Bayes rule in perception, action and cognition

Daniel M. Wolpert and Zoubin Ghahramani
Department of Engineering, University of Cambridge, Trumpington Street, Cambridge

Cognition and intelligent behaviour are fundamentally tied to the ability to survive in an uncertain and changing environment. Our only access to knowledge about the world is through our senses which provide information that is usually corrupted by random fluctuations, termed noise, and may provide ambiguous information about the possible states of the environment. Moreover, when we act on the world through our motor system, the commands we send to our muscles are also corrupted by variability or noise. This combined sensory and motor variability limits the precision with which we can perceive and act on the world. Here we will review the framework of Bayesian decision theory, which has emerged as a principled approach to handle uncertain information in an attempt to behave optimally, and how this framework can be used to understand sensory, motor and cognitive processes.

Bayesian Theory is named after Thomas Bayes (Figure 1) who was an 18th century English Presbyterian minister. He is only known to have published two works during his life, of which only one dealt with mathematics in which he defended the logical foundation of Isaac Newton's methods against contemporary criticism. After Bayes death, his friend Richard Price found an interesting mathematical proof among Bayes' papers and sent the paper to the Editor of the *Philosophical Transactions of the Royal Society* stating "I now send you an essay which I have found among the papers of our deceased friend Mr Bayes, and which, in my opinion, has great merit...." The paper was subsequently published in 1764 as "Essay towards solving a problem in the doctrine of chances." In the latter half of the 20th century Bayesian approaches have become a mainstay of statistics and a more general framework has now emerged, termed Bayesian decision theory (BDT).

Figure 1 Thomas Bayes 1702-1761

Bayesian decision theory incorporates two main components: Bayesian statistics and decision theory:

Bayesian statistics

Bayesian statistics is a framework for making inferences based on uncertain information. The fundamental idea is that probabilities can be used to represent the degree of belief in different propositions and therefore the rules of probability can be used to update beliefs based on new information. For example, imagine you are a doctor and see a patient who you believe may have a particular disease. The strength of this belief can be represented by a real number between 0 and 1, which reflects the probability that we assign to our belief. We can denote this as our prior belief $P(A)$—the probability of event A, that the patient has the disease, is true. Now, imagine we get new information in the form of a positive blood test. Let us denote by B, the event that the blood test is positive. Given this information we would like to update our belief about A, that is, we want to compute our belief in the proposition that the patient has the disease given that we now know he has a positive blood test. We can write this as $P(A|B)$ where the vertical line is shorthand for "given". Bayes theorem provides the rule for computing this:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Here $P(B|A)$ is the probability of observing a positive blood test given that a patient has the disease (termed the likelihood for A), and $P(B)$ is the normalization constant which sums up the probability of observing a positive blood test in this patient under both the possibility that the patient has and doesn't have (indicated by $\bar{A}$) the disease: $P(B) = P(B \mid A)P(A) + P(B \mid \bar{A})P(\bar{A})$.

In its most general sense there is nothing mystical about Bayes rule, it is simply a consistency relationship between probabilities that can be derived straightforwardly. For any two events A and B we can ask what is the probability that both events occur, which is written as $P(A,B)$ where the comma is shorthand for "and". This can be written in two alternative forms, that is the probability of one of the events happening, e.g. $P(A)$, multiplied by the probability of the second event happening given that the first has happened $P(B \mid A)$. This gives three expressions which are all identical: $P(A,B) = P(A \mid B)P(B) = P(B \mid A)P(A)$. By discarding the left hand term and dividing both the others by $P(B)$ gives Bayes rule as above. Although the form of Bayes rule may be intuitive, its implications are extensive and often counter-intuitive.

How does such a rule apply to the brain? If we relabel A as the "state of the world" and B as "sensory input", then Bayes rule becomes applicable to the brain:

$$\overbrace{\mathrm{P(state|sensory\ input)}}^{\text{Posterior}} = \frac{\overbrace{P(\text{sensory input}|\text{state})}^{\text{Likelihood}}\ \overbrace{P(\text{state})}^{\text{Prior}}}{P(\text{sensory input})} \qquad (1)$$

The state of the world could represent something we wish to estimate such as the location of our limbs, the identity of a visual object or the possibility that our poker opponent is bluffing. The input to the system comes from our sensory inputs. Clearly we would like to estimate the probability of different states given out sensory input (left hand side equation 1). Using Bayes rule we can write this as more intuitive components (right hand side). These are $P(\text{state})$, that reflects our *prior* belief in particular states before we receive the sensory information. For example, if we walk out into the street our prior belief that we will encounter a polar bear compared to a white van may be one to a million. Such prior beliefs could be learned through experience of the world. However, once sensory input is received we can use it to update our beliefs. To do this Bayes rule states that the *likelihood* of the state, that is the probability of the sensory input given the hypothesized state should be calculated $\mathrm{P(sensory\ input|state)}$. That reflects how probable it is that we would receive, for example, the current visual input given a polar bear or given a white van. By multiplying the prior by the likelihood and normalizing (scaling this output so that the sum of the probabilities over all possible states sum to one) we can estimate the probability of each state given the sensory input, termed the posterior of the state. This posterior now become our new prior belief and can be further updated based on new sensory input.

Perception and action

We can now illustrate a specific example of where Bayes rule can be used to make optimal estimate of the state of the world. Consider playing tennis and trying to estimate the location where a fast-moving incoming ball will bounce (Figure 2). Bayes rule tells us that to estimate the posterior probability of the ball's location given our visual inputs we need to consider the elements on the right hand side of equation 1. Because vision does not provide perfect information about the ball's location or speed there will be uncertainty about the bounce location from considering the sensory inputs alone. Therefore the likelihood, the probability of the sensory feedback we have received given different possible bounce locations, will have a distribution over the court (shown in red with saturation proportional to likelihood). However, Bayes rule states that we can reduce the uncertainty in this purely sensory estimate if we have a prior over possible bounce locations. For example, over the course of a match with a particular opponent or over our general experience of playing tennis we will find that the ball does not bounce equally often at all locations. Therefore from previous experience of tennis we can have a prior about where our opponent may hit the ball (shown in blue). To compute the posterior for each possible location the value of the prior at that location is multiplied by the value of the likelihood and then normalised to produce the posterior distribution (contours). The most probable location is then the peak of the posterior distribution (purple sphere). Recent experimental data has shown that during motor learning subjects learn the prior of tasks and combine this with sensory evidence in a Bayesian fashion[1].

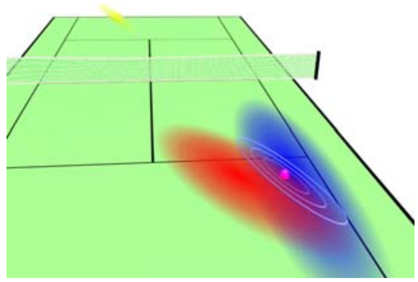

Figure 2: An example of Bayes rule used to estimate the bounce location of an incoming tennis ball. From vision we can estimate the likelihood of different bounce locations (left hand ellipse). Prior knowledge may suggest that balls tend to land in another region (right hand ellipse). Integrating these two sources of information gives the contour ellipses that denote the posterior, and the peak of this (ball in centre of inner ellipse) indicated the most probable bounce location. Models adapted from 3dcafe.com.

In addition to performing Bayesian estimation in motor tasks, Bayes rule underlies many perceptual processes which generate a percept from sensory inputs which are often incomplete, ambiguous or noisy. For example, consider the image of one surface of a die in Figure 3. Interpreting the three dimensional nature of the image is the well known shape-from-shading problem [2]. Most readers will see 4 circular bumps and 2 circular depressions. However, if the entire page is rotated upside down, most readers will now have the percept of 2 circular bumps and 4 circular depressions. Why does the percept flip in this way? The key point is that the image alone is consistent with the circles being bumps or depressions, and the interpretation depends on the location of the light source. As we live in a world where most light sources, such as the sun or a room light, tend to come from above we have a prior that in the absence of further information scenes will be illuminated from above. Given that the brain assumes a light-from-above prior, the shading in Figure 3 is consistent with 4 bumps and 2 depressions but the opposite when turned upside down. Recently it has been shown that this prior is adaptable—if objects such as shown in Figure 3 are illuminated from below and subjects are allowed to touch the bumps and depression (thereby resolving the ambiguity), they can adapt their prior and change their percept when determining hills from valleys using shading[3].

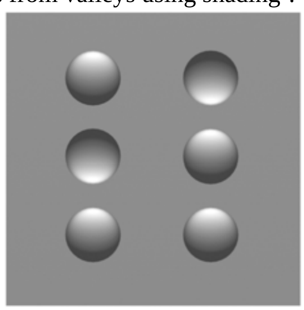

Figure 3. How many of the spots are bumps and how many are holes? If the page is turned upside-down does the answer change?

Visual illusions are in general very informative about brain processing as they are examples of the visual system producing apparently erroneous percepts. However, applying Bayes rule to visual perception has lead to an understanding of the nature of some (but not all) visual illusions. In a Bayesian framework these percepts, while erroneous in the specific case, are nevertheless optimal in a general sense, given uncertainty inherent in the visual input. For example, the pattern of light and shade in Figure 3 is optimally interpreted by the visual system as arising from circular bumps and depressions. Therefore, rather than visual illusions arising from faulty processing, recent studies have shown that several visual illusions arise from a rational system that is designed to make judgements in the face of uncertainty. For example, a recent study has shown that motion illusions can arise as optimal percepts in the face of uncertainty[4].

Cognition

Recently, Bayesian models have been used to model our cognitive ability to learn concepts from a few examples[5]. For example, consider a number game where you are given samples from a set of numbers and asked whether a new number belongs to that set. For example, does 34 belong to the set which includes 31, 33, 35 and 37? If you believe that set contains the odd numbers then the answer would be no, but if you believe the set was all numbers between 30 and 40, then the answer would be yes. By defining a prior over each possible way that the set can be generated: odd numbers, primes numbers, perfect squares, intervals etc. and combining this with the likelihood of observing the data under each of these hypotheses about the sets it is possible to generate the posterior, that is the probability that the new element is a member of the set. This Bayesian formulation allows the remarkably complex patterns of human behaviour in the number game to be predicted. A similar formulation can be used to model other forms of human reasoning, such as the determination of casual relationships.

Decision theory

Bayes rule is an inference method which computes beliefs about the state of the world given sensory input. However, how these beliefs are used to generate decisions or actions is not specified by Bayes rule alone. Decision theory deals with the problem of selecting the decision or action based on our current beliefs. The essence is to minimize the expected loss (or maximizes expected reward/utility) given our beliefs. This loss function quantifies the value of taking each possible *action* for each possible *state* of the world, L(action,state). To choose the best action one simply calculates the expected loss for a given action, that is the loss averaged across the possible states weighted by the degree of belief in the state.

$$\sum_{\text{states}} \mathrm{L(action,state)P(state|sensory\ input)}$$

Where $\sum$ denotes a summation over all possible states. We can then choose the action which has the smallest expected loss. For example, imagine being offered the pufferfish delicacy Fugu at a Japanese sushi restaurant. You know that during the Fugu preparation that the natural poison used by the fish to paralyse predators is sometimes not adequately removed. This leads to around 1 person in 10,000 become seriously ill from the dish which gives a probability of illness if you eat the Fugu of 0.0001. Should you eat the Fugu? The answer depends on your loss function, that is how your rate illness compared to the pleasure of eating the Fugu. Suppose you regard the loss of becoming ill from Fugu as 5,000 and the loss of eating good Fugu as -1 (the negative loss reflects pleasure). Then the expected loss of taking the action to eat the Fugu is

L(eat, bad Fugu) P(bad Fugu) + L(eat, good Fugu) P(good Fugu)

which is 5,000 x 0.0001 - 1 x (1-0.0001) = 0.5-0.9999 = -0.4999. The expected loss if you choose not to eat the Fugu is clearly zero. The action of eating the Fugu gives you the lowest expect loss and you should therefore take the risk.

It is possible to scale this simple concept to the problem of making a sequence of decisions over time to minimize the expected loss over the accumulated lifetime of a human. However, in practise the optimal solutions for such sequential processes are intractable to even the fastest computers. The challenge in recent years has been to develop efficient and accurate approximations to the optimal solution. Recently, a new field of neuroeconomics[6] has emerged which considers neural processing within the framework of Bayesian decision theory.

In summary, evidence has been mounting that the behaviour of the human perceptual, motor and cognitive systems are captured remarkably well by Bayesian models which provide a framework for understanding information processing in the brain.

Acknowledgements: We thank Paul Bays, James Ingram and Lili Tcheang for helpful comments on the paper.

References

1. Kording, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244-7 (2004).
2. Ramachandran, V. S. Perceiving shape from shading. *Sci Am* 259, 76-83 (1988).
3. Ernst, M. O. & Bulthoff, H. H. Merging the senses into a robust percept. Trends Cogn Sci 8, 162-9 (2004).
4. Weiss, Y., Simoncelli, E. P. & Adelson, E. H. Motion illusions as optimal percepts. Nat Neurosci 5, 598-604 (2002).
5. Tenenbaum, J. B. & Griffiths, T. L. Generalization, similarity, and Bayesian inference. Behav Brain Sci 24, 629-40; discussion 652-791 (2001).
6. Glimcher, P. Decisions, Uncertainty, and the Brain : The Science of Neuroeconomics (MIT Press, 2003).