alphen et al., 2001). Each animal was used for only one sequence of experimental protocols. Figure 1 includes, in part, data from mice that went on to experience further training.

Data Analysis

Custom software (by S.G. Lisberger) was used to analyze eye movement traces. All VOR data were normalized by dividing by the mouse's naive VOR, measured at the beginning of the experiment on the seventh day after surgery. Phase changes during learning were measured with respect to each mouse's initial phase. Statistical analyses (Mann-Whitney U test [MWUT], Wilcoxon signed-rank

test [WSRT]) were conducted using StatView (SAS Inst., Cary, NC). Exponential curve fitting was performed using Matlab.

Acknowledgments

Thanks to Stephen Lisberger for supporting early development of mouse VOR techniques and to Eric Knudsen, Stephen Lisberger, Richard Tsien, Dan Madison, Akira Katoh, Rhea Kimpo, Mike Ke, Joyce Liao, Jenifer Austin, and members of the Knudsen, Lisberger, and Tsien labs for helpful discussions. Thanks also to Pam Louderback, Kate Oh, and Kate Wesley for technical help. This research was supported by a McKnight Scholar Award, a Klingenstein Award in Neuroscience, a Sloan Research Fellowship, a Terman Fellowship, and NIH R01 DC04154 to J.L.R. and a Fannie and John Hertz Fellowship to E.S.B.

Received: February 24, 2003
Revised: June 25, 2003
Accepted: August 22, 2003
Published: September 10, 2003

References

Ahn, S., Ginty, D.D., and Linden, D.J. (1999). A late phase of cerebellar long-term depression requires activation of CaMKIV and CREB. Neuron *23*, 559–568.

Albus, J. (1971). A theory of cerebellar function. Math. Biosci. *10*, 25–61.

Caria, M.A., Melis, F., Solinas, A., Tavera, C., and Mameli, O. (2001). Frequency-dependent LTP/LTD in guinea pig Deiters' nucleus. Neuroreport *12*, 2353–2358.

Crepel, F., and Jaillard, D. (1990). Protein kinases, nitric oxide and long-term depression of synapses in the cerebellum. Neuroreport *1*, 133–136.

Fujii, S., Saito, K., Miyakawa, H., Ito, K., and Kato, H. (1991). Reversal of long-term potentiation (depotentiation) induced by tetanus stimulation of the input to CA1 neurons of guinea pig hippocampal slices. Brain Res. *555*, 112–122.

Gonshor, A., and Jones, G.M. (1973). Proceedings: changes of human vestibulo-ocular response induced by vision-reversal during head rotation. J. Physiol. *234*, 102P–103P.

Han, V.Z., Grant, K., and Bell, C.C. (2000). Reversible associative depression and nonassociative potentiation at a parallel fiber synapse. Neuron *27*, 611–622.

Hemart, N., Daniel, H., Jaillard, D., and Crepel, F. (1994). Properties of glutamate receptors are modified during long-term depression in rat cerebellar Purkinje cells. Neurosci. Res. *19*, 213–221.

Ito, M. (1972). Neural design of the cerebellar motor control system. Brain Res. *40*, 81–84.

Ito, M. (1982). Cerebellar control of the vestibulo-ocular reflex—around the flocculus hypothesis. Annu. Rev. Neurosci. *5*, 275–296.

Ito, M., Shiida, T., Yagi, N., and Yamamoto, M. (1974). The cerebellar modification of rabbit's horizontal vestibulo-ocular reflex induced by sustained head rotation combined with visual stimulation. Proc. Jpn. Acad. *50*, 85–89.

Judge, S.J., Richmond, B.J., and Chu, F.C. (1980). Implantation of magnetic search coils for measurement of eye position: an improved method. Vision Res. *20*, 535–538.

Lev-Ram, V., Wong, S.T., Storm, D.R., and Tsien, R.Y. (2002). A new form of cerebellar long-term potentiation is postsynaptic and depends on nitric oxide but not cAMP. Proc. Natl. Acad. Sci. USA *99*, 8389–8393.

Li, J., Smith, S.S., and McElligott, J.G. (1995). Cerebellar nitric oxide is necessary for vestibulo-ocular reflex adaptation, a sensorimotor model of learning. J. Neurophysiol. *74*, 489–494.

Linden, D.J., Dickinson, M.H., Smeyne, M., and Connor, J.A. (1991). A long-term depression of AMPA currents in cultured cerebellar Purkinje neurons. Neuron *7*, 81–89.

Lisberger, S.G. (1996). Motor learning and memory in the vestibulo-ocular reflex: the dark side. Ann. N Y Acad. Sci. *781*, 525–531.

Lisberger, S.G., and Fuchs, A.F. (1978). Role of primate flocculus during rapid behavioral modification of vestibuloocular reflex. I. Purkinje cell activity during visually guided horizontal smooth-pursuit eye movements and passive head rotation. J. Neurophysiol. *41*, 733–763.

Lisberger, S.G., Pavelko, T.A., Bronte-Stewart, H.M., and Stone, L.S. (1994a). Neural basis for motor learning in the vestibuloocular reflex of primates. II. Changes in the responses of horizontal gaze velocity Purkinje cells in the cerebellar flocculus and ventral paraflocculus. J. Neurophysiol. *72*, 954–973.

Lisberger, S.G., Pavelko, T.A., and Broussard, D.M. (1994b). Neural basis for motor learning in the vestibuloocular reflex of primates. I. Changes in the responses of brain stem neurons. J. Neurophysiol. *72*, 928–953.

Luebke, A.E., and Robinson, D.A. (1994). Gain changes of the cat's vestibulo-ocular reflex after flocculus deactivation. Exp. Brain Res. *98*, 379–390.

Marr, D. (1969). A theory of cerebellar cortex. J. Physiol. *202*, 437–470.

McElligott, J.G., Beeton, P., and Polk, J. (1998). Effect of cerebellar inactivation by lidocaine microdialysis on the vestibuloocular reflex in goldfish. J. Neurophysiol. *79*, 1286–1294.

Michnovicz, J.J., and Bennett, M.V. (1987). Effects of rapid cerebellectomy on adaptive gain control of the vestibulo-ocular reflex in alert goldfish. Exp. Brain Res. *66*, 287–294.

Miles, F.A., and Eighmy, B.B. (1980). Long-term adaptive changes in primate vestibuloocular reflex. I. Behavioral observations. J. Neurophysiol. *43*, 1406–1425.

Miles, F.A., and Fuller, J.H. (1974). Adaptive plasticity in the vestibulo-ocular responses of the rhesus monkey. Brain Res. *80*, 512–516.

Miles, F.A., Braitman, D.J., and Dow, B.M. (1980). Long-term adaptive changes in primate vestibuloocular reflex. IV. Electrophysiological observations in flocculus of adapted monkeys. J. Neurophysiol. *43*, 1477–1493.

Montague, P.R., and Sejnowski, T.J. (1994). The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms. Learn. Mem. *1*, 1–33.

Montgomery, J.M., and Madison, D.V. (2002). State-dependent heterogeneity in synaptic depression between pyramidal cell pairs. Neuron *33*, 765–777.

Murashima, M., and Hirano, T. (1999). Entire course and distinct phases of day-lasting depression of miniature EPSC amplitudes in cultured Purkinje neurons. J. Neurosci. *19*, 7326–7333.

Myers, K.M., and Davis, M. (2002). Behavioral and neural analysis of extinction. Neuron *36*, 567–584.

Partsalis, A.M., Zhang, Y., and Highstein, S.M. (1995a). Dorsal Y group in the squirrel monkey. I. Neuronal responses during rapid and long-term modifications of the vertical VOR. J. Neurophysiol. *73*, 615–631.

Partsalis, A.M., Zhang, Y., and Highstein, S.M. (1995b). Dorsal Y group in the squirrel monkey. II. Contribution of the cerebellar flocculus to neuronal responses in normal and adapted animals. J. Neurophysiol. *73*, 632–650.

Pastor, A.M., de la Cruz, R.R., and Baker, R. (1994). Cerebellar role in adaptation of the goldfish vestibuloocular reflex. J. Neurophysiol. *72*, 1383–1394.

Reynolds, T., and Hartell, N.A. (2000). An evaluation of the synapse specificity of long-term depression induced in rat cerebellar slices. J. Physiol. *527*, 563–577.

Robinson, D.A. (1963). Method of measuring eye movement using a scleral search coil in a magnetic field. IEEE Trans. Biomed. Elect. *BME-10*, 137–145.

Robinson, D.A. (1976). Adaptive gain control of vestibuloocular reflex by the cerebellum. J. Neurophysiol. *39*, 954–969.

Salin, P.A., Malenka, R.C., and Nicoll, R.A. (1996). Cyclic AMP mediates a presynaptic form of LTP at cerebellar parallel fiber synapses. Neuron *16*, 797–803.

Schulz, S., Siemer, H., Krug, M., and Hollt, V. (1999). Direct evidence for biphasic cAMP responsive element-binding protein phosphory-

lation during long-term potentiation in the rat dentate gyrus in vivo. J. Neurosci. *19*, 5683–5692.

Sejnowski, T.J. (1977). Statistical constraints on synaptic plasticity. J. Theor. Biol. *69*, 385–389.

Shibuki, K., and Okada, D. (1991). Endogenous nitric oxide release required for long-term synaptic depression in the cerebellum. Nature *349*, 326–328.

Smith, M.R., Nelson, A.B., and Du Lac, S. (2002). Regulation of firing response gain by calcium-dependent mechanisms in vestibular nucleus neurons. J. Neurophysiol. *87*, 2031–2042.

Squire, L.R., and Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. Curr. Opin. Neurobiol. *5*, 169–177.

Staubli, U., and Chun, D. (1996). Factors regulating the reversibility of long-term potentiation. J. Neurosci. *16*, 853–860.

Sutton, M.A., Masters, S.E., Bagnall, M.W., and Carew, T.J. (2001). Molecular mechanisms underlying a unique intermediate phase of memory in aplysia. Neuron *31*, 143–154.

van Alphen, A.M., Stahl, J.S., and De Zeeuw, C.I. (2001). The dynamic characteristics of the mouse horizontal vestibulo-ocular and optokinetic response. Brain Res. *890*, 296–305.

Wang, S.S., Khiroug, L., and Augustine, G.J. (2000). Quantification of spread of cerebellar long-term depression with chemical two-photon uncaging of glutamate. Proc. Natl. Acad. Sci. USA *97*, 8635–8640.