

Changes in Control of Saccades during Gain Adaptation

Vincent Ethier,¹ David S. Zee,² and Reza Shadmehr¹

Departments of ¹Biomedical Engineering and ²Neurology, Johns Hopkins School of Medicine, Baltimore, Maryland 21205

In a typical short-term saccadic adaptation protocol, the target moves intrasaccadically either toward (gain-down) or away (gain-up) from initial fixation, causing the saccade to complete with an endpoint error. A central question is how the motor system adapts in response to this error: are the motor commands changed to bring the eyes to a different goal, akin to a remapping of the target, or is adaptation focused on the processes that monitor the ongoing motor commands and correct them midflight, akin to changes that act via internal feedback? Here, we found that, in the gain-down paradigm, the brain learned to produce a smaller amplitude saccade by altering the trajectory of the saccade. The adapted saccades had reduced peak velocities, reduced accelerations, shallower decelerations, and increased durations compared with a control saccade of equal amplitude. These changes were consistent with a change in an internal feedback that acted as a forward model. However, in the gain-up paradigm, the brain learned to produce a larger amplitude saccade with trajectories that were identical with those of control saccades of equal amplitude. Therefore, whereas the gain-down paradigm appeared to induce adaptation via an internal feedback that controlled saccades midflight, the gain-up paradigm induced adaptation primarily via target remapping. Our simulations explained that, for each condition, the specific adaptation produced a saccade that brought the eyes to the target with the smallest motor costs.

Key words: saccade adaptation; saccade kinematics; forward models; optimal control; computational neuroscience; sensorimotor

Introduction

Saccades, unlike movements like reaching, vocalization, or locomotion, are so brief (rarely lasting >80 ms) that sensory afferents do not play an immediate role in their control. However, this does not imply that saccades are open-loop, ballistic movements. For example, random motor noise (Quaia et al., 2000), repetition of the visual stimulus (Chen-Harris et al., 2008; Golla et al., 2008), certain drugs (Jürgens et al., 1981), expectation of reward (Takikawa et al., 2002), or coupling of the eye movement with an arm movement (Snyder et al., 2002; van Donkelaar et al., 2004) all produce changes in the peak velocities of saccades without necessarily affecting saccade amplitude. That is, although there are many factors that alter the motor commands that initiate saccades, the brain is somehow able to maintain the endpoint accuracy of the saccade. One possibility is that this is because of an internal feedback process that monitors the motor commands and corrects them on-line.

Computational theories in motor control describe forward models as systems that receive a copy of the ongoing motor commands and predict their sensory consequences. The concept is that, if the forward model is accurate, variability in the motor commands that initiate the saccade can be corrected via commands that arrive later during the movement. To maintain accuracy, forward models should constantly learn from movement errors. Indeed, we recently observed that, in a paradigm using

cross-axis adaptation, endpoint errors that were perpendicular to the direction of the saccade produced an adaptive response that included curvature (Chen-Harris et al., 2008). That is, the brain appeared to correct the saccade midflight, consistent with the prediction that the forward model had learned from the endpoint errors.

However, our experiment in cross-axis adaptation also revealed that curvature was only one component of the adaptive response to endpoint errors. Another component was the changes in the motor commands that initiated the saccade. With training, these motor commands gradually aimed the saccade closer to the eventual target, suggesting that part of the adaptive response was akin to target remapping. Therefore, in response to endpoint errors, the brain appeared to use two distinct mechanisms of adaptation: one was similar to target remapping (i.e., initiate a saccade toward an updated goal) and another was similar to adapting the internal feedback (i.e., the forward model).

Compared with the cross-axis paradigm, a far more studied example of saccade adaptation is the in-axis paradigm (McLaughlin, 1967). In this paradigm, the target of the saccade moves intrasaccadically either toward (gain-down) or away (gain-up) from initial fixation. It is not known whether in-axis adaptation is an example of target remapping or an example of changing internal feedback. Here, we show that the specific changes that take place in saccade trajectories suggest that, whereas the mechanism of adaptation in the gain-down paradigm is consistent with changes in the forward model, in the gain-up paradigm the mechanism of adaptation is consistent with target remapping. We suggest that the reason why the brain chooses different mechanisms of adaptation for gain-up versus gain-down is because adapting the internal feedback minimizes

Received July 24, 2008; revised Nov. 5, 2008; accepted Nov. 6, 2008.

This work was supported by National Institutes of Health Grants NS057814 and EY01849.

Correspondence should be addressed to Reza Shadmehr, Johns Hopkins School of Medicine, 410 Traylor Building, 720 Rutland Avenue, Baltimore, MD 21205. E-mail: shadmehr@jhu.edu.

DOI:10.1523/JNEUROSCI.3470-08.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/2813929-09\$15.00/0

“motor costs” in the gain-down paradigm, whereas target remapping minimizes motor costs in the gain-up paradigm.

Materials and Methods

Subjects were recruited from our medical school community. Group 1 subjects trained in a gain-down paradigm and returned later for a control experiment in which the target positions were tailored to their previous adaptation performance. Group 2 subjects trained in a gain-up paradigm and returned later for a similar control experiment. Five subjects participated in each group. There was a break of at least 1 d between the adaptation and control experiments. Subjects gave written consent and the protocol was approved by The Johns Hopkins Institutional Review Board.

Experimental setup

The position of either the right or the left eye was measured with a magnetic-field search-coil system (Robinson, 1963) using directional scleral annuli (Skalar Medical). Raw coil signals were filtered in hardware (90 Hz low-pass butterworth), digitized (1000 Hz), and saved on computer for later analysis. Targets (a laser light of $\sim 1^\circ$ across) were rear-projected using a laser beam controlled by a galvanometer (with a 15° step-response of faster than 10 ms). The beam produced a 2-mm-diameter image on a translucent screen located 1 m in front of the subject. The room was otherwise dark. We stabilized the subject's head with a bite bar of dental impression material.

Experimental paradigms

Saccade adaptation was induced using the standard double-step paradigm (McLaughlin, 1967). After the adaptation experiment, we used the first saccade endpoint on each trial to form a sequence of targets to be shown to the same subject during a later control experiment. The idea of the control experiment was to produce a sequence of targets that allowed the subjects to reproduce saccades of the same amplitude that they made during the previous adaptation protocol, but without intrasaccadic target movement and, therefore, without any stimulus for adaptation. We could then compare the adapted saccades with the control saccades of the same amplitude to see whether saccade velocity and acceleration had changed. The procedures are described in detail below.

Group 1: gain-down. Subjects ($n = 5$) performed 480 gain-down trials equally separated into eight sets of 60 trials. After each set, they were asked to relax for 30 s and close their eyes. The sequence of endpoints that they produced in the adaptation experiment was used to generate the targets for the later control experiment (see below).

Group 2: gain-up. Subjects ($n = 5$) performed 480 gain-up trials equally separated into eight sets of 60 trials. As in group 1, they also were provided with rest periods between sets.

Adaptation trials. Trials began with a fixation target. At a random time in the range of 500–1000 ms, the fixation target was turned off and target T1 appeared at 15° along the horizontal axis. Once the eyes began moving toward T1, T1 was displaced to T2. The displacement was either 5° away from (gain-up) or 5° closer (gain-down) to the initial fixation. The jump was triggered as the eye crossed a virtual 2° window placed around the fixation point, corresponding to ~ 17 ms after saccade onset. Target T2 was kept visible for a random period in the range of 1000–1500 ms. Over time, subjects learned to make a saccade in response to T1 that was smaller (gain-down) or larger (gain-up). The mean intertrial time was 1250 ms.

Control trials. A number of factors can alter saccade velocity without affecting amplitude. For example, saccades to a constant target amplitude at the beginning of a session often have higher velocities than those at the end of the session (Fuchs and Binder, 1983; Chen-Harris et al., 2008; Golla et al., 2008). Here, we anticipated this problem and designed a task in which we could compare an adapted saccade of amplitude x on trial n with a control saccade of the same amplitude, measured on the same trial number, on the same subject, but in a different session.

Trials began with a fixation target. At a random time in the range of 500–1000 ms, the fixation target was turned off and target T1 was displayed at the endpoint of the saccade made by the same subject on the same trial in the adaptation experiment, offset by an amount that reflected the natural hypometria of that subject. If a_{ij} is the saccade end-

point during adaptation, x_{ij} represents target location for subject j on trial i , and h_{ij} is the expected amount of hypometria on that saccade, then $x_{ij} = (1 + h_{ij})a_{ij}$. To compute h_{ij} , we estimated the relationship between saccade amplitude and target eccentricity for each subject by using linear regression on a set of control data in which our subjects made saccades to a variety of target eccentricities in the range of 10 – 20° .

In the control trials, T1 remained at x_{ij} for 700 ms at which point the target moved to T2 (same T2 as in the adaptation experiment). This then became the fixation point for the next trial $i + 1$ in the opposite direction. Each control set consisted of 60 trials, with a mean intertrial interval of 1450 ms. The longer time spent at the final target (700 ms in control trials vs 500 ms in adaptation trials) was necessary to allow subjects in the adaptation paradigm enough time to perform a second, corrective saccade. Between sets, subjects received 30 s of rest.

Error-clamp trials. All experiments began with a baseline condition consisting of 180 error-clamp trials (Ethier et al., 2008). In these trials, the target disappears at saccade initiation and then reappears at the current fixation point after completion of the saccade. In an error-clamp trial, the trial began with a fixation target. At a random time in the range of 500–1000 ms, the fixation target was turned off and target T1 was displayed at 15° horizontal. Once the saccade began, T1 disappeared, and at 500 or 800 ms later (Groups 1 and 2, respectively) the target reappeared at the position where the eye was located at 10 ms before (i.e., 490 or 790 ms). We estimate that, in $>95\%$ of the trials, there were no saccades after this presentation of the target. This suggested that the location of the target coincided with the fixation of the eyes. Because of slight amplitude asymmetry between leftward and rightward saccades, occasionally a drift away from the center developed over time. We restrained eye position inside a $\pm 20^\circ$ range by resetting the fixation point to $\pm 5^\circ$ whenever the eye landed out-of-bounds. The mean intertrial time was 1250 and 1550 ms (for groups 1 and 2, respectively).

Data analysis

Saccade duration was determined by a $20^\circ/\text{s}$ speed threshold at onset and offset of the velocity profile. During the adaptation trials, saccades were rejected if they did not reach a peak velocity $>90^\circ/\text{s}$, if they had a latency <100 ms, if they displayed multiple peaks in their speed profile, or if they were shorter than 50% of the final target displacement. Most subjects had $<5\%$ of their saccades falling under one or more of these criteria, with no subjects exceeding 10% of all saccades. During the error-clamp trials, one subject in group 1 did not follow instructions and his data were removed.

Model

The objective of our modeling was to predict how saccade trajectories might change during adaptation. We considered two potential mechanisms: a remapping of the target (i.e., an explicit change in the goal of the saccade) and a change in the forward model that internally corrected saccade trajectories.

To generate saccade trajectories, we used the same optimal control framework previously used to model saccade trajectories in cross-axis adaptation (Chen-Harris et al., 2008). In the framework, the problem of controlling a saccade is formulated as follows: find the feedback control law that brings the eye to the target as soon as possible, with as little motor command as possible. The origin of this cost function goes back to the work of Harris (1995), in which he showed that saccade undershooting, a common observation in healthy people, suggests a minimization of saccade flight time in the presence of signal-dependent noise (i.e., noise in which the SD grows with the size of the signal). That seminal work, and many that have followed it (Harris and Wolpert, 1998; Harris and Wolpert, 2006), have solved this minimization problem in the context of open-loop control. However, saccades are not open-loop because they benefit from an internal feedback process. Until recently, we did not have a mathematical method to solve feedback control problems in the presence of signal-dependent noise. This changed with the work of Todorov (2005), who provided a closed-form solution for these problems. Here, we applied Todorov's method to a cost function that penalized distance of the fovea to the target, the motor commands of the saccade, and the “flight” time of the saccade. We express this cost at time t for time step Δt as follows:

$$J_t = \mathbf{x}_t^T Q \mathbf{x}_t + \lambda u_t^2 + 2\alpha t \Delta t. \quad (1)$$

In this equation, $\mathbf{x} := [x, \dot{x}, r]$ reflects the “state” of the system (i.e., the position of the eye, the velocity of the eye, and the position of the target). The term u is the motor command (a scalar quantity representing torque in one dimension). The terms in the matrix Q penalize the difference between eye and target, as well as saccade velocity, encouraging the system to bring the eyes to the target as soon as possible. The scalar λ penalizes motor commands (because noise in the motor commands grows with the size) (see below), discouraging large motor commands. Finally, the term α is simply a cost per unit time that penalizes movement duration, encouraging movements to terminate as soon as possible.

The objective is to find the motor commands that minimize the total cost of a saccade $J = \sum J_t$. Motor costs are a part of this total cost and are noted by $J_u = \lambda \sum u_t^2$, indicating that larger commands are penalized more. In principle, the specific trajectory (i.e., position, velocity, and acceleration as a function of time) that minimizes the total cost J depends on a number of factors. For example, if α is small (cost per unit time), saccades will have longer durations. If λ is small (cost of the motor commands), saccades will have shorter durations. The two parameters Q/λ and α are the only values that can be set in our cost function. They were set to the same values reported by Chen-Harris et al. (2008), so that for a broad range of saccade amplitudes, all resulting trajectories closely resembled measured data. We assumed that these parameters were constant for all saccade amplitudes. Furthermore, we assumed that, during adaptation, these parameters did not change.

To find the control system that could bring the eye to any target position while minimizing the total cost J , we considered the dynamics of the oculomotor plant as follows:

$$\mathbf{x}_{t+\Delta} = A \mathbf{x}_t + \mathbf{b}(u_t + \epsilon_t). \quad (2)$$

This equation is a discrete time representation of the oculomotor plant, with matrix A and vector \mathbf{b} having dimensions of 3×3 and 3×1 . We compute them as follows: we begin in continuous time with a second-order linear model of the eye plant with time constants of 224 and 13 ms (Robinson, 1986). In continuous time, A and \mathbf{b} are 2×2 and 2×1 . We then transform the equations to discrete time using a time step of 1 ms using exact solutions to the linear differential equations afforded by matrix exponentials, and then augment their dimensions to include state of the target.

In Equation 2, the term ϵ_t is a random variable that will describe the signal-dependent noise properties of the motor commands. To express this idea in a way that allows us to solve the optimal control problem, we write the following:

$$\epsilon = c|u_t| \phi_t. \quad (3)$$

Here, c is a scalar constant, $|u_t|$ indicates the absolute value of u , and ϕ_t is a normally distributed random variable with mean zero and variance 1 [i.e., $\phi_t \sim N(0,1)$]. Therefore, ϵ_t is a zero mean random variable with a SD that grows linearly as the size of motor command u , a property called signal-dependent noise. The values used for matrices A , \mathbf{b} , and for scalar c were the same as described previously (Chen-Harris et al., 2008).

The solution to our optimal control problem is a feedback controller described as follows:

$$u_t = -G_t \hat{\mathbf{x}}_t. \quad (4)$$

At each time point t during a saccade, a gain G_t is applied to estimate of state $\hat{\mathbf{x}}_t$ to produce motor command u_t [i.e., saccades are assumed to be under control of an internal-feedback loop that provides real-time estimate of state of the eye (via a forward model) (see below)]. The algorithm for finding matrix G is shown by Chen-Harris et al. (2008), and the derivation is available in the last author’s web-based course notes on learning theory, lectures on optimal control. Notice that Equation 4 uses an estimate $\hat{\mathbf{x}}_t$ of the state vector. The reason for this is that the saccade control system does not have access to the true state of the eye during a movement, and consequently has to form an estimate for it. This estimate comes from the forward model:

$$\hat{\mathbf{x}}_{t+1} = \hat{A} \hat{\mathbf{x}}_t + \hat{\mathbf{b}} u_t, \quad (5)$$

where \hat{A} and $\hat{\mathbf{b}}$ are the learner’s estimates of the parameters reflecting the physics of the motor plant. It is reasonable to assume that, in a normal context, the forward model is well calibrated and highly accurate, and so $\hat{A} = A$ and $\hat{\mathbf{b}} = \mathbf{b}$. However, when a sudden error is introduced through an intrasaccadic target step, the learner could assign this error to a failure in the forward model and make the appropriate adjustments to it, or the error could be assigned to an inadequate choice of target goal r (i.e., an incorrect initial motor command), and the learner might therefore modify it for the next trial (i.e., remap the target).

We suggest that this “credit assignment” problem is at the core of how saccades adapt to minimize the endpoint error. In terms of our model, the learner can respond to the error by either remapping the target, in which case $r' = r + \Delta r$, or by adjusting the forward model. The latter was modeled as a change in the effect of the efferent copy of the motor commands on the estimate $\hat{\mathbf{x}}$ of the state vector. The learner can modulate its estimate about the effects of the motor commands on the state. This would happen if, for example, the learner faces a consistent overshoot of the target and starts to believe that the motor commands that are being sent have a stronger effect on the dynamics of the eye than had been initially expected. In general, $\hat{\mathbf{b}}$ is equal to $(1 - \beta)\mathbf{b}$ and is thus parameterized by a factor β to either diminish or increase the impact of u_t in Equation 5.

Results

Saccades during adaptation

Subjects in group 1 participated in a gain-down adaptation paradigm. We recorded the saccade endpoints on each trial and then used the endpoints as target locations for a subsequent control session in the same subject. We hoped that in the control session we would observe saccades of similar amplitude as in the adaptation paradigm. This would then allow us to compare velocities and accelerations under the adapted and normal conditions for saccades of equal amplitude.

Figure 1A displays saccade properties during gain-down adaptation and control sessions. Saccade amplitudes declined with the start of the adaptation block but were well matched by the sequence of amplitudes in the control session. As noted before, the 30 s break between sets produced forgetting (Ethier et al., 2008) (i.e., a transient gain increase after each break). We reproduced this phenomenon and the other major features of adaptation in the control session.

Although saccade amplitudes were similar in the two sessions, trajectories were strikingly different. In general, a gain-down adapted saccade had a smaller peak velocity and a longer duration compared with its “normal” counterpart (i.e., saccade of the same amplitude made in the control session). To examine this difference, we considered the saccade trajectories in the last 40 trials of the last four sets across subjects. Figure 1, B and C, shows the average velocities and accelerations of these saccades (averaged across subjects). The adapted saccade has a smaller peak velocity, smaller acceleration, longer duration, and smaller deceleration than a control saccade of similar amplitude. (Statistical tests of these measures will be presented below.)

Subjects in group 2 participated in a gain-up adaptation paradigm. Saccade amplitudes increased in the adaptation block and were again well matched by the sequence of amplitudes in the control session (Fig. 2A). The 30 s break between sets produced robust forgetting (i.e., a transient gain decrease). However, unlike the gain-down case, here the saccade trajectories in the adaptation session were indistinguishable from the control session. That is, velocity and acceleration profiles of the adapted saccade were nearly identical with unadapted saccades of similar amplitude (Fig. 2B, C).

To perform statistical analysis on the data, we focused on saccade parameters during the last 40 trials of the last four sets, periods in which adaptation had reached a plateau and the transients caused by set breaks had disappeared. Figure 3 shows a summary of the results of the analysis. For the gain-down group, saccade amplitudes did not differ significantly ($F_{(4,1)} = 3.09$; $p > 0.15$) in the adapted and control sessions. Yet in the adapted condition the peak velocity was significantly smaller (242.0 vs 280.2°/s; $F_{(4,1)} = 8.98$; $p < 0.05$), duration was significantly longer (68.8 vs 55.9 ms; $F_{(4,1)} = 12.43$; $p < 0.05$), peak acceleration was significantly smaller (14,800 vs 18,400°/s²; $F_{(4,1)} = 9.74$; $p < 0.05$), and peak deceleration was significantly smaller (−14,600 vs −10,700°/s²; $F_{(4,1)} = 7.90$; $p < 0.05$). Therefore, the gain-down adapted saccades had trajectories that significantly differed from normal saccades. In contrast, in the gain-up paradigm, there were no statistical differences across conditions for any of the saccade parameters. That is, gain-up adapted saccades had trajectories that were indistinguishable from normal saccades.

Saccade latencies

Saccade latency, defined as the time between target presentation and saccade start, was strongly affected by adaptation (Fig. 4). We measured latencies (Fig. 4C) for the same trials that we had computed saccade parameters in the above analysis. In the gain-down group, adapted saccades took much longer to start than normal saccades (214.3 vs 162.9 ms; $F_{(4,1)} = 84.77$; $p < 0.001$). In the gain-up group, the latencies were less affected, but adaptation still introduced a significant increase in the reaction times (203.2 vs 185.3 ms; $F_{(4,1)} = 13.23$; $p < 0.05$). However, this last result is confounded by the fact that, in the gain-up paradigm, latencies were not only higher during adaptation, but they were also higher during baseline (i.e., trials 1–180) (Fig. 4B). Therefore, within-subject variability is likely the cause of the latency differences in the gain-up condition, not adaptation to endpoint errors.

In summary, adaptation in the gain-down condition produced saccades that were “abnormal” in two different ways. (1) The trajectory of the saccade was abnormal. The adapted saccades had reduced peak velocities, reduced accelerations, shallower decelerations, and increased durations with respect to a control saccade of equal amplitude. (2) The reaction times were abnormally long. However, adaptation in the gain-up condition produced saccades that were indistinguishable with respect to control saccades of equal amplitude.

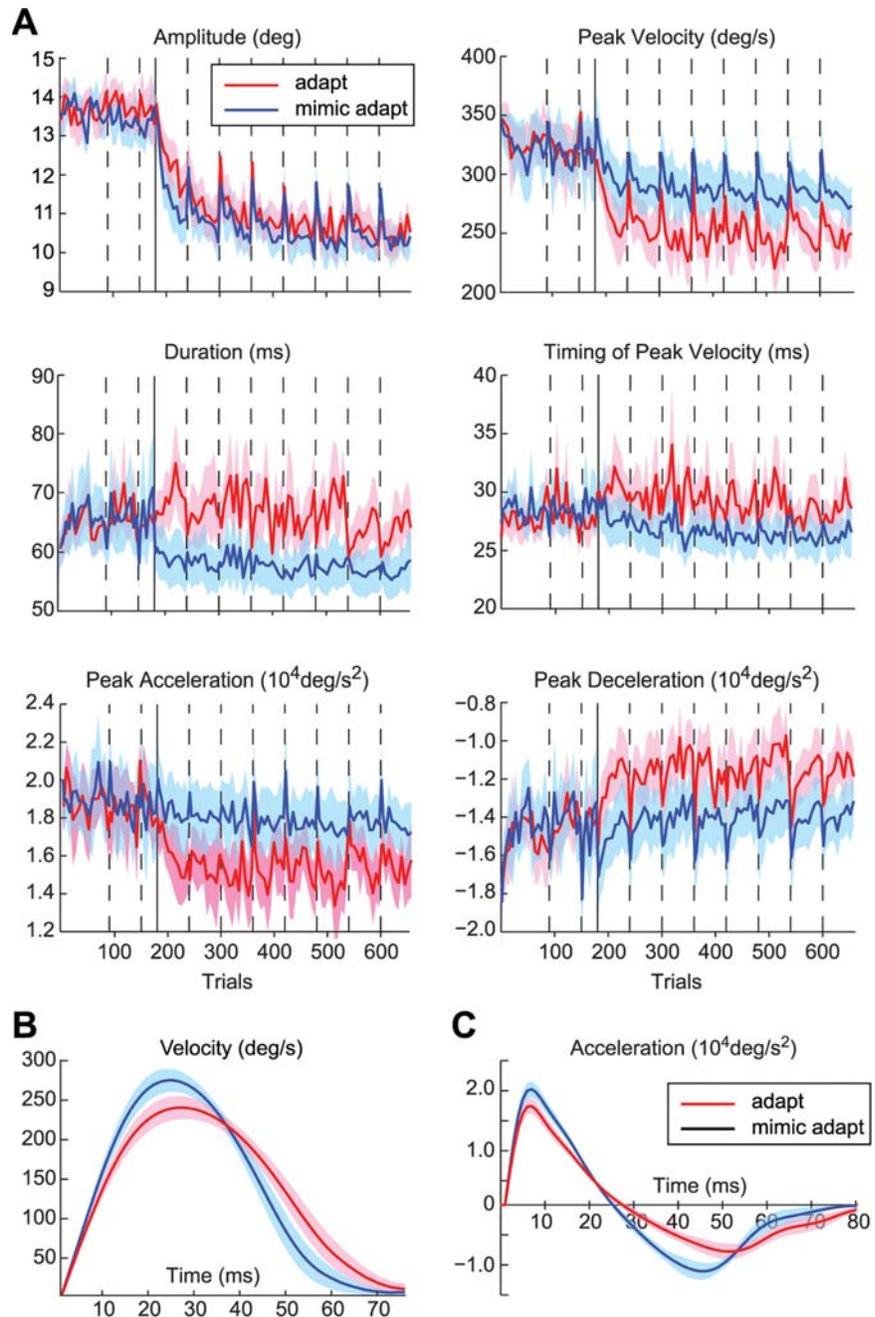


Figure 1. Group average, trial-by-trial saccade kinematics for gain-down adaptation versus corresponding control. The dotted lines represent set breaks, and the solid line divides the baseline (error-clamp) trials from the adaptation/control trials. In red are the mean adaptation kinematics and in black the mean control ($n = 5$). Error bars are SEM. **A**, Saccade amplitude, peak velocity, duration, time with respect to saccade onset at which the velocity reaches its peak, peak deceleration, and peak acceleration. To highlight the within-set structure of the data, bin size was two trials for the first bin, and then four trials, and then six trials for all subsequent bins for each set. **B**, Group average speed profile. This represents the average kinematics for the last 40 trials of each of the last four sets. **C**, Group average acceleration profile.

Inferring the mechanisms of adaptation

What do the specific changes in saccade trajectories in the gain-down paradigm tell us about the adaptive control process? Why should the brain choose a different mechanism of adaptation when it undershoots the target (gain-up) versus when it overshoots (gain-down)?

To explore these questions, we used an optimal control framework to simulate saccades. We considered two hypothetical mechanisms of adaptation: remapping of the target versus

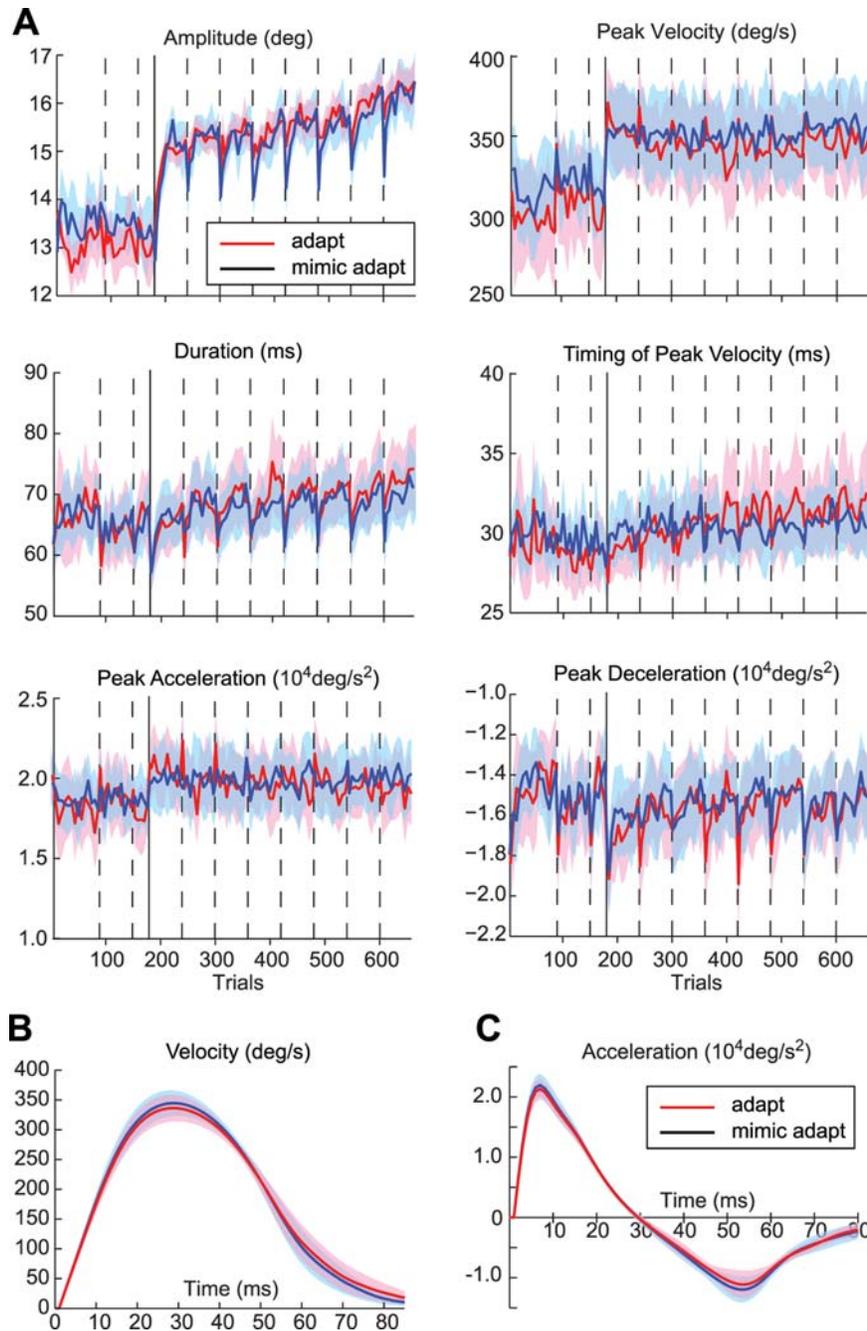


Figure 2. Group average, trial-by-trial saccade kinematics for gain-up adaptation versus corresponding control. The format is same as in Figure 1. **A**, Saccade amplitude, peak velocity, duration, timing at peak velocity, peak deceleration, and peak acceleration. **B**, Group average speed profile. This represents the average kinematics for the last 40 trials of each of the last four sets. **C**, Group average acceleration profile.

changes in an internal feedback that acted as a forward model to internally monitor the motor commands and estimate the current state of the eyes. In principle, remapping of the target should produce saccades that look indistinguishable from control, whereas a change in the internal feedback will make the saccades look abnormal. Our interest was to ask two questions: whether a change in the internal feedback would in fact produce trajectory changes that resembled what we had seen in the gain-down condition; and if so, why should the brain choose remapping in one condition and a change in the internal feedback in the other.

Before adaptation, a 15° horizontal target resulted in a saccade

with a velocity profile indicated by the black line in Figure 5A (control). Let us suppose that, in the gain-down paradigm, we need to produce a 10° saccade in response to a 15° target. To produce this 10° saccade, one can either remap the 15° target to 10° , or produce a 10° saccade through changes in the internal feedback. When we simulated a 10° saccade via target remapping, the duration and peak velocity were, of course, reduced compared with a 15° saccade (Fig. 5A, blue line). This is the trajectory expected for the control condition because remapping is tantamount to displaying a target at a smaller amplitude. However, in response to a 15° target, one can also produce a 10° saccade by changing the forward model (Eq. 5) (i.e., the internal feedback that predicts how the motor commands change the state of the eye). Indeed, when we simulated a 10° saccade by changing the forward model, the resulting saccade (Fig. 5A, red line) had reduced peak velocities, reduced accelerations, and increased durations with respect to target remapping. That is, changing the forward model achieved the same amplitude as did remapping, but with a different trajectory.

Therefore, the simulation explained that if gain-down adaptation was via a forward model, saccades would have a smaller peak velocity but longer duration with respect to control saccades of the same amplitude. It also explained that, with respect to baseline saccades (i.e., 15° saccades before start of adaptation), the adapted saccades would not show a decreased duration despite the fact that the amplitudes were reduced. These changes were precisely what we observed in the experimental data.

Why should one adapt the forward model in the gain-down paradigm and not remap the target? The model provides a possible rationale. Recall that we assume that there is a motor cost $J_u = \lambda \sum_t u_t^2$ associated with saccades (where u_t is the motor command or torque at time t). This cost is not related to energy expenditure. Rather, it is associated with our assumption that noise in the motor command grows with its size, making it a good idea to directly

penalize large motor commands. Previous works on other adaptation paradigms have demonstrated that, in response to a perturbation, movement trajectories in general do not return to a baseline condition, but rather the adapted trajectory is one that reduces endpoint errors while minimizing motor costs (Emken et al., 2007; Izawa et al., 2008).

The motor commands (torques) that are needed to produce the various eye movements are plotted in Figure 5B. The motor cost $\sum u_t^2$ is plotted as a function of time in Figure 5C. We see that, in the gain-down condition, making a 10° saccade with a remapped target has a greater motor cost than making the same

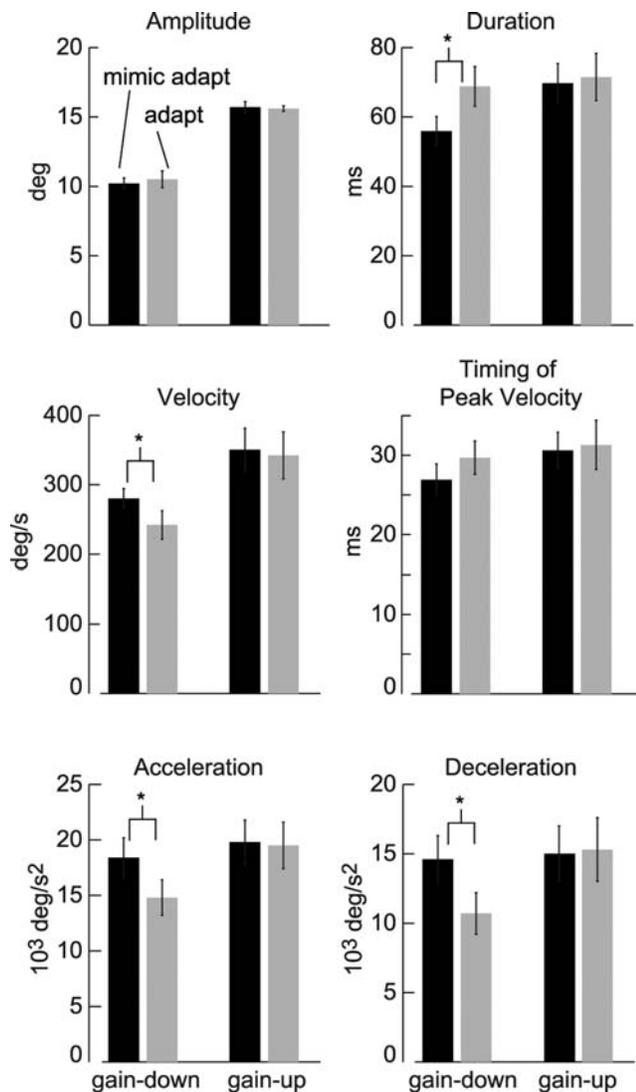


Figure 3. Comparison of saccade parameters in the adaptation versus control condition. Data represent the last 40 trials of the last four sets. Error bars are SEM.

amplitude saccade via a change in the forward model. Therefore, in the gain-down condition, reducing endpoint errors by changing the internal feedback may be a superior policy with respect to target remapping in the sense that it produces saccades that have a smaller motor cost.

It is important, however, to point out that lower motor costs were achieved by stretching the saccade in time compared with a “normal” saccade of equal amplitude. The total cost (Eq. 1) includes terms that penalize saccade duration. Therefore, if we assume that the “value” of time (the term α in Eq. 1) has not changed with respect to preadaptation, these gain-down adapted saccades are in fact suboptimal with respect to normal saccades. This assumption predicts that with additional practice, adapted gain-down saccade trajectories will converge to normal saccades.

Let us now consider the gain-up paradigm. Suppose that, in response to a 15° target, we need to move our eyes to 20°. If we simulate this larger amplitude by remapping a 15° target to 20°, the resulting 20° saccade (Fig. 5A, blue line, gain-up) has, of course, an increased duration and an increased peak velocity with respect to control 15° saccade (Fig. 5A, black line). However, if we simulate this same increased amplitude via a change in the for-

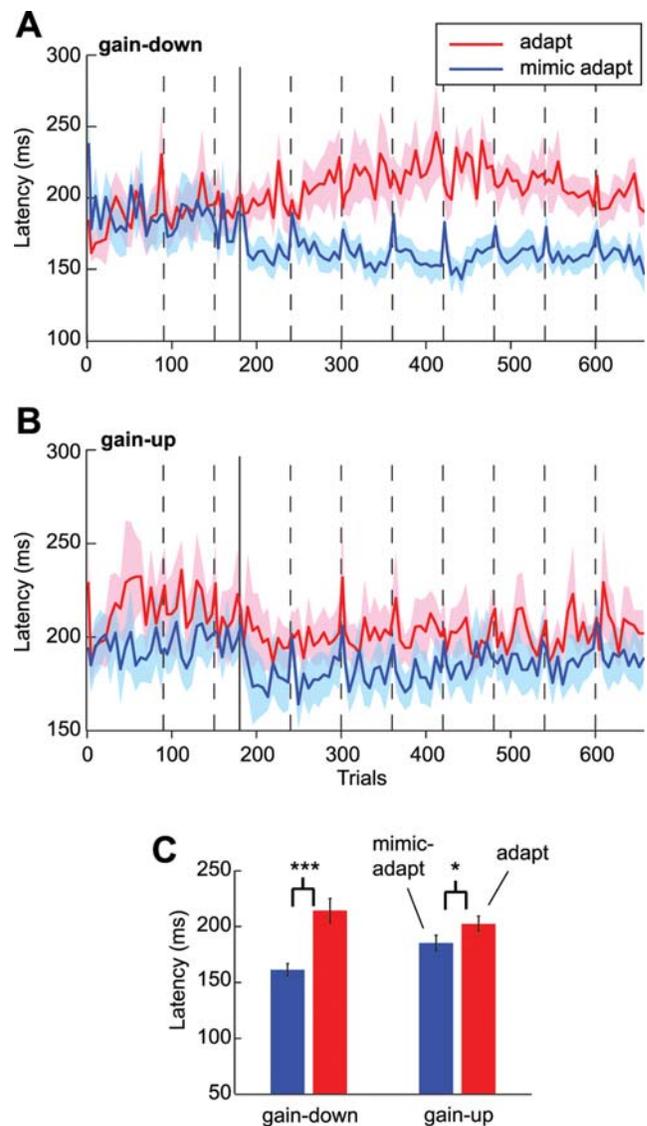


Figure 4. Saccade latencies. The format is the same as in Figure 1. Error bars are SEM. **A, B**, Saccade latencies in the gain-down and gain-up paradigms. The dotted lines represent set breaks, and the solid line marks the end of error-clamp trials. **C**, Data represent the last 40 trials of each of the last four sets. Error bars are SEM. Both adaptation conditions produced significantly longer latencies than the control condition, although the difference was smaller in gain-up.

ward model, the resulting trajectory (Fig. 5A, red line) exhibits increased peak velocities but reduced durations with respect to target remap. Therefore, in the gain-up paradigm if the adaptation is via a forward model, saccades should have a larger peak velocity but shorter duration with respect to control saccades of the same amplitude. However, the remapping of the target is a better policy for gain-up adaptation, because it produces movements that have the least motor costs (Fig. 5C). Indeed, we found that the gain-up adapted saccades were indistinguishable from saccades of equal amplitude during control.

In summary, the model suggested that whereas in the gain-down paradigm adaptation was via changes in a forward model that internally monitored and corrected the eye's movements, in the gain-up paradigm adaptation was via a process akin to target remapping. In both paradigms, minimizing saccade motor costs appeared to play a significant role in the process of adaptation.

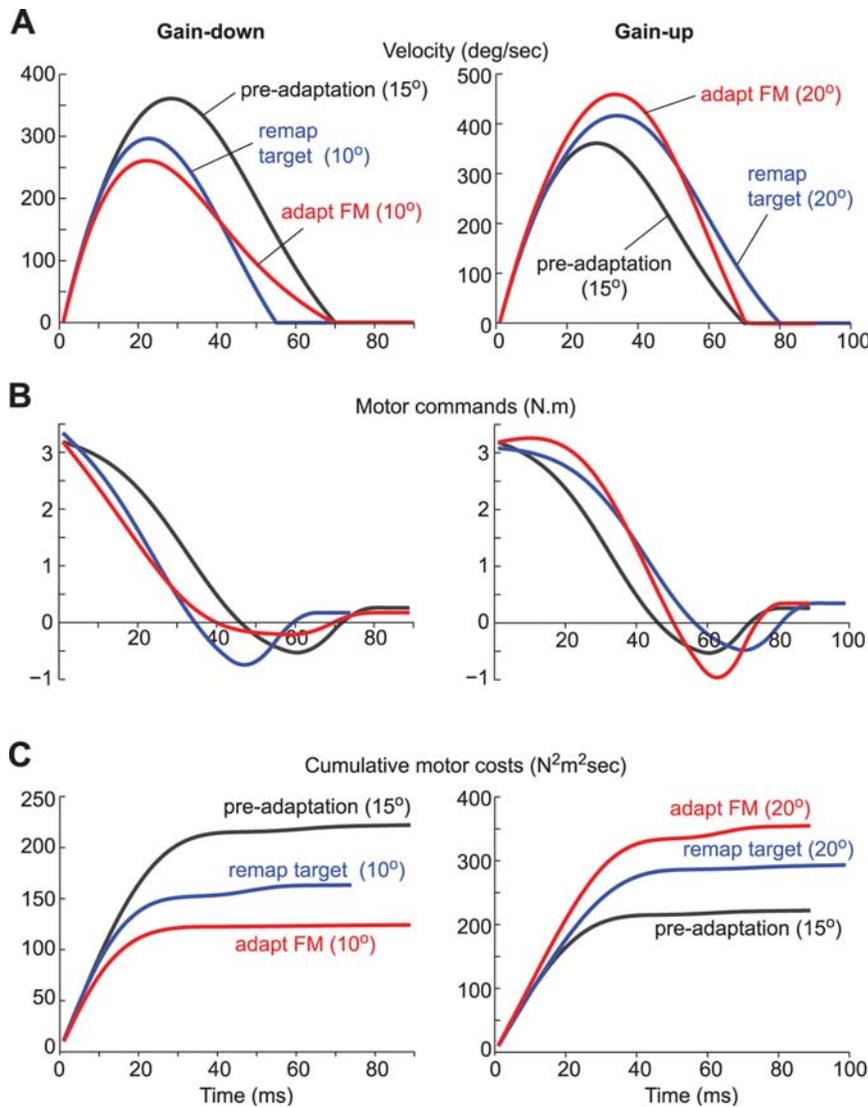


Figure 5. Simulation of adaptive control of saccades. **A**, Gain-down paradigm. The black trace is the speed profile of a 15° control saccade in a baseline condition. The blue trace is the saccade to a remapped target at 10° [i.e., at saccade onset the target is remapped by -5° ($\Delta r = -5^\circ$)]. This is the saccade profile expected in the control condition of group 1. The green trace is the speed profile of a saccade in which the forward model had adapted, producing a 10° saccade in response to a 15° target ($\Delta r = 0^\circ$; $\beta = -0.5$). Gain-up paradigm. The black trace is the speed profile of a 15° control saccade in preadaptation condition. The blue trace is the saccade to a remapped target at 20° ($\Delta r = +5^\circ$; $\beta = 0$). This is the profile expected in the control condition of group 2. The green trace is the speed profile of a saccade in which the forward model had adapted, producing a 20° saccade in response to a 15° target ($\Delta r = 0^\circ$; $\beta = 0.25$). **B**, Corresponding profiles of the motor commands u_t (torque) for the saccades shown in **A**. Units are newton \cdot meters. **C**, Sum of the squared motor commands as a function of time: $\sum_t u_t^2$. Forward model adaptation is the least costly policy to reduce saccade error in a gain-down paradigm. Target remapping is the least costly policy to reduce saccadic error in a gain-up paradigm.

Discussion

Saccadic gain change is a widely studied model of motor learning. In a typical paradigm, the target moves intrasaccadically either toward (gain-down) or away (gain-up) from initial fixation. The endpoint error (Wallman and Fuchs, 1998) affects the way that the brain controls the subsequent saccade (Srimal et al., 2008). A central question is how the motor control system changes in response to movement errors. We considered two potential mechanisms of adaptation: a remapping of the target versus a change in the internal feedback that monitors motor commands and corrects them on-line (i.e., a forward model). In the gain-down paradigm, adapted saccades had reduced peak velocities,

reduced accelerations, shallower decelerations, and increased durations with respect to a normal saccade of equal amplitude. In the gain-up paradigm, however, adapted saccades were indistinguishable from normal saccades. Using an optimal control model of saccades, we found that the changes observed in the gain-down adapted saccades were consistent with changes in a forward model, whereas in the gain-up condition, the process of adaptation was similar to target remapping.

Why should gain-down engage a different adaptation mechanism than gain-up? During adaptation of reaching movements, motor costs appear to have a strong influence on the specific trajectory that the brain chooses for executing the movement: the adapted trajectory is generally one that minimizes the total motor costs of arriving at the target (Emken et al., 2007; Izawa et al., 2008). Here, our simulations showed that, in the gain-down paradigm, a change in the internal feedback was a superior adaptation strategy because it produced saccades that had smaller motor costs than remapping. However, in the gain-up paradigm, remapping had smaller motor costs.

Finally, we found that regardless of the adaptation paradigm, saccade latencies were significantly longer for the adapted saccade compared with a control saccade. This aspect of the movement was not modeled, but, as we argue below, may also be understood in the optimal control framework.

Saccade trajectories in previous adaptation studies

Many studies have examined changes in saccade trajectories (i.e., position, velocity, and acceleration) during gain adaptation. Yet the results have not painted a clear picture. For example, Collins et al. (2008) reported that gain-down adaptation in humans resulted in longer durations and shallower decelerations compared with control saccades of similar amplitude. In contrast, Straube and Deubel (1995) as well as Alahyane and Pélisson (2005)

found no significant changes with respect to control saccades. In monkeys, Fitzgibbon et al. (1986) reported an increase in peak velocity accompanying gain-down adaptation, whereas Frens and Van Opstal (1997) observed a reduction in peak velocities and an increase in durations. In the gain-up paradigm, Straube and Deubel (1995) found that durations were longer and peak velocity was reduced with respect to control saccades. Whereas the results of Collins et al. (2008) and Frens and Van Opstal (1997) are consistent with our observations, results of Alahyane and Pélisson (2005), Straube and Deubel (1995), and Fitzgibbon et al. (1986) are not.

Why is there such disagreement? Many factors can alter sac-

cade velocities without affecting their amplitudes. For example, saccades to a constant target amplitude at the beginning of a session often have higher velocities than those at the end of the session. The reason behind this effect is unknown, but it has been observed in both humans (Fuchs and Binder, 1983; Chen-Harris et al., 2008; Golla et al., 2008) and monkeys (Straube et al., 1997; Barash et al., 1999). Here, we anticipated this problem and designed a task in which we could compare an adapted saccade of amplitude x on trial n with a control saccade of the same amplitude, measured on the same trial number, on the same subject, but in a different session. These “control” saccades accounted for both trial-dependent and amplitude-dependent changes in saccade trajectories.

A disadvantage of our approach was that it required us to first measure the adapted saccades and then subsequently measure the control saccades in a following session. This was required as the sequence of control targets had to be tailored to match the sequence of saccade amplitudes observed for each subject in the adaptation session. Therefore, a session order effect confounded our interpretations. We anticipated this and included baseline trials for both sessions and generally observed similar saccade characteristics.

Different adaptive control mechanisms in gain-down versus gain-up

Our results suggest that, whereas gain-down adaptation is mainly through changes in the internal feedback mechanism that steers saccades midflight, gain-up adaptation is primarily through a remapping of the target. There is evidence for this dissociation from other reports.

If mechanisms of adaptation differ for gain-up versus gain-down, then one might expect to find asymmetric patterns of generalization to untrained targets. For example, Semmlow et al. (1989) found that the generalization field of gain-down was constant in terms of percentage change in the gain of the saccade over the neighboring amplitudes. However, for gain-up, the absolute displacement of the endpoint appeared to be constant for neighboring amplitudes. The asymmetric generalization suggested that adaptation in the gain-up (but not gain-down) paradigm was consistent with target remapping (Semmlow et al., 1989; Wallman and Fuchs, 1998) (but see Noto et al., 1999). These conclusions are consistent with our results.

The superior colliculus and the cerebellum are the major components of the saccadic circuitry likely to be undergoing plasticity during saccade adaptation. There are neurons in the superior colliculus (Takeichi et al., 2007), the nucleus reticularis tegmenti pontis (Takeichi et al., 2005), the fastigial oculomotor region (Inaba et al., 2003; Scudder and McGee, 2003), and the oculomotor vermis (Soetedjo and Fuchs, 2006; Catz et al., 2008) that change their firing pattern during adaptation. It is difficult to pinpoint a single locus where adaptation occurs. However, if the neural substrates of the two mechanisms of adaptation that we proposed are physically separate, it should be possible to inactivate that region and observe adaptation in one paradigm but not the other.

Changes in saccade latency

In a sense, our adapted saccades were always “abnormal” because they tended to exhibit long latencies. Why should saccades in an adaptation paradigm be abnormally slow to initiate? There may be two factors that contribute to this latency effect.

First, Harwood et al. (2008) have suggested that the probability distribution of saccade latency is a function of two costs: the

cost of not foveating the target (e.g., lack of information because of low resolution in the peripheral vision) and the cost of initiating the saccade (e.g., risk of losing track of the target if it suddenly moves while the saccade is underway). That is, the latency of a saccade may itself be a cost function that is being optimized by the brain. This would be consistent with the observation that saccade latencies are smaller for targets of higher value. For example, monkeys have shorter saccade latencies to targets that are rewarded with food (Takikawa et al., 2002). In an adaptation paradigm, the target jumps during the saccade. It is possible that this affects an intrinsic value that the brain might assign to the target. Another possibility, however, is that a jumping target increases the uncertainty associated with its position. We have observed that, in starting a reaching movement, people exhibit systematic increases in reach latencies as target position uncertainty increases (Izawa and Shadmehr, 2008). It is possible that the increased saccade latencies are related to increased uncertainty associated with location of the jumping target and/or the reduced expected value of the movement to that target.

Second, target eccentricity affects saccade latencies: generally, the farther the target, the longer it takes to initiate a saccade (Kalesnykas and Hallett, 1994). For example, we would expect approximately a 1–2 ms increase in latency with a 1° increase in eccentricity (Kalesnykas and Hallett, 1994). In the gain-down paradigm, we conjectured that the mechanism of adaptation was through internal feedback, and not target remapping. That is, in the adaptation trials, the perceived visual stimulus is farther away than in the control trials, which in turn would explain some of the longer latency associated with adapted saccades. The contribution of this effect, however, is at most 10 ms, whereas our data showed latency changes that were ~50 ms. In the gain-up paradigm, we conjectured that the mechanism of adaptation was target remapping. Therefore, based on this principle, there should be no difference in latencies with respect to control in the gain-up paradigm, but an increase in the gain-down paradigm. Whereas we saw a small increase in latency of adapted saccades in the gain-up condition that may have been unrelated to adaptation, we saw a large and robust increase in latencies in the gain-down condition.

Our speculation is that in the gain-down paradigm, both the first and second factors were in play, strongly increasing the latencies of adapted saccades. However, in the gain-up paradigm, only the first factor was in play, resulting in a weaker increase in saccade latencies.

Model limitations

The fundamental assumption of optimal control is that the brain assigns value to stimuli, performs actions toward stimuli that happen to have the largest expected value, and tries to execute movements with trajectories that maximize the expected reward while minimizing motor costs (Todorov and Jordan, 2002; Shadmehr and Krakauer, 2008). Although there is neurophysiological evidence for reward-related signals in the brain (Schultz et al., 1997), we do not know of any evidence for signals that quantify “motor costs” of a movement. Indeed, we do not know the nature of reward functions or the motor cost functions in paradigms such as the one studied here. The quadratic terms in Equation 1 (i.e., our estimate of these functions) are mathematical conveniences that have thus far proven useful predictors of behavior in this and other motor tasks (Liu and Todorov, 2007; Izawa and Shadmehr, 2008; Izawa et al., 2008). Without more knowledge about these functions (Körding et al., 2004), it is not possible to

account for the mechanisms of adaptation of movement trajectories.

For example, the changes in saccade latencies in our study suggest that, in the gain-down condition, the value of the target might be substantially lower than the same target in the control condition. A reduced expected value would produce saccades that are longer in duration and lower in velocity. The resulting saccade would be indistinguishable from one in which the internal feedback had adapted, strongly challenging our conclusions here. Yet, whatever reductions in stimulus value might have occurred during the gain-down paradigm, they probably also occurred during gain-up. The fact that we did not see equivalent trajectory changes in the two conditions with respect to control is some assurance about the robustness of our inference.

References

- Alahyane N, Pélission D (2005) Long-lasting modifications of saccadic eye movements following adaptation induced in the double-step target paradigm. *Learn Mem* 12:433–443.
- Barash S, Melikyan A, Sivakov A, Zhang M, Glickstein M, Thier P (1999) Saccadic dysmetria and adaptation after lesions of the cerebellar cortex. *J Neurosci* 19:10931–10939.
- Catz N, Dicke PW, Thier P (2008) Cerebellar-dependent motor learning is based on pruning a Purkinje cell population response. *Proc Natl Acad Sci U S A* 105:7309–7314.
- Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R (2008) Adaptive control of saccades via internal feedback. *J Neurosci* 28:2804–2813.
- Collins T, Semroud A, Orriols E, Dore-Mazars K (2008) Saccade dynamics before, during and after saccade adaptation in humans. *Invest Ophthalmol Vis Sci* 49:604–612.
- Emken JL, Benitez R, Sideris A, Bobrow JE, Reinkensmeyer DJ (2007) Motor adaptation as a greedy optimization of error and effort. *J Neurophysiol* 97:3997–4006.
- Ethier V, Zee DS, Shadmehr R (2008) Spontaneous recovery of motor memory during saccade adaptation. *J Neurophysiol* 99:2577–2583.
- Fitzgibbon EJ, Goldberg ME, Segraves MA (1986) Short-term saccadic adaptation in the monkey. In: *Adaptive processes in the visual and oculomotor system* (Keller EL, Zee DS, eds), pp 329–333. New York: Pergamon.
- Frens MA, Van Opstal AJ (1997) Monkey superior colliculus activity during short-term saccadic adaptation. *Brain Res Bull* 43:473–483.
- Fuchs AF, Binder MD (1983) Fatigue resistance of human extraocular muscles. *J Neurophysiol* 49:28–34.
- Golla H, Tziridis K, Haarmer T, Catz N, Barash S, Thier P (2008) Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. *Eur J Neurosci* 27:132–144.
- Harris CM (1995) Does saccadic undershoot minimize saccadic flight-time? A Monte-Carlo study. *Vision Res* 35:691–701.
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784.
- Harris CM, Wolpert DM (2006) The main sequence of saccades optimizes speed-accuracy trade-off. *Biol Cybern* 95:21–29.
- Harwood MR, Madelain L, Krauzlis RJ, Wallman J (2008) The spatial scale of attention strongly modulates saccade latencies. *J Neurophysiol* 99:1743–1757.
- Inaba N, Iwamoto Y, Yoshida K (2003) Changes in cerebellar fastigial burst activity related to saccadic gain adaptation in the monkey. *Neurosci Res* 46:359–368.
- Izawa J, Shadmehr R (2008) On-line processing of uncertain information in visuomotor control. *J Neurosci* 28:11360–11368.
- Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. *J Neurosci* 28:2883–2891.
- Jürgens R, Becker W, Kornhuber HH (1981) Natural and drug-induced variations of velocity and duration of human saccadic eye movements: evidence for a control of the neural pulse generator by local feedback. *Biol Cybern* 39:87–96.
- Kalesnykas RP, Hallett PE (1994) Retinal eccentricity and the latency of eye saccades. *Vision Res* 34:517–531.
- Körding KP, Fukunaga I, Howard IS, Ingram JN, Wolpert DM (2004) A neuroeconomics approach to inferring utility functions in sensorimotor control. *PLoS Biol* 2:e330.
- Liu D, Todorov E (2007) Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27:9354–9368.
- McLaughlin S (1967) Parametric adjustment in saccadic eye movements. *Percept Psychophys* 2:359–362.
- Noto CT, Watanabe S, Fuchs AF (1999) Characteristics of simian adaptation fields produced by behavioral changes in saccade size and direction. *J Neurophysiol* 81:2798–2813.
- Quaia C, Paré M, Wurtz RH, Optican LM (2000) Extent of compensation for variations in monkey saccadic eye movements. *Exp Brain Res* 132:39–51.
- Robinson DA (1963) A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans Biomed Eng* 10:137–145.
- Robinson DA (1986) The systems approach to the oculomotor system. *Vision Res* 26:91–99.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593–1599.
- Scudder CA, McGee DM (2003) Adaptive modification of saccade size produces correlated changes in the discharges of fastigial nucleus neurons. *J Neurophysiol* 90:1011–1026.
- Semmlow JL, Gauthier GM, Vercher JL (1989) Mechanisms of short-term saccadic adaptation. *J Exp Psychol Hum Percept Perform* 15:249–258.
- Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. *Exp Brain Res* 185:359–381.
- Snyder LH, Calton JL, Dickinson AR, Lawrence BM (2002) Eye-hand coordination: saccades are faster when accompanied by a coordinated arm movement. *J Neurophysiol* 87:2279–2286.
- Soetedjo R, Fuchs AF (2006) Complex spike activity of Purkinje cells in the oculomotor vermis during behavioral adaptation of monkey saccades. *J Neurosci* 26:7741–7755.
- Srimal R, Diedrichsen J, Ryklin EB, Curtis CE (2008) Obligatory adaptation of saccade gains. *J Neurophysiol* 99:1554–1558.
- Straube A, Deubel H (1995) Rapid gain adaptation affects the dynamics of saccadic eye movements in humans. *Vision Res* 35:3451–3458.
- Straube A, Fuchs AF, Usher S, Robinson FR (1997) Characteristics of saccadic gain adaptation in rhesus macaques. *J Neurophysiol* 77:874–895.
- Takeichi N, Kaneko CR, Fuchs AF (2005) Discharge of monkey nucleus reticularis tegmenti pontis neurons changes during saccade adaptation. *J Neurophysiol* 94:1938–1951.
- Takeichi N, Kaneko CR, Fuchs AF (2007) Activity changes in monkey superior colliculus during saccade adaptation. *J Neurophysiol* 97:4096–4107.
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291.
- Todorov E (2005) Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput* 17:1084–1108.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.
- van Donkelaar P, Siu KC, Walterschied J (2004) Saccadic output is influenced by limb kinetics during eye-hand coordination. *J Mot Behav* 36:245–252.
- Wallman J, Fuchs AF (1998) Saccadic gain modification: visual error drives motor adaptation. *J Neurophysiol* 80:2405–2416.