Vision Research 119 (2016) 82–98

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Sinusoidal error perturbation reveals multiple coordinate systems for sensorymotor adaptation ☆

Todd E. Hudson *, Michael S. Landy

Department of Psychology and Center for Neural Science, New York University, New York, NY 10003, United States

ARTICLE INFO

Article history:
Received 27 August 2015
Received in revised form 21 November 2015
Accepted 2 December 2015

Keywords:
Motor adaptation
Movement error
Plasticity
Internal representation
Coordinate system
Movement encoding
Movement code

ABSTRACT

A coordinate system is composed of an encoding, defining the dimensions of the space, and an origin. We examine the coordinate encoding used to update motor plans during sensory-motor adaptation to center-out reaches. Adaptation is induced using a novel paradigm in which feedback of reach endpoints is perturbed following a sinewave pattern over trials; the perturbed dimensions of the feedback were the axes of a Cartesian coordinate system in one session and a polar coordinate system in another session. For center-out reaches to randomly chosen target locations, reach errors observed at one target will require different corrections at other targets within Cartesian- and polar-coded systems. The sinewave adaptation technique allowed us to simultaneously adapt both dimensions of each coordinate system (x–y, or reach gain and angle), and identify the contributions of each perturbed dimension by adapting each at a distinct temporal frequency. The efficiency of this technique further allowed us to employ perturbations that were a fraction the size normally used, which avoids confounding automatic adaptive processes with deliberate adjustments made in response to obvious experimental manipulations. Subjects independently corrected errors in each coordinate in both sessions, suggesting that the nervous system encodes both a Cartesian- and polar-coordinate-based internal representation for motor adaptation. The gains and phase lags of the adaptive responses are not readily explained by current theories of sensory-motor adaptation.

Motor adaptation is fundamental to the neural control of movement, affording an automatic process to maintain a consistent relationship between motor plans and movement outcomes. That is, adaptation is described as updating an internal mapping between desired motor outcome and motor output (Sanger, 2004; Shadmehr, Smith, & Krakauer, 2010), not a deliberate corrective action. Here, using a method that relies on extremely small perturbations that are unlikely to recruit explicit corrective mechanisms, we revisit one of the central questions of sensory-motor adaptation: coordinate encoding.

A coordinate system is composed of an encoding, defining the dimensions of the space, and an origin. Here, we examine the coordinate encoding used to update motor plans during sensory-motor adaptation rather than, as in most experiments on the topic of coordinate systems (e.g., Henriques & Crawford, 2000; McGuire & Sabes, 2009, 2011; Park, Schlag-Rey, & Schlag, 2006; Ren et al., 2006; Thaler & Todd, 2009a), the origin of that coordinate system. During adaptation the internal mapping from motor plans to movement outcomes must be updated in response to error feedback, and the coordinate system used to encode motor errors has important consequences for the types of errors to which the central nervous system can subsequently adapt. The encoding that has historically been most often proposed/assumed for reach planning and adaptation is a vector-based (polar) code (Bock & Eckmiller, 1986; Cressman & Henriques, 2011; de Graaf, van der Gon, & Sittig, 1996; Georgopoulos, Kalaska, & Massey, 1981; Ghez et al., 1997; Ghilardi et al., 2000; Gordon, Ghilardi, & Ghez, 1994; Hudson & Landy, 2012b; Krakauer, Pine, Ghilardi, & Ghez, 2000; Messier & Kalaska, 1997; Ojakangas & Ebner, 1991; Pesaran, Nelson, & Andersen, 2006; Pine, Krakauer, Gordon, & Ghez, 1996; Tanaka, Sejnowski, & Krakauer, 2009; Taylor & Ivry, 2011; Turnham, Braun, & Wolpert, 2012; Vindras, Desmurget, & Viviani, 2005; Vindras & Viviani, 2002; Wu & Smith, 2013), although an endpoint-based Cartesian encoding has also been advocated (Baud-Bovy & Viviani, 1998; Bedford, 1999; Ghahramani, Wolpert, & Jordan, 1996; Hudson & Landy, 2012b; Thaler & Todd, 2009b; Wu & Smith, 2013).

☆ *Acknowledgement:* NIH Grant EY08266.

* Corresponding author at: New York University Medical Center, Ambulatory Care Center, 220 E38th St, 17th Floor, New York, NY 10016, United States.
E-mail address: teh9@columbia.edu (T.E. Hudson).

http://dx.doi.org/10.1016/j.visres.2015.12.005
0042-6989/

Updated motor plans can produce quite different outcomes when the underlying mapping is coded in polar vs. Cartesian coordinates, even when based on the identical reach and observed error. Consider a calibrated system that has just executed a rightward reach, where feedback is perturbed to indicate a rightward error (Fig. 1A). To maintain calibration, Cartesian- and polar-coded systems update their mappings in ways that have precisely the opposite effect for a subsequent reach to the left. Within a polar-coded system the observed error is an overshoot, which is corrected by reducing the gain of the next reach (Fig. 1B). This leads to rightward compensation for the subsequent leftward reach. The same observed error is compensated by a leftward shift within a Cartesian-coded system (Fig. 1C).

This example relies on the idea that updates in one or the other coordinate system occur in response to error observed in the previous reach, and generalize to other targets in the workspace. Two questions immediately present themselves: (1) Are sensory-motor adaptive responses encoded in Cartesian, polar, or both coordinate systems? (2) Do adaptive responses generalize from one target/workspace location to another? The first question is, of course, primary: it is difficult to define the extent of reach generalization without first knowing the coordinates in which generalization occurs. Although this question has been previously addressed

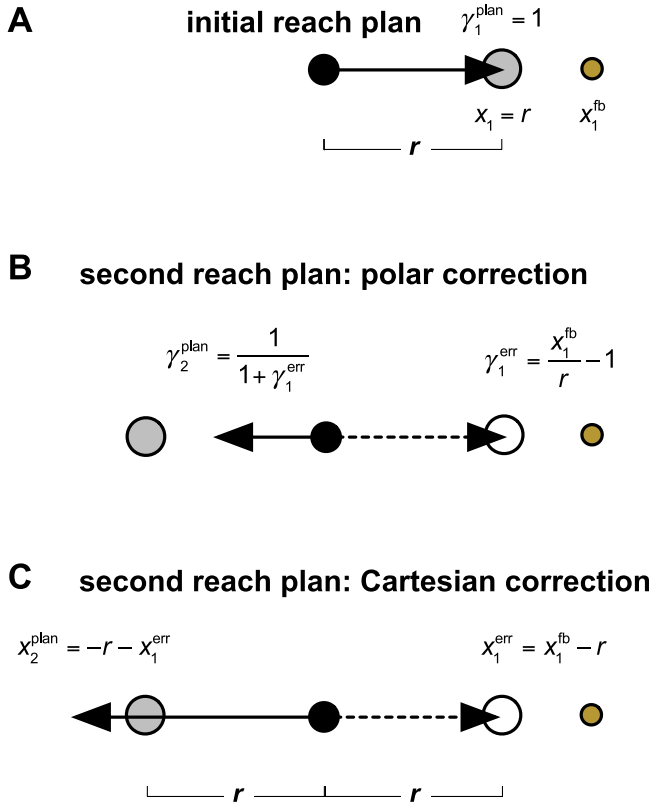

Fig. 1. Hypothetical corrections produced by polar- and Cartesian-coded adaptive systems toward two targets at a distance r from the start position. (A) Original reach plan (arrow) from the center position (black circle) to the target (grey circle). Here, the reach plans for movement trial 1, both polar gain (γ_1^{plan}) and horizontal Cartesian (x_1^{plan}) are assumed calibrated and noiseless. (B) For polar-coded adaptation, the reach error (small, rust-filled circle) from the plan illustrated in A is coded as an overshoot (by a factor of $1+\gamma_1^{err}$, $\gamma_1^{err} > 0$), and the reach on trial 2 (toward the new grey circle) is planned with a smaller gain (of $[1+\gamma_1^{err}]^{-1}$) to completely compensate. (C) For Cartesian-coded adaptation, the endpoint (small, rust-filled circle) is coded as rightward relative to the target in (A) ($x_1^{err} > 0$). The plan for the reach on trial 2 is therefore compensated leftward relative to the target (i.e., $x_2^{plan} = -r - x_1^{err}$ to completely compensate). Compensations produced within these two encodings are in opposite directions relative to the second target position. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(Wu & Smith, 2013), these earlier experiments—and indeed most work involving reach adaptation—suffer from an artifact of the perturbation paradigm used. Most previous research on reach adaptation (e.g., Berniker & Körding, 2011; Krakauer et al., 2000; Mazzoni & Krakauer, 2006; Neva & Henriques, 2013; Pine et al., 1996) involves a large perturbation (usually on the order of 30–50% change in reach gain and 30–60° change in reach direction, corresponding to several cm in the reach workspace and well outside the range of a typical motor error) that is suddenly introduced and, many trials later, just as suddenly removed. The use of sudden, large perturbations makes it difficult to determine the coordinate encoding of adaptation within the motor system; indeed, even gradually introduced perturbations generally continue to be introduced until they match the typical (large) perturbation size (e.g., Hudson, Lackner, & DiZio, 2005; Kagerer, Contreras-Vidal, & Stelmach, 1997), and will therefore be well above threshold by the end of the adaptation period. Such large, obvious perturbations are outside the normal range of simple motor errors and elicit deliberate corrective action (Criscimagna-Hemminger, Bastian, & Shadmehr, 2010; Harris, 1974; Taylor & Ivry, 2012).

In support of the idea of explicit corrections, consider the study of Thoroughman and Taylor (2005), in which a series of increasingly complex force-field perturbations were studied. Their data show a graded response to this set of force fields, such that the easily described ('less complex') fields are more fully compensated than their difficult-to-describe ('complex field') perturbation. These results are consistent with the notion that subjects explicitly corrected for the less complex fields whose patterns were easy to identify and describe, but were less able to explicitly correct for the field whose pattern was too complex to consciously identify. They are also consistent with combinations of explicit and implicit corrections.

As demonstrated by this example, the critical problem with using large perturbations to study adaptation is that we do not know what aspects of 'adapted' motor responses are due to deliberate corrective action designed to counteract an obvious experimenter-imposed perturbation and which are due to automatic adaptive processes. In these circumstances, we cannot be certain whether any adaptive response was generated at all, let alone in what coordinates. Similar problems muddy the interpretation of generalization experiments, in which the typical method is to practice reaching under a large perturbation at a single target location, with subsequent test reaches (often without feedback) to other workspace locations (e.g., Krakauer et al., 2000). If subjects practice perturbed reaches to a single target location and the perturbation is large, the subject has insufficient information to know whether the explicitly planned correction used at the single feedback location should also be used for movements to other targets (i.e., whether reaches to these locations will be modified by the same experimenter-imposed perturbation as reaches to the current target), and must decide to what degree their corrective strategy should be applied to new targets.

To elicit adaptation to unnoticed perturbations, our experiments feature perturbed error feedback following a sinewave defined over consecutive reaches (Fig. 2) based on a method we have developed previously (Hudson & Landy, 2012b). This method allows continuous variation in the size and direction of perturbation in each dimension of a coordinate system, and has the feature that errors in each dimension of the coordinate system may be perturbed at a separate temporal frequency within a single session, allowing the resulting frequency-tagged motor responses to be subsequently analyzed for possible interactions. We have shown that this method elicits measurable adaptation using perturbations that are a fraction the size of those normally used (Hudson & Landy, 2012b). For example, we use a maximum angular perturbation of about 4°, compared with the 30°–90° perturbations used

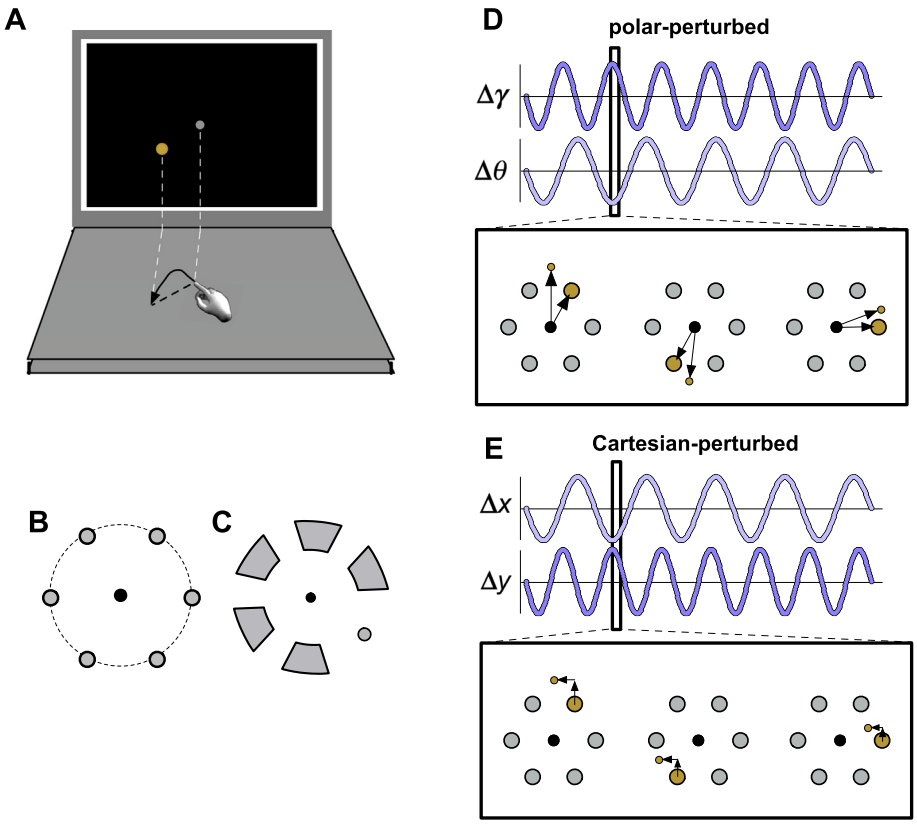

Fig. 2. Stimuli and perturbations. (A) Reach start and target positions displayed on the monitor indicated that reaches were to be made between corresponding positions on the tabletop. (B) Target arrays for the six-target experiment. (C) Target arrays for the generalization experiment. Circles (B and C) indicate target locations that were identical throughout a session. Extended feedback zones (C, generalization experiment) indicate ranges of possible target locations where feedback was provided. Within each of the five feedback zones, target location was selected unpredictably. Feedback of reach endpoints was perturbed following a sinusoidal pattern, as a function of trial number over the course of an experiment. Feedback was provided for all targets in the six-target experiment (B) and for reaches to the five feedback zones in the generalization experiment (C). Endpoint perturbations are encoded either in polar (D) or Cartesian (E) coordinates. Cutouts under each pair of sinewaves represent hypothetical three-trial sequences occurring within the six-target experiment at the time indicated by the dark vertical bar. The target presented in each of the three trials is highlighted in the diagram (rust-filled circles); note that only one target was visible on any given trial. The three-trial segments shown here highlight the fact that sequences of perturbations, even those beginning with the identical perturbation, diverge quickly within the two codes. To highlight the differences between perturbations in the two coordinate systems, perturbations are shown as applied to reaches that landed perfectly on target. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

previously. Use of small, continuously changing perturbations prevents the active, deliberate correction of large perturbations that might otherwise contaminate our results. To allow separate polar- and Cartesian-based perturbations, our experiments feature reaches to a set of targets whose order is randomly selected from a list of six possible target locations located symmetrically around a central start position (i.e., in a six-target extension of Fig. 1).

Our first experiment tests the predicted pattern of adaptation that would be produced by polar-only, Cartesian-only and combined adaptive systems in situations similar to that depicted in Fig. 1. In this experiment perturbation is experienced at each of six target locations separated by 60° (Fig. 2B), so that the adaptation mechanism experiences feedback for reaches in many directions and regions in the workspace. In separate sessions, sinewave perturbation is applied in the gain and angle dimensions of a polar coordinate system and in the x and y dimensions of a Cartesian coordinate system. Although our experiments are symmetric in the two coordinate systems we focus on polar-based updating, which has received the majority of previous attention—and in at least one influential series of papers has been proposed as the only coordinate system used during sensory-motor adaptation (Bock & Eckmiller, 1986; Gordon et al., 1994; Krakauer et al., 2000; Pine et al., 1996). If only polar-encoded errors are used for sensory-motor adaptation, the measured adaptive response to Cartesian perturbation would not mimic the sinusoidal pattern of the perturbation, since a mechanism sensitive to polar-encoded errors would interpret such a sequence as noise. Instead, our results demonstrate robust adaptation to both polar and Cartesian perturbation, suggesting the existence of a dedicated Cartesian-based adaptive mechanism in addition to a polar-encoded mechanism. Our second experiment directly examines generalization of adaptation in both coordinate systems. This experiment makes use of a range of training target locations (where perturbations are experienced; Fig. 2C), the closest of which are separated by at least 45° from a single generalization, or 'no-feedback', target; i.e., the generalization target (small circle, Fig. 2C) sits at the center of a 90° wedge of the workspace where no feedback is provided. Results demonstrate robust adaptation in all dimensions of both coordinate systems at the generalization target. These experiments support the hypothesis that sensorimotor reach adaptation (as opposed to explicit corrections brought on by large experimental perturbations) is dual-coded in Cartesian and polar coordinates. However, the importance of our result is not simply at the level of supporting a broad hypothesis regarding the coordinate encoding of sensory-motor adaptation. We further present evidence that our experimental results (even within a single coordinate system) are not readily accounted for by current models of sensory-motor adaptation. Since these models were primarily developed based on previous large-perturbation experiments, this raises questions regarding the use of data from such experiments to develop

models of sensory-motor adaptation. Finally, our results lead us to argue that many previous experimental results falling under the rubric of adaptation may in fact result from explicit correction to a perceptually obvious perturbation rather than an automatic adaptive process, a notion that has received recent attention (Heuer & Hegele, 2011; Taylor & Ivry, 2011, 2012; Taylor, Krakauer, & Ivry, 2014).

1. Methods

We investigated the coordinate encodings used by the CNS for sensory-motor adaptation by introducing sinewave-perturbed visual feedback, shown on a frontal monitor, of reach endpoints performed on a tabletop (Fig. 2A) in two experiments. In the first 'six-target' experiment, targets were six positions equally spaced around the central starting point (Fig. 2B) and visual feedback was provided following reaches to all six targets. In the second 'generalization' experiment, targets were either chosen from one of five target zones for which feedback was provided, or from a punctate generalization target at which feedback was withheld (Fig. 2C). In each of the two experiments, endpoints were perturbed in one session following polar-coded coordinates (Fig. 2D), and also in Cartesian-coded coordinates in a separate session (Fig. 2E).

1.1. Subjects

Two sets of eight subjects (10 female, median age 25, all but one naive as to the purposes of the experiments) participated in these experiments (2 subjects participated in both experiments). Each subject completed two sessions: one session for Cartesian-perturbed reaches, and one for polar-perturbed reaches, order randomized across subjects. The two sessions of each experiment were completed on separate days. All participants had normal or corrected-to-normal vision, and were members of the Department of Psychology or Center for Neural Science at New York University. Methods were approved by the NYU University Committee on Activities Involving Human Subjects. Subjects gave informed consent prior to testing.

1.2. Apparatus

Subjects were seated in a dimly lit room 42.5 cm away from a frontoparallel transparent polycarbonate screen mounted flush to the front of a 21″ computer monitor (Sony Multiscan G500, 1920 × 1440 pixels, 60 Hz; Sony, Tokyo, Japan). Reach trajectories were recorded using an Optotrak 3-d motion capture system (Northern Digital, Waterloo, Ontario, Canada) with two three-camera heads located above-left and above-right of the subject. Fingertip position was continuously monitored at 200 Hz. The experiment was run using the Psychophysics Toolbox software (Brainard, 1997; Pelli, 1997) and the Northern Digital API (for controlling the Optotrak). Details of the apparatus and method of calibration are the same as used previously (Hudson & Landy, 2012a, 2012b).

1.3. Procedure

Reaches were made from point to point on a tabletop to virtual targets that were presented on a display screen (Fig. 2A). Subjects were required to lift the finger off the tabletop to begin each reach, and finish the reach (re-touching the tabletop) within 300 ms (i.e., subjects made point-to-point reaches, not drawing/tracing movements). Points were awarded to maintain motivation; hits on the target earned subjects one point. Missing the target earned no points, and too-slow reaches were repeated. The fingertip endpoint, target location, and a running average of points per reach were displayed at reach completion, except for reaches to the 'no-feedback' target during the generalization experiment (see *Procedure: Generalization experiment*).

All reaches. Reaches were made to target locations arranged around a start position, i.e., 'center-out' reaches. Reaches were constrained to be completed within 300 ms. Timeouts were signaled mid-reach as the time constraint was surpassed, with an auditory tone and immediate blanking of the stimulus. Timed-out reaches were simply re-run, since no feedback would have been provided to interfere with the adaptive process (see below).

Each trial began with the start position shown on the monitor. When the fingertip was brought to within 2 cm of the tabletop location corresponding to the displayed start position, a fingertip indicator dot was shown (always unperturbed), allowing subjects to cover the location of the start position on the tabletop. Once the start position was covered (fingertip was within the 3.2 mm start circle radius and within 2 mm of the tabletop surface), the target was presented, and the subject was free to begin the movement when ready. At movement initiation (fingertip lifted 2 mm from the tabletop surface and outside the start circle), the start position and fingertip indicator (but not the target) were blanked. When feedback was provided at movement termination (fingertip re-touched the tabletop, i.e., it was below the 2 mm height threshold and velocity had dropped below 40 mm/s), a dot was displayed to indicate fingertip endpoint position. The location of the feedback dot was a combination of the true fingertip endpoint and a perturbation that varied sinusoidally over trials (see *Perturbations*).

Six-target experiment. In the first experiment, each reach was to one of six fixed target locations (one target every 60°, 115 mm from the start position, Fig. 2B) in unpredictable order (all targets equi-probable). Endpoint feedback (perturbed) was provided after all reaches. Subjects completed 180 reaches to each of these six targets for a total of 1080 reaches.

Generalization experiment. In the second experiment, reaches were made within five extended target zones and one punctate target location (Fig. 2C); we refer to the punctate location as the 'generalization target' and the five extended regions in which feedback was provided as 'feedback zones'. Each of the feedback zones represented a range of locations at which a target might be presented (distance range: 103–127 mm, angular range: 30°), where the target locations actually selected are referred to as feedback targets. Subjects completed 144 reaches to each feedback zone and 120 reaches to the generalization target in unpredictable order (shuffling the set of generalization and feedback zones), for a total of 840 reaches. Subjects were provided with endpoint feedback (perturbed) only following reaches to targets presented in one of the five feedback zones. Reaches to the generalization target were never followed by endpoint feedback. The entire configuration of feedback zones and generalization target was rotated by an unpredictable amount for each session. This prevented generalization targets being located along one of the cardinal axes (e.g., directly to the right of the start position) or even in the same position across sessions.

The placement and widths of target zones means that 60% of feedback targets appeared at locations >100° from the generalization target, with the remaining 40% appearing at locations between 45° and 75° from the generalization target. This would prevent robust generalization if the underlying mechanism responsible for updating in adaptation were narrowly tuned to the location of the observed error (e.g., neurons with a 25–30° half-width, such that responses to errors observed 45° away are essentially nonexistent; see Krakauer et al., 2000; Tanaka et al., 2009).

All experimental sessions began with a series of 72 unperturbed practice reaches, in which targets were presented randomly on an

annulus of radius 109–121 mm centered on the start position. Practice reaches were the same as experimental reaches in most respects, except that no points were awarded and reach endpoints were unperturbed. These initial reaches allowed subjects to become familiar with reach and time constraints.

1.4. Perturbations

Although there were two types of perturbation imposed on each subject's endpoints in separate sessions, the procedure was identical in both sessions; the only difference between the two sessions was the pattern of visuomotor perturbation applied to reach endpoint feedback.

Polar-coded perturbation. In polar-coded perturbation sessions, displayed feedback indicating fingertip endpoints was perturbed in a sinewave pattern over trials coded in γ–θ (reach gain and angle, respectively) coordinates, one sinusoid governing γ perturbations and a second sinusoid governing θ perturbations (Fig. 2D). That is, a reach on trial t with actual gain $\gamma(t)$ and direction $\theta(t)$ was displayed as if the gain was $\gamma^{fb}(t) = \gamma(t) + \Delta\gamma(t)$ (and a similar perturbation for direction), with

$$\begin{aligned}\Delta\gamma(t) &= -P_\gamma \sin(2\pi f_\gamma t/T)\\ \Delta\theta(t) &= -P_\theta \sin(2\pi f_\theta t/T),\end{aligned} \tag{1a}$$

where t is the trial number, T is the number of perturbed trials per session, P is the perturbation amplitude and f the corresponding perturbation frequency (in cycle/session). Perturbation frequencies were 7 cycle/session in γ and 5 cycle/session in θ. The amplitudes of the perturbations were equated to the perturbation amplitude used during Cartesian-coded sessions such that, for on-target reach endpoints, the maximum distance between an on-target reach endpoint and endpoint feedback was equal to that used during Cartesian perturbation. The signs of the γ- and θ-perturbation sinusoids were assigned unpredictably at the start of each session. We reverse the sign of the data if necessary to match Eq. (1a) (if the perturbing sinusoid was initially positive-going).

Cartesian-coded perturbation. In Cartesian-coded perturbation sessions, feedback indicating fingertip endpoints was perturbed in a sinewave pattern over trials coded in x–y coordinates (Fig. 2E). In analogy to Eq. (1a), we have

$$\begin{aligned}\Delta x(t) &= -P_x \sin(2\pi f_x t/T)\\ \Delta y(t) &= -P_y \sin(2\pi f_y t/T).\end{aligned} \tag{1b}$$

The perturbation amplitude was 7.5 mm in both x and y. The x-perturbation completed 5 cycles over the course of each experimental session and the y perturbation completed 7 cycles. Again, while the signs of x- and y-perturbation sinusoids were assigned unpredictably at the start of each session, we reverse the sign of the data if necessary to match Eq. (1b).

1.5. Data analysis

The major evidence implicating a particular encoding of reach errors would be compensatory reach errors in that coordinate system. That is, perturbation in the $-x$ or $-y$ directions would be compensated by reach endpoints deviated in the $+x$ or $+y$ direction relative to targets (Cartesian encoding); similarly, perturbation in the positive γ or θ directions would be compensated by reach endpoints deviating in the $-\gamma$ or $-\theta$ directions. That is, for fingertip feedback to land on the target center (assuming no motor noise), the perturbation and corrective response must sum to zero. For most models of adaptive correction (e.g., corrections programmed as a constant proportion of the previous error magnitude), unperturbed reach endpoints will be an inverted, scaled and shifted copy of the perturbing sinusoid defined by Eq. (1), i.e., with the same frequency. Thus, we predict a sequence of compensatory corrections to movement plans of the form $\Delta\gamma(t) = A_\gamma P_\gamma \sin[2\pi f_\gamma(t-\varphi_\gamma)/T]$, where $\Delta\gamma$ is the compensatory response, A_γ is the compensatory response amplitude relative to the perturbation, and φ_γ is the phase lag (in number of reaches). An analogous prediction holds for the other three coordinate dimensions (θ, x, and y).

1.5.1. Signal detection

Following Hudson and Landy (2012b), we formulate the problem of detecting the predicted adaptive response as a problem of Bayesian signal detection. Here the signal refers to the theoretical adaptive response that may be present in the data, and a signal is detected when there is significant positive evidence favoring signal-present over signal-absent models of the data. We ask, for example, whether the raw gain data are well modeled as $\Delta\gamma(t) = A_\gamma P_\gamma \sin[2\pi f(t-\varphi_\gamma)/T] + \beta + \varepsilon(t)$. Here, we allow for a possible sinusoidal response at any frequency f, a possible constant pointing bias β, and Gaussian motor noise ε that is independent over trials with standard deviation σ. Analogous equations are used to model the data for the other three coordinates.

We consider three models of the data. The dual-coding model (M_2) predicts that there will be an adaptive response that is a scaled and shifted version of the perturbing sinusoid in both coordinate systems. We consider two other models of the data: a single-coded model (M_1, in which either $A_\gamma = A_\theta = 0$ or $A_x = A_y = 0$) in which there is a Cartesian-coded or a polar-coded adaptive response, but not both, and a model in which there is no consistent adaptive response within either coordinate system (M_0, in which all response amplitudes are zero).

Evidence (Good & Toulmin, 1968) in dB favoring one model over its competitors (here, M_2 vs. M_1, M_0) is defined as:

$$\lambda = 10\log_{10} O(M_2), \tag{2}$$

where $O(M_2) = \frac{p(M_2|\boldsymbol{D}\cdot\imath)}{p(\sim M_2|\boldsymbol{D}\cdot\imath)} = \frac{p(M_2|\boldsymbol{D}\cdot\imath)}{p(M_0\vee M_1|\boldsymbol{D}\cdot\imath)}$ is the odds ratio favoring the dual-coding hypothesis in the data observed in our N subjects, $\boldsymbol{D} = \{\boldsymbol{d}_1, \ldots, \boldsymbol{d}_N\}$. The probability of the mth model is:

$$\begin{aligned}p(M_m|\boldsymbol{D}\cdot\imath) &\propto p(M_m|\imath)p(\boldsymbol{D}|M_m\cdot\imath)\\ &= p(M_m|\imath)\int d\Theta p(\Theta|M_m\cdot\imath)p(\boldsymbol{D}|\Theta\cdot M_m\cdot\imath),\end{aligned} \tag{3}$$

where Θ is used to represent all model parameters, both theoretically relevant (parameters describing adaptive responses: A, f, and φ) and irrelevant (β and σ), and $\imath$ is any other background information relevant to the problem (e.g., information about the range of possible motor variance values that occur in the normal human population). We use non-informative priors for $p(\Theta|M_m\cdot\imath)$: flat distributions for most parameters, and the Jeffreys prior proportional to $1/\sigma$ for motor noise. The likelihood term $p(\boldsymbol{D}|\Theta\cdot M_m\cdot\imath)$ is straightforward based on the Gaussian noise model, and the model prior assumes all three models are equally likely. The computation of Eq. (3) is performed by numerical integration as described by Hudson and Landy (2012b).

1.5.2. Parameter estimation

The signal detection equations described above do not rely on estimates of the signal/model parameters, but instead integrate over the probability distributions associated with all model parameters. Once this procedure yields evidence in favor of a particular model, it is then logical to estimate the parameters of that model (since prior to model selection, the epistemic validity of those parameters and their estimates as a description of the data is in question). We would therefore like to compute probability distributions over the three parameters of the sinusoidal signal contained in target-relative endpoints: $p(f|\boldsymbol{D}\cdot M\cdot\imath)$, $p(A|f\cdot\boldsymbol{D}\cdot M\cdot\imath)$,

and $p(\varphi|f \cdot \boldsymbol{D} \cdot M \cdot \imath)$. Note that distributions over amplitude and phase are conditional on the response frequency because we are interested in the amplitude and phase of the adaptive response at the response frequency. Here, there is no practical distinction between, e.g., $p(A|f \cdot \boldsymbol{D} \cdot M \cdot \imath)$ and $p(A|\boldsymbol{D} \cdot M \cdot \imath)$, since the only appreciable component of each adaptive response is at this frequency (Figs. 3 and 4).

We first compute the joint posterior distribution over the signal parameters,

$$p(f \cdot A \cdot \varphi|\boldsymbol{D} \cdot M \cdot \imath) \propto p(f \cdot A \cdot \varphi|M \cdot \imath) p(\boldsymbol{D}|f \cdot A \cdot \varphi \cdot M \cdot \imath)$$
$$= p(f \cdot A \cdot \varphi|M \cdot \imath) \prod_i \int d\sigma_i p(\sigma_i|M \cdot \imath) \int d\beta_i p(\beta_i|M \cdot \imath) p(\mathbf{d}_i|\beta_i \cdot \sigma_i \cdot f \cdot A \cdot \varphi \cdot M \cdot \imath), \quad (4)$$

where $\mathbf{d}_i$ is the data for subject i and we allow for different bias and noise parameters for each subject. A closed-form solution for this posterior has previously been derived (Hudson & Landy, 2012b). The 1D marginal distribution for the frequency of the response $p(f|\boldsymbol{D} \cdot M \cdot \imath)$ is computed by numerical integration of the 3D joint posterior. The 1D conditional marginal posteriors for the other two parameters are computed by numerical integration of $p(A \cdot \varphi|\boldsymbol{D} \cdot \hat{f} \cdot M \cdot \imath)$, where $\hat{f}$ is the maximum-likelihood estimate of the frequency.

All probability distributions and parameter estimates were derived from raw (un-binned and un-smoothed) data. Parameter estimates were derived as the mean of the corresponding marginal posterior and are reported as means ±1 standard deviation, as computed from the square root of the second central moment of the corresponding probability distribution. To reduce the visual appearance of noise, however, data plots (Figs. 3A and 4A) were created after first pooling across subjects and binning data and then removing outliers (>3 SD from the best-fit sinusoid).

2. Results

Data for both experiments and both polar and Cartesian-coded perturbations are shown in Figs. 3A and 4A. A compensatory response is easily visible in the data, as the noisy datapoints cluster about a sine wave that is a negative and slightly delayed version of the sinusoidally varying perturbation (solid black curve). This compensatory response is evident for both types of perturbation, for both coordinates of each perturbation, in both experiments, and for both the feedback and generalization targets in the generalization experiment. We first provide a model comparison that verifies this observation. We next examine the details of these compensatory responses, followed by analyses of the implications of our results for models of the encoding of errors for visuo-motor adaptation.

2.1. Model comparison: dual-coding hypothesis

We focus on three hypotheses: (H_0) corrective responses are either absent or are coded using a third coordinate system not considered above (other than Cartesian or polar coordinates); (H_1) only one of the two coordinate encodings is used, resulting in reliable compensation (and generalization) in response to

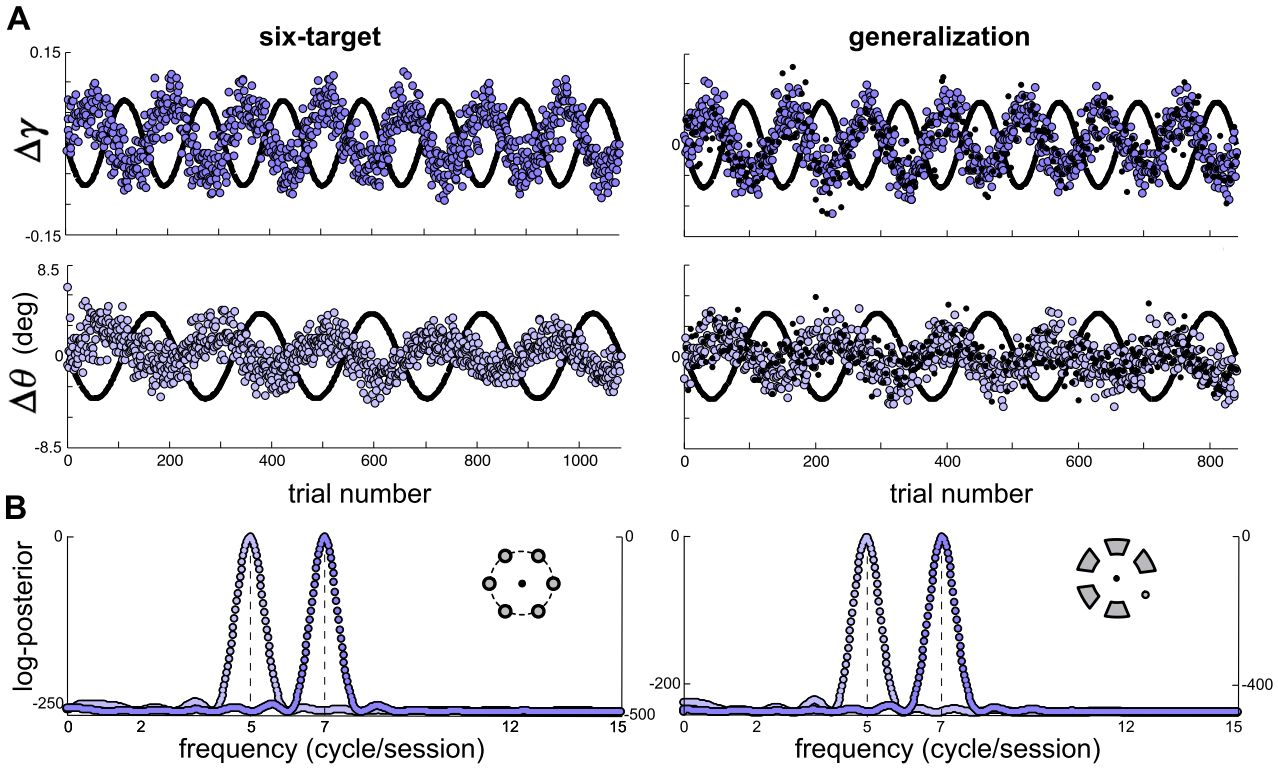

Fig. 3. Average endpoint data and frequency distributions: polar perturbation. (A) Average fingertip endpoint errors are plotted in terms of gain (γ: dark blue) and angle (θ: light-blue). Positive angles correspond to clockwise errors and positive gains correspond to overshoots of the target (i.e., an overall reach gain greater than one). γ and θ perturbations are plotted along with endpoint data as solid black curves. Complete correction would produce data that are the mirror image of the perturbing sinusoid reflected about the horizontal. For the generalization experiment, feedback- and generalization-target data are plotted separately as colored and black clrcles, respectively. Generalization data are plotted in 3-trial bins (i.e., at most 280 colored or black datapoints), and six-target data in 5-trial bins. (B) Log-posteriors over endpoint-error frequency in the perturbed dimensions are plotted in the same colors as data averages. These distributions peak at the respective perturbing frequencies (indicated by vertical dashed lines). Ordinate values correspond to the closer perturbation value (e.g., the ordinate describing data perturbed at 5 cycle/session is on the left). Icons indicate whether posteriors are derived from six-target (left column) or generalization experiments (right column). For the generalization experiment, log posteriors are based on feedback-target data. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

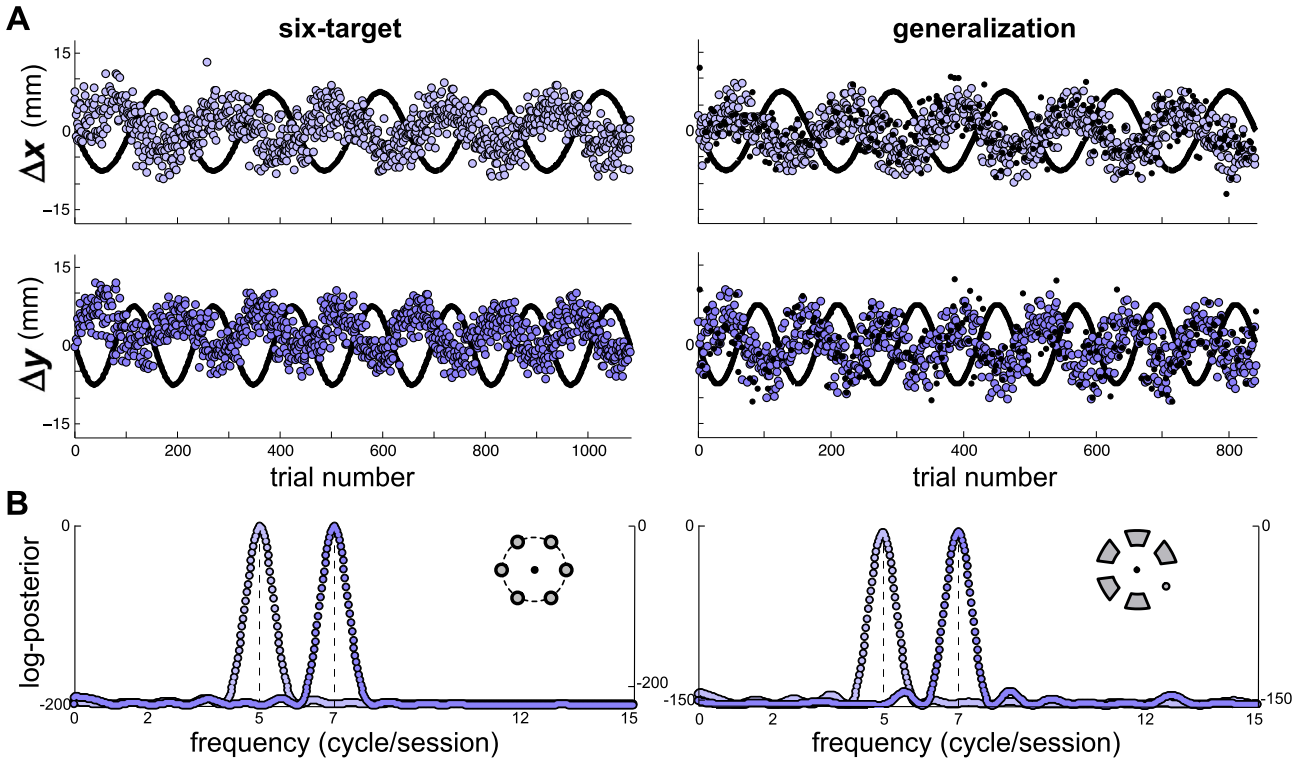

Fig. 4. Average endpoint data and frequency distributions: Cartesian perturbation. (A) Average fingertip endpoint errors are plotted in terms of *x*- (light-blue) and *y*-error (dark blue). (B) Log posteriors over endpoint-error frequency in the perturbed dimensions are plotted in the same colors as data averages, and peak at the perturbing frequencies (indicated by vertical dashed lines). Same plotting conventions as Fig. 3. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

perturbations in only one of the two sessions of each experiment; and (H_2) both coordinate encodings are used, resulting in compensation (and generalization) during both Cartesian- and polar-perturbed sessions of the two experiments. We simultaneously compare these three possible models of the compensatory response to perturbation using a Bayesian approach (Jaynes, 2003) based on an extension of the method of Hudson and Landy (2012b).

Results from polar and Cartesian perturbations are shown in Figs. 3 and 4, respectively. The results of the model comparison based on these data are shown in Table 1 and strongly support the dual-coding hypothesis, H_2. For example, in the six-target experiment there is 2893 dB of evidence in favor of H_2, equivalent to odds favoring H_2 of about 10^{289}:1 over the two alternatives. The evidence is still quite strong at the generalization target of the generalization experiment (363 dB), despite the reduced number of trials contributing to that model comparison.

As an alternative approach to the Bayesian model comparison, we can test whether Fourier power at the perturbing frequencies is significantly greater than zero, i.e., a statistical analog of a test of the difference between H_2 and the alternatives. Bootstrapped null-hypothesis distributions of power at the perturbing frequency were obtained by shuffling the time indices of each dataset 1000 times. Adaptation based on dual-coded movement plans would be expected to produce significant *p*-values for all perturbed dimensions at the perturbing frequency. Bootstrapping results are consistent with Bayesian model comparisons, with significant power at the perturbing frequencies ($p < .01$) in both sessions of both experiments for both coordinates (including separate analysis of the generalization target). The results of both experiments are therefore consistent with the hypothesis (H_2) that the CNS maintains a dual-coded adaptive mapping to compensate shifts, rotations and scalings of sensory-motor space.

Table 1
Evidence in dB for each hypothesis (H_0: no compensation, H_1: compensation in only one coordinate system, H_2: compensation in both) for both experiments and for both types of target in the generalization experiment. Note that negative evidence values indicate evidence against the corresponding hypothesis.

Experiment	Evidence for hypothesis (dB)		
	H_0	H_1	H_2
Six-target	−7156	−2893	2893
Generalization: feedback targets	−10,603	−3701	3701
Generalization: generalization target	−822	−363.1	363.1

2.2. Parameter estimates

Having established above that the adaptive response in all conditions can be characterized by a sinusoid, we next estimate the parameters of the sinusoidal model of the adaptive response. Log-posteriors for the frequency of compensation are shown in Figs. 3B and 4B, and posterior distributions for gain and phase of compensation at the perturbing frequency are shown in Fig. 5. All plots of posterior and log-posterior distributions are normalized to a peak of p = 1.

Estimated gain and phase parameters are given in Table 2 for all conditions and experiments. Estimated frequency values were all within 0.02 cycle/session of the perturbing frequency, evident in the log-posteriors of Figs. 3 and 4. In addition, log-posteriors show no evidence of cross talk between the perturbations within a given session (e.g., there is no secondary peak at 7 cycle/session within the *x*-dimension data of Fig. 4B) or of non-linear distortion products at 2 or 12 cycle/session in any of the log-posteriors. If the secondary frequency is outside the region of uncertainty for the primary response frequency (i.e., at the perturbed frequency), we can test for secondary response components as if this were

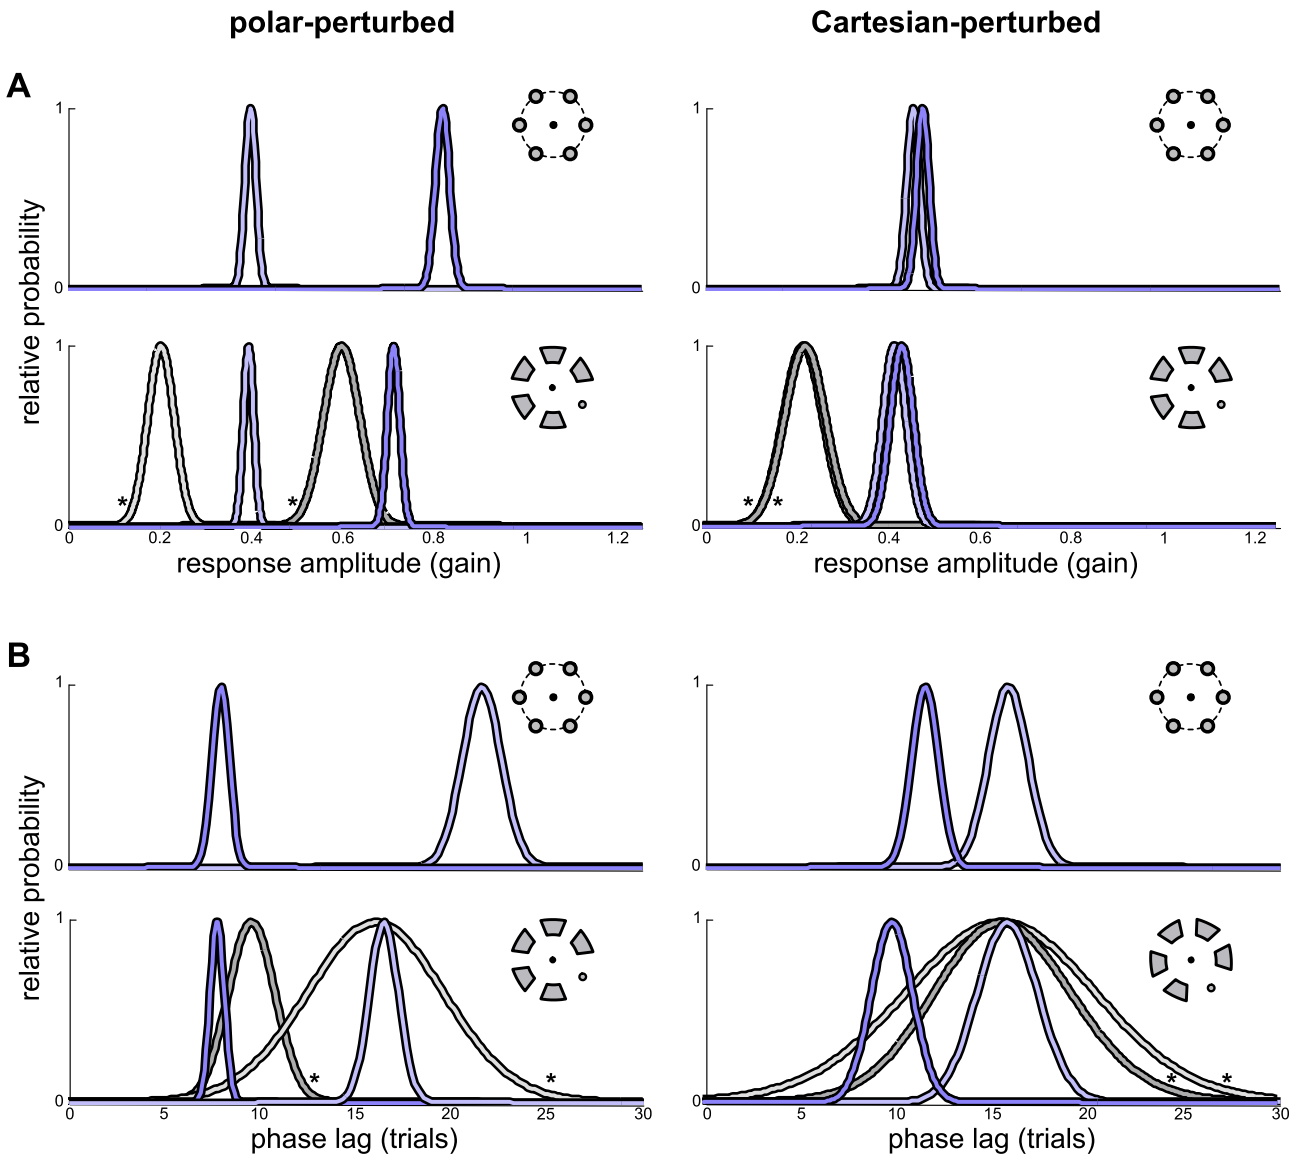

Fig. 5. Relative probability distributions over amplitude and lag of the sinusoidal corrective responses to Cartesian and polar perturbation. (A) Probability distributions over response amplitude, plotted as amplitude gain to facilitate comparison between polar- and Cartesian-perturbation data. (B) Probability distributions over the phase lag of the sinusoidal response to perturbation. Light-colored distributions correspond to θ- and x-dimension data (left and right columns, respectively), and dark-colored distributions correspond to γ- and y-dimension data (left and right columns). In plots derived from generalization-experiment data, feedback-target distributions are plotted in blue, and generalization-target distributions are plotted in grey. In all plots, icons indicate whether distributions are computed from six-target or generalization data. Colors are the same as those used throughout, where the greyscale lines are set to match the corresponding blue hue (i.e., light-grey is the generalization result corresponding to the light-blue feedback-target result). Generalization results are further identified by an asterisk. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the only component present (see Bretthorst, 1988). This condition holds for all plots in Figs. 3 and 4, where the region of uncertainty extends at most ±1 cycle from the perturbing frequency. In all cases of possible secondary response components (off-frequency within a dimension, or the sum or difference of perturbing frequencies), the estimated amplitude at the corresponding frequency was always less than 3% of the perturbing amplitude, which is not statistically reliable. Thus, we find no evidence for cross-talk between the coordinate axes within either encoding. This is the most common assumption made regarding sensory and/or motor adaptation in general, that it is the output of a system that is approximately linear. Indeed, many groups have gone a step further than this, and proposed a particular linear system, such as a Kalman filter (Korenberg & Ghahramani, 2002), a Kalman-inspired mechanism (Burge, Ernst, & Banks, 2008) or, most commonly, a first-order difference mechanism (Scheidt, Dingwell, & Mussa-Ivaldi, 2001; Scheidt & Stoeckmann, 2007; Shadmehr & Mussa-Ivaldi, 1994; Takahashi, Scheidt, & Reinkensmeyer, 2001; Thoroughman & Shadmehr, 2000).

Response gains (Fig. 5A) are approximately equal for the two dimensions of the Cartesian system (within each condition and experiment), but are always smaller in the θ relative to the γ dimension of the polar system. The pattern of gains and phase lags might be generated, for example, if adaptive corrections were updated based on a scaled copy of the error observed in the previous movement (i.e., a first-order difference mechanism) with increased learning rate in the γ dimension relative to the θ dimension.

The adaptation found for polar-perturbed reaches is qualitatively consistent with previous work using a step-function adaptation

Table 2
Parameter estimates for gain of compensation and phase lag (in trials) for each dimension, experiment, and target type.

Dimension		Six-target	Generalization: feedback targets	Generalization: generalization target
γ	Response gain	0.82 ± 0.016	0.72 ± 0.014	0.6 ± 0.04
	Phase lag	8.02 ± 0.46	7.8 ± 0.364	9.6 ± 1.24
θ	Response gain	0.41 ± 0.012	0.39 ± 0.011	0.2 ± 0.03
	Phase lag	21.7 ± 1.03	16.5 ± 0.77	16.1 ± 3.67
x	Response gain	0.45 ± 0.014	0.42 ± 0.015	0.22 ± 0.043
	Phase lag	15.9 ± 1.04	15.8 ± 0.865	15.6 ± 6.11
y	Response gain	0.47 ± 0.014	0.44 ± 0.014	0.22 ± 0.045
	Phase lag	11.5 ± 0.70	10.0 ± 0.588	15.7 ± 3.78

paradigm (i.e., perturbation introduced in a discrete step) with separate perturbations of the γ and θ dimensions during individual sessions (Baud-Bovy & Viviani, 1998; Krakauer et al., 2000; Pine et al., 1996; van den Dobbelsteen, Brenner, & Smeets, 2001; Vindras & Viviani, 2002), in that we find adaptation to occur more rapidly and to display an overall more robust response (both feedback and generalization) in the γ dimension relative to the θ dimension.

2.3. Six-target experiment: polar-only hypothesis

Despite having rejected both versions of a single-coordinate model of the adaptive process (see *Model comparison*), the long history of polar-only models (perhaps the most oft-cited being Krakauer et al., 2000) of the encoding used for sensory-motor adaptation suggests that we take a closer look at these models. In particular, we note that our model comparison tacitly assumes a global update, and therefore a model based on spatially selective updating may fall outside the bounds of the above analysis.

The polar-only hypothesis states that there is a single adaptive mechanism controlling reach adaptation that is based on an underlying polar representation and, by extension, that the sinusoidal modulation of reach adaptation to Cartesian perturbation shown in Fig. 4 is also the result of updating a polar encoding of motor errors. We distinguish three forms of the polar-only hypothesis:

(I) There is a single polar-coded compensation that is applied globally and updated on every trial.
(II) There is a separate polar-coded compensation for each zone of the workspace. On each trial, the current error is used to update all of these compensations. Each zone's compensation is updated by an amount that decreases with increasing distance between the current target and the zone associated with that stored compensation.
(III) There is a separate polar-coded compensation for each zone of the workspace. On each trial, the current error is used to update only that zone's stored compensation.

The robust adaptation to Cartesian perturbations is clearly incompatible with Hypothesis I (global, polar-only adaptation). However, a polar-only encoding of error could possibly result in the successful adaptation to Cartesian perturbation in the six-target experiment (Fig. 4A) if the subject separately adapted to the perturbation at each target. In this section, we derive an estimate of the *generalization tuning function* of Hypothesis II, i.e., the proportion of observed reach error that is corrected on each subsequent reach. It is maximal for repeated reaches to the same target location, and falls off as separation increases (Ghahramani et al., 1996; Krakauer et al., 2000; Turnham et al., 2012). We ask whether a polar-only encoding of errors and updates based on this tuning function can account for the successful compensation of Cartesian perturbation in our data.

We estimate the generalization tuning function based on the data of Krakauer et al. (2000). There, subjects were adapted to rotational or gain perturbation for reaches in one direction, and subsequently tested (without visual feedback) for reaches in multiple directions. In their Figs. 5 and 6A, they report the degree of adaptive compensation as a function of reach direction relative to the trained direction. They plot observed corrections at generalization targets relative to the degree of compensation at the training target. We scaled their tuning function based on the proportional corrective response observed for repeated target presentations, which is equivalent to measuring adaptive responses at one of their training locations (and was on average, mean ± SD across subjects, 0.82 ± .08 for reach gain and 0.69 ± .12 for reach direction). We scaled the tuning functions of Krakauer et al. (2000) using these values, and used the resulting tuning functions to simulate performance in our six-target experiment. The simulation was given the same sequence of reaches and feedback was perturbed in the same manner as for our experimental subjects. The simulation maintained a separate mapping for each of the 6 targets. After each simulated reach, the size of the update applied to the mapping for each target was the product of (1) the size of the error observed in the just-completed reach and (2) the value of the scaled tuning function based on the angular distance between the location of the update and the target of the just-completed reach. Simulated reaches were noiseless; each simulated reach toward a target landed off-target by precisely the current mapping for that target.

Using the scaled tuning functions and procedure described above, we simulated our polar six-target experiment. This simulation predicted essentially complete adaptation to polar perturbation (response gains nearly 1). This prediction is clearly incorrect. We therefore re-scaled the tuning functions to reduce the

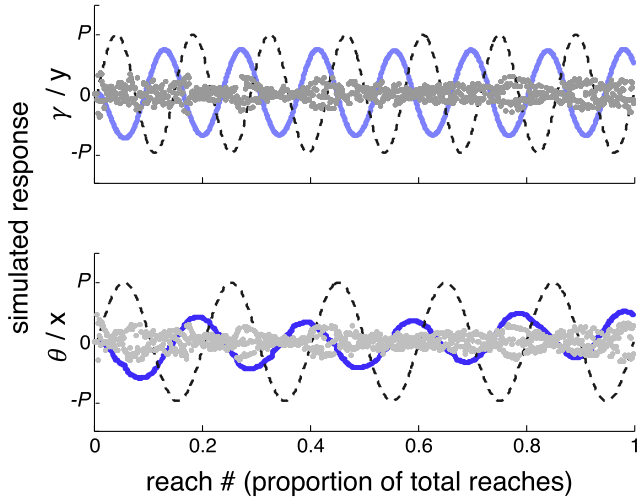

Fig. 6. Generalization tuning predictions. Prediction of average adaptive compensation in the six-target experiment for a model that stores independent, polar-coded compensations for each zone in the workspace with updates based on scaled versions of the generalization tuning functions reported by Krakauer et al. (2000). Predictions are plotted for gain (γ, light blue) and direction (θ, dark blue) compensations in the polar perturbation condition, and for the x (light grey) and y (dark grey) compensations for the Cartesian perturbation condition. Dashed sinusoids represent the perturbation. The same sinusoid is used to represent y- and gain-dimension perturbations (upper plot, data scaled appropriately to the corresponding amplitude of sinusoidal perturbation). Similarly, the same sinusoid is used to represent x- and θ-dimension perturbations (lower plot, data scaled appropriately to the corresponding amplitude of sinusoidal perturbation). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

simulated adaptive response, such that it would reproduce approximately the response amplitudes to polar perturbation measured in the six-target experiment. We used these reduced-amplitude tuning functions in a simulation of both the polar and the Cartesian perturbation condition. These (re-scaled) generalization functions produced the simulated data shown in Fig. 6. Fig. 6A shows simulated reach errors in the gain dimension of the polar session (blue) and the *y* dimension of the Cartesian session (grey). Fig. 6B shows simulated reach errors in the angle dimension of the polar session (blue) and the *x* dimension of the Cartesian session (grey). In both plots, simulated data are scaled to the size of the perturbation in the appropriate units. In the polar simulation, perturbation elicits a clear sinusoidal adaptive response. In the Cartesian simulation, perturbation elicits no noticeable sinusoidal adaptive response. The prediction of the polar-only hypothesis implied by the Cartesian simulation is inconsistent with the observed data (Table 2), which demonstrates an adaptive response whose amplitude gain is as large as that observed in the angle dimension of the polar session.

The generalization tuning function we used in the simulations was based on training at a single target, and the data of Krakauer et al. (2000) indicate that tuning is likely wider with training at multiple targets as is the case in our experiment (see also *Comparison to previous models: radial basis functions* for further exploration of this point). However, the observed pattern in the simulation results (sinusoidal modulation in simulated polar, but not Cartesian, sessions) was robust to changes in both the scaling and in the width of the underlying tuning function. Scaling the tuning functions simply changed the scale of the scatter in the Cartesian session and there was no width of the tuning function that was able to produce sinusoidal modulation in the simulated Cartesian session (and note that a flat tuning function is equivalent to Hypothesis I, and delta-function tuning is the same as Hypothesis III). We therefore conclude that our simulation results are inconsistent with the polar-only model utilizing separate adaptation at each target described by Hypothesis II.

2.4. Generalization experiment: polar-only hypothesis

The generalization experiment allowed us to further examine a polar-only hypothesis, Hypothesis III, in which separate compensations are computed independently at each target location using only reaches to that location to update the current estimated compensation. At each target, this creates a situation in which a polar update can compensate for Cartesian perturbation, and vice versa. As seen in Fig. 1, the difference between Cartesian and polar updating only becomes evident for generalization of compensation to reaches in multiple locations in the workspace. Hypothesis III posits that the presence of significant sinusoidal modulation of adaptive compensation in the Cartesian session of the six-target experiment was due to a polar-only mechanism operating independently at each target.

This spatially-selective version of the polar-only hypothesis has several weaknesses in conception. First, the strict version of this model would require an unrealistically large set of separate corrections to be stored—one per distinct location in the workspace. It would also be unable to adapt under conditions similar to those used for our practice reaches, where target locations are selected randomly on an annulus, because no target location is ever repeated—there is no opportunity to use the adaptive correction learned from the first reach to each target because there is never a second reach to the same target. The generalization experiment addresses this via random selection of target locations within the segmented annulus depicted in Fig. 2C. The generalization experiment also addresses a second prediction of relaxed versions of this hypothesis: adaptation will not generalize to any moderately sized region of the workspace where perturbed reach feedback has not been experienced.

Results from the feedback targets test the first prediction, that no adaptive corrections should be made when targets are not repeated. There is clear sinusoidal correction in all dimensions tested here, which is inconsistent with the (strict) polar-only hypothesis. Results from generalization targets test the second (relaxed) prediction that no adaptive corrections should be made in the 90°-wide section of the workspace where perturbed feedback was never experienced. Generalization-target results are also inconsistent with this prediction, showing strong sinusoidal corrective modulation in both polar and Cartesian sessions.

Finally, we note that sinusoidal adaptive responses of the two dimensions of Cartesian perturbation in the generalization experiment are (1) roughly similar to one another, and (2) are always slightly stronger than that observed in the polar rotation dimension for reaches to both feedback targets and to the generalization target. This is consistent with the results of our six-target experiment.

2.5. Comparison to previous models

As mentioned above (see *Parameter estimates*), the most common models of sensory-motor adaptation involve either a first-order difference mechanism or a Kalman filter. With the exception of Krakauer et al. (2000), however, these models typically do not deal with the issue of generalization. This leaves an ambiguity regarding how these models should be expected to perform under the current experimental conditions, and we therefore implemented the simplest extension of these models, assuming global generalization of updates in the coordinate system corresponding to the applied perturbation. Note that, to keep simulation noise low, simulated motor noise was set to match the lower end of measured endpoint noise in each condition (as measured during unperturbed practice reaches), corresponding to standard deviations of 9 mm for Cartesian dimensions, and 2.5 deg and .045 for angle and gain in the polar conditions.

First-order difference models. We first constructed a first-order difference adapter that corrects a single estimate of the required correction, regardless of the location of the target, but with separate corrections for each polar and Cartesian dimension. This mechanism stores a single correction for each dimension and, on each trial, adjusts the correction by subtracting a portion of the current error (the "correction ratio"). This updated correction is then applied to the subsequent movement. By fitting these adaptation mechanisms separately to each of the four dimensions we study, we are not modeling how non-orthogonal adaptation dimensions are combined. By doing this we ignore any impact of, for example, the effect of adaptation in *x* during the polar session on the effectiveness of the θ adaptation mechanism for compensating our θ perturbation. If we had implemented a model with separate adaptations in all four dimensions (*x*, *y*, γ and θ), summing their individual corrections, this would have injected noise into the process. During polar perturbation, the *x* and *y* adaptation mechanisms would have received random error signals (Fig. 1), and so their corrections would be random as well. However, such combination rules and the effects of combining adaptive corrections from different dimensions are not a component of most models of sensory-motor adaptation, and so were also not added to the modeling below.

To compare this model to our results, we forced the first-order mechanism to reproduce the adaptive gains observed in our experiments by adjusting the correction ratio for each dimension and under both polar and Cartesian perturbations separately. We then ask whether the corresponding phase lags predicted from the model match observed values. Although the 12 observed (Table 2)

and predicted lag values were well-correlated ($r = .82$; $p < .05$), a paired t-test of the same data shows a large and statistically significant discrepancy ($t(11) = 18.9$, $p < .05$); phase lags observed experimentally are shorter than those predicted by a first-order difference mechanism.

To account for savings during relearning of a previously experienced perturbation, an extension of the first-order difference mechanism has been proposed in which there are two states with multiple time constants (Smith, Ghazizadeh, & Shadmehr, 2006). This model assumes the two-state adaptive mechanism has the following form: compensation on each trial is the sum of two separate compensation mechanisms: $x^{t+1} = x_1^{t+1} + x_2^{t+1}$. Each component compensation is updated on each trial, $x_i^{t+1} = A_i x_i^t - B_i \Delta^t$, where parameter A_i determines the forgetting rate, B_i controls the learning rate, and Δ^t is the error observed on trial t. One mechanism adapts quickly and forgets quickly ($A = .59$, $B = .21$), and the other learns and forgets more slowly ($A = .992$, $B = .02$; Smith et al., 2006). Applying this two-state model to our stimuli, we find that predicted amplitudes (6 polar, 6 Cartesian, as in Table 2) are uncorrelated with observed values ($p > .05$), although predicted phases show a correlation with our data ($r(11) = .54$, $p < .05$). Despite a correlation with our phase lags, the two-state model predicts lags that are too short ($t(11) = 2.7$, $p < .05$) relative to our results.

Kalman filter models. The other common model of sensory-motor adaptation is a Kalman filter. We therefore next simulated the response of a basic Kalman filter (Kalman, 1960) designed to estimate the best sensory-motor mapping during sinusoidal perturbation. Again, this model was applied to each of the four dimensions separately. The Kalman filter does not have a free parameter to adjust the adaptive gain; simulations yield 12 response gains and 12 phase lags corresponding to our experimental conditions. The results of these simulations are uncorrelated with our experimental results ($p > .05$ both for response gains and phase lags).

A Kalman filter model, altered to include multiple timescales (as in *First-order difference models*), has also been proposed (Körding, Tenenbaum, & Shadmehr, 2007). In this model of sensory-motor adaptation, a bank of filter elements with a range from fast to slow timescales is assumed, where the predicted motor output is a weighted sum of the outputs of this bank of elements. The weightings assigned to each element in the bank are based on the solution to a credit-assignment problem, where fast-timescale elements are up-weighted when the sensory-motor mapping is perceived to be changing quickly, and slower-timescale elements are favored when the mapping is perceived to change slowly or not at all. Simulations of this model (using code supplied by the authors) applied separately to each dimension under the conditions of our experiments produce response gains and phase lags that are uncorrelated with our results ($p > .05$ in both cases).

We also examined a model based on partial Kalman updating (Burge et al., 2008) with a single timescale. Although Kalman gain, which is based in part on the estimated variance of the mapping (i.e., the currently estimated stability/instability of the mapping), will adjust from trial to trial in a Kalman filter, Burge et al. (2008) propose a model in which they fix the Kalman gain prior to introducing a step perturbation in the mapping (when the filter estimate of the mapping would normally change), and use this model in simulations of adaptation to a step-perturbed sensory-motor mapping. The Kalman gain was set in the same way as in the Burge et al. paper, by allowing the model to learn a gain during 100 practice reaches and fixing the gain at this value before introducing the sinewave perturbation. Simulations of this fixed-gain model in response to sinewave perturbation, applied separately to each dimension, also fail to correlate with the response gains and phase lags observed in our experiments ($p > .05$ in both cases).

Radial basis functions. Finally, we examined the radial basis function model proposed in several recent papers (Tanaka et al., 2009; Taylor, Hieber, & Ivry, 2013) to account for sensory-motor adaptation. Basis function representations provide a flexible encoding of coordinate frames for movement (Pouget & Sejnowski, 1997). For reach planning, they have been proposed with basis functions representing different movement directions (Brayanov, Press, & Smith, 2012; Tanaka et al., 2009; Taylor et al., 2013), velocities and forces (Donchin, Francis, & Shadmehr, 2003; Thoroughman & Shadmehr, 2000).

Tanaka et al. (2009) perform a series of simulations of Krakauer et al. (2000) that are most directly relevant to the current work. The Tanaka et al. (2009), Tanaka, Krakauer, and Sejnowski (2012) radial basis function model differs from adaptation models described previously in several ways, including its use of a distributed representation combining Cartesian-defined basis units and unit activation-strength that depends on the angular separation between target and the to-be-activated network unit. Tanaka et al. (2009) proposed this model to account for the findings of Krakauer et al. (2000), including the adaptation rates during reaches to single and multiple targets under a 30° rotation perturbation and the pattern of post-adaptation generalization.

This model consists of a set of radial basis function (RBF) units representing different potential target positions in reach space relative to the start position. In Tanaka et al.'s implementation and ours, a small number of units represent locations equally spaced in direction about the starting point (all locations at the same distance from the start). Each unit stores a Cartesian reach plan that is meant to reflect the current adaptation state for reaches to its corresponding location. In planning a reach to a target, each unit 'votes' for its own stored reach plan (x–y displacement), where the strength of each unit's vote (the unit's 'activation') is a function of the angular difference between the direction to the target and the unit's location/direction. The executed reach is thus the sum of all units' stored reach plans weighted by their respective activations. The function that determines activation strength is a single-peaked tuning function. After each reach, each unit updates its stored reach plan by subtracting a portion of the 2D Cartesian error observed in the previous reach. The amount of error correction is proportional to the unit's activation strength (the same strength that was used to plan the reach).

This model reflects the flexibility of basis-function representations. If the activation-strength function is infinitely narrow, then the model effectively computes independent adaptive corrections for each unit (i.e., for each target). If the function is very broad, this is equivalent to a single global adaptation. For intermediate widths of the function, it provides a method for implementing the generalization gradient that has been measured experimentally (Krakauer et al., 2000).

In our simulations we follow Tanaka et al. (2009), except that we use a von Mises tuning function (45° full-width at half-height, which nearly perfectly matches the Gaussian tuning curve used by Tanaka et al.) for angular separations instead of a Gaussian, since angle is a wrapped variable. We constructed an RBF network of 15 units (equally spaced around a circle centered on the start position), and selected an update strength parameter that reproduced the exponential fit of Krakauer et al. (2000) to their single-target, angle-perturbation condition (i.e., reaches made to a single target, endpoints perturbed 30°). We then used the network to simulate reaches made during our experiments. Note that while Tanaka and colleagues only use their RBF network to simulate the response to angle perturbation, the network produces a 2D response to perturbation, regardless of the coordinates of the perturbing function. We therefore first generated responses of this network to the polar perturbation conditions of Krakauer et al. (2000) and then also simulated the sinewave perturbations used

in both of our experiments (six-target and generalization), under both polar and Cartesian perturbation.

Krakauer et al. (2000) presented several results relevant to sensory-motor generalization to polar perturbation. Two are most relevant to the current study (and were simulated by Tanaka et al., 2009). The first result is based on varying the number of target locations in a center-out reach paradigm (Fig. 7A). If there were generalization from one target location to another, they predicted an increased learning rate when plotted as return visits to any single target in the array. This is because some learning relevant to that single target will have occurred at other targets in the array, and therefore correction would occur faster than if it were the only target. The model reproduces their finding that learning rates are nearly constant when plotted as the number of visits to a single target, although there was increased learning in the eight-target condition (Fig. 7B). This is not surprising, since only targets in the eight-target condition are more closely-spaced than the 45° tuning width. Their second result is based on a typical generalization experiment, wherein center-out reaches are made to a set of targets, and only a subset of those targets are associated with error feedback. For a single feedback target, they show that when subsequent generalization targets are closer than 45° there is some generalization, but that the generalization response drops to its baseline level (about 20% of the response to the feedback target) when generalization targets are further than 45° from the feedback target (Krakauer et al., 2000). This falloff of generalization response produces a tuning curve with the general shape of those produced by the RBF network, depicted in Fig. 7C.

We note several aspects of the response of the RBF model to step-perturbation before proceeding. First, the RBF model naturally produces a 2D corrective response to 2D error feedback. It will therefore also produce a corrective response when exposed to perturbations of reach distance (such as those used by Krakauer et al., 2000). We show its responses to the analogous perturbations of reach distance in the lower panels of the plot-pairs shown in Fig. 7. Notice also that the RBF model has no baseline activation (i.e., activation drops to zero for reach directions more than 45° away from a given RBF unit), and therefore cannot reproduce the tuning curves of Krakauer et al. (2000), who demonstrate significant baseline corrective responses, even for generalization targets positioned 180° from the feedback target. However, recalling our introductory discussion of the corrective response dictated by a Cartesian-encoded adaptive system (Fig. 1), this RBF network will produce the opposite of the desired corrective response for generalization targets 180° from the feedback target under reach-distance (as opposed to reach-angle) perturbation. This point was also made by Tanaka et al. (2009), who saw no benefit to adding a baseline activation to the activation tuning function used by their network for this reason.

We also simulated the effects of sinewave perturbation with this RBF network. There are several striking elements of the response of the network to sinusoidal perturbation, relative to our experimental results. First, the RBF network produces a consistently weaker response to magnitude perturbation than its response to angular perturbation (amplitude gains of .58 vs. .68 in the six-target simulation, and .55 vs. .66 for feedback targets in the generalization simulation). Note that this is the opposite pattern to our experimental result, and that of every previous experiment comparing adaptive responses to angle and magnitude perturbations, which always find stronger adaptation gain in the distance dimension. Also, phase lags produced by the RBF model are substantially longer than those observed experimentally in the feedback conditions of our experiments (RBF lags average almost 35 trials, as opposed to the approximately 13.5 trials observed here). Most significant, however, is the fact that the RBF network produces no reliable sinusoidal corrective response at the generalization target (simulation of the generalization experiment) to either dimension of the polar perturbation condition. In simulating our Cartesian-perturbation conditions, the network produces a response with adaptation gain of about .79 and .68 at the feedback targets (and phase lags of about 30 reaches in both cases). This is stronger than the response of the network to either polar dimension or the responses of our experimental subjects. This condition also yielded a reliable response at the generalization target, with response gains of .44 and .33 for the x and y dimensions. Finally, notice that amplitude gains are larger than those observed experimentally, in all dimensions of all experiments (that yielded a significantly nonzero response). Clearly, the pattern of responses produced by this network is inconsistent with our experimental results, showing the wrong pattern of adaptive gains, gains and phase lags of incorrect size, and no generalization in the polar perturbation condition.

It is interesting that Tanaka et al. (2009) chose a Cartesian-encoded RBF network to simulate the experimental results of Krakauer et al. (2000), since the latter group explicitly proposes that adaptation and generalization utilize a polar representation. The advantage of a polar-represented RBF network is that the response gains for angle and distance perturbations, as well as baseline activations for each, can be independently specified. We performed simulations with such a network (RBF units store a 2D distance-angle representation of reach plans), where the activation function for the angle dimension is again a von Mises tuning function (45° full-width at half-height, as above) defined over angular separations between each unit and the target, and the activation function for the distance dimension is a von Mises tuning function (60° full-width at half-height, because tuning is wider for distance than angle in Krakauer et al., 2000) also defined over angular separations between each unit and the target. Each activation function was raised (i.e., a von Mises function plus a constant) so that the minimum activation was above 0. Because this network allowed two useful degrees of freedom (baseline activation and learning rate) for fitting the earlier experimental results (Krakauer et al., 2000), we were able to exactly match the output of the network to their single-feedback-target condition with a single generalization target at 180°, both in terms of the timecourse of learning and the extent of generalization. After matching the output of the network to this condition in both polar dimensions, we tested the network in the other conditions of the Krakauer et al. (2000) experiments and our own.

In our previous RBF simulations (with Cartesian-encoded RBF units) we reproduced the first of the two Krakauer et al. (2000) results, showing essentially no generalization (Fig. 7B) as additional feedback targets are added (Fig. 7A) in a feedback-target-only experiment. However, any RBF network that produces no generalization in the first Krakauer experiment will have difficulty also reproducing the second result, which demonstrates graded generalization as additional feedback targets are added to a feedback- plus generalization-target experiment. Thus, while our earlier simulations reproduced the first but not the second Krakauer result, the current simulations show the opposite pattern, reproducing the second but not the first Krakauer result (Fig. 7D and E). In addition, this network's adaptive response to our sinusoidally varying polar perturbation was now much stronger than observed experimentally (amplitude gains of .92 and .88 for distance and angle perturbations in the six-target experiment, and .94 and .51 for the generalization target of the generalization experiment), and there was no reliable adaptive response at the generalization target in either Cartesian dimension of the Cartesian-perturbation condition.

In summary, none of these current models (Kalman, first-order difference, or RBF network) can predict the pattern of observed (sinewave) data, although our augmented version of the multiple-time-constant model of Smith et al. (2006) produces

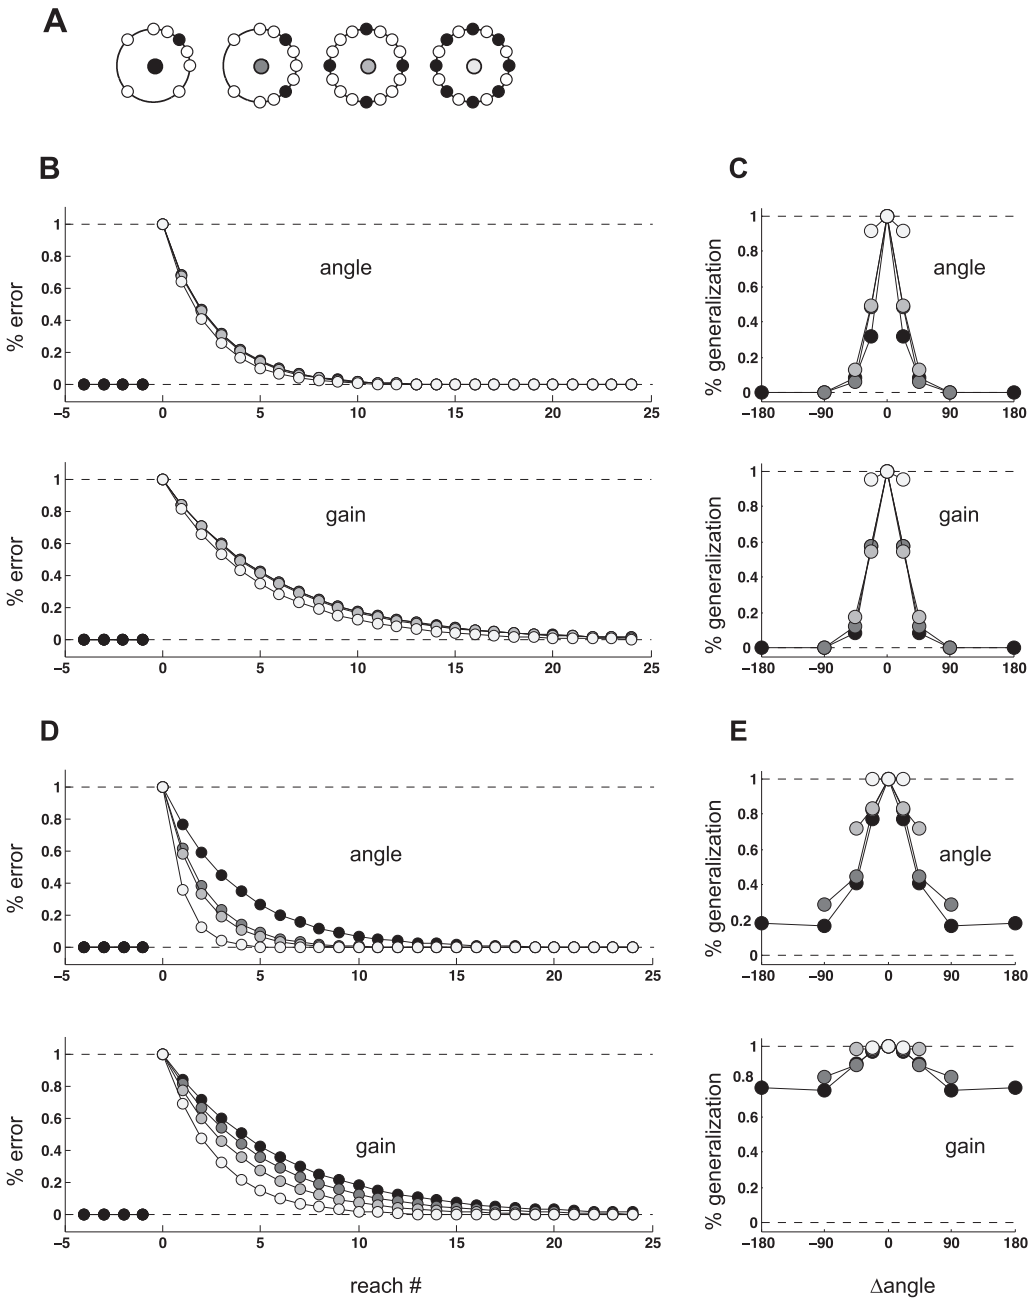

Fig. 7. Responses of radial basis function networks to various perturbations. (A) Configuration of central start position, feedback (filled) and no-feedback (unfilled) targets used by Krakauer et al. (2000) and Tanaka et al. (2009). The greyscale shading of the central start position is matched to the plot symbols in B–E. In all experiments, a step-function perturbation paradigm was employed, consisting of an initial series of unperturbed reaches, followed immediately and without warning by the introduction of a large perturbation (30° rotation or +.5 of reach gain). In one set of experiments, reaches were made to each of the filled patterns of targets in separate sessions. In another set of experiments, 'training' reaches were made to filled target locations and subsequent reaches to unfilled target locations were made in an effort to measure the gain of generalization. (B) Timecourse of the simulated adaptive response to the first experiment type (filled targets only). (C) Percent generalization produced by the RBF network following learning of perturbations at the filled targets. Note that the simulated result must be symmetrical about zero. (D–E) Same as B and C, but using the polar-encoded RBF network. Note that in B–E the upper plot of each panel depicts the response to stepped polar angle perturbation and the lower panel depicts the network response to stepped polar gain perturbation. In panels B and D the decay is always (more or less) faster with a greater number of targets, and in panels C and E generalization is always (more or less) more robust with a greater number of targets.

outputs correlated with the pattern of observed phase lags. The other models examined above either predict phase lags shorter than those observed experimentally (paired *t*-tests: Kalman filter: $t_{11} = 7.5$; two-state model: $t_{11} = 2.7$; multiple-timescale Kalman filter: $t_{11} = 10.9$; partial Kalman updater: $t_{11} = 10.2$; all $p < .05$), or much longer than those observed experimentally (first-order difference mechanism: $t_{11} = -18.9$, $p < .05$). Finally, the RBF model fails to produce a robust response at the generalization target in simulations of our generalization experiment.

3. Discussion

Although there has been a great deal of interest in the mechanisms by which the CNS maintains calibrated motor output in the face of noisy sensory data and imperfect effectors (Cressman & Henriques, 2011; Ghafouri, Archambault, Adamovich, & Feldman, 2002; Shadmehr et al., 2010; Smeets, van den Dobbelsteen, de Grave, van Beers, & Brenner, 2006), the coordinates in which corrective responses to motor errors are encoded

remains unresolved (Ghahramani et al., 1996; Lee, Pesaran, & Andersen, 2011; Pine et al., 1996; van den Dobbelsteen et al., 2001; Vindras & Viviani, 2002). We have previously shown that motor planning during unperturbed reaches is based on both a polar and a Cartesian coordinate encoding, even for the identical reaches (Hudson & Landy, 2012a), and it is clear that these plans can be updated when an imposed perturbation is apparent [when it is large and/or self-imposed, such as when using a tablet-based input device, e.g., Wu and Smith (2013), who used large perturbations of drawing/tracing movements]. Here, we examined whether motor plans in both coordinate systems are updated in response to perturbations that are too small to be noticed, consistent with the notion of an automatic sensory-motor adaptive response as opposed to a deliberate, cognitively mediated correction.

We hypothesized that the CNS might use both Cartesian and polar encodings of movement errors to update future reach plans because, for reaches in random directions, a polar-coded adaptive system would be blind to a uniform shift of sensory-motor space (i.e., errors produced by such a shift would be indistinguishable from zero-mean noise within this encoding); similarly, a Cartesian-coded adaptive system would be blind to a uniform scaling or rotation of sensory-motor space. A dual-coded adaptive mechanism would allow both types of correction to be made. In our experiments we find overwhelming evidence for dual-coded adaptation to motor errors. The predicted adaptive response to both Cartesian and polar-coded perturbations of endpoint feedback is visible even in the average data (Figs. 3A and 4A), and cannot be predicted by either generalization tuning (Fig. 6) or spatial selectivity (given the good response at the generalization target in Fig. 4) of a polar-only adaptive mechanism.

3.1. Results preclude a polar-only adaptive mechanism

Due to the fact that most previous motor-planning research has assumed a polar encoding, three versions of the polar-only hypothesis were considered: purely global compensation (see *Results*, Hypothesis I) and two versions of separate local compensations (Hypotheses II and III). The robust adaptation to Cartesian perturbation refutes Hypothesis I. Simulations of the Cartesian-session adaptive responses predicted for a local mechanism based on the generalization tuning reported by Krakauer et al. (2000) are inconsistent with Hypothesis II. Hypothesis III might allow for compensation during Cartesian-coded perturbation in the six-target experiment, but is inconsistent with the strong compensatory response to the generalization target in the generalization experiment. In its strictest form, Hypothesis III also cannot account for adaptation when targets are not repeated (such as reaches to feedback targets in the generalization experiment).

3.2. Results are inconsistent with radial basis function models

Several recent papers have proposed a radial basis function model of sensory-motor adaptation (Tanaka et al., 2009; Taylor et al., 2013). The radial basis function model is particularly relevant here in that it explicitly represents either the polar distance and angle dimensions or the Cartesian x and y dimensions while simultaneously making use of an activation function based on polar angle. Despite this, two versions of an RBF network failed to reproduce the pattern of our experimental results. In particular, these networks failed to reproduce the generalization phenomena observed here, displaying either no generalization in the polar experiment when network units represented the Cartesian x and y dimensions or no generalization in the Cartesian experiment when network units represented polar distance and angle dimensions.

3.3. Incremental perturbations and cognitive corrections

It is typical for adaptation experiments to use large perturbations. This has the apparent advantage of making a weak adaptive response detectable [see Hudson and Landy (2012b) for a comparison of response detectability under traditional and sinewave-adaptation paradigms]. However, large perturbations not only produce large compensations that are easy to detect experimentally, they are also obvious to experimental subjects. When applying such a large perturbation, it is not usually possible to determine what portion of the ensuing response is due to sensory-motor adaptation and what portion due to an active cognitive response to suprathreshold perturbation (Harris, 1974). Most work involving adaptation of reach extent and direction, for example, involves abrupt changes in movement gains of 30–50% and rotations of 30–90° or more (Imamizu, Uno, & Kawato, 1995; Pine et al., 1996; Salomonczyk, Cressman, & Henriques, 2011; Vindras & Viviani, 2002), where it is typical to find an adaptive response to the imposed step-perturbation that is initially extremely rapid but slows considerably after some portion of the perturbation has been corrected (Criscimagna-Hemminger et al., 2010; Körding et al., 2007; Krakauer et al., 2000).

The current experiments employ incremental perturbations, such that the difference in perturbation from one trial to the next is imperceptibly small. Indeed, even the maximum perturbation amplitude (7.5 mm) is quite small (smaller than the width of the average human index fingertip). By employing slowly varying sinusoidal perturbations, the maximum of the perturbation function occurs only briefly during each experimental session. This method allows us to obtain results that we believe are free from contamination based on cognitive correction effects (indeed one subject, who was a post-doc in our department, during debriefing strongly suggested that our experimental programming be 're-checked and fixed', because he was so convinced that no perturbation had occurred); this is similar to the rationale of previous work where perturbations were introduced in a random walk (Baddeley, Ingram, & Miall, 2003; Burge et al., 2008; Cheng & Sabes, 2007) or in a regular series of small steps (Cameron, Franks, Inglis, & Chua, 2010; Criscimagna-Hemminger et al., 2010; Hudson et al., 2005; Kagerer et al., 1997; Turnham et al., 2012), although in most cases the maximum perturbation was still large enough to be easily detected. We note that one such study (Turnham et al., 2012), comparing ±60° incremental (nonsinusoidal) and randomly-perturbed pre-training on the effect of subsequent 30°-stepped perturbations, showed longer reach initiation times for the incremental learning group—a result that was interpreted as a cognitive effect in their incremental relative to their random perturbation-learning group. We would argue that both groups were likely using cognitively based corrections in response to suprathreshold (30°) stepped perturbations, but the random perturbation-learning group was simply more familiar with (and therefore faster to respond to) large steps in perturbation magnitude than was the incremental perturbation-learning group.

We believe that, in addition to the cognitive contamination due to large perturbations themselves, there is likely a second cognitive effect based on the application of large perturbations during a standard generalization experiment. In the standard generalization experiment, suprathreshold step-perturbation is applied at a single target location for a large number of reaches. During these reaches, subjects are explicitly learning that there is a perturbation being applied at the location of the training target, but have no information regarding whether anything similar is occurring at other locations in the workspace. Thus, when reaches are finally made to other target locations, the adaptive response should not simply be interpreted as generalization of adaptation (an automatic process), but rather a combination of this process and the

generalization of cognitive correction that might result from uncertainty regarding whether the perturbation is expected to persist at other locations in the workspace (Bedford, 1989). Our six-target experiment avoids both contaminating effects by using small perturbations and treating all reach targets identically. Similarly, we attempt to eliminate these effects by employing small perturbations with feedback over three-quarters of the workspace in the generalization experiment. We cannot, however, rule out the notion that another contextual effect could have affected the result at the generalization target to an extent, since its differential treatment (no feedback was provided at that location only) should have eventually become apparent to subjects.

Finally, we note that in our experiments subjects made point-to-point reaches on a tabletop as instructed by a display (Fig. 2A), a method that has been used in much previous work (e.g., Baldeo & Henriques, 2013; Burge et al., 2008; Krakauer et al., 2000, 2004; Neva & Henriques, 2013). In such experiments, a coordinate transformation from visual to motor space is required. Studies of the kinds of coordinate transformations subjects can learn suggest that nonlinear mappings are quite difficult (Bedford, 1989, 1993), and indeed typical input devices implement an affine mapping from input to display space: $\mathbf{v} = \mathbf{x} + \mathbf{Tm}$, where $\mathbf{m}$ is a vector indicating the location of a target in motor space, $\mathbf{v}$ is the corresponding location in visual (display) space, and the transformation includes both a translation $\mathbf{x}$ and a linear transformation $\mathbf{T}$ that is typically limited to a scaling and rotation (i.e., no skew). We cannot say what combination of adaptive and deliberate compensatory mechanisms are involved in this transform, but previous work with large perturbations has shown no evidence that different mechanisms are involved (Messier & Kalaska, 1997) when reaches are instructed on a screen vs. when feedback and hand are co-located.

3.4. Evidence against cross-talk within a code

Our sinewave method allowed us to simultaneously adapt both axes of a perturbed coordinate system: either the orthogonal x and y axes of a Cartesian code, or the γ and θ axes of a polar code. Perturbing the two axes of each code simultaneously at different frequencies allowed us to separate the components of the response due to each dimension of the perturbation, and thus to detect any departure from a strict polar or Cartesian encoding of motor errors. Such a departure would manifest as multiple frequency components within one dimension of the adaptive response (see *Results*). Multiple frequency components should be visible as additional peaks in the log-posterior function (Figs. 3B and 4B), just as they would in a Fourier analysis of the response (see Bretthorst, 1988).

It is clear that secondary frequency components are not present in the response to perturbation (Figs. 3B and 4B). The fact that there is no evidence of secondary peaks in the posterior distributions allows us to rule out non-linear response components within each code. For example, one would interpret additional peaks in the frequency distributions at 2 and 12 cycle/session as evidence of quadratic distortion products due to cross-talk between the two dimensions of the perturbation. This might occur if, e.g., the response to x perturbation was the scaled product of the x and y perturbations, i.e., x-dimension responses were modulated by y-dimension perturbations. Instead, we find only a single peak in each of the posterior distributions, suggesting that adaptive responses follow strict polar and Cartesian encodings that feature independent contributions of the two spatial dimensions in Cartesian space, and similarly independent contributions of the γ and θ dimensions of polar space.

3.5. Computational structure of sensory-motor adaptation

Although these experiments were not designed to uncover the computational structure of adaptive change, we have compared our results to a first-order difference updater (Scheidt et al., 2001; Srimal, Diedrichsen, Ryklin, & Curtis, 2008) and several alternatives, including multi-rate updaters (Körding et al., 2007; Smith et al., 2006), radial-basis-function networks (Tanaka et al., 2009, 2012; Taylor et al., 2013), and a Kalman filter model (Korenberg & Ghahramani, 2002). No extant model predicts the overall pattern or the values of the parameters observed during sinewave adaptation.

In addition to the question of the computations involved in updating motor plans, there is also a question of how plans and updates based on the two coordinate systems are combined. Based on our current and previous results (Hudson & Landy, 2012a), it appears that both polar and Cartesian motor plans are always in operation and contributing to individual reaches. It would therefore seem highly likely that some form of weighted averaging is performed over plans within individual reaches. For example, we have previously suggested that both codes may be responsible for a full movement plan and that torques derived from both codes might be encoded as weighted sets of motor-unit activation strengths to the arm musculature (Hudson & Landy, 2012a). This is similar to the weighted average-based combination proposed by Wu and Smith (2013) based on experiments performed using large perturbations (and therefore likely to involve some combination of cognitive and adaptation-based corrections). The question of the computational structure of the mechanism or mechanisms responsible for sensory-motor adaptation remains unsettled, although a radial-basis-function framework that included units using Cartesian and polar-encoded movement plans would be a good place to start. For example, Pouget and Sejnowski (1997) propose a basis-function network model that allows targets in a mixed RBF representation to be read out in multiple reference frames. However, in this model, output units rely on only one of the reference frames: the coordinate frame tied to the behavior they control. Because a given behavior has access to only one frame of reference, this model, as it is now conceived, could not account for our results. However, we think a model of this type, properly modified to have readouts in multiple coordinate frames, with adaptive corrections learned and applied to each, may have promise for future modeling efforts designed to capture the adaptation results described here.

3.6. Context-specificity

We have shown that both polar- and Cartesian-coded adaptive responses occur during the identical reaching task. The only difference between polar and Cartesian adaptation sessions was the pattern over trials of imposed errors. Nothing in the pattern of reaches or other contextual and task-related parameters differed between the two sessions of either experiment. Thus, in addition to recent work demonstrating task- and context-specific adaptation using multiple reference frames (Bernier & Grafton, 2010; Ghafouri et al., 2002; Scheidt & Ghez, 2007; Thaler and Todd, 2009b), we show that use of multiple coordinate encodings is a general feature of sensory-motor adaptation and does not depend on the behavioral context. This is also consistent with our earlier results (Hudson and Landy, 2012a), which demonstrate dual encoding of reach plans within the identical set of (unperturbed) reaches, precluding any task-dependent segregation inherent in the use of one or the other coordinate system.

References

Baddeley, R. J., Ingram, H. A., & Miall, R. C. (2003). System identification applied to a visuomotor task: Near-optimal human performance in a noisy changing task. *Journal of Neuroscience, 23*, 3066–3075.

Baldeo, R., & Henriques, D. (2013). Dual adaptation to opposing visuomotor rotations with similar hand movement trajectories. *Experimental Brain Research, 227*, 231–241.

Baud-Bovy, G., & Viviani, P. (1998). Pointing to kinesthetic targets in space. *Journal of Neuroscience, 18*, 1528–1545.

Bedford, F. L. (1989). Constraints on learning new mappings between perceptual dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 232–248.

Bedford, F. L. (1993). Perceptual and cognitive spatial learning. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 517–530.

Bedford, F. L. (1999). Keeping perception accurate. *Trends in Cognitive Sciences, 3*, 4–11.

Bernier, P. M., & Grafton, S. T. (2010). Human posterior parietal cortex flexibly determines reference frames for reaching based on sensory context. *Neuron, 68*, 776–788.

Berniker, M., & Körding, K. P. (2011). Estimating the relevance of world disturbances to explain savings, interference and long-term motor adaptation effects. *PLoS Computational Biology, 7*, e1002210.

Bock, O., & Eckmiller, R. (1986). Goal-directed arm movements in absence of visual guidance: Evidence for amplitude rather than position control. *Experimental Brain Research, 62*, 451–458.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436.

Brayanov, J. B., Press, D. Z., & Smith, M. A. (2012). Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. *Journal of Neuroscience, 32*, 14951–14965.

Bretthorst, G. L. (1988). *Bayesian spectrum analysis and parameter estimation*. New York: Springer-Verlag.

Burge, J., Ernst, M. O., & Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of Vision, 8*(4). 20, 1–19.

Cameron, B. D., Franks, I. M., Inglis, J. T., & Chua, R. (2010). Implicit motor learning from target error during explicit reach control. *Experimental Brain Research, 206*, 99–104.

Cheng, S., & Sabes, P. N. (2007). Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *Journal of Neurophysiology, 97*, 3057–3069.

Cressman, E. K., & Henriques, D. Y. (2011). Motor adaptation and proprioceptive recalibration. *Progress in Brain Research, 191*, 91–99.

Criscimagna-Hemminger, S. E., Bastian, A. J., & Shadmehr, R. (2010). Size of error affects cerebellar contributions to motor learning. *Journal of Neurophysiology, 103*, 2275–2284.

de Graaf, J. B., van der Gon, J. J., & Sittig, A. C. (1996). Vector coding in slow goal-directed arm movements. *Perception & Psychophysics, 58*, 587–601.

Donchin, O., Francis, J. T., & Shadmehr, R. (2003). Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: Theory and experiments in human motor control. *Journal of Neuroscience, 23*, 9032–9045.

Georgopoulos, A. P., Kalaska, J. F., & Massey, J. T. (1981). Spatial trajectories and reaction times of aimed movements: Effects of practice, uncertainty, and change in target location. *Journal of Neurophysiology, 46*, 725–743.

Ghafouri, M., Archambault, P. S., Adamovich, S. V., & Feldman, A. G. (2002). Pointing movements may be produced in different frames of reference depending on the task demand. *Brain Research, 929*, 117–128.

Ghahramani, Z., Wolpert, D. M., & Jordan, M. I. (1996). Generalization to local remappings of the visuomotor coordinate transformation. *Journal of Neuroscience, 16*, 7085–7096.

Ghez, C., Favilla, M., Ghilardi, M.-F., Gordon, J., Bermejo, R., & Pullman, S. (1997). Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental Brain Research, 115*, 217–233.

Ghilardi, M., Ghez, C., Dhawan, V., Moeller, J., Mentis, M., Nakamura, T., ... Eidelberg, D. (2000). Patterns of regional brain activation associated with different forms of motor learning. *Brain Research, 871*, 127–145.

Good, I. J., & Toulmin, G. H. (1968). Coding theorems and weight of evidence. *IMA Journal of Applied Mathematics, 4*, 94–105.

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

Harris, C. S. (1974). Beware of the straight-ahead shift – A nonperceptual change in experiments on adaptation to displaced vision. *Perception, 3*, 461–476.

Henriques, D. Y., & Crawford, J. D. (2000). Direction-dependent distortions of retinocentric space in the visuomotor transformation for pointing. *Experimental Brain Research, 132*, 179–194.

Heuer, H., & Hegele, M. (2011). Generalization of implicit and explicit adjustments to visuomotor rotations across the workspace in younger and older adults. *Journal of Neurophysiology, 106*, 2078–2085.

Hudson, T. E., Lackner, J. R., & DiZio, P. (2005). Rapid adaptation of torso pointing movements to perturbations of the base of support. *Experimental Brain Research, 165*, 283–293.

Hudson, T. E., & Landy, M. S. (2012a). Motor learning reveals the existence of multiple codes for movement planning. *Journal of Neurophysiology, 108*, 2708–2716.

Hudson, T. E., & Landy, M. S. (2012b). Measuring adaptation with a sinusoidal perturbation function. *Journal of Neuroscience Methods, 208*, 48–58.

Imamizu, H., Uno, Y., & Kawato, M. (1995). Internal representations of the motor apparatus: Implications from generalization in visuomotor learning. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 1174–1198.

Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge, UK: Cambridge University Press.

Kagerer, F. A., Contreras-Vidal, J. L., & Stelmach, G. E. (1997). Adaptation to gradual as opposed with sudden visuo-motor distortions. *Experimental Brain Research, 115*, 557–561.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, 82*, 35–45.

Körding, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience, 10*, 779–786.

Korenberg, A. T., & Ghahramani, Z. (2002). A Bayesian view of motor adaptation. *Current Psychology of Cognition, 21*, 537–564.

Krakauer, J. W., Ghilardi, M.-F., Mentis, M., Barnes, A., Veystman, M., Eidelberg, D., et al. (2004). Differential cortical and subcortical activations in learning rotations and gains for reaching: A PET study. *Journal of Neurophysiology, 91*, 924–933.

Krakauer, J. W., Pine, Z. M., Ghilardi, M. F., & Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *Journal of Neuroscience, 20*, 8916–8924.

Lee, B., Pesaran, B., & Andersen, R. A. (2011). Area MSTd neurons encode visual stimuli in eye coordinates during fixation and pursuit. *Journal of Neurophysiology, 105*, 60–68.

Mazzoni, P., & Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *Journal of Neuroscience, 26*, 3642–3645.

McGuire, L. M., & Sabes, P. N. (2009). Sensory transformations and the use of multiple reference frames for reach planning. *Nature Neuroscience, 12*, 1056–1061.

McGuire, L. M., & Sabes, P. N. (2011). Heterogeneous representations in the superior parietal lobule are common across reaches to visual and proprioceptive targets. *Journal of Neuroscience, 31*, 6661–6673.

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

Neva, J. L., & Henriques, D. Y. (2013). Visuomotor adaptation and generalization with repeated and varied training. *Experimental Brain Research, 226*, 363–372.

Ojakangas, C. L., & Ebner, T. J. (1991). Scaling of the metrics of visually-guided arm movements during motor learning in primates. *Experimental Brain Research, 85*, 314–323.

Park, J., Schlag-Rey, M., & Schlag, J. (2006). Frames of reference for saccadic command tested by saccade collision in the supplementary eye field. *Journal of Neurophysiology, 95*, 159–170.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Pesaran, B., Nelson, M. J., & Andersen, R. A. (2006). Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron, 51*, 125–134.

Pine, Z. M., Krakauer, J. W., Gordon, J., & Ghez, C. (1996). Learning of scaling factors and reference axes for reaching movements. *Neuroreport, 7*, 2357–2361.

Pouget, A., & Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience, 9*, 222–237.

Ren, L., Khan, A. Z., Blohm, G., Henriques, D. Y., Sergio, L. E., & Crawford, J. D. (2006). Proprioceptive guidance of saccades in eye-hand coordination. *Journal of Neurophysiology, 96*, 1464–1477.

Salomonczyk, D., Cressman, E. K., & Henriques, D. Y. (2011). Proprioceptive recalibration following prolonged training and increasing distortions in visuomotor adaptation. *Neuropsychologia, 49*, 3053–3062.

Sanger, T. D. (2004). Failure of motor learning for large initial errors. *Neural Computation, 16*, 1873–1886.

Scheidt, R. A., Dingwell, J. B., & Mussa-Ivaldi, F. A. (2001). Learning to move amid uncertainty. *Journal of Neurophysiology, 86*, 971–985.

Scheidt, R. A., & Ghez, C. (2007). Separate adaptive mechanisms for controlling trajectory and final position in reaching. *Journal of Neurophysiology, 98*, 3600–3613.

Scheidt, R. A., & Stoeckmann, T. (2007). Reach adaptation and final position control amid environmental uncertainty after stroke. *Journal of Neurophysiology, 97*, 2824–2836.

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

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience, 33*, 89–108.

Smeets, J. B., van den Dobbelsteen, J. J., de Grave, D. D., van Beers, R. J., & Brenner, E. (2006). Sensory integration does not lead to sensory calibration. *Proceedings of the National Academy of Sciences U S A, 103*, 18781–18786.

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biology, 4*, e179.

Srimal, R., Diedrichsen, J., Ryklin, E. B., & Curtis, C. E. (2008). Obligatory adaptation of saccade gains. *Journal of Neurophysiology, 99*, 1554–1558.

Takahashi, C. D., Scheidt, R. A., & Reinkensmeyer, D. J. (2001). Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *Journal of Neurophysiology, 86*, 1047–1051.

Tanaka, H., Krakauer, J. W., & Sejnowski, T. J. (2012). Generalization and multirate models of motor adaptation. *Neural Computation, 24*, 939–966.

Tanaka, H., Sejnowski, T. J., & Krakauer, J. W. (2009). Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. *Journal of Neurophysiology, 102*, 2921–2932.

Taylor, J. A., Hieber, L. L., & Ivry, R. B. (2013). Feedback-dependent generalization. *Journal of Neurophysiology, 109*, 202–215.
Taylor, J. A., & Ivry, R. B. (2011). Flexible cognitive strategies during motor learning. *PLoS Computational Biology, 7*, e1001096.
Taylor, J. A., & Ivry, R. B. (2012). The role of strategies in motor learning. *Annals of the New York Academy of Sciences, 1251*, 1–12.
Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience, 34*, 3023–3032.
Thaler, L., & Todd, J. T. (2009a). The use of head/eye-centered, hand-centered and allocentric representations for visually guided hand movements and perceptual judgments. *Neuropsychologia, 47*, 1227–1244.
Thaler, L., & Todd, J. T. (2009b). The control parameters used by the CNS to guide the hand depend on the visuo-motor task: Evidence from visually guided pointing. *Neuroscience, 159*, 578–598.
Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature, 407*, 742–747.
Thoroughman, K. A., & Taylor, J. A. (2005). Rapid reshaping of human motor generalization. *Journal of Neuroscience, 25*, 8948–8953.
Turnham, E. J., Braun, D. A., & Wolpert, D. M. (2012). Facilitation of learning induced by both random and gradual visuomotor task variation. *Journal of Neurophysiology, 107*, 1111–1122.
van den Dobbelsteen, J. J., Brenner, E., & Smeets, J. B. (2001). Endpoints of arm movements to visual targets. *Experimental Brain Research, 138*, 279–287.
Vindras, P., Desmurget, M., & Viviani, P. (2005). Error parsing in visuomotor pointing reveals independent processing of amplitude and direction. *Journal of Neurophysiology, 94*, 1212–1224.
Vindras, P., & Viviani, P. (2002). Altering the visuomotor gain. Evidence that motor plans deal with vector quantities. *Experimental Brain Research, 147*, 280–295.
Wu, H. G., & Smith, M. A. (2013). The generalization of visuomotor learning to untrained movements and movement sequences based on movement vector and goal location remapping. *Journal of Neuroscience, 33*, 10772–10789.