ADAPTIVE PLASTICITY IN THE HUMAN SACCADE SYSTEM

by

JAMES P. HERMAN

A dissertation submitted to the Graduate Faculty in Biology in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2012

This manuscript has been read and accepted by the
Graduate Faculty in Biology in satisfaction of the
dissertation requirement for the degree of Doctor of Philosophy.

Jonathan B. Levitt

Date

Chair of Examining Committee

Laurel Eckhardt

Date

Executive Officer

Jay A. Edelman, CCNY
Ofer Tchernichovsky, CCNY
Lance M. Optican, National Eye Institute, NIH
David S. Zee, Johns Hopkins Medical School
Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

ADAPTIVE PLASTICITY IN THE HUMAN SACCADE SYSTEM

by

JAMES P. HERMAN

Advisor: Professor Josh Wallman (Jonathan B. Levitt, Chair of Examining Committee)

The rapid point-to-point movements of the eyes called saccades are some of the most commonly made by humans, yet differ from nearly every other type of motor output in that they are completed too quickly to be adjusted during their execution by visual feedback. Yet, saccadic accuracy remains quite high over a lifetime despite inevitable changes to the physical structures controlling the eyes, indicating that the oculomotor system actively monitors and adjusts motor commands to achieve this consistent behavioral production. Indeed, it seems that beyond the ability to compensate for slow, age-related bodily changes, saccades can be modified following traumatic injury or pathology that affects their production, or in response to more short-term systematic alterations to post-saccadic visual feedback in a lab setting. It is, in fact, thought that all of these forms of plasticity rely on the visual detection of accuracy errors by a unified set of mechanisms that support the process known as saccade adaptation. A great deal has been learned about saccade adaptation, as it has been extensively studied as a phenomenon in its own right, as well as being used to explore the process of motor learning in general. However, many fundamental questions about human saccade adaptation remain unanswered, often related to the way that saccade adaptation might operate in the natural environment with substantially more complex visual stimuli than are generally used in the lab. Here, we addressed these questions with (in some cases original) variants and more conventional examples of the frequently used intrasaccadic target step (ISS) paradigm (in which an experimenter causes saccadic error by shifting a target during the movement). By exploring the responses to whole-field ISSs, we have inferred that saccade adaptation might be supported by a trans-saccadic integration mechanism, and may be sensitive to intrasaccadic motion signals. Challenging the oculomotor system by confronting it with multiple post-saccadic targets has revealed that saccade adaptation can occur in a target-identity specific manner, so that even if post-saccadic error varies from trial-to-trial, adaptation seems to reflect the average behavior of the target. At a more basic level, we systematically varied ISSs to determine the lower limits of the oculomotor system's sensitivity to intrasaccadic displacement during adaptation. Also at a more basic level, we looked at the effects of rendering post-saccadic feedback more intermittent during adaptation, finding it to have little effect on the magnitude or rate of adaptive dynamics, similar to other forms of motor learning but somewhat dissimilar from operant conditioning. These experiments also furnished a useful setting to develop and test a novel model of saccade adaptation which more explicitly relies on post-saccadic sensory prediction than previous models, but that is nonetheless in keeping with the ethos of modern motor learning theory. Finally, we found that the establishment and maintenance of a context for saccadic performance by adaptation could be achieved by consistently pairing a target's visual-identity with a specific ISS, extending what had been previously recognized as constituting a cue for contextual motor learning. In general, our results suggest that saccade adaptation is a highly flexible mechanism that not only supports the maintenance of accuracy, but also makes use of a wide range of brain functions to deftly tailor saccadic behavior contingent on task demands.

Acknowledgements

First and foremost, I thank Josh Wallman, my mentor, advisor, enthusiastic supporter, advocate, cheerleader, fellow playful explorer, a source of continual inspiration, and some – productive – frustration. It was a combination of luck and fate that crossed our paths, and I consider myself more than blessed to have had the privilege of sharing the road with such a great man for any period. It saddens me deeply that we can no longer revel in our shared love of life’s great pleasures. He is and always will be a part of who I am; I would not be the man I am today without him. I aspire to the sense of excitement and wonder, the twinkle in the eye, and the rigorous approach to doing science every day that he has shown me. I hope that what I have done here, and what I continue to work on every day would have made him proud.

I am deeply indebted to Annabelle Blangero, and Mark R. Harwood. Mark has been there since day one of my graduate research career, never once refusing to offer his assistance, even just to hold my hand as I solved a problem on my own by talking it through with him; to say nothing of his encyclopedic knowledge of eye movement literature which I relied upon heavily. I know of no more erudite scientist anywhere, a quality I hold in high regard and hope to someday acquire a small fraction of. From the moment she entered the Wallman Lab, Annabelle has been my ardent supporter, defender, and an animated source of encouragement. There is nobody else in my life that both possesses the will and capacity to dig into my work, and has championed it more loudly than she. Mark and Annabelle stepped in as chief thesis editors and advisors without forethought, and kept me going at times when my own determination flagged. This document simply would not exist without their tireless effort.

I thank Laurent Madelain for his invaluable advice and friendship. Without his efforts propelling me forwards and guiding my scientific, professional, and personal paths, I am certain I would be lost. He provided counsel for which I had no source, and for which there is no substitute.

I am extremely grateful to the members of my examining committee, Jonathan Levitt, Jay Edelman, Ofer Tchernichovsky, Lance Optican, and David Zee for welcoming me into the academy as a colleague. Their astute advice, keen observations, and friendly, supportive manner have reaffirmed my love for my chosen profession and helped make my step sure in continuing down this avenue.

I owe many thanks to all the members of the Wallman Lab for cultivating a familial work environment, to say nothing of their willingness to be my “guinea pigs” in countless experiments. They have set a peerless standard, and I am lucky to count them amongst my friends and colleagues.

Two people I cannot thank enough are my parents, William and Joanna Clapps Herman. There are simply no words to describe the kind of support they’ve provided me. They are my models for being, and more than any other, those whom I hope to become. They are the source of my love of learning, the greatest gift I have ever been given and the foundation of my joy. They have also been the fierce defenders of my selfhood, in times when I thought that doing what I love might mean giving up part of who I am, they were always there to tell me otherwise. Very few in my life understand what this journey has meant as they do, and none have contributed so much to its completion. All my successes are dedicated implicitly and explicitly to them.

I also thank four other amazing members of my family: Donna Herman, John and Lucia Mudd, and Peter Mudd. They have showered me with love and affection for as long as I can remember, and they have done so unfailingly. No matter what I have done, or where life has taken my, their unbreakable care and sustenance has allowed me to go out into the world with confidence, knowing that I am not alone. Their affirmation has given me the conviction to be myself wherever I am.

I am lucky and grateful to have friends in my life who have transcended that inadequate term. My adopted brothers and sisters have both buoyed me and helped keep my feet on the ground. I thrive on my interactions with them: they challenge me, stimulate me, excite me, frustrate me, confuse me, and are all a part of who I am today. I cannot say enough good about these wonderful individuals: Adriel G. Mesznik, Jordan R. Seiler, Elizabeth A. Heller, James J. Hamlin, Colby Brin, Diego S. Segalini, Daria Masullo, James A. Carmichael, Christian E. Montegut, Daniel D. Rutimann, Pietrocelli Jiminez, Erik H. Wong, Matthew T. Cushman, Ben Foster, Tori Dodge, and Ilana J. Deluca.

Finally, I thank my incredible wife Heather Leigh. The most vivacious, talented, beautiful, supportive, intelligent, funny person I know. The one that I know I can turn to at any second when I’m in need of assistance. The one who has most intimately shared in my personal struggles and triumphs, and has made it possible for me to make it through the woods of self-doubt and self-persecution.

Table of Contents

Table of Figures

Introduction

General Introduction

As humans are primarily visual animals, it will come as no surprise that amongst the knowledge recorded on papyrus in ancient Egypt was documentation of various disorders of vision and accompanying treatments (Bryan, 1930). Ancient anatomists such as Galen were also aware of the extraocular muscles and their actions (May, 1968). Fascinatingly, however, the rapid discontinuous movements of the eyes known as saccades were not recognized until the pioneering work of Porterfield (1737), who documented fast and slow phases of nystagmus[1] through the use of afterimages. It is now recognized that saccades are one of the most common movements in the human repertoire of behavioral output. They are used to very rapidly (up to 800°/s) orient the relatively small area (1°) of our best vision (the maximally receptor-dense region of the retina called the fovea) to environmental features of interest, facilitating the ability to quickly obtain a high-resolution view of a large part of our world through the small foveal window.

Brief history of eye movement research:

The study of post-rotational nystagmus via afterimages dominated eye movement research prior to the study of reading by Hering (1879), and Lamare (1892). These researchers independently developed techniques that allowed for the detection of saccades by the "distinctive" (Lamare, 1892) or "clapping" (Hering, 1879) sound that accompanied their execution, which was detected via listening devices applied directly to the surface of the eye. For a brief period, movements of the eyes were recorded by attaching devices physically to the eye by using an eye-cup (Delbarre and Delabarre, 1898; Huey, 1898). However, it was the development of the photographic method of recording eye movements (Dodge and Cline, 1901; Dodge, 1903) that truly began a revolution in the study of eye movements (Taylor, 1937). Several others developed similar, relatively easy to use eye tracking systems, and it was soon realized that the dominant

---

[1] Rhythmic, oscillating involuntary eye movement.

pattern of eye movements seen during reading (sequences of saccades and periods of non-movement called fixations) was used in almost every other task humans engage in.

The flurry of research on saccades and fixations that followed the development of photographic eye-tracking methods set the stage for the next era of eye movement research by demonstrating the astoundingly diverse ways that these two simple actions could be used. Indeed, in their now-classic work on the "saccadic main sequence" Bahill and colleagues (1975), note that Dodge and Cline (1901) realized that saccade duration and average velocity are surrogate measures, since duration grows with increasing saccade amplitude. As a whole, the main sequence comprises the set of relationships between saccadic duration, peak velocity, and amplitude (Bahill et al., 1975).

It was the invention of the "scleral search coil system"[2] by Robinson (1963) that allowed for the sort of precise measurements and careful analyses that has characterized the study of saccades in the past 50 years. Robinson (1964) wasted no time, and used his new, high-accuracy recording technique to document: (1) the "pulse-step"[3] temporal profile of muscle tensions used in generating the step-like spatial profile of saccades, and (2) the negligible role played by the inertia of the eye in determining saccadic trajectory. In that early work, Robinson also developed and tested a model of the mechanics of the eye. This sort of basic modeling and characterization of eye mechanics and oculomotor control was extended to include the motor neurons and relevant brainstem nuclei thanks to the fortuitously contemporaneous advent of chronic single neuron recording in the awake behaving animal (Fuchs and Luschei, 1970, 1971; Keller and Robinson, 1971; Robinson, 1970; Schiller, 1970).

Though technological advancements made certain methods possible, it was the control systems approach initially championed by (Robinson, 1975) that fundamentally changed the way that researchers thought about eye movements. Because of their high speed and brief duration (a 10° saccade lasts

---

[2] A device for measuring eye orientation that employs a scleral contact lens embedded with wire coils and externally applied oscillating magnetic fields. Eye orientation is determined by measuring the current induced in the wire by the magnetic fields via leads coming off the lens coils.

[3] The temporal profile of extraocular muscle tension required to generate the step-like change in position achieved by a saccade has at least two components: (1) a tension pulse which rapidly carries the eye to the new position, and (2) a smaller amplitude step-change in tension (relative to the pre-saccade tension) which holds the eye at the new position.

roughly 40ms), it was thought that saccades must be “ballistic,” having a pre-determined trajectory not subject to feedback. Robinson's (1975) control system approach, and its important extension by Zee and colleagues (1976), made possible the realization that saccades must be subject to “local” feedback compensating for a large range of variability in the command signal. Indeed, the elaboration of this model explained a wide range of previously mysterious phenomena (Zee and Robinson, 1979).

The application of control theory to the study of saccades, and other types of movements, has continued to be quite useful. Specifically, this methodology has revealed that signal-dependent-noise[4] is a fundamental determinant of the feedback-control policies underlying the generation of a single movement trajectory (Harris and Wolpert, 1998; Todorov, 2005). The control theory approach to understanding motor control has also extended into attempts to understand motor learning, which we discuss further in the Background section, under ‘How is Error Used?’

Why study saccades?

By this time, each of the factors, still cited today (Sparks, 2002; Leigh and Zee, 2006) that make eye movements an excellent model system for the study of motor control in general had been documented: (1) they can be measured accurately, (2) the eye has limited degrees of freedom in its movement (3 dimensions), (3) there is no variable load on the eye, and thus no stretch reflex as are both present in other motor systems, (4) the neurons participating in oculomotor control are readily accessible, and (5) the computations underlying this control seem readily understandable. To this list, we would add the additional points: (6) the several different types of eye movements, and the maintenance of their performance over time means that there are a multitude of sufficiently distinct yet interrelated areas of study to support a rich and diverse field of eye movement research, and (7) datasets can be gathered quite rapidly, especially thanks to the recent development of infrared camera-based recording techniques (Ebisawa, 1995).

---

[4] Noise with a variance that scales in proportion to the size of the signal it accompanies.

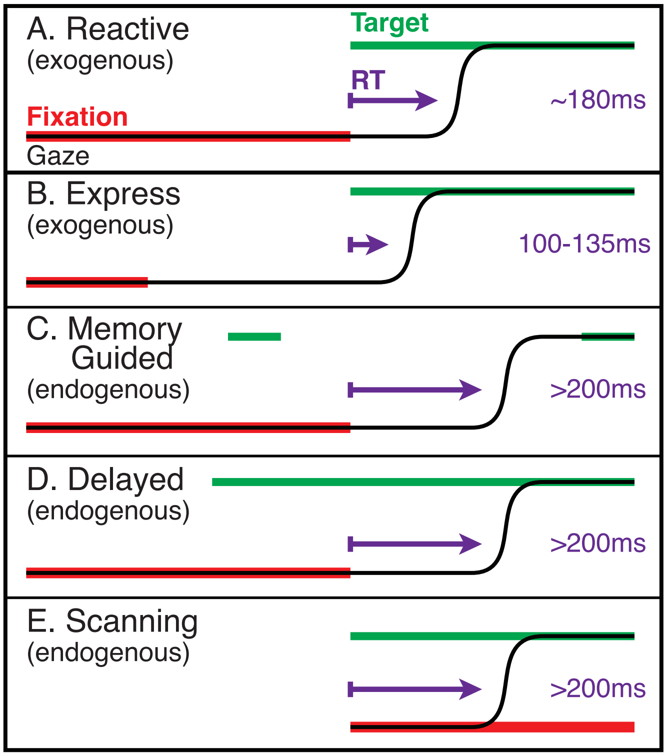

Figure 0.1 Types of Saccade (after the fashion of Hopp & Fuchs, 2004). **A**. Reactive or targeting saccades are triggered exogenously by the sudden onset of a **target**, and have **reaction times** (**RTs**) of approximately 180ms. **B**. Express saccades are triggered exogenously by **target** onset as well, but in order to occur, there must be a delay or gap between **fixation** offset and **target** onset; these can have **RTs** as low as 100-135ms in humans. **C**. Memory guided saccades are triggered endogenously to the remembered location of a previously presented **target** (though given a go-signal by **fixation** offset), and have **RTs** larger than 200ms. **D**. Delayed saccades are triggered also triggered endogenously though through a **fixation**-offset go signal, to an already-present **target**, and have **RTs** greater than 200ms. **E**. Scanning saccades are triggered endogenously (sometimes by the onset of a stimulus array) in order to fixate a new stimulus of interest, and have **RTs** larger than 200ms.

Figure 0.1 - Types of Saccade

Categorization of saccadic eye movements:

One quite general categorization that has been useful in the study of saccades is their partitioning into various types (Figure 0.1). Most broadly, saccades can be triggered endogenously, by a decision to move the eyes to an environmental feature of interest, or exogenously by the sudden appearance of a relevant stimulus. Saccades can also be characterized by their reaction time, the time that elapses between the signal to execute a saccade and their moment of onset. Based on these criteria, the generally recognized categories of saccade are: (A) Reactive or targeting, triggered exogenously by the simultaneous onset of a target and offset of fixation; (B) Express, triggered exogenously by target onset *after* fixation offset (a so-called gap paradigm), (C) Memory guided, triggered endogenously (though given a go-signal by the fixation offset) to the remembered location of a previously presented target; (D) Delayed, also triggered endogenously through the external fixation-off go-signal to an already present target; (E) Scanning, between two continuously present stimuli (and sometimes by the onset of a stimulus array, as in visual search).

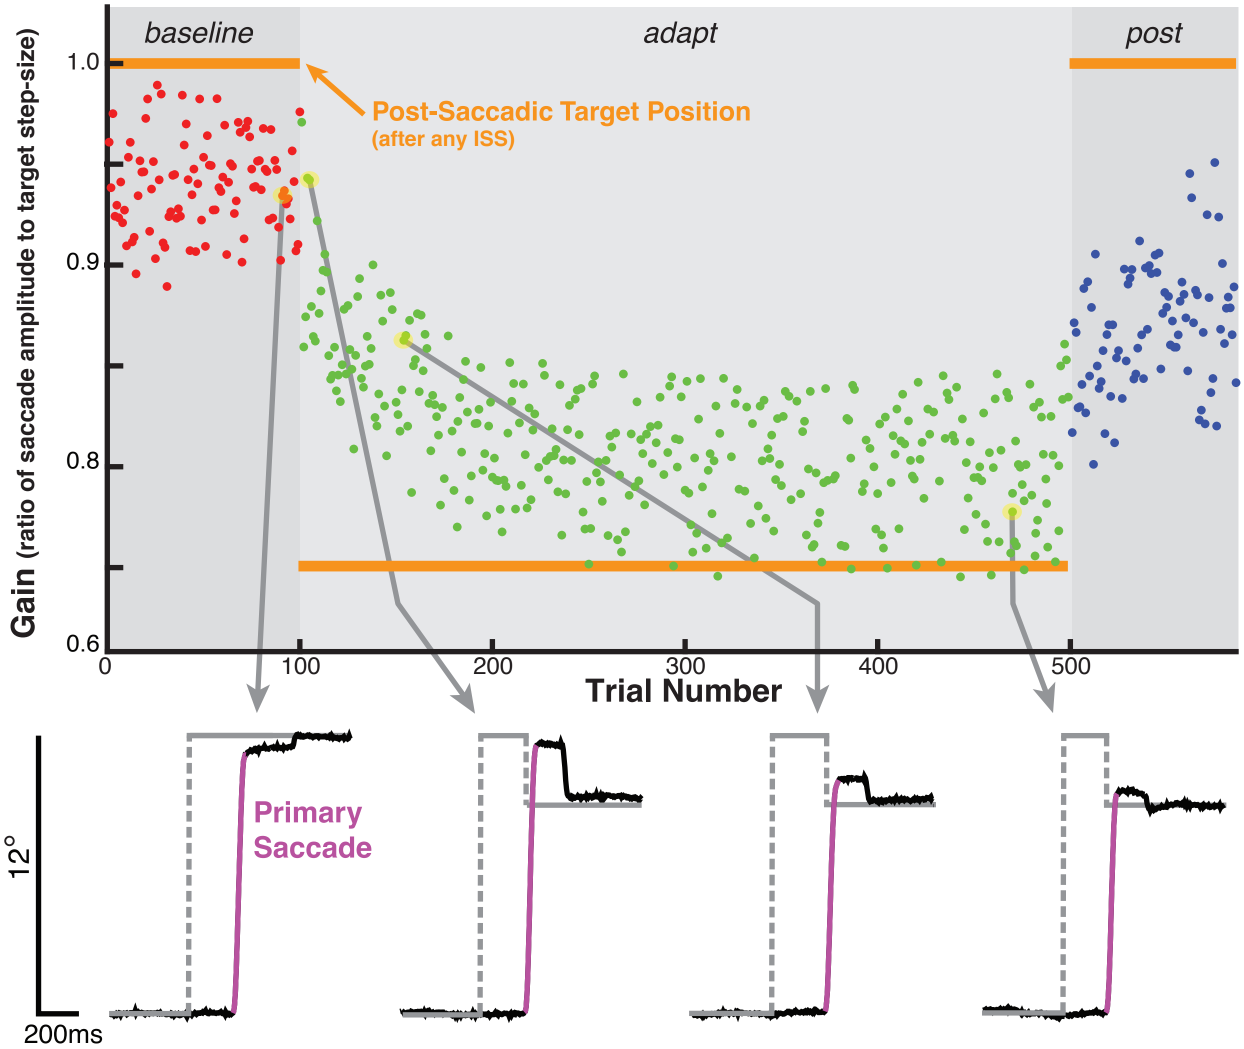

Figure 0.2 The Effects of ISS Adaptation. A typical ISS (intrasaccadic target step) adaptation session is captured by a plot of the primary saccade gain (ratio of saccade amplitude to target step-size) in each trial. Average **baseline-phase saccade gain** is slightly less than 1 indicating that primary saccades generally fall short of the target, even in the absence of any ISSs. **Adapt-phase gain** decreases progressively to an asymptote. Note that adaptation is incomplete, primary saccades do not resume their tendency to undershoot the **post-saccadic target location** by the end of the adapt-phase. **Post-phase gain** progressively increases towards the baseline gain level once ISSs are removed.

Figure 0.2 – The effects of ISS-adaptation

The origins of saccadic adaptation

Amongst the host of advances in research regarding saccades and eye movements in general in the late 1960s and the 1970s were those most germane to this thesis. While it was recognized that there was variation in saccade gain from individual to individual (gain: the ratio of saccade amplitude to target eccentricity or step-size), and control system models included elements capable of capturing such variation (Zee et al., 1976a), others soon realized that both pathological (Kommerell et al., 1976; Abel et al., 1978), and normal saccade gain (McLaughlin, 1967), could be modified by experience. Kommerell, Abel and their collaborators documented patients who made low-gain saccades with one eye, and relatively normal movements with the other. They both found that forcing the patient to view the world through only the weakened eye, over several days, had the effect of restoring relatively normal saccadic gain. Meanwhile, Mclaughlin found that by arranging for each of a series of saccades (to point-like stimuli in the lab) to trigger consistent 1° secondary intra-saccadic target steps (ISSs) back towards the fixation, saccadic gain could be noticeably decreased in just a handful of trials. These were the beginning of the study of what is now known generally as saccade adaptation.

An example of a typical McLaughlin-style, ISS adaptation session is presented in Figure 0.2. At the time, because of the extremely rapid adjustment that McLaughlin demonstrated (on the order of minutes) and the relatively slow adjustment observed by Kommerell and Abel (on the order of days), it was thought that these two cases of gain adjustment were supported by distinct mechanisms. However, it was later shown by Scudder and colleagues (1998) that the two forms of adaptation shared the same time course, when compared under similar conditions. Thus it seemed clear that one source of error must drive the changes in gain observed for saccadic dysmetria in general.

Just before Kommerell and colleagues' (1976) publication, Optican and Robinson, (1980) had begun another important study in which they deliberately weakened extraocular muscles in monkeys, similarly resulting in low saccade gain in one eye, and normal saccade gain in the other. Furthermore, motivated in part by the knowledge that humans with certain forms of cerebellar degeneration show lasting saccade dysmetria (Zee et al., 1976b), Optican and Robinson were able to show that the cerebellum plays a key role in the adaptive plasticity of saccades. A host of subsequent studies have confirmed the importance of the cerebellum in saccade adaptation (for a review, see Robinson, Fuchs, and Noto, 2002).

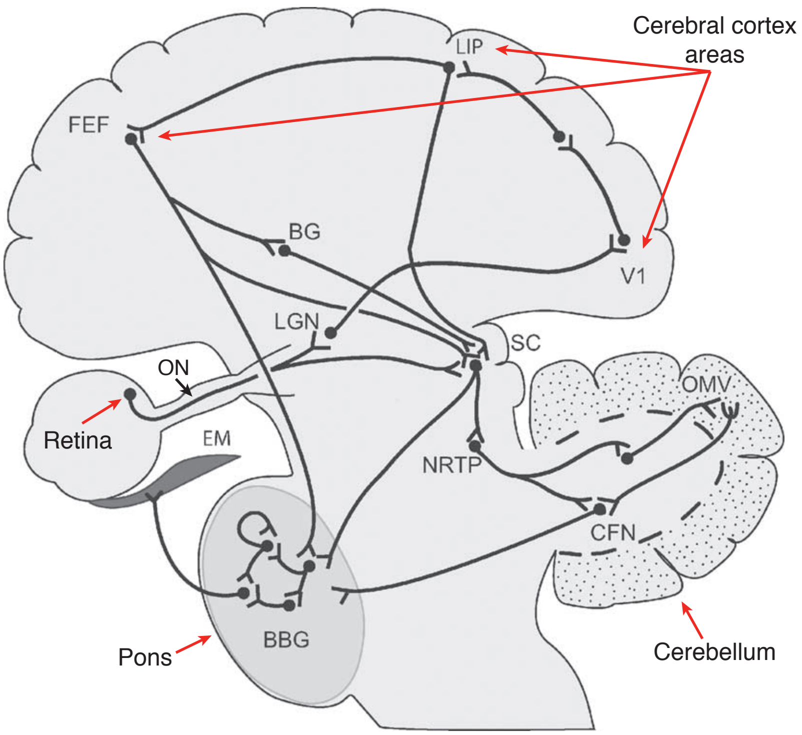

Figure 0.3 Modified from Hopp & Fuchs (2004). A summary of anatomical loci involved in primate saccade system plasticity. Arrows are drawn in red only to distinguish them visually from black inter-areal projections. **ON**: Optic nerve, **LGN**: Lateral geniculate nucleus, **V1**: Primary visual cortex, **LIP**: Lateral intraparietal cortex, **FEF**: Frontal eye fields, **BG**: basal ganglia, **SC**: Superior colliculus, **NRTP**: Nucleus reticularis tegmenti pontis, **CFN**: Caudal fastigial nucleus, **OMV**: Oculomotor vermis, **BBG**: brainstem burst generator, **EM**: Extraocular muscles.

Figure 0.3 – Saccade adaptation relevant anatomy

Neuroanatomy of the oculomotor system

Along with the cerebellum, several areas in the primate brain have been implicated in the plasticity of saccadic eye movements (Figure 0.3; for a review, see Hopp and Fuchs, 2004). The other brain structure that has been most commonly implicated in saccade adaptation is the superior colliculus (SC), a layered midbrain structure whose superficial layers (SCs) receives a direct input from the retina (Moschovakis et al., 1996), while the intermediate (SCi) and deep layers (SCd) project directly to the premotor brainstem burst generator (BBG) in the pons, and indirectly to the cerebellum. Waitzman and colleagues (1988) pointed to the SC as the source for saccadic “motor error,” while Krauzlis and colleagues (1997) suggested it as the source of motor error for both saccades and pursuit. Motor error in this context referred to the driving force behind an eye movement that would align gaze and target, rather than an error signal that would drive changes in saccade gain as a result of dysmetria. However, two more recent studies have demonstrated that consistently stimulating the SC immediately after saccade completion leads to changes in gain (Kaku et al., 2009; Soetedjo et al., 2009). Thus it seems that the SC may serve both of these functions. Further, it seems that activity in the SC in response to eccentric targets does not change during short-term (Mclaughlin-type, ISS-driven) adaptation (Frens and Van Opstal, 1997). Strangely, however, early evidence suggested that saccades evoked by directly stimulating the SC following ISS-adaptation are unmodified by adaptation (Fitzgibbon and Goldberg, 1985). This was subsequently demonstrated to be a methodological issue, however, related to the intensity of stimulation, as Edelman and Goldberg (2002) found that stimulation-evoked saccades indeed displayed the hallmarks of visually-evoked saccades, when stimulation intensity was carefully controlled. These findings implicate the SC as an important player in the adaptive process, but one that sits upstream of the site of adaptive changes.

Further relevant background on saccade adaptation topics specific to the experiments that support this thesis are to be found both in the background section (below), and in the each chapter's introduction. The background section, as well as the general discussion in chapter 6, is divided into 3 subsections with questions as headings; these questions (broadly) motivated all of the experiments described in the chapters that follow. The decision to give each chapter its own introduction was taken because the experiments described in each chapter are quite independent, and further because these chapters have been written with the chief aim of eventual publication of each set of original experiments. Two published articles are presented as well: one in Appendix A, which was followed up by the work in chapter 2, and one as chapter 5.

Background

What is the error signal?

In the absence of pathology, or an experimental perturbation, saccade accuracy stays roughly constant over time (Warabi et al., 1984; Munoz et al., 1998). Ergo, saccadic gain is actively maintained by the brain. Further, Mclaughlin's ISS paradigm can be used to cause adaptive changes that vary with ISS-amplitude (Miller et al., 1981) suggesting the existence of some "error signal" which is instructive as to the direction and magnitude of ongoing saccadic dysmetria. Early on in the study of saccade adaptation, both retinal error (the difference between target and gaze) at primary saccade completion (Optican and Robinson, 1980), and the ensuing corrective saccade(s) required for eventual target fixation (see, for example, Albano and King, 1989), were offered as candidate error signals.

By showing that adaptation can occur in the near absence of corrective saccades, two studies demonstrated that post-saccadic visual information must be the driving force behind saccade adaptation (Wallman and Fuchs, 1998; Noto and Robinson, 2001). However, the conclusion that the error signal must be visual in nature does not amount to a conclusion that retinal error is the error signal determining saccade adaptation. Bahcall and Kowler (2000) found that when they instructed subjects to saccade partway to a target that then underwent a backwards ISS, gain-decrease adaptation occurred, despite the conclusion of each primary saccade in a positive (onwards) retinal error. They explained these results in terms of the ISS introducing a prediction error between the expected and actual visual error after the saccade.

In the study of motor learning at large, there is wide agreement that adaptive modifications of behavioral output are driven by sensory prediction errors (SPEs): the difference between the predicted and actual sensory consequences of a movement (Shadmehr et al., 2010). Given what is known about active internal monitoring of saccade trajectory from the control systems approach (see above), the idea that SPE might govern adaptive changes to saccade gain is well founded. Indeed, beyond the internal monitoring of trajectory, it has been demonstrated that there are single neurons whose pre-saccade activity is predictive of the post-saccadic presence of a stimulus in their response field, both in area LIP

(Duhamel et al., 1992) and the SCi (Walker et al., 1995). The SPE would also unify the findings of Bahcall and Kowler (2000) with those of more conventional studies of saccade adaptation, as the SPE caused by an ISS would be identical regardless of primary saccade amplitude (Figure 0.4).

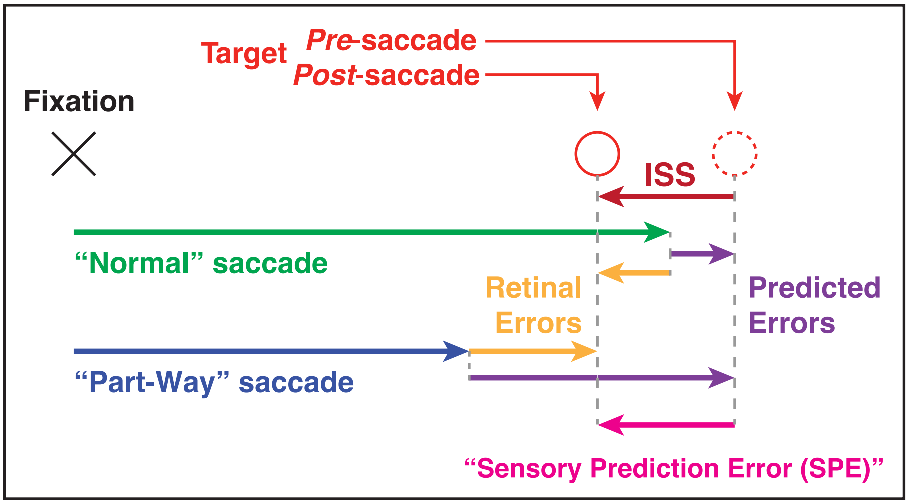

Figure 0.4 Schematic of the **Sensory Prediction Error** (**SPE**) concept. **SPE** does not depend on saccade amplitude: **SPE** is identical whether resulting from a **"Normal" saccade** typical of nearly all studies of saccade adaptation induced by **intrasaccadic target steps** (**ISSs**), or a **"Part-Way" saccade** like those in the adaptation study by Bahcall and Kowler (2000). Despite a large difference in **retinal error**, thanks to offsetting (subtraction) by the **predicted error**, the SPEs are the same. Note that the **SPE** and the **ISS** are identical in magnitude and direction in this heuristic example because noise sources have been ignored for the purposes of illustration. The **SPE** would, in actuality, likely equal the **ISS** only on average.

Figure 0.4 – Schematic of the Sensory Prediction Error concept

A particularly striking example of the SPE demonstrated that its action could override a conscious strategy (Mazzoni and Krakauer, 2006). In this study, regardless of whether subjects were aware of a visuomotor rotation resulting in altered feedback about hand movements, they adapted in response to the rotation. Amazingly, this meant that the group of subjects made aware of the rotation initially made accurate movements as a result of the application of a conscious strategy, but their movements became progressively dysmetric (Figure 0.5). The commonality between the aware and unaware group was that an SPE was present in both, convincingly demonstrating the explanatory power of this idea. Note that this is quite similar to the way that the instruction to saccade "part-way" in the Bahcall and Kowler (2000) study caused different retinal errors, but resulted in the same SPE, here the instructions caused different angular errors, but resulted in the same SPE.

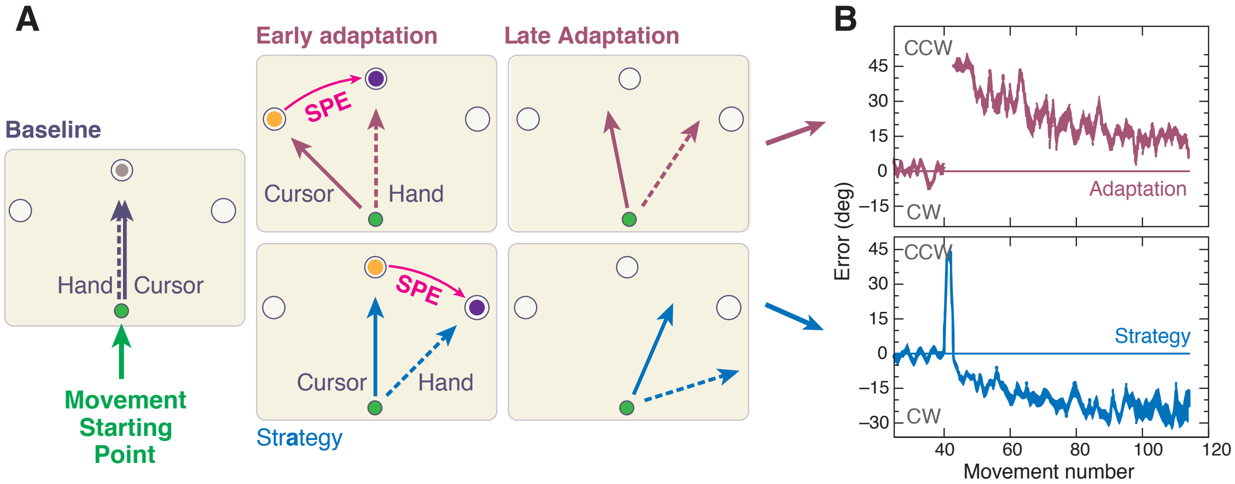

Figure 0.5 A particularly exotic example of SPE in action (modified from Shadmehr et al., 2010). **A**. Experimental schematic. Subjects first underwent a baseline phase with no visuomotor rotation (movement numbers below 40 in B). Next, subjects in the "normal" group underwent visuomotor adaptation: as they attempted to move the cursor coupled to their hand movements towards the top-marker early in adaptation, they were shown cursor feedback rotated counterlockwise by 45°, leading to the eventual adjustment of motor output such that by late in the adaptation, they were aiming more clockwise, towards the right-marker, resulting in feedback which came closer to the top-marker goal. Meanwhile, after a few trials under the visuomotor rotation (spike in B, bottom) the "strategy" group was explicitly instructed about the presence of the rotation, and were instructed that it could be effectively dealt with by aiming for the right-marker. Yet they too showed adaptation, driving angular error to become increasingly negative despite the initial success of the strategy in eliminating it entirely. An explanation is offered by the suggestion that **Sensory Prediction Errror** (**SPE, magenta arrow**) was driving the adaptation in both groups: the **sensory prediction** (**purple dots**) of the cursors position was (at least early in adaptation) tied to the trajectory of the hand, with subjects expecting to find the cursor on the marker towards which they had aimed. However, as a result of the applied visuomotor rotation, the **error feedback** (**orange dots**) from the displayed cursor's position was rotated by 45°, resulting in an **SPE** with equal magnitude and direction in both groups, and an accompanying similar amount and rate of adaptation in both groups. **B**. Angular errors resulting from each of the series of movements produced by the two groups, "normal" group on top, and "strategy" group on bottom.

Figure 0.5 – An exotic example of SPE

Powerful and widely accepted though SPE is, there are subtle differences in the way it is interpreted to operate in saccade adaptation. In their work exploring the basis of "spontaneous recovery" (the reappearance of a learned behavior after having been, seemingly, extinguished) in saccade adaptation, Ethier and colleagues (Ethier et al., 2008a) employed a state-space model to fit their data and explain spontaneous-recovery-like aspects of their observations (the specifics of their explanation is presented in *How is Error Used?*, below). In that model, they presumed that the error signal was a bias subtracted retinal error: retinal error less a constant bias term (y_b) corresponding to the average undershoot of unperturbed baseline saccades ($E = RE - y_b$) The assumption implicit in this definition is that the predicted sensory outcome of each movement is static: the error is only zero if the saccade lands a distance y_b short of the target. In support of this proposition, (Wong and Shelhamer, 2011) found that retinal error minus "inherent hypometria" (their term which corresponds exactly to y_b) best accounted for adaptation results in their constant-error (variable ISS) paradigm[5]. Note that bias-subtracted-retinal-error would predict a gain increase to occur in the saccade part-way paradigm of Bahcall and Kowler (2000) rather than the observed gain decrease. Meanwhile, several other studies have confirmed that some form of prediction must be involved in the evaluation of error (Henson, 1978; Robinson et al., 2003; Collins and Wallman, 2012).

More exotically, two studies have found that gain changes are possible in the complete absence of error. Madelain and colleagues (2011) found that it was possible to use amplitude-dependent reinforcement to drive changes in saccade gain in a paradigm in which the target was always absent at saccade completion. It is somewhat surprising that this work is the only one to have explored the possible role of motivation or task demands in saccade adaptation, given the known importance of the basal ganglia for motor learning in general (Graybiel et al., 1994). Meanwhile, Bonnetblanc and Baraduc (2007) found that simply by extinguishing a target during the saccade, large amplitude movements (>15°) that are initially hypometric, progressively increase in amplitude until error is close to the range observed with a continuously illuminated target.

Finally, (Harris, 1995) proposed that minimizing total-flight-time of the eye en route to a target might principally explain the tendency to undershoot (Becker, 1989a). However, he also proposed in that work that if the oculomotor system monitored total-flight-time over the course of many saccades, changes in this variable could be used to signal the need to adjust saccade gain. Though the former proposition has been more widely recognized than the latter, the Harris model of saccade adaptation based on this idea does reproduce the gain-increase gain-decrease rate asymmetry documented many times over (Miller et

---

[5] In this infrequently used form of ISS adaptation, ISSs are not constant either in degrees (°) or as a percentage of the primary target amplitude, but rather are chosen (based on saccade amplitude) to achieve a specific post-saccadic retinal error.

al., 1981; Deubel et al., 1986; Fitzgibbon et al., 1986; Deubel, 1987; Straube et al., 1997; Scudder et al., 1998; Kroller et al., 1999).

The theme of what constitutes the error signal for saccade adaptation is touched on in every chapter in this dissertation, being one of the principal questions that have motivated the work here contained.

In chapter 1, we consider mechanisms that might support adaptation without an identifiable saccade target. Under such conditions, it is difficult to conceive of an SPE or even a retinal error being computed, yet there are reports of adaptation utilizing such ambiguous stimuli.

In chapter 2, we consider the possibility that predictable or unpredictable target behavior may constitute a component of the signal driving adaptive changes.

In chapter 3, we explore the possible contributions of various putative error signals in our examination of adaptation in response to systematically varied small intra-saccadic steps.

Finally, in chapter 4, we used a paradigm in which the target was frequently extinguished intrasaccadically to examine impacts on learning dynamics.

How are errors detected?

In addressing the question of what constitutes the error signal, it was assumed that localizing the target after the saccade (for the purposes of calculating retinal error, for instance) is trivial. While it is true that in the natural environment, we do sometimes make saccades to lone, trivially discriminated stimuli, it is also common to direct our gaze to one of many similar objects, or to explore a single object, or in fixating portions of less structured visual stimuli such as textures or patterns. In the laboratory, however, saccade adaptation has traditionally been studied using small discrete targets – deemphasizing the potential challenges faced by the oculomotor system in extracting an error signal from the information-dense image it is confronted with upon saccade completion. That is, in most studies of saccade adaptation, it is safe to assume that bottom-up vision constitutes the “actual sensory information” for SPE, what happens when more is required?

Experiments exploring the use of landmarks in trans-saccadic localization have suggested that these additional stimuli can influence the perceptual detection of intra-saccadic target displacement. Deubel (2004) found that intrasaccadically shifting nearby objects along with a target modulated the perceived direction of the target's displacement (at least when the target is extinguished during and briefly after a saccade). However, it is not clear if this type of perceptual intra-saccadic detection is related to that relied upon by the oculomotor system for adaptation. When such targets are not extinguished during and after the saccade, perceptual reports regarding intra-saccadic displacement are much poorer (Deubel et al., 1996). Indeed, it seems that subjects are more apt to report a continuously visible stimulus as intrasaccadically stable regardless of whether it moved or not (Deubel et al., 2010). Also, though extinguishing these targets during and after a saccade facilitates detection of their displacement (Deubel et al., 1996), it also significantly decreases the impact of that displacement on adaptation (Shafer et al., 2000).

If active maintenance of saccadic accuracy functions successfully in the natural environment – as appears possible (Kommerell et al., 1976; Abel et al., 1978) – saccades to objects may represent a uniquely good opportunity for detecting errors in the oculomotor system. Such movements have a well-defined goal, rendering the process of determining saccadic dysmetria relatively straightforward compared to making saccades to unstructured stimuli or within a larger object.

Those studies that have explored adaptation with both a target and a background present have shown that objects do play a significant role in saccade adaptation. In their study of the interactions between target and background during saccade adaptation, Ditterich and colleagues (1999) found that the size of the target determined the extent to which intra-saccadic background shifts contributed to adaptation. When the target was small, shifts of the background had no effect on adaptation, but when the target was a large circle, shifting the background along with the target caused greater adaptation than shifting it in opposition to the target. This result suggests that the process of selecting visual information specific to the target is an active one. In keeping with these findings, Robinson and colleagues (2000) found that adaptation with a point-like target in monkeys was unaffected by background shifts. Interestingly, however, they also found that the mere presence of a background image reduced adaptation magnitude compared to the case of a target alone. This latter finding suggests that the addition of the background may increase the difficulty of localizing the target post-saccadically, or calculating an error associated with it, as delaying the presentation of error is known to reduce its efficacy (Shafer et al., 2000).

The two studies that have explored the possibility of adaptation with a background alone have found that it can indeed take place. Ditterich and colleagues (1999), and Deubel (1991) both found that intra-saccadic shifts of an unstructured background image does cause adaptation. These again suggest that the impotence of background shifts (in contributing to adaptation) in the presence of a target is the result of an active selection of target information (or suppression of non-target information). These studies are somewhat in conflict with other work that has found that intra-saccadic displacement detection is severely impaired (compared to such detection during fixation), perhaps due to saccadic suppression of vision (for example, Bridgeman et al., 1975). However, more in keeping with Ditterich and colleagues (1999) and Deubel (1991) are the findings of Castet and Masson (2000), who found that intra-saccadic velocity detection is possible in the right frequency band.

One of the most intriguing reports on the role of more complex stimuli in saccade adaptation makes the claim that ISSs do not cause adaptation when saccades are made within an object, only when they are made between objects (Collins et al., 2007). Collins and colleagues' (2007) work is supported by work on two-saccade sequences which demonstrates that secondary saccades only compensate for ISSs during inter-object primary saccades, and not during intra-object primary saccades (Vergilino-Perez et al., 2004).

Finally, work from our own lab has demonstrated that when confronted with two targets upon saccade completion (the saccade target and a distractor) adaptive responses are specific to the target, despite the target initially being positioned more remotely (Appendix A - Madelain et al., 2010). We found that this target-specific adaptation was indistinguishable from adaptation with a single target, and that positioning the distractor remotely did not cause adaptation. The finding that this form of adaptation was the same as that of a single target suggests that (as proposed above) making saccades to clearly defined objects does represent a unique opportunity for the oculomotor system, in that adaptation seems to proceed as normal despite the presumed additional visual processing that would be required to calculate an error specific to the target, or ignore the distractor.

In chapter 1, we explore intra-saccadic detection of whole-field stimulus displacement as a means to understand how such displacement might signal saccadic error, addressing the possible impact of stimulus content in terms of the presence of a target-object or high-level features, as well as exploring the role of intra-saccadic versus trans-saccadic vision in constraining the mechanism of displacement detection.

In chapter 2, we directly interrogate the difference between adaptation which may rely more heavily on bottom-up vision post-saccadically versus adaptation which requires top-down vision for error detection.

In chapter 3, we consider the possibility that error detection is related to the statistics of the movements themselves.

In chapter 4, we point out that even if an error is not explicitly detected in a visual sense, it is possible that one is still formally present in the sense of a force which actively drives change in gain in the same fashion that a visually detected error would.

How are errors used?

Though framing saccade adaptation with the question-headings chosen here attempts to address different aspects of the adaptation process independently, it is likely that processes such as the calculation or detection of error and its utilization in driving adaptation are intimately intertwined. For instance, using a point-like target, Havermann and Lappe (2010) found that increasing noise in retinal error by systematically varying it from trial to trial attenuated the adaptation induced (compared to the case with no such additional variation), though the average retinal error was identical between compared conditions. Though it is likely that individual saccadic errors were calculated the same way on each trial, it does not seem to make sense to describe the error signals themselves as being the same between different noise conditions. In the following, while we attempt to address only the way that errors are used in driving adaptation, we recognize that it may be impossible to do so without considering the nature of the error signal or its detection.

Many experiments have demonstrated the ways in which error can be used to adapt gain in a saccade-specific way. Many different types of specificity have been demonstrated: leftward saccades may be adapted independently from rightward saccades (Semmlow et al., 1989), it is possible to adapt different saccade amplitudes in the same direction (Miller et al., 1981; Frens and Van Opstal, 1997), and several studies have emphasized some degree of independence in adapting different saccade types (for example reactive versus voluntary saccades; (Erkelens and Hulleman, 1993; Deubel, 1995a, 1995b; Fujita et al., 2002; Gaveau et al., 2005; but see Fuchs et al., 1996; Deubel, 1999; Hopp and Fuchs, 2002).

It is possible that adaptation specificity (or lack thereof) may fall under the more general aegis of context; for example, perhaps differential adaptation of distinct saccade types is the manifestation of those types of saccadic performance constituting distinct contexts for action. In the motor learning literature writ large, it has been suggested that the cerebellum may support the formation, utilization, and adjustment of distinct sensory predictions through distinct, paired “forward” and “inverse” models (Wolpert et al., 1998). Forward models are used to predict the sensory and motor consequences of movement (position, velocity, et cetera) using the outgoing motor command and current sensory data; inverse models are used to determine the appropriate motor command for a given desired motor consequence. Such paired forward and inverse models could support contextual motor learning if each context has a corresponding pair of, with the forward model being adjusted based on the difference between the forward model prediction of the sensory consequences of a movement and the actual outcome (this is essentially SPE, discussed above). Note that though this formalism has more traditionally been used to represent ideas about feedback corrections that occur during a movement, conceptually it applies equally well to the offline error feedback correction needed to support saccade adaptation (Chen-Harris et al., 2008).

Indeed, work from several labs has demonstrated that saccade adaptation can be sensitive to a variety of contextual cues. These include target eccentricity (Chaturvedi and van Gisbergen, 1999), horizontal and vertical orbital eye position (Shelhamer and Clendaniel, 2002a; Alahyane and Pelisson, 2004; Zimmermann and Lappe, 2011), head orientation (Shelhamer and Clendaniel, 2002b), and perhaps even the acceleration due to gravity . Though Deubel (1995b) found that a target's visual appearance was not an effective contextual cue, we challenged this in our own work (Chapter 5 – Herman et al., 2009), and our own conclusions were recently supported by the finding that visual appearance (as well as an auditory and a more ambient context) can enhance the effects of a head-orientation-defined context (Beaton et al., 2010). What of the use of error within a context, however? Putting aside for a moment the idea that error is used in a specific way, we must also consider how error might be used to adjust saccadic gain in one set of circumstances.

Specifically within the field of saccade adaptation (as noted above), several studies have shown that gain-increase adaptation is slower and less substantial than gain decrease, suggesting perhaps that onwards errors are used differently from backwards errors (Miller et al., 1981; Deubel et al., 1986; Fitzgibbon et al., 1986; Deubel, 1987; Straube et al., 1997; Scudder et al., 1998; Kroller et al., 1999). Recent work suggests that gain-decrease adaptation may act by adjusting the online control of saccades, while gain-increase adaptation more closely resembles a remapping of target position (Ethier et al., 2008b, but see Straube and Deubel, 1995; Alahyane and Pelisson, 2005).

The way that errors vary over time may also influence their impact on adaptation. As mentioned above, Havermann and Lappe (2010) found that increasing spatial noise in retinal error (in a fixed-error, variable-ISS paradigm) decreased induced adaptation, despite mean error remaining constant. However Srimal and colleagues (2008), findings in modeling adaptation with constant or random ISSs suggested that more spatially variable errors are used no differently than those with less variability. Meanwhile, it appears that gradually introduced errors are more effective in driving changes than abruptly introduced errors, at least for gain-increase adaptation (Wong and Shelhamer, 2011).

Indeed, several motor learning studies have suggested differences in the impact of gradually versus abruptly introduced errors (Kagerer et al., 1997; Klassen et al., 2005; Hatada et al., 2006; Michel et al., 2007; Huang and Shadmehr, 2009). This may be because smaller errors are treated differently from larger errors (Wei and Kording, 2009; Marko et al., 2012). It has also been suggested that the use of error is dependent on whether the error is ascribed to an internal or external cause (Berniker and Kording, 2008).

Most generally, it has been suggested that errors are recorded by memory systems with multiple different timescales of “remembering” and “forgetting,” and that error is stored (remembered) and retained (not forgotten) at rates which correspond to the timescale of the error-source (Kording et al., 2007; Joiner and Smith, 2008; Huang and Shadmehr, 2009). That is, memories are stored more slowly and retained for a longer period if errors persist, but are stored and forgotten rapidly if errors are more short-term. However, there are differing interpretations of the method by which the timescale of the error leads to changes in memory rates. One suggestion is that this process is implicit: as the error persists, memory stored at a slower rate has more of an opportunity to build and be retained (Joiner and Smith, 2008). The other is that the nature or statistics of the learning-event leads to an active choice of memory rates which are appropriate to the event; if the brain deems a memory worth remembering, it is retained for a longer period (Huang and Shadmehr, 2009). It is not clear whether these views are, in fact, incompatible, or simply two different ways of interpreting the same phenomenon.

These views about motor memory stem from efforts to model motor learning phenomena with a control theory perspective using the state-space approach. One of the goals of systems control is stable performance in the face of perturbation (Corriou, 2004), making it a natural perspective from which to model saccade adaptation. The state-space representation has an advantage over the more traditional transfer function[6] representation in that the relationship between input and output may be expressed in a way that closely resembles the process being modeled. State-space models have been applied to saccade adaptation data, with fruitful results, on several occasions (Ethier et al., 2008a, 2008b; Srimal et al., 2008). Importantly, these models are used under the assumption that the relevant time-scale is a single trial. Because inter-trial intervals are variable, such models fall under the heading of discrete-time linear control with variable-time sampling. Indeed, the analyses and conventions used in this thesis take

---

[6] A function encapsulating the relationship (in spatial or temporal frequency terms) between input and output. In the case of motor control, input might constitute sensory target-related information, and output the command signal being sent to the muscles.

the same approach: assuming that relevant information about the adaptive process can be gleaned by the temporally irregular observation of individual primary saccade gains. However, it is important to note that adaptive behavior can also be described a continuous-time process, as control theory models of single saccade trajectories explicitly do. Modeling of saccade adaptation is yet in its infancy, with a great many wide open questions as to the approach as well as the implications of findings yielded by modeling. Importantly, few have yet asked whether analytically determined limitations on control system performance apply to the brain's efforts in modifying saccadic behavior. For example, control systems are generally designed to handle perturbations in a single range of frequencies; a consequence of this is that they perform poorly when confronted with perturbations outside this range (Corriou, 2004). Is the brain subject to the same constraint, or might the its capacity for adaptation on many timescales render such a difficulty moot (Drew and Abbott, 2006)? This sort of comparison of known features of linear systems control and the operation of the brain in motor control will be necessary as state-space modeling of saccade adaptation matures.

A final concept that falls distinctly under the heading of error usage but which does not relate specifically to the use of saccadic position errors in driving changes in gain is the finding that the post-saccadic presentation of a moving stimulus can lead to changes in post-saccadic ocular drift without effecting saccade (Optican and Miles, 1985; Kapoula and Robinson, 1986; Deubel, 1991). The importance of this finding is that it shows that the type of error is intimately linked to its usage.

In chapter 2, we weigh in on the impact of errors that are more or less variable over time. We also consider the possibility that the way that errors are used explains the finding that onwards errors are less effective than downwards errors in inducing adaptation. That is, perhaps they are not used in adjusting primary saccade gain, but another aspect of saccadic behavior.

In chapter 3, we consider the possibility that, as in other forms of motor learning, smaller errors are used differently than larger errors, as well as attempting to ascertain what accounts for the seemingly differential impact (or use) of error by different individuals.

In chapter 4, we consider the possibility that the statistics of the learning environment modulates the way that error is used.

Finally, in chapter 5, we address the question of whether a target's visual appearance can serve as a contextual cue for adaptation.

Chapter 1: Objects, motion and correlation in the detection of trans-saccadic displacements

Introduction

To obtain high-resolution visual information from the environment, we orient the maximally sensitive foveal region of the retina towards features of interest using rapid movements of the eyes called saccades. Because saccades aimed at eccentricities beyond 8° are usually hypometric (too small), more than one saccade is frequently required to achieve a desired gaze orientation (Becker, 1989a); these secondary (and tertiary, et cetera) movements are called corrective saccades. While corrective saccades can occur in the absence of visual feedback (Becker and Fuchs, 1969; Hallett, 1978), it is widely accepted that such movements predominantly rely on visual mechanisms (Becker, 1972; Prablanc and Jeannerod, 1975; Deubel et al., 1982).

The nature of that visual mechanism is constrained by the conditions under which the corrective saccades are made. In the simplest conditions (and those most commonly studied in the laboratory), corrective saccades to a lone target may be driven entirely bottom-up; there is no need for trans-saccadic retention or integration of information. Only slightly more complicated is the circumstance of making saccades to a well-defined (and uniquely identifiable) target when more than one object is present: this task demands the trans-saccadic retention of the target's identity, which can be used to relocalize the target anew after each saccade (Appendix A - Madelain et al., 2010). Finally, most demanding is the case of making saccades to a target that is poorly defined or difficult to segregate from its background; under such conditions, it has been suggested that the oculomotor system may rely on a "correlation" mechanism operating on low-level features to achieve the performance observed (Deubel, 1991; Ditterich et al., 1999).

However, these two studies that proposed a correlation mechanism were focused on the adaptive plasticity of saccades (or saccade adaptation), which requires systematic, repeated intra-saccadic target shifts (McLaughlin, 1967; for a review, see Hopp and Fuchs, 2004; Pélisson et al., 2010). Consequently, it is unclear to what extent observed behavior in those experiments relied on the detection of such shifts on a single trial level, versus an accumulated detection of the systematic shift over many trials. Moreover, the structured backgrounds in these studies were shifted intrasaccadically, and early motion signals from low frequency components may have been detectable (Castet and Masson, 2000). Consequently, their corrective saccades and adaptation might have been driven by motion signals, rather than a trans-saccadic correlation mechanism.

To determine the relative importance of object identification, motion signals and cross-correlation to the trans-saccadic detection of stimulus displacement, we manipulated whole-field stimuli in three ways across two experiments: (1) We varied the amount of "objectness" either by explicitly degrading a target so that it resembled its background or by phase-scrambling the entire image; (2) we moved the image intrasaccadically and simultaneously extinguished it during the saccade on half the trials; (3) we flipped the contrast of the image intrasaccadically on half the trials. Our detection metric was the proportion of corrective movements in the direction of the imposed displacement. Our results support a strong role for both object identification and motion signal mechanisms in detection of stimulus displacement, and at most a weak effect of cross-correlation mechanisms. These findings are discussed in the context of saccade adaptation in more naturalistic, cluttered visual scenes.

Methods

Subjects

Informed consent was obtained prior to any recording sessions, and protocols were reviewed and approved by the Institutional Review Board (IRB) of the City College of New York (CCNY), and thus complied with all human-subject protocol requirements.

Four subjects, ages 28-37, participated in Experiment I (two males and two females). Eleven subjects, ages 18-37, participated in Experiment II (five males and six females). All were faculty and students from the CCNY community and had normal or corrected-to-normal vision.

Equipment

Stimuli were displayed on an Iiyama Vision Master Pro 514 CRT display (Oude Meer, Netherlands) at a resolution of 800 x 600 pixels (visible area 41.5 cm x 30.5 cm), and a vertical sync-rate of 200Hz with 8-bit color depth.

Stimulus generation and display, data storage, and overall experimental session orchestration were controlled with a custom interface in LabView (National Instruments, Austin, TX) running in Windows XP (Microsoft Corporation, Redmond, WA) on a Dell PC (Austin, TX).

Eye movements and gaze position were measured and collected by an Eyelink-1000 infrared camera system (SR-Research, Mississagua, Ontario, Canada), which sampled right-eye gaze (pupil - corneal reflection) at 1000Hz. Each participant sat in a darkened room, 57cm from the display, and underwent a nine-point self-paced calibration prior to the start of recording.

General Procedure

We presented images that filled the CRT display. Primary saccades always triggered intra-saccadic displacements of the whole image. Though the display was 800 x 600 pixels, all stimuli were generated at 900 x 600 pixels, and centered on the display so that horizontal translations would always result in a still-filled screen. Online saccade detection was achieved with a 20°/s velocity threshold. This threshold was higher than that used for offline analysis to avoid spurious triggering, and to ensure that saccade velocity was high when intrasaccadic target shifts occurred. The delay from saccade start (as measured offline by a 10°/s threshold) to a stimulus change on the display was 16 ± 3 ms (mean ± SD). On some trials ("Intrasaccadically-extinguished"), stimuli were also replaced at saccade onset with an 800 x 600 gray image, whose luminance was calculated from the mean pixel value of the pre-saccade stimulus being used in the current trial; the shifted image was reilluminated upon saccade completion (velocity below 20°/s threshold).

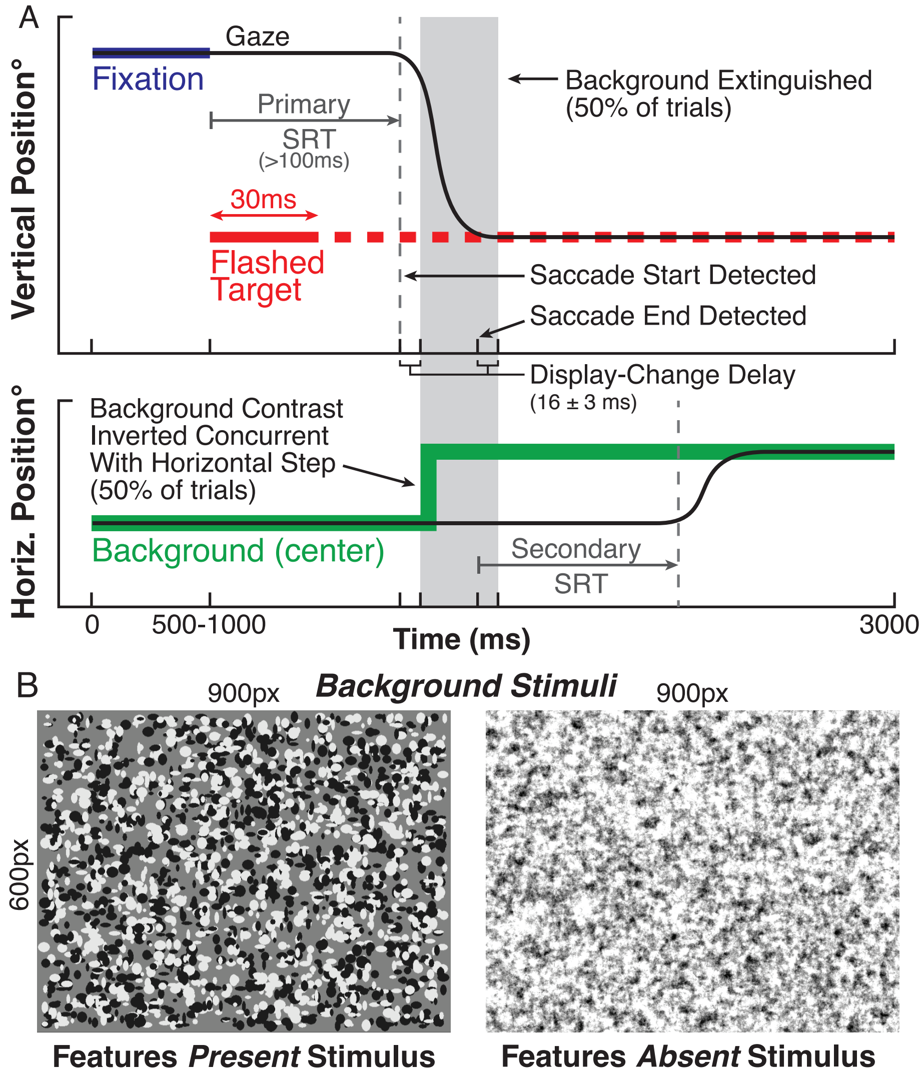

Figure 1.1 Experiment I procedure and stimuli. **A**. Trial structure. Subjects fixated a small (0.6°) **blue cross** extinguished with the onset of a **red target square** flashed briefly 5° below serving as the saccade go-signal. Upon detection of saccade start, the **background** was shifted horizontally (-1°, -0.5°, +0.5°, or +1°). On 50% of trials, the background was extinguished until the end of the saccade (interval indicated by the gray rectangle). On an independent 50% of trials, concurrent with the shift, background contrast was inverted. In trials with both contrast-inversion and intrasaccadic blanking, the modified image was not displayed until saccade completion. Background shifts elicited secondary, corrective saccades as subjects attempted to track the movements of the background. **B**. Examples of the stimuli used in Experiment I. Background stimuli were generated prior to each trial and were thus slightly different for each subject in each trial. A features present stimulus consisted of 2500 randomly placed ellipses (1250 light, and 1250 dark) on a neutral background. A features absent stimulus was generated by phase-scrambling (see Methods) a newly made features present stimulus. All stimuli were 100 pix wider than the display to ensure that horizontal shifts would always result in a filled screen.

Figure 1.1 – Chapter 1 / Experiment I Methods

Contrast-inversion of a given image was achieved simply by subtracting each pixel's intensity from the maximum 8-bit encoding value of 255.

Spatial aspects (positioning of fixation stimuli, horizontal intra-saccadic shift amplitude, et cetera) were chosen prior to all recording sessions and were the same for all subjects. However, all stimuli were generated on the fly prior to each trial, thus there were unique stimuli for each trial experienced by each subject. The recording duration was fixed at 2500ms, meaning a variable period of time passed following the primary saccade, depending on that movement's latency

Experiment I

Subjects were informed prior to the experiment that they would be presented with a fixation cross and that a flashed target would indicate the primary saccade goal and go-signal. They were further informed that their primary saccades would trigger a leftward or rightward shift of the background stimulus, which they should track with their eyes.

The procedure for a single trial is illustrated in Figure 1.1A. Subjects fixated a horizontally centered blue cross (0.6°) with vertical position drawn uniformly from the interval: [4.5°, 5.5°]. Once gaze was stable (within a 2° window of the fixation for 300ms), recording began; 500ms later, the fixation cross was extinguished and, concurrently, a small red target-square (0.3°) was flashed for 30ms, 5° below the fixation cross, serving as the go-signal. Upon detection of saccade start, the stimulus shifted horizontally (-1°, -0.5°, 0.5°, 1°). Two additional, independent, shift-concurrent perturbations were possible: (1) the stimulus was intrasaccadically-extinguished, or (2) the stimulus was contrast-inverted.

Two different stimulus image types were used, which we call: (1) features-present, and (2) features-absent (Figure 1.1B). Each features-present stimulus consisted of 2500 ellipses with positions selected uniformly from two intervals ([0,900] and [0,600]), 1250 dark and 1250 light, on a neutral gray background. Each ellipse had two radii (independently) drawn uniformly from the interval: [3 pixels, 13 pixels]. Each features-absent stimulus was generated by phase-scrambling a freshly generated features-present stimulus. Phase scrambling was achieved using LabView by first decomposing an image into Fourier amplitudes and phase-angles via Fourier transform. Each phase-angle was then modified by randomly adding either -π/2 or + π/2; amplitudes were left intact. A phase-scrambled image was then constructed from the unmodified amplitudes and modified phase-angles via inverse Fourier transform. Note that because the red target was only briefly flashed, there was no explicitly defined post-saccadic target in Experiment I, but this cued location did contain clearly identifiable objects in the features-present case, but not in the features-absent case.

In summary, we randomly interleaved trials in which three stimulus properties were systematically varied (8 conditions, 2 x 2 x 2 design): (1) feature presence (features-present or features-absent), (2) intra-saccadic illumination (continuously-illuminated or intrasaccadically-extinguished), and (3) contrast inversion (contrast-unchanged or contrast-inverted). Additionally, there were 4 possible intra-saccadic shift amplitudes (-1°, -0.5°, 0.5°, and 1°), resulting in 32 unique combinations of stimulus and shift amplitude. Each of these 32 combinations was repeated 10 times, resulting in 320 total trials.

Experiment II

Subjects were informed that they would be presented with a red fixation square, and that at some point, they would notice a visual transient (resulting from a portion of the background being replaced) that would serve as the primary saccade goal and go-signal. As before, they were instructed that their primary saccades would trigger a leftward or rightward shift of the stimulus, which they should track with their eyes.

The procedure for a single trial is illustrated in Figure 1.2. Subjects fixated a horizontally centered blue cross (0.6°) with vertical position drawn uniformly from the interval [-5.5°, -4.5°]. Once gaze was stable (within a 2° window of fixation for 300ms), recording began; 500ms later, the fixation was extinguished, and the 2° x 2° deg area centered 5° vertically upward from fixation was replaced by the degraded target-object. Upon detection of primary saccade start, the scene was concurrently shifted (-1°, +1°) and intrasaccadically-extinguished; upon saccade completion, the scene was reilluminated.

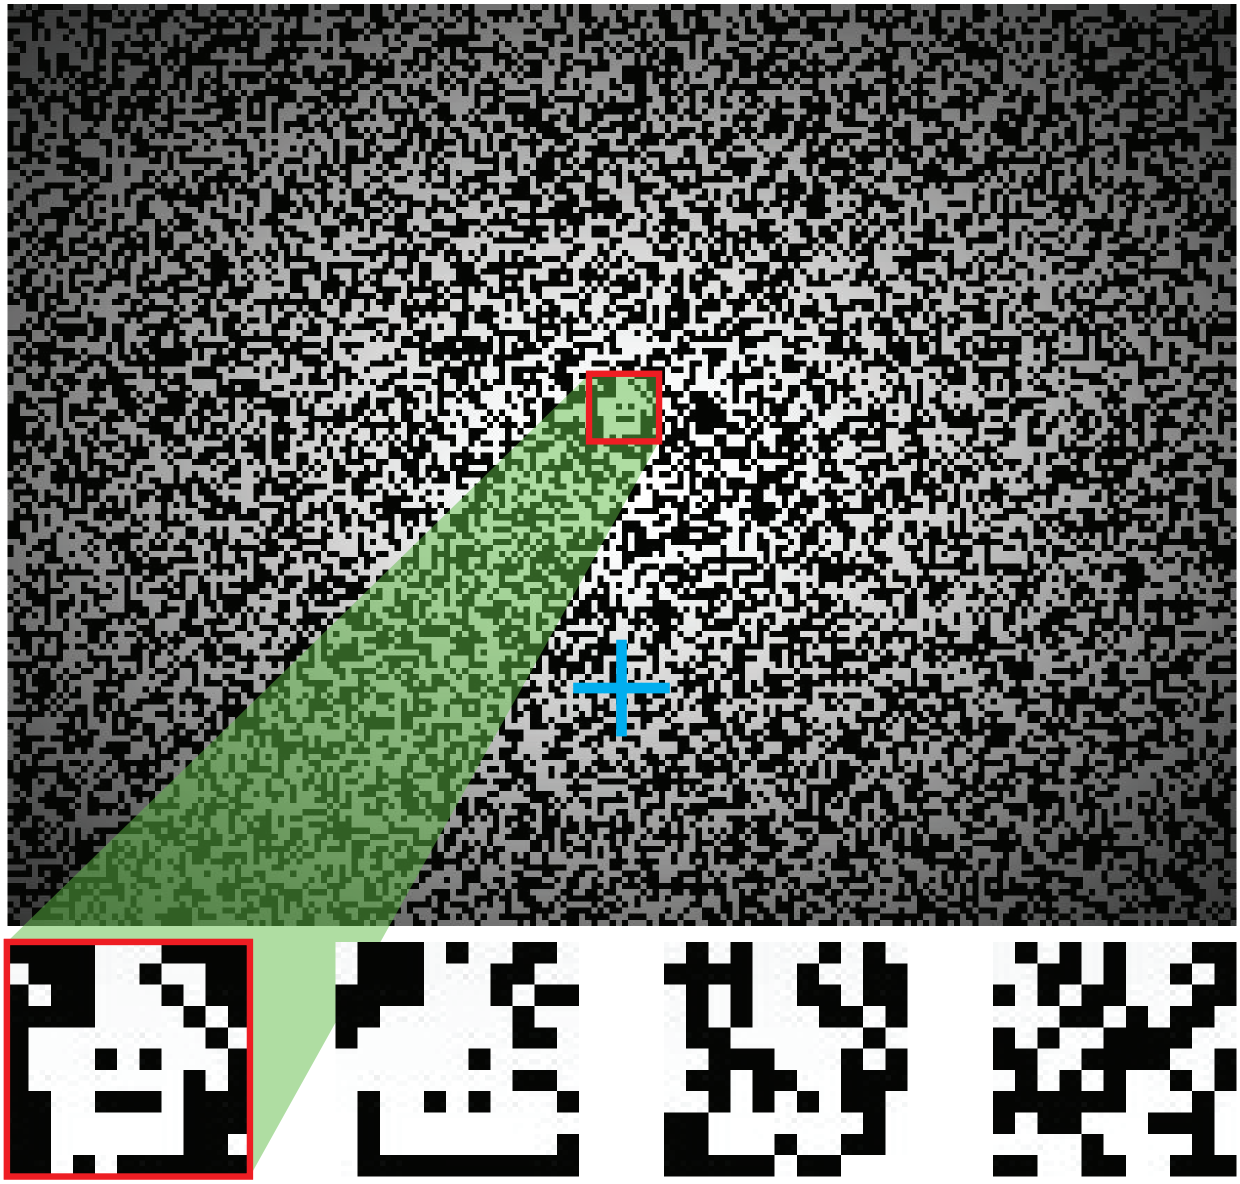

Examples of Partially Degraded Targets by Percent-Degradation

Figure 1.2 Experiment II Stimuli. Subjects fixated a small (0.6°) **blue cross** (color and size modified for the purposes of illustration) on a background of randomly-colored 2-pixel x 2-pixel black and white boxes. Concurrent with the extinction of the fixation, a **portion of the display** (10 x 10 boxes, 2° x 2°, highlighted in **red** again only for the purposes of illustration), was replaced by a partially degraded target-object. The degraded targets presented below the full-screen image are examples, stimuli were always generated prior to each trial. Percent-degradation was achieved by randomly choosing new colors (black or white) for that percentage of the boxes comprising the target region.

Figure 1.2 – Chapter 1 / Experiment 2 Stimuli

The whole-screen image was composed of “boxes” (2 x 2 pixels, 0.2° / side) randomly colored black or white with a Gaussian window applied to obscure the image-edges. The degraded targets were generated from a base image resembling a house or a face (Figure 1.2). Degradation was achieved by selecting the percentage of boxes appropriate to the condition (25%, 50%, 75%, or 100%) and randomly recoloring these black or white. This target-object was easily segregated from the background when minimally degraded, but was indistinguishable at the maximum degradation level (Figure 1.2).

Experiment II was designed to study the post-saccadic detection of displacements of an explicit object, and to what extent its segregation from its background was important. Hence, to exclude low-level motion detecting mechanisms, and focus on the object-identity itself, the whole screen images were intrasaccadically-extinguished on every trial. Image contrast was again inverted on a randomly interleaved 50% of trials. There were a total of 400 trials: 25 trials in each of 16 conditions (4 levels of target degradation x 2 possible contrast inversion states (contrast-unchanged or contrast-inverted) x 2 intra-saccadic shift amplitudes).

Analysis & Statistics

Data were analyzed using a purpose-written interface in Matlab (The Mathworks Inc., Natick, MA). Saccades were first detected automatically using a 10°/s velocity threshold, and confirmed by visual inspection. We discarded a small portion of trials due to blinks, or when the primary movement was smaller than 0.67 of the target eccentricity. We measured the first saccade after this primary movement with a minimum amplitude criterion of 0.4°. This amplitude criterion served to reduce three categories of false measurement: 1) back-and-forth microsaccades, which are typically not “goal-directed”, 2) unavoidable recording noise that mimicked a saccade, 3) small saccades correcting for execution noise in the planned vertical primary saccade. For this last category, it should be noted that the off-axis endpoint error is typically much smaller than in the direction of the saccade vector (van Beers, 2007). Nonetheless, throughout, we term the measured secondary saccades as “corrective”, in the sense of correcting for the intra-saccadic stimulus shifts, which is further justified by the there being a strong correlation between the shifts and saccade amplitudes. Trials without any such corrective detected were excluded from further analysis.

We assessed tracking performance by comparing shift-direction to corrective saccade direction separately for each shift-amplitude (-1°, -0.5°, +0.5°, and +1°), and computing the proportion of corrective saccades in the same direction as the shift. Finding no effect of shift-amplitude on performance, we collapsed across amplitudes and computed a single proportion to quantify performance.

Proportions were used to fit binomial distribution parameters and 95% confidence intervals (CIs) via the Clopper-Pearson method (Johnson et al., 1993). This method, which computes an exact interval based on the continuous distribution function of the binomial distribution, allows the determination of the statistical significance (at a given α-value) of the comparison between a value (between 0 and 1), and the proportion used to fit the CI. For example, the 95% CI fitted to 5 "hits" in 10 trials (5/10) is [0.19,0.81], this indicates that this proportion is significantly different from all values smaller than 0.19 or greater than 0.81. Consequently, when tests of statistical significance relied on this type of comparison (between a value and a CI), we indicate significance or lack thereof in the text merely by quoting the α-value used to compute the CI. Importantly, intervals computed in this way cannot be used to compare two (or more) proportions; to test the statistical significance of such comparisons, we used the χ^2 proportion test (as indicated in the text).

Results

Tracking background shifts

Experiment I examined the ability to track horizontal stimulus shifts, triggered by downward primary saccades, with corrective saccades. We compared tracking performance as 3 stimulus properties were systematically varied (8 conditions, 2 x 2 x 2 design; Figure 1.1). These properties were: features-present or features-absent, continuously-illuminated or intrasaccadically-extinguished, and contrast-unchanged or contrast-inverted intrasaccadically. For a full description of each of these, conditions, see Methods. Briefly, features-present stimuli consisted of many dark and light ellipses on a neutral background while features-absent stimuli were generated by phase-scrambling features-present stimuli. Intrasaccadically-extinguished stimuli were replaced with a uniform-gray field from detection of saccade start to saccade completion; continuously-illuminated were unmanipulated in this regard. Contrast-inverted stimuli had pixel intensity values switched, high-to-low, concurrent with the horizontal shift; contrast-unchanged stimuli pixel values remained unmodified. We have used the relative levels of corrective saccade performance in these conditions as a means to discern the types of mechanisms that might be available to the oculomotor system in tracking stimulus shifts.

Features Present

When the stimulus contained features, performance was high in tracking trans-saccadic stimulus shifts; even when the stimulus was extinguished during the saccade, most corrective saccades were in the direction of the shift (Figure 1.3A, top panel). Across subjects, shifts of the continuously-illuminated stimulus resulted in corrective saccades in the direction of the stimulus shift on 97% (135/139) of trials, compared to 92% (117/127) on intrasaccadically-extinguished trials (Figure 1.3B, blue boxes), a non-significant difference ($\chi2$ proportion test, $\chi2 = 3.32$, $p = 0.07$). Across and within subjects, performance was significantly above chance: the 95% confidence intervals plotted in Figure 1.3B (blue boxes) do not overlap the 50% performance line ($\alpha = 0.05$; see Methods). To probe this result further, we included an additional manipulation: the reversal of stimulus contrast concurrent with its shift. If tracking performance is based on matching local stimulus features across the saccade, this manipulation might be expected to reduce performance to chance levels (50%).

Surprisingly, contrast reversal nearly inverted performance, driving the majority of corrective saccades in the direction opposite the shift (Figure 1.3A, bottom panel). Across subjects, continuously-illuminated shifts drove corrective saccades in the shift direction on 11% (15/137) of trials, compared to 26% (35/136) on Intrasaccadically-extinguished trials, a significant difference (χ^2 proportion test, $\chi^2 = 9.97$, $p < 0.01$; Figure 1.3B - red boxes). Importantly, a trans-saccadic correlation mechanism would not be expected to have significant differences between the continuously-illuminated and intrasaccadically-extinguished conditions. The prediction that this mechanism would produce chance performance was also clearly not upheld, as, across subjects, performance was now significantly below chance ($\alpha = 0.05$); the same was true at the individual level, except subject 2 in the intrasaccadically extinguished condition (Figure 1.3B, second gray box from left). While this near inversion of performance is not consistent with a trans-saccadic comparison of local feature position, it is consistent with the extraction of a motion-energy signal from these stimuli. Moreover, a motion-energy signal would likely be compromised by extinguishing the stimuli intrasaccadically, consistent with the significant differences we found between these conditions. If this was the sole mechanism, why did the intra-saccadic extinction not reduce performance to chance?

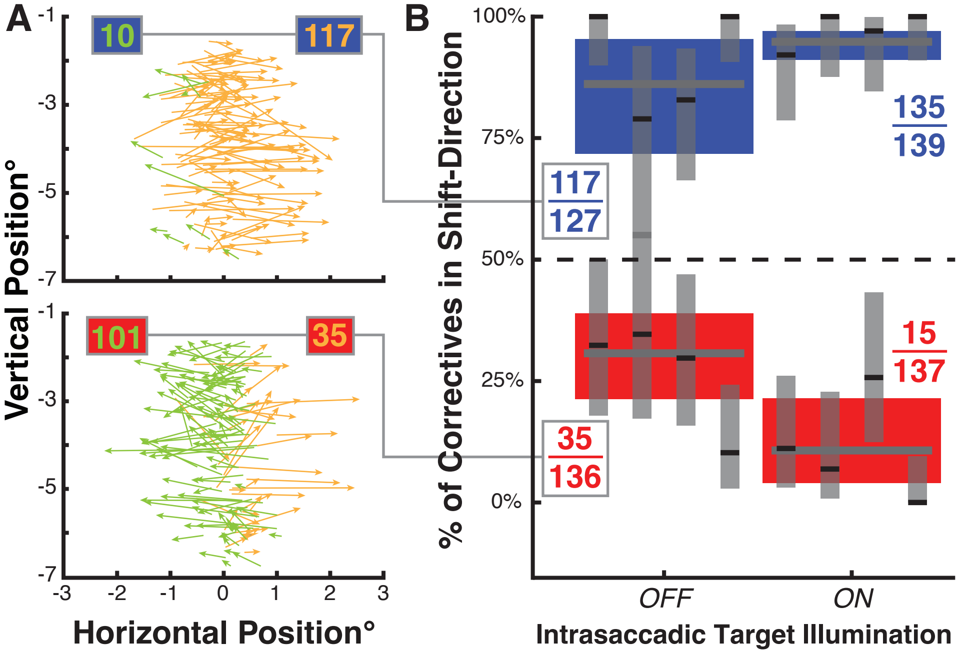

Figure 1.3 Experiment I Features Present Data. **A**. Raw corrective saccade vectors (arrows) pooled across subjects in the intrasacadically-extinguished (top) and intrasaccadically-extinguished + contrast-inverted (bottom) conditions. **Yellow arrows** are movements in the same direction as the stimulus shift, and **green arrows** are against the shift. As indicated by color-filled boxes and grey line segments, top panel data correspond to the upper-left data in B, and bottom panel to the lower left data. **B**. Proportions of Corrective Saccades in Shift Direction. Bernoulli parameter estimates and 95% confidence intervals (CIs) for individual subjects are represented by black line segments and gray rectangles, and were calculated by the Clopper-Pearson method (for details, see Methods). Maximum likelyhood across-subject estimates (gray line segments with **blue boxes for contrast unchanged** and **red boxes for contrast-inverted**) were computed from the product of individual subject probability density functions. Dashed black line represents chance performance (50%). CIs may be used for comparison to a single value: tracking performance in **contrast-unchanged conditions** is significantly greater than chance (at α = 0.05) since **blue CIs** do not overlap the chance-line and in **contrast-inverted conditions** is significantly smaller than chance. Colored fractions indicate the number of correctives in shift-direction (numerator) out of the total (denominator) across subjects.

Figure 1.3 – Chapter 1 / Experiment I Features Present Data

We were curious whether a second feature-dependent mechanism was involved, or whether some low-level stimulus property (such as spatial frequency content of the image) was sufficient to explain our results. Thus, we also included a set of conditions in which the stimuli were instead phase-scrambled versions of the feature-rich stimuli used above (see Methods).

Features Absent

When the stimulus was continuously illuminated, tracking performance with features-absent stimuli was identical to that with features-present stimuli. Both in contrast-unchanged and in contrast-inverted conditions, the proportions of corrective saccades in the shift-direction elicited by features-absent stimuli: 93% (128/137) and 12% (16/130) respectively (Figure 1.4, right part), were indistinguishable from those elicited by the features-present stimuli shown in Figure 1.3: 97% (128/137) and 11% (15/137), ($\chi2$ proportion tests, contrast-unchanged: $\chi2 = 2.10$, $p = 0.15$, contrast-inverted: $\chi2 = 0.12$, $p = 0.73$). This similarity of performance irrespective of feature presence is consistent with a low-level motion energy mechanism operating intrasaccadically. Our results with features-present stimuli suggested an additional mechanism that operates trans-saccadically, is feature presence necessary for its expression?

When features-absent stimuli were extinguished intrasaccadically, tracking performance was strongly affected (Figure 1.4, left). With contrast-unchanged, performance was weaker in this condition (73%, 93/126) than in either the continuously illuminated features-absent (93% 128/137; $\chi^2 = 18.83$, $p < 0.001$) or the intrasaccadically extinguished features-present conditions of Figure 1.3 (92%, 117/127; $\chi^2 = 15.04$, $p < 0.001$). Meanwhile, with contrast-inverted, the performance was close to chance in the intrasaccadically extinguished, features-absent condition (53%, 62/116), which was again significantly different from both the continuously illuminated features-absent (12%, 16/130; $\chi^2 = 47.92$, p = 4.45 x 10-12) and the intrasaccadically extinguished features-present conditions of Figure 1.3 (26%, 35/136; $\chi^2 = 20.31$, p = 6.59 x 10-6). Additionally, though contrast-unchanged performance across subjects was significantly above chance ($\alpha = 0.05$), this was largely due to the strong performance of subject 4 (88%, 35/40) and subject 3 (73%, 24/33), who were each significantly above chance ($\alpha = 0.05$); subjects 1 (65%, 24/37) and 2 (63%, 10/16) displayed performance indistinguishable from chance. The significantly impaired and variable performance in tracking features-absent stimuli trans-saccadically compared to features-present stimuli suggests that any low-level (feature-independent) trans-saccadic shift detection mechanism is weak at best. It seems that trans-saccadic tracking does require the presence of features to be useful on a saccade-by-saccade basis.

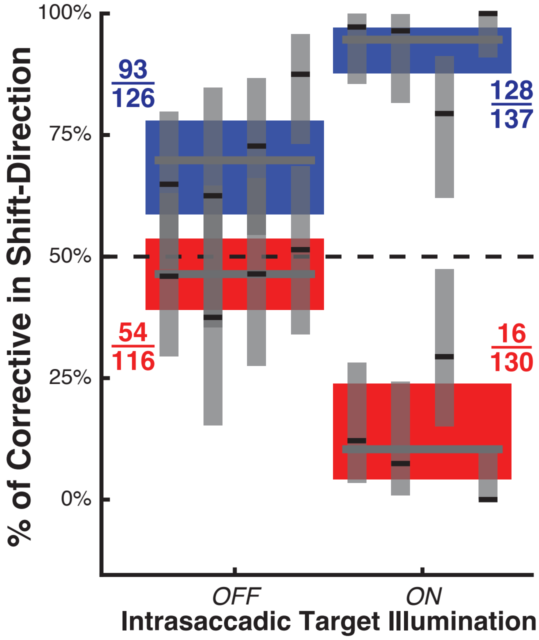

Figure 1.4 Experiment I Features Absent Data. Conventions as in Figure 1.3. Proportions of Corrective Saccades in Shift Direction. Bernoulli parameter estimates and 95% confidence intervals (CIs) for individual subjects are represented by black line segments and gray rectangles. Maximum likelyhood across-subject estimates are represented by gray line segments with **blue boxes for contrast unchanged** and **red boxes for contrast-inverted**. Dashed black line represents chance performance (50%). CIs may be used for comparison to a single value: tracking performance in the **contrast-inverted** + intrasaccadically-extinguished condition is not significantly different from chance (at α = 0.05) since the **red CI** overlaps the chance performance line. Colored fractions indicate the number of correctives in shift-direction (numerator) out of the total (denominator) across subjects.

Figure 1.4 – Chapter 1 / Experiment I Features Absent Data

We were curious whether the expression of this mechanism – which resembles a motion-energy-sensitive one in its response to combined displacement and contrast-inversion – would persist in the presence of an identifiable saccade target embedded in the scene. Is the trans-saccadic shift-detection mechanism we observed a manifestation of tracking whole-field shifts or does it play a role in the tracking of any trans-saccadic displacement?

Tracking combined target/background shifts

Experiment II also examined performance in tracking horizontal stimulus shifts triggered by vertical primary saccades. Primary saccades were triggered by the appearance of a partially degraded object, embedded in an already present background, both of which consisted of black and white boxes. We now compared tracking performance as 2 stimulus properties were systematically varied (8 conditions, 4 x 2): percent degradation, and contrast-unchanged or –inverted (as above). Degradation was systematically varied (25%, 50%, 75%, or 100%) by randomly recoloring (black or white) a percentage of the boxes comprising the object prior to its appearance. In these sessions, stimuli were always intrasaccadically-extinguished.

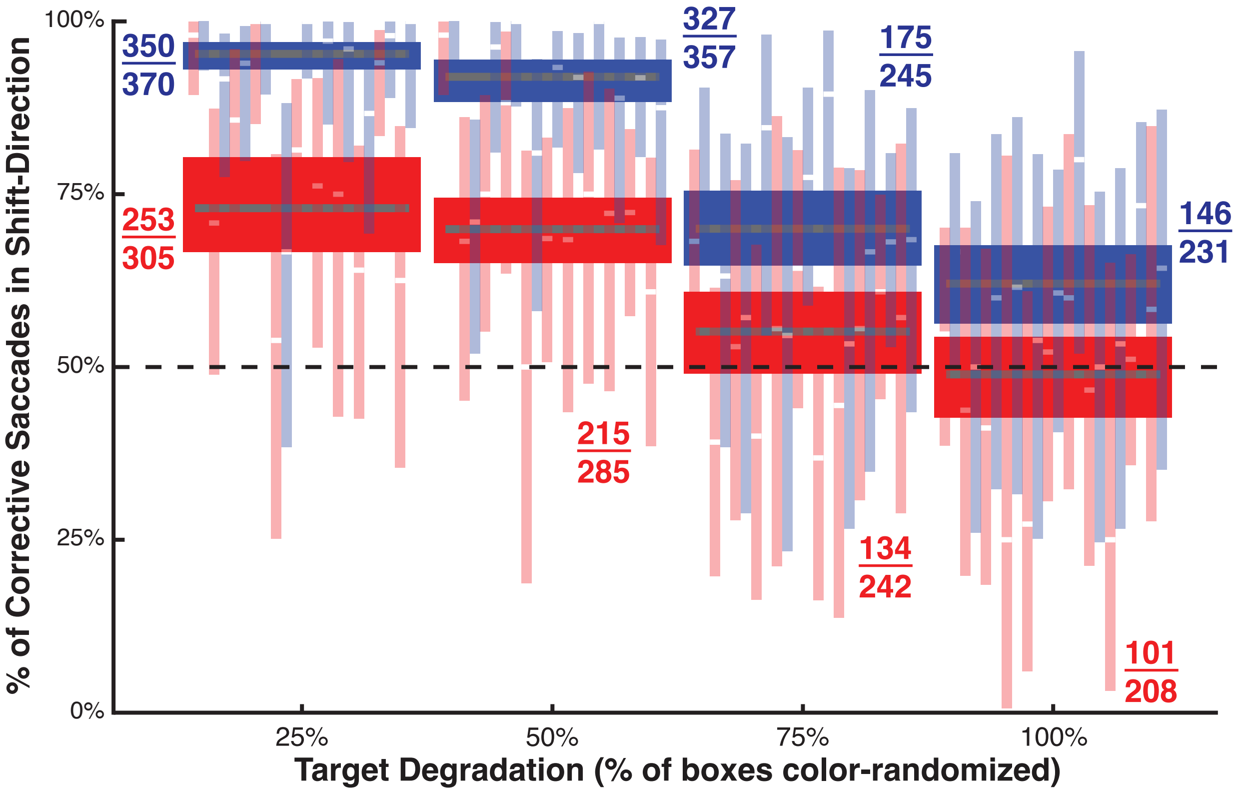

Figure 1.5 Experiment II Data. Proportions of Corrective Saccades in Shift Direction. Bernoulli parameter estimates and 95% confidence intervals (CIs) for individual subjects are represented by semi-transparent white line segments and **blue** (**contrast-unchanged**) or **red** (**contrast-inverted**) rectangles, Across-subject estimates are indicated by gray line segments with solid **blue boxes for contrast unchanged** and **red boxes for contrast-inverted**) were computed from the product of individual subject probability density functions. Dashed black line represents chance performance (50%). CIs may be used for comparison to a single value: for example, unlike the Features Present **contrast-inverted** + intrasaccadically-extinguished condition of Experiment I (Figure 1.3, left, **red box**), tracking performance in the contrast-inverted 25% target degradation condition of Experiment II (above, left) was significantly greater than chace, as indicated by the red CI not overlapping with the chance-line. For a discussion of why contrast inversion led to a reversal of tracking behavior in some conditions of Experiment I but not in Experiment II, see the text. Colored fractions indicate the number of correctives in shift-direction (numerator) out of the total (denominator) across subjects.

Figure 1.5 – Chapter 1 / Experiment II Data

When the saccade target was degraded to the point of resembling the background (100%), performance was similar to that obtained with features-absent stimuli of Experiment I (Figure 1.5). Contrast-unchanged performance in this condition was marginally (though significantly; α = 0.05) above chance (63%, 146/203); contrast-inverted performance was indistinguishable from chance (49%, 101/208; α = 0.05). Comparing these values with those of Experiment I's intrasaccadically-extinguished features-absent condition (Figure 1.4, right), we found the contrast-unchanged (blue) values to be significantly different (63%, 146/203 vs. 73%, 93/126; $\chi^2 = 4.14$, $p = 0.04$), while the contrast-inverted (red) values were not (49% vs. 53%, 62/116; $\chi^2 = 0.71$, $p = 0.40$). However, as noted above, performance with features-absent, contrast-unchanged stimuli was inflated by the performance of one subject. When that subject's data were removed, contrast-unchanged proportions were indistinguishable ($\chi^2 = 0.49$, $p = 0.48$). This finding corroborates the weakness of any putative low-level (not feature-reliant) trans-saccadic displacement detection mechanism suggested by our findings with intrasaccadically-extinguished feature-absent stimuli (Figure 1.4, right).

At the other extreme, when the saccade target was easily identified, the pattern of results resembled neither that of Experiment I's features-present nor features-absent stimuli. Unsurprisingly, with contrast-unchanged (blue), 25% target degradation tracking performance was significantly above chance (Figure 1.5, left; 95%, 353/370, $\alpha = 0.05$), and indistinguishable from that of Experiment I's features-present condition (Figure 1.3, left; 92%, 117/127; $\chi^2 = 1.98$, $p = 0.16$). Surprisingly, however, 25% target degradation tracking performance remained above chance (83%, 253/305; $\alpha = 0.05$) with contrast-inverted (red); significantly smaller than the 95% with contrast-unchanged ($\chi^2 = 28.26$, $p < 0.001$), and dramatically different from the features-present stimuli when contrast-inverted (26%, 35/136; $\chi^2 = 135.90$, $p << 0.001$). Put another way, target presence did not affect tracking performance when contrast was unchanged, but strongly affected performance when contrast was inverted (compared with Experiment I's features-present data). This result suggests that the mechanism that resulted in tracking performance reversal when features-present stimulus contrast was inverted (Figure 1.3, left, red), is not the sole contributor to trans-saccadic shift-detection when an object or features are present. Instead, it seems that when tracking a clearly identifiable object, a distinct mechanism that tracks its position is likely involved. Indeed, such a mechanism need not explicitly compare target positions trans-saccadically, rather the post-saccadic relocalization of a remembered target will suffice to track stimulus shifts in this case.

In summary, Experiment II's results both confirmed those of Experiment I and appeared to reveal another tracking mechanism. Performance with the target degraded such that it resembled the background was indistinguishable from Experiment I's features-absent performance, confirming the weakness of any putative low-level trans-saccadic tracking mechanism. Meanwhile, unlike Experiment I's features-present data where performance was reversed (significantly below chance) by contrast inversion, tracking performance was significantly above chance when the 25% degraded target / background were contrast inverted. This suggests the ability to track intrasaccadic shifts of a clearly defined target even when its contrast before and after the saccade is incongruous; a result more in line with post-saccadic bottom-up orienting to a salient stimulus than with a trans-saccadic stimulus tracking mechanism.

Discussion

Our results are most consistent with the presence of multiple mechanisms all of which may contribute to the tracking of intrasaccadically-displaced stimuli. Our first experiment with continuously illuminated stimuli suggested the involvement of a low-level (insensitive to feature presence) motion-energy sensitive mechanism acting intrasaccadically. Fascinatingly, our results with intrasaccadically-extinguished stimuli from Experiment I seemed most consistent with a mechanism that selectively integrated feature information trans-saccadically into a motion-energy-like signal. Experiment II with a minimally degraded target-object suggested that the presence of such an object likely renders irrelevant the mechanisms listed above, whence the system can instead rely on post-saccadic relocalization. Both sets of experiments concluded that a mechanism based on comparing low-level position information trans-saccadically is quite weak at best. These findings both expand and constrain our understanding of the mechanisms that operate intra- and trans-saccadically to detect stimulus shifts.

Intra-saccadic motion-energy

Our report of intra-saccadic shift detection with continuously-illuminated stimuli is consistent with a previous report which documented intra-saccadic detection of the motion of a low-frequency (0.17 cycles/°) grating via a low-level mechanism (Castet and Masson, 2000). Indeed, the notion that a single stimulus-step can excite motion detection apparati is well accepted. Further, the insensitivity to feature presence we document is in keeping with a characterization of the underlying motion detection mechanism as low-level, as both types of stimuli had identical spatial frequency content, the principal determinant of intra-saccadic motion perception as in Castet and Masson (2000). A novelty in our report is the contrast-inversion contingent reversal of intrasaccadically perceived stimulus motion as assessed by tracking performance (Figure 1.3, 1.4). However, it is well documented elsewhere that concurrent stimulus steps and contrast-inversions result in reversal of perceived motion direction (Anstis, 1970; Anstis and Rogers, 1975). This phenomenon, termed "reverse phi-motion" is understood to be entirely compatible with well-accepted theories of low-level motion detection (Adelson and Bergen, 1985; Chubb and Sperling, 1989). Thus, it seems that our findings in this regard merely elaborate upon existing phenomena.

Trans-saccadic feature integration

Our results demonstrate for the first time (to our knowledge) that visual feature information can be integrated trans-saccadically at the service of detecting whole-field stimulus shifts. That features-present contrast-unchanged shifts resulted in predominantly in-shift-direction corrective saccades and contrast-inverted shifts resulted in predominantly against-shift-direction corrective saccades is again suggestive of a motion-detection-like mechanism. That features-absent contrast-unchanged shifts resulted in significantly but not overwhelmingly in-shift-direction corrective saccades and contrast-inverted shifts resulted in ambiguously directed corrective saccades suggests a high-level motion-detection like mechanism. We propose that these results may be the manifestation of the same mechanism responsible for what has been called "attentive tracking" (Cavanagh, 1992; Verstraten et al., 2000), or "third-order motion" (Lu and Sperling, 1995, 1996). These works describe a motion perception mechanism that requires the endogenous allocation of attention to a particular set of features. Our features-present stimuli would allow for such an allocation, but features-absent stimuli would not. In addition, this mechanism is usually described as "slow", occurring at a longer time-scale than lower-order mechanisms, consistent with the longer time-scale required for trans-saccadic versus intra-saccadic integration. This suggestion does not furnish an immediate explanation for the contrast-inversion contingent reversal of corrective saccade direction observed with intrasaccadically-extinguished stimuli, as no previous work on the proposed mechanism has explored this possibility. Determining whether our results truly represent a manifestation of this same mechanism will require further work.

We must acknowledge a caveat regarding our features-absent stimuli. Although the phase scrambling method we used to generate our features-absent stimuli seemed effective in that we found a strong performance difference with these stimuli compared to features-present, it is possible that some feature information (contours, shapes) was not eliminated by this procedure. It may thus be that these stimuli are better classified as strongly noise masked than strictly features-absent. Further work will be required to determine whether this performance difference was due to feature masking or feature absence.

Trans-saccadic correlation

Two previous works (Deubel, 1991) have suggested that a "correlation" mechanism might facilitate the type of trans-saccadic shift-detection we here observe. However, our results suggest that if such a mechanism exists, it is fairly weak, at best. Those studies used stimuli which were (respectively) quite similar to the stimuli we have utilized here; indeed, our features-present stimuli were inspired by those of Ditterich et al. (1999), and Deubel's (1991) stimulus was similar to both our features-absent (Experiment I) and our 100% degraded (Experiment II) stimuli. However, two important aspects of their methodology make it extremely unlikely for a correlation mechanism to have been responsible for their results. First, as their studies were focused on saccade adaptation, they used stimulus shifts that were consistent over hundreds of trials, which would have facilitated the use of a weak correlation mechanism by allowing for the accumulation of evidence in favor of one shift-direction. Once such a determination had been made, no subsequent shift-detection would have been necessary since the direction remained constant. Second, and more importantly, the stimuli in these studies were not extinguished intrasaccadically. Given that our results, and those of Castet and Masson (2000) convincingly demonstrate the presence of an intra-saccadic low-level motion-detection mechanism, we consider it far more likely that this was the mechanism that facilitated detection (to the degree it was necessary) in those studies.

Trans-saccadic object localization

The effect of contrast-inversion on tracking a continuously visible target-object was significant, but did not resemble the effect of that manipulation in other conditions. As mentioned above, we suggest that tracking in this case does not necessarily rely on either intra- or trans-saccadic shift-detection. Instead, tracking shifts of the target-object (and implicitly the scene as well, as they were shifted in concert) is achievable in these circumstances by simply relocalizing the target after the primary saccade. We propose, then, that tracking the contrast-inverted target-object (and background scene) were achieved by the ability to recognize the object irrespective of any contrast-inversion. Indeed, this object is quite distinct even when considered as merely an area of uniform luminance relative to the background; contrast-inversion would thus leave the target easily identified after the primary saccade. While alternate hypotheses regarding performance in this sub-condition are available, they are highly speculative compared to the relative simplicity of the foregoing description. Any further clarification of the mechanism(s) underlying the behavior in this condition must thus await further experimentation.

Relationship to saccade adaptation

Visual mechanisms that direct corrective saccades, as explored in this chapter, would clearly be expected to also drive saccade adaptation. In the Introduction, we outlined three levels of complexity in correcting for post-saccadic targeting errors: 1) bottom-up pop-out stimuli, such as a single spot on a blank background; 2) a well-defined, localizable target amongst a structured background; 3) a target that is poorly defined or difficult to segregate from background. We have already discussed how our data suggest that adaptation in the last level (Deubel, 1991; Ditterich et al., 1999) was most likely due to intra-saccadic motion-energy detection mechanisms. The same motion-energy detection argument can be made for the adaptation caused by intra-saccadic background shifts when a spot target is extinguished upon saccade (Robinson et al., 2000), because the background was illuminated continuously. In contrast, we have some new evidence of adaptation to a background displacement (target extinguished from saccade onwards) that incorporated intra-saccadic masking of the background, which therefore points towards a trans-saccadic feature integration mechanism in this adaptation.

In brief, in collaboration with Laurent Madelain in France, we have recently finished four adaptation experiments that explored the effects of target and natural-image background shifts. Target and background were displaced trans-saccadically: independently, together, or with the target extinguished upon saccade. As just mentioned, the target-extinguished condition gave significant adaptation to background displacement that could not be detected by a mechanism operating intrasaccadically. In the other three conditions, when there was a clearly defined target post-saccade, the background had no effect, and the adaptation followed the target displacement. This corresponds to the second level in the target localization framework above, and the data are consistent with this chapter's Experiment II that showed the primacy of trans-saccadic object localization when motion-energy mechanisms cannot be invoked. These adaptation conditions are also consistent with previous reports that background displacements have no effect when a target is present (Deubel, 1991; Ditterich et al., 1999; Robinson et al., 2000).

Two novel design elements in these new experiments were that the rectangular target was a translucent color highlighting of the natural scene image, and subjects were required to attend to and report the orientation of the rectangle. Hence, location error signals related to the rectangular object were put into competition with error signals related to the underlying image features. As in Experiment II above, the object identity trumped other features. When the target was present, subjects were successful in filtering out the task-irrelevant information of the underlying image features. This suggests a key role for target selection, and attention in saccade adaptation.

Finally, these most recent data, suggesting a role for target selection in adaptation, are a logical extension of those in the next chapter, which were a first step away from the first level of post-saccadic target localization outlined above. Most saccade adaptation studies require only correcting errors of a single target displacement on a blank background, and so simple retinal error can be used without any explicit target selection. In the real world there are a myriad number of potential error signals from surrounding non-targets, and we asked if similar competing non-targets would interfere with adaptation, or whether these could be successfully ignored.

Chapter 2: Visual selection in saccade adaptation and prediction in correction

Introduction

Saccadic eye movements (or simply “saccades”) are rapid, point-to-point movements that humans make to orient the maximally sensitive foveal region of the retina towards features of interest. Because they are so brief (the duration of a 10° saccade is ~40ms) it is impossible to use visual signals to guide them in-flight, as the earliest of these arrive at relevant oculomotor midbrain structures approximately 40ms after onset (Goldberg and Wurtz, 1972). Consequently, it is thought that the accuracy of saccades is maintained over a lifetime by a mechanism that uses the visual error (difference between fovea and target position) resulting from previous saccades. Indeed, McLaughlin (1967) first noted that surreptitiously displacing a target backwards during each of a sequence of successive saccades both caused those movements to result in negative visual errors, and a progressive reduction in their amplitudes. This amplitude modification, termed saccade adaptation, has since been extensively studied with the intra-saccadic target step (ISS) paradigm introduced by McLaughlin (for a review see Hopp and Fuchs, 2004). Saccade adaptation has been observed in monkeys whose eye muscles have been weakened by tenectomy (Optican and Robinson, 1980), and in humans with weakness due to disease (Kommerell et al., 1976, Abel et al., 1978). Importantly, it appears that the adaptation that results from ISS paradigms is comparable to that resulting from tenectomy, suggesting that a single, visual error mechanism is responsible the maintenance of saccade accuracy (Scudder et al., 1998).

Saccade adaptation is usually studied in the lab with single small (<1°) targets. In this setting, it is clear how the oculomotor system might calculate the visual error resulting from a given saccade: by the retinal location of the bottom-up visual signal from the target. It is no surprise, then, that saccade adaptation is frequently characterized as a passive or automatic mechanism. However, the natural environment is often densely cluttered with salient visual stimuli. It is not readily apparent how visual error might be calculated in such a setting. If saccade adaptation maintains saccade accuracy in everyday life, it must implicitly rely on some – more active – mechanism for post-saccadically identifying and locating the goal of the primary movement.

One clue to the nature of such an active mechanism may lie in the observation that pre-saccadic object selection is a necessary component of saccade adaptation in response to ISSs of discrete targets. Evidence from the study of two-saccade sequences indicates that a secondary movement only compensates for an ISS if the primary movement was made between two objects, rather than within a single object (Vergilino-Perez and Findlay, 2006). It may thus be unsurprising that ISSs can induce adaptation only when a saccade targets a new object, not when it moves the gaze from one portion of an object to another (Collins et al., 2007; but see Lavergne et al., 2011). However, because pre- and post-saccadically presented stimuli were visually identical in these studies, their results do not clarify the criteria by which a post-saccadically confronted stimulus is identified as the pre-saccadic goal.

In a recent paper, we sought to clarify the nature of such post-saccadic identification by confronting individuals with two stimuli after each saccade (Appendix A – Madelain et al., 2010). We used two versions of a modified ISS paradigm: in the first, the pre-saccadic target was simultaneously stepped to a new location and replaced by a distractor; in the second, the pre-saccadic target remained unperturbed and the distractor was presented at the new location. We found that indeed, adaptation occurred selectively in the case when the target was stepped. Our conclusion was that the adaptation we observed relied on the post-saccadic identification of the pre-saccade target-goal. However, because (in a given session) the target always stepped intrasaccadically to a predictable position, it was unclear whether this selection was visual: based on the identity of the target, spatial: based on its predictable location, or some combination of the two.

We hypothesize that the repetitive, spatially consistent ISSs used in most saccade adaptation studies (including our 2010 work in Appendix A) is likely to result in an expectation of finding the target at its post-ISS location following the primary saccade. Further, we propose that such an expectation plays a role both in the adaptation of primary saccades, and in the programming of “corrective” saccades (additional movements that follow most primary saccades).

In order to better understand the putative roles of visual selection and spatial predictability in saccade adaptation and corrective saccade generation, we modified our original paradigm: instead of two possible target locations and stimuli (target and distractor), we used five possible target locations and stimuli (one target and four distractors). This change furnished significantly greater variation in potential post-saccadic target location, and thus the opportunity to manipulate an individual's ability to predict that location. Will unpredictable ISSs (biased either forward or backward) and a target embedded amongst stimuli differentiated only by color and shape result in adaptation; or are such demands too great for the adaptive machinery? If adaptation is observed, is it modulated by the consistency of the ISSs?

We also test the idea that adaptation sessions with consistent ISSs result in expectation- or prediction-based corrective saccades by interleaving a small number of trials with either no ISS – putting the visual location of the target at odds with its expected/predicted location – or in which target identity is eliminated. If ISS consistency results in predictive corrections, we expect that corrective movements on these trials will be erroneously directed towards the expected post-ISS location of the target. We also ask whether the biased – but not predictable – ISSs used in the paradigm described above result in any predictive correction saccades.

Methods

Subjects

Data were collected from 12 subjects, ages 18-35, 6 men, and 6 women, all faculty and students from the City College of New York (CCNY) community with normal or corrected-to-normal vision. Of these, 3 participated in all conditions, and the remaining 9 only some. This study received prior approval by the Institutional Review Board (IRB) of CCNY, and thus complied with all human-subject protocol requirements.

Equipment

Stimuli were displayed on an Iiyama Vision Master Pro 514 CRT display (Oude Meer, Netherlands) at a resolution of 800 x 600 pixels (visible area 41.5 cm x 30.5 cm), and a vertical sync-rate of 200Hz.

Stimulus generation & display, data storage, and overall experimental session orchestration were controlled with custom software made in LabView (National Instruments, Austin, TX) running in Windows XP (Microsoft Corporation, Redmond, WA) on a Dell PC (Austin, TX).

Eye movements and gaze position were measured and collected by an Eyelink-1000 infrared camera system (SR-Research, Mississagua, Ontario, Canada), which sampled right-eye gaze (pupil - corneal reflection) at 1000Hz.

General

Subjects were instructed prior to each session that they would be presented with a target initially on the left half of the screen that would abruptly step to the right. They were told to follow the movements of the target with their eyes, and that at some point during the experiment, other stimuli would appear as well.

Subsequent to receiving instructions, each participant sat in a darkened room, 57cm from the display, and underwent a nine-point self-paced calibration.

Stimuli and Trials

The stimuli used were: a red triangle, a blue circle, a purple square, a yellow hexagon, a green pentagon, and a dark gray octagon (0.5° diameter), presented on a gray background (Figure 2.1A). Only the gray octagon had a unique role, in that it was never used as a fixation or as a target for primary saccades, but was rather only presented after the primary saccade in certain trials (see below). As all of the first five stimuli served as the target in equal amount, we refer to the first five stimuli as the “target stimuli” and the gray octagon as the “neutral stimulus.”

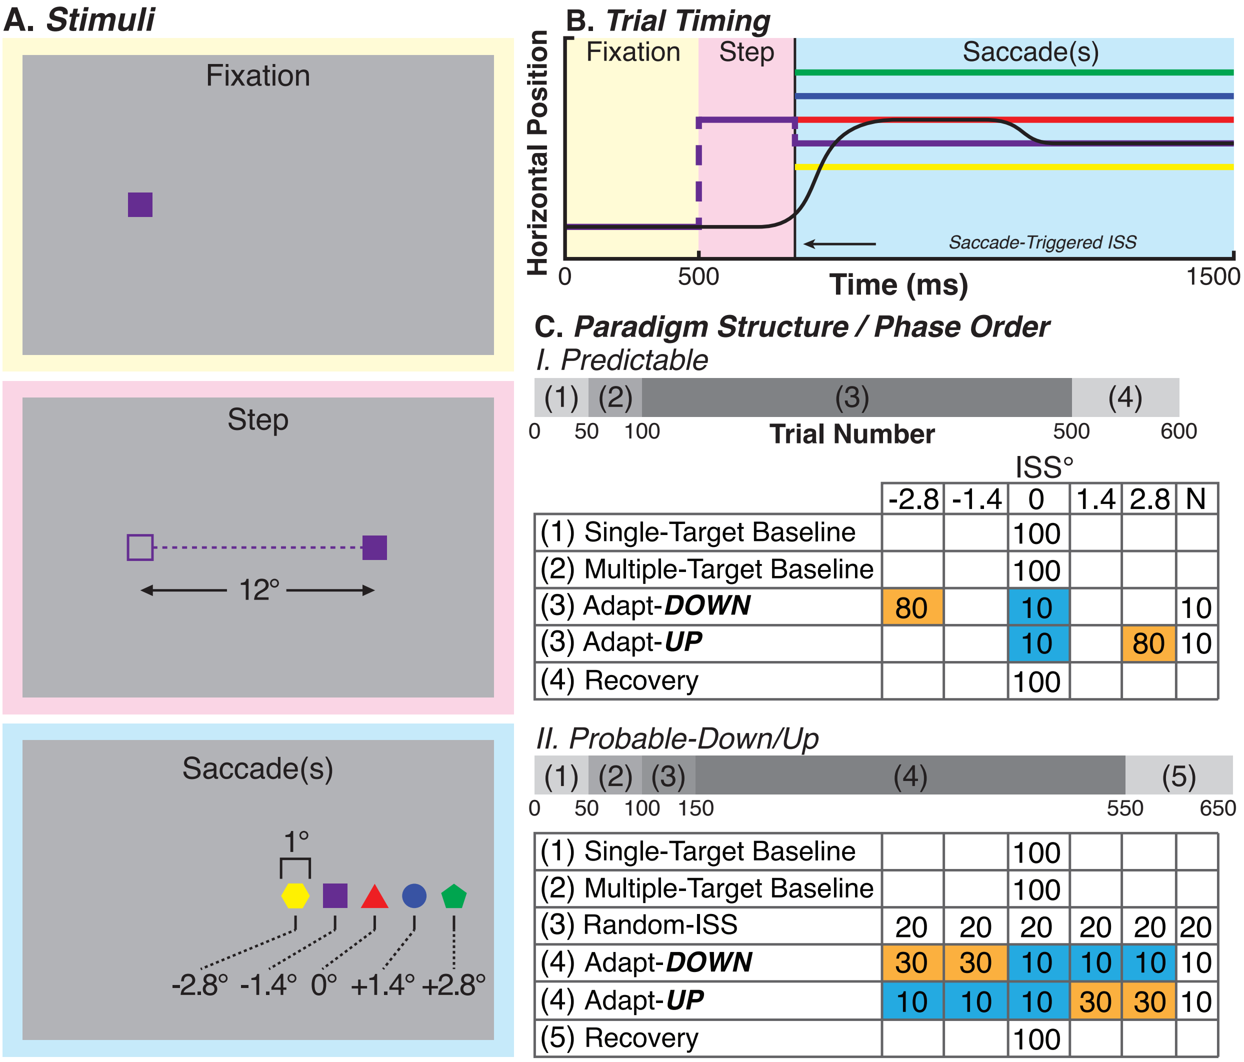

	ISS°					
	-2.8	-1.4	0	1.4	2.8	N
(1) Single-Target Baseline			100			
(2) Multiple-Target Baseline			100			
(3) Adapt-***DOWN***	80		10			10
(3) Adapt-***UP***			10		80	10
(4) Recovery			100			

(1) Single-Target Baseline			100			
(2) Multiple-Target Baseline			100			
(3) Random-ISS	20	20	20	20	20	20
(4) Adapt-***DOWN***	30	30	10	10	10	10
(4) Adapt-***UP***	10	10	10	30	30	10
(5) Recovery			100			

Figure 2.1 Stimuli & Methods. **A**. Spatial representation (not-to-scale) of stimului in a -1.4° ISS multiple-stimulus trial. Border colors correspond to epochs of same shades in B, right. Subjects saw a single target-stimulus at fixation (**purple square**; top) that stepped 12° to the right (middle). The primary saccade triggered an intrasaccadic target step (-1.4° in this example) and the appearance of flanking stimuli appeared (bottom); numbers below stimuli indicate relative horizontal position. **B**. Timing of trial represented in A. **Purple-trace: fixation** / **saccade-target**, black: right-eye gaze, remaining colors correspond to target-stimuli shown in A, bottom. Primary target step always occured 500ms after recording began, and ISS / flanker-onset occured shortly (16 ± 3ms) after saccade start was detected by a 25°/s velocity threshold-crossing. Stimuli remained illuminated until trial conclusion (1500ms). C. Experimental session breakdown by phase. Probable sessions included an extra, "Random ISS," phase (3; bottom) to minimize the detectability of the ISS bias. Trial-type composition tables by phase (numerals on far left). Frequent-ISS trials are highlighted in **orange**, and infrequent in **blue**. Last column header ("N") indicates neutral-stimulus trials (see Methods).

Figure 2.1 – Chapter 2 / Stimuli & Procedures

In all trials, a single target stimulus was presented for fixation on the left (range: -7° to -4°), centered vertically, and stepped 12° to the right (Figure 2.1A). After the primary saccade, there were three possible stimulus configurations (corresponding to the three trial types described below): (1) single-stimulus: the fixation / saccade stimulus alone, (2) multiple-stimulus: a horizontal array of all five target stimuli, equally spaced, and (3) neutral-stimulus: a horizontal array (just as the multiple-stimulus case) but consisting of five copies of the neutral stimulus.

In (1) single-stimulus trials, the target stepped and remained. In (2) multiple-stimulus trials, upon the subject's primary saccade, the primary target stimulus stepped intrasaccadically by -2.8°, -1.4°, 0°, +1.4°, or +2.8° (we refer to these locations as 1, 2, 3, 4, and 5, respectively). Concurrent with the intrasaccadic step (ISS), the remaining four target stimuli were presented (positions chosen at random prior to each trial) at the "other" possible locations. That is, if the primary target stimulus stepped to location 2, locations 1, 3, 4, and 5 were concurrently populated with the remaining target stimuli (Figure 2.1A, B). Note, all ISS trials were multiple-stimulus trials. In (3) neutral-stimulus trials, upon the subject's primary saccade, locations 1, 2, 3, 4, 5 were populated with the neutral stimulus. In all trials, the post-saccadic stimuli remained until recording ceased. Each trial lasted 1500ms with the target step occurring 500ms into the trial; inter-trial intervals ranged from 700-1200ms with a uniform random distribution.

Paradigms

Four paradigms were used: (i) Predictable-Down, (ii) Predictable-Up, (iii) Probable-Down, and (iv) Probable-Up (Figure 2.1C & D). Trial sequence and primary-target positions for each paradigm were fixed prior to all sessions and were identical for all subjects. Importantly, the use of each stimulus as fixation / primary-saccade-target was balanced across trial-types and experimental phases, so that there was no association between a given stimulus and saccade-triggered target behavior.

(i) Predicable-Down had 4 phases: (1) single-stimulus baseline, with 100% (50) single-stimulus trials; (2) multiple-stimulus baseline, 100% (50) multiple-stimulus trials with no (0°) ISSs; (3) adapt, a pseudo-random mixture of 80% (320) frequent-ISS (-2.8°) trials, 10% (40) multiple-stimulus "catch" trials with no (0°) ISSs, and 10% (40) neutral-stimulus trials; and (4) recovery, 100% (100) single-stimulus trials. In (ii) Predictable-Up, the sequencing was identical to (i) but the frequent-ISSs in the adapt phase were +2.8° instead of -2.8°.

(iii) Probable-Down had 5 phases: (1) single-stimulus baseline, with 100% (50) single-stimulus trials; (2) multiple-stimulus baseline, 100% (50) multiple-stimulus trials with no (0°) ISSs; (3) random-ISS, a pseudorandom mixture of 50 ISS trials (20% or n=10 to each ISS position); (4) adapt, a pseudorandom mixture of 60% (240) frequent-ISS trials (-2.8°:120, -1.4°:120), 30% (120) infrequent-ISS trials (0°:40, +1.4°:40, +2.8°:40), and 10% (40) neutral-stimulus trials; (5) recovery, 100% (100) single target trials. In (iv) Probable-Up the sequencing was identical to (iii), but the frequent-ISS trials were onwards (+1.4° and +2.8°) and the infrequent-ISS trials were not (-2.8°, -1.4° and 0°).

Analysis and Conventions

All analyses were performed off-line using software written in Matlab (The Mathworks, Natick, MA). Saccades were detected automatically using a 10°/s velocity threshold, but each saccade was visually checked and corrected if necessary. A small portion (< 5%) of trials were discarded in this initial phase due to blinks or hypometric primary movements (<50% of target eccentricity).

Gain Change Quantification and Statistics

We quantified gain changes starting repeated-measures Analysis of Variance (ANOVA). We then used the Tukey-Kramer method to find the differences (and associated 95% confidence intervals) between all estimated population marginal means (PMMs). For example, in determining the amount of induced adaptation amongst conditions we began with a 2 X 4 X 6 (Phase X Condition X Subject) ANOVA; the two levels in the Phase factor were: (1) the final 50 trials of the adapt phase, and (2) the 50 trials immediately preceding the adapt phase. We then used Tukey-Kramer to compute the difference: (2) – (1), across and within subjects for each condition simultaneously (Figure 2B). Meanwhile, we calculated adaptation as a percentage of the average adapt-phase-ISS similarly, but having first scaled the gain data by the average ISS, and instead using a 2 X 2 X 6 (Phase X Condition X Subject) ANOVA (the 4 conditions now collapsed to just 2: Probable and Predictable).

In determining the effects of the multiple-stimulus baseline and random-ISS phases, we used a single-factor ANOVA with three levels: (1) single-stimulus baseline, (2) multiple-stimulus baseline, and (3) random-ISS phase.

When presenting the results of individual hypothesis tests (as indicated throughout the text) we present p-values and test-statistics. When an ANOVA and subsequent testing was performed, we present the relevant factor's F-statistic and associated p-value once, followed by subsequent-test-specific information. In particular: because the Tukey-Kramer method provides simultaneous confidence intervals (CIs) on all differences amongst PMMs (for a given α), when we mention the significance of individual differences, we specify only the $\alpha = 0.05$ CI(s), and (when relevant) a value for comparison. Similarly, when presenting the results of post-hoc multiple comparison tests of means (gains) or mean ranks (latencies) resulting from ANOVAs, we present only α. Further, when we found no significant difference at $\alpha = 0.05$, we reran the test with increasing alpha values, presenting the smallest alpha value at which the comparison remained non-significant: $\alpha > 0.4$ (for example).

Corrective Saccade Direction

Though frequently it was possible to categorize corrective movements as "towards" or "not towards" the target on the basis of direction alone, there were cases when we made such categorizations on the basis of saccade end-point. For example, in both gain-increase conditions, corrective saccades in 0° ISS trials were predominantly onward (Figure 2.5). In this case, we categorized movements as towards the target (visually-selective) if their end-points were at or below +0.7° (the halfway point between the 0° and +1.4° stimuli).

Results

Primary Saccades - Gain Changes

In general, primary saccade gain changes followed experimental patterns of intrasaccadic target steps (ISSs; Figure 2.2A). The gain changes we observed in the Predictable conditions were qualitatively comparable to our previous result (Appendix A - Madelain et al., 2010). As in more conventional adaptation paradigms, gain changes appeared to proceed more rapidly in gain decrease (Figure 2.2A, I & II), than in gain increase conditions (Figure 2.2A, III & IV). Because the variability of some subject's gains appeared to increase towards the end of adapt-phases in the Probable conditions (for example, Figure 2.2A, III), we compared the variance of the final 50 trials of the adapt phases to the single target baseline (in each condition), using the Two-sample F-test for equal variances. We found significant increases in variance in Predictable-Down (F = 0.58, p<0.001), Probable-Down (F = 0.71, p<0.004), and Predictable-Up (F = 0.77, p<0.03), but not in Probable-Up (F = 0.96, p=0.76).

An Analysis of Variance (ANOVA) revealed significant gain changes during the adapt phase across conditions (F = 183.71, p < 0.001). Significant changes in each condition can be seen in gray boxes representing 95% confidence intervals (CIs) in Figure 2.2B, which do not overlap zero (Tukey-Kramer method, CIs in Figure 2.2B). For the details of our calculations of adaptation magnitude and significance, see Methods. We also found that, as in more conventional paradigms, adaptation was usually incomplete: the change in gain occurring over the course of the adapt-phase was smaller than the average ISS (Tukey-Kramer method, CIs and average ISS values in Figure 2.2B). This was not the case in Predictable-Down, where the amount of adaptation was indistinguishable from the average ISS (Tukey-Kramer method, CI: [0.05, 0.08], average ISS: -0.07). On an individual level, 6/6 subjects adapted significantly in Predictable-Down, with 4/6, 4/6, and 2/6 in Probable-Down, Predictable-Up, and Probable-Up, respectively (Tukey-Kramer method, graphical CIs in Figure 2.2B).

To compare gain changes across conditions, we also expressed changes as a fraction of the average ISS in the adapt-phase. For example, in the Probable-Down condition, ISSs were distributed in the adapt-phase as: 30% -2.8°, 30% -1.4°, 10% 0°, 10% +1.4°, 10% +2.8° (with 10% neutral-stimulus trials); thus average ISS in that condition was (0.3 x -2.8° + 0.3 x -1.4° + 0.1 x 0° + 0.1 x 1.4° + 0.1 x 2.8°) / 12° = -0.07). Interestingly, when represented in this way, we found the average gain changes in the Probable and Predictable conditions to be statistically indistinguishable (Figure 2.2C; ANOVA, Condition effect: F = 0.24, p = 0.62). Hence, the target-specific adaptation we here demonstrate (matching the pattern of target ISSs, unaffected by distractors) does not seem to be modulated by the more regular post-saccadic target location in the Predictable conditions. Although this key finding is somewhat surprising (see Discussion), it suggests that in the present conditions there is a common (and hence presumably visual) selection mechanism underlying the target-specific adaptation.

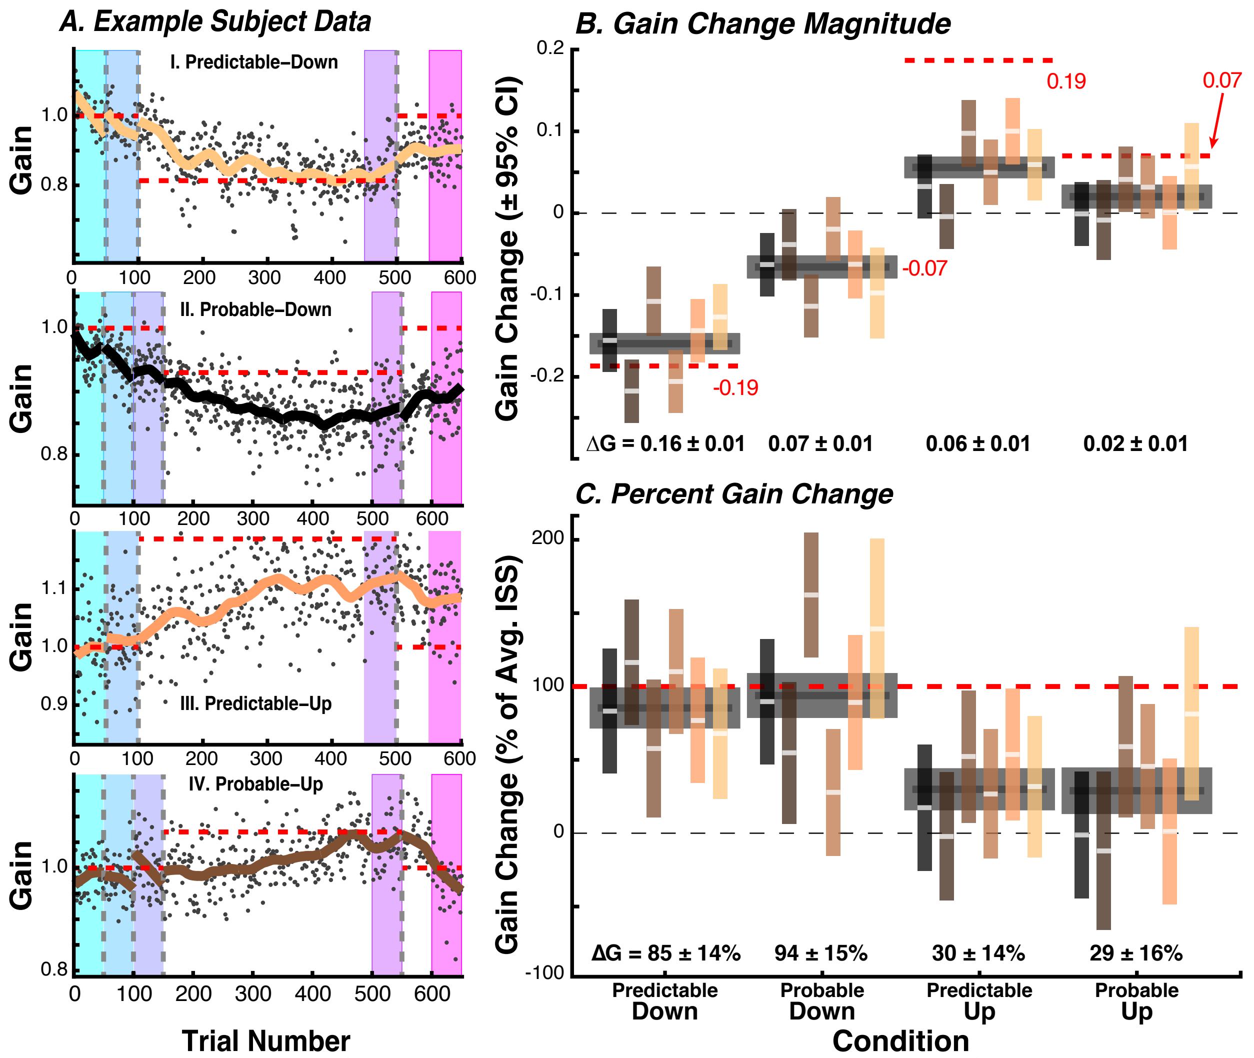

Figure 2.2 Induced Gain Changes. **A**. Example subject data by condition. Curves accompanying individual data-points are robust lowess smooths (40 point windows); curve-color matches colors used for plotting individual subject data in B and C. For example, the data presented for Predictable-Down can be paired with the right-most error-bar in the Predictable-Down group in **B**. Gray and red dashed lines indicate: transitions between phases, and average post-saccadic target position (respectively). Hilighted trial-windows (colored boxes) indicate samples used for gain-change analyses (see Results). B. Gain changes resulting from the adapt phase. These were calculated from estimated population marginal means differences (phase preceding adapt - adapt) after ANOVA using the Tukey-Kramer method (see Methods). Data are plotted both across and within subjects. Large grey boxes are 95% confidence intervals (CIs) on differences across subjects, and dark gray lines are differences; specific values are presented below. Individual subject differences are white lines and colored boxes are accompanying 95% CIs. Red dashed lines and accompanying red numerals indicate the average adapt-phase ISS value. C. Gain changes as a percent of the average adapt-phase ISS were calculated analogously to B by scaling gains by the average ISS prior to an (independent) ANOVA.

Figure 2.2 – Chapter 2 / Induced Gain Changes

Finally, we found that gain varied significantly amongst the single-stimulus baseline, the multiple-stimulus baseline (which contained no ISSs) and the random-ISS phase (present in Probable conditions only; ANOVA: F = 84.93, $p < 0.001$). Gain decreased significantly in the multiple-stimulus baseline by 0.014 (Tukey-Kramer method, CI: [0.009,0.019]), and then in the random-ISS phase by 0.022 (CI: [0.016,0.029]).

Secondary Saccades

The surprising finding that the regularity of post-saccade target location has no influence on adaptation (Fig. 2.2C), led us to conclude that: if our Predictable paradigm's ISSs do influence post-saccadic expectations regarding target location, such expectations do not modulate the adaptation process. Is it the case, then, that no such expectation is formed during adaptation? Or rather that the strength of any such expectation does not play a direct role in adaptation? We sought to further delineate this process by closely examining corrective saccade behavior. For example, can we show that corrective saccades were also purely driven by visual selection, as opposed to expectation, mechanisms?

Overall timing

Comparing conditions exposed a clear effect of target predictability on corrective saccade latency (time from primary saccade completion) during the adapt phase. Principally, the median latency of correctives from both Predictable conditions (181 ± 1.6ms) was significantly smaller than those from both Probable conditions (215 ± 1.7ms; Wilcoxon rank sum test, Z = 26.24, $p<0.001$). As can clearly be seen from the dark gray confidence intervals in Figure 2.3, Predictable-Down (202.5 ± 3.9ms) and Up (174±1.3ms) conditions have significantly smaller latency than their respective Probable conditions (223 ± 2.5ms & 208 ± 2.2ms; Kruskal-Wallis ANOVA, $\chi 2 = 1080.26$, $p < 0.001$, $\alpha = 0.001$). Median latency was also significantly larger in the gain-decrease conditions (215 ± 2.2ms), than in gain-increase (186 ± 1.4ms; Wilcoxon rank sum test, Z = 20.30, $p < 0.001$). Onwards correctives have previously been reported to be shorter latency than backwards (Prablanc and Jeannerod, 1976). Hence it seems both: (a) ISS regularity (which might facilitate the formation of clear expectations about final target location), and (b) the increased likelihood of onwards correctives (present in Up versus Down conditions), resulted in shorter median latencies.

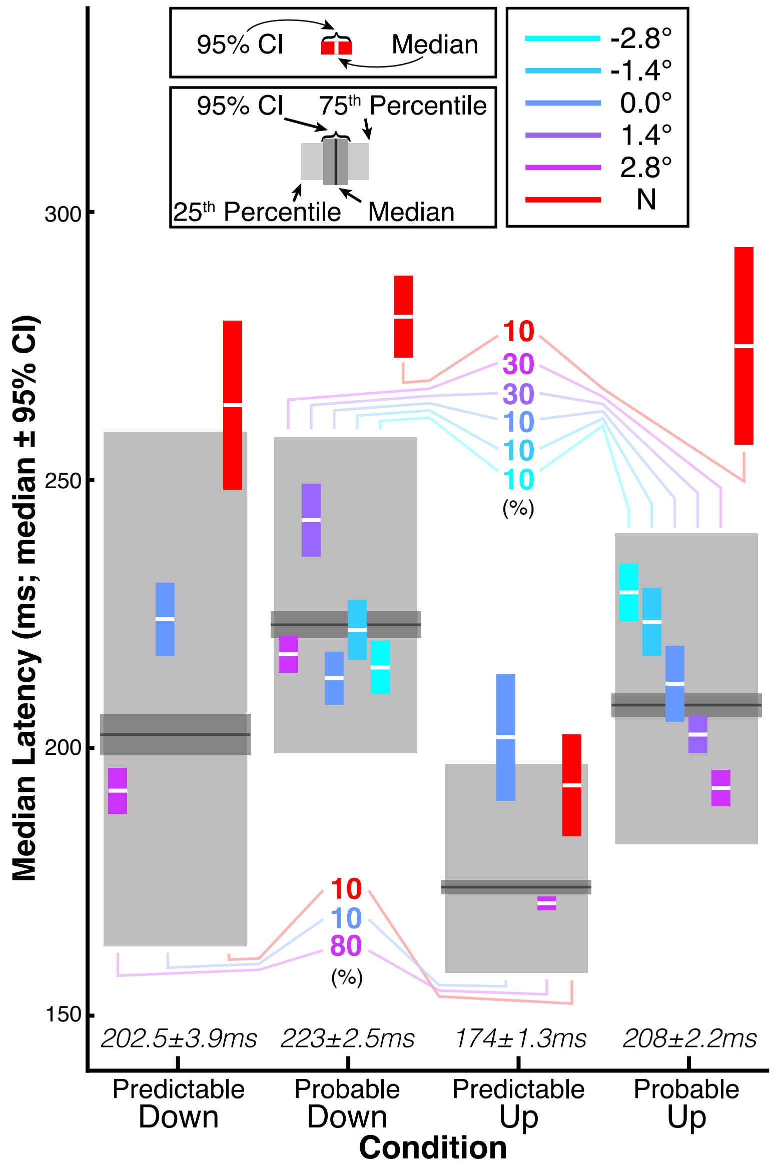

Figure 2.3 The effects of post-saccadic target predictability on corrective saccade latency during adaptation. Latency data from the adapt-phase are plotted across subjects for each adaptation condition (grey boxes), and for each ISS in each condition (colored boxes). Black line segments with surrounding dark grey boxes and light grey boxes represent medians, 95% confidence intervals (CIs), and 25th through 75th percentiles of underlying distributions (respectively), for each condition. White line segments and colored boxes represent medians and 95% CIs for within-condition ISS-values. Black italic numerals below large grey boxes are within-condition medians, and colored numbers correspond to the percentage of adapt-phase trials that contained each ISS value, in each condition.

Within conditions, the more likely post-saccadic target locations also resulted in reduced latencies (Figure 2.3, colored boxes). In both predictable conditions, median latency was significantly shorter in the more frequently (80%) occurring ISS trials as compared to either catch (ISS = 0°) or neutral-stimulus trials (Kruskal-Wallis ANOVA, $\chi 2 = 1565.99$, $p < 0.01$, $\alpha = 0.001$). Similarly, in Probable-Up, median latency in more frequently occurring ISS trials (30% each at +1.4° and +2.8°) was significantly smaller than in less frequently occurring ISS trials (10% at -2.8°, -1.4°; $\alpha = 0.001$); +2.8° but not +1.4° ISS trial latency was also smaller than 0° ($\alpha = 0.001$). The exception to this pattern was in the Probable-Down condition. Not only was median latency in the more frequently occurring -2.8° indistinguishable from the less frequently occurring 0°, +1.4°, and +2.8° ISS trials ($\alpha > 0.88$), but the frequently occurring -1.4° ISS trials had significantly greater latency than other ISS trials ($\alpha = 0.001$). These exceptions can, however, be explained by the joint consideration of corrective timing, direction, and amplitude (see Discussion).

Timing and Direction

Our adaptation data support the idea that post-saccadic retinal error is used in a visually selective manner, uninfluenced by predictability, whereas corrective saccade timing suggests that expectations regarding post-saccadic target location are formed during adaptation. However, we found a further, striking indication of the interplay between visual and expectation-based selection by examining corrective saccade direction. In certain circumstances, we observed clear evidence of correctives that were either "visually-selective" (i.e. to the veridical target location when the target is present) or "expectation-based" (i.e. do not go to the target, instead being influenced by the expected target location). Validating such a distinction was the finding that movements categorized as expectation-based generally occurred earlier than visually-selective ones. Perhaps unsurprisingly, Predictable conditions yielded more expectation-based movements than did Probable. Also, the relative propensity to produce the two categories of corrective saccades depended on adaptation-direction: expectation-based were more likely in gain-increase paradigms, while gain-decrease yielded more visually selective movements. Our categorization was based on both direction and amplitude (see Methods). Importantly, correctives not directed towards the target (in trials with target present) were invariably followed by subsequent movements that aligned gaze and target.

A representative subject's data illustrates our categorization of corrective saccades, and their relative proportions, in *Predictable-Down* (Figure 2.4A). In catch (0° ISS) trials, 75% (30/40) of correctives were "visually-selective" movements (towards the target; red) and 25% (10/40) were "expectation-based," occurring at shorter latency and directed towards the -2.8° location (blue; Figure 2.4A, middle). Neutral-stimulus trials also resulted in a bimodal distribution of corrective directions: 48% were expectation-based (shorter latency, and towards -2.8°) and 52% were visually-selective (longer latency, towards the 0° location occupied by the target prior to the primary saccade; Figure 2.4A, right). (Note, we have chosen to label movements towards 0° in neutral stimulus trials as visually-selective for the sake of simplicity in the remaining Results, and because the original target was at that location and so these may be based on a visual working memory trace; the pattern of results is also consistent with the other visually-selective correctives. But we acknowledge our pragmatic approach to this subset of the data.) Meanwhile, in the most frequent, -2.8° ISS trials, we observed a unimodal distribution of target-directed correctives (Figure 4A, left). However, since the target location and its most frequent location were the same in these trials, we consider it likely that this population contained both visual-selective and (at the shortest latencies) expectation-based movements.

A group level analysis of the *Predictable-Down* condition largely recapitulated the example subject's pattern of results (Figure 2.4B). Subjects made only target-directed correctives in -2.8° ISS trials (Figure 2.4B, left), whereas in catch trials (Figure 2.4B, middle) they made 20% expectation-based correctives and 80% visually-selective correctives (see also Figure 2.6 for proportion data). In neutral-stimulus trials, we found 60% visually-selective correctives, and 40% expectation-based correctives, despite one subject making only visually-selective movements (Figure 2.4B, right). When both categories were present (catch and neutral stimulus trials), median latencies of the two were significantly different across and within subjects (Kruskal-Wallis ANOVA, $\chi 2 = 217.04$, $p < 0.001$, $\alpha = 0.05$), except for one individual in neutral-stimulus trials ($\alpha > 0.3$; Figure 2.4B, right, subject 4).

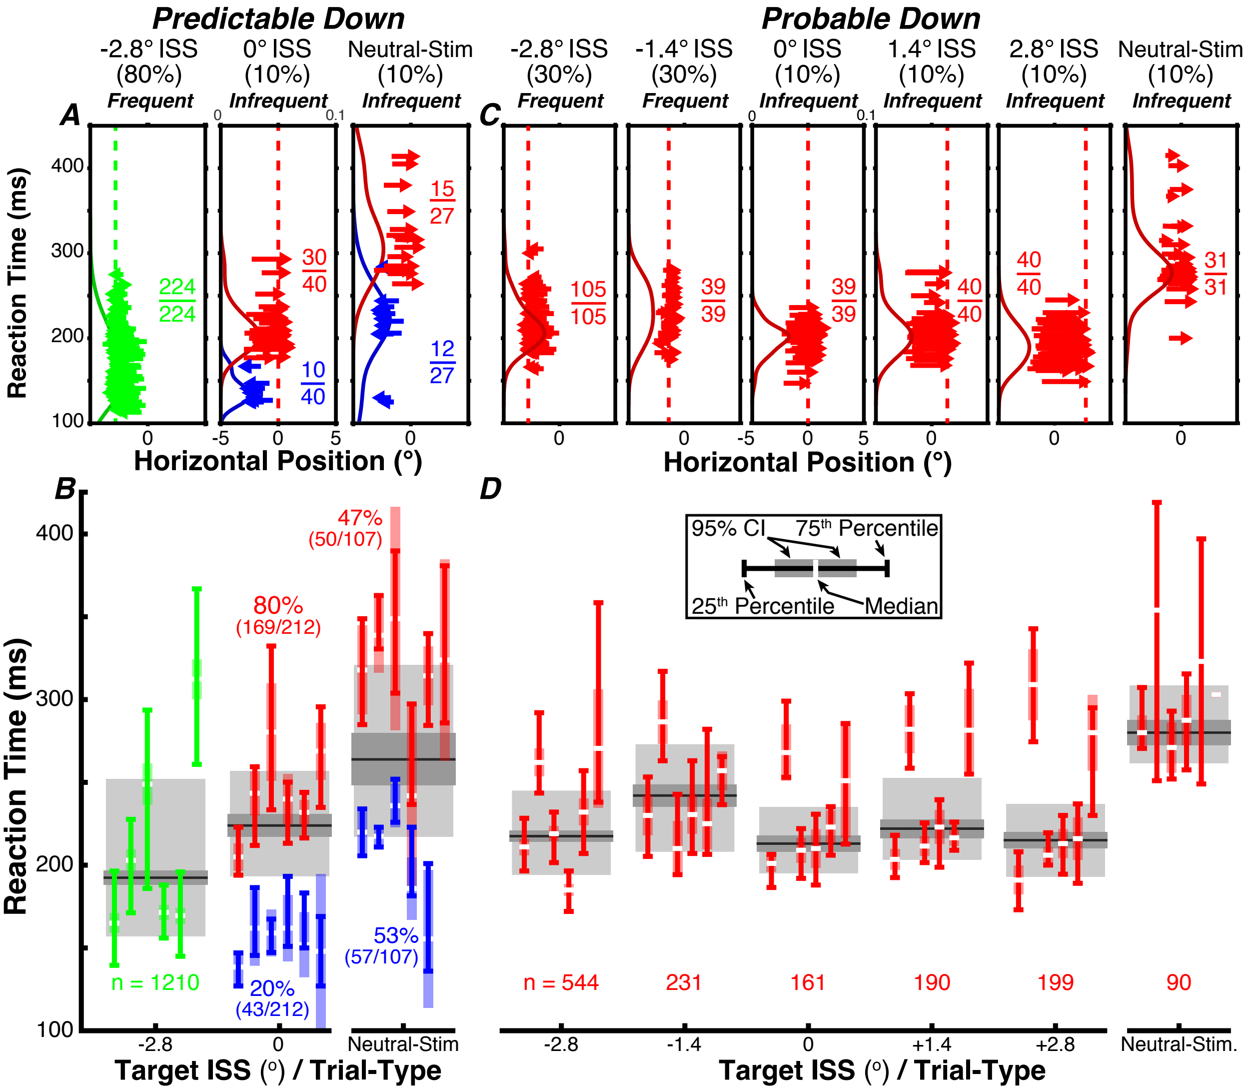

Figure 2.4 Latency Effects on Secondary Saccade Production During Gain-Decrease Adaptation. **A.** Predictable-Down example subject data. Arrows are plotted from secondary saccade horizontal start-to-end (x-axis) at latency (ms) relative to primary saccade completion (y-axis). Curves are latency histograms smoothed with gaussian kernels, and refer to top axis values (kernel width chosen to guarantee a single peak; for visualization purposes only). Vertical dashed lines illustrate post-ISS target position. Separate plots for -2.8° (left) and 0° (center) ISS, as well as neutral-stimulus trials (right). Left: green saccades are directed towards the target; center: blue are directed away from the target, and red towards it; right: blue are towards -2.8° and red are towards 0° stimulus. These data are presented, below, in B as the left-most of each colored set of symbols. **B.** Predictable-Down summary data (across and within subjects). Dark and light gray boxes and black line represent inter-quartile range (25th-75th percentile), 95% median CI, and median (respectively). Colored error-bars, boxes, and white line segments represent corresponding values for individual subjects. Colors correspond to populations illustrated in A, and left-most in each group is from example subject data presented in A. **C.** Probable-Down example subject data. Plotting conventions as in A. As in A (left), green saccades are directed towards the target. Data are again presented as left-most of each set of green symbols in D, below. **D.** Probable-Down summary data. Plotting conventions as in B. Median CI boxes are absent in cases with fewer than 3 data points. Far right: this subject made only one corrective saccade in this trial type, the white line-segment indicates the latency of that movement.

Intriguingly, in *Probable-Down*, we found 100% visually-selective corrective movements. All were directed towards the target (ISS trials) or the target's location prior to the primary saccade (neutral-stimulus trials; Figure 2.4C & D). This trend is evident at an individual level (Figure 2.4C) as well as across subjects (Figure 2.4D). Though one might conclude that the lack of expectation-based correctives in this condition is due to a lack of short-latency movements, we contend the converse: that the lack of short-latency movements is due to the lack of any expectation-based correctives (see Discussion).

Predictable-Up yielded a large proportion of expectation-based saccades, as well as some visually-selective. Across and within subjects (with one exception), median expectation-based corrective latency was again smaller than visually-selective (Kruskal-Wallis ANOVA, $\chi 2 = 205.56$, $p < 0.001$, $\alpha = 0.05$; Figure 2.5B). Perhaps unsurprisingly, in Predictable-Up, correctives in the frequent ISS trials were not clearly divisible, though we categorized them on the basis of landing point (see Corrective Saccade Direction, Methods). Interestingly, Predictable-Up resulted in a greater relative proportion of expectation-based correctives, while Predictable-Down had more visually-selective. For example in catch and neutral-stimulus trials, 75% and 52% (respectively) of a representative subject's correctives were visually-selective in *Predictable-Down* (Figure 4A, middle, right), while 81% and 100% were expectation-based in *Predictable-Up* (Figure 5A, middle, right). This increase in expectation-based correctives was also present across subjects in Predictable-Down (Figure 5B), where we found values of 48% and 93% (catch and neutral-stimulus trials, respectively). As in *Predictable-Down*, expectation-based movements occurred significantly earlier in *Predictable-Up* (when both categories were present):

In Probable paradigms, gain-decrease yielded no expectation-based correctives, and *Probable-Up* yielded comparable proportions to *Predictable-Down*. To detail the similarities between *Probable-Up* and *Predictable-Down*, we selected the conditions in those paradigms that had equal frequencies of presentation (i.e. the 10% cases of Catch, Neutral, and the infrequent ISS trials) and plot the proportion of expectation-based correctives in Figure 2.6, across subjects. In both the infrequent-ISS versus catch trials panel (Figure 2.6, left) and neutral-stimulus trials panel (Figure 2.6, right), there is a graded transition (from right to left) of more expectation-based correctives in Predictable-Up continuing through Probable-Up, Predictable-Down, concluding with completely visually-selective correctives in Probable-Down.

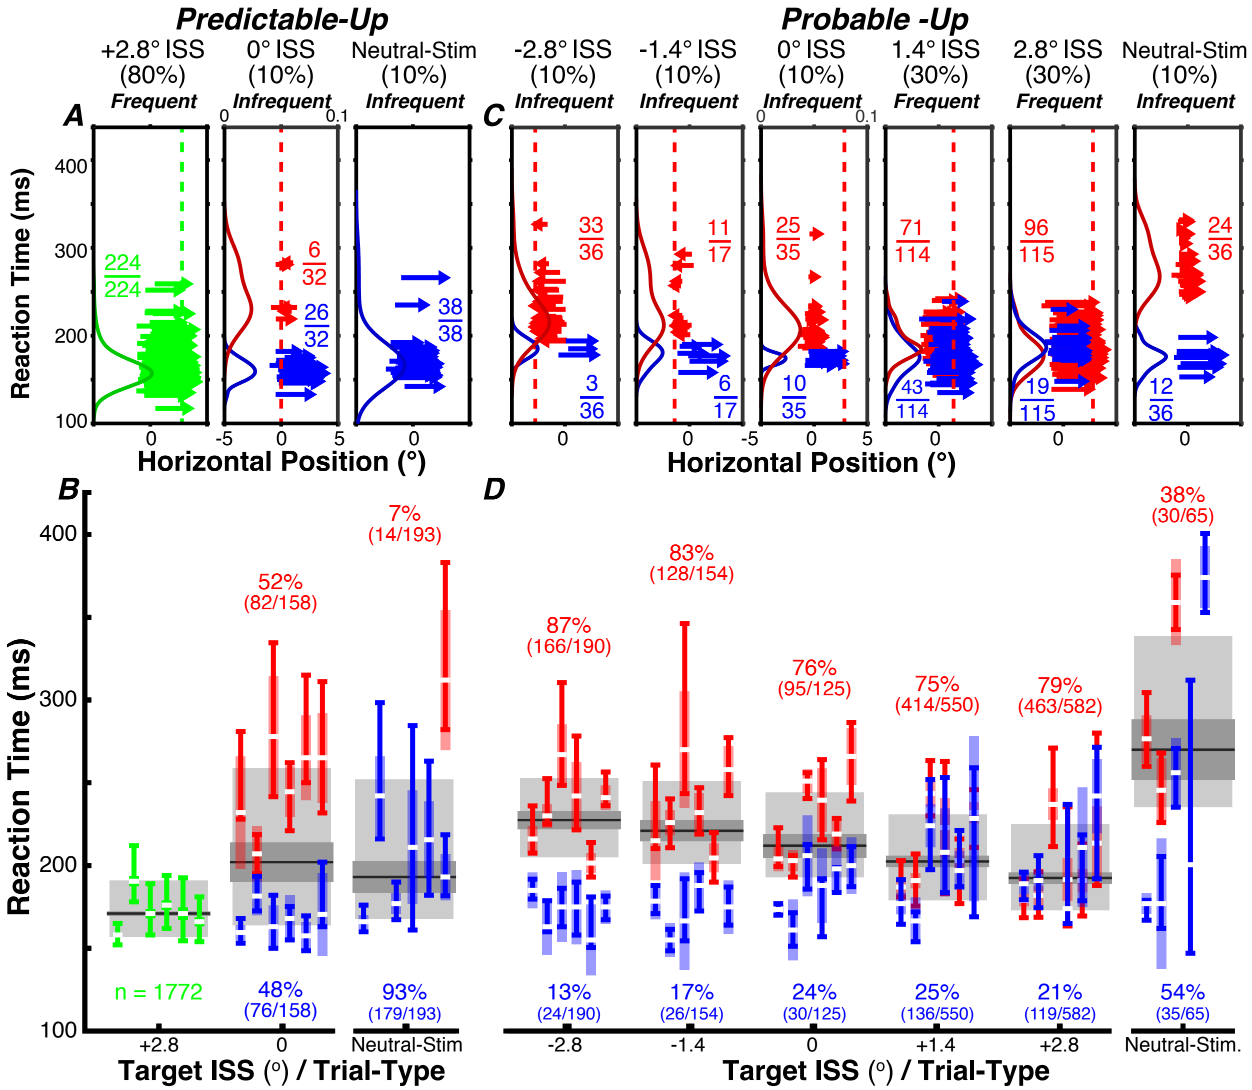

Figure 2.5 Latency Effects on Secondary Saccade Production Accompanying Gain-Increase Adaptation. **A.** Predictable-Up Example subject data. Arrows are plotted from secondary saccade horizontal start-to-end (x-axis) at time of occurance relative to primary saccade completion (y-axis). Accompanying curves are latency histograms smoothed with gaussian kernels, and refer to top axis values (kernel width chosen to guarantee a single peak; for visualization purposes only). Vertical dashed lines illustrate relative post-ISS target position. Colored fractions indicate saccade counts by corrective saccade category (color): target-directed (green), expectation-based (blue) and visually-selective (red). Separate plots are presented for +2.8° (left) and 0° (center) ISS, as well as neutral-stimulus trials (right). This subject's data are also presented, below, in B as the left-most of each colored set of symbols. **B.** Predictable-Down summary data (across and within subjects). Dark and light gray boxes and black line represent inter-quartile range (25th-75th percentile), 95% median CI, and median (respectively); colored error-bars, boxes, and white line segments represent inter-quartile range, CI, and median values for individual subjects. Colors correspond to saccade populations illustrated in A, above. Note: only one subject made secondary saccades towards 0° in neutral stimulus trials, accordingly there is only one red symbol in the neutral-stimulus group, on the right. **C**. Probable-Down example subject data. Colors and plotting conventions as in A. Subject's data are again presented as left-most of each set of colored symbols in D, below. **D.** Probable-Down summary data. Plotting conventions as in B. Median CI boxes are absent in cases with fewer than 3 data points.

Pooling across the three Probable-Up infrequent ISS trials (-2.8°, -1.4°, and 0°), we found 17% (80/469) of correctives were expectation-based, a value statistically indistinguishable from the 20% (43/212) observed in *Predictable-Down* catch trials (Figure 2.6; $\chi 2 = 1.0264$, p = 0.31). Similarly, in neutral-stimulus trials, the proportions of expectation-based correctives in *Probable-Up* (54%, 58/107) and *Predictable-Down* (44%, 35/79) were not distinguishable ($\chi 2 = 1.7823$, p = 0.18). Another similarity between those two conditions was the indistinguishable median latency in *Probable-Up*'s infrequent ISS trials (223 ± 3.75ms) and *Predictable-Down*'s catch trials (224 ± 6.85ms; $\alpha = 0.001$).

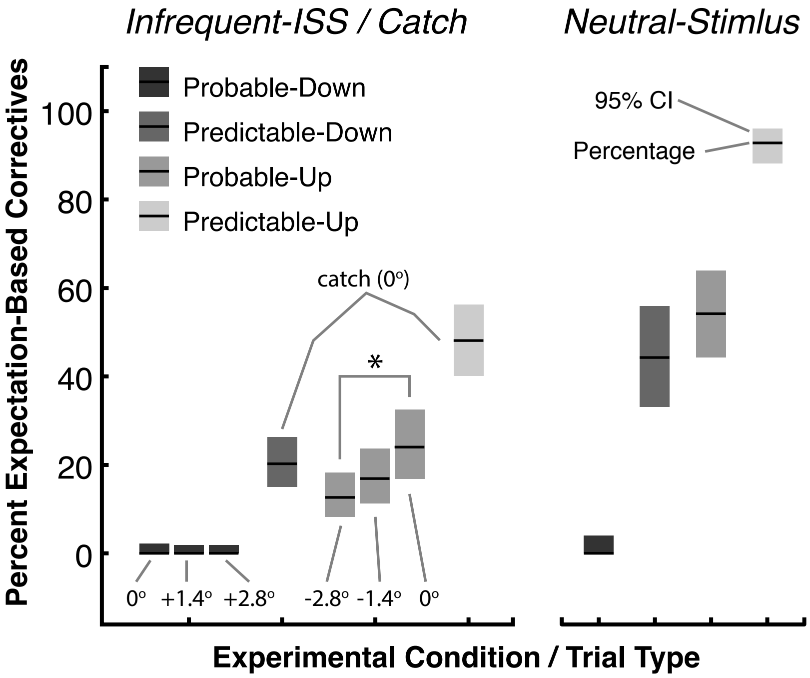

Figure 2.7 Proportions of Expectation-Based Corrective Saccades Varied with Conditions. Raw data are counts of expectation-based and total corrective saccades (in each condition / trial-type), pooled across subjects. Presented are binomial-parameter estimates (maximum likelihood; black line segements) and 95% confidence intervals (CIs; boxes). Where two CI-boxes are non-overlapping, or where indicated with an asterisk, proportions are significantly different. **Left**: Infrequent-ISS (Probable conditions, marked with ISS values) and catch trials (Predictable). **Right**: Neutral-stimulus trials.

Figure 2.6 – Chapter 2 / Proportions of Expectation-Based Corrective Saccades by Condition

In summary: corrective saccade timing and direction data both suggest that *Predictable-Up* resulted in the strongest post-saccadic expectations regarding target location, followed by *Probable-Up* and

Predictable-Down, with no evidence of expectation in *Probable-Down.* The effects of such expectations were: (a) an overall decrease in corrective-saccade latency, and (b) the production of short-latency corrective saccades directed towards more-probable post-saccadic target locations. Across conditions, with target present, expectation-based movements had similar timing: their median latency in infrequent-ISS trials was statistically unaffected by condition ($\alpha > 0.91$). Interestingly, eliminating stimulus identity (color and shape) had a modulatory effect: we still saw evidence of expectation in neutral-stimulus trials, but these correctives had significantly different median latency ($\alpha = 0.01$). We maintain that the overall latency effect (a), results from the increased tendency to produce expectation-based correctives (b), and not vice-versa, a point we address below.

Discussion

Our findings demonstrate that saccade adaptation can be driven purely by visual selection (Probable-Down), and does not require post-saccadic target presentation to be spatially consistent. Somewhat paradoxically, while corrective saccade data suggest that expectations regarding post-saccadic target location can be formed during adaptation (Figures 2.3-2.5), expressing adaptation magnitude as a fraction of average ISS implied that any such expectations do not modulate the adaptation process (Figure 2.2C). In contrast to previous work on extra-retinal signals and corrective saccades (see below), we have shown that visually driven corrective saccade programming can be trumped by the mere expectation of finding a target in an alternate location. Notably, we also found that ISS predictability did not uniformly manifest in expectation-based corrective saccade production, which we suggest is an indication of a preexisting bias in the expectation of post-saccadic target location.

Adaptation

Previously, we demonstrated that when confronted with more than one post-saccadic stimulus – target and distractor – saccade adaptation responds selectively to the target, ignoring the distractor (Appendix A - Madelain et al., 2010). However, the nature of the selectivity (visual or spatial) could not be clarified with that paradigm, because after the saccade the visual targets were always in the same spatial locations (and hence saccade targeting could have relied purely on a prediction of the final spatial location of the target, and not an active selection trial-by-trial).

Active visual selection in saccade adaptation

Here, we dissociated visual and spatial mechanisms by varying the post-saccadic target location trial-by-trial amongst multiple distractors that appeared only upon saccade. Because we randomized target and distractor identities across trials, there was no relationship between target feature and ISS. This is important to emphasize given the finding that saccade adaptation can be context-specific to visual features (Herman et al., 2009), and this design did indeed prevent any such specificity. Unlike Madelain et al. (2010), or any other saccade adaptation paradigm of which we are aware, we have created a situation where active visual-selection is frequently necessary in order to ultimately foveate the target.

Both our adaptation and corrective saccade data showed strong evidence of active visual selection, but expected ISS only modulated corrective saccade direction and latency, and not adaptation magnitude. If, as one might suspect, expectation-based corrective-saccade programming is tantamount to a prediction of primary-saccade position error, why is this not used to drive adaptation? Why wait for a visual error signal, which must be long delayed by comparison, and therefore likely less effective (Bahcall and Kowler, 2000; Shafer et al., 2000; Fujita et al., 2002)? Perhaps post-saccadic uncertainty as to target location leads the oculomotor system to form predictions of post-saccadic target location that are used in corrective saccade programming, but to await visual confirmation of the expected post-saccadic error prior to using it in modifying primary saccade gain.

Ours is not the first study to point out the importance of visual object selection in saccade adaptation (Collins et al., 2007; Lavergne et al., 2011). However, to the best of our knowledge, ours is the first to address the question of how a pre-saccadic goal might be identified and relocated post-saccadically. Namely, by a combination of top-down prediction and bottom-up visual processing. We have revealed the dual involvement of these processes by creating a variety of conflicts between them, and examined the behavioral consequences, something we are not aware of having been undertaken previously. However, that movement outcome prediction is indeed involved in saccade adaptation was recently demonstrated by Wong and Shelhamer (2011), and Collins and Wallman (2012).

Variability in ISS

As mentioned above, we were surprised to learn that the increased ISS regularity in the Predictable conditions did not affect the magnitude of adaptation (Figure 2.2C). One possible explanation for this result (discussed above) is that the nature of our adaptation paradigms – with richer visual stimuli that required a more active approach compared to more conventional adaptation paradigms – encouraged a nuanced utilization of spatially-selected post-saccadic error. That argument aside, this finding is not entirely without precedent: Srimal et al. (2008) compared adaptation in a "steady" condition (with constant downward ISSs) to a "random" condition (with a mix of 1° onward and 2° backward ISSs). Though they did not compare the magnitude of adaptation directly, they did compare a fitted rate of learning in response to negative visual errors between their two conditions and found it to be indistinguishable across subjects. Both our and Srimal et al.'s results suggest the somewhat counterintuitive idea that trial-to-trial variability in retinal errors does not decrease their influence on adaptation. A study which examined the effect of increased visual error variability on adaptation found quite the opposite: that adaptation magnitude was attenuated in proportion to added variability (Havermann and Lappe, 2010). However, the nature of the method used in that study to experimentally control the amount of variation in retinal error was such that variations in final target position could vary by as much as 15° in a single session, dramatically larger than the 3° or 5.6° used by Srimal et al. and ourselves (respectively). Indeed, there were many cases in Havermann and Lappe's experiment where ISSs left the target near its initial fixation position (i.e. 100% step-back). It may, thus, be that this degree of variability has an entirely different effect from the type of variation used by us or by Srimal et al. Further work will be required to resolve these conflicting findings.

The small but significant gain decreases resulting from our multiple-stimulus and random-ISS phases are not entirely surprising. Srimal et al., also found a significant reduction in gain as a consequence of their random adaptation condition. From the perspective of the flight-time-minimization hypothesis put forth by (Harris, 1995), undershooting errors are less costly in terms of in-flight saccade durations than overshooting errors, and minimizing these costs gives an optimal undershoot that shifts towards lower gain values as the endpoint variability increases. Hence, it might be quite sensible to reduce saccade gain in response to random ISSs. This argument is relevant to the gain-decreasing effect of the Multi-Stimulus phase as well: if the array of post-saccadic stimuli was perceived as a single, extended target, a lower gain would be expected.

As has been noted elsewhere (Hopp and Fuchs, 2004), comparing the amount of induced adaptation amongst different paradigms, or even amongst individual subjects within the same paradigm, is quite difficult both because of inter-subject variability and because individual performance can vary from day-to-day. Nonetheless, we feel confident that the amounts of adaptation we have observed are in keeping with results from more conventional adaptation experiments (Figure 2.2B).

Neural substrate of error computation

As noted in the introduction, it is easier to imagine how retinal error calculation – fundamental to the adaptation process – might occur in experiments using single point-like targets. The intermediate layers of the Superior Colliculus (SCi) receives a direct retinal projection, and contains visuomotor cells with response fields arranged in a map that would serve such a purpose. Although the cerebellum is considered the principal anatomical substrate for the physiological changes induced by saccade adaptation (see Hopp and Fuchs, 2004), recent evidence has led to the suggestion that the SC might be the source of the post-saccadic error signal that informs the cerebellum of the need for adjustment to subsequent saccades. Work from two labs independently demonstrated that subthreshold stimulation of the SCi immediately following each of several saccades induces saccade adaptation (Kaku et al., 2009; Soetedjo et al., 2009). Why then, could the SC not participate in the calculation of visual error in more complex scenes such as the natural environment or the present study? In exploring the passive visual sensitivity of single units in the SCi, it has been found that these neurons show little preference over feature dimensions such as color (Marrocco and Li, 1977; Ottes et al., 1987), indicating that the SC is unlikely to be the origin of the type of selectivity we observe. However, recent work exploring the role of SCi neurons in saccade selection in a visual search task has shown that a single-unit's activity reflects the presence of a target-stimulus in its response field as little as 80ms after stimulus-onset – even if the stimulus is not selected for the subsequent saccade (Shen and Paré, 2007). This is well within the 150ms window, described by Shafer et al. (2000), during which visual signals are most effective for inducing adaptation.

Secondary Saccades

Our newly modified paradigm revealed a surprising novelty, that corrective saccade programming can be affected by the frequency with which a post-saccadic target appears at a given location. We propose that the totality of our results is best explained by a combination of factors: (1) a native tendency to execute onward corrective saccades (following primary saccades of the size used here) and (2) the manipulation of expected post-saccadic target location. In particular the observation of a greater proportion of expectation-based than visually-selective correctives in gain-increase paradigms and the reverse tendency in gain-decrease is supportive of this two-factor explanation.

It has long been known that because most saccades undershoot their targets (Becker, 1989b), targets are frequently foveated in a sequence of two or more movements; a further consequence is that secondary movements must predominantly be onwards. Becker and Fuchs (1969) first noted the tendency to execute onwards corrective saccades in the absence of any visual input. Though the authors of that work interpreted their finding to mean that it is possible to pre-program a series of two saccades in order to reach a target, they also noted that: “it is remarkable that the second saccade is always in the same direction as the initial saccade.” That is, when they observed “pre-programmed” correctives, they were exclusively onwards. It has also previously been noted that onwards corrective saccades have significantly smaller median latency than backwards in adults (Prablanc and Jeannerod, 1975), but not in children (Shawkat et al., 1993). We suggest that these past findings are most consistent with our contention that a lifetime of experience executing onward corrective saccades leads to the expectation that the target will lie beyond the fovea at the conclusion of a primary saccade.

Four pieces of evidence from the present work underscore our contention as well: (1) gain-increase paradigms resulted in significantly more expectation-based correctives than did gain-decrease (Figure 2.6); (2) gain-increase neutral-stimulus trials frequently resulted exclusively in onwards corrective saccades (Figure 2.5B & 2.5D); (3) the 60% onwards ISSs of Probable-Up were as effective in promoting expectation-based movements as the 80% backwards ISSs of Predictable-Down (Figure 2.6); (4) our reproduction of Prablanc and Jeannerod's (1975) finding: median latency of (predominantly onwards) correctives in Predictable-Up was significantly smaller than (the predominantly backwards correctives) in Predictable Down (Figure 2.3). Notably, our results are arguably stronger evidence for the involvement of expectation in corrective saccade programming: no previous study has observed non-visually programmed correctives in the presence of an illuminated target.

What we suggest more specifically is that the patterns of ISSs in our conditions had varying degrees of efficacy in modifying an existing post-saccadic expectation of finding the target beyond the fovea. This is reflected in the graded transition we observed from more expectation-based correctives in Predictable-Up continuing through Probable-Up, Predictable-Down, concluding with completely visually-selective correctives in Probable-Down (Figure 2.6). Predictable-Up was most effective because its pattern of ISSs most closely matched preexisting expectations. It required only 60% onwards ISSs in Probable-Up to elicit a comparable proportion of expectation-based correctives as the 80% backwards ISSs of Predictable-Down because the onwards ISSs were more in-line with preexisting expectations. Meanwhile it seems that the 60% backwards ISSs of Probable-Down were effective in eliminating preexisting expectations (of onwards correctives). Indeed, it seemed that the Probable-Down adapt phase abolished an expectation of finding the target beyond the fovea: median onwards corrective saccade latency was significantly smaller prior to the adapt phase (in the random-ISS phase) of Probable-Down than in the adapt-phase (Wilcoxon rank-sum test, $\chi^2 = 3.91$, $p < 0.05$).

Because we argue that the pattern of ISSs in a given condition modified or interacted with individuals' expectations, we were curious if we could see evidence of this building up over trials. However, when we examined the relative frequency of expectation-based or visually-selective movements by trial number, we found no such evidence. In fact, it appeared as though early on in the adapt phase, expectation-based movements were more likely than they were later on in the phase. Nonetheless, it seems probable that there is some way of characterizing the rate of acquisition of the tendency to produce expectation-based correctives. Further work will be required to effectively capture this process.

Secondary Saccade Amplitude

The consistent pattern of our corrective saccade data had one exception, which was noted in the Results. Corrective latencies in response to -1.4° ISSs were larger than to other ISSs in the Probable-Down condition (Figure 2.3). Corrective amplitudes were smallest in this condition (see Figure 2.4C), because adaptation of the primary saccade left only small retinal errors to be corrected. Hence, the increased latencies are unsurprising given that very small saccades are known to have longer latencies (Wyman and Steinman, 1973; Kalesnykas and Hallett, 1994). Moreover, our stimuli (1°) were larger than those typically used in saccade adaptation experiments; this might have increased latencies further, since larger diameter targets also have longer latencies (Madelain et al., 2005; Harwood et al., 2008).

Conclusions and Directions for Further Research

We have for the first time demonstrated that saccade adaptation can rely exclusively on visual target selection. We have further discovered that top-down expectation can sometimes supersede bottom-up stimulus drive in the programming of corrective saccades.

Our study has at least two shortcomings. First, our stimuli appeared at fixed locations relative to the target's primary step location, which may have encouraged cognitive intervention. That is, the target's final position in the array not only caused a retinal error, it also alerted the subject to the ISS: if the target is leftmost in the array after the primary saccade, the subject may determine that the ISS was negative on this basis alone. One way to remedy this situation is to modify stimuli such that the subject is constantly presented with a 1-D array of stimuli that fills the field of vision. This would reduce the likelihood of cognitive involvement, and prevent the use of landmarks described above. Second, we somewhat arbitrarily chose two types of ISS bias to explore, but one can imagine more subtle manipulations that might allow for a more nuanced understanding of the shape of an individual's prior expectations regarding post-saccadic target location.

A major question left open by our study is what features of the target are used by the oculomotor system to achieve the results we have observed. That is, given that we have demonstrated a visual selection mechanism: what aspects of the stimuli are necessary for the functioning of this mechanism. Our stimuli were defined by both shape and color. We suspect that color alone would suffice for the results that we observe. However, a particularly intriguing situation to consider is that in which multiple identically colored but distinctly shaped stimuli are presented after a saccade, which will the system choose? How might the preexisting tendency to onward corrective movements figure into such a situation? Would a more similarly shaped stimulus at a greater eccentricity be ignored? These questions are of fundamental importance to the understanding of post-saccadic visual processing at the service of adaptation.

Chapter 3: End-point variability is not noise for saccade adaptation.

Introduction

Because a small (roughly 1° in diameter), central region of the human retina (the fovea) has the greatest receptor density, the best view of an environmental stimulus is achieved by accurately orienting gaze towards it. We move our gaze from point-to-point using rapid movements called Saccades. Somewhat surprisingly, the accuracy of saccades changes little with age (Warabi et al., 1984), suggesting that movement accuracy is actively monitored and maintained. Indeed, experiments on humans with extraocular muscle paresis (Kommerell et al., 1976; Abel et al., 1978; Optican et al., 1985) and tenectomized monkeys (Optican and Robinson, 1980; Snow et al., 1985) demonstrated an impressive endogenous ability to restore normal saccade amplitudes in a relatively brief period (on the order of days). Meanwhile, Mclaughlin (1967) noted that modification of saccade amplitude can be achieved by an intrasaccadic step (ISS) paradigm, applying experimentally arranged, primary-saccade-triggered secondary target shifts in a series of successive saccade trials.

This maintenance, restoration, or manipulation of saccade size is termed “saccade adaptation.” Since the recognition that adaptive changes achieved by the ISS paradigm are the same as those resulting from tenectomy (Scudder et al., 1998), the ISS paradigm has been used to extensively explore and document the multiplicity of subtle ways that saccade adaptation can display sensitivity (for a review, see Hopp and Fuchs, 2004; Pélisson et al., 2010).

There are two major classes of ISS paradigm. The first is identical to the method used by Mclaughlin (1967): the amplitude of the ISS is fixed (in degrees of visual angle, or as a percentage of the primary target step). The second is a variant that reflects the increasingly widespread belief that saccade adaptation is driven by retinal error (the difference between gaze and target position): the ISS is varied from trial-to-trial to ensure that the primary saccade results in a specific retinal error. Though experiments of the first class are by far more common (Hopp and Fuchs, 2004; Pélisson et al., 2010), experiments of the second class have in some ways been more basic or systematic in their approach (Henson, 1978; Robinson et al., 2003; Havermann and Lappe, 2010; Wong and Shelhamer 2011). Robinson and colleagues (2003) performed an experiment to characterize the adaptive response to a variety of fixed retinal errors (varied systematically in separate sessions), finding that as errors became increasingly negative, adaptation increased, but decreased once error was beyond ~40% of the initial target step. Meanwhile, the one study (of which we are aware) that compared the effects of different ISSs used only 2 subjects, and 2 ISS values (25% & 50%), making their results somewhat difficult to interpret (Miller et al., 1981).

Thus, to more systematically explore the effects of varying ISS size, we characterized gain-decrease adaptation in response to a variety fixed ISS amplitudes (where gain is defined as the ratio of saccade amplitude to primary-target-step size). We chose to only conduct gain-decrease experiments because gain-increase adaptation has been shown to have distinct characteristics (Miller et al., 1981; Deubel et al., 1986; Deubel, 1987; Fitzgibbon et al., 1986; Straube et al., 1997; Scudder et al., 1998; Kroller et al., 1999), and has been suggested to employ distinct mechanisms (Semmlow, 1989; Hernandez et al., 2008; Ethier et al., 2008b). We used ISSs with amplitudes of 1°, 0.75°, 0.5°, 0.25°, 0.1° and 0° (as a control), following 10° primary target steps. Our primary goal was to find the smallest ISS effective for inducing significant adaptation. We were also curious if saccade adaptation, like other forms of motor learning (Wei and Kording, 2009; Criscimagna-Hemminger et al., 2010), would respond differently to smaller errors. Lastly, we wondered whether any inter-subject differences in adaptation magnitude might be related to aspects of baseline saccade performance, like variability, as was the case in a fixed-retinal-error ISS paradigm (Wong and Shelhamer, 2011), and as has been implied in modeling efforts (Ethier et al., 2008a).

Methods

General

Subjects were instructed prior to each session that they would be presented with a small red annulus that would begin in each trial on the left, step to the right at a random time, and that they should track it with their gaze.

Subsequent to receiving instructions, each participant sat in a darkened room, 57cm from the display, and underwent a nine-point self-paced calibration prior to the start of recording.

Stimuli & Procedure

The fixation and target stimulus was a small (0.3°) red annulus, on a dark background.

All target positions were determined prior to all recordings and were identical for each subject.

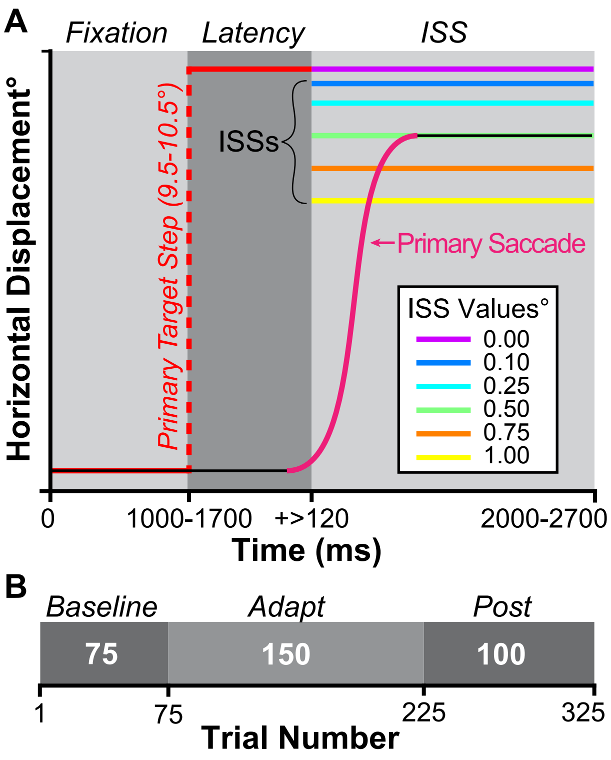

Figure 3.1 Methods **A**. Single-Trial. Primary **target** step occured 1000-1700ms after fixation onset. A minimum of 120ms later, the **primary saccade** occured, triggering the intrasaccadic target step (ISS). ISSs and primary target step are not drawn to scale. Grayscale boxes demarcate temporal windows named above. Fixation: pre-target-step fixation period; Latency: delay between target step and **primary saccade** onset. ISS: intrasaccadic target step and subsequent corrective saccade (not shown) period. **B**. Experimental session structure by phase. Phases are depicted as grascale boxes with phase names above. White numerals in each box indicate the length of the phase in trials, with absolute trial number Indicated on the axis.

Figure 3.1 – Chapter 3 / Methods

We used 2 trial-types: (1) no-ISS and (2) ISS trials (Figure 3.1A). In no-ISS trials, the subject fixated the target on the left portion of the display (initial position drawn uniformly from the interval [-7.5°, -2.5°])

for a random period (drawn uniformly from: [1000ms,1700ms]). The target stepped (amplitude uniformly drawn from [9.5°, 10.5°]), and remained in place until the end of the trial (total trial duration ranged from 2000-3200ms: Total Duration = 1000ms + Fixation Duration). Meanwhile, in ISS trials, the only difference was that upon detection of the subject's primary saccade (velocity > 25°/s), the target stepped intrasaccadically by a value depending on the condition (Figure 3.1A). Possible values were 1°, 0.75°, 0.5°, 0.25°, 0.1° and 0°, thus in the 0°-ISS condition, ISS trials were identical to no-ISS trials.

Each experimental session comprised 3 phases, a 75-trial baseline consisting exclusively of no-ISS trials, a 150-trial adapt phase of ISS trials, and a 100-trial post or recovery phase of no-ISS trials (Figure 3.1B).

Subjects

Informed consent was obtained prior to any recording sessions, and protocols were reviewed and approved by the Institutional Review Board (IRB) of the City College of New York (CCNY), and thus complied with all human-subject protocol requirements.

Data were collected from 6 subjects, ages 25-37, 2 female and 4 male. All were faculty and students from the CCNY community and had normal or corrected-to-normal vision.

Equipment

Stimuli were displayed on an Iiyama Vision Master Pro 514 CRT display (Oude Meer, Netherlands) at a resolution of 800 x 600 pixels (visible area 41.5 cm x 30.5 cm), and a vertical sync-rate of 200Hz with 8-bit color depth.

Stimulus generation and display, data storage, and overall experimental session orchestration were controlled with a custom interface in LabView (National Instruments, Austin, TX) running in Windows XP (Microsoft Corporation, Redmond, WA) on a Dell PC (Austin, TX).

Eye movements and gaze position were measured and collected by an Eyelink-1000 infrared camera system (SR-Research, Mississagua, Ontario, Canada), which sampled right-eye gaze (pupil - corneal reflection) at 1000Hz.

Analysis & Statistics

Data were analyzed using a purpose-written interface in Matlab (The Mathworks Inc., Natick, MA). Saccades were first detected automatically using a 10°/s velocity threshold, and confirmed by visual inspection. A small number of trials (< 5%) were discarded due to blinks or hypometric primary movements (< 50% of target eccentricity).

We quantified gain changes (adaptation and recovery) starting with a repeated-measures Analysis of Variance (ANOVA). We then used the Tukey-Kramer method to find the differences (and associated 95% confidence intervals) between all estimated population marginal means (PMMs). For example, in determining the amount of induced adaptation and recovery amongst conditions we began with a 3 X 6 X 6 (Phase X Condition X Subject) ANOVA; the three levels in the Phase factor were: (1) the final 50 trials of the baseline phase, (2) the final 50 trials of the adapt phase, and (3) the final 50 trials of the recovery phase (these are indicated by the shaded areas in Figure 3.2A). We then used Tukey-Kramer to compute the differences: (2) - (1) for adaptation, and (3) - (2) for recovery, across and within subjects for each condition simultaneously (Figure 3.2B, 3.2C).

Where applicable, we quote test statistics and p-values. However, because the Tukey-Kramer method provides simultaneous confidence intervals (CIs) on all differences amongst PMMs (for a given α), but only provides a boundary on p-values, when we mention the statistical significance of such differences we specify only the α-value. An example would be comparing the amount of adaptation in two conditions: a comparison of two Tukey-Kramer differences (end-of-baseline minus end-of-adapt in each condition) whose significance at the chosen α is ensured when the 95% CIs are non-overlapping; it is in cases of this type (when a p-value cannot be directly computed) that we provide only the α-value.

Stepwise linear regression was computed using the Statistics Toolbox in Matlab.

Results

We measured the propensity of a group of subjects to decrease their primary saccade gain (ratio of saccade amplitude to target step-size) in response to a set of fixed intrasaccadic steps (ISSs) of a point-like target. We sought the smallest ISS that was effective in causing significant adaptation both across and within subjects. We were curious about any differences in this minimum at the subject level: might such differences be related to simple observables such as variability in landing position or the average undershoot of saccades prior to any adaptation? We asked whether subjects responded differently to smaller errors, and, finally, we attempted to determine what potential cause best explained the pattern of adaptive changes we observed.

Minimal ISS to induce adaptation

Across subjects, adaptation occurred with ISSs as small as 0.25° (Figure 3.2B). To determine the magnitude and significance of adaptation we calculated differences in gain between the ends of the baseline and adapt phases (see Methods for details). Comparing group level confidence intervals (CIs) to 0, it is apparent that significant adaptation occurred for all ISS amplitudes greater than 0.1° (Figure 3.2B; $\alpha = 0.05$). Interestingly, comparing CIs to ISS values (red dashed lines in Figure 3.2B) indicates that while adaptation was (on average) complete (indistinguishable from ISS value) for the ISS=0.25° condition ($\alpha = 0.05$), larger ISSs yielded incomplete adaptation (Figure 3.2B; $\alpha = 0.05$). This is consistent with the suggestion that small errors may be used differently than larger errors. From Figure 3.2B, it is also apparent that individuals varied in their responsiveness to different ISSs, we next attempted to understand this variability.

Most subjects showed significant adaptation in response to 0.5° ISSs (Figure 3.2B; $\alpha = 0.05$). However, subject D required an ISS of at least 0.75°, and the smallest effective ISS for subject A was 1° (Figure 3.2B; $\alpha = 0.05$).

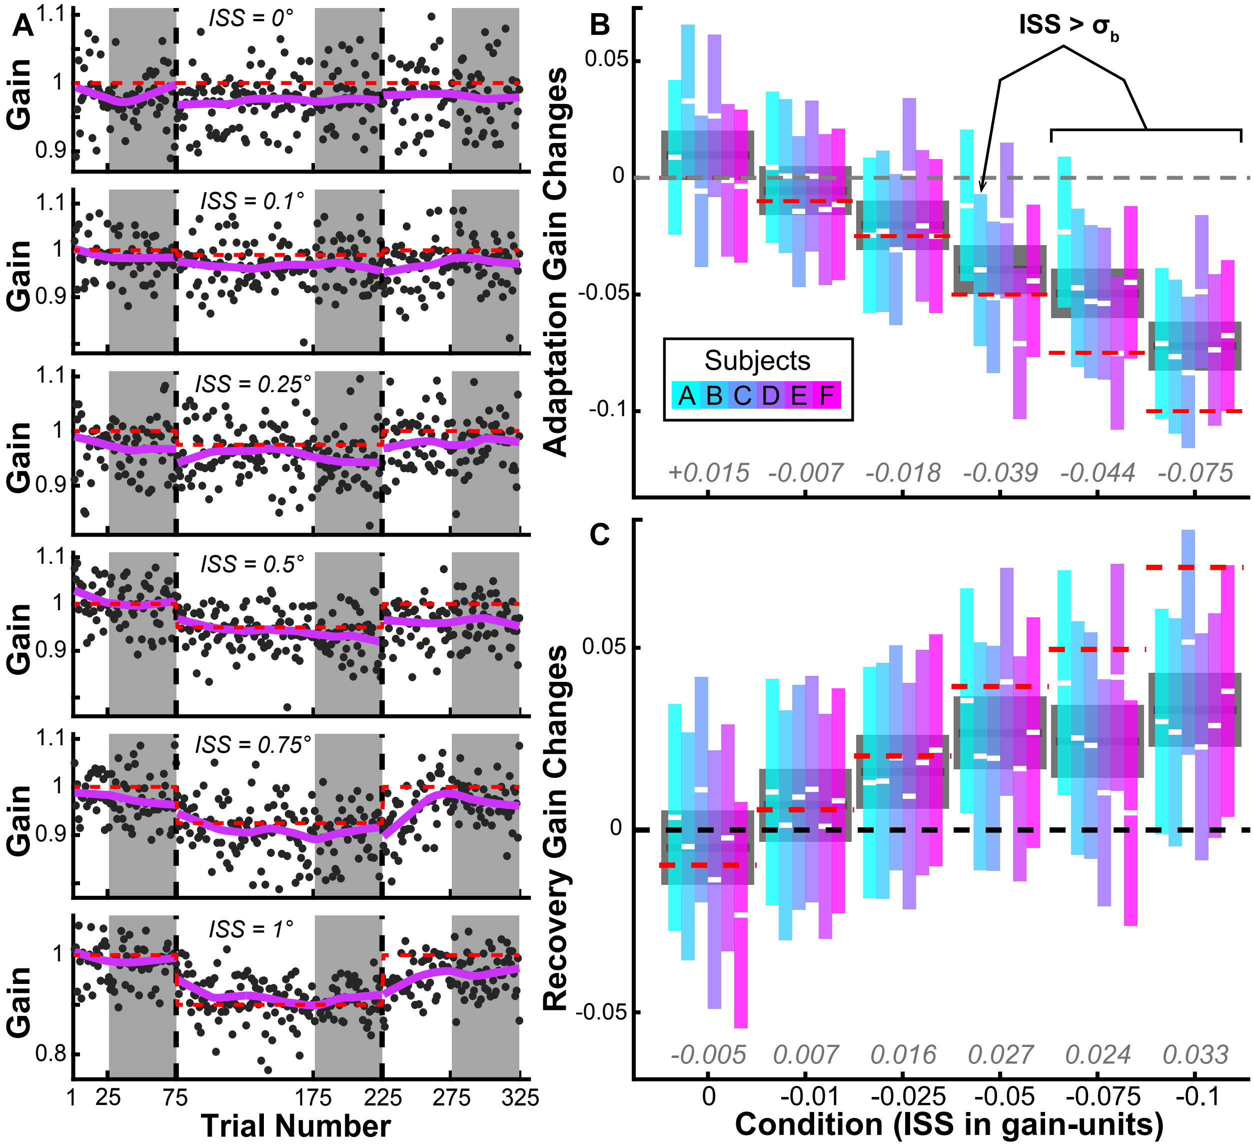

Figure 3.2 Gain Adaptation and Recovery. **A**. Example subject raw data. Primary saccade gains (black dots) are plotted versus trial numbers in ISS condition: 0° at top, down to 1° at bottom. **Red dashed lines** indicate the final target position, and **purple traces** are robust lowess smooths, color corresponds to subject E's **purple boxes** in B and C. Gray regions indicate those used to calculate gain-change magnitudes plotted in B, C. For details on this procedure, see Methods. Briefly, adaptation was calculated by the difference between the means of the final 50 trials of baseline and adapt phases using the Tukey-Kramer method after an ANOVA. **B**. Magnitude of adaptation across and within subjects. Gray boxes and dark gray line segments represent across-subjects adaptation, colored boxes and white line-segments represent individual subject adaptation (mean ± 95% CI: Confidence Interval). Dashed gray line represents 0 change, and may be used to determine significance by comparison with CIs (no overlap meaning significant change). **Red dashed lines** are ISS-values (in gain units), and gray italic numbers indicate across-subjects adaptation magnitude. We also indicate those cases where ISS-value was greater than the subject's baseline variability (σ_b). **C.** Magnitude of recovery. Conventions as in B, except r**ed dashed lines** indicate adaptation magnitude as shown in A and gray italic text indicates magnitude of recovery. Black dashed line represents 0 change.

Relation between adaptation threshold and end-point variability:

The propensity for an individual subject to adapt was unrelated to end-point variability. When ISS > σ_b (the standard deviation of landing position in the final 50 trials of the baseline), significant adaptation always occurred (Figure 3.2B, as indicated; $\alpha = 0.05$). However, when ISS < σ_b significant adaptation still sometimes occurred: in the ISS=0.5° condition, subjects C, E, and F displayed significant adaptation (σ_b= 0.51°, 0.61° and 0.65°, respectively), while A and D did not (σ_b= 0.7° and 0.58°). This is suggestive that end-point variability is not equivalent to target perturbation sensitivity.

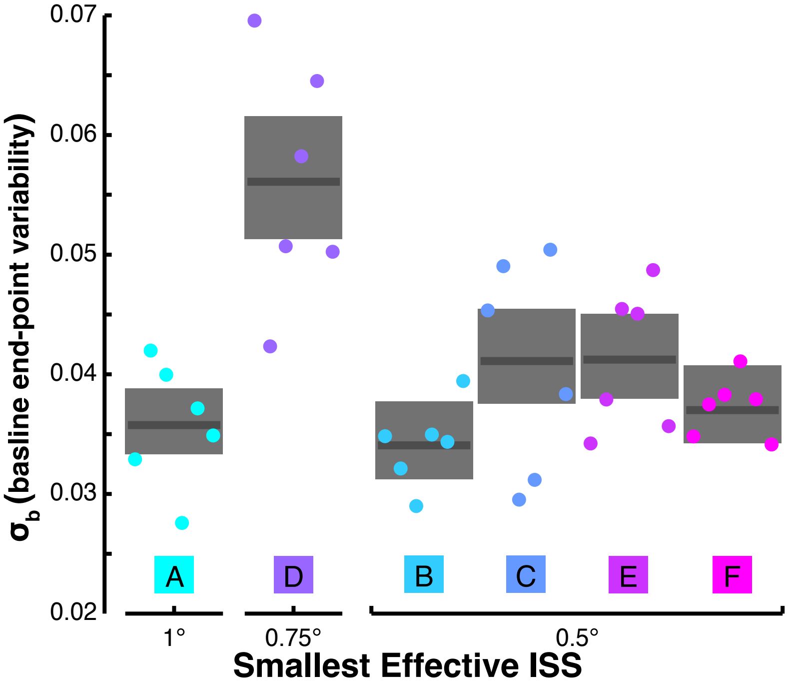

Figure 3.3 End-point variability and smallest effective ISS. We define baseline end-point variability (σ_b) as: the standard deviatino of primary saccade gain during the final 50 trials of the baseline phase. Colored symbols are σ_b calculated for each subject in each condition. Gray line segments are σ_b, calculated by pooling mean-subtracted gains across conditions for each subject ,and grey boxes are bootstrapped confidence intervals from the same pooled data. The smallest effective ISS is the minimum ISS size for which a particular subject showed a significant decrese in gain due to adaptation (see Figure 3.2A).

Figure 3.3 – Chapter 3 / Smallest Effective ISS and End-Point Variability

In order to reveal any trend relating the threshold ISS value to end-point variability, we plotted an estimate of each subject's variability grouped by threshold (Figure 3.3). For an across-conditions estimate of σ_b, we centered and pooled the final 50 trials of each baseline phase (separately for each subject), computed the standard deviation, and used a bootstrapping procedure to calculate a confidence interval. Interestingly, while subject B had the largest end-point variability in gain (0.056) and a higher threshold for adaptation (0.75°), the highest threshold (1°) was displayed by subject A, who had the 2nd smallest end-point variability (0.035). The finding that ISS threshold is not clearly related to end-point variability further strengthens our contention that end-point variability does not mirror ISS sensitivity.

As a final component of this argument, and to explore the question of what does account for adaptive changes in gain, we have included end-point variability along with several other putative predictors of adaptation in a regression analysis.

What best explains overall adaptation results

ISSs best explained the collective pattern of adaptation observed. Several factors were individually assessed for their power to predict adaptation magnitude (calculated as in Figure 3.2B) using linear regression: (1) mean (retinal) error, (2) mean corrective saccade amplitude, (3) ISS – σ_b, (4) ISS – "inherent hypometria" (IH), (5) mean error – IH, and (6) ISS (Figure 3.4). We used the first 25 trials of the adapt phase to compute mean error (relative to the target's location after ISS) and corrective saccade amplitude, since both change quite rapidly during adaptation. The term "inherent hypometria" is taken from Wong and Shelhammer (2011) who found that the best predictor of adaptation was the difference between error and IH; we calculated IH as the average difference between target location and primary saccade end-point, in the final 50 trials of the baseline.

We found ISS to have the largest R^2 (0.77, $p << 0.01$) of any factor examined (Figure 3.4). Interestingly, mean error – IH had the smallest R^2 (0.01, 0.5). Note also that while factors (3) or (4) both had large, significant R^2 values (0.68, and 0.76, respectively, both $p<<0.01$), both were decrements in explanatory power relative to ISS alone. The finding that ISS has the greatest ability (of any individual factor) to explain adaptation is consistent both with the notion that ISS best accounts for adaptation in conventional (consistent ISS) adaptation paradigms, and with our contention that end-point variability does not determine ISS sensitivity. However, this result does not rule out the possibility that factors other than ISS simply play a lesser or more modulatory role in adaptation.

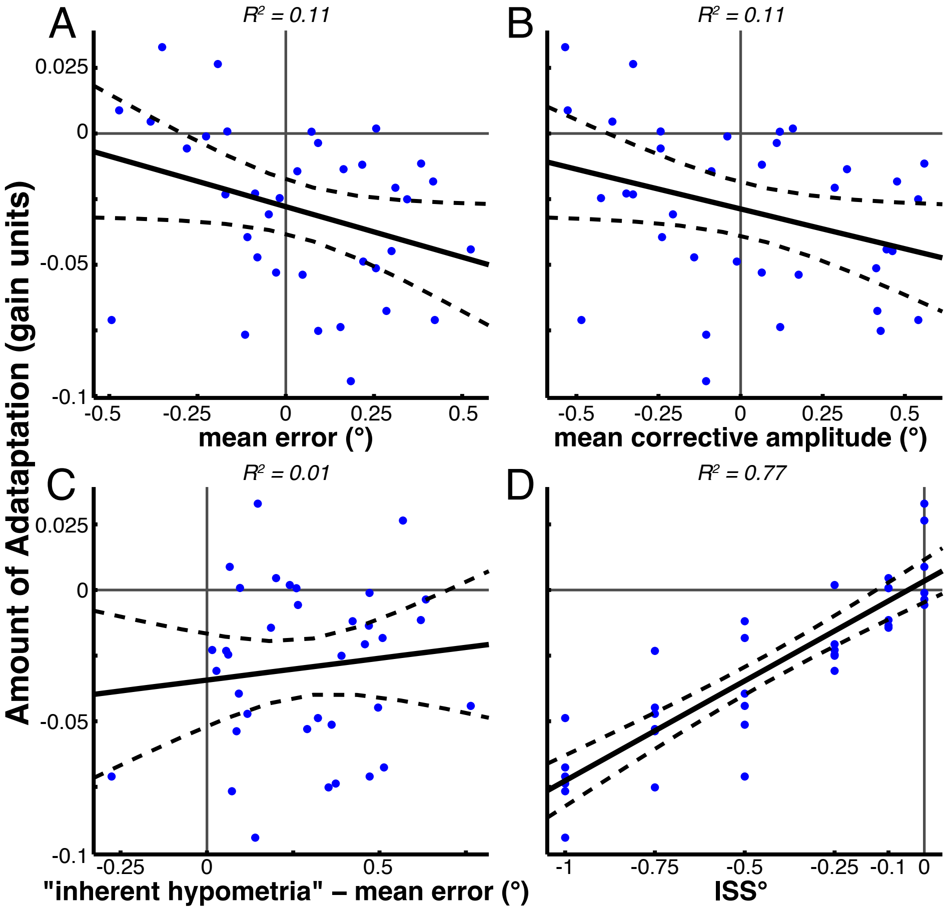

Figure 3.4 Regression Analysis. In all panels: one data point (•) for each subject in each condition; solid black lines are least-squares linear fits, and dashed black lines are prediction confidence intervals, corresponding R^2 values are presented above each plot; solid gray lines are x and y = 0, respectively; vertical axis is the mean amount of adaptation, for details of this calculation, see Methods. **A**. Mean error was calcualted in the first 25 trials of adaptation, since this value will vary most significantly amongst conditions in this portion of the adaptation phase, see text for clarification. **B**. Similarly, mean corrective saccade amplitude was calculated only in the first 25 adapt-phase trials. **C**. The difference between "inherent hypometria" and mean error. Inherent hypometria was calculated the mean end-point error during the final 50 trials of the baseline phase. **D**. Intrasaccadic Step Size (ISS) in degrees (°).

Figure 3.4 – Chapter 3 / Regression Analysis of Adaptation Results

A stepwise linear regression was conducted to determine whether a combination of factors might best account for adaptation magnitude. Only ISS was necessary to explain the amount of observed adaptation (regression coefficient: 0.76, t-stat: 10.69, $p \ll 0.01$), and no additional explanatory power was furnished by any of the remaining terms, which were: (1) ISS size, (2) "inherent hypometria" (IH), (3) mean (retinal error), (4) mean corrective saccade amplitude, (5) σ_b (baseline variability), and (6) ISS/σ_b. The final factor was included because that ratio is a way of representing the relationship between ISS and end-point variability that cannot be captured by simply including the individual factors. This result suggests that adaptation magnitude depends most directly on ISS size; that the other factors that we considered do not meaningfully improve the ability to predict the magnitude of adaptation, or that these other factors contribute to the adaptation process in a nonlinear fashion.

Recovery

Significant recovery occurred whenever significant adaptation had occurred (Figure 3.2C). Similar to the completeness of adaptation, across subjects, we only observed complete recovery in the ISS=0.25° condition (Figure 3.2C; recovery magnitude was not significantly different from the magnitude of adaptation; $\alpha = 0.05$). Recovery appeared to saturate to a greater extent than did adaptation: despite significantly greater adaptation in the ISS=1.0° compared to the ISS=0.75° condition (Figure 3.2B; $\alpha = 0.05$), recovery magnitudes were indistinguishable ($\alpha = 0.05$).

Discussion

Our results are broadly in keeping with that of the one previous study (of which we are aware) that explored the effect of systematically varying ISS size. That study observed that backwards ISSs ranging from 25-50% of the primary target step resulted in proportional gain decreases that reached 60% completeness (Miller et al., 1981). In contrast, our ISSs ranged from 0-10% of the primary target step, and gain changes reached 70% completeness (on average). Interestingly, we found that adaptation was complete for an ISS of 2.5% (of primary step), but that larger ISSs did not result in complete adaptation (Figure 3.2B). This finding is consistent with those of Robinson et al. (2003), who found that when adaptation was induced with fixed post-saccadic retinal errors, smaller errors resulted in proportionally greater adaptation. Wei and Kording (2009) pointed out that a model that adapts more in response to errors attributable to internal versus external factors can account for Robinson et al.'s (2003) findings. In that we have also found small errors to yield proportionally greater adaptation, our results are consistent with this hypothesis.

This present study is part of a growing literature that suggests that retinal error is not the primary signal driving saccade adaptation in humans. We found that ISS was the best predictor of adaptation magnitude, and that the additional factors we considered – retinal error amongst them – did not improve that prediction. Several other studies have highlighted examples where retinal error is a poor predictor of adaptation as well. Bahcall and Kowler (2000) asked subjects to purposefully make saccades partway to a target that was stepped intrasaccadically. Despite a large positive retinal error, this paradigm led to a decrease in saccade gain. As previously noted, Wong and Shelhammer (2011) found that retinal error did not account for the adaptation they observed in their fixed post-saccadic retinal error study. Meanwhile, work from our own lab demonstrates that even when the specific trial-to-trial pattern of retinal errors from one session is exactly reproduced in a subsequent session (with the same subject), but with different ISSs, the magnitude of adaptation differs significantly (Collins and Wallman, 2012). All of these studies suggest that a “prediction error” (the difference between expected and actual post-saccadic target location), drives the adaptation.

If a prediction error does indeed drive saccade adaptation, our results suggest that the predictive mechanism can be quite accurate. We found significant adaptation across subjects with an ISS of 0.25°, smaller than any individual’s σ_b (baseline variability) in any session (Figure 3.3); as well as several examples of individuals who adapted to 0.5° ISSs despite having $\sigma_b > 0.5°$. This implies that accompanying predictions of post saccadic target location must be at least as accurate as the ISS size. In keeping with this conclusion are results from the study of two-saccade sequences. When the second saccade in such a sequence is memory guided, and no visual landmark is available at the conclusion of the first saccade, making an accurate second saccade requires compensation for the natural variability in landing position arising from the first saccade. Several studies have explored the extent of such compensation (Bock et al., 1995; Ditterich et al., 1998; Munuera et al., 2009; Joiner et al., 2010) with differing results. However, it appears that in most cases, compensation can be complete (Ditterich et al., 1998; Joiner et al., 2010). Thus, it is not unreasonable to suggest that individuals can detect deviations from prediction of the scale of our minimum effective ISSs on a trial-by-trial basis.

Our results leave open the question of what accounts for inter-subject differences in minimum effective ISS. That such differences were present in our data suggests that individuals either possess differing capacities for detecting intrasaccadic perturbations, or differing capacities to adaptively respond to them. Unfortunately, the present results do not allow us to disambiguate these two possibilities. It is worth noting however, that both possibilities are consistent with previously mentioned results from two-saccade sequences, in which it is apparent that the ability to compensate in the second-saccade for variability in first-saccade differs amongst individuals (Munuera et al., 2009; Joiner et al., 2010). Further work will be required to clarify this important distinction.

In conclusion, our results support the notion that (when fixed) intrasaccadic target steps are a better predictor of adaptation magnitude than retinal error. In addition, our results suggest that saccade end-point variability is not all "noise." Rather, it seems that the oculomotor system has a non-visual estimate of saccade end-point that is much more accurate than previously thought.

Chapter 4: Effect of intermittent feedback on saccade adaptation and retention

Introduction

Saccades are rapid reorientations of gaze that humans execute approximately 200,000 times per day. Because they are so fast (on the order of tens of milliseconds), they cannot be guided in flight by visual feedback. Rather, the accuracy of each movement must be evaluated post-hoc, and any necessary adjustments incorporated into subsequent saccades. Mclaughlin (1967) first noted that covertly shifting a target during each of several consecutive saccades results in a gradual change in saccade amplitude (in the direction of the shifts). This form of plasticity, known as saccade adaptation, can be expressed in the laboratory in a matter of minutes (over some hundreds of trials in humans, and thousands of trials in monkeys; Hopp and Fuchs, 2004). It is thought that the difference between target and gaze (so called visual or retinal error) is the signal that guides adaptation (Wallman and Fuchs, 1998; Noto and Robinson 2001, but see Wong and Shelhammer, 2011 and Collins and Wallman, 2012). In part due to the seemingly simple nature of this adjustment mechanism, saccade adaptation and has been used extensively as model of motor learning in general.

Several lines of evidence suggest that motor learning is a process that occurs at multiple timescales (Smith et al., 2006; Kording et al., 2007; Ethier et al., 2008a); both learning and forgetting can happen either quickly or slowly. Such ability, it is argued, allows us to flexibly adapt to a host of changes (either internal or external) that occur at a variety of timescales themselves. Thus, when it is necessary to modify motor output in response to a transient disturbance (such as muscle fatigue), the required adjustment can be learned rapidly, and forgotten rapidly when it becomes un-useful. Meanwhile, if such a disturbance persists (say, due to muscle degeneration), an adaptive process that both learns and forgets slowly facilitates more long-term adjustment to movement parameters.

In support of this theoretical description of motor learning, several studies have found that if errors are introduced gradually, the resulting learned behavior is more resistant to decay than if it is introduced abruptly (Kagerer et al., 1997; Klassen et al., 2005; Kluzik et al., 2008; Huang and Shadmehr, 2009). The generalized interpretation, then, is that the learning / retention timescale(s) are not constant, but rather are “matched” to the learning event by the brain.

If enhanced retention of movement patterns with greater long-term utility is a general property of motor learning, decreasing the frequency or rate of error feedback should also lead to more enduring motor memories. That is, if a movement adaptation has been acquired by feedback delivered only on every 10th movement, and allowed to spontaneously wash-out, such a memory should be retained for a longer period than one acquired with feedback given after every movement. Surprisingly, it seems that studies of the effects of intermittent or reduced-rate feedback on motor learning are almost exclusively experiments in motor skill learning (Bilodeau and Bilodeau, 1958; Ho and Shea, 1978; Winstein and Schmidt, 1990). Such work explores the ability to improve motor performance in novel tasks, rather than the ability to adapt in response to experimentally induced perturbations to movement tasks. In general, these studies have found that intermittent feedback (which these studies called “knowledge of results,” abbreviated KR), has little to no effect on learning (skill acquisition), but that retention of learning is greater following intermittent versus uninterrupted KR. The lone adaptive motor learning study (of which we are aware) to document the effects of intermittent feedback also found that it enhanced retention, but only if a second, interfering task was learned prior to retention testing (Overduin et al., 2006). Enhanced retention has also been extensively documented in the field of operant conditioning, wherein a conditioned behavior acquired via an intermittent reinforcement schedule is more resistant to extinction than one acquired with reinforcement on every trial (Pavlov, 1927; Skinner, 1938; Humphreys, 1939; Sheffield, 1949).

In an effort to relate these findings on intermittent feedback (or reinforcement) to what has been learned from more contemporary motor learning work, we conducted two saccade adaptation experiments. In the first, we varied the rate of (error) feedback delivery in a fixed number of trials, and looked at the rate and magnitude of both adaptation and spontaneous recovery. In the second, we instead kept the total number of feedback trials fixed and compared motor learning dynamics between uninterrupted and intermittent feedback cases. To maximize learning effects, we relied exclusively on gain-decrease adaptation (where gain refers to the ratio between saccade amplitude and primary target step size), which is generally of greater magnitude than gain-increase. Also, in hopes of observing passive retention alone, we followed adaptation with a period in which feedback was always withheld to rule out the possibility that recovery from adaptation was driven by the target. Does feedback intermittency affect retention in saccade adaptation, as our speculative reasoning and findings from past learning experiments suggest? What of those past experiments' predicted minor effects on adaptation?

Methods

General

Subjects were instructed prior to each session that they would be presented with a small red annulus that step at random times around the display, and that they should track the movements of this target with their eye movements. They were also instructed that the target might disappear during tracking, and that it would be re-illuminated shortly thereafter.

Subsequent to receiving instructions, each participant sat in a darkened room, 57cm from the display, and underwent a nine-point self-paced calibration prior to the start of recording.

Procedure & Stimuli

The target stimulus was a small (0.3°) red annulus on a gray background.

Online saccade detection was achieved with a 25°/s velocity threshold. This threshold was higher than that used for offline analysis to avoid spurious triggering, and to ensure that saccade velocity was high when intrasaccadic target shifts occurred.

We used an equal mix of leftward and rightward 12° horizontal target steps, with the target beginning at the previous trial's final target location. All target positions were determined prior to all recordings and were identical for each subject.

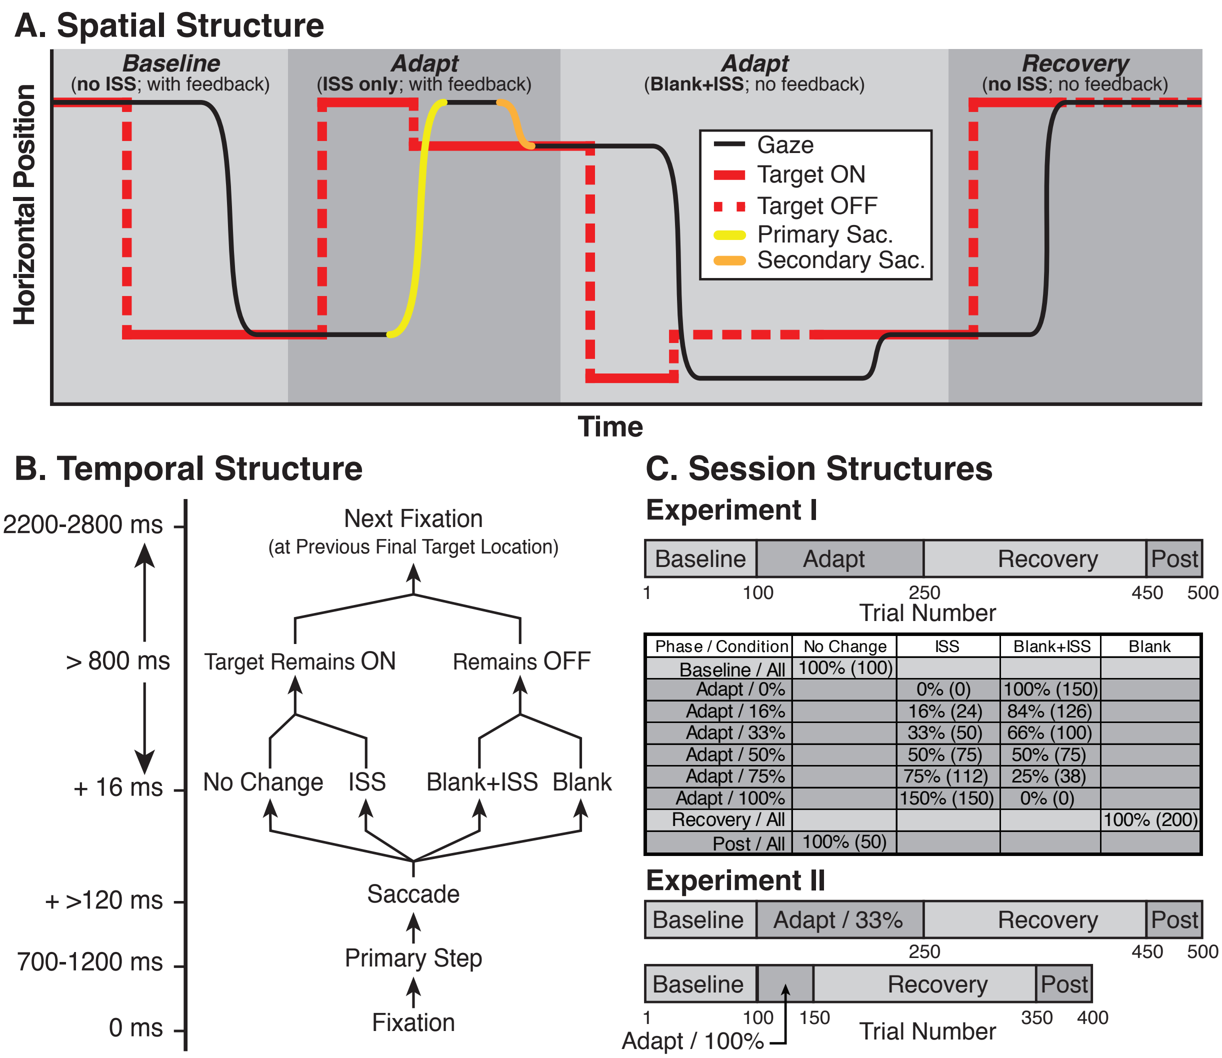

Phase / Condition	No Change	ISS	Blank+ISS	Blank
Baseline / All	100% (100)			
Adapt / 0%		0% (0)	100% (150)	
Adapt / 16%		16% (24)	84% (126)	
Adapt / 33%		33% (50)	66% (100)	
Adapt / 50%		50% (75)	50% (75)	
Adapt / 75%		75% (112)	25% (38)	
Adapt / 100%		150% (150)	0% (0)	
Recovery / All				100% (200)
Post / All	100% (50)			

Figure 4.1 Methods. **A**. Spatial Structure. Each trial began at the final target location of the previous trial, and each session contained an equal number of rightward and leftward 12° primary target steps. Three phases and their four trial-types are shown. Baseline (no ISS): a "normal" reflexive saccade trial, the target abruptly appeared at an eccentric location, triggering a saccade shortly after. Adapt (ISS only): the primary saccade triggered a second, intrasaccadic target step (ISS), in turn often provoking a secondary, corrective saccade; Adapt (Blank+ISS) saccade triggered both target blanking and an ISS, the target was turned on later at the post-ISS location (see B for timing info). Recovery (Blank): saccade triggered target blanking only, no-ISS, the target was later turned on at the same location. **B**. Temporal structure. The primary target step occured 700-1200 ms after the appearance of the fixation. A saccade occured a minimum of 120 ms later, leading to a display change a further 16 ± 3ms after saccade detection. The next fixation appeared a minimum of 800ms later, though this time was typically greater than 1300ms. **C**. Session Structures. Experiment I sessions were 500 trials in length, comprising a 100-trial baseline, 150-trial adapt, 200-trial recovery, and a 50-trial post phase. Conditions differed in the percentage of ISS trials in the adapt phase, with values in the 0-100% range. In Experiment II, the 33% feedback condition was taken from Experiment I. The "consecutive feedback" condition was 400 trials in length, as the adapt phase now contained only 50 trials.

Figure 4.1 – Chapter 4 / Methods

We used 4 trial-types (Figure 4.1A, 4.1B): (1) no-ISS with feedback (intrasaccadic target step) trials, in which the target stepped and remained on, in place; (2) no-ISS without feedback, in which the target stepped and was extinguished upon the subject's primary saccade, turning back on at the stepped location at the end of the trial, (3) ISS trials, in which the target stepped, and then stepped again upon the subject's saccade; (4) BLANK+ISS trials, in which the target stepped, and then was simultaneously extinguished and stepped upon the subject's saccade, turning back on at the post-ISS location at the end of the trial. Each trial lasted 2200-2800ms, for detailed timing information, see Figure 4.1B.

Each session in Experiment I had the same phase-structure (Figure 4.1C): (1) a baseline phase comprising 100 no-ISS trials, (2) a 150-trial adapt phase with a mixture of X ISS trials and (150-X) BLANK+ISS trials, where X varied by percent-feedback condition (100% - 150, 75% - 112, 50% - 75, 33% - 50, 16% - 24, 0% - 0); (3) a recovery phase of 200 no-ISS without feedback trials; (4) a post-phase of 50 no-ISS with feedback trials.

One additional session was recorded for Experiment II (the "uninterrupted feedback" condition), which was identical to the sessions in Experiment I, except that the shortened adapt phase (2) consisted of 50 ISS trials only (Figure 4.1C).

Equipment

Stimuli were displayed on an Iiyama Vision Master Pro 514 CRT display (Oude Meer, Netherlands) at a resolution of 800 x 600 pixels (visible area 41.5 cm x 30.5 cm), and a vertical sync-rate of 200Hz with 8-bit color depth.

Stimulus generation & display, data storage, and overall experimental session orchestration were controlled with a custom interface in LabView (National Instruments, Austin, TX) running in Windows XP (Microsoft Corporation, Redmond, WA) on a Dell PC (Austin, TX).

Eye movements and gaze position were measured and collected by an Eyelink-1000 infrared camera system (SR-Research, Mississagua, Ontario, Canada), which sampled right-eye gaze (pupil - corneal reflection) at 1000Hz.

Subjects

Informed consent was obtained prior to any recording sessions, and protocols were reviewed and approved by the Institutional Review Board (IRB) of the City College of New York (CCNY), and thus complied with all human-subject protocol requirements.

Data were collected from 7 subjects, ages 23-37, participated, 3 female and 4 male. All were faculty and students from the CCNY community and had normal or corrected-to-normal vision.

Analysis & Statistics

Data were analyzed using a purpose-written interface in Matlab (The Mathworks Inc., Natick, MA). Saccades were first detected automatically using a 10°/s velocity threshold, and confirmed by visual inspection. A small number of trials (< 5%) were discarded due to blinks or hypometric primary movements (<50% of target eccentricity).

We quantified gain changes (adaptation and recovery) starting with a repeated-measures Analysis of Variance (ANOVA). We then used the Tukey-Kramer method to find the differences (and associated 95% confidence intervals) between all estimated population marginal means (PMMs). For example, in determining the amount of induced adaptation and recovery amongst conditions we began with a 3 x 6 x 7 (Phase x Condition x Subject) ANOVA; the three levels in the Phase factor were: (1) the final 50 trials of the baseline phase, (2) the final 50 trials of the adapt phase, and (3) the final 50 trials of the recovery phase. We then used Tukey-Kramer to compute the differences: (2) - (1) for adaptation, and (3) - (2) for recovery, across and within subjects for each condition simultaneously (Figure 4.2A, 4.2D). We used a similar method to assess adaptation and recovery rates, now using an 11 x 6 x 7 (within-phase-trial-grouping x condition x subject) repeated measures ANOVA. The 11 levels in the first factor corresponding to the successive groups of trials in the considered phase: for adaptation rate, (1) 86-100, (2) 101-115, (3) 116-130, et cetera. The Tukey-Kramer method was then used to compute differences between the first group and all successive groups (Figure 4.2B, 4.2E). Finally, to compute in successive groups as a percentage of total adaptation, we scaled the data by the difference (11) - (1) prior to performing an additional ANOVA / Tukey-Kramer difference calculation.

Where applicable, we quote test statistics and p-values. However, because the Tukey-Kramer method provides simultaneous confidence intervals (CIs) on all differences amongst PMMs (for a given α), but only provides a boundary on p-values, when we mention the statistical significance of such differences we specify only the α-value. An example would be comparing the amount of adaptation in two conditions, a comparison of two Tukey-Kramer differences (end-of-baseline minus end-of-adapt in each condition) whose significance at the chosen α is ensured when the 95% CIs are non-overlapping; it is in cases of this type (when a p-value cannot be directly computed) that we provide only the α-value.

In our analysis of Experiment II, we plot the progression of adaptation across subjects (Figure 4.3A). This was accomplished by additively aligning (gain in) the first ISS trial in the adapt-phase in each condition to the mean in that trial across subjects and (the three plotted) conditions. We then computed the within-condition average gain and standard deviation in each trial across subjects, plotting that average along with the standard error (Figure 4.3A).

Results

Experiment I

Reducing the percentage of ISS trials with feedback weakly attenuated adaptation magnitude (Figure 4.2A). We quantified the magnitude of adaptation as the difference between gain at the ends of the baseline and adapt phases (for details, see Methods). Significant adaptation was observed for all nonzero %feedback conditions, as can be seen by comparing gray confidence intervals to the green, dashed non-significant difference line at 0 ($\alpha=0.05$). Interestingly, a significant reduction in adaptation magnitude (compared to 100%) was observed only when feedback was delivered on fewer than 50% of trials ($\alpha=0.05$), as can be seen by comparing confidence intervals (Figure 4.2A). Further, adaptation magnitude with 16% feedback (0.06 gain units) was roughly half (0.47) of that induced with 100% feedback (0.13 gain units). Importantly, no significant adaptation was observed in the 0% feedback condition, in keeping with previous findings in monkeys (Shafer et al., 2009). Thus, consistent with previous motor learning

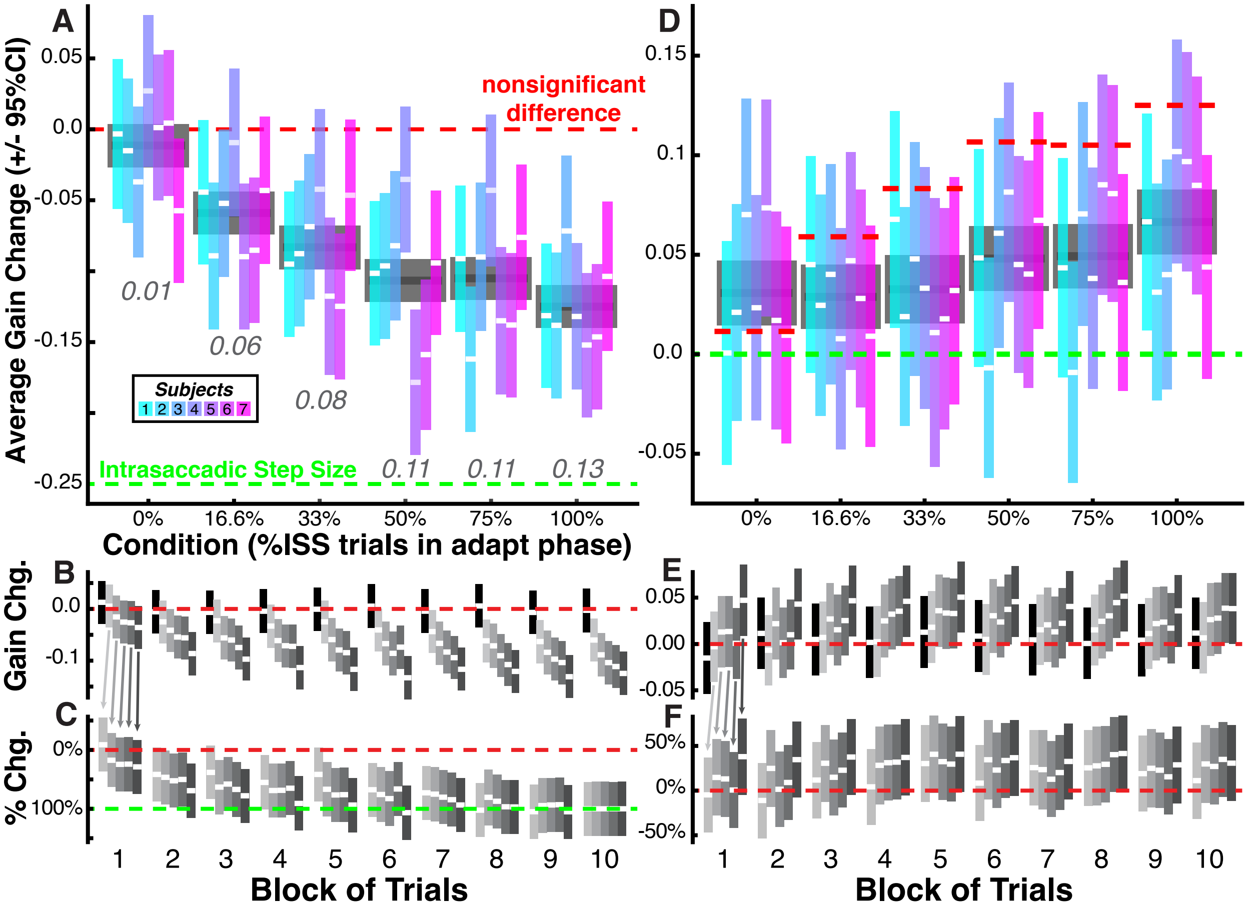

Figure 4.2 Adaptation and Recovery Gain Changes. **A**. Magnitude of Adaptation. For details of computational procedure, see Methods. Briefly: we calculated the difference between the final 50 trials of the baseline and adapt phases using the Tukey-Kramer method after an ANOVA. Gray line-segments and boxes represent mean ± 95% confidence interval (CI) across subjects (means values in italic text below); white line segments and colored boxes represent corresponding individual subject means and CIs. Intrasaccadic step size (-25%) is indicated by the **dashed green line**. **B**. Progression of adaptation-induced gain changes in blocks of 15 trials (gain units). Within each block, conditions run left to right, from 0% feedback (black) to 100% feedback (dark gray), white line segments and shades boxes are again the mean ± 95% CI. **Dashed red line** indicates no adaptation. **C**. As in B, but scaled (divided) by the final adaptation magnitude to give the percentage of total adaptation in each block. Because there was very little adaptation in the 0% condition, black bars were omitted from this panel (gray arrows between CIs in B and C indicate correspondence). **Dashed red line** indicates no adaptation, and **dashed green line** indicates full (total) adaptation. **D**. Magnitude of Recovery, as in A, calculated from differences between the final 50 trials of the adapt recovery phases. Dashed red lines indicate adaptation magnitude from A. **E-F**. As in B-C, but showing recovery progression in blocks of 20 trials. Throughout, significance of gain changes may be determined by comparison to a value of interest. For example: from panel A data, one can see that significant adaptation occured in conditions with %feedback ≥ 16.6 since gray confidence intervals do not overlap the **dashed red no adaptation line**. As another example in panel A: **subject 4** adapted significantly only in the 100% condition as can be seen by the **purple CIs** overlapping the **no adaptation line** in all but the 100% condition, where the **line** and **CI** do not overlap.

Figure 4.2 – Chapter 4 / Experiment I Adaptation and Recovery

studies, and classic operant conditioning experiments, decreased feedback frequency only mildly reduces adaptation magnitude despite a considerable reduction in the number of instances that error was experienced.

Varying percent feedback during adaptation also had a weak effect on the rate of adaptation (Figure 4.2B, C). We compared the average gain change (relative to baseline) in successive groups of 15 trials, in each condition (Figure 4.2B). Plotting the data in this way revealed that only the 16% feedback condition lagged behind the others in reaching significance (α=0.05). We also plotted gain changes during the adapt-phase as a percentage of the maximum gain change achieved in each condition (Figure 4.2C). This analysis revealed that adaptation in the 100% feedback condition reached a value comparable to final adaptation magnitude most rapidly (Figure 4.2C, block 2). Thus, while the 100% feedback condition led to slightly faster, and the 16% feedback condition to slightly slower adaptation, overall the rates of adaptation appeared to be quite similar.

Recovery magnitude was minimal, irrespective of percent feedback (Figure 4.2D). We quantified recovery similarly to adaptation, by the difference between the final 50 trials of the adapt and recovery phases. Though we found significant increases in gain during the recovery period across percent-feedback conditions (α=0.05, Figure 4.2D), the magnitude of recovery was always significantly smaller than the magnitude of adaptation (α=0.05). Somewhat surprisingly, recovery magnitude in the 100% feedback condition was not significantly greater than either the 50% or 75% conditions (α=0.05, Figure 4.2D). This finding suggests that varying percent feedback did not substantially impact the degree to which adaptive changes were retained. However, this analysis does not provide a picture of potential differences in recovery rate over smaller timescales, which we next sought to better characterize.

Varying percent feedback appeared to have an even smaller effect on recovery rate than it did on adaptation rate (Figure 4.2E, F). Plotting recovery magnitude for successive groups of trials (as we did to explore adaptation rate) revealed that recovery magnitude in the 100% feedback condition attained significance much earlier than other conditions (α=0.05, Figure 4.2E). Indeed, measured in 20-trial groupings, only the 100% and 75% conditions had recovered significantly by the end of the recovery phase (α=0.05, Figure 4.2E, trial-group 10). However, when we instead quantified recovery increases as percentages of adaptation magnitude, we saw that no condition reached a percentage significantly different from 0 (α=0.05, Figure 4.2F). Consistent with effects on overall magnitude, it seems that percent feedback has little impact on the retention of adaptation.

Experiment I adaptation results were consistent with those of previous studies in that intermittent feedback only weakly reduced adaptation magnitude and rate, compared to uninterrupted feedback. Recovery results, however, contrasted with past work in that we observed little to no effect of feedback intermittence on retention (recovery) magnitude or rate. A possible explanation for this difference is that we did not simultaneously control for absolute number of feedback trials and feedback-trial frequency.

Experiment II

To contrast the effects of adaptation with different feedback frequencies while keeping the number of feedback trials fixed, we conducted a second experiment. This experiment compared the 33% condition from Experiment I to a session with an adapt phase comprising only 50 feedback trials (other aspects were unchanged). As above, we focused on the magnitude and rate of adaptation and recovery.

While adaptation magnitude was unaffected, consecutive feedback trials led to faster adaptation than did those with no-feedback interspersed (Figure 4.3A). Rather than comparing successive groups of trials (as we did for Experiment I), we plucked out the gains on the 50 feedback trials from the 33% condition (red), and plotted them side-by-side with those from the consecutive condition (blue); we also plotted the first 50 feedback trials from the 100% condition, for comparison (green, Figure 4.3A). It appears that the consecutive feedback trials led to a higher rate of adaptation early-on which slowed asymptotically, while the no-feedback interspersed trials resulted in a lower but perhaps more linear rate of change (Figure 4.3A). When we compared the final 5 trials of the adapt phase across subjects (now including both feedback and non-feedback trials from the 33% condition), we found the average gains to be indistinguishable (t-test, p = 0.63). Though we previously found the 100% condition to result in greater adaptation that the 33% condition, that comparison was made after many more feedback trials in the

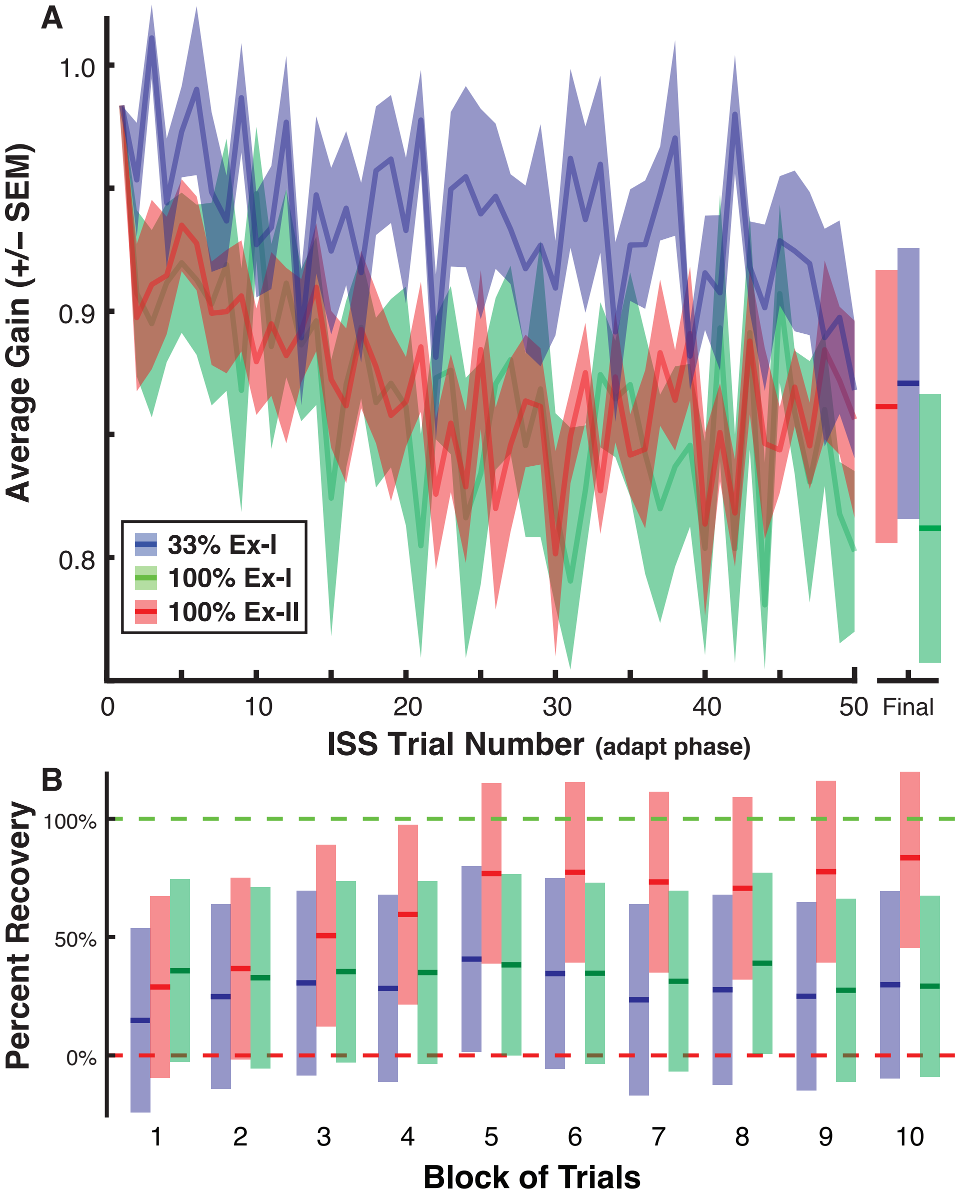

Figure 4.3 Experiment II data. **A**. Comparison of adaptation rate between conditions: the **33%** and **100%** feedback conditions of Experiment I are compared to the **100%** (uninterrupted) feedback condition of Experiment II. We aligned all the ISS trials of the **33%** condition, all the ISS trials of the **100%** condition from Experiment II, and the first 50 ISS trials from the **100%** condition of Experiment I, additively normalized indivudal subjects within conditions to set their first-trial's gain equal, and averaged across subjects, plotting mean ± SE. At far right is the mean gain in the **Final** 5 adaptation trials ± 95% confidence interval **B**. Recovery gain change progression expressed as a percentage of the average magnitude of adaptation, computed in successive blocks of 20 trials (as in Figure 4.2F).

Figure 4.3 – Chapter 4 / Experiment II Adaptation and Recovery

100% condition had occurred. Thus, we consider these adaptation results to be entirely compatible with those of Experiment II.

Recovery from consecutive feedback adaptation was more rapid than from adaptation with no-feedback trials interspersed (Figure 4.3B). Plotting percent recovery (as in Experiment I) in successive groups of trials, we found that recovery from consecutive feedback both reached a value significantly different from 0% ($\alpha = 0.05$), and indistinguishable from 100% ($\alpha = 0.05$) earlier than recovery from 33% feedback (Figure 4.3B). In fact, though 33% feedback transiently displayed a significant percent-recovery (trial-group 5, Figure 4.3B), as previously plotted, ultimately, recovery in this condition was not significant ($\alpha = 0.05$). These recovery results suggest that when the number of feedback trials is controlled for, adaptation with intermittent feedback progresses more slowly, and is retained longer.

The results from Experiment II are more fully in keeping with previous work on intermittent feedback that were those of Experiment I. In contrast to Experiment I, in which we observed no difference in recovery magnitude or rate due to feedback intermittence (Figure 4.2F), we observed a clear effect on both aspects of recovery in Experiment II (Figure 4.3B). Together, these findings suggest that feedback intermittence does affect retention rate, but that rate is also affected by the absolute number of feedback trials.

Discussion

Our data show that decreasing percent-feedback in a fixed-trial-count motor adaptation phase weakly attenuates adaptation magnitude; no significant reduction (relative to 100% feedback) was found until feedback percentage was below 50% (Figure 4.2A). Reduced percent-feedback also appeared to decrease adaptation rate (Figure 4.2C), but had no effect on the rate of the subsequent recovery (Figure 4.2F). However, decreasing percent-feedback while keeping feedback-trial-count constant (variable phase trial-count) appeared to result in slower adaptation (Figure 4.3A), and slower recovery from adaptation (Figure 4.3B).

The finding that reducing percent feedback only weakly reduces adaptation magnitude is consistent with previous work on saccade adaptation, other forms of motor learning, and operant conditioning. The only saccade adaptation study (of which we are aware) closely comparable to ours, examined the extent of adaptation in monkeys in response to intrasaccadic step (ISS) trials intermingled with no-ISS trials (Noto and Robinson, 2001), rather than with no-feedback trials (as we did). Nonetheless, their data also showed little difference in the extent of adaptation induced with 100%, 75%, or 50% ISS trials. Past motor skill learning studies have also concluded that reducing percent feedback does not decrease the extent of learning (Bilodeau and Bilodeau, 1958; Ho and Shea, 1978; Winstein and Schmidt, 1990). Finally, operant conditioning studies have similarly concluded that reducing the rate of reinforcement does not diminish the strength of a conditioned response (Pavlov, 1927; Skinner, 1938; Humphreys, 1939; Sheffield, 1949).

The finding that recovery rate varied in Experiment II and not in Experiment I is in agreement with previous motor learning and operant conditioning experiments. Several motor skill learning studies have conducted experiments similar to Experiment I, in which both percent feedback and the number of feedback trials are changed in concert (Goldstein and Rittenhouse, 1954; Annett, 1959; McGuigan, 1959; Ho and Shea, 1978; Schmidt et al., 1989; Winstein and Schmidt, 1990). In these experiments, when retention was tested in the absence of feedback, no difference in the rate of retention was found. Meanwhile, differences in retention rate have been observed in studies that varied percent feedback while keeping the number of feedback trials constant, similar to Experiment II (Bilodeau and Bilodeau, 1958; Ho and Shea, 1978). Finally, operant conditioning work that has shown an effect of partial reinforcement on the retention of a conditioned behavior also kept the total number of reinforcements constant (Humphreys, 1939; Sheffield, 1949). This perspective suggests that both the number of feedback trials and the rate at which they are delivered are modulators of retention rate. Alternatively, however, since a greater number of feedback trials necessarily means a longer duration training- or learning-period, it is possible that duration rather than feedback-trial-count is the relevant factor.

The proposition that a longer duration of training or a greater number of feedback trials results in more enduring retention is in keeping both with contemporary motor learning theory and experiment, and learning work in general. Indeed, this phenomenon, originally documented by Ebbinghaus (1913) and subsequently termed the “over learning” effect, has been observed in experiments on verbal learning (Luh, 1922), motor skill learning (Melnick, 1971; Schendel and Hagman, 1982), and adaptive motor learning (Joiner and Smith, 2008; Huang and Shadmehr, 2009). Indeed, both of the latter two studies of motor learning used experimental data and state-space modeling to demonstrate that even when additional training does not result in greater learning, it results in greater retention (Joiner and Smith, 2008; Huang and Shadmehr, 2009). Given that we saw no variation in retention resulting from the fixed-duration adapt phase but variable feedback-trial-count of our Experiment I, we must conclude that learning-period-duration and not feedback-trial-count determines retention. Parsimony also dictates that if learning period was the principal determinant of retention in Experiment I, it most likely also was in Experiment II. Thus, we cannot conclude with any certainty that intermittent feedback affected retention in the present work. Further experimentation will be required to reveal any impact of intermittent feedback alone.

In conclusion, we have found that intermittent feedback minimally impacts adaptation, while having little to no effect on learning rate. Meanwhile, intermittent feedback seems to have no effect on the retention of adaptation, instead our results confirm that the duration of the learning period does strongly modulate retention.

Chapter 5: Saccade Adaptation Specific to Visual Context[7]

Introduction

Because saccades are so brief (on the order of tens of milliseconds), they cannot be adjusted on-line through the use of visual information (feedback), which takes 40-50 milliseconds to reach the superior colliculus (Goldberg and Wurtz, 1972). Instead, any inaccuracies in the motor plan must be compensated for after the movement's completion, with corrections being manifested on subsequent saccades. In the laboratory, saccade adaptation can be demonstrated in humans by the double-step-paradigm (McLaughlin 1967), in which a target steps and, when the subject initiates a saccade toward it, the target makes a small second step back towards its starting position, resulting in the eye landing beyond the target, and provoking a second, corrective saccade. Repetitions of this series of target steps and intrasaccadic target back-steps cause a decrease in saccade amplitude so that the initial saccade brings the eye progressively closer to landing on the target. Saccade adaptation has also been observed both in human patients with disease-related weakness of the extra-ocular muscles (Optican et al., 1985) and in monkeys, in which muscles have been experimentally debilitated (Optican and Robinson, 1980; Snow et al., 1985). It is widely held that postsaccadic retinal error (the difference between target and gaze positions) drives the adaptation both in the double-step-paradigm (Wallman and Fuchs, 1998; Noto and Robinson, 2001; Seeberger et al., 2002) and in cases of attenuated muscle strength (Scudder et al., 1998).

Experiments utilizing the double-step-paradigm have shown saccade adaptation to be specific to the vector of the adapted saccade: Adaptation of saccades in one direction (left versus right) does not affect the amplitudes of saccades in the opposite direction (Deubel et al., 1986; Semmlow et al., 1987; Moidell and Bedell, 1988; Frens and van Opstal, 1994; Albano 1996) and large and small saccades can be adapted somewhat independently (Miller et al., 1981; Semmlow et al., 1987). Thus, small saccades in

---

[7] This manuscript appears with permission from the Journal of Neurophysiology, Copyright © 2009 The American Physiological Society (Herman et al., 2009).

one direction can be made larger while large saccades in the same direction can be made smaller, so that the endpoints of these saccades come progressively closer together (Watanabe et al., 2000).

The specificity of saccade adaptation also extends beyond the saccade vector to the saccade type. Adapting saccades that either are reactive responses to novel stimuli, or are voluntary shifts of gaze to existing stimuli, or are guided by memory leads to incomplete and asymmetric transfer to the other two types (Erkelens and Hulleman 1993; Deubel 1995b; Fujita et al. 2002).

Saccade adaptation can also be specific to a sensorimotor context: whether the eye is deviated up or down at the time of a horizontal saccade, or whether the head is tilted to the left or right; providing the subject with proprioceptive data from the eye or from the head muscles and otoliths, respectively (Alahyane and Pélisson 2004; Shelhamer and Clendaniel 2002a,b; Shelhamer et al. 2004). These results do not, however, challenge the view of saccade adaptation as a motor repair mechanism. When the eyes are deviated in any direction, the set of tensions on the extra-ocular muscles is changed and thus the commands sent to the muscles must be changed accordingly in order to achieve a particular movement. Similarly, because gaze position is the sum of head and eye positions, accurate redirection of gaze necessarily depends on head orientation. In this sense, it would be expected that a motor-domain repair mechanism would be sensitive to head and eye position.

However, saccade adaptation may also reflect changes beyond those required for a motor-domain saccade-repair mechanism. Amplitude adaptation of voluntary saccades transfers to arm movements (Cotti et al., 2007). There is also some evidence that directional adaptation of reactive saccades transfers to arm movements and vice-versa (Bock et al., 2008). These examples of transfer are suggestive of the adjustment of a common sensory map. This point is strengthened by the demonstration that changes in visual judgment of position accompany adaptation (Bahcall and Kowler, 1999; Awater et al., 2005; Hernandez et al. 2008). Finally, it has been reported that adaptation of saccades between visually defined objects does not transfer to similar saccades made *within* an object (Collins et al., 2007), suggesting that visual selection is necessary for adaptation effects to reveal themselves.

The flexibility displayed by the saccade system in adaptation may be a manifestation of a generalized learning mechanism. Here we ask whether saccade adaptation can depend on a purely visual context. We used two experimental paradigms: in the first, we trained subjects to selectively decrease their saccade gain (ratio of saccade amplitude to target displacement) to one of two visually distinct targets by stepping back only one of the targets during the subject's saccade to it. In the second experiment, we adapted saccades to one target and assessed to what extent the gain change transferred to the other target, which was presented infrequently and without any retinal error feedback. Our results indicate that a visual stimulus can serve as an effective contextual cue for saccade adaptation.

Methods

General

Subjects viewed stimuli in a darkened room at a distance of 57 cm while on a bite-board to minimize head movements. Stimuli were generated on a computer running VisionWorks (Vision Research Graphics, Durham, NH) and displayed on a 21" monochrome CRT display with a fast phosphor and a vertical refresh rate of 200Hz (Image Systems, Inc., LaFox, IL).

Pupil position was digitized at 240Hz, using an infrared video eye tracking system (ISCAN Inc.; Woburn, MA), controlled by a computer running SuperScope II (GW Instruments, Inc.; Somerville, MA). Immediately preceding each experimental session, a 50-point horizontal calibration was carried out by having the subject fixate a 0.3° target 10 times at each of five randomized screen locations and strike a key to acquire a 200 ms average of pupil-position. Locations were randomized to ensure that there was always a saccade between measurements. A least squares fit to these measurements was used for off-line analysis.

Stimuli

Target steps were horizontal, with amplitudes of 9-11° and directions (leftward or rightward) selected with equal probability under the constraint that the target remain within a range of ± 15°. The first trial began with the target at the center of the screen, and subsequent trials always began with the target at the end-position of the trial that preceded it.

The saccade target in each trial was either a steadily illuminated 17 cd/m^2 0.3° filled circle or one flickering (square wave) at a rate of 5Hz with a temporal contrast of 48.4% (peak: 17 cd/m^2; trough: 5.9 cd/m^2). Both were presented on a 4.7 cd/m^2 background.

Experiment I: stimulus-dependent step-back

Experiment I consisted of a baseline phase, a blocked-presentation phase and a random-presentation phase (Figures 5.1A,B). In each baseline trial, one of the two target types was stationary for 720-1220ms and subsequently stepped to a new position. It then remained illuminated for 450ms, and was briefly extinguished (for 50ms) providing the subject with an opportunity to blink.

To strike a balance between having a sufficient number of switches in target type (to analyze changes in saccade metrics at transitions between trial types) and enough consecutive trials of one type (to establish context), we arranged the trials during the baseline and blocked-presentation phases into blocks each having 3-10 trials of one target type, so that the target type alternated at an average interval of 6.5 trials.

The baseline phase consisted of 7 blocks (43 trials total), 21 trials with one target and 22 with the other. The blocked-presentation phase consisted of blocks of no-step-back trials with one stimulus (like those in the baseline phase) and blocks of step-back trials with the other. In step-back trials, after the target stepped to a new position, when the subject made a saccade to acquire it the target stepped back towards its initial location by 3° and remained there for 450 ms until the blink time. Saccades initiation was determined by eye velocity exceeding a threshold adjusted to optimize detection (on average the backstep was triggered 15.5 ± 2.6 ms (mean ± standard deviation) after the start of the saccade, 29.9 ± 4.1 ms from the end of the saccade). The blocked-presentation phase contained 63 blocks (423 trials total); odd blocks always consisting of non-step-back trials and even blocks consisting of step-back trials. Approximately every 40 trials, 1-minute-long breaks occurred, during which time the subject was instructed to either fixate the target (which remained illuminated) or close his/her eyes (18 breaks total). The random-presentation phase consisted of 40 non-step-back and 40 step-back trials, pseudo-randomly intermixed, which retained the pairing of target and trial type from the blocked-presentation phase. The overall trial ordering was the same for all sessions of experiment I.

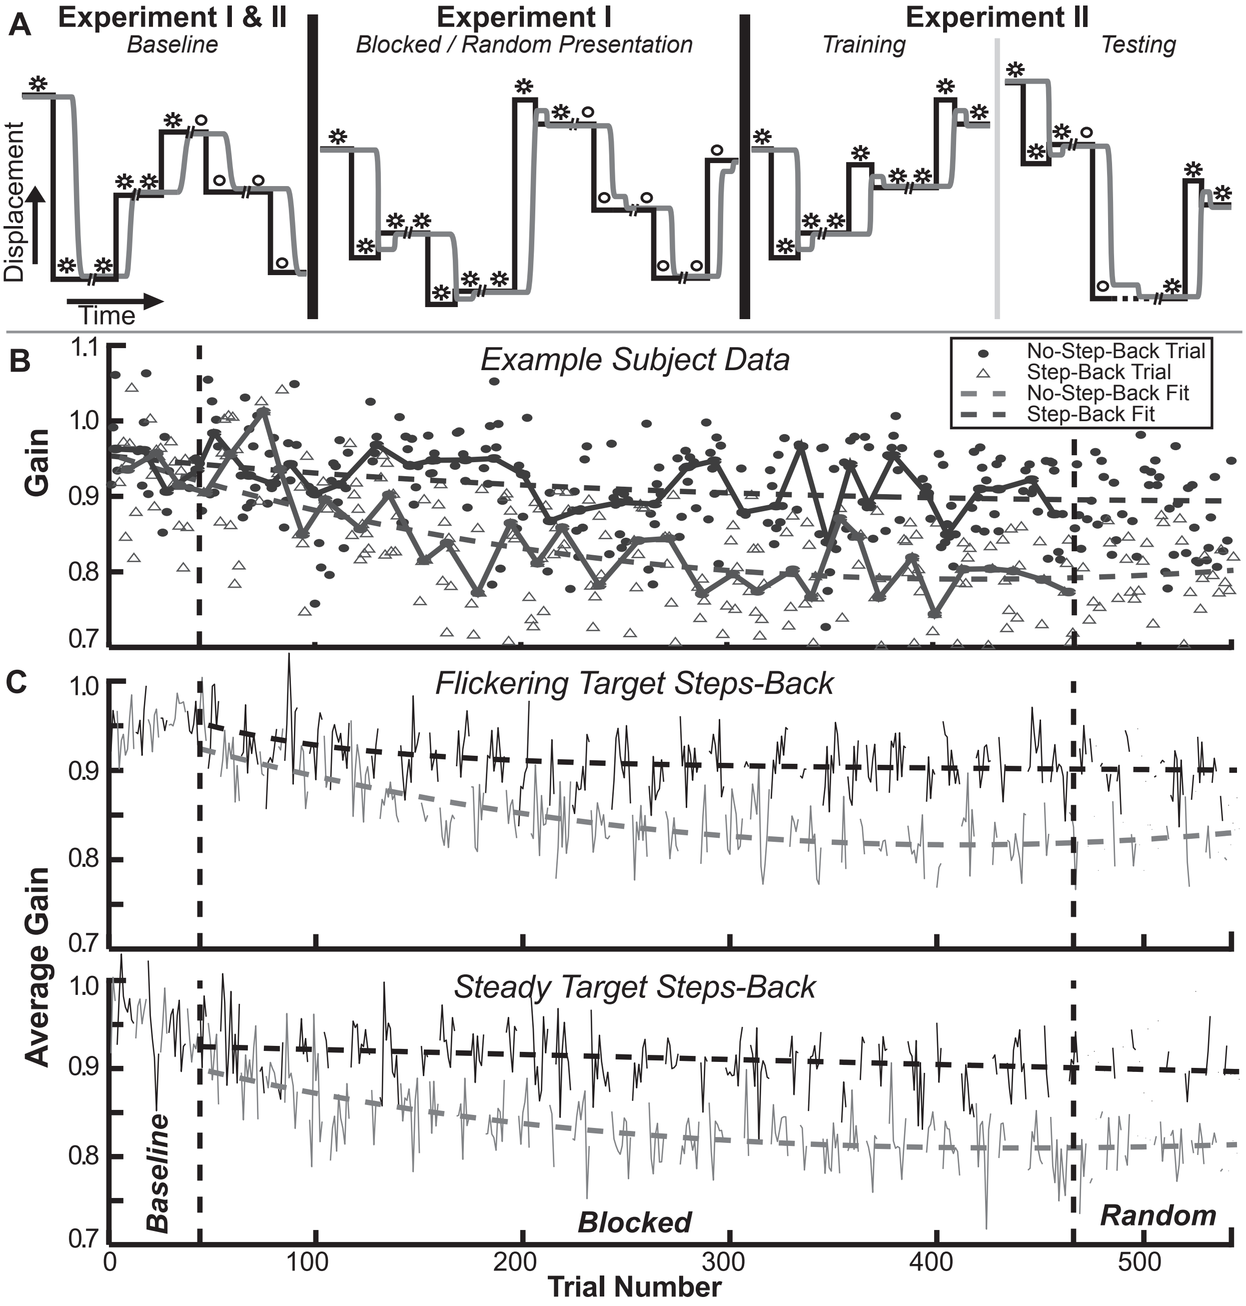

Figure 5.1 Experiment Schematics & Experiment I data. **A**. Left panel: In baseline trials, neither target type stepped back during saccades. Middle panel: In Experiment I, one target (stars) stepped back, while the other (open circle) did not. Right: Experiment II consisted predominantly of step-back trials, plus infrequent trials with the no-step-back stimulus blanked at saccade onset (OAS trials; dashed line) to deprive the subject of feedback. Black trace, target position; gray trace, eye position. Pairs of slashes represent the blink-time. **B**. An example of a single subject's trial-by-trial saccade gain. In this session, the flickering target stepped-back upon saccades to it. Solid traces connect median gains in each mini-block, and dashed traces are double-exponential fits. **C**. Averages across subjects showing a gradually increasing difference between the gains of saccades to targets that stepped back (light traces) versus those that did not (dark traces).

Experiment II: transfer-between stimulus types

Experiment II consisted of baseline, training, and testing phases (Figure 5.1A). The baseline phase consisted of 20 consecutive no-step-back trials with one target type followed by 20 with the other type. Training comprised 150 step-back trials with one of the targets to lower the gain of the saccades. The testing phase consisted of a further 410 trials, a pseudo-random mixture of 370 step-back trials and 40 off-at-saccade (OAS) trials with the other, non-adapted, target. In OAS trials, after the target stepped to a new location and the subject made a saccade to it, the target was extinguished for 450ms. It then reappeared at the same location for 450 ms until the blink-time. The testing phase began with an OAS trial, and an OAS trial was interspersed after approximately every 10 step-back trials (Figure 5.6A). This ordering was the same for all sessions and subjects.

Control experiment: conventional adaptation of flickering target

Because there were two unconventional elements to our experimental design (flickering targets and contextual switching), we ran a control experiment on three subjects (KX, MH, & XZ) with the flickering target presented alone throughout. There were 600 trials total, 100 baseline trials (no step-backs), followed by 400 step-back trials, and a further 100 recovery trials (no step-backs). As before, the primary target movement was a randomized mixture of left and rightward steps, 9-11° in amplitude, with 3° intrasaccadic step-backs.

Subjects

Six experienced subjects, aged 18-35 with normal or corrected vision, performed experiment I (AK, JH, KX, LM, MH & XZ), and 2 of these performed experiment II (KX, MH). Subjects performed each experiment twice, once with each target stepping back. Three of the six subjects (AK, JH & MH) performed experiment I with the flickering target stepping back during the first of two sessions, and the other three (KX, LM & XZ) performed the experiment with the steady target stepping back first. One subject (KX) performed experiment II with the flickering target stepping back during the first of two sessions, while the other (MH) performed experiment II with the steady target stepping back first. There was always at least 24 hours between experimental sessions. Subjects were instructed to follow the target and were told about the break periods and the blink-time. Written consent was obtained from all subjects, and the experimental protocol was approved by the Institutional Review Board of the City College of New York.

Analysis

All analyses were performed using MATLAB (The MathWorks, Inc., Natick, MA). During off-line analysis, saccades were detected automatically using a velocity threshold of 10°/s, and a minimum latency criterion of 100 ms to exclude anticipatory movements. The start and end of each saccade was confirmed and, if necessary, corrected by the experimenter. Approximately 2% of all saccades were excluded from further analysis.

Results

General

In general, our results show that when one target-type (flickering or steady) steps back during saccades and the other does not, saccades to the step-back target develop consistently smaller gain than those to the no-step-back target. Furthermore, if one target-type steps back while the other target type is extinguished during a saccade, no difference in gain is observed between the two targets after adaptation.

Experiment I: stimulus-dependent step-back

The results of experiment I indicate that subjects are able to selectively reduce the gain of saccades made to one of two visually distinct targets. Figure 5.1B shows the gain on each trial for one subject in an session in which the flickering target stepped back (SB) and the steady target did not (NSB). With the onset of the blocked-presentation phase, the gain of saccades to the flickering target decreased at a greater rate than those made to the steady target. The gain of this subject's saccades to the two targets was significantly different in both the blocked-presentation and random-presentation phases (mean difference in gain: blocked-presentation-phase, 0.08±0.1 SD; random-presentation-phase, 0.09±0.1, both $p<0.01$, two-tailed t-test). The gain of saccades to the SB target continued to decrease gradually over the course of the session, despite a small increase during the random-presentation phase; meanwhile, the gain of saccades to the NSB target showed a much slower rate of decline.

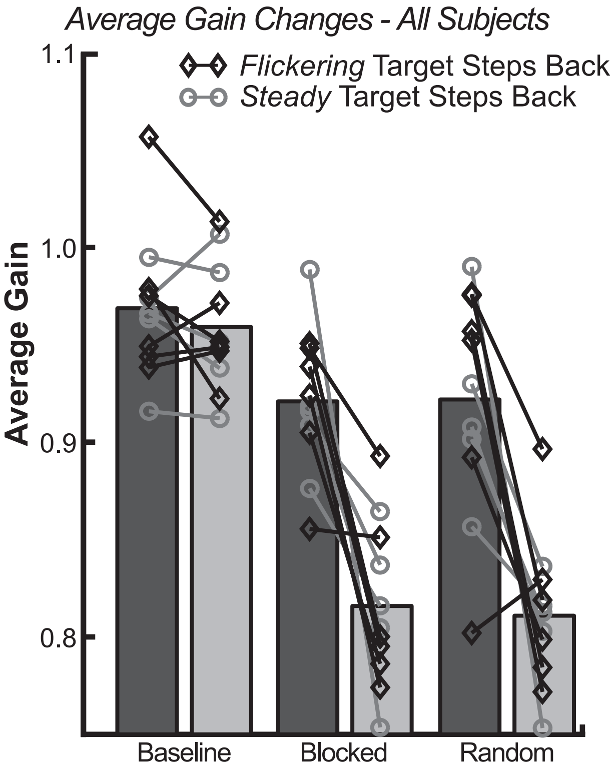

Figure 5.2 Average Gain Differences during Experiment I. Bars represent averages across subjects without normalization; Dark bars: no-step-back trials, light bars: step-back trials. Symbols and lines represent average gains of individual subjects. Data from blocked-presentation phase is from the second half of that phase (trials 212-466).

Figure 5.2 – Chapter 5 / Experiment I Average gain changes during Experiment I

Averaging each trial across subjects (after an additive normalization of each session's data-set to equalize mean baseline gain), we found lower gain for saccades to the target that stepped back, whether it was the flickering target (Figure 5.1C, top) or the steady target (Figure 5.1C, bottom). Specifically, gain was reduced to 85% of baseline for the SB target, and 94% for the NSB target (ratio of the average of the final 20 trials in the random-presentation phase to the average baseline trials, across subjects). Thus, the context-specificity of the adaptation was not absolute. Considering each subject in each target-type condition, the gain was significantly lower to the SB target in 11 out of 12 experimental sessions in both the blocked-presentation and random-presentation phases. On average (across subjects), saccade gain was 0.92±0.04 for the NSB and 0.82±0.04 for the SB target during the second half of the blocked-presentation phase (trials 212-466) and was 0.92±0.06 for the NSB target and 0.81±0.04 for the SB target during the random-presentation phase (Figure 5.2). A mixed-effects Analysis of Variance (ANOVA; four factors: trial type, session phase, experiment type (which stimulus stepped back) and subject) showed that the gain of saccades to the SB target was significantly smaller than those to the NSB target (F = 61.62, $p<0.01$). This analysis also showed that the saccades to the NSB target were significantly smaller during blocked & random presentation phases than they were during the baseline phase ($p<0.01$, Tukey-Kramer post-hoc test). Finally, the same ANOVA revealed no significant effect of which stimulus was the SB and which was the NSB target, on saccade gain (F = 0.56, p = 0.45).

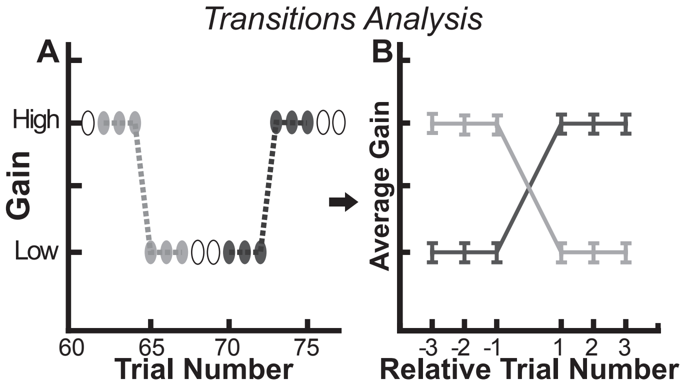

Figure 5.3 Transition Analysis. Schematic of method for examining changes in gain at transitions from one target-type to the other in the blocked presentation phase. **A**. The gain values for 3 trials before and after each transition were extracted; light and dark ovals indicate transitions from no-step-back to step-back and step-back to no-step-back blocks, respectively; unfilled ovals indicate trials not included in transition analysis. In blocks of fewer than 6 trials, some trials contributed to both traces. **B**. Idealized traces resulting from averaging across transitions shown in A.

Figure 5.3 – Chapter 5 / Experiment I Transitions analysis schematic

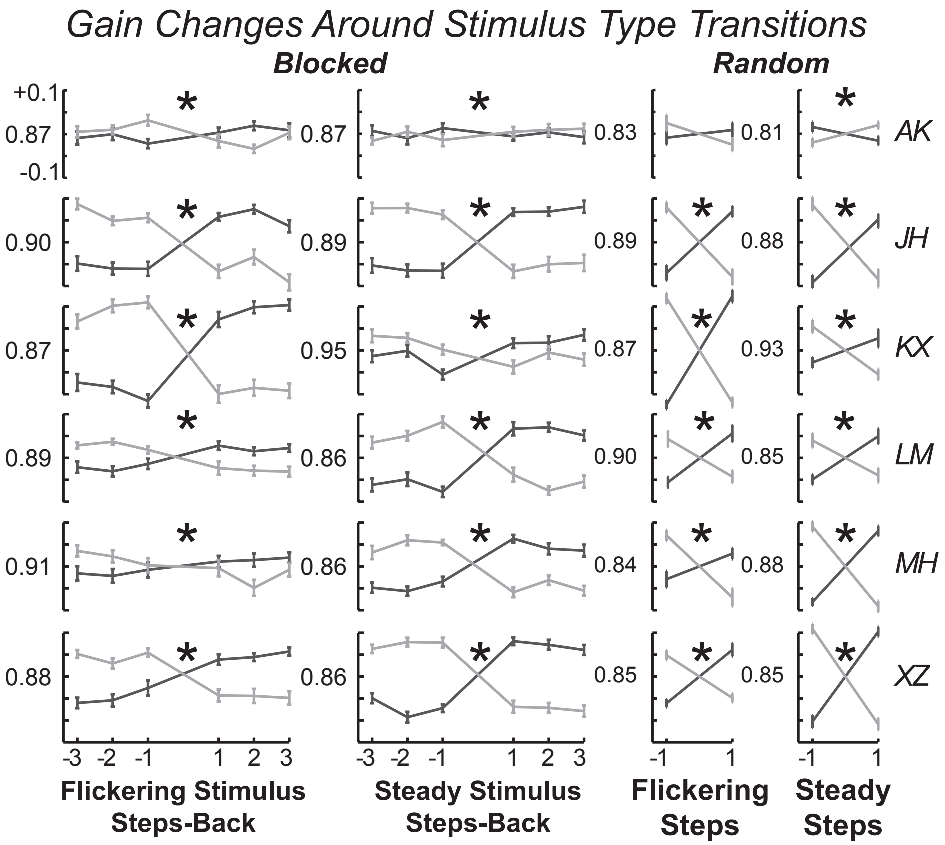

Figure 5.4 Gain Changes Around Transitions. Transitions from step-back to no-step-back (dark) and vice versa (light). In all but one case the gain changes at the two transitions differed significantly, with no difference between cases in which the target that stepped back was flickering (1st & 3rd columns), versus steady (2nd & 4th columns). (*, $p<0.01$ by one-tailed Tukey-Kramer post-hoc test). Note: as indicated in the upper left, each ordinate has the same span (0.2) but has a different central value.

Figure 5.4 – Chapter 5 / Experiment I Average gain changes around transitions

To quantify the purely contextual aspect of the adaptation, independent of the gain changes within each block of trials of one type, we examined the transitions from one stimulus type to the other. For the blocked-presentation phase of each session, we extracted the last 3 trials of a block of one stimulus type and the first 3 trials in the next block, making a vector of length 6 of the saccadic gains (Figure 5.3). Because each block had a minimum of three trials, this kept the sample sizes constant for statistical purposes. This procedure resulted in one set of 32 vectors for the transitions from step-back to no-step-back, and one for the transitions from no-step-back to step-back. For the random-presentation phase, in which the stimulus types were randomly interleaved, we extracted the last trial of one stimulus type and the first trial of the following stimulus type, to avoid having different numbers of points extracted from different transitions (20 transitions of each type). We plotted each subject's average transition vector from SB to NSB and NSB to SB separately for blocked-presentation and random-presentation phases, and for data from sessions in which the flickering target stepped back versus those in which the steady target stepped back (Figure 5.4). Across sessions, during the blocked-presentation phase, we found that when the target type switched, the gain of the first saccade was 0.08 ± 0.024 (mean ± SE) higher to the NSB target, and 0.08 ± 0.022 lower to the SB target. During the random-presentation phase, the gain of the first saccade was 0.1 ± 0.03 higher to the NSB target and 0.12 ± 0.029 lower to the SB target. To test the significance of these gain changes, both within and across subjects, we compared the gain on the trials before a switch to the gain on the trials after a switch with a two-factor ANOVA (four levels in one factor: last step-back trial, first no-step-back trial, last no-step-back trial, first step-back trial, with subjects as a second factor). Using a Tukey-Kramer post-hoc test, we found that the last step-back trial was significantly different from the first no-step back trial and that the last no-step-back trial was significantly different from the first step-back trial for each subject (all $p<0.01$), and across subjects ($p<0.01$), during both the blocked and random-presentation phases.

Gain change on the first trial of each block increased progressively across the adaptation session (Figure 5.5A). To show individual variations we plotted the cumulative sums of these gain changes over the course of each session (Figure 5.5B). In other words, each time that the stimulus switched, we computed the change in gain; summing these up according to the transition type to obtain the cumulative record (step-back to no-step-back or vice-versa). Some subjects seemed to learn the context almost immediately, indicated by a difference in slope of the cumulative sums that began after just a few blocks (Figure 5.5B: JH, XZ). Other subjects required more substantial exposure to the differential adaptation (Figure 5.5B: KX, both columns; LM, left column; MH left column). In some of the traces the slopes start to increase a bit, late in the session, as though the gains to the SB and NSB targets were continuing to diverge (Figure 5.5B: LM, XZ). This is also reflected in the larger transition sizes for some subjects we observed in the random-presentation phase versus the blocked-presentation phase of the experiment (Figure 5.4: MH, XZ).

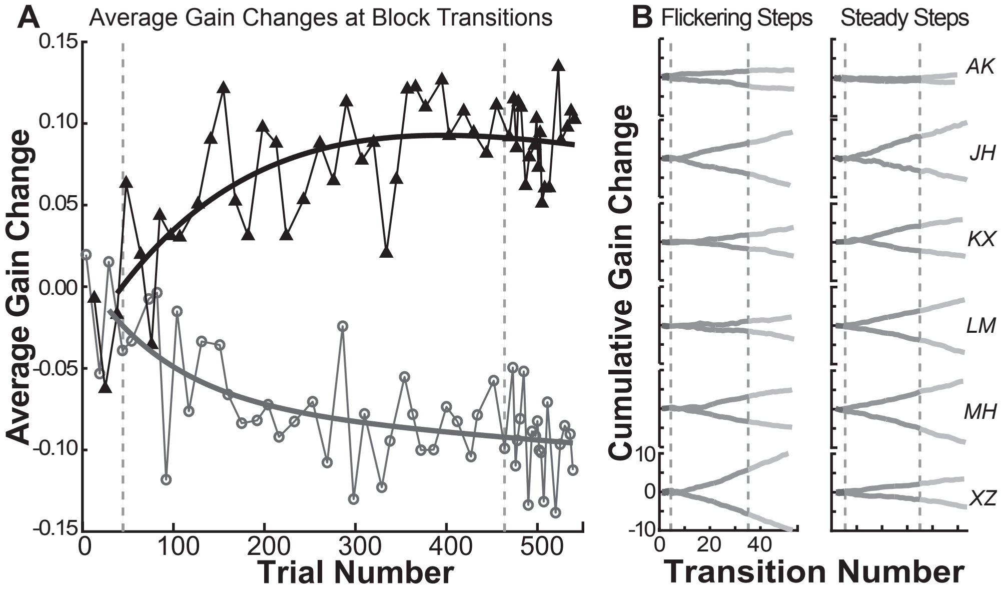

Figure 5.5 Development of Gain Changes at Transitions. **A**. Average gain changes at stimulus transitions; Dark trace/triangles: gain change on switching to a no-step-back trial, light trace/open circles: gain change on switching to a step-back-trial. Curves are double exponential fits. **B.** Cumulative sums for each session show individual differences in the rate of acquisition of fast gain changes at transitions. Vertical dashed lines separate baseline, blocked, and random phases.

Figure 5.5 – Chapter 5 / Experiment I Development of gain changes at transitions

We found it curious that, on average, most of the gain-change between successive blocks occurred abruptly at the first trial of each block. Because these gain-changes were larger than would be expected for any one trial under conventional adaptation, it appears that the gain values reached for each target are stored with little decrement between blocks. To evaluate whether there was any gain change beyond the first trial of each block, we tested how quickly the gain of the first, third and last trial in a block deviated from baseline by using an ANOVA with Tukey-Kramer post-hoc tests as a metric (two factors: first, third and last trials as the first factor; block number as the second factor; data pooled across sessions). The gain of the last trial of a step-back block deviated significantly from baseline after only 4 blocks, and the third trial deviated from baseline after 8 mini-blocks, whereas the first trial deviated only after 11 blocks. We take this as an indication that the adaptation progressed within each block and that it took some time within a session for the fast switching of gain on the first trial of a block to be manifested.

If the reaction time of saccades to one or both of the stimuli were longer during the differential adaptation, it might reflect a cognitive strategy. Furthermore, when two perceptual tasks are interleaved, the reaction time of the first response to the new task is frequently longer than the subsequent ones – the switch cost (Barton et al. 2006; Altmann 2007). In experiment I, reaction times were much shorter than one might expect for cognitively driven saccades (Deubel 1995) and did not significantly differ between adaptation and baseline phases (median = 168 ms ± 29 SD for random-presentation phase vs. 164 ms ± 28 for baseline phase; Tukey-Kramer post-hoc test, p = 0.6). Furthermore, the latency of the first saccade of a new block was, on average, only 5.5 ms ± 2 longer than that of the next saccade, a non-significant change (p = 0.31 t-test).

We did observe a significant difference in the distribution of latencies of saccades made to the steady (165 ms ± 33) versus the flickering (179 ms ± 40) stimulus (F = 458.43, $p < 0.01$). However, because the choice of stimulus used as the step-back target had no statistically significant effect on the experimental outcome, this difference in latency is unlikely to have influenced our main result.

Experiment II: transfer between stimulus types

Experiment I employed stimulus-dependent step-backs, which concurrently guided subjects to decrease the amplitude of their saccades to the SB target and to keep the amplitude of their saccades to the other target unchanged. To determine the necessity of this differential feedback, we measured the transfer of adaptation from the SB to the other target by presenting the other target sporadically, and extinguishing it at saccade onset (OAS) to remove post-saccadic retinal error (Figure 5.1A, right). One subject's results are plotted in Figure 5.6A, in which gains of saccades to the OAS stimulus were indistinguishable from those to the SB stimulus; that is, in the absence of differential feedback, complete transfer occurred. To compare the gains of saccades to the 2 targets, cumulative distributions of OAS trials and step-back trials (after trial 140) were calculated for each subject under conditions in which either the flickering or steady target stepped back. These distributions were compared individually using the Kolmogorov-Smirnov test statistic (Figure 5.6B). In none of the 4 sessions were the gains of saccades to the OAS target greater than those to the SB target. A Receiver Operating Characteristic (ROC) analysis

(Figure 5.6B), confirmed this result. The ROC analysis in this case provides an estimate of the fraction of trials that one would be able to classify, on the basis of a single saccade, as having contained the SB or OAS target. Further, we examined the gains using a two-factor ANOVA with step-back trials vs. OAS trials as one factor and subjects as the second factor. In Tukey-Kramer post-hoc tests, the step-back saccades were not statistically distinguishable from the OAS saccades for each subject (all p>0.5).

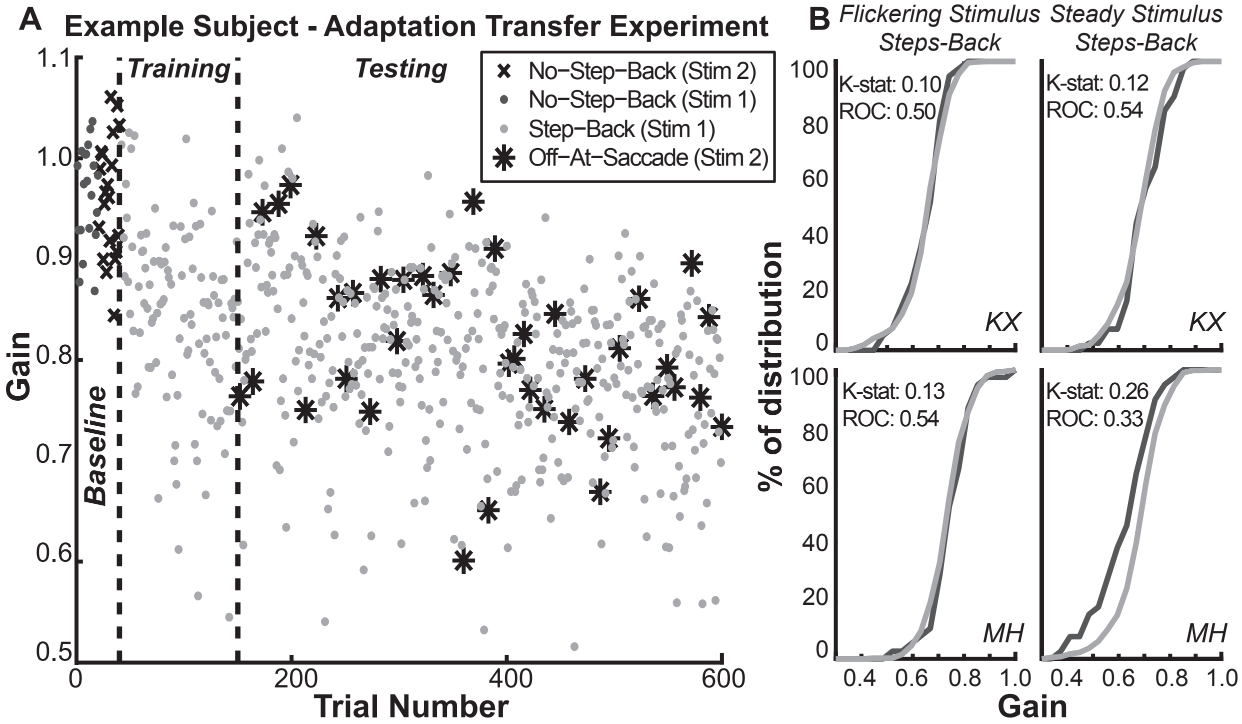

Figure 5.6 Adaptation Transfer Results. **A**. Similar gain changes of saccades to both targets when no-step-back target (stars) is extinguished at saccade onset to eliminate feedback (OAS trials), amid predominant presentation of the step-back-target with feedback (grey dots). Vertical dashed lines separate baseline, training and testing phases. **B**. Comparison of cumulative distributions of the gains of saccades to the step-back (light trace) versus the no-step-back OAS target (dark trace) with Kolmogorov-Smirnoff and Receiver Operating Characteristic (ROC) statistics show nearly complete transfer of adaptation

Figure 5.6 – Chapter 5 / Experiment II Adaptation transfer results

There was no noticeable difference in the degree of adaptation depending on whether the flickering or the steady target served as the SB stimulus; these were also indistinguishable from the control experiment. A single factor analysis of variance comparing gains of saccades (across sessions) from the final 40 trials in each session (including only step-back trials from experiment I, and both step-back and OAS trials from experiment II), showed them to be statistically indistinguishable (Experiment I: flickering:

0.81 ± 0.078 (median gain ± SD), steady: 0.81 ± 0.08, p = 0.32; Experiment II: flickering: 0.79 ± 0.067, steady: 0.81 ± 0.062; Control: 0.8 ± 0.06, F = 0.92, p = 0.44).

To compare the rates of adaptation for flickering and steady targets, we split the first 120 step-back trials into 6 blocks, separately for each session-type (flickering versus steady target stepping back) of experiments I and II and the control experiment of conventional adaptation with the flickering target. We again pooled across subjects after additively normalizing to the mean baseline gains, so that the variances reflected only differences in the course of adaptation and not individual's baseline variances. We found that as gains gradually decreased there were no significant differences between these five groups of SB gains after 20, 40, 60, 80, 100, or 120 trials (2 factor ANOVA; first factor: experiment type, second factor: number of preceding step-back trials; F = 156.95, $p<0.01$), leading us to conclude that the rate of adaptation in these experiments was approximately the same, despite the fact that during experiment I the interspersed NSB trials might have decreased the rate of adaptation for the SB target.

Discussion

Our results show that the same target step can elicit saccades of different gain, depending on the visual attributes of the target, as a consequence of a learned contextual association (Figure 5.2). Furthermore, the results demonstrate that the context-dependent gain difference develops gradually (Figures 5.1A, 5.5A), much as in conventional saccade adaptation experiments, and most of the gain difference is evident on the first saccade after a switch of target type (Figures 5.4, 5.5A), implying that only a small part of the gain difference is due to the progressive gain changes within each block of trials of one target type. Finally, the context dependence of gain is not absolute in that the gain of saccades to the target that does not step back is also reduced (Figures 5.1, 5.2).

Our study of contextual saccade adaptation follows an earlier, unsuccessful attempt to differentially adapt saccade gain (Deubel, 1995a). That study differed from our own in that: (1) the ordering of the presentation of the two stimuli was random throughout; (2) the subject foveated a fixation target which changed into one of the two stimuli at the time of the initial target step; (3) the stimuli were differentiated by color and shape (the targets were a red circle and a green cross). What might account for the difference in the results between these two studies? First, trial structure may be important. If postsaccadic error acts most strongly to change the gain of the immediately following saccade, random target presentation may slow contextual learning since half of the trials would be followed by the other target-type. On the other hand, long blocks of each stimulus would allow the subject to ignore the target-type and be guided by whether or not the target stepped back; that is, the target step-back during blocked-presentation could *itself* act as a context, since each trial predicts with high probability that the following trial will be of the same type. For these reasons, we used short blocks of irregular length. Second, in experiment I, the subject fixated the saccade target for as long as 1.2 seconds prior to the target step. Perhaps a long fixation period provides the subject with time to assign the context-specific gain appropriate to the presented target, whereas in Deubel's paradigm the target step may have initiated the programming of a saccade before the context had been decoded. Finally, it is conceivable that shape and color are simply not effective contextual cues for saccade adaptation, as opposed to the random trial structure in Deubel's experiment having masked learning with these features. Thus the difference between our results and those of Deubel suggests that there are limits to the generality of the associations that saccade adaptation can learn.

One might question whether the adaptation shown in this study is the same as conventional adaptation, or whether it is a manifestation of a conscious targeting strategy, in which the subject explicitly recognizes the contextual association and voluntarily produces hypometric saccades or “aims” his or her saccades to a position short of the target. We are disinclined to accept this explanation, largely because the adaptation develops gradually over hundreds of trials (Figure 5.1C). Even after the first 100 adapt trials the gain continued to fall in 10 of the 12 experimental sessions (in which the slopes of linear regressions of gain against trial number 101-546 were negative and significantly non-zero; $p<0.05$). If a subject were consciously directing saccades to a particular unmarked point on the screen, one might expect that the variability would diminish with practice, but the gain would not change progressively over hundreds of trials. Furthermore, the finding that there was some transfer between the gains to the two target types is more suggestive of incompletely separated contextual adaptation than of cognitive strategy. Indeed, we suspect that the target-specificity we have observed reflects the acquisition of two contexts, one for each target. In experiment II, in which one target type always was extinguished upon saccade initiation, the gain of saccades to that target decreased as much as those to the other target, which stepped back. Therefore, it is likely that in experiment I, subjects were actively maintaining an elevated saccade gain for the NSB target while decreasing gain for the SB target, as a result of their differing postsaccadic consequences.

Despite our contention that it is not highly cognitive, the contextual adaptation we observe clearly uses a more elaborate error signal than simply accumulating retinal error. Although many experiments support the importance of simple retinal error (e.g. Wallman and Fuchs, 1998; Noto and Robinson, 2001; Seeberger et al., 2002), these do not exclude an error signal incorporating a prediction of the retinal error expected, given the saccade executed. Indeed, there is some evidence that predicted retinal error does have a role to play in saccade adaptation (Bahcall and Kowler, 2000), and predicting the sensory consequences of action is, more widely, seen as an important teaching signal in motor adaptation

(Wolpert and Ghahramani, 2000). If the functional utility of motor adaptation is to improve accuracy, then it is plausible that other, higher-level, predictions could also be taken into account.

A theoretical framework for multiple predictions (or 'forward models') has been widely used in the study of arm movement adaptation and has been supported by many empirical data (Kluzik et al., 2008; Wagner and Smith, 2008; Wolpert et al., 1998). In motor learning theories of this type, an inverse model is used to compute the motor commands necessary to achieve a desired trajectory, and a forward model predicts the sensory consequences as the movement occurs. In keeping with the known involvement of the cerebellum in motor learning, it has been suggested that there are multiple, overlapping internal models localized in the cerebellum (Kluzik et al. 2008; Wolpert et al. 1998). Multiple models could facilitate efficient context-dependent motor learning and execution, if the cerebellum used a sensory context signal to flexibly select the model (or sets of models) appropriate for motor execution and/or the models to be modified during motor learning (Wolpert et al. 1998). Although one might argue that this type of contextual flexibility is essential only for motor systems that have to deal with varying loads, adaptation of the vestibulo-ocular reflex, which is known to be mediated by the cerebellum, can be specific to head-tilt and head-rotation frequency, velocity and acceleration (Boyden et al., 2004; 2006). Moreover, the separate adaptabilities of saccades to sudden target steps compared to voluntary saccades to stationary targets (Alahyane et al., 2007; Erkelens and Hulleman, 1993; Deubel, 1995b; Fujita et al., 2002; Gaveau et al., 2005) may reflect separate cerebellar loci (Alahyane, et al. 2008), consistent with at least partially separate internal models for each movement type. In experiment I, we speculate that the lack of consistent trial-to-trial retinal error -- in contrast to conventional saccade adaptation experiments -- led to a persistent dissonance between predictions and observations of that error. This continued dissonance may have coaxed the oculomotor system into searching for another signal that might be predictive of retinal error, exploiting the consistent relationship between stimulus type and target movement. If multiple context-dependent internal models exist, it would not be surprising if some subset of models might be shared, with the effect that adapting one movement would transfer somewhat to the other, as we have found.

Contextual learning might be a general feature of motor adaptation. In the directional adaptation of arm movements, the presence of distinct proprioceptive contexts permits two tasks to alternate without interference (Krakauer et al., 2006). In the case of saccade adaptation as well, distinct proprioceptive contexts prevent increased-gain adaptation from interfering with the progressive course of decreased-gain adaptation (Shelhamer and Clendaniel, 2002a,b; Aboukhalil et al., 2004). In addition, there is an indirect suggestion of a contextual effect on saccade adaptation in that decreased gain can be retained for long intervals (up to five days) following the initial step-back adaptation session, despite the subject presumably having made hundreds of thousands of saccades in between (Alahyane and Pélisson, 2005). Because, in the laboratory, having subjects make saccades without step-backs following adaptation restores "normal" saccade gain in a few hundred trials, the prolonged retention observed may result from the contextual difference between daily life and the laboratory setting – the sights, smells and sounds experienced during the experimental session(s), as well as the taste of the bite-board, and the seated posture of the subject. Alternatively, the retention may be due to the low frequency in daily life of reflexive saccades (driven by abrupt stimulus steps) compared to voluntary saccades, and the partial independence of the adaptation of these two types mentioned above.

Why might it be useful to customize saccadic parameters based on a sensory context? Whatever error signals are employed to guide saccade adaptation, the oculomotor system would benefit from being able to distinguish saccades that were inaccurate because the target had moved during the saccade from those that were inaccurate because of a miscalibration of the saccadic system. A recent model of saccade adaptation posited that large errors are attributed to the target having moved, while small errors are attributed to errors in the saccadic system (Chen-Harris et al., 2008). We conjecture that, from an ethological point of view, it would be more efficient for the oculomotor system to make use of the stimulus context to distinguish between situations in which the targets behave predictably, in which case errors could be attributed to the motoric performance, and those in which saccades are used to track targets moving irregularly, in which case the errors could be attributed to the target motion. In both situations, saccade adaptation might be used, but it would be useful to separate the parameters of the first case, in which saccade adaptation tunes up the motor system for saccades to stationary objects, from the second case, in which saccade adaptation is used to compensate for statistical regularities in the target movement. From this viewpoint, it might not be surprising to find the aforementioned differing adaptabilities for internally and externally motivated saccades.

In conclusion, our results extend the notion that saccades can have two gains for the same target vector under different sensorimotor contexts, by showing that purely visual attributes of target stimuli can define the context. It remains to be determined whether this ability serves a particular oculomotor function or is a manifestation of saccade adaptation, like other forms of learning, being sensitive to any informative context.

General Discussion

The five experimental chapters of this manuscript aimed at investigating the processes underlying saccadic adaptation in human subjects at different levels ranging from trans-saccadic target detection to contextual effects. We can draw a number of conclusions from the combination of our results. First that multiple signals can be used to track a target across saccades but that the presence of an object trumps the other, more low level, mechanisms (chapter 1) and that this feature-sensitive signal defining the object identity can drive adaptation (chapter 2). We have also shown that the feature information about the target is strong enough to be a contextual cue for adaptation (chapter 5). Secondly, we found that the predictability of a target's post-saccadic location was not a facilitative factor in adaptation and therefore can't be the main signal driving adaptation (chapter 2). Neither can retinal error as we demonstrated that ISS size best explained the amount of adaptation (chapter 3) and that this step could be as small as 0.25° and still be effective, even if smaller than subjects' saccade end-point variability (chapter 3). This latter result indicates that the oculomotor system has strong non-visual evidence about the landing position of the eyes and that this information is likely used in the calculation of the error signal to produce a predicted retinal error. Finally, we showed that giving feedback infrequently only minimally impacts the strength and rate of the adaptation and had surprisingly little effect on the retention (chapter 4).

We are going to discuss these results in the light of the pertinent literature, and within the frame of the three questions used to segment the background section of the introduction to this thesis: What is the error signal? How are errors detected? How are errors used?

What is the error signal?

In chapter 2, we adapted subjects' gain both up and down under conditions with more and less predictable post-saccadic target positions, finding that there was greater adaptation with the more predictable target. However, the more predictable condition also had a larger mean intrasaccadic step (ISS). When we instead calculated adaptation as a percentage of that mean ISS, we found that adaptation percentages in more and less predictable conditions to be indistinguishable (separately for gain-increase and decrease; Figure 2.2). This suggests that ISSs alone may constitute the error signal for saccade adaptation. In keeping with this, the step-wise linear regression we performed in chapter 3 revealed that ISSs alone were sufficient to explain the adaptation results we obtained with systematically varied small ISSs. The idea that ISSs might be a candidate for the error signal is not without basis, as pointed out in the introduction, ISSs are on average, equivalent to sensory prediction error (SPE; Figure 0.4). Thus, it is not entirely surprising that ISSs should be able to explain aggregate data in this manner. However, in the experiments of chapter 2, it also seemed that subjects' predictions regarding post-saccadic target position might be changing as a result of adaptation, reflected by their tendency to execute corrective saccades towards likely post-saccadic target locations, even when the target was elsewhere (Figures 2.4 & 2.5). This is somewhat of a contradiction. Another caveat is that it is unclear why gain-increase adaptation should have achieved such a minimal percentage of the average ISS compared to gain-decrease in the data of chapter 2. If ISSs represent the error signal, there is no reason for gain-up / gain-down asymmetry. Of course, this asymmetry would be no better explained by an SPE signal driving adaptation. At present an explanation for the discrepancy between the explanatory power of ISSs and the strong evidence we have that sensory predictions are modified by adaptation is not forthcoming. A full resolution will require an analysis (of the data of chapter 2 in particular) focused on understanding the ongoing relationship between SPEs and ISSs as sensory predictions shift during adaptation.

Because it involved adaptation under conditions with feedback frequently withheld, the experiments of chapter 4 seemed like a potentially well-suited forum for exploring the role of prediction in adaptation. If, indeed, an SPE is driving adaptation, withholding feedback might serve as a useful tool to clarify the details of its calculation. Thus, we implemented two state-space models, with two different SPEs driving adaptation. This class of models has been used previously to reveal and illustrate multiple-timescale motor learning phenomena (Ethier et al., 2008a, Smith et al., 2006), and to explore other training effects on adaptation rate (Srimal et al., 2008). The first model (which we call the BRE-model) was explicitly taken from Ethier and colleagues (2008a), which uses a bias-subtracted retinal error (mentioned above); this form of SPE assumes that prediction is linked only to the mean saccade gain, and the constant term which corresponds to undershoot bias or hypometria (Becker, 1989). The second model (termed the

PRE-model) used the difference between a predicted retinal error, and actual retinal error; this form of SPE assumes that prediction is based on an internal estimate of saccade amplitude, and on the history of retinal errors. Thus, the first model assumes that changes in error prediction during adaptation result directly from changes in gain, while the second model assumes that gain and prediction may change somewhat more independently. For a full description of model implementation, see Appendix B. Briefly, however, both models have “hidden memory” (only indirectly observable), characterized by a retention rate ***A***, and learning rate ***B***; after each primary saccade, the new memory state is calculated by retaining a portion (***A***) of the previous state, and learning (adding) a portion (***B***) of the error resulting from the just-executed saccade.

We found that accurate modeling of observed behavior was provided only by the second, PRE-model. The BRE-model fits captured observed behavior well in some cases, but not consistently, whereas the PRE-model behaved well across subjects and conditions. This was confirmed by subjecting squared errors to an analysis of variance (ANOVA), which revealed a significant effect of model-type corresponding to the PRE-model’s significantly smaller mean squared error ($F=102.08$, $p=5.7x10^{-24}$). The superior performance of the PRE-model suggests that the obligatory coupling of prediction to gain state may be an oversimplification. This is well in-line with the findings of chapter 2, which showed clearly how predictive behavior (as measured by corrective saccade production) and gain changes were strongly uncoupled: gain-increase adaptation was weaker but resulted in more predictive behavior, while gain-decrease adaptation was stronger but resulted in less predictive behavior.

Another point that may account for differential model behavior is that (aside from how putative *changes* in prediction are handled) the BRE-model assumes that predictions are calculated using mean saccade amplitude only, further variation is not taken into account. Our view, encapsulated in the PRE-model and supported by several lines of evidence (Sparks and Mays, 1983; Bock et al., 1995; Ditterich et al., 1998; Sommer and Wurtz, 2004; Collins et al., 2009; Munuera et al., 2009; Joiner et al., 2010), is that the oculomotor system (and thus the putative error prediction mechanism) has access to a substantially more accurate (internal) prediction of saccade amplitude than just the mean. Thus, while the use of bias-subtracted retinal-error represents one extreme: that none of the saccade-amplitude variation around the mean is known to the oculomotor system; we chose the other: assuming the system has perfectly accurate ability to estimate saccade amplitude internally, at the service of computing a predicted retinal error. This is in fact of substantial importance, because this is the term that will effectively set the model's tolerance to ISSs. If the internal amplitude estimate is subject to a great deal of noise, it is impossible for the subject to know whether a small difference between predicted error and actual error is genuine, or results from a corrupted prediction of sensory consequences stemming from that noisy internal estimate (Figure 6.1). This is a clear shortcoming of our PRE-model as it assumes too good of an internal estimate of saccade amplitude (no noise incorporated), and must be addressed in further work.

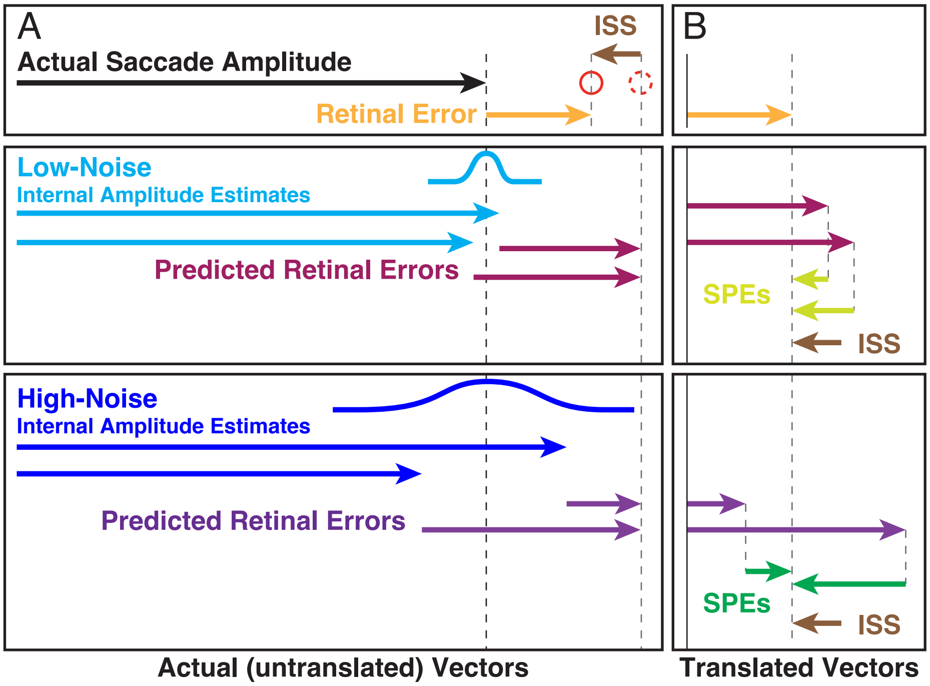

Figure 6.1 Internal Amplitude Estimate Noise-level Sets Sensitivity to Error. A. The **actual saccade amplitude** determines the retinal error, while the internal estimates determine the predicted retinal error values (**magenta for low-noise**, and **violet for high-noise**). A **low-noise internal estimate** results in two similar **predictions of retinal error**, while a **high-noise internal estimate** results in more disparate **predictions of retinal error**. B. Translating the **actual retinal error**, **low-noise prediction** and **high-noise prediction** vectors so that they have the same start-point allows SPE differences to be represented graphically as arrows from the head of each prediction vector to the head of the retinal error vector. Note that the **SPEs resulting from the low noise predictions** are similar to the **ISS**, while the SPEs from the **high-noise predictions** differ substantially from the **ISS**.

Figure 6.1 Internal Amplitude-Estimate Noise and Sensitivity to Error

Though retinal error has long been implicated as the driving force behind saccade adaptation, our modeling results join a growing body of evidence that challenges that view. Two works have been particularly influential in establishing the primacy of retinal error as the driving force behind saccade adaptation (Wallman and Fuchs, 1998; Noto and Robinson, 2001). However, these works asked whether retinal error or corrective saccade amplitude (which are highly correlated in conventional intra-saccadic target step adaptation paradigms) drives adaptation by devising paradigms which eliminated corrective movements, leaving only post-saccadic visual information responsible for any resulting adaptation. Thus, these works are entirely compatible with the notion that deviations from predicted retinal error, rather than retinal error alone, drives adaptation. Two recent papers, however, did directly assess the possibility that retinal error or deviations from retinal error prediction better explains adaptation results under experimental conditions designed to disambiguate these hypotheses (Wong and Shelhammer, 2010; Collins and Wallman, 2012). Both of these works concluded that deviations from predicted retinal error best explained their results. These works also acknowledged, however, the alternate possibility that adaptation was driven by the goal of maintaining an undershoot bias. However, the failure of the BRE-model in fitting our data provides evidence to the contrary since that model was driven by a bias-subtracted retinal error which amounts to being “driven by the goal of maintaining an undershoot bias.”

How are errors detected?

Our results from chapter 1 identified 4 mechanisms that can be used to detect displacement and drive corrective saccades: (1) one intrasaccadically sensitive to a low-level motion signal, (2) one trans-saccadically sensitive to a high-level or attention-based motion signal, (3) a trans-saccadically operating low-level cross-correlation-like mechanism, and (4) a trans-saccadic object or blob translation-detection mechanism. We consider it probable that the adaptation observed by Deubel (1991) and Ditterich and colleagues (1999) was driven by the intrasaccadic detection of a low-level motion signal. However, we cannot conclude that this is the case, as we chose to focus on corrective saccade behavior rather than testing adaptation. We chose to limit our study in this manner because we did not want the repetitive quality of the displacement required for adaptation to act as an additional cue to the detection of such displacement. We thus cannot say with certainty that any of the first 3 mechanisms listed above can be utilized in adaptation, though we do consider it likely. That said, at the time of preparation of this document, along with our colleague Laurent Madelain, we are near completing a manuscript based on experiments which demonstrate that adaptation can occur via high-level trans-saccadic detection of whole-field displacement.

Of the data presented in chapter 2, it is most puzzling that performance was above chance in detecting the trans-saccadic displacement of unstructured stimuli (such as the phase-scrambled stimulus used in experiment I or the 100% degraded target of experiment II), especially when one considers that only one subject displayed such performance reliably. It is doubtful that humans possess the capacity to compute something akin to a low-level cross-correlation, as would presumably be required for such performance. Instead, we conjecture that there may have been occasions on which the stimuli contained a “blob” or some other such feature by chance, which facilitated the displacement detection in those trials, and skewed performance. However, further work specifically exploring the conditions and stimulus parameters that allow for this sort of detection are needed before concluding that the cross-correlation computation is truly beyond human abilities. Likewise, we cannot exclude the possibility that the adaptation seen by Deubel (1991) and Ditterich and colleagues (1999) was facilitated by low-level cross-correlation, as they suggest. We do propose, however, that the true test of this notion requires testing the possibility of inducing adaptation using stimuli which require such a mechanism (unstructured, and extinguished intrasaccadically).

Most impressive amongst the findings of chapter 1 is our documentation of the apparent ability to integrate information across the primary saccade at the service of driving correctives. This ability depended strongly on image content, it was present only in the absence of an explicit target, yet required the presence of features. The possible ramifications of this finding are profound: not only is this the first documented case (of which we are aware) of true trans-saccadic integration, it is also a mechanism that might be crucial for the operation of saccade adaptation in the natural environment. As mentioned in the background section of the introduction, in everyday life, it is somewhat difficult to imagine how a position error signal might be calculated at the completion of a saccade that does not expressly target an object. If this feature-sensitive ability we have documented is the manifestation of a generalized prediction error detection mechanism, it would allow for more than just object-directed saccades to be used in the adaptive maintenance of saccadic accuracy, since the error associated with any saccade would be easily computed as long as the (pre- and post-saccadic) retinal image contained features such as contours or shapes. Future work will be required to determine whether, indeed, adaptive changes can occur via this mechanism, and more generally, to determine precisely what stimulus properties allow its operation.

The results of chapter 2 demonstrate for the first time (that we are aware) that target identity can be used to detect saccadic error. Meanwhile (as discussed in the background section of the introduction), it is understood that target position can be used to detect error, in the sense that bottom-up stimulus information is the likely mode by which errors are detected with lone targets. Our assumption is that when post-saccadic target behavior is predictable, the detection of error relies on position information, rather than relying on target identity in all cases, which seems an unnecessary expenditure of computational resources. What is most surprising about our documentation of target-identity-based error detection is its apparent similarity to bottom-up or position-based detection. As mentioned (in several places), we found that when we expressed adaptation magnitude as a percentage of the average ISS used in each adaptation condition, these percentages were indistinguishable (separately for gain-increase and decrease). This suggests that the errors stemming from ISSs detected via target-identity or via target-position induce adaptation with equivalent strength. This is somewhat surprising, because one might assume that target-identity contingent detection would be slightly more time consuming, and hence the error would be delayed relative to detection via target-position (Shafer et al., 2000). Unfortunately, however we cannot conclude with absolute certainty that there truly were two distinct modes of error detection used in the experiments of chapter 2. As we point out in that chapter's discussion section, it is possible that target identity was the reigning mode of error detection in all of chapter 2's conditions. The best test would be to conduct control sessions with identical patterns of ISSs but a lone target (as were performed in Appendix A). This would ensure that error detection does not rely on target-identity. If these adaptation sessions yielded similar results to those with multiple stimuli, we could then conclude that target-identity and target-position based detection are on equal footing, in terms of their ability to induce adaptation. If they indeed are the same, as we propose, one particularly interesting question (given the findings of Kaku et al., 2009, and Soetedjo et al., 2009) is, what differences or similarities in SC activity might one find as these errors are detected?

Finally, the results of chapter 3, documenting adaptation with small ISSs, suggests that the detection of error by a purely bottom-up mechanism does not depend on amplitude / landing-position variability, or undershoot-bias. This proposition is consistent with the findings of Collins and colleagues (2009), who found that the detection of intrasaccadic target displacement does not depend on landing position, and somewhat inconsistent with the findings of Wong and Shelhamer (2011) who found that adaptation depended partially on “inherent undershoot” (undershoot-bias). Our result is only partly incompatible with that of Wong and Shelhamer (2011) in that they used a constant-error (variable-ISS) paradigm to fix post-saccadic retinal error across subjects, while we used a fixed ISS, resulting in a variable retinal error (in each session). The difference between our own study’s findings and theirs may thus simply be a methodological one. Our view of the fundamental limitation on bottom-up saccadic error detection is, again, outlined in Figure 6.1. As we have discussed, we contend that variability in the internal estimate of saccade amplitude sets the lower limit on the ability to detect errors resulting from ISSs. Furthermore, it is important to recognize that this variability is not the same as variability in the command signal (which would be common to both actual amplitude / landing-position variability and variability in the internal estimate) rather this is noise introduced in the process of converting an efference-copy signal into an estimate of saccade amplitude, and is not a component of the variability in actual landing position:

$$Amp_{actual} = Amp_{desired} + \sigma_{command} + \sigma_{execution}$$

Variability in actual amplitude / landing-position would include variability in the command signal, and an additional component of execution noise unknown to the internal estimate.

$$Amp_{estimated} = Amp_{desired} + \sigma_{command} + \sigma_{estimation}$$

However, as mentioned in the discussion of chapter 3, it is possible that differences in the magnitude of adaptation are due not to differences in the ability to detect error (error sensitivity) but are rather due to differences in the way detected errors are used to induce adaptation. We propose that the most direct way to test whether amplitude-estimate-variability translates into sensitivity to error in adaptation would be a paired double-step and saccade adaptation experiment. In double-step saccade experiments, individuals show differences in their ability to compensate in their second saccade for amplitude variability in the first saccade (Bock et al. 1995; Ditterich et al. 1998; Joiner et al. 2010; Munuera et al. 2009; Sommer and Wurtz 2004; Sparks and Mays 1983). It stands to reason that this ability is related to the quality of their internal amplitude estimate. The proposed experiment then, would use the double-step session to quantify compensation ability, and the experimenter would then determine whether this value was predictive of any gain changes observed in the adaptation session.

How are errors used?

The results of chapter 2 suggest that errors that vary from trial-to-trial are used in the same manner as more consistent errors. This is (again) due to the finding that when expressed as a percentage of the average intra-saccadic step, adaptation magnitude was found to be indistinguishable between conditions using more and less predictable ISSs. This finding is somewhat in conflict with the findings of Havermann and Lappe (2010), who found that variable errors were less effective for adaptation, but in keeping with those of Srimal and colleagues (2008) who found that randomly varying errors are used in the same way as consistent ones. The explanation for the difference between our findings and those of Havermann and Lappe may again lie in the fact that they used a constant-error, variable-ISS paradigm. Indeed, their no-added-noise condition already had considerable variability in ISS in order to keep error constant. Furthermore, in a typical 0° added-noise gain decrease session (constant-error of -3°, primary steps of 15°), ISSs varied (roughly) between -7.5° and -12.5° (Havermann and Lappe, 2010, their Figure 2). This is a dramatic sort of variability in post-saccadic target, and bears little resemblance to our paradigm. Meanwhile, Srimal and colleagues (2008) used a state-space model to fit adaptation both in a condition with variable ISSs (between +1° and -2°), and in one with more consistent ISSs (-2°, primary steps in both conditions were 16°). Consistent with our own findings, they observed that the rate of learning in response to backwards ISSs was indistinguishable between the two conditions. In summary, it seems reasonable to suggest that ISS-induced errors might indeed be used in a similar fashion regardless of variability, within a reasonable range.

Another intriguing possibility regarding the usage of errors was raised by the data presented in chapter 2: that onwards errors are less effective than backwards errors for modifying saccade gain because they are used primarily to modify corrective saccade behavior. We found that gain-increase adaptation was less pronounced than gain-decrease adaptation (as have many others), however we also found that gain-increase adaptation was substantially more likely to result in prediction-based corrective saccades than was gain-decrease (Figure 2.6). We argue in the discussion section of chapter 2 that this reflects an interaction between target predictability and a proposed preexisting proclivity for executing onwards corrective saccades, resulting from the tendency to undershoot (Becker, 1989). In support of this, Wong and Shelhamer (2011) found that (in abrupt but not gradual) gain-increase adaptation, the incidence of onwards correctives increased, which they too suggested was a hallmark of predictive behavior. Our suggestion is not entirely dissimilar from the finding of (Optican and Miles, 1985), amongst others, that post-saccadic drift can be selectively adapted by the post-saccadic presentation of a motion signal. The similarity lies in the suggestion that different components of saccadic behavior can be differentially modified. Our suggestion then, is that looking at the primary saccade alone in detailing the effects of gain-increase ISS-adaptation may mask the potentially strong effects of onwards errors in modifying saccadic behavior, as we have observed. Instead, looking at the sequence of saccades achieved in foveating a target may be more useful. Further work, however, will be required to validate this speculative line of thinking.

While the results of chapter 2 suggested that spatial variability has little impact on the way that errors are used in adaptation, the results of chapter 4 suggest that adaptation may also be resilient to temporal variability in error. We found that intermittent feedback only minimally impacts adaptation magnitude and rate, suggesting that the adaptive system is fairly robust in its ability to contend with such temporal variability. Fascinatingly, adaptation with feedback being delivered on 50% of trials was just as strong as with feedback on 100% of trials (Figure 4.2). Instead, in keeping with other studies of motor learning

(Huang and Shadmehr, 2009; Joiner and Smith, 2008), we found that the duration of the learning period (number of trials in the adapt-phase) modulated the learning and retention rates (Figure B.1). That learning-duration modulated rates irrespective of feedback intermittency (in Chapter 4, Experiment I), suggests that the adaptive system was able to determine that the learning period was ongoing irrespective of feedback intermittency. This is in keeping both with the suggestion that withholding feedback does not amount to withholding error (consistent with an SPE as the error signal), and with the notion that the brain is sensitive to the statistics of the learning environment in general (Behrens et al., 2007). Indeed, one component that we were not able to test in this series of experiments is how the adaptive system estimated the duration of the learning-period. That is, if learning-period duration determines rates, what determines learning-period duration? As mentioned, some authors have suggested this might occur implicitly (Joiner and Smith, 2008). However, it is also possible that the transition between experiencing errors that are more within an expected range (during the baseline) and those outside of the expected range (once the adapt-phase has begun) alerts the adaptive system that a learning-period has been entered. As proposed by Behrens and colleagues (2007), the brain may continually estimate the "volatility" of the learning environment, which would be high both at the onset and offset of a conventional saccade-adaptation experiment, and could be thus used to determine the duration of a learning event. One way to test this idea would be to compare the results of Experiment I as they stand with a new set of experiments in which the baseline-phase had the same intermittent-feedback as the adapt-phase (but without the ISSs); this would have the effect of obscuring the transition from the baseline to the adapt-phase, would it impact adaptation dynamics as well?

Consistent with other motor learning studies, the results of chapter 3 suggest that small errors may be used differently from larger errors in saccade adaptation (Robinson et al. 2003) (Figure 3.2A). We found that adaptation with 0.25° ISSs was complete, unlike most other examples of saccade adaptation (Hopp and Fuchs, 2004). Interestingly, this was also the only condition in which we observed complete recovery as well (Figure 3.2D). However, that adaptation was incomplete with the "larger" ISSs used in this study is likely a feature of the brief adaptation period we used (a mere 150 trials). That is, there is no principled reason to suspect that 0.25° is smaller at a fundamental level than 0.5°, and yet adaptation was incomplete even with this error size. The principle explanation that is usually given to account for the increased efficacy of small errors is that they may be more similar to the scale of error encountered by the system in the homeostatic regulation of motor accuracy in response to small changes in muscle strength or motor-pool innervation (Wei and Kording, 2007). Further work on this subject will be required to understand whether there is some basic size-scale that differentiates “small” errors from “large.”

Finally, in chapter 5, we showed that the visual appearance of a target could be used to define a context for adaptation. We suggested that the errors experienced with a given target served to establish and maintain a contextual state for adaptation, as we found that when competing errors were not paired with one of the two targets, adaptation with the other target transferred completely. We further proposed that any sensory cue might ultimately result in the formation of a motor context if the cue was predictive of the errors experienced in that context. However, it was unclear how this might occur, since others have found visual cues to be ineffective in establishing contexts for saccade adaptation (Deubel, 1991). However, more recently, (Beaton et al., 2010)Beaton and colleagues (2010) found that visual, auditory, and more ambient (the room an adaptation session was conducted in) contexts could all be used to enhance context-specific adaptation effects. This result demonstrates that the system responsible for establishing and maintaining contexts for action has access to a wide range of sensory data indeed.

Conclusion

In conclusion, we have displayed evidence that the saccade adaptation system makes use of trans-saccadic integration, intrasaccadic velocity detection, target-identity-specific error calculation, and a highly sensitive prediction of post-saccadic retinal error all at the service of adjusting of saccadic gain. The two most important contributions of this work have been: (1) a clarification of the limited relationship between ISS adaptation in the lab and adaptive maintenance of saccade gain in the natural environment, and (2) an expansion of what is understood about the relationship between saccade adaptation and other forms of motor learning, and learning in general. Though no direct comparison was made of adaptation in the lab versus the natural environment, we outlined many of the challenges faced by the adaptive system in the natural environment, and have (in some cases) elucidated the ways they are dealt with. We also probed some basic properties of ISS adaptation that have not before been subject to systematic study, finding that in many ways, saccade adaptation shares properties with other forms of motor learning. It is our hope that our efforts to explore saccade adaptation and saccadic behavior under generally more demanding conditions has demonstrated the power of this approach. Further work in this vein will yield a greater understanding of the ways in which our behaviors are molded by our experiences in everyday life, and thus, of the ways in which the brain itself is shaped by the rich diversity of those experiences in general.

Appendix A: Saccade adaptation is unhampered by distractors[8]

Introduction

Saccades are the rapid movements of the eyes used to look around the environment, to read and to react to the sudden movement or appearance of an object. Because they are so rapid - in humans a 10° saccade lasts about 40 to 50 ms (Becker, 1989) - no useful feedback from the visual system, which requires about 40 to 50 ms to reach the superior colliculus (Wurtz & Goldberg, 1972) is available to guide the movement on-line. Moreover, vision is largely impaired during a saccade (Brooks & Fuchs, 1975). Thus, the behavior is often described as ballistic or open-loop movements in which the trajectory is thought to be programmed prior to the movement. Such open-loop behaviors are kept accurate by parametric feedback, meaning that, although the on-going behavior is not adjusted by feedback, the parameters of the control system governing the behavior can be adjusted by the consequences of the behavior. In the case of saccades, much evidence shows that if the eye consistently lands beyond targets at a particular distance in a particular direction, the oculomotor system reduces the amplitude of saccade elicited by a target in that location. Such repair in sensorimotor control has been observed in humans with diseases inducing weakness in extraocular muscles (Kommerel et al 1976, Optican et al 1985), as well as in monkeys (Optican & Robinson, 1980). This plasticity, called saccadic adaptation, emerges through motor learning that continually maintains accuracy in response to new sensorimotor transformations.

In the laboratory, saccadic adaptation has been demonstrated by experiments in which the target is surreptitiously moved while the eye is in flight and hence its vision is impaired (McLauglin, 1967). This double-step paradigm is known to control both the extent and direction of change in amplitude: if the target is repeatedly moved closer to the starting point, so that the eye lands beyond the target, the future saccades become smaller; if the target is moved further from the starting point, so that the eye lands short of the target, the saccades become larger. Saccade adaptation is usually viewed as a servo system, in which an error signal is reduced by a change in the motor output: after the saccade, the mismatch between the position of gaze and the visual target induces a change in the system's parameter (Noto & Robinson, 2001; Wallman & Fuchs, 1998). Moreover, it has been shown that adaptation depends on the

---

[8] This manuscript appears with permission from the Journal of Vision, Copyright © 2010 by the Association for Research in Vision and Ophthalmology (Madelain et al., 2010).

physical and temporal properties of this visual error signal (Bahcall & Kowler, 2000; Shafer et al 2000). One could postulate that, instead, it is the corrective saccade necessary to foveate the target after a dysmetric initial saccade that provides the signal to adjust the amplitude. However, it has been demonstrated that adaptation occurs in the absence of corrective saccades (Bahcall & Kowler, 2000; Noto & Robinson, 2001; Wallman & Fuchs, 1998).

The weight of available evidence indicates that the visual error signal existing when the primary saccade lands drives adaptation (see Hopp & Fuchs, 2004 for a review). To our knowledge, saccadic adaptation has been experimentally studied exclusively by shifting either a single target or the whole visual scene. This creates a single, non ambiguous error signal that depends on both the saccade amplitude and the extent of the experimental displacement. However the ability to selectively induce adaptation when more than one error signal is available is unknown. To address this question we designed a novel paradigm based on the conventional double step: we introduced a new object – a distractor – appearing at the time of the adaptation step, i.e. during the saccade. By controlling the position of both the target and distractor this allowed us to induce two independent position error signals. The panel A of figure A.1 plots a schematic view of this paradigm. In one case (double-step plus distractor, left panel), the target is shifted during the saccade and a distractor appears at the position previously occupied by the target. The saccade would then be too long with respect to the target and too short with respect to the distractor. In the other case (single-step plus distractor, right panel), the target is not stepped back but the distractor appears left with respect to the target. The saccade would then be too short with respect to the target. In each case, two position errors would be present that might drive a change in gain, one being related to the target, the other to the distractor. One could postulate three different outcomes from this situation. First, the gain might remain unchanged either because the oculomotor system uses an average of the two position errors or because it alternatively and randomly selects one or the other error signal. Second, the gain might systematically decrease, regardless of which paradigm is used, because of the tendency of normal saccades to be hypometric: the saccadic system might be preferentially driven by the position error encouraging a gain decrease and ignore the one encouraging an increase in saccade amplitude. Third, the gain might decrease in the double-step plus distractor paradigm and remained unchanged in the single-step plus distractor paradigm, implying that the saccadic system responds to the position error of the target and ignores the one related to the distractor (Figure A.1A). Our results support the third hypothesis: the target, not the distractor, causes changes in saccade amplitude.

Methods

General Methods

The three experiments tested whether adaptation could occur in the presence of two error signals, one originating from the target, and the other from a distractor. Five subjects performed the experiments. Each trial started with a fixation period of 750–1,250 ms (figure A.1B) during which a target was displayed to the left of the center of the screen against a gray background (luminance 40 cd/m²). The initial target position was randomized (either -7, -6, -5 or -4° with respect to the center of the screen). After this fixation period, the target stimulus, stepped by 10 or 12° (each amplitude occurred in 50% of the trials in a pseudo random order) to the right of the fixation cross, and remained visible for 1,500 ms. Subjects were instructed to track the target.

In the double-step plus distractor trials the target stepped back (experiment 1) or forward (experiment 2) at saccade onset by 20% of the first step (2° and 2.4° for a 10 and 12° initial step, respectively) and a distractor appeared at saccade onset, at the position previously occupied by the target (Figure 1B and C top row). In the single-step plus distractor sessions the target did not step back at saccade onset but the distractor appeared either left (experiment 1) or right (experiment 2) to the target (2° and 2.4° for a 10 and 12° initial target step, respectively) (Figure 1B and C bottom row). Conventional double-step paradigms, in which the target is stepped during the saccade, were also carried using the same spatial and temporal parameters to provide a control situation.

In order to eliminate the possibility that the visual properties of the stimuli induced changes in gain each target was used in 50% of the trials in a pseudo random order. The target was either a 0.5 x 0.5° green filled circle or a yellow filled square encompassing 0.5° in diameter. If the target was the green circle, the distractor was the yellow square, and, conversely, in trials in which the target was the yellow square, the green circle was used as the distractor. Therefore, the target was the object visible during the fixation period and the distractor the object appearing at saccade onset.

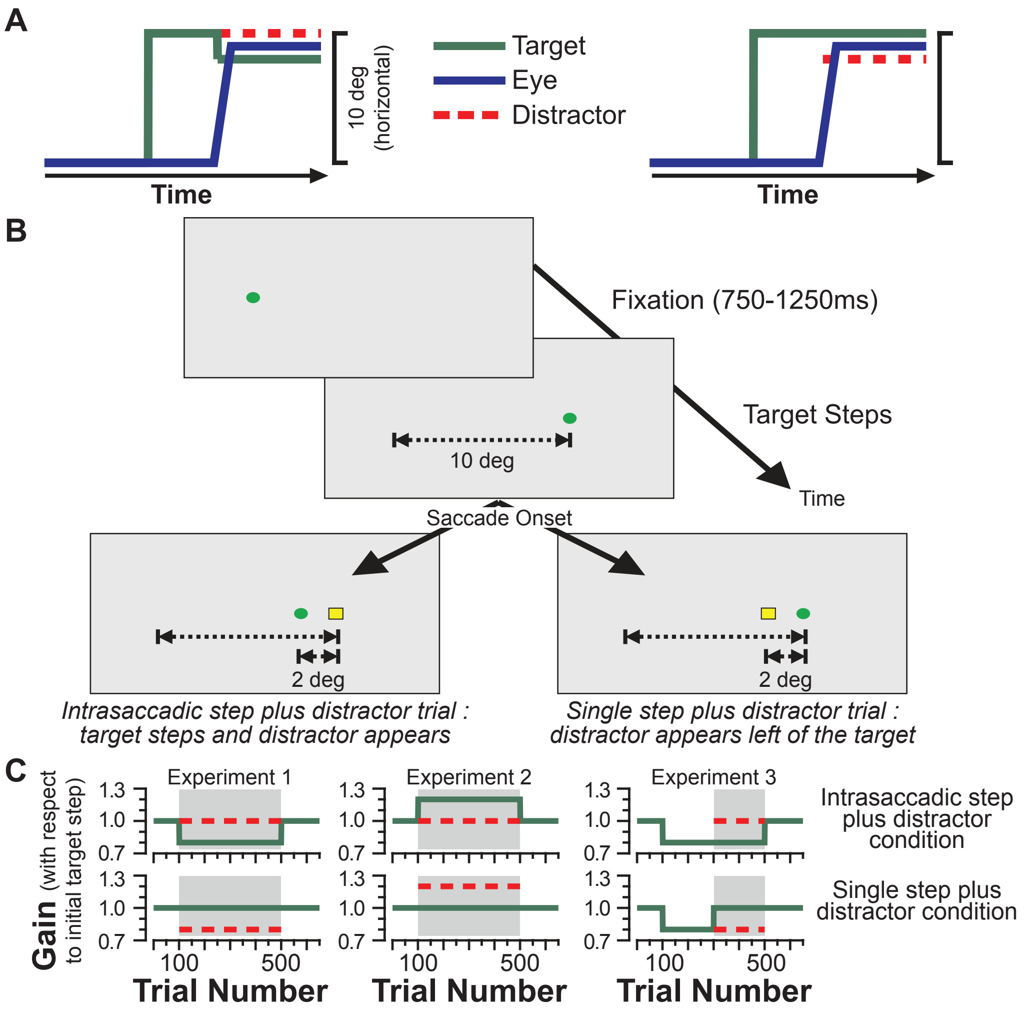

Figure A.1 Methods. **A**. Schematic diagram of the position errors following a 9 deg saccade for a 10 deg initial target step. In the target adaptation paradigm (left), the saccade (blue) would be hypermetric with respect to the target (green solid line) but hypometric with respect to the distractor (red dashed line). In the distractor adaptation paradigm (right), the reverse would be true. Vertical arrows mark saccade onset. **B**. Schematic diagram of the temporal sequence of trials used in target adaptation or distractor adaptation of Experiment 1. After a fixation period (750-1250 ms) the target (green square) steps to the right. As soon as a saccade onset is detected, two stimulus configurations are possible, depending on the paradigm used. In the double-step plus distractor paradigm (left), the target (green disc) steps back by 20% of the initial target step while a distractor (yellow square, color has been changed for legibility) appears at the initial target location. In the single-step plus distractor paradigm (right), the target position remains unchanged while a distractor appears left of the target (see text). **C**. Time course of trials in the three experiments. All sessions began by 100 regular single step trials. In the double-step plus distractor paradigm (upper graphs) or single-step plus distractor paradigm (lower graphs) of experiment one and two, subjects experienced 400 adaptation trials (from trial 101 to 500) followed by 150 regular single step trials. In experiment 3, subjects experienced 200 conventional double steps trials (from trial 101 to 300) followed by 200 experimental trials (from trial 301 to 500), followed by 150 single step trials. Vertical gray lines mark the transition between stages.

Figure A.1 Methods

Experiment 1: Decreased gain adaptation

The first experiment tested whether backward adaptation could occur in the presence of two error signals, one originating from the target, and the other from a distractor. Each subject performed three experimental paradigms: a conventional double-step paradigm, a double-step plus distractor paradigm, and a single step plus distractor paradigm.

For the conventional double-step paradigm a session was divided into three consecutive stages. We first collected 100 regular, pre-adaptation trials in which the target was not stepped back at saccade onset. We then collected 400 experimental trials in which the target stepped back at saccade. In a remaining 150 recovery trials, the target did not step back after its original step.

For the double-step plus distractor paradigm, subjects first performed 100 regular, pre-adaptation trials in which the target was not stepped back at saccade onset. We then collected 400 double step plus distractor trials in which the target stepped back at saccade onset and the distractor appeared at the same time, at the position previously occupied by the target. In a remaining 150 recovery trials, the target did not step back after its original step (Figure A.1B left and C top left).

For the single step plus distractor paradigm, the session structure of the experiment was identical as in the other paradigms. In the experimental trials the target did not step back at saccade onset but the distractor appeared left of the target (Figure A.1B right and C bottom left).

Experiment 2: Increased gain adaptation

The second experiment tested whether forward adaptation could by an error signal originating from a distractor. The design was identical to the one used in experiment 1 except that the in the experimental trials the target (or the distractor) did reappear at saccade onset 2°ree (or 2.4° for a 12 deg initial step) right of the initial position in order to induce an increase in gain (Figure A.1C, middle panels).

Experiment 3: Recovery from decreased gain adaptation

In a third series of experiments we asked whether the presence of the distractor would affect the recovery following backward adaptation (figure A.1C, right). After 100 pre-adaptation trials, adaptation was induced by performing 200 conventional double-step backward trials. After this adaptation period, we collected 200 experimental trials that were either single-step plus distractor trials similar to the one used in experiment 1 (Figure A.1B, right panel) or double-step plus distractor trials (Figure A.1B, left panel). In a remaining 150 recovery trials, the distractor did not appear at saccade onset.

Participants

The participants were 2 naïve subjects, plus three of the authors (subjects S3, S4 and S5). All had normal vision. Participants who were naïve with respect to the goal of the experiments had previous experience in oculomotor experiments. The order of the eight experimental sessions (three in the first experiment, three in the second experiment and two in the third experiment) was randomize across subjects. Subjects usually performed one or two daily sessions.

Written consent was obtained from all subjects and the experimental protocol was approved by the Institutional Review Board of the City College of New York.

Apparatus

Subjects in a darkened room viewed stimuli at a distance of 57 cm while on a bite board to minimize head movements. Stimuli were generated on a computer running VisionSometimesWorks (Vision Research Graphics, Durham, NH) and displayed on a 21-in. CRT display with a vertical refresh rate of 100 Hz. Pupil position was digitized at 240 Hz, using an infrared video eye-tracking system (ISCAN, Woburn, MA), controlled by a computer using the LabView software package (National Instruments). Trigger signals between the visual display computer and the LabView computer allowed us to synchronize data collection to stimulus presentation with 4-ms resolution. Immediately preceding each experimental session, a 50-point horizontal calibration was carried out by having the subject fixate a 0.3° target 10 times at each of five randomized screen locations and the experimenter pressed a key to acquire a 50-ms average of pupil position. Locations were randomized to ensure that there was always a saccade between measurements. An on-line least-squares fit to these measurements was used, and the resulting gain and offset were saved for off-line analysis.

For on-line saccade detection, a real-time algorithm used a velocity criterion to identify the first saccade at least 1000 ms after the target onset. The change in visual display was triggered as soon as the velocity signal exceeded a fixed threshold.

Data analysis

All analyses were performed using MATLAB (The MathWorks, Natick, MA). During off-line analysis, saccades were detected automatically using a velocity threshold of 30°/s and a minimum latency criterion of 100 ms to exclude anticipatory movements. The start and end of each saccade immediately following the target step were confirmed and, if necessary, corrected by the experimenter. Approximately 5% of all saccades were excluded from further analysis.

To show the trends of the adaptation and recovery within individual experiments, we smoothed the records of saccadic gain vs. trial number by means of a Lowess iterative non-linear regression (width=50 trials). To describe the changes in amplitude we computed the mean gain (and corresponding standard deviation) of the last 50 trials at each stage (pre-adaptation, adaptation, recovery). We then compared the mean gains obtained in the last 50 trials of each stage for each paradigm using individual two-ways ANOVA (three levels in one factor: pre-adaptation, adaptation, recovery, three levels in the other factor: conventional double-step, double-step plus distractor, single-step plus distractor, 50 measurements in each degree). We then computed post-hoc all pair-wise t tests using the Sidak corrections.

Results

Our result show that saccadic gain changes according to the target's position and ignores the distractor: in all experiments saccadic adaptation occurred in the double step plus distractor trials but not in the single step plus distractor trials. Moreover the amount of adaptation obtained in the presence of a distractor was not consistently different from what was obtained in the conventional double steps paradigms indicating that the saccadic system is able to selectively adapt its gain by following the relevant position error.

Experiment 1: Decreased gain adaptation

Figure A.2A plots the saccadic gains for each trial in one subject (S1). For the double-step plus distractor (red dots), the gain remained mostly unchanged in the pre-adaptation trials (trial 1-100). In the last 50 pre-adaptation trials the gain averaged 0.91 (SD 0.03), well within the range of gain reported in the literature. During the adaptation trials, the gain quickly decreased to reach an average of 0.75 (SD 0.05) in the last 50 trials. Finally, the saccadic gain increased back during the recovery trials (mean 0.86, SD

0.04), much like one would expect using a conventional adaptation paradigm. In fact, the pattern of gain was very similar in the conventional double-step condition (blue dots) with averages of 0.93 (SD 0.04), 0.78 (0.06) and 0.87 (0.04) for the pre-adaptation, adaptation and recovery stages respectively. In the single-step plus distractor condition (green dots), the results were dramatically different: the gain remained mostly unchanged during the session, averaging 0.92 (SD 0.05), 0.92 (0.04) and 0.90 (0.06) for the last 50 pre-adaptation, adaptation and recovery trials respectively. The other subjects also showed a consistent decrease in saccade amplitude during the experimental trials followed by an increase during the recovery trials in the conventional double-step and double-step plus distractor paradigms, associated with little (S5, figure 2E) or no change in gain for the single-step plus distractor condition.

To evaluate the reliability of the changes in saccadic gain, we compared the gains obtained in the last 50 trials of the pre-adaptation, adaptation and recovery stages for each paradigm using individual two-ways ANOVA (2,2) and post-hoc all pair-wise t tests using the Sidak corrections. In all subjects we found a significant ($p<0.05$) effect of both stages and paradigms, as well as a significant interaction between the two factors. Figure 2F plots the means and standard deviation as well as the results of the corrected *t* tests comparing the experimental trials in each paradigm. A blue star indicates a significant (corrected $p<0.05$) difference between the conventional double-step and the double-step plus distractor paradigms, while a green star indicates a significant difference between the two paradigms involving a distractor. It can be seen that the gain was systematically and significantly lower during adaptation in the target double-step plus distractor paradigm than in the single-step plus distractor condition (0.78 and 0.93 respectively when averaged across subjects). This was not true when comparing the conventional and double-step plus distractor paradigms (gain: 0.77 and 0.78 respectively when averaged across subjects): the gain was significantly higher in the conventional paradigm for two subjects, but significantly lower in one subject and unchanged in the other two subjects.

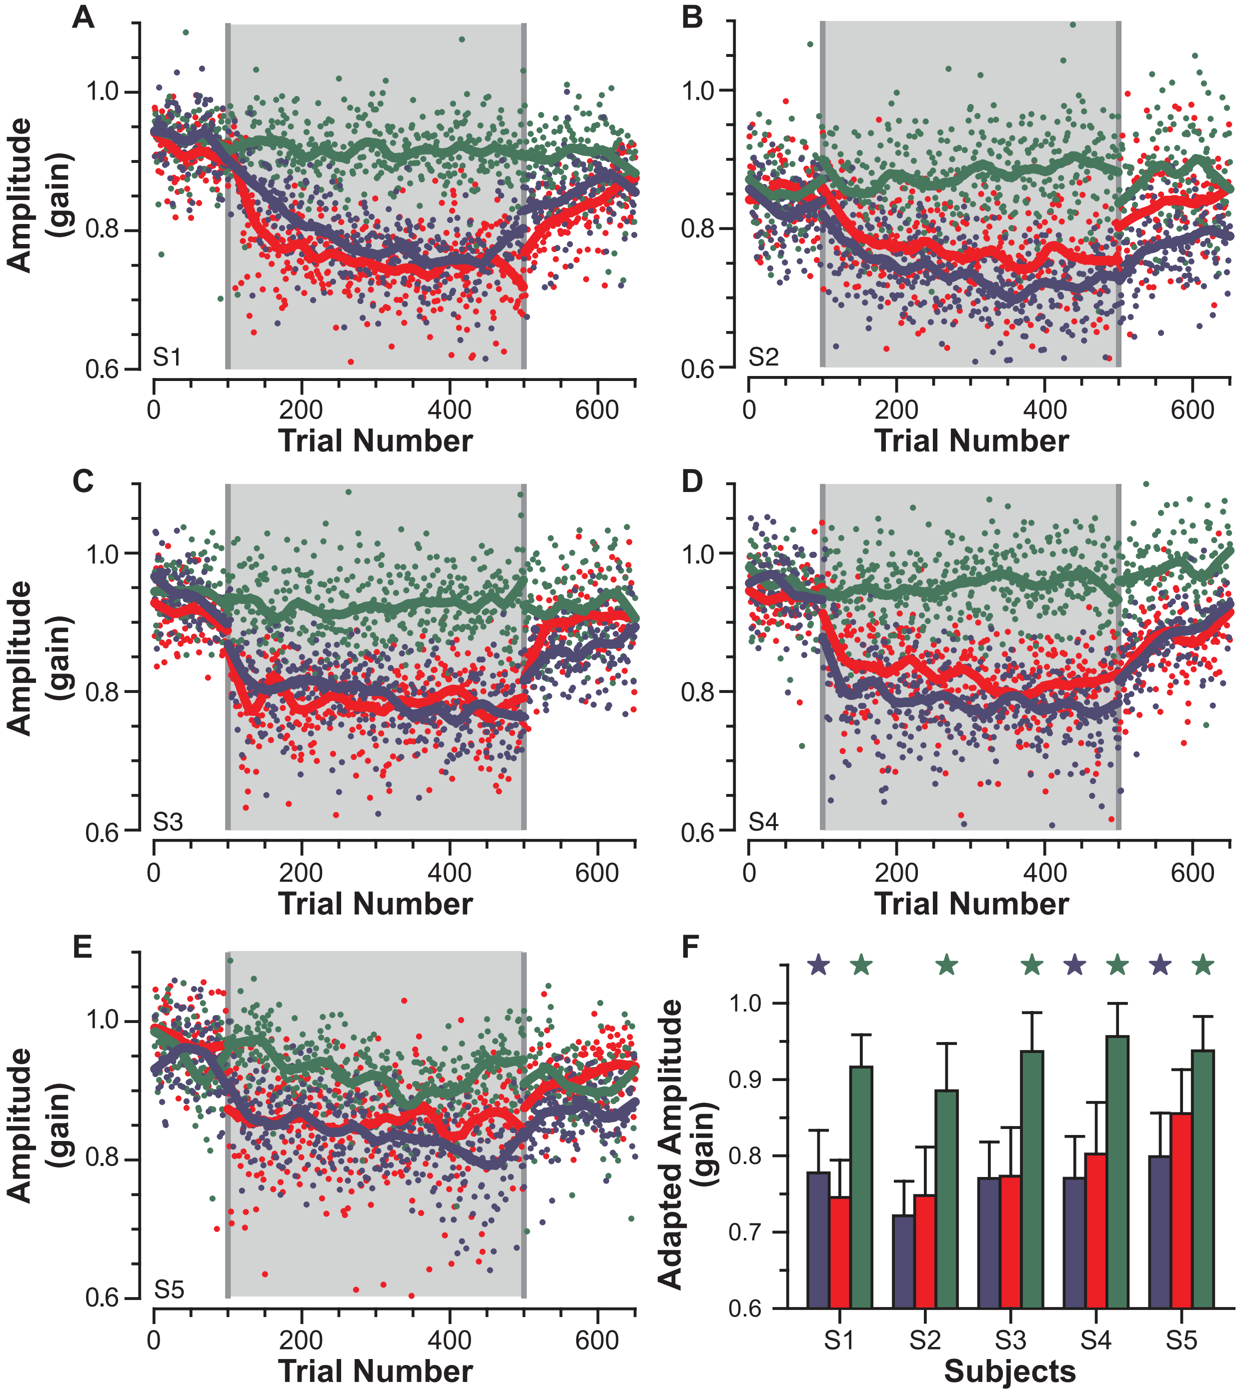

Figure A.2 Gain Decrease Data **A.** Individual saccade amplitudes for all trials in experiment one (subject 1). **Conventional double steps**, **double-step plus distractor**, and **single-step plus distractor** paradigms. Solid lines, corresponding lowess smooths. Vertical gray lines mark the transition between stages. **B**, **C**, **D** and **E**: same as panel A for subjects 2, 3, 4 and 5. **F.** Individual mean gains for the last 50 adaptation trials in the **conventional**, **double-step plus distractor**, and **single-step plus distractor** paradigms. Stars indicate significant differences (pair-wise t-tests $p<0.05$) between mean adapt-phase gains in **conventional** vs. **double-step plus distractor** paradigms (**blue stars**) or **double-step plus distractor** vs. **single-step plus distractor** paradigms (**green stars**).

Figure A.2 Decreased gain adaptation (Experiment 1)

It is interesting to point out that no consistent differences were observed when comparing the baseline trials across paradigms (0.91, 0.91 and 0.92 in the conventional, double-step plus distractor and single step plus distractor paradigms respectively, when averaged across subjects). Finally we looked at the individual differences in the baseline and adaptation trials in each paradigm: we found a systematically significant reduction in gain for the conventional (from 0.91 to 0.77 respectively when averaged across subjects) and double-step plus distractor (from 0.91 to 0.78 respectively when averaged across subjects) paradigms, but no differences for the single-step plus distractor paradigm (0.92 and 0.93 respectively when averaged across subjects). This pattern of results implies that subjects adapted equally well in both the conventional and double-step plus distractor paradigms but did not adapt in the single-step plus distractor condition, revealing the ability of selectively using the target related position error to control changes in saccadic gain.

Experiment 2: Increased gain adaptation

The designs of these experiments were similar to the one used in experiment 1, except that the target end position was set to encourage an increase in saccade amplitude (Figure A.1C). Figure 3A plots the saccadic gains for each trial in one subject (S1). For the double-step plus distractor condition (red dots), the gain remained mostly unchanged in the pre-adaptation trials (mean 0.93, SD 0.04), increased during the adaptation trials (mean 1.03, SD 0.06), and, finally, decreased back during the recovery session (mean 0.94, SD 0.05). The pattern of changes in gain was very similar in the conventional double-step condition (blue dots) with averages of 0.89 (SD 0.06), 1.02 (0.08) and 0.93 (0.06) for the pre-adaptation, adaptation and recovery stages respectively. In the single-step plus distractor condition (blue dots), the gain remained mostly unchanged during the session, averaging 0.94 (SD 0.06), 0.93 (0.08) and 0.93 (0.06) for the last 50 pre-adaptation, adaptation and recovery trials respectively.

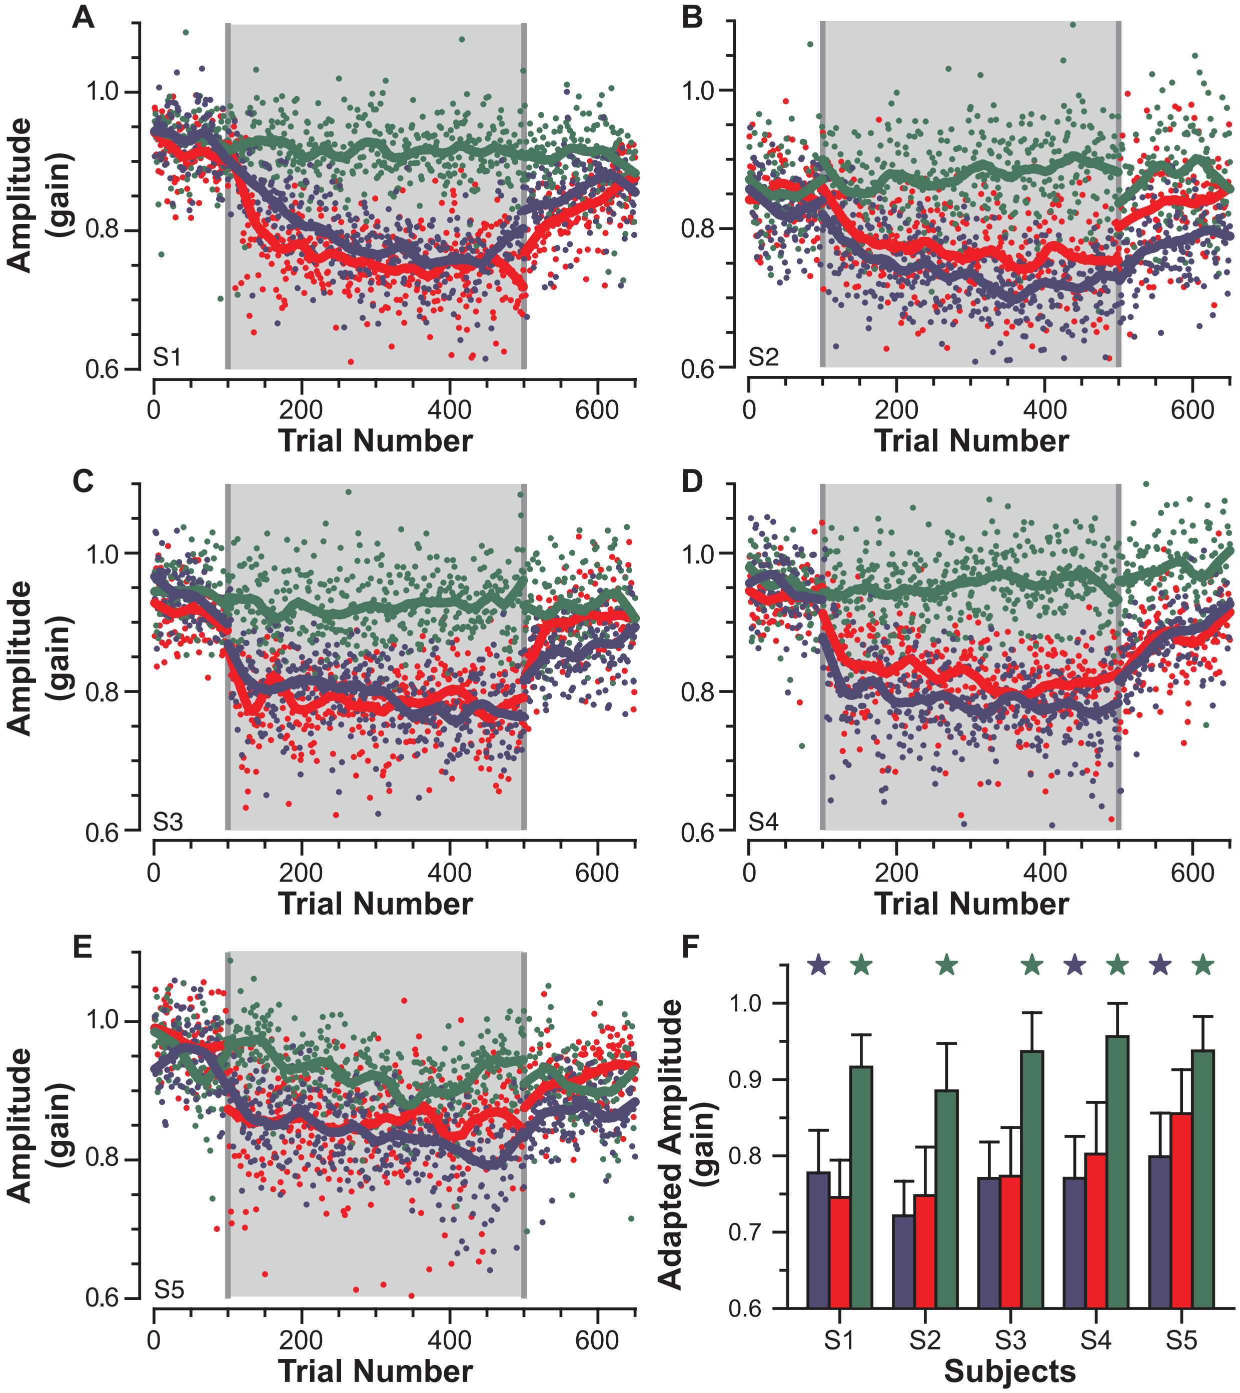

Figure A.2 Gain Decrease Data **A.** Individual saccade amplitudes for all trials in experiment one (subject 1). **Conventional double steps**, **double-step plus distractor**, and **single-step plus distractor** paradigms. Solid lines, corresponding lowess smooths. Vertical gray lines mark the transition between stages. **B**, **C**, **D** and **E**: same as panel A for subjects 2, 3, 4 and 5. **F.** Individual mean gains for the last 50 adaptation trials in the **conventional**, **double-step plus distractor**, and **single-step plus distractor** paradigms. Stars indicate significant differences (pair-wise t-tests $p<0.05$) between mean adapt-phase gains in **conventional** vs. **double-step plus distractor** paradigms (**blue stars**) or **double-step plus distractor** vs. **single-step plus distractor** paradigms (**green stars**).

Figure A.3 Increased gain adaptation (Experiment 2)

Other subjects (except subject 5, Figure A.3E) showed a comparable pattern of results. In all subjects we found a significant (ANOVA (2,2), $p<0.05$) effect of both stages and paradigms, as well as a significant interaction between the two factors, except for subject 5 in whom there was no significant effect of the stages. Figure 3F plots the means and standard deviation as well as the results of the corrected *t* tests comparing the adaptation trials in each paradigm. It can be seen that the gain was systematically and significantly higher during adaptation in the double-step plus distractor paradigm than in the single-step plus distractor condition (0.99 and 0.93 respectively when averaged across subjects). This was not true when comparing the conventional and double-step plus distractor paradigms (gain: 1.01 and 0.99 respectively when averaged across subjects): the gain was not significantly different in the conventional paradigm for four subjects, but significantly higher in subject 5.

Finally we looked at the individual differences in the baseline and experimental trials in each paradigm: we found a significant increase in gain for the conventional (from 0.92 to 1.01 respectively when averaged across subjects) and double-step plus distractor (from 0.92 to 0.99 respectively when averaged across subjects) paradigm, except for subject 5 who did not significantly adapt in the conventional nor in the double-step plus distractor paradigms. We did not find significant differences in the single-step plus distractor paradigm (0.92 and 0.93 respectively when averaged across subjects), except for subject 2 (from 0.84 to 0.88). This pattern of results confirms the results from experiment 1 in that most subjects were able to selectively follow the relevant position error to control changes in saccadic gain.

Experiment 3: Recovery from decreased gain adaptation

The goal of these experiments was to probe the ability to recover from adaptation when an error signal related to a distractor encourages a lower gain while an error signal related to the target should drive a gain close to unity (Figure A.1A, right). Recall that there are now four consecutive stages in each paradigm (figure 1C, right): 100 pre-adaptation trials, 200 conventional double-step trials, 200 recovery trials in which both a target and distractor were present, followed by 150 regular trials. In the single-step plus distractor paradigm the saccades, after 200 conventional adaptation trials, should be hypometric with respect to the target, but not with respect to the distractor. In the double-step plus distractor paradigm, the reverse was true. Given the results from experiments 1 and 2 one might expect to observe an increase in

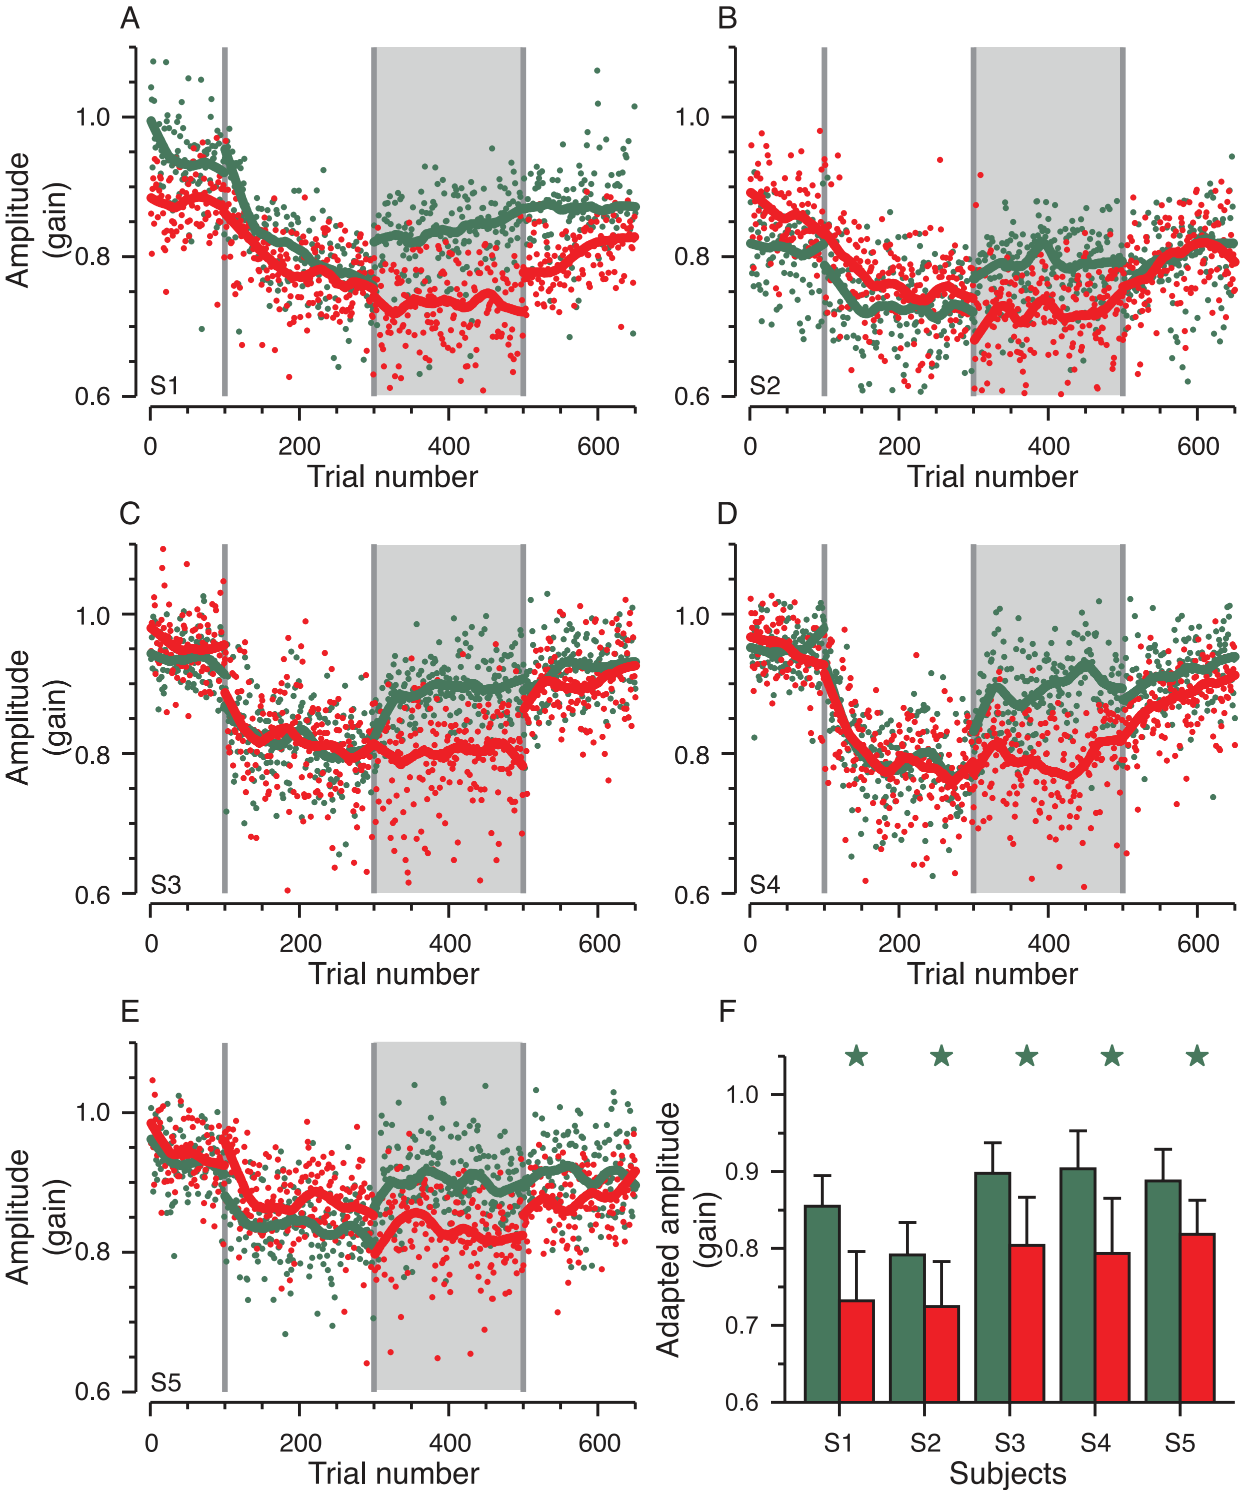

Figure A.4 A, **B**, **C, D** and **E**: Individual saccade amplitudes for all trials in experiment three (subject 1, 2, 3, 4 and 5 respectively). Symbols and lines as in Figure 2, except for the conventional paradigm. F: Individual mean gains for the last 50 recovery trials. Symbols and lines as in Figure A.2 (except for the conventional adaptation paradigm).

Figure A.4 Recovery from decreased gain adaptation (Experiment 3) gain, i.e. a recovery, in the single-step plus distractor condition but not in the double-step plus distractor paradigm.

Figure A.4A plots the saccadic gains for each trial in one subject (S1). For the double-step plus distractor paradigm (red dots), the gain remained mostly unchanged in the pre-adaptation trials (mean 0.88, SD 0.04), decreased during the adaptation trials (mean 0.76, SD 0.07), remained low during the experimental trials (mean 0.73, SD 0.06), and increased during the last 150 regular trials (mean 0.82, SD 0.04). For the single-step plus distractor condition (green dots) the general pattern was comparable, except that the gain increased during the experimental trials: average gain values for each stage were 0.93 (SD 0.06), 0.76 (0.04), 0.85 (0.07) and 0.88 (0.07), respectively.

All other subject showed a similar pattern of results. In all subjects we found a significant (ANOVA (3,1), $p<0.05$) effect of both the stages and paradigms, as well as a significant interaction between the two factors, except for subject 2 in whom there was no significant effect of the paradigms. Figure A.4F plots the means and standard deviation as well as the results of the corrected *t* tests comparing the recovery trials in each paradigm. It can be seen that the gain was systematically and significantly higher in the single-step plus distractor trials than in the double-step plus distractor trials (0.87 and 0.77 respectively when averaged across subjects).

To probe the gain adaptation we looked at the individual differences in the baseline and adaptation trials in each paradigm: we found a significant decrease in gain in all subjects (from 0.91 to 0.78 respectively when averaged across subjects and paradigms). When comparing the adaptation trials and the experimental trials we found a significant increase in gain in all subjects in the single-step plus distractor paradigm (from 0.78 to 0.87 respectively when averaged across subjects) but not in the double-step plus distractor paradigm in which the gain remained either unchanged (S1, S3 and S4) or significantly decreased (S2 and S5), from 0.79 to 0.77 respectively when averaged across subjects. Finally we looked at the individual differences in the recovery and regular trials in each paradigm. The gain did significantly increase in two subjects (S3 and S5) but not in the other three subjects in the single-step plus distractor paradigm (from 0.87 to 0.89 respectively when averaged across subjects) indicating that recovery was almost complete following the experimental trials. In the double step plus distractor condition, we found a systematic and significant increase in gain in all subjects (from 0.77 to 0.87

respectively when averaged across subjects), revealing that saccadic gain recovered only during the last 150 regular trials.

Discussion

Our results show that, in the presence of two conflicting retinal errors, one originating from the target and the other from a distractor, the saccadic system is able to selectively adapt its gain by following the relevant error signal while ignoring the other one.

In the first experiment, designed to encourage a 20% decrease in gain during the adaptation trials, we observed the expected adaptation when using a conventional double steps paradigm. When adding a distractor appearing at the original step position, while the target stepped backward at saccade onset (double-step plus distractor paradigm, Figure A.1), we found that subjects essentially ignored the distractor and adapted to an extent comparable to the adaptation driven by the conventional paradigm. In the single-step plus distractor (Figure A.1) adaptation did not occur. The degree of adaptation we observed in double-step plus distractor paradigm (14% decrease on average) was comparable to the one obtained with conventional double-step adaptation (15% decrease on average), with comparable final gains (0.78 and 0.77 respectively). By comparison, the final gain in the single-step plus distractor trials was 0.93, i.e. a +1% increase when compared to the pre-adaptation stage. It has been observed that the amount of adaptation is usually less than the amount demanded by the size of the adaptation step – in our case -14% and -15% for a -20% step (Miller et al, 1981, Semmlow et al, 1987; Robinson et al, 2003). Subjects were variable both in their pre-adaptation and adapted gains (compare Figures A.2A, B, C, D and E). Such inter-subject and inter-session variability has previously been observed in normal adaptation (Erkelens and Hulleman, 1993).

In the second experiment, designed to induce a 20% increase in saccade amplitude (figure 1C, middle panels), the general pattern of results was comparable to the one obtained in the first experiment: we found a significant gain increase in the conventional double-step and double-step plus distractor paradigms (except for subject 5), but not in the single-step plus distractor paradigm (except for subject 2). However, the extent of adaptation in the double-step plus distractor paradigm was only +8%, almost half of what was observed in the gain decrease experiment. This asymmetry is also a prominent feature of conventional saccadic adaptation; in our control experiment the average gain changed by +11% after forward-step trials, but –15% after back-step trials, although the target stepped by 20% for 400 trials in both cases. The origin of these differences in forward and backward adaptation is unknown but it may be related to saccades being naturally hypometric (Henson, 1978; Straube, Fuchs, Usher, and Robinson, 1997), perhaps because of a tonic tendency for the gain to decrease. It has also been suggested that distinct mechanisms are involved: a reduction of motoric gain in the case of backward adaptation and a remapping of the desired goal for saccades in the case of forward adaptation (Semmlow, Gauthier, and Vercher, 1989; Panouillères et al, JOV, 2008). Interestingly, it has been proposed that recovery following an amplitude-decrease adaptation was a forward adaptation (Deubel et al, 1986). This would explain why, in the third experiment, the gain following the target recovery trials averaged 0.87, a value lower than the one observed during the pre-adaptation trials (0.91). This reveals that saccadic gain did not perfectly recover. It should be pointed out that this value was comparable to the ones observed with the conventional double steps paradigm in which recovered gain averaged 0.86 after 150 recovery trials. Regardless of the exact degree of recovery, the goal of this experiment was to probe whether saccadic amplitude could increase even though a distractor was located as to encourage a low gain (Figure A.1C, right panels). This was clearly the case: the change in gain averaged +12% in the target recovery condition versus -2% in the distractor recovery condition (Figure A.4F).

Results from the three experiments reported here systematically revealed that adaptation – and recovery – is not impaired by the presence of a distractor. Because in our paradigms, the target was defined by being the object presented during the fixation period, not by its specific visual characteristics these results further confirm that adaptation does not depend on the visual properties of the target. Bahcall and Kowler (2000) previously reported that the shape of the adapting target had no influence on changes in gain while Deubel (1995) found a similar result with color. Instead, only a mechanism of error selection might explain the present results. We will discuss first, the general implications for understanding saccadic adaptation, and then the possible mechanism underlying this ability to select a particular error signal.

Effects of error signals on adaptation.

In the laboratory, oculomotor adaptation is conventionally studied by having subjects making saccades to a single simple target surreptitiously stepped during the movement, inducing a position error signal when the saccade lands. It has been proposed that this error signal controls the adjustment of saccade amplitudes over many trials (Noto and Robinson, 2001; Wallman and Fuchs, 1998). This implies that postsaccadic errors are fed back, making this signal obligatory for saccadic adaptation (Srimal et al, 2008). It could therefore be tempting to oversimplify the process involved in saccadic adaptation and view it as a passive change in saccade parameter to minimize position error signals. However, experimental data reveal that the processes underlying adaptation might be more complex, in particular regarding the effect of the position error. For instance the degree of adaptation is not proportional to the retinal error when the target is placed at different distances from saccade endpoint after each movement, and gain decreases occur even when placing the target on the fovea, (Robinson et al., 2000). Moreover, when instructing subjects to make saccades two-thirds of the way to the target and then backstepping the target during saccades, Bahcall and Kowler (2000) observed a decrease in saccadic gain even though the target was beyond the fovea after the saccade. Another support to the complex effects of the error signal comes from studies of contextual effects. We recently found that an identical initial target step could elicit saccades of two different amplitudes depending on the visual properties of the target (Herman et al 2009). In these experiments subjects learned to associate a specific gain with each context. Because step-back trials were randomly interleaved with nonstep-back trials, this implies the ability of selectively assign a particular error signal to drive one of the gains. Stated otherwise, successive error signals were differentially treated according to the target visual properties. Saccadic adaptation is also highly sensitive to proprioceptive context (Alahyane and Pelisson, 2004; Shelhamer and Clendaniel, 2002). Other evidence of contextual learning in saccadic adaptation is found when studying retention after conventional adaptation. In monkeys, some maintenance of adaptation has been reported after 24 hr under normal illumination (Noto, Watanabe, and Fuchs, 1999). In humans, the gain decrease induced by a back-step protocol lasted up to 5 days (Alahyane and Pelisson, 2005). Despite the thousands of saccades made between sessions during which the target did not step, the gain measured in the laboratory did not immediately recovered. A common view of motor learning propose that the error signal resulting from an action is compared to a predicted signal originating from a forward model (Wolpert et al 1998): the resulting signal would be used to adjust the future motor commands. Here motor learning consists essentially in learning the correct prediction of an action, given a particular context. Because various contexts - and hence various predictions - may coexist it follows that a given error signal may be contextualized to usefully drive changes in the motor command.

Our present experiments further challenge the view that saccadic adaptation might results from an adjustment in gain in response to any error signal by probing the effects of simultaneous – instead of successive – error signals. If any error signal could control saccadic adaptation (Figure A.1A), one would expect no adaptation either because of the averaging of the error signals or because of the random selection of one of the two possible error signals. Alternatively, one could postulate that one of the position error signals would systematically be favored by the saccadic system, either because it matches the target position following the initial step or because it encourages the natural tendency of the saccadic system of being hypometric. In the first case no adaptation would have occur in the first and second experiments when using both the double-step plus distractor and single-step plus distractor paradigms. In the latter case, a decrease in gain would have been systematic in experiment one, while no change in gain would have been observed in experiment two. This clearly was not the case, and instead the target-related error signal was able to drive the saccadic gain but the distractor-related error had no effect, supporting the notion that saccadic adaptation depend on the selection of an error signal. In a recent experiment it has also been shown that when the saccade is directed within a previously fixated object, no adaptation occurs but when the saccade is directed to a different object the saccadic gain decreases even though the whole configuration was shifted backward during the saccade in both conditions (Collins et al 2007). The lack of transfer from between-object adaptation to within-object saccades further confirmed these results. The authors suggest that adaptation requires the selection of an object as saccade target, in addition to a position error signal.

Selection of the error signal

It is commonly admitted that selective attention determines which of the various objects present in the environment is to be a target to guide a movement. A large number of studies have been devoted to explore the relations between attention and eye movements. Although it is logically possible that orienting attention and gaze orienting are two independent processes (Posner, 1980), several experiments have shown that a stimulus presented in the period preceding saccade execution is best processed when its location coincides with the saccadic target (Castet el al, 2006; Deubel and Schneider, 1996; Hoffman and Subramaniam, 1995; Kowler, Anderson, Dosher, and Blaser, 1995). Because improvements of perceptual performance are interpreted as the result of selective attention a strong coupling between the preparation of a saccade and a shift of covert selective attention is now widely accepted (Findlay and Gilchrist, 2003), although this coupling is not constant across time (Castet et al 2006, Montagnini and Castet, 2007). Moreover, covert attention and attentional selection of a saccadic target share common mechanisms (Findlay and Gilchrist, 2003; Kustov and Robinson, 1996; Moore, Armstrong, and Fallah, 2003; Nobre, Gitelman, Dias, and Mesulam, 2000; Rizzolatti et al., 1987). Relations between orienting attention and adaptation have been explored, showing that adapting attention leads to a change in saccade endpoint (McFadden et al, 2002) and, conversely, that attention is coupled to the adapted saccade (Doré-Mazars & Collins, 2005, Collines et al., 2010 but see Ditterich et al., 2000). Attention is therefore a natural candidate to account for the selection of a particular error signal following a saccade: by enhancing the target visual processing while inhibiting signals related to the distractor, selective attention would bias the oculomotor system toward the target error signal. In our experiments, this would induce saccadic adaptation in the direction of the target error signal and would explain why the error signal related to the distractor might be ignored. The temporal course of the selective attention as well as the visual characteristics necessary to select the target against the distractor remain unknown. Here the rather long fixation period (750-1250 ms) and the shapes and colors used to construct the stimuli were chosen to minimize the ambiguity between the target and distractor objects. However, these issues deserve further research.

Selective attention, however, is not itself a mechanism, but rather a capability which may be functionally supported by more than one neural mechanism (Desimone and Duncan, 1995). A priori, both selection based on spatial location and selection based on stimulus features (or some combination of the two) are consistent with the results we have observed. That is, the target might be selectively attended to (or the distractor ignored) based on spatial location or visual appearance. However, there are potential constraints on the use of these mechanisms in the context of saccade adaptation, as well as logical conclusions regarding their usage that are worth mentioning here. For example, it is known that postponing the post-saccadic presentation of a target (thus delaying the error computation) reduces the amount of adaptation that can be induced in the double-step paradigm (Fujita et al., 2002; Shafer et al., 2000). This raises the question of the length of time required for selection based on features; if the minimal interval required for such processing was excessively long it might lessen or even prevent adaptation. Since we have seen no reduction in adaptation as compared to the case when a single target is present after each saccade, either this process can proceed with minimal delay, or a spatial-attention mechanism (which is likely to be substantially more rapid) is at work. The use of selection based on spatial location, on the other hand, would require that the adaptation process induce a change in the expected post-saccade location of the target (Collins et al. 2009); or at least a shift in the immediate post-saccadic locus of spatial attention, since after the primary saccade, the target is consistently displaced relative to its pre-saccadic position. Further, selection based on spatial location would, in fact, require an initial feature-based selection. That is, early in the adaptation block, her attention, post-saccadically, is presumably directed towards the target's pre-saccade location. Thus it seems likely that a feature-based mechanism be employed to shift the attentional locus. However, the consistent post-saccadic presentation of the target in one location (relative to pre-saccade) could have the function of shifting the subject's expectation of the post-saccadic target location, and the post-saccadic locus of attention along with it. Thereafter, target selection could be achieved solely based on spatial location. Further exploration of attentional mechanisms and target selection in multiple-stimulus environments are needed to clarify this issue further.

Although, as we discussed, selective attention provides an appealing explanation for our results, we wish to explore another hypothetical alternative. One could postulate that when the saccade lands, the movement has two possible outcomes: either the gaze is closer to the target – and far from the distractor – and the goal of the movement is achieved, or the gaze is far from the target – and closer to the distractor – and the goal is not achieved. By attempting to reproduce the movements that were successful while avoiding the ones that were not, saccade amplitudes would progressively change to compensate for the target's adaptation step, ignoring the distractor. In other words, saccadic adaptation could be explained by a selection process in which those responses that produces a valuable outcome are selected over those that do not. We recently conducted a series of experiments in which changes in gain were controlled by reinforcement contingencies in the absence of a position error signal (Madelain et al 2008), a result compatible with this possibility. Other evidences exist that are compatible with the effects of reinforcement on saccades. For instance, in monkeys saccades to rewarded locations had shorter latencies (Ikeda and Hikosaka, 2003; Lauwereyns et al., 2002; Watanabe et al., 2003) and higher peak velocities than saccades of the same amplitude to non-rewarded locations (Takikawa et al., 2002). In humans, it has also been shown that saccades peak velocities increased when viewing the visual stimulus was rewarding (Xu-Wilson, Zee & Shadmehr, 2009) and we found that saccade latencies can be altered by reinforcement (Madelain et al., 2007).

The present experiments were not designed to probe whether attention or reinforcement learning is responsible for adaptation when facing several error signals, but to test whether adaptation might occur when a distractor is present. In fact, the question whether a change occurs because of attention or because of reinforcement is very difficult to answer (see Maunsell, 2004 for a review) in particular because the tool commonly used to control attention is the manipulation of reward. Nevertheless, efforts to distinguish the relative contribution of attention and reward at the behavioral and neural levels will be a valuable step towards establishing a precise description of neural mechanisms underlying complex behaviors.

Conclusions

These experiments provide strong evidence that saccadic adaptation occurs in the presence of two competing position errors, implying that the saccadic system is able to selectively use the signal originating from the target to change its gain, while ignoring the one related to a distractor. This result further confirms the complex effects of error signals on saccade parameters. Whether the mechanisms underlying saccadic adaptation rely on attention or reinforcement deserves further investigations.

Appendix B: Model

The modeling work presented here is based on data from chapter 4 in which we varied the percent of feedback during the adaptation phase to study the effect of intermittent information on the learning and retention of saccadic adaptation.

Methods

BRE-Model

This model was essentially identical to that used by Ethier et al. (2008a). We thus follow the conventions of their presentation. Saccade production was assumed to rely on a two-component "hidden" memory represented by **X**. The value of this state on the n+1th trial (prior to the primary saccade) depended on the previous (nth) trials' state $\boldsymbol{X^n}$, the error in the previous trial $\boldsymbol{E^n}$, and Gaussian noise $\boldsymbol{\varepsilon_y}$:

$$X^{n+1} = A \cdot X^n + B \cdot E^n + \varepsilon_X \qquad (1)$$

The parameter ***A*** (a diagonal 2 x 2 array) determines how the state will change from trial-to trial due to the passage of time, and ***B*** how the state will change due to error. This memory is hidden because it cannot be observed directly, only indirectly by observing saccade production, which is assumed to depend on the primary target step $\boldsymbol{p^n}$, an inherent bias $\boldsymbol{y_b}$, a weighted sum of the memory states (with unknown weighting, **C**), and execution noise $\boldsymbol{\varepsilon_y}$:

$$y^n = p^n - y_b + C^T \cdot X^n + \varepsilon_y \qquad (2)$$

Where ***c*** represents the relative weight of each state in determining saccade amplitude. Meanwhile, the error on the nth trial is assumed to depend on the primary target step, the intrasaccadic target step (ISS) $\boldsymbol{u^n}$, the saccade amplitude, and the bias: a bias-subtracted retinal error:

$$E^n = p^n + u^n - y^n - y_b \qquad (3)$$

In BLANK+ISS or no-ISS without feedback trials, $\boldsymbol{u^n}$ was chosen so that $\boldsymbol{E^n}$ would equal 0.

Fitting this model first consists of identifying the values of the system parameters ($\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$, and $\boldsymbol{y_b}$). To do this, Ethier et al. (2008a) chose to render this system of equations in a form easily solved by a method called subspace identification (van Overschee and de Moor, 1996; Ethier et al., 2008a). However, the form of our model (presented below) made the application of this method nontrivial. Instead, we used a constrained optimization technique that was equally easily applied to both models. Importantly, prior to this algorithmic switch, we first compared parameter and memory state estimates obtained via the subspace identification / Kalman filter method of Ethier et al. (2008a), to those obtained via our own constrained optimization method; finding no significant difference between these two methodologies, we subsequently used only the optimization technique for both models to keep our treatment of the two systems as similar as possible.

We also explored two variants on the BRE-model. In the first, we added an exponential scaling term to equation (3):

$$E^n = e^{D\cdot(t_{sc}^n - t_{ton}^n)} \cdot (p^n + u^n - y^n - y_b) \qquad (4)$$

Here, **D** determines the rate at which the error decays as a function of the difference between the time of saccade completion $\boldsymbol{t_{sc}}$, and the time that the target was re-illuminated $\boldsymbol{t_{ton}}$. In the second variant, rather than choosing $\boldsymbol{u^n}$ so that $\boldsymbol{E^n}$ would equal 0 in trials without feedback, we set $\boldsymbol{u^n}$ to 0 so that $\boldsymbol{E^n}$ would be computed on the basis of the target's location after the primary step:

$$E^n = p^n - y^n - y_b \qquad (5)$$

PRE-Model

Saccade production was presumed to rely on a hidden memory state $\boldsymbol{\theta^n}$, while error prediction relied on a separate hidden memory state $\boldsymbol{\phi^n}$. The value of these states in the n+1th trial depended on their values in the previous (nth) trial, the error in the previous trial, and Gaussian noise. Importantly, the error signals for the two single-component memory states were distinct, $\boldsymbol{\theta}$ relied on a "difference error" $\boldsymbol{E_d}$ and $\boldsymbol{\phi}$ on retinal error $\boldsymbol{E_r}$:

$$\theta^{n+1} = A \cdot \theta^{n} + B \cdot E_{d}^{n} + \varepsilon_{\theta} \qquad (6)$$

$$\varphi^{n+1} = A \cdot \varphi^{n} + B \cdot E_{r}^{n} + \varepsilon_{\theta} \qquad (7)$$

$$E_{d}^{n} = E_{r}^{n} - E_{p}^{n} \qquad (8)$$

$$E_{r}^{n} = p^{n} + u^{n} - y^{n} \qquad (9)$$

$$E_{p}^{n} = p^{n} + \varphi^{n} - y^{n} \qquad (10)$$

Combining equations (8), (9), and (10), we can see that the difference error is simply the difference between the ISS ($\boldsymbol{u^n}$) and the error memory ($\boldsymbol{\phi^n}$)

$$E_{d}^{n} = u^{n} - \varphi^{n} \qquad (11)$$

Where, as above, $\boldsymbol{p^n}$ represents the nth primary target step, $\boldsymbol{u^n}$ represents the intrasaccadic step, and $\boldsymbol{y^n}$ represents the primary saccade amplitude. Note that equation (8) contains the implicit assumption that prediction is based on a perfectly accurate estimate of saccade amplitude. This choice was made to highlight the extreme opposite viewpoint implicit in the BRE-model, that the accuracy of any internal estimate of saccade amplitude does not include any variation around the mean (see General Discussion for fuller consideration of this distinction). Meanwhile, the PRE-model assumes that target production relies only on the primary target step, the value of $\boldsymbol{\theta^n}$, and execution noise:

$$y^{n} = p^{n} - y_{b} + \theta^{n} + \varepsilon_{y} \qquad (11)$$

For the purposes of simulation, $\boldsymbol{\varepsilon_y}$ was set to 0.5°. This value was obtained by bootstrapping from a pool of mean-subtracted baseline phase saccades from all subjects in all conditions.

Finally, note that although the PRE-model appears more elaborate than the BRE-model, the BRE-model in effect has 9 parameters: $\boldsymbol{A}$ (2) + $\boldsymbol{B}$ (2) + $\boldsymbol{C}$ (2) + $\boldsymbol{X^0}$ (2) + $\boldsymbol{y_b}$ (1) = 9; While the PRE-model only has 4: $\boldsymbol{A}$ (1) + $\boldsymbol{B}$ (1) + $\boldsymbol{\theta^0}$ (1) and $\boldsymbol{\phi^0}$ (1) = 4. However, the PRE-model is presumed to have access to a non-visual estimate of saccade amplitude.

Results

Fitting with the BRE-model yielded mixed results, the learning dynamics evident in some datasets were captured well, while others were not (Figure B.1A). Intriguingly, fit-quality did not depend on percent feedback condition or subject: we found examples of good fits and poor fits irrespective of these factors (Figure B.1A). If, for example, the model provided better fits for the 100% but not other conditions, one might conclude that adaptation is driven by retinal error but that recovery depends on extra factors. That fit-quality did not depend on condition, however, suggested that both adaptation and recovery might depend on some extra-retinal error signal. Prior to including such a signal, however, we first wanted to determine that reasonable modifications to the BRE-model were insufficient to capture the behavior we observed.

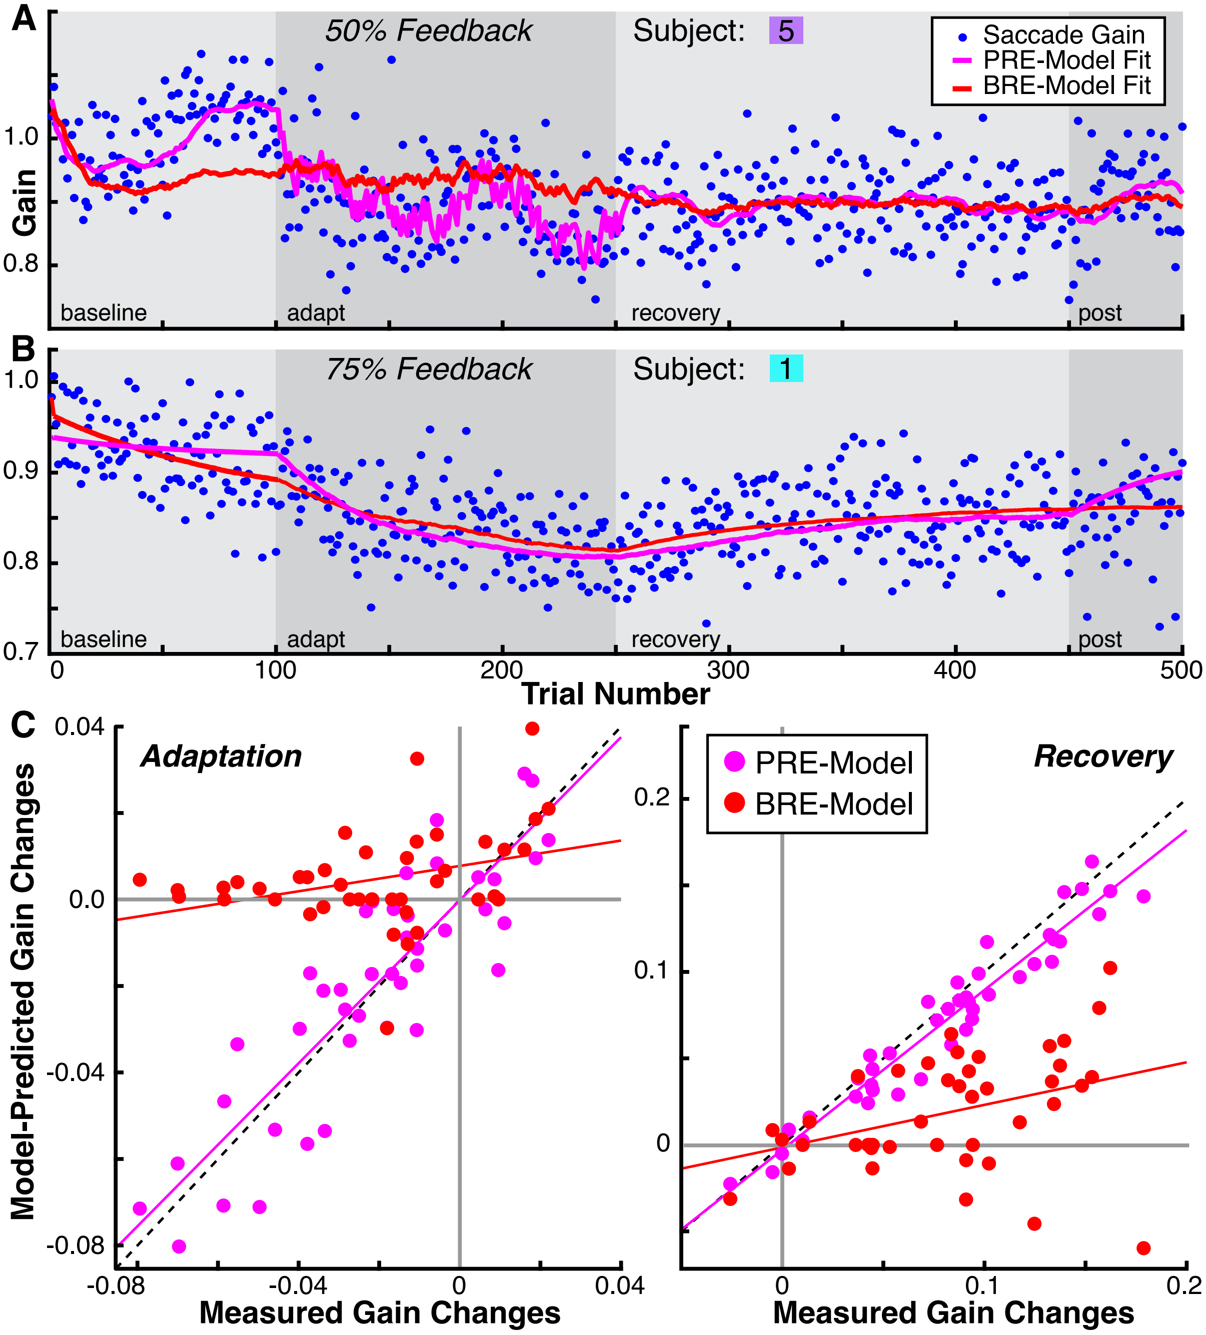

Figure B.1 Comparing Model Performance. A. An Example of Differing Fit Results. Subject 5's data from the 50%-feedback adaptation condition. The **BRE-model fit** (**red**) did not reflect baseline or adapt phase gain dynamics, though matched behavior adequately in recovery and post phases. Meanwhile, the **PRE-model fit** (**magenta**) better captured gain dynamics across phases. B. An Example of Similar Fit Results. Subject 1's data from the 75%-feedback condition. The **BRE-model** and **PRE-model** fits were comparable in this case, with the **BRE-model** capturing trends in all but the post-phase. C. Across subjects and conditions, a comparison of measured gain changes and model-predicted gain changes during adaptation (left) and recovery (right) showed that the **PRE-model's** behavior was far more consistently aligned with measured results than were those of the **BRE-model**. Dashed black lines are lines of identity, and colored lines are least-squares fits for the two models.

Figure B.1 Comparing Model Performance

We considered two modifications to the BRE-model, inspired by suggestions from the literature, neither of which improved fit quality or consistency (data not shown). We first considered allowing the impact of retinal error to decay exponentially from the time of saccade completion until the target was re-illuminated (to serve as the next trial's fixation). Motivated by the findings of Shafer et al. (2000), who pointed out that delaying post-saccadic target presentation (and thus RE) significantly diminished adaptation, this seemed potentially quite relevant since our paradigm did not entirely eliminate RE, but instead delayed it (by a variable period depending on primary saccade latency) for up to 1600ms. However, this change did not substantially modify BRE-model fits. Next, we considered the possibility that error was not 0° on no-feedback trials, but rather that error was computed from the target's position prior to the primary saccade. This possibility was motivated by the classic findings of Becker and Fuchs (1969), amongst others, who noted that subjects could often execute corrective saccades "in the dark" towards the location of a target that was extinguished during the primary saccade. Again, however, this modification was unsuccessful in substantively altering BRE-model fit quality or consistency. These efforts led us to finally consider a model including an additional, extra-retinal error signal.

A state-space model driven by deviations from predicted retinal error (PRE) yielded fits of overall better quality and that were more consistent across conditions (Figure B.1B). This model was structurally similar to the BRE-model, it learned from error at rate ***A***, and retained the memory at rate ***B***. The additional key assumptions of the PRE-model were: (1) retinal error is not discounted by a constant bias, but rather by a predicted retinal error; (2) this predicted retinal error is computed by combining estimates of target eccentricity and saccade amplitude with the history of retinal errors experienced. That is: prior to the introduction of any intrasaccadic target steps (ISSs), predictions of retinal error are based primarily on the estimated difference between target eccentricity and saccade amplitude; after many such ISSs, causing discrepant retinal errors, however, the model incorporates those experiences to form a prediction of retinal error more closely corresponding to the post-ISS target position. A 3 factor ANOVA (condition x subject x model-type) confirmed that the PRE-model had significantly smaller mean squared error than the BRE-model (F=102.08, $p=5.7\times10^{-24}$); post-hoc tests revealed that errors were smaller for each subject and each condition (α=0.05). More qualitatively, using each model's fits to predict the amount of adaptation and recovery for each subject in each condition revealed that the PRE-model yielded values that were substantially closer those measured (Figure B.2B). We again leave the full description of this PRE-model to the Methods section. Briefly, similar to the BRE-model, gain changes in the PRE-model were driven directly by a discounted retinal error. However, the additional key assumptions of this model were: (1) the discounting of retinal error is not by a constant bias, but rather by a prediction of retinal error which is (2) computed by combining a non-visual estimate of saccade amplitude, and the history of retinal errors experienced. That is: prior to any intrasaccadic target steps (ISSs), the predicted retinal error depends only on primary saccade amplitude; after many such ISSs, however, the model now predicts a retinal error more closely corresponding to the post-ISS target position. Note that this malleability is vital, a static prediction is not useful: as mentioned above, augmenting the BRE-model with the ability to compute retinal error on the basis of the pre-saccade target location does not improve its performance. Having concluded that the PRE-model best accounted for our data, we were now able to probe the relationship between model parameters and observed behavior.

Fitted learning rate parameter averages did not correspond well with adaptation rate observations from Experiment I (Figure B.2). In contrast to the observation that adaptation with 100% feedback was fastest and with 16% feedback was slowest, plotting the mean learning rate, ***A*** (see Methods), by condition did not reveal an obvious trend (Figure B.2C). A trend was also lacking in the mean forgetting rate, ***B***, but this was hardly in conflict with recovery observations, which did not suggest a trend. Perhaps this lack of correspondence between adaptation and learning rates was attributable to the structure of Experiment I, with an adapt phase long enough to accommodate considerable slowing of adaptation rate from beginning to end (Figure 4.2C).

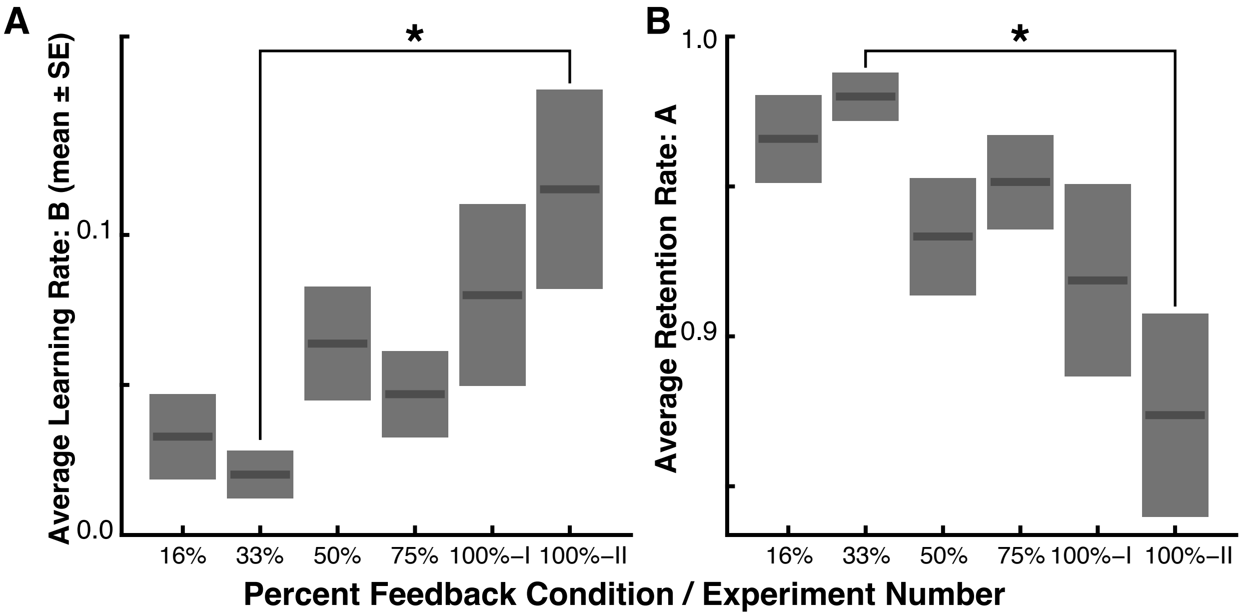

Figure B.2 Fitted Model Parameters. A. Average PRE-model learning rate, “B”, across subjects. Dark grey line segments and surrounding lighter grey boxes represent averages and standard errors across subjects. Larger values indicate a higher learning rate, fitting was constrained to keep B within the interval: [0,1]. Asterisk (*) denotes significantly different values as determined by a post-hoc multiple-comparison test following a 2-factor ANOVA (subject X condition). B. Average PRE-model retention rate ,“A”, across subjects. Larger values denote greater retention of learning, and fitting was again constrained to keep A within the interval: [0,1]. The asterisk (*) denotes a significant difference, from a post-hoc test after a seperate ANOVA.

Figure B.2 Fitted Model Parameters

In contrast, average learning and retention rate parameters fitted to data from Experiment II corresponded well with observed adaptation and recovery rates (Figure B.2). We found learning rate, ***B***, in the consecutive feedback condition (100%-II; 0.12 ± 0.03, Mean ± SE) to be significantly greater than in the 33% feedback condition (0.02 ± 0.01; t-test, p = 0.02). Meanwhile, the retention rate, ***A***, was significantly greater in the 33% feedback condition (0.98 ± 0.01), than in the consecutive feedback condition (0.87 ± 0.03; t-test, p = 0.01). These findings confirm our observations that feedback intermittency modulates learning and retention rates. However, since the model was quite good at capturing Experiment I data, we were puzzled as to why the correspondence between parameters and behavior was evident to a greater degree in Experiment II than I. We thus attempted to identify additional aspects of model performance that might explain the dissonance between model learning rates and observed adaptation rates in Experiment I.

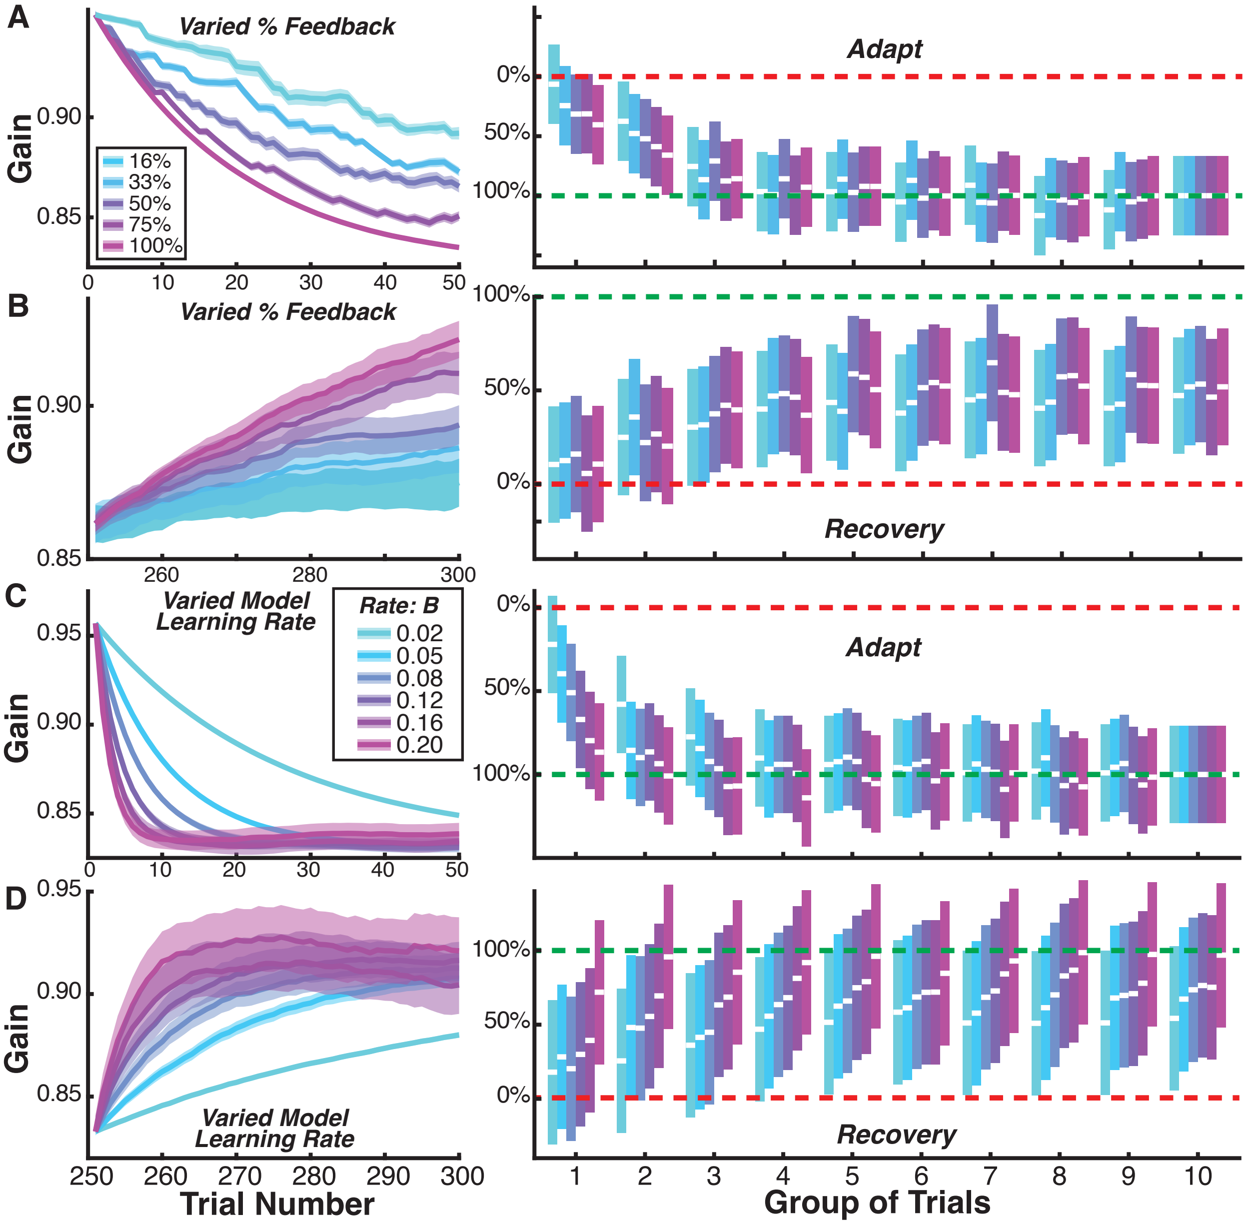

Figure B.3 PRE-model simulation data. **A**. Effects of varying % feedback on adaptation. Parameters were fixed at mean across-subject fit values and 7 runs were stimulated for each % feedback condition (for comparison to real data, comprising 7 subjects). **Left**: first 50 adaptation trials (additively normalized to first trial) averaged across subjects as in Figure 4.3A. **Right**: average percent gain change as in Figure 4.2C. **B**. The same simulation data replotted to examine recovery / retention. **Left**: as in Figure 4.2F. **Right**: as in 4.2F. **C**. The effects of varying model learning rate on adaptation. Other parameters were fixed, and % feedback was fixed (100%), while learning rate was varied (see Legend). 100 simulations were run in order to most clearly illustrate effects. **D**. Varied model learning rate effects on recovery, conventions as in B, above.

Figure B.3 PRE-model simulation data

Through simulation, we found that varying feedback frequency results in the appearance of varied adaptation rate, while varying model learning rate results in the appearance of both varied adaptation and recovery rates (Figure B.3). First, keeping PRE-model parameters fixed at average Experiment I values, we stimulated 7 adaptation sessions in each percent-feedback condition (1 run for each subject). Averaging across runs, and plotting adaptation percentage, as previously (Figure 4.2C), it is evident that this manipulation alone is sufficient to mimic the appearance of a varied rate of adaptation (Figure B.3A). Unsurprisingly (since parameters were fixed), recovery rate did not appear to be directly modulated by varying feedback frequency (Figure B.3B). Also to be expected was the finding that varying PRE-model learning rate (while keeping percent feedback fixed) appeared to directly modulate adaptation rate (Figure B.3C). Note that this plot illustrates the relatively large difference in adaptation rate resulting from the fitted learning rate values of Experiment II's intermittent feedback (0.02) and uninterrupted feedback (0.12) conditions, and the relatively small difference between Experiment I's 100% feedback condition (0.08) and Experiment II's uninterrupted feedback condition (0.12). Meanwhile, somewhat unexpectedly, we found that varying model learning rate also appeared to affect recovery rate (Figure B3.D).

The foregoing examination of model behavior suggests a more coherent interpretation of our results. The observation of variations in adaptation but not recovery rate in Experiment I is consistent with the absence of a trend in fitted model learning and retention rates. Simulation data instead suggests that adaptation rates appeared to vary primarily due to the direct impact of percent feedback on the progression of adaptation, since substantial variation in learning rate would have affected both learning and recovery rates. Meanwhile, observed differences in both adaptation and recovery rates in Experiment II are consistent with differences in fitted learning and retention rates.

In summary, modeling efforts first suggested that a model driven by deviations from predicted retinal error best accounted for our data. This observation was key, as it fundamentally changes the characterization of the retention / recovery period: in the BRE-model, withholding feedback results in a passive decay of memory, but because the PRE-model still treats the withholding of feedback as a deviation from prediction, active learning still occurs on those occasions. Second, in a seeming contradiction, we found that percent feedback did not affect rate parameters in Experiment I, but feedback intermittence seemed to affect rate parameters in Experiment II. Third, simulations showed that there was no conflict, suggesting that feedback intermittence does affect learning and retention rates, but that the duration (in trials) of a period of altered feedback also contributes to rate determination.

References

Abel LA, Schmidt D, Dell'Osso LF, Daroff RB. Saccadic system plasticity in humans. *Ann Neurol* 4: 313–318, 1978.

Adelson EH, Bergen JR. Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A* 2: 284–99, 1985.

Alahyane N, Pelisson D. Eye position specificity of saccadic adaptation. *Invest Ophthalmol Vis Sci* 45: 123–130, 2004.

Alahyane N, Pelisson D. Long-lasting modifications of saccadic eye movements following adaptation induced in the double-step target paradigm. *Learn Mem* 12: 433–443, 2005.

Albano JE, King WM. Rapid adaptation of saccadic amplitude in humans and monkeys. *Invest Opthalmol Vis Sci* 30: 1883–93, 1989.

Annett J. Learning a pressure under conditions of immediate and delayed knowledge of results. *Quarterly J Exp Psychol* 11: 3–15, 1959.

Anstis SM, Rogers BJ. Illusory reversal of visual depth and movement during changes of contrast. *Vis Res* 15: 957–61, 1975.

Anstis SM. Phi movement as a subtraction process. *Vis Res* 10: 1411–30, 1970.

d' Avossa G, Tosetti M, Crespi S, Biagi L, Burr DC, Morrone MC. Spatiotopic selectivity of BOLD responses to visual motion in human area MT. *Nat Neurosci*10: 249–55, 2007.

Bahcall DO, Kowler E. The control of saccadic adaptation: implications for the scanning of natural visual scenes. *Vis Res* 40: 2779–2796, 2000.

Bahill AT, Clark MR, Stark L. The main sequence, a tool for studying human eye movements. *Math Biosci* 24: 191–204, 1975.

Beaton K, Wong A, Shelhamer M. Context-specific adaptation of sensorimotor responses. In: *Annual Meeting of the Society for Neuroscience Abstracts*. San Diego: Society for Neuroscience, 2010.

Becker W, Fuchs AF. Further properties of the human saccadic system: eye movements and correction saccades with and without visual fixation points. *Vis Res* 9: 1247–1258, 1969.

Becker W. The control of eye movements in the saccadic system. *Bibl Opthalmol* 82: 233–43, 1972.

Becker W. The neurobiology of saccadic eye movements. Metrics. *Rev Oculomot Res* 3: 13–67, 1989a.

van Beers RJ. The sources of variability in saccadic eye movements. *J Neurosci* 27: 8757–8770, 2007.

Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat Neurosci* 10: 1214–1221, 2007.

Berniker M, Kording K. Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci* 11: 1454–1461, 2008.

Bilodeau E a, Bilodeau IM. Variable frequency of knowledge of results and the learning of a simple skill. *J Exp Psychol* 55: 379–83, 1958.

Bock O, Goltz H, Belanger S, Steinbach M. On the role of extraretinal signals for saccade generation. *Exp Brain Res* 104: 349–350, 1995.

Bonnetblanc F, Baraduc P. Saccadic adaptation without retinal postsaccadic error. *Neurorep* 18: 1399–402, 2007.

Bridgeman B, Hendry D, Stark L. Failure to detect displacement of the visual world during saccadic eye movements. *Vis Res* 15: 719–22, 1975.

Bryan CP. *The Papyrus Ebers*. London: Geoffrey Bles, 1930.

Castet E, Masson GS. Motion perception during saccadic eye movements. *Nat Neurosci*3: 177–83, 2000.

Cavanagh P. Attention-based motion perception. *Science* 257: 1563–5, 1992.

Chaturvedi V, van Gisbergen JA. Perturbation of combined saccade-vergence movements by microstimulation in monkey superior colliculus. *J Neurophysiol* 81: 2279–2296, 1999.

Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R. Adaptive control of saccades via internal feedback. *J Neurosci* 28: 2804–13, 2008.

Chubb C, Sperling G. Two motion perception mechanisms revealed through distance-driven reversal of apparent motion. *PNAS* 86: 2985–9, 1989.

Collins T, Rolfs M, Deubel H, Cavanagh P. Post-saccadic location judgments reveal remapping of saccade targets to non-foveal locations. *J Vis* 9: 291–9, 2009.

Collins T, Vergilino-Perez D, Beauvillain C, Doré-Mazars K. Saccadic adaptation depends on object selection: evidence from between- and within-object saccadic eye movements. *Brain Res* 1152: 95–105, 2007.

Collins T, Wallman J. The relative importance of retinal error and prediction in saccadic adaptation. *J Neurophysiol* 12: 3342–3348, 2012.

Corriou J-P. *Process Control*. London: Springer-Verlag, 2004.

Delbarre E, Delabarre EB. A Method of Recording Eye-Movements. *Am J Psychol* 9: 572–574, 1898.

Deubel H, Koch C, Bridgeman B. Landmarks facilitate visual space constancy across saccades and during fixation. *Vis Res* 50: 249–59, 2010.

Deubel H, Schneider WX, Bridgeman B. Postsaccadic target blanking prevents saccadic suppression of image displacement. *Vis Res* 36: 985–996, 1996.

Deubel H, Wolf W, Hauske G. Corrective saccades: effect of shifting the saccade goal. *Vis Res* 22: 353–364, 1982.

Deubel H, Wolf W, Hauske G. Adaptive gain control of saccadic eye movements. *Hum Neurobiol* 5: 245–253, 1986.

Deubel H. Adaptivity of Gain and Direction in Oblique Saccades. In: *Eye Movements: Form Physiology to Cognition*, edited by O'Regan JK, Levy-Schoen A. Amsterdam: Elsevier Science Publishers B.V., 1987, p. 181–190.

Deubel H. Adaptive Control of Saccade Metrics. In: *Presbyopia research: from molecular biology to visual adaptation*, edited by Obrecht, Gérard, Stark L. Springer New York, 1991, p. 93–100.

Deubel H. Separate adaptive mechanisms for the control of reactive and volitional saccadic eye movements. *Vis Res* 35: 3529–3540, 1995a.

Deubel H. Is Saccadic Adaptation Context-Specific? In: *Eye Movement Research: Mechanisms, Processes and Applications*, edited by Findlay JM, Kentridge RW, Walker R. Elsevier Science, 1995b.

Deubel H. 2 Separate Mechanisms for the Adaptive Control of Reactive, Volitional, and Memory-Guided Saccadic Eye Movements. In: *Attention and Performance XVII: Cognitive Regulation of Performance: Interaction of Theory and Application*, edited by Gopher D, Koriat A. A Bradford Book, 1999, p. 814.

Deubel H. Localization of targets across saccades: Role of landmark objects. *Visual Cognition* 11: 173–202, 2004.

Ditterich J, Eggert T, Straube A. Fixation errors and timing in sequences of memory-guided saccades. *Behav Brain Res* 95: 205–17, 1998.

Ditterich J, Eggert T, Straube A. Does Visual Background Information Influence Saccade Adaptation. In: *Current oculomotor research: physiological and psychological aspects*, edited by Becker, Wolfgang, Deubel, Heiner, Mergner T. Springer New York, 1999, p. 71–80.

Dodge R, Cline TS. The angle velocity of eye movements. *Psychol Rev* 8: 145–157, 1901.

Dodge R. Five types of eye movement in the horizontal meridian plane of the field of regard. *Am Journ Physiol* 8: 307–329, 1903.

Drew PJ, Abbott LF. Models and properties of power-law adaptation in neural systems. *J Neurophysiol* 96: 826–33, 2006.

Duhamel J-RR, Colby CL, Goldberg ME. The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255: 90–92, 1992.

Ebbinghaus H. *Memory: A Contribution to Experimental Psychology*. New York: Dover Publications, 1913.

Ebisawa Y. Unconstrained pupil detection technique using two light sources and the image difference method. In: *Visualization and Intelligent Design in Engineering and Architecture II*, edited by Brebbia C, Hernandez S. The University of Michigan, 1995, p. 271.

Edelman JA, Goldberg ME. Effect of short-term saccadic adaptation on saccades evoked by electrical stimulation in the primate superior colliculus. *J Neurophysiol* 87: 1915–1923, 2002.

Erkelens CJ, Hulleman J. Selective adaptation of internally triggered saccades made to visual targets. *Exp Brain Res* 93: 157–164, 1993.

Ethier V, Zee DS, Shadmehr R. Spontaneous recovery of motor memory during saccade adaptation. *J Neurophysiol* 99: 2577–2583, 2008a.

Ethier V, Zee DS, Shadmehr R. Changes in control of saccades during gain adaptation. *J Neurosci* 28: 13929–13937, 2008b.

Fitzgibbon E, Goldberg M. Short-term saccadic adaptation does not affect the saccades evoked by electrical stimulation of the monkey superior colliculus. In: *Annual Meeting of the Society for Neuroscience Abstracts*. Dallas: Society for Neuroscience, 1985.

Fitzgibbon EJ, Goldberg ME, Segraves MA. *Short Term Saccadic Adaptation in the Monkey*. Oxford, Eng.: Pergamon, 1986.

Frens MA, Van Opstal AJ. Monkey superior colliculus activity during short-term saccadic adaptation. *Brain Res Bull* 43: 473–483, 1997.

Fuchs A, Luschei E. Firing patterns of abducens neurons of alert monkeys in relationship to horizontal eye movement. *J Neurophysiol* 33: 382–392, 1970.

Fuchs A, Luschei E. The activity of single trochlear nerve fibers during eye movements in the alert monkey. *Exp Brain Res* 26: 78–89, 1971.

Fuchs AF, Reiner D, Pong M. Transfer of gain changes from targeting to other types of saccade in the monkey: constraints on possible sites of saccadic gain adaptation. *J Neurophysiol* 76: 2522–2535, 1996.

Fujita M, Amagai A, Minakawa F, Aoki M. Selective and delay adaptation of human saccades. *Brain Res Cogn Brain Res* 13: 41–52, 2002.

Gaveau V, Alahyane N, Salemme R, Desmurget M. Self-generated saccades do not modify the gain of adapted reactive saccades. *Exp Brain Res* 162: 526–531, 2005.

Goldberg ME, Wurtz RH. Activity of superior colliculus in behaving monkey. I. Visual receptive fields of single neurons. *J Neurophysiol* 35: 542–59, 1972.

Goldstein M, Rittenhouse C. Knowledge of results in the acquisition and transfer of a gunnery skill. *J Exp Psychol* 48: 187–96, 1954.

Graybiel A, Aosaki T, Flaherty A, Kimura M. The basal ganglia and adaptive motor control. *Science* 265: 1826–1831, 1994.

Hallett PE. Primary and secondary saccades to goals defined by instructions. *Vis Res* 18: 1279–1296, 1978.

Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.

Harris CM. Does saccadic undershoot minimize saccadic flight-time? A Monte-Carlo study. *Vis Res* 35: 691–701, 1995.

Harwood MR, Madelain L, Krauzlis RJ, Wallman J. The spatial scale of attention strongly modulates saccade latencies. *J Neurophysiol* 99: 1743–1757, 2008.

Hatada Y, Rossetti Y, Miall RC. Long-lasting aftereffect of a single prism adaptation: shifts in vision and proprioception are independent. *Exp Brain Res* 173: 415–24, 2006.

Havermann K, Lappe M. The influence of the consistency of postsaccadic visual errors on saccadic adaptation. *J Neurophysiol* 103: 3302–10, 2010.

Henson DB. Corrective saccades: effects of altering visual feedback. *Vis Res* 18: 63–67, 1978.

Hering E. Der Raumsinn und die Bewegungen des Auges. In: *Handbuch der Physiologie, Vol. 3, Part 1 Physiology des Gesichtssinnes*, edited by Hermann L. Leipzig: Vogel, 1879, p. 343–601.

Herman JP, Harwood MR, Wallman J. Saccade Adaptation Specific to Visual Context. *J Neurophysiol* 101: 1713–1721, 2009.

Ho L, Shea JB. Effects of relative frequency of knowledge of results on retention of a motor skill. *Percept Mot Skills* 46: 859–66, 1978.

Hopp JJ, Fuchs AF. Investigating the site of human saccadic adaptation with express and targeting saccades. *Exp Brain Res* 144: 538–548, 2002.

Hopp JJ, Fuchs AF. The characteristics and neuronal substrate of saccadic eye movement plasticity. *Prog Neurobiol* 72: 27–53, 2004.

Hopp JJ, Fuchs AF. Amplitude adaptation occurs where a saccade is represented as a vector and not as its components. *Vis Res* 46: 3121–3128, 2006.

Huang VS, Shadmehr R. Persistence of motor memories reflects statistics of the learning event. *J Neurophysiol* 102: 931–40, 2009.

Huey EB. Preliminary Experiments in the Physiology and Psychology of Reading. *Am J Psychol* 9: 575–586, 1898.

Humphreys LG. The effect of random alternation of reinforcement on the acquisition and extinction of conditioned eyelid reactions. *J Exp Psychol* 25: 141–158, 1939.

Johnson N, Kotz S, Kemp A. *Univariate Discrete Distributions*. Hoboken, NJ: Wiley-Interscience, 1993.

Joiner WM, Fitzgibbon EJ, Wurtz RH. Amplitudes and directions of individual saccades can be adjusted by corollary discharge. *J Vis* 10: 22.1–12, 2010.

Joiner WM, Smith MA. Long-term retention explained by a model of short-term learning in the adaptive control of reaching. *J Neurophysiol* 100: 2948–55, 2008.

Kagerer F a, Contreras-Vidal JL, Stelmach GE. Adaptation to gradual as compared with sudden visuo-motor distortions. *Exp Brain Res* 115: 557–61, 1997.

Kaku Y, Yoshida K, Iwamoto Y. Learning signals from the superior colliculus for adaptation of saccadic eye movements in the monkey. *J Neurosci* 29: 5266–5275, 2009.

Kalesnykas RP, Hallett PE. Retinal eccentricity and the latency of eye saccades. *Vis Res* 34: 517–531, 1994.

Kapoula Z, Robinson DA. Saccadic undershoot is not inevitable: saccades can be accurate. *Vis Res* 26: 735–743, 1986.

Keller EL, Robinson DA. Absence of a stretch reflex in extraocular muscles of the monkey. *J Neurophysiol* 34: 908–919, 1971.

Klassen J, Tong C, Flanagan JR. Learning and recall of incremental kinematic and dynamic sensorimotor transformations. *Exp Brain Res* 164: 250–9, 2005.

Kluzik J, Diedrichsen J, Shadmehr R, Bastian AJ. Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J Neurophysiol* 100: 1455–1464, 2008.

Kommerell G, Olivier D, Theopold H. Adaptive programming of phasic and tonic components in saccadic eye movements. Investigations of patients with abducens palsy. *Invest Ophthalmol Vis Sci* 15: 657–660, 1976.

Kording KP, Tenenbaum JB, Shadmehr R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat Neurosci* 10: 779–786, 2007.

Krauzlis RJ, Basso MA, Wurtz RH. Shared motor error for multiple eye movements. *Science* 276: 1693–1695, 1997.

Kroller J, De Graaf JB, Prablanc C, Pelisson D. Effects of short-term adaptation of saccadic gaze amplitude on hand-pointing movements. *Exp Brain Res* 124: 351–362, 1999.

Lamare M. Des mouvements des yeux dans la lecture. *Bulletin et Memoire de la Societe Francaise d'Ophthalmologie* 10: 354–364, 1892.

Lavergne L, Vergilino-Perez D, Lemoine C, Collins T, Doré-Mazars K. Exploring and targeting saccades dissociated by saccadic adaptation. *Brain Res* 1415: 47–55, 2011.

Leigh RJ, Zee DS. *The Neurology of Eye Movements*. 4th ed. New York: Oxford University Press, 2006.

Lu Z, Sperling G. Three systems for visual motion perception. *Curr Dir Psychol Sci* 5: 44–53, 1996.

Lu ZL, Sperling G. Attention-generated apparent motion. *Nature* 377: 237–9, 1995.

Luh CW. The conditions of retention. *Psychological Monographs* 31: 1–87, 1922.

Madelain L, Harwood MR, Herman JP, Wallman J. Saccade adaptation is unhampered by distractors. *J Vis* 10: 29, 2010.

Madelain L, Krauzlis RJ, Wallman J. Spatial deployment of attention influences both saccadic and pursuit tracking. *Vis Res* 45: 2685–2703, 2005.

Madelain L, Paeye C, Wallman J. Modification of saccadic gain by reinforcement. *J Neurophysiol* 106: 219–232, 2011.

Marko MK, Haith AM, Harran MD, Shadmehr R. Sensitivity to prediction error in reach adaptation. *J Neurophysiol* (July 5, 2012). doi: 10.1152/jn.00177.2012.

Marrocco RT, Li RH. Monkey superior colliculus: properties of single cells and their afferent inputs. *J Neurophysiol* 40: 844–60, 1977.

Galen (trans by May MT). *On the usefulness of the parts of the body.* Ithaca: Cornell University Press, 1968.

Mazzoni P, Krakauer JW. An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26: 3642–3645, 2006.

McGuigan F. The effect of precision, delay, and schedule of knowledge of results on performance. *J Exp Psychol* 58: 79–84, 1959.

McLaughlin SC. Parametric adjustment in saccadic eye movements. *Percept Psychophys* 2: 359–362, 1967.

Melnick MJ. Effects of overlearning on the retention of a gross motor skill. *Res Q* 42: 60–9, 1971.

Miller JM, Anstis T, Templeton WB. Saccadic plasticity: parametric adaptive control by retinal feedback. *J Exp Psychol Hum Percept Perform* 7: 356–366, 1981.

Moschovakis AK, Scudder CA, Highstein SM. The microscopic anatomy and physiology of the mammalian saccadic system. *Prog Neurobiol* 50: 133–254, 1996.

Munoz DP, Broughton JR, Goldring JE, Armstrong IT. Age-related performance of human subjects on saccadic eye movement tasks. *Exp Brain Res* 121: 391–400, 1998.

Munuera J, Morel P, Duhamel JR, Deneve S. Optimal sensorimotor control in eye movement sequences. *J Neurosci* 29: 3026–3035, 2009.

Noto CT, Robinson FR. Visual error is the stimulus for saccade gain adaptation. *Brain Res Cogn Brain Res* 12: 301–305, 2001.

Optican LM, Miles FA. Visually induced adaptive changes in primate saccadic oculomotor control signals. *J Neurophysiol* 54: 940–958, 1985.

Optican LM, Robinson DA. Cerebellar-dependent adaptive control of primate saccadic system. *J Neurophysiol* 44: 1058–1076, 1980.

Ottes FP, Van Gisbergen JA, Eggermont JJ. Collicular involvement in a saccadic colour discrimination task. *Exp Brain Res* 66: 465–478, 1987.

Overduin S a, Richardson AG, Lane CE, Bizzi E, Press DZ. Intermittent practice facilitates stable motor memories. *J Neurosci* 26: 11888–92, 2006.

van Overschee P, de Moor BLR. *Subspace identification for linear systems: theory, implementation, applications, Volume 1*. Ann Arbor: Kluwer Academic Publishers, 1996.

Pavlov I (trans. by GV Anrop). *Conditioned reflexes*. London: Oxford University Press, 1927.

Porterfield W. An essay concerning the motions of our eyes. Part I. Of their external motions. *Edinburgh Medical Essays and Observations* 4: 124–294, 1737.

Prablanc C, Jeannerod M. Corrective saccades: dependence on retinal reafferent signals. *Vis Res* 15: 465–469, 1975.

Pélisson D, Alahyane N, Panouillères M, Tilikete C. Sensorimotor adaptation of saccadic eye movements. *Neurosci Biobehav Rev* 34: 1103–1120, 2010.

Robinson DA. A Method of Measuring Eye Movemnent Using a Scieral Search Coil in a Magnetic Field. *IEEE transactions on bio-medical engineering* 10: 137–45, 1963.

Robinson DA. The Mechanics of Human Saccadic Eye Movement. *J Physiol* 174: 245–64, 1964.

Robinson DA. Oculomotor unit behavior in the monkey. *J Neurophysiol* 33: 393–403, 1970.

Robinson DA. Oculomotor control signals. *Basic Mechanisms of Ocular Motility and Their Clinical Implications*. .

Robinson F, Noto C, Watanabe S. Effect of visual background on saccade adaptation in monkeys. *Vis Res* 40: 2359–2367, 2000.

Robinson FR, Noto CT, Bevans SE. Effect of visual error size on saccade adaptation in monkey. *J Neurophysiol* 90: 1235–1244, 2003.

Robinson FRFR, Fuchs AFA, Noto CTCT. Cerebellar Influences on Saccade Plasticity. *Ann NYAS* 956: 155–163, 2002.

Schendel JD, Hagman JD. On sustaining procedural skills over a prolonged retention interval. *J Appl Psychol* 67: 605–610, 1982.

Schiller P. The discharge characteristics of single units in the oculomotor and abducens nuclei of the unanesthetized monkey. *Exp Brain Res*. .

Schmidt RA, Young DE, Swinnen S, Shapiro DC. Summary knowledge of results for skill acquisition: support for the guidance hypothesis. *J Exp Psychol. Learning, memory, and cognition* 15: 352–9, 1989.

Scudder CA, Batourina EY, Tunder GS. Comparison of two methods of producing adaptation of saccade size and implications for the site of plasticity. *J Neurophysiol* 79: 704–715, 1998.

Semmlow JL, Gauthier GM, Vercher JL. Mechanisms of short-term saccadic adaptation. *J Exp Psychol Hum Percept Perform* 15: 249–258, 1989.

Shadmehr R, Smith M a, Krakauer JW. Error correction, sensory prediction, and adaptation in motor control. *Ann Rev Neurosci* 33: 89–108, 2010.

Shafer JL, Noto CT, Fuchs a F. Temporal characteristics of error signals driving saccadic gain adaptation in the macaque monkey. *J Neurophysiol* 84: 88–95, 2000.

Shawkat FS, Harris CM, Wilson J, Taylor DS. Eye movements in children with opsoclonus-polymyoclonus. *Neuropediatrics* 24: 218–23, 1993.

Sheffield VF. Extinction as a function of partial reinforcement and distribution of practice. *J Exp Psychol* 39: 511–26, 1949.

Shelhamer M, Clendaniel R. Sensory, motor, and combined contexts for context-specific adaptation of saccade gain in humans. *Neurosci Lett* 332: 200–204, 2002a.

Shelhamer M, Clendaniel RA. Context-specific adaptation of saccade gain. *Exp Brain Res* 146: 441–450, 2002b.

Shen K, Paré M. Neuronal activity in superior colliculus signals both stimulus identity and saccade goals during visual conjunction search. 7: 1–13, 2007.

Skinner B. *The behavior of organisms*. New York: Appleton-Century, 1938.

Smith MA, Ghazizadeh A, Shadmehr R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4: e179, 2006.

Snow R, Hore J, Vilis T. Adaptation of saccadic and vestibulo-ocular systems after extraocular muscle tenectomy. *Invest Ophthalmol Vis Sci* 26: 924–31, 1985.

Soetedjo R, Fuchs AF, Kojima Y. Subthreshold activation of the superior colliculus drives saccade motor learning. *J Neurosci* 29: 15213–15222, 2009.

Sommer MA, Wurtz RH. What the brain stem tells the frontal cortex. II. Role of the SC-MD-FEF pathway in corollary discharge. *J Neurophysiol* 91: 1403–1423, 2004.

Sparks DL, Mays LE. Spatial localization of saccade targets. I. Compensation for stimulation-induced perturbations in eye position. *J Neurophysiol* 49: 45–63, 1983.

Sparks DL. The brainstem control of saccadic eye movements. *Nat Rev Neurosci* 3: 952–964, 2002.

Srimal R, Diedrichsen J, Ryklin EB, Curtis CE. Obligatory adaptation of saccade gains. *J Neurophysiol* 99: 1554–1558, 2008.

Straube A, Deubel H. Rapid gain adaptation affects the dynamics of saccadic eye movements in humans. *Vis Res* 35: 3451–3458, 1995.

Straube A, Fuchs AF, Usher S, Robinson FR. Characteristics of saccadic gain adaptation in rhesus macaques. *J Neurophysiol* 77: 874–895, 1997.

Taylor EA. *Controlled reading: a correlation of diagnostic, teaching, and corrective techniques*. The University of Chicago press, 1937.

Todorov E. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput* 17: 1084–1108, 2005.

Vergilino-Perez D, Collins T, Dore-Mazars K. Decision and metrics of refixations in reading isolated words. *Vis Res* 44: 2009–2017, 2004.

Vergilino-Perez D, Findlay JM. Between-object and within-object saccade programming in a visual search task. *Vis Res* 46: 2204–16, 2006.

Verstraten FA, Cavanagh P, Labianca AT. Limits of attentive tracking reveal temporal properties of attention. *Vis Res* 40: 3651–64, 2000.

Waitzman DM, Ma TP, Optican LM, Wurtz RH. Superior colliculus neurons provide the saccadic motor error signal. *Exp Brain Res* 72: 649–52, 1988.

Walker MF, Fitzgibbon EJ, Goldberg ME. Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. *J Neurophysiol* 73: 1988–2003, 1995.

Wallman J, Fuchs AF. Saccadic gain modification: visual error drives motor adaptation. *J Neurophysiol* 80: 2405–2416, 1998.

Warabi T, Kase M, Kato T. Effect of aging on the accuracy of visually guided saccadic eye movement. *Ann Neurol* 16: 449–454, 1984.

Wei K, Körding K. Relevance of error: what drives motor adaptation? *J Neurophysiol* 101: 655–64, 2009.

Winstein CJ, Schmidt RA. Reduced frequency of knowledge of results enhances motor skill learning. *J Exp Psychol Learn Mem Cogn* 16: 677–691, 1990.

Wolpert DM, Miall RC, Kawato M. Internal models in the cerebellum. *TICS* 2: 338–347, 1998.

Wong AL, Shelhamer M. Sensorimotor adaptation error signals are derived from realistic predictions of movement outcomes. *J Neurophysiol* 105: 1130–40, 2011.

Wyman D, Steinman RM. Letter: Latency characteristics of small saccades. *Vis Res* 13: 2173–2175, 1973.

Zee DS, Optican LM, Cook JD, Robinson DA, Engel WK. Slow saccades in spinocerebellar degeneration. *Arch Neurol* 33: 243–251, 1976a.

Zee DS, Robinson DA. A hypothetical explanation of saccadic oscillations. *Ann Neurol* 5: 405–414, 1979.

Zee DS, Yee RD, Cogan DG, Robinson DA, Engel WK. Ocular motor abnormalities in hereditary cerebellar ataxia. *Brain* 99: 207–34, 1976b.

Zimmermann E, Lappe M. Eye position effects in oculomotor plasticity and visual localization. *J Neurosci* 31: 7341–8, 2011.