Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

7 Bayesian approaches to modelling action selection

Max Berniker, Kunlin Wei, and Konrad Körding

Summary

We live in an uncertain world, and each decision may have many possible outcomes; choosing the best decision is thus complicated. This chapter describes recent research in Bayesian decision theory, which formalises the problem of decision making in the presence of uncertainty and often provides compact models that predict observed behaviour. With its elegant formalisation of the problems faced by the nervous system, it promises to become a major inspiration for studies in neuroscience.

7.1 Introduction

Choosing the right action relies on our having the right information. The more information we have, the more capable we become at making intelligent decisions. Ideally, we want to know what the current state of the world is, what possible actions can we take in response to it, and what the outcomes of these actions will be. When we choose actions that will most clearly bring about our desired results, we are said to be behaving rationally (see Chapter 2). Equivalently, we could say that rational behaviour is optimal, in that this behaviour executes the best actions for achieving our desired results (see Chapters 3 and 4). Thus behaving rationally is equivalent to solving an optimality problem: what actions should we select to best achieve our goals?

The mathematics that formalises these types of optimality problems is well developed under the study of *decision theory* and *optimal control*. Using these approaches, all possible actions can be assessed in terms of their ability to achieve our goals. Thus the optimal action can be identified and executed. To quantify the relative merit of one action over another, we must define a cost function (sometimes referred to as a loss function). In its most general form, the cost function assigns a value to a choice of action and the resulting outcome, within the context of our goals. In essence, it mathematically quantifies our goals and our strategy for achieving them. Therefore, by selecting actions that minimise this cost, we can best achieve our goals. To illustrate, consider a simplified game of darts. Assume the actions we can choose amongst uniquely specify where the

Modelling Natural Action Selection, eds. Anil K. Seth, Tony J. Prescott and Joanna J. Bryson.

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

dart lands on the board, the outcome of our action. Assuming our goal is to get the highest score possible, our ‘cost’ function is inversely proportional to our score; that is, we achieve the lowest cost by getting the highest score. The optimal choice of action is the one (or any of the many actions, see Chapter 5) that result in the dart landing on the triple 20 (60 points, better than a bull’s eye). Mathematically, we would express this optimal action as,

$$a^o = \arg\max_a \{\text{cost}(outcome(a), a)\} \tag{7.1}$$

Here a is shorthand for the possible actions, and a^o is the optimal action. The above example, while valuable in portraying the general approach to selecting optimal actions, requires more elaboration before we can analyse human motor behaviours. Certainly a more realistic description of the example would include temporal dynamics of our limb and the dart, and perhaps nonlinearities associated with them (we’ll address both below), but first and foremost, we shall argue that we must take into account uncertainty. That is, we’ve assumed we deterministically choose a command and can be sure of its outcome. Yet, most action selection problems are stochastic in nature, and action outcomes are affected by uncertainty. Under more realistic conditions we cannot simply choose the action that lands the dart on the triple 20, since we cannot be certain of this outcome. Instead, we must account for the statistics of the task. To continue the example, if we aim at the centre of the dartboard we are not certain to hit the bull’s eye with the dart (even if we are a seasoned champion). Indeed, if we aim at the centre and throw many times we will end up with a distribution of dart positions. This distribution characterises the likelihood of each outcome, the dart landing at a certain position on the board, given that we aimed at the centre. This distribution helps to define a new optimisation problem: what action minimises the likely cost?

Taking the uncertainty of our aims into account, our new optimal action minimises the so-called expected cost: the average cost when weighted by the probability of the various outcomes.

In the darts example, the best aiming point is a point where we will receive high scores even if we make large mistakes. In fact, both amateur and world-class players are known to adopt a strategy that is well predicted by this approach. Mathematically, we would express this best action as,

$$a^o = \arg\max_a \left\{\sum\nolimits_{outcomes} \text{cost}(outcome(a), a) p(outcome|a)\right\} \tag{7.2}$$

Here $p(outcome|a)$ is the probability of an outcome conditioned on the choice of action, a.

While playing darts, our uncertainty in the dart’s location arises largely from motor noise (Figure 7.1a). However, there are many sensory sources of uncertainty as well. For example, our visual system is noisy and our sense of the location of the dartboard relative to our body is uncertain (Figure 7.1a). Moreover, our proprioceptive system is noisy as well; the orientation of our hand and arm as we release the dart are uncertain (Figure 7.1a). These and many other sources of uncertainty combine to produce variability in the motor outcome given our actions.

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

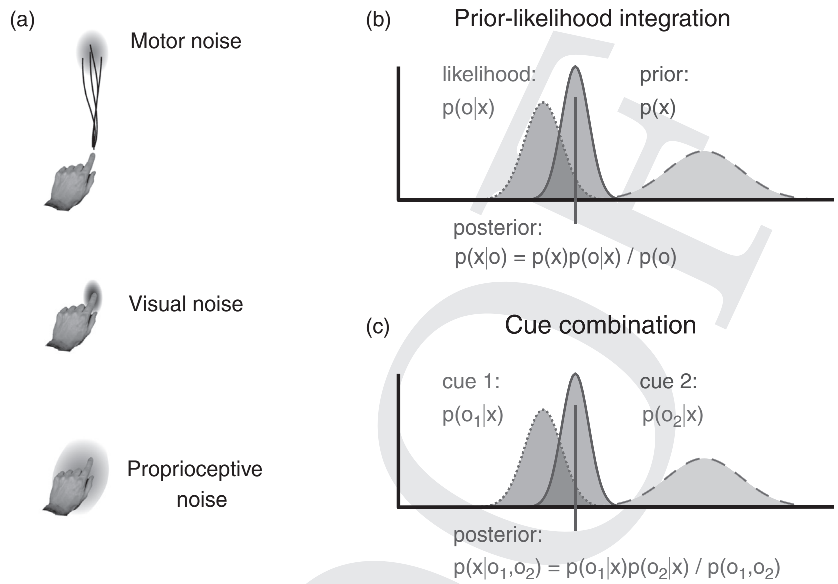

Figure 7.1 (a) Possible sources of noise in the motor system. Motor commands possess noise and result in uncertain movements. Visual and proprioceptive senses too, contain inherent uncertainties. (b) Bayesian integration of a prior and likelihood. The prior, denoted with the dashed curve, represents the probability of a state, x. The likelihood, denoted with the dotted curve, represents the probability of observing the data, o, given x. The posterior is the probability that x is the state, given our observation, o. (c) Bayes' rule applied to cue combinations is mathematically equivalent, only instead of using a prior and likelihood, we integrate two likelihoods.

Pressing our example further, we can examine more sophisticated and realistic action selection problems. For instance, a more sensible description of dart throwing recognises that the task is dynamic. The dart's final position on the board depends on the ballistic trajectory it takes once it has been released from our hand. The motion of the dart up until the moment we release it is dictated by the inertial mechanics of our limb and the force generating properties of our muscles. Clearly, our description of the dart-throwing problem can take on greater and greater levels of detail and physical accuracy. Nevertheless, our problem is still to select the best possible action, albeit an action that varies in time due to the dynamic nature of the problem. The minimisation to be performed at each instant is a sum over both the statistics of the possible outcomes at the current time, and the statistics of future possible outcomes that unfold as a result of the current choice of action.

For a dynamic task such as this, we need to know the state of the system to compute the optimal action. For our dart-throwing example, depending on how we modelled the problem, the state might be the orientation and velocity of our hand. To measure this state we would use sensory feedback. Yet this information is subject to noise and its own statistics. Each of our sensory modalities offers only a limited level of precision, which can vary depending on the situation. Vision is limited under dim light or

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

extra-foveal conditions, hearing becomes unreliable for weak sounds, and proprioception drifts without calibration from vision (Deneve, 2008; Hoyer and Hyvärinen, 2003; Pouget *et al*., 2003; Zemel *et al*., 1998). Therefore, in addition to the uncertainty concerning the outcome of our motor actions, we must cope with the further uncertainty of the state.

Despite the varying complexity of the above examples, we can still formulate the action selection problem as one of an optimisation. We must solve a statistical problem concerning likely motor outcomes conditioned on the actions we choose. The difficulty with this problem, and action selection in general, is that computing the probability of an outcome given an action can be onerous. Doing so demands the specification of many variables that are uncertain and subject to their own statistics. This includes information about our body, the world and how they interact. Integrating all this uncertain information requires a statistical approach. Bayesian integration is the mathematical framework that calculates how uncertain information from multiple sources can be optimally combined. It results in a coherent and maximally accurate estimate of a set of observations, crucial for ultimately selecting the best action.

As explained above, before we can rationally, or optimally, select an action, we must be well informed, i.e., we must know the state of the world, the actions we can take, and their outcomes. In this chapter we focus on the process of constructing a state estimate. A number of recent psychophysics and computational studies have analysed how people and animals integrate uncertain information to make sensorimotor decisions. Below we will discuss some of these findings in the context of action selection. In particular, we discuss how Bayesian integration allows us to combine multiple pieces of information into a single distribution, how we can update this distribution over time as we continue to gain new information about it, and finally, how we can use our observations to update our beliefs about the structure of the world, that is, what processes are responsible for shaping our observations.

7.2 Bayesian estimation

7.2.1 Combining prior knowledge with new evidence

Combining uncertain information to produce a coherent and accurate estimate of our body and the world is integral to action selection. Sometimes the source of this uncertain information is our senses, and we must compare it against what we would expect before it can be of benefit. As an example, consider the task of descending a staircase. Based on our familiarity with walking down stairs, we have strong assumptions for things like the distance between steps, their height and their general shape. These assumptions form a prior over stairs, a belief in their typical properties. Often these priors are strong enough that we feel comfortable taking stairs without even observing them, as when we descend stairs without looking at our feet, or in the dark. Normally though, we'll first observe how far we'll need to step. Vision does not provide perfect measurements however. The visual system provides us with an estimate, or likelihood of the step's height. This likelihood is the probability of having a particular sensory observation for a given stair height. Bayes'

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

rule defines how to combine the prior and the likelihood to make an optimal estimate of the step's height.

Bayes' rule states that the probability of the step's height being value x, given our observation, o, is the product of the prior probability of the stair height and the likelihood, normalised by the probability of the observation. Mathematically, this is expressed as,

$$p(x|o) = p(x)p(o|x)/p(o) \tag{7.3}$$

The distribution produced by Equation (7.3) is known as the posterior probability (this is shown graphically Figure 7.1b). We can also interpret Bayes' rule as the 'optimal' means of combining a prior and a likelihood, as it produces an estimate with minimal uncertainty.

Several studies, using many sensory modalities, have shown that when subjects combine preceding knowledge with new information their behaviour reflects the integration of a prior and likelihood in a manner prescribed by Bayes' rule (Körding, 2007). In a typical study (Körding and Wolpert, 2004), subjects will indicate their estimate of a target's location through a motor task. On each trial, the target's location is drawn from a normal distribution, the prior. Noisy feedback of the target's location is then provided, the likelihood. The distribution used for the prior or the likelihood can be fixed, or vary across subjects as an experimental condition. Bayesian statistics predicts how subjects should combine the likelihood and the prior. These predictions are then compared against human performance, often revealing a high degree of similarity.

These paradigms have been applied to a wide range of topics spanning sensorimotor integration, force estimation, timing estimations, speed estimations, stance regulation, the interpretation of visual scenes, and even cognitive estimates (Chater *et al*., 2006; Knill and Richards, 1996; Körding and Wolpert, 2006; Körding, Ku, *et al*., 2004; Miyazaki *et al*., 2005, 2006; Peterka and Loughlin, 2004; Tassinari *et al*., 2006; Weiss *et al*., 2002). Together these studies demonstrate that people are adept at combining prior knowledge with new evidence in a manner predicted by Bayesian statistics; a necessary first step in action selection.

7.2.2 Combining multiple pieces of information

In many cases it is not new information that is combined with prior knowledge, but rather two or more different pieces of information available to the nervous system, that are combined. For example, we may see and feel an object at the same time (Ernst and Banks, 2002). We can then use our sense of vision and touch to estimate a common property of the object, for instance, its size or its texture. This type of task is what is commonly referred to as *cue combination*; two or more sensory cues are combined to form a common, or joint, estimate. Just as before, Bayesian statistics prescribes how we should combine the likelihoods to compute an optimal estimate from the posterior distribution (Figure 7.1c). Again, only by accurately combining these sensory cues can we optimally select the appropriate action.

Recent studies have found that when combining information from multiple senses, people are also similar to optimal. As an example study, we consider how people combine

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

visual and auditory information to estimate the position of a target. First the precision of visual and auditory perceptions are separately measured for each subject (Alais and Burr, 2004). This is done to characterise the subject's likelihood for the two sensory modalities. Then the precision and accuracy of perception is measured when subjects use both senses. Their performance in these cue combination trials can be predicted using the rules of Bayesian integration, further evidencing people's ability to cope with uncertain information in a statistically optimal manner.

In addition to auditory and visual cues, combinations of other sensory modalities have been analysed in a good number of studies (Ernst and Banks, 2002; Ernst and Bulthoff, 2004; Geisler and Kersten, 2002; Kersten and Yuille, 2003; Knill, 2007; Knill and Richards, 1996; Körding, 2007; Rosas *et al*., 2005; van Ee *et al*., 2003). In all reported cases, cues are combined by the subjects in a fashion that is close to the optimum prescribed by Bayesian statistics.

7.2.3 Credit assignment

The rules of Bayesian statistics reviewed above prescribe how we bring new information to bear upon our beliefs. As just illustrated, at times we are only concerned with estimating a single property. However, sometimes we are concerned with estimating many properties at once. For instance, suppose we hold two blocks in our hand, one sitting atop the other, and we want to estimate their individual weights. Our observation of their combined weight upon our hand is indicative of their individual weights, namely the magnitude of their sum. However, this information is not sufficient to accurately establish their individual weights. To estimate them, we need to solve a credit assignment problem; that is, how does each property (individual weight) contribute to our observation (overall weight). Bayesian statistics also prescribes an optimal solution to this problem.

Our observation of their combined weight forms a likelihood, a probability of having a particular sensory observation of their combined weight, for each value of their individual weights. By combining this with a prior over the object's weight, perhaps based on their size (Flanagan *et al*., 2008), we can compute a posterior distribution of their individual weights. Again, this distribution will provide an optimal estimate of their individual weights, simultaneously solving the credit assignment problem: how much does each block contribute to the total weight I feel? Recent research (Berniker and Körding, 2008) addresses this type of problem by integrating information over time.

7.3 Bayesian estimation across time

In our discussions above we were concerned with estimating what we assumed were static properties, such as the weight of a block or the location of an auditory cue. However, the world is dynamic, and as such its properties and our perception of them are continually changing. Thus we constantly need to integrate newly obtained information with our

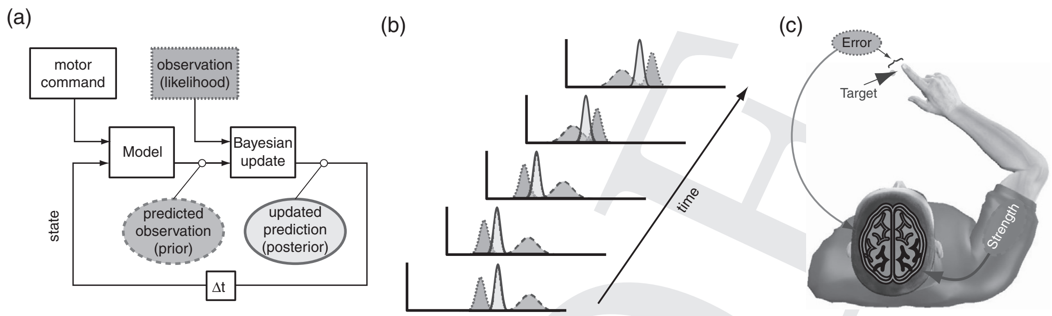

Figure 7.2 (a) Typical procedure for optimally estimating the world's state in modern control theory. A model of the world, combined with a motor command is used to estimate a predicted state (the prior). Observations of the world dictate the likelihood of a particular observation given the current world state. A Kalman filter is used to make a Bayesian update of our belief in the world's current state. (b) This process of Bayesian inference repeats itself at each time step, using the posterior from one time step, as the prior for the following time step. (c) Motor adaptation can be framed as an analogous update procedure. Our prior belief in muscle properties (e.g., muscle strength, labelled in green) is integrated with our observed motor errors to update estimates in our muscle properties.

current beliefs to inform new estimates of the world if we are to behave rationally. This implies that Bayesian integration should take place in an ongoing, or continuous manner.

This ongoing integration of information is an approach taken extensively in modern applications of control theory through the use of a Kalman filter. Kalman filtering is a procedure for using a model and our observations to continuously update our beliefs (Figure 7.2a). At each update the Kalman filter combines a model's estimate of the world's state (the prior) with a measured observation (the likelihood) to update a prediction of the world's current state, represented by a posterior (Figure 7.2b). At any point of time, the posterior from the past defines the prior for the future. This formalism is well developed and used in a wide range of applications from aeronautics to humanoid robotics (Stengel, 1994; Todorov, 2006). Indeed, even the motor control problems of two applications as disparate as controlling a jet and controlling our bodies, share many computational analogies. In both cases, continuously incoming information needs to be assessed to move precisely (albeit on different timescales). Only recently has the methodology of Kalman filtering been applied to make quantitative predictions of human movement behaviour.

As an example of this approach we consider a recent study of motor adaptation (Berniker and Körding, 2008). Since the properties of our bodies change continuously throughout our lives, it is imperative that we monitor these changes if we are to control our motor behaviours accurately and precisely. For example, errors in the perceived strength of our muscles will translate to movement errors. However, we can uses these errors to obtain a likelihood characterising how strong are muscles are. According to Bayesian statistics we should combine this newly obtained information, our motor errors, with our prior beliefs. Accordingly we infer new and improved estimates of our muscles (Figure 7.2c).

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

It has been found that motor adaptation across time can be understood using the predictions of Kalman filtering. For instance, evidence suggests how people estimate the position of their hand (Wolpert *et al*., 1995), adapt to robot-rendered force fields (Berniker and Körding, 2008), and even balance a pole upon their hand (Mehta and Schaal, 2002) can all be explained using this strategy. Taken together, these studies highlight the hypothesis that when people integrate information over time, they seem to do so in a fashion that is consistent with the optimal Bayesian solution. These findings have important consequences for optimal action selection, where finding the best action relies on accurate estimates of the world's state.

7.4 Bayesian estimation of structure

In our discussion above, cue-combination studies have provided evidence that the nervous system combines multiple pieces of information for estimating states in an optimal fashion. However, the brain receives a very large number of sensory cues, and not all cues should be combined; many cues may simply be irrelevant. To make sense of the world then, the nervous system should only combine those sensory cues that are likely to provide information about the same characters of the world.

Consider the example of a ventriloquist: he/she synchronises his/her voice with the movement of the puppet's mouth. Due to visuo-auditory cue combination, the audience experiences the illusion of a talking puppet, or more precisely, that a voice is being emitted from the puppet's mouth. If the ventriloquist's voice is out of sync with the puppet's movements, this illusion will immediately break down. In this example, the temporal proximity of cues is an overriding factor in inducing the merging of cues. Similarly, spatial proximity or disparity can also influence whether we combine different cues into a single percept: if a thundering sound is emitted close to a flash of light, these two cues are readily interpreted as an explosion; in contrast, if the sound is far away from the flash, these two stimuli may be perceived as independent events. The nervous system combines cues that appear to originate from the same source. In other words, it seems that the nervous system estimates the *structure* of the world, i.e., what causes the perceived cues.

Traditional Bayesian analyses of cue combination examine tasks where the experimental cues are close to coincident, and implicitly assume that the cues are caused by the same event, or source. New studies have tested subject performance in situations where two cues are dissimilar from one another. In the simplest case, a visual cue and an auditory cue are simultaneously presented for estimating the location of an object. These two cues can either have the same cause or have different causes. In the same-cause case, these two cues should be combined to form a single estimate; in the different-cause case, these two cues should be processed separately (Figure 7.3a). The optimal estimate should weigh each cause with respect to its likelihood. This likelihood is a function expressing the probability of observing the cues when they arise from the same cause, or not. Not surprisingly, this likelihood depends on spatial and temporal disparity between cues, with increasing disparity between the cues subjects' belief in a

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

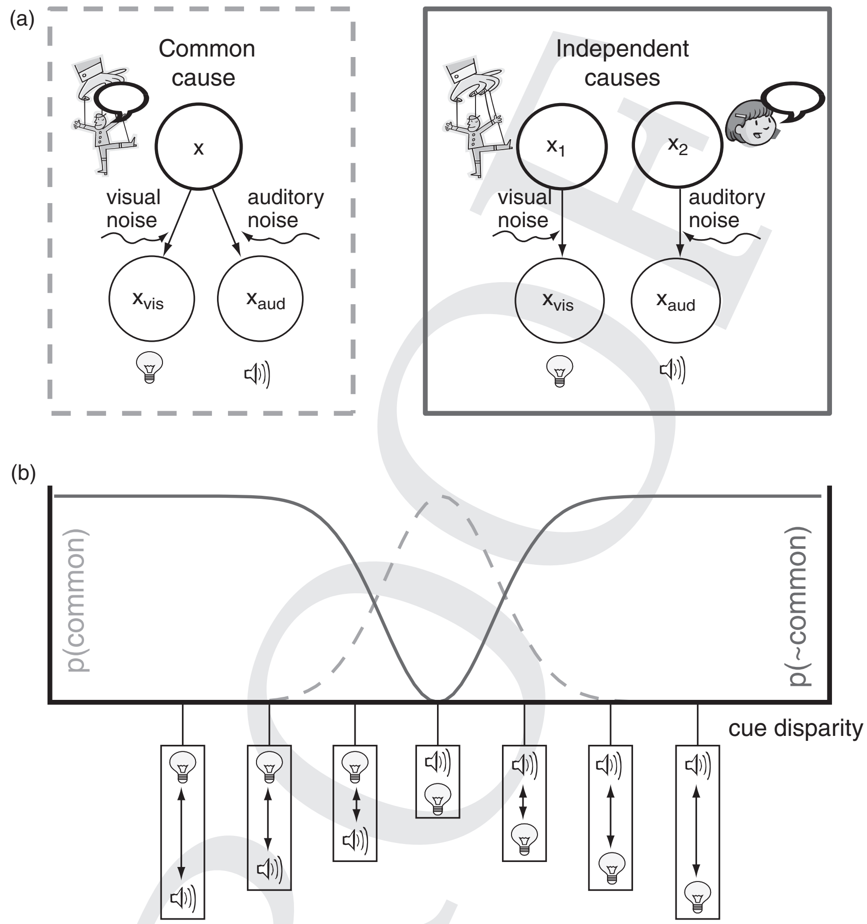

Figure 7.3 Two structural beliefs of the world. (a) If two cues are adequately coincident, subjects perceive them as having a common cause (the dashed box), a phenomenon typified through ventriloquism. In this case, the two cues should be integrated according to Bayes' rule. If the cues are disparate in time or space, subjects perceive them as having independent causes (the solid box). (b) Subject's belief of a common cause. As the spatial disparity of two cues, a voice and a puppet's motions, or a light flash and a tone, is experimentally controlled the belief in a common cause can be manipulated.

common cause decreases (Figure 7.3b). Recent studies have provided support for the predictions of Bayesian models of causal inference for cross-modality cue combination (Körding *et al*., 2006; Shams *et al*., 2005), depth perception (Knill, 2007), and estimation of stimuli numbers (Wozny *et al*., 2008). These studies highlight that the nervous system estimates state with consideration of causal structures among sensory cues in a way that is close to statistical optimum.

One central feature of the human movement system is that it adapts continuously to changes in the environment and in the body. It has been suggested that this adaptation is largely driven by movement errors. However, errors are subject to unknown causal

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

structures. For example, if an error is caused by changes in the motor apparatus, such as strength changes due to muscle fatigue, the nervous system should adapt its estimates of the body. On the other hand, if an error is caused by random external perturbations, the nervous system should not adapt its estimates of the body, as this error is irrelevant. Given the inherent ambiguity in errors, the best choice of adaptive action should be a function of the probability of error being relevant. Bayesian models can predict the optimal adaptation to errors with unknown causal relationship. These predictions have been confirmed in a recent study on motor adaptation in a visuomotor task (Wei and Körding, 2008). It appears that in these sensorimotor tasks the nervous system not only estimates the state but also estimates the causal structure of perceptual information.

The statistical problem that subjects solve in cue combination implicitly involves an inference about the causal structure of the stimuli, e.g., did the flash of light cause the auditory tone, did the tone cause the flash, was there an unknown cause for them both, or were they simply coincidental? The problem faced by the nervous system is thus similar to those studied extensively in psychology and cognitive science that occur in the context of causal induction (Gopnik *et al*., 2004; Griffiths and Tenenbaum, 2005; Michotte, 1963; Sperber and Premack, 1995; Tenenbaum *et al*., 2006). Many experiments demonstrate that people interpret events in terms of cause and effect. The results presented here show that sensorimotor integration exhibits some of the same factors found in human cognition. Incorporating the causal relationships among state variables, it appears the nervous system estimates a structural view of the world. Once again, this process is necessary to optimally selection actions, whether they be for motor, or any other behaviours.

7.5 Bayesian decision theory and inverse decision theory

We have surveyed a great amount of evidence on how people judge, gauge, and generally make sense of their perceptions of the world. We have shown that they estimate the structure of the world, using models to predict how its property values evolve over time, and combine multiple sensory cues or their prior experience, to form a best estimate. Conceptually, however, this is only the first step in action selection. Once we have estimated the state of the world, we need to assess the relative value of all available actions in terms of the statistics of their outcomes. Only then can we choose the action that will best achieve or goals. Mathematically this is expressed as minimising our expected costs (Equation (7.2)). This formalism of choosing optimal actions by combining the statistics of outcomes and the resulting costs falls under the study of *decision theory* and *optimal control*. In all but the simplest cases, computing these optimal commands is very difficult.

Consider one such simple case, that of pea shooting. Those that pursue this hobby blow dried peas through a small hollow tube, or pea shooter. Once the pea leaves the shooter, its fate is dictated by wind, aerodynamics, etc. Assuming you have total control over the position and orientation of your pea tube (and in the World Pea Shooting Championship, it is not unheard of for contestants to use laser targeting for this express

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

purpose) where you aim will dictate a distribution of possible locations the pea will land within the target. Referring again to equation (7.2), choosing the best aiming location requires considering the likely score, a distribution of pea locations will yield, for any given aim point. However, if you are unsure where exactly you are aiming (or your laser is faulty), then you must also integrate your uncertainty in your aiming point into your final decision.

Often, when we are examining the behaviours of a human subject or animal, we want to compare them against what is optimal (see Chapter 3). Assuming we can overcome the difficulties outlined above, we can move forward with our analysis, and compare experimental evidence against our theoretical predictions. Yet, there is a fundamental hurdle in this comparison. For this type of analysis to be constructive, we must know beforehand what it is the experimental subject is trying to achieve; that is, what their cost function is. With this in mind, researchers have designed experiments where the cost of the task is relatively explicit, such as in the dartboard example. In a set of reaching studies not unlike throwing a dart, it was observed that people are remarkably close to the optimal choices prescribed by decision theory (Maloney *et al*., 2006; Trommershäuser *et al*., 2003, 2005). Similar experiments have studied visual tasks (Najemnik and Geisler, 2005, 2008, 2009) and force-producing tasks (Körding and Wolpert, 2004; Körding *et al*., 2004; see also, Part II). This further demonstrates people's abilities to integrate statistical information in a Bayes-optimal manner, not just for estimation, but also for action selection.

In general though, even when subjects are given explicit rewards and penalties, the cost function they actually use is unknown to us. For instance, one may value a 2 point reward more than twice that of a 1 point reward. Fortunately, the mechanisms of decision theory may be inverted and computational techniques can be used to infer not only the costs subjects use, but also their priors and likelihoods to make their decisions. This pursuit is referred to as inverse decision theory. Though this technique has long been employed in economics (Kagel and Roth, 1995), and agent-based modelling (see Part III) only recently has this approach been used in neuroscience and motor control.

Using inverse decision theory, recent studies have examined both implicit cost functions that subjects use when performing motor tasks (Körding *et al*., 2004) and when penalising target errors (Körding and Wolpert, 2004, see Figure 7.4). These inferred cost functions have highly nonlinear and nontrivial forms (Todorov, 2004). These findings highlight a crucial problem in decision theory; good fits to behaviour may be obtained with incorrect cost functions. Inverse decision theory can thus be used as a means of searching for violations in the assumptions we make using the Bayesian approach to decision making.

As mentioned, inverse decision theory also allows for the analysis of subject priors and likelihoods. For example, studies have found that subjects sometimes use a non-Gaussian prior, and underestimate the speed of visual motion (Stocker and Simoncelli, 2006; Weiss *et al*., 2002). While this approach is enticing for future studies, it also has its weaknesses. Inverse decision theory will always yield a cost function, or likelihood, or prior, for which the observed behaviour is optimal. Like other models built to explain data, over-fitting is also a problem for these decision theoretic models. Researchers must

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

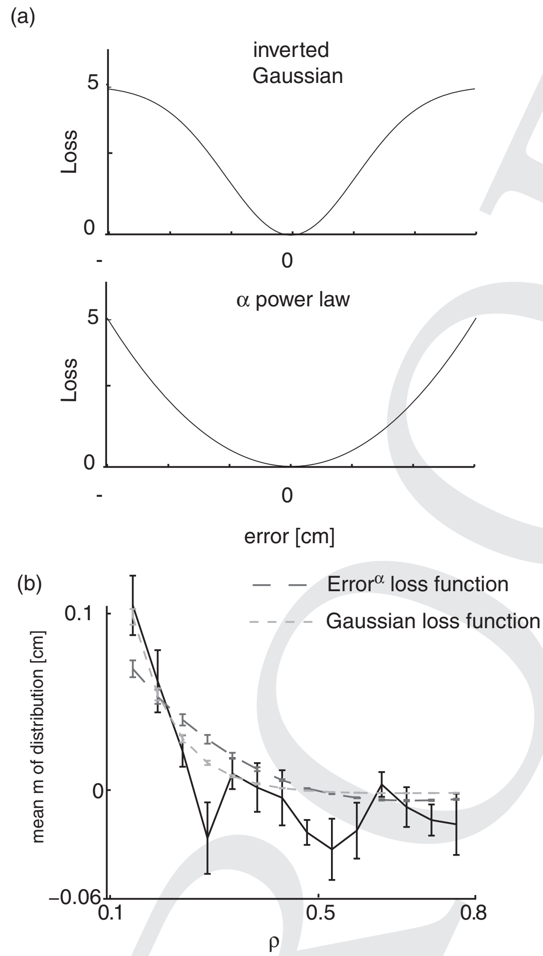

Figure 7.4 Inverse decision theory. (a) cost functions are proposed for penalising motor errors. In this example, the upper cost function penalises errors with an inverted Gaussian, and the lower cost function with a power law. (b) In a motor task the free parameter, the skewness of targets, ρ, is varied to examine subject's performance. The best fits for the two cost functions are overlaid on human data (taken from Körding and Wolpert, 2004).

take steps to ensure this character does not bias their results, or experiments can be designed to independently test these results.

7.6 Bayesian decision theory in dynamical systems: optimal control

7.6.1 The optimal control problem

Up to this point, the actions we've been considering are relatively simple choices, discrete actions that instantly effect outcomes, e.g., where we aim the pea shooter or

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

dart. However, as we pointed out in the introduction, many, if not most, of the choices we take under more realistic conditions are dynamical in nature. Again, a more realistic description of dart throwing considers the forces and motions of the dart, the inertial mechanics of our limb and the force generating properties of our muscles (all of which we can only infer through noisy information). As the state of the dart evolves in time, so too must our actions. Under these more realistic conditions our possible actions move from a range of aiming points, to a list of possible muscles to excite, and the time-history of how these muscles should be excited. Clearly, the relationship between our actions under this scenario and the dart's trajectory and eventual landing place on the board are complicated to say the least.

Under these dynamic conditions, our costs too, are no longer immediately evident, and instead may evolve over time, or may not even be apparent until some delayed period of time after we've taken them. For instance, how hard we throw the dart will influence its path once it leaves our hands and eventually hits its target. Moreover, how hard we throw will influence how quickly we tire, influencing our later throws. How we perform during the dart game may influence our physical and emotional wellbeing for the remainder of the afternoon as well. Taking these considerations into account, one can see how choosing our optimal actions can quickly become a daunting task. Mathematically, we still express our optimal actions in essentially the same form as Equation (7.2), however, now the actions are functions of time and state, and the costs too must be integrated across time. This time- and state-dependent series of actions is often referred to as a policy: a plan of action to control decisions and achieve our desired outcome.

The best possible policy dictates the action we should take at this instant, knowing how it will influence both future states and the actions we can take at that future time. In a sense, to be optimal we need to know the future; we need to be aware of the possible future states and actions that will result from the decision we make now. Under very simple scenarios these future states and actions could be learned through trial and error. However, more generally the future is inferred through models of the world. These models, referred to as forward models, require hypothesis about the nature of the world (Bayesian estimation of structure) and how it evolves (Bayesian estimation across time). Regardless of how the future states are estimated, we need to know how choosing an action based on our current state, will affect future states and the costs associated with them. Mathematically, this is compactly represented by what is referred to as the value function. The value function quantifies the cost of being in any state, and acting according to a given policy for all future times. Not surprisingly, computing this value function is very difficult, as it is akin to seeing into the future. However, once we know this value function, we can relatively easily asses the value of being in any given state, and choose the action that moves us to a state with the best value.

7.6.2 Solution strategies in optimal control

Because this value function is such an integral component to solving for optimal actions, a whole field of numerical techniques has been spawned to achieve this goal (e.g., see Bryson and Ho, 1975; Stuart and Peter, 2003). Difficult to compute analytically, many

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

techniques have been developed to approximate the value function and the accompanying optimal policy. As the problem descriptions get more complicated (e.g., as when we move from choosing where to aim a dart to choosing muscle activation patterns), the likelihood that these techniques will provide an accurate solution diminishes. As such, approximating the optimal actions is often the best one can hope for. Regardless of the difficulty in solving these problems, the same basic principles hold; we represent the necessary statistics of actions and how they impact the statistics of outcomes.

When our observations of the state of the world are relatively certain, we can attempt to approximate the value function directly. There are two widely used approaches. Value iteration is a boot-strapping algorithm, initialised with a naïve value function and repeatedly updated until the estimate is acceptably accurate. Using the current estimate of the value function, the algorithm sweeps across all world states, updating the value of being in that state given that states cost and the value of the best possible future state. The policy is never explicitly computed, instead actions are chosen by evaluating the value function's estimate. In a similar approach, policy iteration uses a current estimate of the value function to compute a policy. This policy is then used to converge on an estimate of the value function. These steps are repeated until both the policy and value function converge. Value and policy iteration techniques are particularly applicable to problems with a finite and discrete number of world states and actions. When our certainty of the world and how our actions will influence it is less certain we must rely on other methods. In reinforcement learning algorithms, rather than using accurate distributions of states and costs, empirical observations are used (Sutton and Barto, 1998). By observing world states and actions taken, an intermediate function is computed that can be used to estimate the optimal policy.

While the above algorithms are very successful for a large class of problems, a more general approach to solving optimal control problems is found through the Hamilton–Jacobi–Bellman equation (or simply the Bellman equation when dealing with discrete time). These equations define the necessary and sufficient conditions for optimality in terms of the value function. Under limited conditions, these equations can be used to find analytical solutions to the optimal policy. For example, the widely popular method of dynamic programming can be used to solve for an optimal policy when the states and actions are discrete and finite. What's more, for linear problems with Gaussian noise, these equations can be used to derive the value function and the optimal policy, the so-called linear quadratic Gaussian regulator (LQG) problem. Under more general conditions, these equations can be used to approximate, perhaps iteratively, the value function and an optimal policy (Stengel, 1994). A wide range of numerical techniques is available to solve problems of optimal control and these techniques generally incorporate Bayesian updating.

7.6.3 Optimal control as Bayesian inference

Interestingly, recent advanced statistical techniques suggest that these traditional methods for computing optimal actions may be supplemented, or even altogether abandoned for what promises to be a truly Bayesian inference formalism. To motivate this new

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

approach, we first point out that in the simplest of cases, it is well known that there is a correspondence between how to choose optimal actions, and how to make optimal inferences. In the linear setting, the optimal estimation problem is the mathematical dual to the optimal control problem (e.g., Stengel, 1994). The defining equations for the solutions to these two problems are identical. Practically speaking, this means the same numerical techniques can be employed to solve either problem. Theoretically, the implication is that choosing an optimal action is a similar problem to optimally estimating the world's state.

Recent work has extended this duality to a larger class of problems. Making some broad assumptions concerning system dynamics and probability distributions of the optimal solutions, this duality has been extended to a large class of nonlinear problems (Todorov, 2008, 2009). There is a correspondence between the distribution of states under optimal conditions, and the solution to the value function. This work further emphasises the implicit connection between optimal actions and optimal inference.

Taking another approach, recent work has neglected a cost function altogether and reformulated the problem of optimal action selection as a problem of optimal Bayesian inference (Friston, 2009; Friston *et al*., 2009; Toussaint, 2009, 2010). This new work replaces the traditional cost function with a desired distribution over states. Actions are selected to reduce uncertainty in our expectations of the world. These new methods, which are purely inference processes, subsume the problem of optimal action selection and optimal inference into a more general Bayesian inference procedure.

7.7 Experimental investigation of learnt statistics

We began this chapter discussing how noise and uncertainty must be considered when choosing our actions. As such, we introduced the Bayesian formalism for making the optimal estimates necessary for choosing our actions. In particular, we reviewed how new evidence should be combined with prior knowledge, and experimental results that demonstrate human subjects appear to do this rationally. We then discussed how these estimates could be updated as evidence is accumulated across time. These estimates could then be used to choose actions according to Bayesian decision theory and optimal control. Having reviewed the necessary theory and background we now present some recent experimental work that examines an important aspect of action selection: representing prior beliefs.

Based on a growing body of evidence, many animals, and humans in particular, appear to make estimates, and choose actions, consistent with the predictions of Bayesian statistics. Implicit in these choices is our prior knowledge; the knowledge we have accumulated throughout life, as well as the knowledge we accumulate during the course of an experimental manipulation, represented as a statistical distribution. While these studies have shown that human subjects can efficiently use prior information, little is known about the way such priors are learned. It is clear priors can change as we accumulate knowledge, however, it is unclear *how* these priors change over time, if they converge to the veridical distribution, and over what timescales they change.

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

We recently set out to examine just these questions. We designed two simple 'coin catching' experiments to examine if subjects could not only accurately estimate a prior, but also the timescale over which it was learned. As many Bayesian experiments rely on a subject's use of an appropriate prior, these results would have an important impact on future studies.

In these experiments we investigated how subjects adapted their expectation of coin locations (a prior) in response to changes in the underlying distribution. Subjects had to estimate the location of a virtual target coin, randomly drawn from a normal distribution. On every trial, subjects were given noisy information of the coins current location, in the form of a single 'cue coin' and were then asked to guess the location of the 'target coin'. To successfully estimate the target coin's location, subjects needed to integrate the coin's likelihood (obtained from the cue coin) with its prior (the distribution of previous target coin locations). By collecting data on where the subjects estimated the target coin to be, we could then estimate the prior used by the subject and analyse its temporal evolution.

In the first experiment, the target coins were drawn from a normally distributed prior. One of two priors was used for each subject; one prior had a relatively narrow distribution, and the other a relatively wide distribution. This first experiment determined whether or not subjects could accurately learn the prior, and if so, the timescale over which it was learned. In the second experiment naïve subjects were recruited to participate in a modified task. Halfway through the experiment, the variance of the prior would switch. The same two variances used in experiment 1 were used. One of these two values was randomly chosen at the start of the experiment and assigned to the prior. After half of the trials the prior's variance would switch to the other value. This allowed us to further probe how subjects learned a prior, and to observe if the learning rate remained constant.

To infer each subject's estimate of the prior's variance, we measured the relative weighting subject's placed on the cue coin relative to the prior's mean, when estimating the target coin's location. This gain, r is a measure of the subject's estimate of the prior (see Section 7.2.1). If subjects believed the prior had a wide distribution, r would be close to 1.0, indicating the cue coin was the best proxy for the target coin's location. Similarly, if the subject's believed the prior had a narrow distribution, the gain r, would be close to 0, indicating the prior's mean was the best proxy for the target coin. In both experiments r was either 0.8 or 0.2. We computed this gain over bins of 10 consecutive trials.

Across all subjects, r took on relatively large values in the first trials, indicating the subject's belief in a 'flat' prior; e.g., the subject's displayed little preference for initially believing the coins would appear in any particular location. However, as the experiment progressed r converged to the correct value. For subjects in the narrow variance group of experiment 1, the data averaged across subjects indicated that approximately 200 trials (20 bins) were required to correctly estimate the variance of the prior (see Figure 7.5c), and on average converged to the correct value. The subjects in wide variance group of experiment 1 essentially began the experiment with the correct gain (see Figure 7.5a). In addition to this analysis, we also inferred the subjects' estimate of the mean of the prior. This too was found to be accurately estimated, albeit on a much faster timescale (data not shown). Using the values inferred for the mean and variance, we were able to

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

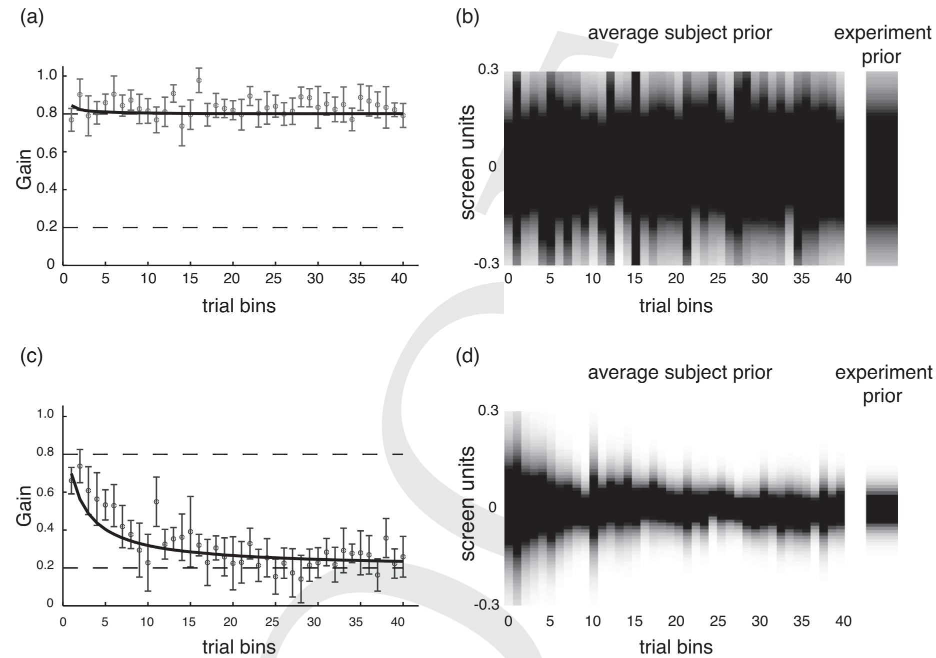

Figure 7.5 Learning a fixed prior. (a) Across-subject averages for the gain, r, in the wide variance group (mean $\pm$ standard error). The bold black line indicates the Bayesian inference model's fits to the experimental data. (b) The inferred average prior as it evolved over the experiment. (c) and (d) Across-subject averages for the gain, r, in the narrow variance group and the inferred prior as it evolved (taken from Berniker *et al*., 2010).

reconstruct the subject's estimated prior during the course of learning (Figure 7.5b and c). This analysis shows that human subjects converge to the correct variance of the prior with a timescale on the order of a hundred trials.

To interpret the subjects' results in a Bayesian framework, we postulated an optimal inference model assigned to the same task, and questioned how it would perform. By observing the cue and target coin locations, the model accurately applied the rules of Bayesian statistics to update a joint distribution of the prior's estimated mean and variance. For this inference, we used an often-assumed normal-scaled inverse gamma distribution. This uses normal distribution to represent a belief in the mean, and an inverse gamma distribution to represent a belief in the variance. The model contains four free parameters that needed to be specified before we could compare its results with those of the subjects. Therefore, we used the data from the first half of the experiment to fit the model's parameters and then proceeded to compute the model's results for the second half of the experiment. With more knowledge about the task than the subjects, the model represents the upper limit in accuracy on what could be observed experimentally.

Just as with the data shown above, during the first 250 trials of experiment 2, the subjects quickly acquired the mean of the prior and the correct variance over a slower timescale, (statistically indistinguishable from the first 250 trials of experiment 1).

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

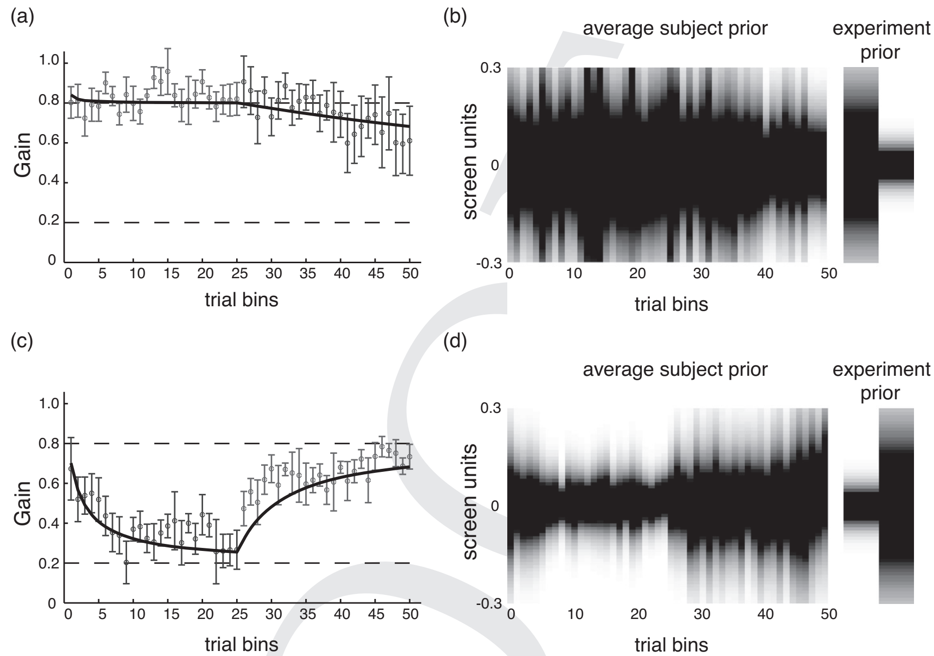

Figure 7.6 Learning a switching prior. (a) Across-subject averages for the gain, r, in the wide variance first group (mean ± standard error). The bold black line indicates the Bayesian inference model's predicted result based on the fit from experiment 1. (b) The inferred average prior as it evolved over the experiment. (c) and (d) Across-subject averages for the gain, r, in the narrow variance first group and the inferred prior as it evolved (taken from Berniker *et al*., 2010).

However, the second 250 trials were distinct, and there were clear differences in the apparent rate at which subjects learned the prior's variance. The learning rate during the second half of experiment 2 was smaller than the apparent learning rate of the first half (or equivalently, the first half of experiment 1). For example, with the wide variance first group, after 500 trials the average behaviour did not yet reflect an accurate estimate for the variance as it did in the first half of the experiment (Figure 7.6a). Inferring the average subject prior using the measured gains and means we could track the prior as it changed over the course of learning (Figure 7.6b). It appeared as if learning the first prior somehow acted to impede subjects' ability to adapt to the second prior. This is especially evident when transitioning from a large variance to a small variance; the change in the subject's prior was gradual (Figure 7.6c and d).

A good fit to experiment 1 was obtained with the Bayesian inference model by minimising the log-likelihood of the subject data (see Figure 7.5a and c). This model was then used to predict the results of experiment 2. The model's behaviour was qualitatively similar to the subjects' (see Figure 7.6a and c, black lines). In particular, note that the inference model correctly predicts a slower learning rate for the second half of the experiment; just as with the subject's data, the model is slow to infer the last half of experiment 2 relative to the first half of experiment 2. Initially, the model is relatively

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

uncertain of the prior and predisposed to estimating large changes in the variance based on the observed distribution of coins. As the experiment progresses, the model's estimate of the variance becomes more certain. By the 250th trial, the model's certainty in the prior's variance makes it insensitive to the new observations of the coin's distribution, now indicating a different variance. As a result, the model is now slow to estimate the new variance, qualitatively similar to the subjects' behaviour.

Overall, we find that subjects can learn an accurate prior, and appear to do so in a rational way. They relatively quickly learn the mean of a distribution (as sample statistics would suggest, the certainty in the mean is higher than the certainty in the variance) and the variance is learned accurately, and in a manner consistent with a Bayesian interpretation. Together this suggests humans can accurately and efficiently learn a prior for making optimal action selections.

7.8 Discussion

Here we have reviewed the results from a wide range of studies that probe cue combination, sensorimotor integration, and motor control using experiments with human volunteers. These studies have found that information such as trial-by-trial error is not sufficient to describe their error. Instead, what has been found is that the statistics of movements play an important role and that quantitative predictions of behaviour can be obtained from the Bayesian formalism.

Estimating the relevant variables needed for motor behaviours happens on many timescales. Traditionally, depending on the timescales involved, different terms are used to refer to this process. Estimations performed on a very fast timescale are usually referred to as information integration; e.g., cue combination or state estimation. Phenomena that require estimations over medium durations are often called motor adaptation; e.g., adapting to a visuomotor disturbance. Estimations over relatively long timescales are usually called motor learning; e.g., learning the structure of a disturbance. These three terms are used widely in different areas of the motor control literature. In this article we have discussed all three phenomena in the same context, highlighting the fact that the same Bayesian ideas can be used to explore many distinct classes of motor behaviour.

Related issues are the phenomena of consolidation and interference. Under some conditions, subjects can acquire a new motor behaviour, and recall the behaviour for later use, even after significant periods of time and with other movements happening in between. The phenomena of preserving a motor behaviour has been termed consolidation (e.g., Brashers-Krug *et al*., 1996). However, it is often found that the act of adapting to one motor behaviour can have adverse effects on our ability to retain a previously adapted motor behaviour; thus one motor behaviour interferes with the consolidation of another. A Bayesian analysis of motor adaptation has already demonstrated how adaptation can lead to new estimates of multiple parameters if they have similar likelihoods (Berniker and Körding, 2008). Thus it could be that adapting to one motor behaviour inadvertently influences the parameters of another. Further, it is not difficult to speculate how prolonged

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

practice with a new motor behaviour increases the certainty in its properties. A Bayesian model would translate this certainty into strong priors, making the behaviour relatively unsusceptible to re-adaptation, a form of consolidation.

As we have shown, the general problem of choosing actions can be modelled as an optimal control problem. This approach has been successfully applied in a number of fields. It is worth noting that different fields sometimes use different formulations to refer to the same problem. As a result, many researchers examine choice selection under the banner of Markov decision processes, while others use the wording of stochastic optimal control. However, the same mathematics, requiring Bayesian statistics, govern both approaches. Moreover, this same formalism is not limited to action selection and motor control, but is also widely used to model cognitive processes and the more general problem of intelligence or even social intelligence (see Part III).

Most Bayesian models discussed above assume that noise sources are Gaussian and that the interactions between the subject and the world have linear dynamics. A set of recent approaches in Bayesian statistics allows a powerful framework well beyond these simple assumptions to be applied. For instance, the structure inference problem is a simple case of non-Gaussian probability distributions. Methods that may be applied to such problems range from linearisation techniques such as extended Kalman filters over variational techniques, to particle filtering methods (Bishop 2006; MacKay, 2003). Clearly more realistic models of the motor system will benefit from further advances in machine learning.

The evidence reviewed for Bayesian integration in the nervous system has been based on psychophysical experiments where human behaviour is measured and the underlying neural processes are unobserved. The Bayesian models proposed are often interpreted as archetypes for the neural computations responsible for the motor behaviours. Recent studies though are looking for evidence of these computations at the neuronal level. Purely computational studies now show how neurons might easily solve some of the needed computations (Ma *et al*., 2006). Other studies are looking for neural correlates of uncertainty (Cisek and Kalaska, 2005; Gold and Shadlen, 2001, 2003; Kiani and Shadlen, 2009) and electrophysiological evidence of multimodal cue combination in the context of uncertainty (Angelaki *et al*., 2004). Future studies might focus on imaging or electrophysiological studies to find more evidence for the neural substrates of these computations (see also, Part II).

The human nervous system clearly has limitations and it would be very surprising if our behaviour were optimal for all tasks. Indeed, many past experiments have found deviations from optimality. For example, in situations requiring economic decisions, human subjects systematically overestimate small probabilities and exhibit so-called framing effects (Kahneman and Tversky, 1979). As a result of these biases it is known that we are not optimal at speculating on, for example, the stock market. The phenomena that we reviewed here where near optimal Bayesian behaviour is observed are all rather low-level phenomena – simple perceptual and sensorimotor tasks. The optimality observed for these phenomena does not appear to carry over to the range of high level decision making tasks that have been characterised in behavioural economics. However, we have yet to find the limits of the Bayesian approach for modelling action selection.

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

7.9 Conclusions

Recognising the inherent uncertainty in motor tasks, many motor control problems have been successfully described by applying the rules of Bayesian statistics to analyse their solutions.

This same method holds promise for many of the outstanding problems in cue combination, motor control, cognitive science, and neuroscience, in general. Bayesian statistics and normative modelling are complementary to descriptive studies of the nervous system. As such, combining the two approaches will provide models that not only predict, but explain how the nervous system solves the problems it is confronted with. This insight into the purpose, and detail of structure, will provide a deeper understanding of the nervous system.

References

Alais, D. and D. Burr (2004). The ventriloquist effect results from near-optimal bimodal integration. *Curr. Biol*. **14**(3): 257–62.

Angelaki, D. E., A. G. Shaikh, A. M. Green, *et al*. (2004). Neurons compute internal models of the physical laws of motion. *Nature* **430**(6999): 560–4.

Berniker, M. and K. Körding (2008). Estimating the sources of motor errors for adaptation and generalization. *Nat. Neurosci*. **11**(12): 1454–61.

Berniker, M., M. Voss, and K. Körding (2010). Learning priors for Bayesian computations in the nervous system. *PLoS One* **5**(9): e12686.

Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Berlin: Springer.

Brashers-Krug, T., R. Shadmehr, and E. Bizzi (1996). Consolidation in human motor memory. *Nature* **382**(6588): 252–5.

Bryson, A. E. and Y. C. Ho (1975). *Applied Optimal Control: Optimization, Estimation, and Control*. London: Taylor and Francis Group.

Chater, N., J. B. Tenenbaum, and A. Yuille (2006). Probabilistic models of cognition: where next? *Trends Cogn. Sci*. **10**(7): 292–3.

Cisek, P. and J. F. Kalaska (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**(5): 801–14.

Deneve, S. (2008). Bayesian spiking neurons I: inference. *Neural Comput*. **20**(1): 91–117.

Ernst, M. O. and M. S. Banks (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**(6870): 429–33.

Ernst, M. O. and H. H. Bulthoff (2004). Merging the senses into a robust percept. *Trends Cogn. Sci*. **8**(4): 162–9.

Flanagan, J. R., J. P. Bittner, and R. S. Johansson (2008). Experience can change distinct size-weight priors engaged in lifting objects and judging their weights. *Curr. Biol*. **18**: 1742–7.

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn. Sci*. **13**(7): 293–301.

Friston, K. J., J. Daunizeau, and S. J. Kiebel (2009). Reinforcement learning or active inference? *PLoS One* **4**(7).

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

Geisler, W. S. and D. Kersten (2002). Illusions, perception and Bayes. *Nat. Neurosci*. **5**(6): 508–10.

Gold, J. I. and M. N. Shadlen (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci*. **5**(1): 10–16.

Gold, J. I. and M. N. Shadlen (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J. Neurosci*. **23**(2): 632.

Gopnik, A., C. Glymour, D. Sobel, *et al*. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychol. Rev*. **111**(1): 3–32.

Griffiths, T. L. and J. B. Tenenbaum (2005). Structure and strength in causal induction. *Cognit. Psychol*. **51**(4): 334–84.

Hoyer, P. O. and A. Hyvärinen (2003). Interpreting neural response variability as Monte Carlo sampling of the posterior. In *Neural Information Processing Systems*, Vol 15. Cambridge, MA: MIT Press.

Kagel, J. H. and A. E. Roth, eds (1995). *The Handbook of Experimental Economics*. Princeton, NJ, Princeton University Press.

Kahneman, D. and A. Tversky (1979). Prospect theory: an analysis of decision under risk. *Econometrica* **XVLII**: 263–291.

Kersten, D. and A. Yuille (2003). Bayesian models of object perception. *Curr. Opin. Neurobiol*. **13**(2): 150–8.

Kiani, R. and M. N. Shadlen (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* **324**(5928): 759.

Knill, D. and W. Richards, eds. (1996). *Perception as Bayesian Inference*, Cambridge: Cambridge University Press.

Knill, D. C. (2007). Robust cue integration: a Bayesian model and evidence from cue-conflict studies with stereoscopic and figure cues to slant. *J. Vis*. **7**(7): 5 1–24.

Körding, K. (2007). Decision theory: what 'should' the nervous system do? *Science* **318**(5850): 606–10.

Körding, K., J. Tenenbaum, *et al*. (2006). A generative model based approach to motor adaptation. In *Neural Information Processing Systems*, Vol 15. Cambridge, MA: MIT Press.

Körding, K. P., I. Fukunaga, I. Howard, *et al*. (2004). A neuroeconomics approach to inferring utility functions in sensorimotor control. *PLoS Biol*. **2**(10): e330.

Körding, K. P., S. P. Ku, and D. M. Wolpert (2004). Bayesian integration in force estimation. *J. Neurophys*. **92**(5): 3161–5.

Körding, K. P. and D. M. Wolpert (2004). The loss function of sensorimotor learning. *Proc. Natl. Acad. Sci. USA* **101**(26): 9839–42.

Körding, K. P. and D. M. Wolpert (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn Sci*. **10**(7): 319–26.

Körding, K. P. and D. M. Wolpert (2004). Bayesian integration in sensorimotor learning. *Nature* **427**(6971): 244–7.

Ma, W. J., J. M. Beck, P. E. Latham, and A. Pouget (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci*. **9**(11): 1432–8.

MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press.

Maloney, L. T., J. Trommershäuser, *et al*. (2006). Questions without words: A comparison between decision making under risk and movement planning under risk. *Integrated Models of Cognitive Systems*. W. Gray. New York, NY, Oxford University Press, pp. 297–314.

Mehta, B. and S. Schaal (2002). Forward models in visuomotor control. *J. Neurophysiol*. **88**(2): 942–53.

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

Michotte, A. (1963). *The Perception of Causality*. London: Methuen.

Miyazaki, M., D. Nozaki, and Y. Nakajima (2005). Testing Bayesian models of human coincidence timing. *J. Neurophysiol*. **94**(1): 395–9.

Miyazaki, M., S. Yamamoto, S. Uchida, and S. Kitazawa (2006). Bayesian calibration of simultaneity in tactile temporal order judgment. *Nat. Neurosci*. **9**(7): 875–7.

Najemnik, J. and W. S. Geisler (2005). Optimal eye movement strategies in visual search. *Nature* **434**(7031): 387–91.

Najemnik, J. and W. S. Geisler (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *J. Vis*. **8**(3): 1–14.

Najemnik, J. and W. S. Geisler (2009). Simple summation rule for optimal fixation selection in visual search. *Vision Res*. **49**: 1286–94.

Peterka, R. J. and P. J. Loughlin (2004). Dynamic regulation of sensorimotor integration in human postural control. *J. Neurophysiol*. **91**(1): 410–23.

Pouget, A., P. Dayan, and R. S. Zemel (2003). Inference and computation with population codes. *An. Rev. Neurosci.* **26**: 381–410.

Rosas, P., J. Wagemans, M. O. Ernst, and F. A. Wichmann (2005). Texture and haptic cues in slant discrimination: reliability-based cue weighting without statistically optimal cue combination. *J. Opt. Soc. Am. A* **22**(5): 801–809.

Shams, L., W. J. Ma, S. Tanaka, *et al*. (2005). Sound-induced flash illusion as an optimal percept. *Neuroreport* **16**(17): 1923–7.

Sperber, D. and D. Premack (1995). *Causal Cognition: A Multidisciplinary Debate*. Oxford: Oxford University Press.

Stengel, R. F. (1994). *Optimal Control and Estimation*. New York: Dover Publications.

Stocker, A. A. and E. P. Simoncelli (2006). Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci*. **9**(4): 578–85.

Stuart, J. R. and N. Peter (2003). *Artificial Intelligence: A Modern Approach*. Harlow, UK: Pearson Education.

Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning*. Cambridge, MA: MIT Press.

Tassinari, H., T. E. Hudson, and M. S. Landy (2006). Combining priors and noisy visual cues in a rapid pointing task. *J. Neurosci*. **26**(40): 10154–63.

Tenenbaum, J. B., T. L. Griffiths, and C. Kemp (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci*. **10**(7): 309–18.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nat. Neurosci*. **7**(9): 907–15.

Todorov, E. (2006). Optimal Control Theory. In *Bayesian Brain*, ed. K. Doya. Cambridge, MA, MIT Press, pp. 269–98.

Todorov, E. (2008). General duality between optimal control and estimation. Proceedings of the 47th IEEE Conference on Decision and Control, Cancun, Mexico, pp. 4286–92.

Todorov, E. (2009). Efficient computation of optimal actions. *Proc. Natl. Acad. Sci. USA* **106**(28): 11478–83.

Toussaint, M. (2009). Robot trajectory optimization using approximate inference. 26th International Conference on Machine Learning (ICML 2009), ACM New York, NY, USA, pp. 449–56.

Toussaint, M. (2010). *A Bayesian View on Motor Control and Planning*. Berlin: Springer.

Trommershäuser, J., S. Gepshtein, L. T. Maloney, M. S. Landy, and M. S. Banks (2005). Optimal compensation for changes in task-relevant movement variability. *J. Neurosci*. **25**(31): 7169–78.

Trommershäuser, J., L. T. Maloney, and M. S. Landy (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *J. Opt. Soc. Am. A* **20**: 1419–33.

Trim: 247mm × 174mm Top: 12.653mm Gutter: 16.871mm
CUUK1505-07 CUUK1505/Seth ISBN: 978 1 107 00049 0 June 23, 2011 11:16

van Ee, R., W. J. Adams, and P. Mamassian (2003). Bayesian modeling of cue interaction: bistability in stereoscopic slant perception. *J. Opt. Soc. Am. A* **20**(7): 1398–406.

Wei, K. and K. P. Körding (2008). Relevance of error: what drives motor adaptation? *J. Neurophys*. **90545**.2008.

Weiss, Y., E. P. Simoncelli, and E. H. Adelson (2002). Motion illusions as optimal percepts. *Nat. Neurosci*. **5**(6): 598–604.

Wolpert, D. M., Z. Ghahramani, and M. I. Jordan (1995). An internal model for sensorimotor integration. *Science* **269**(5232): 1880–2.

Wozny, D. R., U. R. Beierholm, and L. Shams (2008). Human trimodal perception follows optimal statistical inference. *J. Vis*. **8**(3): 1–11.

Zemel, R. S., P. Dayan, and A. Pouget (1998). Probabilistic interpretation of population codes. *Neural Comput*. **10**(2): 403–30.