VU Research Portal

Random walk of motor planning in task-irrelevant dimensions

van Beers, R.J.; Brenner, E.; Smeets, J.B.J.

published in

Journal of Neurophysiology
2013

DOI (link to publisher)

[10.1152/jn.00706.2012](https://doi.org/10.1152/jn.00706.2012)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

van Beers, R. J., Brenner, E., & Smeets, J. B. J. (2013). Random walk of motor planning in task-irrelevant dimensions. *Journal of Neurophysiology*, 109(4), 969-977. <https://doi.org/10.1152/jn.00706.2012>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Random walk of motor planning in task-irrelevant dimensions

Robert J. van Beers, Eli Brenner, and Jeroen B. J. Smeets

MOVE Research Institute Amsterdam, Faculty of Human Movement Sciences, VU University Amsterdam, Amsterdam, The Netherlands

Submitted 15 August 2012; accepted in final form 17 November 2012

van Beers RJ, Brenner E, Smeets JB. Random walk of motor planning in task-irrelevant dimensions. *J Neurophysiol* 109: 969–977, 2013. First published November 21, 2012; doi:10.1152/jn.00706.2012.—The movements that we make are variable. It is well established that at least a part of this variability is caused by noise in central motor planning. Here, we studied how the random effects of planning noise translate into changes in motor planning. Are the random effects independently added to a constant mean end point, or do they accumulate over movements? To distinguish between these possibilities, we examined repeated, discrete movements in various tasks in which the motor output could be decomposed into a task-relevant and a task-irrelevant component. We found in all tasks that the task-irrelevant component had a positive lag 1 autocorrelation, suggesting that the random effects of planning noise accumulate over movements. In contrast, the task-relevant component always had a lag 1 autocorrelation close to zero, which can be explained by effective trial-by-trial correction of motor planning on the basis of observed motor errors. Accumulation of the effects of planning noise is consistent with current insights into the stochastic nature of synaptic plasticity. It leads to motor exploration, which may subserve motor learning and performance optimization.

motor planning; motor learning; motor noise; variability; random walk

THE MOVEMENTS PRODUCED BY biological motor systems are variable. Even if we try to produce exactly the same movement a number of times, every movement will be different (“repetition without repetition”) (Bernstein 1967). Recent studies (Churchland et al. 2006; Sober et al. 2008) have shown that variations in motor behavior correlate with variations in pre-motor activity, which suggests that movement variability originates, at least partly, in central movement planning. Variations in central planning can result from error-driven motor learning, but the observation that planning variations also occur in the absence of errors (Churchland et al. 2006; Sober et al. 2008) suggests that there are also spontaneous variations. Spontaneous variations are likely to result from noise in the process of motor-command generation; we will therefore refer to these variations as effects of planning noise.

We will express the state of motor planning in movement t by the planned aim point $\mathbf{m}_{pl}^{(t)}$ (van Beers 2009), which is the location where the movement would end if it were driven by the centrally generated motor command without being corrupted by noise in the motor periphery. The effect of planning noise $\mathbf{r}_{pl}^{(t)}$ can affect the state of motor planning in different ways. Models for motor variability (Schmidt et al. 1979; van Beers 2007; van Beers et al. 2004; Wing and Kristofferson

1973) assume that the random effect of planning noise in one movement is independent of that in other movements and that it is added to some mean end point $\mathbf{m}$, which itself does not change:

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

When a series of movements to the same target is made, it is, according to this hypothesis, as if the planned aim points are drawn independently from the same distribution. Planned aim points in consecutive movements will therefore be statistically independent of one another.

According to models for trial-by-trial motor learning (Cheng and Sabes 2007; van Beers 2009), there is no fixed mean end point, but the random effect of planning noise is added to the planned aim point of the previous movement:

$$\mathbf{m}_{pl}^{(t)} = \mathbf{m}_{pl}^{(t-1)} + \mathbf{r}_{pl}^{(t)} \quad (2)$$

According to this hypothesis, the random effects of planning noise accumulate when a series of movements to the same target is made, giving rise to a random walk of the planned aim point. Planned aim points in consecutive movements will therefore be correlated positively.

The aim of this study was to determine which of these two hypotheses is correct. This is important, because the way in which our motor system deals with its own noise is fundamental for understanding the neural mechanisms of motor planning and motor learning. We distinguished between the hypotheses by analyzing the serial dependence of repeated, discrete movements. We took advantage of the fact that the output of many motor tasks can be decomposed into a task-relevant and a task-irrelevant component. For the task-relevant component, a possible positive serial dependence could be masked by error-driven planning corrections. In contrast, planning corrections are not required for the task-irrelevant component. This component thus allowed us to determine the serial dependence of repeated movements in the near absence of planning corrections.

MATERIALS AND METHODS

This study consists of three experiments. The first one was conducted specifically for this study. The other two were conducted earlier, for other purposes; we reanalyzed their data here.

Experiment 1. Subjects made reaching movements to line-shaped targets. This is a redundant task, as only the distance between the movement end point and the line matters; as long as the end point is not near or beyond the end of the line, the position parallel to the line is irrelevant (Diedrichsen et al. 2010; Sternad et al. 2010). The component of the end point orthogonal to the target line is therefore task relevant and the component parallel to the line task irrelevant.

Address for reprint requests and other correspondence: R. J. van Beers, MOVE Research Institute Amsterdam, Faculty of Human Movement Sciences, VU Univ. Amsterdam, Van der Boerhorststraat 9, 1081 BT Amsterdam, The Netherlands (e-mail: r.j.van.beers@vu.nl).

We tested six right-handed subjects (aged 24–30 years; five female). All were naive to the purpose of the experiment, had normal or corrected-to-normal vision, and provided written, informed consent. The experiment was conducted in accordance with the Declaration of Helsinki, and it was part of a program that was approved by the Ethics Committee of the Faculty of Human Movement Sciences.

Subjects were seated on a chair and used their right hand to move a stylus on a UD-1825-A drawing tablet (63.5×45.7 cm; Wacom, Vancouver, WA), which recorded the stylus tip position at 200 Hz with a resolution of 0.1 mm. Images were projected by a LCD projector (85 Hz; $1,024 \times 768$ pixels) on a projection screen above the tablet. Subjects looked down onto a mirror that was placed midway between the tablet and the projection screen. They could therefore not see their hand and the stylus, but they saw the projected images at a position that appeared to coincide with the tablet.

At the beginning of a trial, a yellow disk (5 mm diameter) appeared against a black background at a fixed location ~ 35 cm straight ahead of the trunk. To allow subjects to place the stylus quickly and accurately at this starting location and to prevent drift of the felt finger location (Smeets et al. 2006), a red disk (3 mm diameter) was shown at the current stylus location. Once the subject held the stylus still at the starting location, the stylus-location feedback went off, and the target appeared. We used three targets. Each was yellow, 6 cm long, and 4 mm thick and had its center 10 cm from the starting location, 25° to the left of the forward direction. The Forward target (Fig. 1A) was a line aligned with the target direction, the Diagonal target (Fig. 1C) was the same line as the Forward target but rotated 45° clockwise, and the Arc target (Fig. 1E) was part of a circle centered on the starting location.

The task was to move the stylus as accurately as possible to the target in a single movement in the absence of visual feedback. The movement end point was determined online as the first location since the start of the movement at which the stylus location was the same in two consecutive frames. From this moment, the movement end point was shown for 1 s as a red disk (3 mm diameter) alongside the target.

To motivate subjects, a score was displayed based on the distance between the end point and the target line. The score was inversely proportional to the distance from the line, which for this purpose, was assumed to have zero thickness, with a maximum of 100 points that was awarded if the error was <1 mm. After the 1-s interval, the target, the score, and the end point were extinguished, and the cursor displaying the current stylus location appeared to start the next trial.

A series consisted of 200 movements to the same target and took ~ 10 min. Each subject completed one series for each target in a single session, separated by breaks of ~ 30 s. The target order was counter-balanced. Subjects practiced the task for ~ 20 movements before the first series. We did not exclude any trials from the analysis.

Experiment 2. We examined large gaze shifts between visual targets achieved by combined eye-head movements. This task can be performed in different ways, as the contributions of the eye and the head may vary as long as the sum of both matches the target displacement. The task-relevant component was therefore the sum of the head and eye orientations after the movement, whereas the task-irrelevant component was their difference.

We analyzed a data set on large (80° and 100°) eye-head gaze shifts; details about this experiment can be found in the original paper (Tabak et al. 1996). In short, two visual targets were present, either 40° to the left and right of straight ahead, to evoke 80° gaze shifts (six subjects) or 50° to either side to evoke 100° gaze shifts (11 subjects). Head and eye orientations were recorded by scleral coils at 500 Hz. In some trials, a small amplitude ($<1^\circ$) oscillation (frequency: between 6 and 14 Hz, fixed/session) was added to the head orientation via a torque helmet. These trials could be included in the present analysis, because the oscillations hardly influenced the amplitude of the head movement (Tabak et al. 1996). For these trials, eye and head orientation was averaged over one oscillation period when determining the orientations after the gaze shift. End orientations were estimated at 400 ms and 600 ms after gaze-shift onset.

The analysis was done on a time series of 80 movements/subject for the 80° gaze shifts and on a series of 60 movements/subject for the

Fig. 1. Raw data of a representative subject in experiment 1. A: target (gray) and all movement end points (black) for the Forward target. B: task-relevant (open gray squares) and task-irrelevant (filled black circles) components of the end points shown in A plotted as a function of the trial number. For the task-relevant component, a value of 0 corresponds to a point exactly on the line. For the task-irrelevant component, it corresponds to the center of the target. C and D: same as A and B but for the Diagonal target. E and F: same as A and B but for the Arc target. ACF(1), lag 1 autocorrelation function.

100° gaze shifts. One-half of these was leftward, and the other one-half was rightward. A total of 1.6% of the trials was excluded from the analysis because the recording had failed.

Experiment 3. We examined measures of the trajectory of goal-directed movements. To reach a target location, the movement amplitude is important, but the peak velocity is not directly relevant. We analyzed two data sets, one on saccades (van Beers 2007) and one on reaching movements (van Beers 2009). Details about the methods can be found in the original papers. The saccades were from *experiments 1* (leftward saccades of nine amplitudes between 2° and 18°) and 2 (9° saccades in 24 equally spaced directions) from van Beers (2007). In each session, 75 saccades were made from the starting location to a fixed target. Five subjects participated, and 10% of the trials were discarded from the analysis based on criteria detailed in van Beers (2007). The reaching movements were from *experiment 1* in van Beers (2009). Eight subjects completed 24 series of 30 reaching movements to a small visual target. The targets were at 10 cm from the starting location in 24 equally spaced directions. Subjects did not see their hand during the movement, but they received visual feedback about the movement end point immediately after each movement. A total of 0.56% of the trials was excluded from the analysis because the recording had failed.

Autocorrelation. For all experiments and for both task-relevant and task-irrelevant aspects, we calculated the (sample) lag 1 autocorrelation function [ACF(1)] as:

$$\text{ACF}(1) = \frac{\sum_{t=6}^{n-1} x^{(t)} x^{(t+1)} - \frac{1}{n-6} \left(\sum_{t=6}^{n-1} x^{(t)} \right) \left(\sum_{t=6}^{n-1} x^{(t+1)} \right)}{\sum_{t=6}^{n-1} (x^{(t+1)})^2 - \frac{1}{n-6} \left(\sum_{t=6}^{n-1} x^{(t+1)} \right)^2} \quad (3)$$

where $x^{(t)}$ is the movement parameter considered (end point, peak velocity, etc.) of trial number t in the series, and n is the number of trials in the series. The first five movements of a series were not included, because correction for a large error in the first movement in the series could occur in these movements (van Beers 2009); including these movements would lead to an overestimate of the steady-state autocorrelation. The method developed by Marshall (1980) was used to deal with missing values. This method does not involve interpolation but estimates the autocorrelation using only the observed values. For *experiment 2*, autocorrelations were estimated for leftward and rightward gaze shifts separately, and these two estimates were then averaged (per subject). $\text{ACF}_{rel}(1)$ and $\text{ACF}_{irr}(1)$ refer to the ACF(1) of task-relevant and task-irrelevant parameters, respectively.

Random walk plus noise model. To allow us to interpret the magnitude of the observed autocorrelations in the task-irrelevant components in *experiment 1*, we modeled how movement end points are generated if the random effects of planning noise accumulate over movements (the second hypothesis in INTRODUCTION). If we assume that no planning corrections are made, then the temporal evolution of the planned aim points is described by Eq. 2. However, movement end points $\mathbf{x}^{(t)}$ differ from planned aim points as a result of noise added to the control signal in the motor periphery (“execution noise”) (Jones et al. 2002; van Beers et al. 2004). This can be modeled by adding a random effect of execution noise, $\mathbf{r}_{ex}^{(t)}$, to the planned aim point:

$$\mathbf{x}^{(t)} = \mathbf{m}_{pl}^{(t)} + \mathbf{r}_{ex}^{(t)} \quad (4)$$

The resulting model for the end points (Eqs. 2 and 4) is known as a random walk plus noise model (Chatfield 2004; Wei 2006). The estimate of the autocorrelation of a finite time series generated by a random walk plus noise model is smaller than that of the underlying random walk.

This model can easily be extended to include error-driven planning corrections that are made for the task-relevant component. We simply add a correction term that is proportional to the error in the previous movement to the right-hand side of Eq. 2 (van Beers 2009):

$$\mathbf{m}_{pl}^{(t)} = \mathbf{m}_{pl}^{(t-1)} - B\mathbf{e}^{(t-1)} + \mathbf{r}_{pl}^{(t)} \quad (5)$$

where $\mathbf{e}^{(t-1)}$ is the end-point error in the previous movement, and B is the learning rate that determines the size of the correction. The random walk plus noise model can therefore be viewed as a special case of this more-general model, with $B = 0$. Or to model the task-relevant and task-irrelevant components simultaneously, one could use a matrix learning rate, with a zero element for the task-irrelevant component and a nonzero element for the task-relevant component.

We conducted Monte Carlo simulations to determine the lag 1 autocorrelations predicted by the random walk plus noise model for the task-irrelevant component in *experiment 1*. We assumed that planning and execution noise were both white and independent of one another and that both were drawn from zero-mean Gaussian distributions. Following van Beers (2009), we assumed that the planning-noise variance was 21% of the total motor variance (the sum of planning and execution variance). To determine the predicted autocorrelation, we simulated a time series of 200 trials and estimated the lag 1 autocorrelation using Eq. 3. We repeated this 2,000 times (corresponding to 2,000 subjects) and averaged the obtained autocorrelations.

Since in *experiment 2* subjects moved multiple effectors (their head and eyes), the random walk plus noise model is not directly applicable to this experiment. However, by making some assumptions, we could use this model to obtain an estimate of the expected autocorrelation for the task-irrelevant component (the difference of the head and eye orientations after the movement) if the random effects of planning noise accumulate. Specifically, we assumed that 1) both eye and head movements have the same relative levels of planning and execution noise as arm movements; 2) if an error is made in a movement, this error is used to correct planning of the next movement; 3) the magnitude of these planning corrections is such that the lag 1 autocorrelation of the gaze is zero [because this agrees with the data (see below), and it minimizes the gaze variance (van Beers 2009)]; and 4) the eye and the head each account for one-half of the total correction. We used Monte Carlo simulations to determine the lag 1 autocorrelation of the task-irrelevant component predicted by this model. We simulated 2,000 time series with the same lengths as in the experiment and averaged the obtained autocorrelations.

Change-point model. We here consider an alternative explanation for the positive autocorrelations in *experiment 1*. Subjects could perform the task by switching between different positions along the target line and dwelling at each position for a random number of trials, with uncorrelated variability at each position. Such a process is known as a change-point process (Hinkley 1970). To examine whether the subjects' behavior could have resulted from a change-point process, we used the Bayesian algorithm of Adams and MacKay (2007) to estimate the change points (the switches to another location) in the task-irrelevant component of each series. This algorithm has to know the hazard rate (the rate at which change points occur), but we obviously did not know this rate in advance. We therefore first ran the algorithm of Wilson et al. (2010), which is an extension of the Adams and MacKay (2007) algorithm, to estimate the hazard rate. For both algorithms, we assumed that the generative distribution between change points was a Gaussian with unknown mean and known SD, which we derived from the data of the task-relevant component. Since for small targets, the ratio of end-point SD in movement extent and direction is ~ 2.0 (van Beers et al. 2004), we obtained this SD by multiplying the task-relevant SD by 2.0 and 0.5 for the Forward and Arc target, respectively. For the Diagonal target, the value was derived from the task-relevant SDs for the Forward and Arc target as $\sqrt{1/2\sigma_{rel,Forward}^2 + 1/2\sigma_{rel,Arc}^2}$, with $\sigma_{rel,i}^2$ as the variance of the task-relevant component for target i . We further assumed that the prior, for the mean of the Gaussian, was a Gaussian itself with a mean of zero (the middle of the line) and a SD of $10\times$ the SD of the generative Gaussian, but the results did not depend on the exact value of this parameter. For the Wilson et al. (2010) algorithm, we further assumed

that the prior, for the hazard rate, followed a beta distribution with parameters $\alpha = 2$, and $\beta = 50$, but the results were virtually independent of these values.

The Wilson et al. (2010) algorithm produced an estimate of the hazard rate following each of the 200 trials. We used the last of these estimates (typical values: 0.02–0.06) as input for the Adams and MacKay (2007) algorithm. This algorithm determined the probability of each possible value of the run length (the time since the last change point) for every trial. To estimate the actual change points from these probabilities, we started at the last movement, took the most probable run length, worked backward to the start of this run, then went one movement back, took the most-probable run length, and so on, until we reached the first trial. The mean of all responses within a run was used as an estimate of the mean of the generative Gaussian in that run.

We ran Monte Carlo simulations to determine the $ACF(1)$ predicted by the change-point model with the parameters derived in the manner described above. In each simulation, we drew 200 independent random numbers from a zero-mean Gaussian with the SD that we assumed for the generative distribution and added these to the means, as specified by the change-point model. We repeated this 1,000 times, calculated the $ACF(1)$ for each of these, and averaged these to obtain the predicted $ACF(1)$.

We did not consider the change-point model for *experiments 2 and 3*, because this does not seem a sensible model for these experiments. Since these experiments tested natural movements that we make routinely without considering the speed or the relative contributions of different effectors, we believe it is highly unlikely that subjects produced abrupt changes in these aspects. This is confirmed by visual inspection of the raw data (see, for example, Figs. 3A and 4A), in which no abrupt changes are evident.

Statistical analysis. Since the prediction of the random walk plus noise model is that the $ACF_{irr}(1)$ is positive, we used one-tailed *t*-tests to test the alternative that this $ACF(1)$ was smaller than or equal to zero. Note that this is a conservative test, because the sample autocorrelation (Eq. 3) underestimates the actual autocorrelation for a short time series (Kendall 1954; Marriott and Pope 1954). We used one-tailed paired *t*-tests to test whether the $ACF_{rel}(1)$ and $ACF_{irr}(1)$ were different. We used two-tailed paired *t*-tests to compare the observed $ACF_{irr}(1)$ in *experiment 1* with the value predicted by the change-point model.

RESULTS

Experiment 1. Subjects made 200 consecutive reaching movements to each of three line-shaped targets. The Forward target (Fig. 1A) was aligned with the target direction. It is not surprising that the end points exhibited more variability in the task-irrelevant than in the task-relevant component (see Fig. 1A for an example). Figure 1B shows these components as a function of the trial number. The task-relevant component varied near the target line with no obvious relation between consecutive trials. In contrast, the task-irrelevant component followed a different pattern, with most end points relatively close to the previous one and larger changes on a longer time scale. To quantify this behavior, we calculated the $ACF_{rel}(1)$ and $ACF_{irr}(1)$. For this example, the $ACF_{irr}(1)$ was 0.81, and the $ACF_{rel}(1)$ was 0.02. Across all subjects, the $ACF_{irr}(1)$ (0.54 ± 0.10 , mean $\pm$ SE) for the Forward target was significantly greater than zero ($P = 0.002$) and also significantly ($P = 0.001$) larger than the $ACF_{rel}(1)$ (0.03 ± 0.04 ; Fig. 2).

These results are inconsistent with the first hypothesis mentioned in INTRODUCTION, which predicts that the $ACF_{irr}(1)$ will be zero. The positive $ACF_{irr}(1)$ agrees qualitatively with the random walk plus noise model. To find out whether it also agrees quantitatively, we determined the $ACF_{irr}(1)$, predicted

by this model, using Monte Carlo simulations (see MATERIALS AND METHODS for details). These simulations showed that the predicted $ACF_{irr}(1)$ is ~ 0.80 . This closely matches the value for the example of Fig. 1B. However, the mean of all subjects (0.54 ; Fig. 2) is smaller. This difference could be due to the fact that some subjects drifted toward an end of the target line, where they started to make corrections to avoid missing the target. To test whether this is a plausible explanation, we determined for each subject how many movements ended < 1 cm from an end of the line. This number ranged between 2 and 142, and it was negatively correlated ($r = -0.83$; $P = 0.043$) with the $ACF_{irr}(1)$. This confirms that drifts toward line ends can explain why the mean-observed $ACF_{irr}(1)$ was smaller than predicted. It also demonstrates that the random walk plus noise model is a simplification of the true behavior of the task-irrelevant component, because this component is not irrelevant near the ends of the line. We will discuss other reasons why this model is a simplification in DISCUSSION. We nevertheless compared the data with the predictions of this model, because that allowed us to interpret the magnitude of the observed $ACF_{irr}(1)$.

Before we can conclude that the positive $ACF_{irr}(1)$ originates from accumulation of the random effects of planning noise, we will consider an alternative explanation. Subjects could have performed the task by (voluntarily or involuntarily) trying out different positions along the line, perhaps to try whether they performed better when they aimed for a different position along the line. Even if they pointed at each position with uncorrelated variability, the switches to other positions would have caused the overall $ACF_{irr}(1)$ to be positive. Simulations, in which the number and the size of switches were free parameters, showed that there exist values of these parameters that can reproduce both the observed variance and autocorrelation. To test whether such a change-point process could explain the observed, positive $ACF_{irr}(1)$, we used Bayesian change-point detection algorithms (Adams and MacKay 2007; Wilson et al. 2010) to estimate, for each subject, the change points in the time series of the task-irrelevant component. We then compared the autocorrelations predicted by these change-point models with the observed values. This analysis showed that for all subjects, the predicted autocorrelation was smaller than the observed one, and the difference was significant (mean difference $\pm$ SE: 0.08 ± 0.01 ; $P < 0.001$, paired *t*-test). This demonstrates that the task-irrelevant data cannot be the result of a pure change-point process. Note that this does not exclude

Fig. 2. Mean $ACF(1)$ in *experiment 1*. Mean $ACF(1)$ for the task-relevant (gray bars) and task-irrelevant (black bars) components for each of the 3 targets. Error bars: ± 1 SE (between subjects). * $P < 0.05$; ** $P < 0.01$.

the possibility that change points did occur. For instance, the large step following *trial 141* in Fig. 1*B* could be one. The fact that the change-point model predicts too small values for the $ACF_{irr}(1)$ means that a change-point process cannot explain the full magnitude of the $ACF_{irr}(1)$; the $ACF_{irr}(1)$ between change points must have been larger than the value of zero, assumed by a change-point process.

The $ACF_{rel}(1)$ was close to zero (Fig. 2). This is consistent with results on reaches to small, circular targets (van Beers 2009). In that study, it was shown that a zero $ACF(1)$ is a hallmark of optimal trial-by-trial corrections of motor planning on the basis of observed errors that minimize the end-point variance. The finding of a near-zero $ACF_{rel}(1)$ in the present experiment thus suggests that near-optimal planning corrections were made for the task-relevant component.

We tested two other targets. The results for the Diagonal target were very similar to those for the Forward target (see Figs. 1, *C* and *D*, and 2). The $ACF_{irr}(1)$ was again significantly ($P = 0.002$) greater than zero and also significantly ($P = 0.003$) larger than the $ACF_{rel}(1)$. For five of the six subjects, the observed $ACF_{irr}(1)$ was larger than that predicted by the change-point model, but the overall difference was not significant (mean difference $\pm$ SE: 0.00 ± 0.03 ; $P > 0.5$) for this target.

The results for the Arc target (Fig. 1*E*) were somewhat different than those for the other targets. The representative subject (Fig. 1*E*) had approximately the same variability in both components. She probably did make corrections for the task-irrelevant component, giving rise to little variability in any direction and a relatively small $ACF_{irr}(1)$ of 0.12 (Fig. 1*F*). Many other subjects displayed similar behavior. This suggests that many subjects, for some reason, always aimed for approximately the middle of the arc. As a result, the $ACF_{irr}(1)$ was not as large as for the other targets (Fig. 2). Nevertheless, it was significantly ($P = 0.043$) greater than zero and significantly ($P = 0.049$) larger than the $ACF_{rel}(1)$. The change-point model could be rejected for this target, but now, the observed $ACF_{irr}(1)$ was smaller than predicted (mean difference $\pm$ SE: -0.18 ± 0.06 ; $P = 0.035$).

In this experiment, the $ACF_{irr}(1)$ was always positive, which is consistent with accumulation of the effects of planning noise, and it was always larger than the $ACF_{rel}(1)$, which was close to zero. This was found for all targets. Since the targets had different orientations, this demonstrates that the dissociation between $ACF_{rel}(1)$ and $ACF_{irr}(1)$ is not related to anisotropies in the effects of planning and execution noise (Gordon et al. 1994; van Beers 2012; van Beers et al. 2004). To examine the generality of this dissociation, we examined data from other tasks. The motor output of these tasks can also be

decomposed into a task-relevant and a task-irrelevant component, but for different reasons than for the task of *experiment 1*.

Experiment 2. Whereas *experiment 1* focused on a redundant task, we will now consider redundant motor systems, which have more degrees of freedom than are required to perform the task at hand. For such systems, the between-movement variability in the available degrees of freedom can be decomposed into a task-relevant and a task-irrelevant part, which has led to the concept of the uncontrolled manifold (UCM) (Scholz and Schöner 1999). The UCM is composed of all variations of the degrees of freedom that do not affect task performance. Numerous studies [for reviews, see Latash et al. (2002, 2007)] have shown that the (normalized) variance is larger within the UCM than in the orthogonal subspace of variations that do affect task performance. However, little is known about the serial dependence within these two subspaces (Dingwell et al. 2010). We examined the serial dependence for large (80° and 100°) gaze shifts achieved by combined eye-head movements. Since the gaze direction relative to the trunk is the sum of the angles of the eye in the head and of the head on the trunk, only this sum was relevant for performing the task. How this sum was composed of contributions from the eye and the head was not directly relevant. The task-relevant component of the motor output was therefore equal to the sum of the head and eye orientations, whereas the task-irrelevant component (the UCM) equaled their difference.

Figure 3*A* shows a representative example of a series of 100° gaze shifts. A first striking point is that variability was much smaller in the gaze direction than in the head and eye orientations and in the UCM. The traces of the eye and head orientations were almost mirror images of each other, showing that variations in these orientations largely canceled each other. This is consistent with earlier results on the UCM. The $ACF(1)$ of gaze direction in Fig. 3*A* was quite small (0.18). The head and eye orientations and the UCM had larger $ACF(1)$ s (all between 0.31 and 0.34). Since these three $ACF(1)$ s were always very similar, we only compared the $ACF(1)$ of gaze with that of the UCM. Figure 3*B* shows the mean $ACF(1)$ s at two instants. The first was 400 ms after movement onset. At this time, the primary saccade had finished, but the head could still be moving, and a secondary saccade could still be made. The second instant was 600 ms after movement onset. At that time, all eye and head movements had finished. At both instants, the $ACF(1)$ of gaze was close to zero, whereas that of the UCM was positive (0.13 at 400 ms; 0.23 at 600 ms). At 400 ms, the $ACF(1)$ of the UCM just failed to be significantly larger than zero ($P = 0.059$), but it was significantly larger than the $ACF(1)$ of gaze ($P = 0.022$). At 600 ms, the $ACF(1)$ of the

Fig. 3. Results of *experiment 2*. *A*: gaze direction (head plus eye orientation), head and eye orientation, and uncontrolled manifold (UCM; head minus eye orientation), 600 ms after movement onset for a representative subject that made a series of 100° gaze shifts. Orientations for consecutive fixations of the same target are connected by lines. The shown $ACF(1)$ values are means for the left and right target. *B*: mean $ACF(1)$ for gaze direction (gray bars) and UCM direction (black bars), determined 400 ms and 600 ms after movement onset. Error bars: ± 1 SE (between subjects). * $P < 0.05$; ** $P < 0.01$.

UCM was significantly larger than zero ($P = 0.0065$) and also significantly larger than the $ACF(1)$ of gaze ($P = 0.028$).

The values of $ACF_{irr}(1)$ were smaller in this experiment than in *experiment 1*. To understand why, we estimated the $ACF(1)$ that could be expected from a straightforward extension of the random walk plus noise model (see MATERIALS AND METHODS). Simulations revealed that this model predicts an $ACF_{irr}(1)$ of ~ 0.4 . The reason it is smaller than in *experiment 1* (where it was 0.8) is that the time series here was shorter. The observed $ACF_{irr}(1)$ was, however, smaller than 0.4. This could be because some of the assumptions made (see MATERIALS AND METHODS) were violated and/or because the model does not include the effects of online corrections during the movement, but it could also be related to the fact that there are optimal combinations of eye and head movement that minimize gaze variability (Saglam et al. 2011) and effort (Kardamakis and Moschovakis 2009). The eye and head contributions are therefore not truly task irrelevant. It seems reasonable that changes in these contributions are tolerated as long as they are near this optimal combination, but the contributions may be adjusted when they drift too far away from the optimum. This illustrates that the random walk plus noise model is also a simplification for the behavior of the task-irrelevant component in this experiment. However, the positive value of the $ACF_{irr}(1)$ does support the hypothesis that the effects of planning noise accumulate over movements, but some planning corrections may be applied as well. This experiment thus demonstrates that the dissociation between task-relevant and task-irrelevant autocorrelations also applies to redundant motor systems, and it provides additional evidence that the random effects of planning noise accumulate over movements.

Experiment 3. We next considered measures of the trajectory of goal-directed movements. For a goal-directed movement, only the end state is relevant, not how it is achieved. For instance, for a saccadic eye movement to a small target, the saccade end point is task relevant, but trajectory measures, such as the movement speed, are not directly relevant. The saccade amplitude is therefore task relevant, whereas the peak velocity can be considered task irrelevant. We analyzed data from an experiment in which subjects made a series of 75 consecutive saccades from a fixed starting location to a fixed, small target (van Beers 2007). In the example in Fig. 4A, the amplitude varied around its mean value with no obvious relation between consecutive trials; its $ACF(1)$ was 0.15. In contrast, the peak velocity exhibited variations on a longer time scale, giving rise to an $ACF(1)$ of 0.43. Across the whole data set, the $ACF(1)$ of peak velocity was significantly ($P = 0.0005$) greater than zero and significantly larger ($P = 0.001$) than that

of amplitude, which was close to zero (Fig. 4B). We conducted a similar analysis on a series of 30 successive reaching movements to small, circular targets (van Beers 2009); Fig. 4B shows that a similar dissociation was found there. The $ACF(1)$ of peak velocity was significantly ($P = 0.0002$) greater than zero and significantly larger ($P = 0.0003$) than that of amplitude.

For both arm and eye movements, the mean $ACF(1)$ of peak velocity was positive but rather small (~ 0.14). It is not possible to make quantitative predictions for these $ACF(1)$ s using the random walk plus noise model, because amplitude and peak velocity are not independent. Within a set of movements to the same target, movements with a larger amplitude have, on average, a larger speed (Messier and Kalaska 1999; Smeets and Hooge 2003; van Beers 2007). The near-zero $ACF(1)$ of amplitude will therefore tend to reduce the $ACF(1)$ of peak velocity. Another factor that limits this autocorrelation is that there are optimal velocity profiles and peak velocities that minimize the end-point variance (Harris and Wolpert 1998, 2006; van Beers 2008). Just like the eye and head contributions in *experiment 2*, changes in peak velocity may therefore be tolerated within a certain window only. As a result, peak velocity is not completely task irrelevant for goal-directed movements. However, despite these interdependences of peak velocity and amplitude, peak velocity had a positive autocorrelation, which supports the idea that the random effects of planning noise accumulate over movements.

DISCUSSION

We found a systematic dissociation between the serial dependence of task-relevant and task-irrelevant movement components. Whereas the $ACF(1)$ of task-relevant movement parameters was always close to zero, task-irrelevant parameters always had a positive $ACF(1)$. This is a very general phenomenon, as it was found for a redundant task, for a redundant motor system, and for trajectory measures of goal-directed movements. Moreover, it was found for movements of the arm, the eye, and the head, and a similar effect has also been found for leg movements during walking (Dingwell et al. 2010). This suggests that the commonly made assumption (Schmidt et al. 1979; van Beers 2007; van Beers et al. 2004; Wing and Kristofferson 1973)—that variations in successive movements are independent—is incorrect.

We attribute the positive autocorrelation to an accumulation of the random effects of planning noise across movements. In the absence of corrections, this will lead to a random walk of motor planning (Eq. 2). It is, however, unlikely that a pure

Fig. 4. Results of *experiment 3*. A: amplitude and peak velocity of 75 consecutive saccades of a representative subject from a fixed starting location to a target, 9° up and to the left. B: mean $ACF(1)$ for movement amplitude and peak velocity for saccades (Eye; left) and for reaching movements (Arm; right). Error bars: ± 1 SE (between subjects). *** $P < 0.001$.

random walk will occur, because there are many reasons why it can be beneficial to make some corrections, even for what we have been referring to as task-irrelevant parameters. In *experiment 1*, it makes sense to make corrections when approaching the end of the target line. In *experiments 2* and *3*, the task-irrelevant parameters have preferred values as they affect the task performance indirectly by shaping the movement variability. In other situations, task-irrelevant parameters may have preferred values, because these minimize the sensitivity to errors in the control signal (Cusumano and Cesari 2006), or they maximize the error tolerance (Sternad et al. 2011). Corrections in task-irrelevant parameters may also occur, because the motor output does not consist of strictly independent task-relevant and task-irrelevant components [see Dingwell et al. (2010) for an example]. Two other problems of the random walk plus noise model are: 1) that this model predicts that the variance in the planned aim point will grow in proportion to the number of movements executed, which is unrealistic, as drifts are limited by physiological ranges, and 2) that this model does not include the effects of online corrections during the movements, which will also affect the serial dependence of motor planning. Together, there are many reasons why the random walk plus noise model is a simplification of the true behavior of the task-irrelevant component. A more realistic model could include a stationary random walk, such as an Ornstein-Uhlenbeck process (Papoulis 1991), which bounds the variance over time, but this will depend on the specific situation and is beyond the scope of this study. Here, we focused on accumulation of the random effects of planning noise, and although there may have been some control in what we have been referring to as task-irrelevant parameters, the positive ACF(1) that we found for all of these parameters supports the idea that the random planning effects accumulate over movements.

We emphasize that the dissociation between the serial dependence of task-relevant and task-irrelevant movement dimensions does not mean that these two dimensions follow completely different rules or that the random effects accumulate only in the task-irrelevant dimension. Instead, they accumulate in both dimensions (see *Eq. 5*). The only difference between the two dimensions is that there are substantial trial-by-trial planning corrections (i.e., a positive learning rate B) in the task-relevant dimension and little or no planning corrections (i.e., a zero or near-zero B) in the task-irrelevant dimension. In the absence of planning correction, motor planning follows a random walk, leading to a large autocorrelation. In contrast, if trial-by-trial planning corrections are made, then these corrections counteract the random walk, giving rise to a small autocorrelation. The larger the corrections, the smaller the ACF(1) will be (van Beers 2012). If planning corrections are made such that the motor variability is minimized, then the autocorrelation will be zero (van Beers 2009). Our finding that the $ACF_{rel}(1)$ was close to zero in all experiments thus suggests that our motor system generally aims to minimize task-relevant motor variability.

Previous work on the UCM for redundant motor systems (Latash et al. 2002, 2007; Scholz and Schöner 1999) has shown that there is more variability in the task-irrelevant than in the task-relevant subspace of all joint movements. Our results complement this finding by showing that the two subspaces also have different serial dependence. The relation between these findings is that the random-walk behavior can explain

how the large variability within the UCM comes about. This also suggests that the “minimum intervention principle” of optimal feedback control (Todorov and Jordan 2002), which states that deviations from the average movement trajectory are corrected only when they interfere with task performance, applies not only within movements but also between movements.

It may seem counterintuitive that random effects of motor planning accumulate over movements, but it is understandable from the stochastic nature of synaptic plasticity. In motor planning, a certain population of neurons forms an inverse model that generates a motor command given the current and the desired state of the body (Wolpert and Ghahramani 2000). Which motor command is generated is determined by the synaptic strengths between the neurons in the population. These synaptic strengths can change during motor learning, but they also display random variations due to stochastic processes at the cellular level, such as chance coincidences between presynaptic and postsynaptic events, probabilistic spike generation, and probabilistic release of neurotransmitters (Medina and Mauk 1999; Ribrault et al. 2011; Sejnowski 1977; Seung 2003). As a result of these random variations, synaptic strengths will undergo a random walk in the absence of motor-error signals (Medina and Mauk 1999; Sejnowski 1977; Seung 2003) and so will the internal model. When there is a motor-error signal (a difference between expected and observed movement), this error signal may bias the random walk, such that future movements will be more accurate. Complex spikes generated by Purkinje cells in the cerebellum, which are known to be important for motor learning (Gilbert and Thach 1977), could be responsible for this biasing of the random walk (Medina and Lisberger 2008), possibly by interacting with noise in parallel fibers (Kitazawa 2002).

Although the explanation just offered is consistent with our current understanding of the stochastic nature of synaptic plasticity, it is unclear whether the effects of synaptic noise are large enough to explain the observed random walks. An alternative, not mutually exclusive, explanation is that the nervous system may use an active mechanism to produce random walks. It may do so because there are advantages of such a plastic system compared with a system in which motor planning is, once learned, fixed and inflexible. Flexibility is necessary for motor planning to remain accurate in situations where the input-output relation of the motor system changes, such as when muscles fatigue or when carrying an object. In such dynamic conditions, which are the rule rather than the exception, it can be beneficial to continuously explore the space of possible motor commands to find the best command and not settle into local minima. Motor exploration has been demonstrated in the central planning signals of reaching movements of nonhuman primates (Mandelblat-Cerf et al. 2009) and of adult birdsongs (Tumer and Brainard 2007). Similarly, tuning curves of neurons in the primary motor cortex and the supplementary motor area have been observed to exhibit slow, random drifts during repeated reaching movements (Rokni et al. 2007). Accumulation of the effects of planning noise naturally leads to such exploration. Optimization algorithms, such as simulated annealing (Kirkpatrick et al. 1983) and stochastic gradient descent (Seung 2003), are based on the same idea. Motor planning in the brain can therefore be viewed as a neural approximation of such optimization algorithms. As a result, the

brain does not need to perform complicated calculations to find motor commands that minimize some cost function, but accumulation of the effects of planning noise, combined with appropriate error signals that reflect the cost function, will guide planning toward the optimal command. Such a strategy can find the optimal movement time in the presence of a speed-accuracy tradeoff (Brenner and Smeets 2011), and it may even converge to motor commands that minimize movement variability (Kitazawa 2002).

Taking account of random walks is important when interpreting motor-related neural activity. This is especially relevant for brain-computer interfaces, which must deal with nonstationary inputs (McFarland et al. 2006). A better understanding of the random-walk nature of motor planning may inspire the development of algorithms that are better able to filter the irrelevant components from a recorded neural signal so that the relevant component can be extracted more accurately. Accumulation of the effects of planning noise should also be included in models for trial-by-trial motor learning, as a failure to do so will lead to misestimation of the motor-learning process.

Finally, the dissociation between task-relevant and task-irrelevant serial dependence can be used as a tool to determine which aspects of a task a motor system attempts to optimize and which it does not. It is widely assumed that movements are planned to minimize a certain cost function (Todorov 2004), but it is often not known what exactly the cost function is. The serial correlation of candidate elements of the cost function can reveal whether an element is included, as elements that are included can be expected to have an ACF(1) near zero, whereas elements that are not included will have a positive autocorrelation.

ACKNOWLEDGMENTS

We thank Robert Wilson for assistance with running his change-point algorithm.

GRANTS

Support for this research was provided by the Netherlands Organisation for Scientific Research (NWO Grant 453-08-004).

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

R.J.v.B., E.B., and J.B.J.S. conception and design of research; R.J.v.B. performed experiments; R.J.v.B. and J.B.J.S. analyzed data; R.J.v.B., E.B., and J.B.J.S. interpreted results of experiments; R.J.v.B. prepared figures; R.J.v.B. drafted manuscript; R.J.v.B., E.B., and J.B.J.S. edited and revised manuscript; R.J.v.B., E.B., and J.B.J.S. approved final version of manuscript.

REFERENCES

- Adams RP, MacKay DJC. *Bayesian Online Changepoint Detection*. Ithaca, NY: Cornell University Library, arXiv.org, 2007, 0710.3742.
- Bernstein NA. *The Co-ordination and Regulation of Movement*. Oxford: Pergamon, 1967.
- Brenner E, Smeets JBJ. Quickly “learning” to move optimally. *Exp Brain Res* 213: 153–161, 2011.
- Chatfield C. *The Analysis of Time Series. An Introduction*. Boca Raton, FL: Chapman & Hall/CRC, 2004.
- Cheng S, Sabes PN. Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *J Neurophysiol* 97: 3057–3069, 2007.
- Churchland MM, Afshar A, Shenoy KV. A central source of movement variability. *Neuron* 52: 1085–1096, 2006.
- Cusumano JP, Cesari P. Body-goal variability mapping in an aiming task. *Biol Cybern* 94: 367–379, 2006.
- Diedrichsen J, White O, Newman D, Lally N. Use-dependent and error-based learning of motor behaviors. *J Neurosci* 30: 5159–5166, 2010.
- Dingwell JB, John J, Cusumano JP. Do humans optimally exploit redundancy to control step variability in walking? *PLoS Comput Biol* 6: e1000856, 2010.
- Gilbert PFC, Thach WT. Purkinje cell activity during motor learning. *Brain Res* 128: 309–328, 1977.
- Gordon J, Ghilardi MF, Ghez C. Accuracy of planar reaching movements. 1. Independence of direction and extent variability. *Exp Brain Res* 99: 97–111, 1994.
- Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Harris CM, Wolpert DM. The main sequence of saccades optimizes speed-accuracy trade-off. *Biol Cybern* 95: 21–29, 2006.
- Hinkley DV. Inference about the change-point in a sequence of random variables. *Biometrika* 57: 1–17, 1970.
- Jones KE, Hamilton AF, Wolpert DM. Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88: 1533–1544, 2002.
- Kardamakis AA, Moschovakis AK. Optimal control of gaze shifts. *J Neurosci* 29: 7723–7730, 2009.
- Kendall MG. Note on bias in the estimation of autocorrelation. *Biometrika* 41: 403–404, 1954.
- Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science* 220: 671–680, 1983.
- Kitazawa S. Optimization of goal-directed movements in the cerebellum: a random walk hypothesis. *Neurosci Res* 43: 289–294, 2002.
- Latash ML, Scholz JP, Schöner G. Motor control strategies revealed in the structure of motor variability. *Exerc Sport Sci Rev* 30: 26–31, 2002.
- Latash ML, Scholz JP, Schöner G. Toward a new theory of motor synergies. *Motor Control* 11: 276–308, 2007.
- Mandelblat-Cerf Y, Paz R, Vaadia E. Trial-to-trial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *J Neurosci* 29: 15053–15062, 2009.
- Marriott FHC, Pope JA. Bias in the estimation of autocorrelations. *Biometrika* 41: 390–402, 1954.
- Marshall RJ. Autocorrelation estimation of time series with randomly missing observations. *Biometrika* 67: 567–570, 1980.
- McFarland DJ, Anderson CW, Müller KR, Schlögl A, Krusienski DJ. BCI Meeting 2005—workshop on BCI signal processing: feature extraction and translation. *IEEE Trans Neural Syst Rehabil Eng* 14: 135–138, 2006.
- Medina JF, Lisberger SG. Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. *Nat Neurosci* 11: 1185–1192, 2008.
- Medina JF, Mauk MD. Simulations of cerebellar motor learning: computational analysis of plasticity at the mossy fiber to deep nucleus synapse. *J Neurosci* 19: 7140–7151, 1999.
- Messier J, Kalaska JF. Comparison of variability of initial kinematics and endpoints of reaching movements. *Exp Brain Res* 125: 139–152, 1999.
- Papoulis A. *Probability, Random Variables, and Stochastic Processes*. New York: McGraw-Hill, 1991.
- Ribrault C, Sekimoto K, Triller A. From the stochasticity of molecular processes to the variability of synaptic transmission. *Nat Rev Neurosci* 12: 375–387, 2011.
- Rokni U, Richardson AG, Bizzi E, Seung HS. Motor learning with unstable neural representations. *Neuron* 54: 653–666, 2007.
- Saglam M, Lehen N, Glasauer S. Optimal control of natural eye-head movements minimizes the impact of noise. *J Neurosci* 31: 16185–16193, 2011.
- Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT. Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol Rev* 47: 415–451, 1979.
- Scholz JP, Schöner G. The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126: 289–306, 1999.
- Sejnowski TJ. Statistical constraints on synaptic plasticity. *J Theor Biol* 69: 385–389, 1977.
- Seung HS. Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* 40: 1063–1073, 2003.

- Smeets JBJ, Hooge ITC.** Nature of variability in saccades. *J Neurophysiol* 90: 12–20, 2003.
- Smeets JBJ, van den Dobbelaars JJ, de Grave DDJ, van Beers RJ, Brenner E.** Sensory integration does not lead to sensory calibration. *Proc Natl Acad Sci USA* 103: 18781–18786, 2006.
- Sober SJ, Wohlgenuth MJ, Brainard MS.** Central contributions to acoustic variation in birdsong. *J Neurosci* 28: 10370–10379, 2008.
- Sternad D, Abe MO, Hu X, Müller H.** Neuromotor noise, error tolerance and velocity-dependent costs in skilled performance. *PLoS Comput Biol* 7: e1002159, 2011.
- Sternad D, Park SW, Müller H, Hogan N.** Coordinate dependence of variability analysis. *PLoS Comput Biol* 6: e1000751, 2010.
- Tabak S, Smeets JBJ, Collewyn H.** Modulation of the human vestibuloocular reflex during saccades: probing by high-frequency oscillation and torque pulses of the head. *J Neurophysiol* 76: 3249–3263, 1996.
- Todorov E.** Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.
- Todorov E, Jordan MI.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
- Tumer EC, Brainard MS.** Performance variability enables adaptive plasticity of “crystallized” adult birdsong. *Nature* 450: 1240–1244, 2007.
- van Beers RJ.** How does our motor system determine its learning rate? *PLoS One* 7: e49373, 2012.
- van Beers RJ.** Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63: 406–417, 2009.
- van Beers RJ.** Saccadic eye movements minimize the consequences of motor noise. *PLoS One* 3: e2070, 2008.
- van Beers RJ.** The sources of variability in saccadic eye movements. *J Neurosci* 27: 8757–8770, 2007.
- van Beers RJ, Haggard P, Wolpert DM.** The role of execution noise in movement variability. *J Neurophysiol* 91: 1050–1063, 2004.
- Wei WWS.** *Time Series Analysis. Univariate and Multivariate Methods.* Boston: Pearson, 2006.
- Wilson RC, Nassar MR, Gold JI.** Bayesian online learning of the hazard rate in change-point problems. *Neural Comput* 22: 2452–2476, 2010.
- Wing A, Kristofferson A.** Response delays and the timing of discrete motor responses. *Percept Psychophys* 14: 5–12, 1973.
- Wolpert DM, Ghahramani Z.** Computational principles of movement neuroscience. *Nat Neurosci* 3, Suppl: 1212–1217, 2000.