

Statistical decision theory and trade-offs in the control of motor response

JULIA TROMMERSHÄUSER*, LAURENCE T. MALONEY
and MICHAEL S. LANDY

Department of Psychology and Center for Neural Science, New York University, New York, NY, USA

Received 12 June 2002; revised 13 August 2002; accepted 10 September 2002

Abstract—We present a novel approach to the modeling of motor responses based on statistical decision theory. We begin with the hypothesis that subjects are ideal motion planners who choose movement trajectories to minimize expected loss. We derive predictions of the hypothesis for movement in environments where contact with specified regions carries rewards or penalties. The model predicts shifts in a subject's aiming point in response to changes in the reward and penalty structure of the environment and with changes in the subject's uncertainty in carrying out planned movements. We tested some of these predictions in an experiment where subjects were rewarded if they succeeded in touching a target region on a computer screen within a specified time limit. Near the target was a penalty region which, if touched, resulted in a penalty. We varied distance between the penalty region and the target and the cost of hitting the penalty region. Subjects shift their mean points of contact with the computer screen in response to changes in penalties and location of the penalty region relative to the target region in qualitative agreement with the predictions of the hypothesis. Thus, movement planning takes into account extrinsic costs and the subject's own motor uncertainty.

Keywords: Motor planning; statistical decision theory; judgment under risk.

1. INTRODUCTION

Motor responses have consequences. According to Swiss tradition, the tyrant Gessler ordered the marksman Wilhelm Tell to shoot an apple off the head of Tell's own son with a crossbow at 50 paces distance (Fig. 1A). Were Tell successful, he would be released. Should he refuse or miss altogether, both he and his son would be put to death. It is not clear what Gessler intended the outcome to be if Tell had shot and hit the boy. Perhaps Gessler, cruel man that he was, would free Tell to live out his life with the memory of having killed his own son. We shall assume that this was the case.

*To whom correspondence should be addressed. E-mail: trommer@cns.nyu.edu

In Fig. 1A, every point where the arrow might land is associated with a clear penalty or reward. Of course, Tell cannot pick exactly where his arrow will go. He can only choose where to aim, and even if he aims at the exact center of the apple, he knows he risks hitting the boy below. If he aims a bit above the center, he reduces the risk that he will hit the boy, but likely increases the risk that he will miss both apple and boy. The question is, where should Tell aim?

Let us consider an experimental task analogous to Wilhelm Tell's quandary (but with less grave consequences). Following a signal, the subject must rapidly touch a point on a computer screen with his/her fingertip. The items visible on the screen are illustrated in Fig. 1B. If the subject hits within the upper circle on the screen ('the apple'), s/he wins a dollar. If s/he instead hits the circular penalty region just below ('the child'), s/he loses ten dollars. Hitting anywhere else on the screen incurs no penalty (but generates no reward). There is a large penalty for failure to respond within a brief time after the signal.

In the situation just described, where should the subject aim? With a short enough time limit and small enough target regions, the subject cannot be certain that a movement aimed at the center of the reward region will not, instead, end up

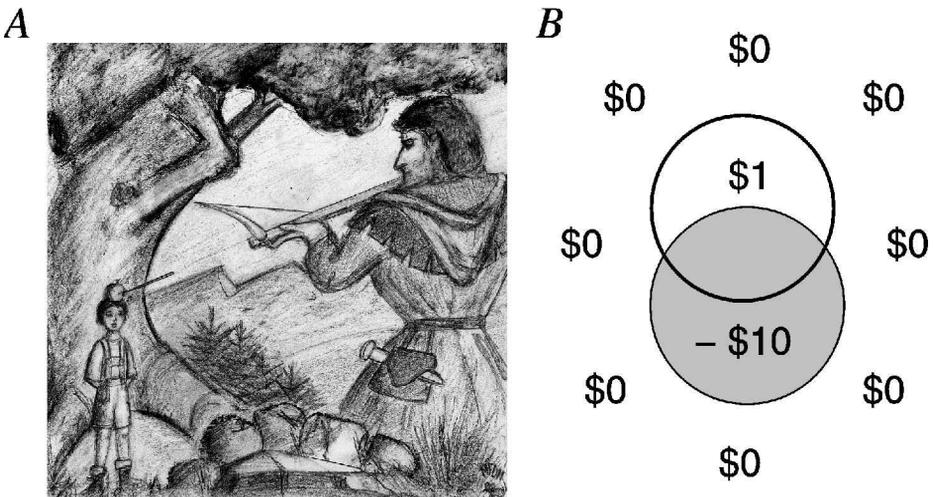


Figure 1. Wilhelm Tell's quandary. (A) Tell, forced to shoot at an apple perched on his son's head, must decide precisely where to aim. Intuitively, this should depend on the consequences associated with hitting the apple, hitting his son, and missing both apple and son (drawing by M. F. Dal Martello). (B) A pointing task analogous to Wilhelm Tell's quandary. Once a signal is given, the subject must quickly reach out and touch a point in the figure. The rewards associated with each point are labeled by region. The region of highest reward is enclosed in the white circle. If the subject hits in the intersection of the target and penalty region, s/he collects both rewards, for hitting the target and the penalty, respectively. The time limit in the task is chosen so that the subject cannot reliably hit within the white circle without some risk of hitting the grey circular region beneath that carries a large penalty. There is a very large penalty for not responding before the time limit expires. What is the optimal point to aim for?

outside the reward region, possibly in the penalty region. Intuition suggests that the wise subject should aim somewhere above the center of the reward region, but it is not at all clear by how much. It is plausible that the values of the rewards and penalties should affect the choice of aim point: the larger the penalty associated with the penalty region, the larger we would expect the subject's aiming bias to be, all else held constant. The experimenter can increase or reduce the uncertainty in the subject's motor response by manipulating the time limit,¹ and it is plausible that altering motor uncertainty should also affect the subject's choice of aim point.

The task just described is representative of a wide range of everyday movement planning tasks where there are (a) externally-imposed costs and benefits associated with the outcomes of actions, and where (b) the uncertainty inherent in carrying out a planned movement may alter the consequences for the planner. In this article we first develop a model, based on statistical decision theory (Blackwell and Girshick, 1954; Ferguson, 1967), that predicts optimal behavior for such 'Wilhelm Tell' tasks. We then report an experiment in which subjects' actual performance is compared to predicted optimal performance. The Theory of Signal Detection (Green and Swets, 1966/1974) is an analysis of detection experiments that models an ideal observer who makes decisions so as to minimize expected loss. We propose an analogous model of movement planning.

2. MODELS OF MOTOR STRATEGY AND MOTOR BEHAVIOR

In planning a goal-directed movement, the motor system is required to pick one of many possible motor programs. Since there is some disagreement in the literature as to exactly what constitutes a motor program, we will use the slightly more inclusive term 'visual-motor strategy' to describe the outcome of movement planning. A strategy includes the choice of the goal trajectory as well as any effect of ongoing visual feedback on the movement. In this section, we review the previous literature on movement planning.

There are many possible trajectories that can satisfy the demands of a particular task. Thus, additional constraints are required to specify a unique solution. It is plausible that the motor system tends to avoid trajectories that lead to 'wear and tear' on the organism. Previous models of motor planning typically emphasize selection of a single deterministic trajectory that minimizes biomechanical costs while achieving the specified goal of the movement. These models differ primarily in the choice of biomechanical cost function.

These biomechanical cost functions include measures of joint mobility (Soechting and Lacquaniti, 1981; Kaminsky and Gentile, 1986), muscle tension changes (Dornay *et al.*, 1996), mean squared rate of change of acceleration (Flash and Hogan, 1985), mean torque change (Uno *et al.*, 1989) and peak work (Soechting *et al.*, 1995). The resulting model predicts that the motor system will select the strategy that minimizes the specified biomechanical cost while achieving the intended goal.

Rosenbaum and colleagues proposed a slightly different model of motor planning. They propose that the motor system computes a goal posture based on a 'database' of stored posture representations (Rosenbaum *et al.*, 1995, 2001). The goal posture is a weighted sum of the stored candidate postures where the weights assigned to each stored posture are chosen on the basis of two criteria: accuracy (How close is the stored posture to fulfilling the goal of the movement?) and efficiency (What are the biomechanical costs of moving from the current posture to the stored posture?). The motor system, given the computed goal posture, then seeks to effect it, using visual and kinesthetic feedback to optimize the requested change in posture. Note that the choice of goal posture involves a trade-off between the accuracy of potential goal postures and the biomechanical costs entailed by change of posture.

Obstacles or other 'extrinsic' objects further constrain the range of movement paths. A number of studies have demonstrated that hand trajectories are altered in the presence of distractors or obstacles (Dean and Brüwer, 1994; Howard and Tipper, 1997; Tipper *et al.*, 1997; Castiello, 2001; Mon-Williams *et al.*, 2001). These findings are typically explained by assuming that trajectories are chosen to avoid spatial overlap of the limb with the obstacle (Rosenbaum *et al.*, 1999). In other obstacle avoidance approaches the inertial properties of the arm are considered as well (Sabes and Jordan, 1997; Sabes *et al.*, 1998). However, the outcome of the proposed computation is still a single, deterministic trajectory that optimizes the trade-off between the goal of the task and biomechanical costs. The optimization now excludes all trajectories that collide with the obstacle since, for a sufficiently solid obstacle, these trajectories are not achievable.

The models just described do not take into account the possible consequences of subjects' motor uncertainty in carrying out tasks. Nor do they allow for consideration of trade-offs among extrinsic costs, imposed as part of the task. There is, for example, no ready way to frame a trade-off between a small probability of collision with an obstacle and a large decrease in the chances of achieving the goal of the task. Nor is it easy to explain why a subject might be more willing to risk collision as the reward associated with successful completion of the task is increased.

There is considerable evidence that the motor system takes its own motor uncertainty into account when planning movements. Consider the task of moving the hand quickly to a target. The task is more difficult (i.e. the landing point is less likely to hit the target) for shorter movement times (the time interval from movement initiation to completion), and for smaller and more distant targets. In fact, subjects show an awareness of these constraints by prolonging their movement time for smaller target diameters as predicted by Fitts law (Fitts, 1954). Thus, subjects take uncertainty into account and select movement times that allow the target to be hit with constant reliability (see Meyer *et al.*, 1988).

Following this observation Harris and Wolpert (1998) suggested that movement trajectories are selected to minimize the variance of the final eye or arm position. They proposed that the underlying determinant of trajectory planning is the mini-

mization of the noise in the neural control signal that activates the muscles during the execution of a motor command and the post-movement period. In their model the arm's position is computed as function of a (deterministic) biomechanical expression and a noisy neural signal, where the noise increases with the magnitude of the neural signal. Optimal trajectories of the eye and arm are then determined by minimizing the total positional variance during the immediate post-movement period (of 500 ms). The idea behind this approach is that the variability in the final position of a saccade or pointing movement is a result of the accumulated deviations of the executed trajectory from the planned trajectory over the duration of the movement. The model managed to describe horizontal saccadic eye movements, hand paths for a set of point-to-point movements, as well as the optimal trajectory in an obstacle avoidance task (Hamilton and Wolpert, 2002).

Note that this approach is based on the notion that the endpoint variability is a consequence of the 'biological noise' in the motor control system and therefore unavoidable. Motor trajectories will always be distributed around the optimal aiming point. If so, does this motor uncertainty affect where we aim?

In the following section, we develop a model for selection of visual-motor strategies that explicitly takes into account motor uncertainty and that allows for trade-offs between the penalty structure imposed by the task and biomechanical costs. The key ideas are as follows.

First, when the motor system selects a visuo-motor strategy, it in effect imposes a probability density on the space of possible movement trajectories that could occur once the motor strategy is executed. This probability density is likely affected by the goal of the movement, the planned duration, the possibility of visual feedback during the movement, previous training, and intrinsic uncertainty in the motor system. We emphasize that the consequences for the subject are completely mediated through this probability density and we can, for the most part, ignore the details of the actual mechanisms that produce and steer the action.

Second, whatever the penalty structure of the task, the penalty incurred by the subject depends only on the motion trajectory that actually occurs. Where Wilhelm Tell aims is important, but where the arrow hits is all that really matters in the end.

Third, the subject acts so as to minimize expected loss as computed from the magnitude of each possible penalty and the probability of incurring it. Note that we use the term 'loss' to describe both gains and losses, treating gains as a 'negative loss'. In our experiments, subjects typically win money and their expected loss is therefore negative (an expected gain).

3. A STATISTICAL DECISION THEORY MODEL OF POINTING MOVEMENTS

A movement trajectory $\tau(t)$ is a function that specifies, for any time t , the position of the entire body at that time. For our purposes, though, it is enough to consider the position of a single fingertip in time and space: $\tau : t \rightarrow (x(t), y(t), z(t))$. We have chosen notation so that the development of the theory is little affected if the

range of $\tau(t)$ is expanded to include a full representation of the position of the arm or of the full body.

We will denote by S a choice of visual-motor strategy. The effect of choosing S is simply to impose a probability density on the set of possible trajectories that could occur when S is executed. For the simple sort of task we will investigate experimentally in a later section, we could think of the choice of S as a choice of aim point or as the choice of a planned trajectory. However, as the word ‘strategy’ suggests, a visual-motor strategy could involve a complex series of goals in space and time. The consequence of any chosen strategy, though, is simply the probability density assigned to possible trajectories, $P(\tau|S)$.

Consider now a target region R in space. For any choice of strategy S , we can compute the probability of passing through the target region (‘hitting the target’). It is simply the probability,

$$P_{hit}(S) = \int_{R_\infty} P(\tau|S) d\tau, \quad (1)$$

where R_∞ is the set of trajectories that pass through R at some point in time after the start of the execution of the visual-motor strategy. If the target region is large compared to the fingertip and no time limit is imposed, it is plausible that the motor system can always pick a strategy S such that the actual trajectory always passes through the target region, i.e. $P_{hit}(S) = 1$.

If we impose a time limit on the task, it is no longer enough to pass through the target region. The trajectory must pass through the target region before the time limit t^* has expired. The probability of hitting the target is now

$$P_{hit}(S) = \int_{R_{t^*}} P(\tau|S) d\tau, \quad (2)$$

where R_{t^*} denotes the trajectories that pass through R after the start of the execution of the visual-motor strategy and before time t^* . Given Fitts’ law (Fitts, 1958), it is plausible that the experimenter can select a time limit and/or reduce target size so that the target will be missed in some cases and, therefore, $P_{hit}(S) < 1$. Our primary concern is with such ‘speeded tasks’ where, on every trial, there is a substantial probability that a movement intended to hit the target region will instead hit another region, incurring a penalty.

If all that mattered were whether or not the subject hit the target, then we could simply choose a visual-motor strategy S^* that maximizes $P_{hit}(S)$ (the choice of S^* need not be unique). However, we should also consider the biomechanical costs associated with different strategies. Further, for the Wilhelm Tell tasks considered above, we must also factor in the possibility of losses incurred by trajectories that pass through penalty regions. It is not enough to maximize the chances of hitting the apple. We must also consider the chances he will hit the boy.

Let us denote the other regions by R_i where i ranges from 1 to N . Then the probability of hitting the i th region before the time limit is

$$P(R_i|S) = \int_{R_{i,t^*}} P(\tau|S) d\tau, \quad (3)$$

where R_{i,t^*} denotes the trajectories that pass through R_i after the start of the execution of the visual-motor strategy and before time t^* . If we change notation slightly, denoting the target region by R_0 , then the equation above summarizes the probability of hitting any region, target or penalty. Let C_i denote the cost associated with hitting region i for i ranging from 0 to N (a ‘reward’ is just a negative ‘cost’). Then the *expected loss* of a movement strategy is calculated as

$$L(S) = \sum_{i=0}^N C_i P(R_i|S). \quad (4)$$

Note that the expected loss includes a possible negative cost from hitting the target region R_0 .

For tasks with an imposed time limit and a penalty for failure to respond before the limit, we should add a term to the expected loss function that reflects this ‘timeout’ penalty. We denote the probability that a task leads to a ‘timeout’ by $P(\text{timeout}|S)$ and the associated cost by C_{timeout} .

Last of all, we need to include biomechanical costs associated with a given movement trajectory due to the ‘intrinsic’ constraints of the arm. We codify these biomechanical costs as a second loss function $B(S)$. This yields a new expected loss function

$$L(S) = \sum_{i=0}^N C_i P(R_i|S) + C_{\text{timeout}} P(\text{timeout}|S) + \lambda B(S), \quad (5)$$

where λ characterizes the trade-off between physical effort and expected reward/penalty that the subject will tolerate.

An optimal visual-motor strategy, according to statistical decision theory is one that minimizes the expected loss function $L(S)$ (Blackwell and Girshick, 1954; Ferguson, 1967; Berger, 1985; see Maloney, 2002). Note that implicit in the form of the previous equation is the fact that subjects can be expected to trade off biomechanical costs and costs imposed externally by the task. As we will see, predicted optimal behavior will be influenced by the presence of penalty target regions, the magnitudes of penalties and rewards, the ‘timeout’ penalty, and changes in motor uncertainty induced by changing the ‘timeout’ time limit.

Note that the form of the final equation does not explicitly depend on how we represent trajectories. However we represent them, all that matters is whether they hit any given penalty region at an appropriate time. Further, while our penalty regions are specified in space, it would be easy to extend them to represent regions

in space-time, allowing us to represent tasks that involve avoiding or seeking objects in a dynamic environment.

It is clear that movement trajectories also depend on visual feedback (for reviews, see Connolly and Goodale, 1999; Desmurget and Grafton, 2000). Depending on timing and viewing conditions, visual feedback may allow the subject to follow a trajectory more accurately. In our formulation, this reduction in motor uncertainty affects only the probabilities of hitting various regions and possibly the probability of timeout.

In the next section, we illustrate the model for the simple task described in the introduction.

4. A THOUGHT EXPERIMENT

Imagine a subject carrying out the simple task, analogous to that faced by Wilhelm Tell, that we discussed above. S/he is asked to touch within a target circle drawn on a computer screen. If the subject hits the target on a trial, s/he will win 100 points. After the experiment is over, the total point winnings are converted to a monetary reward. On all trials, a second ‘penalty’ circle appears to the left of the target, partially overlapping it (Fig. 2). If the subject hits within the penalty circle, s/he will lose 100 points. Note that, as the circles overlap, the subject may hit within both of them and simultaneously incur the specified reward and the specified penalty, receiving zero points. S/he receives the sum of all rewards and penalties earned in a single trial. We force the subject to touch the screen within a specified time limit that is the same on every trial. Late responses incur a very high penalty.

This leaves the subject with the following options:

1. *The target is hit*: The subject receives a reward of C_0 points (as we represent values as *costs*, $C_0 < 0$).
2. *The penalty is hit*: The subject receives a penalty of C_1 points.
3. *Time out*: The subject was too slow and receives a large penalty of $C_{timeout}$.

The first and second options are not mutually exclusive if the target and penalty regions overlap.

So, how does one predict the optimal aiming point that minimizes the cost function (equation (5))? Given our effectively 2-dimensional task, we will designate possible visual-motor strategies S by their resulting mean point of impact on the computer screen (x, y) . We can think of this point as the ‘aim point’ of the subject. For aim point (x, y) , the expected cost is,

$$L(x, y) = C_0 P(R_0|x, y) + C_1 P(R_1|x, y) + C_{timeout} P(timeout|x, y) + \lambda B(x, y). \quad (6)$$

Let us assume that both the probability of a timeout and the biomechanical costs are nearly constant over the limited range of relevant screen locations. To minimize

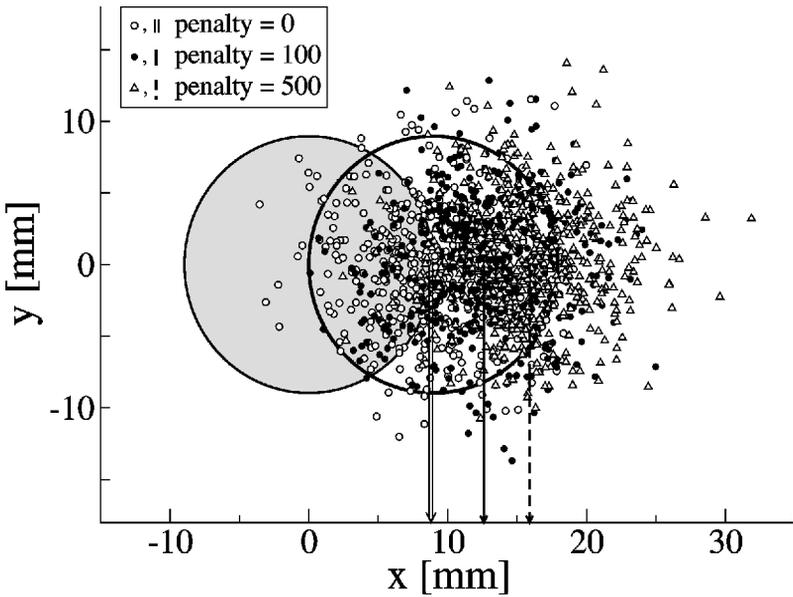


Figure 2. Outcome of the thought experiment. 500 simulated responses per condition for varying penalties. Arrows indicate the optimal aiming point for each possible penalty. The left- and right-hand circles mark the areas of the penalty circle and the target (radius = 9 mm). The response variability ($\sigma = 4.2$ mm) was set so that 90% of the responses hit the target in the zero-penalty condition.

expected costs, we can then ignore the constant timeout and biomechanical penalty terms. Thus, the subject must choose (x, y) so as to minimize,

$$L(x, y) = C_0 P(R_0|x, y) + C_1 P(R_1|x, y). \quad (7)$$

Next, we assume that pointing trajectories are unbiased and distributed around the aim point (x, y) according to a Gaussian distribution,

$$p(x', y'|x, y) = \frac{1}{2\pi\sigma^2} e^{-[(x-x')^2 + (y-y')^2]/2\sigma^2}, \quad (8)$$

where σ indicates the spatial variability of the subject's responses in any direction away from the aim point. The probability of hitting region R_i is

$$P(R_i|x, y) = \int_{R_i} p(x', y'|x, y) dx' dy'. \quad (9)$$

For the stimulus configuration used in our experiment (see Fig. 3) no analytical solution could be found for equation (9). The integral was solved numerically by Monte Carlo integration (Press *et al.*, 1992) and the results were used for minimizing equation (7).

We first consider how the choice of aim point will depend on the penalty C_1 associated with hitting the penalty region. Suppose that C_1 is 0, i.e. there is no penalty associated with hitting the penalty region. Then, we can reasonably expect

that the subject will seek to hit the target as often as possible, given the time limit. The optimal aim point is the center of the target under these conditions. The subject's winnings depend only on motor uncertainty, captured by the parameter σ in the specification of the Gaussian.

Next, suppose that hitting the penalty circle incurs a penalty of 100 points. Given the spatial variability of the pointing responses, the subject may accidentally hit the penalty circle when aiming at the center of the target (see the overlap of the distribution of the filled circles with the shaded penalty area in Fig. 2). When trying to minimize the loss across all trials, it may be preferable to aim at a point which is shifted to the right of the center of the target. Depending on the amount of the penalty, it might be less costly to miss the target once in a while to avoid the risk of incurring so many penalties. Therefore, we expect the optimal aim point to shift farther to the right with increasing penalty (if the target area remains constant, as it does in our experiment).

This scenario is illustrated in Fig. 2, which displays the outcome of a simulation of an optimal subject who minimizes equation (7) for varying penalty values C_1 . When subjects take into account their own response variability *and* the varying penalties, their optimal aim point shifts farther to the right for increasing values of C_1 . For instance, in the case of a penalty of 500 points, the optimal aim point ($x_{\text{opt}} = 16$ mm, $y_{\text{opt}} = 0$ mm) is barely within the area of the target (9 mm radius, target center at 9 mm). Aiming farther to the left increases the penalties due to hitting the penalty circle; aiming farther to the right reduces the chance of collecting points by missing the target.

We now describe an experiment designed to test whether subjects behave in a manner consistent with the theory just outlined. In this thought experiment, we assumed that we knew the subject's motor uncertainty, as described by a Gaussian distribution. In the experiment below, for this simple pointing task, we will estimate subjects' motor uncertainty from the variability in the data. We use that variability estimate to compute subjects' optimal aim point for various choices of penalty using the model developed above. We compare these predictions to subjects' actual performances.

5. METHOD

5.1. Apparatus

Subjects were seated in a dimly lit room in front of a transparent touch screen (AccuTouch from Elo TouchSystems, accuracy $< \pm 2$ mm (std), resolution of 15 500 touchpoints/cm²) mounted vertically in front of a 21-inch computer monitor (Sony Multiscan CPD-G500, 1280 \times 1024 pixels @ 75 Hz). A chin rest was used to control the viewing distance, which was 29 cm in front of the touch screen. The computer keyboard was mounted on the table centered in front of the monitor. The experiment was run using the Psychophysics Toolbox (Brainard, 1997; Pelli,

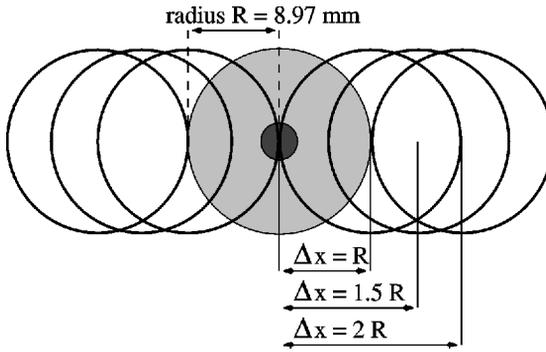


Figure 3. Layout of the stimuli. Stimuli were presented on a black background. The red penalty circle was displayed near the center of the display first (see Fig. 4). The six open circles indicate the six different target positions.

1997) on a Pentium III Dell Precision workstation. A calibration procedure was performed before each session to ensure that the touch screen measurements were geometrically aligned with the visual stimuli.

5.2. Stimuli

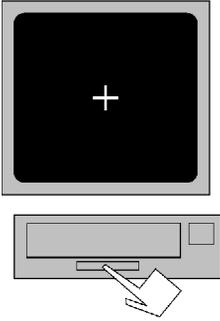
Stimuli were presented on a black background. A red penalty circle was displayed first, followed by a green target. The penalty circle was light red with a small bright red circle indicating the center, and displayed near the center of the screen. To prevent subjects from using preplanned motor movements the whole stimulus configuration was shifted by a randomly chosen amount in each trial; the shifts in x and y were chosen independently from a uniform distribution over the range $\pm 44 \text{ mm}$. The green target was transparent so that the overlap with the penalty circle was visible. The target and penalty circles had radii of 32 pixels/9 mm. The green target appeared at one of six possible positions, horizontally displaced from the penalty circle (Fig. 3).

5.3. Procedure

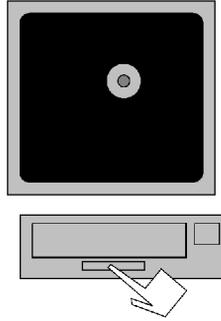
Each trial followed the procedure illustrated in Fig. 4. A fixation cross indicated the start of the trial. The subject was required to depress the space bar of the keyboard with the same finger that s/he would later use to touch the screen. The trial would not begin until the space bar was depressed; the subject was required to hold the space bar down until after the green target appeared. Next, the red penalty circle was displayed, indicating that the subject should be prepared to move shortly. This was followed by the green target after an interval of 500 ms. Only after the appearance of the target was the subject allowed to release the space bar and initiate a movement toward the touch screen (allowing us to measure the time of movement initiation). After the green target was displayed, subjects were required to touch the screen within 750 ms or they would incur a loss of 500 points. If the subject touched the

1. Start of trial:

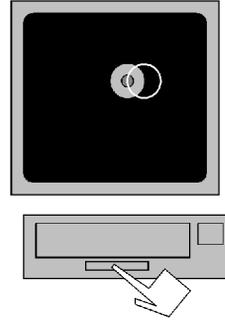
Display of fixation cross (2 sec)

**2. Target display (red):**

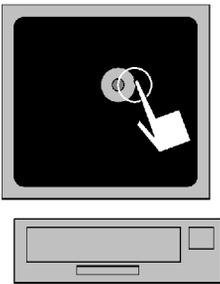
Display of red target 500 ms before display of green target

**3. Target display (green):**

Display of green target for maximal 750 ms

**4. Touch of screen:**

Recording of touch location, reaction and movement time

**5. Feedback on target value:****6. End of trial:**

Display of current cumulative score

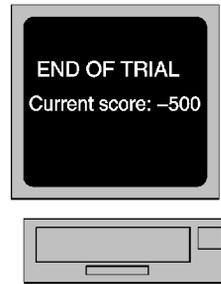


Figure 4. Sequence of events during a single trial.

screen within the area of the red or the green target, the target that was hit (or both) ‘exploded’ graphically. Then, the points awarded for that trial were shown, followed by the subject’s total accumulated points for that session. A hit on the green target gained 100 points (a loss of -100). The penalty for touching the red penalty circle was constant during a block of trials, and could amount to 0, 100 or 500 points. If the screen was touched in the region of overlap between the target and penalty circle, then the reward and penalty were combined. If a subject anticipated the target display, releasing the space bar before, or within 100 ms after, display of the green target, the trial was abandoned and repeated later during the block.

Each block of trials consisted of 10 repetitions of each of the six target locations, presented in random order. A single experimental session consisted of a touch

screen calibration, 30 practice trials, and three blocks of 60 trials corresponding to the three penalty levels, presented in random order.

5.4. Subjects and instructions

6 subjects participated in the experiment. The subjects were four male and two female members of the Department of Psychology at New York University. All participants were right-handed, had normal or corrected-to-normal vision and ranged in age from 27 to 32 years. All subjects had given their informed consent prior to testing and were paid for their participation. All were unaware of the hypotheses under test. Subjects ran two sessions of 210 trials (30 practice trials, 3 blocks of 60 trials) which took approximately 30 minutes per session. Subjects were informed of the payoffs and penalties for each block of trials. All used their right index finger to depress the space bar at the start of a trial and to touch the touch screen. Subjects were told that the overall score over the two sessions would result in a bonus payment of 25 cents per 1000 points so as to motivate fast, accurate responses.

5.5. Data analysis

For each trial we recorded the *reaction time* (the interval from stimulus display to release of the space bar), the *movement time* (the interval from release of the space bar to touching the screen), the *screen position* that was hit and the *score*. Trials in which the subject left the space bar less than 100 ms after display of the green target, or hit the screen more than 750 ms after display of the green target were excluded from the analysis.

During the first session, two subjects reported difficulty with the required time limit of 750 ms. Subjects' responses were more variable in the first session (average spatial variance $\sigma^2 = 82.2 \text{ mm}^2$ in the first session, $\sigma^2 = 16.4 \text{ mm}^2$ in the second session) and scores were significantly lower in the first session (7012 ± 3818 during the first session, 14617 ± 975 during the second). After the second session, all subjects reported that they had no difficulty performing the task. Therefore, we only analyzed the data from the second session.

Each subject contributed approx. 180 data points, i.e. 10 repetitions per condition. Endpoint positions (x_p, y_p) were recorded relative to the center of the red penalty circle (Fig. 3). The six position of the green target were $x_{green,i} = -18, -13.5, -9, 9, 13.5$ and 18 mm and $y_{green,i} = 0$ ($i = 1 \dots 6$). To test whether the recorded endpoint differed from the target center, we also calculated the endpoint relative to the target: $\Delta x_i = x_p - x_{green,i}$ and $\Delta y_i = y_p - y_{green,i}$. A value of $\Delta x_i > 0$ indicates that the recorded endpoint was to the right of the target center; a value of $\Delta y_i > 0$ indicates that the recorded endpoint was above the target center.

Data were analyzed individually for each subject as a 2-factor, repeated measures ANOVA. The factors were the target position and the penalty level. This form of ANOVA was calculated for three dependent measures: Δx_i , reaction time and

movement time. In general, data are reported as mean \pm standard error. As we did not expect Δy_i to differ significantly from zero for any of the conditions, data were analyzed across subjects using a 3-factor repeated measures ANOVA. Two factors (target position and penalty level) varied within subjects, the third factor (subject) tested for differences between subjects. Statistically significant results are printed in italics.

6. RESULTS

1073 recorded endpoints were included in the analysis (i.e. only 7 responses were omitted for being late). As subjects differed significantly in their spatial endpoint variability (Levene test, $F(5, 2140) = 14.194$, $p < 0.001$) the data were analyzed individually for each subject. Results are reported in Table 1.

As displayed in Table 1, reaction and movement times differed between subjects, but were consistent across conditions for each subject (note the small standard deviations in Table 1). The results of the statistical analysis for each subject confirmed that reaction and movement times did not differ significantly across conditions (data not reported here; contact authors for details). Thus, the time limit of 750 ms was short enough to force the subjects to respond rapidly and consistently across conditions.

Our model (Section 3) predicts a shift of the endpoint away from the red penalty circle. This shift should be stronger for target positions closer to the penalty circle, and for higher penalty values. We did not expect any shift in the endpoint in zero-penalty conditions. Due to the symmetry of the stimulus configuration, we also did not expect a vertical shift.

In agreement with what we expected, subjects' endpoints were distributed without significant bias in the y -direction. The deviation in the y -direction, Δy , neither

Table 1.

Experimental results

Subject	σ^2 (mm ²)	Reaction time	Movement time	$F(10, 90)$	p
S1	8.0555	211 \pm 17 ms	412 \pm 23 ms	6.172	$p < 0.001$
S2	15.233	232 \pm 23 ms	426 \pm 19 ms	1.794	$p = 0.073$
S3	15.6025	264 \pm 18 ms	352 \pm 29 ms	4.790	$p < 0.001$
S4	15.9275	277 \pm 17 ms	331 \pm 19 ms	3.096	$p = 0.002$
S5	22.9555	237 \pm 35 ms	356 \pm 17 ms	7.179	$p < 0.001$
S6	18.786	203 \pm 15 ms	403 \pm 12 ms	1.985	$p = 0.044$

Data reported for the six subjects individually; spatial variability (σ^2), reaction and movement times (\pm one standard deviation) computed by averaging across all conditions (~ 180 data points per subject). To test whether movement endpoints shifted away from the penalty region with increasing penalty level and with smaller target distance, a 2-factor, repeated measures ANOVA was performed (see *Data analysis* for details), F - and p -values for the interaction between 'target position' and 'penalty level' are reported on the right.

varied significantly among the six subjects ($F(5, 54) = 2.034$, $p = 0.088$), nor across target positions ($F(5, 50) = 0.555$, $p = 0.734$) or penalty levels ($F(5, 53) = 2.662$, $p = 0.079$).

Figure 5 shows the mean horizontal shifts of the endpoints (Δx) as a function of the target position and penalty, displayed separately for each subject. For all subjects, the endpoint shifted away from the red penalty circle (located at $x = 0$ mm); and this shift was largest for the highest penalty value and the target positions closest to the penalty circle.

The results of the statistical analysis are consistent with these observations. Table 1 contains the F - and p -values for the interaction between target position and penalty value. The interaction between target position and penalty value was significant for five of the six subjects (see Table 1), confirming the expected shift of endpoints with increasing penalty values. The missing significant interaction for subject S2 most likely results from his general tendency to ‘avoid the red target’, which also persisted during the penalty = 0 condition (see middle panel in the top row in Fig. 5).²

Qualitatively, the observed pattern of movement endpoints followed the predictions of our model. To make a quantitative prediction, we again assumed that the recorded endpoints were distributed around the optimal aiming point according to a Gaussian distribution with a width of σ , where this variability is independent of penalty value and target location. This is justified by the lack of effect of experimental conditions on reaction and movement times. Thus, we estimated the endpoint reliability σ by computing the standard deviations σ_x , and σ_y in the x - and y -directions across all conditions separately for each subject. The variabilities σ_x^2 and σ_y^2 did not differ systematically ($F_{\max} < 1.5$ for all subjects), consistent with the assumed distribution in equation (8). The standard deviation σ was estimated separately for each subject by pooling data for horizontal and vertical deviations. The resulting estimates are reported in Table 1. All other parameters required by our model (target size, location and penalty value) are specified by the design of our experiment. Thus, with *no* free model parameters other than σ (computed directly from the data), we were able to calculate a prediction of our model for each subject and experimental condition. The optimal aiming point as predicted by equation (7) was computed for different target locations and penalties. The predicted values of x_{opt} are plotted as the curves in Fig. 5, allowing a direct comparison with the experimental data.

It is clear from Fig. 5 that for all subjects our model compares well with the experimental data. Our model predicts shifts in the movement endpoints that are close to those observed experimentally; and the trends as a function of penalty level and target position are quite similar. Also note that the model predicts a shift of the optimal aiming point in the horizontal direction. In accordance with the model prediction, subjects did not alter their aiming point in the vertical direction throughout the experiment (see above), but only shifted their aiming point away from the penalty area in the horizontal direction.

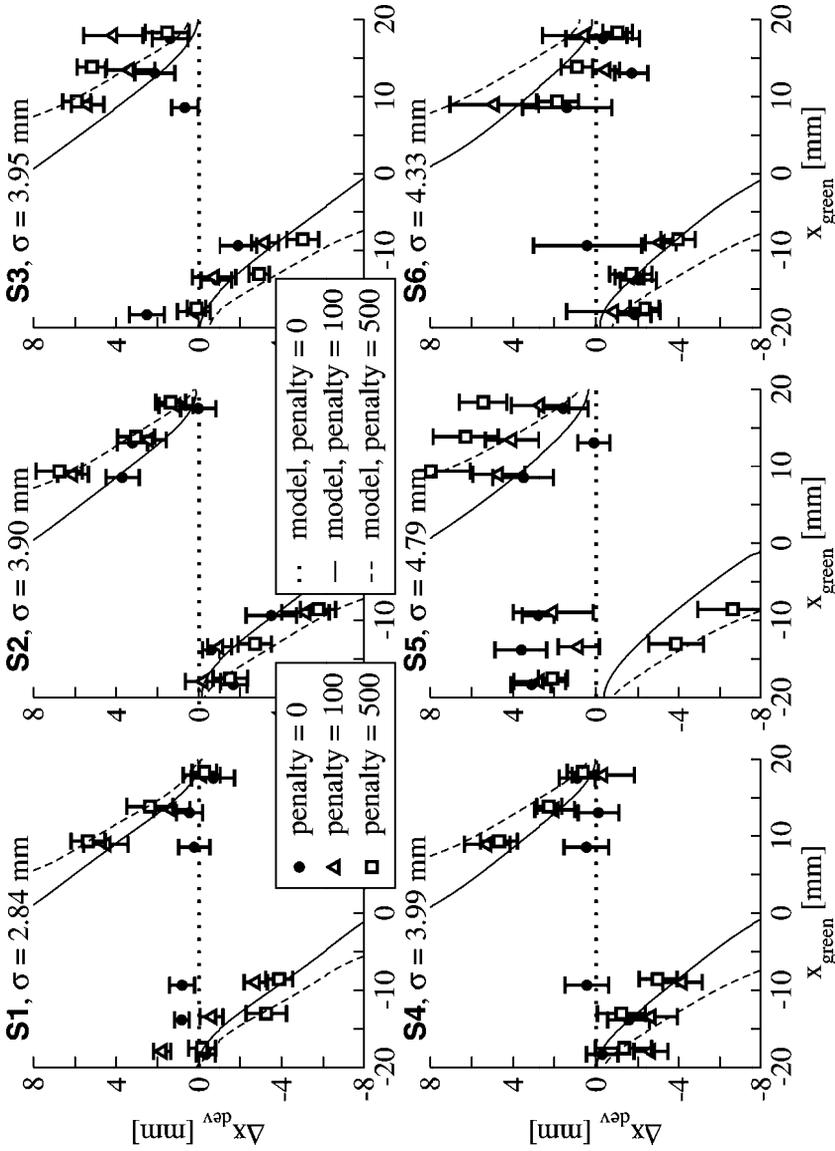


Figure 5. Results. Average horizontal deviations of the movement endpoints from the target as a function of target position and penalty (the penalty circle was located at $x = 0$ mm) for each of the six subjects. Error bars represent plus-or-minus one standard error, computed from 10 data points per condition. Curves represent the predicted optimal shifts based on the model of Section 3 and motor uncertainty estimated individually for each subject using the subject's own variance.

Table 2.

Comparison of the subjects' performance to that of an ideal observer

Subject	Score (experiment)	Optimal performance	Efficiency (%)
S1	15700	17058 \pm 425	<i>92.04%</i>
S2	14400	14887 \pm 668	96.73%
S3	15500	14769 \pm 650	104.95%
S4	14600	14751 \pm 662	98.98%
S5	12700	12625 \pm 804	100.59%
S6	14800	13841 \pm 816	106.93%

Optimal performance as predicted by our model (equation (7)) for each subject, employing the individual variabilities σ^2 as given in Table 1. The model predictions were generated by simulating the actual experiment (100 runs); data reported as average \pm standard deviation. The actual performance is considered to differ significantly if the subject's score differs more than two standard deviations (significant deviations in *italic*).

To compare the subject's performance to that of an 'optimal performer', we simulated the experiment for an 'optimal performer' and compared the overall winnings of the optimal performer to the actual subject's performance (Table 2). Note that subjects with higher variabilities σ^2 (subjects S5 and S6, see Table 1) miss the green target or collect a penalty from accidentally hitting the red target more often. This outcome is consistent with the model: an ideal performer with lower variability will, on average, win more than an ideal performer with higher variability. The model predicts lower overall winnings for more variable subjects. Table 2 shows that all six subjects' performances deviate less than 8% from optimal performance.

7. DISCUSSION

We have presented a model of motor planning based on statistical decision theory. The theory is applicable to tasks where there are explicit costs associated with the outcomes of actions and where the uncertainty inherent in carrying out a planned movement may alter the consequences for the planner. Under these circumstances, an optimal choice of motor strategy is achieved by minimizing an expected loss function (equation (6)) that takes into account motor uncertainty, biomechanical costs, and costs associated with time limits imposed on the mover.

We used our model to predict where subjects should aim in a simple task involving hitting a target region on a computer screen when a second penalty region was also present. The rewards and penalties associated with hitting either region were monetary and the penalties associated with the penalty region and its position relative to the target region were both varied. We measured the subjects' motor uncertainty as part of the experiment and used these measurements to predict the optimal aim point for each location of the penalty region and magnitude of penalty.

We tested these predictions experimentally. Subjects attempted to maximize their monetary reward in a task where they had to tap a touch screen on every trial. We found that subjects shifted their mean impact point away from the middle of the penalty region, with larger shifts for closer penalty circles and higher magnitudes of penalty. These results suggest that humans take both costs and their own movement uncertainty into account in movement planning. Actual performance was in good agreement with the predictions of the model.

It is likely that additional factors such as fatigue, aversion to errors, motivation, and self-confidence shape and influence a subject's performance. If we manage to quantify these factors, we can easily include them into our approach and study how these 'non-physical' factors interact with the subject's desire to make money at the task. It may be objected that such psychological factors cannot be put into units of money, but the fact is that the subject, by his/her actions, is in effect doing just that: translating fatigue, boredom, or self-confidence into dollars by his or her choice of strategy.

An important implication of the cost equation (equation (6)) is that the subject will trade off biomechanical factors against other possible losses and gains. S/he will alter her or his movement strategy so as to risk physical damage or pain if the perceived gain outweighs the risk (something that the professional soccer player knows too well as s/he plays one last game despite bad knees). Models of motor planning that consider only biomechanical factors and a specific target point cannot readily account for such trade-offs.

It is evident that the model we propose is readily expanded to tasks with multiple penalty regions and multiple target regions with more complex limits on timing than the simple timeout rule we impose. In the experiment reported here, we only predicted and measured the endpoint of the subject's motion trajectory, not the full motion trajectory. We emphasize that the theory, as developed above, is intended to predict not just the endpoint of the trajectory, but the trajectory itself. It would be very natural to model tasks involving penalty regions and reward regions distributed arbitrarily in three-dimensional space near the subject. We suggest that games involving rewards and penalties can potentially evoke a far richer range of motor behavior than tasks typically used in the laboratory.

Statistical decision theory (Blackwell and Girshick, 1954; Ferguson, 1967) is applicable to modeling perception and action under risk (Maloney, 2002). Bayesian decision theory is a branch of statistical decision theory that is applicable when there is uncertainty surrounding the state of the external environment. In all of the 'Wilhelm Tell' tasks we have discussed, there is little or no uncertainty concerning the locations of target and penalty regions or costs associated with each. It would not be difficult to extend the scope of the proposed model to include tasks in which the exact location, shape or cost associated with these regions was in doubt.

We could, for example, leave out the visual feedback about the location of penalty targets, by keeping the penalty targets invisible. Over a sequence of trials the subject might then hit the invisible penalty target and will come up with expectations about

the local distribution of ‘dangerous’ areas. The more evidence a subject has about a penalty target in a certain area the more likely s/he will account for this area when choosing his or her aiming point. The resulting Bayesian decision theoretic model of motor planning is a natural counterpart to Bayesian approaches to modeling visual perception (Knill and Richards, 1996; Maloney, 2002; Mamassian *et al.*, 2002). Vetter and Wolpert (2000) have presented a Bayesian model of motor planning by using a probabilistic framework to demonstrate that subjects alter their pointing movements when they change their expectations about the experimental context.

It should be noted that ours is not the only study demonstrating effects of ‘non-physical’ or cognitive factors on performance in speeded motor tasks. Speeded pointing performance depends on perceived effort (Rosenbaum and Gregory, 2002) and is impaired during dual task paradigms requiring overt attention (see Castiello, 1999, for review). It has also been demonstrated that information from preceding trials (de Lussanet *et al.*, 2001, 2002) and expectations about future events, for instance expected changes in the experimental environment (Vetter and Wolpert, 2000), are integrated into the motor plan. Subjects take uncertainty into account when planning movements (e.g. Adolph and Avolio, 2000). Furthermore, Sabes and colleagues (Sabes and Jordan, 1997; Sabes *et al.*, 1998) have evidence of subjects using details of motor uncertainty in planning movements around obstacles. What distinguishes the current study from these is the development of a theory of ideal motor behavior and a demonstration of such behavior in a task where costs and rewards are clearly specified.

Overall, actual behavior matched ‘optimal’ behavior as predicted by the model (Table 2). The subjects’ performance was always within 8% of optimal performance. It is implausible that human performance is precisely optimal in any visual or motor task but much can be learned by determining exactly how subjects fail. By comparing human performance to that of an ideal movement planner, we may be able to demonstrate limitations of movement planning, much as ideal observer analysis has been so successful at probing sites of information loss in the visual system (Geisler, 1989).

But what about Wilhelm Tell? He was indeed an experienced marksman and managed to shoot the apple from his son’s head. While doing so he still considered the possibility of missing his aim and harming his son. That is why he saved a second arrow for Gessler.

Acknowledgements

We thank Allen Ingling and Christopher Currie for technical support. This work was supported by NIH grant EY08266, Human Frontier Science Program grant RG0109/1999-B, and by an Emmy-Noether fellowship (DFG) to J. T. We also thank Karen Adolph, Marty Banks, Huseyin Boyaci, Sergei Gephstein and Jamie Hillis for helpful comments on an earlier draft. We especially thank Maria Felicita Dal Martello for the illustration in Fig. 1.

NOTES

1. The manipulation of the time limit may also reduce visual feedback control of the movement and increase visual uncertainty of target location. For the purposes of this article, each of these effects may be included with movement variability, altering none of the theoretical results or conclusions presented here.

2. In all six subjects, deviations in the x -direction differed significantly across target position, i.e. the first main effect, target position, was significant. Due to the lack of shift for distant targets and zero penalties, there was no significant second main effect of penalty for any of the subjects (data not reported here; contact authors for more details).

REFERENCES

- Adolph, K. E. and Avolio, A. M. (2000). Walking infants adapt locomotion to changing body dimensions, *J. Exper. Psychol.: Human Perception and Performance* **26**, 1148–1166.
- Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, 2nd edn. Springer, New York.
- Blackwell, D. and Girshick, M. A. (1954). *Theory of Games and Statistical Decisions*. Wiley, New York.
- Brainard, D. H. (1997). The psychophysics toolbox, *Spatial Vision* **10**, 433–436.
- Castiello, U. (1999). Mechanisms of selection for the control of hand action, *Trends in Cognitive Science* **3**, 264–271.
- Castiello, U. (2001). The effects of abrupt onset of 2-D and 3-D distractors on prehension movements, *Perception and Psychophysics* **63**, 1014–1025.
- Connolly, J. D. and Goodale, M. A. (1999). The role of visual feedback of hand position in the control of manual prehension, *Exper. Brain Res.* **125**, 281–286.
- Dean, J. and Brüwer, M. (1994). Control of human arm movements in two dimensions: paths and joint control in avoiding simple linear obstacles, *Exper. Brain Res.* **97**, 497–514.
- Desmurget, M. and Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements, *Trends in Cognitive Science* **4**, 423–431.
- Dornay, M., Uno, Y., Kawato, M. and Suzuki, R. (1996). Minimum muscle-tension chance trajectories predicted by using a 17-muscle model of the monkey's arm, *J. Motor Behavior* **2**, 83–100.
- Ferguson, T. S. (1967). *Mathematical Statistics: A Decision Theoretic Approach*. Academic Press, New York.
- Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement, *J. Exper. Psychol.* **47**, 381–391.
- Flash, T. and Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model, *J. Neurosci.* **5**, 1688–1703.
- Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discrimination, *Psychol. Rev.* **96**, 1–71.
- Green, D. M. and Swets, J. A. (1966/1974). *Signal Detection Theory and Psychophysics*. Wiley, New York. Reprinted 1974, Krieger, New York.
- Hamilton, A. F. C. and Wolpert, D. M. (2002). Controlling the statistics of action: obstacle avoidance, *J. Neurophysiol.* **87**, 2434–2440.
- Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning, *Nature* **394**, 780–784.
- Howard, L. A. and Tipper, S. P. (1997). Hand deviations away from visual cues: indirect evidence for inhibition, *Exper. Brain Res.* **113**, 144–152.

- Kaminsky, T. and Gentile, A. M. (1986). Joint control strategies and hand trajectories in multijoint pointing movements, *J. Motor Behavior* **18**, 261–278.
- Knill, D. C. and Richards, W. (1996). *Perception as Bayesian Inference*. Cambridge University Press, Cambridge, UK.
- de Lussanet, M. H., Smeets, J. B. and Brenner, E. (2001). The effect of expectations on hitting moving targets: influence of the preceding target's speed, *Exp. Brain Res.* **137**, 246–248.
- de Lussanet, M. H., Smeets, J. B. and Brenner, E. (2002). The relation between task history and movement strategy, *Behavioral Brain Research* **129**, 51–59.
- Maloney, L. T. (2002). Statistical decision theory and biological vision, in: *Perception and the Physical World*, Heyer, D. and Mausfeld, R. (Eds), pp. 145–189. Wiley, New York.
- Mamassian, P., Landy, M. S. and Maloney, L. T. (2002). Bayesian modeling of visual perception, in: *Probabilistic Models of the Brain: Perception and Neural Function*, Rao, R., Lewicki, M. and Olshausen, B. (Eds), pp. 13–36. MIT Press, Cambridge, MA.
- Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E. and Smith, J. E. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements, *Psychol. Rev.* **95**, 340–370.
- Mon-Williams, M., Tresilian, J. J., Coppard, V. L. and Carson, R. G. (2001). The effect of obstacle position on reach-to-grasp movements, *Exp. Brain Res.* **137**, 497–501.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies, *Spatial Vision* **10**, 437–442.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (1992). *Numerical Recipes in C. The Art of Scientific Computing*, 2nd edn. Cambridge University Press, Cambridge, UK.
- Rosenbaum, D. A. and Gregory, R. W. (2002). Development of a method for measuring movement related effort. Biomechanical considerations and implications for Fitts' law, *Exp. Brain Res.* **142**, 365–373.
- Rosenbaum, D. A., Loukopoulos, L. D., Meulenbrock, R. J., Vaughan, J. and Engelbrecht, S. E. (1995). Planning reaches by evaluating stored postures, *Psychol. Rev.* **102**, 28–67.
- Rosenbaum, D. A., Meulenbrock, R. J., Vaughan, J. and Jansen, C. (1999). Coordination of reaching and grasping by capitalizing on obstacle avoidance and other constraints, *Exp. Brain Res.* **128**, 92–100.
- Rosenbaum, D. A., Meulenbrock, R. J., Vaughan, J. and Jansen, C. (2001). Posture-based motion planning: applications to grasping, *Psychol. Rev.* **108**, 709–734.
- Sabes, P. N. and Jordan, M. I. (1997). Obstacle avoidance and a perturbation sensitivity model for motor planning, *J. Neurosci.* **17**, 7119–7128.
- Sabes, P. N., Jordan, M. I. and Wolpert, D. M. (1998). The role of inertial sensitivity in motor planning, *J. Neurosci.* **18**, 5948–5957.
- Soechting, J. F. and Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man, *J. Neurosci.* **1**, 710–720.
- Soechting, J. F., Buneo, C. A., Herrmann, U. and Flanders, M. (1995). Moving effortlessly in three dimensions: Does Donders' Law apply to arm movement? *J. Neurosci.* **15**, 6271–6280.
- Tipper, S. P., Howard, L. A. and Jackson, S. R. (1997). Selective reaching to grasp: evidence for distractor interference effects, *Visual Cognition* **4**, 1–38.
- Uno, Y., Kawato, M. and Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement: Minimum torque-change model, *Biol. Cybernetics* **61**, 89–101.
- Vetter, P. and Wolpert, D. M. (2000). Context estimation for sensorimotor control, *J. Neurophysiology* **84**, 1026–1034.