Title
Flexible strategies for sensory integration during motor planning

Author
Sober, Samuel Judah

Publication Date
2005

Peer reviewed|Thesis/dissertation

Flexible Strategies for Sensory Integration During Motor Planning

by

Samuel Judah Sober

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

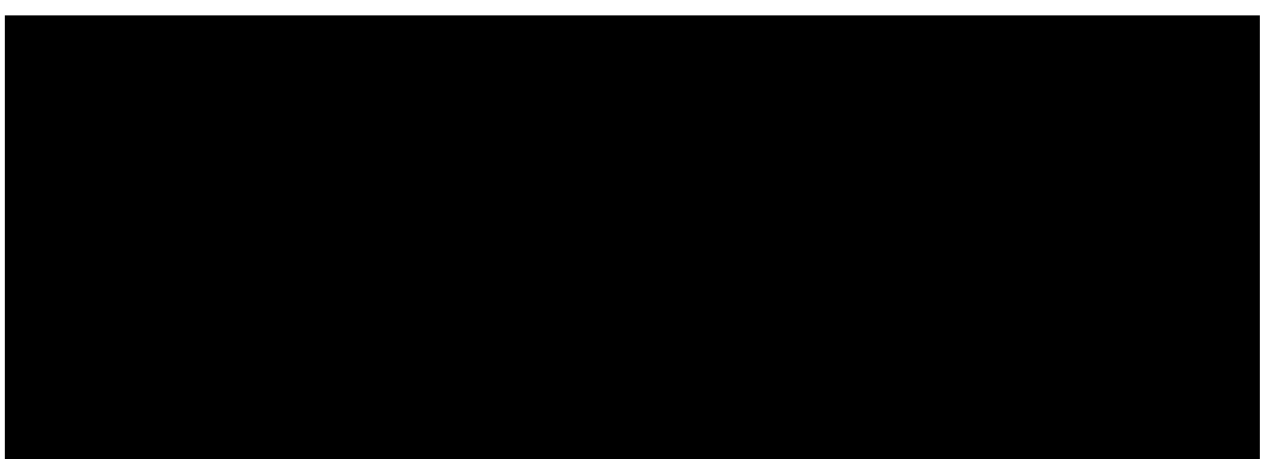

Date University Librarian

Degree Conferred:.....

This page left blank intentionally.

Acknowledgments

I owe my interest in neuroscience to the inspiration and encouragement of many people. First to Allan Berlind, whose thoughtful and clear-headed lectures introduced me to the field. His teaching provided an excellent introduction to the rewards and frustrations of thinking about the brain at a systems level, and I am grateful to have had him as an advisor. I am also thankful for the generosity of Bill Lytton and Sonny Yamasaki, who took a chance on me as an inexperienced summer research intern. They gave me a problem to work on, the freedom to approach it however I chose, and the chance to watch them argue with each other, and in doing so taught me a great deal about the practice of research. I am also indebted to John Kirn and Rick Jensen for supporting my early interest in research.

I am incredibly fortunate to have had Flip Sabes as my graduate advisor. The clarity and depth of his thinking about motor control guided me through my initial readings on the subject, and the rigor with which he approaches experimental problems shaped and strengthened the studies described in this thesis. Additionally, Steve Lisberger and Michael Stryker provided excellent critical advice throughout my time at UCSF, and I am grateful for their take-no-prisoners approach to serving on my thesis committee and (along with Emo Torodov) administering the oral defense.

My family has been a constant source of love, humor, and joy since long before I started this work. I am lucky to have parents who have always encouraged me to pursue whatever I love to do, and luckier still to have had them living so close to me for a year and a half in California. The Sober migration to the Bay Area never fully materialized, and Wisconsin and North Carolina might be far away, but my family always feels close.

My friends have been the best part of my time in San Francisco and have given me so much happiness and support over the last six years. Thanks to Rebecca Blank for forgiving me for her couch, Jeremy Kay for ignoring my shoes long enough to be a great friend, Stephanie Albin and Patrick Haddick for hazy late nights, Dan Engber for more free movies and Vietnamese beef than a good man can bear, Mandy Green for always being there, Summer Halas for raging through town every so often, Erik Olsen for 20 years of friendship, and Eric and Ashley Frankelittle for sharing so many decadent meals and beautiful hikes. And to Ave, for rainstorms and oysters, picnics and insurgency movies, and for making this last year so wonderful.

* * *

Note: The text of Chapter 1 of this dissertation is a reprint of the material as it appears in the Journal of Neuroscience. The text of Chapter 2 is a reprint of the material as it appears in Nature Neuroscience. The co-author listed on both of these publications directed and supervised the research that forms the basis for the dissertation.

Flexible Strategies for Sensory Integration during Reach Planning

Samuel Judah Sober

Abstract

The brain's success in perceiving and interacting with the world depends on its ability to assimilate multiple streams of sensory information. One particularly important sensory integration problem concerns the use of sensory feedback from the body to plan and execute arm movements. Information about arm position is available to both vision and proprioception, and these signals must be integrated in order to plan accurate reaches. Although both modalities can encode arm position, the question of how the brain weights these signals has not been systematically investigated. To address this issue, we developed a psychophysical technique for quantifying the relative reliance on vision versus proprioception during reach planning. Subjects planned and executed reaching movements while receiving virtual visual feedback. By analyzing the reaching errors that resulted from shifting the visual feedback away from the true hand position, we measured sensory integration at two different stages of reach planning. As described in Chapter 1, we found that the first planning stage relies mostly on visual information, whereas the second stage relies more on proprioceptive signals. This difference suggests that the planning computations performed at each stage inform the selection of sensory input. This hypothesis is further supported by the results presented in Chapter 2, which demonstrate that altering the details of the sensorimotor task drives changes in sensory integration. In Chapter 3, we explore the consequences of sensorimotor adaptation on sensory integration, and show that the effects of learning can drive changes at a single stage of reach planning. Lastly, in Chapter 4 we present a set of hypotheses about where and how the two planning stages might be implemented in the primate brain, and suggest a series of experiments designed to confirm the computational role of single neurons.

Stephen G. Lisberger

Table of contents:

Chapter 1: Sensory Integration during Motor Planning 1

Appendix to Chapter 1 36

Chapter 2: Flexible Strategies for Sensory Integration during Motor Planning ... 40

Supplementary Information to Chapter 2 66

Chapter 3: Sensory Integration following Sensorimotor Adaptation 78

Chapter 4: Possible Neural Mechanisms and Future Directions 106

References 135

List of Figures and Tables:

Chapter 1

Figure 1: Misestimation of arm position results in two types of reach errors 4
Figure 2: Data collection and experimental design 5
Figure 3: Two models of feedforward motor planning 9
Figure 4: Predicted shift-induced errors in initial reach direction 19
Figure 5: Data and velocity-command model fit from one subject 20
Figure 6: Initial velocity data averaged across all subjects 21
Figure 7: Best-fit values of α_{MV} and α_{INV} for all subjects and both models 23
Figure 8: Initial reach segments and predicted INV error 27
Figure 9: Effects of visual shifts on reach velocity magnitude 29

Chapter 2

Figure 10: Two types of reaching errors induced by visual shift 42
Figure 11: Predicted changes in Experiments 1 and 2 45
Figure 12: Data collection and task parameters 46
Figure 13: Sample data, Experiment 1 47
Figure 14: Group data, Experiment 1 49
Figure 15: Model fits, Experiment 1 50
Figure 16: Sample and group data, Experiment 2 52
Figure 17: Model fits, Experiment 2 53

Supplementary Information to Chapter 2

Table 1: ANOVA results, Experiment 1 68
Table 2: ANOVA results, Experiment 2 68
Figure 18: Results of intermanual transfer experiment 71
Figure 19: Methods and model fits, gaze direction control experiment 73
Figure 20: Methods and model fits, passive proprioceptive experiment 75

Chapter 3

Figure 21: Task design 84
Figure 22: Aftereffects 93
Figure 23: Data from a representative subject, with-feedback trials 95
Figure 24: Visual shift-induced directional errors, averaged across subjects 96
Figure 25: Reach trajectories from a representative subject, no-feedback trials ...97
Figure 26: Predicted changes in reach direction from a single subject 98
Figure 27: Fit values of α_{INV} in the with- and no-feedback conditions 99

Chapter 4

Figure 28: Planning stages and potential cortical sites 109
Figure 29: Tuning for initial hand position in the parietal cortex 115
Figure 30: Spatial distribution of hand position estimates 117
Figure 31: Proposed recording experiment 120
Figure 32: Psychophysical results from Chapter 2 121
Figure 33: Proposed stimulation experiment 122
Figure 34: Hypothesized costs for coordinate transformations 125
Figure 35: Target modality and sensory integration 127
Figure 36: Visual feedback type and sensory integration 129
Figure 37: Target modality and sensory integration: top-down control 131

Chapter 1: Sensory Integration during Motor Planning

Abstract

When planning goal-directed reaches, subjects must estimate the position of the arm by integrating visual and proprioceptive signals from the sensory periphery. These integrated position estimates are required at two stages of motor planning: first to determine the desired movement vector, and second to transform the movement vector into a joint-based motor command. We quantified the contributions of each sensory modality to the position estimate formed at each planning stage. Subjects made reaches in a virtual reality environment in which vision and proprioception were dissociated by shifting the location of visual feedback. The relative weighting of vision and proprioception at each stage was then determined using computational models of feedforward motor control. We found that the position estimate used for movement vector planning relies mostly on visual input, whereas the estimate used to compute the joint-based motor command relies more on proprioceptive signals. This suggests that when estimating the position of the arm, the brain selects different combinations of sensory input based on the computation in which the resulting estimate will be used.

Sensory channels often provide redundant information, as is the case when both visual and proprioceptive feedback encode the position of the arm. Recent studies suggest that when integrating redundant signals, the brain forms a statistically optimal (i.e. minimum-variance) estimate by weighting each modality according to its relative precision. Minimum-variance models have been shown to account for human performance when subjects integrate vision and touch (Ernst and Banks, 2002), vision and audition (Ghahramani, 1995), and other combinations of sensory input (Welch et al., 1979; Jacobs, 1999; van Beers et al., 1999). These models are appealing because they provide a simple rule by which the brain could minimize errors due to sensory noise.

While these models predict a single, optimal estimate, other lines of research suggest that the brain forms multiple and sometimes inconsistent estimates of environmental variables. For example, studies of patients with temporal or parietal lobe lesions indicate that the brain has independent streams of visual processing for perceptual as opposed to motor tasks (Goodale and Milner, 1992; Milner and Goodale, 1995). Studies of reaching to illusory objects have shown a similar dissociation in normal subjects (Aglioti et al., 1995; Haffenden et al., 2001). These results suggest that sensory signals might be processed differently depending on how they will be

Here we focus on the integration of visual and proprioceptive feedback from the arm before the execution of a reach. This study seeks to quantify how vision and proprioception are combined to estimate arm position and to determine whether the nervous system uses a single criterion (e.g. minimum-variance) to combine the two modalities, or if different combinations of sensory input are selected at each stage of motor planning. Our approach is to displace the visual feedback from the arm before movement onset and use the resulting movement errors to infer the relative weighting given to each sensory modality.

Our analysis relies on the premise that estimates of arm position (a term we use to denote both the position of the fingertip and the angles of the joints) are used in two separate stages of motor planning and on the observation that distinct patterns of movement errors would result from position misestimation at each stage. In the first stage, a desired movement vector in visual (extrinsic) space is computed by subtracting the estimated initial arm position from the target location. Clearly, if this initial position is misestimated, the planned movement vector will be wrong (Rossetti et al., 1995). We will refer to the resulting error pattern, illustrated in Figure 1A and B, as movement vector (MV) error. A second and perhaps less intuitive source of error is the transformation of the extrinsic movement vector into a joint-based (intrinsic) motor command (Ghilardi et al., 1995; Goodbody and Wolpert, 1999). This transformation is equivalent to evaluating an inverse model of the arm (Jordan, 1996) and requires an estimate of the arm's initial position. Position misestimation at this stage of reach planning will also result in movement errors. An example is illustrated in Figure 1C. A leftward shift in estimated arm position causes the subject to choose the wrong motor commands (an extension or flexion of the elbow), resulting in clockwise (CW) errors in initial reach direction (Figure 1D). We will refer to this type of error as inverse model (INV) error. While both MV and INV error result from misestimation of the arm's initial position, the two stages of motor planning may rely on two different position estimates.

The integration of vision and proprioception at these two planning stages has never been independently and simultaneously characterized. Here we show that shifts of visual feedback before movement onset result a combination of the MV and INV error patterns. Fitting the observed errors with simple mathematical models of motor planning allows us to quantify the relative contributions of vision and proprioception to the position estimate used at each planning stage.

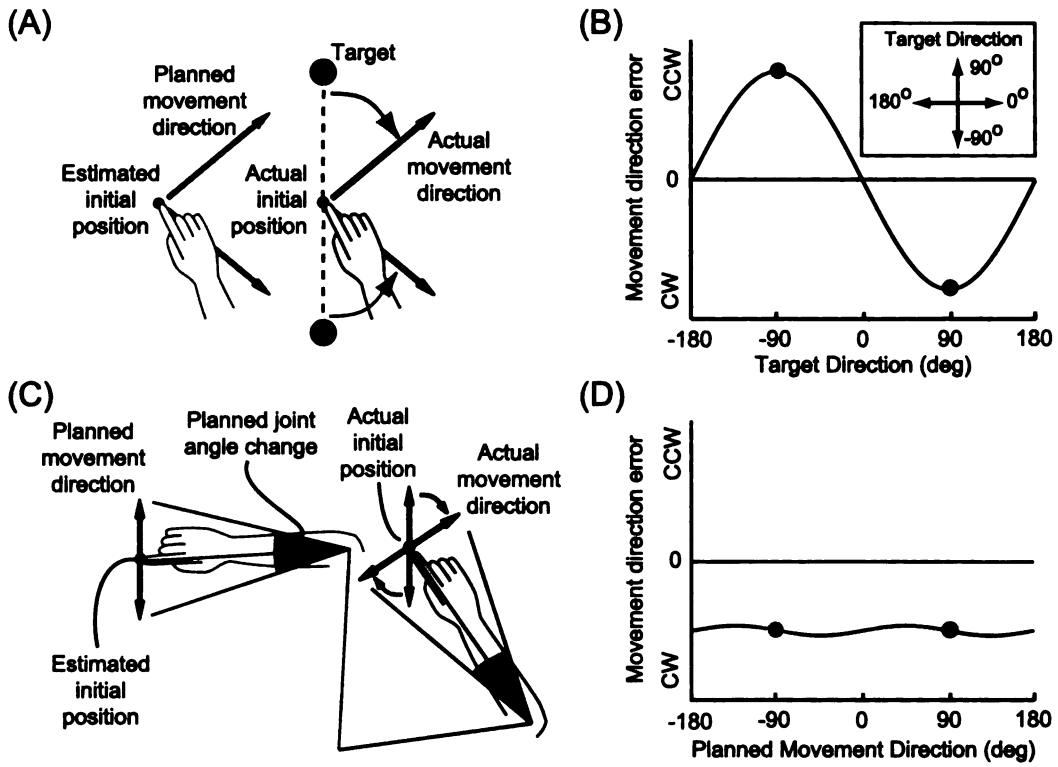

Figure 1: Misestimation of arm position results in two types of reach errors. (A) Errors resulting from a leftward shift in the position estimate used to plan the movement vector. The planned movement directions (gray arrows) differ from the actual hand-to-target directions (dashed lines). The pattern of directional errors (colored arrows) is plotted as a function of target direction in (B). A rightward shift would produce the opposite pattern (see Figure 4B). CW: clockwise; CCW: counterclockwise. (C) Errors resulting from a leftward shift in the position estimate used to transform the desired movement vector into a joint-based motor command. The directions of the achieved movements (black arrows) differ from the planned movement directions (grey arrows). The pattern of errors (colored arrows) is plotted as a function of planned movement direction in (D). The leftward shift shown here produces CW errors for all planned reach directions. A rightward shift would produce CCW errors.

Materials and Methods

Experimental Setup and Data Collection

Seven right-handed subjects (2 female, 5 male) participated in the experiment. Subjects were 26-33 years of age and were healthy with normal or corrected-to-normal vision. All subjects were naive to the purpose of the experiment and were paid for their participation.

The task was performed with the right arm, which rested on a shoulder-height table (Figure 2A). In order to minimize friction, the arm was supported by air sleds (0.73 kg upper arm, 1.18 kg forearm). The wrist was pronate and fixed in the neutral position with a brace, and the index finger was extended in a custom splint that permitted only vertical movement of the digit. Both shoulders were lightly restrained to minimize movement of the torso. This configuration restricted movement of the arm to two degrees of freedom and to a horizontal plane just above the table. Arm position could therefore be expressed interchangeably as x, a two-dimensional vector representing the Cartesian position of the fingertip, or as θ, a two-dimensional vector composed of the shoulder and elbow angles.

Three dowels, which served as tactile start points, were fixed a few centimeters above the plane of movement and could only be reached when the subject raised his or her fingertip from the splint. The visual feedback spot and target circles were presented via a mirror and rear-projection screen such that the images appeared to lie in the plane of the arm. An LCD projector (1024 $\times$ 768 pixels) with a 75 Hz refresh rate was used. A drape prevented vision of the arm, the table, and the dowels. Five infrared-emitting diodes were attached to each subject's arm and torso (Figure 2B). Arm position data were sampled at 120 Hz using an infrared tracking system (OPTOTRAK, Northern Digital, Waterloo, Ontario). Elbow and shoulder

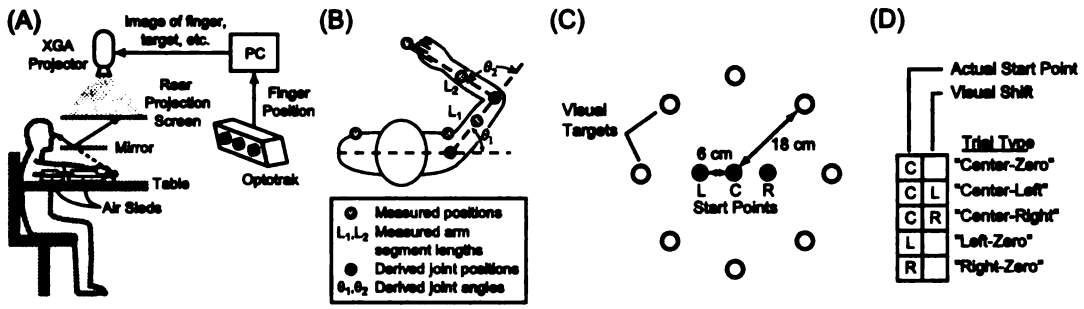

Figure 2: Data collection and experimental design. (A) Side view of the behavioral apparatus. For clarity, the tactile start points (dowels) and drape are not shown. (B) Top view of a subject showing the placement of the infrared markers (gray dots). Joint angles (θ_1 and θ_2) were computed from these five positions and from the measured lengths of the upper and lower arm (L_1 and L_2). (C) Workspace configuration. The dowels marking the three tactile start points (gray dots) were arranged in a line parallel to the subject's left–right axis, and were spaced 6 cm apart. L – left start point, C – center start point, R – right start point. (D) Trial types. Grid entries indicate the tactile start point (left column) and the presence of a leftward or rightward shift (right column) in each trial type.

angles were computed using the five marker positions and the lengths of the upper arm and forearm, which were measured using standard anatomical landmarks.

Task Design

The workspace contained the three start points and eight targets (Figure 2C). The dowels marking the start points were spaced 6 cm apart, and the center dowel was positioned approximately 40 cm from the subject's chest and slightly to the right of midline. The targets were evenly arrayed on a circle of radius 18 cm centered at the middle start point.

The experiment consisted of 160 trials. At the beginning of every trial, text reading "Left", "Center", or "Right" appeared briefly at a random location on the screen, instructing the subject to locate the appropriate dowel with his or her raised index finger. The trial continued when (1) the fingertip was within 1 cm of a point directly below the appropriate dowel and (2) the fingertip was lowered back to the

splint, where it remained for the rest of the trial. At this point, the visual feedback spot (a white circle of radius 3 mm) appeared at the location of the subject's fingertip or displaced to the left or right by 6 cm. Simultaneously, a red target of radius 5 mm appeared at one of the eight target locations. After a variable delay of 500-1500 msec, the target turned green, cueing the subject to begin the reach. Subjects were instructed to reach directly and accurately to the target. When the fingertip had moved 5 mm from its starting point, the feedback spot was extinguished, and the remainder of the reach was performed without visual feedback. The trial ended when the tangential fingertip velocity fell below 1.2 mm/sec. The target remained illuminated for the entire reach.

The experiment was comprised of five trial types (Figure 2D). In Left-Zero, Right-Zero, and Center-Zero trials, reaches were made from each of the three start points with no visual shift. In Center-Left and Center-Right trials, reaches from the center start point were made with 6 cm leftward and rightward visual shifts, respectively. Note that with these shifts, the feedback spot appeared at the locations of the left and right tactile start points. The experiment consisted of four reaches to each of the eight targets under these five conditions in a pseudorandom order, totaling $4 \times 8 \times 5 = 160$ reaches. In order to prevent subjects from adapting to the visual shifts, trials with left shifts, right shifts, and veridical feedback were pseudorandomly interleaved, and no two consecutive trials included the same shift. Adaptation was also unlikely since only two-fifths of the trials included shifts and because the visual feedback (shifted or veridical) was only available at the start point.

A set of 32 familiarization trials preceded the actual experiment. First, a block of 24 Left-Zero, Center-Zero, and Right-Zero trials was performed with continuous visual feedback to acclimate subjects to the task and experimental apparatus. Next, a block of 8 trials was performed with initial feedback only to familiarize subjects

with reaching in the absence of visual feedback. After the experiment was completed, subjects were asked whether they felt that the location of the visual feedback spot ever deviated from the location of their fingertip. All but one subject reported being unaware of any visual shift. The remaining subject (HA) reported that the location of the feedback spot seemed to have been displaced on a small number of trials (fewer than five).

Data Analysis and Model Fitting

Trajectory analysis. Arm position data were smoothed with a low-pass Butterworth filter with a cutoff frequency of 6 Hz, and the fingertip velocity and acceleration were successively computed using numerical differentiation (first differences). We quantified initial reach directions by determining the angle of the instantaneous velocity or acceleration vector at the point along the trajectory at which the tangential velocity first exceeded 40% of its peak value.

Modeling the initial movement direction. In building an explicit model of reach planning, we had to specify which extrinsic and intrinsic variables are used. Behavioral studies have variously suggested that reach planning uses either kinematic (Flash and Hogan, 1985; Atkeson and Hollerbach, 1985) or dynamic (Uno et al., 1989; Gordon et al., 1994a) variables, and neurophysiological findings have been cited to support both hypotheses (Cheney and Fetz, 1980; Georgopoulos et al., 1982; Todorov, 2000). We therefore fit our data twice, using the two models shown in Figure 3. In the velocity command model, the motor command is specified kinematically (as joint angle velocities), whereas in the torque command model the motor command is specified dynamically (as joint torques).

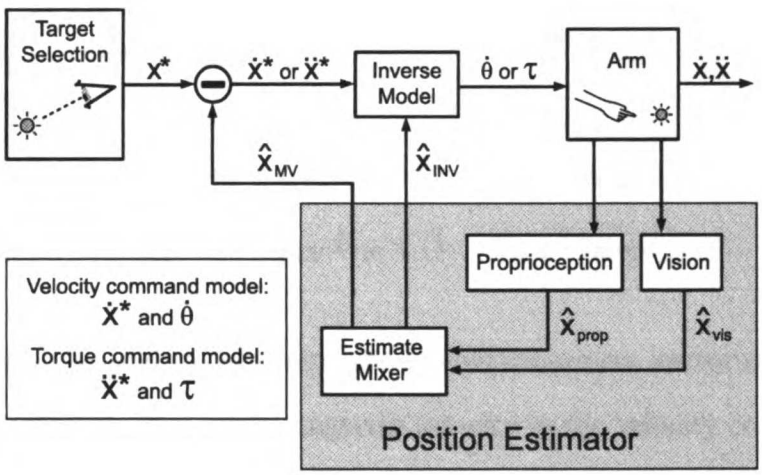

Figure 3: Two models of feedforward motor planning. Arm position estimates $\hat{x}_{\mathrm{MV}}$ and $\hat{x}_{\mathrm{INV}}$ are computed by combining visual and proprioceptive signals. A movement vector describing the desired direction of the initial velocity ($\dot{x}^*$) or acceleration ($\ddot{x}^*$) of the hand is computed by subtracting $\hat{x}_{\mathrm{MV}}$ from the target location x^*. An inverse model transforms this desired extrinsic vector into an intrinsic motor command specifying joint angle velocities ($\dot{\theta}$) or torques (τ), depending on the model being implemented. This transformation makes use of a second position estimate, $\hat{x}_{\mathrm{INV}}$. Finally, the motor command is executed, determining the initial hand trajectory. Note that the loop through the "Vision" and "Proprioception" boxes does not imply feedback control; in these models, the position estimates are only used to plan the plan the initial, feedforward component of a movement

The goal of these models is to understand how visual and proprioceptive signals from the sensory periphery combine to guide the initial, feedforward, component of the reach. In these models, only the initial velocities or accelerations of movements are computed, and feedback control is not modeled. We assume that the central nervous system weights the visual ($\hat{x}_{\text{vis}}$) and proprioceptive ($\hat{x}_{\text{prop}}$) position estimates and adds them to create two estimates of the position of the arm, $\hat{x}_{\text{MV}}$ ('movement vector') and $\hat{x}_{\text{INV}}$ ('inverse model'):

$$\hat{x}_{\text{MV}} = \alpha_{\text{MV}}\,\hat{x}_{\text{vis}} + (1-\alpha_{\text{MV}})\,\hat{x}_{\text{prop}} \quad (1)$$

$$\hat{x}_{\text{INV}} = \alpha_{\text{INV}}\,\hat{x}_{\text{vis}} + (1-\alpha_{\text{INV}})\,\hat{x}_{\text{prop}}. \quad (2)$$

The estimate $\hat{x}_{\text{MV}}$ is used to determine the desired movement vector in both models. This vector specifies desired initial fingertip velocity in the velocity command model and the desired initial fingertip acceleration in the torque command model. The second estimate, $\hat{x}_{\text{INV}}$, is used to convert the desired movement into an intrinsic motor command. This command is expressed as joint velocities in the velocity command model and as joint torques in the torque command model. In both models, therefore, the weighting parameters α_{MV} and α_{INV} characterize sensory integration at the two stages of reach planning.

Velocity Command Model. In the velocity command model, the planned movement vector is defined as the desired initial velocity of the fingertip ($\dot{x}^*$). The direction of this velocity is specified by

$$\measuredangle \dot{x}^* = \measuredangle(x_d^* - \hat{x}_{\text{MV}}) + \omega_d, \quad (3)$$

where $\measuredangle x$ represents the angle of vector x, x_d^* represents the location of target $d \in$

$[1, ..., 8]$, $\hat{x}_{\mathrm{MV}}$ is the estimated hand position defined in Equation 1, and ω_d is an angular offset from the straight line connecting the estimated initial position and target x_d^*. We included the ω_d terms to account for the fact that natural, unperturbed reaching movements are slightly curved (Soechting and Lacquaniti, 1981; Atkeson and Hollerbach, 1985; Uno et al., 1989), resulting in initial reach directions that differ from the target direction. This baseline bias was fit from the Center-Zero trials: each ω_d was set to the average angular difference between the initial velocity direction and the target direction for target x_d^*. Equation 3 does not specify the magnitude of $\dot{x}^*$, as ultimately only the predicted direction of movement will be compared to the data. Note that the pattern of errors in $\angle\dot{x}^*$ resulting from misestimation of $\hat{x}_{\mathrm{MV}}$ is the MV error shown in Figure 1A and B.

Given a desired Cartesian fingertip velocity $\dot{x}^*$, the *ideal* joint angle velocity command would be:

$$\dot{\theta}_{\mathrm{ideal}} = J^{-1}(\theta)\,\dot{x}^*$$

where the Jacobian matrix $J(\theta)$ is the gradient of the fingertip location with respect to the joint angles:

$$J(\theta) = \nabla_\theta K(\theta) = \frac{dx}{d\theta}$$

The kinematics equation $x = K(\theta)$ describes the mapping from joint angles to fingertip locations. Note that since the arm is restricted to planar, two-joint movements, both the kinematics and the Jacobian are invertible. Since the internal inverse model must rely on an estimate of the position of the arm ($\hat{x}_{\mathrm{INV}}$), the implemented motor command will be

$$\begin{aligned} \dot{\theta} &= J^{-1}(\hat{\theta}_{\mathrm{INV}})\,\dot{x}^* \qquad (4) \\ \text{where} \quad \hat{\theta}_{\mathrm{INV}} &= K^{-1}(\hat{x}_{\mathrm{INV}}) \end{aligned}$$

Finally, this joint velocity command is executed, and the arm moves with an initial fingertip velocity determined by the Jacobian (evaluated at the true arm position):

$$\dot{x} = J(\theta)\,\dot{\theta} = J(\theta)\,J^{-1}(\hat{\theta}_{\mathrm{INV}})\,\dot{x}^{*} \tag{5}$$

This model predicts that the initial velocity $\dot{x}$ will be distorted from the desired velocity $\dot{x}^{*}$ if the arm position is misestimated. The matrix $D_{\mathrm{vel}}(\theta, \hat{\theta}_{\mathrm{INV}}) = J(\theta)\,J^{-1}(\hat{\theta}_{\mathrm{INV}})$, which we will call the velocity distortion matrix, determines the INV error (Figure 1C and D) in the velocity command model.

Torque Command Model. In addition to the velocity command model, which assumes that reaches are planned in kinematic coordinates, we also considered a model in which the dynamics of the movement are controlled via the joint torques τ. In the torque command model, the movement vector specifies a desired initial acceleration, $\ddot{x}^{*}$, which is offset by some angle ω_d from the target direction:

$$\measuredangle \ddot{x}^{*} = \measuredangle(x_d^{*} - \hat{x}_{\mathrm{MV}}) + \omega_d \tag{6}$$

The ω_d in this model are determined by measuring the average initial accelerations for reaches to each target in the baseline (Center-Zero) trials.

The ideal joint torque command can be computed from the desired acceleration as follows. The relationship between joint and endpoint acceleration is found by differentiating Equation 5 with respect to time:

$$\begin{aligned} \ddot{x} &= J(\theta)\ddot{\theta} + \frac{d}{dt}\left[J(\theta)\right]\dot{\theta} \\ \ddot{x} &\approx J(\theta)\ddot{\theta} \end{aligned} \tag{7}$$

The approximation in Equation 7 follows from the fact that we are only considering the initial component of the movement, when the magnitude of the angular velocity is small. The relationship between the joint torques and the kinematic variables of the movement is given by the dynamics equations for the planar, two-joint, arm:

$$\begin{aligned} \tau &= I(\theta)\,\ddot{\theta} + H(\theta,\dot{\theta})\,\dot{\theta} \\ \tau &\approx I(\theta)\,\ddot{\theta} \end{aligned} \tag{8}$$

where I is the position-dependent, non-isotropic inertia of the arm, and the H term represents velocity-dependent centripetal forces, joint interaction torques, and damping forces at the joints. Since this latter term varies linearly with respect to joint velocity, it is small at the onset of movement, yielding the approximation of Equation 8. Inverting Equation 7 and combining it with Equation 8, we obtain the ideal joint torque command,

$$\tau_{\text{ideal}} \approx I(\theta)\,J^{-1}(\theta)\,\ddot{x}^{*} \tag{9}$$

However, the true value of the arm position is not available to the nervous system, which must make use of the estimated joint angles, $\hat{\theta}_{\text{INV}}$ when computing the inverse model:

$$\tau \approx I(\hat{\theta}_{\text{INV}})\,J^{-1}(\hat{\theta}_{\text{INV}})\,\ddot{x}^{*}$$

Finally, we can invert Equation 9 to determine the fingertip acceleration that results from a given joint torque command:

$$\ddot{x} \approx J(\theta)\,I^{-1}(\theta)\,\tau$$

Combining the two previous equations, we arrive at an expression for the distortion

of the planned fingertip acceleration:

$$\ddot{x} \approx J(\theta)\, I^{-1}(\theta)\, I(\hat{\theta}_{\mathrm{INV}})\, J^{-1}(\hat{\theta}_{\mathrm{INV}})\, \ddot{x}^{*} \tag{10}$$

The resulting acceleration distortion matrix is given by

$$D_{\mathrm{acc}}(\theta, \hat{\theta}_{\mathrm{INV}}) = J(\theta)\, I^{-1}(\theta)\, I(\hat{\theta}_{\mathrm{INV}})\, J^{-1}(\hat{\theta}_{\mathrm{INV}}),$$

where $I(\theta)$ is the inertia matrix of the arm. The desired acceleration ($\ddot{x}^{*}$) will only be achieved if the arm position estimate $\hat{\theta}_{\mathrm{INV}}$ is correct. Note that Equations 8-10 represent only the instantaneous, initial dynamics of a rigid body model of the arm. Our intention in employing this simplified model is only to show that our results are robust to a consideration of the arm's inertia, and are therefore not dependent on the purely kinematic analysis performed with the velocity command model. In order to fit the torque command model to the data, we used previously published estimations of the inertia matrix $I(\theta)$ (Sabes and Jordan, 1997).

Generating quantitative model predictions. In order to generate quantitative model predictions for comparison to our data, the model target locations (x_d^{*}) were set to the locations of the visual targets, and the arm position variables (x and θ) were set to the measured pre-movement values. Since it is not possible to measure $\hat{x}_{\mathrm{prop}}$ and $\hat{x}_{\mathrm{vis}}$ directly, we assumed that vision and proprioception were unbiased, i.e. that $\hat{x}_{\mathrm{prop}} = x$ and that $\hat{x}_{\mathrm{vis}}$ corresponded to the location of the visual feedback spot. We will consider the possible consequences of sensory biases in the Results and Appendix.

To illustrate the errors predicted by the velocity command model, we used the model to simulate the effects of the shifts used in the actual experiment (Figure

4). In these simulations, $\hat{x}_{\text{prop}}$ and x were set to the location of the center start point, $\hat{x}_{\text{vis}}$ was placed 6 cm to the left or right of x, and the ω_d terms were set to zero. The mean arm length across subjects was used, and various values of α_{MV} and α_{INV} were chosen to demonstrate the influence of these mixing parameters on the predicted errors.

Fitting model predictions to the data. For each subject and model, the values of α_{MV} and α_{INV} were simultaneously fit to a single dataset consisting of all of the Center-Zero, Center-Right, and Center-Left trials for each subject (96 trials total). The weighting parameters α_{MV} and α_{INV} were fit to minimize the squared error between the model predictions and the measured initial movement directions using a general purpose nonlinear regression algorithm (nlinfit in MATLAB, The Mathworks Inc., Natick, MA). Note that only the directions, and not the magnitudes of the initial velocities (velocity command model) or accelerations (torque command model) were compared to the model predictions. Since the parameter space was only two-dimensional, we were able to plot the error surface over a reasonable range of parameter values. These plots were smooth, and no local minima were observed (data not shown), confirming our observation that the fit values did not depend on the initial conditions used in the optimization.

Hypothesis testing and confidence limits. To test the hypothesis that a given position estimate relies on signals from a certain modality, we employed permutation tests (Good, 2000) against the null hypothesis that the estimate relies exclusively on the other modality. First consider a test for whether $\hat{x}_{\text{MV}}$ makes use of visual information. The null hypothesis is $H_0 : \alpha_{\text{MV}} = 0$, i.e. that only proprioception is

used. Rearrangement of Equation 1 yields

$$\hat{x}_{\mathrm{MV}} = \hat{x}_{\mathrm{prop}} + \alpha_{\mathrm{MV}}(\hat{x}_{\mathrm{vis}} - \hat{x}_{\mathrm{prop}}) \qquad (1')$$

The null hypothesis states that $\hat{x}_{\mathrm{MV}}$ is independent of $(\hat{x}_{\mathrm{vis}} - \hat{x}_{\mathrm{prop}})$ in Equation 1'. By substituting Equation 1' for Equation 1 in the models and permuting the trials from which this difference is taken, we broke any existing dependence of $\hat{x}_{\mathrm{MV}}$ on on $\hat{x}_{\mathrm{vis}}$, thereby constructing synthetic datasets which obeyed H_0. By creating 1000 such datasets and fitting α_{MV} to each of them, we created a distribution of synthetic α_{MV} under H_0, which were typically centered around $\alpha_{\mathrm{MV}} = 0$. We rejected H_0 if the α_{MV} fit to the true (unpermuted) dataset was greater than the 95^{th} percentile of the synthetic distribution. We tested the null hypothesis $H_0 : \alpha_{\mathrm{INV}} = 0$ in the same fashion.

Next, we tested whether $\hat{x}_{\mathrm{MV}}$ makes use of proprioceptive information. In this case the null hypothesis is $H_0 : \alpha_{\mathrm{MV}} = 1$, i.e. that only vision is used. A different rearrangement of Equation 1 yields

$$\hat{x}_{\mathrm{MV}} = \hat{x}_{\mathrm{vis}} + (1 - \alpha_{\mathrm{MV}})(\hat{x}_{\mathrm{prop}} - \hat{x}_{\mathrm{vis}}) \qquad (1'')$$

If we define $\beta_{\mathrm{MV}} = (1 - \alpha_{\mathrm{MV}})$ then the null hypothesis can be written $H_0 : \beta_{\mathrm{MV}} = 0$. By substituting Equation 1'' for Equation 1 in the models and permuting the trials from which the difference $(\hat{x}_{\mathrm{prop}} - \hat{x}_{\mathrm{vis}})$ is taken, we broke any existing dependence of $\hat{x}_{\mathrm{MV}}$ on $\hat{x}_{\mathrm{prop}}$, thereby constructing synthetic datasets which obeyed H_0. Using the permutation methods described in the previous paragraph, we then tested whether β_{MV} from the unpermuted dataset was greater than the 95^{th} percentile of the synthetic distribution of β_{MV}. We tested the null hypothesis $H_0 : \alpha_{\mathrm{INV}} = 1$ in the same

fashion.

Finally, we tested whether there was a difference in the relative weighting of vision and proprioception between the two position estimates. To accomplish this, we performed a permutation test comparing each model to a simplified version of itself in which only a single weighting of vision and proprioception is used, $H_0 : \alpha_{MV} = \alpha_{INV}$ and $\hat{x}_{MV} = \hat{x}_{INV}$. This test was implemented by replacing α_{MV} and α_{INV} with a common part α_{comm} and a difference α_{diff}:

$$\alpha_{comm} = \alpha_{MV}$$

$$\alpha_{diff} = \alpha_{INV} - \alpha_{MV}$$

Applying these definitions to Equations 1 and 2, we obtain the following:

$$\hat{x}_{MV} = \hat{x}_{prop} + \alpha_{comm}(\hat{x}_{vis} - \hat{x}_{prop}) \quad (1''')$$

$$\hat{x}_{INV} = \hat{x}_{prop} + \alpha_{comm}(\hat{x}_{vis} - \hat{x}_{prop}) + \alpha_{diff}(\hat{x}_{vis} - \hat{x}_{prop}) \quad (2''')$$

Under H_0, there is no difference between the two original weighting parameters, so $\alpha_{diff} = 0$. This means that in Equation 2''' there would be no dependence on $(\hat{x}_{vis} - \hat{x}_{prop})$ beyond that accounted for in the α_{comm} term. Due to normal statistical variation, however, inclusion of α_{diff} in the model would still improve the fit. We therefore compared the best-fit value of α_{diff} from the real dataset to values obtained from 1000 synthetic datasets in which we permuted the trials from which $(\hat{x}_{vis} - \hat{x}_{prop})$ were taken for the α_{diff} term in Equation 2'''. For the α_{comm} terms, the true values of $(\hat{x}_{vis} - \hat{x}_{prop})$ were used. If the absolute value of the α_{diff} fit to the true data was greater than the 95^{th} percentile obtained from the synthetic datasets, we inferred that the α_{diff} term reflects a real difference between α_{MV} and α_{INV}, and we

rejected H_0. Additionally, we performed a more standard F-test of the "extra sums of squares" obtained by including the second mixing parameter (α_{diff}) in the model (Draper and Smith, 1998).

In order to put confidence limits on the fit values of α_{MV} and α_{INV}, we employed a bootstrapping technique (Efron and Tibshirani, 1993). For each subject, we created 1000 datasets in which the data from every trial were resampled (with replacement) from one of the four trials of the same type and with the same target. The parameters α_{MV} and α_{INV} were then fit to each resampled dataset. The resulting distribution was used to find the confidence ellipses for the fit parameter vectors [α_{MV}, α_{INV}] for that subject.

Results

Errors in initial reach direction. The velocity and torque command models predict the errors in initial movement direction for given values of the weighting parameters α_{MV} and α_{INV}. Figure 4 shows the predictions made by the velocity command model for four sets of parameter values. Similar error patterns are predicted by the torque command model.

Figure 5 shows a typical subject's reach trajectories for trials beginning at the center start point with a leftward visual shift (panel A), no shift (panel B), and a rightward shift (panel C). The shift-induced changes in movement direction were opposite in sign for the two visual shifts. Initial velocity directions for each of these movements are shown in Figure 5D. As was typical, this subject displayed directional biases in the unshifted condition. The ω_d (Equation 3, Methods) were set to the mean of these biases for each direction (Figure 5D, dotted line). Figure 5E shows the

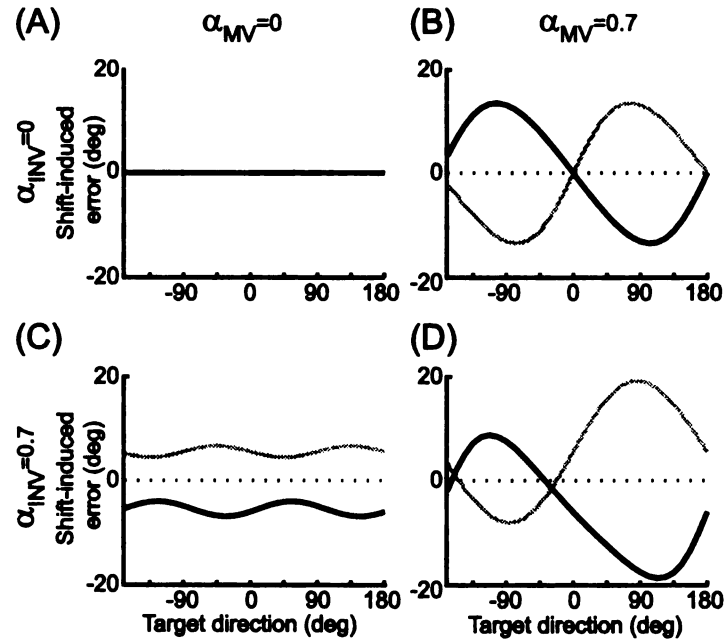

Figure 4: Shift-induced errors in initial reach direction predicted by the velocity command model with various values of α_{MV} and α_{INV}. Each plot shows the predicted errors in initial velocity direction as a function of target direction. Left column: $\alpha_{MV} = 0$; right column: $\alpha_{MV} = .7$; top row: $\alpha_{INV} = 0$; bottom row: $\alpha_{INV} = .7$. Black lines: leftward shift; gray lines: rightward shift. Positive values on the ordinate correspond to CCW errors. Note that (B) shows the effects of MV error alone, (C) shows the effects of INV error alone, and (D) shows their combined effects when both $\hat{x}_{MV}$ and $\hat{x}_{INV}$ are shifted.

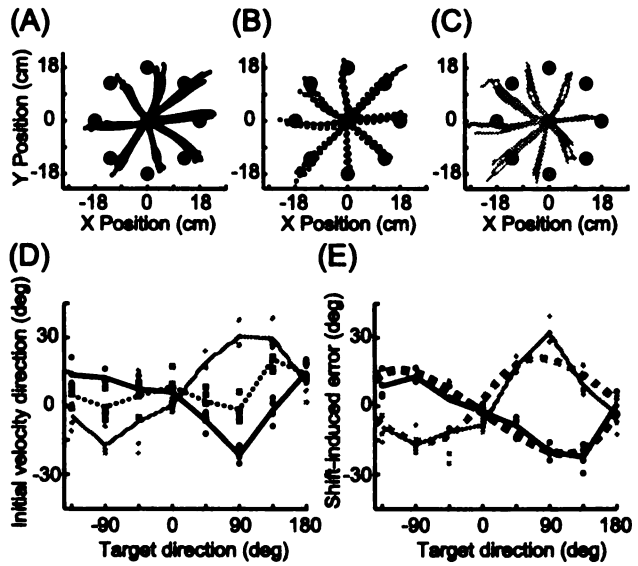

Figure 5: Data and velocity-command model fit from one subject. Movement paths from all Center-Left (A), Center-Zero (B), and Center-Right (C) trials. (D) Initial velocity direction (with respect to target direction) as a function of target direction for Center-Left (•, individual trials; solid black line, mean), Center-Zero (■, dotted line), and Center-Right (•, gray line) trials. (E) Shift-induced error as function of target direction. Dashed lines represent the errors predicted by the best-fit velocity-command model ($\alpha_{\mathrm{MV}} = .97$, $\alpha_{\mathrm{INV}} = .34$). Other symbols as in panel D. Data are from subject HA.

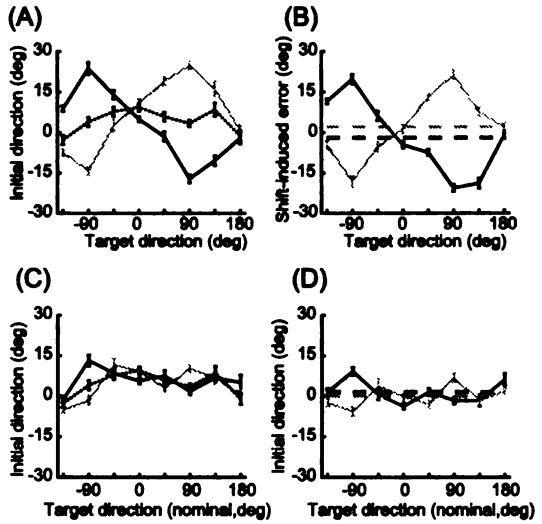

Figure 6: Initial velocity data averaged across all subjects. (A) Initial velocity direction (with respect to target direction) for Center-Left (solid black line), Center-Zero (dotted line), and Center-Right (gray line) trials. (B) Shift-induced errors in initial velocity direction (line colors as in A). (C) Initial velocity direction (with respect to target direction) for Left-Zero (solid black line), Center-Zero (dotted line) and Right-Zero (gray line) trials. (D) Data from (C) after subtraction of the mean Center-Zero directions. Line colors as in (C). Target directions in (C) and (D) are relative to the center start point for ease of comparison. Errorbars in all plots are ±1 s.e. Dashed lines in (B) and (D) indicate means for a given dataset.

velocity command model fit to this subject's shift-induced reach errors, which were computed by subtracting the appropriate ω_d from the Center-Left and Center-Right initial reach directions. The model captures the main features of the observed error pattern ($R^2 = 0.73$). The fit values of the weighting parameters were $\alpha_{\mathrm{MV}} = .97$ and $\alpha_{\mathrm{INV}} = .34$. This suggests that when planning a movement vector this subject relied almost entirely on vision to estimate the position of the hand. In contrast, when computing how this vector should be transformed into a motor command, the subject used a mixed estimate that was 34% visual and 66% proprioceptive. Consistent with these fit values, the data and model fit seen in Figure 5E show an error pattern intermediate between those shown in Figure 4B and D.

Initial velocity data averaged across all subjects are shown in Figure 6. Baseline

directional biases (dotted lines in panels A and C) varied from subject to subject, but were always within 20° of the target direction (mean ± s.e., 7.3° ± 2.0). All but one subject showed significant variation in the values of the baseline bias across target directions (ANOVA, $p < .05$). The shift-induced errors in initial velocity direction are shown in Figure 6B. To highlight the effects of INV error, the mean errors across targets for the two shifts are shown as dashed lines. The separation between these means reflects the rotational (CW-CCW) shifts typical of INV error (see Figure 4C and D).

We examined the initial directions of movements made from the left and right start points in the absence of visual shifts in order to test whether the shift-induced errors were simply due to changes in the visually perceived start point, rather than to misestimation of arm position. Figure 6C and D show that this is not the case. The bimodal pattern reflecting the MV error is absent, and these control data do not show the pattern of CW-CCW shifts seen in the shifted trials (compare the dashed lines in Figure 6B and D).

The initial direction data presented in Figures 5 and 6 were sampled from the point in the reach trajectory at which the tangential velocity first exceeded 40% of its peak value (see Methods). This landmark occurred 125 ± 24 msec after reach onset (mean ± average within-subject s.d.), and nearly always fell within the first centimeter of the reach. Since feedback signals are able to influence reach trajectories starting at about 150 msec (Prablanc and Martin, 1992; Paillard, 1996), it is possible that on some trials the velocity and acceleration at the time of the measurement were influenced by sensory feedback of the earliest portions of the reach. This might be a cause for concern as our models are strictly feedforward. However, using an earlier landmark (the point at which tangential velocity exceeds 20% of peak, which falls 68 ± 15 msec into the reach) yielded nearly identical average values of initial

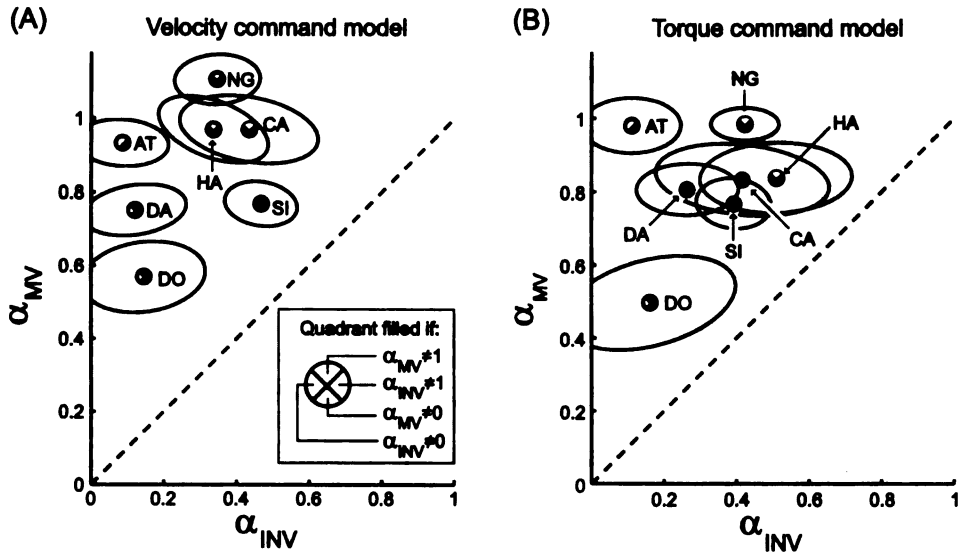

Figure 7: Best-fit values of α_{MV} and α_{INV} for all subjects and both models. Each symbol is divided into quadrants that are filled or empty depending on the results of the specified hypothesis tests ($p < .05$, see inset). Ellipses represent one standard error (bootstrap analysis, see Methods). Two-letter labels identify individual subjects. The dashed line represents $\alpha_{INV} = \alpha_{MV}$.

direction and produced model fits that were not significantly different from those obtained using the 40% criterion (not shown). Because the tangential velocity was very small at the 20% landmark, though, measurements of velocity direction taken at this landmark were significantly more noisy than those taken at the 40% landmark. For this reason, we elected to use the 40% criterion in all of our analyses.

Weighting parameters α_{MV} and α_{INV}. The errors in initial direction were used to fit both the velocity command model and the torque command model. The fit values of α_{MV} and α_{INV} for all subjects are shown in Figure 7. The average values of $[\alpha_{MV}, \alpha_{INV}]$ across subjects were [.87,.28] for the velocity command model and [.82, .33] for the torque command model. The values of α_{MV} indicate that the position estimate used for movement vector planning relied predominately on vision. For every subject and both models, α_{MV} was greater than 0.5, and in every

case we could reject the null hypothesis that $\hat{x}_{\mathrm{MV}}$ relied solely on proprioception ($H_0: \alpha_{\mathrm{MV}} = 0$). The values of α_{INV} suggest that the position estimate used for converting a movement vector into a motor command relied more on proprioception. In all but one case, the fit values of α_{INV} were less than 0.5 (the exception was subject HA, torque command model $\alpha_{\mathrm{INV}} = .51$), and in every case the null hypothesis that $\hat{x}_{\mathrm{INV}}$ relied solely on vision ($H_0: \alpha_{\mathrm{INV}} = 1$) was rejected. Despite these strong biases toward vision and proprioception, however, both position estimates appear to rely on a mixture of sensory inputs: in 7 of 14 cases we could reject the null hypothesis that $\hat{x}_{\mathrm{MV}}$ was purely visual ($H_0: \alpha_{\mathrm{MV}} = 1$), and in 9 of 14 cases we could reject the null hypothesis that $\hat{x}_{\mathrm{INV}}$ was purely proprioceptive ($H_0: \alpha_{\mathrm{INV}} = 0$).

The difference between the fit values of the two weighting parameters suggests that the two position estimates $\hat{x}_{\mathrm{MV}}$ and $\hat{x}_{\mathrm{INV}}$ are indeed distinct quantities. We examined this hypothesis by testing whether the fit values of α_{MV} and α_{INV} differed significantly from each other (see Methods). In 12 out of 14 cases, the permutation test allowed us to reject the null hypothesis that the two parameters were equal (the exceptions were subjects DO and HA, torque command model). These results were confirmed by an F-test ($p < .05$), which agreed with the permutation test in all but a single case (subject CA, torque command model). In the majority of cases, therefore, the two position estimates relied on different combinations of sensory input, reflecting a significant difference between multisensory integration at the two stages of reach planning proposed by our models.

Both the velocity command and torque command models fit the observed data well, and neither performed consistently better across subjects (R^2 values ranged from .63 to .80 for the velocity command model and from .45 to .80 for the torque command model). This similarity suggests that the choice of controlled variable (joint velocities or joint torques) in the model is not critical, and that our results fol-

low from the assumption of a two-stage planning process in which a desired extrinsic movement vector is computed and then converted into an intrinsic motor command. Additionally, our assumption that visual and proprioceptive signals are additively combined in Cartesian space (Equations 1 and 2) did not influence our conclusions. We found nearly identical values of α_{MV} and α_{INV} (all absolute differences less than .003) when we refit the data with an alternate model in which visual and proprioceptive cues were combined in joint angle coordinates, implemented by substituting a $\hat{\theta}$ for each $\hat{x}$ in Equations 1 and 2.

In our models, the position estimates $\hat{x}_{MV}$ and $\hat{x}_{INV}$ are weighted sums of $\hat{x}_{vis}$ and $\hat{x}_{prop}$ (Equations 1 and 2). In other words, we have assumed that each combined position estimate lies on the line that connects the two unimodal estimates and that the distance along that line is determined by the parameter α_{MV} or α_{INV}. However, van Beers et al. (1999) have argued that since individual sensory modalities are more or less reliable along different spatial axes, a simple scalar weighting of two unimodal estimates may not produce the statistically optimal combination of these signals. These authors supported their argument by showing that in some conditions the integrated estimate of arm position lies off the straight line connecting the visual and proprioceptive estimates. Such a finding suggests that $\hat{x}_{MV}$ and $\hat{x}_{INV}$ might vary across the two-dimensional horizontal plane, and that a weighted-sum model might be insufficient. To address this issue, we fit our data with a second alternate model in which $\hat{x}_{MV}$ and $\hat{x}_{INV}$ were free to vary across the horizontal plane. Despite this freedom, the best-fit $\hat{x}_{MV}$ and $\hat{x}_{INV}$ still lay near the line connecting the unimodal estimates and there was no consistent component perpendicular to that line (data not shown). Furthermore, the component along the line agreed with the fits shown in Figure 7. These results validate the original weighted-sum model of Equations 1 and 2 for our data. Note however, that these findings do not necessarily

contradict the model of van Beers et al. (1999), since our study was conducted in a different part of the workspace, and workspace location has been shown to influence the orientations of the unimodal covariance ellipses (van Beers et al., 1998).

In order to fit our models to the data, we have also assumed that the visual and proprioceptive position estimates are unbiased. The analysis described in the preceding paragraph suggests that if any biases exist, they lie principally along the axis parallel to the feedback shift. In fact, such biases could arise for two different reasons. First, the unimodal estimates may be inherently biased, so that $\hat{x}_{\text{vis}}$ and $\hat{x}_{\text{prop}}$ might differ from the locations of the feedback spot and the fingertip, respectively. A second source of bias could arise in the internal transformations of the visual and proprioceptive signals required at each planning stage. For example, comparing $\hat{x}_{\text{prop}}$ to the target location might require computing the Cartesian fingertip location from the proprioceptive signal, while evaluating the inverse model might require a joint-based representation. If these transformations were biased, the true value of $\hat{x}_{\text{prop}}$ may not be the same in Equations 1 and 2. In the Appendix, we show that neither of these types of bias would significantly affect the fit values of α_{MV} and α_{INV}.

Position-dependent changes in the distortion matrix. The empirical measurements of the direction of the initial velocity and acceleration had to be taken after the onset of the reach, at which point the fingertip had moved a small distance from its initial position. On the other hand, when fitting the models to the data, the model distortion matrices were evaluated at the initial position of the fingertip. This simplification would have a negligible effect if the distortion matrix were nearly constant over the initial movement segment. However, if the distortion matrix varied rapidly across the workspace, there would be a marked change in INV error between

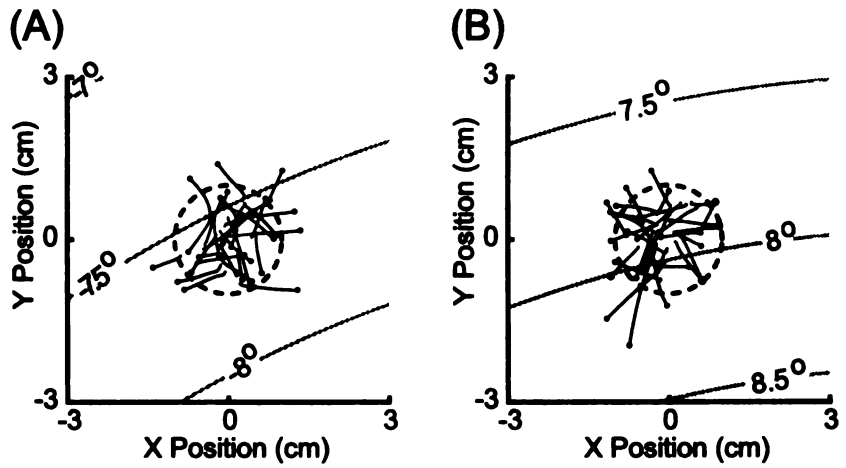

Figure 8: Initial reach segments and predicted INV error. The initial portion of reaches (unfiltered empirical data) are shown for all trials from subject HA in the (A) Center-Left and (B) Center-Right conditions. Black lines represent the path from movement onset to the point at which the tangential velocity first exceeds 40% of the peak velocity (black dots). The large circle (dashed line) represents the center start point window (radius 1 cm). Gray contour lines show the magnitude of the predicted INV errors (velocity command model) due to the distortion matrix $D_{\text{vel}}(\theta, \hat{\theta}_{\text{INV}})$ as a function of arm position, assuming $\alpha_{\text{INV}} = 1$. Positive contour values correspond to CCW errors, negative values signify CW errors.

the initial arm position and the location at which the velocity and acceleration measurements were made. We assessed whether the assumption of a constant distortion matrix significantly affected our results by determining how the predicted INV error in the velocity command model varies over the initial segment of the trajectory (Figure 8). For a given arm position, the INV error (that is, the error introduced via the distortion matrix) depends on three variables: the true arm position, the error in the estimated arm position, and the desired movement direction. For this analysis, we assumed that the arm position estimate was equal to the location of the visual feedback ($\alpha_{\text{INV}} = 1$). This was the conservative choice, as it maximizes the INV error. We also averaged the predicted error over the eight target directions, as the INV error varies little over the desired movement direction (Figure 4C). We then calculated the predicted INV error for each shift direction as a function of the arm's

position for a representative subject and made a contour plot of the results (Figure 8, grey lines). Superimposed on these plots are the initial segments of the same subject's reach trajectories, ending at the point where the velocity and acceleration were measured. Note that the predicted INV error typically varied less than 0.5° over the initial movement segment. In contrast, for this subject the direction of the initial velocity had a within-condition standard deviation of 5.73°. Therefore, any error in model prediction stemming from the assumption of a stationary distortion matrix would be lost in the inherent movement variability.

Magnitude of initial velocity. In our models, α_{MV} and α_{INV} were fit using only the directional component of the error in initial reach velocity or acceleration. Here, we show that the shifts in visual feedback also lead to errors in the magnitude of the initial velocity, and that these errors are consistent with the predictions of the velocity command model. The results for initial acceleration and the torque command model are qualitatively the same.

First consider that both the direction and magnitude of the INV velocity error are determined by the distortion matrix, $D_{\text{vel}}(\theta, \hat{\theta}_{\text{INV}})$, as shown in Equation 5. For the starting location and visual shift directions used in this experiment, the distortion matrices are mostly rotational, i.e. the desired velocity undergoes a rotation but very little scaling. We determined this by evaluating the velocity distortion matrix at the starting location for each subject for both the left and right visual shifts using the best-fit α_{INV}. We then found the velocity direction that yielded the greatest absolute percentage change in the magnitude of the velocity. Across subjects and shift directions, the average maximum scaling was only 0.75% ± 0.44% (mean ± 1 s.d.) of the original length. The model therefore predicts that in our experiment, the INV error should have a negligible effect on the magnitude of the velocity.

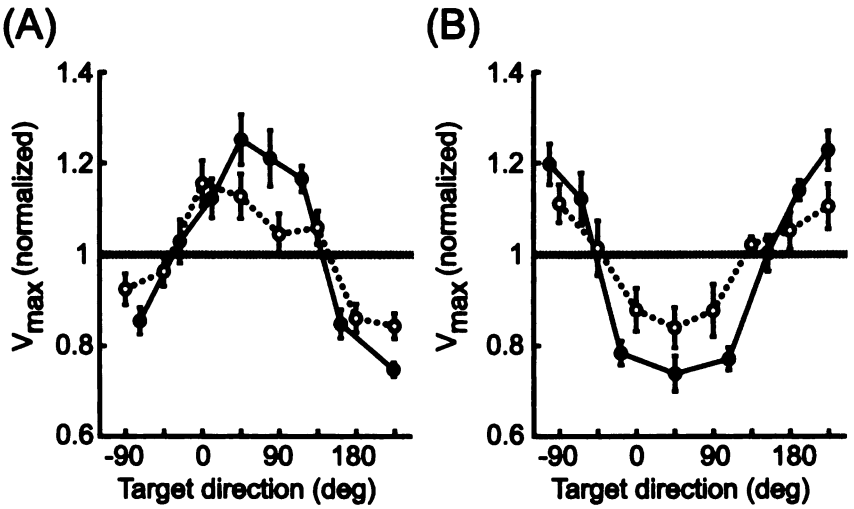

Figure 9: Effects of visual shifts on reach velocity magnitude. Each line plots the mean peak tangential velocity (across subjects) for each target normalized to the mean peak tangential velocity in Center-Zero trials. Target directions are defined as the direction from the fingertip start position to the visual target. (A) Left-Zero (solid line) and Center-Left (dotted line) trials. (B) Right-Zero (solid line) and Center-Right (dotted line) trials. Errorbars are ± 1 S.E.

In contrast, MV error alters both the direction and length of the planned movement vector, and we would expect these variables to influence the magnitude of the planned velocity. This effect can be understood by considering the movements made from the left and right start points in the absence of a visual shift. For each subject and target, the peak velocities for reaches in the Left-Zero and Right-Zero conditions were normalized to the average peak velocity for the same target in the Center-Zero condition. These values are plotted as a function of target direction in Figure 9 (filled symbols). This is the pattern that would be expected in the Center-Left and Center-Right conditions if the estimated arm positions were located at the left and right starting locations, respectively (i.e. if $\alpha_{MV} = 1$). If there were no error in the position estimate ($\alpha_{MV} = 0$), the average peak velocity would be the same as in the Center-Zero condition, and so the normalized values would all be near unity.

As can be seen from the open symbols in Figure 9, the dependence of peak velocity on target direction in the Center-Left and Center-Right trials has the same shape as that seen in the Left-Zero and Right-Zero conditions, but the effect is smaller in size. Qualitatively, this is the pattern of MV errors that the velocity command model would predict for the best-fit values of α_{MV}, which are less than one.

Discussion

By dissociating visual and proprioceptive feedback, we induced errors in the arm position estimates used at two different stages of reach planning. We then modeled the errors in initial movement direction as a function of the misestimation at each stage. Comparison of our models with the experimental data allowed us to quantify these position errors and thus compute the extent to which each planning stage relies on visual feedback. We found that the arm position estimate used for vector planning relies mostly on visual feedback, whereas the estimate used to convert the desired movement vector into a motor command relies more on proprioceptive signals.

This finding agrees with the results of a recent study by Sainburg and colleagues that employed a reaching task very similar to our own (Sainburg et al., 2003). In their experiment, the location of the initial visual feedback was constant across trials while the actual initial location of the fingertip was varied. The authors found that these manipulations did not affect reach direction, suggesting that subjects relied heavily on the visual position signal (which did not change location) when planning movement vectors. Furthermore, an inverse dynamic analysis revealed that subjects generated intrinsic motor commands that took into account the true position of the arm, suggesting that subjects relied mostly on proprioceptive signals when generating motor commands. Although the relative weightings of vision and proprioception were

not quantified, these results are in agreement with our own.

Psychophysical studies of the tradeoff between vision and proprioception in arm position estimation have suggested that each modality is weighted according to its statistical reliability (Howard and Templeton, 1966; Welch and Warren, 1980; van Beers et al., 1999) or depending on the focus of the subject's attention (Warren and Schmitt, 1978; Welch and Warren, 1980). However, while both of these factors may influence multisensory integration, these models provide only a single criterion for weighting the unimodal signals. Such models cannot account for our finding that vision and proprioception are weighted differently at different stages of motor planning.

Our results instead suggest that multisensory integration depends on the computations in which the integrated estimates are used. In order to compute the movement vector, for example, the position of the hand must be compared to that of the target in a common coordinate frame. Transforming signals from one coordinate frame to another presumably incurs errors, either from imperfections in the mapping between them (bias) or due to noise introduced in the additional computation (variance). The effects of these errors on movement control can be reduced by giving less weight to transformed signals. In our experiment, where targets are presented visually, the increased reliance on visual feedback when planning movement vectors would therefore have been advantageous.

A similar argument can explain the predominance of proprioception when transforming the movement vector into a motor command. Computing the inverse model of the arm requires knowledge of the arm's posture. Although our experimental constraints created a one-to-one relationship between joint angles and fingertip location, during more natural, unconstrained movements joint angles cannot be uniquely inferred from visual feedback specifying only the location of the fingertip. Addition-

ally, errors can arise from biases in the coordinate transformation from extrinsic to intrinsic coordinates (Soechting and Flanders, 1989a) and from variance introduced during the computation, as in the first stage of planning. Because of these factors, the reduced reliance on vision at this second planning stage would have been advantageous.

This interpretation is compatible with a modified minimum-variance principle that takes into account the errors introduced by coordinate transformations. The computations performed at each stage require information about different aspects of the position of the arm: when reaching to a visual target, movement vector planning requires only the extrinsic location of the fingertip, whereas computing the inverse model of the arm requires knowing the intrinsic, joint-based posture of the arm. The values of α_{MV} and α_{INV} reflect this difference, since the nervous system relies more heavily on the signals which contain the information necessary to perform the relevant computation and do not need to be transformed.

These conclusions are based primarily on analyses of initial movement direction. However, numerous authors have argued that the planning of reach direction and extent are independent processes (Gordon et al., 1994b; Messier and Kalaska, 1997). If this hypothesis were true, the rules for integrating vision and proprioception might be different for the planning of movement extent and direction. Indeed, such a difference was found in the recent study by Sainburg and colleagues (Sainburg et al., 2003). As noted above, they found that the direction of movements to a given target depended on the location of the visual feedback, and not on the actual position of the arm. In contrast, planning of movement extent appeared to depend on the arm's true position to a greater or lesser degree, depending on the position of the arm relative to the target. Our results do not rule out the possibility that a separate estimate or set of estimates is used to compute movement extent. Nonetheless, we

have shown in Figure 9 that the peak velocity relies on a position estimate located between the positions specified by vision and proprioception, consistent with the results of our analyses of movement direction. Although these data suggest that the planning of reach amplitude and direction might use the same position estimates, this hypothesis would have to be confirmed by a study that better controlled for the various factors affecting reach amplitude.

Our quantification of sensory integration relies on model-based analyses of the empirical data. We must therefore address how sensitive our conclusions are to the details of the model. First, the velocity command and torque command models produce similar estimates of multisensory integration at each planning stage (Figure 7). This shows that our results do not depend critically on the assumption that the nervous system specifies kinematic or dynamic motor commands. Second, the alternate model in which the two signals are combined in intrinsic space produces the same fit values, showing that our results are not sensitive to the assumption that unimodal signals are additively combined in extrinsic space. Third, the alternate model in which α_{MV} and α_{INV} were allowed to vary across the horizontal plane produces results similar to those of the one-dimensional models, demonstrating that our results do not depend on the assumption that vision and proprioception are weighted by a scalar term.

Despite these invariances, however, all of our models make the basic assumption that motor planning involves two stages, each using a separate estimate of arm position. This need not be the case, since the whole planning process could in theory be done in a single computational stage that computes an intrinsic motor command directly from the target location and a single estimate of the initial arm position (e.g., Uno et al., 1989). However, even if motor planning were performed in a single stage, the two types of error described in this paper would still arise.

A single-stage planner receiving unshifted feedback from the arm would determine the motor command appropriate to move the hand from the initial position to the target. The resulting movement constitutes the baseline trajectory. If the arm position estimate is shifted, the motor command will be the one appropriate to move the arm from the incorrect estimated location to the target. This command would achieve the target location if the arm were actually at the estimated position, and we will refer to that hypothetical trajectory as the planned trajectory. The difference between the planned and baseline trajectories is essentially the MV error described above. Since the arm is not at the estimated location, however, when the motor command is executed the resulting reach direction will differ from that of the planned trajectory. This difference is the INV error. If such a single-stage planner were in fact in operation, the MV and INV errors would be due to a single shifted estimate of arm position, and so we would expect that our analyses would find equal values for α_{MV} and α_{INV}. The fact that we found consistent differences between α_{MV} and α_{INV} suggests that planning indeed involves two separate stages.

Many cortical areas that encode pending movements appear to integrate information from multiple sensory modalities (Colby and Duhamel, 1996; Wise et al., 1997; Andersen et al., 1997), and single cortical neurons encoding arm position show varying weightings of visual and proprioceptive feedback from the arm (Graziano, 1999; Graziano et al., 2000). Given these findings, it is tempting to speculate that the two planning stages proposed in this paper might be computed in different cortical areas. Two lines of evidence support the idea that the parietal cortex is involved in the computation of extrinsic movement vectors. Recordings from the intraparietal sulcus have revealed coding of reach direction in retinocentric coordinates (Buneo et al., 2002), suggesting that this area encodes movement vectors but not intrinsic motor commands. Additionally, disruption of neural activity in the posterior

parietal cortex by transcranial magnetic stimulation prevents subjects from making corrective movements during reaching (Desmurget et al., 1999), providing further evidence that this region might help compute the discrepancy between hand position and target location. Studies examining the motor and premotor cortices, on the other hand, suggest a role for these areas in transforming movement vectors into motor commands. Neural activity in these areas during both single- and multi-jointed movements encodes the intrinsic details of the movement in addition to the extrinsic movement vector (Kakei et al., 1999; Scott and Kalaska, 1997; Scott et al., 1997). However, these findings do not represent a complete dissociation between parietal and frontal cortices. For example, Scott et al. (1997) showed that activity in parietal area 5 is also modulated by intrinsic factors. Future studies that use manipulations of sensory feedback to selectively alter the extrinsic movement vector and intrinsic motor commands will be needed to clarify the neural bases of these two computations.

Appendix to Chapter 1: Sensory Integration during Motor Planning

As described in Methods, in determining the weighting parameters α_{MV} and α_{INV} we assumed that the unimodal position estimates $\hat{x}_{vis}$ and $\hat{x}_{prop}$ were unbiased. In this Appendix, we consider the effects of biases in the unimodal position signals on the fit values for the α_{MV} and α_{INV}. We will show that sensory biases along the direction of the visual shift (0° in Figure 1) would not bias our fit values for α_{MV} and α_{INV}.

Given values for the unimodal positions estimates, an integrated position estimate $\hat{x}$ (which could be either $\hat{x}_{MV}$ or $\hat{x}_{INV}$) can be computed from the appropriate weighting parameter α (α_{MV} or α_{INV}), and *vice versa*:

$$\hat{x} = \hat{x}_{prop} + \alpha(\hat{x}_{vis} - \hat{x}_{prop}) \quad (11)$$

$$\alpha = \frac{\hat{x} - \hat{x}_{prop}}{\hat{x}_{vis} - \hat{x}_{prop}} \quad (12)$$

Note that in contrast to the Methods section, in Equation 12 and the rest of this Appendix we treat all position and bias variables as scalar values along the 0° axis. Since we do not have direct access to the internal unimodal estimates $\hat{x}_{vis}$ and $\hat{x}_{prop}$, we previously assumed that these estimates were unbiased, i.e. that $\hat{x}_{prop} = x$, the true location of the arm, and that $\hat{x}_{vis}$ was offset from that location by the visual shift. Here, we consider the possibility that the unimodal estimates are biased by amounts b_{prop} and b_{vis}, respectively:

$$\hat{x}_{prop} = x + b_{prop}$$

$$\hat{x}_{vis} = x \pm \Delta + b_{vis},$$

where the visual shift is $-\Delta$ in Center-Left trials and $+\Delta$ in Center-Right trials.

We first explore how these biases would affect our estimates of α when fit sep-

arately to trials with left and right visual shifts. Consider dividing our data into two sets, one containing Center-Left and Center-Zero trials, and the other containing Center-Right and Center-Zero trials. The integrated position estimates and the best-fit weighting parameters pertaining to the two datasets will be identified by the subscripts L and R. Given values for the unimodal biases, the integrated position estimates can be determined from the true weighting parameter α and Equation 11:

$$\hat{x}_L = x + b_{\text{prop}} + \alpha(-\Delta + b_{\text{vis}} - b_{\text{prop}}) \quad (13)$$

$$\hat{x}_R = x + b_{\text{prop}} + \alpha(+\Delta + b_{\text{vis}} - b_{\text{prop}}) \quad (14)$$

If α_L were fit with the assumption that $b_{\text{prop}} = b_{\text{vis}} = 0$, we would still be able to obtain the true value for $\hat{x}_L$ by evaluating Equation 13 using the same zero-bias assumption and the fit value of α_L. Counterintuitively, this means that we would be able to recover the true value of $\hat{x}_L$ even from a biased fit value for α_L. Next, we can determine the bias in our fit value of α_L by inserting the true value of $\hat{x}_L$ (Equation 13) into a version of Equation 12 that assumes, as do our models, that $\hat{x}_{\text{vis}}$ and $\hat{x}_{\text{prop}}$ are unbiased:

$$\begin{aligned} \alpha_L &= \frac{\hat{x}_L - x}{(x - \Delta) - x} \\ &= \alpha - \left[\alpha b_{\text{vis}} + (1 - \alpha) b_{\text{prop}}\right] / \Delta \end{aligned} \quad (15)$$

Similarly, from the Center-Right trials and Equations 12 and 14 we would find

$$\alpha_R = \alpha + \left[\alpha b_{\text{vis}} + (1 - \alpha) b_{\text{prop}}\right] / \Delta \quad (16)$$

Equations 15 and 16 suggest that the presence of biases in the internal unimodal

position estimates would cause errors in our fit values of α_{MV} and α_{INV}. However, a comparison of these two equations shows that the effects on the Center-Left and Center-Right trials are in the opposite direction, as long as the bias terms do not depend on the trial condition (a reasonable assumption given that the trials are interleaved and that the proprioceptive location, at least, is constant across trials). Therefore, in the complete analysis in which α_{MV} and α_{INV} are fit to all trials, we expect that any bias effects would cancel out.

This observation can be made more rigorous by averaging Equations 15 and 16, giving

$$(\alpha_L + \alpha_R)/2 = \alpha \tag{17}$$

Equation 17 tells us that the average of the parameters fit to the R and L datasets should be equal to the true weighting parameters, regardless of potential biases in the unimodal position signals. For each subject and each model, we fit α_{MV} and α_{INV} to the R and L datasets, and then compared $(\alpha_L+\alpha_R)/2$ to the fit values of α shown in Figure 7. The mean $\pm$ s.d. difference is 0.0067 ± 0.017 with a maximum absolute difference of 0.068. These results show that our assumption of unbiased unimodal signals had a negligible effect on our results. Additionally, note that because these arguments can be applied separately to the parameters α_{MV} and α_{INV}, Equation 17 holds for both parameters even if the biases in a unimodal position estimate are different between Equations 1 and 2. For this reason, biases in the transformations between coordinate frames (see Results) cannot be responsible for our findings.

Chapter 2: Flexible strategies for sensory integration during motor planning

Abstract

When planning goal-directed reaching movements, subjects combine visual and proprioceptive feedback from the arm to form two estimates of the arm's position: one to plan the reach direction, and another to convert that direction into a motor command. These position estimates are based on the same sensory signals but rely on different combinations of visual and proprioceptive input, suggesting that the brain weights sensory inputs differently depending on the computation being performed. Here we show that the relative weighting of vision and proprioception depends on both the sensory modality of the target and the information content of the visual feedback, and that these factors affect the two stages of planning independently. The observed diversity of weightings demonstrates the flexibility of sensory integration, and suggests a unifying principle by which the brain chooses sensory inputs in order to minimize errors arising from the transformation of sensory signals between coordinate frames.

Information about a single environmental cue or variable is often available to more than one sensory modality, raising the question of how the brain weights and combines sensory signals during perception and behavior. Previous theories of how this is accomplished have focused on the statistical properties of the individual sensory signals (Welch et al., 1979; Ghahramani, 1995; Jacobs, 1999; van Beers et al., 1999; Ernst and Banks, 2002). Here we demonstrate that input statistics alone are insufficient to explain sensory integration during motor planning. Rather, we provide evidence that sensory integration is strongly influenced by the cost of performing coordinate transformations on these signals. The results suggest a model in which sensory signals are weighted in order to limit the errors that are introduced when sensory signals are transformed between coordinate frames.

We have previously developed a method for studying the integration of visual and proprioceptive feedback signals from the arm during the planning of visually guided reaches (Sober and Sabes, 2003). Subjects reached to visual targets in a virtual visual feedback environment. The relative weighting of vision and proprioception was quantified by displacing the virtual visual feedback by a fixed vector and measuring the resulting reach errors. Such shifted feedback introduces two types of error: a "Movement Vector" (MV) error that arises when the location of the visual target is compared to the position of the fingertip (Fig. 10a), and an "Inverse Model" (INV) error that arises when the planned direction is converted into an intrinsic (muscle- or joint-based) motor command (Fig. 10b and c). Although errors from both of these stages can affect a single movement, the MV and INV error (Fig. 10d and e, respectively) can be separately inferred from the pattern of shift-induced errors measured across multiple reaches in different directions (Fig. 10f). Furthermore, the magnitude of each type of error is proportional to the reliance on the shifted visual feedback at each planning stage. By fitting these shift-induced error patterns with a simple

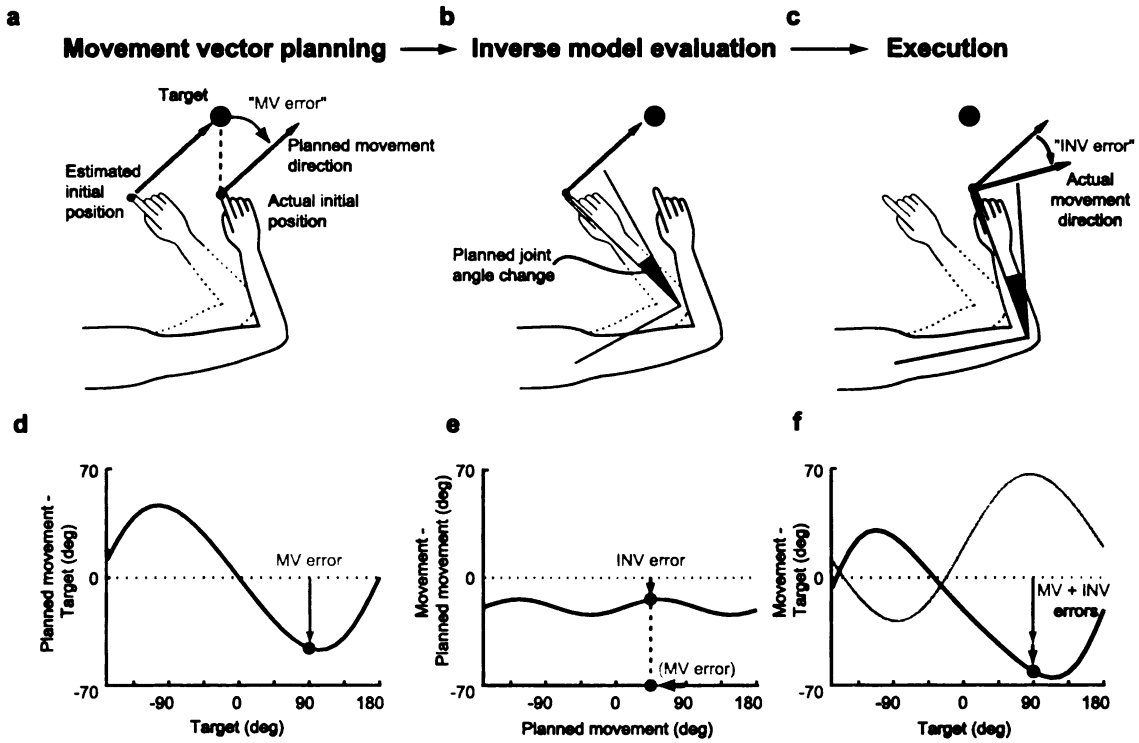

Figure 10: Two types of reaching errors induced by visual shift. **(a)** Error in planned movement direction that results from a leftward error in the position estimate used to plan the movement vector ("MV error"), for example due to a leftward shift of the visual feedback. The planned movement direction for a reach to the target at 90° (blue dot) differs from the true hand-to-target direction, resulting in a clockwise error in planned movement direction (blue arrow). **(d)** The bimodal MV error pattern resulting from a leftward error in estimated position. The blue arrow represents the MV error shown in panel **a**. **(b** and **c)** Error in initial movement direction resulting from a leftward error in the position estimate used to evaluate the inverse model of the arm and compute the motor command ("INV error"). The wrong joint angle displacement (motor command) is planned for the desired movement vector (panel **b**), resulting in a movement error (panel **c**). **(e)** A leftward error in estimated position results clockwise errors for all target directions. **(f)** Measured errors in initial direction are a combined effect of MV and INV errors (black curve). A rightward error in position estimates would produce an error pattern of opposite sign (gray curve). Here and in subsequent figures, 0° refers to rightward vector and positive angles are counterclockwise.

model of motor planning, we determined the relative weighting of visual feedback at each planning stage. We found that the two stages of reach planning rely on very different mixtures of visual and proprioceptive input: movement vector planning relies mostly on visual signals, whereas computation of the intrinsic motor command relies more on proprioceptive signals (Sober and Sabes, 2003).

This result raises the question of why the same sensory signals are weighted so differently at the two stages of reach planning. We suggest that the difference observed in Sober and Sabes (2003) was due to the coordinate frame in which the two computations took place. In that study, subjects reached to visual targets with their index fingers, and the visual feedback before reach onset reflected only the position of the fingertip. Computing the movement vector required an estimate of the current position of the fingertip relative to the visual target, a simple vector subtraction in visual coordinates. In contrast, computation of the motor command required an estimate of the intrinsic state of the arm, which was encoded by proprioceptive afferents but was not explicitly represented by the visual feedback. The sensory signals that were not already in the appropriate coordinate frame for each stage of planning would have had to undergo a transformation in order to be used in the relevant computation. We hypothesize that the weighting of vision and proprioception reflects a strategy of minimizing errors that arise during such coordinate transformations. We tested this idea in two experiments, altering either the coordinate frame of the target or the nature of the visual feedback.

We first examined whether sensory integration during movement vector planning is influenced by the sensory modality of the target (Experiment 1). When reaching to a visual target, the nervous system might rely more on vision because visual feedback from the arm is already in the target's coordinate frame. Before proprioceptive signals can be used, they must be transformed into visual coordinates. This

transformation will incur errors, either due to biases in the mapping between coordinate frames or to variance introduced by the additional computation (Soechting and Flanders, 1989b; McIntyre et al., 2000), and the effect of these errors will scale with the relative weighting of proprioception. Thus, if a proprioceptive target (i.e. the felt location of the other hand underneath the tabletop) were used instead of a visual target, we predict that subjects would rely more on proprioceptive feedback from the reaching arm, since the target would already be in an intrinsic coordinate frame. This change would be reflected as a reduction in the MV error observed when the visual feedback is shifted. Model data illustrating this prediction are shown in Figure 11a.

In Experiment 2, we asked whether the information content of the visual feedback influences the relative weighting of vision and proprioception. As shown in Figure 10b, an estimate of the joint angles is required to compute the appropriate motor command from the desired movement vector. We hypothesize that the previously observed reliance on proprioception during motor command generation was due in part to the fact that the visual feedback (a single spot of light) specified only the position of the fingertip (Sober and Sabes, 2003). The transformation of visual signals (from the fingertip) into proprioceptive coordinates (joint angles) is also expected to incur errors, again due to mapping biases or to the introduction of variance during the computation. Additionally, this transformation is particularly prone to errors since it is generally underconstrained: a single fingertip location can be achieved by many joint angle configurations when the arm's range of motion is unrestricted (Bernstein, 1967). Although the arm is actually restricted to two degrees of freedom in our experiments, the reliance on proprioception in this context likely reflects a general strategy of avoiding this underconstrained transformation of the visual input. If the virtual visual feedback explicitly represented joint angles as

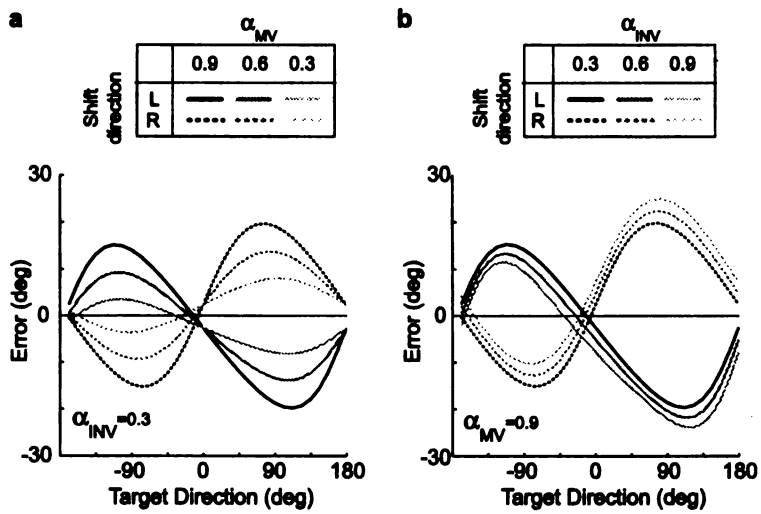

Figure 11: Predicted changes in Experiments 1 and 2. Each panel shows the movement direction errors predicted for different values of α_{MV} and α_{INV}, which represent the relative weighting of visual feedback at the first and second stages of reach planning, respectively (see Methods) . Both panels include the model predictions for ($\alpha_{MV} = .9$, $\alpha_{INV} = .3$), black lines, corresponding to the values expected for reaches to a visual target with fingertip feedback (Sober and Sabes, 2003). **(a)** Proprioceptive targets in Experiment 1 are expected to result in a lower value of α_{MV}. **(b)** Arm feedback in Experiment 2 is expected to result in a higher value of α_{INV}.

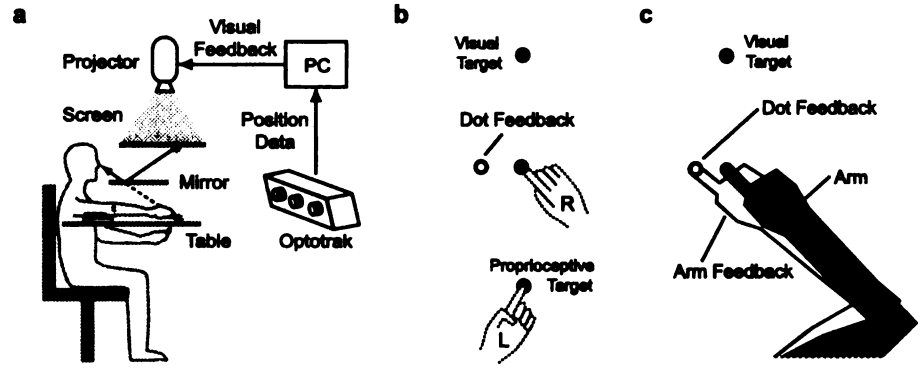

Figure 12: Data collection and task parameters. **(a)** Virtual visual feedback setup. **(b)** In Experiment 1 a visual and proprioceptive target were both present on every trial; the "go" signal cued subjects which reach to make. Visual feedback of the fingertip position was either veridical (gray dot) or shifted leftward (white dot) or rightward by 6 cm. **(c)** In Experiment 2, subjects reached to visual targets with visual feedback that showed either the fingertip or a simple, polygonal rendering of the arm. Visual feedback either reflected the true position of the fingertip or was shifted leftward (white dot, polygon) or rightward by 7 cm.

well as fingertip location, however, this source of uncertainty would be eliminated, and the errors due to the transformation would be reduced. We therefore predict that if the visual feedback included an image of the reaching arm (and hence explicit information about joint angles), subjects would rely more on vision during motor command generation. This in turn would lead to greater INV error when the visual feedback is shifted (model data shown in Fig. 11b).

Results

We tested these two predictions in separate experiments. In both studies, subjects made planar reaching movements on a horizontal table with virtual visual feedback (Fig. 12a). Visual feedback was provided as subjects planned reaches to one of six potential targets, but was extinguished as subjects began the reach.

In Experiment 1, two targets were available on each trial: a visual target, con-

sisting of a spot of light, and a proprioceptive target, provided by the felt position of the left index fingertip (Fig. 12b). The left hand was positioned beneath the table so that the two hands never came into contact. The color of the "go" signal instructed subjects whether to reach to the visual or proprioceptive target. This instruction was randomized across trials. Data from a typical subject are illustrated in Figure 13. For both target types, raw reach trajectories show the effects of shifted visual feedback (Fig. 13a). For example, reaches made with leftward feedback shifts display rightward directional errors, consistent with a leftward error in estimated hand position during movement vector planning. Note also that the effects of the visual shifts are smaller when the subjects reached to proprioceptive targets, consistent with our hypothesis that subjects would rely less on visual feedback in this trial type. Reaction times, movement times, and baseline reach trajectories were qualitatively similar for visual-target and proprioceptive-target reaches, although reaches to proprioceptive targets were more variable and tended to be hypometric (see Supplementary Information).

In quantifying the errors due to the pre-movement exposure to shifted visual feedback, we focused on the direction of movement just after reach onset, as this variable reflects the movement plan and not online feedback corrections. Shift-induced errors were computed by subtracting the initial movement direction on baseline trials from the initial movement direction on trials that included shifts of the visual feedback. The example subject's initial movement directions and shift-induced errors are shown in Figure 13b and c as a function of target and trial type. Note that these data resemble the model predictions of Figure 11a.

Using the simple model of motor planning illustrated in Figure 10, we inferred the relative weighting of vision and proprioception at each planning stage. This was accomplished by fitting two model parameters, α_{MV} and α_{INV}, that represent the

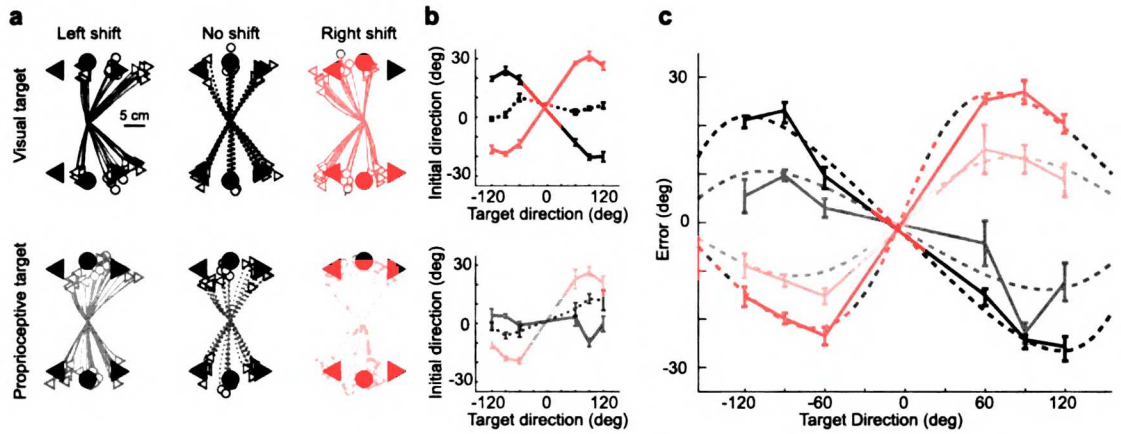

Figure 13: Sample data, Experiment 1. **(a)** Raw reach trajectories for a representative subject. Filled symbols, reach targets; corresponding open symbols, reach endpoints. **(b)** Initial reach direction, with respect to target direction, using the color conventions of panel **a**. **(c)** Reach errors induced by visual feedback shifts, computed by subtracting the baseline directional biases (dotted lines in panel **b**) from the initial directions in the shifted conditions (solid lines in panel **b**). Model fits, dashed lines. Best-fit weighting parameters: $\alpha_{\mathrm{MV}} = .96, \alpha_{\mathrm{INV}} = .32$ for visual-target reaches and $\alpha_{\mathrm{MV}} = .49, \alpha_{\mathrm{INV}} = .19$ for proprioceptive-target reaches.

relative weighting of vision when computing the movement vector and the intrinsic motor command, respectively ($\alpha = 0$ for proprioception only, and $\alpha = 1$ for vision only; see Methods for more details). This analysis was performed separately for the visual- and proprioceptive-target data shown in Figure 13c, and the model fits are shown as dotted lines. For this subject, the best-fit weighting parameters were ($\alpha_{\mathrm{MV}} = .96, \alpha_{\mathrm{INV}} = .32$) for visual-target trials and ($\alpha_{\mathrm{MV}} = .49, \alpha_{\mathrm{INV}} = .19$) for proprioceptive-target trials. This subject therefore relied almost entirely on visual feedback ($\alpha_{\mathrm{MV}} = .96$) when computing the movement vector to visual targets, but relied roughly equally on visual and proprioceptive information ($\alpha_{\mathrm{MV}} = .49$) when planning the movement vector to proprioceptive targets.

This trend was consistent across subjects. Figure 14 shows the initial direction and shift-induced errors averaged across the 7 subjects in Experiment 1. The group averages reflect the same trend observed for the example subject: shifted-induced

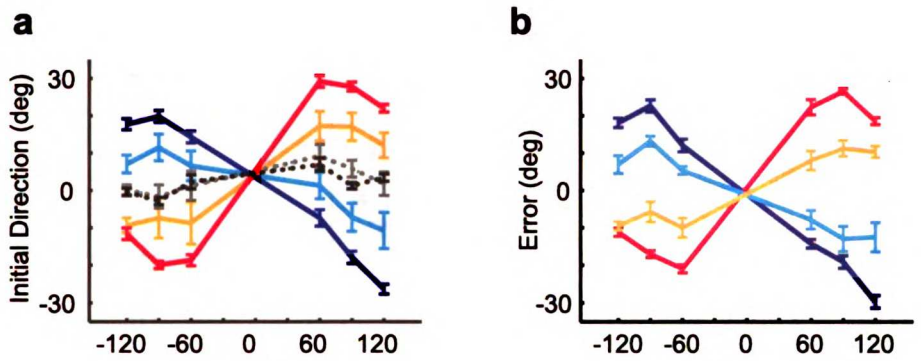

Figure 14: Group data, Experiment 1. **(a)** Average initial direction (±1 SE), with respect to target direction, for visual- and proprioceptive-target reaches. **(b)** Average errors in initial direction (±1 SE). Color conventions as in Figure 13.

errors are larger when reaching to visual targets, revealing a decreased reliance on vision when reaching to proprioceptive targets. This effect was quantified by the fit values of α_{MV} and α_{INV} (Fig. 15a and b). Every subject showed a significant ($p < .05$) reduction in α_{MV} when reaching to proprioceptive targets, as shown in Figure 15c (mean $\alpha_{MV} = .88$, visual-target trials; mean $\alpha_{MV} = .42$, proprioceptive-target trials). The mean change of .46 corresponds to a nearly fivefold increase the use of proprioceptive information in the proprioceptive-target condition. There was also a small reduction in α_{INV} (average change, .17), which was significant in a single subject (the completely filled circle in Fig. 15c). The weighting of feedback signals during movement vector planning therefore depends critically on the sensory modality of the target.

Three issues regarding the design of Experiment 1 warranted control studies. First, although subjects only received shifted visual feedback from the reaching (right) hand in Experiment 1, it is possible that this shifted feedback might have affected the felt position of the target (left) hand. Analogous transfer effects have been

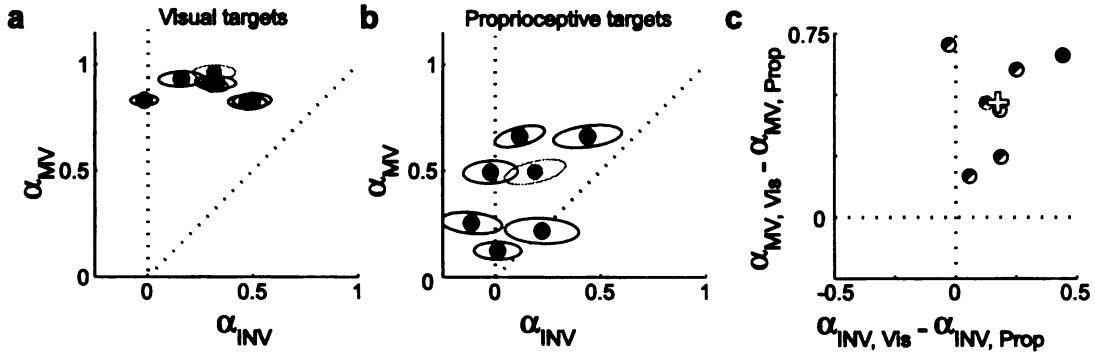

Figure 15: Model fits, Experiment 1. **(a,b)** Fit values of α_{MV} and α_{INV} for reaches to visual and proprioceptive targets, respectively. Ellipses represent 1 SE. Dotted lines, $\alpha_{INV} = \alpha_{MV}$ and $\alpha_{INV} = 0$. **(c)** Differences in model fit between visual- and proprioceptive-target reaches. The top left half of the symbol is filled if the task-dependent change in α_{MV} is significant (permutation test, $p<.05$), and the bottom right half is filled if the change in α_{INV} is significant. +, average across subjects. In all plots, the gray symbols are for the example subject shown in Figure 13.

described in studies of visuomotor adaptation (Kalil and Freedman, 1966; Wallace and Redding, 1979; Dizio and Lackner, 1995; Sainburg and Wang, 2002), although such effects are not always present (Taub and Goldberg, 1973; Cohen, 1973; Choe and Welch, 1974; Kitazawa et al., 1997). In this study, however, the feedback shift was not consistent from trial to trial, and so an adaptive shift in felt position of the left hand is not expected. In a separate experiment, we found that the visual shift protocol used in Experiment 1 does not affect the felt position of the left hand (see Supplementary Information).

Second, eye position was neither measured nor constrained in Experiment 1. This is potentially problematic, since subjects may have determined the target location for proprioceptive reaches using either a visual fixation point or the felt gaze direction, rather than the proprioceptive signals from the left arm. To address this concern, we repeated Experiment 1 with the additional constraint that subjects were required to maintain visual fixation on one of a number of points near the center of the workspace (see Supplementary Information). We found that the task-dependent changes in α_{MV}

and α_{INV} were very similar whether or not gaze was constrained. We conclude that any information about the target location provided by gaze direction had a minimal effect on sensory integration during reach planning.

A third potential confound in the proprioceptive-target condition is that subjects actively positioned the left arm and were required to support that arm against gravity during the planning and execution of the reaching movements. Efference copy therefore may have provided non-proprioceptive information about the position of the left arm. In order to ensure that this factor did not account for our results, we tested subjects on another variation of Experiment 1 in which the left arm was passively moved to the target location and was supported by a second tabletop throughout the trial. As shown in Supplementary Information, we found that the predicted change in α_{MV} was of comparable magnitude when the target arm was passively supported. The observed changes in α_{MV} therefore do not depend on the use of non-proprioceptive cues to the location of the proprioceptive target.

In Experiment 2, we tested the hypothesis that subjects would rely more on vision during inverse model evaluation if the visual feedback represented the arm, and thus the configuration of the joints, rather than just the location of the fingertip. All reaches were to visual targets, and feedback consisted of either a spot specifying fingertip location or a simple virtual image of the arm (Fig. 12c). Figure 16a and b show data from a single subject. The bimodal MV error pattern is similar in the two feedback conditions. Averaging across targets for each shift direction reveals the INV error pattern: reaches made with leftward shifts are biased clockwise (Fig. 10e), while reaches with rightward shifts are biased counter-clockwise. This INV error pattern is more pronounced in the arm-feedback condition in the example subject (compare the dashed lines in Fig. 16a and b), as in the rest of the subject pool (Fig. 16c). These data suggest that subjects rely more on vision when computing

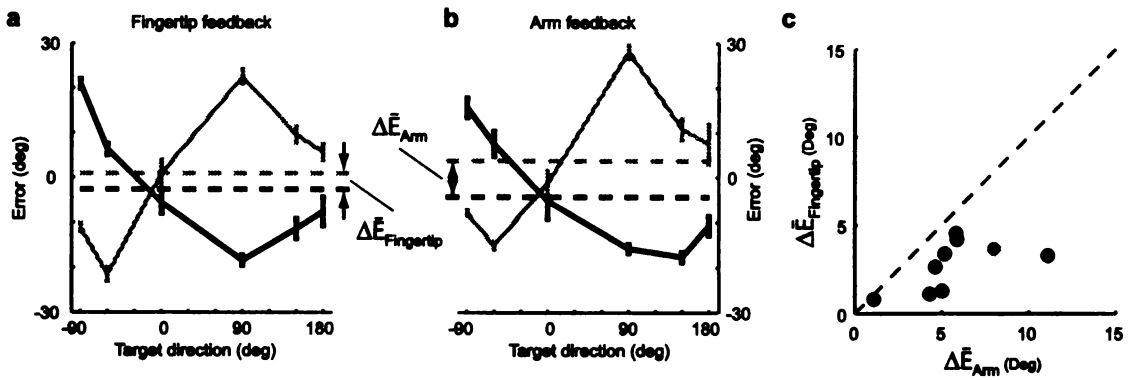

Figure 16: Sample and group data, Experiment 2. **(a,b)** Data from a representative subject in the fingertip-feedback and arm-feedback conditions, respectively. Gray, rightward visual shift; black, leftward visual shift; dashed lines, mean error across the six target directions. Arrows indicate $\Delta\bar{E}$, the difference between the mean errors for rightward- and leftward-shifted visual feedback. Error bars, ±1 SE. **(c)** Differences in mean errors for all subjects. Gray dot, example subject whose data are shown in panels **a** and **b**; dashed line, $\Delta\bar{E}_{\text{Arm}} = \Delta\bar{E}_{\text{Fingertip}}$.

the intrinsic motor command when the visual feedback includes the posture of the arm.

This effect is quantified by fitting the weighting parameters to the data from fingertip- and arm-feedback reaches. Model fits for the example data in Figure 16a and b yielded values of ($\alpha_{\text{MV}} = .84$, $\alpha_{\text{INV}} = .25$) for fingertip-feedback trials and ($\alpha_{\text{MV}} = .79$, $\alpha_{\text{INV}} = .54$) for arm-feedback trials (gray dots in Fig. 17a and b). This subject therefore relied 25% on visual feedback during the second stage of reach planning when the feedback represented only fingertip position. However, when visual feedback consisted of an image of the arm, that contribution increased to 54%. For all subjects, α_{INV} was greater during the arm-feedback trials (mean $\alpha_{\text{INV}} = .24$, fingertip feedback; mean $\alpha_{\text{INV}} = .42$, arm feedback), although this difference was only significant in 3 out of 10 subjects (Fig. 17c). The mean change of .18 corresponds to a 75% increase in the use of visual information in the arm-feedback condition.

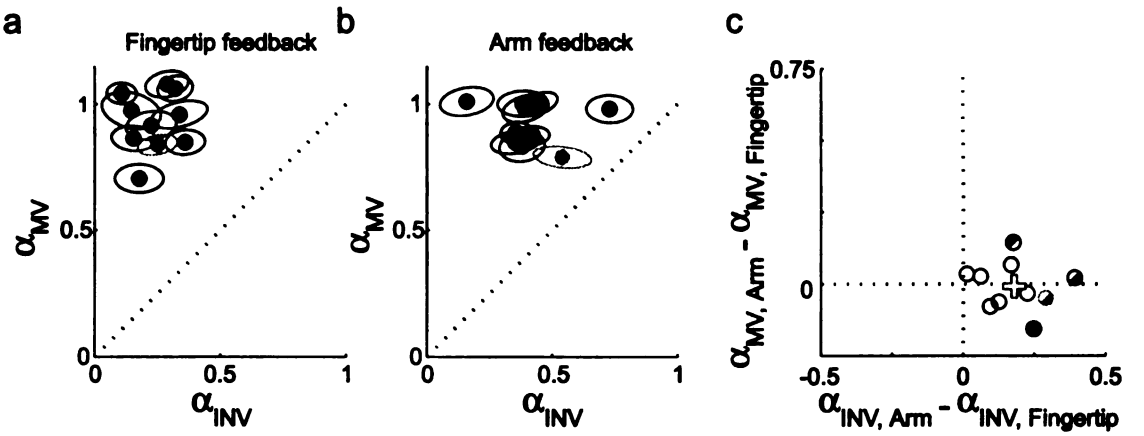

Figure 17: Model fits, Experiment 2. **(a,b)** Fit values of α_{MV} and α_{INV} for fingertip-feedback and arm-feedback trials, respectively. **(c)** Differences in model fit between arm- and fingertip-feedback reaches. In all plots, the gray dot represents the subject whose data is shown in Figure 16a and b. Other plotting conventions as in Figure 15.

Discussion

We have shown that altering the details of a sensorimotor task changes the relative weighting of visual and proprioceptive feedback used to plan the motor response. These changes can occur on a trial to trial basis as task conditions are varied. Furthermore, the task variations used in Experiments 1 and 2 drive different changes in weighting at the two stages of reach planning, suggesting a degree of independence in sensory integration at these stages.

Our results show that sensory integration is not determined solely by the statistics of the sensory input but is significantly influenced by the computations required for task execution. The observed pattern of changes reflects a reduced reliance on signals that must be transformed between visual and proprioceptive coordinate frames. There are several potential costs that could be minimized by avoiding transformed signals, for example the metabolic cost of performing additional neural computations. However, we propose that sensory signals are weighted in order to minimize

errors resulting from inherently noisy coordinate transformations. This points to a view of sensory integration as a set of local, independently controlled processes optimized to improve sensorimotor performance.

In Experiment 1, the value of α_{MV} depended on the sensory modality of the target, suggesting that the brain weights sensory feedback in order to minimize the adverse effects of transforming arm position signals into the coordinate frame of the target. Control studies demonstrated that these effects on α_{MV} were not the result of intermanual transfer of the visual shift, the effects of gaze direction, or the use of an active motor command to position the proprioceptive target. One remaining question is why this effect was not symmetric: when reaching to visual targets visual feedback is weighted almost 90%, whereas proprioception was weighted just over 50% when reaching to the fingertip of the other hand. This asymmetry may be due to the fact that when reaching to a proprioceptive target an additional transformation is required to convert proprioceptive signals from the left arm into the intrinsic coordinates of the right arm.

Another issue that arises from Experiment 1 is why α_{INV} decreased when reaching to proprioceptive targets. This effect was not predicted, since the target location is only used to compute the movement vector (see the schematic of Fig. 10). While this effect was smaller than the change in α_{MV}, the trend was apparent in all but one subject. Parallel changes in α_{MV} and α_{INV} raise the possibility that some factor other than the sensory modality of the target might be modulating sensory integration at both stages. One candidate is attention: when subjects are instructed to preferentially attend a particular sensory modality, that modality is weighted more (Warren and Schmitt, 1978; Welch and Warren, 1980). This raises the possibility that a shift of attentional focus from vision to proprioception could underlie the decreases in both α_{MV} and α_{INV}, despite the fact that the latter change is much

smaller. Note however that Experiment 1 was explicitly designed to reduce the likelihood of such attentional shifts: visual- and proprioceptive-target trials were randomly interleaved, every trial included both a visual and a proprioceptive target, and subjects were not instructed which target to reach to until shortly before the "go" signal (see Methods). Thus, any shift in subjects' attention would have had to occur between the appearance of the target instruction and movement onset. Furthermore, attentional effects alone cannot explain key components of our results: vision and proprioception can each dominate at different planning stages during the same movement (see Figs. 15 and 17), and in Experiment 2 α_{INV} increased without consistent changes in α_{MV} (Fig. 17c).

Experiment 2 demonstrated that the value of α_{INV} depends on the information content of the visual feedback. Proprioception predominates when the feedback represents only fingertip position (mean $\alpha_{\mathrm{INV}} = .25$). We argue that this strategy is chosen in order to minimizes the effects of errors that occur when inferring the posture of the arm from the position of the fingertip. On the other hand, vision and proprioception are weighted nearly equally when visual feedback of the whole arm is available (mean $\alpha_{\mathrm{INV}} = .54$, see also Fig. 17), suggesting that simply rendered, virtual visual feedback of the arm is equally as informative to the brain as proprioception for determining the arm's posture. The fact that the average α_{MV} was nearly identical in the dot and arm feedback conditions (see Results) shows that the type of feedback has a specific effect on the second stage of reach planning. This specificity suggests that the key difference between the two conditions is the information about the configuration of the joints, rather than the size or salience of the visual stimulus.

Our results demonstrate that simple minimum-variance models of sensory integration, in which the weighting of sensory signals is determined solely by their input statistics, are incomplete. The minimum-variance principle has successfully

described sensory integration in a wide variety of tasks (Ghahramani, 1995; Jacobs, 1999; van Beers et al., 1999; Ernst and Banks, 2002), however it is inconsistent with our finding that the same sensory signals can be given very different weightings at different stages of reach planning or at the same stage of planning in different tasks. While the brain does appear to minimize the variability of integrated signals, we propose that this computation takes into account the additional variability that arises during the coordinate transformations required for a given sensorimotor task. Minimum-variance theories must therefore be extended to include the cost of coordinate transformations in computing the optimal sensorimotor strategy.

The experiments presented here demonstrate for the first time that sensory integration during reach planning is a dynamic process driven by the computational demands of the task. These psychophysical phenomena presumably reflect differences in the neural circuits or patterns of activity that are engaged during the performance of each task. Neurons in many of the cortical areas involved in reach planning receive both visual and proprioceptive inputs (Andersen et al., 1997; Wise et al., 1997; Graziano, 1999; Graziano et al., 2000). Our results suggest that these sensory inputs are gated in a task-dependent way. By quantifying the mixture of visual and proprioceptive inputs to a neural population across the task manipulations described above, it may be possible to determine the computation in which a group of neurons is participating. Such studies will help bridge the gap between computational and physiological descriptions of motor planning.

Methods

All subjects were right-handed, 18-34 years of age, and were healthy with normal or corrected-to-normal vision. Subjects were naive to the purpose of the experiments and were paid for their participation. Seven subjects (5 women, 2 men) performed Experiment 1, and 10 different subjects (5 women, 5 men) participated in Experiment 2. There was no overlap in the subject pools for the two experiments.

Experimental Setup and Data Collection

Visual targets and visual feedback of the arm were presented using a virtual-reality display system described previously (Sober and Sabes, 2003) and shown in Figure 12a. Reaches were performed with the right arm, which rested on a shoulder-height table. A custom-built splint immobilized the wrist joint and fixed the right index finger in an extended position. An air sled under the upper arm minimized friction between the arm and the table. The torso was restrained using a vest attached to the back of the subject's chair. This arrangement restricted the arm to two degrees freedom such that arm position could be expressed interchangeably as x, a two-dimensional vector representing the Cartesian position of the fingertip on the table, or as θ, a two-dimensional vector consisting of the elbow and shoulder angles.

General Task Design

In both experiments, reaches began from a start point located on the subject's midline and approximately 40 cm from the chest. Six targets were located on an imaginary circle centered at the start point, although the exact location of the targets differed between the two experiments (see below). At the beginning of each trial subjects received visual feedback from the reaching arm. This visual feedback re-

flected either the true location of the fingertip or a location shifted to the left or right. Rightward, leftward, and null shifts were randomly interleaved to prevent sensorimotor adaptation. The visual feedback disappeared after reach onset once the fingertip moved 5 mm from the start point. Reach endpoints were defined as the first time after reach onset that the tangential velocity of the fingertip dropped below 3 mm/sec.

After completing the experiment, subjects were asked whether the visual feedback had accurately reflected the position of the right arm. All subjects in both experiments reported being unaware of any visual shifts.

Scoring

At the end of each trial, text appeared on the feedback display which informed subjects of their current score and whether their reaching movement had been a "hit" or a "miss." If a reach ended within 2 cm of the target, the trial was judged a "hit" and subjects received 1 point. All trials that included visual feedback shifts were counted as "hits" to avoid providing error feedback and thus driving sensorimotor adaptation. If subjects missed the target by more than 2 cm on an unshifted trial, the trial was judged a "miss" and one point was deducted from the score. In order to reduce variability in planning time and velocity profiles, subjects' score also depended on two timing criteria. Two time intervals were defined. Interval 1 was from the "go" signal until the time when the fingertip had moved 5 mm from the start point. Subjects were allowed a minimum and maximum time of 150 and 750 msec for this interval. Interval 2 was from the end of interval 1 until the first time after the peak tangential velocity that the velocity fell to 10 cm/sec. Subjects were allowed a minimum and maximum time of 250 and 650 msec for this interval. If either of

these timing requirements was not met, subjects received zero points if they had hit the target on that trial and lost one point if they had missed. Text appeared at the end of each trial to inform subjects of any timing errors on that trial. Subjects committed relatively few timing errors. The average error rate across subjects for interval 1 was 3% in Experiment 1 and 2% in Experiment 2. The error rate for interval 2 was 3% in Experiment 1 and 5% in Experiment 2.

Task Design - Experiment 1

In Experiment 1 we varied the sensory modality of the reach target (Fig. 12b). On half of the trials ("visual-target trials"), subjects reached to a visual target consisting of a dot of red light 5 mm in diameter. On the other half of the trials ("proprioceptive-target trials"), subjects reached to the felt location of their left index finger, which had been positioned below the tabletop as described below. On all trials, visual feedback of the right index fingertip consisted of a white spot of radius 5 mm which appeared either at the fingertip or shifted leftward or rightward by 6 cm.

At the beginning of each trial, subjects positioned their right fingertip to place the feedback spot within an 8 mm ring. The feedback dot was visible only when the Y-coordinate of the fingertip (defined as the axis pointing towards/away from the body) was within 8 mm of the Y-coordinate of the start point. Once the starting position was achieved, the ring disappeared and an array of 9 arrows appeared. The arrows were used to guide subjects' left index finger along the bottom surface of the table (Fig. 12a) to the proprioceptive target location without giving explicit visual feedback of the fingertip's position. The location of the arrows remained fixed, but the direction and magnitude were adjusted online, indicating the direction and (scaled) distance from the left fingertip to the proprioceptive target. The positions

of the arrows were randomized across trials and were not correlated with the target location. When the fingertip had moved to within 5 mm of the specified position for the proprioceptive target, the arrows disappeared and the visual target (a red dot) appeared at a location distinct from the start point and the proprioceptive target (Fig. 12b). Visual and proprioceptive targets were selected randomly on each trial from six possible locations ($60°, 90°, 120°, -120°, -90°, -60°$ as shown in Fig. 13a) on an imaginary circle of radius 14 cm, centered at the start point. Target directions in Experiment 1 were chosen to maximize the statistical power of the model when fitting α_{MV}. The visual and proprioceptive targets were always 180° apart, as in the example shown in Figure 12b. After a delay of 500-1500 msec, a target instruction cue appeared in the form of a small colored dot on top of the feedback spot. Subjects reached to the visual target if this dot was red and to the proprioceptive target if the dot was blue. The target instruction cue remained illuminated for 1400-1600 msec and then disappeared. The disappearance of the target instruction cue served as the "go" signal, and cued the subject to begin the reach. When the tangential fingertip velocity fell below 3 mm/sec at the end of the reach, the movement was scored. On trials with unshifted feedback that were scored as misses, subjects were required to make a corrective movement to the target. On reaches to visual targets, this correction was guided with reilluminated visual feedback. On reaches to proprioceptive targets, a field of arrows directed the right fingertip to the left, in the manner described above. Subjects were required to hold the left index fingertip at the proprioceptive target while performing the reaching movement with the right arm. All subjects did so successfully on every trial: mean excursion of the left index fingertip during reaches was < 1 mm in all subjects, and in no case did the fingertip move more than 1 cm.

Eight reaches were made to each of six targets with three possible visual shifts,

totaling $8 \times 6 \times 3 = 144$ reaches in both the visual- and proprioceptive-target conditions. Reaches to visual and proprioceptive targets were randomly interleaved, forming a dataset of 288 trials. A set of 36 familiarization trials preceded the experiment and included all six trial types. In this familiarization period, however, trials that would normally include leftward or rightward visual shifts were performed without visual shifts, but with the start points moved leftward or rightward by 6 cm. The actual start points therefore matched the apparent start points that would have been seen in the presence of a visual shift. Familiarization trials acclimated subjects to the task and familiarized them with making reaches that appeared to originate from three different start points. Familiarization trials were not included in the subsequent analyses.

Task Design - Experiment 2

In Experiment 2 we varied the information content of the visual feedback (Fig. 12c). In half of the trials ("fingertip feedback trials"), the visual feedback from the right arm consisted of a bright white spot (100% gray level on the XGA projector used to present visual images), as in Experiment 1, and feedback was either veridical or shifted right or left by 7 cm. In the other half of trials ("arm-feedback trials"), the visual feedback consisted of a white polygon (60% gray level) in the shape of the subject's arm. The arm feedback was either coextensive with the subject's arm or the fingertip was shifted left or right by 7 cm. When the fingertip was shifted, the arm feedback was displayed with the correct posture given the shifted location of the fingertip and the true position of the subject's shoulder.

At the beginning of a trial, subjects used visual feedback to position their right index fingertip at the start point, as in Experiment 1. In Experiment 2, however,

visual feedback appeared when the Y-coordinate of the fingertip was within 3 cm of starting value. When the visual feedback was within the start window, a visual target consisting of a red dot of radius 5 mm appeared. Target location was selected randomly on each trial from one of six positions: $0°, 90°, 150°, 180°, -90°$, and $-60°$ along an imaginary circle with a radius of 18 cm. Target directions in Experiment 2 were chosen to maximize the statistical power of the model when fitting α_{INV}. After a delay of 500-1500 msec, a small red dot appeared on top of the visual feedback spot (fingertip-feedback trials) or the tip of the animated finger (arm-feedback trials), remained illuminated for 1400-1600 msec, and then disappeared, serving as the "go" signal. At the end of trials with unshifted feedback, the feedback was reilluminated to provide error feedback, as in Experiment 1. Subjects did not use their left hands in Experiment 2.

For each type of visual feedback, eight reaches were made to each of the six targets with the three visual shifts, totaling $8 \times 6 \times 3 = 144$ reaches. The two feedback conditions were presented in blocks. Half of the subjects completed 144 fingertip-feedback trials followed by 144 arm-feedback trials, whereas the other half of the subjects completed the blocks in the reverse order. Block order had no effect on reach kinematics or the results or the model fits. Each block began with 36 familiarization trials as in Experiment 1.

Trajectory analysis

Position data were smoothed with a low-pass Butterworth filter (cutoff frequency of 6 Hz), and fingertip velocity was computed using first-order numerical differentiation. We quantified initial reach directions by determining the angle of the instantaneous velocity at the point when the tangential velocity first exceeded 40% of its peak

value. This landmark fell 82 $\pm$ 18 msec (mean $\pm$ average within-subject SD) after reach onset in Experiment 1 and 94 $\pm$ 19 msec (mean $\pm$ average within-subject SD) after reach onset in Experiment 2.

Modeling the initial movement direction

Work from our and other labs has shown that subjects form two estimates of arm position when planning a reaching movement (Sober and Sabes, 2003; Ghilardi et al., 1995; Rossetti et al., 1995; Goodbody and Wolpert, 1999; Sainburg et al., 2003). We have developed a method for quantifying the relative reliance on vision and proprioception during reach planning (Sober and Sabes, 2003), which we applied to the data collected from Experiments 1 and 2. Briefly, we assume that the two arm position estimates, $\hat{x}_{\mathrm{MV}}$ and $\hat{x}_{\mathrm{INV}}$, are weighted combinations of visual ($\hat{x}_{\mathrm{vis}}$) and proprioceptive ($\hat{x}_{\mathrm{prop}}$) inputs:

$$\hat{x}_{\mathrm{MV}} = \alpha_{\mathrm{MV}}\,\hat{x}_{\mathrm{vis}} + (1-\alpha_{\mathrm{MV}})\,\hat{x}_{\mathrm{prop}} \tag{18}$$

$$\hat{x}_{\mathrm{INV}} = \alpha_{\mathrm{INV}}\,\hat{x}_{\mathrm{vis}} + (1-\alpha_{\mathrm{INV}})\,\hat{x}_{\mathrm{prop}}. \tag{19}$$

The weighting parameters α_{MV} and α_{INV} describe the relative reliance on vision and proprioception used when forming the position estimates.

In the model, the estimate $\hat{x}_{\mathrm{MV}}$ is used to calculate the desired initial fingertip velocity, $\dot{x}^*$ or the "movement vector". The direction of this vector is determined by

$$\measuredangle \dot{x}^* = \measuredangle(x_d^* - \hat{x}_{\mathrm{MV}}) + \omega_d, \tag{20}$$

where $\measuredangle z$ represents the angle of any vector z, x^* represents the target location, ω is an angular offset term that captures directional biases in baseline reaches, and

$d \in [1, ..., 6]$ indexes x^* and ω over the six targets. Note that an error in $\hat{x}_{\text{vis}}$ due to shifted visual feedback will cause errors in $\hat{x}_{\text{MV}}$, and thus in $\measuredangle \dot{x}^*$, and that the magnitude of these errors will scale with α_{MV} (Equation 18 and Fig. 11a).

The desired fingertip velocity $\dot{x}^*$ is then used to compute the motor command, modeled here as a vector of joint angle velocities $\dot{\theta}$:

$$\dot{\theta} = J^{-1}(\hat{\theta}_{\text{INV}})\dot{x}^*, \tag{21}$$

where $\hat{\theta}_{\text{INV}} = K^{-1}(\hat{x}_{\text{INV}})$, is calculated from the inverse kinematics function K^{-1} that describes the mapping from Cartesian fingertip location to joint angles, and $J(\theta) = \frac{dx}{d\theta}$ is the Jacobian of the arm. This planning stage represents an evaluation of an inverse model of the arm (Jordan, 1996)

When the motor command is executed, the fingertip moves with an initial velocity described by

$$\dot{x} = J(\theta)\,\dot{\theta} = \left[J(\theta)\,J^{-1}(\hat{\theta}_{\text{INV}})\right]\dot{x}^*. \tag{22}$$

Note that shifting the visual feedback will cause a discrepancy between $\dot{x}$ and $\dot{x}^*$, since the matrix $J(\theta)\,J^{-1}(\hat{\theta}_{\text{INV}})\dot{x}^*$ will equal the identity matrix only when $\theta = \hat{\theta}_{\text{INV}}$. The size of this INV error increases with increasing α_{INV}, as shown in Figure 11b. Finally, since Equation 20 only specifies the direction of $\dot{x}^*$, only the predicted direction of the initial movement vector in Equation 22 will be fit to empirical data.

Our conclusions rest on the validity of this model and the resulting fit values of the weighting of vision and proprioception. See Sober and Sabes (2003) for a discussion of the validity of the model and demonstrations that the fit values of α_{MV} and α_{INV} are robust across various assumptions about the nature of reach planning.

Fitting model predictions to the data

We fit the two free parameters of the model, α_{MV} and α_{INV}, to the data from each subject and experimental condition. Each fit was performed with a nonlinear regression algorithm (nlinfit in MATLAB, The Mathworks Inc., Natick, MA) that minimized the squared error between the predicted initial movement directions and those measured in shifted feedback trials. The values of $\hat{x}_{\mathrm{vis}}$ and $\hat{x}_{\mathrm{prop}}$ in the model were set to the position of the visual feedback and the true position of the fingertip, respectively, on each trial. The value of x^* was set to the target position. The baseline bias terms (ω_d in Equation 20) were determined from the mean angular difference between the target direction and the initial velocity direction in the unshifted feedback trials.

We used permutation tests (Good, 2000) to evaluate the hypotheses that the values of α_{MV} and α_{INV} differed between experimental conditions. Additionally, we used a bootstrapping technique (Efron and Tibshirani, 1993) to place confidence limits on the fit values of α_{MV} and α_{INV}. Details of these statistical procedures are described in Supplementary Information.

Supplementary Information to Chapter 2: Flexible strategies for sensory integration during motor planning

Effects of target direction and task condition

In order to examine the effects of task parameters on reach kinematics in the absence of visual feedback shifts, we tested whether the target direction or task affected the direction and timing of baseline (unshifted) reaches. Two-way ANOVAs were performed, with target direction as one factor and target type (for Experiment 1) or feedback type (for Experiment 2) as the second. Four dependent variables were tested: baseline directional biases (ω_d in Equation 20), total reach distance, reaction time (RT), and movement time (MT). Reaction time was calculated by finding the interval from the "go" signal until the first time that tangential velocity exceeded 3 mm/sec. Movement time was defined as the latency from this landmark to the first time velocity falls to 3 mm/sec after reaching peak velocity. This latter landmark was defined as the reach endpoint when computing the total reach distance. The corrective movements that followed "miss" trials were ignored. We used the Bonferroni correction (Campbell and Machin, 1993) to compensate for the effects of repeated testing. The threshold for significance was therefore .05/N, where N was the number of separate ANOVAs applied to the data from each experiment. For Experiment 1, N=28 (7 subjects x 4 two-way ANOVAs), and for Experiment 2, N=40 (10 subjects x 4 two-way ANOVAs).

The effects of target direction and target type (visual- vs proprioceptive-target reaches) on various reach parameters are shown in Table S1. Target directions had a strong effect on baseline directional error and reach distance due to small biases in unperturbed reaches. The most salient effect of target type was a tendency for subjects to underreach (i.e. make hypometric movements) to proprioceptive targets. Mean reach distance (averaged across subjects and target directions) was 127 mm for proprioceptive targets and 138 mm for visual targets. This trend resulted in a significant main effect of target type on reach distance in 5 of 7 subjects, and a

	Baseline bias	Reach distance	RT	MT
Main effect of direction	5/7	7/7	-	3/7
Main effect of target type	-	5/7	-	4/7
Interaction	3/7	2/7	-	1/7

Table 1: ANOVA results, Experiment 1. Number of subjects with significant ($p<.05$) effects of target direction, target type (visual vs. proprioceptive), or interactions between direction and type. RT, reaction time; MT, movement time; -, no subjects with significant effect.

	Baseline bias	Reach distance	RT	MT
Main effect of direction	9/10	5/10	-	4/10
Main effect of feedback type	2/10	6/10	2/10	-
Interaction	-	1/10	-	1/10

Table 2: ANOVA results, Experiment 2. Number of subjects with significant ($p<.05$) effects of target direction, feedback type (dot vs. arm), or interactions between direction and type. Conventions as in Table S1.

significant main effect of target type on movement time in 4 of 7 subjects. Additionally, there was greater variability in reach direction for reaches to proprioceptive targets. This trend was consistent across subjects and may have resulted from a greater uncertainty in the location of the proprioceptive target than the position of the visual target (van Beers et al., 1998).

As in Experiment 1, target direction affected baseline directional bias in most subjects in Experiment 2 (Table S2). Although there was a main effect of feedback type on reach distance in 6 of 10 subjects, there was no consistent trend to the sign of this effect (average reach distance was greater for arm-feedback reaches in 5 of 10 subjects and greater for fingertip-feedback trials in the other 5).

In order to demonstrate that the effect of target type on the fit value of α_{MV} in Experiment 1 was not an artifact of the tendency to make hypometric reaches to proprioceptive targets (see Table S1), we fit the data with an alternative version of the model. In the new model, the target position (x^* in Equation 20) was set to the mean reach endpoint in the baseline condition. In other words, we assumed that

reaches to proprioceptive targets were hypometric because of a bias in estimated target position. In this alternate analysis, we observed the same reduction in α_{MV} when reaching to proprioceptive targets as in the original analysis. The change in α_{MV} was significant in all subjects, and had a mean value of .60.

Control for intermanual transfer of feedback shift effects

Experiment 1 showed that MV errors induced by visual shifts were smaller when subjects reached to proprioceptive (as opposed to visual) targets. In the model-based analysis, this difference was assumed to be the result of a reduced reliance on the (shifted) visual feedback in the proprioceptive-target condition. However, there is an alternative explanation for this change in error magnitude: intermanual transfer of the feedback shift effects to the left hand could bias location of the proprioceptive target. In the extreme case, if $\hat{x}_{MV}$ and the felt position of the proprioceptive target were shifted by the same amount and in the same direction, then the fit value of α_{MV} would be equal to zero in proprioceptive-target reaches, since the desired movement direction would still be correct. Such an effect could potentially be responsible for the observed shift in α_{MV}.

To address this possibility, we performed an additional experiment in which six subjects used gaze direction to indicate the location of the proprioceptive target. As in Experiment 1, subjects positioned the right index fingertip at the start point using visual feedback. An arrow field then guided the left index fingertip to a target location under the tabletop. The visual feedback (a white spot of radius 5 mm) was either veridical or displaced to the left or the right by 6 cm. Unlike in Experiment 1, however, only proprioceptive targets were used. Also, the targets were located only at the 60, 90, and 120° positions, due to the constraints of the eye monitoring system. 750 ms after the right index fingertip arrived at the start position, a small green dot

appeared on top of the feedback spot. When the green dot appeared, subjects moved their eyes to fixate the position on the tabletop corresponding to the felt location of the left index fingertip and then tapped the right index fingertip. An infrared eye-tracking system (ISCAN, Burlington, MA) was used to monitor gaze direction. Subjects were allowed as much time as they needed to fixate the left index fingertip, and were instructed to tap their right index finger only when they were sure that they were fixating the proprioceptive target. After the subjects had tapped their right fingertip, the green dot disappeared and a small blue dot appeared in its place. After 1900-2100 msec the blue dot disappeared, signaling the subject to reach to the proprioceptive target.

Figure 18 shows that although shifts of the visual feedback do affect reach endpoints (as in Experiments 1 and 2), they do not affect the felt location of the proprioceptive target. Leftward shifts in visual feedback lead to rightward biases in reach endpoints (blue cross in Figure 18a). If a leftward shift in the feedback from the right hand affected the felt location of the left hand, one would expect to see a leftward shift in the gaze position used to indicate the location of the proprioceptive target (blue cross in Figure 18b). However, no such pattern is seen (Figure 18c). This experiment demonstrates that errors in the arm position estimates induced by visual shifts do not transfer from the right to the left arm.

Control for eye gaze direction effects

A key premise of Experiment 1 is that target information in the proprioceptive-target condition comes from the felt position of the left arm. However, in our initial experiment, gaze was neither monitored not constrained. This presents a potential confound, since gaze signals could provide additional target information or bias the planning of reaching movements. For example, in the proprioceptive-target condi-

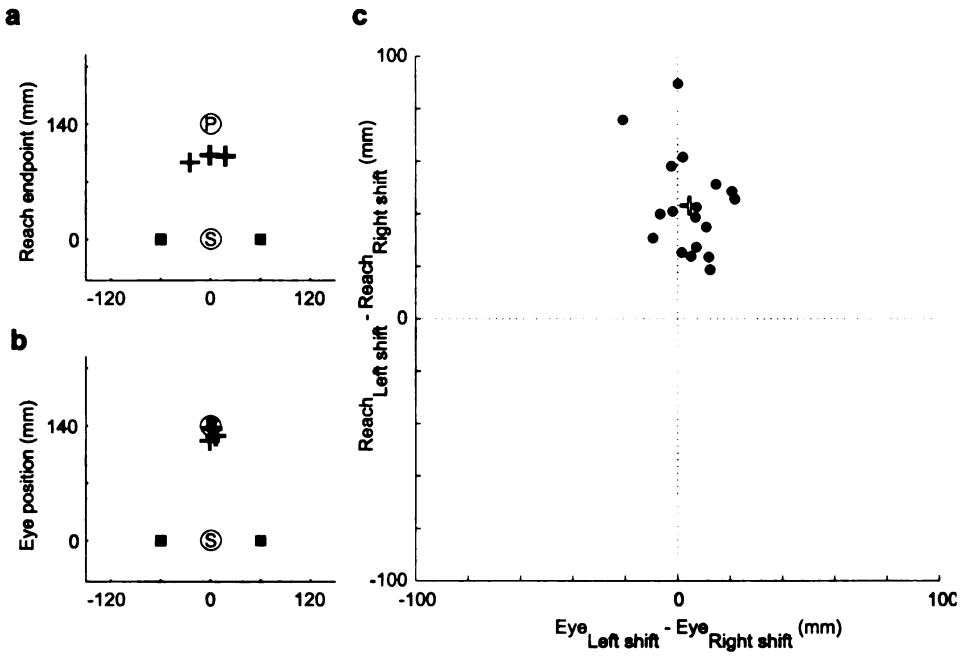

Figure 18: Results of intermanual transfer experiment. Data from reaches to a single target from a representative subject are shown in A and B. A shows the average reach endpoints for reaches made with leftward (blue crosses), zero (black crosses), and rightward (red crosses) shifts of the visual feedback. P, proprioceptive target location, S, reach start point. The blue and red squares show the location of the left- and right-shifted visual feedback, respectively. B shows the average gaze position used to indicate the proprioceptive target location (color conventions as in A). A scatterplot of the group data are shown in C. Values on the abscissa indicate the average left-right separation of gaze positions recorded after leftward and rightward shifts of visual feedback. Ordinate values indicated the average right-left separation of reach endpoints after leftward and rightward visual shifts. Each point represents the values for a single subject and target. The red dot represents the data point obtained from the subject and target shown in A and B, and the white cross indicates the mean value across subjects and targets.

tion, subjects might have been reaching not to the felt position of the left fingertip, but rather to a visual fixation point or to a target specified by extra-retinal signals related to gaze direction. Another potential problem is that subjects subject might have directed their gaze to different parts of the workspace in the visual-target and proprioceptive-target conditions (e.g. fixating visual targets when reaching to them, but fixating the feedback from the reaching hand when reaching to proprioceptive targets). Since gaze direction has been shown to affect the accuracy of reaching movements (Bock, 1986), it is therefore possible that the different patterns of motor error in the two target conditions result from different gaze behaviors.

In order to control for these effects, we ran 5 additional subjects on a modified version of Experiment 1. In each trial, a fixation cross (Figure 19, a-c) appeared as soon as the subject positioned the left index fingertip at the proprioceptive target. After a delay of 1400-1600 msec, the fixation cross turned either red or blue, instructing the subject to reach to the visual or proprioceptive target, respectively. Subjects were required to maintain fixation for the duration of the reaching movement.

Subjects were able to maintain appropriate fixation on most trials. Trials on which the subject broke fixation before tangential fingertip velocity exceeded 40% of its peak value (see Methods) were marked as errors and excluded from analysis. Very few trials had to be excluded (average error rate 3.33%, range $2.08 - 4.51\%$).

The results of this modified experiment were similar to those of the original version. Figure 19d shows that as before, subjects relied mostly on vision when reaching to visual targets, but in every case increased their reliance on proprioception when reaching to proprioceptive targets (Figure 19e). The magnitude of these changes (mean change in $\alpha_{MV} = .50$, mean change in $\alpha_{INV} = .16$, Figure 19f) was in very close agreement with those observed in the original version of Experiment 1 (mean change in $\alpha_{MV} = .46$, mean change in $\alpha_{INV} = .17$). These results show that the

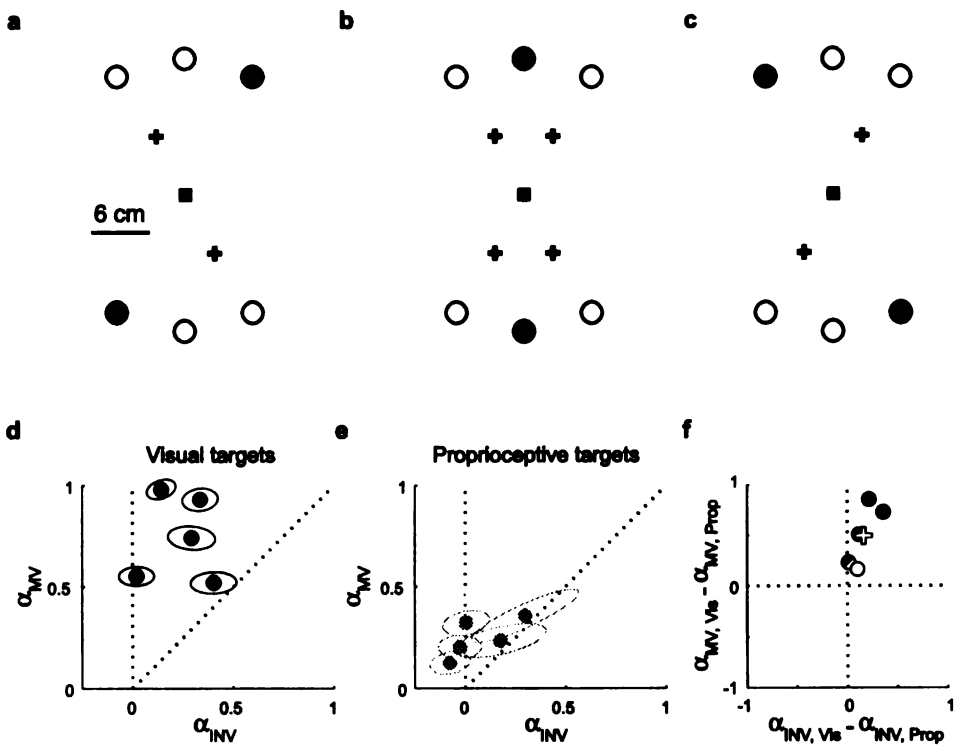

Figure 19: Methods and model fits, gaze direction control experiment. In each trial, subjects fixated a cross that appeared in the workspace. The location of the fixation cross was chosen randomly on each trial from a set of locations (+ symbols in a-c), situated roughly halfway between the start point (black squares) and circle of targets (black dots). Potential fixation points were chosen based on the location of the visual and proprioceptive targets. Two of the four fixation points were used for reaches where the targets were at 60° and 240° (a) and at 120° and 300° (c). When targets were at 90° and 270° (b), all four potential fixation points were used. The results of this experiment are shown in d-f, using the same conventions as in Figure 15 of the main text.

changes in the weighting of vision and proprioception are not due to a difference in gaze strategy between the two conditions, and demonstrate that our results do not change substantially when subjects are not allowed to fixate the proprioceptive target during reach planning.

Although the task-dependent changes in α_{MV} were nearly identical in the original and gaze-constrained versions of Experiment 1, constraining the direction of gaze appears to have affected the absolute value of α_{MV}. In reaches to both visual and proprioceptive targets, the value of α_{MV} was lower when eye position was constrained (compare Figure 15a,b with Figure 19d,e). This trend, which had a magnitude of .14 in the visual-target condition and .17 in the proprioceptive-target condition, fell short of significance in both cases (t-test, p=.12 and p=.14, respectively). This effect is likely due to the fact that the fixation constraint put the visual feedback in the visual periphery at the time of reach initiation. Since peripheral visual signals are noisier than foveal ones (Frisen and Glansholm, 1975; Wright and Johnston, 1983), the finding that subjects weight proprioceptive feedback more heavily in the control study is consistent with a minimum-variance strategy (see Discussion).

Control for active movement of the target arm

In the original version of Experiment 1, proprioceptive targets were specified by the position of the left index fingertip. However, since subjects actively moved the left hand to the target location and supported it against gravity once it was in place, non-proprioceptive motor cues such as efference copy could have contributed to the sensed position of the left hand. We therefore tested 5 additional subjects in a second modified version of Experiment 1 aimed at minimizing non-proprioceptive cues to target location. Specifically, the left arm was passively moved to the proprioceptive target on each trial and was supported by a second tabletop (Figure 20a). The left

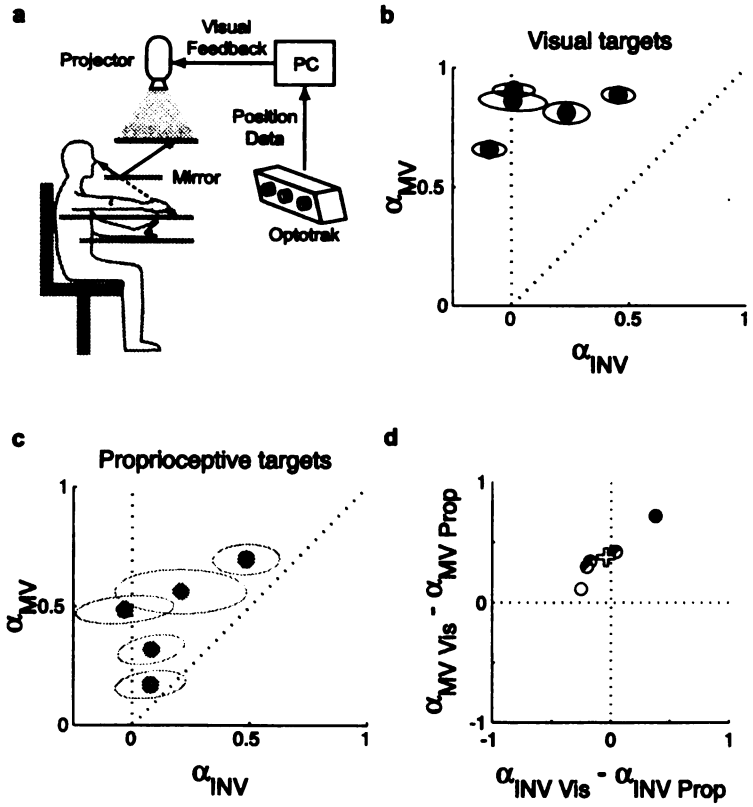

Figure 20: Methods and model fits, passive proprioceptive experiment. In this study, the left arm rested on a sled (not shown) supported by the lower table (a) and was passively moved to the appropriate position to serve as a proprioceptive target. The left index fingertip touched the upper table, which supported the right arm. Results are shown in b-d, plotting conventions as in Figure 15.

forearm rested on a specially designed sled with an adjustable brace that fixed the left wrist and held the left index fingertip in a raised position against the underside of the upper table. The right arm rested on the top of the upper table. Subjects received no information about where to move the left hand. Rather, an experimenter stood to the subject's left and positioned the left index fingertip at the appropriate proprioceptive target on each trial. This positioning was done based on visual cues available to the experimenter but hidden from the subject, who was instructed to relax the left arm while the experimenter positioned it.

The results of this modified experiment are shown in Figure 20b-d. As in the original experiment, all subjects had a larger fit value of α_{MV} in the visual-target condition (mean difference in $\alpha_{MV} = .38$, Figure 20d). This result shows that the pattern of sensory integration during movement vector planning observed in the original experiment did not depend on the subject making active movements with the left (target) arm. The average value of α_{INV} in the modified experiment was very similar in both conditions (mean difference in $\alpha_{INV} = .04$, larger for proprioceptive targets).

Hypothesis testing and confidence limits.

We used a bootstrapping technique (Efron and Tibshirani, 1993) to determine confidence limits on the fit values of α_{MV} and α_{INV}. For each subject and trial type (e.g. visual-target vs. proprioceptive-target trials in Experiment 1), we created 1000 datasets in which each datum was resampled (with replacement) from one of the eight trials with the same reach target and visual shift. The model was fit to each of these resampled datasets, resulting in a distribution of 1000 values of $[\alpha_{MV}, \alpha_{INV}]$ which was used to find the confidence ellipse for the values of α_{MV} and α_{INV} fit to the true (un-resampled) dataset.

To test the hypothesis that the value of α_{MV} or α_{INV} differed between two trial types, we used a permutation test against the null hypothesis that the mixing parameter in question was equal in the two conditions. For example, in order to test the hypothesis that α_{MV} differs when a subject reaches to visual versus proprioceptive targets, we tested against the null hypothesis $H_0 : \alpha_{MV1} = \alpha_{MV2}$, where α_{MV1} is the value of α_{MV} in the visual-target condition and α_{MV2} is the value of α_{MV} in the proprioceptive-target condition. This test is derived from the following

rearrangement of Equations 18 and 19 from the main text:

$$\hat{x}_{\text{MV, i}} = P_i + \alpha_{\text{MV, c(i)}}(V_i - P_i) \quad (18')$$

$$\hat{x}_{\text{INV, i}} = P_i + \alpha_{\text{INV, c(i)}}(V_i - P_i) \quad (19')$$

where V_i and P_i represent the visual and proprioceptive signal on the i^{th} trial, $c(i)$=1 for visual-target trials, and $c(i)$=2 for proprioceptive-target trials. Combining Equations 18′ and 19′ with Equation 20 - Equation 22 from the main text, we can obtain a single expression for the predicted initial velocity in all conditions:

$$\dot{x}_i = \underbrace{J\left(K^{-1}(P_i)\right) J^{-1}\left(K^{-1}\left(P_i + \alpha_{\text{INV, c(i)}}(V_i - P_i)\right)\right)}_{\Phi_i}\left(x_i^* - P_i - \alpha_{\text{MV, c(i)}}(V_i - P_i)\right).$$

This equation can be rewritten using the indicator function $\delta_{c(i)=k}$, which equals 1 if trial i is of type k and zero otherwise:

$$\dot{x}_i = \Phi_i\left[x_i^* - P_i - \left[\alpha_{\text{MV1}}\delta_{c(i)=1} + \alpha_{\text{MV2}}\delta_{c(i)=2}\right](V_i - P_i)\right] \quad (23)$$

According to the null hypothesis H_0 : $\alpha_{\text{MV1}} = \alpha_{\text{MV2}}$, the model will produce the same predictions if the α_{MV1} and α_{MV2} are randomly interchanged across trials. This manipulation is achieved by permuting (randomly re-sorting) the values of $c(i)$ shown explicitly in Equation 23, leaving the un-permuted values of Φ_i. To perform the permutation test, we created 1000 artificial datasets in this manner, fit the model to each dataset, and computed the resulting R^2. If the R^2 value for the true dataset is greater than the 95^{th} percentile of artificial R^2 values, then H_0 can be rejected. The null hypothesis H_0' : $\alpha_{\text{INV1}} = \alpha_{\text{INV2}}$ was tested in a similar fashion.

Chapter 3: Sensory Integration following Sensorimotor Adaptation

Abstract

When planning reaching movements, subjects combine visual and proprioceptive feedback in order to estimate the arm's initial position. This sensory integration problem is solved at two different stages of reach planning: first when computing the desired movement direction, and second when converting that vector into a motor command. In this study, we explore how sensorimotor adaptation affects these two planning stages. Subjects reached to visual targets after adapting to shifted visual feedback. In one trail type, visual feedback consisted of a spot of light that followed the movement of a subject's fingertip. In a second trial type, subjects received no online feedback but pointed to a spot of light prior to the initiation of each reach. In both trial types, subjects planned reaching movements using displaced visual feedback, either because of shifts in the online visual feedback or because of errors in the reach to the initial visual landmark. Analysis of the resulting reach errors suggests that sensorimotor adaptation drives changes in the movement vector planning stage, but that the second stage of reach planning is unaffected. Additionally, we ask whether the relative weighting of vision and proprioception is influenced by the availability of online visual feedback, a factor known to influence the extent of sensorimotor adaptation. We found that on average subjects weighted visual signals equally in the two trial types, suggesting that a stationary visual landmark can influence arm position estimation to the same extent as continuous visual feedback.

INTRODUCTION

When planning a reaching movement, subjects integrate vision and proprioception to estimate the position of the arm. In the course of planning a single movement, two position estimates are formed. First, an integrated estimate is used to plan the desired reach direction, or movement vector. Then, a second position estimate is used to transform the movement vector into a joint- or muscle- based (intrinsic) motor command (Sober and Sabes, 2003). We have demonstrated the flexibility of sensory integration at both of these stages. Manipulations of the reaching task have shown that the reliance on vision versus proprioception at the first planning stage is strongly influenced by the coordinate frame of the target and that sensory integration at the second stage is significantly affected by the information content of the visual feedback (Sober and Sabes, 2005).

In our previous work, we quantified sensory integration using acute shifts of visual feedback to dissociate visual and proprioceptive hand position signals. These shifts lead to errors at both stages of reach planning by displacing the estimated hand position away from the true hand location. At the first stage of reach planning, shifted visual feedback induces errors in the planned movement vector, or "MV error". At the second stage, feedback shifts lead to errors in generating the motor command necessary to achieve the desired movement vector. This inverse model or "INV error" consists of counterclockwise changes in initial movement direction when visual feedback is shifted rightward and clockwise changes when feedback is shifted leftward (Goodbody and Wolpert, 1999). Since the magnitudes of MV and INV error are proportional to the reliance on vision at each planning stage, we were able to infer the weighting of vision and proprioception from the changes in initial reach direction caused by shifted visual feedback (Sober and Sabes, 2003, 2005).

Here we explore how vision and proprioception are combined following sensorimotor adaptation, when a persistent shift of visual feedback leads to lasting changes in motor planning and a realignment of the visual and proprioceptive maps of the workspace (Helmholtz, 1925; Welch, 1986).

We quantify sensory integration following adaptation to shifted visual feedback in order to address three questions. First, we ask whether the process of sensorimotor adaptation differentially affects the two stages of motor planning. If movement vector planning and inverse model evaluation are truly independent planning stages, it seems reasonable that the process of adaptation might affect them independently. By examining the effects of adaptation on MV and INV error, we hope to further demonstrate the modularity of the planning process. Subjects adapted to shifted visual feedback and then performed center-out reaches under two conditions. In one case, subjects made reaches with continuous visual feedback of their fingertip using the same visual feedback shift they experienced during training ("with-feedback reaches"). In a second condition, subjects began by reaching to a spot of light at the center of the workspace. These initial movements were inaccurate, due to the aftereffects of adaptation. This aftereffect led to a dissociation of vision and proprioception analogous to that experienced in with-feedback reaches. From this starting point, subjects then reached to a target in the continued absence of online visual feedback ("no-feedback reaches"). By comparing the magnitudes of MV and INV error in both trial types with the results of our previous studies (in which these errors were evoked in the absence of adaptation), we are able to determine the relative effects of sensorimotor adaptation on the two planning stages.

Second, comparing sensory integration in with- and no-feedback reaches to that found in previous studies allows us to ask whether the reach errors experienced during adaptation have an effect on sensory integration. In both our current and previous

studies, we induced reach errors by altering visual feedback. Error on any given trial is therefore proportional to the reliance on vision. In our previous studies, subjects were given no visual feedback about the accuracy of their reaching movements. In the current study, however, error information is provided in order to drive sensorimotor adaptation. Subjects in the current experiment could therefore use error feedback to drive adaptive changes in sensory integration, decreasing their reliance on vision in order to reduce the size of reaching errors. Using a previously described mathematical model of motor planning (Sober and Sabes, 2003), we quantify the relative reliance on vision and proprioception during motor planning. Comparing the inferred weighting of vision and proprioception in the present study with that found previously allows us to ask whether such a strategy is in fact employed.

Third, we ask whether the presence or absence of online visual feedback influences sensory integration. Studies of sensorimotor adaptation have shown that this factor influences the extent to which visual shifts lead to learning (Foley and Maynes, 1969; Cohen, 1967). We therefore investigate whether the availability of continuous visual feedback affects sensory integration as well. Although no-feedback and with-feedback reaches both began with a discrepancy between visual and proprioceptive signals, visual feedback represented the full motion of the fingertip only in the with-feedback case. Fitting the same model to data from both trial types allows direct comparison of the weighting of vision and proprioception in these two conditions.

METHODS

All subjects (4 female, 5 male) were healthy, right-handed, had normal or corrected-to-normal vision, and ranged from 24-46 years of age. Subjects were paid for their participation and were naive to the purpose of the experiments. Reach targets and visual feedback of the arm were presented using a virtual-reality display system described previously (Sober and Sabes, 2003). Subjects performed all reaching movements with the right arm, which rested on a shoulder-height table. The wrist was immobilized by a splint, which also fixed the right index finger in an extended position. Friction between the right arm and the table was minimized by air sleds under the upper arm and forearm. The left arm rested comfortably below the table, and the left hand held a mouse which subjects clicked as part of certain tasks (see Task Design below). The torso was immobilized with a padded harness, and a drape prevented direct vision of the arm and shoulder. Movement data were collected using an infrared tracking system (OPTOTRAK, Northern Digital, Waterloo, Ontario). Infrared-emitting diodes attached to the arm and shoulders allowed us to compute the angles of the shoulder and elbow joints.

Task Design

The workspace consisted of eight targets equally spaced along an imaginary circle of radius 10 cm surrounding a ninth, central target. Each target (when illuminated) was a spot of radius 5 mm. The central target was positioned approximately 40 cm from the subject's chest and slightly to the right of the midline.

Block design. Each subject participated in two sessions, which were separated by at least 24 hours. Each session was composed of three blocks – Baseline, Workspace Control, and Visual Shift. In Baseline blocks, visual feedback was veridi-

cal. In Visual Shift blocks, the visual feedback was shifted to the right or left (one direction in each session) by 120 mm. In Visual Shift blocks, the visual workspace was also shifted by the same amount as the visual feedback. The effect of this was that the actual workspace (that is, the location of the reaches in table-centered coordinates) was not altered, although the apparent (visual) workspace changed. As a control for this shift in apparent location, Workspace Control blocks were included in which visual feedback was veridical, but the workspace was shifted to match the apparent workspace in that session's Visual Shift block. Blocks consisted of varying combinations of six different trial types (training, with-feedback testing, no-feedback testing, no-feedback transition, visual midline, and proprioceptive midline), which were presented in sets.

Training trials. A training set (Figure 21A) consisted of 49 reaches between the central target and the eight peripheral targets. Each training trial began in the absence of visual feedback with the appearance of one of the nine targets. Subjects were then required to move their index finger towards the target. When the fingertip had moved 75% of the distance to the target, the subject began to receive visual feedback. Visual feedback in this and all other trial types consisted of a white circle of light 6 mm in diameter. Subjects completed their reach in the presence of visual feedback, and the target and feedback spot were extinguished after the spot was held within the target for 150 msec. We selected this schedule of visual feedback after a preliminary study indicated that this method produced more robust adaptation aftereffects than continuously illuminated visual feedback (data not shown).

With-feedback trials. A with-feedback testing set (Figure 21B) consisted of eight center-out reaches to the peripheral targets in pseudorandom order. The trial began when the central target appeared as a green circle. Subjects positioned the feedback spot on the central target and clicked a button on the mouse held in the

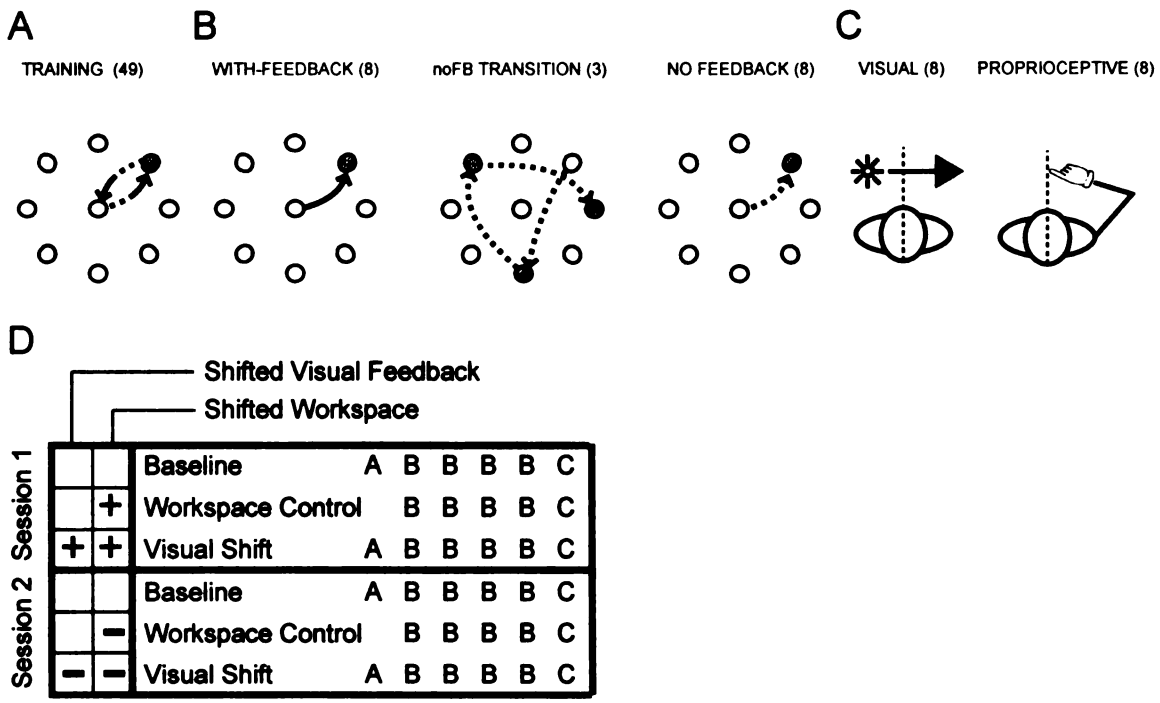

Figure 21: Task design. In each experimental session, subjects performed six different tasks. (A) In a set of training trials, subjects performed 49 reaches and received visual feedback at the end of each movement. (B) During with-feedback trials (8 trials/set), visual feedback was available throughout the movement. A set of 3 no-feedback transition trials followed each set of with-feedback trials. In these trials, subjects reached to three targets in the absence of visual feedback. During no-feedback trials (8 trials/set), subjects reached to the central start point in the absence of visual feedback, and then reached to one of the peripheral targets. (C) In visual midline and proprioceptive midline trials, subjects used only a single sensory modality to estimate their midline on the table surface. (D) Each session included a Baseline, Workspace Control, and Visual Shift block, made up of the trial types described in (A-C), as shown. In Visual Shift and Workspace Control blocks, the nine visual targets were shifted rightward or leftward ("+" or "-") by 120 mm. In Visual Shift blocks, visual feedback from the fingertip was shifted by the same amount. Visual feedback was veridical in Workspace Control blocks.

left hand. A blue target circle then appeared in one of the eight peripheral locations. After a delay of 500-1500 msec, the target turned purple, which cued the subject to reach to the target. The start point disappeared as soon as the movement was initiated. The velocity of the fingertip was monitored online, and the trial ended when either the visual feedback spot had been held at the target for 100 msec or if fingertip speed dropped below 1.2 mm/sec. When either of these conditions were met the target disappeared and the green start point reappeared to begin the next trial.

No-feedback trials. No-feedback testing sets were identical to with-feedback sets except that subjects did not receive visual feedback from the finger. Subjects were therefore required to estimate the location of the central target at the beginning of every trial. The estimated start points of the no-feedback reaches – that is, the location of the fingertip when subjects clicked the mouse button to begin each trial – were used to measure the magnitude of the sensorimotor aftereffects induced by the visual shifts (see Data Analysis below).

No-feedback transition trials. Preliminary studies indicated that when no-feedback testing trials directly followed with-feedback trials, subjects showed an aftereffect magnitude equal to the magnitude of the visual shift (100% adaptation) on the first movement to the central start point. Since adaptation to shifted visual feedback typically has a significantly smaller magnitude (Welch, 1986; Clower and Boussaoud, 1998), this suggested that the subject might be relying on visual feedback from the previous trial to estimate the first no-feedback start point. For this reason, a set of three transition trials were introduced between the with-feedback and no-feedback trial sets. In each transition trial, a pseudorandomly selected target was illuminated and subjects were required to point, without feedback, to their best estimate of the target location and click the mouse button.

Visual and proprioceptive midline trials. At the conclusion of each block, subjects completed two additional sets of single-modality trials intended to determine the contributions of visual and proprioceptive recalibration to the start point estimation aftereffect described above (Figure 21C). In the first set ("visual midline"), a small white dot moved slowly across the workspace. Subjects kept their reaching hand stationary and clicked the mouse button when they judged that the dot was crossing the midsaggital plane (midline). This was repeated eight times, and the dot began its trajectory from alternating sides on consecutive trials. In a second set ("proprioceptive midline"), also consisting of eight trials, subjects pointed at their estimate of the midline at the approximate distance from the body that the white dot had traversed in the visual midline trials. This estimation was done in the absence of visual feedback. Between the proprioceptive midline trials subjects moved their reaching hand to alternating sides of the workspace before making the next estimate. Prior to beginning each experimental session, subjects were familiarized with all six trial types. This familiarization was completed in the absence of any visual shifts.

Data Analysis and Model Fitting

Aftereffect calculations. Aftereffects were measured in three different ways. First, reach aftereffects were quantified by measuring the change in subjects' estimates of the position of the central start point in no-feedback testing trials. Reach aftereffects were separated into a component parallel to the direction of the visual feedback shift and a component in the orthogonal direction. A subject was judged to have a start point aftereffect if there was a significant change between the Visual Shift and Baseline blocks (t-test, $p<.05$) in the component parallel to the visual feedback

shift. Single-modality (visual midline and proprioceptive midline) aftereffects were measured in the same way, except that aftereffects were only assessed along the direction parallel to the shift.

Since the topic of this study was the effects of adaptation on motor planning, and not adaptation *per se*, we applied a strict subject selection criterion based on the aftereffect data. A subject was judged to have adapted successfully if in *both* experimental sessions they showed a significant aftereffect in both the start point *and* one of the two single-modality tests. Since previous work has shown that such single-modality effects accompany prism adaptation (Harris, 1963; Hay and Pick, 1966; Uhlarik and Canon, 1971) we established the second requirement to reject data from subjects that may have been altering their start point estimates as part of a conscious strategy to compensate for the shifted visual feedback or as a result of procedural, task-related learning. If a subject did not meet the two requirements in one or both sessions their data was not included in the final analysis.

Trajectory analysis. Raw trajectory data were smoothed with a low-pass Butterworth filter with a cutoff frequency of 6 Hz. Average trajectories for each target and feedback condition were calculated by designating the first and last times the tangential velocity equaled 10% of the peak velocity for each individual trial. The velocity trace between these landmarks was then resampled using cubic spline interpolation into 50 points evenly spaced in time. The interpolated trajectories were then averaged to compute the mean trajectory for each subject, target and feedback condition. These averaged trajectories were used for display only and were not used in any data analysis.

Initial direction. We quantified initial reach direction by choosing two early landmarks along the reach trajectory and determining the average velocity between them. The first and second landmarks were the first points along the trajectory

where the speed exceeded 40 and 60% of the peak speed. These markers nearly always fell within the first centimeter of the reach path. The results presented here are robust to variations in these landmark percentages between 20-70% of the peak.

Model fitting. The data from each subject were fit with a model of feedforward motor planning. In earlier work, we have developed a technique that allows us to quantify how vision and proprioception are combined at each planning stage (Sober and Sabes, 2003, 2005). In these earlier experiments, subject planned reaching movements with shifted visual feedback. These shifts introduced error at the movement vector and inverse model planning stages (MV error and INV error, respectively). By measuring the magnitude of these errors as a function of target direction, we were able to fit the values of α_{MV} and α_{INV}, corresponding to the extent to which the subject relied on visual feedback at each stage.

In our previous work, the visual shifts varied randomly from trial to trial, and subjects never received feedback about the resulting reach errors (Sober and Sabes, 2003, 2005). As a result, visuomotor adaptation did not occur. In this study, however, subjects received error feedback in the training and with-feedback testing trials. Subjects adapted to the consistent visual shifts and displayed the type of reaching aftereffects described in classic studies of prism adaptation (Helmholtz, 1925; Welch, 1986). As described below, we found that the process of sensorimotor adaptation led to an elimination of MV error. We therefore did not attempt to quantify α_{MV} (see Discussion). However, INV error persists after visuomotor adaptation, and we quantified α_{INV} as in our previous work.

For each subject, data from the with- and no-feedback conditions were fit separately. In the model, visual and proprioceptive signals are integrated to yield a combined estimate of arm position. Arm position is interchangeably expressed in Cartesian workspace coordinates (x) or in terms of joint angles (θ), related by the

kinematics equation $x = K(\theta)$. Note that the kinematics equation is invertible for the restricted, two degree of freedom movements made in this experiment. A mixing term α_{INV} parameterizes the relative contributions of these two modalities to the final estimate:

$$\hat{x}_{\text{INV}} = \alpha_{\text{INV}}\,\hat{x}_{\text{vis}} + (1 - \alpha_{\text{INV}})\,\hat{x}_{\text{prop}}. \tag{24}$$

The direction of the planned movement vector is described by

$$\measuredangle \dot{x}^* = \measuredangle(x_d^* - x_d^{\text{start}}) + \omega_d, \tag{25}$$

where $\measuredangle z$ represents the angle of any vector z, x^* represents the target location, x^{start} represents the start point location, ω is an angular offset term that captures directional biases in baseline reaches, and $d \in [1, ..., 8]$ indexes x^*, x^{start} and ω over the eight targets.

The desired fingertip velocity $\dot{x}^*$ is transformed into a motor command, which we model as a vector of joint angle velocities $\dot{\theta}$:

$$\dot{\theta} = J^{-1}(\hat{\theta}_{\text{INV}})\dot{x}^*, \tag{26}$$

where $\hat{\theta}_{\text{INV}} = K^{-1}(\hat{x}_{\text{INV}})$ and $J(\theta) = \frac{dx}{d\theta}$ is the arm Jacobian. This planning stage describes an evaluation of an inverse model of the arm (Jordan, 1996).

When this command is implemented by the arm, the change in endpoint is related to the change in joint angles by the Jacobian:

$$\dot{x} = J(\theta)d\theta = J(\theta)J^{-1}(\hat{\theta}_{\text{INV}})\dot{x}^* \tag{27}$$

Note that if the initial position estimate is accurate ($\hat{\theta}_{\text{INV}} = \theta$), the arm movement will be along the desired movement vector. If the shifted visual feedback results in a discrepancy between θ and $\hat{\theta}_{\text{INV}}$, however, the initial velocity $\dot{x}$ will differ from the desired velocity $\dot{x}^*$, resulting in INV error.

The model was fit to the data from each subject using α_{INV} as the sole free parameter. The ω_d for each of eight target directions was set to equal the average angular offset in the Baseline block for that target and experimental session. The fit datasets consisted of arm positions at the beginning of each movement ($\hat{x}_{\text{prop}}$ in Equation 24), the visual locations of the start point ($\hat{x}_{\text{vis}}$ in Equation 24), which in with-feedback trials was coextensive with the visual feedback spot, and the initial directions of movement ($\dot{x}$ in Equation 27) for with-feedback and no-feedback testing trial sets in the Baseline and Visual Shift blocks. The value of α_{INV} was fit separately to datasets from with- and no-feedback trials using a nonlinear regression algorithm (MATLAB, The Mathworks Inc., Natick, MA). Each fit dataset consisted of 128 trials (4 blocks $\times$ 8 targets $\times$ 4 repetitions).

Alternate model analysis The above model assumes that the visual and proprioceptive signals combined to estimate arm position in are unbiased. In Sober and Sabes (2003) we show that the model is able to fit the correct value of α_{INV} even in the presence of constant biases of $\hat{x}_{\text{vis}}$ and $\hat{x}_{\text{prop}}$, provided that the biases are the same when visual feedback is shifted leftward and rightward. However, sensorimotor adaptation leads to aftereffects in visual midline and proprioceptive midline trials that differ based on the direction of the visual shift (see Figure 22).

It is unclear whether these behavioral biases reflect changes in $\hat{x}_{\text{vis}}$ and $\hat{x}_{\text{prop}}$ or whether they result from some process separate from those described in the preceding equations. To test how biased values of $\hat{x}_{\text{vis}}$ and $\hat{x}_{\text{prop}}$ might affect the fit value of α_{INV}, we performed an alternate analysis in which Equation 24 was modified to take

into account visual and proprioceptive biases:

$$\hat{x}_{\mathrm{INV}} = \alpha_{\mathrm{INV}}\,(\hat{x}_{\mathrm{vis}}+\Delta_{\mathrm{V_s}})+(1-\alpha_{\mathrm{INV}})\,(\hat{x}_{\mathrm{prop}}+\Delta_{\mathrm{P_s}}). \qquad (24')$$

In this modified equation, $\Delta_{\mathrm{V_s}}$ and $\Delta_{\mathrm{P_s}}$ represent the biases in vision and proprioception resulting from sensorimotor adaptation to visual feedback shifted in direction s. Fitting the values of $\Delta_{\mathrm{V_s}}$ and $\Delta_{\mathrm{P_s}}$ (in addition to α_{INV}) would result in a model with more degrees of freedom than could be constrained by the data. We therefore estimated these terms from the shift-induced changes in subjects' estimates of the midline:

$$\begin{aligned} \Delta_{\mathrm{V_s}} &= -(\bar{\mathrm{V}}_{\mathrm{s}} - \bar{\mathrm{V}}_{\mathrm{bl}}) \\ \Delta_{\mathrm{P_s}} &= -(\bar{\mathrm{P}}_{\mathrm{s}} - \bar{\mathrm{P}}_{\mathrm{bl}}). \end{aligned}$$

The terms $\bar{\mathrm{V}}_{\mathrm{s}}$ and $\bar{\mathrm{P}}_{\mathrm{s}}$ are the mean visual and proprioceptive midline estimates following a shift in direction s, and $\bar{\mathrm{V}}_{\mathrm{bl}}$ and $\bar{\mathrm{P}}_{\mathrm{bl}}$ are the appropriate baseline values. The sign reversal reflects the fact that a leftward shift in a perceptual map will lead to a rightward shift in the estimate of the subject's midline, and vice versa.

Significance testing and confidence intervals. We determined whether each fit value of α_{INV} was significantly different from zero using a permutation test (Good, 2000). A rearrangement of Equation 24 yields

$$\hat{x}_{\mathrm{INV}} = \hat{x}_{\mathrm{prop}} + \alpha_{\mathrm{INV}}(\hat{x}_{\mathrm{vis}} - \hat{x}_{\mathrm{prop}}).$$

Under the null hypothesis $\alpha_{\mathrm{INV}} = 0$, the initial direction of movement is independent of $(\hat{x}_{\mathrm{vis}} - \hat{x}_{\mathrm{prop}})$. By permuting the trial identities from which this difference term is taken, we can break any existing dependence, creating an artificial dataset that

obeys the null hypothesis. By creating 1000 such datasets and fitting α_{INV} to each of them, we created a distribution of α_{INV} under the null hypothesis. We rejected the null hypothesis for the actual dataset if the best-fit α_{INV} was greater than the 95^{th} percentile of this distribution.

Additionally, we used a bootstrapping technique (Efron and Tibshirani, 1993) to put confidence limits on the fit values of α_{INV}. We created 1000 datasets where all data for each trial were resampled (with replacement) from one of the four trials in the same block with the same target direction and feedback condition. The value of α_{INV} was then fit to each synthesized dataset. The resulting distribution of α_{INV} was used to find the confidence limits of the fit α_{INV} from the actual datasets.

Predicted changes across the workspace. To compute the average change in initial reach direction predicted by the model at each point in the workspace (shown in Figure 26), we used the fit values of α_{INV} and the arm Jacobian for a given subject to calculate the matrix $M = J(\theta)J^{-1}(\hat{\theta}_{\text{INV}})$ at an array of points across the workspace. The average predicted change at each workspace location was calculated as the mean angular change induced by M across the unit circle.

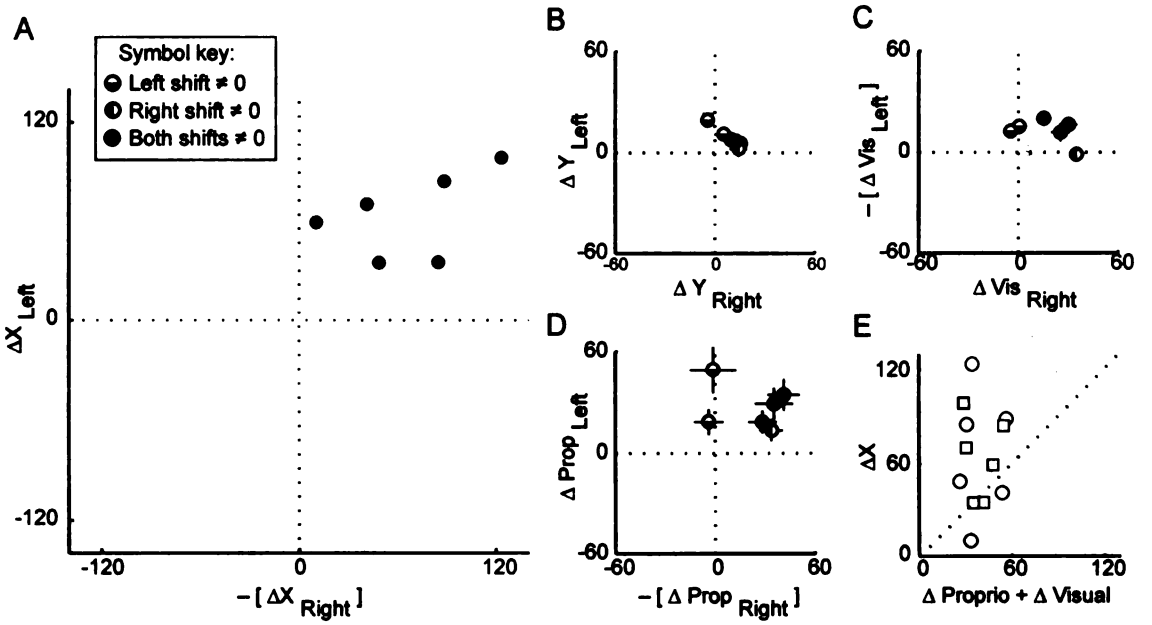

Figure 22: Aftereffects. Each symbol plots the mean ± 1 S.E. aftereffect (Visual Shift minus Baseline) from a single subject after rightward and leftward shifts of visual feedback. Errorbars are hidden by the symbols in many cases. (A) shows the component of start point aftereffects parallel to the direction of the visual feedback shifts. (B) shows the aftereffect component in the orthogonal direction. (C) and (D) show the aftereffects in visual midline and proprioceptive midline trials, respectively. Symbols in (A) through (D) indicate the results of t-tests ($p<.05$) to determine whether aftereffects are significant. Positive values in (A), (C), and (D) indicate rightward aftereffects, however note that values on the abscissa (A and D) and ordinate (C) have been sign-reversed so that data from subjects showing the expected aftereffects fall in the upper right quadrant. (E) shows aftereffect magnitude plotted against the combined magnitudes of the visual and proprioceptive-only aftereffect recorded in a single experimental session. Sessions including leftward visual feedback shifts are plotted as squares, and sessions including rightward shifts are plotted as circles.

RESULTS

Aftereffects All 9 subjects showed significant start point aftereffects in both experimental sessions. In the single-modality trials, however, only 6 subjects showed significant aftereffects in either the visual midline or proprioceptive midline trials in both experimental sessions. Using the criterion outlined in Methods, these 6 datasets were selected for full analysis. Aftereffect data for these selected subjects

are shown in Figure 22. Adaptation to leftward and rightward visual shifts produced large aftereffects along that axis (Figure 22A) in the direction opposite the visual shift. When averaged across the two experimental sessions completed by each subject, the magnitude of start point aftereffects ranged from 29-93% of the magnitude of the visual feedback shift. Small aftereffects were sometimes seen in the direction orthogonal to the axis of shift (Figure 22*B*). Additionally, the single-modality trials revealed a recalibration of the visual and proprioceptive estimates of the subjects' midline, as previously reported (Hay and Pick, 1966; Welch et al., 1974). In visual midline tests, the visually perceived straight-ahead direction was perturbed in the direction of the shift in visual feedback (Figure 22C). Proprioceptive midline trials showed that the felt position of the midline was shifted in the opposite direction (Figure 22D). Figure 22E shows that the magnitude of start point aftereffects was not predicted by the sum of visual midline and proprioceptive midline aftereffects. This finding is consistent with the findings of Welch et al. (1974), although Hay and Pick (1966) found that reach aftereffects were well-predicted by the sum of effects observed on trials using single modalities.

Initial Direction Reach trajectories from a single subject in the with-feedback condition are shown in Figure 23A and B. Leftward and rightward and visual shifts produced clockwise and counterclockwise changes in the initial direction of reach, the pattern typical of INV error. This rotational pattern is not consistent with errors in movement vector planning, which is characterized by a bimodal pattern of changes in initial direction (see Figure 1, Chapter 1). Because subjects received continuous visual feedback throughout the movement, they fully corrected their initial errors in the later parts of the reach trajectory. The clockwise/counterclockwise pattern of changes was consistent over most, but not all, target directions for this subject (Figure 23C).

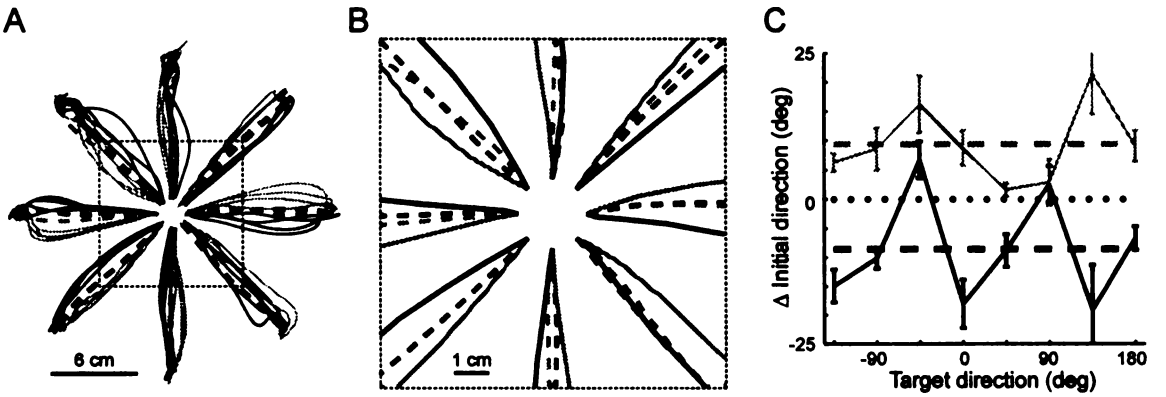

Figure 23: Reach trajectories and initial reach directions from a representative subject, with-feedback trials. Although all trajectories begin at the same central start point, trajectories have been offset by 1 cm in the direction of the target for clarity. Narrow lines in (A) show individual reach trajectories during leftward (black) and rightward (gray) shifts of visual feedback. Heavy dashed lines show mean trajectories in the appropriate Baseline condition. (B), detail of the initial portion of the movement (box in A). Heavy solid lines in (B) show the mean trajectory during leftward and rightward shifts. (C) Shift-induced changes in initial direction (Visual Shift minus Baseline) as a function of target direction. A target direction of 0° corresponds to the rightward target, and the target furthest from the subject is at 90°. Positive and negative values on the ordinate correspond to counterclockwise and clockwise errors, respectively. Errorbars, ± 1 S.E. Dashed lines represent the mean error across target directions for a given shift. Data are from subject tr.

This pattern of shift-induced changes in initial direction was observed in every subject. To investigate whether this pattern was statistically significant, changes in initial reach direction were pooled across the eight target directions and subjected to a t-test. For all subjects, directional changes resulting from the leftward and rightward shifts were significantly different ($p < .05$). Figure 24A shows the mean change in initial reach direction in with-feedback reaches. Rotational errors were not seen in the Workspace-Control blocks (Figure 24*C* and *D*), in which the start point and targets were shifted but visual feedback was veridical. Changes in initial direction in the Workspace Control blocks were not significant in any dataset in the with-feedback condition, and differed significantly in only one of six datasets in the no-feedback condition. These results indicate that the pattern of clockwise and counterclockwise changes observed in the Visual Shift blocks was not due to the change in apparent workspace.

Data from no-feedback trials reflected a similar pattern of changes in initial direction. Subjects began each no-feedback reach by pointing to the visual start point (Figure 25A, open circles) in the absence of visual feedback. Each reach therefore began from a different starting point. Aligning the start points (Figure 25B and C) shows that reaches in the no-feedback condition display the same clockwise-counterclockwise pattern of changes seen in with-feedback reaches (compare to Figure 23A and B).

The direction and magnitude of the disparity between the signals provided by vision (the start point) and proprioception (the hand location) in the no-feedback condition varied from trial to trial. In our model, the position estimate $\hat{x}_{\mathrm{INV}}$ lies along the line connecting the two position signals (Equation 24). Furthermore, the predicted error on any given trial depends on both the actual and estimated hand locations (Equation 27). There was therefore a moderate amount of trial-to-trial

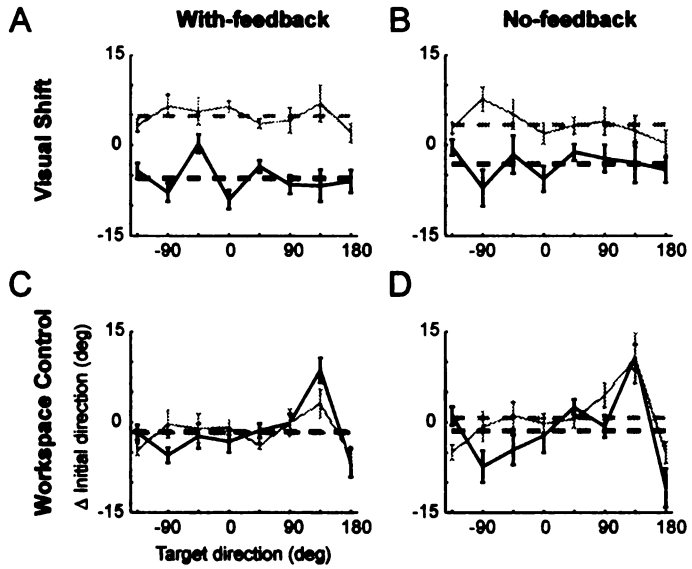

Figure 24: Visual shift-induced directional errors, averaged across subjects. (A) and (B) show the mean ± 1 S.E. change in initial direction in the Visual Shift blocks relative to Baseline blocks. (C) and (D) show the same changes in the Workspace Control blocks relative to the Baseline blocks. Data from the with-feedback condition are shown in parts (A) and (C), and results from the no-feedback trials are shown in (B) and (D). Other plotting conventions as in Figure 23C.

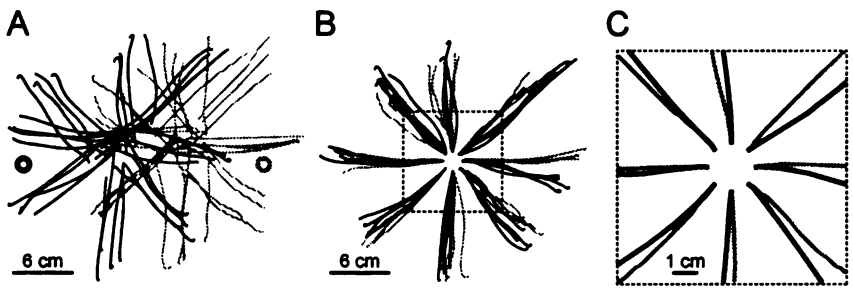

Figure 25: Reach trajectories from a representative subject, no-feedback trials. Trajectory data are shown in (A). Center-out reaches begin after subjects reach to a visual start point (circles) as explained in Methods. Trajectories and visual start points used during leftward and rightward shifts of visual feedback are shown in black and gray, respectively. In (B), reaches to each target are aligned at their starting points. Average initial trajectories (from the area enclosed by the box in B) are shown in (C). Data are from subject wi.

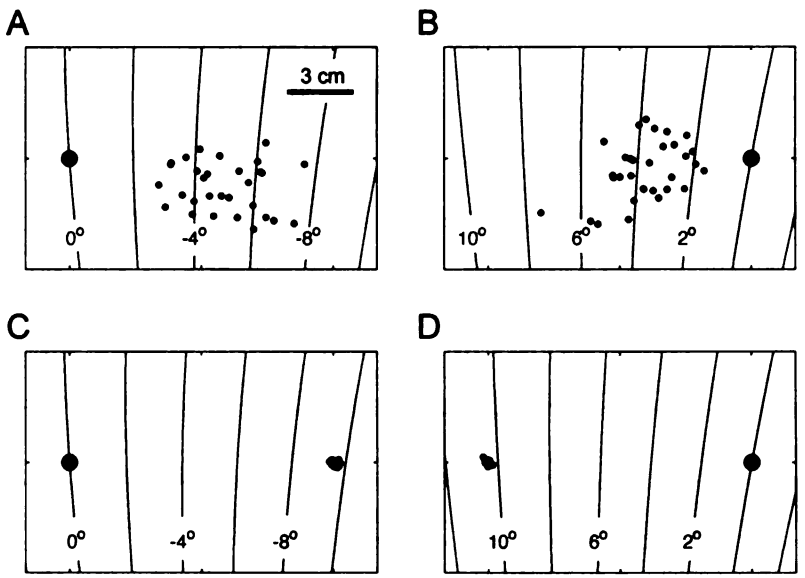

Figure 26: Predicted changes in reach direction from a single subject. The locations of reach start points (black dots) in no-feedback (A and B) and with-feedback (C and D) trials are shown for leftward (A and C) and rightward (B and D) shifts of visual feedback. Gray dots show the locations of the visual start points. Contour lines show the average change in initial reach direction predicted by the model for reaches begun at each location in the workspace. Positive values represent counterclockwise changes. Data and error predictions are from the reach start position data, arm dimensions, and model fits ($\alpha_{\mathrm{INV}} = .82$, with-feedback; $\alpha_{\mathrm{INV}} = .73$, no-feedback) for subject tr.

variability in the predicted change in initial reach direction for no-feedback trials (Figure 26*A* and *B*), whereas the predicted effects of visual shifts were of a fairly constant size in with-feedback trials (Figure 26*C* and *D*). Despite this variation, however, data from no-feedback trials showed similar patterns of directional error as those from with-feedback reaches (Figure 24B). This difference was significant in all 6 subjects.

Model Fitting We fit the model separately to the with- and no-feedback data from each subject as described in Methods. The values of ω_d (average baseline bias) were typically small, indicating that the planned movement vector ($\dot{x}^*$ in Equation 25 - Equation 27) was approximately parallel to the start-target vector. The mean $\pm$ 1 S.D. values of ω_d were $1.3 \pm 4.6°$ in the with-feedback condition and $0.7 \pm 7.8°$ in

the no-feedback condition

Despite the difference in the availability of online visual feedback provided in the two conditions, the model fits showed that subjects used a similar mixture of vision and proprioception in with- and no-feedback trials (Figure 27A). Although some subjects showed substantial differences in sensory integration in the two trial types, the average value of α_{INV} was similar in the two feedback conditions (mean $\alpha_{INV} = .40$, with feedback; $\alpha_{INV} = .45$, no feedback). Permutation tests (see Methods) revealed that all fit values of α_{INV} in both feedback conditions were significantly different from zero. This allows us to categorically reject the hypothesis that subjects were relying exclusively on proprioception when estimating the location of their arm. Fit values of α_{INV} in the two feedback conditions did not correlate with the magnitude of sensorimotor adaptation or the size of single-modality aftereffects. Additionally, we fit the model to data from the three subjects that had initially been eliminated from consideration for failing to meet our criterion for sensorimotor adaptation. These three subjects had substantially lower values of α_{INV} than the six who adapted (mean $\alpha_{INV} = .24$, with feedback; $\alpha_{INV} = .19$, no feedback). However, including these subjects did not qualitatively change our results. As before, values of α_{INV} were comparable across feedback conditions (mean $\alpha_{INV} = .34$, with feedback; $\alpha_{INV} = .35$, no feedback).

So far, our analysis has assumed that visual and proprioceptive signals are unbiased. That is, we have assumed that the visual and proprioceptive signals combined to estimate arm position represent the actual location of visual objects in the workspace and the true position of the fingertip, respectively. We also performed an alternate analysis (Figure 27B, see Methods) that assumed that the single-modality aftereffects shown in Figure 22C and D reflect a recalibration of vision and proprioception that applies to the signals combined to form $\hat{x}_{INV}$. This alternate analysis

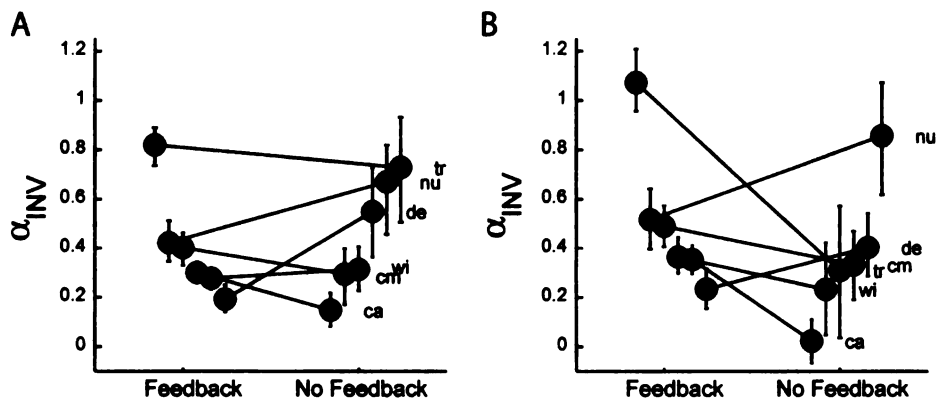

Figure 27: Fit values of α_{INV} in the with- and no-feedback conditions. Errorbars are ±1 S.E. Grey symbols indicate model fits where α_{INV} is not significantly different from zero. The initials at right identify the subjects to whom the values of α_{INV} apply, and black lines connect the values fit to data from the two feedback conditions from the same subject. (A), results of the primary analysis. (B), results of the alternate analysis that takes into account potential biases in vision and proprioception (see Methods).

yielded qualitatively similar results as the original model. We found comparable values for α_{INV} in the two feedback conditions (mean α_{INV} = .51, with feedback, α_{INV} = .36, no feedback). The average difference in α_{INV} between the with- and no-feedback conditions was not significantly different from zero (1-tailed t-test, p=.40). Note that the between-condition difference in α_{INV} is mostly due to a single subject (subject tr), who shows an unusually high value of α_{INV} in the with-feedback condition in both analyses. Therefore, although it is not clear whether the observed unimodal biases apply to $\hat{x}_{vis}$ and $\hat{x}_{prop}$ in Equation 24, including them in the model does not produce qualitatively different results.

DISCUSSION

Work from our own and other labs has described the modular nature of reach planning, in which the brain first plans the overall movement vector and then converts it into an appropriate motor command. Evidence for these two stages has come

from studies in which errors are introduced at the first stage (Rossetti et al., 1995), second stage (Ghilardi et al., 1995; Goodbody and Wolpert, 1999), or at both stages simultaneously (Sober and Sabes, 2003, 2005). Here, we provide further evidence for the modularity of reach planning by showing that sensorimotor adaptation predominantly affects only one of the two stages. Following adaptation to shifted visual feedback, subjects showed none of the MV error seen when visual feedback is shifted acutely (Sober and Sabes, 2003). The training trials, which provided subjects with feedback about their reaching errors and drove sensorimotor adaptation, therefore led to the elimination of MV error, and subjects planned reaching movements roughly parallel to the visually-specified start-target vector (Goodbody and Wolpert, 1999). On the other hand, extensive practice with shifted feedback did not eliminate INV error, which arises at the second stage of reach planning. Sensorimotor adaptation therefore appears to drive substantial changes at the movement vector planning stage, but less so during inverse model evaluation.

In our previous work, we used acute shifts of visual feedback to quantify sensory integration at both stages of motor planning. In addition to quantifying α_{INV}, we used MV error to infer α_{MV}, the weighting of vision and proprioception used during vector planning. In this study, we were unable to quantify α_{MV} because the process of sensorimotor adaptation effectively eliminated MV error. Although it is tempting to speculate that this reflects a complete reliance on vision at the movement vector planning stage, this is unlikely to be the case. Sensorimotor adaptation is composed of several phenomena. First, adaptation shifts the visual and proprioceptive maps of the workspace (Hay and Pick, 1966; Welch et al., 1974). Such changes (Figure 22C and D) presumably alter the perceived positions of the start and target, leading to realignments of the planned movement vector. Second, sensorimotor adaptation causes changes in the reach direction planned on the basis of sensory input (Welch

et al., 1974; Harris, 1963; Hardt et al., 1971). These changes are thought to account for the fact that reach aftereffects are frequently larger than the sum of visual and proprioceptive effects, as was often the case in our results (Figure 22E). Since the absence of MV error results in part from changes affecting the MV planning stage, rather than from changes in sensory integration, measuring MV error does not provide reliable information about α_{MV}. INV error, on the other hand, is not reduced during sensorimotor adaptation and should not be affected by shifts in sensory maps or by systematic alterations of the planned movement vector. We are therefore able to quantify α_{INV} in much the same way as in our previous work.

Despite the changes in movement vector planning that accompany adaptation, sensory integration during inverse model evaluation appears to be unaffected by adaptation. The with-feedback condition presented here is similar to the task presented in Sober and Sabes (2003). In both cases, subjects were exposed to laterally shifted, continuous visual feedback when planning reaching movements. This similarity of experimental design allows us to compare the fit values of α_{INV} in the with-feedback condition to the values of α_{INV} shown in Figure 7A of Chapter 1. In that study, the mean value of α_{INV} was .28. This value is somewhat lower than the mean (α_{INV}=.40) in the with-feedback condition. Therefore, although sensorimotor adaptation drives changes at the first stage of planning that lead to the elimination of MV error, INV error is the present study is as large or larger as in the absence of adaptation. Although the higher mean value of α_{INV} in the current results might reflect a real difference in sensory integration, note the presence of an outlier value of α_{INV} in the with-feedback condition (subject tr, see Figure 27A). If this value is excluded, then the mean value of α_{INV} in the present study is .32, in good agreement with the results of Sober and Sabes (2003).

A second goal of this study was to ask whether the visual error feedback that

drives sensorimotor adaptation affects sensory integration. In order to induce and maintain adaptation, subjects were given visual feedback of their reach errors in training and with-feedback trials. Because visual feedback was shifted, the INV error experienced on any trial was proportional to α_{INV}, the subject's reliance on vision. Subjects could therefore have reduced the size of their INV errors by relying more on proprioception. In contrast, in Sober and Sabes (2003), subjects did not receive any feedback about the accuracy of their reaches. Sensory integration was therefore quantified in the absence of error information, and as such the average value of α_{INV} in this earlier study can be used as a baseline measure of how vision and proprioception are integrated. Since we observed higher values of α_{INV} in this study than in Sober and Sabes (2003) (see above), we can conclude that subjects did not decrease their reliance on vision following adaptation in order to reduce INV error.

Sensory integration at the second stage of reach planning, therefore, does not appear to be a site of plasticity during sensorimotor adaptation. In contrast, we have shown that large changes in α_{INV} are driven by changing the information content of the visual signal from the arm (Sober and Sabes, 2005). Sensory integration during inverse model evaluation is therefore flexible, but driven by different signals than those that cause the observed changes in movement vector planning.

The third aim of this study was to investigate whether the availability of online visual feedback has an impact on the weighting of vision and proprioception. Fitting our model to the data showed that the mean value of α_{INV} is similar in the no-feedback and with-feedback conditions. This indicates that a visual landmark - the spot of light to which subjects initially reached in no-feedback trials - can influence a subject's arm position estimate to the same extent as continuous visual feedback. The most significant factor, therefore, appears to be that the subject treats the

visual object as a proxy for hand position, and not whether the visual cue mimics the movement of the hand. Similarly, Lackner (1974) found that subjects undergo sensorimotor adaptation when a discrepancy exists between the felt position of the hand and the location of a visual landmark (a vertical metal pin). This study provides further evidence that the nervous system can use the location of a stationary object bearing no resemblance to the hand to estimate arm position.

An alternate explanation for the results of our model fits is that sensory integration is strongly influenced by gaze direction. Eye position has been shown to affect the accuracy of reaching to remembered visual targets (Enright, 1995), and the accuracy of reaching is affected by the retinotopic location of the target (Bock, 1986; Henriques et al., 1998). Although eye position was neither measured nor constrained in our studies, it seems likely that subjects' oculomotor behavior was similar in both feedback conditions, since the visual display was nearly identical in both tasks. A possible interpretation of our results, therefore, is that the similar values of α_{INV} might reflect similar oculomotor behavior in the two conditions, due to the potential influence of gaze direction on the hand position estimate $\hat{x}_{\text{INV}}$. This explanation is unlikely to account for our results for two reasons. First, in Sober and Sabes (2003) we show that while the fit values of α_{INV} reflect a reliance on proprioception at the second stage of reach planning, subjects rely mostly on vision at the first stage of planning, when the movement vector is being computed. Since eye position is presumably the same when both computations are taking place, sensory integration cannot be determined by eye position alone. Second, we have shown that when eye position is controlled, sensory integration at a single planning stage can differ dramatically between two different tasks (Sober and Sabes (2005), Supplementary Information). Therefore, although the gaze dynamics during the present experiment are unknown, it seems unlikely that the similar values of α_{INV} result from similar

oculomotor behaviors.

Sensory integration and sensorimotor learning are both processes in which the brain must combine discrepant peripheral signals in order to produce accurate behavior. In the case of sensory integration, properties of the task and the stimulus have been shown to influence the weighting of different modalities (Ernst and Banks, 2002; Sober and Sabes, 2005). In both of these studies, sensory inputs appear to be weighted in ways that reduce the variability of the integrated estimate, thereby improving task performance. In the case of visuomotor learning, visual error feedback is used to drive changes in motor planning and in the calibration of visual and proprioceptive workspace maps. These changes serve both to reduce the apparent motor error and to decrease the discrepancy between vision and proprioception.

Our data suggest a close relationship between adaptation to shifted visual feedback and movement vector planning. Because the visual shifts employed in this study altered only the apparent kinematics of movement (rather than the dynamics of the reaching limb), it is perhaps unsurprising that alterations were seen only during movement vector planning, when the desired reach kinematics are computed. Many studies have shown that exposing subjects to force fields or altering the experienced dynamics of the arm leads to lasting changes in motor planning (Shadmehr and Mussa-Ivaldi, 1994; Dizio and Lackner, 1995). It seems likely that this type of learning reflects changes at the second stage of planning, when the dynamics of the reaching limb must be taken into account. If this is the case, then it seems clear that different types of error signals drive changes at discrete computational stages in the motor planning pathway. Exploring how these errors signals affect different planning modules will provide insight into the brain's ability to adapt to novel tasks and environments.

Chapter 4: Possible Neural Mechanisms and Future Directions

Abstract

Psychophysical studies have shown that the planning of reaching movements consists of two distinct stages. First, the overall direction of movement (movement vector) is planned by combining visual and proprioceptive feedback from the arm with signals specifying the location of the target. Second, the desired movement vector is converted into an intrinsic motor command using visual and proprioceptive information about the arm's posture. Sensory integration at both stages is flexible, and visual and proprioceptive signals appear to be weighted in order to minimize a cost associated with transforming signals between coordinate frames. Here, I review some of the recent literature on the neurophysiology of motor control and consider where in the brain these two planning stages might be performed. Then, I propose a set of experimental approaches that would allow an experimenter to confirm the computational role of single reach-related neurons. Finally, I suggest several models of the cost of coordinate transformations in terms of neural activity, and suggest possible mechanisms by which this cost might influence sensory integration in a range of behavioral tasks.

In the preceding chapters, we have described a mathematical model of reach planning in which sensory signals from the arm and the reaching target are combined to produce motor commands. This model proposes that reach planning proceeds in two separate stages – a first stage in which reach trajectories are planned in the coordinates of the workspace, and a second in which the planned trajectory is converted into a motor command. In the model, visual and proprioceptive inputs are weighted at each stage to produce an estimate of the arm's position. In Chapter 1, we provided evidence that the weighting of these two inputs can vary greatly between the two planning stages, and suggested that this variation is due to differences in the computational problems that the brain must solve at each stage. In particular, the observed pattern of sensory integration is consistent with a strategy in which the brain avoids a cost associated with relying on signals that have been transformed between coordinate frames. In Chapter 2, we tested this hypothesis by varying the parameters of the reaching task. These manipulations led to independent changes in sensory integration at the two stages, and provided further evidence for a two-stage model of reach planning in which sensory integration is strongly influenced by the cost of performing coordinate transformations.

In this chapter, I address how such a two-stage model of motor planning might be instantiated in the primate brain. I begin by considering where in the brain the two planning stages might take place. Although no neurophysiological studies have been explicitly designed with our mathematical model in mind, many researchers have investigated the planning of reach trajectories and the generation of motor commands. This literature is therefore extremely useful in trying to guess the mapping between the proposed planning stages to the many areas of the brain involved in target-directed reaching. The response properties of neurons in certain areas of the brain often correlate well with the general properties of one of the planning stages,

allowing a preliminary mapping of planning computations onto neural structures.

Following this review of the literature, I suggest an experimental approach to identifying the computational roles of single neurons. The proposed approach uses the psychophysical methods developed in earlier chapters to generate quantitative predictions about how neurons involved in specific computations should fire under different task conditions. Such a method could be applied to any area of the brain where firing is affected by hand position during reach planning and could be used to differentiate multiple functional networks within a single area.

Finally, I consider how the cost of performing coordinate transformations might be expressed in terms of the activity of populations of neurons. Based on neural models of how sensory signals are combined to estimate arm position, I consider how this cost might influence sensory integration. While speculative, these model circuits provide a biological account of the psychophysical results described in earlier chapters.

Two Stages of Reach Planning: Expected Neural Correlates

When planning reaching movements, sensory signals from the arm and target are integrated and transformed into the appropriate motor commands. Psychophysical studies have suggested that reach planning consists of two distinct stages (Sober and Sabes, 2003; Rossetti et al., 1995; Goodbody and Wolpert, 1999; Sainburg et al., 2003). These stages, and the computations comprising them, are shown in Figure 28a. During the movement vector planning (MV) stage, visual and proprioceptive signals about the position of the arm are combined to produce $\hat{x}_{\mathrm{MV}}$, an estimate of the hand's initial position (Equation i in Figure 28a). The combination of sensory signals is parameterized by α_{MV}, which describes the percent to which subject rely

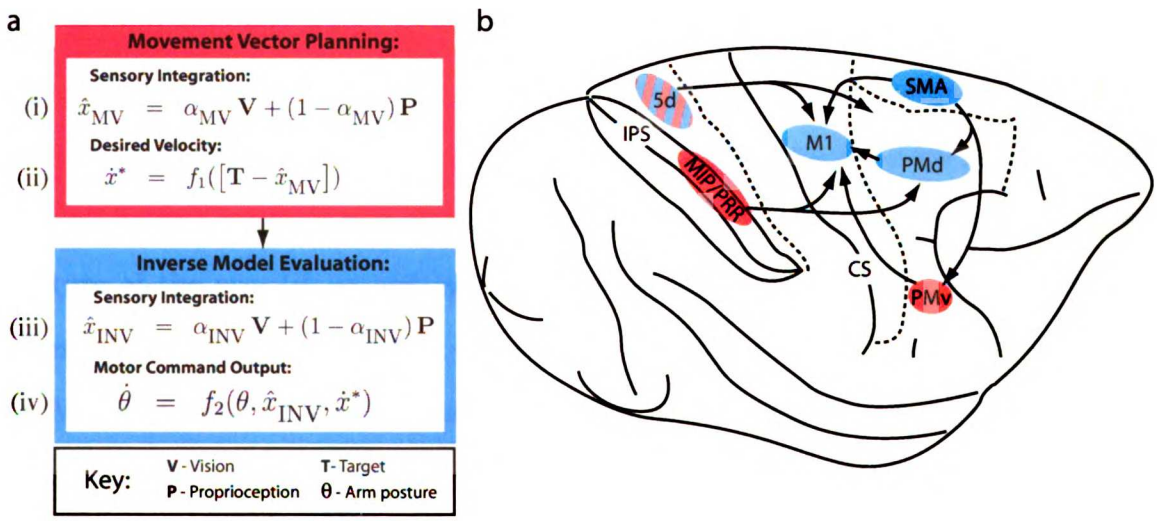

Figure 28: Planning stages and potential cortical sites. (**a**) Reach planning consists of a movement vector planning stage (red) and an inverse model evaluation stage (blue). The computations that make up the two stages are summarized by Equations i-iv. (**b**) Potential cortical loci for the two planning stages. Lateral view of a macaque brain (anterior is to the right) showing several areas involved in goal-directed reaching movements. Areas are colored based on their potential participation in the two planning stages according to the color conventions in (**a**), as explained in the text. Part (**b**) adapted from Wise et al. (1997), additional connections drawn based on Dum and Strick (2002) and Kurata (1991). IPS, intraparietal sulcus; CS, central sulcus; see text for other abbreviations.

on visual feedback. The position estimate $\hat{x}_{\mathrm{MV}}$ is used to plan the desired velocity (Equation ii), according to some function f_1. Note that the only input to this function is $[\mathrm{T} - \hat{x}_{\mathrm{MV}}]$, the vector from the estimated initial position to the target.

In the inverse model evaluation (INV) stage, the desired velocity computed by the MV stage is converted into a motor command. First, visual and proprioceptive signals are again combined to produce an estimate of the arm's position (Equation iii). The weighting of vision and proprioception at this stage is described by α_{INV}. This integrated arm position estimate is then used to generate a motor command, shown here as a velocity in joint-angle space ($\dot{\theta}$), according to some function f_2. Note that unlike the first stage of planning, the output of this function depends on the true posture of the arm (θ) as well as its estimated initial position ($\hat{x}_{\mathrm{INV}}$). The exact forms of functions f_1 and f_2 are described in Chapter 2.

When speculating about the neural basis of these computations, it is important to consider which areas of the brain have the right combination of response properties to be involved in the MV and INV stages. First, the candidate areas must receive both visual and proprioceptive feedback from the arm. Second, the structures must be active during reach planning and/or execution, and lesions of these areas should lead to deficits when performing sensorimotor tasks. Third, the candidate areas must be connected to each other such that sensory signals are passed along to motor output structures. These criteria describe a large number of networks in the brain involving cortical, subcortical, and cerebellar structures. For the purposes of this review, however, I will focus primarily on one particularly well-studied set of interconnected cortical areas spanning the parietal and frontal lobes (Wise et al., 1997).

This parieto-frontal network, shown in Figure 28b, includes several areas within and dorsal to the intraparietal sulcus. These areas, which receive visual as well as proprioceptive inputs, are strongly interconnected with several motor and premotor

areas in frontal cortex (Wise et al., 1997; Dum and Strick, 2002; Kurata, 1991). Reach-related activity has been observed in each of these areas (Kalaska et al., 1990; Scott et al., 1997; Kakei et al., 1999; Li et al., 2001; Buneo et al., 2002), and lesion studies have demonstrated their importance to the execution of reaching movements (Brinkman, 1984; Passingham, 1988; Hoffman and Strick, 1995; Rushworth et al., 1998).

The components of the parieto-frontal circuit are therefore candidates for implementing the MV and INV stages of reach planning. The response properties expected from an area involved in each of these stages can be characterized according to the computations performed during planning. Calculating the movement vector, for example, consists of planning the desired velocity of the hand in extrinsic (workspace) coordinates. The output of this planning stage is determined by the disparity between the estimated position of the hand and the target (Equation ii in Figure 28a), and is not influenced by the change in joint angles or muscular forces needed to achieve the desired kinematics. Neurons that compute the movement vector should therefore code exclusively the extrinsic aspects of the planned movement. That is, they should modulate their activity only when the desired hand trajectory is altered, even if the change in joint angles needed to achieve it is identical, and should show the same response when a given trajectory is achieved with different motor commands. Consistently finding this type of "extrinsic coding" in an area would suggest that that area is involved in the MV planning stage.

Evaluating the inverse model, on the other hand, consists of transforming the extrinsic movement vector into the appropriate intrinsic (arm-based) coordinates (Equation iv in Figure 28a). The motor command necessary to achieve a given hand velocity depends on the posture of the arm and on the muscular forces needed to produce the appropriate change in joint angles. Neurons involved in the INV stage

should therefore modulate their activity to reflect the change in joint angles or muscle contractions needed to achieve the specified movement vector, even when different motor commands are used to achieve the same extrinsic trajectory. Conversely, these cells should show the same response when identical motor commands are used to produce different hand trajectories. Finding this type of "intrinsic coding" would implicate an area in the INV planning stage.

Extrinsic and Intrinsic Coding: a Review of the Literature

Recent work on the physiology of motor control has taken a variety of approaches to the question of extrinsic versus intrinsic coding. One of the most fruitful has been to record neural activity while varying the motor command used to achieve a given movement direction. In one series of experiments, monkeys were trained to make wrist movements with directions specified in extrinsic coordinates (Kakei et al., 1999, 2001). These movements were repeated with the forearm in three different postures, thereby dissociating the extrinsic movement vector - the direction of hand movement - from the intrinsic motor command needed to produce it. Neural recordings in the frontal lobe revealed that cells in the ventral premotor cortex (PMv) are tuned to the extrinsic aspects of the movement, whereas cells in primary motor cortex (M1) showed a mixture of extrinsic and intrinsic (muscle-like) responses (Kakei et al., 2001). These data suggest that PMv might be involved in movement vector planning, whereas M1 might take part in the transformation of the desired direction into an appropriate motor command. Areas shown in Figure 28b are colored to reflect these interpretations.

Data suggesting that the extrinsic and intrinsic aspects of reaching are encoded by PMv and M1 (respectively) also come from experiments in which primate subjects

adapt to distorted visual feedback. By driving sensorimotor adaptation during a tracing task, Schwartz et al. (2004) showed that cells in PMv encoded the trajectory of a cursor representing hand position, suggesting that PMv encodes the apparent movement vector even when it does not reflect the actual motion of the hand. Activity in M1, on the other hand, encoded the true hand trajectory, suggesting that this area computes the intrinsic motor command. Kurata and Hoshi (2002) obtained comparable results using displacing prisms to separate neural activity related to the position of extrinsic (visual) targets from activity encoding the motor commands used to reach to them. They found cells coding in extrinsic and intrinsic coordinates during reach planning in PMv, whereas responses found during planning in M1 were almost exclusively intrinsic.

Evidence for intrinsic coding in M1 and other areas has also been obtained by training monkeys to make reaches with similar hand trajectories but different postures (Scott et al., 1997; Scott and Kalaska, 1997). Variation in the starting posture meant that different changes in joint angles were required to produce the same hand path, dissociating the planned movement vector from the motor command needed to produce it. In M1, dorsal premotor cortex (PMd), and parietal area 5, the majority of cells modulated their activity when different postures were used to produce the same movement. These findings provide more evidence that M1 might be involved in transforming movement vectors into motor commands, and suggests that PMd and area 5 may be as well.

This view of M1 and PMd is supported by the results of Caminiti et al. (1990, 1991). In these studies, monkeys made reaching movements in three different regions of the workspace. Targets were arranged so that the same movement vector was executed from different initial arm positions, requiring different motor commands to produce parallel reach trajectories. These authors found that single cells in both M1

and PMd regularly changed their response properties when the initial position of the reach was varied, providing further evidence that these areas encode the intrinsic aspects of reaching movements.

Other groups have used externally imposed forces to ask whether neural activity reflects the planned trajectory or the motor command necessary to produce it. These studies, which typically involve training subjects to produce similar kinematics while reaching with different loads or force fields, allow the same reach trajectory to be associated with a variety of required muscular forces. Such manipulations produce consistent changes in firing in M1 during reaching, once again showing that this area codes the intrinsic aspects of movement (Li et al., 2001). Even more provocative is the finding that delay-period (premotor) activity in PMd and the supplementary motor area (SMA) smoothly changes from extrinsic to intrinsic prior to reach onset, strongly implicating these areas in the kinematics-to-dynamics transformation characteristic of the INV stage (Padoa-Schioppa et al., 2002). Firing in parietal area 5, however, does not appear to be influenced by external loads (Kalaska et al., 1990), suggesting that this area is not involved in the transformation of desired kinematics to dynamic motor commands. However, the findings of Scott et al. (1997) described above suggest that area 5 codes the angles of the joints in addition to the kinematics of the endpoint, making the relationship between this parietal region and the two stages of reach planning difficult to evaluate.

Neurons elsewhere in the parietal cortex appear to encode the position of the target in extrinsic, eye-centered coordinates. By varying the reach start point, target location, and gaze direction, Batista et al. (1999) showed that cells in the parietal reach region (PRR) are tuned to the position of the target relative to the animal's center of gaze. Cells that code in eye-centered coordinates have similar activity during reaches in different directions and with different start points, provided that

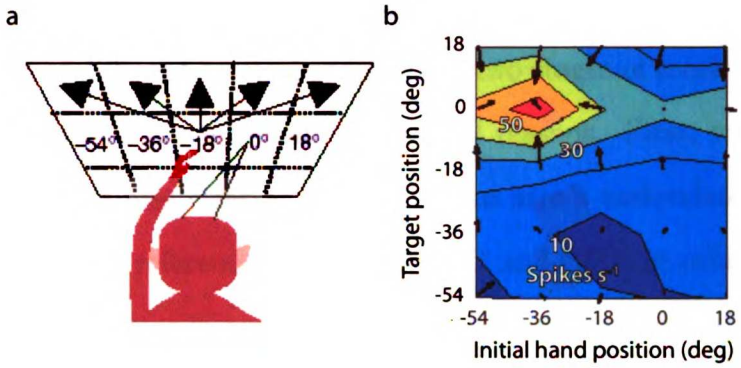

Figure 29: Tuning for initial hand position in the parietal cortex. (**a**) Task schematic. A monkey made reaches from five different start points to five target positions while maintaining gaze fixation at a single position. (**b**) Response of a single neuron recorded in PRR. Numerical labels and color contours show the mean firing rate during the 400 msec epoch centered on reach initiation. Reproduced from Buneo et al. (2002).

target position relative to the center of gaze is constant. It is therefore unlikely that these cells compute the motor command necessary for task execution. On the other hand, the activity of most PRR neurons is not in the hand-centered coordinates of the movement vector, either. However, the hand-centered movement vector could in theory be readily extracted from the activity of a population of PRR neurons. Although neurons in this region are tuned for the location of the target in eye coordinates, the tuning of some cells is scaled (or "gain modulated") by the initial hand position (Figure 29). This scaling provides the information necessary to extract the position of the target in hand-centered coordinates from a population of gain-modulated neurons (Buneo et al., 2002; Salinas and Abbott, 1995). Based on these data and on the anatomical connectivity of PRR, it is tempting to speculate that the gain-modulated eye-centered activity observed in PRR represents a source of input to the MV planning stage. Significantly, neurons in area 5 appear to code in coordinates intermediate between eye and hand (Buneo et al., 2002), suggesting that they might process the output of PRR for further use in motor planning.

Figure 28b shows our predictions, based on the literature reviewed above, about the participation of different cortical areas in the two stages of reach planning. However, these speculations are uncertain for several reasons. First, while the physiological data are sometimes clear as to whether an area's task-related activity is in an extrinsic or intrinsic reference frame, as for M1 and PMv, in other cases (such as area 5) the correspondence is not as apparent. Second, assuming that all cells in a given area code the same way is almost certainly an oversimplification. A single area might be home to many functionally distinct networks, and so assigning an area to a given planning stage based on its average response may not be appropriate. Third, although extrinsic and intrinsic coding are the expected hallmarks of the two stages of reach planning, these criteria are not definitive. Movement vector planning and inverse model evaluation are transformations of sensory inputs, and so it is ambiguous whether activity in a given reference frame reflects the output of a transformation or the transformation itself. Intrinsic coding, for example, might reflect either the transformation of the extrinsic movement vector or the motor signals downstream of it. EMG signals are by definition intrinsic, for example, but it would be a mistake to conclude that the inverse model is being evaluated by the muscles.

Neural Correlates of Sensory Integration: a Suggested Experimental Approach

For these reasons, I propose an experimental approach that does not rely on making distinctions between intrinsic and extrinsic patterns of activity. Rather, I suggest using the psychophysical methods described in Chapters 1 and 2 in combination with neural recordings to correlate neural firing patterns with the fit values of α_{MV} and α_{INV}. Establishing that the tuning of a single neuron reflects the same mixture of

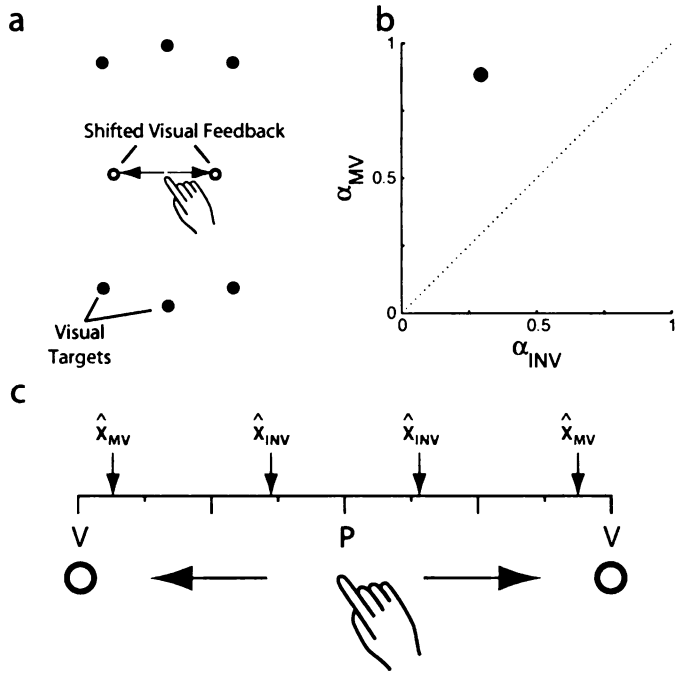

Figure 30: Spatial distribution of hand position estimates. (**a**) Experimental technique used in Experiment 1 of Chapter 2.. Subjects reached to visual targets with leftward (blue arrow) and rightward (red arrow) shifts of visual feedback. (**b**) Mean fit values of α_{MV} and α_{INV}. (**c**) Inferred location of hand position estimates. Vertical arrows show the position of $\hat{x}_{\mathrm{MV}}$ computed from the locations of the visual (V) and proprioceptive (P) signals from the hand and from the fit values of α_{MV} and α_{INV}.

visual and proprioceptive information observed during movement vector planning or inverse model evaluation will provide strong evidence that the cell is involved in that stage of reach planning.

Figure 30a and b show the psychophysical technique used in Experiment 1 of Chapter 2. Subjects planned reaches to visual targets while visual feedback (consisting of a spot of light) was shifted leftward or rightward. By analyzing the reach errors that resulted from these shifts, we fit the values of α_{MV} and α_{INV}. Mean model fits are shown in Figure 30b.

Since shifting the visual feedback dissociates the visual and proprioceptive position signals, the relative weighting of the two inputs determines the position the

integrated position estimate. When estimating the initial hand position during movement vector planning, for example, subjects relied mostly on vision (mean value of $\alpha_{\mathrm{MV}} = .88$, Figure 30b). The estimated initial hand position ($\hat{x}_{\mathrm{MV}}$ in Equation i) is therefore much closer to the visual signal than to the proprioceptive signal, as shown in Figure 30c. When estimating the initial hand position during inverse model evaluation, on the other hand, subjects rely mostly on proprioception (mean $\alpha_{\mathrm{INV}} = .29$). The estimated initial hand position ($\hat{x}_{\mathrm{INV}}$) is therefore closer to the proprioceptive signal.

Dissociating visual and proprioceptive feedback therefore allows the experimenter to shift the arm position estimate used at a given planning stage by a known amount. This is an potentially useful experimental tool, since neurons in several areas of the brain are tuned to initial hand position during reach planning and execution (Caminiti et al., 1990, 1991; Buneo et al., 2002). The neuron shown in Figure 29b, for example, is tuned to both the initial position of the hand and the position of the target.

Although cells tuned to initial hand position have been observed in several areas of the brain, the tuning of such neurons has never been studied when visual and proprioceptive feedback from the hand is dissociated during reach planning. It has therefore not been possible to say whether a neuron tuned to initial hand position is encoding $\hat{x}_{\mathrm{MV}}$ (during movement vector planning) or $\hat{x}_{\mathrm{INV}}$ (during inverse model evaluation). The experiment described below uses shifts of visual feedback to disambiguate the computational role of the recorded neuron.

Figure 31a schematizes a behavioral task similar to that shown in Figure 30. The primate subject would make center-out reaches from a linear array of start points (start position range from – to +, arbitrary units). Veridical (unshifted) visual feedback would be provided during this portion of the experiment. Figure 31b

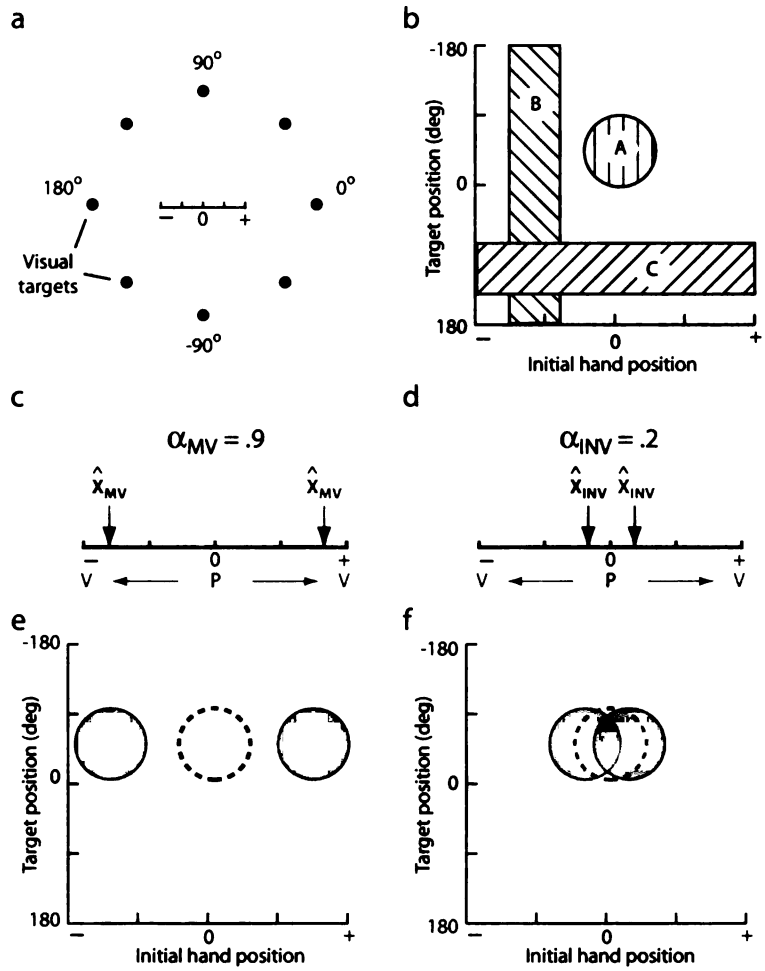

Figure 31: Proposed recording experiment. (**a**) Workspace configuration. Subjects make center-out reaches from a range of start positions. (**b**) Model neural response properties. (**c**) Location of the position estimate $\hat{x}_{\mathrm{MV}}$ following leftward (blue) and rightward (red) shifts of the visual feedback, assuming a fit value of $\alpha_{\mathrm{MV}} = .9$. (**d**) Location of $\hat{x}_{\mathrm{INV}}$, for a fit value of $\alpha_{\mathrm{INV}} = .2$. (**e** and **f**) Expected shift in tuning for model neuron A following leftward (blue) and rightward (red) visual feedback shifts.

shows the tuning profiles of three hypothetical neurons. Neuronal responses would be quantified as a function of starting and target position, as shown. Neurons might be tuned to both the starting and target position (model neuron A, note the similarity to the cell shown in Figure 29b), the initial hand position only (model neuron B) or the target position only (model neuron C).

The response profiles in Figure 31b describe the model neurons' tuning to the initial position of the hand. However, from these data alone it is not clear which of the two position estimates ($\hat{x}_{\mathrm{MV}}$ or $\hat{x}_{\mathrm{INV}}$) is being encoded. This ambiguity will be resolved by having the subject make reaches with visual feedback shifts similar to the ones shown in Figure 30a. This manipulation will yield two important pieces of data. First, we will be able to fit our model of reach planning to the behavioral data, yielding values for α_{MV} and α_{INV}. As shown in Figure 31c and d, these values tell us how far the position estimates $\hat{x}_{\mathrm{MV}}$ and $\hat{x}_{\mathrm{INV}}$ move when the visual feedback is shifted. Second, we will be able to measure shifts in the neuron's tuning profile that result from shifting the visual feedback, and compare the size of these shifts with those expected from a cell encoding a particular position estimate. If a tuning profile encodes $\hat{x}_{\mathrm{MV}}$, for example, then its peak should shift by the same distance as $\hat{x}_{\mathrm{MV}}$, as shown in Figure 31e. If a neuron encodes the location of $\hat{x}_{\mathrm{INV}}$, on the other hand, its peak is predicted to shift by the same distance as $\hat{x}_{\mathrm{INV}}$, as illustrated in Figure 31f. A close correlation between these two types of data – such as that shown between Figure 31c and e or Figure 31d and f – would provide strong evidence of the involvement of the recorded neuron in one of the two stages of reach planning.

This experimental paradigm would allow the identification of other neuronal populations as well. Cells that showed no shift in tuning profile (even when the fit values of α_{MV} and α_{INV} indicate that visual signals are in fact being used) would be categorized as proprioceptive, since they encode the true initial position of the hand.

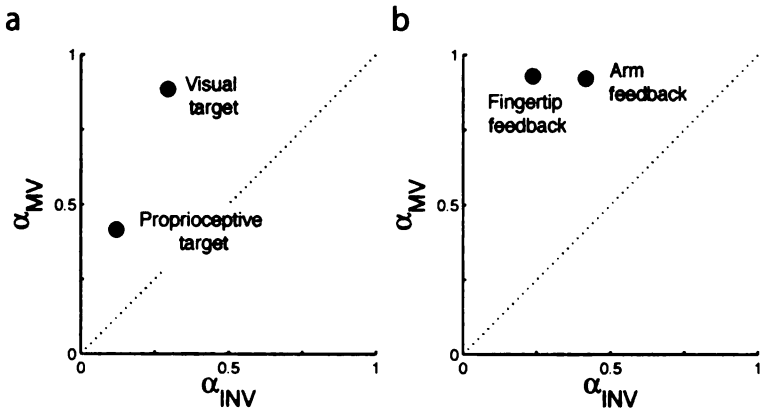

Figure 32: Psychophysical results from Chapter 2. (a) Mean fit values of α_{MV} and α_{INV} from Experiment 1, in which the modality of the reaching target was varied. (b) Mean fit values from Experiment 2, in which the content of the visual feedback was varied.

Conversely, cells with shifts equal in magnitude to the visual shift (even when α_{MV} and α_{INV} indicate a partial reliance on proprioception) would be categorized as purely visual.

This technique therefore constitutes a quantitative method to collect evidence about the computational roles of single neurons. The experiment shown in Figure 31 uses only a single behavioral condition (reaches to visual targets using visual feedback of the fingertip), resulting in a single set of values for α_{MV} and α_{INV}. However, our previous work has shown that altering the coordinate frame of the target or the content of the visual feedback reliably drives stage-specific, quantifiable changes in sensory integration (Figure 32). If preliminary testing suggested that a cell encoded $\hat{x}_{MV}$ or $\hat{x}_{INV}$, testing with a different behavioral task (e.g. switching from visual to proprioceptive targets in order to vary the value of α_{MV}) would allow the correlation between the fit value of α_{MV} and the shift in neural tuning to be further established. Note however that this technique is only viable for cells tuned to initial hand position (such as model neurons A and B in Figure 31b). Cells not tuned to this parameter (such as model neuron C) could not be analyzed with this method.

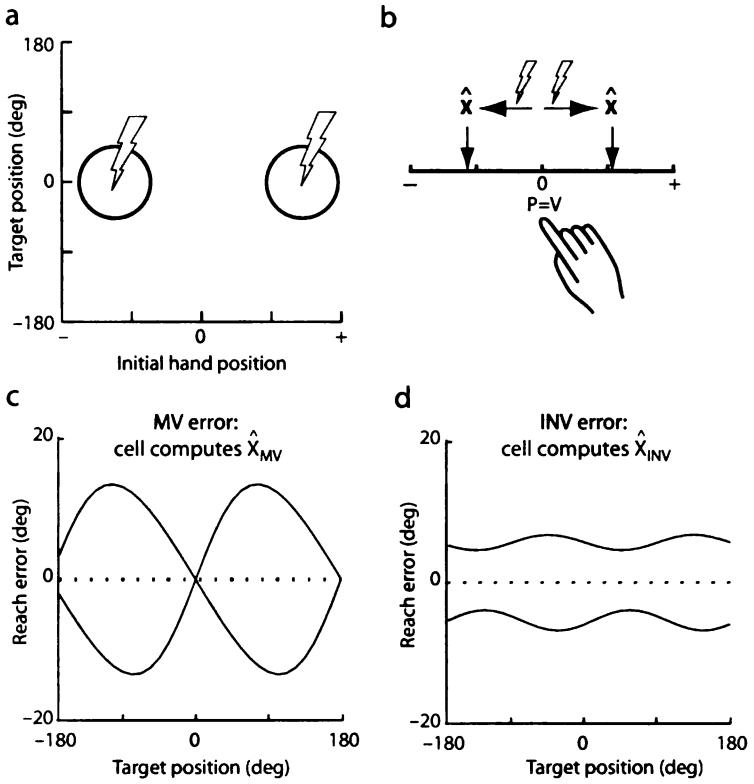

Figure 33: Proposed stimulation experiment. (**a**) Response profiles of two hypothetical neurons, which have peaks to the left (blue) and right (red) of the central initial hand position. (**b**) Expected effect of stimulating the two hypothetical neurons. (**c**) Movement vector error expected if stimulated neurons encode $\hat{x}_{\mathrm{MV}}$ and if stimulation leads to rightward (red) or leftward (blue) shifts in the estimated hand position. (**d**) Inverse model error expected if stimulated neurons encode $\hat{x}_{\mathrm{INV}}$.

Additional data about the neural basis of a particular computation could be gathered by electrically stimulating the brain during reach planning. Figure 33 shows a second proposed experiment, which would be performed after the tuning of a neuron has been established using the procedure shown in Figure 31. After establishing a neuron's response profile (two hypothetical profiles are shown in Figure 33a), stimulating current would be passed through the recording electrode while the subject planned reaching movements with veridical visual feedback. Provided that the stimulation drives activity in both the recorded cell and in other nearby cells with the same tuning, the hand position estimate will be shifted in the direction of the cell's preferred position (Figure 33b). In our previous work, we have shown that errors in $\hat{x}_{\mathrm{MV}}$ and $\hat{x}_{\mathrm{INV}}$ produce distinct patterns of motor error. Stimulating a network of neurons that computes $\hat{x}_{\mathrm{MV}}$ would be predicted to induce movement vector (MV) error when reaches are executed (Figure 33c). Perturbing the activity of cells that compute $\hat{x}_{\mathrm{INV}}$, on the other hand, would introduce inverse model (INV) error, as shown in Figure 33d. See Figure 10 in Chapter 2 for an explanation of the particular shapes of MV and INV error.

For stimulation to yield a consistent pattern of motor error, the stimulating current must drive activity in a population of neurons with similar tuning to the recorded cell. This would be expected to occur if cells with similar preferred initial hand positions were adjacent to each other in the brain. Such topography, however, has not been reported in the cortical areas in which hand-position tuning has been described. However, stimulation could also activate a coherent network of neurons if cells with similar tuning activated each other through lateral connections. Connectivity of this type is thought to exist in visual cortex (Malach et al., 1993). If no such topography or connectivity exists, however, stimulation would not be expected to produce MV or INV error, even if the recorded neuron was indeed involved in

computing the position of $\hat{x}_{MV}$ or $\hat{x}_{INV}$. A failure of stimulation to elicit MV or INV error might therefore be due to the connectivity of the cortical circuit rather than the computational role of the site being stimulated.

The techniques described above could be applied to any neural population that encodes the initial position of the hand as a part of reach planning, and could therefore be used to investigate subcortical and cerebellar structures in addition to the cortical areas shown in Figure 28b. These methods are designed to implicate single cells in particular computations, and could be used to identify separate computational networks that might be mixed within a single area of the brain. Furthermore, the specificity of this technique offers the potential to map the computational steps outlined in Figure 28 onto low-level anatomical features such as cell morphology and cortical layer.

Possible Neural Mechanisms for Flexible Sensory Integration

As discussed in Chapter 2, the dependence of α_{MV} and α_{INV} on the modality of the target and the nature of the visual feedback (Figure 32) is consistent with brain's avoiding a cost associated with transformed sensory signals. A full neural model of sensory integration must explain this cost in terms of neural activity, and several possibilities are discussed below. I propose a model for how this cost affects activity in neural populations that integrate visual and proprioceptive signals, and speculate as to how these factors give rise to the observed task-related effects on α_{MV} and α_{INV}.

The cost of performing a coordinate transformation using neural units could take several forms. Visual and proprioceptive afferents encode the position of the arm in different coordinate frames. These coordinate systems presumably reflect

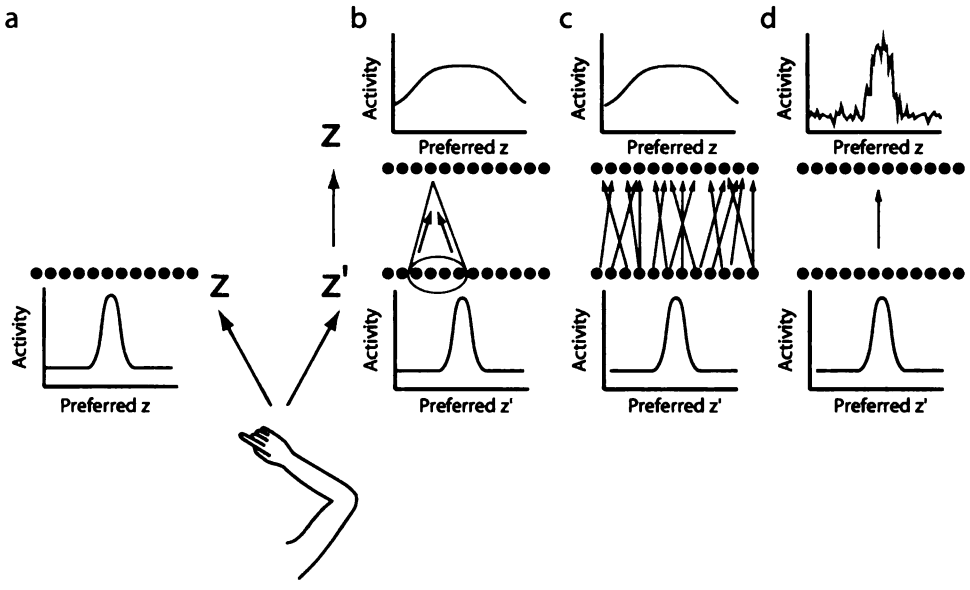

Figure 34: Hypothesized costs for coordinate transformations. (**a**) General case in which arm position can be represented in two coordinate systems, $\mathbf{z}$ and $\mathbf{z}'$. Arm position is encoded by spiking activity in populations of neurons (colored dots) with a range of preferred values of $\mathbf{z}$ (green) or $\mathbf{z}'$ (red). Transforming a neural representation of arm position in $\mathbf{z}'$ into a representation in $\mathbf{z}$ (blue) could change the shape of network activity because of convergence (**b**), random errors in connectivity (**c**), or neural noise (**d**).

the properties of the sensory periphery – visual signals begin in the two-dimensional coordinates of the retina and are used to construct a three-dimensional representation of space, whereas proprioceptive signals have a dimensionality related to the number of joint angles, muscle lengths, or muscle tensions that describe the reaching limb. Figure 34 shows the simplified case in which arm position can be represented in one dimension and is encoded by populations of neurons with preferred tuning in two different coordinate frames, $\mathbf{z}$ and $\mathbf{z}'$. In this schema, coordinate transformations are accomplished via the pattern of synaptic connectivity between populations of neurons (Figure 34 b-d).

There are several ways that transforming the population activity in this manner could alter population activity. First, inputs from neurons with a range of preferred values of might converge onto each output cell, resulting in a broadening (detuning) of the population response (Figure 34b). This would be expected if the mapping between coordinate systems were many-to-one. For example, if joint angles were represented in $\mathbf{z}'$ and the visual location of the fingertip were represented in $\mathbf{z}$, then the fact that many different postures can produce a single fingertip position would be expressed as a convergence of $\mathbf{z}'$ inputs onto $\mathbf{z}$ outputs. Second, even if there were a one-to-one mapping between coordinate frames, random errors in the connectivity between the input and output populations could also cause a broadening in the population response (Figure 34c). A third possibility is that the coordinate transformation would simply add noise to the population code, owing to the stochastic nature of neural transmission and/or imprecision in the strength of the connections between the two populations (Figure 34d). Any of these factors (or some combination of them) would alter the shape of the population activity encoding the position of the arm.

Figure 35 shows a model of how these transformation-induced changes in popu-

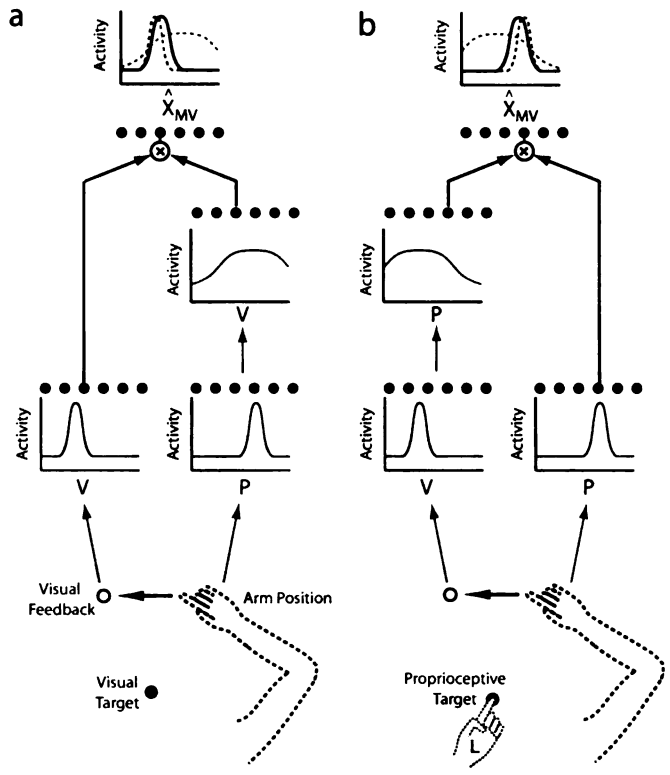

Figure 35: Target modality and sensory integration. (**a**) Model of neural activity in the visual-target condition of Experiment 1 from Chapter 2. Shifted feedback dissociates the visual and proprioceptive encoding of arm position (bottom). Proprioceptive activity is transformed into visual coordinates, resulting in a broadening of the population activity (middle). The two sensory streams converge on an output population representing $\hat{x}_{MV}$ (top). (**b**) Model of neural activity in the proprioceptive-target condition. In this case, visual signals from the arm are transformed into proprioceptive coordinates.

lation activity could affect sensory integration during movement vector planning. In Chapter 2, we showed that subjects rely more on visual feedback from the arm at this stage when reaching to visual targets, but more on proprioception when reaching to proprioceptive targets. We hypothesized that this difference in sensory integration arises from the cost of transforming sensory signals into the coordinate frame of the target. In the example shown in Figure 35a, the visual feedback is shifted leftward as the subject plans a reach to a visual target. The peak activity of the neural population representing arm position in visual coordinates (model neural activity in green) is therefore located to the left of the peak representing arm position in proprioceptive coordinates (red). The transformation of the proprioceptive encoding of arm position into visual coordinates (blue) results in a broadening of the activity profile, as described above. This transformed population activity is then combined with activity encoding arm position in visual coordinates as the two populations converge onto an output population. This output network (top, model neural activity in black) encodes the hand position estimate $\hat{x}_{\mathrm{MV}}$, and multiplies the inputs provided by the two sensory pathways (Salinas and Abbott, 1996; Pouget et al., 2003). Because the proprioceptive population activity has been broadened as a result of being transformed, the peak of the output population activity will be closer to the position originally encoded by visual signals (dashed green line). If the position estimate $\hat{x}_{\mathrm{MV}}$ is located at the peak or center-of-mass of the output population activity, then the fit value of α_{MV} will be close to one, reflecting the proximity of $\hat{x}_{\mathrm{MV}}$ to the visual signal and in agreement with our psychophysical data on visual-target reaches (Figure 32a). Note that in the limit case where the transformed activity has been broadened completely (flattened), the visual signal will dominate completely.

In the case of reaches to proprioceptive targets, on the other hand, we have hypothesized that visual signals are transformed into proprioceptive coordinates. As

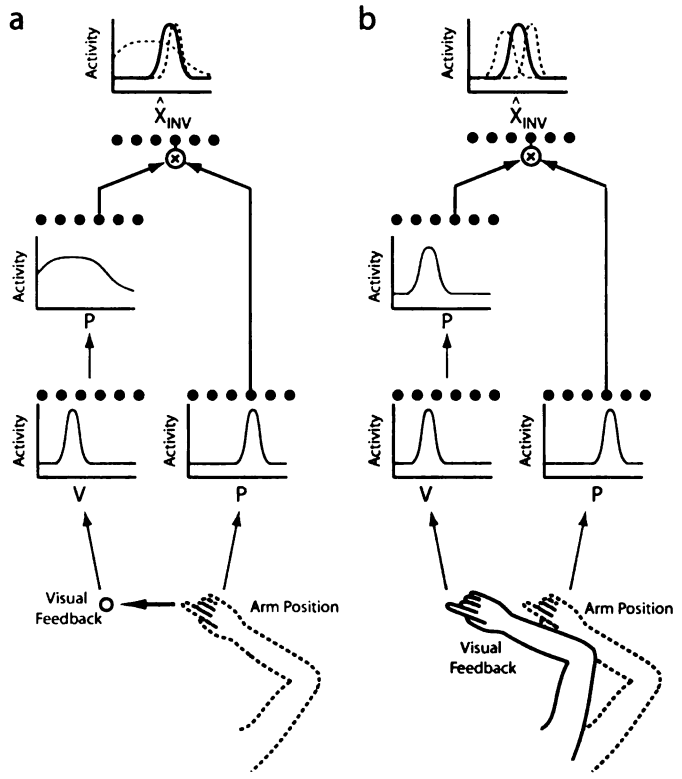

Figure 36: Visual feedback type and sensory integration. (**a**) Model of neural activity in the fingertip-feedback condition of Experiment 2 from Chapter 2. (**b**) Model of neural activity in the arm-feedback condition. Conventions as in Figure 35.

shown in Figure 35b, this transformation might lead to a broadening of the population activity as visually-driven activity is converted into proprioceptive coordinates. When the activity of the transformed signal is combined with the proprioceptive population activity, the peak of the output population encoding of $\hat{x}_{\mathrm{MV}}$ will be closer to the proprioceptive location. The fit value of α_{MV} in this case will therefore be closer to zero, as in the psychophysical results shown in Figure 32a.

Figure 36 shows how the cost of transforming sensory signals could account for the effects of visual feedback type on sensory integration during inverse model evaluation. We have shown that when visual feedback consists of an image of the whole arm, subjects rely more on visual feedback at the INV stage than when only fingertip

position is shown (Figure 32b). We have hypothesized that this difference arises because of the many-to-one nature of the arm's kinematics. Since a single fingertip position can be achieved by many different arm postures, subjects might rely less on vision during fingertip-feedback trials in order to avoid the uncertainty of the visual-to-proprioceptive transformation. This uncertainty cost is reflected in the broadened population activity that results from transforming visual feedback of the fingertip into proprioceptive coordinates (Figure 36a, model neural activity in blue). When this activity pattern converges with the activity encoding the proprioceptive location of the arm (red), the peak of the resulting population activity (black), and therefore the location of $\hat{x}_{\mathrm{INV}}$, will be closer to the position originally encoded by proprioceptive signals (dashed red line). The fit value of α_{INV} will therefore be close to zero, as observed in our psychophysical data.

On the other hand, if the visual feedback shows the whole arm (and thus explicitly specifies the joint angles), the uncertainty introduced by transforming visual signals into proprioceptive coordinates will be reduced. This difference would result in a transformed population activity (Figure 36b, blue) that is less broad than the one resulting from transformed fingertip feedback. This sharper peak would in turn result in $\hat{x}_{\mathrm{INV}}$ being closer to location specified by visual feedback, and therefore an increased value of α_{INV}, consistent with the results shown in Figure 32b.

In the model circuits shown in Figures 35 and 36, the position of the integrated arm position estimate ($\hat{x}_{\mathrm{MV}}$ or $\hat{x}_{\mathrm{INV}}$) is an emergent property of the way that sensory signals are transformed and combined. In this "bottom-up" scheme, the relative weighting of vision and proprioception is determined by the shapes of the visual and proprioceptive population activities, the exact form of the cost of coordinate transformations, and the way in which population activities are combined in the output layer. The integration terms α_{MV} and α_{INV} are therefore not represented

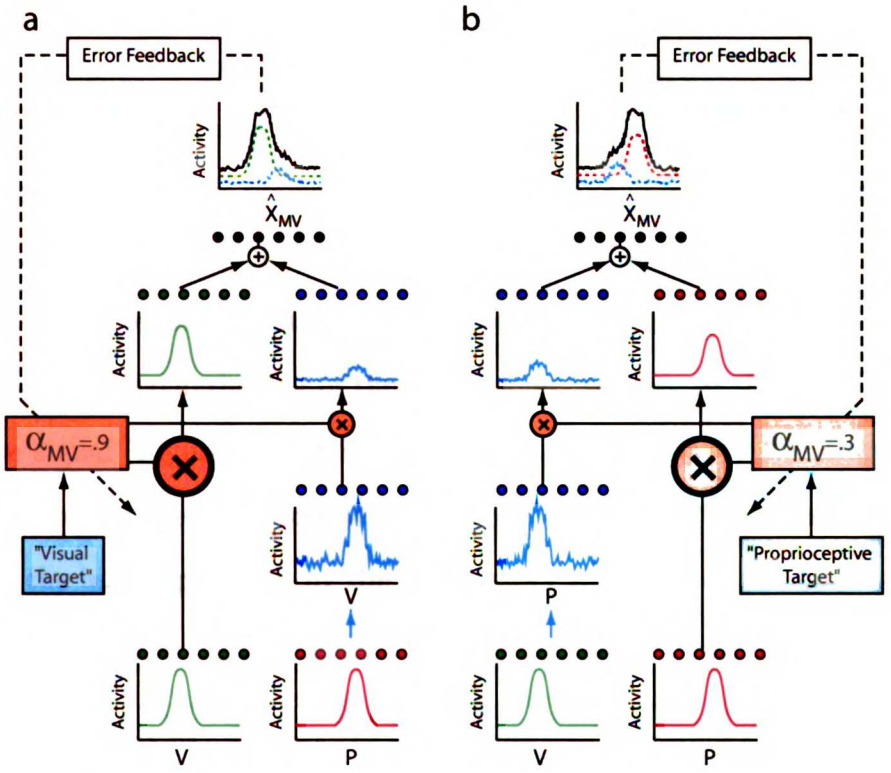

Figure 37: Target modality and sensory integration: top-down control. In this schema, α_{MV} is a learned parameter that explicitly controls sensory integration during reaches to visual (a) and proprioceptive (b) targets. Coordinate transformations increase the noisiness of population activity, leading to output noise (and task error) that scales with the reliance on transformed signals. To reduce this error, a learning mechanism (dashed lines) sets the activity in a network of cells (orange boxes) that influence sensory integration based on task context (blue boxes). These neurons scale the visual and proprioceptive inputs according to the learned value of α_{MV} (orange circles). Other conventions as in Figure 35.

explicitly by neural activity, but simply reflect the way that visual and proprioceptive signals propagate through the network.

In an alternative neural circuit, α_{MV} and α_{INV} might be controlled explicitly by the brain. Figure 37a and b show a model of sensory integration during reaches to visual and proprioceptive targets, respectively. There are several important differences between this schema and the one shown in Figures 35 and 36. First, the cost of transforming a sensory signal is the addition of noise to the population code, instead of a broadening of the activity. Second, the value of α_{MV} is represented by the activity of a network of neurons (orange boxes). This network introduces a multiplicative scaling to the two streams of sensory information such that a high value of α_{MV} leads to a strengthening of the visual signal relative to the transformed proprioceptive signal (Figure 37a), whereas a low α_{MV} strengthens the proprioceptive activity relative to the transformed visual signal (Figure 37b). Third, the output layer sums the inputs provided by the two sensory streams, rather than multiplying them.

The final key feature of this circuit is that the value of α_{MV} is learned based on the error feedback that the brain accumulates from the performance of various sensorimotor tasks. In this "top-down" scheme, the brain learns about the costs of coordinate transformations via task error. For example, when reaching to visual targets (Figure 37a), the brain would learn that the noise in the output population (black) increases rapidly with the reliance on transformed proprioceptive signals (blue). In order to decrease this noise cost, the brain would learn to set a high value of α_{MV}, corresponding to a higher gain for the visual signal than for the transformed proprioceptive signal. For reaches to proprioceptive targets, on the other hand, the brain would learn to minimize noise in the output population by decreasing the reliance on vision (Figure 37b).

Sensory integration – top-down or bottom-up?

The psychophysical methods developed in the first chapters of this thesis allow the relative weighting of vision and proprioception to be quantified with a fair degree of precision in a variety of sensorimotor tasks. However, data gathered with these techniques do not address the fundamental question of whether sensory integration is better characterized as a top-down or bottom-up process. In Chapter 2, we argued that the brain weights sensory inputs in order to minimize the variability of the integrated position estimates. This top-down characterization of sensory integration (similar to the one shown in Figure 37) requires that the brain learn (at least implicitly) the relationship between the mixture of sensory signals and task error. One advantage of this arrangement is its flexibility. For example, parameters such as the shape of the sensory population activity or the form of the transformation cost might vary from task to task. This type of network is explicitly designed to minimize task error (the definition of which might itself vary from task to task) and could therefore perform consistently across these conditions. A disadvantage of this type of architecture, however, is that it must build up a model of the relationship between sensory integration and task error in each behavioral context.

Bottom-up networks, in contrast, are simpler but perhaps less flexible. The performance of the network in Figure 35, for example, is highly dependent on the particular form of the cost of the coordinate transformation and the relative widths of the population activities. Although such networks can provide a simple account of our data, their sensitivity to model parameters suggests that they might not be able to explain sensory integration in a wider variety of contexts.

The model networks shown here represent only a small sample of the potential architectures that could give rise to the observed psychophysical data. The true neu-

ral implementation could be constrained by both psychophysical and physiological studies. The question of top-down versus bottom-up coding might be addressed by behavioral experiments that address the dynamics of sensory integration in new behavioral contexts. Such studies could shed light on whether the weighting of vision and proprioception is learned from experience or hardwired. In parallel with this work, neurophysiological studies could reveal the neural cost of performing coordinate transformations. These experiments might be based on an approach similar to that outlined earlier for identifying the computational role of single neurons. Together, such approaches could clarify the computational and physiological underpinnings of reach planning, and establish the links between them.

Literature Cited

Aglioti, S., DeSouza, J. F. X., and Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, 5(6):679–685.

Andersen, R. A., Snyder, L. H., Bradley, D. C., and Xing, J. (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience*, 20:303–330.

Atkeson, C. G. and Hollerbach, J. M. (1985). Kinematic features of unrestrained vertical arm movements. *Journal of Neuroscience*, 5(9):2318–2330.

Batista, A. P., Buneo, C. A., Snyder, L. H., and Andersen, R. A. (1999). Reach plans in eye-centered coordinates. *Science*, 285(5425):257–260.

Bernstein, N. (1967). *The coordination and regulation of movements.* Pergammon, London.

Bock, O. (1986). Contribution of retinal versus extraretinal signals towards visual localization in goal-directed movements. *Experimental Brain Research*, 64(3):476–482.

Brinkman, C. (1984). Supplementary motor area of the monkey's cerebral cortex: short- and long-term deficits after unilateral ablation and the effects of subsequent callosal section. *J. Neurosci.*, 4(4):918–929.

Buneo, C. A., Jarvis, M. R., Batista, A. P., and Andersen, R. A. (2002). Direct visuomotor transformations for reaching. *Nature*, 416(6881):632–636.

Caminiti, R., Johnson, P. B., Galli, C., Ferraina, S., and Burnod, Y. (1991). Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. *J. Neurosci.*, 11(5):1182–1197.

Caminiti, R., Johnson, P. B., and Urbano, A. (1990). Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *J. Neurosci.*, 10(7):2039–2058.

Campbell, M. J. and Machin, D. (1993). *Medical Statistics: A Commonsense Approach.* John Wiley & Sons, Chichester, 2 edition.

Cheney, P. D. and Fetz, E. E. (1980). Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of Neurophysiology*, 44(4):773–791.

Choe, C. S. and Welch, R. B. (1974). Variables affecting intermanual transfer and decay of prism adaptation. *Journal of Experimental Psychology*, 102(6):1076–1084.

Clower, D. and Boussaoud, D. (1998). Adaptive aftereffect depends on actual, not representational, feedback of limb position. *European Journal of Neuroscience*, 10:6715.

Cohen, M. M. (1967). Continuous versus terminal visual feedback in prism aftereffects. *Percept. Mot. Skills*, 24(3):1295–1302.

Cohen, M. M. (1973). Visual feedback, distribution of practice, and intermanual transfer of prism aftereffects. *Percept. Mot. Skills*, 37(2):599–609.

Colby, C. L. and Duhamel, J. R. (1996). Spatial representations for action in parietal cortex. *Cognitive Brain Research*, 5(1-2):105–115.

Desmurget, M., Epstein, C. M., Turner, R. S., Prablanc, C., Alexander, G. E., and Grafton, S. T. (1999). Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nat. Neurosci.*, 2(6):563–567.

Dizio, P. and Lackner, J. R. (1995). Motor adaptation to coriolis force perturbations of reaching movements: endpoint but not trajectory adaptation transfers to the nonexposed arm. *J. Neurophysiol.*, 74(4):1787–1792.

Draper, N. and Smith, H. (1998). *Applied Regression Analysis*. John Wiley & Sons, New York, 3rd edition.

Dum, R. P. and Strick, P. L. (2002). Motor areas in the frontal lobe of the primate. *Physiol. Behav.*, 77(4-5):677–682.

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, Boca Raton, Florida.

Enright, J. T. (1995). The nonvisual impact of eye orientation on eye-hand coordination. *Vision Research*, 35(11):1611–1618.

Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433.

Flash, T. and Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.*, 5(7):1688–1703.

Foley, J. E. and Maynes, F. (1969). Comparison of training methods in the production of prism adaptation. *Journal of Experimental Psychology*, 81(1):151–155.

Frisen, L. and Glansholm, A. (1975). Optical and neural resolution in peripheral vision. *Investigative Opthamology*, 14(7):528–536.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., and Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 2(11):1527–1537.

Ghahramani, Z. (1995). *Computation and Psychophysics of Sensorimotor Integration.* Thesis/dissertation, Massachusetts Institute of Technology.

Ghilardi, M. F., Gordon, J., and Ghez, C. (1995). Learning a visuomotor transformation in a local area of work space produces directional biases in other areas (vol 73, pg 2535, 1995). *Journal of Neurophysiology*, 74(6):U28–U28.

Good, P. (2000). *Permutation Tests.* Springer Series in Statistics. Springer-Verlag, New York, 2 edition.

Goodale, M. A. and Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1):20–25.

Goodbody, S. J. and Wolpert, D. M. (1999). The effect of visuomotor displacements on arm movement paths. *Experimental Brain Research*, 127(2):213–223.

Gordon, J., Ghilardi, M. F., Cooper, S. E., and Ghez, C. (1994a). Accuracy of planar reaching movements .2. systematic extent errors resulting from inertial anisotropy. *Experimental Brain Research*, 99(1):112–130.

Gordon, J., Ghilardi, M. F., and Ghez, C. (1994b). Accuracy of planar reaching movements .1. independence of direction and extent variability. *Experimental Brain Research*, 99(1):97–111.

Graziano, M. S. A. (1999). Where is my arm? the relative role of vision and proprioception in the neuronal representation of limb position. *Proceedings of the National Academy of Sciences of the United States of America*, 96(18):10418–10421.

Graziano, M. S. A., Cooke, D. F., and Taylor, C. S. R. (2000). Coding the location of the arm by sight. *Science*, 290(5497):1782–1786.

Haffenden, A. M., Schiff, K. C., and Goodale, M. A. (2001). The dissociation between perception and action in the ebbinghaus illusion: Nonillusory effects of pictorial cues on grasp. *Current Biology*, 11(3):177–181.

Hardt, M. E., Held, R., and Steinbach, M. (1971). Adaptation to displaced vision: a change in the central control of sensorimotor coordination. *Journal of Experimental Psychology*, 89(2):229–239.

Harris, C. (1963). Adaptation to displaced vision: Visual, motor, or proprioceptive change? *Science*, 140:812–813.

Hay, J. and Pick, R. (1966). Visual and proprioceptive adaptation to optical displacement of the visual stimulus. *Journal of Experimental Psychology*, 71(1):150–158.

Helmholtz, H. v. (1925). *Treatise on physiological optics*, volume 3. Optical Society of America, Rochester, New York.

Henriques, D. Y. P., Klier, E. M., Smith, M. A., Lowy, D., and Crawford, J. D. (1998). Gaze-centered remapping of remembered visual space in an open- loop pointing task. *Journal of Neuroscience*, 18(4):1583–1594.

Hoffman, D. S. and Strick, P. L. (1995). Effects of a primary motor cortex lesion on step-tracking movements of the wrist. *J. Neurophysiol.*, 73(2):891–895.

Howard, I. P. and Templeton, W. B. (1966). *Human Spatial Orientation.* John Wiley & Sons, London.

Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision Res.*, 39(21):3621–3629.

Jordan, M. (1996). Computational aspects of motor control and motor learning. In Heuer, H. and Keele, S., editors, *Handbook of Perception and Action: Motor Skills.* Academic Press.

Kakei, S., Hoffman, D. S., and Strick, P. L. (1999). Muscle and movement representations in the primary motor cortex. *Science*, 285(5436):2136–2139.

Kakei, S., Hoffman, D. S., and Strick, P. L. (2001). Direction of action is represented in the ventral premotor cortex. *Nat. Neurosci.*, 4(10):1020–1025.

Kalaska, J. F., Cohen, D. A., Prud'homme, M., and Hyde, M. L. (1990). Parietal area 5 neuronal activity encodes movement kinematics, not movement dynamics. *Exp. Brain Res.*, 80(2):351–364.

Kalil, R. E. and Freedman, S. J. (1966). Intermanual transfer of compensation for displaced vision. *Percept. Mot. Skills*, 22(1):123–126.

Kitazawa, S., Kimura, T., and Uka, T. (1997). Prism adaptation of reaching movements: specificity for the velocity of reaching. *J. Neurosci.*, 17(4):1481–1492.

Kurata, K. (1991). Corticocortical inputs to the dorsal and ventral aspects of the premotor cortex of macaque monkeys. *Neurosci. Res.*, 12(1):263–280.

Kurata, K. and Hoshi, E. (2002). Movement-related neuronal activity reflecting the transformation of coordinates in the ventral premotor cortex of monkeys. *J. Neurophysiol.*, 88(6):3118–3132.

Lackner, J. R. (1974). Adaptation to displaced vision - role of proprioception. *Perceptual and Motor Skills*, 38(3):1251–1256.

Li, C. S., Padoa-Schioppa, C., and Bizzi, E. (2001). Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron*, 30(2):593–607.

Malach, R., Amir, Y., Harel, M., and Grinvald, A. (1993). Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc. Natl. Acad. Sci. U. S. A.*, 90(22):10469–10473.

McIntyre, J., Stratta, F., Droulez, J., and Lacquaniti, F. (2000). Analysis of pointing errors reveals properties of data representations and coordinate transformations within the central nervous system. *Neural Comput.*, 12(12):2823–2855.

Messier, J. and Kalaska, J. (1997). Differential effect of task conditions on errors of direction and extent of reaching movements. *Experimental Brain Research*, 115(3):469–478.

Milner, A. D. and Goodale, M. A. (1995). *The Visual Brain in Action.* Oxford Psychology Series. Oxford University Press, Oxford.

Padoa-Schioppa, C., Li, C. S., and Bizzi, E. (2002). Neuronal correlates of kinematics-to-dynamics transformation in the supplementary motor area. *Neuron*, 36(4):751–765.

Paillard, J. (1996). Fast and slow feedback loops for the visual correction of spatial errors in a pointing task: A reappraisal. *Canadian Journal of Physiology and Pharmacology*, 74(4):401–417.

Passingham, R. E. (1988). Premotor cortex and preparation for movement. *Exp. Brain Res.*, 70(3):590–596.

Pouget, A., Dayan, P., and Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, 26:381–410.

Prablanc, C. and Martin, O. (1992). Automatic control during hand reaching at undetected 2-dimensional target displacements. *Journal of Neurophysiology*, 67(2):455–469.

Rossetti, Y., Desmurget, M., and Prablanc, C. (1995). Vectorial coding of movement: vision, proprioception, or both? *J. Neurophysiol.*, 74(1):457–463.

Rushworth, M. F., Johansen-Berg, H., and Young, S. A. (1998). Parietal cortex and spatial-postural transformation during arm movements. *J. Neurophysiol.*, 79(1):478–482.

Sabes, P. N. and Jordan, M. I. (1997). Obstacle avoidance and a perturbation sensitivity model for motor planning. *Journal of Neuroscience*, 17(18):7119–7128.

Sainburg, R. L., Lateiner, J. E., Latash, M. L., and Bagesteiro, L. B. (2003). Effects of altering initial position on movement direction and extent. *J. Neurophysiol.*, 89(1):401–415.

Sainburg, R. L. and Wang, J. (2002). Interlimb transfer of visuomotor rotations: independence of direction and final position information. *Exp. Brain Res.*, 145(4):437–447.

Salinas, E. and Abbott, L. F. (1995). Transfer of coded information from sensory to motor networks. *J. Neurosci.*, 15(10):6461–6474.

Salinas, E. and Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proc. Natl. Acad. Sci. U. S. A.*, 93(21):11956–11961.

Schwartz, A. B., Moran, D. W., and Reina, G. A. (2004). Differential representation of perception and action in the frontal cortex. *Science*, 303(5656):380–383.

Scott, S. H. and Kalaska, J. F. (1997). Reaching movements with similar hand paths but different arm orientations. i. activity of individual cells in motor cortex. *J. Neurophysiol.*, 77(2):826–852.

Scott, S. H., Sergio, L. E., and Kalaska, J. F. (1997). Reaching movements with similar hand paths but different arm orientations .2. activity of individual cells in dorsal premotor cortex and parietal area 5. *Journal of Neurophysiology*, 78(5):2413–2426.

Shadmehr, R. and Mussa-Ivaldi, F. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5):3208–3224.

Sober, S. J. and Sabes, P. N. (2003). Multisensory integration during motor planning. *J. Neurosci.*, 23(18):6982–6992.

Sober, S. J. and Sabes, P. N. (2005). Flexible strategies for sensory integration during motor planning. *Nature Neuroscience*, in press.

Soechting, J. F. and Flanders, M. (1989a). Errors in pointing are due to approximations in sensorimotor transformations. *Journal of Neurophysiology*, 62(2):595–608.

Soechting, J. F. and Flanders, M. (1989b). Sensorimotor representations for pointing to targets in 3- dimensional space. *Journal of Neurophysiology*, 62(2):582–594.

Soechting, J. F. and Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. *J. Neurosci.*, 1(7):710–720.

Taub, E. and Goldberg, L. A. (1973). Prism adaptation: control of intermanual transfer by distribution of practice. *Science*, 180(87):755–757.

Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: a model. *Nature Neuroscience*, 3(4):391–398.

Uhlarik, J. J. and Canon, L. K. (1971). Influence of concurrent and terminal exposure conditions on nature of perceptual adaptation. *Journal of Experimental Psychology*, 91(2):233.

Uno, Y., Kawato, M., and Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement: Minimum torque-change model. *Biol. Cybern.*, 61(2):89–101.

van Beers, R. J., Sittig, A. C., and van der Gon, J. J. D. (1998). The precision of proprioceptive position sense. *Experimental Brain Research*, 122(4):367–377.

van Beers, R. J., Sittig, A. C., and van der Gon, J. J. D. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of Neurophysiology*, 81(3):1355–1364.

Wallace, B. and Redding, G. M. (1979). Additivity in prism adaptation as manifested in intermanual and interocular transfer. *Percept Psychophys.*, 25(2):133–136.

Warren, D. H. and Schmitt, T. L. (1978). On plasticity of visual-proprioceptive bias effects. *Journal of Experimental Psychology-Human Perception and Performance*, 4(2):302–310.

Welch, R. B. (1986). Adaptation of space perception. In Boff, K. R., Kaufman, L., and Thomas, J. P., editors, *Handbook of Perception and Human Performance: Sensory Processes and Perception*, volume 1, pages 24.1–24.45. Wiley.

Welch, R. B., Choe, C. S., and Heinrich, D. R. (1974). Evidence for a 3-component model of prism adaptation. *Journal of Experimental Psychology*, 103(4):700–705.

Welch, R. B. and Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin*, 88(3):638–667.

Welch, R. B., Widawski, M. H., Harrington, J., and Warren, D. H. (1979). Examination of the relationship between visual capture and prism adaptation. *Perception & Psychophysics*, 25(2):126–132.

Wise, S. P., Boussaoud, D., Johnson, P. B., and Caminiti, R. (1997). Premotor and parietal cortex: Corticocortical connectivity and combinatorial computations. *Annual Review of Neuroscience*, 20:25–42.

Wright, M. J. and Johnston, A. (1983). Spatiotemporal contrast sensitivity and visual field locus. *Vision Research*, 23(10):983–989.