

Motor Learning Is Optimally Tuned to the Properties of Motor Noise

Robert J. van Beers^{1,2,*}

¹Department of Physics of Man, Helmholtz Institute, Utrecht University, Padualaan 8, 3584 CH Utrecht, The Netherlands

²Research Institute MOVE, Faculty of Human Movement Sciences, VU University Amsterdam, Van der Boeorchorststraat 9, 1081 BT Amsterdam, The Netherlands

*Correspondence: r.vanbeers@fbw.vu.nl

DOI 10.1016/j.neuron.2009.06.025

SUMMARY

In motor learning, our brain uses movement errors to adjust planning of future movements. This process has traditionally been studied by examining how motor planning is adjusted in response to visuomotor or dynamic perturbations. Here, I show that the learning strategy can be better identified from the statistics of movements made in the absence of perturbations. The strategy identified this way differs from the learning mechanism assumed in mainstream models for motor learning. Crucial for this strategy is that motor noise arises partly centrally, in movement planning, and partly peripherally, in movement execution. Corrections are made by modification of central planning signals from the previous movement, which include the effects of planning but not execution noise. The size of the corrections is such that the movement variability is minimized. This physiologically plausible strategy is optimally tuned to the properties of motor noise, and likely underlies learning in many motor tasks.

INTRODUCTION

Suppose you play darts and aim for the bull's eye in the center of the board but hit the point 10 cm to the right of it. What should you do to do better in the next throw? Your previous throw had a rightward error of 10 cm, so you could try to make a leftward correction of the same size, and thus aim for the point 10 cm left of the bull's eye. Alternatively, you could argue that if you want to hit the bull's eye, you should always aim for it, and thus not make any correction. Or perhaps, the best strategy is to do something in between and aim for instance 5 cm to the left of the bull's eye. It is not obvious which strategy is the best.

This is a very simple example of a motor learning task in which errors in previous movements can be used to modify planning of future movements. Motor learning can take more complicated forms, such as learning a new motor skill like juggling (Shadmehr and Wise, 2005), but virtually all forms of motor learning are based on the movement errors made. The way in which errors are used to modify motor planning has traditionally been studied using visuomotor or dynamic perturbations that disturb the

motor performance (von Helmholtz, 1867; Welch, 1978; Shadmehr and Mussa-Ivaldi, 1994; Shadmehr and Wise, 2005). Subjects participating in such an experiment, however, have to perform two tasks at the same time (Berniker and Kording, 2008; Wei and Körding, 2009). On the one hand, they have to estimate the perturbations and predict the perturbation of the next movement, while on the other hand they have to perform the genuine motor learning task of correcting their motor planning. The results of such studies will therefore reflect both tasks.

The darts example suggests an alternative and cleaner way to identify how the motor system learns from movement errors. Given the natural movement variability and our motor system's propensity to correct for movement errors, there should be a certain relation between the errors in consecutive movements. If this is the case, an analysis of the time-series statistics of natural movement errors provides an excellent opportunity to identify the motor learning strategy in a way that is not confounded by the subject's need to estimate and predict perturbations. This approach is followed here for the well-practiced task of moving the right index fingertip to visual targets.

The darts example demonstrates that it can be difficult to determine how one should correct for an observed error. Errors arise when the brain generates inappropriate motor commands. A motor command is a set of time-varying signals that are sent to the muscles. Given the initial arm posture, a motor command corresponds to a unique spatial location where it will cause the finger to move to. I will denote this location as the *movement endpoint*, or simply *endpoint*. Motor commands can be inappropriate in two distinct ways. First, the central planning of a movement, i.e., the generation of motor commands in premotor and motor areas (Churchland et al., 2006a, 2006b), can be inaccurate as a result of a systematic error in the transformation from the intended target location into a motor command. In this case, the endpoint is systematically biased and does not coincide with the target location. This occurs frequently, as is exemplified by the regular misses of movements made with the unseen hand (Woodworth, 1899). This could be related to the complicated mechanics of the arm and the large number of muscles that have to be controlled, which make it difficult to generate accurate motor commands. Second, stochastic noise in central movement planning (Churchland et al., 2006a, 2006b) and in peripheral movement execution (Jones et al., 2002; van Beers et al., 2004; Faisal et al., 2008), i.e., in the relay of motor commands by motor neurons and in the conversion into mechanical forces in muscles, will cause the actual motor output

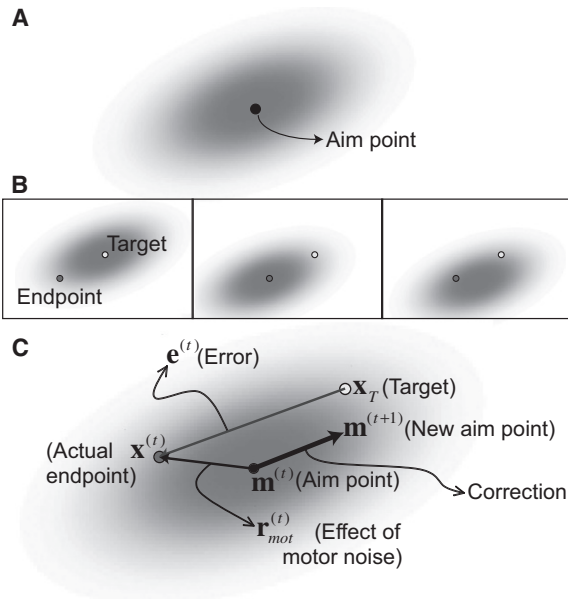


Figure 1. The Effect of Motor Noise and the Aim Point Correction (APC) Model

(A) The effect of motor noise is that, given an ideal (i.e., before motor noise is added) motor command, there is a probability density of movement endpoints (gray cloud). The aim point is the location where the movement would end if the motor command was not corrupted by motor noise.

(B) Three examples of different combinations of planning inaccuracy and effect of motor noise in single movements that lead to the same movement error. In the left plot, planning was perfect (the aim point coincides with the target), and the error is entirely due to motor noise. In the middle plot, the effect of motor noise was accidentally zero, and the error is entirely due to inaccurate planning. The right plot shows an example in which both planning and motor noise have nonzero contributions. In all three plots, the cloud represents the probability density resulting from motor noise, and is centered on the aim point.

(C) Vector diagram of the APC model. The goal is to reach target \mathbf{x}_T . Since movement planning is generally inaccurate, the aim point $\mathbf{m}^{(t)}$ in movement t will generally differ from the target location. As a result of motor noise, the actual endpoint $\mathbf{x}^{(t)}$ will differ from the aim point by a random amount $\mathbf{r}_{mot}^{(t)}$ (thin black arrow). The movement error $\mathbf{e}^{(t)}$ (gray arrow) is the difference between the endpoint and the target location. According to this model, a correction is made by shifting the aim point an amount $-B\mathbf{e}^{(t)}$ (bold black arrow). The correction is proportional to the previous error; learning rate B specifies the fraction of the error that is corrected for. Note that the direction of the correction is not perfect because the brain does not know the contribution of planning inaccuracy to the observed error.

to differ from the intended output. I will refer to *motor noise* as the total amount of noise added to the motor command during both movement planning and execution. The result of this noise is that the actual endpoint will differ from the endpoint if no noise had been added. Consequently, there is a probability density of movement endpoints given a certain ideal (i.e., before motor noise is added) motor command (Figure 1A). I will refer to the mean of this probability density as the *aim point*. When the probability density is Gaussian, this would be the endpoint if the motor command was not corrupted by motor noise.

Motor noise complicates error correction because it makes it impossible to know what the aim point of the previous movement was. There are an infinite number of combinations of planning

inaccuracy and effect of motor noise that could have produced a particular error (Burge et al., 2008). The error could be entirely due to motor noise, entirely to movement planning, or to a combination of both (Figure 1B). In the first case, no correction should be made; in the second case, a large correction would be required, whereas a smaller correction would be necessary in the third case. Our brain, however, does not know the actual contributions of these two sources to a particular error. It only knows the error. How, then, does our brain make corrections?

Theories have been developed for trial-by-trial learning in the presence of visuomotor or dynamic perturbations (Thoroughman and Shadmehr, 2000; Baddeley et al., 2003; Donchin et al., 2003; Diedrichsen et al., 2005; Cheng and Sabes, 2006, 2007; Smith et al., 2006; Burge et al., 2008). Some of these models (Baddeley et al., 2003; Diedrichsen et al., 2005; Cheng and Sabes, 2006, 2007; Burge et al., 2008) are particularly relevant here because they include the effects of motor noise. When the perturbation term is removed from these models, they make predictions for trial-by-trial corrections in the absence of perturbations but in the presence of motor noise. I will refer to the model constructed this way, based on the models of Baddeley et al. (2003), Diedrichsen et al. (2005), and Burge et al. (2008), as the *Aim Point Correction (APC) model*. I will assume highly reliable error feedback, so that the uncertainty herein can be neglected.

The APC model (Figure 1C) is described by the following equations:

$$\mathbf{x}^{(t)} = \mathbf{m}^{(t)} + \mathbf{r}_{mot}^{(t)} \quad (1A)$$

$$\mathbf{e}^{(t)} = \mathbf{x}^{(t)} - \mathbf{x}_T \quad (1B)$$

$$\mathbf{m}^{(t+1)} = \mathbf{m}^{(t)} - B\mathbf{e}^{(t)} \quad (1C)$$

The goal is to reach target location \mathbf{x}_T . Motor planning is generally inaccurate, so that for a particular movement t , the aim point $\mathbf{m}^{(t)}$ does not coincide with the target location. Equation 1A states that the movement endpoint $\mathbf{x}^{(t)}$ equals the sum of the aim point and the effect of motor noise, $\mathbf{r}_{mot}^{(t)}$ (thin black arrow in Figure 1C), which is a random vector drawn from a zero-mean Gaussian with covariance matrix Σ_{mot} . Movement error $\mathbf{e}^{(t)}$ is the difference between the endpoint and the target location (Equation 1B; gray arrow in Figure 1C). According to this model, a correction is made by shifting the aim point in the opposite direction as the error (Equation 1C; bold black arrow in Figure 1C). Learning rate B defines the fraction of the error that is corrected for. Note that all vectors in this paper represent locations that are expressed in extrinsic spatial coordinates.

RESULTS

Observed Learning

To examine whether the APC model describes how our brain corrects for movement errors, I analyzed the statistics of series of movement endpoints. In Experiment 1, eight subjects produced series of 30 arm movements from a fixed start position to a fixed visual target. Each subject produced 24 of such series, all with the same start position, to targets in different directions

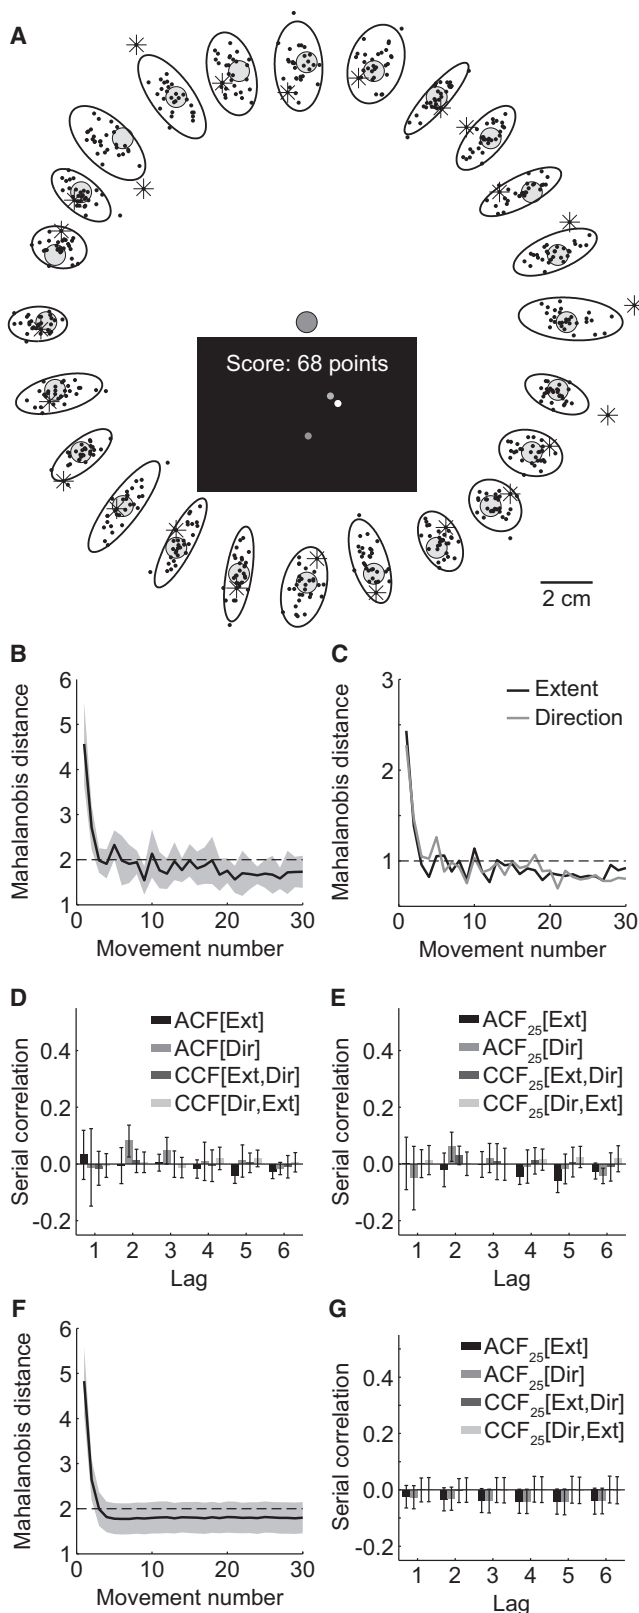


Figure 2. Results of Experiment 1

(A) The start position (dark gray disc in the center), the targets (light gray discs), all the endpoints (small dots) and their 95% confidence ellipses of a representative subject (SG). Asterisks mark the endpoint of the first movement to a target. Inset: the view that subjects had after completion of a movement, with the start position (dark gray), the target (light gray), the endpoint (white), and the score that was determined by the size of the error (discs not to scale). (B) Observed mean Mahalanobis distance of the two-dimensional endpoints as a function of the movement number in the series. The shaded area indicates the across-subjects standard deviation. The dashed line at 2 represents the expected value when all endpoints are drawn independently from an identical two-dimensional Gaussian distribution. (C) Observed mean Mahalanobis distance for the extent and direction components of movement endpoints as a function of movement number. The dashed line at 1 indicates the expected value when all numbers are drawn independently from an identical one-dimensional Gaussian distribution. (D) Observed mean ACFs and CCFs taking into account all 30 endpoints of each series. Error bars denote the across-subjects standard deviation. "Ext" and "Dir" refer to the extent and direction component, respectively. (E) Observed mean ACF₂₅'s and CCF₂₅'s that take into account only the last 25 endpoints of each series. (F) The Mahalanobis learning curve predicted by the PAPC model. The shaded area indicates the across-subjects standard deviation, as predicted by this model. (G) ACF₂₅'s and CCF₂₅'s as predicted by the PAPC model. Error bars denote the across-subjects standard deviation, as predicted by this model.

(Figure 2A). To allow subjects to position their finger quickly and accurately on the start position, a cursor was shown at the finger location when it was within 3 cm from the start position. The visual feedback went off when the finger began to move to the target. Immediately after movement completion, visual feedback of the movement endpoint was given. The endpoint was shown alongside the target, and, to motivate subjects, a score was awarded based on the error (see inset Figure 2A).

All endpoints of a representative subject are shown in Figure 2A. As expected, there is variability in the endpoints, and the mean endpoint in a series is generally close to the target. In many cases, the first movement to a target (indicated by asterisks) was quite inaccurate. No systematic pattern emerged in the direction of these large initial errors, both within and across subjects. They could overshoot the target, undershoot it, and/or their direction could be wrong. Uncertainty in the visual localization of the target is unlikely to be the major source of these errors, since their magnitude (often >2 cm) is much larger than could be expected from localization uncertainty (standard deviation about 6 mm, see Hansen and Skavenski, 1977; van Beers et al., 1998). Biases in visual localization are also unlikely to explain the errors because biases will not produce the wild variations of the errors that occur even between neighboring targets (see the initial errors for the targets in the 10:30 and 11 o'clock directions in Figure 2A). This suggests that the generation of motor commands is the most likely source of the initial errors. Apparently, movement planning can be rather inaccurate the first time one moves to a target.

Large errors were generally made only in the first movement to a target. This indicates that subjects corrected for initial errors in later movements. Learning curves were constructed to quantify the speed of learning. A plot of the error magnitude as a function of the movement number in the series could serve as a learning curve, but since endpoint distributions are anisotropic (Gordon

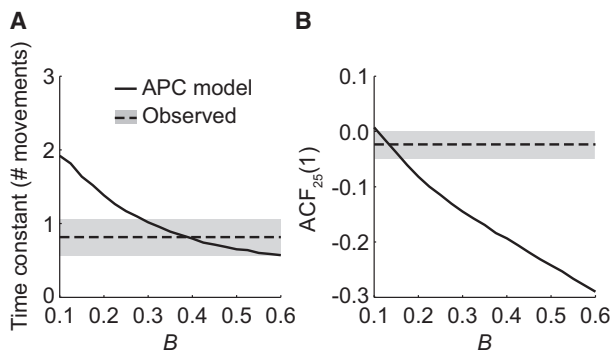


Figure 3. Predictions of the APC Model for Experiment 1

The time constant of the learning curve (A) and $ACF_{25}(1)$ averaged over the extent and direction components (B) predicted by the APC model as a function of B . The observed values are also shown, with shaded areas representing 95% confidence intervals of the mean. In both plots, the uncertainty in the model predictions is so small that their 95% confidence intervals are fully covered by the line.

et al., 1994; van Beers et al., 2004; Figure 2A), such a curve would mainly reflect learning of the movement extent and practically ignore learning of the movement direction. Instead, I plotted the Mahalanobis distance (see [Experimental Procedures](#)), which weights both components equally. It can be interpreted as the squared number of standard deviations that a given endpoint differs from the mean endpoint in its series, while taking the anisotropy into account. Since it is a normalized quantity, it can be averaged across series and subjects, even when their variance differs.

The mean Mahalanobis distance begins large and decreases quickly to a smaller value, which is retained until the end of the series (Figure 2B). The speed of learning was determined by fitting exponentials to each subject's mean Mahalanobis curve (see [Experimental Procedures](#)). This produced a time constant of 0.82 ± 0.25 movements (weighted average across subjects $\pm 95\%$ confidence interval), which suggests that on average 46% of the initial error had been corrected for in the next movement.

It has been proposed that the direction and extent of a movement are planned independently (Gordon et al., 1994; Krakauer et al., 2000). Are errors in these two components corrected differently? Figure 2C shows the mean learning curves for the individual components, where the extent component is defined as the component of the (two-dimensional) endpoints parallel to the vector from the start position to the mean endpoint, and the direction component is orthogonal to this. The mean time constants were 0.67 ± 0.21 and 0.93 ± 0.31 movements for the extent and direction components, respectively, and were not significantly different from each other (two-tailed paired t test, $p > 0.5$) or from the overall time constant (two-tailed paired t tests, both $p > 0.1$).

The time constants alone give insufficient information to test the APC model. The learning curves do not reveal what happens after large errors have been corrected. The APC model predicts that corrections will still be made, even when errors are small. This implies that there will be a certain relation between the endpoints of consecutive movements, or, more general,

between the endpoints of movements separated by a certain lag (number of movements) k . Such relations are quantified by serial correlations. Since the endpoints are two-dimensional vectors, the serial correlations consist of two autocorrelation functions $ACF(k)$, one for each component, and two cross-correlation functions $CCF(k)$ between the components (see [Experimental Procedures](#)). Here, the lag 1 correlations are the most informative ones. The $ACF(1)$ is positive when the endpoints of consecutive movements tend to be close together, whereas it is negative when consecutive endpoints tend to be far apart, on opposite sides of the mean endpoint (see Figure 7C for examples). A zero $ACF(1)$ implies that consecutive endpoints are independent of one another. The $CCF(1)$'s express such relations between the extent component of one endpoint and the direction component of the previous or next one.

All ACFs and CCFs in this experiment are close to zero (Figure 2D). Since the estimation of autocorrelations from short time series is fundamentally biased (Marriott and Pope, 1954; Kendall, 1954), it is impossible to test whether the ACFs differ significantly from zero. It is nevertheless clear that the endpoints of consecutive movements are not strongly dependent on one another.

When the first movement to a target was inaccurate, the partial correction for its error will give a positive contribution to the $ACF(1)$. Error-corrective learning in the "steady state" in which errors are small can therefore be better identified from serial correlations estimated from only the last 25 endpoints of each series. These correlations, denoted by $ACF_{25}(k)$ and $CCF_{25}(k)$, are also close to zero (Figure 2E). The observed $ACF(1)$ and $ACF_{25}(1)$ were not significantly different for the extent and direction components ($p = 0.4$ and $p = 0.3$, respectively, two-tailed paired t tests).

Test of the APC Model

To test whether corrections were made in the way described by the APC model, Monte Carlo simulations were run in which corrections were made according to Equation 1 (see [Experimental Procedures](#)). For the first movement in a series, an additional random vector \mathbf{r}_0 , drawn from a zero-mean Gaussian with covariance matrix Σ_0 , was added to the right hand side of Equation 1A to reflect the difficulty of planning the first movement to a target. The model was evaluated by comparing the predicted and observed time constant and $ACF_{25}(1)$ (averaged over the two components). Learning rate B is the only free parameter of the model (see [Experimental Procedures](#)). The predicted time constant decreases with increasing B (Figure 3A). Around $B = 0.4$, the predicted time constant matches the observed one. However, the model predicts a negative $ACF_{25}(1)$ for this value of B (Figure 3B), which clearly disagrees with the data. Hence, there is no value of B that can reproduce both the time constant and the autocorrelations. This means that the APC model cannot explain how subjects made corrections.

Possible Explanation

A possible explanation for the failure of the APC model is that it applies only to large errors and that subjects make no corrections when errors are small. Such a strategy would lead to rapid correction of large initial errors, followed by a constant, rather

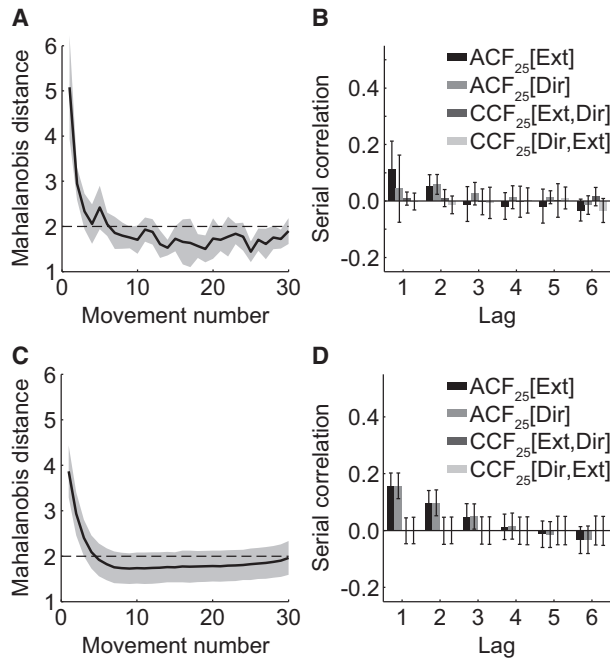


Figure 4. Results of Experiment 2

(A) Observed mean Mahalanobis distance as a function of movement number, plotted in the same format as Figure 2B.
 (B) Observed mean ACF₂₅'s and CCF₂₅'s plotted in the same format as Figure 2E.
 (C) The Mahalanobis learning curve predicted by the PAPC model, plotted in the same format as Figure 2F.
 (D) ACF₂₅'s and CCF₂₅'s as predicted by the PAPC model, plotted in the same format as Figure 2G.

accurate performance. Since according to this strategy no corrections are made in the last 25 movements, the endpoints of these movements would be independent of one another, implying zero autocorrelation. This is exactly what was found. Alternatively, corrections could be made throughout the entire series, but in a different way than is described by the APC model (see below). To distinguish between these possibilities, Experiment 2 was conducted which was identical to Experiment 1 apart from one factor. Now, subjects did not see their actual movement endpoint, but the point midway between the actual endpoint and the target. Subjects were unaware of this manipulation and believed they saw the actual endpoints. If subjects do not make corrections for small errors, their steady-state behavior will be the same as in Experiment 1. The ACF₂₅(1) will therefore be the same in both experiments. In contrast, if subjects do make corrections, their corrections will be too small because they are based on error signals that convey only half the actual error. Making smaller corrections will therefore cause consecutive endpoints to be closer together than in Experiment 1, i.e., the ACF₂₅(1)'s will be larger.

Subjects corrected for large initial errors also in this experiment (Figure 4A). The ACF₂₅(1) (Figure 4B) was larger than in Experiment 1 for both components (extent, $p = 0.004$; direction, $p = 0.040$; one-tailed paired t tests). This demonstrates that subjects make corrections all the time, also when errors are

small. This is consistent with reports that subjects adapt to small incremental perturbations of error feedback when each increment falls within the natural movement variability (Kagerer et al., 1997; Ingram et al., 2000; Klassen et al., 2005; Magescas and Prablanc, 2006).

Planned Aim Point Correction (PAPC) Model

The results suggest that humans make corrections for every movement error, but they do that in a different way than is described by the APC model. I will now present a new model, the *planned aim point correction (PAPC) model*, that is physiologically plausible and that can reproduce the data. The key difference between the two models is the different way in which is dealt with noise that arises during central movement planning (*planning noise*) and during peripheral movement execution (*execution noise*).

According to the APC model, a correction is made by shifting the aim point. A neural implementation of this model would thus require the brain to represent aim points. It is, however, unlikely that aim points are represented. Recall that the aim point is the position where the finger would land had the motor command not been corrupted by noise. However, motor command generation is a stochastic process that produces signals that vary from movement to movement (Churchland et al., 2006a, 2006b). This variability is captured by the term *planning noise*. Planning noise should therefore not be understood as noise that is added to an existing noise-free motor command, but it is present immediately when the motor command is generated. The supposed noise-free motor command simply does not exist. As a result, the aim point is probably not represented.

What signal could the brain use as an alternative for the aim point? The best signals that it could use are actual central planning signals of the previous movement. These central planning signals could be either the actually planned motor command, i.e., the motor command that includes planning noise, or its corresponding endpoint, which I will refer to as the *planned aim point* $\mathbf{m}_p^{(t)}$. The brain can have access to the actually planned motor command if it stored a copy of that command. It is widely assumed that such efference copies of motor commands are made because they are useful for canceling movement-induced changes of sensory signals (von Holst and Mittelstaedt, 1950; Sperry, 1950) and for feed-forward control of movement (Miall and Wolpert, 1996). There is, however, no evidence that efference copies are stored and can be recalled later, which makes this possibility speculative. Alternatively, the brain could store the previous movement's planned aim point. To achieve this, a forward model (Miall and Wolpert, 1996) could use an efference copy of the actual (noisy) motor command to generate an estimate of the planned aim point.

It is not important here whether the brain stores efference copies or planned aim points because both give rise to the same model, the PAPC model. If the efference copy is stored, this copy can be modified, based on the previous error, to make the correction. A modified version of the efference copy then acts as the motor command of the next movement. The modification will give rise to new planning noise in the generation of this command. If planned aim points are stored, the planned aim point of the previous movement can be modified, based

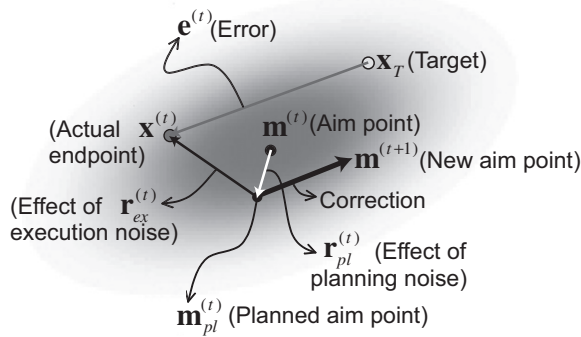


Figure 5. The Planned Aim Point Correction (PAPC) Model

This vector diagram is a modification of Figure 1C and differs in two respects. First, the effect of motor noise is decomposed into two random components, one resulting from movement planning ($r_{pl}^{(t)}$) and one resulting from movement execution ($r_{ex}^{(t)}$). Second, the correction is not made relative to the previous aim point but relative to the previous “planned aim point” $m_{pl}^{(t)} = m^{(t)} + r_{pl}^{(t)}$. Note that the direction of the correction is not perfect because the brain does not know the contribution of planning inaccuracy to the observed error.

on the previous error, to implement the correction. The modified aim point is then used to generate the motor command for the next movement.

Both implementations lead to the same model equations. The planned aim point $m_{pl}^{(t)}$ differs from the hypothetical aim point by the effect of planning noise $r_{pl}^{(t)}$ (white arrow in Figure 5):

$$m_{pl}^{(t)} = m^{(t)} + r_{pl}^{(t)} \quad (2)$$

Here, $r_{pl}^{(t)}$ is a random vector drawn from a zero-mean Gaussian with covariance matrix $\Sigma_{pl} = w \Sigma_{mot}$, where w is the fraction of the total effect of motor noise that is added during planning.

The actual movement endpoint is then found by adding the effect of execution noise $r_{ex}^{(t)}$ (thin black arrow in Figure 5):

$$x^{(t)} = m^{(t)} + r_{pl}^{(t)} + r_{ex}^{(t)} \quad (3)$$

where $r_{ex}^{(t)}$ is a random vector drawn from a zero-mean Gaussian with covariance matrix $\Sigma_{ex} = (1 - w) \Sigma_{mot}$.

According to the PAPC model, corrections are made relative to the previous planned aim point $m_{pl}^{(t)}$ (bold black arrow in Figure 5):

$$m^{(t+1)} = m_{pl}^{(t)} - B e^{(t)} \quad (4)$$

The aim point $m^{(t)}$ can be eliminated from Equations 2, 3, and 4, which is consistent with the idea that aim points are not represented in the brain. The PAPC model is then formulated in terms of planned aim points and takes the form of a linear dynamical system:

$$\begin{aligned} m_{pl}^{(t+1)} &= m_{pl}^{(t)} - B e^{(t)} + r_{pl}^{(t+1)} \\ x^{(t)} &= m_{pl}^{(t)} + r_{ex}^{(t)} \end{aligned} \quad (5)$$

The PAPC model has two free parameters: learning rate B and fraction w that specifies the fraction of the total effect of motor noise that is added during planning. These parameters were estimated for each subject from the observed time constant and $ACF_{25}(1)$'s (see Experimental Procedures). The best estimates

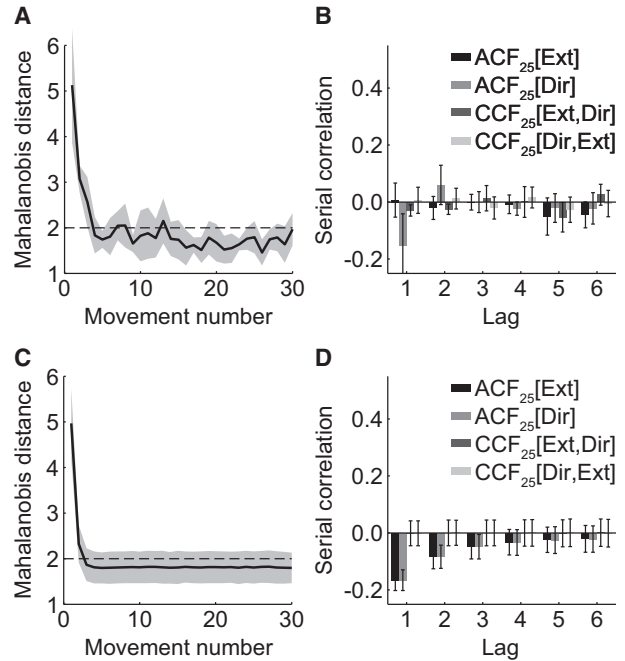


Figure 6. Results of Experiment 3

(A) Observed mean Mahalanobis distance as a function of movement number, plotted in the same format as Figure 2B. (B) Observed mean ACF_{25} 's and CCF_{25} 's plotted in the same format as Figure 2E. (C) The Mahalanobis learning curve predicted by the PAPC model, plotted in the same format as Figure 2F. (D) ACF_{25} 's and CCF_{25} 's as predicted by the PAPC model, plotted in the same format as Figure 2G.

averaged over subjects are (mean \pm SEM): $w = 0.21 \pm 0.03$ and $B = 0.38 \pm 0.04$. The Mahalanobis distance curve and serial correlations found in Experiment 1 are reproduced accurately for these mean values (Figures 2F and 2G; significance of differences between predictions and data: time constant, $p = 0.9$; $ACF_{25}(1)$ of extent and direction, $p = 0.5$ and $p = 0.6$, respectively, two-tailed t tests). The same parameter values also reproduce the results of Experiment 2 quite well (Figures 4C and 4D; significance of differences: time constant, $p = 0.4$; $ACF_{25}(1)$ of extent, $p = 0.3$; only the predicted $ACF_{25}(1)$ of direction is larger than observed, $p = 0.04$, two-tailed t tests).

Experiment 3 was performed to further test the generality of the model. This experiment was similar to Experiment 2, but now the shown errors were not 50% smaller but 50% larger than the actual errors. Again, subjects were not aware of this manipulation. The model predicts a roughly similar learning curve as in Experiments 1 and 2, and negative $ACF_{25}(1)$'s of about -0.17 (see Figures 6C and 6D). The negative autocorrelations arise from the fact that the error signals are larger than the actual errors, which will cause subjects to overcorrect and often overshoot the target (relative to the previous endpoint). Most of these predictions are confirmed by the data. The observed learning curve (Figure 6A) is quite similar to that of the other experiments, and the observed time constant is not significantly different from the predicted value ($p > 0.05$, two-tailed t test). The

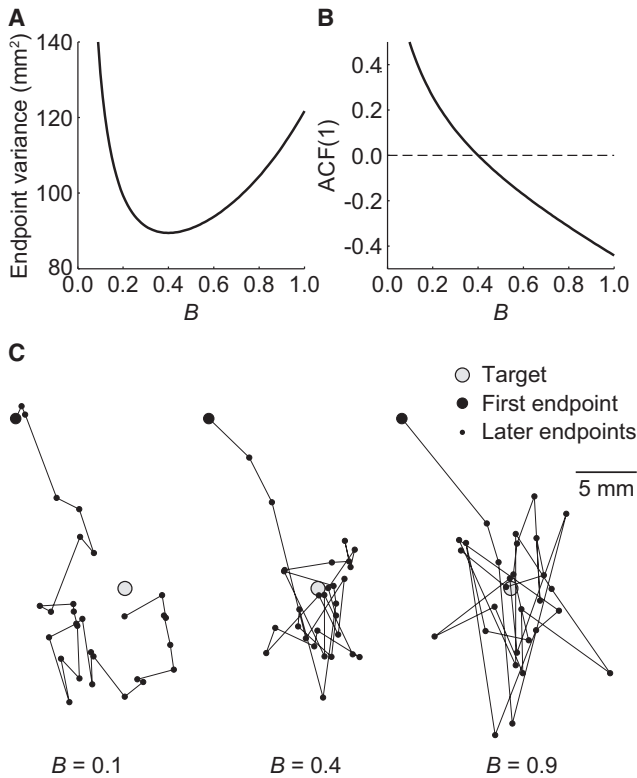


Figure 7. Influence of Learning Rate B in the PACC Model

(A) Theoretical endpoint variance (Equation 6A) as a function of B , for $w = 0.21$, and $\text{Tr}(\Sigma_{\text{mot}}) = 68 \text{ mm}^2$.

(B) Theoretical ACF(1) (Equation 6B) as a function of B , for $w = 0.21$.

(C) “Planned aim points” $\mathbf{m}_{pl}^{(t)}$ in simulated series for $B = 0.1$, $B = 0.4$, and $B = 0.9$ (in all cases, $w = 0.21$). Lines connect endpoints of consecutive movements. The same set of random numbers was used for each B .

ACF₂₅(1) of the direction component is negative (Figure 6B) and closely matches the predicted value ($p = 0.8$, two-tailed t test). However, the ACF₂₅(1) of extent is not negative but close to zero (Figure 6B) and differs significantly from the predicted value ($p = 0.001$, two-tailed t test). Possible explanations of this discrepancy are given in the Discussion.

Finally, it is worth paying attention to the estimated values of the model parameters. The value of w suggests that from the endpoint variability that results from motor noise, about 21% is due to planning and about 79% to execution. The estimated learning rate of 0.38 implies that 38% of each error is corrected for in the planning of the next movement. Why 38%? For this model, the theoretical endpoint variance $\text{Var}(\mathbf{x})$ and ACF(1) are (see Experimental Procedures):

$$\text{Var}(\mathbf{x}) = \frac{w + 2B(1 - w)}{B(2 - B)} \text{Tr}(\Sigma_{\text{mot}}) \quad (6A)$$

$$\text{ACF}(1) = 1 - B - \frac{B(2 - B)(1 - w)}{w + 2B(1 - w)} \quad (6B)$$

where Tr denotes the matrix trace. These functions are plotted as a function of B for $w = 0.21$ in Figures 7A and 7B. The variance

reaches a minimum for $B = 0.40$, which is not significantly different from the value estimated from the data ($p > 0.6$, two-tailed t test). The learning rate used by the brain apparently minimizes the endpoint variance.

The planned aim points $\mathbf{m}_{pl}^{(t)}$ in a simulated series of 30 movements are plotted for three values of B (0.1, 0.4, and 0.9) in Figure 7C. All plots were generated using the same set of random numbers. For small B , the planned aim point changes slowly. Since corrections (term $B\mathbf{e}^{(t)}$ in Equation 5) are small, the changes of the planned aim point are mainly determined by planning noise (term $\mathbf{r}_{pl}^{(t+1)}$). Without the error correction term, the planned aim point would describe a random walk. This explains why the changes of this point are often not directed toward the target and why the autocorrelation is positive and why the variance is large. For large B , the planned aim point changes rapidly. The changes are now mainly determined by the correction term. The correction is, however, generally not in the correct direction (see Figure 5). The direction would be correct if the correction was applied to the previous endpoint, but it is applied to the previous planned aim point. Large corrections are therefore often counterproductive. Consider for instance the case that planning was perfect (i.e., the planned aim point coincided with the target location) but that an overshoot was produced as a result of execution noise. Planning will then be adjusted so as to reduce the movement extent. The expected extent of the next movement will therefore correspond to an undershoot. This explains the negative autocorrelation and the relatively large variance for large B .

For intermediate B (0.4), the deleterious effects of small and large B cancel. The optimal value of B that minimizes the variance is the value for which the autocorrelation vanishes, because any nonzero autocorrelation causes additional variance. The observed small autocorrelations can hence be seen as an indication that the brain uses a strategy that minimizes the endpoint variance. Note that this does not mean that the planned aim point $\mathbf{m}_{pl}^{(t)}$ (or even the hypothetical aim point $\mathbf{m}^{(t)}$) remains constant (Figure 7C). It is impossible to fix these because the brain does not know the contributions of planning inaccuracy, planning noise and execution noise to the observed error. The best it can do is to minimize the variability in the planned aim points, and therefore in the endpoints, and that is achieved when B is optimal.

DISCUSSION

This study used a novel and simple method to study how our brain uses errors in previous movements to adjust planning of future movements. This method does not involve visuomotor or dynamic perturbations but takes advantage of the natural movement variability. In experiments with perturbations, subjects not only adjust their movement planning, but they also have to estimate and predict the perturbations. The actual correction strategy can therefore be better identified from tasks that do not involve perturbations. Time series analysis of movement endpoints in the absence of perturbations proved to be a powerful method.

To evaluate how the quality of the identification in this study compares to that of perturbation studies, one can compare the estimates of the learning rate. The mean estimate of 0.38 in

this study agrees well with the estimates in perturbation studies (Scheidt et al., 2001; Baddeley et al., 2003; Cheng and Sabes, 2007; Scheidt and Stoeckmann, 2007). However, the coefficient of variation of this estimate (the across-subject standard deviation divided by the mean) was only 28% in the present study, whereas it was equal to or greater than 50% in perturbation studies (Cheng and Sabes, 2007; Scheidt et al., 2001). This confirms that the perturbation-free method developed here is more efficient for identifying the learning strategy.

The learning strategy identified here is fundamentally different from the strategy assumed in the APC model (Baddeley et al., 2003; Diedrichsen et al., 2005; Burge et al., 2008) and in models that neglect motor noise altogether (Thoroughman and Shadmehr, 2000; Donchin et al., 2003; Smith et al., 2006). According to the APC model, we make corrections relative to the previous movement's aim point, i.e., the point where the finger would have landed in the absence of motor noise. The brain, however, cannot use this strategy because there is motor noise, as early as in the generation of motor commands in motor and premotor areas (Churchland et al., 2006a, 2006b). As a result, the aim point is unlikely to be represented in the brain. Instead, the brain probably represents the movement endpoint corresponding to the actually planned command, which includes the stochastic noise in movement planning. A useful strategy would therefore be to store this planned aim point, or the motor command, and modify that on the basis of the error when planning a new movement. This idea forms the basis of the PAPC model. Making corrections relative to the planned aim point rather than the aim point leads to an increased autocorrelation of movement endpoints (cf. Figures 3B and 7B), which is necessary to reproduce the data. This increase is a direct consequence of the assumption that corrections are made by modification of central planning signals of the previous movement. The effect of the previous movement's planning noise is then retained, while new noise is added in the planning of the next movement. The effects of planning noise therefore accumulate over movements, giving rise to an increased autocorrelation. As a result, the planned aim point is not constant but displays random changes even in the steady state when errors are small (see Figure 7C). The slow random drifts in the tuning curves of motor cortical neurons that have been observed in this steady state (Rokni et al., 2007) could be a neural correlate of these changes.

Before discussing the PAPC model in more detail, it is important to consider the recent idea that motor adaptation is the combined effect of two distinct processes (Smith et al., 2006). One process learns slowly but retains information well, whereas the other has a large learning rate but poor retention. Should the PAPC model be formulated as a two-state model, rather than the single-state model proposed here? In the two-state models, there are two processes, each having a learning rate and a capacity for retention, specified by retention rate A that defines how quickly the state decays back to its baseline level. Such a baseline level is meaningful in perturbation studies, as it specifies the state prior to the perturbations, but it is not a meaningful concept in the present experiments. Accordingly, the only sensible value for the retention factor is $A = 1$, for both states, implying no decay. However, the two-state model then reduces to a single-state model with $A = 1$ and a learning rate equal to the

sum of the two individual learning rates. Hence, two-state and single-state models are identical in the absence of perturbations. It is, however, possible that a generalization of the PAPC model to a situation with perturbations is better described by a two-state than a single-state model. Future investigations can explore this possibility.

Can alternative models be constructed that can reproduce the data? The difference between the PAPC and the APC model is that the effect of planning noise has been added, which raises the endpoint autocorrelation. A similar effect could be achieved by adding a different signal with positive autocorrelation. Theoretically, it is possible that the estimate of target location or that of movement error is corrupted by noise with a positive autocorrelation. So, removing the planning noise from the PAPC model, and adding positively autocorrelated noise to the estimate of target location or that of movement error could lead to a model that can reproduce the data. However, no studies have estimated the autocorrelations in these noise sources, so assuming them to be positive is somewhat arbitrary. In contrast, the positive autocorrelation arises naturally in the PAPC model as a result of using the previous motor signals to plan the next movement.

How does the PAPC model compare to the currently popular Bayesian models that have been used to account for motor learning and adaptation in response to perturbations (Korenberg and Ghahramani, 2002; Kording et al., 2007; Berniker and Kording, 2008; Burge et al., 2008; Wei and Körding, 2009)? Bayesian models are the most basic of generative models and describe learning at a rather abstract level. In contrast, the PAPC model describes learning at a lower level, in terms of standard constructs and terms from motor control such as motor commands, efference copies, planning noise, and execution noise. By combining these terms in a simple, physiologically plausible way, a linear dynamical system emerges that describes how the brain makes corrections to a motor plan. A linear dynamical system also forms the basis of the Kalman filter, which is a commonly used type of Bayesian model (Korenberg and Ghahramani, 2002; Baddeley et al., 2003; Kording et al., 2007; Burge et al., 2008). Hence, it is probably possible to construct Bayesian models that are described by the same equations as the PAPC model. Thus, although such a Bayesian model and the PAPC model derive from different principles, they are mathematically identical. The PAPC model can therefore be interpreted as a Bayesian model that is not described at the abstract level but at the lower, physiological level. As a result, this model indicates how Bayesian models for motor learning could be implemented in the brain, and as such it could spur new neurophysiological research aimed at identifying the neural implementation of Bayesian models. This applies also to other models that are described at the higher, abstract level, such as the model used by Cheng and Sabes (2006, 2007). The PAPC model is a special case of that more general model, which also includes two distinct noise sources. Whereas Cheng and Sabes (2006, 2007) referred to these as general output (or performance) noise and state (or learning) noise, they are made explicit in the PAPC model as planning and execution noise.

In the PAPC model, corrections are made relative to the previous movement's planned aim point. Two different neural

implementations are possible. The planned aim point could be represented via a stored efference copy of the motor command, or it could be stored directly as estimated by a forward model. The present results cannot distinguish between these possibilities. The first option is speculative as there is no evidence that efference copies are stored. It is however well possible that they are stored as this can be done relatively easily by storing only a few parameters. Motor commands can be encoded as combinations of a small number of muscle synergies, and only a few parameters are needed to represent even a time-varying combination of such synergies (d'Avella et al., 2006). There is in fact little doubt that the brain does store time-varying signals, as songbirds are known to store representations of birdsong (Troyer and Doupe, 2000). Interestingly, the idea of storing efference copies, and later recalling and modifying them as is proposed here, is comparable to the role that efference copies are thought to play in song learning in songbirds (Troyer and Doupe, 2000; Crapse and Sommer, 2008). The alternative way to implement the PAPC model involves a forward model that generates an estimate of the planned aim point. Neurons in the posterior parietal cortex have been found that could serve as a forward model (Mulliken et al., 2008). This implementation of the PAPC model is nevertheless also speculative as there is no evidence that estimates of planned aim points are stored and later modified. Future research is required to determine the actual neural implementation of the model.

This and other studies (Scheidt et al., 2001; Baddeley et al., 2003; Cheng and Sabes, 2007; Scheidt and Stockmann, 2007) suggest that our brain corrects about 38% of each error in the planning of the next movement. The present study is the first that explains this value: it is the learning rate for which the variance in movement endpoints is minimal. The near-zero autocorrelation of movement endpoints can be seen as a hallmark of a variance-minimizing strategy. The value of 38% results from the relative proportions of planning and execution noise. The learning rate can be different in other tasks as these relative proportions can vary between tasks. The question arises why our brain would minimize the endpoint variance. One could argue that motor learning should be more concerned about rapid correction of large errors, for instance by using a larger learning rate (Figure 7C). However, the learning mechanism is not only active after large errors, but it is always active, also in the steady state when errors are small. A larger learning rate would cause the average error in the steady state to be unnecessarily large. In other words, it would make our movements less precise than is physiologically possible, and this is likely to be more harmful to our functioning or to an animal's survival value than being a little slow with correcting large errors. A variance-minimizing strategy could also explain part of the two failures of the PAPC model reported here. The observed lag 1 autocorrelation of the direction component in Experiment 2 and of the extent component in Experiment 3 were closer to zero than predicted. Although subjects were not aware that the shown endpoints were not veridical, their motor systems may have detected that something was wrong (for instance by detecting discrepancies between expected and seen endpoints), and adjusted the learning rate such that the endpoint variance was minimized under the new circumstances.

In the PAPC model, learning rate B is a scalar. The model could easily be generalized to allow B to be a matrix. That would for instance make it possible to have different learning behavior for the extent and direction components. However, neither the time constant nor the autocorrelations differed between these components. The experimental results thus imply that B is a scalar. Another generalization could be to allow B to vary with the size of the error. If, for instance, B is increased for large errors and remains the same for small errors, large errors would be corrected faster, whereas the steady state behavior would remain the same. This is however not consistent with the data because it would lead to a different learning curve. Indeed, a study in which the learning rate was explicitly estimated as a function of the error size suggests that the learning rate is constant for the error sizes in Experiments 1 and 2, whereas for larger errors it is decreased rather than increased (Wei and Körding, 2009). This effect also offers another explanation why the observed autocorrelation of the extent component in Experiment 3 was about zero, and not negative as predicted by the PAPC model. All of these results together suggest that in natural movement behavior (with no perturbations), learning rate B is a scalar that does not depend on the size of an error. The learning strategy is therefore easy to implement. There is no need to know the planning and execution variances. Just correcting about 38% of each error will produce the observed behavior. The actual value of B could be learnt from the experience of repeated movements.

As a by product, this study produced estimates of the relative contributions of movement planning and movement execution to the total amount of motor noise. Although these contributions have not been estimated before, contributions of 21% and 79% of planning and execution, respectively, are consistent with the results of earlier studies on the sources of motor noise (Jones et al., 2002; van Beers et al., 2004; Churchland et al., 2006a, 2006b) and they agree also with the (rougher) estimates of the relative proportions of state and output noise by Cheng and Sabes (2007). The numbers might suggest that planning is a relatively unimportant source. However, planning may have substantial effects on the variability in movement velocity (Churchland et al., 2006a), and it has a sizeable effect on the endpoint autocorrelation (Figure 3B). Taking the effects of planning noise into account is therefore crucial when estimating the learning strategy from the trial-by-trial behavior in any motor learning task.

This study identified the strategy adopted by our brain to correct for movement errors of the unseen hand. A similar strategy could be used in the darts problem mentioned in the Introduction; aiming for the point 4 cm left of the bull's eye could be a good tactic. The importance of the strategy reaches however much further, as this strategy is likely to underlie error correction in many motor tasks. Corrections will be made by modification of planning signals from previous movements in any task that involves repeated movements, and that will most likely happen in accordance with the strategy identified here. The PAPC model could therefore underlie correction of movement planning for the seen hand, with the modification that the error signal will be related to some earlier part of the movement. The strategy could also model motor learning in the presence of

perturbations if an appropriate perturbation term is added. Other generalizations could include movements of other body parts and temporal and rhythmic movement tasks, for which the distinct roles of planning and execution noise are well established (Wing and Kristofferson, 1973). In all of these cases, the probabilistic nature of motor control is likely to explain how our brain makes and keeps movement planning accurate.

EXPERIMENTAL PROCEDURES

Experiments

Five male and three female subjects between 17 and 26 years old participated in Experiment 1, after providing informed consent and with approval from the Institutional Review Board. All subjects reported being right handed and were unaware of the purposes of the study. They made reaching movements with their right hand on a table (98 × 55 cm) while the position of the index fingertip was recorded by an Optotrak Certus system (Northern Digital) at 300 Hz. Subjects did not have direct vision of their arm and the table because they looked in a mirror placed midway between the table and a projection screen (above and parallel to the table). Small colored discs (4 mm radius) projected on the screen against a black background by an LCD projector (resolution 1280 × 720 pixels) defined the start position (purple) and the targets (yellow) and could also indicate the finger location (red or green). The start position was always the same, about 35 cm straight ahead of the waist. The finger position cursor was only shown at the beginning of a trial to allow subjects to place their finger quickly and accurately on the start location and to prevent drift of the perceived finger location throughout an experimental session (Smeets et al., 2006). The finger cursor appeared (red) when the finger was within 3 cm from the start location and turned green when it had been within 0.5 cm for 1 s. At the same moment the target appeared. Subjects were instructed to make a quick, uncorrected movement to the target. The finger cursor went off when the finger speed exceeded 2 cm/s. As a result, subjects received no informative visual feedback about the movement trajectory (<1 mm was shown). The movement endpoint was defined as the location where the finger speed first dropped below 2 cm/s. This location was shown immediately (red disc), along with a score (see inset of Figure 2A), which was determined by the distance from the target. One second later, the target and feedback went off and the next trial started.

An experiment consisted of 24 series of 30 movements each. The targets were at 10 cm from the start location in equally spaced directions. The same target was used for all movements in a series. The target of the first series was randomly chosen straight to the left or right. Each later target direction differed 105 degrees from the previous direction in the counter clockwise direction. Series were separated by breaks of 10 s. Prior to the experiment, each subject practiced the task for several minutes (with a different target than in the first series).

Experiments 2 and 3 were identical to Experiment 1 apart from the fact that not the actual endpoint was shown, but the endpoint that corresponded to an error that was 50% smaller (Experiment 2) or 50% larger (Experiment 3) than the actual error. Subjects were not informed about these manipulations, and postexperimental questioning confirmed that none of them had been aware that the feedback was not veridical. All eight subjects participated in Experiment 2; six of them participated in Experiment 3.

Data Analysis

The data analysis was completely based on two-dimensional movement endpoints. A small fraction of the movements (0.56%, 0.26%, and 0.35% in Experiments 1, 2, and 3, respectively) was discarded from the analysis because the recording had failed.

The Mahalanobis distance $D^{(t)}$ of movement number t in a series was calculated as

$$D^{(t)} = (\mathbf{x}^{(t)} - \bar{\mathbf{x}})^T \Sigma^{-1} (\mathbf{x}^{(t)} - \bar{\mathbf{x}}) \quad (7)$$

where $\mathbf{x}^{(t)}$ is the endpoint of movement t , $\bar{\mathbf{x}}$ and Σ are the mean and covariance matrix of all endpoints in the series, respectively, and T and -1 denote the

matrix transpose and inverse, respectively. Alternatively, one could replace $\mathbf{x}^{(t)} - \bar{\mathbf{x}}$ in Equation 7 by $\mathbf{x}^{(t)} - \mathbf{x}_r$, i.e., take the distances relative to the target position rather than the mean endpoint. This produced slightly larger distance measures but similar time constant estimates. Therefore, only the results using Equation 7 are reported.

Time constants of the learning curves, and their 95% confidence intervals, were estimated for individual subjects using nonlinear least-squares regression. The function $(a - b)\exp(-t/t_c) + b$ was fitted to the Mahalanobis distances averaged over the 24 series, where t is the movement number, a and b are constants and t_c is the time constant. Because the reliability (half the width of the confidence interval) of these estimates varied across subjects, a weighted average over subjects was calculated by weighting each subject's time constant by the inverse of the squared reliability.

The (sample) cross-correlation function $CCF(k)[i, j]$ between components i and j at lag k was calculated as

$$CCF(k)[i, j] = \frac{\sum_{t=1}^{n-k} x_i^{(t)} x_j^{(t+k)} - \frac{1}{n-k} \left(\sum_{t=1}^{n-k} x_i^{(t)} \right) \left(\sum_{t=1}^{n-k} x_j^{(t+k)} \right)}{\sqrt{\sum_{t=1}^{n-k} (x_i^{(t)})^2 - \frac{1}{n-k} \left(\sum_{t=1}^{n-k} x_i^{(t)} \right)^2} \sqrt{\sum_{t=1}^{n-k} (x_j^{(t+k)})^2 - \frac{1}{n-k} \left(\sum_{t=1}^{n-k} x_j^{(t+k)} \right)^2}} \quad (8)$$

where $x_i^{(t)}$ denotes component i of the endpoint of movement t , and n is the number of endpoints considered (30 or 25). The method developed by Marshall (1980) was used to deal with missing values. The (sample) autocorrelation function $ACF(k)[i]$ of component i at lag k was found as: $ACF(k)[i] = CCF(k)[i, i]$.

Model Simulations

Each Monte Carlo simulation consisted of 2000 sets of 24 simulated series of 30 movements each, corresponding to 2000 subjects performing a full experiment. Equation 5 was implemented to simulate a series in which errors were corrected according to the PAPC model. This involved drawing one random vector $\mathbf{r}_0 \sim N(\mathbf{0}, \Sigma_0)$ (i.e., from a zero-mean Gaussian with covariance matrix Σ_0) to reflect the difficulty of planning the first movement to a target, 30 random vectors $\mathbf{r}_{pl}^{(t)} \sim N(\mathbf{0}, \Sigma_{pl})$ to reflect the effect of planning noise, and 30 random vectors $\mathbf{r}_{ex}^{(t)} \sim N(\mathbf{0}, \Sigma_{ex})$ to reflect the effect of execution noise. This implies that, besides learning rate B , the elements of three covariance matrices were free parameters. However, the observations that the Mahalanobis distance, $ACF(1)$ and $ACF_{25}(1)$ were not different for the extent and direction components suggest that Σ_0 , Σ_{pl} , and Σ_{ex} were, up to scaling factors, equal to each other. This justifies the introduction of w to define the relative scaling of the planning and execution noise matrices as $\Sigma_{pl} = w \Sigma_{mot}$ and $\Sigma_{ex} = (1-w) \Sigma_{mot}$. Since the analysis involved normalized quantities only, the exact values of the elements of the covariance matrices are irrelevant. The scaling of Σ_0 was a free parameter, and was chosen as $\Sigma_0 = 4 \Sigma_{mot}$ because then the observed mean Mahalanobis distance of the first movement in Experiment 1 was reproduced. Simulations showed that the resulting time constant, ACF and ACF_{25} were virtually independent of this scaling factor in a wide neighborhood around the value used. As a result, B and w were effectively the only free parameters.

Simulations of the APC model were identical to those of the PAPC model, with w set to 0. Here, B was effectively the only free parameter.

Estimation of Model Parameters

Model simulations were performed for values of B and w between 0.1 and 0.8, and time constants and $ACF_{25}(1)$'s were estimated from these. These estimates were well approximated by third order polynomial regressions as a function of B and w . Parameters B and w were then estimated for each subject individually by finding the parameter values for which the sum of the squares of the normalized difference between observed and predicted (by the regressions) values of the time constant, the $ACF_{25}(1)$ of extent and the $ACF_{25}(1)$ of direction, was minimized. Normalization of the differences was achieved by dividing each difference by the reliability of the observed value. Note that the potentially more powerful method of maximum likelihood estimation, which can be implemented using the expectation-maximization algorithm (Cheng and Sabes, 2006), could not be used for parameter estimation as it produces biased estimates for the short time series used here.

Variance and Autocorrelation Equations

For each component (extent or direction), the model can be reformulated in terms of the errors $e^{(t)}$:

$$e^{(t+1)} = (1 - B)e^{(t)} + r_{pl}^{(t+1)} - r_{ex}^{(t)} + r_{ex}^{(t+1)} \tag{9}$$

The error autocovariance $\gamma(k)$ at lag k then is

$$\gamma(k) = E[e^{(t)}e^{(t+k)}] = (1 - B)\gamma(k - 1) + E[r_{pl}^{(t)}e^{(t-k)}] - E[r_{ex}^{(t-1)}e^{(t-k)}] + E[r_{ex}^{(t)}e^{(t-k)}] \tag{10}$$

where $E[f]$ denotes the expectation of f . For $k = 0$ and $k = 1$, the autocovariance is

$$\begin{aligned} \gamma(0) &= (1 - B)\gamma(1) + \sigma_{pl}^2 + (1 + B)\sigma_{ex}^2 \\ \gamma(1) &= (1 - B)\gamma(0) - \sigma_{ex}^2 \end{aligned} \tag{11}$$

where σ_{pl}^2 and σ_{ex}^2 denote the variance due to planning and execution, respectively. This system of equations has as solutions Equation 6B and

$$\text{Var}(e) = \gamma(0) = \frac{\sigma_{pl}^2 + 2B\sigma_{ex}^2}{B(2 - B)} = \frac{w + 2B(1 - w)}{B(2 - B)}\sigma_{mot}^2 \tag{12}$$

Equation 6A follows from Equation 12 under the assumption that corrections are made independently in the extent and direction components. This assumption is justified by the observations that the cross-correlations were zero and that the time constants, ACF(1)'s and ACF₂₅(1)'s did not differ between the two components.

ACKNOWLEDGMENTS

I thank Pieter Schiphorst and Hans Kolijn for technical assistance; the reviewers for valuable comments and suggestions; and Daniel Wolpert, Jörn Diedrichsen, Casper Erkelens, Raymond van Ee, Jeroen van Boxtel, Tomas Knapen, Dagmar Wismeijer, Jeroen Smeets, and Eli Brenner for inspiring discussions.

Accepted: June 25, 2009
Published: August 12, 2009

REFERENCES

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

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

Burge, J., Ernst, M.O., and Banks, M.S. (2008). The statistical determinants of adaptation rate in human reaching. *J. Vis.* 8, 1–19.

Cheng, S., and Sabes, P.N. (2006). Modeling sensorimotor learning with linear dynamical systems. *Neural Comput.* 18, 760–793.

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

Churchland, M.M., Afshar, A., and Shenoy, K.V. (2006a). A central source of movement variability. *Neuron* 52, 1085–1096.

Churchland, M.M., Yu, B.M., Ryu, S.I., Santhanam, G., and Shenoy, K.V. (2006b). Neural variability in premotor cortex provides a signature of motor preparation. *J. Neurosci.* 26, 3697–3712.

Crapse, T.B., and Sommer, M.A. (2008). Corollary discharge across the animal kingdom. *Nat. Rev. Neurosci.* 9, 587–600.

d'Avella, A., Portone, A., Fernandez, L., and Lacquaniti, F. (2006). Control of fast-reaching movements by muscle synergy combinations. *J. Neurosci.* 26, 7791–7810.

Diedrichsen, J., Hashambhoy, Y., Rane, T., and Shadmehr, R. (2005). Neural correlates of reach errors. *J. Neurosci.* 25, 9919–9931.

Donchin, O., Francis, J.T., and 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. *J. Neurosci.* 23, 9032–9045.

Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303.

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

Hansen, R.M., and Skavenski, A.A. (1977). Accuracy of eye position information for motor control. *Vision Res.* 17, 919–926.

Ingram, H.A., van Donkelaar, P., Cole, J., Vercher, J.L., Gauthier, G.M., and Miall, R.C. (2000). The role of proprioception and attention in a visuomotor adaptation task. *Exp. Brain Res.* 132, 114–126.

Jones, K.E., Hamilton, A.F.de C., and Wolpert, D.M. (2002). Sources of signal-dependent noise during isometric force production. *J. Neurophysiol.* 88, 1533–1544.

Kagerer, F.A., Contreras-Vidal, J.L., and Stelmach, G.E. (1997). Adaptation to gradual as compared with sudden visuo-motor distortions. *Exp. Brain Res.* 115, 557–561.

Kendall, M.G. (1954). Note on bias in the estimation of autocorrelation. *Biometrika* 41, 403–404.

Klassen, J., Tong, C., and Flanagan, J.R. (2005). Learning and recall of incremental kinematic and dynamic sensorimotor transformations. *Exp. Brain Res.* 164, 250–259.

Kording, K.P., Tenenbaum, J.B., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* 10, 779–786.

Korenberg, A.T., and Ghahramani, Z. (2002). A Bayesian view of motor adaptation. *Curr. Psychol. Cogn.* 21, 537–564.

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

Magescas, F., and Prablanc, C. (2006). Automatic drive of limb motor plasticity. *J. Cogn. Neurosci.* 18, 75–83.

Marriott, F.H.C., and Pope, J.A. (1954). Bias in the estimation of autocorrelations. *Biometrika* 41, 390–402.

Marshall, R.J. (1980). Autocorrelation estimation of time series with randomly missing observations. *Biometrika* 67, 567–570.

Miall, R.C., and Wolpert, D.M. (1996). Forward models for physiological motor control. *Neural Netw.* 9, 1265–1279.

Mulliken, G.H., Musallam, S., and Andersen, R.A. (2008). Forward estimation of movement state in posterior parietal cortex. *Proc. Natl. Acad. Sci. USA* 105, 8170–8177.

Rokni, U., Richardson, A.G., Bizzi, E., and Seung, H.S. (2007). Motor learning with unstable neural representations. *Neuron* 54, 653–666.

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

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

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

Shadmehr, R., and Wise, S.P. (2005). *The Computational Neurobiology of Reaching and Pointing* (Cambridge, MA: MIT Press).

Smeets, J.B.J., van den Dobbelaer, J.J., de Grave, D.D.J., van Beers, R.J., and Brenner, E. (2006). Sensory integration does not lead to sensory calibration. *Proc. Natl. Acad. Sci. USA* 103, 18781–18786.

Smith, M.A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* 4, e179. 10.1371/journal.pbio.0040179.

Sperry, R.W. (1950). Neural basis of the spontaneous optokinetic response produced by visual inversion. *J. Comp. Physiol. Psychol.* 43, 482–489.

- Thoroughman, K.A., and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature* *407*, 742–747.
- Troyer, T.W., and Doupe, A.J. (2000). An associational model of birdsong sensorimotor learning. I. Efference copy and the learning of song syllables. *J. Neurophysiol.* *84*, 1204–1223.
- van Beers, R.J., Sittig, A.C., and Denier van der Gon, J.J. (1998). The precision of proprioceptive position sense. *Exp. Brain Res.* *122*, 367–377.
- van Beers, R.J., Haggard, P., and Wolpert, D.M. (2004). The role of execution noise in movement variability. *J. Neurophysiol.* *91*, 1050–1063.
- von Helmholtz, H. (1867). *Handbuch der Physiologischen Optik* (Leipzig, Germany: Leopold Voss).
- von Holst, E., and Mittelstaedt, H. (1950). Das reafferenzprinzip. Wechselwirkungen zwischen zentralnervensystem und peripherie. *Naturwissenschaften* *37*, 464–476.
- Wei, K., and Körding, K. (2009). Relevance of error: what drives motor adaptation? *J. Neurophysiol.* *101*, 655–664.
- Welch, R.B. (1978). *Perceptual Modification. Adapting to Altered Sensory Environments* (New York: Academic Press).
- Wing, A.M., and Kristofferson, A.B. (1973). Response delays and the timing of discrete motor responses. *Percept. Psychophys.* *14*, 5–12.
- Woodworth, R.S. (1899). The accuracy of voluntary movements. *Psychol. Rev. Monogr.* *3 (Suppl.)*, 1–114.