

Special Issue: Probabilistic models of cognition

Bayesian decision theory in sensorimotor control

Konrad P. Körding¹ and Daniel M. Wolpert²

¹Brain and Cognitive Sciences, Massachusetts Institute of Technology, Building NE46-4053, Cambridge, Massachusetts, 02139, USA

²Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, UK

Action selection is a fundamental decision process for us, and depends on the state of both our body and the environment. Because signals in our sensory and motor systems are corrupted by variability or noise, the nervous system needs to estimate these states. To select an optimal action these state estimates need to be combined with knowledge of the potential costs or rewards of different action outcomes. We review recent studies that have investigated the mechanisms used by the nervous system to solve such estimation and decision problems, which show that human behaviour is close to that predicted by Bayesian Decision Theory. This theory defines optimal behaviour in a world characterized by uncertainty, and provides a coherent way of describing sensorimotor processes.

Introduction

The central nervous system (CNS) constantly sends motor commands to our muscles. Determining the appropriate motor command is fundamentally a decision process. At each point in time we must select one particular motor command from the set of possible motor commands. Two components jointly define the decision problem: knowledge of the state of the world (including our own body) and knowledge of our objectives.

The sensory inputs of humans are plagued by noise [1,2] which means that we will always have uncertainty about our hand's true location (Figure 1a). This uncertainty depends on the modality of the sensory input: when we use proprioception to locate our hand we may have more uncertainty about its position compared to when we can see it. Moreover, our muscles produce noisy outputs [3,4] and when we quickly move to a target location (shown as a red $\times$ in Figure 1a) our final hand position will typically deviate from the intended target. Even if our sensors were perfect they would only tell us about the part of the world that we can currently sense. This uncertainty places the problem of estimating the state of the world and the control of our motor system within a statistical framework. Bayesian statistics [5–8] provides a systematic way of solving problems in the presence of uncertainty (see the online article by Griffiths and Yuille associated with this issue: [Supplementary material online](#)).

The approach of Bayesian statistics is characterized by assigning probabilities to any degree of belief about the state of the world (see also [Conceptual Foundations](#) editorial by Chater, Tenenbaum and Yuille).

Bayesian statistics defines how new information should be combined with prior beliefs and how information from several modalities should be integrated. Bayesian decision theory [9–11] defines how our beliefs should be combined with our objectives to make optimal decisions. Understanding the way the CNS deals with uncertainty might be key to understanding its normal mode of operation.

The cost of each movement (such as energy consumed) must be weighed against the potential rewards that can be obtained by moving. In the framework of decision theory a utility function should quantify the overall desirability of the outcome of a movement decision. We should choose a movement so that as to maximize utility. Several recent papers have addressed what functions people optimize with their movements. Understanding what human subjects try to optimize is a necessary step towards a rational theory of movement selection.

The selection of a movement can be described as the rational choice of the movement that maximizes utility according to decision theory (see [Box 1](#)). This approach thus asks why people behave the way they do. An increasing number of laboratories have addressed this question within this framework. Here we review recent studies that find human movement performance to be close to the predictions obtained from optimally combining probability estimates with movement costs and rewards. The approach has the potential to embed human behaviour into a coherent mathematical framework.

Estimation using Bayes rule

We need to estimate the variables that are relevant for our choice of movement. For example, when playing tennis we may want to estimate where the ball will bounce. Because vision does not provide perfect information about the ball's velocity there is uncertainty as to the bounce location. However, if we know about the noise in our sensory system then the sensory input can be used to compute the likelihood – the probability of getting the particular sensory inputs for different possible bounce locations (shown in red in Figure 1b). We can combine this with information that is available over repeated experience of tennis: the position where the ball hits the ground is not

Corresponding author: Körding, K.P. (kording@mit.edu).




Figure 1. Bayesian integration. (a) Perception and movement lead to uncertainty. When we briefly look at our hand we cannot be certain where exactly it is. The resulting uncertainty is sketched as the grey probability distribution around the finger at upper left. When we only feel our hand without looking we might have more uncertainty (below). Right: if we make a fast movement from a starting position to a target we will not always hit the target (red X) but there will be some probability distribution of endpoint position. (b) Example: The other player is hitting the ball. Seeing the ball, we can estimate that it will land in the red region (with a likelihood proportional to the saturation). We have prior knowledge that the ball is likely to land in the green region (with a probability proportional to the saturation). The black ellipses denote the posterior, the region where the Bayesian estimate would predict the ball to land. (c) The experimental set-up in typical movement psychophysics experiments. (d) Human subjects' reliance on the prior as a function of increasing perceptual uncertainty. (e) The inferred prior for the different conditions and subjects. The actual distribution used in the experiment is shown in red. (Data for (d) and (e) replotted from [12]).

uniformly distributed over the court. For example the bounce locations are likely to be concentrated within the confines of the court and the distribution might be highly peaked near the boundary lines where it is most difficult to return the ball. This distribution of positions is called the 'prior' (sketched in green) and could be learned through experience. Bayes Rule defines how to combine prior and likelihood to make an optimal estimate of the bounce location (see Box 2).

Bayesian integration in motor control

Bayes rule makes it clear that to perform optimally we must combine prior knowledge of the statistic of the task

Box 1. Decision theory

Decision theory quantifies how people should choose in the context of a given utility function and some partial knowledge of the world. The expected utility is defined as:

$$E[Utility] \equiv \sum_{\text{possible outcomes}} p(\text{outcome}|\text{action})U(\text{outcome})$$

where $p(\text{outcome}|\text{action})$ is the probability of an outcome given an action and $U(\text{outcome})$ is the utility associated with this outcome. According to decision theory people choose the action so as to maximize the expected value of utility. Choosing according to this criterion is the definition of choosing rationally. Within the framework of economics, numerous problems have been described in these terms. For example, people's decision about the ratio of risky to non-risky assets in their portfolio has been described in terms of people having partial knowledge about their future earnings while maximizing their future utility [59]. Companies' decisions about wages and employment of workers have been modelled in terms of the company having partial information about workers' ability and maximizing profits [60]. Moreover, the decisions of the central bank to increase or decrease interest rates has been modelled in terms of optimally reducing the uncertainty about future inflation [61]. Economics tries to understand both how agents should optimally behave when deciding under uncertainty and how they actually behave in such cases. Bayesian decision making is the systematic way of combining Bayesian estimates of probability with utility functions.

Optimal control aims to solve similar problems where the decision is not just happening at one point of time but a continuous output (such as muscle force). The expected utility changes constantly according to new information coming in. Solutions to this problem typically use the notion of 'cost-to-go' the average integrated cost from a current state to a target state.

with the likelihood obtained from the sensory input. In a recent experiment [12], it was tested whether people use such a strategy. Instead of the bounce location of a tennis ball subjects had to estimate the position of a cursor relative to their hand (Figure 1b). Subjects could use two sources of information: The distribution of displacements over the course of many trials (prior), as well as what they see during the current trial (likelihood). The quality of the visual feedback was also varied, in some cases a ball was shown at the position of the cursor giving precise feedback whereas in other trials a large cloud was shown at the position of the cursor thereby increasing the variability (noise) in the sensory input (see Figure 1d).

In this experiment, the Bayesian estimation process defines the optimal estimate as a weighted combination of the mean location of the prior and the peak of the sensory likelihood (see Box 2). Moreover, it predicts that with increasing noise in the sensory feedback subjects should increase the weight of the prior and decrease the weight of their sensory feedback in their final estimate of the location. Figure 1d shows that this Bayesian strategy is observed. From the data it is possible to infer the prior that people are using – assuming that they use an optimal Bayesian strategy. Figure 1e shows that people used a prior that was very close to the optimal one.

This experiment therefore shows that subjects in this task exhibit a strategy very similar to the one predicted by optimal Bayesian statistics. Some experiments could

Box 2. Bayesian statistics

When we have a Gaussian prior distribution $p(x)$ and we have a noisy observation o of the position that leads to a Gaussian likelihood (red curve, Figure 1) $p(o|x)$ it is possible to use Bayes rule to calculate the posterior distribution (yellow curve, Figure 1; how probable is each value given both the observation and the prior knowledge):

$$p(x|o) = p(o|x) \frac{p(x)}{p(o)}$$

This equation assigns a probability to every possible location. If we assume that the prior distribution $p(x)$ is a symmetric one dimensional Gaussian with variance σ_p^2 and mean $\hat{\mu}$ and that the likelihood $p(o|x)$ is also a symmetric one dimensional Gaussian with variance σ_o^2 and mean o , it is possible to compute the posterior that is then also Gaussian in an analytical way. The optimal estimate $\hat{x}$, that is the maximum of the posterior is:

$$\hat{x} = \alpha o + (1 - \alpha) \hat{\mu}$$

where

$$\alpha = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_o^2}$$

Moreover we can calculate the width of the posterior as $\sigma^2 = \alpha \sigma_o$. The parameter α is always less than 1. This Bayesian approach leads to a better estimate of possible outcomes than any estimate that is only based on the sensory input.




Figure 1. Bayesian integration. The green curve represents the prior and red curve represents the likelihood. The yellow curve represents the posterior, the result from combining prior and likelihood.

even relate Bayesian priors for movement production with properties of motor neurons [13]. Using a prior in a Bayesian way, however, is not only restricted to producing movement trajectories and is seen in both force estimation [14] and timing judgements [15]. These results show that the prior knowledge used by human subjects remains plastic and we can optimally adapt to the statistical properties of a task.

Bayesian integration in perception

In addition to results from sensorimotor integration, other research has addressed how human perception can be described by Bayesian estimation processes ([16,17], see Yuille and Kersten, this issue). Within this framework many illusions and other visual effects can be understood [18] by making general assumptions about priors over possible visual objects [19,20] or the direction of illumination [21,22]. Brightness perception [23], shape perception [24], movement perception [25] and certain illusions in length perception [26] have been shown to arise as optimal percepts in Bayesian models in which subjects have uncertainty about the state of the world based on vision alone and therefore incorporate (reasonable) prior beliefs over the possible states of the world. These studies with different approaches have shown that human perception is close to the Bayesian optimal suggesting the Bayesian process may be a fundamental element of sensory processing.

Bayesian cue combination

Bayesian processes can also be used to understand how cues from two different modalities can be combined into a single estimate. For example if we feel the size of an object and at the same time see this object we may want to combine the information from these two modalities. The computational problem that occurs is that if the two cues need to be combined into a joint estimate, equivalent to the way the prior needs to be integrated with the cue in Bayesian integration. To combine cues the system needs to weigh one cue against the other. Calculating this optimally in a Bayesian way means that the weighing will depend on the relative uncertainties in the cues. Recent studies have shown that we are close to the Bayesian optimal when we combine visual and haptic information to estimate the size of an object [27] or visual and auditory information to estimate the position of a stimulus [28]. Similarly, we can combine multiple cues within a modality such as visual texture and motion or stereo cues into a single depth estimate in a way predicted by Bayesian statistics [29–31].

Recent studies have examined how we can combine proprioceptive information about the location of our hand with visual information of the hand itself [32,33] or how we estimate the position of their hand and the configuration of their joints [34]. These results too may be interpreted in a Bayesian framework as optimal estimation in the presence of unknown alignments of the relevant coordinate systems. Bayesian statistics is a general framework that specifies how we could optimally

Box 3. Utility functions

To describe decision making, economists usually use the concept of a utility function [62,63], a hypothesized function that increases with the desirability of the outcome. Although these concepts arose in economics, utility does not have to be directly related to money. For example state of health, and even altruistic feelings of having helped others or having punished defectors [64] is assumed to influence utility. Mathematically utility is defined as the value that we prescribe to each possible outcome of our decisions:

$$Utility = U(outcome)$$

Utility may have a complex relationship to quantity of rewarding stimuli. For example, if we invest well and double our money, we are now twice as well off with respect to the money we own but our utility may not increase by a factor of two as it may tend to saturate as our wealth increases. This effect is described in the framework of prospect theory [65]. Although various deviations have been found that show that people do not seem to perform statistically optimally, the assumption of optimality can still typically explain much of the observed behaviour. Because utility is such an important concept different fields refer to the same idea using different names. Motor control often uses Loss function or Cost function as a name for the negative of the utility function. Neuroscience often refers to functions optimized by neurons as an Objective function. And within a reinforcement learning framework [66,67] utilities are typically called rewards. Different communities refer to concepts equivalent to utility under different names. In some cases a second function is introduced that characterizes the total cost of a movement ('cost to go') which is the integrated instantaneous cost. Regardless of how they are called, utility functions serve to quantify the relative values of different decision outcomes.

A concept that is often used to study utility is the concept of indifference curves. Consider there are two goods, for example apples and bananas. We can ask how desirable different combinations of

apples and bananas are – that is their utility function. An indifference curve is a curve within this space along which people are equally content, that is have the same utility. If people for example only care about the calories of the food, the utility function would be a straight line (see Figure 1). If people prefer a mixture between different goods to the same amount of just one good, a situation that is common in economics the indifference curve takes on a convex form. Asking people questions about their preferences can reveal these curves. However, from choice alone it is impossible to know the full utility function. When all utilities are scaled by a constant, decisions are unaffected.




Figure 1. Indifference curves. People who care only about the number of calories would exhibit straight indifference curves. People who prefer a mixture would exhibit convex indifference curves.

combine different sources of information into a single estimate. Human performance indicates that in many cases of cue integration they operate very close to this theoretical optimum.

Taken together these studies show that over a wide range of phenomena people exhibit approximately Bayes-optimal behaviour. This makes it likely that the algorithm implemented by the CNS may actually support mechanisms for those kinds of Bayesian computations.

Costs and rewards

To put movement into a rational framework it is necessary to define a function that measures how good or bad the outcome of a particular movement is. This function, often termed *cost* may for example be related to the energy consumed during a movement. In general people should prefer less demanding movements – movements that put less strain on the muscles or movements that can be executed using less energy. We are thus faced with the problem of selecting among the infinite set of possible movements the one that minimizes the cost (see Box 3).

Ethological cost functions

Several different cost functions have been proposed for pointing movements. For example, it has been proposed that people move so that their movements are as smooth as possible [35,36]. Such a cost function can explain many findings about target directed movements. More recently

evidence has been presented that the precision of the movement, rather than smoothness, defines the cost function [37]. This approach provides a more intuitive choice of a cost that is based on accuracy, as well as explaining a range of new behaviours. In these approaches a utility function is assumed, which measures how well a movement is performed.

Several recent studies have shown that when the utility is externally defined, where the outcome of a movement is assigned a monetary reward, subjects quickly learn to move in a way that maximises the potential reward [38,39]. However, there are certainly limits to our ability to perform in an optimal fashion and as the complexity of a task is increased this optimality breaks down [40]. Nevertheless, many movement phenomena can be explained by assuming that people move optimally with respect to a simple utility function. The optimality framework provides both a more compact representation than a description of the behaviour and also addresses why we choose to move the way we do.

Measuring cost functions

The cost function used by the CNS might depend on several movement parameters, such as force magnitudes and force durations. Various hypothesized utility functions predict different choices and thus different indifference lines (Figure 2a; Box 3). In an experimental setting it was addressed which function of force is optimized.




Figure 2. Measuring utility functions. (a) Indifference curves in the force–time space as predicted by different cost functions. The cost is the same along each curve. (b) The cost is inferred from the subject's decisions. The 'hotter' the colour, the less desirable the force. (Adapted from [41]).

Human volunteers had to produce forces of varying magnitude and duration [41]. During each trial subjects had to choose which of two combinations of force magnitude and duration they prefer. From a large number of such choices it is possible to infer the indifference lines in the F–T space (Figure 2a). From these indifference curves it is possible to infer the cost function. The measured cost function (Figure 2b) for this task is relatively complicated and differed from the predictions made by any of the known models. Experiments along these lines (see also [42]) can help address utility function that depend on several parameters. The usefulness of this approach, however, will ultimately depend on the degree of generalization. It will depend on how well new movements can be predicted by cost function measured for one class of movement. We expect that many of the utility function used by human subjects will be understood by an adaptive combination of a relatively small number of simple utility functions.

Models of optimal control: using online feedback

Understanding task statistics, the noise on our sensors and actuators and the utility function allows us to predict optimal behaviour. So far we have discussed these processes applied to discrete decisions chosen from a small number of possible decisions. However, in general we produce a continuous trajectory of movement in response to a contiguous stream of sensory input. The system will thus constantly use feedback to update its movements (Figure 3).

Kalman filter

A large number of studies in the area of optimal control use these ideas to model human behaviour. In situations in which we need to estimate the state of the body as it evolves over time we typically use a Kalman filter [43,44]. Kalman filters are a standard technique used in engineering when the unknown state is to be tracked over time. Some psychophysical experiments [43] tested the hypothesis that people use such a mechanism to estimate how their hand moves in the dark. After each movement they had to estimate where their hand was even though

it was not visible. An optimal Kalman filter – a Bayesian technique for continuously varying problems – produced very similar results if it was assumed that people systematically overestimate the forces they produce. Although initially human subjects became less precise, they then went through a period where they became progressively more precise as a function of the trial duration. This effect can be traced back to the finding that a Kalman filter progressively becomes more precise as additional sensory information comes in. It thus seems that people are able to continuously update their estimates based on information coming in from the sensors in a way predicted by Bayesian statistics.

Optimal feedback control

In many cases people do not just have to estimate the position of their limbs but instead have to choose a strategy by which they will efficiently reach their target, a strategy that is optimizing a cost function. Optimal feedback control [45] is a framework for studying such problems. This approach is identical to




Figure 3. Optimal controllers. In generating a movement, the controller, an optimal decision maker, takes into account both the output of the Bayesian estimation process as well as the utility function. The Bayesian estimator combines inputs from the sensors (for example, about limb positions) with prior knowledge in addition to the efference copy – the signal sent by the CNS to the muscles.

Box 4. Questions for future research

The decision theoretic description of human movement leads to several important questions.

- How is prior information encoded in the CNS? How is it combined with new evidence to generate estimates? Some theoretical studies suggest that uncertainty may be represented by neuromodulators [68]. However, the kind of uncertainty that is associated with virtually any variable may well be represented differently.
- Which approximations does the nervous system use? It may have evolved efficient approximate solutions towards solving problems in the areas of Bayesian statistics, decision making and control. The approximations used by the CNS may inform future algorithm developments.
- Most studies on Bayesian integration are done in very simple cases. Are the mechanisms similar in the case of making large complicated movements using many joints. Movements of, for example, the hand involve many joints. Moreover, typically movements involve feedback and the system needs to estimate its state as it changes over time.
- Are utilities and probabilities represented independently of one another? Can one change either of them while leaving the other intact and predict the behavioural changes?
- Are these mechanisms of decision making shared between high-level decision making in the context of cognitive problems and low-level intuitive decision making for movement.

decision theory where a decision is happening at each point in time. Using optimal feedback control as a model of human performance makes several interesting predictions that have been experimentally verified. Importantly the strategy to specify a desired trajectory and then just use feedback to keep you on that trajectory is suboptimal and there is experimental evidence that people use a strategy predicted by optimal feedback control [46–48]. Similar effects can be seen in cases where information is integrated from one trial to the next [49].

In the typical movements that people execute in their everyday life they usually use visual feedback of their movements [50]. For that reason optimal feedback control may be a good candidate for modelling typical behaviour. Optimal feedback control is easily solved in the case of linear dynamics, quadratic costs and Gaussian noise sources (LQG). However, for more realistic situation of nonlinear system such as our arm, with more complex models of noise and cost it can still be prohibitively hard to derive optimal control laws. Some recent research addresses new approaches for computing such control laws [51–53]. Optimal feedback control will allow predictions of progressively more complicated and interesting human movements derived from Bayesian Decision Theory.

Future directions

The approach of formalizing human decision making as being based on partial uncertainty and utility functions formalizes the problems that are solved by the CNS. There is converging evidence from various communities that Bayesian approaches can serve as a coherent description of human decision making.

The optimal statistical approach to sensorimotor control raises many important questions (see Box 4).

However, many of our movements are in the context of complicated tasks such as social interaction. In such cases a good Bayesian model may be arbitrarily complicated as it involves a Bayesian inference about the state of the world, including the mental state of other people. From the movements of other people we can start inferring their internal state. Although this is theoretically possible [54] it will ultimately involve inferring the utility functions of others which is a computationally hard problem [55]. Novel Bayesian approaches have started to be able to describe how people make causal inference [56], a skill that people are particularly good at. In general, Bayesian inference on complicated real world problems are often still proving prohibitively hard from a computational perspective. Quite possibly the brain is using efficient approximations to Bayesian decision making that still allow it to perform well.

Beyond those algorithmic problems it is also important to consider possible constraints and biases in making inferences that are imposed by the brain. The brain is the substrate that is being used to support Bayesian inference and optimal control [57,58]. The way the brain is built, acquired through the course of evolution will already supply us with some knowledge about what kind of a world to expect. It will thus already define what class of models can be implemented and moreover what kind of inference algorithms the brain will have to use. We should not expect that it will be the Bayesian optimum in all cases. Finding deviations from optimal behaviour may lead to interesting insights into the organization of the CNS. Moreover, the Bayesian approach does not specify a representation of the involved data structures. The algorithm implemented by the CNS should not only support Bayesian calculations but also a systematic way of acquiring a useful representation on which to use these calculations.

In conclusion, Bayesian decision theory predicts many of the properties of the movement system and is a coherent framework in which to think about movement decisions. How the brain solves the underlying inference problems and how it represents its information is an important question for further research (see also Editorial ‘Where next?’ in this issue).

Acknowledgements

We like to thank the German Science Foundation Heisenberg Program for support (KK) as well as the Wellcome grant and the HFSP for financial support.

Supplementary data

Supplementary data associated with this article can be found at [doi:10.1016/j.tics.2006.05.003](https://doi.org/10.1016/j.tics.2006.05.003)

References

- 1 Barlow, H.B. *et al.* (1987) Human contrast discrimination and the threshold of cortical neurons. *J. Opt. Soc. Am. A* 4, 2366–2371
- 2 Scott, S.H. and Loeb, G.E. (1994) The computation of position sense from spindles in mono- and multiarticular muscles. *J. Neurosci.* 14, 7529–7540

- 3 Clamann, H.P. (1969) Statistical analysis of motor unit firing patterns in a human skeletal muscle. *Biophys. J.* 9, 1233–1251
- 4 Matthews, P.B. (1996) Relationship of firing intervals of human motor units to the trajectory of post-spike after-hyperpolarization and synaptic noise. *J. Physiol.* 492, 597–628
- 5 Cox, R.T. (1946) Probability, frequency and reasonable expectation. *Am. J. Phys.* 17, 1–13
- 6 Freedman, D.A. (1995) Some issues in the foundation of statistics. *Found. Sci.* 1, 19–83
- 7 MacKay, D.J.C. (2003) *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press
- 8 Jaynes, E.T. (1986) *Bayesian Methods: General Background*, Cambridge Univ. Press
- 9 Yuille, A. and Bulthoff, H.H. (1996) Bayesian decision theory and psychophysics. In *Perception as Bayesian Inference* (Knill, D. and Richards, W., eds), pp. 123–161, Cambridge University Press
- 10 Brainard, D.H. and Freeman, W.T. (1997) Bayesian color constancy. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 14, 1393–1411
- 11 Russell, S. and Wefald, E. (1991) Principles of metareasoning. *Artif. Intell.* 49, 400–411
- 12 Körding, K.P. and Wolpert, D.M. (2004) Bayesian integration in sensorimotor learning. *Nature* 427, 244–247
- 13 Singh, K. and Scott, S.H. (2003) A motor learning strategy reflects neural circuitry for limb control. *Nat. Neurosci.* 6, 399–403
- 14 Körding, K.P. et al. (2004) Bayesian Integration in force estimation. *J. Neurophysiol.* 92, 3161–3165
- 15 Miyazaki, M. et al. (2005) Testing Bayesian models of human coincidence timing. *J. Neurophysiol.* 94, 395–399
- 16 Knill, D. and Richards, W., eds (1996) *Perception as Bayesian Inference*, Cambridge University Press
- 17 Yuille, A. and Kersten, D. (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* DOI:10.1016/j.tics.2006.05.002
- 18 Geisler, W.S. and Kersten, D. (2002) Illusions, perception and Bayes. *Nat. Neurosci.* 5, 508–510
- 19 Kersten, D. and Yuille, A. (2003) Bayesian models of object perception. *Curr. Opin. Neurobiol.* 13, 150–158
- 20 Kersten, D. et al. (2004) Object perception as Bayesian inference. *Annu. Rev. Psychol.* 55, 271–304
- 21 Adams, W.J. et al. (2004) Experience can change the ‘light-from-above’ prior. *Nat. Neurosci.* 7, 1057–1058
- 22 Brewster, D. (1826) On the optical illusion of the conversion of cameos into intaglios and of intaglios into cameos, with an account of other analogous phenomena. *Edinburgh J. Sci.* 4, 99–108
- 23 Adelson, E.H. (1993) Perceptual organization and the judgment of brightness. *Science* 262, 2042–2044
- 24 Langer, M.S. and Bulthoff, H.H. (2001) A prior for global convexity in local shape-from-shading. *Perception* 30, 403–410
- 25 Stocker, A. and Simoncelli, E. Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci.* (in press)
- 26 Howe, C.Q. and Purves, D. (2002) Range image statistics can explain the anomalous perception of length. *Proc. Natl. Acad. Sci. U. S. A.* 99, 13184–13188
- 27 Ernst, M.O. and Banks, M.S. (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433
- 28 Alais, D. and Burr, D. (2004) The ventriloquist effect results from near-optimal bimodal integration. *Curr. Biol.* 14, 257–262
- 29 Jacobs, R.A. (1999) Optimal integration of texture and motion cues to depth. *Vision Res.* 39, 3621–3629
- 30 Knill, D.C. and Saunders, J.A. (2003) Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Res.* 43, 2539–2558
- 31 Hillis, J.M. et al. (2004) Slant from texture and disparity cues: optimal cue combination. *J. Vis.* 4, 967–992
- 32 van Beers, R.J. et al. (1999) Integration of proprioceptive and visual position-information: An experimentally supported model. *J. Neurophysiol.* 81, 1355–1364
- 33 van Beers, R.J. et al. (1996) How humans combine simultaneous proprioceptive and visual position information. *Exp. Brain Res.* 111, 253–261
- 34 Sober, S.J. and Sabes, P.N. (2005) Flexible strategies for sensory integration during motor planning. *Nat. Neurosci.* 8, 490–497
- 35 Hogan, N. (1984) An organizing principle for a class of voluntary movements. *J. Neurosci.* 4, 2745–2754
- 36 Flash, T. and Hogan, N. (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* 5, 1688–1703
- 37 Harris, C.M. and Wolpert, D.M. (1998) Signal-dependent noise determines motor planning. *Nature* 394, 780–784
- 38 Trommershäuser, J. et al. (2003) Statistical decision theory and the selection of rapid, goal-directed movements. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 20, 1419–1433
- 39 Maloney, L.T. et al. Questions without words: A comparison between decision making under risk and movement planning under risk. In *Integrated Models of Cognitive Systems* (Gray, W., ed.), Oxford University Press (in press)
- 40 Wu, S.W. et al. (2006) Limits to human movement planning in tasks with asymmetric gain landscapes. *J. Vis.* 6, 53–63
- 41 Körding, K.P. et al. (2004) A neuroeconomics approach to measuring human loss functions. *PLoS Biol.* 2, e330
- 42 Körding, K.P. and Wolpert, D. (2004) The loss function of sensorimotor learning. *Proc. Natl. Acad. Sci. U. S. A.* 101, 9839–9842
- 43 Wolpert, D.M. et al. (1995) An internal model for sensorimotor integration. *Science* 269, 1880–1882
- 44 Kalman, R.E. (1960) A new approach to linear filtering and prediction problems. *J. of Basic Engineering* 82D, 35–45
- 45 Todorov, E. (2004) Optimality principles in sensorimotor control. *Nat. Neurosci.* 7, 907–915
- 46 Todorov, E. and Jordan, M.I. (2002) Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5, 1226–1235
- 47 Saunders, J.A. and Knill, D.C. (2005) Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements. *Exp. Brain Res.* 162, 458–473
- 48 Saunders, J.A. and Knill, D.C. (2004) Visual feedback control of hand movements. *J. Neurosci.* 24, 3223–3234
- 49 Baddeley, R.J. et al. (2003) System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *J. Neurosci.* 23, 3066–3075
- 50 Land, M. et al. (1999) The roles of vision and eye movements in the control of activities of daily living. *Perception* 28, 1311–1328
- 51 Todorov, E. (2005) Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput.* 17, 1084–1108
- 52 Brock, O. and Kavraki, L. (2001) Decomposition-based motion planning: a framework for real-time motion planning in high-dimensional configuration spaces. In *IEEE International Conference on Robotics and Automation* (Vol. 2), pp. 1469–1474, ICRA
- 53 Todorov, E. et al. (2005) From task parameters to motor synergies: a hierarchical framework for approximately-optimal feedback control of redundant manipulators. *Journal of Robotic Systems* 22, 669–710
- 54 Wolpert, D.M. et al. (2003) A unifying computational framework for motor control and social interaction. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 358, 593–602
- 55 Abbeel, P. and Ng, A.Y. (2004) Apprenticeship learning via inverse reinforcement learning. In *Twenty-first International Conference on Machine Learning* (Vol. 69), pp. 1–8, ACM International Conference Proceeding Series
- 56 Tenenbaum, J.B. et al. Intuitive theories as grammars for causal inference. In *Causal Learning: Psychology, Philosophy, and Computation* (Gopnik, A. and Schulz, L., eds), Oxford University Press (in press)
- 57 Scott, S.H. (2004) Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* 5, 532–546
- 58 He, J. (1991) Feedback gains for correcting small perturbations to standing posture. *IEEE Transactions on Automatic Control* 36, 322–332
- 59 Guiso, L. et al. (1996) Income risk, borrowing constraints, and portfolio choice. *Am. Econ. Rev.* 86, 158–172
- 60 Gibbons, R. et al. (2005) Comparative advantage, learning, and sectoral wage determination. *J. Labor Econ.* 23, 681–724

- 61 Bernanke, B.S. and Woodford, M. (1997) Dynamic effects of monetary policy. *Journal of Money, Credit and Banking* 29, 653–684
- 62 Bentham, J. (1780) *An Introduction to the Principles of Morals and Legislation*, Clarendon Press
- 63 Bernoulli, D. (1738) Specimen theoriae novae de mensura sortis. *Comentarii academiae scientiarum imperialis Petropolitanae (for 1730 and 1731)* 5, 175–192
- 64 Fehr, E. and Rockenbach, B. (2004) Human altruism: economic, neural, and evolutionary perspectives. *Curr. Opin. Neurobiol.* 14, 784–790
- 65 Kahneman, D. and Tversky, A. (1979) Prospect theory: an analysis of decision under risk. *Econometrica* XLVII, 263–291
- 66 Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, MIT Press
- 67 Smart, W.D. and Kaelbling, L.P. (2002) Effective reinforcement learning for mobile robots. In *International Conference on Robotics and Automation* (Vol. 1), pp. 3404–3410, IEEE
- 68 Yu, A.J. and Dayan, P. (2005) Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692

Five things you might not know about Elsevier

1.

Elsevier is a founder member of the WHO's HINARI and AGORA initiatives, which enable the world's poorest countries to gain free access to scientific literature. More than 1000 journals, including the *Trends* and *Current Opinion* collections, will be available for free or at significantly reduced prices.

2.

The online archive of Elsevier's premier Cell Press journal collection will become freely available from January 2005. Free access to the recent archive, including *Cell*, *Neuron*, *Immunity* and *Current Biology*, will be available on both ScienceDirect and the Cell Press journal sites 12 months after articles are first published.

3.

Have you contributed to an Elsevier journal, book or series? Did you know that all our authors are entitled to a 30% discount on books and stand-alone CDs when ordered directly from us? For more information, call our sales offices:

+1 800 782 4927 (US) or +1 800 460 3110 (Canada, South & Central America)
or +44 1865 474 010 (rest of the world)

4.

Elsevier has a long tradition of liberal copyright policies and for many years has permitted both the posting of preprints on public servers and the posting of final papers on internal servers. Now, Elsevier has extended its author posting policy to allow authors to freely post the final text version of their papers on both their personal websites and institutional repositories or websites.

5.

The Elsevier Foundation is a knowledge-centered foundation making grants and contributions throughout the world. A reflection of our culturally rich global organization, the Foundation has funded, for example, the setting up of a video library to educate for children in Philadelphia, provided storybooks to children in Cape Town, sponsored the creation of the Stanley L. Robbins Visiting Professorship at Brigham and Women's Hospital and given funding to the 3rd International Conference on Children's Health and the Environment.