Addressing Bayesian statistics and Utility functions in psychophysical experiments
(with particular focus on the motor system)

Konrad Paul Körding

3 Cambridge Center
Massachusetts Institute of Technology
Building NE20-388
Cambridge, MA 02139 USA
Tel: (617)452-2548
eMail: konrad@koerding.com

Motivation

The sensors of humans are plagued by noise, and likewise our muscles produce noisy outputs. Even if we had perfect sensors they would only tell us about the part of the world that we can currently sense. Uncertainty that stems from noise or the lack of complete knowledge places estimation of attributes of the worlds and control of our actuators firmly within a statistics framework. Bayesian statistics is the principled way of dealing with such problems.
Much research in the field of machine learning addresses how artificial systems can optimally solve such problems . Bayesian learning and inference are thriving fields and the developed algorithms now allow efficient handling of complicated real world data.
Recently many counterintuitive finding, such as visual illusions, in neuroscience can be understood by considering the brain is performing Bayesian estimation. At the same time studies start to appear where the neural correlate of such statistical processing is analyzed. Therefore the Bayesian approach may provide a framework to understand and macroscopic and microscopic attributes of neural function.
In the studies described in this paper it was analyzed how human performance can be understood in terms of statistically optimal processing. These studies address how the human nervous system can use information to perform optimally in a world full of uncertainties.

Description of the theory

Bayesian statistics can be used to infer the states of variables that are not directly measured, a process called *Bayesian inference*. The way inference is done is the following:

1) The *prior knowledge* about the system is specified. The structure of the problem is defined, specifying which variables depend on which other variables in which way. 2) Bayesian methods are used to infer how probable each state of the unobserved variables is given the values of the observed variables. There are numerous methods available towards this goal. Some methods deliver an approximate answer efficiently while some other methods deliver exact solutions to impler problems.

Numerous articles summarize the mathematical(Cox 1946) and philosophical(Freedman 1995) ideas that are behind Bayesian statistics and how to use them. The main focus of this article is on the concepts behind Bayesian methods that are used to understand the human brain.

Notation:

Upper case letters denote random variables (e.g. the position of my hand), lower case letters denote a particular value that a random variable can take(e.g. 10cm).

$p(A = a)$ is the probability that a random variable A takes on a specific value a, given that there is no additional knowledge about other related variables.

$p(A = a \mid B = b)$ is the probability of observing A=a given that the statistician has the knowledge that the variable B has the value of b (*conditional probability*).

$p(A=a, B=b)$ is the probability that *A*=*a* and *B*=*b*. It can be rewritten to be:

$$p(A=a, B=b) = p(A=a)\,p(B=b \mid A=a)$$

Conditional independence and graphical Bayesian models

Statistical problems become a lot easier once we understand their structure. If we have 3 variables, *A*, *B* and *C* we could for example know that *A* causes *B* and *C* and that *B* and *C* do not interact. As an example, *A* could be the size of an object, *B* is the size as reported by our sense of touch alone and *C* is the size reported by our visual system alone. B and C thus show *conditional independence* on A. As many problems have such structure often graphs are drawn to depict how variables depend on one another. There is a range of variants of such graphs, but for the problems considered here *graphical Bayesian networks* (directed acyclic graphs) are used. If there exists no direct link between two variables *B* and *C* then they are independent conditional on the other variables that are "in-between": p(B=b|C=c,Others=o)=p(B=b|Others=o).

Using Bayes rule: ball positions in tennis

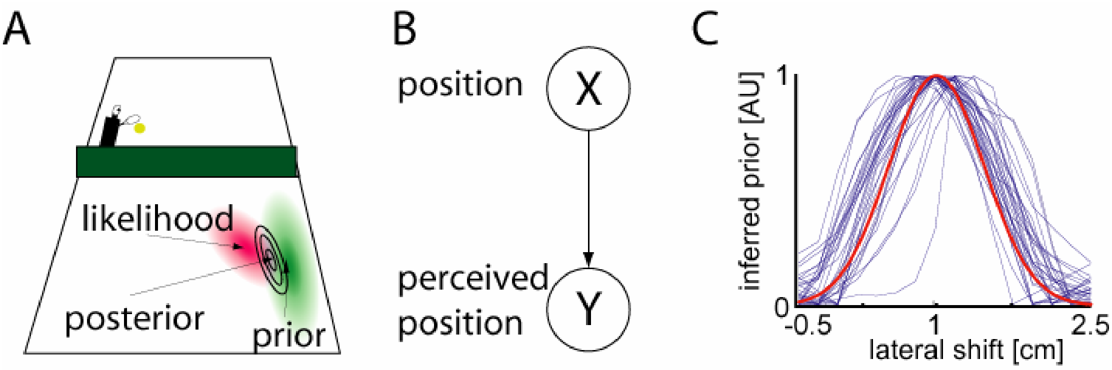

Figure 1: A) Example: The other player is hitting the ball. Seeing the ball we can estimate that it will land in the red region (with a likelihood proportional to the saturation). We have prior knowledge that the ball is likely to land in the green region (with a probability proportional to the saturation). The black ellipses denote the posterior, the region where the Bayesian estimate would predict the ball to land. B) Structure: The position influences the perceived position. C) Human subjects had to estimate a position. From these measurements their priors were inferred (blue lines). The real distribution is shown in red.

If we are playing tennis we want to estimate where the ball will hit the ground (see Fig. 1A). Because vision does not provide perfect information about the balls speed there is uncertainty about where the ball will land. If we, however, played a lot of tennis before we can have knowledge about where our partner is likely to play to. We want to combine this knowledge with what we see to obtain an optimal estimate of where the ball will hit the ground.
This example can be abstracted the following way (Fig. 1B): The physical properties of the ball define the position $X = x$ where the ball will hit the ground. The visual system, however, does not perceive where the ball will really hit the ground but rather some noisy version thereof, *Y*=*y*. Knowing the uncertainties in the visual system we know how likely it is to perceive the ball being at *Y*=*y* if it really is at *X*=*x*'. This is called *likelihood* (p(Y=y|X=x'|)) and is sketched in red in Fig. 1A. Just based on this knowledge we could ignore any other knowledge we might have and our best estimate would be in the middle of the red cloud, at *y*. This procedure however ignores that we can have prior knowledge about the way our partner plays. In particular over the course of many matches the positions where the ball hits the ground will not be uniformly distributed but highly concentrated within the confounds of the court and if our enemy is a good player highly peaked near the boundary lines where it is most difficult to intercept them. This distribution of positions where the ball hits the ground , p(*X*=*x*), is what is called a *prior* and which could be learnt using methods of *Bayesian learning*. We can apply *Bayes Rule* to compute the *posterior* ($p(X = x \mid Y = y)$), the probability of the ball landing taking into account the prior and the new evidence*:*

$p(X = x \mid Y = y) = p(Y = y \mid X = x)\dfrac{p(X = x)}{p(Y = y)}$. Our uncertainty about its position is thus set in terms of probability. We know where we can expect the ball to land with which probability, given everything we know about it.

If we can assume that the prior distribution $p(X = x)$ is a symmetric two dimensional Gaussian with variance σ_p^2 and mean $\hat{\mu}$ and the likelihood $p(Y = y \mid X = x)$ is also a symmetric two dimensional Gaussian with variance σ_V^2 and mean *y* it is possible to compute the optimal estimate $\hat{x}$ as:

$\hat{x}=\alpha y+(1-\alpha)\hat{\mu}$ where $\alpha=\frac{\sigma_p^2}{\sigma_p^2+\sigma_V^2}$. It is thus possible to define the optimal estimate given seen and prior knowledge. But not only is it possible to calculate what the optimal strategy is, it is also possible to calculate how much better the estimate is compared to a strategy ignoring prior knowledge: The variance of the estimate if only the visual feedback is used is σ_V^2 if however the prior is used to the variance is $\frac{\sigma_p^2}{\sigma_p^2+\sigma_V^2}\sigma_V^2$ which is always less than the variance of the non-Bayesian estimate. If the prior has the same variance as the likelihood then the variance of the Bayesian estimate is half the variance of the non-Bayesian estimate.

In a recent experiment (Körding & Wolpert 2004) it was tested if people use such a Bayesian strategy. Subjects had to estimate the value of a one dimensional variable, the displacement of a cursor relative to the hand, in close analogy of estimating the position where the ball will land. This variable was drawn randomly each trial out of a Gaussian distribution defining a *prior* distribution and subjects received extensive training so that they can know the *prior*. In addition they saw the position thus receiving the *likelihood*. The experiment allowed measuring the subject's estimate of the displacement. Out of this data it is possible to infer the prior that people are using. If they ignore the prior information the prior should be flat. The data shown in Fig. 1C (blue) shows that people used a prior that was very close to the optimal one (shown in red). This experiment thus showed that people can use Bayes rule to estimate the state of an important variable.
It has been shown in other experiments that people equally use Bayes rule when they estimate a force (Körding et al 2004 - in press). It is possible to use the same method used here for combining prior knowledge with new evidence to combine information obtained from two different sensors, for example feeling (haptic information) and seeing something and then having to estimate its size. A recent experiment also showed that people use Bayesian methods to combine visual and haptic information (Ernst & Banks 2002).

Kalman Controllers: hand position in the dark

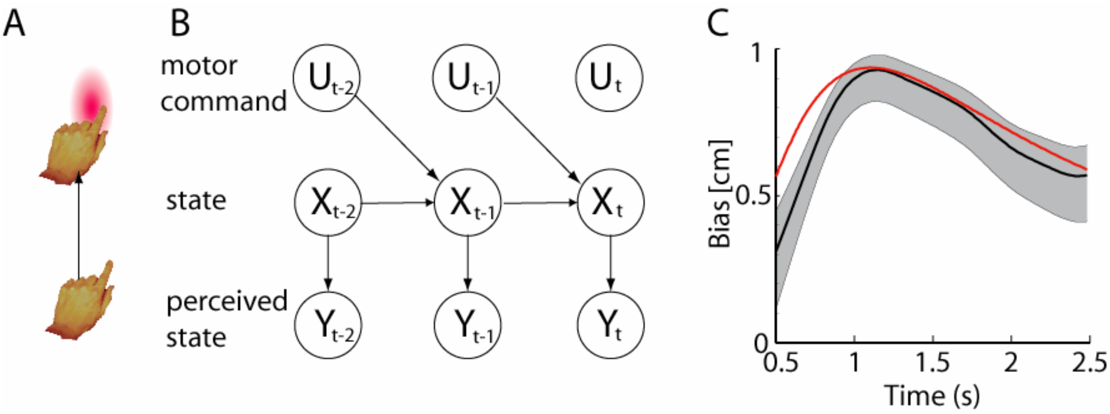

Figure 2: A) Example: The hand is moving in the dark, we want to estimate where it is. The uncertainty about the state of the hand is sketched in red. B) Structure: The state at time t is influenced only by the previous state and the motor command. The perceived state is the state with added noise. C) The error of the human estimates of travelled distance is shown as a function of the duration of the movement. The optimal (assuming overestimated force) Kalman controller predicts the red curve.

If we move our hand in the dark we have uncertainty about its exact position or velocity because our proprioceptive sensors are not perfect (Fig 2A). We are thus faced with the problem of estimating the state **X** of the hand, which is characterized by the position and velocity. From experience people know how the state of the hand changes over time. In particular they know that the position changes proportional to its velocity and that the velocity changes proportional to the force applied, although there might be noise in this process. The subjects thus have *prior* knowledge that they can combine with the *likelihood* provided by their sensors. In this case however they constantly have to update their estimate thus effectively using Bayes rule at every time-step.
The model for the hand is shown in Fig. 2B. Assuming that the real state X of our hand is X=x then the perceived state Y will be some noisy version thereof and U=*u* is the motor command we are sending to our muscles. We know something about the structure of the problem: The state of the hand at time t only depends on the state of the hand at time t-1 and the applied force, this is the *Markov property*. The state of the hand does not explicitly depend on the state of the hand at any but the preceding time. Using Bayes rule twice it is possible to obtain:

$$p(X_t = x_t \mid X_{t-1} = x_{t-1}, Y_t = y_t, U_t = u_t) \approx p(X_t = x_t \mid X_{t-1} = x_{t-1}, U_t = u_t)\, p(Y_t = y_t \mid X = x_t)$$

Where $p(X_t = x_t \mid X_{t-1} = x_{t-1}, M = m_t)$ is the probability of finding oneself in state x_t at time t after having been in state x_{t-1} at *t*-1. People can be expected to have a model for their hand in terms of these variables (called *forward model*) that they acquired from past experience.
Figure 2B shows the graph with the relations between the variables.
Assuming that random variables have n-dimensional Gaussian distributions, the equations obtained describe the optimal *Kalman controller* and it is possible to derive the optimal strategy(Wolpert et al 1995) for predicting the state of the hand:
$\dot{\hat{x}}(t) = \hat{A}\hat{x}(t) + \hat{B}u(t) + K(t)\left[y(t) - \hat{x}(t)\right]$ where $\hat{x}$ is the current optimal estimate, $\dot{\hat{x}}$ is the change of the optimal estimate, $\hat{A}$ is a matrix that characterizes how the hand moves without perturbation, $\hat{B}$ a matrix that characterizes how forces change the state of the hand $K(t)$ is the *Kalman gain* which is a function of the other matrices. The Kalman controller is a generalization of the *Kalman filter* (Kalman 1960) that does not allow a motor signal

In an experiment human volunteers subjects (Wolpert et al 1995) moved their hands in the dark. After each movement they had to estimate where their hand was. Movements of varying temporal duration were done between 500 ms and 2500 ms. Subjects systematically estimated that their hand had moved further than it actually had moved (Fig. 2C gray). An optimal Kalman controller (Fig 2C, red) produced very similar results if it was assumed that people systematically overestimate their forces. For small times the overestimation of distance increased with time. This was due to the overestimated forces. As times however increased, the likelihood becomes more important compared to the prior. That is why the controller becomes better if the movement lasts a long period of time. The optimal controller thus shows very similar errors to those made by human subjects. It thus seems that people are able to continuously update their estimates based on information coming in from the sensors.

Dealing with rewards:

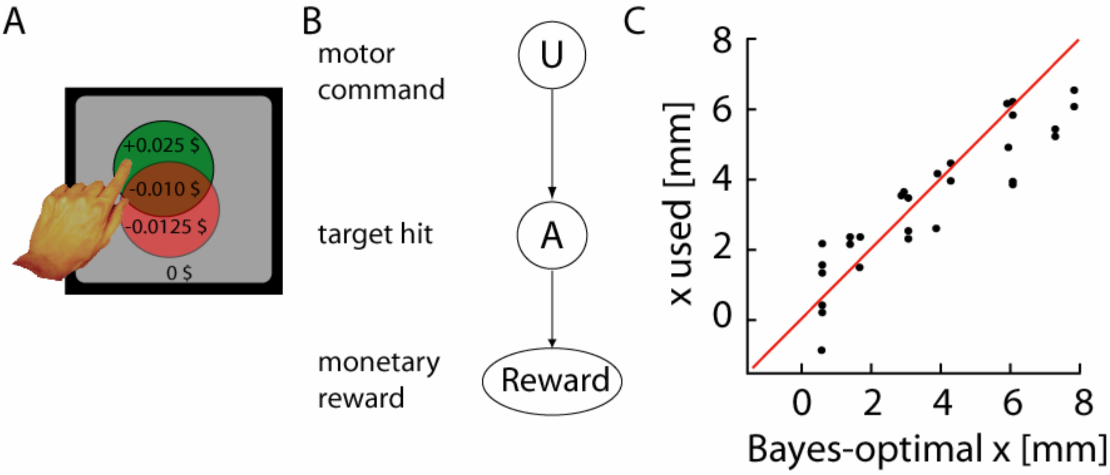

Figure 3: A) Example: People rapidly move their hand to a computer screen to touch circles and gain rewards. The circles are quite small and they have to move so fast that they can not be certain about their final finger position. B) Structure. Depending on the motor commands there is a probability distribution about which circles are hit. Depending on which circles are hit there are monetary rewards. C) The positions towards which people moved are shown against the positions they should have pointed to assuming the use of Bayesian decision theory.

In many situations there are not only probabilities involved but also costs or rewards. Consider for example throwing darts. While the 20 gives us a large number of points the neighbours only give a low number of points (1 and 5). The 19 however has neighbours that give more points (3 and 7). An expert thrower will therefore usually try to hit the triple 20 as she is unlikely to miss it. An intermediate player will try and go for the triple 19 while a novice is best off targeting the middle of the board. Such problems thus ask for a combination between statistical theory and information about rewards, called *Bayesian Decision Theory*.
Figure 3A shows a simple abstraction of such a task that has been used by Trömmershauser and Maloney(Trommershäuser et al 2003). People have to move their hand very fast towards a target, just like in darts, so that uncertainty can not be avoided. There are 2 circles that people can touch. Touching one leads to a monetary gain, touching the other leads to a monetary loss. Depending on the motor command *U* there will be a probability distribution of hitting each target p(*A*=*a*|*U*=*u*). Figure 3B shows the structure of the problem. Optimal behaviour for such tasks can be derived calculating the expected value of the reward:

$\text{Expected Reward} = \int p(A=a|\text{knowledge})\text{Reward}(a)da$ where Reward(a) is the number of points scored that way. The optimal behaviour is then defined as the one that leads to a maximal expected gain. Bayesian decision theory thus emerges naturally as probabilities are combined with rewards.

A recent experiment tested if human subjects are able to use Bayesian decision theory to move optimally(Trommershauser et al 2003). Subjects had to touch a touchscreen monitor very fast that displayed two spheres, one defining a monetary loss, the other one a monetary gain. They then measured where on average subjects touched the screen. Figure 3C shows the position where the screen was touched against the optimal position according to the theory. Subjects are very close to the optimal solution. This demonstrates that people can use Bayesian decision theory to move optimally.

Neuroeconomics: internal rewards

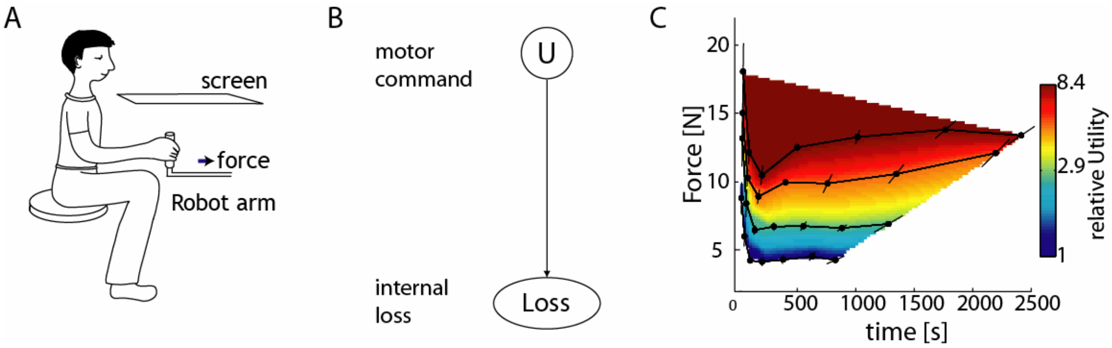

Figure 3: A) The experiment holds a robotic handle that produces forces that the subject has to hold. B) Structure of the problem. Depending on the force and how long it needs to be held there will be some internal loss. C) The loss can be inferred from the subjects decisions. The hotter the color the less desirable the force.

If we move our hands there will not only be a target that we want to reach but at the same time we will try and minimize our effort. Each movement, for example while shaking hands with a robot (Figure 3A) will require energy as we have to produce forces. It can be assumed that people will generally prefer less demanding movements. People are thus faced with the problem of selecting among the huge set of possible movements the one that minimizes their effort or loss. It is assumed that there is a function – Loss – that defines how much loss is associated with a movement. This loss function will depend on various parameters of the movement for example the on the magnitude and the duration that subjects have to hold a force (Figure 3B). If only the magnitude (F) and the duration (t) of a force are considered important then the loss will be a function of these two parameters (Loss=f(F, t)). This function however is defined in our own head and it is not straightforward to measure it. The economists, however, are often faced with similar problems. Assume people have a fixed amount of money and can only buy apples and bananas, how many of each should they buy? An important concept that is used in economics is the concept of a utility function. This function defines how good the outcome of an action, for example a trade is. It is thus just the inverse of the loss. A tool that is often used to study such functions is called indifference curve. A set of combinations of apples and bananas can be equally desirable and people have no specific preference for any specific such combination. They could for example be equally content with having 2 bananas and 3 apples or 3 apples and 2 bananas. Measuring such indifference lines can often elucidate properties of the underlying utility function. Such curves can be inferred by asking people for judgments of preference, for example: Would you rather have 2 bananas and 3 apples or 1 banana and 4 apples). A large number of such preference statements allow inferring the indifference curves and the utility function. The same concept can also be used for the study of loss functions in the sensorimotor system. In an experimental setting people had to produce forces of varying magnitude and duration. They then had to decide which force they had to experience again(Körding et al 2004). From their responses, just like in economics, it is possible to infer the utility function and the indifference lines (Figure 3C). The loss function that people show is relatively complicated for this task. This is in contrast with previous theoretical explorations that have proposed relatively simple utility functions (Harris & Wolpert 1998). The method explained here thus allows directly addressing the properties of the utility function.

Internal rewards and uncertainty:

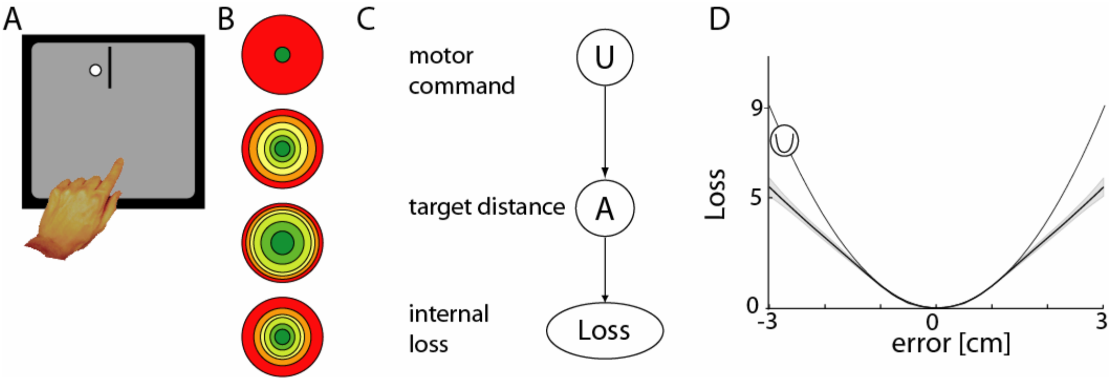

Figure 4: A) Example: White spheres appear on a computer screen the position of which is steered by a hand. The errors and their distribution, however, is controlled by a computer program. B) An important question is how bad an error is compared to another error. Several possible virtual dartboards are sketched. C) Structure of the problem. Depending on the motor commands there is a probability distribution about target distances. D) The loss function inferred in this experiment

If we are playing darts then the dart board will assign a score to each of our throws. When we have to point to a position in space (Figure 4A) we also want to point as precisely as possible. Because our sensorimotor system is not perfect we will nevertheless not always be perfectly precise in pointing. We are nevertheless sometimes content or annoyed about a movement. Computationally, there must be some way how the brain assigns a loss to each movement. In other words there must be a dartboard that assigns a number to each error we are making. An important question is how this internal dartboard looks like and how bad, for example, it is to make an error of 2 cm compared to an error of 1 cm (Figure 4B,C). Measuring the properties of this internal dartboard ask for a combination between statistical theory and information about rewards, called *Bayesian Decision Theory*. To measure the loss function that people have the following experiment was used. Many dots like in Figure 4A appeared on the screen. The average position of the dots was determined by the position of the hand of the subject. The dots varied around this average position in a way that was not symmetric but skewed. In this case different loss functions predict different optimal behaviour. Analyzing the way that people do move it was possible to infer the loss function they are using (Figure 4D). The loss function is approximately quadratic for small errors and significantly less than quadratic for large errors. The human estimation is thus outlier insensitive. Such loss functions are very useful because a couple of outliers does not significantly influence the estimate. Consequently they are often used in artificial techniques and a field called *robust fitting* (Huber 1981) analyzes how such functions could best be built.

Outlook

Bayesian statistics is a large field (c.f. MacKay 2003). Some researchers develop novel algorithms for Bayesian inference. They developed fast algorithms that lead to exact solutions, a classical one is belief propagation(Pearl 1988). There are a large number of novel uses of this algorithm for complicated problems where the graph does have loops called loopy belief propagation which often results in good approximations. As an example, the technologically important issue of efficient data transmission through noisy information channels is best solved using this method (Berrou et al 1993). Bayesian methods are a systematic way of thinking about function approximations and thus supersede the neural networks literature(Neal 1996). Bayesian methods are also the best known methods for speech recognition using *hidden Markov models* (Rabiner 1989) that are a discrete version of the Kalman filter. Bayesian statistics is also moving towards being a major driving factor at understanding the brain. There is ample evidence that the human visual system uses Bayesian methods (Fleming et al 2003, Weiss et al 2002). Reflections of such statistical processing can already be seen in the way that neurons in the primary visual cortex process stimuli (Sharma et al 2003). Reflections of such processing can also be seen in neurons involved into the decision about the direction of movement of a visual stimulus(Gold & Shadlen 2002). There is furthermore evidence that Bayesian statistics is not only used by people to move or perceive. It is indeed used to solve higher level cognitive tasks such as generalization (Tenenbaum & Griffiths 2001), concept learning(Tenenbaum 2000) and the judgement of randomness(Griffiths & Tenenbaum (in press)). Bayesian statistics allows novel questions to be asked in psychophysical and electrophysiological experiments while providing a coherent framework in which information processing can be understood.

Experiments that address the way the brain deals with statistical problems with and without rewards help us understand the algorithm the brain uses to solve its everyday problems. The approach suggests novel experiments as well as novel interpretations of previous experiments and places a broad range of experimental results in a coherent theoretical framework.

References

Berrou c, Glavieux A, Thitimajshima P. 1993. Near Shannon limit error-correcting coding and decoding: turbo-codes. *Proc ICC'93, Geneva, Switzerland* 1: 1064-70

Cox RT. 1946. Probability, Frequency and Reasonable Expectation,. *American Journal of Physics* 17: 1-13

Ernst MO, Banks MS. 2002. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415: 429-33

Fleming RW, Dror RO, Adelson EH. 2003. Real-world illumination and the perception of surface reflectance properties. *J Vis* 3: 347-68

Freedman DA. 1995. Some issues in the foundation of statistics. *Foundations of Science* 1: 19–83

Gold JI, Shadlen MN. 2002. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36: 299-308

Griffiths TL, Tenenbaum JB. (in press). *From algorithmic to subjective randomness.* Presented at Advances in Neural Information

Processing Systems

Harris CM, Wolpert DM. 1998. Signal-dependent noise determines motor planning. *Nature* 394: 780-4.

Huber PJ. 1981. *Robust Statistics*. New York: Wiley

Kalman RE. 1960. A new approach to linear filtering and prediction problems. *J. of Basic Engineering (ASME)* 82D: 35-45

Körding KP, Fukunaga I, Wolpert D. 2004. A neuroeconomics approach to measuring human loss functions. *(submitted)*

Körding KP, Ku SP, Wolpert D. 2004 - in press. Bayesian Integration in force estimation. *Journal of Neurophysiology*

Körding KP, Wolpert DM. 2004. Bayesian integration in sensorimotor learning. *Nature* 427: 244-7

MacKay DJC. 2003. *Information Theory, Inference, and Learning Algorithms*. Cambridge, UK: Cambridge University press, also available online at http://www.inference.phy.cam.ac.uk/mackay/itila/book.html

Neal RM. 1996. *Bayesian Learning for Neural Networks*. New York: Springer-Verlag

Pearl J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*: Morgan Kaufmann

Rabiner L. 1989. A tutorial on Hidden Markov Models and selected applications in speech recognition. *Proc. IEEE* 77: 257-86

Sharma J, Dragoi V, Tenenbaum JB, Miller EK, Sur M. 2003. V1 neurons signal acquisition of an internal representation of stimulus location. *Science* 300: 1758-63

Tenenbaum JB. 2000. *Rules and similarity in concept learning*, MIT Press, Cambridge

Tenenbaum JB, Griffiths TL. 2001. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences* 24: 629-41

Trommershauser J, Maloney LT, Landy MS. 2003. Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A Opt Image Sci Vis* 20: 1419-33

Trommershäuser J, Maloney LT, Landy MS. 2003. Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A*, 20: 1419-33

Weiss Y, Simoncelli EP, Adelson EH. 2002. Motion illusions as optimal percepts. *Nat Neurosci* 5: 598-604

Wolpert DM, Ghahramani Z, Jordan MI. 1995. An internal model for sensorimotor integration. *Science* 269: 1880-2.