Intercepting moving targets: why the hand's path depends on the target's velocity

Eli Brenner* & Jeroen B.J. Smeets

Department of Neuroscience, Erasmus MC, Rotterdam, The Netherlands

ABSTRACT

In order to intercept a moving target one must reach some position at the same moment as the target. Considering that moving towards such a position takes time, it seems obvious that one must determine where one can best intercept the target well in advance. However, experiments on hitting moving targets have shown that the paths that the hand takes when trying to intercept targets that are moving at different velocities are different, even if the targets are hit at the same position. This is particularly evident at high target velocities, which seems strange because the benefit of considering the target's velocity should be largest for fast targets. We here propose that the paths' curvature may intentionally differ for different target velocities in order to maximise the chance of hitting the target. Arriving at the target with a velocity that matches that of the target can reduce the consequence of certain temporal errors. In particular, if the path curves in a way that makes the component of the hand's final velocity that is orthogonal to the hitting direction exactly match the velocity of the target, then no additional error will arise from arriving at the target slightly earlier or later than expected. On the other hand, moving along a curved path is likely to increase the spatial errors. We argue that a compromise between these two influences could account for the differences between paths towards fast and slow targets.

Keywords: Interception, Minimal Jerk, Velocity, Model, Optimal, Arm, Human

1. INTRODUCTION

A seemingly obvious requirement for intercepting moving objects is that one must anticipate the direction and extent of the object's displacement during the time that it takes one to reach it. The accuracy with which one can make such predictions is obviously limited by external factors: even the path of a ball flying through the air is not completely predictable, because it is influenced by factors such as wind. However, for the short time intervals that we will consider in the present paper such limitations can probably usually be ignored. This is certainly the case for the virtual targets that we have used in our experiments. A more fundamental limitation is the accuracy with which the person in question is able to make the judgement. In order to predict where one will hit a target, one must combine information about the target's position and motion with estimates of the timing of one's own movements. Considering the many delays within the human nervous system, the fact that these delays depend on a variety of object properties (e.g. contrast), and the rather limited resolution of judgements of timing (e.g. synchrony), it is surprising how well people can intercept moving targets. So how do they do it? In order to answer this question we have conducted many studies aimed at identifying the information that is used to perform a simple hitting task. An important conclusion from those studies was that people do not use the perceived speed to anticipate where they will hit the target. Instead they rely on a rough estimate based on their recent experience, and gradually adjust the anticipated point of interception as the movement progresses, so that the direction of movement improves as the time across which a prediction has to be made decreases. In the present paper we examine why information about the velocity appears not to be used to predict the point of interception.

1.1. Evidence that visual information about target velocity is not used to anticipate where a target will be hit

There are three reasons for believing that the visually perceived target velocity is not used to plan out where to hit the target. The first reason is that manipulating the apparent speed of the target by moving the background does not influence the hand's path (Brouwer et al. 2002; Smeets & Brenner, 1995). It does influence the speed at which the hand moves towards the target, which suggests that the movement time and the path are planned separately on the basis of

* e.brenner@erasmusmc.nl

Human Vision and Electronic Imaging X, edited by Bernice E. Rogowitz, Thrasyvoulos N. Pappas, Scott J. Daly, Proc. of SPIE-IS&T Electronic Imaging, SPIE Vol. 5666 0277-786X/05/$15

different information (Brenner & Smeets 1996; Smeets & Brenner, 1995). The second reason is that if the target disappears shortly before it is hit, people tend to hit behind fast targets and in front of slow ones (Brouwer et al. 2002). Our explanation for this was that the extent to which people aim ahead of the target depends on an expected target velocity and the remaining time to impact, but not on the perceived target velocity. Since the remaining time decreases as the movement progresses, the error gradually decreases as the hand approaches the target. The direction in which people start moving towards a target appeared not to be influenced by the perceived target velocity, but it was influenced by the target's velocity on the *previous* trial (de Lussanet et al. 2001). This brings us to the third reason for believing that the visually perceived target velocity is not used to plan out where to hit the target. Assuming that the direction in which the hand starts to move is linearly related to the point at which one expects to intercept the target at that moment, one can estimate the expected point of interception at the moment that the hand starts moving by comparing the initial motion of the hand towards moving targets with that towards static targets at various positions. This can be done for targets that are moving at different speeds. Doing so suggested that if the target is moving, subjects aim a fixed distance ahead of it, irrespective of its velocity (Brenner & Smeets 1996).

In a more elaborate study (de Lussanet et al. 2004), we used combinations of target velocities and starting positions that enabled us to directly compare the direction in which the hand starts to move when heading towards targets that are hit at about the same position, but are moving at different velocities. We reasoned that if people correctly anticipate where they will hit the target, the paths will be identical for targets that are hit at the same position, irrespective of their velocities. If people do not anticipate correctly the paths will differ. The results were rather surprising: for low target velocities the paths were the same, indicating that the target's velocity is successfully considered, but for fast target velocities the paths differed considerably. In fact, hitting movements towards fast targets started in the same direction if the targets were at the same place when the movement started, irrespective of the target velocity, indicating that the target's velocity is not considered at all. This finding surprised us, because it is clearly more important to consider the target's velocity if the target is moving fast. Why then do people appear to ignore fast motion?

1.2. Does the direction in which the hand starts to move only depend on where one expects to intercept the target?

The reasoning that led to the surprising asymmetry between fast and slow targets was based on the assumption that the planned path towards a given position does not depend on the target velocity. But is this really so? Irrespective of the path, the speed at which the hand moves towards the target does depend on the target's speed: people move faster when hitting fast targets. Our explanation for this is that getting the timing right is more important if the target is moving fast, because the spatial error that arises from arriving slightly too early (or late) is directly proportional to the target's velocity. Spatial errors become larger when we move faster, irrespective of the target speed. Temporal errors become smaller when we move faster (Schmidt et al. 1979). This too is presumably independent of the target's speed, but the same temporal error gives rise to a larger spatial error if the target is moving faster. Thus the benefit of moving fast increases as the velocity of the target increases (Brenner et al. 2002; Brouwer et al. 2000). We here propose that a similar reasoning could influence the target's path.

In our studies of hitting moving targets, the target always moved on a path that was more or less orthogonal to the direction of the hitting movement (see Figure 1a). The targets were animated spiders that ran rightward across a screen. In this configuration any error in timing the moment of contact with the screen will result in an error in the position of the hit with respect to the position of the target. If one hits the screen too early one will hit ahead of the target, and if one hits too late one will hit behind the target. Of course, in this reasoning whether one is *too* late is related to the moment that the target passes. Thus there can be two reasons for being too late. One could underestimate the time that it will take the target to reach a given position on the screen, or one could misjudge the time that it will take oneself to reach the screen. Although it may seem that these two reasons are equivalent in terms of performance, a consideration of possible movement strategies shows that they are not. The difference lies in the possibilities of improving ones performance by moving on a curved path.

Errors that are caused by misjudging how long it will take to reach the screen can be reduced by moving in a way that shifts the hand laterally at the same velocity as the target near the expected moment of contact with the screen. Thus if the screen is hit earlier than expected it will also be hit earlier on the target's path, and therefore at the same position as the target. The assumption is of course that one is better at aligning ones hand with the target, than at judging the exact time of contact. Of course, if it is the timing of the target reaching a given position that is misjudged this strategy will not help, because one will move consistently behind or ahead of the target. (Instead one may want to approach the target head on, because in that case an incorrect timing will make very little difference, but people could not do this in our experiments). In the present paper we examine whether the curvature in the paths could result from a strategy designed to increase the chance of hitting the target. In particular, could the strange dependence on velocity arise from a set of reasonable assumptions? We answer this question by modelling hitting movement in a manner that allows us to relate the performance to the assumptions.

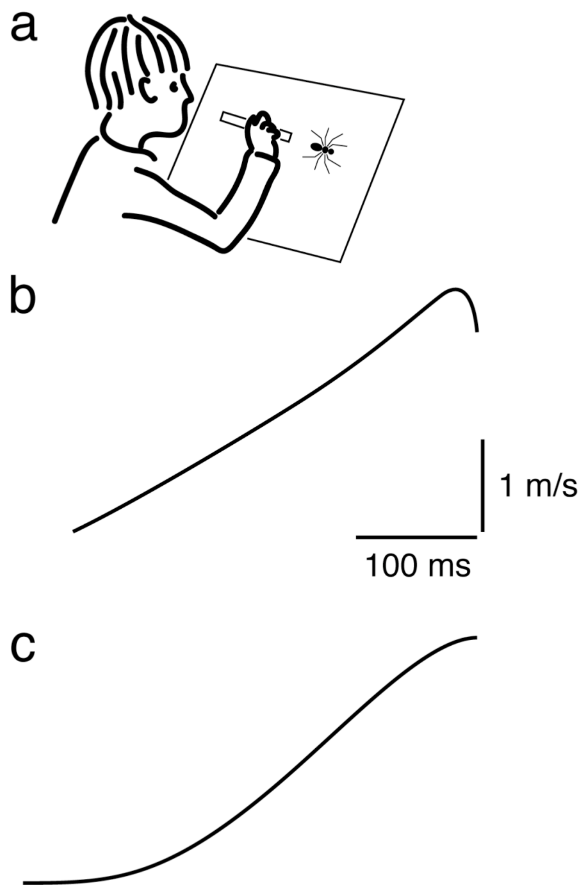

Figure 1a. *Schematic representation of a subject hitting the target (an animated running spider) with a rod.* ***b.*** *The rod's average velocity profile during the hit (taken from the data for static targets in figure 10B of Smeets & Brenner 1995).* ***c.*** *Our simulated velocity profile (equation 1).*

2. MODELLING THE HITTING MOVEMENT

There are many ways in which hitting movements could be modelled. Since we wanted to relate the hand's path to the assumptions in the model, we needed a simple model in which we could easily manipulate the key assumptions. Moreover this model should capture the whole movement, and the movement should be smooth. We therefore decided to use the minimal jerk model (Flash & Hogan 1985) as our starting point. The initial position of the hand was defined as the origin, and the initial velocity and acceleration were both set to zero. The final position was set to be at a distance (*D*) of 40 cm in the main direction of the hand movement and 0 cm in the orthogonal direction, which is the direction in which the target was moving. Furthermore, we assume that the acceleration in both directions is zero at the time of contact. The only value that we will vary is therefore the final velocity of the hand.

For the velocity in the direction of the screen we want a large final value, because our experimental data show that the velocity increases until just before the screen is hit (see Figure 1b). One advantage of ending with a high velocity is that moving fast near the time of contact reduces the influence that errors in judging the distance to the screen have on the timing of the hit. The task in our experiments was always to hit the target as quickly as possible. What we meant by this instruction was that subjects should minimise the movement time (and reaction time) while still hitting the targets. The same movement time can be achieved by minimal jerk trajectories with many patterns of acceleration. We chose the pattern that resulted in the largest final velocity without the hand starting to move in the 'wrong' direction (i.e. away from the screen). Although this choice was rather arbitrary, we considered it satisfactory because it gave us velocity profiles that are quite similar to those measured experimentally. Our choice results in a displacement in the direction of the target that is given by:

$$y_t = \frac{D}{2}\left(5t^4 - 3t^5\right) \qquad (1).$$

Where t is the time that has passed, expressed as a proportion of the movement time. The velocity increases from 0 at t=0 to 5D/2MT at t=1 (Figure 1c). If the hand moves straight towards the target, equation 1 is enough to describe the whole movement. However, we know that for some target velocities the paths curve differently for different target speeds. In the following sections we model an additional lateral component that gives rise to curved paths. We examine the pattern of curvature for lateral components based on several assumptions.

2.1. A model based on constantly heading for the current target position

Before trying to optimise the path with respect to hitting errors, we first consider the simplest possible reason for moving along curved paths. In our previous papers, we explained differences between the paths towards targets that were hit at the same position, but were moving at different speeds, as being caused by failing to account for the target's speed, and therefore initially aiming for the wrong position. If the estimate of the point of interception is continuously improved during the movement, the hand will still hit the target, but it will have moved along a curved path. Figure 2a shows what the path would look like for the extreme case in which the hand is constantly directed towards the actual position of the target. Paths towards a single point of interception are shown for several different target velocities. The paths were determined by numerical simulation. For a static target the above-mentioned strategy obviously leads to a straight path with the velocity profile shown in Figure 1c. If the target is moving, the hand will initially move towards an earlier target position.

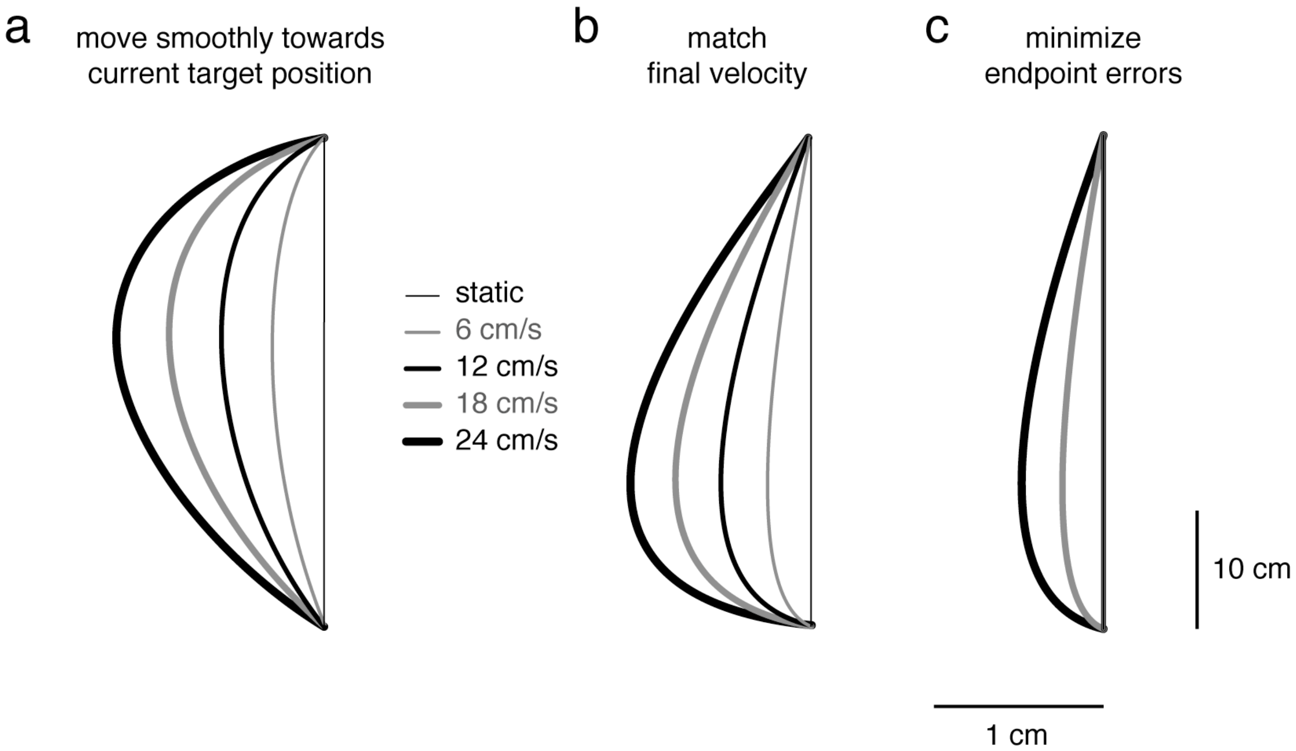

Figure 2. *Paths towards a single position for five different target speeds. The hand starts at the bottom of each curve and moves upwards. The target moves to the right across the top of each figure at the indicated speed. Predicted paths are shown for three different minimal jerk models. Actually all the paths are almost straight, but they have been stretched laterally (see different scales in the two directions) to emphasize the differences.* ***a.*** *At each instant the hand is directed towards the instantaneous target position, so that it is initially directed too far to the left.* ***b.*** *The path is intentionally curved so that the final rightward velocity of the hand exactly matches that of the target.* ***c.*** *The path is only intentionally curved to the extent that the benefit of a curved path outweighs the cost.*

To model the hand's path we combined a movement towards the instantaneous target position (based on equation 1) with a lateral movement. At each instant we determined how the hand would need to move laterally with respect to a straight line towards the target's position at that moment, in order to reach the target at that position with a minimal total squared jerk (jerk is the derivative of acceleration). Initially the optimal solution is obviously to have no lateral motion, so the hand starts moving straight towards the target. However, as the target moves away from the position towards which the hand is heading, so that part of the hand's motion is considered to be lateral (because the directions are defined with respect to the direction of the target), the lowest jerk is found for a path with a non-zero lateral component. This component was determined by finding the value for which the total jerk in the remaining movement of the hand would be minimal if the target did not move any further. For each 1 ms step in our simulation, the initial position, velocity and acceleration were the result of combining all the previous steps (starting at zero for t=0; Henis & Flash, 1995). The anticipated final position was the target's position at that instant. The final acceleration was always assumed to be zero. The final lateral speed was selected to minimise the jerk. The required acceleration at each instant was used to predict what the hand would do during the next ms.

Figure 2a shows how the hand initially moves towards the target and then gradually curves towards the point of interception. This very simple model is not consistent with our finding that people initially aim ahead of the moving target's instantaneous position (Brenner & Smeets, 1996). More importantly, the model clearly cannot explain why people appeared to consider slow velocities more efficiently than fast ones (de Lussanet et al. 2004).

2.2. An intentionally curved path (matching final lateral velocities)

In section 1.2 we proposed that people might intentionally move on a curved path in order to increase their chance of hitting the target if they arrive earlier or later than they expected. In order to minimise the effect of arriving earlier or later than expected one could ensure that the final lateral movement of the hand is identical to that of the target, so that the exact moment of contact is irrelevant. This strategy will also result in subjects initially aiming behind the target, because they must do so in order to be able to move in the same direction as the target at the moment of impact. Using the minimal jerk model to describe such movements is very straightforward, because all one needs to do is to introduce a final lateral velocity that is identical to that of the target. If the target is moving orthogonal to the hand's total displacement, the hand's movement towards the screen is given by equation 1. The hand's initial and final lateral positions are zero, as is the initial velocity and acceleration. In order to move at the same lateral velocity as the target near the moment of contact the hand's final lateral velocity should be equal to the velocity of the target and the final lateral acceleration should be zero. If so, the hand's lateral displacement is given by:

$$x_t = v_{target} \times MT\left(7t^4 - 3t^5 - 4t^3\right) \qquad (2).$$

Where positive values are in the direction of target motion, t is the time that has passed since the movement began (expressed as a proportion of the movement time), and v_{target} and MT are the target velocity and the movement time, respectively.

Figure 2b shows the paths that are predicted by this model. As is evident from equation 2, the curvature is directly proportional to the target's speed. The curvature is also proportional to the movement time. The excursion is *larger* if the hand moves more slowly, because if one moves slowly the path has to curve more to reach the same final lateral velocity in a maximally smooth movement. For our figures we used a movement time of 220ms, because that was the approximate movement time in the experiment that most clearly revealed the phenomenon that we are trying to explain: the different ways of treating fast and slow motion (de Lussanet et al. 2004). We ignored the fact that the movement time differed systematically between subjects and was influenced by the target velocity.

This second model predicts similar paths to those predicted by our first model, although there is a clear difference in the timing of the hand's lateral motion between the paths predicted by the two models. The similarity is not at all self-evident considering that in this model the curvature arises from considering the target's velocity, whereas the model that was described in the previous section gives rise to curved paths by not considering the target motion at all. Most importantly, neither of the models can account for the difference between the ways in which fast and slow targets are treated.

2.3. Minimising errors

In section 2.2 we proposed that people try to end their movement with their hand moving laterally at the same speed as the target. The reason that they would want to do so is that this minimises the influence of arriving slightly earlier or later than anticipated. However, moving on a curved path may also have a price. To understand this, let us consider three possible origins of hitting errors.

One obvious source of errors is misjudging the target's position or velocity with respect to the initial position of the hand. We can distinguish between misjudging the distance to the screen, misjudging where the target's trajectory is on the screen, and misjudging the target's position on that trajectory. In our experimental configuration (see Figure 1a), moving fast can reduce the temporal consequences of errors in estimating the target's distance (i.e. the distance to the screen). Arriving more or less orthogonal to the *screen* will decrease the spatial consequences of misjudging the screen distance, so the optimal path may not be straight towards the *target* (Brenner & Smeets, 1995), but we will ignore this because all the movements that we discuss here end more or less orthogonal to the screen. Errors in localising the target's trajectory on the screen obviously cannot be reduced by moving the hand differently. Moving the hand differently could make it less important to correctly localise the target on its trajectory. If the movement is stopped by contact with the target, then increasing the time spent on the target's path, for instance by moving on a path that brings the hand to hit the target head on, could decrease the errors because it means that even if one misjudges the target's position on its path one is still likely to hit it, though not at the anticipated position. In our experiments such errors in localising the target cannot be reduced by moving along a different path, because the hand was considered to have missed the target if it hit the screen at a position that the target would only reach shortly after the moment of impact.

A second obvious source of errors is misjudging the timing of ones own movements. Misjudging a movement's timing with respect to that of the target may seem to be completely equivalent to misjudging the target's position on its trajectory, which has been discussed in the previous paragraph, but this is not necessarily so. If the errors arise from uncertainty about the time needed to reach the screen (i.e. the target's trajectory), rather than from failures to synchronise the hand with the target, then moving in a way that reduces the influence of variability in the moment of impact could be beneficial. This can be achieved by ensuring that the hand is moving along with the target near the moment of impact. This is the basis for the curved paths that we discussed in the previous section.

The final source of errors that we will consider arises from variability in executing the planned movements. Although this factor may seem to be unavoidable, and therefore irrelevant to our reasoning, it is not, because the errors in executing the movement will depend on the path. If the magnitude of the errors is proportional to the amplitude of the neuronal control signal (Harris & Wolpert, 1998) or the length of the path, then moving on a curved path will increase the errors. The precise relationship between path curvature and endpoint errors is difficult to assess. In particular it is difficult to assess how the signals related to the different directions of movement interact. Obviously the errors on the screen do not only depend on the lateral acceleration of the hand, but also on the acceleration in the direction of the target. However, since the variability is likely to be largest in the direction of the acceleration, the lateral component of the acceleration is likely to give rise to relatively large errors on the screen. In order to explain the difference between the ways in which slow and fast movements are dealt we assume that the combined error is the sum of a component that depends on the acceleration towards the screen and a component that depends on the hand's lateral acceleration. We will discuss this assumption in section 3.1.

In order to evaluate how the path influences the final error we must combine all the above-mentioned sources of errors. Of particular interest are the sources that depend on the velocity of the target or the hand's path. We identified two such sources. One is the spatial variability due to execution errors. Assuming that the contribution of the lateral acceleration (that makes the hand follow a curved path) to the standard deviation in the endpoints is directly proportional to the mean lateral acceleration during the movement, we can determine how the variability will depend on the final lateral velocity of the hand (v_{hand}) by integrating the (absolute value of the) lateral acceleration. Doing so shows that this component of the standard deviation is directly proportional to the final velocity of the hand,

$$SD_1 = c_1 \times |v_{hand}| \quad (3),$$

where c_1 is a constant. For a constant movement time (i.e. acceleration in the direction of the screen) the contribution of the acceleration towards the screen to the standard deviation in the endpoints is a constant, which we can add to the former equation to obtain the total contribution of errors in executing the movement:

$$SD_{exec} = c_1 \times |v_{hand}| + c_2 \quad (4).$$

The second source of errors that depends on the hand's path is misjudging the moment of contact. The variability in timing the moment of contact probably depends on the movement time (Balasubramaniam et al. 2004; Newell et al. 1979; Schmidt et al. 1979), but since the movement time is considered to be constant the spatial error that this introduces is directly proportional to the motion of the target relative to the hand:

$$SD_{timing} = c_3(|v_{hand} - v_{target}|) \quad (5).$$

Finally all other sources of variability are considered to combine into a single constant value that is independent of the velocity of the target or hand:

$$SD_{rest} = c_4 \quad (6).$$

Assuming that the three sources of variability are independent of each other we can combine them to estimate the total variability:

$$SD_{total} = \sqrt{SD^2_{exec} + SD^2_{timing} + SD^2_{rest}} \quad (7).$$

We can then find the value of v_{hand} for which SD_{total} is minimal. This is so when

$$v_{hand} = \frac{c_3^2 \times v_{target} - c_1 c_2}{c_3^2 + c_1^2} \quad (8a),$$

if $c_1 c_2 < c_3^2 \times v_{target}$, and otherwise

$$v_{hand} = 0 \quad (8b).$$

Thus the path towards a given position will not depend on the target velocity until the velocity v_{target} reaches a value of $c_1 c_2 / c_3^2$. After that the influence depends on how much larger c_3 is than c_1. In order for the transition to occur at a target velocity of 12cm/s (de Lussanet et al. 2004), $c_1 c_2 / c_3^2$ must be 12 cm/s. If we use the standard deviation in repeated movements, which is about 10% of the movement time (Brenner et al. 2002), as an estimate of the accuracy with which people can judge the timing of a hit, then for an average movement time of 220ms (see de Lussanet et al. 2004) the value of c_3 will be about 22ms. In our experiments, the standard deviation of hits towards static targets was about 1cm, so assuming that it is mainly caused by errors in execution we can estimate that c_2=1 cm. In that case, $c_1=12\text{x}0.022^2/1=0.0058$ s. For these values of the constants we can determine the paths that correspond with the optimized final velocity of the hand (using equation 8). For static targets and ones moving at 6 or 12 cm/s, the optimal value of v_{hand} is 0 (equation 8b). For targets moving at 18 or 24 cm/s the optimal values are 5.6 and 11.2 cm/s respectively (equation 8a). The corresponding paths are shown in Figure 2c.

2.4. Comparison with real data

Using these values we can model the results of the first experiment that clearly showed a difference between paths for fast but not slow targets (experiment 1 in de Lussanet et al. 2004). Figure 3 shows the main results from that study (data) as well as our model trajectories based on the values given above (model). Note that the constants were selected to fit these data qualitatively, so the fact that the fit is quite good is not too surprising. However, having found a good fit we can examine whether the selected constants are reasonable. If they are, we will not only have support for our model, but will also have a first indication of the contributions of various sources of errors to the overall behavioural variability.

The easiest constant to discuss is c_3. We set this value to be 10% of the movement time in accordance with the variability in real movements. This is only justified if not knowing how long it will take the hand to reach the screen is the main source of temporal errors, rather than synchronising the motion of the hand with that of the target. If this assumption is incorrect, then the value of c3 will be lower, in which case the values of the other constants would also have to be lower to obtain the same quality of fit.

The constant c_2 represents the error that arises when executing a movement straight to the screen. By setting this value to 1cm we attributed most of the measured error of hits towards static targets to this factor. This seems to be a reasonable assumption, but if perceptual localization plays a larger role the only change is that c_2 will be slightly smaller and c_4 will become more important. A lower value of c_2 need not influence the optimal speed of the hand, if the values of the other constants are also lower, but lower values for all the constants except c_4 would reduce the benefit of optimizing the final speed of the hand, because the influence of c_4 on the total final variability is independent of the velocity of the target or of the hand. The value of 1 cm is our estimate for a movement time of 220 ms. Obviously this value depends on the movement time: faster movements are less accurate. If we assume that c_2 is proportional to the total acceleration of the hand during the movement then its value is actually inversely proportional to the movement time (see next paragraph).

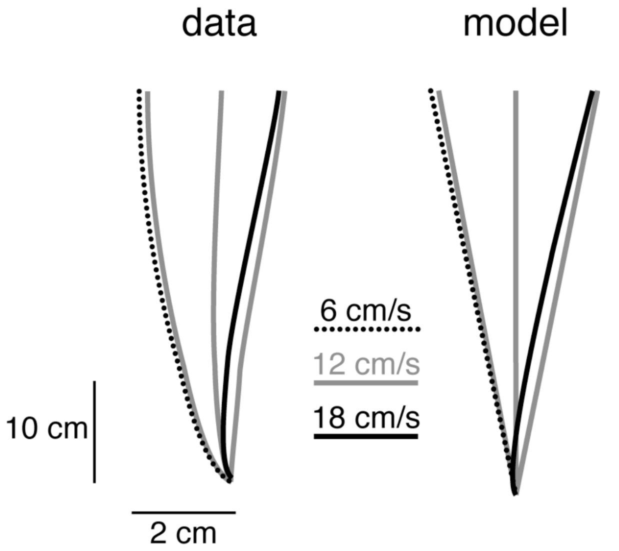

***Figure 3**. Example of the asymmetry between hitting fast and slow targets. The data are redrawn from Figure 3A of de Lussanet et al. 2004. The model is based on the constants given in section 2.3. Different line styles indicate different target velocities. The three grey curves represent targets that appeared, and were therefore also hit, at different positions on the screen.*

The critical value to examine is that of c_1, because the value of c_1 was specifically selected to ensure that the hand would start moving on a curved path from a target velocity of 12cm/s. It is difficult to evaluate whether the value of c_1 is reasonable as such, but its value can be compared to that of c_2 if we are willing to assume that both relate to the same origin of spatial errors: neuromuscular variability that is proportional to the total acceleration during the movement. The total acceleration towards the screen is $5D/2MT$, where D is the distance to the screen (40 cm in our model, although the real value was slightly lower) and MT is the movement time (220ms). The total (absolute) lateral acceleration is $2.024v_{hand}$. Thus the lateral acceleration is about $0.0036v_{hand}$ of the acceleration towards the screen. According to equation 4 the contributions of the variability in the lateral and forward components of the movement of the hand to the total variability are c_1v_{hand} and c_2 respectively, so the lateral component is $0.0058v_{hand}$ of the forward component (c_1/c_2=0.0058 s/cm). Thus the ratio of the 'fit' values of c_1 and c_2 is quite close to what we would predict on the basis of the accelerations. The fact that the ratio is about 60% larger is easy to explain from the fact that the variability in the direction of motion is usually (quite understandably) larger than the variability in the orthogonal direction, so for our lateral errors the influence of the lateral acceleration of the hand is slightly larger than that of the acceleration towards the screen. The value of c_1 in relation to c_2 can therefore be considered as support for our model. Moreover, Brouwer et al. (2000) found that the endpoint error is about 7% larger for targets moving at 18cm/s than for targets moving at 6cm/s, while our model (with this combination of constants and a negligible contribution of c_4) predicts that the endpoint error will be about 6% larger for targets moving at 18cm/s.

Finally, our model may explain why we did not notice the asymmetry between fast and slow targets in our first experiment (Brenner & Smeets 1996; Smeets & Brenner 1995). One factor that was conspicuously different in that study is that people hit the targets much more slowly: the average movement time was over 300 ms. If we ignore the differences between people, and simply increase the value of c_3 from 0.022 to 0.03 (10% of 0.3s), then for the same values of c_1 and c_2 equation 8 predicts that people will start following a curved path from a target velocity of about 6 cm/s, which was the lowest target velocity that we used (apart from static targets). As mentioned above, the value of c_2 is probably inversely proportional to the movement time, in which case people will even start following a curved path from a target velocity of about 4.4 cm/s. Thus we could explain some of the paths in our original study by accepting that people intentionally move in a way that makes the hand curve towards the target, so that they move differently towards static and moving targets (Figure 10A in Smeets & Brenner 1995) and appear to initially be aiming towards an inappropriate position (Figure 3 in Brenner & Smeets 1996).

3. DISCUSSION

Figure 3 clearly shows that it is possible to account for the data that we set out to explain on the basis of the proposed model. In the next section we will discuss the assumptions that we had to make to achieve this. We will pay particular attention to two of the most critical assumptions. However although these assumptions are critical for this particular model, they are not fundamental for the general idea that there may be an advantage of moving along a curved path, and that this advantage may depend on the target's speed. In the previous sections we developed a model that is based on minimizing the hitting errors. However, there may be other benefits of following a curved path. For instance, following the target more closely, by starting one's movement with a bias towards the direction of the target, could be advantageous if the target suddenly changes its direction of motion. We currently have no idea how to quantify such influences, so they are not incorporated in our model. It is uncertain whether such considerations alone could account for the different way in which fast and slow targets are treated, but the possibility that curved (or straight) paths may have other benefits than those mentioned in section 2 should be kept in mind.

3.1. The assumptions

We have made many assumptions. A number of these assumptions are critical for our model whereas others are inconsequential. The most critical and debatable assumption is that the variability in movement endpoints as a result of errors in executing the movement increases linearly with the lateral velocity of the hand, and thus with the curvature of the path, with an intercept that is not zero (equation 4). The reasons for proposing this particular pattern are that it sounds reasonable that the variability should increase linearly with the curvature and that doing so gives our model the properties that we are looking for. This equation would for instance hold if the final error were related to the length of the path. However, we proposed that the magnitude of the error is proportional to the acceleration of the hand rather than the length of the path. If the magnitude of the acceleration were determined at each moment (irrespective of its direction) and then integrated over the duration of the movement, the contribution of the lateral component of the acceleration would be extremely small. Similarly, if the deviation due to the lateral component were proportional to the total lateral acceleration, and completely independent of the deviation due to the motion towards the screen, then equation 4 would be replaced by:

$$SD_{exec} = \sqrt{c_1^2 \times v_{hand}^2 + c_2^2} \qquad (9).$$

In that case the critical term c_1c_2 disappears from equation 8, so that the curvature increases linearly with the target velocity for all target speeds. In order to obtain the effects that we see in our experiments (de Lussanet et al. 2004) the total sum of the separate accelerations in the two directions must determine the error, as in equation 4. If it is indeed noise in the control signals that makes the error depend on the acceleration, as proposed by Harris and Wolpert (1998), then this makes sense because it is clear that the two components serve different functions, and probably even make use of different muscles, so there is no justification in calculating an overall signal amplitude (as in equation 9). However, the precise form of equation 4 remains rather arbitrary.

The second debatable assumption is that people are poor enough at estimating the duration of their own actions to justify moving on a curved path. There are many examples of perceptual studies in which all kinds of timing judgments are shown to be performed quite poorly, but that is not the issue here. In order to predict when one will hit the screen one must predict how long the movement is going to take (and possibly when one will start moving; Schmidt, 1969). A critical assumption for our proposal is that it is this estimate that we are uncertain of. As we already mentioned in section 2.3, errors in judging the temporal relationship between the felt position of the hand and the visible position of the target cannot be reduced by changing the movement path. Clearly it is generally very difficult to determine whether arriving at a target too early or too late is the consequence of misjudging the moment of impact or of misjudging how long the movement will take. Any support for our model could therefore be taken to support the idea that misjudging the duration of the movement is an important source of uncertainty. However, this uncertainty must mainly be in the component towards the screen, because otherwise a curved path will provide no advantage. Perhaps synchronising the hand with the target (laterally) is easier than judging the duration of the movement. In any case this distinction is consistent with considering the two directions of motion as separate but related components, as we did in the previous paragraph.

3.2. Were our previous conclusions all wrong?

An obvious question that arises from the present analysis is whether all our previous conclusions were wrong. This is not necessarily so, because the reasoning in the present explanation is teleological in nature, whereas our previous studies examined the possible causes of the observed behaviour. In particular, de Lussanet et al. (2004) developed a mass-spring model that readily explained the same data. The main difference between the model that we present here and that presented in de Lussanet et al (2004) is that the present model suggests why these particular paths may be advantageous, whereas the previous model suggests how visual information could be related to the lateral acceleration of the hand to obtain these paths. Thus, the curved paths which reduce the average final error may well arise from aiming a fixed distance ahead of moving targets rather than a distance that depends on their velocity (Brenner & Smeets 1996). Similarly the fact that movements end with a substantial velocity in the direction of target motion could be due to a mechanism that relies on relative damping (de Lussanet et al. 2002).

We previously related the fact that people did not follow the same path to a given position for targets that moved at different velocities, to limitations in the use of visual information for guiding behaviour. In particular, we concluded that the perceived speed was ignored. Ignoring the perceived speed could be justified by the finding that the neuronal delays are particularly long for information about motion (Brenner et al. 1998). However, the present investigation shows that the fact that people take different paths to hit slow and fast targets is not necessarily the result of failing to make correct predictions, but the paths may intentionally curve more for fast targets because of the benefits that this has in terms of the final accuracy. Nevertheless, the *mechanism* that the brain uses to achieve the curved path may well be to ignore the speed above a certain value, even if the reason for doing so is not that it is impossible to consider all speeds, but that it is advantageous to only consider the target's speed within a certain range. Some support for relying on mechanisms based on partial information rather than predicting the optimal trajectory in advance can be found in the next section.

3.3. What the model does not explain

The model that we here propose can account for most of the paths that we found, but it does not explain why the path changed at all when the target disappeared (Brouwer et al. 2002), or why the path was not affected by moving the background to make the target appear to move faster or slower (Smeets & Brenner 1995). In order to answer such questions one must consider the mechanisms by which visual information is converted into actions. Moving backgrounds and disappearing targets are not conditions that are likely to have guided the evolution or development of interceptive skills, so such manipulations do not need to be considered when developing a model that is designed to find the trajectories that give an optimal performance. However, such manipulations are clearly useful when searching for the mechanisms that underlie the actual performance.

Subjects hit targets that disappeared just before they were hit differently than they hit ones that did not. This indicates that the hand is guided by a constant interaction between the available perceptual information and the control of the muscles, rather than the trajectory being planned in advance in accordance with an optimal velocity of approach and then executed as planned. The fact that people tended to hit too far *ahead* of slow targets that disappeared, and the fact that people appeared to start their movements towards a position that was ahead of moving targets rather than towards their current position, shows that people do not just move towards the target as proposed in section 2.1. A model such as that proposed in de Lussanet et al. (2002; 2004) can explain these finding as well as the lack of influence of moving the background. Thus that model is more complete. What the present model adds is a possible explanation why such a strange coupling between visual information and muscle activation could evolve.

3.4. Conclusion

The present analysis shows that there may be good reasons for people to follow a curved path when trying to intercept moving targets. The optimal path may be a compromise between the advantages of approaching the target in a particular manner, and the disadvantages of moving along a longer, curved path. In the present paper we examine this compromise for a single kind of hitting task, because this is a task for which we have abundant experimental data. However we expect that the same kind of reasoning will apply to any experimental configuration, so that future studies with other configurations may not only help determine how our interceptive movements are controlled, but also how exactly the variability in the endpoint depends on each of the many sources of errors that we have mentioned in section 2.3 (and perhaps also others that we have missed). The present endeavour nicely demonstrates how different kinds of models can complement each other in furthering our understanding of motor control: while our previous models indicated that certain kinds of visual information guide certain aspects of our actions, and how they do so, the present model suggests why this particular coupling may be beneficial.

REFERENCES

1. Balasubramaniam R, Wing AM, Daffertshofer A (2004) Keeping with the beat: movement trajectories contribute to movement timing. Experimental Brain Research 159:129-34.
2. Brenner E, Smeets JBJ (1995) Moving one's finger to a visually specified position: target orientation influences the finger's path. Experimental Brain Research 105:318-320
3. Brenner E, Smeets JBJ (1996) Hitting moving targets: co-operative control of ‘when’ and ‘where’. Human Movement Science 15:39–53
4. Brenner E, Smeets JBJ, de Lussanet MHE (1998) Hitting moving targets: continuous control of the acceleration of the hand on the basis of the target’s velocity, Experimental Brain Research 122:467–474
5. Brenner E, de Lussanet MHE, Smeets JBJ (2002) Independent control of acceleration and direction of the hand when hitting moving targets. Spatial Vision 15:129–140
6. Brouwer A-M, Brenner E, Smeets JBJ (2000) Hitting moving targets: the dependency of hand velocity on the speed of the target. Experimental Brain Research 133:242–248
7. Brouwer A, Brenner E, Smeets JBJ (2002) Hitting moving objects: is target speed used in guiding the hand? Experimental Brain Research 143:198–211
8. Flash T, Hogan N. (1985) The coordination of arm movements: an experimentally confirmed mathematical model. Journal of Neuroscience 5:1688-1703.
9. Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. Nature 394:780-784
10. Henis E, Flash T (1995) Mechanisms underlying the generation of averaged modified trajectories. Biological Cybernetics 72:407-419
11. de Lussanet M HE, Smeets JBJ, Brenner E (2001) The effect of expectations on hitting moving targets: influence of the preceding target’s speed. Experimental Brain Research 137:246–248
12. de Lussanet MHE, Smeets JBJ, Brenner E (2002) Relative damping improves linear mass-spring models of goal-directed movements. Human Movement Science 21:85–100
13. de Lussanet MHE, Smeets JBJ, Brenner E (2004) The quantitative use of velocity information in fast interception. Experimental Brain Research 157:181–196
14. Newell KM, Hoshizaki LEF, Carlton MJ, Halbert JA (1979) Movement time and velocity as determinants of movement timing accuracy. Journal of Motor Behavior 11:49-58
15. Schmidt RA (1969) Movement time as a determiner of timing accuracy. Journal of Experimental Psychology 79:43-55
16. Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT (1979). Motor-output variability: A theory for the accuracy of rapid motor acts, Psychological Review 86, 415– 451.
17. Smeets JBJ, Brenner E (1995) Perception and action are based on the same visual information: distinction between position and velocity. Journal of Experimental Psychology; Human Perception and Performance 27:77–88