



Review

Decision-theoretic models of visual perception and action

Laurence T. Maloney*, Hang Zhang

Department of Psychology, New York University, United States
 Center for Neural Science, New York University, United States

ARTICLE INFO

Article history:

Received 16 August 2010

Received in revised form 27 September
 2010

Keywords:

Perception
 Action
 Statistical decision theory
 Bayesian decision theory
 Ideal observer models
 Gain function
 Loss function
 Prior
 Likelihood

ABSTRACT

Statistical decision theory (SDT) and Bayesian decision theory (BDT) are closely related mathematical frameworks used to model ideal performance in a wide range of visual and motor tasks. Their elements (gain function, likelihood, prior) are readily interpretable in terms of information available to the observer. We briefly describe SDT and BDT and then review recent work employing them as models of biological perception or action. We emphasize work that employs gain functions and priors as independent or dependent variables.

At one extreme, Bayesian decision theory allows the experimenter to compute ideal performance in specific tasks and compare human performance to ideal (Geisler, 1989). No claim is made that visual processing is in any sense “Bayesian”. At the other extreme, researchers have proposed Bayesian decision theory as a process model of “perception as Bayesian inference” (Knill & Richards, 1996). We end by discussing how possible ideal models are related to imperfect, actual observers and how the “Bayesian hypothesis” can be tested experimentally.

© 2010 Published by Elsevier Ltd.

1. Introduction

Statistical decision theory (SDT) emerged with the publication of Blackwell & Girshick's *Theory of Games and Statistical Decisions* in 1954. An immediate stimulus to its development was the *Theory of Games and Economic Behavior* by von Neumann and Morgenstern (1944/1953) and, like game theory, SDT is normative: it is a mathematical method for selecting optimal actions under conditions of uncertainty. On each of a series of turns in SDT a player gains instantaneous information about an uncertain environment and then selects an action. The choice of action determines whether the player merits reward or incurs punishment.

Bayesian decision theory (BDT) is a special case of SDT. Both methods are widely employed in mathematical statistics (Berger, 1985; Ferguson, 1967; Gelman, Carlin, Stern, & Rubin, 2003; Jaynes, 2003; O'Hagan, 1994) and pattern classification (Duda, Hart, & Stork, 2000). In recent years, BDT has been more and more frequently used in developing models of biological perception and action (Knill & Richards, 1996; Maloney, 2002; Mamassian, Landy, & Maloney, 2002; Yuille & Bülthoff, 1996), in part because its mathttic

tions. The three sets are \mathbf{W} , the states of the world, \mathbf{X} , the possible sensory states, and \mathbf{A} , possible actions (Fig. 1A). On every “turn” the world is in some specific state, $w \in \mathbf{W}$, unknown to the observer. The observer is given access to a sensory state $x \in \mathbf{X}$,

and must decide what action, $a \in \mathbf{A}$, to select. The interpretation of these elements is very flexible. The state of the world may be the distance to a specific object or the intrinsic color of a surface. Actions could include estimates of depth, a motor program specified as a pattern of neural activity over time, or a decision between fight and flight. Signal detection theory (Green & Swets, 1966/1974) is an application

We use upper-case X to denote the particular sensory state available to the observer on a specific occasion and lower-case x to denote sensory states in general, the latter analogous to “the people you know”, the former to “your good friend Dennis” who just walked into your office.

* Corresponding author. Address: Department of Psychology, 6 Washington Place, 2nd Floor, New York, NY 10003, United States. Fax: +1 212 995 4349.
 E-mail address: ltm1@nyu.edu (L.T. Maloney).




Fig. 1. (A) The elements of statistical decision theory. The three vertices correspond to \mathbf{W} , the possible states of the world, \mathbf{X} , the possible sensory states, and \mathbf{A} , the available actions. The three edges correspond to the gain function, $G(a, w)$, the likelihood function, $f(x|w)$, and the decision rule, $d(x)$ where $x \in \mathbf{X}$ denotes a sensory state, $a \in \mathbf{A}$, a particular action and $w \in \mathbf{W}$, a particular state of the world. (B) Equal variance Gaussian signal detection theory. The distribution of the sensory state X depends on the state of the world. The two possible world states are S ("signal present") and \bar{S} ("signal absent") and the distributions are Gaussian with equal variance but differing in mean by d' (Green & Swets, 1966/1974).

of SDT that nicely captures all the key ideas and we will use a particular signal detection example to illustrate the key ideas of SDT (and, later, BDT) as we introduce them.

The states of the world in signal detection theory are just "signal present" and "signal absent" denoted as $\mathbf{W} = \{S, \bar{S}\}$, the sensory states are any real magnitude that we refer to as the *strength* of the signal $x \in \mathbf{X} = \mathbb{R}$ and the possible actions are simply to say "signal present" or "signal absent", denoted as $\mathbf{A} = \{s, \bar{s}\}$.

There are three functions that serve to complete the description of SDT (Fig. 1A). The first is the *likelihood function* $f(x|w)$, the probability density of sensory states contingent on the state of the world which, as written, links the sensory information to the state of the world.² Remarkably, it can be shown that the likelihood function captures all of the sensory information relevant to estimating the state of the world (Berger & Wolpert, 1988; Maloney, 2002), a result known as the Likelihood Principle.

In Fig. 1B we plot the two possible likelihood functions of Gaussian equal variance signal detection theory, one for the world state S ("signal present") and one for the world state \bar{S} ("signal absent"). These are the probability density functions that X may have, depending on the state of the world

$$f(x|\bar{S}) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$$

$$f(x|S) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(x-d')^2}{2}} \quad (1)$$

where d' is the mean of the distribution when the signal is present. One possible value of X is marked on Fig. 1B and, while it could have arisen from either distribution, it seems intuitively plausible that it arose from the world state "signal present".

The second function is the *gain function* $G(a, w)$ that determines the gain or loss experienced by the observer on a particular trial. It is also referred to as *loss function* or *value function* in the literature. Losses are just negative gains and vice versa. A possible gain function for the simple signal detection theory model we consider is tabulated in an inset to Fig. 1B. With this gain function, the signal detection theory observer gains one unit if she correctly names the state of the world and otherwise receives nothing.

The third function is the *decision function* $a = d(x)$ that captures the strategy of any particular SDT observer. The decision function maps the sensory state (the only novel information available on a particular trial) to an action. This modest function is intended to model all of perceptual and cognitive processing. In signal detection theory, the choice of a rule $d(x)$ applied to the signal strength X completely specifies the signal detection observer.

We will add one more function, the prior distribution of states of the world, below and, once we do so, SDT will transmute into

² The likelihood function is often written as $L(w|x) = f(x|w)$ to emphasize that it provides information about possible states of the world given a known sensory state x . We will, however, continue to use $f(x|w)$.

BDT, a special case of SDT where the observer has access to the prior distribution of states of the world. For now though, we will consider what we can say about different choices of decision functions in SDT without a prior. One reason to do so is to develop a better understanding of SDT. A second reason is to examine what we can say about different decision rules even when the prior distribution is not known.

We characterize any decision rule $d(x)$ by evaluating its expected gain in each world state

$$EG[d|w] = \int_{-\infty}^{\infty} G(d(x), w) f(x|w) dx \quad (2)$$

The equation is readily interpreted: the state of the world determines the probability density that each possible sensory state can occur through the likelihood function $f(x|w)$; the decision function maps the sensory state to an action $a = d(x)$, and the observer receives the gain $G(a, w)$, weighted by the likelihood summed across all possible sensory events. In the signal detection theory example, with the gain function shown in Fig. 1B, Eq. (2) can be written as

$$\begin{aligned} EG[d|S] &= \int_{-\infty}^{\infty} G(d(x), S) f(x|S) dx = p[d(X) = S] \\ EG[d|\bar{S}] &= \int_{-\infty}^{\infty} G(d(x), \bar{S}) f(x|\bar{S}) dx = p[d(X) = \bar{S}] \end{aligned} \quad (3)$$

and, in the ordinary terminology of signal detection theory, the two rightmost probabilities are the probability of a “hit” (correctly identifying the signal when present) and the probability of a “correct rejection” (correctly identifying the signal when absent) denoted $p[\text{Hit}]$ and $p[\text{CR}]$, respectively. We can summarize any decision rule by $EG[d|w]$ and, for the signal detection theory example, we can plot this summary as a plot of $EG[d|S]$ versus $EG[d|\bar{S}]$. That is, we plot $p[\text{Hit}]$ versus $p[\text{CR}]$. The range of expected gain on both axes in this case is 0–1 and gain is synonymous with probability correct. We refer to the resulting plot as a *gains plot* and the point plotted for each decision rule as the *gains plot* for that rule. For any rule $d(X)$ we can compute its gains plot but we cannot guarantee that every point on the gains plot has a rule. We plot some examples of gains plots for decision rules as shown in Fig. 2A. The exact location of the gains plot for each rule depends on the likelihood functions in Fig. 1B through Eq. (3).

The rule $d_1(x)$ always chooses the action s , the rule $d_2(x)$ always chooses the action \bar{s} , the rule $d_3(x)$ chooses the action s precisely when $X > 0.5$, and the rule $d_4(x)$ chooses the action s precisely when $X \leq 0.5$. The third rule is intuitively appealing. If the sensory state is greater than the point where the two distributions cross in Fig. 1B, we choose s and otherwise \bar{s} . If, however, the world state is certain to be s , then d_1 will always earn the maximum possible gain. The rule d_4 in contrast, seems perverse, inferior to the others. As we shall see below, it is.

Given any two rules, say d_2 and d_3 , we can mix them probabilistically by deciding to use d_2 with probability q and otherwise d_3 . We denote the resulting mixture rule as $d_5(x)$. The expected gain for $d_5(x)$ mixture rule in world state S is easy to compute. With probability q we execute rule $d_2(x)$ with expected gain $EG[d_2|S]$ and otherwise (with probability $1 - q$) we execute rule $d_3(x)$ with expected gain $EG[d_3|S]$. The overall expected gain for the world state S is just

$$EG[d_5|S] = qEG[d_2|S] + (1 - q)EG[d_3|S] \quad (4)$$

We can similarly compute

$$EG[d_5|\bar{S}] = qEG[d_2|\bar{S}] + (1 - q)EG[d_3|\bar{S}] \quad (5)$$

The gains plot of the mixture rule d_5 corresponds to a point in Fig. 2A plotted that is on the line joining the points for d_2 and d_3 . Its displacement from d_3 along the line is proportional to q . The point for d_5 is plotted on Fig. 2 under the assumption $q = 0.25$.

The shaded region in Fig. 2A contains the plots of $P[\text{Hit}]$ vs. $P[\text{CR}]$, for all possible rules $d(x)$ including mixture rules. The top-

right edge of this region, marked by a heavy blue curve, is the receiver operating characteristic curve (ROC curve) of signal detection theory (Green & Swets, 1966/1974) slightly disguised as we have plotted $P[\text{CR}]$ on the horizontal axis rather than the more familiar probability of a “false alarm”. A false alarm occurs when the decision rule selects s (“signal present”) when the world state is \bar{S} (“signal absent”), and $P[\text{FA}] = 1 - P[\text{CR}]$. If we switched to $P[\text{FA}]$ we would “flip” the plot left to right, restoring the form of the ROC curve that is likely familiar to the reader. In the form we employ, gain increases as we go to the right or up. The unmarked bottom-left side of the region is a sort of anti-ROC curve. If you take any rule on the ROC curve and simply respond s when the rule dictates \bar{s} and vices versa, you get a rule whose gains plot is on the anti-ROC curve. The rule d_4 is the “anti-rule” to d_3 and vice versa. An observer can only do very badly in a signal detection task if he has the capability to do very well.

2.1. Dominance and admissibility

The decision rule d_3 always has a higher expected gain $EG[d_3|w]$ than decision rule d_4 for all states of the world. Consequently, employing d_3 rather than d_4 always leads to a higher expected gain. We say that one decision rule d_a dominates another d_b precisely when

$$EG[d_a|w] \geq EG[d_b|w] \quad (6)$$

for all $w \in \mathbf{W}$ and, for at least one choice of w , the inequality is strict. In Fig. 2B we illustrate dominance graphically. All the rules whose plotted expected gains fall into the rectangular area are dominated by the rule whose gains plot falls at the top-right vertex of the rectangle. A decision rule d that is dominated by another rule is *inadmissible*. A decision rule that is not dominated by any other rule is *admissible*. The admissible rules in Fig. 2A are precisely those that fall on the top-right frontier marked by a heavy blue curve, the ROC curve. The rules d_1, d_2, d_3 are admissible, d_4 is not and d_5 is not. In the signal detection example, any mixture of two rules with $0 < q < 1$ such as d_5 is inadmissible.


2.2. Minimax criterion

Dominance imposes a partial ordering³ on the decision rules. If one decision rule dominates another then the former offers higher expected gain without further consideration of the state of the world. But any admissible rule neither dominates, nor is dominated by, any other admissible rule. We have no obvious way to choose among the rules whose gain plots fall on the heavy blue curve in Fig. 2A. The *minimax criterion* allows us to select a rule that gives the “best worst case”. We score each rule by identifying the worst it can do, its minimum gain. For example, the minimum gain for decision rule d_1 (always say “signal present”) is, of course, 0 when the signal is absent, world state \bar{S} . A minimax rule (there may be more than one) has the “best worst case”, that is the “maximum minimum gain”.⁴ The rule d_3 is a minimax rule, an outcome that is not completely surprising when we consider that: (i) it is admissible and (ii) the gains and loss for correct and incorrect responses are identical in the two world states.

The minimax criterion makes no use about any information that we might have about the relative probability of states of the world and Savage (1954) criticizes it and its “worst case” emphasis as

³ A set of items is a complete order if the ordering is complete (every item is either greater than, less than, or equal to, every other item) and the ordering is transitive (if $a > b$ and $b > c$ then $a > c$). A partial order is transitive but need not be complete. That is, some pairs of items may not be ordered.

⁴ SDT and BDT are typically presented with gain functions replaced by loss functions, a cosmetic change if we think of gains as just negative losses and vices versa. Then a minimax rule minimizes the maximum loss and the origin of the term minimax is evident. We retain the term “minimax” even though we work with gain functions.



unduly pessimistic. Bayesian to make use of information and states of the world embodied

2.3. Priors

The prior $\pi(w)$ is just a probability distribution over the states of the world, \mathbf{W} . Once we have a prior, we can calculate the Bayes gain for each rule

$$BG(d) = \int_{-\infty}^{\infty} EG[d|w] \pi(w) dw \quad (7)$$

The Bayes gain assigns a score to each rule d and consequently we can order rules by their Bayes gain. If there is a rule⁵ that has a greater Bayes gain than all other rules, it is referred to as a *Bayes rule* (there may be more than one). In BDT we choose a Bayes rule over any other rule.

For the signal detection example, we specify the probability that a signal will be present or absent as a 2-vector $[1 - \pi, \pi]'$ where $\pi = \pi(w)$ is just a discrete form of Eq. (7)

$$\pi EG[d|S] + (1 - \pi) EG[d|\bar{S}] \quad (8)$$

we can rewrite in vector form as

$$G(d) = [EG[d|\bar{S}], EG[d|S]] \begin{bmatrix} 1 - \pi \\ \pi \end{bmatrix} \quad (9)$$

the inner product of a gains vector and a prior vector. Consider all the rules that share the same Bayes gain

$$B = [EG[d|\bar{S}], EG[d|S]] \begin{bmatrix} 1 - \pi \\ \pi \end{bmatrix} \quad (10)$$

where B is a constant. Eq. (10) is the equation of a straight line that is perpendicular to the line containing the vector $[1 - \pi, \pi]'$. This observation gives us a graphical method to identify decision rules that have the same Bayes gain. We draw the prior vector $[1 - \pi, \pi]'$ on the gains plot (dashed arrow in Fig. 2C) and then draw lines orthogonal to a line containing the vector (red dashed lines in Fig. 2C). From Eq. (10) we see that points on a single red dashed line have the same Bayes gain (Eq. (10)) and this *equivalent Bayes gain* increases as the red dashed line moves up or to the right. The point where the red line just touches the convex set of possible gains corresponds to the Bayes rule, the rule that maximize Bayes gain. The rule d_1 is a Bayes rule if the prior is $[0, 1]'$. That is, if S occurs with probability 1 and \bar{S} never occur, then the rule d_1 (always respond s) has the highest possible Bayes gain. Similarly, the rule d_2 is a Bayes rule if the prior is $[1, 0]'$. The rule d_3 is a Bayes rule if the prior is $[0.5, 0.5]'$. With a bit of geometric reasoning we see that, in this

simple case, any admissible rule is a Bayes rule for some choice of prior and any Bayes rule is admissible. See [Maloney \(2002\)](#) for discussions on more complex BDT models and the mathematical conclusions we can draw from them.

2.4. Ordering the rules

The literature concerning Bayesian approaches to biological vision is almost entirely concerned with Bayes rules, rules that have the maximum possible Bayes gain. If we think of an organism as embodying a decision rule then it is appealing to think of the Bayes rule for a given prior as specifying the maximum expected gain possible for any organism when that prior is correct. However, the Bayes criterion can also be used to order rules (organisms) that are distinctly sub-optimal. All of the decision rules sharing a single dashed red line in [Fig. 2C](#) have the same Bayes gain and Bayes gain increases as the red dashed line moves to the right. To compare two rules (organisms) we need only determine the line each is on and then determine which line is more to the top-right. We'll return to this point in a later section, *Imperfectly Optimal Observers*.

2.5. The gain function

We simplified the signal detection example by choosing a very simple and symmetric gain function ([Fig. 1B](#), inset). A different choice of gain function would only transform the axes of [Fig. 2C](#) by a linear transformation.⁶ In [Fig. 2D](#), we replot the gains plot if we set $G(s, S) = 0.5$ but keep $G(\bar{s}, \bar{S}) = 1$ and $G(s, \bar{S}) = G(\bar{s}, S) = 0$. The result is a compression by 0.5 along the vertical axis. The prior vector and the equivalent Bayes lines are unaffected and consequently the Bayes rules in 2C are no longer Bayes rules. We have shifted to a rule that puts more emphasis on correctly classifying the absence of a signal. This outcome is intuitive since correctly classifying the absence of a signal is worth twice as much as correctly classifying the presence of a signal.

3. Modeling biological perception and action

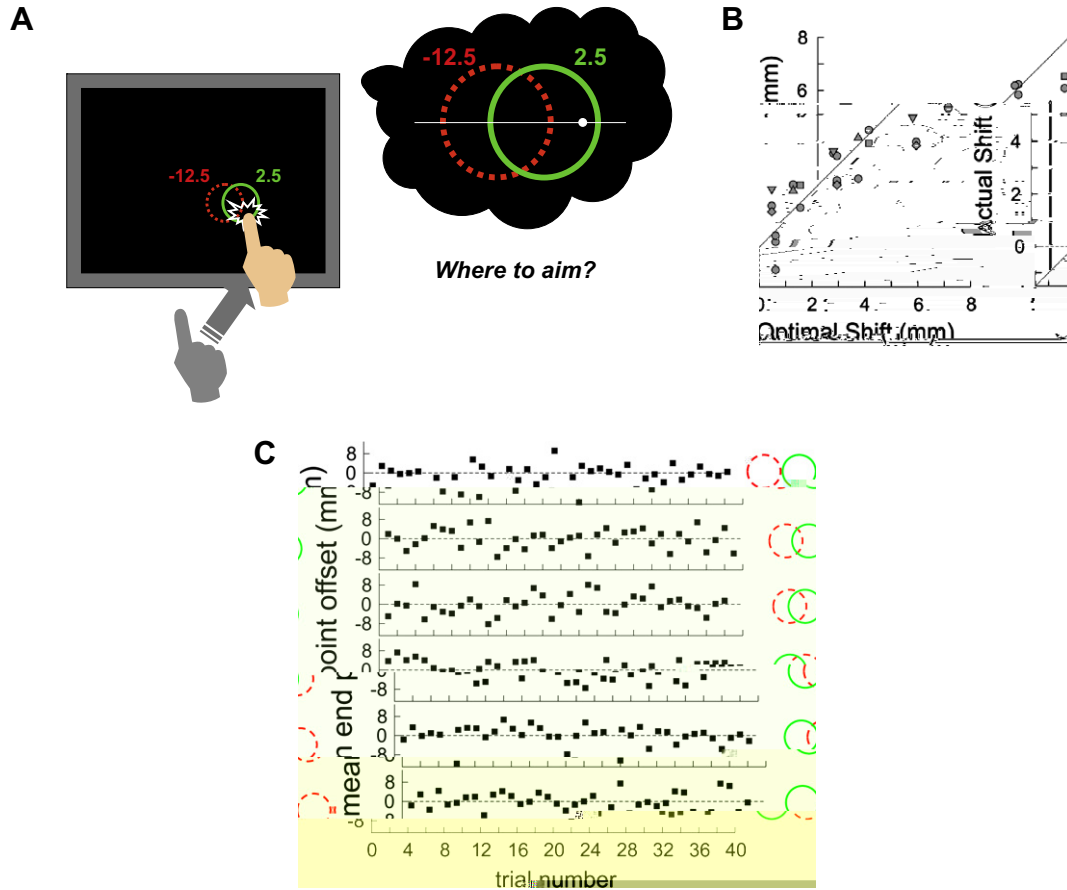


Fig. 3. Asymmetric gain functions in space. (A) A stimulus configuration such as the one shown appears on a computer screen in front of the observer who was instructed to reach out and touch the screen within 700 ms. The gain function is coded by colored circles whose position and relative orientation change from trial to trial. A hit within the solid green circle results in a gain of 2.5 cents, within the dashed red circle, a loss of 12.5 cents. The observer moves rapidly and cannot completely control his movement.

uncertainty spontaneously during training or, less likely, that they knew it before the experiment began. Second, they could combine their knowledge of visuo-motor uncertainty with novel gain functions to choose the aim point that maximizes their expected gain.

The gain function for any task determines the possibilities for reward and punishment; it is remarkable that observers in the tasks of Trommershäuser et al. could come close to maximizing expected gain with such arbitrary choices of gain function.

People do fail in similar tasks when the gain functions are more complex. Wu, Trommershäuser, Maloney, and Landy (2006) pointed out that the stimuli and gain functions used in Trommershäuser et al. (2003a,b) were always symmetric around a line joining the centers of the red and green circles (Fig. 3A, inset). Observers may have deduced that the optimal aim point always lay on this line and this insight may have aided them in planning movement.

Wu et al. (2006) used stimuli with a reward region and two penalties regions differing in magnitude of penalty and found that observers showed patterned failures in selecting aim points. They tended to regress toward the line of symmetry.

3.2. Compensating for altered likelihood functions

Körding and Wolpert (2004) asked observers to reach out and touch a target. Their movement drove a cursor onto a visual target

and nominally the cursor corresponded to the location of their index fingertip. Observers were never allowed to see the hand they reached with. On some trials, the cursor was laterally displaced relative to the actual position of the fingertip. On each movement, the lateral shift of the cursor was randomly drawn from a Gaussian distribution with a mean of 1 cm to the right of the finger and a standard deviation of 0.5 cm. There were four feedback conditions (Fig. 4A). In the σ_0 condition, the position of the cursor was signaled by a white dot whose uncertainty simply reflected the observer's own visuo-motor error. In the σ_M or σ_L conditions, extra uncertainty was introduced by using a cloud of dots with medium (σ_M) or large (σ_L) standard deviation to mark the nominal location of the fingertip. In the final, σ_∞ condition, feedback was withheld. In all conditions, feedback was presented briefly when the fingertip was halfway to the target (Fig. 3A). The endpoint of the reaching movement was presented only for the σ_0 condition.

The question that Körding and Wolpert (2004) addressed was how much the observer should compensate for uncertain visual feedback. Suppose that, on a specific trial, the observer sensed a lateral shift 2 cm to the right. The true lateral shift might be 1.8 cm or 2.2 cm to the right, but the former possibility was more likely than the latter given that the shift was drawn from a prior that was Gaussian with mean 1 cm. Intuitively, the observer's compensation for the 2 cm error should regress toward 1 cm and the degree of regression depends on condition.

This trend was observed in human observers. In Fig. 4B, the mean deviation of the endpoint of the cursor from the target is plotted against the true lateral shift for a typical observer for each of the four conditions. If the sensed lateral shift were fully compensated, the mean deviation should have been 0 and the results for the σ_0 are not far from this limiting case. In contrast, in the σ_∞ condition, the observer failed to compensate or nearly so. As Fig. 4B shows, the higher the uncertainty, the more weight on the prior, the larger the slope. Note that, except for the σ_0 condition, the observer had no opportunity to progressively learn the appropriate weight for a specific feedback condition, because no feedback was provided for the final position of the finger.

Possibly as a consequence, the likelihood functions inferred from the Bayesian observer model did not agree with the actual likelihood functions. The observers' estimates of the standard deviations of the halfway visual feedback were inferred to be 0.67 and 0.8 cm, respectively for the σ_M or σ_L conditions, much smaller than the actual standard deviations, 1 and 2 cm. The picture that emerges is that of an observer whose performance changes from condition to condition in qualitative agreement with BDT but who is effectively using erroneous estimates of likelihood functions.

3.3. Trading information for accuracy

Several studies have focused on whether people could choose the temporal parameter that maximizes expected gain of the movement. In

observers attempted to touch two targets one after another within an overall time limit (Wu, Dal Martello, & Maloney, 2009; Zhang, Wu, & Maloney, 2010). For example, in Fig. 6, the observer first touched the blue circle then the green circle. The two targets corresponded to the same (left) or different (right) rewards. Allocating more movement time to a target would plausibly increase the probability of hitting it and earning the corresponding reward. However, even when the second target was five times more valuable than the first target, observers still allocated slightly more time to the first target.

3.4. Asymmetric gain functions in time

A recent study considered tasks analogous to those of Trommershäuser et al. (2003a,b, 2008) but with gain functions that are specified in the temporal domain (Hudson, Maloney, & Landy, 2008). Before the formal experiment, observers completed extensive training of reaching at specific movement times. During this initial training period, observers attempted to make movements of specified durations to hit targets on a computer touch screen (Fig. 7A). Prescribed times were specified on a time bar and, after every trial, the observer's actual duration was also plotted on the time bar so that observers could compare their time to the prescribed time and improve their training performance.

In the main experiment, observers saw a specification of a temporal gain function. Fig. 7B shows the four temporal gain functions used in the experiment. Their task was to plan a movement to hit the target at a time of their choosing. The planned duration of their movement controlled their expected gain and their performance was compared to performance maximizing expected gain.

The actual versus optimal movement times across the four conditions and all observers are summarized in Fig. 7C. Observers were close to optimal. No obvious trends of learning were identified. Hudson et al. compared observed performance across time to reinforcement learning models and excluded the possibility that such models predict observed performance: "To investigate the possibility that observers used a hill-climbing strategy during the main experiment, instead of maximizing expected gain by taking account of their own temporal uncertainty function and experimentally imposed gain function, we performed a hill-climbing simulation using each observer's temporal uncertainty function. In the simulation, intended duration was moved away from the penalty region by $3 \Delta t$ ms after each penalty and towards the center of the target region by Δt ms for each miss of the target that occurred on the opposite side from the penalty (corresponding to the 3:1 ratio of penalty to reward). The value of Δt was initially set to be relatively large. With each change of direction of step, Δt was reduced by 25% to a minimum step size of 1.5 ms. While this simulation approximately reproduced the final average reach times observed experimentally, it does not provide a good model of observer performance. First, there were significant autocorrelations of reach durations beyond lag zero in the simulation data but not in the experimental data. Second, a learning algorithm would be expected to produce substantially higher σ values during test than




Fig. 6. A sequential movement task. The observer first touched the blue circle then the green circle. The tdm targets corresponded to the same (left) or different (right) rewards. The movement time allocating to a target would increase its probability to be hit and thus its reward be won. However, even when the second target was five times valuable than the first target, observers still stuck to allocate slightly more time to the first target. Figure reproduced from Zhang et al. (2010). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In contrast, Mamassian (2008) found that observers in a timing task also failed to maximize expected gain but in this task observers had no prior training with the task and therefore less information about their own timing uncertainty. Maloney and Mamassian (2009) discuss the possible effect of training on observers' abilities to maximize gain.

3.5. Temporal uncertainty in humans and mice

Balci, Freestone, and Gallistel (2009) used a clever design to probe how well mice could cope with their own temporal uncertainty. Fig. 8A gives an illustration of the task. On any trial a food reward would appear at one of two separate feeding hoppers. There were two type of trials. On short-trials, the reward was delivered at the "S" hopper with a short latency (3 s) after the start of the trial. On long-trials, at the "L" hopper with a long latency (6 s).

If the mouse stayed at the right hopper until the time of reward delivery, it would get the reward and otherwise nothing. The difficulty for the mouse was to decide when to move from the "S" hopper to the "L" hopper and its performance was limited by its own timing uncertainty. If, for example, a short-latency reward was delivered but the mouse had switched from the "S" hopper to the "L" hopper at 2 s, it would lose the reward due to a *premature switch*. The opposite error was called a *late switch*. Due to its uncertainty in estimating elapsed time, the mouse could not completely determine the actual time of switch. If the mouse chose to switch after it judged 3 s had elapsed then would likely incur a considerable risk of an early switch. On many trials the mouse would judge that 3 s had elapsed when in fact less than 3 s had elapsed and short-latency reward was still possible. The mouse would forfeit the possibility of reward on all such trials. The choice of switch time that maximized expected gain was determined by the mouse's temporal uncertainty.

A counterpart task was carried out on human observers in which observers were rewarded in points for pressing down one of two keys at the time of reward delivery. One key offered a short-latency reward, the other a long, and the human observer could only press one key at time. The problem for the human observer was to choose the point in time to switch from the

short-latency reward key to the long-latency reward key. The short- and long-latencies for human observers were 2 s and 3 s.

Both human and mouse observers completed several sessions. The probability of short-trials in a session could vary from 0.1 to 0.9. Based on an observer's temporal uncertainty, Balci et al. used a BDT model to solve out the aimed switch time that maximized expected gain for the specific observer. Fig. 8B and C plot optimal switch time against mean switch time for each observer and probability condition respectively for humans and for mice. The average absolute difference between actual and optimal switch times was 172 ms for humans (5% of the 3-s range), and 436 ms for mice (7% of their 6-s range). For mice, Fig. 8D shows how the mean switch time \bar{T} varied almost in the same way with the short trial probability as the optimal switch time \hat{T}_0 did.

For both human and mice observers, this near optimal performance was unlikely to be the result of reinforcement learning. Analyzing trials across the duration of the experiment, Balci et al. identified no discernible improvement (Fig. 8E).

4. Imperfectly optimal observers

Use of BDT as a benchmark model does not imply that human visuo-motor processing is in any sense Bayesian inference, even when human performance is close to ideal (Maloney & Mamassian, 2009). We can view the experimental studies just described as comparisons of human performance to the performance of a BDT observer with the same sensory and motor limitations as the human observer. Geisler (1989) proposed using statistical models as benchmarks in just this way: "... the ideal discriminator measures all the information available to the later stages of the visual system ..." (Geisler, 1989, p. 30). Thus, we compare human performance to a BDT observer precisely because the BDT observer makes the best use of the available information. This benchmark approach grew out of earlier work by Barlow and colleagues (Barlow, 1972, 1995) and it has proven to be a useful tool in the study of human perception (see, for example, Najemnik & Geisler, 2005).

Nevertheless, suppose that we have benchmarked human performance in a visuo-motor task and it is remarkably close to that of its BDT counterpart and we cannot reject the hypothesis that the BDT observer is an accurate model of human visuo-motor processing. This was the outcome of several of the studies we reviewed above. Are we justified in advancing the BDT observer as a model of the perceptual process, at least for this task?

One evident reason that we cannot is technical: a failure to reject the hypothesis of optimal performance may simply be a Type II Error in statistical terms. It is possible that the null hypothesis of optimality is not true but that our experimental design and statistical analyses failed to detect the discrepancy between human performance and ideal. The underlying problem is that the BDT observer is an idealization akin to the notion of a fair coin that has probability of coming down heads of exactly 0.5. No physical coin is ever perfectly fair, and every biological organism can have an off day. In speaking of a fair coin, Feller (1968) justifies the use of such models: "... we preserve the model not merely for its logical simplicity, but essentially for its usefulness and applicability. In many applications it is sufficiently accurate to describe reality." (Feller, 1968, p. 19). However, if the only "reality" we have to describe is that the human observer, in some visuo-motor tasks, does nearly as well as he can be expected to, then there is no reason to conclude that the elements of BDT correspond to anything in human visuo-motor processing. As every child learns in kindergarten these days, there are many ways to be excellent.

A second, and separate problem with BDT as a process model is that the BDT observer needs access to accurate information about likelihood, gain, and prior. In particular, the prior distribution of a

BDT observer is readily interpreted as claims about the environment and the use of the prior is characteristic of Bayesian approaches. Nakayama and Shimojo (1992) argue that the amount of information in the prior for many simple visual tasks is impossibly large. Maloney (2002) estimates the number of world states (the size of the domain of the prior) for one simple shape from shading task and shows that it is too large to be plausibly learned from experience or represented neurally.

Nakayama & Shimojo's argument is apparently compelling but there is an evident way out of this difficulty. We need only drop the requirement that the visuo-motor system have exact estimates of priors. The resulting model observer follows the computations of BDT may not maximize expected gain because of its erroneous estimates of prior, gain or likelihood function (Maloney, 2002). We refer to them as *imperfectly optimal observers*⁷ echoing the title of Janetos and Cole (1981).

In Fig. 9, for example, we consider the case of an imperfectly

optimal observers by determining which of the two has the higher Bayes gain with respect to the true prior. We can potentially measure the true prior for any task describable by BDT and also measure the observer's prior experimentally (e.g. [Adams, Graf, & Ernst, 2004](#); [Mamassian & Landy, 2001](#)) and work out the cost to the observer of any error in estimating priors. Similarly, [Körding](#)

[and Wolpert \(2004\)](#), described above, estimated likelihood functions from human performance and found they were discrepant from actual likelihood functions.

[Janetos and Cole \(1981\)](#) in an article entitled “imperfectly optimal animals”, pointed out a third problem with using idealized models as models of biological performance. They described two tasks where animals' performance well approximated the performance of an optimal algorithm similar in spirit to BDT as presented here. They then pointed out that, for both tasks, there was a very simple behavioral rule that would approximate the performance of the optimal algorithm. The experimenter might mistakenly conclude that an organism implementing the simple rule was an instantiation of the optimal rule.

5. Testing the Bayesian hypothesis

The Bayesian approach is not a specific falsifiable hypothesis but rather a (mathematical) language that allows us to describe the structure of the environment, the flow of visual processing,

two different experiments, and find that the two estimates are discrepant, then there are serious grounds to question the entire Bayesian enterprise. We refer to this criterion as a *comparison test* and it is evidently a test of whether human behavior is controlled by a system of consistent priors on states of the world.

Similarly, Maloney and Mamassian (2009) describe a different test of BDT that they refer to as a *transfer test*. The test assesses whether the visuo-motor system can store and recall and retrieve priors, likelihood functions, and gain functions independently of one another. Maloney and Mamassian argue that the ability to transfer prior information acquired while learning one task to another task carried out later in the same environment would suggest that prior can be stored and reinstated independent of particular tasks.

Consider, for example, the tasks of Trommershäuser et al. (2003a,b) and that of Hudson et al. (2008). In both cases observers had the potential to learn their own spatial and temporal motor uncertainty in training tasks with gain functions different from those employed in the main experiments. This uncertainty was in effect a prior distribution dictating how a movement aimed at a particular point in space and time would be realized.

In the main part of the experiment, they were challenged with a variety of arbitrary gain functions but the movement and its uncertainty were unchanged. That is, the training task and the main experimental task shared the same prior. The lack of any trends in performance in the main experiments of Trommershäuser et al. (2003a,b) and Hudson et al. (2008), indicate that observers could recall and combine prior information learned during training with novel gain functions. Balci et al. (2009) also found no evident trends in performance. Observers in these tasks successfully transferred prior information from one task to a second, passing the transfer test of Maloney and Mamassian (2009).

6C conclusion

Statistical decision theory (SDT) and Bayesian decision theory (BDT) are mathematical frameworks that are particularly congenial to describing the kinds of tasks biological organisms engage in (Milner & Goodale, 1995). Both theories emphasize the potential gain or loss associated with the outcomes of actions and both emphasize the constraints on action introduced by uncertainty. They provide a natural vocabulary for crafting idealized counterparts to actual observers in order to compare human performance to the best performance possible for the human observer.

In this review we first presented the elements of SDT and BDT and then discussed recent work that systematically manipulates these elements as part of an experimental design. The overall conclusion we can draw is that human observers can exploit arbitrary gain functions imposed on the environment and compensate at least in part for changes in environmental priors. They do not always maximize Bayes gain but, in many experiments, they come remarkably close without obvious pattern in their failures. In other experiments (e.g. Zhang et al., 2010) they fail, sometimes dramatically.

By varying gain functions as an independent variable, we potentially observe a wider range of behavior than we would otherwise observe. Moreover, the pattern of failures and successes observed may aid us in developing accurate process models of human visuo-motor processing.

We also discussed interpretations of BDT as process models of human performance (“perception as Bayesian inference,” Knill & Richards, 1996). We asked in effect whether the elements of BDT (priors, etc.) were useful components of process models of human visuo-motor processing and noted that very simple, non-optimal models can well approximate ideal performance (Janetos & Cole,

1981). We focused on a class of model observers that we referred to as *imperfectly optimal observers* that implement BDT but with possibly erroneous estimates of prior, gain, and likelihood functions. We discussed two methods for testing such whether human performance is captured by an imperfectly optimal observer: comparison tests and transfer tests (Maloney & Mamassian, 2009).

SDT and BDT are fundamentally about combining information about uncertainty and gain so as to maximize the expected gain of the observer. This topic is also central to the study of human decision making. It is interesting to compare human performance in visuo-motor tasks which is often found to be near-optimal to that observed in decision making experiments where decision makers typically do not maximize gain (e.g. Kahneman & Tversky, 2000). The observed deviations are large and patterned, with observers typically showing distortions in their use of both probability and gain. One study (Wu, Delgado, & Maloney, 2009) directly compared human decision making with a mathematically equivalent motor task and found that human observers distort probability in both tasks but that the distortions were markedly different in the motor and “classical” decision tasks.

In summary, the kinds of experiments inspired by SDT/BDT are powerful tools for exploring the limits of human visuo-motor capability and the SDT/BDT framework allows comparison of human performance in apparently different tasks such as decision making and movement planning.

acknowledgments

LTM was supported by the Humboldt Foundation.

References

- Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the ‘light-from-above’ prior. *Nature Neuroscience*, 7(10), 1057–1058.
- Balci, F., Freestone, D., & Gallistel, C. R. (2009). Risk assessment in man and mouse. *Proceedings of the National Academy of Sciences*, 106(7), 2459–2463.
- Barlow, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, 1, 371–394.
- Barlow, H. B. (1995). The neuron doctrine in perception. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 415–435). Cambridge, MA: MIT Press.
- Battaglia, P. W., & Schrater, P. R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *Journal of Neuroscience*, 27(26), 6984–6994.
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. New York: Springer.
- Berger, J. O., & Wolpert, R. L. (1988) (2nd ed.). *The likelihood principle: A review, generalizations, and statistical implications* (Vol. 6). Hayward, CA: Institute of Mathematical Statistics.
- Blackwell, D., & Girshick, M. A. (1954). *Theory of games and statistical decisions*. New York: Wiley.
- Dean, M., Wu, S., & Maloney, L. (2007). Trading off speed and accuracy in rapid, goal-directed movements. *Journal of Vision*, 7(5), 1–12.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2000). *Pattern classification*. New York: Wiley.
- Feller, W. (1968) (3rd ed.). *An introduction to probability theory and its applications* (Vol. 1). New York: Wiley.
- Ferguson, T. S. (1967). *Mathematical statistics: A decision theoretic approach*. New York: Academic Press.
- Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, 96(2), 267–314.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2003). *Bayesian data analysis* (2nd ed.). Boca Raton, FL: Chapman & Hall/CRC.
- Green, D. M., & Swets, J. A. (1966/1974). *Signal detection theory and psychophysics*. New York: Wiley (Reprinted 1974, New York: Krieger).
- Hudson, T. E., Maloney, L. T., & Landy, M. S. (2008). Optimal compensation for temporal uncertainty in movement planning. *PLoS Computational Biology*, 4(7), e10000130.
- Janetos, A. C., & Cole, B. J. (1981). Imperfectly optimal animals. *Behavioral Ecology and Sociobiology*, 9(3), 203–209.
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge, UK: Cambridge University Press.
- Kahneman, D., & Tversky, A. (Eds.). (2000). *Choices, values, and frames*. Cambridge, UK: Cambridge University Press.
- Knill, D. C., & Richards, W. (Eds.). (1996). *Perception as Bayesian inference*. New York: Cambridge University Press.
- Körding, K. (2007). Decision theory: What “should” the nervous system do? *Science*, 318, 606–610.

- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, 35(3), 389–412.
- Maloney, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world: Psychological and philosophical issues in perception* (pp. 145–189). New York: Wiley.
- Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. *Visual Neuroscience*, 26, 147–155.
- Mamassian, P. (2008). Overconfidence in an objective anticipatory motor task. *Psychological Science*, 19(6), 601–606.
- Mamassian, P., & Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research*, 41(20), 2653–2668.
- Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modelling of visual perception. In R. Rao, M. Lewicki, & B. Olshausen (Eds.), *Probabilistic models of the brain: Perception and neural function* (pp. 13–36). Cambridge, MA: MIT Press.
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford, UK: Oxford University Press.
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434(7031), 387–391.
- Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science*, 257(5075), 1357–1363.
- Neisser, U. (1976). *Cognition and reality*. San Francisco: W.H. Freeman & Co.
- von Neumann, J., & Morgenstern, O. (1944/1953). *Theory of games and economic behavior* (3rd ed.). Princeton, NJ: Princeton University Press.
- O'Hagan, A. (1994). *Kendall's advanced theory of statistics. Bayesian inference* (vol. 2). New York: Halsted Press (Wiley).
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A*, 20(7), 1419–1433.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*, 16, 255–275.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, 12(8), 291–297.
- Wu, S.-W., Dal Martello, M. F., & Maloney, L. T. (2009a). Sub-optimal allocation of time in sequential movements. *PLoS One*, 4(12), e8228.
- Wu, S.-W., Delgado, M. R., & Maloney, L. T. (2009b). Economic decision-making under risk compared with an equivalent motor task. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 6088–6093.
- Wu, S.-W., Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2006). Limits to human movement planning in tasks with asymmetric gain landscapes. *Journal of Vision*, 6(1), 53–63.
- Yuille, A. L., & Bülthoff, H. H. (1996). Bayesian decision theory and psychophysics. In D. C. Kniill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 123–162). New York: Cambridge University Press.
- Zhang, H., Wu, S., & Maloney, L. (2010). Planning multiple movements within a fixed time limit: The cost of constrained time allocation in a visuo-motor task. *Journal of Vision*, 10(6), 1–17.