Vision Research 113 (2015) 198–210

Contents lists available at ScienceDirect

Vision Research

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

Quickly making the correct choice

Eli Brenner *, Jeroen B.J. Smeets

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

ARTICLE INFO

Article history:
Received 17 September 2014
Received in revised form 5 March 2015
Available online 23 April 2015

Keywords:
Arm movements
Decision-making
Optimal control
Speed–accuracy trade-off
Tapping
Fitts' law

ABSTRACT

In daily life, unconscious choices guide many of our on-going actions. Such choices need to be made quickly, because the options change as the action progresses. We confirmed that people make reasonable choices when they have to quickly decide between two alternatives, and studied the basis of such decisions. The task was to tap with their finger on as many targets as possible within 2 min. A new target appeared after every tap, sometimes accompanied by a second target that was easier to hit. When there was only one target, subjects had to find the right balance between speed and accuracy. When there were two targets, they also had to choose between them. We examined to what extent subjects switched to the target that was easier to hit when it appeared some time after the original one. Subjects generally switched to the easier target whenever doing so would help them hit more targets within the 2-min session. This was so, irrespective of whether the different delays were presented in separate sessions or were interleaved within one session. Whether or not they switched did not depend on how successful they were at hitting the targets on earlier attempts, but it did depend on the position of the finger at the moment that the easy target appeared. We conclude that people have continuous access to reasonable estimates of how long various movement options would take and of how precise the endpoints are likely to be, given the instantaneous circumstances.

1. Introduction

There are many situations in daily life in which it is crucial to decide fast. For instance, consider deciding whether to cross an intersection before or after an approaching car. You have to be quick to decide whether to speed up or slow down, because the longer you wait the more difficult it will be to do whatever you decide to do. The same applies to deciding whether to take an additional small step before an obstacle or to immediately take a large step across the obstacle, or to deciding whether to try to catch an approaching ball or to make sure it does not hit you. There is evidence that people are good at making all kinds of fast decisions. People very efficiently decide where to aim (Gepshtein, Seydell, & Trommershäuser, 2007; Trommershäuser, Maloney, & Landy, 2003), when to respond (Jarvstad et al., 2012) or to start moving (Battaglia & Schrater, 2007; Faisal & Wolpert, 2009), and how to adjust their on-going movement (Nashed, Crevecoeur, & Scott, 2014) and to what extent (Körding & Wolpert, 2004; Saunders & Knill, 2004, 2005). This has led to the idea that human movements are close to optimal (e.g. Diedrichsen, Shadmehr, & Ivry, 2010; Körding & Wolpert, 2006; Todorov, 2004; Todorov & Jordan, 2002; Trommershäuser, Maloney, & Landy, 2008). However, performance does not always appear to be optimal when subjects have to choose between two targets that they could move to (Jarvstad et al., 2014; Young, Pratt, & Chau, 2008). To evaluate why, we consider how people could know which target to choose.

There are various ways in which near-optimal choices between movements could theoretically be achieved. If the brain has a priori instantaneous access to all the relevant information, including information about the precision with which suitable motor commands can be generated and executed to move to any possible target, and the precise cost function associated with any resulting random errors, picking the best choice should be no problem. However, it is not evident that the brain *could* have such knowledge, given the infinite number of options, and considering that the consequences of motor commands are likely to depend on the circumstances and to change through practice and fatigue. Moreover, having a priori access to such knowledge is inconsistent with evidence that people make systematic errors when moving their unseen finger (Sousa, Brenner, & Smeets, 2011) or an invisible object (Brenner & van Damme, 1999; Smeets et al., 2006) to isolated visual targets in complete darkness, when indicating an object's position on a table by pointing with an unseen hand (van Beers, Sittig, & Denier van der Gon, 1999) or by marking the position with a pen (Schot et al., 2012) under the table, or when

* Corresponding author.
E-mail address: e.brenner@fbw.vu.nl (E. Brenner).

http://dx.doi.org/10.1016/j.visres.2015.03.028
0042-6989/© 2015 Elsevier Ltd. All rights reserved.

moving the hand to a position that is seen through a mirror so that the hand remains invisible (Kuling, Brenner, & Smeets, 2013). People also make consistent timing errors when trying to intercept moving targets with an invisible hand (de la Malla, López-Moliner, & Brenner, 2012), or even if the hand is visible but the timing with respect to the moving target is difficult to judge (Brenner, Canal-Bruland, & van Beers, 2013).

An alternative to having reliable a priori knowledge of all relevant factors is to rely on experience acquired while executing the task in question. For instance, for repeated movements to a single target, near-optimal performance could be achieved by shifting where one aims a little after every miss (Brenner & Smeets, 2011; van Beers, 2009). There are many examples of ways in which performance could be guided by feedback (Cusumano & Dingwell, 2013; Franklin et al., 2008; Krakauer et al., 2006) and the adjustments on subsequent trials depend on more aspects of the feedback than only the success (Schween et al., 2014), probably because the optimal way to respond to feedback is different for different tasks and circumstances (Liu & Todorov, 2007; Todorov, 2004). One way to identify the use of feedback-based strategies is by looking for serial dependencies within the data (Chaisanguanthum, Shen, & Sabes, 2014; Rabbitt, 1966; van Beers, Brenner, & Smeets, 2013). To optimize choices between targets, it would make sense to rely on one's success in hitting similar targets on previous trials.

It has been shown that choices can be adjusted after the movement has started if there is reason to do so (Nashed, Crevecoeur, & Scott, 2014; Voudouris, Smeets, & Brenner, 2013). A simple way to evaluate people's choices between two targets when under time pressure is therefore to examine whether they switch to hitting an easier target if it appears slightly later than the original one. By using a simple hitting task, and only two kinds of targets, we expected to be able to determine performance for each kind of target, and therefore to be able to determine what the optimal choice would be, as well as having enough repetitions of each target to be able to look for serial dependencies. The task was to hit as many targets as possible within 2 min. Whether or not it was advantageous to switch to the easier target depended on how much later it appeared. Subjects had to decide fast, because otherwise they will already have moved too far toward the original target. We found that the average performance was close to optimal. We identified some serial dependencies, but none that could account for our subjects' performance. Subjects' choices did depend on the precise progress of their movement at the moment that the easy target appeared. We conclude that the human brain has access to reasonably good estimates of how quickly and variably the arm can move, and relies on this information to rapidly guide the hand in choosing between possible movements.

2. Methods

2.1. Subjects and instructions

Two groups of ten young adult subjects took part in the experiment (Group A: 5 male, 5 female; Group B: 7 male, 3 female). The two groups performed slightly different sets of conditions, as described in the subsection *conditions*. The subjects were not aware of the purpose of the study, but were told that their goal was to tap with their index finger on as many targets as possible within several 2-min sessions. They were told that there would sometimes be more than one target, in which case it did not matter which they tapped. They were allowed to practice a bit before the experiments began. Each condition was presented in a separate session. There were short breaks between the sessions. The study was carried out in accordance with The Code of Ethics of the World Medical Association (Declaration of Helsinki).

2.2. Stimuli and task

Targets were presented on a big back-projection screen (Techplex 150; 1.25 m wide, 1 m high) that was tilted backwards by 30° (Fig. 1A). The targets were constrained to a central circular region with a diameter of 60 cm. They were projected from behind (InFocus DepthQ Projector; 1280 by 768 pixels; 120 Hz) in an otherwise normally illuminated room (fluorescent illumination). At any moment in time there were either one or two targets on the screen. Subjects stood in front of the screen and tried to tap on targets with their right index finger. The target(s) disappeared and at least one new one appeared as soon as a tap was detected. Subjects received feedback as to whether they hit the target (a tone sounded if their fingertip's position – as defined below – was within the target at the time of the tap). To motivate them to do their best they were shown their score (the number of targets that they had hit within the 2 min) on a high-score list after each session. They were free to stand and move as they pleased. They were warned that they had to lift their finger off the screen and tap on the targets, not to slide their finger across the screen (in which case no tap would be detected so they would not earn any points). They usually tapped the screen briefly so that they lifted their finger before they could have determined the position of the new target. We refer to the combination of one or two new targets appearing on the screen, and the finger moving toward a target and tapping on the screen, as a *trial*.

In the whole experiment there were only two kinds of targets: easy targets and hard targets. The easy targets were disks with a 2 cm radius that were 10 cm from the last target that the subjects tried to hit. The hard targets were disks with a 1 cm radius that

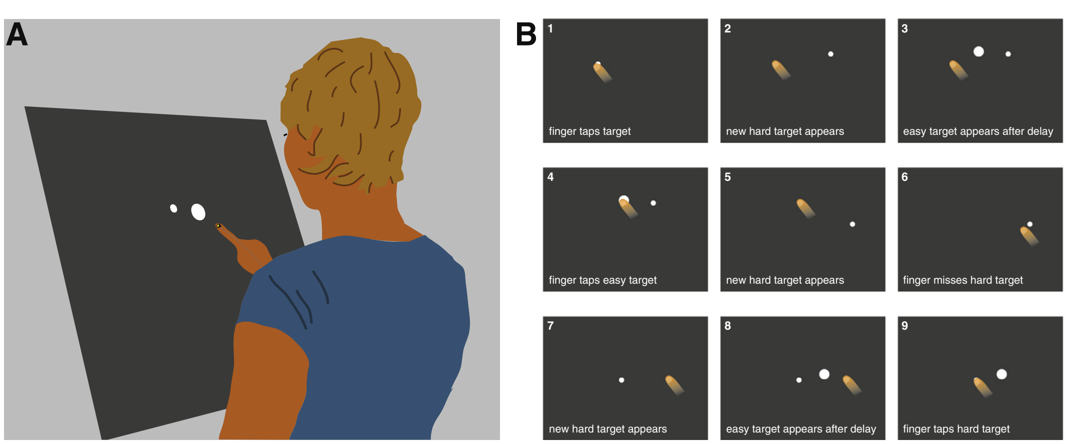

Fig. 1. The task. (A) The subject stood in front of a big screen and tapped as many targets as possible. (B) Possible sequence of events for four taps.

were 20 cm from the last target that the subjects tried to hit. At every moment during the 2-min sessions there could be an easy target on the screen, a hard target on the screen, or both. At least one target appeared as soon as a tap was detected (vertical acceleration threshold of 10 m/s^2 while the finger was less than 5 mm from the screen and more than half way to the target from the position of the previous tap). In some trials of some conditions, a hard target appeared as soon as a tap was detected, and an additional easy target appeared slightly later (Fig. 1B). The delay before such an easy target appeared was our main experimental variable.

The previous target or targets disappeared as soon as a tap was detected. If the finger hit a target, it was obvious that that was the target that had been chosen. If the finger missed both targets, we considered the target closest to the tap to be the one that had been chosen. When determining which was closest, we considered the distances from the targets' nearest edges, not the distances from their centers. New target positions were determined with respect to the center of the target that the subject hit or was trying to hit. The new targets could appear in any direction (with equal probability) except ones that would place either of them outside the above-mentioned 60 cm diameter central part of the screen. Thus, hard targets appeared 20 cm from the previous target that the subject hit or was trying to hit, and easy targets were presented 10 cm from the previous target that the subject hit or was trying to hit. If there were two targets, the easy target appeared close to the line connecting the hard target with the previous target that the subject (was trying to) hit. The easy target was equally likely to be 2.6 cm to either side of this line.

2.3. Measuring finger movements

A marker (an infrared light emitting diode) was attached to the nail of the subject's right index finger. Its position was determined at 500 Hz by an Optotrak 3020 that was placed at about shoulder height to the left of the screen. Before each session, the marker's position was measured when the subject placed his or her finger at four indicated positions on the screen. The relationship between the measured and indicated positions was used to later convert the measured position of the marker on the nail to an estimate of the coordinate of the fingertip with respect to the projected images. This calibration automatically corrected for the fact that the marker was attached to the nail rather than to the tip of the finger. The Optotrak also measured the position of a second marker that was attached to the left side of the screen, and that stopped emitting infrared light for about 10 ms, 1 ms after light fell on a sensor that was placed in the path of the light directed toward the top left corner of the screen. By presenting flashes at the top left corner of the screen whenever a new target appeared, we could tell when new targets appeared with the same 2 ms resolution with which we determined the position of the finger. The precision with which the Optotrak determined the position of the marker on the finger was better than 0.1 mm, but subjects could orient their fingers slightly differently on different trials, and touch the screen with slightly different parts of the finger on different trials, so the precision in determining the position of the fingertip is probably closer to 1 mm.

2.4. Conditions

There were four baseline conditions and eight test conditions. Each group performed three of the four baseline conditions and four or five of the eight test conditions. Both groups performed a baseline condition with only *easy targets* and a baseline condition with only *hard targets*. Group A also performed a baseline condition in which both targets always appeared simultaneously (*both targets*). Instead, group B performed a baseline condition in which the two kinds of targets were randomly interleaved, with only one target present at a time (*interleaved targets*). Each group also took part in a number of test conditions in which there was always a hard target, but 25% of the time an additional easy target appeared after some time delay. For group A the time delays were 0 ms, 50 ms, 100 ms and 150 ms. For group B they were 75 ms, 125 ms and 175 ms. In addition, both groups also performed one test condition (*interleaved delays*) in which delays of 0 ms, 50 ms, 100 ms and 150 ms were interleaved (on the 25% of trials in which a second, easy target appeared). Due to an error when setting up the experiment, group B did not receive the 150 ms delayed targets in the *interleaved delays* condition. Each participant performed the seven or eight conditions in a different random order.

2.5. Analysis

In order to determine whether our subjects chose the optimal target to tap, we had to determine what the best choice would be. To do so we determined the mean time that it took them to hit hard and easy targets (in each condition). This is not just the mean time between the first target appearing and the next tap, because we must also consider the likelihood of hitting the target. For the conditions with only one kind of target, we simply divided the total time (120 s) by the number of hits. For the other conditions, we summed the time between successive taps separately for instances (which we will call trials) in which the subject was trying to hit each kind of target, and divided this sum by the number of times that that kind of target was hit (as explained in Fig. 2). This was done separately for trials in which one target was present and trials in which two targets were present.

Following this procedure, the time taken per hit was always measured from the moment that the first target appeared, so the mean time taken to hit easy targets that appeared after a delay was longer if the delay was longer, even if the time taken from the moment the easy target appeared did not change. With the time taken per hit defined in this manner, the difference between the times for easy and hard targets is a direct measure of the advantage (or disadvantage) of choosing the easy target. We could have estimated the advantage of choosing the easy target from the first two baseline conditions (with only *easy targets* or only *hard targets*) and the delay, but to do so we would have to assume that the delay and the presence of the other target, either on the same trial or on previous trials, does not have any effect.

To determine whether the time taken to hit each kind of target depends on the simultaneous presence of the other target, on the target that was hit on the previous trial, or on the delay before the second (easy) target appeared, we compared the times taken to hit each kind of target across conditions. For the test conditions, we then used values from the session in question to determine the advantage of choosing the easy target (whenever there were enough trials to do so reliably), and related this to our subjects' actual choices. Note that this analysis determines whether subjects made the optimal choice given the way they tried to hit the two kinds of targets. It does not guarantee that they hit each kind of target optimally.

In order to understand how our subjects knew which choices to make, we examined to what extent their choices depended on their performance on previous trials, and to what extent they depended on random variability in the position of the finger when the second target appeared. The details of these analyses will be presented in the results section, when explaining the data. The general idea was to examine whether certain choices were more likely if they were more successful in the near past, and to examine whether the instantaneous circumstances were considered when choosing between targets.

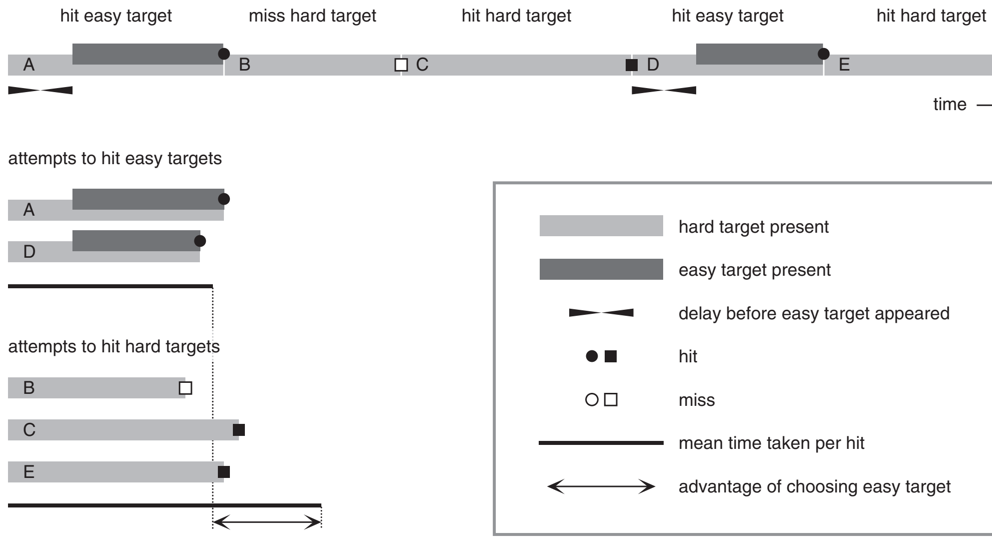

Fig. 2. Determining the advantage of choosing the easy target. For these five taps, the mean time taken per hit is $(T_A + T_D)/2$ for the easy targets and $(T_B + T_C + T_E)/2$ for the hard targets on their own (where T_i is the time between taps on trial i). In the latter case, the sum is divided by 2, rather than by 3, because only two of the three attempts to hit the hard target were successful (trials C and E). If subjects did not always pick the easy target when it was present, we also separately determined the mean time taken per hit when there were two targets and the hard target was chosen.

2.6. Statistics and exclusion of data

We compared our subjects' performance across conditions or targets with paired t-tests or repeated measures analyses of variance whenever we considered it meaningful to do so. We examined whether performance was different from the average performance on trials immediately after misses with one-sample t-tests. The specific tests are mentioned in the results section. In all cases we determined mean values per subject, and then tested for consistency across subjects.

In some conditions it was clearly advantageous for certain subjects to choose either the easy or the hard targets, so they seldom (or never) attempted to hit the other kind of target. The reported values for the mean time taken to hit a target are based on individual means of at least 20 attempts to hit that kind of target. We indicate when this meant that not all subjects contributed to a reported value. Values based on fewer than 20 attempts did contribute to our analysis of the extent to which each subject made the choices that we considered that he or she should have made. There was always enough data to estimate the time taken to hit hard targets, because there were many trials with only hard targets. When the correct choice was to hit the hard target, the time taken to hit easy targets was sometimes estimated from fewer than 20 values. We indicate when this was the case. When evaluating performance after misses we excluded estimates that were based on fewer than 3 instances.

3. Results

In the baseline condition with only easy targets, the mean time between taps was about 363 ms, and about 95.7% of the taps were on target. In the baseline condition with only hard targets, the mean time between taps was about 483 ms, and about 79.3% of the taps were on target. Dividing each subject's mean time between taps by his or her percentage of taps that were on target provides the mean time taken per hit (Table 1).

Table 1
Calculating the time taken per hit in the baseline conditions. Values are means ± standard errors of the subjects' individual (mean) values. For the easy alone and hard alone conditions these values are based on all 20 subjects' data.

Target	Condition	Mean time (ms)	Hit (%)	Time per hit (ms)
Easy	Easy targets	363 ± 6	95.7 ± 0.8	379 ± 5
Easy	Both targets	379 ± 12	94.2 ± 1.3	402 ± 11
Easy	Interleaved targets	398 ± 7	97.9 ± 0.5	407 ± 7
Hard	Hard targets	483 ± 7	79.3 ± 2.0	615 ± 16
Hard	Interleaved targets	492 ± 9	77.7 ± 2.7	640 ± 22

3.1. Choosing or switching between targets has some cost

Mean times taken per hit for various targets in various conditions are shown in Figs. 3 and 4. Fig. 3A shows the values for the easy and hard targets on their own, separately for the two groups of subjects, as well as for the both targets condition in which both targets were always shown simultaneously and the interleaved target condition in which the two kinds of targets were interleaved (but only one target was shown at a time). Obviously, subjects took considerably longer to hit hard targets. When both targets were shown on each trial subjects always chose to hit the easy target. Nevertheless, they took longer to hit the easy target when both targets were shown than when only the easy target was shown (making it a choice reaction time task rather than a simple reaction time task; difference of 23 ms; $t_9 = 3.5$; $p = 0.006$). Presumably the choice itself (identifying the easy target) took some time (Donders, 1868). It also took longer to hit the easy target when easy targets were interleaved with hard ones, than when all targets were easy (difference of 28 ms; $t_9 = 4.0$; $p = 0.003$). On average, it also took longer to hit the hard target when hard targets were interleaved with easy ones, but this effect was not consistent across subjects (difference of 19 ms; $t_9 = 0.8$; $p = 0.4$).

That it took longer to hit targets when they were interleaved, despite not having to choose between targets, shows that not knowing the size of the target in advance has some additional cost even when there is no choice to be made. Fig. 3B shows that a large part of this cost results from having to switch between targets on successive trials. A repeated measures analysis of variance for the baseline condition with interleaved targets, comparing mean times per hit for the factors current target (easy or hard) and previous target (easy or hard), revealed a significant interaction between the current and previous target type ($F_{1,9} = 19$, $p = 0.002$) as well as the obvious significant effect of the current target type ($F_{1,9} = 106$, $p < 0.0001$).

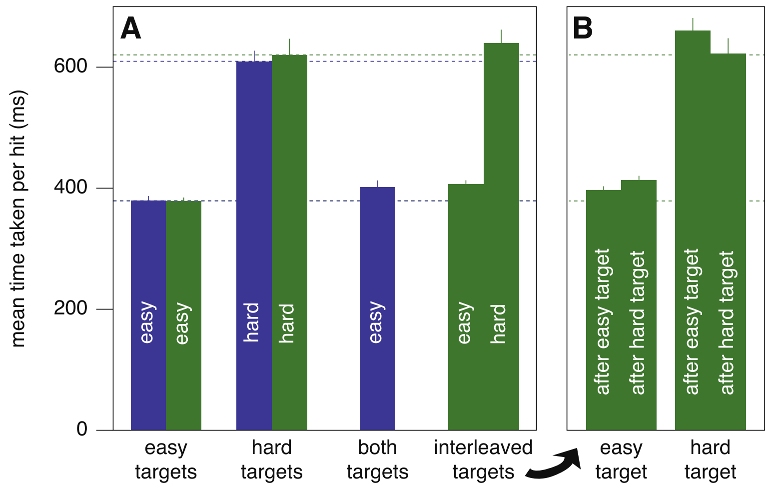

Fig. 3. Performance in the baseline conditions. (A) Overall performance of groups A (blue bars) and B (green bars). The dashed lines indicate the values for the easy and hard targets on their own (to make it easier to compare the conditions). (B) Performance for the interleaved targets split by previous trial type. Error bars are standard errors across subjects.

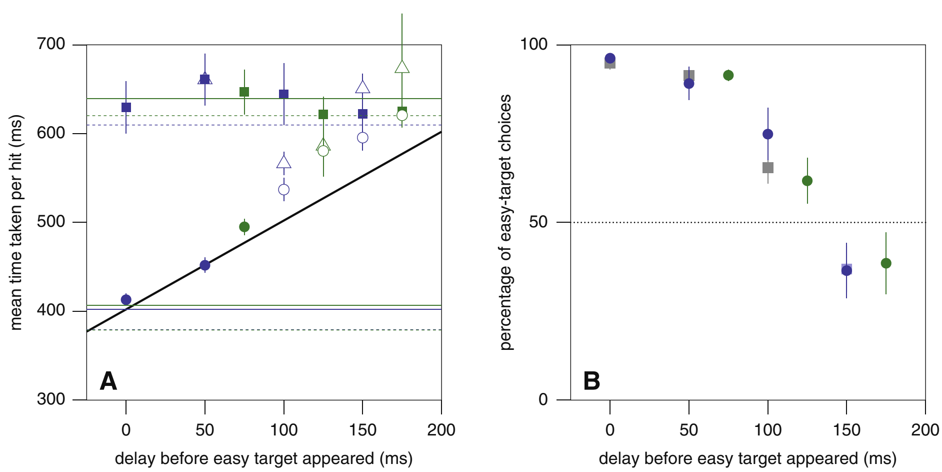

Fig. 4. Performance in the test conditions as a function of when the second, easy target appeared. Colors indicate the two groups of subjects (as in Fig. 3). Error bars are standard errors across subjects. (A) Mean time taken per hit. Circles: performance when both targets were presented and the easy target was chosen. Triangles: performance when both targets were presented and the hard target was chosen. Squares: performance when only the hard target was presented. Open symbols indicate that not all subjects chose the target in question at least 20 times. Dashed lines: time taken in the easy targets and hard targets baseline conditions (as in Fig. 3A). Solid horizontal lines: time taken in the other two baseline conditions. Black diagonal line: expected time taken if subjects switch as soon as the easy target appears (established by adding the delay to the average performance in the baseline condition with both targets). (B) Circles: choices between the targets. Squares: choices in the interleaved condition (gray when averaged across the two groups).

We also determined the time it took to hit easy and hard targets in the test conditions in which an easy target sometimes appeared after a delay (Fig. 4A). These times were determined from the moment the first, hard target appeared. The times were close to what one would expect on the basis of the baseline conditions, except that it took longer to hit the easy targets with longer delays than one might expect on the basis of when they appeared and how long it took to hit easy targets in the baseline conditions (circles above the thick black line for the longer delays). We will discuss this finding in more detail when we evaluate the finger's position at the time that the easy target appeared. When evaluating the advantage of choosing the easy target, we always considered the time it took to hit easy (and hard) targets in the condition in question, so all costs of switching at different times were automatically considered.

3.2. Did subjects make the correct choice?

For short delays, it clearly took less time to hit easy targets than to hit hard targets. Thus it was clearly advantageous to switch to the easy target if an easy target appeared. In accordance with this, subjects usually chose the easy target (Fig. 4B). For long delays, the mean time taken per hit was similar for easy and hard targets, and the two targets were chosen about equally frequently. The fractions of trials in which individual subjects chose to aim for the easy target globally follow the extent to which it was advantageous for those subjects to aim for the easy target (Fig. 5), but there was quite a lot of variability across subjects and conditions.

Even when our analysis suggests that it was clearly advantageous to switch, subjects did not always do so. This could be because some subjects were still learning to make optimal choices under the prevailing circumstances on the basis of feedback obtained during the session. However, it could also arise from random variability in when they started to move or in which direction, which could influence both the choice between the targets and the time taken to hit each. We first examine whether there is any evidence that our subjects' behavior is determined by some simple learning algorithm based on feedback on previous trials, and then turn to a possible influence of random variability across trials.

3.3. Responding to feedback: slowing down after misses

To examine the role of feedback, we first examined whether our subjects slowed down their movements immediately after a miss, as subjects had done in previous experiments (Brenner & Smeets, 2011; Rabbitt, 1966). We found a tendency to slow down after a miss for all targets in our baseline conditions (Fig. 6A), but the tendency was larger for the easy targets, and it was significant for all four easy targets (group A, only easy targets: $t_9 = 2.6$; $p = 0.03$; group B, only easy targets: $t_9 = 3.1$; $p = 0.01$; both targets present but easy target chosen: $t_8 = 2.5$; $p = 0.04$ [one subject had no misses]; easy target in condition with interleaved targets: $t_9 = 2.6$; $p = 0.03$) but not for any of the three hard ones. The increase in movement time after missing targets could imply that subjects initially systematically moved too fast, but if so their performance should have become better as time progressed. The mean number of targets hit within each 5 s period of the session, averaged across all subjects and sessions, was smaller for the very first seconds, but remained quite constant after that, with a slight decrease as the session progressed (at about 2% per minute) that may have been caused by fatigue (Fig. 7A). It was not caused by a systematic increase (or decrease) in the frequency with which subjects chose the easy target (Fig. 7B). Thus, our subjects' behavior did not change dramatically during the course of a session, but there were some subtle fluctuations, at least in movement speed.

3.4. Responding to feedback: choices do not depend on recent success

We next examined whether our subjects were more likely to switch to the second, easier target if they had missed the immediately preceding target (which was usually a hard target) than if they had not. The difference was not significant for any of the delays (seven paired *t*-tests; Fig. 6B). Rather than basing their choice on whether they had hit the preceding target, our subjects might have based their choice on whether they had hit the easy target last time they tried to. We found that they tried to hit the

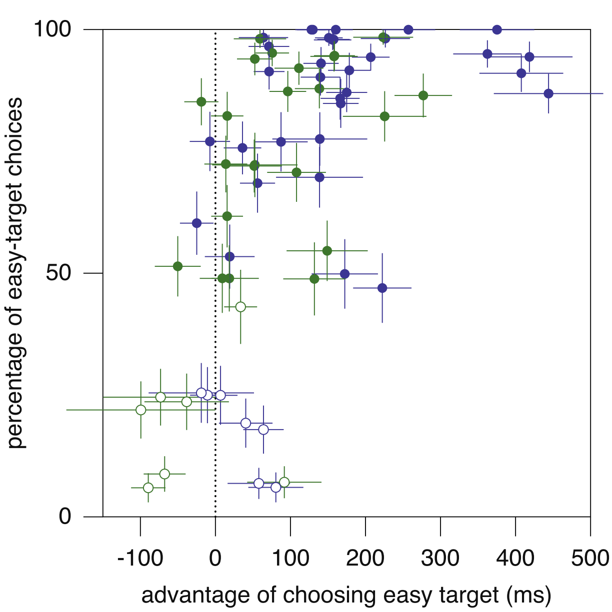

Fig. 5. How good is the choice in the test conditions? Fraction of trials in which subjects switched to hitting the easy target as a function of the calculated advantage of doing so. Each dot represents one subject's data for one condition. Open symbols indicate that the easy target was chosen fewer than 20 times (the lower the number of easy-target choices, the less reliable is our estimate of the advantage of choosing the easy target). Values are means and standard deviations of 10,000 random samples from the underlying data with the same number of taps per sample as in the original data (bootstrapping method).

easy target 82% of the time if they had hit the previous easy target, and 83% of the time if they had missed the previous easy target (averaged across delays; difference not significant). The only indication that having hit or missed a target might have influenced our subjects' subsequent choices was that if our subjects had *not* tried to hit the easy target last time they had the opportunity, they tried to hit the easy target 52% of the time if they had *missed* the hard target instead, but only 45% of the time if they had *hit* the hard target instead ($t_{19} = 2.0$; $p = 0.06$).

Perhaps we found no clear effect of whether our subjects hit or missed the previous target because our subjects also considered how long it took them to hit the target (the optimal movement time is one at which some targets are missed, so missing a target does not necessarily mean that one should choose a different target next time). Moreover, our subjects might not only have considered how they had fared the very last time they were given a choice. We might therefore expect to see our subjects choose the easy target more often when their preceding choices for the easy target were particularly fast and successful, or when their preceding choices for the hard target were particularly slow and unsuccessful. We would expect to see more choices for the hard target if their preceding choices for the hard target were particularly fast and successful, or if their preceding choices for the easy target were slow and unsuccessful.

To see whether this is the case, we subtracted the time taken to try to tap easy targets before an easy target was chosen, from the time taken to try to tap easy targets before a hard target was chosen. We also subtracted the time taken to try to tap hard targets before a hard target was chosen, from the time taken to try to tap hard targets before an easy target was chosen. In both cases we expect a positive value if subjects base their choice on the timing on previous trials. We considered the average of various numbers of preceding trials. The same reasoning was used to examine the influence of having missed the target on preceding trials. We only found any indication that our subjects were basing their choices on their performance on previous trials when the two targets appeared simultaneously (0 ms delay). In that case, we see a steady decline in the positive value of the *average difference in time* with the number of previous trials (Fig. 8A), in accordance with the directly preceding trial being fully responsible for the effect (the directly preceding trial is included in the average when considering more trials). For the 0 ms delay, we also see a tendency to pick the other target if a certain target was missed on the immediately preceding opportunity (Fig. 8B), but the difference was not significant (see 95% confidence intervals). Thus, altogether, feedback during the session appears to have very little direct influence on our subjects' choices.

3.5. Serial dependence: no tendency to stick to a choice

From the above analysis of our data it would appear that most of the differences between the choices that were made and between the movement speeds on different trials were due to random variability (the only exception being the systematic increase in movement time after missing a target). That the choices between targets varied at random during the sessions, rather than the choices usually being the same as the previous ones but interspersed with occasional switches, can be seen from the number of times that our subjects chose the same target as they had last time they were given a choice (Fig. 9A). The fractions of choices were very similar to what one would expect if subjects chose at random, given the number of choices they made for each of the targets (red horizontal lines; for an equal number of easy and hard target choices the fraction would be half the trials; if only the easy target was chosen this would obviously be all the trials). It is evident that our subjects did not have a strong tendency to stick to their choice.

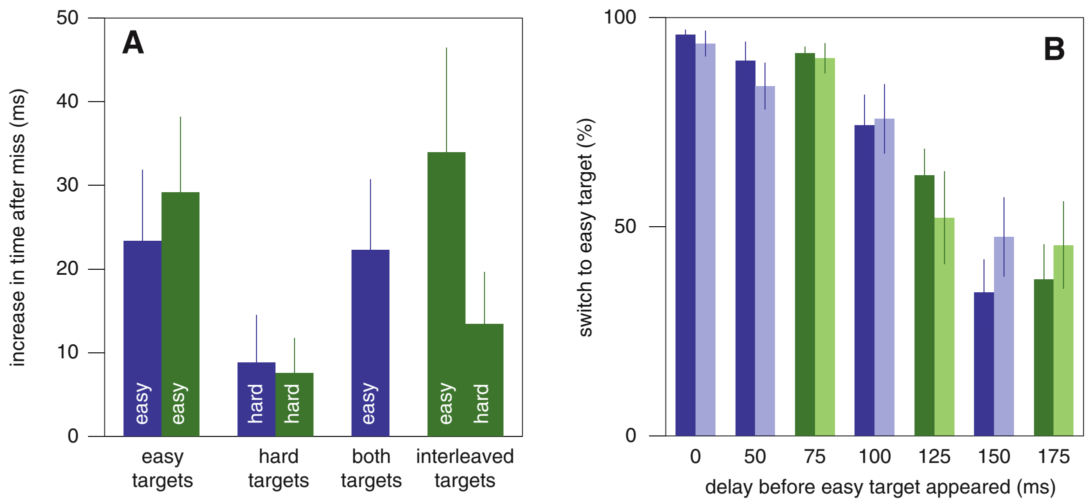

Fig. 6. Responses to errors. (A) Baseline conditions. How much longer the mean movement time is on trials immediately following a miss, than the overall mean movement time for that kind of trial. (B) Test conditions. How often subjects switched to the easy target when given a choice if they had hit the last target before the choice (dark bars) as compared to if they had missed the last target before the choice (bright bars). Error bars are standard errors across subjects.

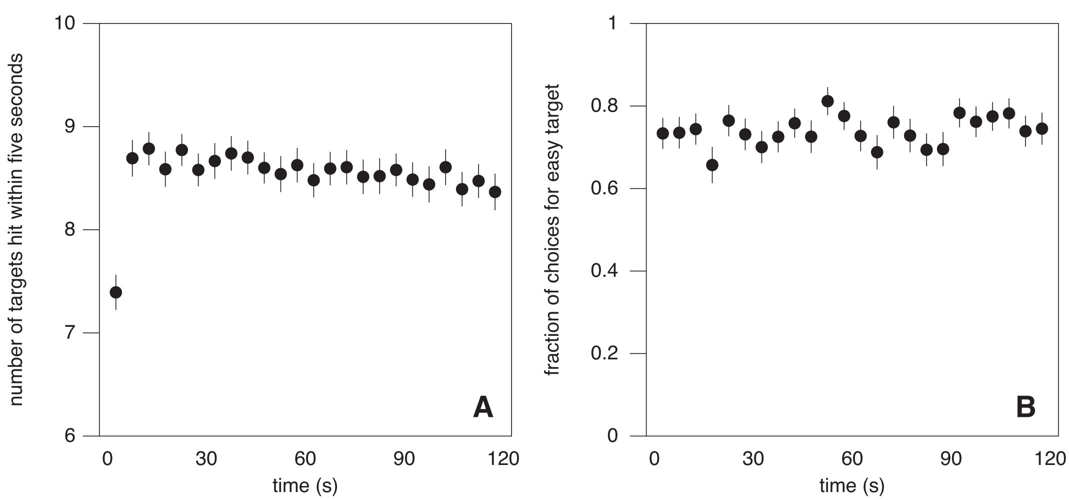

Fig. 7. Time course of performance during a session. Values are averages across all subjects and sessions within consecutive 5 s bins (with standard errors across subjects). (A) Mean number of targets hit. (B) Mean fraction of times that subjects tried to hit the easy target.

3.6. Modest effects of movement direction: target occluded by the arm

In the above analysis, we ignored the fact that the movement was in different directions on different trials. The fraction of trials in which the easy target was chosen depended on the direction of the movement ($F_{11,209} = 3.0$, $p = 0.0009$), but the differences were quite modest (comparisons using paired *t*-tests only revealed significant differences after Bonferroni correction between the two directions to the lower right, 300° and 330°, and between movements to the left and slightly upward and ones to the right and slightly downward, 150° and 330°; Fig. 9B). The time it took to hit targets also depended on the movement direction ($F_{11,209} = 8.6$, $p < 0.0001$), as well as obviously on the kind of target (easy or hard; $F_{1,19} = 206$, $p < 0.0001$). There was a significant interaction ($F_{11,209} = 2.0$, $p = 0.03$), perhaps because it took particularly long to hit hard targets when this meant moving downwards and to the right. It took longer to hit both kinds of targets when doing so involved moving downwards and to the right (pairwise comparisons revealed significant differences between this angle of 300° and all angles except 120°, and between the angle of 330° and both 30° and 60°; Fig. 9C). All these effects are probably the result of the target sometimes being occluded by the arm when presented at the lower right.

3.7. Random variability in the initial part of the movement

Another aspect that differs across trials is the position and motion of the finger at the moment that it could be redirected to a different target. The finger's initial distance from the new targets differed slightly across trials because new targets were placed with respect to the target that the subject was trying to hit, not with respect to where the finger hit the screen. More importantly, the finger tapped the screen briefly, moving away from the screen at about the time at which the new target was presented (shown in Fig. 10A for the *easy targets* condition, for which movements are the fastest), whereas it only started moving consistently toward the targets about 120 ms after the targets appeared (Fig. 10B). Such motion that is not yet directed toward the targets introduces additional variability in the finger's position and motion with respect to the targets.

The most obvious way in which random variability in the position of the finger might influence subjects' choices between the

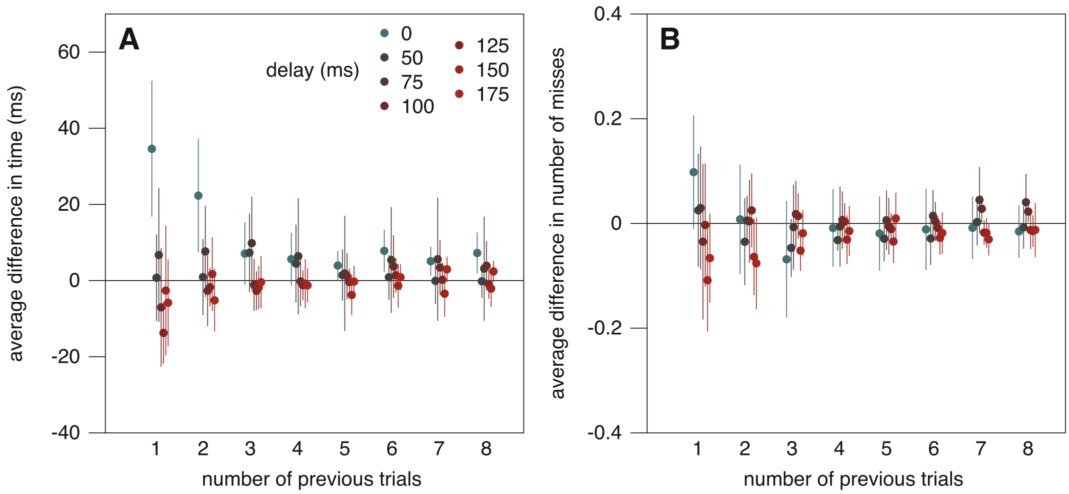

Fig. 8. Extent to which the difference in performance on trials *preceding* attempts to hit easy and hard targets is consistent with the choice that was made (averaged across easy and hard targets, and then across subjects, only considering differences between average values that are based on at least three instances). Points show average values for each delay before the easy target appeared, considering various numbers of previous trials, with 95% confidence intervals across subjects. (A) Difference between the latest times taken to try to hit the two kinds of targets. (B) Difference between the latest number of misses for the two kinds of targets.

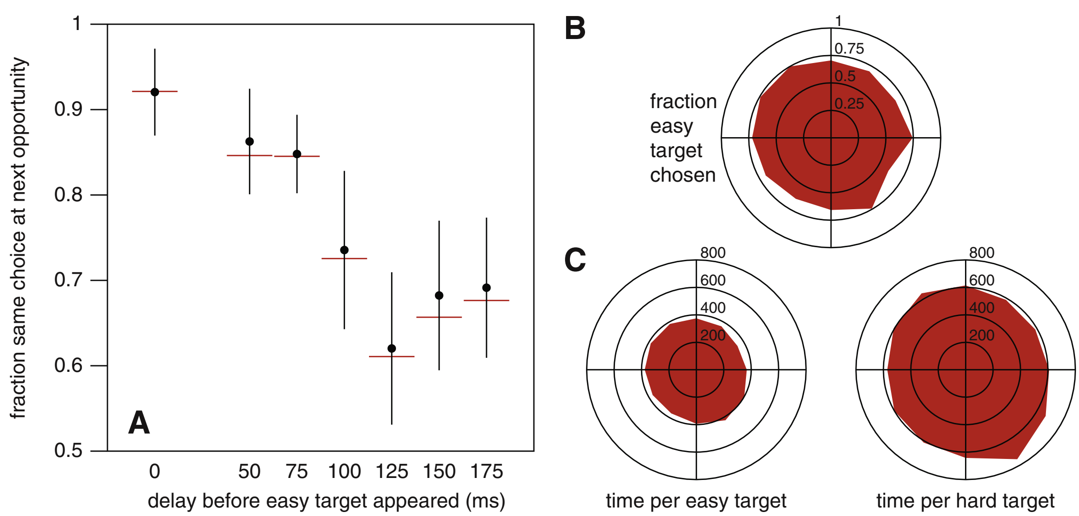

Fig. 9. Possible biases. (A) Extent to which subjects were biased to make the same choice as last time (mean and 95% confidence interval across subjects). The red horizontal lines indicate what one would expect for independent choices on successive trials (given the actual frequencies with which the choices were made). (B) Extent to which the choice depended on the direction of the required movement (30° bins). Data averaged across all delays. (C) Extent to which the mean time taken per hit depended on the direction of the required movement (30° bins). Data for the baseline conditions with only a single type of target. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

two targets is that once the finger has moved substantially toward the hard target it might be less attractive to switch to the easy target. Fig. 11A shows that on the trials in which the easy target was chosen, the finger was still about 20 cm from the hard target when the easy target appeared (dark bars). The finger was slightly closer to the hard target at that time on the trials in which the hard target was chosen (bright bars). The difference was significant for the four largest delays (100 ms: $t_8 = 4.4$, $p = 0.002$; 125 ms: $t_9 = 2.5$, $p = 0.03$; 150 ms: $t_9 = 4.0$, $p = 0.003$; 175 ms: $t_9 = 6.5$, $p = 0.0001$) but not for the other three.

Fig. 11B shows that the finger had only moved slightly more by the time the second target appeared in trials in which the hard target was chosen (not significant for any of the delays), so how far the finger had moved is less important than whether it had moved closer to the hard target. Since the two new targets always appeared in about the same direction from the previous tap, it is not surprising to see that the finger was also closer to the easy target on trials in which subjects tried to hit the hard target (Fig. 11C). Again, the differences were significant for the largest four delays (100 ms: $t_8 = 4.4$, $p = 0.002$; 125 ms: $t_9 = 2.8$, $p = 0.02$; 150 ms: $t_9 = 2.3$, $p = 0.05$; 175 ms: $t_9 = 5.4$, $p = 0.0004$) but not for the other three. Subjects were presumably inclined to switch to the easy target when there was enough time to adjust the finger's movement before it passed the easy target.

3.8. Support for considering the instantaneous circumstances

The fact that the finger's position when the easy target appeared influenced the choice between the two targets suggests that subjects identify the target that they can best aim for on the fly. Not only do they readily switch between targets when a new, easier target appears, but they consider the position of their finger with respect to the targets when doing so. Considering the position of the finger could explain some of the apparent deviations from

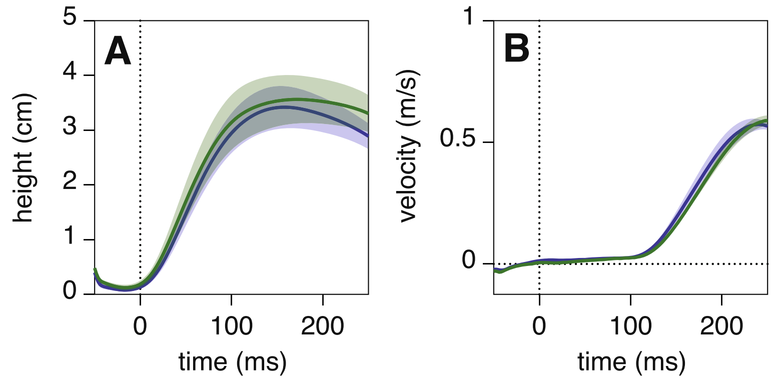

Fig. 10. Kinematics of movements in the easy targets baseline condition. The finger's height (distance from the screen; (A) and velocity in the direction of the target (B) from the moment the target appeared. Curves show means (with between-subject standard errors) for each group (color coding as before). Note that zero on the horizontal axis is the time the new target was presented, not the time the finger tapped the screen. It obviously takes some time to detect the taps and present new targets on the screen. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

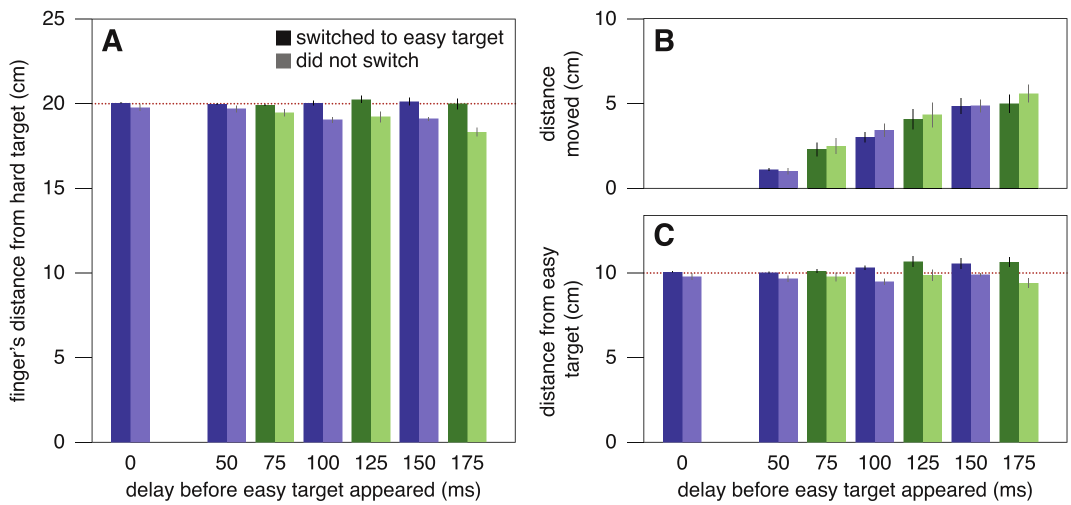

Fig. 11. The finger's position when the easy target appeared (test condition trials in which both targets were presented). Comparison between trials in which subjects tried to hit the easy (dark bars) or hard (bright bars) target. Average three-dimensional distances, with standard errors across subjects. (A) Distance from the hard target. (B) Distance from the finger's position when the hard target appeared (so the values are zero by definition when there is no delay). (C) Distance from the easy target. That this distance is sometimes larger than the original distance to the target (dotted red line) does not mean that subjects initially moved in the wrong direction, because the distance also increases when subjects lift their finger. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

optimal behavior in Fig. 5, because we calculated the advantage of choosing the easy target on the basis of performance in that session without considering that it may, for instance, have systematically taken longer to reach the easy target on trials in which the hard target was chosen, than it actually took on the trials in which the easy target was chosen. To further examine how our subjects' choices depended on the position and motion of their finger near the time at which they made their choice, we examined the position of the finger during the 300 ms after the easy target appeared. Again, we compared trials in which subjects did and did not switch to the easy target. We did so for each delay separately, by simply separately averaging all trials of all subjects in which the easy or hard target was chosen. We only considered trials in which both targets were presented. The position of the finger was determined in three directions: toward the hard target from the position of the previous tap, in the orthogonal direction on the screen (whereby deviations in the direction of the easy target were considered to be positive), and orthogonal to the screen.

There are clear systematic differences in the position and motion of the finger between trials in which the easy and hard target was chosen (Fig. 12). The finger had moved further (thin curves below thick ones in the left column) and was moving faster (steeper slope of thin curves) on trials in which the hard target was chosen. For the longer delays, the easy target also appeared to be more likely to be chosen if the finger's deviation from a straight path to the hard target was in the correct direction (middle column). From the deviations toward the easy target it would appear that it took a bit less than 200 ms to respond to the new target when one chose to move toward that target. For the longer delays, the finger had often already crossed the easy target (dotted lines in left column) by that time, so it is evident that the optimal choice will depend on the position and motion of the finger. Considering that our subjects' choices already appeared to be close to optimal before we demonstrated that they were considering the precise circumstances (Fig. 5), and that our subjects considered the circumstances in a reasonable manner (Fig. 12), it would appear that appropriate choices can be made very quickly and efficiently.

3.9. Interleaving delays is not detrimental

If people really consider the circumstances on each trial, it should not matter whether the different delays are presented in separate sessions, as they were in the conditions that we have been discussing until now, or whether several delays are interleaved

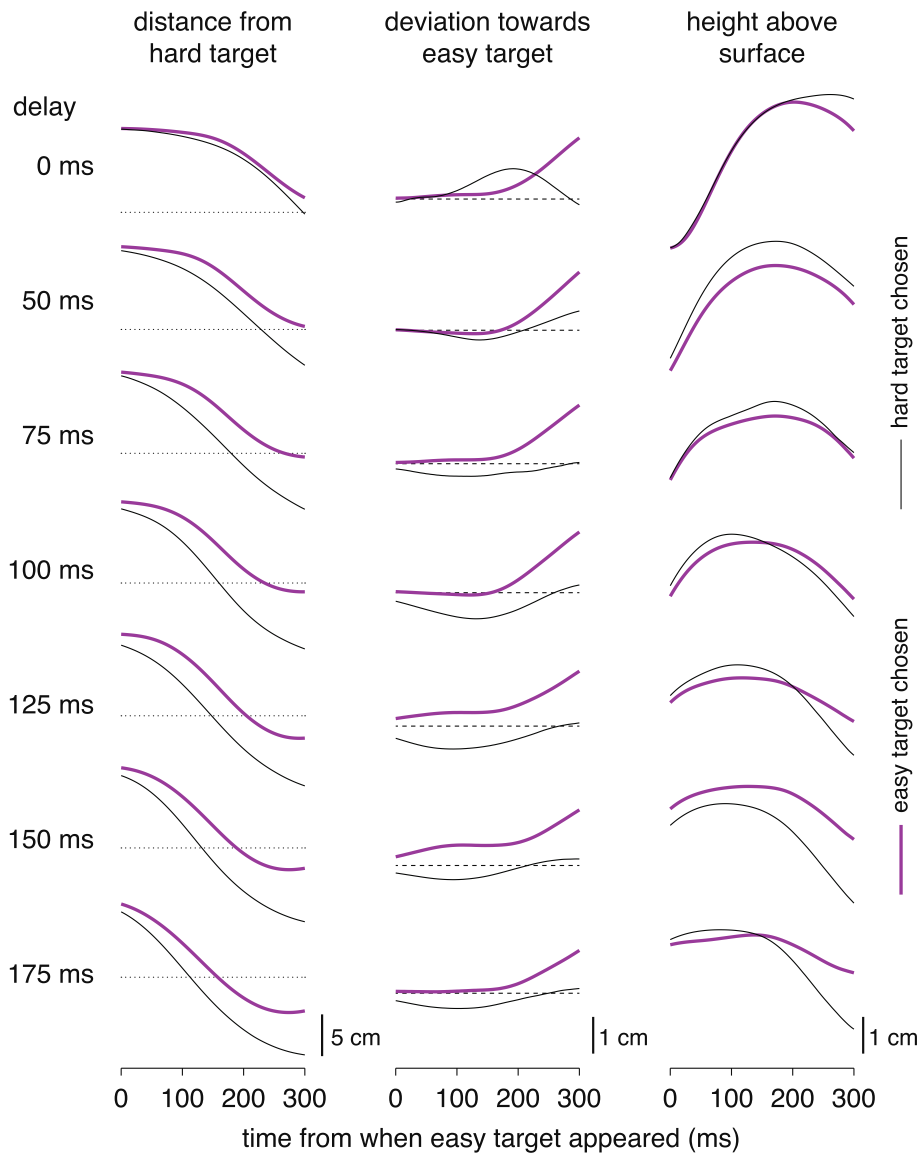

Fig. 12. The position of the finger on trials in which the easy (thick purple curves) and hard (thin black curves) target was chosen, as a function of the time from when the easy target appeared. The columns represent the three different directions (from left to right: along a line connecting the previous tap with the current hard target, in the orthogonal direction on the screen, and orthogonal to the screen). The rows represent the seven test conditions, with different delays between the hard and easy target appearing. The dotted lines in the left column indicate a distance from the target of 10 cm (showing about when the finger crossed the easy target). The dashed lines in the middle column indicate a deviation of zero (meaning that the finger is moving straight toward the hard target). Positive values in the middle column indicate a deviation toward the easy target. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

within one session. Our subjects' average choices were indeed almost identical in both cases for the delays of 0, 50 and 150 ms (see squares in Fig. 4B). On average, they tended to choose the easy target slightly less often when the targets were interleaved when the delay was 100 ms. Individual subjects were quite consistent in their choices under certain conditions (correlation of $r = 0.87$ between fractions of choices for the easy target when each delay was presented in a separate session and when all delays were interleaved; subjects of group A).

4. Discussion

On average, our subjects chose the target that would give them the best score (allow them to hit most targets; Fig. 5). Our subjects' choices did not rely heavily on recent feedback: they did not pick a target particularly often after having missed the other target (Figs. 6B and 8B) or after having been particularly successful when aiming for that target in preceding trials (Fig. 8). Moreover, performance hardly changed when delays were interleaved, although interleaving delays makes it impossible to anticipate in advance how one should respond to the next choice (Fig. 4B). Neither did performance improve systematically during the sessions (Fig. 7). Thus, our subjects did not appear to rely strongly on fine-tuning their choices on the basis of their success.

Our subjects' choices did depend on the position and motion of the finger at the time that the easy target appeared (Fig. 12). The position and motion differed across trials because the finger started moving before it could have been directed toward the first target (Fig. 10A). Such motion was primarily 'upwards', away from the screen, but there was also a component parallel to the screen, in a random direction with respect to where the hard target appeared. This random component, and random variability in the latency to respond to the target appearing and in the movement speed, gave rise to considerable variability in the position of the finger when the easy target appeared. Additional variability was introduced by slowing down after missing a target (Fig. 6A) and by the target sometimes temporarily being hidden by the arm (Fig. 9C). Our subjects considered the consequences of such variability when choosing between the two targets (Figs. 11 and 12).

In order to select the best target considering the circumstances, our subjects must have had access to quite reliable estimates of the time it would take them to reach the targets and of the likelihoods of missing the targets. At first glance, this appears to be inconsistent with slowing down after missing targets (Fig. 6A) and with aiming points constantly being adjusted to feedback about one's errors (e.g. van Beers, 2009). However, people may have quite good estimates of how they would fare, but nevertheless fine-tune them on the basis of new information whenever possible. We may not have found any evidence for fine-tuning the basis of our subjects' choices in response to feedback in this study because the effects of any improvements in judging the time it takes to reach the targets or in judging the likelihoods of missing them under specific circumstances were masked by the effects of random variability in the circumstances (such as the position of the finger when the easy target appeared).

4.1. Are the movements themselves optimal?

If subjects have good estimates of the time it will take them to reach the targets, and of the likelihoods of missing the targets, we would not only expect their choice between the targets to be near-optimal, but also the movement to the chosen target should be near-optimal. Although our subjects' choices matched their performance when aiming for the different targets, it is not certain that their performance for the individual targets was optimal. There is ample evidence that people are close to optimal in determining where to aim. This may be because they know their own capabilities (also see Trommershäuser, Maloney, & Landy, 2003), but near-optimal performance could be achieved by adjusting the movement plan slightly after every error (Brenner & Smeets, 2011), possibly considering the magnitude of the error and its origin (van Beers, 2009). We therefore conducted a control experiment to evaluate whether subjects chose optimal movement times in our task.

5. Control experiment

We repeated the baseline conditions in which there was a single target for five different target sizes. Varying the target size should make subjects move at different speeds, so we can determine the relationship between movement time and endpoint variability (we expect a linear relationship between movement time and the logarithm of the standard deviation in the errors; Plamondon & Alimi, 1997; Welford, 1960). Knowing this relationship, we can determine the optimal movement time for any target size. Beside examining whether our subjects' performance was near-optimal, we also examined whether they relied on an existing estimate of their movement capabilities, or whether they adjusted their movement time in response to feedback, because the results for the original baseline condition were not quite clear about this (subjects only responded significantly to feedback for the easy targets; Fig. 6A).

Fourteen subjects who had not taken part in the original experiment took part in the control experiment. They were told that their goal was to tap with their index finger on as many targets as possible within several separate 90-s sessions. The target was always at a distance of 20 cm from the previous target. Its radius could be 1, 1.5, 2, 3 or 4.5 cm (values chosen on the basis of pilot data to achieve a large range of movement times). Each subject was presented with the targets of different sizes in a different order. The same order was repeated twice, so that each subject tried to hit targets of each size for twice 90 s. One subject did not systematically move faster when the target was larger, so her data were not included in the further analysis. The setup was identical to that of the main experiment.

The movements were in random directions. Since endpoint variability in the direction of the movement is likely to differ from that in the orthogonal direction (Gepshtein, Seydell, & Trommershäuser, 2007; van Beers et al., 1999), we determined the standard deviations separately in these two directions (considering undershooting the target and counter-clockwise errors as positive). Increasing the mean movement time decreases both these measures of endpoint variability (shown for one subject in Fig. 13A). The fit lines allow us to estimate the two standard deviations for any movement time. Knowing a subject's standard deviations and his or her bias (average biases ranged from no bias for the smallest targets to undershooting the target centre by 9 mm for the largest targets), we can calculate the fraction of attempts that are likely to be successful for any movement time and target size (the part of a two-dimensional Gaussian with these standard deviations that is within the target radius). Dividing each movement time by the corresponding fraction gives the expected mean time taken per hit (curves in Fig. 13B; each for a different target size). The movement time for which the mean time taken per hit is minimal, for a given target size, is the optimal movement time.

For long movement times, almost all attempts to hit the targets are successful, irrespective of the target's size, so the time taken per hit is simply the movement time. For short movement times, many attempts to hit the target are unsuccessful, especially if the target is small, so the time taken per hit increases again as the movements become faster. The precise shapes of the curves differ between subjects, because they depend on the relationship between movement time and endpoint variability, and slightly on the magnitude of the bias (subjects probably undershoot large targets' centers, despite this slightly decreasing the fraction of successful hits, because moving less far reduces the movement time), but the general pattern is the same for all subjects.

The points in Fig. 13B show this subject's actual mean movement time and mean time taken per hit for each target size. He hit slightly fewer targets than we calculated that he should have hit for his actual mean movement times (points above curves), suggesting that it is probably not quite true that the errors are distributed according to a two-dimensional Gaussian. More importantly, his mean movement times were not precisely aligned with the troughs in the curves. He moved more slowly than we would consider optimal for the large targets and faster than we would consider optimal for the small targets. This tendency is consistent across subjects, although the mean movement times were quite close to the optimal ones (Fig. 13C).

The systematic differences between the actual and optimal movement times might at least partly arise from errors in calculating the optimal times as a result of us imposing a linear relationship between movement time and log endpoint variability for movement times of 250 ms and more, which is probably not entirely justified. Although most deviations from this linear relationship are modest, only one of the 26 standard deviations for the largest target (smallest movement time) was below the fit line, so the disadvantage of moving very fast is larger than we assumed. Consequently, the advantage of moving very slowly is probably also less large than we assumed. Thus the optimal movement speed is probably actually a bit slower for large targets and faster for small targets, in accordance with our subjects' performance.

To examine whether the differences between the optimal movement times that we calculated for the different subjects can account for some of the variability between their actual movement times, we examined whether the optimal and actual movement

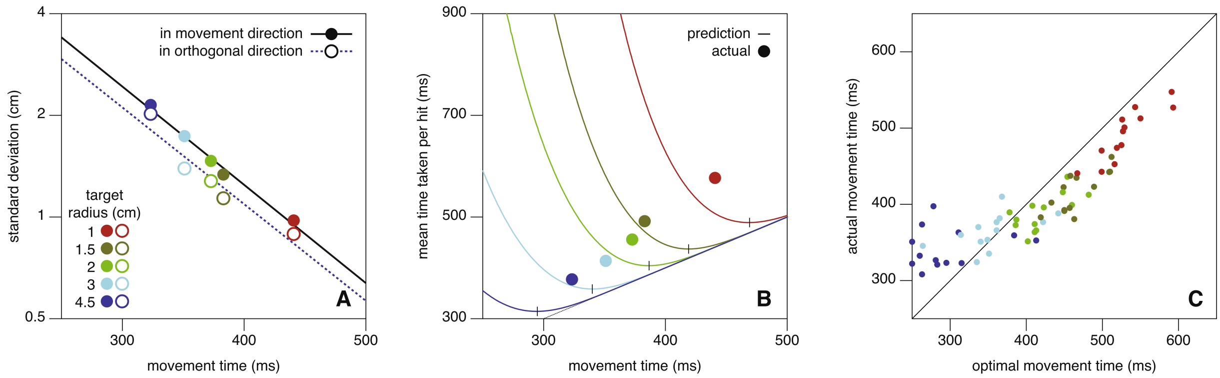

Fig. 13. Is our subjects' movement time optimal? (A) The relationship between movement time and (log) endpoint variability for one subject of the control experiment, both for variability in the direction of the movement and for variability in the orthogonal direction. The colors indicate the target size. (B) How this subject's mean time taken per hit is expected to depend on the movement time for each target size considering the fit lines in A (curves). Small vertical lines indicate the positions on the curves for which the movement times were optimal. The dots show the subject's actual mean movement time and time taken per hit. (C) The actual movement times of all subjects for each target size, as a function of their optimal movement times (the calculated optimal movement times for the two leftmost points were unrealistically small, so they were assigned a value of 250 ms). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

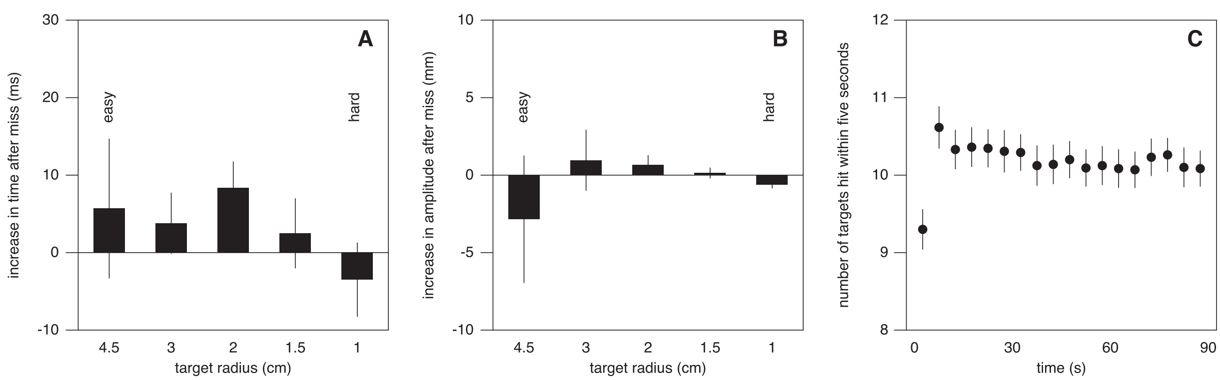

Fig. 14. Changes during the control experiment. How much longer the mean movement time (A) and amplitude (B) is on trials immediately following a miss, than the overall mean movement time or amplitude for that kind of trial (with standard errors across subjects). (C) Mean number of targets hit within consecutive 5 s bins during a session (averaged across all subjects, target sizes and replications, with standard errors across subjects).

times were correlated across subjects. We did this separately for each target size (points of each color in Fig. 13C). The correlations were $r = 0.20$ ($t_{11} = 0.7$, $p = 0.5$) for a target radius of 4.5 cm, $r = 0.55$ ($t_{11} = 2.2$, $p = 0.05$) for a target radius of 3 cm, $r = 0.68$ ($t_{11} = 3.0$, $p = 0.01$) for a target radius of 2 cm, $r = 0.77$ ($t_{11} = 4.0$, $p = 0.002$) for a target radius of 1.5 cm, and $r = 0.88$ ($t_{11} = 6.2$, $p < 0.0001$) for a target radius of 1 cm. Thus, subjects adjusted their movement times to their own capabilities, despite the systematic deviations from what we estimated to be optimal. This was least evident for the largest target (for which two of our calculated optimal movement times were unrealistically short), possibly because endpoint variability is no longer the main factor to limit performance for extremely fast movements (energetic considerations, pain and fatigue may also be considered).

In the baseline conditions of the main experiment we saw that subjects increased their movement time after missing targets (Fig. 6A). In this experiment the movement times also appeared to increase after having missed targets (Fig. 14A), but the increase was only significant (one-tailed t-tests) when the target radius was 2 cm ($t_{12} = 2.4$; $p = 0.02$). As in the main experiment, the magnitude of the increase in movement time appeared to be larger for larger targets, but the differences were not significant ($F_{4,48} = 0.9$, $p = 0.5$). The magnitude of the increase in movement time would be larger after missing large targets if it were the result of reducing the tendency to undershoot the centre of the target after a miss (making the movement longer; our subjects mainly undershot the centre of large targets), but we found no evidence that the movement amplitude increased after a target was missed for any target radius (one-tailed t-tests), or that such an increase depended on the target's size ($F_{4,48} = 0.6$, $p = 0.7$; Fig. 14B). Finally, as in the main experiment (Fig. 7A), we see no systematic improvement in performance during the session (Fig. 14C).

Averaged across all subjects and target sizes, our calculations suggest that our subjects hit about 88% of the number of targets that they could have hit. We already mentioned several reasons why this value may be underestimating our subjects' ability to optimize their movements, both of which question the assumptions that we made in order to estimate optimal performance (Gaussian distribution of errors; linear relationship between log standard deviation of errors and movement time). In our task, the number of misses that should be tolerated depends on how much shorter the movement time is in return, because the goal is to hit as many targets as possible within a fixed amount of time. If one moves faster one will attempt to hit more targets, but will miss a higher fraction of them. Previous studies suggesting that movement speed is close to optimal (Harris & Wolpert, 1998; Tanaka, Krakauer, & Qian, 2006) are more difficult to interpret, because it is unclear to what extent one or two misses can be tolerated in order to adhere to the request to move as fast as possible. Jarvstad et al. (2014) tried to circumvent this issue by limiting the available time for each movement, rather than asking subjects to be as fast as possible. They reasoned that the optimal behavior in this case is to maximize endpoint accuracy by making full use of the available time, irrespective of target size. Their subjects still made faster and less precise movements for larger targets, which they interpreted as evidence that their subjects' performance was not optimal. However, although moving too slowly was not penalized, other than that the trial was repeated, the subjects probably experienced arriving too late as a failure. They may therefore have moved faster for larger targets because doing so made them more certain to arrive in time, with a minimal increase in the probability of missing the target. This kind of reasoning is consistent with externally imposed time limits even influencing performance when it is the preferred timing that is imposed (Zhang, Wu, & Maloney, 2010).

6. Conclusions

Our results suggest that subjects have good estimates of their movement capabilities and use them efficiently, both to optimize their movement time within a task (control experiment, Fig. 13C), and – most importantly – when choosing between two actions. We found that subjects made reasonable choices (Figs. 4 and 5), and that they considered the precise circumstances when doing so (Figs. 11 and 12) rather than basing their decisions on their success on previous trials (Figs. 6B and 8). Subjects do presumably fine-tune their estimates on the basis of feedback (Figs. 6A and 14A), but their initial estimates are good enough to make near-optimal choices between the targets. Consequently, performance does not improve during the session (except at the very start, which is explained by the very first movement taking exceptionally long; Figs. 7A and 14C).

Acknowledgments

We thank Stefan Hagebeek, Robia Alphonse, Wietse Krab and Bob Koomen for collecting the data and assisting with the preliminary analysis.

References

Battaglia, P. W., & Schrater, P. R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *Journal of Neuroscience, 27*, 6984–6994.
Brenner, E., Canal-Bruland, R., & van Beers, R. J. (2013). How the required precision influences the way we intercept a moving object. *Experimental Brain Research, 230*, 207–218.
Brenner, E., & Smeets, J. B. J. (2011). Quickly 'learning' to move optimally. *Experimental Brain Research, 213*, 153–161.
Brenner, E., & van Damme, W. J. (1999). Perceived distance, shape and size. *Vision Research, 39*, 975–986.
Chaisanguanthum, K. S., Shen, H. H., & Sabes, P. N. (2014). Motor variability arises from a slow random walk in neural state. *Journal of Neuroscience, 34*, 12071–12080.
Cusumano, J. P., & Dingwell, J. B. (2013). Movement variability near goal equivalent manifolds: Fluctuations, control, and model-based analysis. *Human Movement Science, 32*, 899–923.
de la Malla, C., López-Moliner, J., & Brenner, E. (2012). Seeing the last part of a hitting movement is enough to adapt to a temporal delay. *Journal of Vision, 12(10)*(4), 1–15.
Diedrichsen, J., Shadmehr, R., & Ivry, R. B. (2010). The coordination of movement: Optimal feedback control and beyond. *Trends in Cognitive Science, 14*, 31–39.
Donders, F. C. (1868). Over de snelheid van psychische processen. Onderzoekingen gedaan in het Physiologisch Laboratorium der Utrechtsche Hoogeschool, 1868–1869. *Tweede reeks, II*, 92–120.
Faisal, A. A., & Wolpert, D. M. (2009). Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. *Journal of Neurophysiology, 101*, 1901–1912.
Franklin, D. W., Burdet, E., Tee, K. P., Osu, R., Chew, C. M., Milner, T. E., et al. (2008). CNS learns stable, accurate, and efficient movements using a simple algorithm. *Journal of Neuroscience, 28*, 11165–11173.
Gepshtein, S., Seydell, A., & Trommershäuser, J. (2007). Optimality of human movement under natural variations of visual-motor uncertainty. *Journal of Vision, 7*(5), 13.1–18.
Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature, 394*, 780–784.
Jarvstad, A., Hahn, U., Warren, P. A., & Rushton, S. K. (2014). Are perceptuo-motor decisions really more optimal than cognitive decisions? *Cognition, 130*, 397–416.
Jarvstad, A., Rushton, S. K., Warren, P. A., & Hahn, U. (2012). Knowing when to move on: Cognitive and perceptual decisions in time. *Psychological Science, 23*, 589–597.
Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*, 244–247.
Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Science, 10*, 319–326.
Krakauer, J. W., Mazzoni, P., Ghazizadeh, A., Ravindran, R., & Shadmehr, R. (2006). Generalization of motor learning depends on the history of prior action. *PLoS Biology, 4*, e316.
Kuling, I. A., Brenner, E., & Smeets, J. B. J. (2013). Proprioception is robust under external forces. *PLoS One, 8*, e74236.
Liu, D., & Todorov, E. (2007). Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *Journal of Neuroscience, 27*, 9354–9368.
Nashed, J. Y., Crevecoeur, F., & Scott, S. H. (2014). Rapid online selection between multiple motor plans. *Journal of Neuroscience, 34*, 1769–1780.
Plamondon, R., & Alimi, A. M. (1997). Speed/accuracy trade-offs in target-directed movements. *Behavioral and Brain Science, 20*, 279–303.
Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology, 71*, 264–272.
Saunders, J. A., & Knill, D. C. (2004). Visual feedback control of hand movements. *Journal of Neuroscience, 24*, 3223–3234.
Saunders, J. A., & Knill, D. C. (2005). Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements. *Experimental Brain Research, 162*, 458–473.
Schot, W. D., Brenner, E., Sousa, R., & Smeets, J. B. J. (2012). Are people adapted to their own glasses? *Perception, 41*, 991–993.
Schween, R., Taube, W., Gollhofer, A., & Leukel, C. (2014). Online and post-trial feedback differentially affect implicit adaptation to a visuomotor rotation. *Experimental Brain Research, 232*, 3007–3013.
Smeets, J. B. J., van den Dobbelsteen, J. J., de Grave, D. D. J., van Beers, R. J., & Brenner, E. (2006). Sensory integration does not lead to sensory calibration. *Proceedings of the National academy of Sciences of the United States of America, 103*, 18781–18786.
Sousa, R., Brenner, E., & Smeets, J. B. J. (2011). Judging an unfamiliar object's distance from its retinal image size. *Journal of Vision, 11(9)*(10), 1–6.
Tanaka, H., Krakauer, J. W., & Qian, N. (2006). An optimization principle for determining movement duration. *Journal of Neurophysiology, 95*, 3875–3886.
Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience, 7*, 907–915.
Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience, 5*, 1226–1235.
Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 20*, 1419–1433.
Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Science, 12*, 291–297.
van Beers, R. J. (2009). Motor learning is optimally tuned to the properties of motor noise. *Neuron, 63*, 406–417.
van Beers, R. J., Brenner, E., & Smeets, J. B. J. (2013). Random walk of motor planning in task-irrelevant dimensions. *Journal of Neurophysiology, 109*, 969–977.
van Beers, R. J., Sittig, A. C., & Denier van der Gon, J. J. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of Neurophysiology, 81*, 1355–1364.
Voudouris, D., Smeets, J. B. J., & Brenner, E. (2013). Ultra-fast selection of grasping points. *Journal of Neurophysiology, 110*, 1484–1489.
Welford, A. T. (1960). The measurement of sensory-motor performance: Survey and reappraisal of twelve years' progress. *Ergonomics, 3*, 189–230.
Young, S. J., Pratt, J., & Chau, T. (2008). Choosing the fastest movement: Perceiving speed–accuracy tradeoffs. *Experimental Brain Research, 185*, 681–688.
Zhang, H., Wu, S. W., & Maloney, L. T. (2010). Planning multiple movements within a fixed time limit: The cost of constrained time allocation in a visuo-motor task. *Journal of Vision, 10*(6), 1.