

Role of uncertainty in sensorimotor control

Robert J. van Beers, Pierre Baraduc and Daniel M. Wolpert*

Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London, Queen Square, London WC1N 3BG, UK

Neural signals are corrupted by noise and this places limits on information processing. We review the processes involved in goal-directed movements and how neural noise and uncertainty determine aspects of our behaviour. First, noise in sensory signals limits perception. We show that, when localizing our hand, the central nervous system (CNS) integrates visual and proprioceptive information, each with different noise properties, in a way that minimizes the uncertainty in the overall estimate. Second, noise in motor commands leads to inaccurate movements. We review an optimal-control framework, known as ‘task optimization in the presence of signal-dependent noise’, which assumes that movements are planned so as to minimize the deleterious consequences of noise and thereby minimize inaccuracy. Third, during movement, sensory and motor signals have to be integrated to allow estimation of the body’s state. Models are presented that show how these signals are optimally combined. Finally, we review how the CNS deals with noise at the neural and network levels. In all of these processes, the CNS carries out the tasks in such a way that the detrimental effects of noise are minimized. This shows that it is important to consider effects at the neural level in order to understand performance at the behavioural level.

Keywords: noise; uncertainty; motor control; motor planning; multisensory integration

1. INTRODUCTION

Neural signals are noisy and this limits the amount of information that a signal can contain. The noise at the single neuronal level translates into uncertainty and variability at the higher, cognitive level. For instance, noise in sensory information about an object’s location leads to uncertainty in the position at which the object is perceived, and noise in motor commands results in movement inaccuracy and variability. Noise at the neural level therefore has direct effects at the cognitive level. In this paper, we explore these effects. We mainly use goal-directed arm movements as an example to illustrate the effects.

Goal-directed movements require several different processing steps. First, the target and the hand have to be localized. Second, motor commands have to be determined that can bring the hand to the target position. Finally, the motor commands have to be sent to the arm muscles, resulting in a movement. Neural noise is present at all of these stages and a main theme of this review is that the strategy used to achieve these processes is the one that minimizes the detrimental effects of neural noise. The emerging view therefore is that the neural control of movements, but possibly of other tasks as well, has evolved to maximize fitness. Through natural selection, movement control may have been ‘improved’ until performance was limited by biophysical constraints, thereby reaching a global or local optimum. Clearly, neural noise is one of the

important constraints. Therefore, some of the CNS’s functioning at the higher, cognitive level can be understood from the properties of the underlying activity at the lower, neural level.

In this review, we first describe how noise in sensory signals limits spatial perception and how information from different sensory modalities is combined so as to minimize the perceptual uncertainty. We then describe a control framework called TOPS, which formulates how goal-directed movements are planned on the basis of the expected effect of noise in motor commands. We show how the nervous system combines noisy sensory signals and noisy motor output signals during a movement to allow the state of the body to be estimated, and show how these mechanisms can be used to reduce the uncertainty about whether sensory information is due to changes induced by one’s own action or by events in the outside world. Finally, we explain how uncertainty can be encoded and decoded by neural networks.

2. SENSORY NOISE

If we want to make a goal-directed reaching movement, the nervous system first has to have spatial information about the target and the hand. Their positions are important, as are their orientation, size, shape, etc. Sensory information is used to estimate these quantities.

The amount of information that sensory signals convey about the outside world and about the state of one’s body is limited. At the neural level, this is reflected by noise in the neural signals. Neural noise can lead to two kinds of imperfections that can be quantified at the higher level: accuracy, which refers to constant errors, and precision, which refers to variable errors and uncertainty. Here, we

* Author for correspondence (wolpert@hera.ucl.ac.uk).

One contribution of 14 to a Discussion Meeting Issue ‘The physiology of cognitive processes’.

mainly address limitations on the precision because these are a direct consequence of neural noise.

Precision is limited by the properties of the sensory receptors. For instance, the size of photoreceptors in the retina sets a limit to visual acuity. Because continuous signals in the outside world must be coded by discrete spikes, noise may be added at later processing stages, which further reduces the precision. Spatial and geometric factors can also play an important role in shaping the characteristics of the precision at the perceptual level. This will be illustrated for visual and proprioceptive information about the position of one's hand.

Visual and proprioceptive localization have been studied in the horizontal plane at waist level (van Beers *et al.* 1998). It was found that, for both modalities, precision depends on the position and, for a given position, the precision also depends on the direction considered. For visual localization, precision decreases with increasing distance from the observer and localization is less precise in depth than in azimuth. This reflects that it is more difficult for the visual system to judge distance than direction (Foley & Held 1972), which is also partly due to the geometry of the setting in which subjects look slantwise down on the table. It is different for proprioceptive localization. Precision decreases with increasing distance from the shoulder, but here, localization is more precise in depth than in azimuth. This is mainly a geometric effect (van Beers *et al.* 1998). Assuming that proprioceptive signals reflect joint angles, the observed pattern can be understood by transforming these signals from joint angles to hand positions (known as the kinematic transformation). Consider, for instance, an almost extended arm (see figure 1). Uncertainty in both the shoulder and elbow angles translates into uncertainty in hand position mainly in azimuth, with less uncertainty in the distance from the shoulder. According to such principles, the 2D precision should vary with arm posture, and this has been observed experimentally (van Beers *et al.* 1998). The precision of both visual and proprioceptive localization thus simply reflects noise in the sensory signals that is translated into the perceptual domain. This is, in essence, also true for auditory localization, which is most precise directly in front of and behind the observer and decreases monotonically towards the periphery (Fedderson *et al.* 1957). This follows from how properties of wave sounds moving around the head vary with direction, but the situation here is more complicated because it also depends on the frequency and, especially for estimating elevation, on the spectral shape of the sound (Middlebrooks & Green 1991).

When localizing the hand on the basis of simultaneous visual and proprioceptive information, the question arises as to how the CNS integrates the information from different sources. This issue of multisensory integration is of fundamental importance because in everyday life there is an abundance of information from various senses, and combining this into a single percept is one of the complex tasks that our CNS accomplishes continuously. The mechanism can be understood from the non-uniform precision of visual and proprioceptive localization, as explained above. When the right hand is in front of and to the left of the body, the visual and proprioceptive precision ellipses are approximately orthogonal to one another




Figure 1. Illustration of the relationship between noise in proprioceptive signals about shoulder and elbow angles and the resulting uncertainty in finger localization. Ellipses represent the precision of proprioceptive localization. The narrower an ellipse in a certain direction, the more precise is the localization. For the almost extended right arm shown on the left, noise in information about the shoulder (dashed lines) and elbow angles (dotted lines) results in uncertainty in approximately the same direction. This results in a relatively high precision in depth, and low precision in azimuth. For the more flexed right arm shown on the right, noise in information about shoulder and elbow angles has effects in different directions. In addition, the effect of shoulder noise is much less than for the extended arm because the distance between the hand and the shoulder is smaller. Consequently, the precision ellipse is smaller and has a different orientation.

(figure 2*a*). In general, visual and proprioceptive localization have different constant errors (Warren & Schmitt 1980); therefore the visual and proprioceptive ellipses have been plotted at different locations in figure 2*b*. One can now ask where, and with what precision, is the hand localized if both senses provide information simultaneously? There are several possibilities. First, one modality could completely dominate, and the information from the other modality may not be used at all. Second, the CNS could calculate the arithmetic mean of the positions sensed by each modality. Finally, the CNS could calculate a more complicated average.

Absolute dominance of one modality would imply that the hand is localized exactly the same as if the other modality were not present. Simple averaging, irrespective of the direction-dependent precision, would mean that it is localized on the straight line between the centres of the visual and proprioceptive ellipses. In the case of an arithmetic mean, the percept would be exactly halfway between the two centres, whereas it could be at other places on the straight line were a weighted mean calculated (the larger a modality's weight, the closer it will be localized to the centre of that ellipse). The optimal weighted mean (i.e. the one producing the smallest variance) is illustrated by the dashed circle in figure 2*b*. This seems a reasonable




Figure 2. Optimal integration of visual and proprioceptive information about the position of the right hand. (a) Top view of a subject. For this hand position, the ellipses representing the precision of visual and proprioceptive localization (not to scale) are approximately orthogonal. (b) The visual and proprioceptive ellipses have been plotted at different locations to reflect that vision and proprioception generally have different biases. If integration of visual and proprioceptive information would amount to weighted averaging, irrespective of the direction-dependent precision, the seen hand would be localized on the straight (dashed) line. The circle labelled 'suboptimal integration' represents the best (i.e. producing the smallest variance) way in which this can be done. However, optimal integration does take the direction-dependent precision into account, and this predicts that localization of the seen hand will be as indicated by the ellipse labelled 'optimal integration'. This is the smallest ellipse that can be obtained, indicating that this method minimizes the uncertainty in the overall estimate.

way to fuse the information, but in 2D space there is a better way. The maximum-likelihood estimate (Ghahramani *et al.* 1997), based on the 2D precision of visual and proprioceptive localization, is illustrated by the smallest ellipse in figure 2b (van Beers *et al.* 1999). Surprisingly, the centre of this ellipse is not on the straight line between the centres of the visual and proprioceptive ellipses. This can be understood from the orientation of these ellipses. The fusion can be considered as a more complicated weighted averaging in which the weights vary with direction. For instance, in the direction of the major axis of the proprioceptive ellipse, vision is more precise than proprioception and therefore the visual weight is larger than the proprioceptive weight. In the direction orthogonal to that, however, the proprioceptive weight is larger. This explains why the ellipse lies off the straight line.

There is strong experimental support for this model. The predictions for both the constant errors (the mean lying off the straight line) and the variable errors (the variance being smaller than can be expected from any direction-independent weighted mean) have been confirmed experimentally (van Beers *et al.* 1996, 1999). Recently, it has also been tested directly whether the weighting varies with direction (van Beers *et al.* 2002). The visual and proprioceptive weights have been estimated many times using a classical method. In the paradigm of prism adaptation (Welch 1978), subjects view their hand through optical prisms that displace the visual field, inducing a conflict between vision and proprioception. This leads to adaptation of the visual and proprioceptive mappings in order to resolve or reduce the conflict. The magnitudes of the adaptation in the two modalities are a measure for their weights. With normal vision, proprioceptive adaptation has usually been found to be larger

than visual adaptation (Welch & Warren 1986). Because the dominant modality will adapt least, this finding has led to the belief that, for spatial localization, vision dominates proprioception. This finding is in accord with the model, because traditional prism adaptation creates a conflict between vision and proprioception in azimuth, which is in a direction in which localization is 'best' for vision and 'worst' for proprioception. The model predicts that the relative weighting and thus the relative adaptation varies with direction. In line with this prediction, it was found that the relative visual adaptation in depth was larger than in azimuth. More than that, proprioception was found to be weighted more heavily than vision in depth.

There is evidence that the CNS uses similar mechanisms to integrate other types of information. For instance, visual and auditory localization (Ghahramani *et al.* 1997) and visual texture and motion cues to depth (Jacobs 1999) are combined according to the 1D equivalent of the mechanism explained above. Therefore, in integration, all of the available information is used and the CNS fuses it in a way that minimizes the uncertainty in the overall estimate. This indicates that the way in which the CNS integrates information from different sources can only be understood from the effects of uncertainty, and therefore of noise at the neural level.

3. MOTOR NOISE

We now return to goal-directed arm movements. When the hand and the target have been localized, the movement can be planned. Movement planning involves determining the motor commands that will produce the intended movement. A movement is usually specified at a high, symbolic level, such as: 'pick up that glass and drink from it'. However, the motor system works at a low, much

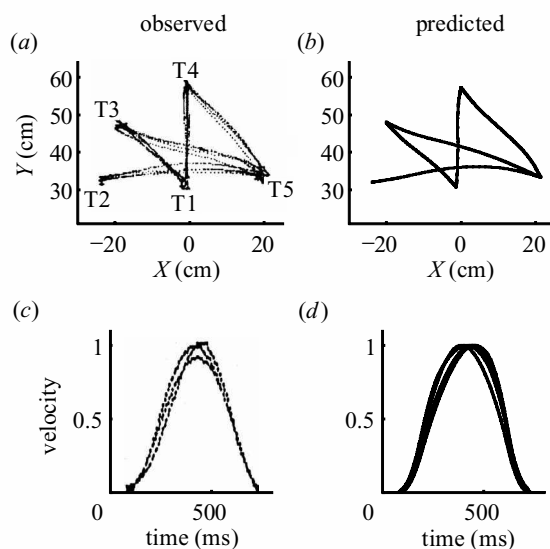


Figure 3. Comparison of observed and predicted trajectories for goal-directed arm movements in the horizontal plane. (a) Observed hand paths for five different point-to-point movements (from Uno *et al.* 1989). The origin of the coordinate system is on the shoulder. X and Y directions represent the transverse and sagittal axes, respectively. (b) The optimal hand paths as predicted by the TOPS framework for the same movements as in (a). (c) Observed velocity profiles for movements from T1 to T3 in (a). The other movements had similar bell-shaped velocity profiles. (d) Velocity profiles of all the predicted movements shown in (b), normalized to have a maximum velocity of 1.

more detailed level; it specifies temporal profiles of muscle activations. There is a large gap between these high- and low-level specifications, and the question is how the nervous system bridges this gap. We present a control framework called TOPS, which formulates how goal-directed movements specified at the high level are planned on the basis of properties of the motor system at the low level (Harris & Wolpert 1998).

Given a task, the motor system usually has an infinite number of ways in which to achieve it, due to redundancy in the motor system. This can be illustrated if we consider the simple task of moving the hand from one point in space to another. Such a movement can be made with a range of durations and an infinite number of paths between the start and final hand locations. Given the path, the hand could move along it with infinitely many speed profiles and, for each point on the path, the hand can be placed there with a infinite set of different arm configurations. Similarly, the arm can be held in a given posture stiffly, with opposing muscles co-contracting, or with low co-contraction levels. Finally, the same muscle tension can be generated by different patterns of neural firing. Motor planning can be considered as the problem of selecting one solution from the infinity of possibilities.

Despite the redundancy, humans produce very stereotyped movements. Hand trajectories are, in general, very smooth, with approximately straight paths (compared with the corresponding trajectories in joint space) and bell-shaped velocity profiles (Morasso 1981; see also figure 3). Certain movements, however, have a consistent but small amount of curvature (Atkeson & Hollerbach 1985). The faster a movement is executed, the worse is its precision;

this relation is characterized by Fitts' law (Fitts 1954). For curved movements, such as in writing and drawing, there is a relation between the curvature and the movement speed, known as the two-thirds power law (Lacquaniti *et al.* 1983). It is interesting to note that saccadic eye movements are also stereotyped, but in a different way. Saccades have approximately symmetrical velocity profiles, although for longer saccades velocity profiles are skewed towards the end of the movement (Collewyn *et al.* 1988). In addition, there is a relationship between the duration and peak velocity of a saccade and its amplitude, known as the main sequence (Bahill *et al.* 1975). Finally, the ocular system has three degrees of freedom but effectively uses them as if there were only two, because the eye position defines the torsion, which is known as Donders' law (Donders 1848) and Listing's law (von Helmholtz 1867).

Stereotypical patterns of movement are, therefore, the end result of motor learning or evolution, indicating that movements may have been optimized to maximize fitness (Harris 1998). Movements are 'improved' until the performance is limited by biophysical constraints, thereby reaching an optimum. The question then is in what way are arm and eye movements optimal? The idea taken from optimal-control engineering is to define a cost function. The cost can be determined for every possible movement and the optimal movement is the one with the lowest cost. Movement planning then amounts to selecting the movement with the lowest cost. For eye movements, it has been proposed (Enderle & Wolfe 1987) that the cost is movement time. The rationale was that vision is very degraded during a saccade; keeping movement time to a minimum would minimize the time for which we are deprived of vision. For arm movements, the observed smoothness has led to the idea that the cost function is the mean-squared jerk (the temporal derivative of the acceleration) of the hand (Hogan 1984; Flash & Hogan 1985), or the mean-squared rate of change of the joint torques (Uno *et al.* 1989). Although these cost functions predicted the observed movements quite well, there were some problems associated with them. First, they seem arbitrary quantities; it is not clear why it would be advantageous to minimize jerk. Second, these quantities are difficult to calculate for the CNS. Third, why would the costs be different for the eye and the arm? The ideal cost would be similar for all systems, simple to compute and have some evolutionary advantage.

In TOPS, the cost is the variability in movement endpoints, or, in other words, the expected movement error. This seems to be a sensible quantity because the goal of a movement is to reach the target and this cost directly represents how well this is achieved. In addition, it can be applied to all systems and is easy to compute, because the nervous system usually gets direct feedback about movement errors.

The other assumption in the TOPS framework is that the endpoint variability is due to noise in the motor commands. It assumes signal-dependent noise, which is the standard deviation in the motor command signal is proportional to its magnitude (constant coefficient of variation). This is an important assumption. Constant noise, for instance, would not work because that would predict that the faster a movement is executed, the more precise it will be. This is the converse of the empirical

Fitts' law. Signal-dependent noise, however, predicts Fitts' law because a faster movement requires larger motor commands and larger commands are noisier. A smaller endpoint variability can thus be obtained by using smaller motor commands, which results in a longer movement time. The assumption of signal-dependent noise is supported by the empirical finding that the standard deviation in isometric force production is proportional to the mean force (Schmidt *et al.* 1979; Meyer *et al.* 1988). It has been shown (Jones *et al.* 2002) that the basic physiological organization of the motor-unit pool, such as the range of twitch amplitudes and the range of recruitment thresholds, is responsible for this relationship.

The idea behind TOPS is that the CNS aims to minimize the consequences of noise in the motor system. Movement planning uses the redundancy of the motor system to reduce the endpoint variability. Movements predicted by this principle correspond to actual movements. For example, it predicts approximately straight finger trajectories and bell-shaped velocity profiles that match very well with observed trajectories (see figure 3). Optimal trajectories are inherently smooth, because abrupt changes in the trajectory would require large motor commands that would carry more noise than motor commands for smoother trajectories. In addition, Fitts' law and the two-thirds power law are also predicted by TOPS. This framework also works very well for rapid eye movements; for instance, it reproduces saccadic velocity profiles and the main sequence. Moreover, it was recently shown that it also reproduced observed trajectories for the more complicated situation in which both the eye and the head move simultaneously in order to fixate a peripheral target (Wolpert & Harris 2001).

All of these findings strongly support the TOPS framework. TOPS has more predictive power than the various other cost functions proposed for arm and eye movements because it can easily be generalized to other types of movements. The framework is a biologically plausible theoretical underpinning for both eye and arm movements with no need to construct highly derived signals such as jerk to estimate the cost of a movement. Instead, in the TOPS framework, variance of the final position is the cost and this cost is almost directly available to the nervous system via feedback of movement inaccuracy or its consequences, such as time spent in making corrective movements (Meyer *et al.* 1988; Harris 1995). There is no need explicitly to calculate the cost of different movements because the optimal trajectory could be learned from the experience of repeated movements. In addition, the TOPS framework highlights the important effects of noise and uncertainty, and the methods that the CNS has developed to minimize their effects. This stresses the need to take into account the properties at the low, neural level to understand human behaviour at the high, cognitive level.

4. SENSORIMOTOR NOISE

Information about hand position is necessary to plan a goal-directed arm movement. However, due to motor noise, actual movements generally differ from the intended ones, yet we usually reach the target. This is because feedback is used during the movement to make the necessary corrections. This indicates that, during the

movement itself, it is also important to have an estimate of the hand's position. It has been described above how visual and proprioceptive information is combined to estimate the position of a static hand. During the course of a movement, however, another source of information can be used as well: a copy of the motor commands sent to the muscles. This efference copy (Sperry 1950; von Holst & Mittelstaedt 1950; Festinger & Canon 1965) can be used to predict the consequence of the motor commands. This prediction is made by an internal forward model (Wolpert *et al.* 1995; Miall & Wolpert 1996); a system that mimics the causal flow of a process by predicting its next state given the current state and the motor command. Due to noise and possible inaccuracies in the forward model, the predicted new state will, in general, have a finite precision and accuracy.

In estimating the state of the hand, the CNS faces another integration problem: how does it combine sensory inflow about hand position with a prediction of the position based on motor outflow, when the information from both sources is imperfect? This is a problem that has been solved in the engineering field of optimal state estimation. For linear systems, the solution is known as the Kalman filter (Kalman & Bucy 1961; figure 4). This is a linear dynamical system that integrates the motor outflow and the sensory inflow in a way that minimizes the uncertainty in the overall estimate. This optimal integration is achieved at any moment throughout the movement. In the feed-forward path, a forward model predicts the next state on the basis of the efference copy of the motor command, thereby simulating the dynamics of the arm. In the feedback path, the sensory feedback is compared with the sensory feedback predicted by a second forward model, one that mimics the behaviour of the sensory system. The difference between predicted and actual sensory feedback is used to correct the state estimate of the (first) forward model. Because the reliability of the estimate will vary during the movement, the optimal weighting varies over time. This is determined by the time-varying Kalman gain, which determines the optimal weighting to minimize the uncertainty in the overall estimate.

The Kalman filter has been used to model the estimation of hand position immediately after a movement based on motor outflow and proprioceptive feedback (Wolpert *et al.* 1995). Experimental data showed that both the constant and variable errors initially increased monotonically with movement duration. After about 1 s, both remained constant or decreased slightly. The Kalman filter reproduced this pattern. During the early part of the movement, when the current state estimate was accurate, the position predicted by the forward model was weighted heavily. Later, however, the estimate became less reliable, and the weighting shifted smoothly towards the feedback process. Models based purely on sensory feedback or on motor outflow could not reproduce these patterns, thereby providing evidence that information from both sources was used to estimate hand position, in order to reduce the combined effect of sensory and motor noise.

One problem in the estimation of state is that, due to feedback delays, the sensory signals about the state (position and velocity) of the arm will inevitably lag behind the actual state. It takes about 200 ms for visual feedback to influence an ongoing movement (Keele &

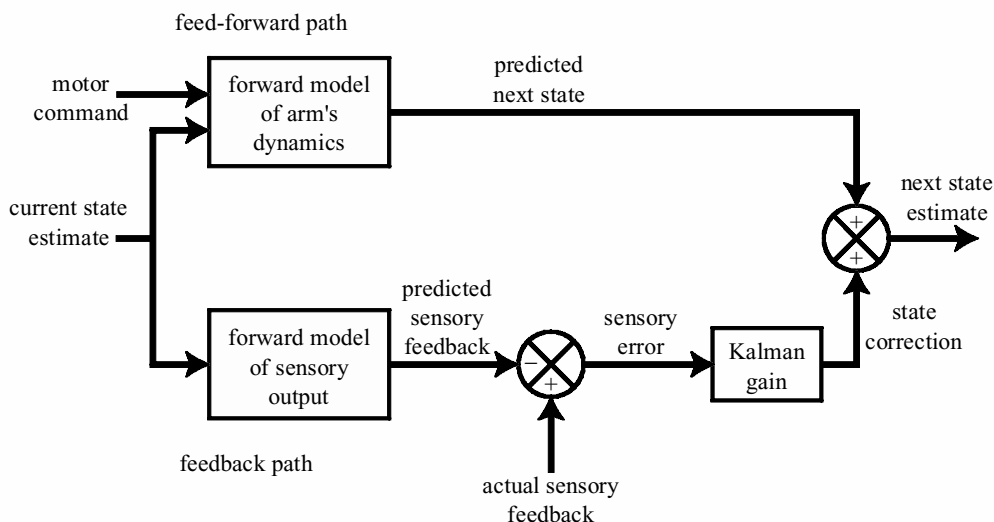


Figure 4. Schematic representation of the Kalman filter that estimates the next state given the motor command and the current state estimate. The model consists of two paths. In the (upper) feed-forward path, a forward model predicts the next state on the basis of the current state estimate and an efference copy of the motor command. In the (lower) feedback path, the actual sensory feedback is compared with the feedback predicted by a forward model of the sensory output (based on the current estimate). The difference between actual and predicted sensory feedback is the sensory error and is used to correct the forward model's state estimate. The Kalman gain determines how the outcome of both paths is weighted.

Posner 1968; Georgopoulos *et al.* 1981; van Sinderen *et al.* 1988), mainly because the photoreceptors in the retina are slow. The corresponding delay for proprioception is somewhat shorter, around 120 ms (Jeannerod 1988). These delays are long compared with the duration of a typical goal-directed movement, which is a serious problem for feedback control. One model developed to deal with time delays is the Smith predictor (Miall *et al.* 1993). This model also uses a forward model of the dynamics of the motor system, but in addition to that it also models the delays present in the system.

It is difficult to prove that the nervous system follows strategies identical to models such as the Kalman filter and the Smith predictor. However, the key ingredient in these architectures is the forward model, and there is strong evidence that the nervous system uses forward models to predict the consequences of motor actions. This has been demonstrated convincingly for the situation in which one could be uncertain about whether sensory information is due to changes induced by one's own action or by a change in the outside world. It is often impossible to distinguish between these two possibilities on the basis of the sensory information only, because similar sensory information can arise for external events or self-generated movements. However, it is possible to distinguish between them when the consequences of one's own actions as predicted by a forward model are taken into account.

This has been studied extensively for tickling sensation. It is well known that healthy humans cannot tickle themselves and that self-administered tactile stimuli feel less ticklish than externally administered tactile stimuli (Weiskrantz *et al.* 1971). This could be explained by the use of forward models (Blakemore *et al.* 1999; figure 5). When motor commands have been sent, the forward model predicts the resulting sensory feedback, which, in the case of tickling movements, would include tactile signals. The predicted feedback is compared with the actual sensory feedback and when these correspond to each other




Figure 5. Model used to determine whether sensory information arises from external influences or from a self-produced movement. A forward model predicts the sensory feedback based on an efference copy of the motor command. This prediction is then compared with the actual sensory feedback to produce the sensory discrepancy signal. If there is little or no sensory discrepancy, the sensory information is likely to be the result of a self-generated movement. If a large discrepancy is detected, the sensory information is more likely to be the result of external events.

they cancel each other, resulting in the absence of a tickle sensation. However, when one is tickled by someone else, the forward model does not predict any tactile feedback. The comparison with the actual feedback therefore results in a large sensory discrepancy, giving rise to a strong tickle sensation.

To test this hypothesis, Blakemore *et al.* (1999) asked subjects to move their left hand and to rate the tickle sensation of stimuli that a robot applied to their right hand. The movement of the robot could be coupled to the movement of the subject's left hand. When the robot's movement corresponded exactly to the movement of the left hand, as if the left hand held an object that tickled the

right hand, tickle sensation was lower than when the robot tickled the subject when the subject made no movement. Next, movements were tested in which the motion of the left hand determined the robot movement, but now the relationship between action and its consequences were varied parametrically. The ticklishness rating increased systematically with the discrepancy between the action and its predicted consequence, either due to introducing a time delay or a spatial rotation between the motion of the left hand and the displacement of the robot that tickled the right hand. This directly supports the forward model hypothesis because it shows that the tickle sensation increases with the discrepancy between the sensory feedback predicted by a forward model and the actual sensory feedback.

A functional magnetic resonance imaging study (Blakemore *et al.* 1998) showed that more activity was found in bilateral secondary somatosensory cortex and the anterior lobe of the right cerebellum for externally produced compared with self-produced stimuli. The increase in somatosensory cortex activity probably reflects the increased ticklishness sensation. The observed effect for the cerebellum could reflect the discrepancy between predicted and actual sensory feedback, indicating that the forward model that predicts the sensory consequences of motor commands resides in the cerebellum. This is also suggested by the finding that the cerebellar activity correlates with the delay between hand movement and tactile stimulation (Blakemore *et al.* 2001).

This work has important consequences at the cognitive level. It could be argued that a defect in the central self-monitoring, as described above, might underlie delusions of control experienced by people with schizophrenia (Frith 1992). These patients move their limbs but claim that they are being moved by external agents. To examine whether this symptom may be due to lack of prediction of the consequences of action, Blakemore *et al.* (2000) used a similar paradigm to the previous studies. The results showed that patients with such symptoms did not rate self-produced stimuli as less ticklish than externally produced stimuli, whereas control subjects did rate the self-produced stimuli as less ticklish. This indicates that an impairment in the 'self-monitoring' mechanism, as implemented as a forward model, could cause thoughts or actions to become isolated from the sense of will normally associated with them, leading to symptoms associated with schizophrenia. This highlights the importance of the mechanisms that the nervous system has developed to reduce the uncertainties about whether sensory information is due to self-induced changes or to changes in the outside world.

5. NEURAL NOISE AND NETWORK MODELS

In this section, we review noise at the neural and network level. In general, the neural mechanisms by which the CNS copes with uncertainty are still poorly understood. However, recent experimental and computational studies have begun to shed light on both the sources of neural noise and its control. Combined electrophysiological recordings and theoretical analyses have delineated important sources of neuronal noise, while neural network

modelling has indicated the existence of links between neural tuning and connectivity and noise control.

Noise is present at every level of the sensorimotor chain. In sensors, errors can arise through an incorrect estimation of sensor characteristics (e.g. temperature, metabolic state, sensor position or activation). Noise is also present in effectors: an incorrect estimate of limb position or muscle fatigue will translate into a motor error. In both proprioceptors and muscles, unavoidable instantaneous noise is added by the transduction between a continuous mechanical signal and a discrete sequence of spikes, through nonlinear dynamical systems (Read & Siegel 1996). The operation of neural networks also generates variability in the signals that are propagated. Synaptic variability (Allen & Stevens 1994; Tsodyks & Markram 1997), local firing synchronization (Stevens & Zador 1998) and chaotic network dynamics (van Vreeswijk & Sompolinsky 1996) have been shown to increase the variance of spike trains, although the isolated neuron is itself extremely reliable (Mainen & Sejnowski 1995). In general, this neural noise can be described by a Poisson distribution (with the notable exception of motor neurons where noise is close to Gaussian (Gomez *et al.* 1986)).

The increase in noise due to the complex architecture of neural networks is compensated by the redundancy of information representation. Indeed, sensory or motor signals are often coded by assemblies of neurons, each of which is tuned to (i.e. discharges maximally for) a given set of signal parameters. This distribution makes the system immune to synaptic failure or to the death of individual neurons. When tuning distribution is uniform (e.g. in primary visual cortex), population activity allows robust information representation if noise is not too correlated between neurons (Abbott & Dayan 1999). Moreover, it has been suggested that a population can encode more than one stimulus or response, thus enabling an internal representation of uncertainty (Zemel *et al.* 1998).

A common problem for neurophysiologists is to decode the information embedded in population activity. Information theory shows how the precision of information retrieval is bounded by the noise. The variance of the best unbiased estimator will necessarily exceed a minimal value called the 'Cramér-Rao bound' (Cox & Hinckley 1974). Pouget and collaborators have described a neural network of broadly tuned neurons that can at the same time recover the information encoded in another population with near-optimal precision (Pouget *et al.* 1998; Deneve *et al.* 1999) and re-encode it in a noise-reduced pattern. Thus, this type of neuronal architecture can clean up the noise before further processing.

The optimal estimation of the information embedded in the collective firing of neurons is also the solution to multisensory integration when sensory inputs from different modalities must be integrated in a common representation. We have seen in the preceding section that humans seem optimally to estimate the position of their hand when they can use both vision and proprioception. That is, they compute the most likely place according to the probability distributions associated with proprioceptive and visual inputs in isolation. This is precisely how an optimal estimator of the corresponding collective neural activity would behave. Thus, principles close to those exposed by Pouget *et al.* (1998) are possibly in operation. However,

neurons involved in the localization task are not tuned to a specific value of gaze or arm position, but respond monotonically with gaze or hand position (Hepp & Henn 1985; Helms Tillery *et al.* 1996). In that case, it has been shown that a very simple network can compute an estimate whose variance is close to the Cramér-Rao bound (Guigon & Baraduc 2002). Moreover, this type of network has the capability of learning a new sensory-sensory congruence, e.g. during prism adaptation.

This last point reminds us that neural networks do not only process sensory information to get the best estimate of a raw input, but in general process it to produce an adapted motor response, be it scratching your nose or making a saccade to the right of these words. Sensorimotor transformation schemes that make use of redundancy and broad tuning (Salinas & Abbott 1995; Baraduc *et al.* 2001) are a good solution to minimize the detrimental effects of noise and seem ubiquitous in the CNS (Georgopoulos *et al.* 1982; Kalaska *et al.* 1983; Crutcher & DeLong 1984; Fortier *et al.* 1989). Recently, Todorov (2002) has shown that cosine tuning is an optimal way to minimize motor errors in the presence of signal-dependent noise.

Although these studies have provided clues, how neuronal noise influences the computational schemes in the CNS is a question that has just begun to be investigated.

6. CONCLUSIONS

We have reviewed a number of activities that the CNS has to perform in order to execute a goal-directed arm movement. We described spatial localization, integration of information from multiple sensory modalities, movement planning and integration of sensory and motor information during movement. The way in which the CNS carries out each of these activities can be understood from the idea that there is noise and uncertainty in the sensory and motor systems, and that the CNS tries to minimize the detrimental effects of this noise. The noise is present at the low, neural level, whereas the way in which tasks are carried out is usually visible at the high, behavioural level, for instance as hand trajectories. This indicates that it is often important to consider effects at the neural level in order to understand performance at the behavioural level. This is true for execution of goal-directed arm movements, but in a more general way, this idea could hold for a much wider class of cognitive tasks.

R.v.B. is supported by the Wellcome Trust. P.B. is supported by an EU Marie Curie training fellowship. This work was supported by the Wellcome Trust, the Human Frontiers Science Programme, McDonnell Foundation and the BBSRC.

REFERENCES

- Abbott, L. F. & Dayan, P. 1999 The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101.
- Allen, C. & Stevens, C. F. 1994 An evaluation of causes for unreliability of synaptic transmission. *Proc. Natl Acad. Sci. USA* **91**, 10 380–10 383.
- Atkeson, C. G. & Hollerbach, J. M. 1985 Kinematic features of unrestrained vertical arm movements. *J. Neurosci.* **5**, 2318–2330.
- Bahill, A. T., Clark, M. R. & Stark, L. 1975 The main sequence, a tool for studying human eye movements. *Math. Biosci.* **24**, 191–204.
- Baraduc, P., Guigon, E. & Burnod, Y. 2001 Recoding arm position to learn visuomotor transformations. *Cerebr. Cortex* **11**, 906–917.
- Blakemore, S. J., Wolpert, D. M. & Frith, C. D. 1998 Central cancellation of self-produced tickle sensation. *Nat. Neurosci.* **1**, 635–640.
- Blakemore, S. J., Frith, C. D. & Wolpert, D. M. 1999 Spatio-temporal prediction modulates the perception of self-produced stimuli. *J. Cogn. Neurosci.* **11**, 551–559.
- Blakemore, S. J., Smith, J., Steel, R., Johnstone, E. C. & Frith, C. D. 2000 The perception of self-produced stimuli in patients with auditory hallucinations and passivity experiences: evidence for a breakdown in self-monitoring. *Psychol. Med.* **30**, 1131–1139.
- Blakemore, S. J., Frith, C. D. & Wolpert, D. M. 2001 The cerebellum is involved in predicting the sensory consequences of action. *Neuroreport* **12**, 1879–1884.
- Collewijn, H., Erkelens, C. J. & Steinman, R. M. 1988 Binocular coordination of human horizontal saccadic eye-movements. *J. Physiol.* **404**, 157–182.
- Cox, D. R. & Hinkley, D. V. 1974 *Theoretical statistics*. London: Chapman & Hall.
- Crutcher, M. D. & DeLong, M. R. 1984 Single cell studies of the primate putamen. II. Relations to direction of movement and pattern of muscular activity. *Exp. Brain Res.* **53**, 244–258.
- Deneve, S., Latham, P. E. & Pouget, A. 1999 Reading population codes: a neural implementation of ideal observers. *Nat. Neurosci.* **2**, 740–745.
- Donders, F. C. 1848 Beitrag zur Lehre von den Bewegungen des menschlichen Auges. *Anat. Physiol. Wiss.* **1**, 105–145.
- Enderle, J. D. & Wolfe, J. W. 1987 Time-optimal control of saccadic eye-movements. *IEEE Trans. Biomed. Engng* **34**, 43–55.
- Fedderson, W. E., Sandel, T. T., Teas, D. C. & Jeffress, L. A. 1957 Localization of high-frequency tones. *J. Acoust. Soc. Am.* **29**, 988–991.
- Festinger, M. L. & Canon, L. K. 1965 Information about spatial location based on knowledge about efference. *Psychol. Rev.* **72**, 373–384.
- Fitts, P. M. 1954 The information capacity of the human motor system in controlling the amplitude of movements. *J. Exp. Psychol.* **47**, 381–391.
- Flash, T. & Hogan, N. 1985 The coordination of arm movements: an experimentally confirmed model. *J. Neurosci.* **5**, 1688–1703.
- Foley, J. M. & Held, R. 1972 Visually directed pointing as a function of target distance, direction, and available cues. *Percept. Psychophys.* **12**, 263–268.
- Fortier, P. A., Kalaska, J. F. & Smith, A. M. 1989 Cerebellar neuronal activity related to whole arm-reaching movements in the monkey. *J. Neurophysiol.* **62**, 198–211.
- Frith, C. D. 1992 *The cognitive neuropsychology of schizophrenia*. London: Lawrence Erlbaum Associates.
- Georgopoulos, A. P., Kalaska, J. F. & Massey, J. T. 1981 Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty and change in target location. *J. Neurophysiol.* **46**, 725–743.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R. & Massey, J. T. 1982 On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* **2**, 1527–1537.
- Ghahramani, Z., Wolpert, D. M. & Jordan, M. I. 1997 Computational models of sensorimotor integration. In *Self-organization, computational maps, and motor control* (ed. P.

- Morasso & V. Sanguineti), pp. 117–147. Amsterdam, The Netherlands: Elsevier.
- Gomez, C., Canals, J., Torres, B. & Delgado-Garcia, J. M. 1986 Analysis of the fluctuations in the interspike intervals of abducens nucleus neurons during ocular fixation in the alert cat. *Brain Res.* **381**, 401–404.
- Guigon, E. & Baraduc, P. 2002 A neural model of perceptual-motor alignment. *J. Cogn. Neurosci.* **14**, 538–549.
- Harris, C. M. 1995 Does saccade undershoot minimize saccadic flight-time? A Monte-Carlo study. *Vis. Res.* **35**, 691–701.
- Harris, C. M. 1998 On the optimal control of behaviour: a stochastic perspective. *J. Neurosci. Meth.* **83**, 73–88.
- Harris, C. M. & Wolpert, D. M. 1998 Signal-dependent noise determines motor planning. *Nature* **394**, 780–784.
- Helms Tillery, S. I., Soechting, J. F. & Ebner, T. J. 1996 Somatosensory cortical activity in relation to arm posture: non-uniform spatial tuning. *J. Neurophysiol.* **76**, 2423–2438.
- Hepp, K. & Henn, V. 1985 Iso-frequency curves of oculomotor neurons in the rhesus monkey. *Vis. Res.* **25**, 493–499.
- Hogan, N. 1984 An organizing principle for a class of voluntary movements. *J. Neurosci.* **4**, 2745–2754.
- Jacobs, R. A. 1999 Optimal integration of texture and motion cues to depth. *Vis. Res.* **39**, 3621–3629.
- Jeannerod, M. 1988 *The neural and behavioural organization of goal-directed movements*. Oxford: Clarendon.
- Jones, K. E., Hamilton, A. & Wolpert, D. M. 2002 The sources of signal dependent noise during isometric force production. *J. Neurophysiol.* (In the press.)
- Kalaska, J. F., Caminiti, R. & Georgopoulos, A. P. 1983 Cortical mechanisms related to the direction of two dimensional arm movements: relations in parietal area 5 and comparison with motor cortex. *Exp. Brain Res.* **51**, 247–260.
- Kalman, R. E. & Bucy, R. S. 1961 New results in linear filtering and prediction. *J. Basic Engng (ASME)* **83D**, 95–108.
- Keele, S. W. & Posner, M. I. 1968 Processing of visual feedback in rapid movements. *J. Exp. Psychol.* **77**, 155–158.
- Lacquaniti, F., Terzuolo, C. A. & Viviani, P. 1983 The law relating kinematic and figural aspects of drawing movements. *Acta Psychologica* **54**, 115–130.
- Mainen, Z. F. & Sejnowski, T. J. 1995 Reliability of spike timing in neocortical neurons. *Science* **268**, 1503–1506.
- Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E. & Smith, J. E. K. 1988 Optimality in human motor performance: ideal control of rapid aimed movements. *Psychol. Rev.* **98**, 340–370.
- Miall, R. C. & Wolpert, D. M. 1996 Forward models for physiological motor control. *Neural Netw.* **9**, 1265–1279.
- Miall, R. C., Weir, D. J., Wolpert, D. M. & Stein, J. F. 1993 Is the cerebellum a Smith predictor? *J. Mot. Behav.* **25**, 203–216.
- Middlebrooks, J. C. & Green, D. M. 1991 Sound localization by human listeners. *A. Rev. Psychol.* **42**, 135–159.
- Morasso, P. 1981 Spatial control of arm movements. *Exp. Brain Res.* **42**, 223–227.
- Pouget, A., Zhang, K. C., Deneve, S. & Latham, P. E. 1998 Statistically efficient estimation using population coding. *Neural Comput.* **10**, 373–401.
- Read, H. L. & Siegel, R. M. 1996 The origins of aperiodicities in sensory neuron entrainment. *Neuroscience* **75**, 301–314.
- Salinas, E. & Abbott, L. F. 1995 Transfer of coded information from sensory to motor networks. *J. Neurosci.* **15**, 6461–6474.
- Schmidt, R. A., Zelaznik, H., Hawkin, B., Frank, J. S. & Quinn, J. T. 1979 Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol. Rev.* **86**, 415–451.
- Sperry, R. W. 1950 Neural basis of the spontaneous optokinetic response produced by visual inversion. *J. Comp. Physiol. Psychol.* **43**, 482–489.
- Stevens, C. F. & Zador, A. M. 1998 Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.* **1**, 210–217.
- Todorov, E. 2002 Cosine tuning minimizes motor errors. *Neural Comput.* **14**, 1233–1260.
- Tsodyks, M. V. & Markram, H. 1997 The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl Acad. Sci. USA* **94**, 719–723.
- Uno, Y., Kawato, M. & Suzuki, R. 1989 Formation and control of optimal trajectories in human multijoint arm movements: minimum torque-change model. *Biol. Cybern.* **61**, 89–101.
- van Beers, R. J., Sittig, A. C. & Denier van der Gon, J. J. 1996 How humans combine simultaneous proprioceptive and visual position information. *Exp. Brain Res.* **111**, 253–261.
- van Beers, R. J., Sittig, A. C. & Denier van der Gon, J. J. 1998 The precision of proprioceptive position sense. *Exp. Brain Res.* **122**, 367–377.
- van Beers, R. J., Sittig, A. C. & Denier van der Gon, J. J. 1999 Integration of proprioceptive and visual position-information: an experimentally supported model. *J. Neurophysiol.* **81**, 1355–1364.
- van Beers, R. J., Wolpert, D. M. & Haggard, P. 2002 When feeling is more important than seeing in sensorimotor adaptation. *Curr. Biol.* **12**, 834–837.
- van Sinderen, J. F., Denier van der Gon, J. J. & Gielen, C. C. A. M. 1988 Conditions determining early modification of motor programmes in response to changes in target location. *Exp. Brain Res.* **71**, 320–328.
- Van Vreeswijk, C. & Sompolinsky, H. 1996 Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726.
- von Helmholtz, H. 1867 *Handbuch der physiologischen optik, band 3*. Leipzig, Germany: Voss.
- von Holst, E. & Mittelstaedt, H. 1950 Das Reafferenzprinzip Wechselwirkungen zwischen Zentralnervensystem und Peripherie. *Naturwissenschaften* **37**, 464–476.
- Warren, D. H. & Schmitt, T. L. 1980 Intermodal organization: a methodological localization study. *Percept. Motor Skills* **50**, 1111–1118.
- Weiskrantz, L., Elliot, J. & Darlington, C. 1971 Preliminary observations of tickling oneself. *Nature* **230**, 598–599.
- Welch, R. B. 1978 *Perceptual modification*. New York: Academic.
- Welch, R. B. & Warren, D. H. 1986 Intersensory interactions. In *Handbook of perception and human performance vol. 1: sensory processes and perception* (ed. K. R. Boff, L. Kaufman & J. P. Thomas), pp. 25–1–25–36. New York: Wiley.
- Wolpert, D. M. & Harris, C. M. 2001 Optimal saccadic control with motor noise predicts head fixed and head free saccadic trajectories and main sequence. *Soc. Neurosci. Abstr.* **27**, 71.36.
- Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. 1995 An internal model for sensorimotor integration. *Science* **269**, 1880–1882.
- Zemel, R. S., Dayan, P. & Pouget, A. 1998 Probabilistic interpretation of population codes. *Neural Comput.* **10**, 403–430.

GLOSSARY

CNS: central nervous system

TOPS: task optimization in the presence of signal-dependent noise