

Quickly ‘learning’ to move optimally

Eli Brenner · Jeroen B. J. Smeets

Received: 14 February 2011 / Accepted: 28 June 2011 / Published online: 12 July 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract People take account of the variability in their movements in a near-optimal manner in various visuo-motor tasks. Is knowledge of one’s variability needed for such near-optimal performance, or could it arise from responding to one’s success in previous attempts in some simple manner? We asked subjects to move a pen back and forth across a tablet to make a cursor move as quickly as possible between two targets. The cursor had to stop within the targets. Task difficulty was varied between blocks. Part of the variation in difficulty was explicit (three target sizes) whereas the rest had to be discovered during the movements (two mappings between the movements of pen and cursor). In all cases, subjects sped up after stopping within a target and slowed down after failing to do so. We interpret this as evidence that explicit knowledge of one’s variability is not necessary for performing close to optimally.

Keywords Fitts’ law · Movement timing · Motor learning · Speed-accuracy trade-off · Optimal performance

Introduction

There is variability in even the simplest of human arm movements. This is inevitable because there are limitations to the spatial and temporal resolution with which target positions and one’s own posture can be judged from the

sensory information, and limitations on how precisely the hand can be made to move in a chosen manner. People are aware of such variability and pick trajectories accordingly. For instance they veer away from obstacles and approach target surfaces orthogonally (Brenner and Smeets 2007). Many have argued that people pick the *optimal* strategy, meaning that people move in the manner that gives the best performance considering the variability in the movements and the task constraints (Diedrichsen et al. 2010; Körding and Wolpert 2004, 2006; Todorov and Jordan 2002; Trommershäuser et al. 2003, 2008). But how do people know how to achieve this? Do they develop reliable estimates of their own variability for a wide range of movements and combine these estimates with the costs and benefits of the likely movement outcomes given all possible strategies under the pertaining conditions? This would imply that they know the relationship between trajectory, speed and variability for a very wide range of movements.

In the current study we explore an alternative strategy that might explain the close-to-optimal performance in repetitive movements: quickly adapting each movement to success or failure during the previous movement. We do so for the relationship between speed and accuracy in a simple reciprocal aiming task (Fitts 1954) in which people move a pen across a tablet to move a cursor back and forth between two targets on a screen. The perceived size of the targets and the distance between them could be combined with knowledge about the variability in the motion of the cursor to select the appropriate movement speed. However, subjects are unlikely to directly move at the optimal speed, because the variability in the motion of the cursor depends on the mapping between pen and cursor. We augment the influence of such mapping by sometimes using a linear mapping between pen and cursor, and sometimes a non-linear mapping that makes the cursor move faster *between*

E. Brenner (✉) · J. B. J. Smeets
Research Institute MOVE, Faculty of Human Movement
Sciences, VU University, van der Boechorststraat 9,
1081 BT Amsterdam, The Netherlands
e-mail: e.brenner@vu.nl

the targets and more slowly *near* the targets (following Fernandez and Bootsma 2004, 2008).

There are two ways in which people could adapt the speed of their movements to the different mappings. They could use their experience during the task to estimate their own variability under the prevailing circumstances, or they could directly adjust their speed on the next trial on the basis of whether they hit or missed the target. In both cases experience with the specific conditions will result in changes in movement speed. Moreover, in both cases the average speed will change gradually if the circumstances call for a faster or slower movement than one had initially expected. However, the extent to which such changes depend on whether or not the target was hit in the previous trial differs between the two explanations. If people speed up after a hit and slow down after a miss, we should be able to see that the movement time on each trial depends on the success in the previous trial. If subjects do not respond to recent success, but adjust the speed of their movements until their estimate of the variability in their actions is compatible with the requirements of the task, the movement time on each trial should hardly depend on the success in the immediately preceding trial because the estimate of variability should not change much with a single additional trial. Note also that in order to know how to adjust the movement speed on the basis of one's estimated variability, it is not enough to estimate one's variability at a single speed; one also has to know one's variability for other movement speeds.

Methods

We used a sub-set of the conditions of the study by Fernandez and Bootsma (2004). Ten subjects each performed one block of trials for each of six conditions in random order. The data was collected during a methodology course at the Faculty of Human Movement Sciences.

Task and conditions

The subjects sat about 70 cm from a computer screen (CRT; 47.8 by 30.4 cm; 1,096 by 686 pixels; 160 Hz) in a normally illuminated room. They were encouraged to place the graphics tablet (12 by 9 cm) at a comfortable position. They were informed that the point at which the cursor changes direction will be taken as the endpoint of the movement and that they should not lift the stylus off the tablet's surface. Their task was to move the cursor (a 4.4 mm diameter red disc) as quickly as possible between two circular targets (19.2, 38.4 or 76.8 mm diameter black discs, 307.2 mm apart, on a white background). They were to move back and forth between the targets irrespective of

whether they succeeded to land within the target (thus, if they missed a target they were to continue by moving to the other target). They did not receive explicit feedback about their performance in terms of whether they hit the target, but they could see the cursor at all times so they could easily judge their success.

For each target size the task was performed once for each of the two mappings (Fig. 1a). With one mapping the cursor's position on the screen (c , measured in mm from the screen centre) was proportional to the pen's position on the tablet (p , measured in mm from the centre of the tablet).

$$c = 3.98p \quad (1)$$

With the other mapping, the relationship between lateral pen position (p) and horizontal cursor position (c) was non-linear, so that it became easier for subjects to stop within the targets (the relationship between sagittal pen position and vertical cursor position was linear, following Eq. 1. The horizontal position on the screen is given by

$$c = \left(\frac{1.2}{1 + e^{-5\frac{3.98p}{D}}} + 0.4 \frac{3.98p}{D} - 0.6 \right) D \quad (2)$$

where D is the lateral distance of the centre of the target from the screen centre (153.6 mm). This non-linear mapping is equivalent to subjects suddenly having much less lateral variability in their movements when they are near the targets. The values in (1) and (2) are chosen in such a way that for both mappings the centre of the targets on the screen corresponds to the same position on the tablet. Subjects were not told in advance whether a block of movements would be performed with the linear or the non-linear mapping. Subjects started each block by bringing the cursor to one of the targets. They then made 50 movements (25 to the right and 25 to the left). This was repeated for each of the six conditions (3 target sizes; 2 mappings) in random order. There was a short break (about half a minute) between the blocks.

Data analysis

Movements were segregated at the moments that the pen reached its extreme lateral positions. We considered the measured positions of the pen at those moments to be where the movements ended. Since the centre of the cursor followed the measured position of the pen, about half the cursor had to overlap with the target for the movement to be considered to have ended within the target. We determined whether each movement ended within or outside the target (labelled hit and miss respectively), and how long it had taken to reach this endpoint from the time one had reached the previous endpoint. For each subject and condition we then determined the average movement time and the (lateral) standard deviations of the endpoints for each of

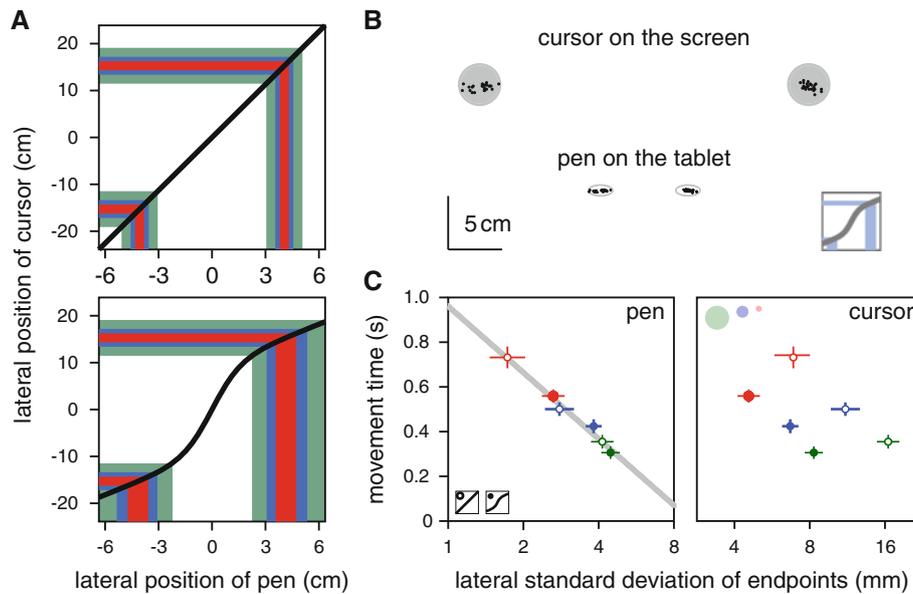


Fig. 1 Variation on Fernandez and Bootsma's (2004) rapid pointing task, in which subjects had to move a cursor back and forth as quickly as possible between two targets by moving a pen across a graphics tablet. **a** Two mappings between the horizontal coordinates of pen and cursor. Normally, motion of the cursor is a linearly scaled version of that of the pen (*top panel*). People move faster if more of the space on the tablet is attributed to the region on the screen containing the targets and less to the space between the targets (*lower panel*). The three target sizes are shown in *different colours*. Note that the non-linear mapping makes a larger part of the tablet fall 'within' each target on the screen. **b** Example showing the endpoints (*black dots*) of the pen and the cursor for a trial with the non-linear mapping. The

grey areas represent the visible targets on the screen. The *ellipses* represent the regions on the tablet that correspond with these targets. **c** The relationship between the movement times and the standard deviations in the lateral positions of the endpoints of the pen and cursor. *Open* and *solid* symbols: linear and non-linear mapping. Movements were faster with the non-linear mapping, but the relationship between the speed and accuracy of the pen did not change (the points fall on a single *line* when expressed in terms of endpoints of the pen: *grey line* with a slope of -0.3 for the relationship between movement time in seconds and the binary logarithm of the standard deviation of the endpoints in mm). *Error bars* indicate standard errors across subjects

the two targets. The latter measure was determined both for the pen on the tablet and for the cursor on the screen (these are not equivalent because of the different mappings; Fig. 1b) and the values for the two targets were averaged. Relating the average movement time to the variability in performance allows us to evaluate the benefit of the non-linear mapping.

Since we expect subjects to move faster when a larger part of the tablet is devoted to the region near the targets (Fernandez and Bootsma 2004), but subjects did not know the mapping in advance, we expected to see changes in the movement time across consecutive trials within a sequence. We therefore averaged the movement time across subjects for consecutive movements (within each condition).

If subjects adjust their movement speed directly to their success on individual trials, they will slow down (increase their movement time) whenever they miss the target and speed up (decrease their movement time) whenever they hit the target. To evaluate whether they do so we determine the average change in movement time after a hit and after a miss (for every subject and mapping). We assess the reliability of effects across subjects with a repeated measures

analysis of variance with factors *mapping* and *previous performance* (*hit* or *miss*). Even for completely random variations in movement time across trials, the movement time could appear to increase after a miss if subjects miss more often when they happen to move faster. If so, subjects will not only move more slowly on the movements after the ones on which they missed the target, but also on movements before the ones in which they missed the target. We can therefore distinguish between a correlation of speed with accuracy and actively responding to success on the previous trial by also determining the average change in movement time *before* a hit or a miss.

To confirm that the adjustments to the movement speed are effective in regulating the endpoint variability, we divided each subject's movements in each condition into the faster half and the slower half, and determined the standard deviation within each half. To evaluate whether adjustments to the speed are different for the initial part of the movement than for the final part, and whether this depends on the mapping, we determined the proportion of the movement time that was spent reaching the midline (half way between the targets) and related this to the movement time itself.

We also examined whether subjects adjust the movement amplitude to errors on previous trials. We did so by comparing how far the pen stopped from the centre of the tablet (in the lateral direction) on attempts to hit each of the two targets. For each trial we subtracted this distance for the trial before the trial in question from this distance for the trial after the trial in question. We averaged these differences between the next and previous trial separately for trials on which subjects hit, overshoot or undershot the target. Misses were divided into overshoots and undershoots on the basis of the horizontal component being larger or smaller than the distance to the target centre. The reliability of changes in movement amplitude were assessed with a repeated measures analysis of variance.

Results

There were 500 movements for each condition (10 subjects, 50 movements each). When the mapping between the pen on the tablet and the cursor on the screen was linear, the pen missed the target on 137, 73 and 34 movements for the small, middle and large target. For the non-linear mapping, it missed the target on 76, 28 and 9 movements for the small, middle and large target. There were considerable differences between subjects in the number of errors they made. These differences are related to the subjects' average movement times: subjects who moved faster made more errors (negative correlation across subjects between movement time and number of errors: $r = -0.68$).

Figure 1c shows the relationship across conditions between variability in the movement endpoints and movement time. Not surprisingly, movements were faster and more variable for larger targets. Movements were also faster when stopping within the target was made easier by the non-linear mapping (solid symbols lower than corresponding open symbols). This increase in movement speed resulted in a corresponding increase in variability of the *pen* on the tablet (solid symbols 'in line' with open ones). The non-linear mapping changed the relationship between speed and accuracy in terms of variability of the *cursor* on the screen, which of course is what this non-linear mapping was meant to do.

Although the non-linear mapping made subjects move faster than they did for the linear mapping, they did not move so much faster as to miss the target on as many trials as for the linear mapping. Figure 2a shows how the average movement time declined with consecutive movements. To investigate whether this decline was driven by seeing whether one succeeded to end the previous movement within the target, we examined the change in movement time after hits and misses. Figure 2b shows that subjects tended to decrease their movement time after a hit and to

increase it after a miss (on average by 5 ± 1 ms and 42 ± 10 ms respectively; mean \pm standard error across subjects' mean values; significant main effect of previous performance: $F_{1,9} = 33$; $P = 0.0003$).

The increase in movement time after a miss did not only influence the next movement. The movement time changed back gradually: five movements after missing the target the average movement time was still 23 ± 11 ms higher than when the target was missed (which is about what one would expect for four 5 ms decreases after a 42 ms increase in movement time). Subjects tended to change their movement time less with the non-linear mapping (average changes of -3 and 34 ms after a hit and a miss) than with the linear mapping (average changes of -8 and 50 ms), but the interaction between the factors *mapping* and *previous performance* in the analysis of variance was not significant ($F_{1,9} = 1.1$; $P = 0.31$). Even if we were to take the non-significant differences between the magnitudes of the changes for the two mappings seriously, they could partly be explained by the change being proportional to the movement time, because the average movement time was longer with the linear mapping (open symbols in Fig. 1c). The average change after a miss was 7% with the non-linear and 9% with the linear mapping.

The average duration of movements after misses is longer than the average duration of the movements leading to the misses themselves (Fig. 2b). This is not just because misses mainly occur on trials in which one moved fast, because if that were the case we would find similarly longer durations for the movements before the ones leading to the misses, which we do not (Fig. 2c). Thus Fig. 2 shows that subjects speed up if they perform well and slow down if they miss the target.

Figure 3a confirms that movement speed is the primary determinant of the standard deviation in the endpoints (Harris and Wolpert 1998; Schmidt et al. 1979): the relationship between movement speed and endpoint variability is similar when comparing the faster and slower half of the movements within each condition, to when comparing the different conditions. Figure 3b shows that the time spent traversing the first half of the distance between the targets decreased less when the movement time decreased than did the time spent traversing the second half of the distance. The relationship between the proportion of time spent on the first half of the distance and the total movement time was independent of the mapping.

This study was designed to study movement speed, but the movements obviously also had to end within the targets, so we can look at movement amplitude in a similar manner as for movement speed. Only one subject undershot the target at least 5 times for both mapping conditions, and only six subjects did so for both mapping conditions together, so we ignored the mapping condition and only

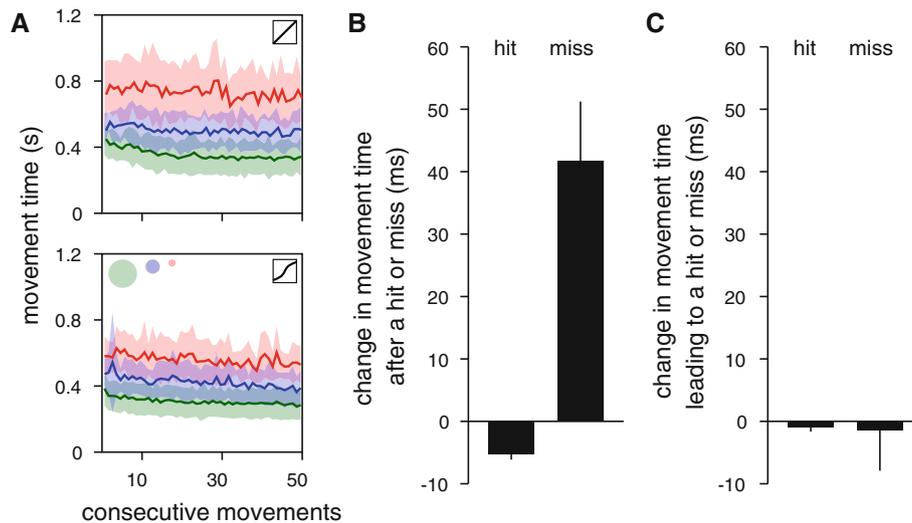
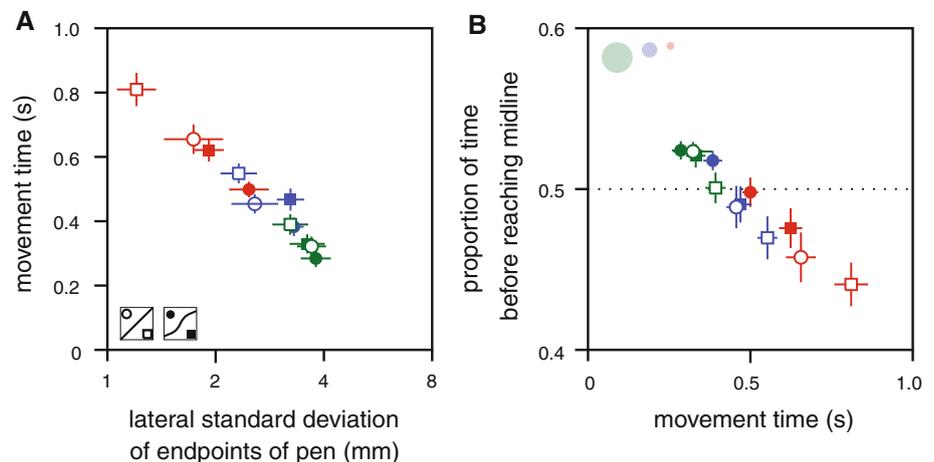


Fig. 2 Learning the adequate response. **a** Changes in movement time during sets of movements. *Red, blue and green* curves represent movements between the small, intermediate and large targets, respectively. The *shaded areas* are standard errors across subjects. **b** Influence that hitting or missing the target has on the movement time on the next trial. The average changes in movement time

between the movements that resulted in a hit or miss and the next movements were determined separately for each subject (irrespective of the condition), and then averaged across subjects. *Error bars* are standard errors across subjects. **c** Change in movement time with respect to the previous movement when the target was hit or missed. Other details as in **b**

Fig. 3 Differences between conditions follow the same relationship as differences between the faster (*circles*) and slower (*squares*) half of the trials within each condition. Coding of mapping and target size as in Fig. 1c. **a** Relationship between speed and accuracy. **b** For longer movement times, the time spent after reaching the midline increased more than did the time that it took to reach the midline, irrespective of the mapping



considered the latter six subjects for our analysis of movement amplitude. Analysis of variance revealed a significant effect of previous performance (*hit, overshoot or undershoot*; $F_{2,10} = 7.9$; $P = 0.009$). For movements that ended within the target, the average difference between the lateral eccentricity of the pen's endpoints on the next and on the previous movement to that target was slightly smaller than zero (-0.1 mm). For movements that ended beyond the target (overshoots), the average decrease in eccentricity was larger (-0.3 mm). For movements that ended before the target (undershoots), there was an *increase* in eccentricity (1.1 mm). This suggests that the endpoint is adjusted on the basis of success on previous trials in a similar manner as is the movement time.

Discussion

Reinterpreting Fernandez and Bootsma (2004)

Our experiment replicates the results of Fernandez and Bootsma (2004) in showing that subjects move their hand faster with the non-linear mapping. Fernandez and Bootsma interpreted this in terms of being able to use visual feedback more effectively because the cursor is moving more slowly when it is close to the target. We propose an alternative explanation: subjects moved faster with the non-linear mapping because they missed fewer targets so they kept speeding up. The following three paragraphs explain why we favour this explanation.

By plotting movement time against variability on the *tablet* we show that the relationship between speed and accuracy of the hand does not depend on the mapping (Fig. 1c). If seeing the cursor move more slowly near the target with the non-linear mapping had helped move the hand more accurately, the variability of the *pen* would have been lower for an equivalent movement time for the non-linear mapping. It is not, so there is no reason to believe that the non-linear mapping improved the visual guidance of the ongoing movement of the pen.

Figure 1c shows that longer movement times are associated with less endpoint variability when comparing the mean performance for the six conditions. Figure 3a shows that the same relationship between movement time and endpoint variability holds for a comparison between faster and slower trials within each condition. A negative correlation between movement time and number of errors when comparing individual subjects' average performance suggests that a similar relationship holds across subjects. Thus, endpoint variability depends on how quickly the hand moves (perhaps for fundamental reasons related to producing the movements; Harris and Wolpert 1998).

Comparing the results shown in Fig. 2b and c shows that the fluctuations in movement time are not just random variability, with a correlation between moving slowly and moving accurately. Combining random variability in movement speed with a larger endpoint variability for faster movements could account for our having found slower movements after misses (Fig. 2b), because if the likelihood of a movement being a miss is higher when subjects move faster than average, the movement after a miss is likely to be slower. However, following the same reasoning, the previous movement is just as likely to be slower than average, which is clearly not the case (Fig. 2c). Thus we can conclude that the fluctuations in movement time are at least partly caused by subjects responding to their success or failure in hitting the target. They may sometimes misjudge whether they hit the target, and probably respond in a more graded manner than described above, with larger changes when they clearly hit or miss the target than when the movement ends close to the target's border, but this does not change the overall reasoning.

Speed and accuracy

Subjects missed fewer targets with the non-linear mapping, probably because the difference between their initial movement time and how fast they could move without missing the targets was larger (Fig. 2a). However, the movement times for the two mappings appear to differ from the very first trial. Perhaps subjects started adjusting their movement speed before what we call the first trial, because the mapping was already present when moving to

the starting point. Alternatively, there may be *some* benefit of seeing the cursor close to the target slightly longer (as proposed in Fernandez and Bootsma 2004), especially for the smallest targets for which the time left to make adjustments may not be too short (Brenner and Smeets 2003).

We show that subjects immediately adjusted their movement speed to whether they hit or missed the target (Fig. 2b). Similar adjustments to reaction times have previously been demonstrated in choice-response tasks (Rabbitt 1966). We found a steady decline in movement time during each set of movements (Fig. 2a). For the sets of movements with the non-linear mapping we expect the movement time to gradually decrease because subjects cannot know in advance that it will be easier to hit those targets. The decrease in movement time during sets of movements with a 'normal' relationship between pen and cursor movements suggests that subjects start more cautiously than necessary, but it is also possible that additional information is obtained by making the movements, such as proprioceptive information about the positions of the targets (Smeets et al. 2006). Including additional information improves precision, and will thus allow subjects to move faster without making more errors.

We asked subjects to move as fast as possible with movements that end within the targets. We did not specify how important it was to end within the targets on all trials, or how the cost of moving fast (in terms of the risk of missing the target) was to be assessed in relation to the benefit of moving faster (in terms of complying to the instruction). Part of the variability between subjects in our study is undoubtedly due to them assessing such costs and benefits differently. That the time step after a miss is larger than the time step after a hit suggests that subjects are not willing to accept missing many targets in order to be faster, in accordance with the instructions, but we cannot be more precise than that. Several studies have used tasks that make the costs and benefits less ambiguous, so that they allow one to better judge how closely performance approaches optimality. Could adjustments based on success on previous trials lead to near-optimal trade-offs in such tasks? In the next paragraphs we will evaluate such a procedure in relation to an elegantly simple pointing task with explicit costs (Trommershäuser et al. 2003, 2008).

Reinterpreting Trommershäuser et al. (2003)

We use one example from Trommershäuser, Maloney and Landy's extensive set of experiments to illustrate how people could quickly learn to perform close to optimally without any elaborate considerations based on estimates of their own variability. In the experiments, subjects made rapid pointing movements towards a screen. They gained

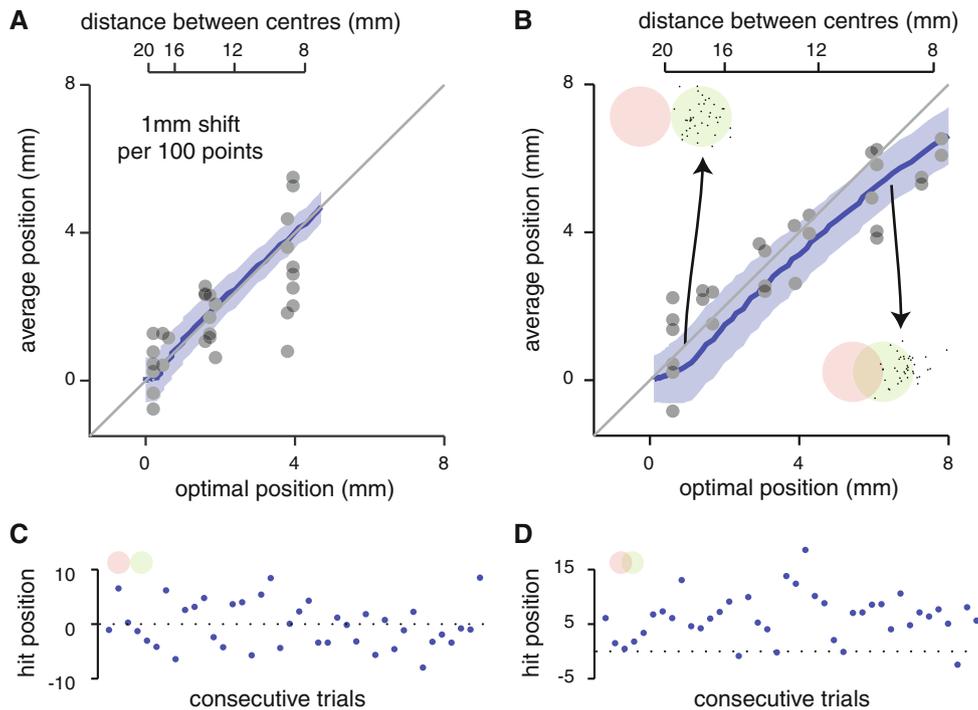


Fig. 4 Simulating performance in Trommershäuser, Maloney and Landy’s rapid pointing task with a simple algorithm that does not rely on knowing one’s variability (Eq. 3). The task is to quickly hit a green 18 mm diameter target disk to gain 100 points. Accidentally hitting a red penalty disk gives a loss of either 100 points (a) or 500 points (b). The separation between the disks is varied and the average position that is hit is compared with the optimal position to aim for. A positive position is away from the target centre with respect to the penalty region. The grey points are human pointing data from Figure 5 of Trommershäuser et al. (2003). The blue curves represent simulations for an aiming point that is updated in accordance with Eq. 3, and pointing movements with a horizontal and vertical standard deviation of 4 mm with respect to the aiming point. We used 1 mm per 100 points for g and l , so both are 1 mm when hitting the penalty disk gives a loss of 100 points (a), but l is 5 mm when the penalty is 500 points (b). Although the simulations were done in two

dimensions, only the component in the direction of a line through the centres of the two disks is shown, because there were no systematic effects in the orthogonal direction. Each position on the curve (and the standard deviation represented by the shaded area) is based on 10,000 simulations of 40 trials each (the number of trials for each separation in the experiments). The initial aiming point for each simulation is along the line through the centres of the two disks, half way across the part of the target that does not overlap with the penalty area. The curve connects simulations for separations between 7 and 22 mm. The insets in b show the hit positions on single simulations for separations of 9 and 18 mm. Positions in the relevant direction on consecutive trials of these two simulations are shown in c and d respectively. Lag-one autocorrelations in such data were negligible: -0.01 ± 0.15 and 0.00 ± 0.16 (mean \pm standard deviation) for separations of 9 and 18 mm, respectively

points by touching the screen within an indicated region. They lost points if they touched a second indicated region or if they moved too slowly. As in our study, their participants had to consider how fast to move as well as where to aim (which will be the average hit position), but their study was designed to examine how the aiming point was determined, so we model that aspect. This makes the modelling more closely related to our secondary finding that participants adjusted their movement amplitude in response to the outcome of the previous movement, than to our main finding that they adjusted the speed of the movement in response to the outcome of the previous movement. In accordance with Trommershäuser, Maloney and Landy’s experimental design, we consider how the separation between the target and the penalty region influences where subjects aim to hit the screen (see Fig. 4).

We assume that people shift their aiming point in response to having missed the target or hit the penalty disk (whereby the shift depends on the direction of the error).

Our model can be written as

$$\vec{a}_{n+1} = \vec{a}_n - g \frac{\vec{h}_n - \vec{t}}{|\vec{h}_n - \vec{t}|} H(|\vec{h}_n - \vec{t}| - r_t) + l \frac{\vec{h}_n - \vec{p}}{|\vec{h}_n - \vec{p}|} H(r_p - |\vec{h}_n - \vec{p}|) \tag{3}$$

where \vec{a}_n represents the coordinates of the aiming point on movement n , \vec{h}_n is the position that was hit on movement n , $H(x)$ is the Heaviside function (zero for $x < 0$ and one for $x \geq 0$), r_t and r_p are the radii of the target and penalty disks (both 9 mm), \vec{t} is the target centre, \vec{p} is the centre of the

penalty disk, and g and l are the magnitudes of the shifts in response to gains and losses, respectively.

Figure 4a and b show that simply shifting the aiming point in response to errors in this manner could give performance that is close to optimal, and very similar to the results of Trommershäuser et al. (2003). When the disks overlap considerably, people even appear to veer less from the 500-point penalty disk than would be optimal, in accordance with the model's performance (right side of Fig. 4b). The magnitude of the shifts is not critical (similar results are obtained for 0.5 or 2 mm shifts per 100 points), but the simulation is closest to human performance when the ratio between the magnitudes of responses to missing the target and hitting the penalty region resembles the ratio between the number of points obtained or lost.

Since the shifts in the aiming point are small in comparison with the variability they do not directly become apparent as shifts when one plots sequences of trials (Fig. 4c, d), which explains why Trommershäuser et al. did not observe shifts in the aiming point, especially if one considers that we used the worst initial aiming point that is likely to be chosen: aiming for the centre of the region that gives a positive reward, completely ignoring the fact that missing it on the side of the penalty region is worse than doing so on the other side. In the present study we also found modest spatial shifts (about 0.3 and 1.1 mm; see last paragraph of the “Results” section) in comparison with the spatial variability (several mm; Fig. 1c).

Generalising these findings

Our experiment suggests that people adjust their movement plan in response to feedback in quite a simple manner. Figure 4 shows that people could achieve near optimal performance in repetitive movements by making small, simple adjustments to their movement plan in response to feedback (for a more elaborate mechanism see van Beers 2009). Similar simple adjustments could probably lead to the trade-off between the time to plan one's action and the time to perform it (Battaglia and Schrater 2007; Faisal and Wolpert 2009): one could respond earlier whenever one fails to reach the place that one intended to reach (indicating a performance error) and later whenever one does succeed to do so but the movement is nonetheless unsuccessful (indicating a planning error). Such a strategy is likely to lead to performance that is not too far from optimal, and individual human subjects also only perform near-optimally (see systematic errors in Figure 3 of Faisal and Wolpert 2009).

The main advantage of relying on a mechanism such as the one that we propose is that there is no need to estimate one's precision. It is relatively simple to imagine how the brain could estimate sensory uncertainties (for which

optimal performance has often been found; e.g. Hillis et al. 2004; Muller et al. 2007). If the activity of different cells each indicate the likelihood of a certain value of the parameter of interest, then the cells with the highest activation will identify the most likely value, and the width of the distribution of active cells will indicate how certain one can be of this value (Ma et al. 2006). Estimating motor uncertainty is conceptually not that simple, because at least some of the variability is presumably introduced after the signals leave the brain (Jones et al. 2002). Thus it would have to be learnt through experience.

Of course estimates of movement variability could be learnt through experience, and then used to judge how to best perform the task, but we believe that subjects may directly learn how to best perform the task (for a more elaborate example of how this could occur see Franklin et al. 2008). Although learning one's variability would appear to be more flexible and more reliable, it is also more complicated, so unless the estimated variability can be used for many tasks under many circumstances it may not be worth the additional effort. We conclude that since near-optimal performance in many motor tasks could easily be achieved with simple learning strategies, it is premature to conclude that the brain computes the optimal solution by explicitly considering motor variability.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Battaglia PW, Schrater PR (2007) Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *J Neurosci* 27:6984–6994
- Brenner E, Smeets JBJ (2003) Fast corrections of movements with a computer mouse. *Spat Vis* 16:365–376
- Brenner E, Smeets JBJ (2007) Flexibility in intercepting moving objects. *J Vis* 7:14
- Diedrichsen J, Shadmehr R, Ivry RB (2010) The coordination of movement: optimal feedback control and beyond. *Trends Cogn Sci* 14:31–39
- Faisal AA, Wolpert DM (2009) Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. *J Neurophysiol* 101:1901–1912
- Fernandez L, Bootsma RJ (2004) Effects of biomechanical and task constraints on the organization of movement in precision aiming. *Exp Brain Res* 159:458–466
- Fernandez L, Bootsma RJ (2008) Non-linear gaining in precision aiming: making Fitts' task a bit easier. *Acta Psychol* 129: 217–227
- Fitts PM (1954) The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47:381–391
- Franklin DW, Burdet E, Tee KP, Osu R, Chew CM, Milner TE, Kawato M (2008) CNS learns stable, accurate, and efficient

- movements using a simple algorithm. *J Neurosci* 28:11165–11173
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784
- Hillis JM, Watt SJ, Landy MS, Banks MS (2004) Slant from texture and disparity cues: optimal cue combination. *J Vis* 4:967–992
- Jones KE, Hamilton AFD, Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88:1533–1544
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247
- Körding KP, Wolpert DM (2006) Bayesian decision theory in sensorimotor control. *Trends Cogn Sci* 10:319–326
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438
- Muller CMP, Brenner E, Smeets JBJ (2007) Living up to optimal expectations. *J Vis* 7:2
- Rabbitt PM (1966) Errors and error correction in choice-response tasks. *J Exp Psychol* 71:264–272
- Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT Jr (1979) Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol Rev* 86:415–451
- Smeets JBJ, van den Dobbelaere JJ, de Grave DDJ, van Beers RJ, Brenner E (2006) Sensory integration does not lead to sensory calibration. *Proc Natl Acad Sci USA* 103:18781–18786
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235
- Trommershäuser J, Maloney LT, Landy MS (2003) Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A* 20:1419–1433
- Trommershäuser J, Maloney LT, Landy MS (2008) Decision making, movement planning and statistical decision theory. *Trends Cogn Sci* 12:291–297
- van Beers RJ (2009) Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63:406–417